PREPARED FOR THE U.S. DEPARTMENT OF ENERGY,
UNDER CONTRACT DE-AC02-76CH03073

PPPL-3600
UC-70

PPPL-3600

A Parallel Two-fluid Code for Global Magnetic Reconnection Studies

by

J.A. Breslau and S.C. Jardin

August 2001

PPPL
PRINCETON PLASMA PHYSICS LABORATORY

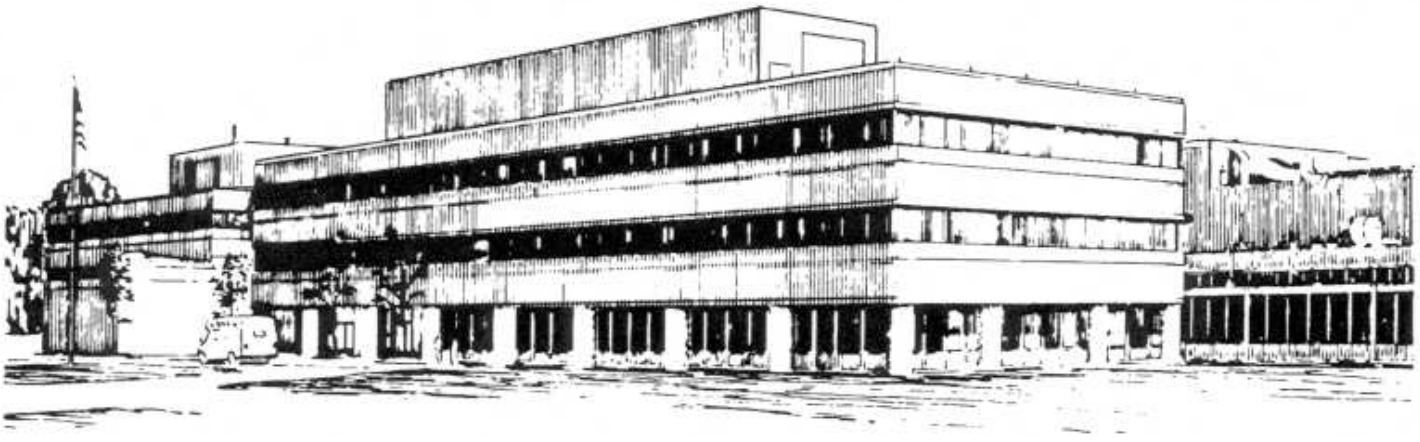

PRINCETON PLASMA PHYSICS LABORATORY
PRINCETON UNIVERSITY, PRINCETON, NEW JERSEY

PPPL Reports Disclaimer

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

Availability

This report is posted on the U.S. Department of Energy's Princeton Plasma Physics Laboratory Publications and Reports web site in Calendar Year 2001. The home page for PPPL Reports and Publications is: http://www.pppl.gov/pub_report/

DOE and DOE Contractors can obtain copies of this report from:

U.S. Department of Energy
Office of Scientific and Technical Information
DOE Technical Information Services (DTIS)
P.O. Box 62
Oak Ridge, TN 37831

Telephone: (865) 576-8401

Fax: (865) 576-5728

Email: reports@adonis.osti.gov

This report is available to the general public from:

National Technical Information Service
U.S. Department of Commerce
5285 Port Royal Road
Springfield, VA 22161

Telephone: 1-800-553-6847 or
(703) 605-6000

Fax: (703) 321-8547

Internet: http://www.ntis.gov/ordering.htm

A Parallel Two-Fluid Code for Global Magnetic Reconnection Studies

J. A. Breslau and S. C. Jardin
Princeton Plasma Physics Laboratory

E-mail: jbreslau@pppl.gov, jardin@pppl.gov

This paper describes a new algorithm for the computation of two-dimensional resistive magnetohydrodynamic (MHD) and two-fluid studies of magnetic reconnection in plasmas. It has been implemented on several parallel platforms and shows good scalability up to 32 CPUs for reasonable problem sizes. A fixed, non-uniform rectangular mesh is used to resolve the different spatial scales in the reconnection problem. The resistive MHD version of the code uses an implicit/explicit hybrid method, while the two-fluid version uses an alternating-direction implicit (ADI) method. The technique has proven useful for comparing several different theories of collisional and collisionless reconnection.

Key Words: extended MHD, magnetic reconnection

1. INTRODUCTION

Magnetic reconnection is a universal phenomenon. The annihilation of magnetic field lines with the attendant conversion of magnetic field energy to thermal and kinetic energy occurs both in astrophysical regimes such as solar flares, accretion disks, and magnetospheric substorms [1] and in laboratory plasmas, such as tokamaks and reversed field pinches [2]. In spite of its ubiquity, there is much about the magnetic reconnection process that is still not well understood, and it is consequently the subject of a great deal of active research and debate [3].

A typical plasma undergoing magnetic reconnection is characterized by a spatial separation into two regions. One is an outer, "ideal" region in which magnetic field gradients are small and the evolution of the plasma is governed by the non-dissipative equations of ideal MHD. The reconnection itself occurs in an inner, "diffusion" region in which the magnetic field reverses direction over a short distance, resulting in a high current density and necessitating the inclusion of one or more non-ideal terms that become important where spatial gradients are steep. Although reconnection occurs in a small, localized region, it results in topological

changes in the magnetic field structure that strongly affect the evolution of the global plasma.

Previous two-dimensional numerical studies of reconnection have typically taken one of two approaches to deal with this separation of scales. They have either restricted their attention to the vicinity of the diffusion region [4]–[9]; or they have advanced both regions using the equations of reduced MHD [10]–[13]. The former approach has the disadvantage of requiring the specification of boundary conditions for the mass inflow and outflow rates. If these rates are determined by the diffusion layer physics, it is difficult to determine them from these local studies. The global approach avoids this problem but has typically required the omission of potentially important physical effects, such as compressibility and out-of-plane fields and flows, in the interest of computational efficiency.

Ideally, one would like to model the global reconnecting plasma using a complete set of compressible MHD equations on a mesh fine enough to resolve the details of the inner reconnection region. The parallel code described in this paper does this, solving the two-fluid MHD equations to time-evolve a model system for magnetic reconnection consisting of merging spheromaks or flux tubes with varying relative helicities. Efficient parallelization and non-uniform mesh spacing allow it to solve for the global plasma accurately while simultaneously providing high resolution of the narrow current sheet structures.

The following section outlines the considerations that went into the numerical modeling of the reconnection problem, and the decisions that were made. Following that, we present the results of some convergence checks. We conclude with some examples of useful results.

2. MODELING THE RECONNECTION PROBLEM

2.1. The System of Equations

The two-fluid MHD equations being solved consist of the following:
continuity:

$$\frac{\partial \rho}{\partial t} + \nabla \cdot (\rho \mathbf{v}) = 0 \tag{1}$$

force balance:

$$\rho \left(\frac{\partial \mathbf{v}}{\partial t} + \mathbf{v} \cdot \nabla \mathbf{v} \right) = \mathbf{J} \times \mathbf{B} - \nabla p + \nu \nabla^2 \mathbf{v} \tag{2}$$

low-frequency Maxwell's equations:

$$\frac{\partial \mathbf{B}}{\partial t} = -\nabla \times \mathbf{E} \tag{3a}$$

$$\mathbf{J} = \nabla \times \mathbf{B} \tag{3b}$$

$$\nabla \cdot \mathbf{B} = 0 \tag{3c}$$

generalized Ohm's law:

$$\mathbf{E} + \mathbf{v} \times \mathbf{B} = \eta \mathbf{J} + \frac{\mathbf{J} \times \mathbf{B} - \nabla p_e}{ne} \tag{4}$$

electron and ion pressure equations:

$$\frac{\partial p_e}{\partial t} + \mathbf{v}_e \cdot \nabla p_e = -\frac{5}{3} p_e \nabla \cdot \mathbf{v}_e + \frac{2}{3} \left[\eta |\mathbf{J}|^2 + \nabla \cdot \left(\frac{m_i \kappa_e}{Z+1} \nabla \frac{p_e}{\rho} \right) - Q \right] \tag{5}$$

$$\frac{\partial p_i}{\partial t} + \mathbf{v} \cdot \nabla p_i = -\frac{5}{3} p_i \nabla \cdot \mathbf{v} + \frac{2}{3} \left[-\nu \mathbf{v} \cdot \nabla^2 \mathbf{v} + \nabla \cdot \left(\frac{m_i \kappa_i}{Z+1} \nabla \frac{p_i}{\rho} \right) + Q \right]. \tag{6}$$

Here ρ is the mass density, $\mathbf{v}_e = \mathbf{v} - \mathbf{j}/ne$ is the electron fluid velocity, $p = p_e + p_i$ is the sum of the electron and ion fluid pressures, ν is a scalar viscosity, η is a scalar resistivity, $\kappa_{e,i}$ are scalar thermal conductivities, and Q represents the transfer of heat between the two species as specified by Braginskii [14].

