

Multi-region approach to free-boundary three-dimensional tokamak equilibria and resistive wall instabilities

Cite as: Phys. Plasmas **23**, 056114 (2016); <https://doi.org/10.1063/1.4948722>

Submitted: 09 February 2016 • Accepted: 19 April 2016 • Published Online: 20 May 2016

N. M. Ferraro,  S. C. Jardin,  L. L. Lao, et al.



View Online



Export Citation



CrossMark

ARTICLES YOU MAY BE INTERESTED IN

Axisymmetric simulations of vertical displacement events in tokamaks: A benchmark of M3D-C¹, NIMROD, and JOEUK

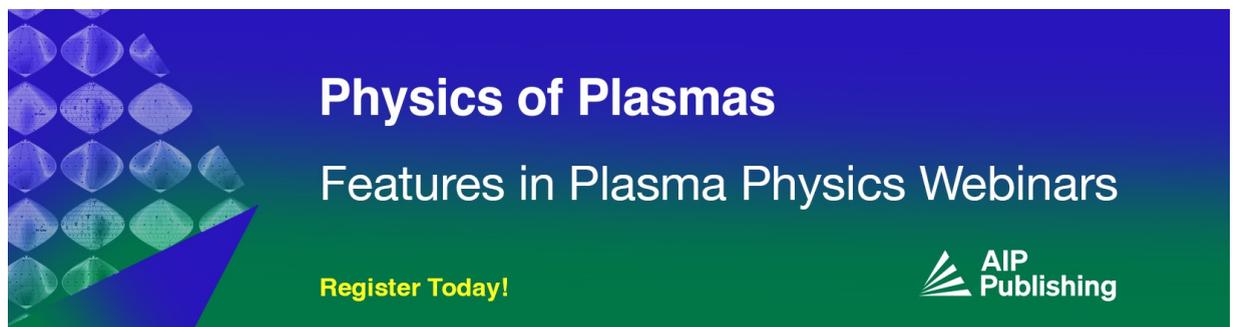
Physics of Plasmas **27**, 022505 (2020); <https://doi.org/10.1063/1.5127664>

Modelling of NSTX hot vertical displacement events using M3D-C1

Physics of Plasmas **25**, 056106 (2018); <https://doi.org/10.1063/1.5016348>

Magnetic flux pumping in 3D nonlinear magnetohydrodynamic simulations

Physics of Plasmas **24**, 102511 (2017); <https://doi.org/10.1063/1.4990704>



Physics of Plasmas
Features in Plasma Physics Webinars

Register Today!



Multi-region approach to free-boundary three-dimensional tokamak equilibria and resistive wall instabilities

N. M. Ferraro,^{1,a),b)} S. C. Jardin,² L. L. Lao,¹ M. S. Shephard,³ and F. Zhang³

¹General Atomics, La Jolla, California 92186, USA

²Princeton Plasma Physics Laboratory, Princeton, New Jersey 08543, USA

³Rensselaer Polytechnic Institute, Troy, New York 12180, USA

(Received 9 February 2016; accepted 19 April 2016; published online 20 May 2016)

Free-boundary 3D tokamak equilibria and resistive wall instabilities are calculated using a new resistive wall model in the two-fluid M3D-C1 code. In this model, the resistive wall and surrounding vacuum region are included within the computational domain. This implementation contrasts with the method typically used in fluid codes in which the resistive wall is treated as a boundary condition on the computational domain boundary and has the advantage of maintaining purely local coupling of mesh elements. This new capability is used to simulate perturbed, free-boundary non-axisymmetric equilibria; the linear evolution of resistive wall modes; and the linear and nonlinear evolution of axisymmetric vertical displacement events (VDEs). Calculated growth rates for a resistive wall mode with arbitrary wall thickness are shown to agree well with the analytic theory. Equilibrium and VDE calculations are performed in diverted tokamak geometry, at physically realistic values of dissipation, and with resistive walls of finite width. Simulations of a VDE disruption extend into the current-quench phase, in which the plasma becomes limited by the first wall, and strong currents are observed to flow in the wall, in the SOL, and from the plasma to the wall.

Published by AIP Publishing. [<http://dx.doi.org/10.1063/1.4948722>]

I. INTRODUCTION

The macroscopic evolution and stability of tokamak plasmas may be strongly influenced by the interaction between the plasma and surrounding conducting structures. Disruptive instabilities, which often involve substantial changes to the magnetic field at the location of the resistive wall, are particularly sensitive to the presence of these structures. Several fluid codes include the ability to model a thin wall of finite resistivity (typically called a “resistive wall”) surrounding the plasma, usually by shaping the computational domain boundary like the physical wall and applying appropriate boundary conditions. We explore an alternative method here, in which the resistive wall and the vacuum region external to the wall are included as spatially resolved regions within the computational domain. This method, implemented in the M3D-C1 code and described in detail in Section II, has some advantages over the boundary condition method, including the property that it maintains local coupling of mesh elements.

Here, we demonstrate the capability of this new model to address a variety of applications for which the resistive wall plays a crucial role. External kinks that are stable in the presence of a perfectly conducting wall may become unstable when the wall has finite resistivity. These resistive wall modes (RWMs) are an important instability which often set the upper limit for achievable plasma pressure in tokamaks. In Section III, we validate the implementation of the resistive

wall model in M3D-C1 by benchmarking against an analytic solution for a RWM. This benchmarking is done over a wide range of wall resistivities and widths, spanning both the inertial and resistive wall regimes of the instability.

Stationary plasma equilibria are generally insensitive to the presence of a resistive wall, since the lack of time variation implies the absence of inductively driven currents in the wall. (Time independence implies the absence of a toroidal loop voltage, which is driven by a time-varying flux.) However, the presence of a perfectly conducting boundary close to the plasma in a numerical model will prevent the relaxation of the plasma to its proper stationary equilibrium. We show how the resistive wall model here may be used to obtain accurate time-independent axisymmetric and non-axisymmetric equilibria with M3D-C1. The capability to calculate the non-axisymmetric equilibria is of considerable interest, because the application of non-axisymmetric fields to the plasma is observed to have a significant effect on the transport and stability properties of a tokamak plasma. These effects may include the suppression of edge-localized modes (ELMs), enhanced particle transport (“pump-out”), and torquing of the plasma. The phenomenology of these effects varies among tokamaks and among different parameter regimes within a tokamak. This complexity is due in part to the fact that the plasma generally responds strongly to the applied fields by amplifying and suppressing various spectral components of the magnetic field. Thus, the magnetic equilibrium in the presence of the plasma is generally very different from the sum of the equilibrium (axisymmetric) magnetic field and the applied non-axisymmetric fields. Progress toward a complete physical understanding and a predictive modeling capability of the observed effects of

Note: Paper GI2 5, Bull. Am. Phys. Soc. **60**, 110 (2015).

^{a)}Invited speaker.

^{b)}Present address: Princeton Plasma Physics Laboratory, Princeton, New Jersey 08543, USA. Electronic mail: nferraro@pppl.gov.

non-axisymmetric fields therefore requires a capability to calculate both the resulting non-axisymmetric magnetic geometry and also the changes in transport that are induced by the perturbed magnetic geometry. Both aspects have been the object of significant effort in recent years.

In Section IV, we describe modifications to the M3D-C1 Grad-Shafranov solver necessary to obtain axisymmetric equilibria suitable for linear stability calculations with M3D-C1. This is an important step in obtaining accurate stability and response predictions, since directly interpolating the Grad-Shafranov solution from separate equilibrium codes often leads to unacceptably large errors in the toroidal current density.

It is demonstrated in Section V that perturbed free-boundary non-axisymmetric equilibria in the presence of applied non-axisymmetric fields may be calculated using the resistive wall model. This is done by calculating the time-independent solution to the extended-MHD equations, linearized about the axisymmetric equilibrium obtained using the Grad-Shafranov solutions described in Section IV. These solutions differ from previously published solutions in that now, using the resistive wall model, the magnetic perturbations due to response currents in the plasma are allowed to extend beyond the wall. This leads to quantitative differences in the calculated internal response of the plasma. Validation of this capability against experimental data was done as part of a recent experimental campaign among US tokamaks and has been published separately.¹

Finally, the resistive wall model is applied to the linear, nonlinear, and early current-quench phases of an axisymmetric vertical displacement event (VDE). Vertical displacement events are axisymmetric, vertical motions of the plasma caused by a loss of vertical stability. These events result in the intersection of the plasma with the plasma-facing walls of the tokamak, the loss (“quench”) of the thermal energy and current in the plasma, and the uncontrolled termination (“disruption”) of the discharge.² The term “VDE” typically refers to the case when vertical stability is lost during normal operation of the tokamak, but vertical stability may also be lost as a secondary instability after a thermal quench resulting from a different event. In either case, the vertical motion of the plasma into the wall may result in large thermal loads on the wall and may drive currents in the surrounding conducting structures that may cause significant electromagnetic forces.³ In Section VI, the axisymmetric evolution of a VDE is simulated for a wide range of wall resistivities. VDE calculations of this type, in which the resistive wall is treated as a region within the computational domain, have been done previously using the TSC code.^{4,5} Unlike the TSC calculations, here we do not prescribe a temperature or width for the “Halo” region; rather, here the currents and temperature in this region evolve according to the same extended-MHD equations that govern the plasma evolution. It is shown that the maximum current in the wall, and hence the maximum force on the vessel, is largely independent of the wall resistivity. The timescale of the instability increases with the wall conductivity, however, and therefore the total impulse on the vessel also increases with the wall conductivity. A spike in

the toroidal current is also observed when the plasma comes into contact with the wall.

