Proceedings of ISSS-6 (2001): 1–8 © Copernicus Gesellschaft 2001

ISSS–6

Adaptive Mesh Refinement MHD for Global Simulations

Tamas I. Gombosi[1], Gábor Tóth[1], Darren L. De Zeeuw[1], Kenneth G. Powell[2], and Quentin F. Stout[3]

[1]Department of Atmospheric, Oceanic and Space Sciences, The University of Michigan, Ann Arbor, Michigan, USA
[1]Department of Aeropace Engineering, The University of Michigan, Ann Arbor, Michigan, USA
[3]Department of Electrical Engineering and Computer Science, The University of Michigan, Ann Arbor, Michigan, USA

Abstract. Techniques that have become common in aerodynamics codes have recently begun to be implemented in space-physic codes, which solve the governing equations for a compressible plasma. These techniques include high-resolution upwind schemes, block-based solution-adaptive grids and domain decomposition for parallelization. While some of these techniques carry over relatively straightforwardly from aerodynamics to space physics, space physics simulations pose some new challenges. This paper gives a brief review of the state-of-the-art in modern space-physics codes, including a validation study of several of the techniques in common use. A remaining challenge is that of flows that include regions in which relativistic effects are important; some background and preliminary results for these problems are given.

1 Governing Equations

The governing equations for an ideal, non-relativistic, compressible plasma may be written in a number of different forms. In primitive variables, the governing equations, which represent a combination of the Euler equations of gasdynamics and the Maxwell equations of electromagnetics, may be written as:

$$\begin{aligned}
&\frac{\partial \rho}{\partial t} + \mathbf{u}\cdot\nabla\rho + \rho\nabla\cdot\mathbf{u} = 0\\
&\rho\frac{\partial \mathbf{u}}{\partial t} + \rho\mathbf{u}\cdot\nabla\mathbf{u} + \nabla p - \mathbf{j}\times\mathbf{B} = 0\\
&\frac{\partial \mathbf{B}}{\partial t} + \nabla\times\mathbf{E} = 0\\
&\frac{\partial p}{\partial t} + \mathbf{u}\cdot\nabla p + \gamma p\nabla\cdot\mathbf{u} = 0
\end{aligned} \tag{1}$$

where the current density $\mathbf{j}$ and the electric field vector $\mathbf{E}$ are related to the magnetic field $\mathbf{B}$ by *Ampère's law* and *Ohm's law*, respectively:

$$\mathbf{j} = \frac{1}{\mu_0}\nabla\times\mathbf{B} \tag{2}$$

$$\mathbf{E} = -\mathbf{u}\times\mathbf{B} \tag{3}$$

For one popular class of schemes, the equations are written in a form in which the gasdynamic terms are put in divergence form, and the electromagnetic terms in the momentum and energy equations are treated as source terms. This gives:

$$\begin{aligned}
&\frac{\partial \rho}{\partial t} + \nabla\cdot(\rho\mathbf{u}) = 0\\
&\frac{\partial(\rho\mathbf{u})}{\partial t} + \nabla\cdot(\rho\mathbf{u}\mathbf{u} + p\mathbf{I}) = \mathbf{j}\times\mathbf{B}\\
&\frac{\partial \mathbf{B}}{\partial t} + \nabla\times\mathbf{E} = 0\\
&\frac{\partial E_{gd}}{\partial t} + \nabla\cdot(\mathbf{u}(E_{gd}+p)) = \mathbf{j}\cdot\mathbf{E}
\end{aligned} \tag{4}$$

where E_{gd} is the gasdynamic total energy, given by

$$E_{gd} = \frac{p}{\gamma-1} + \rho\frac{\mathbf{u}\cdot\mathbf{u}}{2} \tag{5}$$

The fully conservative form of the equations is

$$\frac{\partial \mathbf{U}}{\partial t} + (\nabla\cdot\mathbf{F})^{\mathrm{T}} = 0\,, \tag{6}$$

where $\mathbf{U}$ is the vector of conserved quantities

$$\mathbf{U} = \begin{pmatrix} \rho \\ \rho\mathbf{u} \\ B \\ E_{mhd} \end{pmatrix} \tag{7}$$

and $\mathbf{F}$ is a flux diad,

$$\mathbf{F} = \begin{pmatrix} \rho\mathbf{u} \\ \rho\mathbf{u}\mathbf{u} + \left(p + \frac{\mathbf{B}\cdot\mathbf{B}}{2\mu_0}\right)\mathbf{I} - \frac{1}{\mu_0}\mathbf{B}\mathbf{B} \\ \mathbf{u}\mathbf{B} - \mathbf{B}\mathbf{u} \\ \mathbf{u}\left(E_{mhd} + p + \frac{\mathbf{B}\cdot\mathbf{B}}{2\mu_0}\right) - \frac{1}{\mu_0}(\mathbf{u}\cdot\mathbf{B})\,\mathbf{B} \end{pmatrix}^{\mathrm{T}} \tag{8}$$

Correspondence to: Tamas Gombosi (tamas@umich.edu)

where E_{mhd} is the magnetohydrodynamic energy, given by

$$E_{mhd} = \frac{p}{\gamma - 1} + \rho\frac{\mathbf{u}\cdot\mathbf{u}}{2} + \frac{\mathbf{B}\cdot\mathbf{B}}{2\mu_0} \tag{9}$$

Godunov (1972) showed that the fully conservative form, Equation 6, is not symmetrizable. The symmetrizable form may be written as

$$\frac{\partial \mathbf{U}}{\partial t} + (\nabla\cdot\mathbf{F})^{\mathrm{T}} = \mathbf{Q}\ , \tag{10}$$

where

$$\mathbf{Q} = -\nabla\cdot\mathbf{B}\begin{pmatrix} 0 \\ \frac{\mathbf{B}}{\mu_0} \\ \mathbf{u} \\ \frac{\mathbf{u}\cdot\mathbf{B}}{\mu_0} \end{pmatrix} \tag{11}$$

Vinokur (1996) separately showed that Equation 10 can be derived starting from the primitive form, if no stipulation is made about $\nabla\cdot\mathbf{B}$ in the derivation. *Powell* (1994) showed that this symmetrizable form can be used to derive a Roe-type approximate Riemann solver for solving the MHD equations in multiple dimensions.

