

ABSTRACT

Title of dissertation: FAST SOLVERS
AND UNCERTAINTY QUANTIFICATION
FOR MODELS OF MAGNETOHYDRODYNAMICS

Edward Phillips, Doctor of Philosophy, 2014

Dissertation directed by: Professor Howard Elman
Department of Computer Science
Institute for Advanced Computer Studies

The magnetohydrodynamics (MHD) model describes the flow of electrically conducting fluids in the presence of magnetic fields. A principal application of MHD is the modeling of plasma physics, ranging from plasma confinement for thermonuclear fusion to astrophysical plasma dynamics. MHD is also used to model the flow of liquid metals, for instance in magnetic pumps, liquid metal blankets in fusion reactor concepts, and aluminum electrolysis. The model consists of a non-self-adjoint, nonlinear system of partial differential equations (PDEs) that couple the Navier-Stokes equations for fluid flow to a reduced set of Maxwell's equations for electromagnetics.

In this dissertation, we consider computational issues arising for the MHD equations. We focus on developing fast computational algorithms for solving the algebraic systems that arise from finite element discretizations of the fully coupled MHD equations. Emphasis is on solvers for the linear systems arising from algorithms such as Newton's method or Picard iteration, with a main goal of developing

preconditioners for use with iterative methods for the linearized systems. In particular, we first consider the linear systems arising from an exact penalty finite element formulation of the MHD equations. We then draw on this research to develop solvers for a formulation that includes a Lagrange multiplier within Maxwell's equations. We also consider a simplification of the MHD model: in the MHD kinematics model, the equations are reduced by assuming that the flow behavior of the system is known. In this simpler setting, we allow for epistemic uncertainty to be present. By mathematically modeling this uncertainty with random variables, we investigate its implications on the physical model.

FAST SOLVERS AND UNCERTAINTY QUANTIFICATION FOR
MODELS OF MAGNETOHYDRODYNAMICS

by

Edward G. Phillips

Dissertation submitted to the Faculty of the Graduate School of the
University of Maryland, College Park in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
2014

Advisory Committee:
Professor Howard Elman, Chair/Advisor
Professor Ricardo Nochetto
Professor Dianne O'Leary
Professor David Levermore
Professor James Baeder, Dean's Representative

© Copyright by
Edward G. Phillips
2014

Acknowledgments

I would like to thank my advisor, Howard Elman, for his guidance and support throughout the course of my graduate career. His mentorship has helped me to mature as an academic and provided me with an excellent foundation for my research.

I am also grateful to my collaborators at Sandia National Laboratories: John Shadid, Eric Cyr, and Roger Pawlowski. Their input and encouragement, as well as the code and computational resources that they have made available to me, have been integral to this work.

I also thank Ricardo Nochetto, Dianne O’Leary, David Levermore, and James Baeder for serving on my committee and contributing their insights on my work. I am especially grateful to Ricardo for his continued support throughout my time at the University of Maryland.

Finally, I would like to thank my family, friends, and my always encouraging partner, Cameron, for supporting me throughout my education.

Table of Contents

List of Tables	v
List of Figures	vii
1 Introduction	1
1.1 Overview of Magnetohydrodynamics	1
1.2 Numerical Methods for Magnetohydrodynamics	2
1.2.1 Fully Coupled Solution Strategy	3
1.2.2 Decoupled Solution Strategies	4
1.3 Modeling Uncertainty	5
1.4 Outline of Dissertation	7
2 Problem Statement and Survey of Existing Approaches	9
2.1 The MHD Equations	9
2.2 A Stable Exact Penalty Finite Element Formulation	14
2.3 A Stable Lagrange Multiplier Finite Element Formulation	17
2.4 Preconditioners for the Discretized MHD Equations	21
2.4.1 Block Preconditioners for Coupled Linear Systems	22
2.4.2 Existing Preconditioners for Fully Coupled MHD	24
2.5 Uncertainty Quantification	27
3 Block Preconditioners for an Exact Penalty MHD Formulation	31
3.1 A Block Preconditioner for the 2D Picard System	33
3.1.1 An Approximation for X	36
3.1.2 An Approximation for Y	37
3.2 A Block Preconditioner for the Newton System	43
3.3 Computational Results	47
3.3.1 MHD Lid Driven Cavity	48
3.3.1.1 Picard Iteration	52
3.3.1.2 Newton's Method	54
3.3.1.3 Robustness with Respect to Mesh Refinement	56
3.3.2 Hartmann Flow	60
3.4 Conclusion	63

4	Block Preconditioners for a Lagrange Multiplier MHD Formulation	65
4.1	Block Preconditioners for the Picard System	67
4.1.1	Preconditioners for \mathcal{M}_P	69
4.1.1.1	Mass Augmentation	70
4.1.1.2	Grad Div Augmentation	72
4.1.2	An Approximation of X_P	75
4.1.2.1	The Two-Dimensional Case	77
4.1.2.2	The Three-Dimensional Case	78
4.1.3	A Commutator for \hat{Y}_P	82
4.1.4	Computational Tasks	84
4.2	Block Preconditioners for the Newton System	86
4.2.1	Preconditioners for \mathcal{M}_N	87
4.2.2	Approximations of X_N and Y_N	89
4.2.3	Computational Results	92
4.2.4	Two-Dimensional MHD Lid Driven Cavity	94
4.2.4.1	Picard Iteration	96
4.2.4.2	Newton's Method	100
4.2.5	Three-Dimensional MHD Lid Driven Cavity	104
4.2.5.1	Picard Iteration	106
4.2.5.2	Newton's Method	108
4.3	Conclusion	110
5	A Stochastic Approach to Uncertainty in the Equations of MHD Kinematics	112
5.1	A Finite Element Formulation	114
5.1.1	Example Problem: Hartmann Flow	119
5.1.2	Example Problem: MHD Eddy	120
5.2	MHD Kinematics with Uncertain Data	121
5.2.1	Uncertain Velocity	123
5.2.1.1	Test Problem 1: Irrotational Fluctuations	124
5.2.1.2	Test Problem 2: Incompressible Flow	127
5.2.2	Uncertain Resistivity	129
5.2.2.1	Test Problem 3: Piecewise constant β	129
5.2.2.2	Test Problem 4: β as a truncated KL expansion	134
5.2.2.3	Test Problem 5: Uncertain Resistivity in 3D	137
5.3	Linear Solvers for the Discretized Kinematics System	141
5.3.1	Analysis of Eigenvalues	142
5.3.2	Numerical Results	145
5.4	Conclusion	147
6	Summary and Conclusions	149
	Bibliography	152

List of Tables

3.1	Definitions of discrete operators as they correspond to continuous operators.	32
3.2	Iterations required for convergence of the nonlinear iteration. ‘×’ indicates no convergence within twenty iterations. ‘*’ indicates convergence required backtracking.	51
3.3	Average GMRES iterations required for convergence with $\mathcal{P}_{P,\alpha}$ on the Picard linearization of the MHD lid driven cavity problem with $\alpha = 1$ and $\alpha = \alpha_*$	52
3.4	Computed values of α_* for the second Picard iteration.	53
3.5	Average GMRES iterations required for convergence with $\mathcal{P}_{P,\alpha,\gamma}$ on the Newton linearization of the MHD lid driven cavity problem.	54
3.6	Computed values of γ_* at the second Newton step.	55
3.7	Average GMRES iterations required for convergence with $\mathcal{P}_{P,\alpha}$ on the Picard linearization of the Hartmann flow problem.	62
3.8	Average GMRES iterations required for convergence with $\mathcal{P}_{P,\alpha,\gamma}$ on the Newton linearization of the Hartmann flow problem.	62
4.1	Definitions of discrete operators as they correspond to continuous operators.	66
4.2	Number of nonlinear iterations needed to obtain convergence for the two-dimensional MHD lid driven cavity problem with $R = 100$ on a 100×100 element mesh.	95
4.3	Average GMRES iterations required for convergence with $\mathcal{P}_{P,Q}$ and $\mathcal{P}_{P,GD}$ on the Picard linearization of the two-dimensional MHD lid driven cavity problem with various values of k . ‘×’ indicates no convergence due to breakdown of the preconditioner. Bold entries correspond to the choice $k = \frac{S}{R_m}$	97
4.4	Average GMRES iterations and time in seconds required for convergence with $\mathcal{P}_{P,Q}$ and $\mathcal{P}_{P,GD}$ on the Picard linearization of the two-dimensional MHD lid driven cavity problem with various values of R_m	99

4.5	Average GMRES iterations required for convergence with $\mathcal{P}_{N,Q}$ and $\mathcal{P}_{N,GD}$ on the Newton linearization of the two-dimensional MHD lid driven cavity problem with various values of k . ‘ \times ’ indicates no convergence due to breakdown of the preconditioner. Bold entries correspond to the choice $k = \frac{S}{R_m}$	101
4.6	Average GMRES iterations and time in seconds required for convergence with $\mathcal{P}_{N,Q}$ and $\mathcal{P}_{N,GD}$ on the Newton linearization of the two-dimensional MHD lid driven cavity problem with various values of R_m	103
4.7	Number of nonlinear iterations needed to obtain convergence for the three-dimensional MHD lid driven cavity problem with $R = 100$ on a $40 \times 40 \times 40$ element mesh.	105
4.8	Average GMRES iterations required for convergence with $\mathcal{P}_{P,Q}$ and $\mathcal{P}_{P,GD}$ on the Picard linearization of the three-dimensional MHD lid driven cavity problem with various values of β	106
4.9	Values of β_* at the second step of a Picard iteration.	106
4.10	Average GMRES iterations required for convergence with $\mathcal{P}_{N,Q}$ and $\mathcal{P}_{N,GD}$ on the Newton linearization of the three-dimensional MHD lid driven cavity problem with various values of β	109
5.1	Preconditioned GMRES iteration counts.	146

List of Figures

3.1	Streamlines for the MHD lid driven cavity problem with $R = 5000$ and $R_m = 0, 0.1, 0.3, 0.4, 5, 10, 20, 30$. The four latter cases are zoomed in to $[0, 1] \times [0.8, 1]$	49
3.2	Magnitude of the divergence of \vec{B}_h versus h for $R = 64$ and $R_m = 64, 256, 1024$	50
3.3	Mesh refinement results for the MHD lid driven cavity problem for the Picard iteration. Parameters are $R = 256, R_m = 256$ on the top and $R = 256, R_m = 1024$ on the bottom. Average GMRES iterations on the left and average linear solve time per nonlinear iteration on the right.	57
3.4	Mesh refinement results for the MHD lid driven cavity problem for Newton's method. Parameters are $R = 256, R_m = 256$ on the top and $R = 256, R_m = 1024$ on the bottom. Average GMRES iterations on the left and average linear solve time per nonlinear iteration on the right.	58
3.5	Magnitude of the error in \vec{u}_h and \vec{B}_h and the divergence of \vec{B}_h versus h . 61	61
4.1	Velocity streamlines for the two-dimensional MHD lid driven cavity problem with $R = 100, S = 1$, and various values of R_m	94
4.2	Weak parallel scaling results for the two-dimensional MHD lid driven cavity problem with Picard linearization. Average GMRES iterations on the left and average linear solve time per nonlinear iteration on the right.	99
4.3	Weak parallel scaling results for the two-dimensional MHD lid driven cavity problem with Newton's method. Average GMRES iterations on the left and average linear solve time per nonlinear iteration on the right.	103
4.4	Velocity streamlines at cross-sections $z = 0$ and $z = \frac{1}{4}$ for the three-dimensional MHD lid driven cavity problem with $R = 100, S = 1$, and various values of R_m	104

4.5	Weak parallel scaling results for the three-dimensional MHD lid driven cavity problem with Picard linearization. Average GMRES iterations on the left and average linear solve time per nonlinear iteration on the right.	107
4.6	Weak parallel scaling results for the three-dimensional MHD lid driven cavity problem with Newton's method. Average GMRES iterations on the left and average linear solve time per nonlinear iteration on the right.	109
5.1	Velocity profiles (a), induced magnetic fields (b), and error $\ \vec{B}_h - \vec{B}\ _2$ (c) for the Hartmann problem.	118
5.2	Velocity profile and induced magnetic fields for deterministic MHD eddy problem.	120
5.3	Profiles of u_x and u_y along the line $x = 0$ for two random instances with $\sigma^2 = 5.0 \times 10^{-3}$, together with the mean profile.	125
5.4	Profiles of B_x and B_y along the line $x = 0$ for two random instances with $\sigma^2 = 5.0 \times 10^{-3}$, together with the deterministic solution obtained from \vec{u}_0	126
5.5	Profiles of B_x and B_y for Test Problem 1, plotted along the line $x = 0$ for the mean $\mu(\vec{B}_h)$ with $\sigma^2 = 5.0 \times 10^{-3}$, 6.0×10^{-3} , and 7.0×10^{-3}	127
5.6	Profiles of B_x and B_y for Test Problem 2, plotted along the line $x = 0$ for the mean $\mu(\vec{B}_h)$ with $\sigma^2 = 5.0 \times 10^{-3}$, 6.0×10^{-3} , and 7.0×10^{-3}	128
5.7	Domain partitionings considered in Test Problem 3.	130
5.8	Instances of \vec{B}_h obtained with partitioning P_2 for Test Problem 3.	130
5.9	Instances of \vec{B}_h obtained with partitioning P_3 for Test Problem 3.	131
5.10	Deterministic and mean magnetic field lines compared for Test Problem 3.	132
5.11	Euclidean norm of standard deviation $\ \sigma(\vec{B}_h)\ _2$ for Test Problem 3.	133
5.12	Random realizations of η for Test Problem 4.	135
5.13	Deterministic and mean magnetic field lines compared for Test Problem 4.	136
5.14	Euclidean norm of standard deviation $\ \sigma(\vec{B}_h)\ _2$ for Test Problem 4.	136
5.15	Deterministic fields lines (top), mean field lines with partitioning P_1 (middle), and mean field lines with partitioning P_2 at cross sections $z = -\frac{1}{4}$ (left), $z = 0$ (middle), and $z = \frac{1}{4}$ (right) for Test Problem 5.	138
5.16	$\ \sigma(\vec{B}_h)(\vec{x})\ _2$ at cross sections for partitionings P_1 (top) and P_2 (bottom) for Test Problem 5.	139

Chapter 1: Introduction

This dissertation is concerned with computational issues that arise in the modelling of magnetohydrodynamics (MHD). In particular, we focus on developing fast computational algorithms for solving the algebraic systems that arise from finite element discretizations of the fully coupled MHD equations. Additionally, we mathematically model the effects of epistemic uncertainty on the reduced kinematics model of MHD.

1.1 Overview of Magnetohydrodynamics

The MHD model describes the flow of electrically conducting fluids in the presence of magnetic fields. For different ranges of physical parameters, the MHD model can describe such materials as plasma, liquid metal, and brine. MHD has been used to model plasmas in applications such as plasma confinement for thermonuclear fusion and astrophysical plasma dynamics governing the behavior of stars and solar wind [33]. Liquid metal applications of the MHD model include magnetic pumps, liquid metal blankets in fusion reactor concepts, aluminum electrolysis, and the geodynamo [18, 48].

The MHD model is governed by a system of non-self-adjoint, nonlinear partial

differential equations (PDEs) obtained from coupling the Navier-Stokes equations for fluids to Maxwell's equations for electromagnetics. The two sets of equations are coupled by the Lorentz force, which governs the effect of a magnetic field on fluid flow, and the appearance of the fluid velocity in Ohm's law, which accounts for the influence of hydrodynamics on the electric current. When fully coupled, the equations model the complex behavior of MHD materials. Because both hydrodynamic and electromagnetic effects play a strong role in this system, MHD dynamics can span over a large range of length- and time-scales. This complex physical behavior makes the MHD equations difficult to solve and necessitates the development of robust, accurate numerical methods of approximating their solution.

1.2 Numerical Methods for Magnetohydrodynamics

The MHD equations can be numerically approximated by discretizing in space and time and linearizing the equations. Many spatial discretization methods have been applied to the MHD equations, including finite difference, finite volume, and boundary element methods. In this dissertation, we focus on discretizations of the MHD equations resulting from the finite element method (FEM). We will consider linearizations of the equations resulting from fixed point iterations of Picard type and Newton's method.

When discretized and linearized, the approximation of the solution to the MHD equations is reduced to solving a series of large linear systems. In general, these systems are sparse, highly indefinite, and non-symmetric. Furthermore, the

strong physical coupling present in the original system of PDEs is implicit in these algebraic systems. For these reasons, the development of efficient solution strategies for these linear systems is imperative for MHD simulations. These strategies should account for the coupling between hydrodynamic and electromagnetic phenomena and be robust over a range of physical parameters.

1.2.1 Fully Coupled Solution Strategy

One solution strategy is to develop solvers for the fully coupled linear systems arising from the MHD equations. Because these systems are large and sparse, iterative methods, such as the generalized minimum residual (GMRES) and biconjugate gradient with stabilization (BiCGStab) methods, are a natural choice of solver. For fast convergence, iterative methods applied to these systems should be paired with robust, efficient preconditioners. The challenge of the fully coupled solution strategy is then to develop preconditioners that effectively account for the complex physical processes implicit in the discretized MHD equations. The coupling between hydrodynamic and electromagnetic phenomena must then be accounted for in the preconditioner.

This strategy has the advantage of requiring only one linear solve for each iteration of a nonlinear iteration. This quality lends itself to obtaining fast steady-state solutions. The speed of the algorithm is governed by the rate of convergence of the nonlinear iteration and the cost of the linear solves.

1.2.2 Decoupled Solution Strategies

Alternatively, the MHD system can be split into its fluid and its electromagnetic parts, and each part can be solved separately. By assuming that the electromagnetic behavior is known, the hydrodynamic behavior is modeled by the Navier-Stokes equations with the Lorentz force appearing as a forcing term. If the flow has little effect on the magnetic field (e.g. in MHD propulsion applications such as magnetic pumps and rail guns [18]), this system alone may be a good approximation the full MHD system. The model then simulates the effects of a known magnetic field on a fluid's flow behavior.

If the fluid behavior is treated as known, the electromagnetic behavior is modeled by Maxwell's equations with the velocity appearing as a constant in Ohm's law. This is a linear system of PDEs referred to as the kinematics equations for MHD. If the electromagnetics have little influence over the hydrodynamics (e.g. in kinematic dynamo theory [45, 47]), the full MHD system may be approximated by the kinematics equations. The model then simulates the magnetic field induced by a prescribed fluid flow profile, providing a means for assessing whether a given flow can sustain dynamo action.

In order to solve the fully coupled MHD equations, operator splitting methods [3, 43, 57] iterate between the hydrodynamic system and the electromagnetic system. In these algorithms, the solution of one system is used as the data for the other system. By decomposing the fully coupled MHD system into its component parts, these methods have the benefit of requiring solving linear systems that are not

only smaller than the original system but are also derived from a single physical process. Existing solvers for the Navier-Stokes equations and Maxwell's equations can be applied to these linear systems. Thus, the physical coupling plays no role in developing iterative solvers for operator splitting methods; instead the coupling is accounted for by the outer iteration between the two systems. While the linear systems arising in operator splitting algorithms are smaller and less complex than fully coupled linear systems, errors may be introduced due to operator splitting and many outer iterations may be required to obtain accurate solutions to stationary problems.

1.3 Modeling Uncertainty

If one wanted to model an actual physical system by the kinematics model of MHD, this would require specifying the flow behavior of the system. Specifically, the velocity field would need to be measured throughout the physical domain. In general, the velocity field may be measured at a sampling of discrete points, and it may be approximated for the rest of the domain. Due to the inaccurate specification of the flow behavior, the electromagnetic behavior predicted by the kinematics model may not represent the physical situation very well. This highlights a fundamental issue in the modeling of real-life physical systems: epistemic uncertainty in the specification of physical quantities can influence the physical phenomena predicted by a mathematical model.

In this thesis, we focus specifically on epistemic uncertainty in the MHD kine-

matics model. Treating the hydrodynamics as known may introduce uncertainty because the flow properties of the fluid may not be known on the interior of the domain. There are also aspects of the physical model that suggest the importance of small-scale uncertainty. For instance, the large-scale mean flow of the earth's outer core cannot account for the magnitude of the earth's magnetic field. In geodynamo theory, it is proposed that small-scale turbulent behavior can give rise to a large-scale magnetic field through the α -effect [18]. This suggests that small fluctuations of a velocity field may produce large changes in a magnetic field.

Additionally, the distribution of material properties may be uncertain in physical applications. When multiple fluids are present, such as when multiple liquid metals are mixing together, the magnetic resistivity is not homogeneous throughout the domain and may vary over orders of magnitude [29]. Because the resistivity can have a strong influence on such physical systems, including changing the topology of the magnetic field, uncertainty in the distribution of the resistivity may produce different magnetic effects than predicted by simulations using the mean resistivity.

The mathematical modeling of uncertain quantities such as these is a primary concern in the field of uncertainty quantification. When uncertain quantities are input data for systems of PDEs, one method for quantifying the effect of epistemic uncertainty is to model the input data as random variables [31]. This provides a mathematical framework in which the solution to the system of PDEs is also a random variable. Statistical properties of the solution, such as mean and variance, can then be obtained from the model. Thus, by following this methodology and deriving stochastic expressions for hydrodynamic and material properties within the

MHD kinematics equations, the effects of epistemic uncertainty on electromagnetic properties can be modeled.

1.4 Outline of Dissertation

In this dissertation, we consider computational issues arising in both fully coupled and decoupled approaches to the MHD model. In particular, we focus on developing fully coupled solvers for the discretized MHD equations and on quantifying uncertainty in the MHD kinematics equations. We begin by introducing the MHD equations in Chapter 2. We review existing finite element formulations of the MHD equations, paying particular attention to an exact penalty formulation [36] and a Lagrange multiplier formulation [59]. We discuss existing fully coupled solution strategies for the MHD equations as well as existing literature on uncertainty quantification for systems of PDEs.

Chapters 3 and 4 concern the development of fast, robust fully coupled solvers for the linear systems arising from finite element discretizations of the MHD system. In particular, we develop preconditioners for iterative methods applied to these linear systems. In Chapter 3, we focus on the linear systems arising from the exact penalty formulation, and in Chapter 4, we focus on the Lagrange multiplier formulation. The goal of these chapters is to introduce algorithmically scalable preconditioners that perform well over a range of physical parameters and to demonstrate the effectiveness of these preconditioners on a number of test problems.

In Chapter 5, we develop a numerical method for the MHD kinematics equa-

tions that accounts for epistemic uncertainty. We present stochastic expressions for the input data and investigate their effects on the electromagnetics predicted by the model on a series of test problems. Chapter 6 presents some concluding remarks.

Chapter 2: Problem Statement and Survey of Existing Approaches

2.1 The MHD Equations

The MHD equations are obtained from a coupling of the Navier-Stokes equations for fluid flow and Maxwell's equations for electromagnetics. Restricting the model to the case of a single incompressible, homogeneous fluid, the equations at a steady state can be written

$$\rho \vec{u} \cdot \nabla \vec{u} - \nu \Delta \vec{u} + \nabla p - \vec{j} \times \vec{B} = \vec{f}, \quad (2.1a)$$

$$\nabla \cdot \vec{u} = 0, \quad (2.1b)$$

$$\nabla \times \left(\frac{1}{\mu} \vec{B} \right) = \vec{j}, \quad (2.2a)$$

$$\nabla \cdot (\varepsilon \vec{E}) = \rho_c, \quad (2.2b)$$

$$\nabla \times \vec{E} = \vec{0}, \quad (2.2c)$$

$$\nabla \cdot \vec{B} = 0, \quad (2.2d)$$

$$\vec{E} + \vec{u} \times \vec{B} = \eta \vec{j}. \quad (2.3)$$

The unknowns here are the fluid velocity \vec{u} , the pressure p , the magnetic induction \vec{B} , the electric field \vec{E} , and the current density \vec{j} . The charge density ρ_c can be regarded as an auxiliary variable that can be obtained from \vec{E} through (2.2b). The external forcing term \vec{f} in the momentum equation is assumed to be known. The

physical parameters prescribed in these equations are the fluid density ρ , the kinematic viscosity ν , the electric permittivity ε , the magnetic permeability μ , and the magnetic resistivity η . In the single fluid case, each of these parameters is treated as a fixed constant throughout the domain.

Note that (2.1a) – (2.1b) are the incompressible Navier-Stokes equations, (2.2a) – (2.2d) are a reduced form of Maxwell’s equations, and (2.3) is Ohm’s law, which is included to complete the system. The hydrodynamics and the electromagnetics are coupled through the effect of the Lorentz force $\vec{j} \times \vec{B}$ on the motion of the fluid as well as the influence of the term $\vec{u} \times \vec{B}$ on the current density.

On a bounded domain $\Omega \subset \mathbb{R}^d$, the MHD system is completed by supplying boundary conditions for \vec{u} , \vec{B} , and \vec{E} . We consider the boundary conditions

$$\vec{u} = \vec{g}, \tag{2.4a}$$

$$\vec{B} \times \vec{n} = \vec{q} \times \vec{n}, \tag{2.4b}$$

$$\vec{E} \cdot \vec{n} = k \tag{2.4c}$$

on $\partial\Omega$. (Alternatively, we could prescribe $\vec{B} \cdot \vec{n}$ and $\vec{E} \times \vec{n}$. Both choices of boundary conditions can be motivated from a physical perspective depending on the physics of the particular problem [29]. In Section 2.3, we will consider a finite element formulation that uses edge elements to discretize \vec{B} . In this setting, it is convenient to prescribe the tangential component $\vec{B} \times \vec{n}$ on the boundary, so we consider only (2.4) in formulating the problem.)

The equations (2.1a)–(2.3) hold in either two or three dimensions. In three dimensions, the vector operations are defined in the usual ways. In two dimensions,

we define extensions of the cross product and curl to apply to two-dimensional vectors and scalars. The cross product between two vectors $\vec{v} = (v_x, v_y)$ and $\vec{w} = (w_x, w_y)$ is a scalar field defined as

$$\vec{v} \times \vec{w} := v_x w_y - v_y w_x, \quad (2.5)$$

which defines the curl of a 2D vector to be

$$\nabla \times \vec{v} := \begin{pmatrix} -\frac{\partial}{\partial y} & \frac{\partial}{\partial x} \end{pmatrix} \vec{v}. \quad (2.6)$$

We also define the cross product of a vector \vec{v} with a scalar ϕ as the vector

$$\vec{v} \times \phi := \begin{pmatrix} v_y \phi \\ -v_x \phi \end{pmatrix}, \quad (2.7)$$

which defines the curl of a scalar to be

$$\nabla \times \phi := \begin{pmatrix} \frac{\partial}{\partial y} \\ -\frac{\partial}{\partial x} \end{pmatrix} \phi. \quad (2.8)$$

We can simplify the MHD system by substituting (2.2a) into (2.1a) and (2.3), and then substituting (2.3) into (2.2c), to obtain the commonly used form

$$\rho \vec{u} \cdot \nabla \vec{u} - \nu \Delta \vec{u} + \nabla p + \vec{B} \times (\nabla \times \frac{1}{\mu} \vec{B}) = \vec{f}, \quad (2.9a)$$

$$\nabla \cdot \vec{u} = 0, \quad (2.9b)$$

$$\nabla \times (\frac{\eta}{\mu} \nabla \times \vec{B}) - \nabla \times (\vec{u} \times \vec{B}) = \vec{0}, \quad (2.9c)$$

$$\nabla \cdot \vec{B} = 0. \quad (2.9d)$$

Now the MHD equations are expressed as a coupling of the three unknowns \vec{u} , p , and \vec{B} . These equations are completed using only boundary conditions (2.4a) and (2.4b). Then \vec{j} and \vec{E} can be recovered from (2.2a) and (2.3) respectively.

In the reduced MHD equations (2.9), the hydrodynamics and electromagnetics of the system are clearly expressed as two coupled systems of PDEs with constraints. The momentum equation (2.9a) with the incompressibility constraint (2.9b) form the Navier-Stokes equations with the Lorentz force $\vec{B} \times (\nabla \times \frac{1}{\mu} \vec{B})$ coupling these equations to the electromagnetic equations. The induction equation (2.9c) with the solenoidal constraint (2.9d) compose the equations of MHD kinematics. When we consider the decoupled kinematics model, we refer to the equations (2.9c) - (2.9d).

Now we can nondimensionalize the system by introducing the new quantities

$$\vec{u}_*(\xi) := \frac{\vec{u}(L\xi)}{\bar{u}}, \quad (2.10a)$$

$$\vec{B}_*(\xi) := \frac{\vec{B}(L\xi)}{\bar{B}}, \quad (2.10b)$$

$$p_*(\xi) := \frac{p(L\xi)\bar{u}^2}{\rho}, \quad (2.10c)$$

$$\vec{f}_* := \frac{\vec{f}\bar{L}}{\bar{u}^2}, \quad (2.10d)$$

where \bar{u} , \bar{B} , and \bar{L} are characteristic values for the velocity, magnetic field, and length respectively. Then, the stationary MHD equations can be written in terms of the starred quantities (omitting the stars from our notation for simplicity) as

$$\vec{u} \cdot \nabla \vec{u} - \frac{1}{R} \Delta \vec{u} + \nabla p + S \vec{B} \times (\nabla \times \vec{B}) = \vec{f}, \quad (2.11a)$$

$$\nabla \cdot \vec{u} = 0, \quad (2.11b)$$

$$\frac{1}{R_m} \nabla \times (\nabla \times \vec{B}) - \nabla \times (\vec{u} \times \vec{B}) = \vec{0}, \quad (2.11c)$$

$$\nabla \cdot \vec{B} = 0, \quad (2.11d)$$

where the nondimensional scalars

$$R = \frac{\rho \bar{u} \bar{L}}{\nu}, \quad (2.12a)$$

$$R_m = \frac{\mu \bar{u} \bar{L}}{\eta}, \quad (2.12b)$$

$$S = \frac{\bar{B}}{\mu \rho \bar{u}^2} \quad (2.12c)$$

are the (fluid) Reynolds number, the magnetic Reynolds number, and the coupling coefficient, respectively. The system of PDEs is nonlinear due to the fluid convection $\vec{u} \cdot \nabla \vec{u}$ as well as the nonlinear coupling in the Lorentz force $S(\nabla \times \vec{B}) \times \vec{B}$ and the magnetic convection $-\nabla \times (\vec{u} \times \vec{B})$. Nondimensionalizing \vec{g} by \bar{u} and \vec{q} by \bar{B} , we complete system (2.11) with the same nondimensional boundary conditions (2.4a) and (2.4b). It is important to observe that (2.11) is overdetermined in that it is a system of $2d + 2$ equations in only $2d + 1$ unknowns. Many strategies exist for incorporating the solenoidal condition (2.11d) into the other three equations to ensure the solvability of the system.

One such strategy is to add a penalty term to the induction equation in its weak formulation which implicitly enforces the solenoidal condition. We refer to this strategy as an exact penalty formulation of the MHD equations. Stable mixed finite element [36] and stabilized equal order finite element [28] formulations based on this strategy have been proposed. We discuss the stable formulation in more detail in Section 2.2.

Another strategy is to introduce a new unknown, a Lagrange multiplier, into the system. By adding the gradient of a Lagrange multiplier to the induction equation with appropriate boundary conditions, the number of unknowns can be made equal

to the number of equations without changing the solution. We refer to strategies of this type as Lagrange multiplier formulations. Stable mixed finite element [59] and stabilized equal order finite element [14, 15] as well as discontinuous Galerkin [40] formulations have been proposed for this form of the equations. We consider the stable formulation in Section 2.3.

A generalization of Lagrange multiplier strategies is the divergence cleaning strategy [19], in which properties of the system of PDEs can be changed by adding different functionals of the Lagrange multiplier to equations (2.11c) and (2.11d). Another strategy for enforcing the solenoidal condition is the vector potential approach [44, 60], in which \vec{B} is defined as the curl of another unknown, i.e. $\vec{B} = \nabla \times \vec{A}$. Because $\nabla \cdot \nabla \times \vec{A} = 0$, this implicitly enforces $\nabla \cdot \vec{B} = 0$.

2.2 A Stable Exact Penalty Finite Element Formulation

In this section, we detail the exact penalty formulation as introduced in [36]. In order to use the exact penalty weak formulation of the MHD equations, we restrict our study to the case where Ω is a bounded convex polyhedron. In this setting, note that $H(\text{curl}, \Omega) \cap H(\text{div}, \Omega)$ is embedded in $(H^1(\Omega))^d$ [16]. Furthermore, to simplify the statement of the weak formulation, we consider homogeneous Dirichlet boundary conditions, i.e. $\vec{g} = \vec{q} = \vec{0}$. Then, we consider weak solutions $\vec{u} \in \mathbf{H}_0^1(\Omega)$, $p \in L^2(\Omega)$, $\vec{B} \in \mathbf{H}_\tau^1(\Omega)$ to the MHD equations (2.11), where $\mathbf{H}_0^1(\Omega) = (H_0^1(\Omega))^d$ and $\mathbf{H}_\tau^1(\Omega) = \{\vec{v} \in (H^1(\Omega))^d \mid \vec{v} \times \vec{n} = 0\}$. Defining the space $W = \mathbf{H}_0^1(\Omega) \times L^2(\Omega) \times \mathbf{H}_\tau^1(\Omega)$ and representing the solution as $U = (\vec{u}, p, \vec{B})$ and the test function as $V = (\vec{v}, q, \vec{C})$,

we state the exact penalty nonlinear weak formulation as: Find $U \in W$ such that

$$\mathcal{N}_{EP}(U, V) = \langle \vec{f}, \vec{v} \rangle, \quad \forall V \in W, \quad (2.13)$$

where the nonlinear form \mathcal{N}_{EP} is

$$\begin{aligned} \mathcal{N}_{EP}(U, V) &= (\vec{u} \cdot \nabla \vec{u}, \vec{v}) + \frac{1}{R}(\nabla \vec{u}, \nabla \vec{v}) - (p, \nabla \cdot \vec{v}) + (q, \nabla \cdot \vec{u}) \\ &\quad + S(\vec{v} \times \vec{B}, \nabla \times \vec{B}) - S(\vec{u} \times \vec{B}, \nabla \times \vec{C}) \\ &\quad + \frac{S}{R_m}(\nabla \times \vec{B}, \nabla \times \vec{C}) + \frac{S}{R_m}(\nabla \cdot \vec{B}, \nabla \cdot \vec{C}). \end{aligned} \quad (2.14)$$

(We use the subscript EP to indicate that the form \mathcal{N}_{EP} is associated with the exact penalty formulation.) Every term in this form except the final one is obtained from multiplying (2.11a) by \vec{v} , (2.11b) by q , and (2.11c) by $S\vec{C}$, integrating by parts, and summing. The term $\frac{S}{R_m}(\nabla \cdot \vec{B}, \nabla \cdot \vec{C})$ is the exact penalty term, included to weakly enforce the solenoidal constraint (2.11d). This follows from the fact that if $\vec{B} \in \mathbf{H}^1(\Omega)$ and Ω is a bounded convex polyhedron, then there exists a scalar $c \in H^2(\Omega)$ such that

$$\nabla \cdot \nabla c = \nabla \cdot \vec{B}, \quad (2.15)$$

with $c = 0$ on $\partial\Omega$ and $\nabla c \in \mathbf{H}_\tau^1(\Omega)$ [36]. Letting $V = (\vec{0}, 0, \nabla c)$, we obtain from (2.13) that $(\nabla \cdot \vec{B}, \nabla \cdot \vec{B}) = 0$, and hence that (2.11d) is enforced almost everywhere in Ω .

Linearizing the nonlinear weak formulation (2.13) leads to a set of systems of the form

$$\mathcal{B}^{(n)}(\delta U, V) = R^{(n)}, \quad \forall V \in W, \quad (2.16a)$$

$$U^{(n+1)} = U^{(n)} + \delta U, \quad (2.16b)$$

where $R^{(n)}$ is the nonlinear residual and $\mathcal{B}^{(n)}$ is a bilinear form defined by the linearization method. Two common linearization techniques are Picard iteration and Newton's method. A version of Picard iteration that leads to a coercive weak form gives the bilinear form

$$\begin{aligned} \mathcal{B}_{EP,P}^{(n)}(U, V) &= (\vec{a} \cdot \nabla \vec{u}, \vec{v}) + \frac{1}{R}(\nabla \vec{u}, \nabla \vec{v}) - (p, \nabla \cdot \vec{v}) + (q, \nabla \cdot \vec{u}) \\ &\quad + S(\vec{v} \times \vec{b}, \nabla \times \vec{B}) - S(\vec{u} \times \vec{b}, \nabla \times \vec{C}) \\ &\quad + \frac{S}{R_m}(\nabla \times \vec{B}, \nabla \times \vec{C}) + \frac{S}{R_m}(\nabla \cdot \vec{B}, \nabla \cdot \vec{C}), \end{aligned} \quad (2.17)$$

where

$$\vec{a} = \vec{u}^{(n)}, \quad \vec{b} = \vec{B}^{(n)} \quad (2.18)$$

are the previous velocity and magnetic field in the Picard iteration. This results in linear forms of the fluid convection $(\vec{a} \cdot \nabla \vec{u}, \vec{v})$, the Lorentz force $S(\vec{v} \times \vec{b}, \nabla \times \vec{B})$ and the magnetic convection $-S(\vec{u} \times \vec{b}, \nabla \times \vec{C})$. It is known that the bilinear form $\mathcal{B}_{EP,P}^{(n)}$ is continuous and that an associated inf-sup stability condition is satisfied on W [36]. Thus, when a unique solution to (2.13) exists, there exists a unique solution to each linear problem in the above Picard iteration. Furthermore, the nonlinear iteration converges to the unique solution to (2.13) from any initial iterate.

Newton's method takes as its bilinear form

$$\begin{aligned} \mathcal{B}_{EP,N}^{(n)}(U, V) &= (\vec{a} \cdot \nabla \vec{u}, \vec{v}) + (\vec{u} \cdot \nabla \vec{a}, \vec{v}) + \frac{1}{R}(\nabla \vec{u}, \nabla \vec{v}) - (p, \nabla \cdot \vec{v}) + (q, \nabla \cdot \vec{u}) \\ &\quad + S(\vec{v} \times \vec{b}, \nabla \times \vec{B}) + S(\vec{v} \times \vec{B}, \nabla \times \vec{b}) - S(\vec{u} \times \vec{b}, \nabla \times \vec{C}) \\ &\quad + \frac{S}{R_m}(\nabla \times \vec{B}, \nabla \times \vec{C}) + \frac{S}{R_m}(\nabla \cdot \vec{B}, \nabla \cdot \vec{C}) - S(\vec{a} \times \vec{B}, \nabla \times \vec{C}). \end{aligned} \quad (2.19)$$

It can be shown that this system is well-posed and Newton's method converges

provided that the initial iterate is close enough to the unique solution of (2.13) [36].

