"Part of the inhumanity of the computer is that, once it is competently programmed and working smoothly, it is completely honest."

ISAAC ASIMOV

5

Multidimensional Smoothed Particle Magnetohydrodynamics

5.1 Introduction

In more than one spatial dimension errors associated with the non-zero divergence of the magnetic field need to be taken into account in any numerical MHD scheme. There are two distinct issues to be addressed. The first is the treatment of terms proportional to $\nabla \cdot \mathbf{B}$ in the MHD equations (in particular in the formulation of the induction equation and the magnetic force). The second is the maintenance of the $\nabla \cdot \mathbf{B} = 0$ constraint. Note that a solution to the latter problem does not necessarily resolve the former, since maintaining $\nabla \cdot \mathbf{B} = 0$ in a particular numerical discretisation does not guarantee that it is zero in all discretisations.

Perhaps the first to address these issues in a numerical context were Brackbill and Barnes (1980), where it was noted that using a conservative formulation of the magnetic force could cause a supposed steady state setup to change because of the small but non-zero component of magnetic force directed along the field lines due to the monopole term. This error can have serious consequences even though the proportional error in the magnetic field is small. As discussed in §4.4, in SPMHD the force parallel to the field can have catastrophic consequences, leading to numerical instability under some circumstances. Brackbill and Barnes (1980) approached the problem by preferring a non-conservative formulation of the momentum equation which guarantees that the magnetic force is exactly perpendicular to the field. Such an approach has also been used successfully in an SPMHD context by several authors (e.g. Benz, 1984; Meglicki et al., 1995; Byleveld and Pongracic, 1996; Cerqueira and de Gouveia Dal Pino, 2001), however the numerical simulations of shocks seems to require the exact conservation of momentum in order to provide the correct jump conditions at shock fronts (which means, at the very least, the discrete formulation should be based on continuum equations which conserve momentum exactly even with a

non-zero magnetic divergence[1]). This issue of neglect or inclusion of divergence terms extends further to the formulation of the induction and energy equations. In the formulation used by Brackbill and Barnes (1980), magnetic flux and energy are conserved exactly, but the conservation of momentum is sacrificed. More recently, this question has been re-addressed by Powell et al. (1999) in the light of the second issue, namely how best to maintain the divergence constraint without resorting to expensive divergence cleaning procedures. The approach taken by Powell et al. (1999) was to include source terms in the equations which allow the divergence errors to be propagated appropriately by the flow. In the Powell et al. (1999) approach, momentum, energy and magnetic flux conservation (in a volume sense) are sacrificed, although it seems that this does not have too severe consequences with the '8-wave' Riemann solver Powell et al. used for the simulation of shocks (however we do not find this to be the case in SPMHD). The equation set used by Powell et al. (1999) and its effect on the propagation of divergence errors is discussed below (§5.2.1). More recently, however, it has been pointed out by both Janhunen (2000) (by considering the presence of monopoles in Maxwell's equations) and Dellar (2001) (from relativistic considerations) that a consistent formulation of the MHD equations in the presence of magnetic monopoles should retain both the conservation of momentum and energy. In §4.3.2 we were able to verify that the set of MHD equations derived by Janhunen (2000) and Dellar (2001) indeed form a consistent set by deriving the SPMHD equations from a variational principle which uses the SPH form of the induction equation as a constraint in order to derive the momentum and energy equations. Similarly it can also be shown that the formalism used by Brackbill and Barnes (1980), in which the conservation of flux is retained but the conservation of momentum and energy are not, is also consistent (although undesirable due to both the non-conservation and the effects on the propagation of divergence errors). Furthermore the derivation given in §4.3.2 was for the discrete SPMHD equations, ensuring consistency in both the continuum and discrete forms. This consistent set of equations and the consequences for the propagation of divergence errors has already been discussed (albeit briefly) in Chapter 4. Further discussion and comparison with the Powell et al. (1999) approach is given in §5.2.1 and examined numerically in §5.3.2.