II. MODEL IMPLEMENTED IN M3D-C1

The implementation of the resistive wall model in M3D-C1 differs somewhat from the method used both in comparable parallel, nonlinear MHD codes, including M3D,⁶ JOEK,⁷ and NIMROD,^{8,9} and serial, linear MHD codes such as MARS.¹⁰ In those codes, the resistive wall is implemented as a boundary condition in which the component of the magnetic field tangential to the simulation domain boundary at each point on the boundary is related to the normal component of the field at all boundary points: $B_i^{tan} = M_{ij}B_j^{norm}$. The matrix M , usually calculated using a Green’s function approach, depends only the geometry and conductivity of external conducting structures and on any externally applied currents (e.g., from magnetic coils); as long as these are time-independent, M will also be constant. An advantage of this method is that M may be calculated using a separate code such as VACUUM,¹¹ STARWALL,⁷ or GRIN¹² that contains an accurate, and possibly very complicated, description of the external conducting structures and currents, sparing the MHD code from these details. A disadvantage for parallel codes is the non-locality of the boundary condition. Since M couples together all of the degrees of freedom on the boundary, this necessitates communication among non-adjacent domains in domain-decomposed parallelization schemes. This will adversely affect the scalability of both the construction and solution of matrix equations, especially if the boundary conditions are implemented implicitly.

M3D-C1 avoids the problems associated with non-locality by including the resistive wall and the surrounding vacuum region as distinct regions within its computational domain. This limits, to some extent, the complexity of the external conducting structures that can be modeled. In particular, the geometry of the resistive wall must be axisymmetric, because the mesh in M3D-C1 is composed of triangular prisms and cannot vary toroidally at the present time. Properties of the wall, such as the resistivity, may vary arbitrarily.

The magnetic field is explicitly split into the part due to currents in the plasma and wall, \vec{B}^p , and the part due to currents in external coils, \vec{B}^x

$$\vec{B} = \vec{B}^p + \vec{B}^x. \quad (1)$$

The finite element representation of the external field \vec{B}^x is calculated from a description of external coils using Green’s functions. Since the current density of the external field never enters into the dynamical equations, the current density in the code is taken to be $\vec{J} = \nabla \times \vec{B}^p / \mu_0$.

An example of a M3D-C1 mesh with a resistive wall region using the wall geometry of the DIII-D tokamak is shown in Figure 1. There are three regions: a “plasma” region, which includes the entire region enclosed by the resistive wall; a “resistive wall” region; and an “external vacuum” region, which extends from the outer boundary of the resistive wall to the computational domain boundary.

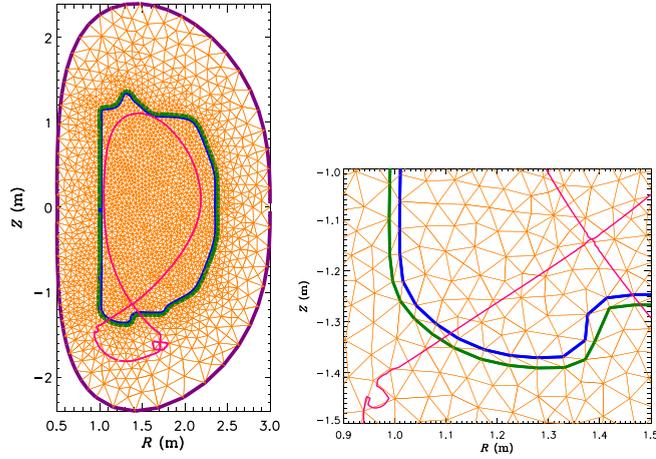


FIG. 1. A mesh in which the resistive wall region is an approximation of the DIII-D first wall is shown. The magenta line shows the magnetic separatrix. The simulation domain boundary, vacuum-wall boundary, and wall-plasma boundary are shown in purple, green, and blue, respectively. *Right*: a smaller area of the mesh near the lower divertor, showing the discretization of the resistive wall region.

A different set of equations is solved in each mesh region. In the plasma region, the full two-fluid model is implemented

$$\frac{\partial n_i}{\partial t} + \nabla \cdot (n_i \vec{v}) = 0, \quad (2)$$

$$m_i n_i \left[\frac{\partial \vec{v}}{\partial t} + \vec{v} \cdot \nabla \vec{v} \right] = \vec{J} \times \vec{B} - \nabla p - \nabla \cdot \Pi_i, \quad (3)$$

$$\begin{aligned} \frac{\partial p}{\partial t} + \vec{v} \cdot \nabla p + \Gamma p \nabla \cdot \vec{v} &= (\Gamma - 1)(\eta J^2 - \nabla \cdot \vec{q} - \Pi_i : \vec{v}) \\ &+ \frac{1}{n_e e} \vec{J} \cdot \left(\nabla p_e - \Gamma \frac{\nabla n_i}{n_i} p_e \right), \end{aligned} \quad (4)$$

$$\begin{aligned} \frac{\partial p_e}{\partial t} + \vec{v} \cdot \nabla p_e + \Gamma p_e \nabla \cdot \vec{v} &= (\Gamma - 1)(\eta J^2 - \nabla \cdot \vec{q}_e) \\ &+ \frac{1}{n_e e} \vec{J} \cdot \left(\nabla p_e - \Gamma \frac{\nabla n_i}{n_i} p_e \right), \end{aligned} \quad (5)$$

$$\frac{\partial \vec{B}^p}{\partial t} = -\nabla \times \vec{E}, \quad (6)$$

where

$$\Pi = -\mu(\nabla \vec{v} + \nabla \vec{v}^t) + \Pi_i^{\parallel} + \Pi_i^{\wedge}, \quad (7)$$

$$\vec{E} = \eta \vec{J} - \vec{v} \times \vec{B} + \frac{1}{n_e e} (\vec{J} \times \vec{B} - \nabla p_e), \quad (8)$$

$$\vec{q} = -\kappa \nabla (T_e + T_i), \quad (9)$$

$$\vec{q}_{e,i} = -\kappa \nabla T_e - \kappa_{\parallel} \frac{\vec{B} \vec{B}}{B^2} \cdot \nabla T_e. \quad (10)$$

The implementation of the parallel and ion gyroviscosity, Π_i^{\parallel} and Π_i^{\wedge} , is described in more detail in Ref. 13, and these terms are not included in the results presented here. Note that the open field-line region between the last-closed flux

surface (LCFS) and the resistive wall, where the plasma has low density and high resistivity, is also treated using these equations.

The resistive wall region only includes the resistive Faraday's law

$$\frac{\partial \vec{B}^p}{\partial t} = -\nabla \times (\eta_w \vec{J}^p). \quad (11)$$

In the external vacuum region, the only constraint is that the field remains current-free

$$\nabla \times \vec{B}^p = 0 \quad (12)$$

(M3D-C1 uses a vector potential formulation of the magnetic field, so $\nabla \cdot \vec{B} = 0$ is satisfied manifestly.)

Two types of time-steps are implemented in M3D-C1: a split, semi-implicit method, in which the fluid velocity and magnetic field are advanced separately, and a θ -implicit method in which all variables are advanced together.¹³ In both cases, the magnetic field in the wall and external vacuum region is advanced simultaneously with the magnetic field in the plasma; in fact, there is no distinction between these fields in the code.

Although the \vec{v} , p , p_e , and n_i fields extend into the resistive wall and external vacuum region, their values in these regions do not enter into the dynamical equations and therefore do not affect the solution. In M3D-C1, these fields are set to a constant value in these regions, and boundary conditions on these fields are set at the interface between the resistive wall region and the plasma region. In the calculations presented here, no-slip, no-normal-flow ($\vec{v} = 0$) conditions are enforced on the velocity, and uniform Dirichlet conditions are set on n_i , p , and p_e . There is no boundary condition on \vec{B}^p at this interface; rather, the boundary condition that \vec{B}^p is constant is enforced on the computational domain boundary (the boundary enclosing the external vacuum region).

III. RESISTIVE WALL MODE

In this section, we validate the implementation of the resistive wall model in M3D-C1 by comparing the calculated linear growth rate of a resistive wall mode in a straight cylinder against an analytic solution derived by Liu *et al.*¹⁴ The equilibrium is given by

$$J_z(r) = \begin{cases} J_0 & r < r_0 \\ 0 & r > r_0, \end{cases} \quad (13)$$

$$B_z(r) = B_0, \quad (14)$$

$$\rho(r) = \begin{cases} \rho_0 & r < r_0 \\ 0 & r > r_0, \end{cases} \quad (15)$$

with no equilibrium rotation, where z is the axial coordinate. The perturbed fields are taken to have the form

$$\delta \vec{B} = \nabla \psi \times \hat{z}, \quad (16)$$

$$\delta \vec{v} = \nabla \phi \times \hat{z}, \quad (17)$$

$$\psi = \tilde{\psi}(r) e^{i(m\theta - nz/R_0) + \gamma t}, \quad (18)$$

$$\phi = \tilde{\phi}(r)e^{i(m\theta - nz/R_0) + \gamma t}, \quad (19)$$

where $2\pi R_0$ is the axial period of the cylinder. The region $r < r_0$ is treated using the ideal MHD equations; $r > r_0$ is treated using the “vacuum-like” equation $\nabla_{\perp}^2 \psi = \nabla^2 \psi - \partial_z^2 \psi = 0$. A thin conformal wall having thickness d and resistivity η_w is present at $r = r_w$.

It is important to note here that the treatment of the vacuum region in this derivation is not strictly physical. Representing the vacuum field using Equation (16) implies a current density in the vacuum

$$\delta \vec{J} = \frac{1}{\mu_0} \nabla \times \delta \vec{B} = -\frac{1}{\mu_0} \left(\hat{z} \nabla_{\perp}^2 \psi + \frac{in}{R_0} \nabla_{\perp} \psi \right), \quad (20)$$

which requires $\psi = 0$ in a true vacuum. Instead, Liu *et al.* use the “vacuum-like” equation $\nabla_{\perp}^2 \psi = 0$ in the region $r > r_0$, which eliminates the axial current density but allows perpendicular (to z) current density in the “vacuum” region. In reality, these perpendicular currents would be eliminated by perturbations to the toroidal field, but those are not included explicitly in the model here.

