

Additive Schwarz-based Fully Coupled Implicit Methods for Resistive Hall Magnetohydrodynamic Problems [★]

S. Ovtchinnikov ^a, F. Dobrian ^{b,c},
X.-C. Cai ^{a,1}, and D. E. Keyes ^c

^a*Department of Computer Science, University of Colorado, Boulder, CO 80309*

^b*Department of Computer Science, Old Dominion University, Norfolk, VA 23529*

^c*Department of Applied Physics & Applied Mathematics, Columbia University,
New York, NY 10027*

Abstract

A parallel, fully coupled, nonlinearly implicit Newton-Krylov-Schwarz algorithm is proposed for the numerical simulation of a magnetic reconnection problem described by a system of resistive Hall magnetohydrodynamics equations in slab symmetry. A key component of the algorithm is a restricted additive Schwarz preconditioner defined for problems with doubly periodic boundary conditions. We show numerically that with such a preconditioned nonlinearly implicit method the time step size is no longer constrained by the CFL number or the convergence of the Newton solver. We report the parallel performance of the algorithm and software on machines with thousands of processors.

Key words: restricted additive Schwarz preconditioner, Newton-Krylov, fully implicit, parallel computing, magnetohydrodynamics

[★] The research was supported in part by the Department of Energy under cooperative agreements DE-FC02-01ER25479 and DE-FC02-04ER25595, and in part by the National Science Foundation under grants CCR-0219190, CNS-0420873, ACI-0305666, and CCF-03-52334.

¹ Corresponding author: cai@cs.colorado.edu

1 Introduction

Magnetic reconnection is an important process in magnetically confined plasmas and has been observed both in space plasmas and in laboratory experiments [2,13,24]. Experimental investigations of the magnetic reconnection are often difficult and numerical simulations are becoming increasingly useful for reconnection studies. Mathematically, the reconnection can be described by a system of magnetohydrodynamics equations (MHD) [17], where plasma is treated as a conducting fluid satisfying the Navier-Stokes equations coupled with the Maxwell's equations. Solving the MHD equations numerically is a challenge because of the complex, not yet fully understood behavior of the solution. The system admits phenomena such as Alfvén waves and their instabilities, and one of the intrinsic features of the system is the formation of a singular current density sheet [29], which is linked to the reconnection of magnetic field lines.

During the process of magnetic reconnection, the magnetic field configuration undergoes a topological rearrangement that results in conversion of magnetic energy into kinetic flow energy and heat over a relatively short period of time. In the strictly ideal MHD limit, magnetic field lines are “frozen” into the plasma and magnetic reconnection cannot occur. Capturing the change of the magnetic field topology requires a more general model than the ideal MHD [25]. In this paper we focus on a resistive Hall MHD model for the magnetic reconnection.

In order to simulate this multi-scale, multi-physics phenomenon, a robust solver is needed to deal with the high degree of nonlinearity and the non-smooth, nearly singular behavior in the system. One of the popular approaches to the numerical solution of the MHD system is based on the splitting of the system into two parts, where equations for the current and the vorticity are advanced in time, and the corresponding potentials are obtained by solving Poisson-like equations in a separate step. In such an explicit approach, to satisfy the Courant-Friedrichs-Lewy (CFL) condition, the time step may become very small, especially in the case of fine meshes.

Recently several implicit methods have been introduced in order to overcome the time step size issue. For example in [9,10], implicit nonlinear solvers are proposed as solution methods for systems of reduced and reduced Hall MHD equations in two-dimensional space. The solver is a matrix-free Newton-Krylov method with a physics-based preconditioner. Excellent results with large time steps are reported based on single processor calculations. In [26], a fully parallel, conservative, nonlinearly implicit numerical method is proposed for the integration of the single-fluid resistive MHD system of equations, where a variant of the matrix-free Newton-Krylov method, without preconditioning,

is used in conjunction with an adaptive time integration scheme and variable spatial discretization accuracy. In this approach, the time step is not restricted by the CFL, but is restricted to some extent by the nonlinear solver as the time step needs to be cut when the unpreconditioned GMRES is unable to solve the Jacobian system.

In our paper, we develop a fully parallel nonlinearly implicit method based on a high-order time integration scheme, an inexact Newton’s method with an explicitly calculated Jacobian matrix, and a Krylov subspace linear solver with a restricted additive Schwarz preconditioner defined on overlapping subdomains. We show numerically that this approach allows large time steps that are not restricted by the CFL or the linear/nonlinear solvers and works well on machines with thousands of processors. The MHD system that we consider has four equations: two linear elliptic and two nonlinear and time dependent. We take a “fully coupled” approach such that no operator splitting is applied to the system of MHD equations. Note that most of existing approaches are based on some form of operator splitting which considers the unknowns of a single physical variable across all mesh points as a subproblem.

We first apply a second- or third-order implicit time integration scheme, and then, to guarantee the nonlinear consistency, we use a Newton-Krylov-Schwarz algorithm to solve the large sparse nonlinear system of algebraic equations containing all physical variables at every time step. In a Newton-Krylov-Schwarz algorithm, a system is solved by applying outer Newton iterations, whose Jacobian systems are solved with a preconditioned Krylov subspace method, where the preconditioning is accomplished via the parallel Schwarz technique. An explicit algorithm is used to compare the results obtained with the implicit approach. Both implementations are based on the PETSc (Portable Extensible Toolkit for Scientific computation) [1] library thus providing a convenient test bed for investigations of parallel properties of the algorithms. We focus on scalability studies on fine meshes and on machines with thousands of processors.

The remainder of this paper is organized as follows. In Section 2, we discuss the model MHD problem and provide some useful definitions. Spatial and temporal discretizations are covered in Section 3. The details of algorithms are described in Sections 4 and 5, and numerical results are reported in Section 6. The paper is concluded in Section 7.

2 Model MHD Problem

The system of equations we model can be derived starting from the momentum transfer equations. Following [3,15], we can write

$$nm_e \left(\frac{\partial \mathbf{V}_e}{\partial t} + (\mathbf{V}_e \cdot \nabla) \mathbf{V}_e \right) = -\nabla p_e - ne(\mathbf{E} + \mathbf{V}_e \times \mathbf{B}) + \nu_e \nabla^2 \mathbf{V}_e + nen\mathbf{j}, \quad (2.1)$$

for the species of electrons and

$$nm_i \left(\frac{\partial \mathbf{V}_i}{\partial t} + (\mathbf{V}_i \cdot \nabla) \mathbf{V}_i \right) = -\nabla p_i + ne(\mathbf{E} + \mathbf{V}_i \times \mathbf{B}) + \nu_i \nabla^2 \mathbf{V}_i - nen\mathbf{j}, \quad (2.2)$$

for the species of ions. In (2.1) and (2.2) the plasma is considered to be quasi neutral, ions are singly charged, ion/electron number density is n , the resistivity is η , the ion/electron viscosity is given by $\nu_{i,e}$, \mathbf{E} is the electric field strength, \mathbf{B} is the magnetic induction, \mathbf{j} is the current density, $m_{i,e}$ is the ion/electron mass, $\mathbf{V}_{i,e}$ is the ion/electron velocity, and $p_{i,e}$ is the ion/electron pressure. Additionally, we can introduce

$$ne(\mathbf{V}_i - \mathbf{V}_e) \equiv \mathbf{j}. \quad (2.3)$$

Maxwell's equations enter the picture via the following three equations:

$$\nabla \cdot \mathbf{B} = 0, \quad (2.4)$$

$$\nabla \times \mathbf{B} = \mu_0 \mathbf{j}, \quad (2.5)$$

$$\nabla \times \mathbf{E} = -\frac{\partial \mathbf{B}}{\partial t}. \quad (2.6)$$

(2.1) – (2.6) provide a full description of the plasma, given certain assumptions on $p_{i,e}$. The incompressibility condition

$$\nabla \cdot \mathbf{V}_{i,e} = 0 \quad (2.7)$$

is added if the plasma is considered incompressible.

