SIAM J. SCI. COMPUT.
Vol. 39, No. 6, pp. A2993–A3013

PRECONDITIONERS FOR MIXED FINITE ELEMENT DISCRETIZATIONS OF INCOMPRESSIBLE MHD EQUATIONS*

MICHAEL WATHEN†, CHEN GREIF†, AND DOMINIK SCHÖTZAU‡

Abstract. We consider preconditioning techniques for a mixed finite element discretization of an incompressible magnetohydrodynamics (MHD) problem. Upon discretization and linearization, a 4-by-4 nonsymmetric block-structured linear system needs to be (repeatedly) solved. One of the principal challenges is the presence of a skew-symmetric term that couples the fluid velocity with the magnetic field. We propose a preconditioner that exploits the block structure of the underlying linear system, utilizing and combining effective solvers for the mixed Maxwell and the Navier–Stokes subproblems. We perform a spectral analysis for an ideal version of the preconditioner, and develop and test a practical version of it. Large-scale numerical results for linear systems of dimensions up to 10^7 in two and three dimensions validate the effectiveness of our approach.

Key words. incompressible magnetohydrodynamics, mixed finite element methods, saddle-point linear systems, preconditioners, Krylov subspace methods, Schur complement

AMS subject classifications. 65F08, 65F10, 65F15, 65N22, 74S05

DOI. 10.1137/16M1098991

1. Introduction.

Incompressible magnetohydrodynamics (MHD) describes the behavior of electrically conductive incompressible fluids (liquid metals, plasma, salt water, etc.) in an electromagnetic field [1, 6]. It has a number of important applications in technology and industry, along with geophysical and astrophysical applications. Some such applications are electromagnetic pumping, aluminum electrolysis, the Earth's molten core, and solar flares.

Following the formulation in [7, 19], we consider the steady-state incompressible MHD model equation in mixed form:

$$-\nu\,\Delta \boldsymbol{u} + (\boldsymbol{u}\cdot\nabla)\boldsymbol{u} + \nabla p - \kappa\,(\nabla\times\boldsymbol{b})\times\boldsymbol{b} = \boldsymbol{f} \qquad \text{in } \Omega, \tag{1.1a}$$

$$\nabla\cdot\boldsymbol{u} = 0 \qquad \text{in } \Omega, \tag{1.1b}$$

$$\kappa\nu_m\,\nabla\times(\nabla\times\boldsymbol{b}) + \nabla r - \kappa\,\nabla\times(\boldsymbol{u}\times\boldsymbol{b}) = \boldsymbol{g} \qquad \text{in } \Omega, \tag{1.1c}$$

$$\nabla\cdot\boldsymbol{b} = 0 \qquad \text{in } \Omega. \tag{1.1d}$$

Here, Ω is a bounded Lipschitz polygonal or polyhedral domain of $\mathbb{R}^d$ for $d = 2, 3$ with boundary $\partial\Omega$. The unknowns are the velocity $\boldsymbol{u}$, the hydrodynamic pressure p, the magnetic field $\boldsymbol{b}$, and the Lagrange multiplier r associated with the divergence constraint on the magnetic field. The functions $\boldsymbol{f}$ and $\boldsymbol{g}$ represent external forcing terms. Equations (1.1) are characterized by three dimensionless parameters: the hydrodynamic Reynolds number $\mathrm{Re} = 1/\nu$, the magnetic Reynolds number $\mathrm{Rm} = 1/\nu_m$, and the coupling number κ. For further discussion of these parameters and their typical values, we refer the reader to [1, 6] and the references therein.

*Submitted to the journal's Methods and Algorithms for Scientific Computing section October 17, 2016; accepted for publication (in revised form) September 27, 2017; published electronically December 19, 2017.

http://www.siam.org/journals/sisc/39-6/M109899.html

†Department of Computer Science, The University of British Columbia, Vancouver, BC, V6T 1Z4, Canada (mwathen@cs.ubc.ca, greif@cs.ubc.ca).

‡Department of Mathematics, The University of British Columbia, Vancouver, BC, V6T 1Z2, Canada (schoetzau@math.ubc.ca).

For simplicity, we consider homogeneous Dirichlet boundary conditions:

$$\boldsymbol{u} = \boldsymbol{0}, \quad \boldsymbol{n} \times \boldsymbol{b} = \boldsymbol{0}, \quad r = 0 \quad \text{on } \partial\Omega, \tag{1.2}$$

with $\boldsymbol{n}$ being the unit outward normal on $\partial\Omega$. We note that extensions to inhomogeneous and other boundary conditions can easily be developed but are omitted.

By taking the divergence of the magnetostatic equation (1.1c), we obtain the Poisson problem

$$\Delta r = \nabla \cdot \boldsymbol{g} \quad \text{in } \Omega, \quad r = 0 \quad \text{on } \partial\Omega. \tag{1.3}$$

Since $\boldsymbol{g}$ is divergence-free in physically relevant applications, the multiplier r is typically zero, and its primary purpose is to ensure stability; see also [7].

In the case of liquid metals, the ratio between magnetic viscosity, ν_m, and fluid viscosity, ν, tends to be small. For example, mercury has a ratio of about 10^{-7} with $\nu_m \approx 10^4 - 10^5$, $\nu \approx 10^{-2} - 10^{-4}$, and $\kappa \approx 10^2 - 10^9$; cf. [1]. In [1] the authors define *strong coupling* for $10^2 \leq \kappa \leq 10^9$. For more discussion of the physical parameters, we refer the reader to [1, 6, 18, 21].

Incompressible MHD problems have been extensively studied in the context of various discretizations and formulations. However, the development of preconditioned iterative solutions to large-scale MHD problems has been rather limited so far. The recent work presented in [17] considers preconditioners for an exact penalty formulation of a stationary MHD model. The approach is based on utilizing effective preconditioners for the separate subproblems, and then incorporating the coupling terms.

The discretizations considered in [17] are based on standard nodal elements for the approximation of the magnetic fields. Nodal discretizations are effective in smooth settings, but they often fail to capture singularities in nonsmooth domains; see [19] and the references therein. Our proposed method uses edge elements for the magnetic field; these are more natural for the approximation of the curl-operator and can be applied in a seamless fashion, without penalty or stabilization, to problems that involve nonsmooth domains, singularities, and other challenging settings.

The goal in this paper is to design a scalable numerical solution procedure for problem (1.1)–(1.2). We consider mixed finite element methods based on Taylor–Hood elements for the fluid variables and the lowest order Nédélec element pair for the magnetic unknowns on triangular or tetrahedral meshes; these discretizations belong to the class of conforming mixed methods considered and analyzed in [19]. A first attempt at developing block preconditioning techniques in the context of this work was done in [12, 20], where the above-mentioned finite element framework was established. In this paper we significantly expand that work and derive new Schur complement–based preconditioners, combining state-of-the-art solvers for the Navier–Stokes and Maxwell subproblems. We perform a spectral analysis (valid in two and three dimensions) and present several large-scale numerical experiments to test the viability of our approaches.

An outline of the rest of the paper follows. In section 2, we introduce the finite element discretization, consider decoupling schemes for the nonlinear iterations, and discuss the resulting matrix systems. In section 3 we review relevant preconditioners for the Navier–Stokes and Maxwell subproblems. Our new preconditioning approach for the fully coupled MHD discretization is introduced and analyzed in section 4. We also give details of a scalable implementation of the proposed preconditioners in this section. In section 5 we show a series of numerical results in two and three dimensions. Finally, we offer some concluding remarks in section 6.

TABLE 1.1
Notation for block systems.

Notation	Meaning
$\mathcal{K}$	coefficient matrix
$\mathcal{M}$	preconditioner

TABLE 1.2
Notation for superscripts.

Superscript	Meaning	Used for
NS	Navier–Stokes	$\mathcal{K}$, $\mathcal{M}$
S	Stokes	$\mathcal{K}$, $\mathcal{M}$
MX	Maxwell	$\mathcal{K}$, $\mathcal{M}$
MHD	Magnetohydrodynamics	$\mathcal{K}$, $\mathcal{M}$

TABLE 1.3
Notation for subscripts.

Subscript	Meaning	Used for
I	ideal	$\mathcal{M}$
P	practical	$\mathcal{M}$
S	Schur complement based	$\mathcal{M}$

Notation. In this paper we will follow the notation rules given in Tables 1.1 to 1.3. We will denote closely related matrices (either permutations of the original or an extra off-diagonal block) with a tilde.

2. Discretization.

In this section we specify the mixed finite element discretization for the incompressible MHD model (1.1)–(1.2), the nonlinear Picard iteration, decoupling approaches for the nonlinear problem, and the resulting linear systems arising in each iteration step.

2.1. Mixed finite element approximation.

Our mixed finite element approximation is based on the variational formulation for (1.1)–(1.2) introduced and analyzed in [19]. It consists in finding a weak solution $(\boldsymbol{u}, p, \boldsymbol{b}, r)$ in the standard Sobolev spaces

$$
\begin{aligned}
\boldsymbol{u} \in \boldsymbol{V} &= \big\{\boldsymbol{v} \in H^1(\Omega)^d \,:\, \boldsymbol{v} = \boldsymbol{0} \text{ on } \partial\Omega\big\},\\
p \in Q &= \{\, q \in L^2(\Omega) \,:\, (q,1)_\Omega = 0\},\\
\boldsymbol{b} \in \boldsymbol{C} &= \big\{\boldsymbol{c} \in L^2(\Omega)^d \,:\, \nabla \times \boldsymbol{c} \in L^2(\Omega)^{\bar{d}},\ \boldsymbol{n} \times \boldsymbol{c} = \boldsymbol{0} \text{ on } \partial\Omega\big\},\\
s \in S &= \{r \in H^1(\Omega) \,:\, r = 0 \text{ on } \partial\Omega\}.
\end{aligned}
\tag{2.1}
$$

Here and in the following, we write $(\cdot,\cdot)_\Omega$ for all L^2-inner products, and use $\bar{d} = 2d - 3$ to define the curls simultaneously for two- and three-dimensional vector fields [7].

