

AN AUGMENTED LAGRANGIAN PRECONDITIONER FOR THE MAGNETOHYDRODYNAMICS EQUATIONS AT HIGH REYNOLDS AND COUPLING NUMBERS*

FABIAN LAAKMANN[†], PATRICK E. FARRELL[†], AND LAWRENCE MITCHELL[‡]

Abstract. The magnetohydrodynamics (MHD) equations are generally known to be difficult to solve numerically, due to their highly nonlinear structure and the strong coupling between the electromagnetic and hydrodynamic variables, especially for high Reynolds and coupling numbers. In this work, we present a scalable augmented Lagrangian preconditioner for a finite element discretization of the **B-E** formulation of the incompressible viscoresistive MHD equations. For stationary problems, our solver achieves robust performance with respect to the Reynolds and coupling numbers in two dimensions and good results in three dimensions. We extend our method to fully implicit methods for time-dependent problems which we solve robustly in both two and three dimensions. Our approach relies on specialized parameter-robust multigrid methods for the hydrodynamic and electromagnetic blocks. The scheme ensures exactly divergence-free approximations of both the velocity and the magnetic field up to solver tolerances. We confirm the robustness of our solver by numerical experiments in which we consider fluid and magnetic Reynolds numbers and coupling numbers up to 10,000 for stationary problems and up to 100,000 for transient problems in two and three dimensions.

Key words. magnetohydrodynamics (MHD), multigrid, augmented Lagrangian

MSC codes. 65N55, 65N30, 65F10, 65F08

DOI. 10.1137/21M1416539

1. Introduction. In this work, we consider the incompressible viscoresistive magnetohydrodynamics (MHD) equations on a simply connected polytopal domain $\Omega \subset \mathbb{R}^d$, $d \in \{2, 3\}$. In the stationary three-dimensional setting, we investigate the formulation

$$(1.1a) \quad -\frac{2}{\text{Re}} \operatorname{div} \varepsilon(\mathbf{u}) + \mathbf{u} \cdot \nabla \mathbf{u} + \nabla p + S \mathbf{B} \times (\mathbf{E} + \mathbf{u} \times \mathbf{B}) = \mathbf{f},$$

$$(1.1b) \quad \operatorname{div} \mathbf{u} = 0,$$

$$(1.1c) \quad \mathbf{E} + \mathbf{u} \times \mathbf{B} - \frac{1}{\text{Re}_m} \operatorname{curl} \mathbf{B} = \mathbf{0},$$

$$(1.1d) \quad \operatorname{curl} \mathbf{E} = \mathbf{0},$$

$$(1.1e) \quad \operatorname{div} \mathbf{B} = 0.$$

Here, $\mathbf{u} : \Omega \rightarrow \mathbb{R}^3$ denotes the velocity, $p : \Omega \rightarrow \mathbb{R}$ the fluid pressure, $\mathbf{B} : \Omega \rightarrow \mathbb{R}^3$ the magnetic field, $\mathbf{E} : \Omega \rightarrow \mathbb{R}^3$ the electric field, Re the fluid Reynolds number, Re_m the magnetic Reynolds number, S the coupling number, $\mathbf{f} : \Omega \rightarrow \mathbb{R}^3$ a source term, and $\varepsilon(\mathbf{u}) = \frac{1}{2}(\nabla \mathbf{u} + \nabla \mathbf{u}^\top)$. The system is completed with the boundary conditions

$$(1.2) \quad \mathbf{u} = \mathbf{0}, \quad \mathbf{E} \times \mathbf{n} = \mathbf{0}, \quad \mathbf{B} \cdot \mathbf{n} = 0 \quad \text{on } \partial\Omega,$$

*Submitted to the journal's Computational Methods in Science and Engineering section April 30, 2021; accepted for publication (in revised form) June 6, 2022; published electronically August 15, 2022.

<https://doi.org/10.1137/21M1416539>

Funding: The work of the first author was supported by the EPSRC Centre for Doctoral Training in Partial Differential Equations: Analysis and Applications, grant EP/L015811/1. The work of the second author was supported by EPSRC grants EP/V001943/1 and EP/R029423/1.

[†]Mathematical Institute, University of Oxford, Oxford, UK (fabian.laakmann@maths.ox.ac.uk, patrick.farrell@maths.ox.ac.uk).

[‡]NVIDIA Corporation, Santa Clara, CA USA (lmitchell@nvidia.com).

where \mathbf{n} is the unit outer normal vector. The above formulation based on the electric and magnetic fields was first rigorously analyzed by Hu, Qiu, and Shi [31].

In two dimensions, the electric field is a scalar field and hence denoted as E . The curl operators and cross products are interpreted depending on whether the arguments are scalar or vector valued corresponding to the definitions

$$(1.3) \quad \operatorname{curl} \mathbf{B} = \partial_x B_2 - \partial_y B_1, \quad \mathbf{curl} E = \begin{pmatrix} \partial_y E \\ -\partial_x E \end{pmatrix}$$

and

$$(1.4) \quad \mathbf{u} \times \mathbf{B} = u_1 B_2 - u_2 B_1, \quad \mathbf{B} \times E = \begin{pmatrix} B_2 E \\ -B_1 E \end{pmatrix}.$$

Moreover, the boundary conditions for the electric field change to $E = 0$ on $\partial\Omega$ in two dimensions.

Other formulations include the current density $\mathbf{j} = \mathbf{E} + \mathbf{u} \times \mathbf{B}$ [32] as an unknown or eliminate the electric field using (1.1c). In addition to the stationary case, we also consider the time-dependent version of (1.1), where the time-derivatives $\frac{\partial \mathbf{u}}{\partial t}$ and $\frac{\partial \mathbf{B}}{\partial t}$ are added to (1.1a) and (1.1d), respectively, with suitable initial conditions $\mathbf{u}(\mathbf{x}, 0) = \mathbf{u}_0(\mathbf{x})$ and $\mathbf{B}(\mathbf{x}, 0) = \mathbf{B}_0(\mathbf{x}) \forall \mathbf{x} \in \Omega$. Note that MHD models neglect displacement currents $\frac{\partial \mathbf{E}}{\partial t}$ [27, sect. 1.5].

The main contribution of this work is to provide block preconditioners for linearizations of (1.1) with good convergence even at high Reynolds and coupling numbers. The performance relies on the following three (novel) approaches:

- 1) We consider a fluid-Reynolds-robust augmented Lagrangian preconditioner for an $\mathbf{H}(\operatorname{div}) \times L^2$ -discretization of the Navier–Stokes equations that relies on a specialized multigrid method.
- 2) We introduce a new monolithic multigrid method for the electromagnetic block.
- 3) We discover that using the outer Schur complement which eliminates the (\mathbf{u}, p) block instead of the (\mathbf{E}, \mathbf{B}) block has crucial advantages for ensuring robustness for high parameters.

Furthermore, we show that our preconditioners extend in a straightforward manner to the time-dependent version of (1.1). This has the substantial advantage that the choice of the time-stepping scheme is no longer restricted by the ability to solve the linear systems. In particular, it allows for the use of fully implicit methods for high Reynolds numbers and coupling parameters.

An important point for discretizations is the enforcement of the magnetic Gauss' law $\operatorname{div} \mathbf{B} = 0$ in the weak formulation, achieved in most cases by a nonphysical Lagrange multiplier r [47]. However, in general, a Lagrange multiplier only enforces the divergence constraint in a weak sense, which can cause severe problems for the discretization and numerical simulations [13, 19]. For the \mathbf{B} - \mathbf{E} formulation (1.1) Hu, Qiu, and Shi [31] show that both a Lagrange multiplier and an augmented Lagrangian term lead to a pointwise preservation of Gauss' law with appropriate choices of discrete spaces. In this work, we consider the latter approach by replacing (1.1d) with

$$(1.5) \quad \frac{1}{\operatorname{Re}_m} \nabla \operatorname{div} \mathbf{B} + \mathbf{curl} \mathbf{E} = \mathbf{0}.$$

The literature proposes numerous numerical schemes and preconditioning strategies for the numerical solution of the different formulations. The most common approach is based on block preconditioners in both the stationary [35, 40, 41, 53, 52]

and time-dependent [16, 18, 42] cases. Here, the main challenges are to find suitable approximations of one or more Schur complements and robust linear solvers for the inner auxiliary problems. Phillips et al. [42] simplify the Schur complement by the use of vector identities and approximate the remaining parts based on a spectral analysis. They report iteration counts for a stationary three-dimensional lid-driven cavity problem up to $Re = Re_m = 100$. A similar approach is used by Wathen and Greif in [52] where they construct an approximate inverse block preconditioner by sparsifying a derived formula for the exact inverse and drop low order terms. Here, results for Hartmann numbers $Ha = \sqrt{SRe_mRe}$ up to 1,000 are reported for stationary problems. Other approaches include fully coupled geometric [1, 2] and algebraic [49, 50] monolithic multigrid methods. In [1], Adler et al. present results for a two-dimensional Hartmann problem for parameters up to $Re = Re_m = 64$.

However, the performance of most of these preconditioners deteriorates significantly for high Reynolds and coupling numbers. To the best of our knowledge, a practical robust preconditioner for the stationary MHD equations has not yet been proposed. The common problem for high magnetic Reynolds numbers and coupling numbers for the stationary case is that all available Schur complement approximations become less accurate for Newton-type linearizations, causing the linear solver to fail to converge. Conversely, Picard-type linearizations can allow an exact computation of the Schur complement but fail to converge in the nonlinear iteration.

In this work, we consider two different linearizations. The first is the Picard iteration proposed by Hu, Qiu, and Shi [31]. We compute an approximation to the outer Schur complement of the arising block system and introduce a robust linear solver for the different blocks. This scheme works well for small magnetic Reynolds numbers but the nonlinear iteration fails to converge for higher Re_m , as anticipated in the analysis of [31]. The second is a full Newton linearization, which converges well for high Reynolds numbers and coupling numbers. However, our approximation of the Schur complement deteriorates slightly for high parameters.

Ma et al. [36] have developed Reynolds-robust preconditioners for the time-dependent MHD equations that are based on norm-equivalent and field-of-values equivalent approaches. To the best of our knowledge, their strategy does not extend to the stationary case; in general, the time-dependent case offers crucial advantages for the development of robust solvers. For example, Ma et al. treat complicated terms like the hydrodynamic convection term $\mathbf{u} \cdot \nabla \mathbf{u}$ explicitly in the time-stepping scheme, which can cause problems for convection-dominated problems and does not apply in the stationary case. The discretization of the time derivative causes mass matrices with a scaling of $1/\Delta t$, where Δt denotes the time step size, to appear in the block matrix on the diagonal blocks. As we will see also in our numerical results for the time-dependent problems, these extra terms dominate the scheme for small Δt and hence simplify the development of robust solvers.

Most applications are in the regime of high Reynolds and coupling numbers and hence it is of great interest to build robust solvers with respect to these parameters. For liquid metals, the fluid Reynolds number Re tends to be much larger than Re_m . For example, the flow of liquid mercury is characterized by a ratio of 10^7 between these two constants; typical values in aluminium electrolysis are $Re_m = 10^{-1}$ and $Re = 10^5$ [27]. High magnetic Reynolds numbers occur on large length scales, as in geophysics and astrophysics. The magnetic Reynolds number of the outer Earth's core is in the range of 10^3 and of the sun is in the range of 10^6 [20]. Magnetic Reynolds numbers between $10^1 - 10^3$ have been used in several dynamo experiments that investigate planetary magnetic fields [37]. The coupling number S is around 10^0 for aluminium

electrolysis [27] and Armero and Simo [4] define strong coupling for S in the range of $10^2 - 10^9$.

The remainder of this work is outlined as follows. In section 2, we derive an augmented Lagrangian formulation for (1.1) and describe the finite element discretization and linearization schemes. In section 3, we introduce block preconditioners for these schemes, present a calculation of the corresponding (approximate) Schur complements, and describe robust linear multigrid solvers for the different blocks. Numerical examples and a detailed description of the algorithm are presented in section 4.

2. Formulation, linearization, and discretization.

2.1. An augmented Lagrangian formulation. We modify (1.1) by introducing two augmented Lagrangian terms: $-\gamma \nabla \operatorname{div} \mathbf{u}$ for $\gamma > 0$ is added to (1.1a), and (as previously discussed) $-1/\operatorname{Re}_m \nabla \operatorname{div} \mathbf{B}$ is added to (1.1d). Note that both terms leave the continuous solution of the problem unchanged. We use the first term to control the Schur complement of the fluid subsystem [11, 24] and the second term to enforce the divergence constraint $\operatorname{div} \mathbf{B} = 0$, as shown in [31, Thm. 9].

Following these modifications, we consider the following system,

$$(2.1a) \quad -\frac{2}{\operatorname{Re}} \operatorname{div} \varepsilon(\mathbf{u}) + \mathbf{u} \cdot \nabla \mathbf{u} - \gamma \nabla \operatorname{div} \mathbf{u} + \nabla p + S \mathbf{B} \times (\mathbf{E} + \mathbf{u} \times \mathbf{B}) = \mathbf{f},$$

$$(2.1b) \quad \operatorname{div} \mathbf{u} = 0,$$

$$(2.1c) \quad \mathbf{E} + \mathbf{u} \times \mathbf{B} - \frac{1}{\operatorname{Re}_m} \operatorname{curl} \mathbf{B} = \mathbf{0},$$

$$(2.1d) \quad -\frac{1}{\operatorname{Re}_m} \nabla \operatorname{div} \mathbf{B} + \operatorname{curl} \mathbf{E} = \mathbf{0},$$

subject to the boundary conditions (1.2). For convenience, we consider homogeneous boundary conditions in this section but all the results extend in a straightforward manner to inhomogeneous boundary conditions. However, there are subtle technicalities for the implementation of the degrees of freedom in the finite element method in the inhomogeneous case, which are explained in detail in section 4.2.