The above equations can be viewed as the general description of plasma in the magnetohydrodynamics formalism. Special workable systems of equations may be obtained by employing reduction procedures, where certain relative magnitude ordering schemes are introduced for such physical quantities as pressure and the magnetic field induction. For example in [14], where a magnetized, two-species (electron and ion), quasi neutral plasma with singly charged ions of mass m_i and a constant, uniform number density n_0 is considered, the reduction procedure is used to obtain systems of equations corresponding to “high-”, “low-” and “zero-” β approximations with dimensionless β denoting the conventional plasma parameter [14]. In the derivation a normalization scheme is adopted [15], such that all lengths are measured in terms of some

scale a , all magnetic fields are measured in terms of some scale B_0 , all velocities in terms of characteristic Alfvén speed $V_A = B_0/\sqrt{\mu_0 n_0 m_i}$, and all pressures in terms of B_0^2/μ_0 . By introducing the stream functions ϕ and ψ , one can write $\mathbf{B} = \nabla\psi \times \hat{\mathbf{z}} + B_z \hat{\mathbf{z}}$ and $\mathbf{V}_i = \nabla\phi \times \hat{\mathbf{z}} + V_z \hat{\mathbf{z}}$, where \mathbf{B} is the magnetic induction, \mathbf{V}_i is the ion velocity, B_z and V_z are the magnetic field and ion velocity in z direction, respectively. Furthermore, other quantities such as vorticity and the current density are expressed as $U = \nabla^2\phi$ and $j = -\nabla^2\psi$ using the stream functions. In this paper we work with a model MHD problem described in [4,18], which also corresponds to the “zero- β ” case in [14]:

$$\left\{ \begin{array}{l} \nabla^2\phi = U \\ \nabla^2\psi = \frac{1}{d_e^2}(\psi - F) \\ \frac{\partial U}{\partial t} + [\phi, U] = \frac{1}{d_e^2}[F, \psi] + \nu\nabla^2 U \\ \frac{\partial F}{\partial t} + [\phi, F] = \rho_s^2[U, \psi] + \eta\nabla^2(\psi - \psi^0), \end{array} \right. \quad (2.8)$$

where U is the vorticity, F is the canonical momentum, ϕ and ψ are the stream functions for the vorticity and current density, respectively, ν is the plasma viscosity, η is the normalized resistivity, $d_e = c/\omega_{pe}$ is the inertial skin depth, $\rho_s = \sqrt{T_e/T_i}\rho_i$ is the ion sound Larmor radius, and ψ^0 is the equilibrium component of ψ that is defined below. The current density is obtained via $j = (F - \psi)/d_e^2$. The Poisson bracket is defined as: $[A, B] \equiv (\partial A/\partial x)(\partial B/\partial y) - (\partial A/\partial y)(\partial B/\partial x)$.

For the case we wish to study every variable in the system is assumed to be the sum of an equilibrium and a perturbation component; i.e. $\phi = \phi^0 + \phi^1$, $\psi = \psi^0 + \psi^1$, $U = U^0 + U^1$, and $F = F^0 + F^1$, where $\phi^0 = U^0 = 0$, $\psi^0 = \cos(x)$, and $F^0 = (1 + d_e^2)\cos(x)$ are the equilibrium components [24]. After substitutions, we arrive at the following system for the perturbed variables:

$$\left\{ \begin{array}{l} \nabla^2\phi^1 = U^1 \\ \nabla^2\psi^1 = \frac{1}{d_e^2}(\psi^1 - F^1) \\ \frac{\partial U^1}{\partial t} + [\phi^1, U^1] = \frac{1}{d_e^2}[F^1, \psi^1] + \nu\nabla^2 U^1 + \frac{1}{d_e^2} \left(\frac{\partial\psi^1}{\partial y} F_{eqx} + \frac{\partial F^1}{\partial y} B_{eqy} \right) \\ \frac{\partial F^1}{\partial t} + [\phi^1, F^1] = \rho_s^2[U^1, \psi^1] + \eta\nabla^2\psi^1 + \left(\frac{\partial\phi^1}{\partial y} F_{eqx} + \rho_s^2 \frac{\partial U^1}{\partial y} B_{eqy} \right), \end{array} \right. \quad (2.9)$$

where $F_{eqx} = -(1 + d_e^2)\sin(x)$ and $B_{eqy} = \sin(x)$. The system is defined on a rectangular domain $\Omega \equiv [l_x, l_y] \equiv [2\pi, 4\pi]$, and doubly periodic boundary conditions are assumed. For initial conditions, we use a nonzero initial perturbation in ϕ^1 and a zero initial perturbation in ψ^1 . The exact form of the perturbation follows after some useful definitions. The aspect ratio is $\epsilon = l_x/l_y$. The

perturbation's magnitude is scaled by $\xi = 10^{-4}$. We define $\tilde{d}_e = \max\{d_e, \rho_s\}$ and $\gamma = \epsilon \tilde{d}_e$. For the initial value of the ϕ perturbation we use

$$\phi^1(x, y, 0) = \begin{cases} \xi \frac{\gamma}{\epsilon} \operatorname{erf}\left(\frac{x}{\sqrt{2}\tilde{d}_e}\right) \sin(\epsilon y) & \text{if } 0 \leq x < \frac{\pi}{2} \\ -\xi \frac{\gamma}{\epsilon} \operatorname{erf}\left(\frac{x - \pi}{\sqrt{2}\tilde{d}_e}\right) \sin(\epsilon y) & \text{if } \frac{\pi}{2} \leq x < \frac{3\pi}{2} \\ \xi \frac{\gamma}{\epsilon} \operatorname{erf}\left(\frac{x - 2\pi}{\sqrt{2}\tilde{d}_e}\right) \sin(\epsilon y) & \text{if } \frac{3\pi}{2} \leq x \leq 2\pi. \end{cases} \quad (2.10)$$

Other field quantities are set as: $U^1(x, y, 0) = \nabla^2 \phi^1(x, y, 0)$ and $F^1(x, y, 0) = \psi^1(x, y, 0) - d_e \nabla^2 \psi^1(x, y, 0)$. From now on, we drop the superscript and assume that the four fields ϕ , ψ , U and F represent the perturbed components only. In the (2.9) two of the equations are time independent and linear; another two equations are time dependent and nonlinear. Solving (2.9) numerically is the focus of this paper.

3 Discretizations

In this section we describe some nonlinearly implicit discretization schemes for solving (2.9). We cover the computational domain with a uniform $\{M_x, M_y\}$ mesh, and the corresponding mesh sizes are $h_x = 2\pi/M_x$ and $h_y = 4\pi/M_y$, respectively. Standard second-order central finite difference methods are used for all the spatial derivatives. For the time derivatives, we use a fixed time step Δt throughout the computation and we implement two time integration schemes of order 2 and 3. At time level t^k , we denote the grid values of the unknown functions $\phi(x, y, t)$, $\psi(x, y, t)$, $U(x, y, t)$, and $F(x, y, t)$, as $\phi_{i,j}^k$, $\psi_{i,j}^k$, $U_{i,j}^k$, and $F_{i,j}^k$.

For each grid point, the equilibrium values are given as:

$$F_{eqx}(i, j) = -(1 + d_e^2) \sin(h_x i), \quad (3.1)$$

$$B_{yeq}(i, j) = \sin(h_x i). \quad (3.2)$$

Equation (3.1) corresponds to the F 's x -equilibrium field component and equation (3.2) corresponds to B 's y -equilibrium component. At time level k , the stream function components are expressed as (ignoring the k index):

$$R_\phi^k(i, j) \equiv \frac{\phi_{i+1,j} - 2\phi_{i,j} + \phi_{i-1,j}}{h_x^2} + \frac{\phi_{i,j+1} - 2\phi_{i,j} + \phi_{i,j-1}}{h_y^2} - U_{i,j} = 0 \quad (3.3)$$

and

$$R_\psi^k(i, j) \equiv \psi_{i,j} - d_e^2 \left(\frac{\psi_{i+1,j} - 2\psi_{i,j} + \psi_{i-1,j}}{h_x^2} + \frac{\psi_{i,j+1} - 2\psi_{i,j} + \psi_{i,j-1}}{h_y^2} \right) - F_{i,j} = 0. \quad (3.4)$$

The spatial components of the other two equations are discretized as

$$\begin{aligned}
R_U^k(i, j) \equiv & \left[\left(v_{xp} \frac{U_{i,j} - U_{i-1,j}}{h_x} + v_{xm} \frac{U_{i+1,j} - U_{i,j}}{h_x} + v_{yp} \frac{U_{i,j} - U_{i,j-1}}{h_y} + v_{ym} \frac{U_{i,j+1} - U_{i,j}}{h_y} \right) \right. \\
& - \left(B_{xp} \left(\frac{F_{i,j} - F_{i-1,j}}{h_x} + F_{eqx} \right) + B_{xm} \left(\frac{F_{i+1,j} - F_{i,j}}{h_x} + F_{eqx} \right) \right. \\
& + \left. B_{yp} \left(\frac{F_{i,j} - F_{i,j-1}}{h_y} \right) + \left. B_{ym} \left(\frac{F_{i,j+1} - F_{i,j}}{h_y} \right) \right) / d_e^2 \\
& \left. - \nu \left(\frac{U_{i+1,j} - 2U_{i,j} + U_{i-1,j}}{h_x^2} + \frac{U_{i,j+1} - 2U_{i,j} + U_{i,j-1}}{h_y^2} \right) \right] \quad (3.5)
\end{aligned}$$

and

$$\begin{aligned}
R_F^k(i, j) \equiv & \left[\left(v_{xp} \left(\frac{F_{i,j} - F_{i-1,j}}{h_x} + F_{eqx} \right) + v_{xm} \left(\frac{F_{i+1,j} - F_{i,j}}{h_x} + F_{eqx} \right) \right. \right. \\
& \left. \left. + v_{yp} \frac{F_{i,j} - F_{i,j-1}}{h_y} + v_{ym} \frac{F_{i,j+1} - F_{i,j}}{h_y} \right) \right. \\
& - \left(B_{xp} \left(\frac{U_{i,j} - U_{i-1,j}}{h_x} \right) + B_{xm} \left(\frac{U_{i+1,j} - U_{i,j}}{h_x} \right) \right. \\
& \left. + B_{yp} \left(\frac{U_{i,j} - U_{i,j-1}}{h_y} \right) + B_{ym} \left(\frac{U_{i,j+1} - U_{i,j}}{h_y} \right) \right) \rho_s^2 \\
& \left. - \eta \left(\frac{\psi_{i+1,j} - 2\psi_{i,j} + \psi_{i-1,j}}{h_x^2} + \frac{\psi_{i,j+1} - 2\psi_{i,j} + \psi_{i,j-1}}{h_y^2} \right) \right]. \quad (3.6)
\end{aligned}$$

