

Towards a Scalable Fully-Implicit Fully-coupled Resistive MHD Formulation with Stabilized FE Methods^{*}

J.N. Shadid¹, R.P. Pawlowski¹, J.W. Banks², L. Chacón³, P.T Lin¹,
R.S. Tuminaro¹

Abstract

This paper explores the development of a scalable, nonlinear, fully-implicit stabilized unstructured finite element (FE) capability for 2D incompressible (reduced) resistive MHD. The discussion considers the implementation of a stabilized FE formulation in context of a fully-implicit time integration and direct-to-steady-state solution capability. The nonlinear solver strategy employs Newton-Krylov methods, which are preconditioned using fully-coupled algebraic multilevel preconditioners. These preconditioners are shown to enable a robust, scalable and efficient solution approach for the large-scale sparse linear systems generated by the Newton linearization. Verification results demonstrate the expected order-of-accuracy for the stabilized FE discretization. The approach is tested on a variety of prototype problems, including both low Lundquist number (e.g., an MHD Faraday conduction pump and a hydromagnetic Rayleigh-Bernard linear stability calculation) and moderately-high Lundquist number (magnetic island coalescence problem) examples. Initial results that explore the scaling of the solution methods are presented on up to 4096 processors for problems with up to 64M unknowns on a CrayXT3/4. Additionally, a large-scale proof-of-capability calculation for 1 billion unknowns for the MHD Faraday pump problem on 24,000 cores is presented.

1 Introduction

The magnetohydrodynamics (MHD) model describes the dynamics of charged fluids in the presence of electromagnetic fields. MHD models are used to describe important phenomena in the natural world (e.g., solar flares, astrophysical magnetic field generation, Earth's magnetosphere interaction with the solar wind) and in technological applications (e.g., spacecraft propulsion, magnetically confined plasma

^{*} This work was partially supported by DOE NNSA ASC Algorithms effort, the DOE Office of Science AMR program at Sandia National Laboratory under contract DE-AC04-94AL85000, Lawrence Livermore National Laboratory under contract DE-AC52-07NA27344, and Oak Ridge National Laboratory under contract DE-AC05-00OR22725

¹ Sandia National Laboratories, Albuquerque, New Mexico, 87185

² Lawrence Livermore National Laboratory, Livermore, California, 94551

³ Oak Ridge National Laboratory, Oak Ridge, TN 37831

for fusion energy devices such as tokamak reactors, and plasma dynamics in pulsed reactors such as Z-pinch devices) [1]. The mathematical basis for the continuum modeling of these systems is the solution of the governing partial differential equations (PDEs) describing conservation of mass, momentum, and energy, augmented by Maxwell’s equations for the electric and magnetic field. This system of PDEs is non-self adjoint, strongly coupled, highly nonlinear, and characterized by multiple physical phenomena that span a very large range of length- and time-scales. These interacting, nonlinear multiple time-scale physical mechanisms can balance to produce steady-state behavior, nearly balance to evolve a solution on a dynamical time scale that is long relative to the component time-scales, or can be dominated by just a few fast modes. These characteristics make the scalable, robust, accurate, and efficient computational solution of these systems over relevant dynamical time scales of interest (or to steady-state solutions) extremely challenging.

For multiple-time-scale systems, fully-implicit methods can be an attractive choice that can often provide unconditionally-stable time integration techniques [2,3]. The stability of these methods, however, comes at a cost, as these techniques generate large and highly nonlinear sparse systems of equations that must be solved at each time step. In the context of MHD, the dominant computational solution strategy has been the use of explicit [4–9] and partially implicit methods that include implicit-explicit [10–14], semi-implicit [15–19], and operator-splitting [20–22] time integration methods. With the exception of fully-explicit strategies, which are limited by stability restrictions to follow the fastest component time scale, all these temporal integration methods include some implicitness to enable a more efficient solution of MHD systems. Such implicitness is aimed at removing one or more sources of numerical stiffness in the problem, either from parabolic diffusion or from fast wave phenomena. While these types of techniques currently form the basis for most production-level resistive MHD simulation tools (see e.g. [19,13,23]), a number of outstanding numerical and computational issues remain. These include conditional stability limits, operator-splitting-type errors, and limited temporal orders of accuracy [17,18,24].

Recently, considerable progress has been made in the development of fully-implicit formulations that attempt to robustly and accurately integrate these systems while following the dynamical time-scales of interest [20,25,11,10,26–31]. In [20] for example, a nonlinear implicit MHD solver is proposed based on a implicit-operator-split (IOS) approach. The IOS algorithm employs Krylov solvers for required inversions in each split step, and is iterated upon in a Gauss-Seidel manner to achieve some degree of nonlinear consistency. However, for large implicit time steps, large numerical errors are reported to be possible with this algorithm in transient calculations unless enough nonlinear iterations are taken [20]. In [11,10], incomplete-LU-preconditioned Krylov methods are employed to invert the linearly implicit (and also implicit-explicit) set of MHD equations in the context of the Versatile Advection Code. CPU speedups vs. explicit of ~ 40 are reported [10]. An unpreconditioned Newton-Krylov solver for 3D compressible MHD is explored in [26], which also reports order-of-magnitude speedups vs. explicit approaches for fine enough grids. More recently, the same researchers have developed an “operator-based” parallel preconditioner for 3D MHD, based on directional splitting of the implicit operator followed by a characteristic decomposition of the resulting directional PDE operators [32]. The work in [27] explores a Newton-Krylov-Schwarz parallel approach for the reduced Hall MHD model, where gains of an order of magnitude with respect to explicit approaches and good parallel scalability are reported. Finally, [28–31] develop optimal “physics-based” preconditioning strategies for a fully implicit Newton-Krylov treatment of 2D and 3D extended MHD, and report excellent parallel scalability and algorithmic speedups ranging from one to two orders of magnitude.

In the specific context of stabilized finite element methods used for spatial discretizations for MHD, Salah et. al. [33] developed a formulation for the constant resistivity magnetic induction equation with a given velocity field. They use a Lagrange multiplier to enforce the solenoidal constraint on

the magnetic field, \mathbf{B} and thus develop a system with four unknowns, the three components of \mathbf{B} and the Lagrange multiplier, r . This system requires compatible spaces for \mathbf{B} and r that respect the Ladyzhenskaya-Babuska-Brezzi (LLB) condition, that is analogous to the velocity and pressure spaces (\mathbf{v}, p) in the Stokes flow system (see e.g. [34,35]). An inconsistent stabilized formulation [36] following Brezzi and Pitkaranta [37] is used to allow equal-order interpolation for \mathbf{B} and r . These methods were subsequently extended to an incompressible, constant resistivity MHD system, with both a Lagrange multiplier formulation for \mathbf{B} as described above, and a vector potential formulation [38]. The solution of the resistive MHD system in [33,38] uses an outer decoupled nonlinear solution strategy. This decoupled strategy solves the flow and magnetics system separately in each sub-step and couples the system by the outer iteration that can sub-cycle the component solves. An ILU preconditioned Newton-Krylov type solver is used for the flow equations and a direct sparse solver is used for the linear magnetics equation. Codina and Silva [39] developed a stabilized FE formulation for resistive MHD, in the curl form of the equations with constant properties, and develop stability parameters that handle the velocity-pressure coupling, the solenoidal constraint by a Lagrange multiplier method, and streamline upwind Petrov-Galerkin (SUPG) like terms that control oscillation due to convection effects. A fixed point nonlinear solution is employed to resolve the nonlinearities and no mention is made of the component linear solver(s) that are used. These authors present a coercivity result for the system that enables development of the stabilization parameters. Other studies of stabilized FE methods applied to resistive MHD systems include Gerbeau [40], who considered the coercivity of a stabilized FE formulation for steady state systems and employed a fixed-point (Picard) type nonlinear solution strategy, and Lankalapalli et. al. [41], who developed a vorticity-streamfunction vector potential formulation that uses a SUPG FE discretization and a fixed-point nonlinear solver with a GMRES iterative method.

The current study complements previous work by providing a robust, efficient, fully-coupled stabilized FE formulation for resistive MHD with solution methods that enable both fully-implicit transient and direct-to-steady-state solutions. Our solution method relies on inexact Newton-Krylov methods [42,43] to solve the resulting large-scale nonlinear algebraic systems. For preconditioning, we compare well-known variable-overlap, additive, one-level Schwarz domain-decomposition methods [44,45] with a relatively new algebraic multilevel technique employing a graph-based aggressive-coarsening aggregation method applied to the nonzero block structure of the Jacobian matrix [46,47]. The algebraic multilevel method effectively uses corrections that are computed by a sequence of coarse operators to accelerate the convergence of the iterative Krylov method on the fine mesh. Employing a multilevel preconditioner is intended to enable the development of scalable solution methods for MHD.

In this study, we focus on a 2D incompressible resistive MHD formulation with an associated internal energy equation. This formulation is suitable for high-Lundquist-number reduced MHD models [48–50] as well as low magnetic-Reynolds-number liquid metal MHD [51,52]. As in standard reduced MHD, the magnetic field dynamics is described in terms of a single component of the vector potential (which is also the magnetic flux function). Unlike standard reduced MHD, however, we choose to describe the flow in primitive variables and solve directly for the flow velocity-vector and pressure. We discretize the resulting MHD system using a simplified form of a consistently stabilized FE method [36], based on the general approach of Hughes et. al. (see [53] and references therein). The stabilized formulation circumvents the LBB condition for a compatible discretization of the saddle-point problem arising from discretization of the incompressible MHD equations using mixed finite element formulations. In particular, the stabilized FE formulation allows for equal-order interpolation of the incompressible MHD equations while avoiding spurious pressure modes and reducing oscillations for convection-dominated problems. In the current context of incompressible resistive MHD systems, an additional benefit of stabilized FE is the use of equal-order interpolation for all quantities. This simplifies the data structures of the parallel unstructured FE code, and the linear algebra interface for the iterative solution methods

that are employed.

The remainder of this paper is organized as follows. Section 2 presents our formulation of the resistive reduced MHD model. The stabilized FE formulation of the governing 2D resistive MHD equations is presented in Section 3. In Section 4 a brief overview of the fully-implicit Newton-Krylov solution method is presented with a discussion of the domain decomposition and multilevel preconditioners. Section 5 presents representative order-of-accuracy verification studies and performance, scaling and simulation results for some illustrative resistive MHD problems. Most of these problems are formulated in the low-Hartmann number regime, as our solution method is optimal for parabolic and elliptic problems. However, as discussed earlier, MHD becomes truly multiscale in the high-Lundquist-number regime. For this reason, we have also included in this study a high-Lundquist number test, the island coalescence problem. Future work will explore the implementation of physics-based preconditioning ideas to deal with the strongly hyperbolic MHD regime [28–31]. In physics-based preconditioning, the hyperbolic MHD system is effectively parabolized by a block factorization, thus rendering the associated systems amenable to the algebraic multigrid iterative technology demonstrated in the present study. Finally, Section 6 closes with a few conclusions.

2 Reduced MHD Model Equations

Our base MHD model is the one-fluid visco-resistive MHD system [1]. This model provides a continuum description of charged fluids in the presence of electromagnetic fields. Formally, visco-resistive MHD augments the Navier-Stokes fluid description with a Lorentz force term $\mathbf{J} \times \mathbf{B}$ in the momentum equation and a Joule heating source term $\eta \|\mathbf{J}\|^2$ in the specific internal energy equation. The system is closed with Faraday’s law, Ampere’s law, and the solenoidal constraint for the magnetic field. The resulting system of equations is:

$$\frac{\partial(\rho\mathbf{v})}{\partial t} + \nabla \cdot [\rho\mathbf{v} \otimes \mathbf{v} - \mathbf{T}] - \mathbf{J} \times \mathbf{B} - \mathbf{f}(T) = \mathbf{0} \quad (1)$$

$$\frac{\partial\rho}{\partial t} + \nabla \cdot [\rho\mathbf{v}] = 0 \quad (2)$$

$$\frac{\partial(\rho e)}{\partial t} + \nabla \cdot [\rho\mathbf{v}e + \mathbf{q}] - \mathbf{T} : \nabla\mathbf{v} - \eta \|\mathbf{J}\|^2 = 0. \quad (3)$$

$$\frac{\partial\mathbf{B}}{\partial t} - \nabla \times (\mathbf{v} \times \mathbf{B}) + \nabla \times \left(\frac{\eta}{\mu_0} \nabla \times \mathbf{B}\right) = \mathbf{0}. \quad (4)$$

Here \mathbf{v} is the plasma velocity; ρ is the ion mass density; \mathbf{T} is the stress tensor; \mathbf{B} is the magnetic field; T is the plasma temperature; $\mathbf{J} = \frac{1}{\mu_0} \nabla \times \mathbf{B}$ is the current density; \mathbf{f} is a body force term that we have allowed to vary with temperature (as in our example of thermal buoyancy driven flow in Sec. 5.2); e is the plasma specific internal energy; η is the resistivity; and μ_0 is the magnetic permeability of free space. In its simplest form, the MHD equations are closed with simple constitutive equations for the stress tensor, \mathbf{T} , and the heat flux \mathbf{q} ,

$$\mathbf{T} = -P\mathbf{I} + \mathbf{\Pi} = -\left(P + \frac{2}{3}\mu(\nabla \cdot \mathbf{v})\right)\mathbf{I} + \mu[\nabla\mathbf{v} + \nabla\mathbf{v}^T], \quad (5)$$

$$\mathbf{q} = -\chi\nabla T, \quad (6)$$

and a caloric equation of state, $e = \hat{C}_P T$. Here, P is the plasma pressure, $\mathbf{\Pi}$ is the viscous stress tensor, \mathbf{I} is an identity tensor, μ is the plasma viscosity, χ is the thermal conductivity, and \hat{C}_P is the specific heat at constant pressure.

For the purposes of this study, we focus on a 2D geometry in the incompressible limit ($\nabla \cdot \mathbf{v} = 0$). This limit is useful in the modeling of various applications such as low-Lundquist-number liquid-metal MHD flows [51,52], and high-Lundquist-number, large-guide-field fusion plasmas [48–50]. In the 2D incompressible regime, it can be shown that the in-plane and out-of-plane dynamics decouple (i.e., B_z , v_z , with z the ignorable direction, do not impact the evolution of the system in the $x - y$ plane). As a result, the system in (1) – (4) simplifies considerably, as it can be expressed in terms of a few scalar quantities like the vorticity component in the ignorable direction, the in-plane streamfunction, and the poloidal flux (or, alternatively, the vector potential component in the ignorable direction) [48–50]. For our implementation, however, it is of interest to keep a primitive description of the fluid flow, and to enforce the incompressibility constraint explicitly (as is often done in the CFD community [see [53] and references therein]). Thus, we preserve a 2D form of (1), and we replace the continuity equation (2) by $\nabla \cdot \mathbf{v} = 0$.

In regards to the magnetic field evolution equation, we replace (4) by an evolution equation for the vector potential component in the ignorable direction, $\mathbf{A} = (0, 0, A_z)$, which reads:

$$\frac{\partial A_z}{\partial t} + \mathbf{v} \cdot \nabla A_z - \frac{\eta}{\mu_0} \nabla^2 A_z + E_z^0 = 0. \quad (7)$$

This equation is completely equivalent to (4) in two dimensions, but exhibits a standard convection diffusion form that is more convenient for developing stabilized FE formulations.

Satisfying the solenoidal involution $\nabla \cdot \mathbf{B} = 0$ in the discrete representation to machine precision is a topic of considerable current interest in both structured and unstructured finite-volume and unstructured finite-element contexts (see e.g. [9,54,55]). It should be noted that the discrete FE representation of (7), despite describing the vector potential dynamics, does not automatically enforce the solenoidal constraint. In the case of linear nodal elements in 2D, this property can be shown to hold on element interiors. However the non-continuity of the normal derivatives at element edges in a C^0 FE approximation allows violation of this condition on the skeleton of the FE mesh, which is a set of measure zero. This implies that the divergence condition holds point-wise on element interiors and in an L_2 sense over any finite sub-region of the domain. In addition, it can be shown that at element edges $\mathbf{B} \cdot \hat{\mathbf{n}}$ is continuous, a condition that is also required on any surface in the continuous problem [56].