The weak formulation of (2.1) seeks $\mathcal{U} := (\mathbf{u}, p, \mathbf{E}, \mathbf{B}) \in \mathcal{Z} := \mathbf{V} \times Q \times \mathbf{R} \times \mathbf{W}$ with

$$(2.2) \quad \mathbf{V} := \mathbf{H}_0^1(\Omega), \quad Q := L_0^2(\Omega), \quad \mathbf{R} := \mathbf{H}_0(\operatorname{curl}, \Omega), \quad \mathbf{W} := \mathbf{H}_0(\operatorname{div}, \Omega).$$

In two dimensions, the space for the electric field is scalar valued and can be identified with $R := H_0^1(\Omega)$. The weak formulation is to find $\mathcal{U} \in \mathcal{Z}$ such that for all $\mathcal{V} := (\mathbf{v}, q, \mathbf{F}, \mathbf{C}) \in \mathcal{Z}$ and $\mathcal{F} = (\mathbf{f}, 0, \mathbf{0}, \mathbf{0})$ there holds

$$(2.3) \quad \mathcal{R}(\mathcal{U}, \mathcal{V}) := \mathcal{N}(\mathcal{U}, \mathcal{V}) - (\mathcal{F}, \mathcal{V}) = 0$$

with

$$(2.4) \quad \begin{aligned} \mathcal{N}(\mathcal{U}, \mathcal{V}) = & \frac{2}{\operatorname{Re}} (\varepsilon(\mathbf{u}), \varepsilon(\mathbf{v})) + (\mathbf{u} \cdot \nabla \mathbf{u}, \mathbf{v}) + \gamma (\operatorname{div} \mathbf{u}, \operatorname{div} \mathbf{v}) \\ & - (p, \operatorname{div} \mathbf{v}) + S(\mathbf{B} \times \mathbf{E}, \mathbf{v}) + S(\mathbf{B} \times (\mathbf{u} \times \mathbf{B}), \mathbf{v}) \\ & - (\operatorname{div} \mathbf{u}, q) \\ & + (\mathbf{E}, \mathbf{F}) + (\mathbf{u} \times \mathbf{B}, \mathbf{F}) - \frac{1}{\operatorname{Re}_m} (\mathbf{B}, \operatorname{curl} \mathbf{F}) \\ & + \frac{1}{\operatorname{Re}_m} (\operatorname{div} \mathbf{B}, \operatorname{div} \mathbf{C}) + (\operatorname{curl} \mathbf{E}, \mathbf{C}). \end{aligned}$$

All boundary integrals that result from integration by parts vanish because of the choice of the boundary conditions (1.2).

Note that \mathbf{W} and \mathbf{R} are chosen from the de Rham complex [5]

$$(2.5) \quad \mathbb{R} \xrightarrow{\text{id}} H_0^1(\Omega) \xrightarrow{\text{grad}} \mathbf{H}_0(\text{curl}, \Omega) \xrightarrow{\text{curl}} \mathbf{H}_0(\text{div}, \Omega) \xrightarrow{\text{div}} L_0^2(\Omega) \xrightarrow{\text{null}} 0,$$

which is exact for the simply connected domains we consider. The corresponding complex in two dimensions is given by

$$(2.6) \quad \mathbb{R} \xrightarrow{\text{id}} H_0(\text{curl}, \Omega) \xrightarrow{\text{curl}} \mathbf{H}_0(\text{div}, \Omega) \xrightarrow{\text{div}} L_0^2(\Omega) \xrightarrow{\text{null}} 0.$$

This ensures that formulation (2.1) enforces the divergence constraint $\text{div } \mathbf{B} = 0$ and $\text{curl } \mathbf{E} = \mathbf{0}$. To see this, we test (2.3) with $\mathcal{V} = (\mathbf{0}, 0, \mathbf{0}, \text{curl } \mathbf{E})$ and conclude that $\text{curl } \mathbf{E} = \mathbf{0}$. Here, \mathcal{V} is a valid test function because the above exact sequence implies that $\text{curl}(\mathbf{R}) = \mathbf{W}$. Similarly, testing with $\mathcal{V} = (\mathbf{0}, 0, \mathbf{0}, \mathbf{B})$ results in $\text{div } \mathbf{B} = 0$.

2.2. Linearization: Newton and Picard. The Newton linearization of (2.3) for the initial guess $\mathcal{U}^n = (\mathbf{u}^n, p^n, \mathbf{E}^n, \mathbf{B}^n)$ is to find an update $\delta\mathcal{U}$ such that

$$(2.7) \quad \mathcal{N}_N(\delta\mathcal{U}, \mathcal{U}^n, \mathcal{V}) = \mathcal{R}(\mathcal{U}^n, \mathcal{V}) \quad \forall \mathcal{V} \in \mathcal{Z},$$

$$(2.8) \quad \mathcal{U}^{n+1} = \mathcal{U}^n + \delta\mathcal{U},$$

with the weak form of the nonlinear residual $\mathcal{R}(\mathcal{U}^n, \mathcal{V})$ evaluated at \mathcal{U}^n and

$$(2.9) \quad \begin{aligned} \mathcal{N}_N(\delta\mathcal{U}, \mathcal{U}^n, \mathcal{V}) = & \frac{2}{\text{Re}} (\varepsilon(\delta\mathbf{u}), \varepsilon(\mathbf{v})) + (\mathbf{u}^n \cdot \nabla \delta\mathbf{u}, \mathbf{v}) + (\delta\mathbf{u} \cdot \nabla \mathbf{u}^n, \mathbf{v}) \\ & + \gamma(\text{div } \delta\mathbf{u}, \text{div } \mathbf{v}) - (\delta p, \text{div } \mathbf{v}) \\ & + S(\mathbf{B}^n \times \delta\mathbf{E}, \mathbf{v}) + S(\delta\mathbf{B} \times \mathbf{E}^n, \mathbf{v}) \\ & + S(\mathbf{B}^n \times (\delta\mathbf{u} \times \mathbf{B}^n), \mathbf{v}) + S(\delta\mathbf{B} \times (\mathbf{u}^n \times \mathbf{B}^n), \mathbf{v}) \\ & + S(\mathbf{B}^n \times (\mathbf{u}^n \times \delta\mathbf{B}), \mathbf{v}) \\ & - (\text{div } \delta\mathbf{u}, q) \\ & + (\delta\mathbf{E}, \mathbf{F}) + (\mathbf{u}^n \times \delta\mathbf{B}, \mathbf{F}) + (\delta\mathbf{u} \times \mathbf{B}^n, \mathbf{F}) \\ & - \frac{1}{\text{Re}_m} (\delta\mathbf{B}, \text{curl } \mathbf{F}) \\ & + \frac{1}{\text{Re}_m} (\text{div } \delta\mathbf{B}, \text{div } \mathbf{C}) + (\text{curl } \delta\mathbf{E}, \mathbf{C}). \end{aligned}$$

The bilinear form for the Picard iteration that we consider is given by

$$(2.10) \quad \begin{aligned} \mathcal{N}_P(\delta\mathcal{U}, \mathcal{U}^n, \mathcal{V}) = & \mathcal{N}_N(\delta\mathcal{U}, \mathcal{U}^n, \mathcal{V}) - S(\delta\mathbf{B} \times \mathbf{E}^n, \mathbf{v}) - S(\mathbf{B}^n \times (\mathbf{u}^n \times \delta\mathbf{B}), \mathbf{v}) \\ & - S(\delta\mathbf{B} \times (\mathbf{u}^n \times \mathbf{B}^n), \mathbf{v}) - (\mathbf{u}^n \times \delta\mathbf{B}, \mathbf{F}). \end{aligned}$$

Note that in contrast to the definition of the Picard iteration in [31], we do not scale the term $(\text{curl } \delta\mathbf{E}, \mathbf{C})$ with S/Re_m and consider the full Newton linearization of the advection term $(\mathbf{u} \cdot \nabla)\mathbf{u}$. The advantage of the Picard linearization (2.10) in comparison to the Newton linearization (2.9) is that it allows an exact Schur complement computation in two dimensions and converges well for high Re . However, its major disadvantage is the failure of nonlinear convergence for high Re_m .

2.3. Discretization. For a finite element discretization, we seek $\mathcal{U}_h := (\mathbf{u}_h, p_h, \mathbf{E}_h, \mathbf{B}_h) \in \mathcal{Z}_h := \mathbf{V}_h \times Q_h \times \mathbf{R}_h \times \mathbf{W}_h$ such that

$$(2.11) \quad \mathcal{N}(\mathcal{U}_h, \mathcal{V}_h) = (\mathcal{F}, \mathcal{V}_h) \quad \forall \mathcal{V}_h \in \mathcal{Z}_h.$$

We choose Raviart–Thomas elements of degree k \mathbb{RT}_k [44] for \mathbf{W}_h , Nédélec elements of first kind $\mathbb{NED1}_k$ [38] for \mathbf{R}_h in 3D, and continuous Lagrange elements \mathbb{CG}_k for R_h in 2 dimensions (2D). Note that these elements belong to a discrete subcomplex of (2.5)

$$(2.12) \quad \mathbb{CG}_k \xrightarrow{\text{grad}} \mathbb{NED1}_k \xrightarrow{\text{curl}} \mathbb{RT}_k \xrightarrow{\text{div}} \mathbb{DG}_{k-1} \xrightarrow{\text{null}} 0,$$

and of (2.6)

$$(2.13) \quad \mathbb{CG}_k \xrightarrow{\text{curl}} \mathbb{RT}_k \xrightarrow{\text{div}} \mathbb{DG}_{k-1} \xrightarrow{\text{null}} 0.$$

This implies that we enforce $\text{div } \mathbf{B}_h = 0$ and $\text{curl } \mathbf{E}_h = \mathbf{0}$ pointwise with the same proof as for the continuous case. These identities also hold for inhomogeneous boundary conditions, since the interpolation operator $\mathcal{I}_{\mathbf{W}_h}^h$ into the Raviart–Thomas space satisfies, for all divergence-free $\mathbf{B} \in \mathbf{H}_0(\text{div}, \Omega)$ [12, Prop. 2.5.2],

$$(2.14) \quad \text{div}(\mathcal{I}_{\mathbf{W}_h}^h \mathbf{B}) = 0.$$

Moreover, following [21], we add the following stabilization term to address the problem that the Galerkin discretization of advection-dominated problems can be oscillatory [22],

$$(2.15) \quad \sum_{K \in \mathcal{M}_h} \frac{1}{2} \int_{\partial K} \mu h_{\partial K}^2 [[\nabla \mathbf{u}_h]] : [[\nabla \mathbf{v}_h]] \, ds.$$

Here, $[[\nabla \mathbf{u}_h]]$ denotes the jump of the gradient, $h_{\partial K}$ is a function giving the facet size, and μ is a free parameter that is chosen according to [15].

Note that a fully robust discretization should also include a stabilization term for the magnetic field \mathbf{B} in the case of dominating magnetic advection. The literature does not propose many stabilization types for this problem. The most promising work by Wu and Xu [54] uses the so-called SAFE-scheme for stabilization which is based on an exponential fitting approach. While the original SAFE-scheme is only a first order method, it can be extended to higher order as shown in [55]. We aim to include this stabilization in future work.

For the hydrodynamic part, we consider the $\mathbf{H}(\text{div}) \times L^2$ -conforming element pair $\mathbb{BDM}_k \times \mathbb{DG}_{k-1}$ with the Brezzi–Douglas–Marini element \mathbb{BDM}_k of order k [14, 39]. This discretization ensures that $\text{div } \mathbf{u}_h = 0$ holds pointwise since $\text{div } \mathbf{V}_h \subset Q_h$. Additionally, it exhibits pressure robustness, i.e., the error estimates do not degrade for high Reynolds numbers [33].

Since the discretization is nonconforming, we must consider a discontinuous Galerkin formulation of the hydrodynamic advection and diffusion terms [26, sect. 7]. We denote by $\mathcal{F}_h = \mathcal{F}_h^i \cup \mathcal{F}_h^\partial$ all facets of the triangulation, which consists of the interior facets \mathcal{F}_h^i and the Dirichlet boundary facets \mathcal{F}_h^∂ . We assign to each $F \in \mathcal{F}_h$ its diameter h_F and unit normal vector \mathbf{n}_F . The jump and average operators across a facet are denoted by $[[\cdot]]$ and $\{\{\cdot\}\}$, respectively, and are defined as $[[\Phi]] = \Phi^+ - \Phi^-$ and $\{\{\Phi\}\} = \frac{1}{2}(\Phi^+ + \Phi^-)$. The penalization parameter is chosen as $\sigma = 10k^2$. Inhomogeneous

boundary data are described by \mathbf{g}_D . We then add the following bilinear forms to (2.11):

$$\begin{aligned}
 a_h^{DG}(\mathbf{u}_h, \mathbf{v}_h) = & -\frac{2}{\text{Re}} \sum_{F \in \mathcal{F}_h} \int_F \{ \{ \varepsilon(\mathbf{u}_h) \} \} \mathbf{n}_F \cdot \llbracket \mathbf{v}_h \rrbracket \, ds \\
 & -\frac{2}{\text{Re}} \sum_{F \in \mathcal{F}_h} \int_F \llbracket \mathbf{u}_h \rrbracket \cdot \{ \{ \varepsilon(\mathbf{v}_h) \} \} \mathbf{n}_F \, ds \\
 (2.16) \quad & +\frac{1}{\text{Re}} \sum_{F \in \mathcal{F}_h} \frac{\sigma}{h_F} \int_F \llbracket \mathbf{u}_h \rrbracket \cdot \llbracket \mathbf{v}_h \rrbracket \, ds \\
 & -\frac{1}{\text{Re}} \sum_{F \in \mathcal{F}_h^\partial} \frac{\sigma}{h_F} \int_F \mathbf{g}_D \cdot \mathbf{v}_h \, ds + \frac{2}{\text{Re}} \sum_{F \in \mathcal{F}_h^\partial} \int_F \mathbf{g}_D \cdot \varepsilon(\mathbf{v}_h) \mathbf{n}_F \, ds,
 \end{aligned}$$

$$\begin{aligned}
 c_h^{DG}(\mathbf{u}_h, \mathbf{v}_h) = & \frac{1}{2} \sum_{F \in \mathcal{F}_h^i} \int_F \llbracket (\mathbf{u}_h \cdot \mathbf{n}_F + |\mathbf{u}_h \cdot \mathbf{n}_F|) \mathbf{u}_h \rrbracket \cdot \llbracket \mathbf{v}_h \rrbracket \, ds \\
 (2.17) \quad & +\frac{1}{2} \sum_{F \in \mathcal{F}_h^\partial} \int_F (\mathbf{u}_h \cdot \mathbf{n}_F + |\mathbf{u}_h \cdot \mathbf{n}_F|) \mathbf{u}_h \cdot \mathbf{v}_h \, ds \\
 & +\frac{1}{2} \sum_{F \in \mathcal{F}_h^\partial} \int_F (\mathbf{u}_h \cdot \mathbf{n}_F - |\mathbf{u}_h \cdot \mathbf{n}_F|) \mathbf{g}_D \cdot \mathbf{v}_h \, ds.
 \end{aligned}$$

Hu, Qiu, and Shi prove in [31, Theorem 4] that (2.11) is well-posed and has at least one solution. The solution is unique for suitable source and boundary data. While the well-posedness and convergence of the Newton iteration remains an open problem, Hu, Qiu, and Shi prove that the Picard iteration converges to the unique solution of (2.11) if both $\text{Re}^2 \|f\|_{-1}$ and $\text{Re} \text{Re}_m^{\frac{3}{2}} \|f\|_{-1}$ are small enough.

