DOE/SC-0073

Fusion Energy Sciences Advisory Committee

Panel on Integrated Simulation and Optimization of the Magnetic Fusion Systems

November 2002

U.S. Department of Energy
Office of Science

FUSION SIMULATION PROJECT

Integrated Simulation & Optimization of Fusion Systems

Final report of the FESAC ISOFS Subcommittee

ISOFS (Integrated Simulation & Optimization of Fusion Systems) Subcommittee Members:

Jill Dahlburg, General Atomics (Chair)
James Corones, Krell Institute, (Vice-Chair)
Donald Batchelor, Oak Ridge National Laboratory
Randall Bramley, Indiana University
Martin Greenwald, Massachusetts Institute of Technology
Stephen Jardin, Princeton Plasma Physics Laboratory
Sergei Krasheninnikov, University of California - San Diego
Alan Laub, University of California - Davis
Jean-Noel Leboeuf, University of California - Los Angeles
John Lindl, Lawrence Livermore National Laboratory
William Lokke, Lawrence Livermore National Laboratory
Marshall Rosenbluth, University of California - San Diego
David Ross, University of Texas - Austin
Dalton Schnack, Science Applications International Corporation

EXECUTIVE SUMMARY

Fusion is potentially an inexhaustible energy source whose exploitation requires a basic understanding of high-temperature plasmas. The development of a science-based predictive capability for fusion-relevant plasmas is a challenge central to fusion energy science, in which numerical modeling has played a vital role for more than four decades. A combination of the very wide range in temporal and spatial scales, extreme anisotropy, the importance of geometric detail, and the requirement of causality which makes it impossible to parallelize over time, makes this problem one of the most challenging in computational physics. Sophisticated computational models are under development for many individual features of magnetically confined plasmas and increases in the scope and reliability of feasible simulations have been enabled by increased scientific understanding and improvements in computer technology. However, full predictive modeling of fusion plasmas will require qualitative improvements and innovations to enable cross coupling of a wider variety of physical processes and to allow solution over a larger range of space and time scales. The exponential growth of computer speed, coupled with the high cost of large-scale experimental facilities, makes an integrated fusion simulation initiative a timely and cost-effective opportunity.

Worldwide progress in laboratory fusion experiments provides the basis for a recent FESAC recommendation to proceed with a burning plasma experiment (see FESAC Review of Burning Plasma Physics Report, September 2001). Such an experiment, at the frontier of the physics of complex systems, would be a huge step in establishing the potential of magnetic fusion energy to contribute to the world's energy security. An integrated simulation capability would dramatically enhance the utilization of such a facility and lead to optimization of toroidal fusion plasmas in general. This science-based predictive capability, which was cited in the FESAC integrated planning document (IPPA, 2000), represents a significant opportunity for the DOE Office of Science to further the understanding of fusion plasmas to a level unparalleled worldwide.

The ISOFS Subcommittee recommends that a major initiative be undertaken, referred to here as the Fusion Simulation Project (FSP). The purpose of the initiative is to make a significant advance within five years toward the ultimate objective of fusion simulation: to predict reliably the behavior of plasma discharges in a toroidal magnetic fusion device on all relevant time and space scales. By its very nature in enabling more comprehensive modeling, the FSP will lead to a wealth of insights not realizable previously, with new understanding in areas as diverse as wall interaction phenomena, the effects of turbulence on long time confinement, and implications of plasma self heating in advanced tokamak operating regimes. The long-term goal is in essence the capability for carrying out 'virtual experiments' of a burning magnetically confined plasma, implying predictive capability over many energy-confinement times, faithful representations of the salient physics processes of the plasma, and inclusion of the interactions with the external world. Since confidence in the ability to predict is ultimately based on code performance against experimental data, a vigorous and ongoing validation regime must also be a critical element of this project.

The characteristics of fusion plasmas make the goal extremely challenging. These characteristics include the presence of multiple time scales, ranging over fourteen orders of magnitude, and multiple spatial scales, ranging over eight orders of magnitude. The linear algebraic systems that must be solved are often ill-conditioned. The computational domains are geometrically complex, and the solutions severely anisotropic. In many cases, the physics approximations are not completely understood, and hence the simulation equations are unclear. The underlying physics is coupled with essential nonlinearities. Taken in isolation, approaches have been developed or are under investigation for each of these challenges. However, an integrated simulation for fusion plasmas will present *all* of these features simultaneously.

Success of this project will require coordinated and focused advances in *fusion physics* (to further develop the underlying models and elucidate their mathematical basis), *applied mathematics* (to further develop suitable algorithms for solving the mathematical models on the appropriate computer architecture, and to define frameworks within which these algorithms may be easily assembled and tested), and *computer science* (to provide an architecture for integrated code development and use, and to provide analysis and communication tools appropriate for remote collaboration). Strong collaborations, forged across these disciplines and among fusion scientists working in different topical areas, will be an essential element of the program. In addition, the Fusion Simulation Project will require significant improvements in computational and network infrastructure, including enhancements to shared resources as well as to local or topical computing centers. Because of the complexity of the FSP, the planning process should continue into CY2003. We recommend a staged approach: beginning with clarification of the physics issues, accompanied by efforts to address algorithmic issues and followed by clarification of architectural issues.

The necessary core expertise for the FSP is resident in several units within the DOE Office of Science. Primary among these are the ongoing fusion experimental and theoretical research and development activities within the Office of Fusion Energy Sciences, the applied mathematics development activities within the Office of Advanced Scientific Computing, the recently developed SciDAC initiative, and materials science research in the Office of Basic Energy Sciences.

To achieve its goals, the FSP is envisioned as proceeding through three five-year phases in which successively more complex and disparate phenomena will be integrated. During the first five years, the project will concentrate on specific physics integration issues that are expected to deliver significant scientific insights in their own right, but are also prototypical of the integration issues faced by the whole initiative. Each Focused Integration Initiative (FII) will concentrate on developing a predictive modeling capability for a specific programmatically important scientific problem and will begin to develop and gain experience with relevant mathematical tools, new algorithms, and computational frameworks. During the second five-year period the project will undertake larger and more comprehensive integration activities and take them to the next level of development. During the final five-year period, the focus will be on comprehensive integration. There will be links among all the physics components of the project. To provide a tradeoff between computational efficiency and physical fidelity there will be multiple levels of description of many of the physical processes.

Verification and validation are critical components of the FSP. To succeed, an integral feature of this initiative must be an intensive and continual close coupling between the simulation efforts and experiments. The phenomena in magnetic fusion devices, the equations describing them, and the interactions among the various critical phenomena are sufficiently complex that developing the most effective approximations and establishing that the models have the required accuracy can only be accomplished by continual iteration and testing against experimental data.

Funding for the FSP must be at a level adequate to accomplish the project goals. The successful NNSA Accelerated Strategic Computing Initiative (ASCI) Level 1 University Centers Program, funded at $25M/yr, provides an appropriate example of the level of resources required. A preliminary assessment of the challenges and complexity of possible FIIs indicates that they would be comparable to that of each of the five ASCI University Level 1 Center Programs. We further estimate that four-five such FIIs will be required to cover all the critical science areas which must eventually go into the final integrated simulation code. Further refinement of the costs and timelines will be carried out as the FSP is developed. Through the course of the project, we envision that funding would be approximately equally allocated between the DOE OFES and OASCR research elements. Because this initiative rests entirely on a progressing science base, and will for successful execution attract and retain junior researchers committed to the goals of fusion energy sciences, it is paramount that FSP funing be new rather than redirected from present critical areas.

TABLE OF CONTENTS

Executive Summary

I. Background 5

II. Overview and Recommendations 7
- a. Goals: 5,10,15 Year Overview 7
- b. The Focused Integrated Initiative (FII) Approach 8
- c. FSP Project Size and Scale 8
- d. Infrastructure 9
- e. Due Attention to Governance 10
- f. The Need to Continue the Planning Process 11

III. The FSP Scientific Elements 12
- a. Introduction 12
- b. Focused Integration Initiatives 14
- c. FII Examples 16
- d. Insights 19
- e. Computational Science 20
- f. Verification and Validation 26
- g. Infrastructure Requirements 29

IV. The FSP Path 33
- a. Project Roadmap 33
- b. The FSP and other Office of Science Activities 35
- c. Roadmap Implementation: Project Governance & Management 36

V. Summary 39

VI. Acknowledgements 41

Attachment: ISOFS Charge Letter 42

Appendix. Overview of Fusion Science for the FSP . . . . . Vol 2

I. BACKGROUND

In February 2002, the DOE Office of Science asked the Fusion Energy Sciences Advisory Committee (FESAC) to assist in defining a major new initiative to be sponsored jointly by the Office of Fusion Energy Sciences (OFES) and the Office of Advanced Scientific Computing Research (OASCR). The goal of this initiative, the Fusion Simulation Project (FSP), is to create a comprehensive set of theoretical fusion models, an architecture for bringing together the disparate physics models, combined with the algorithms andcomputational infrastructure that enable the models to work together. The required funding level for the FSP is expected to be on the order of $100M spread over five to six years. A FESAC ISOFS (Integrated Simulation and Optimization of Fusion Systems) Subcommittee, with members from the fusion, applied mathematics and computer science communities, was constituted to generate a plan for moving forward with the FSP. The ISOFS Subcommittee membership is listed on the cover page of this document. The 2002 timeline of the ISOFS Subcommittee is shown in Fig.I.1.

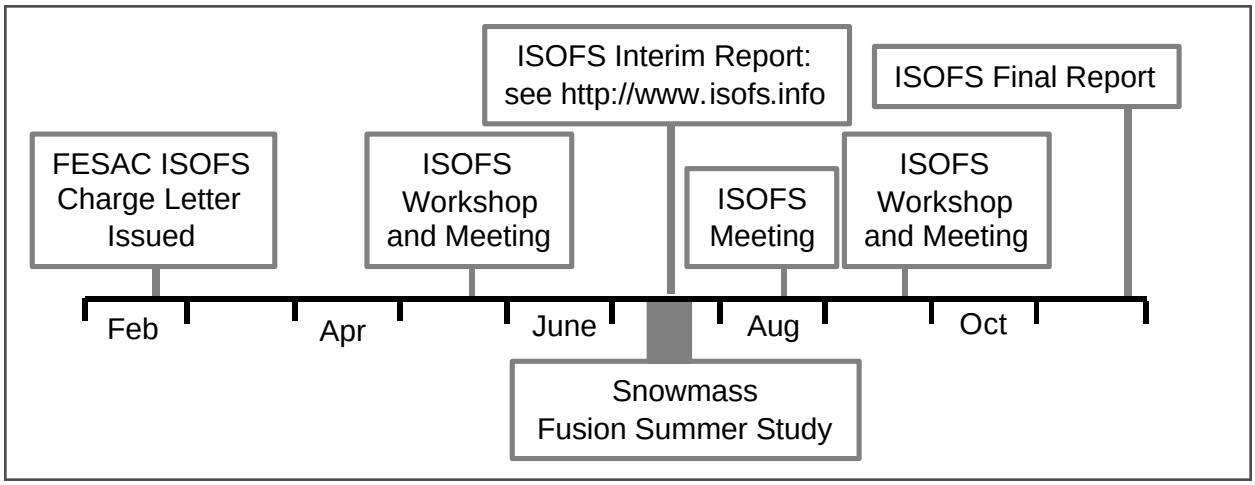

Figure I.1. *FESAC ISOFS Subcommittee activities timeline for 2002. ISOFS Workshop presentations and discussion may be found at: http://www.isofs.info .*

Impetus and fundamental interest for the FSP initiative primarily comes from the goal to develop an attractive fusion energy source. The fossil fuels that underpin the United States economy cannot be relied upon to carry our nation into the 22nd Century. Oil and gas are non-renewable resources feeding a rapidly growing global energy appetite. There is also the threat of global climate change due to the burning of fossil fuels.

In the summer of 2002, fusion physicists met at the Snowmass Fusion Summer Study to plan the next stage of research towards the ultimate goal of fusion energy. The 2002 Snowmass Development Pathway Subgroup discussed the major next step plasma physics facilities in the fusion International Portfolio Approach that are required for this goal. These include advanced tokamak and non-tokamak physics facilities, a burning plasma facility(s), a Fusion Plasma Simulator (FPS), and a strong core program. In particular, the FPS is envisioned to be an integrated research tool that contains comprehensive coupled self-consistent models of all important plasma phenomena that would be used to guide experiments and be updated with ongoing experimental results. Most importantly, the

FPS would serve as an intellectual integrator of physics phenomena in advanced tokamak configurations, advanced stellarators and tokamak burning plasma experiments. It would integrate the underlying fusion plasma science with the Innovative Confinement Concepts, thereby accelerating progress. Development of a facilities class FPS capability is estimated to be a fifteen-year, $400M activity. The FSP that is discussed in detail in this report is a first five-year stage of the ultimate FPS. The need for this kind of integrated simulation capability is recognized in the preliminary report of the FESAC Development Path Subcommittee charged with identifying the requirements for the start of operation of a fusion energy demonstration power plant in 35 years.

The workshops, meetings, and regular correspondence of the ISOFS Subcommittee resulted in the vision for the FSP that is described in this report. This final report of the ISOFS Subcommittee provides the response to the FESAC ISOFS charge letter of February 22, 2002; a copy of the letter is in the report Attachment. The report Appendix, an overview of frontier fusion science, addresses aspects of the charge and also provides a self-contained reference summary of fusion science in the context of this initiative. Responses to particular questions contained in the ISOFS charge letter are as follows:

- **What is the current status of integrated computational modeling and simulation?**
  Appendix Section VIII with additional detail in Appendix Sections III-VII.

- **What should be the vision for integrated simulation of toroidal confinement fusion systems?**
  Sections IIa,b, IIIb,c, and Appendix Section IB and IX.

- **What new theory and applied mathematics are required for simulation and optimization of fusion systems?**
  Section IIIe, and Appendix Sections III-VII and IX.

- **What computer science is required for simulation and optimization of fusion systems?**
  Section IIIe.

- **What are the computational infrastructure needs for integrated simulation of fusion systems?**
  Sections IId and IIIg.

- **How should integrated simulation codes be validated, and how can they best be used to enable new scientific insights?**
  Sections IIId,f, and Appendix Section I, IX, and X.

We note that this document contains a refined response to the first two charges above, building upon the initial response in the July 12, 2002 ISOFS interim report.

The FSP computational undertaking represents a significant opportunity and a significant challenge to fusion research, which has always been at the forefront of advanced scientific computing. Integrating fusion computer codes for full-system fusion simulations will require even greater research collaboration among fusion physicists, and applied mathematicians and computer scientists dedicated to putting fusion energy on the power grid. Creating the computational resources to simulate fusion will do more than substantially advance fundamental science. We will give ourselves the ability to see our energy future, and then build it.

II. Overview and Recommendations

The Fusion Simulation Project (FSP) described in this document is designed to provide an integrated simulation and modeling capability for magnetic fusion confinement systems. The FSP is the detailed response to findings of the FESAC Integrated Program Planning Activity (IPPA 2000), which identified the requirement for enhanced simulation for predicting the performance of externally controlled confinement systems. It is recognized that this goal can only be met through extensive and sustained collaborations between fusion scientists, and applied mathematicians and computer scientists. We note that these challenges are coming at a time of increasing opportunity between these groups, recognized in large measure by the DOE Office of Science SciDAC projects, and that the FSP will be able to further the momentum well-fostered by SciDAC. Hence,

We recommend that a major initiative be undertaken, here referred to as the Fusion Simulation Project (FSP), to create a comprehensive set of theoretical fusion models, combined with the algorithms required to realize them and an architecture and computational infrastructure that enable them to work together.

The purpose of the FSP is to make a significant advance toward the ultimate objective of fusion simulation: to predict in detail the behavior of any discharge in a toroidal magnetic fusion device on all important time and space scales. This is in essence the capability for carrying out 'virtual experiments' of burning, magnetically confined plasmas. This requires faithful representations of the salient physical processes individually and their interactions with the external world (sources, control systems and bounding surfaces), leading to a predictive capability over many energy-confinement times,

a. Goals: 5,10,15 Year Overview

The goal of the FSP is to produce a comprehensive fusion simulation tool (the Fusion Plasma Simulator) by the year 2020. This tool will play an essential role in the development path for fusion energy. It will effectively serve as an *intellectual integrator* of physics phenomena in advanced tokamak configurations, advanced stellarators and tokamak burning plasma experiments. In order to achieve this overarching goal, the project will proceed through three five-year phases in which successively more complex and disparate phenomena will be integrated together. We describe this process briefly here and in more detail in Section III.

During the first five years, the project will concentrate on specific high-profile physics integration issues that are considered to be the most critical, and are also prototypical of the integration issues faced by the whole initiative. We expect to gain new scientific insights during this period. We will also develop the mathematical frameworks for the project and gain experience with computational frameworks and new algorithms.

During the second five-year period the project will undertake larger and more comprehensive integration activities and bring them to the next level of development. The mathematical framework will be expanded to include a wider range of integrated phenomena and will become standardized for the project. The project will develop a unified computational framework that aids in managing the increasing complexity, as well as integrating such aspects as advanced graphics and user interface. New algorithms will continue to be developed and refined as needed.

During the final five-year phase, the focus will be on comprehensive integration. There will be links among all the physics components of the project. To provide a tradeoff between computational efficiency and physical fidelity there will be multiple levels of description of many of the physical processes. The simulation capabilities will be extensively exercised, and comprehensive comparisons between the simulation and experiment will take place. The capability will be used to guide experiments and be updated with ongoing experimental results.

b. The Focused Integrated Initiative (FII) Approach

Fusion computations at varying degrees of integration have already led to significant insights pertaining to the physics mechanisms underlying the performance of plasmas confined in a range of toroidal magnetic configurations. We expect that the FSP will lead to new surprises coming from more comprehensive models and emerging from enhanced synergy between theory, experiments and modeling. This integration initiative provides a tremendous opportunity to garner new insights by addition of new physics to the plasma models and by enabling more comprehensive models through integration.

In order to realize integration from the beginning of the project, we recommend that the FSP commence with programmatic teams, subsets of the full FSP that we term Focused Integration Initiatives (FIIs). We describe the FIIs in detail in Section III*b*. The goal of each FII team is the solution of a compelling problem in fusion science physics that requires integrated simulation. The FIIs should be multi-disciplinary and multi-institutional, and by their research should integrate subsets of the full breadth of fusion fundamentals and applications of varying complexity using selected algorithms and interoperable software. The traditional modeling elements that structure our understanding of fusion plasmas include: plasma sources; turbulence; extended MHD; 1.5D (one and one-half dimensional) transport; and fusion materials. Each FII should cut across and integrate two or more of these traditional elements, to provide physics integration both spatially and temporally, with a guiding focus of a single overarching scientific question or topic that satisfies the criterion of importance to the fusion program. The community will be invited to define overarching FII themes through the proposal process.

As we envision it, each FII will focus on achieving predictive modeling capability for the particular fusion science problem it has elected to address. In order to develop a critical mass of research with adequate intellectual vibrancy, and to encourage development path risk and opportunity, the FSP should be initially comprised of 4-5 FII units. Primary to each of the FII activities must be verification of the accuracy of the new integrated model developed within the FII, and validation of the model with experimental data. Verification and validation — critical components for the FIIs — imply non-trivial supporting access to experiments, experimental data and diagnostics. We thus recommend close coupling of each FII research team with relevant experiments, and that the development of a reliable experimental predictive capability should be a substantive part of each FII.

c. FSP Project Size and Scale

We strongly recommend that within the five-year time frame specifically considered to be the FSP, the initiative should be carried out at a scale such that certain computational goals can be achieved:

1) Robust computational modules are developed in each of the selected FII areas representing the state-of-the-art in physics content, numerical methods, and computational science methods, enabling efficient incorporation into the integration framework.

2) Approaches are developed for the fundamental problems of disparate time or space scales, and coupling of models of processes having different dimensionalities.

3) An initial inter-operable code capability that allows for three-dimensional geometry is available for widespread testing as a research tool.

4) The effectiveness of the integration approach is demonstrated by application to interpreting experimental data, and testing the validity of various physics models.

This initiative rests entirely on a progressing science base. Therefore it is paramount that FSP funding be new rather than redirected from present, critical areas. We fully support the assessment of the importance of the core fusion program that was stated in the September 2002 Burning Plasma Strategy Report: 'The core program is … essential to the successful and full exploitation of the burning plasma program. Predictions on the confinement, stability properties and dynamics of plasmas in the burning regime have all come from the intense experimental, modeling and theoretical efforts of the core program. The underpinnings of any burning plasma experiment therefore fundamentally rests on the foundation of knowledge that has come from the core program. Moving forward with a burning plasma experiment requires experimental scientists, engineers, and theorists and computational scientists from this core to design experiments and interpret the results.'

Further, funding for the FSP must be at a level adequate to accomplish the FSP goals. To derive an adequate funding profile that will enable a critical mass of research, we use the successful $25M/year DOE Accelerated Strategic Computing Initiative (ASCI) Level 1 University Centers Program as an example. In that program, each strategic element or Center is receiving $4-5M/year for each of 5-10 years. This indicates that each FSP FII team should be initiated with funding of about $4-5M/year, and that the FSP will require approximately $20M/year for each year of the project. Through the course of the project, funding should be approximately equally allocated from the OFES and OASCR research elements.

d. INFRASTRUCTURE

The FSP will be integrated into broader fusion science and simulation activities. For example, access to experiments, experimental data and diagnostics are critical to the success of the initiative. Likewise, reliable access to suitable computing facilities will be required, including data storage and networking, and also collaborative tools. None of these items are included in the budget for the FSP as envisioned in this report. However we do stress here the need to supply a variety of computational platforms. Computational infrastructure that includes platforms of extremely high capability, and also high performance networks, will certainly be needed to achieve the goals of the FSP. The FSP will push the available envelopes of both sustained performance and long-distance researcher collaboration from the outset and as the project moves forward. The graphic below (Fig. II.1) illustrates fusion simulation performance projections in the context of past accomplishments. Storage and networking needs for the future simulation activities can be deduced from the increasing memory requirements necessary for the simulations.

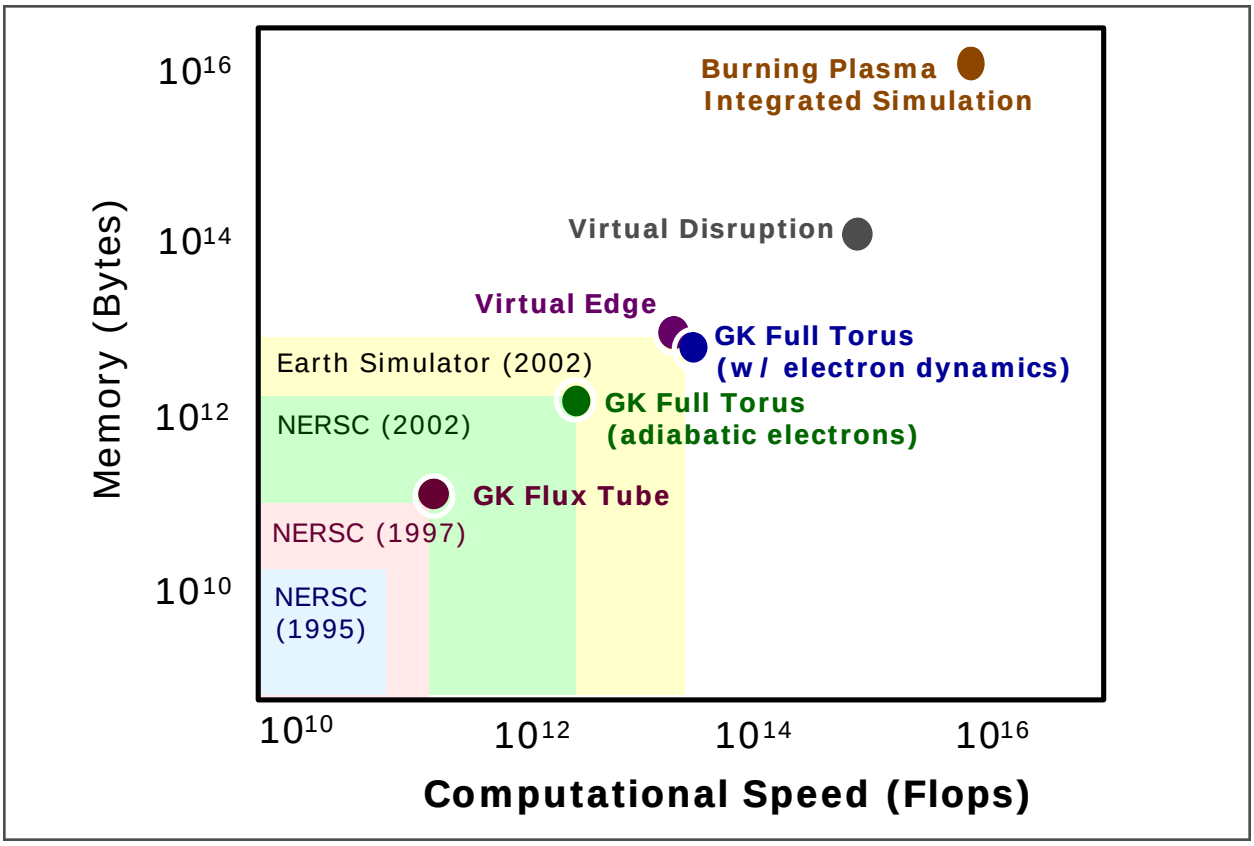

Figure II.1. *Computational requirements for fusion simulations.*

Further, it is essential to realize that many aspects of the project will require readily available computational cycles in real time for program development and debugging. To assure high levels of productivity by researchers, the latter are profoundly critical for success along with the ultra-scale simulation capability.

e. Due Attention to Governance

To bring these disparate components together will require the dedicated skills of many accomplished physicists, applied mathematicians and computer scientists. There is no doubt that the sociology of the FSP will be a challenge. On the one hand, a strong fusion physics effort is required, involving a number of institutions and the relevant theory, simulation, and experimental communities, each of which will bring a required degree of intellectual independence. On the other hand, setting priorities and a considerable amount of central direction will be essential, for the reason that the FSP must be a coordinated, goal driven activity. Even more challenging will be effective integration of first-rate computer scientists and applied mathematicians as full partners with fusion physicists in this venture. The issue of project governance includes also the establishment of an effective cooperative arrangement between and within the two sponsor entities, OFES and OASCR, and clear delineations of working relationships with other initiatives and activities in the DOE such as the Office of Science SciDAC, the OFES fusion experiments, and OASCR computing resources. Harmonizing all of these elements, particularly in light of robust institutional competition (which is a strength of the DOE laboratory system), will require innovative, flexible management in the field and at headquarters.

Early success will require planning, leadership, and likely new management approaches. Should success be achieved, the FSP could provide a template for other large-scale cross-disciplinary computational initiatives for the future. It is our view that the next paradigm shift in problem solving ability from large scale computation may be fueled by this and other comparable major collaborative computational projects that are now ongoing (e.g., the Community Climate Systems Model [see, e.g., Kiehl, 23 May, 2002, http://www.isofs.info]). The sustained effort that may be required to coordinate this project to a successful outcome is well balanced by this potential. Both the investigation of the new fusion science that will be enabled, and also the possible outcomes achievable from the novel investigation of high-end computational science paradigms, well justify a substantial degree of thoughtful planning from the outset.

f. The Need to Continue the Planning Process

As noted above, two workshops on the FSP were held in 2002, which brought together fusion scientists, applied mathematicians and computer scientists with an interest in the FSP. The September 2002 meeting had two major goals: obtaining technical input for this report and establishing and enhancing contact among the participating communities. The level of intellectual energy and enthusiasm for the FSP activity was very high at the September meeting, and more than 100 technical researchers participated. The stage is clearly set for broad participation from all relevant sectors. While the current report addresses the relevant strategic technical issues, further thinking must be done to develop a working program plan. Because of the complexity of the FSP, this planning process is staged: *first* clarification of the physics issues; *next* clarification of the algorithmic issues; and, *finally* clarification of architectural issues. This planning process is ongoing and overlapping in time, and we expect that it will continue through the life of the FSP.

For success, we believe that the FSP planning process that has now begun should continue during 2003. Several sorts of activities should be considered:

- focused technical workshops that continue to broaden participation among fusion physicists and applied mathematicians and computer scientists;
- small working groups that begin to clarify and define the software architecture, including documenting requirements;
- venues for the clarification of needed collaborative tools; continued integration of the outputs of the above by the ISOFS Subcommittee — or whichever future organization DOE decides to enfranchise in this role — into a detailed planning document that will lead to a suitable FSP proposal call;
- as technical planning becomes more refined, activities that provide more accurate budget estimates for the duration of the FSP; and,
- attention to new and ongoing international activities in these areas with a goal of fostering collaboration where feasible.

III. The FSP Scientific Elements

a. Introduction

Central to the understanding of fusion plasmas are fusion experiments. A toroidal fusion experiment, for example the tokamak shown in Fig.III.1a, consists of an inner plasma of ionized gas confined by a configuration of magnetic fields.

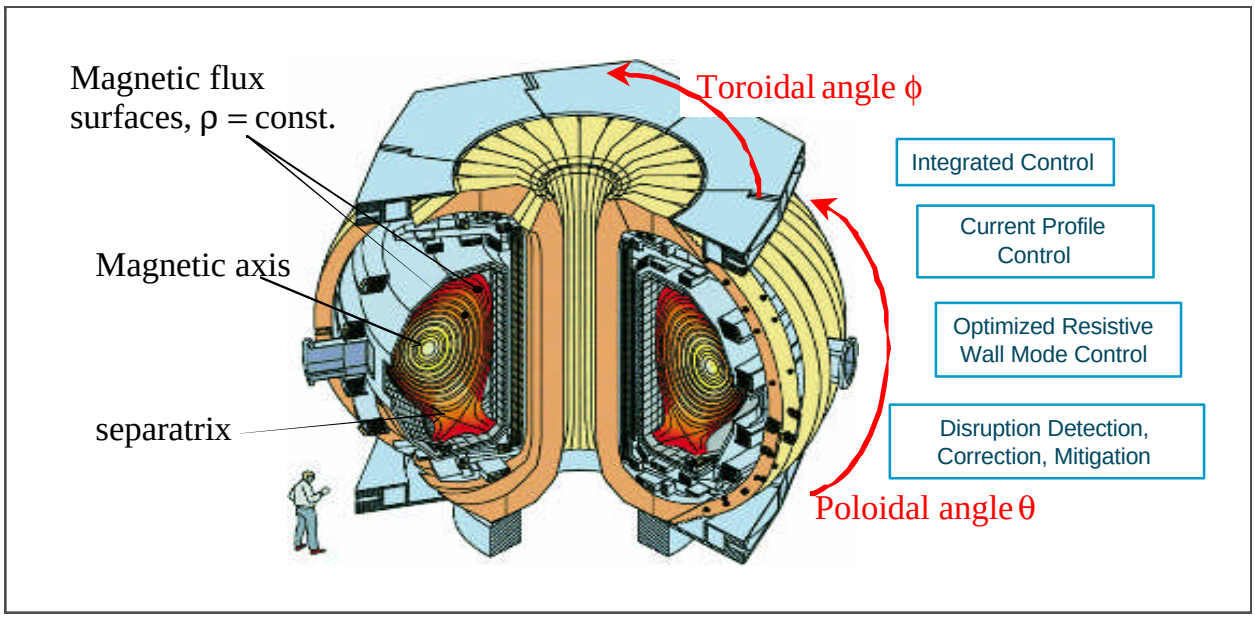

Figure III.1a. *Cutaway view of an advanced tokamak, DIII-D.*

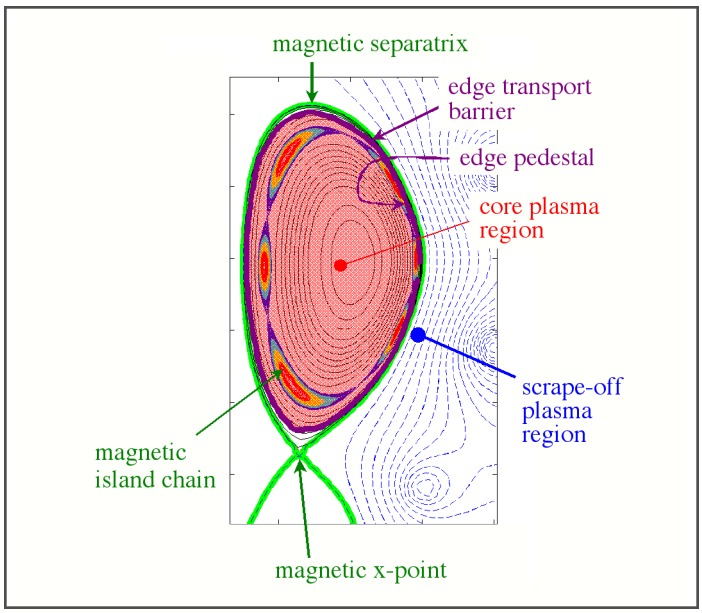

Figure III.1b. *Key plasma and magnetic regions of a typical tokamak plasma, shown as a computed cross-section.*

Plasma containment results from the formation of closed, nested magnetic flux surfaces and the tendency of the individual plasma particles, ions and electrons, to move along magnetic field lines and thus remain close to the flux surfaces. Loss of confinement or transport results from the drifting of particles across these surfaces or from the breakup of the surfaces themselves; see Fig.III.1b. The walls of the device form a vacuum chamber, which is in turn surrounded by the main magnetic coils and the various devices for diagnosing the plasma behavior and for injecting particles, energy, and momentum. Ultimately, the balance of these sources with a wide range of loss mechanisms, together with large-scale instabilities that can disrupt the plasma, determine the performance of the machine. These processes and the models used to describe them are overviewed in some detail in the Appendix to this report.

