

Semi-Implicit Stellarator Magnetohydrodynamics with Nodal Spectral Elements

**University of Wisconsin-Madison
Center for Plasma Theory and Computation
Report UW CPTC 25-1**

Carl Sovinec,¹ Sanket Patil,² and Vincent Cauilan¹

¹Department of Nuclear Engineering & Engineering Physics

²Department of Physics

University of Wisconsin-Madison
Madison, WI 53706-1609

November 15, 2025

NOTICE

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government, nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

Semi-Implicit Stellarator Magnetohydrodynamics with Nodal Spectral Elements

C. R. Sovinec,¹ S. A. Patil,² and J. V. Cauffman¹

University of Wisconsin-Madison, ¹Dept. of Nuclear Engineering & Engineering Physics and

²Dept. of Physics

Nonlinear time-dependent computation of macroscale dynamics in stellarators is motivated by laboratory results showing the possibility of robust operation in conditions where magnetohydrodynamic (MHD) modes are linearly unstable. A new formulation of semi-implicit MHD computation for toroidally shaped magnetic confinement systems uses 2D nodal spectral elements over the poloidal plane and Fourier representation over a generalized toroidal angle. Geometric mappings and steady-state (equilibrium) fields are expanded in the same 3D representation as the time-evolved fields to model non-axisymmetric configurations. For accuracy at large timestep, the semi-implicit operator is based on the ideal-MHD energy integral using 3D pressure and magnetic fields. The nodal spectral elements allow numerical convergence through either h -refinement or p -refinement. Our implementation (NIMSTELL) with the continuous H^1 expansion of magnetic-field components and diffusive divergence control is a generalization of the NIMROD code [JCOMP **195**, 355 (2004)]. The NIMSTELL implementation is verified linearly and nonlinearly on resonant ideal interchange, where convergence from the stable side results from the stabilization method used in NIMROD [JCOMP **319**, 61 (2016)]. Optionally, NIMSTELL may use an $\mathbf{H}(\text{curl})$ representation for vector potential, and both magnetic representations are verified with respect to results from JOEKE [Phys. Plasmas **29**, 063901 (2022)] on linear and nonlinear magnetic tearing in the W7-A rotating-ellipse configuration. Application of the existing vector-potential implementation to interchange shows that it needs a minimum level of electrical resistivity to avoid numerical noise for a given level of spatial resolution. Solving the algebraic systems from the implicit parts of the time advance is facilitated by including the Fourier components of stellarator mode families in each preconditioning operation.

1. Introduction

One of the many important applications of numerical magnetohydrodynamics (MHD) computation is the study of macroscale dynamics in magnetically confined plasma. The MHD model simplifies significant aspects of high-temperature plasma dynamics, but it provides critical insight on instabilities driven by spatial inhomogeneity and the consequences of instability. Moreover, different levels of MHD computation satisfy a variety of physical research needs. Computations of potential energy in the linearized ideal MHD system address stability at minimal cost [1]. Eigenmode computations for linear ideal [2] and nonideal [3], [4] behavior produce the spectra of frequencies and growth rates and eigenfunctions. In MHD's complete form, time-dependent nonlinear computations are used to study whether instabilities saturate in relatively benign topology changes [5], [6], relax the configuration to more stable configurations [7], [8], or lead to disruptive evolution [9], [10], [11], [12].

The majority of MHD computations for magnetic confinement to date have been applied to configurations where electrical current flowing in the plasma provides the necessary rotational transform, *i.e.*, twist, of the magnetic field to confine charged particles in toroidal configurations.

Stellarators and torsatrons are toroidally asymmetric configurations, where the magnetic twist results from shaped external coils. This approach to magnetic confinement avoids free energy from spatial variation of current density running along magnetic field in the plasma [13]. Nonetheless, macroscale dynamics driven by plasma pressure and residual plasma current can be important. Here, we present development for time-dependent simulations of stellarators and torsatrons.

An open research question that motivates MHD computation for stellarators is what limits the plasma pressure for a given magnetic field. Unlike the symmetric tokamak configuration that relies on plasma current, nonlinear evolution from linear instability often remains non-disruptive in stellarators. The largest attained ratio of average plasma pressure and magnetic pressure, $\langle\beta\rangle$, of more than 4.5% was achieved in the Large Helical Device (LHD), whose design was not optimized for MHD [14]. The helical coil currents were distributed to produce a large plasma aspect ratio, reducing the Shafranov shift that moves the central part of the plasma toward the outboard side of the torus. The high- $\langle\beta\rangle$ LHD discharges exceed the Mercier linear ideal-MHD stability limit [15] with seemingly little effect, possibly due to LHD’s relatively large magnetic shear [16]. Relaxation-type computations with the HINT code show loss of outer flux surfaces with increasing Shafranov shift [17], [18]. Similarly, the Wendelstein 7-AS experiment achieved MHD-quiescent operation at $\langle\beta\rangle = 3.4\%$ [17]. Appreciable Shafranov shift was observed in its high- $\langle\beta\rangle$ operation, and relaxation computations with the PIES code [19] indicate chaotic field-line topology over the outer part of the plasma [17]. Earlier operation in W7-AS included an initial inward-shifted configuration to allow large Shafranov shift. The magnetic well property, which enhances linear MHD stability [13], is reduced with the inward shift, but the experiment only produced non-disruptive MHD relaxation events that lost approximately 10% of the stored thermal energy [20].

Time-dependent MHD computations for stellarators can help address this and other important questions. Early contributions include computation of two-fluid effects on ideal and resistive ballooning using the finite-element version of the original M3D code [21]. A study of partial pressure-profile collapses in inward-shifted LHD discharges iterated between the resistive reduced-MHD dynamical evolution with the NORM code and updates of the equilibrium profile [6] from VMEC [22]. The role of nonlinear resistive ballooning in LHD core density collapses was investigated with the explicit MIPS code [23]. More recent time-dependent stellarator MHD codes are redevelopments—to varying extent—of existing macroscale simulation codes for toroidally symmetric configurations. The stellarator version of the M3D-C1 code [24] applies its 3D finite-element representation to vary the shape of the poloidal plane over the toroidal angle [25]. It retains the potential representation of the M3D equations for full MHD evolution. It has been applied to pressure-induced MHD [26] and magnetic reconnection events associated with electron cyclotron current drive [27] in the W-7X experiment and to interchange dynamics in an optimized quasi-axisymmetric (QA) configuration [28]. Generalizing JOEKE [29], [30] required theoretical development of its “ansatz based” reduced MHD [31]. The implementation retains JOEKE’s Fourier representation over the toroidal angle and applies it to geometric information in addition to the physical fields [32]. Recent work adds parallel flow evolution to the implementation and extends benchmarks and validation to the W7-AS configuration [33]. Both M3D-C1 and JOEKE use implicit time-stepping to avoid numerical stability limitations.

Analogous to these developments for M3D-C1 and JOEKE, the implementation described here, NIMSTELL, is a stellarator variant of the full-MHD NIMROD code [34], which has been applied to tokamaks, reversed-field pinches, spheromaks, and other magnetic confinement and space-plasma configurations. The NIMROD implementation assumes that the geometry and prescribed steady-state fields are symmetric with respect to a periodic coordinate, which is typically taken as

the toroidal angle. Previous application of NIMROD to stellarator configurations include the study of pressure-driven topology evolution in a straight stellarator [35], evolution of the rotational-transform profile subject to Ohmic current drive in the Compact Toroidal Hybrid (CTH) experiment [36], and flux-surface loss and chaotic-field transport in a 10-field period torsatron [37]. A study of sawtoothing from Ohmic current drive in CTH investigates the influence of external transform on current-driven dynamics and quantifies the numerical cost of representing stellarator fields with an axisymmetric mesh [38]. These applications satisfy NIMROD’s geometric limitation, but many modern stellarator experiments are not tractable with NIMROD. Besides computational efficiency, rigid external coils cannot be represented within NIMROD’s computational domain. The coil fields must be represented through magnetic-field boundary conditions along a symmetric surface that lies between the plasma surface and the physical coils, and such surfaces do not exist in modern experiments and designs. The development of NIMSTELL alleviates these restrictions through Fourier expansions of geometry and steady-state fields over a generalized toroidal coordinate.

With respect to the other recent stellarator MHD implementations, NIMSTELL adds a unique combination of numerical methods that are described in this article. NIMSTELL solves the non-reduced MHD system and has spectral spatial representation capability with its 2D nodal spectral elements [39], [40] combined with Fourier series for the periodic coordinate. It also uses a semi-implicit time-advance of primitive variables for flow-velocity and magnetic field, though vector potential is an option for NIMSTELL. For comparison, M3D-C1 solves non-reduced systems with potential fields for flow-velocity and for magnetic field and uses reduced quintic triangles for the poloidal plane and Hermite cubics over the toroidal angle [25], [41]. Having continuous first derivatives in its spatial representation allows differentiation up to fourth order for its potential formulation. JOEKE is typically used to solve reduced MHD systems. It uses 2D Bezier elements that can have high-order geometric continuity for the poloidal plane and 1D Fourier series for the toroidal angle [30], [42]. Both M3D-C1 and JOEKE use implicit time advances [43], [30]. Their consistent implicit formulations allow accuracy at larger time-step than NIMROD’s semi-implicit methods, but there is a tradeoff in the size and cost of the algebraic systems that need to be solve at each step [44]. The leapfrog-based semi-implicit method of NIMROD and NIMSTELL solves one physical field at a time. We expect this property to be helpful for stellarator applications, as it is for others, in reducing the size of matrices used for preconditioning.

The remainder of this article is organized as follows: The MHD modeling is described in Section 2. Numerical methods are presented in Section 3. Test applications are presented in Section 4, and conclusion are discussed in Section 5. Additional topics related to NIMSTELL’s vector-potential option are covered in the appendices.

2. Mathematical model

In this article, we focus on solving the time-dependent equations of MHD. Our model is non-ideal in the sense that there is resistive electric field in Ohm’s law and viscous dissipation in the plasma fluid-velocity equation. For magnetic confinement, we include the important extensions of anisotropic thermal conduction and viscous stress, where the orientation of the anisotropy follows the magnetic field.

The nonlinear evolution equations for the plasma-fluid model solved by NIMSTELL are:

$$\frac{\partial n}{\partial t} + \nabla \cdot (n\mathbf{V} - D_n \nabla n) = 0, \quad (1)$$

$$\rho \left(\frac{\partial}{\partial t} + \mathbf{V} \cdot \nabla \right) \mathbf{V} = \nabla \cdot \left[\frac{1}{\mu_0} \mathbf{B}\mathbf{B} - \mathbf{I} \left(\frac{B^2}{2\mu_0} + p \right) - \underline{\Pi} \right], \quad (2)$$

$$\frac{nk_B}{\Gamma-1} \left(\frac{\partial}{\partial t} + \mathbf{V} \cdot \nabla \right) T = -nk_B T \nabla \cdot \mathbf{V} - \nabla \cdot \mathbf{q} + Q. \quad (3)$$

Mass density ρ and electron particle density n are related by $\rho = mn$ with m being the effective ion mass for a given charge state. When the charge state is 1 and electrons and ions have the same temperature T , plasma pressure is $p = 2nk_B T$. The Γ -factor in Eq. (3) is the adiabatic index, and D_n is a numerical particle diffusivity that helps avoid numerical noise.

This single-fluid, single-temperature model is closed by the stress tensor $\underline{\Pi}$ and the heat-flux density $\mathbf{q}$, and Q represents a heat-source density. The conductive heat-flux closure for Eq. (3),

$$\mathbf{q} = -(\kappa_{\parallel} - \kappa_{\text{iso}}) \hat{\mathbf{b}} \hat{\mathbf{b}} \cdot \nabla k_B T - \kappa_{\text{iso}} \nabla k_B T, \quad (4)$$

is a computationally tractable approximation for high-temperature magnetized plasma that represents anisotropy through distinct parallel $\kappa_{\parallel}$ and perpendicular $\kappa_{\perp}$ thermal conductivities with respect to the magnetic direction vector $\hat{\mathbf{b}} = \mathbf{B}/B$. Similarly, a viscous closure for Eq. (2) uses

$$\underline{\Pi} = -\rho \left[\frac{1}{2} \nu_{\parallel} \hat{\mathbf{b}} \cdot \underline{\mathbf{W}} \cdot \hat{\mathbf{b}} (3\hat{\mathbf{b}}\hat{\mathbf{b}} - \underline{\mathbf{I}}) + \nu_{\text{iso}} \underline{\mathbf{W}} \right], \quad (5)$$

where $\underline{\mathbf{W}} = \nabla \mathbf{V} + \nabla \mathbf{V}^T - (2/3) \underline{\mathbf{I}} \nabla \cdot \mathbf{V}$ is the traceless rate of strain tensor and $\nu_{\parallel}$ and ν_{iso} are viscous diffusivities. Numerical particle diffusion can cause momentum and energy conservation errors. It is possible to add corrections to Eqs. (2-3) to preserve these conservation properties, as is done in NIMROD [45] and other codes.

NIMSTELL has options for solving either the magnetic field components or the magnetic vector potential in its time advance. The former solves the combined Faraday's/Ampere's equation with the resistive-MHD Ohm's law for electric field:

$$\frac{\partial}{\partial t} \mathbf{B} = \nabla \times \left(\mathbf{V} \times \mathbf{B} - \frac{\eta}{\mu_0} \nabla \times \mathbf{B} \right) + \kappa_b \nabla \nabla \cdot \mathbf{B}, \quad (6)$$

where η is the electrical resistivity¹ and $\mu_0 \mathbf{J} = \nabla \times \mathbf{B}$. The last term in (6) provides a diffusive correction to numerical divergence error [46], and it has proven effective in applications when used in combination with NIMROD's high-order spatial representation and flux-conserving boundary conditions [34]. NIMSTELL's alternative is to advance the vector potential using the resistive-MHD Ohm's law in the potential representation for electric field:

$$\frac{1}{\mu_0} \nabla \times \nabla \times \mathbf{A} - \underline{\sigma} \cdot \left(\mathbf{V} \times \nabla \times \mathbf{A} - \frac{\partial}{\partial t} \mathbf{A} - \nabla \chi \right) = 0. \quad (7)$$

Here, $\nabla \times \mathbf{A} = \mathbf{B}$, and $\underline{\sigma}$ is the electrical conductivity tensor, which may have anisotropy with respect to $\hat{\mathbf{b}}$. The scalar field χ is a potential that is used to set a gauge condition. The analysis performed in Appendix A leads us to use $\nabla^2 \chi = d_A \nabla \cdot \mathbf{A}$ to damp the divergence of the vector potential and approximately satisfy the Coulomb gauge.