The MHD eigensystem arising from Equation 6 or Equation 10 leads to eight eigenvalue/eigenvector pairs. The eigenvalues and associated eigenvectors correspond to an entropy wave, two Alfvén waves, two magnetofast waves, two magnetoslow waves, and an eighth eigenvalue/eigenvector pair that depends on which form of the equations is being solved. This last pair has a zero eigenvalue in the fully conservative case, and an eigenvalue equal to that associated with the entropy wave, in the symmetrizable case. The expressions for the eigenvectors, and the scaling of the eigenvectors, are more intricate than in gasdynamics (*Roe and Balsara*, 1996).

2 Solution Techniques

Because the MHD equations are a system of hyperbolic conservation laws, many of the techniques that have been developed for the Euler equations can be applied relatively straightforwardly. In particular, the high-resolution finite-volume approach (*van Leer*, 1979) (i.e. approximate Riemann solver + limited interpolation scheme + multi-stage time-stepping scheme) is perfectly valid. The Rusanov/Lax-Friedrichs approximate Riemann solver can be applied directly; no knowledge of the eigensystem of the MHD equations is required, other than the fastest wave speed in the system. A Roe-type scheme can be constructed, but requires more work, because of the complexity of the eigensystem. In addition, an HLLE-type Riemann solver has been derived by *Linde* (1998); it is less dissipative than the Rusanov/Lax-Friedrichs scheme, but more robust and less computationally intensive than the Roe scheme. Whichever approximate Riemann solver is chosen to serve as the flux function, standard interpolation schemes and limiters can be used to construct a finite-volume scheme.

One added difficulty in solving the MHD equations is that the MHD energy has three components: internal, magnetic and kinetic. Thus, as in gasdynamics, flows with substantially more kinetic energy than internal energy can lead to positivity problems when computing the pressure. Also, in contrast to gasdynamics, regions in which the magnetic field is large can yield similar problems. Conservative and positive HLL-type schemes for MHD have been described by *Janhunen* (2000). Another alternative, due to *Balsara and Spicer* (1999b), is to use a hybrid scheme: both the conservative energy equation and the entropy equations are solved. Close to shock waves the energy equation is used to obtain the correct weak solution, at other places the more robust and positive entropy equation can be used. A variant of this technique has been implemented in our code.

3 Controlling $\nabla\cdot\mathbf{B}$

Another way in which the numerical solution of the MHD equations differs from that of the gasdynamic equations is the constraint that $\nabla\cdot\mathbf{B} = 0$. Enforcing this constraint numerically, particularly in shock-capturing codes, can be done in a number of ways, but each way has its particular strengths and weaknesses. Only a brief overview is given below; each of the schemes discussed below is explained more fully in the references cited, and Tóth has published a numerical comparison of many of the approaches for a suite of test cases (*Tóth*, 2000).

Brackbill and Barnes (1980) first proposed using a Hodge-type projection to the magnetic field. This approach leads to a Poisson equation that must be solved each time the projection takes place:

$$\nabla^2\phi = \nabla\cdot\mathbf{B} \tag{12}$$

$$\mathbf{B}_{projected} = \mathbf{B} - \nabla\phi \tag{13}$$

The resulting projected magnetic field is divergence-free on a particular numerical stencil, to the level of error of the solution of the Poisson equation. While it is not immediately obvious that the use of the projection scheme in conjunction with the fully conservative form of the MHD equations gives the correct weak solutions, Tóth has proven this to be the case (*Tóth*, 2000). The projection scheme has several advantages, including the ability to used standard software libraries for the Poisson solution, its relatively straightforward extension to general unstructured grids, and its robustness. It does, however, require solution of an elliptic equation at each projection step; this can be expensive, particularly on distributed-memory machines.

Powell (*Powell*, 1994; *Powell et al.*, 1999) first proposed an approach based on the symmetrizable form of the MHD equations, Equation 10. In this approach, the source term on the right-hand side of Equation 10 is computed at each time step, and included in the update scheme. Discretizing this form of the equations leads to enhanced stability and accuracy, however, there is no stencil on which the divergence is identically zero. In most regions of the flow, the divergence source term is small. However, near discontinuities, it is not guaranteed to be small. In essence, the inclusion of the

source term changes what would be a zero eigenvalue of the system to one whose value is u_n, the component of velocity normal to the interface through which the flux is computed. The scheme is typically referred to as the eight-wave scheme; the eighth wave is corresponds to propagation of jumps in the normal component of the magnetic field. The eight-wave scheme can be thought of as a hyperbolic or advective approach to controlling $\nabla \cdot \mathbf{B}$; symmetrizable form of the equations, Equation 10, are consistent with the passive advection of $\nabla \cdot \mathbf{B}/\rho$. The eight-wave scheme is computationally inexpensive, easy to add to an existing code, and quite robust. However, if there are regions in the flow in which the $\nabla \cdot \mathbf{B}$ source term (Equation 11) is large, the conservation errors can create problems.