The resulting resistive wall mode dispersion relation remains useful for benchmarking, but care must be taken to ensure that the numerical treatment of the vacuum region is consistent with this “vacuum-like” equation. This means that either variations to the toroidal field must be allowed by allowing a different representation of the magnetic field in the vacuum region or perpendicular currents must be allowed in the “vacuum” region. Since M3D-C1 represents the \vec{B} and \vec{v} fields using a flux / potential formulation, it is simple to restrict the form of the perturbations to the two-field model assumed in Equations (16) and (17);¹⁵ we choose to use this reduced model here to ensure that the analytic result is strictly applicable. M3D-C1 uses the same representation of the magnetic field throughout its domain, including in the “vacuum” regions, and therefore we must allow perpendicular currents in the “vacuum” regions in order to compare with the Liu *et al.* result. We do this by modifying the form of the resistive diffusion operator in M3D-C1 to neglect the damping of the perpendicular components of the current density. (This modified form is only used for benchmarking in this section and is not generally used in the code.) We note that the failure of a recent benchmarking activity¹⁶ using JOREK to find good quantitative agreement with the Liu *et al.* result might be due to JOREK representing the magnetic field between the plasma and resistive wall using Equation (16) and not appropriately modifying the resistivity there to allow perpendicular currents.

The analytic result of Liu *et al.*, obtained using the thin-wall approximation, is the following dispersion relation:

$$\frac{\nu}{m - nq_0} - \frac{1}{1 - \frac{\gamma\tau_w}{\gamma\tau_w + \mu} \left(\frac{r_0}{r_w}\right)^{2\mu}} = \frac{(\gamma\tau_A)^2}{2} \frac{q_0^2}{(m - nq_0)^2}, \quad (21)$$

where $\tau_w = \mu_0 r_w d / 2\eta_w$, $\tau_A = \sqrt{\mu_0 \rho_0} R_0 / B_0$, $q_0 = 2B_0 / (R_0 \mu_0 J_0)$, $\mu = |m|$, and $\nu = m / \mu$. In the ideal-wall limit, where $\gamma\tau_w \gg \mu$, Equation (21) reduces to

$$(\gamma\tau_A)^2 = \frac{2(m - nq_0)^2}{q_0^2} \left[\frac{\nu}{m - nq_0} - \frac{1}{1 - (r_0/r_w)^{2\mu}} \right]. \quad (22)$$

When the equilibrium is unstable in this limit, the instability is more properly termed an external kink. In the no-wall limit, where $\gamma\tau_w \ll \mu$, Equation (21) reduces to

$$(\gamma\tau_A)^2 = \frac{2}{q_0^2} (m - nq_0) [\nu - (m - nq_0)]. \quad (23)$$

In the limit in which the growth rate is sufficiently slow that inertia does not play a role, $\tau_w \gg \tau_A$, Equation (21) reduces to

$$\gamma\tau_w = -\mu \frac{1 - \nu(m - nq_0)}{1 - \nu(m - nq_0) - (r_0/r_w)^{2\mu}}. \quad (24)$$

M3D-C1 is run here using straight cylinder geometry, as is assumed by Liu *et al.*, rather than using high-aspect ratio toroidal geometry. We use the same equilibrium used in a recent JOREK benchmark¹⁶ with $R = 10$ m, $r_0 = 0.8825$ m, $B_0 = 1$ T, and $q_0 = 1.1$. The density varies from $\rho_0 = 10^{20}$ m⁻³ in the plasma to 10^{17} m⁻³ in the region between the plasma and the wall, following a tanh profile in the normalized poloidal flux Ψ centered at $\Psi = 1$ and having width $\Delta\Psi = 0.01$. Resistivity is taken to be discontinuous with $\eta = 2.745 \times 10^{-8}$ Ω m when $r < r_0$ and 2.745×10^2 Ω m when $r_0 < r < r_w$.

First, we consider the ideal wall limit. Here, no resistive wall model is present; the wall is modeled using perfectly conducting boundary conditions at the simulation domain boundary. The radius of the conducting wall is varied between $r_w = 1.4$ and $r_w = 2.5$ m. The M3D-C1 results, shown in Figure 2, are found to be in good agreement with the analytic result, Equation (22). The small difference between the calculated growth rate and the analytic solution is mainly due to the finite radial resolution (about 1 mm here) at the plasma-vacuum interface, where the analytic solution has a singular current layer.

Next, we consider the case with a resistive wall. Here, we include a resistive wall region from $r = r_w = 1.2$ m to $r = r_w + d = 1.22$ m. Since the ideal-wall external kink mode

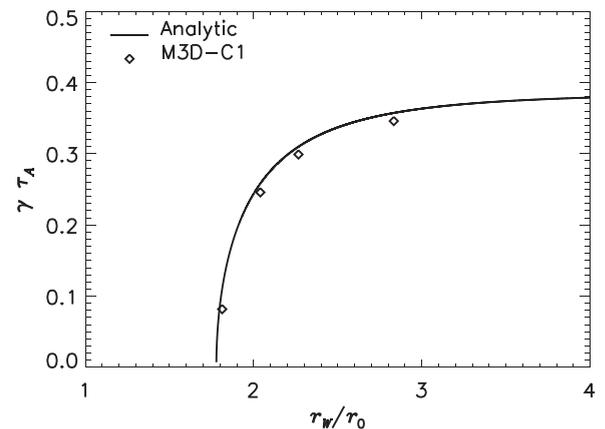


FIG. 2. The growth rate of the external kink as a function of the radius of a conformal conducting wall.

is stable at this value of r_w , the instability found here is a RWM. The computational domain boundary is placed at $r=4$, which is far enough out so as not to noticeably affect the growth rate. The growth rate is also calculated over a range of η_w . These growth rates are compared with the analytic thin-wall solution, Equation (21), in Figure 3. Again, excellent agreement is found through both the resistive ($\tau_A/\tau_w \ll 1$) and inertial ($\tau_A/\tau_w \gg 1$) regimes.

Finally, we demonstrate the capability of this method to resolve currents in the resistive wall by calculating the resistive wall mode growth rate for walls in which the thin wall limit is invalid. In Figure 4, the calculated growth rate, with d ranging from 2 cm to 50 cm, is compared both with the thin wall solution, Equation (21), and with the general analytic solution

$$\frac{\nu}{m - nq_0} - \frac{1}{1 - F\left(\frac{r_0}{r_w}\right)^{2\mu}} = \frac{(\gamma\tau_A)^2}{2} \frac{q_0^2}{(m - nq_0)^2}, \quad (25)$$

where

$$F = \frac{I_{\mu-1}(\rho_b)K_{\mu-1}(\rho_a) - I_{\mu-1}(\rho_a)K_{\mu-1}(\rho_b)}{I_{\mu-1}(\rho_b)K_{\mu+1}(\rho_a) - I_{\mu+1}(\rho_a)K_{\mu-1}(\rho_b)}. \quad (26)$$

Here, $\rho_a = \sqrt{2\gamma\tau_w r_w/d}$, $\rho_b = (1 + d/r_w)\rho_a$, and I and K represent modified Bessel functions of the first and second kinds. The derivation of Equation (25) is outlined in Appendix A. Note that the general result begins to deviate appreciably from the thin-wall result in this case when $d/r_w \gtrsim 0.1$. The M3D-C1 solution is in good agreement with the general result.

We note that in these calculations, packing of the mesh significantly improves the computational efficiency of these calculations. The meshes used in this section use strongly anisotropic mesh packing. Typical mesh elements have radial and poloidal extents of roughly 1 mm and 2 cm, respectively, near the plasma-vacuum boundary $r = r_0$ and in the wall. The mesh elements are significantly larger in the region between the plasma and the resistive wall and in the external vacuum region. A section of the mesh with $d = 5$ cm is shown in Figure 5.

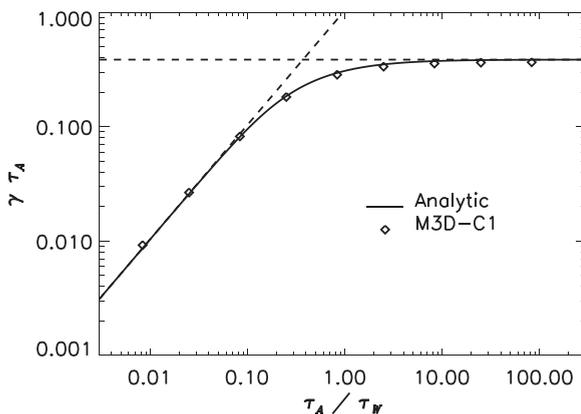


FIG. 3. The RWM growth rate calculated by M3D-C1 (symbols) as a function of the resistive diffusion time of the wall τ_w , compared with the thin-wall solution (solid line).

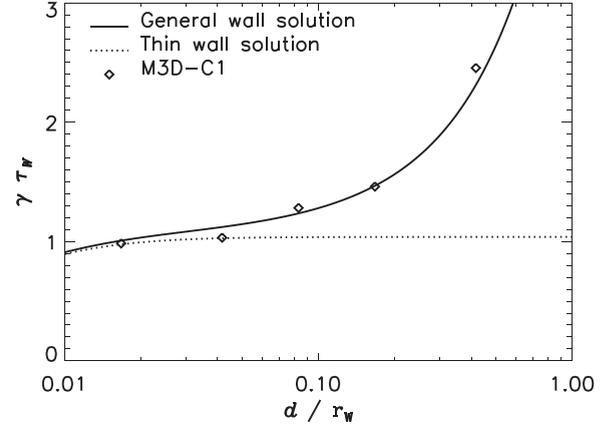


FIG. 4. The RWM growth rate calculated by M3D-C1 (symbols) as a function of the wall thickness d , compared with the general analytic solution (solid line) and the thin-wall solution (dotted line). Note that τ_w is proportional to d , so the growth rate in physical units is decreasing with d .