Like NIMROD and other codes (see Ref. [47], for example), NIMSTELL's time advance is for nonlinear perturbations about a prescribed steady state without an amplitude ordering. This is not represented in Eqs. (1-7) for brevity. However, for example, the implemented continuity equation solves

$$\frac{\partial n}{\partial t} + \nabla \cdot (n_s \mathbf{V} + n \mathbf{V}_s + n \mathbf{V}) = \nabla \cdot D_n \nabla n, \quad (8)$$

where n_s and $\mathbf{V}_s$ are prescribed steady particle density and flow-velocity profiles that are assumed to satisfy time-independent equations, possibly including external sourcing. Applying NIMSTELL to

¹ Later, we also use " η " as an element coordinate, but meanings should be clear from context.

linear time-dependent problems is achieved by not including the nonlinear terms, such as $n\mathbf{V}$ in Eq. (8). The separated steady fields are also convenient for nonlinear simulations of experimentally determined conditions, but one needs to consider physical consistency carefully if the prescribed fields do not satisfy the steady-state limit of solved evolution equations. Alternatively, a user may set the prescribed fields to zero or other steady states, such as vacuum magnetic field, and represent the rest in the evolved fields. NIMSTELL has preprocessing capability to load stellarator equilibria from the DESC code [48] into its steady-state fields with NIMSTELL's geometric mapping based on the shape of the DESC-computed flux surfaces. For slab and cylindrical test cases, steady-state fields and geometry are generated within NIMSTELL's preprocessor.

There is flexibility in the boundary conditions and initial conditions for NIMSTELL time-dependent computations. Particle density and temperature may be prescribed with Dirichlet conditions, or particle flux and conductive thermal flux may be imposed. Typically, flow-velocity is set to zero at the boundaries, though it is not a requirement of the model. When solving for vector potential with the $H(\text{curl})$ representation described in Section 3, our computations apply homogeneous Dirichlet conditions to the tangential components of $\mathbf{A}$ and to the scalar field χ to represent electrically conducting walls. However, our basis vectors for the magnetic representation do not generally align with 3D stellarator geometries, so we do not directly impose Dirichlet conditions on $\hat{\mathbf{n}} \cdot \mathbf{B}$. Instead, it is practical to include numerically implicit damping of $\hat{\mathbf{n}} \cdot \mathbf{B}$ in a weighted surface term—see Section 3.4.

3. Numerical methods

The primary distinction between time-dependent codes for stellarators and those for toroidally symmetric systems lies in their spatial representations. Nonetheless, we first briefly summarize NIMSTELL's temporal discretization.

3.1 Semi-implicit advance

The numerical methods used to advance NIMSTELL's solution fields over time are largely the same as those used in NIMROD. Flow-velocity is staggered one-half step in time relative to the particle density, temperature, and the magnetic field. MHD waves are stabilized at arbitrarily large timestep values (Δt) by the semi-implicit (SI) operator $\mathcal{L}_{\text{si}}$ in the flow-velocity equation [49], [50]. It modifies Eq. (2) without loss of numerical consistency with the semi-discrete form of the advance of $\mathbf{V}$ from time level j to $j + 1$ being

$$\begin{aligned} & \rho^{j+\frac{1}{2}} \left(1 + \frac{\Delta t}{2} \mathbf{V}^j \cdot \nabla \right) \Delta \mathbf{V} + \frac{\Delta t}{2} \rho^{j+\frac{1}{2}} \Delta \mathbf{V} \cdot \nabla \mathbf{V}^j + f_v \Delta t \nabla \cdot \underline{\underline{\Pi}}(\Delta \mathbf{V}) + c_{\text{si}} \Delta t^2 \mathcal{L}_{\text{si}}(\Delta \mathbf{V}) \\ & = \Delta t \nabla \cdot \left[\frac{1}{\mu_0} \mathbf{B} \mathbf{B} - \mathbf{I} \left(\frac{B^2}{2\mu_0} + p \right) \right]^{j+\frac{1}{2}} - \Delta t \rho^{j+\frac{1}{2}} \mathbf{V}^j \cdot \nabla \mathbf{V}^j - \Delta t \nabla \cdot \underline{\underline{\Pi}}(\mathbf{V}^j) . \end{aligned} \quad (9)$$

Like NIMROD, NIMSTELL's semi-implicit operator is based on the linear ideal-MHD force operator, which draws from XTOR's formulation [51]. Applying the SI method to an underlying staggered leapfrog scheme, instead of XTOR's predictor-corrector, leads to wave propagation without dissipative numerical error [50], [44], so numerical smoothing must arise from dissipative terms in the system of equations. The anisotropies in NIMSTELL's SI operator are with respect to its 3D fields, which requires algebraic coupling among toroidal Fourier components in the advance of $\mathbf{V}$. The benefits of using a 3D SI operator has previously been considered for some NIMROD computations [38], but NIMROD computations typically use 2D SI operators to minimize Fourier coupling in the algebraic solves. In nonlinear computations, NIMSTELL's SI operator is updated over time to preserve its accuracy. Also, as indicated in Eq. (9), the viscous operator is implicit with

numerical centering coefficient f_v , $0 \leq f_v \leq 1$, and advection is expanded one-half step from the previous time level [44].

3.2 Spatial representation

We consider 3D toroidal domains Ω by creating a non-overlapping partition of any poloidal plane into 2D finite elements and smoothly extending them over ζ . The union of the partition includes the domain boundary $\partial\Omega$. Physical coordinates for each element have an isoparametric mapping by using the spatial expansion that is used for the expansions of n , T , and each component of $\mathbf{V}$, and for $\mathbf{B}$ when Eq. (6) is solved. With biquadratic and higher-order finite elements, element edges in a constant- ζ surface may be curved. Within each element e ,

$$R_e(\xi, \eta, \zeta) = \sum_k \alpha_k(\xi, \eta) [R_{e_{k0}} + \sum_{n=1}^N R_{e_{kn}} e^{in\zeta} + cc.], \quad (10a)$$

$$Z_e(\xi, \eta, \zeta) = \sum_k \alpha_k(\xi, \eta) [Z_{e_{k0}} + \sum_{n=1}^N Z_{e_{kn}} e^{in\zeta} + cc.], \text{ and} \quad (10b)$$

$$\phi_e(\xi, \eta, \zeta) = \sum_k \alpha_k(\xi, \eta) [\phi_{e_{k0}} + \sum_{n=1}^N \phi_{e_{kn}} e^{in\zeta} + cc.] + \zeta \quad (10c)$$

with the element coordinates $0 \leq \xi, \eta \leq 1$ and $0 \leq \zeta \leq 2\pi$. The notation “cc.” refers to the complex conjugate of the previous term. The element bases $\alpha_k(\xi, \eta)$ for quadrilateral elements are the tensor products of the 1D Lagrange polynomials, $\lambda_i(\xi)$ and $\lambda_j(\eta)$, that have nodes at the points used for Gauss-Legendre-Lobatto (GLL) numerical integration. These polynomials form a space that is numerically equivalent to a finite expansion in Legendre polynomials while having the convenience of a nodal representation where coefficients represent function values at the node points [39], [40]. Ensuring function-value continuity along element interfaces just requires sharing coefficient values. For comparison, NIMROD’s mapping does not have an expansion in ζ . Apart from slab-geometry tests, we have not yet explored the potential benefits of distorting the physical angle ϕ over the generalized angle ζ , so the $\phi_{e_{k0}}$ and $\phi_{e_{kn}}$ coefficients are 0 in the computations presented in Section 4. The expansion in Eq. (10) is also applicable to topologically toroidal but geometrically straight configurations through a different choice of Jacobian and metric computations.

The conductive heat-flux closure (4) makes the temperature equation (3) second-order with respect to spatial derivatives. This implies that the temperature field is in the H^1 Hilbert space, where the function value and its first derivatives are square-integrable. The H^1 space is also appropriate for the components of $\mathbf{V}$ when using the viscous closure (5) and for particle density when including the artificial particle diffusion in Eq. (1). When solving Eq. (6) with positive values of resistivity η and divergence-error diffusivity κ_b , the magnetic field components are also in H^1 . Numerically, this representation allows nonzero magnetic divergence error, but with biquadratic and higher-order expansions, the diffusive correction term [46] is effective for magnetic-confinement applications [34]. It is also convenient to have all solution fields in the same space of functions.

Apart from the 3D geometry, the other significant difference between NIMSTELL and NIMROD is the option to solve the vector potential. This option uses the quadrilateral Nedelec edge elements [52] that are widely applied for computational electromagnetics to facilitate numerical convergence in the presence of sharp corners. Figure 1 shows a sketch of edge elements to help visualize the expansion. The union of these elements form a subspace of the $\mathbf{H}(\text{curl})$ space, where the vector functions and their curl are square-integrable. We combine 2D edge elements for the poloidal plane with 1D finite Fourier series over the generalized toroidal angle. In curvilinear-coordinate terms, the edge elements are expansions in reciprocal basis vectors of the mapping for a given element. Within each element, but without explicitly showing the Fourier expansion in ζ ,

$$\begin{aligned} \mathbf{A}_e(\xi, \eta, \zeta, t) = & \sum_{k,l} A_{e,\xi_{k,l}}(\zeta, t) \zeta_k(\xi) \lambda_l(\eta) \nabla \xi + \sum_{k',l'} A_{e,\eta_{k',l'}}(\zeta, t) \lambda_{k'}(\xi) \zeta_{l'}(\eta) \nabla \eta \\ & + \sum_{k'',l''} A_{e,\zeta_{k'',l''}}(\zeta, t) \lambda_{k''}(\xi) \lambda_{l''}(\eta) \nabla \zeta. \end{aligned} \quad (11)$$

The λ_k functions again represent 1D nodal polynomials based on GLL integration. However, as discussed in Appendix A, we consider expansions of different degree for $\mathbf{A}$ relative to those for the other physical fields. The ζ_l functions are nodal polynomials based on Gauss-Legendre (GL) integration, where there are no nodes at the endpoints of the interval. The Nedelec edge elements have ζ_l functions of one degree lower than those of the λ_k functions. With the endpoint coefficients for the λ_k bases of Eq. (11) being shared across element borders, tangential components of $\mathbf{A}(R, Z, \phi)$ are continuous regardless of the degree of the polynomial bases [53]. Another property of edge elements is that the normal component of their curl is continuous across element borders—see Appendix B. This satisfies the requirement of $\mathbf{B} \cdot \hat{\mathbf{n}}$ continuity across any interface. That $\nabla \cdot \mathbf{B} = 0$ anywhere within an element follows directly from $\mathbf{B} = \nabla \times \mathbf{A}$ with $\mathbf{A}$ being defined everywhere within each $\mathbf{H}(\text{curl})$ element. Thus, this optional representation eliminates numerical $\nabla \cdot \mathbf{B}$ error to machine roundoff level.

When it is used, the NIMSTELL computation of $\mathbf{A}$ proceeds by solving Eq. (7) and the gauge relation in weak form. A step of the time advance determines the change in vector potential $\Delta \mathbf{A} \in \mathbf{A}_{h,N,p}$ and the scalar $\chi \in X_{h,N,p}$ such that

$$\begin{aligned} & \int dVol \left\{ f_\eta \frac{\eta_s}{\mu_0} (\nabla \times \mathbf{G}) \cdot \nabla \times \Delta \mathbf{A} + \mathbf{G} \cdot [\tilde{\sigma}_\delta \hat{\mathbf{b}} \hat{\mathbf{b}} + \tilde{\sigma}_\perp \mathbf{I}] \cdot \left(\nabla \chi + \frac{1}{\Delta t} \Delta \mathbf{A} \right) - \frac{\tilde{\sigma}_\perp}{2} \mathbf{G} \cdot \mathbf{V} \times (\nabla \times \Delta \mathbf{A}) \right. \\ & \quad \left. + \nabla \vartheta \cdot [\nabla \chi - f_d d_A \Delta \mathbf{A}] \right\} \\ & = \int dVol \left\{ \tilde{\sigma}_\perp \mathbf{G} \cdot \mathbf{V} \times (\nabla \times \mathbf{A}) - \frac{\eta_s}{\mu_0} (\nabla \times \mathbf{G}) \cdot \nabla \times \mathbf{A} + d_A \nabla \vartheta \cdot \mathbf{A} \right\} \end{aligned} \quad (12)$$

for all $\mathbf{G} \in \mathbf{A}_{h,N,p}$ and all $\vartheta \in X_{h,N,p}$. Here, $\mathbf{A}_{h,N,p}$ is the product subspace of 2D $\mathbf{H}(\text{curl})$ elements and the 1D Fourier exponential bases, and $X_{h,N,p}$ is the continuous H^1 subspace of 2D nodal spectral elements of the same degree as λ_k and the 1D Fourier bases. The mesh spacing is represented by h , the polynomial degree is represented by p , and the largest toroidal Fourier wavenumber is N . We define a resistive scaling factor η_s so that the electrical conductivities in Eq. (12) are normalized; $\sigma_\perp = \eta_s^{-1} \tilde{\sigma}_\perp$ and $\sigma_\parallel = \eta_s^{-1} \tilde{\sigma}_\parallel$. Also, $\tilde{\sigma}_\delta \equiv \tilde{\sigma}_\parallel - \tilde{\sigma}_\perp$. Although temporal indices are suppressed, the vector $\mathbf{A}$ is from level $j + 1/2$, and with the staggering, $\mathbf{V}$ is at level $j + 1$. The numerical coefficients f_η and f_d define the centering for resistive diffusion and for the gauge relation.

The low-resolution 2D mesh of elements shown in Fig. 2 is one that is used for the cylindrical interchange computation of Section 4.1. Topologically, it is typical of meshes used with NIMSTELL. The outer part is a polar mesh but near the axis, it transitions to a rectangular grid. We have not programmed triangular elements or degenerate quadrilateral elements, which would be needed to use a strictly polar mesh. The transitioning mesh also avoids extremely small element dimensions. When using DESC for stellarator equilibria, the outer topologically polar mesh is shaped to align with equilibrium magnetic flux surfaces. As also indicated by Fig. 2, the isoparametric mapping allows NIMSTELL to curve element sides.

3.3 Numerical interchange stabilization

When a numerical mesh is constructed and a set of basis functions is chosen, the spectral-element/Fourier representation for dependent variables yields a particular subspace of the H^1 space. The time-dependent computation then determines, by some means, the best set of

functions from a given subspace at each step. However, it is well known that difficulties arise in hyperbolic problems when attempting to model configurations with insufficient numerical resolution for a given set of dissipation coefficients, see Refs. [54], [55], and [39] for example. In these conditions, finite-element “stabilization” methods [56], [57] have proven effective for many applications, including MHD [58], [59], [60], [61].




Figure 1. Sketches of the reference element with (ξ, η) element coordinates (left) and the $\nabla \xi$ (center) and the $\nabla \eta$ (right) basis vector functions after mapping to physical coordinates.




Figure 2. Finite-element mesh of the lowest-resolution cylindrical interchange computation. The elements are curved and transition from a polar grid (12 radial $\times$ 16 azimuthal) to a rectangular grid of 16 elements at the axis.

