Center for Tokamak Transients Simulations

Contents

1. Introduction and Background

Disruption refers to the premature termination of tokamak plasma discharges through sudden loss of macroscopic stability and energy confinement. Plasma disruptions are sometimes called the “Achilles Heel” of the tokamak, because they interrupt operation and release enough energy to damage large devices. While the tokamak is presently the leading concept for a commercial fusion power plant because of its stability and exceptional confinement properties, progress in disruption control is critical for the tokamak reactor concept to be realized.

Disruption control includes prediction, avoidance, and mitigation, and all three aspects have been major parts of the world-wide tokamak experimental program for decades. Much progress has been made in demonstrating high performance disruption-free operation over 10s of seconds or more [1]. Experimental techniques have also been developed to determine when a disruption is imminent, and mitigation techniques have been developed using “killer pellets”, massive gas injection, and/or shattered pellet injection.

However, experiments on today’s devices will never fully inform us about disruptions in a reactor-scale tokamak such as ITER, now being constructed in France. The plasma current and magnetic field in ITER will be 5-10 times the value in today’s machines, and the energy associated with the magnetic field will be over 100 times greater. High fidelity computer simulations can help bridge the gap between today’s experiments and the next generation experiments such as ITER. In particular, they can be used to study fundamental questions, such as when linear instability implies the onset of disruption and why locking of magnetic perturbations to fixed external structures so frequently leads to disruption. They can be used to characterize known disruptive dynamics to improve understanding of their consequences. They can also be used to help engineer mitigation systems that are being developed to protect ITER and follow-on experiments. However, our models must be carefully validated using today's experiments to establish confidence in their predictive capabilities for these important problems.

There are many possible causes for disruptions, and their phenomenology can involve many effects [2,3] over multiple temporal and spatial scales; however, they necessarily involve large-scale macroscopic dynamics. Whether natural or due to external perturbation, the initial instability typically effects a fast "thermal quench" (TQ), when the magnetic configuration loses confinement or radiative losses overwhelm local heating. This is followed by a relatively slow "current quench" (CQ) that transfers large electrical currents into structural components and/or generates a population of relativistic "runaway" electrons which eventually impact the wall.

To date, the most comprehensive descriptions for disruption dynamics are based on the magnetohydrodynamic (MHD) equations. Some intrinsically kinetic effects (such as runaway electron formation), atomic physics effects (such as impurity radiation), and closure terms (such as thermal conduction) are necessary for disruption modeling. We call these extended MHD (or XMHD) models. Applicable numerical implementations treat at least part of these complex systems implicitly to avoid severe time-step limitations. However achieving sufficient efficiency to run high-fidelity simulations of the size required for the next generation of fusion experiments on the latest computer hardware poses mathematical and algorithmic challenges.

Here, we propose a new center to develop and apply numerical simulation to disruption control in tokamaks. We will study nonlinear tokamak stability with the aim of establishing where safe operation can be achieved at or beyond linear instability thresholds. We will apply advanced modeling to the TQ phase to understand what controls the balance between radiation, conduction, and convection in the loss of thermal energy. Modeling of the CQ will address vertical stability, the development of asymmetric currents, structural loading, and the generation of runaway electrons. We will incorporate recent developments for kinetic effects in simulations of neoclassical tearing modes and resistive wall modes, instabilities that precipitate disruption through strong interactions with plasma flow. For mitigation, we will develop new models of shattered pellets and incorporate these models in XMHD simulations that will investigate their ability to penetrate and radiatively cool tokamak plasma. This set of simulation studies represents a higher degree of physics integration than previous tokamak MHD modeling efforts.

The approach of our center is to build upon two XMHD codes maintained by the PIs, NIMROD [4] and M3D-C1 [5], and to couple them to other codes as appropriate and required to model and understand different aspects of plasma disruption. These codes have been used widely to model the onset and nonlinear evolution of many global instabilities in tokamaks including sawteeth [5-8], tearing modes [7,9], energetic particle modes [10,11], edge-harmonic oscillations [12,13], disruptions [14,15], disruption mitigation [16,17], and vertical displacement events (VDEs) [18]. Having two codes that solve nominally the same sets of equations but have widely different spatial representations and algorithms has proven to be indispensable for code validation. It will also enhance confidence when simulating important disruption dynamics in new physical regimes and when making predictions for mitigations systems. Thus, as described in Section 2, we propose simulations studies with an intentional measure of redundancy using both codes. However, detailed parameter scans will normally be undertaken by one code.

Both codes use implicit methods to avoid time-step restrictions when solving nonlinear XMHD systems of equations as initial-value problems [4,19]. The M3D-C1 code solves the equations in a potential representation that preserves the divergence constraint on magnetic field (**B**) and fully decouples parallel vorticity, parallel flow, and compressive dynamics in the large aspect-ratio limit [20]. The potential-based system is solved using C^1 triangular elements in the poloidal plane and Hermite cubic polynomials in the toroidal angle. NIMROD solves the primitive-field (**B**, **V**, *n*, and *T*) XMHD equations using 2D nodal spectral elements and finite Fourier series for the third periodic coordinate. Simulation capabilities for resistive-wall effects, energetic particles, parallel kinetics, and radiation are relevant to disruption simulation and are briefly described where appropriate in Section 2.

Both NIMROD and M3D-C1 are HPC codes that routinely utilize 10-50 thousand processors for the MHD parts alone [5,21,22]. M3D-C1 has benefited from being part of the NERSC Exascale Science Applications Program (NESAP) and has made progress on optimizing it's kernels for the KNL processor. However, the problems that will be addressed in this proposal, especially those involving the kinetic closures discussed in sections 2.3-2.4, will require substantially more processors and better performance than what is now available.

It follows that an important component of this proposal is the inclusion of and close coordination with a team of experts in several areas of computer science and applied mathematics. Performance optimization, improvement of parallel solvers and development of novel preconditioning strategies will improve efficiency on the latest generation of massively parallel computers. Alternative discretizations and mesh adaptivity will be investigated to improve the stability, accuracy, and efficiency of the XMHD codes. Modern data management and software engineering practices will be introduced into the project. These developments will not be done in isolation, but to fulfill the project needs in developing increasingly complex simulation codes to run on increasingly complex and powerful hardware. Also, results on disruption onset, mitigation, and mechanical loading will be evaluated in the context of uncertainty quantification.

Members of our physics team have a long history of pioneering tokamak physics studies with NIMROD and M3D-C1. Besides publication of results in refereed journals and conference reports, presentations by our team have been made via recent invited talks at the American Physical Society Division of Plasma

Physics meeting [23-26], the annual "Theory and Simulation of Disruption Workshop" [27] and the meetings of the International Tokamak Physics Activity (ITPA) in MHD, Control, and Disruptions, of which two PIs are U.S. representatives (V. Izzo and S. Jardin). Members of our team have also received several direct contracts from the ITER Organization to perform specific disruption-related calculations required for design.

Several of the PIs of this proposal were also responsible for completing the FY2013 Department of Energy, Office of Science, Office of Fusion Energy Sciences (OFES) Theory Milestone target on Disruption Physics. The successful completion of this high-profile milestone involving the prediction of disruption forces in ITER and runaway electron confinement during massive gas injection (MGI) was accepted by (OFES) and mention of this appeared in the congressional bill funding OFES for the following year.

2. Proposed Research and Methods

The focus of our center is to develop the most advanced 3D macroscopic stability models of a tokamak and to apply these to a number of different issues relevant to minimizing the occurrence and effects of plasma disruptions. Each issue has significant theoretical and numerical challenges, but with the coordinated cross-disciplinary team proposed here, we expect to make enough progress to inform and guide experimental studies.

2.1 Ideal MHD Driven Disruptions

In this section we discuss activities related to disruptions caused by violating an ideal MHD stability threshold, such as exceeding the pressure or current limit. The TQ and CQ phenomena discussed here are also present in the more general class of disruptions discussed in the following sections.

2.1.1 Prediction and Avoidance of Disruptions

An idea that is gaining acceptance is to use real-time equilibrium reconstruction and linear stability analysis to guide the control systems in ITER to avoid disruptions [28]. The knowledge of proximity to a stability boundary could be used, for example, to adjust the aiming and power levels in the heating and current-drive systems so the pressure and current profiles evolve in such a way that they do not enter the unstable region.

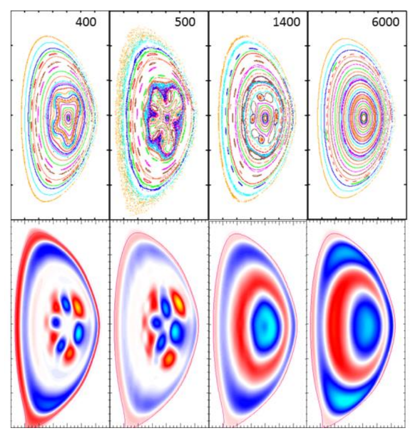

Figure 1 Poincare plots (top) and change in temperature (bottom) from the start of the calculation at 4 times [29]

However, it is well known that crossing some linear instability boundaries does not lead to a disruption. Examples of this include the occurrence of repetitive internal kink modes (sawtooth cycles) or edge localized modes. The control algorithm needs to be sophisticated enough to recognize that linear stability boundaries for these are not disruptive stability boundaries.

In addition, it is known that if pressure-driven instabilities are localized enough, in some cases they will just lead to locally enhanced transport that regulates the further increase of the pressure and likewise does not lead to disruption. Examples of this have been discussed in [29] and [30] and we illustrate one calculation here.

Shown in Figure 1 is a NSTX plasma discharge 124379 at time 0.64 s. Increasing the central neutral beam heating causes an internal (4,3) mode to go unstable near the q = 1.33 surface. This instability distorts the magnetic surfaces in such a way that parallel thermal conductivity acts to reduce the pressure in the center of the discharge to the point where it becomes linearly stable and the magnetic surfaces reform. It can be seen that the surfaces initially deform, then become stochastic in the center, but eventually completely heal and the configuration returns to axisymmetry to a high degree by the final time t=6000 τ_A. (Note that because of the change in the shift of the magnetic axis, the final temperature snapshot shows in-out distortion). The net effect of the localized 3D MHD instability was to increase the effective thermal transport in the center of the discharge.

Figure 1 illustrate a "soft β-limit" that does not lead to a disruption. Other calculations where the linear stability limit was exceeded over a larger volume do lead to a hard disruption [14,30]. We plan to perform extensive simulations of this sort with the goal of coming up with a criteria of when linear stability calculations can be used to predict disruptions, and to compare this with recent nonlinear analytic theory [31] and with high-β NSTX and DIII-D experimental data. We expect that plasma rotation will also play a role in this criterion and will also investigate when active probing of the plasma can be used to detect nearness to a disruptive stability limit.

