


Center for Simulation of Wave Interactions with Magnetohydrodynamics


Abstract

The proposed work will build upon the successes of the DOE SciDAC programs by taking several of the most advanced fusion computer codes, combining these to provide a unique tool in the worldwide fusion program, and implementing them on the Leadership-Class and other computing facilities. This will be done in concert with the computer science and math communities in such a way that the software framework we form will lay a foundation for even more powerful and comprehensive simulations in the future. Intense radio-frequency waves provide an essential control technique for fusion grade tokamak plasmas. The frequency and phasing of these waves can be adjusted using precision control so as to keep the plasma in a quiescent state, significantly improving its quality as a confinement device and greatly reducing the likelihood of a catastrophic disruption. The work described here is aimed at providing an essential design tool for such a control system. The framework will allow for the treatment of two regimes of potential plasma global instabilities: fast phenomena and slow phenomena. In the first, the wave



Simulation of an internal reconnection event using the M3D code.

Source: W. Park et al., *Nuclear Fusion* 43, 483 (2003)



Simulation of conversion of fast waves to both ion Bernstein waves and slow ion cyclotron waves using the AORSA2D code.

Source: E. F. Jaeger et al., *Phys. Rev. Lett.* 90, 195001 (2003)

fields act only indirectly on the instabilities by modifying the background equilibrium, while in the second they act directly by competing with the instability drive terms. We will work with the community to develop standards and documented APIs for fusion application modules, which will allow the incorporation of many of today's fusion application codes into this framework.

Center for Simulation of Wave-Plasma Interactions

SC Program Announcement: LAB 05-11, Scientific Discovery through Advanced Computing: Fusion Simulation Prototype Centers

Laboratory: Oak Ridge National Laboratory

Principal Investigator: D. B. Batchelor

Title: Senior Research Staff Member

Telephone: (865) 574-1288

Fax: (865) 576-7926

e-mail: batchelordb@ornl.gov

Mailing address: P.O. Box 2008
Oak Ridge, TN 37831-6169

Laboratory Official: Stanley L. Milora

Title: Director, Fusion Energy Division

Telephone: (865) 574-0988

Fax: (865) 576-7926

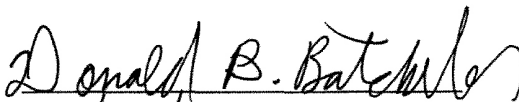
e-mail: milorasl@ornl.gov

Requested Funding:	Year 1	Year 2	Year 3	Year 4	Year 5
	2 million	2 million	2 million	2 million	2 million

Use of human subjects in proposed projects: No

Use of vertebrate animals in proposed project: No

PI Signature:



Date:

3/23/05

Official Signature:



Date:

3/23/05

List of Co-Principal Investigators from Other Institutions

Institution: Princeton Plasma Physics Laboratory
Principal Investigator: S. Jardin
Telephone: (609) 243-2635
e-mail: sjardin@pppl.gov
Mailing address: P. O. Box 451
Princeton, NJ 08543

Institution: University of Indiana
Principal Investigator: R. Bramley
Telephone: (812) 855-7790
e-mail: bramley@cs.indiana.edu
Mailing address: 107 S. Indiana Avenue
Bloomington, IN 47405-7000

Institution: Columbia University
Principal Investigator: D. Keys
Telephone: (212) 854-1120
e-mail: kd2112@columbia.edu
Mailing address: Plasma Physics Laboratory
202 S.W. Mudd Bldg, MC 4701
New York, NY 10027

Institution: Massachusetts Institute of Technology
Principal Investigator: J. Ramos
Telephone: (617) 253-2051
e-mail: ramos@psfc.mit.edu
Mailing address: RM NW16-238
77 Massachusetts Avenue
Cambridge, MA 02139

Institution: General Atomics
Principal Investigator: D. Schissel
Telephone: (858) 455-3387
e-mail: schissel@fusion.gat.com
Mailing address: 3550 General Atomics Court
San Diego, CA 92121-1194

Institution: Science Applications International Corporation
Principal Investigator: D. Schnack
Telephone: (858) 826-6021
e-mail: schnackd@saic.com
Mailing address: 10260 Campus Point Drive
San Diego, CA 92121

List of Co-Principal Investigators from Other Institutions (cont.)

Institution: CompX
Principal Investigator: R. W. Harvey
Telephone: (858) 509-2131
e-mail: bobh@compxco.com
Mailing address: P. O. Box 2672
Del Mar, CA 92014-5672

Institution: New York University
Principal Investigator: H. Strauss
Telephone: (212) 998-3268
e-mail: strauss@CIMS.nyu.edu
Mailing address: 251 Mercer Street
New York, NY 10012

Institution: Lawrence Berkeley National Laboratory
Principal Investigator: P. Colella
Telephone: (510) 642-2652
e-mail: pcolella@lbl.gov
Mailing address: Department of Mechanical Engineering
6119 Etcheverry Hall
University of California
Berkeley, CA 94720-1740

Institution: University of Wisconsin
Principal Investigator: C. Hegna
Telephone: (608) 263-0810
e-mail: hegna@cae.wisc.edu
Mailing address: ERB 517
1500 Engineering Drive
Madison, WI 53706-1687

2. Executive summary

Comprehensive plasma simulations are essential to the development of fusion energy. Not only does simulation serve to advance science by allowing us to evaluate and test basic theory through comparison with experiments, it also has a direct economic benefit by maximizing the productivity of experimental facilities and by supporting design decisions for new facilities. For a device like ITER, such decisions can have multi-billion dollar consequences. The need for a comprehensive simulation capability for fusion devices has been emphasized by the fusion research community and by several program evaluation and planning activities [1, 2]. As stated in the Fusion Simulation Project (FSP) Steering Committee report: “The ultimate goal of the FSP is to simulate the behavior of toroidal magnetic fusion devices on all important time and space scales, and to account for the interactions of all relevant processes. It will serve as an integrated ‘standard model’ for the community, permitting the exploration of new multi-scale physics and mathematics. The driving focus of the project is to contribute to design decisions, experimental planning and performance optimization of the International Thermonuclear Experimental Reactor (ITER)”.

One of the most fundamental observations in fusion devices is that performance is often limited by the appearance of unstable plasma motions that can degrade plasma containment or even terminate the plasma discharge, with the potential for damage to the containment device. These instabilities can be modeled using extended magnetohydrodynamics (MHD). A second observation is that high-power radio frequency (RF) electromagnetic waves can be used to influence plasma stability – sometimes producing instability, and sometimes reducing or eliminating instability. A capability to understand and predict the effects of RF waves on plasma stability would be of significant scientific and economic benefit.

We propose to create a new and unique capability for modeling the interactions between high-power electromagnetic waves and MHD stability, by building an innovative end-to-end ultrascale computational system that integrates emerging mathematical methods, software component and data management systems, and several of the world's most accurate and capable fusion energy codes. This will be carried out by a cross-disciplinary, multi-institutional center for Simulation of Wave Interactions with Magnetohydrodynamics (SWIM). The broad scientific objectives of the center are to:

- **Improve our understanding of interactions that both RF wave and particle sources have on extended-MHD phenomena, and to substantially improve our capability for predicting and optimizing the performance of burning plasmas.**

and to:

- **Develop an integrated computational system for treating multi-physics phenomena with the required flexibility and extensibility to serve as a prototype for the FSP, address mathematics issues related to the multi-scale, coupled physics of RF waves and extended MHD, and to optimize the integrated system on high performance computers.**

It is useful to distinguish two regimes of phenomena addressed by extended MHD models in the presence of wave and particle sources:

1. **Fast phenomena** in which the unstable plasma motions are so rapid that energy deposition, currents, flows, and energetic particle populations driven by RF waves or other sources do not directly influence the rapid motion. Instead the RF acts indirectly over a longer timescale to drive the slow evolution of the plasma equilibrium and profiles, effectively setting the initial conditions for the fast MHD event. An example of such a phenomenon is the crash phase of the sawtooth oscillation, described more fully in Section 3. After the rapid event is over, the plasma is left in a new, relaxed state that may be subject to further slow evolution driven by RF and other sources.
2. **Slow phenomena** such that energy deposition, currents, or flows driven by RF waves or other sources do operate on timescales that directly influence the dynamics of the unstable motions. The perturbations of the plasma state by the instability in turn influence the energy, currents or flows driven by the RF. An example of such a phenomenon is the neoclassical tearing mode (NTM, also described more fully in section 3), which is an important consideration for the performance of burning plasma such as that of ITER.

The first regime above is the traditional one in which RF sources can be thought of as affecting the global plasma pressure and current profile and creating an energetic particle population. Models have focused primarily on axisymmetric systems with closed, nested flux surfaces. The second regime addresses some very important plasma instabilities having localized spatial regions near rational flux surfaces within which RF, the kinetic evolution of the velocity distribution, and extended MHD phenomena are all tightly coupled. Furthermore, the assumption of closed, nested flux surfaces breaks down through the appearance of magnetic islands. Coupling the RF-driven evolution of the plasma distribution function into the dynamical equations of the extended MHD model can be considered as an extension of the general concept of fluid closure. Such coupling will require new physics analysis and will likely require new mathematical and algorithmic techniques to solve the resulting equations stably and efficiently. This is just the kind of coupled, multi-scale, non-linear scientific issues that Focused Integration Initiatives were conceived as addressing.

The SWIM project will consist of three elements.

1. **Development of a computational platform** that will allow efficient coupling of the full range of required fusion codes or modules, that is flexible enough to allow exploration of various physics models and solution algorithms, that permits convenient user access and access to experimental data, and that is robust to evolving physics, code development, and developments in computer hardware.
2. **A physics campaign addressing long timescale discharge evolution in the presence of sporadic fast MHD events.** This will involve developing within the platform an integrated Evolving Equilibrium Model (EEM) that is interfaced to the 3D non-linear extended MHD codes. The EEM will contain modules to calculate wave propagation and absorption in all relevant frequency regimes, modules to calculate the modification of the plasma velocity distribution from sources (RF, neutral injection and α particles), modules to calculate profile and magnetic evolution (assuming closed flux surfaces), linear MHD stability models and reduced models of non-linear MHD events. The primary physics focus will be to allow the development of optimized burning plasma

scenarios and to improve the understanding of how RF can be employed to achieve long-time MHD stable discharges and control sawtooth events.

- 3. A physics campaign for modeling the direct interaction of RF and extended MHD for slowly growing modes.** This will require development of new approaches to closure for the fluid equations and the interfacing of RF modules directly with the extended MHD codes and with code modules that implement the fluid closures. The RF field solvers and Fokker Planck modules for evolving the distribution function will be generalized to non-axisymmetric magnetic fields containing magnetic islands as generated by the extended MHD codes. The primary physics focus will be to improve the understanding of how RF can be employed to control neoclassical tearing modes.

Appropriate modules will be optimized for the fastest supercomputers available, including those at the National Leadership Computational Facility (NLCF) and National Energy Research Supercomputer Center (NERSC). To carry out coupled calculations of this nature in parameter regimes relevant to burning plasma will require substantial increases in computer resources. Straightforward extrapolation of present calculations put the computational requirements in the petaflop range, thus necessitating a significant use of NLCF level of computation as well as a staged increase in the complexity of models and a continuing program of algorithm improvement.

Our approach will be to build on previous and ongoing investments in SciDAC, the base programs in Office of Fusion Energy Science and Office of Advanced Scientific Computing Research, and laboratory internal R&D projects. The proposed project will unite experts in the magnetic fusion and the computer science and enabling technologies communities and will be a five-year effort costing \$2M per year.

3. Background

In this section we present a brief discussion on the plasma physics and computer science issues that are the main focus of this proposal. This presentation will emphasize those projects which will constitute the core of the foundations for the present proposal. It would not be possible to make effective progress on a set of objectives as ambitious as those as set out in this proposal without building on expertise and a base of codes and theory from previous and ongoing DOE projects. Investigators on this proposal include leaders and participants in the critical ongoing SciDAC projects—the Center for Simulating Wave-Plasma Interaction (CSWPI) and the Center for Extended MHD Modeling (CEMM) funded by the Office of Fusion Energy Science (OFES), and the Fusion Collaboratory, and several of the ISICs funded by the Office of Advanced Scientific Computing Research (OASCR). We will leverage these and a significant new investment by OFES and OASCR to develop a new predictive transport capability based on reduced dimensionality models. We will also draw on work supported by ORNL Laboratory Directed R&D (LDRD) funding toward integrated fusion modeling and beginning the process of porting and optimizing needed codes to the National Leadership Computing Facility (NLCF).

3.1 Concepts of plasma confinement in a tokamak

In a tokamak (Figure 1) high temperature plasma is created from isotopes of hydrogen, confined in a toroidal “magnetic bottle,” and heated by an electrical current flowing in the plasma and by the injection of high-energy neutral particle beams (NBI) and/or RF waves.

We refer to the *toroidal* (the long way around the torus) and *poloidal* (the short way around the torus) directions. A quiescent discharge is completely symmetrical in the toroidal direction (axisymmetric).

The magnetic fields in a tokamak are primarily in the toroidal direction (into the paper). If we follow a given magnetic field line, it will form, closed nested magnetic surfaces. These flux surfaces form the basis for magnetic confinement.

On certain *rational surfaces*, the field lines will close on themselves after just a few transits the long way and the short way around the torus. If the lines close after just one transit around in each direction, we call it the $q = 1$ surface. The *sawtooth oscillation* is an instability that is confined to the interior of this surface.

The *neoclassical tearing mode* (NTM) is associated with magnetic islands at one of the outer rational surfaces and is described in the following sections. Other classes of modes that we will consider are *Edge Localized Modes* (ELMs) that are caused by the large pressure and current gradients near the plasma edge, or separatrix, and *Energetic Particle Modes* (EPM) that can occur when the plasma has an energetic particle component which resonates with a normal mode of the plasma torus.

3.2 Stability of tokamaks

If the pressure gradient or the local electric current become too large, the normally quiescent tokamak plasma can become unstable, resulting in growing, sometimes violent perturbations.

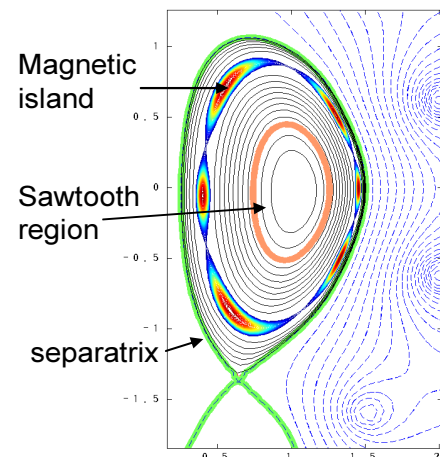


Figure 1: Poloidal cross section of a tokamak

An important and often observed example of such an instability is the relaxation phenomenon called the “sawtooth oscillation.”

The Sawtooth Oscillation.

The sawtooth cycle consists of two distinct phases: (1) The slow sawtooth ramp phase; a quiescent period during which the plasma density and temperature increase approximately linearly with time in the plasma core due to plasma heating and the plasma remains largely axisymmetric and, (2) a rapid crash phase in which a 3D disturbance develops, causing the density and temperature within the $q = 1$ surface to flatten, resulting in a return of the plasma to an axisymmetric state. This sequence is usually repetitive with the ramp phase typically being much longer than the crash time. The crash results in an outward transport of energy and energetic particles that will affect the fusion burn in an ignited device such as ITER. The sawtooth crash is sometimes associated with initiating other types of instability, for example NTMs and ELMs.




Figure 2: Experimentally, the sawtooth period and amplitude have been extended by the application of RF power.[3].

Experimentally, the sawtooth period and amplitude have been extended (and reduced) by the application of RF power. [3] (see Figure 2) This interaction between RF and the instability involved predominantly the generation of an energetic ion component that provides a stabilizing effect. These experiments have more recently been extended [4] to show that RF induced currents driven near the $q = 1$ surface can be used to shorten the sawtooth periods and reduce their amplitude. A predictive model for this would be of great use since these smaller sawtooth oscillations are much less likely to initiate the NTM described in the next section. Tools for simulating each of the components of this interaction have been and are continuing to be developed by the CSWPI and CEMM SciDAC centers. The combining of these tools is one of the central foci of this proposal.

The Neoclassical Tearing Mode.

The NTM involves one or more magnetic islands developing at an outer rational surface. Experimentally, the NTM is thought to limit the obtainable pressure in many long pulse tokamak discharges. [5–8]. The SciDAC CEMM extended-MHD codes are capable of modeling the essential features of these modes in the absence of RF. This is a high-priority area of active research within the CEMM project and we expect the realism of these models to improve in the near future [9,10].




Figure 3: Complete suppression of an $m/n=3/2$ NTM by ECCD in the presence of continued uncoupled sawteeth: DIII-D, [11]

Where NTMs cannot be avoided, active stabilization is required. Electron-cyclotron current drive (ECCD) has been experimentally demonstrated to stabilize tearing modes (Figure 3)[11]. The present plan is for ITER to have a ECCD capability for this purpose, but there is considerable uncertainty about the needed power level, control requirements, and deposition geometry.

The SciDAC CEMM and CSWPI centers have developed, and are continuing to improve, the components needed for an integrated simulation of RF stabilization of the NTM. The combining of these components into a usable, validated, simulation and design tool is a second major focus of this proposal.