Recently, several approaches have been developed that have combined a Riemann-solver-based scheme with constrained-transport approach. The constrained-transport approach of *Evans and Hawley* (1988) treated the MHD equations in the gasdynamics/electromagnetic-split form of Equation 4. The grid used was a staggered one, and the $\nabla \cdot \mathbf{B} = 0$ constraint was met identically, on a particular numerical stencil. *Dai and Woodward* (1998) and *Balsara and Spicer* (1999a) modified the constrained-transport approach by coupling a Riemann-solver-based scheme for the conservative form of the MHD equations, Equation 6 with a constrained-transport approach for the representation of the magnetic field. In their formulations, this required two representations of the magnetic field: a cell-centered one for the Godunov scheme, and a face-centered one to enforce the $\nabla \cdot \mathbf{B} = 0$ condition. *Tóth* (2000) subsequently showed that these formulations could be recast in terms of a single cell-centered representation for the magnetic field, through a modification to the flux function used. Advantages of the conservative constrained-transport schemes include the fact that they are strictly conservative and that they meet the $\nabla \cdot \mathbf{B} = 0$ constraint to machine accuracy, on a particular stencil. Their primary disadvantage is the difficulty in extending them to general grids. *Tóth and Roe* (2000) made some progress on this front; they developed divergence-preserving prolongation and restriction operators, allowing the use of conservative constrained-transport schemes on h-refined meshes. However, they also showed that the conservative constrained-transport techniques lose their $\nabla \cdot \mathbf{B}$-preserving properties if different cells are advanced at different physical time rates. This rules out the use of local time-stepping. Thus, while for unsteady calculations the cost of the conservative constrained-transport approach is comparable to the eight-wave scheme, for steady-state calculations (where one would typically use local time-stepping), the cost can be prohibitive.

Some of the most recent work on the $\nabla \cdot \mathbf{B} = 0$ constraint has been related to modifying the eight-wave approach by adding a source term proportional to $\nabla(\nabla \cdot \mathbf{B})$ so that the the divergence satisfies an advection-diffusion equation, rather than a pure advection equation. This technique, due to *Linde and Malagoli* (2000) referred to as diffusive control of $\nabla \cdot \mathbf{B}$, has the same advantages and disadvantages as the eight-wave approach. It is not strictly conservative, but appears to keep the level of $\nabla \cdot \mathbf{B}$ lower than the eight-wave approach does. In other recent work by *Dedner et al.* (2001), a generalized Lagrange-multiplier method has been proposed, incorporating the projection approach, the eight-wave approach, and the diffusive-control approach into a single framework.

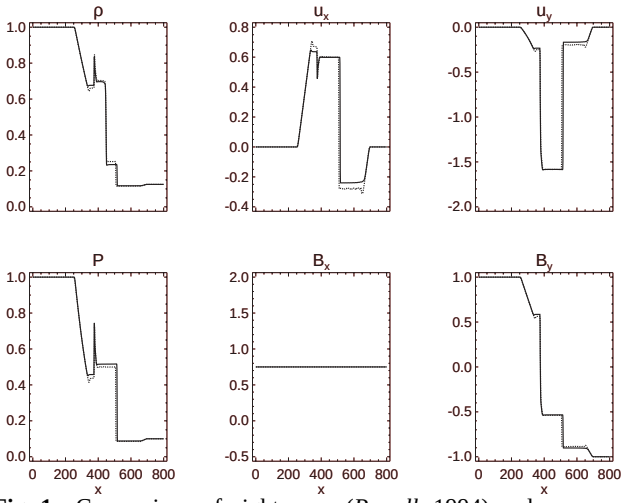

Fig. 1. Comparison of eight-wave (*Powell*, 1994) and conservative constrained transport (*Balsara and Spicer*, 1999b) schemes (solid line) with non-conservative constrained transport (*Evans and Hawley*, 1988) scheme (dotted line) for Brio-Wu plasma-shock-tube problem

4 Validation Studies

In this section, validation studies are presented that compare the solution methods and $\nabla \cdot \mathbf{B}$ control techiques cited above.

The first test cases are plasma-shock-tube problems. In Figure 1, the results of a one-dimensional plasma-shock-tube problem known as the *Brio and Wu* (1988) problem are presented for three schemes: the eight-wave scheme, based on the symmetrizable form of the equations, the conservative constrained transport scheme, based on the fully conservative form of the equations, and the non-conservative constrained transport scheme, based on the gasdynamic/electromagnetic split form of the equations. For this problem, the results of the eight-wave scheme (*Powell*, 1994) and the conservative constrained transport scheme (*Balsara and Spicer*, 1999b) are indistinguishable, and are shown by the solid line. The non-conservative constrained transport scheme (*Evans and Hawley*, 1988) results, shown by the dotted line, display errors as large as 20%, particularly in the velocity. This is not surprising, of course; the various jumps in the Brio-Wu problem correspond to the Rankine-Hugoniot conditions *for a plasma*, which differ from those of a gasdynamic shock. Because the equations are not discretized in a divergence form, substantial errors are expected in the presence of non-zero magnetic fields.

One-dimensional plasma-shock-tube problems such as the Brio-Wu problem are popular validation cases for base schemes, but do not test the $\nabla \cdot \mathbf{B} = 0$ constraint techniques. This is

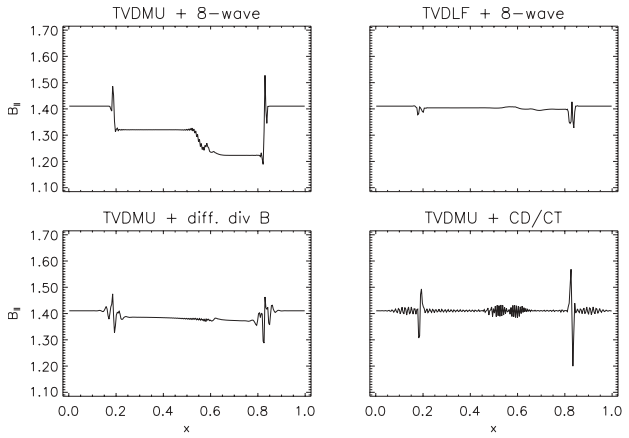

Fig. 2. Comparison of Various $\nabla \cdot \mathbf{B}$ Control Techniques on Rotated Plasma-Shock-Tube Problem