Finally, the desire to use low-order FE basis functions for the vector potential necessitates the consideration of the representation of the Lorentz force term in the momentum equation and the Joule heating term in the energy equation. In the context of the momentum equation, as presented in (1), the straightforward use of the vector potential term yields a second order operator $\mathbf{J} \times \mathbf{B} = (\frac{1}{\mu_0} \nabla \times \nabla \times \mathbf{A}) \times (\nabla \times \mathbf{A})$. Due to the nonlinear nature of the term and the higher order derivatives on \mathbf{J} this term is difficult to integrate by parts to sufficiently reduce the order of the weak form operator. This same issue is present in our 2D formulation, where we now have $J_z = -\frac{1}{\mu_0} \nabla^2 A_z$. To avoid this difficulty, an alternate divergence form (obtained using $\nabla \cdot \mathbf{B} = 0$) is employed to compute the Lorentz force as

$$\mathbf{J} \times \mathbf{B} = \nabla \cdot \left[\frac{1}{\mu_0} \mathbf{B} \otimes \mathbf{B} - \frac{1}{2\mu_0} \|\mathbf{B}\|^2 \mathbf{I} \right]. \quad (8)$$

However, such a reformulation is not possible for the Joule heating source term in the specific internal energy (3). Presently, this term is approximated by computing the curl of a field, $\hat{\mathbf{B}}$ that is obtained from an L_2 projection of the piecewise discontinuous approximation obtained from $\mathbf{B} = \nabla \times \mathbf{A}$.

3 A Stabilized Finite Element Formulation for 2D Resistive MHD

Table 1 presents the governing equations, for momentum, total mass, temperature, and vector potential in residual notation. The continuous PDE problem, defined by the 2D resistive MHD equations in Table 1, is approximated by a stabilized FE formulation. This formulation allows for stable equal-order velocity-pressure interpolation and provides for convection stabilization, as described below.

We employ stabilized FE methods to avoid stability and algorithmic limitations of mixed Galerkin FE formulations. In particular, in a mixed Galerkin FE formulation of the momentum-continuity equations of the Navier-Stokes part of the MHD system, there is a stability requirement that the discrete spaces satisfy the the Ladyzhenskaya-Babuska-Brezzi (LBB) stability condition (see e.g. [34] or [35]). This condition prevents the use of equal-order finite element spaces, defined with respect to the same partition of the computational domain in finite elements. In addition the linearization of the nonlinear mixed Galerkin FE formulation also leads to indefinite linear systems, which are more difficult to solve by iterative methods (see the Appendix - Section 8.2). Finally, an additional difficulty is that the mixed Galerkin formulation is prone to instabilities for highly convected flows, even if the LBB condition is satisfied by the finite element spaces.

Consistently stabilized⁴ finite element methods for Navier-Stokes address the issues above by using a combination of properly weighted residuals of the governing balance equations. These methods simultaneously relax the incompressibility constraint and add streamline-diffusion (and sometimes nonlinear discontinuity-capturing-type operators) to the weak equations to limit oscillations in highly convected flows [57–61]. The specific stabilized FE formulations employed in this study are shown in Table 2. The intrinsic-time-scale stability parameters ($\hat{\tau}_m$, $\hat{\tau}_T$, and $\hat{\tau}_{A_z}$) are based on the formulations of Hughes and Mallet [62], Shakib [63], Hughes [64], and Tezduyar [65] for Navier-Stokes with an adaptation of the stabilized formulation of Codina and Hernandez-Silva [39] for a resistive MHD system. The definition of the stabilization parameters are provided in the Appendix in Table 8 for momentum, temperature, and the vector potential.

It should be noted that our stabilized FE formulation, as presented, has some caveats. Firstly, the definition of the least squares operators is a simplification of proposed formulations for system advection-diffusion-type equations [62]. And secondly, it features no stabilization contribution for strong source terms, and no nonlinear discontinuity-capturing terms (see [53] and references therein).

4 Fully-implicit Fully-coupled Solution by Parallel Newton-Krylov Methods

4.1 Fully-implicit Time Integration and Direct to Steady-state Solutions

For stiff (multiple-time-scale) PDE systems such as MHD, fully-implicit methods are an attractive choice that can often provide unconditionally-stable time integration techniques. These methods can be designed with various types of stability properties that allow robust integration of multiple timescale systems without the requirement to resolve the stiff modes of the system (which are not of interest since they do not control the accuracy of time integration [2,3]). In this study we employ a second-order A-stable implicit midpoint rule [2,3].

The result of a fully-implicit and direct-to-steady-state solution technique is the development of very large-scale, coupled highly nonlinear system(s) that must be solved. Therefore, these techniques place a heavy burden on both the nonlinear and linear solvers and require robust, scalable, and efficient nonlinear solution methods. In this study, Newton-based iterative nonlinear solvers [66] are employed

⁴ Consistent in the sense that the exact solution to the PDE equation satisfies the weak-form residual equations.

Momentum	$\mathbf{R}_m = \rho \frac{\partial \mathbf{v}}{\partial t} + \rho(\mathbf{v} \cdot \nabla \mathbf{v}) + \nabla \cdot \left(-\frac{1}{\mu_0} \mathbf{B} \otimes \mathbf{B} - \mathbf{\Pi} + (P + \frac{1}{2\mu_0} \ \mathbf{B}\ ^2) \mathbf{I} \right) - \mathbf{f}$
Total Mass	$R_P = \nabla \cdot \mathbf{v}$
Temperature	$R_T = \rho \hat{C}_p \frac{\partial T}{\partial t} + \rho \hat{C}_p (\mathbf{v} \cdot \nabla T) + \nabla \cdot \mathbf{q} - \eta \ \mathbf{J}\ ^2$
2D Vector Potential Eq.	$R_{A_z} = \frac{\partial A_z}{\partial t} + \mathbf{v} \cdot \nabla A_z - \frac{\eta}{\mu_0} \nabla^2 A_z + E_z^0$ $\mathbf{B} = \nabla \times \mathbf{A}; \quad \mathbf{A} = (0, 0, A_z)$

Table 1

Residual form of governing low Mach number resistive MHD equations with the 2D form of the vector potential evolution equation in advection-diffusion form. The primitive variables are the velocity vector \mathbf{u} , the hydrodynamic pressure P , the temperature T , and the A_z component of the vector potential in 2D. \hat{C}_p is the specific heat at constant pressure.

Momentum	$\mathbf{F}_{m,i} = \int_{\Omega} \Phi \mathbf{R}_{m,i} d\Omega + \sum_e \int_{\Omega_e} \rho \hat{\tau}_m (\mathbf{v} \cdot \nabla \Phi) \mathbf{R}_{m,i} d\Omega$
Total Mass	$F_P = \int_{\Omega} \Phi R_P d\Omega + \sum_e \int_{\Omega_e} \rho \hat{\tau}_m \nabla \Phi \cdot \mathbf{R}_m d\Omega$
Temperature	$F_T = \int_{\Omega} \Phi R_T d\Omega + \sum_e \int_{\Omega_e} \rho \hat{C}_p \hat{\tau}_T (\mathbf{v} \cdot \nabla \Phi) R_T d\Omega$
Z-component Vector Potential	$F_{A_z} = \int_{\Omega} \Phi R_{A_z} d\Omega + \sum_e \int_{\Omega_e} \hat{\tau}_{A_z} (\mathbf{v} \cdot \nabla \Phi) R_{A_z} d\Omega$

Table 2

Stabilized finite element formulation of transport/reaction PDEs, where the residual equations R_i are presented in Table 1 and the stabilization parameters $\hat{\tau}_i$ are defined in the Appendix in Table 8. Here Φ is a global weighting function used to formally define the weak form. Here the sum \sum_e indicates the integrals are taken only over element interiors Ω_e and integration by parts is not performed.

to solve the challenging nonlinear systems that result in this application. These solvers can exhibit quadratic convergence rates independently of the problem size when sufficiently robust linear solvers are available. For the latter, we employ Krylov iterative techniques. A Newton-Krylov (NK) method [67,42] is an implementation of Newton's method in which a Krylov iterative solution technique is used to approximately solve the linear systems that are generated at each step of Newton's method. Specifically, to solve the nonlinear system $\mathbf{F}(\mathbf{U}) = \mathbf{0}$, we seek a zero of $\mathbf{F} : R^N \rightarrow R^N$ where $\mathbf{U} \in R^N$ is a current approximate solution. The Krylov iterative solver is applied to the linearized residual equation

$$\mathbf{J}_k \mathbf{s}_{k+1} = -\mathbf{F}_k, \quad (9)$$

where \mathbf{J}_k is the Jacobian matrix and \mathbf{F}_k is the nonlinear residual, both of which are evaluated at the previous Newton step solution \mathbf{U}_k . The solution for the Newton direction vector, \mathbf{s}_{k+1} , is used to update the previous solution in the sequence as $\mathbf{U}_{k+1} = \mathbf{U}_k + \theta \mathbf{s}_{k+1}$, where θ is a step length reduction or backtracking parameter. Backtracking algorithms improve the robustness of the nonlinear solver by scaling the Newton correction vector by the parameter, θ , chosen to ensure a sufficient reduction of the nonlinear residual before the step is accepted [66,68].

For efficiency, an inexact Newton method [69–71] is usually employed, whereby one approximately solves (9) by choosing a forcing term η_{k+1} and stopping the Krylov iteration when the inexact Newton condition

$$\|\mathbf{F}_k + \mathbf{J}_k \mathbf{s}_{k+1}\| \leq \eta_{k+1} \|\mathbf{F}_k\| \quad (10)$$

is satisfied. In general, nonlinear residual information is used to determine the forcing η_{k+1} . In our Newton’s method implementation, a constant value is used for η_{k+1} .

In this study, the Jacobian matrix \mathbf{J}_k that is used for the Jacobian-vector products in the Krylov solvers, and as the basis for computing the preconditioners described in Section 4.2, is developed from automatic differentiation (AD) techniques. These methods are applied to the programmed functions representing the weak form residuals outlined in Table 2 by employing the SACADO package from the Trilinos framework [72]. The specific form of this discrete system is discussed in the Appendix in Section 8.2, along with the impact of the block structure of the Jacobian and the implications for the design of solution methods for this system of equations.

4.2 Schwarz Domain Decomposition and Multilevel Preconditioners

For the considered class of linear systems described above, convergence is not achieved without preconditioning due to ill-conditioning in the underlying matrix equations [73]. In the current work, we consider Schwarz domain decomposition preconditioners, where the basic idea is to decompose the computational domain Ω into overlapping subdomains Ω_i and then assign each subdomain to a different processor [44,45]. One application of the algorithm consists of solving on subdomains and then combining these local solutions to construct a global approximation throughout Ω . The i^{th} subdomain problem is usually defined by enforcing homogeneous Dirichlet boundary conditions on the subdomain boundary, $\partial\Omega_i$. In the minimal overlap case, the algebraic Schwarz method corresponds to block Jacobi where each block contains all degrees of freedom (DOFs) residing within a given subdomain. Convergence is typically improved by introducing overlap, which can be done recursively. Incomplete factorization, $ILU(k)$, is employed to approximate the solution of the local Dirichlet problems and avoid the large cost of direct factorization [73,74]. We note that the one-level preconditioner is black-box in that the overlapping subdomain matrices are constructed completely algebraically.

One possible drawback of the one-level Schwarz method is its locality. A single application of the algorithm transfers information between neighboring sub-domains. This implies that many repeated applications are required to combine information across the entire domain. Thus, as the number of subdomains increases, the convergence rate deteriorates for standard elliptic problems due to the lack of global coupling [44]. The convergence rate also deteriorates as the number of unknowns per subdomain increases when $ILU(k)$ is used for a subdomain solver. The results that we present in Section 5.2 will demonstrate such deterioration for the specific problems of interest in our study. To improve algorithmic performance, coarse levels can be introduced to approximate global coupling [75,76]. The use of a coarse mesh to accelerate the convergence of a one-level Schwarz preconditioner is similar in principle to the use of a sequence of coarser meshes in multigrid methods [77,78].

In this paper, only algebraically generated coarse levels are considered. These are significantly easier to implement and integrate with a complicated unstructured simulation than geometric coarse grids [75,76,79]. In particular, there is no need to represent complex geometric features on all levels, e.g. faces, edges, and corners to define the domain boundary. There is also no requirement to interact with the simulation’s geometric data structures. Most algebraic multigrid methods (AMG) associate a graph with the matrix system being solved. Graph vertices correspond to matrix rows for scalar PDEs, while for PDE systems it is natural to associate one vertex with each nodal block of unknowns, e.g. velocities and pressures at a particular grid point. A graph edge exists between vertex i and j if

there is a nonzero in the block matrix which couples i 's rows with j 's columns or j 's rows with i 's columns. In some situations, it may be advantageous to omit edges if all entries within the coupling block are small [80]. In this study, METIS and ParMETIS [81] are used to group fine mesh vertices into aggregates so that each aggregate effectively represents a coarse mesh vertex. These graph partitioning packages subdivide the matrix graph so that each partition has no more nodes than a user supplied parameter and that each partition is somewhat spherically shaped. This graph partitioning is then applied recursively until the user specified number of levels has been achieved. Once the coarse mesh is determined, an initial grid transfer is constructed corresponding to piecewise constant interpolation. The grid transfer matrix, P , contains only zeros and ones. In the scalar PDE case, P_{ij} equals one only if the i^{th} fine grid point has been assigned to the j^{th} aggregate. Within a PDE system, the grid transfer is a block system with an identity matrix for the $(i, j)^{th}$ block if the i^{th} fine grid point has been assigned to the j^{th} aggregate. This initial grid transfer can then be improved by smoothing the corresponding basis functions [82,83]. In this study, we employ a Petrov-Galerkin smoothed aggregation algorithm as implemented in the ML multilevel packages in Trilinos and described in [84].

Finally we note that we orient the graph partitioning algorithm so that they generate somewhat larger aggregates than those typically used in standard smoothed aggregation. This aggressive coarsening significantly reduces the number of unknowns between consecutive levels. This generally limits the total number of levels (≤ 5) which we find better suited for parallel computations [46,85]. Additionally, larger aggregates are consistent with using a sub-domain solver based on Schwarz/ $ILLU(k)$ which in the multigrid context corresponds to a somewhat heavyweight smoother (compared to Gauss-Seidel often used in standard multigrid). That is, one can coarsen more aggressively when a more substantial smoother is employed. The same $ILLU(k)$ algorithm is used as a smoother on each level and on the coarsest level the KLU [86,87] sparse direct solver is employed.

The multilevel implementation described above is provided by ML [75,88]. ML has been employed successfully in a number of applications that include stabilized FE discretizations of Navier-Stokes and transport-reaction type systems [89,79,47], as well as drift-diffusion systems for semi-conductor modeling [90,91]. Aztec, Ifpack, KLU, ML, KLU, and Zoltan are available through the Trilinos framework [72].

5 Results and Discussion

In this study we present two general classes of results. The first class of results, presented in Section 5.1, are for assessment of the accuracy and order-of-accuracy of the stabilized FE formulation described above by a comparison with analytic solutions. The problems in this section include, (1) a flux expulsion problem that qualitatively evaluates the accuracy of the FE formulation in solving the vector potential transport equation on a fully unstructured mesh; (2) a steady-state MHD duct flow with strong coupling between the momentum and the Lorentz forces that provides a detailed assessment of the order-of-accuracy of the solution and a discussion of the effectiveness of the SUPG convection stabilization; and (3) a transient MHD Rayleigh flow and Alfvén wave propagation problem that allows evaluation of the space-time accuracy and order-of-accuracy of the stabilized FE formulation.

The second class of results, presented in Section 5.2, are intended to present initial representative results for the robustness, efficiency and scalability of the fully-coupled AMG preconditioned NK solution method described above. The problems in this section include, (1) a MHD duct flow that models an idealized MHD Faraday pump and is used to study the scalability of the solvers in the context of a direct-to-steady-state solution method, (2) a classic MHD Hydromagnetic Rayleigh-Bernard problem provides a test of the robustness of the direct-to-steady-state solver and also presents another comparison of the accuracy of the formulation with an analytic linear stability result by Chandrasekhar, and

(3) a transient magnetic island coalescence problem, of recent scientific interest, for which we present initial results for the scalability of the solution methods.