At large Δt -values, the second-order derivatives on $\Delta \mathbf{V}$ in NIMROD’s semi-implicit operator enhance the coercivity of the computation for $\Delta \mathbf{V}$, itself. However, without further stabilization, the time-advance converges on resonant interchange from the unstable side, meaning that numerically computed growth rates are larger at low spatial resolution than their converged values. As noted in Ref. [62], this is problematic for nonlinear time-dependent computations when physical modes saturate but growing numerical modes arise. For NIMROD, we previously developed specialized

stabilization terms to make the representation converge from the stable side when using equal-order expansions of all physical fields [63].

The parabolic form of these terms is implemented in NIMSTELL for use with its equal-order $H1$ representation of magnetic-field components. At each step the flow-velocity advance finds $\Delta \mathbf{V} \in \mathbf{V}_{h,N,p}$ and $\sigma, \varpi \in F_{h,N,p}$ such that

$$\begin{aligned}
& \int dVol \left\{ \rho \mathbf{w}^* \cdot \Delta \mathbf{V} + \Delta t \mathbf{w}^* \cdot \left(\frac{\rho}{2} \mathbf{V} \cdot \nabla \Delta \mathbf{V} + \frac{\rho}{2} \Delta \mathbf{V} \cdot \nabla \mathbf{V} \right) + \Delta t (v^* \sigma + \mu^* \varpi) \right. \\
& + C_{\text{si}} \Delta t^2 \left[\frac{1}{\mu_0} \nabla \times (\mathbf{w}^* \times \mathbf{B}) \cdot \nabla \times (\Delta \mathbf{V} \times \mathbf{B}) + \Gamma p (\nabla \cdot \mathbf{w}^*) (\nabla \cdot \Delta \mathbf{V}) \right] \\
& - \frac{C_{\text{si}} \Delta t^2}{2} \left[\mathbf{w}^* \cdot \mathbf{J} \times \nabla \times (\Delta \mathbf{V} \times \mathbf{B}) + \Delta \mathbf{V} \cdot \mathbf{J} \times \nabla \times (\mathbf{w}^* \times \mathbf{B}) - ((\nabla \cdot \mathbf{w}^*) \Delta \mathbf{V} + (\nabla \cdot \Delta \mathbf{V}) \mathbf{w}^*) \cdot \nabla p \right] \\
& - f_v \Delta t \nabla \mathbf{w}^H : \underline{\Pi}(\Delta \mathbf{V}) \} \\
& = \int dVol \left\{ \Delta t \left[\mathbf{w}^* \cdot \mathbf{J} \times \mathbf{B} + p \nabla \cdot \mathbf{w}^* + \nabla \mathbf{w}^H : \underline{\Pi}(\mathbf{V}) - \rho \mathbf{V} \cdot \nabla \mathbf{V} \cdot \mathbf{w}^* \right] \right. \\
& \left. + \sqrt{d_d P_c \Delta t} (v^* \nabla \cdot \mathbf{V} - \sigma \nabla \cdot \mathbf{w}^*) + \sqrt{d_v \Delta t} \mathbf{B} \cdot (\mu^* \nabla \times \mathbf{V} - \varpi \nabla \times \mathbf{w}^*) \right\} \quad (13)
\end{aligned}$$

for all $\mathbf{w} \in \mathbf{V}_{h,N,p}$ and $v, \mu \in F_{h,N,p}$. Here, $\mathbf{V}_{h,N,p}$ is the continuous $H1$ subspace. We also define $P_c \equiv B^2 + \Gamma p$ as the total pressure for compressional waves. Setting homogeneous essential conditions on $\mathbf{V}$ along the domain surface eliminates surface terms that would result from integration by parts. As described in [63], $F_{h,N,p}$ is an incomplete space of discontinuous polynomials in the meshed plane with just the p -th Legendre polynomial along at least one element side. Including the $\sigma \nabla \cdot \mathbf{w}^*$ and $\varpi \nabla \times \mathbf{w}^*$ projections in the $\Delta \mathbf{V}$ -equation does not affect numerical consistency, since the projections onto the incomplete $F_{h,N,p}$ space vanish with increasing resolution of sufficiently smooth functions. The dimensionless coefficients d_d and d_v are of order 1. In the limit of small viscosity, the σ -projection provides a response to $\nabla \cdot \mathbf{V}$ at the scale of the nodes, and the ϖ -projection provides a response to parallel vorticity at that scale, which yields the desirable converge from the stable side. In comparison with Eq. (12) of Ref. [34] and Eq. (20) of Ref. [63], NIMSTELL uses complete 3D fields in the interchange-stabilizing projections and in its semi-implicit operator.

Comparing the flow-velocity equation for the two magnetic-field options, the contribution of the Lorentz force density on the right side of (13) is expressed as implemented in NIMROD and in NIMSTELL for the $H1$ representation of magnetic field components. With components of $\mathbf{B}$ in $H1$, current density may be evaluated from $\nabla \times \mathbf{B}$ at the numerical integration points. For the $\mathbf{H}(\text{curl})$ representation of vector potential, we use the divergence of the Maxwell stress tensor, as shown in Eqs. (2) and (9). This divergence is integrated by parts when projected onto the test vector $\mathbf{w}$, since differentiation of $\mathbf{A}$ is limited to a single curl operation. However, $\mathbf{J}$ and $\nabla \mathbf{B}$ are needed to update the semi-implicit operator in nonlinear computations. Thus, when using $\mathbf{A}$ in $\mathbf{H}(\text{curl})$ in nonlinear computations, an extra projection of $\mathbf{B}$ components onto an $H1$ expansion is performed during each time advance for this purpose. The interchange-stabilization terms are not used when selecting the vector-potential option.

3.4 Magnetic-field surface term

The $H1$ representation of magnetic-field components uses cylindrical unit vectors $(\hat{\mathbf{R}}, \hat{\mathbf{Z}}, \hat{\boldsymbol{\phi}})$ with respect to the geometric axis through the center of a torus. For stellarator configurations, the outward surface normal vector $\hat{\mathbf{n}}$ generally has three nonzero cylindrical components. Taking ξ to be the element coordinate that is tangent to the domain boundary in an element that is adjacent to the boundary, $\hat{\mathbf{n}} \propto \partial \mathbf{R} / \partial \xi \times \partial \mathbf{R} / \partial \zeta$. In NIMSTELL, the mapping has a Fourier expansion over ζ , so it

is not convenient to impose $\hat{\mathbf{n}} \cdot \mathbf{B} = 0$ along the surface through essential conditions on components of $\mathbf{B}$ without setting all components to 0, which is too restrictive. To approximate the $\hat{\mathbf{n}} \cdot \mathbf{B} = 0$ Dirichlet condition, we instead add a surface term that rapidly damps $\hat{\mathbf{n}} \cdot \mathbf{B}$ along the surface. The formulation for advancing $\mathbf{B}$ is to find $\Delta \mathbf{B} \in \mathbf{V}_{h,N,p}$ such that

$$\begin{aligned} & \int dVol \left\{ \mathbf{C}^* \cdot \Delta \mathbf{B} + \Delta t \nabla \times \mathbf{C}^* \cdot \left(f_\eta \frac{\eta}{\mu_0} \nabla \times \Delta \mathbf{B} - \frac{1}{2} \mathbf{V} \times \Delta \mathbf{B} \right) + f_{\kappa_b} \Delta t \kappa_b (\nabla \cdot \mathbf{C}^*) (\nabla \cdot \Delta \mathbf{B}) \right\} \\ & + \oint dS [(1 + d_s)(\hat{\mathbf{n}} \cdot \mathbf{C}^*)(\hat{\mathbf{n}} \cdot \Delta \mathbf{B})] \\ & = \Delta t \int dVol \left\{ \nabla \times \mathbf{C}^* \cdot \left(\mathbf{V} \times \mathbf{B} - \frac{\eta}{\mu_0} \nabla \times \mathbf{B} \right) - \Delta t \kappa_b (\nabla \cdot \mathbf{C}^*) (\nabla \cdot \mathbf{B}) \right\} - \oint dS [d_s(\hat{\mathbf{n}} \cdot \mathbf{C}^*)(\hat{\mathbf{n}} \cdot \mathbf{B})] \quad (14) \end{aligned}$$

for all $\mathbf{C} \in \mathbf{V}_{h,N,p}$, the same subspace used for components of $\mathbf{V}$. For $d_s > 1$ but not so large as to make the resulting matrix ill-conditioned, the surface term penalizes $\hat{\mathbf{n}} \cdot \mathbf{B}$ along the surface. In practice, many of our computations set $d_s \sim 10$. We remark that $\hat{\mathbf{n}}$ in toroidally symmetric computations also generally does not align with either $\hat{\mathbf{R}}$ or $\hat{\mathbf{Z}}$. Setting $\hat{\mathbf{n}} \cdot \mathbf{B} = 0$ along the domain surface of NIMROD computations is accomplished through local constraint relations at each finite-element node along the surface. The distinction is that with $\hat{\mathbf{n}} \cdot \hat{\mathbf{R}}$ and $\hat{\mathbf{n}} \cdot \hat{\mathbf{Z}}$ independent of ϕ , the NIMROD constraints are imposed separately for each toroidal Fourier component of $\mathbf{B}$.

We also contrast this surface term for the $H1$ representation of magnetic components with the essential condition imposed on the tangential components of $\mathbf{A}$ for conducting surfaces when Eq. (12) is used. Staying with the convention of the logical coordinate η being constant along the domain surface for adjacent elements, $\partial \mathbf{R} / \partial \xi$ and $\partial \mathbf{R} / \partial \zeta$ are tangent to the surface. With the $\mathbf{H}(\text{curl})$ bases, setting $\hat{\mathbf{n}} \times \mathbf{A} = 0$, is accomplished by setting the expansion coefficients $A_{e,\xi_{k,l}}$ and $A_{e,\zeta_{k,l}}$ in Eq. (11) to 0 for nodes along the surface, for all Fourier components, and for all time. Setting the coefficients of χ to zero, or any unique value, along the surface then ensures $\hat{\mathbf{n}} \times \mathbf{E} = 0$.

3.5 Algebraic solves

NIMSTELL uses preconditioned conjugate gradients (CG) and generalized minimal residual (GMRES) Krylov-space methods to solve the linear algebraic systems for the implicit and semi-implicit time-advance equations. With 3D geometry, the algebraic systems have coupling among the Fourier components, even for linearized systems of equations. Because the complete matrix is too large to form in general, NIMSTELL computes the implicit operations on the Krylov-space vectors through the “matrix-free” method. This method rearranges the order of the summations for numerical integration and matrix-vector products. The associated computational cost is performing finite-element operations at each step of the Krylov algorithm. Preconditioning applies approximate inverses at each step to accelerate convergence and is critical for making Krylov methods practical for large ill-conditioned systems [64]. The NIMSTELL implementation of CG and GMRES is similar to NIMROD’s, but the computational infrastructure to use the more general data layout for coefficients of $\mathbf{H}(\text{curl})$ expansions has been added.

Preconditioning operations in NIMROD are based on inverting non-overlapping submatrices (“blocks”) that cover different levels of coupling over the poloidal-plane but are diagonal in the toroidal Fourier basis. Typically, the most effective approach is to solve the sparse matrices that represent coupling over the entire (poloidal) plane of elements using the parallel SuperLU_DIST library solver [65]. The strong coupling among Fourier components from 3D stellarator geometry makes the block Fourier diagonal (bFd) approach ineffective for general NIMSTELL application. To address this, the bFd method has been expanded to include multiple toroidal Fourier components in the blocks that are solved by the direct solver. The programming allows the blocks to either cover

bands of Fourier components or user-defined families of Fourier components. The implementation allows overlap of the blocks, meaning that a Fourier component can be included in more than one block. The preconditioner averages the results of the block solves for the Fourier components that are included in more than one block for additive-Schwarz preconditioning [64].

Using preconditioner blocks that span continuous ranges of toroidal Fourier components is effective for single field-period ($N_p = 1$) stellarator configurations, but the effectiveness decreases with increasing N_p values. This can be appreciated from the illustrations shown in Fig. 3. The information on matrix-element scaling represents an $N_p = 5$ stellarator, and Fig. 3a shows that using preconditioning blocks for continuous ranges wastes computational effort and memory on negligible couplings. However, using blocks that include just the strongest couplings that result from the symmetry of the configuration (Fig. 3b) is effective. Table I compares the performance of these block preconditioning choices on a tearing-mode computation for an $N_p = 5$ configuration with elliptical cross section, the W7-A case of Ref. [32] and Section 4.2. For this linear computation, the judicious blocks of n -values that are coupled by the symmetry of the geometry and the equilibrium is as effective as using one matrix covering all Fourier components while requiring seven times less memory and allowing faster computation with the smaller matrices.




Figure 3. Illustration of the magnitude of matrix elements for the $N_p = 5$ case, toroidal resolution of $0 \leq n \leq 10$, and different preconditioner block choices. The allocation of preconditioning blocks in a) represents two stride=3 (“st3”) sets of blocks with overlap (“ov”) of 2 (red) or 5 (green). The allocation of preconditioning blocks in b) is user-prescribed according to the families (red, green, and violet) of n -values that have either $n_1 + n_2$ or $n_1 - n_2$ being a multiple of N_p , where n_1 and n_2 are two of the same family.

Table I. Comparison of Krylov-space solver performance on the $N_p = 5$ tearing-mode computation when preconditioning with the different Fourier block decompositions shown schematically in Fig. 3.

Block Composition (n)	Iteration Count	Memory Use
$\{0-4\}, \{3-7\}, \{6-10\}$	$\sim 15,000$	550 GB
$\{0-7\}, \{3-10\}$	$\sim 2,000$	1.4 TB
$\{0-10\}$	2	2.4 TB
$\{1,4,6,9\}, \{2,3,7,8\}, \{0,5,10\}$	2	350 GB

4. Test applications

The applications presented in this article test resonant interchange in cylindrical geometry, magnetic tearing in the toroidal W7-A geometry, and magnetic tearing in a quasi-axisymmetric configuration. We also demonstrate convergence on anisotropic thermal conduction in a straight stellarator configuration.

