Available online at www.sciencedirect.com
SCIENCE DIRECT®
Journal of Computational Physics 195 (2004) 355–386
JOURNAL OF COMPUTATIONAL PHYSICS
www.elsevier.com/locate/jcp

Nonlinear magnetohydrodynamics simulation using high-order finite elements ☆

C.R. Sovinec [a,*], A.H. Glasser [b], T.A. Gianakon [b], D.C. Barnes [b], R.A. Nebel [b], S.E. Kruger [c], D.D. Schnack [c], S.J. Plimpton [d], A. Tarditi [e], M.S. Chu [f], the NIMROD Team

[a] *Department of Engineering Physics, University of Wisconsin-Madison, 1500 Engineering Drive, Madison, WI 53706-1609, USA*
[b] *Los Alamos National Laboratory, Los Alamos, NM 87545, USA*
[c] *Science Applications International Corporation, San Diego, CA 92121, USA*
[d] *Sandia National Laboratory, Albuquerque, NM 87185, USA*
[e] *Advanced Space Propulsion Laboratory, National Aeronautics and Space Administration-Johnson Space Center, Houston, TX 77050, USA*
[f] *General Atomics Corporation, San Diego, CA 92138, USA*

Received 30 December 2002; received in revised form 29 August 2003; accepted 6 October 2003

Abstract

A conforming representation composed of 2D finite elements and finite Fourier series is applied to 3D nonlinear non-ideal magnetohydrodynamics using a semi-implicit time-advance. The self-adjoint semi-implicit operator and variational approach to spatial discretization are synergistic and enable simulation in the extremely stiff conditions found in high temperature plasmas without sacrificing the geometric flexibility needed for modeling laboratory experiments. Growth rates for resistive tearing modes with experimentally relevant Lundquist number are computed accurately with time-steps that are large with respect to the global Alfvén time and moderate spatial resolution when the finite elements have basis functions of polynomial degree (p) two or larger. An error diffusion method controls the generation of magnetic divergence error. Convergence studies show that this approach is effective for continuous basis functions with $p \geqslant 2$, where the number of test functions for the divergence control terms is less than the number of degrees of freedom in the expansion for vector fields. Anisotropic thermal conduction at realistic ratios of parallel to perpendicular conductivity ($\chi_{\parallel}/\chi_{\perp}$) is computed accurately with $p \geqslant 3$ without mesh alignment. A simulation of tearing-mode evolution for a shaped toroidal tokamak equilibrium demonstrates the effectiveness of the algorithm in nonlinear conditions, and its results are used to verify the accuracy of the numerical anisotropic thermal conduction in 3D magnetic topologies.

AMS: 65M60; 76W05; 76X05

☆ This research is supported by the Office of Science, US Department of Energy.
* Corresponding author. Present address: Department of Engineering Physics, University of Wisconsin-Madison, 1500 Engineering Drive, Madison, WI 53706-1609, USA. Tel.: +1-608-263-5525; fax: +1-608-265-2438.
E-mail address: sovinec@engr.wisc.edu (C.R. Sovinec).
URL: http://nimrodteam.org.

0021-9991/$ - see front matter
doi:10.1016/j.jcp.2003.10.004

Keywords: Magnetohydrodynamic simulation; Finite element; Semi-implicit; Anisotropic diffusion

1. Introduction

High temperature magnetized plasmas are characterized by extremely anisotropic properties relative to the direction of the magnetic field. Perpendicular motions of charged particles are constrained by the Lorentz force, while relatively unrestrained parallel motions lead to rapid transport along magnetic field lines. The orientation and distribution of fluid-like motions of the electrically conducting plasma then determine the degree of restoring force arising from the bending and compression of magnetic flux tubes. When collective motions are able to avoid these restoring forces while releasing available free energy, magnetohydrodynamic (MHD) instability results. As an unstable perturbation grows to finite amplitude, it may induce a nonlinear evolution of the system that includes significant (and sometimes catastrophic) changes in thermal energy and particle confinement. The behavior is often complex, so that analysis must rely on simulation, but the large anisotropies relative to the distorted magnetic field present challenging conditions for numerical methods. For example, numerical truncation errors associated with rapid parallel thermal conduction produce artificial heat transport that leads to qualitative errors in the simulated energy confinement when using low-order representations.

The anisotropies also lead to a wide range of time-scales for different physical effects. For typical conditions in magnetically confined plasmas, parallel thermal conduction is the fastest process in the system. Alfvén-wave propagation occurs on a longer time-scale, followed by sound-wave propagation. Perpendicular thermal conduction and particle diffusion occur on longer time-scales, and global magnetic field diffusion (from non-zero resistivity) is the slowest process. Topology-changing magnetic reconnection occurs on a hybrid time-scale between Alfvénic propagation and global resistive diffusion, and the associated subsonic flows are nearly incompressible, so numerical simulation of this behavior must deal with extreme stiffness resulting from relatively fast wave propagation and parallel thermal conduction. Simulating the behavior of the system is therefore related to various aspects of the numerical simulation of electromagnetics, incompressible fluid dynamics, convective heat transfer, and linear ideal MHD.

Numerical resolution of magnetohydrodynamic anisotropy leading to singular behavior in ideal conditions has been achieved in linear computations by using specialized low-order discretization methods. These methods require solving the displacement vector in the components of a fixed magnetic-flux coordinate system, aligning the numerical mesh with the equilibrium magnetic field, and using different finite element basis functions in the parallel and perpendicular directions [1,2]. For nonlinear simulation, this approach is less compelling. Nonlinear evolution often forms regions with distinct magnetic topology, such as helical islands or regions of magnetic stochasticity embedded in nested flux surfaces. Either occurrence would present formidable challenges for (1) an adaptive meshing algorithm to preserve alignment with the complicated magnetic field and (2) an arrangement of particular basis functions to match the adaptive mesh. Furthermore, a basis function expansion tailored to a particular set of equations may not be suitable for other physical models. For example, discontinuous finite element representations of velocity field components cannot be applied to a system with viscous dissipation without resorting to non-conforming or more complicated mixed approximations. Since closure relations for fluid models remains an active area of research in plasma theory, a specialized discretization will have limited usefulness for a simulation code that is intended to have flexibility in the equations that it solves.

An alternative is to use a numerical representation that has a high rate of spatial convergence. While a number of high-order approximations are possible for simple configurations, the ability to represent a realistic geometry is important for analyzing laboratory data. High-order finite difference methods there-

fore have limited applicability, and the nonlinear character of high-order finite volume methods [3] (designed for accuracy with discontinuous solutions) is not suited for conditions where stiff linear behavior and resolution of narrow dissipation layers is important. The finite element method provides a better approach for nonlinear fusion MHD, where dissipation terms ensure smoothness with sufficient resolution. The convergence rate realized by the finite element method is then controlled by the degree (p) of the polynomial basis functions, relatively independent of geometry and mesh spacing irregularities. In addition, a general finite element implementation can achieve convergence by increasing p with a fixed mesh [4], which constitutes a spectral method.

Applying the finite element method to time-dependent systems leads to separate variational problems for each equation in a marching algorithm if the implicit terms are based on self-adjoint differential operators. Standard analysis can then be used to estimate convergence with respect to mesh spacing when the representation satisfies two conditions: first, the space of piecewise polynomials (S_h) of degree p must be composed of admissible functions, which in our application means that a functional is finite for all finite-valued functions in S_h and that S_h only includes functions that satisfy the Dirichlet boundary conditions. Second, the explicit terms in the marching algorithm, i.e. the "data" for each variational problem, must remain square-integrable functions throughout the evolution. (Chapter 1 of [5] provides a concise mathematical background.) Given these conditions, the analysis tells us that the finite element solution ($\bar{u}$) to a variational problem is the function in S_h with the least "strain energy" error [5], i.e.

$$a(u-\bar{u},u-\bar{u}) \leqslant a(u-\bar{v},u-\bar{v}) \quad \text{for all } \bar{v} \in S_h, \tag{1}$$

where u is the best solution among all admissible functions. Then, knowing that the finite element solution is a better approximation in terms of the strain energy than the interpolate function, which is also in S_h, we eventually arrive at relations for convergence rates [5],

$$|u-\bar{u}|_0 \leqslant K_0 h^{p+1} |u|_{p+1} \tag{2}$$

and

$$|u-\bar{u}|_1 \leqslant K_1 h^{p} |u|_{p+1}, \tag{3}$$

where h characterizes the possibly irregular mesh spacing, $|u|_s$ is the norm of the sth derivative of u, and K_0 and K_1 are independent of h. [The estimates (2) and (3) are for the relevant special case of second-order partial differential equations.] For a time-advance that solves for different fields sequentially, there is a unique strain energy for each equation, and the set of minimization problems is solved at each time-step.

While the finite element representation allows high-order accuracy without restricting geometry, it introduces other challenges. Besides implementation complications, it is well known from incompressible fluid modeling that continuous finite element representations of vector components cannot reproduce a divergence constraint exactly. Furthermore, ensuring convergence to a divergence-free space requires special attention. For plasma modeling, this issue arises with the zero-magnetic-monopole constraint and with nearly divergence-free velocity distributions associated with many unstable MHD modes. A straightforward approach for approximating the magnetic divergence constraint is to add the diffusive term $\kappa_{\mathrm{div}\,b}\nabla\nabla\cdot\mathbf{B}$ to Faraday's law [6–8]. This leads to a method that is related to divergence cleaning techniques for finite difference and finite volume methods [9] and to penalty function methods for finite elements [10].

Here, we report on this application of the finite element method to nonlinear non-ideal MHD, and its implementation in the NIMROD code (non-ideal magnetohydrodynamics with rotation, open discussion) [7]. The objective of the NIMROD project (http://nimrodteam.org) is to achieve accurate and flexible modeling of nonlinear electromagnetic activity in computational domains that are realistic for a variety of

laboratory plasmas. Unlike most previous efforts for nonlinear modeling of high temperature plasmas [11–14], we have avoided spatial representations that restrict the geometry in the poloidal domain. The present NIMROD implementation has the parameter p selected at run-time, which is more general than either the finite element implementation reported in [15] or the earlier NIMROD implementation [7], which used linear and bilinear elements only. This feature has proven useful for exploring the performance of different basis functions in actual applications, and our findings confirm that using $p > 1$ is essential for modeling anisotropies and for satisfying the magnetic divergence constraint. We have restricted our attention to periodic configurations with a 2D boundary, so the finite Fourier series representation with pseudospectral computations of nonlinear terms [16] is applied.

The separation of time-scales in high temperature plasmas is manifest mathematically as stiffness in the non-ideal MHD model, and this is an equally important consideration for numerical simulation. The dominant part of the stiffness can be described through the linear properties of the system at any given time, since propagating shocks do not occur on these slow time-scales. The stiffness makes explicit methods impractical, but semi-implicit methods [17,18] are well suited for these conditions. The semi-implicit operator considered here is based on the linear ideal MHD force operator, as recommended in [13], but the symmetric component of the solution's Fourier expansion is incorporated into the equilibrium fields. (Geometrically, we use "symmetry" with respect to the periodic coordinate, e.g. the toroidal direction for toroidal geometries, which is represented by the finite Fourier series. However, "symmetry" is also used in the mathematical context of symmetric matrices.) In addition, the Laplacian operator used for stabilizing nonlinear pressures has a dynamic coefficient that depends on the non-symmetric part of the solution. This approach makes the algorithm suitable for simulations where the fields evolve significantly from their initial equilibrium configuration, while retaining the accuracy reported in [13]. Furthermore, since each advance in the marching algorithm has a self-adjoint operator for its implicit terms, and positive eigenvalues can be ensured, the requirements for a variational approach to spatial discretization are met. In many cases, there is no implicit dependence among Fourier components, so the resulting algebraic systems have sparse matrices. For equations that have implicit coupling in all three directions, the Fourier representation leads to an algebraic system that includes convolutions among Fourier components.

The NIMROD code has been written for parallel computation on distributed-memory computers with communication routines from the message passing interface (MPI) library (http://www.mpi-forum.org). Standard mesh decomposition techniques with point-to-point communication work well for the finite element representation of the poloidal plane, where overlap of basis functions is local. Coupling in the periodic direction occurs through fast fourier transforms (FFTs) and algebraic operations on a uniform grid over this coordinate. Here, swapping from Fourier-based decomposition to spatially based decomposition via collective communication is used to maintain scalability. Computationally, the most demanding part of the algorithm is the solution of the linear systems. A recent change in the NIMROD code is the use of the parallel distributed-memory version of the SuperLU software library (http://acts.nersc.gov/superlu/) to apply sparse direct-solve methods to the systems that do not have Fourier convolutions. For the systems with coupling among Fourier components, NIMROD has a matrix-free conjugate gradient solve that calls SuperLU routines to invert sparse approximations of the complete matrices as a preconditioning step.

The organization for the remainder of this article is as follows. Section 2 describes the magnetofluid equations solved by NIMROD, and Section 3 presents the discretization techniques that have been applied. In Section 4, we use a resistive linear MHD benchmark to show convergence properties in stiff conditions and to demonstrate performance with respect to the divergence constraint. We also present NIMROD results on a quantitative test of anisotropic thermal conduction. A sample nonlinear simulation that brings together MHD stiffness and anisotropic energy transport is presented in Section 5. In Section 6, we further discuss the properties of the algorithm that are observed in the test results and make comparisons with ideal

MHD eigenvalue computation and incompressible fluid modeling. Conclusions are given in Section 7. The Appendix describes our implementation of regularity conditions for simply connected (topologically cylindrical) configurations.

2. Equations

Resistive MHD is the simplest model capable of reproducing global electromagnetic behavior observed in many laboratory and natural plasmas. For long time-scales, where important nonlinear evolution occurs, it is often necessary to include diffusion and conduction terms, since transport processes act on similar time-scales. The non-ideal model considered in this paper is resistive MHD with anisotropic thermal conduction, kinematic viscous dissipation, particle density diffusion, and the numerical diffusion of magnetic divergence error. Separating terms that represent a steady solution (denoted by the "ss" subscript), this non-ideal MHD model is

$$\frac{\partial \mathbf{B}}{\partial t} = -\nabla \times \mathbf{E} + \kappa_{\mathrm{div}\,b} \nabla \nabla \cdot \mathbf{B}, \tag{4a}$$

$$\mathbf{E} = -\mathbf{V} \times \mathbf{B} - \mathbf{V}_{\mathrm{ss}} \times \mathbf{B} - \mathbf{V} \times \mathbf{B}_{\mathrm{ss}} + \eta \mathbf{J}, \tag{4b}$$

$$\mu_0 \mathbf{J} = \nabla \times \mathbf{B}, \tag{4c}$$

$$\frac{\partial n}{\partial t} + \nabla \cdot (n\mathbf{V} + n_{\mathrm{ss}}\mathbf{V} + n\mathbf{V}_{\mathrm{ss}}) = \nabla \cdot D \nabla n, \tag{4d}$$

$$\begin{aligned}(\rho + \rho_{\mathrm{ss}}) &\left(\frac{\partial \mathbf{V}}{\partial t} + \mathbf{V} \cdot \nabla \mathbf{V} + \mathbf{V}_{\mathrm{ss}} \cdot \nabla \mathbf{V} + \mathbf{V} \cdot \nabla \mathbf{V}_{\mathrm{ss}} \right) + \rho \mathbf{V}_{\mathrm{ss}} \cdot \nabla \mathbf{V}_{\mathrm{ss}} \\ &= \mathbf{J} \times \mathbf{B} + \mathbf{J}_{\mathrm{ss}} \times \mathbf{B} + \mathbf{J} \times \mathbf{B}_{\mathrm{ss}} - \nabla p + \nabla \cdot \nu(\rho + \rho_{\mathrm{ss}}) \nabla \mathbf{V} + \nabla \cdot \nu \rho \nabla \mathbf{V}_{\mathrm{ss}},\end{aligned} \tag{4e}$$

$$\begin{aligned}\frac{n + n_{\mathrm{ss}}}{\gamma - 1} &\left(\frac{\partial T}{\partial t} + \mathbf{V} \cdot \nabla T + \mathbf{V}_{\mathrm{ss}} \cdot \nabla T + \mathbf{V} \cdot \nabla T_{\mathrm{ss}} \right) + \frac{n}{\gamma - 1} \mathbf{V}_{\mathrm{ss}} \cdot \nabla T_{\mathrm{ss}} \\ &= -\frac{p}{2} \nabla \cdot \mathbf{V} - \frac{p_{\mathrm{ss}}}{2} \nabla \cdot \mathbf{V} - \frac{p}{2} \nabla \cdot \mathbf{V}_{\mathrm{ss}} - \nabla \cdot \mathbf{q} + Q,\end{aligned} \tag{4f}$$

where $\mathbf{E}$ is the electric field, $\mathbf{B}$ is the magnetic induction, $\mathbf{V}$ is the particle flow velocity, $\mathbf{q}$ is the heat flux vector, Q is the heat source density, and γ is the ratio of specific heats. The units are MKS, except that the Boltzmann constant has been absorbed into temperature. The particle number density n and mass density ρ are related through the mass per ion (m_i), and total pressure and temperature follow the ideal gas relation, $p = 2nT$, assuming quasineutrality ($n_e \cong n_i = n$) and rapid thermal equilibration among ions and electrons. Eqs. (4a)–(4f) represent the modified Faraday's law, the resistive MHD Ohm's law, the low-frequency limit of Ampere's law, particle conservation, flow velocity evolution, and temperature evolution, respectively. The particle diffusion term is necessary for simulations over transport time-scales, where physical effects beyond MHD influence the number density profile. Its implementation is only phenomenological, because the particle flux should be consistent with the product of the number density and the flow velocity. Finding a better representation of the particle transport is important, but it is beyond the scope of the present effort.