3.3 The elements of fusion simulation

A comprehensive fusion simulation must cover a vast range of timescales. The particular timescales associated with different plasma phenomena determine the appropriate plasma models to use and have led to development of various plasma simulation sub-disciplines—RF wave simulation, micro-turbulence simulation, extended MHD simulation, and transport simulation. Figure 4 summarizes these relationships for plasma parameters characteristic of a burning plasma

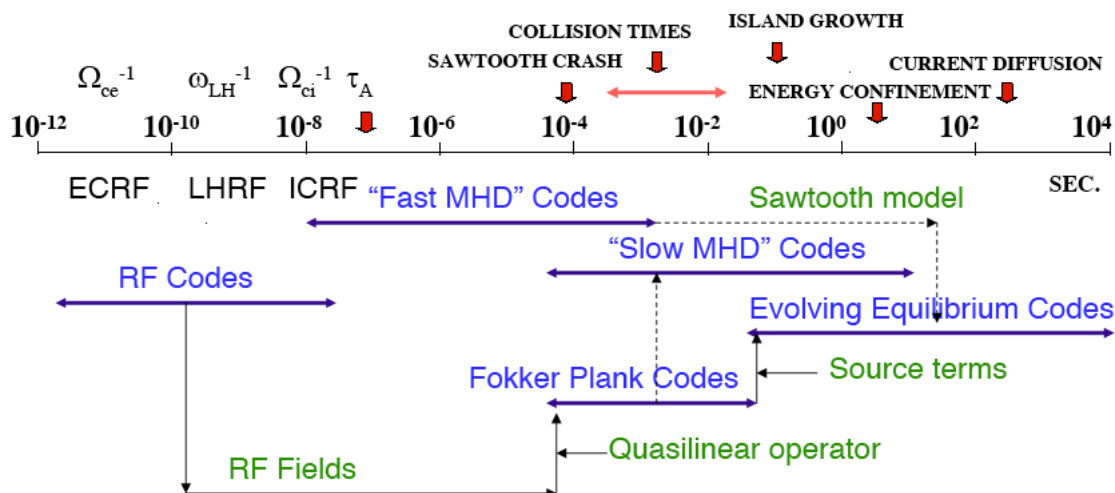


Figure 4. Typical timescales for a burning plasma experiment with $B = 5 \text{ T}$, $R=6 \text{ m}$, $n_e = 10^{20} \text{ m}^{-3}$, $T = 10 \text{ keV}$. Different plasma models are applied to phenomena in different timescale regimes

experiment. RF waves on the fastest timescales, the electron or ion gyro-period (10^{-11} s and 10^{-8} s respectively), are exploited for plasma heating and current drive. The slowest timescale dealt with in this proposal is that required for inductively driven current to diffuse throughout the plasma, several minutes in an ITER scale device. An important intermediate timescale is the time for global transport of energy, or confinement time ($\sim 1 \text{ s}$). Although the RF wave period is the fastest timescale, the wave amplitudes are small enough that global modifications they produce in the plasma distribution functions usually develop on the much slower, transport timescale, although spatially localized modifications due to cyclotron-resonant interactions can develop much more rapidly. Extended MHD models deal with phenomena that are slow compared to the ion gyro-period but range in timescale from the crash of a sawtooth oscillation ($\sim 10^{-4} \text{ s}$) to the slow growth of magnetic islands associated with instabilities like the neoclassical tearing mode ($\sim 0.1 \text{ s}$).

The starting point for the theory of hot magnetized plasmas is the kinetic equation, a form of the Boltzmann equation that describes the evolution in time of the distribution of plasma particles of species j in a six-dimensional phase space of position and velocity, $f_j(\mathbf{x}, \mathbf{v}, t)$.

$$\frac{\partial f_j}{\partial t} + \mathbf{v} \cdot \nabla f_j + \frac{q_j}{m_j} [\mathbf{E} + \mathbf{v} \times \mathbf{B}] \cdot \nabla_{\mathbf{v}} f_j = C(f_j) + S_j(\mathbf{x}, \mathbf{v}, t), \quad (1)$$

where q_j and m_j are respectively the charge and mass of particle species j , \mathbf{E} is the electric field and \mathbf{B} is the magnetic field. The second term on the left describes convection in position space, the third term describes convection in velocity space due to electromagnetic forces on the charged plasma particles. The first term on the right describes redistribution of particles in velocity space due to collisions and the final term represents any particle sources. There is an equation of this sort for each plasma species including electrons and usually several different ion species. The different kinetic equations are nonlinearly coupled through the contribution of each species to the electromagnetic field and through collisions. The electric charge, $q(\mathbf{x}, t)$, and current, $\mathbf{J}(\mathbf{x}, t)$, that act as sources for the electromagnetic field in Maxwell's equations are obtained by taking velocity moments of $f_j(\mathbf{x}, \mathbf{v}, t)$.

$$\begin{aligned} \mathbf{J}(\mathbf{x}, t) &= \sum_j q_j n_j(\mathbf{x}, t) \mathbf{V}_j(\mathbf{x}, t) \equiv \sum_j q_j \int d^3 \mathbf{v} \mathbf{v} f_j(\mathbf{x}, \mathbf{v}, t), \\ q(\mathbf{x}, t) &= \sum_j q_j n_j(\mathbf{x}, t) \equiv \sum_j q_j \int d^3 \mathbf{v} f_j(\mathbf{x}, \mathbf{v}, t) \end{aligned} \quad (2)$$

$$\mathbf{J}_p^1(\mathbf{x}) = \sigma \circ \mathbf{E}_1 = \sum_j \int d\mathbf{x}' dt' \mathbf{K}(f_j^0, \mathbf{x}, t, \mathbf{x}', t') \circ \mathbf{E}_1(\mathbf{x}', t')$$

where \mathbf{V}_j is the mean velocity of species j and n_j is the particle density of species j .

A solution of the time-dependent kinetic equation is rarely feasible in the full 6D phase space. Instead one works with approximate forms that have been averaged over space or time, or one works with equations derived by taking velocity-space moments of Eq. (1). We will be concerned in this proposal with four plasma theory disciplines that derive from this system of equations: RF plasma wave theory, Fokker Planck theory that describes the time average evolution of $f_j(\mathbf{x}, \mathbf{v}, t)$, extended MHD theory that treats plasma motion from a fluid standpoint with extensions to include some kinetic effects, and transport theory that describes the slow evolution of global plasma pressure and magnetic field. The timescales over which these disciplines are applicable are related to the various physical timescales in Figure 4.

Simulation of RF waves and their plasma interactions

Since the RF wave period is by far the fastest timescale in the system, the fields and distribution function can be separated into a time-average, or equilibrium, part $(\mathbf{E}_0, \mathbf{B}_0, f_j^0)$ that is slowly varying, and a rapidly oscillating part, $(\mathbf{E}_1(\mathbf{x})e^{-i\omega t}, \mathbf{B}_1(\mathbf{x})e^{-i\omega t}, f_j^1(\mathbf{x}, \mathbf{v})e^{-i\omega t})$ where ω is the frequency of the RF power source. For our applications, the time-harmonic wave fields are small compared to the equilibrium fields and we may linearize Eq. (1) with respect to these amplitudes. Solving the linearized equation gives the rapidly varying part $f_j^1(\mathbf{x}, \mathbf{v})$ in terms of the equilibrium distribution, $f_j^0(\mathbf{x}, \mathbf{v})$, and the rapidly varying component of the electromagnetic field. This solution then allows us to relate the plasma current induced by the wave fields, \mathbf{J}_p^1 to the fields through a nonlocal, integral conductivity operator acting on the wave field.

$$\mathbf{J}_p^1(\mathbf{x}) = \sigma \circ \mathbf{E}_1 = \sum_j \int d\mathbf{x}' dt' \mathbf{K}(f_j^0, \mathbf{x}, t, \mathbf{x}', t') \circ \mathbf{E}_1(\mathbf{x}', t') \quad (3)$$

The Maxwell's wave equation that must be solved reduces to a generalization of the Helmholtz equation,

$$\nabla \times \nabla \times \mathbf{E} = \mathbf{J}_p + \mathbf{J}_{ant} \quad + \text{boundary conditions} \quad (4)$$

The source for the waves is an externally driven antenna current, \mathbf{J}_{ant} , localized near the plasma edge. The interaction takes place in a bounded domain on which are imposed appropriate boundary conditions determined by the shape of the fusion device.

Three wave frequency regimes are important for fusion applications [12,13]. The waves in the ion cyclotron range of frequencies (ICRF), (20 MHz to 140 MHz) resonate with the ions at the ion cyclotron frequency, ω_i , or its harmonics, and can interact with the motion of the electrons that is parallel to \mathbf{B} . Because of their long wavelength, an accurate calculation of ICRF waves requires solution of the full-wave equation, Eq. (3).

Lower hybrid waves (1 GHz to 4 GHz) can interact with ions and electrons but are used primarily to drive electron current. A geometrical optics approximation to the wave equation is usually adequate for calculation of lower hybrid wave propagation.

Waves in the electron cyclotron range of frequencies (ECRF), (120 GHz to 240 GHz) resonate at the electron cyclotron frequency, ω_e , or its harmonics. Propagation is well described by geometrical optics with ray tracing algorithms.

Solving the wave equation for magnetized plasma in 2D or 3D is a terascale computing problem. The CSWPI has advanced the development of two complementary wave solvers which we intend to apply in SWIM: the AORSA2D/AORSA3D codes [14-16] and the TORIC [17,18] code. The All-Orders Spectral Algorithm (AORSA) solvers are based on an integral equation formulation with no restriction on wavelength relative to the Larmor radius or on wave harmonic. The TORIC code is based on an expansion of the conductivity operator to second order in the smallness of the Larmor radius relative to the wavelength perpendicular to the magnetic field. This converts the integral conductivity operator to an approximate differential form. Conductivity modules needed by the wave solvers were extended to allow for arbitrary, non-Maxwellian distribution functions [19-21]. Effective use is made of the IBM SP computer at the National Energy Research Scientific Computing Center (NERSC) as well as the ORNL IBM SP, SGI Altix, and more recently Cray X1 computers. AORSA has been scaled successfully up to 2048 processors on the NERSC Seaborg computer. Computing speed efficiency for the matrix factorization scales linearly with the number of processors, and there is no indication of saturation caused by communication between processors. AORSA has been nominated to be included in the FY05 Energy Research Joule Report on Software Effectiveness.

Ray tracing, as required for ECRF or lower hybrid is considerably less computationally intensive since it is only necessary to solve a set of ordinary differential equations that describe the ray trajectory and power absorption along the ray. SWIM will employ the GENRAY [22] code for ray tracing applications. These codes (AORSA, TORIC, GENRAY) are the "RF codes" of Fig 4.

Fokker Planck theory

The response of the plasma distribution function, $f_0(\mathbf{x}, \mathbf{v}, t)$, on timescales slower than the RF period is obtained from a time-averaged form of Eq. (1), referred to as the quasilinear Fokker-Planck equation,

$$\frac{df_0}{dt} = Q(\mathbf{E}^1, f_0) + C(f_0) + S(\mathbf{x}, \mathbf{v}) \quad (5)$$

Here, the time derivative includes particle drift motion and, the RF quasilinear operator, $Q(\mathbf{E}^1, f_0)$ [23], is quadratic in the RF field $\mathbf{E}(\mathbf{x}, t)$ and describes wave-induced velocity-space diffusion of f_0 . Velocity moments of the quasilinear operator can also be calculated to give instantaneous macroscopic plasma responses such as local power deposition, $W_{rf}(\mathbf{x})$, momentum sources, $\mathbf{R}_{rf}(\mathbf{x})$, or ponderomotive force, $\mathbf{F}_{rf}(\mathbf{x})$. In the absence of sources the collision operator relaxes f_0 to a local Maxwellian distribution. However quasilinear diffusion, particle or energy sources, or gradients in macroscopic plasma quantities drive f_0 away from Maxwellian.

There are three approaches to solving Eq. (5) that employ different numerical techniques and different time or orbit averages to reduce the dimensionality – finite differences, expansion in orthogonal polynomials, and direct simulation using Monte Carlo particle methods. SWIM will include code representatives from each of the three approaches in order to exploit their relative advantages and to demonstrate flexibility of the SWIM architecture.

We will employ the CQL3D code [24] which solves a time dependent bounce averaged form of Eq. (5) by finite difference methods in a 2D velocity space with a 1D radial spatial coordinate. It interfaces with ray tracing and neutral beam deposition packages. CQL3D includes quasilinear diffusion and collisions and has the capability to include radial derivatives of the bounce averaged distribution functions. The effect of radial drift motion averages out of the lowest-order bounce-averaged equation solved by the CQL3D FP code. The existing ad-hoc models for the radial diffusive motion which exists after bounce-averaging need to be improved to give a physics-based model for radial diffusive transport and RF modifications of the bootstrap current in CQL3D.

The DKES (Drift Kinetic Equation Solver) code [25] solves a time independent form of Eq. (5) by velocity space expansion in orthogonal polynomials and without bounce averaging. DKES does not presently include quasilinear RF diffusion or energy scattering. To treat neoclassical effects where deviation from a Maxwellian is modest only low-order polynomials in energy are needed in velocity space. It may be necessary to increase this number if quasilinear distortion of the distribution function is significant. DKES is widely used to calculate radially-local neoclassical transport properties of stellarators and can therefore analyze nonlocal (on a magnetic surface), three-dimensional transport effects in the presence of coherent islands. These features complement the capabilities of CQL3D.

Monte Carlo methods such as DELTA5D [26] allow solutions in high dimensional phase space and correctly model the drift of particle orbits across flux surfaces. The method can straightforwardly include both Coulomb pitch angle and energy scattering as well as an RF quasilinear diffusion. In addition, RF heating effects on the distribution, non-local closures, radial transport effects and unclosed flux surfaces can be taken into account. The Monte Carlo approach is scalable to massively parallel systems, but there are challenges related to the

computational performance. First, there is the need to run with a sufficiently large number of particles to resolve narrow boundary layers that may be present in the distribution function, especially for the collisionalities of ITER. Next, a number of self-consistency relations (e.g., momentum conservation, equivalency between test particle and background plasmas, ambipolar transport, etc.) need to be prescribed. These may require higher-level iterations outside the main particle stepping loop.

To obtain self-consistency of the wave fields with the distribution, and to follow the time evolution of the plasma, a procedure of iteration of Eqs. (3) through (5) is necessary. This iteration greatly compounds the computational demands of the RF calculation. This loop has now been closed using full-wave fields from the AORSA code to generate the quasilinear operator and the CQL3D Fokker Planck code to calculate the distribution function. More detailed descriptions of the RF wave solvers and Fokker Planck codes to be employed is contained in Appendix A.

Equations of MHD and extensions

Fluid descriptions of the plasma are obtained by taking velocity moments of the kinetic equations, Eq. (1). The *first two* even moments yield the conservation equations for mass and energy. We take the sum of the first odd moments of the electron and ion equations, and the first electron odd moment, neglecting the (small) electron inertial terms, to obtain the fluid momentum equation and the generalized Ohm's law that determines the electric field, thus:

$$mn \frac{d\mathbf{V}}{dt} = -\nabla(p_e + p_i) + \mathbf{J} \times \mathbf{B} - \nabla \cdot (\Pi_e + \Pi_i) + \mathbf{F}_{rf} \quad (6)$$

$$\mathbf{E} = -\nabla \times \mathbf{B} + \frac{1}{ne} (\mathbf{J} \times \mathbf{B} - \nabla p_e - \nabla \cdot \Pi_e + \mathbf{R} + \mathbf{R}_{rf}) \quad (7)$$

Here, \mathbf{R} is the electron-ion Coulomb friction term, and Π_e, Π_i are the stress tensors for the electrons and ions. The \mathbf{F}_{rf} and \mathbf{R}_{rf} are new terms that arise from the moments of the quasilinear operator $Q(f)$ and account for the direct RF/MHD interaction. The density, temperature and flow velocity are determined dynamically by the macroscopic evolution equations. However, there are more unknowns than equations, and *closure relations* must be developed that describe the stress and heat flux tensors in terms of the independent MHD variables, n, V, p, \mathbf{J} , and \mathbf{B} and their derivatives. Specifically the traceless part of the stress tensor and the heat flux vector, along with the velocity moments of the Coulomb collision, $C(f)$, and RF quasilinear scattering, $Q(f)$, operators are the closure variables that must be determined from kinetic theory. Closure relations can be thought of as finding solutions for the distribution function in terms of the MHD variables from a kinetic equation that contains the needed physics.

The closure problem is a challenging issue of its own right in conventional extended-MHD theory without RF interactions. The determination of stress tensor (or viscosity tensor) and the heat fluxes constitute the classic MHD closure problem. There are a number of different approaches to closure currently being developed within the MHD community and the SciDAC CEMM center is evaluating the relative merits of these competing approaches. The inclusion of RF effects in MHD through the quasilinear operator represents new ground in the theory of closures in plasma physics.

CEMM is presently focusing on understanding the physics of sawteeth, tearing modes, and several other plasma instabilities. They have two widely applied Extended MHD codes, NIMROD [27] and M3D[28], and have developed an exploratory adaptive mesh refinement code AMRMHD[29]. These are the codes labeled "Slow MHD codes" and "Fast MHD codes" in

Figure 4. The same codes are used for both “Slow” and “Fast”, but with a much larger implicit timestep for the Slow MHD that does not resolve the fast timescales, but enables long-time integrations. Both NIMROD and M3D scale on existing computers to several hundreds of processors, and work is underway within CEMM to improve the scaling so that the use of several thousand processors can become routine.

It is a stated mission of CEMM to “begin some integration activities that will prototype those required for the Fusion Simulation Project”. For these reasons, it is natural to build upon the successes of CEMM to form the basis for the Extended MHD part of the SWIM project.

Transport theory and 1.5-D transport modeling

To treat the slow plasma evolution due to transport processes the inertial terms from the momentum equation can be neglected [30, 31]. In addition, the effects of the off-diagonal terms in the stress tensors are negligible over long, transport timescales allowing us to replace the momentum equation with the MHD equilibrium equation, $\nabla p = \mathbf{J} \times \mathbf{B}$, where $p = p_e + p_i$ is the total fluid pressure. This equilibrium state is uniquely determined by the profiles of the plasma pressure, the safety factor $q(\psi)$, and the boundary conditions.

Furthermore, the plasma densities and temperatures can be averaged over the magnetic surfaces to eliminate the timescales associated with the fast parallel transport processes within these surfaces allowing one to obtain a self-contained system of 1-D equations for the perpendicular transport of the surface averaged quantities. The details of the flux surface geometry are captured in a set of surface averaged metric quantities describing the enclosed volume, surface area, enclosed current, etc. The resulting system, which consists of an evolving 2D MHD equilibrium together with a system of 1-D transport equations is sometimes called a 1.5-D description. This reduction allows us to integrate the remaining MHD evolution equations for the magnetic fields and thermodynamic quantities for long times (compared with the Alfvén wave transit time) without being hampered by a time-step restriction determined by those waves, a tremendous gain in computational efficiency. This forms the basis for the “Evolving Equilibrium Code Suite” described in Section 4.1.2.1.

The transport equations in the 1.5-D codes relate the time rates of change of density, momentum, and energy, to the balance of sources (i.e. due to neutral beam injection, RF heating, or alpha heating), losses (radiation, charge exchange), and cross-field transport. In “analysis” applications, the rates of change are taken from experimental measurements, the sources and losses are modeled, and the transport is inferred. In “predictive” applications, a transport model, based typically on neoclassical and/or turbulent transport theory, is added to the source/loss models, and the time rates of change of the plasma fluid moments are predicted.

PTRANSP is a two-year multi-institutional project begun in March 2005 to perform a significant upgrade of the evolving equilibrium, 1.5D, TRANSP code by implementing several new predictive capabilities. TRANSP is perhaps the world’s most widely used tokamak analysis code, but is normally used in an analysis mode rather than a predictive mode. It is being upgraded by adding new free-boundary equilibrium and non-linear equation solver options. Some of these options will be obtained from the Tokamak Simulation Code (TSC) by breaking it into modules. The SWIM project plans to build on this national activity to provide the Evolving Equilibrium Code suite that is described in Section 4.