It is widely recognized that the complexity of the dynamics of fusion experimental systems is such that the development of computational models to understand their behavior is critical. Numerical modeling activities in magnetic fusion research are providing important physics understanding and routinely stretching the limits of available computational resources. However, crosscutting issues crucial to the further development of these models require a qualitative change in approach. In particular, there are two fundamental issues that have to date inhibited the integration of different fusion physics areas: the coupling of phenomena at disparate time scales and the necessity of coupling models of different spatial dimensionality.

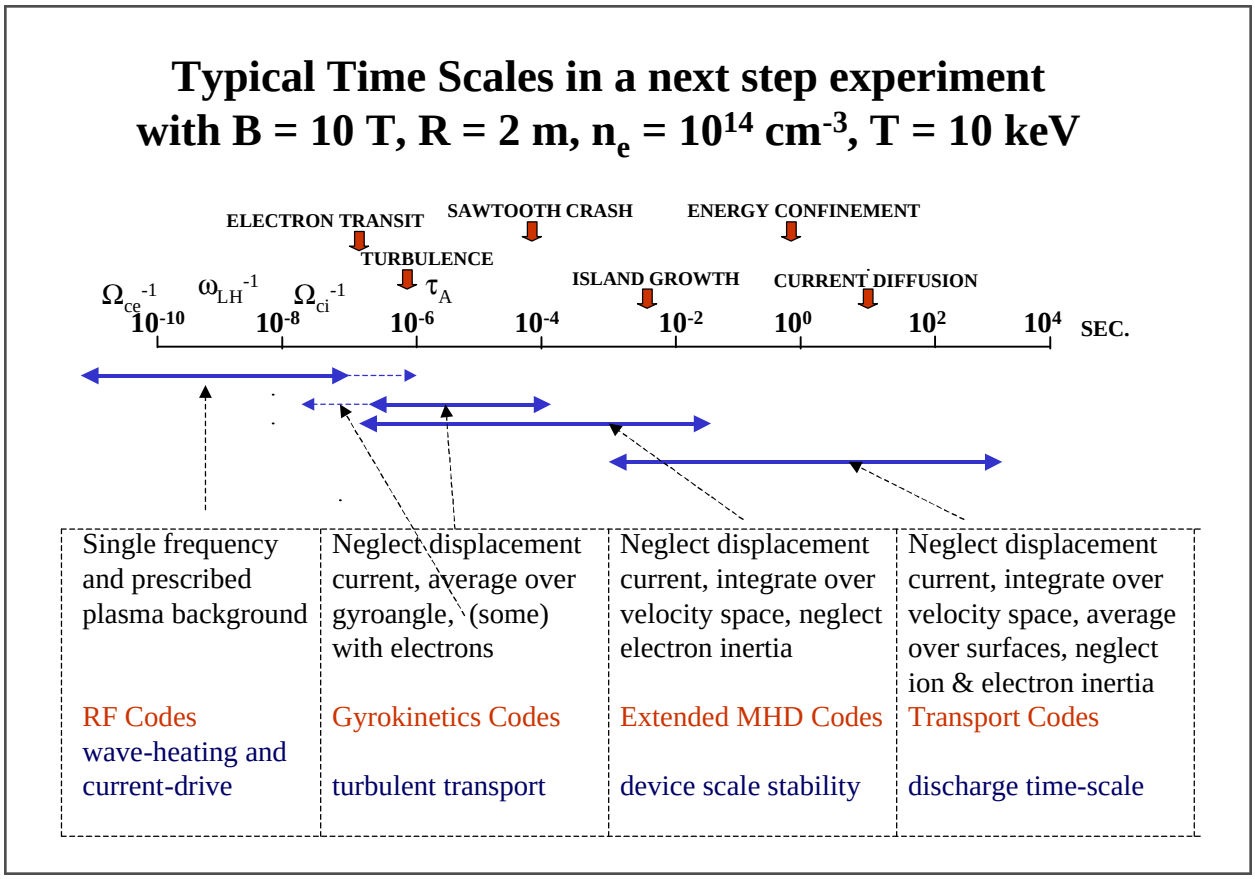

Figure III.2. *Summary of four major fusion timescales.*

A summary of the theoretically defined time scales in experimental fusion devices, and the numerical modeling presently used to investigate physics phenomena in these various regimes, is given in Fig.III.2. (See also Appendix Secs. IB and VII.) The ability to understand and predict the dynamics of high temperature fusion-relevant plasmas, in these regimes, and in the more integrated systems that will be required for further advances, is a formidable physics challenge that is central to the goals of the fusion energy sciences.

Because of the complexities, the goal of establishing a predictive simulation capability for the integrated simulation and optimization of magnetic fusion systems will require an unprecedented degree of collaboration and cooperation across many diverse areas in science and technology. For example, modern tokamaks are hot enough for the individual ions and electrons that comprise the plasma to be virtually collisionless in the direction parallel to the magnetic field, yet their ensemble can exhibit fluid-like behavior on relatively long time scales. The development of computationally tractable mathematical models that can accurately and efficiently capture simultaneously kinetic and fluid effects and describe their evolution and interaction on experimentally relevant time scales is necessary for obtaining a true predictive capability. Such an effort will require coordinated and focused advances in *fusion physics* (to further develop the underlying models and elucidate their mathematical basis), *applied mathematics* (to further develop suitable algorithms for solving the mathematical models on the appropriate computer architecture, and to define frameworks within which these algorithms may be easily assembled and tested), and *computer science* (to provide an architecture for integrated code development and use, and to provide analysis and communication tools appropriate for remote collaboration). *We emphasize that we view fusion physics, applied mathematics and computer science as fundamental to the FSP, and that healthy, focused, and sufficiently funded programs in these areas are essential to the success of the initiative.*

b. Focused Integration Initiatives

The large scale of fusion integrated simulation ultimately called for in the 2020 Fusion Plasma Simulator is unprecedented. At this time it is impossible to define precisely the technical details by which this capability will be achieved. However, we can define a program structure that will promote a variety of technical approaches to integrated simulation while retaining the desired focus on fusion science technical results. By encouraging diverse approaches to integration we will maximize the creativity of the scientific community, and we expect that one or two of these initial approaches will eventually emerge to become adopted throughout the initiative.

First, we outline the important aspects of FSP integration. As background and as noted above, two fundamental issues are common to many fusion physics integration areas: coupling of phenomena at disparate time and spatial scales, and coupling of models of different spatial dimensionality. To solve these generic problems and achieve the integration we are seeking, strong collaboration and advances in physics, applied mathematics, and computer science will be required. Seamless disciplinary collaboration will be an essential element of the program. To this end, constitutive elements of the FSP must be both large enough to encompass a critical mass of multidisciplinary researchers and also small enough to enable team environments.

Further, an intensive and continual close coupling between the calculations and fusion-relevant experiments must be a central feature of this initiative. Phenomena in magnetic fusion devices, the equations describing them, and their mutual interactions, are all sufficiently complex that developing the most effective approximations and establishing when the models have the desired accuracy can only be accomplished by continual iteration and testing of the models with experimental data. A

continual process of testing and iteration is required to advance both modeling and the characterization of experimental results.

From these objectives flow a number of requirements that the integration design must satisfy:

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

- *It must be flexible.*
  - Only the needed physics modules required for a given study should be interconnected.
  - It must be robust to changes in physics paradigms. For example a traditional diffusive transport model will be inadequate if non-local effects turn out to be essential.
  - It must be interpretive as well as predictive. That is, it must be possible to make use of both experimental information such as profiles, and predicted information such as source rates, to interpret other needed quantities such as transport coefficients.
  - It must support choice in appropriate level of description for any of the modules in a particular study. It must allow for three dimensional (3D) effects but also be capable of lower level one dimensional (1D) and two dimensional (2D) models where appropriate.

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

- *It must complement existing research.*
  - The project must provide value to the individuals involved in basic physics research, who may themselves be doing large-scale computation.
  - It must not impose significant overhead (computational or human) on the use and development of the separate physics modules. It must provide needed services so as to be of value even to the user of a single module.

Above all, the integrated capability must technically enable fusion science.

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

To achieve these goals, we believe it is necessary from the beginning to organize the project around major subsets of the whole integration problem, which pieces we term Focused Integration Initiatives (FIIs). The goal of an FII is the solution of an overarching problem in fusion physics that requires integrated simulation. The community will be invited to define the FIIs through the proposal process. As we envision it, each FII will focus on achieving a predictive modeling capability for a particular scientific problem. The FIIs will be implemented by multi-disciplinary, multi-institutional teams. The participants will be free to define a unique technical approach for each FII. Specific technical areas that should be addressed include:

Mathematical Models: Development of mathematical models to be included in the integrated simulation, including their underlying theoretical basis and ranges of physical validity.

Algorithms: Development of the appropriate algorithms for solving the equations of the mathematical models, including consistency, stability and convergence properties, the formulation and implementation of realistic boundary conditions, and performance on advanced computer architectures.

Frameworks: The definition and development of software tools specific to the physical models, mathematical models and algorithms that will enable rapid prototyping on a variety of architectures.

Performance: The development of tools to analyze and predict the performance of the models and algorithms on emerging architectures.

Verification and Validation: The definition and development of software to enable validation of integrated models with other models, and with experimental data.

Data Manipulation, Storage, and Analysis: The development of tools for the efficient storage and analysis of data produced by the integrated simulations.

Collaboration: The definition and development of tools that enable remote collaboration and project management.

Each FII should include approximately equal contributions from fusion physics, and from computational physics and computer science. The details of the management structure can be uniquely defined within each individual FII, although each FII will interface with other FIIs in the FSP by means of an overall coordinating body. It is likely that certain individual researchers will actively participate in and contribute to several FIIs.

Initially, the selected FIIs will likely pursue a variety of approaches to integrated simulation. Some approaches will work better than others, and it seems inevitable that in time a consensus will emerge that one or two architectures should be adopted throughout the FSP. At a decision point on this topic, it is expected that all project participants should be fully enabled to continue in the FSP.

c. FII Examples

In order to clarify the concept of the FIIs, we here provide some candidate examples. We emphasize that these are not to be thought of as exclusive, since the actual FIIs will be defined by the community through further planning activities as well as through the proposal and peer review processes. We consider four candidate FIIs, as shown in Fig. III.3.

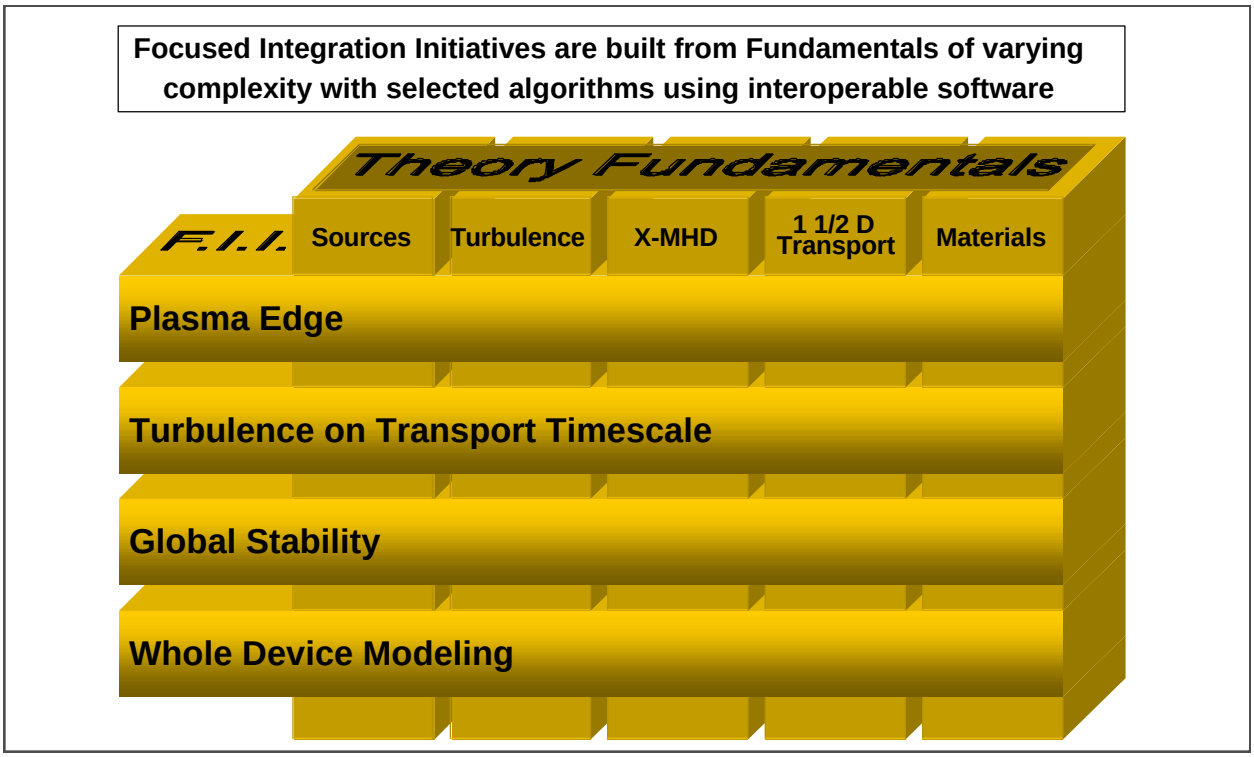

Figure III.3: *Focused Integration Initiatives cut across all the traditional fusion disciplines.*

1. FII Example: The Plasma Edge

Background: The boundary or edge-plasma of a fusion device plays a vital role in device operation. The edge plasma system extends from the top of the pedestal to a few microns inside the confinement device surface. The edge is generally considered to be the region where substantial multidimensional variations can occur in the plasma, neutral particle, and magnetic equilibrium quantities. In addition, owing to the lower plasma temperature and proximity to material surfaces, neutral gases, sputtered impurities, and atomic line-radiation can become important components. There is thus a rich variety of physics and a wealth of potential interactions that can take place in this region.

Comparative purpose: Four plasma edge elements are thought to be key to successful operation of an MFE fusion device: (1) predicting conditions and properties of the pedestal energy transport barrier just inside the magnetic separatrix; (2) understanding plasma/wall interactions for particle recycling and wall lifetime from high energy fluxes; (3) controlling tritium inventory including co-deposition; and (4) controlling wall impurity production and transport into the plasma. All of these elements are being encountered to some extent now in long pulse discharges in operational devices, and they will be encountered fully in a burning plasma experimental device in the ten year timeframe. A number of models of varying sophistication exist to describe these processes. Some models already provide a level of coupling, e.g., hydrogen transport, recycling neutrals, impurity sputtering, and impurity transport codes. However, many of the constituent models need improvement, and more inclusive couplings are required to self-consistently predict the edge-plasma behavior. See Fig.III.1b for a pictorial of a tokamak cross section, and refer to the Appendix for details regarding the edge plasma region.

Overarching theme: An early overarching issue for the edge region could be work toward a good understanding of what controls the suppression of plasma turbulence to produce a transport barrier in the pedestal region (#1 above), and its associated impact on plasma profiles and stability. The ability to predict the behavior of the edge pedestal barrier is essential for projecting the net fusion output of MFE devices. Presently, the key parameter that is believed to control the core fusion output is the plasma temperature at the top of the pedestal; this parameter is now either extrapolated from existing experiments, or assumed. Subsequent edge plasma modeling could focus on including detailed models of plasma-wall interactions, and also look to couple with core physics inside the pedestal region.

2. FII Example: Turbulence on Transport Timescales

Background: The nature of the problem to be considered in an FII on this topic can be summarized as follows: the 'anomalous' transport of mass, energy, and angular momentum in toroidal MFE devices is dominated by fluxes driven by plasma turbulence. Further, while there is a significant disparity of scales, especially timescales, this is a highly coupled system.

Comparative purpose: An objective of an FII in this area would be to bridge the range of temporal and spatial scales so as to compute the full system self-consistently, *as opposed to* just computing 3D fine-scale turbulence with fixed background profiles, or computing 1D transport with highly reduced theoretical or empirical models of the turbulent fluxes, as is often done at present.

Overarching theme: A single overarching science issue and goal is the self-consistent calculation of core temperature and density profiles from first-principles physics. An initial (easier) focus could be to determine steady-state confinement. Subsequent time evolution on the transport timescale is conceptually no more difficult but is more computationally demanding. The achievement of the steady-state goal would, as a side benefit, enable optimization studies. Important issues like simulation of both steady-state and time-dependent versions of internal transport barriers are subsets of this overall goal.

3. FII Example: Global Stability

Background*:* Global stability issues play a central role in determining the optimal operating regime of fusion devices, and in describing their time evolution. It is well known that under some operating conditions, an experimental discharge can spontaneously transform from a symmetrical stable system exhibiting good confinement into one that exhibits symmetry-breaking oscillations and poor confinement or becomes unstable and disruptively quenches.

Comparative purpose: At relatively low temperatures, global stability dynamics is well described as a resistive magnetofluid. Solutions of this model are complicated by a wide separation of space and time scales, and by the inherent high degree of anisotropy that occurs in a toroidally confined magnetized plasma. At the higher temperatures that occur in modern tokamaks, kinetic effects both parallel and perpendicular to the magnetic field introduce important physical processes that can affect the global magnetohydrodynamic (MHD) evolution of the plasma. Presently, mathematical and computational models that include some kinetic effects while retaining the computational tractability of the fluid model are collectively called *extended MHD*. While good progress has been made to date many of the approximations are adapted for the problems and resources at hand — and are not prescribed from first principles.

Overarching theme: The first-principles coupling of the transport and kinetic turbulence models to address the issues of global fusion plasma stability at all relevant temperatures and densities and inclusive of regions extending beyond the central core is a formidable problem requiring integrated modeling as envisioned by the FSP. The overarching science issue of this FII could be the predictive calculation of the onset and evolution of global symmetry-breaking events such as sawtooth oscillations, neoclassical tearing modes, disruptions, edge localized modes, and resistive wall modes, as perhaps triggered by edge plasma effects. A goal for this FII could be the development of a robust predictive capability for fusion device optimization.

4. FII Example: Whole Device Modeling

Background: The distinguishing feature of a whole device modeling FII would be that from the outset it would provide a model of the entire device for the whole discharge timescale. Through its ability to allow understanding of such coupled effects, a whole, or integrated device model should connect theory to experiment, facilitate model validation, allow offline development and exploration of new operating regimes, and amplify the knowledge which can be extracted from experimental results. Such capabilities are increasingly crucial to the development of economically attractive fusion reactors, maximizing the efficient use of experiments, and accelerating design of new or optimized devices with high-confidence validated models.

Comparative purpose: Because of the scope of whole device modeling, existing models are at present necessarily very simple. The state-of-the-art of whole-device complete-shot modeling is represented by an array of 1D transport codes, described in the Appendix. The 1D codes have many features that would be required for a final product whole device model. They employ a formal separation of time-scales between the rapid (Alfven) time on which 2D magnetic equilibria are established, and the much slower time on which heat, particles, and angular momentum, are transported as 1D surface functions across the magnetic surfaces. They also incorporate many features of a truly integrated device model (IDM): a hierarchy of models to describe particular aspects of physics, with trade-offs between speed and accuracy; connection to experimental databases; and, predictive and interpretive modes.

Overarching theme: In essence, whole device modeling is a quintessentially integrated activity. It is envisioned that simple models for all relevant aspects of a whole fusion experimental device would exist in the model, and would be capable of being replaced by more complete and accurate models as they become available and/ or as warranted by the application. Problems to which a whole device modeling capability could be applied include global validation with experiment, development of new or improved experimental diagnostics, or simulation of a proposed new machine on transport timescales. It should also be possible for a whole device modeling code to serve as the 1D transport solver throughout the development of any of the new couplings in other FIIs. From this perspective, an FII initiative in the whole device modeling area would naturally overlap with other FIIs.

d. INSIGHTS

In the past, computational modeling has contributed greatly to insights regarding the behavior of magnetically confined plasmas. We fully expect that the FSP will lead to new surprises coming from more comprehensive models and emerging from enhanced synergy between theory, experiments and modeling. As we start on the road to burning plasmas, some areas ripe for integration have been identified above as FIIs. These include edge physics, turbulence on transport time scales, and global stability, with contributions to the understanding of major and minor disruptions, plasma control, and effective rf heating mechanisms, among others.

By their very nature of enabling more comprehensive modeling, the FIIs will lead to insights not realizable previously. Regarding edge physics, overall transport and confinement are apparently determined by the height of the temperature pedestal at the plasma edge. It is expected that coupled and complex models of particle and heat transport, neutral and impurity fluxes, and edge gradient-induced MHD instabilities and turbulence in a single computational edge framework will pin down which of these mechanisms — either by itself or combined with another — regulates the pedestal height. Similarly, turbulence on transport time scales is a daunting physics and computational task. It is nevertheless deemed feasible at several levels, each exploiting separation of space and time scales appropriately. The result of high-confidence integrated modeling of turbulence and transport might be the discovery computationally of new favorable operating modes, with the ultimate outcome being the determination of transport from first principles. With respect to global stability, integration will facilitate extensions to MHD computations beyond the conventional ideal and resistive models, and may provide a way to control MHD activity that is as effective nonlinearly as it is linearly for realistic toroidal plasmas. Moreover, the inclusion of minority ion species with non-Maxwellian populations will enable extended MHD models to take on burning plasmas.

Perhaps the greatest innovation afforded by integrated modeling will be realized for burning plasma studies. It is well established that the grand challenge in the world fusion program is a burning plasma experiment. Such an experiment is a necessary predecessor to a practical power demo plant because, by its very nature, a burning plasma presents a new category of technical issues. With self-heating as the dominant plasma heating mechanism, new plasma processes and effects will arise. The high flux of energetic particles will impact the plasma and produce a rich source of wall interaction phenomena. Most importantly, all of these effects will be strongly coupled and must be understood and managed in an integrated fashion to insure the stability and success of the experiment. The ultimate Fusion Plasma Simulator will be targeted to model these processes and their consequences, thereby providing the essential insights to guide experimental programs, optimize machine design, provide information for fusion demo devices, and deepen our understanding of the science.

e. Computational Science

1. Overview of Computational Mathematics Opportunities

Fundamental to the mandate of a program in integrated simulation of fusion systems is that simulation with any subset of components becomes routine. Bringing interacting components to a state of self-consistency, and then performing experimental computational science by studying the behavior of the resulting integrated system as internal parameters or external forcings are varied, implies a multiplicity of nests of iteration over the components. In this environment, 'brute force' techniques for the individual topical analyses making up the inner loops of the integrated simulation have untenable costs in computational complexity and storage. Among the opportunities presented by the FSP are those of developing optimal discretizations and optimal solution techniques for fusion systems, and of insuring that all known techniques of potential value are propagated into the fusion context from related fields in computational physics and computational mathematics.

To appreciate the importance of optimal discretizations, namely discretizations that adapt to resolve the most physics for the memory available, or all of the required physics in the least memory, one need only consider the 'curse' of dimensionality. A doubling of resolution in one dimension of the six-dimensional phase-space for the Boltzmann equation, which is at the heart of much of fusion simulation, requires a 64-fold increase in the amount of memory, assuming that enhanced resolution is propagated in a uniform way throughout phase space. An optimal discretization will tune the

discretization locally to achieve a global error bound at minimum cost. This can be achieved via a gridding technique or by means of an approximation technique built on the grid, or (preferably) both. As another example, using an optimal iterative method for a sparse matrix solve as compared with a classical direct method is equivalent, in the cost of solving a Poisson problem on a cube with 100 degrees of freedom on a side, to replacing a 1 Mflop/s computer with a 100 Tflop/s computer — and much cheaper than the hardware-only solution even if some rewriting of data structures is required. Back-of-envelope scenarios for these and many other fusion-relevant numerical problems emphasize the infeasibility of stepping from departmental clusters to terascale systems without a concurrent research program in optimal algorithms for massive fusion simulations, and they underline the proverb: 'I would rather have today's algorithms on yesterday's computers, than *vice versa*.'

Each of the topical areas in the FSP individually present characteristics that are extremely challenging. These include the presence of multiple time scales, ranging over fourteen orders of magnitude, and multiple spatial scales, ranging over eight orders of magnitude. In many cases, the underlying physics is often coupled with essential nonlinearities, and hence reasonable simulation equation closures are the subject of research. Once closures are decided, the solutions are severely anisotropic and the computational domains are often geometrically complex, resulting among other issues in linear algebraic systems that can be sparse and ill-conditioned.

Taken in isolation, there are approaches that have been developed for most numerical challenges as they have arisen in fusion and in other application areas. Such approaches include stiff integrators to handle problems with a wide range of time scales, adaptive meshing to optimally place resolution where it is needed most, optimal order linear solvers, physics-based preconditioners, and sensitivity analysis tools. However, an integrated simulation will present *all* of these features simultaneously, as well as additional troublesome characteristics, such as nonlocal operators, inherent physical instabilities that must be resolved numerically, and which cannot be confused with potential *numerical* instabilities, high (i.e. greater than three) dimensionality for both dependent and independent variable spaces, mixed dimensional code components, and mixed continuum-particle models, based on different but physically co-located meshes.

While the combination of problem characteristics for integrated simulation of a fusion plasma presents the applied mathematics and computational science communities with possibly their greatest challenge yet, it also presents these communities with a magnificent opportunity to explore new methodologies on problems of visibility, usefulness, and external impact. The FIIs will clearly require new modes of thinking and operation for the physicists, applied mathematicians, and computer scientists involved. For example, it is unlikely that a single speciality code can provide the base to which all others should adapt, and ultimately what has worked until now may have to be completely re-thought and redesigned. Furthermore, no transformations arranged by computer science tools alone, such as a peer-to-peer software framework to couple existing codes through their inputs and outputs, will be able to provide generality of application, ease of use, and acceptable computational performance. New algorithms, especially new discretizations and new physics-adapted multilevel preconditioners, will undoubtedly be required.

Research opportunities within an FII that will be shared by the applied mathematics, computational science and fusion science communities include:

1. **Meshing:** New methods for dealing with complex geometries via unstructured and multicomponent meshes. This includes general meshing tools for tori and topologically toroidal geometries, using both fully structured and hybrid structured-unstructured meshes in the poloidal planes. Recent developments in mesh generation, including capabilities for generating hybrid and embedded-boundary Cartesian meshes, will come into play in this research.
2. **Discretization:** Advanced discretizations of differential or integral operators using high-order or solution-specific schemes. The extreme anisotropy present in many situations of relevance to fusion science implies that PDE discretizations must be designed that respect the orders of magnitude differences in transport along vs. across magnetic flux surfaces.
3. **Local refinement techniques***:* Locally refined meshing and discretization techniques that might be determined either adaptively or statically. While tokamaks have relatively fixed and well-defined geometries, the solution isosurfaces have dynamically convoluted and folding geometries.
4. **Linear, nonlinear and conservation law solver technology:** The FSP will require optimal order solvers for linear and nonlinear systems, hybrid continuum-particle solvers, fast curl-curl solvers, and stiff method-of-lines solvers (for integrating compressive Alfven waves in the poloidal field or both compressional and shear Alfven waves to follow slower dynamically relevant timescales more efficiently). Multilevel methods will need to be adapted to the afore-mentioned anisotropy. Also required are hyperbolic conservation law integrators, and nonlinearly consistent iterative methods for coupled physics with essential two-way finite amplitude nonlinearities.
5. **Transfer of field or particle data between representations:** An FII will require techniques for handling the coupling of code components by identifying natural representations that allow transfer of physical quantities. For example, in order for a PIC (particle-in-cell) code and a finite element code to interchange data requires more than unit conversion, and will be one area which cannot be accomplished efficiently and accurately without involving computer scientists, and will be required to translate a field representation from one discretization to another, possibly co-located in the same domain.
6. **Data management, interpretation and visualization***:* Interpretation of results entails multiple numerical and computer science research issues: checks for conservation and discrete satisfaction of continuous properties, visualization, advanced post-processing, and data mining.

Following success on the direct problem of multiphysics simulation in the early years of the initiative, collaborative work with applied mathematicians would pursue sensitivity analysis, stability, design and control of experiments, parameter identification, data assimilation, experimental validation, and computational steering. These ends must be considered in the early stages of software design, however, to ensure that there is a path to the ultimate goal of scientific discovery and the computational optimization of a full burning plasma device.

Achieving FSP goals involves issues that are generic to a wide range of emerging computational science problems involving other fields of physics and engineering. Solutions can likely be leveraged across FIIs and from other similar activities such as the DOE SciDAC program. Capabilities which should be expected from the collaboration include understanding a range of algorithmic and modeling options and their tradeoffs (memory versus time, interprocess communications versus redundant computations, etc.). The simulator should be able to try a range of reasonable options from different sources easily without recoding or even recompiling. Error estimates should be

automatically provided from meshing, discretization, and iterative methods, and performance feedback provided from solvers. The collaboration should eventually help code users to spend more time pushing back the limits of physics understanding, with less time spent in coding and developing mesh generation tools and solvers. On the other hand, the code users in an FII must be willing to experiment with novel algorithms and software methodologies. These collaborations must begin early in the planning phases of an FII, with agreement on achieving the research goals of all stakeholders.

2. Overview of Computer Science Issues

Any FII must provide a software methodology and framework for designing, building, maintaining, and validating the software needed for integrated simulations. A first step an FII must make is to identify the *architectures* needed, defined by the IEEE 610.2 spec as "the structure of the components, their relationships, and the principles and guidelines governing their design and evolution over time." At least three major models are possible for fully integrated architectures starting from the current set of individual topical codes:

1. **Peer-to-peer model:** Existing codes are adapted to communicate directly with each other, but otherwise operate as separate processes. This is the approach used in systems like University of Utah's SciRun, Purdue's PUNCH project, and Indiana University's XCAT. This is a distributed software components model, and maximizes the ability for individual codes to continue to be developed independently, at the cost of large file or data transfers over the network during a simulation.

2. **Single executable model:** Existing codes are subsumed as procedures in a single new executable. This might be done by starting with an existing code (e.g., transport or extended MHD) and then adding on other capabilities step by step. Another approach is to *refactor* existing codes by decomposing them into constituent parts and rebuild a new single code systematically designed from the ground up.

3. **Hybrid model:** Some existing codes are integrated together into single executables, but then they interoperate as peers with other FSP modules. The codes integrated together might be topical codes that require intimate coupling because of data exchange requirements, while the separate modules have less intensive communication requirements.

Any software architecture proposed for an FII must define clearly what the functional modules and components are, indicate how interfaces between modules are defined, and provide a work flow model for how a user will ultimately build simulations from the modules. The definition of modules should follow from the chosen intellectual and mathematical integrated framework, but should also reflect two counterbalancing forces. First, physicists need to continue the full spectrum of the fundamental physics in their individual topics areas, all of which are in a rapid state of development. This implies that some upgrade path is needed that allows the scientists involved to continue running and developing their individual codes during the development of the FSP. In the limiting case, this would argue for a full peer-to-peer model. The second force is the need (driven by limited resources) to identify shared modules usable by multiple topics codes: meshing algorithms and linear and nonlinear solvers are possible candidates identified in the applied mathematics section. In the limiting case, this would argue for a fully refactored system with single executable. More generally, the components and modules in an FII architecture need to be defined at multiple

levels of granularity. At the highest level, general functional modules should be identified, which might consist of (modified) existing physics codes for transport, MHD, sources, etc. At a finer grained level, the functions/routines from which to build codes should be identified, e.g., a toolbox of solvers, meshers, discretizers, and data converters.

An FII will also have a *data architecture*, the model of all data needed to support the research: the types of data and data objects, how they are described and defined, and their relationships. Verification and validation are critical components for the FSP framework, and supporting access to experiments, experimental data and diagnostics implies the need for well-defined data systems. Metadata ("data about data") mechanisms will be needed; an integrated simulation may span multiple geographically distributed machines as well as multiple codes, and tracking the results of a simulation will need data systems which can locate all related outputs and allow multiple users to attach annotations to the data. Access policies and mechanisms will also need to be defined: which users get access to which data; who has write versus read permissions; and, what if any security protocols are required to protect the integrity of the data.