The steady-state terms make the system of equations suitable for nonlinear computations of deviations from a time-independent solution of the same physics model. We note that this is conceptually similar to

linear MHD stability analysis of a solution to the force balance equation ($\mathbf{J}\times\mathbf{B}=\nabla p$), but linear MHD is often used to analyze equilibria that evolve on transport time-scales. In contrast, computations of nonlinear perturbations over long time-scales require steady-state fields that are time-independent solutions of the complete system. For example, the steady-state may have non-zero electric field ($-\mathbf{V}_{ss}\times\mathbf{B}_{ss}+\eta\mathbf{J}_{ss}\neq\mathbf{0}$), but it is assumed to be curl-free and is not computed with the terms in Eq. (4b) that influence the evolution of the perturbed magnetic field through Eq. (4a). Separating steady-state terms in the equations adds complexity to the coding, but it improves numerical accuracy in simulations where the perturbations are small relative to the steady part of the fields [14]. There are also practical benefits for analyzing MHD activity. Fitting equilibrium MHD solutions to data from laboratory measurements is now common experimental practice. Solving the nonlinear evolution of perturbations about a fitted equilibrium provides a powerful analysis tool without the need for complete information regarding the sources that sustain the equilibrium profiles of current, plasma flow, internal energy density, and particle density. Since NIMROD assumes a domain that is symmetric in the periodic coordinate, only symmetric steady-state fields are considered. The perturbed fields are fully 3D, however.

Thermal transport in Eq. (4f) can be modeled as local anisotropic diffusion with separate coefficients for the parallel and perpendicular directions [19]

$$\mathbf{q}=-n[\chi_{\parallel}\hat{\mathbf{b}}\hat{\mathbf{b}}+\chi_{\perp}(\mathbf{I}-\hat{\mathbf{b}}\hat{\mathbf{b}})]\cdot\nabla T, \tag{5}$$

where $\hat{\mathbf{b}}\equiv\mathbf{B}/|\mathbf{B}|$ is the local magnetic direction vector – terms for the separated steady-state fields have been suppressed for clarity. In high temperature plasmas, $\chi_{\parallel}$ may be many orders of magnitude larger than $\chi_{\perp}$, which presents numerically challenging conditions when $\hat{\mathbf{b}}$ is not aligned with the mesh (see Section 4.2). The source term Q in (4f) represents the sum of ohmic ($\eta\mathbf{J}^2$) and viscous ($\nu\rho\nabla\mathbf{V}^{\mathrm{T}}:\nabla\mathbf{V}$) heating.

The boundary conditions considered here for Eqs. (4a)–(4f) are Dirichlet conditions for the normal component of $\mathbf{B}$, for T, and for all components of $\mathbf{V}$ along the bounding surface. For the tangential component of $\mathbf{B}$ and for n, fluxes are specified as natural boundary conditions via surface integrals in the variational form of the equations. Here, the respective flux densities are $\hat{\mathbf{n}}\times\mathbf{E}$ and $-D\nabla n$.

The model represented by Eqs. (4a)–(4f) can be extended to include two-fluid effects, non-local effects of rapid particle streaming at arbitrary collisionality [20], neoclassical effects [21], and ion kinetic effects [22] that are important in the dynamics of many high temperature plasmas. The spatial representation described herein provides a basis for the numerical development of these advanced models, in addition to its utility for the non-ideal MHD model.

3. Numerical methods

3.1. Time-advance

The numerical approach we have used for Eqs. (4a)–(4f) combines the solution efficiency of a semi-implicit time-advance with the geometric flexibility and accuracy of a general finite element method for spatial representation. We arrive at our numerical system of equations by first applying temporal discretization to Eqs. (4a)–(4f). The velocity field values are defined at integer time indices, whereas the remaining fields are defined at half-integer time indices. This creates a leap-frog scheme, and the semi-implicit operator is used in the velocity advance to eliminate time-step restrictions associated with oscillatory behavior. The stabilizing truncation error in this algorithm is dispersive but not dissipative [23], which is an important consideration for simulating conditions where the physical dissipation terms are small.

Our semi-implicit operator consists of two parts, as in [13]. The first includes terms that stabilize wave propagation about the symmetric fields, and the second part stabilizes wave propagation when a significant

non-symmetric component of the solution develops. The first part is derived from the method of differential approximation [24] by considering the ideal portion of the system, which describes oscillatory behavior and ideal linear MHD instabilities. After removing the dissipative and heating terms, the temperature and continuity equations are equivalent to the adiabatic pressure relation

$$\frac{\partial p}{\partial t} = -\mathbf{V}\cdot\nabla p - \gamma p \nabla\cdot\mathbf{V}.$$

Thus, the differential approximation technique is applied to the ideal equations for pressure, magnetic field, and flow velocity.

Applying the approach of [24] for generic wave equations, the differential approximation of an implicit numerical time-advance for the linear ideal MHD equations is

$$\rho_0\frac{\partial \mathbf{V}}{\partial t} - \theta\Delta t\left[\frac{1}{\mu_0}\left(\nabla\times\frac{\partial \mathbf{B}}{\partial t}\right)\times\mathbf{B}_0 + \mathbf{J}_0\times\frac{\partial \mathbf{B}}{\partial t} - \nabla\frac{\partial p}{\partial t}\right] = \frac{1}{\mu_0}(\nabla\times\mathbf{B})\times\mathbf{B}_0 + \mathbf{J}_0\times\mathbf{B} - \nabla p, \tag{6a}$$

$$\frac{\partial \mathbf{B}}{\partial t} - \theta\Delta t\nabla\times\left(\frac{\partial \mathbf{V}}{\partial t}\times\mathbf{B}_0\right) = \nabla\times(\mathbf{V}\times\mathbf{B}_0), \tag{6b}$$

$$\frac{\partial p}{\partial t} + \theta\Delta t\left(\frac{\partial \mathbf{V}}{\partial t}\cdot\nabla p_0 + \gamma p_0\nabla\cdot\frac{\partial \mathbf{V}}{\partial t}\right) = -(\mathbf{V}\cdot\nabla p_0 + \gamma p_0\nabla\cdot\mathbf{V}), \tag{6c}$$

where θ is the centering parameter ($0\leqslant\theta\leqslant 1$) and $\mathbf{V}_0\cong\mathbf{0}$ is assumed so that $\mathbf{B}_0$, $\mathbf{J}_0$, and p_0 satisfy the static force balance equation, $\mathbf{J}_0\times\mathbf{B}_0 = \nabla p_0$. Differentiating Eq. (6a) with respect to time and eliminating $\mathbf{B}$ and p produces the wave equation

$$\rho_0\frac{\partial^2\mathbf{V}}{\partial t^2} - \theta^2\Delta t^2\mathbf{L}\left(\frac{\partial^2\mathbf{V}}{\partial t^2}\right) = 2\theta\Delta t\mathbf{L}\left(\frac{\partial \mathbf{V}}{\partial t}\right) + \mathbf{L}(\mathbf{V}), \tag{7}$$

where $\mathbf{L}$ is the self-adjoint linear ideal MHD force operator

$$\mathbf{L}(\mathbf{V}) = \frac{1}{\mu_0}\{\nabla\times[\nabla\times(\mathbf{V}\times\mathbf{B}_0)]\}\times\mathbf{B}_0 + \mathbf{J}_0\times\nabla\times(\mathbf{V}\times\mathbf{B}_0) + \nabla(\mathbf{V}\cdot\nabla p_0 + \gamma p_0\nabla\cdot\mathbf{V}). \tag{8}$$

The wave equation (7) can also be expressed as the system

$$\rho_0\frac{\partial \mathbf{V}}{\partial t} - \theta^2\Delta t^2\mathbf{L}(\partial\mathbf{V}/\partial t) = \frac{1}{\mu_0}(\nabla\times\mathbf{B})\times\mathbf{B}_0 + \mathbf{J}_0\times\mathbf{B} - \nabla p + 2\theta\Delta t\mathbf{L}(\mathbf{V}), \tag{9a}$$

$$\frac{\partial \mathbf{B}}{\partial t} = \nabla\times(\mathbf{V}\times\mathbf{B}_0), \tag{9b}$$

$$\frac{\partial p}{\partial t} = -(\mathbf{V}\cdot\nabla p_0 + \gamma p_0\nabla\cdot\mathbf{V}). \tag{9c}$$

For oscillatory modes, the eigenvalues of $\mathbf{L}$ are negative, so that the $-\theta^2\Delta t^2\mathbf{L}(\partial\mathbf{V}/\partial t)$ term on the left-hand side of (9a) effectively adds wavenumber-dependent inertia, while the $2\theta\Delta t\mathbf{L}(\mathbf{V})$ term on the right side introduces dissipation [24]. For growing modes, the eigenvalues of $\mathbf{L}$ are positive, but there is a finite maximum eigenvalue [25].

As discussed in [24], we may devise a numerical scheme based on the alternative differential approximation, Eqs. (9a)–(9c). First, we use the freedom to drop the Δt terms on the right side of (9a) before

discretizing (the equations remain consistent with ideal linear MHD in the limit of small Δt) to avoid numerical dissipation in stable modes. We then stagger $\mathbf{B}$ and p in time from $\mathbf{V}$ to obtain a leap-frog scheme that is numerically stabilized by the $-\Delta t^2\mathbf{L}$ operator, which acts on changes in $\mathbf{V}$. The resulting method is similar to the semi-implicit methods described in [13,17,18], where all fields are predicted and then corrected, resulting in some level of numerical dissipation. Here, the leap-frog aspects are more closely related to the method described in [23], where a von Neumann stability analysis for homogeneous equilibria shows that the magnitude of the numerical amplification factor for the stable modes of $\mathbf{L}$ is unity, i.e. free of numerical dissipation, as long as the θ^2 coefficient (denoted C_0, henceforth) is at least 1/4. For unstable physical modes, the scheme correctly reproduces growth, but Δt must be less than the inverse of the growth rate of the fastest mode to avoid a singularity in the time-derivative terms.

Two modifications of this operator are applied to improve its effectiveness for nonlinear simulations. First, we relax the definition of $\mathbf{L}$ to include the symmetric part of the solution, in addition to the steady-state fields, in $\mathbf{B}_0$, $\mathbf{J}_0$, and p_0, so that the eigenvectors of the operator always correspond closely to the linear modes of the system, which is important for accuracy [23]. Though the combined fields may not be in static force balance, in practice they usually represent a state that is near equilibrium, and the operator can be symmetrized explicitly in its weak form. The second modification, which accounts for the second part of the semi-implicit operator, is to include the isotropic Laplacian operator with a small coefficient to ensure stability as geometrically non-symmetric pressures develop in nonlinear simulations. The coefficient is computed dynamically from the "nonlinear pressure"

$$p_{nl}(R,Z) \equiv \max_{\varphi}\left|\frac{\mathbf{B}^2(R,Z,\varphi)}{\mu_0} + \gamma p(R,Z,\varphi) - \frac{\mathbf{B}_0^2(R,Z)}{\mu_0} - \gamma p_0(R,Z)\right|,$$

which determines the largest variation in the magnetoacoustic wave speed due to asymmetries in φ, the periodic coordinate. This semi-implicit operator is closely related to the one discussed in [13], but the dynamically updated coefficients adapt to fields that change in time. Updating coefficients with the evolution implies re-computing matrices and their factors, but this is done on an as-needed basis rather than at every time-step.

In addition to wave propagation, the time-advance algorithm must address the numerical aspects of advection. For magnetically confined plasmas, we usually encounter flow speeds that are significantly less than the largest wave speeds, so limiting the time-step to satisfy the Courant–Friedrichs–Lewy condition [26] is not prohibitive in many cases of interest. Thus, predictor/corrector steps can be combined with the semi-implicit leap-frog algorithm to stabilize flow without introducing low-order numerical dissipation associated with wave propagation [27]. The complete marching algorithm is comprised of a sequence of operations that is described symbolically by

$$\begin{aligned}
&A_{\mathbf{V}}(\mathbf{V}^{\mathrm{j}}, n^{\mathrm{j}+1/2}, T^{\mathrm{j}+1/2}, \mathbf{B}^{\mathrm{j}+1/2}; \mathbf{V}^{\mathrm{j}}) \mapsto \mathbf{V}_{\mathrm{pre}},\\
&A_{\mathbf{V}}(\mathbf{V}^{\mathrm{j}}, n^{\mathrm{j}+1/2}, T^{\mathrm{j}+1/2}, \mathbf{B}^{\mathrm{j}+1/2}; \mathbf{V}_{\mathrm{pre}}) \mapsto \mathbf{V}^{\mathrm{j}+1},\\
&A_{n}(\mathbf{V}^{\mathrm{j}+1}, n^{\mathrm{j}+1/2}; n^{\mathrm{j}+1/2}) \mapsto n_{\mathrm{pre}},\\
&A_{n}(\mathbf{V}^{\mathrm{j}+1}, n^{\mathrm{j}+1/2}; n_{\mathrm{pre}}) \mapsto n^{\mathrm{j}+3/2},\\
&A_{T}(\mathbf{V}^{\mathrm{j}+1}, n^{\mathrm{j}+3/2}, T^{\mathrm{j}+1/2}, \mathbf{B}^{\mathrm{j}+1/2}; T^{\mathrm{j}+1/2}) \mapsto T_{\mathrm{pre}},\\
&A_{T}(\mathbf{V}^{\mathrm{j}+1}, n^{\mathrm{j}+3/2}, T^{\mathrm{j}+1/2}, \mathbf{B}^{\mathrm{j}+1/2}; T_{\mathrm{pre}}) \mapsto T^{\mathrm{j}+3/2},\\
&A_{\mathbf{B}}(\mathbf{V}^{\mathrm{j}+1}, \mathbf{B}^{\mathrm{j}+1/2}; \mathbf{B}^{\mathrm{j}+1/2}) \mapsto \mathbf{B}_{\mathrm{pre}},\\
&A_{\mathbf{B}}(\mathbf{V}^{\mathrm{j}+1}, \mathbf{B}^{\mathrm{j}+1/2}; \mathbf{B}_{\mathrm{pre}}) \mapsto \mathbf{B}^{\mathrm{j}+3/2},
\end{aligned}$$

where superscripts denote the time-level for each field ($t^{\mathrm{j}} = \mathrm{j}\Delta t$ for constant-Δt computations and $t^{\mathrm{j}+1/2} = t^{\mathrm{j}} + \Delta t/2$), and the "pre" subscript denotes a prediction. The last argument of each operation indicates whether advective terms (such as $-\rho\mathbf{V}\cdot\nabla\mathbf{V}$ that appears in $A_{\mathbf{V}}$) are computed from the solution at the previous time-level for a predictor step, or from the predicted field for a corrector step. (Details for each advance are provided below in Eqs. (12)–(15).) The choice of predictor/correct advection over upwind methods simplifies the implementation with the finite element representation.

Advancing the semi-implicit leap-frog scheme with predictor/corrector advection requires the solution of algebraic systems for each advance in the marching algorithm. Besides the semi-implicit operator, which is part of the $A_{\mathbf{V}}$) operation, the spatial representation described in Section 3.2 leads to mass matrices, and dissipation terms are computed implicitly. Using implicit dissipation is particularly important for thermal conduction, where parallel transport is typically the fastest behavior in the system. Wave propagation is also much faster than nonlinear tearing behavior. Consequently, the matrices for advancing velocity and temperature are ill-conditioned. These linear systems must be solved with sufficient numerical precision to accurately reproduce eigenvectors associated with small eigenvalues, because they represent the slow and physically relevant behavior. In the other equations, the implicit dissipation terms typically have small coefficients and introduce no computational penalty; the mass matrices already necessitate solution of algebraic systems.

3.2. Spatial representation

A finite-dimensional spatial representation is achieved through a basis function expansion and a weak form of the marching equations that is equivalent to a collection of variational problems. The choice of basis functions and the selection of physical fields to expand are central issues for this approach. Using 2D Lagrange-type finite elements enables representation of arbitrarily shaped regions of the poloidal plane, and the basis functions have a sufficient level of continuity for a conforming approximation of the non-ideal MHD Eqs. (4a)–(4f), (5). For the remaining direction, which is periodic, the finite Fourier series is an appropriate expansion. We express the collection of variational problems in cylindrical coordinates (R, Z, φ) for toroidal and cylindrical geometries or in Cartesian coordinates (x, y, z) for straight configurations with a periodic z-coordinate. Nonuniform meshing in the physical poloidal coordinates (R, Z or x, y) is accomplished through mappings from element coordinates [5].