The single-fluid (resistive MHD) equations consist of (1)–(4) and the sum of (5) and (6) in the limit $1/ne \to 0$ so that the right-hand-side of (4) contains only the resistive term, and $\mathbf{v}_e \to \mathbf{v}$. A still simpler, but less physical, system can be arrived at by eliminating (1), (5), and (6), holding plasma density and pressure constant in time and space. Solutions of all three sets of equations are implemented by versions of the code described in this paper.

2.2. Geometry

Cylindrical. The code was first developed to look at axisymmetric spheromak merging. The default version therefore uses cylindrical coordinates R, ϕ, and z, treating $\hat{\phi}$ as the ignorable direction. Each row in the non-uniform mesh corresponds to a particular value of z, and each column to a particular value of R.

The most general form for a magnetic field satisfying (3c) in an axisymmetric system is in terms of two scalar functions $\psi(R,z)$ and $g(R,z)$:

$$\mathbf{B} = \nabla\phi \times \nabla\psi + g\nabla\phi. \tag{7}$$

Note that this follows from the introduction of an axisymmetric magnetic vector potential of the form

$$\mathbf{A} = \nabla\phi \times \nabla\Omega - \psi\nabla\phi, \tag{8}$$

where g and Ω are related by

$$g = \Delta^\star_\perp \Omega \equiv \frac{\partial^2 \Omega}{\partial R^2} - \frac{1}{R}\frac{\partial \Omega}{\partial R} + \frac{\partial^2 \Omega}{\partial z^2}. \tag{9}$$

With the choice of the poloidal flux ψ and toroidal field amplitude $G \equiv g/R$ as the two field variables to advance in time, the $\nabla \cdot \mathbf{B} = 0$ condition remains satisfied automatically. The current density can then be calculated as

$$\mathbf{J} = \nabla \times \mathbf{B} = \nabla g \times \nabla\phi + \Delta^\star_\perp \psi \nabla\phi. \tag{10}$$

Combining (3a), (3b), and (4), the time evolution of these new field variables is given by

$$\frac{\partial \psi}{\partial t} + \mathbf{v}_e \cdot \nabla \psi = \eta \Delta^\star_\perp \psi \tag{11}$$

$$\frac{\partial G}{\partial t} = \hat{\phi} \cdot \nabla \times \left(\mathbf{v}_e \times \mathbf{B} - \eta \nabla \times \mathbf{B} + \frac{\nabla p_e}{ne} \right) \tag{12}$$

Cartesian. A Cartesian version of this code also exists, in which $\hat{R} \to \hat{x}$ and $\hat{\phi} \to \hat{y}$. Here the magnetic field is written

$$\mathbf{B} = \hat{y} \times \nabla \psi + G \hat{y}, \tag{13}$$

while the current density is given by

$$\mathbf{J} = \nabla G \times \hat{y} + \nabla^2 \psi \, \hat{y}. \tag{14}$$

Equations (11) and (12) still hold, with the substitutions $\Delta^\star_\perp \to \nabla^2$ and $\hat{\phi} \to \hat{y}$ respectively.

2.3. Boundary Conditions

Conducting walls. We regard the plasma as being confined within a metal box with perfectly conducting walls. The wall boundary conditions for the magnetic field variables ψ and G are determined by the requirement that the tangential electric field vanish at the plasma/conductor boundary. Eq. (4) implies that the tangential current must vanish there as well. The component of the in-plane current density parallel to the walls is accordingly set to zero by zeroing the appropriate (first) derivative of G. In the Cartesian case, this is equivalent to simply imposing even symmetry of G at the top and side boundaries, while in the toroidal case, the boundary conditions at the sides are dependent to some extent on the radial position of the walls. The out-of-plane current density parallel to the walls must be zeroed at these boundaries as well, which means zeroing $\Delta^\star_\perp \psi$. Again, in the Cartesian case this just imposes odd symmetry across conducting boundaries, while the toroidal case has an R-dependence in the boundary condition at the sides.

Toroidal axis. In the toroidal case, the right boundary is always a wall, but the left boundary may either be a wall (at $R > 0$) or the symmetry axis (at $R = 0$). The boundary conditions on velocity at $R = 0$ in cylindrical geometry are that the radial and toroidal velocities should go to zero, while v_z has even symmetry. The axis boundary conditions on the other variables are obtained by Taylor expansions about $R = 0$. Neither G nor ψ has a zeroth-order term in its expansion (they are odd in R), so both vanish on axis. The density and pressure are even in R. The density, for example, can be written (at a fixed value of z)

$$\rho(R) = \rho(0) + c_1 R^2 + c_2 R^4 + \ldots \tag{15}$$

with c_1 and c_2 constants independent of R. Its value at the axis can therefore be calculated to second order from the two values to the right:

$$\rho(0) \approx \frac{4\rho(\Delta R) - \rho(2\Delta R)}{3}. \tag{16}$$

The same extrapolation is used for the pressure.

Midplane symmetry. The resistive MHD version of the code incorporates the assumption of reflection symmetry about the midplane of the reconnection process in order to gain a free two-fold increase in speed. The bottom boundary of the mesh is not a wall but a reflection boundary. ψ, ρ, p, and v_R have even symmetry about this boundary; v_z and v_ϕ have odd symmetry there. The toroidal field G can have either even or odd symmetry about the midplane or be identically zero. These three cases correspond to co-helicity, counter-helicity, or null helicity reconnection respectively (see Section 4.1).

Left-right symmetry. The Cartesian MHD case contains an additional degree of reflection symmetry, about a line passing vertically through the X-point, perpendicular to the midplane. The left boundary is accordingly made a reflection boundary, with ψ, G, ρ, p, and v_z even and v_R and v_ϕ odd. This gives another doubling of execution speed.

Hall symmetry. The new terms introduced in the adjustment to a two-fluid code destroy the two reflection symmetries just mentioned. However, the co- and null helicity cases in Cartesian geometry possess a new type of symmetry of their own: a double reflection, or, equivalently, a 180° rotation. That is, the left half of the mesh is identical to the right half rotated through 180 degrees. It is therefore possible to make use of this symmetry to once again restrict computation to half the mesh.

2.4. Discretization on a Non-Uniform Mesh

The reconnecting system is expected to contain at least two different spatial scales of interest: the global scale and the scale of the current sheet thickness. The most natural way in which to attempt to resolve systems with multiple scales is to vary the spacing of the mesh on which the equations are discretized, packing the cells tightly in the regions containing fine structure and leaving them more coarse in the regions with little spatial variation. One way to accomplish this is to use adaptive mesh spacing, in which the mesh size and shape change dynamically in response to the steepness of gradients in the evolving system. For the systems under consideration here, however, the inherent symmetry makes it possible to predict in advance the location and orientation of the current sheet where the highest resolution will be needed: it will be a horizontal line at the midplane of the mesh. This makes it possible to use a much simpler scheme in which a particular spatially varying mesh spacing is chosen at the beginning of a simulation and then left constant over the course of the run.