2.1.2 3D Modeling of the Thermal Quench

The thermal quench (TQ) refers to the rapid loss of plasma pressure and temperature during the disruptive transient. In today's high performance discharges, the temperature typically decreases rapidly from its pre-quench value of several kilo-electron-volts (keV) to 10s of electron volts (eV). The timescale for this can be as short as 1 ms in a large tokamak; hundreds or thousands of times faster than the pre-quench energy confinement time. The TQ may be caused by one or more 3D global instabilities that destroy magnetic surfaces or advect the plasma into the wall, ruining the confinement properties of the device. TQ also results when an accumulation of high-Z impurities causes radiative collapse, or when impurities are injected for mitigation. The rapid heat loss during the TQ potentially produces damaging thermal loads on surrounding material surfaces. The sudden decrease in plasma temperature causes a sudden increase in resistivity that leads to the subsequent current quench (CQ) and associated large electric fields that may generate relativistic beams of runaway electrons.

High plasma density also causes TQ in tokamaks [32]. Beyond an empirically determined limit, which is proportional to the plasma current divided by system size, disruption becomes extremely likely. The empirical scaling of the density limit is robust, and yet the physical mechanism triggering the disruption has remained elusive. The phenomenon involves impurity density, radiation, energy transport, turbulence, and MHD instability. How the physical mechanisms are coupled and what primary mechanism is responsible for the observed scaling are open questions that we will investigate through simulation. Two leading theories are based on radiative cooling of the plasma by impurities; one considers radiative collapse that propagate inward from the edge[33], and the other is based on radiative cooling within magnetic islands that causes explosive island growth [9,34,35]. A goal of this proposal is to clarify the physics underlying this apparently universal phenomenon.

At present the detailed mechanism of how heat is lost during the TQ in a non-density-limit disruption is also poorly understood. Parallel heat transport along chaotic temporally evolving magnetic field lines certainly plays a role, as does a concomitant impurity influx from the surrounding structures. Presently, it is not clear whether a fluid (extended MHD) model with a diffusive form for anisotropic thermal conduction and evolving impurity species (including radiation) can quantitatively model a TQ in a large tokamak or whether a more fundamental kinetic approach (closure) is required. To investigate this, we will further develop our drift-kinetics computations to incorporate the relevant heat-flux and stress closures, as described in Sect. 2.3.1 for modeling magnetic islands. Further developing the PIC capabilities of the two codes [10,36] for this purpose will also be evaluated.

Plasma-surface interaction during the TQ leads to the generation of impurities and must be included as a boundary source term in the MHD modeling. Effects from neutral dynamics are also likely and should be evaluated. Since this topic relates to the edge modeling area, opportunities exist for incorporating the plasma/wall interaction models developed there. This will be considered as part of a WDM modeling effort in the later years of this proposal.

2.1.3 3D Modeling of the Current Quench

The CQ typically follows the thermal quench in a few milliseconds. For the majority of the CQ, the plasma resistivity has been shown to be in agreement with the classical Spitzer value for the low temperature post-TQ plasma [37]. However, it is also well documented that immediately following the TQ, the current profile flattens faster than can be explained by relaxation due to classical or neoclassical resistive processes in an axisymmetric plasma [38]. This rapid flattening, combined with either magnetic flux or magnetic energy conservation, is believed to lead to the characteristic "current spike" at the start of the disruption.

Several researchers have shown that a current spike of the correct magnitude can be obtained in a 2D disruption simulation by suddenly inserting a helicity-conserving hyper-resistivity term in Ohm's Law. [39-41] This technique was useful for explaining the current spike in TFTR and JT-60U experiments, but is not suitable for predictive modeling as it requires knowledge of when the current spike occurred to turn it on in the code. We also note that current spikes due to MHD activity have been observed in 3D NIMROD disruption mitigation simulations with a fixed boundary and a distributed current-drive source [42], and that alternative explanations for the current spike have been proposed [43].

One of the goals for this proposal is to accurately model the current quench in predictive 3D extended MHD simulations without *ad hoc* modification of the transport coefficients and/or hyper-resistivity to match the experimental traces. Details of the current quench are important as it is in this phase that the majority of the plasma current is transferred to the vessel, both inductively and through conduction on the open field lines (halo current). This is also the phase when large populations of runaway currents can be generated. Since the current quench follows the thermal quench (when the plasma suddenly cools), using realistic values of the Lundquist number should be feasible.

The formation of large populations of high-energy electrons during a disruption is an important concern for ITER [44,45]. The PIs of this proposal have experience modeling runaway electron formation in 2D (axisymmetric) simulations [46] and in modeling the confinement of runaways in 3D disruption simulations [47,48]. While another SciDAC group, the Simulation Center for Runaway Electron Avoidance and Mitigation (SCREAM), is focused on developing improved models for runaway electrons during disruptions, our initial focus will be to perform stand-alone free-boundary M3D-C1 and NIMROD simulations without a runaway electron model. Once we are convinced that we have a valid 3D model of a disrupting plasma (that is numerically stable but physically unstable), we will work with SCREAM to couple our codes to runaway electron models that they are developing.

Sub-gridscale models for highly nonlinear events: Plasma disruptions are notoriously difficult to model. The nature of the disruption itself means that there are multiple physically unstable modes interacting which keep driving shorter wavelength modes until any numerical representation eventually runs out of resolution. It has been recognized since the earliest simulations that it can be difficult to distinguish between physical and numerical instabilities when modeling plasma disruptions [49,50].

The challenge is to incorporate new dissipative terms in the equations we solve that act in a physical way to mimic the effect of the shortest wavelength modes without distorting the long wavelength modes that are driving the disruption. To some extent, finite Larmor radius (FLR) effects will provide smoothing, but this is likely not enough. One approach is to generalize the concept of "artificial viscosity" that was introduced in the 1950s to allow calculation of shock propagation by spreading out the shock front over several computational zones in a way that does not affect the shock speed [51].

Another approach is to follow the formalism of Miura for a Smagorinsky-type sub-grid-scale model in MHD turbulence simulations [52]. A third approach is to add a physics-based anomalous electron viscosity term [53,54] in regions where the field lines are stochastic and local velocity gradients are very large. This needs to be implemented in a physics-justifiable way such that it is negligible for quiescent plasmas but is active in turbulent disrupting plasmas [38]. Evaluating these different approaches through experimental validation will be an objective of this proposed task.

2.2 Vertical Displacement Events (VDE) and Resistive Wall Modes (RWM)

VDEs and RWMs involve the self-consistent interaction of the tokamak plasma and surrounding conducting structures such as the vacuum vessel. Recent improvements to the M3D-C1 and NIMROD resistive models give these codes unique capabilities to model the axisymmetric and non-axisymmetric evolution of the plasma during disruptions. Section 2.3 includes a section on the nonlinear effect of the wall on the locking and growth of low-n tearing modes.

2.2.1 Resistive Wall Models

M3D-C1: In M3D-C1, the resistive wall is now modeled as a spatially resolved region in which only the equation $\dot{\mathbf{B}} = -\nabla \times (\eta_{wall}\,\mathbf{J})$ is solved (Fig. 2). This differs from the more typical method of modeling the resistive wall as a boundary condition in that it does not couple all mesh elements adjacent to the wall. Thus, our approach may improve parallel scaling, as well as allowing the modeling of a wall of arbitrary thickness. Verification of M3D-C1 against analytic solutions of resistive wall modes finds excellent agreement spanning both the inertial and resistive-wall limits of the mode [18].

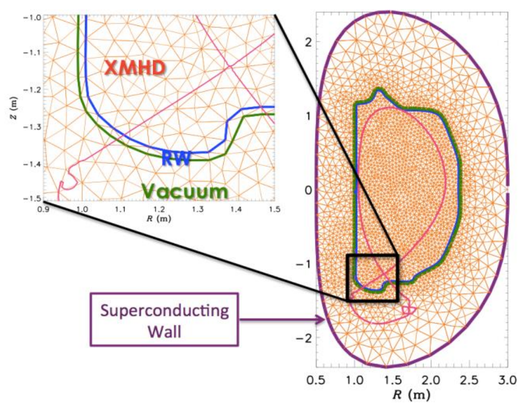

Figure 2: An example of a low-resolution M3D-C1 mesh with a resistive wall. The mesh contains three regions: the Extended-MHD (XMHD) region enclosed by the blue curve; the Resistive Wall (RW) region between the green and blue curves; and the Vacuum region, between the green curve and the computational domain boundary (purple curve). The wall region is only one mesh element wide in this figure, but this is not a requirement of the method. [18]

NIMROD: Two approaches for resistive-wall effects have been implemented in NIMROD. Both use the thin-wall approximation for the immediate interface surrounding the XMHD region. Similar to the original M3D-C0 code, one [55] uses a Green's function computation of the vacuum response, solved by the GRIN code [56], and the interface coupling is temporally explicit. The other uses a meshed representation of the external vacuum region [57], where magnetic field is advanced with a diffusion equation, and the coupling is implicit. The second approach admits both thin-wall and thick-wall modeling, provided that the wall geometry is axisymmetric.

2.2.2 Vertical Displacement Events

The largest transient electromagnetic forces on the tokamak vessel occur when axisymmetric control of the plasma current position is lost, and the plasma current rapidly moves toward the wall. This results during normal operation by exceeding a stability threshold in elongation, for example, or after a thermal quench in which the inductance of the plasma rapidly changes [2,3]. In either case, we refer to the subsequent axisymmetric motion of the plasma as a VDE. The potential for significant harm to large tokamaks warrants comprehensive analysis, including large-scale simulation, of all forms of VDEs.

Asymmetries often develop during VDEs. In particular, when the edge safety factor q_a drops during (or before) the VDE—due to scraping-off by the wall or by impurity cooling of the edge—the plasma becomes unstable to an external kink or RWM [58]. This results in asymmetric forces on the vessel, which pose a special threat to the integrity of the vessel [59-61]. Asymmetric forces are more difficult to brace against than axisymmetric forces, and they can rotate, potentially in mechanical resonance with the vessel structures.

Although the fundamental dynamics of VDEs are macroscopic (MHD-like), there are a number of complications for numerical modeling. First, the gross motion of the plasma torus induces current and transfers energy to an open-field "halo" region that surrounds the torus [62]. The process is slow relative to Alfvén-wave propagation, so the amount of current induced in the halo is constrained by force-balance [63]. However, current density through surfaces may be limited by sheath effects. In addition, current paths through vessel components can be complex and depend on whether plasma contact is able to short gaps between conductors [64]. Also, radiation has a significant role in edge plasma physics during disruptions [3]. Comparison of disruptive behavior in the Joint European Torus (JET) with its old carbon wall and with its new ITER-like beryllium/tungsten wall [65] shows how sensitive these processes are to wall materials. To address the challenges associated with VDEs, we propose to enhance our extended-MHD models and to validate them against existing and forthcoming experimental data. We also propose to use the models to explore methods of disruption mitigation that minimize non-axisymmetric wall forces while safely avoiding runaway electrons.