4.1 Linear resonant interchange

Resonance for low-frequency dynamics is the condition that helical wavefronts align with the equilibrium magnetic field so that the wavenumber vector, $\mathbf{k} = mr^{-1}\hat{\theta} + 2\pi nL_z^{-1}\hat{z}$ in cylindrical geometry, satisfies $\mathbf{k} \cdot \mathbf{B} = 0$. To verify our implementation of the numerical stabilization terms in Eq. (13) and to demonstrate performance of the vector-potential option, we consider the same cylindrical interchange instability from Ref. [2] that was used in Ref. [63]. The parameterized 1D equilibrium profile has uniform axial field and mass density with $B_z = 1$ and $\rho = 1$ after normalization. Axial current density leads to the cylindrical safety-factor profile of

$$q(r) = \frac{2\pi r B_z(r)}{L_z B_\theta(r)} = 2\pi \frac{1+c_2^2 r^2}{L_z c_1}$$

over $0 \leq r \leq 1$, where q is the number of axial passes per azimuthal pass of a field-line and L_z is the length of the periodic cylinder. The Suydam parameter for resonant ideal interchange stability in these profiles is

$$D_s(r) \equiv -\frac{2\mu_0}{rB_z^2} \left(\frac{q}{dq/dr} \right)^2 \frac{dP}{dr} = \frac{c_1^2}{(1+c_2^2 r^2)c_2^4 r^2},$$

and the criterion for stability is $D_s \leq 1/4$. For parameters with $c_1 = 4/7$ and $c_2 = 10/7$, the helical mode with azimuthal wavenumber $m = 4$ and axial wavenumber $n = -1$ for $L_z = 3.523$ is resonant along the cylindrical surface defined by $r = 0.371$ where $D_s = 0.443$. Our reference value of the growth rate is $\gamma_{\text{eig}} = 6.771 \times 10^{-3} \tau_a^{-1}$, where $\tau_a = a\sqrt{\mu_0 \rho}/B_z$. This value is computed from a 1D ideal-MHD eigenvalue solver [63] applied to magnetically aligned components of the displacement vector [2] for the selected helical mode. We have used Richardson extrapolation to refine the eigenvalues ($-\gamma_{\text{eig}}^2$) computed with 1280 and 2560 linear elements over r .

All physical dissipation parameters and D_n are effectively 0 in the NIMSTELL computations considered here. To avoid artificial magnetic reconnection, we also use a relatively small normalized value of $\kappa_b = 10^{-5}$. With the chosen Δt values of 5 for the mesh-spacing scan and 10 for the polynomial-degree scan, diffusion for divergence correction per step is on the spatial scale of 10^{-2} . As a check, reducing κ_b to 10^{-6} in biquartic computations with 128 azimuthal elements over the azimuthal angle decreases the growth rate by 0.2%. Although this form of error control has greater effect at coarser resolution (for example, 0.7% at 64 azimuthal elements), it does not adversely affect our convergence tests. With respect to temporal resolution, spatially resolved results for $\Delta t = 5$ and 10 agree to within 0.015%; at $\Delta t = 40$ the error is only 0.4% for this ideal mode. For the stabilization terms in Eq. (13), we use the standard NIMROD coefficients of $d_d = 1$ and $d_v = 0.1$.

Our results show that NIMSTELL converges from the stable side when refining either with respect to mesh spacing or with respect to the polynomial degree of the spectral-element basis functions. The error shown in Fig. 4 is $(\gamma_{\text{eig}} - \gamma_{\text{NIM}})\tau_a$, and all errors are positive, hence the NIMSTELL growth rates are smaller than the converged value. The meshes have the same distribution of elements shown in Fig. 2. The horizontal axis of Fig. 4a indicates the number of elements in the azimuthal direction of the polar region with the number of elements in the radial direction varying proportionately. A biquartic computation with the 208-element mesh of Fig. 2 (16 over the

azimuthal angle) does not produce a growing mode over $10^4 \tau_a$. A mesh with four times the number of elements is the lowest-resolution computation of Figs. 4a-b. Figure 4a also shows three results computed with NIMROD. NIMSTELL yields the same growth rates to at least three significant digits, confirming NIMSTELL's implementation of the interchange-stabilization terms. We comment that even for an axisymmetric mesh and equilibrium, the implementations differ. The NIMSTELL computations minimally need Fourier components $0 \leq n \leq 1$ and use Fourier transforms for numerical integrals on the ζ -grid, whereas NIMROD computes linear results for $n = 1$ alone without Fourier transforms. Another NIMSTELL computation shown in Fig. 4a has its mesh twisting a full turn over ζ . In this case, the equilibrium and the mesh have $n > 0$ components. The eigenmode shape only twists 1/4 of a turn over 2π in ζ , so the mesh and the mode are intentionally misaligned. The solution then has Cartesian x and y vector components in the $n = 2$ and $n = 4$ Fourier components, and the scalar fields and axial vector components are in $n = 3$. Its growth rate matches the NIMSTELL value with the axisymmetric mesh.

The eigenfunction for this mode is characterized by its sharp features at the resonance (see Fig. 15b of Appendix A). Conditions are close to marginal stability, and only inertia keeps components of flow-velocity from being discontinuous. Thus, the convergence properties of the representation are not optimal for the eigenmodes. For sufficiently smooth functions, C^0 biquartic elements will achieve fifth-order convergence—see Theorem 3.3 of Ref. [54]. The lower rate observed from Fig. 4a can be expected from the nearly singular behavior of the eigenfunction. The roughly linear behavior of error in the polynomial-degree scan plotted on the semi-log scale of Fig. 4b is suggestive of geometric convergence [40], but one should be skeptical with the limited range of results. In fact, it is surprising that this spectral-resolution scan yields the accuracy that it does, given the application. We also emphasize that although the location of the resonant surface is known *a priori*, no mesh packing is used to represent application of NIMSTELL without such prior knowledge.




Figure 4. Convergence of NIMSTELL on the $m = 4, n = 1$ cylindrical interchange for $L_z = 3.523$ using the $H1$ magnetic-component option. Convergence is shown with respect to (a) mesh spacing for biquartic elements and (b) polynomial degree of the spectral-element basis functions for a mesh of 832 elements (32 over θ). Results from NIMROD and a single NIMSTELL computation with a twisting mesh are also plotted in (a), together with a power-law scaling of mesh spacing.

We find that the same set of ideal-MHD computations are not tractable with NIMSTELL's present vector-potential implementation. With $\eta = 0$ in this ideally unstable case, numerical noise grows at $\gamma\tau_a \approx 0.03$ with biquadratic elements and meshes of either 32 or 64 elements over θ . A minimum amount of electrical resistivity is apparently needed for a given mesh. To show this, we consider the resistive version of the ideally unstable conditions and resistive interchange in conditions that are ideally stable. The latter has $c_1 = 1/2$, $c_2 = 5/4$, and $L_z = 2\pi$ to make the (3,1) mode resonant at $r = 0.566$, where $D_s = 0.213$; noise also grows in an $\eta = 0$ computation with this equilibrium. In these computations, we vary η starting with a relatively large value of 10^{-4} and a coarse mesh and decrease η until the growth rate is inaccurate or produces a noisy result. The mesh is then refined, and the resistivity scan is continued.

As shown in Fig. 5a for the ideally unstable case, computations using biquartic elements with quintic λ_k for the vector- and scalar-potential expansion (11) and 64 elements over the azimuthal direction produce reliable results for $\eta \geq 10^{-7}$. At $\eta = 10^{-8}$, the numerical eigenmode has noise. Computations with 128 elements over θ are reliable through $\eta = 10^{-8}$, but produce rapidly growing noise at $\eta = 10^{-9}$. A computation with 256 elements over θ at $\eta = 10^{-9}$ produces the expected localized eigenmode with γ being 1.2% larger than the reference ideal value. In these ideally unstable conditions, resistive reconnection at the resonant surface allows faster than ideal growth rates at large- η [66]. In the ideally stable conditions (Fig. 5b), the resolved computations follow the scaling of $\gamma \sim \eta^{1/3}$ that is expected from asymptotic theory for resistive interchange [67]. However, we again find cases where the resistive diffusion is insufficient to avoid noise at a given level of spatial resolution.




Figure 5. Resistive scaling of NIMSTELL on the $m = 4, n = 1$ cylindrical interchange using the $\mathbf{H}(\text{curl})$ vector-potential option. Results are for (a) the ideally unstable ($m = 4, n = 1$) $L_z = 3.523$ case and (b) the ideally stable ($m = 3, n = 1$) $L_z = 2\pi$ case. The two series of computations have the same basis functions but different meshes, as indicated by the number of elements over θ .

Where the interchange computations with vector potential do not produce noise, under-resolved computations grow more slowly than converged results. Momentarily ignoring the noisy results, this option also converges from the stable side, which is consistent with the numerical eigenvalue

analysis with 1D meshing discussed in Appendix A. We infer that the 1D analysis is not capable of identifying the issue that leads to growing noise. A plausible explanation is simply that the resistive diffusion term, the first term on the left of Eq. (12), is needed to make the $\mathbf{H}(\text{curl})$ vector-potential formulation coercive. Appendix A describes numerical eigenmodes that do not couple to flow or excite other response for ideal MHD with this representation, and their existence is related. While the damping added for the gauge condition contributes to coercivity and removes other numerical modes in the ideal-MHD analysis, it is not sufficient for reliable numerical algebra when resistive diffusion is minimized.

To check nonlinear evolution with the $H1$ magnetic-field representation, we compare NIMSTELL and NIMROD results on the profile with $c_1 = 4/7$ and $c_2 = 10/7$, but $L_z = 4\pi/3$, which has its marginal ideal stability point at $r = 0.466$. The (3,1) mode is resonant at $r = 0.264$ and is the only $n = 1$ mode that is ideally unstable. For $n = 2$, the (6,2) and (7,2) modes are ideally unstable. Although larger- n wavenumbers have more unstable modes, the smaller wavelengths lead to smaller growth rates with viscous dissipation in the system, and they are driven primarily by nonlinear coupling that takes energy from the dominant modes. Both computations use the numerical stabilization method described in Section 3.3. We set $S = \tau_r/\tau_a = \mu_0 a^2/\eta\tau_a = 10^8$ and $\text{Pm} = \mu_0 v_{\text{iso}}/\eta = 10$. There is also a small level of thermal conduction with κ_{iso} set to $n k_B (\Gamma - 1) v_{\text{iso}}$. The particle density is initially uniform, so the pressure gradient is in the temperature profile. For numerical purposes, we set $D_n = v_{\text{iso}}$, and the normalized $\kappa_b = 0.01$. The two computations use the same mesh of 31×40 elements over an annular region surrounding 100 at the axis. The basis functions are biquintic, and the Fourier expansion over z retains $0 \leq n \leq 21$. Also, although it is not a steady state for all equations, we retain the MHD equilibrium in the steady-state fields of our expansion, as if unspecified sources maintain it. Symmetric ($n = 0$) distortions develop in the computations through nonlinear effects, not through dissipation.

Figure 6a-b compares the NIMROD and NIMSTELL magnetic and kinetic fluctuation energies for $0 \leq n \leq 4$. The traces lie on top of each other through the initial nonlinear saturation phase and begin to deviate after that. Larger- n harmonics have a similar level of agreement. Figure 6c shows the total pressure in the $z = 0$ plane at $t = 500 \tau_a$. After this time, the evolution becomes turbulent. Quantitative discrepancies then accumulate, possibly due to the iterative solver tolerance, despite the fact that the two codes apply the same numerical methods.




Figure 6. Comparison of fluctuation energies (a-b) from NIMROD and NIMSTELL computations for $1 \leq n \leq 4$ of the nonlinear interchange computation and (c) the sum of perturbed and equilibrium pressure in the $z = 0$ plane from the NIMSTELL computation at $t = 500 \tau_a$. For reference, the energy in the equilibrium magnetic field is 6.8 in this normalized configuration.

4.2 Resistive tearing

We have verified NIMSTELL on linear resistive tearing in cylindrical geometry, but here we focus on toroidal stellarator computations. The first configuration represents the W7-A stellarator [68], where the elliptical cross section rotates over the toroidal angle with the toroidal field-period (N_p) of five. The aspect ratio is large with the toroidal major radius $R = 1.99$ m and major and minor axes of the ellipse being 0.1178 m and 0.09131 m. The experiment had an Ohmic transformer, and toroidal current was induced in the plasma in some discharges. We use the profile specifications from Ref. [32], as applied for the JOREK code, repeated here for convenience:

$$I_n(\psi_{tn}) = 3\psi_{tn} - 3\psi_{tn}^2 + \psi_{tn}^3 \text{ and}$$

$$P(\psi_{tn}) = P_a - (P_a - P_b)\psi_{tn},$$

where ψ_{tn} is the normalized toroidal magnetic flux function that varies from 0 on the magnetic axis to 1 at the plasma surface, I_n is the toroidal current normalized by 17.5 kA, and for this low plasma- β case $P_a = 1$ Pa and $P_b = 0.01$ Pa. Here, we consider just one set of dissipation coefficients: $\eta = 1.938 \times 10^{-6} \Omega\text{m}$, $\text{Pm} = 0.001$, $v_{\parallel}/v_{\text{iso}} = 1000$, and $(\Gamma - 1)\kappa_{\text{iso}}n/v_{\text{iso}} = D_n/v_{\text{iso}} = 1$; a resistivity-scaling study will be presented in a forthcoming paper on benchmarking with results from other codes. We also set n_s to the uniform value of $2 \times 10^{20} \text{ m}^{-3}$, and $m_i = 1.67 \times 10^{-27} \text{ kg}$. With $RB_\phi = 4.68 \text{ Tm}$ at the plasma surface, the relevant Alfvén time is $\tau_A = 0.548 \mu\text{s}$. Given the specifications from Ref. [32] being in physical units, our NIMSTELL computations for this application are in SI units. Our DESC computation for the equilibrium and geometry uses harmonics of N_p to $N = 40$, though the shape of the equilibrium is described by $n = 0$ and $n = 5$. For DESC's Zernike polynomials, the radial resolution has $L = 14$, and the azimuthal resolution has $M = 12$.

Computed growth rates from scans of spatial and temporal resolution for linear computations are shown in Fig. 7. The poloidal resolution is checked with a mesh of 26 radial $\times$ 48 azimuthal elements surrounding an inner region of 144 elements, and there is packing near the $\iota = 1/2$ surface. Figure 7a shows the results from scanning the degree of the element basis functions for both magnetic representations. The horizontal axis is the polynomial degree used for the $\mathbf{V}$ -expansion for all computations, and the continuous λ_k polynomials for the $\mathbf{H}(\text{curl})$ representation for $\mathbf{A}$ and for χ is of one degree larger. Results for the $H1$ magnetic representation include computations with and without the stabilization terms of Eq. (13), again setting $d_d = 1$ and $d_v = 0.1$ when used. Except where noted, the $H1$ magnetic computations have $\kappa_b = 10^5$. At degree 2, the growth rate with the stabilization term is $3.1 \times 10^{-3} \tau_A^{-1}$ and is not shown in the figure. The computations at degrees 3 and 4 have growth rates within 0.1% of each other but are 3% larger than the computations with vector potential. A bicubic computation with the $H1$ magnetic representation and κ_b reduced to 10^4 (not shown) produces a growth rate within 0.7% of the vector-potential results, and the divergence error of 0.008 measured by $\int (\nabla \cdot \mathbf{B})^2 dVol / \int B^2 dVol$ is larger than the $H1$ results with $\kappa_b = 10^5$ by a factor of two. The scans of toroidal resolution in Fig. 7b with element bases of degree 3 and the $H1$ magnetic representation show variation when the maximum component is varied from $N = 15$ to $N = 21$ but not when increased further to $N = 31$. With $N \geq 21$, the magnetic divergence error of the $H1$ representation is approximately 100 times smaller than at $N = 15$. Thus, the diffusive correction does not appreciably enhance magnetic reconnection with $\kappa_b = 10^5$ when spatial resolution is sufficient. When scanning Δt with adequate spatial resolution (Fig. 7c), results from all three sets of computations are similar. Empirically, the temporal convergence rate is approximately 2.

