FUSION SIMULATION PROGRAM PLANNING PROJECT WORKSHOP SUMMARY REPORT

A Report from the Executive Planning Committee for the Future *Fusion Simulation Program*

on the Workshop Held in Boulder, CO on March 15 through 18, 2010

June 28, 2010

Table of Contents

Table of Figures

1 Introduction

The focus of the current two-year planning project for the Fusion Simulation Program (FSP) is to produce a plan for enabling scientific discovery of important new plasma phenomena with associated understanding that emerges only upon integration. The goal of the FSP is to deliver a suite of predictive integrated fusion energy science (FES) simulation capabilities that are properly validated against experiments in regimes relevant for producing practical fusion energy. This requires developing program and management plans for addressing the scientific and computational gaps and producing a living scientific roadmap that identifies compelling deliverables needed for progress on the science drivers deemed most important for magnetic fusion energy (MFE) research.

Overall, an effective FSP planning process requires cognizance of the strategic importance of delivering some improved nearer-term software capabilities to the user community while also engaging in the longer-term development needed to accelerate progress on the more formidable scientific challenges facing MFE. In order to do so, it is critically important to engage experts from the FES simulation, theory, and experimental communities and from the applied math and computer science communities.

The FSP Planning Workshop, which was held in Boulder, CO on March 15th through 18th, 2010, was accordingly motivated to:

(1) clarify and keep current the "science development roadmap" with associated computational infrastructure needs identified;
(2) have collaborative discussions involving science drivers, physics components, computational frameworks, V&V, etc., to assess software needs/gaps, to set priorities, and to help identify promising future approaches; and
(3) plan for how best to test and subsequently provide appropriate user support of future FSP tools/products.

1.1 Workshop Goals

The workshop was structured to address a set of core technical issues confronting the FSP definition effort by engaging a broad section of our community. Specific goals provided to attendees were to:

- Identify approaches, designs, requirements, scope for the Fusion Simulation Program.
- Identify the infrastructure (physics components and integration needs) needed to address the science drivers.
- Identify models for FSP integration.
- Identify requirements and models for integrated data management including data organization, metadata, data access, and namespace management.
- Identify the needs and possible environments for production computing, including user support, job submission and monitoring, hardware, installed software.
- Determine the sets of software packages that can be shared throughout the FSP.
- Determine any development needed uniquely for any specific science driver.
- Gather information (e.g., where to obtain software, use cases) as input to prototyping process to evaluate proposed technologies.

1.2 Workshop Outcome

The FSP management team was pleased with the productivity and engagement of community members in this workshop. The outcome is this report which describes the requirements for FSP, summarizes the current state of software applicable to the FSP and lays out future directions. The report identifies a set of computational infrastructure needed by all science drivers as well as those required by particular science drivers. It identifies extant multiple approaches and gaps – requirements for which there is currently no solution. Future work, including additional targeted workshops and design reviews will be aimed toward developing a detailed program plan, with schedules, resource requirements and milestones.

The workshop agenda and various presentations can be reviewed at: https://ice.txcorp.com/trac/fspfrmwrkplan/wiki/FspPlanningAgenda

2 Science Drivers

2.1 Introduction

The Science Drivers are a set of compelling scientific problems chosen to focus the FSP's design and initial implementation. They could also be described as a set of evolving use cases. Further, the drivers will define and exercise the required range of technical capabilities and lead to useful simulation tools for the broader fusion community.

At the time of the workshop, each group had defined a "science development roadmap", broad plans for development in each of the science areas. These are defined in terms of the scientific capabilities required and do not refer to particular implementations or codes. Achieving the capabilities outlined in the roadmaps will require coordinated efforts by the FSP in partnership with theory, experiments, and the Scientific Discovery through Advanced Computing Program (SciDAC) Program. Work is underway to outline requirements for advanced physics components, computational frameworks and experimental validation for each of the drivers – the workshop was an important step in the process.

There are six Science Drivers defined, each will be discussed in the following sections of this report.

1. Boundary Layer: includes turbulence, atomic physics and plasma-wall interactions
2. Pedestal: transport barrier, profile structure, relaxation mechanisms (ELMs)
3. Core Profiles: nonlinear turbulence and MHD
4. Wave particle interactions: includes fusion products and RF
5. Disruption avoidance, detection and mitigation
6. Whole device modeling

While articulated as separate tasks now, over time the science driver paths will begin to merge. For example, the pedestal will need to be integrated with Boundary Physics and the Core Profiles efforts and the wave-particle models will need to include the effects of micro-turbulence. From the start, we envision full integration (whole device modeling) at various, but increasing levels of physics fidelity. To start, reduced models will be required for many phenomena, while over time better models will need to be made available. The FSP will need to support development of a range of models, balancing fidelity and computational speed and frameworks should support flexible mix of models employing different levels of accuracy, computation on widest range of platforms. An important part of each science driver's development will be the production of reduced models suitable for whole device modeling.

2.2 Integrated Boundary Layer, Divertor, and Plasma-Wall Interactions

There appeared to be general agreement at the meeting on the overall roadmap[1] presented in the Monday plenary session, though there were some additions and modifications advocated as detailed in a later section. Also, the sequencing of roadmap items was discussed, especially fluid and kinetic models and plasma-material interactions. Because the roadmap presently has 8 steps, it is difficult to communicate concisely to the larger

[1] Listed on the FSP Science Driver Wiki pages: http://fspscidri.web.lehigh.edu/index.php/Main_Page#Integrated_Boundary_Layer_.28SOL.29.2C_Divertor.2C_Plasma_Wall_Interactions

fusion community the importance of the Boundary issues. Consequently, some of the group's discussions focused on highlighting the big science questions that need to be answered and then worked on the staged approach, i.e., the roadmap, that should be taken to answer them.

2.2.1 Scientific goals

Major science questions were summarized as

1. Heat loads to material surfaces. Includes scaling of peak heat fluxes and profiles in L-mode, H-mode-between-ELMs, ELMs, and disruptions.
2. Hydrogenic/helium particle transport and flows. Includes particle fluxes to walls, pumping, fueling via recycling and gas injection (& pellets?). Fueling and flows provided particle and momentum sources to pedestal.
3. Impurity generation and transport. Includes physical & chemical sputtering, sheaths/RF, blobby impurity transport, mixing materials, & intrusion to core.
4. Tritium recycling, transport, and retention in materials. Related to other questions, but list separately because of its high importance.
5. Wall/divertor material modification. Includes surface evolution from erosion, redeposition with mixed materials, and dust generation and transport.

The roadmap discussions focused on two periods, 2-5 years and 5-10 years. For the 2-5 year period, the emphasis is on 2D and 3D fluid models of scrape-off layer (SOL) transport and turbulence to understand and predict heat fluxes and edge flows; some portion of the plasma inside the magnetic separatrix would be included. Validation would be done with devices operating in the collisional fluid regime. Toward the end of the period, impurities would be included with initial PMI models. Thus, this first period encompasses the fluid limit of the first 3 original roadmap items (see end of this document).

In the 5-10 year time frame, the emphasis would shift to kinetic models (4D and 5D) of the SOL/separatrix regions to understand how the same issues of heat and particle transport of hydrogen and impurities changes in hotter, less collisional edge plasmas expected in ITER and other high-power devices. This activity would begin in the initial time phase where the focus is on exploiting fluid models, but given the inherent complexity of kinetic models, especially in the boundary region with large gradients and turbulence fluctuations, it will have a longer development time. The kinetic models can assess the coupling between the SOL and pedestal owing to large ion drift-orbits and possible prompt ion loss. This second period combines the first 4 roadmap items, and the 5th roadmap item (3D magnetic equilibrium, walls, RF sheaths).

The need for a substantial effort in first-principles plasma-wall models was discussed. While advanced development of this activity might be carried out without initial close coupling to the plasma/neutral models for the SOL, basic models of the processes are needed for the two phases described above. Consequently, the PMI model development should be carried out under the Boundary Science Drive. As this is a very complex topic, if additional funding becomes available to accelerate progress (say a SciDAC project), it was emphasized that issues associated with plasma-wall interactions are typically localized to the near surface, say in the range of a few mm's, and these issues differ considerably from bulk material damage from neutrons. Thus, funding for "materials research" needs to clearly distinguish these two areas.

2.2.2 Additional details on a plasma/neutral approach

As mentioned above, the initial 2-5 year target begins with work on hydrogen-only fluid (2D, 3D) and kinetic (4D, 5D) models for the boundary region with the expectation that the fluid models, being simpler (lower dimensionality), would obtain physics results first that could be validate against devices in collisional operating regimes. Furthermore, the kinetic models should be able to reproduce the fluid results with a proper collision model, providing a verification test. The simultaneous fluid/kinetic approach was not explicitly part the Wiki roadmap (kinetic plasma model was item #4, while fluid model was #1). Both fluid and kinetic codes exist, and while fluid codes are arguably considerably more mature, putting the kinetic development on hold for any period of time seems unwarranted. Here basic neutral models with a static particle-recycling model would be used. Turbulent transport would need to be extended or coupled to long transport timescales, likely at the end of this stage. The stage would begin to address science questions 1 and 2 above.

The second stage would add more detailed plasma-wall interaction models, such as a time-dependent recycling model that accounts for wall uptake (pumping) and out gassing depending on conditions such as wall temperature. Improved neutral models would be incorporated. Also, the ability to include slowing evolving 2D MHD equilibria (later 3D in the 5-10 year frame) should be included to initiate studies of discharge ramp-up and ramp-down phases (not explicitly included in the Wiki roadmap). This stage would increase the fidelity of answers to science questions 1 and 2.

The third stage would include initial models for physical and chemical sputtering of impurities. The impurity species would be added to the fluid and kinetic codes as they have completed the previous stages. A key science question to answer here is #3, especially the transport of impurities in the intermittent, ”blob” turbulence.

2.2.3 Overview of PMI modeling approach (more detail in Appendix B)

Plasma-material interactions pose an immense scientific challenge and are one of the most critical issues in magnetic confinement fusion research. The demands on plasma-facing materials in a steady-state fusion device include extreme particle and thermal fluxes. These energetic fluxes have pronounced impacts on the topology and chemistry of the near-surface region of the material, which influence the plasma sheath potentials and subsequent threat spectra. The material evolution is also inherently multiscale in time and are likely controlled by diffusional phenomena that are influenced by the high heat loads and subsequent thermal (and stress) gradients into the material, as well as by defect micro/nanostructures induced by both the ion and neutron particle irradiation. Tables for physical and chemical sputtering on well-characterized materials are known and will be used initially. Also, models are developing that include the time-dependence of particle recycling at the surface that can be implemented. Beyond these basic models, an number of more complex processes that involve plasma modification of the surface of the material must be addressed as summarized in Appendix B.

Component functionality required for 2-5 year goal

1. 2D MHD equilibrium
2. Wall/divertor geometry
3. Mesh
4. 3D plasma turbulence – (two-fluid); rho_i-scale turbulence
5. Field model; EM (A||)
6. 2D plasma profile evolution (long 3D simulations or coupled to 2D)
7. 2D neutral transport
8. Atomic rates (table vs. n, T)
9. Basic dynamic wall model

Component functionality added for 5-10 year goal (begun earlier)

1. 5D plasma turbulence – (2 velocity dimensions); rho_i-scale turbulence
2. 4D plasma profile evolution (long 3D evolution or coupled to 2D)
3. 3D MHD equilibrium
4. 3D neutral model

2.2.4 Overall gaps, issues, and missing items in original roadmap

1. Transport simulations with evolving 2D MHD equilibria
2. Method of long-time SOL transport simulation for strong, intermittent transport – i.e., 3D filamentary “blob” transport.

3. Accurate and efficient algorithm for nonlinear Fokker-Planck collision operator allowing transition from long to short mean-free paths.
4. How much of the region inside the separatrix must be included, e.g., what is the birth-region of blobs, what is the impact of ion orbit loss on heat fluxes and Er, what is the coupling between SOL flows and core rotation, and how do neutrals extending across the separatrix fuel the pedestal?
5. Fundamental models of plasma-wall interactions that describe the evolution of the material under large particle and heat fluxes. Also, the processes by which hydrogen (tritium) is retained in materials and how can it be released.
6. Efficient coupling between a kinetic neutral model (typically Monte Carlo) and the plasma could be substantially improved for transport timescales. Also, hydrogenic radiation transport in high-density, optically thick regimes needs further development.

2.2.5 Connections to other Science Drivers

- There is a strong connection to the Pedestal Science Driver. While a number of physics issues can probably be studied separately, others will likely require a close coupling (see #4 in gaps for details). Also, ELMs from the pedestal are a major concern of unacceptably large heat fluxes to materials
- SOL plasma interaction with RF antennas is also a major issue for the RF and Energetic Particles Science Driver. RF can heat and modify the SOL plasma and large RF sheaths can cause strong sputtering. There should be a cross-cutting activity here. Energetic particle loss impacts divertor/wall heat loads.
- The Disruption Science Driver again involves strong heat fluxes to material surfaces. These very intense interactions may be studied by some special tools, but there is a close similarity to modeling ELM heat fluxes.
- The Whole Device Model Science Driver needs a model of the SOL and heat fluxes; these fluxes are a central concern for operating new high-power devices.

2.3 Pedestal

2.3.1 Scientific goals

The practical goal for pedestal research is to achieve operation with a steady, 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 predict the onset of edge barriers, to predict the structure of the barrier (particularly the pressure at the top of the pedestal) , to predict the nature of the pedestal relaxation and to identify and optimize methods for reducing transient heat deposition on material surfaces. Since the pedestal height sets a critical boundary condition for overall plasma performance, accurate pedestal modeling is essential for an overall predictive capability for fusion plasmas.

2.3.2 Challenges and scientific gaps

Modeling the pedestal faces a number of significant challenges. There is substantial overlap in the temporal and spatial scales for the relevant physics and though it covers only a small region of the plasma, it spans a wide range in dimensionless parameters. For example, the bottom of the pedestal will usually be collisional while the top remains collisionless. Perturbations are often large in amplitude and can span the pedestal region – far from the valid regimes for most expansions which posit small perturbations on a nearly constant background. Models must contend with self-stabilization mechanisms, bifurcations and, in all likelihood, operation very close to marginal stability. Finally, the pedestal is influenced by atomic physics, the interactions with neutrals, impurities and atomic radiation.

There are fundamental experimental observations that cannot yet be modeled representing important areas where substantial scientific progress is needed to achieve goals described above. These include:

- L-H transition, particularly in terms of input power
- Prediction of large scale Er and plasma rotation
- The wide variety of ELM types and non-ELM H-modes
- Heat and particle loads from large Type I ELMs

2.3.3 Possible roadmap for development of pedestal modeling

There are a number of computational approaches to pedestal modeling which can be applied with increasing physics fidelity but also with increasing challenge to theory and computation. The simplest are models like EPED1 and are based on linear calculations of instability threshold and strong separation of scales. The next level of advancement would be 3D fluid or gyrofluid models. These only approximate some of the important kinetic effects, but can run over a wider range of scales than the kinetic codes and capture some of the finite-n physics. 4/5D Drift-kinetic/gyro-kinetic simulations come next in the hierarchy. Full f codes can overcome the small-perturbation limit inherent in df approaches but will require new theoretical formulation which will also be needed to move away from high-n approximations. 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. This general outline leads to a corresponding development roadmap with four major steps:

1. Linear models for pedestal structure

This step would begin with existing models that solve for static (time averaged) pedestal structure via linear stability analysis for 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 (ETG) and neoclassical effects. This analysis typically requires hundreds or thousands of independent stability calculations with trial equilibria. The key issues are robustness, error checking and automation. Extensive comparison with experimental data sets will need to be carried out. Once validated, these sorts of simulations can be readily included into whole device models.

2a. Dynamic evolution of pedestal profiles with quasi-linear models

Dynamic models would be constructed using quasi-linear calculation of transport fluxes along with heat and particle sources. An approach which has been successful for core modeling is to combine accurate linear gyrokinetic stability calculations with fluxes based on parametric scans of nonlinear turbulence codes (e.g. TGLF, MM, etc.). This approach is currently being generalized and extended into the pedestal. Neoclassical transport, including eventually 3D geometric effects, would also be modeled. Particle sources would require coupling to models for neutral transport and pellet fueling – and later to a more complete model of the boundary plasma, recycling, impurity sources, etc. These models could be linked eventually to 3D fluid turbulence simulations for the boundary plasmas and for the edge of L-mode plasmas and gyrokinetic models for the core to provide more complete profile simulations suitable for use in whole device models. All of these models would need to be 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 which limit or control the pedestal pressure gradients. These would include spontaneous plasma behavior (ELMs of various types, EHO, QCM, etc.) and active control through pellets, RMP, EMP, 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 (NIMROD, M3D) or kinetic-fluid codes like BOUT++. 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.

3. Direct Multi-Scale Simulation

The prior computational stages 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 scale separation, to determine when and how it breaks down and to assess the consequences. Theoretical progress will be required to formulate models that properly treat kinetics at finite-n scales. 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.

2.3.4 Summary of required elements of an improved physical model

The following physics components would be required to carry out the program outlined above.

- Realistic geometry (near edge, separatrix) including realistic boundary conditions, & regional coupling (to core, to SOL, divertor plate, wall)
- Reduced pedestal structure models
- Linear MHD
- Linear electromagnetic gyrokinetics (EM GK)
- Neoclassical (including dB)
- Fluid turbulence (separatrix, L-mode, dB)
- Nonlinear EM GK (near edge, electron and ion scales, dB)
- Nonlinear EM GK (cross separatrix, electron and ion scales, dB)
- Transport model for updating profiles
- Sources and sinks
    - Neutrals
    - Beam
    - RF
    - Pellets (*kinetic treatment with large perturbations)
    - Radiation
- Nonlinear extended MHD (across separatrix, dB)
- *Full Fokker-Planck nonlinear collision operator
- *Finite-n kinetic code (6D or extended 5D or kinetic-fluid)

**requires substantial new theoretical development*

2.3.5 Key Physics Gaps

- Edge GK with open/closed fields, full EM perturbations, realistic Boundary Conditions, multi-species, e-i scales
- Formalism for finite-n kinetic models
    - 6D GK, extended GK, higher order with existing GK, kinetic-fluid
- Code to implement finite-n kinetic
- Full coupling to core and boundary plasma

2.4 Prediction of Core Profiles

2.4.1 Scientific goals

The scientific goal is a validated transport model, which reliably predicts, for each plasma species, profiles of density, temperature as well as plasma rotation and current and their evolution on the transport time-scales. That is, it would encompass all the phenomena that sets 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 need 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.

A broad user community would be interested in the models developed under this science driver. These would include "analysts", who work with the experimental and theory communities to validate transport models; researchers engaged in analyzing experimental data or planning experiments, theorists studying micro-turbulence and its interaction with other physical processes (e.g. NTMs) and designers and planners for future devices, especially ITER.

2.4.2 Current physics and modeling status

Plasma microturbulence is the dominant mechanism for the transport of particles, momentum, and energy across magnetic surfaces. In the plasma core, it is believed that this is well-described by the coupled gyrokinetic/Maxwell equations. Still at issue is the level of "fidelity" for solution of these equations that are required to accurately predict plasma transport . A large number of codes have been written which have solved these equations at varying levels of approximation both within the US (GYRO, GEM, GS2, GTC, GTS, XGC1, etc.) and internationally (GENE,GKW, GYSELA, EURTERPE, GT5D, GKV, etc.). Neoclassical (drift-kinetic) transport is often the dominant mechanism for transport of plasma current. Other neo-classical effects (the ion heat flux) can be important in particular regimes and regions of the plasma. Neoclassical theory is well developed, and has been reduced to a set of coupled transport equations for density, rotation, and temperature profiles in many limits (e.g., NCLASS). Additional kinetic effects may be important near the magnetic axis or inside transport barriers. While gyro-kinetic and drift-kinetic physics may dominate profile structure in quiescent plasmas, MHD in the form of sawteeth, tearing modes, ELMs or "soft" b limits are often important in real plasmas. These phenomena are generally treated independently today, though they might be expected to interact strongly in many cases.

2.4.3 Challenges and approaches

Solving the core profile problem will require self-consistent, global solutions of micro and macro nonlinear dynamics on transport time scales including the effects of 3D field structures on turbulence. It will need to deal with mesoscale phenomena (between gyro-orbit and device size), that overlap with MHD scales. To reach the transport time scales, a model must solve the fundamental problem of disparate time scales. To simulate turbulence dynamics, the GK/Maxwell equations must be integrated on GK time scale ($v_{th}/R \sim \mu s$) which is much smaller than confinement times which are on the order of several seconds for reactor-like devices. The dynamic spatial range from ρ_i to a can be addressed on existing computers for short-time simulations. More demanding is the ability to span from the ρ_e to a scales though the utility of this has been a subject to much dispute. FSP codes can address these general issues in three ways at increasing levels of difficulty and fidelity:

1. 1.5D transport models, which describe the transport resulting from plasma microturbulence using (local or non-local) transport equations are integrated on the transport time scale.
2. Local gyrokinetic models, where gyrokinetic/Maxwell equations are integrated on a representative set of flux surfaces for sufficient time allowing the plasma microturbulence to reach a statistical steady-state on

the gyrokinetic time scale and yielding particles, momentum, and energy fluxes, to be employed to advance density, rotation, and temperature on the transport time-scale. This scheme has been implemented (at least) twice (TGYRO, TRINITY) with some validation of results against experiments on DIII-D and JET. While these implementations employed local (flux-tube) GK codes, the underlying idea can probably be made to work with global GK codes [see, for example, Shestakov, Cohen, Crotinger, Lodestro and Xu, JCP 186, 399 (2003)]

3. Global gyrokinetic models, coupled gyrokinetic/Maxwell equations are integrated over the entire volume of the plasma core (and thus addressing mesoscale phenomena). New theoretical or numerical formulation will be needed to be extend to the transport time scale

It has also been suggested to further develop gyro-fluid models which are more computationally tractable for full-radius, long-time simulations. Numerical and algorithmic advancements will be necessary to achieve the necessary computational parallelism required however.