Our previous validation work compared results of the M3D-C0 code to data from JET. M3D-C0 simulations of JET disruption shot 71985 found qualitative agreement with measurements of the time history of the vertical displacement, the amplitude and time history of the halo current, the toroidal current asymmetry, and the toroidal rotation [66]. The Noll relation [67] between asymmetric wall force and vertical current moment was verified in the simulations, as well as toroidal flux asymmetry [68]. These computations will be repeated with M3D-C1 and NIMROD with higher resolution and more realistic values of plasma Lundquist number (*S)* and wall Lundquist number (S_{wall}).

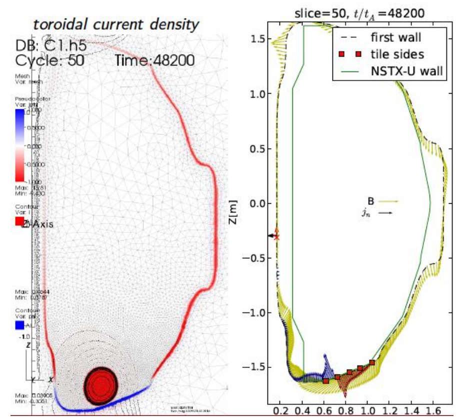

Figure 3 Left: the toroidal current density in an M3D-C1 simulation of a VDE in NSTX. The eddy currents in the wall are clearly visible. Right: The current density normal to the wall (i.e. Halo current) at the same point in the simulation. Red and blue arrows represent current into and out of the wall, respectively. Yellow arrows represent the direction of the magnetic field.

Initial simulations of VDEs for realistic values of S and S_{wall} have been carried out using M3D-C1 in both DIII-D and NSTX geometry. These simulations are initialized with a vertically unstable Grad-Shafranov equilibrium, and evolved through the current quench. To reduce the computational cost of the simulation, an axisymmetric simulation is run and periodically tested for linear non-axisymmetric instabilities. When non-axisymmetric instabilities are found, the fully three-dimensional calculation is started. The calculated toroidal current density and halo currents during the late axisymmetric phase of a VDE in NSTX geometry are shown in Figure 3. When edge scrape-off yields $q_a \sim 2$, asymmetric instabilities begin to develop. In the cases studied, the rate at which these instabilities develop is comparable to the rate at which q_a is dropping. Therefore, the plasma is still predominantly axisymmetric when $q_a \sim 1$, at which point a 1/1 mode rapidly grows. This causes a thermal quench, after which the plasma current returns to near-axisymmetry and quickly decays to zero.

The details of this sequence of events are found to depend sensitively on the thermal profiles in the plasma and in the halo region. In the case of a “cold-VDE,” in which the plasma loses its thermal energy before going vertically unstable, the q profile is observed to rise due to the rapid decay of the plasma current. Whether q_a drops sufficient to cause additional asymmetric instabilities depends on the relative rate of the current decay and the scraping-off rate. In the cases where asymmetric instabilities develop, their growth rates are found to be strongly dependent on the resistivity of the halo region.

Accurate modeling of asymmetry during VDEs therefore requires accurate modeling of the halo region. Reference [63] provides analytical estimates for the width of the halo current and its relation to kink amplitude. It also provides analysis for the voltage along the halo current path. One objective of the proposed study is to apply nonlinear simulation to kink-unstable VDEs and make comparisons with the analytical results. We will further investigate how transport effects in the relatively cold open-field-line halo plasma influence its behavior. We will also investigate how contact with a conducting structure and the presence of the halo region influence kink stability of the plasma torus.

Validation will be a major focus of the proposed research. Data from NSTX, DIII-D, and JET regarding the timescale of the VDE and measurements of the magnitude, location, asymmetry, and rotation of halo currents will also be compared with M3D-C1 and NIMROD simulations.

2.2.3 Resistive Wall Modes

Resistive wall modes (RWMs) are external kinks that have a positive growth rate due to the finite resistivity of the vessel. In the ideal-MHD limit, resistive wall modes should be unstable whenever the no-wall external kink is unstable. However, it is common for tokamak plasmas to run stably above the no-wall limit. Recent work with kinetic-MHD models has been successful in reproducing the observed linear stability [69,70] . However, these models are limited in that they are restricted to linear perturbations with purely ideal-MHD magnetic response, and they use simplified kinetic models that exclude potentially important effects (such as the effect of the perturbed electrostatic potential on particle orbits).

First, we will calculate the effect of non-ideal physics already present in the extended-MHD model on the stability of resistive wall modes. It has been noted that in the presence of dissipation, plasma rotation may be sufficient to completely stabilize the RWM [71]. However, the effects of most of the physical dissipative terms in extended-MHD, including the large parallel ion viscosity and parallel thermal conductivity, have not yet been evaluated in this context. Furthermore, non-dissipative terms such as the gyroviscosity (representing finite Larmor radius effects) and diamagnetic drifts have been

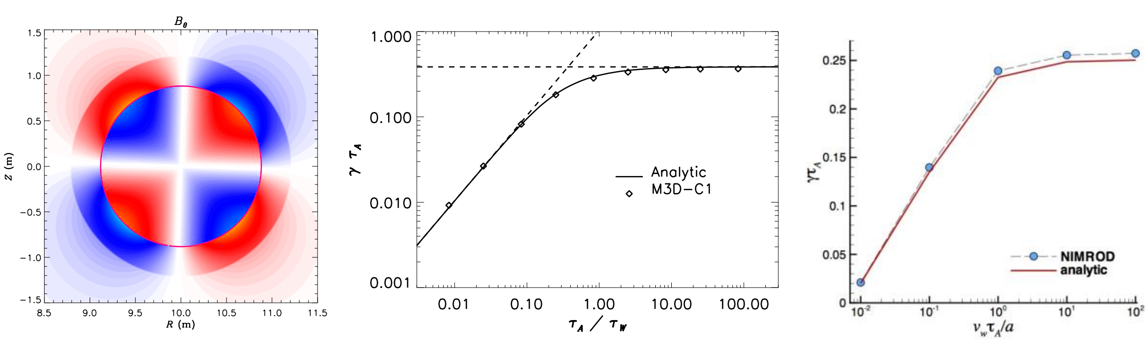

Figure 4—*Left*: the poloidal component of the perturbed field in two M3D-C1 simulations of a RWM in the resistive-wall limit ($\tau_A/\tau_W \ll 1$). The resistive wall in both cases is a circular annulus centered at $(R, Z) = (10, 0)$ extending from $r = 0.9$ to $r = 1.2$. *Center*: the linear growth rate of the RWM as calculated by M3D-C1 (symbols), compared to the analytic solution (solid line). *Right*: analogous linear RWM verification of the NIMROD thin-wall/meshed-vacuum implementation [57].

shown to stabilize other slowly growing modes [72], and the effect of these terms on the RWM will be evaluated.

Second, we propose to develop more complete and self-consistent models of kinetic corrections to linear MHD RWM stability by coupling the M3D-C1 code to a linear-perturbation extension of the DK4D code[73] (under development) and using NIMROD's internal drift kinetic equation solver [74] coupled to its fluid model. (Similar and/or equivalent kinetic couplings are described more in Sec. 2.3.1.) In particular, we will perform linear stability calculations of RWMs that will be verified against reduced models and existing kinetic-MHD solvers (e.g., MARS-K [70], MISK [69,75]). Furthermore, we will validate the model against experimental data and perform predictive simulations of RWM stability for existing tokamaks as well as ITER.

Third, once the linear predictions of these models have been satisfactorily validated, we will carry out nonlinear calculations of resistive wall modes, focusing on whether and how the RWM precipitates a thermal quench. It is not known, for example, whether the thermal quench is due to the breakup of the magnetic surfaces, or through the kinking of the plasma into the wall. Answering this simple question could inform both control and mitigation strategies for when a RWM is detected.

In addition, we note that the M3D-C1 and NIMROD codes both have a PIC capability [10,36] that is able to model "fishbone" modes caused by energetic particles interacting with the (1,1) mode. A stretch-goal of this proposal is to apply this capability to model the experimentally observed interaction of off-axis fishbone-like instabilities with the resistive wall mode [76,77], as well as other modes. The relative advantages of PIC vs continuum modeling of kinetic effects will be explored for RWM and other MHD modes.

2.3 Neoclassical Tearing Modes and Mode Locking

Tearing modes (TMs) are one of the leading causes of disruptions in tokamaks [78-80]. Experimentally, they are either triggered by other magnetic perturbations, or arise "spontaneously", presumably by exceeding a linear threshold. In cases where they are triggered, this apparent meta-stability is believed to be due to the effect of the island on the bootstrap current; this type of nonlinearly unstable mode is called a neoclassical tearing (NTM). Most often, but not always, TMs with small island widths rotate at the plasma rotation rate. However, they can slow down due to coupling with a resistive wall [81,82], or by a viscous torque created by the perturbed ion orbits [83,84]. Eventually the island may stop rotating and become a "locked mode". Disruptions often occur soon after entering this locked state, although how this disruption occurs (or whether it will occur) is not well understood.

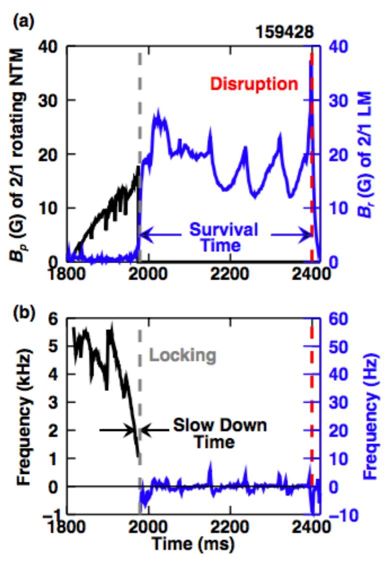

Figure 5 After growth and saturation, NTMs typically slow down, lock, and cause disruptions over relatively long time scales. [80]

NTMs are distinguished from classical TMs [85] by the flattening of the pressure profile across the magnetic island and subsequent reduction in the bootstrap current [86]. Thus, numerical modeling of NTMs requires sufficiently accurate closures for the heat flux and the electron parallel stress in order to compute temperature equilibration about an island and the perturbation to the bootstrap current, respectively. The ion parallel stress produces a torque from the neoclassical toroidal viscosity (NTV) [87,88] that becomes pertinent when studying transition to a locked state. A self-consistent calculation of these terms requires solving the drift kinetic equation (DKE) in a Chapman-Enskog-like form [89,90] .

Preliminary discretizations of these DKE equations have been implemented and verified in the time-dependent, axisymmetric DK4D code [73] and within the NIMROD framework itself [74]. The computational cost, and the algorithmic advances required for full application of these closures, suggests that approximate closures [91-94] may also be useful in the near term for preliminary explorations of the large NTM parameter space.