In this case, Newton's method converges quadratically.

Analogous results are proven when the finite element method is applied and W is replaced with a finite dimensional subspace W_h . In this case, the major difference is that the inf-sup condition must be satisfied on W_h , i.e. we require there to exist a constant $\beta > 0$ such that

$$\inf_{q \in W_{\mathbf{p}}} \sup_{\vec{v} \in W_{\mathbf{u}}} \frac{(\nabla \cdot \vec{v}, q)}{\|\vec{v}\|_1 \|q\|_0} \geq \beta, \quad (2.20)$$

where $W_{\mathbf{p}}$ is the discrete space for p and $W_{\mathbf{u}}$ is the discrete space for \vec{u} . Note that \vec{B} does not appear in this condition, and it is exactly the same condition required for stability of the discrete Navier-Stokes equations [24]. Thus, there is no restriction on the discrete space chosen to approximate \vec{B} , and any stable element pair for the Navier-Stokes equations (e.g. Q_2 - Q_1 Taylor-Hood elements) can be used for \vec{u} and p . For ease of implementation, we will discretize \vec{B} using the same finite-dimensional space used for \vec{u} (e.g. Q_2 - Q_1 - Q_2 elements for \vec{u} - p - \vec{B}).

2.3 A Stable Lagrange Multiplier Finite Element Formulation

While the exact penalty formulation holds for convex domains Ω , it may not capture important features of the magnetic induction in non-convex domains. In general, \vec{B} may have regularity below $(H^1(\Omega))^d$. Because the nodal elements used to approximate \vec{B} in the exact penalty formulation form a finite-dimensional basis for $(H^1(\Omega))^d$, the method can converge to a magnetic field that may not capture physical singularities induced by reentrant corners [16]. One approach to account

for these singularities is to let the unknown \vec{B} reside in the less regular Sobolev space $H(\text{curl}, \Omega)$. In [58, 59], this approach is applied to the MHD equations in the form

$$\vec{u} \cdot \nabla \vec{u} - \frac{1}{R} \Delta \vec{u} + \nabla p + S \vec{B} \times (\nabla \times \vec{B}) = \vec{f}, \quad (2.21a)$$

$$\nabla \cdot \vec{u} = 0, \quad (2.21b)$$

$$\frac{S}{R_m} \nabla \times (\nabla \times \vec{B}) - S \nabla \times (\vec{u} \times \vec{B}) - \nabla r = \vec{0}, \quad (2.21c)$$

$$\nabla \cdot \vec{B} = 0. \quad (2.21d)$$

These are essentially the same equations as (2.11), except for the additional term $-\nabla r$ included in the induction equation. This introduces the new unknown r , a Lagrange multiplier for the electromagnetics, sometimes referred to as the magnetic pseudo-pressure. With this additional unknown, (2.21) form a well-posed system of $2d + 2$ equations in $2d + 2$ unknowns. Notice that taking the divergence of (2.21c) yields $-\Delta r = 0$. Thus, with the boundary condition $r = 0$ on $\partial\Omega$, r must be identically zero. Hence, if the system is completed with the boundary conditions

$$\vec{u} = \vec{g}, \quad (2.22a)$$

$$\vec{B} \times \vec{n} = \vec{q} \times \vec{n}, \quad (2.22b)$$

$$r = 0 \quad (2.22c)$$

on $\partial\Omega$, the augmented MHD equations (2.21) admit the same solutions \vec{u} , p , and \vec{B} as system (2.11).

If we let $\vec{q} = \vec{0}$, the electromagnetic pair (\vec{B}, r) can be assumed to reside in the space $H_0(\text{curl}, \Omega) \times H_0^1(\Omega)$, where $H_0(\text{curl}, \Omega) = \{\vec{v} \in H(\text{curl}, \Omega) | \vec{v} \times \vec{n} =$

$\vec{0} \times \vec{n}$ on $\partial\Omega$ is the subspace of $H(\text{curl}, \Omega)$ with zero tangential component on the boundary. These spaces have been used previously to pose a mixed finite element formulation for Maxwell's equations in mixed form where stability and well-posedness for this pair was proved [20]. As in the exact penalty formulation, the hydrodynamic pair (\vec{u}, p) is assumed to reside in the traditional space $(H_0^1(\Omega))^d \times L^2(\Omega)$ if $\vec{g} = \vec{0}$.

Thus, we define the solution $U = (\vec{u}, p, \vec{B}, r)$, the test function $V = (\vec{v}, q, \vec{C}, s)$, and the solution space $W = (H_0^1(\Omega))^d \times L^2(\Omega) \times H_0(\text{curl}, \Omega) \times H_0^1(\Omega)$. Then a nonlinear weak formulation can be posed as: Find $U \in W$ such that

$$\mathcal{N}_{LM}(U, V) = \langle \vec{f}, \vec{v} \rangle, \quad \forall V \in W, \quad (2.23)$$

where the nonlinear form \mathcal{N}_{LM} is

$$\begin{aligned} \mathcal{N}_{LM}(U, V) &= (\vec{u} \cdot \nabla \vec{u}, \vec{v}) + \frac{1}{R}(\nabla \vec{u}, \nabla \vec{v}) - (p, \nabla \cdot \vec{v}) + (q, \nabla \cdot \vec{u}) \\ &\quad + S(\vec{v} \times \vec{B}, \nabla \times \vec{B}) - S(\vec{u} \times \vec{B}, \nabla \times \vec{C}) \\ &\quad + \frac{S}{R_m}(\nabla \times \vec{B}, \nabla \times \vec{C}) - (\nabla r, \vec{C}) + (\nabla s, \vec{B}). \end{aligned} \quad (2.24)$$

This form is obtained from multiplying (2.21a) by \vec{v} , (2.21b) by q , (2.21c) by \vec{C} , and (2.21d) by $-s$, integrating by parts, and summing. Note that by leaving the term $(\nabla r, \vec{C})$ and integrating $-(\nabla \cdot \vec{B}, s)$ by parts to obtain $(\nabla s, \vec{B})$, we ensure that no additional regularity is required of the magnetic field beyond $\vec{B} \in H(\text{curl}, \Omega)$. The formulation (2.23) is continuous and coercive, and the spaces $(H_0^1(\Omega))^d \times L^2(\Omega)$ and $H_0(\text{curl}, \Omega) \times H_0^1(\Omega)$ satisfy inf-sup conditions for both the hydrodynamic pair (\vec{u}, p) and the electromagnetic pair (\vec{B}, r) . Thus, a solution to (2.23) exists and is unique for sufficiently small data R, R_m, S , and \vec{f} [59].

The nonlinear formulation can be linearized to obtain an iteration of the form

(2.16). A Picard linearization yields the bilinear form

$$\begin{aligned}
\mathcal{B}_{LM,P}^{(n)}(U, V) &= (\vec{a} \cdot \nabla \vec{u}, \vec{v}) + \frac{1}{R}(\nabla \vec{u}, \nabla \vec{v}) - (p, \nabla \cdot \vec{v}) + (q, \nabla \cdot \vec{u}) \\
&\quad + S(\vec{v} \times \vec{b}, \nabla \times \vec{B}) - S(\vec{u} \times \vec{b}, \nabla \times \vec{C}) \\
&\quad + \frac{S}{R_m}(\nabla \times \vec{B}, \nabla \times \vec{C}) - (\nabla r, \vec{C}) + (\nabla s, \vec{B})
\end{aligned} \tag{2.25}$$

while Newton's method yields

$$\begin{aligned}
\mathcal{B}_{LM,N}^{(n)}(U, V) &= (\vec{a} \cdot \nabla \vec{u}, \vec{v}) + (\vec{u} \cdot \nabla \vec{a}, \vec{v}) + \frac{1}{R}(\nabla \vec{u}, \nabla \vec{v}) - (p, \nabla \cdot \vec{v}) + (q, \nabla \cdot \vec{u}) \\
&\quad + S(\vec{v} \times \vec{b}, \nabla \times \vec{B}) + S(\vec{v} \times \vec{B}, \nabla \times \vec{b}) - S(\vec{u} \times \vec{b}, \nabla \times \vec{C}) \\
&\quad + \frac{S}{R_m}(\nabla \times \vec{B}, \nabla \times \vec{C}) - S(\vec{a} \times \vec{B}, \nabla \times \vec{C}) - (\nabla r, \vec{C}) + (\nabla s, \vec{B}),
\end{aligned} \tag{2.26}$$

where

$$\vec{a} = \vec{u}^{(n)}, \quad \vec{b} = \vec{B}^{(n)}. \tag{2.27}$$

Both Picard and Newton linearizations converge to the unique solution of the non-linear problem when a unique solution exists. The Picard iteration converges for any initial guess, and Newton's method converges for initial guesses sufficiently close to the exact solution [59].

When the weak formulation is posed on a finite-dimensional subspace of W , the finite-dimensional hydrodynamic spaces $W_{\mathbf{u}} \subset (H^1(\Omega))^d$ and $W_{\mathbf{p}} \subset L^2(\Omega)$ and the finite-dimensional electromagnetic spaces $W_{\mathbf{B}} \subset H(\text{curl}, \Omega)$ and $W_{\mathbf{r}} \subset H^1(\Omega)$

must satisfy the two inf-sup stability conditions

$$\inf_{q \in W_{\mathbf{p}}} \sup_{\vec{v} \in W_{\mathbf{u}}} \frac{(\nabla \cdot \vec{v}, q)}{\|\vec{v}\|_1 \|q\|_0} \geq \beta_{\mathbf{u}} > 0, \quad (2.28a)$$

$$\inf_{s \in W_{\mathbf{r}}} \sup_{\vec{C} \in W_{\mathbf{B}}} \frac{(\nabla s, \vec{C})}{\|\vec{C}\|_{curl} \|s\|_1} \geq \beta_{\mathbf{B}} > 0. \quad (2.28b)$$

Condition (2.28a) is again the stability condition obtained for the Navier-Stokes equations, and any stable element pair for the Navier-Stokes equations (e.g. Q_2 - Q_1 nodal elements) can be used to approximate \vec{u} and p . The electromagnetic stability condition (2.28b) is satisfied if Nédélec's edge elements of the first kind [50] are used to discretize \vec{B} and standard nodal elements are used to discretize r (e.g. first order edge elements for \vec{B} and Q_1 elements for r). Nédélec's edge elements are constructed to span a finite-dimensional subspace of $H(curl, \Omega)$. Furthermore, with $W_{\mathbf{B}}$ and $W_{\mathbf{r}}$ defined in this way, the condition $\nabla s \in W_{\mathbf{B}}$ holds for all $s \in W_{\mathbf{r}}$ [50]. Thus, we can set $\vec{C} = \nabla s$ in (2.28b) to obtain

$$\inf_{s \in W_{\mathbf{r}}} \sup_{\vec{C} \in W_{\mathbf{B}}} \frac{(\nabla s, \vec{C})}{\|\vec{C}\|_{curl} \|s\|_1} \geq \inf_{s \in W_{\mathbf{r}}} \frac{\|\nabla s\|_0^2}{\|\nabla s\|_0 \|s\|_1} \geq \frac{1}{\gamma}, \quad (2.29)$$

where γ is a Poincaré constant depending only on the domain Ω and satisfying $\|s\|_1 \leq C \|\nabla s\|_0$. This implies that (2.28b) holds and the inf-sup constant $\beta_{\mathbf{B}}$ depends only on Ω .

2.4 Preconditioners for the Discretized MHD Equations

After discretization, a linear system

$$\mathcal{A}\mathbf{x} = \mathbf{f} \quad (2.30)$$

corresponding to the linearized problem (2.16) must be solved at each step of a nonlinear iteration in order to obtain a discrete solution to the original nonlinear problem. Here, \mathbf{x} is a vector containing the coefficients of δU in terms of the finite element basis and \mathbf{f} contains the coefficients of the residual $R^{(n)}$. The structure of the matrix \mathcal{A} depends on the bilinear form derived from a particular weak formulation and linearization scheme as well as the finite elements chosen for the discretization.

In general \mathcal{A} is a large, sparse, nonsymmetric, and indefinite matrix. Because \mathcal{A} is large and sparse, iterative Krylov subspace methods (e.g. GMRES [56]) should be particularly effective for solving system (2.30). The performance of these algorithms is influenced by the distribution of the eigenvalues of \mathcal{A} , converging fastest when there are few distinct (clusters of) eigenvalues. The indefinite nature of \mathcal{A} suggests that preconditioning is necessary for fast convergence of such methods; that is, rather than applying the iterative method directly to system (2.30), it is applied to the system

$$(\mathcal{A}\mathcal{P}^{-1})\mathbf{y} = \mathbf{f}, \quad (2.31)$$

where $\mathbf{x} = \mathcal{P}^{-1}\mathbf{y}$. The challenge is then to choose a preconditioner \mathcal{P} such that the action of \mathcal{P}^{-1} is inexpensive and the preconditioned operator $\mathcal{A}\mathcal{P}^{-1}$ has a desirable features, for example, tightly distributed eigenvalues.

2.4.1 Block Preconditioners for Coupled Linear Systems

Before detailing preconditioners that have been developed specifically for the discretized MHD equations, we first provide background on a general class of pre-

conditioners developed for coupled linear systems. A linear system with two coupled vector unknowns \mathbf{x} and \mathbf{y} can be written in the generalized saddle-point form

$$\begin{pmatrix} A & B \\ C & -D \end{pmatrix} \begin{pmatrix} \mathbf{x} \\ \mathbf{y} \end{pmatrix} = \begin{pmatrix} \mathbf{f} \\ \mathbf{g} \end{pmatrix} \quad (2.32)$$

(for a review of solution methods for systems of this type, see [5]). Both the discretized incompressible Navier-Stokes equations and the discretized Maxwell's equations in mixed form can be written in the form (2.32) when blocked in terms of their physical unknowns (\mathbf{u}, \mathbf{p} and \mathbf{B}, \mathbf{r} , respectively). Depending on the particular formulation, the discretized MHD equations can also be expressed in this form when the unknowns are blocked appropriately.

The motivation behind block preconditioning is to take advantage of the block structure of the linear system (2.32) to decouple the system into its component parts. For example, if A is nonsingular, a classic block preconditioning strategy [49] is to consider the block LU decomposition

$$\begin{pmatrix} A & B \\ C & -D \end{pmatrix} = \begin{pmatrix} I & 0 \\ CA^{-1} & I \end{pmatrix} \begin{pmatrix} A & B \\ 0 & -S \end{pmatrix}, \quad (2.33)$$

where S is the Schur complement

$$S = D + CA^{-1}B. \quad (2.34)$$

A Krylov subspace method preconditioned with the block upper triangular factor will converge in at most two iterations [49]. Hence, an approximation of the block upper triangular factor should be an effective preconditioner. The advantage of such preconditioners is that the form of the preconditioner decouples the two unknowns

\mathbf{x} and \mathbf{y} . The original coupling of the system is algebraically embedded in the Schur complement S . The challenge in developing a preconditioner is then to effectively approximate the action of the inverse of the operator S .

The block form of the preconditioner as well as the strategy for approximating S generally depend upon the particular application from which the linear system arises. For example, many strategies have been developed to approximate the Schur complement arising in the discretized Navier-Stokes equations, including purely algebraic approximations and approximations based on assumptions that discrete operators commute in the same way that the corresponding differential operators do (see [22] for a review of these methods). For Maxwell's equations discretized in mixed form, the block corresponding to A in (2.32), a discretization of the operator $\nabla \times \nabla \times$, is singular and the system does not admit the factorization (2.33). Consequently, preconditioners that augment A to make it nonsingular have been proposed [35, 63, 64].

2.4.2 Existing Preconditioners for Fully Coupled MHD

A growing body of recent research is concerned with the development of preconditioners for linear systems arising from various formulations and discretizations of the fully coupled MHD equations. Several references have used Schwarz domain decomposition methods as preconditioners for the full MHD system [17, 44, 55, 60]. These methods exploit parallelism by constructing and solving smaller problems (e.g. with an incomplete LU factorization) on a set of overlapping subdomains.

The local solutions are then combined to construct a global approximation. While Schwarz methods are easily parallelized, these preconditioners tend to lead to increased iteration counts as the number of subdomains increases. However, it has been demonstrated that using such domain decomposition methods as smoothers for algebraic multigrid produce an effective preconditioner for the full MHD system [44, 60]. This approach was applied to a stabilized equal-order finite element vector potential formulation of the incompressible MHD equations. Multigrid has also been applied to the fully coupled linear system arising in a first-order system least-squares (FOSLS) finite element discretization of the incompressible MHD equations [1].

Another approach is to use operator splitting techniques as preconditioners. For example, split preconditioners were proposed for the viscoresistive MHD equations posed in cylindrical coordinates for tokamak applications [55]. In one such preconditioner, the three-dimensional MHD problem is approximated by ignoring the coupling between the three coordinate directions and thereby reducing the system to three one-dimensional problems. Another preconditioner takes advantage of the physics of the tokamak and splits the system into two-dimensional poloidal and one-dimensional toroidal subsystems.

A series of references has employed parabolic reformulations of the MHD equations to develop preconditioners [11–13]. This approach has been applied to a reduced MHD formulation, a compressible viscoresistive formulation, and the compressible Hall MHD model. The idea is to reformulate the hyperbolic MHD system as a diagonally dominant parabolic system so that classically smoothed multigrid

is effective on the reformulated system. The parabolization of the system is effectively equivalent to performing a block LDU decomposition on the discretized MHD equations and preconditioning the linear system by the block diagonal factor. This preconditioning strategy is made practical by approximating the Schur complement that appears. For compressible MHD, this has been done by either approximating the system at the small flow limit [11] or by assuming that operators appearing within the Schur complement commute [12].

Both preconditioners based on operator splitting and parabolic reformulation can be viewed as variants on the block preconditioning strategy. Another block preconditioner based on an approximate block factorization motivated by operator splitting was developed for a vector potential formulation of the incompressible MHD equations [17]. This preconditioner decouples the full system into one system accounting for the effect of the velocity on the magnetic field and another system corresponding to the Navier-Stokes equations. Another set of block preconditioners has been proposed for the inductionless MHD equations in which the unknowns are \vec{u}, p , the current density \vec{j} , and the electric potential ϕ [2]. These preconditioners are developed by considering LU decompositions of the discretized MHD equations with different orderings of the unknowns. The Schur complements are approximated by selectively incorporating or ignoring the effects of coupling terms.

In this dissertation, we continue in the line of block preconditioners for MHD by developing preconditioners for the exact penalty and Lagrange multiplier formulations detailed in Sections 2.2 and 2.3. We consider block LU decompositions of the associated linear systems that embed the coupling between the hydrody-

dynamic and electromagnetic processes into Schur complements. Building on preconditioning ideas that were introduced for the discretized Navier-Stokes and Maxwell's equations, we develop approximations to these Schur complements that attempt to account for the effects of coupling. While other block preconditioning techniques have made approximations that decouple the two physical processes, by focusing on accurately approximating the effects of coupling, we seek to develop preconditioners that perform particularly well for strongly coupled MHD systems.

2.5 Uncertainty Quantification

Uncertainty quantification is a growing field in scientific computing focused on developing models to simulate physical systems in which uncertainty may be present. This uncertainty may be due to either the inherent irregularity of the physical processes involved, meaning that a deterministic description is impossible, or a lack of data about the physical phenomenon [31]. For example, a physical system to which uncertainty quantification has been applied extensively is flow through a porous medium such as soil (see, e.g., [30] and the references therein). In this setting, properties of the soil medium, such as the hydraulic conductivity, are generally not known with certainty as it is impractical to take the appropriate measurements at an exhaustive number of points in the domain. Instead, properties of the medium are modeled as random fields for which there are many realizations.

If a deterministic model exists when material properties are assumed to be known and not random, the uncertain physical system can be modeled by a stochas-

tic version of this model. For example, flow through a porous medium is modeled by a diffusion equation. When the hydraulic conductivity of the soil is treated as uncertain, the governing equation is a diffusion equation with random data, i.e.

$$-\nabla \cdot (a(\vec{x}, \omega) \nabla u(\vec{x}, \omega)) = 0, \quad (2.35)$$

for \vec{x} in the physical domain \mathcal{D} and ω in the probabilistic sample space Ω [30]. The diffusivity a is then a random field and the solution u , because it depends on a , is also random.

Two main questions for problems of this type are then how to represent the random data to be both reflective of the uncertainty in the physical system and computable, and how to obtain statistical data about the quantities of interest. The first question generally depends on the problem. For the random diffusion equation, several references (e.g. [23, 65, 66]) express the diffusivity as a weighted sum of identically distributed and independent random variables, i.e.

$$a(\vec{x}, \omega) = a(\vec{x}, \vec{y}) = a_0(\vec{x}) + \sum_{i=1}^M a_i(\vec{x}) y_i(\omega), \quad (2.36)$$

where $\vec{y} = (y_1(\omega), \dots, y_M(\omega))$ is a vector of random variables (a truncated Karhunen-Loève expansion [31], for instance, takes this form). Alternatively, the diffusivity can be expressed as a polynomial chaos expansion [30].

Once the physical uncertainty is expressed as random data, various computational methods can be used to obtain statistical data about quantities of interest. Three of the most popular strategies are Monte-Carlo methods, stochastic Galerkin methods, and stochastic collocation methods [65]. In a Monte-Carlo simulation [10], the data is repeatedly sampled, and for each realization, a deterministic problem

is solved for the quantity of interest. Statistical data is then approximated by the solution statistics of the sample. For example, at realization j , the data $a(\vec{x}, \vec{y}^j)$ is generated, and (2.35) is solved for u_j . Then, after K realizations, $E(u)$ can be approximated by the sample mean $\frac{1}{K} \sum_{j=1}^K u_j$. Traditional Monte-Carlo converges slowly at a rate of $\frac{1}{\sqrt{K}}$ but independently of the number of random variables in the expansion (2.36). Monte-Carlo methods also require only the solution of deterministic problems for different realizations of the data.

Stochastic Galerkin methods [31] treat each random variable in the expansion (2.36) as an unknown and apply a finite element-like methodology to solve a weak version of the stochastic problem. Each random variable is projected onto a discrete space spanned by a basis of polynomials. Discretizing both the spatial domain and the stochastic domain, the random PDE reduces to a single linear system, its size being a product of the number of spatial degrees of freedom and the number of stochastic degrees of freedom. Because the number of stochastic degrees of freedom depends on the number of random variables M and the degree of the polynomial space, this linear system can be very large for large values of M .

Collocation methods [23,65] combine the strengths of both Monte-Carlo methods and stochastic Galerkin methods. Solving deterministic problems at a finite number of realizations of the data provides values of u at distinct points in the space of random variables. Using polynomial interpolation, u can be interpolated as a function of \vec{y} to obtain a finite dimensional approximation. This interpolant can then be used to obtain statistical data. If sparse grids are used to provide the collocation points, a relatively small number of deterministic problems need to be

solved to generate the interpolation values, but the number of collocation points increases with the number of random variables.

These types of uncertainty quantification techniques have been applied to several physical systems. To our knowledge, there has been no published study investigating the effects of uncertain data on either the fully coupled MHD system or the kinematics model of MHD.

Chapter 3: Block Preconditioners for an Exact Penalty MHD Formulation

In this chapter, we consider the linear systems obtained when the MHD equations are discretized according to the exact penalty finite element formulation introduced in Section 2.2. Discretizing the linearized formulation (2.16) with a stable $\vec{u} - p$ element pair, we obtain a sequence of linear systems of the form

$$\mathcal{A}\mathbf{x} = \mathbf{f}, \tag{3.1}$$

where $\mathbf{x} = (\mathbf{u}, \mathbf{p}, \mathbf{B})$ contains the coefficients of the discrete solution δU , \mathbf{f} is the discrete nonlinear residual, and \mathcal{A} is the discretization of the weak form. (We will represent vector coefficients in boldface.) The structure of \mathcal{A} depends on the bilinear form associated with the linearization. The matrices resulting from the Picard and Newton linearizations can be written in block form as

$$\mathcal{A}_P = \begin{pmatrix} F & B^t & Z \\ B & 0 & 0 \\ -Z^t & 0 & A \end{pmatrix}, \quad \mathcal{A}_N = \begin{pmatrix} F + \tilde{F} & B^t & Z + \tilde{Z} \\ B & 0 & 0 \\ -Z^t & 0 & A + \tilde{A} \end{pmatrix}, \tag{3.2}$$

respectively, where the component matrices derive from continuous operators as in Table 3.1. Note that the extra terms \tilde{F} , \tilde{Z} , and \tilde{A} are due to the additional coupling in the Newton weak form, corresponding to the terms in the right-hand

Discrete	Continuous	Interpretation	Approximate Norm
$A\mathbf{B}$	$-\frac{S}{R_m}\Delta\vec{B}$	Magnetic diffusion	$\frac{S}{R_m h^2}$
$\tilde{A}\mathbf{B}$	$-S\nabla\times(\vec{a}\times\vec{B})$	Magnetic convection	$\frac{S\ \vec{a}\ }{h}$
$-Z^t\mathbf{u}$	$-S\nabla\times(\vec{u}\times\vec{b})$	Magnetic convection	$\frac{S\ \vec{b}\ }{h}$
$Z\mathbf{B}$	$S\vec{b}\times(\nabla\times\vec{B})$	Lorentz force	$\frac{S\ \vec{b}\ }{h}$
$\tilde{Z}\mathbf{B}$	$S\vec{B}\times(\nabla\times\vec{b})$	Lorentz force	$S\ \nabla\times\vec{b}\ $
$F\mathbf{u}$	$\vec{a}\cdot\nabla\vec{u}-\frac{1}{R}\Delta\vec{u}$	Fluid convection-diffusion	$\frac{\ \vec{a}\ }{h}+\frac{1}{Rh^2}$
$\tilde{F}\mathbf{u}$	$\vec{u}\cdot\nabla\vec{a}$	Fluid convection	$\ \nabla\vec{a}\ $
$B^t\mathbf{p}$	∇p	Pressure gradient	$\frac{1}{h}$
$-B\mathbf{u}$	$\nabla\cdot\vec{u}$	Divergence	$\frac{1}{h}$

Table 3.1: Definitions of discrete operators as they correspond to continuous operators.

side of the identity

$$\mathcal{B}_{EP,N}^{(n)} - \mathcal{B}_{EP,P}^{(n)} = (\vec{u} \cdot \nabla \vec{a}, \vec{v}) + S(\vec{v} \times \vec{B}, \nabla \times \vec{b}) - S(\vec{a} \times \vec{B}, \nabla \times \vec{C}). \quad (3.3)$$

For either linearization, \mathcal{A} is a large, sparse matrix. Thus, for efficiency, a preconditioned iterative method should be considered for solving the systems (3.1). Because \mathcal{A} is nonsymmetric and indefinite, we use preconditioned GMRES for these solves.

3.1 A Block Preconditioner for the 2D Picard System

We consider preconditioning \mathcal{A}_P using a strategy based on approximating Schur complements that generalizes techniques commonly employed for discretizations of the Navier-Stokes equations. To motivate our preconditioning strategy, we consider the block LU decomposition of \mathcal{A}_P ,

$$\mathcal{A}_P = \begin{pmatrix} I & 0 & 0 \\ BF^{-1} & I & 0 \\ -Z^t F^{-1} & Z^t F^{-1} B^t X^{-1} & I \end{pmatrix} \begin{pmatrix} F & B^t & Z \\ 0 & X & -BF^{-1}Z \\ 0 & 0 & Y \end{pmatrix}, \quad (3.4)$$

where

$$X = -BF^{-1}B^t, \quad (3.5)$$

$$Y = A + Z^t F^{-1}Z + Z^t F^{-1}B^t X^{-1}BF^{-1}Z. \quad (3.6)$$

It is easy to show that the minimum polynomial for the block lower triangular factor $L = \mathcal{A}_P U^{-1}$ is $(I - L)^3$. The minimum polynomial is cubic, which implies that if we could use U as a right preconditioner for \mathcal{A}_P , then preconditioned GMRES would converge in at most three iterations [49]. In practice, it is infeasible to apply the

action of U^{-1} exactly. Hence, we construct preconditioners by developing techniques for approximating the actions of the inverses of the matrices on the block diagonal of the upper triangular factor. Thus, we consider the preconditioner

$$\mathcal{P}_P = \begin{pmatrix} \hat{F} & B^t & Z \\ 0 & \hat{X} & -B\hat{F}^{-1}Z \\ 0 & 0 & \hat{Y} \end{pmatrix}, \quad (3.7)$$

where “hatted” operators indicate approximations. The convection-diffusion operator F can be handled well by multigrid, and many effective approximations exist for the pressure Schur complement X arising in discretizations of the Navier-Stokes equations [22]. The new difficulty is the Schur complement Y associated with the magnetic field. The nesting of multiple inverse operators as well as the summing of several terms within Y presents an additional challenge in developing expressions for \hat{Y} .

Note that the structure of \mathcal{P}_P in (3.7) derives from the ordering used for the components of \mathbf{x} , $(\mathbf{u}, \mathbf{p}, \mathbf{B})$. If the components are reordered, then \mathcal{A}_P has a different block structure, as do the resulting block LU decompositions. We consider only reorderings in which the rows and columns have the same ordering, so that square blocks remain on the diagonal. The only two other orderings that permit block LU decompositions are $(\mathbf{u}, \mathbf{B}, \mathbf{p})$ and $(\mathbf{B}, \mathbf{u}, \mathbf{p})$. Like $(\mathbf{u}, \mathbf{p}, \mathbf{B})$, the ordering $(\mathbf{u}, \mathbf{B}, \mathbf{p})$ gives rise to Schur complements that are nested, multi-term, and for this reason, we will not pursue this ordering further. We note that preconditioning a system similar in structure to that obtained from the $(\mathbf{u}, \mathbf{p}, \mathbf{B})$ ordering has been studied from another perspective in [17].

The ordering $(\mathbf{B}, \mathbf{u}, \mathbf{p})$ gives the expressions below. (We use the same notation for the complete matrices as above despite the change in ordering. Thus, for the remainder of this chapter, $\mathcal{A}_P, \mathcal{A}_N, X$, and Y refer to operators arising from the $(\mathbf{B}, \mathbf{u}, \mathbf{p})$ ordering.)

$$\mathcal{A}_P = \begin{pmatrix} A & -Z^t & 0 \\ Z & F & B^t \\ 0 & B & 0 \end{pmatrix}, \quad \mathcal{A}_N = \begin{pmatrix} A + \tilde{A} & -Z^t & 0 \\ Z + \tilde{Z} & F + \tilde{F} & B^t \\ 0 & B & 0 \end{pmatrix}. \quad (3.8)$$

We then have the block LU decomposition

$$\mathcal{A}_P = \begin{pmatrix} I & 0 & 0 \\ ZA^{-1} & I & 0 \\ 0 & BX^{-1} & I \end{pmatrix} \begin{pmatrix} A & -Z^t & 0 \\ 0 & X & B^t \\ 0 & 0 & Y \end{pmatrix}, \quad (3.9)$$

where

$$X = F + ZA^{-1}Z^t, \quad (3.10)$$

$$Y = -BX^{-1}B^t. \quad (3.11)$$

The performance of a preconditioner based on this factorization is now contingent upon developing effective approximations for the operators X and Y . Observe that X here can be viewed as a perturbed convection-diffusion operator. In Section 3.1.1, we will expand on the nature of the perturbation $ZA^{-1}Z^t$. In this light, we argue that strategies developed for the Navier-Stokes Schur complement $-BF^{-1}B^t$ can be employed to approximate Y . This will be demonstrated in Section 3.1.2.

3.1.1 An Approximation for X

Consider the analogue of (3.10) in the continuous space, i.e., where the discrete operators are replaced with their corresponding continuous operators from table 3.1. Then X can be viewed as an approximation to the continuous operator

$$\mathcal{F} + \mathcal{K}, \quad (3.12)$$

where

$$\mathcal{F}\vec{u} := \vec{a} \cdot \nabla \vec{u} - R^{-1} \Delta \vec{u} \quad (3.13)$$

is the convection-diffusion operator and

$$\begin{aligned} \mathcal{K}\vec{u} &:= S\vec{b} \times \left\{ \nabla \times \left(-\frac{S}{R_m} \Delta \right)^{-1} [S \nabla \times (\vec{u} \times \vec{b})] \right\} \\ &= -SR_m \vec{b} \times \nabla \times \Delta^{-1} \nabla \times (\vec{u} \times \vec{b}) \end{aligned} \quad (3.14)$$

is an operator resulting from the coupling between \vec{u} and \vec{B} . For two-dimensional problems, we have the identity

$$-\Delta_p c = \nabla \times \nabla \times c \quad (3.15)$$

for any scalar functions c , where we use the subscript p to indicate the scalar Laplacian as opposed to the vector Laplacian Δ . Furthermore, the Laplacian and the two-dimensional curl operator commute; that is,

$$\Delta \nabla \times c = \nabla \times \Delta_p c. \quad (3.16)$$

Replacing c with $\Delta_p^{-1} c$, this yields

$$\nabla \times \Delta_p^{-1} c = \Delta^{-1} \nabla \times c. \quad (3.17)$$

Together with (3.15), this relation implies that

$$\nabla \times \Delta^{-1} \nabla \times c = -c. \quad (3.18)$$

Applying this identity to the expression (3.14) where $c = \vec{u} \times \vec{b}$, we obtain

$$\mathcal{K}\vec{u} = SR_m \vec{b} \times (\vec{u} \times \vec{b}). \quad (3.19)$$

A discretization of this operator can now easily be constructed and we obtain an approximation for X of the form

$$X \approx \hat{X} := F + K, \quad (3.20)$$

where K is a discretization of \mathcal{K} . K can be written as

$$K = SR_m Q_{\mathbf{u}} \begin{pmatrix} \text{diag}(b_y^2) & \text{diag}(-b_x b_y) \\ \text{diag}(-b_x b_y) & \text{diag}(b_x^2) \end{pmatrix}, \quad (3.21)$$

where $Q_{\mathbf{u}}$ is the velocity mass matrix, b_x and b_y are the x and y components of \vec{b} , and $\text{diag}(c)$ is a diagonal matrix containing the values of the function c at each degree of freedom in the discrete domain. These diagonal matrices can be easily constructed by taking the discrete values of b_x and b_y from the previous Picard iterate. The approximation to X can thus be regarded as a convection-diffusion operator perturbed by a scaled mass matrix. This is similar in structure to a time-dependent convection-diffusion operator, and preconditioning methods for F , such as multigrid, will extend well to approximate the action of X^{-1} .

3.1.2 An Approximation for Y

Now consider the discrete operator Y of (3.11). We will proceed by considering relationships among commutators for the continuous differential operators

corresponding to our discrete operators. This strategy was presented for the Navier-Stokes equations [22], where an approximation to the Schur complement $BF^{-1}B^t$ is needed. It was observed that the divergence and convection-diffusion operators approximately commute, i.e.

$$\nabla \cdot \mathcal{F} \approx \mathcal{F}_p \nabla \cdot, \quad (3.22)$$

where \mathcal{F}_p is a convection-diffusion operator defined on the pressure space, referred to as the pressure convection-diffusion operator. This approximation holds when the convection direction \vec{a} is smooth. A discrete version of (3.22) is given by

$$Q_p^{-1} B Q_u^{-1} F \approx Q_p^{-1} F_p Q_p^{-1} B, \quad (3.23)$$

where Q_p is the pressure mass matrix and F_p is the discrete analogue of \mathcal{F}_p . We will discuss how to construct F_p later. Through algebraic manipulation of the approximation (3.23), the Schur complement approximation

$$BF^{-1}B^t \approx Q_p F_p^{-1} (B Q_u^{-1} B^t) \quad (3.24)$$

is obtained, where the operator $B Q_u^{-1} B^t$ can be treated as a discrete Laplacian operator [24]. The approximation (3.24) is desirable because it circumvents the nested nature of the exact Schur complement so that the action of its inverse entails only matrix multiplications and solves with simple operators.

We could make a similar assumption about a commutator on the operator \mathcal{K} ,

$$\nabla \cdot \mathcal{K} \approx \mathcal{K}_p \nabla \cdot \quad (3.25)$$

where \mathcal{K}_p is an analogue to \mathcal{K} on the pressure space. Then the methodology above

could be applied directly to the operator $F + K$, and we would obtain an approximation analogous to (3.24)

$$B(F + K)^{-1}B^t \approx Q_{\mathbf{p}}(F_{\mathbf{p}} + K_{\mathbf{p}})^{-1}(BQ_{\mathbf{u}}^{-1}B^t) \quad (3.26)$$

However, unlike \mathcal{F} , which is a diagonal operator, \mathcal{K} is a coupled 2D operator, and for this reason, a direct analogue to \mathcal{K} on the pressure space does not exist in general. In fact, there exists an operator \mathcal{K}_p such that (3.25) holds only when $SR_m\vec{b} \approx \vec{0}$. Thus, however \mathcal{K}_p is defined, the error associated with (3.25) must be at least proportional to $SR_m\|\vec{b}\|$. Based on this observation, we propose a continuous Schur complement operator of the form

$$\nabla \cdot (\mathcal{F} + \mathcal{K}) \approx (\mathcal{F}_p + \alpha\mathcal{K}_p)\nabla \cdot, \quad (3.27)$$

where the scalar parameter α can be regarded as a relaxation parameter; that is, when (3.25) does not hold, α can be taken to be small to “relax” the error associated with this approximation. Discretizing this relation, we obtain

$$Q_{\mathbf{p}}^{-1}BQ_{\mathbf{u}}^{-1}X \approx Q_{\mathbf{p}}^{-1}(F_{\mathbf{p}} + \alpha K_{\mathbf{p}})Q_{\mathbf{p}}^{-1}B, \quad (3.28)$$

where $K_{\mathbf{p}}$ is the discrete analogue to \mathcal{K}_p . Note that if $F_{\mathbf{p}} + \alpha K_{\mathbf{p}}$ is dominated by $F_{\mathbf{p}}$ then this operator commutes with the divergence operator with little error. Assuming the norms of $F_{\mathbf{p}}$ and $K_{\mathbf{p}}$ to be approximately the same as those of F and K , this condition holds when

$$\alpha H^2 h^2 \|\vec{b}\|^2 \ll 1 + Rh\|\vec{a}\|, \quad (3.29)$$

where $H = \sqrt{SR_m}$ is the Hartmann number. From this, we can see that α must be small when $Hh\|\vec{b}\|$ is relatively large and α can be large if $Rh\|\vec{a}\|$ is relatively

large. Furthermore, for any problem parameters, α can be large if h is made small enough. Hence, adequate mesh refinement can allow us to take $\alpha = 1$, resulting in the approximation (3.26).

