

INITIAL PERFORMANCE OF FULLY-COUPLED AMG AND APPROXIMATE BLOCK FACTORIZATION PRECONDITIONERS FOR SOLUTION OF IMPLICIT FE RESISTIVE MHD

John N. Shadid*, Eric C. Cyr*, Roger P. Pawlowski*, Ray S. Tuminaro*, Luis Chacón†, and Paul T. Lin*

*Sandia National Laboratories,
Albuquerque, NM
{jnshadi, eccyr, rppawlo, rstumin, ptlin}@sandia.gov

†Oak Ridge National Laboratory
Oak Ridge, TN
chaconl@ornl.gov

Key words: Resistive MHD, Implicit Time Integration, Preconditioned Newton-Krylov, AMG, Block Factorization, Physics-based, Stabilized FE

Abstract. *This brief paper explores the development of scalable, nonlinear, fully-implicit solution methods for a stabilized unstructured finite element (FE) discretization of the 2D incompressible (reduced) resistive MHD system. The discussion considers the 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 (AMG) techniques and a new approximate block factorization (ABF) preconditioner. The intent of these preconditioners is to enable robust, scalable and efficient solution approaches for the large-scale sparse linear systems generated by the Newton linearization. We present results for the fully-coupled AMG preconditioner for two prototype problems, a low Lundquist number MHD Faraday conduction pump and moderately-high Lundquist number incompressible magnetic island coalescence problem. For the MHD pump results we explore the scaling of the fully-coupled AMG preconditioner for up to 4096 processors for problems with up to 64M unknowns on a CrayXT3/4. Using the island coalescence problem we explore the weak scaling of the AMG preconditioner and the influence of the Lundquist number on the iteration count. Finally we present some very recent results for the algorithmic scaling of the ABF preconditioner.*

This work was partially supported by the DOE Office of Science AMR program at Sandia National Laboratory under contract DE-AC04-94AL85000, and Oak Ridge National Laboratory under contract DE-AC05-00OR22725.

1 INTRODUCTION

The magnetohydrodynamics (MHD) model describes the dynamics of charged fluids in the presence of electromagnetic fields. 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. The resistive MHD model 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 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^{4, 1}. 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 and partially implicit methods that include implicit-explicit, semi-implicit, and operator-splitting time integration methods (see⁵ and the references contained therein). 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. 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^{24, 25, 45, 34, 31, 7, 5, 6}.

This study complements previous work by exploring the development of a robust, efficient, fully-coupled stabilized FE formulation for incompressible 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 to solve the resulting large-scale nonlinear algebraic systems. For preconditioning, we compare well-known variable-overlap, additive, one-level Schwarz domain-decomposition methods³³, an algebraic multilevel (AMG) technique employing a graph-based aggressive-coarsening aggregation method^{28, 40}, and an initial approximate block factorization (ABF) technique that utilizes the above AMG solver as a sub-block physics solver. The AMG and ABF preconditioners effectively use 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, isothermal, resistive MHD formulation. This formulation is suitable for reduced MHD models^{44, 20, 14} as well as low magnetic-

Reynolds-number liquid metal MHD^{29, 10}. 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 (FE), based on the general approach of Hughes et. al. (see¹³ and references therein).

2 REDUCED MHD MODEL EQUATIONS

Our base MHD model is the one-fluid visco-resistive MHD system¹⁸. This model provides a continuum description of charged fluids in the presence of electromagnetic fields. Formally, isothermal visco-resistive MHD augments the Navier-Stokes fluid description with a magnetic stress term in the momentum 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:

Momentum Conservation:

$$\mathbf{R}_{\mathbf{u}} = \rho \frac{\partial \mathbf{u}}{\partial t} + \rho(\mathbf{u} \cdot \nabla \mathbf{u}) - \nabla \cdot (\mathbf{T} + \mathbf{T}_{\mathbf{M}}) = 0 \quad (1)$$

Total Mass Conservation:

$$R_P = \frac{\partial \rho}{\partial t} + \nabla \cdot (\rho \mathbf{u}) = 0 \quad (2)$$

Magnetics Evolution Equation:

$$\mathbf{R}_{\mathbf{B}} = \frac{\partial \mathbf{B}}{\partial t} - \nabla \times (\mathbf{u} \times \mathbf{B}) + \nabla \times \left(\frac{\eta}{\mu_0} \nabla \times \mathbf{B} \right) = \mathbf{0}. \quad (3)$$

Where

$$\begin{aligned} \mathbf{T} &= -P\mathbf{I} + \mathbf{\Pi} = -P\mathbf{I} - \frac{2}{3}\mu(\nabla \cdot \mathbf{u})\mathbf{I} + \mu[\nabla \mathbf{u} + \nabla \mathbf{u}^T] \\ \mathbf{T}_{\mathbf{M}} &= \frac{1}{\mu_0} \mathbf{B} \otimes \mathbf{B} - \frac{1}{2\mu_0} \|\mathbf{B}\|^2 \mathbf{I} \end{aligned}$$

In these equations the unknowns are the velocity vector \mathbf{u} , the hydrodynamic pressure P , and the magnetic flux, \mathbf{B} . The transport properties, ρ, μ, η, μ_0 , are respectively, the density, dynamic viscosity, magnetic resistivity, and μ_0 is the magnetic permeability of free space. Constitutive equations define the Newtonian stress tensor, \mathbf{T} . Ampere's law neglecting the displacement current provides the plasma current, $\mathbf{J} = 1/\mu_0 \nabla \times \mathbf{B}$, and the tensor, $\mathbf{T}_{\mathbf{M}}$, is the magnetic stress tensor.