3.4 Challenge of coupling multiscale/multi-physics, mathematics issues

A numerical approximation to a set of time-dependent partial differential equations must be numerically stable. Established techniques exist for determining the numerical stability of relatively simple equations (such as the von-Neumann stability analysis, total variation

diminishing, or the Kreiss matrix stability criterion), but there is no general, practical technique for analyzing the numerical stability of a complex system of nonlinear coupled equations such as we are dealing with here. Stability of each individual component is no guarantee that the coupled system is stable.

An example of this related to Ohm's Law and determination of the electric field appeared in the initial attempts to couple the evolving equilibrium Tokamak Simulation Code (TSC) [32] with the Lower-hybrid Simulation Code (LSC) [33]. Given an electric field, LSC computes the RF induced current, $J_{RF}^n(E^n)$. TSC would then compute the parallel electric field at timestep n from the Ohm's Law equation:

$$E^n = \eta \left[J - J_{RF}^n(E^n) \right] \quad (8)$$

However, since the RF induced current, $J_{RF}^n(E^n)$, itself depends on the electric field that it is being used to compute, it was found that straightforward iteration of Eq. (8) led to an instability that could only be made stable by having the LSC code also return the Jacobian derivative, $\partial J_{RF} / \partial E$, so that Eq. (8) was replaced by the semi-implicit system:

$$E^n = \eta \left[J - J_{RF}^n(E^{n-1}) - \frac{\partial J_{RF}}{\partial E} (E^n - E^{n-1}) \right] \quad (9)$$

The formulation of Eq. (9), provided stable and convergent answers, whereas an iteration based on Eq. (8) did not. We anticipate that this and other similar techniques will be necessary to stabilize the coupled calculations described in the Section 4.

3.5 Computational science background

SWIM will leverage the products of the SciDAC program's ISICs for mathematics and computer science. Even so, an FSP introduces unique challenges that require fusion modeling-specific research. Mesh and geometry interfaces from TSTT [33], solver libraries for discretizations of PDEs from TOPS [34], and adaptive mesh refinement from APDEC [35] will all be leveraged. However, like any complex application, SWIM will require custom preconditioned iterative solvers (assembled from TOPS components), conservative interpolation methods for accurate and effective exchange of data between different discretizations, and significant analysis and testing before multilevel methods like AMR or multigrid can be turned over for production use. Important initial work has been performed, such as the AMRMHD code [29], the experimental coupling of RF and transport codes in ORNL's LDRD project, and formal definition of some code interfaces. SWIM will need to implement a fuller integration of physics to include transport and modular closure relations.

The current state of software systems and High Performance Computing (HPC) methodologies in fusion energy simulations ranges from community efforts like the National Transport Code Collaboration (NTCC) with coding, portability, and documentation standards, to single site systems with rudimentary documentation and legacy approaches to management. SWIM project proposes to leverage the existing systems as much as possible while updating them and mapping the ones needed for an integrated MHD simulation to next-generation high performance platforms.

The existing codes for RF, MHD, and transport described earlier are actually mostly code systems, using conditional compilation and complex run-time scripting to target different problems and physics within each domain. Partly this has been driven by an attempt to add some cross-area physics, but typically as add-ons within the code. Many of the existing codes are multilingual and use Fortran77, Fortran90, C, C++, and a variety of scripting systems: Perl, Python, IDL, and shell scripts, which any framework will need to support initially. Each code

system is in near-daily use and undergoing constant modification, enhancements, and bug fixes. Multiple researchers at distributed remote sites are performing modifications simultaneously and the typical code coherence mechanism is to have one code manager vetting and approving changes. Transitioning to a comprehensive software harness must support this evolution even as it develops an integrated framework. This will enable plasma research productivity while creating new capabilities, and avoid integrating a code system that is frozen at a past version.

Extant software management practices in fusion simulations range from a disciplined use of CVS (Concurrent Versions System) [36] tagging and branching, to ad hoc informal methods including simply sending copies to remote sites. The ad hoc route requires significant human effort to assure that copies do not uncontrollably diverge with time. Code configuration and launch mechanism across the range of contemporary codes is idiosyncratic and uses large numbers of scripts and much file prestaging by hand. One transport code alone has over 50 scripts and 43 makefiles for this purpose. Almost none use autoconfiguration tools. Existing codes in RF, MHD, and transport also vary widely in their use of HPC libraries and contemporary systems like PetSC, Trilinos, or even the Basic Linear Algebra Subroutines. Significant performance gains are likely simply from a more systematic use of DOE-developed and maintained computational toolsets.

Data management and formal interface definitions are important in large-scale integrated simulations, but have not played a significant role in the development of most existing software in the fusion community. Current practice in fusion simulation data management primarily uses file naming conventions, applied on a per-code basis. There is no coherent overarching data management scheme that supports the distributed development and execution environment required for the SWIM project and the FSP. Likewise, with the limited efforts to date at coupling plasma simulations codes, little need has arisen for the kind of formal programmatic interface specifications that will be required to allow flexible coupling of simulations codes and addition of new modules as they evolve. Other simulation efforts [37a] have created data management and formal interface methods and SWIM will leverage and adapt those when possible.

In its simplest form, data management consists of one administrative domain and one physical data location. In practice, as data collections have grown, they have become spread over both administrative domains and physical domains. To be useful to the scientific community, users must be able to find desired data based on metadata attributes such as the content, quality, condition, source, history, and other characteristics of the data. Once found, data need to be efficiently moved between large storage locations or between programs and storage. Efficient data movement includes replication, caching, and bulk data access. Finally, the coupling of computations with operations on data resources introduces new optimization problems, and this functionality becomes critical as data repositories grow in size.

An interoperable extensible framework for plasma simulations has four components: a community-developed code discipline, a formal interface definition methodology to clarify and consolidate cross-physics interactions, careful targeting of selected modules for ultrascale performance, and an overall data management solution. SWIM will draw on the large investment the DOE has already made in those areas in building the required software platform.

3.6 Experience in porting and optimizing on NLCF from other fields

Since the Center for Computational Sciences (CCS) at Oak Ridge National Laboratory was selected in 2002 by the Department of Energy as the site for the Cray X1 scalable vector supercomputer, codes that address important scientific problems in climate, fusion, biology, nanoscale materials and astrophysics have been ported and optimized to it [38]. On the NLCF vectorization uses fine scale parallelism to significantly improve performance over scalar

systems. In computational chemistry the Dynamical Cluster Approximation – Quantum Monte Carlo code simulates high-temperature cuprate superconductors and has achieved Cray X1 performance 20-50 times better than on the IBM p690 parallel system. This code performs many matrix-vector operations and has benefits from the Cray’s high memory bandwidth. The Parallel Ocean Program in climate studies runs 50% faster than on Japan’s Earth Simulator and four times faster than on an IBM p690 cluster. Code properties that facilitate good Cray X1 optimization include: data structures that avoid indirect references, tightly nested loops, large problem components that do not completely fit into the cache; and parallelization methods adaptable to the Cray-specific co-array Fortran and UPC libraries. For codes more suited for massive parallelism than vectorization the NLCF will also include a Cray Red Storm system by Summer 2005. The experience of computer scientists in ORNL’s Center for Computational Science (CCS) will be valuable in gaining high performance for appropriate SWIM modules on NLCF computers.

3.7 The NTCC code library

The National Transport Code Collaboration (NTCC) code modules library is an initial step in the direction of a Fusion Simulation Project. It is a library of widely used code modules, most of them fusion specific, that have been re-engineered to meet certain coding, portability, and documentation standards. A description of this can be found at <http://w3.pppl.gov/NTCC>.

The Fusion Simulation Project will make use of many of the modules in the NTCC library, but will extend the interoperability properties of these and other modules by standardizing interfaces between modules, developing tools to make complete codes out of the modules, and extending this concept to parallel modules

3.8 The Fusion Collaboratory:

The National Fusion Collaboratory Project is funded by the SciDAC program to develop a persistent infrastructure to enable scientific collaboration for all aspects of magnetic fusion research. Initiated in late 2001, this project builds on past collaborative work performed within the U.S. fusion community and adds the component of computer science research done within OASCR. The project is a collaboration itself, uniting fusion scientists and computer scientists from seven institutions to form a coordinated team. The vision is that experimental and simulation data, computer codes, analysis routines, visualization tools, and remote collaboration tools are to be thought of as network services. This view represents a fundamental paradigm shift for the fusion community. In this model, an application service provider (ASP) provides and maintains software or data resources as well as the necessary hardware resources.

The project has deployed the national Fusion Energy Sciences Grid (FusionGrid) that is a system for secure sharing of computation, visualization, and data resources over the Internet. The Grid’s resources are protected by a shared security infrastructure including strong authentication to identify users and fine-grain authorization to allow stakeholders to control their resources. In this environment, access to services is stressed rather than data or software portability. The code TRANSP running at PPPL and the code GATO running at DIII-D have been released on FusionGrid with more codes scheduled for release in the next year.

3.9 ORNL LDRD Projects directed the toward the Fusion Simulation Project

ORNL is investing in two Laboratory Director Research and Development (LDRD) projects aimed at promoting the Fusion Simulation Project and at enabling fusion codes to make optimum use of NLCF computers. They are in the second year of a two-year LDRD project entitled “Comprehensive Fusion Simulation: Component-Based Software Engineering and Evolutionary Time Advancement.” The project is using the Common Component Architecture (CCA) tools to

integrate a model transport simulation code with the high resolution All Orders Spectral Algorithm (AORSA) code. The dual goals of this project are to (1) explore computer science and multi-scale math issues associated with a new, efficient iterative algorithm for spectral wave solutions in time evolving plasma, while (2) assessing and gaining experience with CCA [39] as a standard component architecture for high performance computing. Interfaces and data flow between components is expressed using the System Interface Description Language (SIDL) implemented by the Babel language interoperability tool [40] and the components are implemented (in Fortran 77, Fortran 90, C++, C, Python, or Java) to conform to the SIDL interfaces. A prototype simulation has been developed and an iterative algorithm based on using an LU factorization as a preconditioner successfully tested for AORSA on a problem with time evolving densities and temperatures.

A second LDRD project entitled “Terascale Computations of Multiscale Magnetohydrodynamics for Fusion Plasmas,” was funded in October 2004 as a collaboration with PPPL. The goals of this project are to port and optimize the M3D nonlinear initial value MHD code and the DELTA5D (3 real space + 2 velocity space dimensions) stellarator Monte Carlo transport code to the NLCF Cray X1. The DELTA5D code will be adapted to run in the M3D coordinate system and accept time-dependent field data from M3D. DELTA5D will then be configured to calculate nonlocal electron and ion closure relations (i.e., plasma stress tensor, ohm’s law) for the MHD equations used in M3D. The merged system of codes will be applied to the neoclassical tearing mode problem. The computer science component of this project will focus on improving the performance of the two codes on the NLCF computers. For M3D this involves developing new sparse matrix solvers that efficiently use the vectorization and multi-streaming characteristics of the Cray processors. Improvements in the DELTA5D code will be based upon using the high memory bandwidth and efficient gather-scatter algorithms of the Cray.

4. Proposed research

4.1. Detailed project description

SWIM will create a new and unique capability for modeling the interactions between high power electromagnetic waves and MHD stability. This will be accomplished by building an innovative end-to-end high performance computational system that integrates emerging mathematical methods, software component and data management systems, and some of the world's most accurate and capable fusion energy codes.

This system will satisfy stringent design requirements so that it is capable of scaling to the even more complex coupled simulations envisioned in the Fusion Simulation Project. In particular it will allow efficient coupling of the full range of required physics codes or modules. It will be ported to and perform efficiently on the fastest supercomputers available. It will be flexible enough to allow exploration of various physics models and solution algorithms and to permit configuration for various control flows in the simulation architecture. It will permit convenient user access and access to experimental data. It will be robust to evolving physics, code development, and developments in computer hardware. Abstract interfaces and data structures will be defined with the goal that any new module that adheres to the abstract design can be included in the structure.

To test and demonstrate this modularity we will implement at least two of each of the major high performance code types within each of the CEMM and the CSWPI SciDAC projects, as well as several codes and modules that have been developed in other parts of the fusion program.

To test and demonstrate the flexibility to allow different control flow in the simulation and to demonstrate robustness to physics evolution we will carry out physics campaigns addressing both fast MHD events and slowly growing MHD modes. These involve different modes of interaction between RF and MHD and imply different interconnections between components and different control flow.

The campaign addressing long timescale discharge evolution in the presence of sporadic fast MHD events will involve developing within the framework an integrated Evolving Equilibrium Model (EEM) that is interfaced to the 3D non-linear extended-MHD codes. The EEM will contain modules to calculate wave propagation and absorption in all relevant frequency regimes, modules to calculate the modification of the plasma velocity distribution from sources (RF, NBI and fusion alpha-particles), modules to calculate plasma profile evolution and magnetic evolution assuming closed flux surfaces, linear MHD stability models and reduced models of non-linear MHD events. This program can largely be accomplished by interfacing existing, autonomous codes, with some re-factoring and code modification. Such a system will contribute to the development of optimized burning plasma scenarios and will also facilitate the optimum use of terascale extended MHD codes and terascale RF codes for a variety of studies. The primary physics focus will be to improve the understanding of how RF can be employed to achieve long-time MHD stable discharges, with an emphasis on controlling sawtooth events. However, the tools developed and the software infrastructure will also be useful for the study of other fast MHD events such as energetic particle modes or edge localized modes. The simulation architecture for this campaign, including the EEM is illustrated in Fig. 5.

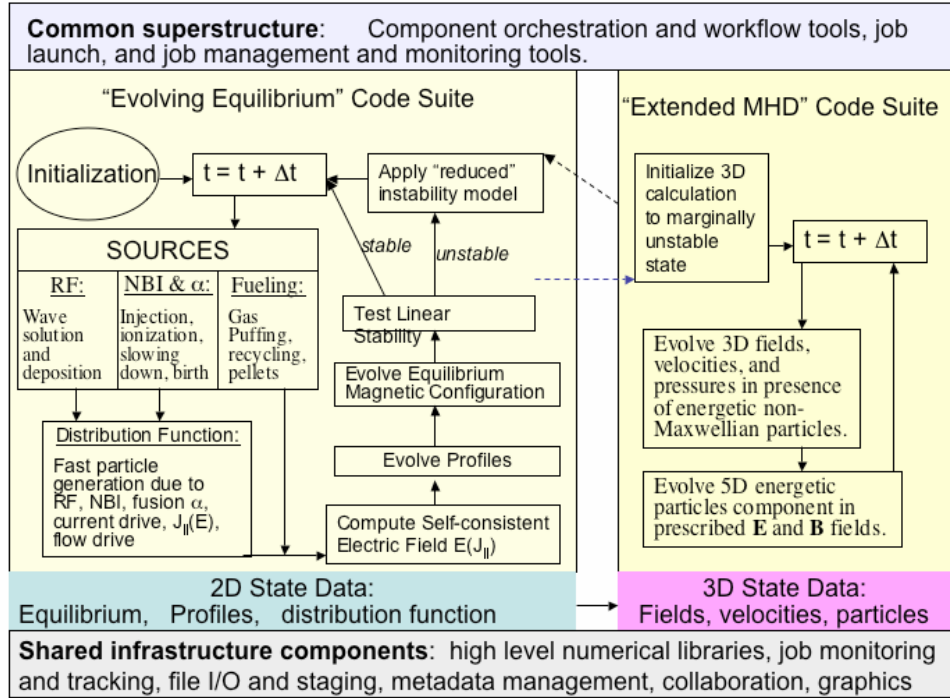


Figure 5: Evolving Equilibrium Model and Extended-MHD Code architecture

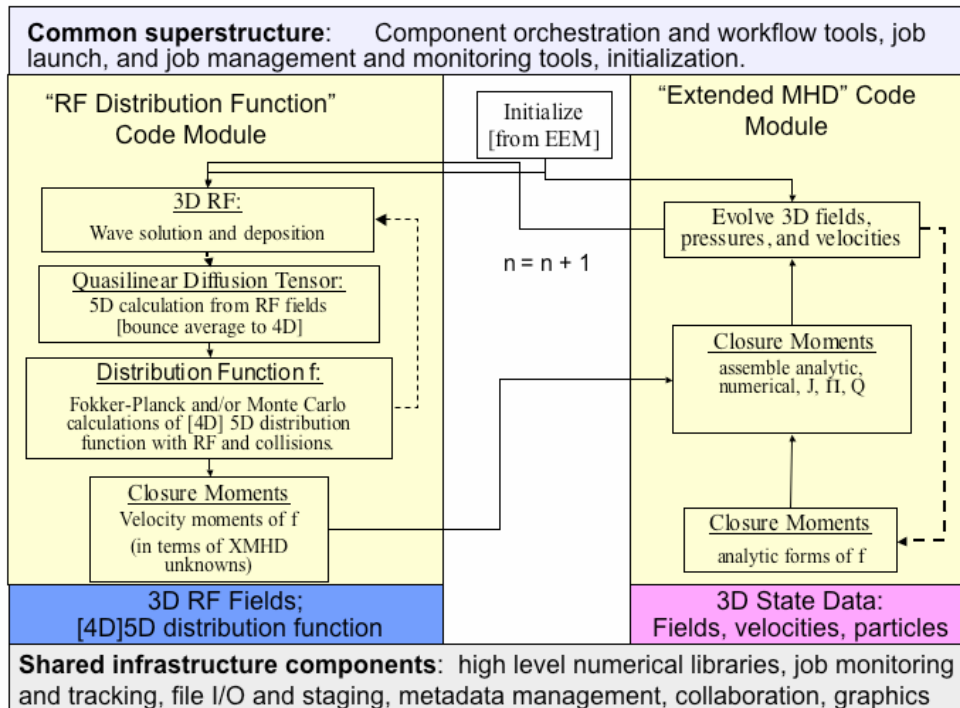


Figure 6: Elements of a simulation framework for slowly evolving MHD phenomena

The campaign for modeling the direct interaction of RF and extended MHD for slowly growing modes will require interfacing RF modules directly with the extended MHD codes and with code modules that implement fluid closures for the MHD equations. The RF field solvers and Fokker Planck modules for evolving the distribution function will be generalized to non-axisymmetric magnetic fields containing magnetic islands as generated by the extended MHD codes. To properly create such coupling and to provide flexibility to explore various closure models and solution algorithms, we will have to go beyond existing autonomous codes.