Through straightforward algebraic manipulation of (3.28), we obtain

$$Y = -BX^{-1}B^t \approx -Q_{\mathbf{p}}(F_{\mathbf{p}} + \alpha K_{\mathbf{p}})^{-1}BQ_{\mathbf{u}}^{-1}B^t. \quad (3.30)$$

We define $F_{\mathbf{p}}$ using a strategy introduced to develop the Least-Squares Commutator (LSC) Schur approximation to the Schur complement for the Navier-Stokes equations [25]. That is, $F_{\mathbf{p}}Q_{\mathbf{p}}^{-1}$ is computed one row at a time, where each row is obtained by solving the least-squares problem

$$\min ||[BQ_{\mathbf{u}}^{-1}F]_{j*} - [F_{\mathbf{p}}Q_{\mathbf{p}}^{-1}]_{j*}B||_{Q_{\mathbf{u}}^{-1}}, \quad (3.31)$$

so that (3.23) is enforced in a least-squares sense row by row. This results in the definition

$$F_{\mathbf{p}} := BQ_{\mathbf{u}}^{-1}FQ_{\mathbf{u}}^{-1}B^t(BQ_{\mathbf{u}}^{-1}B^t)^{-1}Q_{\mathbf{p}}. \quad (3.32)$$

Applying the same strategy for $K_{\mathbf{p}}$ to enforce the discrete version of (3.25), we obtain

$$K_{\mathbf{p}} := BQ_{\mathbf{u}}^{-1}KQ_{\mathbf{u}}^{-1}B^t(BQ_{\mathbf{u}}^{-1}B^t)^{-1}Q_{\mathbf{p}}. \quad (3.33)$$

Substituting these definitions into (3.30) gives the approximation

$$Y \approx \hat{Y}_{\alpha} := -BQ_{\mathbf{u}}^{-1}B^t[BQ_{\mathbf{u}}^{-1}(F + \alpha K)Q_{\mathbf{u}}^{-1}B^t]^{-1}BQ_{\mathbf{u}}^{-1}B^t. \quad (3.34)$$

In practice, $Q_{\mathbf{u}}$ can be replaced by $\text{diag}(Q_{\mathbf{u}})$ for ease of computation [62]. Then the action of \hat{Y}^{-1} requires two solves with $B(\text{diag}(Q_{\mathbf{u}}))^{-1}B^t$ which is a sparse (scaled)

Laplacian and is easy to handle with multigrid. Observe that when $\alpha = 1$, this approximation corresponds to applying the LSC methodology directly to the operator $X = F + K$. Thus, the choice of $\alpha = 1$ is an intuitive choice but does not correct for the commutation error on K . We will return to the preconditioner resulting from the choice $\alpha = 1$ in Section 3.3.

In an effort to develop intuition for a good choice of α , we consider the impact of the approximation when it is applied to a single component of a Fourier series, i.e. where p is of the form

$$p = e^{i\vec{\omega}\cdot\vec{x}}, \quad (3.35)$$

corresponding to the vector frequency $\vec{\omega}$. We would like to choose α so that the effect of the exact Schur complement Y on p is comparable to the effect of the Schur complement approximation \hat{Y}_α . Translating this condition to the continuous setting, we require

$$\nabla \cdot (\mathcal{F} + \mathcal{K})^{-1} \nabla p \approx \Delta_p [\nabla \cdot (\mathcal{F} + \alpha \mathcal{K}) \nabla]^{-1} \Delta_p p. \quad (3.36)$$

This gives

$$\begin{aligned} \alpha \approx & \|\vec{\omega}\|^4 \frac{\|\vec{\omega}\|^4 + H^2(\vec{b} \cdot \vec{\omega})^2 + R^2(\vec{a} \cdot \vec{\omega})^2}{\left[\|\vec{\omega}\|^4 + H^2(\vec{b} \cdot \vec{\omega})^2 \right]^2 + R^2(\vec{a} \cdot \vec{\omega})^2 \|\vec{\omega}\|^4} \\ & + i \frac{RH^2(\vec{a} \cdot \vec{\omega})(\vec{b} \cdot \vec{\omega})^2 \|\vec{\omega}\|^2}{\left[\|\vec{\omega}\|^4 + H^2(\vec{b} \cdot \vec{\omega})^2 \right]^2 + R^2(\vec{a} \cdot \vec{\omega})^2 \|\vec{\omega}\|^4}. \end{aligned} \quad (3.37)$$

If we restrict α to be a real constant, we should choose $\alpha \approx \Re(\alpha(\vec{\omega}))$, i.e

$$\alpha = \|\vec{\omega}\|^4 \frac{\|\vec{\omega}\|^4 + H^2(\vec{b} \cdot \vec{\omega})^2 + R^2(\vec{a} \cdot \vec{\omega})^2}{\left[\|\vec{\omega}\|^4 + H^2(\vec{b} \cdot \vec{\omega})^2 \right]^2 + R^2(\vec{a} \cdot \vec{\omega})^2 \|\vec{\omega}\|^4} \quad (3.38)$$

for a particular $\vec{\omega}$.

In general, our experiments have shown that the value of α is insensitive to the direction chosen for the Fourier mode $\vec{\omega}$ as long as the magnitude of $\vec{\omega}$ is chosen properly. We take $\vec{\omega}$ to be in the direction of \vec{a} , as \vec{a} can be considered an approximation to \vec{u} , the direction of fluid flow. With this choice, the expression (3.38) reduces to

$$\alpha \approx \frac{||\vec{\omega}||^4 + H^2||\vec{b}||^2||\vec{\omega}||^2 \cos^2(\theta) + R^2||\vec{a}||^2||\vec{\omega}||^2}{[||\vec{\omega}||^2 + H^2||\vec{b}||^2 \cos^2(\theta)]^2 + R^2||\vec{a}||^2||\vec{\omega}||^2} \quad (3.39)$$

where θ is the angle between \vec{a} and \vec{b} , and this approximation is entirely determined by the problem parameters and $||\vec{\omega}||$. The discretization of the problem should set a bound on the magnitude of any Fourier mode resolved by the mesh. That is, the most oscillatory Fourier mode should oscillate with period proportional to the mesh size on the pressure space h_p . Thus, we propose setting $||\vec{\omega}|| = \frac{1}{h_p}$, yielding

$$\alpha \approx \alpha_* := \frac{1 + H^2 h_p^2 ||\vec{b}||^2 \cos^2(\theta) + R^2 h_p^2 ||\vec{a}||^2}{[1 + H^2 h_p^2 ||\vec{b}||^2 \cos^2(\theta)]^2 + R^2 h_p^2 ||\vec{a}||^2}. \quad (3.40)$$

The value of α_* can be computed from the mesh Hartmann number Hh and the mesh Reynolds number Rh , both of which are readily available, and the quantities $||\vec{a}||$, $||\vec{b}||$, and $\cos(\theta)$. In this study, we use the average values of the latter quantities over the domain from the previous iterate in the nonlinear iteration. Defined this way, α_* follows the trends implied by (3.29); that is, α_* is small if $H^2 h^2 ||\vec{b}||^2$ dominates $R^2 h^2 ||\vec{a}||^2$ and α_* is close to 1 if $R^2 h^2 ||\vec{a}||^2$ dominates $H^2 h^2 ||\vec{b}||^2$. Furthermore, as h is refined, α_* approaches 1.

With approximations to both X and Y , we can now write the form of our

preconditioner for the linear systems arising from a Picard iteration

$$\mathcal{P}_{P,\alpha} = \begin{pmatrix} \hat{A} & -Z^t & 0 \\ 0 & \hat{X} & B^t \\ 0 & 0 & \hat{Y}_\alpha \end{pmatrix}, \quad (3.41)$$

where \hat{X} is an approximation to X as defined in (3.20) and \hat{Y}_α is an approximation to Y as defined in (3.34).

3.2 A Block Preconditioner for the Newton System

Now consider the block LU factorization for the (reordered) matrix arising from Newton's method,

$$\mathcal{A}_N = \begin{pmatrix} I & 0 & 0 \\ (Z + \tilde{Z})(A + \tilde{A})^{-1} & I & 0 \\ 0 & B\tilde{X}^{-1} & I \end{pmatrix} \begin{pmatrix} A + \tilde{A} & -Z^t & 0 \\ 0 & \tilde{X} & B^t \\ 0 & 0 & \tilde{Y} \end{pmatrix}, \quad (3.42)$$

where

$$\tilde{X} = F + \tilde{F} + (Z + \tilde{Z})(A + \tilde{A})^{-1}Z^t, \quad (3.43)$$

$$\tilde{Y} = -B\tilde{X}^{-1}B^t. \quad (3.44)$$

Although \tilde{X} does not simplify in its continuous form to an operator that can be explicitly constructed, because the Picard matrix and its associated block LU decomposition can be regarded as an approximation of the Newton matrix, we regard the $\mathbf{u}\text{-}\mathbf{B}$ coupling term $(Z + \tilde{Z})(A + \tilde{A})^{-1}Z^t$ as a modification of the analogous term for the Picard iteration $ZA^{-1}Z^t \approx K$. To illustrate this, consider the identity

$$(Z + \tilde{Z})(A + \tilde{A})^{-1}Z^t = ZA^{-1}Z^t + E, \quad (3.45)$$

where E is the perturbation defined to be

$$E := (\tilde{Z} - ZA^{-1}\tilde{A})(A + \tilde{A})^{-1}Z^t. \quad (3.46)$$

Considering the approximate norms of the components of E , we find that

$$\|E\| = O\left(SR_m h \|\vec{b}\| \frac{\|\nabla \times \vec{b}\| + R_m \|\vec{a}\| \|\vec{b}\|}{1 + R_m h \|\vec{a}\|}\right). \quad (3.47)$$

Because the magnitude of E is proportional to h , E can be effectively neglected if the mesh is refined enough. We develop an approximation to E for the cases when the mesh is not refined enough to neglect it. This approximation is important when R_m and $R_m h$ are both large.

Translating each discrete operator in (3.46) and simplifying using the relation (3.15) produces a discrete analogue \mathcal{E} of E defined as

$$\mathcal{E}\vec{u} := S \left[\cdot \times (\nabla \times \vec{b}) - R_m \vec{b} \times (\vec{a} \times \cdot) \right] \left[\frac{1}{R_m} \Delta + \nabla \times (\vec{a} \times \cdot) \right]^{-1} \nabla \times (\vec{u} \times \vec{b}). \quad (3.48)$$

Assuming that the vectors \vec{a} and \vec{b} are smooth allows us to approximate (3.48) by

$$\mathcal{E}\vec{u} \approx -SR_m \vec{b} \times (\vec{a} \cdot \nabla)_p \left(-\frac{1}{R_m} \Delta + \vec{a} \cdot \nabla \right)_p^{-1} (\vec{u} \times \vec{b}), \quad (3.49)$$

where we use the subscript p here to denote that the operator functions on the scalar pressure space. To use (3.49) to construct an approximation to E we make the approximation

$$(\vec{a} \cdot \nabla)_p \approx \left(-\frac{1}{R_m} \Delta + \vec{a} \cdot \nabla \right)_p \quad (3.50)$$

when R_m is large. This is the important case to consider because the $\mathbf{u}\text{-}\mathbf{B}$ coupling is strongest when R_m is large, hence necessitating a good approximation to E . Based

on this observation, we make the approximation

$$(\vec{a} \cdot \nabla)_p \approx (1 - \gamma) \left(-\frac{1}{R_m} \Delta + \vec{a} \cdot \nabla \right)_p, \quad (3.51)$$

where the new parameter γ is included to correct for any error resulting from assumption (3.50) and is assumed to be close to 0 when R_m is large. Given (3.51), we obtain the approximation

$$\mathcal{E} \approx (\gamma - 1) S R_m \vec{b} \times (\vec{u} \times \vec{b}) = (\gamma - 1) \mathcal{K}, \quad (3.52)$$

and E can be approximated by

$$E \approx (1 - \gamma) K. \quad (3.53)$$

Notice that when E is negligible, we should set $\gamma = 1$, but another value of γ may be needed if R_m is large. Using the approximation (3.53) and recalling that $-ZA^{-1}Z^t \approx K$, we have the approximation

$$\tilde{X} \approx \hat{X}_\gamma := F + \tilde{F} + \gamma K. \quad (3.54)$$

As we did for the Picard iteration, we can use $p = e^{i\vec{\omega} \cdot x}$ to gain some intuition about the choice of γ . Evaluating (3.51) with this choice of p , we find that

$$\gamma(\vec{\omega}) \approx \frac{|\vec{\omega}|^4}{|\vec{\omega}|^4 + R_m^2 (\vec{a} \cdot \vec{\omega})^2} - i \frac{R_m (\vec{a} \cdot \vec{\omega})}{|\vec{\omega}|^4 + R_m^2 (\vec{a} \cdot \vec{\omega})^2}. \quad (3.55)$$

Restricting γ to be real gives

$$\gamma \approx \frac{|\vec{\omega}|^4}{|\vec{\omega}|^4 + R_m^2 (\vec{a} \cdot \vec{\omega})^2} \quad (3.56)$$

for some $\vec{\omega}$.

We can learn more about a good choice of $\vec{\omega}$ by considering a discrete version of (3.50),

$$N_{\mathbf{p}} \approx (1 - \gamma)(A_{\mathbf{p}} + N_{\mathbf{p}}), \quad (3.57)$$

where $N_{\mathbf{p}}$ is a discretization of $(\vec{a} \cdot \nabla)_p$ and $A_{\mathbf{p}}$ is a discretization of $-\frac{1}{R_m} \Delta_p$. This can be rewritten as

$$\gamma I \approx A_{\mathbf{p}}(A_{\mathbf{p}} + N_{\mathbf{p}})^{-1}. \quad (3.58)$$

From this we can obtain the order of magnitude of γ to be

$$\gamma = O\left(\frac{1}{1 + R_m h_p \|\vec{a}\|}\right). \quad (3.59)$$

If we let

$$\vec{\omega} = \vec{a} \sqrt{\frac{R_m}{h_p \|\vec{a}\|}}, \quad (3.60)$$

then (3.56) satisfies (3.59). This choice of Fourier mode is reasonable from a physical perspective. It is chosen in the direction of flow \vec{a} and for a given h_p , its magnitude is proportional to $\sqrt{R_m \|\vec{a}\|}$, the width of characteristic layers appearing for the convection-diffusion operator $(\vec{a} \cdot \nabla - \frac{1}{R_m} \Delta)_p$. Furthermore, as the mesh is refined, higher frequency modes can be resolved, and $\|\vec{\omega}\|$ increases to reflect this. Hence, we use the approximation

$$\gamma \approx \gamma_* := \frac{1}{1 + R_m h_p \|\vec{a}\|}, \quad (3.61)$$

which depends only on the mesh magnetic Reynolds number $R_m h_p$ and the quantity $\|\vec{a}\|$ which can be taken as the average over the domain.

Using \hat{X}_γ as defined in (3.54) to approximate \tilde{X} and \hat{Y}_α as defined in (3.34)

to approximate \tilde{Y} , we have the preconditioner for the Newton system

$$\mathcal{P}_{N,\alpha,\gamma} := \begin{pmatrix} \widehat{A + \tilde{A}} & -Z^t & 0 \\ 0 & \hat{X}_\gamma & B^t \\ 0 & 0 & \hat{Y}_\alpha \end{pmatrix}. \quad (3.62)$$

We note that the choice of α should incorporate γ . It is easy to show that, according to the same arguments from Section 3.1.2, α should be defined as

$$\alpha \approx \alpha_*(\gamma) := \frac{1 + \gamma H^2 h_p^2 \|\vec{b}\|^2 \cos^2(\theta) + R^2 h_p^2 \|\vec{a}\|^2}{[1 + \gamma H^2 h_p^2 \|\vec{b}\|^2 \cos^2(\theta)]^2 + R^2 h_p^2 \|\vec{a}\|^2}. \quad (3.63)$$

3.3 Computational Results

In this section, we evaluate the performance of the preconditioners for solving a set of benchmark problems. Our implementation is in the Trilinos framework developed at Sandia National Laboratories [38] using the Teko package to construct the block preconditioners and a GMRES Krylov solver from AztecOO [37]. For component solves, we use algebraic multigrid from the ML package [27], with incomplete factorization smoothers coming from IFPACK. Specifically, for solves on the velocity and magnetic spaces we use one V-cycle of AMG with two pre- and post-sweeps of a Schwarz domain decomposition smoother with ILU(0) on each subdomain and one level of overlap between subdomains. For solves on the pressure space we use one V-cycle of AMG with five sweeps of a Gauss-Seidel smoother. This AMG technology has been demonstrated to be algorithmically scalable for both an equal order stabilized finite element formulation of the full MHD system and as a component solve in physics-based preconditioners [17, 60]. All problems were run on the Red Sky

computer at Sandia National Laboratories. For all problems we use a stable Q_2 - Q_1 finite element pair for the velocity and pressure and Q_2 elements for the magnetic field to match the velocity. We use uniform grids for every problem. The relative residual tolerance is 10^{-5} for the nonlinear iteration and 10^{-6} for the GMRES inner iteration. We apply GMRES without any restarting. When considering the performance of the preconditioners, reported linear iteration counts are averaged over all nonlinear iterations up to a maximum of twenty nonlinear iterations. We consider two two-dimensional test problems, a lid driven cavity in the presence of a magnetic field (adapted from the fluid problem in [24]) and the Hartmann flow problem (as detailed in [29]). To explore the robustness of our preconditioning strategies, we test their performance over a range of R and R_m on these problems. We set $S = 1$ and let the variation in R_m account for different degrees of fluid-magnetic coupling.

3.3.1 MHD Lid Driven Cavity

The lid driven cavity problem models the flow of a conducting fluid in the square cavity $[0, 1] \times [0, 1]$, driven by the motion of its lid from left to right with the magnetic field $(-1, 0)$ imposed parallel to the lid. No-flow conditions are imposed on the bottom, left, and right sides of the cavity, and the horizontal velocity of the lid is prescribed to be 1. The tangential component of the magnetic field is specified on the boundary $\partial\Omega$ as

$$\vec{B} \times \vec{n} = (-1, 0) \times \vec{n}. \quad (3.64)$$

For $R < R_c \approx 7800$, the cavity flow problem with no magnetic field has

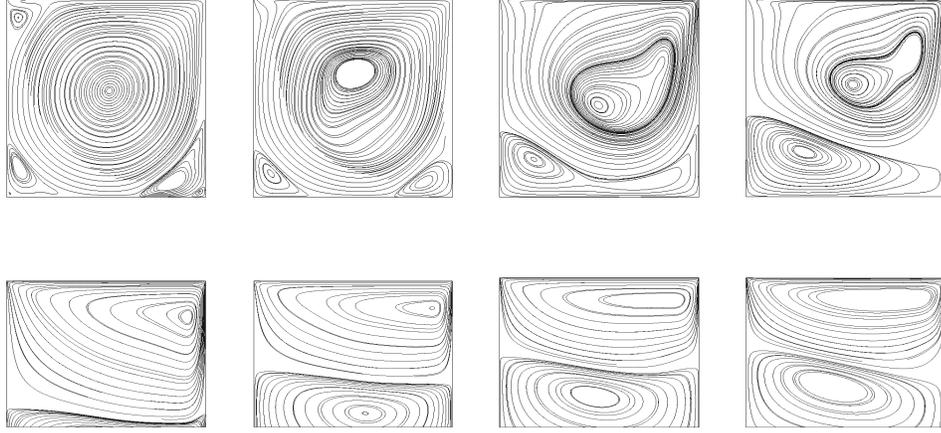


Figure 3.1: Streamlines for the MHD lid driven cavity problem with $R = 5000$ and $R_m = 0, 0.1, 0.3, 0.4, 5, 10, 20, 30$. The four latter cases are zoomed in to $[0, 1] \times [0.8, 1]$.

a stable solution dominated by one large eddy in the center of the cavity with smaller secondary eddies in the corners [61]. Imposing the magnetic field leads to weakening of the flow and, for stronger fields, merging of the secondary eddies. As the magnetic field is strengthened (that is, as R_m increases), the height of the primary eddy decreases as an effect of the increasing magnitude of the Lorentz force. Furthermore, other horizontal eddies of decreasing height develop, stretching from the left wall to the right wall and stacking on top of each other, with the number of horizontal eddies increasing with R_m . As a result of the braking effect of the Lorentz force, the flow in the lower part of the domain is almost stagnant for large R_m . Using solutions obtained from our code, this behavior is shown in Figure 3.1. Streamlines are plotted for the case where $R = 5000$ with increasing R_m . This set of problems is equivalent to those considered in [61], and our results are qualitatively very similar, with the same number and height of eddies appearing.

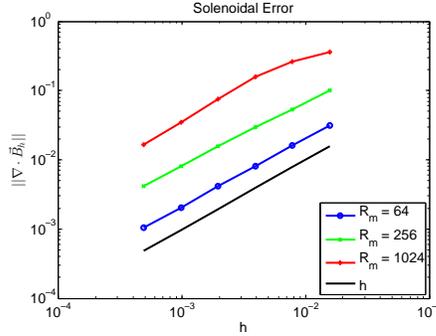


Figure 3.2: Magnitude of the divergence of \vec{B}_h versus h for $R = 64$ and $R_m = 64, 256, 1024$.

To demonstrate that the exact penalty formulation weakly enforces the solenoidal condition, we have plotted the magnitude of the divergence of the computed magnetic field in Figure 3.2. Because we have found that $\|\nabla \cdot \vec{B}_h\|_{L^2(\Omega)}$ is independent of the fluid Reynolds number, we only show results for $R = 64$. It can be seen from the figure that this quantity is approximately proportional to $R_m h$. This is consistent with results obtained in the context of Marder cleaning for Maxwell's equations [46]. In this context, the penalty term $\frac{S}{R_m}(\nabla \cdot \vec{B}, \nabla \cdot \vec{C})$ is analogous to a pseudo-current term added to Maxwell's equations scaled by $\frac{S}{R_m}$. For time-dependent problems, the scaling of the pseudo-current term determines the rate at which the error in the solenoidal condition vanishes.

Table 3.2 shows the number of steps required by each nonlinear scheme on a uniform 512×512 element mesh, starting from a zero initial guess. As R_m increases, the fluid-magnetic coupling strengthens and the nonlinear problems become somewhat more difficult to solve. For entries in the table with asterisks, convergence was achieved using a backtracking strategy; see [52]. Newton's method was slightly

		Picard				Newton			
R	R_m	1	64	256	1024	1	64	256	1024
	1		2	×	×	×	2	3	3
64		2	7	×	×	2	4	3	3
256		2	4	7*	×	2	4	5	3
1024		2	3	4	11*	2	3	6	16*
4096		2	3	4	5	2	2	7	7*

Table 3.2: Iterations required for convergence of the nonlinear iteration. ‘×’ indicates no convergence within twenty iterations. ‘*’ indicates convergence required backtracking.

		$\alpha = 1$				$\alpha = \alpha_*$			
R	R_m	1	64	256	1024	1	64	256	1024
	1		36	45	44	78	36	45	44
64		35	42	47	84	35	42	47	77
256		36	42	60	109	36	42	57	90
1024		44	44	89	193	44	43	77	142
4096		68	60	123	291	68	58	91	174

Table 3.3: Average GMRES iterations required for convergence with $\mathcal{P}_{P,\alpha}$ on the Picard linearization of the MHD lid driven cavity problem with $\alpha = 1$ and $\alpha = \alpha_*$.

more robust than Picard iteration. Although backtracking alone was not sufficient to make the Picard iteration fully robust for these examples, since our emphasis is on the linear solvers, we did not pursue this further.¹

3.3.1.1 Picard Iteration

In this section, we consider the performance of the preconditioner $\mathcal{P}_{P,\alpha}$ on the MHD lid driven cavity problem for a range of parameters, using a 512×512 element

¹This type of behavior is often sensitive to the initial guess and also to the properties of the globalization methods used; see [52] for other ways to enhance robustness. We also note that poor performance of the nonlinear solver is related to under-resolution of the mesh. For the problematic cases in Table 3.2, both the Picard iteration and Newton’s method converge on finer meshes without use of globalization methods.

$R \backslash R_m$	1	64	256	1024
1	1.000	1.000	0.998	0.992
64	1.000	0.970	0.889	0.667
256	0.998	0.895	0.677	0.339
1024	0.996	0.777	0.407	0.123
4096	0.998	0.777	0.272	0.046

Table 3.4: Computed values of α_* for the second Picard iteration.

mesh. To study the impact of the parameter α on the preconditioner, we compare the performance of $\mathcal{P}_{P,\alpha}$ with $\alpha = \alpha_*$ as defined in (3.40) and $\alpha = 1$, corresponding to the LSC preconditioner applied directly to X . GMRES iteration counts with these two preconditioners for various choices of R and R_m are reported in Table 3.3. The values of α_* used at the second nonlinear step are reported in Table 3.4. (Note that for the first step, $\vec{a} = \vec{0}$, and therefore $\alpha_* = 1$ independent of R and R_m .) From Table 3.3, it is clear that the automatically computed $\alpha = \alpha_*$ leads to improved performance compared to LSC preconditioning. The differences are minimal for small R and R_m (the iteration counts are identical for $R = 1$ and $R_m = 1$), but they become substantial for larger R and R_m . For example, for $R = 4096$, $R_m = 1024$, the counts for $\alpha = 1$ and $\alpha = \alpha_*$ are 291 and 174, respectively, a 40% reduction. Table 3.4 also shows that α_* is close to 1 for small R and R_m . These results are consistent with the observation made in Section 3.1 that when the fluid-magnetic

	$\alpha = 1, \gamma = 1$				$\alpha = \alpha_*, \gamma = 1$				$\alpha = \alpha_*, \gamma = \gamma_*$			
$R \backslash R_m$	1	64	256	1024	1	64	256	1024	1	64	256	1024
1	36	42	47	68	36	42	47	68	36	42	47	66
64	34	46	69	171	34	47	69	163	34	47	66	117
256	37	50	99	299	37	50	94	228	37	49	88	164
1024	44	50	169	584	44	50	133	288	44	50	121	249
4096	68	56	208	603	68	54	142	419	68	54	137	364

Table 3.5: Average GMRES iterations required for convergence with $\mathcal{P}_{P,\alpha,\gamma}$ on the Newton linearization of the MHD lid driven cavity problem.

coupling is strong, the validity of the approximation for the commutator (3.25) is weaker; the “relaxed” variants (3.27)–(3.28) improve performance. We also found that an optimal choice of α determined by a brute-force search produced iteration counts essentially the same as for $\alpha = \alpha_*$. Thus, the choice of parameter given by (3.40) determines an essentially automated version of the preconditioner (3.28).

3.3.1.2 Newton’s Method

Here, we consider the performance of the preconditioner $\mathcal{P}_{N,\alpha,\gamma}$ on the linear systems arising from Newton’s method applied to the MHD lid driven cavity problem. We consider a 512×512 mesh with R between 1 and 4096 and R_m between 1 and 1024. We compare three preconditioners: the fully parameterized precondi-

$R \backslash R_m$	1	64	256	1024
1	0.999	0.939	0.794	0.491
64	0.999	0.940	0.796	0.493
256	0.999	0.940	0.796	0.495
1024	0.999	0.940	0.796	0.493
4096	0.999	0.940	0.796	0.493

Table 3.6: Computed values of γ_* at the second Newton step.

tioner $\mathcal{P}_{N,\alpha_*,\gamma_*}$ with γ_* from (3.61) and α_* from (3.63); the preconditioner $\mathcal{P}_{N,\alpha_*,1}$ parameterized only by α with α_* from (3.40); and the unparameterized preconditioner $\mathcal{P}_{N,1,1}$. $\mathcal{P}_{N,\alpha_*,1}$ corresponds to applying the strategy derived for the Picard iteration to the Jacobian system, as though the fluid-magnetic coupling E is negligible. GMRES iteration counts are reported in Table 3.5 and the values of γ_* are reported in Table 3.6.

These results show that the (doubly) parameterized preconditioner for the Jacobian systems nearly always exhibits enhanced performance, with the most significant improvements occurring in regimes where the spatial resolution is weakest (large R or R_m). The impact of the new parameter is comparable to that of α , and we emphasize that, as above, the enhanced performance is obtained using an automated strategy for choosing parameters.

3.3.1.3 Robustness with Respect to Mesh Refinement

In Sections 3.3.1.1 and 3.3.1.2, we explored the robustness of our preconditioning strategies with respect to the physical parameters R and R_m , on a fixed mesh. We are also interested in their performance as the mesh is refined. We investigate this on the two problems corresponding to $R = 256, R_m = 256$ and $R = 256, R_m = 1024$. Because the multigrid component solves depend on the parallel architecture used, we also use this as an opportunity to study the parallel scalability of the preconditioners. That is, the Schwarz-ILU smoother for the multigrid solves decomposes the domain into as many subdomains as there are processors. Hence, we want to make sure that increasing the number of processors does not cause the performance of the full preconditioners to degrade. Toward both of these ends, we increase the number of processors as we refine the mesh so that the number of unknowns per processor remains approximately constant. That is, we perform a weak parallel scaling study by considering 64×64 , 128×128 , 256×256 , 512×512 , 1024×1024 , and 2048×2048 element discretizations on 1, 4, 16, 64, 256, and 1024 processors respectively, keeping the number of unknowns per processor at approximately 70,000.

We report average iteration counts and computation times per nonlinear step for these experiments in Figure 3.3 for the Picard iteration and Figure 3.4 for Newton's method. Newton's method converges on all grids considered for $R = 256, R_m = 256$ but only on the three most refined grids for $R = 256, R_m = 1024$. For the Picard iteration, we compare the preconditioner \mathcal{P}_{P,α_*} with α_* from (3.40) with the unparameterized preconditioner $\mathcal{P}_{P,1}$. As a frame of reference for parallel scalability, we

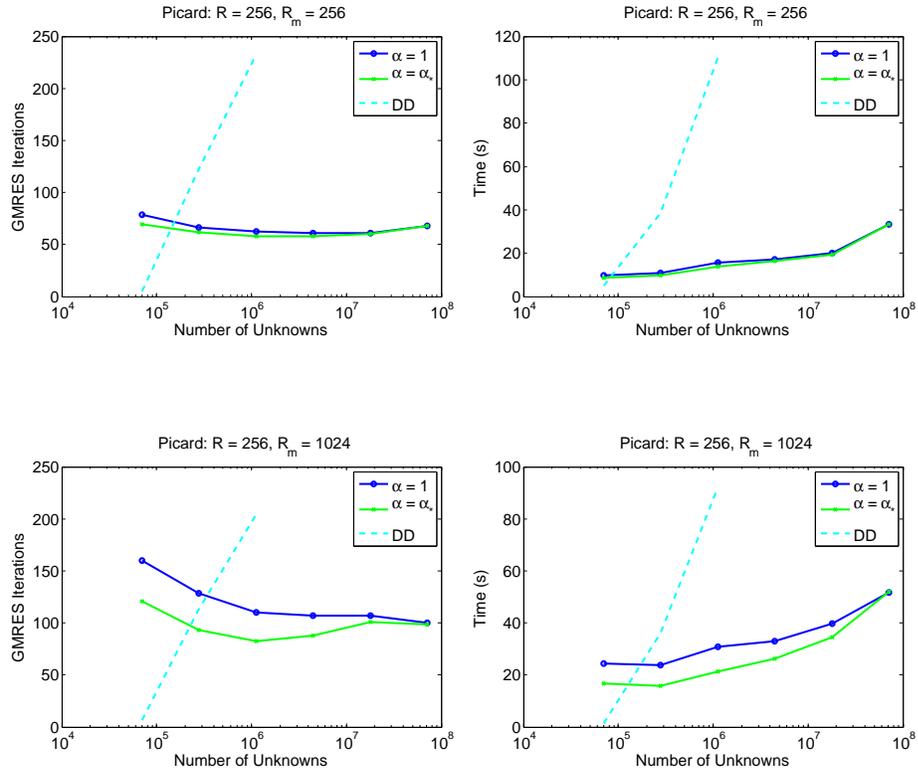


Figure 3.3: Mesh refinement results for the MHD lid driven cavity problem for the Picard iteration. Parameters are $R = 256$, $R_m = 256$ on the top and $R = 256$, $R_m = 1024$ on the bottom. Average GMRES iterations on the left and average linear solve time per nonlinear iteration on the right.

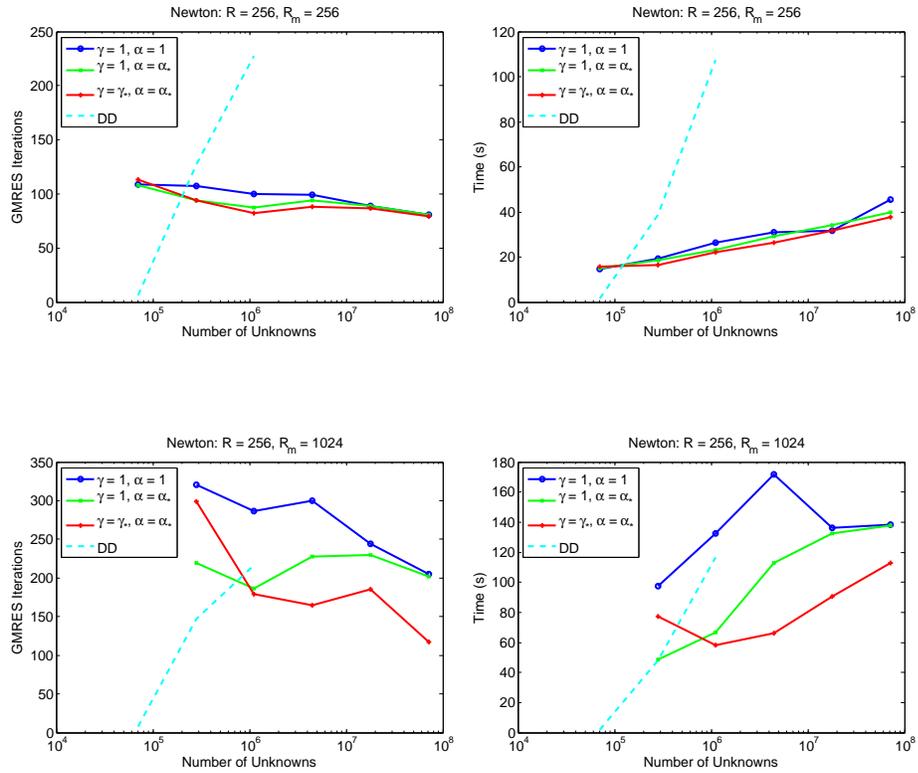


Figure 3.4: Mesh refinement results for the MHD lid driven cavity problem for Newton's method. Parameters are $R = 256, R_m = 256$ on the top and $R = 256, R_m = 1024$ on the bottom. Average GMRES iterations on the left and average linear solve time per nonlinear iteration on the right.

also compare these preconditioners to a pure domain decomposition preconditioner using a SuperLU [21] ILUTP factorization with a drop tolerance of 0.0001 and a zero pivot threshold of 0.01 (labeled DD in the scaling plots). Thresholding and partial pivoting are necessary here because of the zero block on the diagonal for the pressure space. From the plots, it can be seen that both \mathcal{P}_{P,α_*} and $\mathcal{P}_{P,1}$ are robust with respect to mesh refinement, with iteration counts that are nearly constant or decreasing as the number of unknowns grows. The parameterized preconditioner also performs uniformly as well or better than the unparameterized preconditioner, the benefit being more pronounced for problems with fewer unknowns or larger R_m . The iteration counts for the two preconditioners become similar as h is refined because α_* approaches 1 as h decreases. The domain decomposition preconditioner is not competitive and does not scale well. In fact, for all the tests run on more than 16 processors, GMRES runs out of memory before convergence with this preconditioner. While iteration counts tend to decrease, we see a slight increase in computation time as the mesh is refined. This appears to be due to increased communication costs in the component AMG solves as more processors are added, and we expect that performance can be enhanced by fine-tuning these solves.

For Newton’s method, we compare $\mathcal{P}_{N,\alpha_*,\gamma_*}$, the fully parameterized preconditioner, with $\mathcal{P}_{N,\alpha_*,1}$, the preconditioner parameterized only for the matrix arising from Picard iteration and with $\mathcal{P}_{N,1,1}$, the unparameterized preconditioner. The plots for $R = 256, R_m = 1024$ contain only three data points because Newton’s method diverges for this problem on the three coarsest meshes. We use γ_* from defined in (3.61). Iteration counts generally decrease for all three preconditioners as

h is refined. Since both α_* and γ_* converge to 1 as h decreases, the performances of the three preconditioners become essentially the same as the problem size increases. In all but the least refined cases, the fully parameterized preconditioner outperforms the others, especially for problems with large R_m . The results then show the importance of the parameter γ in keeping iteration counts low.

3.3.2 Hartmann Flow

The Hartmann flow problem describes the flow of a conducting fluid through a channel in the presence of a transverse magnetic field. We consider the channel $[-1/2, 1/2]^2$ and the transverse field $(0, 1)$. With appropriate boundary conditions, this problem has the explicit analytic solution $\vec{u} = (u_x, 0)$, $\vec{B} = (B_x, 1)$ with

$$u_x = \frac{GR(\cosh(H/2) - \cosh(Hy))}{2H \sinh(H/2)}, \quad (3.65a)$$

$$B_x = \frac{G(\sinh(Hy) - 2 \sinh(H/2)y)}{2S \sinh(H/2)}, \quad (3.65b)$$

$$p = -Gx - SB_x^2/2, \quad (3.65c)$$

where G can be any scalar. Given R and R_m , we choose G so the maximum magnitude of \vec{u} is normalized to 1. In our implementation, we prescribe the analytic value of \vec{u} and $\vec{B} \times \vec{n} = (0, 1) \times \vec{n}$ on the boundary.