because, in one dimension, the constraint that $\nabla \cdot \mathbf{B} = 0$ degenerates to the constraint that $B_x = \text{constant}$. However, rotating a one-dimensional problem so that the discontinuities run oblique to the grid yields a problem that can test the $\nabla \cdot \mathbf{B}$ constraint technique. In Figure 2, results are presented for the component of the magnetic field parallel to the direction of motion of the waves (i.e. the analog to B_x in the one-dimensional case). The exact solution for this quantity is a constant ($B_\parallel = \sqrt{2}$); the numerical results differ depending on the $\nabla \cdot \mathbf{B}$ constraint technique. The largest error, on the order of 10%, comes from using the eight-wave scheme in conjunction with the Roe approximate Riemann solver. This is due to the $\nabla \cdot \mathbf{B}$ source term, which is not small in the region of the fast magnetosonic shock in this case, and leads to conservation errors. Surprisingly, the eight-wave scheme in conjunction with the Rusanov/Lax-Friedrichs approximate Riemann solver yields errors that are an order of magnitude smaller. The diffusive $\nabla \cdot \mathbf{B}$ control technique, used in conjunction with the Roe approximate Riemann solver, yields errors on the order of 2%; the conservative constrained-transport technique, yields results that are centered on the correct value, but somewhat oscillatory. It should be noted that only $B_\parallel$ is shown here, in part because the errors in other variables are much smaller: the differences in $B_\perp$, pressure, density, and $u_\parallel$ among the schemes are two orders of magnitude smaller than those in $B_\parallel$; the differences in $u_\perp$ among the schemes are more than one order of magnitude smaller than those in $B_\parallel$.

The third validation case is one that is more representative of space-physics calculations. It represents a quasi-steady interaction of the solar wind with Earth's magnetic field. The boundary condition upstream of Earth is a steady plasma flow, with: a density of 5 molecules per cubic centimeter, a temperature of 180,000 K, a velocity of 400 kilometers per second, pointed directly outward from the Sun, and a magnetic field of 5 nanoTesla, pointed northward. Earth (including its atmosphere through the ionosphere) is represented as a conducting sphere with an embedded, non-tilted, non-rotating magnetic dipole. These conditions are a simplification of the real situation, in which the flow from the Sun would be unsteady, and the Earth's intrinsic magnetic field more complicated. The calculation is carried out on a three-dimensional, solution-adaptive grid, on a parallel machine. Details of the approach are given by *Groth et al.* (1999). The code can be run first-order or as a second-order MUSCL scheme, using any combination of the solvers and $\nabla \cdot \mathbf{B}$ control techniques mentioned above.

Figure 3 shows the effect of $\nabla \cdot \mathbf{B}$ control technique. Four methods — the eight-wave, diffusive control, projection and conservative constrained transport techniques — are compared, using a second-order MUSCL scheme with a Rusanov solver, and a grid with a smallest cell size of a quarter R_E. Although the grid for this case is relatively coarse, the various $\nabla \cdot \mathbf{B}$ control techniques lead to results that differ by only 1–2%. The relative cost depends on implementation, but the eight-wave and diffusive-control techniques are the least expensive, the projection scheme somewhat more (because of the elliptic step each time the magnetic field is projected) and the constrained transport is substantially more expensive (approximately a factor seven over the eight-wave scheme) because of the inability to use local time-stepping in this steady problem.

5 Semi-Relativistic Plasmas

While the solar-wind speed remains non-relativistic in the solar system, the intrinsic magnetic fields of several planets in the solar system are high enough, and the density of the solar wind low enough, that the Alfvén speed,

$$V_A = \sqrt{\frac{\mathbf{B} \cdot \mathbf{B}}{\mu_0 \rho}} \tag{14}$$

can reach appreciable fractions of the speed of light. In the case of Jupiter, the Alfvén speed in the vicinity of the poles is of order ten! Even Earth has a strong enough intrinsic magnetic field that the Alfvén speed reaches twice the speed of light in Earth's near-auroral regions.

For these vicinities, solving the non-relativistic ideal MHD equations does not make sense. Having waves in the system propagating faster than the speed of light, besides being non-physical, causes a number of numerical difficulties. However, solving the fully relativistic MHD equations is overkill. What is called for is a semi-relativistic form of the equations, in which the flow speed and acoustic speed are non-relativistic, but the Alfvén speed can be relativistic. A derivation of these semi-relativistic equations from the fully relativistic equations is given in *Gombosi et al.* (2001); the final result is presented here.

The semi-relativistic ideal MHD equations are of the form

$$\frac{\partial \mathbf{U}_{sr}}{\partial t} + (\nabla \cdot \mathbf{F}_{sr})^{\mathrm{T}} = 0 \tag{15}$$