Recent analysis of DIII-D data [80] shows that the time scales for a born-rotating but ultimately locked 2/1 NTM are ~200 ms in the quasi-stationary state, ~15 ms for the locking dynamics, remains locked for ~300 ms until a disruption terminates the discharge. The long-time scales associated with each phase of the tearing mode evolution are encouraging for active feedback [95] and mitigation techniques [96] used to avoid the disruption. However, these long timescales present challenges to numerical modeling. Our plan therefore is to pursue a strategy that decomposes the work into the physics of each experimental phase.

2.3.1 Kinetic-MHD Stability of NTMs:

Modeling of NTM growth and saturation places a large emphasis on the development and verification of DKE closures. Current verification exercises of DKE discretizations have been limited to axisymmetric geometry in benchmarks with codes such as NEO [97]. These efforts [73,74] have focused on efficient and accurate discretizations of velocity space. In the presence of 3D perturbations, it will be even more challenging to resolve velocity space. Thus, considerable work is needed to extend algorithms and advance computational approaches to model DKE closures when simulating 3D macroscopic instabilities.

The DKE closures we propose to solve are fully self-consistent with the fluid equations being solved by the extended-MHD codes [89,90]. This significant theoretical advance, when fully implemented, is equivalent to enabling the extended MHD codes to become 5D drift kinetic codes while taking advantage of the investments already made in algorithms and coding infrastructure. In going from a 3D initial-value code to a 5D initial value code will require significant increases in resources. Investments made here prepare these codes for exascale.

The DKE closure for M3D-C1 will be provided by coupling the MHD code to a fully three-dimensional extension of DK4D [73], which is under development. We will focus on enhancing the efficiency of the upgraded DKE solver to make it appropriate for high-performance computing applications by exploring advanced algorithms, data structures, and meshing techniques. Furthermore, we will develop a method to couple the DKE solver to M3D-C1 that allows for efficient simulation of the hybrid kinetic-MHD problems, including the NTM work proposed here and the RWM simulations discussed in section 2.2.3. We will devote considerable effort to suppressing and/or preventing numerical instabilities that have made initial attempts at similar couplings of MHD solvers to the axisymmetric DK4D challenging. This work will be synergistic with the ongoing efforts of the General Atomics Theory and Computational Sciences Group to develop initially the non-axisymmetric DKE solver using methods similar to those previously developed for DK4D.

Alternatively, the current implementation of the 5D DKE within NIMROD uses finite elements for the poloidal spatial coordinates and the velocity-space pitch-angle coordinate, a Fourier representation of the toroidal direction, and non-classical polynomials for the velocity-space speed coordinate. This extension does not exploit the structure of the 5D equations in the linear solves so there are opportunities for optimization. Presently, the process of reducing the number of degrees of freedom by eliminating interior contributions from the finite elements of the global system solve is not applied in the pitch-angle-velocity-space direction. We plan to develop this capability with the existing infrastructure for immediate application. In the longer term, our plan is to implement 3D finite element infrastructure into the NIMROD framework to achieve further performance optimization for the DKE solves.

Both the NIMROD and M3D-C1 approaches will be verified, first on axisymmetric problems and then for test problems with 3D magnetic perturbations. After this, we will perform self-consistent NTM simulations that begin with a saturated, conventional tearing mode as a seed-island initial condition. By then turning on the DKE closures, the initially saturated mode will be driven to larger amplitude through a process of temperature equilibration inside the island and subsequent loss of the neoclassical bootstrap current. The results of these simulations will be compared to the modified Rutherford equation and/or other reduced models [92] for verification.

2.3.2 Locking of NTMs in the presence of resistive walls and field errors:

When considering whether a tearing mode should trigger preemptive disruption mitigation, it is crucial to know whether the tearing mode is likely to lock. At present, a quantitative predictive model for this does not exist. Therefore, the process through which a saturated, rotating NTM slows down and locks to the wall requires further investigation. The 3D mode introduces new torques on the plasma from the NTV and Maxwell torques from field errors and the drag on the resistive wall. The computation of the Maxwell torques are straight-forward and resistive-wall modeling has been addressed in NIMROD and M3D-C1 in section 2.2.1. The computation of the NTV torque requires closure terms from the ion DKE which produces different challenges relative to the electron DKE. For the electron DKE the fastest time scale is set by the electron parallel speed which is faster than the Alfven speed. Therefore, implicit methods, which enable the success of extended MHD codes, will be necessary to handle the DKE-closures. Alternatively, the solution to the ion DKE is impacted by finite Larmor radius effects (e.g. finite width bounce orbits and ion orbit loss) which require suitable representations of velocity space. These challenges represent an exciting and fruitful area for collaboration between the physicists and applied mathematicians on this project.

We have tested the capabilities of NIMROD and M3D-C1 for simulating magnetic-island locking by running 2D computations with the visco-resistive MHD model. The 1D equilibria have magnetic shear with a large guide-field, and the configurations are linearly stable. Reconnection is driven by imposing magnetic perturbations, i.e. field errors, along the boundary. Figure 6 shows that the two sets of linear computational results on reconnected field in cylindrical geometry agree over the relevant range of flow values; they also reproduce the nearly $\pi/2$ change in phase predicted by analytical theory [98]. The quantitative discrepancies with theory may be due to the difficulty of varying parameters without crossing from one asymptotic regime to another, whereas the analytical relation is evaluated with the layer response time for one regime. We also made direct comparisons of NIMROD's computed nonlinear saturation with the quasi-linear analysis of Ref. [98], modified for slab geometry. As also shown in Fig. 6, the sharp change in net plasma flow at the resonance location, i.e. locking, is reproduced as the imposed perturbation is varied at fixed flow-rate [99]. These linear and nonlinear results enhance confidence that our models can represent the important effects of locking in disruption simulations.

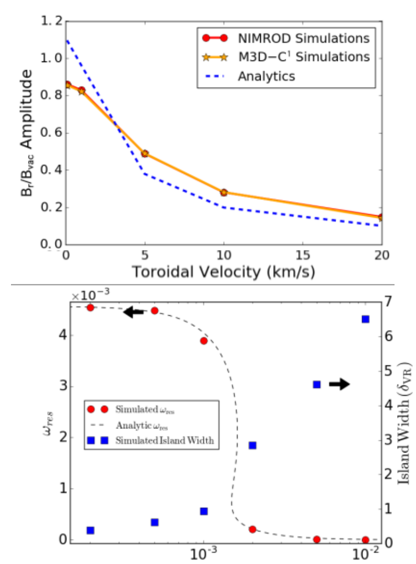

Figure 6. Comparison of linear cylindrical results from NIMROD, M3D-C1, and theory as flow is varied (top) and comparison of nonlinear NIMROD results with quasi-linear theory [99] as perturbation amplitude is varied (bottom).

An important physics topic to address with this modeling is the scaling of the locking to burning plasma conditions. Future burning plasma devices, such as ITER, will have a large moment of inertia relative to the injected torque from the neutral beams. As such, the expected rotation is anticipated to be much slower than present devices. Quantifying the scaling of the torques to burning plasma conditions and the impact of this slower rotation on the probability that an NTM locks to the wall is essential.

2.3.3 Growth of locked modes and how they cause disruptions:

The mechanism by which locked modes cause disruptions is poorly understood. The growth of the mode after locking is possibly due to the resistive diffusion of flux through the wall. This picture is consistent with disruptivity scaling with internal inductance divided by the edge q [80] – near the ideal stability boundary delta-prime is positive and the tearing mode is more unstable in the presence of a resistive wall [81]. As the island grows, it interacts with other poloidal harmonics (3/1, 4/1, etc.) and the overlap causes stochasticity. Whether this will lead to a thermal quench that causes a disruption is not understood. Because the likelihood of disruption is higher if the plasma is located near the wall [80], edge interactions may be important. Following the nonlinear growth of a locked island to understand how this might lead to a disruption is one of the goals of this work.

2.4 Disruption Mitigation

Effective disruption mitigation is critical for ITER and future burning plasma devices in order to avoid damages due to extreme localized heating, deposition of high-energy relativistic electrons, and large electromechanical forces. The U.S. is the lead international partner in the development of a Disruption Mitigation System (DMS) for ITER. Disruption mitigation strategies involve the injection of large quantities of impurities so that the TQ is dominated by radiative rather than conducted heat loss, with MHD activity playing a significant role in the TQ evolution. Mitigation simulations and modeling necessarily require impurity radiation, ionization, and recombination, neutral dynamics and transport, pellet ablation, and MHD activity. A leading candidate system for ITER DMS is the Shattered Pellet Injection (SPI) system [100]. A prototype DMS using SPI, which is being tested on the DIII-D tokamak, has been largely successful [101,102]. New SPI disruption mitigation experiments are being planned on JET and other tokamaks. The ablation of the tiny pellet fragments with small spatial scales in the vicinity of the fragments must be coupled to the macroscopic MHD evolution. Integrated simulations and modeling across multiple physical phenomena with disparate spatial and temporal scales are therefore necessary to extrapolate the DIII-D and other tokamak results to the much higher temperatures and magnetic fields of ITER.

Our proposed work includes first the construction of a self-consistent SPI fragment plume model based on a continuum mechanics approach [103] and development of numerical algorithms to track the plumes. Secondly, the SPI plume model and the tracking algorithms will be implemented into the FronTier-MHD code that contains a comprehensive pellet physics model with electron kinetics, atomic processes, ionization, and adaptive meshing [104-106]. Stand-alone simulations and validation tests will be carried out using FronTier-MHD. Lastly, the local FronTier-MHD code will be integrated with the global NIMROD and M3D-C1 full MHD codes and their validations against DIII-D and other tokamak SPI experiments and applications toward ITER projection will be performed.

2.4.1 Construction of SPI Plume Model and Development of Tracking Algorithms

The use of pellet injection methods for pre-emptive plasma thermal quenches in ITER is gaining increased attention for mitigating plasma disruptions [100,107]. In SPI, a large pellet composed of a homogeneous mixture of frozen neon and deuterium is injected at high velocity inside a "breaker tube" causing pellet fragmentation. A stream of smaller fragments called the "debris plume" travels from the site of impact to the plasma boundary. The frontal "sacrificial" section of the plume cools the plasma enabling deeper penetration of pellet fragments in the rear section. This fragmentation process is illustrated in Fig. 7. It is this inherent coupling of penetrating fragment plume and cooling background tokamak plasma that justifies and necessitates a rigorous integrated modeling approach in order to advance the state-of-the-art in the use of SPI for disruption mitigation in large tokamaks such as ITER.

