SIAM J. SCI. COMPUT.
Vol. 38, No. 6, pp. B1009–B1031

BLOCK PRECONDITIONERS FOR STABLE MIXED NODAL AND EDGE FINITE ELEMENT REPRESENTATIONS OF INCOMPRESSIBLE RESISTIVE MHD*

EDWARD G. PHILLIPS†, JOHN N. SHADID†, ERIC C. CYR†, HOWARD C. ELMAN‡, AND ROGER P. PAWLOWSKI†

Abstract. The scalable iterative solution of strongly coupled three-dimensional incompressible resistive magnetohydrodynamics (MHD) equations is very challenging because disparate time scales arise from the electromagnetics, the hydrodynamics, as well as the coupling between these systems. This study considers a mixed finite element discretization of a dual saddle point formulation of the incompressible resistive MHD equations using a stable nodal (Q2/Q1) discretization for the hydrodynamics and a stable edge-node discretization of a reduced form of the Maxwell equations. This paper presents new approximate block factorization preconditioners for this system which reduce the system to approximate Schur complement systems that can be solved using algebraic multilevel methods. These preconditioners include a new augmentation-based approximation for the magnetic induction saddle point system as well as efficient approximations of the Schur complements that arise from the complex coupling between the Navier–Stokes equations and the Maxwell equations.

Key words. magnetohydrodynamics, preconditioners, mixed finite elements

AMS subject classifications. 76W05, 65F08, 65M22

DOI. 10.1137/16M1074084

1. Introduction.

The magnetohydrodynamics (MHD) model describes the flow of electrically conducting fluids in the presence of magnetic fields. From a continuum approximation perspective, the model is governed by a system of non-self-adjoint, nonlinear partial differential equations obtained from coupling the Navier–Stokes equations and a reduced form of the Maxwell equations. Because both hydrodynamic and electromagnetic effects play a strong role in this system, MHD can span a large range of length and time scales. This complex physical behavior makes the MHD equations difficult to solve and necessitates the development of robust, accurate numerical methods to approximate their solution. For this reason, much effort has been invested into the development of scalable preconditioners for the linear systems arising in fully implicit discretizations of the MHD equations [3, 4, 5, 6, 21, 22, 26, 30, 31, 33].

This study focuses on preconditioning the linear systems arising from a particular finite element formulation of the primitive variable incompressible MHD equations [29]. This formulation poses the equations as a coupling of two subsystems of saddle point type. The electromagnetic saddle point system is posed in terms of the magnetic induction, which is defined on edge elements, and a Lagrange multiplier, defined on

*Submitted to the journal's Computational Methods in Science and Engineering section May 9, 2016; accepted for publication (in revised form) August 22, 2016; published electronically November 17, 2016.

http://www.siam.org/journals/sisc/38-6/M107408.html

Funding: This work was partially funded by the U.S. Department of Energy under grant DE-SC0009301 and by the Department of Energy Office of Science ASCR Applied Math program at Sandia National Laboratory. Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

†Center for Computing Research, Sandia National Laboratories, Albuquerque, NM 87185 (egphill@sandia.gov, jnshadi@sandia.gov, eccyr@sandia.gov, rppawlo@sandia.gov).

‡Department of Computer Science and Institute for Advanced Computer Studies, University of Maryland, College Park, MD 20742 (elman@cs.umd.edu).

nodes. The hydrodynamic saddle point system is composed of the incompressible Navier–Stokes equations with the coupling between the fluid velocity and the fluid pressure. These two subsystems are coupled by the Lorentz force and a magnetic convection term.

The approach of this study is to develop block preconditioners that exploit the inherent block structure of the fully coupled linear system generated by both Newton and fixed point linearizations of the the nonlinear discrete systems. These types of preconditioners develop scalable algebraic multigrid methods (AMG) for challenging systems by employing AMG-based preconditioners that can effectively handle both nodal ($H(grad)$-conforming) and edge-element ($H(curl)$-conforming) finite element approximations. To this end, we draw on block preconditioning strategies that have been developed for the discretized Maxwell equations in saddle point form and the discretized Navier–Stokes equations while accounting for the coupling between the two subsystems. By applying block factorizations to the system matrix, the effects of coupling can be embedded and localized in algebraic Schur complements, thus moving the action of off-diagonal operators onto the block diagonal. The effectiveness of a preconditioner is then determined by how well such Schur complements are approximated. By analyzing the fully coupled system from a block perspective, the strength of coupling, and time scales of physical mechanisms can be employed in the design of the approximate block factorization. In this process, results from our previous work on developing preconditioners for a mixed finite element discretization of an exact-penalty formulation for MHD [26] are generalized to the current mixed incompressible fluid integration (Q2/Q1) and edge-element **B**-field and nodal (Q1) Lagrange multiplier finite element formulation [29] to develop effective Schur complement approximations. An alternative approach is to use specialized monolithic multigrid methods (see [1]), as broadly applicable nodal monolithic multigrid methods (e.g., [30, 31]) do not extend directly to mixed discretizations.

The remainder of this paper is organized as follows. In section 2, the problem is outlined with a presentation of the resistive incompressible MHD system as well as a discussion of the finite element discretization that is employed. This section also presents the basic form of the proposed preconditioner which involves an approximation of the electromagnetics and fluid subsystems. Sections 3 and 4 focus on the two saddle point subblock approximations. The performance of the preconditioners is then demonstrated on a variety of test problems in section 5. Section 6 closes with a few observations and conclusions.

2. The discrete problem and approach to preconditioning. The primitive variable nondimensional resistive MHD system is given by

$$\frac{\partial \vec{B}}{\partial t} + \frac{1}{Re_m}\nabla \times \nabla \times \vec{B} - \nabla \times (\vec{u} \times \vec{B}) + \nabla r = \vec{0}, \tag{1a}$$

$$\nabla \cdot \vec{B} = 0, \tag{1b}$$

$$\frac{\partial \vec{u}}{\partial t} - \frac{1}{Re}\Delta \vec{u} + \vec{u} \cdot \nabla \vec{u} + \nabla p + S\vec{B} \times \nabla \times \vec{B} = \vec{0}, \tag{1c}$$

$$\nabla \cdot \vec{u} = 0 \tag{1d}$$

with appropriate boundary conditions [14]. The unknowns are the magnetic induction $\vec{B}$, a Lagrange multiplier r, the fluid velocity $\vec{u}$, and the fluid pressure p. The Lagrange multiplier, r, is a nonphysical degree of freedom included so that the solenoidal involution (1b) can be enforced explicitly as a constraint. The nondimensional parameters Re, Re_m, and S are the (fluid) Reynolds number, the magnetic Reynolds number, and the coupling coefficient. These nondimensional parameters are defined

in terms of physical constants and characteristic scales as

$$Re = \frac{\rho \bar{u} \bar{L}}{\nu}, \quad Re_m = \mu \sigma \bar{L} \bar{u}, \quad S = \frac{\bar{B}^2}{\mu \rho \bar{u}^2}, \tag{2}$$

where ρ, ν, μ, and σ are the density, viscosity, permeability, and conductivity, and $\bar{u}, \bar{L}$, and $\bar{B}$ are the characteristic velocity, length scale, and magnetic field strength [13].

Briefly, (1a)–(1b) form a reduced version of the low frequency Maxwell equations, where Faraday's law has been used with the standard definition of the generalized resistive Ohm's law to form an evolution equation for the magnetic induction [14]. Equations (1c)–(1d) compose the incompressible Navier–Stokes equations with the Lorentz force $-S\vec{B} \times \nabla \times \vec{B}$ as a source. The magnetic convection, $\nabla \times (\vec{u} \times \vec{B})$ and the Lorentz force act as a nonlinear coupling between the electromagnetic and hydrodynamic subsystems.

A semidiscrete (backward Euler) and linearized version of the continuous system (1) is presented in (3):

$$\tfrac{1}{\Delta t}\vec{B}' + \tfrac{1}{Re_m}\nabla \times \nabla \times \vec{B}' - \delta \nabla \times (\vec{a} \times \vec{B}') - \nabla \times (\vec{u}' \times \vec{b}) + \nabla r' = -R_{\vec{B}}, \tag{3a}$$

$$\nabla \cdot \vec{B}' = -R_r, \tag{3b}$$

$$\tfrac{1}{\Delta t}\vec{u}' - \tfrac{1}{Re}\Delta \vec{u}' + \vec{a} \cdot \nabla \vec{u}' + \delta \vec{u}' \cdot \nabla \vec{a} + \nabla p' + S\vec{b} \times \nabla \times \vec{B}' + \delta S \vec{B}' \times \nabla \times \vec{b} = -R_{\vec{u}}, \tag{3c}$$

$$\nabla \cdot \vec{u}' = -R_p. \tag{3d}$$

This system represents the linearization with either Newton's method ($\delta = 1$) or a Picard linearization ($\delta = 0$), and needs to be solved at each step of a nonlinear iteration at each time step. The current values of the unknowns $\vec{u}$ and $\vec{B}$ are denoted by $\vec{a}$ and $\vec{b}$ in the nonlinear iteration, R_* refers to the corresponding nonlinear residual for each unknown, and $\vec{u}'$, p', $\vec{B}'$, r' are the increments of the dependent variables solved for in the fixed point iteration in defect correction form. In the subsequent discussion the prime superscript is omitted for economy of notation. When a solution to the nonlinear time-discretized problem exists, the Picard iteration converges to that solution for any initial guess and Newton's method converges for initial guesses sufficiently close to the exact solution [29].

2.1. Form of the discrete system. Following the finite element discretization proposed in [29], we discretize $\vec{B}$ with curl-conforming edge elements and r with appropriate nodal elements for a stable discretization of the Maxwell subsystem. The hydrodynamic unknowns can be discretized with any element pair that is stable for the Navier–Stokes equations. Throughout this work, we use first order Nédélec elements of the first kind [24] for $\vec{B}$, bilinear (Q_1) elements for r, and Taylor-Hood (Q_2-Q_1) elements for $\vec{u}$ and p. This leads to a stable discretization of the MHD equations [29]. The convergence of this discretization has been demonstrated in the literature [28]. After discretizing in space with finite elements, each system of equations representing the unknowns of the finite element approximation to (3) is represented by a discrete linear system of the form

$$\mathcal{A}\mathbf{x} = R, \tag{4}$$

where $\mathbf{x} = (\mathbf{B}, \mathbf{r}, \mathbf{u}, \mathbf{p})$ is a vector containing the coefficients of $\vec{B}, r, \vec{u}$, and p in the finite element basis. The system matrix $\mathcal{A}$ can be written in block form as

$$\mathcal{A} = \begin{pmatrix} \mathcal{M}_{\mathbf{B}} & -\mathcal{Z}_{\mathbf{u}}^t \\ \mathcal{Z}_{\mathbf{B}} & \mathcal{M}_{\mathbf{u}} \end{pmatrix}, \tag{5}$$

where $\mathcal{M}_{\mathbf{B}}$ is the magnetic induction saddle point matrix, $\mathcal{M}_{\mathbf{u}}$ is the Navier–Stokes saddle point matrix, and $\mathcal{Z}_{\mathbf{u}}^t$ and $\mathcal{Z}_{\mathbf{B}}$ contain the coupling terms. Using the notation summarized in Table 1, each of these blocks is a block 2×2 matrix, i.e.,

$$\mathcal{M}_* = \begin{pmatrix} \frac{1}{\Delta t} Q_* + F_* & D^t \\ D & 0 \end{pmatrix}, \quad \mathcal{Z}_* = \begin{pmatrix} Z_* & 0 \\ 0 & 0 \end{pmatrix}. \tag{6}$$

Throughout the text, Q_*, F_*, will refer to the mass matrix and convection-diffusion stiffness matrix for unknown $*$. The operators $Z_{\mathbf{B}}$ and $Z_{\mathbf{u}}^t$ are discretizations of the Lorentz force and the off-diagonal component of the magnetic convection. The gradient and divergence operator are denoted as D^t and D for each unknown as well. We note that the discrete gradient and divergence are different matrices in the hydrodynamic and electromagnetic subsystems. We omit subscripts on these operators for economy of notation.