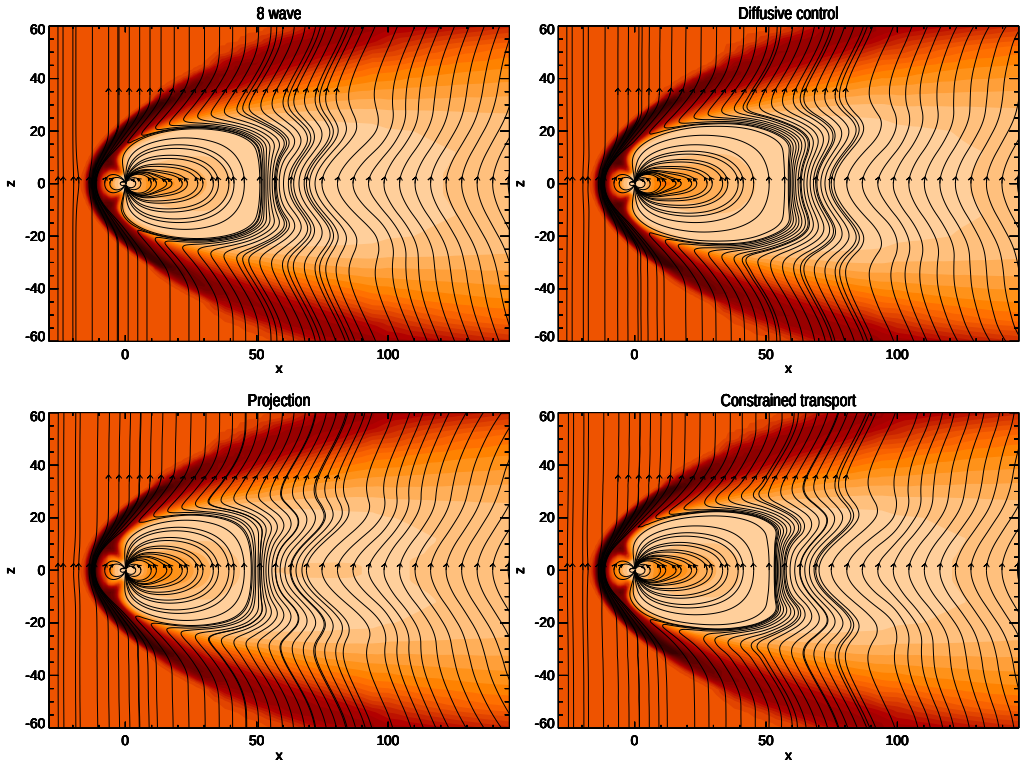

Fig. 3. Comparison of $\nabla \cdot \mathbf{B}$ Control Techniques for Earth's Magnetosphere Simulation

where the state vector, $\mathbf{U}_{sr}$, and the flux diad, $\mathbf{F}_{sr}$, are

$$\mathbf{U}_{sr} = \begin{pmatrix} \rho \\ \rho\mathbf{u} + \frac{1}{c^2}\mathbf{S}_{\mathrm{A}} \\ \mathbf{B} \\ \frac{\rho\mathbf{u}\cdot\mathbf{u}}{2} + \frac{p}{\gamma-1} + e_{\mathrm{A}} \end{pmatrix} \tag{16}$$

$$\mathbf{F}_{sr} = \begin{pmatrix} \rho\mathbf{u} \\ \rho\mathbf{u}\mathbf{u} + p\mathbf{I} + \mathbf{P}_{\mathrm{A}} \\ \mathbf{u}\,\mathbf{B} - \mathbf{B}\,\mathbf{u} \\ \left(\frac{\rho\mathbf{u}\cdot\mathbf{u}}{2} + \frac{\gamma p}{\gamma-1}\right)\mathbf{u} + \mathbf{S}_{\mathrm{A}} \end{pmatrix}^{\mathrm{T}} \tag{17}$$

In the above,

$$\begin{aligned} \mathbf{S}_{\mathrm{A}} &= \frac{1}{\mu_0}(\mathbf{E}\times\mathbf{B}) \\ e_{\mathrm{A}} &= \frac{1}{2\mu_0}\left(B^2 + \frac{1}{c^2}E^2\right) \\ \mathbf{P}_{\mathrm{A}} &= e_{\mathrm{A}}\mathbf{I} - \frac{1}{\mu_0}\mathbf{B}\,\mathbf{B} - \frac{1}{\mu_0 c^2}\mathbf{E}\,\mathbf{E} \end{aligned}$$

are the Poynting vector, the electromagnetic energy density, and the electromagnetic pressure tensor, respectively. The electric field $\mathbf{E}$ is related to the magnetic field $\mathbf{B}$ by Ohm's law, Equation 3.

This new system of equations has wave speeds that are limited by the speed of light; for strong magnetic fields, the modified Alfvén speed (and the modified magnetofast speed) asymptote to c. The modified magnetoslow speed asymptotes to a, the acoustic speed. This property offers the possibility of a rather tricky convergence-acceleration technique, first suggested by *Boris* (1970); the wave speeds can be lowered, and the stable time-step thereby raised, by artificially lowering the value taken for the speed of light.

The equations above are valide in physical situations in which $V_A > c$. A slight modification yields a set of equations, the steady-state solutions of which are independent of the value taken for the speed of light. Defining the true value of the speed of light to be c_0, to distinguish it from the artificially lowered speed of light, c, the equations are:

$$\frac{\partial \mathbf{U}_{sr}}{\partial t} + (\nabla \cdot \mathbf{F}_{sr})^{\mathrm{T}} = \mathbf{Q}_{c_0} \tag{18}$$

where the state vector, $\mathbf{U}_{sr}$, and the flux diad, $\mathbf{F}_{sr}$, are as defined above, and the new source term is

$$\mathbf{Q}_{c_0} = \frac{1}{\mu_0}\left(\frac{1}{c_0^2} - \frac{1}{c^2}\right)\mathbf{E}\nabla\cdot\mathbf{E} \tag{19}$$

An implementation of the semi-relativistic equations has been made. It is based on the Rusanov/Lax-Friedrichs approximate Riemann solver; the Roe scheme for the semi-relativistic equations would be quite a mess, due to the complicated expressions for the eigenvalues and eigenvectors. The eight-wave scheme is used to control $\nabla \cdot \mathbf{B}$.

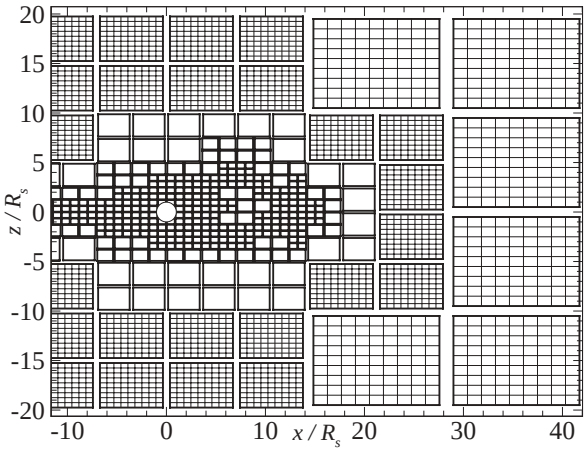

Fig. 4. Adapted grid in the simulation of a coronal mass ejection. The Figure shows a 2D cross section of the 3D grid structure (cutting through the CME).

6 Block-Based AMR on Cartesian Mesh

Keeping in mind the desire for high performance on massively parallel computer architectures, a block-based adaptive mesh refinement (AMR) technique is used in our code (called BATS-R-US). The governing equations are integrated to obtain volume-averaged solution quantities within rectangular Cartesian computational cells. The cells are embedded in regular structured blocks of equal sized cells. The blocks are all self-similar. Although each block occupies the same amount of space in memory, the blocks may occupy different sized volumes in physical space.