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^{29, 10}, and high-Lundquist-number, large-guide-field fusion plasmas^{44, 20, 14}. 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 equations (1) – (3) 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)^{44, 20, 14}. 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 [see¹³ and references therein]). Thus, we preserve a 2D form of equation (1), and we replace the continuity equation (2) by $\nabla \cdot \mathbf{v} = 0$.

In regards to the magnetic field evolution equation, we replace Eq. (3) by an evolution equation for the vector potential component in the ignorable direction, $\mathbf{A} = (0, 0, A_z)$, which reads:

$$\frac{\partial A_z}{\partial t} + \mathbf{v} \cdot \nabla A_z - \frac{\eta}{\mu_0} \nabla^2 A_z + E_z^0 = 0. \quad (4)$$

This equation is equivalent to equation (3) in two dimensions.

3 A STABILIZED FE FORMULATION FOR 2D RESISTIVE MHD AND THE DISCRETE SYSTEM

Table 1 presents the governing equations in convected form, for momentum, total mass, 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.¹⁹). 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. 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 stabilized 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 to the weak equations to limit oscillations in highly convected flows¹³.

Momentum	$\mathbf{R}_m = \rho \frac{\partial \mathbf{v}}{\partial t} + \rho(\mathbf{v} \cdot \nabla \mathbf{v}) + \nabla \cdot \left(-\frac{1}{\mu_0} \mathbf{B} \otimes \mathbf{B} - \boldsymbol{\Pi} + (P + \frac{1}{2\mu_0} \ \mathbf{B}\ ^2) \mathbf{I} \right)$
Total Mass	$R_P = \nabla \cdot \mathbf{v}$
Z-Vector Potential	$R_{A_z} = \frac{\partial A_z}{\partial t} + \mathbf{v} \cdot \nabla A_z - \frac{\eta}{\mu_0} \nabla^2 A_z + E_z^0$ $\mathbf{B} = \nabla \times \mathbf{A}; \quad \mathbf{A} = (0, 0, A_z)$

Table 1: Residual form of governing 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 , and the A_z component of the vector potential in 2D.

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	$F_P = \int_{\Omega} \Phi R_P d\Omega + \sum_e \int_{\Omega_e} \rho \hat{\tau}_m \nabla \Phi \cdot \mathbf{R}_m d\Omega$
Z-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. Here Φ is a global weighting function used to formally define the weak form. The sum \sum_e indicates the integrals are taken only over element interiors Ω_e and integration by parts is not performed.

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 Mallet²³ and Shakib⁴³ for Navier-Stokes with an adaptation of the stabilized formulation of Codina and Hernandez-Silva⁹ for a resistive MHD system. Details of this formulation can be found in⁴⁰.

The discrete form of the matrix equations that results from the stabilized FE discretization of the governing balance PDEs is

$$\begin{bmatrix} \mathbf{F} & \mathbf{B}^T & \mathbf{Z} \\ \mathbf{B} & \mathbf{C} & \mathbf{0} \\ \mathbf{Y} & \mathbf{0} & \mathbf{D} \end{bmatrix} \begin{bmatrix} \Delta \mathbf{v} \\ \Delta \mathbf{p} \\ \Delta \mathbf{A}_z \end{bmatrix} = \begin{bmatrix} -\mathbf{F}_v \\ -\mathbf{F}_p \\ -\mathbf{F}_{A_z} \end{bmatrix}. \quad (5)$$

In this representation, the vectors, $\Delta \mathbf{v}$, $\Delta \mathbf{p}$, $\Delta \mathbf{A}_z$, contain the Newton updates to the nodal velocities, pressures and vector potential respectively. The block matrix, \mathbf{F} , corre-

sponds to the combined discrete transient, convection, diffusion and stress terms acting on the unknowns $\Delta \mathbf{v}$; \mathbf{B}^T , corresponds to the discrete gradient operator; \mathbf{Z} the Lorentz force operator; \mathbf{B} , the divergence operator; \mathbf{C} , corresponds to the discrete “pressure Laplacian” type operator that is generated by the pressure stabilization¹³; \mathbf{Y} , is a vector mass-matrix type operator scaled by the gradient components of A_z ; and \mathbf{D} is a combined discrete transient, convection, diffusion operator acting on $\Delta \mathbf{A}_z$. The vectors \mathbf{F}_v , \mathbf{F}_P , and \mathbf{F}_{A_z} contain the right hand side residuals for Newton’s method. The existence of the weak form Laplacian matrix, \mathbf{C} , 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{C} helps to enable the fully-coupled 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^{41, 40}.

4 FULLY-IMPLICIT FULLY-COUPLED SOLUTION BY NEWTON-KRYLOV METHODS

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 integration¹). The implicit time integration methods used in this study is a second-order A-stable implicit midpoint rule.

The result of a fully-implicit or 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 solvers¹² are employed 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. The robustness and efficiency of these methods rely on effective preconditioning methods. We describe the methods employed in this study below.