For explicit finite different schemes (those in which the values of the variables at the old time are used to approximate the spatial derivatives that are used to advance to the new time), the maximum time step size Δt for numerical stability is proportional to $(\Delta x)^2$ for parabolic terms, which contain second derivatives. The

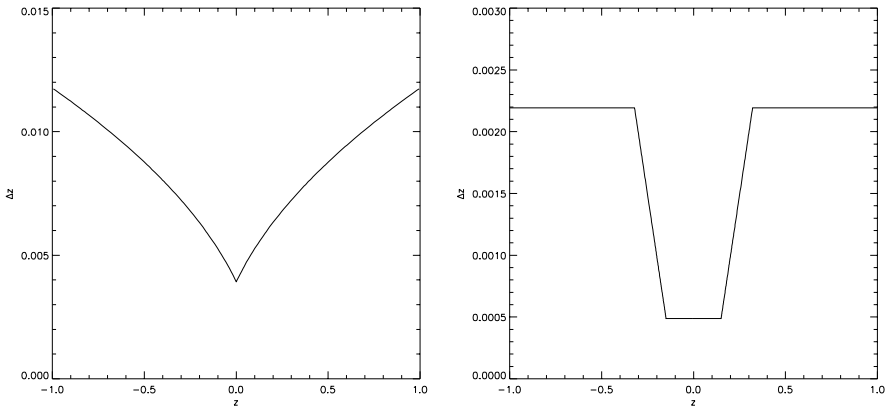

a. b.

FIG. 1. The two different mesh spacing schemes used by MRC. **a.** Linear increase of spacing with row number. **b.** Geometric increase of spacing with row number outside constant-spacing fine zone.

presence of any regions of very small grid spacing can thus unduly restrict the time step required for the implementation of an explicit scheme. In contrast, *implicit* difference schemes (those in which the values of the variables at the advanced time are computed self-consistently based at least in part on evaluation of the spatial derivatives at the advanced time) are generally numerically stable for all time step sizes regardless of mesh spacing. The drawback of implicit schemes is that they require matrix inversions at every time step, making them more time-consuming per step than their explicit counterparts.

A compromise between these two types of schemes is possible that is particularly conducive to parallelization (see section 2.6). Since less fine structure is expected in the $\hat{R}$ (or $\hat{x}$) direction than perpendicular to it, a relatively coarse uniform mesh (about 100 zones across) is imposed in that direction and solved explicitly. It is known that the $\hat{z}$ direction *will* possess fine structure in the form of the midplane current sheet, so the mesh spacing is made fine in that region, increasing in size outside of it. This direction is then solved implicitly. Thus there is no restriction on the time step based on the spacing in this direction. (See section 2.5 below for more on this splitting technique).

The single-fluid MHD version of the code determines the z spacing by a scheme in which the size of the smallest zone, always located at the midplane, is specified as a fraction of the size the zones would have if the mesh were uniform. The spacing is then incremented by a constant amount for each successive zone moving away from the midplane, with the constant chosen in such a way as to guarantee that the requested number of total zones just fills the problem domain (Figure 1a).

In the two-fluid version, we instead specify the fraction of the total number of zones that is to be packed into a region of constant spacing centered on the midplane and extending out to $z = \pm 0.15$. Outside this region, the spacing increases by 1%

per zone until no more increase is needed to fill the remaining space with the remaining zones, at which point the spacing again becomes uniform (Figure 1b).

2.5. Numerical Integration

Implicit/explicit hybrid methods.

The time advance method is motivated by the need to maintain second-order accuracy in time and space, to allow a reasonable time step that is not restricted by the smallest zone spacing in the reconnection layer, and to be consistent with efficient parallel implementation.

Let n be the time index, j the R (or x) coordinate index, and k the z coordinate index. We define a state vector $\mathbf{U}$ as follows:

$$\mathbf{U} = [\rho, \psi, G, v_R, v_\phi, v_z, p_e, p_i] \tag{17}$$

and adopt the notation that $\mathbf{U}^n_{j,k}$ represents the solution at time level t^n and spatial location (R_j, z_k). The time advance equations (1)–(6) can be written in the linearized form

$$\frac{\partial \mathbf{U}}{\partial t} + \mathsf{A} \cdot \frac{\partial \mathbf{U}}{\partial x} + \mathsf{B} \cdot \frac{\partial^2 \mathbf{U}}{\partial x^2} + \mathsf{C} \cdot \frac{\partial \mathbf{U}}{\partial z} + \mathsf{D} \cdot \frac{\partial^2 \mathbf{U}}{\partial z^2} = 0, \tag{18}$$

where A, B, C, and D are the 8×8 Jacobian matrices.

In the single-fluid MHD version of our code, we advance (18) using the hybrid explicit/implicit method

$$\mathbf{U}^{n+1/2} = \mathcal{L} \cdot \mathbf{U}^n - \frac{\Delta t}{2} \Big[\mathsf{A} \cdot \Delta^1_x \mathbf{U}^n + \mathsf{B} \cdot \Delta^2_x \mathbf{U}^n + \theta (\mathsf{C} \cdot \Delta^1_z \mathbf{U}^{n+1/2} + \mathsf{D} \cdot \Delta^2_z \mathbf{U}^{n+1/2}) + (1-\theta) \left(\mathsf{C} \cdot \Delta^1_z \mathbf{U}^n + \mathsf{D} \cdot \Delta^2_z \mathbf{U}^n \right) \Big] \tag{19a}$$

$$\mathbf{U}^{n+1} = \mathbf{U}^n - \Delta t \Big[\mathsf{A} \cdot \Delta^1_x \mathbf{U}^{n+1/2} + \mathsf{B} \cdot \Delta^2_x \mathbf{U}^n + \theta (\mathsf{C} \cdot \Delta^1_z \mathbf{U}^{n+1} + \mathsf{D} \cdot \Delta^2_z \mathbf{U}^{n+1}) + (1-\theta) \left(\mathsf{C} \cdot \Delta^1_z \mathbf{U}^n + \mathsf{D} \cdot \Delta^2_z \mathbf{U}^n \right) \Big] \tag{19b}$$

where the operator Δ^m_w denotes the second order finite difference approximation to the mth derivative with respect to w. The "implicitness" coefficient θ, logically restricted to the range $0 < \theta < 1$, is generally chosen to be $1/2$ for second-order accuracy in time. The operator $\mathcal{L}$ is defined such that

$$(\mathcal{L} \cdot \mathbf{U}^n)_{j,k} = \begin{pmatrix} \rho_{j,k} \\ \psi_{j,k} \\ G_{j,k} \\ \frac{1}{2}(v_{R\,j-1,k} + v_{R\,j+1,k}) \\ v_{\phi\,j,k} \\ \frac{1}{2}(v_{z\,j-1,k} + v_{z\,j+1,k}) \\ p_{e\,j,k} \\ p_{i\,j,k} \end{pmatrix}^n \tag{20}$$

I.e., for the two poloidal velocity components v_R and v_z, the explicit piece of (19) has the character of the Lax-Wendroff method, while for all other components of $\mathbf{U}$, it behaves like the simpler (and less dissipative) predictor-corrector method. For dissipation coefficients $\nu \geq 10^{-2}$ and $\eta \geq 10^{-5}$, this scheme has been found to remain numerically stable for tens of thousands of time steps.

The Alternating Direction Implicit method.