The computational grid is composed of many self-similar blocks. In regions that require increased resolution, a block is refined by dividing it into eight identical octants. In regions that are deemed over-resolved, the refinement process is reversed and eight blocks are coarsened and coalesced into a single block. Multiple physics-based refinement criteria are used to direct the coarsening and division of blocks. A hierarchical tree-like data structure is used to keep track of mesh refinement and the connectivity between solution blocks (*Stout et al.*, 1997; *Powell et al.*, 1999).

An example of a 2-D cut through a 3-D grid, taken from a calculation of a coronal mass ejection, is shown in Figure 4 (*Groth et al.*, 2000). Grids like those shown in Figure 4 go a long way towards resolving the disparate scales in a problem. Each level of refinement in the grid introduces cells that are smaller by a factor two in each dimension from those one level higher in the grid. Typical calculations have 10-15 levels of refinement; some calculations have more than 20 levels of refinement. In the case of 20 levels of refinement, the finest cells on the mesh are more than one million times smaller in each dimension than the coarsest cells on a mesh.

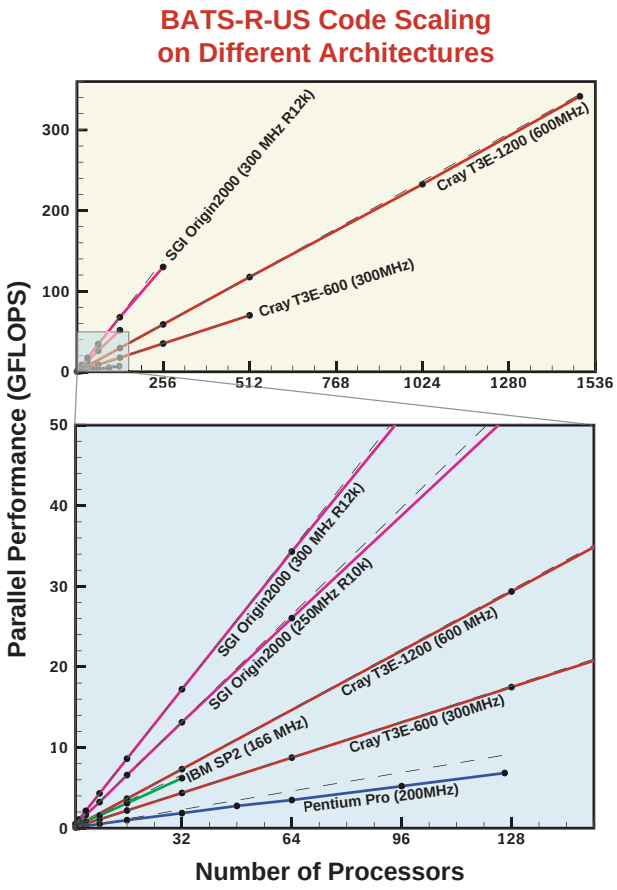

Fig. 5. Parallel scaling of BATS-R-US on various architectures. Black dashed lines represent perfect scaling from single node performance.

7 Parallel Implementation

The parallel block-based AMR solver was designed from the ground up with a view to achieving very high performance on massively parallel architectures. The underlying upwind finite-volume solution algorithm, with explicit time stepping, has a very compact stencil and is therefore highly local in nature. The hierarchical data structure and self-similar blocks make domain decomposition of the problem almost trivial and readily enable good load-balancing, a crucial element for truly scalable computing. A natural load balancing is accomplished by simply distributing the blocks equally amongst the processors. The self-similar nature of the solution blocks also means that serial performance enhancements apply to all blocks and that fine grain parallelization of the algorithm is possible. The parallel implementation of the algorithm has been carried out to such an extent, that even the grid adaptation is performed in parallel.

Other features of the parallel implementation include the use of FORTRAN 90 as the programming language and the message passing interface (MPI) library for performing the interprocessor communication. Use of these standards greatly enhances the portability of the code and leads to very good serial and parallel performance. The message passing is per-

formed in an asynchronous fashion with gathered wait states and message consolidation such that it typically accounts for less than 3-5% of processor time.

Implementation of the algorithm has been carried out on Cray T3E supercomputers, SGI and Sun workstations, on Beowulf type PC clusters, on SGI shared-memory machines, on a Cray T3D, and on several IBM SP2s. BATS-R-US nearly perfectly scales to 1,500 processors and a sustained speed of 342 GFlops has been attained on a Cray T3E-1200 using 1,490 PEs. For each target architecture, simple single-processor measurements are used to set the size of the adaptive blocks. The scaling of BATS-R-US on various architectures is shown in Figure 5.

8 Time to Solution

Since a major goal of global space plasma simulations is the creation of a predictive space weather tool, wallclock time to solution is a paramount issue. In particular, a predictive model must run substantially faster than real time. From the starting point – the observation of a solar event, to the ending point – post-processing the data from a simulation based on the initial conditions derived from the observations, a simulation must be accomplished rapidly to be of use.

The main limitation of the present generation of global space plasma codes is the explicit time stepping algorithm. Explicit time steps are limited by the Courant-Friedrichs-Lewy (CFL) condition, which essentially ensures that no information travels more than a cell size during a time step. This condition represents a non-linear penalty for highly resolved calculations, since finer grid resolution not only results in more computational cells, but also in smaller time steps.

In global MHD simulations of space plasmas the CFL condition is controlled by two factors: (1) the smallest cell size in the simulation, and (2) the fast magnetosonic speed in high magnetic field, low plasma density regions. In a typical magnetosphere simulation with a smallest cell size of about 0.25 $R_{\rm E}$ the CFL condition limits the time step to about 10^{-2} s. This small step is primarily controlled by the high fast magnetosonic speed (due to the high Alfvén speed) in the near-Earth region.