Instead, it will be necessary to perform some code development and to re-factor some of the codes into module sets that can be more tightly coupled and controlled by the computational framework. The primary physics focus will be to improve the understanding of how RF can be employed to control neoclassical tearing modes. However, the tools developed will also contribute to understanding other slowly evolving, non-axisymmetric phenomena, and could also be applied to simulate the sawtooth event if the assumptions involved in the Evolving Equilibrium description break down. The simulation architecture for this campaign, including the EEM is illustrated in Figure 6.

The computational science research entailed by SWIM includes defining a software architecture (a set of components and their interactions), engineering a suite of complex multi-physics codes to have scalable performance on widely differing architectures, and building an integrated data arena that supports component interactions, automatic file management, and an distributed shared metadata system. Solvers specific to plasma physics need to be designed and implemented, and the numerical and systems issues involved in passing complex physics objects with different representation paradigms (finite element, particle-in-cell, AMR, finite differences) between physics modules need to be resolved.

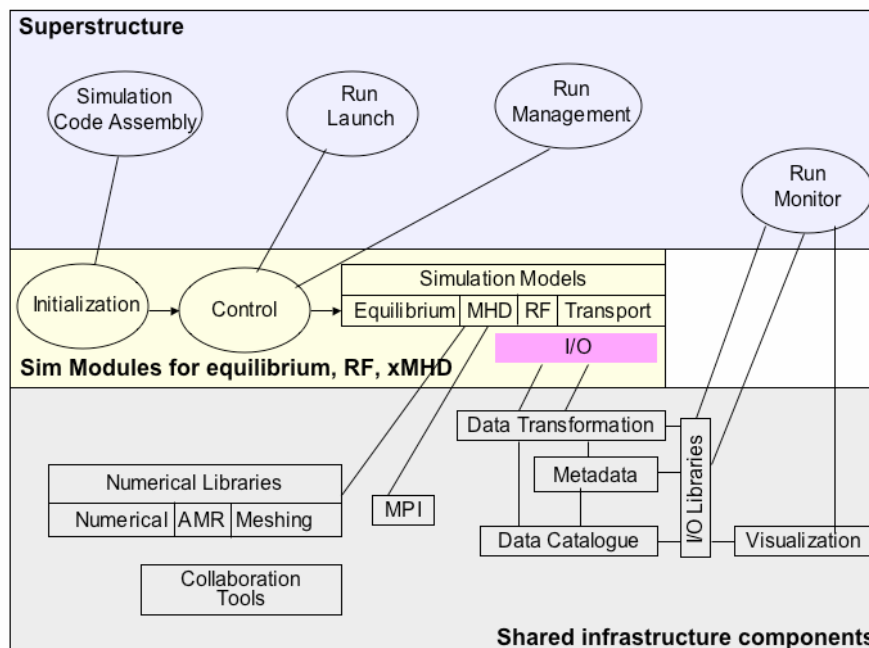


Figure 7: Functionality is in the control superstructure, the simulation modules, and the supporting infrastructure.

4.1.1 Elements of work

The tasks necessary to achieve these goals are divided into multiple elements of work, to be carried out in parallel when possible. Each element requires the combined efforts of computer scientists, mathematicians, and physicists, and each element is described in more detail in the indicated section.

1. The general architecture for the computational system will be defined. This involves analyzing the fast and slow MHD simulations categorizing the functionality into three categories: control superstructure, simulation level, and shared infrastructure [Section 4.1.2]
2. Specific components and interfaces for fast MHD will be defined. This involves a detailed analysis of the data flow, communications, and control requirements of this class of simulations. All participants must agree on standard interfaces between computational elements, so that different implementations can easily be swapped in and out. The process for modifying and extending these standards will be developed and agreed to. [Section 4.1.3]
3. The physics components will be modified and improved to incorporate the interaction between RF and extended MHD, or in the case of the fast MHD campaign between RF and profile evolution. This entails new analysis, coding and testing of these improvements. [Section 4.1.4]
4. The architecture will be extended to encompass the requirements of the slow MHD campaign. Slow MHD will require more extensive and deeper modification and re-factoring of existing codes and will require preliminary integration testing and experimentation before defining interfaces. [Section 4.1.5]
5. An implementation of the architecture will be created, first supporting the fast MHD campaign, and then the slow MHD, with the physics codes and modules modified to support the standard interfaces and control protocols [Section 4.1.6], and basic shared superstructure and infrastructure components created [Section 4.1.7] including a system to manage the data (source codes, I/O files, experimental data for validation and verification, performance analysis and results, simulation metadata).
6. Identification and rectification of performance bottlenecks will be carried out at all stages, targeting both MPP and vector computers. This will require first optimizing chosen existing codes and later inter-component, overall optimization. [Section 4.1.8]
7. SciDAC math centers tools will be adapted and implemented, efficient preconditioners for fusion simulations investigated, and mathematical issues related to stability, accuracy and performance analyzed. [Section 4.1.9]
8. At each stage the scientific validity of the simulation will be established by benchmarking of the multiple codes and by comparison with experiment and carrying out new research within our two physics campaigns of fast and slow extended MHD. This element also encompasses commitment to engage the wider fusion simulation community and develop a user base as early in the project as feasible [Section 4.1.10].

4.1.2 *The SWIM computational architecture*

A well-defined architecture and implementation plan is fundamental to coordinating a large-scale geographically distributed multidisciplinary effort like SWIM. Designing a software architecture requires a detailed understanding of both the physics involved and how it gets done computationally and therefore requires close collaboration of physicists, mathematicians, and computer scientists. To provide long-term flexibility and extensibility, the SWIM architecture is based on principles of modularity and reusability. The functionality provided by modules (or *components* in modern software terminology) of any sophisticated simulation code can be divided into three basic categories, depicted in Fig. 7.

- The control *superstructure*, which orchestrates the simulation and provides capabilities to interface with external tools for job launch and monitoring, and connect to remote data repositories, and manage simulation workflows. SWIM alone has several workflows to support; the FSP will have many more.
- The domain-specific (in this case plasma physics) *simulation modules* which individually or in concert (controlled by the superstructure), perform the desired simulations. The majority of simulation modules will be derived from existing physics codes. Simulation components for the fast MHD evolving equilibrium code will be relatively coarse grained and will follow closely the general outlines of the current standalone codes. For the slow MHD campaign, we anticipate finer-grained components which will often require refactoring of existing codes at a deeper level. Some new simulation modules will also have to be developed, particularly for the MHD closure models.
- The supporting *infrastructure*, which provides shared computational and numerical services for use by multiple simulation modules and superstructure elements. Infrastructure components are highly varied, including physics-related modules which are used by multiple simulation modules, numerical libraries, and other purely computational services. The majority of SWIM infrastructure components will be based on existing libraries and tools. New development will be primarily capabilities related to data and metadata management.

This tripartite architecture has been used in other large-scale software systems, including the NWChem parallel computational chemistry package [41, 42], and has proven to be a powerful conceptual tool to insure maximum flexibility and code reuse. The terminology of superstructure and infrastructure is borrowed from the Earth Systems Modeling Framework (ESMF) [43] although the SWIM architecture differs in some ways from ESMF.

Our intent is to specify the SWIM architecture at the highest possible levels of abstraction in order to provide maximum flexibility in how it is actually realized in software. We will evaluate SWIM's requirements and detailed architecture against a variety of available computational frameworks. An upcoming workshop bringing together fusion researchers and framework developers [44] will provide input to the decision process. However our current expectation is that Common Component Architecture (CCA) [38,45,46] with which Bramley and Bernholdt have extensive experience, will provide the required flexibility and extensibility and is the leading candidate for this role.

4.1.3 Interfaces in the SWIM computational architecture

The three-way classification gives a way of organizing the components, but an architecture also requires defining how components *interact* with each other and with the outside world in terms of the flow of data, communications, and control. To insure flexibility and extensibility, these interactions must be abstracted and standardized within the community, so that different implementations of a given functionality can easily be swapped in or out.

We will begin by establishing the process by which these standards can be defined, modified, and extended, drawing on the experience of other scientific communities in developing common interfaces, such as the work of the SciDAC TOPS [34] and TSTT [33] center in establishing common interfaces for linear solvers and unstructured mesh representations, and standards such as the Message Passing Interface (MPI) and High-Performance Fortran (HPF). The process will be followed systematically to define the standard interfaces needed throughout the SWIM architecture, adopting existing interfaces wherever feasible. Interfaces will be expressed using the Scientific Interface Definition Language [47] (SIDL), which provides a programming language independent form of expression, and define key data structures as necessary. A variety of approaches are available to represent data structures and file formats, such as the netCDF Markup Language [48] or Common Data Language [49], the HDF5 Data Description Language (ASCII [50] or XML [51]), the Grid Forum's Data Format Description Language [52] or the

Fusion Simulation Markup Language [53]. We will evaluate the various approaches in this context and use it to express data structures and file formats associated with the various components. The mathematical types, frequency of exchange, and sizes of many of the data objects that need to be defined for simulation components are estimated in Appendix B.

Our initial focus for the interface design work will be the components required for the fast MHD campaign.

4.1.4 Modification and improvement of physics components

Fast MHD Campaign

We anticipate that most physics components for the campaign on fast MHD events can be based on existing codes with little additional physics development. However, new development will be needed to obtain solutions to the Fokker Planck solutions that include more accurate models of RF-driven transport for energetic particles. Each of the three complementary approaches to Fokker Planck solution requires some development.

- Simplified models for the radial diffusive motion and bootstrap current effects presently in CQL3D [24] will be replaced with physics-based models
- A velocity-space dimension to account for energy scattering will be added to DKES [25] to include the effects of RF quasilinear diffusion. Within the EEM, where the equilibria are considered axisymmetric, this can be done without increasing the overall dimensionality of DKES.
- For Monte Carlo codes [26] the performance issues related to obtaining adequate particle statistics will be addressed on the NLCF computers. Methods of maintaining consistency relations related to conservation of macroscopic momentum will be investigated. By using appropriate “background” distributions, it should be possible to compute distributions which are consistent with the macroscopic MHD flows, although this remains to be demonstrated.

Slow MHD Campaign

Unlike the fast MHD campaign, there is no clear separation of timescales between the plasma response to RF and the evolution of the extended-MHD variables. However, because the RF and bounce times are short, we can split time the advancement of an instantaneous MHD state and the determination of the self-consistent RF electric field and particle distribution function as indicated schematically in Figure 6. Moments of the distribution function, modified so as make them functions of the independent, fluid-like extended MDH variables, couple the RF to MHD module and the MHD state variables are used to determine the RF. Predictor corrector techniques may be required to enhance stability. Proceeding with this campaign on slowly growing modes will require three primary areas of physics development:

- Extension of wave solvers to operate with the perturbed (3D, time dependent) plasma state,
- Development of modified MHD closure models to include the effects of RF in order to couple the results of the RF-driven distribution function back to extended-MHD.
- Improved Fokker Plank solvers to incorporate essential 3D geometric features such as magnetic islands.

Extension of wave solvers

Interfacing RF to the 3D plasma state requires essentially the same data as specified for the EEC except now in 3D; and it must be obtained from the extended-MHD solver rather than from a 2D equilibrium solver. In the initial phases, we will focus on ray tracing using the GENRAY code for ECCD. In the longer term, we will adapt the full-wave codes, which will enable us to study

minority current drive or lower hybrid current drive. The AORSA3D code is capable of calculating ICRF wave deposition in these non-axisymmetric geometries but is computationally intensive.

Modified closure relations for extended MHD including RF interactions

CEMM is exploring a range of closure models to explore the physics of slow MHD in the absence RF. We will build on that effort with an emphasis on how RF affects the closure physics. Hybrid closures which employ direct solutions for the RF-driven distribution function and required RF modifications to the “background” components are the focus of our proposed research.

Hybrid Closure through direct solutions of the distribution function A direct solution for the distribution function of the plasma species that is most strongly coupled to the RF is needed. Moments of this distribution function can then be used to update corresponding fluid terms in the extended-MHD equations. As described in Section 3.4, the parallel RF-driven current must be “inverted” to provide the electric field as a function of the current to allow developing an Ohm’s Law that can be used to eliminate the electric field from the extended-MHD equations.

CEMM Closures To specify the closure variables, we will investigate the different methods considered in the CEMM effort, modified by the contribution of the RF quasilinear operator. These include the following: local collisional closures, where the standard short mean-free-path expansion of the distribution function about a Maxwellian is used [54,55,56]; parallel integral (non-local) neoclassical kinetic closures, where the parallel viscosity and parallel heat fluxes are obtained as solutions of a free-streaming drift-kinetic equation [25,57]; collisionless parallel local closures, where local dynamic evolution equations for the parallel viscosity and parallel heat fluxes are considered [58,59]; analytic neoclassical closures, where the parallel viscosity is given by its neoclassical expression in a closed-magnetic-surface near-equilibrium [60].

Improved Fokker Plank solvers

Much of this work will overlaps with a similar task in the fast MHD campaign. However the presence of 3D structures substantially complicates a rigorous solution of the Fokker Planck equation, even when cross-flux-surface transport is neglected. Specific additions include the need for considering 3D modifications imposed by magnetic islands and the need for high spatial accuracy because derivatives of the parallel electric field are required for the extended-MHD codes. Specific considerations for the CQL3D code include taking the bounce average along the field line at the island center, but with power deposited on the actual island flux surfaces and the volume averages taken between incremental volumes of the helical flux surfaces of the island.

4.1.5 Extending the SWIM computational architecture to slow MHD

From a computational viewpoint, slow MHD simulations differ from fast primarily in the frequency of exchange of large data objects between components, and potentially by using domain modules that are newly defined from refactorization of existing physics codes. Both will require a much deeper and finer-grained integration of physics modules. This will require new interfaces, and probably some revisions to increase the generality of interfaces that were developed focusing on the fast MHD needs. This is a natural part of the process of interface design, and the reason that the *first* task in the interface design effort is to define the process by which these standards are decided *and modified*. Throughout this process, we will also be looking for opportunities to reuse code by refactoring simulation modules and pushing the part of the functionality that can be used by multiple components into the infrastructure layer. The process for developing slow MHD components and interfaces will be fundamentally the same as that used by fast MHD, but will lag somewhat behind it, so that the fast MHD architecture can be treated as a prototype, and testing and experimentation can help in both defining and implementing the additional capabilities.

4.1.6 Initial architecture implementation work plan

The SWIM project will utilize software packages which are undergoing continual development and use by a larger community. Our strategy will be to work directly with the main stream of development and insure that modifications introduced when code is turned into a SWIM component will also work appropriately when the package is used separately from SWIM. We have identified physics code “managers” for the first five key applications we will be using – individuals who have been involved in the development of the standalone package, and who are also participating in SWIM:

- M3D: Steve Jardin
- NIMROD: Dalton Schnack
- AORSA: Don Batchelor
- CQL3D: Bob Harvey
- TRANSP: Doug McCune

The first stage is for each code manager to create a test and benchmarking suite for the code. The suites must be extensive since existing codes have a relatively large number of potential execution paths which exercise different parts of the code. Together with Bramley and Bernholdt, the managers will define and agree upon common software aids for configuration and source code management [61-66] test harnesses [67-70], and bugtracking [71,72] that will be used across the SWIM project. Evolution and changes in tools will be implemented only with the consensus of the software process team.

The second stage (contemporaneous with the first) is definition of the required interfaces for physics simulation level components in fast MHD, as described above and will involve all project members.

The third stage, once the components and their interfaces are suitably defined, is creation or adaptation of auxiliary components essential for fast MHD. We will define file I/O components first, because that will be the mechanism for exchanging data between simulation level components initially, allowing rapid prototyping and identification of potential numerical stability or physics problems before making larger changes in the codes. Using file I/O will be superseded by “live” component connections by the end of the first year. Another component that will be essential is a simple job harness for composing a SWIM application from components. Initially this will be a simple script in Python, PERL, or IDL, with a more sophisticated user interface using Web portals [73-76] coming later. Preliminary infrastructure components for handling mesh mediation and data transformation will also be created in this phase.

Later stages will consist of refining the architecture, following the processes previously defined for establishing and evolving interfaces, and development and adaptation of additional supporting components.

4.1.7 Required superstructure and infrastructure Components

Most of the infrastructure components in SWIM will be leveraged from existing DOE and other activities, either directly using other products or by adapting them to SWIM’s needs. While many details of the superstructure and infrastructure will only be defined as part of the design process at the beginning of the project, we can outline some of the aspects which are already clear.

Software interfaces are primarily used for interactions between computational components, but SWIM must also support data access, transfer, and management. The requirements include efficient transfer of large simulation datasets, a metadata catalogue allowing rapid data discovery, and the ability to seamlessly tie into existing experimental data repositories.

Physically, SWIM's data management system will be distributed over multiple computer systems at several sites, yet it will be presented to the user as one logical construct to facilitate ease-of-use. SWIM's shared infrastructure components will include the ability for job launch, real-time monitoring, and management. Large data sets produced by SWIM will require visualization tools, and the SWIM project's geographical distribution needs team members to work together effectively while in separate locations, including shared coding and desktop application sessions.

Data management, job management, visualization and collaboration are being addressed by DOE projects (some listed in Appendix C) and SWIM will leverage those. Significant work is required to create an integrated data management system even from the existing base, and other work must be done to adapt and harden the other utilities for fusion simulations; if necessary in later years of the project resources may be directed to work on expanding them to meet emerging needs.

4.1.8 Performance engineering

Portable performance across a broad spectrum of platforms is important to SWIM, and we will make a significant investment in insuring high performance. The first phase of this effort will involve the creation of testing and benchmarking suites for the various codes in order to facilitate identification of performance hotspots. The initial performance engineering work will target M3D, NIMROD, AORSA, CQL3D, TRANSP, with other codes added later. This phase of the work is shared with correctness test suites in the architecture design task (Sect. 4.1.6).

The second phase is the porting and building of each code on three major target architectures: workstation, MPP cluster, and vector. In particular portability will be targeted at each of the SWIM sites, including the National Leadership Computing Facility at ORNL. A test harness that allows regression testing of the codes on all of the architectures will be chosen and implemented.

The third phase is detailed performance analysis of the codes on each architecture, using standard tools. Simple basic block (function and subroutine) profiling will identify where the time is spent in each code module (likely to be different on different architectures). Then tools such as Jumpshot, Vampir, Paraver, Dimemas, and Paradyn can then be used to provide more detailed loop- and statement-level analysis of particular hotspots. Because of our close connections with the developers, the TAU performance toolkit [77] will be used throughout the SWIM project to instrument code for performance analysis. Multiple platforms and tools are essential to develop a balanced and detailed view of performance problems and suggestions for enhancements. In particular, parallel and other I/O is often highly site-dependent and so requires a broad base for analysis. Our performance engineering work will draw upon the extensive base of experience developed across many platforms by the ORNL Scientific Applications Group (of which Fahey is a member). We will also be able to draw on the newly established Cray Supercomputing Center of Excellence (CSCE) for special assistance in porting and tuning for the Cray systems. The CSCE is located at ORNL and is headed by John Levesque.