2.4.4 A unifying theme: model validation

Common to all three approaches is the same scientific questions: "Over what parameter regimes can we reliably predict transport phenomena with computational transport models? " Making model validation the central theme leads to healthy iteration between experimental measurements, theoretical insight and code development. It provides a criterion (model validity) which we all accept for selecting one model, or one approach (1.5D/local/global) over another. The FSP could provide a vehicle for a greatly expanded transport validation program

Validation encompasses a set of scientific "Grand Challenges" associated with core turbulence and transport including:

- Origin of intrinsic rotation including the effects of SOL flows on the L/H transition and SOL flows and the role of flow dynamics in RWM stabilization.
- The dynamics of NTM growth and saturation with self-consistent turbulence including 3D magnetic equilibria (and providing a route toward stellarator transport).
- The stochasticity of the B-field on the micro-scale and its impact on turbulent transport.
- ETB and ITB formation and back transition including hysteresis effects.
- The formation of density profiles and the physics of the density limit.
- The locality or non-locality of turbulent transport.
- The nature of coupling between core and edge turbulence and transport.
- Turbulent transport near the beta-limit including soft vs. hard beta collapse.
- The role of energetic particles and turbulence interactions in setting confinement for a burning plasma.
- The regimes of importance for ρ_e-scale fluctuations in tokamak confinement.

2.4.5 Scientific roadmap

1. Start with a detailed comparison of current first-principles and reduced transport models through experimental validation. Address discrepancies in models of electron transport and momentum transport (by inclusion of appropriate GK formulations for calculation of momentum transport, radial electric field).
2. Evaluate current approaches to full radius and long-time simulations via local models. Produce reduced models for whole device simulations.
3. Develop approach to extract information from global simulations to allow prediction of profile evolution for temperature, density and momentum.
4. Treat mesoscale phenomena on transport time scales. Model interactions with neoclassical tearing modes (NTM) & other MHD and micro-turbulence - Include treatment of evolving 3D equilibria.
5. Characterize and incorporate boundary interactions between the core and pedestal including fluctuations and flows.

The requirements for code components and coupling algorithms would include:

- A description of particle, momentum and energy sources
- A description of the equilibrium magnetic field. We note that the equilibrium field evolves on the resistive time scale which is typically much longer than the transport time scale. The description should include a free-boundary equilibrium calculation so that the FSP plasma is able to interact with the tokamak's PF coil system.
- Coupling to the pedestal/edge. At a minimum, this would be a scalar boundary condition (e.g., density, temperature, rotation at the top of the pedestal) The need for further significant improvement in the characterization of core/edge coupling is recognized.
- A description of neoclassical transport .
- A sawtooth model to simulate transient profile evolution within the q=1 surface and to launch heat pulses at $q \geq 1$ generated by sawteeth .

2.5 Wave-Particle Interactions and RF

2.5.1 Summary of scientific issues

Superthermal particles including fusion products (alpha particles born at 3.5 MeV) and those resulting from RF and Beam heating and current drive are necessary elements in fusion plasmas. To achieve the required level of performance and to avoid localized heat loads on material surfaces, these particles must be thermalized with minimal energy loss. At the same time, the fast particles represent potent source of free energy for instabilities. The key challenges is a self-consistent description of the phase space distribution on long time scales (energy confinement or slowing-down) which are orders of magnitude longer than time scales for underlying Alfvenic wave-particle interactions. RF propagation and damping in the presence of these energetic particles along with the strong nonlinearities and mutual coupling to plasma transport through pressure, velocity and current profiles and fluctuation spectra must be addressed.

2.5.2 Introduction

The science of wave-particle interactions in tokamak plasmas encompasses the physics of collective instabilities induced by energetic particles and the physics of radio-frequency (RF) waves used for plasma heating and localized control of the pressure and current profiles. It is important to note that progress is needed in self-consistently integrating the RF physics and collective phenomena in order to predict the performance of fusion plasmas with RF heating in the presence of significant alpha and neutral beam populations.

For RF waves, the key scientific questions are whether waves in the ion cyclotron range of frequency will heat the core plasma efficiently and whether RF energy will be dissipated in the plasma edge through linear and nonlinear RF plasma interactions. In order to make progress on these questions, the key scientific challenges are to develop understanding and well validated simulation capability in (i) the physics of edge plasma interactions and associated dissipation mechanisms; (ii) the self-consistent evolution of multiple suprathermal ionic species with the RF waves in the plasma core and (iii) the stability of the plasma in the presence of suprathermal particle distributions and the resultant particle transport when instabilities are excited.

For collective instabilities, the central question is whether multiple unstable Alfvenic modes driven by suprathermal particles (alphas, beams, RF) in a burning plasma will lead to significant redistribution and loss of energetic particles. In order to address this issue the scientific challenge is to develop simulation capability that self consistently evolves the fast particle distribution with the collective instabilities on transport time scales and to integrate this into whole device simulation codes such as TRANSP.

Not only can externally imposed RF waves interact with energetic particle populations already existing in plasma due to fusion reactions and neutral beam injection, they also directly generate nonthermal populations of ions and electrons. These interactions can be both destabilizing and stabilizing and it affects the basic MHD properties of a burning plasma. A central need of the fusion program and a critical need for ITER is the development of an interpretive and predictive simulation capability to describe the conditions for the onset and consequence on transport time scales of collective instabilities induced by suprathermal particles generated by a variety of mechanisms.

2.5.3 User community and programmatic applications

An ability to quantitatively and accurately describe a whole-discharge involving the known sources and sinks of energetic particles (D-T reactions, RF waves, neutral beams), and their interactions with the spectrum of collective fast ion and background plasma instabilities will be of great utility to the entire fusion community. The research users will be (a) computational plasma physicists who will develop the predictive codes and virtual diagnostics, (b) fusion scientists who will attempt to validate the predictions of these models and use them to refine and plan experiments on existing tokamaks and ITER; (c) the theoretical community who will need to understand the limitations of numerical models and develop improved models.

For RF physics, the experimental community needs to focus on improved methods for diagnosing the range of dissipation mechanisms in the plasma edge suspected of limiting the effectiveness of RF plasma heating while the simulation community needs to develop virtual diagnostics to compare to the measurements. In the plasma core, additional efforts are required for a more comprehensive measurement of the RF field penetration into the plasma core and the resulting suprathermal particle distributions.

For collective instabilities, the understanding of nonlinear phenomena will require high-resolution measurements of the phase space distribution of the energetic particles as well as continued improvements in the measurement of the internal mode properties.

2.5.4 Current physics and modeling status

Although a great deal of progress has been made in understanding the linear instabilities in the MHD limit, comparatively less progress has been made in understanding the resulting transport in the presence of such instabilities.

A great deal of emphasis has been placed on the development of particle simulation codes, either particle-in-cell (PIC) or hybrid-MHD kinetic, which are very good for describing the fast nonlinear development of energetic particle driven MHD modes. However, these codes have an intrinsic difficulty in predicting behavior on longer transport time scales. For example, none of the state-of-the-art codes yet predict the fast ion distribution on a slowing down time scale under the influence of plasma instabilities. This is true both for RF heated and non-RF heated plasmas. In addition, the MHD description does not capture the finite FLR effects that are expected to modify the instability leading to important damping mechanisms in burning plasmas. For this reason, efforts have been made to develop gyrokinetic simulation capability as well as analytic and numerical two fluid and kinetic models.

In the area of RF wave-particle interactions, continuum descriptions (both full-wave and ray tracing) have been successful in describing core wave propagation in minority heating and mode conversion regimes. Zero orbit width Fokker Planck codes have been used to describe the evolution of nonthermal electron distributions and in some cases nonthermal ion distributions generated by RF waves. However finite ion orbit effects have only recently been treated in the RF wave – fast ion interaction by including an RF acceleration operator in Monte Carlo orbit codes. Wave propagation and absorption at high harmonics of the ion cyclotron frequency has been more problematic to describe successfully in large part because of the tendency of those waves to also interact with fast ion populations such as neutral beam ions, that are already present in plasma. The interaction of RF waves with the edge plasma is still in its infancy in that most models employ simplified boundary conditions that do not account for the complicated geometry of actual launching structures and nonlinear effects such as RF sheath generation or parametric decay instability. Thus the integrated edge – core RF problem has not yet been successfully carried out.

LIST of existing codes: M3 D, Wisconsin Code (developed as LASL), GYRO, Irvine code, GEM (Colorado), NOVA, AEGIS (IFS), Spong's code, Gorolenkov's quasi-linear code, GTC. AORSA (ORNL), TORIC (IPP and MIT and PPPL), CQL3D (CompX), ORBIT RF (GA), sMC (ORNL), TOPICA (Torino), LHEAF (MIT), VORPAL (TechX).

2.5.5 Roadmap

A. Roadmap for developing the science of RF wave-particle interactions:

1. At present the integrated edge to core RF problem requires theoretical development in order to advance the state of the art. This is because the best approaches for treating the complicated edge and launcher geometry are generally finite element methods (FEM), whereas the best treatments of the core absorption are usually spectral basis codes, as the dielectric operator becomes algebraic in k-space. Theoretical work is needed to determine if the two approaches can be combined; if the edge can be treated accurately in a spectral code or alternatively if the plasma conductivity operator can be formulated in 2D and 3D in an appropriate FEM basis.
2. Solving the RF edge to core problem above should lead to significant progress in describing the linear coupling of RF waves in the ion cyclotron range of frequencies (ICRF) and the lower hybrid range of frequencies (LHRF). However, further theoretical work will be needed to account for nonlinear effects, due to the presence of RF power in the edge, such as parametric decay instability, RF sheath formation, and local modification of the edge plasma via RF induced ponderomotive forces.
3. In the core plasma, theoretical work is needed to describe how ICRF and LHRF waves drive toroidal plasma rotation and plasma flows and their subsequent effect on transport. In particular, existing theoretical models for plasma rotation via mode converted ICRF waves have been found to grossly under estimate experimental observations when implemented in simulation codes. Similarly, it is not clear theoretically how much of the observed increase in toroidal plasma rotation that occurs with ICRF and LHRF power is due to the direct injection of momentum versus RF induced changes in the intrinsic rotation mechanism of the plasma.
4. In the core plasma, theoretical work is now on-going and must be continued on the development of kinetic closure relations for the hierarchy of MHD equations in order to describe the stabilization of neo-classical tearing modes via localized RF current generation and the stabilization of sawteeth and RWMs via ICRF generated energetic particle distributions in the presence of alphas and beam ions.
5. Finally the inclusion of finite ion orbit width effects in both minority ICRF heating and at high harmonics of the ion cyclotron frequency should be revisited theoretically. At present, full wave solvers and Monte Carlo codes are used to treat this problem in a somewhat brute force fashion whereby statistical particle lists are passed from the orbit codes to the full-wave solvers and a 4-D quasilinear diffusion coefficient is passed from the field solver to the orbit code. Alternate approaches need to be developed whereby finite ion orbit effects can be included in continuum codes that solve the Fokker Planck equation directly.

B. Roadmap for collective instabilities driven by energetic particles:

1. At present the collective fast particle instabilities in experimentally relevant geometries are mostly described by ideal MHD codes although a great deal of progress has been made in recent years to include thermal and fast particle effects on the mode structure and damping. In order to quantitatively understand and predict excitation thresholds in burning plasma conditions, additional effort is required in the development of two fluid, gyrokinetic and fully kinetic treatments of Shear Alfvénic Wave instabilities. In the next 2 years significant progress can be made in the description of FLR effects on Shear Alfven Wave instabilities in full plasma geometry using two fluid and gyrokinetic models.

2. Solving for the nonlinear evolution and transport of fast ions in a field of Alfvénic instabilities using the known sources of fast ions as the boundary condition is essential for predictive modeling of ITER plasmas. This effort must proceed in parallel with improvements in the linear description of the instabilities. It is unrealistic to assume that fully nonlinear codes can address this need for whole discharge modeling in the next 4-5 years. However, reduced models, taking as a starting point the linear eigenmodes generated by linear mode solvers, can address the key wave-particle nonlinearities on transport time scales. Such reduced models can be developed in a 2-3 year time frame. The limitation of such descriptions is that the nonlinear modification of the mode structure cannot be addressed in the initial stages, however the mode amplitude and frequency evolution (bursting/chirping) and resultant transport of particles can be addressed on the required energy confinement time or slowing down time scale. Such models can then be integrated into whole discharge simulation codes to address the impact of instabilities on losses and discharge evolution. As advances are made in the linear description of the instabilities, then these can be readily integrated into the reduced model. On the 3-5 year time frame, the reduced models can also incorporate wave-wave nonlinearities, such as zonal flows that can then impact on thermal transport. This effort would require close interaction with experiment for validation of theoretical predictions of mode amplitudes and fast ion transport. It will also require the development of numerical diagnostics for the resolution of the phase space dynamics of the suprathermal particles.
3. On a 3-4 year time frame, reduced models can be incorporated into RF codes to evolve the fast particle distribution under the simultaneous influence of RF fields and Alfvénic instabilities. This is important in regimes where a spectrum of Shear Alfvén waves are expected under RF heating conditions. This activity is complimentary to existing efforts under the SWIM project to address the interaction of RF waves and fast ions with transient MHD events such as the sawtooth crash.
4. In parallel with the above two activities, fundamental advances are needed in the theoretical formulation and numerical implementation of fully nonlinear codes to capture the evolution of the fast ion distribution on transport time scales. At present fully nonlinear codes are not capable of integrating beyond several periods of the instability and that is only for modes that are strongly unstable. In realistic systems near marginal stability, where the resonance widths in phase space are very narrow, much higher phase space resolution is needed to capture the mode dynamics near threshold. In addition new methods need to be developed to capture the dynamics of large-scale transport events, such as avalanches, where there is no spatial scale separation between the perturbed distribution and the system size and where the perturbations can be a large fraction of the initial distribution. In the 3-5 year time frame, fully nonlinear codes can be used to validate reduced models on short time scales and they can be used for more detailed analysis of certain plasma states that our output from a whole device simulation code using reduced models. On longer time scales, beyond 5 years, advances in fully nonlinear codes may begin to address systems on transport time scales.

2.6 Disruption Prediction, Avoidance, Mitigation and Effects

2.6.1 Scientific issues

The main scientific goals are to 1) obtain an improved predictive capability for the onset of disruptions to aid in avoidance and to aid in the development of algorithms for triggering disruption mitigation; 2) 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 might enable the robust operation of tokamaks by allowing more aggressive operating regimes and by enabling faster recovery from off-normal events. The effects of disruptions include severe heat loads, JxB forces and run-away electron generation. 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; 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.

2.6.2 Introduction

A disruption is the rapid termination of plasma current and stored energy. Disruptions may be triggered by a variety of instabilities and are often associated with operational limits for plasma current, pressure or density. The triggering may be through ideal or resistive MHD modes sometimes driven by equilibrium profiles but also through loss of plasma control or injection of impurities. Whatever the sequence of events, the final steps involve coupling to external kink modes, break up and 3D distortion of magnetic surfaces and rapid transport of plasma energy. Development of improved avoidance techniques and of appropriate mitigation triggers will require analysis of disruption data to try to identify all of the causes of disruptions, and modeling of disruption onset to confirm the identification of the causes and to determine the requirements for avoiding or triggering mitigation of disruptions. To simulate the dynamics of disruptions, three distinct phenomena need to be modeled, a thermal quench during which a global instability causes the loss of most plasma kinetic energy resulting in a large drop in temperature; a current quench in which the now resistive plasma transfers current to the vessel walls or to superthermal electrons; and generation of runaway electron populations driven by the large voltage induced from the rapidly decaying plasma current. The thermal quench presents a large, transient heat load to the first wall typically in a localized manner and can melt or evaporate significant material. The current quench results in large JxB forces, typically in a non-axisymmetric pattern, threatening the mechanical integrity of the vacuum vessel and other in-vessel components. The runaway electrons, which in theory could carry a substantial fraction of the original current, have the potential to create significant localized damage to first wall components.

2.6.3 Current status and key gaps

Extended MHD Codes have the ability to model the initial stages of rapidly growing MHD instabilities, and to provide a guide to the conditions under which such instabilities are triggered. Progress is being made on the generation of stochastic magnetic field caused by instability growth; rapid parallel transport of plasma in stochastic field; physics of the open field lines region and currents induced using a resistive wall model. Relatively crude models for impurity radiation and transport have been incorporated to simulate disruption mitigation experiments. For detailed modeling of the wall response to transient heat and particle loads, the HEIGHTS package is available, consisting of coupled codes for computing plasma transient deposition on surfaces, vapor formation, radiation transport, atomic data, MHD, and surface thermal conduction and hydraulics. Validation of these models is not being carried out in a systematic manner though triggering of disruptions by ideal MHD and loss of vertical control is relatively well studied.

Models need improvement in many areas:

- Disruptions triggered by tearing modes, resistive wall modes and sawteeth are not well understood and are challenged by the disparate times scales introduced by these non-ideal phenomena. Simulations will need to include kinetic effects through gyro-kinetic or hybrid approaches beginning with studies in static magnetic fields, progressing to investigation of self-consistent growth of magnetic islands then working up to full kinetic studies in stochastic magnetic fields. A key issue to be addressed here is the effect of self-consistent 3D electric fields on the behavior of low collisionality plasmas in regions of mixed stochastic field lines and magnetic islands. Rotation is another key piece of physics that needs inclusion.
- No 3D fluid code exists which includes the effects of neutrals, impurities and wall interactions. The required steps involve 1) improved boundary conditions including models for plasma sheaths and the electromagnetic properties of the material wall, 2) the transport and atomic physics of an arbitrary number of impurity species in 3D fields, 3) a fluid model for neutrals interacting with the plasma and the wall, 4) coupling to reduced wall models, 5) coupling to more sophisticated wall models (e.g. HEIGHTS). This area of development has strong overlap with the needs of the boundary physics science driver.

- A computationally tractable theory for coupling Fokker-Plank solutions for runaway electron back into the MHD solutions is needed. These also require more complicated models for the collision operator including relativistic effects, knock-on electrons, etc. The coupling of MHD to Fokker-Plank calculations needs to include electric and magnetic-field induced radial transport and to take into account the effects of runaway electrons on plasma resistivity and field evolution.
- There is a need to produce reduced models which can be used in conjunction with free boundary transport codes. Codes currently have difficulties solving through the thermal and current quench. It is worth noting that the effects of radiation impurity are likely to help alleviate this situation, especially for the thermal quench phase. The computed fields do become more symmetric as the simulation progresses.

2.6.4 Roadmap

The science development roadmap requires capabilities which are necessary to answer a set of seven scientific questions.

1. How well can we predict the onset of a disruption and what strategies are available to avoid their development?
2. What are the effects of runaway electrons and what is the impact of operating regimes on their generation?
3. How does impurity transport affect disruption dynamics, and how do we use this information to mitigate the effects?
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?
7. What are the best plasma models for simulating plasma disruptions?

These are described in somewhat more detail below:

1. Disruption Prediction and Avoidance: Develop models of the plasma, including 3D equilibrium reconstruction (capable of handling magnetic islands) in the predisruption phase, to understand the causes of disruptions, avoid disruptions or to trigger mitigation. For this purpose it will be helpful to develop a "2.5D" transport code, capable of following 3D plasmas on transport time scales. Many other elements of these calculations are connected to efforts required for other science drivers. One of the critical elements for this is accurate prediction of core profiles requiring integration of nonlinear turbulence and nonlinear MHD over the full (core) radius and on long time scales. Another element, the use of RF for control of tearing modes and sawteeth, is largely covered in the wave-particle driver. These and other control strategies will likely be integrated through "whole device" modeling. The FSP will need to define how these activities will mesh with work to be carried out under the disruptions driver.
2. Generation of Runaway Electrons: Runaway electrons have been modeled in two ways: First, using the Fokker-Planck code CQL3D with crude models for radial transport due to stochastic field but fairly a complete model of electron equations of motion and second by integrating a simplified electrons equation of motion in extended MHD code with three-dimensional magnetic fields and no feedback of these electrons on the MHD dynamics. In the near term, theoretical development is required to produce a computationally tractable model for runaway electron feedback onto the MHD equations similar to that used for energetic ions. These models would take magnetic fields from MHD simulations for use in the Fokker-Plank codes. On a similar time scale, development could start on developing models suitable for integrating runaway electron modeling into transport codes through parameterization of magnetic fields for various disruption scenarios. Other reduced models are possible (e.g., Zakharov's Kadomtsev-Pogutse shell model extensions). For the longer term, it will be necessary to solve for the distribution function of runaway

electrons in 5D space using the drift kinetic equation (DKE) and to integrate this other elements of the calculation.