4.1 Schwarz domain decomposition preconditioners

For the considered class of linear systems described above, convergence is not achieved without preconditioning due to ill-conditioning in the underlying matrix equations³⁶. This paper considers 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 processor³³. 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 factorization³⁶. 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 coupling³³. The convergence rate also deteriorates as the number of unknowns per subdomain increases when $ILU(k)$ is used for a subdomain solver. To improve algorithmic performance, coarse levels can be introduced to approximate global coupling^{33, 47, 41} and produce a multilevel preconditioner. 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 methods⁴⁶.

4.2 Fully-coupled AMG preconditioners

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^{47, 41, 42}. 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, pressure and vector potential 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 small³⁵. In this study, METIS and ParMETIS²⁶ are used to group fine graph vertices into aggregates so that each aggregate effectively represents a coarse graph 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^{49, 48}. In this study, we employ both a non-smoothed (NSA) and Petrov-Galerkin smoothed aggregation (PGSA) algorithm as implemented in the ML multilevel package in Trilinos and described in^{37, 39}.

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^{38, 37}. Additionally, larger aggregates are consistent with using a sub-domain solver based on Schwarz/ $ILU(k)$ 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 $ILU(k)$ algorithm is used as a smoother on each level and on the coarsest level the KLU¹¹ sparse direct solver is employed.

4.3 Approximate block factorization preconditioners

In this study we also consider preconditioners based on approximate block factorization (ABF) methods^{8, 30, 2, 15}. ABF preconditioners for systems, carefully consider the spectral properties of the component block operators and the properties of the approximate Schur complement operators that are employed. Through this linear algebraic view of preconditioning, a simplified system of block component equations is developed that encodes a specific “physics-based” decomposition for many applications. An important goal of the ABF approach is to produce scalar or vector sub-systems to which AMG methods are more effectively applied as optimal solvers.

To motivate our initial ABF preconditioner of the resistive MHD system consider the block LU factorization of the Jacobian in Eqn. (5):

$$\begin{bmatrix} \mathbf{F} & \mathbf{B}^T & \mathbf{Z} \\ \mathbf{B} & \mathbf{C} & \mathbf{0} \\ \mathbf{Y} & \mathbf{0} & \mathbf{D} \end{bmatrix} = \begin{bmatrix} \mathbf{I} & \mathbf{0} & \mathbf{0} \\ \mathbf{B}\mathbf{F}^{-1} & \mathbf{I} & \mathbf{0} \\ \mathbf{Y}\mathbf{F}^{-1} & -\mathbf{Y}\mathbf{F}^{-1}\mathbf{B}^T\mathbf{S}^{-1} & \mathbf{I} \end{bmatrix} \begin{bmatrix} \mathbf{F} & \mathbf{B}^T & \mathbf{Z} \\ \mathbf{0} & \mathbf{S} & -\mathbf{B}\mathbf{F}^{-1}\mathbf{Z} \\ \mathbf{0} & \mathbf{0} & \mathbf{W} \end{bmatrix} \quad (6)$$

where

$$\mathbf{S} = \mathbf{C} - \mathbf{B}\mathbf{F}^{-1}\mathbf{B}^T \quad (7)$$

$$\mathbf{W} = \mathbf{D} - \mathbf{Y}\mathbf{F}^{-1}(\mathbf{I} + \mathbf{B}^T\mathbf{S}^{-1}\mathbf{B}\mathbf{F}^{-1})\mathbf{Z}. \quad (8)$$

This factorization has two Schur complements. First, Eq. (7) is recognized as the fluid Schur complement derived from the Navier-Stokes equations¹⁵. While Eq. (8) is Maxwell’s equation Schur complement with an embedded contribution from the fluid Schur complement. The nested nature of \mathbf{W} presents a challenge for effective approximation. Unfortunately reordering the blocks of the incompressible MHD system does not alleviate the problem of nested Schur complements.

To address this issue we have developed an initial approach where the fluid coupling and the velocity/magnetics coupling are applied in sequence by an operator-splitting approach (a type of approximate factorization)¹⁶. Decomposing the Jacobian into a sum of fluid flow and Maxwell’s equation components by applying the operator-split approximation we have

$$\begin{bmatrix} \mathbf{F} & \mathbf{B}^T & \mathbf{Z} \\ \mathbf{B} & \mathbf{C} & \mathbf{0} \\ \mathbf{Y} & \mathbf{0} & \mathbf{D} \end{bmatrix} \approx \begin{bmatrix} \mathbf{F} & \mathbf{0} & \mathbf{Z} \\ \mathbf{0} & \mathbf{I} & \mathbf{0} \\ \mathbf{Y} & \mathbf{0} & \mathbf{D} \end{bmatrix} \begin{bmatrix} \mathbf{F}^{-1} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{I} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{I} \end{bmatrix} \begin{bmatrix} \mathbf{F} & \mathbf{B}^T & \mathbf{0} \\ \mathbf{B} & \mathbf{C} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{I} \end{bmatrix} = \begin{bmatrix} \mathbf{F} & \mathbf{B}^T & \mathbf{Z} \\ \mathbf{B} & \mathbf{C} & \mathbf{0} \\ \mathbf{Y} & \boxed{\mathbf{Y}\mathbf{F}^{-1}\mathbf{B}^T} & \mathbf{D} \end{bmatrix}. \quad (9)$$

The composite matrix on the right is the effective numerical approximation generated by the operator-split approximate factorization. It should be noted that this approximation, when used as a preconditioner, requires the solution to both the magnetics/velocity coupling and the Navier-Stokes system. Because of the relatively simple 2×2 structure of these systems the resulting Schur complements are more easily approximated. The magnitude and the behavior of the resulting splitting error is currently under analysis and numerical investigation.