5.1 Representative Verification and Order-of-Accuracy Results

In this subsection, detailed numerical verification results are presented for a set of 2D resistive MHD problems that admit analytic solutions. Two nonlinear convergence criteria are used to ensure that the numerical solution error is below discretization error. The first is a sufficient reduction in the relative nonlinear residual norm, $\|F_k\|/\|F_o\| < 10^{-2}$. In general, this requirement is easily satisfied. The second convergence criterion is based on a sufficient decrease of a weighted norm of the Newton update vector. The latter criterion requires that the correction, $\Delta\chi_i^k$, for any variable, χ_i , is small compared to its magnitude, $|\chi_i^k|$, and is given by

$$\sqrt{\frac{1}{N_u} \sum_{i=1}^{N_u} \left[\frac{|\Delta\chi_i|}{\varepsilon_r |\chi_i| + \varepsilon_a} \right]^2} < 1,$$

where N_u is the total number of unknowns, ε_r is the relative error tolerance between the variable correction and its magnitude, and ε_a is the absolute error tolerance of the variable correction, which sets the magnitude of components that are to be considered to be numerically zero. In this study, the relative-error and absolute-error tolerance are 10^{-8} and 10^{-10} respectively. Each linear system in Newton's method is solved to a level of $\eta = 10^{-8}$.

5.1.1 Flux Expulsion Problem

This problem admits an analytical solution, which is commonly used for verification of resistive MHD implementations [92,38]. The problem consists of an infinitely long, rotating cylindrical conductor immersed in a conducting medium in a uniform magnetic field $(B_0, 0, 0)$. The cylinder rotates with the velocity field $\mathbf{v} = (-\omega_0 y, \omega_0 x, 0)$, where ω_0 is a constant rate of rotation. The flow is zero, $\mathbf{v} = \mathbf{0}$, outside the cylinder. The dependent variable is $\mathbf{A} = (0, 0, A_z)$. The geometry is taken as a cylinder of radius, $r_0 = 0.2$ within a computational domain $\Omega = [-2, 2] \times [-1, 1]$. The analytic solution for the vector potential is given, in terms of Bessel functions as:

$$A_z = \Im\mathbf{m}[B_0 f(r) e^{i\theta}],$$

where

$$f(r) = r + \frac{C}{r}, \quad r > r_0$$

$$f(r) = DJ_1(pr), \quad r \leq r_0$$

and

$$p = \frac{(1-i)k_0}{\sqrt{2}}, \quad k_0 = \frac{\sqrt{Re_m}}{r_0}$$

$$C = \frac{r_0[2J_1(pr_0) - pr_0 J_0(pr_0)]}{pJ_0(pr_0)}, \quad D = \frac{2}{pJ_0(pr_0)}.$$

The computational solution is produced on an unstructured mesh, a coarse example of which is presented in Figure 1. The boundary conditions on all edges of the computational domain are taken from the analytic solution. As the rotational rate is increased in the conductor, the linear A_z (uniform B_x) field is distorted. At high rotation rates, the field is expelled from the cylinder and the effects of diffusion modify the flux field in the region about the rotating conductor. Figure 1 shows contours of A_z as a function of the magnetic Reynolds number, $Re_m = \omega_0 r_0^2 / \eta$. In this study, the resistivity $\eta = 1.0$

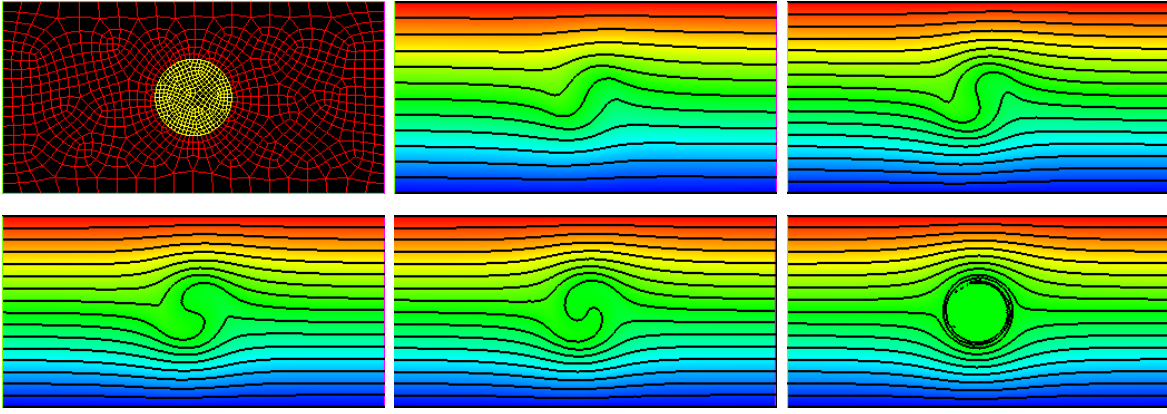


Fig. 1. Flux expulsion verification problem. The top left figure shows a coarse computational unstructured mesh $\Omega = [-1, 1] \times [-\frac{1}{2}, \frac{1}{2}]$. The remaining figures show filled-color contour plots of numerical solution of A_z for $Re_m = 6, 12, 24, 48, 12288$ from the top center image to the lower right image. Red values are high and blue values are low in each image.

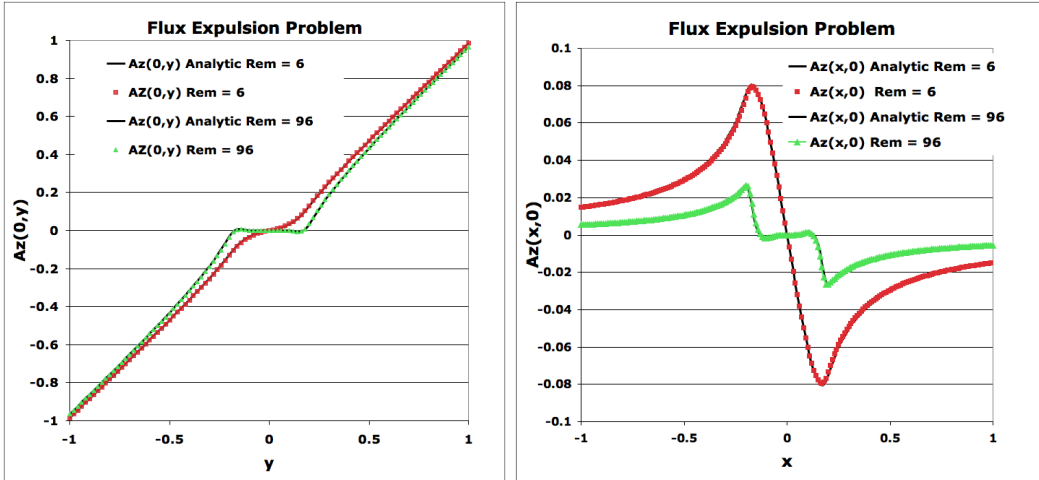


Fig. 2. Flux expulsion verification problem. Figures show profiles of analytic solution of $A_z(0, y)$ (left) and $A_z(x, 0)$ (right) along with numerically computed solutions from the vector potential solution in $\Omega = [-2, 2] \times [-1, 1]$ with 20K unstructured quad elements. The red squares are for $Re_m = 6$ and the green triangles are for the more highly convected case of $Re_m = 96$, the black line is the analytic solution due to Moffat [92].

and the rotational rate, ω_0 , is selected to achieve the desired magnetic Reynolds number. The noisy solution for the highest Reynolds number case is due to the highly convected flow being solved on a very coarse unstructured mesh. The SUPG-type stabilization does decrease the unphysical oscillations but it is not a monotonicity-preserving method. Section 5.1.2 presents a more careful study of the effectiveness of the SUPG stabilization for controlling unphysical oscillations due to convection effects. A comparison of the analytical solution for the flux expulsion problem with the numerical solution is provided in Figure 2, where $A_z(x, 0)$ and $A_z(0, y)$ are shown in a region about the conductor. Figure 3 presents the results for the post-processed variables $B_x(0, y)$, $B_y(0, y)$ and $J_z(0, y)$ computed on the 20K element unstructured mesh. There has been no significant optimization of the mesh to resolve the highly localized induced magnetic field and plasma current internal layers near the rotating inner cylinder. These comparisons indicate a very good agreement between the computed solutions and the exact analytical solution of Moffat [92] on a reasonably coarse unstructured mesh.

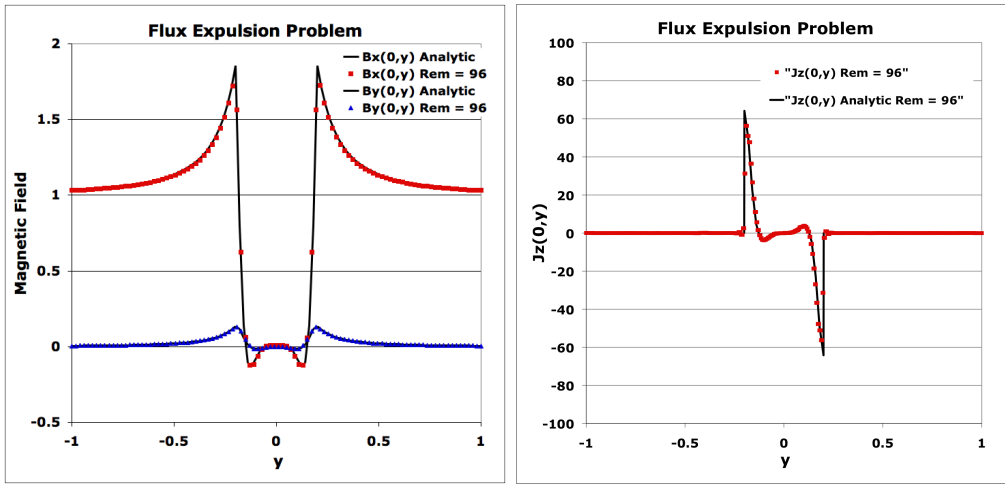


Fig. 3. Flux expulsion verification problem. Figures show profiles of analytic solution of $B_x(0, y)$ and $B_y(0, y)$ (left) along with $J_z(0, y)$ (right) along with numerically computed solutions from the vector potential solution in $\Omega = [-2, 2] \times [-1, 1]$ with 20K unstructured quad elements. The black line is the analytic solution due to Moffat [92].

5.1.2 A Modified Hartmann Flow Vector Potential Problem

This problem also admits an analytical solution [51] which is often used for MHD verification [33,93]. In the classical form, the solution is one-dimensional, with dependent variables $(\mathbf{v}, P, \mathbf{B})$. The geometry is taken as a square box with $(x, y) \in [-L, L] \times [-L, L]$ with an assumed pressure gradient that drives the flow of $\frac{\partial P}{\partial x} = -G_0$. The analytic asymptotic solution is of the form $\mathbf{B} = (B_x, B_0, 0)$ and $\mathbf{v} = (v_x, 0, 0)$, where $B_y = B_0$ is an applied external magnetic field whose tension retards the flow. The solution is given by

$$v_x = -\frac{\rho G_0 Ha}{\mu_0 B_0^2} \left[\frac{\cosh(Ha) - \cosh(y Ha/L)}{\sinh(Ha)} \right]$$

$$B_x = -\frac{B_0 Re_m}{Ha} \left[\frac{\sinh(y Ha/L) - (y/L) \sinh(Ha)}{\cosh(Ha) - 1} \right].$$

In vector potential form, the solution for the modified Hartmann problem reads:

$$A_z = -B_0 x - \frac{G_0 y^2}{2B_0} + \frac{G_0}{B_0} [Ha \cosh(y Ha) \operatorname{csch}(Ha)],$$

and must be sustained by an external electric field with

$$E_z^0 = \frac{G_0}{B_0} [Ha \coth(Ha) - 1].$$

In Figure 4, plots of the spatial variation of the numerical solution for v_x and A_z are presented along with the post-processed quantity B_x . The 1D nature of v_x and B_x are evident along with the 2D variation of the vector potential A_z . The order of accuracy of the numerical formulation is computed by comparison with the analytic solution. As before, the analytic boundary conditions are used on computational boundaries. In Figure 5, profiles from the computed numerical solution are compared with the analytic solution for v_x and B_x at $x = 0$ for various values of $Re = U(2L)/\nu = Re_m = \mu_0 U(2L)/\eta$, where U is the maximum x-direction velocity, and the Hartmann number is $Ha = B_0(2L)/\sqrt{\rho\nu\eta}$. In this study, we have taken $L = 1$, $\rho = 1$, $\nu = \eta = 1$, $\mu_0 = 1$. The values of G_0 and B_0 are then selected to produce the desired Reynolds and Hartmann numbers. Figure 5 shows excellent agreement with the analytical solution. Figure 6 shows a detailed spatial convergence study for a mesh with spacing

$\Delta x = \Delta y$. The expected second-order accuracy for v_x and A_z , which are interpolated with linear nodal elements, is confirmed by these results. B_x , which is computed by taking the curl of the vector potential and then reconstructed at the nodes by the a lumped mass L_2 projection (as stated in Section 2), also shows second-order convergence. Formally, derivative quantities computed from linear nodal basis functions are only first-order accurate. However, on a uniform mesh, the lumped mass projection recovers a second-order accurate approximation, demonstrating super-convergence [94]. In general, super-convergence behavior does not occur for fully-unstructured meshes.




Fig. 4. Modified Hartmann flow verification problem, $Re = Re_m = 1$ and $Ha = 1$. Figures show color plots of spatial variation of numerical solution of v_x (left), A_z (center), and B_x (right) to visualize the computed solution that corresponds to the analytic solution. Red values are high and blue values are low in each image.

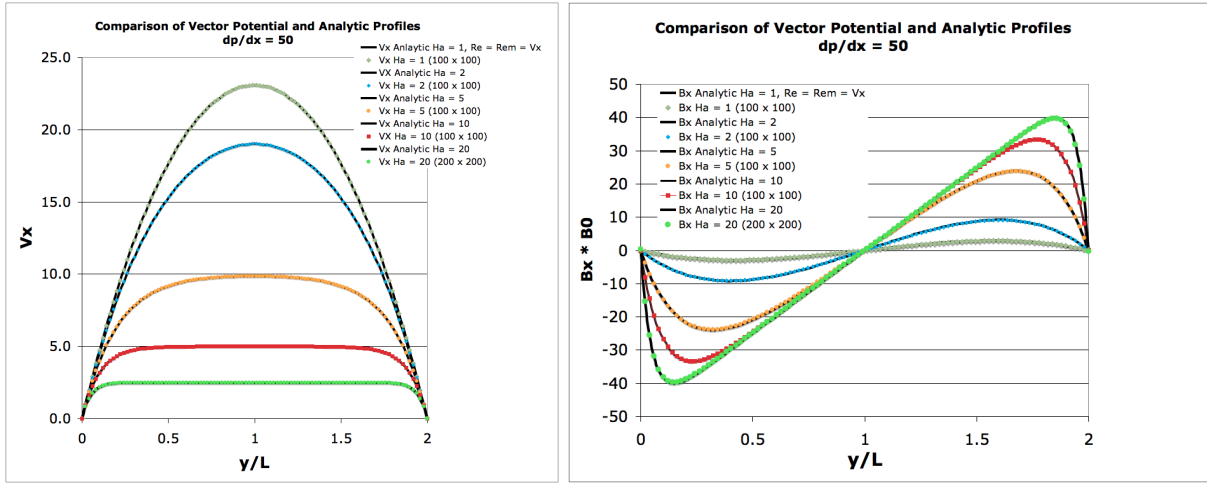


Fig. 5. Modified Hartmann flow verification problem, $Re = Re_m$. Figures show profiles of analytic solution of v_x (left) and $B_0 B_x$ (right) along with numerically computed solutions at $x = 0.0$ from the vector potential solution.