IV. FREE-BOUNDARY AXISYMMETRIC TOROIDAL EQUILIBRIA

MHD codes such as M3D-C1 often start from an axisymmetric initial condition—which may be a reconstruction of an experimental discharge—either as the initial value for a non-linear calculation or as the equilibrium for an eigenfunction or linear response calculation. In the case of a static, axisymmetric plasma, Equation (3) reduces in toroidal (R, φ, Z) coordinates to the well-known Grad-Shafranov equation

$$\nabla \cdot \left(\frac{\nabla \psi}{R^2} \right) = -\mu_0 p'(\psi) - \frac{1}{R^2} F F'(\psi), \quad (27)$$

where $\vec{B} = \nabla \psi \times \nabla \varphi + F \nabla \varphi$.

Many codes are available for solving this equation for ψ given $p(\psi)$ and $F(\psi)$ profiles and boundary conditions on ψ , which may be obtained through experimental measurements. Unfortunately, three issues prevent M3D-C1 from using the solutions from these codes. The first is that the domain of most of these Grad-Shafranov codes often does not extend far enough outside of the plasma to be useful for M3D-C1, which requires the solution all the way to the outer domain boundary, which may be far outside the physical vacuum vessel. Second, even those that could be extended to include the full M3D-C1 domain (EFIT,¹⁷ for example) do not output the field from the plasma separately from the field from the coils; this distinction is necessary in the M3D-C1

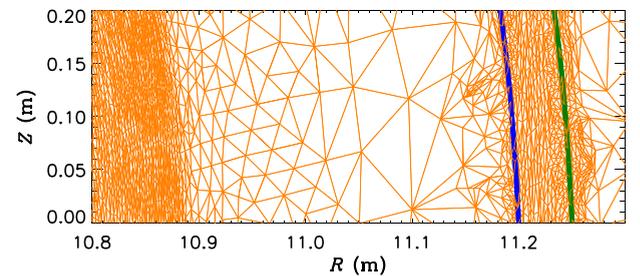


FIG. 5. A section of the mesh used for the resistive wall mode calculations with $d = 5$ cm in M3D-C1. The center of the plasma region is at $(R, Z) = (10, 0)$ m, the plasma-vacuum interface is at $r_0 = 0.8825$ m, and the resistive wall extends from $r = 1.2$ to 1.25 m.

formulation. The third issue is that a very high degree of accuracy in the solution is often required in the Grad-Shafranov solution for well-converged linear eigenfunctions to be obtained, especially in diverted, high-confinement mode (H-mode) plasmas.^{18,19} Therefore, it is preferable to obtain the weak solution to the Grad-Shafranov equation on the M3D-C1 finite element basis rather than to interpolate the result from a separate code that uses a different discretization scheme.

For these reasons, a Grad-Shafranov solver is implemented as part of M3D-C1.²⁰ $p(\psi)$ and $F(\psi)$ profiles may be read from common equilibrium file formats (such as EFIT's "eqdsk" format). This solver has been updated to allow solutions in the case where the poloidal field coils are inside the computational domain, by separating the field due to the coils from the field due to the plasma, as described in Section II. The solver uses a Picard iteration scheme, which can be unstable when the stabilizing wall is moved far from the plasma.¹⁷ Therefore, a proportional feedback controller on the magnetic axis is added to some of the coil currents to ensure vertical stability during the Grad-Shafranov solve. (This feedback is not active during the MHD evolution.)

One practical consideration that arises when using p profiles from typical equilibrium formats is that these profiles typically do not extend beyond the last closed flux surface (LCFS). One straightforward way to handle this in M3D-C1 is to treat the equilibrium density and temperature as uniform outside the last closed flux surface. However, this choice conflicts with the reality of the experiment, in which the density and temperature continue to decrease beyond the LCFS. In order to allow for such profiles in M3D-C1, the option to allow pressure gradients beyond the separatrix is also implemented. In practice, the following is typically done: the density n and electron temperature T_e profiles as a function of ψ are provided from experimental data, using an equilibrium reconstruction to map from real space to ψ . M3D-C1 then determines the pressure outside the LCFS by assuming that the ion temperature T_i outside the LCFS, which is often poorly constrained by empirical data, is uniform: $p(\psi) = n(\psi)[T_e(\psi) + T_i^{SOL}]$. (In contrast, inside the LCFS T_i is obtained from the given p , T_e , and n_e profiles.) This pressure profile is then used in the M3D-C1 Grad-Shafranov solver. The advantage of this method is a more realistic treatment of the region outside the LCFS; the disadvantage is that this sometimes introduces noticeable differences between M3D-C1's equilibrium geometry and the geometry of the source solution due to the change in force balance at the edge.

An example of an equilibrium solution from the M3D-C1 Grad-Shafranov solver for a DIII-D discharge is shown in Figure 6. The poloidal field coil ("F-coil") currents and the p' and FF' profiles have been taken from an EFIT reconstruction. From the figure, it can be seen that the poloidal field coils are in the external vacuum region of the computational domain.

V. FREE-BOUNDARY NON-AXISYMMETRIC TOROIDAL EQUILIBRIA

Here, we demonstrate the capability to calculate the perturbed magnetic geometry in the presence of externally

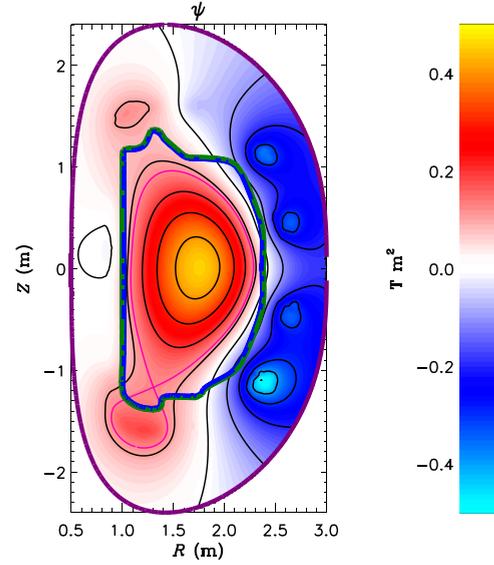


FIG. 6. The total poloidal flux per radian in the equilibrium solution calculated by the M3D-C1 Grad-Shafranov solver for a DIII-D discharge.

applied magnetic fields using the new resistive wall model in M3D-C1. This capability is unique in that it simultaneously allows: a free-boundary response; a two-fluid model; and a diverted magnetic geometry in which plasma current and pressure may extend up to, or even beyond, the LCFS.

M3D-C1 calculations have previously been compared with measurements of non-axisymmetric response in the edge of an H-mode discharges in the DIII-D tokamak in which perturbations having toroidal mode number $n=3$ were applied using the DIII-D I-coils.^{21–23} These calculations were done using a superconducting boundary roughly approximating the DIII-D first wall. At this boundary, the perturbed magnetic field was initialized to be the value due to the perturbing coils only and held fixed. Thus, the magnetic field due to response currents in the plasma was constrained to vanish at the boundary. This boundary condition is not generally appropriate in these experiments, in which the perturbing field is oscillated at 1–10 Hz, which is much

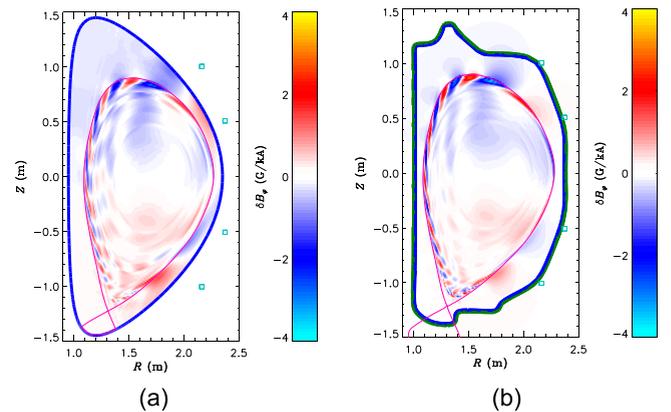


FIG. 7. The total perturbed toroidal field δB_ϕ at $\phi = 0$, in the simulation with a close conducting wall (left) and with a resistive wall (right). The cyan squares indicate the locations of the I-coils, which are applying an even-parity $n=3$ perturbation. The toroidal field from the I-coils at this toroidal angle is zero, so the depicted perturbed fields are entirely due to response currents in the plasma.

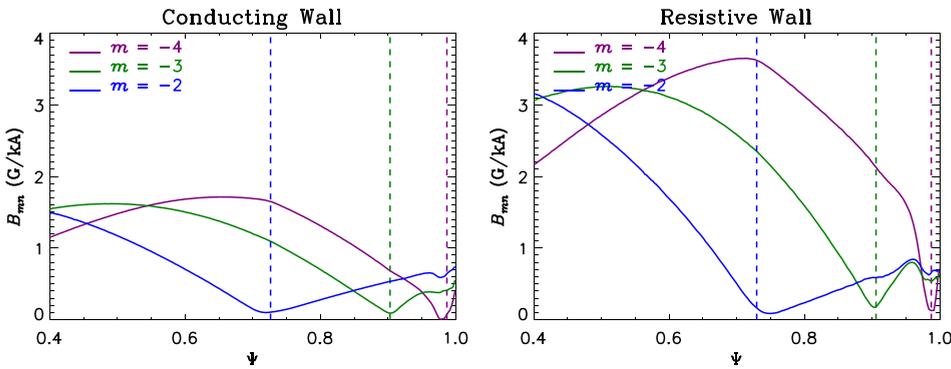


FIG. 8. The spectrum of the perturbed radial field in the $n=1$ plasma response in a conducting-wall calculation (left) and a resistive-wall calculation (right).

lower than the frequency associated with the resistive wall diffusion time ($1/2\pi\tau_W \sim 50$ Hz in DIII-D).