4.3.1 Inverting the Fluid and Magnetic/Velocity Systems

The preconditioner defined by Eq. (9) requires an approximate inverse of the 2×2 systems

$$\begin{bmatrix} \mathbf{F} & \mathbf{B}^T \\ \mathbf{B} & \mathbf{C} \end{bmatrix} \quad \text{and} \quad \begin{bmatrix} \mathbf{F} & \mathbf{Z} \\ \mathbf{Y} & \mathbf{D} \end{bmatrix}. \quad (10)$$

Preconditioning for the Navier-Stokes equations (the matrix on the left) has been an active area of study. The approach we have taken is to use the Pressure Convection-Diffusion (PCD) preconditioner¹⁵. This uses a block factorization of the Navier-Stokes system and approximates the inverse of the fluid Schur complement (see Eq. (7)) using a commuting argument.

For the magnetics/velocity coupling we employ a similar 2×2 block LU factorization. A preconditioner utilizing a block factorization based on the ordering in Eq. (10) requires that the inverse of the Schur complement

$$\mathbf{P} = \mathbf{D} - \mathbf{Y}\mathbf{F}^{-1}\mathbf{Z} \quad (11)$$

is approximated. Similar to⁶ (with the exception that they use a \mathbf{B} field formulation) we assume the following continuous commuting condition

$$\nabla A_z \cdot \left(\frac{\partial}{\partial t} + \vec{w} \cdot \nabla - \nu \nabla^2 \right) \approx \left(\frac{\partial}{\partial t} + \vec{w} \cdot \nabla - \nu \nabla^2 \right) \nabla A_z \cdot. \quad (12)$$

This assumption motivates the discrete commuting condition $\mathbf{Y}\mathbf{Q}_u^{-1}\mathbf{F} \approx \mathbf{D}\mathbf{Q}_a^{-1}\mathbf{Y}$, where \mathbf{Q}_a and \mathbf{Q}_u are the vector potential and velocity mass matrices respectively. Rearranging the discrete condition and substituting into Eq. (11) gives the approximate Schur

complement

$$\begin{aligned} \mathbf{P} &\approx \hat{\mathbf{P}} = \mathbf{D} - \mathbf{Q}_a \mathbf{D}^{-1} \mathbf{Y} \mathbf{Q}_u^{-1} \mathbf{Z} \\ &= \mathbf{Q}_a \mathbf{D}^{-1} (\mathbf{D} \mathbf{Q}_a^{-1} \mathbf{D} - \mathbf{Y} \mathbf{Q}_u^{-1} \mathbf{Z}). \end{aligned} \tag{13}$$

The inverse of $\hat{\mathbf{P}}$ requires that the operator

$$\mathbf{D} \mathbf{Q}_a^{-1} \mathbf{D} - \mathbf{Y} \mathbf{Q}_u^{-1} \mathbf{Z}$$

is explicitly formed (using lumped approximations of the inverse mass matrices) and “inverted” using one multigrid V-cycle.

4.4 Solution method software used in this study

The multilevel implementation described above is provided by ML^{47, 17}. The approximate block factorization methods are implemented in the Teko package that is available through the Trilinos framework²¹.

5 REPRESENTATIVE RESULTS

5.1 Scaling of DD and fully-coupled AMG preconditioners

The first example 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 -direction⁴⁰. 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 in the range of $x \in [2.5, 7.5]$ while outside of this range it 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^{22, 50}, and then the flow exits with a parabolic profile. The simple geometry of this problem facilitates scalability studies as different mesh sizes can be easily generated. A sample of the filled contours of A_z and B_y are presented in Figure 1.




Figure 1: Contour plots of A_z (left) and B_y (right) for the idealized Faraday conduction MHD pump.

The second example is a transient island coalescence problem that consists of a perturbed Harris sheet magnetic field configuration^{3, 27, 40} that introduces two magnetic islands in the plasma as initial conditions for the island coalescence problem. The structure




Figure 2: Contour plots for the island coalescence driven magnetic reconnection at times $t = 0.0$ on the left and 10.0 on the right. The images show isolines of the magnetic potential A_z and filled contours of the current J_z . The formation of the x-point and the thin current sheet is clearly evident.

of this perturbation can be seen in the initial condition plot at time $t = 0$ of Figure 2 (left) with iso-lines of A_z . The combined magnetic field produced by the two magnetic islands produces Lorentz forces that pull the islands together. 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.^{3, 32}). The right image in Figure 2 shows an iso-line plot of A_z and filled color contours of the the plasma current J_z during the reconnection event at time, $t = 10.0$. Clearly evident is the formation of the x-point between the islands, 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^{3, 27}.