Now let the domain Ω be divided into regular meshes $\mathcal{T}_h = \{K\}$ consisting of triangles ($d = 2$) or tetrahedra ($d = 3$) with mesh size h. We consider Taylor–Hood elements for $(\boldsymbol{u}, p)$ and the lowest order Nédélec pair for $(\boldsymbol{b}, r)$. The corresponding

finite element spaces are

$$
\begin{aligned}
\boldsymbol{V}_h &= \{\, \boldsymbol{u} \in \boldsymbol{V} \,:\, \boldsymbol{u}|_K \in \mathcal{P}_2(K)^d,\, K \in \mathcal{T}_h \,\},\\
Q_h &= \{\, p \in Q \cap H^1(\Omega) \,:\, p|_K \in \mathcal{P}_1(K),\, K \in \mathcal{T}_h \,\},\\
\boldsymbol{C}_h &= \{\, \boldsymbol{b} \in \boldsymbol{C} \,:\, \boldsymbol{b}|_K \in \boldsymbol{R}_1(K),\, K \in \mathcal{T}_h \,\},\\
S_h &= \{\, r \in S \,:\, r|_K \in \mathcal{P}_1(K),\, K \in \mathcal{T}_h \,\}.
\end{aligned} \tag{2.2}
$$

Here $\mathcal{P}_k(K)$ denotes the space of polynomials of total degree at most k on K, and $\boldsymbol{R}_1(K) = \{\boldsymbol{a} + \boldsymbol{b} \times \boldsymbol{x} : \boldsymbol{a} \in \mathbb{R}^{\overline{d}},\, \boldsymbol{b} \in \mathbb{R}^d\}$ denotes the linear edge element space on K in terms of the position vector $\boldsymbol{x}$ on K.

With the finite element spaces in place, our mixed finite element approximation to problem (1.1)–(1.2) reads as follows: find $(\boldsymbol{u}_h, p_h, \boldsymbol{b}_h, r_h) \in \boldsymbol{V}_h \times Q_h \times \boldsymbol{C}_h \times S_h$ such that

$$
\begin{aligned}
A(\boldsymbol{u}_h, \boldsymbol{v}) + O(\boldsymbol{u}_h; \boldsymbol{u}_h, \boldsymbol{v}) + C(\boldsymbol{b}_h; \boldsymbol{v}, \boldsymbol{b}_h) + B(\boldsymbol{v}, p_h) &= (\boldsymbol{f}, \boldsymbol{v})_\Omega,\\
B(\boldsymbol{u}_h, q) &= 0,\\
M(\boldsymbol{b}_h, \boldsymbol{c}) - C(\boldsymbol{b}_h; \boldsymbol{u}_h, \boldsymbol{c}) + D(\boldsymbol{c}, r_h) &= (\boldsymbol{g}, \boldsymbol{c})_\Omega,\\
D(\boldsymbol{b}_h, s) &= 0
\end{aligned} \tag{2.3}
$$

for all $(\boldsymbol{v}, q, \boldsymbol{c}, s) \in \boldsymbol{V}_h \times Q_h \times \boldsymbol{C}_h \times S_h$. Here, we denote by $(\cdot, \cdot)_\Omega$ the inner product in $L^2(\Omega)^d$. The bilinear forms A, B, M, and D are given by

$$
\begin{aligned}
A(\boldsymbol{u}, \boldsymbol{v}) &= \nu(\nabla \boldsymbol{u}, \nabla \boldsymbol{v})_\Omega, & B(\boldsymbol{u}, q) &= -(\nabla \cdot \boldsymbol{u},\, q)_\Omega,\\
M(\boldsymbol{b}, \boldsymbol{c}) &= \kappa \nu_m (\nabla \times \boldsymbol{b}, \nabla \times \boldsymbol{c})_\Omega, & D(\boldsymbol{b}, s) &= (\boldsymbol{b}, \nabla s)_\Omega.
\end{aligned}
$$

Moreover, the trilinear forms O and C are defined as

$$
\begin{aligned}
O(\boldsymbol{w}; \boldsymbol{u}, \boldsymbol{v}) &= \big((\boldsymbol{w} \cdot \nabla)\boldsymbol{u}, \boldsymbol{v}\big)_\Omega + \frac{1}{2}\big((\nabla \cdot \boldsymbol{w})\boldsymbol{u}, \boldsymbol{v}\big)_\Omega,\\
C(\boldsymbol{d}; \boldsymbol{v}, \boldsymbol{b}) &= \kappa \left(\boldsymbol{v} \times \boldsymbol{d}, \nabla \times \boldsymbol{b}\right)_\Omega.
\end{aligned}
$$

The forms A, B, and O are associated with the variational formulation of the incompressible Navier–Stokes subsystem, M and D are associated with that of the Maxwell subsystem in mixed form, and C is the coupling form which combines the two subproblems into the full MHD system. We note that the convection form O appears in a standard skew-symmetric fashion. As a consequence, the discretization (2.3) is energy-stable without violating consistency.

The discrete problem (2.3) falls into the class of conforming mixed discretization studied in [19]. Hence, it is stable and has a unique solution for small data (i.e., for sufficiently large ν, ν_m, κ and forcing terms $\boldsymbol{f}$ and $\boldsymbol{g}$ with sufficiently small L^2-norms). Moreover, we have optimal-order error estimates in natural norms, both for smooth and nonsmooth solutions. In particular, the strongest singularities of the curl-curl operator in nonconvex domains are correctly captured and resolved.

Remark 2.1. The pairs $\boldsymbol{V}_h \times Q_h$ and $\boldsymbol{C}_h \times S_h$ form standard and optimally convergent mixed discretizations for the fluid and magnetic equations in isolation. However, the approximation properties of these pairs are not properly matched for the fully coupled system (2.3). Specifically, the optimal order $\mathcal{O}(h^2)$ of the Taylor–Hood spaces $\boldsymbol{V}_h \times Q_h$ (for the H^1-norm velocity errors and the L^2-norm pressure errors) are potentially reduced due to the coupling with the lower-order magnetic spaces $\boldsymbol{C}_h \times S_h$

in (2.2). Nonetheless, we have chosen to work with the spaces in (2.2) due to computational considerations and availability of fast solvers. In particular, we avoid the need for stabilized or enriched fluid elements and are able to use the well-established auxiliary space preconditioner [10] for the lowest-order Nédélec pair.

2.2. Picard iteration. A common choice for dealing with the nonlinearity within the incompressible Navier–Stokes equations in isolation is to perform Picard or Oseen iterations [4]. We adapt this approach for the fully coupled MHD system and linearize around the current velocity and magnetic fields. Hence, given a current iterate $(\boldsymbol{u}_h, p_h, \boldsymbol{b}_h, r_h)$, we solve for updates $(\delta\boldsymbol{u}_h, \delta p_h, \delta\boldsymbol{b}_h, \delta r_h)$ and introduce the next iterate by setting

$$\begin{aligned} \boldsymbol{u}_h &\to \boldsymbol{u}_h + \delta\boldsymbol{u}_h, \quad p_h \to p_h + \delta p_h, \\ \boldsymbol{b}_h &\to \boldsymbol{b}_h + \delta\boldsymbol{b}_h, \quad r_h \to r_h + \delta r_h. \end{aligned}$$

The updates $(\delta\boldsymbol{u}_h, \delta p_h, \delta\boldsymbol{b}_h, \delta r_h) \in \boldsymbol{V}_h \times Q_h \times \boldsymbol{C}_h \times S_h$ are found by solving the Picard system

$$\begin{aligned} A(\delta\boldsymbol{u}_h, \boldsymbol{v}) + O(\boldsymbol{u}_h; \delta\boldsymbol{u}_h, \boldsymbol{v}) + C(\boldsymbol{b}_h; \boldsymbol{v}, \delta\boldsymbol{b}_h) + B(\boldsymbol{v}, \delta p_h) &= R_u(\boldsymbol{u}_h, \boldsymbol{b}_h, p_h; \boldsymbol{v}), \\ B(\delta\boldsymbol{u}_h, q) &= R_p(\boldsymbol{u}_h; q), \\ M(\delta\boldsymbol{b}_h, \boldsymbol{c}) + D(\boldsymbol{c}, \delta r_h) - C(\boldsymbol{b}_h; \delta\boldsymbol{u}_h, \boldsymbol{c}) &= R_b(\boldsymbol{u}_h, \boldsymbol{b}_h, r_h; \boldsymbol{c}), \\ D(\delta\boldsymbol{b}_h, s) &= R_r(\boldsymbol{b}_h; s) \end{aligned} \tag{2.4}$$

for all $(\boldsymbol{v}, q, \boldsymbol{c}, s) \in \boldsymbol{V}_h \times Q_h \times \boldsymbol{C}_h \times S_h$. Note that this system is linearized around $(\boldsymbol{u}_h, \boldsymbol{b}_h)$. The right-hand side linear forms correspond to the residual at the current iteration $(\boldsymbol{u}_h, p_h, \boldsymbol{b}_h, r_h)$ and are defined by

$$\begin{aligned} R_u(\boldsymbol{u}_h, \boldsymbol{b}_h, p_h; \boldsymbol{v}) &= (\boldsymbol{f}, \boldsymbol{v})_\Omega - A(\boldsymbol{u}_h, \boldsymbol{v}) - O(\boldsymbol{u}_h; \boldsymbol{u}_h, \boldsymbol{v}) \\ &\quad - C(\boldsymbol{b}_h; \boldsymbol{v}, \boldsymbol{b}_h) - B(\boldsymbol{v}, p_h), \\ R_p(\boldsymbol{u}_h; q) &= -B(\boldsymbol{u}_h, q), \\ R_b(\boldsymbol{u}_h, \boldsymbol{b}_h, r_h; \boldsymbol{c}) &= (\boldsymbol{g}, c)_\Omega - M(\boldsymbol{b}_h, \boldsymbol{c}) + C(\boldsymbol{b}_h; \boldsymbol{u}_h, \boldsymbol{c}) - D(\boldsymbol{c}, r_h), \\ R_r(\boldsymbol{b}_h; s) &= -D(\boldsymbol{b}_h, s) \end{aligned} \tag{2.5}$$

for all $(\boldsymbol{v}, q, \boldsymbol{c}, s) \in \boldsymbol{V}_h \times Q_h \times \boldsymbol{C}_h \times S_h$. For small data, the iteration (2.4) converges for any initial guess [19].