We have used the following notations in the above equations:

$$v_x = -\frac{\phi_{i,j+1} - \phi_{i,j-1}}{2h_y}, \quad (3.7)$$

$$v_y = \frac{\phi_{i+1,j} - \phi_{i-1,j}}{2h_x}, \quad (3.8)$$

$$v_{xp} = v_{xm} = \frac{v_x}{2}, v_{yp} = v_{ym} = \frac{v_y}{2}, \quad (3.9)$$

and

$$B_x = \frac{\psi_{i,j+1} - \psi_{i,j-1}}{2h_y}, \quad (3.10)$$

$$B_y = -\frac{\psi_{i+1,j} - \psi_{i-1,j}}{2h_x} + B_{yeq}(i, j), \quad (3.11)$$

$$B_{xp} = B_{xm} = \frac{B_x}{2}, \quad (3.12)$$

$$B_{yp} = B_{ym} = \frac{B_y}{2}. \quad (3.13)$$

We note that above spatial discretization is second-order accurate. Two of the four equations in (2.9) are time dependent. As suggested in [21], temporal discretization of higher order often provides for a better solution accuracy in numerical simulations of fluid dynamics problems, despite the fact that a lower spatial discretization order is used. For instance, the temporal discretization may be of order 3 or 4, while the order of the spatial discretization is 2. In this paper we employ backward differentiation formulas (BDF) [20] for the implicit methods and Adams formula [20] for the explicit version of our algorithm. Using multistep formulas requires the solutions from previous time steps. For example, the third-order BDF for system (2.9) with the help of (3.3) – (3.6) results in:

$$\left\{ \begin{array}{l} G_\phi^{k+1}(i, j) \equiv R_\phi^{k+1}(i, j) = 0 \\ G_\psi^{k+1}(i, j) \equiv R_\psi^{k+1}(i, j) = 0 \\ G_U^{k+1}(i, j) \equiv \frac{1}{6\Delta t} \left(11U_{i,j}^{k+1} - 18U_{i,j}^k + 9U_{i,j}^{k-1} - 2U_{i,j}^{k-2} \right) - R_U^{k+1}(i, j) = 0 \\ G_F^{k+1}(i, j) \equiv \frac{1}{6\Delta t} \left(11F_{i,j}^{k+1} - 18F_{i,j}^k + 9F_{i,j}^{k-1} - 2F_{i,j}^{k-2} \right) - R_F^{k+1}(i, j) = 0, \end{array} \right. \quad (3.14)$$

where (U^k, F^k) , (U^{k-1}, F^{k-1}) , (U^{k-2}, F^{k-2}) are solutions for (U, F) obtained at k , $k-1$ and $k-2$ time steps, respectively. Therefore, a high-order BDF requires sufficient solution history to be accumulated at the beginning of the time integration process. In our approach, the time integration process begins, starting from an initial guess at $t = 0$, with a BDF of order one (backward Euler), gradually increasing the BDF order up to a desired value as more and more solution history becomes available. For comparison purposes, we also implement an explicit scheme:

$$\left\{ \begin{array}{l} R_\phi^{k+1}(i, j) = 0 \\ R_\psi^{k+1}(i, j) = 0 \\ U_{i,j}^{k+1} - U_{i,j}^k - \Delta t \left(\frac{3}{2}R_U^k(i, j) - \frac{1}{2}R_U^{k-1}(i, j) \right) = 0 \\ F_{i,j}^{k+1} - F_{i,j}^k - \Delta t \left(\frac{3}{2}R_F^k(i, j) - \frac{1}{2}R_F^{k-1}(i, j) \right) = 0. \end{array} \right. \quad (3.15)$$

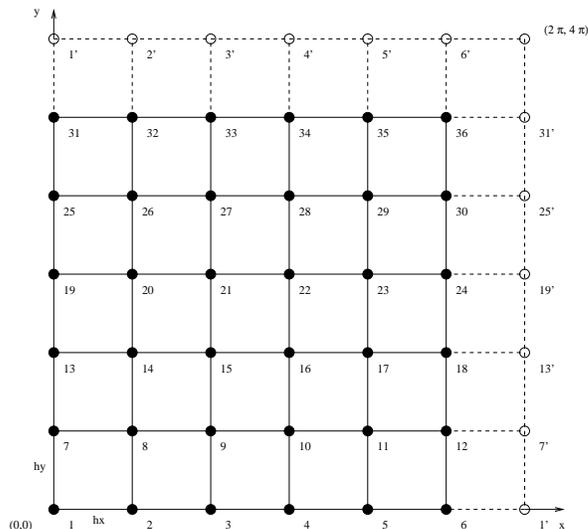


Fig. 1. Rectangular domain $\Omega = [0, 2\pi) \times [0, 4\pi)$ with doubly periodic boundaries and a sample 6×6 mesh. The solid circles indicate genuine mesh points while empty circles denote virtual points that correspond to the boundary mesh points on the opposite side of the domain.

The system of PDEs (2.9) is defined on a rectangular domain and doubly periodic boundary conditions are used [18]. There are several ways to numerically implement the doubly periodic boundary conditions [22], and we take the so-called “wraps around” approach that is provided as an option in [1], in which the top, the bottom as well as the left and the right boundaries are glued together. Fig. 1 helps to illustrate the point by showing an example of the 6×6 mesh. The standard five-point-stencil discretization of the PDE centered at the mesh point number 8 requires corresponding information from four “neighboring” points: 2, 9, 14 and 7. For the discretization centered at the boundary mesh point 7 the four “neighboring” points are 1, 8, 13 and 12. In some sense, a system with the periodic boundary conditions is more global than the one with Dirichlet or Neumann boundaries and this “globalization” may have an impact on the performance of the domain decomposition [28,30] method used to solve the system of equations.

4 One-level Newton-Krylov-Schwarz Methods

At each time step, we compute the solution by solving a system of nonlinear algebraic equations $G(E) = 0$, which is obtained by putting the finite difference equations (3.14) in a certain order. For some algorithms, the orderings of unknowns and the finite difference equations are not important, but for our algorithm to work, it is crucial to order them in the “fully coupled” fashion.

More precisely, we define

$$E = (\phi_{11}, \psi_{11}, U_{11}, F_{11}, \phi_{21}, \psi_{21}, U_{21}, F_{21}, \dots)^T,$$

and

$$G = (G_\phi(1, 1), G_\psi(1, 1), G_U(1, 1), G_F(1, 1), \\ G_\phi(2, 1), G_\psi(2, 1), G_U(2, 1), G_F(2, 1), \dots)^T,$$

where the notations of (3.14), ignoring the superscript, are used. With this kind of ordering the four physical variables at the same mesh points are always together throughout the computations. The mesh points are ordered subdomain by subdomain for the purpose of parallel processing. The system is solved with a one-level Newton-Krylov-Schwarz (NKS) [6,7], which is a general purpose parallel algorithm for solving systems of nonlinear algebraic equations. The Newton iteration is given as: $E_{k+1} = E_k - \lambda_k J(E_k)^{-1} G(E_k)$, $k = 0, 1, \dots$, where E_0 is a solution obtained at the previous time step, $J(E_k) = G'(E_k)$ is the Jacobian at E_k , and λ_k is the steplength determined by a linesearch procedure [12]. Due to doubly periodic boundary conditions, the Jacobian has a one-dimensional null-space that is removed by projecting out a constant [23]. The accuracy of the Jacobian solve is determined by some $\eta_k \in [0, 1)$ and the condition $\|G(E_k) + J(E_k)s_k\| \leq \eta_k \|G(E_k)\|$. The overall algorithm can be described as follows:

- (1) Inexactly solve the linear system $J(E_k)s_k = -G(E_k)$ for s_k using a preconditioned GMRES(30) [27].
- (2) Perform a full Newton step with $\lambda_0 = 1$ in the direction s_k .
- (3) If the full Newton step is unacceptable, backtrack λ_0 using a backtracking procedure until a new λ is obtained that makes $E_+ = E_k + \lambda s_k$ an acceptable step.
- (4) Set $E_{k+1} = E_+$, go to step 1 unless a stopping condition has been met.