The system matrix $\mathcal{A}$ is generally large and sparse as well as nonsymmetric and indefinite. Hence, to solve the system (4) a preconditioned GMRES method [27] is employed. The challenge is then to develop a computationally efficient preconditioner for $\mathcal{A}$. Ideally, this preconditioner should lead to small linear iteration counts independent of the mesh size h and should be robust over changes in the parameters $\Delta t, Re, Re_m$, and S. To develop preconditioners for this system, the block structure of $\mathcal{A}$ as it appears in (5) is exploited. We take advantage of this structure both to algebraically localize the effects of coupling in Schur complement approximations and to provide a framework where we can leverage previous work on preconditioning the discretized Maxwell and Navier–Stokes equations.

We focus on block upper triangular preconditioners motivated by the block LU decomposition

$$\mathcal{A} = \begin{pmatrix} \mathcal{I} & 0 \\ \mathcal{Z}_{\mathbf{B}} \mathcal{M}_{\mathbf{B}}^{-1} & \mathcal{I} \end{pmatrix} \begin{pmatrix} \mathcal{M}_{\mathbf{B}} & \mathcal{Z}_{\mathbf{u}}^t \\ 0 & \mathcal{X}_{\mathbf{u}} \end{pmatrix}, \tag{7}$$

TABLE 1

Definitions of discrete operators as they correspond to continuous operators. The parameter δ is zero for a Picard linearization and one for Newton's method. (Note that the notation for discrete operators is the same for both linearizations; if the distinction is necessary, it should be clear from context.) Although the discrete gradient and divergence operators are different matrices in the electromagnetic and hydrodynamic subsystems, we use the same notation for both to simplify notation.

Discrete	Continuous	Approximate magnitude of operator
$F_{\mathbf{B}}\mathbf{B}$	$\frac{1}{Re_m}\nabla \times \nabla \times \vec{B} - \delta \nabla \times (\vec{a} \times \vec{B})$	$\frac{1}{Re_m h^2} + \delta \frac{\|\|\vec{a}\|\|}{h}$
$D^t\mathbf{r}$	∇r	$\frac{1}{h}$
$D\mathbf{B}$	$-\nabla \cdot \vec{B}$	$\frac{1}{h}$
$F_{\mathbf{u}}\mathbf{u}$	$\vec{a} \cdot \nabla \vec{u} + \delta \vec{u} \cdot \nabla \vec{a} - \frac{1}{Re}\Delta \vec{u}$	$\frac{\|\|\vec{a}\|\|}{h} + \delta \|\|\nabla \vec{a}\|\| \frac{1}{Reh^2}$
$D^t\mathbf{p}$	∇p	$\frac{1}{h}$
$D\mathbf{u}$	$\nabla \cdot \vec{u}$	$\frac{1}{h}$
$Z_{\mathbf{u}}^t\mathbf{u}$	$\nabla \times (\vec{u} \times \vec{b})$	$\frac{\|\|\vec{b}\|\|}{h}$
$Z_{\mathbf{B}}\mathbf{B}$	$S\vec{b} \times (\nabla \times \vec{B}) + \delta S \vec{B} \times (\nabla \times \vec{b})$	$\frac{S\|\|\vec{b}\|\|}{h} + \delta S \|\|\nabla \times \vec{b}\|\|$

where the hydrodynamic Schur complement $\mathcal{X}_{\mathbf{u}}$ is defined as

$$\mathcal{X}_{\mathbf{u}} := \mathcal{M}_{\mathbf{u}} - \mathcal{Z}_{\mathbf{B}}\mathcal{M}_{\mathbf{B}}^{-1}\mathcal{Z}_{\mathbf{u}}^{t}. \tag{8}$$

It is well known that if the exact upper triangular factor is used as a preconditioner for $\mathcal{A}$, then GMRES converges in exactly two iterations [23]. Hence, we consider preconditioners of the form

$$\mathcal{P} = \begin{pmatrix} \hat{\mathcal{M}}_{\mathbf{B}} & -\mathcal{Z}_{\mathbf{u}}^{t} \\ 0 & \hat{\mathcal{X}}_{\mathbf{u}} \end{pmatrix}, \tag{9}$$

where hats indicate approximations. Application of this preconditioner requires approximations of the actions of $\mathcal{M}_{\mathbf{B}}^{-1}$ and $\mathcal{X}_{\mathbf{u}}^{-1}$. Note that an alternative preconditioner can be obtained if the unknowns are reordered with the hydrodynamic degrees of freedom listed first. Then the Schur complement would appear on the electromagnetic degrees of freedom. We use the hydrodynamic Schur complement, because $\vec{B}$ is discretized with edge elements and we prefer to keep operators on this space as simple as possible. Empirically, we have found that component solves using multigrid on edge elements tend to be more difficult and sensitive to problem parameters than those on nodal elements.

The development of a practical preconditioner proceeds by constructing easily computable approximations of $\mathcal{M}_{\mathbf{B}}^{-1}$ and $\mathcal{X}_{\mathbf{u}}^{-1}$. First, an approximation for $\mathcal{M}_{\mathbf{B}}$ motivated by existing preconditioners for the discretized Maxwell equations is considered. Here, $\mathcal{X}_{\mathbf{u}}$ is approximated as a perturbation of the Navier–Stokes saddle point system. Then block preconditioning techniques developed for the original Navier–Stokes system are modified to propose an approximation of $\mathcal{X}_{\mathbf{u}}^{-1}$. The Maxwell system and the perturbed Navier–Stokes system are considered independently; that is, the choice of approximation for $\mathcal{M}_{\mathbf{B}}$ does not affect the definition of $\mathcal{X}_{\mathbf{u}}$. The approximations for the Picard linearization and Newton's method are considered simultaneously, differences in the approaches being noted as they arise.

3. Magnetic induction saddle point. In this section, the saddle point system generated by the electromagnetic equations with the reduced form of Maxwell's equations is considered in the form

$$\mathcal{M}_{\mathbf{B}} = \begin{pmatrix} \frac{1}{\Delta t}Q_{\mathbf{B}} + F_{\mathbf{B}} & D^{t} \\ D & 0 \end{pmatrix}. \tag{10}$$

The system is associated with a transient convection-diffusion operator on the magnetic induction, $\mathbf{B}$, and enforces the solenoidal involution as a constraint with the introduction of a Lagrange multiplier, $\mathbf{r}$. In general, the (0,0) block of $\mathcal{M}_{\mathbf{B}}$ is a difficult operator for component solvers such as multigrid to handle, and special care needs to be taken with this block to obtain a scalable preconditioner for the whole system. The difficulties with this operator stem from the operator $F_{\mathbf{B}}$ having a large null space. Discretizing the continuous relation

$$\nabla \cdot (\tfrac{1}{Re_m}\nabla \times \nabla \times -\delta \nabla \times (\vec{a} \times \cdot)) = 0 \tag{11}$$

yields

$$(Q_{\mathbf{r}}^{-1}D)(Q_{\mathbf{B}}^{-1}F_{\mathbf{B}}) \approx 0, \tag{12}$$

which implies that the null space of $F_{\mathbf{B}}$ is of dimension at least $\dim(\mathbf{r})$. If Δt is sufficiently large, then $\frac{1}{\Delta t}Q_{\mathbf{B}}+F_{\mathbf{B}}$ is dominated by $F_{\mathbf{B}}$, and the singularity of this operator can present significant difficulties for a component solver. In this case, certain high frequency components of the fine-grid error are not well represented in the residual of $\frac{1}{\Delta t}Q_{\mathbf{B}} + F_{\mathbf{B}}$, and this results in a degradation in the performance of traditional multigrid [19]. For this reason, preconditioners for $\mathcal{M}_{\mathbf{B}}$ should not only have a block structure that yields good iteration counts, but should also ameliorate the difficulties associated with the (0,0) subsolve.

3.1. Mass augmentation. Motivated by the preconditioners proposed in [15] (for a system equivalent to the Picard case with $Re_m \equiv 1$) and expanded upon in [25] (for the steady-state Newton case), we propose a block diagonal approximation of the form

$$\mathcal{M}_{\mathbf{B}} \approx \hat{\mathcal{M}}_{\mathbf{B},M} := \begin{pmatrix} (\frac{1}{\Delta t}+k)Q_{\mathbf{B}} + F_{\mathbf{B}} & 0 \\ 0 & \frac{1}{k}L_{\mathbf{r}} \end{pmatrix}, \tag{13}$$

where $k > 0$ is a real constant and $L_{\mathbf{r}}$ is a discrete Laplacian on the magnetic Lagrange multiplier space. Following the eigenvalue analyses of [15, 25], it can be shown that the preconditioned matrix $\mathcal{M}_{\mathbf{B}}\hat{\mathcal{M}}^{-1}_{\mathbf{B},M}$ has eigenvalues

$$\lambda = 1, -\frac{k\Delta t}{1+k\Delta t}, \tag{14}$$

each with multiplicity $\dim(\mathbf{r})$.

For a Picard linearization, the remaining $\dim(\mathbf{B}) - \dim(\mathbf{r})$ eigenvalues are real and can be bounded as

$$\frac{\frac{1}{\Delta t} + \frac{\alpha}{1-\alpha}Re_m}{\frac{1}{\Delta t} + \frac{\alpha}{1-\alpha}Re_m + k} \leq \lambda \leq 1, \tag{15}$$

where α is a coercivity constant independent of Re_m and the mesh size h. This further implies the bounds

$$\frac{\alpha}{\alpha + kRe_m(1-\alpha)} \leq \lambda \leq 1, \quad \frac{1}{1+k\Delta t} \leq \lambda \leq 1. \tag{16}$$

Hence, if we set either $k = \frac{1}{\Delta t}$ or $k = \frac{1}{Re_m}$, we obtain a parameter-independent bound. A good choice for k is then $k = \frac{1}{Re_m}$ so that the mass matrix scales with the magnetic diffusion term resulting from the discretization of $\frac{1}{Re_m}\nabla \times \nabla\times$. Then the performance of the component solver will not degrade if the diffusion term dominates the time-derivative term.

For Newton's method, we have not established a bound on the final $\dim(\mathbf{B}) - \dim(\mathbf{r})$ eigenvalues of $\mathcal{M}_{\mathbf{B}}\hat{\mathcal{M}}^{-1}_{\mathbf{B},M}$, but the empirical results of [25] indicate that the presence of the convection term in the (0,0) block does not significantly affect the performance of the preconditioner form. This effect was shown for the steady-state form of $\mathcal{M}_{\mathbf{B}}$ with exact component solves, even when the magnetic Reynolds number is large enough that convection dominates. Consequently, we suggest that the same approximation $\hat{\mathcal{M}}_{\mathbf{B},M}$ be used for both Picard and Newton linearizations.

Multigrid algorithms have been developed specifically for operators of the form $\frac{1}{\Delta t}Q_{\mathbf{B}} + F_{\mathbf{B}}$ when $F_{\mathbf{B}}$ arises from a Picard linearization (i.e., the eddy current formulation of the Maxwell equations) [2, 18, 19]. These algorithms are designed to

preserve the null space of the curl-curl operator (i.e., the space of discrete gradients) at each level of a coarse-grid hierarchy. Prolongators obtained from an auxiliary nodal problem are used to develop prolongators for the edge discretization. While these multigrid algorithms have been demonstrated to be effective and scalable for the Picard linearization, it is difficult to say whether they should perform well for Newton's method as the null space of $F_{\mathbf{B}}$ is not the same for both linearizations. In the absence of an alternative multigrid method for Newton's method, we apply this type of eddy current multigrid routine for both cases.