For the Newton linearization (2.9), we must solve the following linear system at each step,

$$(2.18) \quad \begin{bmatrix} \mathcal{F} + \mathcal{D} & \mathcal{B}^\top & \mathcal{J} & \tilde{\mathcal{J}} + \tilde{\mathcal{D}}_1 + \tilde{\mathcal{D}}_2 \\ \mathcal{B} & \mathbf{0} & \mathbf{0} & \mathbf{0} \\ \mathcal{G} & \mathbf{0} & \mathcal{M}_E & \tilde{\mathcal{G}} - \frac{1}{\text{Re}_m} \mathcal{A} \\ \mathbf{0} & \mathbf{0} & \mathcal{A}^\top & \mathcal{C} \end{bmatrix} \begin{bmatrix} x_{\mathbf{u}} \\ x_p \\ x_{\mathbf{E}} \\ x_{\mathbf{B}} \end{bmatrix} = \begin{bmatrix} \mathcal{R}_{\mathbf{u}} \\ \mathcal{R}_p \\ \mathcal{R}_{\mathbf{E}} \\ \mathcal{R}_{\mathbf{B}} \end{bmatrix},$$

where $x_{\mathbf{u}}$, x_p , $x_{\mathbf{E}}$, and $x_{\mathbf{B}}$ are the coefficients of the discretized Newton corrections and $\mathcal{R}_{\mathbf{u}}$, \mathcal{R}_p , $\mathcal{R}_{\mathbf{E}}$, and $\mathcal{R}_{\mathbf{B}}$ the corresponding nonlinear residuals. The correspondence between the discrete and continuous operators is illustrated in Table 2.1. We have chosen the notation that operators that include a tilde are omitted in the Picard linearization (2.10). Moreover, we introduce $\eta \in \{0, 1\}$ to distinguish between the stationary ($\eta = 0$) and transient ($\eta = 1$) cases.

For the time-dependent equations, we concentrate here on the implicit Euler method, but the following computations are straightforward to adapt to other implicit multistep methods. We use the same finite element discretization as in the stationary case. Note that in the transient case, the equation

$$(2.19) \quad \partial_t \mathbf{B} + \text{curl} \mathbf{E} = \mathbf{0}$$

immediately implies $\text{div} \mathbf{B} = 0$ if the initial condition satisfies $\text{div} \mathbf{B}_0 = 0$, and this remains true on the discrete level up to solver tolerances; see [30, Theorem 1] for

TABLE 2.1

Overview of operators. Terms that include a tilde are dropped in the Picard iteration. The stationary and transient cases are distinguished by $\eta \in \{0, 1\}$.

Discrete	Continuous	Weak form
$\mathcal{F}\mathbf{u}$	$\frac{\eta}{\Delta t} \mathbf{u} - \frac{2}{\text{Re}} \text{div } \boldsymbol{\varepsilon}(\mathbf{u}) + \mathbf{u}^n \cdot \nabla \mathbf{u}$ $+ \mathbf{u} \cdot \nabla \mathbf{u}^n - \gamma \nabla \text{div } \mathbf{u}$	$\frac{\eta}{\Delta t} (\mathbf{u}, \mathbf{v}) + \frac{2}{\text{Re}} (\boldsymbol{\varepsilon}(\mathbf{u}), \boldsymbol{\varepsilon}(\mathbf{v})) + (\mathbf{u}^n \cdot \nabla \mathbf{u}, \mathbf{v})$ $+ (\mathbf{u} \cdot \nabla \mathbf{u}^n, \mathbf{v}) + \gamma (\text{div } \mathbf{u}, \text{div } \mathbf{v})$
$\mathcal{D}\mathbf{u}$	$S\mathbf{B}^n \times (\mathbf{u} \times \mathbf{B}^n)$	$S(\mathbf{B}^n \times (\mathbf{u} \times \mathbf{B}^n), \mathbf{v})$
$\mathcal{J}\mathbf{E}$	$S\mathbf{B}^n \times \mathbf{E}$	$S(\mathbf{B}^n \times \mathbf{E}, \mathbf{v})$
$\tilde{\mathcal{J}}\mathbf{B}$	$S\mathbf{B} \times \mathbf{E}^n$	$S(\mathbf{B} \times \mathbf{E}^n, \mathbf{v})$
$\tilde{\mathcal{D}}_1\mathbf{B}$	$S\mathbf{B} \times (\mathbf{u}^n \times \mathbf{B}^n)$	$S(\mathbf{B} \times (\mathbf{u}^n \times \mathbf{B}^n), \mathbf{v})$
$\tilde{\mathcal{D}}_2\mathbf{B}$	$S\mathbf{B}^n \times (\mathbf{u}^n \times \mathbf{B})$	$S(\mathbf{B}^n \times (\mathbf{u}^n \times \mathbf{B}), \mathbf{v})$
$\mathcal{M}_{\mathbf{E}}\mathbf{E}$	\mathbf{E}	(\mathbf{E}, \mathbf{F})
$\mathcal{G}\mathbf{u}$	$\mathbf{u} \times \mathbf{B}^n$	$(\mathbf{u} \times \mathbf{B}^n, \mathbf{F})$
$\tilde{\mathcal{G}}\mathbf{B}$	$\mathbf{u}^n \times \mathbf{B}$	$(\mathbf{u}^n \times \mathbf{B}, \mathbf{F})$
$\mathcal{A}\mathbf{B}$	$\text{curl } \mathbf{B}$	$(\mathbf{B}, \text{curl } \mathbf{F})$
$\mathcal{C}\mathbf{B}$	$\frac{\eta}{\Delta t} \mathbf{B} - \frac{1}{\text{Re}_m} \nabla \text{div } \mathbf{B}$	$\frac{\eta}{\Delta t} (\mathbf{B}, \mathbf{C}) + \frac{1}{\text{Re}_m} (\text{div } \mathbf{B}, \text{div } \mathbf{C})$
$\mathcal{A}^\top \mathbf{E}$	$\text{curl } \mathbf{E}$	$(\text{curl } \mathbf{E}, \mathbf{C})$
$\mathcal{B}^\top p$	∇p	$-(p, \text{div } \mathbf{v})$
$\mathcal{B}\mathbf{u}$	$-\text{div } \mathbf{u}$	$-(\text{div } \mathbf{u}, q)$

a proof for implicit Euler which can be extended to other multistep methods in a straightforward manner, provided all starting values are divergence-free. Hence, the augmented Lagrangian term $-\frac{1}{\text{Re}_m} \nabla \text{div } \mathbf{B}$ is no longer necessary to enforce the divergence constraint and could therefore be omitted. Nevertheless, we retain it in our scheme since we employ the identity

$$(2.20) \quad \frac{1}{\text{Re}_m} \text{curl curl } \mathbf{u} - \frac{1}{\text{Re}_m} \nabla \text{div } \mathbf{u} = -\frac{1}{\text{Re}_m} \Delta \mathbf{u}$$

in our derivation of Schur complement approximations below.

3. Derivation of block preconditioners. We now consider block preconditioners for (2.18). The inverse of a 2×2 block matrix can factorized as [10, 22]

$$(3.1) \quad \begin{pmatrix} \mathcal{M} & \mathcal{K} \\ \mathcal{L} & \mathcal{N} \end{pmatrix}^{-1} = \begin{pmatrix} \mathcal{I} & -\mathcal{M}^{-1}\mathcal{K} \\ \mathbf{0} & \mathcal{I} \end{pmatrix} \begin{pmatrix} \mathcal{M}^{-1} & \mathbf{0} \\ \mathbf{0} & \mathcal{S}^{-1} \end{pmatrix} \begin{pmatrix} \mathcal{I} & \mathbf{0} \\ -\mathcal{L}\mathcal{M}^{-1} & \mathcal{I} \end{pmatrix}$$

provided the top-left block \mathcal{M} and the Schur complement $\mathcal{S} = \mathcal{N} - \mathcal{L}\mathcal{M}^{-1}\mathcal{K}$ are invertible. Since the Schur complement is usually a dense matrix, the main task is to find a suitable approximation $\tilde{\mathcal{S}}$ for the Schur complement \mathcal{S} as well as efficient solvers for \mathcal{M} and $\tilde{\mathcal{S}}$.

In sections 3.1 and 3.2 we derive approximations of the Schur complements for two different block elimination strategies. We briefly introduce the theory of parameter-robust multigrid relaxation in section 3.3, and then describe the multigrid methods that we use to solve the top-left block and the Schur complement approximations in sections 3.4 and 3.5.

Both block preconditioners we consider gather the variables as $(\mathbf{E}_h, \mathbf{B}_h)$ and (\mathbf{u}_h, p_h) . They differ in the order of block elimination: the first takes the Schur complement that eliminates (inverts) the $(\mathbf{E}_h, \mathbf{B}_h)$ block, while the second takes the Schur complement that eliminates the (\mathbf{u}_h, p_h) block. The first choice appears several times in the literature [35, 42], while it seems that the second choice has not

yet been investigated. As we will see, for small Re_m and S both preconditioners perform similarly, while for more difficult parameter regimes the second choice notably outperforms the first. We therefore recommend the second strategy and mainly report numerical results for this choice. Nevertheless, we also investigate the first option, both for comparison and because it allows a much more detailed description of the Schur complement. In 2D it even allows an exact computation of the Schur complement. The two strategies are compared in section 4.4.1 below.

3.1. Outer Schur complement eliminating the $(\mathbf{E}_h, \mathbf{B}_h)$ block. Reordering (2.18) for convenience, we consider

$$(3.2) \quad \left[\begin{array}{cc|cc} \mathcal{M}_{\mathbf{E}} & \tilde{\mathcal{G}} - \frac{1}{\text{Re}_m} \mathcal{A} & \mathcal{G} & \mathbf{0} \\ \mathcal{A}^\top & \mathcal{C} & \mathbf{0} & \mathbf{0} \\ \hline \mathcal{J} & \tilde{\mathcal{J}} + \tilde{\mathcal{D}}_1 + \tilde{\mathcal{D}}_2 & \mathcal{F} + \mathcal{D} & \mathcal{B}^\top \\ \mathbf{0} & \mathbf{0} & \mathcal{B} & \mathbf{0} \end{array} \right] \begin{bmatrix} x_{\mathbf{E}} \\ x_{\mathbf{B}} \\ x_{\mathbf{u}} \\ x_p \end{bmatrix} = \begin{bmatrix} \mathcal{R}_{\mathbf{E}} \\ \mathcal{R}_{\mathbf{B}} \\ \mathcal{R}_{\mathbf{u}} \\ \mathcal{R}_p \end{bmatrix}.$$

In the following, we refer to the Schur complement of the 4 x 4 matrix as the outer Schur complement, while we call the Schur complements of the resulting 2 x 2 blocks inner Schur complements. The outer Schur complement eliminating the $(\mathbf{E}_h, \mathbf{B}_h)$ block is given by

$$(3.3) \quad \mathcal{S}^{(\mathbf{E}, \mathbf{B})} = \begin{bmatrix} \mathcal{F} + \mathcal{D} & \mathcal{B}^\top \\ \mathcal{B} & \mathbf{0} \end{bmatrix} - \begin{bmatrix} \mathcal{J} & \tilde{\mathcal{J}} + \tilde{\mathcal{D}}_1 + \tilde{\mathcal{D}}_2 \\ \mathbf{0} & \mathbf{0} \end{bmatrix} \begin{bmatrix} \mathcal{M}_{\mathbf{E}} & \tilde{\mathcal{G}} - \frac{1}{\text{Re}_m} \mathcal{A} \\ \mathcal{A}^\top & \mathcal{C} \end{bmatrix}^{-1} \begin{bmatrix} \mathcal{G} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{bmatrix}.$$

We simplify $\mathcal{S}^{(\mathbf{E}, \mathbf{B})}$ by applying the identity (3.1) to the top-left electromagnetic block

$$(3.4) \quad \mathcal{M} = \begin{bmatrix} \mathcal{M}_{\mathbf{E}} & \tilde{\mathcal{G}} - \frac{1}{\text{Re}_m} \mathcal{A} \\ \mathcal{A}^\top & \mathcal{C} \end{bmatrix}.$$

This results in

$$(3.5) \quad \mathcal{S}^{(\mathbf{E}, \mathbf{B})} = \begin{bmatrix} \mathcal{F} + \mathcal{D} - \mathcal{J} \mathcal{M}_{1,1}^{-1} \mathcal{G} - (\tilde{\mathcal{J}} + \tilde{\mathcal{D}}_1 + \tilde{\mathcal{D}}_2) \mathcal{M}_{2,1}^{-1} \mathcal{G} & \mathcal{B}^\top \\ \mathcal{B} & \mathbf{0} \end{bmatrix}$$

with

$$(3.6)$$

$$\mathcal{M}_{1,1}^{-1} = \mathcal{M}_{\mathbf{E}}^{-1} + \mathcal{M}_{\mathbf{E}}^{-1} \left(\tilde{\mathcal{G}} - \frac{1}{\text{Re}_m} \mathcal{A} \right) \left(\mathcal{C} - \mathcal{A}^\top \mathcal{M}_{\mathbf{E}}^{-1} \left(\tilde{\mathcal{G}} - \frac{1}{\text{Re}_m} \mathcal{A} \right) \right)^{-1} \mathcal{A}^\top \mathcal{M}_{\mathbf{E}}^{-1}$$

and

$$(3.7) \quad \mathcal{M}_{2,1}^{-1} = - \left(\mathcal{C} - \mathcal{A}^\top \mathcal{M}_{\mathbf{E}}^{-1} \left(\tilde{\mathcal{G}} - \frac{1}{\text{Re}_m} \mathcal{A} \right) \right)^{-1} \mathcal{A}^\top \mathcal{M}_{\mathbf{E}}^{-1}.$$

We precondition $\mathcal{S}^{(\mathbf{E}, \mathbf{B})}$ for both linearizations in the stationary case by

$$(3.8) \quad \tilde{\mathcal{S}}^{(\mathbf{E}, \mathbf{B})} = \begin{bmatrix} \mathcal{F} + \mathcal{D} & \mathcal{B}^\top \\ \mathcal{B} & \mathbf{0} \end{bmatrix},$$

and in the transient case by

$$(3.9) \quad \tilde{\mathcal{S}}_{\alpha}^{(\mathbf{E}, \mathbf{B})} := \begin{bmatrix} \mathcal{F} + \alpha \mathcal{D} & \mathcal{B}^\top \\ \mathcal{B} & \mathbf{0} \end{bmatrix}, \quad \alpha = \frac{\Delta t}{\Delta t + \text{Re}_m h^2 + \delta \text{Re}_m h \|\mathbf{u}^n\|_{L^2} \Delta t}.$$

In the following, we motivate this choice of preconditioners and emphasize the cases in which these Schur complement approximations are exact. Therefore, we mainly follow [42], but adapt the computations for our formulation which includes the electric field \mathbf{E} instead of a Lagrange multiplier r .