Use of standard high-performance libraries like Trilinos, PETSc, and ACTS tools whenever possible in the codes will also be promoted within SWIM. This allows us to take advantage of the more extensive development and optimization such packages receive due to their large user bases. An initial analysis indicates many opportunities to improve code performance this way.

The preliminary phases of our effort target intra-component performance. A new issue that SWIM introduces is *overall composed application performance*. Even well-tuned individual components may exhibit new performance issues when they are coupled together. They may be

due to load imbalance, the fact that algorithms may not perform equally well on all architectures, or other factors. Possible solutions we will explore include static load balancing through more effective allocations of processors to various components of the simulation, running simulations distributed across multiple systems to permit each component to run on its optimal architecture, and dynamic reassignment of components to processors using tools such as Zoltan [78]. Analysis of performance in this setting will rely on TAU, which provides unique capabilities to capture component-based computing performance across multiple architectures in a single run. [79]

4.1.9 Numerical mathematics issues

Utilizing SciDAC ISIC outcomes. A large amount of groundwork has been laid for the SWIM project by existing SciDAC Applied Math projects. TOPS software for large-scale algebraic systems that arise from finite discretizations of PDE systems is already employed in M3D and NIMROD at the highly standardized but relatively low level of sparse linear systems. TSTT is building a variety of sophisticated, scalable mesh generation and discretization systems. Equally important, TSTT provides interfaces to multiple discretization schemes and could provide the ones needed to exchange data between, e.g., nodal-value finite difference methods and higher-order finite elements. APDEC uses finite-difference methods on hierarchical structured grids combined with block-structured adaptive mesh refinement (AMR), and multilevel schemes like AMR are critical for achieving fine spatial resolution where required without overwhelming memory resources. While the products of these centers are already in some SWIM codes they will require extensions to meet SWIM needs.

Application-specific Solvers. Even when employing libraries like PETSc that richly integrate algebraic and multilevel preconditioners, superior convergence may be obtained by an application-specific preconditioner, and PETSc recognizes this by allowing users to register preconditioners that allow call-backs to exogenous application subroutines. Many MHD codes currently solve multiple scalar elliptic PDE's, which are ideal for multilevel methods. However for adaptive unstructured grids, the recurrences within standard smoothers and preconditioners do not vectorize well because indirect addressing prevents efficient deep pipelining. SWIM will therefore explore for vector architectures sparse approximate inverse (SPAI) preconditioners [80], which have the added cost of solving large numbers of smaller least-squares problems to create the preconditioner. SPAI will be effective if the cost of creating the preconditioner can be amortized over several solves, and this requires evaluation in the context of specific SWIM applications.

Multiphysics Coupling. For the coupling of SWIM codes into multiphysics packages, the standard approach is mathematically equivalent to nonlinear Gauss-Seidel iteration between the state variables of the component codes. For instance, the outputs of one code may be mapped into the inputs of another code (typically in the form of coefficients or source terms), and two or more codes may then be iterated to consistency, to attain an equilibrium state or within each timestep of a transient problem. Such iteration often requires under-relaxation in practice and determination of convergence must be made without the benefit of a global norm. Once a nonlinear Gauss-Seidel iteration of this type is set up, however, it is natural to consider a global Newton iteration on the union of the state variables between all of the component codes. It is possible to define such an iteration without requiring that a global Jacobian matrix be created and operated upon; instead the action of the Jacobian on the global state vector can be obtained by calls to the component codes with special inputs. The resulting action is an inner iteration in a Jacobian-free Newton-Krylov method for the coupled system. It can be preconditioned in a variety of means, linear and nonlinear, including the method of nonlinear Schwarz, which has been explored successfully in simple problems by TOPS. Since the natural evolution from multiple codes to a single multiphysics code, as will be required by the FSP involves reuse of as much of proven existing codes as possible, this investigation within the SWIM project holds great infrastructural promise

Transferring Data Between Discretizations. Each existing physics code has embedded assumptions about the underlying grid. TSTT's common mesh interfaces can partly address the computer science problems this creates, but in an integrated simulation those assumptions can lead to discretization errors and a nonconservative overall model. Minor differences in representing the same fields might cause numerical instability in the fully coupled simulation. Another danger is artificial numerical diffusion from iterated cross-interpolation. Accurately interpolating data, remediating mesh inconsistencies, and computing consistent gradients between the codes is essential. Individual physics modules should not be required to perform this task, since this is an issue of coupling. Implementing these tasks as components will be a unique challenge in SWIM that may have far-reaching payoff for other coupled simulations.

4.1.10 Scientific validation and execution research program

Scientific validity of the simulation must be established by comparison with analytic theories in appropriate limits, by benchmarking between codes, and by comparison with experiment. Our strategy of including multiple realizations of each major component will greatly improve our ability to carry out meaningful comparisons between different code modules under the same operating conditions. The access to experimental data afforded by connecting to the fusion grid and utilizing the existing TRANSP utilities will improve our ability to validate the simulation through comparison with experiment and to support the experimental programs.

The primary focus of the campaign on fast events will be improving the understanding of how RF can be employed to achieve long-time MHD stable discharges, with an emphasis on the use of RF power for practical control sawteeth and using RF to probe the physics of the sawtooth. Outstanding questions we will focus on in this area include:

- How can we influence the period and inversion radius of the sawtooth oscillation by modifying the plasma current and pressure profiles by RF and by creating an energetic particle component?
- Under what conditions will the sawtooth instability trigger the onset of a metastable island (seed island for a neoclassical tearing mode) or lead to a disruption, and can this be prevented by the application of RF?
- How can the “reduced sawtooth model” now in use [81] be extended to account for RF effects?

Other important fast MHD events are also influenced by the presence of RF power such as the various energetic particle instabilities such as fishbones and toroidal Alfvén eigenmodes (TAE). In addition RF current drive at the plasma edge is an important consideration in controlling Edge Localized Modes (ELMs) that arise in high performance discharges due to steep pressure and current gradients at the edge. We will address these to the extent that resources allow, but in any case the SWIM simulation capability will be valuable to other researchers in the community for modeling these phenomena.

We will strive to make the Evolving Equilibrium Model suite the “community standard” for developing optimized burning plasma scenarios. We will have an outreach program to engage modelers not presently in the SWIM project to become code users, and co-developers. A measure of success for our project will be in how successful we are in drawing in outside users, and how productive they are in performing scenario optimizations.

The primary focus of the campaign on slow MHD events will be on the use of ECH power for practical control of neoclassical tearing modes and using RF to probe the physics of the NTM. Outstanding questions we will focus on in this area include:

- How does the saturated island width depend on the power, steering accuracy and modulation of the ECRH power?
- What power level, and feedback approaches are needed to fully stabilize the mode?
- Can detailed coupled RF/MHD calculations be used to obtain more complete and qualitative understanding than the results from the semi-analytic Rutherford equation models? This will lead to improved reduced models for NTM effects that will be included in the EEM.

Success in this campaign will contribute to understanding other slowly evolving phenomena that overlap the RF quasilinear timescale. Particularly important would be understanding the effect RF on the evolution of the 3D plasma present during the sawtooth ramp phase for cases where the sawtooth reconnection is incomplete and a residual magnetic island remains, or when other assumptions utilized in the evolving equilibrium description are not realized.

4.1.11 Summation

The work described in this proposal builds upon the successes of the DOE SciDAC programs by taking several of the most advanced fusion computer codes, combining them to provide a unique tool in the worldwide fusion program, and implementing them on the Leadership-Class and other computing facilities. The software architecture defined will lay a foundation for even more powerful and comprehensive simulations in the future.

We have identified two regimes of potential plasma global instabilities in the presence of strong RF: fast phenomena and slow phenomena. In the first, the wave fields act only indirectly on the instabilities by modifying the background equilibrium, while in the second they act directly by competing with the instability drive terms. The two physics campaigns will target these two classes of instability, and the software architecture will be general enough to accommodate both of these.

The fast MHD campaign will utilize a 2D evolving equilibrium model, and will only invoke the 3D codes when an instability threshold is crossed. This results in an enormous increase in computational efficiency, allowing long-time integration on the plasma current-diffusion timescale. The slow MHD campaign will involve full 3D coupling between the RF and extended MHD codes. This requires the development of new closure techniques to couple the effects of RF into the moment equations used to advance the fluid variables.

Eight work elements have been defined in section 4.1.1 and subsequent sections. This software engineering approach is required in order to carry out the ambitious objectives of this proposal. It should also provide a baseline project structure for future software engineering projects in the FSP of even greater complexity.

Throughout this process we will work with the community to develop standards and documented APIs for fusion application modules which will allow the incorporation of many of today's fusion application codes into this framework.

From a broader perspective, the project described here is a huge step for fusion science. For the first time, we are merging two major, mature, mainline efforts: Plasma-Wave interaction and Global Stability; to produce a software package of a scale in power and complexity that is unprecedented in our field. Just the commitment to take this inevitable next step is forcing us to deal with such difficult issues as fluid closures and how they are influenced by RF and stability physics. The common vision of what needs to be accomplished will drive the project, and unite the RF and extended MHD researchers with those from computer science and mathematics. This is breaking new ground in fusion science research, and is destined to lead to an accomplishment that we can all be proud of, and will help propel fusion science to the next level.

4.2 Project Schedule, Milestones and Deliverables

Work Element	Year 1	Year 3	Year 5
1. Definition of general architecture	<ul style="list-style-type: none"> Complete abstract architecture specification 	<ul style="list-style-type: none"> Extensions for slow MHD 	<ul style="list-style-type: none"> Generalizations to an FSP
2. Definition of components and inter- faces for fast MHD	<ul style="list-style-type: none"> Adopt interface standardization procedure. 	<ul style="list-style-type: none"> Fast MHD component interfaces completed 	<ul style="list-style-type: none"> Completed
3. Modification of physics components	<ul style="list-style-type: none"> Complete analysis for CQL3D and DKES extensions Theoretical formulation of closure techniques for the different numerical approaches 	<ul style="list-style-type: none"> Complete modifications of FP codes for fast MHD Modification of MHD codes to use new closures 	<ul style="list-style-type: none"> Complete development of RF/MHD closures, extension of RF/FP components to 3D Comparative studies of results using different RF closures
4. Architecture extended for slow MHD campaign	<ul style="list-style-type: none"> Analysis of required functionality for additional components 	<ul style="list-style-type: none"> Interfaces for slow MHD adapted 	<ul style="list-style-type: none"> Fully integrated slow MHD simulations
5. Architecture implementation	<ul style="list-style-type: none"> Software methodology chosen and implemented M3D, NIMROD, and AORSA codes ported to MPP platforms 	<ul style="list-style-type: none"> Interfaces implemented. Initial cross-module data transforms Addition of data management and job management modules 	<ul style="list-style-type: none"> Fully integrated simulations, including distributed data management
6. Performance engineering	<ul style="list-style-type: none"> Performance analysis on basic block and loop levels for M3D, NIMROD, AORSA, CQL3D Scalar code optimizations 	<ul style="list-style-type: none"> Vectorization to port modules to NLCF Analysis and enhancement of slow MHD modules 	<ul style="list-style-type: none"> Network transfer analysis Cross-module performance analysis and optimization
7. Math analysis	<ul style="list-style-type: none"> Analysis of using ISIC outcomes 	<ul style="list-style-type: none"> HPC libraries introduced Begin to implement application-specific solvers 	<ul style="list-style-type: none"> Cross-module numerical analysis
8. Validation and physics research campaigns	<ul style="list-style-type: none"> Develop individual test suites for use on different platforms. Preliminary fast MHD runs 	<ul style="list-style-type: none"> Complete EEM release for community testing Produce a publication on the physics of RF inducement of sawteeth Design optimized scenarios involving sawtooth physics 	<ul style="list-style-type: none"> Produce a publication on the physics of RF stabilization on the NTM Complete fast and slow MHD research campaigns, Begin new campaigns based on energetic particles and or ELMs

4.3 Description of work assignments

Oak Ridge National Laboratory (ORNL) is the lead institution for SWIM and will carry out the management plan as in Section 6. The ORNL will be involved with all aspects of the technical work, particularly: RF physics (Batchelor, Berry, Jaeger); Monte Carlo solutions (Hirshman, Spong); architecture design and implementation (Bernholdt, Elwasif) and performance engineering (Fayhey)

The Princeton Plasma Physics Laboratory (PPPL) will be involved with all aspects of the technical work, particularly: MHD physics (Jardin, Park, Fu, Breslau, Chin); transport physics (McCune); architecture design and implementation (McCune, Klasky) data management (Klasky) work will concentrate on the building of the EEC suite, implementing the new closure models in the M3D code, the M3D applications involving sawtooth and neoclassical tearing modes, and on certain aspects of data management and code monitoring.

SAIC (Schnack) will provide management for the NIMROD team and participate in interface design and adaptation of NIMROD

Indiana University (Bramley) will be responsible for architecture design and overall coordination of the computer science activities.

The Massachusetts Institute of Technology will participate in development and testing of extended-MHD closure models in the presence of RF (Ramos), and M3D physics (Sugiyama)

Columbia University (Keyes) will have primary responsibility for identifying where existing HPC libraries can be used, and also oversee investigation of application-specific solvers, and numerical analysis for multi-physics coupling

CompX (Harvey) will be responsible for physics development and application relative to the CQL3D code and will participate in interface definitions

General Atomics (Schissel) will have primary responsibility for file I/O definition mechanisms will oversee the investigation of testing and extending MDSplus to use parallel data streams, and NFC tools that aid in job management and monitoring

University of Wisconsin (Hegna) will participate in the analysis of MHD closure relations including RF and NTM physics.

New York University (Strauss) will participate in M3D physics

5. Budget summary

Institution	Primary work elements	Year 1	Year 2	Year 3	Year 4	Year 5	Total
ORNL	all	760	760	760	760	760	4975
PPPL	all	585	585	585	585	585	2925
U. Indiana	1,2,4,5,8	110	81	84	87	90	455
MIT	3,8	50	50	50	50	50	250
GA	1,2,4,5	150	75	75	75	75	450
SAIC	2,3,4,5,8	40	40	40	40	40	200
Columbia U.	3,5,7,8	40	55	55	53	50	253
CompX	2,3,5,8	75	75	75	75	75	375
Subcontracts and Students	2,3,5,7,8	200	275	275	275	275	1300

In order to support inclusion of codes into SWIM not specifically developed by present SWIM participants, to allow management flexibility for responding to new developments during the project and to support students and post docs a budget category Subcontracts and Students is included amounting to about 13% of the total budget. This appears in the budget section under ORNL as subcontracts.

6. Management Plan

Project management team: D. Batchelor will be the Project Principal Investigator (PI) with three Co-Principal investigators – S. Jardin (PPPL), R. Bramley (U. Indiana), D. Keyes (Columbia U.)

The PI and Co-PIs will have dual responsibilities to coordinate both within a scientific discipline and also for a project element:

Jardin –Extended MHD and fast MHD campaign

Bramley – Computer Science and Architecture development

Keyes – Applied Mathematics and Algorithms.

Batchelor – RF-Plasma Interactions and slow MHD campaign

In addition to these project responsibilities, the Co-PIs also are responsible for coordinating SWIM activities with corresponding communities, particularly the relevant SciDAC projects, including CEMM (Jardin), CSWPI (Batchelor), Computer Science ISICS (Bramley), and Mathematics ISICS (Keyes). In addition Batchelor will be responsible for coordination with NLCF and NERSC.

Advisory Committee: An advisory committee (AC) will be appointed by the PI and Co-PIs to ensure that the wider community has a voice in setting project direction and priorities. This committee will contain a representation from the fusion experimental and modeling community to help us ensure that the interface standards adopted will be extensible to a broad range of community codes.

Roles and Responsibilities: The PI will have responsibility for overall project coordination with authority over resource allocations. The PI will have responsibility for formal interactions with DOE. To the maximum extent possible, actions will be based on consensus among the PI and Co-PIs. PI and Co-PIs will coordinate technical work across the project in their discipline and element of responsibility and assign priorities for work on particular tasks. Task Co-PIs will plan the tasks for project elements, monitor the progress of team participants. The PI and Co-PIs will jointly identify project issues and develop plans to address them. Input from the AC will be sought as appropriate.

Resource Management: Funding project participants for the first year of the project will be as presented in this proposal. A budget item for new subcontracts has been established in order to optimize progress as the project develops and research needs evolve. Subcontract allocations will be made by the PI/Co-PI management team. Timely out year budget planning will be developed by this team in order provide input to DOE and to project participants for budget proposals. Project participants will provide quarterly estimates of spending to the PI.

Project Meetings and Reporting. Project wide meetings will be approximately three times a year. At least one will be a dedicated project meeting involving the advisory committee. The other meetings may be in association with technical conferences or combined with other SciDAC CEMM or CSWPI meetings. The PI/Co-PIs will hold periodic remote collaboration meetings in the beginning weekly and later at least every four weeks depending on project status. A monthly highlight report will be developed along with a budget status summary.

Outreach and Interaction with the Computer Science, Applied Mathematics and Plasma Physics communities: Through a program of presentations at technical conferences, workshops, highlight reports, and a project webpage, especially during the first phase, the

project will involve the community in the development of fusion-specific standards for data types and code interfaces and will establish a technical presence in order to allow adoption of its models and codes. The project will work to advance project goals by encouraging work in related research areas. The participants will directly involve students and post docs in the research. The PI's have a strong history of effectively engaging graduate students in similar projects.

7. Description of facilities, resources, and personnel

Computing resource requirements

The applications being developed by SWIM will require extensive computer resources. The actual resources available will determine the allowed resolution and simulation time, which in turn determines what range of physical parameters can be realistically modeled. The CEMM and Plasma Wave projects together used over 10 million node-hours in FY2004, mostly at NERSC and ORNL. SWIM would expect to use comparable resources in early years and scale up to several times that on a NLCF scale facility in later years. We anticipate that a request for proposals to be issued very soon for projects to use the NLCF, with allocations expected to start October, 1, 2005. Although the results of the present competition may not be decided yet we intend to proceed with a proposal to NLCF for computer time, data storage, and additional specialized performance engineering support.

In addition, we are planning on purchasing a 16 processor SGI Altix 350 which will be located at PPPL and will be used by the project for code development and for fast-turnaround debugging and model testing runs. Our experience indicates that this is a good environment for developing a parallel code suite as a preliminary stage to moving codes to more scalable systems. All members of the SWIM project will have an account and equal access to this machine.