Choosing flow velocity, magnetic field, particle number density and temperature as the fields to expand, our finite-dimensional solution space ($S_{h,N,\mathrm{p}}$) is the product space composed of all functions $\mathbf{v} \in \mathbf{V}_{h,N,\mathrm{p}}$, $\mathbf{b} \in \mathbf{B}_{h,N,\mathrm{p}}$, $n \in n_{h,N,\mathrm{p}}$, and $T \in T_{h,N,\mathrm{p}}$ that satisfy the essential conditions for the system, i.e. the respective Dirichlet boundary conditions discussed in Section 2. The subscripts denote the measure of the poloidal mesh spacing (h), the largest Fourier index (N), and the polynomial degree of the finite element basis functions (p). These parameters identify a particular space $S_{h,N,\mathrm{p}}$ from the family of all such spaces. Members of the $\mathbf{V}_{h,N,\mathrm{p}}$ and $\mathbf{B}_{h,N,\mathrm{p}}$ spaces have the expansion

$$\mathbf{A}_{h,N,\mathrm{p}}(R,Z,\varphi) = \sum_{\mathrm{i},v} a_{\mathrm{i}v,\mathrm{n}=0}\bar{\alpha}_{\mathrm{i}v\mathrm{n}=0} + \sum_{\mathrm{i},v,\mathrm{n}} (a_{\mathrm{i}v\mathrm{n}}\bar{\alpha}_{\mathrm{i}v\mathrm{n}} + a^{*}_{\mathrm{i}v\mathrm{n}}\bar{\alpha}^{*}_{\mathrm{i}v\mathrm{n}}), \tag{10a}$$

while members of $n_{h,N,\mathrm{p}}$ and $T_{h,N,\mathrm{p}}$ have the expansion

$$F_{h,N,\mathrm{p}}(R,Z,\varphi) = \sum_{\mathrm{i}} \mathrm{f}_{\mathrm{i},\mathrm{n}=0}\alpha_{\mathrm{i},\mathrm{n}=0} + \sum_{\mathrm{i},\mathrm{n}} (\mathrm{f}_{\mathrm{in}}\alpha_{\mathrm{in}} + \mathrm{f}^{*}_{\mathrm{in}}\alpha^{*}_{\mathrm{in}}). \tag{10b}$$

The vector and scalar basis functions in Eqs. (10a), (10b) are

$$\bar{\alpha}_{\mathrm{i}v\mathrm{n}} \equiv \hat{\mathbf{e}}_v(\varphi)\psi_{\mathrm{i}}(\xi_1,\xi_2)\exp(\mathrm{in}\varphi) \tag{11a}$$

and

$$\alpha_{\mathrm{in}} \equiv \psi_{\mathrm{i}}(\xi_1, \xi_2)\exp(\mathrm{in}\varphi), \tag{11b}$$

where ψ_{i} is a 2D polynomial basis function of degree p in the element coordinates ξ_1 and ξ_2. The Fourier components have indices $\mathrm{n} = 0, 1, \ldots, N$, and the direction vectors have $v = R, Z, \varphi$ for cylindrical coordinates or $v = x, y, z$ for Cartesian coordinates. Henceforth, a subscripted capital such as $\mathbf{V}_{h,N,\mathrm{p}}$ denotes a function space of the form (10a) or (10b) that is characterized by h, N, and p, whereas the small-case $\mathbf{v}$ denotes an individual function in $\mathbf{V}_{h,N,\mathrm{p}}$; an exception is made for temperature functions (T) to avoid confusion with time.

The inverse of the transformation $R(\xi_1, \xi_2), Z(\xi_1, \xi_2)$ within each finite element is implied in Eqs. (10a), (10b). For many simulations, we use a topologically polar mesh of quadrilateral elements (for example, see Fig. 1(a)), where the left side of the logically rectangular mesh is mapped to the (R, Z) coordinates of the magnetic axis of the steady-state fields. In cases with relatively uniform mesh spacing, we define the transformation with bicubic splines of R and Z in global mesh coordinates that coincide with the local element coordinates within each quadrilateral element, except for an offset that is unique to each element. For bilinear and biquadratic ψ_{i} ($\mathrm{p} = 1, 2$, respectively), this mapping is superparametric, i.e. the mapping is of higher order than the representation of the solution fields, and a sufficient condition for convergence is not met [5]. However, for simulations with smoothly varying mesh spacing, we find better accuracy than with lower-order mappings for the same mesh. We also expand the steady-state fields with bicubic splines in these cases. The splines are susceptible to overshoot with strong mesh packing, however, because derivatives with respect to the logical coordinates change abruptly. Where strong mesh packing is applied, we use isoparametric mappings for R and Z, and the steady-state fields are interpolated with polynomials of the same degree in the element coordinates.

The physical coordinates in Eqs. (10a), (10b) have been expressed as cylindrical coordinates for toroidal and cylindrical geometry. Taking $R \to x, Z \to y, \varphi \to 2\pi z/L_z$ makes the representation suitable for computing in Cartesian coordinates where boundary conditions at $z = 0$ and $z = L_z$ are periodic. Terms involving derivatives with respect to the periodic coordinate and those resulting from cylindrical curvature have been implemented to allow computation with either coordinate system. The implementation of

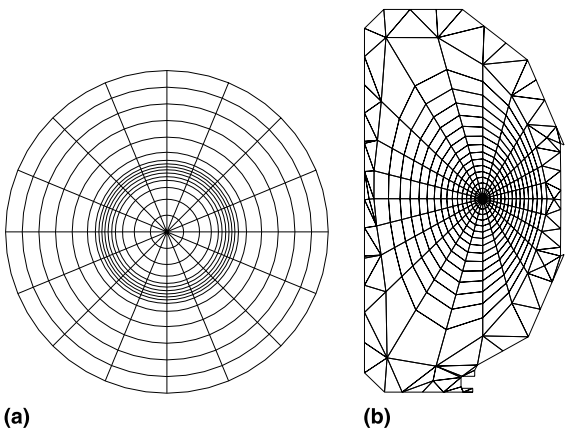

Fig. 1. Finite element meshes illustrating (a) radial packing with piecewise bicubic mapping from a logically rectangular mesh of quadrilateral elements to a topologically polar arrangement of nodes and (b) combining regions of triangular and quadrilateral elements, using linear mappings, to align with equilibrium magnetic flux surfaces in a tokamak interior while matching to a realistic experimental cross-section at the boundary.

regularity conditions for cylindrical configurations (where the domain includes $R = 0$) is discussed in Appendix A.

Using test functions from the same space as the solution fields, $\{\mathbf{w}^{\mathrm{j}}, \mathbf{c}^{\mathrm{j}+1/2}, q^{\mathrm{j}+1/2}, \Theta^{\mathrm{j}+1/2}\} \in S_{h,N,\mathrm{p}}$, produces a Galerkin approximation that is equivalent to a variational problem for each step in our time-advance. Starting with flow velocity, denoting with $\Delta\mathbf{v}$ either the predictor increment $\Delta\mathbf{v}_{\mathrm{pre}}$ or the corrector increment $\Delta\mathbf{v}_{\mathrm{cor}}$, and suppressing the steady-state fields for simplicity, we find $\Delta\mathbf{v} \in \mathbf{V}_{h,N,\mathrm{p}}$ that satisfy

$$\begin{aligned}
&\int \mathrm{d}\mathbf{x}\Bigg\{ \rho^{\mathrm{j}+1/2}\mathbf{w}^* \cdot \Delta\mathbf{v}. + C_0\Delta t^2\left[\frac{1}{\mu_0}\nabla\times(\mathbf{w}^*\times\mathbf{B}_0)\cdot\nabla\times(\Delta\mathbf{v}\times\mathbf{B}_0) + \gamma p_0(\nabla\cdot\mathbf{w}^*)(\nabla\cdot\Delta\mathbf{v})\right] \\
&\quad - \frac{C_0\Delta t^2}{2}[\mathbf{w}^*\cdot\mathbf{J}_0\times\nabla\times(\Delta\mathbf{v}\times\mathbf{B}_0) + \Delta\mathbf{v}\cdot\mathbf{J}_0\times\nabla\times(\mathbf{w}^*\times\mathbf{B}_0)] \\
&\quad + \frac{C_0\Delta t^2}{2}[(\nabla\cdot\mathbf{w}^*)\Delta\mathbf{v}\cdot\nabla p_0 + (\nabla\cdot\Delta\mathbf{v})\mathbf{w}^*\cdot\nabla p_0] \\
&\quad + C_1 p_{nl}\Delta t^2(\nabla\mathbf{w}^*)^{\mathrm{T}} : (\nabla\Delta\mathbf{v}) + g\Delta t\rho^{\mathrm{j}+1/2}\nu(\nabla\mathbf{w}^*)^{\mathrm{T}} : (\nabla\Delta\mathbf{v})\Bigg\} \\
&\qquad = \int \mathrm{d}\mathbf{x}\,\Delta t\Bigg\{ -\rho^{\mathrm{j}+1/2}\mathbf{w}^*\cdot(\bar{\mathbf{v}}\cdot\nabla\bar{\mathbf{v}}) + \frac{1}{\mu_0}\mathbf{w}^*\cdot(\nabla\times\mathbf{b}^{\mathrm{j}+1/2})\times\mathbf{b}^{\mathrm{j}+1/2} \\
&\qquad\quad - \mathbf{w}^*\cdot\nabla p^{\mathrm{j}+1/2} - \rho^{\mathrm{j}+1/2}\nu(\nabla\mathbf{w}^*)^{\mathrm{T}} : (\nabla\mathbf{v}^{\mathrm{j}})\Bigg\}
\end{aligned} \tag{12}$$

for all $\mathbf{w} \in \mathbf{V}_{h,N,\mathrm{p}}$. The new flow velocity is then $\mathbf{v}^{\mathrm{j}+1} = \mathbf{v}^{\mathrm{j}} + \Delta\mathbf{v}_{\mathrm{cor}}$. In (12), $p^{\mathrm{j}+1/2}$ is treated as a nodal quantity, i.e. coefficients of $n^{\mathrm{j}+1/2}$ and $T^{\mathrm{j}+1/2}$ are multiplied and $p^{\mathrm{j}+1/2}$ is interpolated from the resulting products. In addition, the predictor/corrector advection uses $\bar{\mathbf{v}} = \mathbf{v}^{\mathrm{j}}$ for the predictor step and $\bar{\mathbf{v}} = \mathbf{v}^{\mathrm{j}} + f\Delta\mathbf{v}_{\mathrm{pre}}$ for the corrector step with the centering coefficient f. The parameter g is used to control the temporal differencing of the dissipation terms, and we consider implicit differencing with $0.5 \leqslant g \leqslant 1$. For particle number density, we have

$$\int \mathrm{d}\mathbf{x}\{q^*\Delta n + Dg\Delta t(\nabla q^*)\cdot(\nabla\Delta n)\} = \int \mathrm{d}\mathbf{x}\,\Delta t\{(\nabla q^*)\cdot\bar{n}\mathbf{v}^{\mathrm{j}+1} - D(\nabla q^*)\cdot(\nabla n^{\mathrm{j}+1/2})\} \tag{13}$$

for all $q \in n_{h,N,\mathrm{p}}$, where $\bar{n} = n^{\mathrm{j}+1/2}$ for the predictor step and $\bar{n} = n^{\mathrm{j}+1/2} + f\Delta n_{\mathrm{pre}}$ for the corrector step. For the temperature advance, we have

$$\begin{aligned}
&\int \mathrm{d}\mathbf{x}\Bigg\{ \frac{n}{\gamma-1}\Theta^*\Delta T + g\Delta t(\nabla\Theta^*)\cdot n[\chi_\parallel\hat{\mathbf{b}}\hat{\mathbf{b}} + \chi_\perp(\mathbf{I}-\hat{\mathbf{b}}\hat{\mathbf{b}})]\cdot\nabla\Delta T\Bigg\} \\
&\qquad = \int \mathrm{d}\mathbf{x}\,\Delta t\Bigg\{ -\frac{n}{\gamma-1}\Theta^*\mathbf{v}^{\mathrm{j}+1}\cdot\nabla\bar{T} - n\bar{T}\Theta^*\nabla\cdot\mathbf{v}^{\mathrm{j}+1} - (\nabla\Theta^*)\cdot n[\chi_\parallel\hat{\mathbf{b}}\hat{\mathbf{b}} + \chi_\perp(\mathbf{I}-\hat{\mathbf{b}}\hat{\mathbf{b}})]\cdot\nabla T^{\mathrm{j}+1/2} + \Theta^* Q\Bigg\}
\end{aligned} \tag{14}$$

for all $\Theta \in T_{h,N,\mathrm{p}}$. Finally, for the magnetic advance, we have

$$\begin{aligned}
&\int \mathrm{d}\mathbf{x}\Bigg\{ \mathbf{c}^*\cdot\Delta\mathbf{b} + g\Delta t\frac{\eta}{\mu_0}(\nabla\times\mathbf{c}^*)\cdot(\nabla\times\Delta\mathbf{b}) + g\Delta t\kappa_{\mathrm{div}\,b}(\nabla\cdot\mathbf{c}^*)(\nabla\cdot\Delta\mathbf{b})\Bigg\} \\
&\qquad = \int \mathrm{d}\mathbf{x}\,\Delta t\Bigg\{ (\nabla\times\mathbf{c}^*)\cdot(\mathbf{v}^{\mathrm{j}+1}\times\bar{\mathbf{b}}) - \frac{\eta}{\mu_0}(\nabla\times\mathbf{c}^*)\cdot\nabla\times\mathbf{b}^{\mathrm{j}+1/2} - \kappa_{\mathrm{div}\,b}(\nabla\cdot\mathbf{c}^*)(\nabla\cdot\mathbf{b}^{\mathrm{j}+1/2})\Bigg\} \\
&\qquad\quad - \Delta t\int \mathrm{d}\mathbf{s}\times\mathbf{E}\cdot\mathbf{c}^*
\end{aligned} \tag{15}$$

for all $\mathbf{c} \in \mathbf{B}_{h,N,\mathrm{p}}$, where the surface term represents the influence of an applied electric field.

The semi-implicit operator occupies most of the left-hand side of Eq. (12), and it includes the Laplacian part for stabilizing wave propagation in geometrically non-symmetric states arising from nonlinear dynamics. For conditions of interest, $p_{nl} \ll \mathbf{B}_0^2/\mu_0 + \gamma p_0$, and accuracy is not sensitive to the value of C_1 if it is large enough ($C_1 \geqslant 1/4$) for numerical stability. The terms with coefficient $C_0\Delta t^2/2$ result from the ideal MHD operator $\mathbf{L}$ defined in Eq. (8), but they are symmetrized explicitly for conditions where the 0-subscript fields are not in equilibrium, as discussed above. Symmetrizing ensures that the semi-implicit operator has real eigenvalues, since the finite element method then produces a Hermitian matrix by construction. If Δt does not exceed the inverse of the growth rate of the most unstable mode of the ideal MHD system, the resulting matrix is also positive-definite. This condition can always be satisfied in initial value computations, and it certainly will be in a temporally converged calculation.

The forces on the right-hand side of Eq. (12) are computed from separate nodal fields for $\mathbf{B}$ and p, unlike the stabilizing corrections to these forces that appear through the semi-implicit operator. To our knowledge, the disparate representation of implicit and explicit terms does not have negative consequences; though, it does influence convergence properties (see Section 6). Early versions of the NIMROD algorithm were based on von Neumann analysis of the differencing equivalent to using bilinear finite elements [7,28]. We found that the numerical dispersion relation for waves in an infinite uniform equilibrium has the shear and compressional branches decoupled to all orders in h and Δt when velocity and magnetic field are discretized; this could not be achieved for formulations based on currents and potentials. Thus, the impact of the inconsistent representation of implicit and explicit terms is strongly dependent on how the system is formulated. (In contrast, second-order operators in finite difference and finite volume methods are usually constructed from first-order operators, avoiding inconsistency. However, preserving the symmetry of complicated operators like $\mathbf{L}$ in general geometry is difficult.)

Nonlinear terms and coefficients that depend on φ require products of Fourier series expansions. We apply a pseudospectral method [16], using the FFT to find data on a uniform grid over the periodic coordinate; however, the Fourier representation is padded with zero coefficients at high wavenumbers to prevent aliasing from quadratic nonlinearities [29]. Algebraic operations are performed on the periodic grid to construct the needed terms, followed by a transform of the result to obtain its Fourier decomposition. To allow computations involving spatial derivatives of the expanded fields (like ∇T), the transforms and pseudospectral computations are performed at the quadrature points for numerical integration. The appearance of φ-dependent coefficients in the left sides of the equations, like the mass density in the flow velocity advance and the magnetic direction vector in the thermal conduction of the temperature advance, leads to convolution matrices that are dense in the Fourier component index. We solve these systems with a matrix-free iterative method, in order to use FFTs in a direct computation of the matrix-vector product, instead of computing convolutions explicitly. For magnetic fusion plasmas, the non-symmetric ($n > 0$) Fourier components of ρ are small and do not have a significant effect on the flow velocity evolution Eq. (4e). The option of dropping the associated small terms expedites computation, since it allows solving $N + 1$ independent 2D linear systems for each velocity update instead of solving one coupled 3D linear system.

The mathematical symmetries that exists in the weak form of the temporal advance and the caveat that Δt is small enough so that all eigenvalues of the left-hand side of Eq. (12) are positive imply equivalence between Eqs. (12)–(15) and a set of variational problems. Furthermore, the solution space $S_{h,N,p}$ is admissible, because all terms on the left-hand sides of Eqs. (12)–(15) are integrable and the essential conditions are enforced. The representation is therefore a conforming approximation, and we can identify the left-hand side of each of Eqs. (12)–(15) as the respective strain energy. We then expect spatial convergence rates that increase with the polynomial degree of the basis function, p, according to Eqs. (2) and (3). However, the terms on the right-hand sides of Eqs. (12)–(15) are produced during the course of the temporal advance. If a calculation tends to create fields that cannot be resolved smoothly, assumptions used in deriving the convergence-rate relations are violated, and globally high-order discretization is not effective.