3.2. Grad-div augmentation. In this section, we present an approximation of $\mathcal{M}_{\mathbf{B}}$ based on augmenting the (0,0) block with a term of the form $kD^tQ_{\mathbf{r}}^{-1}D$. Similar preconditioners for a system equivalent to the steady-state Picard system were proposed in [34, 35]. This operator is analogous to the continuous operator $-k\nabla\nabla\cdot$, and if k is taken to be equal to $\frac{1}{Re_m}$, then this operator together with the magnetic diffusion operator $\frac{1}{Re_m}\nabla\times\nabla\times$ forms the vector Laplacian $-\frac{1}{Re_m}\Delta$. We leave the parameter $k \geq 0$ variable for generality. In practice, we replace $Q_{\mathbf{r}}$ by its diagonal so that the augmentation term $kD^tQ_{\mathbf{r}}^{-1}D$ is simple to compute. The resulting approximation is similar in form to the steady-state preconditioners presented in [34, 35] and can be shown to be equivalent as $\Delta t \to \infty$.

We motivate the approximation with the block decomposition

$$\mathcal{M}_{\mathbf{B}} = \begin{pmatrix} I & -kD^tQ_{\mathbf{r}}^{-1} \\ 0 & I \end{pmatrix} \begin{pmatrix} \frac{1}{\Delta t}Q_{\mathbf{B}} + F_{\mathbf{B}} + kD^tQ_{\mathbf{r}}^{-1}D & D^t \\ D & 0 \end{pmatrix} \tag{17}$$

$$= \begin{pmatrix} I & -kD^tQ_{\mathbf{r}}^{-1} \\ 0 & I \end{pmatrix} \begin{pmatrix} I & 0 \\ D(\frac{1}{\Delta t}Q_{\mathbf{B}} + F_{\mathbf{B}} + kD^tQ_{\mathbf{r}}^{-1}D)^{-1} & I \end{pmatrix} \times \begin{pmatrix} \frac{1}{\Delta t}Q_{\mathbf{B}} + F_{\mathbf{B}} + kD^tQ_{\mathbf{r}}^{-1}D & D^t \\ 0 & X_{\mathbf{r}} \end{pmatrix}, \tag{18}$$

where the Schur complement on the Lagrange multiplier $X_{\mathbf{r}}$ is defined as

$$X_{\mathbf{r}} := -D(\tfrac{1}{\Delta t}Q_{\mathbf{B}} + F_{\mathbf{B}} + kD^tQ_{\mathbf{r}}^{-1}D)^{-1}D^t. \tag{19}$$

In this decomposition, we first augment the (0,0) block of $\mathcal{M}_{\mathbf{B}}$ in (17) and then perform a block LU decomposition of the augmented saddle point matrix in (18). Observing that the system has been written as $\mathcal{M}_{\mathbf{B}} = \mathcal{U}_0\mathcal{L}\mathcal{U}_1$, we note that $\mathcal{M}_{\mathbf{B}}(\mathcal{U}_0\mathcal{U}_1)^{-1} = \mathcal{U}_0\mathcal{L}\mathcal{U}_0^{-1}$ is a similarity transformation of $\mathcal{L}$ and hence has the same minimum polynomial as $\mathcal{L}$. Thus, if $\mathcal{U}_0\mathcal{U}_1$ were used as a preconditioner for $\mathcal{M}_{\mathbf{B}}$, GMRES would converge in exactly two iterations. Multiplying out $\mathcal{U}_0\mathcal{U}_1$, we obtain the preconditioner

$$\mathcal{M}_{\mathbf{B},GD} := \begin{pmatrix} \frac{1}{\Delta t}Q_{\mathbf{B}} + F_{\mathbf{B}} + kD^tQ_{\mathbf{r}}^{-1}D & D^t(I - kQ_{\mathbf{r}}^{-1}X_{\mathbf{r}}) \\ 0 & X_{\mathbf{r}} \end{pmatrix}. \tag{20}$$

For this preconditioner to be useful, we need computable approximations of the actions of $X_{\mathbf{r}}$ and $X_{\mathbf{r}}^{-1}$. We develop such approximations through a strategy similar to the commutator-based approximations in [9]. As noted above, $DQ_{\mathbf{B}}^{-1}F_{\mathbf{B}} \approx 0$, which implies the following approximation:

$$DQ_{\mathbf{B}}^{-1}(\tfrac{1}{\Delta t}Q_{\mathbf{B}} + F_{\mathbf{B}} + kD^tQ_{\mathbf{r}}^{-1}D) \approx DQ_{\mathbf{B}}^{-1}(\tfrac{1}{\Delta t}Q_{\mathbf{B}} + kD^tQ_{\mathbf{r}}^{-1}D) \tag{21}$$

$$= (\tfrac{1}{\Delta t}Q_{\mathbf{r}} + kDQ_{\mathbf{B}}^{-1}D^t)Q_{\mathbf{r}}^{-1}D. \tag{22}$$

Some straightforward algebraic manipulation of this expression yields

$$X_{\mathbf{r}} \approx -Q_{\mathbf{r}}(\tfrac{1}{\Delta t}Q_{\mathbf{r}} + kDQ_{\mathbf{B}}^{-1}D^t)^{-1}(DQ_{\mathbf{B}}^{-1}D^t). \tag{23}$$

The operator $DQ_{\mathbf{B}}^{-1}D^t$ is approximately a discrete scalar Laplacian on the magnetic Lagrange multiplier space [10], and we approximate it by an explicit discretization of the scalar Laplacian on $\mathbf{r}$, $L_{\mathbf{r}}$. Hence, we have

$$X_{\mathbf{r}} \approx -Q_{\mathbf{r}}(\tfrac{1}{\Delta t}Q_{\mathbf{r}} + kL_{\mathbf{r}})^{-1}L_{\mathbf{r}}, \quad X_{\mathbf{r}}^{-1} \approx -\tfrac{1}{\Delta t}L_{\mathbf{r}}^{-1} - kQ_{\mathbf{r}}^{-1}. \tag{24}$$

Thus, we can approximate the actions of $X_{\mathbf{r}}$ and $X_{\mathbf{r}}^{-1}$ using only matrix-vector multiplications and solves with $Q_{\mathbf{r}}, L_{\mathbf{r}}$, and $\frac{1}{\Delta t}Q_{\mathbf{r}} + kL_{\mathbf{r}}$. Traditional multigrid has been demonstrated to work particularly well on each of these operators, and in practice a diagonal approximation suffices for $Q_{\mathbf{r}}^{-1}$ [10]. We thus define the approximation

$$\hat{\mathcal{M}}_{\mathbf{B},GD} := \begin{pmatrix} \frac{1}{\Delta t}Q_{\mathbf{B}} + F_{\mathbf{B}} + kD^tQ_{\mathbf{r}}^{-1}D & D^t(\frac{1}{\Delta t}Q_{\mathbf{r}} + kL_{\mathbf{r}})^{-1}(\frac{1}{\Delta t}Q_{\mathbf{r}} + 2kL_{\mathbf{r}}) \\ 0 & -Q_{\mathbf{r}}(\frac{1}{\Delta t}Q_{\mathbf{r}} + kL_{\mathbf{r}})^{-1}L_{\mathbf{r}} \end{pmatrix}. \tag{25}$$

At steady state, this approximation reduces to

$$\hat{\mathcal{M}}_{\mathbf{B},GD} = \begin{pmatrix} F_{\mathbf{B}} + kD^tQ_{\mathbf{r}}^{-1}D & 2D^t \\ 0 & -\frac{1}{k}Q_{\mathbf{r}} \end{pmatrix}, \tag{26}$$

and this is the preconditioner referred to as H_1 in [35], where it was proven that the steady version of $\mathcal{M}_{\mathbf{B}}$ preconditioned by this matrix has exactly one eigenvalue $\lambda = 1$. It can be shown that the same is true for the transient version. See Appendix A for details.

4. Perturbed Navier–Stokes saddle point. In this section, an approximation of the action of the inverse of $\mathcal{X}_{\mathbf{u}}$ as defined in (8) is developed. First, an analysis of the operator $\mathcal{Z}_{\mathbf{B}}\mathcal{M}_{\mathbf{B}}^{-1}\mathcal{Z}_{\mathbf{u}}^t$, viewed as a perturbation of the Navier–Stokes saddle point matrix, is presented. We then extend preconditioning strategies proposed for the discretized Navier–Stokes equations to develop an upper triangular approximation of $\mathcal{X}_{\mathbf{u}}^{-1}$. The analysis presented in this section is similar to that introduced in [26] in which an analogous perturbed Navier–Stokes system was approximated as part of an exact penalty formulation of the MHD equations. The current work generalizes that analysis from the two-dimensional steady-state setting to the three-dimensional time dependent problem.

4.1. Approximating the perturbation. Applying the definitions of $\mathcal{Z}_{\mathbf{B}}$ and $\mathcal{Z}_{\mathbf{u}}^t$ from (6), the perturbation operator can be written as

$$\mathcal{Z}_{\mathbf{B}}\mathcal{M}_{\mathbf{B}}^{-1}\mathcal{Z}_{\mathbf{u}}^t = \begin{pmatrix} Z_{\mathbf{B}}(\mathcal{M}_{\mathbf{B}}^{-1})_{0,0}Z_{\mathbf{u}}^t & 0 \\ 0 & 0 \end{pmatrix}. \tag{27}$$

The (0,0) block of $\mathcal{M}_{\mathbf{B}}^{-1}$ can be computed by explicitly inverting the decomposition (18) with any value of k. For simplicity consider $k = 1$; this then yields

$$(\mathcal{M}_{\mathbf{B}}^{-1})_{0,0} = (\tfrac{1}{\Delta t}Q_{\mathbf{B}} + F_{\mathbf{B}})^{-1}(\tfrac{1}{\Delta t}Q_{\mathbf{B}} + F_{\mathbf{B}} + D^tX_{\mathbf{r}}^{-1}D)(\tfrac{1}{\Delta t}Q_{\mathbf{B}} + F_{\mathbf{B}})^{-1}. \tag{28}$$

With this expression for $(\mathcal{M}_{\mathbf{B}}^{-1})_{0,0}$, we proceed by analyzing the continuous operator that corresponds to $Z_{\mathbf{B}}(\mathcal{M}_{\mathbf{B}}^{-1})_{0,0}Z_{\mathbf{u}}^t$. It can be shown that this continuous operator simplifies to

$$\kappa := [S\vec{b}\times(\nabla\times\cdot) + \delta S\cdot\times(\nabla\times\vec{b})][\tfrac{1}{\Delta t}I + \tfrac{1}{Re_m}\nabla\times\nabla\times - \delta\nabla\times(\vec{a}\times\cdot)]^{-1}\nabla\times(\cdot\times\vec{b}). \tag{29}$$

Details are shown in Appendix B.

Notice that κ can be discretized as $Z_{\mathbf{B}}(\frac{1}{\Delta t}Q_{\mathbf{B}} + F_{\mathbf{B}})^{-1}Z_{\mathbf{u}}^t$. As this expression still contains an embedded inverse, we seek approximations that can simplify this expression. We first assume that a discretization of the term $\delta S\cdot\times(\nabla\times\vec{b})$ is negligible compared to a discretization of $S\vec{b}\times(\nabla\times\cdot)$. This assumption is trivial if $\delta = 0$, and was argued for in the Newton case in [26]. This assumption essentially posits that $\vec{b}$ is a smooth function and thus does not have steep gradients. The assumption yields the approximation

$$\kappa \approx S\vec{b}\times\left\{\nabla\times[\tfrac{1}{\Delta t}I + \tfrac{1}{Re_m}\nabla\times\nabla\times - \delta\nabla\times(\vec{a}\times\cdot)]^{-1}\nabla\times\right\}(\cdot\times\vec{b}). \tag{30}$$

We focus on simplifying the expression between the braces,

$$\nabla\times[\tfrac{1}{\Delta t}I + \tfrac{1}{Re_m}\nabla\times\nabla\times - \delta\nabla\times(\vec{a}\times\cdot)]^{-1}\nabla\times \tag{31}$$

in the above expression, by considering the limiting cases when magnetic diffusion is negligible and when magnetic diffusion dominates.