This project will require significant storage on intermediate and archival timescales. We estimate 12 Tbytes for source codes, metadata, data exchanged between components, simulation results, and post-run analysis artifacts. Some of this will be supplied by NERSC and the NLCF, and in particular the parallel I/O files must be staged onto those sites before and after runs. Some storage and cluster resources will also be made available at Indiana University and PPPL.

8. Other current and pending support

ORNL participants are also funded from the Theory program of OFES, the base program of OASCR, and from SicDAC

PPPL participants receive other support from OFES.

SAIC participants receive additional funding from OFES, NASA, and NSF

University of Indiana participants receive additional funding from OASCR

Columbia participants receive additional funding from OASCR

MIT participants receive additional funding from OFES and SciDAC

CompX participants receive additional funding from OFES and SciDAC

General Atomics participants receive additional funding from OFES and SciDAC

University of Wisconsin participants receive additional funding from OFES

New York University participants receive additional funding from OFES and SciDAC

References

- [1] “Fusion Simulation Project: Integrated Simulation and Optimization of Magnetic Fusion Systems,” *J. Fusion Energy* **20**, 135 (2001).
- [2] “Report of the Fusion Simulation Project Steering Committee,” *J. Fusion Energy* **23**(1) (2004).
- [3] L. Eriksson, A. Mueck, O. Sauter et al., “Destabilization of Fast-Ion-Induced Long Sawteeth by Localized Current Drive in the JET Tokamak,” *Phys. Rev. Lett* **92**, 235004 (2004).
- [4] J. P. Graves, K. I. Hopcraft, R. O. Dendy et al., “Sawtooth Evolution during JET Ion-Cyclotron Resonance Heated Pulses,” *Phys. Rev. Lett.* **84**, 1204 (2000).
- [5] R. J. LaHaye, L. L. Lao, E. J. Strait et al., “High beta tokamak operation in DIII-D limited at low density collisionality by resistive tearing modes,” *Nucl. Fusion* **37**, 397 (1997).
- [6] Z. Chang, J. D. Callen et al., “Observation on nonlinear neoclassical pressure-gradient-driven tearing modes in TFTR,” *Phys. Rev. Lett.* **74**, 4663 (1995).
- [7] D.bA. Gates, B. Loyd, A. Morris et al., “Neoclassical islands on COMPASS-D,” *Nucl. Fusion* **37**, 1593 (1997).
- [8] H. Zohm, G. Gantenbein, A. Gude et al., “Neoclassical tearing modes and their stabilization by electron cyclotron current drive in ASDEX Upgrade” *Phys. Plasmas* **8**, 2009 (2001).
- [9] D. P. Brennan et al., “A Mechanism for Tearing Mode Onset Near Ideal Stability Boundaries,” *Phys. Plasmas* **10**, 1643 (2003).
- [10] L. Sugiyama and W. Park, “A nonlinear two-fluid model for toroidal plasmas,” *Phys. Plasmas* **7**, 4644 (2000).
- [11] R. Lahay et al., *Phys. Plasmas* (2002).
- [12] T. H. Stix, *Waves in Plasmas* (American Institute of Physics, New York, 1992).
- [13] M. Brambilla, *Kinetic Theory of Plasma Waves*, (Clarendon Press, Oxford, 1998).
- [14] E. F. Jaeger, L. A. Berry, E. D’Azevedo et al., *Phys. Plasmas* **9**, 1873 (2002).
- [15] E. F. Jaeger, L. A. Berry, J. R. Myra et al., *Phys. Rev. Lett.* **90**, 195001-1 (2003).
- [16] E. F. Jaeger, L. A. Berry, E. D’Azevedo, D. B. Batchelor, and M. D. Carter, *Phys. Plasmas* **8**, 1573 (2001).
- [17] J. C. Wright, P. T. Bonoli, M. Brambilla et al., *Phys. Plasmas* **11**, 2473 (2004).
- [18] M. Brambilla and T. Krucken, *Nucl. Fusion* **28**, 1813 (1988).

- [19] R. J. Dumont, C. K. Phillips and D. N. Smithe, "ICRF wave propagation and absorption in plasmas with non-thermal populations," in *Controlled Fusion and Plasma Physics*, 26B, paper P-5.051 (29th EPS Conference, Montreux, Jun. 17-21, 2002).
- [20] R. J. Dumont, C. K. Phillips, and D. N. Smithe, "Effects of non-Maxwellian Plasma Species on ICRF Propagation and Absorption in Toroidal Magnetic Confinement Devices," in *Radio Frequency Power in Plasmas: 15th Topical Conference*, AIP Conf. Proc. 694 (American Institute of Physics, New York, 2003), p. 439.
- [21] R. J. Dumont, C. K. Phillips, and D. N. Smithe, "Effects of non-Maxwellian Species on Electromagnetic Wave Propagation and Absorption in Magnetically Confined Plasmas," accepted for publication in *Phys. Plasmas* (2005).
- [22] A. P. Smirnov and R. W. Harvey, *Bull. Am. Phys. Soc.* **40**, p. 1837, abstract 8p35, (1995); CompX report CompX-2000-01 (2001).
- [23] C. F. Kennel and F. Engelmann, *Phys. Fluids* **9**, 2377 (1966).
- [24] R. W. Harvey and M. G. McCoy, in *Proceedings of the IAEA Technical Committee Meeting on Advances in Simulation and Modeling of Thermonuclear Plasmas* (IAEA, Montreal, 1992), available through USDOC, NTIS No. DE9300962.
- [25] Variational bounds for transport coefficients in three-dimensional toroidal plasmas W. I. Van Rij and S. P. Hirshman, *Phys. Fluids B* **1**, 563 (1989).
- [26] DELTA5D references may be found at <http://qps.fed.ornl.gov/pvr/document.htm>, Section 6.2.3.
- [27] C. R. Sovinec, A. H. Glasser et al., "Nonlinear magnetohydrodynamics simulation using high-order finite elements," *J. Comput. Phys.* **195**, 355-386 (2004).
- [28] W. Park, E. V. Belova, G. Y. Fu, et al., "Plasma simulation studies using multilevel physics models," *Phys. Plasmas* **6**, 1796 (1999).
- [29] R. Samtaney, P. Colella, S. Jardin, *Computer Physics Communications* (2004).
- [30] H. Grad and J. Hogan, "Classical Diffusion in a Tokamak," *Phys. Rev. Lett.* **24**, 1337 (1970).
- [31] S. P. Hirshman and S. C. Jardin, "Two-dimensional transport of tokamak plasmas," *Phys. Fluids* **22**, 731 (1979).
- [32] S. C. Jardin, N. Pomphrey, J. Delucia, "Dynamic modeling of transport and positional control of tokamaks," *J. Comput. Phys.* **66**, 481-507 (1986).
- [33] D. Ignat, E. Valeo, S. C. Jardin, "Dynamic Modeling of Lower-Hybrid Current Drive," *Nuclear Fusion* **34**, 837 (1994).
- [33] Terascale Simulation Tools and Technology, <http://www.tsst-scidac.org/intro/index.html>.

- [34] Terascale Optimal PDE Simulations, <http://www-unix.mcs.anl.gov/scidac-tops/>.
- [35] An Algorithmic and Software Framework for Applied Partial Differential Equations (APDEC), <http://davis.lbl.gov/APDEC/>
- [36] CVS Project Home, <<https://ccvs.cvshome.org/>><https://ccvs.cvshome.org/>.
- [37] (a) Chris Hill, Cecelia DeLuca, V. Balaji, Max Suarez, Arlindo da Silva, and the ESMF Joint Specification Team, "The Architecture of the Earth System Modeling Framework," *Computing in Science and Engineering*, **6** (1), January/February 2004; (b) G. Toth, O. Volberg, A. J. Ridley, T. I. Gombosi, D. DeZeeuw, K. C. Hansen, D. R. Chesney, Q. F. Stout, K. G. Powell, K. Kane, R. Oehmke, "A Physics-Based Software Framework For Sun-Earth Connection Modeling," *Multiscale Coupling of Sun-Earth Processes* (proceedings of the Conference on the Sun-Earth Connection, Kona, Hawaii, 2004), A. T. Y. Lui, Y. Kamide, and G. Consolini, editors. Elsevier Publ. Co., Amsterdam, The Netherlands, pp. 383-397; (c) Kelvin K. Droegemeier, V. Chandrasekar, Richard Clark, Dennis Gannon, Sara Graves, Everette Joseph, Mohan Ramamurthy, Robert Wilhelmson, Keith Brewster, Ben Domenico, Theresa Leyton, Vernon Morris, Donald Murray, Beth Plale, Rahul Ramachandran, Daniel Reed, John Rushing, Daniel Weber, Anne Wilson, Ming Xue, Sepideh Yalda, "Linked environments for atmospheric discovery (LEAD): Architecture, technology roadmap and deployment strategy," 21th Conf. on Interactive Info. Processing Systems for Meteorology, Oceanography, and Hydrology, San Diego, CA, Amer. Meteor. Soc.
- [38] "Cray X1 Evaluation Status Report," ORNL/TM-2004/13, January 29, 2004.
- [39] <http://www.cca-forum.org/>
- [40] <http://www.llnl.gov/CASC/components/babel.html>
- [41] NWChem Home Page, <http://www.emsl.pnl.gov/docs/nwchem/nwchem.html>
- [42] Ricky A. Kendall, Edo Aprà, David E Bernholdt, Eric J. Bylaska, Michel Dupuis, George I. Fann, Robert J. Harrison, Jailin Ju, Jeffrey A. Nichols, Jarek Nieplocha, T. P. Straatsma, Theresa L. Windus, and Adrian T. Wong, "High Performance Computational Chemistry; Overview of NWChem a Distributed Parallel Application," *Comp. Physics Comm.* **128**, 260 (2000).
- [43] Chris Hill, Cecelia DeLuca, V. Balaji, Max Suarez, Arlindo da Silva, and the ESMF Joint Specification Team, "The Architecture of the Earth System Modeling Framework," *Computing in Science and Engineering*, **6**(1), January/February 2004.
- [44] Joint ORNL/Indiana University Workshop on Computational Frameworks for Fusion, May 3-5 2005 at Oak Ridge National Laboratory.
- [45] Rob Armstrong, Dennis Gannon, Al Geist, Katarzyna Keahey, Scott Kohn, Lois McInnes, Steve Parker, and Brent Smolinski, "Toward a Common Component Architecture for High-Performance Scientific Computing," *Proceedings of the 1999 Conference on High Performance Distributed Computing*.

- [46] David E. Bernholdt, Benjamin A. Allan, Robert Armstrong, Felipe Bertrand, Kenneth Chiu, Tamara L. Dahlgren, Kostadin Damevski, Wael R. Elwasif, Thomas G. W. Epperly, Madhusudhan Govindaraju, Daniel S. Katz, James A. Kohl, Manoj Krishnan, Gary Kumfert, J. Walter Larson, Sophia Lefantzi, Michael J. Lewis, Allen D. Malony, Lois C. McInnes, Jarek Nieplocha, Boyana Norris, Steven G. Parker, Jaideep Ray, Sameer Shende, Theresa L. Windus, and Shujia Zhou, “A Component Architecture for High-Performance Scientific Computing,” Intl. J. High-Perf. Computing Appl. (2005), submitted to ACTS Collection special issue.
- [47] Dahlgren, Tamara and Epperly, Thomas and Kumfert, Gary, Babel User’s Guide, CASC, Lawrence Livermore National Laboratory, Version 0.8, January 2003, Livermore, CA.
- [48] The NetCDF Users’ Guide, <http://my.unidata.ucar.edu/content/software/netcdf/docs/netcdf/index.html>.
- [49] The netCDF Markup Language (NcML, <http://www.vets.ucar.edu/luca/netcdf/>).
- [50] A User’s Guide for HDF5 Release 1.4.5, <http://hdf.ncsa.uiuc.edu/HDF5/doc/H5.user.html>.
- [51] The HDF5 XML Information Page, <http://hdf.ncsa.uiuc.edu/HDF5/XML/>.
- [52] Data Format Description Language Working Group, <https://forge.gridforum.org/projects/dfdl-wg/>.
- [53] FSML: Fusion Simulation Mark-up Language for Interoperability of Data and Analysis Tools, http://www.txcorp.com/projects/index.php#projects_FSML.
- [54] S. I. Braginskii in Reviews of Plasma Physics, Vol.1, M. A. Leontovich ed. (Consultants Bureau, New York, 1965) p. 205.
- [55] A. B. Mikhailovskii and V. S. Tsypin, Plasma Phys. **13**, 785 (1971).
- [56] P. J. Catto and A. N. Simakov, Phys. Plasmas **11**, 90 (2004).
- [57] E. D. Held, J. D. Callen, C. C. Hegna, C. R. Sovinec, T. A. Gianakon and S. E. Kruger, Phys. Plasmas **11**, 2419 (2004).
- [58] J. J. Ramos, Phys. Plasmas **10**, 3601 (2003).
- [59] J. J. Ramos, “Fluid formalism for collisionless magnetized plasmas,” to appear in Phys. Fluids (2005).
- [60] C. C. Hegna, Phys. Plasmas **5**, 1767 (1998) and references therein.
- [61] CVS Project Home, <https://ccvs.cvshome.org/>.
- [62] Bitkeeper – The Scalable Distributed Software Configuration Management System, <http://www.bitkeeper.com>.

- [63] Aegis 4.20, <http://aegis.sourceforge.net/>.
- [64] An Introduction to JAM, <http://freetype.sourceforge.net/jam/index.html>.
- [65] CMZ Online Documentation, <http://wwwcmz.web.cern.ch/wwwcmz>.
- [66] SCONS: A Software Construction Tool, <http://www.scons.org/>.
- [67] DejaGnu – GNU Project – Free Software Foundation (FSF):
<http://www.gnu.org/software/dejagnu/>.
- [68] Junit, Testing Resources for Extreme Programming,
<http://www.junit.org/index.htm>.
- [69] The Expect Home Page, <http://expect.nist.gov/>.
- [70] High Level System Testing Environment, <http://atomicobject.com/haste/>.
- [71] Home::Bugzilla::bugzilla.org, <http://www.bugzilla.org/>.
- [72] GNATS: The Gnu Bugtracking System, <http://www.gnu.org/software/gnats/>.
- [73] D. Gannon, G. Fox, M. Pierce, B. Plale, G. von Laszewski, C. Severance, J. Hardin, J. Alameda, M. Thomas, J. Boisseau, “Grid Portals: A Scientist’s Access Point for Grid Services, Sept. 19 2003, GGF working document. Also see <http://www.collab-ogce.org/nmi/index.jsp>.
- [74] M. Thomas, S. Mock, J. Boisseau, M. Dahan, K. Mueller, D. Sutton, “The GridPort Toolkit Architecture for Building Grid Portals,” Proceedings of the 10th IEEE Intl. Symp. On High Perf. Dist. Comp. Aug 2001.
- [75] D. P. Schissel et al., “Building the US National Fusion Grid: results from the National Fusion Collaboratory Project,” Fusion Engineering and Design **71** 245-250 (2004).
- [76] S. Flanagan et al., “A General Purpose Data Analysis Monitoring System with Case Studies from the National Fusion Grid and the DIII-D MDSplus between Pulse Analysis System,” Fusion Engineering and Design **71**, 263 (2004).
- [77] TAU Portable Profiling,
<http://www.cs.uoregon.edu/research/paracomp/tau/tautools/>.
- [78] Karen Devine and Erik Boman and Robert Heaphy and Bruce Hendrickson and Courtenay Vaughan, “Zoltan Data Management Services for Parallel Dynamic Applications,” Computing in Science and Engineering, **4**(2), pp. 90-97 (2002).
- [79] S. Shende, A. D. Malony, C. Rasmussen, M. Sottile, “A Performance Interface for Component-Based Applications,” Proc. International Workshop on Performance Modeling, Evaluation, and Optimization of Parallel and Distributed Systems, IPDPS’03, IEEE Computer Society, 278, 2003.

- [80] M. Grote and T. Huckle, "Effective Parallel Preconditioning with Sparse Approximate Inverses," Proc. Seventh SIAM Conf. on Parallel Process. for Sci. Comp., February 1995.
- [81] F. Porcelli, D. Boucher, M. Rosenbluth, Plasma Phys. and Cont. Fusion **38**, 2163 (1996).

Biographical Sketches:

Oak Ridge National Laboratory

Donald B. Batchelor is the Plasma Theory Group Leader at Oak Ridge National Laboratory. He has been at ORNL since receiving his Ph.D. in Physics from the University of Maryland in 1976, and has more than 70 refereed publications. After joining ORNL, he began working on the Elmo Bumpy Torus and developed one of the first geometrical optics codes for a relativistic plasma in the vicinity of cyclotron resonances. In 1984, he formed a group to study a broad range of RF problems including ion cyclotron heating in complex geometries. He was instrumental in initiating the development of the ORION, PICES, and RANT3D RF codes. These models have been central to the design and analysis of RF heating experiments including TFTR at Princeton, DIII-D at General Atomics, and JET, as well as for ITER. Don is the principal investigator on the DOE SciDAC project on wave-plasma interactions, which successfully completed an initial three-year award, and is now continuing on a three-year renewal. He has been active in advancing the role of computer simulation in fusion through his participation in both FESAC and DOE advisory panels. His publication in *Physics Today* of an article on Fusion Simulation is indicative of his impact in the area. Other research interests include the development of innovative plasma confinement systems. He led the development of a Small-Aspect Ratio Stellarator-Tokamak Hybrid (SMARTH) that has since developed into the Quasi-Poloidal Stellarator (QPS) experiment at ORNL, an element of the U.S. Compact stellarator program. Dr. Batchelor has organized several international conferences. He is a fellow of the American Physical Society.

Lee A. Berry is a Senior Research Staff Member in the Fusion Energy Division's Theory Group at Oak Ridge National Laboratory. His research experience covers a wide range of experimental and theoretical plasma physics topics including tokamak, bumpy torus, and stellarator experiments and theory. His current plasma research includes stellarator experiments (Deputy Project Manager for the Quasi-Poloidal Stellarator at ORNL), and RF-Plasma theory and computational modeling through the SciDAC RF-Plasma Interactions project. Contributions have been in the areas of neutral beam heating, bumpy torus experiments and design, stellarator design and analysis, fusion technology, semiconductor process tool experiments and modeling, RF flow drive, and algorithm development for the AORSA full-wave code. He received his Ph.D. from the University of California at Riverside in 1970 and has been at ORNL since then. He is a fellow of the American Physical Society, and was awarded the DOE Distinguished Service Award for his leadership on the ORNL bumpy torus program. He has served on a variety of Division of Plasma Physics Committees and on the Office of Science's Fusion Energy Science Advisory Panel, as well as subpanels under FESAC. He is currently Chair of the Coalition for Plasma Physics, and a member of the UCLA Plasma Science and Technology Institute Ad Hoc Committee.