In step 1 above we use a right-preconditioned GMRES to solve the linear system; i.e., the vector s_k is obtained by approximately solving the linear system $J(E_k)M_k^{-1}(M_k s_k) = -G(E_k)$, where M_k^{-1} is the one-level additive Schwarz preconditioner. To formally define M_k^{-1} , we need to introduce a partition of Ω . We first partition the domain into non-overlapping subdomains Ω_l , $l = 1, \dots, N$, as in Figure (2). In order to obtain an overlapping decomposition of the domain, we extend each subregion Ω_l to a larger region Ω'_l , i.e., $\Omega_l \subset \Omega'_l$. Only simple box decompositions are considered in this paper – all subdomains Ω_l and Ω'_l are rectangular and made up of integral numbers of fine mesh cells. The size of Ω_l is $H_x \times H_y$. The subdomain Ω'_l is not always simply connected as shown in Figure 2, however, if we pull all the pieces together so that the subdomain is “connected”, we can more conveniently define the size of Ω'_l as $H'_x \times H'_y$, where the H 's are chosen so that the overlap, δ , is uniform

in the number of fine grid cells all around the perimeter, i.e.,

$$\delta = (H'_x - H_x)/2 = (H'_y - H_y)/2 \quad (4.1)$$

for every subdomain. Note that the boundary subdomains are also extended all around their perimeters because of the doubly periodic physical boundary condition. On each extended subdomain Ω'_l , we construct a subdomain preconditioner B_l , whose elements are $B_l^{i,j} = \{J_{ij}\}$, where the node indexed by (i, j) belongs to Ω'_l . The entry J_{ij} is calculated with finite differences

$$J_{ij} = 1/(2\alpha)(G_i(E_j + \alpha) - G_i(E_j - \alpha)), \quad (4.2)$$

where $0 < \alpha \ll 1$ is a constant and a graph coloring techniques is used to efficiently compute the sparse Jacobian [1,11]. We also compute the entries of the Jacobian matrix using analytical expressions for comparison purposes. Homogeneous Dirichlet boundary conditions are used on the subdomain boundary $\partial\Omega'_l$. The classical additive Schwarz preconditioner [8] can be written as

$$M_k^{-1} = \sum_{j=1}^N (R_j)^T B_j^{-1} R_j. \quad (4.3)$$

Let n be the total number of mesh points and n'_l the total number of mesh points in Ω'_l . Then, R_l^δ is an $n'_l \times n$ block matrix that is defined as: its 4×4 block element $(R_l^\delta)_{i,j}$ is an identity block if the integer indices $1 \leq i \leq n'_l$ and $1 \leq j \leq n$ belong to a mesh point in Ω'_l , or a block of zeros otherwise. The R_l^δ serves as a restriction matrix because its multiplication by a block $n \times 1$ vector results in a smaller $n'_l \times 1$ block vector by dropping the components corresponding to mesh points outside Ω'_l . The $n'_l \times n$ block matrix R_l^0 is defined similarly with the exception that its application to a $n \times 1$ vector also zeros components that correspond to the mesh points laying in $\Omega'_l \setminus \Omega_l$. In this paper we use the left restricted additive Schwarz preconditioner [5] which is given by:

$$M_k^{-1} = \sum_{j=1}^N (R_j^0)^T B_j^{-1} R_j^\delta, \quad (4.4)$$

where δ is the size of the inter subdomain overlap. The restricted additive Schwarz preconditioner can be interpreted as using the residuals from the points in the overlapping region, while discarding the computed values in the overlapping region. For the problem we are considering the restricted additive preconditioner performs much better than the classical additive Schwarz preconditioner, especially when the number of processors is large. Various inexact additive Schwarz preconditioners can be constructed by replacing the matrices B_l in (4.4) with convenient and inexpensive to compute matrices, such as those obtained with incomplete factorizations. In this paper we employ the *LU* factorization.

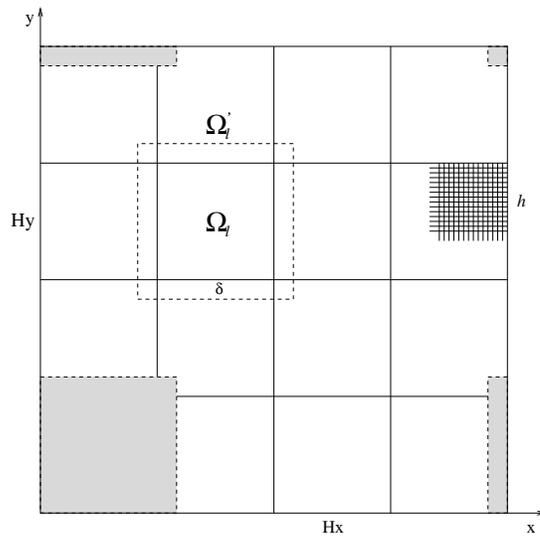


Fig. 2. Decomposition of domain Ω with an overlap δ . The solid lines indicate the partition of the domain into non-overlapping substructures Ω_l of size $H_x \times H_y$ and the dashed rectangle indicates an overlapping subdomain Ω'_l . The filled rectangles show the formation of an extended boundary subdomain with the doubly periodic domain boundary $\partial\Omega$ and the incomplete fine mesh of solid lines illustrates underlying uniform subintervals with mesh size h .

If the classical Schwarz preconditioner is applied to symmetric positive definite systems resulting from a discretization of elliptical problems, then the condition number κ of the preconditioned system satisfies $\kappa \leq C(1 + H/\delta)/H^2$ for the one-level method and $\kappa \leq C(1 + H/\delta)$ for the two-level method, where C is independent of h , H and δ . In the above formulation H is the effective subdomain diameter. The factor $1/H^2$, associated with the number of subdomains on the fine level, lends itself to an increase in numbers of iterations with the increase in the total numbers of subdomains. This increase can be justified by the need to exchange information between distant subdomains. The use of a coarse mesh facilitates the exchange and, potentially, stabilizes the number of iterations. Our problem, however, is of a mixed elliptic/parabolic type, where the growth of the linear iterations numbers does not follow exactly the theoretical predictions for elliptical problems. One consequence of such a deviation is that the growth of the linear iterations numbers is not severe and an introduction of the traditional coarse space does not seem necessary to attain a good parallel performance.

5 Explicit Method

To verify the accuracy of the implicit solution, as well as to compare the parallel performance of the implicit method, we implement an explicit method, in which the equations for the vorticity and the canonical momentum are

first advanced in time and, to guarantee temporal consistency, corresponding equations are solved to update ϕ and ψ . For the time integration we use a second-order explicit Adams formulas albeit modified to allow for adaptive time stepping, in which the time step Δt is determined by the CFL condition. In our current implementation the CFL condition is estimated based on the maximum speed of the signal propagation on a spatial mesh. In particular, from the values of ϕ we compute the spatial components of the plasma's in-plane velocity

$$\mathbf{v} = e_z \times \nabla \phi. \quad (5.1)$$

Then, values of v_x and v_y which are the in-plane velocity components in x and y directions, respectively, are used to compute the CFL estimates:

$$\text{CFL}_x = \left| \frac{v_x \Delta t}{h_x} \right|, \quad \text{and} \quad \text{CFL}_y = \left| \frac{v_y \Delta t}{h_y} \right|, \quad (5.2)$$

where Δt is the proposed time step size and h_x and h_y are the mesh intervals in x and y directions. The CFL estimate is then chosen as the maximum value of estimates from (5.2) over all mesh points. Similar estimates are used for the in-plane values of the magnetic fields that relate to ψ values via:

$$\mathbf{B} = B_0 e_z + \nabla \psi \times e_z. \quad (5.3)$$

If the computed CFL number is

$$\text{CFL} < 0.9, \quad (5.4)$$

then the time step size of Δt is accepted and the temporal updates (the second order explicit Adams), using the same notations as in (3.15), are computed by:

$$U_{i,j}^{k+1} = U_{i,j}^k + \Delta t \left(\frac{3}{2} R_U^k(i, j) - \frac{1}{2} R_U^{k-1}(i, j) \right), \quad (5.5)$$

$$F_{i,j}^{k+1} = F_{i,j}^k + \Delta t \left(\frac{3}{2} R_F^k(i, j) - \frac{1}{2} R_F^{k-1}(i, j) \right).$$

Otherwise, the time step size Δt is reduced until (5.4) is satisfied. Equations (5.5) are, *de facto*, used if no time step resize is required. However, if the time step size needs to be reduced, then the following equations are used:

$$U_{i,j}^{k+1} = U_{i,j}^k + \Delta t R_U^k(i, j) + \frac{(\Delta t)^2}{2\Delta t_1} \left(R_U^k(i, j) - R_U^{k-1}(i, j) \right), \quad (5.6)$$

$$F_{i,j}^{k+1} = F_{i,j}^k + \Delta t R_F^k(i, j) + \frac{(\Delta t)^2}{2\Delta t_1} \left(R_F^k(i, j) - R_F^{k-1}(i, j) \right),$$

where Δt_1 is the size of the previous time step and $\Delta t \neq \Delta t_1$. Equations (5.6) become equations (5.5) assuming $\Delta t = \Delta t_1$.