3. Effects of Impurities, Radiation and Neutrals: Integration of MHD with models for impurity radiation and transport are necessary to calculate heat loads from the thermal quench and to simulate runaway electron mitigation. A first step would be to create standardized libraries for the relevant atomic physics cross-sections with a uniform API for use by MHD and PMI codes. This should be a cross-cutting area useful to several science drivers. Standardized verification and validation cases will need to be developed in collaboration with members of the edge community and experimental teams.
4. Plasma Material Interactions: The extended MHD codes have primitive wall interfaces compared what is available in the edge community and the most developed PMI model, HEIGHTS, has relatively primitive plasma models. To begin bridging the gap, we would collaborate with the boundary physics community to implement sheath boundary conditions and reduced wall models, developed as part of the boundary physics science driver, into the disruption codes and verify the implementation. At the same time, fluxes from MHD simulations would be used by wall models to improve/verify their plasma-wall physics. In the longer term more complete integration of these models would be carried out. Many detailed physics issues, for example, implementation of "kinetic boundary conditions" have similarity to developments needed by the edge transport community.
5. Structural Effects: Forces on the wall are due to conductive and inductive currents in comparable measure. Extended MHD analysis useful and has shown qualitative agreement with experiments, but has limits. For example the wall used in current models is axisymmetric but the real wall is a complicated 3D structure and could lead to localized forces much larger than those calculated. ITER engineers would like real forces on the structures described by their CAD models. In the near term, structural forces in non-axisymmetric walls could be calculated using 3D fields from extended MHD simulation performed with symmetric walls. In the longer term, self-consistent models will be needed and should include heat transfer, electrical, and structural calculations.
6. Modeling of Delivery Systems for Disruption Mitigation Systems: While every effort will be made to avoid disruptions on large reactor-scale devices, it is considered essential to have a realistic and reliable strategy for detecting and mitigating the effects of disruptions when they do occur. Models for disruption detection diagnostics and for disruption mitigation actuators will be critical. Current models of gas jet injection are fairly simplistic and while more complicated models for pellet and gas jet deliver exist (P. Parks) they are difficult to implement. More tractable models will need to be developed and verified. Calculations for gas mitigation may require integration with standard CFD (computational fluid dynamics) codes.
7. Improved Fundamental Models: Further improvements in fundamental plasma and wall models are expected as part of the base theory and computation programs and by other FSP science drivers. This raises the question whether there is something about the physics of disruptions that requires more accurate models. For example:
   - Will electron transport in a stochastic field be well-described by the drift-kinetic equations (DKE)?
   - Will ion transport in a stochastic field be well-described by DKE? (Because of orbit size, this is more problematic, especially for energetic particles).
   - Self-consistent models of the complicated 3D ambipolar potential in regions of mixed stochastic field lines and islands will need to be developed.
   - What is the best model for disruptions beyond an extended MHD model? For example will it be necessary to develop a kinetic-MHD hybrid code.
   - A near term step will be to evaluate gyrokinetics in plasmas with static magnetic islands and stochastic fields and to explore whether the imperfect field alignment of the grids in stochastic magnetic fields is an issue for gyrokinetic codes.

2.6.5 Summary

Systematic modeling of the plasma evolution in tokamak experiments leading up to the full range of different types of disruptions encountered in tokamaks will be needed to gain an understanding of the possibilities for disruption avoidance and for triggering mitigation. And while significant progress has been made under the base program toward better understanding of disruption dynamics, several important questions are still unresolved. Most of the weaknesses can be attributed to lack of integration among the disparate disciplines which are required. Future work under the FSP, can strongly leverage developments required by other Science Drivers for example in modeling of runaway electrons; radiation, impurities and neutrals; plasma wall interactions; integration with engineering calculations and mitigation actuators. Validation will have to be pursued in a much more systematic manner and move from the qualitative to quantitative level.

Summary of desired features for plasma model

- Validity in long-wavelength regimes
- Inclusion of stochastic fields and convergence without a field-aligned grid
- Free-boundary (beyond separatrix)
- Resistive wall
- Impurities transport and radiation
- Proper treatment of sheath boundary conditions
- Radiation transport (both losses and ability to solve in optically thick regime)
- Accurate model electron transport
- Accurate model ion transport

2.7 Whole Device Modeling

2.7.1 Scientific goals

The ultimate goal of the Whole Device Modeling (WDM) science driver is to provide a comprehensive model of the whole device, modeling plasma from the magnetic axis to the wall, for the entire discharge, integrating all relevant physics, and including conducting structures, magnetic materials, and magnets. The WDM is used to perform simulations of existing experiments and future devices, like ITER, integrating the multi-physics of the fusion plasma devices and identifying important "integration" physics issues.

The user community for the WDM tool is broad. Experimentalists use WDM codes to interpret experiments (such as TRANSP or ONETWO) by reading experimental data and enforcing the governing transport equations, although most of these users are removed from the code development. Experimentalists and modelers may use a predictive WDM tool to reproduce an experiment and then extrapolate to new experiments. Modelers may use the WDM tool to simulation discharges on future devices (like ITER), based more heavily on models and less on experimental/empirical bases, and this community is far more concerned with the contents of the WDM code and its development. The WDM tool is often used for testing of physics models either individually, or in self-consistent combinations against experimental data. The WDM tool can provide data for other analyses not included in the WDM, it can be used to develop plasma control approaches and optimize plasma configurations, and even become a tool for between shot discharge analysis.

The development of a WDM tool will closely follow the progress of other science drivers, which provide new models (or components) to be implemented into the WDM structure. Based on present progress in the proto-FSP projects (SWIM, FACETS, CPES), the overall view of the WDM would be a framework which contains many components representing the various physics models that are integrated together to form a whole device model. The framework is the interface to which the user comes in contact, and is the driver for the simulation. This framework and the various connections between it and the components, as well as between the components themselves, are dictated by needs of the physics description, but may also depend on optimization of the computing environment.

The functions of the WDM tool include:

1. A platform for predictive simulations, interpretive analysis of experiments, and experimental data comparison for validation
2. A framework that allows coupled components (physics models) and multi-scale integration for time-dependent simulations
3. Multiple models of the same physics at varying levels of physics fidelity
4. Accommodation of production running (large numbers of runs of low to medium computational needs) and high fidelity demonstrations (few runs of large computational needs)

The most basic components in the WDM for the core plasma simulation are an equilibrium solver and transport evolution solver, coupled with sources/sinks, coupled with models for the transport parameters, and coupled to numerous additional physics models. Initially the equilibrium solver will be 2-dimensional, and the transport equations will be 1-dimensional (radial). At first it is expected that legacy codes will be utilized for many components in the WDM, however, it is recommended that isolation of the best models, from among various existing transport evolution codes (e.g. TRANSP/PTRANSP, Corsica, TSC, ONETWO, etc.) be done soon. An additional component that is considered necessary is a "state of the plasma" description at a given time slice, providing the minimum information required by virtually all physics models, for example the plasma equilibrium geometry and thermal species transport (temperature, density, momentum, and current) profiles. An interface between existing experimental data (generally stored in a database structure, like MDSplus) and the WDM framework is necessary to facilitate the comparison of model simulations with the experiments, and is preferably established with multiple devices to accelerate cross-device comparisons.

The models that are used for everything from a heating and current drive source to the bootstrap current can have varying levels of physics fidelity, that is, all models are not first principles complete at all spatial and time scale descriptions of the physics phenomena, but more often are "reduced" or lower fidelity models (some approximations have been made). It is important to recognize that the WDM tool is viewed as an evolution over time as the various models become better and the computational platforms become faster, allowing even better models, and so forth. For faster simulations, the lower fidelity models are preferable, however, these can be tested by occasional high-fidelity model evaluations at time-slices during or after a simulation. The model to model validation is an important part of growing the WDM tool to be both flexible and representative of our state of the art simulation capability.

The WDM is the likely tool for examining the validation of a physics model against experimental data. This validation can be performed for a single physics model or done in combination with other physics models, where anything not modeled would be prescribed by the experimental data. A particularly new aspect provided by the WDM would be time-dependent validation, as opposed to the common single time-slice validation often performed today. In addition, validation would be done across multiple tokamak devices allowing much more stringent tests of physics. This validation process would benefit significantly by the development of appropriate synthetic diagnostics.

The WDM tool itself will have varying levels of fidelity based on the user's needs. These needs are largely determined by turn-around time, with faster turn-around requiring the use of lower fidelity models, and vice versa.

A list of typical (desirable) physics models (or components) in a WDM are:

- 2D equilibrium solver
- 1D transport solver for (temperature, density, momentum, and current)
- ICRF/LH/NB/EC/alpha heating, current drive, and momentum sources

- Pellet/NB/gas injection/wall particle sources, and divertor pumping/wall sinks
- Transport models (empirical, GLF23, MMM08, TGLF, GYRO, …)
- Bootstrap current model (NCLASS, Sauter, Neo)
- Pedestal model (empirical, ballooning, EPED, ELITE)
- Radiation (bremsstrahlung, line, cyclotron)
- Fusion reactions
- MHD models for NTM's, RWM's, sawteeth
- Fast particle models for *AE's and other energetic particle modes
- Poloidal field coils, conducting structures, and feedback systems
- Neutrals/atomic physics
- Scrape-off-layer plasma (SOL/divertor)
- Plasma wall interaction physics

2.7.2 Roadmap

1. Implement reduced plasma models for all relevant phenomena in quiescent, axisymmetric equilibrium, including off-line verification and validation of reduced models against experiments and high-fidelity codes.
2. Establish mechanisms for coupling to high-fidelity models running on parallel architectures during time-dependent simulations. Start with turbulence models.
3. Implement tight-coupling between core and edge plasmas including pedestal, scrape-off layer and plasma-wall interactions.
4. Include high-fidelity models for interactions between fast and thermal particles, waves, instabilities and turbulence.
5. Implement 3D free boundary equilibrium that can handle magnetic islands, stochastic regions, RF, nuclear and atomic physics modules.
6. Include nonlinear extended MHD models for disruptions, sawteeth, ELMs, etc.

3 Software Integration and Support

3.1 Introduction

This section is a summary of the results of the FSP March Planning Workshop for the requirements and process in the areas of Software Integration and Support. Software integration includes the composition of software, both advanced physics components and utilities for job preparation and results analysis. Software support includes the infrastructure for software development as well as the infrastructure needed for supporting users of the software. We break these down into the four areas of (1) Physics Composition, (2) Task Composition, (3) Development Processes, and (4) Production Computing, which are defined in more detail below. In the last two sections we discuss some of the results on process: how to select from competing solutions and how to manage change as the FSP proceeds.

Software integration is discussed as two areas: (1) Physics Composition [also known as High-Performance Computing Composition or HPCC] is the composition of Physics Components (physics computational modules) in a manner suitable for execution on a parallel computer, including anything from a multi-core, single CPU machine to a Leadership Class Facility (LCF), such as a Cray XT5 or IBM Blue Gene. In terms of batch computing as done on LCFs, Physics Composition is the composition that applies after being released from the batch queue to begin computing on the LCF and prior to job completion. (2) Task Composition (or workflow) covers both the utilities and composition software for everything surrounding this, from concept to research discovery and its reporting and includes items such as job preparation and visualization. That is, Task

Composition includes the composition up to submission to the queue and after job completion. It can also include composition that involves multiple job submissions, e.g., as for optimization with high-performance computing in the inner loop.

Software support is also discussed as two areas: (3) Development Processes includes all aspects of the support needed for developers to carry out their work. It includes items such as revision control, collaboration methods and technologies (e.g., developer mailing lists and wikis), regression testing bug tracking, development environments including cross-platform build systems and package management. It also includes documentation, such as API documentation and/or architecture specifications that are needed by developers. (4) Production Computing includes all aspects of user support, including providing an easy environment for computing, responding to user problems, job monitoring, user help mailing lists, and triage for user questions (which may ultimately need to be referred to developers). It also includes deployment of the software packages developed from Physics Components, Physics Composition, and Task Composition.

3.2 Physics Composition

Physics composition is the assembling of individual physics components and the accompanying coupling software (coupling algorithms, data interpolation, etc.) necessary for the execution of an integrated physics simulation and necessary for meaningful verification and validation of such simulations.

At the FSP March Planning Workshop, the requirements for successful physics composition were discussed in the context of some of the individual science drivers. From these discussions a number of generic, cross-cutting requirements emerged; these are summarized below. There are additional requirements for individual drivers. We anticipate that more detailed analysis of the science drivers, as planned for the remainder of the FSP Definition Project, will uncover additional requirements and also provide more specificity to the requirements already identified. The requirements listed were identified as necessary or at least highly desirable for composing integrated physics simulations. The list includes requirements on individual physics components and requirements on the composition software (or framework).

3.2.1 Requirements on Physics Components

Components should come with documentation of sufficient detail so that component users can determine

- The equations solved by the component along with the solution methodology and order of error in temporal and spatial discretization;
- The limits of validity, whether imposed by the physics approximations, the algorithms, or other factors;
- The architecture and/or internals so that modifications needed for FSP physics composition (e.g. to implement implicit coupling) can be made if need be by someone other than the component developer.
- The parallel capabilities of the components, including how many processors can be used for a problem of a given size, any required processors counts (e.g., multiples of 64 or another number)

Components must be "componentized": they must be provided with an agreed-upon interface (e.g. through wrapper code) that provides, for a given component type (functionality), agreed upon ingredients:

- Common input and output data quantities
- Common input and output variable names
- Common methods

In addition there must be agreed-upon definitions for what constitutes a component; in particular, the agreed-upon level of granularity. In particular, an answer must be provided to the question, "is it a sub-component of another component, or a component" – which will typically depend on the reusability of the (sub)component in question.

- Components should be provided in source form so that they can be modified as needed for bug fixes or to fit within the physics composition framework, whether the framework requires minimal modifications (e.g., changes to I/O) or more extensive modifications.
- Components should come with a suite of test problems for each component along with accepted results. Tests for individual components should exercise the features of the component, along with results. Tests should be repeatable at the numerical precision level, thus implying at minimum control over message patterns and quasi-random number sequences. This control must be exposed to the framework; in particular a component should be able to accept a specified seed from the framework to enable insensitivity of coupled simulations to small parameter changes (see composition software requirements below).
- Components should provide for a mode in which their results are insensitive to small changes of initial parameters. In the case where a component has multiple chains of random events in a component, a technique must be provided to ameliorate the consequences of possible changes in the number of random number invocations in a particular chain. This requires some technique within the component such as combing [Monte Carlo Methods, Vol. I, Mal Kalos and P. Whitlock, Wiley-Interscience (1986)] to reduce the noise.
- Components should provide provenance, i.e. sufficient information to achieve reproducible results, including source code version, versions of external libraries, and compilers and their versions.
- Components should be able to check-point restart and integrate forward precisely as if they had not been restarted.
- Components should be able to revert to the state prior that that of the current time step (as needed for implicit coupling).
- Components should be able to exit gracefully and provide an error code.
- Components should allow for specifiable input and output file names.
- Components should provide output in a widely used, cross-platform binary format (e.g., NetCDF or HDF5), and the output should come with sufficient metadata for automatic visualization of the data.
- Components should not have hard-wired input/output; in particular there should be settable log files (versus e.g. sending output to stdout).
- Components should have any embedded graphics disabled, but with data necessary to produce component graphics included in what is available to the framework.
- Components need to work on a common set of (preferably multiple) platforms and come with a cross-platform build system for building them on multiple platforms.

Physics component precision should not be set by a specification on its compilation line, but rather, precision should be specified in the source code.

3.2.2 Requirements on High-Performance Computational (Physics) Composition Software

- The composition software must include infrastructure that supports tight (in-memory) two-way coupling, and algorithms that achieve this coupling implicitly. A number of examples have emerged that illustrate the need for implicit coupling, and there is growing recognition throughout the fusion community that this will be required for many of the physics integrations contemplated for FSP. Some examples include:
  - Coupling of turbulence and transport: the transport equations need to be advanced implicitly to avoid severe timestep constraints. Specific solutions have been demonstrated for flux-tube (local) continuum gyrokinetics and transport; one approach (with two variations) for global (nonlocal) turbulence has been demonstrated for simplified fluid codes but not yet for a full-physics gyrokinetic code. It is an open question how this need can be met for PIC codes regardless of locality and for any code when timescale separation fails (likely the case for blob-dominated edge transport, and for the core as well if avalanche phenomena and/or coupling to mesoscale MHD is important).

  - Coupling of Monte Carlo neutrals to plasma transport
  - Coupling of core and edge transport
  - Plasma transport in the presence of an evolving magnetic field (an outstanding challenge in the case of particle-based transport).
  - Coupling of extended MHD or two-fluid models to kinetic codes for closure and sub-grid effects.
- Coupling algorithms that have applicability to more than one type of coupling should be made available as components.
- The Physics Composition Software should be able to test components individually and provide the data needed for testing components outside of the framework with any native drivers.
- Metrics for convergence of coupled simulations must be provided. Standards, and implementation of these standards, are needed to determine when a coupled simulation has converged. This is a particularly challenging issue when one or more of the coupled components is noisy (e.g. from a stochastic process); that leads to the following additional requirement:
- The composition software must provide for control over random numbers used by components (and any that might be introduced in couplers). The ability to do meaningful verification and regression testing of coupled simulations with components that utilize random numbers requires that the "random" numbers be reproducible. Under component requirements it was noted that components need to have (and expose to the framework) an option to run with frozen random seeds; the framework needs to be able to exercise this option. Moreover, the solution of a coupled simulation should be insensitive to small changes of parameters, which could conceivably change the number of iterations of a component in a timestep; to achieve insensitivity, the framework must be able to provide (and the components accept) consistent seeds at the start of each timestep regardless of the number of iterations at the preceding timestep.
- The composition software should provide for conservative data interpolation between different grid representations. This is particularly important (required for consistency) when there is frequent data exchange between components operating on different grids (volumetric data coupling), for example in coupling an evolving MHD equilibrium to a core transport solver. Other examples include transport+turbulence, microturbulence+macro-MHD, MHD+RF, plasma transport+neutral transport. It is particularly challenging when one or both of the components operates with an adaptive grid.
- The composition software must provide access to software that can efficiently refactor the parallel decomposition of data, between the different parallel decompositions that may be used by different components (the "MxN problem"). Different components may be decomposed differently because of differing degrees of granularity, different number of parallelized dimensions, decomposition by groups of particles (for PIC) versus domain decomposition for fluids, etc. The software suite must provide for solutions brought in by individual component developers as well as solutions developed for FSP.
- Physics composition software should be able to check-point restart as well as orchestrate the check-point restart of all components in a computation.
- Tests for compositions should exercise intended use cases for the combined suite of components. If there are components utilizing random number generation, a provision for specifying running with frozen seeds must be provided.
- The composition software will need documentation for
  - the physics integrations that can be computed,
  - the coupling algorithms, including data interpolation, component execution model (sequential, concurrent),
  - the parallel decomposition,
  - component interfaces, including the definitions of quantities and units,
  - sufficient documentation of composition software internals that modifications can be made by someone other than the component developer.

3.3 Task Composition (Workflow)

As noted in the introduction to this section, Task Composition (or Workflow) covers both the utilities and composition software for everything surrounding Physics Composition, from concept to research discovery and its reporting and includes items such as job preparation and visualization. That is, in the context of LCFs, Task Composition includes the composition up to submission to the queue and after job completion. In this section we describe the requirements in Task Composition for the FSP that came out of the March Planning Workshop. In the process, we also describe some of the applicable technologies that have been generated by the proto-FSPs.

3.3.1 Requirements for Task Composition (Workflow)

- Program-wide workflow software. This will have to be evaluated by the users to see what is needed, as there are multiple approaches of differing complexities. For example, the FACETS proto-FSP generally develops their workflows in bash, while CPES uses the Kepler scientific workflow engine, which works from a workflow description in XML and is written in Java. There are advantages and disadvantages to the more programming oriented approaches using bash or python versus the more rigid data driven applications like Kepler. For example, Kepler has provenance collection mechanisms, but it is not clear how to use it on LCF machines, which have batch queue systems. Thus, it seems advisable at this stage to ask that FSP support two technologies, one more restrictive and one more programmable, until redundancies are addressed, e.g., through user choices.
- Provenance, or the recording the history of generating a particular element of data. This is the equivalent to the experimentalist's logbook information.
  - The ProtoFSPs have developed different approaches. 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. This indicates that provenance may be needed in several places, including the original data files written by the Physics Components and Composition and also captured and recorded by the workflow.
  - The next stage of sophistication there is to provide a common provenance storage mechanism. This will need to be integrated with the overall Integrated Data Management (IDM). One example is a Catalog in the IDM, with a defined Schema for describing general provenance.
  - To facilitate provenance collection, there is a need for a Provenance API that FSP tools would use to register provenance information. This API will need a binding to the most common workflow languages, like python, bash, C, C++, and Fortran. 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.
- Input file preparation, staging and checking
  - 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), PlasmaState (TRANSP/SWIM NetCDF format), "Fit files" (files that describe profiles used in EFIT fits needed to get profiles across separatrix. Input data needs to be placed in suitable places for codes to access this information at run time, this is staging. The Facets Composer 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.
  - The next stage of development in 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 the DataMover and the 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.