We continue two of the linear computations nonlinearly with the same physical parameters, running each through saturation of the tearing mode. Both have toroidal resolution $0 \leq n \leq 21$, given that the linear results do not change when increasing N to 31. One case has the $H1$ magnetic representation with all basis functions having polynomial degree 3. Because plasma pressure is extremely small, we do not include the numerical-interchange stabilization terms. The second case uses the $\mathbf{H}(\text{curl})$ vector potential representation with polynomial degree of 2 for all basis functions except the continuous cubic λ_k in the $\mathbf{H}(\text{curl})$ expansion of $\mathbf{A}$ and for the scalar-potential χ . Figure 8 shows the evolution of the magnetic fluctuation energies for each computation, starting from the last part of the linear computations. The Fourier components that are not part of the $n = 1, 4, 6, 9, \dots$ linear mode family increase rapidly when the nonlinear terms are included at 1.2 ms, and the nonlinearly driven $n = 2$ component has larger energy than the $n > 1$ components of the linear tearing mode family soon thereafter. The peak $n = 1$ magnetic fluctuation energies are 0.094 J at 3.0 ms in the $H1$ magnetic-field case and 0.11 J at 3.2 ms in the lower-resolution $\mathbf{H}(\text{curl})$ vector-potential case. The peak energy levels agrees with JOREK's 0.1 J level that is evident in Fig. 9 of Ref. [32]. The resulting magnetic islands shown at the end of each simulation in Fig. 9 are also consistent with each other, apart from the narrow ($m = 5, n = 3$) island in the $H1$ result. They are also consistent with Fig. 10 of Ref. [32], although a different toroidal angle is shown. Both NIMSTELL computations indicate that the magnetic fluctuation energies decay slowly after reaching their maximum amplitude. We infer that this slow decay results from the profiles adjusting to the presence of the magnetic island on a resistive-transport timescale, which is approximately 9 ms for the resistivity value used in the computations.




Figure 7. NIMSTELL-computed linear growth rates for the W7-A tearing computation, varying (a) the polynomial degree of the representation with N from the Fourier expansion set to 15 and $\Delta t = 1.8 \tau_a$, (b) N with the polynomial degree of the $\mathbf{V}$ -expansion set to 3 and $\Delta t = 1.8 \tau_a$, and (c) Δt with $N = 21$ and the degree of $\mathbf{V}$ set to 3. The λ_k bases of the $\mathbf{H}(\text{curl})$ expansion for $\mathbf{A}$ is one degree larger than that of the $\mathbf{V}$ -expansion.




Figure 8. Evolution of magnetic fluctuation energy spectrum through nonlinear saturation for the NIMSTELL W7-A tearing computations run with (a) the $H1$ magnetic-field representation and (b) the $\mathbf{H}(\text{curl})$ vector-potential representation. The steady equilibrium field is not included in the traces.




Figure 9. Magnetic Poincaré plots at $t = 4$ ms for the NIMSTELL W7-A tearing computations with (a) the $H1$ magnetic-field representation and (b) the $\mathbf{H}(\text{curl})$ vector-potential representation.

A final tearing-mode case demonstrates a nonlinear application of NIMSTELL for a quasi-symmetric stellarator. The geometry and profile are a modification of the optimized $N_p = 2$ quasi-axisymmetric stellarator equilibrium of Ref. [69]. Starting from the DESC team's example input for this equilibrium [70], we add the toroidal current profile $I(\psi_{tn}) = 19\psi_{tn} - 13\psi_{tn}^2 + 3.8\psi_{tn}^3$ [kA], which raises the rotational transform profile to cross the $\iota = 1/2$ surface, as shown in Fig. 10a. This change affects the fine level of quasi-symmetry described in the reference, but the errors are of order 10^{-3} . We also add small uniform plasma pressure of 1 Pa for NIMSTELL's semi-implicit operator. The shape of the configuration is shown through the plot of the magnitude of $\mathbf{B}_s$ at the plasma surface in Fig. 10b. Based on $B_s \sim 1$ T, $\rho_s = 6.35 \times 10^{-7}$ kg/m³, and $R = 1.2$ m, $\tau_A = 0.78$ μ s. With $\eta = 1.94 \times 10^{-8}$ Ω m and aspect ratio of six, the Lundquist number ($S \equiv \tau_r/\tau_A$) is 3.3×10^6 . We again set $\text{Pm} = 0.001$, $\nu_{\parallel}/\nu_{\text{iso}} = 1000$, and $(\Gamma - 1)\kappa_{\text{iso}}n/\nu_{\text{iso}} = D_n/\nu_{\text{iso}} = 1$. The $H1$ magnetic-field representation is used with a 34×48 polar mesh, packed at the $\iota = 1/2$ surface, surrounding an inner region of 144 elements. All fields are expanded in biquartic elements, toroidal resolution is $0 \leq n \leq 21$, and the numerical-interchange stabilization terms are not used for this case without a pressure gradient. The computation is run linearly to establish the tearing mode, and nonlinear terms are activated at 50 ms. As in the W7-A computation, Fourier components that are not part of the linear mode family are driven rapidly after the nonlinear terms are included (Fig. 11a). The width of the saturated magnetic island (Fig. 11b) is significant with respect to the minor radius of the stellarator.²

² We remind readers that the initial state is a modification of the optimized result of Ref. [69], where the ι -profile does not cross $1/2$.




Figure 10. Optimized and modified rotational transform profiles (a) and $|\mathbf{B}_s|$ at the plasma surface (b) for the quasi-axisymmetric NIMSTELL tearing computation.




Figure 11. Nonlinear results on the quasi-axisymmetric tearing computation showing (a) evolution of magnetic fluctuation energies and (b) magnetic topology via Poincaré surface of section for $t = 76$ ms.

4.3 Anisotropic thermal conduction

To demonstrate numerical convergence for conditions of anisotropic transport, we consider anisotropic thermal conduction with fixed magnetic field without flow. We first have NIMSTELL compute a helically symmetric current-free magnetic field as a magnetostatics problem in a straight periodic geometry, and then use it to solve a time-dependent heat transport computation until steady conditions are reached. Although an analytical solution with anisotropic thermal conduction is not known for this configuration, the computed central temperature provides a measure of the artificial perpendicular heat transport at a given resolution when using a large $\kappa_{\parallel}/\kappa_{\perp}$ -ratio [34].

The analytical form of the magnetic configuration is defined by the magnetic potential Φ_m with $\mathbf{B} = \nabla\Phi_m$. The single-mode solution to $\nabla^2\Phi_m = 0$, plus a secular axial component, is

$$\Phi_m(r, \theta, z) = CI_l \left(\frac{N_p r}{R} \right) \cos \left(l\theta - \frac{N_p z}{R} \right) + z,$$

where $I_l(x)$ is the l -th order modified Bessel function of the first kind, N_p is the axial-period number for an equivalent major radius of R , and C is a constant that sets the magnitude of the field relative to the normalized secular component. We use a configuration with $l = 2$, $R/N_p = 4$, and C is set so that the amplitude of B_r is 0.1 at $r = 1$. The closed magnetic-flux region is elongated in the r - θ plane with a separatrix having x -points at $r = 3.46$. The magnetostatics problem determines $\mathbf{A}$ and χ from a modified version of the implemented Eq. (12) such that $\nabla \times \nabla \times \mathbf{A} + C_d \mathbf{A} + \nabla \chi = 0$ with $\nabla^2 \chi = \nabla \cdot \mathbf{A}$ in the interior of the domain. The parameter C_d is for numerical solvability and is set to 10^{-6} . Boundary conditions impose $\hat{\mathbf{n}} \times (\nabla \times \mathbf{A} - \nabla \Phi_m) = 0$ and $\chi = 0$ with $\hat{\mathbf{n}}$ being the surface normal. The heat transport problem solves Eqs. (3) and (4) with $\mathbf{V} = 0$, $Q = 4$, $\kappa_\perp = 1$, $\kappa_\parallel = 10^6$, and $T = 0$ along the boundary. Rapid parallel transport maintains the open-field temperature at its boundary value. Numerically, the polynomial degree of λ_i for the $\mathbf{H}(\text{curl}) \mathbf{A}$ and χ is one larger than that used for T .

We consider three sets of computations. One has a cylindrical domain of radius 4, and the other two have elliptical domains with semi-major axis of 4 and ellipticity of 0.96 that are aligned with the helical closed-flux region. Of the two shaped computations, one has its plane of elements twist with the closed-flux region. The geometry and magnetic potential in this set are fully represented by the $n = 0$ and $n = 1$ components of the Fourier expansion for the axial coordinate. An entire 2π rotation of the elliptical shape is required, so the configuration has $N_p = 2$ and $R = 8$. With this domain, the heat transport problem is 2D, and the temperature is fully computed in NIMSTELL's $n = 0$ component. The other set of shaped computations maintains the same azimuthal orientation of the elements within the poloidal plane but distorts their poloidal positions to follow the elliptical shaping. These computations represent a single axial period with $N_p = 1$ and $R = 4$, and the geometry, magnetic potential, and temperature have broad axial spectra. The two types of shaped meshing are illustrated in Fig. 12. None of the three sets of computations has a mesh that is precisely aligned with the magnetic flux surfaces, but the elliptically shaped meshes have better alignment than the circular mesh. The center of each mesh transitions to a distorted rectangle, as used for all MHD computations reported above.

Computations with the twisting elliptical mesh have two Fourier components, *i.e.* $N = 1$, and converge with increasing resolution of the r - θ plane. The results in Fig. 13a show an h -convergence series with bicubic elements for T with the number of elements in the radial and azimuthal directions varied simultaneously. It also shows a p -convergence series with a 28×32 annular mesh surrounding the central region of 64 elements and varied polynomial degree. The mesh is not tailored to match the shape of the separatrix, and the h -series produces an empirical convergence rate of 2, due to the boundary layer at the separatrix. Nonetheless, the results are accurate even with relatively coarse resolution. The results shown in Fig. 13b for the cylindrical and elliptically shaped non-twisting configurations have biquintic elements for T and a 31×40 annular mesh surrounding a central region of 100 elements. Resolution is only varied for the Fourier representation over the axial coordinate. Overall, the results are far less accurate than those obtained with the twisting mesh because the Fourier expansion must reproduce the axial variation of $\mathbf{A}$ and T . Elliptical shaping improves accuracy at low resolution, relative to using the cylindrical mesh, but it is not as effective as twisting the mesh to reduce the dimensionality of the problem.




Figure 12. Plots of meshing and computed temperature for the anisotropic heat transport problem. The twisting type of mesh is shown in (a), and the distorting mesh is shown in (b). The plotted mesh lines do not show the curvature of element sides. The axial divisions are for plotting purposes only, and half of the $N_p = 2$ domain is plotted in (a). The major axis of the elliptical mesh is aligned with the orientation of the x -points.




Figure 13. Results on central temperature for the heat transport problem, as computed with the (a) twisting elliptical mesh and (b) without twist for the elliptically shaped and circular meshes. The two curves in (a) show convergence with respect to polynomial degree (“p”) for 960 elements and with respect to mesh resolution (“h”) for bicubic basis functions for T . For the h -series, the horizontal axis shows the base-2 logarithm of the number of elements in the azimuthal direction. Poloidal resolution is held fixed in (b) as resolution over the axial direction is varied.

5. Discussion and conclusions

When using the $H1$ magnetic representation, the NIMSTELL numerical methods follow those used in NIMROD. Nonetheless, the modifications from the NIMROD implementation are extensive. The physical differential operations in the poloidal plane and the Jacobian of the mapping are independent of the Fourier projections in NIMROD. Derivatives are fully formed in the Fourier

representation, and nonlinear products are formed separately through fast Fourier transforms (FFTs) to a mesh over the toroidal angle. In contrast, forming differential operations in NIMSTELL involves partial derivatives with respect to the element coordinates ξ and η and the generalized toroidal angle ζ in the Fourier representation that are transformed to the ζ -mesh, where differentiation is completed through products with the reciprocal basis vectors ∇u_i . The reciprocal basis vectors and Jacobian are formed on the ζ -mesh because it is the coordinates $\mathbf{R}(\xi, \eta, \zeta)$ that have Fourier expansions. The operations on the ζ -mesh account for the reciprocal bases needed for the test functions, in addition to those for the basis functions, before applying FFTs back to the Fourier representation.

Another important extension for practical application of NIMSTELL is having preconditioning operations include multiple Fourier components according to the symmetry of a configuration. The coupling of Fourier components from the shaping of modern stellarators is too significant to rely on iteration without Fourier coupling in the preconditioning, as is typically used for nonlinear NIMROD computations. Nonetheless, covering complete mode families may be infeasible in large production problems. For NIMROD and NIMSTELL, we typically use the SuperLU_DIST sparse direct solver [65] for our preconditioning operations. This has proven effective and practical for NIMROD and smaller NIMSTELL applications. However, preconditioning matrices can be too large to factor, despite our recent memory-related optimization, such as using blocked symbolic operations. An example is the vector-potential computation with $N = 31$ in Fig. 7b. Whereas the memory of three of the National Energy Research Scientific Computing Center’s Perlmutter CPU nodes (512 GB for two AMD EPYC 7763 CPUs [71]) accommodates the preconditioning matrices covering the three mode families for all fields in the magnetic-component computation with $N = 31$, that is not the case for the $N = 31$ vector-potential computation. For that computation, we split each mode family into two bands with overlap of 2 Fourier components, which increases the iteration count for the $\mathbf{V}$ -advance from two to hundreds, even with a loosened tolerance, and the computational time per step is more than 50 times larger than the $H1$ computation. For large problems, it may be more effective to include all Fourier components in mode families using incomplete factorization, which remains to be explored.

The option to use the $\mathbf{H}(\text{curl})$ vector representation is distinct for NIMSTELL. The aim of avoiding magnetic divergence error altogether prompted the analysis shown in Appendix A. Its results on the tearing-mode tests of Section 4.2 agree quantitatively with those of the $H1$ magnetic representation linearly and nonlinearly. Because the $\mathbf{H}(\text{curl})$ implementation uses polynomial basis functions that are of one degree larger than those for the other fields and simultaneously solves a scalar for the gauge condition, there is an associated computational cost. For the $N = 21$ linear computations of Fig. 7b, the computational time per step of the $\mathbf{H}(\text{curl})$ computation is 80% larger than that of the $H1$ computation without stabilization terms and 50% larger than the one with stabilization terms. It is possible that the damping method for the gauge can be applied with a time-split advance, thereby reducing the algebraic system size for $\Delta \mathbf{A}$. Another computational point is that our semi-implicit method uses $\nabla \mathbf{B}$ from the evolving fields in nonlinear computations, so our nonlinear vector-potential computations perform an extra $H1$ projection of magnetic components at each step. The greater concern, raised in Section 4.1, for the NIMSTELL vector-potential representation is the minimum level of resistivity needed for a given mesh when plasma pressure is important. A finite-element stabilization for the vector-potential advance [72] may help, though application to magnetic-confinement should avoid artificial magnetic reconnection, which we also require of our divergence cleaning with magnetic representations. Investigation of vector-potential stabilization methods for NIMSTELL is left for future work.