An architecture is implemented as a *computational framework*. The computational framework includes how modules are linked together, the 'run-time system' which provides communications and control between modules, and lifecycle control (starting, stopping, killing parts of the computation). A computational framework might take the form of a *problem-solving environment* (PSE), which is the full set of utilities and tools needed to set and solve a range of problems from a particular domain. A PSE often includes a graphical user interface, a way of describing problems in the language natural to the problem domain, and specialized post-run analyses that hide complexity from the user.

3. Computer systems issues an FII must address

The central goal for an FII is the complete integrated simulation of an overarching fusion physics problem, with the eventual goals of predictive simulation of a burning plasma and parameter optimization that can lead to more efficient magnetic confinement devices. Accomplishing these goals will require addressing several important but straightforward computer science issues. From a software engineering point of view, an excellent proposal will include the following items.

Requirements analysis: A process is needed to identify the components and capabilities necessary to accomplish both the short and long term goals. Both envisioned scenarios and more formal 'use cases' could be helpful in deriving the requirements. A project like the Fusion Plasma Simulator will have evolving requirements over its lifespan, and the process used for changing requirements needs to be specified.

Sample requirements could include:

1. **Computer languages:** Will the computational framework be required to support multiple computer languages, or will all components be required to have an interface in a single language?
2. **Code ownership**: Will the components be "owned" by their creators, or be community-owned, or be required to be open-source? Will commercial software be used, and if so what licensing will be necessary to assure long-term viability of the proposed FII?
3. **Platform dependence:** Will components be required to run on particular operating systems and hardware platforms? How are those chosen and what support will be needed for porting and testing between platforms if more than one is chosen?

4. **Performance requirements:** An integrated simulation system will not run faster than the slowest of its components. What are the performance requirements, and how will they be expressed?

Survey of existing systems: Many frameworks, systems, and architectures are currently under development and being used for high-performance scientific computing. Examples include CACTUS for computational astrophysics, the DOE's Common Component Architecture, Argonne's PETSc libraries of linear and nonlinear solvers, the National Transport Code Consortium, the University of Utah's SciRun framework, the Community Climate Modeling System. Identification of what can be utilized from these and other similar projects, ranging from design to codes, and what if any deficiencies they have for the FSP, will enable leveraging these existing code bases.

Basic software maintenance: An FII will likely span multiple laboratories, developers, and geographically distributed sites. The code development as well as the final product will be shared, so formal systems for software development and maintenance are required. Some version control system like RCS, CVS, BitKeeper, or SCCS can help in coordinating distributed development, and keeping archival versions of previous releases. Some formal bug tracking tools should be used since the software is likely to be under rapid parallel development by separated code teams. A framework for unit and regression testing will allow automated testing and notification of stakeholders in the FSP of potential problems from updates. Software maintenance is a critical infrastructure for successful development and deployment.

Development path: A migration plan must be provided that indicates how to move from the current standalone codes in different topical areas to an integrated physics framework. This path needs to reflect the requirements of participating code users to continue producing research with their codes during the development of an FII framework. Software engineering research has shown that there is typically a 50% higher cost to develop components to the high quality standards needed for re-use and sharing; however, once a core of usable and useful components is available they can raise expectations and draw in other developers.

Flexibility and extensibility: Each FII should have a plan for tracking and using software utilities and components developed by other FIIs. Eventual interoperability and shorter-term shared module development need to be identified and exploited whenever possible. The problem domain that the FSP framework handles must be explicitly stated.

Data models: A description of the data that needs to be shared or communicated between components of the architecture at runtime must be provided. How is the data described programmatically (e.g., using HDF5 descriptors or XML schema), and how is the data model extended to unforeseen future data interactions? In addition to static information about data objects and how they are defined, each FII needs to provide estimates of how often interacting components need to exchange data, the sizes of the data objects in those interchanges, and what if any data mediators (for interpolation, unit conversion, coordinate transformation, etc.) are required. Related to this is a networking requirements analysis, describing what must be transferred over local or wide area networks, the network capability of the participating sites, and what parts of an integrated simulation will require special high-speed connections, quality of service guarantees, or special security protocols.

While the items above are basic requirements for any FII proposal, additional desiderata might include:

- The ability to work hierarchically with the components of the architecture: an expert in extended MHD should be supported by the framework in assembling a MHD solver with variant capabilities or to explore new methods. At the same time, an expert in RF sources should be able to use a framework-provided 'default' MHD component without becoming an MHD expert.

- Rapid prototyping capabilities, or the ability to run selected components in a lower-fidelity mode for quick tests and proof of concept simulations: this also refers to the ability to quickly compose, compile, and launch 'what-if' scenarios using the framework.

- Collaboration support, such as human interactions via videoconferencing, shared code development and distribution tools, and a shared testbed of hardware and software used by everyone on the same FII, or across multiple FIIs: this underlies the concept of virtual FII centers of research.

- Data analysis tools, that can be used across physics regimes, mesh and discretization techniques (PIC, AMR, finite elements), and disparate codes: this includes visualization but may also include statistical summaries, consistency checkers, etc.

These additional considerations are more generally characteristic features of problem-solving environments.

While an overall governance structure will be mandated for all of the FIIs within the FSP, as described elsewhere in this report, a successful proposal should also be required to address some local governance issues related to the computational science infrastructure proposed. Software version control, bug tracking, and code configuration and maintenance methodologies are of little help unless all the stakeholders use the proposed system. Users will be required to follow some standards, but if too onerous they will be ignored. Particularly in situations where stakeholders are geographically dispersed, a priori agreements need to be worked out on timely responses to issues reported, and the level of support that individual component suppliers are expected to provide.

f. Verification and Validation Requirements

Since the goal of the FSP is to build models capable of accurate prediction, it must be in a position to assess the reliability or accuracy of these models at all phases throughout the project. Assessment of predictive models has been divided into two distinct activities: *verification*, which assesses the degree to which a code correctly implements the chosen physical model and *validation*, which assesses the degree to which a code describes the real world. The former is essentially a mathematical problem (in a broad sense) while the latter is essentially a physical problem. *Overall, the goal of verification and validation is an assessment of the extent to which a simulation represents true system behavior sufficiently to be useful.*

As documented in IPPA 2002 and the 2002 Snowmass Fusion Summer Study documents, predictive capability based on scientific understanding is a key goal of the fusion energy sciences program. The accuracy of our predictions, when mapped to the fusion energy mission, has significant economic consequences. Reactor scale devices are expensive, and uncertainty in the underlying science requires extra margin in their design. Formal verification and validation regimes have been defined and applied to 'high consequence' applications like national defense, environmental protection and nuclear power and in some cases linked to the regulatory schemes for the systems in question. While

we can learn much from work of this rigor, we need to introduce a validation and verification governance regime appropriate to the goals and scope of the FSP. Cost/benefit/risk tradeoffs will need to be made as the project management allocates manpower and other resources. It is also important to recognize that verification and validation is an iterative process carried out over the life of a project, not a one-time test.

1. Verification

The verification process attempts to identify and quantify errors in the computational model and its solution. As such it must logically precede validation. Sources of error include algorithms, numerics, spatial or temporal gridding, coding errors, language or compiler bugs, convergence difficulties and so forth.

The most powerful tool for verification of an individual model is comparison with analytic solutions to the same conceptual (physical) model. This is not always easy since analytic solutions are often only possible in very simple regimes. Comparison between codes is also useful, pointing out the importance of maintaining diversity and breadth in the code library. Codes that use radically different approaches to their solutions provide the most thorough tests. Internal checks for consistency and convergence are, of course also essential, by changing gridding, timesteps, and sometimes even solution algorithms.

For an integrated suite of physics models, verification can present more of a challenge, since there is typically no analytic solution available and there may not be other existing coupled computational models. But there are various options, depending on the type of coupling. First, when available, is comparison with other coupled computational models (either completely independent code efforts, or multiple approaches to coupling implemented within a given code). Also, in some cases a coupled approach can be compared with a 'brute-force' direct simulation of the same physics. For example, a coupling of turbulence and transport can be benchmarked against a background-evolving turbulence code for test problems where the timescales are not too disparate, or a coupling of one-dimensional core transport and 2- or 3D edge transport can be compared with an edge transport calculation that extends all the way into the core. Or, a code that integrates different kinds of physics into a single set of equations can be operated in limits where one kind of physics is expected to dominate and then compare with an existing code that calculates the dominant physics.

In all of these approaches it is helpful to operate within a computational problem solving environment that facilitates side-by-side execution and comparison of multiple computational approaches.

2. Validation

A successful validation regime must begin with planning. An FII must clearly define the goals of the predictive code - what is driving the need for the calculation? Since validation is not a mathematical process, it is only really meaningful in the context of a well-defined application, such as an overarching area of assessment as defined by an FII. The validation regime attempts to assess quantitatively the ability of a code to predict and to define the boundary between acceptable and unacceptable extrapolations; this defines applications for which the code can be trusted. Next, the FII participants must identify the critical issues, design real and numerical experiments, specify metrics and define assessment criteria. Planning is the place where resources are balanced against other elements of the project. Since validation necessarily involves experimental groups who are

outside the simulation project and funded independently, contacts should begin at the onset of FII planning. Critical diagnostics may need to be developed and deployed as part of the overall program, and synthetic diagnostics developed and employed in the simulations. New diagnostic techniques can also be discovered by this process.

The principal validation activity is the design, execution, and analysis of dedicated experiments, both real and numerical. Comparison with historical data from existing archives is useful but almost certainly not sufficient. The crucial comparisons are those designed to test important features of the model and evaluate critically and quantitatively. The experiments must be designed to collect essential data for comparison, particularly initial and boundary conditions needed by the code. Experiments should challenge the codes in fundamental ways and explicitly test the model's assumptions. A hierarchy of measurements and comparisons of increasing difficulty should be established, for example progressing from global to local variables and from steady state to transient conditions.

Experimental design needs to be collaborative between the code project and experimental team. Typically it will require use of simulation as part of that design. In this manner the experiments can be optimally useful and can stress critical parameters and measurements. Use of the codes at this stage help build the collaborative environment, tools and working methods that will be necessary during the measurement and analysis phases. The groups must form a team for the purposes of validation. The team should not have the goal of proving the code is correct, but dispassionately evaluating its status. At the same time, the experimental team must be frank and forthcoming about limitations and errors in the experimental data. The availability and quality of data is a critical need for the validation program and raises the very large issues of error analysis and experimental data validation. Typically evaluation of random errors is straightforward while estimation of systematic errors is not. Often the latter is no more than a 'seat of pants' estimate. Although experiments must be developed collaboratively, independence should be maintained in data collection and analysis.

Each FII, and ultimately the FSP as a whole, must define metrics by which the comparison between the code and experiments are to be evaluated. The metrics need to take account of all sources of error: experimental, both random and systematic; and, assumptions and approximations in the model as well as convergence or numerical errors. Using these metrics, the FII team then must make an assessment of its status. The assessment is essentially a statement of confidence in the code in a particular area and confidence in the ability of the code to extrapolate or predict. (The importance of the latter suggests greater weight be given to predictive tests rather than postdictive comparisons.) Standard statistical analyses for hypothesis testing may be appropriate for this task.

Finally, the process and results need to be well documented. This should include a description of the experiment, the full set of experimental data and metadata, the assumptions, parameters, inputs and outputs from the code, a description of the analysis procedures and error estimates, along with the metrics and assessment. This should be kept as part of the documentation of the FII. Specific validation requirements in topical areas are given in the Appendix.

g. INFRASTRUCTURE REQUIREMENTS

We have discussed the FII concepts, the architecture issues, and verification and validation of the FII developing capabilities. All of this hinges on resources available to produce results, i.e., the computational capabilities available to the FSP. This section highlights some of the salient features of this fundamental project need that is external to the resources of the FSP.

1. Computational Resources

To realize the benefits of the developing simulation capability will require three rather distinct types of computing resources:

1. interactive or rapid turnaround resources for short to moderate times, at all ranges of relevant memory (e.g., including largest numbers of processors for short times) for purposes of code development and debugging, testing of physics formulations, code components, and numerical methods;

2. substantial computer resources for very long periods of time for production runs and parameter surveys. In this case a figure of merit is the number of usable flops available over the course of a year, not the maximum achievable flops, and need not necessarily be on the largest, fastest machines; and,

3. the largest memory, fastest processor machines to allow exploration of extremely challenging physics regimes having high resolution requirements, large Reynolds number, high dimensionality and the like, to push a verified, validated computational capability into a regime that is wholly new.

Each of the fundamental areas of fusion theory are now pushing the limits of computation of the types listed as 2) and 3) above. For example, the key challenge in performing Extended MHD computations relevant to the hot plasmas of modern fusion experiments is to increase the dimensionless parameter characterizing inverse plasma collisionality, the Lundquist number, S. Present Extended MHD calculations have achieved 18 Gflop/sec (GF) on 384 processors of an IBM SP3. This performance limits both the accessible Lundquist number (~10^7) and the problem time (~1 msec). These values are several orders of magnitude less than are required to accurately simulate present fusion experiments. It is estimated that a 1000-fold increase to 20 Tflop/sec (TF) sustained performance could allow values of S approaching 10^9 and the problem time to approach a tenth of a second or more, enabling validation of the mathematical models and comparison with present experiments. Further extensions into the 100s of TF regime would likely be needed to treat some key problems for burning plasma devices.

At the present time roughly 10^{-3} s of a turbulent discharge can be modeled at minimal spatial grid resolution with 120 hours on 128 processors on the NERSC IBM SP (115 MF sustained performance per processor $\rightarrow$ 1.5 TF hours). This time needs to be increased by a factor of 10-100 to address transport time scales. Furthermore additional physics associated with kinetic electrons (which necessitates an increase in computing resources ~50 to 100) and electromagnetic coupling must be included in the models in order to allow a quantitative understanding. Codes in the SciDAC Plasma Microturbulence Project have recently added this physics capability, but presently there are insufficient computer resources to carry out the scientific studies. It is estimated that 10^3 to 10^4 TF hours are required to include kinetic electron dynamics for transport time scales.

A gyrokinetic edge code simulation would require a capability in the 20 TF (sustained performance) range to simulate up to nominal background relaxation timescales. Full-shot simulation, or simulations requiring coupling to the largest spatial scales, e.g. for edge-localized modes, would require at least an additional order of magnitude.

Codes which solve the full, hot plasma wave equation in 2D and 3D, to all orders in Larmor radius divided by scale length and including all cyclotron harmonics have been developed under the SciDAC program to study wave heating, current drive and plasma flow drive. These codes scale well and have achieved efficiencies ~40% relative to theoretical maximum using 1600 processors on the NERSC Seaborg machine. High resolution solutions in 2D to study fast wave mode conversion typically require a few TF hours per toroidal mode calculated. A full antenna spectrum of ~50 toroidal modes then would require ~100 TF hours. A low-resolution solution in 3D for fast wave propagation required ~ 30 TF hours.

As we seek to integrate the disparate plasma models, it is realistic to expect that the level of fundamental physics detail that can be incorporated will be dictated by the available computer resources. Two to three orders of magnitude increase in effective computing may be required to achieve the program goals. This increase can come from several sources: more, and more problem efficient, computer hardware; improved physics analysis resulting in more accurate reduced models; and, improved algorithmic and mathematical methods.

2. Network and Storage

Requirements for network connectivity and mass storage are driven by the vast quantities of data that will be produced by the proposed simulations and by the geographical distribution of participants and computational resources. Precise quantitative predictions are difficult since the frequency with which simulations will be performed, the amount of data generated, the amount of that data that will need to be stored or transferred, how quickly after a simulation data will be needed and how many sites/people will use the data are all uncertain at this point. In part, these depend on the FIIs chosen, and in part on other issues such as researcher proximity. Still, even under the most prosaic imagined situations, rough estimates yield numbers which are large enough to warrant serious attention. Over the next five years, three dimensional, non-linear MHD and turbulence codes may be generating on the order 1 PByte per simulation. We may further hypothesize that, integrated over the entire project, it would be desirable for major simulations, each representing a full, integrated experimental "shot" to be completed on the order of once a week. (This also is practically possible. Consider: 1000x1000x100 spatial zones, 10 variables per zone, 10^6 time steps, and assuming 100 floating-point operations per variable per space-time point, such a calculation would take no more than a few percent of 50 TFlops machine to complete in a week.) By this estimate, aggregate rates in the 10's of PByte per year should be planned for. It is not sufficient to simply archive this data; scientific progress will be linked to our ability make effective use of it.

Meeting these challenges will require a highly capable network, massive storage infrastructures, along with advanced middleware and network services to "glue" them together. All of these elements are required resources for access by the FIIs. Simply moving the estimated simulation data once would require dedicated links at several Gbps. Distributing the data to multiple sites over a wide area network, while technically feasible in the project's time frame, is probably not an economically reasonable approach. Instead, the researchers should plan on moving only that part of the data necessary for visualization or post-run analysis. Assuming a data set for local analysis or visualization is 10-20 GBytes (a reasonable guess for the RAM capacity of a workstation in the next five years), and that 10 or 20 seconds is a reasonable waiting time, one calculates a requirement for burst transfer rates

up to 1 Gbps. Alternately, post-processing or visualization engines could be co-located with the data store and the results streamed to end users. An HDTV stream requires about 50 Mbps, a large display wall, fed with uncompressed data could use 1 Gbps. Thus the network requirement may be estimated conservatively as one that supports a large number of users (perhaps 50-60) located at most of U.S. fusion sites each transferring bursts of 1 Gbps at duty cycles of a few percent.

Data storage in a reliable and robust repository (or repositories) is another formidable challenge. In analogy with experimental data, results generated for major simulations will be of archival quality. The estimates shown above, suggest that data will accumulate at 10's of PBytes per year. As with the network requirements, these rates are not technically insurmountable but do require significant thought as every decision concerning archives of this magnitude will have a serious impact on the ability of the team to carry out the scientific program and on the economics of the project. Architectural decisions include where to store the data, whether to centralize or distribute the archive, whether to support data replication at remote sites, how to integrate post-simulation analysis, how to integrate with experimental archives, and how to guarantee data integrity and consistency. With such a large quantity of data, automated mechanisms for constructing data digests, databases or summaries and new and improved ways to efficiently mine the repositories to extract knowledge will need to be developed.

With distributed resources for software development, computing, storage, analysis and visualization, the project will require advanced network middleware that supports distributed computing and collaboration. At the same time the apparently conflicting requirements for transparency and security in a widely distributed environment point up the need for efficient and effective network services. Central management of Public Key Infrastructure (PKI) or equivalent technologies using 'best practices' and providing around the clock support is essential. It is equally essential that the user authentication framework adopted is such that common policy can be negotiated among the collaborating sites. Mutually agreed upon tools and protocols for resource authorization is also important. For such a large user base and with the need for close collaboration with experimental groups on validation tasks, global directory and naming services may be a key technology and may help to anchor the wealth of distributed metadata. A hierarchical infrastructure with well-managed 'roots' can provide the necessary glue for many collaborative activities. A global name service could also solve the longstanding problem for our field of variable name translation between codes or experiments. Since users, including partners at universities, private companies and international sites are interested only in end-to-end performance, real-time network performance monitoring and problem resolution tools which work across administrative domains will be essential. Finally, a host of collaboration services including teleconferencing, distributed applications and remote visualization will be required.

3. Summary of infrastructure requirements

High-end infrastructure requirements are summarized in the table below. In addition, the project will need local infrastructure consisting of medium scale computing clusters for development, testing, and post-processing; powerful desktop and visualization systems; medium scale storage systems and well provisioned local area networks.

	NOW	SOON (1-3 years)	LATER (> 3 years)
Computers	~4Tflop/s ~2TByte memory	10's Tflop/s TBytes memory	100's Tflop/s 10's TByte memory
Storage	~50TByte/year	~1 PByte/year	10's PByte/year
Networks	0.1 Gbps to desktop 0.62 Gbps backbone	1 Gbps to the desktop 10 Gbps backbone	>1 Gbps to the desktop 100 Gbps backbone

Table III.1. *High-end fusion simulation infrastructure requirements.*

IV. THE FSP PATH

a. PROJECT ROADMAP

The goals of the FSP both near, five-years, and longer term, ten-years and fifteen-years, are ambitious. To meet them careful but flexible planning coupled with a well-designed program governance and program management structure is required. To put governance and management requirements in context, Fig.IV.1 below provides a summary roadmap of the FSP goals.

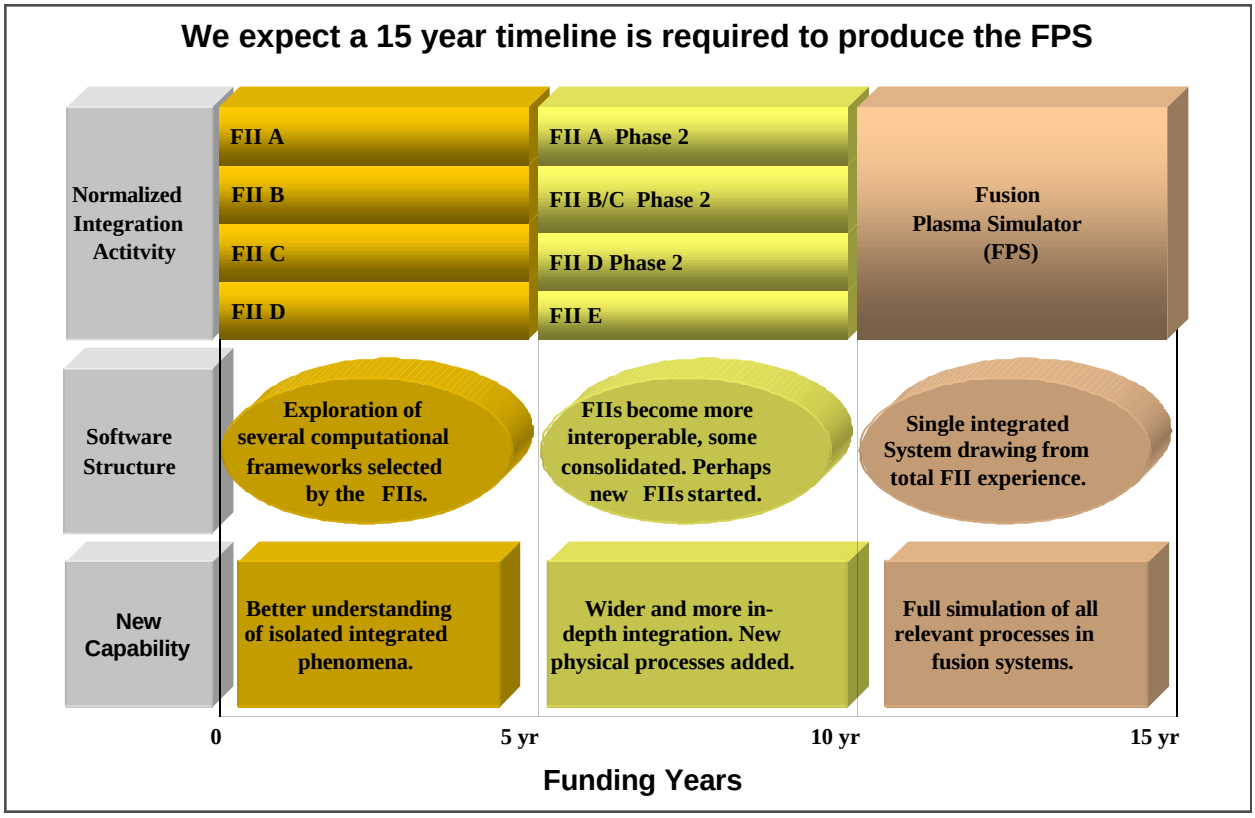

Figure IV.1: *Fifteen-year roadmap for the Fusion Simulation Project*

The plan shows a fifteen year timeline, with significant value and specific milestones at the end of five and ten years. As noted in the 2002 Snowmass Fusion Summer Study, the full extent of the Fusion Plasma Simulator project is expected to require funding on the order of $0.4B throughout the fifteen year period.

The three major phases of the project are described in the following paragraphs.

First five years:
During the first five years, we will initiate several Focused Integration Initiatives (FIIs) as described in Section III. These will be concentrated on specific high-profile physics integration issues that are considered to be the most critical, and are also prototypical of the integration issues faced by the whole initiative. Each FII will be quasi-independent of the others, but there will be efforts made at

coordination between initiatives, looking towards future integrations. We expect at the end of the five-year period to have substantial new capability in each of the FIIs and to gain new scientific insights during this time frame.

Each of the FIIs will be expected to develop a computational framework best fitted to its task. Since there is not universal agreement on the best way to solve many of these integration problems, the computational framework also needs to allow rapid prototyping of different solution techniques over some range of hardware architectures.

New computational algorithms will be developed to treat the unique mathematical problems present in fusion science. For example, many of these arise from the presence of the strong magnetic field, which adds an extreme anisotropy to the plasma and results in temporal and spatial anisotropy of particle motion. Also, the mathematical equations describing plasma waves are higher order parallel to the field than across it. These physical effects lead to sparse matrices with peculiar properties, to the need for very specialized gridding techniques, and to the need to deal with "stiff" equations on disparate time scales.

5-10 years:
During the second five-year period, we expect several things will occur. One is that we will begin to combine select FIIs into larger and more comprehensive integration activities. Another is that we will introduce new FIIs as required. A third is that we will take select FIIs to the next level of development.

During this period, there will also be a comparative reassessment of the issue of computational frameworks. We expect the frameworks to grow in maturity, integrating such things as advanced graphics and user interfaces, and also that some down-selection and solidification will occur. A unified system for effectively managing the increasing complexity of the project will become a priority. We also envision that improved algorithms will enter the project as the nature of the couplings of the new integration phenomena becomes clearer. For example, some new algorithms may achieve greater efficiency by combining individual components rather than by treating each component as a block in a high-level algorithmic diagram.

10-15 years:
During the final five-year period, the focus will be on comprehensive integration. There will be a link between all the physics components of the project. The mathematical framework will largely be in place for the integration.

Part of the challenge of this final integration will be to include multiple levels of description of the same phenomena. For example, there would be an option for plasma equilibrium to be computed either in the 2D axisymmetric approximation or fully in 3D, including the effects of small magnetic islands. Plasma rotation could be included in either of these calculations or neglected. Each level represents a tradeoff between computational efficiency and physical fidelity.

As the integration proceeds into this final phase, we expect the simulation capabilities to be exercised more, and to have new and more comprehensive comparisons between the simulation and experiment. In many cases this will lead to a validation of the model, but we expect that in some cases it will serve to identify shortcomings or inadequacies of the model that will be subsequently addressed. This final integration will succeed only if all the fundamental components have been adequately addressed and if the component integration is carried out correctly.

We expect that the Fusion Plasma Simulator produced by this fifteen year project will be a living software system that will continue to grow and be modified many years into the future. It will serve an invaluable role as an intellectual integrator of many experimental results and approaches, and will be heavily relied upon to reach decisions regarding the development path of fusion energy.

b. The FSP and other Office of Science Activities

We have stressed throughout this report that the FSP will reach across disciplinary boundaries in order to bring together all relevant expertise required to develop an integrated simulation capability for magnetic fusion systems. This expertise is resident in several units within the DOE Office of Science. Primary among these are the activities within the Office of Fusion Energy Sciences and the Office of Advanced Scientific Computing. Other activities are and will continue to be directly relevant to the success of the FSP. Within the DOE, these include the recently developed SciDAC initiative and materials sciences research within the Office of Basic Energy Sciences.

The SciDAC (Scientific Discovery through Advanced Computation) is one of the new, and critical strategic programs within the Office of Science. One of the SciDAC program's principal goals has been to assemble interdisciplinary teams and collaboratories to develop the necessary state-of-the-art mathematical algorithms and software, supported by appropriate hardware and middleware infrastructure, to use terascale computers effectively to advance fundamental research in science that is central to the DOE mission. Substantial success has been achieved towards this goal. The SciDAC success provides an argument for the timeliness of the FSP. Perhaps no area of science is more central to the SciDAC mission than fusion, and five projects were launched under SciDAC auspices in FY 2001 to develop and improve the physics models needed for integrated simulations of plasma systems to advance fusion energy sciences.

One of these projects, the Center for Extended Magnetohydrodynamic Modeling (CEMM), has been able to speed up an Extended MHD modeling code through synergistic interactions with applied mathematicians on another SciDAC team (Terascale Optimal PDE Simulations). The resulting algorithmic improvements have already decreased running times by a factor of two, and further exploitation of certain matrix structures will yield even more improvement.

Another SciDAC fusion project coordinates a multi-institution program on 'Numerical Computation of Wave-Plasma Interactions in Multi-Dimensional Systems.' An applied mathematician on this team was able to recognize and exploit Kronecker product structure in some of the equations underlying simulations within this project. The introduction of this and other such insights into the fusion sciences context has led to codes that are now running two to ten times as fast as previous versions. More detailed physics can now be introduced earlier in the modeling and simulation process, thereby greatly accelerating the pace and scope of the science that can be explored.

A third SciDAC has given rise to the U.S. Fusion Grid (http://www.fusiongrid.org), which is now being tested on DIII-D and C-Mod scientific data analysis. Developed under the auspices of the National Fusion Collaboratory SciDAC project, the Fusion Grid presently combines experimental data that is stored on servers at MIT, GA, and PPPL with a computational code located at PPPL to provide greatly improved data analysis throughput rates combined with instant access to the latest versions of a PPPL numerical code. The Fusion Grid is thus beginning to provide a collaborative fusion research environment that is transcending geography.

These and other SciDAC projects, even though currently funded at only a modest level, can be thought of as part of a pilot program. As such, SciDAC will ultimately have succeeded best if it spawns major application-specific initiatives precisely like the one being proposed in this document. Early SciDAC success stories provide compelling proof-of-concept evidence to strongly suggest that appropriately funded interdisciplinary teams, focused on a full-scale integrated program, will successfully deliver a greatly enhanced simulation capability to fusion energy sciences. Such a capability is absolutely essential for realizing our nation's goal of commercially viable fusion power in a realistic timeframe.

Further, there are significant needs in materials modeling of fusion device hardware which must be met to provide a complete FSP predictive capability. This is work that would be carried out in collaboration with the Office of Basic Energy Sciences (OBES). It would range from basic theory of fundamental material processes on the atomic, mesoscopic and continuum levels to the simulation of complex surface and bulk phenomena. Surface processes of interest include sputtering and other erosion mechanisms, implantation, re-deposition and co-deposition of tritium and surface restructuring, and roughening among many surface problems of interest. Bulk processes including crystal lattice displacement damage, the creation of atomic and cluster defects, microstructure evolution, dimensional instabilities, and a variety of embrittlement processes, will also need attention. It is significant that the basic approach to multi-scale modeling of materials, e.g. the passage from atomistic simulations to mesoscopic simulations to continuum simulations, is consistent with and complementary to the multi-scale, multi-year paradigms for the FSP. It would be beneficial for a mutual working relationship to develop between the FSP and OBES to complement those between OFES and OASCR.

We anticipate that the relationship between the FSP and other Office of Science activities will continue to evolve and mature during the life of the FSP. Experience has shown that many advances in basic science have been achieved in the pursuit of goal driven activities such as the FSP. We strongly believe that this pattern will emerge again in the context of this project, thus benefiting both the goals of enhanced energy production and the advancement of fundamental science.

c. Roadmap Implementation: Project Governance & Management

The FSP initiative will be focused, highly interdisciplinary, and will involve a significant number of people. For these reasons it is extremely important that careful attention be given to governance of the project. The governance structure needs to effectively balance the professional requirements of the creative and individualistic people who will carry out the work with the programmatic needs for focus and timely delivery of results. In addition the structure has to work effectively with the two DOE programs offices, OASCR and OFES, that will support and manage the initiative.

There will be two elements of guidance and oversight needed to reach the technically complex goals indicated in the roadmap, Fig. IV-1, on any of the indicated time scales. The first element, project governance, is the process of coming to the best possible technical judgments when evaluating options to reach project goals. This would include agreements about software architectures, selections of emphasis for physics fundamentals, the down selection of FIIs, and a multitude of other issues of this sort. In addition, there are issues of project management: actual implementation of the broad technical decisions across the FSP, e.g. software standards, tracking of progress, issues of accountability, organization of project reviews, assisting in the representation of the project, and etc.