Finally, we proceed to illustrate the effect of the convection stabilization by the SUPG-type terms for a strongly convecting flow. The geometry is a rectangular channel with $(x, y) \in [0, 5] \times [-1, 1]$. A very coarse mesh of 10×20 elements is employed. The inlet has velocity $v_x(0, y) = V_0$ with $V_0 = 100, 200, 400$ and as above, we have taken $L = 1, \rho = 1, \nu = \eta = 1, \mu_0 = 1$. The boundary conditions on A_z consist of no-flux (natural) boundary conditions on the $y = -1, 1$ boundaries, and $A_z(0, y) = -5$ and $A_z(5, y) = 5$. The outflow conditions on the momentum equation are taken as the surface integral terms that naturally arise from integration by parts of the momentum equation in Table 1. In the surface integrals, the viscous stress is neglected (a standard traction-free outflow boundary condition assumption), and the magnetic stress and fluid pressure are integrated values that are computed just

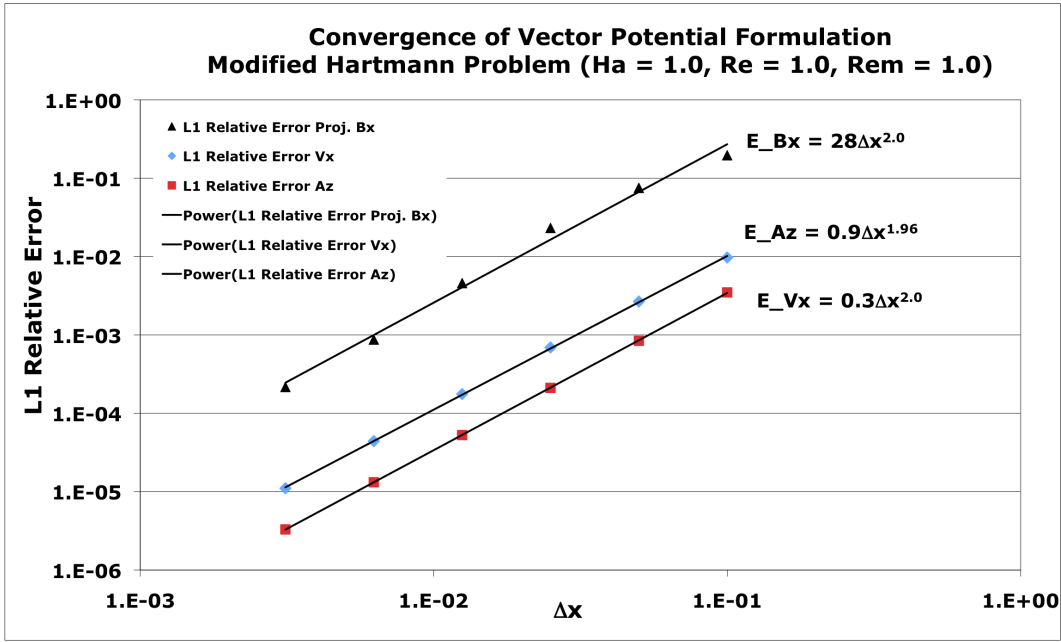


Fig. 6. Modified Hartmann flow verification problem, $Re = Re_m$. Figure shows L1 relative error of the computed numerical solution for v_x , A_z and the derived quantity B_x from the vector potential solution. The relative error is computed point-wise to the exact solution.

inside the exit boundary. These approximate outflow boundary conditions attempt to truncate the domain while still providing a reasonable test of SUPG's ability to control oscillations in the presence of strong convecting flows on coarse meshes. For a cell based $Re^c = U\Delta x/\nu = Re_m^c = \mu_0 U\Delta x/\eta > 1$, the Galerkin discretization can become unstable and produce oscillations. Figure 7 clearly illustrates such nonphysical behavior for large flows. As expected, the SUPG stabilization helps to control and localize these oscillations. Note that, for large velocities, the Dirichlet condition at $x = 5$ results in a very sharp flow boundary layer. This layer is not a numerical artifact, and in fact demonstrates the SUPG discretization is adequately capturing the physics. While there still appears to be some very slight outflow boundary condition effects, seen in the velocity profile for the stabilized case at the last two points, this most likely is due to the use of an inaccurate lumped mass projection to recover the high-order derivatives that are used to reconstruct terms in the PDE residual for the stabilized formulation. This approximate projection technique has problems at boundaries. A consistent mass projection is being explored.

5.1.3 MHD Rayleigh Flow and Alfvén Wave Propagation Problem

This prototype problem also features an analytical solution [51] which has been used for MHD verification [33]. In the classical form, this is a 1D transient problem that has an infinite plate bounding a conducting fluid in a semi-infinite domain above $y = 0$. There is an externally applied magnetic field in the y -direction with magnitude B_0 . The plate is initially at rest and then is suddenly set in motion with a velocity $v_x = U$. A viscous boundary layer flow is developed where the velocity profile is modified from the classical Rayleigh flow profile [95] by the existence of the magnetic field. As the velocity profile is developed, a self-induced magnetic field in the x -direction, B_x , is developed and an Alfvén wave with velocity $A_0 = B_0/\sqrt{\mu_0\rho}$ propagates into the fluid. In the numerical solution, the infinite half-space is approximated as a square box with $(x, y) \in [0, 5] \times [0, 5]$. For the case where the magnetic Prandtl number $Pr_m = \nu/\eta = 1$, the analytic asymptotic solution is of the form $\mathbf{B} = (B_x, B_0, 0)$ and

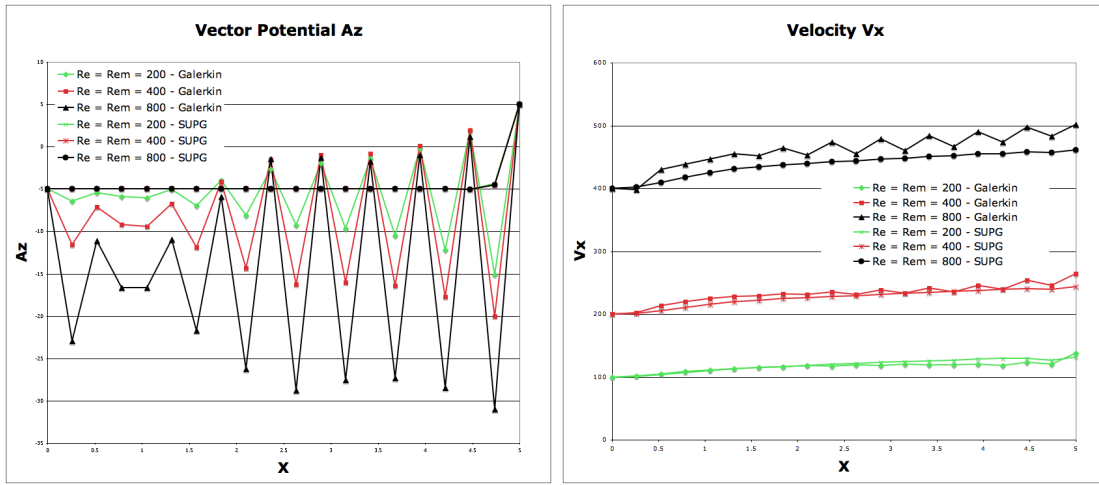


Fig. 7. Test of the SUPG-type stabilization of convective effects. Figures show profiles of A_z (left) and v_x (right). The Hartmann number for this case is 4. Note that in the left figure the SUPG results are indistinguishable for all Re_m .

$\mathbf{v} = (v_x, 0, 0)$, and is given by

$$v_x = \frac{1}{4}U \left[e^{-\frac{A_0 y}{d}} \left(1 - \operatorname{erf} \left(\frac{y - A_0 t}{2\sqrt{dt}} \right) \right) - \operatorname{erf} \left(\frac{y - A_0 t}{2\sqrt{dt}} \right) \right] \\ + \frac{1}{4}U \left[e^{\frac{A_0 y}{d}} \left(1 - \operatorname{erf} \left(\frac{A_0 t + y}{2\sqrt{dt}} \right) \right) - \operatorname{erf} \left(\frac{A_0 t + y}{2\sqrt{dt}} \right) + 2 \right] \\ B_x = -\frac{1}{4}e^{-\frac{A_0 y}{d}} \left(-1 + e^{\frac{A_0 y}{d}} \right) U \sqrt{\mu\rho} \left(\operatorname{erfc} \left(\frac{y - A_0 t}{2\sqrt{dt}} \right) + e^{\frac{A_0 y}{d}} \operatorname{erfc} \left(\frac{A_0 t + y}{2\sqrt{dt}} \right) \right).$$

In terms of the vector potential, the solution (\mathbf{v}, P, A_z) reads:

$$A_z = -B_0 x + \frac{U\sqrt{dt}\sqrt{\mu\rho}}{2\sqrt{\pi}} \left(e^{-\frac{(y-A_0 t)^2}{4dt}} - e^{-\frac{(y+A_0 t)^2}{4dt}} \right) \\ + \frac{U}{4} \frac{\sqrt{\mu\rho}}{A_0} (d + A_0^2 t) \left(\operatorname{erf} \left(\frac{A_0 t - y}{2\sqrt{dt}} \right) - \operatorname{erf} \left(\frac{A_0 t + y}{2\sqrt{dt}} \right) \right) \\ - \frac{U}{4} \frac{\sqrt{\mu\rho}}{A_0} e^{-\frac{A_0 y}{d}} \left(d + A_0 e^{\frac{A_0 y}{d}} y \right) \operatorname{erfc} \left(\frac{y - A_0 t}{2\sqrt{dt}} \right) \\ - \frac{U}{4} \frac{\sqrt{\mu\rho}}{A_0} \left(d e^{\frac{A_0 y}{d}} - A_0 y \right) \operatorname{erfc} \left(\frac{A_0 t + y}{2\sqrt{dt}} \right),$$

and must be sustained by an external electric field with $E_z^0 = \frac{B_0 U}{2}$.

The parameters in this problem are taken as in [33] to be $U = 1.0 \text{ m/s}$, $B_0 = 1.4494e^{-4} \text{ kg/s}^2 A$, $\rho = 0.4e^{-4} \text{ kg/m}^3$, $\mu_0 = 1.256636e^{-6} \text{ kg m/s}^2 A^2$, $\eta = 1.256636e^{-6} \text{ kg m}^3/\text{s}^3 A^2$, and $\mu = 0.4e^{-4} \text{ kg/m s}$ with $d = \eta/\mu_0 = \mu/\rho = 1$. The initial conditions and boundary conditions are defined by the analytical solution above. In Figures 8 and 9, line plots of the analytic and numerical solution for v_x , and A_z , as well as the post-processed quantity B_x , are presented at various discrete times. In these plots, the Hartmann layer in v_x near the plate surface at $y = 0$, and the plateau region influenced by the magnetic field is evident. These figures show excellent agreement with the analytical solution. Figure 10 shows a detailed spatial and temporal order-of-accuracy study. It should be noted that the temporal accuracy computations displayed on the right begin at a CFL = 40 based on the Alfvén velocity. These results

indicate that the second order in time method is already achieving the expected order of accuracy. The expected second-order accuracy for v_x and A_z , which are interpolated with linear nodal elements, is confirmed by these results. Again, as in the last verification problem, the results for B_x exhibit second-order accuracy due to super-convergence. The figure also exhibits the expected order-of-accuracy for the backward Euler (BE, first order) and Crank-Nicolson (mid-point, second order) temporal integration schemes.

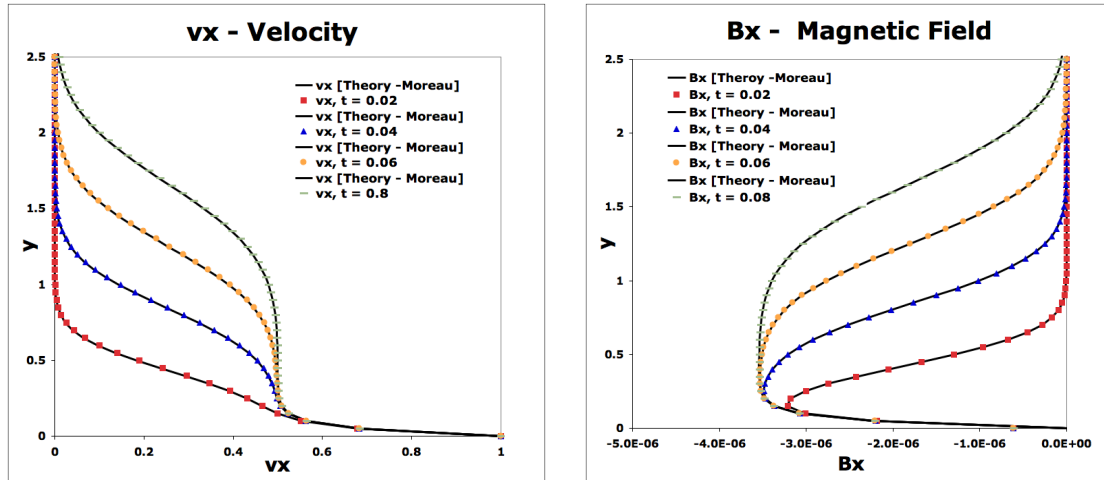


Fig. 8. Modified Rayleigh flow and Alfvén wave verification problem. Figures show plots of x -velocity (left) and x -magnetic field (right). The mesh has 50×250 elements and the time step is 5×10^{-4} time units.

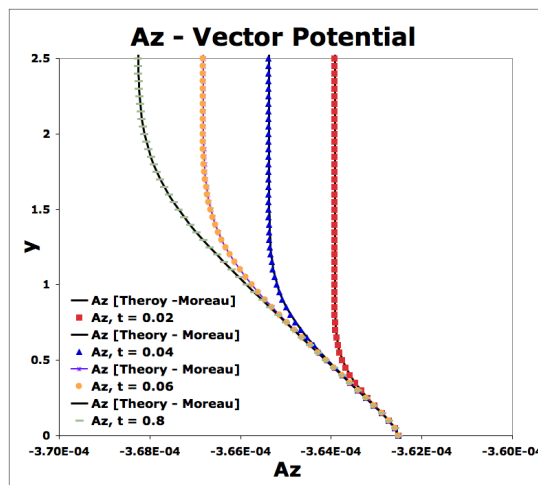


Fig. 9. Modified Rayleigh flow and Alfvén wave verification problem. Figure shows plots of A_z . The mesh has 50×250 elements and the time step is 5×10^{-4} time units.

5.2 Prototype Resistive MHD Problems and Linear Solver Performance

This section briefly describes three prototype resistive MHD problems that are intended to challenge the stabilized FE 2D vector potential formulation and the fully-implicit, fully-coupled solution methods presented in this paper. They are also intended to assess the parallel performance (in a weak scaling sense) of these solution methods to very large problem sizes on thousands of processors. One of the tests demonstrates the solvers in the large-Lundquist-number regime.

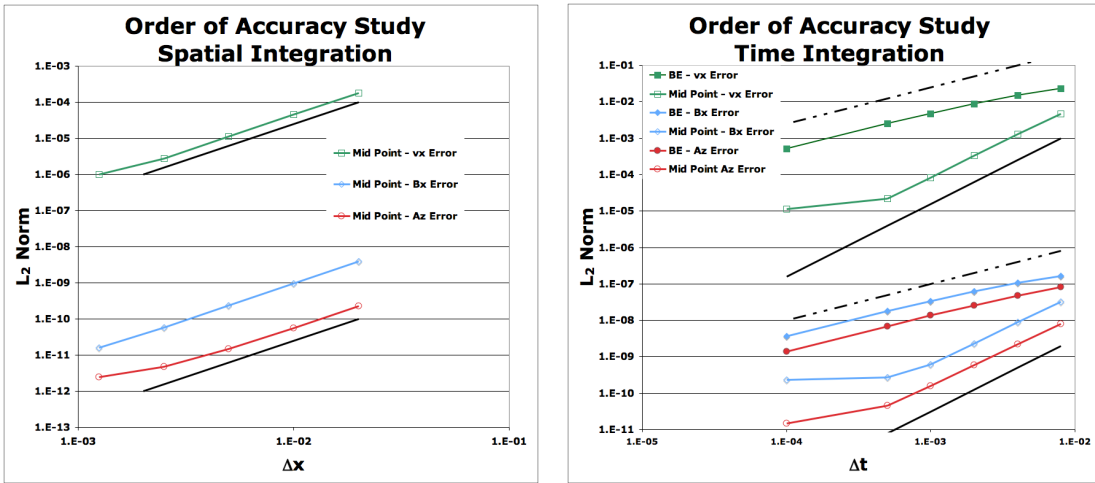


Fig. 10. Order of accuracy study for modified Rayleigh flow and Alfvén wave verification problem. A spatial convergence study for A_z , v_x , B_x , is shown on the left with a comparison to second order reference lines for $\Delta t = 1.0 \times 10^{-4}$. The expected order-of-accuracy is obtained followed by a region where the error begins to plateau as the temporal integration error begins to dominate. Figure on the right shows convergence of the computed numerical solution for the backward Euler and midpoint rule time integrator for $\Delta x = 5.0 \times 10^{-3}$. In the midpoint rule results, the error also begins to plateau as the error in the spatial discretization begins to dominate at small time step sizes.