Here, we employ the resistive wall implementation to consider the opposite extreme: a time-independent calculation where the fields from the plasma may freely penetrate the wall. In this limit, the resistive wall plays essentially no physical role since eddy currents will not be present. The presence of the wall is numerically useful here, however, both in that it reduces the region in which the full extended-MHD system is solved, and it avoids potential numerical issues associated with the direct interaction of the plasma with the perturbing coils, which are just outside the resistive wall in this case. For these calculations, the mesh is packed anisotropically, with elements in the pedestal region having radial dimensions of a few millimeters. Spitzer resistivity is used, and the equilibrium T_e , n_e , and $E \times B$ rotation frequency are taken to match the experimental profiles.

Four non-axisymmetric equilibria are calculated here: two are calculated using a close conducting wall, as was done in the previous calculations, and two were done using a resistive wall in order to approximate the free-boundary response. In both cases, the $n=1$ and $n=3$ response to even-parity (up-down symmetric) fields from the DIII-D I-coils are calculated. The perturbed toroidal field in the $n=3$ calculations is shown in Figure 7. It can be seen that the perturbed magnetic field does not penetrate the conducting wall but freely penetrates the resistive wall. The penetration of the field allows direct comparison with experimental magnetic probe measurements which are placed on the wall. Such a comparison was done as part of the 2014 Fusion Energy Science Joint Research Target for several codes, including M3D-C1 using this new capability. The results of this validation exercise, which generally found good

agreement among codes and between the codes and data, have been published separately.¹

The perturbed radial field can be decomposed into poloidal Fourier components using straight field-line magnetic coordinates (ψ, θ, φ)

$$\delta B_{mn}(\psi) = \frac{(2\pi)^2}{A} \iint \frac{\delta \vec{B} \cdot \nabla \psi}{\vec{B} \cdot \nabla \theta} e^{i(m\theta - n\varphi)} d\theta d\varphi. \quad (28)$$

This representation is useful because it clarifies the distinction between two types of response: “kink” response, which is a distortion of the plasma that preserves the topology of the magnetic field and is the only response permitted in the ideal-MHD model; and “tearing” response, which implies the formation of magnetic islands. (Note that “kink” response is not necessarily equivalent to a kink mode; interchange modes will also imply a kink response.) The magnitude of B_{mn} is plotted as a function of the normalized poloidal flux in Figures 8 (for $n=1$) and 9 (for $n=3$). In the case of the $n=1$ response, the B_{mn} spectrum shows a global response, with significant perturbations over much of the plasma radius. In contrast, the $n=3$ response is much more highly localized to the edge, which is indicative of a driven peeling-ballooning mode. For both $n=1$ and $n=3$, the magnitude of the perturbed fields is generally greater in the resistive-wall case than the conducting-wall case; indeed, for the $n=1$, the resistive-wall response is roughly twice that of the conducting-wall case. This behavior can be understood through the well-known effect of wall stabilization. When the close conducting wall is removed, the plasma moves closer to marginal stability and will be more strongly excited by a given perturbation. This effect is stronger for $n=1$ fields, which fall off less strongly with distance than $n=3$ fields, and therefore is more strongly influenced by the presence of the wall.

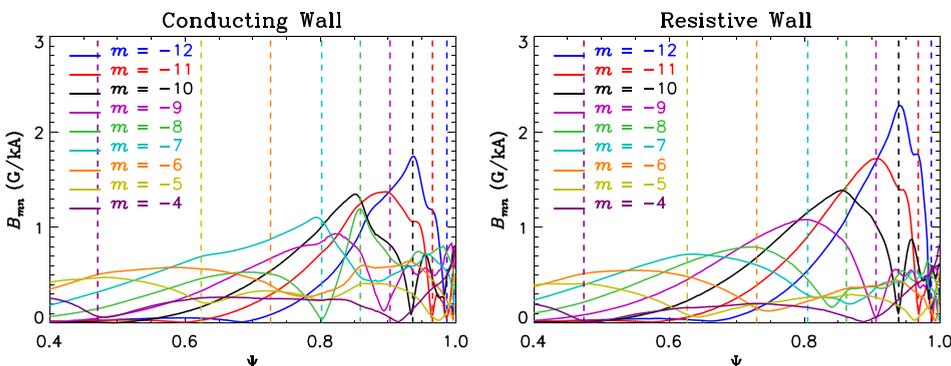


FIG. 9. The spectrum of the perturbed radial field in the $n=3$ plasma response in a conducting-wall calculation (left) and a resistive-wall calculation (right).

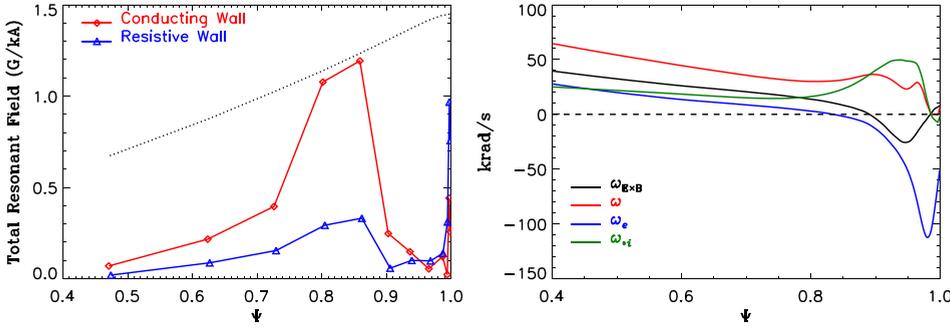


FIG. 10. *Left*: The total resonant field at each mode-rational surface in the $n=3$ response from $m=4$ at $\Psi \approx 0.47$ to $m=12$ at $\Psi \approx 1$ for both the conducting wall (red diamonds) and resistive wall (blue triangles) models. The dotted line shows the value in the absence of the plasma response (i.e., the “vacuum” field). *Right*: The electric drift ($\omega_{E \times B}$), ion rotation (ω), electron rotation (ω_e), and ion diamagnetic (ω_{*i}) frequencies.

The resonant field components of B_{mn} , which are the components at each mode-rational surface for which $m = nq$, imply the presence of magnetic islands and are therefore a measure of the tearing response of the plasma. In Figures 8 and 9, the mode-rational surface associated with each poloidal mode number m is indicated with a vertical dashed line. The magnitude of the resonant field components in the $n=3$ response calculations is shown in Figure 10. The resonant components of the total field are generally found to be smaller than those of the applied field alone (indicated by the dotted line in Figure 10), which implies that the plasma is acting to exclude (“screen”) magnetic islands. The resonant fields are generally largest in the range $\Psi = 0.8$ – 0.85 , which is where the electron rotation in this equilibrium passes through zero, as shown in Fig. 10. This is consistent with the previous studies using a two-fluid model in which it was shown that the tearing response is maximum where the electron rotation is small.²¹ It is also qualitatively consistent with theory in simplified geometry predicting that electron rotation damps tearing modes.²⁴ Interestingly, the screening tends to be weaker in the conducting-wall case.

VI. VERTICAL DISPLACEMENT EVENTS

Both linear and nonlinear, axisymmetric simulations are conducted here with M3D-C1, using an EFIT reconstruction of a vertically unstable DIII-D discharge as initial conditions. In this discharge, a thermal quench was initiated intentionally by injecting a “killer” pellet of neon into the plasma, dramatically raising the radiated power. Vertical stability control of the plasma was lost shortly (~ 2 ms) thereafter, resulting in a current quench and termination of the discharge. The objective of these simulations is not to validate the model through quantitative comparison with experimental measurements, but simply to explore the qualitative dynamics of the vertical instability, wall currents, and subsequent current quench as a function of the resistivity of the wall, η_w . For simplicity, several of the modeling choices used here deviate from the experiment; in particular, the resistive wall here takes the shape of the DIII-D first wall in 2015 (which differs from the wall that was present when the experiment was conducted), and the highly conducting vacuum vessel, which lies several centimeters outside of the first wall, is not included in the model. Instead, the first wall is taken to be conducting structure setting the flux diffusion timescale. Furthermore, these calculations are done at fairly low resolution, with a typical mesh element ~ 10 cm² in the plasma region, in order to complete a large number of nonlinear calculations over a wide range of parameters.

The pressure profile in the reconstructed equilibrium is centrally peaked, with a central electron temperature $T_e \approx 2.4$ keV. To simulate the radiative thermal collapse, an anomalously large perpendicular thermal conductivity is included, with $\kappa \approx 1.54 \times 10^{22}$ m⁻¹ s⁻¹. Given the electron density profile, $\chi_{\perp} = \kappa/n_e$ ranges from roughly 100 m²/s in the core to 800 m²/s in the edge. This results in a collapse of the temperature profile on a timescale of roughly 300 μ s. The resistivity profile is assumed to follow the Spitzer form and evolves with T_e . The reason for including the thermal collapse in these simulations is to observe, in the nonlinear calculations, the redistribution of the current as the plasma cools and the resistivity changes; quantifying other important effects, such as the formation of runaway electrons and the distribution and magnitude of the radiated power, would require a more sophisticated model of the thermal collapse and is not considered here.

A. Linear results

Axisymmetric, linear simulations were run for a large range of η_w , from $\eta_w = 1.9 \times 10^{-8}$ to 19Ω m. (For reference, the resistivity of copper at 20 °C is roughly $1.7 \times 10^{-8} \Omega$ m.) These calculations were initialized with small random perturbations imposed on the equilibrium, and the least-stable mode was obtained by integrating in time until the growth rate ($\gamma = \partial_t \ln K / 2$, where K is the kinetic energy) became independent of time. Only $n=0$ (axisymmetric)

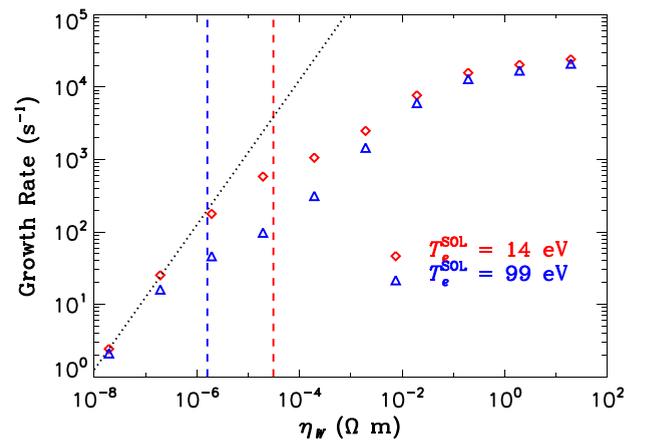


FIG. 11. The linear growth rate γ of the least-stable $n=0$ eigenmode, as a function of η_w . Growth rates were calculated using $T_e^{SOL} = 14$ eV (red diamonds) and 99 eV (blue triangles). The resistivity of the open field line region is indicated by the vertical dashed lines and is consistent with T_e^{SOL} . The dotted black line shows $\gamma \propto \eta_w$.