Adaptive techniques, such as the *hp* finite-element method [30], may be better suited for these conditions. For high temperature plasmas, numerical accuracy requires resolution of the smallest spatial features (tearing layers), so we restrict attention to parameters where all length-scales can be resolved.

The numerical treatment of the magnetic divergence constraint is another central issue for accurate simulation. Re-expressing Eq. (15) with $g=1$ as

$$\int d\mathbf{x}\left\{\mathbf{c}^*\cdot\mathbf{b}^{j+3/2}+\Delta t\frac{\eta}{\mu_0}(\nabla\times\mathbf{c}^*)\cdot(\nabla\times\mathbf{b}^{j+3/2})+\Delta t\kappa_{\mathrm{div}\,b}(\nabla\cdot\mathbf{c}^*)(\nabla\cdot\mathbf{b}^{j+3/2})\right\}$$
$$=\Delta t\int d\mathbf{x}(\nabla\times\mathbf{c}^*)\cdot(\mathbf{v}^{j+1}\times\bar{\mathbf{b}})-\Delta t\oint d\mathbf{s}\times\mathbf{E}\cdot\mathbf{c}^* \tag{16}$$

shows that $\Delta t\kappa_{\mathrm{div}\,b}$ has the role of a Lagrange multiplier for the constraint $(\nabla\cdot\mathbf{b}^{j+3/2})^2=0$ in the variational problem for $\mathbf{b}^{j+3/2}$. If it were necessary to use arbitrarily large values of the product $\Delta t\kappa_{\mathrm{div}\,b}$, our continuous solution space $\mathbf{B}_{h,N,p}$ would not approach a meaningful divergence-free representation in the limit of $h\to 0$, because the formulation does not satisfy *divergence-stability* (see [31] and references therein). As described below in Section 6, the lack of divergence-stability in this case results from imposing too many constraints through the numerical calculation of $(\nabla\cdot\mathbf{b}^{j+3/2})^2$ for the finite number of degrees of freedom in the space [32]. Alternatively, if the value of $\Delta t\kappa_{\mathrm{div}\,b}$ is too small, the constraint is not imposed. In either limit, the "strain energy" represented by the left-hand side of Eq. (16) is a poor norm for choosing the best available solution. For time-dependent problems like the ones considered here, arbitrarily large values of $\Delta t\kappa_{\mathrm{div}\,b}$ are not required to control the generation of error per time-step. The convergence studies presented in Section 4 show that acceptable results are achieved routinely for basis function with $p\geqslant 2$.

Regarding practical considerations, the poloidal mesh is divided into structured blocks of quadrilateral elements and unstructured blocks of triangular elements (see Fig. 1(b)). This organization facilitates domain decomposition for parallel computation and adds geometric flexibility. At this time, the implementation of triangular elements in NIMROD is incomplete (the ψ_i in triangular elements are restricted to linear basis functions), so the results described below consider computations with quadrilateral elements only.

4. Benchmarks and convergence rates

The performance of a numerical algorithm for magnetic fusion applications should be examined in conditions that are sufficiently stiff and anisotropic to represent laboratory plasmas. Since stiffness associated with the rapid propagation of MHD waves arises primarily from linear terms, the linear resistive tearing mode described below is an important benchmark for large-Δt performance. The highly localized nature of the eigenfunction also exercises the treatment of magnetic field divergence error and non-uniform meshing. The second test problem, presented in Section 4.2, provides a quantitative benchmark of anisotropic thermal conduction.

4.1. Linear tearing mode

The domain for our resistive MHD benchmark is a straight cylinder with periodic ends. For a selected helical perturbation ($\sim e^{im\theta+i2\pi nz/L_z}$, where m and n are fixed integers, and L_z is the cylinder length), there exists a concentric cylindrical surface within the domain where the perturbation has constant phase along the equilibrium magnetic field lines, which lie within the surface. The linear MHD response to the perturbation is a resonance (due to anisotropy) such that flows will be local to this surface. However,

resistivity, inertia, and viscosity prevent singular behavior by smoothing spatial scales that are small relative to global length-scales. We have chosen cylindrical geometry for the test to allow comparison of numerical results with an analytic dispersion relation that is valid in the limit of vanishing resistivity. For comparison, we determine the analytic eigenvalue, the matching parameter Δ' resulting from singular perturbation [33,34], by integrating the Euler–Lagrange equations for the helical perturbation [35] in the regions outside the tearing layer. In the pressureless limit, the growth rate for asymptotically small resistivity is then computed from the dispersion relation [34]

$$\gamma = \left(\frac{2\Gamma(5/4)}{\pi\Gamma(3/4)}\Delta'\right)^{4/5}\left[\frac{m^2}{R^2q^4}\left(\frac{\mathrm{d}q}{\mathrm{d}r}\right)^2\right]^{1/5}\left(\frac{B_z}{\sqrt{\mu_0\rho}}\right)^{2/5}\left(\frac{\eta}{\mu_0}\right)^{3/5}, \tag{17}$$

where q is the "safety factor" or magnetic winding number ($2\pi rB_z/L_zB_\theta$ in a periodic cylinder) using equilibrium values at the resonant surface radius (r_s, where $q(r_s) = -\mathrm{m}/\mathrm{n}$). The equilibrium we consider is the pressureless paramagnetic pinch [36] with normalized on-axis current density ($\mu_0 aJ/B$, where a is the cylinder radius) set to 3. The q profile varies from 1.2 on axis to 0.19 at $r = a$ for an aspect ratio $L_z/2\pi a = 5/9$, and resonance for the m = 1, n = −1 perturbation occurs at $r = 0.3859a$ (see Fig. 2). Solving the Euler–Lagrange equations for this equilibrium and resonant surface yields $\Delta' = 6.679$. This value is verified with Fig. 3 of [37] after changing normalization ([37] has J normalized to unity on axis, and a is varied).

The NIMROD computations use the finite element mesh to represent the r–θ plane of the straight periodic cylinder with Fourier representation for the axial direction, so the calculations solve for the Cartesian components of **V** and **B**. The meshes are circular–polar with grid lines running along constant θ-values with uniform spacing and along constant r-values with non-uniform spacing to allow packing near the resonant surface. An example is the 16×16 mesh of bicubic elements with isoparametric mapping shown in Fig. 1(a). The radial mesh spacing as a function of radial cell index is based on the local q-value by defining a discrete cumulative distribution

$$f_\mathrm{i} = \sum_{\mathrm{j}=1}^{\mathrm{i}} 1 + A_p \exp\left\{-\frac{[q(r_\mathrm{j}) - q(r_s)]^2}{W_p^2[q(0) - q(a)]^2}\right\} \quad \text{for i} = 1, 2, \ldots, N_\xi,$$

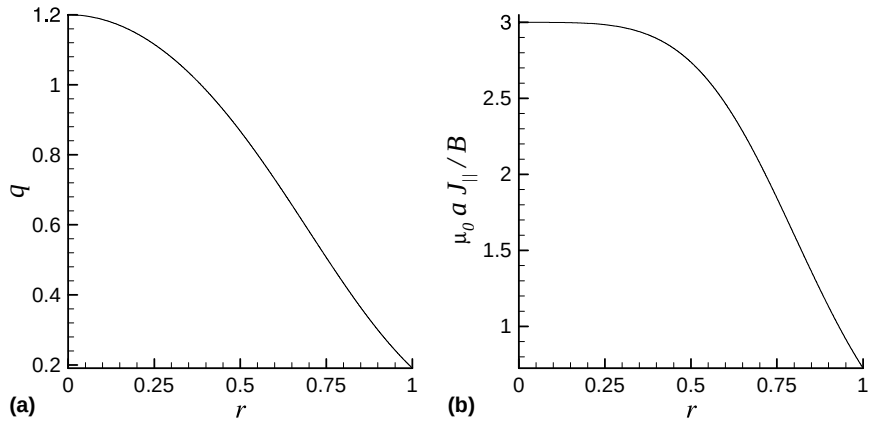

Fig. 2. Profiles of equilibrium *safety factor* or magnetic winding number ($q = 2\pi rB_z/L_zB_\theta$ for the cylinder with periodic length L_z) (a) and normalized parallel current density (b) for the tearing-mode benchmark.

where A_p and W_p are dimensionless parameters that control the magnitude and extent of packing, and r_j, $j = 1, 2, \ldots, N_\xi$ are cell-center locations of a preliminary uniform mesh. We use the f_i-values to define a continuous piecewise linear function of radius that increases from zero to f_1 over the first cell in the uniform mesh, from f_1 to f_2 over the second cell, and so on, reaching f_{N_ξ} at the right side of the mesh. Vertices of the packed mesh are then identified by the radii where the piecewise linear function has the values $k f_{N_\xi}/N_\xi$ for $k = 0, 1, 2, \ldots, N_\xi$. Results for the tearing mode have been computed with $W_p = 0.075$, $5 \leqslant A_p \leqslant 12$, and meshes ranging from 8×8 (with bicubic elements) to 256×256 (with bilinear elements). The resulting mesh spacing changes too abruptly to avoid overshoot with cubic splines, so the mapping and equilibrium field data are interpolated with the same basis functions used for the solution space. For numerical integration, the tests have been completed with nine Gaussian quadrature points per element for bilinear elements, 16 for biquadratic, and 25 for bicubic, which is an additional point per direction relative to what is normally used.

The calculations are run as initial value problems, but only linear terms are included in the time-advance, so the behavior is independent of the perturbation amplitude. The initial flow velocity perturbation is chosen to be smooth and to have non-zero curl to excite the tearing instability, but otherwise, it is arbitrary. The value of kinematic viscosity is chosen to be sufficiently small as to have no significant effect on the computed growth; through experimentation this condition is found to be $Pm \equiv \nu\mu_0/\eta 10^{-3}$ for this mode. We fix the mass density and equilibrium magnetic field to set the Alfvén speed ($v_A \equiv B/\sqrt{\mu_0\rho}$) to 1 m/s on axis, and with $a = 1$ m, the Lundquist number ($S \equiv \mu_0 a v_A/\eta$) is numerically equivalent to the inverse of the electrical diffusivity.

The essential features of the tearing mode are: (1) adherence to the asymptotic analytic scaling $S^{-3/5}$ evident in Eq. (17) and (2) near-singular behavior of the eigenfunction in the vicinity of the resonant surface. Fig. 3 displays computed growth rates on a logarithmic scale to show the asymptotic behavior at large S-values. At the smaller S-values, the tearing layer extends over non-negligible variations in the equilibrium, and the behavior is more diffusive than what is assumed in the asymptotic analytic calculations of [33,34]. The NIMROD results for $S = 10^5 - 10^6$ have been computed with a 32×32 mesh of bicubic elements with $A_p = 5$. At $S = 10^7$, a 48×48 mesh of bicubic elements with $A_p = 8$ resolves the more localized eigenfunction. At $S = 10^8$, a larger mesh of biquadratic elements proves more tractable, and resolution to within 5% of the analytic growth rate is achieved with a 144×144 mesh with $A_p = 12$.

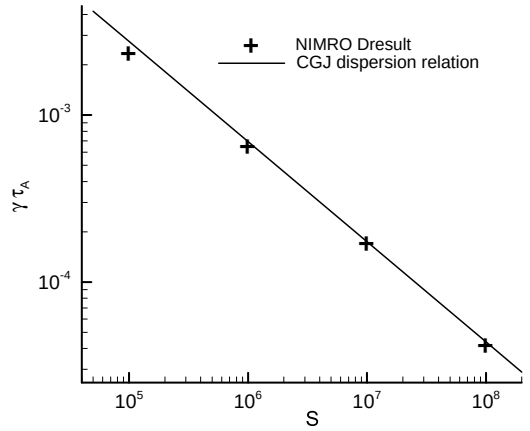

Fig. 3. Comparison of a NIMROD-computed growth-rate scaling for the cylindrical tearing mode with the asymptotic analytic dispersion relation, Eq. (17).

Flow velocity components of the eigenfunction for $S = 10^6$ computed with the 32×32 mesh of bicubic elements show the localized response associated with the resonant surface (see Fig. 4). Although the growth rate is converged with respect to spatial resolution and at $\Delta t = 100\tau_A$ is accurate to within 2% of the temporally converged value, there are azimuthal variations in the axial velocity projection evident at the scale of the mesh (Fig. 4(c)). These variations are reduced when the computation is performed with more elements in the azimuthal direction or by reducing the value of Δt, so the fully converged solution with $p = 3$ is free of the error. However, similar computations with (a) a 48×48 mesh of biquadratic elements, (b) a 24×24 mesh of biquartic elements, and (c) a 20×20 mesh of biquintic elements – all with roughly the same amount of data as the 32×32 bicubic computation – show no azimuthal variations (see Fig. 4(d)), and at $\Delta t = 100\tau_A$, there is only a 0.3% variation among the computed growth rates.

Spatial convergence properties with respect to the mode growth rate at $S = 10^6$ for biquadratic and bicubic elements are shown in Fig. 5. For each calculation, the numbers of elements in the radial and azimuthal directions are identical, and the mesh-packing parameters A_p and W_p are kept fixed as the number

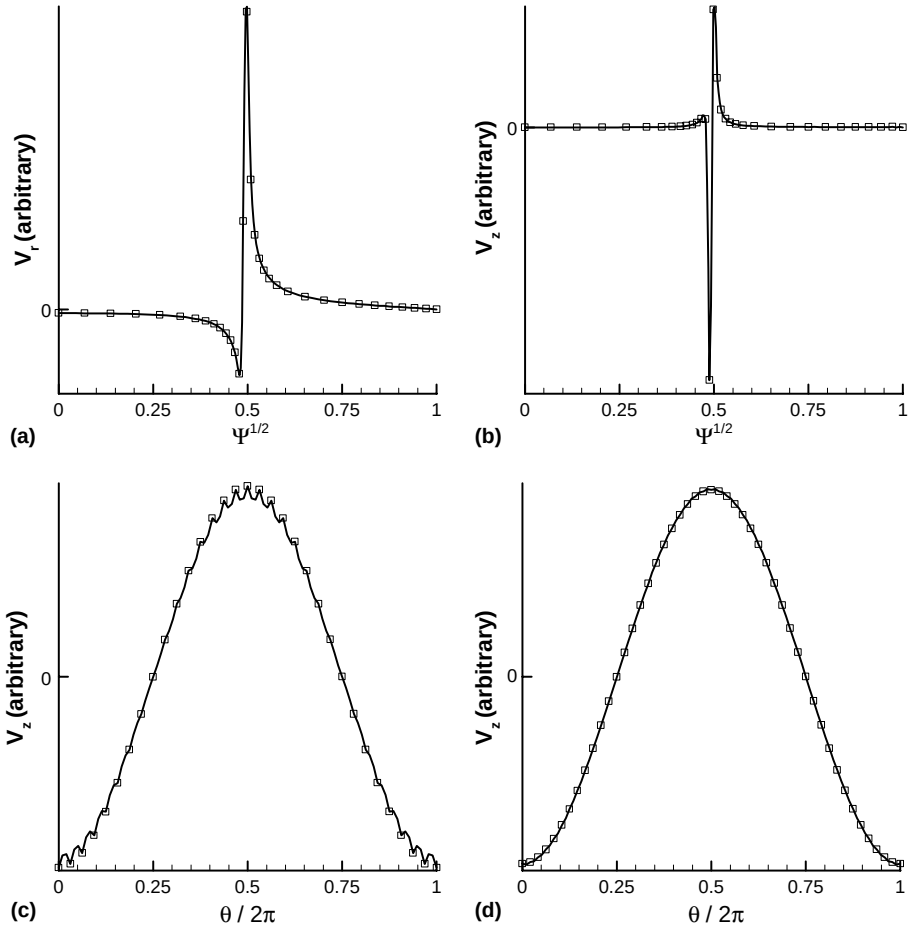

Fig. 4. Eigenfunctions for the cylindrical tearing-mode computed with $\Delta t = 100\tau_A$. Profiles are plotted in the radial coordinate ($\sqrt{\Psi}$) in (a)–(b), where Ψ is the normalized poloidal flux function for the equilibrium, and in the azimuthal angle at a fixed radial coordinate just outside the resonance in (c)–(d). Frames (a)–(c) show the eigenfunction resulting from a 32×32 mesh of bicubic elements, whereas (d) shows the result from a 48×48 mesh of biquadratic elements. Boxes indicate the locations of element-boundary nodes.