Many other approaches to the second issue are also possible. Maintenance of constraints similar to the divergence-free condition for the magnetic field is important not only for MHD problems, but also for incompressible flows (where $\nabla \cdot \mathbf{v} = 0$) and especially in algorithms for numerical relativity, since Einstein's equations can be written in a form corresponding closely to the Maxwell equations[2]. Many possible methods have been proposed for dealing with this problem, each with their own advantages and disadvantages. Perhaps the simplest is to explicitly evolve a vector potential $\mathbf{A}$, from which the magnetic field is derived by taking the curl, guaranteeing that the divergence is zero. The major disadvantage of this approach is that the computation of the force terms involves second derivatives of the evolved variable ($\mathbf{A}$), which in general can be significantly less accurate. One advantage of using the vector potential is that the conservation of magnetic helicity $\mathbf{A} \cdot \mathbf{B}$ can be monitored (§4.2.2), which is particularly important for dynamo and reconnection problems often encountered in Solar physics (e.g. Brandenburg, 2001).

Brackbill and Barnes (1980) proposed a simple projection scheme to 'clean up' the magnetic field at

[1] For example, none of the results obtained on the shock tube tests given in §4.6.3 could be obtained to *any* degree of satisfaction using a formalism based on a non-conservative momentum equation (such as those given in §4.3.5), although the formalism proposed by Morris (1996) (§4.4.2) can be made to give reasonable results since firstly it is based on a conservative form of the continuum equations and secondly at least conserves momentum exactly for isotropic forces

[2] In the case of the Einstein equations, there are six evolution equations and four constraint equations, similar to the two evolution equations and two constraint equations in the Maxwell equations.

each timestep, an approach which is now commonly used in many grid-based MHD codes (e.g. Balsara, 1998). Similar schemes have been used in SPH for incompressible flows (e.g. Cummins and Rudman, 1999). The disadvantage of this approach is that it involves the solution of a Poisson equation which is computationally expensive. Another approach used in grid-based MHD codes is the so-called 'constrained transport' method pioneered by Evans and Hawley (1988) in which differences of the magnetic field across the grid cell are constructed in such a way as to maintain the divergence free condition exactly. Such methods work very well, however cannot be used in SPH because of the absence of a spatial grid (although perhaps some divergence-free interpolation could be devised). A comparison between several of these schemes with the source term approach of Powell et al. (1999) and a projection method for finite difference codes has been recently presented by Tóth (2000). Although not all of the schemes are applicable in an SPH context, many of the numerical tests presented in this chapter are taken from this paper. More recently Dedner et al. (2002) have proposed a method for cleaning the magnetic field which is significantly faster than the projection method by explicitly adding a constraint propagation equation which is coupled to the evolution equation for the magnetic field. This equation propagates the divergence error in a hyperbolic (ie. wave-like) manner away from its source. Adding a small diffusion term means that the error is rapidly reduced to zero.

In §5.2 we investigate several of these approaches to maintaining the $\nabla \cdot \mathbf{B} = 0$ constraint which are applicable in an SPH context, namely the source term approach discussed in the previous chapter (§5.2.1), projection methods (§5.2.2) and the Dedner et al. approach (§5.2.3). The algorithm is then benchmarked, as in the one dimensional case, against a wide range of multidimensional test problems used to test recent grid-based MHD codes (§5.3). The tests involve the propagation of an initially non-zero magnetic divergence (§5.3.2), nonlinear Alfvén waves (§5.3.3), two dimensional shock tubes (§5.3.5), interacting shocks and the transition to turbulence (§5.3.7) and two dimensional spherically symmetric blast waves (§5.3.6).

5.2 Divergence correction techniques

5.2.1 Source term approach

As discussed in §4.2.1 the induction equation can be written in the 'conservative' form

$$\frac{\partial \mathbf{B}}{\partial t} = -\nabla \times (\mathbf{v} \times \mathbf{B}), \tag{5.1}$$

$$= \nabla \cdot (\mathbf{vB} - \mathbf{Bv}). \tag{5.2}$$

which explicitly conserves the volume integral of the flux (4.17). In Lagrangian form this is given by

$$\frac{d\mathbf{B}}{dt} = -\mathbf{B}(\nabla \cdot \mathbf{v}) + (\mathbf{B} \cdot \nabla)\mathbf{v} + \mathbf{v}(\nabla \cdot \mathbf{B}) \tag{5.3}$$