We propose to address first the structure and kinematics of the fragment plume analytically using an approach based on continuum mechanics [103]. Then, using the global properties of the solid plume component we will construct a self-consistent model for the neutral-source distribution within the plasma due to the propagation and ablation of plume fragments. A theory that can predict the distribution of fragment sizes resulting from the shattering process will then be developed. The problem of extracting the size distribution from a pellet-fracturing model is complicated. In the proposed work, we shall assume for convenience single-sized spherical pellet fragments, or possibly various analytical ones that can be used for input into the FronTier pellet ablation code [104-106]. Specifically, the evolution of the jet plume before entering the plasma will be described from the standpoint of continuum mechanics. For long extended jet-like plumes we will employ a 1-D continuity equation within the paraxial approximation to model the stretching or elongation of the plume due to velocity dispersion as it travels from the site of impact inside the breaker tube to the boundary of the plasma.

From this information, we will devise and test various numerical algorithms that will enable us to track each pellet in the plume as it travels through the plasma until it is fully ablated. The field-line flow of the cold, dense, ionized ablation trail deposited in the plasma will be tracked by the FronTier code [103] as described below until it can be integrated into the NIMROD and M3D-C1 MHD codes to follow its cross-field transport due to magnetic turbulence. These codes at present use a simple Spitzer-Harm diffusive heat flow model to describe the background plasma cooling. However, the diffusive model only applies when the background plasma is collisional. We will therefore, develop analytical models for the kinetic effects that must be considered in NIMROD and M3D-C1 when treating electron heat transport in the ablation cloud, similar to what was done in [107]. We will also develop new simulation tools in FronTier so that the ablation of pellets shielded by the ablation trail of its fellow pellets can be taken into account.

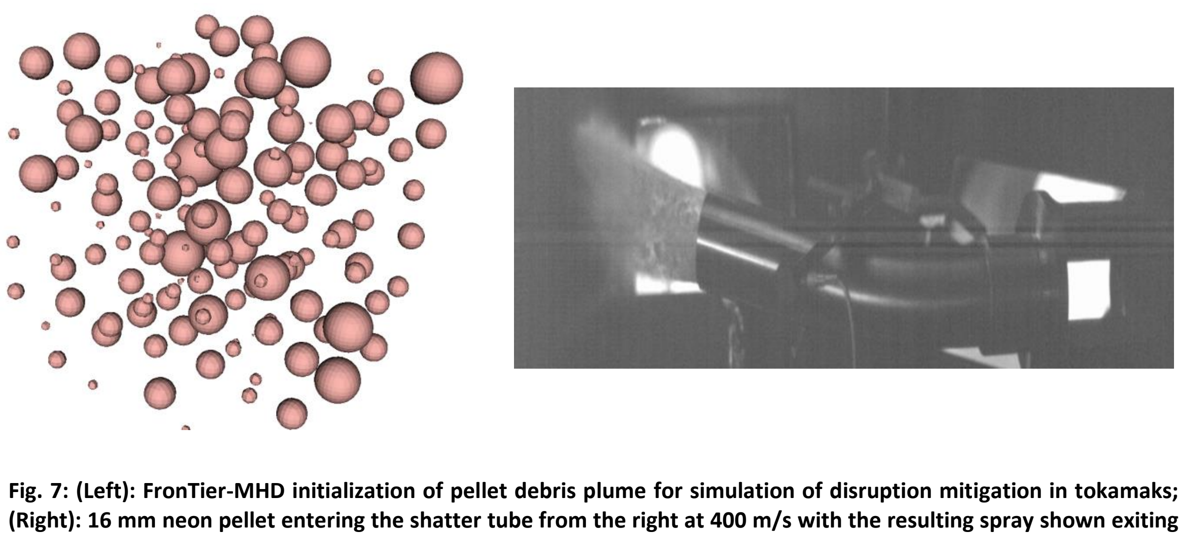

Fig. 7: (Left): FronTier-MHD initialization of pellet debris plume for simulation of disruption mitigation in tokamaks; (Right): 16 mm neon pellet entering the shatter tube from the right at 400 m/s with the resulting spray shown exiting about 1 ms after the pellet entered the tube [100].

2.4.2 Multi-Scale Model for Disruption Mitigation Simulation

We propose to develop multi-scale simulation tools for plasma disruption mitigation processes, which have scales ranging from those of ablation with atomic processes to tokamak device scales. Specifically, we will develop detailed 3D numerical model for the ablation of pellet fragments based on the FronTier–MHD code [104,105] and perform validations and simulations of the structure and kinematics of the injected pellet debris plume. In the second phase of the project, we will couple the FronTier-based pellet ablation code with the NIMROD and M3D-C1 codes and perform studies of plasma disruption mitigation for ITER.

FronTier Local Model for SPI

The 3D pellet ablation code based on FronTier-MHD will individually track a large number of spherical pellets with non-uniform sizes. The code capability also includes the tracking of arbitrary-shaped pellet fragments. The heat deposition of hot electrons, streaming along magnetic fields lines, will be accurately computed for each individual pellet fragment, accounting for the screening of pellets by upstream fragments. The code also supports accurate models for equation-of-state (EOS) in the presence of atomic physics (multi-level ionization / recombination) [108]. The local FronTier ablation model will incorporate all relevant physical effects, such as the cloud charging due to the kinetics of plasma electrons, atomic processes in the partially ionized gas, conductivity model including ionization by electron impact, etc. We will develop and implement in the FronTier-MHD code a novel iterative method for computing local magnetic field in the ablation cloud without numerical expense of full solutions to the resistive MHD equations. For initial simulations, the FronTier-based pellet code will be loosely coupled to tokamak simulations by incorporating a typical time-dependent toroidal magnetic field in the location of the plume and cutoff shielding length due to the curvature of the field from NIMROD and M3D-C1. The model will undergo a comprehensive validation simulation program using data from DIII-D.

Multi-Scale Coupling Techniques: Algorithms and software for a two-way coupling of the local pellet ablation model based on FronTier-MHD with the global NIMROD and M3D-C1 codes will be developed. The coupled codes, executed via a single driver loop, will exchange data during runtime. The coupling algorithms will employ a conservative remapping method that projects physical states in the ablation cloud such as 3D distributions of density, temperature, and conductivity into simulation domains of a global code. In the other direction, background plasma properties and the magnetic field in FronTier will be updated using data of the global codes. The main challenge of the coupling is related to a dynamic identification of a physical domain in each global code, containing the ablating pellet debris cloud, to be mapped into the FronTier-based ablation code. As solutions of the local and global codes will reside on overlapping domains, another challenge is associated with the construction of a smooth, blended solution in the global codes. The location of the overlapping region will be identified based on conductivity values of the ablated material and the magnetic Reynolds number.

Work on the code coupling will leverage methods and libraries under development at other SciDAC centers. We will expensively investigate problems of accuracy and stability of coupling. Before we proceed with the two-way coupling implementation of the production codes in 3D, we will evaluate our main ideas using 1D tests. For this purpose, we will use FronTier's 1D model for the spherically symmetric pellet ablation, a 1D approximation of a global MHD code, and a 1D Lagrangian MHD code "Pressure Relaxation Lagrangian" (PRL) developed at General Atomics [109]. The PRL code has been used for the study of the parallel expansion dynamics and the cross-field drift of the ablation cloudlet. Multiscale coupling in 1D will allow us to test and optimize weight functions and the projection and reconstruction operators, study the behavior of the coupled solution, and detect and resolve possible sources of errors.

2.4.3 NIMROD SPI Simulations and Modeling

With SPI as the primary disruption mitigation (DM) strategy for ITER and experimental data currently available only from one tokamak [101,102], SPI modeling and validation against DIII-D data is a paramount. Individual small pellet shards from an SPI injector will barely penetrate the ITER pedestal, but the SPI concept relies in part on the collective effect of a stream of pellet shards where the leading shards pre-cool the plasma and allow the trailing shards to penetrate more deeply. For this reason, SPI penetration can only be predicted in a simulation that includes a dynamic plasma model. NIMROD simulations of massive gas injection (MGI) [110-112] have further shown that MHD effects such as flux surface destruction will play a major role in the physics of the advancing cold front, so that a fully coupled MHD-SPI/ablation model is needed to predict SPI effectiveness for ITER. With its extensively demonstrated radiation and atomic physics package for DM simulations [110-115] (Fig. 8), the NIMROD code is ideally suited to pursue further DM modeling for SPI scenarios.

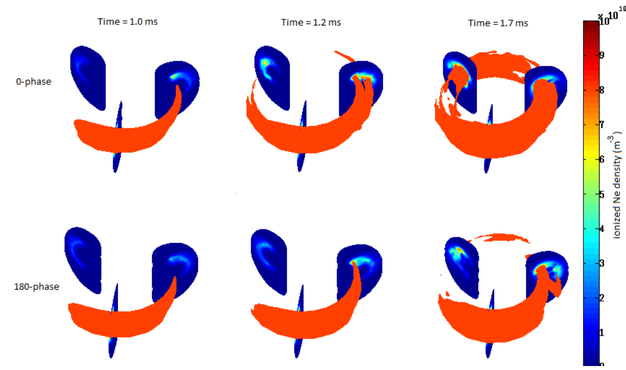

Figure 8: Simulations of MGI in DIII-D with NIMROD showing changes in impurity propagation due to MHD instabilities that alter magnetic topology [115].

In the first stage of the project, before NIMROD/FronTier coupling is complete, scoping and sensitivity studies of SPI will be done with an existing analytic SPI model in NIMROD. Once the code coupling is completed, more realistic simulations using the full model will be pursued. Key validation parameters will include the total radiated energy fraction and current quench duration, and their dependence on the initial pellet size and low-Z/high-Z mix ratio (e.g. D2/Ne mixed pellets, as have been tested on DIII-D [102]). Further modeling will examine the influence of the SPI injection angle, particularly the difference in MHD activity resulting from injection normal to the flux surfaces versus injection at a glancing angle to the q=2 surface, as geometry considerations for ITER DM [116] may require. Beyond validation of DIII-D experiments, the modeling will explore a greater parameter space than the fixed design of the current SPI system allows, including systematic variation of post-shattered shard size, and consideration of additional atomic species. Finally, validation of the effects of increasing machine size and stored energy will be the most critical step for confident extrapolation to ITER. Depending on availability of JET SPI data in the near term, validation against JET data would also be pursued.

2.4.4 M3D-C1 SPI Simulations and Modeling

The M3D-C1 code presently has a basic pellet module [117] and an "average ion" impurity radiation model [118]. During the course of this proposal, we plan to upgrade these models. The SPI model will be obtained and implemented as described in section 2.4.2 in conjunction with the FronTier code. A radiation model that transports and evolves individual impurity charge states and their ionization and recombination [118,119] will be implemented, substantially increasing the accuracy of the radiation treatment. The routines for doing this have previously been installed in the axisymmetric TSC code [120] and the implementation in M3D-C1 will be verified using that. Since the implementation of the radiation capability in M3D-C1 will be largely independent of routines used in NIMROD, and since M3D-C1 uses finite elements in the toroidal direction while NIMROD uses Fourier analysis, we expect these two codes will be able to provide valuable verification checks on one another when applied to the same problem.