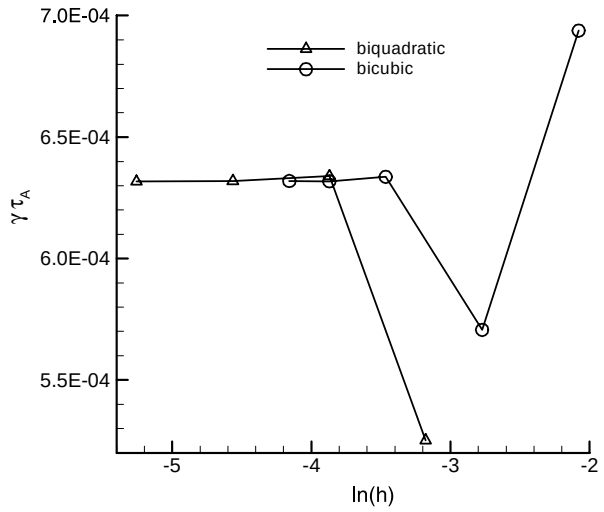

Fig. 5. Convergence of the computed cylindrical tearing-mode growth rate with mesh spacing for biquadratic and bicubic finite elements with $S = 10^6$, and $\Delta t = 100\tau_A$. The parameter h is the inverse of the number of quadrilateral finite elements in each of the radial and azimuthal directions.

of elements is varied. Clearly, convergence to within 1–2% is quite rapid with $p \geqslant 2$ basis functions. In comparison, the growth rate for a 256×256 bilinear mesh with $A_p = 10$ and otherwise similar parameters (not shown in Fig. 5) is in error by more than 25%. Given that the temporal resolution is sufficient for the analytical dispersion relation (17) to describe the growth of the numerical solution from time-step to time-step, we expected to see errors decreasing like h^p by the following argument: The spatial distribution of the computed eigenfunction enters Eq. (17) through the matching parameter Δ'. Furthermore, the Δ'-value computed from the numerical solution can be described as $\Delta'_a + E(h)$, where the error, E, goes to zero as h is reduced and the computed Δ' approaches its analytical value, Δ'_a. For small h, the resulting growth rate $\gamma(h) \sim [\Delta'_a + E(h)]^{4/5} \cong \Delta'^{4/5}_a[1 + (4/5)E(h)/\Delta'_a]$, so the error in the computed growth rate is proportional to the error in the computed matching parameter. Noting that $\Delta' = [(\mathrm{d}b_r/\mathrm{d}r)_{r_s+} - (\mathrm{d}b_r/\mathrm{d}r)_{r_s-}]/(b_r)_{r_s}$ depends on derivatives of the eigenfunction on either side of the resonant surface, we expect to observe $\gamma(h) - \gamma(0) \sim h^p$, i.e. the rate given by finite element analysis for the convergence of first derivatives, Eq. (3). The results for biquadratic and bicubic elements show faster convergence in this test. For example, the biquadratic series of computations for 48, 96, and 192 elements per direction shows $\gamma(h) - \gamma(0) \sim h^{3.2}$.

Performance with respect to the magnetic divergence constraint is more easily related to finite element analysis. In Fig. 6, we plot the 2-norm of the error vs. h on a log–log scale for the biquadratic and bicubic calculations represented in Fig. 5 and for three bilinear computations. As h is decreased, the convergence rate for each basis approaches the value of p, consistent with Eq. (3). In all of these cases, $\Delta t = 100$ and $\kappa_{\mathrm{div}\,b} = 0.1$, where the value of $\kappa_{\mathrm{div}\,b}$ has been chosen to achieve an acceptable error level for the computation with the coarsest mesh, the 8×8 mesh of bicubic elements.

Since the diffusivity $\kappa_{\mathrm{div}\,b}$ is numerical, a result is not converged unless it is insensitive to the $\kappa_{\mathrm{div}\,b}$-value. Therefore, achieving this independence readily as h is reduced is a desirable property for the algorithm. To determine the sensitivity in the tearing-mode calculations, we have varied $\kappa_{\mathrm{div}\,b}$ in computations with different basis functions. The resulting growth rate and magnetic divergence error for a 128×128 bilinear mesh, a 48×48 biquadratic mesh, and a 32×32 bicubic mesh are plotted in Fig. 7. The broad range of $\kappa_{\mathrm{div}\,b}$-values producing the same growth rate for the biquadratic and bicubic cases provides confidence that the error diffusion approach leads to a good strain energy norm for the magnetic advance when $p \geqslant 2$. In

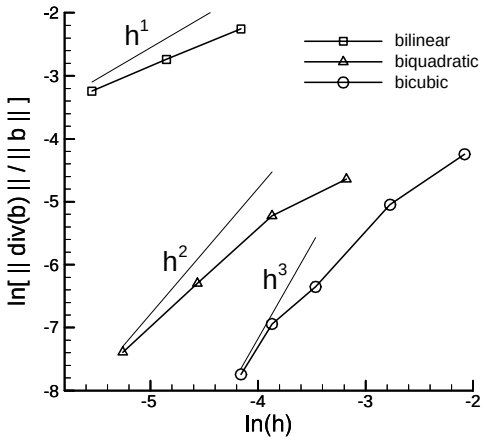

Fig. 6. Logarithm of the normalized magnetic divergence error, $a\sqrt{\int \mathrm{d}\mathbf{x}(\nabla\cdot\mathbf{b})^2/\int \mathrm{d}\mathbf{x}\,\mathbf{b}^2}$, in the computed cylindrical tearing-mode eigenfunctions at $S = 10^6$ and $\Delta t = 100\tau_A$, plotted with respect to mesh spacing for bilinear, biquadratic, and bicubic finite elements. The exponential dependences, h^p, expected of the error in a spatial derivative of the solution, according to Eq. (3), are shown for comparison.

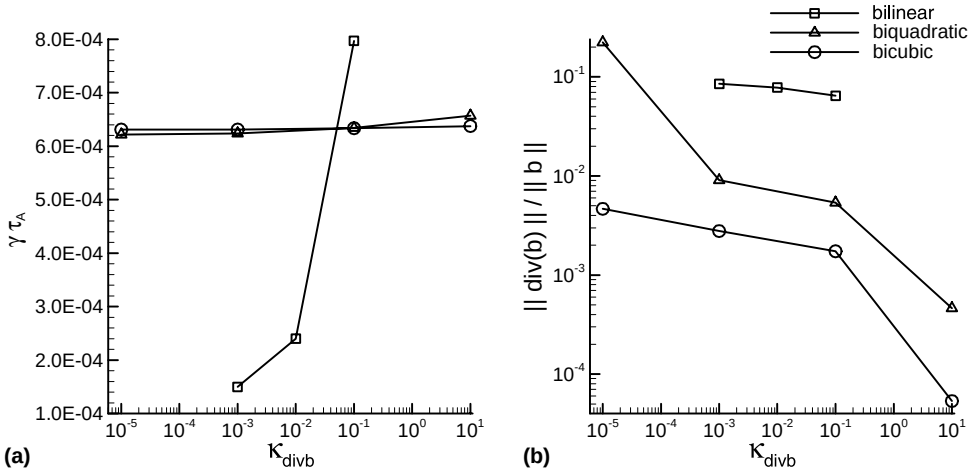

Fig. 7. Growth rate (a) and normalized magnetic divergence error (b) in the computed cylindrical tearing-mode eigenfunctions at $S = 10^6$ and $\Delta t = 100\tau_A$, plotted with respect to the error diffusivity parameter, $\kappa_{\mathrm{div}\,b}$, for 128×128 bilinear, 48×48 biquadratic, and 32×32 bicubic finite element meshes.

contrast, the sensitivity of the bilinear result to the $\kappa_{\mathrm{div}\,b}$-value implies proximity between conditions where the error diffusion term is insufficient to control the error and conditions where the term imposes too many constraints. However, we note that while the performance of bilinear elements is poor in this test, they have been used effectively in simulations with larger levels of physical dissipation.

The last set of computations for the tearing-mode problem considers a range of time-step values to examine temporal convergence properties. The computed growth rates for $S = 10^6$ are plotted in Fig. 8 as a function of $\gamma_0\Delta t$, where γ_0 is the converged value, with dissipation terms evaluated as centered and forward approximations with respect to the time-step (setting the parameter g of Eqs. (12)–(15) to 0.5 and 1, respectively). All of the results shown in Fig. 8 are within 10% of the converged value, but it is possible to

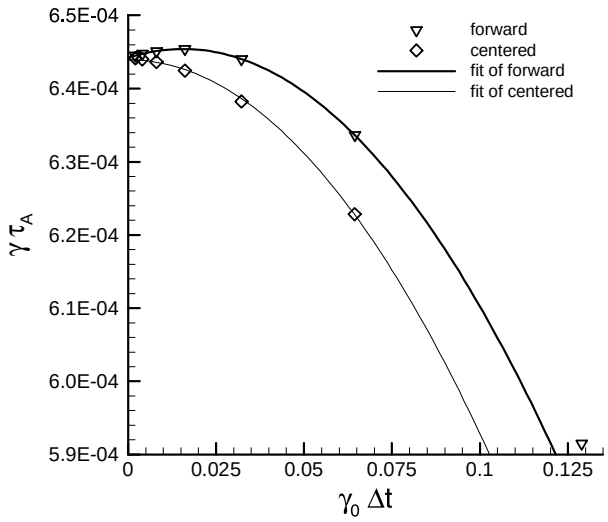

Fig. 8. Convergence of the computed cylindrical tearing-mode growth rate with time-step for forward and centered approximations of the dissipation terms. The spatial representation is a 32×32 mesh of bicubic finite elements, and the S-value is 10^6. The horizontal axis is normalized with the converged growth rate γ_0, and the polynomial fits $6.44 \times 10^{-4} + 1.54 \times 10^{-4}(\gamma_0 \Delta t) - 4.94 \times 10^{-3}(\gamma_0 \Delta t)^2$ and $6.44 \times 10^{-4} - 5.11 \times 10^{-3}(\gamma_0 \Delta t)^2$ for the forward and centered approximations, respectively, have been computed with the numerical results from $\gamma_0 \Delta t \leqslant 0.0644$.

distinguish different asymptotic behavior for the two approximations. Results with the centered approximation are well fit by the quadratic $\gamma_{\text{centered}} \cong 6.44 \times 10^{-4} - 5.11 \times 10^{-3}(\gamma_0 \Delta t)^2$, demonstrating second-order convergence, whereas a linear term is needed to fit the non-monotonic behavior of the forward approximation, $\gamma_{\text{forward}} \cong 6.44 \times 10^{-4} + 1.54 \times 10^{-4}(\gamma_0 \Delta t) - 4.94 \times 10^{-3}(\gamma_0 \Delta t)^2$. The linear term dominates the error in the forward approximation only for $\gamma_0 \Delta t \leqslant 0.03$, and a transition to quadratic behavior occurs where the computed growth rate is still quite accurate. Thus, the truncation error from the dissipation terms has only a small effect on the accuracy in this representative calculation, where the physical conditions are nearly dissipation-free. Temporal convergence is primarily determined by the numerical method used for the large ideal terms – the leap-frog method stabilized by the semi-implicit operator. Since the semi-implicit operator enters with a Δt^2 coefficient (see Section 3.1), the method retains the second-order accuracy of the basic leap-frog method. Forward approximation of the dissipation terms is routinely used in nonlinear NIMROD simulations to provide damping for all wavenumbers that are represented, unlike time-centered dissipation.

4.2. Anisotropic thermal conduction

Eq. (5) for diffusive heat flux reproduces rapid equilibration along magnetic field lines and relatively slow energy transport across magnetic flux surfaces when the ratio of thermal conductivities, $\chi_\parallel/\chi_\perp$, is large. In numerical computations with this model, truncation errors in the temperature gradient are multiplied by the parallel conductivity, and the resulting heat flux errors tend to produce artificial perpendicular transport that can be attributed to the misalignment of $\mathbf{B}$ and the computed ∇T. For our representation, temperature is expanded in the form of (10b), and continuity at the interfaces between elements is not enforced for spatial derivatives. Thus, ∇T is a piecewise continuous vector field that in general has discontinuity along the element interfaces. With a continuous and therefore different representation of magnetic field, such as an expansion in the form of (10a), numerical computations of parallel thermal

equilibration do not reproduce $\mathbf{B}\cdot\nabla T = 0$ everywhere unless the distribution of $\mathbf{B}$ is restricted. (For example, the gradient of the expanded temperature field can satisfy $\mathbf{B}\cdot\nabla T = 0$ everywhere if $\mathbf{B}$ is uniform.) The most challenging conditions for numerical computation arise when $\chi_\parallel/\chi_\perp \gg 1$ and $\mathbf{B}$ has a general distribution with its direction varying across the mesh – conditions that often occur in the simulation of nonlinear MHD activity in high-temperature plasmas. Here, we devise a test that measures the effect of artificial heat transport in finite element computations and use it to show that high-order elements can satisfy $\mathbf{B}\cdot\nabla T = 0$ with sufficient accuracy for realistic $\chi_\parallel/\chi_\perp$-ratios.

While it is possible to evaluate the convergence of $\mathbf{B}\cdot\nabla T$ in a sequence of numerical computations, the impact of the truncation error on thermal transport when computing with realistic $\chi_\parallel/\chi_\perp$-ratios is the more pertinent issue for time-dependent MHD simulations with evolving temperature and pressure profiles. Thus, our test problem has been devised so that the effective perpendicular thermal conductivity, including conduction from numerical errors, can be easily measured from the resulting temperature distribution. The domain is the unit square, $-0.5 \leqslant x \leqslant 0.5$, $-0.5 \leqslant y \leqslant 0.5$, and homogeneous Dirichlet boundary conditions are imposed on T along the entire boundary. The source $Q = 2\pi^2\cos(\pi x)\cos(\pi y)$ is used in the temperature evolution equation to drive the lowest eigenmode of the configuration, and a charge current density di-

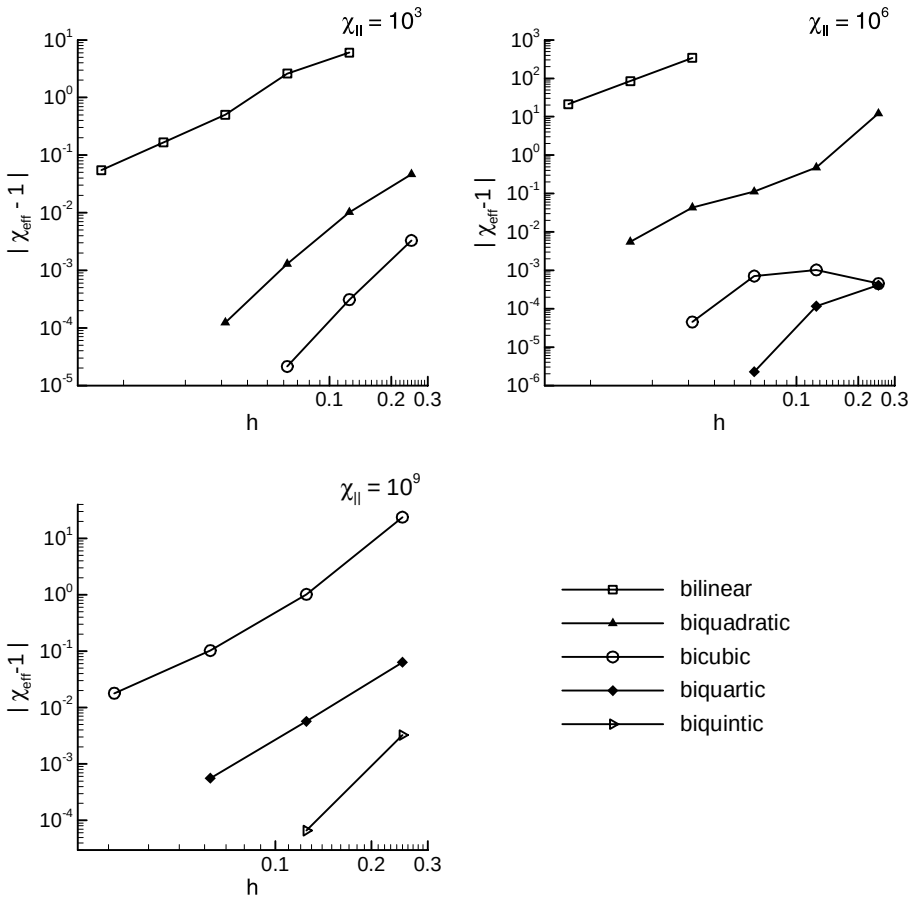

Fig. 9. Error in the effective perpendicular diffusivity resulting in the anisotropic thermal conduction test for the indicated values of parallel diffusivity. Results are plotted with respect to the (uniform) mesh spacing for bilinear through biquintic finite elements, as labeled.

rected out of the x–y plane is induced by an electric field that has the same spatial dependence as the heat source. An extremely large mass density prevents MHD motions, so that diffusive behavior dominates. Analytically, the resulting magnetic field is everywhere tangent to the contours of constant temperature in the solution for isotropic ($\chi_\parallel = \chi_\perp$) thermal conduction, $T(x,y) = \chi_\perp^{-1}\cos(\pi x)\cos(\pi y)$, so anisotropic conditions with $\chi_\parallel \gg \chi_\perp$ have the same solution. In numerical computations, transport that is artificially enhanced by truncation errors in the heat flux vector can lead to a maximum temperature that is less than $\chi_\perp^{-1}$, even when the alignment of $\mathbf{B}$ and the computed ∇T appears acceptable. By setting the $\chi_\perp$-coefficient to unity in the calculations, we arrange the problem so that the computed steady-state value of $T^{-1}(0,0)$ provides a direct measure of the resulting effective perpendicular conductivity including truncation error. As a guide, errors of order 10^{-2} would normally be considered acceptable for nonlinear simulations. Note that the magnetic field direction varies within the domain, so computations with a simple rectangular mesh provide a rigorous test.