2.3. Decoupling. Let us consider two important cases where simplifications to the Picard iteration (2.4) can be used. We introduce the following variants, referred to as *magnetic decoupling* and *complete decoupling*.

As mentioned in the introduction, [1] discusses the notion of strong coupling, according to the value of κ: cases where $\kappa < 100$ are considered to have weak coupling. Otherwise, the problem is treated as one with strong coupling. We have found this to be useful from a computational point of view, too. For $\kappa < 100$ we may converge to a solution by taking the coupling terms explicitly; i.e., we omit them in (2.4). Therefore, for a given solution $(\boldsymbol{u}_h, p_h, \boldsymbol{b}_h, r_h)$, neglecting the coupling terms in (2.4) results in solving for the updates $(\delta\boldsymbol{u}_h, \delta p_h, \delta\boldsymbol{b}_h, \delta r_h) \in \boldsymbol{V}_h \times Q_h \times \boldsymbol{C}_h \times S_h$ such that

$$\begin{aligned} A(\delta\boldsymbol{u}_h, \boldsymbol{v}) + O(\boldsymbol{u}; \delta\boldsymbol{u}_h, \boldsymbol{v}) + B(\boldsymbol{v}, \delta p_h) &= R_u(\boldsymbol{u}_h, \boldsymbol{b}_h, p_h; \boldsymbol{v}), \\ B(\delta\boldsymbol{u}_h, q) &= R_p(\boldsymbol{u}_h; q), \\ M(\delta\boldsymbol{b}_h, \boldsymbol{c}) + D(\boldsymbol{c}, \delta r_h) &= R_b(\boldsymbol{u}_h, \boldsymbol{b}_h, r_h; \boldsymbol{c}), \\ D(\delta\boldsymbol{b}_h, s) &= R_r(\boldsymbol{b}_h; s), \end{aligned} \tag{2.6}$$

where R_u, R_p, R_b, and R_r are as defined in (2.5). We call this *magnetic decoupling* (MD).

When we have both weak coupling and small convection terms in the system (2.4), the simplest strategy is to take all the nonlinear terms explicitly. This is the simplest technique, as it removes all nonlinear terms. For a given solution $(\boldsymbol{u}_h, p_h, \boldsymbol{b}_h, r_h)$, removing the coupling and convection terms in (2.4) results in solving for the updates $(\delta\boldsymbol{u}_h, \delta p_h, \delta\boldsymbol{b}_h, \delta r_h) \in \boldsymbol{V}_h \times Q_h \times \boldsymbol{C}_h \times S_h$ such that

$$\begin{aligned} A_h(\delta\boldsymbol{u}_h, \boldsymbol{v}) + B(\boldsymbol{v}, \delta p_h) &= R_u(\boldsymbol{u}_h, \boldsymbol{b}_h.p_h; \boldsymbol{v}), \\ B(\delta\boldsymbol{u}_h, q) &= R_p(\boldsymbol{u}_h; q), \\ M(\delta\boldsymbol{b}_h, \boldsymbol{c}) + D(\boldsymbol{c}, \delta r_h) &= R_b(\boldsymbol{u}_h, \boldsymbol{b}_h, r_h; \boldsymbol{c}), \\ D(\delta\boldsymbol{b}_h, s) &= R_r(\boldsymbol{b}_h; s), \end{aligned} \tag{2.7}$$

where again R_u, R_p, R_b, and R_r are given in (2.5). We call this *complete decoupling* (CD).

2.4. The linear systems. For the matrix representation of (2.4)–(2.5), we introduce the basis function for the finite element spaces in (2.2):

$$\begin{aligned} \boldsymbol{V}_h &= \operatorname{span}\langle\boldsymbol{\psi}_j\rangle_{j=1}^{n_u}, \qquad Q_h = \operatorname{span}\langle\alpha_i\rangle_{i=1}^{m_u}, \\ \boldsymbol{C}_h &= \operatorname{span}\langle\boldsymbol{\phi}_j\rangle_{j=1}^{n_b}, \qquad S_h = \operatorname{span}\langle\beta_i\rangle_{i=1}^{m_b}. \end{aligned}$$

The aim is to find the coefficient vectors $u = (u_1, \ldots, u_{n_u}) \in \mathbb{R}^{n_u}$, $p = (p_1, \ldots, p_{m_u}) \in \mathbb{R}^{m_u}$, $b = (b_1, \ldots, b_{n_b}) \in \mathbb{R}^{n_b}$, and $r = (r_1, \ldots, r_{m_b}) \in \mathbb{R}^{m_b}$ of the finite element functions $(\boldsymbol{u}_h, p_h, \boldsymbol{b}_h, r_h)$ in terms of the chosen bases. To this end, we define the following stiffness matrices and load vectors:

$$\begin{aligned} A_{i,j} &= A(\boldsymbol{\psi}_j, \boldsymbol{\psi}_i), && 1 \le i, j \le n_u, \\ B_{i,j} &= B(\boldsymbol{\psi}_j, \alpha_i), && 1 \le i \le m_u,\ 1 \le j \le n_u, \\ M_{i,j} &= M(\boldsymbol{\phi}_j, \boldsymbol{\phi}_i), && 1 \le i, j \le n_b, \\ D_{i,j} &= D(\boldsymbol{\phi}_j, \beta_i), && 1 \le i \le m_b,\ 1 \le j \le n_b, \\ f_i &= (\boldsymbol{f}, \boldsymbol{\psi}_i)_\Omega, && 1 \le i \le n_u, \\ g_i &= (\boldsymbol{g}, \boldsymbol{\phi}_i)_\Omega, && 1 \le i \le n_b. \end{aligned}$$

We define the stiffness matrices for the two nonlinear forms, O and C, with respect to the current finite element iterates $\boldsymbol{u}_h \in \boldsymbol{V}_h$ and $\boldsymbol{b}_h \in \boldsymbol{C}_h$ and their associated coefficient vectors u and b as

$$\begin{aligned} O(u)_{i,j} &= O(\boldsymbol{u}_h; \boldsymbol{\psi}_j, \boldsymbol{\psi}_i), && 1 \le i, j \le n_u, \\ C(b)_{i,j} &= C(\boldsymbol{b}_h; \boldsymbol{\psi}_j, \boldsymbol{\phi}_i), && 1 \le i \le n_b,\ 1 \le j \le n_u. \end{aligned}$$

We denote by (u, p, b, r) and $(\delta u, \delta p, \delta b, \delta r)$ the coefficient vectors associated with $(\boldsymbol{u}_h, p_h, \boldsymbol{b}_h, r_h)$ and $(\delta\boldsymbol{u}_h, \delta p_h, \delta\boldsymbol{b}_h, \delta r_h)$, respectively. Then it can be readily seen that the Picard iteration (2.4) amounts to solving the matrix system

$$\begin{pmatrix} A + O(u) & B^T & C(b)^T & 0 \\ B & 0 & 0 & 0 \\ -C(b) & 0 & M & D^T \\ 0 & 0 & D & 0 \end{pmatrix} \begin{pmatrix} \delta u \\ \delta p \\ \delta b \\ \delta r \end{pmatrix} = \begin{pmatrix} r_u \\ r_p \\ r_b \\ r_r \end{pmatrix}, \tag{2.8}$$

with

$$
\begin{aligned}
r_u &= f - Au - O(u)u - C(b)^T b - B^T p,\\
r_p &= -Bu,\\
r_b &= g - Mu + C(b)b - D^T r,\\
r_r &= -Db.
\end{aligned}
\tag{2.9}
$$

At each nonlinear iteration, the right-hand side vectors and matrices $O(u)$ and $C(b)$ in (2.8), (2.9) must be assembled with the solution coefficient vectors (u, p, b, r) of the current iterate. Here, the matrix A is symmetric positive definite, $O(u)$ is nonsymmetric, and $-C(b)$, $C(b)^T$ appear in a skew-symmetric fashion. We also note that M is symmetric positive semidefinite, with nullity m_b corresponding to the dimension of the scalar space of the discrete gradients; see [15]. In what follows, we shall often omit the dependence of $O(u)$ and $C(b)$ on u and b, respectively, and write O and C.

The linear system associated with the *magnetic decoupling* scheme in (2.6) then is

$$
\left(\begin{array}{cc|cc}
A+O & B^T & 0 & 0\\
B & 0 & 0 & 0\\
\hline
0 & 0 & M & D^T\\
0 & 0 & D & 0
\end{array}\right)
\begin{pmatrix} \delta u\\ \delta b\\ \delta p\\ \delta r \end{pmatrix}
=
\begin{pmatrix} r_u\\ r_b\\ r_p\\ r_r \end{pmatrix},
\tag{2.10}
$$

with the right-hand side quantities as defined in (2.9). While still nonsymmetric, the system decouples into a Navier–Stokes block and a Maxwell block, thus allowing the problems to be solved with well-known specifically designed preconditioners, and possibly in parallel.

The linear system connected with the *complete decoupling* scheme in (2.7) is

$$
\left(\begin{array}{cc|cc}
A & B^T & 0 & 0\\
B & 0 & 0 & 0\\
\hline
0 & 0 & M & D^T\\
0 & 0 & D & 0
\end{array}\right)
\begin{pmatrix} \delta u\\ \delta b\\ \delta p\\ \delta r \end{pmatrix}
=
\begin{pmatrix} r_u\\ r_b\\ r_p\\ r_r \end{pmatrix},
\tag{2.11}
$$

again with the right-hand side quantities as defined in (2.9). The system is now symmetric and decouples into a linear Stokes problem and a Maxwell problem. Then we may apply MINRES to both of the subsystems.