Stephen. P. Hirshman received his Ph.D. in 1976 from the Massachusetts Institute of Technology, and is presently a Research Staff Scientist at Oak Ridge National Laboratory. He was a post doctoral research associate at Princeton Plasma Physics Laboratory from 1976 to 1978 and began work at ORNL in 1978. He combines a unique set of mathematical, computational, and plasma physics skills and has made seminal contributions to a number of plasma physics issues including formulation of neoclassical transport theory based on moment equation expansions, computation of transport coefficients for 3D systems (the DKES code), magnetic equilibration in 3D systems (the VMEC code) and stellarator design (the Stellopt code). The Stellopt code was a unique development that allows design of 3D fusion confinement systems, specifically stellarators. The National Compact Stellarator Experiment (NCSX) at PPPL and the Quasi-Poloidal Experiment at ORNL could not have been designed without Steve's contributions. He is a Fellow of the American Physical Society.

David E. Bernholdt is a Senior R & D Staff Member in Computer Science at ORNL. Bernholdt's formal training is in chemistry, with a B.S. in 1986 from the University of Illinois and Ph.D. in 1993 from the University of Florida. He was a postdoctoral fellow at Pacific Northwest National Laboratory 1993–1995. From 1995–2000, he was at Syracuse University, as an Alex Nason Fellow and a Senior Research Scientist in the Northeast Parallel Architectures Center, and a Research Assistant Professor in the Chemistry Department. Bernholdt's primary research focus is on means to improve the performance and productivity of HPC software developers and their software. Some of his activities relevant to this proposal include: co-PI of the Common Component Architecture effort and its lead for User Outreach and Applications Integration, co-PI of the "Comprehensive Fusion Simulation: Component-Based Software Engineering and Evolutionary Time Advancement" project, ORNL lead for the Earth System Grid project, and driving the deployment of Access Grid and related collaborative technologies at ORNL. He also has extensive experience with large-scale scientific software systems, particularly as one of the original designers and developers of the NWChem parallel computational chemistry package.

Eduardo F. D'Azevedo is group leader for the Computational Mathematics Group at the Computer Science and Mathematics Division, Oak Ridge National Laboratory. He is co-author of one book and 19 refereed publications. His current research includes: optimal mesh generation, vectorized iterative solver, out-of-core dense linear solvers, and application specific preconditioners. He received his Ph.D. in 1989 in the Faculty of Mathematics (Department of Computer Science) from the University of Waterloo, Ontario, Canada. He held an ORISE postdoctoral fellowship from 1990–1991, and has been a research staff member at ORNL since 1991. He is a member of the Society for Industrial and Applied Mathematics.

Mark R. Fahey is presently Senior Research Staff in the Scientific Application Support Group within the Center for Computational Sciences (CCS) at ORNL. He is the primary CCS liaison for the fusion researchers funded by the DOE. Mark has ported and optimized several fusion codes to a variety of parallel platforms as part of his CCS work. Mark is the Applications SIG chair for the Cray User Group. He was from 2001–2003 a Research Scientist as part of the Joint Institute for Computational Science at ORNL.

Prior to joining ORNL, Dr. Fahey was a Research Scientist at the Engineering Research and Development Center in Vicksburg, MS, from 1998 to 2001. There he was part of the Computational Migration Group (CMG) and later became the Director of the CMG. Dr. Fahey received his B.A. from St. Norbert College in 1992 and his Ph.D. from the University of Kentucky in 1999, where he was received the Trustee's Distinguished Scholarship (4 years) and the University of Kentucky Center for Computational Sciences fellowships (3 years), respectively. His research in scientific computing and numerical linear algebra ranged topics covering sparse and dense matrix computations, eigenvalue problems, iterative solvers, probability theory, bilinear forms, and numerical integration techniques with demonstrated applicability to problems in chemistry and physics. He also helped write routines that were included in the 3.0 release of LAPACK.

Wael Elwasif obtained his Ph.D. in Computer Science from the University of Tennessee, Knoxville, in 2004. He joined the computer science and mathematics division, ORNL in 2000, where he conducted research in heterogenous distributed computing and component models for high-performance computing. Dr. Elwasif's research interests include parallel and distributed computing, component models for large-scale High-Performance Computing (HPC), HPC language interoperability, and large-scale parallel programming models.

Wayne A. Houlberg is a Senior Research Staff member in the Fusion Energy Division of Oak Ridge National Laboratory. He has conducted research on many aspects of transport in tokamaks and stellarators, including the development of: methods for coupling MHD equilibria with time-dependent transport codes; models and interfaces for many types of fueling, heating, radiation, and non-inductive current sources; models for neoclassical transport; and models for bifurcations of transport properties associated with breaking intrinsic ambipolarity. Since 2001, he has served as Chair of the International Tokamak Physics Assessment Topical Group on Confinement Databases and Modeling. He has been a visiting scientist at the Joint European Torus (JET) project in England (1984, 1987), and at the JT-60U project at JAERI-Naka (2000), and actively collaborated on the analysis of many other national and international tokamak and stellarator experiments. He received his Ph.D. in Nuclear Engineering from the University of Wisconsin-Madison in 1977.

E. Fred Jaeger is a senior research staff member in the Plasma Theory Section of the Fusion Energy Division, Oak Ridge National Laboratory (ORNL). He has co-authored over 61 refereed journal articles, seven invited papers, and 63 papers published in major conference proceedings, and has served on several Department of Energy review panels. His research has covered a broad range of topics in plasma physics theory and computation with applications to fusion plasmas, plasma processing of materials, gas lasers, and space plasmas. Most recently, he has developed tera-scale parallel simulations of radio frequency interactions in tokamak and stellarator plasmas. These simulations, developed under SciDAC, have become valuable tools for understanding electromagnetic wave propagation and heating in 2-D and 3-D plasmas. He received his Ph.D. in Nuclear Engineering from the University of California in 1970, and held postdoctoral fellowships at the National Center for Atmospheric Research, Boulder from

1971-72, and at the Joint Institute for Laboratory Astrophysics, Boulder, from 1973-74. He has been at ORNL since 1974.

Donald A. Spong is a senior research staff member and group leader for stellarator plasma theory in the Fusion Energy Division of Oak Ridge National Laboratory. He is an author on about 100 publications. His current plasma research interests include particle-based closure relations for MHD problems, stellarator neoclassical transport using Monte Carlo and fluid moments methods, fast particle-destabilized Alfvén modes in stellarators, and applications of scientific visualization to plasma physics problems. Previously he has worked in the areas of stellarator optimization, gyrofluid and neoclassical closures for resistive MHD and fast particle instability problems, alpha particle physics, kinetic ballooning modes, plasma processing of semiconductors, hot electron instabilities, and runaway electrons in tokamaks. He has served on the Program Advisory Committee for the National Energy Research Supercomputing Center, the ITER Expert Group on Energetic Particles, and as a reviewer of Plasma Theory proposals for the Office of Fusion Energy, DOE. He received his Ph.D. in Nuclear Engineering (plasma physics and controlled fusion) from the University of Michigan, Ann Arbor, in 1976. He has received the 2001 ORNL Research Team Award for development of the Quasi-Poloidal Stellarator, the 1994 Martin Marietta Author of the Year Publication Award, and the 1988 Martin Marietta publications Award.

Indiana University

Randall Bramley received a Master's degree in computational mathematics and then a Ph.D. in computer science at the University of Illinois-Urbana/Champaign in 1989. After three years as research scientist at the Center for Supercomputing Research and Development at UIUC, he moved to Indiana University where he is currently an associate professor of Computer Science, adjunct professor in Informatics, creator and director of the multidisciplinary Scientific Computing Program, and a senior research scientist in the Pervasive Technologies Labs at Indiana University. He currently has collaborations in x-ray diffractometry, astronomical photometry, geophysics, molecular modeling, bioinformatics, and biophysics. His research interests have ranged from numerical analysis to systems development, particularly in distributed scientific computing, direct access to scientific instruments through web services, and large-scale data management. He has co-edited one book on problem solving environments, has over 40 refereed publications, and is a founding member of the DOE Common Component Architecture Forum.

University of California at Berkeley

Dr. Phillip Colella received his B.A. (1974), M.A. (1976) and Ph.D. (1979) from the University of California at Berkeley, all in applied mathematics. Dr. Colella has been a staff scientist at the Lawrence Berkeley National Laboratory and at the Lawrence Livermore National Laboratory, and from 1989 to 1995, was a Professor in the Mechanical Engineering Department at the University of California at Berkeley. He is

currently a Senior Staff Scientist and Group Leader for the Applied Numerical Algorithms Group in the Computing Sciences Directorate at the Lawrence Berkeley National Laboratory. His research has been in the area of numerical methods for partial differential equations, with contributions in high-resolution finite-difference methods, adaptive mesh refinement, volume-of-fluid methods for irregular boundaries, and programming language and library design for parallel scientific computing. He has also applied numerical methods in a variety of scientific and engineering fields, including shock dynamics, low-Mach number and incompressible flows, combustion, porous media flows, and astrophysical flows. Honors and awards include the IEEE Sidney Fernbach Award for high-performance computing in 1998, the SIAM/ACM prize (with John Bell) for computational science and engineering in 2003, and election to the U.S. National Academy of Sciences in 2004.

CompX

Robert (R.W.) Harvey is an expert in fusion energy computations, specializing in in-depth development and application of computational models for interpretation of auxiliary heating and transport in fusion energy plasmas. Dr. Harvey conceived and led development of several major computer codes used for physics interpretation of fusion energy experiments, while working at General Atomics, San Diego for twenty years, and then at CompX, Del Mar, California, which he founded in 1995. For the past twenty years, he has focused primarily on the development, application, and dissemination of the CQL3D collisional/rf-quasilinear Fokker-Planck code for simulation of microwave and neutral beam heating in tokamak fusion energy plasmas. The code has led to many collaborations, nationally and internationally, for interpretation of experiments, playing a direct role in more than one hundred research publications. Dr. Harvey has contributed to design of the International Tokamak Experimental Reactor, primarily in regard to requirements for electron cyclotron current drive systems to stabilize neoclassical tearing modes. He has also worked as a leader of a GA/Russian Theory Collaboration Program (1989-1995), an Associate Editor of Physics of Fluids (1986-1988), and a lecturer at the University of California, San Diego. Dr. Harvey is an author or co-author on 184 papers. He obtained his Ph.D. in Applied Physics at the University of California, San Diego, 1973.

Princeton Plasma Physics Laboratory

Stephen C. Jardin is a Principal Research Physicist at the Princeton University Plasma Physics Laboratory. He is presently Co-Head of the Computational Plasma Physics Group at PPPL, and MHD Physics leader for the Theory Department. He has been Lecturer with Rank of Professor in the Princeton University Astrophysics Department since 1986. He holds a BS in Engineering Physics from the University of California, a MS (Physics) and MS (Nuclear Engineering) from MIT, and a Ph.D. in Astrophysics from Princeton University (1976). He was the primary developer of several widely used MHD equilibrium, stability, and transport codes including the Tokamak Simulation Code (TSC). He is currently the lead Principle Investigator for the SciDAC Center for Extended Magnetohydrodynamic Modeling, which was just renewed for another 3 years

in the summer of 2004. He holds four U.S. patents, has had over 150 refereed publications in plasma physics, and has supervised six Princeton University Ph.D. students. He is a member of Phi Beta Kappa and a Fellow of the American Physical Society. He has held key positions in several fusion device design teams including those for S-1, PBX-M, CIT, BPX, and TPX. He was the ARIES Physics leader from 1992-2000, and is presently a U.S. representative on the International Tokamak Physics Activity group on MHD, Plasma Control, and Disruptions. He is a member of the NERSC Executive Committee Users Group, is Chair of the NERSC Program Advisory Committee, and was chair of the National Transport Code Collaboration Program Advisory Committee. He has served on recent FESAC subcommittees on the Integrated Simulation and Modeling of Fusion Systems and to review the Inertial Fusion Energy program.

A software engineer and applied mathematician by training, **Douglas McCune** has over 25 years of experience in computational software development at the Princeton Plasma Physics Laboratory. He is the main author and creator of TRANSP, an integrated software package for tokamak fusion plasma simulation, which has been used to analyze and validate experimental results on tokamaks around the world. The TRANSP code pioneered the use of detailed Monte Carlo fast-ion simulation for accurate prediction of the energy content and fusion reactivity of MHD-quiescent tokamak core plasmas, given thermal plasma temperature and density measurements. In recent years, Douglas McCune has co-lead the Computational Plasma Physics Group at PPPL. He has played a leading role in numerous successful projects involving collaborative sharing and integration of software, such as the National Transport Code Collaboration Modules Library project, <http://w3.pppl.gov/NTCC>, and the SciDAC Collaboratory's Fusion Grid project, <http://www.fusiongrid.org>, by which a production TRANSP facility is made available to both PPPL and remote users as a computational service, producing over 2500 runs in 2004. Douglas McCune was named a PPPL Distinguished Engineering Fellow in 2001.

Guo-yong Fu is a Principal Research Physicist at the Princeton University Plasma Physics Laboratory. His main research area includes energetic particle physics and MHD stability of toroidal plasmas. He has authored over 60 refereed publications in plasma physics. He received a Ph.D. in physics from The University of Texas at Austin in 1988. After a one-year postdoctoral fellowship at the Institute of Fusion Studies in Austin, Texas, he worked on MHD stability in stellarators at the center for plasma physics research (CRPP) in Lausanne, Switzerland, from 1989 to 1991. He joined PPPL in 1992. He is currently active in the area of MHD stability in compact stellarators and nonlinear dynamics of energetic particle-driven fishbone and TAE. He won the Kaul Foundation Prize for excellence in plasma physics research in 1998.

Wonchull Park is a Principal Research Physicist at the Princeton University Plasma Physics Laboratory. He has a Ph.D. from Columbia University, and is a Fellow of American Physical Society. His primary field of research is 3D numerical simulation studies of plasmas, using multilevel physics models, MHD, Two-fluids, and various Particle/Fluid hybrid models. He has authored or coauthored over 100 refereed publications in plasma physics

Columbia University

David E. Keyes is the Fu Foundation Professor of Applied Mathematics at Columbia University and the Acting Director of the Institute for Scientific Computing Research at LLNL. He earned a B.S.E. in Mechanical Engineering from Princeton in 1978 and a Ph.D. in Applied Mathematics from Harvard in 1984. He post-doc'ed in Computer Science at Yale and taught there for eight years prior to joining Old Dominion University and the Institute for Computer Applications in Science & Engineering at the NASA Langley Research Center in 1993. He moved to Columbia in 2003. Keyes is author or co-author of over 100 publications in computational science and engineering, numerical analysis, and computer science. He works at the algorithmic interface between parallel computing and the numerical analysis of partial differential equations, across a spectrum of aerodynamic, geophysical, and chemically reacting flows. Newton-Krylov-Schwarz parallel implicit methods, introduced in a paper he co-authored in 1993, are now widely used throughout engineering and computational physics, and have been scaled to thousands of processors. Keyes has been recognized with a Gordon Bell Prize for High-Performance Computing and a National Science Foundation Presidential Young Investigator Award. He is currently co-editor of Int J High Performance Computing Applications and SIAM's Computational Science & Engineering book series, and serves on the SIAM Council.

University of Wisconsin

Chris C. Hegna is an Associate Professor of Engineering Physics at the University of Wisconsin. His primary field of research is theoretical plasma physics with an emphasis on the area of plasma confinement using magnetic fields. He is the acting director of the University's Center for Plasma Theory and Computation. He is the author or co-author of over 70 refereed publications. Hegna's current research includes; the role of macroscopic instabilities in high-temperature plasmas; nonideal and nonlinear magnetohydrodynamic instabilities; kinetic theory modifications to fluid-like descriptions of plasmas; plasma dynamics in non-symmetric magnetic systems; the role of magnetic geometry, symmetry and topology on plasma instabilities, turbulence and transport properties; self-organization properties of plasmas; anomalous current and momentum transport; magnetic reconnection in laboratory and astrophysical plasmas and dynamo processes. He received his Ph.D. in Applied Physics from Columbia University in 1989, and has been employed at the National Institute for Fusion Science (Nagoya, Japan), UKAEA Technologies, Culham Laboratory (Abingdon, England) and Columbia University, working at the Princeton Plasma Physics Laboratory. He is a fellow of the American Physical Society.

Massachusetts Institute of Technology

Jesus Ramos is a Principal Research Scientist at the Plasma Science and Fusion Center of the Massachusetts Institute of Technology. He has a Ph.D. in Physics from the "Universidad Complutense" of Madrid, Spain. He has made extensive contributions to the magnetohydrodynamic theory of fusion plasmas and has participated in the design teams of Alcator C-Mod, CIT, ARIES, TPX, and ITER. He has authored or co-authored 56 refereed publications.

Linda E. Sugiyama is a principal research scientist in the Research Laboratory of Electronics at the Massachusetts Institute of Technology. She has worked in many areas of the theory and numerical simulation of magnetically confined plasmas, including equilibrium, stability, transport, auxiliary heating, and nonlinear simulation. She first proposed and, with W. Park, developed the two-fluid MH3D-T code for axisymmetric confined plasmas, based on the existing MH3D MHD nonlinear code developed at PPPL. These codes formed the basis of the M3D project. She received a B.S. from the University of Wisconsin in 1975, and a Ph.D in Applied Mathematics from the Massachusetts Institute of Technology in 1980. Since then, she has worked at MIT, with brief sojourns at other institutions.