To study convergence properties, the conduction problem is run to steady state with $\chi_\parallel/\chi_\perp$-ratios of 10^3, 10^6, and 10^9 with a range of mesh sizes and basis function p-values. Numerical integration for the finite elements is performed with the standard number of Gaussian quadrature points for a given basis (4 for $\mathrm{p} = 1$, 9 for $\mathrm{p} = 2$, etc.). The resulting error in perpendicular diffusivity, $|T^{-1}(0,0) - 1|$, is plotted in Fig. 9. The accuracy and convergence rate improve substantially with p for this problem, where the solution is a smooth function of position. Convergence rates approach the values predicted by Eq. (2) for $\chi_\parallel/\chi_\perp$-ratios of 10^3 and 10^6. For $\chi_\parallel/\chi_\perp = 10^9$, the obtained convergence rates are slightly less than the predictions. Nonetheless, we find that elements with $\mathrm{p} \geqslant 3$ can meet a sufficient level of accuracy in these extreme but laboratory-plasma-relevant conditions, whereas bilinear elements struggle at $\chi_\parallel/\chi_\perp = 10^3$ and are entirely inadequate at $\chi_\parallel/\chi_\perp = 10^6$. A realistic application including 3D magnetic topology is considered in the following section and confirms the effectiveness of the high-order spatial representation.

5. Nonlinear tearing evolution

As an example of a nonlinear simulation in stiff conditions with large anisotropy, we consider a resistive tearing mode in a toroidal MHD equilibrium with non-circular cross-section, tokamak safety-factor profile, and aspect ratio $R/a = 3$ (see Fig. 10). A vanishingly small value of plasma-beta ($\beta \equiv 2\mu_0 p/B^2$) has been chosen to prevent stabilization of the current-driven mode [38]. In these conditions, the internal energy evolution serves as a measure of confinement properties, but it does not play a role in the MHD activity. The mode, while in its linear stage, is then similar to the cylinder mode described in Section 4.1. The primary distinguishing feature is coupling among poloidal harmonics due to toroidal geometry and the shaped cross-section. Responses that are resonant at surfaces with different rational q-values are coupled if they have the same toroidal Fourier index, n. Other parameters for the simulation are: $n_{\mathrm{ss}} = 10^{20}$ m^{-3}, $\tau_{\mathrm{A}} = 1$ μs, $S = 10^6$, $Pm = 0.1$, $\chi_\perp = 42$ m^2 s$^{-1} = 100\eta/\mu_0$, and $\chi_\parallel = 4.2\times 10^7$ m^2 s^{-1}. Here, the Alfvén time is defined as $\tau_{\mathrm{A}} \equiv q(0)R\sqrt{\mu_0\rho}/B_{\varphi_{\mathrm{vac}}}$, where the denominator is the value of the corresponding vacuum toroidal magnetic field at the geometric center of the cross-section. The numerical particle diffusivity is set to the same value as the perpendicular thermal diffusivity, $D = \chi_\perp$, and for controlling divergence error, $\kappa_{\mathrm{div}\,b} = 100$ m^2 s^{-1}.

Since the tearing mode is the only MHD instability of the equilibrium, we first run a linear computation for the $\mathrm{n} = 1$ toroidal Fourier harmonic. The resulting eigenmode, plotted in Fig. 11, shows coupling from the dominant $\mathrm{m} = 2$ poloidal harmonic to the $\mathrm{m} = 3$ and $\mathrm{m} = 4$ harmonics, and the computed growth rate is $4.72 \times 10^{-4}\tau_{\mathrm{A}}^{-1}$. The nonlinear simulation has toroidal resolution $0 \leqslant \mathrm{n} \leqslant 2$, and the $\mathrm{n} = 1$ eigenmode from the linear computation is used as the initial condition with its amplitude adjusted to create a small but

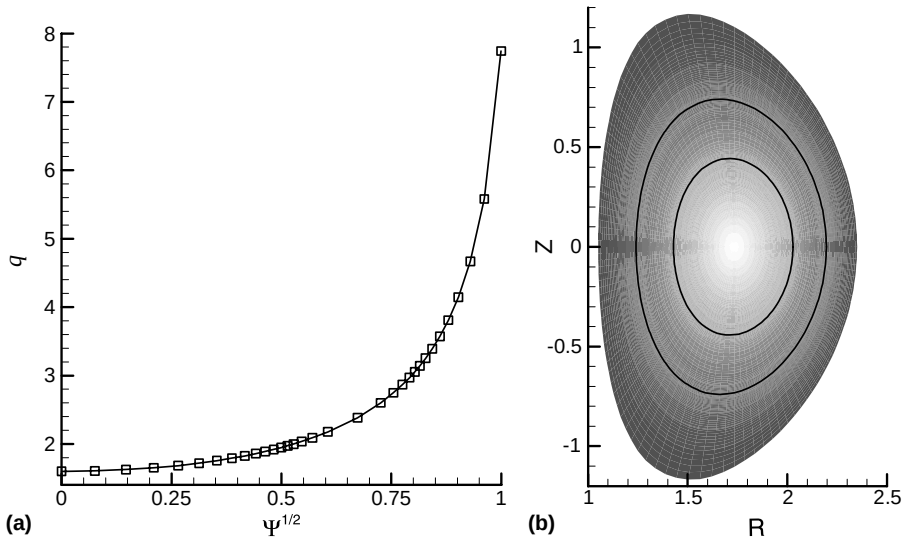

Fig. 10. Steady-state distributions used in the nonlinear simulation. The safety factor profile ($q = \mathrm{d}\Phi/\mathrm{d}\Psi$, where Φ is the toroidal magnetic flux function) is shown in (a) with boxes indicating element boundary locations in the topologically polar mesh. The distribution of $\sqrt{\Psi}$ in the poloidal plane is plotted in (b) with the heavy lines indicating the locations of the $q = 2$ and $q = 3$ surfaces.

finite-sized magnetic island. Both computations (linear and nonlinear) use a 32×32 mesh of biquartic elements ($p = 4$) with moderate packing at the $q = 2$ and $q = 3$ surfaces (see Fig. 10(a)). The time-step in the linear computation is $\Delta t = 2\tau_A$, and in the nonlinear simulation its value is allowed to increase by a factor of two during the simulation. The boundary conditions described in Section 2 imply that the MHD dynamics reproduce fixed-boundary behavior in this configuration where there is no vacuum region surrounding the conducting plasma.

In the nonlinear simulation, the growth of the mode is immediately slowed from the exponential time-dependence that characterizes linear behavior. This is observed from Fig. 12(a) through the non-constant slope of magnetic perturbation energy evolution plotted on a semi-log scale. The result is consistent with analytic theory in that the island width (proportional to the fourth root of perturbation energy) is predicted to have linear-in-time growth starting when the helical island chain extends beyond the resistive tearing layer [39]. Here, the linear time-dependence of the island width occurs for $t < 12$ ms, as shown in Fig. 12(b), and the slope is within 33% of the value given by the analytical relation $\mathrm{d}w/\mathrm{d}t = 1.22\Delta'\eta/\mu_0$ [40], where Δ' has been estimated from the cylindrical dispersion relation, Eq. (17), using a growth rate calculated from the same toroidal equilibrium but with reduced viscosity. Over a time-scale that is long relative to the energy transport time-scale, $a^2/\chi_\perp$, the free energy in the equilibrium current density profile is expended, and a 3D steady state is achieved. The simulation also shows that the coupling of harmonics illustrated in Fig. 11(b) leads to a secondary magnetic island chain at the $q = 3$ surface. Thus, the final state shown in Fig. 12(c) has two sets of helical magnetic surfaces that are embedded in nested toroidal surfaces.

Changes in the temperature profile due to the presence of a magnetic island can lead to nonlinear neoclassical effects in tokamaks [41,42], so accurate modeling of island thermodynamics is also important for tokamak simulation studies. Whether anisotropic heat conduction affects the temperature profile in the presence of the island depends on the balance of diffusion in the parallel and perpendicular directions [43]. The length-scale for parallel conduction is effectively infinite at the island separatrix, since the magnetic field-lines reconnect on themselves after a finite number of transits and do not trace a complete helical surface. However, flattening of the temperature profile occurs within the island when magnetic reconnection decreases the parallel length-scale enough so that parallel conduction occurs at a rate that is com-

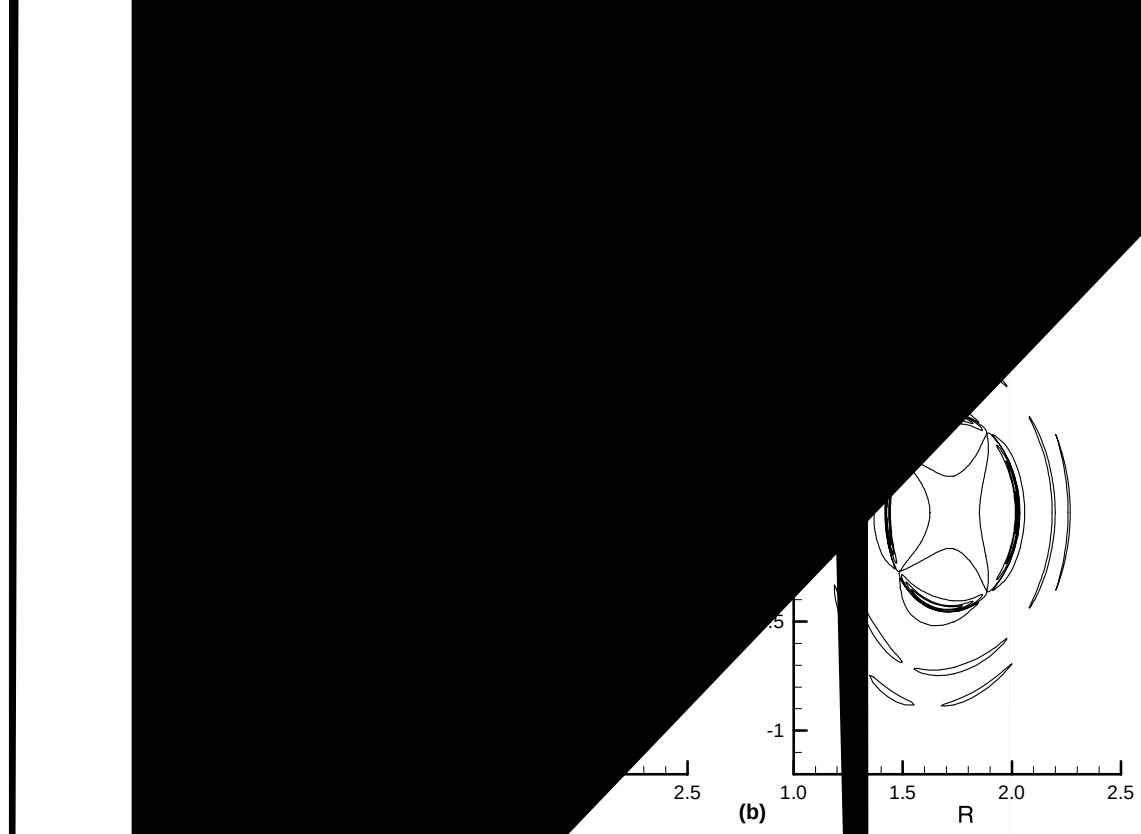

Fi , $Pm = 0.1$ linear tearing mode th s unstable in the equilibrium shown in Fig. 10. The po isocontours of the toroidal compone of the perturbed ideal electric field ($\hat{\varphi} \cdot \mathbf{v} \times \mathbf{B}_{ss}$) are sh

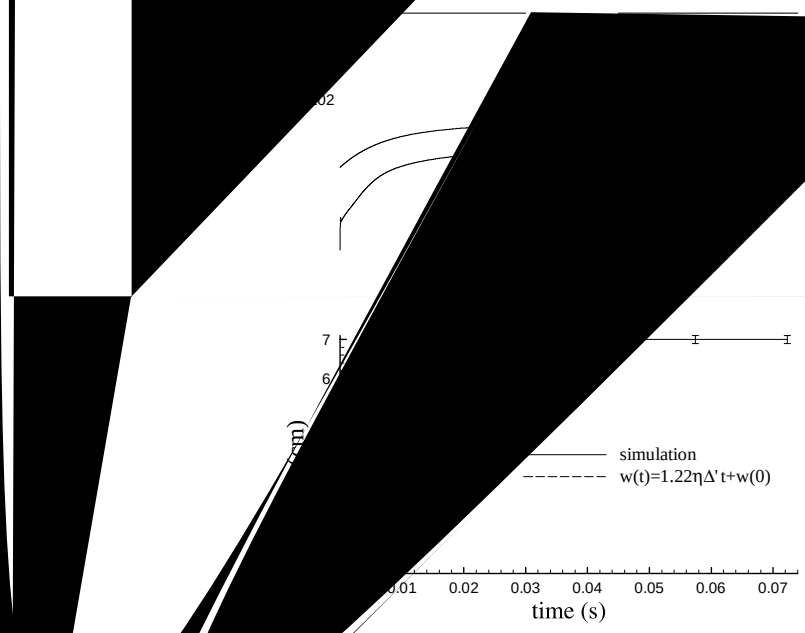

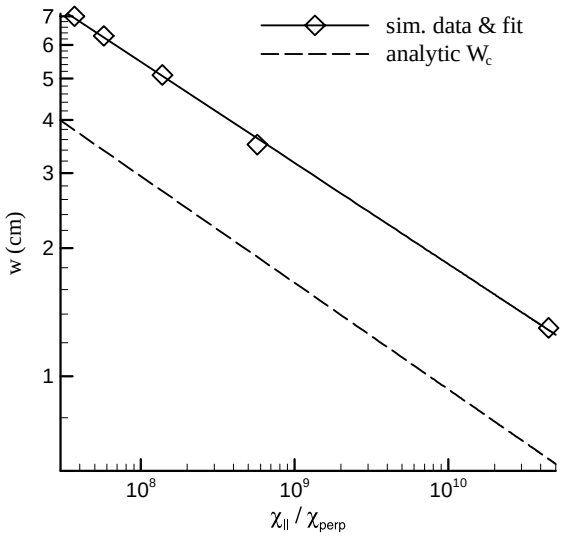

Fig. 13. Magnetic island width on the outboard side of the $q = 2$ surface that produces an inflection in the temperature profile, plotted vs. the ratio of thermal diffusivity coefficients. The solid line is the power-law fit of the simulation results, $w = 433(\chi_\parallel/\chi_\perp)^{-0.237}$ cm, and the dashed line is the analytic scaling of the critical width $W_c \equiv (\chi_\parallel/\chi_\perp)^{-1/4}(8R_0q/\mathrm{n}q')^{1/2}$ where the parallel and perpendicular diffusion times match in cylindrical geometry [43]. The latter is evaluated as an approximate for the toroidal configuration by averaging $q' = \mathrm{d}q/\mathrm{d}r$ from the inboard and outboard sides of the equilibrium; this yields $W_c = 295(\chi_\parallel/\chi_\perp)^{-1/4}$ cm.

petitive with perpendicular conduction, i.e. $\chi_\parallel/L_\parallel^2 \cong \chi_\perp/L_\perp^2$. Since the parallel length-scale within the island is inversely proportional to the island width (for island widths that are small in comparison to the length-scale of the equilibrium magnetic shear), and the perpendicular length-scale is proportional to the island width, the critical island width required to affect the temperature is expected to follow $W_c \sim (\chi_\parallel/\chi_\perp)^{-1/4}$ [43].

To test whether the NIMROD algorithm reproduces the theoretical dependence, we use the magnetic field configuration from five different times in the nonlinear simulation and run thermal-conduction-only computations with gradually increasing $\chi_\parallel$ in each configuration. Recording the $\chi_\parallel/\chi_\perp$-ratio required to produce an inflection of the temperature profile at the resonant surface as a function of island width then permits comparison. (The alternative of running a series of nonlinear MHD simulations with different $\chi_\parallel/\chi_\perp$-ratios would require far more computation.) The simulation result for the island-width scaling, $w \sim (\chi_\parallel/\chi_\perp)^{-0.24}$, is in good agreement with the analytic scaling of [43], and even the numerical coefficients are comparable, as illustrated in Fig. 13. The discrepancy reflects the fact that the numerically observed w and W_c are different quantities. The analytic relation has been derived as a scaling argument to distinguish small- and large-island-width behavior by identifying conditions where the parallel and perpendicular diffusion times match. It is not a precise relation for the condition recorded from the simulations, which is the inflection of the T-profile. The analytic relation has also been derived for cylindrical geometry and does not account for any toroidal effects that influence the island geometry. In fact, the simulation results provide empirical evidence supporting the application of the analytic scaling to toroidal configurations.

6. Discussion

The test results presented in Sections 4 and 5 demonstrate favorable convergence properties in stiff and anisotropic conditions when the degree of the polynomial basis functions, p, is 2 or larger. For $p = 1$, the poor performance can be anticipated from the occurrence of "spectral pollution" in ideal MHD eigenvalue

calculations [1,2] and from the divergence-stability consideration in steady incompressible fluid computations [31,32]. However, while our approach to spatial representation is generally related to the methods used in these applications, there are unique aspects in both the non-ideal MHD application and the algorithm. Here, we discuss how the unique aspects contribute to the favorable performance in the time-dependent MHD computations with $p \geqslant 2$.