3. Review of preconditioning techniques for the subproblems. The linear systems in section 2.4 are associated with a few important subproblems, as discussed. These are the Navier–Stokes, Stokes, and Maxwell problems. In this section we review preconditioners for each of these subproblems.

3.1. Fluid flow preconditioner. For the *magnetic decoupling* scheme (2.10), the governing equations for the fluid flow are the Navier–Stokes equations. Their associated discretized and linearized operator is given by

$$
\mathcal{K}^{\mathrm{NS}} = \begin{pmatrix} F & B^T\\ B & 0 \end{pmatrix},
\tag{3.1}
$$

with $F = A + O$. One of the principal preconditioning approaches in the literature is based on approximations to the Schur complement. Using [4, 16], we look at preconditioners of the form

$$
\mathcal{M}_{\mathrm{I}}^{\mathrm{NS}} = \begin{pmatrix} F & B^T\\ 0 & -S \end{pmatrix},
\tag{3.2}
$$

where $S = BF^{-1}B^T$. In practice, the leading block, F, and the Schur complement, S, are approximated by linear operators that are easier to invert. Two common choices for an approximation to the Schur complement, S, are the least squares commutator and pressure-convection diffusion. We use the pressure-convection diffusion approximation developed in [4], which has proven to be robust with respect to viscosity, different choices of mixed finite elements, and type of mesh triangulation (i.e., squares or triangles in 2D). The approximation is based on

$$S = BF^{-1}B^T \approx A_p\, F_p^{-1} Q_p, \tag{3.3}$$

where the matrix A_p is the pressure Laplacian, F_p is the pressure convection-diffusion operator, and Q_p is the pressure mass matrix:

$$\begin{aligned}
(A_p)_{i,j} &= (\nabla\alpha_j, \nabla\alpha_i)_\Omega, && 1 \le i,j \le m_u,\\
(F_p)_{i,j} &= \nu(A_p)_{i,j} + (\boldsymbol{u}_h \cdot \nabla\alpha_j, \alpha_i)_\Omega, && 1 \le i,j \le m_u,\\
(Q_p)_{i,j} &= (\alpha_j, \alpha_i)_\Omega, && 1 \le i,j \le m_u,
\end{aligned}$$

where $\boldsymbol{u}_h \in \boldsymbol{V}_h$ is the given velocity field in the current iteration step. Note that A_p and F_p are well-defined since we work with continuous pressure elements. An effective implementation of this preconditioning approach is discussed in section 4.3.

For the *complete decoupling* scheme (2.11), the governing equations for the fluid flow are the Stokes equations. Their discrete form is given by the matrix

$$\mathcal{K}^{\mathrm{S}} = \begin{pmatrix} A & B^T \\ B & 0 \end{pmatrix}. \tag{3.4}$$

Here we opt to use a standard block diagonal (i.e., with B^T zero in (3.2)) and symmetric positive definite preconditioner of the form

$$\mathcal{M}_{\mathrm{P}}^{\mathrm{S}} = \begin{pmatrix} A & 0 \\ 0 & \widehat{S} \end{pmatrix}.$$

A natural choice for $\widehat{S}$ is $\widehat{S} = \frac{1}{\nu} Q_p$, where Q_p is the pressure mass matrix in (3.3). It is important here to note that the system is completely decoupled; hence, using a symmetric positive definite preconditioner allows use of MINRES and therefore short recurrences within the Krylov solver. This preconditioner allows one to obtain mesh-independent convergence rates; see [4].

3.2. Maxwell preconditioner. A key part of each decoupling scheme is an efficient preconditioner for the discrete Maxwell subsystem, whose associated matrix is given by

$$\mathcal{K}^{\mathrm{MX}} = \begin{pmatrix} M & D^T \\ D & 0 \end{pmatrix}. \tag{3.5}$$

In [8], it was shown that an ideal block diagonal positive definite preconditioner is

$$\mathcal{M}_{\mathrm{I}}^{\mathrm{MX}} = \begin{pmatrix} M + D^T L^{-1} D & 0 \\ 0 & L \end{pmatrix}. \tag{3.6}$$

Here L is the scalar Laplacian on S_h defined as

$$L_{i,j} = (\nabla\beta_j, \nabla\beta_i)_\Omega, \qquad 1 \le i,j \le m_b. \tag{3.7}$$

Using this preconditioner yields precisely two distinct eigenvalues, 1 and -1; hence a symmetric preconditioned Krylov solver such as MINRES will converge within two iterations in the absence of roundoff errors.

Remark 3.1. Given that the MHD problem is nonsymmetric, and hence we would have to use a nonsymmetric solver anyway, one may be tempted to ask whether it would make sense to incorporate the (1,2) block of the coefficient matrix (3.5) into the preconditioner, namely, replace (3.6) by

$$\widetilde{\mathcal{M}}_{\mathrm{I}}^{\mathrm{MX}} = \begin{pmatrix} M + D^T L^{-1} D & D^T \\ 0 & L \end{pmatrix}.$$

Interestingly, it turns out that there is no advantage in doing so in terms of eigenvalue distribution and, consequently, the convergence of the iterative solver. The preconditioned eigenvalues of $(\widetilde{\mathcal{M}}_{\mathrm{I}}^{\mathrm{MX}})^{-1}\mathcal{K}^{\mathrm{MX}}$ are 1 and $\frac{1\pm\sqrt{5}}{2}$; that is, there are three of them, whereas the block diagonal preconditioner (3.6) gives rise to two eigenvalues. Therefore, the additional computational cost entailed in a matrix-vector product with D^T does not translate into a benefit in terms of iteration counts.

Inverting the (1,1) block of the preconditioner,

$$M_L = M + D^T L^{-1} D, \tag{3.8}$$

is typically computationally prohibitive. Let X be the scalar mass matrix on $\boldsymbol{C}_h$, defined as

$$X_{i,j} = (\boldsymbol{\phi}_j, \boldsymbol{\phi}_i)_\Omega, \qquad 1 \le i, j \le n_b. \tag{3.9}$$

Then it has been shown in [8, Theorem 3.3] that M_L and $M + X$ are spectrally equivalent. Hence, we may use the preconditioner

$$\mathcal{M}_{\mathrm{P}}^{\mathrm{MX}} = \begin{pmatrix} M_X & 0 \\ 0 & L \end{pmatrix}, \qquad \text{where} \qquad M_X = M + X. \tag{3.10}$$

The preconditioner $\mathcal{M}_{\mathrm{P}}^{\mathrm{MX}}$ still has rather attractive spectral properties; in particular, the preconditioned operator $(\mathcal{M}_{\mathrm{P}}^{\mathrm{MX}})^{-1}\mathcal{K}^{\mathrm{MX}}$ has the two eigenvalues 1 and -1 with algebraic multiplicity m_b each, and the rest of the eigenvalues are bounded by a constant independent of the mesh size; cf. [8]. We thus will use $\mathcal{M}_{\mathrm{P}}^{\mathrm{MX}}$ as a preconditioner for the Maxwell subproblem. As discussed in section 4.3, we shall use certain approximations to M_X and L to achieve maximal scalability with respect to computing time and problem size.

4. Preconditioners for the MHD system. In this section, we propose a preconditioning approach for the discrete MHD system (2.8), which is based on keeping the coupling matrix C in the preconditioner, and on applying the preconditioners discussed in sections 3.1 and 3.2 to the Navier–Stokes and Maxwell subsystems. Leaving the coupling terms in and applying these known preconditioners to each of the subproblems in the MHD system yields the ideal preconditioner:

$$\mathcal{M}_{\mathrm{I}}^{\mathrm{MHD}} = \left(\begin{array}{cc|cc} F & B^T & C^T & 0 \\ 0 & -S & 0 & 0 \\ \hline -C & 0 & M_L & 0 \\ 0 & 0 & 0 & L \end{array}\right),$$

where M_L and S are defined as the Schur complement approximations in (3.8) and (3.3), respectively.

4.1. Reordering. We note that by reordering the blocks in $\mathcal{M}_{\mathrm{I}}^{\mathrm{MHD}}$, such that the solution vector is of the form (u, b, p, r), we obtain a 2×2 block triangular preconditioner of the form

$$\widetilde{\mathcal{M}}_{\mathrm{I}}^{\mathrm{MHD}} = \left(\begin{array}{cc|cc} F & C^T & B^T & 0 \\ -C & M_L & 0 & 0 \\ \hline 0 & 0 & -S & 0 \\ 0 & 0 & 0 & L \end{array}\right). \tag{4.1}$$

If we are to use this preconditioner, then we must reorder the linear system accordingly:

$$\left(\begin{array}{cc|cc} F & C^T & B^T & 0 \\ -C & M & 0 & D^T \\ \hline B & 0 & 0 & 0 \\ 0 & D & 0 & 0 \end{array}\right) \begin{pmatrix} \delta u \\ \delta b \\ \delta p \\ \delta r \end{pmatrix} = \begin{pmatrix} r_u \\ r_b \\ r_p \\ r_r \end{pmatrix}, \tag{4.2}$$

with the right-hand side quantities as defined in (2.9). Let us denote by $\mathcal{K}^{\mathrm{MHD}}$ the coefficient matrix defined in (4.2). From this point on, we consider the reordered system.

The computational bottleneck is solving systems associated with the matrix

$$\begin{pmatrix} F & C^T \\ -C & M_L \end{pmatrix} \tag{4.3}$$

in the $(1,1)$ block matrix of (4.1). To invert the matrix in (4.3), we apply a block triangular preconditioner based on the Schur complement given by

$$\begin{pmatrix} F + M_C & C^T \\ 0 & M_L \end{pmatrix}, \qquad \text{where} \qquad M_C = C^T M_L^{-1} C.$$

Using the above approximation in $\widetilde{\mathcal{M}}_{\mathrm{I}}^{\mathrm{MHD}}$ yields

$$\mathcal{M}_{\mathrm{S}}^{\mathrm{MHD}} = \begin{pmatrix} F + M_C & C^T & B^T & 0 \\ 0 & M_L & 0 & 0 \\ 0 & 0 & -S & 0 \\ 0 & 0 & 0 & L \end{pmatrix}. \tag{4.4}$$

To analyze the spectral properties of $\mathcal{M}_{\mathrm{S}}^{\mathrm{MHD}}$, we refer to vectors $b \in \mathrm{null}(M)$ as discrete gradients. With the discrete Helmholtz decomposition, it follows that for each $b \in \mathrm{null}(M)$ there is a unique $r \in \mathbb{R}^{m_b}$ such that $b = Gr$ for a discrete gradient matrix $G \in \mathbb{R}^{n_b \times m_b}$; cf. [8, section 2]. Hence, $\dim(\mathrm{null}(M)) = m_b$. The following result holds true.