For the simplification of the outer Schur complement $\mathcal{S}^{(\mathbf{E},\mathbf{B})}$ we must find approximations for

$$(3.10) \quad \mathcal{K}_1 := \mathcal{D} - \mathcal{J}\mathcal{M}_{1,1}^{-1}\mathcal{G} \quad \text{and} \quad \mathcal{K}_2 := -(\tilde{\mathcal{J}} + \tilde{\mathcal{D}}_1 + \tilde{\mathcal{D}}_2)\mathcal{M}_{2,1}^{-1}\mathcal{G}.$$

Note that the first summand of $\mathcal{J}\mathcal{M}_{1,1}^{-1}\mathcal{G}$ is $\mathcal{J}\mathcal{M}_{\mathbf{E}}^{-1}\mathcal{G}$ which equals \mathcal{D} . Hence, \mathcal{K}_1 simplifies to the second summand of $\mathcal{J}\mathcal{M}_{\mathbf{E}}^{-1}\mathcal{G}$, i.e.,

$$(3.11) \quad \mathcal{K}_1 = -\mathcal{J}\mathcal{M}_{\mathbf{E}}^{-1} \left(\tilde{\mathcal{G}} - \frac{1}{\text{Re}_m} \mathcal{A} \right) \left(\mathcal{C} - \mathcal{A}^\top \mathcal{M}_{\mathbf{E}}^{-1} \left(\tilde{\mathcal{G}} - \frac{1}{\text{Re}_m} \mathcal{A} \right) \right)^{-1} \mathcal{A}^\top \mathcal{M}_{\mathbf{E}}^{-1} \mathcal{G}$$

which corresponds on a continuous level to

$$(3.12) \quad -S\mathbf{B}^n \times \left(\left(\delta \mathbf{u}^n \times \cdot - \frac{1}{\text{Re}_m} \mathbf{curl} \right) \left(\frac{\eta}{\Delta t} I - \frac{1}{\text{Re}_m} \Delta - \delta \mathbf{curl}(\mathbf{u}^n \times \cdot) \right)^{-1} \mathbf{curl}(\mathbf{u} \times \mathbf{B}^n) \right),$$

where \cdot denotes a placeholder for the input of the corresponding operators. Moreover, we have used $\delta \in \{0, 1\}$ to distinguish between the Picard ($\delta = 0$) and Newton ($\delta = 1$) linearizations. In the discrete counterpart (3.11), the matrix arising in the Picard iteration is made by dropping all terms with a tilde.

The continuous expression for \mathcal{K}_2 is given by

$$(3.13) \quad \delta S (\cdot \times \mathbf{E}^n + \cdot \times (\mathbf{u}^n \times \mathbf{B}^n) + \mathbf{B}^n \times (\mathbf{u}^n \times \cdot)) \left(\frac{\eta}{\Delta t} I - \frac{1}{\text{Re}_m} \Delta - \delta \mathbf{curl}(\mathbf{u}^n \times \cdot) \right)^{-1} \mathbf{curl}(\mathbf{u} \times \mathbf{B}^n).$$

Note that \mathcal{K}_2 vanishes for the Picard iteration.

3.1.1. The two-dimensional case. For the Picard linearization, expression (3.11) simplifies to \mathcal{D} in the stationary case. This follows immediately from the two-dimensional analogue of (2.20) and the identity [41]

$$(3.14) \quad \mathbf{curl}(-\Delta)^{-1} \mathbf{curl} \varphi = \varphi$$

which implies for our structure-preserving discretization that

$$(3.15) \quad \mathcal{A}(\mathcal{C} + \mathcal{A}^\top \mathcal{M}_{\mathbf{E}}^{-1} \mathcal{A})^{-1} \mathcal{A}^\top = \mathcal{M}_{\mathbf{E}}.$$

That means in the two-dimensional stationary case the outer Schur complement for the Picard iteration is exactly given by (3.8), i.e., the Navier–Stokes block with the linearized Lorentz force.

In the transient case, the Schur complement for the Picard linearization can no longer be calculated exactly. The behavior of the Schur complement now depends on which of the terms $\frac{1}{\Delta t} I$ and $\frac{1}{\text{Re}_m} \Delta$ dominates in (3.12). If $\frac{1}{\Delta t}$ is small in comparison to $\frac{1}{\text{Re}_m h^2}$, a good approximation of (3.12) is given, as in the stationary case, by $\tilde{\mathcal{S}}^{(\mathbf{E},\mathbf{B})}$. If $\frac{1}{\Delta t}$ dominates over $\frac{1}{\text{Re}_m h^2}$, (3.12) is approximately given by

$$(3.16) \quad S\mathbf{B}^n \times \left(\frac{1}{\text{Re}_m} \mathbf{curl} \left(\frac{1}{\Delta t} I \right)^{-1} \mathbf{curl}(\mathbf{u} \times \mathbf{B}^n) \right).$$

Hence, its magnitude can be approximated by $\frac{S\|\mathbf{B}^n\|^2\Delta t}{\text{Re}_m h^2} \ll 1$ for moderate coupling numbers and therefore we neglect this term by using the approximation $\begin{bmatrix} \mathcal{F} & \mathcal{B}^T \\ \mathcal{B} & \mathbf{0} \end{bmatrix}$ for the Schur complement in this case.

To also include the intermediate regime, we use the approximation of Phillips et al. [42] who suggest using (3.9). The expression for α interpolates between the above-mentioned dominating cases, since $\alpha \approx 0$ if $\frac{1}{\Delta t} \gg \frac{1}{\text{Re}_m h^2}$ and $\alpha \approx 1$ if $\frac{1}{\Delta t} \ll \frac{1}{\text{Re}_m h^2}$.

For the Newton linearization, the simplification of $\mathcal{S}^{(\mathbf{E}, \mathbf{B})}$ is not straightforward, but our numerical tests suggest that $\tilde{\mathcal{S}}^{(\mathbf{E}, \mathbf{B})}$ and $\tilde{\mathcal{S}}_\alpha^{(\mathbf{E}, \mathbf{B})}$ are acceptable preconditioners for $\mathcal{S}^{(\mathbf{E}, \mathbf{B})}$ in the stationary and transient cases, deteriorating only for high S and Re_m .

This can be explained by the fact that for small Re_m or Δt the terms $\frac{1}{\text{Re}_m} \mathbf{curl}$ and $\frac{\eta}{\Delta t} I - \frac{1}{\text{Re}_m} \Delta$ dominate over $\delta \mathbf{u}^n \times \cdot$ and $\delta \mathbf{curl}(\mathbf{u}^n \times \cdot)$ in (3.12). Remember that the terms that include a δ do not appear in the Picard iteration and were hence neglected in the previous derivation for the Picard iteration. Moreover, the term \mathcal{K}_2 is not included in our preconditioner for the Newton scheme which should deteriorate the performance for large S .

3.1.2. The three-dimensional case. The main difficulty in three dimensions is that the identity (3.14) no longer holds. Therefore, $\tilde{\mathcal{S}}^{(\mathbf{E}, \mathbf{B})}$ is not the exact outer Schur complement for the Picard linearization in the stationary case. In [42] the same approximation from the two-dimensional case is used in three dimensions. Based on the argument for the two-dimensional case in the previous subsection, we expect this approximation to work well when the term Δt dominates and to deteriorate in the other cases, especially in the stationary case. The three-dimensional performance of this preconditioner could be substantially improved with a better approximation of $\mathbf{curl} \Delta^{-1} \mathbf{curl}$ than a scaled identity.

We briefly comment on the main part of the outer Schur complement in the stationary case, given by

$$(3.17) \quad S\mathbf{B}^n \times [\mathbf{curl} \Delta^{-1} \mathbf{curl}(\mathbf{u} \times \mathbf{B}^n)].$$

As shown in [40, Chapter 4] one can rewrite $\mathbf{curl} \Delta^{-1} \mathbf{curl}$ as $I - \nabla \Delta_r^{-1} \nabla \cdot$, where Δ_r denotes a scalar Laplacian. These two representations show that the operator is the identity on divergence-free functions and maps curl-free functions to zero. Hence, this operator corresponds to the orthogonal L^2 -projection of a vector field onto its divergence-free part, which we denote by \mathbb{P} . Thus, the weak form of (3.17) block is given by

$$(3.18) \quad S(\mathbb{P}(\mathbf{u} \times \mathbf{B}^n), \mathbb{P}(\mathbf{v} \times \mathbf{B}^n)).$$

The key challenge is then to find a sparse approximation of (3.18). We do not further address this challenge here and focus instead on the outer Schur complement that eliminates the (\mathbf{u}_h, p_h) block.

3.2. Outer Schur complement eliminating the (\mathbf{u}_h, p_h) block. The outer Schur complement eliminating the (\mathbf{u}_h, p_h) block is given by

$$(3.19) \quad \mathcal{S}^{(\mathbf{u}, p)} = \begin{bmatrix} \mathcal{M}_{\mathbf{E}} & \tilde{\mathcal{G}} - \frac{1}{\text{Re}_m} \mathcal{A} \\ \mathcal{A}^\top & \mathcal{C} \end{bmatrix} - \begin{bmatrix} \mathcal{G} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{bmatrix} \begin{bmatrix} \mathcal{F} + \mathcal{D} & \mathcal{B}^\top \\ \mathcal{B} & \mathbf{0} \end{bmatrix}^{-1} \begin{bmatrix} \mathcal{J} & \tilde{\mathcal{J}} + \tilde{\mathcal{D}}_1 + \tilde{\mathcal{D}}_2 \\ \mathbf{0} & \mathbf{0} \end{bmatrix}.$$

Formula (3.1) implies that the outer Schur complement for the Newton iteration is given by

$$(3.20) \quad \mathcal{S}^{(\mathbf{u},p)} = \begin{bmatrix} \mathcal{M}_{\mathbf{E}} - \mathcal{G}\mathcal{N}_{1,1}^{-1}\mathcal{J} & \tilde{\mathcal{G}} - \frac{1}{\text{Re}_m}\mathcal{A} - \mathcal{G}\mathcal{N}_{1,1}^{-1}(\tilde{\mathcal{J}} + \tilde{\mathcal{D}}_1 + \tilde{\mathcal{D}}_2) \\ \mathcal{A}^\top & \mathcal{C} \end{bmatrix},$$

where

$$(3.21) \quad \mathcal{N}_{1,1}^{-1} = (\mathcal{F} + \mathcal{D})^{-1} - (\mathcal{F} + \mathcal{D})^{-1}\mathcal{B}^\top(-\mathcal{B}(\mathcal{F} + \mathcal{D})^{-1}\mathcal{B}^\top)^{-1}\mathcal{B}(\mathcal{F} + \mathcal{D})^{-1}.$$

For this strategy, further simplifications of the Picard or Newton linearizations are not straightforward. Our numerical results in the next section show that

$$(3.22) \quad \tilde{\mathcal{S}}^{(\mathbf{u},p)} = \begin{bmatrix} \mathcal{M}_{\mathbf{E}} & \tilde{\mathcal{G}} - \frac{1}{\text{Re}_m}\mathcal{A} \\ \mathcal{A}^\top & \mathcal{C} \end{bmatrix}$$

works very well as a preconditioner for both schemes. Indeed, in contrast to the previous order of elimination, this approximation works qualitatively the same in two and three dimensions. We expect the approximation to deteriorate in the stationary case for very high Re_m , since the missing term $-\mathcal{G}\mathcal{N}_{1,1}^{-1}(\tilde{\mathcal{J}} + \tilde{\mathcal{D}}_1 + \tilde{\mathcal{D}}_2)$ in the Schur complement approximation gains more influence in comparison to $-\frac{1}{\text{Re}_m}\mathcal{A}$. We also observe this numerically in the next section.

However, a good approximation of the outer Schur complement is maintained for high coupling numbers S , which will clearly be seen in our numerical results in section 4.4.1 below. This behavior is perhaps explained by the fact that $\mathcal{N}_{1,1}^{-1}$ also includes a factor S in the inverse of $(\mathcal{F} + \mathcal{D})$, which cancels the factor of S in the matrices $\mathcal{J}, \tilde{\mathcal{J}}, \tilde{\mathcal{D}}_1$, and $\tilde{\mathcal{D}}_2$.

To use these block preconditioners in practice, we must develop robust preconditioners for the electromagnetic and hydrodynamic subsystems.

3.3. Parameter-robust relaxation. The equations in the hydrodynamic and electromagnetic blocks become difficult to solve in the parameter regimes of interest at high Reynolds and coupling numbers due both to the nonsymmetric linearized advection and Lorentz force terms, and the addition of the symmetric positive semi-definite (SPSD) augmented Lagrangian terms. Standard multigrid methods are known to perform poorly for these kinds of problems.

The key components for a robust multigrid method for the SPSPD augmented Lagrangian terms are a parameter-robust relaxation method, that efficiently damps error modes in the kernel of the singular operators, and a kernel-preserving prolongation operator, as revealed in the seminal work of Schöberl [46]. The nonsymmetric terms are more troublesome, but numerical results have shown [24, 25] that subspace correction methods can still perform well for the Navier–Stokes equations at high Reynolds numbers.

A recent summary of the theory of robust relaxation methods can be found in [23]. Briefly, we consider the multigrid relaxation methods in the framework of subspace correction methods [56]. These decompose a (finite-dimensional) trial space V as

$$(3.23) \quad V = \sum_i V_i,$$

where the sum is not necessarily direct. The parallel subspace correction method applied to a linear variational problem $a(u, v) = (f, v) \quad \forall v \in V$ computes for an

initial guess u^k a correction δu_i to the error $e = u - u^k$ in each subspace V_i by solving

$$(3.24) \quad a(\delta u_i, v_i) = (f, v_i) - a(u^k, v_i) \text{ for all } v_i \in V_i,$$

and sets $u^{k+1} = u^k + \sum_i w_i \delta u_i$ for damping parameters w_i . A rigorous statement regarding the properties the decomposition (3.23) and the considered bilinear form a must fulfill to yield a robust relaxation method can be found in [46, Thm. 4.1]. A key property is that the kernel \mathcal{N} of the SPSD terms is decomposed over the subspaces, i.e.,

$$(3.25) \quad \mathcal{N} = \sum_i (V_i \cap \mathcal{N}).$$

This property means that it must be possible to write any kernel function as the sum of kernel functions in the subspaces V_i . This implies that the subspaces V_i must be at least rich enough to support nonzero kernel functions. The choice of the space decomposition (3.23) is often made by consideration of the discrete Hilbert complexes underpinning the discretization.