The first problem is a steady-state MHD duct flow configuration representing an idealized Faraday conduction pump. The MHD pump induces flow by the action of an applied magnetic and electric field. The second problem is a coupled thermal-convection buoyancy-induced flow that is modulated by an externally applied magnetic field and produces internally generated fields. The corresponding linear stability problem is the classical hydromagnetic Rayleigh-Bernard problem. The final example is a driven magnetic reconnection problem where a Fadeev magnetic field equilibrium [96], which features islands embedded in a Harris current sheet, undergoes a transient reconfiguration of the magnetic field. The computational timing results presented in this section were obtained on the Red Storm Cray XT3/4 computer at Sandia National Laboratories. The Newton convergence criterion is specified in the discussion of each specific problem. In all these examples, the linear systems generated by Newton’s method have been left-scaled by a diagonal matrix that contains the inverse of the row sum of absolute values of all entries in the row.

5.2.1 An Idealized MHD Faraday Pump

As an illustration of the parallel performance of the one-level additive Schwarz domain decomposition and multilevel preconditioners, a weak scalability study is presented for a 2D idealized Faraday MHD conduction pump. This problem models an MHD pump that induces flow in a conducting fluid by applying an external magnetic field in the y -direction and an electric field in the z -direction. The domain is $\Omega = [0, 10] \times [-1, 1]$. There are no-slip fluid velocity conditions applied on the upper and lower surfaces with natural boundary conditions for the system applied at both the inlet and outlet of the domain. On the lower and upper surfaces, a constant external magnetic field $\mathbf{B} = (0, B_0, 0)$ is applied by application of a linear variation in A_z in the x -direction in the range of $x \in [2.5, 7.5]$, while outside of this range the magnetic field is zero. A constant electric field, E_z^0 is applied in the z -direction. The interaction of these fields produces a Lorentz force that pulls fluid in from the $x = 0$ boundary with a parabolic profile, contorts the velocity field into a common “M” profile for these types of flows [97–99], and then the flow exits with a parabolic profile (see Figure 11 and 12). The simple geometry of this problem facilitates scalability studies as different mesh sizes can be easily generated.

In our numerical study, we briefly consider the effect of ILU fill-in for the one-level preconditioner and the benefit of employing multilevel methods with coarse operator solves as described above. The Krylov method is a non-restarted GMRES technique. This choice eliminates the degradation of convergence that can occur from restarting of the GMRES iteration, and allows the scalability of the preconditioners to be addressed. For the one-level DD preconditioner, an incomplete factorization ILU(k) sub-domain solver was used with $k = 1, 3, 7$. For the 3-level preconditioner, the fine and medium meshes use an ILU(1) smoother and the coarsest problem was solved by the KLU sparse direct solver. For this problem, the inexact Newton forcing term is taken as $\eta_{k+1} = 10^{-4}$.




Fig. 11. Color plots of spatial distribution of A_z (upper left), velocity vector magnitude (lower left), B_x (upper right) and B_y (lower right) for idealized Faraday conduction MHD pump. The red color is high and the blue color is low for each variable.

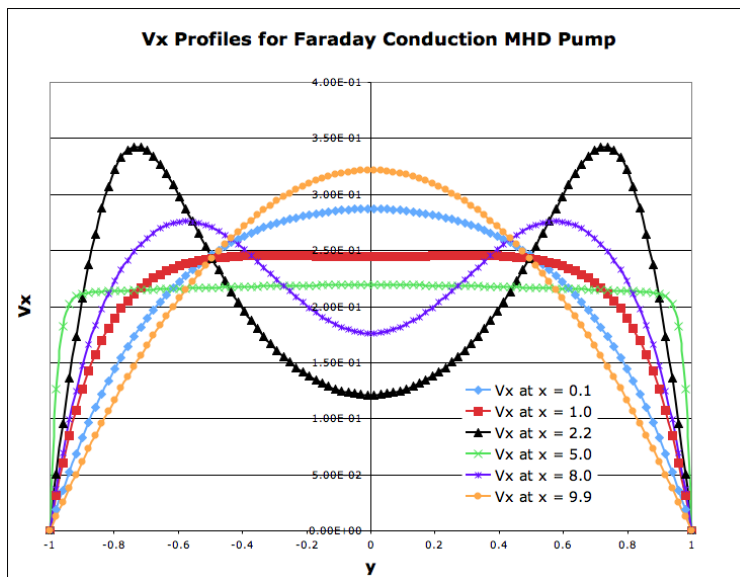


Fig. 12. V_x velocity profiles for the developing flow in an idealized Faraday MHD pump. The locations of the profiles are shown graphically in Figure 11.

As an initial illustration of parallel efficiency, we consider the weak scaling of the one-level DD ILU preconditioner as presented in Table 3. In this study, the 16 processor case solves the problem on a 800×80 mesh. The weak scaling study keeps the work per processor fixed as the problem size is increased. This study is for a low $Re = Re_m = 0.7$ flow with $Ha = 1$. Figure 13 graphically presents the parallel and algorithmic scaling of the one- and three-level preconditioners for the MHD Faraday pump presented in Table 3. Figure 13 (left) summarizes the results for the average iteration count per Newton step as a function of problem size. As the number of unknowns, N (as well as the number of processors, P , in this scaled study), is increased, the number of iterations to convergence for the one-level schemes increases significantly: roughly $N^{1/2}$ in two dimensions. Note that an optimal convergence property,

proc	fine unks	1-level ILU(1)		1-level ILU(3)		1-level ILU(7)		3-level V(1,1)	
		avg its/	time	avg its/	time	avg its/	time	avg its/	time
		Newton step	(sec)	Newton step	(sec)	Newton step	(sec)	Newton step	(sec)
16	256K	136[9]	9.8	86[8]	8.6	62[9]	9.7	47[8]	7.6
64	1M	313[9]	22.8	198[9]	15.1	154[9]	14.8	80.1[9]	9.3
256	4M	714[11]	82.4	701.4[11]	44.9	392[8]	36.3	136[9]	16.9
1024	16M	1583[10]	318	1158[12]	191	888[10]	123	129[10]	16.4
4096	64M	2667[20]	750	1951[16]	488	1766[14]	429	147[13]	31.9

Table 3

Comparison of 1-level and 3-level PGSA method for MHD Faraday pump example problem. The 3-level non-smoothed aggregation method uses an aggregation size of 80 and an ILU(1) smoother on the fine and medium meshes with the KLU direct solver on the coarsest problem. The table entry above for 136[9] indicates the number of GMRES iterations followed by the number of Newton steps in brackets.

i.e., an iteration count independent of problem size, is roughly obtained for the 3-level preconditioner. On the coarsest level, a serial sparse matrix direct solver, KLU, was used to factor the coarse matrix. Since the fine grid smoother is highly parallel and the fine grid work per processor is roughly constant, the cost of producing the coarse grid problem and executing the direct solve (KLU) on the increasingly larger coarse grid causes an increase in the CPU time for the larger problems. While this CPU time increase is non-optimal, the 3-level method remains significantly faster (a factor of about 10–20x) than the corresponding one-level methods. The CPU time growth due to the coarse grid solve can be mitigated by using either approximate coarse grid methods (e.g. [76]) or more levels.

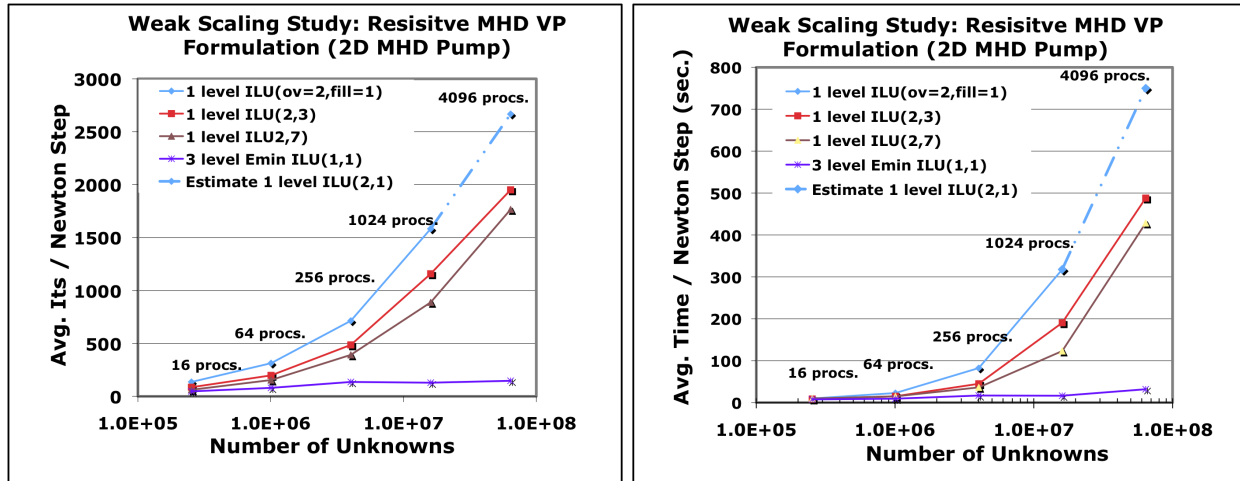


Fig. 13. Weak scaling results for MHD Faraday pump problem. Comparison of 1-level and 3-level PGSA method. The scaling of the average number of iterations per Newton step (left) and the average CPU time per Newton step (right).

As a proof-of-capability of the multilevel preconditioner, we provide an initial demonstration of the solution of a very large scale problem on tens-of-thousands of processors. Specifically, a problem with 1.05 billion unknowns for the MHD pump is solved on the Cray XT3/4 using 6,000 nodes with 4 cores per node for a total of 24,000 cores. For this problem size, the one-level method cannot be applied. The

details of this calculation are presented in Table 4. These results are very encouraging. They indicate convergence of the Newton iteration in $O(10)$ iterations and of the linear solver in $O(100)$, values that are not too far from the more complete weak scaling study of Table 3 that employs three levels and a larger aggregation size of 80. Finally the ability of the multilevel preconditioned Newton-Krylov solver to scale to 1+ billion unknowns on 24,000 cores is an indication that the underlying solution strategy is reasonably robust and scalable for this low Re_m prototype problem.

cores	Fine Mesh Level 0 Unkns.	Intermed. Level 1 Unkns.	Intermed. Level 2 Unkns.	Coarse Level 3 Unkns.	Newton Its.	Avg. No. Linear Its. / Newton	Total Time (min.)
24,000	1.05 billion	23.3M	.5M	11.2K	18	86	33

Table 4

Details of 1.05 billion unknown MHD Faraday pump calculation. The simulation used a mesh of size $51, 200 \times 5, 120$ with roughly 260M quad elements. The 4-level multilevel PGSA preconditioner used an aggregation size of 45 at each level to produce the coarsened operators.

5.2.2 A Classical Hydromagnetic Rayleigh-Bernard Stability Computation

This resistive MHD prototype problem consists of a buoyancy induced thermal convection flow that is modulated by the existence of an externally applied magnetic field. It combines the classical Rayleigh-Bernard buoyancy induced thermal convection flow problem with an externally applied \mathbf{B} field. The magnetic field induces Maxwell stresses that add additional stabilizing effects to the traditional damping provided by the viscous forces in Navier-Stokes. These coupled mechanisms are, for example, critical components of large-scale geo-dynamo simulations that model the time dependent behavior of the Earth’s magnetic field (e.g.[100,1]). This problem solves for the unknowns (\mathbf{v}, P, T, A_z) from the system outlined in Table 1 where the density variation is modeled by a Boussinesq approximation [100]. This approximation uses a constant density in all the terms of the balance equations with the exception of the momentum equation source term. In this term, a linear approximation is used with $\mathbf{f}(T) = \rho_0 \mathbf{g} + \beta(T - T_0) \mathbf{g}$, where ρ_0 and T_0 are a reference density and a reference temperature respectively, and β is the thermal expansion coefficient.

The rectangular domain for this problem is $\Omega = [0, 10] \times [0, 1]$. There are no-slip fluid velocity conditions applied on the upper and lower surfaces with natural boundary conditions for the system applied at both the left and right boundary of the domain. A temperature difference is maintained in the vertical direction by holding the lower surface to a high temperature T_H and the upper surface at the lower temperature T_C that is separated by a distance d . This temperature difference $\Delta T = (T_H - T_C)$ produces an unstable density stratification that interacts with gravity in the negative y -direction. In the classical Rayleigh-Bernard problem, flow is induced when the non-dimensional Rayleigh number, $Ra = \rho^2 \hat{C}_p g \beta \Delta T d^3 / (\mu \chi)$, becomes sufficiently large to have buoyancy effects overwhelm the stabilizing viscous forces [100]. Here d is a characteristic length scale aligned with the gravity vector. In the case of the hydromagnetic Rayleigh-Bernard (HMRB) problem, a stabilizing constant external magnetic field $\mathbf{B} = (0, B_0, 0)$ is applied by application of a linear variation in A_z in the x -direction on the lower and upper surfaces. For a thermal convecting flow at fixed Rayleigh number, convective flow will be damped when the non-dimensional magnetic field strength is increased beyond a critical Chandrasekhar number, $Q = Ha^2$. Beyond this limit, the Maxwell stresses existing in the curved magnetic field lines shut down the convective cellular flow. A typical stable nonlinear solution where cellular flow and fields exist is presented in Figure 14.

As an example of the robustness and efficiency of the fully-coupled multi-level preconditioned NK solution technology, we present the results in Figure 15 for convergence of the nonlinear Newton iteration (left image) and the preconditioned Krylov linear solver (right image). In this study we have taken $\rho = \nu = \mu_0 = \hat{C}_p = d = \chi = \beta = 1$ with the Ra set by adjusting ΔT and Q set by B_0 . The convergence of the outer fully-coupled nonlinear Newton solver in Figure 15 (left) indicates that, as the Chandrasekhar number increases, the nonlinearity of the problem increases and the convergence of the nonlinear iteration becomes more difficult. Without the use of backtracking techniques [68], convergence is not obtained past $Q = 4$. Beyond $Q = 11$, a direct-to-steady-state solution was not obtained and the use of continuation and/or transient solution methods would need to be considered if a direct-to-steady-state solution is required for larger Chandrasekhar numbers. In the context of the linear solver convergence, Figure 15 (right) indicates that the multilevel preconditioned Krylov method converges substantially faster to the solution in the first two Newton steps. In addition, the ability of the 3-level methods to provide a sufficient linear residual decrease in the sub-problems for Newton's method allows for a more robust iterative nonlinear solver. This is indicated in Table 5 for the hydromagnetic Rayleigh-Bernard problem. Note that the one level DD preconditioner failed to solve this problem in less than 40 Newton steps. For this case, the linear solver failed to reach its convergence criterion of $\eta = 10^{-3}$ in 2000 iterations for 36 out of the 40 Newton steps. We have seen many such examples where a direct-to-steady-state nonlinear calculation fails for very ill-conditioned large linear systems when an ineffective preconditioner causes the iterative linear solver to converge too slowly to be practical, or does not converge at all. For this reason, effective scalable preconditioners, such as the multilevel preconditioner described in this study, are not only required for efficiency but also often for robustness in solving very large-scale problems.




Fig. 14. Color plots for the hydromagnetic Rayleigh-Bernard stability type calculation. The images show the stable nonlinear flow and fields at $Ra = 2500$ and $Q = 9$ for temperature (upper left), V_y (center left), V_x (lower left), and J_z with the magnetic field vectors (upper right), B_y (center right), and B_x (lower right). At $Q = 9$ the critical Rayleigh number is ≈ 1922 . The red values are high and blue values are low in each image.

As an additional verification of the accuracy of the resistive MHD formulation presented above, we present a comparison of the computed critical Rayleigh number from simulation, and the theoretical value in Chandrasekhar [100]. In this study, a computational eigenvalue analysis of the linear stability problem is computed as in [79] with the results summarized in Table 6, where the accuracy of the formulation is clearly apparent. It should be noted that by employing the direct-to-steady-state solution capability based on Newton-Krylov-type techniques as presented in this study, additional advanced solution methods such as parameter continuation, bifurcation tracking, and automated linear stability analysis algorithms can be effectively developed to analyze the complex nonlinear solution spaces for

the HMRB problem.