THEOREM 4.1. *The matrix $(\mathcal{M}_{\mathrm{S}}^{\mathrm{MHD}})^{-1}\mathcal{K}^{\mathrm{MHD}}$ has an eigenvalue $\lambda = 1$ with algebraic multiplicity of at least $n_b + n_c$, where n_c is the dimension of the nullspace of $C = C(b)$ and an eigenvalue $\lambda = -1$ with algebraic multiplicity of at least m_b. The dimension of the nullspace of C is $n_c = n_u - n_b + m_b$. The corresponding eigenvalue-eigenvector $(\lambda, \boldsymbol{x})$ pairs are*

$$\lambda = 1, \quad \boldsymbol{x}^T = (u_c^T, b^T, (-S^{-1}Bu_c)^T, (L^{-1}Db)^T),$$

with $u_c \in \mathrm{null}(C)$ and $b \in \mathbb{R}^{n_b}$ arbitrary, and

$$\lambda = -1, \quad \boldsymbol{x}^T = (0, b_g^T, 0, (-L^{-1}Db_g)^T),$$

with $b_g = Gr$ a discrete gradient for $r \in \mathbb{R}^{m_b}$.

Proof. The corresponding eigenvalue problem is

$$\begin{pmatrix} F & C^T & B^T & 0 \\ -C & M & 0 & D^T \\ B & 0 & 0 & 0 \\ 0 & D & 0 & 0 \end{pmatrix} \begin{pmatrix} u \\ b \\ p \\ r \end{pmatrix} = \lambda \begin{pmatrix} F + M_C & C^T & B^T & 0 \\ 0 & M_L & 0 & 0 \\ 0 & 0 & -S & 0 \\ 0 & 0 & 0 & L \end{pmatrix} \begin{pmatrix} u \\ b \\ p \\ r \end{pmatrix}.$$

The four block rows of the generalized eigenvalue problem can be written as

$$(1-\lambda)(Fu + B^T p + C^T b) - \lambda C^T (M + D^T L^{-1} D)^{-1} C u = 0, \tag{4.5}$$

$$-Cu + (1-\lambda) Mb - \lambda D^T L^{-1} D b + D^T r = 0, \tag{4.6}$$

$$Bu = -\lambda S\, p, \tag{4.7}$$

$$Db = \lambda L r. \tag{4.8}$$

If $\lambda = 1$, (4.5) is satisfied if

$$C^T (M + D^T L^{-1} D)^{-1} C u = 0.$$

This only happens when $u \in \mathrm{Null}(C)$. Using u_c to denote a nullspace vector of C, we can then simplify (4.7) to

$$p = -S^{-1} B u_c.$$

Equation (4.8) leads to $r = L^{-1} D b$. If this holds, (4.6) is readily satisfied. Therefore, $(u_c^T, b^T, (-S^{-1} B u_c)^T, (L^{-1} D b)^T)$ is an eigenvector corresponding to $\lambda = 1$. There exist n_c linearly independent such vectors u and n_b linearly independent such vectors b. Hence, it follows that $\lambda = 1$ is an eigenvalue with algebraic multiplicity of at least $n_u + m_b$.

If $\lambda = -1$, (4.8) leads to $r = -L^{-1} D b$. Substituting it into (4.6), we obtain $Cu = Mb$. If $b = b_g$ is a discrete gradient, then $Mb = 0$ and $C^T b = 0$. If we take $u = 0$, then $Cu = 0$ and the requirement $Cu = Mb$ is satisfied. If $u = 0$ and $b = b_G$ is a discrete gradient, (4.5) becomes $B^T p = 0$. Since B has full row rank, this implies $p = 0$. Therefore, if $b = b_g$ is a discrete gradient, then $(0, b_g^T, 0, (-L^{-1} D b_g)^T)$ is an eigenvector corresponding to $\lambda = -1$. There are m_b discrete gradients. Therefore $\lambda = -1$ is an eigenvalue with algebraic multiplicity at least m_b. ∎

Remark 4.2. In the spirit of Remark 3.1, looking at (4.4) one may ask whether incorporating D^T into the preconditioner, such that

$$\widetilde{\mathcal{M}}_{\mathrm{S}}^{\mathrm{MHD}} = \begin{pmatrix} F + M_C & C^T & B^T & 0 \\ 0 & M_L & 0 & D^T \\ 0 & 0 & -S & 0 \\ 0 & 0 & 0 & L \end{pmatrix},$$

may generate a slightly better eigenvalue distribution for the preconditioned system. Consistently with Remark 3.1, it turns out that doing so does not practically generate an improvement of the eigenvalue distribution. From Table 4.1, we see that both preconditioned systems yield exactly the same number of eigenvalues. However, $(\mathcal{M}_{\mathrm{S}}^{\mathrm{MHD}})^{-1} \mathcal{K}^{\mathrm{MHD}}$ only has 2 distinct eigenvalues, whereas $(\widetilde{\mathcal{M}}_{\mathrm{S}}^{\mathrm{MHD}})^{-1} \mathcal{K}^{\mathrm{MHD}}$ has 3 distinct eigenvalues. Thus, the insertion of D^T does not theoretically decrease the number of iterations for a Krylov subspace method to converge; this has been confirmed by numerical experiments. Since incorporating D^T into the preconditioner slightly increases the cost of a single iteration, we opt to use $\mathcal{M}_{\mathrm{S}}^{\mathrm{MHD}}$ rather than $\widetilde{\mathcal{M}}_{\mathrm{S}}^{\mathrm{MHD}}$.

TABLE 4.1
Algebraic multiplicity of eigenvalues for preconditioned matrices associated with $\mathcal{M}_S^{\mathrm{MHD}}$ and $\widetilde{\mathcal{M}}_S^{\mathrm{MHD}}$. Note that $n_c = n_u - n_b + m_b$.

	$(\mathcal{M}_S^{\mathrm{MHD}})^{-1}\mathcal{K}^{\mathrm{MHD}}$	$(\widetilde{\mathcal{M}}_S^{\mathrm{MHD}})^{-1}\mathcal{K}^{\mathrm{MHD}}$
1	$n_b + n_c$	n_u
-1	m_b	0
$\frac{1+\sqrt{5}}{2}$	0	m_b
$\frac{1-\sqrt{5}}{2}$	0	m_b
Total	$n_u + 2m_b$	$n_u + 2m_b$

TABLE 4.2
Orders of matrix entries for relevant discrete operators.

Discrete operator	Order
G	$\mathcal{O}(h^{-1})$
L	$\mathcal{O}(h^{-2})$
D	$\mathcal{O}(h^{-1})$

4.2. From an ideal to a practical preconditioner. We now consider further simplifications of $\mathcal{M}_S^{\mathrm{MHD}}$ in (4.4) to make the preconditioner computationally feasible. Effective sparse approximations are required for the relevant Schur complements that arise. We use the approximations in sections 3.1 and 3.2 for S and M_L. For approximating M_C, we follow a similar approach to that taken in [17, section 3.1]. For a given magnetic field $\boldsymbol{b}$, let $\mathcal{C}_{\boldsymbol{b}}$ be the continuous differential operator analogue of $M_C = C^T M_L^{-1} C$:

$$\begin{aligned}\mathcal{C}_{\boldsymbol{b}}\boldsymbol{u} &= \kappa\Big(\nabla\times\big((\kappa\nu_m\nabla\times\nabla\times+\nabla\Delta^{-1}\nabla\cdot)^{-1}\kappa\nabla\times(\boldsymbol{u}\times\boldsymbol{b})\big)\Big)\times\boldsymbol{b}\\ &= \kappa^2\nabla\times\big(\kappa\nu_m\nabla\times\nabla\times+\nabla\Delta^{-1}\nabla\ \cdot\big)^{-1}\nabla\times(\boldsymbol{u}\times\boldsymbol{b})\times\boldsymbol{b}.\end{aligned} \tag{4.9}$$

The discretization of (4.9) is

$$\mathcal{C}_b u = \kappa^2 G\big(\kappa\nu_m G^T G + D^T L^{-1} D\big)^{-1} G^T (u\times b)\times b, \tag{4.10}$$

where G is a discrete curl matrix and u and b are vectors of velocity and magnetic coefficients, respectively.

In Table 4.2 we state the order of the discrete differential operators that are involved in (4.10). We observe that the discrete curl-curl matrix, $G^T G$, contains entries of magnitude $\mathcal{O}(h^{-2})$, and $D^T L^{-1} D$ is order $\mathcal{O}(1)$. Thus for small h and moderate values of κ and ν_m the curl-curl matrix will be the dominant term. Therefore we consider

$$\mathcal{C}_b u \approx \kappa\nu_m^{-1} G\big(G^T G\big)^{-1} G^T (u\times b)\times b.$$

Furthermore, $G\big(G^T G\big)^{-1} G^T$ is an orthogonal projector onto the range space of G^T; hence, it acts as an identity operator within that space. We therefore use the approximation

$$\mathcal{C}_b u \approx \kappa{\nu_m}^{-1} b\times(u\times b). \tag{4.11}$$

From (4.11), $\mathcal{C}_b$ can be approximated by a scaled mass matrix determined by the coefficients of the magnetic field b. We thus approximate M_C by a scaled mass matrix,

TABLE 4.3
Orders of matrix entries for relevant discrete operators.