After the time advancement is done, two linear equations in (3.15) are solved. In the linear solves, one-level, right-sided restricted additive Schwarz preconditioner with LU on all subdomains is used and the preconditioned systems are solved with GMRES.

6 Numerical Experiments

We use Portable Extensible Toolkit for Scientific computation (PETSc), developed at Argonne National Laboratory [1], for our implementations of the algorithms discussed in previous sections. The codes are written in a hostless manner and allow easy switching between different numbers of processors. Each processor is assigned one subdomain, and the information pertaining to the interior of a subdomain is uniquely owned by that processor. In the implicit algorithm implementation, the processor stores subvectors and a block of the Jacobian matrix associated with an extended subdomain. At the beginning of every nonlinear iteration, the E -dependent local blocks of the Jacobian, as well as the preconditioning matrices, are computed. The preconditioning matrices are factored, and the upper and lower triangular parts are stored. After a solution of each subproblem is obtained, those portions that lie within the overlapping regions are sent to neighboring processors to complete collective operations. Participating processors communicate with each other by message passing using MPI [19].

Similarly in the explicit implementation, each processor stores subvectors and blocks of linear operators associated with the linear systems. Subtracting the nonlinear solver functionality from the implicit implementation discussed above, the remaining procedures are identical in both implicit and explicit implementations. Testing platforms include a beowulf cluster and an IBM BlueGene/L system.

To illustrate the model behavior, we choose nominal values of the inertial skin depth $d_e = 0.08$ and the ion sound Larmor radius $\rho_s = 0.24$. The normalized resistivity and viscosity are $\eta = 10^{-3}$ and $\nu = 10^{-3}$, respectively. The time in the system is normalized to the Alfvén time $\tau_A = \sqrt{4\pi n m_i} l_x / B_{y0}$, where B_{y0} is the characteristic magnitude of the equilibrium magnetic field and l_x is the macroscopic scale length [18]. The computational domain Ω is uniformly partitioned into rectangular meshes up to 1980×1980 in size. The stopping conditions for the iterative processes are given as follows:

- relative reduction in the nonlinear solve: $\|G(E_k)\| \leq 10^{-7} \|G(E_0)\|$,
- absolute tolerance in the nonlinear solve: $\|G(E_k)\| \leq 10^{-7}$,
- relative reduction in the Jacobian solve: $\|r_k\| \leq 10^{-10} \|r_0\|$,
- absolute tolerance in the Jacobian solve: $\|r_k\| \leq 10^{-7}$,

- GMRES restart value: 30.

The same stopping conditions are used for the explicit runs except, of course, for the nonlinear solver parameters.

Typical solutions are shown in Fig. 3, Fig. 4 and Fig. 5 at simulated times $t = 100\tau_A$, $t = 200\tau_A$ and $t = 290\tau_A$, respectively, and corresponding current density contours are given in Fig. 6. The initial perturbation in ϕ produces a feature-rich behavior in ψ , U , and F . The four variables in the system evolve at different rates: ϕ and ψ evolve at a slower rate than F and U . For $\eta = 10^{-3}$ and $\nu = 10^{-3}$ we observe an initial slow evolution of the four fields and the current density profiles up to time $100\tau_A$ and the solution blows up at time near $290\tau_A$. In the middle of the domain the notorious “X” structure is developed, as can be seen in the F contours, where the magnetic flux is reconnected. Similar reconnection areas are developed on the boundaries of the domain due to the periodicity of boundary conditions and the shape of the initial ϕ perturbation. In the reconnection regions sharp current density peaks Fig. 7(a) are formed. In the Fig. 7(b) a time evolution of the maximum of the absolute value of current density is shown. The value of $\log_{10}|j|_{max}$ is plotted as a function of time given in the Alfvén units. We can see that given $\eta = \nu = 10^{-3}$ for the normalized resistivity and viscosity, the current density peak height increases by several orders of magnitude over the time period of about $300\tau_A$. At time $280\tau_A$ the system is ushered into a highly nonlinear phase of evolution, which is characterized by the fast growth of current density peaks and rapid topological transformations of the current density surface. The smooth profiles of the current density surface split into sharp, teeth-like features in the x direction, while still preserving the smooth profile in the y direction. Fig. 8 helps to illustrate the system’s behavior. Eventually, as the numerical time integration proceeds further into the highly nonlinear phase, the implicit algorithm encounters convergence difficulties.

The onset time of the highly nonlinear phase as well as the exact features of current density surfaces depend strongly on the physical parameters. The ratio ρ_s/d_e effectively dominates the behavior of the system. The larger the ratio is, the faster it becomes nonlinear and the harder it is to solve. For $\rho_s/d_e = 1.0$, reducing the values of normalized resistivity η and viscosity ν prolongs the linear evolution phase. As the ratio is increased, the dependence on η and ν values becomes non-monotonic and further studies are in order to fully characterize the details of the system’s evolution. Keeping the $\rho_s/d_e = 3.0$ and changing the values for resistivity and viscosity allow additional insight into the processes taking place in the system. As can be observed in Fig. 11 (a), where the $\log_{10}|j|_{max}$ is plotted for different values for resistivity and viscosity, while keeping the ρ_s/d_e constant, and Fig. 11 (b), where the $\log_{10}|j|_{max}$ is plotted for different ratios ρ_s/d_e , while keeping values for resistivity and viscosity constant at $\eta = \nu = 10^{-2}$, there is evidence to support the claim that

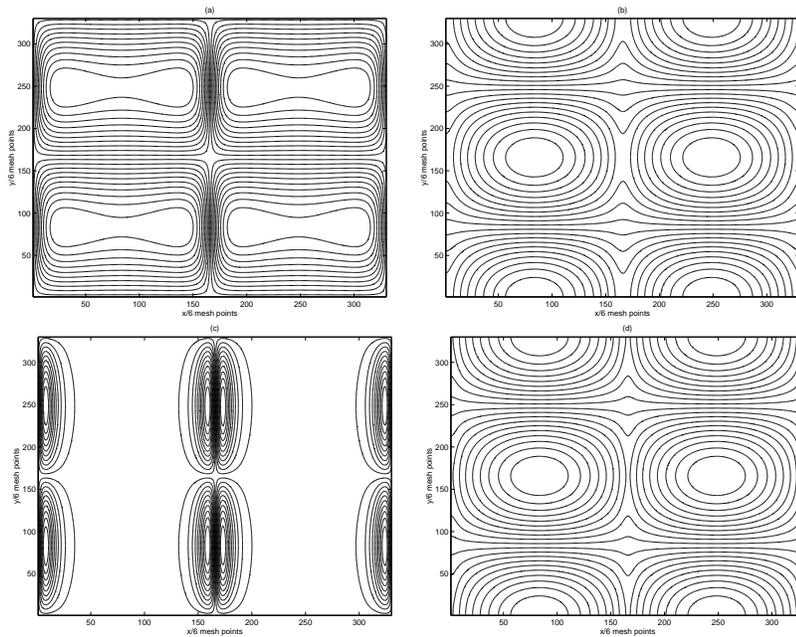


Fig. 3. Contour plots of ϕ (a), ψ (b), U (c), and F (d). The results are obtained on 1980×1980 mesh, $\Delta t = 1.0\tau_A$, time $t = 100\tau_A$, $\eta = 10^{-3}$, $\nu = 10^{-3}$, implicit time stepping.

the system has a pseudo steady state solution for $\eta = \nu = 10^{-2}$, where only minor changes occur after simulated time $200\tau_A$. However, no pseudo steady state solutions have been observed for other values of η and ν .

Using the nominal values of $\eta = 10^{-3}$, $\nu = 10^{-3}$ and $\rho_s/d_e = 3.0$, we compare solutions obtained by our implicit method with these obtained with the explicit method. Fig. 9 shows that the third-order implicit method allows for much larger time steps and produces a solution that is very close to the solution obtained with the explicit algorithm, where the size of the time step is determined by the CFL constraint. Fig. 10, where compute times are plotted for both explicit and implicit methods using different meshes and varying numbers of processors, demonstrates the superior performance of the implicit algorithm over the explicit one in the cases of fine meshes.