5.1.1 Weak scaling of steady-state solvers: MHD Faraday conduction 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. In this study we consider the weak scaling of the one-level DD ILU preconditioner for various levels of fill-in, and the multilevel NSA preconditioner. For 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 Reynolds number, $Re = \rho UL/\mu$, and magnetic Reynolds number, $Re_m = \mu_0 UL/\eta$ flow ($Re = Re_m = 0.7$) with a Hartmann number $Ha = B_0 L/\sqrt{\mu\eta} = 1$. In this study, we have taken $L = 2$, $\rho = 1$, $\mu/\rho = \eta = 1$, $\mu_0 = 1$. E_0^z and B_0 are then selected to produce the desired max velocity U to set the $Re = Re_m$ and Ha . The Krylov method is a non-restarted GMRES technique to allow only the parallel scalability of the preconditioners to be addressed. For the one-level DD preconditioner, an incomplete factorization ILU(k) sub-domain solver was used with $k = 1, 3, 7$ with an overlap of 2. For the 3-level preconditioner, the fine and medium meshes use an ILU(1) smoother with 2 levels of overlap and the coarsest problem was solved by the KLU sparse direct solver.

Figure 3 graphically presents the parallel and algorithmic scaling of the one- and three-level preconditioners for the MHD Faraday pump. Figure 3 (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, that is 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 loss of CPU time scaling is non-optimal, it must be pointed out that the 3-level method is still significantly faster (a factor of about 10–20x) than the corresponding one-level methods. To mitigate this growth of CPU time for the coarse grid solve, either approximate coarse grid methods can be used (e.g.⁴¹) or more levels could be employed.

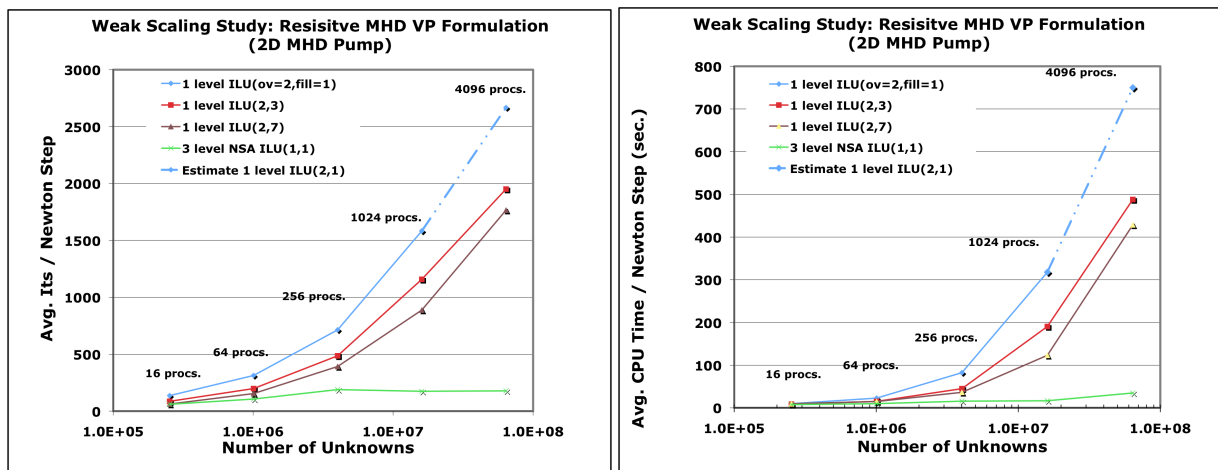


Figure 3: Weak scaling results for MHD Faraday pump problem. The scaling of the average number of iterations per Newton step (left) and the average time per Newton step (right). This time per Newton step is the sum of the linear solve time and time to construct the Jacobian (about 4.9 seconds). These results were obtained on a Cray XT3/4 system.

5.1.2 Weak scaling of transient solvers: Driven magnetic reconnection.

A preliminary study of the weak scaling for the graph-based aggressive coarsening multilevel preconditioner for the fully-implicit transient solution of the island coalescence problem at a resistivity of $\eta = 10^{-3}$ is presented. 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 in²⁷ these choices imply that the resistivity $\eta = 1/S$, where S is the Lundquist number that is defined as $S = \mu_0 L V_A / \eta$, with $V_A = B_0 / \sqrt{\mu_0 \rho}$ the Alfvén velocity. In this case the Lundquist number, $S = 10^3$. The test consists of time steps of size $\Delta t = 0.1$ time units

with an integration carried out to $t = 10.0$, the results are averaged over these 100 time steps. The maximum stiff wave $CFL = 51.2$ in this case, based on the Alfvén wave speed.

Table 3 presents the results for the 2-level PGSA method with an aggregation size of 80. The columns in the table are: average Newton steps per time step, average GMRES iterations per Newton step, average linear system solve time per Newton step in seconds, average GMRES iterations per time step, and average linear solve time per time step (including preconditioner setup time) in seconds.

The multilevel method is seen to limit the growth in the linear iteration count as the problem size is increased. In this case for an increase in problem size by a factor of 64 the multilevel preconditioner has an increase in iteration count of about a factor of 2. While the scaling 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.

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 3: Weak scaling of the block AMG preconditioned Newton-Krylov solver for the stabilized resistive MHD formulation on the island coalescence problem. The smoother in the PGSA method is an ILU(1) and it is a V(1,1) cycle and the aggregation size is 80. In this study the linear solver convergence criteria is 10^{-3} for each step of the Newton solver.