Matrix	Implementation method
Q_p	diagonal of αQ_p, where $\alpha = 0.75$
A_p	single AMG V-cycle
$F + Q_s$	single AMG V-cycle
$M + X$	AMG method developed in [10]
L	single AMG V-cycle

which we denote by Q_S, and whose elements are

$$(Q_S)_{i,j} = \kappa\nu_m^{-1}(\boldsymbol{b}_h \times \boldsymbol{\psi}_j, \boldsymbol{b}_h \times \boldsymbol{\psi}_i)_\Omega, \qquad 1 \le i, j \le n_u. \tag{4.12}$$

Combining the sparse Schur complement approximations in (3.3), (3.10), and (4.12) in the MHD preconditioner (4.4) gives the approximate preconditioner

$$\mathcal{M}_{\mathrm{P}}^{\mathrm{MHD}} = \begin{pmatrix} F + Q_S & C^T & B^T & 0 \\ 0 & M + X & 0 & 0 \\ 0 & 0 & -A_p\, F_p^{-1} Q_p & 0 \\ 0 & 0 & 0 & L \end{pmatrix}. \tag{4.13}$$

In the transition from the ideal preconditioner (4.1) to the practical preconditioner (4.13), some spectral clustering is inevitably lost. But, as we show, the spectral structure of the preconditioned matrix associated with $\mathcal{M}_{\mathrm{P}}^{\mathrm{MHD}}$ is still very appealing. We illustrate this in section 5.

4.3. Implementation. So far we have introduced the matrix systems with possible preconditioners, but have not discussed practical implementation considerations. One of our main goals is to provide a fully scalable solution method. To this end, we will consider mesh-independent solvers for the separate block matrices within the preconditioners. Table 4.3 outlines the methods we use to solve the systems associated with Q_p, A_p, $F + Q_s$, $M + X$, and L, which are the block diagonal matrices in (4.13). Let us provide a few additional comments:

1. For solving systems involving Q_p, we have experimentally observed that scaling by a multiplicative scalar α smaller than 1 results in a slight decrease of iteration counts, especially in the 3D case. In our numerical experiments we provide results that correspond to using $\alpha = 0.75$.
2. For solving systems involving $M + X$, the method developed in [10] aims to overcome issues with standard algebraic multigrid (AMG) methods for the discrete curl-curl operator by using an auxiliary space approach [22] for $H(\mathrm{curl})$ finite element discretizations of elliptic problems. The construction of the auxiliary space multigrid preconditioner relies on three additional matrices: the discrete Nédélec interpolation operator $P \in \mathbb{R}^{n_b \times d\, m_b}$ (where d is the spatial dimension), the discrete gradient operator $G \in \mathbb{R}^{n_b \times m_b}$, and mass matrix defined on the scalar space S_h as Q; see [11, 13].

5. Numerical results. This section examines the efficiency of our preconditioning approaches to the MHD model (1.1)–(1.2). In all experiments, unless otherwise stated, we use a 2-norm absolute tolerance of 1e-4 for the nonlinear solver and relative error of 1e-5 for both the Krylov subspace solver and the auxiliary space multigrid [10].

All numerical experiments have been carried out using the finite element software `FEniCS` [14] in conjunction with the `PETSc4PY` package (Python interface for `PETSc` [2, 3]) and the multigrid package `HYPRE` [5].

We test our methods on problems with inhomogeneous Dirichlet boundary conditions in the hydrodynamic variables, even though the analysis has been carried out solely for the homogeneous Dirichlet case. Other boundary conditions may be handled by our finite element framework, and our preconditioning approaches can be extended accordingly.

In the subsequent tables we use the following notation: $\text{time}_{\text{solve}}$ is the average linear system solve time, time_{NL} is the total nonlinear solve time, it_{NL} is the number of nonlinear iterations, $\text{it}_{\text{av}}^{\text{I}}$ is the average number of FGMRES iterations (applying iterative solution of the preconditioned linear systems), and $\text{it}_{\text{av}}^{\text{D}}$ is the average number of FGMRES iterations when direct methods are used to solve the preconditioned linear systems.

Mesh sequences. In our tests, we consider sequences of uniformly refined simplicial grids (i.e., triangles in 2D and tetrahedra in 3D). We define ℓ to be the mesh level, such that there are 2^ℓ edges along each boundary. In our tables, we usually show the grid level ℓ in the first column. We also note that DoF (in the second column) refers to the total number of degrees of freedom of the linear system.

Stopping criteria. Throughout, we enforce the following nonlinear stopping criteria for the updates:

$$\|\delta u\|_2 + \|\delta p\|_2 + \|\delta b\|_2 + \|\delta r\|_2 < \text{tol}_{\text{NL}},$$

where $\|\cdot\|_2$ is the absolute error in the 2-norm of a vector and $\text{tol}_{\text{NL}} = $ 1e-4.

Initial guess tolerance. For all numerical experiments, we formulate the initial guess by iteratively solving a Stokes problem and mixed Maxwell problem in isolation. We choose a tight Krylov 2-norm relative tolerance as 1e-10 to ensure that the approximations of the inhomogeneous boundary conditions are sufficiently accurate; see [20, section 2.5]. This is to ensure the accuracy of the initial solution on the boundaries since subsequent Picard iterations are solved for homogeneous boundary conditions, and hence any errors in the initial guess will be carried throughout the Picard iterations.

5.1. 2D smooth solution. The first example considered is a simple domain with a structured mesh. We use a unit square domain, $\Omega = [0,1]^2$. We take $\nu = \kappa = 1$, $\nu_m = 10$; then the source terms $\boldsymbol{f}$, $\boldsymbol{g}$ and inhomogeneous Dirichlet boundary conditions on $\partial\Omega$ are defined from the following analytical solution:

$$\begin{aligned}
\boldsymbol{u}(x,y) &= \begin{pmatrix} xy\exp(x+y) + x\exp(x+y) \\ -xy\exp(x+y) - y\exp(x+y) \end{pmatrix}, \\
p(x,y) &= \exp(y)\sin(x), \\
\boldsymbol{b}(x,y) &= \begin{pmatrix} \exp(x+y)\cos(x) \\ \exp(x+y)\sin(x) - \exp(x+y)\cos(x) \end{pmatrix}, \\
r(x,y) &= x\sin(2\pi x)\sin(2\pi y).
\end{aligned}$$

To illustrate the eigenvalue distribution, we compute the eigenvalues of the preconditioned matrix $\left(\widetilde{\mathcal{M}}_{\text{I}}^{\text{MHD}}\right)^{-1}\mathcal{K}^{\text{MHD}}$ in Figure 5.1(a) and $(\mathcal{M}_{\text{P}}^{\text{MHD}})^{-1}\mathcal{K}^{\text{MHD}}$ in Figure 5.1(b). We note that there are no imaginary parts of the eigenvalues for Figure 5.1(a). From Figure 5.1(b) we see that we still observe strong clustering of eigenvalues around 1 and -1 with only a few complex conjugate pairs. Thus, the spectral structure is still rather appealing in terms of eigenvalue clustering. Therefore, the preconditioner $\mathcal{M}_{\text{P}}^{\text{MHD}}$ seems a viable approximation to $\widetilde{\mathcal{M}}_{\text{I}}^{\text{MHD}}$.

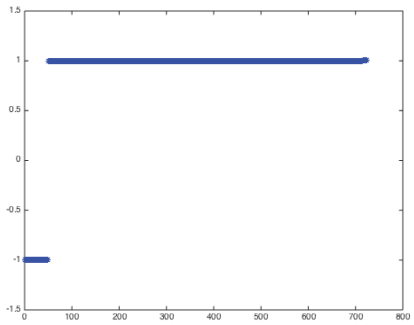

(a) Eigenvalues of the preconditioned matrix $\left(\widetilde{\mathcal{M}}_{\mathrm{I}}^{\mathrm{MHD}}\right)^{-1}\mathcal{K}_{\mathrm{P}}^{\mathrm{MH}}$ associated with the ideal preconditioner.

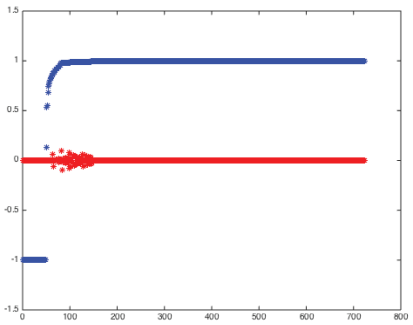

(b) Eigenvalues of the preconditioned matrix $(\mathcal{M}_{\mathrm{P}}^{\mathrm{MHD}})^{-1}\mathcal{K}_{\mathrm{P}}^{\mathrm{MH}}$ associated with the practical preconditioner. The eigenvalues in this case are complex; the blue curves represent their real parts, and the red curves represent their imaginary parts.

FIG. 5.1. *Eigenvalues of preconditioned matrices for the smooth solution given in this section. The number of degrees of freedom for these matrices is* 724.

TABLE 5.1

2D smooth: Number of nonlinear iterations and number of iterations to solve the MHD system with $\mathrm{tol}_{\mathrm{NL}} = 1e\text{-}4$, $\kappa = 1$, $\nu = 1$, *and* $\nu_m = 10$.

ℓ	DoF	$\mathrm{time}_{\mathrm{solve}}$	$\mathrm{time}_{\mathrm{NL}}$	$\mathrm{it}_{\mathrm{NL}}$	$\mathrm{it}_{\mathrm{av}}^{\mathrm{I}}$	$\mathrm{it}_{\mathrm{av}}^{D}$
4	3,556	0.33	2.7	7	24.4	20.1
5	13,764	1.11	9.2	7	25.9	20.4
6	54,148	4.48	37.2	7	27.1	20.9
7	214,788	20.32	166.4	7	28.4	21.4
8	855,556	94.29	762.0	7	31.3	21.8
9	3,415,044	486.53	3835.0	7	34.3	-
10	13,645,828	2231.71	17944.6	7	34.0	-

To test the scalability of our method, we considered a uniformly refined sequence of meshes. The results are presented in Table 5.1. The iterations appear to remain fairly constant with the increasing mesh level for both $\mathrm{it}_{\mathrm{av}}^{D}$ and $\mathrm{it}_{\mathrm{av}}^{\mathrm{I}}$.