Taking the divergence of this equation, we have

$$\frac{\partial}{\partial t}(\nabla \cdot \mathbf{B}) = 0, \tag{5.4}$$

showing that the constraint $\nabla \cdot \mathbf{B} = 0$ enters the MHD equations as an initial condition. However allowing magnetic monopoles resulting from $\nabla \cdot \mathbf{B} \neq 0$ to evolve appropriately within the flow can prevent the build up of unphysical numerical effects associated with their presence and can therefore reduce the need for computationally expensive divergence cleaning procedures. Thus Powell (1994) (see Powell et al. 1999) suggested that the conservative forms of the MHD equations should contain source terms to ensure that these errors are propagated out by the flow. With this in mind, Powell (1994) added source terms to the momentum, energy and induction equations, which take the (Lagrangian) form

$$\frac{d\mathrm{v}^i}{dt} = \frac{1}{\rho}\frac{\partial S^{ij}}{\partial x^j} - \frac{B^i}{\rho}\frac{\partial B^j}{\partial x^j}, \tag{5.5}$$

$$\frac{de}{dt} = -\frac{1}{\rho}\frac{\partial (\mathrm{v}_i S^{ij})}{\partial x^j} - \frac{\mathrm{v}_i B^i}{\rho}\frac{\partial B^j}{\partial x^j}, \tag{5.6}$$

$$\frac{dB^i}{dt} = B^j \frac{\partial \mathrm{v}^i}{\partial x^j} - B^i \frac{\partial \mathrm{v}^j}{\partial x^j}, \tag{5.7}$$

where as in the previous chapter the stress tensor is defined as

$$S^{ij} = -P\delta^{ij} + \frac{1}{\mu_0}\left(B^i B^j - \frac{1}{2}B^2\delta^{ij}\right). \tag{5.8}$$

Taking the divergence of (5.7) shows that the divergence errors in this formalism evolve according to

$$\frac{\partial}{\partial t}(\nabla \cdot \mathbf{B}) + \nabla \cdot (\mathbf{v}\nabla \cdot \mathbf{B}) = 0, \tag{5.9}$$

which has the same form as the continuity equation for the density (where in this case we have a density of magnetic monopoles, $\nabla \cdot \mathbf{B}$). This therefore implies that the total volume integral of $\nabla \cdot \mathbf{B}$ across the simulation is conserved and hence that the *surface* integral of the flux (4.18) is conserved. As discussed in §4.2.2 the conservation of this quantity is a far more important physically than the conservation of the volume integral (4.17).

The disadvantage of using (5.5)-(5.7) is that exact conservation of momentum and energy is sacrificed, which proves to be important for shock-type problems. Correspondingly it can lead to incorrect jump conditions at shock fronts (Tóth, 2000). More recently it has been shown by Janhunen (2000) and Dellar (2001) that the correct formulation of the MHD equations in the presence of monopoles should *not* violate the conservation of momentum and energy, giving

$$\frac{d\mathrm{v}^i}{dt} = \frac{1}{\rho}\frac{\partial S^{ij}}{\partial x^j}, \tag{5.10}$$

$$\frac{de}{dt} = -\frac{1}{\rho}\frac{\partial (\mathrm{v}_i S^{ij})}{\partial x^j}, \tag{5.11}$$

$$\frac{dB^i}{dt} = B^j \frac{\partial \mathrm{v}^i}{\partial x^j} - B^i \frac{\partial \mathrm{v}^j}{\partial x^j}. \tag{5.12}$$

Note that the induction equation (5.12) is the same as in Powell's method and therefore the manner in which the divergence errors evolve (5.9) is exactly the same. We have shown in §4.3.2 that equations (5.10) and (5.12) are indeed consistent with each other by deriving the SPH form of (5.10) from a variational principle which uses the SPH form of (5.12) as a constraint.