Consider the case where magnetic diffusion ($\frac{1}{Re_m}\nabla\times\nabla\times$) is negligible compared to convection ($\frac{1}{\Delta t}I - \delta\nabla\times(\vec{a}\times\cdot)$). Based on the approximate magnitudes reported in Table 1, this corresponds to the condition

$$\tfrac{1}{Re_m h^2} \ll \tfrac{1}{\Delta t} + \tfrac{\delta||\vec{a}||}{h} \tag{32}$$

in the discrete setting. A discrete version of operator (31) has magnitude of the order

$$\tfrac{1}{h^2}\left[\tfrac{1}{\Delta t} + \tfrac{1}{Re_m h^2} + \tfrac{\delta||\vec{a}||}{h}\right]^{-1} \approx \tfrac{1}{h^2}\left[\tfrac{1}{\Delta t} + \tfrac{\delta||\vec{a}||}{h}\right]^{-1}. \tag{33}$$

Applying (32) to the factor $\frac{1}{h^2}$ in this expression, we obtain

$$\tfrac{1}{h^2}\left[\tfrac{1}{\Delta t} + \tfrac{\delta||\vec{a}||}{h}\right]^{-1} \ll Re_m\left[\tfrac{1}{\Delta t} + \tfrac{\delta||\vec{a}||}{h}\right]\left[\tfrac{1}{\Delta t} + \tfrac{\delta||\vec{a}||}{h}\right]^{-1} = Re_m. \tag{34}$$

Hence, the whole operator (31) is negligible compared to $Re_m I$ when magnetic diffusion is negligible.

When diffusion dominates, the operators $\frac{1}{\Delta t}$ and $\delta\nabla\times(\vec{a}\times\cdot)$ can be considered negligible compared to $\frac{1}{Re_m}\nabla\times\nabla\times$. It can be seen that the operator (31) should be approximately $Re_m I$ in this case, as something close to the curl-curl operator is inverted while two curls are applied. To see this more rigorously, observe that, through techniques similar to those in Appendix B, the expression in (31) can be rewritten as

$$Re_m I - Re_m[\tfrac{1}{\Delta t}I - \delta\vec{a}\times\nabla\times][\tfrac{1}{\Delta t}I + \tfrac{1}{Re_m}\nabla\times\nabla\times - \delta\vec{a}\times\nabla\times]^{-1}. \tag{35}$$

The second term here, when discretized, has magnitude of the order

$$Re_m\left[\tfrac{1}{\Delta t} + \tfrac{1}{Re_m h^2} + \tfrac{\delta||\vec{a}||}{h}\right]^{-1}\left[\tfrac{1}{\Delta t} + \tfrac{\delta||\vec{a}||}{h}\right]. \tag{36}$$

In the discrete setting, the assumption that magnetic diffusion dominates corresponds to the condition

$$\tfrac{1}{Re_m h^2} \gg \tfrac{1}{\Delta t} + \tfrac{\delta||\vec{a}||}{h}, \tag{37}$$

and applying this to (36), we obtain

$$Re_m\left[\tfrac{1}{\Delta t} + \tfrac{1}{Re_m h^2} + \tfrac{\delta||\vec{a}||}{h}\right]^{-1}\left[\tfrac{1}{\Delta t} + \tfrac{\delta||\vec{a}||}{h}\right] \ll Re_m. \tag{38}$$

Hence, the remainder in expression (35) is negligible compared to $Re_m I$, and (31) can be approximated by $Re_m I$ when diffusion dominates.

Based on the above observations, we make the approximation

$$\nabla \times [\tfrac{1}{\Delta t} I + \tfrac{1}{Re_m} \nabla \times \nabla \times - \delta \nabla \times (\vec{a} \times \cdot)]^{-1} \nabla \times \approx \gamma Re_m I, \tag{39}$$

where γ is a parameter between 0 and 1. We argue that γ should be approximately 1 when magnetic diffusion dominates, and γ should be approximately 0 when magnetic diffusion is negligible, since the operator should be much smaller in magnitude than $Re_m I$ in this case. A formula for γ that demonstrates this behavior in its limits is given by

$$\gamma := \frac{\Delta t}{\Delta t + Re_m h^2 + \delta Re_m h ||\vec{a}|| \Delta t}. \tag{40}$$

At steady state, this reduces to

$$\gamma = \frac{1}{1 + \delta Re_m h ||\vec{a}||}, \tag{41}$$

which is equivalent to an analogous parameter obtained in [26]. Using this approximation, κ can be approximated by

$$\kappa \approx \gamma S Re_m \vec{b} \times (\cdot \times \vec{b}). \tag{42}$$

Now we can approximate the discrete perturbation operator $Z_{\mathbf{B}}(\mathcal{M}_{\mathbf{B}}^{-1})_{0,0} Z_{\mathbf{u}}^t$ by $\gamma K_{\mathbf{u}}$, where $K_{\mathbf{u}}$ is a discretization of κ as defined in (42). The discrete operator can be written as a coupled scaling of the velocity mass matrix, i.e.,

$$K_{\mathbf{u}} := S Re_m Q_{\mathbf{u}} \begin{pmatrix} \operatorname{diag}(b_y^2 + b_z^2) & -\operatorname{diag}(b_x b_y) & -\operatorname{diag}(b_x b_z) \\ -\operatorname{diag}(b_x b_y) & \operatorname{diag}(b_x^2 + b_z^2) & -\operatorname{diag}(b_y b_z) \\ -\operatorname{diag}(b_x b_z) & -\operatorname{diag}(b_y b_z) & \operatorname{diag}(b_x^2 + b_y^2) \end{pmatrix}. \tag{43}$$

Then the perturbed Navier–Stokes saddle point matrix is approximated by

$$\mathcal{X}_{\mathbf{u}} \approx \begin{pmatrix} \hat{X}_{\mathbf{u}} & D^t \\ D & 0 \end{pmatrix}, \tag{44}$$

where $\hat{X}_{\mathbf{u}}$ is defined as

$$\hat{X}_{\mathbf{u}} := \tfrac{1}{\Delta t} Q_{\mathbf{u}} + F_{\mathbf{u}} + \gamma K_{\mathbf{u}}. \tag{45}$$

Observe that since the electromagnetic saddle point system is solved independently of the hydrodynamics, the fluid-magnetic coupling of the MHD equations is only expressed in the operator $\gamma K_{\mathbf{u}}$. Thus, the importance of coupling effects to the system matrix can be quantified by comparing the sizes of $\frac{1}{\Delta t} Q_{\mathbf{u}} + F_{\mathbf{u}}$ and $\gamma K_{\mathbf{u}}$; that is, if $||\frac{1}{\Delta t} Q_{\mathbf{u}} + F_{\mathbf{u}}|| \gg ||\gamma K_{\mathbf{u}}||$, then coupling effects are negligible. This condition holds when

$$Ha^2 h^2 ||\vec{b}||^2 \Delta t^2 \ll (\Delta t + Reh^2 + Reh||\vec{a}||\Delta t)(\Delta t + Re_m h^2 + \delta Re_m h ||\vec{a}|| \Delta t), \tag{46}$$

where $Ha = \sqrt{SReRe_m}$ is the Hartmann number, which is a nondimensional measure of the coupling in the problem.

By multiplying out the right-hand side, and considering each term independently, we obtain the following sufficient conditions for satisfaction of (46):

$$SRe_m||\vec{b}||^2\Delta t \ll 1, \tag{47}$$

$$Hah||\vec{b}|| \ll 1, \tag{48}$$

$$\sqrt{S}\tfrac{||\vec{b}||\Delta t}{h} \ll 1, \tag{49}$$

$$S||\vec{b}||^2 \ll \delta||\vec{a}||^2. \tag{50}$$

Each of these is a condition under which coupling is unimportant in a preconditioner. Condition (47) can be regarded as a CFL-like condition on the perturbation operator $SRe_m\vec{b}\times(\cdot\times\vec{b})$. It implies that the physics associated with the fluid-magnetic coupling are integrated below the explicit stability limit. Condition (48) essentially says that the mesh Hartmann number is small, meaning that the fluid-magnetic coupling is overresolved in space. Condition (49) is another CFL-like condition, this one for the Alfvén wave speed [14], which is $\sqrt{S}||\vec{b}||$ for this version of the MHD equations. Condition (50) is a discretization-independent statement that the velocity field dominates the magnetic field. This implies that coupling is essentially one directional, from hydrodynamics to electromagnetics. Practically, all of the conditions are fairly strict and limit the degree to which fluid and magnetic behavior interact. We therefore expect realistic implicit MHD simulations that are integrated at time scales of interest not to satisfy any of the restrictions (47)–(50). Hence, the inclusion of the term $\gamma K_{\mathbf{u}}$ is imperative for an effective preconditioner.

4.2. A commutator for the pressure Schur complement. The perturbed Navier–Stokes matrix admits the block LU decomposition

$$\mathcal{X}_{\mathbf{u}} \approx \begin{pmatrix} I & 0 \\ D\hat{X}_{\mathbf{u}}^{-1} & I \end{pmatrix}\begin{pmatrix} \hat{X}_{\mathbf{u}} & D^t \\ 0 & X_{\mathbf{p}} \end{pmatrix}, \tag{51}$$

where the pressure Schur complement is defined as

$$X_{\mathbf{p}} := -D\hat{X}_{\mathbf{u}}^{-1}D^t. \tag{52}$$

We can use an approximation of the upper triangular factor as an approximation of $\mathcal{X}_{\mathbf{u}}$. The operator $\hat{X}_{\mathbf{u}}$ is structurally similar to the traditional Navier–Stokes convection-diffusion operator since it differs only by a zero-order perturbation. Thus, we expect solvers that perform well on $F_{\mathbf{B}}$, such as AMG, to behave similarly on $\hat{X}_{\mathbf{u}}$.

What remains is to develop an efficiently computable approximation of $X_{\mathbf{p}}$. As $\hat{X}_{\mathbf{u}}$ is of the same structure as the perturbed convection-diffusion operator in [26], we use the pressure Schur complement approximation introduced there for $X_{\mathbf{p}}$. This approximation is an adaptation of the least-squares commutator approximation originally developed for the discretized Navier–Stokes equations (see [9, 10, 11] and the references therein). Following the Navier–Stokes literature, we assume that the convection-diffusion operator and the divergence operator approximately commute, i.e.,

$$\nabla\cdot(\tfrac{1}{\Delta t}I - \tfrac{1}{Re}\Delta + \vec{a}\cdot\nabla) \approx (\tfrac{1}{\Delta t}I_p - \tfrac{1}{Re}\Delta_p + (\vec{a}\cdot\nabla)_p)\nabla\cdot, \tag{53}$$

where a subscript p indicates that the operator is defined on the scalar pressure space. This relationship holds as long as the convection direction $\vec{a}$ is smooth. Discretizing this approximation yields

$$[Q_{\mathbf{p}}^{-1}D][Q_{\mathbf{u}}^{-1}(\tfrac{1}{\Delta t}Q_{\mathbf{u}} + F_{\mathbf{u}})] \approx [Q_{\mathbf{p}}^{-1}(\tfrac{1}{\Delta t}Q_{\mathbf{p}} + F_{\mathbf{p}})][Q_{\mathbf{p}}^{-1}D], \tag{54}$$

where $F_{\mathbf{p}}$ is a discretization of the pressure convection-diffusion operator. This approximation can be algebraically manipulated to yield a computable approximation of $DF_{\mathbf{u}}^{-1}D^t$ and the action of $(DF_{\mathbf{u}}^{-1}D^t)^{-1}$. To do the same for $D\hat{X}_{\mathbf{u}}^{-1}D^t$, we need a commutator for the perturbation operator $K_{\mathbf{u}}$. It can be shown that there exists a continuous operator κ_p such that

$$\nabla \cdot \kappa \approx \kappa_p \nabla \cdot \tag{55}$$

only when $SRe_m||\vec{b}||^2 \approx 0$, since κ couples the components of the divergence operator so that $\nabla \cdot \kappa$ cannot, in general, be written as a scalar perturbation of the divergence. Hence, a relationship of the form (55) is in general not true. To remedy this, we introduce an additional degree of freedom α as in [26], letting

$$\nabla \cdot \kappa \approx \alpha \kappa_p \nabla \cdot \tag{56}$$

The parameter α can be chosen appropriately to improve the commutator-based approximation of $X_{\mathbf{p}}$. Discretizing (56) yields

$$DQ_{\mathbf{u}}^{-1}K_{\mathbf{u}} \approx \alpha\gamma K_{\mathbf{p}}Q_{\mathbf{p}}^{-1}D. \tag{57}$$