In the case where several institutions or groups have co-equal technical shareholding status in the FSP, e.g. if the 3-5 multi-institutional FIIs are enfranchised as recommended in this report, the above-drawn distinction has important functional implications. For this circumstance, a sketch of a proposed governance structure is provided below. An analogous set of issues has been addressed by the Community Climate Systems Model (CSSM) activity; see http://www.ccsm.ucar.edu . While there are significant differences between the nature of the science involved in the CCSM and the initiative discussed here, there nonetheless are sufficient similarities that the CCSM activity can help suggest an optimal structure. The organizational chart suggested is:

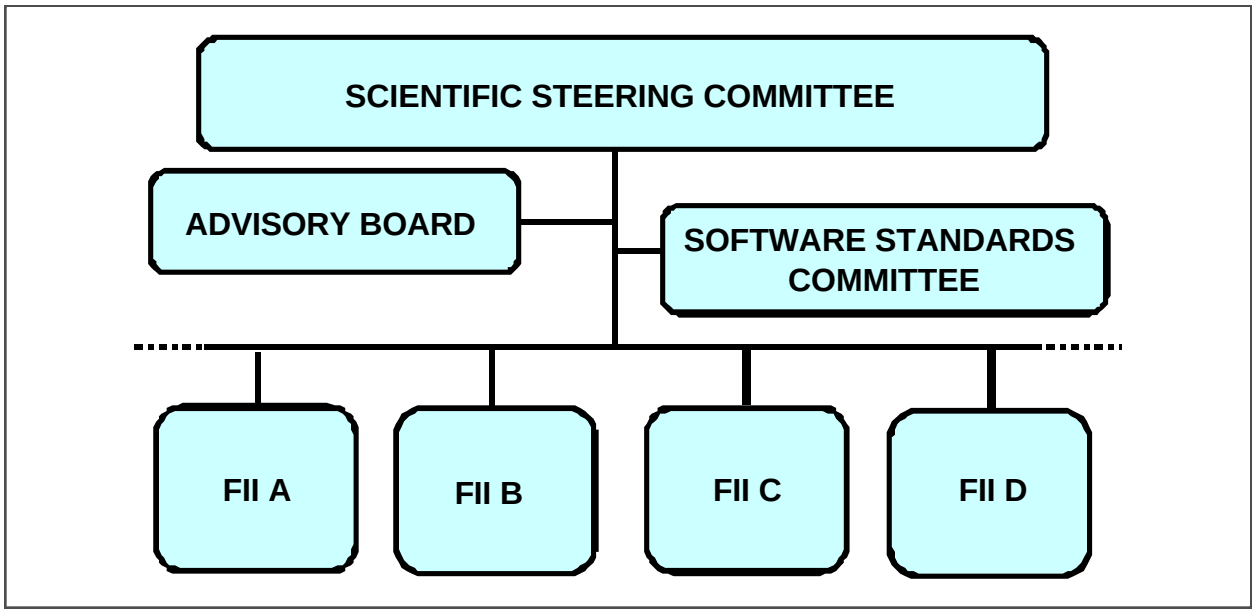

Figure IV.2. *Organizational Chart for the Fusion Simulation Project.*

The functions of these organizational groups are:

Scientific Steering Group (SSG): This group provides the overall scientific direction and vision for the project. It provides oversight and coordination of scientific activities. It is the key group for assuring that integration is effected. A primary function of this group is outreach at the technical level. The SSG ensures the verification and validation function of the FSP. The SSG will also need to work closely with the program management on the topic of resource allocation issues.

Advisory Board: This group is made up of people with scientific breadth that are not directly engaged in the FSP. The group will provide scientific and management advice to both the SSG and program management. A fundamental role of the Advisory Board is to address resource adequacy and FSP collaboration throughout the Office of Science.

Software Standards Committee (SSC): This committee is comprised of representatives from each of the FIIs. It is critical that some level of standards and common practice be made across the FSP with respect to software and collaborative tools. The SSC will work to assure the maximum realistic uniformity in software choices throughout the project. The SSC ensures that each FII has a plan for tracking and using software utilities and components developed by other FIIs.

Focused Integration Initiatives (FIIs): Each FII will have the responsibility of carrying out the overarching research plan to which it is committed. Each FII group oversees the scientific direction for its integrated simulation, including determination of required fundamental research, and coupling with experiment. This is where the real work gets done, in fusion science, computational science, and in the cross-disciplinary activities that involve verification and validation.

V. SUMMARY

The goal of the Fusion Simulation Project (FSP) is to develop the computational capacity to perform integrated simulations of toroidal magnetic confinement devices and provide a validated predictive capability. The panel envisions a program proceeding through three five-year phases, the first of which is detailed in this report and would be comprised of focused integration initiatives (FIIs). This development will be made feasible by close coupling of the integration initiative research with ongoing core program activities in theory, experiment, computer science and applied math carried out under the auspices of DOE OFES and OASCR. Our vision for the initiative is detailed in Section III of this report, with the path discussed in Section IV, and a fifteen-year overview roadmap delineated in Fig. IV.1.

Numerical modeling has played a vital role in magnetic fusion for most of its history, with increases in the scope and reliability of simulation enabled by advances in hardware and numerics and through improvements in basic theory. Knowledge gained by this approach has covered the entire range of problems in the fusion energy sciences. A summary of the current state of fusion plasma simulation can be found in the Appendix, where theoretical issues are also surveyed. Some progress in integrated modeling has been made as well, leading to important insights on topics as diverse as major disruptions, turbulence regulation by flows, and the design of compact stellarators.

Achieving the goals of the FSP will require significant collaborative advances in physics, applied mathematics and computer science. The wide range of temporal and spatial scales, extreme anisotropies and complex geometry, make this problem among the most challenging in computational physics. The numerical challenges for fusion simulation are outlined in Sections II*d* and III*e*. Methods for simulating phenomena coupled over disparate space and time scales, and over different dimensionality will require qualitative improvements, innovations and strong collaborations across all of the constituent disciplines. This disciplinary integration will be an essential element of the project. The project must develop software methodologies and frameworks for designing, building, maintaining, and validating the simulation software. Computer science issues raised by this initiative include: the choice of an architecture for interconnecting code modules; data models; performance monitoring and optimization; provision for flexibility and extensibility; and, tools for enabling human collaborations over long distances. These and related topics are also discussed in section III*e*.

Assessment of predictive models has been divided into two distinct activities: verification, which assesses the degree to which a code correctly implements the chosen physical model, and validation, which assesses the degree to which a code describes the real world. Overall, the goal of verification and validation is an assessment of the extent to which a simulation represents true system behavior sufficiently to be useful. Verification is particularly difficult for integrated models where analytic solutions may not be available in any regime. The validation process puts a premium on close collaboration between computational and experimental groups. To succeed, a central feature of this initiative must be an intensive and continual close coupling between the simulations efforts and experiments. The requirements for verification and validation are summarized in section III*f*.

The Fusion Simulation Project will require significant improvements in computational infrastructure. These include advances at major computational facilities which are shared across the Office of Science community as well as deployment and enhancements to local or topical computing centers. The simulations envisioned here will also produce truly prodigious quantities of data and will require

investments in advanced storage systems at all levels. With geographically dispersed resources and researchers, the wide-area network becomes a crucial element in the computing environment, with associated collaborative tools and protocols. Timely upgrades to the communication network and local infrastructure will be required. Of particular concern is connectivity to university or international partners. Infrastructure requirements are detailed in Section III*g*.

New funding necessary for the success of the FSP is presently estimated at approximately $20M per year for each of five years. To achieve the greatest productivity, this new research should be split between OFES and OASCR, with fusion scientists funded by OFES, and applied mathematicians, computer scientists, and the computational toolkits provided under the auspices of OASCR. This joint undertaking represents a significant opportunity for the DOE Office of Science to create a capability that will advance the understanding of fusion energy to a level unparalleled worldwide.

VI. ACKNOWLEDGEMENTS

We acknowledge helpful contributions to this report from members of the fusion and applied mathematics communities, with particular thanks to speakers at the May 23 ISOFS workshop, and to the speakers and participants of the September 17-18 FSP workshop. We thank the support staff without whom this activity would not have been possible, including Lucille Kilmer, Deanne Eggers, Ron Winther, Beulah Koz, and Marcia Freels. We also thank the Theory Coordinating Committee (TCC) for their letter, and the members of the PSACI PAC for the concepts and suggestions provided in their 18 June 2002 letter to this Subcommittee, and most particularly appreciate their enthusiasm for the Fusion Simulation Project. We are grateful for the excellent and substantive technical contributions: in the computational science sections of this report, by David Brown and David Keyes; and, in the fusion sections, by Jeff Candy, Ron Cohen, Nasr Ghoniem, Greg Hammett, Wayne Houlberg, David Humphreys, William Nevins, Ron Stambaugh, and Ron Waltz.

Attachment: ISOFS Charge Letter

February 22, 2002

Professor Richard D. Hazeltine, Chair
Fusion Energy Sciences Advisory Committee
Institute for Fusion Studies
University of Texas at Austin
Austin, TX 78712

Dear Professor Hazeltine:

This letter provides a charge to the Fusion Energy Sciences Advisory Committee (FESAC) to assist the Office of Fusion Energy Sciences (OFES) in preparing a roadmap for a joint initiative with the Office of Advanced Scientific Computing Research (OASCR). Recent reports, such as the FESAC report "Priorities and Balance within the Fusion Energy Sciences Program," the "Report of the Integrated Program Planning Activity" (IPPA), and the NRC report "An Assessment of the Department of Energy's Fusion Energy Sciences Program," have identified a predictive understanding as a measure of the quality of the science and the maturity of the knowledge base of a field. The IPPA report lists several challenging10-year objectives for the fusion program, including "develop fully integrated capability for predicting the performance of externally-controlled systems including turbulent transport, macroscopic stability, wave-particle physics, and multi-phase interfaces." This objective, as well as several other IPPA objectives related to innovative confinement configurations, will require significantly enhanced simulation and modeling capability. Therefore, the goal of this initiative should be to develop an improved capacity for **Integrated Simulation and Optimization of Fusion Systems**.

The initiative should be planned as a 5-6 year program, which would build on the improved computational models of fundamental processes in plasmas that are being developed in the base theory program and in the SciDAC program. Rough estimates are that an integrated simulation initiative would require a total funding level of about $20 million per year, with funding for the plasma scientists provided by OFES and funding for the applied mathematicians, computer scientists, and computational resources provided by OASCR. Thus, the roadmap should include not only human resources but also computer and network resources.

Please carry out the preparation of the roadmap using experts outside of FESAC membership, as necessary, including experts recommended by the Advanced Scientific Computing Advisory Committee. The sub-panel of experts should obtain community input through a series of workshops covering at least the following questions:

- What is the current status of integrated computational modeling and simulation?
- What should be the vision for integrated simulation of toroidal confinement fusion systems?

- What new theory and applied mathematics are required for simulation and optimization of fusion systems?
- What computer science is required for simulation and optimization of fusion systems?
- What are the computational infrastructure needs for integrated simulation of fusion systems?
- How should integrated simulation codes be validated, and how can they best be used to enable new scientific insights?

The ultimate product should be a roadmap document similar to the one developed for the Genomes to Life Initiative (http://www.doegenomestolife.org/roadmap/index.html). Please conduct a workshop on the first two questions above and provide a summary document with overall program goals and objectives, major program deliverables, and a brief description of the OFES and OASCR funded elements of the program by July 15, 2002, so that OFES would be able to include a description of the program in the FY 2004 OMB budget request. Please complete work on the final roadmap by December 1, 2002, in order to provide the detailed information needed by OFES and OASCR to develop detailed program plans, program announcements and grant solicitations.

I appreciate the time and energy that members of FESAC and FESAC sub-panels have provided to the continuing efforts to develop program plans and roadmaps for the OFES program. I am confident that the Committee's recommendations on a roadmap for **Integrated Simulation and Optimization of Fusion Systems** will form a sound basis for beginning a joint OFES/OASCR program.

Sincerely,

James F. Decker
Acting Director
Office of Science

Front Cover Attribution List

(Appears on back cover)

B. Carreras, W. Dorland, L.Garcia
G. Hammett, G. Kerbel, C. Kim
J-N. Leboeuf, V. Lynch, W. Park
S. Parker, B. Rogers, C. Sovinec
D. Spong, R. Sydora, NCSX Team
W7-X Team

Fusion Simulation Project (FSP)

Integrated Simulation and Optimization of Fusion Systems

Appendix. Overview of Fusion Science for the FSP

November 17, 2002

I. Introduction – Fusion Science Insights and Challenges
- A. Insights from numerical simulation
- B. Simulation challenges

II. Fusion Science Basic Concepts
- A. Basic equations
- B. Reduced descriptions

III. Microscopic Simulation
- A. Introduction and summary
- B. Neoclassical transport
- C. Turbulence and anomalous transport
- D. Non-local phenomena
- E. Challenges for turbulence simulation

IV. Macroscopic Simulation
- A. Introduction and summary
- C. Status
- B. Equations
- D. Challenges

V. Plasma Edge Physics and Plasma-Wall Interaction
- A. Introduction
- B. Properties of the edge plasma
- C. Computational challenges and codes for the edge region
- D. Whole Device Modeling

VI. Equilibrium Evolution and Confinement
- A. Introduction
- B. Transport equations
- C. Uses of transport codes
- D. Challenges of transport modeling

VII. External Sources
A. Introduction
B. Neutral beam injection
C. Edge particle sources
D. Radio-frequency wave heating and current drive
E. Summary: the role of source models in comprehensive integrated simulation
VIII. Fusion Simulation Capabilities Status
IX. Focused Integration Initiatives – Challenges and Opportunities
A. Plasma edge
B. Turbulence on transport timescales
C. Global stability
D. Whole device modeling
X. Validation Requirements in Fusion Topical Areas
A. Transport
B. MHD
C. Radio-frequency heating and current drive
D. Edge/scrape-off-layer/divertor
XI. Background – Orbits and Instability Mechanisms
XII. Glossary

I. Introduction – Fusion Science Insights and Challenges

This appendix summarizes the status of fusion science, emphasizing its theoretical and computational aspects. We wish to demonstrate first the scope and complexity of the physics and mathematics, second the progress that has been made up to now, and third the great challenges and opportunities before us. We discuss the several areas of fusion physics and computation that are referred to in the main report, beginning with a brief introduction to the basic equations and concepts in Sec. II. We follow this with more detailed surveys of the main areas of concern in Secs. III-VII. In Sec. VIII we summarize the 50 or so principal codes in use today. In Secs. IX and X we explore the focused integration initiative examples and validation requirements that are introduced in the main body of the report. In Sec. X we give a very brief review of some elementary plasma physics concepts. Finally, in Sec. XI, there is a glossary. Words introduced in *italics* are defined there.

Before going into the various topics, let us begin by examining our ultimate goal from the point of view of past accomplishments in computational plasma physics and some challenges that we expect will lead to future insights.

Numerical simulation and modeling, at varying stages of integration, have contributed greatly to insights regarding the behavior of magnetically confined plasmas. This has resulted from the continual interaction among computation, theory, and experiment. As the Fusion Simulation Project (FSP) develops, and wider ranges of physical processes are integrated into the simulations, there will be greater and greater reliance on massive computations. At the same time, as pointed out in the main text, the theoretical interpretation and experimental validation efforts will continue to be crucial to the overall success of the project.

We present a few examples to illustrate how this process works and how computation has contributed and will contribute to our insight and understanding.

A. Insights from numerical simulation

1. Global stability and major disruptions in tokamaks

A *major disruption* is the name given to an abrupt abnormal termination of a *tokamak* discharge with total loss of plasma current and confinement. This is thought by many to be the major obstacle that could potentially upset the development path of the tokamak as a practical energy source. Besides terminating the discharge, the disruption has the adverse effects of producing large thermal loads on the *divertor* and large electromagnetic forces on the vessel wall. (See cutaway view in Fig. III.1 of the main report.) Disruptions also have the potential of accelerating electrons to multi-MeV energies, and these can cause additional severe damage to the first wall.

Computer simulations have been crucial in identifying specific sequences of events that can cause disruptions. An excellent example of this are the simulations that explained the major disruption in the record-setting *TFTR* discharge that produced 10 MW of fusion power. [E. Fredrickson, W. Park, et al, "High-beta disruption in tokamaks", Phys. Rev. Lett, **75** 1763

(1995)]. It was shown that under the conditions of that discharge, it was energetically favorable for the plasma interior to deform into a long-wavelength helical structure. In the local regions of the torus where the *unfavorable curvature* of the helical structure aligned with the curvature of the torus itself, the plasma pressure could very easily cause instability. It was the interaction of this localized *ballooning* instability with the interior long-wavelength helical structure that destroyed the nested *magnetic surfaces* and precipitated the disruption. Figure A1 depicts a simulation of this process, with perturbed field lines and *poloidal* cross sections highlighted in colors representing mode amplitude.

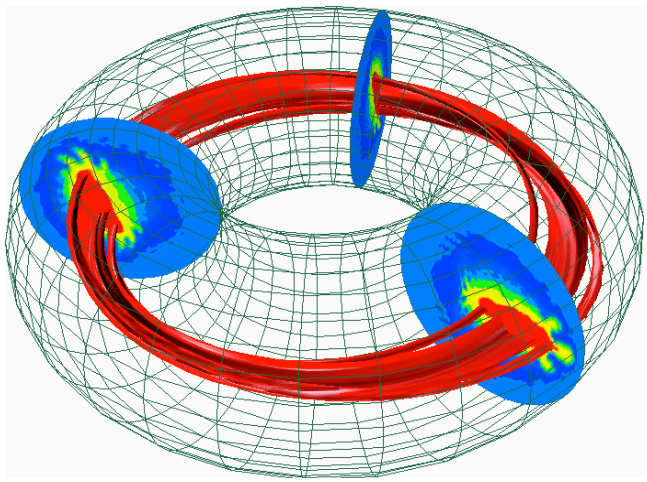

Fig. A1. An internal helical structure interacting with the toroidal geometry causes localized instability that can lead to a major disruption.
[Courtesy Wonchull Park, PPPL]

There are other causes of the major disruption, which we are just beginning to understand and are finding ways to control. Integrated simulations of the kind planned in this project will be instrumental in identifying and understanding these mechanisms. (See Sec. IV.)

2. Design of Three Dimensional Toroidal Configurations

Another example with already some degree of integration is the design of three-dimensional (3-D) stellarator configurations. *Stellarators* are advanced toroidal magnetic confinement devices that utilize 3-D plasma shaping to achieve steady-state operation and to optimize plasma confinement and stability. Among these, the design of compact (or low aspect-ratio) stellarators is one of the new, challenging research areas of interest in magnetic confinement. The challenge is that maintaining good transport and stability properties at low *aspect ratio* is generally more difficult than at high aspect ratio.

However, careful control of the 3-D shape appears to offer the flexibility to achieve good transport and stability characteristics. This insight has been realized through the development of sophisticated optimization codes that numerically explore 30-40 dimensional parameter spaces (the stellarator's outer flux surface shape is described in terms of 30-40 Fourier modes) and arrive at shapes that minimize a range of desired equilibrium, transport and

stability physics target functions [A. S. Ware et al., "High-beta equilibria of drift-optimized compact stellarators", Phys. Rev. Lett. **89**, 125503 (2002)]. Typical stability targets are the *Mercier criterion*, resistive interchange, high-n ballooning, and kink mode stability.

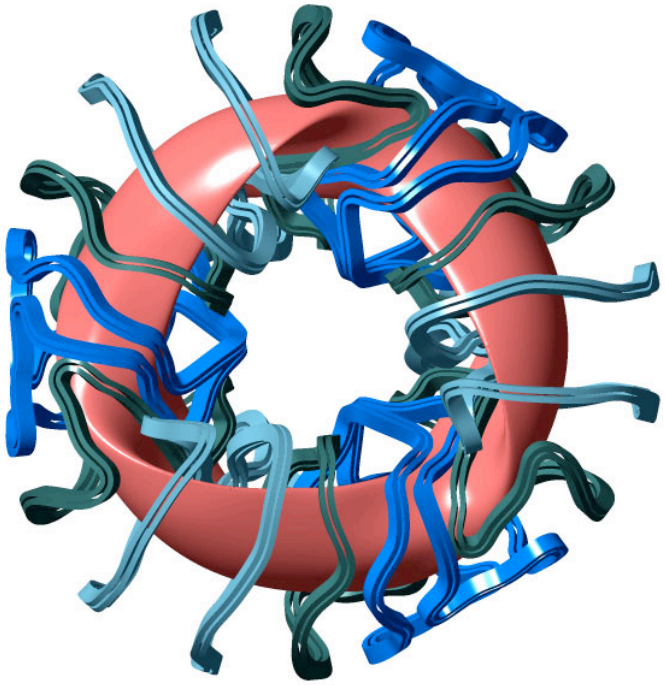

Fig. A2. Quasi-axisymmetric National Compact Stellarator Experiment (NCSX). [Courtesy G. H. Neilson, PPPL]

These optimization procedures have led to unexpectedly robust configurations such as the quasi-omnigeneous or quasi-poloidal compact stellarator (QPS) being designed at Oak Ridge National Laboratory and the quasi-axisymmetric National Compact Stellarator eXperiment (NCSX) device [http://www.pppl.gov/ncsx/] currently in the design phase at Princeton Plasma Physics Laboratory. Figure A2 is a plan view of NCSX, showing the outer closed flux surface and the main magnetic coils. The integration initiative could lead to the inclusion of more sophisticated and more compute-intensive particle-transport, 3-D stability, and coil-reconstruction modules directly in the optimization. This could yield yet more surprising and robust three-dimensional configurations suitable for reactor-grade plasma confinement.

3. Turbulence Regulation by Self-Generated Flows

A major step in tokamak research was the experimental realization that the High Mode (*H-Mode*) of confinement in tokamaks was accompanied by a sudden increase in the primarily poloidal rotation of the plasma in the region of steep pressure gradient at the plasma's edge. Almost contemporaneously, early computer calculations were showing that turbulence could spontaneously generate large radial scale flows, which formed a *transport barrier* [A. Hasegawa and M. Wakatani, " Self-organization of electrostatic turbulence in a cylindrical plasma", Phys. Rev. Lett. **59**, 1581 (1987)]. This insight was later confirmed experimentally with the discovery of internal barriers, confirmed computationally with the almost universal

observation of Reynolds-stress-generated flows in fluid, *gyrofluid* and *gyrokinetic* simulations, and put on sound theoretical foundations by several thorough analytical studies.

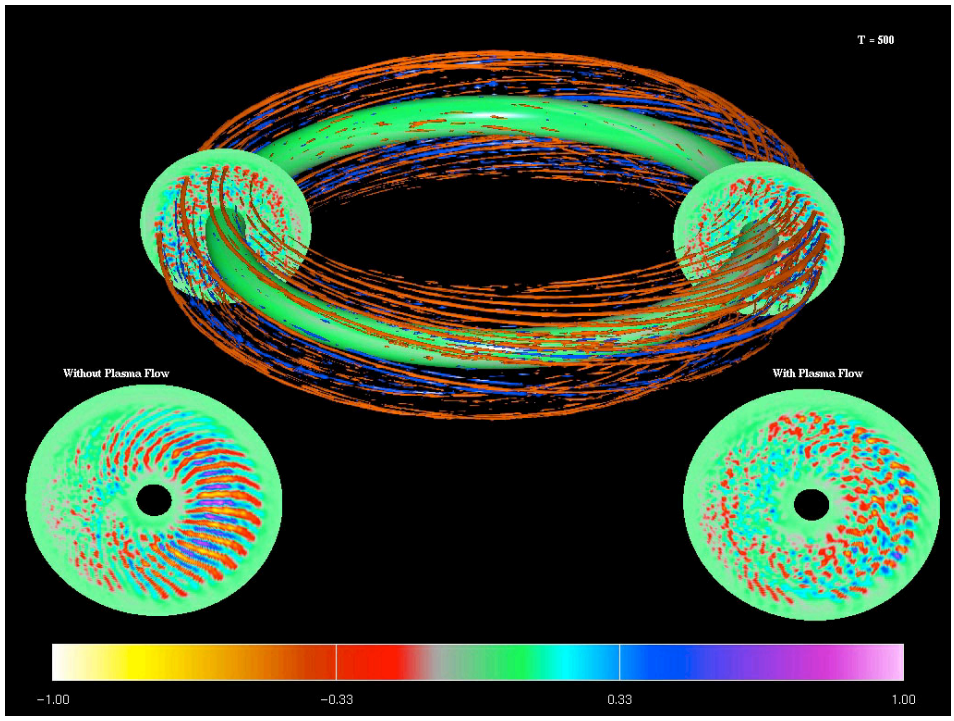

Fig. A3. Computational evidence of flow regulation of turbulence [Courtesy Z. Lin, PPPL/UC-Irvine]

In particular, the computations showed the effectiveness of these flows to regulate the turbulence in terms of scale size, fluctuation levels, diffusivities and therefore the transport. Recent theoretical effort motivated by differences in the characteristics of the turbulence obtained from fluid and kinetic models led to the identification of a residual flow damping mechanism, which in turn gave rise to an improvement of the computational methods. Further theoretical development has also led to the discovery of long-lived, fine radial scale flows nonlinearly generated by the turbulence or *zonal flows* whose properties and consequences were confirmed in the details by the computational models. Figure A3 represents a microturbulence simulation. Poloidal cross sections at bottom left and right show flow patterns without and with zonal flows [Z. Lin et al., "Turbulent Transport Reduction by Zonal Flows: Massively Parallel Simulations," Science **281**, 1835 (1998)].

We anticipate that these kinds of insights coming from the synergy among experiment, theory, and computation, exemplified by the role of flow in regulating the turbulence, will be one of the distinguishing characteristics of the integration initiative. (See Secs. IV-VI.)

4. Full-wave modeling of radio frequency heated, multidimensional plasmas

A major challenge in understanding waves for fusion is that they can dramatically change character after launch, i.e. undergo *mode conversion.* At a given frequency a uniform

magnetized plasma supports several different kinds of waves, with very different characters with respect to polarization, speed and wavelength, such as fast *magnetosonic* waves, slow ion-cyclotron waves and *ion Bernstein waves* (IBW). (See Sec.VII.) In a non-uniform plasma these different waves can convert at isolated spatial locations, the classic example being conversion of the fast magnetosonic wave to an IBW at the 2-ion hybrid resonance.

Understanding this process in the realistic 2- and 3-dimensional geometries of toroidal magnetic confinement devices has required computational modeling and only recently has it become possible to fully resolve the mode conversion layer, for example with the All Order Spectrum Algorithm or AORSA code [E. F. Jaeger et al., "Advances in full-wave modeling of radio frequency heated, multidimensional plasmas", Phys. Plasmas **9**, 1873 (2002)].

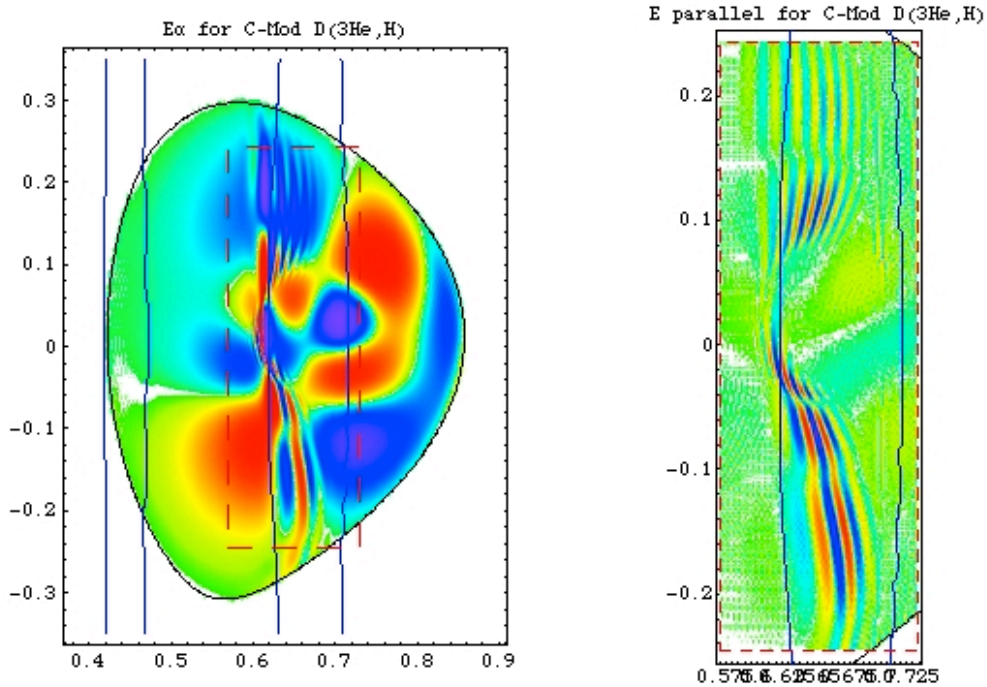

Fig. A4 Mode conversion to fast ion cyclotron wave revealed by AORSA computations. [Courtesy D. Batchelor, ORNL]

The surprise is that these realistic-geometry and well resolved calculations show that with full poloidal field the fast magnetosonic wave converts dominantly to a slow ion-cyclotron wave, not an IBW. This had been hinted at by F. Perkins who did an approximate 1-D calculation in 1977 showing that both types of conversion could, in principle, occur [F. W. Perkins, " Heating Tokamaks via the ion-cyclotron and ion-ion hybrid resonances", Nuclear Fusion **17**, 1197 (1977)]. But the 1-D model did not show that one could actually get a wave in a real 2-D plasma, launched from a real antenna, to the conversion layer with the right conditions for conversion. Nor did it tell how important each process would be. The 2-D calculation gives the complete, quantitative picture, but the 1-D calculation gives the qualitative paradigm by which one understands the complete results. Figure A4 shows wave

amplitudes in a cross section of the *C-Mod* tokamak, with details in the conversion layer expanded in the right-hand diagram.

As with the other examples cited, RF calculations with ever improving resolution and ever tighter coupling to more comprehensive plasma models such as those contemplated as part of this integration initiative are expected to lead to even more surprises.

B. Simulation Challenges

1. Modeling edge physics

Overall transport and confinement are determined to a great extent by the height of the temperature *pedestal* at the plasma edge. This temperature is itself governed by phenomena that start at the wall, and propagate through the *separatrix* region and into the edge region proper. It is expected that coupled and complex models of particle and heat transport, neutral and impurity fluxes, edge-gradient induced *magnetohydrodynamic (MHD)* instabilities, and turbulence in a single computational edge framework will be required to pin down which of these mechanisms, either by itself or in combination with another, regulates the pedestal height. (See Sec. IXA.) Coupling to core plasma models will probably start with the integrated edge models providing the pedestal temperature as a boundary condition for the core before more complete integration becomes feasible through a combination of physics, algorithmic, and hardware advances.

2. Modeling turbulence on transport time scales

An integrated model of turbulence on transport time scales is a daunting physics and computational task. It is nevertheless deemed feasible at several levels, each exploiting separation of space and time scales appropriately. (See Sec. IXB.) One approach consists of iterative solutions of the macroscopic transport equations with occasional updates of the profiles and diffusivities transmitted to and from the microscopic turbulence calculations. Another entails direct coupling of transport equations to gyrofluid-type microscopic models using a physics-based, yet reduced, description of the turbulence. The result of high confidence integrated modeling of turbulence and transport might be the discovery, through computations, of new favorable operating modes, such as a new H mode, with the ultimate outcome being the determination of transport from first principles.

3. Coupling of electrons and ions in core turbulence

Integration will enable examination of the combined role of turbulence on space and time scales associated with the electrons and ions. (See Sec. IIA.). To date realistic and tractable numerical simulations of each exist separately. The expected advances in integration algorithms to bridge disparate time and space scales will enable incorporation of the dynamics of each species with realistic mass ratios, and for realistic scale sizes with respect to experiments, into one computational model that covers the full range of relevant turbulence. The benefits of these computations will be to determine the extent to which

electron and ion turbulence overlap in terms of spectral range and to what extent one range of scales influences and/or regulates the transport associated with the other.

4. Extended MHD

Integration will facilitate extensions to MHD computations beyond the conventional ideal and resistive models. (See Sec. IV.) Routine incorporation of two-fluid effects will determine for instance whether *whistler wave* physics included through the *Hall term* will play as important a role in phenomena leading to plasma disruptions as it does for magnetic reconnection in space and astrophysical plasmas. Similarly, the routine inclusion of diamagnetic effects associated with the non-scalar nature of the ion pressure tensor will reveal whether the pressure-induced plasma rotation provides a way to control MHD activity which is as effective nonlinearly as it is linearly for realistic toroidal plasmas. Moreover, the inclusion of minority ion species with non-Maxwellian populations either as gyrofluids or full-fledged particle species will enable extended MHD models to describe the essential macroscopic physics of burning plasmas. Such features will provide new intuition regarding the effect of energetic particles on large-scale MHD instabilities and vice-versa.

5. The burning plasma experiment

Perhaps the greatest advances afforded by integrated modeling will be realized for burning plasmas. The grand challenge in the world fusion program is the burning plasma experiment, a necessary predecessor to a practical power reactor because, by its very nature, it presents a new category of technical issues. All other fusion experiments obtain the required plasma temperatures by applying energy sources from outside the plasma. A burning plasma – a self-sustaining energy source – maintains them through the action of the ongoing fusion reactions, which produce energetic alpha particles and high-energy neutrons. The charged alphas are trapped by the magnetic field and transfer their kinetic energy to the plasma, thus maintaining the temperatures necessary for fusion to continue.