Because an analytic solution exists for this problem, we can compute the errors $\|\vec{u}_{exact} - \vec{u}_h\|_{L^2(\Omega)}$ and $\|\vec{B}_{exact} - \vec{B}_h\|_{L^2(\Omega)}$. These errors and the error in the solenoidal condition $\|\nabla \cdot \vec{B}_h\|_{L^2(\Omega)}$ are plotted in Figure 3.5. From this figure, it is clear that all three quantities are proportional to h^3 . As with the MHD lid driven cavity problem, $\|\nabla \cdot \vec{B}_h\|_{L^2(\Omega)}$ is also related to R_m , depending to a lesser degree on R . Consequently,

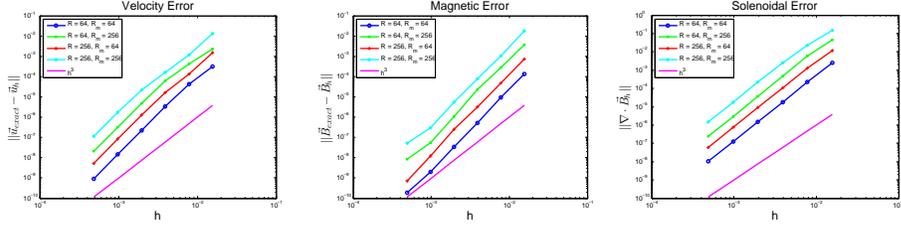


Figure 3.5: Magnitude of the error in \vec{u}_h and \vec{B}_h and the divergence of \vec{B}_h versus h .

the errors in \vec{u} and \vec{B} depend similarly on R_m and R .

We have tested the preconditioner $\mathcal{P}_{P,\alpha}$ on the linear systems resulting from the Picard linearization of the Hartmann flow problem on a fixed 512×512 element mesh. We found that nonlinear iterations did not converge for $R_m > 256$ on this mesh, so we do not consider $R_m = 1024$ for this problem. In Table 3.7, preconditioned GMRES iteration counts are compared for the choices $\alpha = 1$ and $\alpha = \alpha_*$ as defined in (3.40). In this case, the difference between the two preconditioners is marginal for all parameters.

We have also considered the performance of the preconditioner $\mathcal{P}_{N,\alpha,\gamma}$ on the linear systems resulting from Newton's method on a 512×512 element mesh. We compare $\mathcal{P}_{N,1,1}$, $\mathcal{P}_{N,\alpha_*,1}$, and $\mathcal{P}_{N,\alpha_*,\gamma_*}$ in Table 3.8. Here, we see the parameter γ having a much more profound effect on iteration counts than the parameter α . The choice of $\gamma = \gamma_*$ over $\gamma = 1$ leads to large improvements for moderate values of R_m . For instance, we can see a 42% improvement in iteration count for $R = 4099$, $R_m = 256$.

We note here that the performance of the preconditioners with respect to mesh refinement for the Hartmann flow problem follows the same trends as for the MHD

		$\alpha = 1$			$\alpha = \alpha_*$		
R	R_m	1	64	256	1	64	256
	1		38	41	40	38	41
64		31	37	40	31	37	40
256		29	35	50	29	35	49
1024		36	37	61	36	37	61
4096		32	79	81	32	79	81

Table 3.7: Average GMRES iterations required for convergence with $\mathcal{P}_{P,\alpha}$ on the Picard linearization of the Hartmann flow problem.

		$\alpha = 1, \gamma = 1$			$\alpha = \alpha_*, \gamma = 1$			$\alpha = \alpha_*, \gamma = \gamma_*$		
R	R_m	1	64	256	1	64	256	1	64	256
	1		39	43	137	39	43	137	39	43
64		31	50	144	31	50	144	31	48	99
256		29	49	124	29	49	123	29	46	81
1024		36	55	141	36	55	141	36	52	89
4096		32	158	206	32	158	205	32	113	118

Table 3.8: Average GMRES iterations required for convergence with $\mathcal{P}_{P,\alpha,\gamma}$ on the Newton linearization of the Hartmann flow problem.

lid driven cavity problem. For this reason, we do not include a figure analogous to Figures 3.3 and 3.4.

3.4 Conclusion

We have presented a family of block preconditioners for the linear systems arising from both Picard and Newton linearizations of the exact penalty formulation of the stationary MHD equations. For the Picard linearization, we chose an ordering $(\mathbf{B}, \mathbf{u}, \mathbf{p})$ that leads to a simple Schur complement on the velocity space that can be approximated by simplifying the continuous operator corresponding to it. This results in approximating the velocity Schur complement by a perturbed convection-diffusion operator. By modifying the Least-Squares Commutator preconditioner for the Navier-Stokes equations, we have developed an approximation to the nested Schur complement on the pressure space. The preconditioner resulting from direct application of the LSC methodology to the perturbed convection-diffusion operator can be improved by “relaxing” assumptions on the existence of small commutators. Using the parameterized assumption (3.51), similar strategies can be applied to Newton’s method.

We have presented an automated method for choosing the two parameters α and γ based on arguments from Fourier analysis. This method relies only on data which is readily available in a nonlinear iteration. Our numerical studies on the Picard linearization have demonstrated that the automated choice of α leads to significant improvement over $\alpha = 1$. For Newton’s method, we have shown that

both α and γ are important for decreasing iteration counts. The preconditioners are robust with respect to mesh refinement, their performances improving as the mesh is refined. In terms of parallel scalability, these preconditioners compare very well against a preconditioner based on domain decomposition.

Chapter 4: Block Preconditioners for a Lagrange Multiplier MHD Formulation

In this chapter, we consider the linear systems obtained when the MHD equations are discretized according to the Lagrange multiplier finite element formulation introduced in Section 2.3. As in Chapter 3, discretization of the linearized formulation (2.16) leads to a sequence of linear systems of the form

$$\mathcal{A}\mathbf{x} = \mathbf{f}. \tag{4.1}$$

Because the Lagrange multiplier formulation includes the additional unknown r , the vector \mathbf{x} is defined $\mathbf{x} = (\mathbf{u}, \mathbf{p}, \mathbf{B}, \mathbf{r})$. Again \mathbf{f} is the discrete nonlinear residual, and \mathcal{A} depends on the bilinear form associated with the linearization. The matrices resulting from the Picard and Newton linearizations can be written in block form as

$$\mathcal{A}_P = \begin{pmatrix} F & B^t & Z & 0 \\ B & 0 & 0 & 0 \\ -Z^t & 0 & A & D^t \\ 0 & 0 & D & 0 \end{pmatrix}, \quad \mathcal{A}_N = \begin{pmatrix} F + \tilde{F} & B^t & Z + \tilde{Z} & 0 \\ B & 0 & 0 & 0 \\ -Z^t & 0 & A + \tilde{A} & D^t \\ 0 & 0 & D & 0 \end{pmatrix}, \tag{4.2}$$

where the component matrices are defined in Table 4.1.

Discrete	Continuous	Interpretation	Approximate Norm
\mathbf{AB}	$\frac{S}{R_m} \nabla \times \nabla \times \vec{B}$	Magnetic diffusion	$\frac{S}{R_m h^2}$
$\tilde{\mathbf{A}}\mathbf{B}$	$-S \nabla \times (\vec{a} \times \vec{B})$	Magnetic convection	$\frac{S \ \vec{a}\ }{h}$
$-\mathbf{Z}^t \mathbf{u}$	$-S \nabla \times (\vec{u} \times \vec{b})$	Magnetic convection	$\frac{S \ \vec{b}\ }{h}$
$\mathbf{Z}\mathbf{B}$	$S \vec{b} \times (\nabla \times \vec{B})$	Lorentz force	$\frac{S \ \vec{b}\ }{h}$
$\tilde{\mathbf{Z}}\mathbf{B}$	$S \vec{B} \times (\nabla \times \vec{b})$	Lorentz force	$S \ \nabla \times \vec{b}\ $
$\mathbf{F}\mathbf{u}$	$\vec{a} \cdot \nabla \vec{u} - \frac{1}{R} \Delta \vec{u}$	Fluid convection-diffusion	$\frac{\ \vec{a}\ }{h} + \frac{1}{Rh^2}$
$\tilde{\mathbf{F}}\mathbf{u}$	$\vec{u} \cdot \nabla \vec{a}$	Fluid convection	$\ \nabla \vec{a}\ $
$\mathbf{B}^t \mathbf{p}$	∇p	Pressure gradient	$\frac{1}{h}$
$-\mathbf{B}\mathbf{u}$	$\nabla \cdot \vec{u}$	Divergence	$\frac{1}{h}$
$-\mathbf{D}^t \mathbf{r}$	∇r	Pseudo-pressure gradient	$\frac{1}{h}$
$\mathbf{D}\mathbf{B}$	$\nabla \cdot \vec{B}$	Divergence	$\frac{1}{h}$

Table 4.1: Definitions of discrete operators as they correspond to continuous operators.

These matrices are similar in structure to those obtained from the exact penalty formulation but feature significant differences. Whereas the exact penalty led to block 3×3 linear systems, the inclusion of the Lagrange multiplier \mathbf{r} leads to block 4×4 systems. These systems feature two subsystems of saddle point form, the fluid 2×2 block in the upper left, which corresponds to the linearized discrete Navier-Stokes equations, and the electromagnetic 2×2 block in the lower right, which corresponds to a mixed discretization of Maxwell's equations. Because a discrete version of the solenoidal condition $\nabla \cdot \vec{B} = 0$ is enforced explicitly in the constraint $D\mathbf{B} = \mathbf{0}$, the magnetic diffusion operator A does not include a penalty but instead reflects the operator $\frac{S}{R_m} \nabla \times \nabla \times$. In this case, A is a singular matrix whose nullspace is the span of the discrete gradient (i.e. $AQ_{\mathbf{B}}^{-1}D^t = 0$ and $DQ_{\mathbf{B}}^{-1}A = 0$) [20]. Similarly $A + \tilde{A}$ is singular as $DQ_{\mathbf{B}}^{-1}(A + \tilde{A}) = 0$.

Because \mathcal{A}_P and \mathcal{A}_N are structurally similar to the matrices arising from the exact penalty formulation of the MHD equations, we will focus on extending the ideas developed in Chapter 3 to develop preconditioners for \mathcal{A}_P and \mathcal{A}_N in this chapter. We focus first on the linear systems arising from a Picard linearization and then build on these results for the linear systems arising from Newton's method.

4.1 Block Preconditioners for the Picard System

As in Chapter 3, we are interested in developing block preconditioners for the matrices \mathcal{A}_P and \mathcal{A}_N . We focus first on developing a preconditioner \mathcal{P}_P for the Picard matrix \mathcal{A}_P . As with the exact penalty formulation, we have found it helpful

to reorder the unknowns with electromagnetic unknowns first, $(\mathbf{B}, \mathbf{r}, \mathbf{u}, \mathbf{p})$, yielding

$$\mathcal{A}_P = \begin{pmatrix} A & D^t & -Z^t & 0 \\ D & 0 & 0 & 0 \\ Z & 0 & F & B^t \\ 0 & 0 & B & 0 \end{pmatrix}. \quad (4.3)$$

Because A is a singular operator, \mathcal{A}_P does not admit a 4×4 block LU decomposition in this form. In order to perform a block LU decomposition, we consider the block 3×3 matrix obtained from blocking the electromagnetic unknowns together; that is, we define the 2×2 electromagnetic subsystem

$$\mathcal{M}_P = \begin{pmatrix} A & D^t \\ D & 0 \end{pmatrix}, \quad (4.4)$$

which is a discretization of Maxwell's equations in mixed form. It has been shown that when Maxwell's equations in this form are discretized with Nédélec elements for \vec{B} and nodal elements for r , the discretization is stable and well posed, and thus that \mathcal{M}_P is nonsingular [20]. Then, we can compute the block LU decomposition

$$\begin{aligned} \mathcal{A}_P &= \begin{pmatrix} \mathcal{M}_P & -Z^t & 0 \\ Z & F & B^t \\ 0 & B & 0 \end{pmatrix} \\ &= \begin{pmatrix} I & 0 & 0 \\ Z\mathcal{M}_P^{-1} & I & 0 \\ 0 & BX_P^{-1} & I \end{pmatrix} \begin{pmatrix} \mathcal{M}_P & -Z^t & 0 \\ 0 & X_P & B^t \\ 0 & 0 & Y_P \end{pmatrix}, \end{aligned} \quad (4.5)$$

where $\mathcal{Z} = \begin{pmatrix} Z & 0 \end{pmatrix}$ and the Schur complements X_P and Y_P are defined as

$$X_P := F + \mathcal{Z}\mathcal{M}_P^{-1}\mathcal{Z}^t, \quad (4.6)$$

$$Y_P := -BX_P^{-1}B^t. \quad (4.7)$$

Motivated by this decomposition, we investigate preconditioners of the form

$$\mathcal{P}_P = \begin{pmatrix} \hat{\mathcal{M}}_P & -\mathcal{Z}^t & 0 \\ 0 & \hat{X}_P & B^t \\ 0 & 0 & \hat{Y}_P \end{pmatrix}. \quad (4.8)$$

The challenge is then to develop effective expressions for $\hat{\mathcal{M}}_P$, \hat{X}_P , and \hat{Y}_P .

4.1.1 Preconditioners for \mathcal{M}_P

Several block preconditioners have been developed for the discretized Maxwell's equations in mixed form [34, 35, 42, 63, 64]. In this literature, the only linear systems considered correspond to the case where $\frac{S}{R_m} = 1$, but the results are easily generalized to other values of $\frac{S}{R_m}$. In this section, we discuss two block-diagonal preconditioners for the Maxwell system [35, 64]. Both preconditioners work by augmenting the singular (1,1) block to make it nonsingular and then approximating the Schur complement associated with the Lagrange multiplier \mathbf{r} . For each preconditioner, we present an analysis of the eigenvalues of the preconditioned system $\mathcal{M}_P\hat{\mathcal{M}}_P^{-1}$.

4.1.1.1 Mass Augmentation

For the first preconditioner we consider [35], the (1,1) block is augmented with a scaled mass matrix, $kQ_{\mathbf{B}}$, yielding $A + kQ_{\mathbf{B}}$. We have introduced the parameter $k > 0$, which does not appear in [35], to allow for flexibility in the augmentation. The Schur complement associated with the augmented system is $D(A + kQ_{\mathbf{B}})^{-1}D^t$. This operator can be simplified using the relation

$$DQ_{\mathbf{B}}^{-1}(A + kQ_{\mathbf{B}}) = kD, \quad (4.9)$$

which implies that

$$D(A + kQ_{\mathbf{B}})^{-1}D^t = \frac{1}{k}DQ_{\mathbf{B}}^{-1}D^t. \quad (4.10)$$

The operator $DQ_{\mathbf{B}}^{-1}D^t$ is spectrally equivalent to a discrete Laplacian on the space associated with \mathbf{r} , $L_{\mathbf{r}}$ [24]. Thus, the Schur complement can be approximated by $\frac{1}{k}L_{\mathbf{r}}$, and using this approximation, we define the block diagonal preconditioner

$$\hat{\mathcal{M}}_{P,Q} := \begin{pmatrix} A + kQ_{\mathbf{B}} & 0 \\ 0 & \frac{1}{k}L_{\mathbf{r}} \end{pmatrix}. \quad (4.11)$$

The eigenvalues of the preconditioned system $\mathcal{M}_P \hat{\mathcal{M}}_{P,Q}^{-1}$ have been analyzed for the case $\frac{S}{R_m} \equiv 1, k = 1$ in [35]. We extend the analysis from [35] to the case where $\frac{S}{R_m}$ and k are both arbitrary positive constants. In particular, we are interested in the eigenvalues λ of the generalized eigenvalue problem

$$\begin{pmatrix} A & D^t \\ D & 0 \end{pmatrix} \begin{pmatrix} \mathbf{B} \\ \mathbf{r} \end{pmatrix} = \lambda \begin{pmatrix} kQ_{\mathbf{B}} + A & 0 \\ 0 & \frac{1}{k}L_{\mathbf{r}} \end{pmatrix} \begin{pmatrix} \mathbf{B} \\ \mathbf{r} \end{pmatrix}. \quad (4.12)$$

Defining $n = \dim(\mathbf{B})$ and $m = \dim(\mathbf{r})$, this has a total of $n + m$ eigenvalues. From the bottom row of (4.12), we obtain $\mathbf{r} = \frac{k}{\lambda} L_{\mathbf{r}}^{-1} D\mathbf{B}$. Substituting this into the top row of (4.12) gives

$$\lambda A\mathbf{B} + kD^t L_{\mathbf{r}}^{-1} D\mathbf{B} = \lambda^2(kQ_{\mathbf{B}} + A)\mathbf{B}. \quad (4.13)$$

Through a discrete Hodge decomposition, \mathbf{B} can be written as the sum of its discrete curl-free part \mathbf{B}_A and its discrete divergence-free part \mathbf{B}_D (i.e. $\mathbf{B} = \mathbf{B}_A + \mathbf{B}_D$, where $A\mathbf{B}_A = A^t\mathbf{B}_A = 0$ and $D\mathbf{B}_D = 0$). Then (4.13) can be rewritten as

$$\lambda A\mathbf{B}_D + kD^t L_{\mathbf{r}}^{-1} D\mathbf{B}_A = \lambda^2 kQ_{\mathbf{B}}(\mathbf{B}_A + \mathbf{B}_D) + \lambda^2 A\mathbf{B}_D. \quad (4.14)$$

Let the norm induced by a symmetric positive definite matrix M be denoted $\|\cdot\|_M = \langle M\cdot, \cdot \rangle^{1/2}$. Taking the inner product of (4.14) with \mathbf{B}_A and using the relations

$$\langle Q_{\mathbf{B}}\mathbf{B}_A, \mathbf{B}_D \rangle = \langle Q_{\mathbf{B}}\mathbf{B}_D, \mathbf{B}_A \rangle = 0, \quad (4.15a)$$

$$\langle D^t L_{\mathbf{r}}^{-1} D\mathbf{B}_A, \mathbf{B}_A \rangle = \|\mathbf{B}_A\|_{Q_{\mathbf{B}}}^2, \quad (4.15b)$$

proven in [35], we have

$$k\|\mathbf{B}_A\|_{Q_{\mathbf{B}}}^2 = \lambda^2 k\|\mathbf{B}_A\|_{Q_{\mathbf{B}}}^2. \quad (4.16)$$

Because there are at least m linearly independent vectors satisfying $\mathbf{B}_A \neq 0$, this means that (4.12) has eigenvalues $\lambda = \pm 1$ each with multiplicity at least m .

Insight into the remaining $n - m$ eigenvalues can be obtained by taking the inner product of $\hat{\mathbf{B}}_D$ with (4.14), yielding

$$\lambda k\|\mathbf{B}_D\|_{Q_{\mathbf{B}}}^2 = (1 - \lambda)\|\mathbf{B}_D\|_A^2. \quad (4.17)$$

From this equation, it is clear that $0 \leq \lambda \leq 1$. These eigenvalues can be further bounded using the discrete coercivity condition [35], which can be written as

$$\frac{R_m}{S} \|\mathbf{B}_D\|_A^2 \geq \alpha \left(\frac{R_m}{S} \|\mathbf{B}_D\|_A^2 + \|\mathbf{B}_D\|_{Q_B}^2 \right), \quad (4.18)$$

where the constant $\alpha \in (0, 1)$ is independent of both the mesh and the parameters S and R_m . Applying this inequality to (4.17), we obtain the bound

$$\frac{\alpha}{\alpha + k \frac{R_m}{S} (1 - \alpha)} \leq \lambda \leq 1 \quad (4.19)$$

for the remaining $n - m$ eigenvalues.

Because α is independent of the mesh, the eigenvalues of the preconditioned matrix $\mathcal{M}_P \hat{\mathcal{M}}_{P,Q}^{-1}$ are bounded independent of h as long as k is independent of h . If we set $k = \frac{S}{R_m}$, then the lower bound in (4.19) reduces to α , making the preconditioner independent of the problem parameters R_m and S . Furthermore, as t decreases, the lower bound approaches 1, meaning that the eigenvalues cluster closer together. On the other hand, as k decreases, $A + kQ_B$ is more dominated by the singular operator A , and the component solve for this block may become more difficult. In fact, A dominates $A + kQ_B$ as long as $k \ll \frac{S}{R_m h^2}$, and if k is large enough that A does not dominate, then the lower bound in (4.19) must be $O(\alpha h^2)$. We will discuss practical considerations regarding component solves in more depth in Section 4.1.4.

4.1.1.2 Grad Div Augmentation

The next preconditioner we consider [64] arises from an application of the augmented Lagrangian methodology to Maxwell's equations. This methodology comes

from the idea of augmenting a constrained system so that its solution remains the same but the numerical properties of the system are changed. This methodology has been employed in many settings of optimization and constrained boundary value problems [26, 32, 51]. Augmented Lagrangians have also been used to develop effective block preconditioners for linear systems arising from both Maxwell's equations [63, 64] and the Navier-Stokes equations [6, 7].

In the context considered here, we augment the (1,1) block of \mathcal{M}_P by the operator $D^t W^{-1} D$, where W is a symmetric positive definite matrix. This corresponds to scaling the second row of \mathcal{M}_P by $D^t W^{-1}$ on the left and adding it to the first row. Because $D\mathbf{B} = 0$, this does not change the solution of the discrete system. Furthermore, if $A\mathbf{B} = 0$, then $D\mathbf{B} \neq 0$. Thus, the augmented (1,1) block $A + D^t W^{-1} D$ is nonsingular. Then a block LU decomposition of the augmented matrix exists, and the Schur complement is $D(A + D^t W^{-1} D)^{-1} D^t$. The Schur complement can be simplified by observing that

$$DQ_{\mathbf{B}}^{-1}(A + D^t W^{-1} D) = (DQ_{\mathbf{B}}^{-1} D^t) W^{-1} D, \quad (4.20)$$

which implies that

$$D(A + D^t W^{-1} D)^{-1} = W(DQ_{\mathbf{B}}^{-1} D^t)^{-1} DQ_{\mathbf{B}}^{-1} \quad (4.21)$$

and furthermore that

$$D(A + D^t W^{-1} D)^{-1} D^t = W. \quad (4.22)$$

This analysis then suggests a preconditioner of the form

$$\hat{\mathcal{M}}_{P,AL} := \begin{pmatrix} A + D^t W^{-1} D & 0 \\ 0 & W \end{pmatrix}. \quad (4.23)$$

The performance of this preconditioner is governed by the eigenvalues of the generalized eigenvalue problem

$$\begin{pmatrix} A & D^t \\ D & 0 \end{pmatrix} \begin{pmatrix} \mathbf{B} \\ \mathbf{r} \end{pmatrix} = \lambda \begin{pmatrix} A + D^t W^{-1} D & 0 \\ 0 & W \end{pmatrix} \begin{pmatrix} \mathbf{B} \\ \mathbf{r} \end{pmatrix}. \quad (4.24)$$

Observing that $\mathbf{r} = \frac{1}{\lambda} W^{-1} D \mathbf{B}$ from the second row, this system can be reduced to

$$A \mathbf{B} + \frac{1}{\lambda} D^t W^{-1} D \mathbf{B} = \lambda (A + D^t W^{-1} D) \mathbf{B}. \quad (4.25)$$

Writing $\mathbf{B} = \mathbf{B}_A + \mathbf{B}_D$ in terms of its discrete Hodge decomposition, this equation can be rewritten as

$$A \mathbf{B}_D + \frac{1}{\lambda} D^t W^{-1} D \mathbf{B}_A = \lambda A \mathbf{B}_D + \lambda D^t W^{-1} D \mathbf{B}_A. \quad (4.26)$$

Taking the inner product of (4.26) with \mathbf{B}_A yields

$$\frac{1}{\lambda} \|D \mathbf{B}_A\|_{W^{-1}}^2 = \lambda \|D \mathbf{B}_A\|_{W^{-1}}^2. \quad (4.27)$$

Because there are at least m linearly independent vectors satisfying $\mathbf{B}_A \neq 0$, this means that (4.24) has eigenvalues $\lambda = \pm 1$ each with multiplicity at least m . The remaining eigenvalues are obtained by taking the inner product of (4.26) with \mathbf{B}_D , which yields

$$\|\mathbf{B}_D\|_A^2 = \lambda \|\mathbf{B}_D\|_A^2. \quad (4.28)$$

This gives $\lambda = 1$ with multiplicity $n - m$. Thus, the preconditioned system $\mathcal{M}_P \hat{\mathcal{M}}_{P,AL}^{-1}$ has n eigenvalues of $\lambda = 1$ and m eigenvalues of $\lambda = -1$, and this result is independent of the mesh size h , the parameters S and R_m , and the matrix W .

In this study we consider the case where $W = \frac{1}{k}Q_{\mathbf{r}}$, yielding a preconditioner of the form

$$\hat{\mathcal{M}}_{P,GD} := \begin{pmatrix} A + kD^tQ_{\mathbf{r}}^{-1}D & 0 \\ 0 & \frac{1}{k}Q_{\mathbf{r}} \end{pmatrix}. \quad (4.29)$$

Then the augmentation operator $kD^tQ_{\mathbf{r}}^{-1}D$ corresponds to the continuous operator $-k\nabla\nabla\cdot$. For a practical preconditioner, we replace $Q_{\mathbf{r}}$ by its diagonal $\bar{Q}_{\mathbf{r}}$. Although the choice of k should not affect the performance of the block preconditioner, the choice $k = \frac{S}{R_m}$ may be advantageous for the component solve. With this choice, the discrete operator $A + kD^tQ_{\mathbf{r}}^{-1}D$ is a discrete version of the continuous operator $\frac{S}{R_m}(\nabla \times \nabla \times -\nabla\nabla\cdot) = -\frac{S}{R_m}\Delta$. This is attractive since multigrid is known to work well for the (vector) Laplacian.

4.1.2 An Approximation of X_P

We turn our attention now to developing an approximation for X_P as defined in (4.6). We focus first on simplifying the term $\mathcal{Z}\mathcal{M}_P^{-1}\mathcal{Z}^t$. We can derive an expression for \mathcal{M}_P^{-1} by observing that

$$\mathcal{M}_P = \begin{pmatrix} I & -D^tQ_{\mathbf{r}}^{-1} \\ 0 & I \end{pmatrix} \begin{pmatrix} A + D^tQ_{\mathbf{r}}^{-1}D & D^t \\ D & 0 \end{pmatrix}. \quad (4.30)$$

As shown above, the augmented matrix $A + D^tQ_{\mathbf{r}}^{-1}D$ is nonsingular, so the second factor in (4.30) admits a block LU decomposition. Recalling from (4.22) that the Schur complement associated with the augmented system simplifies to $-Q_{\mathbf{r}}$, we

obtain

$$\mathcal{M}_P = \begin{pmatrix} I & -D^t Q_r^{-1} \\ 0 & I \end{pmatrix} \begin{pmatrix} I & 0 \\ D(A + D^t Q_r^{-1} D)^{-1} & I \end{pmatrix} \begin{pmatrix} A + D^t Q_r^{-1} D & D^t \\ 0 & -Q_r \end{pmatrix}. \quad (4.31)$$

Each term can then be inverted explicitly, yielding

$$\mathcal{M}_P^{-1} = \begin{pmatrix} (A + D^t Q_r^{-1} D)^{-1} & (A + D^t Q_r^{-1} D)^{-1} D^t Q_r^{-1} \\ 0 & -Q_r^{-1} \end{pmatrix} \begin{pmatrix} I & 0 \\ -D(A + D^t Q_r^{-1} D)^{-1} & I \end{pmatrix} \begin{pmatrix} I & D^t Q_r^{-1} \\ 0 & I \end{pmatrix}. \quad (4.32)$$

Using this expression for \mathcal{M}_P^{-1} , we can obtain

$$\mathcal{Z} \mathcal{M}_P^{-1} \mathcal{Z}^t = \mathcal{Z} (A + D^t Q_r^{-1} D)^{-1} A (A + D^t Q_r^{-1} D)^{-1} \mathcal{Z}^t. \quad (4.33)$$

We proceed now by considering the continuous operator corresponding to $\mathcal{Z} \mathcal{M}_P^{-1} \mathcal{Z}^t$ as defined in (4.33). This is the strategy we took to approximate the velocity Schur complement for the exact penalty formulation in Chapter 3. Replacing each discrete operator on the right-hand side of (4.33) with its corresponding continuous operator from Table 4.1, we have that $\mathcal{Z} \mathcal{M}_P^{-1} \mathcal{Z}^t$ is a discretization of

$$\mathcal{K} := \frac{S^3}{R_m} \vec{b} \times \left[\nabla \times \left(\frac{S}{R_m} \nabla \times \nabla \times -\nabla \nabla \cdot \right)^{-1} \nabla \times \right]^2 (\cdot \times \vec{b}). \quad (4.34)$$

This operator behaves differently in two and three dimensions, so we consider the two cases separately.

4.1.2.1 The Two-Dimensional Case

In two dimensions, the following relation holds:

$$\nabla \times \left(\frac{S}{R_m} \nabla \times \nabla \times -\nabla \nabla \cdot \right) = -\frac{S}{R_m} \Delta_r \nabla \times . \quad (4.35)$$

This is obtained from the facts that $\nabla \times \nabla = 0$ and that $\nabla \times \nabla \times = -\Delta_r$ when applied to a scalar function (see Section 2.1 for the two definitions of the curl in 2D; the subscript r in Δ_r indicates a scalar Laplacian, as opposed to the vector Laplacian Δ). Equation (4.35) implies that

$$\nabla \times \left(\frac{S}{R_m} \nabla \times \nabla \times -\nabla \nabla \cdot \right)^{-1} = -\frac{R_m}{S} \Delta_r^{-1} \nabla \times , \quad (4.36)$$

which, in turn, yields

$$\nabla \times \left(\frac{S}{R_m} \nabla \times \nabla \times -\nabla \nabla \cdot \right)^{-1} \nabla \times = -\frac{R_m}{S} \Delta_r^{-1} \nabla \times \nabla \times = \frac{R_m}{S} . \quad (4.37)$$

Substituting this into the definition of \mathcal{K} , we obtain

$$\mathcal{K} = SR_m \vec{b} \times (\cdot \times \vec{b}) \quad (4.38)$$

in two dimensions. This is exactly the same operator that appeared in the velocity Schur complement for the exact penalty formulation in Section 3.1.1. Hence, we know that this operator is readily discretized, and its corresponding discrete operator can be written as

$$K := SR_m Q_{\mathbf{u}} \begin{pmatrix} \text{diag}(b_y^2) & -\text{diag}(b_x b_y) \\ -\text{diag}(b_x b_y) & \text{diag}(b_x^2) \end{pmatrix} . \quad (4.39)$$

Given this discrete operator K , the Schur complement on the velocity space \hat{X}_P can be approximated by

$$\hat{X}_P := F + K. \quad (4.40)$$

As we know from Chapter 3, \hat{X}_P is similar in structure to a time-dependent convection-diffusion operator or a convection-diffusion-reaction operator, so we can expect that multigrid will perform well to approximate the action of \hat{X}_P^{-1} in the application of the preconditioner.

4.1.2.2 The Three-Dimensional Case

In three dimensions, the operator $\left[\nabla \times \left(\frac{S}{R_m} \nabla \times \nabla \times -\nabla \nabla \cdot \right)^{-1} \nabla \times \right]^2$ can be simplified by observing that

$$\left(\frac{S}{R_m} \nabla \times \nabla \times -\nabla \nabla \cdot \right)^{-1} \nabla \times \nabla \times = \frac{R_m}{S} I + \frac{R_m}{S} \left(\frac{S}{R_m} \nabla \times \nabla \times -\nabla \nabla \cdot \right)^{-1} \nabla \nabla \cdot. \quad (4.41)$$

Now, observe that

$$\left(\frac{S}{R_m} \nabla \times \nabla \times -\nabla \nabla \cdot \right) \nabla = -\nabla \Delta_r, \quad (4.42)$$

which implies that

$$\left(\frac{S}{R_m} \nabla \times \nabla \times -\nabla \nabla \cdot \right)^{-1} \nabla = -\nabla \Delta_r^{-1}. \quad (4.43)$$

Substituting of this expression into (4.41) yields

$$\left(\frac{S}{R_m} \nabla \times \nabla \times -\nabla \nabla \cdot \right)^{-1} \nabla \times \nabla \times = \frac{R_m}{S} I - \frac{R_m}{S} \nabla \Delta_r^{-1} \nabla \cdot. \quad (4.44)$$

Using this relation, we can simplify the operator $\left[\nabla \times \left(\frac{S}{R_m} \nabla \times \nabla \times -\nabla \nabla \cdot \right)^{-1} \nabla \times \right]^2$ to give

$$\left[\nabla \times \left(\frac{S}{R_m} \nabla \times \nabla \times -\nabla \nabla \cdot \right)^{-1} \nabla \times \right]^2 = \frac{R_m}{S} \nabla \times \left(\frac{S}{R_m} \nabla \times \nabla \times -\nabla \nabla \cdot \right)^{-1} \nabla \times . \quad (4.45)$$

Using the fact that the commutator

$$\nabla \times \left(\frac{S}{R_m} \nabla \times \nabla \times -\nabla \nabla \cdot \right) - \left(\frac{S}{R_m} \nabla \times \nabla \times -\nabla \nabla \cdot \right) \nabla \times \quad (4.46)$$

is exactly zero, we obtain that the inverse operator $\left(\frac{S}{R_m} \nabla \times \nabla \times -\nabla \nabla \cdot \right)^{-1}$ also commutes with the curl, so that

$$\left(\frac{S}{R_m} \nabla \times \nabla \times -\nabla \nabla \cdot \right)^{-1} \nabla \times = \nabla \times \left(\frac{S}{R_m} \nabla \times \nabla \times -\nabla \nabla \cdot \right)^{-1}. \quad (4.47)$$

With this relation, we have that

$$\begin{aligned} \nabla \times \left(\frac{S}{R_m} \nabla \times \nabla \times -\nabla \nabla \cdot \right)^{-1} \nabla \times &= \nabla \times \nabla \times \left(\frac{S}{R_m} \nabla \times \nabla \times -\nabla \nabla \cdot \right)^{-1} \quad (4.48) \\ &= \frac{R_m}{S} \left[I + \nabla \nabla \cdot \left(\frac{S}{R_m} \nabla \times \nabla \times -\nabla \nabla \cdot \right)^{-1} \right]. \end{aligned}$$

This can be further simplified using the relation

$$\nabla \cdot \left(\frac{S}{R_m} \nabla \times \nabla \times -\nabla \nabla \cdot \right) = -\Delta_r \nabla \cdot, \quad (4.49)$$

which implies that

$$\nabla \nabla \cdot \left(\frac{S}{R_m} \nabla \times \nabla \times -\nabla \nabla \cdot \right)^{-1} = -\nabla \Delta_r^{-1} \nabla \cdot . \quad (4.50)$$

Furthermore, it can be shown that

$$I - \nabla \Delta_r^{-1} \nabla \cdot = \Delta^{-1} (\Delta - \nabla \nabla \cdot) = -\nabla \times \Delta^{-1} \nabla \times, \quad (4.51)$$

using the fact that $-\Delta = \nabla \times \nabla \times -\nabla \nabla \cdot$. Thus, we can rewrite

$$\nabla \times \left(\frac{S}{R_m} \nabla \times \nabla \times -\nabla \nabla \cdot \right)^{-1} \nabla \times = -\frac{R_m}{S} \nabla \times \Delta^{-1} \nabla \times. \quad (4.52)$$

Using results (4.45) and (4.45), \mathcal{K} can be written as

$$\mathcal{K} = SR_m \vec{b} \times (\nabla \times \Delta^{-1} \nabla \times) (\cdot \times \vec{b}) \quad (4.53)$$

or

$$\mathcal{K} = SR_m \vec{b} \times (I - \nabla \Delta_r^{-1} \nabla \cdot) (\cdot \times \vec{b}) \quad (4.54)$$

in the three-dimensional setting

As written, neither expression (4.53) nor expression (4.54) can be used directly to approximate X_P because inverse operators appear in both of them. We now focus on approximations of the operator $\nabla \times \Delta^{-1} \nabla \times = I - \nabla \Delta_r^{-1} \nabla \cdot$ that can produce a viable approximation of \mathcal{K} . Toward this end, we note that if $\nabla \cdot (\vec{u} \cdot \vec{b}) \approx 0$, then $(I - \nabla \Delta_r^{-1} \nabla \cdot) (\vec{u} \cdot \times \vec{b}) \approx \vec{u} \times \vec{b}$. Consequently, \mathcal{K} can be approximated by $SR_m \vec{b} \times (\cdot \times \vec{b})$, which is the same as what we obtained in the two-dimensional case. On the other hand, if $\nabla \times (\vec{u} \times \vec{b}) \approx 0$, then $\mathcal{K} \approx 0$. Thus, we have two conflicting approximations which depend on the character of the quantity $\vec{u} \times \vec{b}$. To account for both of these extremes, we propose introducing a parameter $\beta \in [0, 1]$ and approximating

$$\mathcal{K} \approx \beta SR_m \vec{b} \times (\cdot \times \vec{b}), \quad (4.55)$$

such that $\beta \approx 1$ if $\nabla \cdot (\vec{u} \times \vec{b}) \approx 0$ and $\beta \approx 0$ if $\nabla \times (\vec{u} \times \vec{b}) \approx 0$.

While there is no clear way to approximate $\nabla \cdot (\vec{u} \times \vec{b})$, the expression $\nabla \times (\vec{u} \times \vec{b})$ is the convection term appearing in the induction equation (2.21c), and we can

approximate its size based on qualities of the MHD equations. Thus, we focus on deriving an expression for β such that $\beta \approx 0$ when $\nabla \times (\vec{u} \times \vec{b}) \approx 0$ and $\beta \rightarrow 1$ as $\nabla \times (\vec{u} \times \vec{b})$ increases. Note that $\nabla \times (\vec{u} \times \vec{b}) = 0$ in the diffusion-dominated limit where $R_m \rightarrow 0$ or when $\vec{b} = \vec{0}$. Furthermore, $\nabla \times (\vec{u} \times \vec{b})$ becomes more prominent in the convection-dominated limit where $R_m \rightarrow \infty$ or $|\vec{b}| \rightarrow \infty$. Hence, we want β to satisfy $\beta \rightarrow 0$ when the induction equation is diffusion-dominated and $\beta \rightarrow 1$ when the induction equation is convection-dominated. In the discrete setting, diffusion-dominated corresponds to the case where $\|A\| \gg \|Z^t\|$. This occurs when $\frac{S}{R_m h^2} \gg \frac{S|\vec{b}|}{h}$. Similarly, convection-dominated corresponds to $\frac{S}{R_m h^2} \ll \frac{S|\vec{b}|}{h}$. Hence, an expression for β that satisfies the desired conditions is

$$\beta_* := \frac{R_m h |\vec{b}|}{1 + R_m h |\vec{b}|}. \quad (4.56)$$

In the diffusion-dominated limit, $R_m h |\vec{b}| \ll 1$, so $\beta_* \approx 0$, and in the convection-dominated limit, $R_m h |\vec{b}| \gg 1$, so $\beta_* \approx 1$.