In the two-fluid case, the retention of the Hall term in Ohm's law [those terms in Equation (4) proportional to $1/ne$] makes the hybrid scheme just presented impractical. This can be seen by considering the linearized form taken by the Cartesian versions of (11) and (12) for small perturbations about a constant field $B_0\hat{x}$ on time scales sufficiently rapid to freeze out the ion motion:

$$\frac{\partial\tilde{\psi}}{\partial t} = \frac{B_0}{ne}\frac{\partial\tilde{G}}{\partial x} \tag{21}$$

$$\frac{\partial\tilde{G}}{\partial t} = -\frac{B_0}{ne}\frac{\partial^3\tilde{\psi}}{\partial x^3} \tag{22}$$

These equations have traveling wave solutions (known as "whistler" waves) whose dispersion relation

$$\omega^2 = \frac{k^4 B_0^2}{n^2 e^2} \tag{23}$$

indicates a phase velocity that varies inversely with wavelength. Because the mesh contains perturbations with wavelengths as small as $2\Delta x$, where Δx is the minimum spacing, the Courant-Friedrichs-Lewy condition on the time step in an explicit solution of these equations is

$$\Delta t \leq \frac{ne(\Delta x)^2}{\pi B_0}. \tag{24}$$

The quadratic dependence on Δx is restrictive even in the vicinity of the poloidal field null; as B_0 increases, *any* explicit scheme will clearly become impractical. Thus a fully implicit treatment is called for.

The most straightforward method would be to evaluate all terms in the difference scheme at the advanced time level. For a grid of linear size N, this would result in an $(8N^2)\times(8N^2)$ banded matrix with band size $8N$. Instead, we base our solution on the Alternating-Direction Implicit (ADI) Method [16]. This is a splitting method, in which one alternates between treating the derivatives in one direction implicitly and the other explicitly each for half a time step. As with the hybrid scheme above, each matrix equation solved advances just one column (or row) of the mesh, yielding block tridiagonal arrays of manageable size. Schematically, the algorithm is as follows:

$$\mathbf{U}^{n+1/2} = \mathbf{U}^n - \frac{\Delta t}{2}\left(\mathsf{A}\cdot\Delta_x^1\mathbf{U}^n + \mathsf{B}\cdot\Delta_x^2\mathbf{U}^n \right.$$

$$\left. + \mathsf{C}\cdot\Delta_z^1\mathbf{U}^{n+1/2} + \mathsf{D}\cdot\Delta_z^2\mathbf{U}^{n+1/2}\right) \tag{25a}$$

$$\mathbf{U}^{n+1} = \mathbf{U}^{n+1/2} - \frac{\Delta t}{2}\left(\mathsf{A}\cdot\Delta_x^1\mathbf{U}^{n+1} + \mathsf{B}\cdot\Delta_x^2\mathbf{U}^{n+1}\right.$$
$$\left. + \mathsf{C}\cdot\Delta_z^1\mathbf{U}^{n+1/2} + \mathsf{D}\cdot\Delta_z^2\mathbf{U}^{n+1/2}\right) \tag{25b}$$

Note that both equations are now half-steps.

The state vector $\mathbf{U}$ is augmented in this version with a ninth element, $J_\phi/R = \Delta_\perp^\star\psi/R^2$ (or $J_y = \nabla^2\psi$) in order to allow implicit evaluation of third spatial derivatives arising from the Hall terms in (12) without adding two extra bands to the block tridiagonal matrix.

The ADI method has the same linear numerical stability properties as the fully implicit method, i.e., there is no restriction on the size of the time step. In its practical implementation, however, the step must still be kept fairly small for purposes of accuracy. Even this fails to fully solve the numerical stability problems arising from the inclusion of the Hall term. There remains a tendency for short-wavelength numerical noise in the out-of-plane current density to grow quickly and swamp the physical solution. This problem is addressed by the addition of artificial higher order dissipation terms to the poloidal magnetic field and momentum equations. These terms are treated explicitly, but have time-dependent coefficients that are constrained to be just stable for all time step sizes. We utilize a "hyper-resistivity" term of the form

$$\frac{\partial\psi}{\partial t} = \ldots - \frac{2(\Delta z)^4}{ne\ \Delta t}\left(\frac{\partial^4\psi}{\partial x^4} + \frac{\partial^4\psi}{\partial z^4}\right). \tag{26}$$

Note that the hyper-resistivity vanishes when $1/ne \to 0$ and it is not needed. We also employ a "hyper-viscosity" of the form

$$\frac{\partial\mathbf{v}_p}{\partial t} = \ldots - \frac{0.2(\Delta z)^4}{\Delta t}\left(\frac{\partial^4\mathbf{v}_p}{\partial x^4} + \frac{\partial^4\mathbf{v}_p}{\partial z^4}\right), \tag{27}$$

where the subscript $p \in R, z$ represents the velocity components in the symmetry plane. Finally, an increased physical viscosity (separate from the artificial hyper-viscosity) is added along the walls, taking the form

$$\nu(R,z) = \nu_0 + \nu_w \sum_{walls} \exp\left(-\frac{d^2}{\lambda^2}\right), \tag{28}$$

where ν_0 is the "background" viscosity of the original code, ν_w is the elevated value, d represents the perpendicular distance of the point (R,z) from the given wall, and λ is the characteristic width of the viscous boundary region. Typically, $\lambda = 0.05$ and $\nu_w = 0.1$, while for comparison, ν_0 falls in the range 10^{-4} to 0.1.

The time step in the single-fluid implicit/explicit hybrid code is limited by the parabolic terms in the explicit (R or x) dimension; the dominant restriction is typically $\Delta t \le (\Delta x)^2/2\eta$. In contrast, the two-fluid ADI version's maximum step size is free of this restriction. However, we find that when the Hall term is non-zero, the maximum practical time step is proportional to $\Delta z/c$, where c is the fast magnetosonic wave speed. This is because the hyper-resistivity and -viscosity coefficients, each effectively proportional to Δt^{-1}, must be sufficiently large to stabilize

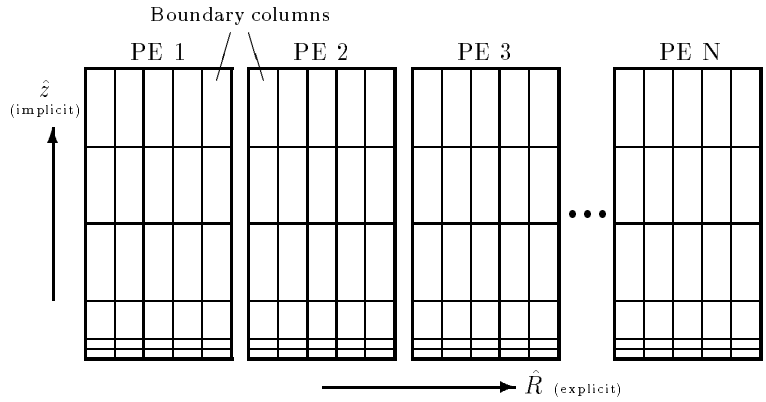

FIG. 2. Schematic view of the 1D domain decomposition used by MRC.

the advancement scheme against numerical noise originating from the Hall term. These terms limit the constant of proportionality in this Courant-like condition to approximately 0.3 in practice. Nevertheless, the ADI scheme makes it possible to use step sizes as much as 80 times larger than what would be allowed under the explicit restrictions.

2.6. Parallelization

Domain decomposition. MRC makes use of a one-dimensional domain decomposition, dividing the mesh into strips that extend vertically from top to bottom. This is illustrated schematically in Figure 2. The figure makes it evident that the bottom of the mesh falls at the midplane of the physical domain. The number of adjacent columns to be assigned to each processing element is determined at runtime by simply dividing the total number of columns in the mesh by the number of processes in the run (with provision for assigning different numbers of columns to different processes in the event that this ratio is not an integer). In addition to its proper columns, each process has one or more "ghost" columns on either side of its segment of the mesh; these ghost columns store the values in the adjoining columns located on the neighboring processors, which must be obtained by point-to-point communications.

When combined with the implicit/explicit hybrid algorithm described above, this scheme has several advantages:

- The implicit steps, which require the direct solution of large sparse banded matrix equations, are all done on data distributed in the vertical direction in a single column. Since any one column is contained entirely within a single process, this most time-consuming operation (occupying more than 80% of the total runtime in single-processor benchmarks) can be done without any interprocess communication. This makes favorable scaling possible.