With self-heating as the dominant plasma heating mechanism, new plasma processes and effects will arise. The high flux of energetic particles will impact the plasma wave structure, alter the plasma pressure and current profiles, and produce, through alpha-particle and neutron collisions, a rich source of wall interaction phenomena. Most seriously, all these effects will be strongly coupled and must be understood and managed in an integrated fashion to insure the stability and success of the experiment. The fully integrated fusion simulator will be targeted to model these processes and their consequences, thereby providing the essential insights to guide experimental programs, optimize machine design, and deepen our understanding of the fundamental science.

II. Fusion Science Basic Concepts

The fundamental problem in fusion science is to understand and control the physical processes that determine the balance between heating and fueling, on the one hand, and loss of confinement, on the other. The losses can come from macroscopic processes leading to

breakup of the nested magnetic flux surfaces and disruptions or from microscopic processes, leading to transport across the flux surfaces. Figure 5A shows a tokamak cross-section, schematically illustrating some of the processes that will be discussed in the following sections. A cutaway view of this device is shown in Fig. III.1 of the main report.

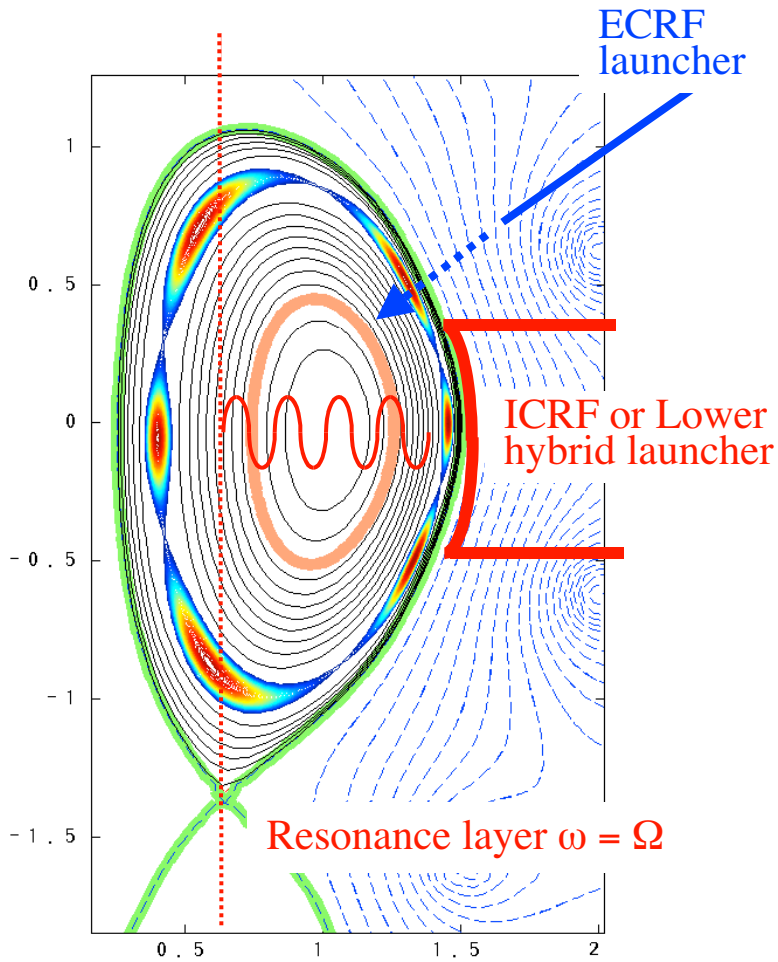

Fig. A5. Tokamak cross-section with launchers, a schematic wave and resonance layer. Here $\Omega = \omega_c$. Also shown is a chain of magnetic islands.

We begin in Part A by introducing the basic equations and concepts of plasma physics, which are the foundation of the subsequent sections, and which give rise to the time-scale chart appearing as Fig. III.2 of the main report. Reproducing this chart in Part B, we briefly outline these time scales.

A. Basic equations

The time-dependent processes governing the stability and confinement of the plasma are governed by the Maxwell-Boltzmann system of equations, which we write in MKS units. We

begin with the distribution function $f_a(\mathbf{x},\mathbf{v},t)$, which is the density of charged particles of species a in a six-dimensional phase space. This function evolves in time, t, according to the Boltzmann equation,

$$\frac{\partial f_a}{\partial t} + \mathbf{v}\cdot\nabla f_a + \frac{q_a}{m_a}\left[\mathbf{E} + \mathbf{v}\times\mathbf{B}\right]\cdot\nabla_{\mathbf{v}} f_a = C(f_a). \tag{A-1}$$

The characteristics of the left-hand side are the particle orbits, the second term, $\mathbf{v}\cdot\nabla f_a$, being the convection in space and the third term, $(q_a/m_a)(\mathbf{E}+\mathbf{v}\times\mathbf{B})\cdot\nabla_v f_a$, being the acceleration in velocity space arising from the *Lorentz force,* where **E** and **B** are the electric and magnetic fields, and q_a and m_a are the charge and mass of species a. Solution of the equation can begin with integration along the characteristics. (A brief Introduction to these is given in Sec. XI.) This is complicated, however, both by the right-hand side of the equation, which represents all sources, sinks, and collisional dissipation, and by the nonlinearity of the whole system. That is, the fields themselves include not only the externally generated **B**, but also spontaneous perturbations and injected waves that depend on the particle distributions through the Maxwell equations,

$$\nabla\times\mathbf{B} = \mu_0\mathbf{J} + \mu_0\varepsilon_0\frac{\partial\mathbf{E}}{\partial t}, \qquad \nabla\cdot\mathbf{B} = 0, \tag{A-2}$$

$$\nabla\times\mathbf{E} = -\frac{\partial\mathbf{B}}{\partial t}, \qquad \nabla\cdot\mathbf{E} = \rho/\varepsilon_0, \tag{A-3}$$

where **J** and ρ are the sums of individual species current and charge densities,

$$\mathbf{J}(\mathbf{x},t) = \sum_a \mathbf{J}_a(\mathbf{x},t) = \sum_a q_a \int d^3\mathbf{v}\mathbf{v} f_a(\mathbf{x},\mathbf{v},t), \tag{A-4}$$

$$\rho(\mathbf{x},t) = \sum_a \rho_a(\mathbf{x},t) = \sum_a q_a \int d^3\mathbf{v} f_a(\mathbf{x},\mathbf{v},t). \tag{A-5}$$

Theoretical and computational challenges. Equations (A-1)-(A-3) generate an enormous array of phenomena. We shall touch upon the more important of these in the following sections. The sources of this complexity are:

1. The complex structure of the confining magnetic fields, which give rise to multiple classes of orbits even in the absence of perturbations,
2. The long-range nature of the electromagnetic forces, which give rise to collective phenomena, both externally driven and spontaneously occurring waves and instabilities.
3. The great range of space and time scales and extreme anisotropy of both individual particle motions and plasma waves, extending from the size of the device (a few meters) and the duration of the discharge (exceeding several minutes in a recent record-setting experiment) all the way down to the electron *gyroradius* and *gyrofrequency*, and

4. Dissipation, i.e., collisions among ion species and electrons, as well as those with neutrals (sources and sinks), all buried in the term $C(f_a)$ on the right-hand side of Eq. (A-5). Even when the collision term itself is negligibly small, dissipation occurs and plays a major role, owing to *wave-particle resonances*, and nonlinear decorrelation of waves.

An electron in a uniform magnetic field **B** orbits the field line in a perpendicular plane (Part B, below) with gyrofrequency $\omega_{ce} = eB/m_e \approx 5\times10^{11}$ s^{-1} and with gyroradius $\rho_{ce} = v/\omega_{ce} \approx 6\times10^{-5}$ m. (These numerical values are for a typical plasma temperature of 3 keV and a typical magnetic field strength of B = 3 T.) Frequencies and wavelengths for the highest range of radio-frequency wave heating are of the same order. (Ions gyrate with frequency lower and radius larger by m_e/m_i and $\sqrt{m_i/m_e}$, respectively.) The same electron moves freely parallel to the field, but in the curved and spatially varying fields of a torus, it is subject to additional slower drifts and accelerations on the whole device scale, of order $2\pi R \approx 10$ m in present-day machines, where R is the *major radius*.

Thus, our phenomena extend over a range of 9-10 orders of magnitude in time and, perpendicular to the magnetic field, 4-5 orders in space. For much of our work, we can also average over the gyro-motion, reducing our equation to five phase-space dimensions (*gyrokinetic theory*), and increasing the time step by another order of magnitude. Even so, one estimate of the coverage needed for a full-discharge simulation would require 10^{11} phase-space points at 10^8 time steps. Not only does this far exceed our present capability, but also we would not be able to store or analyze the output from such a calculation.

Actual plasma computations, therefore, are carried out by means of ordering schemes leading to approximations that focus on particular scales and are valid for particular phenomena. Beginning with Eqs. (A-1)-(A-3), the equations to be solved are obtained by averaging over the finer space and time scales and applying some prescription for closure. For example, in a toroidal device, particles drift away from magnetic surfaces at a rate slower than the gyration by factors of order $\rho^* = \rho_{ci}/a$, where a is the minor radius. Relatively infrequent collisions or decorrelations then allow transport losses described by a random walk with diffusivity of the form $D \sim (\delta x)^2/\delta t$ where δx and δt, the characteristic step size and collision time, exceed the shortest scales given above. *If* D *and related quantities can be simply characterized,* then the resulting transport equations are solved on still longer time and space scales. Whole-device simulation currently proceeds in this manner, but the "if" in the preceding sentence is a big one.

Much progress has been made in the major subfields of plasma theory and computation, which we describe in the following sections. Figure III.2 in the main report, repeated as Fig. A6 in the next section, shows a chart of the principal frequency ranges and the phenomena associated with them.

The great challenge and opportunity before us is to further develop and combine these in a tractable way, leading to whole-device simulation that contains complete and reliable models of all the relevant physics.

B. Reduced descriptions

The codes currently in use are each based on different approximations to the full set of plasma equations (A-1) to (A-3) that isolate phenomena in a more restricted range of frequencies. Together, they span the ranges of frequency described above. A summary of the four major code groups presently in use and the timescales being addressed by them is given in Fig. A6.

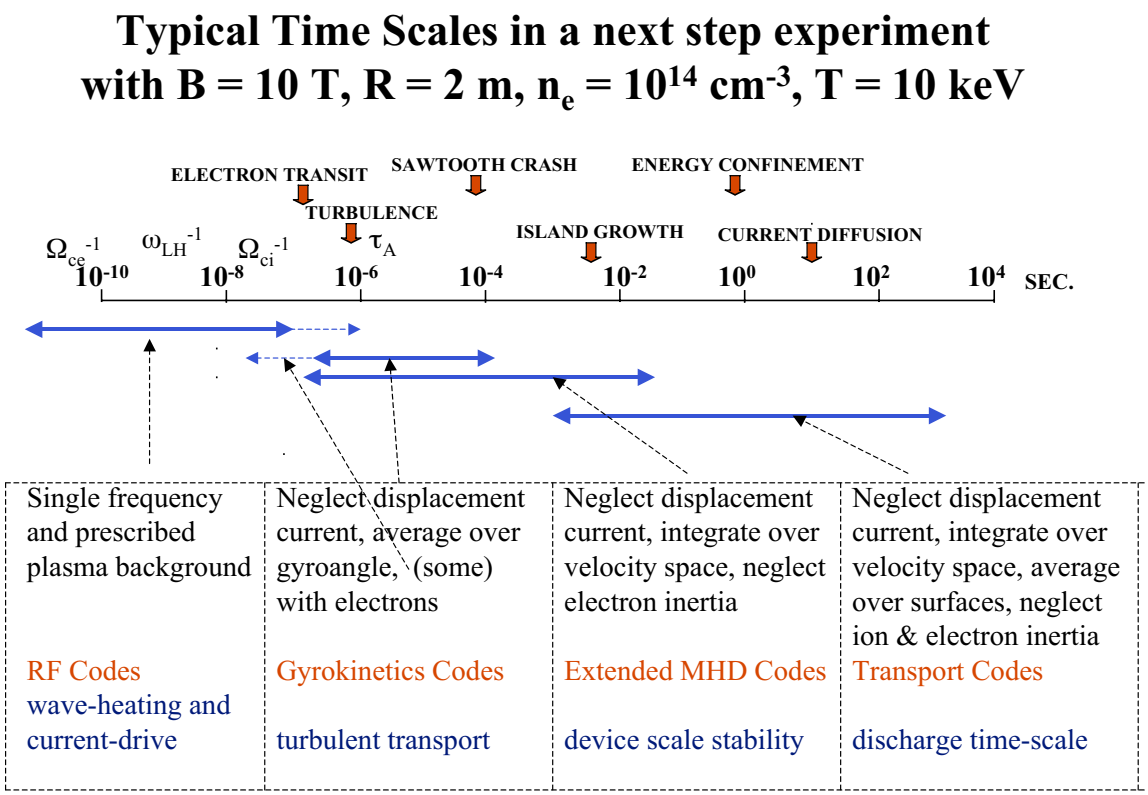

Fig A6. Summary of the four major code groups and the timescales being addressed.

The RF codes (external sources) address frequencies of order the ion cyclotron frequency, Ω_{ci}, and above, up to the electron cyclotron frequency Ω_{ce}. (This notation is used interchangeably with ω_{ci} and ω_{ce}.) They assume a single, fixed frequency set by the oscillator, and solve for the linear and quasi-linear response of a given background plasma to the imposed electromagnetic fields at this frequency. They are used to calculate wave-heating and current-drive by radio-frequency (RF) sources.

The gyrokinetics codes (microscopic modeling) are based on an analytic averaging of the full kinetic equation to eliminate the fast gyro-orbit frequencies from the equations. They and the

other codes described below also neglect the displacement current term in Maxwell-Equation (A-2) to remove light waves from the system. These reductions allow them typically to take time steps about 10 times longer than the ion cyclotron frequency (whose motion is analytically averaged over) and are normally run for 10^3 to 10^4 time steps to calculate stationary turbulent fluctuation levels. Recent additions to these codes to include some electron timescale phenomena bring in the electron transit time, which lowers the maximum time step.

The Extended MHD codes (macroscopic modeling) are based on taking velocity moments (or weighted integrals over velocity space) of Eq. (A-1) to obtain fluid-like equations that describe large-scale phenomena more efficiently. They aim to resolve phenomena occurring on the Alfven transit time, τ_A, although most codes are at least partially implicit to avoid a strict restriction on the time step based on this. These codes can normally run 10^4 to 10^5 time steps to address MHD phenomena such as sawteeth and island growth.

The 2-D transport codes (equilibrium evolution and confinement) further simplify the equations by eliminating the Alfven waves and by averaging plasma properties over 2-D magnetic surfaces. These codes are very efficient and can take many long time steps to model the entire discharge. However, the 2-D edge transport codes need to resolve the parallel dynamics and use a time-step based on the parallel sound-wave propagation near the edge region.

The following sections describe each of these code groups in more detail.

III. Microscopic simulation

A. Introduction and summary

This section deals with the direct calculation of the microscopic processes that affect the quality of a magnetic field configuration as a thermal insulator. Loss of plasma particles and energy across field lines results from at least three categories of microscopic phenomena: diffusion and convection based on individual orbits and collisions (classical, neoclassical), anomalous diffusion and convection from turbulent microinstabilities (usually thought to be dominant), and phenomena that are instantaneous and non-local.

- Neoclassical codes for both 2- and 3-D configurations are available subject to certain approximations. Additional work will be required to implement fully the existing theories and to account for additional effects in the inter-operative environment.
- Micro-turbulence-driven anomalous transport will be dealt with on three levels: a) fine-scale stand-alone gyrokinetic simulations of core plasma turbulence will continue to be developed, interpreted by theory, and compared with experiment to firmly validate the fundamental theory and benchmark the other descriptions; b) reduced simulations will be coupled directly to the transport equation solvers. In some cases, it may be possible to couple directly the full turbulence simulations with the transport equation solvers (a topic of one of our focused initiatives); and c) algebraic models of transport will continue to be developed, again informed by

theory, experiment, and the just-described simulations. When supported by turbulence simulations for selected cases, they will provide the most rapid parameter scans.

- Work will continue on development of models of observed rapid and non-local phenomena (e.g. avalanches, radiative transport, or global magnetic interactions) that do not fit into the diffusive-convective approach. The architecture must include provision for these from the outset, so they can be incorporated when available.

A measure of success for these efforts and their validation will be the ability to predict and model a transport barrier, a region of steep gradients where turbulence is suppressed. This will involve integrated treatment of the radial electric field as also mentioned in Secs. V and VI.

B. Neoclassical transport

At the most basic level, transport is a random-walk process, in which a particle drifts slowly away from a magnetic flux surface and from time to time suffers a collision that transfers it to a new orbit. The resulting diffusivities have the form $D,\chi \sim (\delta x)^2 / \delta t$ where δx and δt are the characteristic step size and collision time. The detailed task of *neo-classical theory* is to compute the flux-surface averaged fluxes (see Sec. VI, Eqs. (A-24)-(A-25)) by solving a gyro-averaged form of the Boltzmann equation (A-1), called the *drift-kinetic equation*. Closure is provided in an ordering in which the collision operator is dominant, and the step size (drift orbit) is small compared to the plasma gradient scale length. (*Classical transport*, in which the step size is of the order of the gyroradius ($\delta x \sim \rho_a$), is small and most often neglected in fusion calculations.)

Neoclassical theory is highly developed in the limit that finite orbit-width effects (violation of the above ordering) and energy scattering are unimportant. Codes for both 2- and 3-D configurations are available subject to these approximations. Additional theoretical work and extensive code development will be necessary in order to account for a number of important phenomena associated with neoclassical theory such as: neoclassical transport when barriers turn off the anomalous transport, *bootstrap* effects (neoclassical currents driven by gradients), non-local orbit effects such as *potato orbits* (orbits near the plasma center), self-consistent electric field effects in non-axisymmetric configurations, impurity transport etc. Additional work will be required to fully implement the existing theories in the "interoperative" environment, particularly in 3-D configurations (stellarators) and in tokamaks with transport barriers. To treat the mutual interaction of islands or stochastic magnetic fields with neoclassical driven transport and currents will require extensive further development.

C. Turbulence and anomalous transport

Gradients of plasma density and temperature provide the free energy to drive micro-instabilities that lead to anomalous transport. These waves typically propagate with phase velocities perpendicular to the field of order $v_a^* = \rho_a v_a \, d\ln n_e / dr$ (see Sec. XI, Eq. (A-30)), or $v_{Ta}^* << v_a$ where the temperature gradient replaces the density gradient. This gives rise to the characteristic *diamagnetic frequencies* $\omega_a^* = k_\perp v_a^*$ and $\omega_{Ta}^* = k_\perp v_{Ta}^*$ in a wave $\tilde{\phi} \sim \exp(i\mathbf{k}_\perp \cdot \mathbf{x})$,

where $k_\perp \sim \rho_a^{-1}$, for a = i (ion waves) or a = e (electron waves). The waves are known generically as *drift waves*, particular examples of which are the *ion temperature-gradient* (ITG) modes and *electron temperature-gradient* (ETG) modes.

To evaluate these we begin with the Maxwell-Boltzmann system, Eqs. (A-1)-(A-5), neglecting displacement current and charge density (*quasi-neutrality*). Averaging over the gyro-angle, we obtain the *guiding center* distribution $F_a(\mathbf{x},\mathbf{v},t) = F_{0a}(\rho,W)\left(1+q_a\tilde{\phi}/T_a\right)+\tilde{h}_a(\mathbf{x},W,\mu,t)$ in a 5-dimensional phase space, where the velocity variables are the magnetic moment $\mu = mv_\perp^2/2B$ and kinetic energy $W = mv^2/2 = \mu B + mv_\parallel^2/2$. It is convenient to separate out the background distribution $F_{0a}(\rho,W)$ and the so-called *adiabatic part* $F_{0a}(\rho,W)q_a\tilde{\phi}/T_a$. (The latter is actually the lowest order solution in one important limit.) The remaining part of the distribution, $\tilde{h}_a$ obeys the *gyrokinetic equation,* in which the convective term in (A-1) is replaced by a similar term with gyro-averaged particle drifts replacing the velocity **v**:

$$\frac{\partial \tilde{h}_a}{\partial t} + \left(\mathbf{v}_{\chi a} + \mathbf{v}_{da} + v_\parallel \hat{\mathbf{b}}\right)\cdot\nabla\tilde{h}_a = -\mathbf{v}_{\chi a}\cdot\nabla f_{0a} - q_a\frac{\partial f_{0a}}{\partial W}\frac{\partial\tilde{\chi}}{\partial t} \qquad \text{(A-6)}$$
$$+ \text{ collisions } + \text{ sources/sinks,}$$

where $\hat{\mathbf{b}}$ points in the direction of the equilibrium magnetic field, $\mathbf{v}_{da}$ is the curvature- and grad-B drift, and the **E**×**B** drift is combined with transport along perturbed magnetic fields lines and the perturbed grad-B drift as

$$\mathbf{v}_{\chi a} = \hat{\mathbf{b}}\times\nabla\chi_a/B, \qquad \text{(A-7)}$$

where

$$\chi_a = \left(\left\langle\tilde{\phi} - v_\parallel\tilde{A}_\parallel\right\rangle_g + \left\langle\mu\tilde{B}_\parallel/q_a\right\rangle_g\right)\Big/B. \qquad \text{(A-8)}$$

Here, we represent the fluctuating fields in terms of the potentials $\left(\tilde{\phi},\tilde{A}_\parallel\right)$ plus $\tilde{B}_\parallel$ and denote the gyro-average by the angle brackets $\langle\ldots\rangle_g$. The principal nonlinearity of Eq. (A-6) is the perturbed convective term $\mathbf{v}_{\chi a}\cdot\nabla\tilde{h}_a$.

To illustrate the method and make contact with theory, we note that the gyrokinetic ordering requires a separation of scales that lends itself to an eikonal representation, i.e., a perpendicular expansion of fluctuations as $\exp(iS)$ where $\nabla S = \mathbf{k}_\perp$ may be interpreted as the perpendicular wave vector. In this case, Eq. (A-8) may be evaluated as

$$\tilde{\chi}_a = J_0(b_a)\left(\tilde{\phi} - v_\parallel\tilde{A}_\parallel\right) + J_1(b_a)/b_a\left(\mu\tilde{B}_\parallel/q_a\right). \qquad \text{(A-9)}$$

The Bessel functions J_0 and J_1 have argument $b_a = k_\perp v_\perp/\omega_{ca}$. Of course, in a non-uniform medium the perpendicular wave number is really a differential operator. (In order to evaluate these functions efficiently, the codes we will describe make extensive use of numerical

Fourier transformation and its inverse.) In particular, the gyro-averaged wave potential is given by

$$\left\langle \tilde{\phi} \right\rangle_a (\mathbf{x}) = \frac{1}{2\pi} \oint d\zeta \tilde{\phi}(\mathbf{x} + \rho_a) = J_0(k_\perp \rho_a) \tilde{\phi}(\mathbf{x}), \tag{A-10}$$

where ζ is the gyro-angle, as illustrated in Fig. A7.

Finally, the Maxwell equations are recast in potential form, e.g.,

$$\nabla^2 A_\parallel = -\mu_0 \sum_a q_a \int d^3v v_\parallel J_0(b_a) \tilde{h}(a), \tag{A-11}$$

where the right-hand side is the parallel current.

We draw attention to the fact that solving the *continuum* Eq. (A-6) is equivalent to following gyro-averaged orbits of an ensemble of discrete particles,

$$F_a = \sum_{i=1,N} w_i(t) \delta(\mathbf{x} - \mathbf{x}_i(t)) \delta(\mathbf{v} - \mathbf{v}_i(t)), \tag{A-12}$$

with suitable particle smoothing, Monte Carlo treatment of collisions, sources and sinks.

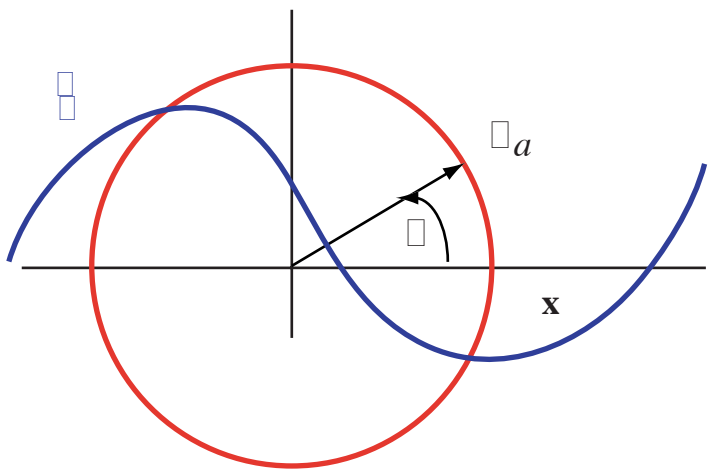

Fig. A7 Gyro-orbit average of perturbed potential.

The system consisting of Eq. (A-6) and the Maxwell equations, Eqs. (A-2)-(A-3) (neglecting the displacement current,) drives turbulence on scales ranging from the electron to the ion gyroradius, as noted at the beginning of this section. The linear and nonlinear growth and damping mechanisms have been well delineated theoretically, and a number of codes are able to compute the evolution and saturation of the waves *and the resulting flux-surface-averaged transport fluxes* (see Sec. VI). The codes are of four types. Each may be either *flux-tube* (localized perpendicular to the field) or *global* (covering a major fraction of the plasma radius). Also, each may be either *continuum* (solving Eq. (A-6) on a fixed Eulerian grid), or

particle-in-cell (or PIC, following an ensemble of gyro-averaged particles in a Lagrangian formulation) as in Eq. (A-12). All four types of code have advantages and disadvantages.

The SciDAC numerical turbulence project supports a two-by-two matrix of codes:

	Continuum	PIC
Flux tube Global	GS2 GYRO	SUMMIT GTC

There are both continuum and PIC because continuum is currently the most developed (kinetic electrons, important even for waves on the ion gyro-scale, and perturbed magnetic fields) while PIC may be ultimately more efficient. A flux tube is described in coordinates that are extended along field lines but are localized in perpendicular directions. There are high-resolution *flux-tube codes* to support multiple space and time scales (electron and ion gyro-physics) and validation of local turbulence physics (rapid parameter scans). *Global codes*, which provide calculations over a major fraction of the plasma radius, are needed to account for extended profile effects and are most likely the best bet for coupling to the integrated simulation.

The plasma turbulence is quasi-2-dimensional because of the plasma anisotropy. The current state-of-the art in global PIC simulations with GTC, without including full electron dynamics or magnetic perturbations, is the following. The resolution requirement along **B** is determined by the equilibrium structure. The structure across the field is determined by the microstructure of the turbulence, $\sim \rho_{ci}$ for *ion temperature-gradient* (ITG) *modes*. This requires $\sim 64 \times \left(a/\rho_{ci}\right)^2 \sim 2 \times 10^8$ grid points and 8 particles per spatial grid point, where a is the minor radius of the machine. This leads to $\sim 1.6\times10^9$ particles and 1 terabyte of RAM for 600bytes/particle. This resolution is currently achievable. The time scale is on the order of $a/v_i \sim$ 1 μsec for 10 time steps. To simulate many correlation times corresponding to a simulation of a few ms requires about 90 hours of IBM SP (SEABORG) time at 4×10^{-9} sec/particle-timestep. This is heroic, but the discharge time actually simulated is much smaller than the plasma equilibration time.

The situation is similar for continuum codes. GYRO, in particular, has the most complete physics to date for ion gyro-scale physics: ions and electrons (trapped and passing), magnetic perturbations and collisions, real geometry, nearly full radius, finite $\rho^* = \rho_{ci}/a$ with profile and $\mathbf{E}\times\mathbf{B}$ shear stabilization and toroidal velocity-shear drive. Simulation of a single radial slice ($\Delta r/a$ = 0.3) of DIII-D plasma for about 1ms takes five 24-hour submissions on 128 processors on SEABORG. Scaling from DIII-D to ITER and taking full radius requires a factor of 15. The 1ms rescales to 2ms but it is still a long way from 3 sec confinement times. This reflects the fact that going beyond the previous state of the art from ions-only simulations to simulations with electron dynamics requires an order of magnitude jump in computing power because the required time step is at least 10 times smaller. Output from GYRO is shown in Fig. A8, showing the strong anisotropy.

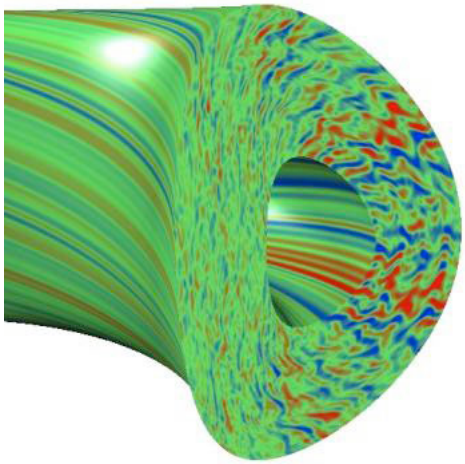

Fig. A8. Ion temperature-gradient turbulence in GYRO simulation. Extended structures along the field lines are evident. [Courtesy J. Candy, GA]

D. Non-local phenomena

We cannot leave this discussion without taking note of observed phenomena in tokamaks that seem to defy the diffusion-convection picture. These include extremely rapid (i.e. faster than the transport time scale) propagation of heat pulses (or "cold" pulses) resulting from sawtooth crashes, impurity injection at the edge, or the H-mode transition itself. There are two current approaches to explain these observations. First, there are empirical techniques based on the concepts of *self-organized criticality* (SOC) and *avalanches*. Empirical models have given good results, but the underlying physics has not been elucidated. Second, the instabilities leading to turbulence usually have fairly well defined thresholds, e.g., a critical temperature gradient. A rapid pulse may occur if the plasma is perturbed from its marginal condition. In this case, the underlying physics is presumed known, but quantitative calculations have not been performed. The simulation project must be prepared to incorporate developments in this area.

E. Challenges for turbulence simulations

The bottom line is that we are currently far from the achievement of turbulence simulations with all the relevant physics on a scale suitable for integration with transport calculations. One estimate, presented at the may 23 ISOFS Workshop, is that we are six orders of magnitude from a solution based on current computational methods and computers, and that we will make up only four of these orders in the next few years by advances in computer technology and currently envisioned schemes for exploiting the time-scale separation (see Sec. IXB). Among the computational and applied mathematical challenges are:

- Continuum kernels solve an advection/diffusion equation on a 5-D grid; we therefore need: linear algebra and sparse matrix solvers (LAPAC, UMFPAC, BLAS), and distributed array redistribution algorithms.
- Particle-in-Cell kernels advance particles in a 5-D phase space and need: efficient "gather/scatter" algorithms that avoid cache conflicts and provide random access to field quantities on a 3-D grid.

- Continuum and Particle-in-Cell kernels perform elliptic solves on 3-D grids (often mixing Fourier techniques with direct numerical solves).
- Other Issues are portability between computational platforms, characterizing and improving computational efficiency, distributed code development, and expanding the user base.

Finally, we note that while core turbulent transport is extremely important (ability to predict internal barriers, for example), edge turbulence, which has the same difficulties described here and more (Sec. VIII), is critical, for the edge pedestal is the greatest source of uncertainty for reactor predictions.

IV. Macroscopic simulation

A. Introduction and summary

Coupled magnetohydrodynamics (MHD) and Maxwell equations play a central role in modern fusion plasma theory. First, MHD determines 2-D or 3-D magnetic equilibrium with nested toroidal magnetic flux surfaces, which are crucial for magnetic plasma confinement. Next, MHD describes both thresholds and nonlinear dynamics of device-scale plasma instabilities (so called MHD instabilities). Very often these instabilities set the limits of the performance of fusion devices.

Plasma dynamics can be completely described by the evolution of the distribution function $f_\alpha(\mathbf{r},\mathbf{v},t)$, for each particle species α, given by each species plasma kinetic equation, together with the self-consistent evolution of the electric and magnetic fields, given by the Maxwell equations, as given in Eqs. (A-1)-(A-6). Solving these equations for space scales and timescales characteristic of large-scale instabilities in confined plasmas is computationally impractical. The Extended MHD approach is to reduce the dimensionality of the problem, by multiplying the kinetic equation by successive powers of the particle velocity $\mathbf{v}$ and integrating over velocity space. If the underlying distribution functions have nice properties, such as a close-to-Maxwellian velocity distribution, the resulting moment equations have fluid-like properties. They are more tractable theoretically and computationally, although formidable problems may still arise.

Magnetic fusion devices are very rich in MHD activity, some relatively benign, some leading to catastrophic disruptions. Some of these are known as *sawtooth oscillations*, *tearing* and *ballooning instabilities*, and *resistive wall modes*, whose general features and some quantitative predictions can be given in terms of ideal or resistive MHD. As higher plasma temperatures are reached, more and more kinetic effects are required to be included.