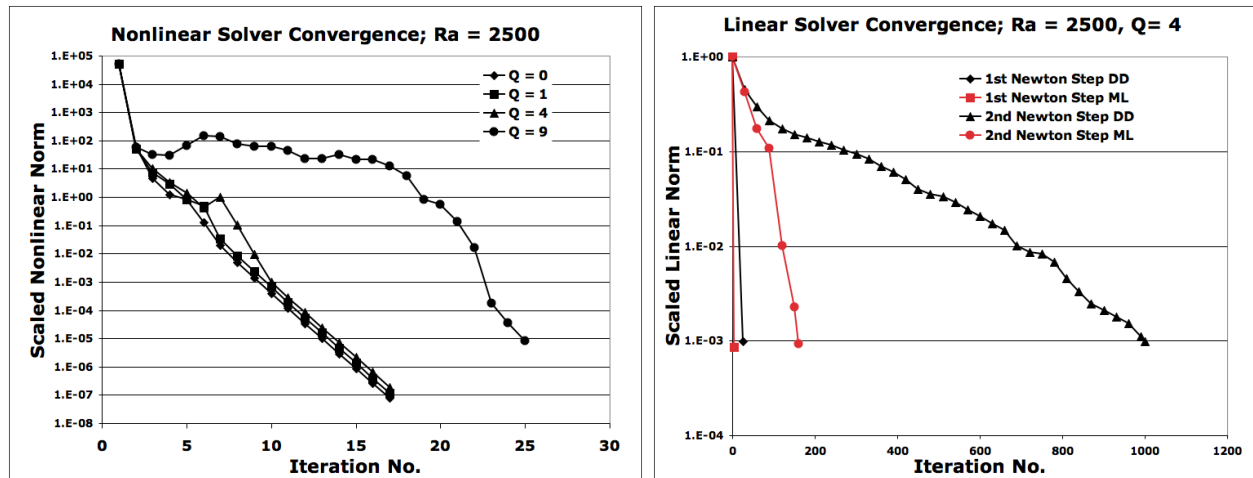


Fig. 15. Hydromagnetic Rayleigh-Bernard problem. A convergence plot of the nonlinear scaled norm in Newton’s method with the 3-level ML PGSA preconditioner (left) and the linear scaled norm in non-restarted GMRES for DD and ML preconditioner (right).

proc	fine grid size	fine grid unknowns	1-level ILU		3-level V(1,1) ILU-ILU-KLU			
			avg its/ Newt step	total time (sec)	medium unkns size	coarse unkns size	avg its/ Newt step	total time (sec)
2048	500×5000	12.5M	1910[40]	> 7200*	412450	13745	115[17]	226

Table 5

Comparison of convergence for the 1-level and 3-level PGSA multilevel preconditioner for Hydromagnetic Rayleigh-Bernard problem. *The 1 level DD preconditioner failed to solve this problem in less than 40 Newton steps.

Q	Ra^*	Ra_{cr} [Chandrasekhar[100]]	% error
0	1707.77	1707.8	0.002
10^1	1945.78	1945.9	0.006
10^2	3756.68	3757.4	0.02

Table 6

Comparison of the computed critical Rayleigh number, Ra^* , by a computational linear stability eigenvalue analysis with theory, Ra_{cr} , from Chandrasekhar [100]. The mesh is a rectangular 200×2000 mesh.

5.2.3 A Transient Simulation of Driven Magnetic Reconnection: Island Coalescence

As a final test, the numerical and computational performance of the implicit stabilized FE resistive MHD formulation presented here is applied to a driven magnetic reconnection example, the island coalescence problem. This is an example of a current, scientifically relevant, and computationally challenging application of resistive MHD. Unlike earlier examples in this study, this problem explores the moderately-high Lundquist number regime, in which the MHD equations are dominated more by hyper-

bolic couplings. Accordingly, we do not expect our current preconditioning approach, which is tailored to parabolic and elliptic PDEs, to perform optimally. Nevertheless, the solver has demonstrated to be reasonably robust, even in this regime. We believe that improving on these results will require the implementation of physics-based preconditioning ideas, as proposed in [28–30].

Magnetic reconnection is a fundamental process whereby a sheared magnetic field is topologically altered via some dissipation mechanism, resulting in a rapid conversion of magnetic field energy into plasma energy and significant plasma transport. Magnetic reconnection dominates the energetics and dynamics of many space and laboratory plasmas, and is at the root of explosive phenomena such as solar flares, coronal mass ejections, plasmoid ejection from the Earth’s magnetotail, and major disruptions in magnetic fusion energy (MFE) experiments [101]. However, plasmas in all the above-mentioned cases are known to have very small electrical resistivities, which cannot explain the observed reconnection timescales. This theoretical conundrum has drawn significant attention of the theoretical and numerical simulation plasma physics community for the last 50 years [101].

The computational challenges of studying magnetic reconnection phenomena are many. The presence of such small resistivities implies the formation of very thin current sheets, which need to be adequately resolved to capture the phenomena accurately. Temporally, time scales for magnetic energy build-up and release scale like a power of the Lundquist number (reciprocal of the normalized resistivity), and are therefore very long compared to normal-mode time scales. Therefore, the problem is intrinsically multiscale in time and space. Addressing the spatial resolution aspect requires mesh adaptation, while the temporal aspect is a textbook case for the use of implicit methods. Recently, the development of an efficient, implicit resistive-MHD capability in a mapped, structured-mesh finite-volume context has led to new computational scientific results that have verified the Sweet-Parker [101,102] asymptotic scaling of the reconnection rate in the very small resistivity regime. The key to the computational verification of this regime is the ability to effectively time-integrate the multiple-time-scale resistive MHD system [102].

The island coalescence equilibrium is described by the Fadeev solution [96], which features islands embedded in a Harris current sheet. The structure of this equilibrium can be seen in the upper left plot of Figure 16 with iso-lines of A_z . The combined magnetic field produced by the two magnetic islands produces Lorentz forces that pull the islands together, and at finite resistivity the islands coalesce (join) to form one island. Figure 16 shows three iso-line plots of A_z and filled color contours of the the plasma current J_z during the reconnection event. Clearly evident is the formation of the x-point between the islands (see upper right image), the development of a thin current sheet at that same x-point location, and the movement of the center of the islands (o-points) towards the x-point [101,102]. The dynamics of island coalescence changes as a function of resistivity. For larger resistivities, the x- and o-points monotonically approach each other. For low resistivities, fluid-plasma pressure builds up as the islands approach and a sloshing or bouncing of the o-point position is encountered that leads to lower reconnection rates (for more details on the physics see e.g. [101,103]). Next, a very brief description of the island coalescence setup is presented, with details provided in [102].

This problem solves for the unknowns (\mathbf{V}, P, A_z) . The initial conditions for the island coalescence problem consists of zero fluid velocities ($\mathbf{v}^0 = 0$), and a Fadeev magnetic equilibrium [96]. The description of this problem is in terms of the magnetic vector potential. The initial conditions, A_z^0 , and the resulting balancing plasma fluid pressure, P^0 are given by

$$A_z^0(x, y, 0) = \delta \ln \left[\cosh \left(\frac{y}{\delta} \right) + \epsilon \cos \left(\frac{x}{\delta} \right) \right], \quad (11)$$

$$P^0(x, y, 0) = P_0 + \frac{[1 - \epsilon^2]}{2 \left[\cosh \left(\frac{y}{\delta} \right) + \epsilon \cos \left(\frac{x}{\delta} \right) \right]^2}, \quad (12)$$




Fig. 16. Plots for the island coalescence driven magnetic reconnection at times $t = 0.0, 9.0, 10.0, 12.6$ computed on the unstructured mesh above. The images show isolines of the magnetic potential A_z and filled contours of the current J_z . The red values are high and blue values are low in each image.

where $\delta = 1/(2\pi)$ and $P_0 = 1.0$. To assure that the initial condition is a resistive equilibrium, an external applied electric field in the z -direction, E_z^0 , of the form

$$E_z^0(x, y) = \frac{\eta [1 - \epsilon^2]}{\delta \left[\cosh\left(\frac{y}{\delta}\right) + \epsilon \cos\left(\frac{x}{\delta}\right) \right]^2} \quad (13)$$

is included. Using symmetry as in [102], a $1/4$ computational domain $[0, 1] \times [0, 1]$ is considered. Symmetry boundary conditions are used on the $x = 0$, $x = 1$ and $y = 0$ boundaries. These conditions enforce zero normal velocity and zero tangential stress for the momentum equation and no flux or Neumann conditions for A_z . On the upper $y = 1$ boundary, there is a zero tangential stress, the vector potential A_z is set by (11), and $P_0 = 1$. In this study, we have also taken $\rho = 1$, $\nu = \eta = 1$, $\mu_0 = 1$ and using the spacing of the o-points we have $L = 1$. As in [102] these choices imply that the resistivity $\eta = 1/S$, where S is the Lundquist number and is defined as $S = \mu_0 L V_A / \eta$, with $V_A = B_0 / \sqrt{\mu_0 \rho}$ the Alfvén velocity.

Our interest in this brief study is to reproduce the results previously carried out with the fully-implicit mapped FV capability in [102], where the peak magnetic reconnection rate was computed as a function of the fluid resistivity for a range of resistivities, spanning various reconnection regimes. The peak magnetic reconnection rate is found as the maximum value of the time derivative of the vector potential at the x-point, i.e. $\Psi_t \equiv \partial A_z / \partial t$ for all time. Figure 17 (left) presents the result of a set of stabilized FE simulations to determine the reconnection rate at the x-point for 3 different resistivity values that illustrate the different nature of the reconnection process as a function of η . Clearly evident in this figure is the peak value for each time history and the oscillatory (or sloshing) behavior of the reconnection rate for lower resistivities as in [102]. In Figure 17 (right) the scaling of the peak magnetic reconnection rate as a function of the resistivity is presented. The Lundquist number varies from 50 to 10^5 . The FE results are seen to compare very well with the results of Knoll and Chacon [102]. In this plot, results are presented for non-uniform structured meshes of size 130,000 FE nodes with a distribution of 512×256 in (x, y) . In this mesh, the FE elements are clustered by the x-point and geometrically increase in size away from this point, which roughly corresponds to the meshes used in [102]. In addition, results are also presented from an unstructured mesh with 40,000 FE nodes, which highly resolves the x-point region with a uniform mesh and then transitions to a coarse unstructured mesh away for the x-point. A plot of an unstructured mesh and the iso-lines of A_z is presented in Figure 18. It is clear from this study that the stabilized FE results closely correspond to the mapped FV results for both the structured and unstructured meshes. In addition, the existence of the theoretical Sweet-Parker slow magnetic reconnection regime where $\Psi_t \approx \sqrt{\eta}$ is also confirmed. A power law fit to the FE results indicate a 0.48 power scaling over this interval.

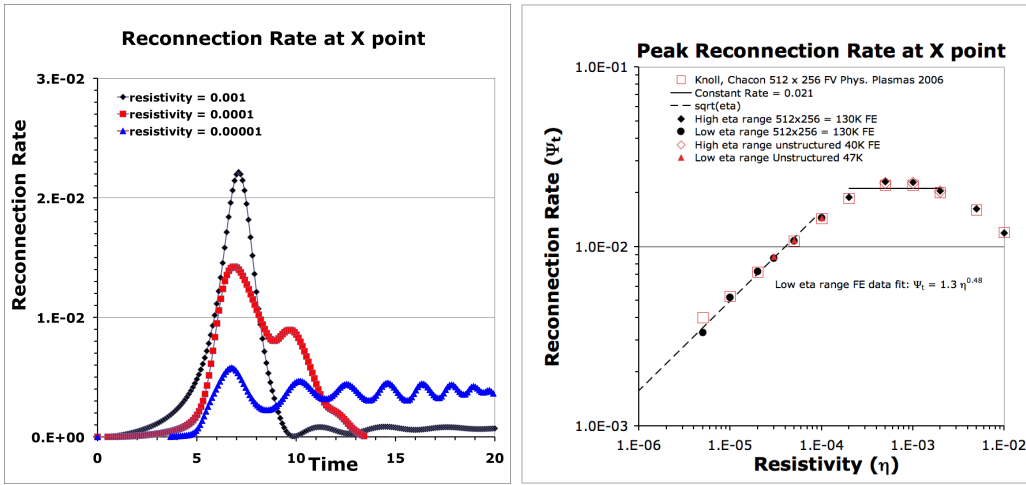


Fig. 17. Left: Time history of the reconnection rate at the x-point in the island coalescence problem for various resistivities. Oscillations in the reconnection rate are evident for lower resistivities as a consequence of the sloshing of magnetic islands. Right: Peak resistive magnetic reconnection rate for the island coalescence problem. Graph includes data for stabilized FE formulation and a comparison to FV resistive MHD data from Knoll and Chacon [102].

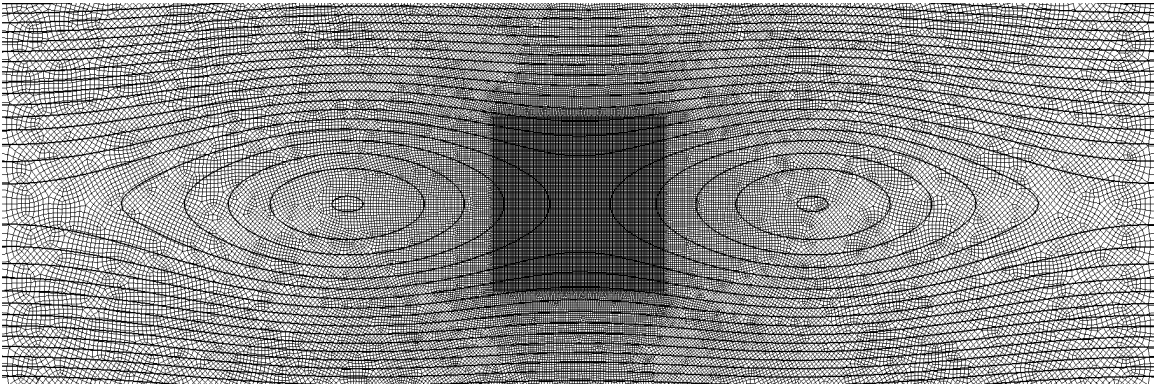


Fig. 18. Detail of unstructured mesh for island coalescence problem. This image also has iso-lines of A_z represented at time $t = 9.0$.

A preliminary study of the scaling for the multilevel preconditioner for the transient solution of the island coalescence problem at a resistivity of $\eta = 10^{-3}$ is presented in Table 7. Here, factoring out the increase in nonlinear iterations per time step, the general block aggregation based algebraic multilevel preconditioner appears to be reasonably effective at keeping the growth in iterations per time step under control as the mesh size is increased. While the scaling of the multilevel preconditioned Newton-Krylov method is not optimal with problem size, the increase in the number of linear iterations per Newton step is gradual with the problem size, and represents a reasonable step towards a scalable algebraic multilevel method.

6 Conclusions

This paper has presented a performance study of an unstructured fully-implicit stabilized FE formulation for 2D incompressible reduced resistive MHD, both in regards to the accuracy of the discretization and the efficiency of the associated solver. The solution methods used in this formulation are based on a fully-coupled Newton-Krylov approach employing algebraic multilevel preconditioners. The study includes verification and order-of-accuracy results, which confirm the expected convergence

Procs	Mesh	N_{unks}	Newton/ Δt	Gmres/ Newton	Time/ Newton	Gmres/ Δt	Time/ Δt
1	64×64	16K	3.9	4.4	2.1	17.2	8.1
4	128×128	64K	4.6	5.8	2.6	26.7	11.9
16	256×256	0.25M	4.9	6.3	2.9	30.9	14.2
64	512×512	1M	6.2	8.8	4.0	54.6	24.6

Table 7

Weak scaling of the block AMG preconditioned Newton-Krylov solver for the stabilized resistive MHD formulation on the island coalescence problem. The time step is $\Delta t = 0.1$ time units. The smoother in the Petrov-Galerkin smoothed aggregation method is an ILU(1) and it is a V(1,1) cycle. In this study the linear solver convergence criteria is 10^{-3} for each step of the Newton solver.

rates for the numerical representation. As an illustration of the robustness, scalability, and efficiency of the solution techniques advocated in this paper, performance results have been presented for both low Lundquist number (MHD duct flow, a hydromagnetic Rayleigh-Bernard linear stability calculation), and moderately-high Lundquist number (magnetic island coalescence) problems. The efficiency and scalability results of this study are very encouraging, and underscore the robustness of the fully-coupled Newton-Krylov nonlinear solver. In particular, the results clearly demonstrate the improved convergence properties of block aggressive coarsening, fully-coupled parallel multilevel preconditioners over more standard parallel additive-Schwarz domain-decomposition methods. Future work will focus on the implementation of physics-based preconditioning ideas to deal more effectively with the large Lundquist number regime, and the extension of this work to a full 3D compressible resistive MHD model.