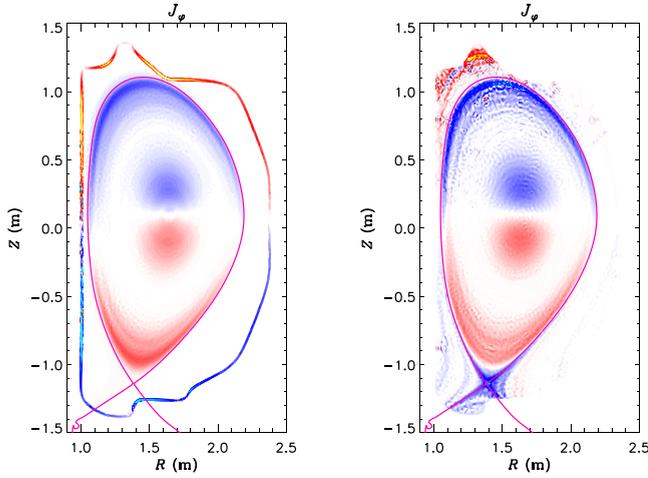


FIG. 12. The perturbed toroidal current density in the fastest-growing linear $n=0$ eigenmode, for (left) $\eta_w = 1.9 \times 10^{-8} \Omega \text{ m}$, in which the response currents are driven in the wall, and (right) $\eta_w = 19 \Omega \text{ m}$, in which the response currents form preferentially in the open field-line region of the plasma. In both cases, $\eta^{SOL} \approx 3.1 \times 10^{-5} \Omega \text{ m}$.

perturbations are considered here. The linear growth rate of the least-stable $n=0$ mode is shown as a function of η_w in Figure 11 for two different choices of the scrape-off layer (SOL) temperature T_e^{SOL} : 14 eV and 99 eV. In both cases, the equilibrium temperature of the entire open field-line region is uniform at this value, and the resistivity of this region, η^{SOL} , is the Spitzer value for the given temperature.

As in the simple resistive wall mode considered in Section III, the linear VDE is found to transition between a resistive-wall regime at low η_w to an inertial regime at high η_w . The growth rate of the VDE is shown in Figure 11 as a function of η_w . In contrast to the simple resistive wall mode, these VDE calculations also exhibit an intermediate resistive-SOL regime, in which strong response currents are driven in the SOL. This occurs in the intermediate regime where $\eta_w > \eta^{SOL}$ and $\tau_w > \tau_A$. The value of η^{SOL} is indicated by the vertical dashed lines in Figure 11. This intermediate

regime is unlikely to be physically important, since η^{SOL} is likely to be much greater than η_w in typical devices.

The perturbed toroidal current density in the eigenmode in the resistive-wall and inertial asymptotic regimes ($\eta_w \ll \eta^{SOL}$ and $\eta_w \gg \eta^{SOL}$) is shown in Figure 12. Since this figure represents a linear eigenmode, the amplitude is arbitrary. Note that the eigenmode is nearly identical inside the separatrix in both cases, but the response currents outside the separatrix depend strongly on η_w . (The graininess of the solution of the large η_w case is a consequence of the low resolution, together with the fact that no artificial method for smoothing the current density, like hyper-resistivity, has been employed here. The current density is the second derivative of the vector potential, which is the field being evolved in M3D-C1,²⁵ and is therefore not guaranteed to be continuous across element boundaries.) In the inertial regime, the response currents flow primarily on the surface of the plasma and are in the counter- I_P direction on the surface that is moving towards the wall. These counter- I_P currents have been termed ‘‘Hiro’’ currents in some instances.²⁶ Taking red to indicate positive (co- I_P) current perturbation, the eigenfunction in Figure 12 would be consistent with a downward-moving plasma. In the resistive-wall regime, these induced currents are driven primarily in the wall.

B. Nonlinear results

Axisymmetric, nonlinear simulations were run with $\eta_w = 1.94 \times 10^{-2}$, 1.94×10^{-3} , 1.94×10^{-4} , 1.94×10^{-5} , and $1.94 \times 10^{-6} \Omega \text{ m}$. In most of the calculations, the temperature at the wall was held fixed at approximately 99 eV. In order to gauge the importance of the edge temperature and thermal conductivity on the vertical instability and current quench, an additional calculation was done with an edge temperature of 65 eV and with χ_\perp reduced a factor of 10 below the baseline case.

In these simulations, the plasma is observed to move downward, as shown in Figure 13. The vertical displacement proceeds on a timescale which depends on the resistivity of the wall, η_w . The time-evolution of the vertical position of

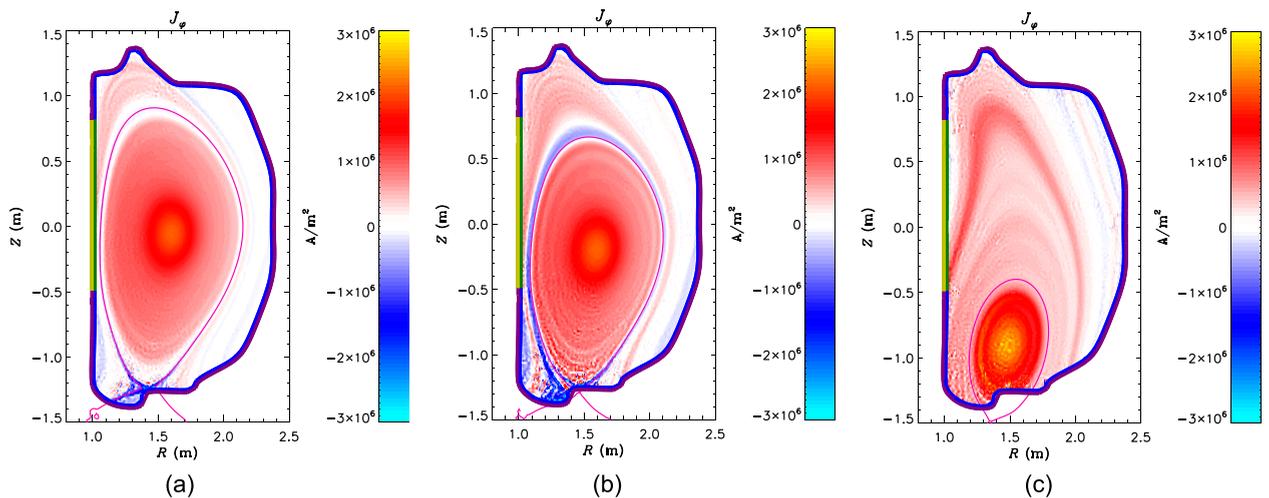


FIG. 13. The toroidal current density in a nonlinear axisymmetric VDE simulation with $\eta_w = 1.94 \times 10^{-3} \Omega \text{ m}$ at $t = 2.27, 2.59,$ and 2.92 ms .

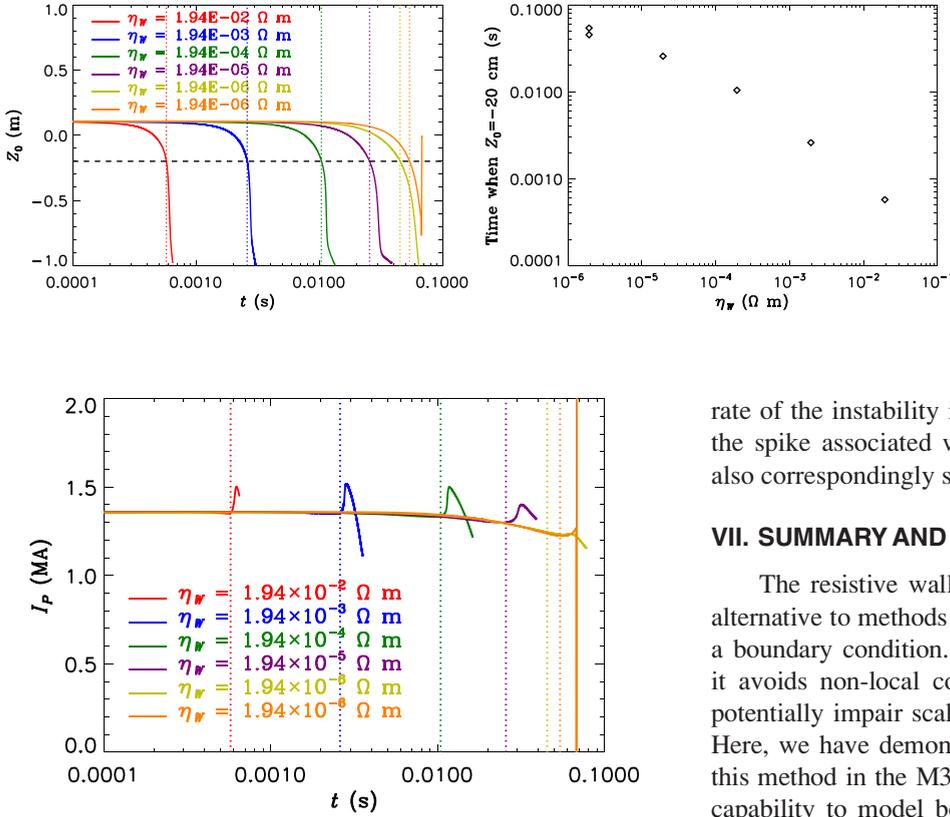


FIG. 15. The toroidal current in the plasma region as a function of time, for different values of η_w . The vertical lines indicate the times at which $Z_0 = -20$ cm.

the magnetic axis, Z_0 , is shown for all of these cases in Figure 14. Although the growth rate of the instability for these values of η_w is roughly linear with η_w^{-1} , the timescale of the macroscopic displacement of the plasma, as measured by the time it takes for the magnetic axis to reach $Z_0 = -20$ cm (roughly 25 cm below its original location), scales more closely with $\eta_w^{-1/2}$. This timescale is evidently not strongly dependent on η^{SOL} ; the timescale between the $T^{SOL} = 65$ eV (orange in Figure 14) and $T^{SOL} = 99$ eV (yellow) cases is less than 20%, despite nearly a factor-of-two difference in η^{SOL} .