Next we consider an initial study of the influence of the Lundquist number on the iteration count of the linear solver with the fully-coupled AMG preconditioner on a 512×512 mesh. The results for this study are presented in Table 4. From Table 4 it is clear that as the dissipation in the system is decreased the convergence rate of the underlying fully-coupled AMG preconditioner decreases and more iterations are required. In the range of Lundquist number from 10^2 to 10^4 this increase was by a factor of $\approx 3x$. This behavior might be expected since our current fully-coupled AMG preconditioning approach is tailored to parabolic and elliptic PDEs. Therefore optimal performance in the higher Lundquist number regimes, on moderately refined meshes, would be surprising. Nevertheless, the solver has demonstrated to be reasonably robust, even in this regime as is evidenced by our ability to pursue computations up to $S = 10^5$ on appropriately refined mesh⁴⁰. We believe that improving on these results will require the implementation of physics-based preconditioning ideas, as proposed in⁵. We are currently pursuing a more extensive study of the higher Lundquist number behavior of these preconditioners.

Lundquist No.	Newt/ Δt	iter/Newt	prec time/N	Advance time/N	iter/ Δt	Advance time/ Δt
10^2	5.5	3.6	3.2	7.5	19.8	41.1
10^3	4.7	6.8	3.3	7.5	32.0	35.4
10^4	4.3	10.3	3.8	8.1	44.2	34.5

Table 4: Study of the influence of the Lundquist number in the iteration count of the PGSA preconditioner for mesh of size 512×512 . The multilevel preconditioner is a 2-level method with $ILU(3)$ and 2 levels of overlap with an aggregation size of 40. Times are in seconds.

5.2 Initial Results for ABF preconditioners

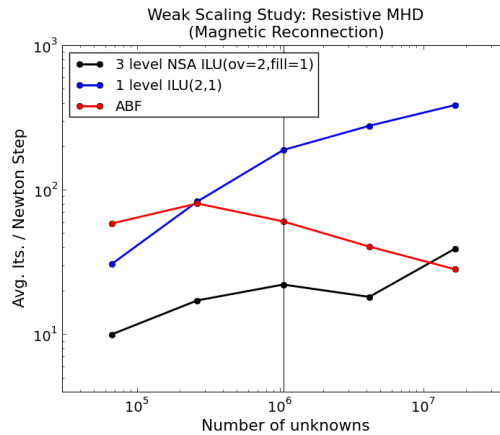


Figure 4: Constant CFL = 52.1 weak scaling for the magnetic reconnection problem on 1, 4, 16, 64 and 256 processors. The thin vertical line near 10^6 corresponds to $\Delta t = 0.1$

Finally we present a preliminary result comparing the performance of the abstract block factorization preconditioner presented in Eq. (9), to the aggressive coarsening and domain decomposition preconditioners. This study of the ABF preconditioner is a weak scaling study that uses a fixed mesh of size 128×128 on each processor, for the 256 processor case there are 16.8M unknowns. In addition for this study we select a fixed $CFL = 51.2$ that corresponds to the previous study above. The domain decomposition preconditioner uses an $ILU(1)$ sub-domain solver with an overlap of 2 and the fully-coupled ML preconditioner uses a 3 level NSA method with $ILU(1)$ smoother with 2 levels of overlap. The ABF preconditioner uses two different multigrid solvers to perform the subsolves required by the factorization. For subsolves involving velocity unknowns (in particular \mathbf{F}^{-1}) a NSA method with an $ILU(1)$ smoother with an overlap of 2 is used. Both the fluid and the magnetic/velocity Schur complements also require a subsolve. For these a multigrid method, using smoothed aggregation for projection and $ILU(1)$ with

overlap of 2 as the smoother, is applied.

Fig. 4 presents results from a study of the ABF preconditioner for a Lundquist number of $S = 10^4$. This preliminary study shows at least two trends. First for large time step sizes, that correspond to $\Delta t = 0.4$ and $\Delta t = 0.2$ for the 128×128 and 256×256 mesh case, the ML and DD preconditioners appear to be more effective. Based on our preliminary analysis of the factorization error incurred by the splitting approximation this behavior may be able to be explained by further analysis. In addition, it should be noted that these time step sizes are probably too large as the time-scale for the reconnection event essentially requires time steps that are on the order of 0.1 time units, or less, as used in the study above. The second clear trend in this preliminary data is a bit more interesting. That is that the iteration count for the ABF preconditioner is decreasing as a function of the mesh spacing (time step size as well since this is a constant CFL study). Clearly this behavior is in contrast to the DD and fully-coupled ML preconditioners. We are currently pursuing more extensive studies of the ABF preconditioner to more clearly understand this interesting behavior.

6 CONCLUSIONS

This brief paper has presented a preliminary performance study of an unstructured fully-implicit stabilized FE formulation for 2D incompressible reduced resistive MHD. The solution methods used in this formulation are based on a fully-coupled Newton-Krylov approach employing a fully-coupled AMG technique and a new ABF as preconditioners. As an illustration of the robustness, scalability, and efficiency of the solution techniques, performance results have been presented for a MHD duct flow, and a magnetic island coalescence problem. The preliminary efficiency and scalability results of this study are encouraging. 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. In addition, the first result for the ABF preconditioner appears encouraging although this method needs significantly more analysis and evaluation. Future work will focus on more extensive studies of these 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.