3.4. Solver for the hydrodynamic block. In order to implement the block factorization (3.1) as the outer preconditioner, we need a solver for the Navier–Stokes subsystem. To do this, we will apply ideas of parameter-robust multigrid relaxation described in section 3.3, albeit without a theoretical guarantee of success. The variational statement of the PDE we wish to solve is

$$(3.26) \quad \begin{aligned} \frac{2}{\text{Re}}(\varepsilon(\mathbf{u}), \varepsilon(\mathbf{v})) + (\mathbf{u}^n \cdot \nabla \mathbf{u}, \mathbf{v}) + (\mathbf{u} \cdot \nabla \mathbf{u}^n, \mathbf{v}) + \gamma(\text{div } \mathbf{u}, \text{div } \mathbf{v}) \\ + S(\mathbf{B}^n \times (\mathbf{u} \times \mathbf{B}^n), \mathbf{v}) - (p, \text{div } \mathbf{v}) = (\mathbf{f}, \mathbf{v}) \quad \forall \mathbf{v} \in \mathbf{H}_0^1(\Omega), \\ -(\text{div } \mathbf{u}, q) = 0 \quad \forall q \in L^2(\Omega). \end{aligned}$$

This corresponds to the standard Newton linearization of the Navier–Stokes equations with an augmented Lagrangian term, plus the linearization of the Lorentz force \mathcal{D} . We follow the approach of [29, 25, 24] to solve this system. The first idea is to use the augmented Lagrangian term $-\gamma \nabla \text{div } \mathbf{u}$ to approximate the inner Schur complement of the hydrodynamic block by choosing a large γ , e.g., $\gamma \approx 10^4$. One can show [8, Thm. 3.2] that the inner Schur complement of the augmented system $\tilde{\mathcal{S}}_{\text{NS}}$ satisfies

$$(3.27) \quad \tilde{\mathcal{S}}_{\text{NS}}^{-1} = \mathcal{S}_{\text{NS}}^{-1} - \gamma \mathcal{M}_p^{-1},$$

where $\mathcal{S}_{\text{NS}}^{-1}$ denotes the Schur complement of the system without the augmented Lagrangian term and \mathcal{M}_p denotes the pressure mass matrix. Therefore, for large γ the pressure mass matrix scaled by $-1/(1/\text{Re} + \gamma)$ is a good approximation for $\tilde{\mathcal{S}}_{\text{NS}}$. As the discretization considered in this work uses discontinuous pressures, the pressure mass matrix is block-diagonal and hence directly invertible. In the transient case $\mathcal{S}_{\text{NS}}^{-1}$ can be further approximated by the inverse of the stationary Schur complement plus an extra term $-\Delta t L_p^{-1}$ [28], where L_p corresponds to the Poisson problem for p with Neumann boundary conditions. In our numerical examples this extra term makes little difference as we only consider time steps $\frac{1}{\Delta t} \ll \gamma$, and we therefore neglect it.

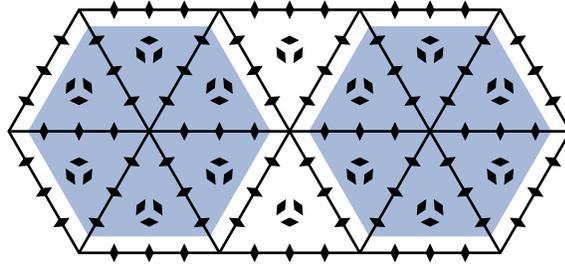


FIG. 3.1. Star patch for \mathbf{BDM}_2 -elements.

Since the augmented Lagrangian term has a large kernel that consists of all solenoidal vector fields, a robust multigrid scheme as described in section 3.3 must be used to solve the augmented momentum block. For the $\mathbf{H}(\text{div}) \times L^2$ -conforming discretization the *star iteration* [25, section 4] can be used as a robust relaxation method. The subspace decomposition is defined as

$$(3.28) \quad \mathbf{V}_i = \{\mathbf{v} \in \mathbf{V}_h : \text{supp}(\mathbf{v}) \subset K_i\},$$

where K_i is the patch of elements sharing the vertex i in the mesh. Example patches are shown in Figure 3.1. Since we use a structure-preserving discretization, the properties of the de Rham complexes (2.12) and (2.13) imply that (3.28) fulfils the kernel decomposition property (3.25). This property was also used in [6] to construct a robust smoother for the $\mathbf{H}(\text{div})$ and $\mathbf{H}(\text{curl})$ Riesz maps and in [29] for the Stokes equations. In this case we may employ the standard prolongation operator induced by the finite element discretization, because the uniformly refined mesh hierarchy we consider is nested.

The velocity block further includes terms given by the convection-diffusion term $(\mathbf{u} \cdot \nabla)\mathbf{u}$, the linearization of the Lorentz force $S\mathbf{B}^n \times (\mathbf{u} \times \mathbf{B}^n)$, and the stabilization term (2.15). Numerical experiments in [24] and in the next section 4 show that these terms only degrade the performance of the preconditioner at high Reynolds and coupling numbers. As we have mentioned before, these somewhat surprising numerical observations are not backed up by theory since these terms do not fit in the framework of section 3.3, and applying geometric multigrid methods to problems with strong advection typically requires special care. The kernel of the stabilization $\mathcal{ST}(\mathbf{u}, \mathbf{v})$ consists of all C^1 vector fields. Therefore, the stabilization term slightly degrades the performance of the solver, but the impact is not very significant as the factor $\mu h_{\partial K}^2$ is small.

3.5. Solver for the electromagnetic block. The weak formulation of the electromagnetic block is given by

$$(3.29) \quad (\mathbf{E}, \mathbf{F}) - \frac{1}{\text{Re}_m} (\mathbf{B}, \text{curl } \mathbf{F}) + \delta (\mathbf{u}^n \times \mathbf{B}, \mathbf{F}) = 0 \quad \forall \mathbf{F} \in \mathbf{H}_0(\text{curl}, \Omega),$$

$$\frac{\eta}{\Delta t} (\mathbf{B}, \mathbf{C}) + (\text{curl } \mathbf{E}, \mathbf{C}) + \frac{1}{\text{Re}_m} (\text{div } \mathbf{B}, \text{div } \mathbf{C}) = (\mathbf{f}, \mathbf{C}) \quad \forall \mathbf{C} \in \mathbf{H}_0(\text{div}, \Omega).$$

Recall that $\eta, \delta \in \{0, 1\}$ distinguish between the stationary ($\eta = 0$) and transient ($\eta = 1$) cases and the Picard ($\delta = 0$) and Newton ($\delta = 1$) linearizations. Eliminating \mathbf{E} , this corresponds to a mixed formulation of

$$\begin{aligned}
 (3.30) \quad & \frac{\eta}{\Delta t} \mathbf{B} + \frac{1}{\text{Re}_m} (\mathbf{curl} \mathbf{curl} \mathbf{B} - \nabla \text{div} \mathbf{B}) + \delta \mathbf{curl}(\mathbf{u}^n \times \mathbf{B}) = \mathbf{f} \text{ in } \Omega, \\
 & \mathbf{B} \cdot \mathbf{n} = 0 \text{ on } \partial\Omega, \\
 & \frac{1}{\text{Re}_m} \mathbf{curl} \mathbf{B} - \delta \mathbf{u}^n \times \mathbf{B} = \mathbf{0} \text{ on } \partial\Omega.
 \end{aligned}$$

For the Picard linearization, this problem simplifies to the mixed formulation for the standard vector Laplace problem [7] with boundary conditions $\mathbf{B} \cdot \mathbf{n} = \mathbf{curl} \mathbf{B} = \mathbf{0}$ on $\partial\Omega$. Chen et al. [17] propose a Schur complement solver and Arnold, Falk, and Winther [7, section 10] propose a norm-equivalent block-diagonal preconditioner for the mixed formulation. We also found that the star multigrid solver applied monolithically to the electromagnetic block (3.30) results in an efficient solver and employ this solver in our numerical examples. All of the solvers described show Re_m -robust performance.

In contrast, the presence of the additional term $\mathbf{curl}(\mathbf{u}^n \times \mathbf{B})$ in the Newton linearization, which has a nontrivial kernel, makes the problem almost singular for high Re_m in the stationary case and hence requires a special multigrid method. Unfortunately the troublesome term $\mathbf{curl}(\mathbf{u}^n \times \mathbf{B})$ is not symmetric and thus does not fit the available analytical framework of Schöberl. Our considerations on this point are therefore necessarily heuristic. Some insight may be gained by employing the vector identity

$$(3.31) \quad \mathbf{curl}(\mathbf{A} \times \mathbf{B}) = (\mathbf{B} \cdot \nabla) \mathbf{A} - (\mathbf{A} \cdot \nabla) \mathbf{B} + \mathbf{A}(\nabla \cdot \mathbf{B}) - \mathbf{B}(\nabla \cdot \mathbf{A})$$

to rewrite (3.30) to

$$(3.32) \quad \frac{\eta}{\Delta t} \mathbf{B} - \frac{1}{\text{Re}_m} \Delta \mathbf{B} - (\mathbf{B} \cdot \nabla) \mathbf{u}^n + (\mathbf{u}^n \cdot \nabla) \mathbf{B} - \mathbf{u}^n (\nabla \cdot \mathbf{B}) - \mathbf{B} (\nabla \cdot \mathbf{u}^n).$$

The last term $-\mathbf{B}(\nabla \cdot \mathbf{u}^n)$ vanishes since we exactly enforce $\nabla \cdot \mathbf{u}^n = 0$ in each step. The terms $-(\mathbf{B} \cdot \nabla) \mathbf{u}^n + (\mathbf{u}^n \cdot \nabla) \mathbf{B}$ are reminiscent of the Newton linearization of the advection term $(\mathbf{u} \cdot \nabla) \mathbf{u}$ of the Navier–Stokes equation, for which it has been demonstrated that a star multigrid method is effective [25]. Numerical experiments with these approaches applied monolithically do indeed yield a robust solver for the stationary and transient cases in 2D, and in the transient case in three dimensions for sufficiently small Δt . We observe in our numerical tests that in three dimensions the solver breaks down for $\text{Re}_m \approx 700$ for a stationary lid-driven cavity problem.

4. Numerical results. In this section, we present numerical results for the Picard and Newton linearization described in the previous sections. We investigate three test problems: the stationary and transient version of a lid-driven cavity problem and a transient island-coalescence problem. The numerical results were produced on ARCHER2, the UK national supercomputer, which consists of 5,860 compute nodes each built of two AMD Zen2 7742 processors with 64 2.25 GHz cores and 256 GB of memory.

4.1. Algorithm details. The algorithm is implemented in Firedrake [43] and uses the solver packages PETSc [9] and PCPATCH [23]. The latter includes the implementation of the multigrid relaxation method described in section 3. It is well known that the convergence of the nonlinear scheme depends heavily on the initial guess and might fail to converge for high Reynolds numbers with poor initial guesses. To circumvent this problem we perform continuation in the Reynolds numbers and

coupling number, for the stationary problems. In the presented tables we always apply continuation to the variable in the column first and use each solution as the starting point for the continuation over the rows. We use the steps 1, 100, 1,000, 5,000, 10,000 for S and 1,500, 1,000, 3,000, 5,000, 7,000, 9,000, 10,000 for Re and Re_m . The reported nonlinear and linear iteration numbers correspond to the final solve in the continuation; however, the extra cost for the continuation should be kept in mind for stationary problems. For time-dependent problems, continuation is not necessary.

We use flexible GMRES [45] as the outermost Krylov solver since we apply GMRES in the multigrid relaxation. Moreover, we apply a block upper triangular preconditioner [10]

$$(4.1) \quad \mathcal{P} = \begin{pmatrix} \mathcal{I} & -\tilde{\mathcal{M}}^{-1}\mathcal{K} \\ \mathbf{0} & \mathcal{I} \end{pmatrix} \begin{pmatrix} \tilde{\mathcal{M}}^{-1} & \mathbf{0} \\ \mathbf{0} & \tilde{\mathcal{S}}^{-1} \end{pmatrix}$$

to (2.18), where we denoted (2.18) here as $\begin{pmatrix} \mathcal{M} & \mathcal{K} \\ \mathcal{L} & \mathcal{N} \end{pmatrix}$. We also investigated a full block-LDU preconditioner without notable improvements in terms of iteration counts, which fits with the recent theoretical results in [51].

Both the block matrix \mathcal{M} and the outer Schur complement approximation $\tilde{\mathcal{S}}^{(\mathbf{u},p)}$ are inverted approximately with two iterations of preconditioned FGMRES (denoted $\tilde{\mathcal{M}}^{-1}$ and $\tilde{\mathcal{S}}^{-1}$, respectively). The former uses the block preconditioner for the hydrodynamic block described in section 3.4, the latter the monolithic multigrid method described in section 3.5. In the numerical results we focus on taking the outer Schur complement that eliminates the hydrodynamic block, except for one case in section 4.4.1. In both multigrid methods we use six preconditioned GMRES iterations as the smoother on each level and the direct solver MUMPS [3] to solve the problem on the coarsest grid. Since this relaxation is quite expensive, convergence in a very small number of outer iterations is required for efficiency. See Figure 4.1 for a graphical representation of the solver.

We have chosen relative and absolute tolerances of 10^{-10} and 10^{-6} for the nonlinear solver and 10^{-7} and 10^{-7} for the outermost linear solver, measured in the Euclidean norm. We use the $\mathbf{H}(\text{div}) \times L^2$ -conforming elements $\mathbb{BDM}_2 \times \mathbb{DG}_1$ for (\mathbf{u}_h, p_h) . Moreover, we apply $\mathbb{CG}_2 \times \mathbb{RT}_2$ elements for (E_h, \mathbf{B}_h) in 2D and $\mathbb{NED}1_2 \times \mathbb{RT}_2$ elements for $(\mathbf{E}_h, \mathbf{B}_h)$ in 3 dimensions. When we consider a manufactured solution we always subtract $\int_{\Omega} p \, dx$ from the pressure to fix the average of p to be zero.

For time-dependent problems, we apply the second order, L-stable BDF2 method with a fixed time step. We compute the first time step with Crank–Nicolson to provide the second starting value for BDF2. For the transient lid-driven cavity problem, we use a time step of $\Delta t = 0.01$ and a final time of $T = 0.1$. We did not choose a higher final time T because of budget limitations. However, we confirmed that the reported numbers are representative for higher T by computing the solution for a few parameters until $T = 1$ without noticeable changes in the iteration counts. Moreover, we confirm the efficiency for more time steps in the island coalescence problem where we iterate in the finest run until $T = 15$ in 2400 time steps.