- FSP will have to support some of the few I/O libraries that provide sufficient support for computing at scale. Those include HDF5, NetCDF, and newly being developed ADIOS. While it is not expected that further work on these libraries will be supported, support may be needed for simplified APIs to add FSP required metadata for visualization, provenance, and other items.
- FSP will need to support tools for taking the data files on the LCFs and migrating it, in whole or in part, to the Integrated Data Management system.
- Both non-interactive (batch, automatic) and interactive (exploratory) data analysis and graphics generation are required by the Fusion Simulation Program.
    - Non-interactive analysis and visualization are expressed as scripts in the existing tool scripting languages (matplotlib, VisIt scripts, IDL, MatLab), 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.
    - Interactive data exploration and analysis can make use of either the scripting tools mentioned above or tools like VisIt and AVS, which have graphical user interfaces.
    - For monitoring (below) and because not all data will likely be put into the Integrated Data Management system, the visualization tools should be able to work directly on FSP generated data files.
    - For visualizing data after it has been moved to the Integrated Data Management system, the visualization tools will need to work with the IDM access API.
    - The proto FSPs have shown two approaches to monitoring. FACETS uses a VisIt back end compute core to generate graphics to send to the Facets Composer monitoring and visualization tool. CPES uses a collection of vector graphics routines to forward image data to a Flash™ embedded graphics viewport in the eSimMon monitoring service
    - The FSP should work towards limiting the number of visualization tools supported to minimize development requirements, but at the same time FSP will have to recognize that the multiple visualization packages have unique capabilities, and so there will have to be support for a number of packages.

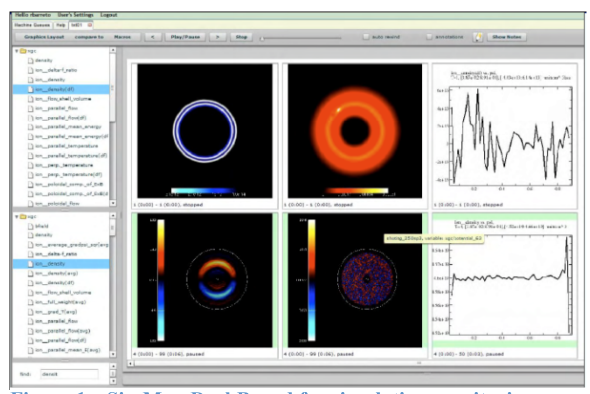

Figure 1. eSimMon DashBoard for simulation monitoring

    - The FSP is not expected to create visualization tools but instead to support the use of existing tools in the FSP. This includes the development of IDL/VisIt/matplotlib APIs and bindings for both the data files on LCFs and for the data within the Integrated Data Management system.
- 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. 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 could be subsumed in FSP needing a dynamic model capability. The need is there, we just need an agreement as to what we call this capability.
  - There was the suggestion 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.

Figure 2. Facets Composer for input file preparation and analysis, job submission and monitoring, and visualization.

- Monitoring: As a workflow is executing, 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 the Facets Composer for this work.
  - 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.
  - Future development within 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.
- Verification and Validation: While specifically not a part of the charter for this workshop, 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.
  - 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.

3.4 Development Processes and Infrastructure

One of the main products of FSP is software. Development processes and supporting infrastructure are critical for producing fusion simulation software that is robust, reliable, usable, maintainable, and able to exploit a range of hardware platforms. Based on community input on the current state of development processes and the needs and requirements for FSP, we have identified the following development process areas and begun outlining possible solutions whose goal is to make the adoption and practice of these processes easier and more efficient. The overall current level of adoption of standard software development processes is sporadic and approaches vary among projects within the fusion community. In the initial year of FSP, many of the current software developers will be adapting to the new process requirements. To make this transition smoother, FSP will place a high priority on the creation of examples of best practices, templates, reference materials, and supporting software infrastructure

3.4.1 Requirements for Development Processes and Infrastructure:

- FSP should support software for collaborative development.
  - Code design meetings should vet all major proposed revisions.
  - Code review mechanisms are needed to ensure compliance with best practices.
  - Revision control: (also known as source code management) is the management of changes to source, code, documents, and other information stored in files. In addition to tracking changes, revision control systems commonly provide automated merging of nonconflicting changes made by development teams. Revision control has the following requirements:
    - Allow multiple repositories to be combined hierarchically.
    - Support sets of federated (linked) repositories (as an alternative to having them centrally located).
    - Anonymous read-only access for non-export control codes (optional) and non-OTP access in general; controlled access to export control codes.
    - Available on MacOS, Linux, and the LCFs
    - Ability to generate development activity statistics
    - Able to take actions upon software commits, such as running tests or emailing software changes.
  - Symbol collision avoidance methods, such as identifier prepending or namespacing will be required.
  - Software access for developers: FSP should provide a central directory of software that includes references to a canonical location (URL) of each top-level package (component or set of components). 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
  - Software standards are needed to facilitate coding and code reading by all.
  - API documentation is needed for all code. Consistency with the software should be ensured by employing a documentation generation system such as Doxygen
  - Communication mechanisms are needed for distributed teams. These include mailing lists and web collaboration tools.
  - Cross platform build systems are required for varying developers to work on different machines.
  - Package management systems should facilitate building of all needed third party libraries.
- User documentation is needed. Such documentation should document all inputs and outputs as well as the use of the software on various platforms. Users should be able to access expert help through a mailing list.
- Performance measurement will be required for all components and for the framework. Automated performance measuring tools can be used effectively within the build system.
- Regular software releases should be available in both binary and source form
  - Binary releases should be available for a standard set of platforms.
  - Source code releases should be available with a cross-platform build system.
  - FSP-wide versioning conventions are needed. E.g., all packages follow X.Y.Z where changes in interfaces are reflected by incrementing Y, and bug fixes are reflected in incrementing Z, while major redesign or reimplementation is reflected in X increments.
  - Backwards compatibility should be required. When new software is proposed for replacing functionality, it should provide a backward-compatible interface (in addition to any new interfaces), e.g., providing ADIOS functionality under an HDF5 interface, much like ATLAS is a drop-in replacement for BLAS and LAPACK.
  - User documentation should accompany the releases.
  - A minimum frequency of releases, e.g., annual, should be established.
- Quality assurance mechanisms should be in place

    - Static code verification will be needed for identifying potential problems.
    - Unit tests should be used for testing subsystems.
    - Regression tests should be required for each new feature.
    - A testing system will be needed for

- Invocation of all tests
- Reporting of test results
- Determining the changes that led to test failure.
    - Bug tracking will need to be supported. The FSP should try to settle on a single system.
    - Prioritization mechanisms should be established for dealing with bug reports and feature requests.

3.5 Production Computing

Production computing refers to the use of computational software for research applications by end users. In most cases the users are research scientists, but not the same people who wrote the software. From such users, there is an expectation that software used within its documented range of validity will perform robustly and produce accurate results. Where this fails, the user expects access to expert help.

Production computing is usually run on carefully built versions of software, which includes careful selection of versions of components in the case of multi-physics software such as an FSP whole device model (WDM). There may be hundreds or thousands of runs executed with each production version of the software.

Production computing is of great importance to the FSP as it is its point of contact with the end users who are the ultimate FSP customer base. It is in production use that FSP has an impact on the wider research community and the advancement of the science.

In this section we itemize the requirements for robust production computing that have come out of the FSP March Planning Meeting.

3.5.1 Requirements for Production Computing:

- Access to high-performance, capability platforms
- Dedicated, stable, high capacity platforms for software execution:
    - Typically with support for remote service invocation or job submission.
    - An open, liberal policy for access to production services by interested researchers in the user community.
    - Optimized for rapid turnaround of user tasks.

- Staff to support the conversion of research computational applications to robust production computational applications. Such staff will likely not be research scientists and will have to have a different reward system.
- End user oriented documentation (tutorials, FAQs, reference guides, etc.).
- Hands on user training, especially for new or heavily revised software.
- Supporting users for preparation of input, submission of runs, and visualization and/or analysis of output (software discussed above in Task Composition) for both capability and capacity computing.
- Support hundreds to thousands of runs on each production version.
- Support users not expert in every physics component of the production simulation—especially the case for multi-physics FSP WDM applications.
- Support typical research applications:
    - Model validation against experimental data.
    - Analysis of experimental data.
    - Predictive simulation for design or planning of experiments.

    - Sensitivity studies to supplement experimental data analysis, or, quantification of model uncertainty (UQ).

- Support monitoring of production job status and progress.
- Support for run management, such as user decision to halt, restart, or delete a queued or running job.
- Archival of run results in a persistent, publicly accessible data storage service (the Integrated Data Management system), including suitable provenance records, user comments, etc.
- Staff to support users for production software operation, including:
    - Build and installation of production software on specific supported platforms
    - Trouble-shooting and analysis of simulation failures; updates to production version software for bug fixes; rapid response where feasible.
    - Triage: ability to identify and refer out more difficult problems to component experts—with implication that each component of production grade multi-physics simulations must also be supported independently.
    - User assistance in matters of understanding of inputs and outputs.

3.6 Change management and process

As the Fusion Simulation Program proceeds, the area of Software Integration and Support will have to change and adapt to the research needs of the Fusion Energy Sciences Program, personnel changes, and changes in available hardware and software. Examples of reasons for change include

Science driver priorities are reflected in the efforts in development of different component and/or composition software. As discoveries are made, the science driver priorities will change, and this will lead to modifications of the development priorities.

Personnel will definitely change over the course of a 15-year program. How does one design a program that can institutionalize the knowledge? How are new personnel (developers, managers, analysts) brought into the program? What is the process of rebudgeting when priority changes lead to a need for more of the expertise at a particular institution?

Computer architectures will change over the course of the program. We saw the revolution of the rise of distributed memory computing over the last 15 years. We are now seeing a revolution in the use of many-core processors, such as in GPU computing. In the next 15 years one can easily imagine another revolutionary change. At what point does the FSP port existing software to r commission new software for fundamentally different platforms?

Third-party software, such as solver libraries and data analysis tools, will be critical to FSP. How can FSP interact with developers of performance-critical third-party libraries to ensure that they provide optimized libraries on key platforms?

Adapting to change will require decisions at multiple levels. FSP should define the different levels of change management processes according to the scope and potential impact of the change. At the very least, one could state that:

- Limited-impact decisions should be made locally by developers.
- Significant changes in software scope or significant redesign should have FSP-wide review.

Further, as the FSP moves forward, it is expected that multiple, competing tools will be proposed for adoption, whether these be tools for software composition, data analysis, development, user support or other. Selection of the best tools for FSP will require a process that includes a needs analysis and an analysis of resource requirements. The needs analysis will require understanding the user needs, but it should also allow for leadership through introduction of tools for more productive ways of working not yet envisioned by the users. The resource analysis must lead to an understanding of the time to introduction and the staffing required for production of first implementation as well as the staffing required for maintenance and user support.

Management must therefore have mechanisms in place that provide for input from all of the stakeholders, leading scientists in the field (that could come, e.g., from national and international committees and/or facilities), the users, those responsible for supporting the users, and those responsible for implementing and maintaining integration software, and those responsible for supporting the computational and experimental facilities. One approach is to have a Change Review Board, with representation from the full set of stakeholders. The Change Review Board would study the proposed changes and new solutions to determine costs and benefits and ultimately make recommendations to management. Such studies could include asking for assessments of resource requirements from the various developer groups. Requests for proposals to supply the envisioned software could be a mechanism for cost-effectively obtaining the software.

4 Advanced Physics Components

4.1 FSP Component strategy in the planning stage

The FSP component strategy in the planning stage is motivated by the FSP scope and deliverables in the execution phase. The expected FSP scope and deliverables in the execution phase will be first explained. This will be followed by the component strategy in the planning stage.

During the projected 15 years FSP execution phase, the FSP scope and deliverables are guided by science drivers (SD). The FSP will develop a common framework and software infrastructure on which all physics software addressing individual or multiple science drivers will be built. The most important deliverables of FSP are computational software release which will be used to resolve specific science drivers. These would involve integration of components using physics and computational coupling schemes tailored for individual science driver. Such SD-specific integrated software will undergo verification and uncertainty quantification by the developers before being released to the production services, which have the responsibility of enabling the usage of the released software by non-developers. By accessing the same data management services as the experiments, validation and uncertainty quantification will be carried out on these SD-specific software releases. The outcome will be fed back to the further development of the physics components and integration schemes. This iterative process with clearly defined deliverables in terms of SD-specific software release integrates the development, verification, production services, and validation, which are essential for the continuing improvement of the physics fidelity of the FSP software.

The component strategy in the FSP planning stage aims at developing a component implementation plan which will be carried out in the execution phase to fulfill FSP deliverables. The strategy has the following ingredients. First, the required computational physics capability will be articulated for each of the science drivers. This capability requirement will be factorized into physics components. Taking into account the additional requirements from the framework consideration, one arrives at the FSP component specification for the primary components. By comparing the requirement with existing capabilities in the SciDAC and base programs, we will identify the key gaps. It is anticipated that there will be two kinds of gaps which require separate solutions. The first type is modest and can be adequately addressed by improving the candidate components drawn from the existing capabilities. The task for the planning stage would be to assess the required improvements, along with the necessary FSP resources. The second type refers to those gaps which are significantly beyond the current mature capabilities. The planning team would need to consult with the existing exploratory research and identify promising approaches and directions. These will be articulated in terms of research opportunities which potentially motivate new FSP component initiatives. The final component execution/implementation plan will thus have detailed plans to address both types of gaps. Carrying out the component implementation plan in the

FSP execution phase will produce the SD-specific software to be subjected to a vigorous V&V. Together with experiments, they will lead to discovery science enabling fusion energy.

4.2 Carrying out the component strategy in the planning stage:

To effectively carry out the component strategy in the FSP planning stage, we have a work plan which emphasizes community engagement. The steps in our work plan are as follows. The first step is to develop the component and coupling (CC) requirements for each science driver. By coupling, we include both integrated physics models and the computational schemes which integrate different physics components over space and time. Based on the CC requirements, the component team will call for SD-specific component/coupling (CC) design from leading computational teams/groups in the fusion community. In parallel, the component team develops a set of component questionnaire and solicits community input on component candidates by topical fusion science area. These input, along with the community input on CC design by science drivers, will be critically evaluated by FSP component working groups which consist of both FSP component team member and leading experts in the field. The objective is to produce a set of prioritized CC design in response to every science driver, in tandem with the parallel effort to assess the overarching framework/integration requirements and to produce FSP software infrastructure design.

The overall work plan described above leads to two sets of specific tasks, which are carried out by engaging the fusion community. The first set of tasks focuses on individual science drivers. For each science driver, we will identity the integrated physics model which is required to resolve the science driver, to the best knowledge of leading experts in the field. A component factorization will be performed with both the components and the coupling scheme identified. The component requirements will then be spelled out, in all three areas, namely physics, computational, and software. These readily lead to the assessment of readiness and gaps. Similarly the requirements on coupling/integration scheme, which integrates the components to resolve a science driver, are articulated. Two areas of obvious importance are data exchange and time synchronization between components. Again, a clear resolution of the requirements will facilitate the readiness and gaps analysis.

The second set of tasks cuts across the six science drivers under consideration and potentially additional ones in future discussions. The specific goals are to identify those components which are (1) common to all science drivers; (2) shared by multiple but not all science drivers; and (3) unique to individual science drivers. For those FSP common components and the ones shared by multiple science drivers, there is the opportunity to port them into the FSP software infrastructure. The planning team will devise the development plan for that. For those primary components unique to individual science driver and those shared by multiple science drivers, but cannot be readily provided by upgrading existing capabilities, the planning team will identify the FSP component common challenges. These will form the basis of a near-, mid-, and long-term R & D strategy for FSP components.

4.3 Accomplished tasks at the FSP planning workshop

Following the presentations summarizing the challenges of the six provisional science drivers, eleven design proposals were presented to address the science drivers in the FSP planning workshop. The charge to the proposal teams is to lay out how the physics challenges are broken down into computable components which are coupled to resolve the integrated physics models. These input are provided by the three proto-FSP centers, five fusion SciDAC centers, and three community code projects funded by the base program. These design proposals were discussed in the context of individual science drivers in the follow-up breakout group discussion by science drivers. The component breakout groups took on the task next day. There are three component working groups tackling two science drivers each. The primary objectives of the working groups are (1) to establish component specification for individual science driver in the context of its likely coupling scheme, and (2) across six science drivers, to identify common components and common or science-driver-unique challenges.

The three component working groups completed the component factorization and the underlying coupling scheme for all six science drivers. These were done for both near term and long term perspectives. The appendices I-III give a summary of the discussion of the working groups and their findings for each of the six science drivers. Our expectation is to continue community engagement to refine and update these findings and the FSP designs.

4.4 Summary findings

Based on the component working groups' report, we can briefly highlight some summary findings of the FSP planning workshop. There are two sets, corresponding to the near term and long term perspectives. The near term perspective addresses the issue of FSP readiness and emphasizes what can be delivered by FSP in a two to five year time frame. The long term perspective targets the direction of R & D which in 10 years, would result in a level of sophistication and fidelity otherwise impossible.

Near term perspective: On the issue of readiness for FSP, the working groups find that there is a solid base of existing (component) capabilities and credible integration schemes to produce meaningful integrated software to tackle every science driver within the first five years of the FSP execution phase. 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. There is a consensus that the early delivery of a FSP-ware provides an excellent platform for verification and (in)validation.

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

Long term perspective: in a 10 year time frame, the working groups envision a FSP capability far beyond what is currently feasible. The substantial investment afforded by FSP over a 10 year development cycle, enables the acceleration of predictive fusion science at a level unprecedented in the U.S. Program. It is also recognized that this acceleration of progress is necessary for U.S. to optimally utilize ITER to explore DEMO physics. In terms of specific component development, the working groups' findings suggest that all science drivers converge to common component R & D needs in key areas. Two examples are the 2D equilibrium and transport solver from B axis to wall and the self-consistent 3D plasma equilibrium and transport solver from B axis to wall. On the coupling front, imbedded calculation, which is called out for essentially all science drivers, poses similar challenges to coupling/framework.

The working groups' findings clearly established that the R & D thrusts in the physics integration area are on converging paths. For example, the core transport capability moves to include the edge plasmas; the edge transport capability moves to include the core plasma transport. Similarly, the extended MHD modeling moves to include the effect of gyrokinetic turbulent transport, while the gyrokinetic transport modelers plans to include the impact of low-n to medium-n magnetic activities.

4.5 Conclusions

The essential conclusions of the component working groups are: (1) the fusion community is ready for FSP with ready deliverables in the first two to five years that represent significant improvements over the current state of the art; (2) the long term prospect of FSP is inspiring.

On the second point, it is specifically recognized that the challenges to achieve predictive simulation for FSP are daunting. Because of the converging paths, success of any one of them make a successful FSP. It is also important to note that we have gotten some time to work on them, which is almost a necessity considering the extreme challenges. In terms of the FSP product, the long term vision is that different science drivers are increasingly being tackled by fewer but highly integrated (physics-wise) components. The long term prospect of FSP is intellectually appealing and cleaner in terms of numbers of components/coupling schemes involved. The other five science drivers, if successfully resolved, converge nicely into the Whole Device Modeling science driver.

5 Integrated Data Management

5.1 Introduction

This document summarizes presentations and discussions carried out during the FSP workshop held in Boulder in March 2010. As much as possible it is meant to represent a consensus of the participants in the "Integrated Data Management" breakout group listed in section G. It is by no means a design document or a design specification, but does attempt to capture FSP requirements, to outline general principles to be followed and to suggest approaches that might best meet FSP needs.

Motivation: 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.

Scope: This integrated data management topic includes all data stored or used as part of the fusion simulation program excepting data 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.

5.2 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. Beyond this are a set of "non-functional" requirements which 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 agreed on by workshop participants are:

- A unified view of all FSP data.
- Remote access to data without requirements for remote logins or file copies.
- Ease of use for data providers and consumers aimed at supporting a large community of non-specialists.
- 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.

The non-functional requirements, along with the approach described below, are driven by 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 range of customers include 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 "n2" problem where a large number of application and groups of data items must be customized to each other. It also emphasizes the importance of designing software systems that can evolve and adapt with new technologies. 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 which must be supported.

5.3 Approach

5.3.1 Main elements of FSP integrated data management system

We envision two main elements which 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 which 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.

Data Catalog

The data catalog provides a consistent, coherent and global view for all FSP data. 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 C.2 below.

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.

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 which represent the data in a more compact form. This could be resampled or processed data or graphics of any kind.

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.

Structured Data

The details and mechanisms for data storage are, to a large extent, of little importance to users. More critical is that the system build on a strong data abstraction – that is, that the user view of the data be simple and well-structured – independent of the underlying implementation. Unnecessary complexity should be hidden from users as far as possible. Well designed structures can promote ease of use, support self-description and extend the useful life of data. There are many data models to do this, but generally speaking, hierarchical and relational data models would appear to be appropriate. Well structured data supports reusable application whose behavior is driven by data and data structures, for example by presenting menus, labeling plots or providing help information found on the fly. Specific instances could include applications which pre-process experimental data for code inputs, post-process or visualize code outputs. Workflow tools that prepare code inputs and controls could obviously benefit from this approach, offering users choices defined by available data and validating inputs against a set of rules (which are also stored data.)

Data Access

To simplify data access , the general approach suggested here would be 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 fully qualified name (to be defined) obviating the need of users to know where or how the data is actually stored. 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.

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 it's 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.

5.3.2 FSP Data centers

The physical realization of the data systems for the FSP also deserve 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.

5.4 Relation between data management, workflows and production computing

The consensus at the workshop was that the vision for integrated data management described here was entirely consistent with the needs of computational workflow (task composition) and for production computing. Workflow which 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. For the purpose considered here, 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 it's 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.

5.5 Summary of major challenges

At the workshop a set of important challenges were identified. These will need to be addressed at some level of detail in the next six to twelve months as part of the project definition.