Then, using (54) and (57), we obtain the approximation for $X_{\mathbf{p}}$

$$X_{\mathbf{p}} \approx Q_{\mathbf{p}}(\tfrac{1}{\Delta t}Q_{\mathbf{p}} + F_{\mathbf{p}} + \alpha K_{\mathbf{p}})^{-1}(BQ_{\mathbf{u}}^{-1}B^t). \tag{58}$$

Now $F_{\mathbf{p}}$ and $K_{\mathbf{p}}$ need only be defined. We define each operator by solving the least-squares problem

$$\min_{(Y_{\mathbf{p}}Q_{\mathbf{p}}^{-1})_{j*}} ||(DQ_{\mathbf{u}}^{-1}Y_{\mathbf{u}})_{j*} - (Y_{\mathbf{p}}Q_{\mathbf{p}}^{-1})_{j*}D||^2_{Q_{\mathbf{u}}^{-1}}, \tag{59}$$

row by row, where Y_* is either F_* or K_*. Then $X_{\mathbf{p}}$ can be approximated as

$$X_{\mathbf{p}} \approx \hat{X}_{\mathbf{p}} := -(DQ_{\mathbf{u}}^{-1}D^t)[DQ_{\mathbf{u}}^{-1}(\tfrac{1}{\Delta t}Q_{\mathbf{u}} + F_{\mathbf{u}} + \alpha\gamma K_{\mathbf{u}})Q_{\mathbf{u}}^{-1}D^t]^{-1}(DQ_{\mathbf{u}}^{-1}D^t). \tag{60}$$

In practice, we replace $Q_{\mathbf{u}}^{-1}$ by a diagonal approximation. Then an application of $\hat{X}_{\mathbf{p}}^{-1}$ requires only matrix-vector products and two solves with the scalar Laplacian-like operator $DQ_{\mathbf{u}}^{-1}D^t$.

Fourier analysis can be employed to choose α to best satisfy $\hat{X}_{\mathbf{p}} \approx -DX_{\mathbf{u}}^{-1}D^t$. The analysis of [26] is extended to yield the choice

$$\alpha = \frac{(Reh^4 + \Delta t)\left[Reh^4 + \Delta t + \gamma Ha^2h^2\Delta t\left(||\vec{b}||^2 - \frac{||\vec{a}\times\vec{b}||^2}{||\vec{a}||^2}\right)\right] + Re^2h^2||\vec{a}||^2\Delta t^2}{\left[Reh^4 + \Delta t + \gamma Ha^2h^2\Delta t\left(||\vec{b}||^2 - \frac{||\vec{a}\times\vec{b}||^2}{||\vec{a}||^2}\right)\right]^2 + Re^2h^2||\vec{a}||^2\Delta t^2}. \tag{61}$$

This value can be automatically generated using only the problem parameters Re, Re_m, and S, the discretization parameters h and Δt, and the previous values of $\mathbf{u}$ and $\mathbf{B}$. At steady state, this expression reduces to

$$\alpha = \frac{1 + \gamma Ha^2h^2\left(||\vec{b}||^2 - \frac{||\vec{a}\times\vec{b}||^2}{||\vec{a}||^2}\right) + Re^2h^2||\vec{a}||^2}{\left[1 + \gamma Ha^2h^2\left(||\vec{b}||^2 - \frac{||\vec{a}\times\vec{b}||^2}{||\vec{a}||^2}\right)\right]^2 + Re^2h^2||\vec{a}||^2}. \tag{62}$$

With these approximations for $X_{\mathbf{u}}$ and $X_{\mathbf{p}}$, we can write an approximation for the perturbed Navier–Stokes saddle point system as

$$\mathcal{X}_{\mathbf{u}} \approx \hat{\mathcal{X}}_{\mathbf{u}} := \begin{pmatrix} \hat{X}_{\mathbf{u}} & D^t \\ 0 & \hat{X}_{\mathbf{p}} \end{pmatrix} \tag{63}$$

with $\hat{X}_{\mathbf{u}}$ and $\hat{X}_{\mathbf{p}}$ defined as in (45) and (60).

5. Numerical experiments. This section empirically investigates the performance of block preconditioners of the form

$$\mathcal{P} = \begin{pmatrix} \hat{\mathcal{M}}_{\mathbf{B}} & -\mathcal{Z}_{\mathbf{u}}^t \\ 0 & \hat{\mathcal{X}}_{\mathbf{u}} \end{pmatrix}, \tag{64}$$

where $\hat{\mathcal{X}}_{\mathbf{u}}$ is defined in (63) and for $\hat{\mathcal{M}}_{\mathbf{B}}$ either the mass matrix augmented approximation $\hat{\mathcal{M}}_{\mathbf{B},M}$ defined in (13) or the grad-div augmented approximation $\hat{\mathcal{M}}_{\mathbf{B},GD}$ defined in (25) is used. These preconditioners are referred to as $\mathcal{P}_M$ and $\mathcal{P}_{GD}$, depending on the Maxwell approximation.

The section proceeds by showing how the preconditioners perform for a series of test problems, considering how the preconditioners depend on the linearization method (Picard iteration or Newton's method), the nondimensional parameters Re, Re_m, and S, and the discretization parameters h and Δt. The effectiveness and algorithmic scalability of these preconditioners is investigated over a range of parameters. In all of the following results, a relative tolerance of 10^{-4} is used for all nonlinear iterations and a relative tolerance of 10^{-3} is used as the stopping criterion for preconditioned GMRES unless otherwise noted. For steady problems, linear iterations are averaged over all nonlinear iterations. For transient problems, backward Euler time marching with a constant time step is employed, and nonlinear and linear iterations are averaged over all instances up to time $t = 2$.

The implementation is in the Trilinos framework [17] using the Teko package [7] to construct the block preconditioners and GMRES from AztecOO [16, 20]. For component solves, algebraic multigrid from the ML package [12] is employed, with incomplete factorization smoothers coming from IFPACK. Traditional smoothed aggregation AMG is employed for all component solves except the mass matrix augmented operator $(\frac{1}{\Delta t} + k)Q_{\mathbf{B}} + F_{\mathbf{B}}$ for which we use the eddy current multigrid routine [19] as implemented in ML. In all cases, one V-cycle is used. For scalar solves, Gauss–Seidel smoothing is used, and for $\mathbf{u}$ and $\mathbf{B}$ solves, either a Gauss–Seidel smoother or a domain decomposition smoother using an ILU(0) factorization on each subdomain is used. The more expensive ILU smoothers are employed for convection-dominated operators (that is, the velocity operator for the MHD generator test problem and both vector operators for the Kelvin–Helmholtz problem described below) as these smoothers lead to more effective overall performance in these cases.

5.1. Lid driven cavity. This test problem is an MHD variant of the classical hydrodynamic lid driven cavity problem [10]. The domain $[-0.5, 0.5]^3$ is considered with the fluid driven by the boundary condition $\vec{u} = (1, 0, 0)$ on the top. No slip boundary conditions ($\vec{u} = \vec{0}$) are imposed on the other walls. A magnetic field is imposed from the right to left by setting the tangential component of $\vec{B}$ as $\vec{B} \times \vec{n} = (-1, 0, 0) \times \vec{n}$ on each wall. Both a transient formulation and a steady-state formulation of this test problem are considered.

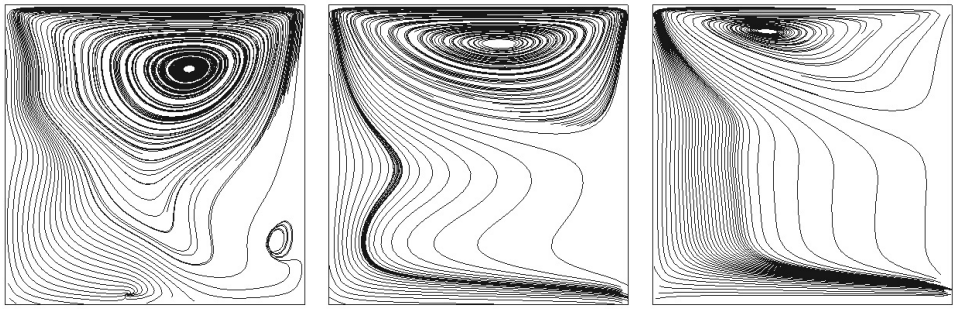

FIG. 1. *Streamlines of $\vec{u}$ in the cross section at $z = 0$ for the steady lid driven cavity problem with $Re = 100$ and $Re_m = 1$ (left), 10 (center), 100 (right).*

For relatively small Reynolds number ($Re < 1000$), the three-dimensional hydrodynamic cavity has a steady-state solution, allowing the steady formulation to be considered in this regime. When no time derivatives are present, the equations can be rescaled so that $S = 1$ and all coupling is accounted for by Re and Re_m. Simulations of the two-dimensional lid driven cavity [26, 32] have demonstrated that without a magnetic field, one large recirculation develops in the domain, but as the strength of the magnetic field increases (i.e., Re_m increases), the flow is affected and multiple recirculation regions with individual eddies develop. The number of eddies increases with both greater Re and greater Re_m, indicating that coupling effects for this problem depend strongly on both fluid and magnetic Reynolds numbers. The effect of increased coupling is shown for the three-dimensional lid driven cavity in Figure 1. As in the two-dimensional case, the large eddy is pushed upward as Re_m increases.

5.1.1. Transient results. For the transient MHD equations, the Alfvén CFL is given by $CFL_{Alfven} = \sqrt{S}\frac{\Delta t}{h}$. If $CFL_{Alfven} \leq 1$, then the Alfvén wave is captured time accurately, whereas if $CFL_{Alfven} > 1$, then the Alfvén wave is a stiff mode in the linear system and requires implicit time integration to be resolved [6]. Because the Alfvén wave results from the coupling between hydrodynamics and electromagnetics, it is important to capture the off-diagonal coupling when $CFL_{Alfven} > 1$. The Hartmann number $Ha = \sqrt{SReRe_m}$ is a measure of the degree of physical coupling in the linear system such that the electromagnetics and hydrodynamics are more strongly coupled for larger Ha. Hence, we use CFL_{Alfven} and Ha as quantities against which to test the robustness of solvers.

The performance of the proposed preconditioners on Picard and Newton linearizations is detailed in Tables 2 and 3. The corresponding nonlinear iteration counts are reported in Table 4. To obtain these results, we fix $Re = 100$ and $\frac{\Delta t}{h} = 1$. We obtain different values of CFL_{Alfven} and Ha by varying S and Re_m. These results were run on a fixed $32 \times 32 \times 32$ mesh. These results demonstrate a slight dependence on CFL_{Alfven} and Ha, with more iterations required for more strongly coupled problems, but generally low iteration counts for both preconditioners. In Figure 2, we show parallel scaling results for these preconditioners for Newton's method applied to the case with $CFL_{Alfven} = 10, Ha = 100$. We use a sequence of meshes with $\frac{\Delta t}{h} = 1$ and $h = \frac{1}{8}, \frac{1}{16}, \frac{1}{32}, \frac{1}{64}, \frac{1}{128}$ on 1, 8, 64, 512, and 4096 processors such that there are approximately 20,000 unkowns per processor. As a benchmark, we compare the results against a domain decomposition preconditioner using a SuperLU domain solver [8] on each subdomain. The results demonstrate very flat scaling in both linear iteration

TABLE 2
Average number of linear iterations per nonlinear step required for convergence of the transient lid driven cavity problem with Picard linearization. $\mathcal{P}_M$ indicates the mass augmented preconditioner and $\mathcal{P}_{GD}$ indicates the grad-div augmented preconditioner.

	$\mathcal{P}_M$			$\mathcal{P}_{GD}$		
CFL_{Alfven} \ Ha	1	10	100	1	10	100
1	20	13	17	30	23	22
10	31	20	43	35	32	49
100	51	33	52	46	36	74

TABLE 3
Average number of linear iterations per nonlinear step required for convergence of the transient lid driven cavity problem with Newton's method.