2.5 Computational and Applied Math Issues

The emphasis of the CTTS computational and applied math efforts are to advance the discretization technologies as well as ensure the computational performance and scalability of the M3D-C1 and NIMROD codes, including linkage with FronTier, as executed on latest generation of DOE massively parallel computers. These efforts will include improving the discretization methods and equation solvers to support scalable solution; adaptive methods for ensuring accurate solutions; and support for complete plasma simulations as part of validation and UQ. The application of state-of-the-art data and software management methods will support the fusion research community.

2.5.1 Solver Developments for the Extended MHD Equations

Both M3D-C1 and NIMROD simulate 3D implicit extended MHD models requiring the solution of repeated linear systems as the simulation proceeds. These linear systems have complicated sparsity structure and are extremely ill-conditioned due to the fact that the underlying multi-physics/multiscale PDE equations are extremely stiff, non-self adjoint, and characterized by physical phenomena spanning a wide range of time- and length-scales. Although efforts on physics based [5] and discretization based [121] transformations and equation scaling substantially reduces the condition number, the level of ill-conditioning continues to pose significant challenges to the equations solver algorithms, and consequently the computational time for linear solutions often dominates the total simulation time. Presently in the production codes, the most robust solver framework uses a configuration with nested parallelism – combining a parallel Krylov solver, such as GMRES, with a powerful block Jacobi preconditioner, for which each preconditioning block is solved concurrently by the parallel direct solver SuperLU_DIST [122] on a subset of processors. To meet the demands for higher simulation resolution and scalability, and deployment of new HPC systems with KNL and GPU nodes, we propose new solver developments in the following areas.

Scaling up triangular solution: Triangular solution usually takes less than 5% of the factorization time, so it has received little attention in performance optimization [123]. However, in this usage context, the direct solver is not used for a one-off solution, but is used as a preconditioner. Therefore for each factorization performed, many triangular solutions (preconditioner applications) are needed in the iterative steps. To increase both intra- and inter-node parallelism, we will replace the message-driven asynchronous execution by an elimination-tree-guided execution using efficient DAG scheduling. The new scheme will avoid large numbers of outstanding MPI messages overflowing MPI buffers and degrading performance. We will also employ a selective inversion technique to change the sequential substitution procedure into a more parallel matrix-vector product. After these MPI related optimization, we will add OpenMP threading support to further intra-node parallel efficiency.

Leverage Trilinos/Kokkos package for next generation architecture performance portability: Kokkos is a Trilinos package [124,125] that includes a number of essential linear algebra kernels upon which iterative solvers and preconditioners can be built. By leveraging Kokkos, the goal is that solvers can be written with one implementation that compiles/runs on multiple architectures (including on GPUs and multi-core/many-cores) and is highly performant (e.g., can leverage memory and architecture-specific features). To leverage Kokkos the M3D-C1 and NIMROD linear system matrices must be adapted to Trilinos' Tpetra stack. Performance comparisons and trade-offs will be documented across different architecture configurations between the existing linear solver capabilities and the Trilinos/Kokkos ones.

Deploying novel hierarchical matrix preconditioners: Another strategy to overcome the bottlenecks of the traditional factorization is to explore novel approximate factorization algorithms that offer asymptotically lower complexity, sometimes linear. We will employ various hierarchical matrix algebra methods to apply low-rank compression to off-diagonal blocks during factorization. This type of "inexact" direct solver is particularly attractive while being used in the preconditioning context, whereby a tolerance-controlled compression provides great flexibility in the trade-off between the size of the factors, scalability and the iterative solution time. The use of hierarchical matrices has been under consideration for years, but there is almost no production-quality software that can handle large-scale problems. Our new STRUMPACK package [126,127] fills the gap between algorithm advances and parallel software. STRUMPACK uses hierarchically semi-separable matrix structures (HSS) for compression. Our initial evaluation of M3D-C1 matrices shows that it achieves better time and memory performance than SuperLU_DIST. In this work, we will perform an analytical study of the underlying MHD equations and discretized linear systems, determine the best partitioning and ordering schemes and obtain estimates of the off-diagonal rank patterns in order to construct an optimal preconditioner. We will also incorporate other formats/algorithms from the H-matrix family, such as the hierarchically off-diagonal low-rank (HODLR) format to see if these offer computational advantages for M3D-C1 and NIMROD over hierarchical solvers based on HSS matrices. The HODLR investigations will build upon an ongoing Trilinos project to develop highly parallel hierarchical preconditioners.

Structure Leveraging Preconditioners: The most frequently used meshes in M3D-C1 and NIMROD are effectively generated by first developing a 2D poloidal plane mesh and then sweeping this 2D mesh in the toroidal direction to obtain a 3D mesh. It has long been recognized that this structure can be exploited to develop discrete representations and to design preconditioners. For example, NIMROD employs FFTs in the toroidal direction. Within the block Jacobi/SuperLU_DIST strategy described above, each block typically corresponds to each poloidal slice. One limitation of this blocking strategy is that toroidal coupling is ignored in the preconditioner. This leads to slow convergence when the number of poloidal planes is large and coupling in the toroidal direction is significant. Therefore, we intend to augment our preconditioners in ways that capture some coupling in the toroidal direction. One possibility is the hierarchical matrix context. In particular, low rank approximations can be used within the preconditioner to approximate toroidal couplings between poloidal planes. That is, partitioning, ordering, and rank choices for the hierarchical factorization can be oriented so that higher rank approximations are used within poloidal planes (mimicking elements of the current block Jacobi scheme) while low rank approximations could be used to approximate coupling between poloidal planes to accelerate convergence over the block Jacobi scheme. A second possibility is a multigrid algorithm that only coarsens (i.e., semi-coarsens) in the toroidal direction and uses the block Jacobi/SuperLU_DIST scheme as a multigrid smoother. A typical multigrid preconditioner generates coarse approximations that are used to accelerate iterative solver convergence. Our aim is to construct an algebraic multigrid preconditioner [128,129] that requires little user input to build these coarse approximations. A number of semi-coarsening multigrid algorithms have been developed to address anisotropic problems on structured grids [130-132] and often leverage structure to develop grid transfers between the different resolution approximations using only matrix entries [133-135]. We intend to investigate multigrid algorithms where grid transfers can leverage structure in only one mesh direction along the lines of more recently proposed semi-coarsening multigrid methods [131,136,137]. A key challenge will be to adapt multigrid algorithms to the matrices produced by M3D-C1 and NIMROD which are discretizations of high order PDEs based on high order basis functions or FFTs. Here, ideas along the lines of [138] will be considered. The new multigrid methods will be implemented in MueLu [139] and will leverage a number of existing kernels.

2.5.2 Finite Element, Code Coupling, and Meshing Developments

The M3D-C1 code is already built on the Parallel Unstructured Mesh Infrastructure (PUMI)[140,141] developed through the SciDAC program and employs a set of specifically tailored mesh generation procedures [121,142] now being used by several fusion plasma simulation groups.

Equation forms and finite element discretizations: M3D-C1 uses three “annihilation projections” applied to the momentum equations. These projections separate the dynamics of the system along the three dominant waves of the system and help in designing optimal preconditioners. These projections can also be interpreted as a Petrov-Galerkin method applied to the original equations. In this context, the role of other Petrov-Galerkin methods, such as those derived from the variational multiscale method [143] as applied to the stream-function/potential representation of the resistive MHD equations used in M3D-C1 will be explored to determine whether they can further improve the numerical performance. Note that this is distinct from the VMS method applied to the resistive MHD equations [144] that does not employ the stream-function/potential representation which has some benefits for tokamak applications, although it requires higher order derivatives. In contrast to M3D-C1, NIMROD uses primitive MHD variables and employs a Fourier basis in the periodic direction and higher-order finite element basis in other directions. The application of the VMS method to further improve the stability of this choice of basis functions will also be explored.

The equations of hydrodynamics and resistive MHD yield discretized systems with poor numerical conditioning. One means of improving the numerical conditioning is to “regularize” the equations. The Leray and the Leray-alpha models are particularly interesting for use in deriving regularized MHD equations and will be considered [145].

During a disruption simulation it is necessary to preserve positivity of quantities such as pressure, which is not guaranteed in a typical discretization method. Techniques that enable this through simple variable transformations and the imposition of variational inequalities [146] will be explored.

The choice of basis functions can strongly influence the numerical solution properties (e.g., positivity) as well as numerical conditioning. Two aspects that become important are selection or construction of different basis functions, and implementation that seamlessly allows different choices of basis. A specific class of basis functions that is of interest includes Bernstein polynomials (in Bernstein–Bézier form) and those derived from them including rational forms with high-order continuity (C1 or more) including triangular elements. These polynomials possess attractive numerical properties (stable and efficient evaluation, pointwise positivity, etc), allow for stable transformations to other polynomials [147]. They readily fit into a FE code by exploiting generalized Bézier decomposition (e.g., see [148] and [149). This class of basis functions directly fits into the directional setup exploited in the M3D-C1 code based on triangular wedge elements that are a tensor product of a poloidal plane triangle times a toroidal direction function. Such a directional setup has been explored in anisotropic resolution with high-order continuity in the wall-normal direction for boundary layer flows [150].

Mesh adaptation and solution transfer: The effective control of the mesh discretization errors in tokamak plasma simulations requires the application of adaptive analysis methods. The components of an adaptive analysis past a fixed mesh analysis are estimation of the mesh discretization errors, determination of the adapted mesh distribution, adapting the mesh to the desired mesh distribution and transfer of solution fields onto the adapted mesh. Our previous efforts employed the PUMI parallel mesh adaptation procedures [140,141] in conjunction with a basic error indicator [151]. Proposed work in this area will focus on improved error estimation taking advantage of recent advances in error estimation for MHD simulations [152,153] and development of anisotropic mesh distributions[152]. Consideration will also be given to the potential application of goal oriented estimation for quantities of interest for plasma disruptions simulations including taking appropriate care to use the error estimates to predict the meshes adaptive mesh distribution [154]. Since the PUMI mesh adaption procedures employ generalized local mesh modification operations we can support highly efficient and accurate solution transfer [155] including consideration of conservation requirements [156]. To support the coupling of FronTier-MHD with M3D-C1 and NIMROD, the solution transfer procedures just outlined will be extended by the addition of parallel procedures to find candidate elements to search for containment of points for solution transfer.