The total toroidal current enclosed by the wall (not including the toroidal current in the wall) is shown as a function of time in Figure 15. Interestingly, a spike in the toroidal current is observed almost immediately after the magnetic axis reaches $Z_0 = -20$ cm, which is very close to the time that the plasma first contacts the lower divertor in all cases. This spike is reminiscent of the spikes in the plasma current observed during disruptions in experiments;²⁷ this is discussed further in Section VII. The spike in these simulations is due to the elimination of the counter- I_P response currents on the surface of the plasma and in the wall when the plasma comes into contact with the wall. Physically, this is due to the quenching of the counter- I_P voltage on the leading edge of the plasma when the wall stops the vertical motion of the plasma edge. For the parameters considered here, the toroidal response currents in the plasma edge are generally dominant over those in the wall. At higher wall conductivity, the

rate of the instability is slower, and the surface currents and the spike associated with the plasma touching the wall are also correspondingly smaller.

VII. SUMMARY AND DISCUSSION

The resistive wall model presented here is an attractive alternative to methods in which the resistive wall is treated as a boundary condition. One of its primary advantages is that it avoids non-local coupling of mesh elements, which can potentially impair scalability, especially in implicit methods. Here, we have demonstrated a successful implementation of this method in the M3D-C1 code, which now has the unique capability to model both linear and nonlinear resistive wall instabilities using a two-fluid model and with resistive walls of arbitrary thickness. The capability to model thick resistive walls is necessary for modeling ITER.

Good agreement with the thin-wall analytic theory of Liu *et al.*¹⁴ is found for a resistive wall mode in straight cylindrical geometry, for wall resistivities ranging from the resistive-wall to inertial limits. Good agreement is also found for walls of varying thickness by extending Liu *et al.*'s theory to allow for walls of arbitrary thickness. It is noted that if the region between the plasma and the wall is treated resistively (as opposed to as a true vacuum)—as it is in most extended-MHD codes—and if the perturbed magnetic field in that region is represented as a perturbed poloidal flux as in Liu *et al.*, then a modified form of the resistivity must be used in order to recover the result of Liu *et al.* This is because the reduced model of the field does not permit a nontrivial current-free solution. It is possible that this detail is responsible for the discrepancy in a previous comparison¹⁶ between JOEYK and the Liu *et al.* result. In MHD stability codes in which this region is a true vacuum, this issue does not arise; the perturbed magnetic field there will simply deviate from the reduced form by additionally perturbing the toroidal field in order to maintain a current-free solution.

The resistive wall model in M3D-C1 also facilitates the calculation of axisymmetric and perturbed non-axisymmetric free-boundary equilibria. Although a superconducting domain boundary is present, this boundary can be moved arbitrarily far from the plasma, in which limit the free-boundary solution is obtained. For truly time-independent solutions, as are presented here, eddy currents are not present in the wall, and a resistive wall model is not needed in principle. However, sectioning the mesh into a plasma region and vacuum region does have the advantages of reducing the extent

FIG. 14. *Left*: the vertical position of the magnetic axis, Z_0 as a function of time, for various values of η_w . The final case (orange) is the “reduced transport” case. Vertical lines indicate the time at which $Z_0 = -20$ cm. *Right*: The time at which $Z_0 = -20$ cm, as a function of η_w . The uppermost data point is the “reduced transport” case.

of the domain in which the full extended-MHD equations are solved, as well as obviating any potential issues arriving from having a numerical plasma at the location of magnetic coils. Furthermore, a simple modification of the linearized equations will allow solutions given perturbations with a non-zero frequency, in which eddy currents will be present and may play an important role.

The free-boundary perturbed non-axisymmetric equilibria calculated here are found to have some quantitative differences with calculations of the type presented in the previous publications, in which the superconducting domain boundary is relatively close to the plasma. The free-boundary solution is found to have a larger kinking response in this case, especially for $n=1$, which is consistent with the idea that the plasma becomes easier to excite as it moves closer to marginal stability. Similar results have been found both experimentally and with ideal-MHD by varying the plasma pressure.²⁸ Unexpectedly, the total resonant fields in the $n=3$ response tend to be larger in the conducting wall case than in the free-boundary case. This may provide an explanation for why previous M3D-C1 calculations, which had a conducting wall close to the plasma, typically found less screening than similar free-boundary calculations with MARS.

Axisymmetric modes, including the VDE, are strongly affected by the presence of the resistive wall. The VDE simulations done here have demonstrated the capability to follow a disruption in realistic geometry into the current quench phase. Although the calculations presented here are done at low resolution and not meant to be quantitatively comparable to experiment, they exhibit behavior that is likely qualitatively similar to VDEs observed in experiments. During the early phase of the vertical instability, the response currents arising in the plasma edge and in the wall are calculated, as is the dependence of their relative strengths on the resistive wall time. Even over five orders of magnitude in the resistive wall time, the maximum current density arising in the wall, and hence the maximum force on the wall, is found to vary only by roughly a factor of 2 (cf. Figure 16). However, since the timescale of the instability is strongly dependent on the wall time, the total impulse delivered to the wall increases with the wall conductivity.

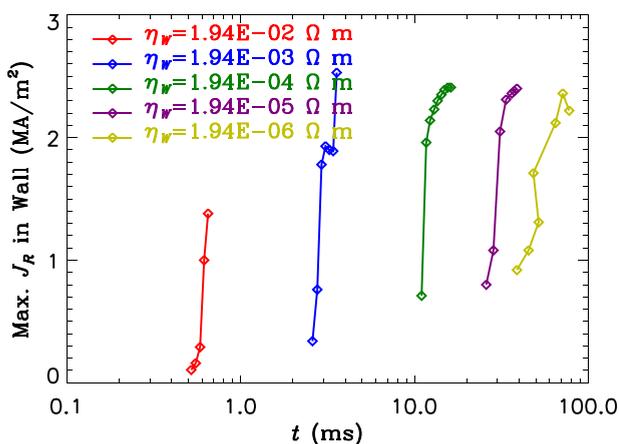


FIG. 16. The maximum radial current density in the resistive wall as a function of time, for various values of η_w .

The spike in the toroidal current enclosed by the wall is also strikingly similar to experimental observations.²⁷ In experiments, this spike is often attributed to the rapid redistribution of the plasma current density during the thermal quench. However, in these calculations, the spike is clearly associated with the plasma contacting the wall; even though there is a redistribution of current density during the thermal quench, no significant change to the total current is observed during that time. This is actually consistent with some experimental results in which the current spike was observed not to occur until well after the thermal quench began.²⁷ Still, it is not clear that mechanism responsible for the current spikes in these simulations is the sole or primary mechanism for those observed in experiments. In some experiments, for example, the spike is observed before equilibrium reconstructions show the plasma undergoing a significant displacement.²⁹ It is possible that in these instances, the spike seen in experiments is related to forces not present in the model considered here. In particular, TSC simulations exhibit a current spike during the thermal quench when hyper-resistivity is included.³⁰

After the current spike in these calculations, q_{edge} is found to drop as the plasma is scraped off by the first wall, consistent with the expectations. Inevitably, this will lead to a kink instability, which is not modeled in these axisymmetric simulations. A full accounting of currents and forces in the wall will require modeling the onset and evolution of such a mode, and recent calculations imply that these non-axisymmetric forces may strongly depend on the wall time.^{31,32} The capability demonstrated here will improve studies of this type by allowing a more scalable computation that can span transport timescales at realistic parameters. This will be a primary focus of future work.

Simulations of several other important tokamak phenomenon will be enabled by this wall model in addition to those applications presented here. In particular, it will allow simulations of mode-locking, a process in which a rotating magnetic perturbation (such as a magnetic island) induces eddy currents in the wall which result in an electromagnetic torque on the plasma. This torque may brake and eventually halt the rotation of the plasma, which often leads to a disruption. While this phenomenon is well understood conceptually, quantitative predictive models of locking and the subsequent inception of a disruption have not yet been demonstrated.

ACKNOWLEDGMENTS

We thank T. Strait and D. Shirake for helpful discussions regarding experimental observations of disruptions. This work was supported by the U.S. Department of Energy Grant Nos. DE-FG02-95ER54309, DE-FC02-04ER54698, DE-AC02-09CH11466, and DE-SC0006617.