There are several ways to increase the time step in an MHD simulation and thus decrease the time to solution. Here we discuss only two of the potential solutions: the so-called "*Boris* correction" (discussed in section 5), and implicit time stepping.

In BATSRUS we have a number of time stepping algorithms implemented. The simplest and least expensive scheme is a multistage explicit time stepping, for which the time step is limited by the CFL stability condition. We have also implemented an unconditionally stable fully implicit time stepping scheme based on (**??**). The second order implicit time discretization (BDF2) requires the solution of a non-linear system of equations for all the flow variables. This can be achived by the Newton-Krylov-Schwarz approach: a Newton iteration is applied to the non-linear equations; a parallel Krylov type iterative scheme is used to solve the linear systems; the convergence of the Krylov solver is accelerated with a Schwarz type preconditioning. In our implementation the Krylov solver is BiCGSTAB, and a modified block incomplete LU (MBILU) preconditioner is applied on a block by block basis. Since every block has a simple Cartesian geometry, the preconditioner can be implemented very efficiently. The resulting implicit scheme requires about 20-30 times more CPU time per time step than the explicit method, but the physical time step can be 1000 to 10000 times larger. This implicit algorithm has a very good parallel scaling due to the matrix free evaluation of the Jacobian and the block by block application of the preconditioner.

In BATS-R-US, we combine explicit and implicit time stepping. Magnetosphere simulations include large volumes where the Alfvén speed is quite low (tens of km/s) and the local CFL number allows large explicit time steps (tens of seconds to several minutes). In these regions implicit time stepping is a waste of computational resources. Since the parallel implicit technique we use is fundamentally block based we only treat those blocks implicitly where the CFL condition would limit the explicit time step to less than the selected time step (typically ~ 10 s). Needless to say, this combined explicit-implicit time stepping represents more computational challenges (such as separate balancing of explicit and implicit blocks). Overall, this solution seems to be a very promising option, but other potential avenues need to explored before one makes a final decision about the most efficient time-stepping algorithm for space MHD simulations.

9 Applications

BATS-R-US has been extensively applied to global numerical simulations of the Sun-Earth system (*Gombosi et al.*, 2000; *Groth et al.*, 2000), the coupled terrestrial magnetosphere-ionosphere (*Gombosi et al.*, 1998; *Song et al.*, 1999, 2000), and the interaction of the heliosphere with the interstellar medium (*Linde et al.*, 1998). In addition, it has also been successfully applied to a host of planetary problems ranging from comets (*Gombosi et al.*, 1996; *Häberli et al.*, 1997), to Mercury (*Kabin et al.*, 2000), Venus (*Bauske et al.*, 1998), Mars (*Liu et al.*, 1999), Saturn (*Hansen et al.*, 2000), to planetary satellites (*Kabin et al.*, 1999, 2001).

Acknowledgements. This work was supported by DoD MURI grant F49620-01-1-0359, NSF KDI grant NSF ATM-9980078, NSF CISE grant ACI-9876943, and NASA IASRP grant NAG5-9406. G. Tóth is partially supported by a postdoctoral fellowship (D 25519) from the Hungarian Science Foundation (OTKA).

References

Balsara, D. S., and D. S. Spicer, A staggered mesh algorithm using high order Godunov fluxes to ensure solenoidal magnetic fields in magnetohydrodynamic simulations, *J. Comput. Phys.*, *149*, 270–292, 1999a.

Balsara, D. S., and D. S. Spicer, Maintaining pressure positivity in magnetohydrodynamic simulations, *J. Comput. Phys.*, *148*, 133–148, 1999b.

Bauske, R., A. F. Nagy, T. I. Gombosi, D. L. De Zeeuw, K. G. Powell, and J. G. Luhmann, A three-dimensional MHD study of solar wind mass loading processes at Venus: Effects of photoionization, electron impact ionization, and charge exchange, *J. Geophys. Res.*, *103*(A10), 23,625–23,638, 1998.

Boris, J. P., A physically motivated solution of the Alfvén problem, Tech. Rep. NRL Memorandum Report 2167, Naval Research Laboratory, Washington, D.C., 1970.

Brackbill, J., and D. Barnes, The effect of nonzero $\nabla \cdot \mathbf{B}$ on the numerical solution of the magnetohydrodynamic equations, *J. Comput. Phys.*, *35*, 426–430, 1980.

Brio, M., and C. C. Wu, An upwind differencing scheme for the equations of ideal magnetohydrodynamics, *J. Comput. Phys.*, *75*, 400–422, 1988.

Dai, W., and P. R. Woodward, A simple finite difference scheme for multidimensional magnetohydrodynamic equations, *J. Comput. Phys.*, *142*, 331, 1998.

Dedner, A., F. Kemm, D. Kröner, C.-D. Munz, T. Schnitzer, and M. Wesenberg, Hyperbolic divergence cleaning for the MHD equations, *J. Comput. Phys.*, 2001, submitted.

Evans, C. R., and J. F. Hawley, Simulation of magnetohydrodynamic flows: A constrained transport method, *Astrophysical Journal*, *332*, 659–677, 1988.

Godunov, S. K., Symmetric form of the equations of magnetohydrodynamics (in Russian), in *Numerical Methods for Mechanics of Continuum Medium*, vol. 1, pp. 26–34, Siberian Branch of USSR Acad. of Sci., 1972.

Gombosi, T. I., D. L. De Zeeuw, R. M. Häberli, and K. G. Powell, Three-dimensional multiscale MHD model of cometary plasma environments, *J. Geophys. Res.*, *101*(A7), 15233–15253, 1996.

Gombosi, T. I., D. L. De Zeeuw, C. P. T. Groth, K. G. Powell, and P. Song, The length of the magnetotail for northward IMF: Results of 3D MHD simulations, in *Physics of Space Plasmas*, edited by T. Chang, and J. R. Jasperse, vol. 15, pp. 121–128, MIT Press, Cambridge, Mass., 1998.