Use of the generalized toroidal angle in NIMSTELL is largely unexplored at this time, meaning that we have taken $\zeta = \phi$ in almost all computations, including those presented here. We foresee optimization of the mapping to achieve distinct but likely overlapping benefits of improving numerical convergence and minimizing Fourier coupling in the algebraic systems. Depending on the application, optimization may involve a combination of rotating the plane of elements, as considered in our thermal-conduction tests of Section 4.3, and modifying how ϕ is related to ζ .

A more foundational aspect of preprocessing for stellarator applications is that NIMSTELL and other stellarator MHD codes rely on external equilibrium solutions. For NIMSTELL, that is a separate computational step that is performed with DESC [48], and equilibria previously computed with VMEC [73] can be handled through DESC’s ability to read VMEC output and re-solve the force balance. For NIMSTELL computations that retain the interpolated external equilibrium in the steady-state fields, it is important for those fields to satisfy force-balance accurately. That accuracy depends on both the numerical quality of the external equilibrium computations and on the resolution of the NIMSTELL representation.

The computations presented in Section 4 show that the $H1$ magnetic representation with diffusive correction of magnetic divergence error is suitable for stellarator applications. Although the Fourier representation does not allow direct application of $\mathbf{B} \cdot \hat{\mathbf{n}} = 0$ boundary conditions through essential conditions on the $H1$ space, implicit damping of $\mathbf{B} \cdot \hat{\mathbf{n}}$ through surface integrals is effective and does not adversely affect the conditioning of the algebraic system for $\Delta \mathbf{B}$. The semi-implicit operator computed from 3D fields stabilizes the temporally staggered leapfrog method in strongly shaped configurations and allows separate advances of each physical field. Also, the numerical-interchange stabilization method from our earlier NIMROD work yields convergence from the stable side in the presence of bad curvature. Although only confirmed here with cylindrical tests, results on pressure-driven MHD modes in stellarator configurations will be presented in a forthcoming benchmarking article. Including multiple Fourier components in our preconditioning operations, according to the symmetry of a given application, has proven to be important for reducing Krylov-space iteration to a practical level. As for the vector-potential option, it requires further numerical development to be used reliably in all applications. It is possible that solutions to the apparent coercivity problem at small resistivity values already exist and will be explored in future work, along with further solver and mapping improvements.

6. Acknowledgments

The first author is pleased to acknowledge discussions with former students Dr. Brian S. Cornille and Mr. Colin M. Guilbault that contributed to early aspects of the NIMSTELL effort. Funding for this work has been provided by U.S. Dept. of Energy awards DE-SC0018642 and DE-FG02-99ER54546. Award DE-SC0024548 for the HiFiStell Project also funds this work, and prior to April 1, 2025, Princeton University provided subcontract SUB0000763 for the University of Wisconsin-Madison effort.

This research uses resources of the National Energy Research Scientific Computing Center (NERSC), a Department of Energy User Facility, NERSC award FES-ERCAP 32189. Our research also uses computational resources at the University of Wisconsin-Madison’s Center for High Throughput Computing (2006), doi:10.21231/GNT1-HW21.

7. Appendix A. Analysis of MHD representation with vector potential

The performance of different numerical formulations for inhomogeneous plasma profiles cannot be fully described through analytics. Thus, we use a numerical eigenvalue solver with a flexible numerical representation over the radial coordinate while assuming $\exp(im\theta + ikz)$ over the azimuthal and axial directions for a chosen (m, k) component. This code had previously been used to develop the numerical interchange stabilization method described in Section 3.3. The code allows rapid investigation of different formulations for meaningful test cases. Here, we use it to investigate the $\mathbf{H}(\text{curl})$ vector-potential representation with various formulations for setting the gauge condition and to compare the influence of changing the degree of the polynomial basis functions for potentials relative to those for the flow velocity and the scalar pressure. Because the element expansion is only over one coordinate and a time-advance algorithm is not represented, the eigenvalue computations provide relevant but incomplete information for NIMSTELL. The algebraic eigenvalue problems resulting with the different numerical formulations are solved with appropriate routines from the LAPACK library [74].

Using vector potential in the linear ideal-MHD system without equilibrium flow, and assuming time-dependence of the form $\exp(-i\omega t)$, leads to the eigenvalue problem,

$$-i\omega\rho_0\mathbf{V} = \nabla \cdot \left[\frac{1}{\mu_0}\mathbf{B}_0(\nabla \times \mathbf{A}) + \frac{1}{\mu_0}(\nabla \times \mathbf{A})\mathbf{B}_0 - \mathbf{I}\left(\frac{1}{\mu_0}\mathbf{B}_0 \cdot \nabla \times \mathbf{A} + p\right) \right], \quad (\text{A1})$$

$$-i\omega p = -\mathbf{V} \cdot \nabla p_0 - \Gamma p_0 \nabla \cdot \mathbf{V}, \text{ and} \quad (\text{A2})$$

$$-i\omega\mathbf{A} = \mathbf{V} \times \mathbf{B}_0 - \nabla\chi \quad (\text{A3})$$

with boundary conditions $rA_\theta = A_z = V_r = 0$ for an ideal conductor at $r = a$. The equilibrium fields are indicated by the “0” subscripts. The first-order form of the system represents how the equations are solved in time-dependent computations, and ideal MHD is the numerically challenging limit of vanishing dissipation. An alternative system uses the total pressure p_T instead of the sum of separate plasma and magnetic pressures in (A1), and (A2) is replaced by the equation for the total pressure,

$$-i\omega p_T = \mathbf{B}_0 \cdot \nabla \mathbf{V} \cdot \mathbf{B}_0 - \left(\frac{B_0^2}{\mu_0} + \Gamma p_0\right) \nabla \cdot \mathbf{V} - \mathbf{V} \cdot \nabla \left(\frac{B_0^2}{2\mu_0} + p_0\right). \quad (\text{A4})$$

When using $\mathbf{H}(\text{curl})$ representations for $\mathbf{A}$, the polynomial expansion for A_r is discontinuous in r and of one degree smaller than those for rA_θ and A_z , which are continuous. We also consider computations that use an $H1$ representation for all components of $\mathbf{A}$.

Three different gauge formulations have been investigated with the eigenvalue computations. The first is the Weyl gauge with $\chi = 0$ and $\mathbf{H}(\text{curl})$ representations for $\mathbf{A}$. This formulation would reduce the computational cost of time-dependent 3D simulations by avoiding the extra field χ . The second formulation is described by $\nabla^2\chi = d_A\nabla \cdot \mathbf{A}$. It leads to the Coulomb gauge by damping $\nabla \cdot \mathbf{A}$. This property is shown by taking the divergence of the ideal part of (A3) and replacing $\nabla^2\chi$ with $d_A\nabla \cdot \mathbf{A}$ to produce $-i\omega\nabla \cdot \mathbf{A} = \nabla \cdot (\mathbf{V} \times \mathbf{B}_0) - d_A\nabla \cdot \mathbf{A}$. For $\mathbf{A}$ in the $\mathbf{H}(\text{curl})$ space, this formulation requires integrating the weak form of the χ -equation by parts. A third formulation also leads to the Coulomb gauge, but it diffuses $\nabla \cdot \mathbf{A}$ by setting $\chi = -D_A\nabla \cdot \mathbf{A}$. Repeating the substitutions produces $-i\omega\nabla \cdot \mathbf{A} = \nabla \cdot (\mathbf{V} \times \mathbf{B}) + \nabla^2 D_A\nabla \cdot \mathbf{A}$. Here, the additional condition $\oint_{\partial\Omega} d\mathbf{S} \cdot \mathbf{A} = 0$ is needed to remove numerical monopoles in $\mathbf{A}$. This formulation may be used with $\mathbf{A}$ in $\mathbf{H}(\text{curl})$, provided that the $\nabla \cdot \mathbf{A}$ projection is integrated by parts. Alternatively, if $\mathbf{A}$ is in $H1$ and the $\nabla \cdot \mathbf{A}$ projection is not integrated by parts but the $\nabla\chi$ projection is, then χ may be of the L^2 space. This would be convenient for eliminating the coefficients of χ in terms of those for $\mathbf{A}$ when solving the algebraic system. Taking the limit of large d_A or D_A enforces the Coulomb gauge, strictly, in the damping and diffusive formulations.

Basic numerical properties of these three formulations are evident from eigenvalue results for equilibria with uniform $\mathbf{B}_0$ in the axial direction and with uniform p_0 and ρ_0 . The analytical spectrum of this system includes a degenerate torsional Alfvén spectrum, the sound wave as the accumulation point of slow modes, and the fast Sturmian modes [13]. We choose a normalized configuration with B_{0z} , ρ_0 , Γ , cylinder radius a , and μ_0 set to unity and $p_0 = 1/2$. For $m = 1$, $k = 2$, for example, the Alfvén spectrum is at $|\omega| = 2$, the first two fast modes are at $|\omega| = 3.088$ and 6.875 , the sound wave is at $|\omega| = \sqrt{4/3}$, and the largest-frequency slow mode is at $|\omega| = 1.245$. There are no modes below the sound accumulation point, between the highest-frequency slow wave and the Alfvén spectrum, and between the Alfvén spectrum and the first fast mode. A comparison of computed $|Re\{\omega\}|$ for the three gauge formulations with either plasma or total pressure is shown in Fig. 14a. Three elements are used over r for each computation, and p , V_r , V_θ , and V_z are cubic within each element and continuous across element interfaces for their $H1$ space. For the Weyl and damping formulations, $\mathbf{A}$ is in $\mathbf{H}(\text{curl})$ with cubic A_r and quartic rA_θ and A_z , and the potential χ for the damping formulation is quartic and in $H1$. For the dissipative formulation, we show results where A_r , A_θ , and A_z are in $H1$ and quartic, χ is in $H1$ and cubic, and the $\nabla \cdot \mathbf{A}$ projection is integrated by parts.

For all three gauge formulations, results computed with total pressure instead of plasma pressure reproduce the analytical frequencies, accurately. Computations with plasma pressure produce two numerical modes between the Alfvén spectrum and the first fast mode. The computed eigenfunctions (not shown) indicate that these numerical modes are noisy in r and have relatively little coupling to $\mathbf{A}$. For the Weyl and damping formulations with $\mathbf{A}$ in $\mathbf{H}(\text{curl})$, there are 16 numerical modes having $Re\{\omega\} = 0$. There is no damping of these non-oscillatory modes in the Weyl formulation. With the damping formulation to control $\nabla \cdot \mathbf{A}$, 11 of these modes have $Im\{\omega\} = -100$, which is the applied d_A -value. The number 11 corresponds with the degrees of freedom in the quartic $H1$ expansion of χ after considering its boundary condition and the regularity condition at $r = 0$. The remaining 5 non-oscillatory modes with the damping formulation have eigenfunctions where the amplitudes of p and $\mathbf{V}$ are near roundoff. These eigenfunctions are solenoidal perturbations of $\mathbf{A}$ at the mesh scale that do not couple to the fluid and are not damped in the absence of resistive dissipation.

The diffusive formulation with $D_A = 100$, χ in $H1$, and the projection of $\nabla \cdot \mathbf{A}$ integrated by parts is also accurate for the uniform equilibrium, regardless of whether $\mathbf{A}$ is in $H1$ (shown) or in $\mathbf{H}(\text{curl})$ (not shown). With $\mathbf{A}$ in $H1$, there is one more non-oscillatory mode. In contrast, computations with $\mathbf{A}$ in $H1$, χ in L^2 , and projections of $\nabla \cdot \mathbf{A}$ not integrated by parts fail to produce an Alfvén spectrum. Moreover, the imaginary parts indicate growing modes, and a numerical MHD formulation that produces growing modes for a spatially uniform equilibrium is not acceptable. This last formulation may require the space of $\mathbf{A}$ and χ functions to satisfy the divergence-stability condition [75].

Using expansions of $\mathbf{A}$ with the continuous λ_k polynomials of the same degree as those for p and $\mathbf{V}$ would reduce computational costs in time-dependent applications. Thus, we also consider results with cubic expansions for all fields except A_r , which is quadratic. Figure 2b shows that regardless of whether plasma or total pressure is used, the $\mathbf{H}(\text{curl})$ computations with damping have oscillatory modes below the sound wave and between the slow and Alfvén modes. There are fewer non-oscillatory modes, but one of these modes has $Im\{\omega\} > 0$ when the system includes plasma pressure. The computation with $\mathbf{A}$ in $H1$, the diffusive gauge formulation, and total pressure has an accurate real spectrum, apart from the non-oscillatory modes.




Figure 14. Two comparisons of $|Re\{\omega\}|$ from the eigenvalue computations with different gauge and pressure formulations for the uniform equilibrium and $m = 1$, $k = 2$. Results for $\mathbf{A}$ of one polynomial degree greater than p and $\mathbf{V}$ are shown in (a), and results for rA_θ and A_z of the same degree as p and $\mathbf{V}$ are shown in (b). The column labeled “Ana” shows reference analytical values. The computational results are for the Weyl formulation (“W”), damped formulation with $\mathbf{A}$ in $\mathbf{H}(\text{curl})$ (“d”), dissipative formulation with $\mathbf{A}$ and χ in H^1 (“D”), and dissipative formulation with χ in L^2 (“L2”). Labels with “pt” indicate the use of total pressure from Eq. (A4). The dash-dot lines show the slow-mode accumulation point, and the dashed lines show the Alfvén spectrum.

Other information that can be obtained from the ideal-MHD eigenvalue computations is whether a formulation converges on local interchange from growth rates that are too small (“stable side”) or too large (“unstable side”). The preference in early ideal-MHD eigenvalue computations was to converge from the unstable side to avoid missing growing modes [2]. However, Lütjens and Luciani note that time-dependent nonlinear computations need to avoid the growth of unresolved or analytically stable modes [62]. This prompted the development of the specific stabilization scheme for NIMROD, described in Section 3.3, to make it converge on resonant interchange from the stable side.