7 Acknowledgments

The authors wish to thank Gary Hennigan and Rob Hoekstra for their collaborative effort in developing the simulation code framework that is used as the software development platform. We also thank Dana Knoll for helpful discussions on the MHD equations and magnetic reconnection phenomena, Pavel Bochev for discussions on finite element representations for MHD and John Evans for helpful discussions on stabilization techniques.

8 Appendix

8.1 Stabilization Parameters

The specific definition of the intrinsic-time-scale stability parameters are provided in Table 8 for momentum, thermal energy, and the vector potential. The specific form of the intrinsic time scales τ 's are an adaptation of the quadratic form by Shakib [63] for the Navier-Stokes equations and the Lorentz form stabilization operator from Codina and Hernandez-Silva [39] for a resistive MHD system. In the form by Shakib, the estimate of the inverse constant for linear nodal elements is taken as $C_1 = 3$ [63] for the diffusive operator. Additional forms for the intrinsic-time-scale stability parameters are discussed in [53] and the references contained therein for Navier-Stokes. The contribution to the intrinsic-time-scale operator for the Lorentz force term uses the same form as in [39] with $C_2 = 10$. We have run limited numerical studies with the choice of $C_1 = 1 - 10$ and $C_2 = 5 - 20$ as well as alternate forms for

Momentum	$\hat{\tau}_m = \left[\left(\frac{2\rho}{\Delta t} \right)^2 + \rho^2 \mathbf{v} \mathbf{G}_c \mathbf{v} + C_1^2 \mu^2 \ \mathbf{G}_c\ + C_2^2 \ B\ \sqrt{\ \mathbf{G}_c\ } \right]^{-\frac{1}{2}}$
Thermal Energy	$\hat{\tau}_T = \left[\left(\frac{2\rho C_p}{\Delta t} \right)^2 + (\rho C_p)^2 \mathbf{v} \mathbf{G}_c \mathbf{v} + C_1^2 \lambda^2 \ \mathbf{G}_c\ \right]^{-\frac{1}{2}}$
Z-component Vector Potential	$\hat{\tau}_{A_z} = \left[\left(\frac{2}{\Delta t} \right)^2 + \mathbf{v} \mathbf{G}_c \mathbf{v} + C_1^2 \eta^2 \ \mathbf{G}_c\ \right]^{-\frac{1}{2}}$

Table 8

Definition of stabilization parameters used in stabilized equations, which use the contravariant metric tensor \mathbf{G}_c (Equation (14)) to define an element-level streamwise length scale. In this study $C_1 = 3$ and $C_2 = 10$.

the Lorentz force intrinsic-time-scale operator contribution. The results indicate negligible variation in problems for which we carry out order-of-accuracy studies. In all the examples in this study, the most sensitive case was the island coalescence problem. This is not surprising since this problem has the least amount of natural diffusion and uses a rather coarse nonuniform mesh of size 512×256 . In our limited studies, for a resistivity $\eta = 10^{-4}$ the results for the maximum of reconnection rate, Ψ_t , varied by $< 0.75\%$.

It should be pointed out that the multidimensional effect of convection is incorporated into the stability parameters by the use of the contravariant metric tensor, \mathbf{G}_c (Equation (14)), of the transformation from local element coordinates $\{\zeta_\alpha\}$ to physical coordinates $\{x_i\}$.

$$[\mathbf{G}_c]^{ij} = \frac{\partial \zeta_\alpha}{\partial x_i} \frac{\partial \zeta_\alpha}{\partial x_j}. \quad (14)$$

Shakib [63] considers the one dimensional limiting case of this multidimensional definition for the advection-diffusion equation and presents a comparison with the original SUPG technique.

8.2 Brief Overview of Discrete Systems of Equations

To give context to the discussion of solution methods and linear algebra techniques that we employ in the linear solution methods above, we present a brief discussion of the structure of the equations that result from the FE discretization of the weak form of the resistive MHD equations. In this discussion, the Newtonian stress tensor is expanded to include the hydrodynamic pressure, P , and the viscous stress tensor term, $\mathbf{\Pi}$. The resulting stabilized form of the total mass residual equation in expanded form is given by

$$F_P = \int_{\Omega} \Phi \left[\frac{\partial \rho}{\partial t} + \nabla \cdot (\rho \mathbf{v}) \right] d\Omega + \sum_e \int_{\Omega_e} \rho \hat{\tau}_m \nabla \Phi \cdot \left[\rho \frac{\partial \mathbf{v}}{\partial t} + \rho (\mathbf{v} \cdot \nabla \mathbf{v}) + \nabla \cdot \left(-\frac{1}{\mu_0} \mathbf{B} \otimes \mathbf{B} - \mathbf{\Pi} + \left(P + \frac{1}{2\mu_0} \|\mathbf{B}\|^2 \right) \mathbf{I} \right) \right] d\Omega. \quad (15)$$

This expansion includes the weak form of a Laplacian operator acting on pressure,

$$L = \sum_e \int_{\Omega_e} \rho \hat{\tau}_m \nabla \Phi \cdot \nabla P d\Omega, \quad (16)$$

which is produced by the stabilized formulation of the total mass conservation equation.

A finite element (FE) discretization of the stabilized equations gives rise to a system of coupled, nonlinear, non-symmetric algebraic equations, the numerical solution of which can be very challenging. These equations are linearized using an inexact form of Newton’s method as described in Section 4.1. A formal block matrix representation of these discrete linearized equations is given by

$$\begin{bmatrix} \mathbf{K} & \mathbf{G} \\ \mathbf{D} & \mathbf{L} \end{bmatrix} \begin{bmatrix} \hat{\mathbf{u}}' \\ \hat{\mathbf{P}}' \end{bmatrix} = - \begin{bmatrix} \mathbf{F}_u \\ \mathbf{F}_P \end{bmatrix}. \quad (17)$$

where the block diagonal contribution of the stabilization procedure has been highlighted by a specific ordering. In this representation, the vector, \mathbf{u}' , contains the Newton updates to the nodal solution variables, $(\hat{\mathbf{v}}, \hat{\mathbf{T}}, \hat{\mathbf{A}}_z)$, with the exception of the nodal pressures, $\hat{\mathbf{P}}'$. The block matrix, \mathbf{K} , corresponds to the combined discrete transient, convection, diffusion and stress terms acting on the unknowns $\hat{\mathbf{u}}'$; the matrix, \mathbf{G} , corresponds to the discrete gradient operator; \mathbf{D} , the divergence operator; and the matrix, \mathbf{L} , corresponds to the discrete “pressure Laplacian” operator discussed above. The vectors \mathbf{F}_u and \mathbf{F}_P contain the right hand side residuals for Newton’s method. The existence of the well conditioned nonzero matrix, \mathbf{L} , in the stabilized FE discretization is in contrast to Galerkin methods using mixed interpolation that produce a zero block on the total mass continuity diagonal. The existence of the block matrix \mathbf{L} helps to enable the solution of the linear systems with a number of algebraic and domain decomposition type preconditioners that rely on non-pivoting ILU type factorization, or in some cases methods such as Jacobi or Gauss-Seidel as sub-domain solvers [104,76].

The difficulty of producing robust and efficient preconditioners for Galerkin (as well as centered finite difference and non-staggered finite volume) formulations has motivated the use of many different types of decoupled solution methods. Often, transient schemes combine semi-implicit methods with fractional-step (operator splitting) approaches or use fully-decoupled solution strategies. In these cases, the motivation is to reduce memory usage and to produce a simplified equation set for which efficient solution strategies already exist. Unfortunately, these simplifications place significant limitations on the broad applicability of these methods. For example, fractional-step methods such as pressure projection [105–107] and operator splitting [108] require time step limitations based on the explicit part of the time integration process as well as on the stability and accuracy associated with the decoupled physics [109–114]. This restriction can severely limit the step size, and direct-to-steady-state simulations with these methods are not possible. Fully-decoupled solution strategies (e.g., the SIMPLE [115], SIMPLER [116], and PISO [117] class of methods) use a successive substitution (or Picard) iteration to simplify the coupled systems of equations. Nonlinearities at each time step are resolved by an outer nonlinear iteration. Unfortunately, while this technique should improve time step limitations, time steps are frequently reduced to facilitate the nonlinear iteration. Convergence of these decoupled methods can often be problematic. In particular, the nonlinear iteration has only a linear rate of convergence and in practice can often exhibit very slow convergence. In addition, since all the equations have been decoupled artificially, this strategy can sometimes result in non-convergence for difficult problems in which the essential coupling of the physics has been violated (see for example [118,119], and the references contained therein). A detailed presentation of the characteristics of current solution methods is far beyond the scope of this brief overview. The intent of our method of fully-coupling the resistive MHD PDEs in the nonlinear solver is to preserve the inherently strong coupling of the physics with the goal to produce a more robust solution methodology. Preservation of this strong coupling, however, places a significant burden on the linear solution procedure to solve the fully coupled algebraic systems.

Finally it should be noted that in our actual linear algebra implementation we use a specific ordering of the unknowns locally at each FE node with each degree of freedom ordered consecutively. A single coupled matrix problem, $\mathbf{J}\mathbf{s} = -\mathbf{F}$, is solved at each Newton step with sophisticated algebraic domain

decomposition and multilevel preconditioned Krylov methods as described in Section 4.2.

References

- [1] H. Goedbloed, S. Poedts, Principles of Magnetohydrodynamics with Applications to Laboratory and Astrophysical Plasmas, Cambridge Univ. Press, 2004.
- [2] K. E. Brenan, S. L. Campbell, L. R. Petzold, Numerical Solution of Initial-Value Problems in Differential-Algebraic Equations, SIAM, Classics in Applied Math, 1996.
- [3] U. M. Ascher, L. R. Petzold, Computer Methods for Ordinary Differential Equations and Differential-Algebraic Equations, SIAM, 1998.
- [4] W. Dai, P. R. Woodward, On the divergence-free condition and conservation laws in numerical simulations for supersonic magnetohydrodynamic flows, *Astrophys. J.* 494 (1998) 317.
- [5] W. Dai, P. R. Woodward, A simple finite difference scheme for multidimensional magnetohydrodynamic equations, *J. Comput. Phys.* 142 (1998) 331.
- [6] D. Ryu, F. Miniati, T. W. Jones, A. Frank, A divergence-free upwind code for multi-dimensional magnetohydrodynamics flows, *Astrophys. J.* 509 (1998) 244.
- [7] D. S. Balsara, D. S. Spicer, A staggered mesh algorithm using high order Godunov fluxes to ensure solenoidal magnetic fields in magnetohydrodynamics simulations, *J. Comput. Phys.* 149 (1999) 270–292.
- [8] D. S. Balsara, Divergence-free adaptive mesh refinement for magnetohydrodynamics, *J. Comput. Phys.* 174 (2001) 614–648.
- [9] G. Tóth, The $\nabla \cdot \mathbf{B} = \mathbf{0}$ constraint in shock-capturing magnetohydrodynamics codes, *J. Comput. Phys.* 161 (2000) 605–652.
- [10] G. Tóth, R. Keppens, M. A. Botchev, Implicit and semi-implicit schemes in the Versatile Advection Code: numerical tests, *Astron. Astrophys.* 332 (1998) 1159–1170.
- [11] R. Keppens, G. Tóth, M. A. Botchev, A. V. D. Ploeg, Implicit and semi-implicit schemes: algorithms, *Int. J. Numer. Meth. Fl.* 30 (1999) 335–352.
- [12] A. Y. Aydemir, D. C. Barnes, An implicit algorithm for compressible three-dimensional magnetohydrodynamic calculations., *J. Comput. Phys.* 59 (1) (1985) 108 – 19.
- [13] W. Park, J. Breslau, J. Chen, G. Y. Fu, S. C. Jardin, S. Klasky, J. Menard, A. Pletzer, B. C. Stratton, D. Stutman, H. R. Strauss, L. E. Sugiyama, Nonlinear simulation studies of tokamaks and STS., *Nucl. Fusion* 43 (6) (2003) 483 – 9.
- [14] S. C. Jardin, J. A. Breslau, Implicit solution of the four-field extended-magnetohydrodynamic equations using high-order high-continuity finite elements., *Phys. Plasmas* 12 (5) (2005) 056101.
- [15] D. S. Harned, W. Kerner, Semi-implicit method for three-dimensional compressible magnetohydrodynamic simulation, *J. Comput. Phys.* 60 (1985) 62–75.
- [16] D. S. Harned, D. D. Schnack, Semi-implicit method for long time scale magnetohydrodynamic computations in three dimensions, *J. Comput. Phys.* 65 (1986) 57–70.
- [17] D. D. Schnack, D. C. Barnes, D. S. Harned, E. J. Caramana, Semi-implicit magnetohydrodynamic calculations, *J. Comput. Phys.* 70 (1987) 330–354.

- [18] D. S. Harned, Z. Mikic, Accurate semi-implicit treatment of the Hall effect in magnetohydrodynamic computations, *J. Comput. Phys.* 83 (1989) 1–15.
- [19] C. R. Sovinec, A. H. Glasser, T. A. Gianakon, D. C. Barnes, R. A. Nebel, S. E. Kruger, D. D. Schnack, S. J. Plimpton, A. Tarditi, M. S. Chu, the NIMROD team, Nonlinear magnetohydrodynamics simulation using high-order finite elements, *J. Comput. Phys.* 195 (1) (2004) 355–386.
- [20] A. Hujeirat, IRMHD: an implicit radiative and magnetohydrodynamical solver for self-gravitating systems, *Mon. Not. R. Astron. Soc.* 298 (1998) 310–320.
- [21] A. C. Robinson, C. J. Garasi, Three-dimensional Z-pinch wire array modeling with ALEGRA-HEDP, *Comput. Phys. Commun.* 164 (2004) 408–413.
- [22] A. C. Robinson, et. al., ALEGRA: An arbitrary Lagrangian-Eulerian multimaterial, multiphysics code, in: *AIAA 2008-1235 46th AIAA Aerospace Sciences Meeting and Exhibit*, Reno, NV, 2008.
- [23] G. Toth, R. Keppens, Versatile advection code, <http://www.phys.uu.nl/~toth/>.
- [24] D. A. Knoll, L. Chacón, L. Margolin, V. A. Mousseau, On balanced approximations for the time integration of multiple time scale systems, *J. Comput. Phys.* 185 (2003) 583–611.
- [25] A. Hujeirat, R. Rannacher, On the efficiency and robustness of implicit methods in computational astrophysics, *New Astron. Revv* 45 (2001) 425–447.
- [26] D. R. Reynolds, R. Samtaney, C. S. Woodward, A fully implicit numerical method for single-fluid resistive magnetohydrodynamics., *J. Comput. Phys.* 219 (1) (2006) 144 – 62.
- [27] S. Ovtchinnikov, F. Dobrian, X.-C. Cai, D. Keyes, Additive Schwarz-based fully coupled implicit methods for resistive Hall magnetohydrodynamic problems, *J. Comput. Phys.* 225 (2007) 1919 – 1936.
- [28] L. Chacón, D. A. Knoll, J. M. Finn, Implicit, nonlinear reduced resistive MHD nonlinear solver, *J. Comput. Phys.* 178 (1) (2002) 15–36.
- [29] L. Chacón, D. A. Knoll, A 2D high- β Hall MHD implicit nonlinear solver, *J. Comput. Phys.* 188 (2) (2003) 573–592.
- [30] L. Chacón, An optimal, parallel, fully implicit Newton-Krylov solver for three-dimensional visco-resistive magnetohydrodynamics, *Phys. Plasmas* 15 (2008) 056103.
- [31] L. Chacón, Scalable solvers for 3D magnetohydrodynamics, *J. Physics: Conf. Series* 125 (2008) 012041.
- [32] D. R. Reynolds, R. Samtaney, C. S. Woodward, Operator-based preconditioning of stiff hyperbolic systems, *SIAM J. Sci. Comput.* in press.
- [33] N. B. Salah, A. Soulaïmani, W. G. Habashi, M. Fortin, A conservative stabilized finite element method for the magnetohydrodynamics equations, *Int. J. Num. Meth. Fluids* 29 (1999) 535–554.
- [34] F. Brezzi, On existence, uniqueness and approximation of saddle-point problems arising from Lagrange multipliers, *RAIRO Model. Math. Anal. Numer.* 21 (1974) 129–151.
- [35] M. Gunzburger, *Finite Element Methods for Viscous Incompressible Flows*, Academic Press, Boston, 1989.
- [36] T. Barth, P. Bochev, M. Gunzburger, J. Shadid, A taxonomy of consistently stabilized finite element methods for the Stokes problem, *SIAM J. Sci. Comput.* 25 (5) (2004) 1585–1607.