	$\mathcal{P}_M$			$\mathcal{P}_{GD}$		
CFL_{Alfven} \ Ha	1	10	100	1	10	100
1	22	18	20	26	20	21
10	31	22	46	30	28	49
100	44	32	51	29	31	62

TABLE 4
Average number of nonlinear iterations required for convergence of the transient lid driven cavity problem per time step.

	Picard			Newton		
Re \ Re_m	1	10	100	1	10	100
1	2	3	3	2	2	2
10	2	2	2	2	2	2
100	3	2	2	3	2	2

count and computation time for both preconditioners. We omit scaling results for the Picard linearization as they are very similar.

5.1.2. Steady-state results. For the steady-state lid driven cavity, we fix $S = 1$ and vary Re and Re_m to obtain different degrees of coupling. Preconditioning results are shown in Tables 5 and 6 for Picard and Newton linearizations with corresponding nonlinear iteration counts reported in Table 7. These results demonstrate a stronger dependence on Re_m than Re for both preconditioners, especially in the Newton case. While the linear iterations associated with the Picard iteration are smaller for large Re_m, the nonlinear solver is not as robust as Newton's method in these cases. The results also demonstrate superior performance of $\mathcal{P}_{GD}$ for Newton's method with large Hartmann. This behavior is consistent with the eddy current solver used in $\mathcal{P}_M$ having been developed for the case where convection is not represented in $F_{\mathbf{B}}$. Hence, we see degradation in the behavior of $\mathcal{P}_M$ when the magnetic system is convection dominated. Parallel scaling is explored for $Re = 100, Re_m = 10$ in Figure 3. As in the transient case, we see very good scaling, with superior performance of $\mathcal{P}_{GD}$ for the largest simulations.

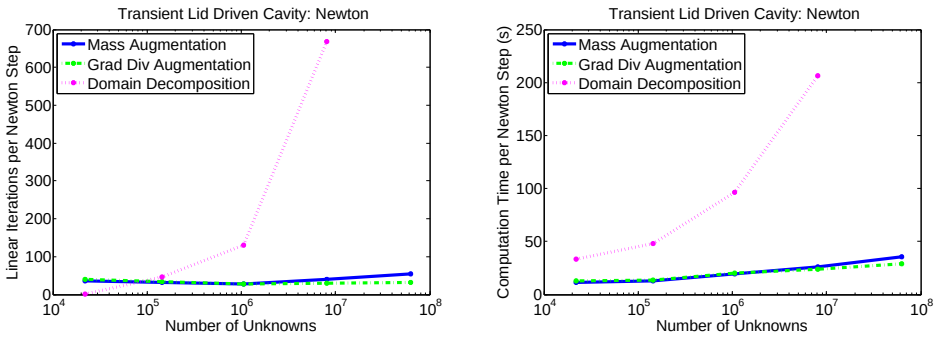

FIG. 2. *Weak scaling for the steady lid driven cavity problem with Newton's method* ($CFL_{Alfven} = 10, Ha = 100$).

TABLE 5

Average number of linear iterations per nonlinear step required for convergence of the steady lid driven cavity problem with a Picard linearization. When nonlinear convergence is not achieved, average linear iterations are averaged over the first 20 *steps of the nonlinear iteration (indicated by* $\times$ *in Table* 7*).*

	$\mathcal{P}_M$			$\mathcal{P}_{GD}$		
Re \ Re_m	1	10	100	1	10	100
1	17	18	26	16	16	23
10	26	25	26	20	31	24
100	39	26	44	24	29	38

TABLE 6

Average number of linear iterations per nonlinear step required for convergence of the steady lid driven cavity problem with Newton's method.

	$\mathcal{P}_M$			$\mathcal{P}_{GD}$		
Re \ Re_m	1	10	100	1	10	100
1	21	21	39	18	19	35
10	27	27	61	22	24	54
100	37	36	94	26	27	75

TABLE 7

Number of nonlinear iterations required for convergence of the steady lid driven cavity problem. $\times$ *indicates that the nonlinear solver did not converge within* 20 *iterations.*

	Picard			Newton		
Re \ Re_m	1	10	100	1	10	100
1	3	7	$\times$	3	3	4
10	3	7	$\times$	3	4	11
100	3	6	16	3	3	8

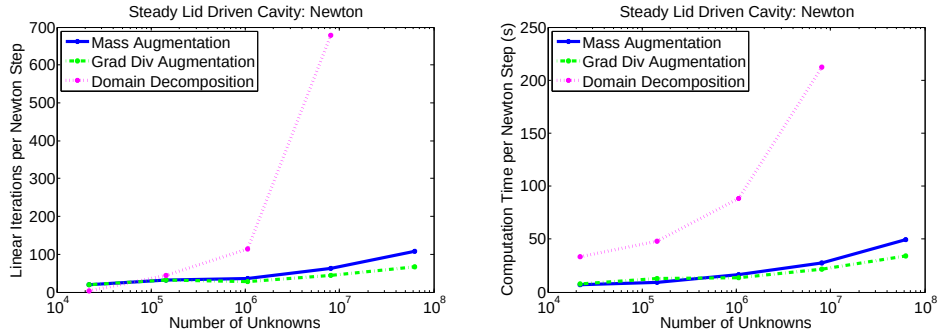

FIG. 3. *Weak scaling for the steady lid driven cavity problem with Newton's method ($Re = 100, Re_m = 10$).*

5.2. MHD generator. This is a steady test problem in which duct flow is used to generate an electric current. We take as our domain the channel $[0, 16]\times[0, 1]\times[0, 1]$. The fluid is forced in from the left through the boundary condition $\vec{u} = (1, 0, 0)$ on the left wall with natural outflow conditions on the right. No slip conditions are supplied on the other walls. An external magnetic field of magnitude B_0 is applied smoothly between $x = x_{on}$ and $x = x_{off}$ according to the function

$$B_y(x) = \tfrac{B_0}{2}[\tanh(\tfrac{x-x_{on}}{\delta}) - \tanh(\tfrac{x-x_{off}}{\delta})], \tag{65}$$

where δ controls how quickly the field transitions from off to on. We take $B_0 = 1, \delta = 0.1, x_{on} = 4$, and $x_{off} = 6$. We impose $\vec{B} = (0, B_y(x), 0)$ by setting the tangential component $\vec{B} \times \vec{n} = (0, B_y(x), 0) \times \vec{n}$. A solution is plotted for the parameters $Re = 500, Re_m = 1, S = 1.25$ (i.e., $Ha = 25$) in Figure 4. From the figure, we can see that the fluid flow is perturbed by the magnetic field and the magnetic field is bent in the flow direction to induce a current ($J = S\nabla \times \vec{B}$) in the z-direction.

For this problem a Picard iteration converges very slowly whereas Newton's method converges in either two or three iterations. In Figure 5, we present a weak scaling study of the preconditioners applied to a Newton linearization of the MHD generator problem. We use a uniform grid with $h = \frac{1}{4}, \frac{1}{8}, \frac{1}{16}$, and $\frac{1}{32}$ on 2, 16, 128, and 1024 processors such that there are approximately 35,000 unknowns per processor. We see near flat scaling for both preconditioners with respect to linear iteration counts and moderate growth in computation time as the mesh is refined. Again, the grad-div augmented preconditioner demonstrates better performance than the mass matrix augmented preconditioner for the largest simulations.

5.3. Hydromagnetic Kelvin–Helmholtz (HMKH) instability. The HMKH problem is a two-dimensional test problem on the domain $[0, 2] \times [0, 1]$. The problem consists of a shear-layer flow that results in a Kelvin–Helmholtz instability [14]. This is induced by an initial condition with $\vec{u} = (U_x, 0)$ in the top half of the domain and $\vec{u} = (-U_x, 0)$ in the bottom half. Additionally, a sheared magnetic field given by the initial condition $\vec{B} = (B_x \tanh(y/\delta), 0)$ is applied. We let $Re = Re_m = 1000$ and $S = 1$ with $U_x = 1.5$ and $B_x = 1$ such that the magnetic field allows the instability to develop. The time evolution of the solution is depicted in Figure 6. It can be seen that a vortex develops in the fluid and the magnetic field is deformed along with the

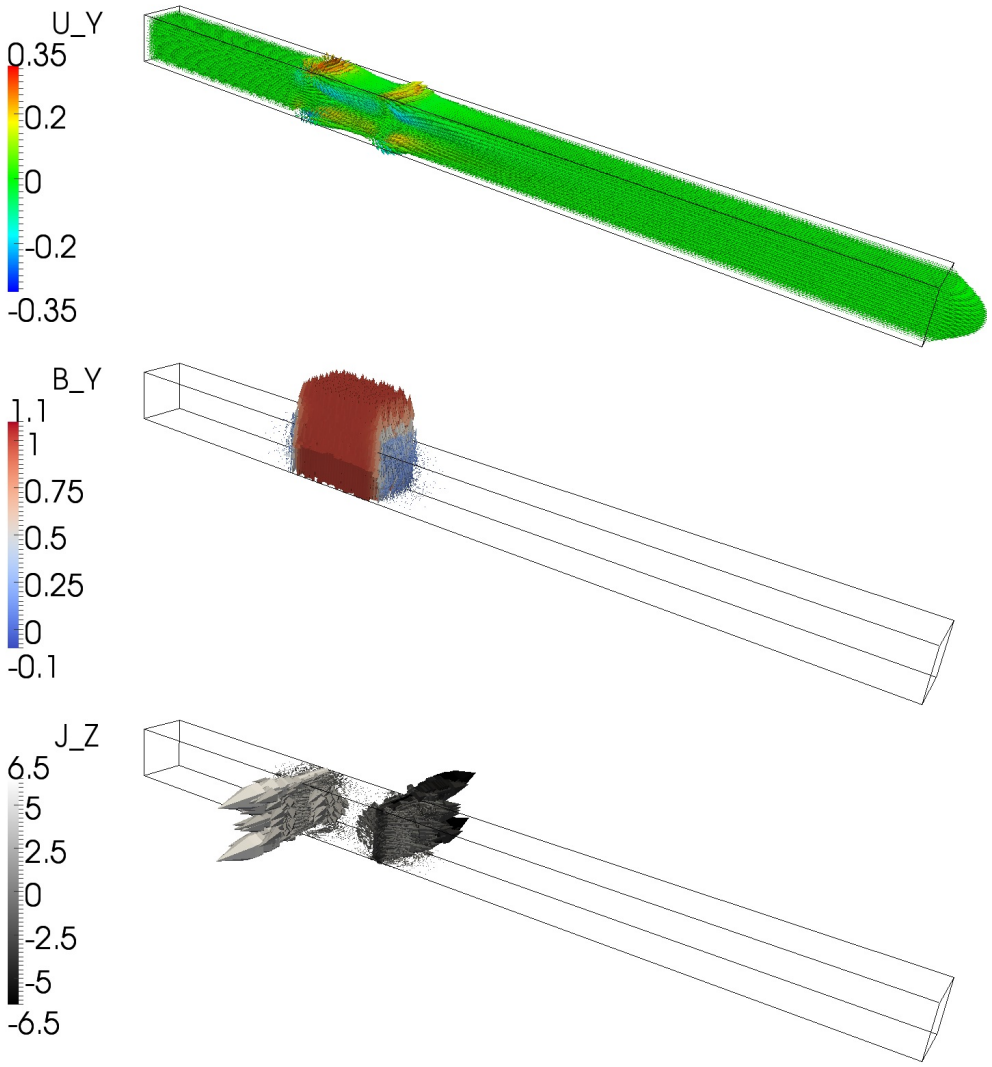

FIG. 4. *Plots of the velocity field, magnetic field, and induced current for the MHD generator problem.*

flow. The fluid CFL number for these parameters is $CFL_u = 1.5\frac{\Delta t}{h}$ while the Alfvén CFL is slightly lower at $CFL_{Alfven} = \frac{\Delta t}{h}$.