Next, we look at some of the machine-dependent properties of the algorithm. Our main focus is on the compute time scalability, which is an important quality in evaluating parallel algorithms. Again, the nominal values for physical parameters are chosen for our experiments. Table 1 shows parallel performance results obtained at $t = 100\tau_A$, $t = 200\tau_A$ and $t = 280\tau_A$. The restart sequence allows runs at specified simulated times, while preserving the temporal discretization order. For instance, in the top part of the Table 1 the code is restarted at $t = 100\tau_A$ and run for ten time steps of size $\Delta t = 1.0\tau_A$ to stop at $t = 110\tau_A$. The compute time as well as numbers of linear and nonlinear iterations are recorded. Similar techniques are also used at $t = 200\tau_A$ and

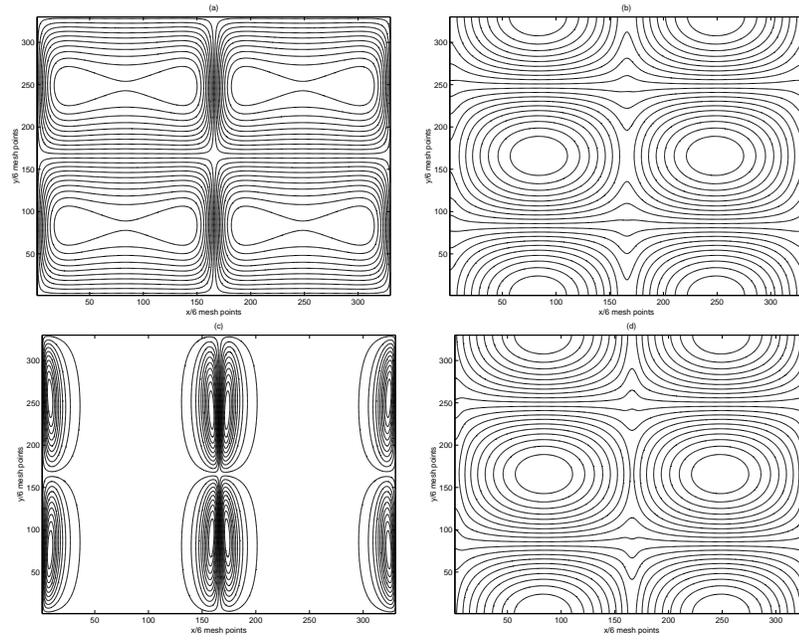


Fig. 4. Contour plots of ϕ (a), ψ (b), U (c), and F (d). The results are obtained on 1980×1980 mesh, $\Delta t = 1.0\tau_A$, time $t = 200\tau_A$, $\eta = 10^{-3}$, $\nu = 10^{-3}$, implicit time stepping.

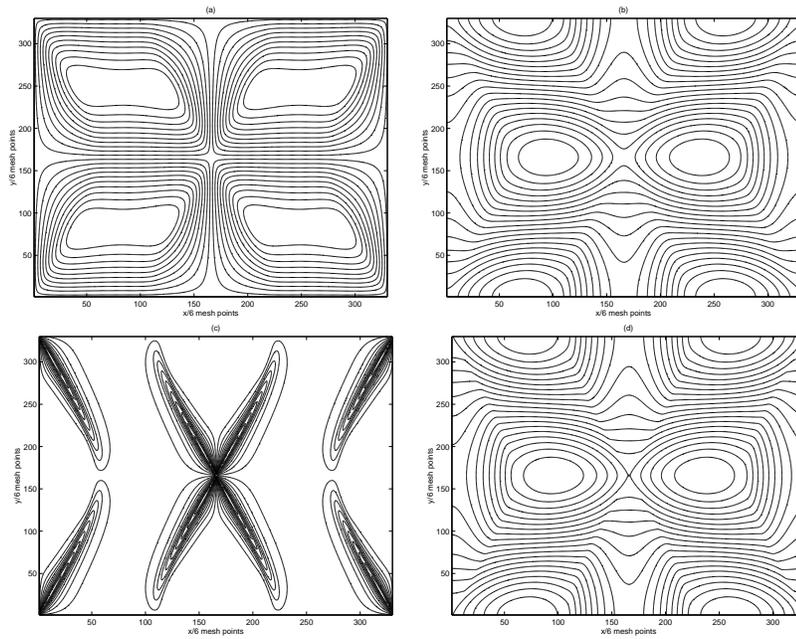


Fig. 5. Contour plots of ϕ (a), ψ (b), U (c), and F (d). The results are obtained on 1980×1980 mesh, $\Delta t = 1.0\tau_A$, time $t = 290\tau_A$, $\eta = 10^{-3}$, $\nu = 10^{-3}$, implicit time stepping.

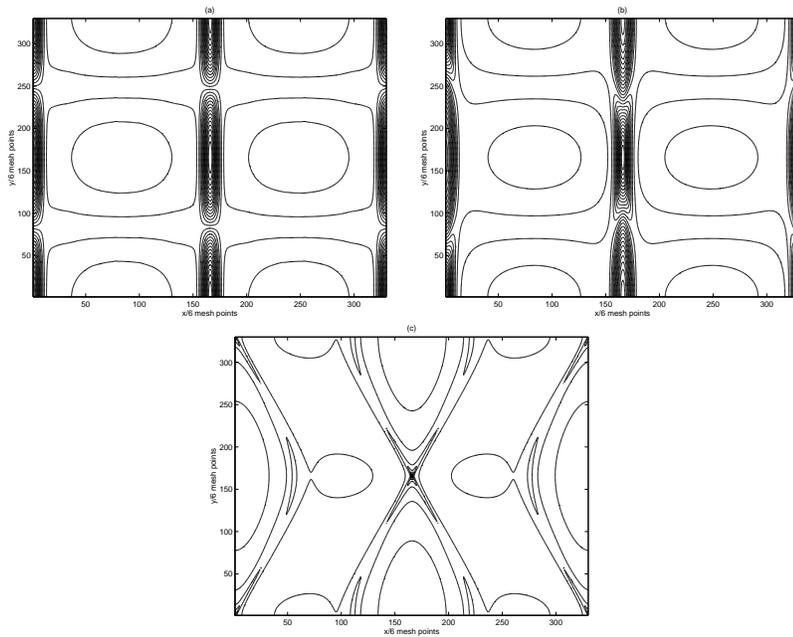


Fig. 6. Contour plots of the current density at $t = 100\tau_A$ (a), $t = 200\tau_A$ (b), and $t = 290\tau_A$ (c). The results are obtained on 1980×1980 mesh, $\Delta t = 1.0\tau_A$, $\eta = 10^{-3}$, $\nu = 10^{-3}$, implicit time stepping.

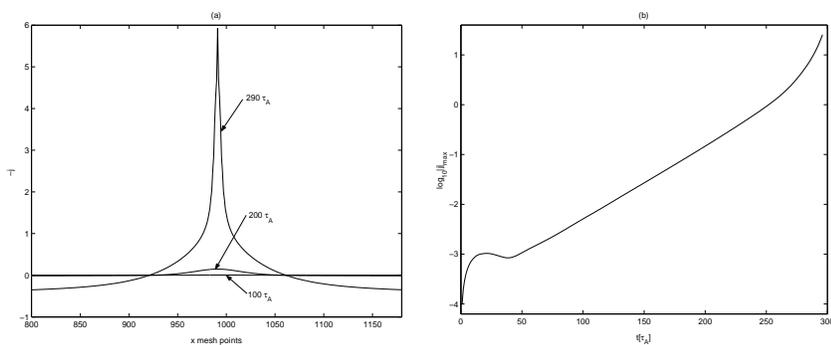


Fig. 7. (a) Time evolution of current density peaks in the reconnection region, 1980×1980 mesh, $\eta = 10^{-3}$, $\nu = 10^{-3}$, $\Delta t = 1.0\tau_A$, $np = 225$ with inter-domain overlap $\delta = 12$, implicit algorithm. (b) Logarithm of $|j|_{max}$ as a function of time obtained over all mesh points with the implicit method ($\Delta t = 1.0\tau_A$) on 1980×1980 mesh with $\eta = 10^{-3}$ and $\nu = 10^{-3}$, $np = 225$.

$t = 280\tau_A$. The restarting functionality allows code runs with varying number of processors.

For a given number of subdomains, the inter-domain overlapping size is chosen to minimize the computing time. That is, referring again to Table 1, in the case of $np = 900$ at $t = 280\tau_A$ the inter-domain overlap $\delta = 13$ results in the compute time of 742.8 seconds, while larger and smaller overlap sizes result in longer compute times. This optimal overlap size varies for different numbers of subdomains effectively ensuing a range of the overlapping sizes $\delta \in [7, 16]$.