5.2.2 Projection methods

A common approach to the divergence problem is to clean up the magnetic field at regular intervals via the *projection method* (e.g. Brackbill and Barnes, 1980). The basic idea is to decompose the magnetic field into a curl and a gradient (which can be done unambiguously for any vector field) according to

$$\mathbf{B}^* = \nabla \times \mathbf{A} + \nabla\phi. \tag{5.13}$$

Taking the divergence of this expression results in the Poisson equation

$$\nabla^2\phi = \nabla \cdot \mathbf{B}^*, \tag{5.14}$$

which can then be solved for the scalar quantity ϕ. The magnetic field is then corrected according to

$$\mathbf{B} = \mathbf{B}^* - \nabla\phi. \tag{5.15}$$

The major disadvantage with this approach is that the solution of the Poisson equation (5.14) is computationally expensive, scaling as $\mathscr{O}(N^2)$. In an astrophysical SPH context this may be offset somewhat by the fact that the Poisson equation for the gravitational field is usually solved using a tree code (e.g. Hernquist and Katz 1989; Benz et al. 1990) which scales as $\mathscr{O}(N\log N)$. However there are some subtle difficulties with this approach, which we outline below.

Projection schemes for incompressible flow in SPH have been implemented by Cummins and Rudman (1999), the results of which are applicable to the present case. The important point, also discussed by Tóth (2000) is that for the projection step to reduce the divergence to zero (ie. to provide an *exact* projection) requires that the discrete version of (5.14) is satisfied exactly. This means that the operator used to evaluate the divergence term on the right hand side of (5.14) should be the same as the divergence operator used in the evaluation of the ∇^2 on the left hand side and that the gradient operator used in the evaluation of ∇^2 should be the same as that used in 5.15. Cummins and Rudman (1999) approach this problem by calculating the ∇^2 using SPH operators, solving the Poisson equation by matrix inversion. Good results were also obtained using an approximate projection (ie. where the divergence operators on the left and right hand side differ). In this scheme Cummins and Rudman (1999) used the SPH evaluation of the Laplacian given in §3.2.4, similar to that used in the artificial dissipation terms (4.80)-(4.85). The Poisson equation is then solved by inverting the resulting matrix equation.

The solution of (5.14) by direct summation (of which the tree code is an approximation), uses the exact solution to the Poisson equation (5.14) given by

$$\phi(\mathbf{r}) = \int G(|\mathbf{r}-\mathbf{r}'|)\nabla \cdot \mathbf{B}(\mathbf{r}')\mathrm{d}V(\mathbf{r}'), \tag{5.16}$$

where $G(|\mathbf{r}-\mathbf{r}'|)$ is the Green's function, given by

$$\begin{aligned} G(r) &= \frac{1}{2\pi}\ln r + \text{const}, \\ G(r) &= -\frac{1}{4\pi r}, \end{aligned} \tag{5.17}$$

in two and three dimensions respectively. The gradient needed in the correction step can be calculated directly, giving (in three dimensions)

$$\nabla\phi(\mathbf{r}) = -\frac{1}{4\pi}\int \frac{\nabla\cdot\mathbf{B}(\mathbf{r}')}{|\mathbf{r}-\mathbf{r}'|^3}(\mathbf{r}-\mathbf{r}')\mathrm{dV}(\mathbf{r}'). \tag{5.18}$$

In SPH we replace the volume element $\rho\mathrm{dV}$ with the mass per SPH particle and write the integral as a summation according to

$$\nabla\phi_a = -\sum_b m_b \frac{(\nabla\cdot\mathbf{B})_b}{4\pi\rho_b}\frac{(\mathbf{r}_a-\mathbf{r}_b)}{|\mathbf{r}_a-\mathbf{r}_b|^3}. \tag{5.19}$$

Since we still retain the freedom to choose the discrete operator used to evaluate $\nabla\cdot\mathbf{B}$ at each particle, it becomes clear that the solution by direct summation will only provide an *approximate* projection, since (5.14) is not discretely satisfied. This approximate solution will be degraded further when implemented using a tree code. A further disadvantage of the projection method for many of the problems considered in this paper is that it is somewhat complicated to implement in the case of periodic boundary conditions. Despite these subtleties the projection method based on the Green's function solution is found to give a substantial reduction in the divergence errors in a single step (§5.3.2).

The projection step is implemented in this thesis as follows: For a given magnetic field on the particles, the divergence is calculated using (5.31). The correction to the field is then calculated by a direct summation using (5.19) (with the Green's function appropriate to the number of spatial dimensions) and subtracted accordingly. Using the timestepping scheme described in §3.6 the correction is made to the magnetic field $\mathbf{B}^0$ at the beginning of the timestep. This means that the divergence is calculated in a separate loop to the usual force calculation.