1. 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.
2. 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.
3. 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.
4. Universal Access Layer: This critical piece of software must be highly capable and easy to use.
5. 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, distributed computing and other strategies for improving overall performance will need to be investigated.
6. 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.
7. Governance and decision processes: The management processes that oversee the development and maintenance of the data management system need to be defined.
8. Change management: Some structured approach must be in place to plan, document and control the evolution of the system over time.

The overall system outlined here is moderately complex with many moving parts. A good doctrine for design would be to begin with modest implementations and build in capabilities as needed over time.

5.6 Next steps

Once these general principles and approaches are agreed to, we will need to begin moving toward concrete designs to implement them. The steps required include:

- Finalize this workshop report and identify any outstanding issues
- Factor overall problem into a set of discrete, well-defined elements
- Gather candidate solutions and develop conceptual designs
- Assess through discussion and conceptual design reviews
- Prepare draft of Integrated Data Management write-up for FSP final report
- For design elements which require them, develop prototype solutions, assess these and report.

- Finalize design and document
- FSP design review

Appendices

Appendix A.

Participants in the breakout session and contributors to the Science Driver write-ups.

Boundary Plasma
T. Rognlien (LLNL) – discussion leader
L. Berry (ORNL)
J. Brooks (Purdue)
R. Cohen (LLNL)
C-S Chang (NYU)
M. Dorr (LLNL)
D. Knoll (LANL)
Brian LaBombard (MIT)
S. Krasheninnikov (UCSD)
M. Greenwald (MIT)
J. Shadid (SNL)
D. Stotler (PPPL)
Dennis Whyte (MIT)

Pedestal
Phil Snyder (GA) – discussion leader
C.S. Chang (NYU)
J. Cummings
Jeffrey Hittinger (LLNL)
S. Klasky
S. Parker
L. Sugiyama
X. Q. Xu (LLNL)

Core Profiles
Bill Nevins (LLNL) – discussion leader
Pat Diamond (UCSD)
Bill Dorland (U. Md.)
Jeff Candy (GA)
Greg Hammett (PPPL)
Chris Holland (UCSD)
Scott Parker (U. Co.)

Wave particles
Raffi Nazikian (PPPL) – discussion leader
Herb Berk (U.Tx)
Paul Bonoli (MIT)
Guoyong Fu (PPPL)
Cynthia Phillips (PPPL)
Don Spong (ORNL)

Disruptions
S. Kruger Tech-X – discussion leader
J. Breslau (PPPL)
V. Chan (GA)
R. Harvey (CompX)
A. Hassanein (Purdue)
E. Hollmann (UCSD)
D. Humphreys (GA)
V. Izzo (UCSD)
S. Jardin (PPPL)
Y. Petrov (CompX)
A. Pigarov (UCSD)
S. Putvinski ITER
A. Reiman (PPPL)
V. Sizyuk (Purdue)
T. Sizyuk (Purdue)
D. Stotler (PPPL)
H. Strauss (NYU)
D. Whyte (MIT)

Whole Device Modeling
Chuck Kessel (PPPL) – discussion leader
Don Batchelor (ORNL)
Glenn Batman (Lehigh)
Jeff Candy (GA)
Julian Cummings (Cal Tech)
Tom Casper (ITER)
Ammar Hakim (TechX)
Jon Kinsey (GA)
Arnold Kritz (Lehigh)
Doug McCune (PPPL)
Allan Reiman (PPPL)
Sveta Shasharina (TechX)

Participants in Core Transport and Whole Device Modeling Science Drivers

J. Candy
T. Casper
J. Cummings
A. Hakim
C. Kessel
D. McCune
W. Nevins
A. Reiman

Participants in Integrated Data Management Breakout

Tom Casper (ITER)
Martin Greenwald (MIT)
Chris Holland (UCSD)
Doug Hudson (ORNL)
Scott Klasky (ORNL)
Doug McCune (PPPL)
Bill Nevins (LLNL)
Svetlana Shasharina (TechX)
Arie Shoshani (LBNL)
Linda Sugiyama (MIT)
Brian Van Straalen (LBNL)

Participants in Disruption and Wave-Particle Science Drivers

H.L. Berk
V.S. Chan
G.Y. Fu
S.E. Kruger
R. Nazikian
A. Reiman
D. Spong
L. Sugiyama

Appendix B.1: Science Drivers Original roadmap

1. Develop/extend 3D fluid turbulence/2D transport
2. Include basic models for plasma/wall interactions
3. Include impurity and radiation transport
4. Include kinetic plasma transport (4D) and turbulence (5D)
5. Include 3D magnetic equilibrium, wall geometry, initial RF sheaths
6. Include first-principles plasma-wall models
7. Include detailed RF antenna/wave effects
8. Include improved radiation transport and atomic physics

Appendix B.2 Development of First-Principles Plasma-Wall Interaction Models

(S. Krasheninnikov and B. Wirth)

OPPORTUNITY AND CURRENT STATUS

It is acknowledged that plasma-material interactions pose an immense scientific challenge and are one of the most critical issues in magnetic confinement fusion research [1]. The demands on plasma-facing materials in a steady-state fusion device include extreme particle and thermal fluxes. These energetic fluxes have pronounced impacts on the topology and chemistry of the near-surface region of the material, which influence the plasma sheath potentials and subsequent threat spectra. These evolutions are also inherently multiscale in time and are likely controlled by diffusional phenomena that are influenced by the high heat loads and subsequent thermal (and stress) gradients into the material, as well as by defect micro/nanostructures induced by both the ion and neutron particle irradiation. The complexity of plasma materials surface interactions are schematically illustrated in Figure 1. This complexity is further underscored by the fact that the plasma and materials surface are strongly coupled to each other, mediated by an electrostatic and magnetic sheath, despite the vastly different physical scales for surface (~ nm) versus plasma (~ mm) processes.

For example, the high probability (> 90%) of prompt local ionization and re-deposition for sputtered material atoms means that surface material in contact with the plasma is itself a plasma-deposited surface, not the original ordered material. Likewise, the recycling of hydrogenic plasma fuel is self-regulated through recycling processes involving the near-surface fuel transport in the material and the ionization sink action of the plasma. The intense radiation environment (ions, neutrons, photons) ensures that the material properties are modified and dynamically coupled to the plasma materials surface interaction processes. Some of the most critical plasma materials interaction issues include: i) the net erosion of plasma-facing surfaces; ii) net tritium fuel retention in surfaces; iii) H isotope and material mixing in the wall; and iv) the minimization of core plasma impurities. Furthermore, the plasma-material surface boundary plays a central role in determining the fusion performance of the core plasma. However, while it is widely accepted that the plasma-surface interface sets a critical boundary condition for the fusion plasma, predictive capabilities for PSI remain highly inadequate.

Gaining understanding and predictive capabilities in this critical area will require addressing simultaneously complex and diverse physics occurring over a wide range of lengths (angstroms to meters) and times (femtoseconds to days and beyond to operating lifetimes), as graphically represented in Figure 1. The lower time and length scales correspond to individual ion implantation and sputtering, which occurs at or near the material surface, in addition to a range of ionization and recombination processes of the sputtered neutrals and ions in the near surface sheath. At intermediate length and time scales, a wealth of physical processes are initiated, including diffusion of the now implanted ionic/neutral species, the possibility of chemical sputtering processes at the surface, the formation of gas bubbles, surface diffusion driving surface topology changes and phonon scattering by radiation defects that reduces the thermal conductivity of the material. At longer length and time scales, additional phenomena such as long-range material transport in the plasma, re-deposition of initially sputtered surface atoms, amorphous film growth and hydrogenic species diffusion into the bulk material and permeation become important. This broad palette of physical phenomena will require development not only of detailed physics models and computational strategies at each of these scales, but algorithms and methods to strongly couple them in a way that can be robustly validated. While present research is confined to each of these scales, or pioneering ways to couple two or more of them, the current approaches already push the state-of-the-art in technique and available computational power. Therefore, simulations spanning multiple scales needed for ITER, DEMO, etc., will require extreme-scale computing platforms and integrated physics and computer science advances.

BASIC SCIENCE CHALLENGES AND RESEARCH NEEDS

Multiple scientific challenges exist to develop the capability to predict the response of plasma facing materials to the extreme thermal and particle fluxes expected in the magnetic fusion energy environment. These challenges involve deciphering the complex and diverse physics that occur over inherently multiscale length and timescales and include i) hydrogenic fuel retention and control, ii) surface morphology and film evolution, and iii) sheath physics and power transmission. Of course, there are obvious synergistic interactions even amongst these four identified challenges.

Hydrogenic fuel retention

Practical considerations, based on tritium retention, fuel inventory and erosion management, lead to a design preference for the use of tungsten as the divertor material, and possibly also the entirety of the plasma facing components, in future fusion reactors. In fact, the 'all metal' ITER design option involves replacing the current graphite design with tungsten strike points that would operate at temperatures above 1000 K. In such environments, tungsten will experience high heat loads and exposure to various hydrogen and helium isotopes and impurity species, possibly including beryllium, carbon and argon (which is added to enhance radiative plasma losses). To date, only limited experience exists with solid tungsten in magnetic fusion confinement devices, while recent laboratory based plasma exposure experiments involving hydrogen and helium implantation on tungsten have demonstrated a very rich, and as yet poorly understood surface response behavior.

Baldwin and co-workers have used the PISCES facility at UCSD to investigate the response of W following exposure to mixed deuterium – beryllium, deuterium – helium or deuterium – helium – beryllium plasmas and observe dramatically different surface response depending on the surface temperature and implantation conditions [2,3]. At temperatures below about 600°C, limited blistering was observed and in fact, the presence of some beryllium in the plasma was observed to be beneficial in providing suppression of blister formation. However, for temperatures above about 700°C, the formation of low melting temperature W-Be intermetallic alloy second phase regions was observed, which is believed to accelerate surface degradation. Further and perhaps more importantly, exposure of tungsten surfaces to helium plasmas at temperatures above about 700°C leads to the formation of a low density, tungsten 'fuzz', which consists of "amorphous W nano-rod like structures" [Baldwin 08], which are shown in the SEM photograph in Figure 2. The thickness of the 'fuzz' layer exhibits an Arrhenius temperature dependence with an activation energy around 0.7 eV, and is thereby consistent with a diffusional mechanism. Interestingly, the nanoscaled fuzz is not observed to form under lower temperature helium exposure, even at significantly larger implantation fluences. Combined these observations indicate that tungsten surface response to mixed constituent plasmas involves complex erosional and re-deposition processes coupled with defect and diffusional mechanisms.

This example begins to demonstrate the complexity of tungsten surface response, and the wide range of anticipated operating parameter space (operating temperature, plasma constituents and impurities, etc) for tungsten based plasma facing components in magnetic fusion reactors. Furthermore, and perhaps more importantly, this example raises a number of scientific questions regarding the synergisms between hydrogen and helium implantation in materials and the corresponding interactions between hydrogen and helium with atomic defects and diffusive mechanisms leading to surface topology changes in plasma facing components and mechanical property changes in bulk materials that impact not only irradiated materials but also fuel cells, hydrogen storage media and the plasma facing thruster materials in plasma propulsion devices. Specifically, the key questions are:

1. What are the controlling kinetic processes (e.g., defect and impurity concentrations, surface diffusion, etc.) responsible for the formation of a nanoscale 'fuzz' on tungsten surfaces subject to high temperature He plasmas?
2. What exposure conditions (e.g., a phase boundary map of temperature, dose, dose rate, impurities) lead to nanoscale 'fuzz' or other detrimental surface evolution?
3. How much tungsten mass loss occurs into the plasma as a result of nanoscale 'fuzz' formation? And, finally, how can this surface evolution be mitigated?
4. What are the controlling He – defect and hydrogen/deuterium/tritium interaction mechanisms that influence hydrogen permeation and retention?
5. What plasma impurities increase sputtering yields of tungsten? What mitigation measures are possible to reduce tungsten mass loss?

Surface and film dynamics

Key questions relate to whether the essential underlying physical and chemical kinetic processes and rates (e.g., diffusion, molecular bonding, surface release, amorphization, etc.) be accurately understood and modeled as a function of the controlling exposure conditions? Furthermore, how does the evolving nano-to-microstructure at or near the surface impact average material properties of erosion, material mixing, plasma fuel retention, thermal properties and secondary electron emission? These multiple questions are essential to developing fundamental knowledge to enable high fidelity multiscale computer modeling of the interface layer between the plasma and bulk materials. Indeed, the near surface film evolution is strongly dictated by the plasma conditions, yet the film provides feedback to the plasma itself through particle removal (sputtering) and fueling processes.

Plasma sheath surface coupling

During a typical plasma discharge in current tokamaks, the hydrogen species inventory in the wall strongly changes during the discharge, primarily, due to continuous implantation of plasma particles, which is a very fast process. Moreover, transport of hydrogenic species within the wall includes slower processes, for example, bulk diffusion and co-deposition. The interaction of these multi-timescale processes within the wall material with the edge plasma is what determines the fuelling boundary condition for a confined plasma since the overall wall fuel inventory is much larger (100-1000x) than the plasma particle inventory. Along with plasma parameter variation during a discharge, the fast and slow processes in the wall make the simulated plasma transport inherently time-dependent, spatially two-dimensional, and strongly non-linear. Correspondingly multi-scale time-dependent non-linear implicit solvers are under development. While current modeling capability, for example in the WALLPSI code, do couple the plasma and surface to determine sheath effects, there is a significant need to incorporate more complex phenomena such as multi-component materials, plasma impurity species and helium ash, along with surface compositional and morphological evolution.

References:

1. G. Federici, C.H. Skinner, J.N. Brooks, J.P. Coad, C. Grisolia, A.A. Haasz, A. Hassanein, V. Philipps, C.S. Pitcher, J. Roth, W.R. Wampler, D.G. Whyte, "Plasma-material interactions in current tokamaks and their implications for next step fusion reactors", Nuclear Fusion, 41 (2001) 1967.
2. M.J. Baldwin et al "The effect of high-fluence mix-species (D, He, Be) plasma interactions with W," 18th Int. Conf. on Plasma Surface Interactions, May 2008 & accepted for publication in Journal of Nuclear Materials (2009).
3. Y. Ueda, et al. "The effects of tungsten surface conditions on carbon deposition," 18th Int. Conf. on Plasma Surface Interactions, May 2008 & accepted for publication in Journal of Nuclear Materials (2009).

Figures:

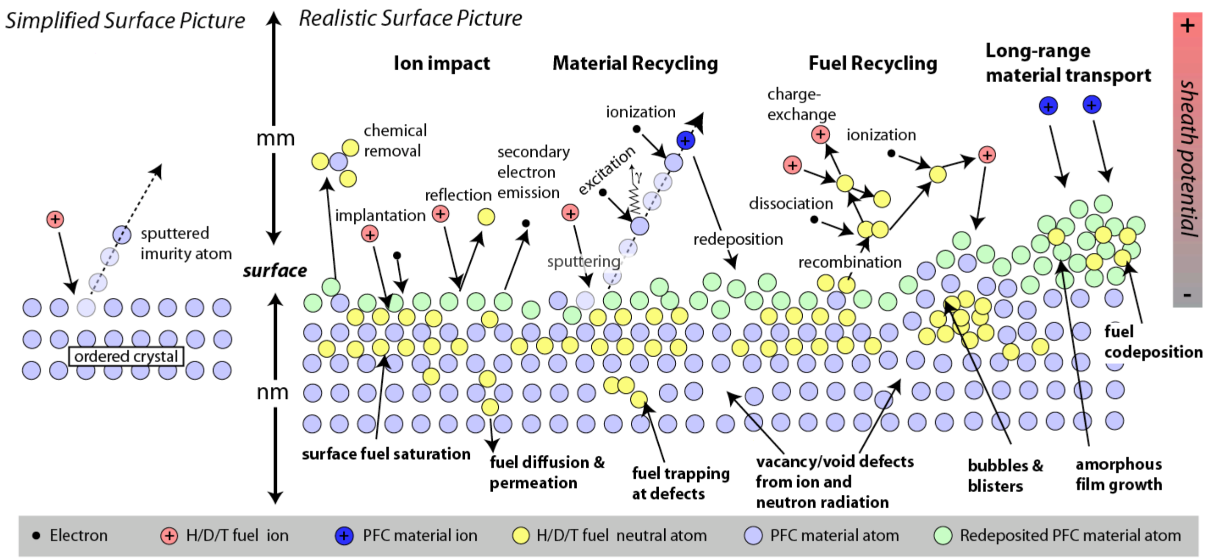

Figure 1. Schematic illustration of the complex, synergistic and inherently multi-scale surface interactions occurring at the material surface in a realistic magnetic fusion plasma environment, in contrast to the simplified picture of ion-induced sputtering from an ordered material.

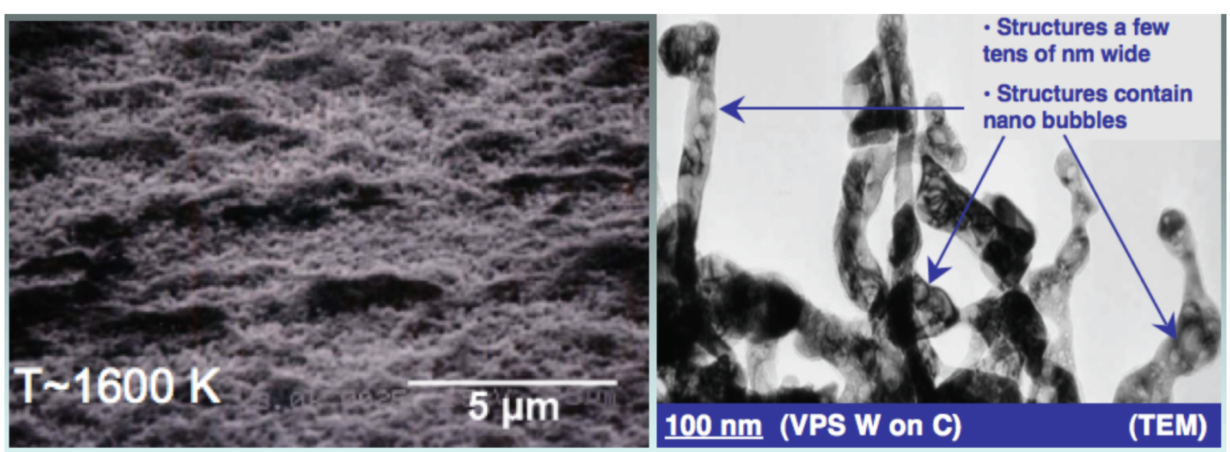

Figure 2. Example of W response to He plasma bombardment at ~1200°C, as reproduced from Ref. [2,3]. (left) Sub-micron surface structures reduce W reflectivity. (right) TEM picture of complex microstructures.

Appendix C: Core Transport and Whole Device Modeling Science Drivers

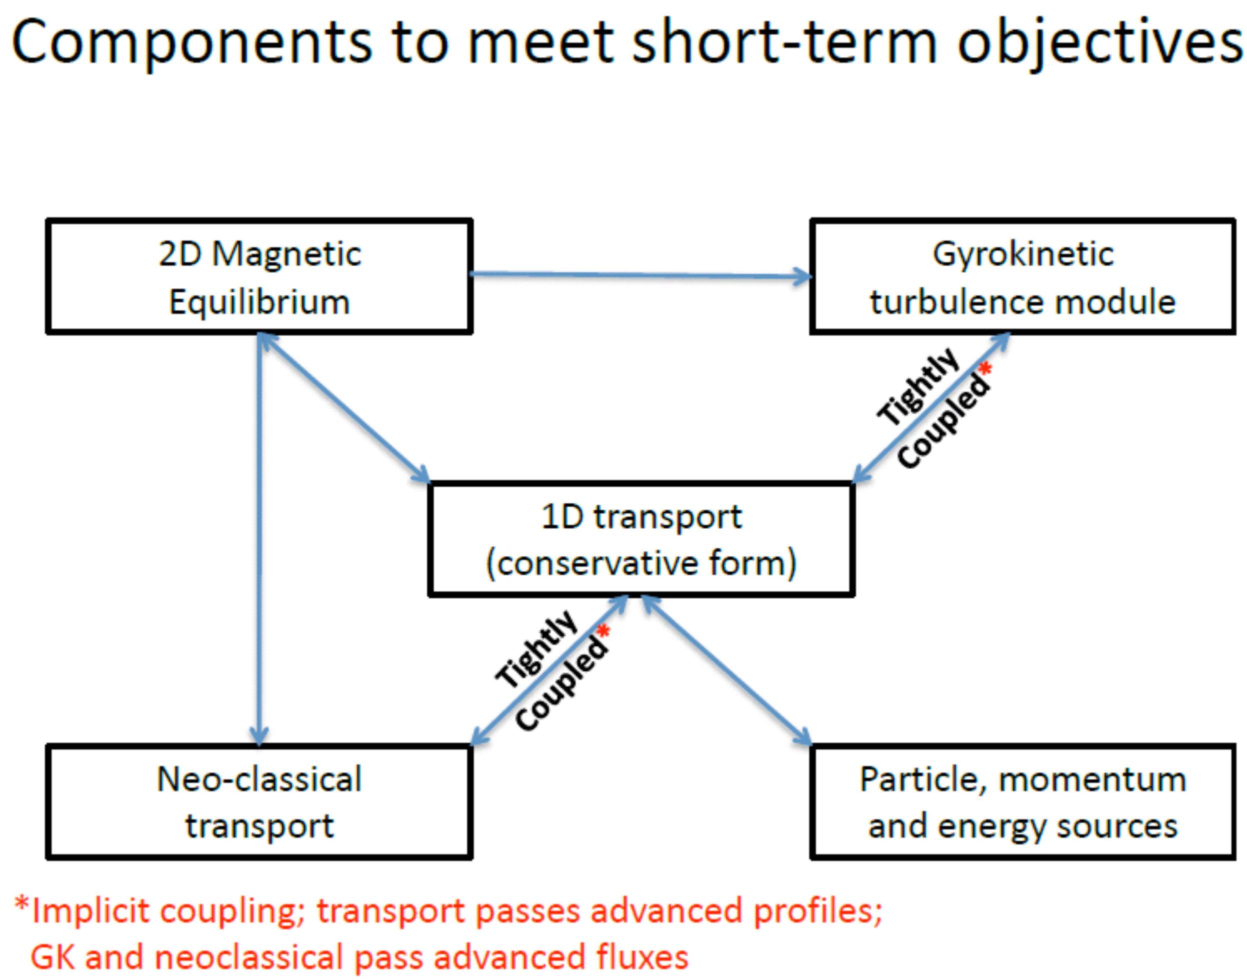

Appendix C.1: Core Transport

The short-term objective (2-5 yr) is develop a modeling capability that includes

1. turbulent transport (TGLF, GYRO, GEM, GS2, IFS-PPPL)
2. neoclassical transport (NCLASS, analytic, NEO)
3. sources (NUBEAM, [AORSA], theory, measured)
4. equilibrium (EFIT, many others)

There should be, roughly, both "first-principles" and "reduced" descriptions of all physics. Using these modules in accordance with the standard delta-f transport theory, we can evolve core temperature, density and electric field.