4.2. Interpolating boundary data. The theory from the previous sections has been formulated for homogeneous boundary conditions, but the generalization is straightforward for nonhomogeneous boundary conditions. However, there is a subtle technicality in the implementation if one wants to enforce the divergence constraint $\text{div } \mathbf{B}_h = 0$ pointwise. Strong boundary conditions are enforced in a finite element code by interpolating the given boundary data onto the corresponding finite element space. If the interpolation of the boundary values \mathbf{g} were exact, identity (2.14) would

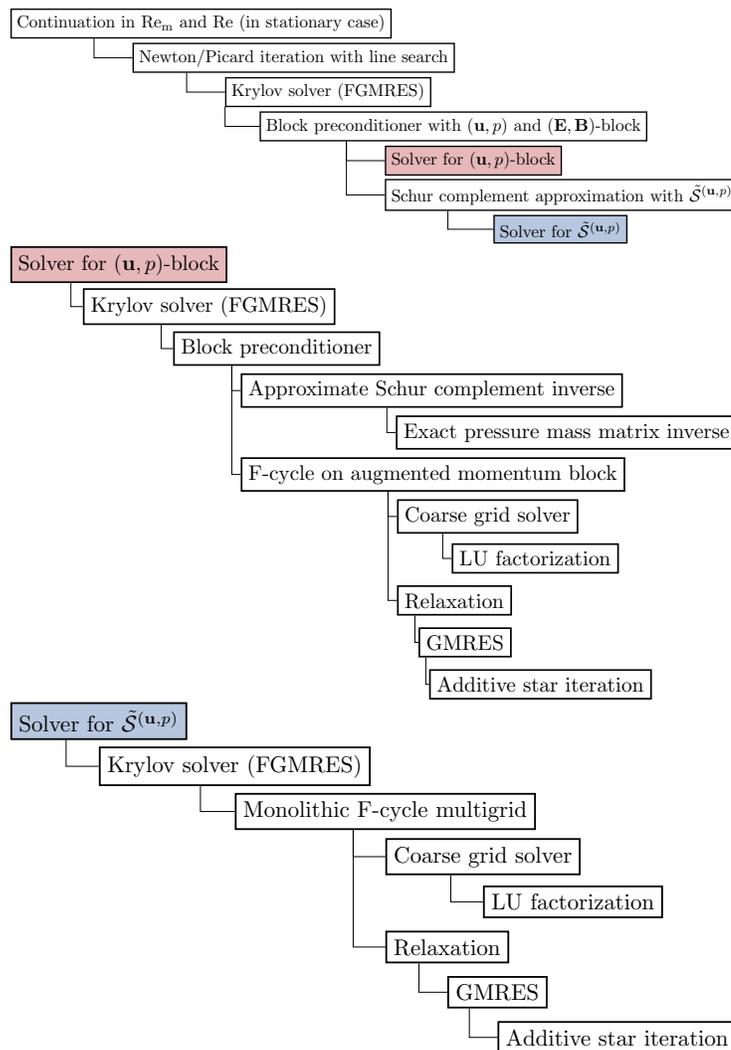


FIG. 4.1. Graphical outline of the solver.

imply that $\text{div } \mathbf{B}_h = 0$ holds. However, the degrees of freedom for the interpolation are moments and are usually implemented by a quadrature rule whose quadrature degree is based on the polynomial degree of the finite element space. If \mathbf{g} is a nonpolynomial expression, this quadrature rule might not interpolate the boundary condition exactly and therefore one loses the property that $\text{div } \mathbf{g}_h = 0$ on $\partial\Omega$ holds exactly.

To circumvent this problem we use high order quadrature rules for the evaluation of the degrees of freedom to ensure that the interpolation is exact up to machine precision. In Figure 4.2 we have illustrated the effect of the quadrature degree on the enforcement of the divergence constraint. We have used the method of manufactured solutions for a smooth problem to compute $\|\text{div } \mathbf{B}_h\|_0$ for different quadrature degrees. Moreover, we have plotted the L^2 -norm over $\partial\Omega$ of the interpolation of the divergence-free function \mathbf{B} into the \mathbb{RT}_2 space. One can clearly observe that a quadrature degree of 2 for \mathbb{RT}_2 elements is not sufficient to enforce $\text{div } \mathbf{B}_h = 0$ up to machine

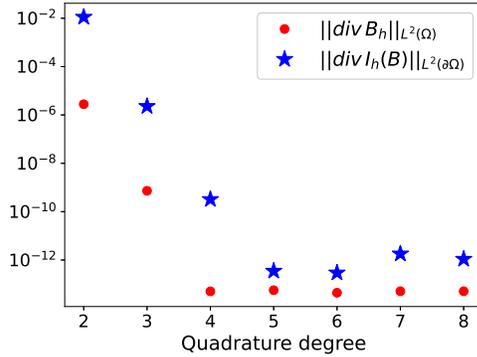


FIG. 4.2. L^2 -norm of the divergence of the solution \mathbf{B}_h and the interpolant of the boundary condition for different quadrature degrees in the evaluation of the degrees of freedom for the Raviart–Thomas space.

precision. A higher quadrature degree preserves the divergence of the boundary data more accurately and leads to the pointwise enforcement of $div \mathbf{B}_h = 0$.

4.3. Two-dimensional results.

4.3.1. Stationary lid-driven cavity in 2D. First, we consider a lid-driven cavity problem posed over $\Omega = (-1/2, 1/2)^2$ for a background magnetic field $\mathbf{B}_0 = (0, 1)^\top$ which determines the boundary conditions $\mathbf{B} \cdot \mathbf{n} = \mathbf{B}_0 \cdot \mathbf{n}$ on $\partial\Omega$, and set $\mathbf{f} = \mathbf{0}$ [36]. We impose the boundary condition $\mathbf{u} = (1, 0)^\top$ at the boundary $y = 0.5$ and homogeneous boundary conditions elsewhere. The problem models the flow of a conducting fluid driven by the movement of the lid at the top of the cavity. The magnetic field imposed orthogonal to the lid creates a Lorentz force that perturbs the flow of the fluid.

For the multigrid hierarchy we use a coarse mesh of 16×16 cells and six levels of refinement in 2D resulting in a 1024×1024 grid with 73.4 million DoFs. For both linearizations we observe fairly constant Krylov iteration counts for Re and S in the range of 1 to 10,000 in Table 4.1. In terms of the nonlinear convergence, the Picard linearization sometimes takes slightly more iterations than the Newton linearizations, with slightly better linear iteration numbers.

As mentioned earlier, our scheme does not include a stabilization for high magnetic Reynolds numbers. However, we have verified that our solutions do not exhibit oscillations in this regime. A plot of the streamlines for different Re and Re_m can be found in Figure 4.3. One can clearly observe the phenomenon that for high magnetic Reynolds numbers the magnetic field lines are advected with the fluid flow. Iteration counts for varying Re and Re_m are displayed in Table 4.2.

TABLE 4.1

(Nonlinear iterations) Average outer Krylov iterations per nonlinear step for the stationary lid-driven cavity problem in 2D.

$S \backslash Re$	Picard			Newton		
	1	1,000	10,000	1	1,000	10,000
1	(3) 5.3	(4) 3.5	(3) 4.3	(2) 6.5	(4) 3.5	(3) 4.3
1,000	(4) 3.5	(3) 4.7	(2) 8.5	(2) 5.5	(3) 4.7	(2) 6.5
10,000	(3) 5.0	(3) 4.3	(2) 7.0	(2) 6.5	(2) 6.0	(2) 7.0

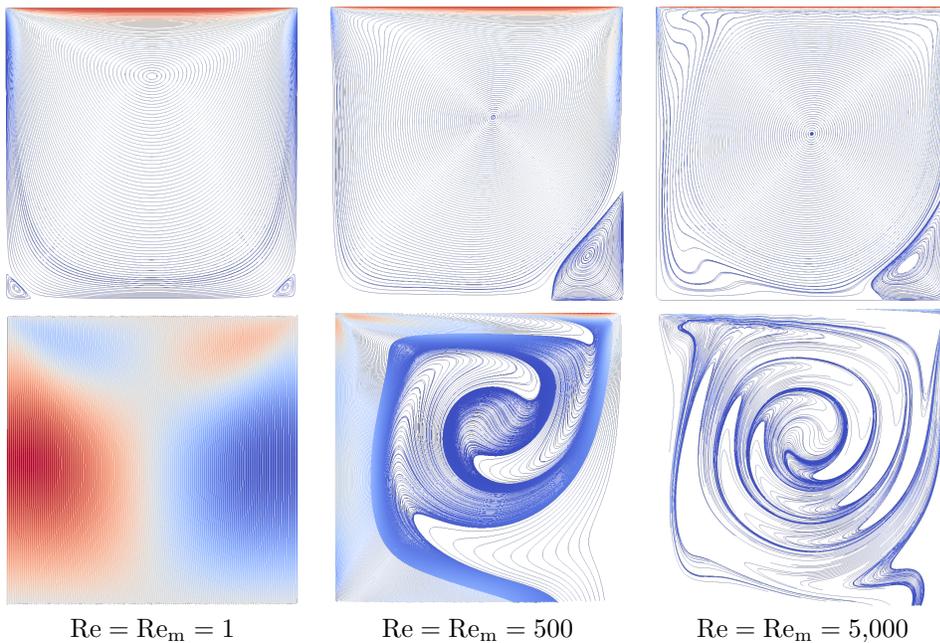


FIG. 4.3. Streamlines for the two-dimensional stationary lid-driven cavity problem for \mathbf{u} (upper row) and \mathbf{B} (lower row).

TABLE 4.2

Iteration counts for the stationary lid-driven cavity problem in 2D with $\mathbf{H}(\text{div}) \times L^2$ -discretization for different Re_m and Re .

$\text{Re}_m \backslash \text{Re}$	Picard			Newton		
	1	1,000	10,000	1	1,000	10,000
1	(3) 5.3	(4) 3.5	(3) 4.3	(2) 6.0	(3) 4.3	(3) 4.3
1,000	-	-	-	(2) 4.5	(3) 3.0	(3) 3.0
10,000	-	-	-	(2) 4.5	(4) 5.5	(3) 5.7

For the Picard linearization we observe that the nonlinear scheme already fails to converge for a magnetic Reynolds number of 100. The poor nonlinear convergence of the Picard iteration for high Re_m even with continuation was previously observed for other formulations [40, 41].

For the Newton linearization the linear iterations increase slightly since the approximation of the Schur complement $\mathcal{S}_N^{(\mathbf{u}, p)}$ by $\tilde{\mathcal{S}}^{(\mathbf{u}, p)}$ becomes less accurate for high Reynolds numbers. On the other hand, the number of nonlinear iterations remains fairly constant which seems to indicate that the linear solver for the $(\mathbf{E}_h, \mathbf{B}_h)$ block described in section 3.5 works very well for high Re_m in 2D.

4.3.2. Scott–Vogelius discretization for (\mathbf{u}_h, p_h) . Thus far we have considered an $\mathbf{H}(\text{div}) \times L^2$ discretization for the hydrodynamic variables. For comparison, in this subsection we include results for Scott–Vogelius elements [48], i.e., $(\text{CG}_k)^d \times \text{DG}_{k-1}$ elements. A fluid-Reynolds-robust preconditioner for this element pair was recently developed in [24]. While this conforming discretization does not require stabilization terms to weakly enforce continuity, it is only stable on certain types

TABLE 4.3

Iteration counts for the stationary lid-driven cavity problem in 2D with Scott–Vogelius elements for different Re_m and Re .

$Re_m \backslash Re$	Picard			Newton		
	1	1,000	10,000	1	1,000	10,000
1	(2) 4.0	(2) 2.5	(2) 8.5	(2) 4.0	(2) 2.5	(3) 9.7
1,000	-	-	-	(5) 1.8	(3) 3.0	(2) 4.0
10,000	-	-	-	(8) 5.2	(4) 6.2	(2) 5.5

of meshes. For this reason, the mesh hierarchy is barycentrically refined, and the specialized multigrid method of [24] exploits this structure. This barycentric refinement ensures stability for polynomial order $k = d$ [58].

The results are shown in Table 4.3. We observe that the Krylov iteration counts are in general similar for the Scott–Vogelius element, making this an attractive alternative for those wishing to employ conforming schemes. However, one must keep in mind that the work per Krylov iteration is substantially higher for this element, due to the use of larger patches in the so-called macrostar relaxation method.

To summarize the two-dimensional stationary results, both schemes considered provide Reynolds-robust solvers and also perform very well for large coupling numbers S . The Picard iteration is unsuitable for high magnetic Reynolds number because the nonlinear iteration fails to converge. The Newton scheme performs well for high Re_m with good nonlinear convergence and a slight increase in the linear iteration numbers.

4.3.3. Time-dependent lid-driven cavity problem in 2D. We next consider the time-dependent lid-driven cavity problem. We choose the same boundary conditions and right-hand side as in the stationary case. The numerical results for varying S and Re are shown in Table 4.4. As in the stationary case, the Krylov iteration counts remain almost constant for the two linearizations. We notice that the Picard iteration fails to converge for high S and Re_m for the chosen $\Delta t = 0.01$. However, we tested that one can get the Picard iteration to converge in most cases by choosing a smaller Δt in the first time steps. We do not report these results here to keep the tables consistent.

Table 4.4 also shows iteration counts for varying Re and Re_m . The linear solver is robust for most parameter values, with iteration counts only increasing for $Re = 1$ and $Re_m = 100,000$.

For completeness, we also study the case of high Re_m and S at the same time in Table 4.4, which we expect to be the most challenging case. Again the slight increase of the Krylov iterations in the Newton iteration is due to inaccurate outer Schur complement approximation. However, the solvers perform very well, considering the difficulty of the problem.

4.3.4. Time-dependent island-coalescence problem in 2D. Next, we consider a two-dimensional island-coalescence problem to demonstrate the effectiveness of our method for a physically relevant model that shows behavior which is unique to MHD problems. Furthermore, we report results for a weak scalability test with up to 256 processors and 18.3M DoFs to examine the performance of our algorithm.

The island-coalescence problem is used to model magnetic reconnection processes in large aspect ratio tokamaks. For a strong magnetic field in the toroidal direction, the flow can be described in a two-dimensional setting by considering a cross section of the tokamak. We consider the same problem as in [1, sect. 4.2]. The domain

TABLE 4.4
Iteration counts for the transient lid-driven cavity problem in 2D.