- Because the execution time of the matrix solution algorithm is independent of the contents of the matrix, the time each process spends on the implicit portion

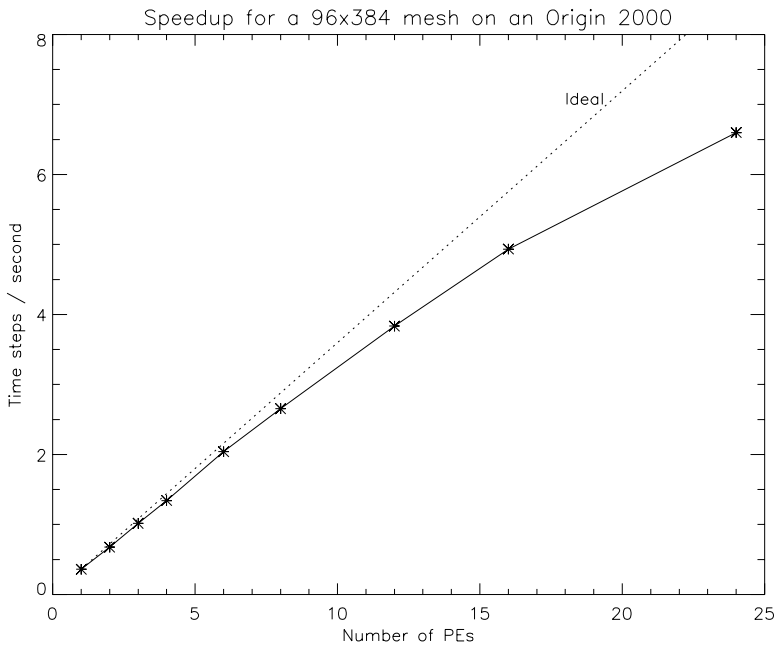

FIG. 3. Parallel speedup of the magnetic reconnection code on an Origin 2000. Comparison with Amdahl's law indicates a parallel fraction of approximately 99%.

of the time step is simply proportional to the number of columns it contains. The assignment of the same number of columns to each process is therefore sufficient to achieve good load balancing.

• The explicit half-time-step is second-order in space, so that the evaluation of derivatives in the horizontal direction at any point requires information only about the nearest neighbors to that point. It is therefore necessary to maintain and communicate only a single ghost column on each side of each process, keeping the amount of data to be communicated reasonably small compared to the total amount in the mesh.

• This communication is itself parallelizable in the sense that PE 3 can send its right boundary information to PE 4 at the same time that PE 1 is sending the corresponding information to PE 2, and so on. The communication time therefore does not increase with increasing N.

The major drawback of this type of domain decomposition is that it limits the maximum number of processors that can be used to no more than the number of columns; in practice, two or three columns per processor are necessary for reasonable efficiency. This places upper limits on the scalability of the code for any constant problem size.

Parallelism and ADI.

1D domain decomposition provides acceptable parallel efficiency for the MHD version of the code (Figure 3), but requires modification for the two-fluid version. With the numerical scheme now implicit in both directions, all the data in the mesh must be broken up into blocks, transposed, and redistributed across all processors

twice each time step in order to ensure that the matrix solution is always done over data residing on a single PE.

The subroutines that perform the system-wide transposition of the global data array exist in two versions. The first performs all its operations at once, in between implicit half steps. Ghost columns or rows are then swapped, and the next half step is performed. The other version represents an attempt to boost the efficiency of the expensive vector all-to-all operation by making use of non-blocking communication. In the non-blocking version of the code, each column (or row) is broken up and sent as soon as it has been computed. Although the send command is initiated immediately, it is not required to be completed until the following column (or row) has been computed. This non-blocking approach has been found to work well on low-bandwidth networks of distributed PCs, where it reduces overall runtime by as much as 20%. However, on highly optimized supercomputers with very high interprocess bandwidth, and on shared-memory systems, it yields little improvement.

In either case, the overall difference in execution speed between the hybrid and ADI versions of the code is about a factor of two. This indicates that the processing bottleneck in both cases is the implicit matrix solution, which, in the ADI case, has replaced the explicit stepping routine. The mesh transposition, whether blocking or non-blocking, does not add significantly to the runtime.

3. CONVERGENCE STUDIES

3.1. Single-fluid code

Figure 4 shows the time behavior of the rate of change of poloidal flux ψ evaluated at the X-point for runs with identical physical parameters but varying mesh spacings and time steps (see Section 4 below for details on the configurations of these runs). As the figure makes clear, neither a halving of the time step nor a doubling of the mesh resolution from their base values has any noticeable effect on the behavior of the reconnection rate $\dot{\psi}(t)$.

3.2. Two-fluid code

The introduction of hyper-resistivity and hyper-viscosity into the two-fluid code makes convergence a more serious and complex issue for its results. Even for many of the $1/ne \rightarrow 0$ cases, which ought to mimic the single-fluid results, it is not possible to achieve numerical stability without retaining some hyper-viscosity coefficient. However, because $\Delta t \propto \Delta z$ in the ADI code, the magnitudes of the hyper-viscosity and hyper-resistivity coefficients decrease as

$$C_{HYPER} \propto \frac{(\Delta z)^4}{\Delta t} \propto (\Delta z)^3. \tag{29}$$

The effects of these artificial terms may thus be isolated by performing a convergence study in zone size. Figure 5 shows the results of one such study, with $\eta = 10^{-3}$, $\nu = 10^{-1}$, and $\chi \equiv \rho/ne = 0.02$. The reconnection time t_R is evidently linear in Δz as the latter approaches zero when the Hall term is turned on. The linear convergence seems to be a consequence of the first-order truncation error in the finite difference approximation for non-uniform mesh spacing. This result nevertheless shows that the effect of the numerical dissipation terms on the behavior

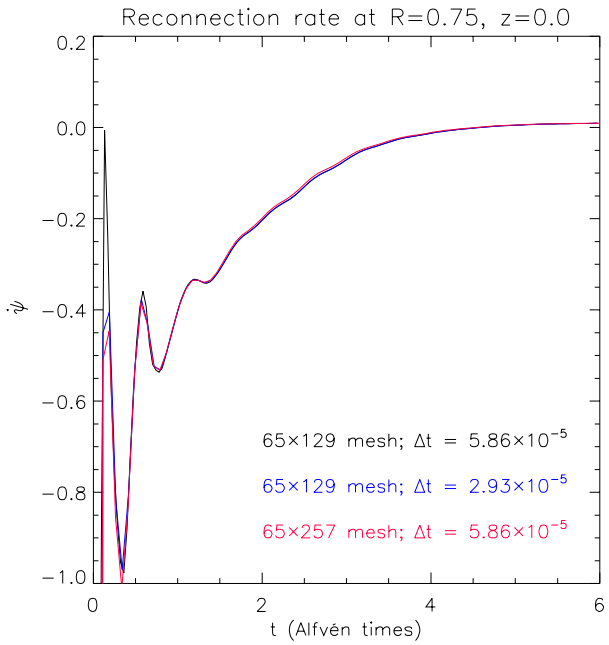

FIG. 4. Time history of reconnection rate $\dot{\psi}$ at X-point for co-helicity study in Cartesian geometry showing convergence of single-fluid code results with Δz and Δt.