We show weak scaling results for the preconditioners for a fixed time step of $\Delta t = \frac{1}{80}$ in Figure 7. We consider uniform meshes with $h = \frac{1}{20}, \frac{1}{40}, \frac{1}{80}, \frac{1}{160}$, and $\frac{1}{320}$ on 8, 32, 128, 512, and 2048 processors such that the fluid CFL ranges from $\frac{3}{8}$ to 6 on the finest mesh, while the diffusion time scale remains more restrictive. The results show only a slight increase in iterations for the grad-div augmented preconditioner with a growth of approximately 2.5 times as the problem size increases by a factor of 256. Furthermore, a CFL of 6 is fairly large for the finest mesh and iteration count is moderate. The mass matrix augmented preconditioner, on the other hand,

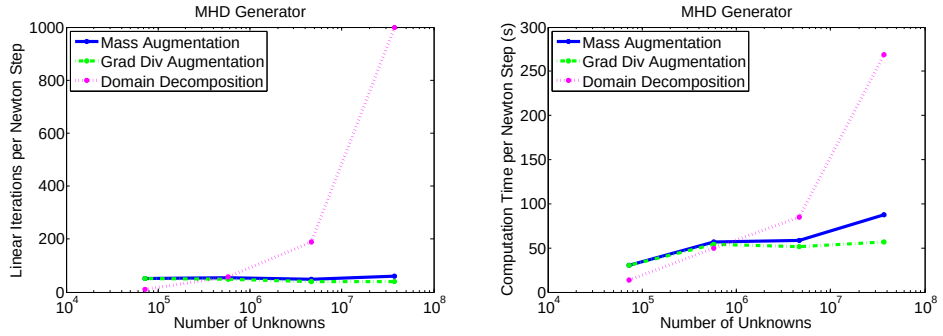

FIG. 5. *Weak scaling for the MHD generator problem* ($Ha = 25$).

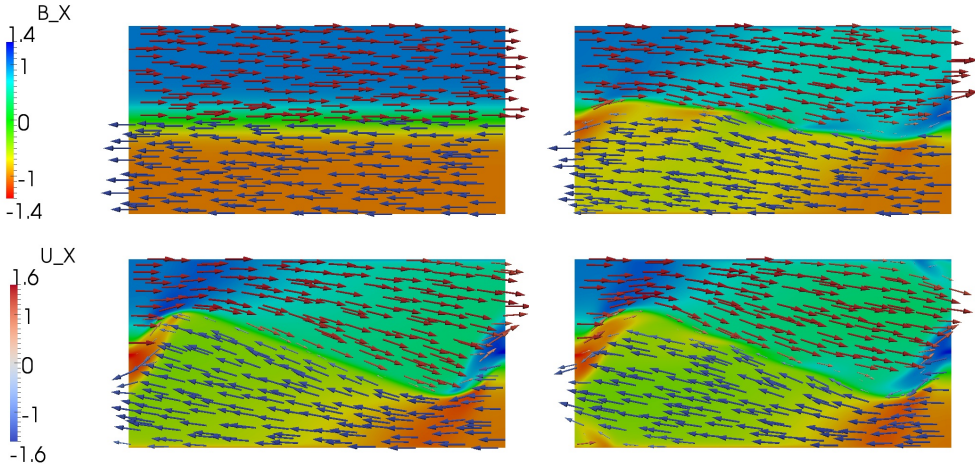

FIG. 6. *Snapshots of the HMKH problem at* $t = 0, 2, 4,$ *and* 6.

does not perform very well with iteration counts growing larger than 500 on the 2048 processor simulation. This behavior appears to be due to the magnetic Reynolds number being fairly large for this problem; that is, with $Re_m = 1000$, the magnetic block is convection dominated and the eddy current multigrid employed for this block fails. A fixed CFL study with $CFL_u = 3$ is plotted in Figure 8. In this case, mass augmentation always failed, but the grad-div preconditioner performs very well with both iteration count and computation time decreasing as the problem is refined.

6. Conclusion. In this work, we presented new block preconditioners for a dual saddle point formulation of the MHD equations with mixed finite elements. These preconditioners segregate the Maxwell and Navier–Stokes subsystems to elucidate a block 2×2 structure, each block diagonal term itself being of 2×2 saddle point structure. This allowed for the approximation of the magnetic induction saddle point system independent of the hydrodynamics.

A new grad-div augmented solver was developed for the magnetic induction saddle point system, general enough to handle time-dependent and steady formulations,

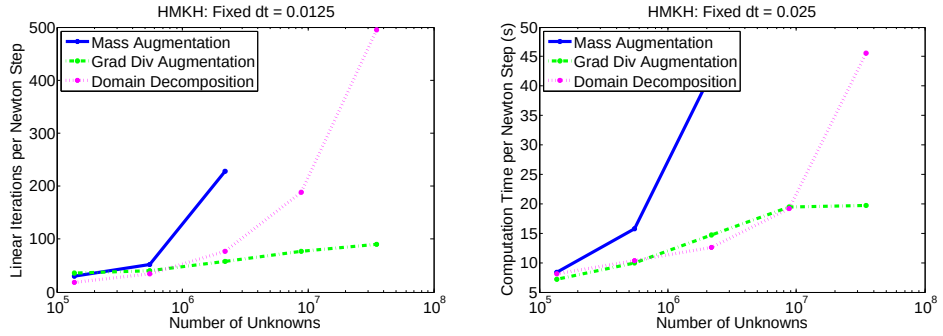

FIG. 7. *Weak scaling for the HMKH problem with a fixed time step.*

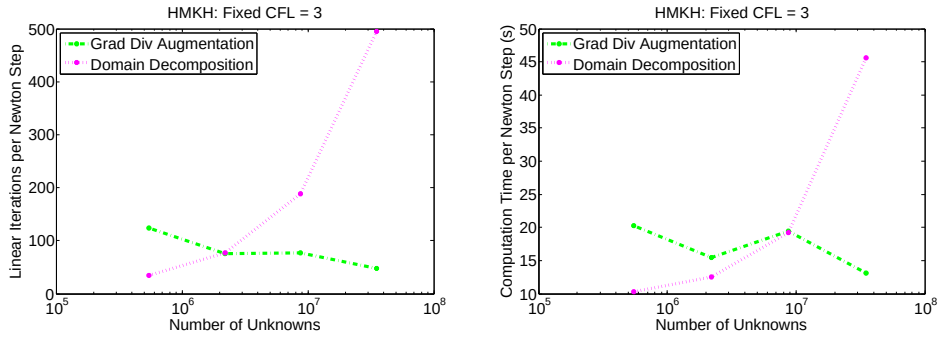

FIG. 8. *Weak scaling for the HMKH problem with a fixed CFL.*

large values of Re_m, as well as the convection operator arising from a Newton linearization of the MHD equations. An eigenvalue analysis was provided that agrees with the scalable and robust performance of this approximation. This new approximation has the advantage that it requires only traditional AMG for the subsolve with the magnetic field, and it compares well against a previously proposed mass augmented preconditioner with edge-based AMG. The grad-div augmented approximation demonstrated more robust convergence than mass augmentation for Newton's method.

To account for the electromagnetic-hydrodynamic coupling, the Schur complement on the hydrodynamic variables was approximated as a perturbation of the discrete Navier–Stokes system. Generalizing the two-dimensional steady-state analysis of [26], a relaxed commutator strategy was employed to develop a pressure Schur complement approximation that captures the velocity-pressure coupling of the perturbed Navier–Stokes system. Combining this strategy with the two augmentation-based Maxwell approximations, we obtained a block 4×4 upper triangular preconditioner for the fully coupled MHD system.

Computational results have demonstrated the effectiveness of the preconditioners. Robustness of the fluid-magnetic coupling mechanisms to physical parameters and the Alfvén CFL number was shown. Algorithmic scalability was demonstrated on a number of challenging test problems up to 2048 processors. The grad-div augmentation

method has proven to be more robust than mass matrix augmentation, especially when the magnetic Reynolds number is large.

Appendix A. Eigenvalue bounds for grad-div Maxwell approximation. The performance of the preconditioner $\hat{M}_{\mathbf{B},GD}$ for $M_{\mathbf{B}}$ is determined by the generalized eigenvalue problem

$$\begin{pmatrix} \frac{1}{\Delta t}Q_{\mathbf{B}} + F_{\mathbf{B}} & D \\ D & 0 \end{pmatrix} \begin{pmatrix} \mathbf{B} \\ \mathbf{r} \end{pmatrix} = \lambda \begin{pmatrix} \frac{1}{\Delta t}Q_{\mathbf{B}} + F_{\mathbf{B}} + kD^tQ_{\mathbf{r}}^{-1}D & D^t(\frac{1}{\Delta t}Q_{\mathbf{r}} + kL_{\mathbf{r}})^{-1}(\frac{1}{\Delta t}Q_{\mathbf{r}} + 2kL_{\mathbf{r}}) \\ 0 & -Q_{\mathbf{r}}(\frac{1}{\Delta t}Q_{\mathbf{r}} + kL_{\mathbf{r}})^{-1}L_{\mathbf{r}} \end{pmatrix} \begin{pmatrix} \mathbf{B} \\ \mathbf{r} \end{pmatrix}. \tag{66}$$

In this section, it is shown that all of the eigenvalues of this problem are exactly one.

The bottom row of (66) yields $\mathbf{r} = -\frac{1}{\lambda}(\frac{1}{\Delta t}L_{\mathbf{r}}^{-1} + kQ_{\mathbf{r}}^{-1})D\mathbf{B}$. Substituting this into the top row and simplifying gives

$$\left[(1-\lambda)(\tfrac{1}{\Delta t}Q_{\mathbf{B}} + F_{\mathbf{B}}) - \tfrac{1-\lambda}{\lambda}\tfrac{1}{\Delta t}D^tL_{\mathbf{r}}^{-1}D - \tfrac{(1-\lambda)^2}{\lambda}kD^tQ_{\mathbf{r}}^{-1}D\right]\mathbf{B} = 0. \tag{67}$$

Using a discrete Hodge decomposition, $\mathbf{B}$ can be written as the sum of its curl-free part $\mathbf{B}_C$ and its divergence-free part $\mathbf{B}_D$ [15]. The curl-free part is in the left null space of $F_{\mathbf{B}}$ (i.e., $\mathbf{B}_C^tF_{\mathbf{B}} = 0$) and the divergence-free part satisfies $D\mathbf{B}_D = 0$. Given this notation, some properties proven in [15] can be written as

$$\mathbf{B}_D^tQ_{\mathbf{B}}\mathbf{B} = \mathbf{B}_D^tQ_{\mathbf{B}}\mathbf{B}_D, \tag{68}$$

$$\mathbf{B}_C^tQ_{\mathbf{B}}\mathbf{B} = \mathbf{B}_C^tQ_{\mathbf{B}}\mathbf{B}_C, \tag{69}$$

$$\mathbf{B}_C^tD^tL_{\mathbf{r}}^{-1}D\mathbf{B}_C = \mathbf{B}_C^tQ_{\mathbf{B}}\mathbf{B}_C. \tag{70}$$

Left multiplying (67) by $\mathbf{B}_D^t$ and using these relations yields

$$(1-\lambda)\mathbf{B}_D^t(\tfrac{1}{\Delta t}Q_{\mathbf{B}} + F_{\mathbf{B}})\mathbf{B} = 0. \tag{71}$$

As there are $\dim(\mathbf{B}) - \dim(\mathbf{r})$ linearly independent vectors $\mathbf{B}_D$ satisfying $D\mathbf{B}_D = 0$, this implies that $\lambda = 1$ with multiplicity $\dim(\mathbf{B}) - \dim(\mathbf{r})$. Left multiplying (67) by $\mathbf{B}_C^t$ and multiplying through by $-\lambda$, we obtain

$$(1-\lambda)^2(\tfrac{1}{\Delta t}\mathbf{B}_C^tQ_{\mathbf{B}}\mathbf{B}_C + k\mathbf{B}_C^tD^tQ_{\mathbf{r}}^{-1}D\mathbf{B}_C) = 0. \tag{72}$$

Observe that $\frac{1}{\Delta t}\mathbf{B}_C^tQ_{\mathbf{B}}\mathbf{B}_C + k\mathbf{B}_C^tD^tQ_{\mathbf{r}}^{-1}D\mathbf{B}_C$ is a positive real number. Since there are $\dim(\mathbf{B})$ linearly independent vectors satisfying $\mathbf{B}_C^tF_{\mathbf{B}} = 0$, we obtain $\lambda = 1$ with multiplicity $2\dim(\mathbf{r})$. In total, we have $\lambda = 1$ with multiplicity $\dim(\mathbf{B}) + \dim(\mathbf{r})$, so all eigenvalues of the preconditioned system are one. This result is independent of Δt and k.