Multi-Scale Coupling Algorithms: An important deliverable of this research is software for multiscale simulations of plasma disruption mitigation that couples the local FronTier-MHD model for a single pellet or SPI with the tokamak-scale codes NIMROD and M3D-C1. Such a two-way coupling presents several applied math, algorithmic, and software development challenges. We will develop solution transfer libraries for conservative remapping of solutions that reside in different spaces of physics states on overlapping geometric domains, meshed and discretized using Cartesian meshes and finite volumes / finite differences in FronTier and finite element meshes and discretization combined with spectral methods in global MHD codes. Another challenge is associated with solution and error estimators that will dynamically identify physical domains in global codes to be refined in a numerical resolution and physics model sense by the local ablation code. Special attention will be devoted to the study of stability and possible artifacts of coupled simulations well as the parallel load balance and performance of the coupled codes. This research will also leverage methods and tools developed by other SciDAC Centers.

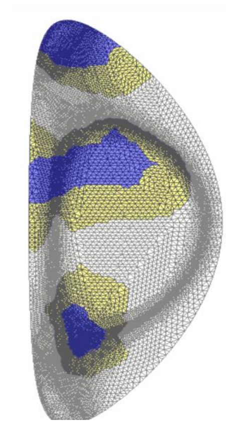

Figure 9: Three of sixteen partitions of a mesh with ten layers of ghosted entities on each partition.

Parallel support for particle based operations: To better capture the physics across the scales of importance to the plasma behavior, it is desirable to combine mesh-based PDE discretizations with particle (PIC) methods. The PUMI mesh infrastructure is being extended to effectively support PIC calculations on distributed unstructured meshes [36].

Since a PIC particle push requires no communication so long as both the starting and ending mesh entity are local, the first PUMI development was a multiple layer ghosting mechanism (see Figure 9). The second, which is ongoing, is a fast adjacency-based search with fast parametric inversion to find the element a particle is within. Since particles do not traverse many elements in one push operation, using mesh adjacencies will limit the number of elements that need to be checked through parametric inversion. Key to efficiency of this process will be memory layout and data. It has already been determined that replacing the current parametric inversion with one that does more floating point operations, but has better data access, is three times faster. Determining the optimal adjacencies and data layout is an area that must be addressed. The third PUMI development is fast operations to sum particle contributions in forming the right hand side vector as needed to solve for the updated fields needed for the next particle push. Based on limited profiling done to date, it is clear that node parallel performance and consideration of GPU acceleration are important areas to be addressed.

An area of future development is maintaining load balance. Since different mesh distributions are optimal for the particle push and PDE field solve load balance, one must either repartition the mesh between each of those steps, or determine a single partition that best meets the needs of both. Since the time required for complete repartitioning and associated data motion at the frequency of these steps is likely to be too large, methods to define a best overall single partitioning, including the potential for a hierarchic partition, will be developed. The fact that the particle distribution can change and the mesh can be adapted for discretization error control, the computational and communication loads will become imbalanced. Thus the application of occasional fast dynamic load balancing at both the level of the distribution of the mesh to nodes and to processes and/or threads on node is highly desirable. Fast dynamic load balancing methods that explicitly account for the multi-criteria associated with mesh-based particle methods can take full advantage of mesh adjacencies [157].

2.5.3 Performance Optimization

Over the next 5 years, DOE is expected to deploy 5-7 supercomputers in the 30-300 PF range. As traditional CPU architectures lack the energy efficiency required to keep the system power requirements under a reasonable 10-20MW limit, it is all but assured these systems will be built from either many core (Intel Knights Landing) or GPU-accelerated (NVIDIA) node architectures. One set of developments to more effectively use these node architectures has been array based mesh PUMI topology [158] and parallel mesh adaptation on many core [159] and GPU-accelerated [160] systems.

Each of the many core or GPU-accelerated node architectures in the upcoming systems is likely to include a hierarchical memory architecture that pairs high-bandwidth (but low capacity) memory (e.g. MCDRAM or HBM) with high-capacity (but low bandwidth) memory (e.g. DDR). Although these technologies existed in part or in whole in the previous generation of MIC (KNC) and GPU-accelerated supercomputers, recent advances in architecture and programming model have begun to address their productivity challenges. For example, both Knights Landing and Kepler/Pascal GPUs can now present their respective hierarchical memory architectures to the user as a single address space with the fast memory acting as a cache (direct-mapped, hardware cache on KNL and a runtime-managed page cache on GPUs). This eliminates the complexity of micromanaging data movement and data locality into scratchpads in favor of the well-established field of managing locality for a cache hierarchy.

We expect these new architectures to present a number of computational challenges for NIMROD and M3D-C1. The most pressing of these is expected to emerge in the sparse factorization and sparse triangular solve routines where endemic complex memory access patterns, irregular computational patterns, and complex data dependencies are antithetical to the trends in processor architecture (massive thread/data parallelism, wide SIMD, smaller on-chip caches, profound HBM cache effects, and limited total memory). Moreover, additional performance issues associated with manycore and GPU architectures may emerge in NIMROD and M3D-C1 particularly as massive thread- and GPU- parallelism demand routines that nominally consume 1-5% be efficiently threaded (Amdahl's law) [161].

For NIMROD and M3D-C1, we will initially target the manycore-based supercomputers at LBL and ANL (Cori and Aurora) which are based on the similar KNL and KNH processors [162,163]. To that end, we will perform detailed profiling of NIMROD and M3D-C1 on the KNL-based Cori as a function of MPI vs. OpenMP parallelism on a node to identify performance hotspots as well as any threading bottlenecks. For the identified hotspots, we will develop performance models to determine whether the performance limitations are implementation artifacts (performance is less than what the architecture is capable of) or algorithmic (performance is at the architecture limit and any improvements must be enabled via algorithmic changes) [164-166]. We will work with the SuperLU/STRUMPACK developers to ensure both triangular solve and factorization run efficiently given the massive thread parallelism, wide SIMD, and limited memory capacity of the Knights processors. Moreover, given Cori has 9K KNL nodes and Aurora will have 48K KNH nodes, we must strive to substantially improve the multi-node scalability of SuperLU and STRUMPACK.

Once the efficacy of using an alternate solver in NIMROD/M3D-C1 has been established, we will work with their developers to maximize performance on the Knights platforms. Finally, we will work with the SuperLU/STRUMPACK, M3D-C1, and NIMROD developers to ensure these applications run efficiently on the GPU-accelerated supercomputers at ORNL (Titan/Summit) [167,168]. Unfortunately, this is a potentially disruptive change as exploiting the most modern and forward looking features of the GPU ecosystem often requires extensive use of CUDA and/or C++.

2.5.4 Uncertainty Quantification

In simulating the onset of disruptions, their mitigation and their consequences for the vessel walls, there are a number of uncertainties, which result in uncertainties in the simulation predictions. It is important that uncertainties in simulation output be quantified for two reasons. First, in the proposed validation studies, these uncertainties determine both how much discrepancy between experiments and simulations is expected, and how strong any validation test is [169]. Second, when simulation predictions are to be used to inform design or operation decisions (e.g. for ITER), the uncertainties are needed to ensure appropriate margins and to identify when more information is needed [170]. For this reason we will pursue an uncertainty quantification effort as part of the proposed project, focused on three simulation objectives. These are described briefly below.

Disruption Onset: The prediction and identification of disruption onset are critical issues for the success of ITER and future tokamak fusion reactors. While the relationship between the crossing of a linear stability threshold and the onset of a disruption is complex due to the possibility of nonlinear saturation, the determination of linear stability is a crucial component of disruption prediction and identification. In real-time, identifying disruption onset is complicated by uncertainties in the experimental equilibrium reconstruction (e.g., uncertainty in the pressure and current profiles) as well as the details of the physics model implemented (e.g., details of the resistive wall surrounding the plasma). The effects of these uncertainties on the prediction of linear instabilities will be determined for given equilibria. Further, to support eventual real time disruption detection, techniques for generating inexpensive and effective surrogates will be sought.

Disruption Mitigation: Key metrics for the effectiveness of disruption mitigation strategies (DMS) include the radiation fraction, the poloidal and toroidal radiation peaking factors, and current quench time. The uncertainty in the prediction of these quantities will be determined in the simulation of candidate DMS. A DMS strategy of particular interest is shattered pellet injection (SPI). Simulations of SPI tests conducted in the DIII-D tokamak suffer from uncertainties in plasma parameters such as viscosity, resistivity and thermal conductivity, and from uncertainties in the post-shattering shard sizes and distributions. Sensitivities of the mitigation metrics to these uncertainties will be determined for use in validation assessments against measurements in DIII-D. Also of interest here are the uncertainties that arise in the modeling of the spectroscopic observables measured in the experiments given the thermodynamic and hydrodynamic outputs of the simulations.

Wall Forces and Currents: When a disruption results in the loss of vertical position control, currents induced in the vessel walls result in significant structural loads due to the Lorentz force. A computational model to predict these forces is being validated with controlled experiments on NSTX and DIII-D, but there are a number of uncertainties that affect validation comparisons. These include the uncertainties in the reconstructed initial equilibrium (from measurements), the modeling of transport in both the core plasma and the scrape-off layer, and the modeling of three-dimensional effects in two-dimensional simulations. The impact of these uncertainties on wall force predictions will be determined, and used in the validation of the models against experiments and when possible 3D simulations, and ultimately to inform the uncertainties in ITER predictions.

To conduct the uncertainty studies described above, the primary software tool will be the Dakota toolkit [171], which provides access to a wide range of UQ algorithms. The integration of UQ software tools with simulation codes such as M3D-C1 and NIMROD will be through the OMFIT [172] integrated modeling framework, which provides a standardized interface to run the code and can orchestrate post-processing analysis. One of the primary challenges arises because the computational models are relatively expensive to evaluate, so that the number of samples of uncertain inputs one can evaluate is limited. Several strategies to address this will be employed, including input space dimension reduction through sensitivity analysis and the development of model surrogates. Finally, the dimension reduction of the uncertain input space will identify those inputs that are most important to determine good uncertainty estimates for and possibly seek ways to reduce uncertainty in.

2.5.5 Simulation Data Management and Fault Tolerance

The application of validation and UQ processes requires the simulation results come with providence metadata indicating the versions of hardware and software, physical parameters set and the simulation control options used [173]. As indicated previously, CTTS plans to take full advantage of the One Modeling Framework for Integrated Tasks (OMFIT)[172,174], being developed as part of the Advanced Tokamak Modeling (AToM) SciDAC [175], to integrate M3D-C1 and NIMROD simulation capabilities and results as part of whole device simulations, and to link with experimental data from magnetically confined thermonuclear fusion experiments. Efforts will be carried out to define and support the simulation metadata needed and to have that information added to the M3D-C1 and NIMROD simulation results. In addition to OMFIT, consideration will be given to investigating planned additions to ADIOS for supporting simulation providence metadata. Further developments will require procedures to support data processing procedures which search and operate on simulation data in support of validation and UQ operations [176,177]. As is common with large-scale simulation codes, M3D-C1 and NIMROD already employ checkpoint restart methodologies to minimize the influence of systems faults for long running simulations. However, we plan to investigate if new ADIOS checkpointing capabilities and technologies such as burst buffers can improve the fault tolerance for our codes.