Gombosi, T. I., D. L. De Zeeuw, C. P. T. Groth, K. G. Powell, and Q. F. Stout, Multiscale MHD simulation of a coronal mass ejection and its interaction with the magnetosphere-ionosphere system, *J. Atmos. Solar-Terr. Phys.*, *62*, 1515–1525, 2000.

Gombosi, T. I., G. Tóth, D. L. D. Zeeuw, K. C. Hansen, K. Kabin, and K. G. Powell, Semi-relativistic magnetohydrodynamics and physics-based convergence acceleration, *J. Comput. Phys.*, 2001, submitted.

Groth, C. P. T., D. L. De Zeeuw, K. G. Powell, T. I. Gombosi, and Q. F. Stout, A parallel solution-adaptive scheme for ideal magnetohydrodynamics, in *Proc. 14th AIAA Computational Fluid Dynamics Conference*, Norfolk, Virginia, AIAA Paper No. 99-3273, 1999.

Groth, C. P. T., D. L. De Zeeuw, T. I. Gombosi, and K. G. Powell, Global 3D MHD simulation of a space weather event: CME formation, interplanetary propagation, and interaction with the magnetosphere, *J. Geophys. Res.*, *105*(A11), 25,053 – 25,078, 2000.

Häberli, R. M., T. I. Gombosi, D. L. DeZeuuw, M. R. Combi, and K. G. Powell, Modeling of cometary x-rays caused by solar wind minor ions, *Science*, *276*, 939–942, 1997.

Hansen, K. C., T. I. Gombosi, D. L. DeZeeuw, C. P. T. Groth, and K. G. Powell, A 3D global MHD simulation of Saturn's magnetosphere, *Adv. Space Res.*, *26*(10), 1681–1690, 2000.

Janhunen, P., A positive conservative method for magnetohydrodynamics based on HLL and roe methods, *J. Comput. Phys.*, *160*, 649–661, 2000.

Kabin, K., T. I. Gombosi, D. L. DeZeeuw, K. G. Powell, and P. L. Israelevich, Interaction of the Saturnian magnetosphere with Titan: Results of a 3D MHD simulation, *J. Geophys. Res.*, *104*(A2), 2451–2458, 1999.

Kabin, K., T. I. Gombosi, D. L. De Zeeuw, and K. G. Powell, Interaction of Mercury with the solar wind, *Icarus*, *143*, 397–406, 2000.

Kabin, K., M. R. Combi, T. I. Gombosi, D. L. De Zeeuw, K. C. Hansen, and K. G. Powell, Io's magnetospheric interaction: an MHD model with day-night asymmetry, *PSS*, *49*, 337–344, 2001.

Keppens, R., G. Tóth, M. A. Botchev, and A. van der Ploeg, Implicit and Semi-Implicit Schemes: algorithms, *Int. J. for Num. Meth. in Fluids 30*, 335–352, 1999.

Linde, T. J., A three-dimensional adaptive multifluid MHD model of the heliosphere, Ph.D. thesis, Univ. of Mich., Ann Arbor, 1998.

Linde, T., and A. Malagoli, On local $\nabla \cdot \mathbf{B}$ control in ideal MHD simulations, *J. Comput. Phys.*, 2000, submitted.

Linde, T. J., T. I. Gombosi, P. L. Roe, K. G. Powell, and D. L. De Zeeuw, The heliosphere in the magnetized local interstellar medium: Results of a 3D MHD simulation, *J. Geophys. Res.*, *103*(A2), 1889–1904, 1998.

Liu, Y., A. F. Nagy, C. P. T. Groth, D. L. De Zeeuw, T. I. Gombosi, and K. G. Powell, 3D multi-fluid MHD studies of the solar wind interaction with Mars, *Geophys. Res. Lett.*, *26*(17), 2689–2692, 1999.

Powell, K. G., An approximate Riemann solver for magnetohydrodynamics (that works in more than one dimension), Tech. Rep. 94-24, Inst. for Comput. Appl. in Sci. and Eng., NASA Langley Space Flight Center, Hampton, Va., 1994.

Powell, K. G., P. L. Roe, T. J. Linde, T. I. Gombosi, and D. L. D. Zeeuw, A solution-adaptive upwind scheme for ideal magnetohydrodynamics, *J. Comput. Phys.*, *154*(2), 284–309, 1999.

Roe, P. L., and D. S. Balsara, Notes on the eigensystem of magnetohydrodynamics, *SIAM J. Appl. Math.*, *56*(1), 57–67, 1996.

Song, P., D. L. De Zeeuw, T. I. Gombosi, C. P. T. Groth, and K. G. Powell, A numerical study of solar wind–magnetosphere interaction for northward IMF, *J. Geophys. Res.*, *104*(A12), 28,361–28,378, 1999.

Song, P., T. Gombosi, D. De Zeeuw, K. Powell, and C. P. T. Groth, A model of solar wind - magnetosphere - ionosphere coupling for due northward IMF, *Planet. Space Sci.*, *48*, 29–39, 2000.

Stout, Q. F., D. L. De Zeeuw, T. I. Gombosi, C. P. T. Groth, H. G. Marshall, and K. G. Powell, Adaptive blocks: A high-performance data structure, in *Proc. Supercomputing'97*, 1997.

Tóth, G., R. Keppens, and M. A. Botchev, Implicit and semi-implicit schemes in the Versatile Advection Code: numerical tests, Astron. & Astrophys. *332*, 1159–1170, 1998.

Tóth, G., The $\nabla \cdot \mathbf{B}$ constraint in shock capturing magnetohydrodynamic codes, *J. Comput. Phys.*, *161*, 605–652, 2000.

Tóth, G., and P. L. Roe, Divergence- and curl-preserving prolongation and restriction formulae, *J. Comput. Phys.*, 2000, submitted.

van Leer, B., Towards the ultimate conservative difference scheme. V. A second-order sequel to Godunov's method, *J. Comput. Phys.*, *32*, 101–136, 1979.

Vinokur, M., A rigorous derivation of the MHD equations based only on Faraday's and Ampére's laws, Presentation at LANL MHD Workshop, 1996.