REFERENCES

- [1] ASCHER, U. M., AND PETZOLD, L. R. *Computer Methods for Ordinary Differential Equations and Differential-Algebraic Equations*. SIAM, 1998.
- [2] BENZI, M., GOLUB, G. H., AND LIESEN, J. Numerical solution of saddle point problems. *Acta Numerica 14* (2005), 1–137.

- [3] BISKAMP, D. *Magnetic Reconnection in Plasmas*. Cambridge University Press, Cambridge, UK, 2000.
- [4] BRENAN, K. E., CAMPBELL, S. L., AND PETZOLD, L. R. *Numerical Solution of Initial-Value Problems in Differential-Algebraic Equations*. SIAM, Classics in Applied Math, 1996.
- [5] CHACÓN, L. An optimal, parallel, fully implicit Newton-Krylov solver for three-dimensional visco-resistive magnetohydrodynamics. *Phys. Plasmas* 15 (2008), 056103.
- [6] CHACÓN, L. Scalable solvers for 3D magnetohydrodynamics. *J. Physics: Conf. Series* 125 (2008), 012041.
- [7] CHACÓN, L., KNOLL, D. A., AND FINN, J. M. Implicit, nonlinear reduced resistive MHD nonlinear solver. *J. Comput. Phys.* 178, 1 (2002), 15–36.
- [8] CHOW, E., AND SAAD, Y. Approximate inverse techniques for block-partitioned matrices. *SIAM J. Sci. Comput.* 18, 6 (1997), 1657–1675.
- [9] CODINA, R., AND HERNANDEZ-SILVA, N. Stabilized finite element approximation of the stationary magneto-hydrodynamics equations. *Computational Mechanics* 38, 4-5 (2006), 344–355.
- [10] DAVIDSON, P. A. *An Introduction to Magnetohydrodynamics*. Cambridge Univ. Press, 2001.
- [11] DAVIS, T. *Direct methods for sparse linear systems*. SIAM, Philadelphia, PA, 2006.
- [12] DENNIS, JR., J. E., AND SCHNABEL, R. B. *Numerical Methods for Unconstrained Optimization and Nonlinear Equations*. Series in Automatic Computation. Prentice-Hall, Englewood Cliffs, NJ, 1983.
- [13] DONEA, J., AND HUERTA, A. *Finite Element Methods for Flow Problems*. John Wiley, 2002.
- [14] DRAKE, J. F., AND ANTONSEN, T. M. Nonlinear reduced fluid equations for toroidal plasmas. *Phys. Fluids* 27, 4 (1984), 898–908.
- [15] ELMAN, H., HOWLE, V. E., SHADID, J. N., SHUTTLEWORTH, R., AND TUMINARO, R. S. A taxonomy of parallel multi-level block preconditioners for the incompressible navier-stokes equations. *JCP.* 227 (2008), 1790–1808.
- [16] FLETCHER, C. A. J. *Computational Techniques for Fluid Dynamics*, vol. Vol I and II of *Computational Physics*. Springer-Verlag, Berlin Heidelberg, 1988.

- [17] GEE, M., SIEFERT, C., HU, J., TUMINARO, R., AND SALA, M. ML 5.0 smoothed aggregation user's guide. Tech. Rep. SAND2006-2649, Sandia National Laboratories, Albuquerque, NM, 87185, 2006.
- [18] GOEDBLOED, H., AND POEDTS, S. *Principles of Magnetohydrodynamics with Applications to Laboratory and Astrophysical Plasmas*. Cambridge Univ. Press, 2004.
- [19] GUNZBURGER, M. *Finite Element Methods for Viscous Incompressible Flows*. Academic Press, Boston, 1989.
- [20] HAZELTINE, R. D., KOTSCHENREUTHER, M., AND MORRISON, P. J. A four-field model for tokamak plasma dynamics. *Phys. Fluids* 28, 8 (1985), 2466–2477.
- [21] HEROUX, M., BARTLETT, R., HOWLE, V., HOEKSTRA, R., HU, J., KOLDA, T., LEHOUCQ, R., LONG, K., PAWLOWSKI, R., PHIPPS, E., SALINGER, A., THORNQUIST, H., TUMINARO, R., WILLENBRING, J., AND WILLIAMS, A. An Overview of Trilinos Project. *ACM Trans. Math. Software* 31, 3 (2005), 397–423.
- [22] HUGHES, M., PERICLEOUS, K. A., AND CROSS, M. The numerical modelling of dc electromagnetic pump and brake flow. *Appl. Math. Modelling* 19 (1995), 713–723.
- [23] HUGHES, T., AND MALLET, M. A new finite element formulation for computational fluid dynamics: III. The generalized streamline operator for multidimensional advective-diffusive systems. *Comput. Meth. Appl. Mech. Engrg.* 58 (1986), 305–328.
- [24] HUJEIRAT, A. IRMHD: an implicit radiative and magnetohydrodynamical solver for self-gravitating systems. *Mon. Not. R. Astron. Soc.* 298 (1998), 310–320.
- [25] HUJEIRAT, A., AND RANNACHER, R. On the efficiency and robustness of implicit methods in computational astrophysics. *New Astronomy Reviews* 45 (2001), 425–447.
- [26] KARYPIS, G., AND KUMAR, V. Parallel multilevel k-way partitioning scheme for irregular graphs. *ACM/IEEE Proceedings of SC96: High Performance Networking and Computing* (1996).
- [27] KNOLL, D. A., AND CHACÓN, L. Coalescence of magnetic islands, sloshing, and the pressure problem. *Phys. Plasmas* 13, 3 (2006), 32307 – 1.
- [28] LIN, P. T., SALA, M., SHADID, J. N., AND TUMINARO, R. S. Performance of fully-coupled algebraic multilevel domain decomposition preconditioners for incompressible flow and transport. *Int. J. Num. Meth. Eng.* 67, 9 (2006), 208–225.
- [29] MOREAU, R. *Magnetohydrodynamics*. Kluwer, Dordrecht, 1990.
- [30] MURPHY, M. F., GOLUB, G. H., AND WATHEN, A. J. A note on preconditioning for indefinite linear systems. *SIAM J. Sci. Comput.* 21 (2000), 1969–1972.