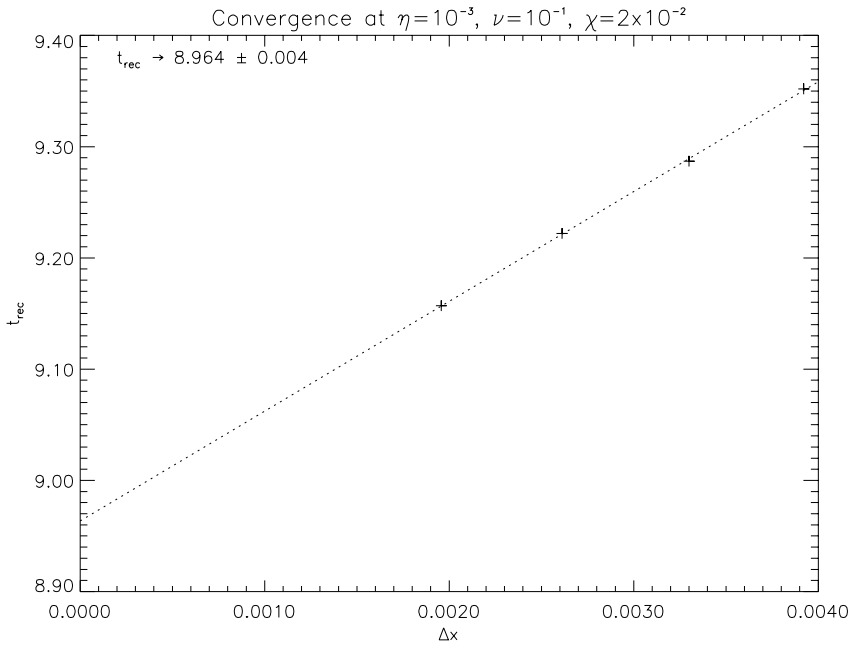

FIG. 5. Convergence of reconnection time with mesh spacing in the two-fluid study.

of the two-fluid code is predictable and can be scaled out by comparing as few as two cases with the same physical parameters.

4. RESULTS

4.1. Benchmarks

First proposed in the 1950s in an attempt to explain coronal heating by solar flares [17], the Sweet-Parker model of magnetic reconnection is still the basic theory of the field, providing a common point of departure for later refinements. A stationary, two-dimensional model based on the equations of resistive MHD, it describes null-helicity reconnection across an X-shaped field null as a purely resistive process taking place in a long, flat, rectangular current sheet whose dimensions determine the reconnection rate. Under a simple set of assumptions, momentum and mass conservation arguments predict that the sheet thickness, and thus the plasma inflow rate, should scale in proportion to the square root of the local resistivity.

Initial conditions.

When referring to spheromak or flux tube reconnection, the terms "co-helicity", "counter-helicity", and "null helicity" refer to the relative orientations of the prevailing toroidal (or out-of-plane) magnetic fields of the two merging islands. In the null helicity case, as in the analytical model, there is no toroidal field. In the co-helicity case, strong parallel toroidal fields tend to reduce the plasma $\beta(\equiv 2p/B^2)$ and to resist compression. Counter-helicity reconnection involves oppositely directed toroidal fields which must annihilate at the mid-plane as reconnection proceeds. The toroidal currents must be parallel in all cases in order for the in-plane magnetic fields to be oppositely directed, a prerequisite for reconnection. The constant-pressure single-fluid study described in this section looked at co-helicity reconnection only.

We begin the simulation with the system in magnetostatic equilibrium so that the possible effects of large initial force imbalances do not influence the results. A global Taylor equilibrium [18] is computed with *three* parallel spheromaks possessing toroidal current in alternating directions. With constant resistivity, this initial state would not undergo reconnection; its magnetic fields would simply decay in place on the resistive time scale. In order to enable fast reconnection studies, the central island holding the other two apart must be destroyed rapidly. This is accomplished by setting an artificially high value of resistivity in the region occupied by this island. Since the sign of ψ alternates along with the sign of the current and is positive in the region to be destroyed, the form chosen for the resistivity is the function

$$\eta(\psi) = \begin{cases} \eta^-, & \psi \le 0 \\ \eta^- + (\eta^+ - \eta^-)\left\{1 - \exp\left[-\left(\frac{\psi}{0.075}\right)^2\right]\right\}, & \psi > 0 \end{cases} \tag{30}$$

where the constant η^- is the normalized resistivity (inverse Lundquist number) of the bulk plasma, typically ranging from 10^{-5} to 10^{-3}, and η^+ is the artificially elevated resistivity, generally set to 10^{-1}. The Gaussian form of the resistivity function for positive ψ ensures continuity of the derivatives of the two branches of the function across $\psi = 0$. With the resistivity in this form, the central island

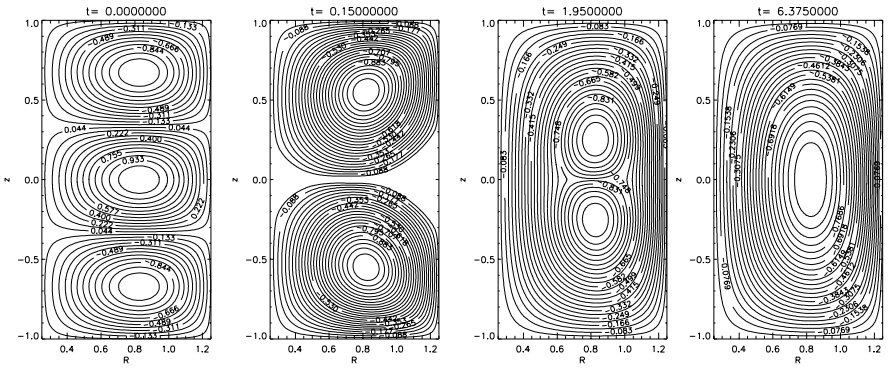

FIG. 6. Poloidal flux contours during successive phases of co-helicity reconnection. Toroidal geometry.

straddling the midplane reflection boundary is only a transient structure, vanishing quickly and allowing the reconnection of the remaining islands to proceed.

Results.

Because these simulations follow the global merging problem from the initial three-island state to the final, reconnected single-island equilibrium, a steady state as described in the Sweet-Parker theory is never reached. Instead, the system consistently undergoes an irreversible change in global field line topology, with all flux private at the beginning of each run and public at the end. This change occurs in several distinct identifiable phases, as illustrated in Figure 6:

1. The central island decays away rapidly due to its artificially high resistivity. Its O-point field null on the midplane is replaced by an X-point, signaling the beginning of the reconnection process.

2. The two remaining islands reconnect at the midplane. Private flux becomes public and plasma in the current sheet is accelerated outward, where it eventually hits the wall and is diverted back around. A quasi-steady state resembling the Sweet-Parker picture forms during this phase. Reconnection ends when the X-point is replaced once again by a new O-point as the merging is completed and a single island is formed.

3. The new low-energy equilibrium decays in place on the resistive time scale. This phase is of little interest for the purposes of this study, so runs are generally terminated at its onset.

The ability of the code to resolve both global and small-scale structures is illustrated in Figure 7, which shows the reconnecting plasma at a single instant in time at successive $10\times$ magnifications of the vertical scale. Both the global islands and the narrow current sheet are clearly visible.

The total reconnection time, as characterized by the time taken from the initial formation of the central X-point to its final disappearance, shows a clear and relatively uniform power law scaling with η at each value of the viscosity, as shown in

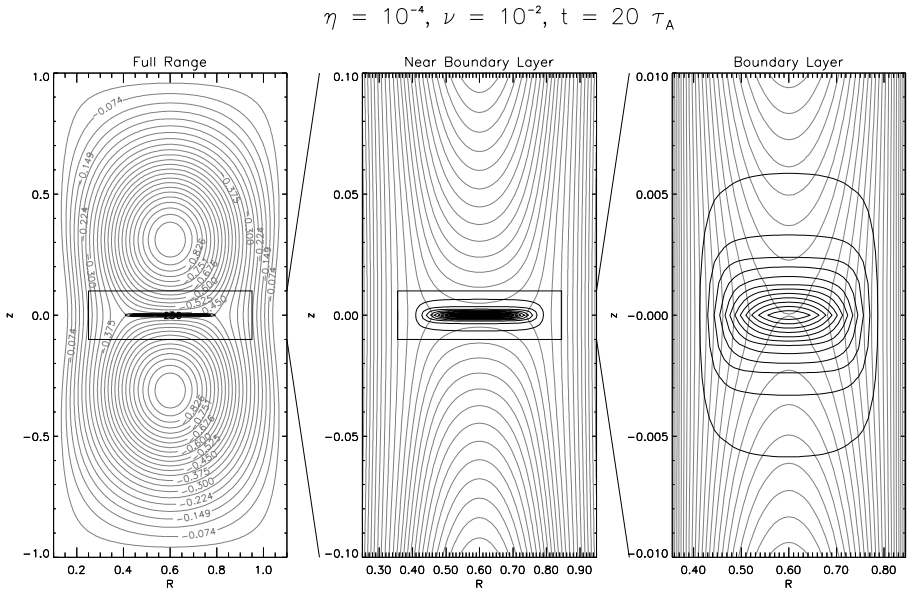

FIG. 7. Poloidal flux contours (gray) and toroidal current density contours (black) during fast null helicity reconnection in Cartesian geometry. The two plots on the right are expanded views of the boxed regions in the plots to their immediate left.