An alternative projection scheme can be implemented by solving for the vector potential $\mathbf{A}$. That is, we take the curl of (5.13) to obtain

$$\nabla\times\mathbf{B}^* = \nabla(\nabla\cdot\mathbf{A}) - \nabla^2\mathbf{A}. \tag{5.20}$$

Choosing the Gauge condition $\nabla\cdot\mathbf{A}=0$, we obtain a Poisson equation for the vector potential in terms of the current density $\mathbf{J}=\nabla\times\mathbf{B}^*/\mu_0$

$$\nabla^2\mathbf{A} = -\mu_0\mathbf{J} \tag{5.21}$$

with solution

$$\mathbf{A}(\mathbf{r}) = \int G(|\mathbf{r}-\mathbf{r}'|)\mathbf{J}(\mathbf{r}')\mathrm{dV}(\mathbf{r}'). \tag{5.22}$$

Taking the curl, we obtain an equation for the corrected magnetic field in terms of the current density, which in three dimensions is given by

$$\mathbf{B} = \nabla\times\mathbf{A} = -\frac{\mu_0}{4\pi}\int \frac{\mathbf{J}(\mathbf{r}')\times(\mathbf{r}-\mathbf{r}')}{|\mathbf{r}-\mathbf{r}'|^3}\mathrm{dV}(\mathbf{r}'). \tag{5.23}$$

which is simply Biot-Savart's Law. In SPH form this is given by

$$\mathbf{B}_a = -\sum_b m_b \frac{(\nabla \times \mathbf{B}^*)_b \times (\mathbf{r}_a - \mathbf{r}_b)}{4\pi\rho_b |\mathbf{r}_a - \mathbf{r}_b|^3}. \tag{5.24}$$

This method could be useful in an SPH context in situations where several disconnected regions exist containing strong magnetic currents (such as in two isolated neutron stars). By solving (5.23), the corrected magnetic field is determined from the current density, resulting in a knowledge of the magnetic field at any point in space. This approach was in fact used as the basis for the very first SPMHD algorithm implemented by Gingold and Monaghan (1977). As a divergence cleaning method, we find that the results are very similar to those obtained using (5.19), although at a slightly higher computational cost since the Poisson equation (5.21) is solved for a vector quantity rather than a scalar, giving (up to) three summations in (5.24) as opposed to just one in (5.19).

Finally it is worth commenting on the possibility of using iterative methods for solving the Poisson equation (5.14), although there is not the time or space available to investigate these ideas further in this thesis. The main point is that divergence errors usually arise in a simulation as short wavelength errors, typically of opposite sign. Obtaining the full solution to the Poisson equation (using the Green's function or otherwise) is computationally expensive because both the long and short wavelength components must be accounted for. This is perhaps best illustrated by the multigrid methods which explicitly tackle the problem in this manner by using simple iterative schemes such as the Jacobi or Gauss-Seidel methods (which are good at removing the short wavelength errors) on a progressively coarser heirarchy of grids (thus acting on progressively longer wavelength errors). The solution by direct summation (5.19) is slow because the (small) contribution from distant neighbours must be accounted for (which is accelerated in the tree code by treating groups of distant particles as single entities). However, since for the purposes of divergence cleaning we are interested in eliminating mainly the short wavelength errors, performing simple iterations on the Poisson equation expressed using SPH operators may give satisfactory results with a much lower computational expense. Furthermore an approximate solution to a specified accuracy (which may be achieved in just a few iterations) is all that is really required from the cleaning procedure, rather than the full, exact solution. A similar point has been made by Tóth (2000). An iterative solution to the Poisson equation (5.14) can be obtained by solving a diffusion equation of the form

$$\frac{\partial \phi}{\partial t} = \nabla^2 \phi - \nabla \cdot \mathbf{B} \tag{5.25}$$

via a relaxation method (Press et al., 1992). Methods for solving diffusion equations implicitly using iterative procedures have been recently developed for use in SPH by Whitehouse and Bate (2004) and Monaghan (1997a) and it may be possible to apply these ideas to the divergence cleaning problem.