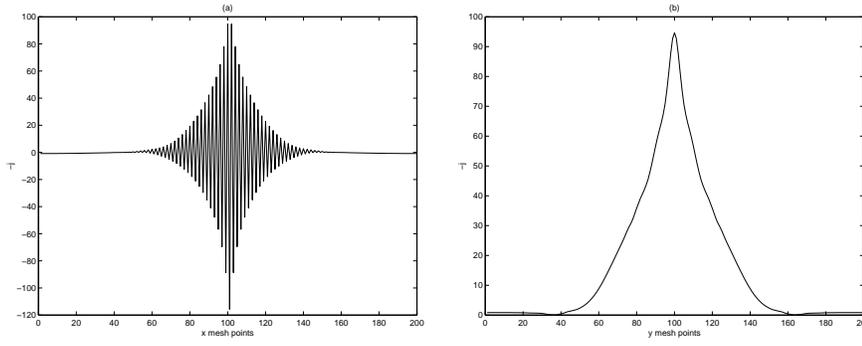


Fig. 8. The highly nonlinear phase of the system's evolution with current density peaks profiles in x direction (a) and y direction (b). The results are obtained on 200×200 mesh, $\eta = 10^{-3}$, $\nu = 10^{-3}$, $\Delta t = 1.0\tau_A$, $np = 1$, implicit algorithm.

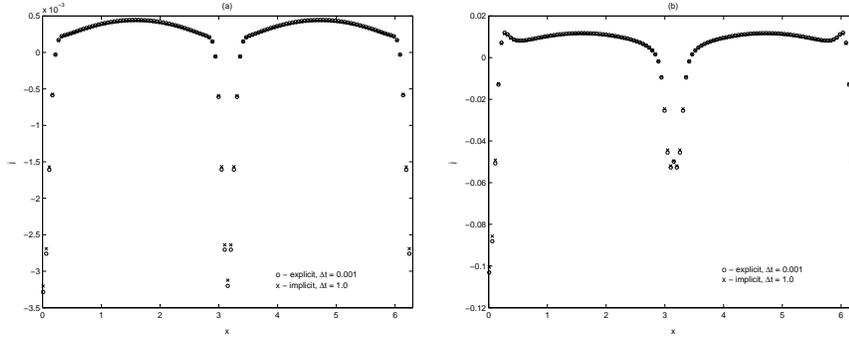


Fig. 9. Current density profiles comparisons obtained with explicit and implicit methods, 600×600 mesh, $\eta = 10^{-3}$, $\nu = 10^{-3}$, $\rho_s = 0.24$, $d_e = 0.08$, the implicit time step size $\Delta t = 1.0\tau_A$ and the explicit time step size $\Delta t = 0.001\tau_A$. (a) $t = 100\tau_A$. (b) $t = 200\tau_A$.

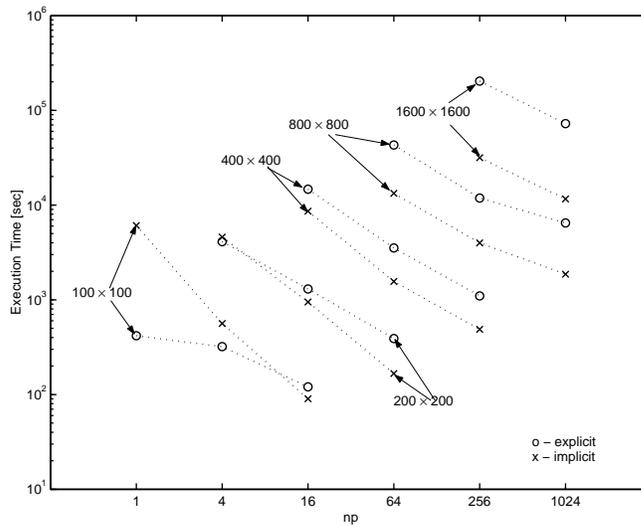


Fig. 10. Execution time comparison between explicit and implicit methods on 100×100 , 200×200 , 400×400 , 800×800 and 1600×1600 meshes obtained with $np = 1, 4, 16, 64, 256, 1024$ processors.

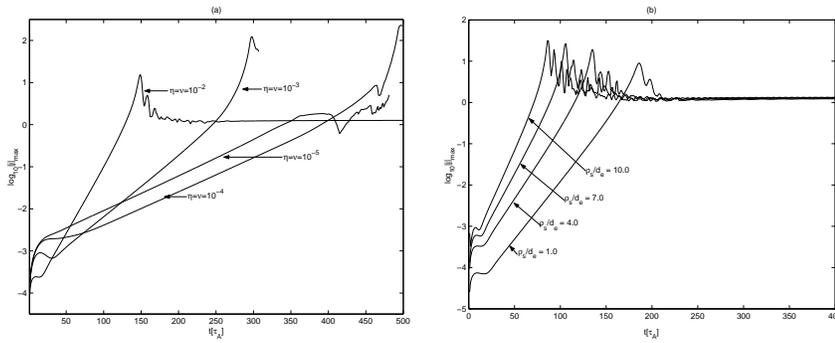


Fig. 11. Current density maxima as a function of time, 600×600 mesh, $\Delta t = 1.0\tau_A$, $np = 25$, implicit algorithm. (a) – constant ratio $\rho_s/d_e = 3.0$ with $\eta = \nu \in [10^{-5}, 10^{-2}]$. (b) – constant $\eta = \nu = 10^{-2}$ and $\rho_s/d_e \in [1.0, 10.0]$.

We show the parallel speedup curves in Fig. 12, where the speedup value is calculated as $t(225)/t(np)$, which gives a ratio of time needed to solve the problem with 225 processors to the time needed to solve the problem with np processors on 1980×1980 mesh. A comparison to the ideal speedup is made in every plot. The speedup plots are reported at simulated times $t = 100\tau_A$, $t = 200\tau_A$ and $t = 280\tau_A$. It is observed that the compute time scalability deteriorates as the problem becomes more and more nonlinear. Indeed, at $t = 100\tau_A$ the speedup is nearly linear. However, for $t = 200\tau_A$ and $t = 280\tau_A$ the experimental speedup becomes sublinear for $np > 400$. Some of the speedup degradation can be attributed to the insufficient problem size pointing out that the problems sizes about $15M$ unknowns are not large enough for $np > 1000$ and the communication cost takes its toll on the overall performance. Larger problem sizes may sufficiently improve speedup values.

7 Conclusions and Future Work

A parallel, nonlinearly implicit method is used for solving a model MHD problem. The proposed fully coupled implicit scheme with the third-order temporal discretization allows much larger time steps than the explicit method, while still preserving the solution accuracy. Moreover, the fully coupled implicit scheme demonstrates superior compute time performance when compared to the explicit scheme. This performance of the implicit method is attained on fine meshes with a large number of processors. It is also observed that the one-level Schwarz preconditioner is sufficient for times up to $t = 150\tau_A$, where the execution time speedup is nearly perfect (Fig. 12(a)) and the deterioration of iteration numbers is mild. However, for the latter half of the simulation ($150\tau_A - 300\tau_A$), the one-level method needs some improvement.

Future continuation of this work may include solutions of the MHD problem on finer meshes with a larger number of processors and with multilevel versions

Table 1

Scalability with respect to the number of subdomains, one-level algorithm, 1980×1980 mesh, LU factorization for all subproblems, $\Delta t = 1.0\tau_A$, 10 time steps taken at simulation times $t = 100\tau_A$, $t = 200\tau_A$ and $t = 280\tau_A$. For every global domain partition the inter-domain overlap $\delta \in [7, 16]$ such that the compute time is minimized over the range of test cases, where the $\delta \in [5, 20]$. Physical parameters are given as follows: $\eta = \nu = 10^{-3}$, $d_e = 0.08$ and $\rho_s = 0.24$. The measurements are conducted with the number of processors $np = 225, 324, 400, 484, 900, 1936, 2025$ and each processor is assigned one subdomain.