In particular, this capability will include

- *Er* evolution according to strong-rotation formulation
- validated core simulation of DIII-D L-mode
- demonstrate flexible execution with swapping in and out of components
- post-hoc linear MHD stability test
- effects of energetic particle-turbulence coupling

A fundamental new functionality needed, but which is probably not realizable on the 5-year timescale, is a generalized kinetic model of low-n macroscopic phenomena (some theory which brings together GK and MHD into a unified whole, still ignoring fast time-scales which are ignorable). This functionality is needed to address: NTM physics, moderate-n Alfven modes, RWM, RMP and sawteeth in a more rigorous way. Also, an improved model of Er (diamagnetic-level rotation) evolution is required to describe intrinsic rotation.

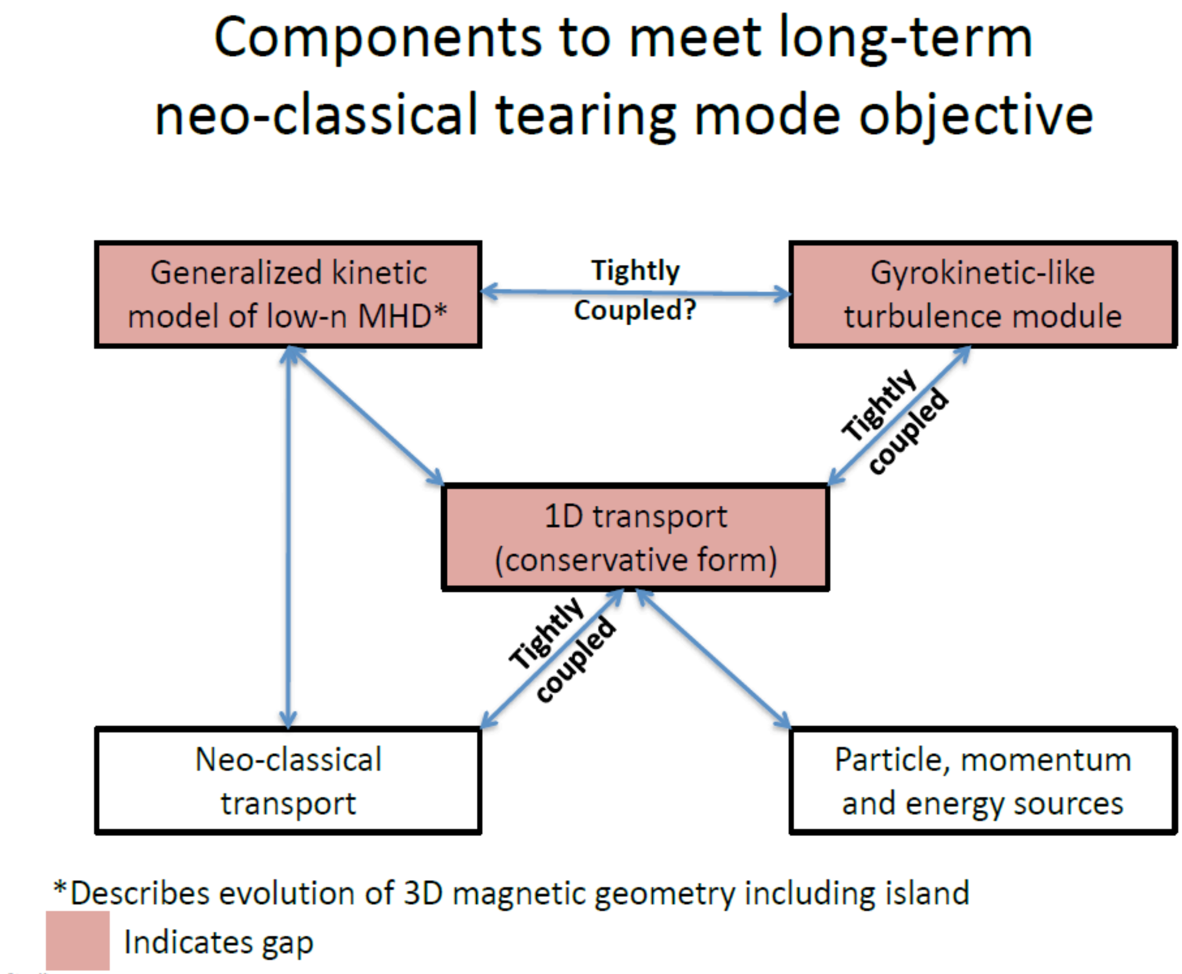

The first core transport diagram is a simplified diagrammatic description of the coupling described in 1-4 above. The second core transport diagram is a suggestion, of indeterminate credibility, regarding how the picture might change in the long-term to address the gap associated with unification of MHD and gyrokinetics. For example, it is not at all clear whether MHD, neoclassical and turbulence are separable processes in general. More precisely, neoclassical and GK are probably not separable if equilibrium and fluctuation scales are not separable. Moreover, if MHD and GK are unified, then there will be only one box for MHD+GK. So, it is possible that GK+MHD+neoclassical may be one tightly coupled module.

Appendix C.2: Whole Device Modeling

The WDM tool itself will have varying levels of fidelity based on the user's needs. These needs are largely determined by turn-around time, with faster turn-around requiring the use of lower fidelity models, and vice versa.

A list of typical (desirable) physics models (or components) in a WDM are:

1. 2D equilibrium solver
2. 1D transport solver for (temperature, density, momentum, and current)

3. ICRF/LH/NB/EC/alpha heating, current drive, and momentum sources
4. Pellet/NB/gas injection/wall particle sources, and divertor pumping/wall sinks
5. Transport models (empirical, GLF23, MMM08, TGLF, GYRO, …)
6. Bootstrap current model (NCLASS, Sauter, Neo)
7. Pedestal model (empirical, ballooning, EPED, ELITE)
8. Radiation (bremsstrahlung, line, cyclotron)
9. Fusion reactions
10. MHD models for NTM's, RWM's, sawteeth
11. Fast particle models for *AE's and other energetic particle modes
12. Poloidal field coils, conducting structures, and feedback systems
13. Neutrals/atomic physics
14. Scrape-off-layer plasma (SOL/divertor)
15. Plasma wall interaction physics

As an example, any one of these physics model areas could break out into more detail, say for the heating, current drive, and momentum sources (where this level of detail would appear in the appropriate science driver area):

1. ICRF, GENRAY (ray-tracing), TORIC (full wave), AORSA (full wave all orders), a Fokker Planck solver such as CQL3D, ORBIT-RF (Monte Carlo finite orbit effects)
2. LH, GENRAY/adjoint (ray-tracing, 1D Fokker Planck), GENRAY/CQL3D (ray-tracing, 2D Fokker Planck), TORICLH/CQL3D (full wave, 2D Fokker Planck)
3. NB, NUBEAM (Monte Carlo orbit following)
4. EC, TORAY(ray-tracing, 1D FP), GENRAY/CQL3D (ray-tracing, 2D FP with momentum conserving corrections)

Several possible science driven deliverables for the 2-5 year time frame can be identified for the WDM, based on existing physics modules and credible development time for the WDM framework (this is not exhaustive), and these include:

1. Development of the WDM tool composed of 2D equilibrium solver, 1D transport solver, an array of source models, transport models, and additional physics modules.
2. Provide experimental interpretive capability equivalent of TRANSP or ONETWO within the framework.
3. Fokker Planck treatment of multiple fast particle species (ICRF minority, NB, and fusion alphas) in an ITER discharge simulation.
4. Core transport model fidelity examination, with GLF23/MMM08, TGLF, and GYRO (for example).
5. Development of a reduced model for the redistribution of fast NB particles by Alfven eigenmodes.
6. Initial demonstration of connecting the core plasma to SOL plasma model with ELM evolution.
7. Comparison of reduced sawtooth trigger model (Porcelli) with NOVA-K stability calculations including multiple fast particle species in ITER.

A diagram of the deliverable #2 is shown in Fig. 1, showing the particular emphasis of the task on modeling multiple fast ion species coupled to the basic WDM framework. The tight coupling of the transport models to the transport solver is shown, along with the 2D equilibrium and 1D transport solver main driver. The basic information required of the various components in the WDM framework is represented by the plasma state. The tight inter-coupling of the source models that is required for the multiple fast species treatment is also shown. Only a few of the additional physics models are shown. It is recognized that within the WDM the new fast particle modeling for the heating and current drive sources facilitates additional model improvement for the Porcelli sawtooth model, which can be extended to include fast species other than alpha particles, and allows the direct comparison through the NOVA-K stability code (a higher fidelity model) which requires the distribution function information that the source modeling now provides.

Possible science driven deliverables for the 5-10 year time frame can be identified for the WDM, HOWEVER, it is expected that these efforts will need to begin early in the FSP program because of the required development time:

1. Implement tight coupling between the core and edge plasma, including the pedestal, SOL, and plasma wall interactions, and may require a 2D transport solver structure in the core to facilitate coupling to the SOL plasma, in order to model the ELM pulses in ITER.
2. Show the nonlinear evolution of fast particle distribution under the influence of Alfven eigenmode instabilities in ITER with high safety factor and reverse magnetic shear
3. Implement 3D fixed and/or free boundary equilibrium solver that can be used to directly include magnetic islands and stochastic regions, with particular focus on the MHD interactions like neo-classical tearing modes, sawteeth, resistive wall modes and edge localized modes.
4. Implement nonlinear extended MHD models for plasma instabilities and disruptions.

Shown in Fig. 2 is the diagram of the deliverable #1 showing the particular emphasis of the task to model the pedestal region, the SOL plasma, the neutrals and plasma material interactions (PMI) at the divertor surface (and first wall). The pedestal model connects directly to the transport solver since it sets the boundary condition for the core transport, and also connects to the SOL since it provides the source for particles and energy. There is coupling of the SOL plasma, neutrals and PMI in the divertor since this provides the consistency condition between the plasma parameters at the divertor surface and pump opening and those at the upstream mid-plane. This modeling provides an accurate model for the source of impurities, and would encourage improved modeling of impurity transport in the core plasma. The 1D transport solver may not provide the best interface conditions for the 2D SOL plasma and so a 2D core transport solver would be considered.

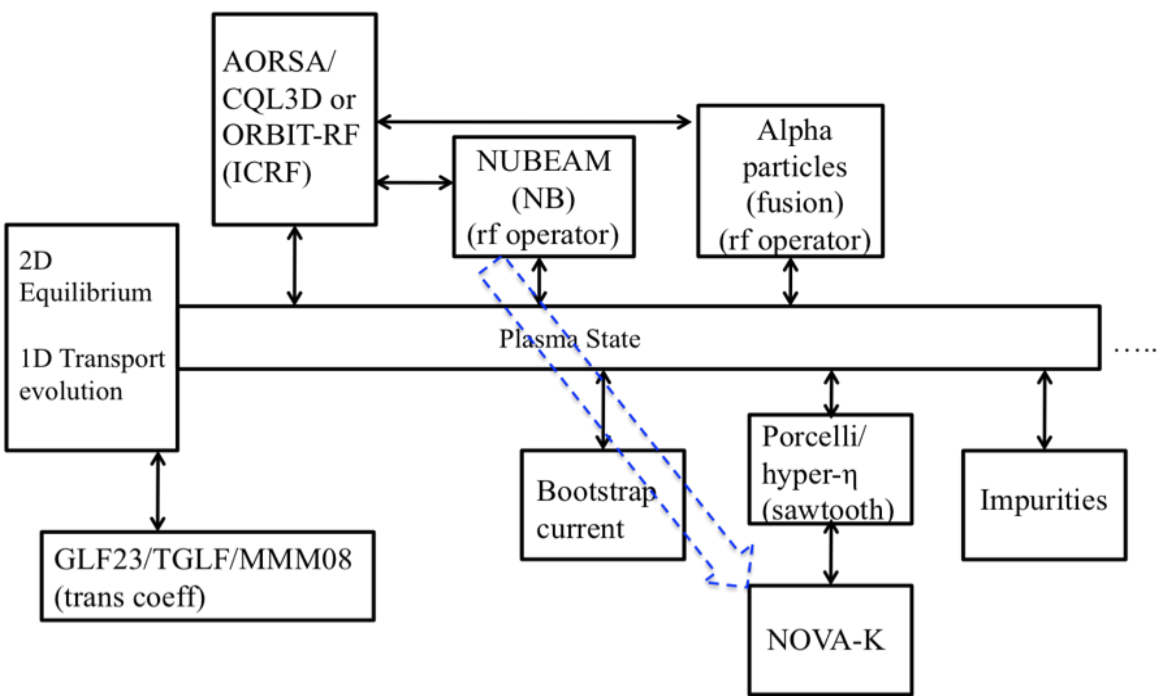

Figure 3. Short term deliverable flow diagram for the whole device modeling for the Fokker Planck treatment of all fast particle species (ICRF minority, NB, and fusion alphas) in an ITER discharge simulation.

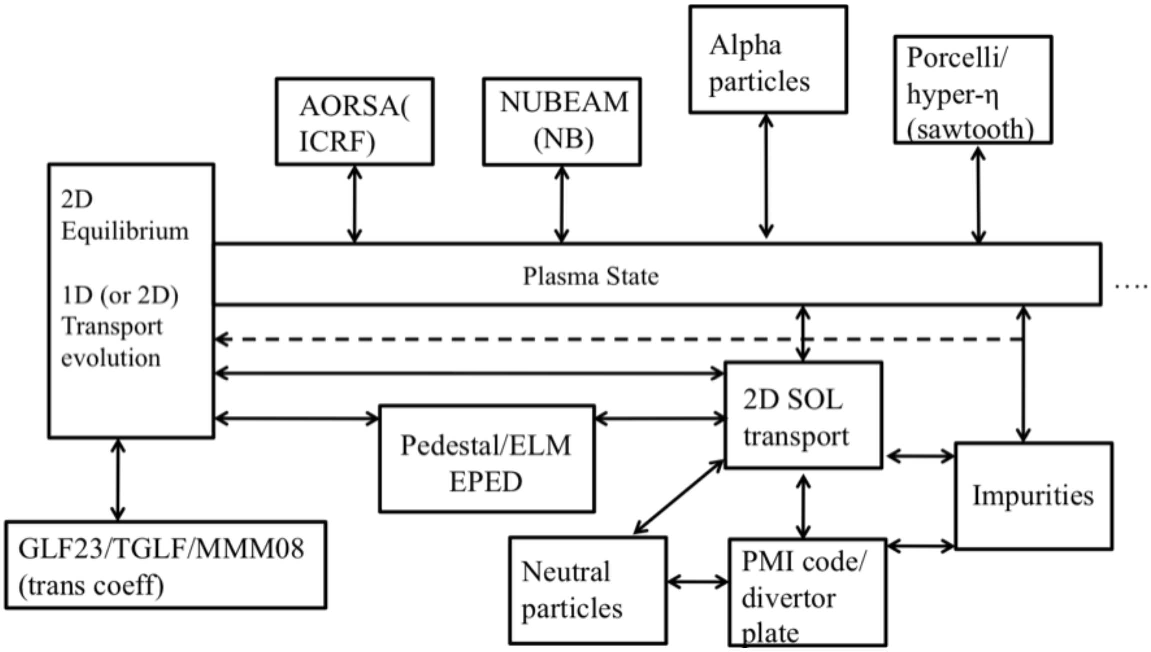

Figure 4. Long term deliverable flow diagram for the whole device modeling for the implementation of tight coupling between the core and plasma edge, including the pedestal, the SOL, and the plasma wall interactions, in order to model the ELM pulses in ITER.

Appendix C.3: Disruption and Wave-Particle Science Drivers

Key science challenges for the Wave-Particle Science Driver

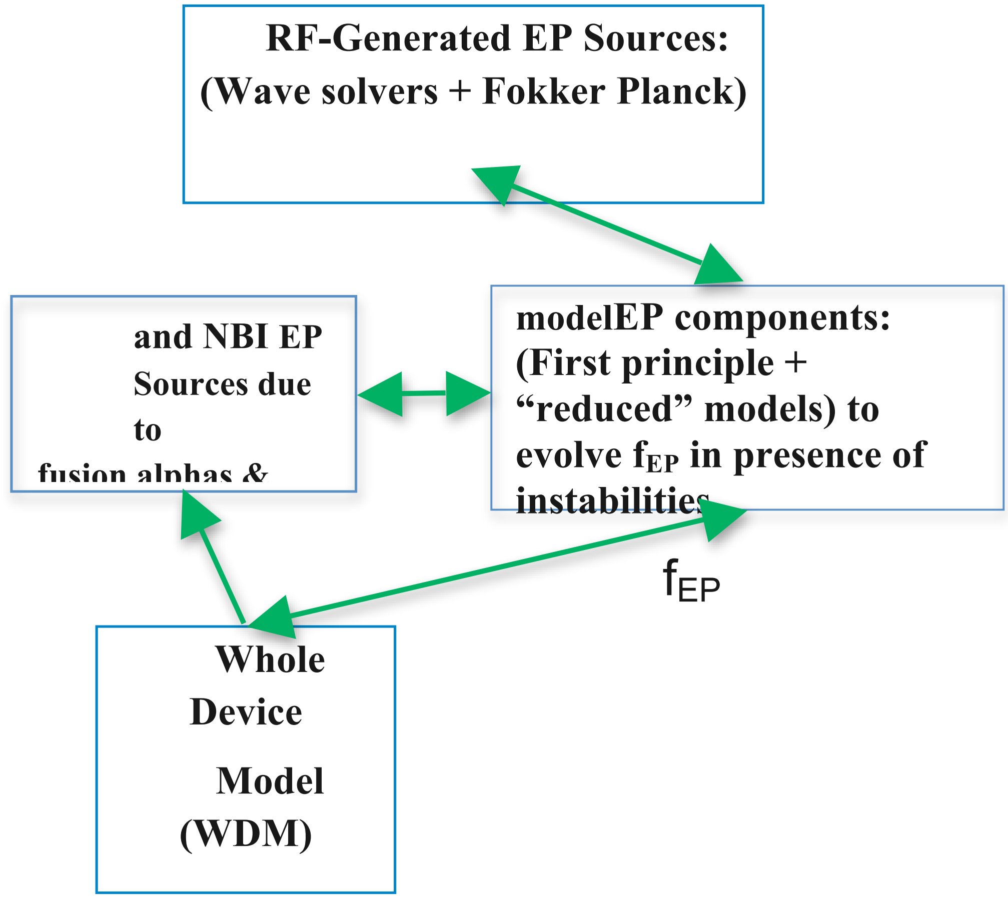

Figure 5. (WP-1) Coupling scheme for components of the wave-particle science driver assuming a 2-5 year time window.

The Wave-Particle Science Driver has the goal of studying whether collective fast particle phenomena can significantly affect discharge evolution, fusion performance and facility integrity in alpha dominated burning plasmas, such as the ITER device. This Science Driver has the further goal of investigating if radio-frequency (RF) waves can couple into the core of a fusion plasma and be used to effectively control the plasma in the presence of multiple energetic particle species such as due to fusion alpha particles or neutral beam injection (NBI) and edge dissipation mechanisms associated with non-linear formation of RF sheaths or parametric decay instability (PDI).

Component factorization for the Wave-Particle Science Driver (2-5 year version)

Component Functionality:

Energetic particles (EP) have two sources in the coupling scheme shown in Fig. WP-1. The RF-generated energetic particle sources consist of wave propagation models (either full-wave or ray tracing) that iterate with a Fokker Planck code to self-consistently evolve non-thermal electron and ion distributions. These coupled RF sources span the ion cyclotron (IC), lower hybrid (LH), and electron cyclotron (EC) range of frequencies. The other sources of energetic particles are due to fusion products such as alpha particles and fast ions due to neutral beam injection (NBI). These sources are passed to an EP component that can consist of either first principle or reduced models that evolve the EP particle distribution (fEP) (in both space and time) in the presence of Alfvenic type instabilities such as Toroidal Alfven Eigenmodes (TAE's). It is expected that reduced physics models will be used in the EP component over the shorter term (~ 2 years) and that first principle simulations will be employed at the mid-term (~ 5 years). The Whole Device Model (WDM) performs a transport evolution of the background plasma profiles for density, temperature, as well as solving an evolution equation for the poloidal magnetic field.