- [37] F. Brezzi, J. Pitkaranta, On the stabilization of finite element approximations of the stokes problem, *Efficient Solutions of Elliptic Systems*, Vieweg Notes on Numerical Fluid Mechanics, Vieweg, Wiesbaden, W. Hackbusch (ed.) 10 (1984) 11–19.
- [38] N. B. Salah, A. Soulaïmani, W. G. Habashi, A finite element method for magnetohydrodynamics, *Comput. Meth. Appl. M.* 190 (2001) 5867–5892.
- [39] R. Codina, N. Hernandez-Silva, Stabilized finite element approximation of the stationary magneto-hydrodynamics equations, *Comput. Mech.* 38 (4-5) (2006) 344–355.
- [40] J.-F. Gerbeau, A stabilized finite element method for the incompressible magnetohydrodynamic equations, *Numer. Math.* 87 (2000) 83–111.
- [41] S. Lankalapalli, J. Flaherty, M. Shephard, H. Strauss, An adaptive finite element method for magnetohydrodynamics, *J. Comput. Phys.* 225 (2007) 363–381.
- [42] P. N. Brown, Y. Saad, Convergence theory of nonlinear Newton–Krylov algorithms, *SIAM J. Optimiz.* 4 (1994) 297–330.
- [43] S. Eisenstat, H. Walker, Globally convergent inexact Newton methods, *SIAM J. Optimiz.* 4 (1994) 393–422.
- [44] B. Smith, P. Bjorstad, W. Gropp, *Domain Decomposition: Parallel Multilevel methods for elliptic partial differential equations*, Cambridge University Press, 1996.
- [45] A. Quarteroni, A. Valli, *Domain Decomposition Methods for Partial Differential Equations*, Oxford University Press, Oxford, 1999.
- [46] M. Sala, J. N. Shadid, R. S. Tuminaro, An improved convergence bound for aggregation-based domain decomposition preconditioners, *SIAM J. Matrix Anal. A.* 27 (3) (2006) 744–756.
- [47] P. T. Lin, M. Sala, J. N. Shadid, R. S. Tuminaro, Performance of fully-coupled algebraic multilevel domain decomposition preconditioners for incompressible flow and transport, *Int. J. Numer. Meth. Eng.* 67 (9) (2006) 208–225.
- [48] H. R. Strauss, Nonlinear, 3-dimensional magnetohydrodynamics of noncircular tokamaks, *Phys. Fluids* 19 (1) (1976) 134–140.
- [49] R. D. Hazeltine, M. Kotschenreuther, P. J. Morrison, A four-field model for tokamak plasma dynamics, *Phys. Fluids* 28 (8) (1985) 2466–2477.
- [50] J. F. Drake, T. M. Antonsen, Nonlinear reduced fluid equations for toroidal plasmas, *Phys. Fluids* 27 (4) (1984) 898–908.
- [51] R. Moreau, *Magnetohydrodynamics*, Kluwer, Dordrecht, 1990.
- [52] P. A. Davidson, *An Introduction to Magnetohydrodynamics*, Cambridge Univ. Press, 2001.
- [53] J. Donea, A. Huerta, *Finite Element Methods for Flow Problems*, John Wiley, 2002.
- [54] A. Dedner, F. Kemm, D. Kroner, C.-D. Munz, T. Schnitzer, M. Wesenberg, Hyperbolic divergence cleaning for the MHD equations, *J. Comput. Phys.* 175 (2002) 645–673.
- [55] L. Chacón, A non-staggered, conservative, $\nabla \cdot \mathbf{B} = \mathbf{0}$, finite-volume scheme for 3D implicit extended magnetohydrodynamics in curvilinear geometries, *Comput. Phys. Commun.* 163 (2004) 143–171.
- [56] J. D. Jackson, *Classical Electrodynamics*, 2nd Edition, John Wiley & Sons, 1975.

- [57] T. Hughes, A. N. Brooks, A multidimensional upwind scheme with no cross-wind diffusion, in T.J.R. Hughes (ed.), *Finite Element Methods for Convection Dominated Flows*, AMD Vol. 34, ASME, New York 19-35.
- [58] A. N. Brooks, T. Hughes, Streamline upwind/petrov-galerkin formulations for convection dominated flows with particular emphasis on the incompressible Navier-Stokes equations, *Comput. Meth. Appl. M.* 32 (1982) 199–259.
- [59] T. Hughes, A. Brooks, A theoretical framework for Petrov-Galerkin methods with discontinuous weighting functions: Application to the streamline-upwind procedure, in: R. G. et al (Ed.), *Finite Elements in Fluids*, Vol. 4, J. Willey & Sons, 1982, pp. 47–65.
- [60] T. J. R. Hughes, L. P. Franca, G. M. Hulbert, A new finite element formulation for computational fluid dynamics: VII. the Galerkin/Least-Squares method for advective-diffusive equations, *Comput. Meth. Appl. M.* 73 (1989) 173–189.
- [61] T. Hughes, L. Franca, M. Balestra, A new finite element formulation for computational fluid dynamics: V. Circumventing the Babuska-Brezzi condition: A stable Petrov-Galerkin formulation of the Stokes problem accommodating equal-order interpolations, *Comput. Meth. Appl. M.* 59 (1986) 85–99.
- [62] T. Hughes, M. Mallet, A new finite element formulation for computational fluid dynamics: III. The generalized streamline operator for multidimensional advective-diffusive systems, *Comput. Meth. Appl. M.* 58 (1986) 305–328.
- [63] F. Shakib, *Finite element analysis of the compressible Euler and Navier-Stokes equations*, Ph.D. thesis, Division of Applied Mathematics, Stanford University (1989).
- [64] T. Hughes, Multiscale phenomena: Green’s functions, the Dirichlet-to-Neumann formulation, subgrid scale models, bubbles and the origins of stabilized methods, *Comput. Meth. Appl. M.* 127 (1995) 387–401.
- [65] T. Tezduyar, Stabilized finite element formulations for incompressible flow calculations, *Adv. Appl. Mech.* 28 (1992) 1–44.
- [66] J. E. Dennis, Jr., R. B. Schnabel, *Numerical Methods for Unconstrained Optimization and Nonlinear Equations*, Series in Automatic Computation, Prentice-Hall, Englewood Cliffs, NJ, 1983.
- [67] P. N. Brown, Y. Saad, Hybrid Krylov methods for nonlinear systems of equations, *SIAM J. Sci. Stat. Comp.* 11 (1990) 450–481.
- [68] R. P. Pawlowski, J. N. Shadid, J. P. Simonis, H. F. Walker, Globalization techniques for Newton-Krylov methods and applications to the fully-coupled solution of the Navier-Stokes equations, *SIAM Rev.* 48 (2006) 700–721.
- [69] R. S. Dembo, S. C. Eisenstat, T. Steihaug, Inexact Newton methods, *SIAM J. Numer. Anal.* 19 (1982) 400–408.
- [70] S. C. Eisenstat, H. F. Walker, Globally convergent inexact Newton methods, *SIAM J. Optimiz.* 4 (1994) 393–422.
- [71] S. C. Eisenstat, H. F. Walker, Choosing the forcing terms in an inexact Newton method, *SIAM J. Sci. Comput.* 17 (1996) 16–32.
- [72] 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, An Overview of Trilinos Project, *ACM Trans. Math. Software* 31 (3) (2005) 397–423.
- [73] Y. Saad, *Iterative Methods for Sparse Linear Systems*, SIAM, 2003.

- [74] O. Axelsson, *Iterative Solution Methods*, Cambridge University Press, New York, 1994.
- [75] R. Tuminaro, C. Tong, J. Shadid, K.D.Devine, D. Day, On a multilevel preconditioning module for unstructured mesh Krylov solvers: two-level Schwarz, *Comm. Num. Method. Eng.* 18 (2002) 383–389.
- [76] J. Shadid, R. Tuminaro, K. Devine, G. Henningan, P. Lin, Performance of fully-coupled domain decomposition preconditioners for finite element transport/reaction simulations, *J. Comput. Phys.* 205 (1) (2005) 24–47.
- [77] U. Trottenberg, C. Oosterlee, A. Schüller, *Multigrid*, Academic Press, London, 2001.
- [78] W. L. Briggs, V. E. Henson, S. McCormick, *A Multigrid Tutorial*, Second Edition, SIAM, Philadelphia, 2000.
- [79] J. N. Shadid, A. G. Salinger, R. P. Pawlowski, P. T. Lin, G. L. Hennigan, R. S. Tuminaro, R. B. Lehoucq, Stabilized FE computational analysis of nonlinear steady state transport/reaction systems, *Comput. Meth. Appl. M.* 195 (2006) 1846–1871.
- [80] J. Ruge, K. Stüben, Algebraic multigrid (AMG), in: S. F. McCormick (Ed.), *Multigrid Methods*, Vol. 3 of *Frontiers in Applied Mathematics*, SIAM, Philadelphia, PA, 1987, pp. 73–130.
- [81] G. Karypis, V. Kumar, Parallel multilevel k-way partitioning scheme for irregular graphs, *ACM/IEEE Proceedings of SC96: High Performance Networking and Computing*.
- [82] P. Vaněk, J. Mandel, M. Brezina, Algebraic multigrid by smoothed aggregation for second and fourth order elliptic problems, *Computing* 56 (1996) 179–196.
- [83] P. Vaněk, M. Brezina, J. Mandel, Convergence of algebraic multigrid based on smoothed aggregation, *Numer. Math.* 88 (2001) 559–579.
- [84] M. Sala, R. Tuminaro, A new Petrov-Galerkin smoothed aggregation preconditioner for nonsymmetric linear systems, *SIAM J. Sci. Comput.* 31 (1) (2008) 143–166.
- [85] M. Sala, J. Hu, R. Tuminaro, ML 3.1 smoothed aggregation user’s guide, Tech. Rep. SAND2004-4819, Sandia National Laboratories (September 2004).
- [86] T. Davis, *Direct methods for sparse linear systems*, SIAM, Philadelphia, PA, 2006.
- [87] M. Sala, Amesos 2.0 reference guide, Tech. Rep. SAND2004-4820, Sandia National Laboratories (September 2004).
- [88] M. Gee, C. Siefert, J. Hu, R. Tuminaro, M. Sala, ML 5.0 smoothed aggregation user’s guide, Tech. Rep. SAND2006-2649, Sandia National Laboratories, Albuquerque, NM, 87185 (2006).
- [89] P. T. Lin, M. Sala, J. N. Shadid, R. S. Tuminaro, Performance of a geometric and an algebraic multilevel preconditioner for incompressible flow and transport, in: *Computational Mechanics: WCCM VI in conjunction with APCOM’04*, Tsinghua Univ. Press, Springer-Verlag, 2004.
- [90] P. T. Lin, J. N. Shadid, M. Sala, R. Tuminaro, G. L. Hennigan, R. J. Hoekstra, Performance of a parallel algebraic multilevel preconditioner for stabilized finite element semiconductor device modeling, *J. Comput. Phys.* 228 (17) (2009) 6079–6616.
- [91] P. T. Lin, J. N. Shadid, R. S. Tuminaro, M. Sala, Performance of a Petrov-Galerkin algebraic multilevel preconditioner for finite element modeling of the semiconductor device drift-diffusion equations, *Int. J. Numer. Meth. Eng.* (2010) in Press.
- [92] H. K. Moffat, *Magnetic Field Generation in Electrically Conducting Fluids*, Cambridge Univ. Press, 1983.

- [93] R. Codina, N. Hernandez-Silva, Stabilized finite element approximation of the stationary magnetohydrodynamics equations, *Comput. Mech.* 38 (2006) 344–355.
- [94] M. Ainsworth, J. T. Oden, *A posteriori error estimation in finite element analysis*, John Wiley, 2000.
- [95] G. K. Batchelor, *An Introduction to Fluid Mechanics*, Cambridge Univ. Press, 1967.
- [96] V. M. Fadeev, I. F. Kvartskhava, N. N. Komarov, Self-focusing of local plasma currents., *Nucl. Fusion* 5 (3) (1965) 202 – 209.
- [97] M. Hughes, K. A. Pericleous, M. Cross, The numerical modelling of dc electromagnetic pump and brake flow, *Appl. Math. Model.* 19 (1995) 713–723.
- [98] S. L. L. Verardi, J. R. Cardoso, M. C. Costa, Three-dimensional finite element analysis of MHD duct flow by the penalty function formulation, *IEEE T. Magn.* 37 (5) (2001) 3384–3387.
- [99] N. B. Morley, M. S. Tillack, *Magnetohydrodynamics*, chapter in the *Standard Handbook for Electrical Engineers*, 15th Edition, McGraw Hill, 2006.
- [100] S. Chandrasekhar, *Hydrodynamic and Hydromagnetic Stability*, Oxford Univ. Press, 1961.
- [101] D. Biskamp, *Magnetic Reconnection in Plasmas*, Cambridge University Press, Cambridge, UK, 2000.
- [102] D. A. Knoll, L. Chacón, Coalescence of magnetic islands, sloshing, and the pressure problem., *Phys. Plasmas* 13 (3) (2006) 32307 – 1.
- [103] E. Priest, T. Forbes, *Magnetic Reconnection: MHD Theory and Applications*, Cambridge Univ. Press, 2006.
- [104] J. Shadid, A fully-coupled Newton-Krylov solution method for parallel unstructured finite element fluid flow, heat and mass transfer simulations, *Int. J. CFD* 12 (1999) 199–211.
- [105] J. B. Bell, P. Colella, H. M. Glaz, A second-order projection method for the incompressible Navier–Stokes equations, *J. Comput. Phys.* 85 (1989) 257.
- [106] A. J. Chorin, A numerical method for solving incompressible viscous problems, *J. Comput. Phys.* 2 (1967) 12.
- [107] P. M. Gresho, On the theory of semi-implicit projection methods for viscous incompressible flow and its implementation via a finite element method that also introduces a nearly consistent mass matrix. part 1: Theory, *Int. J. Num. Meth. Fluids* 11 (1990) 587–620.
- [108] E. S. Oran, J. P. Boris, *Numerical Simulation of Reactive Flow*, Cambridge Univ. Press, Cambridge, 2001.
- [109] R. Codina, Pressure stability in fractional step finite element methods for incompressible flow, *J. Comput. Phys.* 170 (2001) 112–140.
- [110] J. L. Guermond, L. Quartapelle, On stability of convergence of projection methods based on pressure poisson equation, *Int. J. Numer. Meth. Fl.* 26 (1998) 1039–1053.
- [111] G. E. Karniadakis, M. Israeli, S. A. Orszag, High-order splitting methods for the incompressible Navier–Stokes equations, *J. Comput. Phys.* 97 (1991) 414–443.
- [112] J. B. Perot, An analysis of the fractional step method, *J. Comput. Phys.* 108 (1993) 51–58.
- [113] A. Quateroni, F. Saleri, A. Veneziani, Factorization methods for the numerical approximation of Navier–Stokes equations, *Comput. Meth. Appl. M.* 188 (2000) 505–526.

- [114] J. C. Strikwerda, Y. S. Lee, The accuracy of the fractional step method, *SIAM J. Numer. Anal.* 37 (1) (1999) 37–47.
- [115] S. V. Patankar, D. A., A calculation procedure for heat, mass and momentum transfer in three dimensional parabolic flows, *Int. J. Heat Mass Tran.* 15 (1972) 1787–1806.
- [116] S. V. Patankar, *Numerical heat transfer and fluid flow*, Hemisphere Pub. Corp, New York, 1980.
- [117] R. Issa, Solution of the implicitly discretized fluid flows equations by operator splitting, *J. Comp. Phys* 62(1) (1986) 40–65.
- [118] G. B. Deng, J. Piquet, P. Queutey, M. A. Visonneau, A new fully coupled solution of the Navier–Stokes equations, *Int. J. Numer. Meth. Fl.* 19 (1994) 605–639.
- [119] G. B. Deng, J. Piquet, X. Vasseur, M. A. Visonneau, A new fully coupled method for computing turbulent flows, *Comput. Fluids* 30 (2001) 445–472.