5.2.3 Hyperbolic divergence cleaning

Dedner et al. (2002) examine alternative divergence cleaning procedures. In their paper (see also Munz et al., 2000), they derive a general constrained formulation of the MHD equations, from which formalisms can be derived to give divergence cleaning which is elliptic (involving the solution of a Poisson equation), parabolic (in which the divergence errors are diffused away) and hyperbolic (where the diver-

gence errors are propagated away from their source at a characteristic speed). The projection method described above is an elliptic approach, the main disadvantage to which is the substantial computational cost involved in the solution of the Poisson equation. The parabolic approach was found to be severely limited in scope due to the timestep restrictions imposed by the Courant condition[3]. The hyperbolic approach was found to be particularly effective, especially when combined with a parabolic term such that divergence errors are both transported and diffused. It is this approach that we outline below in an SPH context.

The basic idea is to introduce an additional scalar field ψ, which is coupled to the magnetic field by a gradient term in the induction equation,

$$\frac{d\mathbf{B}}{dt} = -\mathbf{B}(\nabla\cdot\mathbf{v}) + (\mathbf{B}\cdot\nabla)\mathbf{v} - \nabla\psi. \tag{5.26}$$

Note that our induction equation maintains the consistent treatment of divergence terms discussed above. The variable ψ is then calculated by adding an additional constraint equation, which for the combined hyperbolic/parabolic approach is given by

$$\frac{d\psi}{dt} = -c_h^2(\nabla\cdot\mathbf{B}) - \frac{\psi}{\tau}. \tag{5.27}$$

Neglecting the second term on the right hand side of (5.27) gives an equation for ψ which is purely hyperbolic. This implies that divergence errors are propagated in a wave-like manner away from their source with characteristic speed c_h (for more details we refer the reader to the Dedner et al. (2002) paper). The second term on the right hand side is a parabolic term which causes ψ to decay exponentially to zero with e-folding time τ (this is easily seen by neglecting the hyperbolic term and solving the resulting ordinary differential equation for $\psi(t)$). Since it is desirable for the divergence errors to be propagated at the maximum possible rate (within the timestep constraint imposed by the Courant condition), c_h should be set equal to the maximum signal propagation speed. For simplicity we calculate this as

$$c_h = \sqrt{\frac{\gamma P}{\rho} + \frac{1}{2}\frac{B^2}{\mu_0\rho}}, \tag{5.28}$$

where the maximum value over all of the particles is used. The decay timescale τ is given by

$$\frac{1}{\tau_a} = \frac{\sigma c_h}{h_a}, \tag{5.29}$$

where σ is a dimensionless parameter which determines the decay timescale. Setting $\sigma = 0$ gives a purely hyperbolic correction. A value of $\sigma = 0.2$ would imply that ψ will have decayed significantly after the divergence errors have propagated approximately 5 smoothing lengths. In §5.3.2 we examine in detail the effects of varying the value of σ. We find that values of σ in the range $0.05 - 0.2$ generally give the best results, giving a good balance between the hyperbolic (fast but non-diffusive) and parabolic (diffusive but slow-acting) effects. In practise some diffusion is also added by the artificial resistivity terms (§4.5).

[3] an equivalent approach in SPMHD is to use an artificial resistivity in order to diffuse away divergence errors. This has been used, for example, by Morris (1996) and Hosking (2002)

The gradient term in the induction equation is calculated using a simple SPH estimate (§3.2.3)

$$\nabla\psi_a = \frac{1}{\rho_a}\sum_b m_b(\psi_b - \psi_a)\nabla_a W_{ab}. \tag{5.30}$$

Similarly the divergence of the magnetic field is calculated using

$$(\nabla\cdot\mathbf{B})_a = \sum_b m_b(\mathbf{B}_a - \mathbf{B}_b)\cdot\nabla_a W_{ab}. \tag{5.31}$$