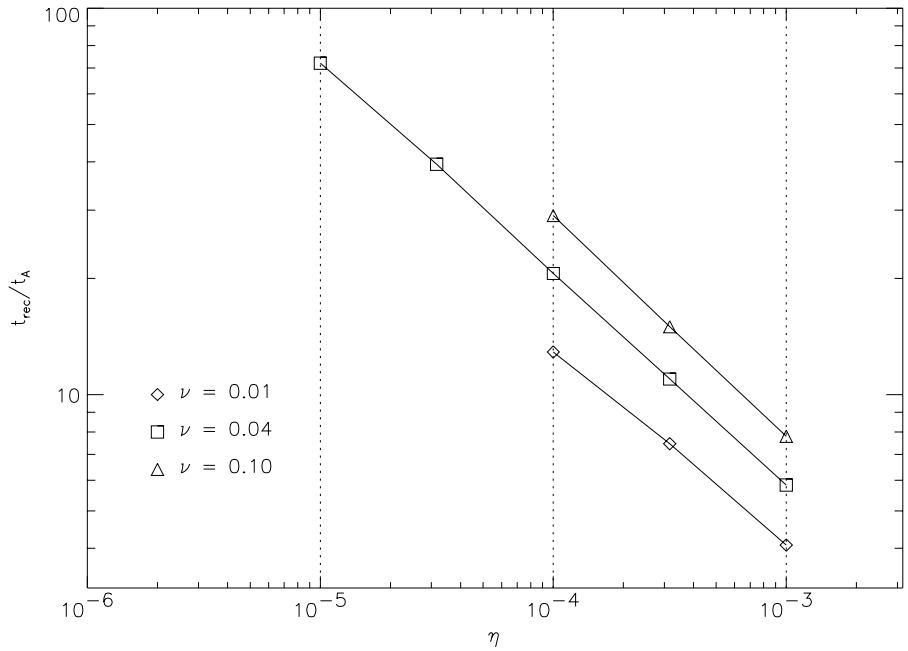

FIG. 8. Scaling of the total reconnection time with resistivity and viscosity; single-fluid, constant-pressure code.

Figure 8's log-log plot. The power law exponent in each case lies between 0.5 and 0.6, consistent with the Sweet-Parker theory.

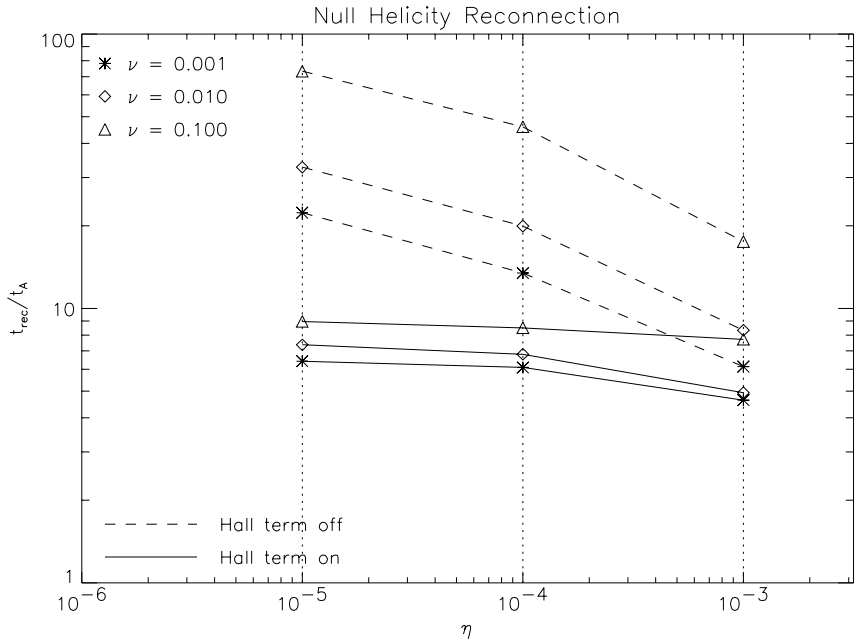

FIG. 9. Comparison of MHD and two-fluid rate scalings of null helicity reconnection. Cartesian geometry.

4.2. Two-Fluid Effects

Progress requires the exploration of physics beyond resistive MHD. The two-fluid code, run with the same initial conditions as the single-fluid study, yielded dramatically different results. Whistler wave physics, entering through the Hall term in (4), greatly accelerated the reconnection rate and broke its dependence on resistivity, as illustrated in Figure 9.

The high resolution provided by MRC in the central region eased the challenge of explaining this new result, supporting a theory of collisionless reconnection first put forth by Biskamp, Schwarz, and Drake [19]: There is a separation of scales between the electron and ion fluids in the new system, with the magnetic field lines tied only to the electrons within a distance of one ion skin depth from the X-point; the ions are effectively unmagnetized within this region. Reconnection therefore occurs within a small region governed by the equations of electron MHD, resulting in super-Alfvénic outflows at the whistler speed; the actual field-line breaking and reconnecting takes place on a still smaller scale in which the hyper-resistivity term dominates (Figure 10). The unmagnetized ions simply flow across the field lines, following the fast-streaming electrons to maintain charge neutrality.

5. CONCLUSIONS

Historically, the prevailing 2D incompressible models of collisional magnetic reconnection have been inadequate, either because of a failure to reproduce the geometry and dynamics observed in physical systems or because of internal inconsistencies. Numerical attempts to adjudicate among them have also been fraught with problems: insufficient resolution, insufficient physical models, and unrealistic

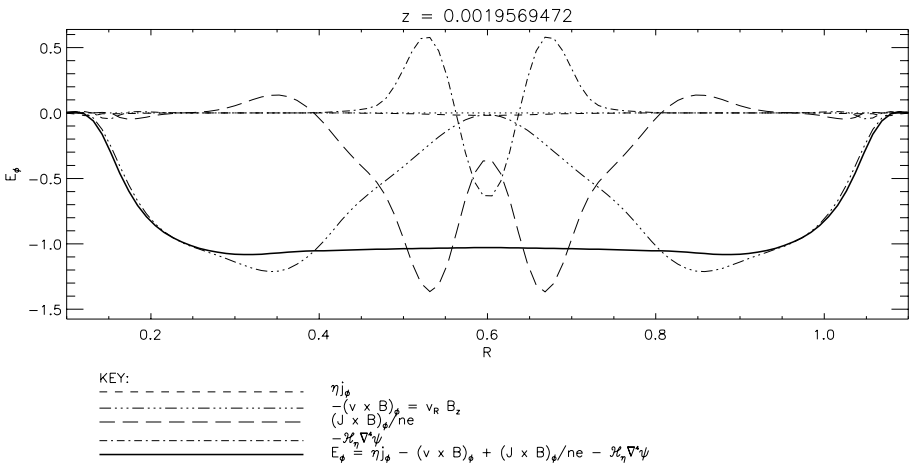

FIG. 10. Out-of-plane electric field along the midplane during null helicity Hall reconnection, by component. $\eta = 10^{-4}$, $\nu = 10^{-2}$. Cartesian geometry.

boundary conditions that attempt to prejudge reconnection layer geometry and merging rates. The Magnetic Reconnection Code is one of the first to address all these issues simultaneously. Its combination of the following attributes appears to be unique among reconnection codes:

- It solves the full two-fluid equations, including the Hall term, electron pressure gradient term, evolving density and electron and ion pressures, and varying out-of-plane magnetic fields.
- It does so on a fine, non-uniform mesh that fully resolves both the global reconnection problem of merging flux cylinders and the local diffusion region geometry.
- It operates at all values of β, including $\beta \sim 1$, appropriate to astrophysical plasmas.
- It utilizes a parallel scalable implicit algorithm that enables it to accomplish these tasks efficiently.