$t = 100\tau_A$			
np	Compute Time [sec]	Total Nonlinear Iterations	Linear/Nonlinear
225	716.8	10	42.2
324	457.8	10	46.1
400	342.3	10	56.0
484	284.7	10	56.4
900	151.7	10	63.0
1936	101.0	10	96.5
2025	90.8	10	98.4
$t = 200\tau_A$			
225	965.5	11	65.6
324	754.3	12	53.6
400	465.7	10	68.7
484	427.8	10	74.0
900	276.8	12	67.7
1936	218.6	14	162.4
2025	189.9	13	149.5
$t = 280\tau_A$			
225	2473.1	24	113.5
324	1691.9	24	127.7
400	1359.6	24	135.1
484	1185.0	25	141.3
900	742.8	25	181.0
1936	514.8	27	226.6
2025	504.8	26	244.3

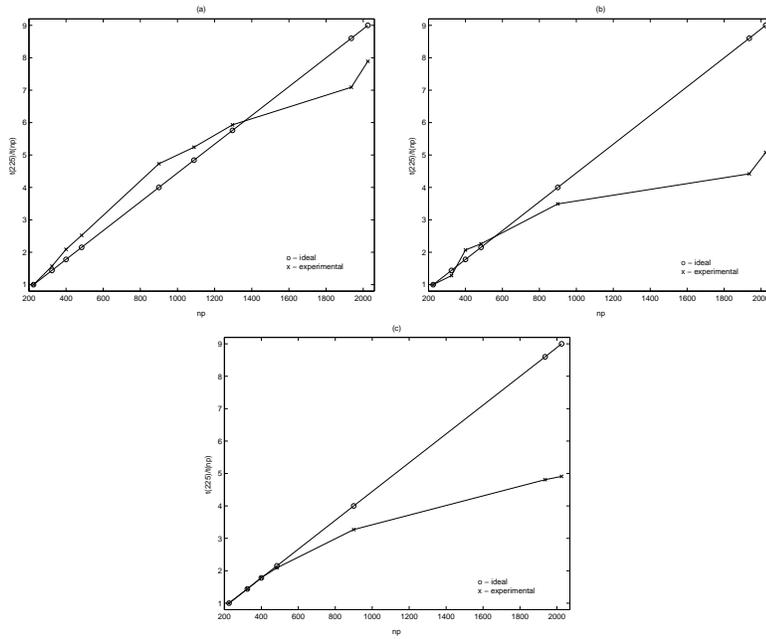


Fig. 12. Speedup curves $t(np = 225)/t(np)$ as compared to the ideal speedup for 1980×1980 mesh, LU factorization for all subproblems, $\Delta t = 1.0\tau_A$, 10 time steps at times $t = 100\tau_A$ (a), $t = 200\tau_A$ (b) and $t = 280\tau_A$ (c). Physical parameters are: $\eta = \nu = 10^{-3}$, $d_e = 0.08$ and $\rho_s = 0.24$. The straight line in the graphs shows the ideal speedup.

of the algorithm. Longer time integration with various η and ν values, as well as higher ρ_s to d_e ratios, may be helpful in the further understanding of the fully coupled implicit algorithm for the numerical solutions of MHD problems.

The importance of mesh adaptivity in resolving the near singular behavior of the current density in the final stages of the simulation has been quantified in [16] for the case of h -type Cartesian adaptive mesh refinement. We plan to extend our fully implicit parallel approach to r -type adaptive mesh refinement, by drawing points into the near singular layer and introducing additional metric terms into the governing equations, without changing the logical organization of the data or load balancing.

Acknowledgments

We would like to thank the PETSc team for technical support and Dr. Kai Germaschewski and Prof. Scott Parker for useful discussions.

References

- [1] S. Balay, K. Buschelman, V. Eijkhout, W. D. Gropp, D. Kaushik, M. Knepley, L. C. McInnes, B. F. Smith, and H. Zhang, *PETSc Users Manual*, ANL-95/11 - Revision 2.3.1, Argonne National Laboratory, (2006).
- [2] D. Biskamp, *Nonlinear Magnetohydrodynamics*, Cambridge University Press, New York, 1993.
- [3] S. I. Braginskii, *Reviews of Plasma Physics*, volume 1, Consultants Bureau, New York, 1965.
- [4] E. Cafaro, D. Grasso, F. Pegoraro, F. Porcelli, and A. Saluzzi, *Invariants and geometric structures in nonlinear Hamiltonian magnetic reconnection*, Phys. Rev. Lett. 80(20) (1998) 4430–4433.
- [5] X.-C. Cai and M. Sarkis, *A restricted additive Schwarz preconditioner for general sparse linear systems*, SIAM J. Sci. Comput. 21 (1999) 792-797.
- [6] X.-C. Cai, W. D. Gropp, D. E. Keyes, R. G. Melvin, and D. P. Young, *Parallel Newton-Krylov-Schwarz algorithms for the transonic full potential equation*, SIAM J. Sci. Comput. 19 (1998) 246-265.
- [7] X.-C. Cai, W. D. Gropp, D. E. Keyes, and M. D. Tidriri, *Newton-Krylov-Schwarz methods in CFD*, Proceedings of the International Workshop on the Navier-Stokes Equations, Notes in Numerical Fluid Mechanics, R. Rannacher, eds. Vieweg Verlag, Braunschweig, (1994).
- [8] X.-C. Cai, *Additive Schwarz algorithms for parabolic convection-diffusion equations*, Numer. Math. 60 (1990) 41-62.
- [9] L. Chacón, D. A. Knoll, and J. M. Finn, *An implicit, nonlinear reduced resistive MHD solver*, J. Comput. Phys. 178 (2002) 15–36.
- [10] L. Chacón and D. A. Knoll, *A 2D high- β Hall MHD implicit nonlinear solver*, J. Comput. Phys. 188 (2003) 573–592.
- [11] T. F. Coleman and J. J. Moré, *Estimation of sparse Jacobian matrices and graph coloring problems*, SIAM J. Numer. Anal. 20 (1983) 187-209.
- [12] J. E. Dennis and R. B. Schnabel, *Numerical Methods for Unconstrained Optimization and Nonlinear Equations*, SIAM, Philadelphia, 1996.
- [13] A. Edwards, D. Campbell, W. Engelhardt, H. Fahrback, R. Gill, R. Granetz, S. Tsuji, B. Tubbing, A. Weller, J. Wesson, and D. Zasche, *Rapid collapse of a plasma sawtooth oscillation in the JET tokamak*, Phys. Rev. Lett. 57(2) (1986) 210–213.
- [14] R. Fitzpatrick, *Scaling of forced magnetic reconnection in the Hall-magnetohydrodynamical Taylor problem with arbitrary guide field*, Phys. Plasmas 11(8) (2004) 3961–3968.

- [15] R. Fitzpatrick, *Scaling of forced magnetic reconnection in the Hall-magnetohydrodynamic Taylor problem*, Phys. Plasmas 11(3) (2004) 937–946.
- [16] K. Germaschewski, A. Bhattacharjee, R. Grauer, D. Keyes, and B. Smith, *Using Krylov-Schwarz methods in an adaptive mesh refinement environment*, Proceedings of the Chicago Workshop on Adaptive Mesh Refinement Methods, Lecture Notes in Computational Science and Engineering, Springer 41 (2003) 115–124.
- [17] R. J. Goldston and P. H. Rutherford, *Introduction to Plasmas Physics*, IOP Publishing, 1995.
- [18] D. Grasso, F. Pegoraro, F. Porcelli, and F. Califano, *Hamiltonian magnetic reconnection*, Plasma Phys. Control Fusion 41 (1999) 1497–1515.
- [19] W. Gropp, E. Lusk, and A. Skjellum, *Using MPI: Portable Parallel Programming with the Message-Passing Interface*, The MIT Press, Cambridge, 1994.
- [20] E. Hairer, S. Nørsett, and G. Wanner, *Solving Ordinary Differential Equations I*, Springer-Verlag, New York, 1993.
- [21] G. Jothiprasad, D. Mavriplis, and D. Caughey, *Higher-order time integration schemes for the unsteady Navier-Stokes equations on unstructured meshes*, J. Comput. Phys. 191 (2003) 542–566.
- [22] O. A. Ladyzhenskaya, *The Boundary Value Problems of Mathematical Physics*, Springer-Verlag, New York, 1985.
- [23] J. Lewis and R. Rehm, *The numerical solution of a nonseparable elliptic partial differential equation by preconditioned conjugate gradients*, J. Res. Nat. Bur. Standards 85 (1980), no. 5, 367–390.
- [24] M. Ottaviani and F. Porcelli, *Nonlinear collisionless magnetic reconnection*, Phys. Rev. Lett. 71(23) (1993) 3802–3805.
- [25] M. Ottaviani and F. Porcelli, *Fast nonlinear magnetic reconnection*, Phys. Plasmas 2(11) (1995) 4104–4117.
- [26] D. R. Reynolds, R. Samtaney, and C. S. Woodward, *A fully implicit numerical method for single-fluid resistive magnetohydrodynamics*, J. Comput. Phys. (2006).
- [27] Y. Saad, *Iterative Methods for Sparse Linear Systems*, SIAM, 2003.
- [28] B. Smith, P. Bjørstad, and W. Gropp, *Domain Decomposition. Parallel Multilevel Methods for Elliptic Partial Differential Equations*, Cambridge University Press, New York, 1996.
- [29] H. R. Strauss and D. W. Longcope, *An adaptive finite element method for magnetohydrodynamics*, J. Comput. Phys. 147 (1998) 318–336.
- [30] A. Toselli and O. Widlund, *Domain Decomposition Methods – Algorithms and Theory*, Springer-Verlag, Berlin, 2005.