Appendix B. Simplification of the continuous perturbation operator. Using the expression (28) for $(\mathcal{M}_{\mathbf{B}}^{-1})_{0,0}$ and the correspondences summarized in Table 1, we obtain that the continous operator associated with $(\mathcal{M}_{\mathbf{B}}^{-1})_{0,0}Z_{\mathbf{u}}^t$ is

$$\begin{aligned} &[\tfrac{1}{\Delta t}I + \tfrac{S}{Re_m}\nabla\times\nabla\times - \delta\nabla\times(\vec{a}\times\cdot)]^{-1} \\ &\quad\circ\left[\tfrac{1}{\Delta t}I + \tfrac{S}{Re_m}\nabla\times\nabla\times - \delta\nabla\times(\vec{a}\times\cdot) - \tfrac{1}{\Delta t}\nabla\chi_r^{-1}\nabla\cdot\right] \\ &\quad\circ\left[\tfrac{1}{\Delta t}I + \tfrac{S}{Re_m}\nabla\times\nabla\times - \delta\nabla\times(\vec{a}\times\cdot)\right]^{-1}\nabla\times(\cdot\times\vec{b}), \end{aligned} \tag{73}$$

where $\chi_r = -\nabla\cdot[\frac{1}{\Delta t}I + \frac{S}{Re_m}\nabla\times\nabla\times - \delta\nabla\times(\vec{a}\times\cdot)]^{-1}\nabla$. Observe that the magnetic convection-diffusion operator and the curl operator commute as

$$\begin{aligned}(74)\qquad &\left[\tfrac{1}{\Delta t}I + \tfrac{S}{Re_m}\nabla\times\nabla\times - \delta\nabla\times(\vec{a}\times\cdot)\right]\nabla\times\\ &= \nabla\times\left[\tfrac{1}{\Delta t}I + \tfrac{S}{Re_m}\nabla\times\nabla\times - \delta\vec{a}\times(\nabla\times\cdot)\right],\end{aligned}$$

which implies that

$$\begin{aligned}(75)\qquad &\left[\tfrac{1}{\Delta t}I + \tfrac{S}{Re_m}\nabla\times\nabla\times - \delta\nabla\times(\vec{a}\times\cdot)\right]^{-1}\nabla\times\\ &= \nabla\times\left[\tfrac{1}{\Delta t}I + \tfrac{S}{Re_m}\nabla\times\nabla\times - \delta\vec{a}\times(\nabla\times\cdot)\right]^{-1}.\end{aligned}$$

Using (B) followed by $\nabla\cdot\nabla\times = 0$ and (B), expression (73) reduces to

$$(76)\qquad \left[\tfrac{1}{\Delta t}I + \tfrac{S}{Re_m}\nabla\times\nabla\times - \delta\nabla\times(\vec{a}\times\cdot)\right]^{-1}\nabla\times(\cdot\times\vec{b}).$$

Then $Z_{\mathbf{B}}(\mathcal{M}_{\mathbf{B}}^{-1})_{0,0}Z_{\mathbf{u}}^t$ corresponds to

$$(77)\qquad [S\vec{b}\times(\nabla\times\cdot) + \delta S\cdot\times(\nabla\times\vec{b})]\left[\tfrac{1}{\Delta t}I + \tfrac{1}{Re_m}\nabla\times\nabla\times - \delta\nabla\times(\vec{a}\times\cdot)\right]^{-1}\nabla\times(\cdot\times\vec{b}).$$

REFERENCES

[1] J. H. Adler, T. R. Benson, E. C. Cyr, S. P. MacLachlan, and R. S. Tuminaro, *Monolithic multigrid methods for two-dimensional resistive magnetohydrodynamics*, SIAM J. Sci. Comput., 38 (2016), pp. B1–B24.

[2] P. B. Bochev, C. J. Garasi, J. J. Hu, A. C. Robinson, and R. S. Tuminaro, *An improved algebraic multigrid method for solving Maxwell's equations*, SIAM J. Sci. Comput., 25 (2003), pp. 623–642.

[3] L. Chacón, *An optimal, parallel, fully implicit Newton-Krylov solver for three-dimensional viscoresistive magnetohydrodynamics*, Phys. Plasmas, 15 (2008), 056103.

[4] L. Chacón, *Scalable parallel implicit solvers for 3D magnetohydrodynamics*, J. Phys. Conf. Ser., 125 (2008), 012041.

[5] L. Chacón, D. Knoll, and J. Finn, *An implicit, nonlinear reduced resistive MHD solver*, J. Comput. Phys., 178 (2002), pp. 15–36.

[6] E. C. Cyr, J. N. Shadid, R. S. Tuminaro, R. P. Pawlowski, and L. Chacón, *A new approximate block factorization preconditioner for two-dimensional incompressible (reduced) resistive MHD*, SIAM J. Sci. Comput., 35 (2013), pp. B701–B730.

[7] E. C. Cyr, J. N. Shadid, and R. S. Tuminaro, *Teko: A block preconditioning capability with concrete example applications in Navier–Stokes and MHD*, SIAM J. Sci. Comp., to appear.

[8] J. W. Demmel, S. C. Eisenstat, J. R. Gilbert, X. S. Li, and J. W. H. Liu, *A supernodal approach to sparse partial pivoting*, SIAM Journal on Matrix Analysis and Applications, 20 (1999), pp. 720–755.

[9] H. Elman, V. Howle, J. Shadid, R. Shuttleworth, and R. Tuminaro, *A taxonomy and comparison of parallel block multi-level preconditioners for the incompressible Navier-Stokes equations*, J. Comput. Phys., 227 (2008), pp. 1790–1808.

[10] H. Elman, D. Silvester, and A. Wathen, *Finite Elements and Fast Iterative Solvers with Applications in Incompressible Fluid Dynamics*, Oxford University Press, New York, 2005.

[11] H. Elman and R. Tuminaro, *Boundary conditions in approximate commutator preconditioners for the Navier-Stokes equations*, Electron. Trans. Numer. Anal., 35 (2009), pp. 1068–9613.

[12] M. Gee, C. Siefert, J. Hu, R. Tuminaro, and M. Sala, *ML* 5.0 *Smoother Aggregration User's Guide*, Technical. report SAND2006-2649, Sandia National Laboratories, Albuzuerque, NM, 2006.

[13] J.-F. Gerbeau, C. L. Bris, and T. Lelièvre, *Mathematical Methods for the Magnetohydrodynamics of Liquid Metals*, Oxford University Press, Oxford, 2006.

[14] J. Goedbloed and S. Poedts, *Principles of Magnetohydrodynamics with Applications to Laboratory and Astrophysical Plasmas*, Cambridge University Press, Cambridge, 2004.
[15] C. Greif and D. Schötzau, *Preconditioners for the discretized time-harmonic Maxwell equations in mixed form*, Numer. Linear Algebra Appl., 14 (2007), pp. 281–297.
[16] M. Heroux, *AztecOO User Guide*, Technical. report SAND2004-3796, Sandia National Laboratories, Albuquerque, NM, 2004.
[17] M. Heroux, R. Bartlett, V. Howle, R. Hoekstra, J. Hu, T. Kolda, R. Lehoucq, K. Long, R. Pawlowski, E. Phipps, A. Salinger, H. Thornquist, R. Tuminaro, J. Willenbring, A. Williams, and K. Stanley, *An overview of the Trilinos project*, ACM Trans. Math. Software., 31 (2005), pp. 397–423.
[18] R. Hiptmair and J. Xu, *Nodal auxiliary space preconditioning in H(Curl) and H(Div) spaces*, SIAM J. Numer. Anal., 45 (2007), pp. 2483–2509.
[19] J. J. Hu, R. S. Tuminaro, P. B. Bochev, C. J. Garasi, and A. C. Robinson, *Toward an h-independent algebraic multigrid method for Maxwell's equations*, SIAM J. Sci. Comput., 27 (2006), pp. 1669–1688.
[20] S. Hutchinson, L. Prevost, J. N. Shadid, C. Tong, and R. S. Tuminaro, *Aztec User's Guide Version* 2.0, Technical. report Sand99-8801J, Sandia National Laboratories, Albuzuerque, NM, 1999.
[21] R. Keppens, G. Tóth, A. Botchev, and A. Ploeg, *Implicit and semi-implicit schemes: Algorithms*, Internat. J. Numer. Methods. Fluids, 30 (1999), pp. 335–352.
[22] P. Lin, J. Shadid, R. Tuminaro, M. Sala, G. Hennigan, and R. Pawlowski, *A parallel fully-coupled algebraic multilevel preconditioner applied to multiphysics PDE applications: Drift-diffusion, flow/transport/reaction, resistive MHD*, Internat. J. Numer. Methods. Fluids, 64 (2010), pp. 1148–1179.
[23] M. F. Murphy, G. H. Golub, and A. J. Wathen, *A note on preconditioning for indefinite linear systems*, SIAM J. Sci. Comput., 21 (2000), pp. 1969–1972.
[24] J.-C. Nédélec, *Mixed finite elements in* $\mathbb{R}^3$, Numer. Math., 35 (1980), pp. 315–341.
[25] E. Phillips and H. Elman, *A stochastic approach to uncertainty in the equations of MHD kinematics*, J. Comput. Phys., 284 (2015), pp. 334–350.
[26] E. G. Phillips, H. C. Elman, E. C. Cyr, J. N. Shadid, and R. P. Pawlowski, *A block preconditioner for an exact penalty formulation for stationary MHD*, SIAM J. Sci. Comput., 36 (2014), pp. B930–B951.
[27] Y. Saad and M. H. Schultz, *GMRES: A generalized minimal residual algorithm for solving nonsymmetric linear systems*, SIAM J. Sci. Stat. Comput., 7 (1986), pp. 856–869.
[28] A. Schneebeli and D. Schötzau, *Mixed finite elements for incompressible magnetohydrodynamics*, C. R. Acad. Sci. Paris Sërïes I Math, 337 (2003), pp. 71–74.
[29] D. Schötzau, *Mixed finite element methods for stationary incompressible magnetohydrodynamics*, Numer. Math., 96 (2004), pp. 771–800.
[30] J. Shadid, R. Pawlowski, J. Banks, L. Chacón, P. Lin, and R. Tuminaro, *Towards a scalable fully-implicit fully-coupled resistive MHD formulation with stabilized FE methods*, J. Comput. Phys., 229 (2010), pp. 7649–7671.
[31] J. N. Shadid, R. P. Pawlowski, E. C. Cyr, R. S. Tuminaro, L. Chacon, and P. D. Weber, *Scalable implicit incompressible resistive MHD with stabilized FE and fully-coupled Newton-Krylov-AMG*, Comput. Methods Appl. Mech. Engrg., 304 (2016), pp. 1–25.
[32] V. Shatrov, G. Mutschke, and G. Gerberth, *Three-dimensional linear stability analysis of lid-driven magnetohydrodynamic cavity flow*, Phys. Fluids, 15 (2003), pp. 2141–2151.
[33] G. Tóth, R. Keppens, and A. Botchev, *Implicit and semi-implicit schemes in the Versatile Advection code: Numerical tests*, Astronom. Astrophys., 332 (1998), pp. 1159–1170.
[34] S.-L. Wu, T.-Z. Huang, and C.-X. Li, *Modified block preconditioners for the discretized time-harmonic Maxwell equations in mixed form*, J. Comput. Appl. Math., 237 (2013), pp. 419–431.
[35] S.-L. Wu, T.-Z. Huang, and L. Li, *Block triangular preconditioner for static Maxwell equations*, Comput. Appl. Math., 30 (2011), pp. 589–612.