A numerical approximation of the ideal-MHD linear force operator, $\mathbf{L}$ from Eq. (8), appears in our semi-implicit time-advance algorithm and in computations of ideal MHD eigenvalue problems [1]. However, it plays a different role in the two applications. Observe that an ideal MHD eigenvalue problem can be defined by considering $C_0\Delta t^2 \rightarrow \lambda$ as a free parameter in our velocity advance, Eq. (12), with $C_1 = v = 0$ and dropping all terms on the right side. The remaining Eqs. (13)–(15), and the solution space for n, T, and $\mathbf{B}$ would not be used, so the numerical spectrum is determined by the approximations of $\mathbf{L}$ and the inertial term that result from the basis functions used for $\mathbf{V}_{h,N,p}$. In contrast, the purpose of the $\mathbf{L}$-operator appearing in our time-advance algorithm is to add selective numerical dispersion to the leap-frog method. Here, the $\mathbf{L}$-operator does not, in itself, determine the result of linear computation, because the algorithm has separate computations of n, T, and $\mathbf{B}$, and the solution space is correspondingly larger. As a means to extend numerical stability beyond the Δt-limitations arising in a purely explicit time-advance, semi-implicit operators only need to approximate the combined operation of the explicit terms appearing on the right sides of Eqs. (12)–(15) – see Eq. (20) of [23]. Thus, it is possible to use approximations of $\mathbf{L}$ that are unacceptable from the ideal MHD eigenvalue standpoint, provided that errors resulting from the spectrum of the semi-implicit operator diminish rapidly as time-dependent results are converged. Although we have not analyzed the spectrum of our semi-implicit operator, the results presented in Section 4 show favorable convergence properties for $p \geqslant 2$.

A more readily apparent distinction from the ideal MHD problem is that the non-ideal MHD system (4a)–(4f) is higher-order as a system of partial differential equations, due to the dissipative terms. To ensure that the dissipative terms, like $\int d\mathbf{x}\, g\Delta t\rho v(\nabla \mathbf{w}^*) : (\nabla \Delta \mathbf{v})$ in Eq. (12), are square-integrable, hence, to be admissible in a conforming approximation, the basis functions must be continuous. The alternative is to use a *mixed method* [44] with additional equations and finite-dimensional spaces for the spatial derivatives themselves, but this is achieved at the expense of solving larger linear systems.

While the algorithms for ideal MHD eigenvalue problems and time-dependent non-ideal MHD can differ in important ways, general properties that lead to effective computation are similar because the ideal terms also dominate the behavior of stiff time-dependent MHD systems. The most important properties [1,2,45] that are needed for the ideal part of the algorithm are (1) resolution of $\mathbf{B}\cdot\nabla$ in singular layers and (2) accurate approximation of nearly incompressible flows and the magnetic divergence constraint [9]. Regarding the $\mathbf{B}\cdot\nabla$ calculation, we have quantitatively examined its convergence properties with the anisotropic thermal conduction tests of Sections 4.2 and 5. For the ideal MHD contributions, the context is different ($\mathbf{B}\cdot\nabla\mathbf{V}$ and $\mathbf{B}\cdot\nabla\mathbf{B}$), but the convergence properties of all spatial derivatives within elements follow the behavior described by Eq. (3). The results presented in Sections 4.1 and 4.2 provide confidence that continuous basis functions with $p \geqslant 2$ are satisfactory in this regard and that $p \geqslant 3$ accommodates severe anisotropy. A unique benefit of using generic high-order basis functions, which are not specialized for different vector components in a magnetic coordinate system, is that convergence properties are not lost in nonlinear computations when the magnetic direction vector changes significantly.

Regarding the divergence constraint and compressibility, when $p \geqslant 2$, the finite-element part of Eq. (10a) is closely related to continuous vector-field expansions that are used in 2D viscous incompressible fluid computation. However, the error diffusion method is not one of the standard methods for enforcing incompressibility. To compare the error diffusion method with the standard methods, consider introducing an auxiliary scalar variable in Eq. (16) for the divergence error and a separate constraint equation. In this mixed method, the magnetic advance now solves for $\mathbf{b}^{j+3/2}$ and X that satisfy

$$\int \mathrm{d}\mathbf{x}\left\{\mathbf{c}^*\cdot\mathbf{b}^{\mathrm{j}+3/2}+\Delta t\frac{\eta}{\mu_0}(\nabla\times\mathbf{c}^*)\cdot(\nabla\times\mathbf{b}^{\mathrm{j}+3/2})-(\nabla\cdot\mathbf{c}^*)X\right\}$$
$$=\int \mathrm{d}\mathbf{x}\,\Delta t(\nabla\times\mathbf{c}^*)\cdot(\mathbf{v}^{\mathrm{j}+1}\times\bar{\mathbf{b}})-\Delta t\int \mathrm{d}\mathbf{s}\times\mathbf{E}\cdot\mathbf{c}^*, \tag{18a}$$

$$\int \mathrm{d}\mathbf{x}\left\{\frac{\Xi X}{\lambda}+\Xi\nabla\cdot\mathbf{b}^{\mathrm{j}+3/2}\right\}=0, \tag{18b}$$

for all $\mathbf{c}\in\mathbf{B}_{h,N,\mathrm{p}}$ and for all $\Xi\in X_{h,N,\mathrm{p}'}$, where $X_{h,N,\mathrm{p}'}$ is a finite-dimensional space for the additional scalar X. There is no differentiation of the auxiliary scalar, so its representation only needs to be piecewise continuous to satisfy the requirements for a conforming approximation. This method can be related to the projection method of Brackbill and Barnes [9], but solving Eqs. (18a), (18b) simultaneously with a large value of λ prevents the formation of monopoles, whereas projection removes them after the magnetic field is advanced. Numerical analysis of finite elements for steady incompressible fluid applications proves that it is possible to find $X_{h,N,\mathrm{p}'}$ for continuous representations of $\mathbf{b}^{\mathrm{j}+3/2}$ with p>1 such that the product space of $\{\mathbf{B}_{h,N,\mathrm{p}},\mathbf{X}_{h,N,\mathrm{p}'}\}$ satisfies divergence-stability [31,46]. Convergence to a divergence-free vector field is then assured even in the limit of $\lambda\to\infty$, which is comparable to taking the limit $\Delta t\kappa_{\mathrm{div}\,b}\to\infty$.

If one were to replace (18b) with the local relation $X=-\lambda\nabla\cdot\mathbf{b}^{\mathrm{j}+3/2}$, substituting X into (18a) recovers Eq. (16) with $\Delta t\kappa_{\mathrm{div}\,b}\to\lambda$, but this changes the numerical character of the finite element solution. The space represented by $\{(\nabla\cdot\mathbf{b})|\mathbf{b}\in\mathbf{B}_{h,N,\mathrm{p}}\}$ is not among the $X_{h,N,\mathrm{p}'}$ spaces that satisfy divergence-stability in combination with continuous representations of $\mathbf{B}_{h,N,\mathrm{p}}$, because it imposes too many constraints [32]. If the approximation is so over-constrained that the matrix resulting from $\int \mathrm{d}\mathbf{x}\,\lambda(\nabla\cdot\mathbf{c}^*)(\nabla\cdot\mathbf{b}^{\mathrm{j}+3/2})$ is invertible, the physical terms in (18a) would not affect the solution unless λ has a small value. The *penalty method* described in [10] uses this form of the constraint relation, but *selective reduced numerical integration*, i.e. intentionally inaccurate numerical integration, of the constraint terms ensures that the matrix resulting from $\int \mathrm{d}\mathbf{x}\,\lambda(\nabla\cdot\mathbf{c}^*)(\nabla\cdot\mathbf{b}^{\mathrm{j}+3/2})$ is singular. Ref. [47] shows that in some cases, reduced numerical integration is identical to using a mixed method that satisfies divergence-stability.

In our time-dependent computations without selective reduced integration, poor performance of the error diffusion technique results from over-constraining the computation when the value of $\Delta t\kappa_{\mathrm{div}\,b}$ is chosen to be too large for a given continuous representation of magnetic field. The increasing range of acceptable $\Delta t\kappa_{\mathrm{div}\,b}$-values with polynomial degree (p), illustrated by the results shown in Fig. 7(a), reflects better separation of the longitudinal and solenoidal parts of the expanded vector field as the number of degrees of freedom in each element are increased. This increasing separation implies that the matrix from $\int \mathrm{d}\mathbf{x}\,\lambda(\nabla\cdot\mathbf{c}^*)(\nabla\cdot\mathbf{b}^{\mathrm{j}+3/2})$ becomes singular as p is increased from unity, so the constraint term does not dominate the physical terms when $\Delta t\kappa_{\mathrm{div}\,b}$ is finite.

We can also assess the longitudinal/solenoidal separation by counting the dimensionality of the spaces $\mathbf{B}_{h,0,\mathrm{p}}$ and $\{(\nabla\cdot\mathbf{b})|\mathbf{b}\in\mathbf{B}_{h,0,\mathrm{p}}\}$ as functions of p in a specific example: a 2D rectangular mesh of m Lagrange elements in each direction has $(m\mathrm{p}+1)^2$) nodes or $2(m\mathrm{p}+1)^2$ coefficients for an arbitrary vector field with two components, which is relevant for the $\mathrm{n}=0$ part of our computations. With Dirichlet boundary conditions for the normal component along the entire boundary as essential conditions, the dimension of $\mathbf{B}_{h,0,\mathrm{p}}$ is then $2(m^2\mathrm{p}^2-1)$. Finding the dimensionality of $\{(\nabla\cdot\mathbf{b})|\mathbf{b}\in\mathbf{B}_{h,0,\mathrm{p}}\}$ is a little more complicated, because the subspaces of scalars from $\partial b_x/\partial x$ and from $\partial b_y/\partial y$ intersect but are not identical. With Cartesian components and the mesh aligned with the axes of the coordinate system, the scalar field formed by $\partial b_x/\partial x$ consists of discontinuous piecewise polynomials of degree $\mathrm{p}-1$ in the x-direction and continuous piecewise polynomials of degree p in the y-direction. In general, this field can be described by a nodal polynomial expansion with $m\mathrm{p}(m\mathrm{p}+1)$ nodes, but the boundary conditions on $\mathbf{B}$ constrain path integrals across the x-dimension, $\int_{x_{\min}}^{x_{\max}}\mathrm{d}x\,\partial b_x/\partial x$, so there are only $m^2\mathrm{p}^2-1$ degrees of freedom. The scalar field of

$\partial b_y/\partial y$ similarly has $m^2\mathrm{p}^2-1$ degrees of freedom. The two discontinuous scalar fields share all polynomials that are continuous and of degree $\mathrm{p}-1$ in both directions and that satisfy the path integral constraints in both directions. Therefore, the intersection is described by an expansion with $m^2(\mathrm{p}-1)^2$ nodes. Adding the dimensions of the discontinuous spaces and subtracting the dimension of their continuous intersection, we find $\{(\nabla\cdot\mathbf{b})|\mathbf{b}\in\mathbf{B}_{h,0,\mathrm{p}}\}$ to have dimension $m^2\mathrm{p}^2+2m^2\mathrm{p}-m^2-2$, which is the number of constraints imposed when using members of this space as test functions for the divergence constraint. The ratio of degrees of freedom in $\mathbf{B}_{h,0,\mathrm{p}}$ to the number of constraint equations is then

$$\frac{2m^2\mathrm{p}^2-2}{m^2\mathrm{p}^2+2m^2\mathrm{p}-m^2-2}\cong\frac{2\mathrm{p}^2}{\mathrm{p}^2+2\mathrm{p}-1},$$

where the approximation holds for large m. Although this result has been derived for a special mesh, we expect that the large-m approximation is general. For $\mathrm{p}=1$, the approximate ratio is unity, making computational results very sensitive to the value of $\Delta t\kappa_{\mathrm{div}\,b}$. For $\mathrm{p}=2$ and larger, the ratio exceeds unity, and for large p, it approaches 2, the optimal ratio in 2D computations [32].

For 3D computations, the constraint can be addressed separately for each Fourier component, because the divergence operator is linear. Considerations for the $\mathrm{n}=0$ Fourier component are identical to those given above. For all other Fourier components, the number of additional degrees of freedom due to the third dimension is equivalent to the number of nodes in the representation of b_z, which is $(m\mathrm{p}+1)^2$ in the simple mesh used above. The number of test functions, and hence the number of constraints, also increases. Here, the derivative $\partial b_z/\partial z$ is just an algebraic operation in the Fourier representation, $2\pi\mathrm{i}\mathrm{n}b_z/L_z$, so the space of test functions includes polynomials that are continuous and of degree p in both x and y. All of the possible continuous functions (of degree p in one direction and $\mathrm{p}-1$ in the other) resulting from $\partial b_x/\partial x$ and $\partial b_y/\partial y$ are contained in the larger continuous space associated with $2\pi\mathrm{i}\mathrm{n}b_z/L_z$. Thus, the number of constraints for Fourier indices with $\mathrm{n}>0$ is $(m\mathrm{p}+1)^2$ plus the number of nodes that allow discontinuity in expansions for $\partial b_x/\partial x$ and $\partial b_y/\partial y$, $2(m\mathrm{p}+1)(m-1)$. The ratio of degrees of freedom to the number of constraint equations for $\mathrm{n}>0$ is then

$$\frac{3m^2\mathrm{p}^2+2m\mathrm{p}-1}{m^2\mathrm{p}^2+2m^2\mathrm{p}+2m-1}\cong\frac{3\mathrm{p}^2}{\mathrm{p}^2+2\mathrm{p}}\quad\text{for large } m.$$

Here, again, the ratio exceeds unity for $\mathrm{p}>1$ and large m, and it approaches the optimal ratio of 3 for 3D computations in the limit of large p. (Section 4.4 of [32] describes similar conclusions for quadrilateral and brick elements, but the number of constraints is determined by the accuracy of the numerical integration. The ratios derived above are based on exact integration and the combined finite element/Fourier representation.)

The increasing insensitivity to $\Delta t\kappa_{\mathrm{div}\,b}$-values as p is increased, as demonstrated by the results in Fig. 7(a), is consistent with what we have found above regarding the dimensionality of $\mathbf{B}_{h,N,\mathrm{p}}$ and the number of constraints imposed by the divergence cleaning term. The ratios of degrees of freedom for $\mathrm{n}=0$ and $\mathrm{n}>0$ approach their optimal values as p is increased, so we can expect increasing separation of expanded longitudinal and solenoidal fields. For computations where the ratios are greater than unity but not optimal, selecting finite $\Delta t\kappa_{\mathrm{div}\,b}$-values avoids over-constraining the magnetic advance. We have found that setting $\Delta t\kappa_{\mathrm{div}\,b}/h^2\sim\mathrm{O}(1)$ or $\mathrm{O}(10)$, i.e. diffusing the error over the element dimension at each time-step, enforces the constraint sufficiently in most of our applications.

The numerical issues for compressibility of flow are similar to the considerations for the magnetic constraint. Although the equations we solve are compressible, the anisotropies of the MHD system lead to very different responses between shearing and compression, and compressive behavior tends to equilibrate on time-scales that are fast in comparison to resistive tearing behavior [33,34]. The numerical operator $\mathbf{L}$ appearing on the left-hand side of Eq. (12) contains the terms

$$C_0\Delta t^2\left[\frac{B_0^2}{\mu_0}(\nabla\cdot\mathbf{w}^*_\perp)(\nabla\cdot\Delta\mathbf{v}_\perp)+\gamma p_0(\nabla\cdot\mathbf{w}^*)(\nabla\cdot\Delta\mathbf{v})\right], \tag{19}$$

where the first term arises from motion perpendicular to $\mathbf{B}_0$. Since the coefficients can be very large in comparison to others in Eq. (12) – the ratio $\Delta t^2(\gamma p_0+B_0^2/\mu_0)/\rho_0$ is the square of the distance traveled by the fastest wave in the MHD system in a time-step – the terms in (19) restrict compressibility, especially for perpendicular motions. Since $\mathbf{B}_{h,N,\mathrm{p}}$ and $\mathbf{V}_{h,N,\mathrm{p}}$ share the same set of basis functions, the numerical arguments regarding the dimensionality of the space and the number of constraints are also applicable to compressibility. However, C_0 is set by numerical stability requirements for the semi-implicit advance, so the only freedom in controlling the magnitude of the compressive terms is through the Δt-value for the time-step. For example, the error displayed in Fig. 4(c) eigenfunction computed for the cylindrical tearing mode with a 32×32 mesh of bicubic elements decreases with h, but it also decreases with Δt. With reduced Δt, the accuracy of the semi-implicit operator is relatively less important, while reducing h leads to better resolution of the anisotropy.

Regarding computational performance, the sparse direct solver library, SuperLU, has provided a significant improvement over iterative methods, which is attributable to the ill-conditioning of the matrices when Δt is large. As examples of current performance, the 32×32 bicubic, 24×24 biquartic, and 20×20 biquintic linear computations of the cylindrical tearing mode considered in Section 4.1 each take approximately 2.5 s per time-step on one processor of a 2 GHz Intel Pentium IV-based workstation. When the mesh is increased to 48×48 in the bicubic computation, it takes 9.3 s per time-step. Running the nonlinear computation discussed in Section 5, which has three Fourier components and uses matrix-free iterative solves for the temperature advance, on the IBM-SP3 at the National Energy Research Supercomputing Center (http://www.nersc.gov) takes 13.1 s per step on 12 processors, 7.7 s per step on 27 processors, and 5.9 s per step on 48 processors.