	Picard			Newton		
$S \setminus \text{Re}$	1	10,000	100,000	1	10,000	100,000
1	(2.0) 3.0	(2.2) 3.6	(3.1) 3.3	(1.6) 3.6	(2.2) 3.6	(3.1) 3.3
1,000	(3.0) 4.0	(3.0) 3.3	(2.7) 3.0	(2.1) 4.7	(2.2) 3.9	(2.3) 3.3
10,000	-	-	-	(2.2) 6.5	(2.5) 5.0	(2.3) 5.6

	Picard			Newton		
$\text{Re}_m \setminus \text{Re}$	1	10,000	100,000	1	10,000	100,000
1	(2.0) 3.0	(2.2) 3.6	(3.1) 3.3	(1.6) 3.6	(2.2) 3.6	(3.1) 3.3
10,000	-	-	-	(2.0) 3.1	(2.3) 3.6	(3.1) 3.3
100,000	-	-	-	(2.2)10.9	(3.0) 3.2	(3.3) 3.2

	Picard			Newton		
$\text{Re}_m \setminus S$	1	1,000	10,000	1	1,000	10,000
1	(2.0) 3.0	(3.0) 4.0	-	(1.6) 3.6	(2.1) 4.7	(2.2) 6.5
1,000	(3.0) 2.5	-	-	(2.0) 3.1	(2.2) 5.6	(2.8)11.0
10,000	-	-	-	(2.0) 3.1	(2.2) 6.3	(3.1)11.8

$\Omega = (-1, 1)^2$ results from the unfolding of an annulus in the cross-sectional direction where the left and right edges are mapped periodically. The equilibrium solution for $k = 0.2$ is given by

$$\mathbf{u}_{eq} = \mathbf{0}, \quad p_{eq}(x, y) = \frac{1 - k^2}{2} \left(1 + \frac{1}{(\cosh(2\pi y) + k \cos(2\pi x))^2} \right),$$

$$\mathbf{B}_{eq}(x, y) = \frac{1}{\cosh(2\pi y) + k \cos(2\pi x)} \begin{pmatrix} \sinh(2\pi y) \\ k \sin(2\pi x) \end{pmatrix}, \quad E_{eq} = \frac{1}{\text{Re}_m} \text{curl } \mathbf{B}_{eq} - u_{eq} \times B_{eq},$$

which results in right-hand sides $\mathbf{f} = \mathbf{0}$ and \mathbf{g} for (1.1c) given by

$$(4.2) \quad \mathbf{g} = \frac{-8\pi^2(k^2 - 1)}{\text{Re}_m(\cosh(2\pi y) + k \cos(2\pi x))^3} \begin{pmatrix} \sinh(2\pi y) \\ k \sin(2\pi x) \end{pmatrix}.$$

The initial condition for \mathbf{B}_{eq} is given by perturbing it for $\varepsilon = 0.01$ with

$$(4.3) \quad \Delta \mathbf{B} = \frac{\varepsilon}{\pi} \begin{pmatrix} -\cos(\pi x) \sin(\pi y/2) \\ 2 \cos(\pi y/2) \sin(\pi x) \end{pmatrix}.$$

The authors believe that the reported $\Delta \mathbf{B}$ in [1] includes a typographical error, as it is not divergence-free, and amended the second component appropriately. Therefore, the problem setup is not exactly identical and hence we might see slightly different solutions. The reconnection rate can be computed as the difference between $\text{curl } \mathbf{B}$ evaluated at $(0, 0)$ at the current time and the initial time, divided by $\sqrt{\text{Re}_m}$. Note

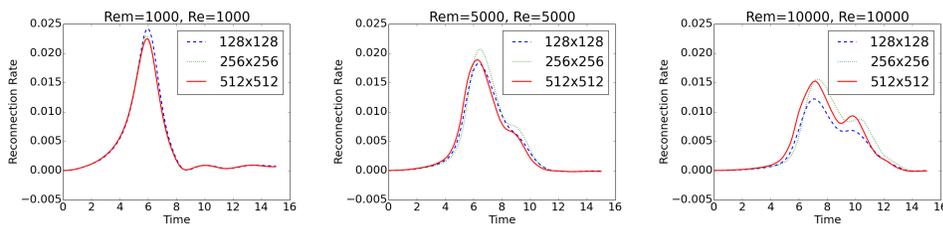


FIG. 4.4. Reconnection rates for the island-coalescence problem.

that in our formulation $\mathbf{B} \in \mathbf{H}(\text{div}, 0)$ and therefore we apply the curl weakly by solving a problem for $j_0 \in \mathbf{H}_0(\text{curl}, \Omega)$ such that

$$(4.4) \quad (j_0, k) = (\mathbf{B}, \text{curl } k) \quad \forall k \in \mathbf{H}_0(\text{curl}, \Omega).$$

In order to make the point evaluation of j_0 at $(0,0)$ well-posed we project j_0 to the space $\mathbb{CG}1$ as in [1].

Figure 4.4 shows the reconnection rates for $\text{Re} = \text{Re}_m = 1,000$, $\text{Re} = \text{Re}_m = 5,000$, and $\text{Re} = \text{Re}_m = 10,000$ for three different spatial and temporal resolutions. We have fixed a coarse grid of 16×16 cells and compute results for three (1.1M DoFs), four (4.6M DoFs), and five (18.4M DoFs) levels of refinement. For the three levels of refinement, we chose $\Delta t = 0.025$ and halved it with each refinement. We iterated until $T = 15$ which results in 2400 time steps for the finest resolution.

One can observe a decreasing height of the peak for increasing Reynolds numbers and the so-called “sloshing” [34] effect that results in further peaks after the main peak with higher Reynolds numbers. Convergence for our considered meshes can be observed for $\text{Re} = \text{Re}_m = 1,000$ and $\text{Re} = \text{Re}_m = 5,000$ while a further refinement is needed for $\text{Re} = \text{Re}_m = 10,000$. Nevertheless, our finest grid results match the results of [1, Figure 4] where finer meshes of up to 2560×2560 cells and $\Delta t = 0.0016$ have been considered. For example, our finest result for $\text{Re} = \text{Re}_m = 10,000$ clearly reproduces the second peak in the reconnection rate.

Furthermore, we performed a weak parallel scaling test for nine different combinations of the Reynolds numbers. We chose a coarse grid of 16×16 cells with three (1.1M DoFs), four (4.6M DoFs), and five (18.4M DoFs) levels of refinement. All tests were performed with 16 cores per node on 1, 4, and 16 nodes resulting in 16, 64, and 256 cores for the different refinements. We observed (not reported here) that scaling over the nodes with a fixed number of cores per node provides better results than increasing the number of cores per node. This seems to indicate that our code is mainly limited by the memory bandwidth. Furthermore, we ensured that the numbers of cores used in our simulations evenly divide the number of cells in the 16×16 coarse grid to minimize load imbalances. For an optimal scaling of the patch smoother in the multigrid relaxation, the number of patches (and hence vertices) per processor should also be evenly balanced, but this was not implemented.

In Table 4.5, we report the average runtimes per linear iteration rather than the total runtime to take into account that the numbers of linear and nonlinear iterations change slightly between the different refinements. The runtimes only show a slight increase the more cores being used and hence indicate good weak scaling of our method.

As for the lid-driven cavity problem, we observe excellent robustness of the linear and nonlinear iteration counts with respect to the Reynolds numbers. Both linear

TABLE 4.5

Average time per linear iteration in minutes for the two-dimensional island-coalescence problem.

	128 × 128 on 16 cores			256 × 256 on 64 cores			512 × 512 on 256 cores		
Re _m \ Re	1	1,000	10,000	1	1,000	10,000	1	1,000	10,000
1	0.13	0.12	0.13	0.14	0.14	0.14	0.17	0.17	0.17
1,000	0.13	0.12	0.12	0.14	0.13	0.13	0.15	0.14	0.15
10,000	0.12	0.11	0.12	0.14	0.13	0.13	0.15	0.15	0.15

TABLE 4.6

(left) Iteration counts for the stationary lid-driven cavity problem in 3 dimensions for the Newton linearization. (right) Iteration counts for the stationary lid-driven cavity problem in 3 dimensions for taking the outer Schur complement that eliminates the $(\mathbf{E}_h, \mathbf{B}_h)$ block.

Using $\tilde{\mathcal{S}}^{(\mathbf{u}, p)}$ for order $(\mathbf{u}, p, \mathbf{E}, \mathbf{B})$				Using $\tilde{\mathcal{S}}^{(\mathbf{E}, \mathbf{B})}$ for order $(\mathbf{E}, \mathbf{B}, \mathbf{u}, p)$			
S \ Re	1	1,000	10,000	S \ Re	1	1,000	10,000
1	(3) 6.0	(3) 7.0	(4) 20.0	1	(3) 6.0	(4) 14.7	(-) > 50
1,000	(3) 7.3	(2) 9.5	(2) 6.5	1,000	(3) 12.7	(-) > 50	(-) > 50
10,000	(3) 9.0	(2) 13.0	(2) 12.5	10,000	(3) 20.0	(-) > 50	(-) > 50

and nonlinear solvers converge in either 1 or 2 iterations in the investigated ranges of Re and Re_m. We therefore do not report a table here that shows each iteration count.

4.4. Three-dimensional results. In three dimensions, we observe in general that the stationary problems are harder to solve for high parameters than in 2D. We believe that the following three points are mainly responsible for this behavior. First of all, the discretization of the electric field changes from a scalar-valued CG-function to a vector-valued NED1-function with tangential boundary conditions. Moreover, the kernel of the term $\mathbf{curl}(\mathbf{u}^n \times \mathbf{B})$ is larger in three dimensions which degrades the performance of the monolithic solver for the $(\mathbf{E}_h, \mathbf{B}_h)$ block for high Re_m. Furthermore, the grids we consider are much coarser than in 2D because of computational costs.

4.4.1. Stationary lid-driven cavity problem in three dimensions. We adapt the two-dimensional lid-driven cavity problem to three dimensions by considering the domain $\Omega = (-1/2, 1/2)^3$ and the boundary conditions $\mathbf{u} = (1, 0, 0)^\top$ on the boundary $y = 0.5$ and $\mathbf{u} = (0, 0, 0)^\top$ on the other faces. The background magnetic field $\mathbf{B}_0 = (0, 1, 0)^\top$ determines the boundary conditions for \mathbf{B} . We use a coarse grid of $8 \times 8 \times 8$ cells with 3 levels of refinement which results in a $64 \times 64 \times 64$ grid with 59.1 million DoFs. For the three-dimensional problem, we only investigate the Newton linearization as we have seen in 2D that the Newton iteration outperforms the Picard iteration in nearly all cases. The results on the left in Table 4.6 show a good control over the linear iteration numbers for the lid-driven cavity problem, where the case of $S = 1$ and Re = 10,000 seems to be the most challenging case.

On the right of Table 4.6 we report a comparison to taking the outer Schur complement that eliminates the $(\mathbf{E}_h, \mathbf{B}_h)$ block. As mentioned in section 3, we can clearly see that this choice performs worse for high values of Re and S where no convergence in 50 linear iterations was reached. We observed similar behavior for unreported experiments on transient and two-dimensional problems.

TABLE 4.7

Iteration counts for the transient lid-driven cavity problem in 3 dimensions for the Newton linearization.

$S \backslash \text{Re}$	1	10,000	100,000	$\text{Re}_m \backslash \text{Re}$	1	10,000	100,000
1	(2.1) 7.3	(3.2) 2.1	(3.3) 2.0	1	(2.1) 7.3	(3.2) 2.1	(3.3) 2.0
1,000	(3.0) 8.6	(3.3) 2.8	(3.5) 2.6	10,000	(2.5) 7.1	(3.2) 2.0	(3.3) 2.0
10,000	(4.0) 11.3	(4.0) 7.0	(4.0) 6.2	100,000	(3.0) 15.1	(3.2) 2.6	(3.3) 2.0

$\text{Re}_m \backslash S$	1	100	1,000
1	(2.1) 7.3	(2.1) 7.3	(3.0) 8.6
1,000	(3.0) 7.6	(3.0) 6.8	(3.1) 9.5
10,000	(2.5) 7.1	(3.1) 7.1	(3.2) 9.7

We do not include a full table for high Re_m , as in this case the monolithic multigrid solver cannot deal with the term $\text{curl}(\mathbf{u}^n \times \mathbf{B}^n)$ that occurs in the Newton linearization. As in 2D, this term is crucial for the convergence of the nonlinear iteration. For Newton, the iteration counts increase very slightly from $\text{Re} = \text{Re}_m = 1$ by 8.0 Krylov iterations per nonlinear step to 10.0 iterations for $\text{Re}_m = 500$ and $\text{Re} = 1$ and fails to converge for higher Re_m . We want to emphasize that in this case the failure of convergence is indeed caused by the inner multigrid method and not by an inaccurate outer Schur complement approximation. To the best of our knowledge, preconditioning methods that robustly treat the vector Laplace operator with an additional $\text{curl}(\mathbf{u}^n \times \mathbf{B})$ term in three dimensions are not known, and we intend to investigate this problem further in future work.

4.4.2. Time-dependent lid-driven cavity problem in three dimensions.

Finally, we consider the time-dependent version of the three-dimensional lid-driven cavity problem, which was also investigated in detail in [36]. The numerical results in Table 4.7 show good control of the iteration counts and the linear iteration numbers only notably increases for very high values of S . Moreover, we observe robust convergence of the monolithic multigrid solver for the $(\mathbf{E}_h, \mathbf{B}_h)$ block for high Re_m . As mentioned before in section 3.5, this can be explained by the fact that the problem does not become nearly singular for high Re_m due to the extra mass matrix. Therefore, the fact that the kernel of $\text{curl}(\mathbf{u}^n \times \mathbf{B})$ is not fully captured by the multigrid method has less influence.

5. Conclusion and outlook.

We have presented scalable block preconditioners for an augmented Lagrangian formulation of the incompressible MHD equations that exhibit parameter-robust iteration counts in most cases. We described how to control the outer Schur complement of two linearization types and introduced a special monolithic multigrid method to solve the electromagnetic block. This method is fully Re_m -robust in 2D; in three dimensions, it is able to efficiently compute results for higher parameters than was previously possible. Furthermore, our solvers allow the use of fully implicit methods for time-dependent problems. We aim to include stabilization techniques for high magnetic Reynolds numbers in future work and further investigate how to develop a robust multigrid method for the problem including the term $\text{curl}(\mathbf{u}^n \times \mathbf{B})$. This would enable a more robust solver for the most difficult case of stationary problems in three dimensions at high magnetic Reynolds numbers.

Acknowledgment.

The authors would like to thank Kaibo Hu for many useful suggestions and discussions.

Downloaded 08/16/22 to 163.1.203.243 . Redistribution subject to SIAM license or copyright; see https://pubs.siam.org/terms-privacy

Code availability. The code that was used to generate the numerical results and all major Firedrake components have been archived on [57].