Component requirements:

RF Components:

The full-wave RF components solve Maxwell's equations under the assumption of rapidly oscillating, time harmonic wave fields, which reduces the system to a generalized Helmholtz wave equation,

$$-\nabla\times\nabla\times\mathbf{E}+\frac{\omega^2}{c^2}\left(\mathbf{E}+\frac{i}{\omega\varepsilon_0}\mathbf{J}_p\right)=-i\omega\mu_0\mathbf{J}_{ant} \qquad \text{(WP-1)}$$

where J_{ant} is an externally driven antenna current, localized near the plasma edge, that acts as a source for the waves. The fluctuating plasma current J_p can be derived directly from the rapidly varying part of the distribution function f_s^1(r,v). In general, J_p is a non-local, integral operator on the wave electric field and is expressed as:

$$\mathbf{J}_p(\mathbf{r},t)=\sum_s\int d\mathbf{r}'\int_{-\infty}^{t}dt'\sigma\left(f_s^0(E),\mathbf{r},\mathbf{r}',t,t'\right)\cdot\mathbf{E}(\mathbf{r}',t') \qquad \text{(WP-2)}$$

where s(f_s^0, r, r¢, t, t ¢) is the "plasma conductivity kernel." The long time scale response of the plasma distribution function is obtained from the bounce averaged Fokker-Planck equation which can be written in form:

$$\begin{aligned}
&\frac{\partial}{\partial t}(\lambda f_0)=\nabla_{\mathbf{u_0}}\cdot\boldsymbol{\Gamma}_{\mathbf{u_0}}+\langle\langle S\rangle\rangle,\\
&\lambda=u_{//0}\tau_B, \qquad \text{(WP-3)}\\
&\nabla_{\mathbf{u}}\cdot\boldsymbol{\Gamma}_{\mathbf{u}}=C(f_0)+Q(\mathbf{E},f_0),
\end{aligned}$$

where ⟨⟨S⟩⟩ is a source of particles, τB is the bounce time, $C(f_0)$ is the momentum conserving Balescu-Lenard collision operator, and Q(E, f0) is the RF quasilinear operator due to Kennel and Engelmann. In the shorter term (2-5 years) it is required that the wave propagation component describe linear edge-to-core coupling of ICRF and LHRF waves in realistic 3-D launcher & vessel geometry. Core wave absorption physics in the ICRF and LHRF regimes will also have to accurately describe the RF wave – fast ion interaction with finite ion orbit width effects properly included. Here fast ions would include fusion alpha particles, ions from NBI, and fast ions generated by the ICRF power itself. It is also required that the ECRF wave induced flux will be incorporated self-consistently in the MHD closure hierarchy, although the electron distribution function in this case can be assumed to be close to thermal (Maxwellian).

Energetic Particle Components:

During the near term (2-5 years) it will be expected that the EP component will be capable of describing the evolution of EP distributions due to Alfvénic/acoustic instabilities and other macroscopic MHD modes and collisions for a whole device simulation, using lower dimensional models. Also during this period, first principle simulations of the nonlinear evolution of Alfvénic/acoustic instabilities and EP distributions in experimentally relevant conditions on the mode saturation time scale will be required.

Candidate codes for the components

1. Energetic particle components:
   - Nonlinear extended MHD and gyrokinetic: NIMROD, M3D-K, GKM, GTC, GYRO, GEM, TAEFL
   - Linear extended MHD: AE3D-K
   - Linear kinetic MHD: NOVA-K
   - Monte Carlo PIC: DELTA5D
2. Wave-particle components:
   - Wave propagation: AORSA, TORIC, GENRAY
   - Fokker Planck: CQL3D, ORBIT RF, sMC
   - Antenna - Coupling: TOPICA, RANT3D
3. Whole Device Model: TSC, PTRANSP, TRANSP, CORSICA

Coupling scheme for the Wave-Particle Components:

The coupling scheme for components of the wave-particle science driver in Fig. WP-1 is relatively straightforward. A Whole Device Model (WDM) advances the background plasma profiles and MHD equilibrium that are used to calculate the sources of energetic particles due to fusion products, neutral beam injection, and RF waves. These initial energetic particle distribution functions are nonlinearly evolved in the presence of Alfvenic instabilities using both first principle and reduced models. The evolved EP particle distribution is fed back to the RF component where the RF wave – particle interaction is then re-evaluated. Similarly, fEP is fed back to the WDM component where its affect on plasma profiles and stability is included in the time advance of the transport equations.

Gap Issues:

Several near term gap issues that need to be addressed for the EP component include the need to simulate instability near marginal stability, the need to develop synthetic diagnostics, especially those for assessing dynamic losses of energetic particles to the wall, and the need to include realistic particle distribution functions in the EP components. The primary gap issues that must be addressed for the RF Component are a way to combine the core wave solution with a 3-D description of the RF launcher and vessel and the development of a combined core wave solver plus continuum Fokker Planck with finite orbit effects accounted for (improvement of bounce average resonant description from zero orbit width to finite orbit width). Also in the nearer term, RF and energetic particle distributions must be coupled with Alfven instabilities and an implicit numerical scheme needs to be developed for inclusion of ECRF wave induced flux in the MHD closure hierarchy.

Component factorization for the Wave-Particle Science Driver (10 year period):

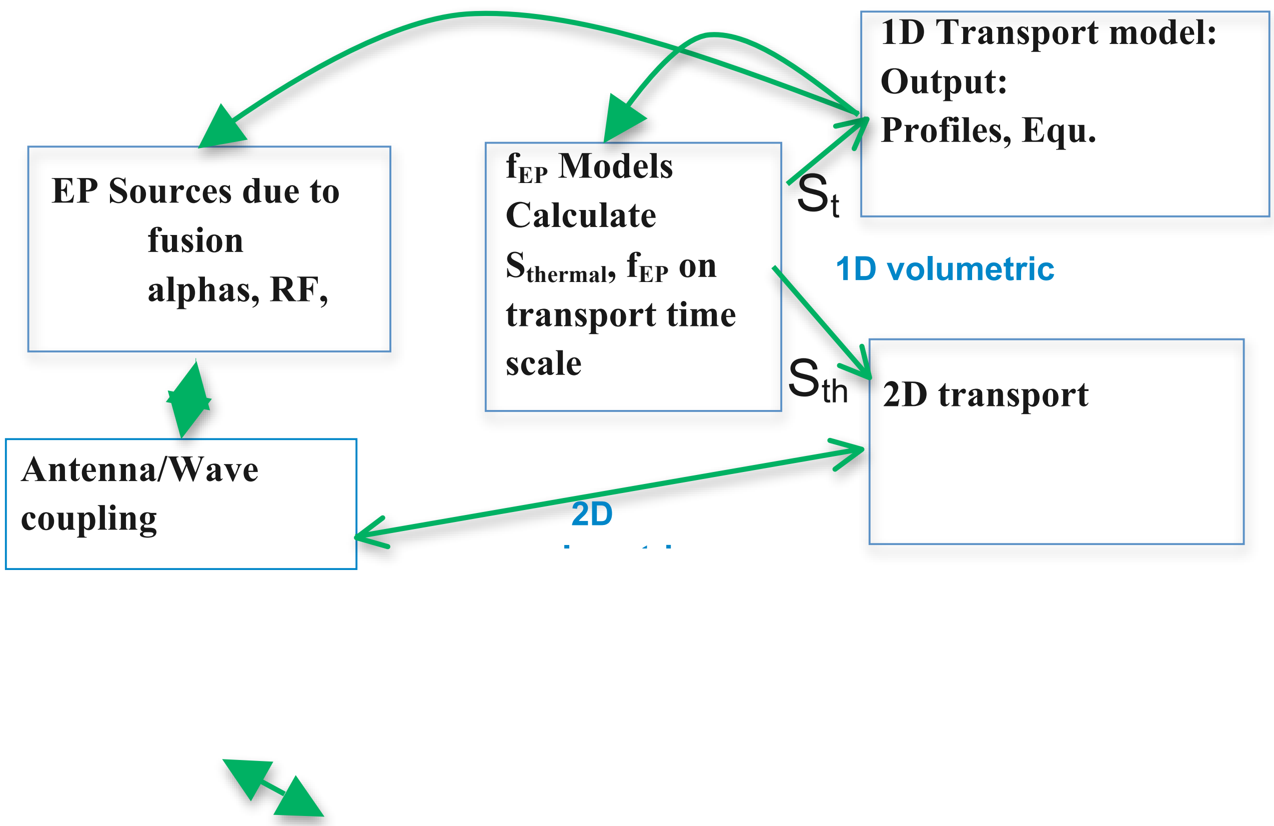

Component functionality and requirements:

For the coupling scheme shown in Fig. WP-2, the functionality of EP sources due to fusion alphas, RF, and NBI are much the same as in the section above on Wave-Particle component functionality. However the antenna / wave coupling component must now be capable of evaluating nonlinear effects such as RF sheath formation and parametric decay instability (PDI) of the RF pump wave. Also it will be necessary to include ICRF generated energetic particle tails in the MHD closure hierarchy so that modification of sawteeth via ICRF can be studied. The EP component will be required to simulate the nonlinear evolution of EP driven Alfvenic/acoustic instabilities with macroscopic MHD in the presence of RF on transport time scales. Also, the EP component will be required to predict fast ion transport and mode saturation levels and effects on the macroscopic MHD in burning plasmas. The component for edge transport in Fig. WP-2 will need to include sources of heat due to nonlinear RF dissipation mechanisms or EP losses and evolve the edge plasma accordingly.

Candidate codes for the components

1. Energetic particle components:
   - Nonlinear extended MHD and gyrokinetic: NIMROD, M3D-K, GKM, GTC, GYRO, GEM, TAEFL
   - Linear extended MHD: AE3D-K
   - Monte Carlo PIC: DELTA5D

2. Wave-particle components:
   - Wave propagation: AORSA, TORIC, GENRAY
   - Fokker Planck: CQL3D, ORBIT RF, sMC
   - Non-linear RF: VORPAL
   - Antenna - Coupling: TOPICA, RANT3D

3. Whole Device Model: TSC, PTRANSP, TRANSP, CORSICA

Coupling scheme for the Wave-Particle Components:

The coupling scheme for components of the wave-particle science driver in Fig. WP-2 is as follows. Again, the Whole Device Model (WDM) advances the background plasma profiles and MHD equilibrium that are used to calculate the sources of energetic particles due to fusion products, neutral beam injection, and RF waves. These initial energetic particle distribution functions are again nonlinearly evolved in the presence of Alfvenic instabilities using both first principle and reduced models, but this time on the transport time scale. The evolved EP particle distribution is fed back to the RF component where the RF wave – particle interaction is then re-evaluated. Similarly, fEP is fed back to the WDM component where its affect on plasma profiles and stability is included in the time advance of the transport equations. An additional level of complex coupling is now present however as the effect of EP losses and edge heat dissipation due to nonlinear sheaths or PDI are included in an edge transport model. As the SOL and boundary are evolved by the edge transport code, these changes must be fed back to the RF wave coupling calculation so that can be recomputed.

Gap Issues:

The main gap in the EP component for achieving the 10 year goals will be extending the EP codes to 3D equilibria. The primary gap issues that must be addressed for the RF Component will be how to include nonlinear RF sheath and PDI effects in the edge RF description. Also a scheme must be formulated for inclusion of ICRF generated ion tails in the MHD closure hierarchy so that sawtooth stabilization via ICRF generated tails can be simulated.

Common Components:

Although we did not discuss the Disruption Impact Science Driver in this write up, we did identify in our working group and FSP presentation a common set of components between the Disruption Impact and Wave-Particle Science Drivers which were the following:

1. Transport, Sources, equilibrium:

   TRANSP, PTRANSP, TSC, CORSICA, ONETWO

   PIES, EFIT, JSOLVER, VMEC, QSOLVER, TEQ, TOQ

2. (ii) Extended kinetic MHD:

   NIMROD, M3D

3. (iii) Global gyrokinetic codes:

   GTC, GEM, GYRO

4. (iv) Fokker Planck codes:

   CQL3D

<u>Grand challenges in component of possible relevance to other SD's:</u>

One of the grand challenges of the Wave-Particle Science Driver will be the development of an edge RF coupling component that includes nonlinear effects due to RF sheath formation and PDI. This Grand Challenge will be important to the Science Driver on Integrated Boundary Layer (SOL), Divertor, and Plasma Wall Interactions. The effects of RF power dissipation due to sheath formation and PDI in the edge would be expected to impact the boundary layer and plasma wall. Likewise, as plasma parameters in the SOL evolve, this would be expected to influence conditions for the onset of these nonlinear RF effects.

Disruption science driver:

The scientific goal of the disruption science driver is to understand the dynamics of mitigated and unmitigated disruption in order to understand how to limit their effects. The scientific impact is to enable the robust operation of tokamaks by allowing more aggressive operating regimes and by enabling faster recovery from off-normal events.

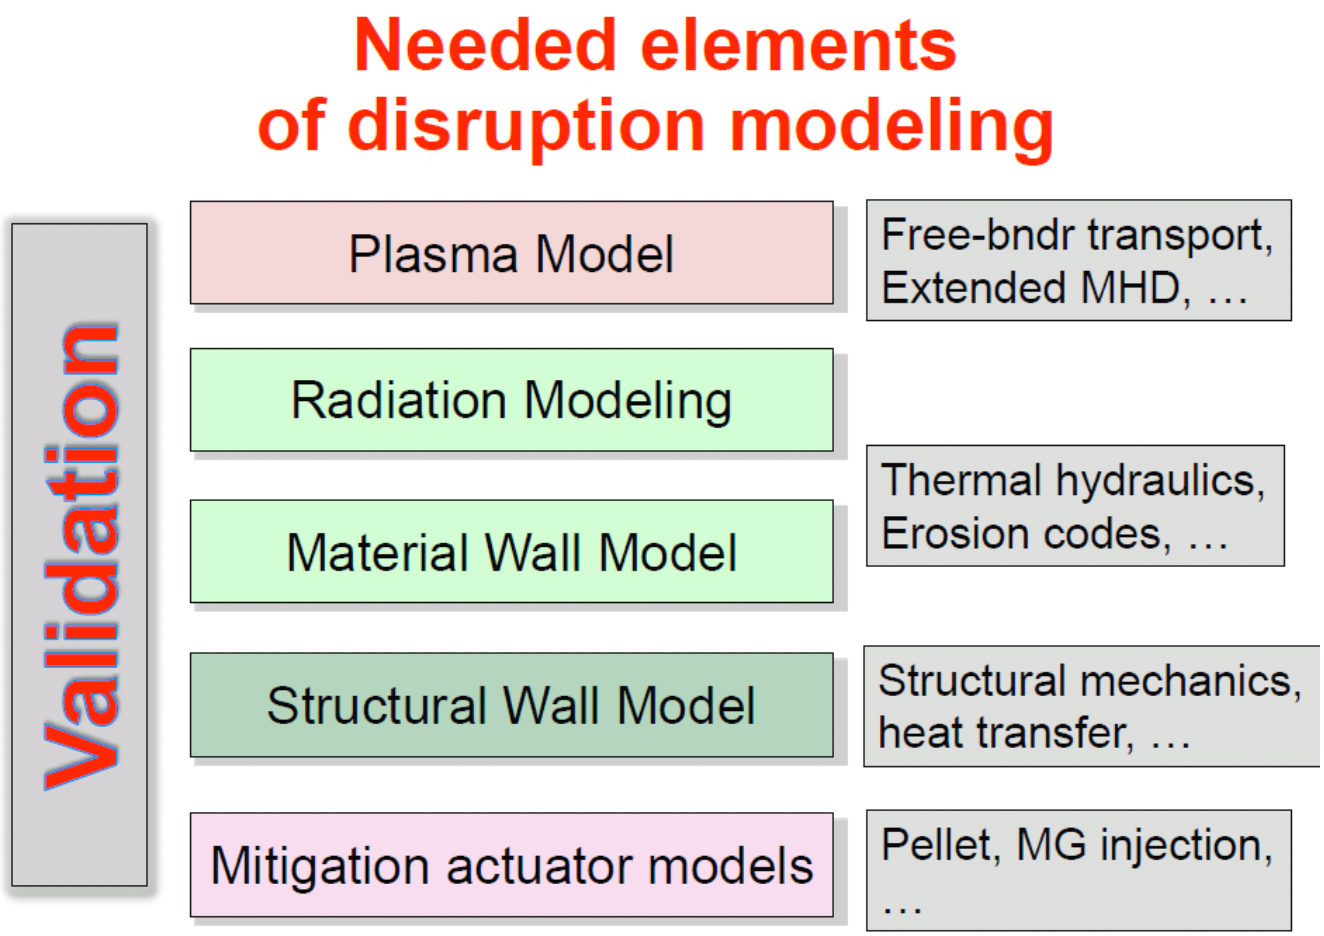

Figure 7. Disruption-1: Required elements to model tokamak disruption.

Disruption modeling requires a number of elements as indicated in the diagram. In the near term, these elements can be integrated as follows. The primary engine is an extended MHD component, which has imbedded runaway electron model and impurity/radiation model. It takes MHD equilibrium and volumetric 1D profiles from a whole device model, while provides the plasma and runaways profiles to a Fokker-Planck runaway electron propagation component. Similarly the profiles and the electromagnetic fields will be communicated to the components which model the structural mechanics of the wall and the materials response to plasma irradiation.

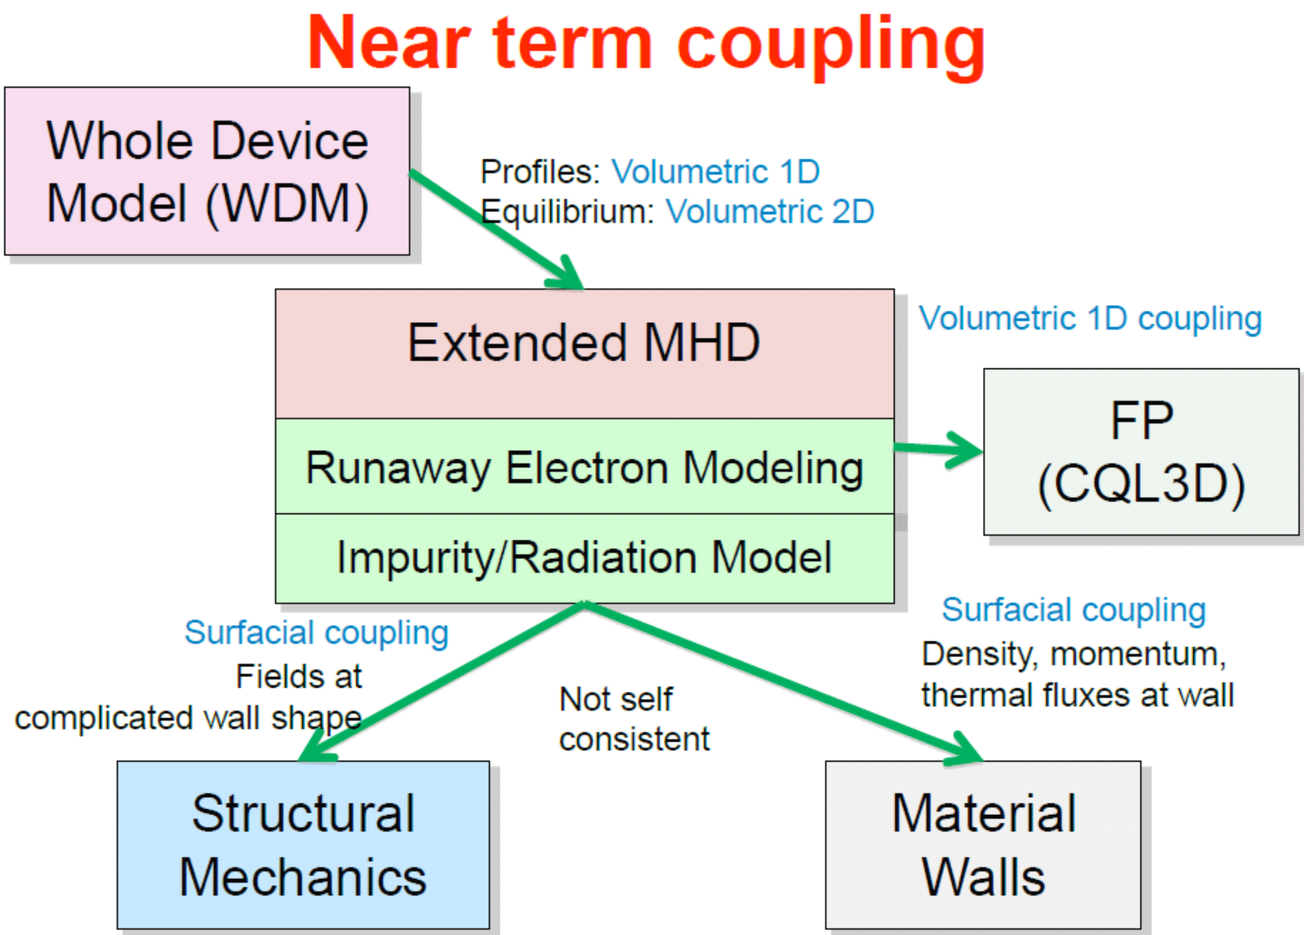

Figure 8. Near-term coupling scheme

In the mid-term, there is the opportunity to carrying long time-scale modeling of disruption by coupling 3D equilibrium calculation to the averaged transport equations in 3D magnetic field, and hence implicitly advance the plasma profile and magnetic field.