SAIC

Dalton Schnack is a computational physicist with over 30 years experience in the analytic and numerical solution of nonlinear, multidimensional problems in hydrodynamics and magnetohydrodynamics (MHD). He has authored or co-authored many papers in the fields of linear and nonlinear resistive MHD, and computational methods related to such problems. He has extensive experience in the supercomputing environment. He is actively involved in studying the nonlinear MHD properties magnetic fusion experiments and the solar corona, and in the highly nonlinear (turbulent) properties of the Navier-Stokes and MHD equations. After receiving his B.S., Dr. Schnack worked for over seven years (1967-1973) for the Pratt and Whitney Division of United Technologies Corporation as a Senior Scientific Programmer/Analyst. There he worked on computational problems of steady flow in nozzles and supersonic exhaust jets, and the performance of axisymmetric compressors. During this time, he completed work on an M.S. in physics, which included a thesis on ultra short optical pulse propagation. Dr. Schnack did his doctoral research at Lawrence Livermore Laboratory under Prof. John Killeen where he began his interest in nonlinear MHD processes in fusion plasmas. After graduation, he served as a staff physicist in the computational physics group at the National Magnetic Fusion Energy (MFE) Computer Center, a supercomputer network funded by the Department of Energy. In 1980, he joined the fusion theory group at Los Alamos National Laboratory where he worked on problems relevant to the reversed-field pinch and compact torus experiments. In July 1982, Dr. Schnack joined the Applied Plasma Physics and Technology Division of SAIC, and in 1996 was appointed Director of the Center for Energy and Space Science. He is presently Principal Investigator for two grants with the U. S. Department of Energy. He is actively involved in research

related to the nonlinear fluid dynamics of advanced magnetic fusion devices, and the nonlinear properties of the MHD equations. Dr. Schnack is a member of Phi Kappa Phi, national scholastic honor society, and is a Fellow of the American Physical Society. He is also a member of the American Geophysical Union and the Solar Physics Division of the American Astronomical Society. He has been an active participant in the functions of the international fusion program for many years. Dr. Schnack has co-authored over 70 refereed publications, and one book.

New York University

Hank Strauss is a Research Professor in the Magnetofluid Dynamics Division of the Courant Institute of Mathematical Sciences, New York University. He has a Ph.D. from the University of Texas. He is a fellow of the American Physical Society. He has worked in many areas of theoretical and computational plasma physics, particularly in magnetohydrodynamics. For the last several years, he has been highly involved in the M3D project.

General Atomics

David P. Schissel is the Manager of the Data Analysis Applications Group at the DIII-D National Fusion Facility and the Director of the Advanced Imagery Laboratory at General Atomics. He has over 20 years experience in fusion plasma physics and management with over 75 fusion-related publications from plasma confinement research to advanced computer science research. For the DIII-D National Fusion Facility, he is responsible for coordinating computer hardware and software resources to support the scientific staff's data analysis requirements. As the principal investigator of the USDOE SciDAC funded National Fusion Collaboratory Project, he is responsible for coordinating advanced computer science research among seven distinct institutions to deliver a persistent collaboration infrastructure to transform fusion research. As principal investigator for basic confinement studies on the DIII-D tokamak, he has been responsible for creating and implementing experimental proposals to study the global energy confinement and local energy transport properties of neutral beam heated fusion plasmas. He has been a visiting scientist at the JET Joint Undertaking Tokamak (Culham Laboratory) in the United Kingdom and at Tore Supra Tokamak (Cadarache Laboratory) in France. He received his B.S. degree in Nuclear Engineering from the University of Wisconsin Madison in 1979, his M.S. degree in plasma physics from the Massachusetts Institute of Technology in 1982, and is a fellow of the American Physical Society.

Appendix A: Description of Principal Computer Codes

Extended MHD

NIMROD The NIMROD code solves the nonlinear time-dependent extended-MHD equations with numerical methods that excel for extremely stiff and anisotropic systems. To address the multiple timescales of nearly dissipation-free high-temperature plasmas, fluid moments and electromagnetic fields are advanced in time with a semi-implicit approach. [27]. The complete linear ideal-MHD energy integral is used as the semi-implicit operator to preclude artificial coupling of normal modes at large time-step, and integration is based on the symplectic leap-frog scheme to avoid numerical dissipation. The spatial domain is represented with two-dimensional finite elements for the arbitrarily shaped poloidal plane and finite Fourier series for the toroidal direction. Lagrange polynomials of arbitrary degree serve as the basis functions for the finite elements, and production simulations are typically run with bicubic or biquartic elements (fourth- or fifth-order accurate, respectively) to resolve the extreme anisotropies in magnetized plasmas. For parallel computations, the three-dimensional domain is divided into blocks of finite elements and layers of Fourier components. All communication (both point-to-point and collective) is accomplished with routines in the MPI library [<http://www.mpi-forum.org>]. The non-local effects of free-streaming plasma particles are modeled by accumulating kinetic contributions to fluid-closure terms along characteristics of the drift-kinetic equation, which is solved in the basis of the pitch-angle scattering operator. For modeling fast-particle effects on macroscopic modes, a simulation-particle-based approach has been adapted from M3D for NIMROD's high-order finite elements. The NIMROD code runs on a variety of hardware platforms from laptops with Linux or the Mac OSX operating system to massively parallel architectures with distributed memory. The source code for NIMROD has been publicly available for more than five years, and information for its use and modification is provided on the website <http://nimrodteam.org>.

M3D The M3D code [28] (or multilevel 3D) is a massively parallel nonlinear 3D extended MHD code that makes no assumptions regarding the axisymmetry of the boundaries so that it is equally applicable to stellarators and to tokamaks. 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 includes resistive MHD, two-fluid, hybrid, and fully kinetic particles. M3D uses a stream function/potential representation for the magnetic vector potential and velocity that has been designed to minimize spectral pollution and lead to well-conditioned sparse matrix inversions. Parallel thermal conduction is simulated with the "artificial sound" method. [84] The solution algorithm is partially implicit in that only the most time-step limiting terms including the compressional Alfvén wave and field diffusion terms are implemented implicitly. The three dimensional mesh in M3D facilitates the resolution of multi-scale spatial structures, such as reconnection layers and the representation of fully three-dimensional boundaries that occur in evolving the free boundary of a tokamak or in a stellarator. The mesh uses unstructured triangular finite elements in the poloidal section that can be of arbitrary order. A fully 3D parallel domain decomposition is used with a bias to minimizing the inter-processor communication during the elliptic solves. The PETSc library is used to provide high-performance portable sparse-matrix solvers.

AMRMHD The Princeton/LBL AMRMHD code [29] is an adaptive mesh refinement 3D MHD code built within the Chambo framework. The hyperbolic terms are evaluated using an un-split second-order Godunov method based on the symmetrized 8-wave formulation of the equations. All variables are cell-centered with the solenoidal property of the magnetic field enforced by the application of a Hodge projection on the face-centered magnetic fields used in the flux calculation. The parabolic terms are treated with a semi-implicit method that leads to variable coefficient Helmholtz equations which are solved using a multi-grid method. The method is designed to faithfully follow the multiple characteristics of the problem, yielding a faithful representation of highly anisotropic multiple-scalelength phenomena. Computational speed-up due to AMR depends on the problem, but are more than 100 for some applications. A pellet mass source has been added that allows us to follow the pellet ablation process as it is injected into the tokamak.

CEMM Simulation Codes

	NIMROD	M3D	AMRMHD*
Poloidal discretization	High order quadrilateral finite elements	Triangular linear finite elements	Structured adaptive grid
Toroidal discretization	Pseudospectral	Finite difference	Structured adaptive grid
Time integration	Semi-implicit	Partially implicit	Partially implicit and time adaptive
Enforcement of $\nabla \cdot \mathbf{B} = 0$	Error Diffusion	Vector Potential	Projection Method
Libraries	SuperLU (LBL)	PETSc (ANL)	CHOMBO (LBL)
Sparse Matrix Solver	Direct and Conjugate Gradient	GMRES and ICCG	Conjugate Gradient
Pre-conditioner	Direct solve of approximate matrices	Incomplete LU	Multigrid

*Exploratory project

RF Solvers

GENRAY GENRAY [22] is a general ray tracing code for the calculation of electromagnetic wave propagation and absorption in the geometrical optics approximation. It provides a solution of the ray tracing equations in general non-axisymmetric geometry, although work to date is with axisymmetric equilibria with added toroidal perturbations. Several alternative dispersion functions, D , are provided in order to ray trace for Electron Cyclotron (EC), Lower Hybrid (LH), and Ion Cyclotron Range of Frequencies (ICRF) waves. Current drive is calculated based on Maxwellian distribution functions. Results are coupled to the CQL3D code to provide input for calculation of the RF quasi-linear diffusion coefficients. The code requires solving six ordinary differential equations and will not pose a significant load on computer resources.

AORSA The ORNL All Orders Spectral Algorithm (AORSA) code [14] set solves the full set Maxwell's equations for the time varying RF electric field using a plasma conductivity kernel that is valid for all cyclotron harmonics and for all values of the perpendicular wave vector times the cyclotron radius. Fourier modes in cylindrical coordinates, (R, Z, ϕ) , the major radius, vertical height, and toroidal angle, respectively, are used as basis functions for solutions in 1D, 2D, or 3D. The linear equations for the electric field are generated by collocation, and, because of the integral nature of the plasma response, are dense. For Maxwellian velocity distributions, the plasma conductivity can be expressed in terms of analytic functions, while for non-Maxwellians, a 2D numerical integral must be performed at each collocation point. The parallel ScaLAPACK library is used to factor this matrix. Depending on the dimensionality and resolution, there are from 10^3 to $<10^6$ complex unknowns. The electric field solution is then used to calculate the 4D (for 2D spatial resolution) RF quasilinear operators. An average along the magnetic field lines then calculates the 3D form required as input to the CQL3D code which then evolves the distribution function.

TORIC The perpendicular plasma conductivity in the 2D (axisymmetric tokamak geometry) TORIC [17,18] kernel is a Taylor expansion in the perpendicular wavelength and includes terms through second order in the (assumed) small parameter of perpendicular wavelength times the ion cyclotron radius. Flux coordinates are used to express the radial dependence in terms of a single variable with a finite-element basis set. Fourier modes are used to expand the dependence on poloidal angle. This poloidal dependence allows an algebraic expression for the plasma conductivity parallel to the equilibrium magnetic field. This formulation of the plasma conductivity describes interactions through the second cyclotron harmonic. ICRH at the first and second harmonic are typical applications, with lower hybrid applications possible at for very high resolution. The resulting linear equations are block tridiagonal. For typical resolutions, both the block size and number of blocks are from a few hundred to several thousand. This matrix is factored using a parallel (or serial for small problems) algorithm that saves the LU factorization of each block to disk, and then generates a solution from this decomposition using a variation of the Thomas algorithm.

RF Simulation Codes

	AORSA	TORIC	GENRAY
discretization	Fourier modes in 1-3D	Finite elements in radius, Fourier modes in poloidal angle (2D)	ODEs in 3D
parallelization	MPI	MPI	MPI
Libraries	LAPACK, ScaLAPACK, blas, blac, fftpack	LAPACK, ScaLAPACK, blas, blac	LAPACK, netCDF, PGPLOT, blas
Matrix solver	ScaLAPACK	ScaLAPACK	Not applicable

Fokker-Planck Codes

CQL3D CQL3D is a multi-species, fully relativistic, bounce-averaged, time-dependent Fokker-Planck equation solver [24] that includes both Coulomb and quasilinear RF diffusion. It is 2D in momentum space with a radial coordinate for noncircular plasmas. The code is run in combination with lower hybrid, fast wave, electron cyclotron and electron Bernstein wave ray tracing or full-wave (AORSA) RF data, the FREYA neutral beam deposition package, and a given toroidal electric field, thereby providing a general model for the distortion of the electron and ion distribution functions resulting from auxiliary heating and current drive injected from the plasma periphery. The distributions are taken to be toroidally symmetric and independent of azimuthal angle about the ambient magnetic field. Radial drifts are neglected. With the bounce-average, account is taken of variations as a function of (non-circular) radial coordinate, poloidal angle, and two momentum-space directions. A kinetic bootstrap current calculation is included. The code may be run with separate 2D momentum space solves on each flux surface, or in 3D mode including radial transport according to prescribed diffusion and pinch terms. Solution is by splitting for the 3D equations, alternating between direct solve of the implicit 2D-in-V equations and in the radial coordinate. An implicit sparse matrix solve of the full set of 3D equations is now being developed to improve stability and convergence as well as an MPI parallelization.

DKES The Drift Kinetic Equation Solver (DKES) code solves a time independent form of the drift kinetic equation (not bounce averaged) by velocity space expansion in orthogonal polynomials and. This code is widely applied for transport in 3D magnetic configurations, but it presently does not include quasilinear RF diffusion or energy scattering [25]. Although this code requires magnetic flux surfaces, they need not be nested and islands can be treated. High resolution spatial calculations can be performed by inverting a dense block-tridiagonal system of equations.

DELTA5D DELTA5D solves the drift kinetic equation using direct simulation Monte Carlo techniques.[26] Conceptually, the approach includes all relevant physics, but, depending on the problem, at the expense of the computational time needed to obtain adequate particle statistics. This code evolves particle trajectories by solving four coupled ordinary differential equations per particle derived from a guiding center Hamiltonian (i.e., averaged over the fast cyclotron motion). Particle orbits are followed through a five-dimensional phase space consisting of three coordinate directions and two

velocity dimensions (energy and the pitch angle between the particle velocity and the magnetic field). Collisions are included by periodically stopping the orbits, using a stochastic operator to randomly change the particle direction and energy, and then restarting the orbit integration. Although Monte Carlo particle simulation of plasmas is conceptually simple, there are a number of differences from Monte Carlo applications in other fields such as neutron transport or molecular dynamics that make plasma simulations computationally intensive. These include the curved trajectories and reflections that characterize charged particle orbits in spatially inhomogeneous magnetic fields and the long-range, small-angle scattering nature of the collisional process. Also, for processes occurring on relatively long transport timescales, energy conservation is of importance and more accurate integration methods need to be used than in the case of particle-based micro-turbulence simulations. DELTA5D uses the LSODE11 variable step solver, with options including fourth-order Runge Kutta and Gill's method. The code is efficiently parallelized using MPI.

Appendix B: Data Types Being Passed between Modules

TABLE B1

No	Data Type	Standard notation for function and dependent variables	Now computed by	Requires Data Type
1.	2D Plasma Equilibrium	$\Psi(R,Z,t), J(R,Z,t)$	TSC-A,H,TEQ	2,3,13,14
		$R(\psi,\theta,t), Z(\psi,\theta,t)$	JSOLVER	1,2,3
2.	Equilibrium metric quantities	$\left\langle \frac{ \nabla\psi ^2}{R^2} \right\rangle, \left\langle \frac{1}{R^2} \right\rangle, \dots$	TSC-B	1
3.	Surface averaged profiles	$p_i(\psi,t), n_i(\psi,t), q(\psi,t) \dots$	TSC-D	2,4
4.	Surface averaged H & CD	$\langle S^{NB} \rangle(\Psi,t), \langle J_{\square}^{NB} \rangle(\Psi,t), \langle S^{IH} \rangle(\Psi,t), \langle J_{\square}^{IH} \rangle(\Psi,t), \langle \partial J_{\square}^{IH} / \partial E \rangle(\Psi,t) \dots$	NUBEAM, LSC, TSC-E	1,3,8
5.	Surface averaged particle	$\langle S_j \rangle(\Psi,t)$	TSC-C, DEGAS 2	1,3
6.	Surface averaged velocity-	$D_j(V_{\perp}, V_{\square}, \psi, t), j = 1, 4$	AORSA, TORIC	1,3
		$S_{DEP,j}(\psi, \varepsilon, \mu, t), S_{CX,j}(\psi, \varepsilon, \mu, t)$	NUBEAM	1,3
7.	Ray Tracing for EC, EBW, FW, LH	R rays \times S steps/ray \times Q quantities per time step	GENRAY	1,3
8.	Surface Averaged Distribution Function	$f_{0i,e}(u_0, \theta_0, \psi, t)$	CQL3D,	1,6,7
		$n_e^{RA}(\psi, t), J^{RA}(\psi, t)$	TSC-F	
9.	Monte Carlo Distribution Function	$\{\vec{V}, \vec{R}, t\}$	NUBEAM	1,3
		$f_j(\psi, \theta, \varepsilon, \mu, t)$		
10.	Linear Stability	$\xi_n(\psi, \theta), n=1, N$	DCON, PEST	1,3
			NOVA-K	1,3,8 or 9
11.	3D Plasma State (fluid)	$P(\psi, \theta, \phi, t), B(\psi, \theta, \phi, t), V(\psi, \theta, \phi, t)$	M3D, AMRMHD	1,3,10,12
		$P_n(\psi, \theta, t), B_n(\psi, \theta, t), V_n(\psi, \theta, t)$	NIMROD	
12.	3D Perturbed Distribution Function	$\{w, V_{\square}, V_{\perp}, \vec{R}, t\}_j$	M3D, NIMROD, DELTA5D	9a,11
13.	Coils & Vessel	$I_i(t)$	TSC-G	1,14

NOTE: TSC-A ...TSC-H are the different modules within the TSC code (see table SCJ-2)

Appendix C: Useful Existing Utility Components

Several existing utilities are available in scientific computing and some are candidates to be used in SWIM.

Data Management: The MDSplus data acquisition and management system combined with a relational database for metadata is used at the three large U.S. experimental facilities for a complete data management solution, and will be evaluated for use in SWIM. The lack of parallel I/O capabilities in MDSplus are a known concern, but recent work by the NFC project to integrate the parallel I/O capability of Globus' XIO utility, a derivative of GridFTP, is showing promise as a way to maintain the existing MDSplus API while providing efficient data transfer capabilities. Other distributed data solutions include SRB/MCAT and at a higher level the OGSA-DAI distributed data access interface has several implementations being built. For metadata management most systems now store metadata separately from the actual data. This allows finding desired data based on attributes without having to access and manipulate the large data sets it references. Metadata attributes can include historical information (provenance), a complete set of computational artifacts (provenience), or other information like textual comments entered by a scientist. Systems include electronic notebooks like that already used in the fusion experimental community, the multi-lab Electronic Notebook, the Scientific Annotation Middleware and Obsidian.

Job Management: Distributed systems need the ability for job launch, real-time monitoring, and management. Most national computational facilities use some batch management system like PBS or SLURM, but distributed frameworks like SWIM codes must run on a variety of platforms and locations. Portals technology allows users to monitor and manage jobs from any web browser and has been successful in the GriPhyN grid computing effort and in weather modeling. The FusionGrid Monitor, Particle Physics Data Grid and the Earth Systems Grid provide some job handling capabilities that could be used in SWIM.

Visualization and collaboration: Producing a large quantity of data necessitates the ability to visualize results. There has been a significant amount of previous work by the fusion community on visualization. The NFC project using the open source SCIRun has created a fusion-based power application for visualization of NIMROD data, and the CEMM project is using the commercial application AVS Express to visualization both M3D and NIMROD data. The geographically distributed nature distributed teams like SWIM necessitates the need for team members to work together effectively while in separate locations. Such effective work has collaboration needs beyond standard teleconferencing. The NFC project has experience deploying collaborative tools that allow both human interactions (audio and video) as well as shared applications. These systems are already installed in the major U.S. fusion laboratories, and can now be deployed on desktop and laptop systems at a minimal cost.