Thus, motivated by approximation (4.55), we approximate the Schur complement X_P in three dimensions by

$$\hat{X}_P := F + \beta K, \quad (4.57)$$

where K is a discretization of the operator $SR_m \vec{b} \times (\cdot \times \vec{b})$. We expect β_* to be a good choice for β and will investigate this further in Section 4.2.5. As in the two-dimensional setting, K takes the form of a scaled mass matrix and can be written

$$K := SR_m Q_{\mathbf{u}} \begin{pmatrix} \text{diag}(b_y^2 + b_z^2) & -\text{diag}(b_x b_y) & -\text{diag}(b_x b_z) \\ -\text{diag}(b_x b_y) & \text{diag}(b_x^2 + b_z^2) & -\text{diag}(b_y b_z) \\ -\text{diag}(b_x b_z) & -\text{diag}(b_y b_z) & \text{diag}(b_x^2 + b_y^2) \end{pmatrix}. \quad (4.58)$$

Hence, \hat{X}_P again has a desirable structure, and we expect multigrid to perform well in approximating \hat{X}_P^{-1} .

4.1.3 A Commutator for \hat{Y}_P

In both the two- and three-dimensional settings, the velocity Schur complement takes the form

$$\hat{X}_P = F + \beta K, \quad (4.59)$$

where β is one in 2D and $\beta \in [0, 1]$ in 3D. This is the same form as the velocity Schur complement expression used for the exact penalty formulation in Chapter 3 (see Section 3.1.1). Hence, we follow the strategy developed in Section 3.1.2 to make the approximation

$$Y_P \approx -B(F + \beta K)^{-1}B^t. \quad (4.60)$$

Specifically, we make the continuous approximation

$$\nabla \cdot (\mathcal{F} + \beta \mathcal{K}) \approx (\mathcal{F}_p + \alpha \beta \mathcal{K}_p) \nabla \cdot, \quad (4.61)$$

where $\alpha \in [0, 1]$ is a relaxation parameter included to compensate for the fact that the commutator $\nabla \cdot \beta \mathcal{K} - \beta \mathcal{K}_p \nabla \cdot$ is small only when $\beta SR_m \|\vec{b}\| \approx 0$. A discretization of (4.61) yields the relation

$$Q_{\mathbf{p}}^{-1} B Q_{\mathbf{u}}^{-1} (F + \beta K) \approx Q_{\mathbf{p}}^{-1} (F_p + \alpha \beta K_p) Q_{\mathbf{p}}^{-1} B, \quad (4.62)$$

which, through some algebraic manipulation gives

$$-B(F + \beta K)^{-1}B^t \approx Q_{\mathbf{p}}(F_{\mathbf{p}} + \alpha \beta K_{\mathbf{p}})^{-1}BQ_{\mathbf{u}}^{-1}B^t. \quad (4.63)$$

As in Chapter 3, we define the operators $F_{\mathbf{p}}$ and $K_{\mathbf{p}}$ following the strategy introduced to develop the Least-Squares Commutator (LSC) Schur complement approximation for the Navier-Stokes equations [25]. That is, $F_{\mathbf{p}}Q_{\mathbf{p}}^{-1}$ and $K_{\mathbf{p}}Q_{\mathbf{p}}^{-1}$ are computed one row at a time, where each row is obtained by solving the least-squares problems

$$\min\| [BQ_{\mathbf{u}}^{-1}F]_{j*} - [F_{\mathbf{p}}Q_{\mathbf{p}}^{-1}]_{j*}B \|_{Q_{\mathbf{u}}^{-1}}, \quad (4.64a)$$

$$\min\| [BQ_{\mathbf{u}}^{-1}K]_{j*} - [K_{\mathbf{p}}Q_{\mathbf{p}}^{-1}]_{j*}B \|_{Q_{\mathbf{u}}^{-1}}. \quad (4.64b)$$

This results in the definitions

$$F_{\mathbf{p}} := BQ_{\mathbf{u}}^{-1}FQ_{\mathbf{u}}^{-1}B^t(BQ_{\mathbf{u}}^{-1}B^t)^{-1}Q_{\mathbf{p}}, \quad (4.65a)$$

$$K_{\mathbf{p}} := BQ_{\mathbf{u}}^{-1}KQ_{\mathbf{u}}^{-1}B^t(BQ_{\mathbf{u}}^{-1}B^t)^{-1}Q_{\mathbf{p}}. \quad (4.65b)$$

Substituting these definitions into (4.63) gives the approximation

$$-B(F + \beta K)^{-1}B^t \approx -BQ_{\mathbf{u}}^{-1}B^t[BQ_{\mathbf{u}}^{-1}(F + \alpha\beta K)Q_{\mathbf{u}}^{-1}B^t]^{-1}BQ_{\mathbf{u}}^{-1}B^t. \quad (4.66)$$

In practice, we can replace $Q_{\mathbf{u}}$ by its diagonal $\bar{Q}_{\mathbf{u}}$ for ease of computation [62]. Then the action of the inverse of this approximation requires only two solves with the scaled Laplacian operator $B\bar{Q}_{\mathbf{u}}^{-1}B^t$. The Fourier analysis used to specify α for the exact penalty formulation (see Section 3.1.2) applies directly here. Thus, we propose using the value

$$\alpha_* := \frac{1 + \beta H^2 h_{\mathbf{p}}^2 \frac{(\vec{a} \cdot \vec{b})^2}{\|\vec{a}\|^2} + R^2 h_{\mathbf{p}}^2 \|\vec{a}\|^2}{\left[1 + \beta H^2 h_{\mathbf{p}}^2 \frac{(\vec{a} \cdot \vec{b})^2}{\|\vec{a}\|^2}\right]^2 + R^2 h_{\mathbf{p}}^2 \|\vec{a}\|^2} \quad (4.67)$$

for the parameter α . This value can be generated automatically based on the problem parameters R , R_m , and S , as well as the mesh size for the pressure variable $h_{\mathbf{p}}$ and the previous iterates in the nonlinear iteration \vec{a} and \vec{b} .

4.1.4 Computational Tasks

With expressions for $\widehat{\mathcal{M}}_P$, \widehat{X}_P , and \widehat{Y}_P , we have two block preconditioners for the Picard matrix \mathcal{A}_P , which we write

$$\mathcal{P}_{P,Q} = \begin{pmatrix} \widehat{A + kQ_{\mathbf{B}}} & 0 & -Z^t & 0 \\ 0 & \frac{1}{k}\widehat{L}_{\mathbf{r}} & 0 & 0 \\ 0 & 0 & \widehat{X}_P & B^t \\ 0 & 0 & 0 & \widehat{Y}_P \end{pmatrix}, \quad (4.68)$$

$$\mathcal{P}_{P,GD} = \begin{pmatrix} \widehat{A + kD^t\bar{Q}_{\mathbf{r}}^{-1}D} & 0 & -Z^t & 0 \\ 0 & \frac{1}{k}\bar{Q}_{\mathbf{r}} & 0 & 0 \\ 0 & 0 & \widehat{X}_P & B^t \\ 0 & 0 & 0 & \widehat{Y}_P \end{pmatrix}, \quad (4.69)$$

where $\widehat{A + kQ_{\mathbf{B}}}$, $\widehat{L}_{\mathbf{r}}$, and $\widehat{A + kD^t\bar{Q}_{\mathbf{r}}^{-1}D}$ are approximations of $A + kQ_{\mathbf{B}}$, $L_{\mathbf{r}}$, and $A + kD^t\bar{Q}_{\mathbf{r}}^{-1}D$. Because $\bar{Q}_{\mathbf{r}}$ is diagonal, the action of its inverse can be applied explicitly. The operator \widehat{X}_P in both cases is an approximation of $F + \beta K$, where $\beta \equiv 1$ in 2D and $\beta \in [0, 1]$ in 3D. The operator \widehat{Y}_P is then defined as

$$\widehat{Y}_P := -\widehat{B\bar{Q}_{\mathbf{u}}^{-1}B^t} [B\bar{Q}_{\mathbf{u}}^{-1}(F + \alpha\beta K)\bar{Q}_{\mathbf{u}}^{-1}B^t]^{-1} \widehat{B\bar{Q}_{\mathbf{u}}^{-1}B^t}, \quad (4.70)$$

where $\widehat{B\bar{Q}_{\mathbf{u}}^{-1}B^t}$ is an approximation of $B\bar{Q}_{\mathbf{u}}^{-1}B^t$. The application of both preconditioners require the action of the inverse of the perturbed convection-diffusion matrix $F + \beta K$ and the scaled Laplacian $B\bar{Q}_{\mathbf{u}}^{-1}B^t$. Furthermore, $\mathcal{P}_{P,Q}$ requires the action of the inverse of the Laplacian $L_{\mathbf{r}}$. In practice, we replace all these operators these by one V-cycle of algebraic multigrid.

The remaining computational task in applying these preconditioners is approximating the inverse of the operators $A + kQ_{\mathbf{B}}$ and $A + kD^t\bar{Q}_{\mathbf{r}}^{-1}D$. Because of the prominence of the singular operator A in these expressions, standard multigrid using traditional smoothers and coarsening algorithms may not necessarily perform well. That is, traditional multigrid smoothers such as Jacobi and Gauss-Seidel smoothers rely on errors being well represented in the residual of the discrete operator [41], and this may not be the case for $A + kQ_{\mathbf{B}}$ and $A + kD^t\bar{Q}_{\mathbf{r}}^{-1}D$. To see this, we decompose an error e into its curl-free and divergence-free parts as $e = e_A + e_D$, where $Ae_A = 0$ and $De = 0$. Then the residuals can be written as

$$r_Q := (A + kQ_{\mathbf{B}})e = Ae_D + kQ_{\mathbf{B}}(e_A + e_D), \quad (4.71)$$

$$r_{GD} := (A + kD^t\bar{Q}_{\mathbf{r}}^{-1}D)e = Ae_D + kD^t\bar{Q}_{\mathbf{r}}De_A. \quad (4.72)$$

If $k \ll \frac{S}{R_m h^2}$, then $\|A\| \gg \|kQ_{\mathbf{B}}\|$ and $r_Q \approx Ae_D$. As discussed in Section 4.1.1.1, we require $k \ll \frac{S}{R_m h^2}$ in order for the eigenvalues of $\mathcal{M}_P \hat{\mathcal{M}}_{P,Q}^{-1}$ to be mesh-independent. Thus, in the context of the preconditioner $\mathcal{P}_{P,Q}$, $r_Q \approx Ae_D$, meaning that the curl-free part of the error e_A is not represented in the residual r_Q . Several sophisticated multigrid algorithms have been proposed for operators with the same structure as $A + kQ_{\mathbf{B}}$ (see, e.g., [39, 41, 54] and the references therein). These operators arise from the discretization of weak formulations of the form

$$(\nabla \times \vec{B}, \nabla \times \vec{C}) + (\sigma \vec{B}, \vec{C}) \quad (4.73)$$

with edge elements for \vec{B} . An algebraic multigrid method was proposed for such operators by generating a coarse-grid hierarchy that preserves the kernel of the curl-curl operator on each level [54]. The prolongation operator for this hierarchy can

be derived from the prolongator used in coarsening a nodal discretization of the auxiliary scalar weak formulation

$$(\nabla r, \nabla q) + (\sigma r, q). \quad (4.74)$$

Improved prolongation operators have since been proposed that lead to better scalability with respect to mesh refinement [8, 41].

If $k \ll \frac{S}{R_m}$, then $\|A\| \gg \|kD^t\bar{Q}_r\|$ and $r_{GD} \approx Ae_D$. Similarly, if $k \gg \frac{S}{R_m}$, then $\|A\| \ll \|kD^t\bar{Q}_r\|$ and $r_{GD} \approx kD^t\bar{Q}_rDe_A$. Thus, unless k is comparable to $\frac{S}{R_m}$, r_{GD} does not represent both components of the error. Because neither the performance of $\hat{\mathcal{M}}_{P,GD}$ as a preconditioner for \mathcal{M}_P nor the approximation of $X_{P,GD}$ depend on k , we can safely choose k in preconditioner $\mathcal{P}_{P,GD}$ such that both components of the error e are represented in the residual r_{GD} . A convenient choice of k is then $k = \frac{S}{R_m}$, so that $\|A\| \approx \|kD^t\bar{Q}_r\|$. In fact, when $k = \frac{S}{R_m}$, $A + kD^t\bar{Q}_r^{-1}D$ approximates the vector Laplacian $-\frac{S}{R_m}\Delta = \frac{S}{R_m}(\nabla \times \nabla \times -\nabla\nabla\cdot)$. In this case, traditional multigrid should work well to approximate the inverse of $A + kD^t\bar{Q}_r^{-1}D$.

4.2 Block Preconditioners for the Newton System

Many of the same strategies we used to develop preconditioners for the Picard matrix \mathcal{A}_P can be applied to the Newton matrix \mathcal{A}_N , but they should be adapted to account for the presence of the additional operators \tilde{F} , \tilde{A} , and \tilde{Z} . As we did for the Picard matrix in Section 4.1, we begin by reordering the unknowns as $(\mathbf{B}, \mathbf{r}, \mathbf{u}, \mathbf{p})$

and blocking the electromagnetic unknowns together so we can write

$$\mathcal{A}_N = \begin{pmatrix} \mathcal{M}_N & -\mathcal{Z}^t & 0 \\ \mathcal{Z} + \tilde{\mathcal{Z}} & F + \tilde{F} & B^t \\ 0 & B & 0 \end{pmatrix}, \quad (4.75)$$

where \mathcal{M}_N is defined to be

$$\mathcal{M}_N = \begin{pmatrix} A + \tilde{A} & D^t \\ D & 0 \end{pmatrix}. \quad (4.76)$$

This corresponds to a discretization of Maxwell's equations in mixed form including the additional convection operator $S\nabla \times (\vec{a} \times \vec{B})$. A block LU decomposition of \mathcal{A}_N then suggests a block preconditioner of the form

$$\mathcal{P}_N = \begin{pmatrix} \hat{\mathcal{M}}_N & -\mathcal{Z}^t & 0 \\ 0 & \hat{X}_N & B^t \\ 0 & 0 & \hat{Y}_N \end{pmatrix}, \quad (4.77)$$

where $\hat{\mathcal{M}}_N$ is a preconditioner for \mathcal{M}_N , and

$$\hat{X}_N \approx F + \tilde{F} + (\mathcal{Z} + \tilde{\mathcal{Z}})\mathcal{M}_N^{-1}\mathcal{Z}^t, \quad (4.78)$$

$$\hat{Y}_N \approx -B\mathcal{X}_N^{-1}B^t. \quad (4.79)$$

We proceed by discussing how to adapt each of the approximations $\hat{\mathcal{M}}_P$, \hat{X}_P , and \hat{Y}_P from the Picard preconditioner for use with the Newton system.

4.2.1 Preconditioners for \mathcal{M}_N

Because we are not aware of any preconditioners developed specifically for Maxwell's equations with a convection term, we adapt the two preconditioners dis-

cussed above for the matrix \mathcal{M}_N . Specifically, we propose the two preconditioners

$$\hat{\mathcal{M}}_{N,Q} = \begin{pmatrix} A + \tilde{A} + kQ_{\mathbf{B}} & 0 \\ 0 & \frac{1}{k}L_{\mathbf{r}} \end{pmatrix}, \quad (4.80)$$

$$\hat{\mathcal{M}}_{N,GD} = \begin{pmatrix} A + \tilde{A} + kD^tQ_{\mathbf{r}}^{-1}D & 0 \\ 0 & \frac{1}{k}Q_{\mathbf{r}} \end{pmatrix}. \quad (4.81)$$

In each case, the only difference from the Picard case is that we have included the convection term \tilde{A} in the (1,1) block of the preconditioner.

The eigenvalues of the preconditioned system $\mathcal{M}_N\hat{\mathcal{M}}_{N,Q}^{-1}$ can be analyzed by modifying the arguments of [35] that we discussed in Section 4.1.1.1. Using the fact that if $A\mathbf{B}_A = 0$, then $\mathbf{B}_A^t\tilde{A} = 0$ it can be shown that $\mathcal{M}_N\hat{\mathcal{M}}_{N,Q}^{-1}$ has eigenvalues $\lambda = \pm 1$ each with multiplicity m . The remaining $n - m$ eigenvalues satisfy the relationship

$$\lambda k \|\mathbf{B}_D\|_{Q_{\mathbf{B}}}^2 = (1 - \lambda) \left(\|\mathbf{B}_D\|_A^2 + \langle \mathbf{B}_D, \tilde{A}\mathbf{B}_D \rangle \right) \quad (4.82)$$

with \mathbf{B}_D satisfying $D\mathbf{B}_D = 0$. Because no coercivity condition is guaranteed for the operator \tilde{A} , this cannot be simplified further as we did with (4.17), but we expect the remaining eigenvalues to behave well. We analyze a generalization of this preconditioner in more depth for the MHD kinematics system in Chapter 5. In Section 5.3, we empirically demonstrate that the ideal version of this preconditioner is mesh-independent for a class of systems that includes \mathcal{M}_N .

The eigenvalues of the preconditioned system $\mathcal{M}_N\hat{\mathcal{M}}_{N,GD}^{-1}$ can be analyzed by modifying the arguments of [64] that we discussed in Section 4.1.1.2. Using the fact that if $A\mathbf{B}_A = 0$, then $\mathbf{B}_A^t\tilde{A} = 0$, the same argument as in Section 4.1.1.2 shows

that $\mathcal{M}_N \hat{\mathcal{M}}_{N,GD}^{-1}$ has n eigenvalues $\lambda = 1$ and m eigenvalues $\lambda = -1$.

4.2.2 Approximations of X_N and Y_N

The major difference in developing preconditioners for the Picard and Newton systems arises in approximating X_N because the extra convection term \tilde{A} is embedded in the exact Schur complement $F + \tilde{F} + (\mathcal{Z} + \tilde{\mathcal{Z}})\mathcal{M}_N^{-1}\mathcal{Z}^t$. Since this was also the case for the exact penalty formulation in Chapter 3, we build on the techniques explored in Section 3.2 to approximate X_N here. Keeping with this strategy, we assume that $\tilde{\mathcal{Z}}$ does not play a prominent role in the velocity Schur complement and focus on approximating the operator $\mathcal{Z}\mathcal{M}_N^{-1}\mathcal{Z}^t$.

Using a decomposition analogous to (4.32) for \mathcal{M}_N , we can compute

$$\mathcal{Z}\mathcal{M}_N^{-1}\mathcal{Z}^t = Z(A + \tilde{A} + D^t Q_r^{-1} D)^{-1} (A + \tilde{A}) (A + \tilde{A} + D^t Q_r^{-1} D)^{-1} Z^t. \quad (4.83)$$

As in Section 4.1.2, we simplify this discrete operator by considering its corresponding continuous operator

$$\begin{aligned} \tilde{\mathcal{K}} := & S^2 \vec{b} \times \nabla \times \left[\frac{S}{R_m} \nabla \times \nabla \times -S \nabla \times (\vec{a} \times \cdot) - \nabla \nabla \cdot \right]^{-1} \left[\frac{S}{R_m} \nabla \times \nabla \times -S \nabla \times (\vec{a} \times \cdot) \right] \\ & \left[\frac{S}{R_m} \nabla \times \nabla \times -S \nabla \times (\vec{a} \times \cdot) - \nabla \nabla \cdot \right]^{-1} \nabla \times (\cdot \times \vec{b}) \end{aligned} \quad (4.84)$$

In the two-dimensional setting, the following relationship holds:

$$\left[\frac{S}{R_m} \nabla \times \nabla \times -S \nabla \times (\vec{a} \times \cdot) - \nabla \nabla \cdot \right] \nabla \times = -\nabla \times \left(\frac{S}{R_m} \Delta_r + S \vec{a} \times \right). \quad (4.85)$$

Using this fact, it can be shown that

$$\tilde{\mathcal{K}} = S^2 \vec{b} \times \Delta_r \left(\frac{S}{R_m} \Delta_r + S \vec{a} \times \nabla \times \right)^{-1} (\cdot \times \vec{b}) \quad (4.86)$$

in the two-dimensional setting. Employing a similar strategy to that developed in Section 3.2, we can simplify this expression by approximating

$$\Delta_r \left(\frac{S}{R_m} \Delta_r + S\vec{a} \times \nabla \times \right)^{-1} \approx \gamma, \quad (4.87)$$

where $\gamma \in [0, 1]$. We should have $\gamma \approx 1$ when the operator $\frac{S}{R_m} \Delta_r + S\vec{a} \times \nabla \times$ is diffusion-dominated and $\gamma \approx 0$ when it is convection-dominated. A discretization of $\frac{S}{R_m} \Delta_r + S\vec{a} \times \nabla \times$ defined for the unknown \mathbf{r} is diffusion-dominated when $\frac{S}{R_m h_{\mathbf{r}}^2} \gg \frac{S}{\|\vec{a}\|} h_{\mathbf{r}}$, where $h_{\mathbf{r}}$ is the mesh size on the magnetic pseudo-pressure space. Thus, an expression for γ with the desired behavior is

$$\gamma_* := \frac{1}{1 + R_m h_{\mathbf{r}} \|\vec{a}\|}, \quad (4.88)$$

which is exactly the same expression derived for the exact penalty formulation in Section 3.2.

In the three-dimensional setting, we can use similar techniques to those in Section 4.1.2.2 to show that

$$\left[\frac{S}{R_m} \nabla \times \nabla \times - S \nabla \times (\vec{a} \times \cdot) \right] \left[\frac{S}{R_m} \nabla \times \nabla \times - S \nabla \times (\vec{a} \times \cdot) - \nabla \nabla \cdot \right]^{-1} = -\nabla \times \Delta^{-1} \nabla \times. \quad (4.89)$$

Furthermore, it can be shown that

$$\nabla \times \left[\frac{S}{R_m} \nabla \times \nabla \times - S \nabla \times (\vec{a} \times \cdot) - \nabla \nabla \cdot \right]^{-1} = \left[\frac{S}{R_m} \nabla \times \nabla \times - S \vec{a} \times \nabla \times - \nabla \nabla \cdot \right]^{-1} \nabla \times. \quad (4.90)$$

Thus, the three-dimensional version of $\tilde{\mathcal{K}}$ reduced to

$$\tilde{\mathcal{K}} = S^2 \vec{b} \times \left[\frac{S}{R_m} \nabla \times \nabla \times - S \vec{a} \times \nabla \times - \nabla \nabla \cdot \right]^{-1} \nabla \times \nabla \times (-\nabla \times \Delta^{-1} \nabla \times) (\cdot \times \vec{b}). \quad (4.91)$$

As in Section 4.1.2.2, we can approximate $(-\nabla \times \Delta^{-1} \nabla \times) (\vec{u} \times \vec{b})$ by $\beta (\vec{u} \times \vec{b})$ with $\beta \in [0, 1]$, and a good choice of β given by β_* as defined in (4.56). Then we have

$$\tilde{\mathcal{K}} \approx \beta S^2 \vec{b} \times \left[\frac{S}{R_m} \nabla \times \nabla \times - S \vec{a} \times \nabla \times - \nabla \nabla \cdot \right]^{-1} \nabla \times \nabla \times (\cdot \times \vec{b}). \quad (4.92)$$

We can make a similar approximation to (4.87) here and assume that

$$\left[\frac{S}{R_m} \nabla \times \nabla \times - S \vec{a} \times \nabla \times - \nabla \nabla \cdot \right]^{-1} \nabla \times \nabla \times \approx \gamma \frac{R_m}{S} I, \quad (4.93)$$

where γ should be chosen as

$$\gamma_* := \frac{1}{1 + R_m h_{\mathbf{r}} \|\vec{a}\|}. \quad (4.94)$$

Because the approximation (4.93) is defined for vector arguments, we do not use the scalar mesh size $h_{\mathbf{r}}$ here. Furthermore, because $-\nabla \nabla \cdot$ arose originally from an arbitrary augmentation of the Maxwell system, we do not incorporate this term in the expression for γ_* .

Given the above analysis, we have the approximation

$$\tilde{\mathcal{K}} \approx \beta \gamma \frac{S}{R_m} \vec{b} \times (\cdot \times \vec{b}) \quad (4.95)$$

in both the two- and three-dimensional cases, with $\beta = 1$ in 2D. Then, an approximation for the velocity Schur complement X_N can be written as

$$\hat{X}_N := F + \tilde{F} + \beta \gamma K, \quad (4.96)$$

where K is defined as in (4.39) and (4.58).

Given that the Newton velocity Schur complement \hat{X}_N is of the same structure as the Picard velocity Schur complement \hat{X}_P , we use the same approximation \hat{Y}_P

for the pressure Schur complement in both cases. As in Section 3.2, we only need to adjust the parameter α_* to incorporate γ . This then defines two preconditioners for the Newton system which we write

$$\mathcal{P}_{N,Q} = \begin{pmatrix} A + \widehat{\tilde{A}} + kQ_{\mathbf{B}} & 0 & -Z^t & 0 \\ 0 & \frac{1}{k}\hat{L}_{\mathbf{r}} & 0 & 0 \\ 0 & 0 & \hat{X}_N & B^t \\ 0 & 0 & 0 & \hat{Y}_N \end{pmatrix}, \quad (4.97)$$

$$\mathcal{P}_{N,GD} = \begin{pmatrix} A + \widehat{\tilde{A}} + kD^t\bar{Q}_{\mathbf{r}}^{-1}D & 0 & -Z^t & 0 \\ 0 & \frac{1}{k}\bar{Q}_{\mathbf{r}} & 0 & 0 \\ 0 & 0 & \hat{X}_N & B^t \\ 0 & 0 & 0 & \hat{Y}_N \end{pmatrix}. \quad (4.98)$$

The operator \hat{X}_N in both cases is an approximation of $F + \tilde{F} + \beta\gamma K$, where $\beta \equiv 1$ in 2D, $\beta \in [0, 1]$ in 3D, and $\gamma \in [0, 1]$. The operator \hat{Y}_N is defined as

$$\hat{Y}_N := -\widehat{B\bar{Q}_{\mathbf{u}}^{-1}B^t} [B\bar{Q}_{\mathbf{u}}^{-1}(F + \tilde{F} + \alpha\beta\gamma K)\bar{Q}_{\mathbf{u}}^{-1}B^t]^{-1} \widehat{B\bar{Q}_{\mathbf{u}}^{-1}B^t}. \quad (4.99)$$

4.2.3 Computational Results

In this section, we evaluate the performance of the preconditioners for solving a set of benchmark problems. Our implementation is in the Trilinos framework developed at Sandia National Laboratories [38] using the Teko package to construct the block preconditioners and a GMRES Krylov solver from AztecOO [37]. For component solves, we use algebraic multigrid from the ML package [27], with incomplete factorization smoothers coming from IFPACK. Specifically, for solves on

the velocity spaces we use one V-cycle of AMG with two pre- and post-sweeps of a Schwarz domain decomposition smoother with ILU(0) on each subdomain and one level of overlap between subdomains. For solves on the pressure and magnetic pseudo-pressure spaces we use one V-cycle of AMG with five sweeps of a Gauss-Seidel smoother. We reiterate that, because of the large null space of A , the augmented magnetics blocks tend to be more challenging to approximate than the other operators appearing in our preconditioners. Thus, we use more computationally intense solvers for the operators associated with \mathbf{B} . For the mass matrix augmented blocks $A+kQ_{\mathbf{B}}$ and $A+\tilde{A}+kQ_{\mathbf{B}}$, we use an auxiliary-space algebraic multigrid as discussed in Section 4.1.4. Specifically, we use ten V-cycles of the multigrid solver detailed in [41] as it is implemented in ML. For the grad div augmented blocks $A+kD^t\bar{Q}_{\mathbf{r}}^{-1}D$ and $A+\tilde{A}+kD^t\bar{Q}_{\mathbf{r}}^{-1}D$, we use five V-cycles of AMG with two pre- and post-sweeps of a Schwarz domain decomposition smoother with ILU(2) on each subdomain and one level of overlap between subdomains. The number of V-cycles for each of these was tuned experimentally to balance the number of linear iterations and the cost of an application to minimize total computation time.

All problems were run on the Red Sky computer at Sandia National Laboratories. For all problems we use a stable Q_2 - Q_1 finite element pair for the velocity and pressure. We use first order Nédélec elements for the magnetic field and Q_1 elements for the Lagrange multiplier r . We use uniform grids for every problem. The relative residual tolerance is 10^{-5} for the nonlinear iteration and 10^{-6} for the GMRES inner iteration. We apply GMRES without any restarting. When considering the performance of the preconditioners, reported linear iteration counts are averaged over all

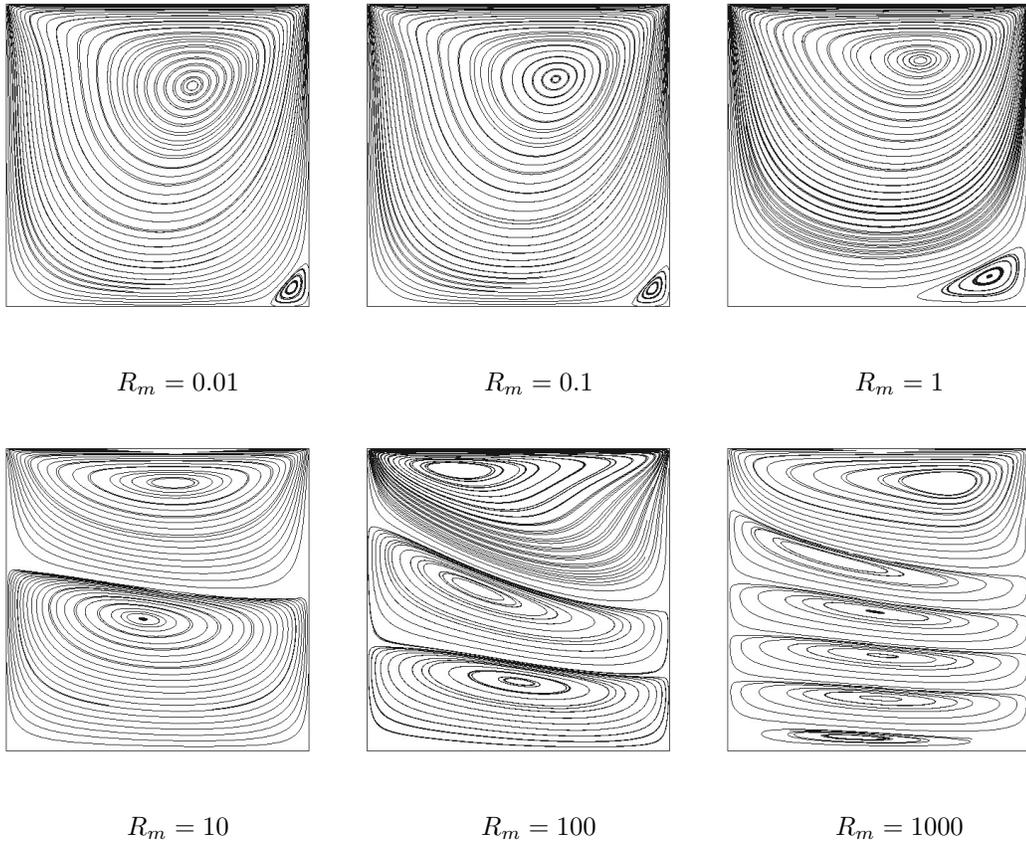


Figure 4.1: Velocity streamlines for the two-dimensional MHD lid driven cavity problem with $R = 100$, $S = 1$, and various values of R_m .

nonlinear iterations up to a maximum of fifty nonlinear iterations. We consider two- and three-dimensional versions of a lid driven cavity in the presence of a magnetic field (adapted from the fluid problem in [24]) as our test problems.

4.2.4 Two-Dimensional MHD Lid Driven Cavity

In this section, we explore the robustness of our preconditioners on the two-dimensional lid driven cavity as detailed in Section 3.3.1. We consider the square domain $[-\frac{1}{2}, \frac{1}{2}]^2$. We impose no-flow conditions on the bottom, left, and right sides

R_m	0.01	0.1	1	10	100	1000
Picard	4	4	6	8	27	> 50
Newton	3	4	4	4	7	22

Table 4.2: Number of nonlinear iterations needed to obtain convergence for the two-dimensional MHD lid driven cavity problem with $R = 100$ on a 100×100 element mesh.

of the domain and a horizontal velocity of 1 on the top. The tangential component of the magnetic field is specified on the boundary $\partial\Omega$ as

$$\vec{B} \times \vec{n} = (-1, 0) \times \vec{n}. \quad (4.100)$$

All results in this section were obtained on a fixed 100×100 element mesh, unless otherwise noted. Whereas we considered multiple values of the fluid Reynolds number R in Chapter 3, we consider only one value $R = 100$ here. We have observed that similar trends in preconditioning results can be seen for different values of R . Consistent with the results of Chapter 3, larger Reynolds numbers tend to lead to larger iteration counts, but with approximately the same dependence on the magnetic Reynolds number R_m . Furthermore, we keep the coupling coefficient S constant at $S = 1$. The degree of coupling in the problem can then be controlled by varying R_m , as illustrated in Figure 4.1. This figure shows that for small values of R_m , the velocity solution is approximately the same as what would be obtained for the pure fluid lid driven cavity problem. As R_m increases, the coupling between the hydrodynamics and electromagnetics becomes stronger and the solution devi-

ates more from the pure fluid case. The number of nonlinear iterations needed to obtain these results are reported for both a Picard iteration and Newton's method in Table 4.2.

4.2.4.1 Picard Iteration

In this section, we consider the performance of the two preconditioners $\mathcal{P}_{P,Q}$ and $\mathcal{P}_{P,GD}$ on the linear systems arising in a Picard iteration. Because the choice of the parameter α appearing in the pressure Schur complement approximation (4.70) was analyzed at length for the exact penalty formulation in Chapter 3, we consider here only the case where $\alpha = \alpha_*$, as defined in (4.67). In our experience with this benchmark problem, consistent with the findings of Chapter 3, the choice of $\alpha = \alpha_*$ tends to result in better iteration counts than $\alpha = 1$.

To assess the dependence of the two preconditioners on the augmentation parameter k , we report iteration counts with several values of k in Table 4.3. For each preconditioner, we consider a small, medium, and large value of the magnetic Reynolds number (i.e. $R_m = 0.01, 1, 100$). From the table, it is clear that the performance of $\mathcal{P}_{P,Q}$ depends strongly on the parameter k . As k increases, the number of iterations needed for convergence also increases. This is consistent with the eigenvalue bound (4.19) for $\mathcal{M}_P \hat{\mathcal{M}}_{P,Q}^{-1}$, which shows that as k increases, the lower bound on the eigenvalues of the preconditioned system approaches zero. Furthermore, this effect becomes more pronounced as R_m increases. Some increase in iteration counts can be seen for $\mathcal{P}_{P,Q}$ when k is small. This seems to be due to the fact that when k

		$\mathcal{P}_{P,Q}$			$\mathcal{P}_{P,GD}$		
k	R_m	0.01	1	100	0.01	1	100
	0.01		48	36	40	×	×
0.1		37	29	36	×	31	46
1		35	28	62	×	31	×
10		36	36	152	39	40	×
100		46	71	413	39	×	×
1000		57	133	944	49	×	×

Table 4.3: Average GMRES iterations required for convergence with $\mathcal{P}_{P,Q}$ and $\mathcal{P}_{P,GD}$ on the Picard linearization of the two-dimensional MHD lid driven cavity problem with various values of k . ‘×’ indicates no convergence due to breakdown of the preconditioner. Bold entries correspond to the choice $k = \frac{S}{R_m}$.

is small, the augmented operator $A+kQ_{\mathbf{B}}$ is more dominated by the singular matrix A and thus more difficult for the multigrid solver. While using $\mathcal{P}_{P,Q}$ with $k = \frac{S}{R_m}$ does not always give the best results (e.g. $k = 0.1$ converges faster than $k = 0.01$ for $R_m = 100$), this choice does tend to give close to optimal iteration counts. Thus, while fine tuning k may result in a slight improvement in performance, the choice of $k = \frac{S}{R_m}$ results in a convenient and effective preconditioner.

For the preconditioner $\mathcal{P}_{P,GD}$, the choice of $k = \frac{S}{R_m}$ is always best among the values in Table 4.3. Values of k close to $\frac{S}{R_m}$ also tend to give good results, but when k is not close to $\frac{S}{R_m}$ the preconditioner breaks down. This appears to be because when k is not close to $\frac{S}{R_m}$, the two terms in the residual r_{GD} , defined in (4.72), are not balanced. Thus, one component of the error, either e_A or e_D , is not represented in the residual, and the multigrid solver we use for this block no longer works.

Based on these observations, we set $k = \frac{S}{R_m}$ in comparing $\mathcal{P}_{P,Q}$ and $\mathcal{P}_{P,GD}$. The two preconditioners are compared for a range of R_m in Table 4.4. In terms of iteration counts, $\mathcal{P}_{P,Q}$ is a slightly better preconditioner for moderate R_m , while $\mathcal{P}_{P,GD}$ is better for very small or very large R_m , but in general both preconditioners are comparable. The most significant difference is at $R_m = 1000$ where $\mathcal{P}_{P,GD}$ requires 64% of the iterations that $\mathcal{P}_{P,Q}$ does. In terms of computation time, the application $\mathcal{P}_{P,Q}$ is always cheaper than $\mathcal{P}_{P,GD}$ on this mesh.

To compare the preconditioners in terms of parallel scalability, we perform a weak scaling study for the case of $R_m = 100$. That is, we consider 50×50 , 100×100 , 200×200 , 400×400 , and 800×800 element discretizations on 1, 4, 16, 64, and 256 processors respectively, keeping the number of unknowns per processor

	$\mathcal{P}_{P,Q}$		$\mathcal{P}_{P,GD}$	
R_m	Iterations	Time (s)	Iterations	Time (s)
0.01	46	11	39	13
0.1	36	8	38	12
1	28	6	31	11
10	25	6	28	10
100	40	9	34	13
1000	85	21	54	23

Table 4.4: Average GMRES iterations and time in seconds required for convergence with $\mathcal{P}_{P,Q}$ and $\mathcal{P}_{P,GD}$ on the Picard linearization of the two-dimensional MHD lid driven cavity problem with various values of R_m .