5.2. 2D smooth solution parameter tests. We next test the performance of the three nonlinear iteration schemes *Picard* (P), *magnetic decoupling* (MD), and *complete decoupling* (CD) introduced in section 2. The convergence of the nonlinear iterations is likely to be affected by the parameter setup of the problem, i.e., by the values of the fluid viscosity (ν), the magnetic viscosity (ν_m), and the coupling number (κ). By varying κ and ν, we examine the robustness of the three schemes with respect to these parameters. Nonlinear iterations that do not converge are denoted by "-" in the tables.

Viscosity test. As a first test we consider varying the fluid viscosity, ν, for $\mathrm{tol}_{\mathrm{NL}} = 1e\text{-}4$, $\kappa = 1$, and $\nu_m = 10$. The nonlinear iteration results are shown in Table 5.2, and the average linear solve times are shown in Table 5.3.

As the fluid viscosity (ν) decreases, the fluid flow equations (1.1a) and (1.1b) become more convection-dominated. Thus we see that the (CD) scheme breaks down for small ν, since the convection term in this decoupling scheme is taken explicitly.

TABLE 5.2
Number of nonlinear iterations for various values of ν with $\kappa = 1$ and $\nu_m = 10$.

		$\nu = 1$			$\nu = 0.1$			$\nu = 0.01$		
ℓ	DoF	(P)	(MD)	(CD)	(P)	(MD)	(CD)	(P)	(MD)	(CD)
4	3,556	5	5	9	7	7	-	11	11	-
5	13,764	5	5	9	15	7	-	11	11	-
6	54,148	5	5	9	7	7	-	13	11	-
7	214,788	5	5	9	7	7	-	11	11	-
8	855,556	5	5	9	7	7	-	11	11	-

TABLE 5.3
Average linear solver time for various values of ν with $\kappa = 1$ and $\nu_m = 10$.

		$\nu = 1$			$\nu = 0.1$			$\nu = 0.01$		
ℓ	DoF	(P)	(MD)	(CD)	(P)	(MD)	(CD)	(P)	(MD)	(CD)
4	3,556	0.04	0.04	0.03	0.08	0.04	-	0.04	0.04	-
5	13,764	0.18	0.17	0.14	0.97	0.18	-	0.19	0.18	-
6	54,148	0.90	0.91	0.67	0.95	0.92	-	2.05	0.93	-
7	214,788	4.48	4.48	3.49	4.68	4.57	-	5.02	4.48	-
8	855,556	25.52	22.81	16.67	25.63	23.03	-	25.01	22.61	-

TABLE 5.4
Number of nonlinear iterations for various values of κ with $\mathrm{tol}_{\mathrm{NL}} = 1e\text{-}5$, $\nu = 1$, and $\nu_m = 10$.

		$\kappa = 1$			$\kappa = 10$			$\kappa = 100$		
ℓ	DoF	(P)	(MD)	(CD)	(P)	(MD)	(CD)	(P)	(MD)	(CD)
4	3,556	5	5	9	7	10	17	8	-	-
5	13,764	5	5	9	7	10	18	8	-	-
6	54,148	5	5	9	7	10	18	8	-	-
7	214,788	5	5	9	7	10	18	8	-	-
8	855,556	5	5	9	7	10	18	8	-	-

TABLE 5.5
Average linear solver time for various values of κ with $\nu = 1$ and $\nu_m = 10$.

		$\kappa = 1$			$\kappa = 10$			$\kappa = 100$		
ℓ	DoF	(P)	(MD)	(CD)	(P)	(MD)	(CD)	(P)	(MD)	(CD)
4	3,556	0.06	0.06	0.04	0.06	0.06	0.03	0.06	-	-
5	13,764	0.23	0.23	0.21	0.24	0.25	0.19	0.24	-	-
6	54,148	1.12	1.08	0.80	1.01	1.05	0.79	0.98	-	-
7	214,788	4.80	4.92	3.65	5.05	4.92	3.74	5.52	-	-
8	855,556	25.21	25.93	18.03	25.49	26.03	18.49	26.75	-	-

On the other hand, the Picard and (MD) schemes perform similarly. We note that for smaller ν both (P) and (MD) have trouble converging without a sufficiently refined mesh. From Table 5.3 we see that the direct solve for the Picard (P) system and magnetic decoupling are very similar, but as expected the complete decoupling solve time is shorter.

Coupling number test. The next parameter test examines the effects of the coupling terms in the three nonlinear iteration schemes. We expect the Picard scheme to outperform the other schemes for large values of κ. The results in Table 5.4 show that this is indeed the case. Both the (MD) and (CD) schemes completely break down for $\kappa \geq 100$. This is the point at which the Picard iteration (P) becomes the most viable option. Altogether, as expected, the full Picard iteration is more robust than (CD) and (MD). The timing results in Table 5.5 show a trend similar to the timing results for the viscosity in Table 5.3.

TABLE 5.6
2D L-shaped domain: Number of nonlinear iterations and number of iterations to solve the MHD system with $\mathrm{tol}_{\mathrm{NL}} = 1e\text{-}4$, $\kappa = 1$, $\nu = 1$, *and* $\nu_m = 10$. *The iteration was terminated before completion for* $\ell = 9$ *due to the computation reaching the prescribed time limit.*

ℓ	DoF	$\mathrm{it}_{\mathrm{NL}}$	$\mathrm{it}_{\mathrm{av}}^{D}$
5	12,880	5	24.4
6	51,678	5	26.0
7	203,712	5	27.4
8	809,705	5	29.6
9	3,219,082	-	-

5.3. 2D smooth solution on L-shaped domain. To further test the robustness of our preconditioning technique, we will consider nonconvex domains. We first consider a problem with a smooth solution in the L-shaped domain $\Omega = (-1,-1)^2 \setminus ([0,0) \times (1,1])$. We prescribe the same analytical solution as in section 5.1. The results in Table 5.6 show very good scalability with respect to mesh refinement for direct solves of the preconditioner. On the other hand, we report that when using iterative inner solvers with a reasonably loose tolerance, the iterations dramatically deteriorate and scalability is lost. We speculate that the AMG solver cannot handle well the nonconvexity of the domain. In such cases it may be necessary to apply a rather strict convergence tolerance.

5.4. 2D singular solution on L-shaped domain. We next consider the model singular solution from [7] on the L-shaped domain $\Omega = (-1,1)^2 \setminus ([0,1) \times (-1,0])$. That is, taking $\nu = \kappa = 1$ and $\nu_m = 10$, we set the forcing terms and the boundary conditions such that the analytic solution is given by the strongest corner singularities of the underlying Stokes and Maxwell operators subject to the boundary conditions (1.2) on the two inner sides meeting at the reentrant corner. In polar coordinates (ρ, ϕ) at the origin, the fluid solution $(\boldsymbol{u}, p)$ is then of the form

$$\boldsymbol{u}(\rho,\phi) = \rho^{\lambda} \boldsymbol{\Psi}_u(\phi), \qquad p(\rho,\phi) = \rho^{\lambda-1} \Psi_p(\phi), \tag{5.1}$$

with the singular exponent $\lambda \approx 0.54448373678246$ and where $(\boldsymbol{\Psi}_u, \Psi_p)$ are smooth functions in the angle ϕ. Similarly, the magnetic pair $(\boldsymbol{b}, r)$ (with $\nabla \cdot \boldsymbol{b} = 0$ and $\nabla \times \boldsymbol{b} = 0$) is of the form

$$\boldsymbol{b}(\rho,\phi) = \rho^{-1/3} \boldsymbol{\Psi}_b(\phi), \qquad r \equiv 0. \tag{5.2}$$

Detailed expressions can be found in [7, section 5.2]. Using Nédélec elements allows us to properly capture this singular solution, whereas applying standard nodal elements for $\boldsymbol{b}$ fails to do so. As with the smooth solution on an L-shaped domain in section 5.3, we only consider direct applications of the preconditioner. The results are shown in Table 5.7. We draw similar conclusions with respect to mesh independence. Overall, the iterations show good scalability.

5.5. 2D Hartmann flow. As a final 2D numerical example, we consider the 2D Hartmann flow problem, which involves a steady unidirectional flow in the channel $\Omega = (0,10) \times (-1,1)$ under the constant transverse magnetic field $\boldsymbol{b}_D = (0,1)$ on $\partial\Omega$. We impose the analytical solution given in [7, section 5.3] with Dirichlet boundary conditions for $\boldsymbol{u}$ on the entire boundary $\partial\Omega$. The MHD solution then takes the form

$$\begin{aligned} &\boldsymbol{u}(x,y) = (u(y), 0), \quad p(x,y) = -Gx + p_0(y), \\ &\boldsymbol{b}(x,y) = (b(y), 1), \quad r(x,y) \equiv 0. \end{aligned} \tag{5.3}$$

TABLE 5.7
2D singular solution on L-shaped domain: Number of nonlinear iterations and number of iterations to solve the MHD system with $\text{tol}_{\text{NL}} = 1e\text{-}4$, $\kappa = 1$, $\nu = 1$, *and* $\nu_m = 10$.

ℓ	DoF	it_{NL}	it_{av}^D
3	740	4	13.8
4	2,724	4	14.5
5	10,436	4	15.8
6	40,836	4	17.5
7	161,540	4	18.5
8	642,564	4	20.0
9	2,563,076	4	21.8

TABLE 5.8
2D Hartmann flow: Number of nonlinear iterations and number of iterations to solve the MHD system with $\text{tol}_{\text{NL}} = 1e\text{-}4$, $\kappa = 1$, $\nu = 1$, *and* $\nu_m = 1000$.