REFERENCES

- [1] J. H. ADLER, T. R. BENSON, E. C. CYR, P. E. FARRELL, S. P. MACLACHLAN, AND R. S. TUMINARO, *Monolithic multigrid methods for magnetohydrodynamics*, SIAM J. Sci. Comput., 43 (2021), pp. S70–S91, <https://doi.org/10.1137/20m1348364>.
- [2] 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, <https://doi.org/10.1137/151006135>.
- [3] P. R. AMESTOY, I. S. DUFF, J.-Y. L'EXCELLENT, AND J. KOSTER, *A fully asynchronous multifrontal solver using distributed dynamic scheduling*, SIAM J. Matrix Anal. Appl., 23 (2001), pp. 15–41, <https://doi.org/10.1137/S0895479899358194>.
- [4] F. ARMERO AND J. 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](https://doi.org/10.1016/0045-7825(95)00931-0).
- [5] D. N. ARNOLD, *Finite Element Exterior Calculus*, CBMS-NSF Regional Conf. Ser. in Appl. Math. 93, SIAM, Philadelphia, 2018, <https://doi.org/10.1137/1.9781611975543>.
- [6] D. N. ARNOLD, R. S. FALK, AND R. WINTHER, *Multigrid in $H(\text{div})$ and $H(\text{curl})$* , Numer. Math., 85 (2000), pp. 197–217, <https://doi.org/10.1007/pl00005386>.
- [7] D. N. ARNOLD, R. S. FALK, AND R. WINTHER, *Finite element exterior calculus, homological techniques, and applications*, Acta Numer., 15 (2006), pp. 1–155, <https://doi.org/10.1017/S0962492906210018>.
- [8] C. BACUTA, *A unified approach for Uzawa algorithms*, SIAM J. Numer. Anal., 44 (2006), pp. 2633–2649, <https://doi.org/10.1137/050630714>.
- [9] S. BALAY, S. ABHYANKAR, M. F. ADAMS, J. BROWN, P. BRUNE, K. BUSCHELMAN, L. DALCIN, V. ELJKHOUT, W. D. GROPP, D. KARPEYEV, D. KAUSHIK, M. G. KNEPLEY, D. A. MAY, L. C. McINNES, R. T. MILLS, T. MUNSON, K. RUPP, P. SANAN, B. F. SMITH, S. ZAMPINI, H. ZHANG, AND H. ZHANG, *PETSc Users Manual*, Technical Report ANL-95/11 - Revision 3.15, Argonne National Laboratory, Argonne, IL, 2021.
- [10] M. BENZI, G. H. GOLUB, AND J. LIESEN, *Numerical solution of saddle point problems*, Acta Numer., 14 (2005), pp. 1–137, <https://doi.org/10.1017/S0962492904000212>.
- [11] M. BENZI AND M. A. OLSHANSKII, *An augmented Lagrangian-based approach to the Oseen problem*, SIAM J. Sci. Comput., 28 (2006), pp. 2095–2113, <https://doi.org/10.1137/050646421>.
- [12] D. BOFFI, F. BREZZI, AND M. FORTIN, *Mixed Finite Element Methods and Applications*, Springer, Berlin, 2013, <https://doi.org/10.1007/978-3-642-36519-5>.
- [13] J. BRACKBILL AND D. BARNES, *The effect of nonzero $\nabla \cdot \mathbf{B}$ on the numerical solution of the magnetohydrodynamic equations*, J. Comput. Phys., 35 (1980), pp. 426–430, [https://doi.org/10.1016/0021-9991\(80\)90079-0](https://doi.org/10.1016/0021-9991(80)90079-0).
- [14] F. BREZZI, J. DOUGLAS, AND L. D. MARINI, *Two families of mixed finite elements for second order elliptic problems*, Numer. Math., 47 (1985), pp. 217–235, <https://doi.org/10.1007/bf01389710>.
- [15] E. BURMAN AND P. HANSBO, *Edge stabilization for the generalized Stokes problem: A continuous interior penalty method*, Comput. Methods Appl. Mech. Engrg., 195 (2006), pp. 2393–2410, <https://doi.org/10.1016/j.cma.2005.05.009>.
- [16] L. CHACÓN, *Scalable parallel implicit solvers for 3D magnetohydrodynamics*, J. Phys. Conf. Ser., 125 (2008), 012041, <https://doi.org/10.1088/1742-6596/125/1/012041>.
- [17] L. CHEN, Y. WU, L. ZHONG, AND J. ZHOU, *MultiGrid preconditioners for mixed finite element methods of the vector Laplacian*, J. Sci. Comput., 77 (2018), pp. 101–128, <https://doi.org/10.1007/s10915-018-0697-7>.
- [18] 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, <https://doi.org/10.1137/12088879x>.
- [19] W. DAI AND P. R. WOODWARD, *On the divergence-free condition and conservation laws in numerical simulations for supersonic magnetohydrodynamical flows*, Astrophys. J., 494 (1998), pp. 317–335, <https://doi.org/10.1086/305176>.
- [20] C. DAVIES, M. POZZO, D. GUBBINS, AND D. ALFÈ, *Constraints from material properties on the dynamics and evolution of earth's core*, Nature Geosci., 8 (2015), pp. 678–685, <https://doi.org/10.1038/ngeo2492>.

- [21] J. DOUGLAS AND T. DUPONT, *Interior penalty procedures for elliptic and parabolic Galerkin methods*, in Computing Methods in Applied Sciences, Vol. 58, Springer, Berlin, 1976, pp. 207–216, <https://doi.org/10.1007/BFb0120591>.
- [22] H. ELMAN, D. SILVESTER, AND A. WATHEN, *Finite Elements and Fast Iterative Solvers*, Oxford University Press, Oxford, 2014, <https://doi.org/10.1093/acprof:oso/9780199678792.001.0001>.
- [23] P. E. FARRELL, M. G. KNEPLEY, L. MITCHELL, AND F. WECHSUNG, *PCPATCH: Software for the topological construction of multigrid relaxation methods*, ACM Trans. Math. Softw., 47 (2021), pp. 1–22, <https://doi.org/10.1145/3445791>.
- [24] P. E. FARRELL, L. MITCHELL, L. R. SCOTT, AND F. WECHSUNG, *A Reynolds-robust preconditioner for the Scott–Vogelius discretization of the stationary incompressible Navier–Stokes equations*, SMAI J. Comput. Math., 7 (2021), pp. 75–96, <https://doi.org/10.5802/smai-jcm.72>.
- [25] P. E. FARRELL, L. MITCHELL, AND F. WECHSUNG, *An augmented Lagrangian preconditioner for the 3D stationary incompressible Navier–Stokes equations at high Reynolds number*, SIAM J. Sci. Comput., 41 (2019), pp. A3073–A3096, <https://doi.org/10.1137/18m1219370>.
- [26] N. R. GAUGER, A. LINKE, AND P. W. SCHROEDER, *On high-order pressure-robust space discretizations, their advantages for incompressible high Reynolds number generalised Beltrami flows and beyond*, SMAI J. Comput. Math., 5 (2019), pp. 89–129, <https://doi.org/10.5802/smai-jcm.44>.
- [27] J.-F. GERBEAU, C. L. BRIS, AND T. LELIÈVRE, *Mathematical Methods for the Magnetohydrodynamics of Liquid Metals*, Oxford University Press, Oxford, 2006, <https://doi.org/10.1093/acprof:oso/9780198566656.001.0001>.
- [28] T. HEISTER AND G. RAPIN, *Efficient augmented Lagrangian-type preconditioning for the Oseen problem using grad-div stabilization*, Internat. J. Numer. Methods Fluids, 71 (2013), pp. 118–134, <https://doi.org/10.1002/fld.3654>.
- [29] Q. HONG, J. KRAUS, J. XU, AND L. ZIKATANOV, *A robust multigrid method for discontinuous Galerkin discretizations of Stokes and linear elasticity equations*, Numer. Math., 132 (2015), pp. 23–49, <https://doi.org/10.1007/s00211-015-0712-y>.
- [30] K. HU, Y. MA, AND J. XU, *Stable finite element methods preserving $\nabla \cdot B = 0$ exactly for MHD models*, Numer. Math., 135 (2016), pp. 371–396, <https://doi.org/10.1007/s00211-016-0803-4>.
- [31] K. HU, W. QIU, AND K. SHI, *Convergence of a B-E based finite element method for MHD models on Lipschitz domains*, J. Comput. Appl. Math., 368 (2020), 112477, <https://doi.org/10.1016/j.cam.2019.112477>.
- [32] K. HU AND J. XU, *Structure-preserving finite element methods for stationary MHD models*, Math. Comp., 88 (2018), pp. 553–581, <https://doi.org/10.1090/mcom/3341>.
- [33] V. JOHN, A. LINKE, C. MERDON, M. NEILAN, AND L. G. REBHOLZ, *On the divergence constraint in mixed finite element methods for incompressible flows*, SIAM Rev., 59 (2017), pp. 492–544, <https://doi.org/10.1137/15m1047696>.
- [34] D. A. KNOLL AND L. CHACÓN, *Coalescence of magnetic islands, sloshing, and the pressure problem*, Phys. Plasmas, 13 (2006), 032307, <https://doi.org/10.1063/1.2173515>.
- [35] L. LI AND W. ZHENG, *A robust solver for the finite element approximation of stationary incompressible MHD equations in 3D*, J. Comput. Phys., 351 (2017), pp. 254–270, <https://doi.org/10.1016/j.jcp.2017.09.025>.
- [36] Y. MA, K. HU, X. HU, AND J. XU, *Robust preconditioners for incompressible MHD models*, J. Comput. Phys., 316 (2016), pp. 721–746, <https://doi.org/10.1016/j.jcp.2016.04.019>.
- [37] S. MOLOKOV, R. MOREAU, AND K. MOFFATT, *Magnetohydrodynamics*, Springer, Dordrecht, Netherlands, 2007, <https://doi.org/10.1007/978-1-4020-4833-3>.
- [38] J.-C. NÉDÉLEC, *Mixed finite elements in \mathbb{R}^3* , Numer. Math., 35 (1980), pp. 315–341, <https://doi.org/10.1007/BF01396415>.
- [39] J.-C. NÉDÉLEC, *A new family of mixed finite elements in \mathbb{R}^3* , Numer. Math., 50 (1986), pp. 57–81, <https://doi.org/10.1007/bf01389668>.
- [40] E. G. PHILLIPS, *Fast Solvers and Uncertainty Quantification for Models of Magnetohydrodynamics*, Ph.D. thesis, University of Maryland, College Park, MD, 2014.
- [41] 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>.
- [42] E. G. PHILLIPS, J. N. SHADID, E. C. CYR, H. C. ELMAN, AND R. P. PAWLOWSKI, *Block preconditioners for stable mixed nodal and edge finite element representations of incompressible resistive MHD*, SIAM J. Sci. Comput., 38 (2016), pp. B1009–B1031, <https://doi.org/10.1137/16M1074084>.
- [43] F. RATHGEBER, D. A. HAM, L. MITCHELL, M. LANGE, F. LUPORINI, A. T. T. MCRAE, G.-T. BERCEA, G. R. MARKALL, AND P. H. J. KELLY, *Firedrake: Automating the finite*

- element method by composing abstractions*, ACM Trans. Math. Software, 43 (2016), 24, <https://doi.org/10.1145/2998441>.
- [44] P. A. RAVIART AND J. M. THOMAS, *A mixed finite element method for 2-nd elliptic problems*, in Mathematical Aspects of Finite Element Methods, Lecture Notes in Math., Springer, Berlin, 1977, pp. 292–315, <https://doi.org/10.1007/bfb0064470>.
- [45] Y. SAAD, *A flexible inner-outer preconditioned GMRES algorithm*, SIAM J. Sci. Comput., 14 (1993), pp. 461–469, <https://doi.org/10.1137/0914028>.
- [46] J. SCHÖBERL, *Robust Multigrid Methods for Parameter Dependent Problems*, Ph.D. thesis, Johannes Kepler Universität Linz, Linz, Austria, 1999.
- [47] 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>.
- [48] L. R. SCOTT AND M. VOGELIUS, *Conforming Finite Element Methods for Incompressible and Nearly Incompressible Continua*, in Large Scale Computations in Fluid Mechanics, Vol. 22 (Part 2), American Mathematical Society, Providence, RI, 1985, pp. 221–244.
- [49] 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, <https://doi.org/10.1016/j.jcp.2010.06.018>.
- [50] J. SHADID, R. PAWLOWSKI, E. CYR, R. TUMINARO, L. CHACÓN, AND P. 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, <https://doi.org/10.1016/j.cma.2016.01.019>.
- [51] B. S. SOUTHWORTH, A. A. SIVAS, AND S. RHEBERGEN, *On fixed-point, Krylov, and 2×2 block preconditioners for nonsymmetric problems*, SIAM J. Matrix Anal. Appl., 41 (2020), pp. 871–900, <https://doi.org/10.1137/19m1298317>.
- [52] M. WATHEN AND C. GREIF, *A scalable approximate inverse block preconditioner for an incompressible magnetohydrodynamics model problem*, SIAM J. Sci. Comput., 42 (2020), pp. B57–B79, <https://doi.org/10.1137/19m1255409>.
- [53] M. WATHEN, C. GREIF, AND D. SCHÖTZAU, *Preconditioners for mixed finite element discretizations of incompressible MHD equations*, SIAM J. Sci. Comput., 39 (2017), pp. A2993–A3013, <https://doi.org/10.1137/16m1098991>.
- [54] S. WU AND J. XU, *Simplex-averaged finite element methods for $H(\text{grad})$, $H(\text{curl})$, and $H(\text{div})$ convection-diffusion problems*, SIAM J. Numer. Anal., 58 (2020), pp. 884–906, <https://doi.org/10.1137/18m1227196>.
- [55] S. WU AND L. T. ZIKATANOV, *On the Unisolvence for the Quasi-polynomial Spaces of Differential Forms*, preprint, <https://arxiv.org/abs/2003.14278> (2020).
- [56] J. XU, *Iterative methods by space decomposition and subspace correction*, SIAM Rev., 34 (1992), pp. 581–613, <https://doi.org/10.1137/1034116>.
- [57] *Software Used in ‘An Augmented Lagrangian Preconditioner for the Magnetohydrodynamics Equations at High Reynolds Number,’* <https://doi.org/10.5281/zenodo.5879212> (2022).
- [58] S. ZHANG, *A new family of stable mixed finite elements for the 3D Stokes equations*, Math. Comput., 74 (2005), pp. 543–555, <https://doi.org/10.1090/s0025-5718-04-01711-9>.