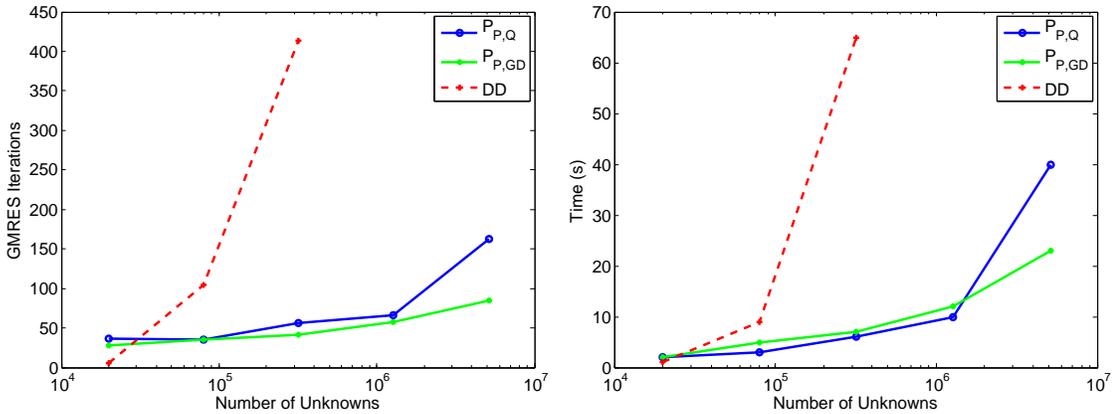


Figure 4.2: Weak parallel scaling results for the two-dimensional MHD lid driven cavity problem with Picard linearization. Average GMRES iterations on the left and average linear solve time per nonlinear iteration on the right.

at approximately 20,000. We report average iteration counts and computation times per nonlinear step for these experiments in Figure 4.2. The results are compared against a pure domain decomposition preconditioner using a SuperLU [21] ILUTP factorization with a drop tolerance of 0.0001 and a zero pivot threshold of 0.01 (labeled DD in the scaling plots). From the plots, it can be seen that the grad div augmented preconditioner $\mathcal{P}_{P,GD}$ scales fairly well, with only slight growth in iteration count and computation time as the mesh is refined. The mass matrix augmented preconditioner $\mathcal{P}_{P,Q}$ does not perform quite as well, but its performance is superior to the domain decomposition preconditioner.

4.2.4.2 Newton's Method

In this section, we consider the performance of the preconditioners $\mathcal{P}_{N,Q}$ and $\mathcal{P}_{N,GD}$ on the linear systems arising in Newton's method. Again, because both γ and α were studied in Chapter 3, we set $\gamma = \gamma_*$ and $\alpha = \alpha_*$ as defined in (4.88) and (4.67).

We investigate the dependence of the two preconditioners on the parameter k in Table 4.5. As with the Picard iteration, the grad div augmented preconditioner $\mathcal{P}_{N,GD}$ achieves its lowest iteration counts with $k = \frac{S}{R_m}$ and breaks down if k is not close to $\frac{S}{R_m}$. For the mass matrix augmented preconditioner $\mathcal{P}_{N,Q}$, we again see a degradation in performance as k increases, but we can now also see breakdown in the preconditioner when k is too small. This is especially problematic in the case of $R_m = 100$ since the preconditioner breaks down for the ideal choice $k = 0.01$.

		$\mathcal{P}_{N,Q}$			$\mathcal{P}_{N,GD}$		
k	R_m	0.01	1	100	0.01	1	100
	0.01		49	×	×	×	×
0.1		38	32	×	×	33	53
1		36	30	×	×	33	×
10		38	35	138	40	42	×
100		47	58	359	39	×	×
1000		58	105	873	50	×	×

Table 4.5: Average GMRES iterations required for convergence with $\mathcal{P}_{N,Q}$ and $\mathcal{P}_{N,GD}$ on the Newton linearization of the two-dimensional MHD lid driven cavity problem with various values of k . ‘×’ indicates no convergence due to breakdown of the preconditioner. Bold entries correspond to the choice $k = \frac{S}{R_m}$.

This effect seems to be because we are using a solver designed for $A + kQ_{\mathbf{B}}$ to approximate the action of the inverse of $A + \tilde{A} + kQ_{\mathbf{B}}$. The multigrid algorithm designed in [41] creates a coarse-grid hierarchy that preserves the null space of A , but because \tilde{A} does not have the same null space, this methodology does not extend to $A + \tilde{A} + kQ_{\mathbf{B}}$ if \tilde{A} dominates. It appears that the solver works when k is large, because in that case the nonsingular operator $kQ_{\mathbf{B}}$ is non-negligible. However, as noted in Section 4.1.1.1, when $kQ_{\mathbf{B}}$ is not negligible compared to A , the lower bound in (4.19) must be $O(\alpha h^2)$, meaning that the eigenvalues of $\mathcal{M}_N \hat{\mathcal{M}}_{N,Q}^{-1}$ must depend on the mesh. Thus, in order to use preconditioner $\mathcal{P}_{N,Q}$ for systems with large R_m , the parameter k must be very large and depend on the mesh. This makes $\mathcal{P}_{N,Q}$ an impractical preconditioner that does not scale if an auxiliary-space Maxwell solver is used to approximate the action of the inverse of $A + \tilde{A} + kQ_{\mathbf{B}}$. We are unaware of solvers designed for matrices of the form $A + \tilde{A} + kQ_{\mathbf{B}}$, and the arguments of Section 4.1.4 imply that traditional multigrid will not work for this matrix. Hence, the construction of a practical version of $\mathcal{P}_{N,Q}$ is beyond the scope of this study.

With the choice $k = \frac{S}{R_m}$, we compare the two preconditioners in Table 4.6. We again see that $\mathcal{P}_{N,Q}$ cannot be used for problems with large magnetic Reynolds number. The grad div augmented preconditioner, on the other hand, performs fairly well with comparable iteration counts to the Picard preconditioner $\mathcal{P}_{P,GD}$ up to $R_m = 100$. In general, $\mathcal{P}_{N,GD}$ performs very well regardless of R_m . The result for $R_m = 1000$ is still favorable considering that this is a very strongly coupled problem and a Picard iteration does not converge within 50 steps on this problem. Weak parallel scaling results are plotted for $\mathcal{P}_{N,GD}$ on the $R_m = 100$ problem in Figure 4.3,

	$\mathcal{P}_{P,Q}$		$\mathcal{P}_{P,GD}$	
R_m	Iterations	Time (s)	Iterations	Time (s)
0.01	48	12	37	13
0.1	39	9	34	13
1	29	7	32	12
10	×	×	29	10
100	×	×	43	15
1000	×	×	104	41

Table 4.6: Average GMRES iterations and time in seconds required for convergence with $\mathcal{P}_{N,Q}$ and $\mathcal{P}_{N,GD}$ on the Newton linearization of the two-dimensional MHD lid driven cavity problem with various values of R_m .

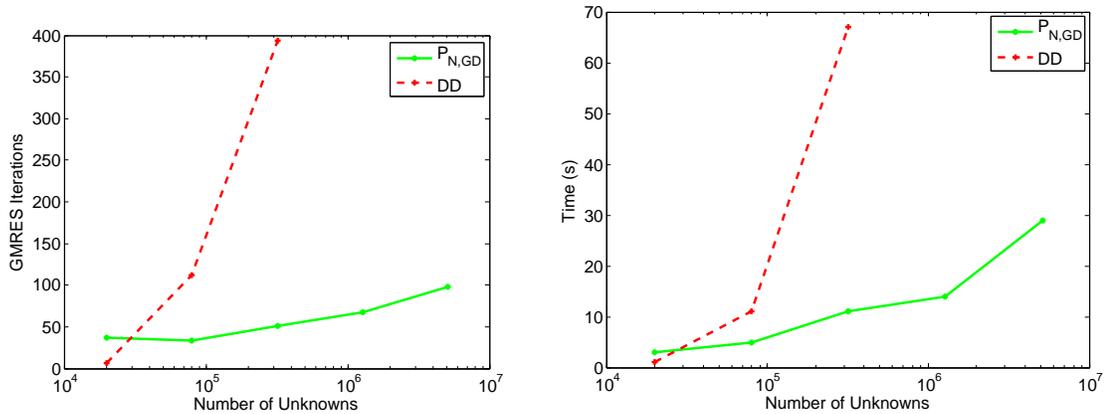


Figure 4.3: Weak parallel scaling results for the two-dimensional MHD lid driven cavity problem with Newton's method. Average GMRES iterations on the left and average linear solve time per nonlinear iteration on the right.

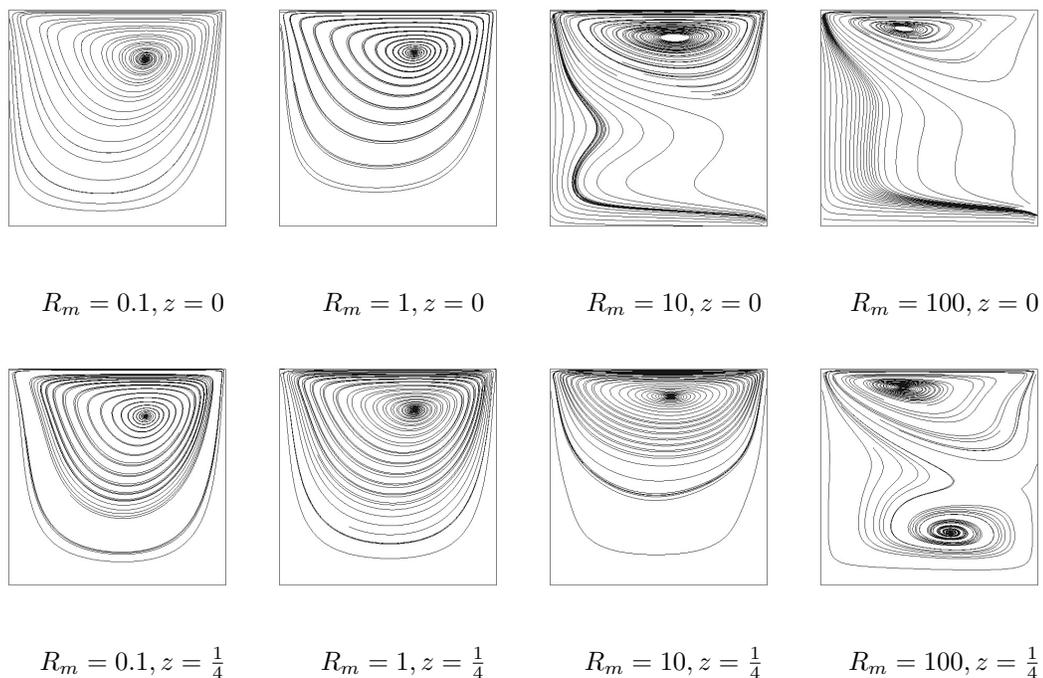


Figure 4.4: Velocity streamlines at cross-sections $z = 0$ and $z = \frac{1}{4}$ for the three-dimensional MHD lid driven cavity problem with $R = 100, S = 1$, and various values of R_m .

where one can see similar scaling behavior as was demonstrated by $\mathcal{P}_{P,GD}$ on the Picard system.

4.2.5 Three-Dimensional MHD Lid Driven Cavity

In this section, we investigate the performance of our preconditioners on a three-dimensional lid driven cavity problem. We consider the domain $[-\frac{1}{2}, \frac{1}{2}]^3$ with no-flow conditions imposed on the bottom, front, back, left, and right sides of the domain and a horizontal velocity of 1 on the top. An external magnetic field is

R_m	0.01	0.1	1	10	100
Picard	2	2	3	5	16
Newton	2	2	3	3	5

Table 4.7: Number of nonlinear iterations needed to obtain convergence for the three-dimensional MHD lid driven cavity problem with $R = 100$ on a $40 \times 40 \times 40$ element mesh.

applied by imposing

$$\vec{B} \times \vec{n} = (-1, 0) \times \vec{n} \quad (4.101)$$

on the boundary. We consider the case where $R = 100$ and $S = 1$ with various values of R_m to control the degree of physical coupling. All results in this section were obtained on a fixed $40 \times 40 \times 40$ element mesh, unless otherwise noted. Some solutions to this problem are plotted in Figure 4.4. The number of nonlinear iterations needed to obtain these results are reported for both a Picard iteration and Newton’s method in Table 4.7. Similar trends to the two-dimensional case can be seen in Figure 4.4. For low magnetic Reynolds number, the solution is dominated by one large vortex, and as R_m increases, the vortex is pushed upward in the domain. The character of streamlines at the two cross-sections $z = 0$ and $z = \frac{1}{4}$ is very similar for small R_m , but as R_m increases, three-dimensional effects become more prominent, and the streamlines at the two cross-sections become less similar. We do not consider R_m larger than 100 in the three-dimensional setting because this requires more resolution than we have the computational resources for in this study.

		$\mathcal{P}_{P,Q}$			$\mathcal{P}_{P,GD}$		
R_m	β	0	1	β_*	0	1	β_*
	0.01		31	31	31	25	25
0.1		28	28	28	23	23	23
1		25	22	24	25	24	25
10		61	43	44	54	35	39
100		282	152	99	243	81	61

Table 4.8: Average GMRES iterations required for convergence with $\mathcal{P}_{P,Q}$ and $\mathcal{P}_{P,GD}$ on the Picard linearization of the three-dimensional MHD lid driven cavity problem with various values of β .

4.2.5.1 Picard Iteration

In this section, we consider the performance of the two preconditioners $\mathcal{P}_{P,Q}$ and $\mathcal{P}_{P,GD}$ on the linear systems arising in a Picard iteration. We are interested primarily in the parameter β here. Letting $\alpha = \alpha_*$ and $t = \frac{S}{R_m}$, we investigate

R_m	0.01	0.1	1	10	100
β_*	0.0001	0.0012	0.0123	0.1111	0.5556

Table 4.9: Values of β_* at the second step of a Picard iteration.

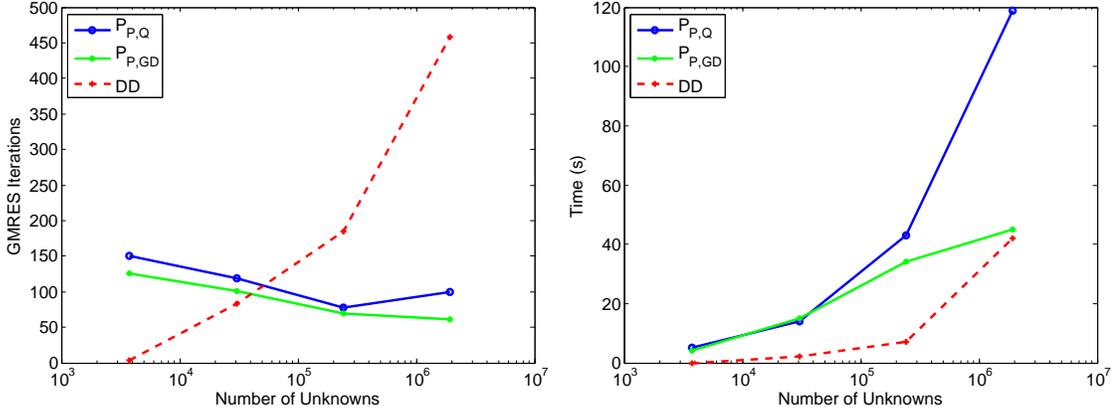


Figure 4.5: Weak parallel scaling results for the three-dimensional MHD lid driven cavity problem with Picard linearization. Average GMRES iterations on the left and average linear solve time per nonlinear iteration on the right.

the performance of $\mathcal{P}_{P,Q}$ and $\mathcal{P}_{P,GD}$ for three different values of β in Table 4.8. Motivated by the analysis of Section 4.1.2.2, we use $\beta = 0$, which is the best choice if $\nabla \times (\vec{u} \times \vec{b}) = 0$; $\beta = 1$, which is the best choice if $\nabla \cdot (\vec{u} \times \vec{b}) = 0$; and $\beta = \beta_*$ as defined in (4.56), which we derived to make the preconditioner robust. The particular values used for β_* at the second nonlinear step are reported in Table 4.9. From Table 4.8, it is clear that $\beta = 0$ never leads to better convergence than either $\beta = 1$ or $\beta = \beta_*$. The choice $\beta = 1$ yields slightly better performance than $\beta = \beta_*$ for $R_m = 1$ and $R_m = 10$, but for the strongly coupled case of $R_m = 100$, $\beta = \beta_*$ yields a dramatic improvement over $\beta = 1$, with a 35% reduction in iterations for $\mathcal{P}_{P,Q}$ and a 25% reduction for $\mathcal{P}_{P,GD}$. Because β_* gives comparable results to $\beta = 1$ for moderate R_m and improves performance for difficult problems with large R_m , we consider β_* the best of the three values of β considered.

In Figure 4.5, we investigate the weak parallel scalability of the two precon-

conditioners with $\beta = \beta_*$ for the case of $R_m = 100$. We compute with $5 \times 5 \times 5$, $10 \times 10 \times 10$, $20 \times 20 \times 20$, and $40 \times 40 \times 40$ finite element discretizations on 1, 8, 64, and 512 processors, respectively. This keeps approximately 4,000 degrees of freedom on each processor. The results are again compared against a pure domain decomposition preconditioner using a SuperLU ILUTP factorization. In terms of iteration count, both $\mathcal{P}_{P,Q}$ and $\mathcal{P}_{P,GD}$ and scale very well, with iteration counts decreasing as the number of processors increases. In terms of computation time, we see growth for both preconditioners with $\mathcal{P}_{P,GD}$ performing better. Because the iteration counts scale well, we expect that better scaling can be achieved for the computation time by fine-tuning the component solves. While the domain decomposition preconditioner is competitive in terms of computation time, its poor scaling indicates that $\mathcal{P}_{P,Q}$ and $\mathcal{P}_{P,GD}$ will be more advantageous for more refined problems.

4.2.5.2 Newton's Method

In this section, we consider the performance of the preconditioners $\mathcal{P}_{N,Q}$ and $\mathcal{P}_{N,GD}$ on the linear systems arising in Newton's method. Table 4.10 shows iteration counts for the two preconditioners with the three values $\beta = 0, 1$, and β_* . The results are consistent with the results for the two-dimensional Newton problem in that $\mathcal{P}_{N,Q}$ breaks down for large R_m while $\mathcal{P}_{N,GD}$ performs well for all values of R_m considered. Table 4.10 also reiterates what we found for the choice of β in the three-dimensional Picard problem; that is, both $\beta = 1$ and $\beta = \beta_*$ tend to be good choices, but $\beta = \beta_*$ leads to improved performance for large R_m . Weak scaling results are

		$\mathcal{P}_{N,Q}$			$\mathcal{P}_{N,GD}$		
$R_m \backslash \beta$		0	1	β_*	0	1	β_*
		0.01	31	31	31	25	25
0.1	23	23	23	23	23	23	
1	26	23	26	27	25	27	
10	×	×	×	50	34	37	
100	×	×	×	298	83	72	

Table 4.10: Average GMRES iterations required for convergence with $\mathcal{P}_{N,Q}$ and $\mathcal{P}_{N,GD}$ on the Newton linearization of the three-dimensional MHD lid driven cavity problem with various values of β .

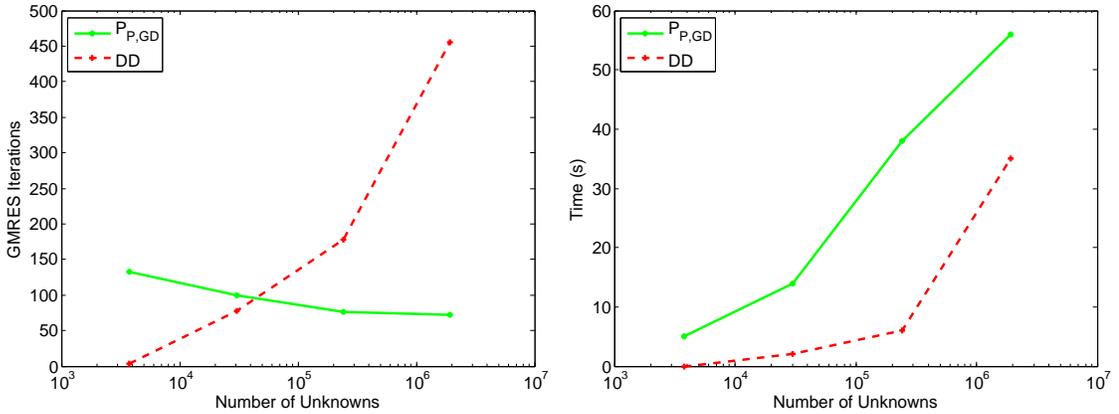


Figure 4.6: Weak parallel scaling results for the three-dimensional MHD lid driven cavity problem with Newton's method. Average GMRES iterations on the left and average linear solve time per nonlinear iteration on the right.

plotted in Figure 4.6 for $\mathcal{P}_{N,GD}$ with $\beta = \beta_*$ and $R_m = 100$. This figure again shows that iteration counts decrease as the mesh is refined. While the computation time required by $\mathcal{P}_{N,GD}$ is greater than that required by a domain decomposition preconditioner at this scale, it demonstrates better scaling and can be improved by the tuning component solvers.

4.3 Conclusion

In this chapter, we have presented a series of block preconditioners for the linear systems arising from both Picard and Newton linearizations of a Lagrange multiplier formulation of the stationary MHD equations. We employed two preconditioners for Maxwell's equations in mixed form to handle the electromagnetics saddle point matrix associated with the unknowns \vec{B} and r . The first preconditioner augments the singular curl-curl matrix A with a scaled mass matrix $kQ_{\mathbf{B}}$. We used an existing auxiliary-space multigrid method as a solver for the operator $A + kQ_{\mathbf{B}}$. The second Maxwell preconditioner augments A with the augmented Lagrangian operator $kD^t\bar{Q}_{\mathbf{r}}^{-1}D$ and standard AMG can be used for the augmented operator. We experimentally investigated the choice of the scaling parameter k and concluded that the choice of $k = \frac{S}{R_m}$ is convenient and effective for both preconditioners applied to the Picard system. When applied to the Newton system, mass matrix augmentation requires that k depend on the mesh size h , and $\mathcal{P}_{N,Q}$ becomes an impractical preconditioner. $\mathcal{P}_{N,GD}$, on the other hand, performs well on linear systems arising from both Picard and Newton linearizations.

We extended the analysis of Chapter 3 to develop approximations to the Schur complements on the velocity and pressure spaces, X and Y . In the two-dimensional setting, we could use exactly the same approximations as we did for the exact penalty formulation. In the three-dimensional setting, the dimension-dependent character of the curl operator forced us to make an additional approximation within the velocity Schur complement. We parameterized this approximation with the new parameter β and provided a formula for a choice β_* based on whether the induction equation is convection- or diffusion-dominated. Our experimental results demonstrated that β_* is an effective choice for β , especially for very large magnetic Reynolds number. Our results showed that both preconditioners are robust with respect to mesh refinement and the magnetic Reynolds number for the Picard system in both two and three dimensions. The grad div augmented preconditioner $\mathcal{P}_{N,GD}$ is scalable and robust for two- and three-dimensional Newton systems. The grad div augmented preconditioners are appealing because they rely only on standard multigrid solvers (we used ILU and Gauss-Seidel smoothers for this study). The mass matrix augmented preconditioners required a non-standard coarsening strategy for the block associated with \vec{B} . The version proposed for the Newton system may become practical if a solver is developed for operators of the type $A + \tilde{A} + kQ_{\mathbf{B}}$.

Chapter 5: A Stochastic Approach to Uncertainty in the Equations of MHD Kinematics

In this chapter, we consider the kinematics equations, which govern the influence of the fluid flow on the magnetic field. These equations constitute a component block within fully coupled MHD simulations such as the Lagrange multiplier formulation discussed in Section 2.3 and Chapter 4. Furthermore, solution of these equations is also required in operator splitting techniques that alternate between solving the Navier-Stokes equations and the kinematics equations [57]. The kinematics equations are also of particular interest in the field of kinematic dynamo theory, in which the ratio of the Lorentz force to inertia is assumed to be small [45]. In this case, the velocity can be prescribed, and the generation of the magnetic energy induced by the flow can be studied. Kinematic simulations can be used to model MHD generators, in which plasmas act as conductors to generate electric currents, as well as natural dynamos such as the sun and the geodynamo. They are of primary interest in investigating whether a given flow profile can sustain dynamo action.

When the velocity field is prescribed, this simplifies the MHD equations, but it may also introduce some epistemic uncertainty into the model. The flow prop-

erties of the fluid may not be known on the interior of the domain. Additionally, there are aspects of the physical model that motivate incorporation of small-scale uncertainty. For instance, the large-scale mean flow of the earth's outer core cannot account for the magnitude of the earth's magnetic field. In geodynamo theory, it is proposed that small-scale turbulent behavior can give rise to a large-scale magnetic field through the α -effect [18]. Furthermore, the distribution of material properties may be uncertain in physical applications. When multiple fluids are present, such as when multiple liquid metals are mixing together, the magnetic resistivity will not be homogeneous throughout the domain and may vary over orders of magnitude. Because the resistivity can have a strong influence on such physical systems, including changing the topology of the magnetic field, we are interested in how uncertain heterogeneous distributions of the resistivity may affect the induced magnetic energy.

In this chapter, we explore these issues by mathematically simulating uncertainty in both the velocity field and the resistivity within the MHD kinematics model. In this model, we treat the uncertain quantities as random fields correlated in space. We will obtain mean and variance data through Monte-Carlo simulation. In addition, because each Monte-Carlo trial requires the solution of linear systems with randomly varying dynamics, we develop and explore efficient and robust solvers for discrete kinematics systems.

Thus, we consider two issues: the impact of uncertainty of velocity and resistivity on statistical properties of the magnetic fields modeled by the equations of MHD kinematics, together with efficient computational algorithms for computing

these quantities. The remainder of the chapter is structured as follows. In Section 5.1, we will derive a finite element formulation for the deterministic equations of MHD kinematics. Section 5.2 is devoted to the incorporation of uncertainty into the model. In this section, we describe a means of modeling the uncertainty in both the resistivity and the velocity field and apply the model to representative test problems. In Section 5.3, we propose, analyze, and test a block preconditioner for solving the linear systems arising in our model. Finally, we will draw conclusions in Section 5.4.

5.1 A Finite Element Formulation

The steady-state kinematics of MHD are governed by Maxwell's equations

$$\nabla \times \left(\frac{1}{\mu} \vec{B} \right) = \vec{j}, \quad (5.1a)$$

$$\nabla \cdot (\varepsilon \vec{E}) = \rho_c, \quad (5.1b)$$

$$\nabla \times \vec{E} = \vec{0}, \quad (5.1c)$$

$$\nabla \cdot \vec{B} = 0, \quad (5.1d)$$

and Ohm's law

$$\vec{j} = \sigma(\vec{E} + \vec{u} \times \vec{B}), \quad (5.2)$$

on a domain $\mathcal{D} \subset \mathbb{R}^d, d = 2$ or 3 (plus appropriate boundary conditions). The unknowns here are the magnetic induction \vec{B} , the electric field \vec{E} , and the current density \vec{j} ; the charge density ρ_c can be regarded as an auxiliary variable obtained after computing \vec{E} . We will regard the fluid velocity \vec{u} as given. For many appli-

cations, the electric permittivity ε and the magnetic permeability μ do not vary significantly for different fluids [29], so we let ε and μ be fixed scalar constants over the whole domain. However, because the heterogeneities of the electric conductivity σ can be large for different liquid metals, we consider σ to be a prescribed, not necessarily constant function on \mathcal{D} .

We consider the boundary conditions

$$\vec{B} \times \vec{n} = \vec{q}, \quad (5.3a)$$

$$\vec{E} \cdot \vec{n} = k, \quad (5.3b)$$

on $\partial\mathcal{D}$. We choose these conditions over the alternative (prescribing $\vec{B} \cdot \vec{n}$ and $\vec{E} \times \vec{n}$) because the requirement on the tangential component of \vec{B} is then the natural Dirichlet condition for the curl-conforming edge elements employed to discretize \vec{B} .

A standard simplification of equations (5.1) and (5.2) is obtained by eliminating the variables \vec{j} and \vec{E} , yielding the following equations for the kinematics of MHD in terms of \vec{B} :

$$\nabla \times \left(\frac{\eta}{\mu} \nabla \times \vec{B} \right) - \nabla \times (\vec{u} \times \vec{B}) = \vec{0}, \quad (5.4a)$$

$$\nabla \cdot \vec{B} = 0, \quad (5.4b)$$

on \mathcal{D} , where $\eta = 1/\sigma$ is the magnetic resistivity. A boundary condition such as (5.3a) is required to complete this system. After \vec{B} is obtained from solving equations (5.4), \vec{E} , \vec{j} , and ρ_c can be recovered.

As stated, the equations (5.4) are over-determined because there are $d + 1$ equations in d unknowns. In order to make the system well-defined without changing

the solution \vec{B} , we introduce a Lagrange multiplier r (we refer to this variable as the magnetic pseudo-pressure), and consider the equations

$$\nabla \times \left(\frac{\eta}{\mu} \nabla \times \vec{B} \right) - \nabla \times (\vec{u} \times \vec{B}) + \nabla r = \vec{0}, \quad (5.5a)$$

$$\nabla \cdot \vec{B} = 0, \quad (5.5b)$$

with the boundary conditions (5.3a) and $r = 0$ on $\partial\mathcal{D}$. It can be shown that (5.5) admits the same solution \vec{B} as (5.4) by taking the divergence of equation (5.5a). This yields $\Delta r = 0$ on \mathcal{D} , which, with the zero Dirichlet condition on r , implies that $r = 0$ on \mathcal{D} .

For developing a weak formulation for this problem, we consider the spaces

$$V_0 = \{\vec{C} \in H(\text{curl}, \mathcal{D}) \mid \vec{C} \times \vec{n} = \vec{0} \text{ on } \partial\mathcal{D}\}, \quad (5.6a)$$

$$V_{\vec{q}} = \{\vec{C} \in H(\text{curl}, \mathcal{D}) \mid \vec{C} \times \vec{n} = \vec{q} \text{ on } \partial\mathcal{D}\}, \quad (5.6b)$$

$$Q_0 = H_0^1(\mathcal{D}). \quad (5.6c)$$

Multiplying the equations (5.5) by test functions $\vec{C} \in V_0$ and $s \in Q_0$, and integrating by parts, we obtain the following weak formulation: Find $(\vec{B}, r) \in V_{\vec{q}} \times Q_0$ such that

$$a(\vec{B}, \vec{C}) + c(\vec{B}, \vec{C}) + b(\vec{C}, r) = 0, \quad (5.7a)$$

$$b(\vec{B}, s) = 0, \quad (5.7b)$$

for all $(\vec{C}, s) \in V_0 \times Q_0$, where the bilinear forms are defined as

$$a(\vec{B}, \vec{C}) = \left(\frac{\eta}{\mu} \nabla \times \vec{B}, \nabla \times \vec{C} \right), \quad (5.8a)$$

$$c(\vec{B}, \vec{C}) = - \left(\vec{u} \times \vec{B}, \nabla \times \vec{C} \right), \quad (5.8b)$$

$$b(\vec{B}, s) = (\vec{B}, \nabla s). \quad (5.8c)$$

By integrating by parts the divergence of the magnetic induction \vec{B} , we limit the required regularity of \vec{B} . This allows the model to include magnetic fields \vec{B} with strong singularities arising from re-entrant corners, which is not feasible for $\vec{B} \in H^1(\Omega)^d$ [16].

We discretize the domain into a shape-regular partition \mathcal{T}_h of quadrilaterals or hexahedra $\{K\}$. Letting $\mathcal{P}_\ell(K)$ be the space of polynomials of degree ℓ on K and $\mathcal{N}_\ell(K)$ the space of Nédélec vector polynomials of the first kind [50] (with $\mathcal{P}_{\ell-1}(K)^d \subset \mathcal{N}_\ell(K) \subset \mathcal{P}_\ell(K)^d$), we consider the finite dimensional spaces

$$V_0^h = \{\vec{C}_h \in V_0 \mid \vec{C}_h|_K \in \mathcal{N}_\ell(K), K \in \mathcal{T}_h\}, \quad (5.9a)$$

$$V_{\vec{q}}^h = \{\vec{C}_h \in V_{\vec{q}} \mid \vec{C}_h|_K \in \mathcal{N}_\ell(K), K \in \mathcal{T}_h\}, \quad (5.9b)$$

$$Q_0^h = \{s_h \in Q_0 \mid s_h|_K \in \mathcal{P}_\ell(K), K \in \mathcal{T}_h\}. \quad (5.9c)$$

Then the discrete formulation is as follows: Find $(\vec{B}_h, r_h) \in V_{\vec{q}}^h \times Q_0^h$ such that

$$a(\vec{B}_h, \vec{C}_h) + c(\vec{B}_h, \vec{C}_h) + b(\vec{C}_h, r_h) = 0, \quad (5.10a)$$

$$b(\vec{B}_h, s_h) = 0, \quad (5.10b)$$

for all $(\vec{C}_h, s_h) \in V_0^h \times Q_0^h$.

Let \mathbf{B} be the vector containing the coefficients of \vec{B}_h with respect to a basis for V_0^h , and let \mathbf{r} be the vector containing the coefficients of r_h with respect to a

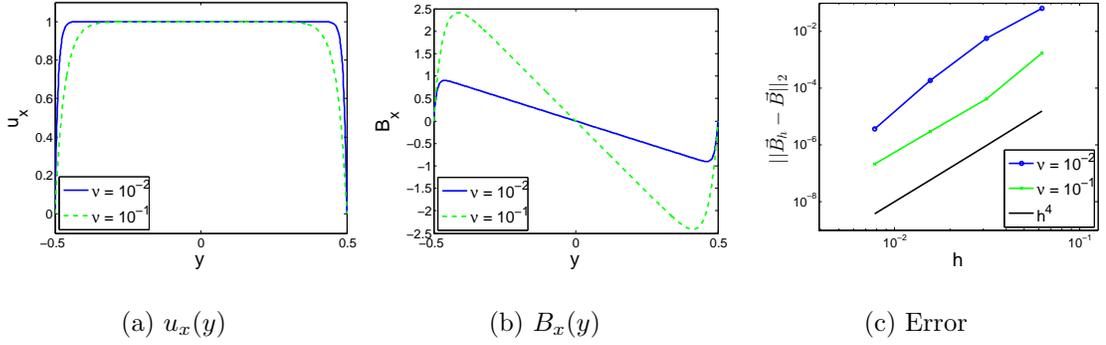


Figure 5.1: Velocity profiles (a), induced magnetic fields (b), and error $\|\vec{B}_h - \vec{B}\|_2$ (c) for the Hartmann problem.

basis for Q_0^h . Then, the finite element solution of the weak formulation (5.10) can be computed by solving a linear system of saddle point structure,

$$\begin{pmatrix} A + N & D^t \\ D & 0 \end{pmatrix} \begin{pmatrix} \mathbf{B} \\ \mathbf{r} \end{pmatrix} = \begin{pmatrix} f \\ 0 \end{pmatrix}. \quad (5.11)$$

In this equation, f includes boundary data, A is a discretization of the magnetic diffusion operator $\nabla \times (\frac{\eta}{\mu} \nabla \times \cdot)$, N is a discretization of the magnetic convection operator $\nabla \times (\vec{u} \times \cdot)$, D is the discrete (negative) divergence operator, and D^t is the discrete gradient.

We now demonstrate some physical aspects of these equations by applying them to two deterministic example problems. All simulations throughout the chapter are implemented using the deal.II finite element library [4] with first order Nédélec elements for \mathbf{B} and bilinear elements for \mathbf{r} (i.e. $\ell = 1$ in (5.9)).

5.1.1 Example Problem: Hartmann Flow

The Hartmann problem [18] is a classic two-dimensional test problem modeling the flow of an electrically conducting fluid through a channel in the presence of an externally applied transverse magnetic field. We pose this problem on the domain $[-0.5, 0.5]^2$ in the presence of the external magnetic field $\vec{B} = (0, 1)$. The coupled MHD equations admit the exact analytic solution

$$\vec{u} = \left(\frac{\cosh(H/2) - \cosh(Hy)}{\cosh(H/2) - 1}, 0 \right), \quad (5.12a)$$

$$\vec{B} = \left(H\nu \frac{\sinh(Hy) - 2 \sinh(H/2)y}{\cosh(H/2) - 1}, 1 \right), \quad (5.12b)$$

for this problem, where $H = \sqrt{\frac{1}{\nu\eta}}$ is the Hartmann number and ν is the kinematic viscosity of the fluid. Representative images of the components of \vec{u} and \vec{B} in the x -direction are plotted in Figure 5.1 for $\eta = 10^{-2}$ and $\mu = 1$ with $\nu = 10^{-1}$ and 10^{-2} . The plots show that smaller viscosity leads to thinner boundary layers in both \vec{u} and \vec{B} and that the magnitude of the induced magnetic field increases with the viscosity. This simple example demonstrates that fairly small changes in the velocity field can lead to large changes in the magnetic field.

We pose a kinematic version of the Hartmann problem by prescribing the velocity defined by (5.12a) over the domain \mathcal{D} and imposing $\vec{B} \times \vec{n} = (0, 1) \times \vec{n}$ on the boundary $\partial\mathcal{D}$. The exact solution \vec{B} to this kinematic problem is then given by (5.12b). Applying the finite element formulation (5.10), we obtain the approximation \vec{B}_h to \vec{B} . To validate the deterministic finite element formulation, we plot the convergence of the error $\|\vec{B}_h - \vec{B}\|_2$ as h is refined in Figure 5.1.

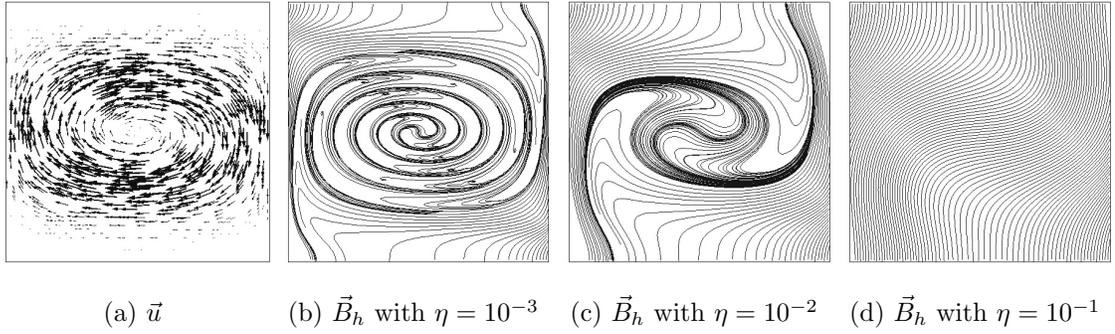


Figure 5.2: Velocity profile and induced magnetic fields for deterministic MHD eddy problem.