We again consider the cylindrical profiles described in Section 4.1 and $m = 4$ modes with $k = -1.5$ and $k = -1.784$, which are resonant at $r = 0.507$ and $r = 0.371$ with D_s parameters of 0.200 and 0.443, respectively. Results on the fastest-growing mode from computations with the Weyl and damping formulations for the analytically unstable $k = -1.784$ wavenumber are shown in Fig. 15a. These computations have $\mathbf{A}$ in $\mathbf{H}(\text{curl})$ with quartic rA_θ and A_z and cubic H^1 expansions for p and $\mathbf{V}$. The elements have uniform spacing over r to demonstrate convergence without prior knowledge of localized behavior for a particular mode. Apart from the 20-element computations, where the damping formulation produces a weakly growing mode and the Weyl computation does not, the growth rates from the two sets of computations agree to within 0.001%. Growth rates from the 160- and 320-element computations agree to within 0.23%. Part of the eigenfunction for the unstable mode from the 80-element computation with damping is plotted in Fig. 15b to show the extent of radial localization. The phase of the mode with respect to $\exp(im\theta + ikz)$ is arbitrary. The observed convergence from the stable side when using $\mathbf{A}$ of greater polynomial degree than p and $\mathbf{V}$ is consistent with the analysis in Ref. [62] for their lower-order representations of displacement and magnetic flux with reduced-MHD.




Figure 15. Results for the $m = 4, k = -1.784$ unstable interchange mode that is resonant at $r = 0.371$. Frame (a) shows numerical convergence on growth rates for the Weyl and damping formulations with $d_A = 1$ for the latter. Frame (b) shows eigenfunction components of the 80-element damping computation.

Figure 16 shows results for the analytically stable $k = -1.5$ wavenumber from four sets of computations. The Weyl formulation does not produce a growing mode in the computations with 160 or fewer elements. The damping formulation with $\mathbf{A}$ in $\mathbf{H}(\text{curl})$ and $d_A = 1$ produces weakly growing modes at all resolution levels, and the 320-element growth rate is slightly faster than the one from the Weyl formulation. Other formulations do not perform as well. Computations with the diffusive formulation, $D_A = 1$, and $\mathbf{A}$ in $\mathbf{H}(\text{curl})$ produce modes that grow faster than the damping formulation. When the $H1$ expansion is used for $\mathbf{A}$, the computations produce overstable modes that grow very rapidly. Other formulations also perform poorly. Computations with damping and $\mathbf{A}$ in $\mathbf{H}(\text{curl})$ but of the same polynomial degree as p and $\mathbf{V}$ produce a non-oscillatory mode with $Im\{\omega\} = 0.14$. Computations with total pressure instead of plasma pressure also find growing modes with $Im\{\omega\}$ of order 0.1. Computations with $\mathbf{A}$ in $H1$, χ in L^2 , and the $\nabla \cdot \mathbf{A}$ projection not integrated by parts produce overstable modes with $Im\{\omega\}$ of order 1. These results show that none of the formulations reproduce a strictly stable spectrum when the conditions are close to marginal stability, but the Weyl and damping formulations with $\mathbf{A}$ in $\mathbf{H}(\text{curl})$ and λ_k of one degree greater than the bases for p and $\mathbf{V}$ predict the weakest numerical growth. The equilibrium profile has $D_s \rightarrow \infty$ for $r \rightarrow 0$, which makes these computations challenging for all formulations. We have also found that adding the interchange stabilization terms to eigenmode computations with vector potential does not improve the numerical behavior.




Figure 16. Results for the $m = 4$, $k = -1.5$ stable interchange mode that is resonant at $r = 0.443$. The symbols and traces show results with the Weyl, damping, and diffusive formulations with $\mathbf{A}$ in the spaces indicated in the figure legend.

Because NIMSTELL is intended to allow p -refinement, we also consider sets of 40-element computations that scan the polynomial degree of the basis functions. They have $\mathbf{A}$ in $\mathbf{H}(\text{curl})$ with the continuous λ_k of one degree larger than the $H1$ expansions for p and $\mathbf{V}$, and the number of Gaussian integration points in each element is increased with the polynomial degree to ensure accurate projections. The computations use the damping formulation with $d_A = 1$. Although p -refinement is not recommended for modeling localized behavior, Fig. 17 shows that accurate results can be obtained for the unstable $k = -1.784$ wavenumber, and the computations avoid growing solutions for the analytically stable $k = -1.5$ wavenumber to a greater extent than with h -refinement. For comparison, the unstable-mode computation with quintic polynomials has fewer degrees of freedom than the 80-element computation with cubic polynomials, and the quintic computation is slightly more accurate.

For the reasons discussed in Section 4.1, the NIMSTELL code allows the user to choose either the new $\mathbf{H}(\text{curl})$ expansion for the vector potential or the $H1$ expansion of magnetic-field components. The latter is used in NIMROD, and for NIMSTELL, it is implemented with the 3D coordinate mapping. Eigenmode analysis for the $H1$ $\mathbf{B}$ representation is reported in Ref. [63]. Findings from that research are relevant to the use of $H1$ $\mathbf{B}$ in NIMSTELL computations.




Figure 17. Convergence results with p -refinement for 40-element computations of the stable and unstable wavenumbers considered previously. The horizontal axis shows the degree of the expansions used for A_r , p , and V , and the $H(\text{curl})$ elements have rA_θ and A_z of one degree larger.

Appendix B. Numerical continuity of $\nabla \times \mathbf{A}$ across element interfaces

The $H1$ continuity of the mapping and the shared coefficients for the GLL expansions imply that the representation of $\mathbf{A}$ in Eq. (11) has continuous tangential components and that its curl has continuous normal components [52]. Given the non-standard NIMSTELL combination of 2D $H(\text{curl})$ elements with a 1D Fourier expansion, we prove these properties in this appendix.

Each element has its own mapping, *i.e.*, its own transformation between element coordinates and physical coordinates. Nonetheless, the continuity across elements implies that the tangent basis vectors that run along the sides of a 2D element, $\partial \mathbf{R} / \partial \eta$ and $\partial \mathbf{R} / \partial \zeta$ at $\xi = 0, 1$ for example, are continuous. The continuity of the tangential components of $\mathbf{A}$ then follows from $\nabla u_i \cdot \partial \mathbf{R} / \partial u_j = \delta_{ij}$ and the shared coefficients of the expansion. For example,

$$\mathbf{A}_e \cdot \frac{\partial \mathbf{R}}{\partial \eta} \Big|_{\xi=1} = \sum_{k', l'} A_{e, \eta_{k', l'}}(\zeta) \lambda_{k'}(1) \varsigma_{l'}(\eta), \quad (\text{B1})$$

where the dot product selects coefficients that are shared with the adjacent element, regardless of whether the two sets of element coordinates have the same orientation in physical space.

For the continuity of the normal component of $\nabla \times \mathbf{A}$ along element interfaces, again consider the representation at $\xi = 1$. Since the Jacobian is

$$J = \frac{\partial \mathbf{R}}{\partial \xi} \cdot \frac{\partial \mathbf{R}}{\partial \eta} \times \frac{\partial \mathbf{R}}{\partial \zeta} = \frac{1}{\nabla \xi \cdot \nabla \eta \times \nabla \zeta}, \quad (\text{B2})$$

and $\nabla \xi$ is parallel to $\frac{\partial \mathbf{R}}{\partial \eta} \times \frac{\partial \mathbf{R}}{\partial \zeta}$, the unit normal along this surface is

$$\hat{\mathbf{n}}|_{\xi=1} = \frac{1}{|\nabla \xi|} \nabla \xi \Big|_{\xi=1} = \frac{J}{\left| \frac{\partial \mathbf{R}}{\partial \eta} \times \frac{\partial \mathbf{R}}{\partial \zeta} \right|} \nabla \xi \Big|_{\xi=1}. \quad (\text{B3})$$

The normal component of $\nabla \times \mathbf{A}_e$ has contributions from the second and third sums in Eq. (11):

$$\nabla \times \mathbf{A}_e \cdot \hat{\mathbf{n}}|_{\xi=1} = \left[\sum_{k'',l''} A_{e,\zeta_{k'',l''}}(\zeta) \lambda_{k''}(1) \dot{\lambda}_{l''}(\eta) - \sum_{k',l'} \dot{A}_{e,\eta_{k',l'}}(\zeta) \lambda_{k'}(1) \zeta_{l'}(\eta) \right] \frac{1}{\left| \frac{\partial \mathbf{R}}{\partial \eta} \times \frac{\partial \mathbf{R}}{\partial \zeta} \right|} \Big|_{\xi=1}, \quad (\text{B4})$$

where the dot above a function represents its derivative. The only differentiation in (A4) is along the 1D basis functions that lie within the interface, which are the same for adjacent elements. Also, the coefficients are those for an endpoint of the GLL $\lambda_k(\xi)$ polynomials, which are shared for adjacent elements. Thus, apart from a minus sign for outward normal vectors relative to each element, the above expression matches that of the adjacent element. The continuity properties for tangential components of $\mathbf{A}$ and for the normal component of $\nabla \times \mathbf{A}$ hold regardless of the degree of each polynomial expansion in Eq. (11), and the two terms in (B4) are of the same degree in η when the GLL $\lambda_{l''}(\eta)$ is of one degree larger than the GL $\zeta_{l'}(\eta)$.

References

- [1] A. H. Glasser, “The direct criterion of Newcomb for the ideal MHD stability of an axisymmetric toroidal plasma,” *Phys. Plasmas*, vol. 23, no. 7, p. 072505, July 2016, doi: 10.1063/1.4958328.
- [2] R. Gruber, *Finite element methods in linear ideal magnetohydrodynamics*. Berlin ; New York: Springer-Verlag, [1985].
- [3] E. Strumberger and S. Günter, “CASTOR3D: linear stability studies for 2D and 3D tokamak equilibria,” *Nucl. Fusion*, vol. 57, no. 1, p. 016032, Nov. 2016, doi: 10.1088/0029-5515/57/1/016032.
- [4] A. Bondeson, G. Vlad, and H. Lütjens, “Resistive toroidal stability of internal kink modes in circular and shaped tokamaks,” *Phys. Fluids B*, vol. 4, no. 7, pp. 1889–1900, July 1992, doi: 10.1063/1.860041.
- [5] H. Lütjens, J. F. Luciani, and X. Garbet, “Nonlinear three-dimensional MHD simulations of tearing modes in tokamak plasmas,” *Plasma Phys. Control. Fusion*, vol. 43, no. 12A, p. A339, Nov. 2001, doi: 10.1088/0741-3335/43/12A/326.
- [6] K. Ichiguchi, S. Sakakibara, S. Ohdachi, and B. A. Carreras, “Numerical magnetohydrodynamic analysis of Large Helical Device plasmas with magnetic axis swing,” *Plasma Phys. Control. Fusion*, vol. 55, no. 1, p. 014009, Dec. 2012, doi: 10.1088/0741-3335/55/1/014009.
- [7] D. D. Schnack, E. J. Caramana, and R. A. Nebel, “Three-dimensional magnetohydrodynamic studies of the reversed-field pinch,” *Phys. Fluids*, vol. 28, no. 1, pp. 321–333, Jan. 1985, doi: 10.1063/1.865151.
- [8] C. R. Sovinec, J. M. Finn, and D. del-Castillo-Negrete, “Formation and sustainment of electrostatically driven spheromaks in the resistive magnetohydrodynamic model,” *Phys. Plasmas*, vol. 8, no. 2, pp. 475–490, Feb. 2001, doi: 10.1063/1.1335585.
- [9] B. Carreras, H. R. Hicks, J. A. Holmes, and B. V. Waddell, “Nonlinear coupling of tearing modes with self-consistent resistivity evolution in tokamaks,” *Phys. Fluids*, vol. 23, no. 9, pp. 1811–1826, Sept. 1980, doi: 10.1063/1.863206.
- [10] S. E. Kruger, D. D. Schnack, and C. R. Sovinec, “Dynamics of the major disruption of a DIII-D plasma,” *Phys. Plasmas*, vol. 12, no. 5, p. 056113, Apr. 2005, doi: 10.1063/1.1873872.

- [11] D. Pfefferlé, N. Ferraro, S. C. Jardin, I. Krebs, and A. Bhattacharjee, “Modelling of NSTX hot vertical displacement events using M3D-C1,” *Phys. Plasmas*, vol. 25, no. 5, p. 056106, Apr. 2018, doi: 10.1063/1.5016348.
- [12] H. R. Strauss, B. C. Lyons, and M. Knolker, “Resistive wall tearing mode disruptions in DIII-D and ITER tokamaks,” *Phys. Plasmas*, vol. 29, no. 11, p. 112508, Nov. 2022, doi: 10.1063/5.0107048.
- [13] J. P. Freidberg, *Ideal magnetohydrodynamics*. New York: Plenum Press, [1987] ©1987, 1987. Accessed: Sept. 23, 2025. [Online]. Available: <https://search.library.wisc.edu/catalog/999583347102121>
- [14] O. Motojima *et al.*, “Extended steady-state and high-beta regimes of net-current free heliotron plasmas in the Large Helical Device,” *Nucl. Fusion*, vol. 47, no. 10, p. S668, Sept. 2007, doi: 10.1088/0029-5515/47/10/S15.
- [15] C. Mercier, “A necessary condition for hydromagnetic stability of plasma with axial symmetry,” *Nucl. Fusion*, vol. 1, no. 1, p. 47, Sept. 1960, doi: 10.1088/0029-5515/1/1/004.
- [16] K. Y. Watanabe *et al.*, “Effects of global MHD instability on operational high beta-regime in LHD,” *Nucl. Fusion*, vol. 45, no. 11, p. 1247, Oct. 2005, doi: 10.1088/0029-5515/45/11/004.
- [17] A. Weller *et al.*, “Significance of MHD Effects in Stellarator Confinement,” *Fusion Science and Technology*, vol. 50, no. 2, pp. 158–170, Aug. 2006, doi: 10.13182/FST06-A1231.
- [18] Y. Suzuki, K. Y. Watanabe, and S. Sakakibara, “Theoretical studies of equilibrium beta limit in LHD plasmas,” *Phys. Plasmas*, vol. 27, no. 10, p. 102502, Oct. 2020, doi: 10.1063/5.0015106.
- [19] A. H. Reiman and H. S. Greenside, “Numerical solution of three-dimensional magnetic differential equations,” *Journal of Computational Physics*, vol. 75, no. 2, pp. 423–443, Apr. 1988, doi: 10.1016/0021-9991(88)90121-0.
- [20] A. Weller *et al.*, “Survey of magnetohydrodynamic instabilities in the advanced stellarator Wendelstein 7-AS,” *Phys. Plasmas*, vol. 8, no. 3, pp. 931–956, Mar. 2001, doi: 10.1063/1.1346633.
- [21] H. R. Strauss, L. E. Sugiyama, G. Y. Fu, W. Park, and J. Breslau, “Simulation of two fluid and energetic particle effects in stellarators,” *Nucl. Fusion*, vol. 44, no. 9, p. 1008, Aug. 2004, doi: 10.1088/0029-5515/44/9/010.
- [22] S. P. Hirshman and J. C. Whitson, “Steepest-descent moment method for three-dimensional magnetohydrodynamic equilibria,” *Phys. Fluids*, vol. 26, no. 12, pp. 3553–3568, Dec. 1983, doi: 10.1063/1.864116.
- [23] N. Mizuguchi, Y. Suzuki, and N. Ohyaibu, “Nonlinear dynamics of a collapse phenomenon in heliotron plasma with large pressure gradient,” *Nucl. Fusion*, vol. 49, no. 9, p. 095023, Sept. 2009, doi: 10.1088/0029-5515/49/9/095023.
- [24] N. M. Ferraro, S. C. Jardin, L. L. Lao, M. S. Shephard, and F. Zhang, “Multi-region approach to free-boundary three-dimensional tokamak equilibria and resistive wall instabilities,” *Phys. Plasmas*, vol. 23, no. 5, p. 056114, May 2016, doi: 10.1063/1.4948722.