MRC has confirmed what many previous studies have concluded: in the resistive MHD domain, in which the plasma is highly collisional, the Sweet-Parker model applies. There is a long, thin current sheet, with width $\sim B/\nabla B$ and thickness $\propto \eta^{1/2}$. Reconnection is slow, with $v_{in}/v_A \propto \eta^{1/2}$ as well. Adding the Hall term to the standard fluid treatment of the plasma profoundly alters the magnetic reconnection dynamics. For physically realistic values of the ion skin depth, it provides fast, resistivity-independent collisionless reconnection across microscopic current sheets. The whistler-mediated model of collisionless reconnection provides a good explanation for this phenomenon and is validated by this study.

ACKNOWLEDGMENT

It is a pleasure to acknowledge useful discussions with Drs. R. Kulsrud, W. Park, J. Drake, M. Yamada, and H. Ji. This work was supported by USDOE Contract No. DE-AC020-76-CH03073. One of us (JB) was supported by the Fusion Energy Sciences Fellowship Program, administered by the Oak Ridge Institute for Science and Education for the United States Department of Energy.

REFERENCES

1. R. M. Kulsrud, Magnetic reconnection in a magnetohydrodynamic plasma, *Phys. Plas.* **5**, 5 (1998), p. 1599.
2. H. P. Furth, J. Killeen, and M. N. Rosenbluth, Finite-resistivity instabilities of a sheet pinch, *Phys. Fluids* **6**, 4 (1963), p. 459.
3. J. Birn, J. F. Drake, M. A. Shay, *et al.*, GEM magnetic reconnection challenge, *J. Geophys. Res. – Space* **106**, A3 (2001), p. 3715.
4. J. C. Stevenson, Numerical studies of magnetic field annihilation, *J. Plas. Phys.* **7**, 2 (1972), p. 293.
5. M. Ugai and T. Tsuda, Magnetic field-line reconnexion by localized enhancement of resistivity. Part 1. Evolution in a compressible MHD fluid, *J. Plas. Phys.* **17**, 3 (1977), p. 337.
6. T. Hayashi and T. Sato, Magnetic reconnection: acceleration, heating, and shock formation, *J. Geophys. Res.* **83**, A1 (1978), p. 217.
7. L. C. Lee and Z. F. Fu, Multiple X Line Reconnection. 1. A criterion for the transition from a single X line to a multiple X line reconnection, *J. Geophys. Res.* **91**, A6 (1986), p. 6807.
8. C. Anderson and F. Jamitzky, A numerical investigation of magnetic reconnection, *J. Plas. Phys.* **55**, 3 (1996), p. 431.
9. D. A. Uzdensky and R. M. Kulsrud, Two-dimensional numerical simulation of the resistive reconnection layer, . *Phys. Plas.* **7**, 10 (2000), p. 4018.
10. B. V. Waddell, et. al., Non-linear growth of the m=1 tearing mode, *Nucl. Fus.* **16**, 3 (1976), p. 528.
11. D. Biskamp and H. Welter, Coalescence of magnetic islands, *Phys. Rev. Lett.* **44**, 16 (1980), p. 1069.
12. W. Park, D. A. Monticello, and R. B. White, Reconnection rates of magnetic fields including the effects of viscosity, *Phys. Fluids* **27**, 1 (1984), p. 137.
13. R. G. Kleva, J. F. Drake, and F. L. Waelbroeck, Fast reconnection in high temperature plasmas, *Geophys. Res. Lett.* **2**, 1 (1995), p. 23.
14. S. I. Braginskii, Transport processes in a plasma, in *Reviews of Plasma Physics*, edited by M. A. Leontovich (Consultants Bureau, New York, 1965), Vol. 1, p. 205.
15. R. D. Richtmyer and K. W. Morton, *Difference Methods for Initial-Value Problems, 2nd Edition* (Interscience Publishers, New York, 1967), p. 360.
16. N. Gershenfeld, *The Nature of Mathematical Modeling* (Cambridge University Press, Cambridge, UK, 1999), p. 84.
17. E. N. Parker, Sweet's mechanism for merging magnetic fields in conducting fluids, *J. Geophys. Res.* **62** (1957), p. 509.
18. J. B. Taylor, Relaxation of toroidal plasma and generation of reverse magnetic fields, *Phys. Rev. Lett.* **33**, 19 (1974), p. 1139.
19. D. Biskamp, E. Schwarz, and J. F. Drake, Two-fluid theory of collisionless magnetic reconnection, *Phys. Plas.* **4**, 4 (1997), p. 1002.

External Distribution

Plasma Research Laboratory, Australian National University, Australia
Professor I.R. Jones, Flinders University, Australia
Professor João Canalle, Instituto de Fisica DEQ/IF - UERJ, Brazil
Mr. Gerson O. Ludwig, Instituto Nacional de Pesquisas, Brazil
Dr. P.H. Sakanaka, Instituto Fisica, Brazil
The Librarian, Culham Laboratory, England
Library, R61, Rutherford Appleton Laboratory, England
Mrs. S.A. Hutchinson, JET Library, England
Professor M.N. Bussac, Ecole Polytechnique, France
Librarian, Max-Planck-Institut für Plasmaphysik, Germany
Jolan Moldvai, Reports Library, MTA KFKI-ATKI, Hungary
Dr. P. Kaw, Institute for Plasma Research, India
Ms. P.J. Pathak, Librarian, Insitute for Plasma Research, India
Ms. Clelia De Palo, Associazione EURATOM-ENEA, Italy
Dr. G. Grosso, Instituto di Fisica del Plasma, Italy
Librarian, Naka Fusion Research Establishment, JAERI, Japan
Library, Plasma Physics Laboratory, Kyoto University, Japan
Research Information Center, National Institute for Fusion Science, Japan
Dr. O. Mitarai, Kyushu Tokai University, Japan
Library, Academia Sinica, Institute of Plasma Physics, People's Republic of China
Shih-Tung Tsai, Institute of Physics, Chinese Academy of Sciences, People's Republic of China
Dr. S. Mirnov, TRINITI, Troitsk, Russian Federation, Russia
Dr. V.S. Strelkov, Kurchatov Institute, Russian Federation, Russia
Professor Peter Lukac, Katedra Fyziky Plazmy MFF UK, Mlynska dolina F-2, Komenskeho Univerzita, SK-842 15 Bratislava, Slovakia
Dr. G.S. Lee, Korea Basic Science Institute, South Korea
Mr. Dennis Bruggink, Fusion Library, University of Wisconsin, USA
Institute for Plasma Research, University of Maryland, USA
Librarian, Fusion Energy Division, Oak Ridge National Laboratory, USA
Librarian, Institute of Fusion Studies, University of Texas, USA
Librarian, Magnetic Fusion Program, Lawrence Livermore National Laboratory, USA
Library, General Atomics, USA
Plasma Physics Group, Fusion Energy Research Program, University of California at San Diego, USA
Plasma Physics Library, Columbia University, USA
Alkesh Punjabi, Center for Fusion Research and Training, Hampton University, USA
Dr. W.M. Stacey, Fusion Research Center, Georgia Institute of Technology, USA
Dr. John Willis, U.S. Department of Energy, Office of Fusion Energy Sciences, USA
Mr. Paul H. Wright, Indianapolis, Indiana, USA

The Princeton Plasma Physics Laboratory is operated
by Princeton University under contract
with the U.S. Department of Energy.

Information Services
Princeton Plasma Physics Laboratory
P.O. Box 451
Princeton, NJ 08543

Phone: 609-243-2750
Fax: 609-243-2751
e-mail: pppl_info@pppl.gov
Internet Address: http://www.pppl.gov