5.1.2 Example Problem: MHD Eddy

The physical effects of the resistivity η are demonstrated by a two-dimensional benchmark problem considered in [45], which models the effect of an eddy on a magnetic field. We prescribe the velocity field

$$\vec{u}(x, y) = \begin{pmatrix} \frac{\cos(\pi x)}{\pi} 32y(1 - 4y^2)^3 \\ -\sin(\pi x)(1 - 4y^2)^4 \end{pmatrix} \quad (5.13)$$

on the domain $\Omega = [-\frac{1}{2}, \frac{1}{2}]^2$ and a vertical magnetic field on the boundary with the condition $\vec{B} \times \vec{n} = (0, 1) \times \vec{n}$ on $\partial\mathcal{D}$. Figure 5.2 demonstrates the effect of the resistivity on the induced magnetic field \vec{B} for three values of η on a 64×64 element mesh. In this figure, we have plotted the velocity profile defined by (5.13) as well as the magnetic field lines for the solution \vec{B}_h with $\eta = 10^{-3}, 10^{-2}$, and 10^{-1} . The figure demonstrates that two competing physical processes are at play in the kinematics model. First, the boundary condition corresponds to an external magnetic field applied to the domain. With infinite resistivity, the velocity field

plays no role in the kinematics equations, so the solution is determined solely by the boundary conditions. In this case, this results in the uniform vertical magnetic field $\vec{B} = (0, 1)$. For large resistivity, the solution is dominated by this process. For $\eta = 10^{-1}$ for instance, the solution appears to be a perturbation of the field $\vec{B} = (0, 1)$. The second physical process is governed by the effect of the velocity field on the magnetic field. For small resistivity, the kinematic equations are dominated by the convective term $\nabla \times (\vec{u} \times \vec{B})$ which tends to pull the magnetic field lines in the direction of the velocity field. As the resistivity approaches zero, the topology of the magnetic field then approaches that of the velocity field. For this problem, this means that the magnetic field lines should look more like concentric ellipses as η decreases. This is demonstrated for $\eta = 10^{-3}$, where the magnetic field is nearly “frozen” in the fluid in the center of the domain. Hence, these simulations show that qualitative characteristics of the magnetic topology can indicate the relative resistivity of the system. The more the field lines appear to be pulled by the velocity field (i.e. for this problem, the more swirling in the magnetic field), the smaller the resistivity.

5.2 MHD Kinematics with Uncertain Data

Because the equations of MHD kinematics may involve uncertain quantities, we propose a formulation that incorporates uncertainty. In particular, we consider the cases where the resistivity η or the velocity field \vec{u} are uncertain. (Because the permeability μ does not vary significantly in applications, we will take it as fixed.

Equivalently, we can regard any uncertainty in μ as being absorbed into η .) In this section, we present mathematical representations for each quantity that incorporate uncertainty by treating $\eta = \eta(\omega)$ and $\vec{u} = \vec{u}(\omega)$ as random variables. This yields bilinear forms depending on random variables, i.e.

$$a(\vec{B}, \vec{C}, \omega) = \left(\frac{\eta(\omega)}{\mu} \nabla \times \vec{B}, \nabla \times \vec{C} \right), \quad (5.14a)$$

$$c(\vec{B}, \vec{C}, \omega) = - \left(\vec{u}(\omega) \times \vec{B}, \nabla \times \vec{C} \right), \quad (5.14b)$$

defining a stochastic weak formulation. Thus, by the Doob-Dynkin Lemma, the solution (\vec{B}_h, r_h) is also a random variable defined on the same sample space [53]. Hence, each realization of $\eta(\omega)$ and $\vec{u}(\omega)$ yields the weak formulation: Find $(\vec{B}_h(\omega), r_h(\omega)) \in V_{\vec{q}}^h \times Q_0^h$ such that

$$a(\vec{B}_h(\omega), \vec{C}_h, \omega) + c(\vec{B}_h(\omega), \vec{C}_h, \omega) + b(\vec{C}_h, r_h(\omega)) = 0, \quad (5.15a)$$

$$b(\vec{B}_h(\omega), s_h) = 0, \quad (5.15b)$$

for all $(\vec{C}_h, s_h) \in V_0^h \times Q_0^h$. The solution of any realization of this problem can be obtained by solving a linear system of the form

$$\begin{pmatrix} A(\omega) + N(\omega) & D^t \\ D & 0 \end{pmatrix} \begin{pmatrix} \mathbf{B}(\omega) \\ \mathbf{r}(\omega) \end{pmatrix} = \begin{pmatrix} f \\ 0 \end{pmatrix}. \quad (5.16)$$

Given this framework, we can employ a Monte-Carlo simulation to obtain statistical properties of \vec{B}_h . We repeatedly generate independent random instances of $\eta(\omega)$ and $\vec{u}(\omega)$ and solve for $\vec{B}_h(\omega)$. We then estimate the mean and standard deviation of \vec{B}_h by the (pointwise) sample mean, which we denote $\mu(\vec{B}_h)(\vec{x})$, and the sample standard deviation, which we denote $\sigma(\vec{B}_h)(\vec{x})$. A canonical error estimate for Monte-Carlo

simulation [10] states that for N trials, the (pointwise) stochastic error for each component of $\mu(\vec{B}_h)$ satisfies

$$|E(B_i(\vec{x}, \omega)) - (\mu(\vec{B}_h)(\vec{x}))_i| \approx 2 \frac{(\sigma(\vec{B}_h)(\vec{x}))_i}{\sqrt{N}} \quad (5.17)$$

with 95% confidence. Thus, we obtain with 95% confidence the error result

$$\|E(\vec{B}(\vec{x}, \omega)) - \mu(\vec{B}_h)(\vec{x})\|_2 \approx 2 \frac{\|\sigma(\vec{B}_h)(\vec{x})\|_2}{\sqrt{N}}. \quad (5.18)$$

5.2.1 Uncertain Velocity

In this section, we consider the case where fluctuations are allowed in the velocity field. We assume that a mean flow, \vec{u}_0 , is known and represent the fluctuations by a random variable, $\vec{u}_*(\omega)$, with mean zero. We thus express \vec{u} as the sum of its deterministic and random parts as

$$\vec{u}(\omega) = \vec{u}_0 + \vec{u}_*(\omega). \quad (5.19)$$

Rather than letting each component of \vec{u}_* be an independent scalar random variable, we derive two expressions for \vec{u}_* from assumptions about physical properties of the fluid, either that the fluctuation is irrotational or that the fluid is incompressible. This results in a natural coupling of the components of \vec{u}_* . We explore the effects of an uncertain velocity field on the MHD kinematics system in Sections 5.2.1.1 and 5.2.1.2 by applying fluctuations of these types to the Hartmann flow problem detailed in Section 5.1.1.

5.2.1.1 Test Problem 1: Irrotational Fluctuations

If the random fluctuations of the fluid are irrotational ($\nabla \times \vec{u}_* = 0$), then \vec{u}_* is a conservative vector field and can be written as the gradient of a scalar potential ϕ , i.e.

$$\vec{u}_* = \nabla \phi. \quad (5.20)$$

Under this assumption, only the random scalar field ϕ needs to be specified in order to define \vec{u} . We assume the potential field to vary continuously and to be spatially correlated. These assumptions are satisfied if we assume ϕ to be a stationary random field with the covariance function defined by

$$C(\vec{x}, \vec{y}) = \sigma^2 e^{-\frac{\|\vec{x}-\vec{y}\|_2}{\ell}}. \quad (5.21)$$

Here, σ^2 is the variance and ℓ is a correlation length. Clearly, the covariance is greatest when the Euclidean distance between the points \vec{x} and \vec{y} is small. In effect, this covariance function generates fluctuations in ϕ on a scale proportional to ℓ .

If ϕ has mean zero, then ϕ can be approximated by a truncated Karhunen-Loève (KL) expansion [31]

$$\phi(\vec{x}, \xi) \approx \sum_{i=1}^M \sqrt{\lambda_i} \phi_i(\vec{x}) \xi_i, \quad (5.22)$$

where $\phi_i(\vec{x})$ and λ_k are the eigenfunctions and eigenvalues of C . We will assume the random variables $\{\xi_i\}$ to be independently and uniformly distributed in $[-\sqrt{3}, \sqrt{3}]^M$. We choose M large enough to capture 95% of the total variance [9]; that is,

$$\sum_{i=1}^M \lambda_i > 0.95 |\mathcal{D}| \sigma^2. \quad (5.23)$$

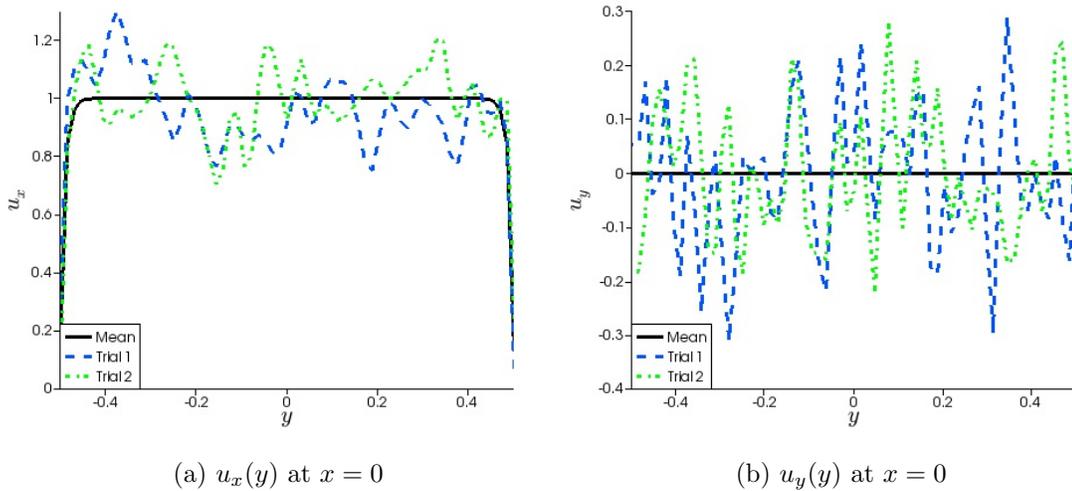


Figure 5.3: Profiles of u_x and u_y along the line $x = 0$ for two random instances with $\sigma^2 = 5.0 \times 10^{-3}$, together with the mean profile.

We note that the correlation length affects this requirement, with small ℓ leading to large M .

We let the mean velocity profile \vec{u}_0 be given by the deterministic Hartmann profile (5.12a) and introduce fluctuations by letting \vec{u} be defined by (5.19) and (5.22) using the covariance function (5.21). We let $\eta \equiv 10^{-2}$, $\nu = 10^{-2}$, and $\ell = 0.1$. This correlation length corresponds to fairly small-scale fluctuations in the velocity field. We compare the effects of three choices for the variance, $\sigma^2 = 5.0 \times 10^{-3}$, 6.0×10^{-3} , and 7.0×10^{-3} . The increase in σ^2 corresponds to an increase in the magnitude of the fluctuations. We present results for this problem discretized on a 64×64 element mesh. Sample random instances of \vec{u} with $\sigma^2 = 5.0 \times 10^{-3}$ are plotted in Figure 5.3. In this figure, the profiles of the components of \vec{u} , u_x and u_y , are plotted along the line $x = 0$ and compared to the mean values $(u_0)_x$ and $(u_0)_y$. The corresponding

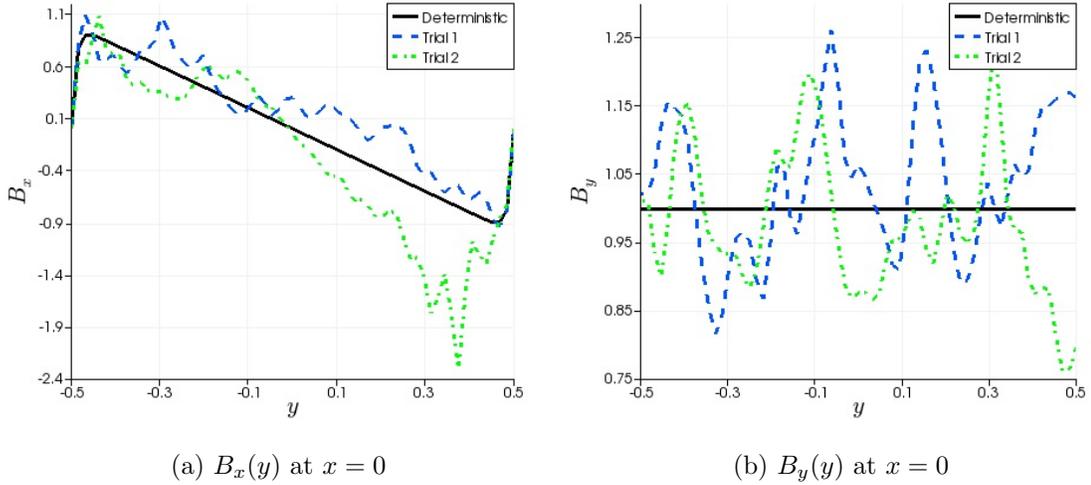


Figure 5.4: Profiles of B_x and B_y along the line $x = 0$ for two random instances with $\sigma^2 = 5.0 \times 10^{-3}$, together with the deterministic solution obtained from \vec{u}_0 .

solutions \vec{B}_h are plotted in Figure 5.4. These are compared to the deterministic solution obtained with $\vec{u} \equiv \vec{u}_0$. Both the random data \vec{u} and the corresponding solutions demonstrate high frequency oscillations around the deterministic profiles.

Some results for the Monte-Carlo simulation after 10,000 trials are plotted in Figure 5.5. In this figure, the profiles of the mean solution $\mu(\vec{B}_h)$ are plotted along the line $x = 0$ for the three values of σ^2 . This is compared to the deterministic solution \vec{B}_h with $\vec{u} \equiv \vec{u}_0$. The Euclidean norm of the pointwise variance of the solution $\|\sigma(\vec{B}_h)(\vec{x})\|_2$ is bounded by 0.28 for $\sigma^2 = 5.0 \times 10^{-3}$, 0.38 for $\sigma^2 = 6.0 \times 10^{-3}$, and 0.51 for $\sigma^2 = 7.0 \times 10^{-3}$. By (5.18), this implies that the maximum (pointwise) stochastic error for this problem is approximately 0.01. From Figure 5.5, it can be seen that on average, irrotational fluctuations in the velocity field result in a slight growth in the magnitude of the induced magnetic field as compared to the

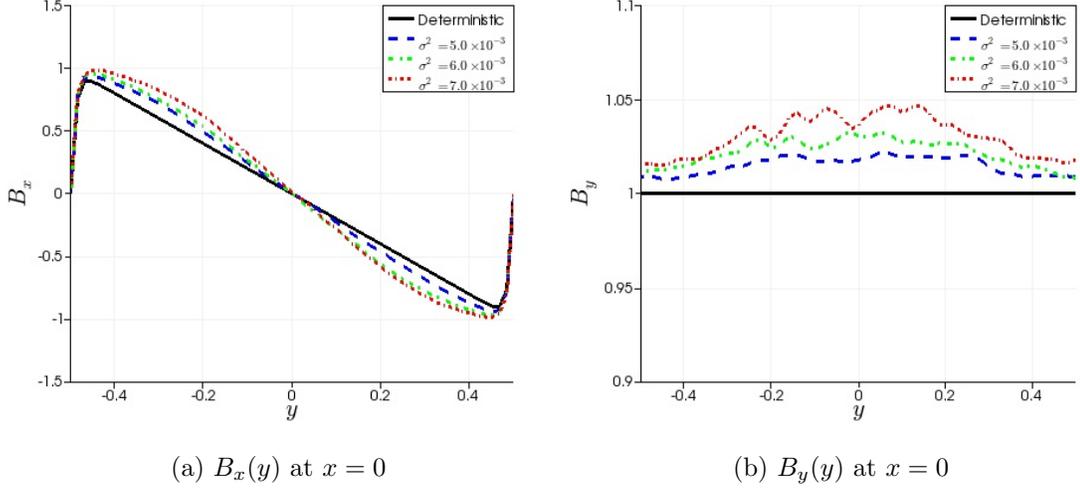


Figure 5.5: Profiles of B_x and B_y for Test Problem 1, plotted along the line $x = 0$ for the mean $\mu(\vec{B}_h)$ with $\sigma^2 = 5.0 \times 10^{-3}$, 6.0×10^{-3} , and 7.0×10^{-3} .

deterministic case. This growth increases as σ^2 increases.

5.2.1.2 Test Problem 2: Incompressible Flow

If the fluid is assumed to be incompressible (i.e. $\nabla \cdot \vec{u} = 0$) and \vec{u}_0 is incompressible, then \vec{u}_* must also be incompressible (see (5.19)). Thus, in this case, we can prescribe \vec{u}_* to be the curl of a potential ϕ . In two dimensions, ϕ is a scalar, and in three dimensions ϕ is a vector. We consider only the 2D case in this study.

Then, \vec{u}_* can be written as

$$\vec{u}_* = \left(\frac{\partial \phi}{\partial y}, -\frac{\partial \phi}{\partial x} \right), \quad (5.24)$$

and the random variable \vec{u}_* can be computed from the scalar random variable ϕ . As above, we let ϕ be defined by a KL expansion (5.22) with the covariance function C .

Again, we let $\eta \equiv 10^{-2}$, $\nu = 10^{-2}$, and $\ell = 0.1$ and consider $\sigma^2 = 5.0 \times 10^{-3}$, $6.0 \times$

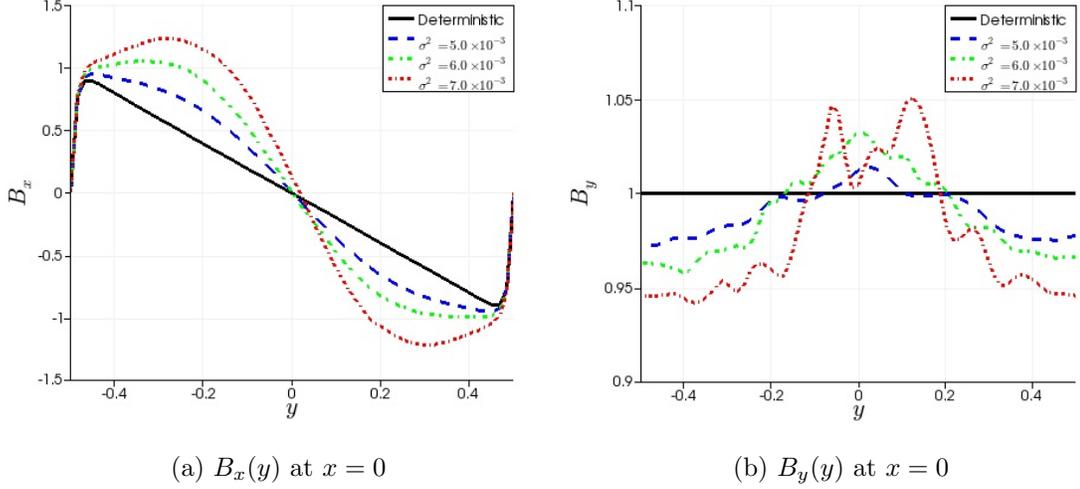


Figure 5.6: Profiles of B_x and B_y for Test Problem 2, plotted along the line $x = 0$ for the mean $\mu(\vec{B}_h)$ with $\sigma^2 = 5.0 \times 10^{-3}$, 6.0×10^{-3} , and 7.0×10^{-3} .

10^{-3} , and 7.0×10^{-3} .

Mean solution profiles obtained from the Monte-Carlo simulation after 10,000 trials are plotted in Figure 5.6. The normed standard deviation $\|\sigma(\vec{B}_h)(\vec{x})\|_2$ is bounded by 0.49 for $\sigma^2 = 5.0 \times 10^{-3}$, 1.23 for $\sigma^2 = 6.0 \times 10^{-3}$, and 3.17 for $\sigma^2 = 7.0 \times 10^{-3}$, corresponding to a maximum stochastic error of about 0.06. Compared to the case of irrotational fluctuations, both the standard deviation of the solution and the magnitude of the induced magnetic field are greater when non-zero vorticity is permitted in the fluctuations. The difference in the magnitude of the magnetic field is fairly significant between the two test problems, suggesting that fluid vorticity plays a large role in generating magnetic fields. Thus, when small-scale rotational behavior is present in a fluid, simulations based on the mean flow of the fluid may not capture important magnetic effects.

5.2.2 Uncertain Resistivity

In this section, we consider the case where just the resistivity η is a random field over the domain. This is motivated by the fact that multiple fluids may be present in a physical system and there may be some epistemic uncertainty in the distribution of the fluids throughout the domain. In practical applications, the resistivity can range over orders of magnitude between two fluids. For example, in aluminum electrolysis, the resistivity of liquid aluminum is approximately $4.0 \times 10^{-3} \Omega m$, while the resistivity of the fluid bath from which the aluminum is reduced is approximately $2.9 \times 10^{-7} \Omega m$ [29]. We propose defining

$$\eta(\vec{x}, \omega) = 10^{\beta(\vec{x}, \omega)}, \quad (5.25)$$

where β is a random scalar field yet to be specified. This expression both emphasizes the variability in the order of magnitude of η and guarantees that $\eta > 0$. We first use this to investigate the effects of uncertain resistivity on the MHD eddy problem discussed in Section 5.1.2, for two different choices of β . We then consider a three-dimensional extension of the MHD eddy problem

5.2.2.1 Test Problem 3: Piecewise constant β

In this section, we assume that the domain is occupied by multiple immiscible fluids with different resistivities. In this setting, we let β be a piecewise constant scalar field over the domain. We partition the domain \mathcal{D} into n subdomains $P = \{\mathcal{D}_k\}_{k=1}^n$ and let β be a constant on each of these subdomains. If we assume the

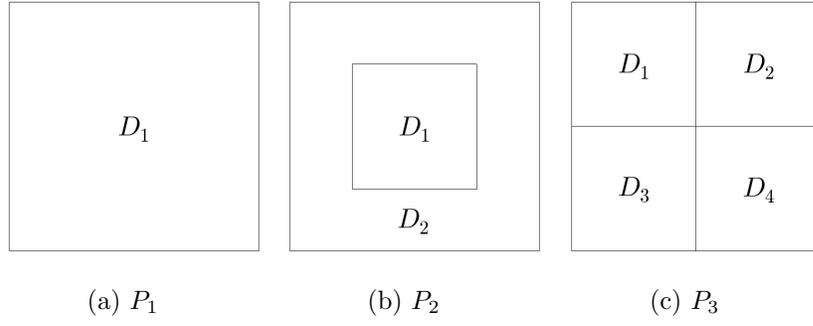
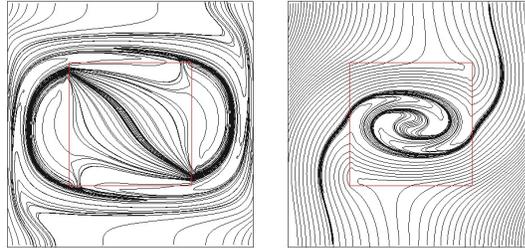


Figure 5.7: Domain partitionings considered in Test Problem 3.



(a) $\eta = 10^{-1}$ on D_1 (b) $\eta = 10^{-3}$ on D_1
 $\eta = 10^{-3}$ on D_2 $\eta = 10^{-1}$ on D_2

Figure 5.8: Instances of \vec{B}_h obtained with partitioning P_2 for Test Problem 3.

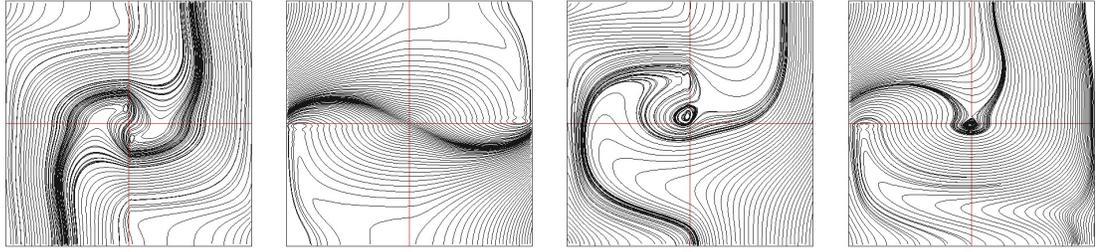
resistivity to be uncertain on each of the subdomains, then we can let $\beta(\cdot, \vec{\xi})|_{\mathcal{D}_k} = \xi_k$ where $\vec{\xi} = [\xi_1, \dots, \xi_n]^t$ is a random vector.

We investigate the effect of a piecewise constant resistivity by considering the MHD eddy problem, defining \vec{u} by (5.13) on the domain $[-\frac{1}{2}, \frac{1}{2}]^2$. We consider three partitionings of the domain:

$$P_1 = \{[-\frac{1}{2}, \frac{1}{2}]^2\}, \tag{5.26}$$

$$P_2 = \{[-\frac{1}{4}, \frac{1}{4}]^2, [-\frac{1}{2}, \frac{1}{2}]^2 \setminus [-\frac{1}{4}, \frac{1}{4}]^2\}, \tag{5.27}$$

$$P_3 = \{[-\frac{1}{2}, 0) \times [0, \frac{1}{2}], [0, \frac{1}{2}] \times [0, \frac{1}{2}], [-\frac{1}{2}, 0) \times [-\frac{1}{2}, 0), [0, \frac{1}{2}] \times [-\frac{1}{2}, 0)\}, \tag{5.28}$$



(a) $\eta = 10^{-3}$ on D_1, D_4 (b) $\eta = 10^{-3}$ on D_2, D_3 (c) $\eta = 10^{-3}$ on D_1, D_3 (d) $\eta = 10^{-3}$ on D_3, D_4
 $\eta = 10^{-1}$ on D_2, D_3 $\eta = 10^{-1}$ on D_1, D_4 $\eta = 10^{-1}$ on D_2, D_4 $\eta = 10^{-1}$ on D_1, D_2

Figure 5.9: Instances of \vec{B}_h obtained with partitioning P_3 for Test Problem 3.

as shown in Figure 5.7. We let each ξ_i be independently and uniformly distributed in the interval $[-1.0, -3.0]$. Defining η by (5.25), we obtain $E(\eta) \equiv \eta_0 \approx 2.1 \times 10^{-2}$ independent of the partitioning. With partitioning P_1 , the resistivity is a random constant over the domain. Thus, the solutions plotted in Figure 5.2 are representative instances of magnetic fields that may be induced for this partitioning. With partitioning P_2 , the resistivity in the center of the domain may differ from the resistivity near the boundaries, and this can result in in different kinds of behavior than seen for a constant resistivity. Two examples of the kind of behavior we may obtain are plotted in Figure 5.8. In the first example, the resistivity is much larger in the center subdomain, and one can see that the behavior of the magnetic field is more characteristic of a larger resistivity in the center. Near the boundary of \mathcal{D} , the field lines are similar to those obtained with $\eta = 10^{-3}$ on the entire domain. Around the interface of the two subdomains, the character of the field lines shifts. The second example in Figure 5.8 shows the opposite case, where the resistivity is smaller in the center subdomain. In this case, one can see behavior characteristic of a large

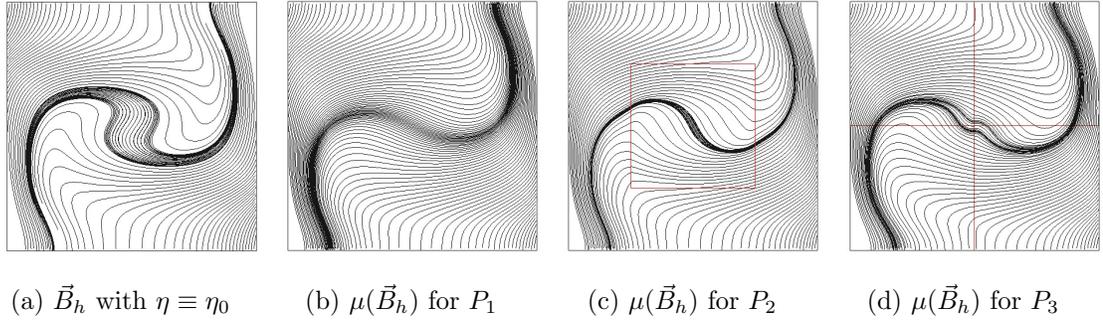


Figure 5.10: Deterministic and mean magnetic field lines compared for Test Problem 3.

resistivity near the boundary of \mathcal{D} and behavior characteristic of a small resistivity in the center. Figure 5.9 depicts some examples of magnetic fields that can result from partitioning P_3 . As with partitioning P_2 , it can be seen that the character of the field lines in a particular subdomain are characteristic of the resistivity on that subdomain, and at subdomain interfaces there is a shift in the character of the magnetic field. Figures 5.8 and 5.9 demonstrate not only that a discontinuous resistivity can produce profoundly different solutions than a constant resistivity, but also that the resistivity in a particular region can be approximated by considering the character of the field lines in that region.

Some results for the Monte-Carlo simulation after 10,000 trials on a 64×64 element mesh are shown in Figures 5.10 and 5.11. The field lines for the deterministic case where $\eta = E(\eta(\omega))$ are compared to those obtained from the mean $\mu(\vec{B}_h)(\vec{x})$ in Figure 5.10, and the norm of the standard deviation $\|\sigma(\vec{B}_h)\|_2$ is plotted in Figure 5.11. Note that the probability that the resistivity is greater than the mean resistivity ($Pr(\eta > \eta_0) \approx 0.33$) is less than the probability that the resistivity is less

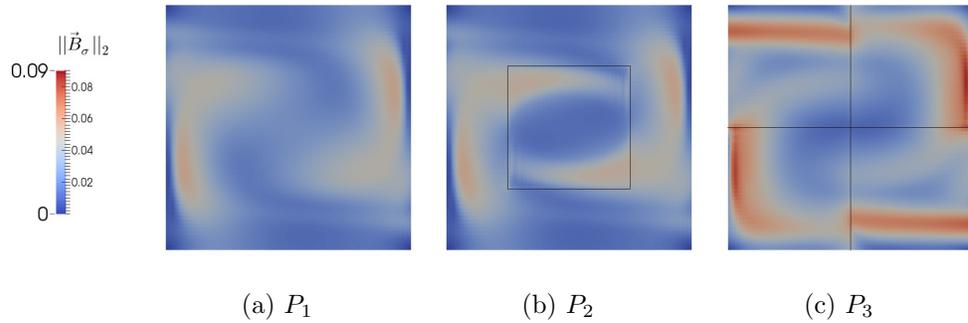


Figure 5.11: Euclidean norm of standard deviation $\|\sigma(\vec{B}_h)\|_2$ for Test Problem 3.

than the mean resistivity ($Pr(\eta < \eta_0) \approx 0.67$). Despite this, the mean field lines in each example resemble ones that would arise with $\eta > \eta_0$. Because this effect occurs for partitioning P_1 , it appears that it is due primarily to the variability of η . When the resistivity is allowed to vary for this test problem, the mean magnetic field is dominated by the qualities of magnetic fields induced by larger resistivities. Figure 5.10 shows that the large-scale behavior of the mean magnetic field is very similar for the three partitionings, although some small-scale differences are present. For partitioning P_2 , the behavior in the center subdomain is consistent with a slightly smaller resistivity than in the rest of the domain. With partitioning P_3 , the field lines appear to correspond to a larger resistivity around the origin where the four subdomains meet. The results suggest that variability in the resistivity has a more significant effect on the mean magnetic field than the presence of multiple fluids with different resistivities. Although multiple resistivities may result in random instances of \vec{B} that differ significantly from a constant resistivity, as shown in Figures 5.8 and 5.9, the means do not differ dramatically from solutions for constant resistivities, with the main differences arising at subdomain interfaces.

Figure 5.11 shows the norm of the standard deviation of \vec{B}_h over the domain. Because $\|\sigma(\vec{B}_h)(\vec{x})\|_2 \leq 0.09$, we know from (5.18) that the stochastic error is bounded as

$$\|E(\vec{B}_h(\vec{x}, \omega)) - \mu(\vec{B}_h)(\vec{x})\|_2 \leq 2 \frac{0.09}{\sqrt{10,000}} = 1.8 \times 10^{-3} \quad (5.29)$$

with 95% confidence. From the figure, it can be seen that the standard deviation is affected by the partitioning of the domain. In this case, the standard deviation increases as the number of subdomains increases. The average standard deviation over the domain for partitionings P_1 , P_2 , and P_3 is approximately 2.9×10^{-2} , 3.0×10^{-2} , and 4.1×10^{-2} . Furthermore, the standard deviation tends to be larger along the subdomain interfaces.

5.2.2.2 Test Problem 4: β as a truncated KL expansion

In this section, we assume that the resistivity of the system varies continuously over space. This situation may arise when fluids can blend together, such as when different liquid metals are combined into an alloy. This scenario can be modeled by supplying β with the spatially correlated covariance function C of (5.21). If β has mean $\beta_0(\vec{x})$, then the truncated KL expansion for β can be written as

$$\beta(\vec{x}, \xi) \approx \beta_0(\vec{x}) + \sum_{i=1}^M \sqrt{\lambda_i} \beta_i(\vec{x}) \xi_i. \quad (5.30)$$

We assume the random variables $\{\xi_i\}$ to be independently and uniformly distributed in $[-\sqrt{3}, \sqrt{3}]^M$ and choose M large enough to capture 95% of the total variance.

We consider the domain $[-\frac{1}{2}, \frac{1}{2}]^2$ and prescribe the velocity field (5.13) with η defined by (5.25). We let β have mean $\beta_0 = -2.0$ and variance $\sigma^2 = 0.4$. This defines

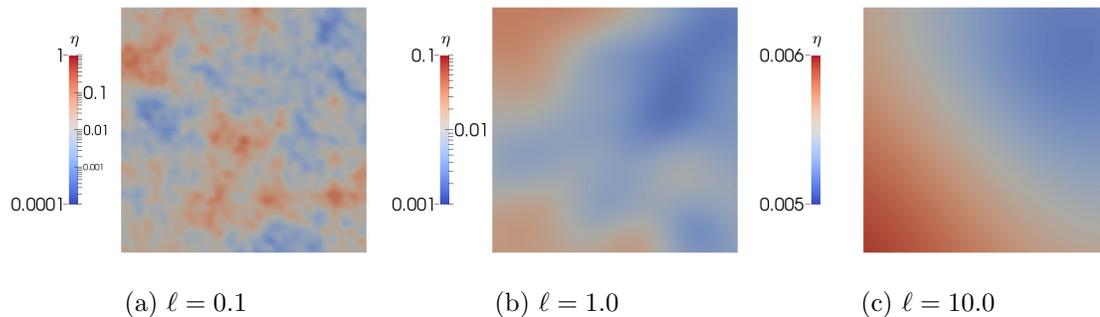
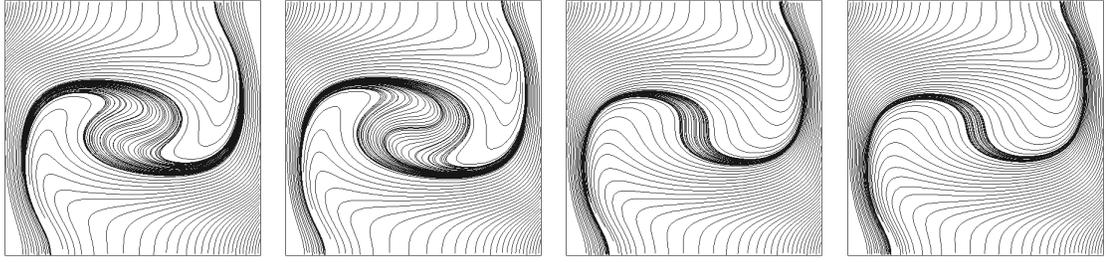


Figure 5.12: Random realizations of η for Test Problem 4.

the mean of η to be $E(\eta(\omega)) = \eta_0 \approx 1.5 \times 10^{-2}$ (note that $\eta_0 \neq 10^{E(\beta)}$). We discretize the problem on a 64×64 element mesh. We consider three choices of correlation length $\ell = 0.1, 1.0$, and 10.0 resulting in truncated KL expansions of length 1834, 39, and 2. Representative realizations of β are plotted in Figure 5.12 for these choices. The plots demonstrate that as ℓ increases the resistivity varies over a smaller range for a given realization. Realizations with $\ell = 0.1$, thus, represent heterogeneous systems in which fluids with disparate resistivities have not been mixed well. As ℓ increases the realizations become more homogeneous in resistivity, corresponding to later stages in a mixing process. With $\ell = 10.0$, η varies over a very small range in a particular instance, and in this respect it is similar to the constant resistivity obtained for Test Problem 3 with partitioning P_1 . Thus, large correlation lengths correspond to an uncertain final resistivity at the end of a mixing process.

Figures 5.13 and 5.14 show some results of the Monte-Carlo simulation after 10,000 trials. Figure 5.13 shows the magnetic field lines associated with the mean of \vec{B}_h with each value of ℓ . These results are compared with the solution to the deterministic problem for $\eta \equiv \eta_0$. It can be seen that with $\ell = 0.1$, the mean field



(a) \vec{B}_h with $\eta \equiv \eta_0$ (b) $\mu(\vec{B}_h)$ with $\ell = 0.1$ (c) $\mu(\vec{B}_h)$ with $\ell = 1.0$ (d) $\mu(\vec{B}_h)$ with $\ell = 10.0$

Figure 5.13: Deterministic and mean magnetic field lines compared for Test Problem

4.