- [31] OVTCHINNIKOV, S., DOBRIAN, F., CAI, X.-C., AND KEYES, D. Additive Schwarz-based fully coupled implicit methods for resistive Hall magnetohydrodynamic problems. *J. Comput. Phys.* 225 (2007), 1919 – 1936.
- [32] PRIEST, E., AND FORBES, T. *Magnetic Reconnection: MHD Theory and Applications*. Cambridge Univ. Press, 2006.
- [33] QUARTERONI, A., AND VALLI, A. *Domain Decomposition Methods for Partial Differential Equations*. Oxford University Press, Oxford, 1999.
- [34] REYNOLDS, D. R., SAMTANEY, R., AND WOODWARD, C. S. A fully implicit numerical method for single-fluid resistive magnetohydrodynamics. *J. Comput. Phys.* 219, 1 (2006), 144 – 62.
- [35] RUGE, J., AND STÜBEN, K. Algebraic multigrid (AMG). In *Multigrid Methods*, S. F. McCormick, Ed., vol. 3 of *Frontiers in Applied Mathematics*. SIAM, Philadelphia, PA, 1987, pp. 73–130.
- [36] SAAD, Y. *Iterative Methods for Sparse Linear Systems*. SIAM, 2003.
- [37] SALA, M., HU, J., AND TUMINARO, R. ML 3.1 smoothed aggregation user’s guide. Tech. Rep. SAND2004-4819, Sandia National Laboratories, September 2004.
- [38] SALA, M., SHADID, J. N., AND TUMINARO, R. S. An improved convergence bound for aggregation-based domain decomposition preconditioners. *SIAM J. Matrix Analysis* 27, 3 (2006), 744–756.
- [39] SALA, M., AND TUMINARO, R. A new Petrov-Galerkin smoothed aggregation preconditioner for nonsymmetric linear systems. *SIAM J. Sci. Comput.* 31, 1 (2008), 143–166.
- [40] SHADID, J., PAWLOWSKI, R., BANKS, J. W., CHACON, L., LIN, P., AND TUMINARO, R. Towards a scalable fully-implicit fully-coupled resistive MHD formulation with stabilized FE methods. *Submitted to Journal of Computational Physics*.
- [41] SHADID, J., TUMINARO, R., DEVINE, K., HENNINGAN, G., AND LIN, P. Performance of fully-coupled domain decomposition preconditioners for finite element transport/reaction simulations. *J. Comput. Phys.* 205, 1 (2005), 24–47.
- [42] SHADID, J. N., SALINGER, A. G., PAWLOWSKI, R. P., LIN, P. T., HENNINGAN, G. L., TUMINARO, R. S., AND LEHOUCQ, R. B. Stabilized FE computational analysis of nonlinear steady state transport/reaction systems. *Comp. Meth. Applied Mech. Eng.* 195 (2006), 1846–1871.

- [43] SHAKIB, F. *Finite element analysis of the compressible Euler and Navier-Stokes equations*. PhD thesis, Division of Applied Mathematics, Stanford University, 1989.
- [44] STRAUSS, H. R. Nonlinear, 3-dimensional magnetohydrodynamics of noncircular tokamaks. *Phys. Fluids* 19, 1 (1976), 134–140.
- [45] TÓTH, G., KEPPENS, R., AND BOTCHEV, M. A. Implicit and semi-implicit schemes in the Versatile Advection Code: numerical tests. *Astron. Astrophys.* 332 (1998), 1159–1170.
- [46] TROTTEBERG, U., OOSTERLEE, C., AND SCHÜLLER, A. *Multigrid*. Academic Press, London, 2001.
- [47] TUMINARO, R., TONG, C., SHADID, J., K.D.DEVINE, AND DAY, D. On a multilevel preconditioning module for unstructured mesh Krylov solvers: two-level Schwarz. *Comm. Num. Method. Eng.* 18 (2002), 383–389.
- [48] VANĚK, P., BREZINA, M., AND MANDEL, J. Convergence of algebraic multigrid based on smoothed aggregation. *Numerische Mathematik* 88 (2001), 559–579.
- [49] VANĚK, P., MANDEL, J., AND BREZINA, M. Algebraic multigrid by smoothed aggregation for second and fourth order elliptic problems. *Computing* 56 (1996), 179–196.
- [50] VERARDI, S. L. L., CARDOSO, J. R., AND COSTA, M. C. Three-dimensional finite element analysis of MHD duct flow by the penalty function formulation. *IEEE TRANSACTIONS ON MAGNETICS* 37, 5 (2001), 3384–3387.