B. Equations

On the time and space scales that the electrons and ions maintain local charge neutrality, the lowest order moment equations for the electron and ion species can be added together to form a set of equations for a plasma fluid with a density $\rho = (M_i\, n_i + m_e\, n_e)$, fluid velocity $\mathbf{v}$ =

$(M_i\, n_i\, \mathbf{v_i} + m_e\, n_e\, \mathbf{v_e}) / (M_i\, n_i + m_e\, n_e) \sim \mathbf{v}_i$ and pressure $p = p_i + p_e$. (Note that the symbol ρ is being used here for the mass density. The same symbol has elsewhere been used to represent charge density and normalized minor radius.) The displacement current can be neglected in Ampere's law (since $\nabla \bullet \mathbf{J} = 0$), eliminating electromagnetic radiation and electrostatic oscillations. The MHD equations can be written in a general form, in MKS units, as the low-frequency Maxwell equations,

$$\frac{\partial \mathbf{B}}{\partial t} = -\nabla \times \mathbf{E} \quad , \qquad \nabla \times \mathbf{B} = \mu_0 \mathbf{J} \quad , \qquad \nabla \cdot \mathbf{B} = 0 \quad , \tag{A-13}$$

the continuity equation,

$$\frac{\partial \rho}{\partial t} = -\nabla \cdot \rho \mathbf{v} \quad , \tag{A-14}$$

the total momentum equation,

$$\rho\left(\frac{\partial \mathbf{v}}{\partial t} + \mathbf{v} \cdot \nabla \mathbf{v}\right) = -\nabla \cdot \mathbf{P} + \mathbf{J} \times \mathbf{B} \quad , \tag{A-15}$$

and the energy equation,

$$\frac{\partial p}{\partial t} + \mathbf{v} \cdot \nabla p = -\frac{5}{3} p \nabla \cdot \mathbf{v} - \frac{2}{3}\left(\Pi : \nabla \mathbf{v} - \nabla \cdot \mathbf{q} + Q\right). \tag{A-16}$$

The energy equation (A-16) assumes that the ratio of specific heats $\gamma = \gamma_e = \gamma_i = 5/3$. In these equations, μ_0 is the permeability of free space, n is the number density, ρ is the mass density, $\mathbf{v}$ is the center of mass velocity, $\mathbf{B}$ is the magnetic flux density, $\mathbf{E}$ is the electric field, $\mathbf{J}$ is the current density, p is the scalar pressure, $\mathbf{q}$ is the heat flux, η is the electrical resistivity, the stress tensor is $\mathbf{P} = p\mathbf{I} + \mathbf{\Pi}$, where $\mathbf{I}$ is the unit tensor and $\mathbf{\Pi}$ is the traceless part of the stress tensor, and Q is other heat sources and sinks.

In general, the electron motion decouples from the bulk fluid motion, although the two are related by $\mathbf{v} = \mathbf{v_e} + \mathbf{J}/ne$. For near-Maxwellian distribution functions, for example, decoupling can occur due to the effects of a non-negligible ion *Larmor radius* (*finite Larmor radius* or *FLR*), which is still small relative to the system size. The equation for the electric field, known as Ohm's law, comes from the electron momentum equation. Ignoring terms of order m_e/M_i, it is

$$\mathbf{E} = -\mathbf{v} \times \mathbf{B} + \eta \mathbf{J} + \frac{1}{ne}\mathbf{J} \times \mathbf{B} - \frac{1}{ne}\nabla \cdot \mathbf{P}_e - \frac{1}{ne}\rho_e\left[\frac{\partial \mathbf{v_e}}{\partial t} + \mathbf{v_e} \bullet \nabla \mathbf{v_e}\right] \tag{A-17}$$

It requires a pressure or temperature equation for the electrons,

$$\frac{\partial p_e}{\partial t} + \mathbf{v}_e \cdot \nabla p_e = -\frac{5}{3} p_e \nabla \cdot \mathbf{v}_e - \frac{2}{3}\left(\Pi_e : \nabla \mathbf{v}_e - \nabla \cdot \mathbf{q}_e + Q_e\right) \quad , \qquad \text{(A-18)}$$

Additional detail, still within the confines of a two-fluid moment description, can be obtained by keeping the anisotropies relative to the confining magnetic field, such as the two pressures $p_\perp$ and $p_{||}$ and/or the heat fluxes. The above equations then refer to the average quantities $p_j = (p_{j||}+2p_{j\perp})/3$, etc.

To close the system, expressions for the higher order moments $\mathbf{\Pi}$ and $\mathbf{q}$ must be obtained independently, from solutions to the kinetic equation. At high collisionality, these are the usual collisional viscous stress tensor and the heat flux (proportional to the local velocity and temperature gradients, respectively), and this leads to the *Braginskii equations*. At lower collisionality or long mean free path, these terms contain non-local kinetic effects. Proper closure becomes a complex question that must take into account details of the confinement configuration. (In toroidal systems, these non-local geometrical effects have been addressed in a flux-surface-averaged sense by neoclassical theory.) Unfortunately also, in this limit there is no single, unambiguous way to define the set of "two-fluid" or FLR terms, so that models depend upon a mixture of theoretical and practical considerations (see R. D. Hazeltine and J. D. Meiss, Plasma confinement, Addison Wesley, 1992.)

Another approach that is being pursued is the so-called hybrid approach where a distribution of particles is used to provide closure to the fluid equations. There are 2 categories of the hybrid approach that are being used: pressure coupling and current coupling. Even within these categories, there are several approaches.

In the *pressure-coupling scheme*, the distribution of particles is used to calculate the pressure tensor, and then the velocity is advanced in time from Eq. (A-15). The pressure tensor has been calculated differently by different researchers: from the gyroviscous stress tensor in terms of gyrofluid moments; by using approximations to the gyroviscous stress tensor given in terms of Bessel functions in Fourier-Ballooning space; by using a particle Hall-MHD closure that uses test particles to compute the off-diagonal elements of the pressure tensor; or by calculating the stress tensor directly from summing moments of Gyrokinetic particles.

In the *current coupling scheme*, the ion current is calculated directly from the particles, and no fluid ion equations of motion is needed. This has been implemented using the full equation of motion for the particles. This method is relatively straightforward but inefficient. There is also discussion of implementing current coupling using gyrokinetic particles, in which case the polarization current must be dealt with.

C. Status

Currently two state-of-the-art 3-D codes are being supported by the SciDAC Center for Extended Magnetohydrodynamic Modeling (CEMM) project; M3D and NIMROD. These are both focused on the modeling of linear and non-linear phenomena in fusion experiments that require Extended-MHD descriptions.

The NIMROD code solves the primitive form of the plasma fluid-model in axisymmetric toroidal, cylindrical, or periodic-linear geometry with arbitrary poloidal cross-sectional shape. (The geometry must have an ignorable periodic coordinate, but the simulated dynamics are fully three-dimensional.) The user selects which terms are retained in Ohm's law through an input parameter. The semi-implicit numerical method is used to advance the solution from initial conditions. This avoids severe time step restrictions associated with wave-like normal modes of the system, sound, Alfven, and whistler waves—while avoiding numerical dissipation. For accuracy at time steps that are orders of magnitude larger than explicit stability limits, the semi-implicit operator for mass motions is based on the linearized ideal MHD energy integral. Matrix inversion is accomplished by parallel preconditioned Krylov methods, which is the most computationally demanding part of the time advance. Performance is therefore dependent on the effectiveness of the preconditioner.

The spatial representation of NIMROD is an important feature of the code. NIMROD uses a combination of logically quadrilateral and triangular finite elements for the poloidal plane and pseudospectral collocation for the periodic direction. The polynomials used for finite element basis functions are selected by the user for optimal efficiency, and poloidal mesh lines need not be orthogonal. For many fusion problems, accuracy is improved by aligning grid lines with the equilibrium flux surfaces inside the separatrix. The grid can also be packed around low order rational flux surfaces to efficiently resolve the small spatial scales that arise at high Lundquist number S. Triangular meshing outside the last closed flux surface allows complicated, realistic boundary shapes.

The M3D code is a parallel code that is especially suited for geometries with inherently three-dimensional boundaries, e.g. stellarators, but can also be used to simulate axisymmetric devices. M3D consists of two parts, a mesh module and a physics module. The mesh module contains the grid, implementation of differential and integral operators, I/O, and inter-processor communication. The physics module handles time advancement of the equations and contains a hierarchy of physics levels that can be invoked to resolve increasingly complete phase-spaces, and therefore provide increasing realism. The module includes resistive MHD, two-fluid, and kinetic particles. Electrons are represented as a fluid with an approximate fluid closure. M3D uses a stream function/potential representation for the magnetic vector potential and velocity that has been designed to minimize spectral pollution. Parallel thermal conduction is simulated with the "artificial sound" method. The solution algorithm is quasi-implicit in that only the most time-step limiting terms including the compressional Alfven wave and field diffusion terms are implemented implicitly, with explicit time stepping used for the remaining terms.

A three-dimensional mesh is utilized to facilitate the resolution of multi-scale spatial structures, such as reconnection layers and to accommodate fully three-dimensional boundary conditions that occur in stellarators or the evolving free boundary of a tokamak bounded by a separatrix. The mesh uses unstructured, 3-D piecewise-linear triangular finite elements in the poloidal sections. The domain decomposition consists of slicing the toroidal geometry into a set of poloidal planes with each poloidal plane further partitioned into equal area patches. One or more of the poloidal patches are assigned to each processor. The fluid part of each time step consists of uncoupled 2-D scalar elliptic equations that are solved

concurrently within each poloidal plane. The PETSc library has been used extensively to provide a portable, efficient parallel implementation for the elliptic equations that need solution at each time step. These are solved with a Krylov accelerated iterative scheme that uses the overlapped Schwarz method for preconditioning. This leads to excellent parallel scalability.

D. Challenges

These codes require high resolution and many time steps to give an accurate representation of modern fusion experiments. To see this, recall that resistivity effects are characterized by a resistive diffusion time scale, $\tau_R \sim \mu_0 L^2/\eta$, which is much larger than the Alfvén-wave transit time, $\tau_A \sim L\sqrt{\rho\mu_0}/B$. (Here $L \sim 1\,m$ is the spatial scale of the device). In fusion machines the Lundquist number, $S = \tau_R/\tau_A$, is of the order of 10^8. As a result, even though resistive effects determine the physics of the process, they actually become important only within some very narrow layers. Proper resolution of these layers as well as strong anisotropy of plasma properties along and across the magnetic field imposes serious computational challenges.

We can estimate the computational resources required to carry out the necessary simulations for parameters typical of present and proposed experiments as shown in Table A-I.

TABLE A-I. Typical dimensionless parameters for present and proposed experiments

parameter	name	CDXU	NSTX	CMOD	DIII-D	FIRE	ITER
R (m)	Major radius	0.3	0.8	0.6	1.6	2.0	5.0
T_e [keV]	Elec. Temp.	0.1	1.0	2.0	2.0	10	10
β	Plasma beta	0.01	0.15	0.02	0.04	0.02	0.02
$S^{1/2}$	Inv resis. length	200	2600	3000	6000	20,000	60,000
$(\rho^*)^{-1} = Ba/T^{1/2}$	Ion number	40	60	400	250	500	1,200
$a/\lambda e$	Recip. norm e skin depth	250	500	1000	1000	1500	3000

Estimates based on explicit time-stepping with no grid refinement. Let us first estimate the computational requirements for a 3-D calculation with uniform zoning of size Δx and a explicit time-stepping scheme based on the CFL criteria for the poloidal Alfven wave, i.e. $\Delta t = \Delta x / V_{AP}$. For a 3-D mesh of linear dimension N, i.e., N^3 mesh total mesh points, it would take N time steps to calculate one Alfvén wave transit time $\tau_A = a / V_{AP}$. Typical ideal and resistive MHD instabilities would grow on the timescales $T_{IDEAL} \sim \beta^{-1/2}\,\tau_A$ and $T_{RESIS} \sim S^{1/2}\,\tau_A$, requiring about $\beta^{-1/2}$ N and $S^{1/2}$ N time-steps, respectively. Thus, the total number of space-time points required to compute an ideal or resistive instability would be about $\beta^{-1/2}\,N^4$ (ideal) and $S^{1/2}\,N^4$ (resistive). The plasma beta, denoted by β, is the ratio of plasma pressure to magnetic pressure. Its value is an important measure of confinement quality.

Current experience shows that with real performance of about 100 Mflops/processor and of order 1000 processor-hours, we can compute a problem with 100^3 mesh points for 10^4 time

steps, using a complex fluid model with the compressional wave and field terms implemented implicitly. This is about 10^{10} space-time points in 3×10^{14} operations. This is easily sufficient resolution and time steps to calculate an ideal mode in CDX-U, and is nearly sufficient to study the initial growth phase of a nonlinear tearing mode. We see this since the number of linear mesh points is comparable to: the linear tearing-layer width, $S^{-1/2}$, the ratio of the system size to the ion Larmor radius $1/\rho*$, and the ratio of the system and the ratio of the system size to the electron collisionless skin depth, a/λ_e. These are the relevant lengths that enter the two-fluid Extended MHD model.

The question is: what type of computer power is needed to study this physics in a larger, hotter device with a stronger magnetic field? We see from Table I that depending on what scale length needs to be resolved, the number of mesh points in a linear direction will increase by about an order of magnitude as we go from CDX-U to DIII-D. The increase in the total number of space-time points would be the fourth power of this factor times another scaling factor that will between unity (for ideal scaling) to about 10 (for resistive scaling). Thus, the number of space-time points required would increase anywhere from 10^4 to 10^5. Running on a 10 Teraflops (delivered) computer for 3 days would correspond to about 3×10^{18} floating point operations which would be about 10^4 times greater than what was quoted above for what is available to us today, so a full DIII-D calculation might be feasible with this hardware increase alone.

Grid refinement, implicit time stepping and improved algorithms. It is straightforward to see that the above scaling estimates can be gross overestimates if we take into account improved algorithms and meshing. For example, for a field-line following mesh and with adaptive mesh refinement, the total number of mesh points should only have to grow linearly as we go to larger machines and higher resolution, rather than cubic. With implicit time differencing, the time step will not have to decrease nearly as fast as linear with zone size. More efficient solvers can give an additional factor. Thus, with these computational improvements, some of which have already been implemented to varying degree in M3D and NIMROD, we can realistically expect to be able to calculate modes using Extended MHD models in DIII-D, NSTX, and CMOD in the time frame of this project, and even FIRE and ITER calculations might be within reach. Of course, it also may not be necessary to simulate the exact parameters of a machine if we can determine scaling relations from doing a series of calculations at reduced parameters.

V. Plasma Edge Physics and Plasma-Wall Interaction

A. Introduction

The edge plasma, which bridges the hot plasma core and the material wall (see Fig. A9) plays a crucial role in both overall plasma confinement and plasma-wall interactions. Some examples of the impact of the edge plasma are:

1. Changes in edge plasma parameters can lead to dramatic improvement in core plasma confinement in the H-mode via the formation of transport barriers, regions with

reduced plasma turbulence. The height of the pedestal of this barrier plays a crucial role in the performance of a fusion reactor;

2. Practically all magnetic fusion devices have a density limit, which may be due to strong anomalous plasma transport at the edge;
3. There is a strong, but poorly understood, influence of neutrals and wall conditions on plasma confinement, e.g., one of the best shots in the TFTR tokamak was enabled by lithium conditioning of the walls;
4. Heat load on first wall, which is determined by edge plasma conditions, is a serious issue for reactor-relevant conditions;
5. Wall sputtering, transport of ions and neutrals, including hydrocarbons, in the edge plasma, and deposition processes caused significant accumulation of tritium in the first wall of both TFTR and JET tokamaks.

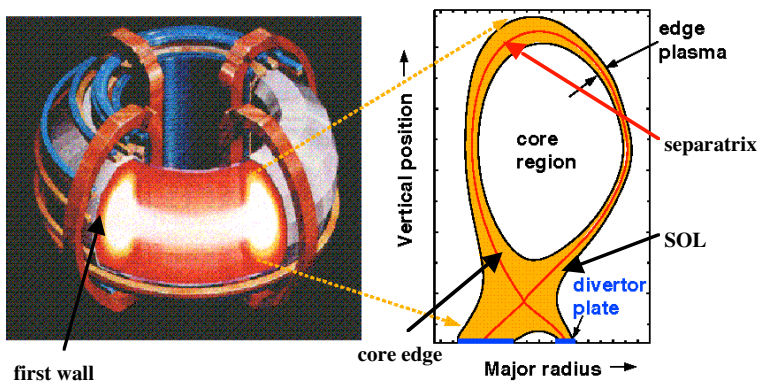

Fig. A9. Schematic view of edge plasma region in tokamak

The plasma edge region has many of the same problems as the core, but it also has a number of attributes that make it crucially distinct from the core. In particular, in toroidal magnetic fusion devices, such as tokamaks, field lines in the core of the device lie, at least approximately, on closed toroidally shaped flux surfaces, giving rise to good confinement of particles and heat. However, on the periphery of such devices, there inevitably exists a *scrape-off layer* (SOL), where magnetic field lines are in direct contact with material surfaces. Due to competition of parallel and perpendicular plasma transport, two-dimensional effects are strong in the SOL. In addition, so-called *divertor configurations* are common to most of the large toroidal devices, in which extra field coils are added to make the magnetic field intersect material surfaces at a location relatively remote from the main plasma volume, as shown in Fig. A9. In divertor configurations the magnetic *separatrix* divides closed magnetic field lines in the core-edge region from open ones in the SOL.

In the core plasma, the separation of scales is both a blessing and a curse. On the one hand, as in gyrokinetic theory, an ordering exists that allows one to solve the equations by averaging over smaller scales and ignoring variations in the longer scales. On the other hand, the same scale separation contributes to the difficulty of a first-principles whole-device simulation. In the edge, by contrast, the spatial scales are compressed, tending to invalidate the ordering

schemes. Gyrokinetic ordering, for example, can break down, i.e. $\rho * \sim 1$, and the collisionality varies along the field from the long mean-free-path to the short mean-free-path regime.

Thus, the solution of the fundamental equations becomes much more difficult, but if solutions are obtained, they should be readily assimilated into the global simulation–or become a microcosm of the global simulation. Indeed, one of the focused integration initiatives (FIIs) might be devoted to this entire region.

B. Properties of the edge plasma

Plasma near the separatrix, being impacted by fast parallel transport to material surfaces, tends to have steep radial gradients in temperature and density and to be relatively cold. The low temperature, coupled with proximity to bounding surfaces, results in a relatively high concentration of neutral gas and impurities. These properties lead to a relatively strong role for atomic physics processes: ionization, recombination, excitation, and radiative transport.

There is general consensus, supported by both analytic theory and numerical simulations of anomalous transport, that the poloidal **E×B** flows, spontaneously generated by the nonlinear dynamics, play a central role in regulating the saturation level of the turbulence and the resulting cross-field transport. Strong shear of this poloidal flow tears apart convective eddies and reduces the turbulence level and cross-field transport by forming the H-mode transport barrier. Due to the interplay between significantly different physics on open and closed magnetic flux surfaces, strong shear of the radial electric field and **E×B** poloidal flow arises, and, as a result, an H-mode transport barrier may be formed somewhat inside the separatrix. Reduction of anomalous transport at the barrier causes steepening of the plasma temperature and density profiles in this region. As a result, strong and repetitive MHD modes can develop causing ELM bursts. It is believed that the MHD modes responsible for ELMs are the so-called ballooning and peeling modes driven by plasma pressure and electric current gradients.

With increasing plasma density, the plasma particle flux to the divertor targets starts to decrease as the *detached divertor regime* is being formed. It is rather well understood and shown with both simplified analytic models and sophisticated 2-D plasma transport codes such as UEDGE that plasma-neutral coupling and atomic-physics effects, including impurity radiation and plasma recombination, play key roles in establishing this regime. However, there are strong indications that cross-field plasma transport also plays an important role here.

Disruption of the discharge for densities above the *density limit* is believed to be due to the mixed effects of enhanced plasma transport to the first wall and thermal collapse of the plasma due to impurity and hydrogen radiation. Interestingly, both the DBM (Univ. Maryland) and BOUT (Lawrence Livermore National Lab.) edge-plasma turbulence codes show trends in plasma transport enhancement at high plasma densities that are somewhat similar to that seen in experiments.

We note that in both improved confinement H-mode and standard L-mode there is rather strong interaction of plasma with first wall material surfaces resulting in both sputtering and (re- co-) deposition. These are rather complex processes involving ion transport and neutral-particle transport in edge plasmas, chemistry associated with both heavy particle interactions and interactions with electrons, and surface effects (implantation, collision cascades, deposition, adsorption, desorption, diffusion, etc.). In all regimes the first wall is a huge reservoir of hydrogen isotopes, which plays a dominant role in neutral gas recycling.

C. Computational challenges and codes for the edge region

Fusion plasmas in general have a large span of spatial and temporal scales, from the electron gyroradius and cyclotron frequency to the device size and the energy confinement time. This large span makes simulation a substantial challenge. In the hot core plasma, there is often a wide separation between the space and time scales characterizing turbulent fluctuations and those characterizing evolution of the equilibrium. This scale separation facilitates the use of separate simulation codes to describe turbulence and transport. However, in the edge region, this scale separation can become small or even non-existent, leading to the challenge of combining a wide range of turbulence and transport scales into a single, large-scale simulation. In addition to the plasma-physics scales, the presence of important atomic physics processes introduces new length and time scales, which create a range larger than that for the core. Ionization, recombination, and charge-exchange rate coefficients for hydrogen and impurities are in the range of 10^{-14} - 10^{-13} m^3/s, which for typical densities of 10^{20} m^{-3} give time scales of 10^{-6} - 10^{-5} s. On the other hand, near-unity recycling from the hydrogen-saturated material surfaces, where each incident ion results in a neutral hydrogen atom injected back into the edge plasma, yields a long time scale, 10^{-1} s, for establishment of equilibrium profiles.

The edge region of fusion devices generally has substantially lower plasma temperatures than the core, resulting in Coulomb collisional mean-free-paths parallel to **B** being much less than the *connection length*, i.e., the parallel distance traveled in making one poloidal revolution. Furthermore, the gradient scale-lengths and observed turbulent wavelengths perpendicular to **B** are usually greater than the ion gyroradius. Because of the spatial localization provided by collisions and **B**, fluid models have been adopted to give the basic description for both turbulence and transport simulations. However, there is growing concern in the community that fluid models are not adequate for plasma conditions in the H-mode pedestal region, and therefore they should be replaced with more accurate, but much more complex, kinetic models.

The principal edge-plasma turbulence codes within the U.S. community are presently BOUT and DBM. These are 3-D turbulence codes dealing with a fluid plasma description based on the electromagnetic Braginskii equations for plasma vorticity, density, electron and ion temperatures and parallel momentum. The BOUT code is non-local, can describe a magnetic X-point geometry on both sides of the separatrix, and is based on a toroidal segment simulation volume, while the DBM code is a flux-tube code (presently without an X-point) and well suited to parametric studies. With sources added in the core-edge region and sinks

in the SOL, BOUT has begun to follow some short-time profile evolution in response to the turbulence.

The numerical algorithms used in BOUT consist of finite-difference equations in 3-D where the resulting ordinary differential equations for the time dependence of each cell variable are advanced with the fully implicit Newton-Krylov solver PVODE. BOUT is written in C. The implicit integration increases the time step by a factor of 3-6 without preconditioning. Parallelization is obtained by domain decomposition in the direction along **B,** utilizing MPI for message passing. Because of weak coupling between the domains, scaling with the number of processors is essentially linear up to 120, but using 64 processors is more typical, and simulations have been done on various parallel platforms (IBM SP, T3E, SUN, DEC and Linux clusters).

The DBM code is written in Fortran 90 and solves the reduced Braginskii equations in a general flux-tube magnetic geometry, without (as yet) a magnetic X-point. It utilizes fourth-order spatial finite differencing and a second-order-accurate trapezoidal leapfrog time advance. The communication routines are MPI-based and at present, like BOUT, involve 1-D domain decomposition.

In addition, the edge plasma community has extensive experience with kinetic neutral transport and related atomic physics via DEGAS-2 (PPPL) and TNG (UCSD), as well as the 2-D coupled fluid plasma/neutral transport via UEDGE (LLNL), all of which provide an excellent base for planned extensions. UEDGE is an implicit time-dependent code capable of very long-time simulations of profile evolution that includes the important neutral particle sources from recycling surfaces and gas puffing, plus impurity species.

However, even though there is a consensus that neutrals can be an important ingredient in both edge-plasma turbulence and in the formation of the H-mode barrier, so far edge plasma turbulence codes are lacking the proper physics that would take these effects into consideration.

To describe erosion of the first wall under normal operation conditions the REDEP/WBC package is often used in the U.S. edge plasma community. The WBC is a Monte Carlo code, which computes impurity-atom and ion motion at the kinetic description level, including sputtering and deposition processes, and both elastic and inelastic impurity-plasma collisions. The code is very time consuming to run; therefore it is often run separately, and its output used as input to the REDEP code, which uses cruder models for impurity transport. As an input for plasma parameters the REDEP/WBC package uses either experimental data or UEDGE modeling results. We note that the cross-field impurity-ion transport built into this package is rather rudimentary

VI. Equilibrium Evolution and Confinement

A. Introduction

Plasma transport or confinement will be a major component of the integrated simulation project. Processes described in the preceding section determine the sources and sinks that drive the plasma fueling, heating, and rotation. In the following section we discuss the confinement or transport mechanisms that balance these. First, however, let us set up the framework in which these operate. With exceptions to be discussed below, the plasma equilibrium consists of nested flux surfaces. Because plasma parameters equilibrate rapidly along the magnetic field, they are nearly constant on these surfaces, and thus transport takes place principally across magnetic surfaces, and can be described in one spatial dimension, denoted by a generalized *radius* or flux label ρ. The term *1-1/2 D transport* refers to the fact that ρ must be related geometrically to the 2-D equilibrium. These *1-1/2 D* transport codes are central to today's efforts at integrated simulation. We will come back to this point later.

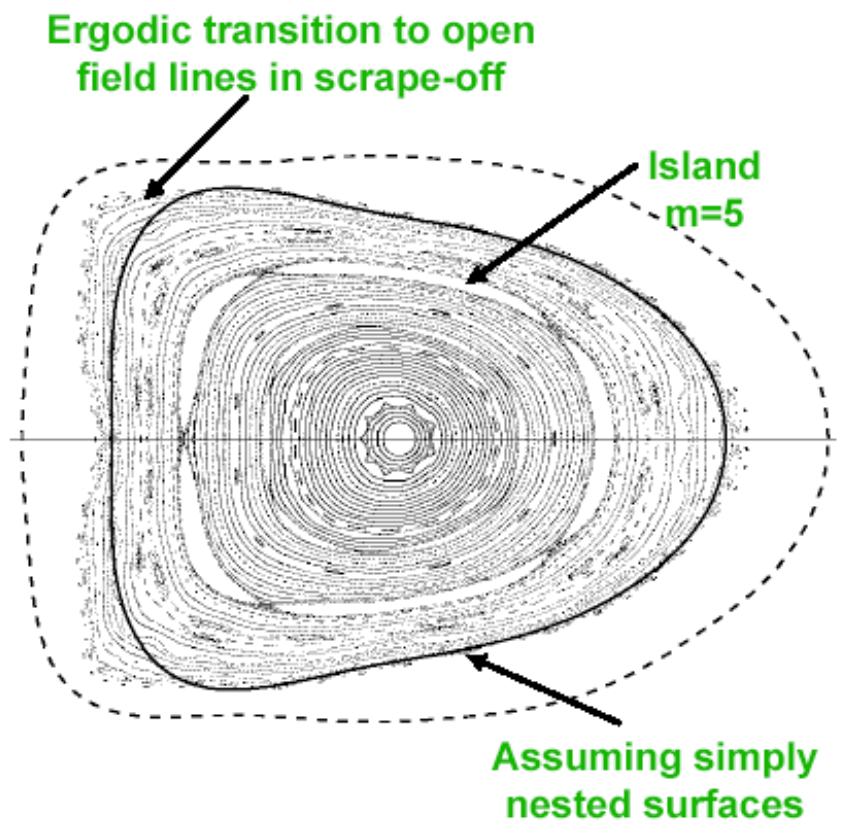

Fig. A10 PIES equilibrium of a low aspect-ratio stellarator [Courtesy J. Lyon and W. Houlberg, ORNL].

The breakdown of the above argument gives rise to a physics and computational challenge that will be central to the integration initiative. Islands and ergodic regions can form, owing to MHD instabilities and field errors in tokamaks, as well as in 3-D equilibria in stellarators. Figure A10 depicts a cross-section of a low aspect-ratio stellarator equilibrium showing

islands and ergodic regions. (This is a puncture plot with each dot representing the passage of a magnetic field line.) These regions evolve on the transport time scale.

B. Transport equations

Despite these caveats, tokamak transport is reliably described by 1-D surface-averaged transport equations in most circumstances. For stellarators, a design goal is to minimize the externally driven islands and stochastic regions so that similar transport studies are approximately valid.

Construction of a set of 1-D equations makes use of characteristic time scale separation. The fastest or Alfvén time scale (~μs) is assumed to establish the basic magnetic geometry. This is computed with MHD equilibrium codes that are free from the Alfvén time scale so this timescale is effectively eliminated from the problem. Likewise, particle densities and temperatures are assumed to equilibrate rapidly along the magnetic field lines, so that they can be considered 1-D functions of the flux coordinate function, thus eliminating the fast parallel transport time scale from the problem. The 1-D functions evolve on timescales characteristic of cross-field transport, which is one of the slowest timescales under consideration. This timescale establishes the profiles of thermodynamic quantities, density, temperature, and angular momentum. Careful selection of variables is required to take advantage of this timescale separation.

The geometry of the flux surfaces is specified by solving the MHD equilibrium equations: $\mathbf{J}\times\mathbf{B}=\nabla p$, $\nabla\times\mathbf{B}=\mu_0\mathbf{J}$, and $\nabla\cdot\mathbf{B}=0$. In two-dimensional geometry, appropriate for tokamaks without magnetic islands, flux surfaces are known to exist, and this leads to a nonlinear partial differential equation, the *Grad-Shafranov equation*, for the poloidal *magnetic flux* Ψ as a function of the cylindrical coordinates *R, Z* in the poloidal plane. Different techniques are used for *fixed boundary* calculations where the plasma/boundary interface is specified, and *free boundary* calculations where the plasma/boundary interface is determined self-consistently from the actual coils that produce the confining magnetic field. In three dimensions, there is an additional complexity in that there is no guarantee that the nested flux surfaces exist. The VMEC 3-D equilibrium code assumes the existence of these surfaces, whereas the PIES 3-D equilibrium code does not.

Figure A5 (or A10), *without* the islands or ergodic regions, is a typical tokamak (or stellarator) result. The radial variable $\rho(\Psi)$ is usually defined in terms of the plasma volume or magnetic flux (integral of $\mathbf{B}\cdot\hat{\mathbf{e}}$ over a surface perpendicular to $\hat{\mathbf{e}}$). For example, $\rho=a(\Psi/\Psi_a)$, $\rho=a(V/V_{\text{tot}})^{1/2}$ or $\rho=a(\Phi/\Phi_{\text{tot}})^{1/2}$, where a is the nominal minor radius of the device, V is the volume and Φ is toroidal flux within a flux surface.