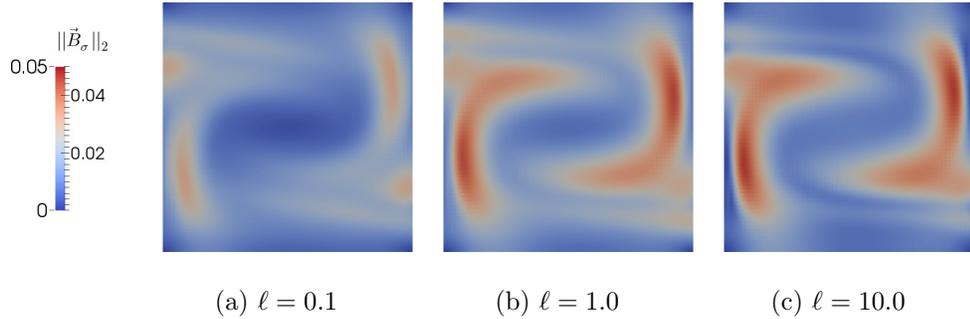


Figure 5.14: Euclidean norm of standard deviation $\|\sigma(\vec{B}_h)\|_2$ for Test Problem 4.

lines are very similar to the field lines obtained from the deterministic problem. As ℓ increases, the field lines differ more from the deterministic solution. In fact, the behavior exhibited for larger ℓ appears to result from a resistivity greater than the mean resistivity η_0 ; that is, as ℓ increases, the magnetic field lines are drawn more toward the infinite resistivity solution $\vec{B} = (0, 1)$.

Figure 5.14 shows the norm of the standard deviation of \vec{B}_h over the domain. Because $\|\sigma(\vec{B}_h)(\vec{x})\|_2 \leq 0.05$, the stochastic error is bounded above by 1.0×10^{-3} with 95% confidence. From this figure, it can be seen that the standard deviation

increases as ℓ increases. For small ℓ , the mean solution is not only more similar to the deterministic solution, but the standard deviation is also smaller. This suggests that the induced magnetic field is more responsive to variations in the resistivity of homogeneous systems than to small-scale variations in the resistivity of heterogeneous systems. Although the resistivity can vary over a very large range in a heterogeneous system, the fluctuations within random instances have little effect on the induced magnetic field. On the other hand, profound differences from the deterministic case are seen when the resistivity of a homogeneous system is uncertain. This is consistent with the results obtained for Test Problem 3, where we found that variability in the resistivity had stronger effects on the mean solution than the presence of multiple resistivities.

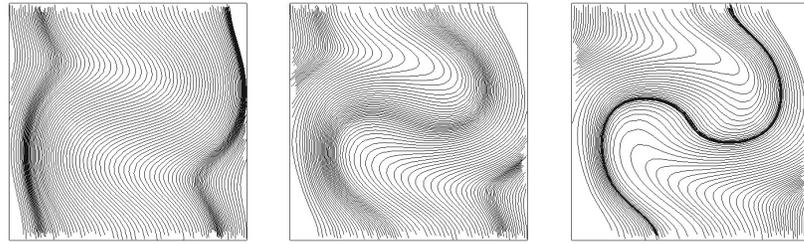
5.2.2.3 Test Problem 5: Uncertain Resistivity in 3D

In this section, we examine the behavior of solutions obtained for three-dimensional models with uncertain piecewise constant resistivity. We take as our domain the cube $[-\frac{1}{2}, \frac{1}{2}]^3$, and consider a generalization of the velocity field (5.13) in which the component in the z -direction is identically 1,

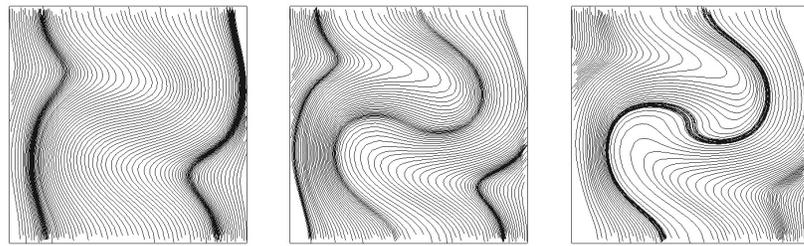
$$\vec{u}(x, y, z) = \begin{pmatrix} \frac{\cos(\pi x)}{\pi} 32y(1 - 4y^2)^3 \\ -\sin(\pi x)(1 - 4y^2)^4 \\ 1 \end{pmatrix}, \quad (5.31)$$

which describes a swirling flow in the z -direction. Again, we prescribe a magnetic field in the y -direction for the boundary condition, $\vec{B} \times \vec{n} = (0, 1, 0) \times \vec{n}$.

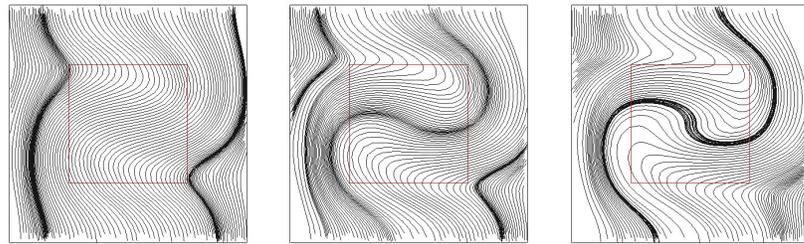
We define the resistivity by (5.25) using piecewise constant β on the partition-



(a) \vec{B}_h at $z = -\frac{1}{4}$ (b) \vec{B}_h at $z = 0$ (c) \vec{B}_h at $z = \frac{1}{4}$



(d) $\mu(\vec{B}_h)$ at $z = -\frac{1}{4}$ for P_1 (e) $\mu(\vec{B}_h)$ at $z = 0$ for P_1 (f) $\mu(\vec{B}_h)$ at $z = \frac{1}{4}$ for P_1



(g) $\mu(\vec{B}_h)$ at $z = -\frac{1}{4}$ for P_2 (h) $\mu(\vec{B}_h)$ at $z = 0$ for P_2 (i) $\mu(\vec{B}_h)$ at $z = \frac{1}{4}$ for P_2

Figure 5.15: Deterministic fields lines (top), mean field lines with partitioning P_1 (middle), and mean field lines with partitioning P_2 at cross sections $z = -\frac{1}{4}$ (left), $z = 0$ (middle), and $z = \frac{1}{4}$ (right) for Test Problem 5.

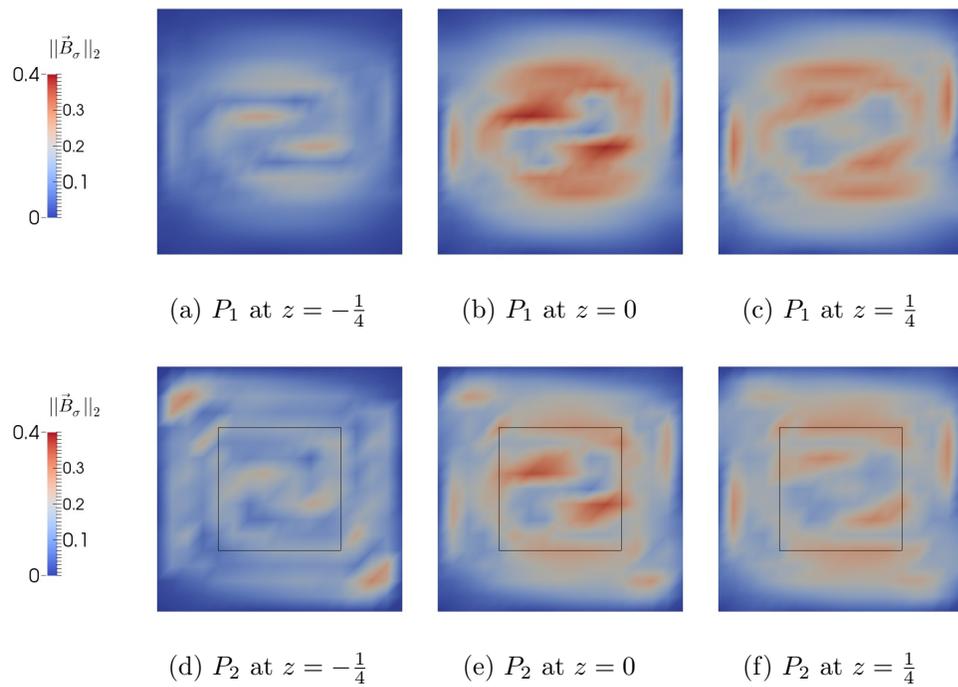


Figure 5.16: $\|\sigma(\vec{B}_h)(\vec{x})\|_2$ at cross sections for partitionings P_1 (top) and P_2 (bottom) for Test Problem 5.

ings

$$P_1 = \{[-\frac{1}{2}, \frac{1}{2}]^3\}, \quad (5.32)$$

$$P_2 = \{[-\frac{1}{4}, \frac{1}{4}]^2 \times [-\frac{1}{2}, \frac{1}{2}], ([-\frac{1}{2}, \frac{1}{2}]^2 \setminus [-\frac{1}{4}, \frac{1}{4}]^2) \times [-\frac{1}{2}, \frac{1}{2}]\}. \quad (5.33)$$

These are extensions of partitionings P_1 and P_2 from Test Problem 3 in the z -direction. On each subdomain \mathcal{D}_k , we let $\beta(\cdot, \vec{\xi})|_{\mathcal{D}_k} = \xi_k$ where ξ_k is uniformly distributed in $[-1.0, -3.0]$. The mean resistivity in this case is $\eta_0 \approx 2.2 \times 10^{-2}$. Field lines of the sample mean computed from 1,000 Monte-Carlo simulations on a $16 \times 16 \times 16$ element mesh are compared to deterministic field lines for $\eta \equiv \eta_0$ in Figure 5.15. Field lines are computed in the x - y cross sections at $z = -\frac{1}{4}, 0$, and $\frac{1}{4}$. From the deterministic field lines, it can be seen that near the bottom of the domain, the magnetic field is more dominated by the infinite resistivity solution $\vec{B} = (0, 1, 0)$, but as z increases the velocity field has a larger effect on the magnetic field. Comparing to the two-dimensional case, this is as if the resistivity decreases as z increases. We see the same effect for the mean solution. Unlike for the 2D test problems, the mean magnetic field looks like a deterministic solution with $\eta < \eta_0$. Comparing the mean field lines obtained for the two different partitionings, differences due to the multiple resistivities present in partitioning P_2 are very slight. As in 2D, this shows that the effects of variability in the resistivity are stronger than the effects of discontinuities in the resistivity. Cross sections of the norm of the standard deviation are plotted in Figure 5.16. This shows that variability in the magnetic field is greatest at the center cross section $z = 0$. Furthermore, the figure demonstrates that, as in Test Problem 3, a greater number of subdomains leads to

a larger standard deviation that is distributed throughout more of the domain.

5.3 Linear Solvers for the Discretized Kinematics System

In Monte-Carlo simulation, for each realization of the uncertain quantities, a linear system of the form $\mathcal{A}x = b$ must be solved, where

$$\mathcal{A} = \mathcal{A}(\omega) = \begin{pmatrix} A(\omega) + N(\omega) & D^t \\ D & 0 \end{pmatrix}. \quad (5.34)$$

Because many realizations are required to produce accurate statistical results, it is imperative that the linear solver be efficient and robust over random variations in the parameters η and \vec{u} . The linear systems are sparse, nonsymmetric and indefinite, and, depending on the level of spatial refinement, they can be very large. Hence, a preconditioned iterative method such as preconditioned GMRES is a natural choice of solver for these systems. Motivated by the results of [49], many successful block preconditioners have been developed for solving similar saddle point systems. Following this line of research, we develop a generalization of a preconditioner proposed in [35] for the time-harmonic Maxwell equations to be used for the discretized kinematics equations.

The system studied in [35] can be considered a special case of (5.5) in which $\frac{\eta}{\mu} \equiv 1$ and $\vec{u} \equiv 0$. If we let \hat{A} be a discretization of the unscaled magnetic diffusion operator $\nabla \times \nabla \times$, then the coefficient matrix of the resulting linear system can be written

$$\hat{\mathcal{A}} = \begin{pmatrix} \hat{A} & D^t \\ D & 0 \end{pmatrix}. \quad (5.35)$$

The preconditioner developed in [35] is of the form

$$\hat{\mathcal{P}} = \begin{pmatrix} Q_{\mathbf{B}} + \hat{A} & 0 \\ 0 & L_{\mathbf{r}} \end{pmatrix}, \quad (5.36)$$

where $Q_{\mathbf{B}}$ is the mass matrix for \mathbf{B} and $L_{\mathbf{r}}$ is a discrete Laplacian on the magnetic pseudo-pressure space. A generalization of this preconditioner for use with the system \mathcal{A} is

$$\mathcal{P}_k = \begin{pmatrix} kQ_{\mathbf{B}} + A + N & 0 \\ 0 & \frac{1}{k}L_{\mathbf{r}} \end{pmatrix}, \quad (5.37)$$

where $k > 0$ is a constant to be specified.

5.3.1 Analysis of Eigenvalues

We give a complete analysis of this preconditioner for the case where $N \equiv 0$. The performance of preconditioner \mathcal{P}_k for system \mathcal{A} is governed by the eigenvalues λ of the generalized eigenvalue problem

$$\begin{pmatrix} A & D^t \\ D & 0 \end{pmatrix} \begin{pmatrix} \mathbf{B} \\ \mathbf{r} \end{pmatrix} = \lambda \begin{pmatrix} kQ_{\mathbf{B}} + A & 0 \\ 0 & \frac{1}{k}L_{\mathbf{r}} \end{pmatrix} \begin{pmatrix} \mathbf{B} \\ \mathbf{r} \end{pmatrix}. \quad (5.38)$$

Defining $n = \dim(\mathbf{B})$ and $m = \dim(\mathbf{r})$, this has a total of $n + m$ eigenvalues. From the bottom row of (5.38), we obtain $\mathbf{r} = \frac{k}{\lambda}L_{\mathbf{r}}^{-1}D\mathbf{B}$. Substituting this into the top row of (5.38) gives

$$\lambda\mathbf{A}\mathbf{B} + kD^tL_{\mathbf{r}}^{-1}D\mathbf{B} = \lambda^2(kQ_{\mathbf{B}} + A)\mathbf{B}. \quad (5.39)$$

Through a discrete Hodge decomposition, \mathbf{B} can be written as the sum of its discrete curl-free part \mathbf{B}_A and its discrete divergence-free part \mathbf{B}_D (i.e. $\mathbf{B} = \mathbf{B}_A + \mathbf{B}_D$, where

$A\mathbf{B}_A = A^t\mathbf{B}_A = 0$ and $D\mathbf{B}_D = 0$). Then (5.39) can be rewritten as

$$\lambda A\mathbf{B}_D + kD^tL_r^{-1}D\mathbf{B}_A = \lambda^2kQ_{\mathbf{B}}(\mathbf{B}_A + \mathbf{B}_D) + \lambda^2A\mathbf{B}_D. \quad (5.40)$$

Let the norm induced by a symmetric positive definite matrix X be denoted $\|\cdot\|_X = \langle X\cdot, \cdot \rangle^{1/2}$. Taking the inner product of (5.40) with \mathbf{B}_A and using the relations

$$\langle Q_{\mathbf{B}}\mathbf{B}_A, \mathbf{B}_D \rangle = \langle Q_{\mathbf{B}}\mathbf{B}_D, \mathbf{B}_A \rangle = 0, \quad (5.41a)$$

$$\langle D^tL_r^{-1}D\mathbf{B}_A, \mathbf{B}_A \rangle = \|\mathbf{B}_A\|_{Q_{\mathbf{B}}}^2, \quad (5.41b)$$

proven in [35], we have

$$k\|\mathbf{B}_A\|_{Q_{\mathbf{B}}}^2 = \lambda^2k\|\mathbf{B}_A\|_{Q_{\mathbf{B}}}^2. \quad (5.42)$$

Because there are at least m linearly independent vectors satisfying $\mathbf{B}_A \neq 0$, this means that (5.38) has eigenvalues $\lambda = \pm 1$ each with multiplicity at least m .

Insight into the remaining $n - m$ eigenvalues can be obtained by taking the inner product of $\hat{\mathbf{B}}_D$ with (5.40), yielding

$$\lambda k\|\mathbf{B}_D\|_{Q_{\mathbf{B}}}^2 = (1 - \lambda)\|\mathbf{B}_D\|_A^2. \quad (5.43)$$

From this equation, it is clear that $0 \leq \lambda \leq 1$. These eigenvalues can be further bounded using the discrete coercivity condition. In [35], this condition is written in terms of the unscaled norm $\|\cdot\|_{\hat{A}}$ as

$$\|\mathbf{B}_D\|_{\hat{A}}^2 \geq \alpha (\|\mathbf{B}_D\|_{\hat{A}}^2 + \|\mathbf{B}_D\|_{Q_{\mathbf{B}}}^2), \quad (5.44)$$

where the constant $\alpha \in (0, 1)$ is independent of the mesh. It can be shown that

$$\frac{\eta_m}{\mu}\|\mathbf{B}_D\|_{\hat{A}}^2 \leq \|\mathbf{B}_D\|_A^2 \leq \frac{\eta_M}{\mu}\|\mathbf{B}_D\|_{\hat{A}}^2, \quad (5.45)$$

where $\eta_m = \min_{\vec{x} \in \mathcal{D}} \{\eta(\vec{x})\}$ and $\eta_M = \max_{\vec{x} \in \mathcal{D}} \{\eta(\vec{x})\}$, from which we can obtain the coercivity condition in terms of the scaled norm $\|\cdot\|_A$:

$$\|\mathbf{B}_D\|_A^2 \geq \alpha \left(\frac{\eta_m}{\eta_M} \|\hat{\mathbf{B}}_D\|_A^2 + \frac{\eta_m}{\mu} \|\mathbf{B}_D\|_{Q_B}^2 \right). \quad (5.46)$$

Applying this inequality to (5.43), we obtain

$$\frac{1}{1 + \frac{k\mu}{\alpha\eta_m} \frac{\eta_M - \alpha\eta_m}{\eta_M}} \leq \lambda \leq 1. \quad (5.47)$$

The constant $\frac{\eta_M - \alpha\eta_m}{\eta_M}$ is necessarily smaller than 1; thus, we can write the weaker bound

$$\frac{1}{1 + \frac{k\mu}{\alpha\eta_m}} < \lambda \leq 1. \quad (5.48)$$

If we let $k = \frac{\eta_m}{\mu}$, this bound depends only on the coercivity constant α . This dependence on α is similar to that obtained in [35] for $\hat{\mathcal{A}}$ preconditioned by $\hat{\mathcal{P}}$. Because α is independent of the mesh, letting $k = \frac{\eta_m}{\mu}$ defines a preconditioner which should be robust with respect to both mesh refinement and variations in the resistivity.

When $N \neq 0$, much of the same analysis applies. Because \mathbf{B}_A is curl-free, N satisfies $\langle N\mathbf{B}, \mathbf{B}_A \rangle = 0$. Given this relationship, the presence of N does not affect the two eigenvalues $\lambda = \pm 1$ with multiplicity m . However, the remaining eigenvalues are approximated by

$$\left(\lambda \left(\frac{\eta_m}{\mu} + k \right) - k \right) \|\mathbf{B}_D\|_{Q_B}^2 = (1 - \lambda) \left(\|\mathbf{B}_D\|_A^2 + \langle N\mathbf{B}, \mathbf{B}_D \rangle \right). \quad (5.49)$$

Because $\langle N\mathbf{B}, \mathbf{B}_D \rangle$ can become negative, it is difficult to say more about these eigenvalues, but in practice, this preconditioner proves to be effective. This will be demonstrated in the following section.

5.3.2 Numerical Results

Because the number of preconditioned GMRES iterations required for convergence depends on the linear system obtained from the random data $\vec{u}(\omega)$ and $\eta(\omega)$, we can regard these iteration counts as a random variable. Consequently, we obtain an estimate of the mean of the number of iterations required for convergence using the sample mean from a Monte-Carlo simulation. The GMRES iteration continues until the stopping criterion

$$\|b - \mathcal{A}x\| \leq 10^{-8}\|b\| \tag{5.50}$$

is satisfied. We compute the average number of iterations to reach this criterion over the Monte-Carlo simulation. We use the preconditioner \mathcal{P}_k defined by (5.37) with $k = \frac{\eta_m}{\mu}$. Direct methods are used to solve the subsidiary systems corresponding to the blocks $kQ_{\mathbf{B}} + A + N$ and $\frac{1}{k}L_{\mathbf{r}}$ in \mathcal{P}_k . While direct methods are viable for the small problems investigated in this study, we note that effective multigrid solvers have been developed for systems similar in form to the block $kQ_{\mathbf{B}} + A + N$ (see e.g. [39, 41] and the references therein).

Average iteration counts for each test problem are reported in Table 5.1, each on three different meshes. These results demonstrate that the preconditioner is highly effective for both problems with fluctuations in the velocity field (Test Problems 1 and 2) and those with heterogeneous resistivities (Test Problems 3, 4, and 5). Furthermore, the preconditioner is robust with respect to the mesh, with average iteration counts decreasing as the mesh is refined.

Test Problem	Mesh	16×16	32×32	64×64
	1 ($\sigma^2 = 6.0 \times 10^{-3}$)		5.99	5.95
2 ($\sigma^2 = 6.0 \times 10^{-3}$)		5.78	5.77	5.61
3 (P_2)		4.58	3.93	3.19
4 ($\ell = 1.0$)		4.98	4.32	3.05

	$4 \times 4 \times 4$	$8 \times 8 \times 8$	$16 \times 16 \times 16$
5 (P_2)	5.16	4.88	4.47

Table 5.1: Preconditioned GMRES iteration counts.

5.4 Conclusion

We have presented a numerical method for simulating the kinematics of MHD when either the velocity field or the magnetic resistivity of the fluid is uncertain, applying the method to several steady-state test problems. We have modeled the effect of random perturbations in the velocity field on the induced magnetic field. In particular, we have demonstrated that, on average, fluctuations with non-zero vorticity have a large global effect on the induced magnetic field. This supports the theory that small-scale turbulent flow is necessary for dynamo action. These results also suggest that simulations based on mean flow may underpredict the magnitude of magnetic fields. We have also demonstrated that uncertainty in the distribution of the resistivity can result in different magnetic topologies than in the deterministic case. Our results show that this effect is most pronounced when the resistivity in large regions of the domain is uncertain. On the other hand, we have found that, on average, the induced magnetic field is largely insensitive to small-scale fluctuations in the resistivity, even when these fluctuations vary over several orders of magnitude throughout the domain.

In this study, we have introduced several stochastic models for uncertain quantities in the context of MHD kinematics. By expressing the resistivity as a piecewise constant field or a truncated KL expansion, we allow for this quantity to be modeled by a vector of independent random variables. We have also proposed physically motivated expressions for the velocity field that not only couple its components in a natural way but also require the construction of only one random field in two

dimensions. We have employed Monte-Carlo simulation to obtain mean and variance data, but the stochastic expressions introduced here can be used directly in more sophisticated uncertainty quantification methods such as stochastic Galerkin, stochastic collocation, and quasi-Monte-Carlo methods.

In addition, we have developed a preconditioner for the discrete kinematics equations which is robust over variations in both the resistivity and the velocity field. This is important because many linear systems need to be solved in order to obtain accurate probabilistic distributions of the solution \vec{B}_h . Because this preconditioner is mesh independent, it allows for the possibility of larger-scale MHD kinematics simulations in both two and three dimensions. Furthermore, this preconditioner can be useful in fully coupled MHD models in which the resistivity and velocity field can fluctuate due to the coupling of the kinematics equations to the Navier-Stokes equations.

Chapter 6: Summary and Conclusions

With an increasing demand for high resolution MHD simulations, a growing body of research has centered around the development of fully coupled numerical models of MHD. While fully coupled methods allow for fast direct-to-steady-state solutions, they require the solution of large, sparse linear systems. These systems are particularly difficult to solve when hydrodynamic and electromagnetic phenomena are strongly coupled. This dissertation has sought to develop robust, scalable block preconditioners for iterative methods applied to the linear systems arising from discretizations of the fully coupled MHD equations.

In particular, we considered two finite element formulations for stationary, viscoresistive, incompressible MHD. In Chapter 3, we considered an exact penalty formulation, which enforces the divergence free constraint on the magnetic induction through a penalty term on the induction equation. While this formulation applies only to convex domains, it results in a simplified system where the unknowns are \vec{u} , p , and \vec{B} . Building on commutator-based strategies introduced for the discretized Navier-Stokes equations, we developed a family of preconditioners for both Picard and Newton linearizations of this formulation. We ensured that these preconditioners are effective on strongly coupled systems by algebraically embedding the effects

of coupling into Schur complements and accurately approximating these operators. By including “relaxation” parameters and tuning them with Fourier analysis, we proposed an automatic way of improving these approximations. The quality of these preconditioners was demonstrated for a range of parameters on a series of two-dimensional test problems.

Chapter 4 focused on a formulation that introduces a Lagrange multiplier into the induction equation and uses Nedelec elements \vec{B} , thereby allowing nonconvex domains. Because this formulation has the added complication of including the saddle point system associated with the mixed form of Maxwell’s equations as a subsystem, we incorporated preconditioners developed for Maxwell’s equations. Extending the strategies introduced for the exact penalty formulation in Chapter 3, we developed effective block preconditioners for this formulation. We showed that these preconditioners are both robust with respect to the magnetic Reynolds number as well as scalable on two- and three-dimensional test problems.

We also investigated the effects of uncertain problem parameters on the kinematics model of MHD. In Chapter 5, we developed a numerical method for modelling the effects of uncertain velocity fields and resistivities on an induced magnetic field. By considering either irrotational or incompressible fluctuations in the velocity, we reduced the problem of supplying an uncertain velocity field to supplying a random scalar potential over the domain. Modelling this potential by a truncated KL expansion, we demonstrated that small-scale fluctuations in the velocity can cause large-scale changes in the mean of the induced magnetic field. We modelled an uncertain resistivity with both piecewise constant random fields and truncated KL

expansions and demonstrated that the mean magnetic field is more sensitive to uncertainty in the resistivity of a homogeneous fluid than to spatial variations in the resistivity of a heterogeneous fluid, even when those variations are very large.

Bibliography

- [1] J. Adler, M. Brezina, T. Manteuffel, S. McCormick, J. Ruge, and L. Tang. Island coalescence using parallel first-order system least squares on incompressible resistive magnetohydrodynamics. *SIAM Journal on Scientific Computing*, 35:S171–S191, 2013.
- [2] S. Badia, A. Martín, and R. Planas. Block recursive LU preconditioners for the thermally coupled incompressible inductionless MHD problem. *submitted to Elsevier*, 2013.
- [3] S. Badia, R. Planas, and J. Gutiérrez-Santacreu. Unconditionally stable operator splitting algorithms for the incompressible magnetohydrodynamics system discretized by a stabilized finite element formulation based on projections. *International Journal for Numerical Methods in Engineering*, 93:302–328, 2013.
- [4] W. Bangerth, T. Heister, G. Kanschat, et al. `deal.II Differential Equations Analysis Library, Technical Reference`. <http://www.dealii.org>.
- [5] M. Benzi, G. Golub, and J. Liesen. Numerical solution of saddle point problems. *Acta Numerica*, 14:1–137, 2005.
- [6] M. Benzi and M. Olshanskii. An augmented Lagrangian-based approach to the Oseen problem. *SIAM Journal on Scientific Computing*, 28:2095–2113, 2006.
- [7] M. Benzi, M. Olshanskii, and Z. Wang. Modified augmented lagrangian preconditioners for the incompressible navier-stokes equations. *International Journal for Numerical Methods in Fluids*, 66:486–508, 2011.
- [8] P. Bochev, C. Garasi, J. Hu, A. Robinson, and R. Tuminaro. An improved algebraic multigrid method for solving Maxwell’s equations. *SIAM Journal on Scientific Computing*, 25:623–642, 2003.
- [9] J. Brown Jr. Mean square error in series expansions of random functions. *Journal of the Society of Industrial and Applied Mathematics*, 8:28–32, 1960.

- [10] R. Caflisch. Monte Carlo and quasi-Monte Carlo methods. *Acta Numerica*, 7:1–49, 1998.
- [11] L. Chacón. An optimal, parallel, fully implicit Newton-Krylov solver for three-dimensional viscoresistive magnetohydrodynamics. *Physics of Plasmas*, 15:056103, 2008.
- [12] L. Chacón. Scalable parallel implicit solvers for 3D magnetohydrodynamics. *Journal of Physics: Conference Series*, 125:012041, 2008.
- [13] L. Chacón, D. Knoll, and J. Finn. An implicit, nonlinear reduced resistive MHD solver. *Journal of Computational Physics*, 178:15–36, 2002.
- [14] R. Codina and N. Hernández. Stabilized finite element approximation of the stationary magneto-hydrodynamics equations. *Computational Mechanics*, 38:344–355, 2006.
- [15] R. Codina and N. Hernández. Approximation of the thermally coupled MHD problem using a stabilized finite element method. *Journal of Computational Physics*, 230:1281–1303, 2011.
- [16] M. Costabel and M. Dauge. Weighted regularization of Maxwell equations in polyhedral domains. *Numerische Mathematik*, 93:239–277, 2002.
- [17] E. Cyr, J. Shadid, R. Tuminaro, R. Pawlowski, and L. Chacón. A new approximate block factorization preconditioner for 2D incompressible (reduced) resistive MHD. *SIAM Journal on Scientific Computing*, 35:B701–B730, 2013.
- [18] P. Davidson. *An Introduction to Magnetohydrodynamics*. Cambridge University Press, 2001.
- [19] A. Dedner, F. Kemm, D. Kröner, C.-D. Munz, T. Schnitzer, and M. Wesenberg. Hyperbolic divergence cleaning for the MHD equations. *Journal of Computational Physics*, 175:645–673, 2002.
- [20] L. Demkowicz and L. Vardapetyan. Modeling of electromagnetic absorption/scattering problems using hp-adaptive finite elements. *Computer Methods in Applied Mechanics and Engineering*, 152:103–124, 1998.
- [21] J. Demmel, S. Eisenstat, J. Gilbert, X. Li, and J. Liu. A supernodal approach to sparse pivoting. *SIAM Journal on Matrix Analysis and Applications*, 2(3):720–755, 1999.
- [22] H. Elman, V. Howle, J. Shadid, R. Shuttleworth, and R. Tuminaro. A taxonomy and comparison of parallel block multi-level preconditioners for the incompressible Navier-Stokes equations. *Journal of Computational Physics*, 227:1790–1808, 2008.

- [23] H. Elman, C. Miller, E. Phipps, and R. Tuminaro. Assessment of collocation and Galerkin approaches to linear diffusion equations with random data. *International Journal for Uncertainty Quantification*, 1:19–34, 2011.
- [24] H. Elman, D. Silvester, and A. Wathen. *Finite Elements and Fast Iterative Solvers with Applications in Incompressible Fluid Dynamics*. Oxford University Press, 2005.
- [25] H. Elman and R. Tuminaro. Boundary conditions in approximate commutator preconditioners for the Navier-Stokes equations. *Electronic Transactions on Numerical Analysis*, 35:1068–9613, 2009.
- [26] M. Fortin and R. Glowinski. *Augmented Lagrangian Methods: Applications to the Numerical Solution of Boundary-Value Problems*. Elsevier Science Publishers B.V., 1983.
- [27] M. Gee, C. Siefert, J. Hu, R. Tuminaro, and M. Sala. ML 5.0 smoother aggregation user’s guide. Technical Report SAND2006-2649, Sandia National Laboratories, 2006.
- [28] J.-F. Gerbeau. A stabilized finite element method for the incompressible magnetohydrodynamics equations. *Numerische Mathematik*, 87:83–111, 2000.
- [29] J.-F. Gerbeau, C. Le Bris, and T. Lelièvre. *Mathematical Methods for the Magnetohydrodynamics of Liquid Metals*. Oxford University Press, 2006.
- [30] R. Ghanem. Scales of fluctuation and the propagation of uncertainty in random porous media. *Water Resources Research*, 34:2123–2136, 1998.
- [31] R. Ghanem and P. Spanos. *Stochastic Finite Elements: A Spectral Approach*. Dover Publications, 2003.
- [32] R. Glowinski and P. Le Tallec. *Augmented Lagrangian and Operator-Splitting Methods in Nonlinear Mechanics*. Society for Industrial and Applied Mathematics, 1989.
- [33] J. Goedbloed and S. Poedts. *Principles of Magnetohydrodynamics with Applications to Laboratory and Astrophysical Plasmas*. Cambridge University Press, 2004.
- [34] J. Gopalakrishnan, J. Pasciak, and L. Demkowicz. Analysis of a multigrid algorithm for time harmonic Maxwell equations. *SIAM Journal on Numerical Analysis*, 42:90–108, 2005.
- [35] C. Greif and D. Schötzau. Preconditioners for the discretized time-harmonic Maxwell equations in mixed form. *Numerical Linear Algebra with Applications*, 14:281–297, 2007.

- [36] M. Gunzburger, A. Meir, and J. Peterson. On the existence, uniqueness, and finite element approximation of solutions of the equations of stationary, incompressible magnetohydrodynamics. *Mathematics of Computation*, 56:523–563, 1991.
- [37] M. Heroux. AztecOO user guide. Technical Report SAND2004-3796, Sandia National Laboratories, 2004.
- [38] M. Heroux, R. Bartlett, V. Howle, R. Hoekstra, J. Hu, T. Kolda, R. Lehoucq, K. Long, R. Pawlowski, E. Phipps, A. Salinger, H. Thornquist, R. Tuminaro, J. Willenbring, A. Williams, and K. Stanley. An overview of the Trilinos project. *ACM Transactions on Mathematical Software*, 31:397–423, 2005.
- [39] R. Hiptmair and J. Xu. Nodal auxiliary space preconditioning in $H(\text{Curl})$ and $H(\text{Div})$ spaces. *SIAM Journal on Numerical Analysis*, 45:2483–2509, 2007.
- [40] P. Houston, D. Schötzau, and X. Wei. A mixed DG method for linearized incompressible magnetohydrodynamics. *Journal of Scientific Computing*, 40:281–314, 2009.
- [41] J. Hu, R. Tuminaro, P. Bochev, C. Garasi, and A. Robinson. Toward an h-independent algebraic multigrid method for Maxwell’s equations. *SIAM Journal on Scientific Computing*, 27:1669–1688, 2006.
- [42] Q. Hu and J. Zou. Substructuring preconditioners for saddle-point problems arising from Maxwell’s equations in three dimensions. *Mathematics of Computation*, 73:35–61, 2004.
- [43] A. Hujeirat. IRMHD: An implicit radiative and magnetohydrodynamical solver for self-gravitating systems. *Monthly Notices of the Royal Astronomical Society*, 298:310–320, 1998.
- [44] P. Lin, J. Shadid, R. Tuminaro, M. Sala, G. Hennigan, and R. Pawlowski. A parallel fully-coupled algebraic multilevel preconditioner applied to multiphysics PDE applications: drift-diffusion, flow/transport/reaction, resistive MHD. *International Journal for Numerical Methods in Fluids*, 64:1148–1179, 2010.
- [45] J. Liu, S. Tavener, and H. Chen. ELLAM for resolving the kinematics of two-dimensional resistive magnetohydrodynamic flows. *Journal of Computational Physics*, 227:1372–1386, 2007.
- [46] B. Marder. A method for incorporating Gauss’ law into electromagnetic PIC codes. *Journal of Computational Physics*, 68:48–55, 1987.
- [47] P. Mininni, A. Alexakis, and A. Pouquet. Shell-to-shell energy transfer in magnetohydrodynamics, Part II: Kinematic dynamo. *Physical Review E*, 72:046302, 2005.
- [48] R. Moreau. *Magnetohydrodynamics*. Kluwer Academic Publishers, 1990.

- [49] M. Murphy, G. Golub, and A. Wathen. A note on preconditioning for indefinite linear systems. *SIAM Journal on Scientific Computing*, 21:1969–1972, 2000.
- [50] J.-C. Nédélec. Mixed finite elements in \mathbb{R}^3 . *Numerische Mathematik*, 35:315–341, 1980.
- [51] J. Nocedal and S. Wright. *Numerical Optimization*. Springer-Verlag, 1999.
- [52] R. Pawlowski, J. Shadid, J. Simonis, and H. Walker. Globalization techniques for Newton-Krylov methods and applications to the fully-coupled solution of the Navier-Stokes equations. *SIAM Review*, 48(4):700–721, 2006.
- [53] M. Rao and R. Swift. *Probability Theory with Applications*. Springer, 2006.
- [54] S. Reitzinger and J. Schöberl. An algebraic multigrid method for finite element discretizations with edge elements. *Numerical Linear Algebra with Applications*, 9:223–238, 2002.
- [55] D. Reynolds, R. Samtaney, and H. Tiedeman. A fully implicit Newton-Krylov-Schwarz method for tokamak magnetohydrodynamics: Jacobian construction and preconditioner formulation. *Computational Science & Discovery*, 5:014003, 2012.
- [56] Y. Saad and M. Schultz. GMRES: A generalized minimal residual algorithm for solving nonsymmetric linear systems. *SIAM Journal on Scientific and Statistical Computing*, 7:856–869, 1986.
- [57] N. Ben Salah, A. Soulaïmani, and W. Habashi. A finite element method for magnetohydrodynamics. *Computational Methods in Applied Mechanics and Engineering*, 190:5867–5892, 2001.
- [58] A. Schneebeli and D. Schötzau. Mixed finite elements for incompressible magneto-hydrodynamics. *Comptes Rendus de l’Académie des Sciences - Series I*, 337:71–74, 2003.
- [59] D. Schötzau. Mixed finite element methods for stationary incompressible magneto-hydrodynamics. *Numerische Mathematik*, 96:771–800, 2004.
- [60] 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. *Journal of Computational Physics*, 229:7649–7671, 2010.
- [61] V. Shatrov, G. Mutschke, and G. Gerberth. Three-dimensional linear stability analysis of lid-driven magnetohydrodynamic cavity flow. *Physics of Fluids*, 15:2141–2151, 2003.
- [62] A. Wathen. Realistic eigenvalue bounds for the Galerkin mass matrix. *IMA Journal on Numerical Analysis*, 7:449–457, 1987.

- [63] S.-L. Wu, T.-Z. Huang, and C.-X. Li. Modified block preconditioners for the discretized time-harmonic maxwell equations in mixed form. *Journal of Computational and Applied Mathematics*, 237:419–431, 2013.
- [64] S.-L. Wu, T.-Z. Huang, and L. Li. Block triangular preconditioner for static Maxwell equations. *Computational & Applied Mathematics*, 30:589–612, 2011.
- [65] D. Xiu and J. Hesthaven. High-order collocation methods for differential equations with random inputs. *SIAM Journal on Scientific Computing*, 27:1118–1139, 2005.
- [66] D. Xiu and J. Shen. Efficient stochastic Galerkin methods for random diffusion equations. *Journal of Computational Physics*, 228:266–281, 2009.