APPENDIX A: RWM DISPERSION RELATION FOR WALL OF ARBITRARY THICKNESS

Liu *et al.*¹⁴ derive the dispersion relation for a resistive wall mode in straight cylindrical geometry using a two-field model and a step-function equilibrium, as described in Section III. They show that in this model, the poloidal flux must satisfy the equation

$$\nabla_{\perp}^2 \psi = 0, \quad (\text{A1})$$

everywhere, and is subject to the jump condition

$$\left[\frac{r\psi'}{\psi} \right]_{r=r_0} = -\frac{2m}{m-nq_0} + \frac{q_0^2(\gamma\tau_A)^2}{(m-nq_0)^2} \frac{r_0\psi'(r_0)}{\psi(r_0)}, \quad (\text{A2})$$

at the plasma-vacuum interface ($r=r_0$). Additionally, the thin wall approximation yields a second jump condition at $r=r_w$

$$\left[\frac{r\psi'}{\psi} \right]_{r=r_w} = 2\gamma\tau_w. \quad (\text{A3})$$

Equation (A1) implies a solution of the form

$$\psi = \begin{cases} A_p r^{\mu} & r < r_0 \\ A_v r^{\mu} + B_v r^{-\mu} & r_0 < r < r_w \\ B_x r^{-\mu} & r_w < r, \end{cases} \quad (\text{A4})$$

which, together with the jump conditions and enforcing continuity of $\psi(r_0)$ and $\psi(r_w)$, yields the thin-wall dispersion relation equation (21).

In the general case, not considered by Liu *et al.*, the jump condition Equation (A3) is replaced by Faraday's law $\partial \vec{B} / \partial t = -\nabla \times (\eta_w \vec{J})$ in the wall region $r_w < r < r_w + d$. Using the modified form of the resistivity appropriate for the two-field model, as discussed in Section III, Ohm's law reduces to $\gamma\psi = \eta_w \nabla_{\perp}^2 \psi$. The general solutions to this equation are modified Bessel functions $I_{\mu}(kr)$ and $K_{\mu}(kr)$, where $k^2 = \gamma/\eta_w$. (Note that for the standard form of the resistivity, in which poloidal current is damped, Faraday's law would yield $\gamma\psi = \eta_w \nabla^2 \psi$ and the general solution would have $k^2 = \gamma/\eta_w + n^2/R^2$.) Thus, ψ is of the form

$$\psi = \begin{cases} A_p r^{\mu} & r < r_0 \\ A_v r^{\mu} + B_v r^{-\mu} & r_0 < r < r_w \\ A_w I_{\mu}(kr) + B_w K_{\mu}(kr) & r_w < r < r_w + d \\ B_x r^{-\mu} & r_w + d < r. \end{cases} \quad (\text{A5})$$

Taking this form of ψ together with the jump condition Equation (A2), and enforcing continuity of $\psi(r_0)$, $\psi'(r_w)$, and $\psi'(r_w + d)$, yields the general dispersion relation, Equation (25), with $\rho_a = kr_w$. This reduces to Equation (21) in the limit $d \ll r_w$.

APPENDIX B: COMPUTATIONAL EXPENSE

The computational resources used by the calculations presented here vary from roughly 8 cpu-h for high-resolution linear perturbed equilibrium calculations to over 300 cpu-h for transport-timescale nonlinear axisymmetric VDE calculations. High-resolution meshes in two dimensions may have 25k mesh nodes or more; with six fields and twelve degrees of freedom per field per mesh node, one of these calculations may have nearly 2M degrees of freedom. Parallelization is accomplished using Message Passing Interface (MPI) for domain decomposition; typically 8–32 domains (with one domain per MPI process) are used in a calculation. The

inclusion of two-fluid or other non-ideal effects does not strongly impact the computational expense of these calculations, in which the matrix equations are solved using direct LU-factorization methods (both SuperLU_dist³³ and MUMPS³⁴ have been used successfully on different hardware systems), which are relatively insensitive to the stiffness of the system.

In marked contrast, fully nonlinear 3D calculations with M3D-C1 (not presented here), which typically have 24 or more toroidal planes, often require on the order of 100k cpu-h, depending on the problem. Typically, between 1k and 10k computational cores are used in these calculations. Scalability and memory considerations require iterative matrix solution methods to be used in these cases. Therefore, the additional stiffness incurred by including two-fluid effects increases the computational cost of 3D nonlinear calculations. A more complete description of the numerical methods used by M3D-C1 can be found in a publication by Jardin *et al.*³⁵

- ¹J. King, E. Strait, S. Lazerson, N. Ferraro, N. Logan, S. Haskey, J.-K. Park, M. Lanctot, Y. Liu, R. Nazikian, M. Okabayashi, C. Paz-Soldan, D. Shiraki, and A. Turnbull, *Phys. Plasmas* **22**, 072501 (2015).
- ²F. Schuller, *Plasma Phys. Controlled Fusion* **37**, A135 (1995).
- ³E. Strait, L. Lao, J. Luxon, and E. Reis, *Nucl. Fusion* **31**, 527 (1991).
- ⁴R. Sayer, Y.-K. Peng, S. Jardin, A. Kellman, and J. Wesley, *Nucl. Fusion* **33**, 969 (1993).
- ⁵S. Miayamoto, A. Isayama, I. Bandyopadhyay, S. Jardin, R. Khayrutdinov, V. Lukash, Y. Kusama, and M. Sugihara, *Nucl. Fusion* **54**, 083002 (2014).
- ⁶H. R. Strauss, A. Pletzer, W. Park, S. Jardin, J. Breslau, and L. Sugiyama, *Comput. Phys. Commun.* **164**, 40 (2004).
- ⁷M. Hölzl, P. Merkel, G. Huysmans, E. Nardon, E. Strumberger, R. McAdams, I. Chapman, S. Günter, and K. Lackner, *J. Phys.: Conf. Ser.* **401**, 012010 (2012).
- ⁸A. H. Glasser, C. R. Sovinec, R. A. Nebel, T. A. Gianakon, S. J. Plimpton, M. S. Chu, D. D. Schnack, and NIMROD Team, *Plasma Phys. Controlled Fusion* **41**, A747 (1999).
- ⁹A. L. Becerra, C. C. Hegna, C. R. Sovinec, S. E. Kruger, J. R. King, and S. A. Sabbagh, in *APS Meeting Abstracts* (2014), p. 8047P.
- ¹⁰Y. Q. Liu, A. Bondeson, C. M. Fransson, B. Lennartson, and C. Breitholtz, *Phys. Plasmas* **7**, 3681 (2000).
- ¹¹M. Chance, *Phys. Plasmas* **4**, 2161 (1997).
- ¹²A. Pletzer and H. Strauss, *Comput. Phys. Commun.* **182**, 2077 (2011).
- ¹³N. M. Ferraro and S. C. Jardin, *J. Comput. Phys.* **228**, 7742 (2009).
- ¹⁴Y. Liu, R. Albanese, A. Portone, G. Rubinacci, and F. Villone, *Phys. Plasmas* **15**, 072516 (2008).
- ¹⁵S. C. Jardin, J. Breslau, and N. Ferraro, *J. Comput. Phys.* **226**, 2146 (2007).
- ¹⁶R. McAdams, "Non-linear magnetohydrodynamic instabilities in advanced tokamak plasmas," Ph.D. thesis (University of York, 2014).
- ¹⁷Q. Ren, M. Chu, L. Lao, and R. Srinivasan, *Plasma Phys. Controlled Fusion* **53**, 095009 (2011).
- ¹⁸B. J. Burke, S. E. Kruger, C. C. Hegna, P. Zhu, P. B. Snyder, C. R. Sovinec, and E. C. Howell, "Edge localized linear ideal magnetohydrodynamic instability studies in an extended-magnetohydrodynamic code," *Phys. Plasmas* **17**, 032103 (2010).
- ¹⁹N. M. Ferraro, S. C. Jardin, and P. B. Snyder, *Phys. Plasmas* **17**, 102508 (2010).
- ²⁰S. C. Jardin, *J. Comput. Phys.* **200**, 133 (2004).
- ²¹N. M. Ferraro, *Phys. Plasmas* **19**, 056105 (2012).
- ²²M. Shafer, E. Unterberg, A. Wingen, D. Battaglia, T. Evans, N. Ferraro, J. Harris, D. Hillis, and R. Nazikian, *Phys. Plasmas* **21**, 122518 (2014).
- ²³A. Wingen, N. Ferraro, M. Shafer, E. Unterberg, T. Evans, D. Hillis, and P. Snyder, *Nucl. Fusion* **54**, 064007 (2014).
- ²⁴R. Fitzpatrick and T. C. Hender, *Phys. Fluids B* **3**, 644 (1991).
- ²⁵N. M. Ferraro and S. C. Jardin, *Phys. Plasmas* **13**, 092101 (2006).
- ²⁶L. Zakharov, S. Galkin, S. Gerasimov, and JET-EFDA Contributors, *Phys. Plasmas* **19**, 055703 (2012).
- ²⁷J. Wesson, D. Ward, and M. Rosenbluth, *Nucl. Fusion* **30**, 1011 (1990).

- ²⁸M. J. Lanctot, H. Reimerdes, A. M. Garofalo, M. S. Chu, Y. Q. Liu, E. J. Strait, G. L. Jackson, R. J. La Haye, M. Okabayashi, T. H. Osborne, and M. J. Schaffer, *Phys. Plasmas* **17**, 030701 (2010).
- ²⁹R. Granetz, I. Hutchinson, J. Sorci, J. Irby, B. LaBombard, and D. Gwinn, *Nucl. Fusion* **36**, 545 (1996).
- ³⁰B. Merrill, S. Jardin, M. Ulrickson, and M. Bell, *Fusion Eng. Des.* **15**, 163 (1991).
- ³¹H. Strauss, R. Paccagnella, and J. Breslau, *Phys. Plasmas* **17**, 082505 (2010).
- ³²H. Strauss, R. Paccagnella, J. Breslau, L. Sugiyama, and S. Jardin, *Nucl. Fusion* **53**, 073018 (2013).
- ³³X. S. Li and J. W. Demmel, *ACM Trans. Math. Software* **29**, 110 (2003).
- ³⁴P. R. Amestoy, C. Ashcraft, O. Boiteau, A. Buttari, J.-Y. L'Excellent, and C. Weisbecker, *SIAM J. Sci. Comput.* **37**, A1451 (2015).
- ³⁵S. C. Jardin, N. Ferraro, J. Breslau, and J. Chen, *Comput. Sci. Discovery* **5**, 014002 (2012).