The transport equations themselves are obtained by taking velocity moments of the Boltzmann equation (A-1) and computing appropriate *flux-surface averages*,

$$\langle A\rangle=\frac{1}{V'(\rho)}\int_0^{2\pi}d\phi\int_0^{2\pi}d\theta\sqrt{g(\rho,\theta,\phi)}A(\rho,\theta,\phi), \qquad \text{(A-19)}$$

$$V'(\rho) = \int_0^{2\pi} d\phi \int_0^{2\pi} d\theta \sqrt{g(\rho,\theta,\phi)}, \tag{A-20}$$

where $V' = dV/d\rho$, g is the Jacobian of the coordinate transformation and A is the function to be averaged. The resulting time-dependent continuity and energy transport equations are of the form:

$$\frac{\partial n_a}{\partial t} = -\frac{1}{V'}\frac{\partial}{\partial \rho}\left(V'\langle \Gamma_a \cdot \nabla\rho\rangle\right) + \langle S_{pa}\rangle, \tag{A-21}$$

$$\begin{aligned}\frac{3}{2}\frac{\partial(n_a T_a)}{\partial t} = &-\frac{1}{V'}\left[V'\left\langle\left(\mathbf{Q}_a + \frac{3}{2}T_a\Gamma_a\right)\cdot\nabla\rho\right\rangle\right] - n_a T_a \frac{\partial}{\partial\rho}\left(\frac{\langle\Gamma\cdot\nabla\rho\rangle}{n_a}\right)\\ &-\langle(\nabla\mathbf{u}_a):\Pi_a\rangle + \sum_{a'}\Delta Q_{aa'} + \langle S_{Ea}\rangle,\end{aligned} \tag{A-22}$$

where $\langle\Gamma_a \cdot \nabla\rho\rangle$ and $\langle\mathbf{Q}_a \cdot \nabla\rho\rangle$ are transport fluxes, $\langle S_{pa}\rangle$ and $\langle S_{Ea}\rangle$ are sources of particles and energy, $\langle(\nabla\mathbf{u}_a):\Pi_a\rangle$ is the viscous heating, and $Q_{aa'}$ is the energy exchange between species. There are similar equations for momentum balance. Additional constraints are that fast energy exchange between ions forces their temperatures to be nearly equal and that quasi-neutrality, $\nabla\cdot\mathbf{J} = 0$, determines *ambipolarity* constraints,

$$n_e = \sum_{a=\text{ions}} Z_a n_a,\ \Gamma_e = \sum_{a=\text{ions}} Z_a \Gamma_a \tag{A-23}$$

The fluxes, viscous heating, and sources are inputs to the code and are discussed in more detail in the preceding and following Appendices. Kinetic theory is used to provide *closure*, whereby Γ_a, $\mathbf{Q}_a$, and Π_a are expressed in terms of the thermodynamic variables n_a, T_a and their gradients. For example, the particle flux is usually expressed in the *diffusive-convective* form

$$\langle\Gamma_a \cdot \nabla\rho\rangle = -\left\langle|\nabla\rho|^2\right\rangle D_a \frac{\partial n_a}{\partial\rho} + \left\langle|\nabla\rho|^2\right\rangle^{1/2} n_a u_{na}, \tag{A-24}$$

$$\langle\mathbf{Q}_a \cdot \nabla\rho\rangle = -\left\langle|\nabla\rho|^2\right\rangle n_a \chi_a \frac{\partial T_a}{\partial\rho} + \left\langle|\nabla\rho|^2\right\rangle^{1/2} n_a T_a u_{qa} \tag{A-25}$$

where $\left\langle|\nabla\rho|^2\right\rangle$ is a geometrical factor, D_a and χ_a are the particle and heat diffusivities, and u_{na} and u_{qa} are the convective velocities. The latter may themselves be represented in terms of other plasma gradients and referred to as *off-diagonal elements*.

The magnetic fluxes also evolve. This is typically described in a Grad-Hogan scheme, in which the toroidal flux is taken as the reference frame and is updated with infrequent calls to the MHD equilibrium equation. The poloidal flux evolves relative to the toroidal flux and is described by resistive diffusion, making use of Faraday's law (see Eq. (A-3)) and Ohm's law in a flux function form that relates the toroidal current density as the secondary of a

transformer to the external loop voltage. The time scale is governed by the parallel electrical resistivity.

C. Uses of transport codes

Codes that solve the transport equations as described here form the core of current state-of-the-art integrated modeling. These codes have proliferated over the years to serve a multitude of purposes, and each achieves good results within its range of validity. We distinguish two basic modes of operation. In the *interpretive* mode, all (available) plasma-parameter profiles and their time derivatives are inferred from experimental data, including the equilibrium geometry, which is calculated in an MHD code using magnetic data as constraints. The sources, such as energy deposition, are likewise calculated using the measured profiles. The code then solves (A-21) and (A-22) to calculate the transport fluxes or diffusivities, Eqs. (A-24)-(A-25), which are used to determine empirical confinement scalings and the like. In the *predictive mode*, on the other hand, models for the diffusivities or fluxes are employed, and the code calculates the time evolution of Eqs. (A-21) (A-22) to determine the plasma profiles, which can then be compared with experiment.

The major interpretive codes in use today, TRANSP and ONETWO, contain detailed models of all the principal sources and sinks that are relevant to tokamaks. Each can be run, alternatively, in the predictive mode. Indeed, in the absence of one or more pieces of information, a subset, such as the current or poloidal-field evolution, may be simulated within an interpretive run.

Other codes, such as BALDUR, WHIST, CORSICA, or TSC, are simulation codes whose main use is testing transport models or performing design studies. These carry varying levels of detail outside the particular feature being tested. For example, a core thermal-transport model may be tested by running only the temperature evolution equation, holding the density fixed, setting a boundary condition just inside the edge, and taking the source as calculated in an experimental interpretive run. The goal would be to test the temperature profile evolution without complicating distractions. Similar simulations in which all of the density, temperature, and momentum equations for each species plus current penetration could be run, while still using simple boundary conditions. At the other end of the spectrum, one might wish to explore reactor designs, studying wall conditions, connection to external circuits etc., while employing a very simple plasma model based on empirical scaling. Each of these problems requires different functionality in the code to make efficient use of the available computer time.

D. Challenges of transport modeling

Flexibility. A principal challenge of the simulation initiative is to preserve present functionality within a flexible configuration that allows the user to select from the full range of options and efficiently run his/her case, from the very simplest model to the most detailed full integration.

Rotation and radial electric field. The radial electric field, poloidal and toroidal rotation, and the pressure gradient are all coupled through the radial force balance. Both the physics and the computation of these processes require continued development and will form an essential aspect of integrated simulation. For example, bifurcations are observed experimentally, in which the plasma jumps suddenly to an alternate state, often forming a transport barrier. Enhanced diagnostics and developing theoretical models imply scenarios in which sheared $\mathbf{E}\times\mathbf{B}$ flow suppresses turbulence, while turbulence can drive locally sheared poloidal rotation. Also, $\partial E_\rho / \partial \rho$ squeezes orbits and reduces collisional transport, while toroidal rotation enhances MHD stability. Toroidal rotation generally appears diffusive and governed by turbulent transport but can appear even in the absence of apparent torque. Other issues are whether we should treat E_ρ or poloidal rotation as the independent variable, and whether we can express bifurcation threshold conditions in terms of macroscopic quantities. Internal barriers are strongly localized and move in time, indicating the need for dynamic gridding algorithms for better resolution. Finally, in stellarators, the ambipolarity constraints are expected to govern E_ρ, and in tokamaks, broken toroidal symmetry can do the same.

Transport in the edge and SOL. The basic equations in the edge and SOL are much the same as in the core, but the computational challenges are even greater. Physics and computational challenges include formulation of both turbulent and collisional transport models in steep gradient regions, formulation of the coupling/transition between the closed field lines (slow radial transport timescale) and open field lines (faster parallel timescale), and accommodation of multiple timescale phenomena e.g., *edge localized modes* (ELMs) of the H-mode layer. (See Sec. VIII.)

Non-nested surfaces. Finally, one of the most difficult challenges of integrated core modeling is to extend all the considerations of this and the adjacent two sections to take into account the islands and stochastic regions illustrated in Fig. A10

VII. External Sources

A. Introduction

External systems that add mass, momentum, or energy to a plasma are essential tools in the successful efforts to obtain good performance both in present experiments and in any future experiments and reactors. At present, the external sources include beams of neutral atoms that can carry energy, particles, and angular momentum across magnetic fields; radio frequency waves (RF) whose interactions with a plasma can be used for heating, current drive or flow drive; high-speed pellets of frozen fuel gas that can deliver particles deep into the plasma core; and gas fueling which supplies particles to the plasma edge. In a burning plasma experiment or reactor, the fusion-produced alpha particles, which have some properties in common with beam-injected particles, will provide the principal energy source.

External sources are important, not only to provide fuel (D and T in a reactor) and energy (heating the plasma to 10 keV or more), but also to provide essential elements of control. The

plasma current, pressure, and rotation-velocity profiles all are crucial to plasma performance and can be influenced by these sources. We also have sources such as impurities and gas from plasma-facing components, which can have undesirable consequences, as well as losses from radiation.

The computational capabilities related to neutral beam injection are very well developed and are presently integrated into many codes. The work could be easily ported to codes developed in new initiatives. Wave-plasma interactions at RF frequencies (from ion to electron cyclotron frequencies) are the subject of intense ongoing research including fusion SciDAC activity. However, the scope of needed work, and the ability to provide interactive coupling with other plasma codes, extend far beyond the SciDAC activity. This area includes many distinct problems and will be an essential element of future plasma prediction, interpretation, and control schemes. Work on the ablation and subsequent transport and deposition of fuel from injected pellets is in a relatively early phase of development but certainly amenable to computation. However the physics of fueling in general is not in a satisfactory state at present and would benefit from basic theory studies of particle transport.

B. Neutral beam injection

A number of well understood physics processes are involved in neutral beam injection (NBI) at energies in the neighborhood of 100 keV. The capture of neutral atoms by ionization and charge exchange is modeled by Monte Carlo calculations. A thermal neutral atom formed by charge exchange may escape or be re-ionized. The orbits of fast charged particles are followed and their slowing down and scattering are computed in a Fokker-Planck model solved by Monte Carlo. Secondary processes such as re-neutralization of energetic ions by charge exchange are also calculated.

For tokamaks the TRANSP neutral injection package is the most widely used. It has recently been extracted from the integrated interpretive modeling code TRANSP as a separate module for the National Transport Code Collaboration (NTCC). See the web site http://w3.pppl.gov/NTCC/NUBEAM/UserGuide. Much of interface code is automatically generated by Python script. Input and output is passed in f90-derived data types, employing internal 2-D flux coordinates (ρ, θ), where ρ is the normalized minor radius, and θ is the poloidal angle (the short way around in Fig. III.1 of the main report). Outlines for conversion among other common coordinate representations are included. Inputs consist of MHD equilibrium quantities (magnetic field geometry, plasma pressure in space), profiles of density and temperature of a large number of particle species, descriptions of injected neutral beam geometry and characteristics, atomic physics data and cross sections, wall geometry, and controls for code operation and inputs to auxiliary models. Outputs are radial profiles of power deposition, driven current, ion/electron source rate, and rotation, 2-D profiles of neutral and fast ion density, and the fast ion distribution in four phase-space dimensions (ρ, θ, E, $v_{\parallel}/v$).

The issues for this kind of modeling are more computational than physics: increasing speed, increasing resolution in space and velocity space, and extension to 3-D equilibria for stellarators and other non-axisymmetric configurations.

An entire separate field of study, however, concerns possible instabilities driven by populations of high-energy particles. These are seen experimentally and considerable progress has been made in explaining the observations by theory and computation. These are important because they can lead to anomalous losses of fast particles and reduce heating efficiency.

Alpha particles produced at 3.5 MeV in D-T fusion reactions have many of the properties of high-energy beam particles, resulting in *alpha heating* and including the prospect of driving instabilities, and can be modeled similarly. The chief differences are the much higher alpha energies and their isotropic initial distribution. Because large numbers of alphas will not be produced until a burning plasma is actually achieved, it is extremely important to include this physics in our integrated computation and in this way anticipate burning plasma performance.

C. Edge particle sources

Modeling of edge particle sources is a discipline unto itself and must be included in the edge plasma studies. (See Sec VIII.) Edge conditions in fusion devices have a tremendous influence on the bulk plasma behavior. Neutral particles in the edge also influence other physics processes besides fueling. For example, charge exchange of plasma ions with edge neutrals exerts torque affecting plasma rotation, and neutrals participate in edge instabilities. Edge particle sources are inherently three-dimensional. Several modeling codes presently exist. Most are completely 3-D Monte Carlo (EIRENE, DEGAS). A fast neutral code (NUT), based on zonal integration has recently been placed in the NTCC module library.

D. Radio-frequency wave heating and current drive

Besides inductive Ohmic heating, common to all tokamaks (but usually absent from stellarators), there are three important wave frequency regimes to consider, each associated with a particle or plasma resonance. These are:

- *Electron cyclotron range of frequencies* (ECRF), $\omega \approx \omega_{ce}$.
  f ~ 100 GHz ($\tau \sim 10^{-11}$ sec), $\lambda \sim 0.3$ cm, where λ is the wavelength. The launcher is far from the plasma; the waves propagate in free space and can be computed with geometrical optics.
- *Lower hybrid range of frequencies* (LH), $\omega_{ce} << \omega << \omega_{ci}$.
  f ~ 0.5 - 5 GHz ($\tau \sim 10^{-10}$ sec), $\lambda \sim 1$ cm. The launcher is near the plasma; the waves do not propagate in vacuum but are usually computed with geometrical optics.
- *Ion cyclotron range of frequencies* (ICRF), $\omega \approx \omega_{ci}$
  f ~ 100 MHz ($\tau \sim 10^{-8}$ sec), $\lambda \sim 10$ cm. The launcher is near the plasma, the waves do not propagate in vacuum, and they require solution of wave equation. Alfvén waves at frequencies $\omega << \omega_{ci}$ have also been considered for heating and current drive.

Waves can have strong interactions at localized regions in space due to various kinds of resonance. Either the wave velocity matches the particle velocity $v_{wave} = \omega/k = v_{particle}$, and the particle sees a steady, accelerating electric field, or the wave frequency in a frame moving with the particle (Doppler effect) matches a harmonic of the particle cyclotron

frequency, $\omega - k\mathrm{v}_{particle} = \ell\omega_c$, and the gyrating particle sees a steady component of the electric field. The effect is an energy and momentum kick each time the particle passes through resonance. These kicks accumulate over time to produce energy gain and directed velocity resulting in bulk plasma heating, very energetic tail populations, electric current and fluid drift or flow. In Fig. A5, we have shown a tokamak cross-section with schematic launchers, a wave, and a resonance layer. (Also depicted is a chain of magnetic islands. See Sec. IV.) Similar resonances also occur throughout the plasma when spontaneous turbulence arises.

This illustrates an important difficulty in solving the basic equations, Eqs. (A-1)-(A-3), namely that for perturbations from equilibrium, *the distribution function* $f_a(\mathbf{x},\mathbf{v},t)$ *can be highly structured, requiring very fine resolution in some regions of both configuration and velocity space*.

Let us consider the full-wave calculations as they apply, for example to ICRF heating at frequency ω. The wave propagation and plasma response are described by Eqs. (A-1)-(A-5) with displacement current included. We have now returned full circle to the highest frequency ranges of the whole integrated simulation. For this problem, we exploit the separation of time scales and the small amplitude of the waves to linearize the fields $\mathbf{B}(\mathbf{r},t) = \mathbf{B}_0(\mathbf{r}) + \mathbf{B}_1(\mathbf{r})e^{-i\omega t}$, $\mathbf{E}(\mathbf{r},t) = \mathbf{E}_1(\mathbf{r})e^{-i\omega t}$, and apply the *quasi-linear* method of solution to the plasma response for each species, $f_a = f_0 + f_1 e^{-i\omega t} + f_2$, where f_1 is the fast response producing the wave current $\mathbf{J}$ in Eq. (4.4), and f_2 is the slow time scale response that gives the power deposition and equilibrium evolution. We give a schematic representation of the equations, in which the fast part of Eq. (A-1) reduces to

$$i\omega f_1 + \mathrm{L}_1\{f_1\} = \mathrm{L}_2\{\mathbf{E}_1, f_0\} + \mathrm{L}_3\{\mathbf{B}_1, f_0\}, \qquad \text{(A-26)}$$

where L_1-L_3 are complicated *linear* operators. The slow background response then is the solution to

$$\frac{\partial f_2}{\partial t} + \mathrm{L}_1\{f_2\} = \mathrm{Q}_1\{f_1, \mathbf{E}_1\} + \mathrm{Q}_2\{f_1, \mathbf{B}_1\} + C(f_2), \qquad \text{(A-27)}$$

where Q_1 and Q_2 are *quadratic* or bi-linear operators, proportional to $f_1\mathbf{E}_1$ and $f_1\mathbf{B}_1$, where only the "zero-frequency" components are retained. The set is completed with boundary conditions that connect to the vacuum region and antennas that drive currents $\mathbf{J}_{\mathrm{ant}}e^{-i\omega t}$. To solve these equations, a spectral representation is employed in which the plasma is treated as locally uniform with Fourier expansions in three dimensions.

The wave code with the least restrictive approximations is the All Orders Spectral Algorithm (AORSA). Discretizing using the method of collocation yields equations for the Fourier amplitudes, a gigantic, dense linear system. On the NERSC computer SEABORG, a 3,000-processor IBM SP, a typical run requires 8 hours of processor time at approximately 1.7 teraflops, and memory of 750Mbytes/processor or 1,200 Gbytes total. Inputs are the MHD

equilibrium quantities (magnetic field geometry, plasma pressure in space), profiles of density and temperature of a large number of particle species, distribution functions for non-Maxwellian species (slowing-down distributions or superpositions of Maxwellians), and wave parameters, descriptions of antenna geometry etc. The outputs are Fourier amplitudes of wave fields and plasma response obtained by post-processing: profiles of power deposition, driven current etc., obtained from the moments of f (evolution of $f_0 + f_2$). Calculating wave sources is in itself a task of integrated modeling, as this background evolution, on the same time scale as transport processes, in turn affects the wave absorption. This full integration in calculating wave sources is not presently achieved. A SciDAC project is making progress on integration, but the scope is far short of what is needed

E. Summary: the role of source models in comprehensive integrated simulation

For whole device modeling, source models covering a range in levels of description will be essential elements of any comprehensive device model. We need a 3-D solution of the drift kinetic equation including RF effects, radial transport, and finite orbit width effects.

In focused integration initiatives, source modeling will be essential. A few examples are:

- Particle source modeling is already a key element of edge studies. There are strong interactions between the edge plasma and RF launching structures – RF fields perturb the edge plasma, and edge plasma characteristics influence the wave spectrum launched. This requires integration of RF models (particularly the antenna model) with an arsenal of edge modeling codes.
- In island formation, heating by RF waves at rational surfaces and energetic particles produced during RF heating can stabilize island growth. RF driven currents can either stabilize or destabilize island growth. This requires integration of RF models with MHD and transport near the island (See Sec. VII).
- To understand the role of radial electric fields on transport and flows in stellarators, we need 4-D or 5-D solution of the drift kinetic equation including RF effects.

VIII. Fusion Simulation Present Capabilities Status

This section summarizes the current status of integrated computational modeling and simulation of toroidal confinement fusion systems. The intent in this section in the present document is to provide a general perspective of the status of this very active and mature field. This section responds to the explicit request in the subcommittee's charge letter to report on the status of fusion simulation capabilities, and is the context from which the FSP is defined. The present fusion simulation capabilities form a significant part of the critical underpinning of the FSP.

There are over 50 major toroidal physics design and analysis codes being maintained by the magnetic fusion community. The major multi-user codes are depicted in Figure A11, which shows how they divide into groups, and indicates with arrows the flow of information from one code group to another.

The axisymmetric free boundary equilibrium codes solve the force balance equation by calculating the poloidal magnetic flux in cylindrical coordinates for given pressure and current profile parameterizations. These can be used to define the boundary for the inverse equilibrium codes. There are also two major fully 3-D equilibrium codes in use: VMEC, which assumes the existence of good magnetic surfaces *a priori*, and works in a coordinate system based on these, and PIES, which calculates the existence of surfaces as part of the solution, if they exist.

The collection of linear macroscopic stability codes maintained by the MFE community is quite mature and can assess the stability properties of a given equilibrium with respect to both ideal and non-ideal (resistive) MHD, including the effects of an energetic particle component.

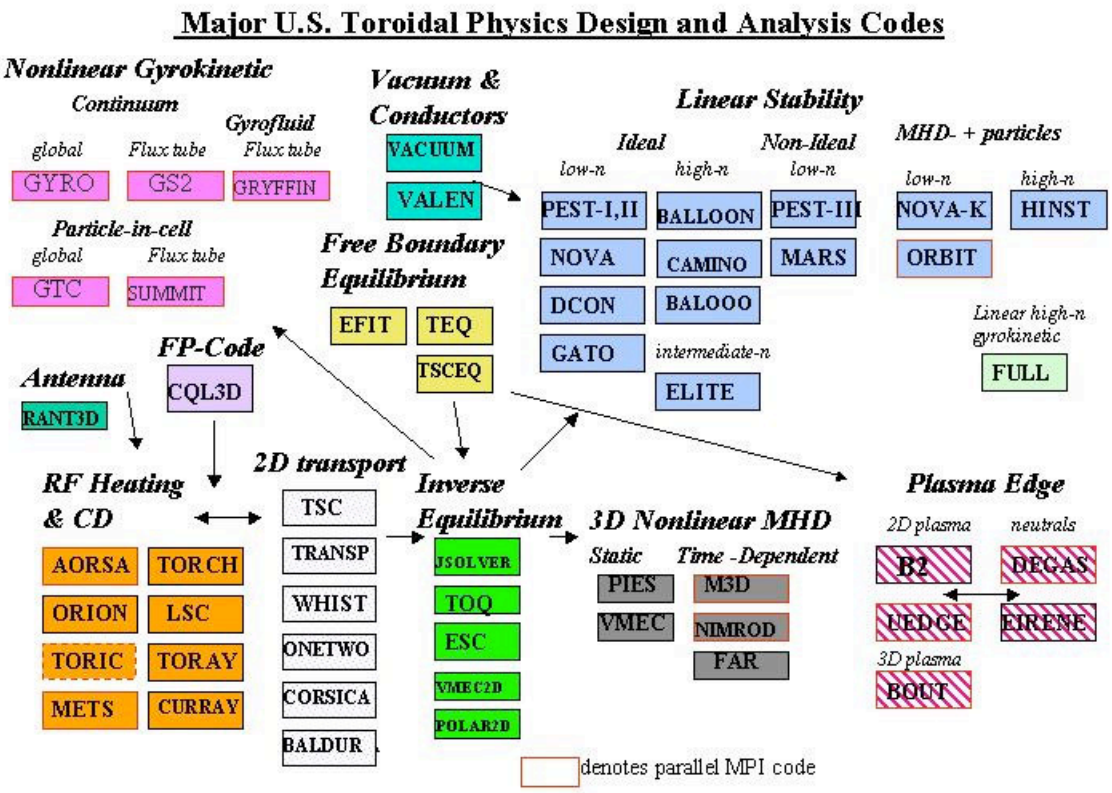

Fig A11. Major U.S. toroidal physics design and analysis codes used by the plasma physics community.

The nonlinear codes fall into four major groupings. In descending order of the frequencies addressed, these are the: 1) RF Heating and Current Drive codes, 2) the Nonlinear Gyrokinetic codes, 3) 3-D Nonlinear Extended MHD codes, and 4) the 2-D Transport codes.

The RF Heating and Current Drive codes calculate the propagation of electromagnetic waves of a given frequency through prescribed background plasma, including reflection and

absorption. The codes are of two major types: ray-tracing (or geometrical optics), and full wave (global solution). There are also depicted antenna and Fokker-Planck codes, which are closely coupled with the RF codes, and provide boundary conditions and background distribution functions. The RF codes are designed to calculate the instantaneous heating and current-drive profiles for a given plasma equilibrium subject to a given RF oscillator source and antenna.

The gyrokinetic codes are based on an analytic reduction of the full 6-dimensional plus time plasma distribution function obtained by averaging over the rapid gyro-motion of ions in a strong magnetic field, and by neglecting the displacement current in the Maxwell equations to remove "light waves" from the system. These codes are appropriate for studying 3-D turbulent transport in a background system with fixed profiles.

The 3-D nonlinear Extended MHD codes are based on taking velocity moments of the Boltzmann equation to yield 3-D magneto-fluid equations for the evolution of the average plasma velocity, density, and pressure, along with a closure procedure. These codes are appropriate for describing global stability phenomena such as sawteeth oscillations, magnetic island evolution, and plasma disruptions.

The 2-D transport codes presently form the core capability in our community for integrated modeling. There are six major codes, with considerable overlap, that exist largely for historical reasons. These codes are all based on the Grad-Hogan evolving equilibrium description where the inertial terms in the momentum equation are neglected and the remaining MHD equations are averaged over the flux surfaces, where they exist.

The 2-D transport codes are all very modular. They are each a collection of equilibrium modules, transport modules, solvers, and source and sink modules representing Neutral Beam Injection (NBI) and RF heating, pellet and gas injection, impurity radiation, and the effects of saturated MHD activity such as sawteeth and islands. These codes have recently benefited from the National Transport Code Collaboration (NTCC), which has formed a modules library so that modules taken from individual codes can be exchanged and shared (see, e.g., w3.pppl.gov/NTCC for more details on this). While these codes address integrated modeling, the individual modules represent simplified reduced descriptions of the full three-dimensional physical phenomena being modeled.

A summary of these four major code groups and the timescales being addressed by them has been given in Fig. A6. The RF codes address frequencies of order the ion cyclotron frequency, Ω_{ci} , and above, up to the electron cyclotron frequency Ω_{ce}. (Again, recall the interchangeable notation for these two quantities, ω_{ci} and ω_{ce}.) The gyrokinetics codes typically take time steps about 10 times longer than the ion cyclotron frequency, whose motion is analytically averaged over, and are normally run for 10^3 to 10^4 time steps to calculate stationary turbulent fluctuation levels. Recent additions to these codes to include some electron timescale phenomena bring in the electron transit time, which lowers the maximum time step. The Extended MHD codes need to resolve phenomena occurring on the Alfven transit time, τ_A, although most codes are at least partially implicit to avoid a strict restriction on the time step based on this. These codes can normally run 10^4 to 10^5 time steps

to address MHD phenomena such as sawteeth and island growth. The 2-D transport codes are very efficient and can take many long time steps to model the entire discharge. However, the 2-D edge transport codes need to resolve the parallel dynamics and use a time-step based on the parallel sound-wave propagation near the edge region.

The calculations now being performed with the gyrokinetics codes, the Extended MHD codes, and the RF codes, are straining the limits of the existing computing capabilities and capacities. For example, recent attempts by the core-turbulence gyrokinetics codes to include both electron and ion dynamics in a self-consistent simulation require upwards of 10^4 processor-hours (over one processor-year) on the IBM SP3 at NERSC to generate one result for a set of fixed background profiles. Similar times are required by the Extended MHD codes to calculate the growth and self-consistent saturation of a neoclassical tearing mode. Thus, we can take solace in the fact that the capability is mostly in place, but must deal with the fact that the computational requirements for a fully integrated 3-D comprehensive simulation capability are truly daunting.

Examples of fundamentally important experimental phenomena that involve 3-D physical processes that cross theoretical boundaries and thus cannot adequately be addressed by the present suite of above-described codes include:

Pedestal physics – A description of the transport barrier that forms in the region of the plasma between the core and the edge, and of the associated edge localized relaxation events;
Long time scale profile evolution – A way to self-consistently evolve the global profiles of plasma temperature and density on the energy-confinement time scale from turbulent transport and in the presence of magnetic islands and other MHD phenomena.
Edge transport: A description of long-mean-free-path particle and heat transport outside the closed magnetic flux surfaces, on the open field lines that impact the first wall or divertor and involve multi-phase physics
Self-consistent heating and current drive: A fundamental model of the interaction of Radio Frequency (RF) waves with plasma in the presence of plasma turbulence.
Sawtooth phenomena – Internal MHD-type modes in the hot core of tokamak plasmas for which fast ion and kinetic effects are clearly relevant experimentally but are only beginning to be addressed computationally.
Island physics: incorporation of the effect of 3-D island formation on equilibrium evolution and turbulent transport

These are but some of the important problems to be addressed by the integrated simulation initiative.

IX. Focused Integration Initiatives – Challenges and Opportunities

The following Subsections A-D expand upon the descriptions of the candidate focused Integration Initiatives (FIIs) described in the main text. These are based on reports from working groups at the September 23-24, 2002 Community Workshop, each concentrating on a particular FII.

A. Plasma Edge

The boundary or edge-plasma of fusion devices plays a vital role in their operation. The edge-plasma region is generally considered to be the region where substantial two-dimensional (2-D) or 3-D variations can occur in the plasma, neutral particle, and magnetic equilibrium quantities. In addition, owing to the lower plasma temperature and proximity to material surfaces: neutral gas, sputtered impurities, and atomic line-radiation can become important components. Thus, a rich variety of potential interactions can take place in this region.

The ISOFS edge discussion group identified four elements that are key to successful operation of an MFE fusion device:

(1) predicting conditions and properties of the pedestal energy transport barrier just inside the magnetic separatrix;
(2) understanding plasma/wall interactions for particle recycling and wall lifetime from high energy fluxes;
(3) controlling tritium inventory including co-deposition; and
(4) controlling wall impurity production and transport into the plasma.

A number of models of varying sophistication exist to describe these processes. Some models already provide a level of coupling, e.g., hydrogen transport, recycling neutrals, impurity sputtering, and impurity transport codes. However, many of the constituent models need improvement, and much more inclusive couplings are required to self-consistently predict the edge-plasma behavior.

In setting priorities for the edge region, it was felt that the short-term (five-year timescale) work should focus on a better understanding of what controls the suppression of plasma turbulence to produce a transport barrier in the pedestal region (#1 above), and its associated impact on plasma profiles. The ability to predict the behavior of the edge pedestal barrier is essential for projecting the net fusion output of MFE devices as discussed extensively at the Snowmass 2002 Fusion Study Workshop. Presently, the key parameter that is believed to control the core fusion output is the plasma temperature at the top of the pedestal; this parameter is now either extrapolated from existing experiments or assumed.

The development of a predictive model of the pedestal requires two key advances. The first is the inclusion of kinetic effects in the 3-D turbulence simulations such that the physics model can span the region from the nearly collisionless portion at the core side of the pedestal to the more collisional region near and outside the magnetic separatrix. Presently, turbulence codes fall into two classes: one class is the collisional fluid codes allowing three spatial dimensions and including the strong magnetic shear region in the edge; the other class is the kinetic particle or continuum codes using three spatial and two velocity-space dimensions to capture the physics of plasma with long mean-free paths, but allowing only very simple plasma geometry. These codes are very compute-intensive, requiring parallel computers with run times of many hours or days (> 2000 node-hours). The computational challenge here is to

obtain a generalized code or coupling of codes that can treat the full range of collisionality through the pedestal for geometries relevant to the plasma edge.

Another component needed for the pedestal is the description of periodic edge profile relaxations called Edge Localized Modes (ELMs), which are believed to be caused by long wavelength MHD instabilities. The transport barrier is intermittently interrupted by these modes. The peak heat flux to material surfaces from such events is a serious concern for a reactor-sized device. Inclusion of such modes in the microturbulence codes should be considered, but requires a larger spatial mesh and an extension of the physical model to capture this MHD phenomenon. An alternative approach is to use a separate code for the turbulence from MHD modes, or to use a linear stability code to enforce marginal stability of the plasma pressure profile for long-time transport evolution.

A second key advance thought possible in the five-year time frame is to couple the turbulent plasma fluxes with a transport calculation of the profile evolution. Such coupling is the heart of a predictive model. This is the plasma-edge analogue of the "Turbulence on the Transport Timescale" FII described in the next section. The turbulence drives plasma fluxes that, together with sources such as neutral ionization, largely determine the plasma profiles; since these plasma gradients provide the free energy to drive the turbulence, it is essential that the turbulence and profile evolution be coupled together. Plasma-wall interactions of sputtering and recycling can make the profile evolution time substantially longer than the turbulence saturation time. Devising stable and efficient coupling schemes is a major task here, which may have spin-off to related couplings for other regions of the fusion plasma and beyond. Experiments and turbulence simulations often show large, intermittent transport events that need to be properly characterized.

Beyond the five-year time frame, edge-plasma modeling should work to develop and couple more detailed models of the plasma-wall interactions. These include models for sputtering, recycling and re-deposition, and transport of neutral gas impurities and radiation. In addition, the edge model should look inward to couple with core physics inside the pedestal region. A vision is that a successful edge model could eventually incorporate the core region models for a whole-device simulation.

B. Turbulence on Transport Timescales

The nature of the problem to be considered can be summarized as follows: The "anomalous" transport of mass, energy, and angular momentum in toroidal MFE devices is dominated by fluxes driven by plasma turbulence. There is a significant disparity of scales, especially timescales. This is a highly coupled system: the plasma pressure and densities are very nearly constant on magnetic flux surfaces, and can thus be considered to be one dimensional (1-D) functions of the flux coordinate ψ; one-dimensional transport equations can be constructed to describe the evolution of these profiles, whose gradients drive turbulence, while the turbulence (3-D, anisotropic) produces the fluxes that drive the evolution of these 1-D profiles.

The object of this Focused Integration Initiative (FII), is to bridge the range of temporal and spatial scales so as to compute the above coupled system self-consistently, *as opposed to* just computing 3-D fine-scale turbulence with fixed background profiles, or computing 1-D transport with highly reduced theoretical or empirical models of the turbulent fluxes.

The single overarching science issue and goal is the self-consistent calculation of global confinement from first-principles physics. The initial (easier) focus would be to determine steady-state confinement; following time evolution on the transport timescale is conceptually no more difficult but is more computationally demanding. The achievement of the steady-state goal would, as a side benefit, enable optimization studies. Important issues like simulation of both steady-state and time-dependent versions of internal transport barriers are subsets of the overall goal.