2.5.6 Software Engineering

M3D-C1 software repository will be updated to GitHub [178] with project wiki, forums and document storage. Developers will interact with users through GitHub's flexible issue tracking system and a discussion forum that will include wiki-based and regular electronic conferencing. Key development branches will be monitored by a code leadership team for each code that will ensure merges are executed on a regular basis consistent with the goals of that development group and the project needs. NIMROD will continue with its current svn-based repository system.

Although much of the development of the new physics capabilities will be by dedicated groups, there will be specific concentrated efforts in the form of agile scrums [179] that will bring a diverse set of developers together for a targeted development. The obvious activities well suited to a scrum are the introduction of new performance methodologies into portions of the code and when multiple components need to be properly integrated to address a multi-physics consideration.

A CTest/CDash [180] project will be created for each code with scripts that download and build the software, and execute compliance and regression tests. Because of the complexity of some expected regression tests, we anticipate having both daily and weekly regression tests in which the daily tests will be unit tests and tests on small data sets. The weekly test will include more extensive cases that that couple multiple components or simply must be larger to be meaningful. The regression tests will be run on multiple HPC systems covering those being used by the M3D-C1 and NIMROD user communities.

To support the user community [172], mentorship and training will be provided for new users. New users will be encouraged to submit bug reports, initiate issues and, as interest and expertise grows, join in code development. Because of the level of expertise required to effectively model complex tokamak plasma physics, specific care will be taken to interact with potential users to ensure they understand the workflows and associated usage cases.

2.6 Integration with WDM and Other Centers

The primary emphasis of this proposal is to improve the capabilities, efficiencies, and scalability of NIMROD and M3D-C1 and to validate these new capabilities over a range of disruptive phenomena. Developing a model of SPI with FronTier and coupling this with the MHD codes is an important new capability described in section 2.4.2 that will make these codes uniquely able to model mitigation experiments realistically. Section 2.1.3 discusses coupling to components developed by the SCREAM SciDAC for modeling runaway electron production and evolution during a disruption. We have also interfaced our codes to OMFIT [172] to ease comparison with experimental data and for performing scans as required by UQ procedures. The OMFIT coupling will also facilitate coupling NIMROD and M3D-C1 to a Whole Device Modeling (WDM) framework such as the AToM SciDAC (already accomplished). This also enables ready coupling to a boundary/materials Center, for example, to supply heat fluxes and receive impurity influxes during a disruption. Previous activities associated with the SWIM SciDAC have demonstrated the ability to couple an MHD code with an RF heating code [95,22] as a tool for modeling disruption prevention. In summary, the work described in this proposal should go far towards producing robust and validated MHD components that could be used within a WDM framework.

3. Timetable of Activities

Ideal MHD Driven Disruptions

Year 1	Develop criteria for when locally exceeding β-limit leads to a disruption
Year 2	Develop validated model that reproduces thermal quench in an ideal MHD disruption
Year 3	Develop validated model for current quench that reproduces current spike and decay times
Ys 4-5	Interface M3D-C1 and NIMROD with runaway electron model as developed by SCREAM

Vertical Displacement Events

Year 1	Benchmark NIMROD and M3D-C1 for axisymmetric VDE in toroidal geometry
	Parametric studies of influence of halo-region properties and compare with analysis
Year 2	Benchmark NIMROD and M3D-C1 for non-axisymmetric VDE in toroidal geometry
	Incorporate sheath effects in VDE computations
Year 3	Study effect of non-axisymmetric walls on VDE
Ys 4-5	Validation studies with DIII-D, NSTX, JET
	Study of wall forces in mitigated and unmitigated VDE

Resistive Wall Modes

Year 1	Benchmark NIMROD and M3D-C1 for linear RWM in toroidal geometry
Year 2	Explore effect of rotation and two-fluid effects on RWM stability
Year 3	Nonlinear studies of RWM --- How does RWM precipitate thermal quench?
Ys 4-5	Explore kinetic effects on stability with NIMROD, M3D-C1/DK4D: Compare with MARS-K
	Explore disruptions caused by energetic particles (fishbone modes) interacting with RWM

Neoclassical Tearing Modes

Year 1	Identify suitable for NTM/locked mode disruptions on DIII-D for modeling
	Implement Ramos-form of DKE closures into NIMROD and M3D-C1.

	Investigate Maxwell torques induced by error fields in the presence of tearing modes
Year 2	Benchmark M3D-C1 and NIMROD with DKE closure about fixed magnetic island geometry
	Work with Δ'>0 cases to produce a saturated TM as an initial state for DKE NTM calculations
	Use Δ'>0 case to study growth of non-rotating magnetic island in presence of a resistive wall
	Investigate resistive-wall torques induced by error fields in the presence of tearing modes
Year 3	Model NTM evolution using DKE closures inc. temp equilib. and perturbation to BS current
	Study side-band induced stochasticity and edge effects in island in presence of resistive wall
	Investigate NTV torques with DKE closures on the mode from field errors
Ys 4-5	Understand the locking of NTMs from NTV, field errors and the drag on the resistive wall
	Investigate hypotheses on how locked modes grow and cause disruptions

Disruption Mitigation by Shattered Pellets

Year 1	Construct SPI plume model and develop tracking algorithms
	Develop 3D local pellet ablation model for FronTier-MHD and perform single-pellet tests
	Perform SPI scoping and sensitivity studies using NIMROD with an existing analytic SPI model
	Implement full ionization/recombination/radiation model in M3D-C1
Year 2	Implement pellet debris plumes into FronTier-MHD and test tracking algorithms.
	Perform SPI simulations and validation tests using FronTier-MHD and DIII-D experimental data
	Develop analytic kinetic heat flow models for use with NIMROD and M3D-C1
	Complete SPI scoping studies using NIMROD and M3D-C1 with an existing analytic SPI model.
Year3	Develop algorithms for coupling of FronTier-MHD pellet ablation with NIMROD and M3D-C1
	Test multiscale coupling algorithms using 1D FronTier pellet code and 1D PRL MHD code
	Start multiscale integration of Frontier-MHD with NIMROD and M3D-C1
Ys 4-5	Perform test of multiscale coupling of FronTier-MHD pellet ablation w. NIMROD and M3D-C1.
	Studies of accuracy, convergence, and stability, conservative properties of coupling algorithms
	Validation tests using FronTier-MHD/NIMROD and M3D-C1 and DIII-D experimental data
	Perform extensive simulations of DIII-Data within UQ program
	Perform validation simulations using JET data as available
	Perform simulations of SPI applied to ITER

Computer Science and Applied Math

Year 1	Adapt Trilinos solver using SuperLU_DIST, speed SuperLU_DIST's triangular solve
	Determination of potential for variational multiscale methods, improved error estimator
	Port M3D-C1 and NIMROD KNL, analyze MPI vs OpenMP, develop performance models
	Formulate UQ problems, identifying and characterizing input uncertainties
	Put up-graded software engineering procedure in place
Year 2	Evaluate Trilinos/Kokkos package for M3D-C1 and NIMROD, analyze and develop strategies for equation orderings and partitionings for STRUMPACK hierarchical HSS solvers
	Upgraded parallel mesh with particles, solution transfer during adaptation with conservation
	Intra-node performance optimization of M3D-C1 and NIMROD on Cori/Aurora
	Integrated appropriate UQ tools with OMFIT and perform initial sensitivity analyses
	Adding metadata to simulation case studies as needed to support validation and UQ
Year3	Hybrid multigrid with toroidal semi-coarsening, investigate HSS hierarchical preconditioners
	Bézier basis functions, mesh adaptation on GPU's, mesh-to-mesh transfer to FronTier and other codes, positivity preservation
	Inter-node performance optimization of M3D-C1 and NIMROD on Cori/Aurora
	Formulation and testing of multi-fidelity surrogate UQ approaches

	Regression test upgrade to include two levels of tests
Ys 4-5	Algebraic multigrid methods for high order PDE systems for M3D-C1 and NIMROD, compare various solver/preconditioning options and optimize set for M3D-C1 and NIMROD
	Anisotropic mesh adaptation, dynamic load balancing for mesh and particles, scalability of adaptive mesh with particles
	Port of M3D-C1 and NIMROD on the GPU-accelerated Summit. Performance optimization on Cori/Aurora and Summit
	Application of UQ tools to validation analysis
	Software and simulation metadata support improvements for validation and user community

4. Project Management Plan

We will have "Code Development" leaders who are responsible for all code modifications and "Task Leaders" who are responsible for the milestones being met. S. Jardin is the overall lead PI and is responsible for the progress reports, meeting organization, and overall functioning of the project. On the physics side, the code development leaders are: C. Sovinec (NIMROD), N. Ferraro (M3D-C1), and R. Samulyak (FronTier). On the CS and Math side, M. Shephard is responsible for the overall implementation of software packages into the physics codes and for implementing the software engineering. We plan to hold meetings twice a year to report on progress and discuss issues that arise, and smaller video conferences more frequently. The task members and task leaders (underlined) are as follows:

Ideal MHD Driven disruption modeling: <u>Jardin</u>, Kruger, Zhu, postdoc
VDE and RWM: <u>Sovinec</u>, <u>Ferraro</u>, King, Strauss, Breslau, Lyons, Held, Zhu
NTM: <u>Kruger</u>, Held, King, Lyons, Ramos
Mitigation: <u>Lao</u>, Izzo, Parks, Samulyak, Jardin, postdoc
STRUMPACK based solvers: <u>Li</u>, Williams, King, Chen
Trilinos solvers: <u>Tuminaro</u>, Jardin, Chen, postdoc
Performance: <u>Williams</u>, Chen, Seol, Held, Ferraro, postdoc
Finite elements and Meshing: <u>Shephard</u>, Oberai, Sovinec, Jardin, postdoc
Data Management and Software Engineering: <u>Seol</u>, Chen, Kruger, Ferraro, Lyons, Samulyak
Uncertainty Quantification: <u>Moser</u>, Lyons, Ferraro, Sovinec, postdoc

5. Project Objectives

The high level objectives of our center are the following:

- Develop better understanding and improved predictive capability of when crossing a linear stability boundary will likely lead to a disruption.
- Provide a tool for calculating vessel forces, both axisymmetric and non-axisymmetric, for worst case VDE and other disruptions in tokamaks.
- Develop increased understanding and improved predictive capability for locked mode disruptions and how best to avoid them.
- Develop, verify, and validate 3D models for disruption mitigation by shattered pellets that can be used to design and optimize a system for future tokamaks.
- Increase the efficiency and scalability of the M3D-C1 and NIMROD codes so that they can model disruptions in ITER scale devices using the most powerful DOE HPC hardware.
- Demonstrate code-coupling via FronTier/(NIMROD,M3D-C1) coupling and by interfacing these codes within a WDM framework.