ℓ	DoF	$\text{time}_{\text{solve}}$	time_{NL}	it_{NL}	it_{av}^I	it_{av}^D
2	1,212	0.66	1.58	2	18.0	13.5
3	4,500	1.71	3.82	2	17.5	13.0
4	17,316	5.13	11.04	2	17.5	12.5
5	67,908	21.06	44.73	2	18.5	12.5
6	268,932	97.16	204.36	2	19.0	13.0
7	1,070,340	447.66	935.03	2	19.0	12.5
8	4,270,596	921.90	1001.37	1	8.0	7.0
9	17,060,868	1459.82	1778.95	1	3.0	-

The exact solution is given by (5.3) with

$$u(y) = \frac{G}{\nu \text{Ha} \tanh(\text{Ha})} \left(1 - \frac{\cosh(y\text{Ha})}{\cosh(\text{Ha})}\right),$$
$$b(y) = \frac{G}{\kappa} \left(\frac{\sinh(y\text{Ha})}{\sinh(\text{Ha})} - y\right),$$
$$p_0(y) = -\frac{G^2}{2\kappa} \left(\frac{\sinh(y\text{Ha})}{\sinh(\text{Ha})} - y\right)^2,$$

where $\text{Ha} = \sqrt{\frac{\kappa}{\nu\nu_m}}$ is the Hartmann number. We impose inhomogeneous Dirichlet boundary conditions from the exact solutions.

The results are reported in Table 5.8. We observe that for the 2D Hartmann flow example the solver appears to accomplish an excellent degree of scalability. For the last two mesh levels the preconditioner approximation becomes better in terms of lower iterations, it_{av}^I. More exploration of the conditioning of the problem and the norm of the (preconditioned) residual is needed to fully understand the reason for this; the large size of the problems presents a challenge in fully exploring this.

5.6. 3D smooth solution. We next consider a 3D example with a smooth solution on $\Omega = [0,1]^3$. Let $\nu = \kappa = 1$, $\nu_m = 10$ and let the analytical solution be

TABLE 5.9
3D smooth solution: Number of nonlinear iterations and number of iterations to solve the MHD system with $\mathrm{tol}_{\mathrm{NL}} = 1e\text{-}4$, $\kappa = 1$, $\nu = 1$, *and* $\nu_m = 10$.

ℓ	DoF	$\mathrm{time}_{\mathrm{solve}}$	$\mathrm{time}_{\mathrm{NL}}$	$\mathrm{it}_{\mathrm{NL}}$	$\mathrm{it}_{\mathrm{av}}^{\mathrm{I}}$	$\mathrm{it}_{\mathrm{av}}^{D}$
1	527	0.03	0.9	4	18.8	18.0
2	3,041	0.22	3.5	3	26.7	22.3
3	20,381	1.77	26.6	3	37.0	24.7
4	148,661	22.11	237.0	3	40.7	26.0
5	1,134,437	206.43	2032.7	3	44.3	-
6	8,861,381	2274.28	19662.0	3	50.0	-

given by

$$\boldsymbol{u}(x,y,z) = \begin{pmatrix} -xy\exp(x+y+z) + xz\exp(x+y+z) \\ xy\exp(x+y+z) - yz\exp(x+y+z) \\ -xz\exp(x+y+z) + yz\exp(x+y+z) \end{pmatrix},$$

$$p(x,y,z) = \exp(x+y+z)\sin(y),$$

$$\boldsymbol{b}(x,y,z) = \begin{pmatrix} -\exp(x+y+z)\sin(y) + \exp(x+y+z)\sin(z) \\ xy\exp(x+y+z) - yz\exp(x+y+z) \\ -\exp(x+y+z)\sin(x) + \exp(x+y+z)\sin(y) \end{pmatrix},$$

$$r(x,y,z) = \sin(2\pi x)\sin(2\pi y)\sin(2\pi z).$$

Then the source terms $\boldsymbol{f}$ and $\boldsymbol{g}$ and inhomogeneous boundary conditions are defined from the analytical solution. The corresponding results are shown in Table 5.9. We observe good scalability when we consider direct solves for the preconditioner. However, the average iterations degrade when using multigrid methods to solve the preconditioner.

6. Conclusions. We have introduced a block-structured preconditioning technique for a linearized steady-state incompressible MHD model, discretized by Taylor–Hood elements for the fluid variables and by the lowest-order Nédélec pair for the magnetic unknowns. The aim was to develop a preconditioner that would yield a scalable solution method. To this end, we have derived effective Schur complement sparse approximations. Spectral analysis demonstrates good eigenvalue clustering properties.

We have generated a large-scale code base, utilizing both the finite element software package `FEniCS` [14] and the linear algebra software `PETSc` [3, 2] to test our proposed preconditioner. Results have been presented for problems in the excess of 10 million degrees of freedom.

Our 2D results are scalable. For 3D problems we observe weak scalability in some of the settings; that is, the number of iterations does not quite stay constant as the mesh is refined. Resolving this may require resorting to other finite element formulations, or turning to more robust inner solvers.

Further developments may include the derivation of other nonlinear solvers that have faster convergence properties than the Picard iterations we considered in this paper. One such solver is Newton's method, which converges quadratically near the solution [9, 17]. The implementation may be challenging due to the energy-stability formulation of the convection term and the computational cost of computing the Jacobian.

Finally, we note that it may be possible to reduce the overall computational time

by loosening the linear solver's tolerance and applying an inexact solution procedure. This remains as an item for future work.

Acknowledgment. We would like to thank two anonymous referees for their helpful comments and suggestions.

REFERENCES

[1] F. Armero and J. C. Simo, *Long-term dissipativity of time-stepping algorithms for an abstract evolution equation with applications to the incompressible MHD and Navier-Stokes equations*, Comput. Methods Appl. Mech. Engrg., 131 (1996), pp. 41–90, https://doi.org/10.1016/0045-7825(95)00931-0.

[2] S. Balay, M. F. Adams, J. Brown, P. Brune, K. Buschelman, V. Eijkhout, W. D. Gropp, D. Kaushik, M. G. Knepley, L. C. McInnes, K. Rupp, B. F. Smith, and H. Zhang, *PETSc User's Manual*, Technical Report ANL-95/11 - Revision 3.4, Argonne National Laboratory, 2013, http://www.mcs.anl.gov/petsc.

[3] S. Balay, M. F. Adams, J. Brown, P. Brune, K. Buschelman, V. Eijkhout, W. D. Gropp, D. Kaushik, M. G. Knepley, L. C. McInnes, K. Rupp, B. F. Smith, and H. Zhang, *PETSc web page*, http://www.mcs.anl.gov/petsc (2014).

[4] H. C. Elman, D. J. Silvester, and A. J. Wathen, *Finite Elements and Fast Iterative Solvers: With Applications in Incompressible Fluid Dynamics*, 2nd ed., Oxford University Press, 2014, https://doi.org/10.1093/acprof:oso/9780199678792.001.0001.

[5] R. D. Falgout and U. Yang, hypre*: A library of high performance preconditioners*, in Computational Science ICCS 2002, Lecture Notes in Comput. Sci. 2331, Springer, Berlin, Heidelberg, 2002, pp. 632–641, https://doi.org/10.1007/3-540-31619-1_8.

[6] J. F. Gerbeau, C. Le Bris, and T. Lelièvre, *Mathematical Methods for the Magnetohydrodynamics of Liquid Metals*, Oxford University Press, 2006, https://doi.org/10.1093/acprof:oso/9780198566656.001.0001.

[7] C. Greif, Li, D. Schötzau, and X. Wei, *A mixed finite element method with exactly divergence-free velocities for incompressible magnetohydrodynamics*, Comput. Methods Appl. Mech. Engrg., 199 (2010), pp. 2840–2855, https://doi.org/10.1016/j.cma.2010.05.007.

[8] 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, https://doi.org/10.1002/nla.515.

[9] M. D. Gunzburger, A. J. Meir, and J. S. Peterson, *On the existence, uniqueness, and finite element approximation of solutions of the equations of stationary, incompressible magnetohydrodynamics*, Math. Comp., 56 (1991), pp. 523–563, https://doi.org/10.2307/2008394.

[10] 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, https://doi.org/10.1137/060660588.

[11] T. V. Kolev and P. Vassilevski, *Parallel auxiliary space AMG for $H(curl)$ problems*, J. Comput. Math., 27 (2009), pp. 604–623, https://doi.org/10.4208/jcm.2009.27.5.013.

[12] D. Li, *Numerical Solution of the Time-Harmonic Maxwell Equations and Incompressible Magnetohydrodynamics Problems*, Ph.D. thesis, University of British Columbia, 2010, https://doi.org/10.14288/1.0051989.

[13] D. Li, C. Greif, and D. Schötzau, *Parallel numerical solution of the time-harmonic Maxwell equations in mixed form*, Numer. Linear Algebra Appl., 19 (2012), pp. 525–539, https://doi.org/10.1002/nla.782.

[14] A. Logg, K. A. Mardal, and G. N. Wells, eds., *Automated Solution of Differential Equations by the Finite Element Method*, Lect. Notes Comput. Sci. Eng. 84, Springer, Heidelberg, 2012, https://doi.org/10.1007/978-3-642-23099-8.

[15] P. Monk, *Finite Element Methods for Maxwell's Equations*, Oxford University Press, 2003, https://doi.org/10.1093/acprof:oso/9780198508885.001.0001.

[16] 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, https://doi.org/10.1137/S1064827599355153.

[17] 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, https://doi.org/10.1137/140955082.

[18] P. H. Roberts, *An Introduction to Magnetohydrodynamics*, Longmans, London, 1967.

[19] D. Schötzau, *Mixed finite element methods for stationary incompressible magneto–hydrodynamics*, Numer. Math., 96 (2004), pp. 771–800, https://doi.org/10.1007/s00211-003-0487-4.

[20] M. Wathen, *Iterative Solution of a Mixed Finite Element Discretisation of an Incompressible Magnetohydrodynamics Problem*, Master's thesis, University of British Columbia, 2014, https://doi.org/10.14288/1.0135538.

[21] X. Wei, *Mixed Discontinuous Galerkin Finite Element Methods for Incompressible Magnetohydrodynamics*, Ph.D. thesis, University of British Columbia, 2011, https://doi.org/10.14288/1.0080512.

[22] J. Xu, *The auxiliary space method and optimal multigrid preconditioning techniques for unstructured grids*, Computing, 56 (1996), pp. 215–235, https://doi.org/10.1007/BF02238513.