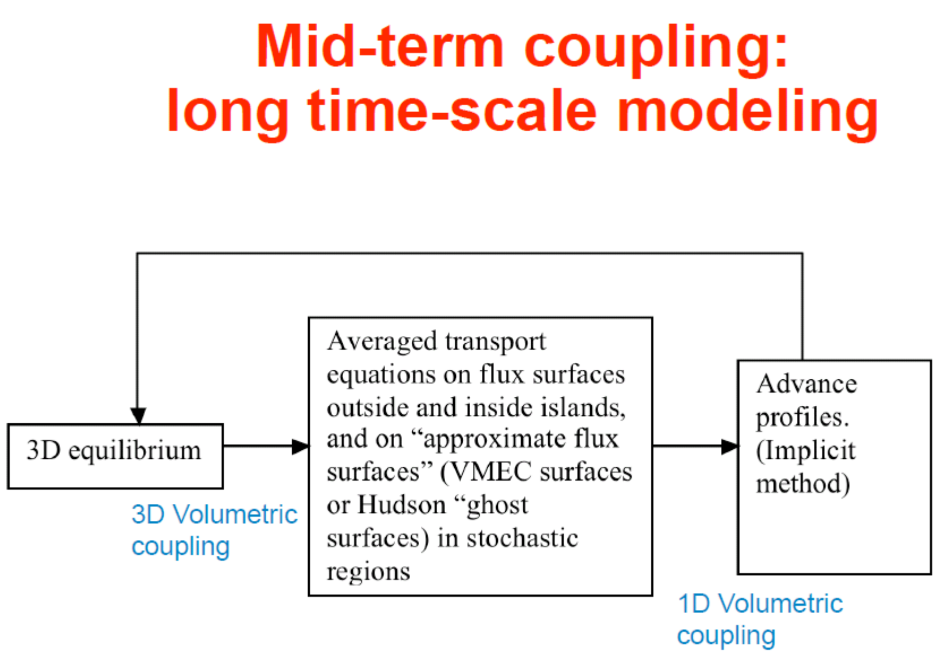

Long-term coupling

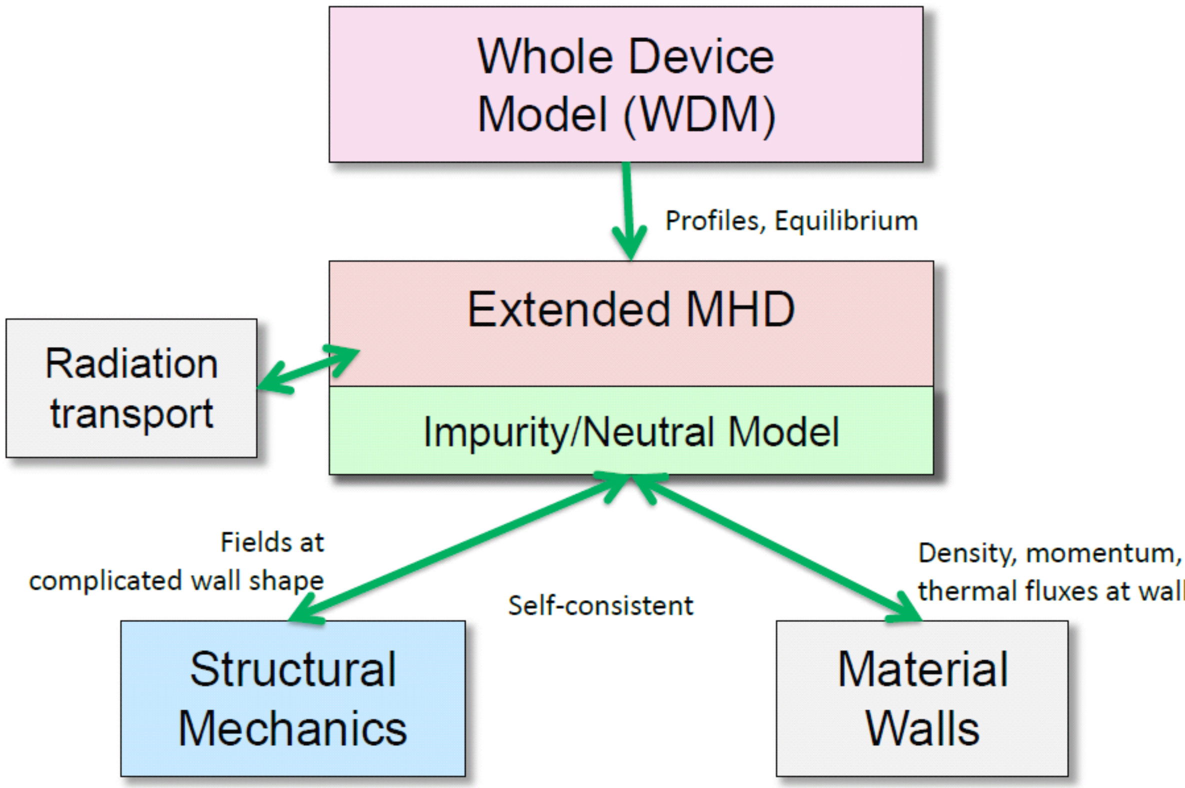

In the long-term perspective, the essential elements remain the same, but the coupling is two ways. Namely the dynamics of runaway will be fed back to the plasma evolution. Not only the plasma impacts the wall, the wall response to the plasma will also be taken into account when evolving the plasma. The component also has much greater physics functionality. For example, the extended MHD component now computes the kinetic MHD in 3D field in the presence of gyrokinetic modeling of the thermal quench.

Example advanced component: Kinetic MHD in 3D field for Modeling Thermal Quench

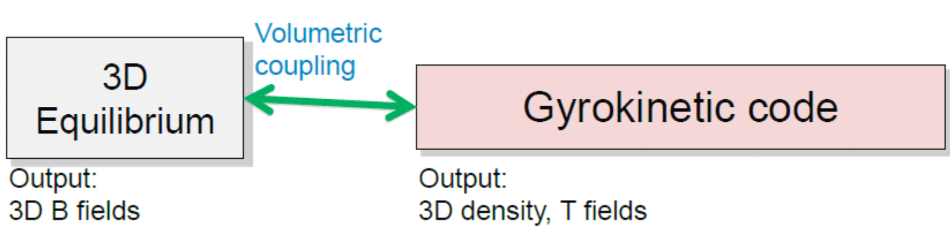

Early: Single-helicity island
Eventually: Stochastic field

Appendix D: Pedestal and Edge Science Drivers

Due to the complexity of the edge physics, the edge component itself needs to be on an edge integration framework to combine the edge subcomponent codes in multi scale. First of all, the non-thermal equilibrium nature of the scrape-off physics, the orbit loss effect in the pedestal region, and the significance of the neoclassical dynamics in self-organizing interaction with the turbulence call for the full-function (full-f) multiscale kinetic simulation. The kinetic simulation time scale is on the order of several collision times (up to 10 ms in DIII D edge). Secondly, simulation of the edge localized mode is an essential element of the edge physics, whose study requires a two-fluid or MHD simulation up to hundreds of Alfven time scale (equivalently up to 0.5 ms in DIII-D). Thirdly, the neutral and impurity transport with the proper atomic physics effect included is another vital component of the edge physics, whose transport time scale is between the toroidal charged particle dynamics time (or turbulence correlation time) and the MHD dynamics time scale (Alfven time). The spatial scale length in the edge physics also has a wide dynamic range: from ion or electron gyroradius (less than 0.1 mm to a few mm) for the kinetic turbulence physics, to the machine size (>m) for the neoclassical and MHD physics.

The complication in the edge simulation component is elevated by the diverted magnetic field geometry and the existence of the material wall. Diverted magnetic field contains separatrix and X-point. The material wall, which bounds the scrape-off plasma, is arbitrarily shaped, electrically grounded usually, absorbs the charged particles, and emits neutral particles. The magnetic separatrix and the material wall do not permit the conventional “core” physics codes, especially the kinetic codes, to study the edge physics. As a result, new edge kinetic codes need to be developed (e.g., XGC0, XGC1, Tempest, COGENT). Only a very small number of the existing fluid/MHD codes can be used as edge subcomponents (e.g., M3D, NIMROD, BOUT) for simulation in a realistic edge geometry.

Due to the highly limited number of currently available edge codes, which can handle the magnetic separatrix and material wall, we may want to include all the “claimed” edge codes in this report and categorize them into three areas according to the applicability regions:

- Codes crossing the magnetic separatrix (e.g., BOUT/BOUT++, XGC0, XGC1, Tempest, Cogent, NIMROD, M3D, UEDGE, DEGAS2, GTNEUT, DIVIMP, etc)
- Codes on closed magnetic field lines (GEM, Elite, EPED, Gyro, GS2, GTC, GTS, NEO, GTC-Neo, etc)
- Codes on open magnetic field lines and/or in contact with the material wall (SOLT, OEDGE, GRETIN, HEIGHTS, etc)

The code names listed above are only for examples. A more complete list of the edge codes provided to the FSP edge component team is attached at the end of this section. The nonlocal nature of the edge physics (large banana excursion, turbulence mesoscale length, and MHD scale length are comparable to the radial scale length of the edge plasma), calls for the necessity of the volumetric coupling without the radial interfacial boundaries in the edge code integration. Thus, in the long term FSP phase, all the edge subcomponent codes will need to be residing across the magnetic separatrix from the core to the wall. Reliable grid interpreters are expected to be important to accommodate different grids used by different physics codes.

The physics phenomena which need to be studied by the edge component framework include, but not limited to, H-mode transition, edge pedestal buildup, edge pedestal shape, edge localized mode stability and crash, plasma flow, heat and particle load on the wall, neutral particle transport, impurity transport, radiation, and edge plasma interaction with RF antenna. The plasma-wall interactions and the atomic processes occur at much faster and finer scales than the edge plasma physics scales, and are beyond the scope of this section. The edge plasma subcomponents should include the plasma-wall interaction and the atomic physics in the form of numerical data table.

More specifically, the experimental findings which need to be understood and validated by the edge component framework, and to be extended to predictions for ITER, include

- Strong core heating induces H-mode transition and builds edge pedestal.
- Core fusion quality goes up together with the edge pedestal height.
- Core plasma rotation increases as the edge rotation increases, leading lower H-mode transition power and to more stable plasma operation.
- As the edge pedestal grows steep, edge localized modes (ELMS) appear and destroys the edge pedestal.
- ELMs degrade the fusion quality in the core and damage the material wall
- Resonant magnetic perturbations (RMPs) driven by external coil array show signs of ELM mitigation in DIII-D and JET.
- Even in the absence of ELMs, the heat load on the divertor plate tends to be localized, which could prematurely damage the ITER's divertor wall.
- Neutrals and impurities, generated by plasma-material interaction, penetrate into the core plasma and significantly affect the pedestal buildup, the plasma fueling, and the fusion quality.
- Radiative loss of the edge plasma energy can significantly influence the divertor heat load.
- Edge plasma can be detached from the divertor plates, reducing the heat load.

Figure 9 shows diagram for a 2-5 years short term goal, which can be achieved by the edge component framework. The backbone of the edge integration framework used in the figure is a flux-flux driven, kinetic, multi-species transport modeling code, due to the non-thermal equilibrium nature of the edge physics as discussed earlier. This short term goal is compromised by the use of spatially separated electromagnetic turbulence codes: i.e., electromagnetic fluid codes in the SOL and electromagnetic kinetic or fluid codes in the core, both of them being pushed toward the separatrix surface from each side without being too close to it. Spatial connection of the nonlocal physics across the separatrix surface will be an issue. An edge electromagnetic kinetic code which can simulate the plasma across the magnetic separatrix needs to be aggressively developed and implemented into this short term integration framework. Other gaps are listed, too, in the diagram. Other short term compromises accepted in Figure 9 are the use of the linear stability codes inside the magnetic separatrix surface to separate the ELM onset criteria from a global nonlinear ELM crash phenomena, and the limiting the ELM control to the RMP method only. Coupling between the subcomponents codes are not yet volumetric in the short-term goal because of the local simulation of the turbulence and the ELM onset physics. After the global edge kinetic codes are available in electromagnetic modes, the kinetic physics will be volumetric. Recommendable data-flow methods (file-based, memory-based, or one-executable based) are shown in the diagram depending upon the amount and frequency of the data to be transferred in the coupled simulations and the structure of the subcomponent codes. For example, DEGAS2 has been integrated into XGC0 as one executable, but M3D-mpp

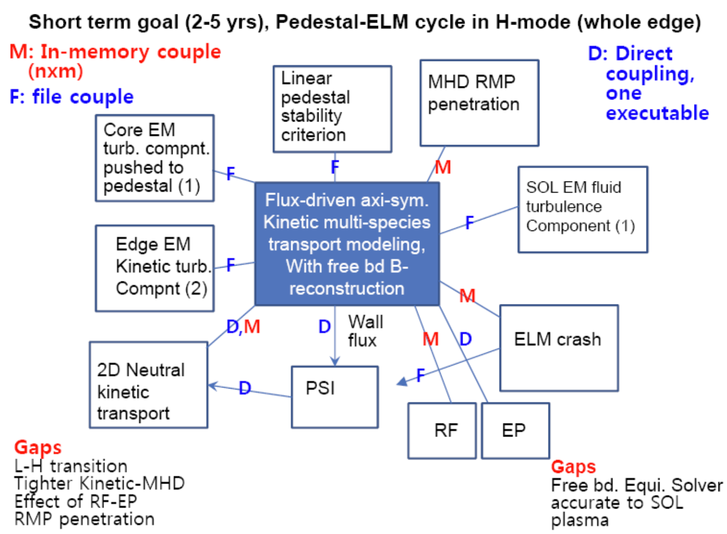

Figure 9. An example short term (2-5 years) goal, starting with the regionally separated local electromagnetic turbulence codes, followed by global edge electromagnetic kinetic code

and XGC0 are independent codes in the integration through memory-based coupling due to the large, real-time exchange of the 3D MHD fluctuation data.

Figure 10. An example the long term (10 years) goal, with the inclusion of the volumetric coupling between all the subcomponents and a more general ELM shows a longer term (10 years) goal of the edge component framework, which now uses the electromagnetic kinetic turbulence codes and more complete ELM codes whose role is both ELM onset and nonlinear crash in the global volume domain from the core to the wall across the magnetic separatrix. The coupling between the subcomponents will then be truly volumetric, without any interfaces between the core and scrape-off plasmas. Pellet and other ELM control physics will be integrated to the edge framework in this phase.

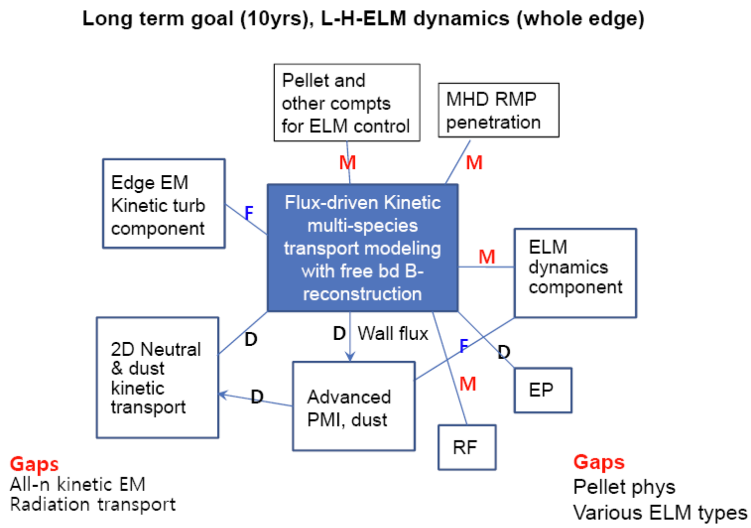

Figure 10. An example the long term (10 years) goal, with the inclusion of the volumetric coupling between all the subcomponents and a more general ELM

Similar diagrams can be drawn if one wishes to study the pedestal physics or the scrape-off physics separately without considering its connection with the other regions. These will be rather short term program in the FSP time scale. One example is a static EPED-like approach in the pedestal.

Some of the edge physics phenomena discussed here are known to be connected to the core physics. Large scale turbulence, MHD and RF phenomena belong to such examples. These large scale nonlocal physics can nonlinearly interact with the small scale local physics to feedback to each other. Thus, recognition of the common components across the different volume or physics space and getting these common components on solid ground will be an important task of FSP. From the edge point of view, the ELM physics, RMP physics, free boundary equilibrium solver, neutral particle fueling, neoclassical tearing modes, and pellet physics are other common component examples.

Rigorous validation activities by experimentalists will be a necessary part of these short and long term edge simulation goal activities. The validity of the localized turbulence models and ELM models used in the early-phase can only be tested against the experimental validation. Such validation activities may also render some incite in the improvement of the reduced transport models.

The following table summarizes the edge subcomponent codes, which have been provided to the edge component team up to the present time. The list is expected to grow as more codes are identified and reported to the FSP team.

Components used in edge (crossing separatrix)

- Free boundary Equilibrium and magnetic reconstruction
  - EFIT (1.1) (L. Lao), (2D) serial
  - TEQ (1.2) (L. Lodestro), (2D) fixed or f ree bd, realistic coils, serial
  - M3D_omp (1.2) (H. Strauss, G. Park), (2D) free bd, model coils, OpenMP,
  - ESC (1.2) (L. Zakharov) (2D), free bd, serial
  - PIES (1.3) (PPPL) (3-D), free bd, being parallelized
- Mesh generation (2-D)
  - UEDGE (3.1) flux surface following
  - Carre (3.1) flux surface following
- Linear stability
  - 2DX (Lodestar, LLNL) linear eigenvalue code, including separatrix (small scale)
- MHD (1.3) (all 3D)
  - NIMROD (C. Sovinec), extended, medium scale
  - M3D_mpp (L. Sugiyama), extended, medium scale
  - M3D-C1 (S. Jardin), extended, expected to be medium-high scale (~10K)
  - PIXIE (L. Chacon), extended, separatrix not yet, but yes in future
  - BOUT++ (LLNL / U. York), reduced MHD, medium to large scale
- Neoclassical & anomalous transport (4.5,4.2) (2-D, 2V)
  - XGC0 (G. Park, C.S. Chang), Lagrangian, small to large scale
  - TEMPEST (X. Xu), Eulerian, medium to large
- Fluid turbulence (4.1) (3D)
  - BOUT / BOUT++ (X. Xu, M. Umansky, U. York), small to moderate scale
- Kinetic turbulence (4.3) (includes neoclassical) (3-D, 2V)
  - XGC1 (S. Ku, C.S. Chang), Lagrangian, extreme scale
  - Cogent (R. Cohen), Eulerian, under development, medium to large scale
- Transport modeling (4.2) (2-D)
  - UEDGE (T. Rognlien), (2-D) small scale with MPI
- Neutral transport
  - DEGAS2 (5.2) (D. Stotler) (3D, 3V), Monte Carlo, small to medium scale
  - UEDGE (5.1) (T. Rognlien) (2D) fluid neutral, small scale with MPI
  - GTNEUT (5.1) (W. Stacey), (2D) multi-group, small scale
- Dust transport
  - DUSTT (R. Smirnov) (3-D, 3V), PIC code (parallilzable), small scale
- Impurity transport
  - DIVIMP (P. Stangeby) (3-D, 3V), PIC code

Components used in edge (on closed field lines)

- Fixed boundary equilibrium codes (1.2)
  - TOQ, JSOLVER…
- Pedestal linear stability (all serial)
  - Elite (P. Snyder)
  - Gato (A.D. Turnbull), fixed or free bd,
  - EPED model (P. Snyder)
  - PEST (J. Manickam)
  - DCON (A. Glasser)
  - Full (G. Rewoldt)
- Core EM turbulence code mimicking edge pedestal (4.3) (MPI)
  - GEM (S. Parker), global, extended, medium to large scale
  - Gyro (J. Candy), global, extended, medium scale
  - GS2 (W. Dorland and M. Barnes), local, medium to large scale
  - GTC / GTS (Z. Lin / W. Wang), global, large scale
- Core neoclassical code mimicking edge pedestal (4.5)
  - Neo (E.A. Belli), small scale
  - GTC-Neo (W. Wang), medium to large scale
- Transport modeling (4.2)
  - BALDUR (G. Bateman), small scale
  - ONETWO (H. St. John), small scale
  - PTRANSP (D. McCune), medium scale
  - CORSICA (L. LoDestro), medium scale

Components used in edge (on open field lines)

- Near surface physics / sheath codes
  - BPH-3D (J. Brooks)
  - WBC / REDEP (J. Brooks)
  - HEIGHTS (A. Hassanein)
- Materials
  - Tables (e.g., from TRIM-SP) (10.1)
  - Wall-PSI (10.2) (A. Pigarov)
  - AMD (10.5)
  - HEIGHTS (10.2) (A. Hassanein)
- Atomic physics
  - Tables (9.1, 9.2) (e.g., from ADAS, collisional radiative models)
- SOL transport
  - OEDGE (P. Stangeby) (analysis code)
- Fluid turbulence
  - SOLT (Lodestar) reduced model, slab geometry
- Radiation transport (7.3)
  - CRETIN (H. Scott)
  - HEIGHTS (A. Hassanein)

Figure 11. Components Lists