7. Conclusions

We have described an algorithm that combines a variational spatial representation with a semi-implicit time-advance to achieve flexibility and accuracy for application to non-ideal MHD. The marching algorithm is considered a set of variational problems, and the hyperbolic character of the nonlinear PDE system is brought out in a sequence of complete advances. The temporal and spatial techniques benefit from each other through their symmetry characteristics. The time-advance stabilizes the propagation of waves at large time-step by introducing an implicit self-adjoint differential operator, and the finite element approach ensures that the matrices resulting in the fully discretized system are Hermitian. Conversely, the variational approach to spatial discretization provides the required accuracy, and the self-adjoint semi-implicit operator allows us to create a variational form of the velocity-advance equation. A more general Galerkin approach may be useful for treating either ion or electron flows implicitly.

The benchmark cases presented in Section 4 and the nonlinear simulation presented in Section 5 demonstrate the effectiveness of the algorithm. The resistive tearing calculations show that a modest number of finite elements with $p > 1$, sufficient mesh packing, and a large time-step can reproduce the subtle force balances associated with MHD anisotropy. For example, the computation with a 16×16 mesh of bicubic elements and $\Delta t = 100\tau_A$, which is nearly 10^5 times greater than the limit for an explicit computation with the same spatial representation, finds a growth rate that is within 12% of the converged result for $S = 10^6$ and $Pm = 10^{-3}$. The anisotropic thermal conduction test in simple geometry shows that sufficient accuracy can be achieved to resolve parallel and perpendicular transport properties in realistic conditions without aligning the grid to the magnetic field; efforts to align the grid will further increase accuracy. The simulation discussed in Section 5 demonstrates performance with respect to slowly growing nonlinear MHD activity,

and the comparison between numerical and analytic results on the magnetic island width required for temperature profile modification confirms that the modeling of anisotropic diffusion in 3D magnetic topologies is accurate.

The geometric flexibility of the algorithm makes it suitable for many applications in magnetic confinement fusion. The nonlinear tearing evolution illustrates conditions encountered while using NIMROD to simulate neoclassical tearing modes and high-beta disruptions in tokamaks [21,48], where accurate anisotropic diffusion is critical. In combination with a temperature-dependent resistivity, the accurate modeling of anisotropic diffusion permits us to address nonlinear free-boundary tokamak computations, where Ohmic heating leads to large electrical conductivity in the region of closed magnetic flux surfaces only [49]. NIMROD is also being used to simulate nonlinear magnetic relaxation in alternate configurations, such as spheromaks [50–52] and reversed-field pinches [49,53], where separation of time-scales tends to be less extreme than in tokamak plasmas, but the behavior often includes evolution to MHD turbulence. Although numerical issues associated with relaxation simulations have not been discussed in this paper, the flexibility to address many different topics with one code has been a goal since the inception of the NIMROD project. Additional geometric flexibility will be achieved with further development of triangular elements. (For simulating experimental configurations without geometric symmetry, the numerical algorithm can be implemented with finite elements in all three directions.)

Further development of the algorithm is proceeding along two general paths. First, we continue to make numerical refinements for the non-ideal MHD model described here. We expect to improve the existing predictor/corrector treatment of flow with regard to accuracy and efficiency in extreme (but not shocked) conditions. In addition, we will investigate selective numerical integration for the compressibility terms in the semi-implicit operator. The second path of development concerns expanding the algorithm to solve more realistic models for high temperature plasmas. The NIMROD implementation is designed to have flexibility in the equations that it solves, and the modularity facilitates efforts to improve numerical methods for more realistic plasma models. Some development has already been completed in the area of two-fluid effects [7], and we are presently working to improve accuracy at the large time-steps needed for nonlinear fusion studies. We are also adding kinetic effects [20–22] that have a strong influence on the MHD-like behavior of nearly collisionless plasmas.

Acknowledgements

The authors wish to acknowledge Curtis Bolton for initiating the NIMROD project and James Callen and Chris Hegna for many valuable discussions regarding MHD-like physics in high performance tokamaks. The first author also wishes to thank David Malkus, Michael Plesha, and Olivier Sauter for assistance with relevant concepts in finite element analysis and the anonymous referees for their many constructive comments. Some of the simulations reported here were performed at the National Energy Research Scientific Computing Center at Lawrence Berkeley National Laboratory.

Appendix A

Several applications require simply connected, topologically cylindrical domains. For these cases, we use the finite element representation for the R–Z plane with one side of the mesh lying along the Z-axis. Physical fields and their partial derivatives must have unique values at the axis, which leads to a set of regularity conditions for the Fourier components in the limit of $R \to 0$. The conditions are derived with a 2D Taylor series expansion of an arbitrary function of Cartesian x and y coordinates with origin at $R = 0$ in a constant-Z plane. Substituting $R(e^{i\varphi} + e^{-i\varphi})/2$ and $R(e^{i\varphi} - e^{-i\varphi})/2i$ for the Cartesian x and y, respectively, while

leaving the expansion coefficients in terms of Cartesian derivatives, determines the appropriate functional form for each Fourier component in the limit of $R \to 0$. For scalars we have

$$\lim_{R\to 0} F(R,\varphi) = f_0(R^2) + \sum_{\mathrm{n}=1}^{N} R^n [f_{\mathrm{n}}(R^2) + f_{\mathrm{n}}^*(R^2)] \tag{A.1}$$

for the finite Fourier series, where $f_{\mathrm{n}}, \mathrm{n} = 0, 1, \ldots, N$ are polynomial functions of their argument. For vectors, we have

$$\lim_{R\to 0} A_{R,\varphi}(R,\varphi) = R a_{R_0,\varphi_0}(R^2) + \sum_{\mathrm{n}=1}^{N} R^{|\mathrm{n}-1|} [a_{R_{\mathrm{n}},\varphi_{\mathrm{n}}}(R^2) + a^*_{R_{\mathrm{n}},\varphi_{\mathrm{n}}}(R^2)] \tag{A.2}$$

where $a_{R_{\mathrm{n}}}$ and $a_{\varphi_{\mathrm{n}}}$ are polynomial functions. The limiting behavior of $A_z(R,Z)$ is the same as for scalars. The radial and azimuthal expansions must also satisfy

$$a_{\varphi_1}(0) = \mathrm{i} a_{R_1}(0), \tag{A.3}$$

so that the vector has a unique direction at $R = 0$.

Since conditions (A.1)–(A.3) apply in the limit of $R \to 0$, discrete representations in R need only consider the leading behavior for each Fourier component index. Conditions where the polynomial expansion goes to zero at $R = 0$ are applied as essential conditions on the solution space, like Dirichlet boundary conditions. Satisfying condition (A.3) for $\mathrm{n} = 1$ vector components is more complicated, since the R and φ components of a vector are computed simultaneously according to the algebraic system resulting from the weak form of the implicit terms in a given equation. Our approach is to first compute the matrix elements for the coefficients of the spatial representation (denoted by $\hat{a}$) in the usual manner without considering (A.3). Then for each node located on the Z-axis, we change to sum and difference coefficients,

$$\hat{a}_+ \equiv \frac{\hat{a}_{R_1} + \mathrm{i}\hat{a}_{\varphi_1}}{2},$$
$$\hat{a}_- \equiv \frac{\hat{a}_{R_1} - \mathrm{i}\hat{a}_{\varphi_1}}{2},$$

in the algebraic system. If $(\underline{\underline{M}})_{\mathrm{c}}$ denotes the two columns of the matrix corresponding to the a_{R_1} and a_{φ_1} elements in the algebraic vector of unknown coefficients, the variable change modifies these columns to

$$(\underline{\underline{M}})_c \Rightarrow (\underline{\underline{M}})_c \begin{pmatrix} 1 & 1 \\ -\mathrm{i} & \mathrm{i} \end{pmatrix}.$$

The number of rows of the linear system is then reduced by taking a linear combination of the two rows (r) corresponding to $\hat{a}_+$ and $\hat{a}_-$,

$$(\underline{\underline{M}})_r \Rightarrow (1 - \mathrm{i})(\underline{\underline{M}})_r \text{ and } (b)_r \Rightarrow (1 - \mathrm{i})(b)_r,$$

and the regularity condition (A.3) is enforced at the given node by setting an essential condition on the sum coefficient, $\hat{a}_+ = 0$. The operations are repeated for each node along the axis, and the resulting matrix retains the Hermitian property of the original.

For the $\mathrm{n} = 0$ component of scalars, the $\mathrm{n} = 0$ Z-component of vectors, and the $\mathrm{n} = 1$ R- and φ-components of vectors, the leading behavior of the polynomial expansion in R is that the slope vanishes in the limit $R \to 0$. For each equation in the time advance, this behavior is enforced by the respective strain energy without any modification to the spatial representation or the resulting linear system. For example if the $\mathrm{n} = 0$ Fourier component in the expansion of B_z has a non-zero radial derivative at small R, computations

of the n = 0 part of the local J_φ will be non-zero, leading to resistive diffusion or radial forces that generate flows to reduce the derivative. The weak form used in the finite element approach avoids singular terms at $R = 0$ that need special treatment in a finite difference or finite volume scheme, where the conditions are explicitly applied to the differences [29]. As another example of the finite element implementation, consider radial derivatives in the n = 1 R-component of magnetic field near $R = 0$. Since the axial n = 1 component is set to 0 on axis, and condition (A.3) is satisfied, the n = 1 part of magnetic divergence reduces to $\partial b_r/\partial r$, and the error diffusion term in the strain energy for the magnetic advance will tend to eliminate any non-zero derivative. The conditions are realized as natural conditions, even though there is no corresponding surface integral.

References

[1] R. Gruber, J. Rappaz, Finite Element Methods in Linear Ideal Magnetohydrodynamics, Springer, Berlin, 1985, pp. 79–117.
[2] L.M. Degtyarev, S.Yu. Medvedev, Comput. Phys. Commun. 43 (1986) 29.
[3] C.-W. Shu, High order ENO and WENO schemes for computational fluid dynamics, in: T.J. Barth, H. Deconinck (Eds.), High Order Methods for Computational Physics, Springer, Berlin, 1999.
[4] I. Babuska, B.A. Szabo, I.N. Katz, SIAM J. Numer. Anal. 18 (1981) 515.
[5] G. Strang, G.J. Fix, An Analysis of the Finite Element Method, Prentice-Hall, Englewood Cliffs, NJ, 1973.
[6] B. Marder, J. Comput. Phys. 68 (1987) 48.
[7] A.H. Glasser, C.R. Sovinec, R.A. Nebel, T.A. Gianakon, S.J. Plimpton, M.S. Chu, D.D. Schnack, and the NIMROD Team, Plasma Phys. Control. Fusion 41 (1999) A747.
[8] C.-D. Munz, P. Omnes, R. Schneider, E. Sonnendrücker, U. Voß, J. Comput. Phys. 161 (2000) 484.
[9] J.U. Brackbill, D.C. Barnes, J. Comput. Phys. 35 (1980) 426.
[10] T.J.R. Hughes, W.K. Liu, A. Brooks, J. Comput. Phys. 30 (1979) 1.
[11] L.A. Charlton, J.A. Holmes, H.R. Hicks, V.E. Lynch, B.A. Carreras, J. Comput. Phys. 86 (1990) 270.
[12] W. Park, D.A. Monticello, Nucl. Fusion 30 (1990) 2413.
[13] K. Lerbinger, J.F. Luciani, J. Comput. Phys. 97 (1991) 444.
[14] A.M. Popov, V.S. Chan, M.S. Chu, Y.Q. Liu, B.W. Rice, A.D. Turnbull, Phys. Plasmas 8 (2001) 3605.
[15] H.R. Strauss, W. Park, Phys. Plasmas 5 (1998) 2676.
[16] D. Gottlieb, S.A. Orzag, Numerical analysis of spectral methods: theory and application, in: Regional Conference Series in Applied Mathematics, vol. 26, SIAM, Philadelphia, 1977.
[17] D.S. Harned, W. Kerner, J. Comput. Phys. 60 (1985) 62.
[18] D.S. Harned, D.D. Schnack, J. Comput. Phys. 65 (1986) 57.
[19] S.I. Braginskii, Transport processes in a plasma, in: Reviews of Modern Physics, vol. 205, Consultants Bureau, New York, 1965.
[20] E.D. Held, J.D. Callen, C.C. Hegna, C.R. Sovinec, Phys. Plasmas 8 (2001) 1171.
[21] T.A. Gianakon, S.E. Kruger, C.C. Hegna, Phys. Plasmas 9 (2002) 536.
[22] S.E. Parker, Y. Chen, C.C. Kim, Bull. Am. Phys. Soc. 47 (9) (2002) 163.
[23] D.D. Schnack, D.C. Barnes, Z. Mikić, D.S. Harned, E.J. Caramana, J. Comput. Phys. 70 (1987) 330.
[24] E.J. Caramana, J. Comput. Phys. 96 (1991) 484.
[25] J.P. Freidberg, Ideal Magnetohydrodynamics, Plenum Press, New York, 1987, pp. 244–246.
[26] R. Courant, K.O. Friedrichs, H. Lewy, Math. Ann. 100 (1928) 32.
[27] R. Lionello, Z. Mikić, J.A. Linker, J. Comput. Phys. 152 (1999) 346.
[28] A.H. Glasser, C.R. Sovinec, Numerical analysis of the NIMROD formulation, International Sherwood Fusion Theory Conference, April 28–30, Madison, Wisconsin (Available from http://www.nimrodteam.org/presentations/Sherwood97/num_anal.pdf.).
[29] D.D. Schnack, D.C. Baxter, J. Comput. Phys. 55 (1984) 485.
[30] C. Schwab, hp-FEM for fluid flow simulation, in: T.J. Barth, H. Deconinck (Eds.), High Order Methods for Computational Physics, Springer, Berlin, 1999.
[31] M.D. Gunzburger, Mathematical aspects of finite element methods for incompressible viscous flows, in: Finite Elements Theory and Application, Springer, New York, 1988.
[32] T.J.R. Hughes, The Finite Element Method, Prentice-Hall, Englewood Cliffs, NJ, 1987, pp. 209–217.
[33] H.P. Furth, J. Killeen, M.N. Rosenbluth, Phys. Fluids 6 (1963) 459.
[34] B. Coppi, J.M. Greene, J.L. Johnson, Nucl. Fusion 6 (1966) 101.
[35] W.A. Newcomb, Ann. Phys. 10 (1960) 232.

[36] R.J. Bickerton, Proc. Phys. Soc. London 72 (1958) 618.
[37] D.C. Robinson, Nucl. Fusion 18 (1978) 939.
[38] A.H. Glasser, J.M. Greene, J.L. Johnson, Phys. Fluids 19 (1976) 567.
[39] P.H. Rutherford, Phys. Fluids 16 (1973) 1903.
[40] D. Biskamp, Nonlinear Magnetohydrodynamics, Cambridge University Press, Cambridge, UK, 1993, pp. 107–114.
[41] T.A. Gianakon, J.D. Callen, C.C. Hegna, Phys. Plasmas 3 (1996) 4637.
[42] C.C. Hegna, Phys. Plasmas 5 (1998) 1767.
[43] R. Fitzpatrick, Phys. Plasmas 2 (1995) 825.
[44] F. Brezzi, A survey of mixed finite element methods, in: Finite Elements Theory and Application, Springer, New York, 1988.
[45] H. Lutjens, J.F. Luciani, Comput. Phys. Commun. 95 (1996) 47.
[46] J.M. Boland, R.A. Nicolaides, SIAM J. Numer. Anal. 20 (1983) 722.
[47] D.S. Malkus, T.J.R. Hughes, Comput. Methods. Appl. Mech. Eng. 15 (1978) 63.
[48] S.E. Kruger, D.D. Schnack, D.P. Brennan, T.A. Gianakon, C.R. Sovinec, Nonlinear MHD dynamics of tokamak plasmas on multiple time scales, Nucl. Fusion, submitted.
[49] C.R. Sovinec, T.A. Gianakon, E.D. Held, S.E. Kruger, D.D. Schnack, and the NIMROD Team, Phys. Plasmas 10 (2003) 1727.
[50] J.M. Finn, C.R. Sovinec, D. Del-Castillo-Negrete, Phys. Rev. Lett. 85 (2000) 4538.
[51] C.R. Sovinec, J.M. Finn, D. Del-Castillo-Negrete, Phys. Plasmas 8 (2001) 475.
[52] R.H. Cohen, H.L. Berk, B.I. Cohen, T.K. Fowler, E.B. Hooper, L.L. LoDestro, E.C. Morse, L.D. Pearlstein, T.D. Rognlien, D.D. Ryutov, C.R. Sovinec, S. Woodruff, Nucl. Fusion 43 (2003) 1220.
[53] P. Martin, L. Marrelli, G. Spizzo, P. Franz, P. Piovesan, I. Predebon, T. Bolzonella, S. Cappello, A. Cravotta, D.F. Escande, L. Frassinetti, S. Ortolani, R. Paccagnella, D. Terranova, B.E. Chapman, D. Craig, S.C. Prager, J.S. Sarff, P. Brunsell, J.–A. Malmberg, J. Drake, Y. Yagi, H. Koguchi, Y. Hirano, R.B. White, C. Sovinec, C. Xiao, R.A. Nebel, D.D. Schnack, and the RFX, MST, EXTRAP T2R, and TPE-RX teams, Overview of quasi single helicity experiments in reversed field pinches, Nucl. Fusion, accepted for publication.