- [25] Y. Zhou, N. M. Ferraro, S. C. Jardin, and H. R. Strauss, "Approach to nonlinear magnetohydrodynamic simulations in stellarator geometry," *Nucl. Fusion*, vol. 61, no. 8, p. 086015, July 2021, doi: 10.1088/1741-4326/ac0b35.
- [26] Y. Zhou, K. Aleynikova, C. Liu, and N. M. Ferraro, "Benign Saturation of Ideal Ballooning Instability in a High-Performance Stellarator," *Phys. Rev. Lett.*, vol. 133, no. 13, p. 135102, Sept. 2024, doi: 10.1103/PhysRevLett.133.135102.
- [27] Y. Zhou, K. Aleynikova, and N. M. Ferraro, "Nonlinear magnetohydrodynamic modeling of current-drive-induced sawtooth-like crashes in the W7-X stellarator," *Phys. Plasmas*, vol. 30, no. 3, p. 032503, Mar. 2023, doi: 10.1063/5.0136654.
- [28] A. M. Wright and N. M. Ferraro, "Investigating nonlinear magnetohydrodynamics in an optimized, reactor-scale quasi-axisymmetric stellarator," *Phys. Plasmas*, vol. 31, no. 8, p. 082509, Aug. 2024, doi: 10.1063/5.0215594.
- [29] G. T. A. Huysmans, "External kink (peeling) modes in x-point geometry," *Plasma Phys. Control. Fusion*, vol. 47, no. 12, p. 2107, Oct. 2005, doi: 10.1088/0741-3335/47/12/003.
- [30] M. Hoelzl *et al.*, "The JOEKE non-linear extended MHD code and applications to large-scale instabilities and their control in magnetically confined fusion plasmas," *Nucl. Fusion*, vol. 61, no. 6, p. 065001, May 2021, doi: 10.1088/1741-4326/abf99f.
- [31] N. Nikulsin, M. Hoelzl, A. Zocco, K. Lackner, and S. Günter, "A three-dimensional reduced MHD model consistent with full MHD," *Phys. Plasmas*, vol. 26, no. 10, p. 102109, Oct. 2019, doi: 10.1063/1.5122013.
- [32] N. Nikulsin *et al.*, "JOEKE3D: An extension of the JOEKE nonlinear MHD code to stellarators," *Phys. Plasmas*, vol. 29, no. 6, p. 063901, June 2022, doi: 10.1063/5.0087104.
- [33] R. Ramasamy *et al.*, "Nonlinear MHD modeling of soft β limits in W7-AS," *Nucl. Fusion*, vol. 64, no. 8, p. 086030, June 2024, doi: 10.1088/1741-4326/ad56a1.
- [34] C. R. Sovinec *et al.*, "Nonlinear magnetohydrodynamics simulation using high-order finite elements," *Journal of Computational Physics*, vol. 195, no. 1, pp. 355–386, Mar. 2004, doi: 10.1016/j.jcp.2003.10.004.
- [35] M. G. Schlutt, C. C. Hegna, C. R. Sovinec, E. D. Held, and S. E. Kruger, "Self-consistent simulations of nonlinear magnetohydrodynamics and profile evolution in stellarator configurations," *Phys. Plasmas*, vol. 20, no. 5, p. 056104, Apr. 2013, doi: 10.1063/1.4802834.
- [36] M. G. Schlutt, C. C. Hegna, C. R. Sovinec, S. F. Knowlton, and J. D. Hebert, "Numerical simulation of current evolution in the Compact Toroidal Hybrid," *Nucl. Fusion*, vol. 52, no. 10, p. 103023, Oct. 2012, doi: 10.1088/0029-5515/52/10/103023.
- [37] T. Bechtel, "Stellarator Beta Limits with Extended MHD Modeling Using NIMROD," PhD, University of Wisconsin-Madison, 2021.
- [38] N. A. Roberds *et al.*, "Simulations of sawtoothing in a current carrying stellarator," *Phys. Plasmas*, vol. 23, no. 9, p. 092513, Sept. 2016, doi: 10.1063/1.4962990.
- [39] M. O. Deville, P. F. Fischer, and E. H. Mund, *High-order methods for incompressible fluid flow*. Cambridge: Cambridge University Press, 2002.
- [40] J. P. Boyd, *Chebyshev & Fourier spectral methods*. Berlin ; New York: Springer-Verlag, [1989].

- [41] S. C. Jardin, J. Breslau, and N. Ferraro, “A high-order implicit finite element method for integrating the two-fluid magnetohydrodynamic equations in two dimensions,” *Journal of Computational Physics*, vol. 226, no. 2, pp. 2146–2174, Oct. 2007, doi: 10.1016/j.jcp.2007.07.003.
- [42] S. J. P. Pamela, G. T. A. Huijsmans, and M. Hoelzl, “A generalised formulation of G-continuous Bezier elements applied to non-linear MHD simulations,” *Journal of Computational Physics*, vol. 464, p. 111101, Sept. 2022, doi: 10.1016/j.jcp.2022.111101.
- [43] N. M. Ferraro and S. C. Jardin, “Calculations of two-fluid magnetohydrodynamic axisymmetric steady-states,” *Journal of Computational Physics*, vol. 228, no. 20, pp. 7742–7770, Nov. 2009, doi: 10.1016/j.jcp.2009.07.015.
- [44] C. R. Sovinec and J. R. King, “Analysis of a mixed semi-implicit/implicit algorithm for low-frequency two-fluid plasma modeling,” *Journal of Computational Physics*, vol. 229, no. 16, pp. 5803–5819, Aug. 2010, doi: 10.1016/j.jcp.2010.04.022.
- [45] K. J. Bunkers and C. R. Sovinec, “The influence of boundary and edge-plasma modeling in computations of axisymmetric vertical displacement,” *Phys. Plasmas*, vol. 27, no. 11, p. 112505, Nov. 2020, doi: 10.1063/5.0023604.
- [46] C.-D. Munz, P. Omnes, R. Schneider, E. Sonnendrücker, and U. Voß, “Divergence Correction Techniques for Maxwell Solvers Based on a Hyperbolic Model,” *Journal of Computational Physics*, vol. 161, no. 2, pp. 484–511, July 2000, doi: 10.1006/jcph.2000.6507.
- [47] A. M. Popov, V. S. Chan, M. S. Chu, Y. Q. Liu, B. W. Rice, and A. D. Turnbull, “Nonlinear three-dimensional self-consistent simulations of negative central shear discharges in the DIII-D tokamak,” *Phys. Plasmas*, vol. 8, no. 8, pp. 3605–3619, Aug. 2001, doi: 10.1063/1.1380235.
- [48] D. W. Dudt and E. Kolemen, “DESC: A stellarator equilibrium solver,” *Phys. Plasmas*, vol. 27, no. 10, p. 102513, Oct. 2020, doi: 10.1063/5.0020743.
- [49] D. S. Harned and W. Kerner, “Semi-implicit method for three-dimensional compressible magnetohydrodynamic simulation,” *Journal of Computational Physics*, vol. 60, no. 1, pp. 62–75, Aug. 1985, doi: 10.1016/0021-9991(85)90017-8.
- [50] D. D. Schnack, D. C. Barnes, Z. Mikic, D. S. Harned, and E. J. Caramana, “Semi-implicit magnetohydrodynamic calculations,” *Journal of Computational Physics*, vol. 70, no. 2, pp. 330–354, June 1987, doi: 10.1016/0021-9991(87)90186-0.
- [51] K. Lerbinger and J. F. Luciani, “A new semi-implicit method for MHD computations,” *Journal of Computational Physics*, vol. 97, no. 2, pp. 444–459, Dec. 1991, doi: 10.1016/0021-9991(91)90008-9.
- [52] J. C. Nedelec, “Mixed finite elements in $\mathbb{R}^3$,” *Numer. Math.*, vol. 35, pp. 315–341, 1980, doi: <https://doi.org/10.1007/BF01396415>.
- [53] S. Zaglmayr, “High Order Finite Element Methods for Electromagnetic Field Computation,” PhD, Johannes Kepler University Linz, 2006.
- [54] G. Strang and G. J. Fix, *An analysis of the finite element method*. Wesley, MA: Wesley-Cambridge Press, 1997.
- [55] L. P. Franca, G. Hauke, and A. Masud, “Revisiting stabilized finite element methods for the advective–diffusive equation,” *Computer Methods in Applied Mechanics and*

- Engineering*, vol. 195, no. 13, pp. 1560–1572, Feb. 2006, doi: 10.1016/j.cma.2005.05.028.
- [56] A. N. Brooks and T. J. R. Hughes, “Streamline upwind/Petrov-Galerkin formulations for convection dominated flows with particular emphasis on the incompressible Navier-Stokes equations,” *Computer Methods in Applied Mechanics and Engineering*, vol. 32, no. 1, pp. 199–259, Sept. 1982, doi: 10.1016/0045-7825(82)90071-8.
 - [57] T. Barth, P. Bochev, M. Gunzburger, and J. Shadid, “A Taxonomy of Consistently Stabilized Finite Element Methods for the Stokes Problem,” *SIAM Journal on Scientific Computing*, vol. 25, no. 5, pp. 1585–1607, 2004, doi: 10.1137/S1064827502407718.
 - [58] A. I. Nesliturk and M. Tezer-Sezgin, “The finite element method for MHD flow at high Hartmann numbers,” *Computer Methods in Applied Mechanics and Engineering*, vol. 194, no. 9, pp. 1201–1224, Mar. 2005, doi: 10.1016/j.cma.2004.06.035.
 - [59] D. Sondak and A. A. Oberai, “Large eddy simulation models for incompressible magnetohydrodynamics derived from the variational multiscale formulation,” *Phys. Plasmas*, vol. 19, no. 10, p. 102308, Oct. 2012, doi: 10.1063/1.4759157.
 - [60] D. Sondak, J. N. Shadid, A. A. Oberai, R. P. Pawlowski, E. C. Cyr, and T. M. Smith, “A new class of finite element variational multiscale turbulence models for incompressible magnetohydrodynamics,” *Journal of Computational Physics*, vol. 295, pp. 596–616, Aug. 2015, doi: 10.1016/j.jcp.2015.04.035.
 - [61] J. N. Shadid, R. P. Pawlowski, E. C. Cyr, R. S. Tuminaro, L. Chacón, and P. D. Weber, “Scalable implicit incompressible resistive MHD with stabilized FE and fully-coupled Newton–Krylov-AMG,” *Computer Methods in Applied Mechanics and Engineering*, vol. 304, pp. 1–25, June 2016, doi: 10.1016/j.cma.2016.01.019.
 - [62] H. Lütjens and J. F. Luciani, “A class of basis functions for non-ideal magnetohydrodynamic computations,” *Computer Physics Communications*, vol. 95, no. 1, pp. 47–57, May 1996, doi: 10.1016/0010-4655(96)00006-9.
 - [63] C. R. Sovinec, “Stabilization of numerical interchange in spectral-element magnetohydrodynamics,” *Journal of Computational Physics*, vol. 319, pp. 61–78, Aug. 2016, doi: 10.1016/j.jcp.2016.04.063.
 - [64] Y. Saad, *Iterative methods for sparse linear systems*, 2nd ed. Philadelphia, PA: Society for Industrial and Applied Mathematics, 2003.
 - [65] X. S. Li and J. W. Demmel, “SuperLU_DIST: A scalable distributed-memory sparse direct solver for unsymmetric linear systems,” *ACM Trans. Math. Softw.*, vol. 29, no. 2, pp. 110–140, June 2003, doi: 10.1145/779359.779361.
 - [66] F. Ebrahimi, S. C. Prager, and C. R. Sovinec, “Resistive–ideal transition of pressure-driven instabilities in current-carrying plasmas beyond the Suydam criterion,” *Phys. Plasmas*, vol. 9, no. 6, pp. 2470–2473, June 2002, doi: 10.1063/1.1481860.
 - [67] B. Coppi, J. M. Greene, and J. L. Johnson, “Resistive instabilities in a diffuse linear pinch,” *Nucl. Fusion*, vol. 6, no. 2, p. 101, June 1966, doi: 10.1088/0029-5515/6/2/003.
 - [68] G. Grieger, H. Renner, and H. Wobig, “Wendelstein stellarators,” *Nuclear Fusion*, vol. 25, no. 9, p. 1231, Sept. 1985, doi: 10.1088/0029-5515/25/9/040.
 - [69] M. Landreman and E. Paul, “Magnetic Fields with Precise Quasisymmetry for Plasma Confinement,” *Phys. Rev. Lett.*, vol. 128, no. 3, p. 035001, Jan. 2022, doi: 10.1103/PhysRevLett.128.035001.

- [70] *DESC Github repository*. [Online]. Available:
<https://github.com/PlasmaControl/DESC/tree/master>
- [71] *NERSC Perlmutter system*. [Online]. Available:
<https://docs.nersc.gov/systems/perlmutter/architecture/>
- [72] J. N. Shadid, R. P. Pawlowski, J. W. Banks, L. Chacón, P. T. Lin, and R. S. Tuminaro, “Towards a scalable fully-implicit fully-coupled resistive MHD formulation with stabilized FE methods,” *Journal of Computational Physics*, vol. 229, no. 20, pp. 7649–7671, Oct. 2010, doi: 10.1016/j.jcp.2010.06.018.
- [73] S. P. Hirshman, U. Schwenn, and J. Nührenberg, “Improved radial differencing for three-dimensional magnetohydrodynamic equilibrium calculations,” *Journal of Computational Physics*, vol. 87, no. 2, pp. 396–407, Apr. 1990, doi: 10.1016/0021-9991(90)90259-4.
- [74] E. Anderson et al., *LAPACK Users’ Guide*. in Software, Environments, and Tools. Society for Industrial and Applied Mathematics, 1999. doi: 10.1137/1.9780898719604.
- [75] M. D. Gunzburger, *Mathematical aspects of finite element methods for incompressible viscous flows*. Hampton, VA : Institute for Computer Applications in Science and Engineering, NASA Langley Research Center ; [Springfield, Va. :] : [National Technical Information Service, distributor], [1986], 1986.