There are a number of candidate approaches for achieving the goal. These include:

(1) direct coupling of the transport and turbulence equations;
(2) an expert system working together with a smart database of past turbulence simulation results and/or analytic models for transport coefficients, which detects when transport has moved outside of the domain of applicability of the database and triggers new turbulence simulations to extend the domain of applicability;
(3) developing a viable gyrofluid closure and directly integrating the resulting fluid equations; and
(4) projecting long-timescale evolution from moments of the extrapolation of particle weights in a gyrokinetic particle-in-cell code.

This candidate FII entails many disciplines in physics and computation. It also entails a number of generic integration issues, including: coupling of disparate timescales; coupling of descriptions with different dimensionality; interfacing to experiments for validation; the need for a flexible framework/environment to support experimentation with multiple approaches; visualization; expert systems to detect failure of an integration component; software standards; and database management. All of these issues present opportunities for collaboration between the theoretical/computational plasma physics and the computer science/math communities.

Such an initiative would also need a strong theory component. Most importantly, there needs to be a parallel effort in analytic turbulence theory to provide crosschecks and understanding. Theory will also be needed for the gyrofluid approach and other formalism extensions, and for development of synthetic diagnostics.

There are clearly opportunities for further linkages and progress toward a full integrated simulation. Turbulence and MHD dynamics may need to talk to one another directly as opposed to via the intermediary of a transport code, but this coupling may be a logical extension of the transport-turbulence coupling activity, as the necessary flows of information are similar. Another logical extension is coupling of the turbulence equations to real sources. But, perhaps most importantly, if the transport-timescale part of transport-turbulence coupling is handled by an existing or new (developed under the initiative) full-device

integrated modeling/transport code, this activity immediately has access to many other developments in whole-device integration.

C. Global Stability

Global stability issues play a central role in determining the optimal operating regime of fusion devices, and in describing their time evolution. It is well known that under some operating conditions, an experimental discharge can spontaneously transform from a symmetrical stable system exhibiting good confinement into one that exhibits symmetry-breaking oscillations and poor confinement or becomes unstable and disruptive. These events are known as sawtooth oscillations, Neoclassical Tearing Modes (NTM), disruptions, Edge Localized Modes (ELMs), or Resistive Wall Modes (RWM), among others. The predictive calculation of the onset and evolution of these events is the over-arching scientific question in this FII.

The long time evolution of tokamak discharges is often described by nested magnetic flux surfaces and two-dimensional axisymmetric force balance coupled with diffusive transport of the one dimensional surface averaged thermodynamic and field quantities. The transport coefficients used in the diffusive model are typically analytic approximations to the results of three-dimensional, sub-grid-scale turbulence that is described by kinetic theory. The properties of the turbulence are determined by the axisymmetric plasma profiles, which in turn are affected by the transport coefficients. The first-principles coupling of the transport and kinetic turbulence models is a formidable problem requiring integrated modeling as described in the previous section.

However, the force balance assumed in the transport model can occasionally become unstable and yield dynamics that evolve on an intermediate time scale that is much shorter than the transport time scale, but much longer than the time scale of the kinetic turbulence. These motions result in three-dimensional magnetic perturbations that break both the underlying axisymmetry of the overall configuration and the nested topology of the magnetic field. Their nonlinear evolution can strongly affect the confinement properties of the magnetoplasma system, and can re-arrange the plasma profiles on time scales faster than that described by diffusive transport.

At relatively low temperatures, these dynamics are well described by *resistive magnetohydrodynamics* (MHD), a mathematical model in which the plasma is treated as an electrically conducting fluid subject to electromagnetic body forces. Solutions of this model are complicated by a wide separation of space and time scales, and by the inherent high degree of anisotropy that occurs in a magnetized plasma. At the higher temperatures that occur in modern tokamaks, kinetic effects both parallel and perpendicular to the magnetic field introduce important physical processes that can affect the global MHD evolution of the plasma. The challenge is to develop mathematical and computational models that include these kinetic effects while retaining the computational tractability of the fluid model. Such models are collectively called *extended MHD*.

There have been two approaches to developing extended MHD models, as described in Section IV. One is to take analytic moments of the higher dimensional drift kinetic equation to obtain expressions for the higher order fluid *closures* (e.g., the heat flux and the components of the stress tensor) that capture the kinetic effects within the fluid model. The other is to solve the kinetic equation by sub-cycling within a time step of the MHD model. Velocity moments of the resulting distribution function can be taken numerically to calculate the required closures. While this approach may be sufficient for minority (or possibly even majority) ion species, it difficult to envision it being applied to electron dynamics because of their small mass and rapid dynamics. There it is likely that we must continue to rely on analytic closures.

A full modeling of kinetic effects on MHD evolution, and capturing the effects of the MHD relaxation of the profiles within longer time scale transport calculations, requires an integrated simulation. For example, consider the long time scale modeling of a burning plasma configuration. The axisymmetric profiles evolve in response to the turbulent transport and the sources of mass, momentum, and energy from RF antennas and neutral beams, which are in turn affected by the evolution of the profiles. Throughout the simulation the stability of the configuration can be monitored. The stability is be affected by the profiles and the presence of the energetic alpha particles that are produced by the fusion reactions. When the configuration becomes unstable it can be used as the initial conditions for an extended MHD simulation. This model would predict the dynamics of the three-dimensional magnetic perturbations and determine their effect on the underlying profiles. In a burning plasma, this calculation would require integrated models for majority ion and electron damping, and for minority (alpha particle) dynamics. The resulting modified profiles could then be returned to the transport calculation for further evolution.

To make progress on this Focused Integration Initiative would require the talents of theoretical and computational plasmas physicists, applied mathematicians, and computer scientists. The existing Extended MHD codes need to be enhanced to include higher resolution and a more rigorous mathematical framework for coupling MHD events with both microscopic effects and with long timescale evolution of profiles, eventually including such effects as plasma rotation and edge effects. The applied math community will be called upon to provide novel methods to deal with the time and spatial scale separation and with the extreme anisotropy; for example moving or adaptive grids, high-order or spectral elements, and nonlinear equation solvers. Since there is no agreed upon “best” method for solution of the Extended MHD equations or for the couplings, there is a need for a computational framework that allows rapid prototyping. Many of the problems faced by this initiative are generic to the larger integration problem. Indeed, since MHD phenomena are intermediate between very fast and very slow phenomena, this FII is a microcosm of the entire project.

D. Whole Device Modeling

The distinguishing feature of the Whole-Device Modeling (WDM) FII is that from the outset it will provide a model of the entire device for the whole discharge timescale. Because of this scope, many of the models of individual systems are necessarily very simple. It is envisioned

that these simple models will be capable of being replaced by more complete and accurate models as they become available and as is warranted by the application.

The state-of-the-art of whole-device complete-shot modeling at present is represented by the national array of 1-1/2-D transport codes described in Secs. VI and VIII. Those codes have many features that would be required for a final product WDM. They employ a formal separation of time-scales between the rapid (Alfvén) time on which 2-D magnetic equilibria are established, and the much slower time on which heat, particles, and angular momentum, are transported as 1-D surface functions across the magnetic surfaces. They also incorporate many features of a WDM: a hierarchy of models to describe particular aspects of physics, with trade-offs between speed and accuracy; connection to experimental databases; predictive and interpretive modes. Some are also equipped with sophisticated user interfaces.

There are at least three distinct thrusts in a WDM FII. The first is to extend the accuracy and reach of the physics modules available in the existing codes; the second is to extend these codes to fully 3-D geometry so they are applicable to stellarators; and the third is the development of a suitable modern computing framework architecture that would allow this effort to couple to the fruits of the other FIIs for the final WDM code.

The first of these involves the development and application of algorithms for coupling "best physics" modules to the surface-averaged long-timescale transport equations described in Sec. VI. The scientific goals of such couplings are many and diverse: the accurate prediction of heating, fueling, current drive, and confinement on the transport timescale; the ability to account accurately for sawteeth, tearing, and wall modes, in order to assess, avoid, or control them; increased understanding of the complex interaction between the plasma edge and core confinement; and the ability to model and develop machine feedback and control systems within a simulation code environment. Many of these capabilities exist already at some level, and this activity can be viewed as extending these and taking them to the next level. Extensions would include incorporating rotation and non-Maxwellian distributions into the equilibrium, and self-consistently incorporating multiple heating and current-drive systems into the formulation.

The second major thrust involves coupling 3-D plasma equilibria and device geometry to 1-D surface-averaged transport and the many source and boundary modules that make up a Whole Device Modeling code. This can be viewed as an extension of the present 2-D equilibria + 1-D surface-averaged transport codes, but it is a non-trivial extension. The calculation of 3-D equilibrium is itself a research topic, and issues of existence of surfaces, magnetic islands, and field line stochasticity must be dealt with on both a fundamental mathematical and a practical computational level. It will also be necessary to extend current methods to include plasma rotation and other non-ideal effects. Developing a system to define the complex 3-D input geometry and to manage the extensive output is itself a challenge.

The third major goal of a WDM FII is the design of the code architecture for the final WDM product. This is obviously critical in many respects. Although it may not speed science results within the five-year timeframe of the FIIs, it is important for the success of the global

initiative that the framework itself gets a timely start. This thrust, in particular, will rely heavily on computer scientists to develop an extensible framework that meets the many needs of the projects.

By its nature, WDM is an integrated activity, and initiatives in this area will overlap other FIIs. While development and testing of particular couplings will be carried out as narrowly focused projects in the other FIIs, integration into a WDM code is a practical requirement for a close connection to experiment, or for an ability to simulate a proposed new machine on transport timescales. It is also possible for a WDM code to serve as the 1-D transport solver throughout the development of any of the new couplings. This suggests coordinating the WDM FII closely with the turbulence, the MHD, and the edge FIIs.

X. Validation Requirements in Fusion Topical Areas

What follows are examples of phenomena that one could anticipate being important for validation tests of fusion simulation codes in various topical areas, with particular regard to possible Focused Integration new capabilities. This is not a statement of the outstanding physics issues - but a list of some calculable and measurable (at least in principle) quantities that are important for testing. They are described beginning with the most general and global to the most detailed and local and most stringent. This ordering also reflects the difficulty in making such measurements.

A. Transport

The coarsest level of agreement between a transport simulation and an experiment would be to match the total stored energy given the input power and other machine parameters. This is only an interesting comparison if the model has few if any free parameters that have been calibrated against existing databases. One should not expect the scaling laws themselves to be derivable as the engineering parameters used are only proxies for the relevant physics variables. A higher level of agreement could be assessed by comparison of profiles given fluxes or models for the sources. At a minimum, one is interested in the width and height (or gradient) of the edge pedestal and core temperature gradients. Eventually this comparison must be extended to all transport channels, ion and electron thermal, particle (including impurities) and momentum transport. Recent work by ITER expert groups have shown however that a wide range of codes, based on physical, semi-empirical, or wholly empirical models can achieve essentially identical performance in matching temperature profiles for a moderately large set of discharges from a variety of machines. This suggests that agreement at this level is insufficient to validate a particular code. More challenging are comparisons of transient behavior including thresholds and dynamics of transport barriers. Many widely observed features of transient transport cannot be explained by the current generation of models. Ultimately, transport models **must** be validated at the level of turbulence and turbulence dynamics - comparisons that will be paced by the development and deployment of fluctuation diagnostics and analysis techniques. Adding to the difficulties is the prediction that fluctuations vary significantly over the poloidal cross section. Synthetic diagnostics,

numerical analogies to the experimental diagnostics need to be developed and adapted to simulation code outputs.

B. MHD

The first test of MHD codes is their ability to reproduce experimental stability limits or operational boundaries for ideal and resistive instabilities, including kinks and ballooning modes, edge localized modes (ELMs) and other edge relaxation phenomena, and resistive wall modes. Especially interesting are modification to the stability boundaries from non-ideal or non-linear effects. A second parameter for comparison is the growth rates of unstable modes, first in their linear phase then in the non-linear phases including the calculation of the saturation levels for these modes. More generally one can look for agreement with the computed eigenvectors and eigenvalues including non-linear mode structures. For example, one should be able to predict various ELM types and their non-linear extent. Large-scale dynamics of disruptions can be compared including halo and eddy current distributions in real machine geometry, runaway populations and so forth. Dedicated experiments should also test explicitly the extensions to ideal MHD including neoclassical, two-fluid, flow, finite Larmor radius and other kinetic effects and non-linear interactions with profiles and transport.

C. Radio-Frequency Heating and Current Drive

Validation of RF models is challenging, as the important quantities are particularly difficult to measure. Testing begins with global quantities like overall heating and current drive efficiencies and proceeds to comparisons with deposition profiles for heat, current, and flow velocity. Beyond these measures, an essential element is the verification of RF waves inside the plasmas. To the extent possible, the two-dimensional fields of wave amplitude and wave number should be measured and compared to code predictions. The position of mode conversion layers should be verified along with the propagation of the outgoing waves. Antenna/edge-plasma interactions and the influence of plasma fluctuations on launched waves will challenge both simulations and measurements. Finally, one will need to compare codes with experiments that test the models of wave-particle interactions. For this, one would need to measure velocity-space distributions along with the wave fields and plasma profiles.

D. Edge/Scrape-Off-Layer/Divertor

Edge modeling involves a wide range of physical effects including those common with the core like transport and MHD as well as neutral dynamics, atomic processes and plasma-material interactions. Fortunately the diagnostic challenge is not quite as severe in the edge plasma as it is in the core. Assessment of edge transport models is similar to that in the core with comparisons involving profile and fluctuation measurements, though with perhaps greater emphasis on poloidal variation. The energy source from the core to the edge is expected to vary strongly around the poloidal circumference leading to predictions of flows and potentials that must be verified. As in the core, profiles from energy, particle, impurity, and momentum transport should all be compared along with the appearance of self-generated flows. Unlike the core, the sources and sinks for particles, including impurities, are not well

characterized. The position and processes for these sources needs to be included in models and tested by experiments. Deposition of impurities and co-deposition of hydrogenic species is a particularly important issue for benchmarking. The role of neutrals is complex, affecting all transport channels through both classical and turbulent processes. The computed three-dimensional distribution of neutrals should be verified and the role of various neutral transport mechanisms tested. A basic model for fueling including sources, neutral dynamics, particle transport needs to be developed and compared with experiments. Finally, important interactions with MHD physics in phenomena like ELMs and the tokamak density limit provide particularly stringent tests for integrated models.

XI. Background – Orbits and instability mechanisms

Here, we explain some of the basics of plasma motion embodied in Eq. (A-1). For a more complete explanation, one should consult any plasma physics textbook. The basic equation, describing the motion of a charged particle in electric and magnetic fields, reflecting the force term in Eq. (A-1) is

$$m\frac{d\mathbf{v}}{dt} = q\left(\mathbf{E} + \mathbf{v}\times\mathbf{B}\right), \tag{A-28}$$

where q is the particle charge. The basic motion in a uniform magnetic field is gyration about the field line with frequency $\omega_c = qB/m$ and radius $\rho_c = v/\omega_c$. Additional forces and non-uniformities produce drifts in a direction perpendicular to both the force and the magnetic field. Thus, an electric field **E** or a gradient ∇B perpendicular to **B** reduces the orbital radius on one-half cycle and increases it on the other, resulting in

$$\mathbf{v}_E = \frac{\mathbf{E}\times\mathbf{B}}{B^2}, \text{ and } \mathbf{v}_d = \frac{v_\perp^2}{2\omega_c}\frac{\mathbf{B}\times\nabla B}{B^2}, \tag{A-29}$$

the **E**×**B** *drift* and the *grad B drift*, respectively. If the magnetic field is curved, a similar expression is obtained for the curvature drift, which is proportional to $v_\parallel^2$. In addition to these particle drifts there are fluid drifts in confined plasmas (i.e., plasmas with radial density and temperature gradients), owing to the fact that at any one radius, there are more particles moving in one direction perpendicular to the gradient than in the reverse direction. This yields a net fluid velocity known as the *diamagnetic drift*

$$\mathbf{v}_a^* = \frac{\rho_a v_a \mathbf{B}\times\nabla n_a}{B n_a}, \tag{A-30}$$

where v_a is the thermal velocity of species a. A similar drift proportional to the temperature gradient is also important in the kinetic theory. For the orderings employed in plasma computation, an important feature of all these drifts is that for thermal particles they are smaller than the thermal velocity by a factor $\rho_a/L << 1$, where L is a macroscopic length, e.g. of order ρ^*. All plasma waves, instabilities, and transport phenomena are related in some way to these drifts, which are illustrated in Fig. A12.

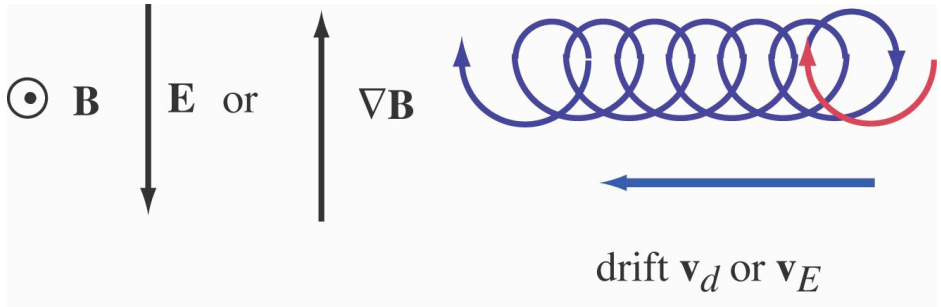

Fig. A12. Ion drifts in electric and magnetic fields.

The existence of waves and instabilities is closely related to the fact that the grad B drift and the diamagnetic drift have opposite signs for ions and electrons (we are taking ρ_a and ω_{ca} to have the sign of q_a), while the **E×B** drift is the same for each species. For example, Fig. A13 provides a simple picture of instability in *bad curvature*. If a perturbation forms in a region of high pressure gradient, then curvature and grad B drifts cause a charge separation. The resulting electric field drives **E×B** drifts that amplify the perturbation.

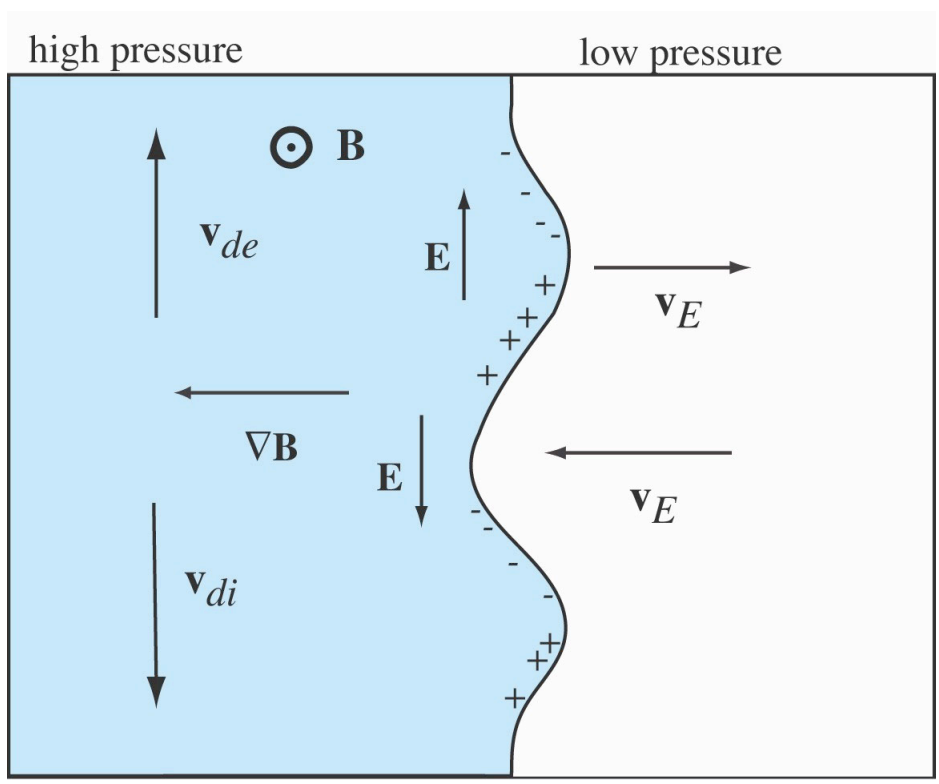

Fig. A13. Bad curvature instability.

XI. Glossary

Alfvén wave: A plasma wave, which involves bending or compression of the magnetic field. The Alfvén time scale is the interval for an Alfvén wave to traverse the plasma. Magnetosonic waves are a type of Alfvén wave.

Alpha heating: In a fusion power plant, energetic alpha particles and neutrons are created by the fusing of deuterium and tritium nuclei. As a charged particle, the alpha particle is unable to cross the confining magnetic field and gives up its energy as heat to the plasma.

Ambipolarity: Mass transport in plasmas is ambipolar, that is the fluxes of electrons and ions are virtually the same. The densities of negatively and positively charged particles that compose a plasma are almost in perfect balance, leaving the plasma essentially neutral (termed quasi-neutrality.)

Aspect ratio: In a toroidal device, the ratio of the major radius to the minor radius.

Avalanche. A sudden macroscopic event in which energy or particles can be distributed across the medium. Often described by simple mathematical models.

Bad curvature: see Curvature

Ballooning: A local instability, which can develop in the tokamak when the plasma pressure exceeds a critical value. It is analogous to the unstable bulge that develops on an over-inflated pneumatic inner tube.

Banana: The shape of a trapped particle orbit (banana orbit) projected onto a poloidal plane. The shape results from drifts away from a magnetic surface.

Beta: The ratio of plasma pressure to magnetic pressure. An essential dimensionless parameter for magnetized plasmas. Denoted by the symbol β.

Boltzmann equation: An equation of motion in phase space (position and velocity). The Boltzmann equation (and its collisionless version, the Vlasov equation), is the starting point for much of plasma physics and is the fundamental equation for kinetic theory.

Bootstrap current: Currents driven by collisional transport effects in toroidal plasmas (see neoclassical).

Braginskii equations: Plasma fluid equations separately describing electron and ion motion. Fluid equations are obtained by calculating velocity moments of the Boltzmann equation and specifying a closure condition.

Closure: Mathematical scheme by which a hierarchy of equations is truncated. This usually involves expressing higher moments in velocity space in terms of lower moments.

C-Mod: A compact high field tokamak experiment

Confinement: Property of magnetic fields in preventing loss of energy and particles; degree or measure of this property, as in "confinement time," for example.

Curvature: In toroidal configurations, field lines are necessarily curved, giving rise to particle drifts. In *bad curvature* or *unfavorable curvature* regions, these drifts give rise to instability (see Fig. A13). In *good curvature* regions, the drifts are stabilizing, owing to the reversal of the relative direction of curvature and pressure gradient. The *grad B drift* acts in a similar way.

Cyclotron frequency; same as gyrofrequency.

Detached plasma: Cold dense plasma in the divertor chamber, separated from the walls by neutral gas. This is characteristic of a particular mode of divertor operation.

Diamagnetic drifts, waves, etc: Refers to plasma dynamics driven by the effect of a plasma pressure gradient in a magnetic field.

DIII-D: A medium scale, strongly shaped tokamak experiment.

Disruption: The abrupt termination of a tokamak plasma through the growth of large amplitude MHD instabilities. Control of disruptions is a critical problem for this type of confinement device.

Divertor: Region outside the plasma core with open field lines leading to a chamber some distance from the plasma.

Drift: The motion of a particle or fluid when subjected to a force or gradient perpendicular to the magnetic field. The drift is perpendicular to both the field and the other force. Examples are *curvature drift*, *grad-B drift*, *E×B* drift, and *diamagnetic drift*. See Appendix A8.

Drift wave: A type of plasma wave arising from the presence of density and temperature gradients across magnetic field lines. Turbulence of various types of drift waves is believed to be responsible for anomalous transport in toroidal plasma experiments. *ITG* and *ETG modes* are examples of drift waves.

ECRF: Electron Cyclotron Range of Frequencies – Refers to radio frequency heating and current drive using waves close to the electron cyclotron frequency.

ELM: Edge Localized Mode – a phenomenon that relaxes the pressure gradient of a plasma device operating in H or High confinement mode.

Equilibrium: Usually refers to MHD equilibrium – a steady state solution of the MHD equations. While confined plasmas may be approximately in local thermal equilibrium, they are not in global thermodynamic equilibrium.

ETG: Electron Thermal Gradient (modes) – A type of "micro" plasma instability that may be responsible for turbulent transport by very short (electron cyclotron radius) fluctuations.

Field lines and flux surfaces: Imaginary lines marking the direction of a force field. These map out surfaces (to which plasma particles are approximately constrained) called flux surfaces.

Finite Larmor radius: FLR. A mathematical expansion or approximation in which terms of order the Larmor radius (gyroradius) divided by a macroscopic scale are retained.

FIRE: A proposed burning plasma experiment, which would operate with very high magnetic fields.

FLR: See finite Larmor radius.

Flux: See Magnetic flux, Transport flux

Good curvature: See Curvature.

Grad-Shafranov equation: A steady-state MHD equation whose solution yields axisymmetric equilibria.

Gyrofluid: A set of fluid equations that treat ions and electrons separately and retain certain kinetic effects through their closure conditions.

Gyrofrequency: the frequency at which a particle gyrates around a magnetic field line; also called *cyclotron frequency*.

Gyrokinetic: Kinetic equations derived from the Boltzmann equation, which is averaged over the fast cyclotron (gyro) motion of the plasma particles reducing the number of spatial dimensions from 6 to 5.

Gyroradius: the radius at which a particle gyrates about the magnetic field; also called *Larmor radius*.

Hall term: A term in the extended Ohm's law in MHD, proportional to the crossed current and magnetic field. The presence of this term, often neglected in simpler descriptions, leads to Whistler waves and can strongly influence the rate of magnetic reconnection.

H-Mode: High (confinement) Mode – an experimental regime with a transport barrier or pedestal at the edge of the plasma.

IBW: See Ion Bernstein Wave.

ICRF: Ion Cyclotron Range of Frequencies - Refers to radio frequency heating, current and flow drive using waves close to the ion cyclotron frequency.

Ion Bernstein Wave (IBW): – short wavelength electrostatic waves in the ion cyclotron range of frequencies. Generated in the plasma by mode conversion.

Ion cyclotron wave: A plasma wave that propagates at frequencies above the ion cyclotron frequency. See ICRF.

IPPA: Integrated Program Planning Activity – the current roadmap for the fusion energy program.

Island – Magnetic Island: A three dimensional magnetic structure arising when closed flux surfaces are perturbed by magnetic fields from plasma fluctuations or from errors in the vacuum fields.

ITB: Internal Transport Barrier – a regime of strongly reduced transport over some part of the core plasma.

ITER: International Thermonuclear Experimental Reactor – A burning plasma experiment proposed as the next step in the international fusion program.

ITG: Ion Thermal Gradient (modes) – A type of "micro" plasma instability that may be responsible for turbulent transport by medium scale (ion cyclotron radius) fluctuations.

JET: Joint European Torus – a large Tokamak experiment run jointly by a European Community consortium.

Kinetic: Refers to aspects of plasma physics related to velocity distributions of electrons and ions.

Kink Instability: A macroscopic (large scale) instability driven by plasma currents.

Larmor radius: Another term for the cyclotron orbit radius of electrons or ions in a magnetic field.

LH: Lower Hybrid – RF waves of frequency intermediate between the electron and ion gyrofrequencies and used for heating and current drive.

L-Mode: Low (confinement) Mode – the baseline for confinement in magnetic fusion devices, dominated by strong turbulent transport.

Lorentz force. The basic electromagnetic force on a particle, $q(\mathbf{E} + \mathbf{v} \times \mathbf{B})$, where q and $\mathbf{v}$ are the particle charge and velocity, and $\mathbf{E}$ and $\mathbf{B}$ are the electric and magnetic fields.

Lundquist number: The ratio of the resistive time scale to the Alfvén time scale. Plays a role for MHD waves analogous to the Reynolds number in fluid dynamics.

Magnetic flux: The integral over a surface area of the magnetic field perpendicular to the surface. Examples are the poloidal flux Ψ and the toroidal flux Φ.

Magnetic surface: A two-dimensional closed toroidal surface on which magnetic field lines lie. In an axisymmetric device, such as a tokamak, the equilibrium consists of nested magnetic surfaces.

Magnetosonic wave: Plasma sound wave, in which the magnetic field is compressed along with the gas.

Major radius: Distance from the center line of the torus to the magnetic axis or center of the plasma cross section.

Maxwellian: Velocity distribution characteristic of local thermal equilibrium: $f \sim \exp(-mv^2/2T)$ where T is the local temperature. (See Equilibrium).

Mercier criterion: A mathematical criterion describing the local stability of MHD modes.

MHD: MagnetoHydroDynamics – A fluid description of magnetized plasmas.

Minor radius: distance from the center of the plasma cross section to the plasma edge.

Mode: 1) a plasma wave that has, at least approximately, constant wavenumber and frequency, or can otherwise be characterized as a normal mode of some wave equation. 2) a mode of operation of a device, such as L-mode or H-mode.

Mode conversion: The process by which waves of one type are converted to another, occurring in a region where the frequency and wavelength of both waves match.

NBI: Neutral Beam Injection – A common method of plasma heating which employs intense beams of neutral atoms.

NCSX: National Compact Stellarator eXperiment; a low aspect-ratio stellarator now in the design phase.

Neoclassical: The theory of collisional transport in toroidal geometry.

NTM: Neoclassical Tearing Modes – Resistive MHD modes driven unstable by bootstrap currents.

Ohmic: Characteristic of dissipation due to plasma resistivity; Ohmic heating, in which a current is driven inductively in the plasma.

Pedestal: Region of steep gradients at plasma edge

PIC: Particle In Cell – a numerical method for solving kinetic equations by following representative particles moving in self-consistent fields. (Cell refers to the grid used in the computation.)

Poloidal: The short way around in a torus–denoted by θ (see Fig. III.1 of the main report).

Potato: The shape of a particle orbit (potato orbit) close to the magnetic axis – a distorted *banana*, which see.

QPS: Quasi-Poloidal Stellarator: a low aspect-ratio stellarator now in the design phase.

Quasi-linear: In plasma theory, an expansion in powers of wave amplitude in which the fast motion is treated linearly and zero-frequency quadratic terms determine the slow background evolution.

Resonance. 1) A singularity in the electromagnetic plasma wave equations arising from the variation of physical parameters in space. 2) A singularity in the particle distribution function arising from the matching of wave and particle velocities in particular regions of phase space, a wave-particle resonance.

Reynolds Stress: The advection or transport of momentum by turbulence.

RF: Radio Frequency – refers to methods of heating, flow or current drive depending on the interaction of plasmas with externally launched waves.

$\rho*$: Pronounced "rho-star", the ratio of the ion gyroradius to the minor radius. Usually $\rho*$ is much smaller than unity and is used as an expansion parameter. See Gyrokinetic.

RWM: Resistive wall mode – an MHD mode driven unstable by finite wall resistivity (as opposed to zero resistivity of the ideal MHD theory).

Sawtooth: Sawtooth oscillation, a minor disruption in which core field lines reconnect flattening the temperature profile, followed by a slow reheat of the central plasma. This occurs cyclically.

Self Organized Criticality – (SOC). A characteristic of turbulent or chaotic systems where large scale or long lasting structures arise spontaneously and are responsible for significant amounts of transport.

Separatrix: The boundary surface between regions of a plasma on closed and open field lines. Since there is no confinement in the direction parallel to magnetic field lines, the plasma on open field lines is very cold. (Typically less than a million degrees K)

SOC: See Self Organized Criticality

SOL: Scrape-Off Layer – region of plasma existing on open field lines outside separatrix.

Stellarator: Unlike most magnetic confinement devices, stellarators employ helical magnetic fields formed by external coils and is intrinsically three dimensional.

Tearing modes: Resistive MHD modes driven by plasma currents.

TFTR: Tokamak Fusion Test Reactor – a large (former) magnetic confinement experiment.

Tokamak: A magnetic confinement device relying on a large toroidal current carried by the plasma itself. (The acronym is Russian standing for TOroidalnaya KAmera Magitnaya Katushka or toroidal chamber and magnetic coil.)

Toroidal: 1) Characteristic of a machine that is shaped like a torus or doughnut. 2)The long way around in a torus-denoted by ϕ (See Fig. III.1 of the main report).

Transport Barrier: Region of strongly reduced transport - typically a bifurcated state. (See ITB, H-mode, and Pedestal).

Transport flux: The total energy or particle flow per unit area passing through a surface element.

Unfavorable curvature: see Curvature

Vlasov equation: The collisionless version of the Boltzmann equation.

Wave-particle resonance. See Resonance 2).

WDM – Whole Device Modeling

Whistler wave: A high frequency plasma wave with a nonlinear dependence of frequency and wave number. Occurring naturally in the ionosphere, whistler waves were first detected by radio operators during World War I.

X-Point: A point in the plasma cross section where the poloidal magnetic field is zero; occurs at one or more points along the *separatrix*

Zonal flows: Flows which are constant on flux surfaces but with strong radial variation. These flows are self-generated by plasma turbulence for which it is also an important non-linear saturation mechanism.