Superfast cleaning

This type of divergence cleaning is most effective when the dynamics in the simulation occur at speeds lower than the fastest wave speed. In this case the divergence cleaning is able to propagate and diffuse the divergence errors faster than they are created in the flow. For the same reason this method is also more effective for codes using a single timestep rather than individual particle timesteps, since the divergence cleaning can take advantage of the 'slack' in the timestep criterion (using individual particle timesteps c_h would be different for each particle). For simulations where divergence errors are generated very quickly (e.g. for problems involving strong shocks) the timescale for removal of the error using the cleaning described above can be too slow to prevent significant errors in the dynamics. One possibility for such problems is to use 'superfast' cleaning, that is to increase the wave speed c_h beyond the maximum imposed by the timestep condition. An operator splitting procedure could then be used to solve the constraint propagation separately between timesteps. For example, having determined the need for extra cleaning by some error criterion, we would then solve the following system of equations in a series of smaller steps which are fractions of the full timestep:

$$\frac{d\mathbf{B}}{dt} = -\nabla\psi \tag{5.32}$$

$$\frac{d\psi}{dt} = -(c_h^*)^2(\nabla\cdot\mathbf{B}) - \frac{\psi}{\tau^*}. \tag{5.33}$$

In the above c_h^* is some multiple of c_h (where the multiplication factor determines the number of substeps necessary – for example using twice the fastest wave speed would require two substeps) and τ^* is the corresponding decay timescale. Note that during the substeps the particles are fixed, such that the neighbour lists do not have to be reconstructed. All that is required is to find the updated estimates of $\nabla\psi$ and $\nabla\cdot\mathbf{B}$ at each substep. The usual induction equation would then be evolved through the full timestep, adding the result to the magnetic field which has been evolved through the constraint substeps.

5.3 Numerical tests

The main issue to be addressed in 2D and 3D problems is the non-zero divergence of the magnetic field. In the SPH context it also allows us to estimate the extent to which the artificial dissipation spuriously affects the numerical results. Again there is a substantial literature of multi-dimensional MHD problems which have been used to test grid-based MHD codes (e.g. Dai and Woodward 1994; Ryu et al. 1995; Balsara 1998; Dai and Woodward 1998; Tóth 2000) and we consider several of these problems here.

5.3.1 Implementation

The implementation of the SPMHD equations used for the multidimensional tests is almost identical to that used in the one dimensional case (§4.6). The density is calculated by summation, the total energy equation is used (although results are indistinguishable using the thermal energy equation in nearly all cases) and the magnetic field is evolved using (4.20) (or using (5.26) when using the hyperbolic cleaning). In the shock tube tests we use unsmoothed initial conditions. The artificial dissipative terms, except where otherwise indicated are implemented using the jump in total magnetic energy (§4.5.1) but as in the one dimensional case the viscosity term uses only the velocity component along the line joining the particles (4.80). Artificial viscosity and thermal conductivity are applied using the switches discussed in §3.5.2 whilst the artificial resistivity term is applied uniformly using $\alpha_B = 1$. A major difference between the simulations presented here and those in the previous chapter is that the anticlumping approach was not found to be uniformly successful in eliminating the tensile instability for all of the problems considered (in particular for the Alfvén wave test only a narrow range of parameters would produce stable results). Furthermore this term was found to result in spurious extra numerical noise, particularly in the shock tube tests. For this reason we have eliminated the tensile instability by simply subtracting the constant component of the magnetic field from the gradient term (§4.4.4). However all of the test cases have also been run using the stable Morris formulation of the magnetic force (§4.4.2) and show very little difference in the results.

Error estimates

Various estimates can be made of the error produced in the simulation by any non-zero magnetic divergence. Monitoring these quantities over the course of a simulation thereby gives some measure of the magnitude of the error produced by $\nabla \cdot \mathbf{B}$. The most common approach in SPH implementations to date has been to monitor the dimensionless quantity

$$\frac{h\nabla \cdot \mathbf{B}}{|\mathbf{B}|} \tag{5.34}$$

and ensure that it remains small (typically < 0.01) over most of the simulation, where h is the SPH smoothing length and the divergence is calculated using (5.31). This provides some measure of the relative error in the magnetic field but no indication of how much influence this error has in the dynamics. For this reason it is also useful to measure the relative error in the total force caused by a non-zero divergence,

$$E_{force} = \frac{\mathbf{f}_{mag} \cdot \mathbf{B}}{|\mathbf{f}||\mathbf{B}|} \tag{5.35}$$

where $\mathbf{f}_{mag}$ is the magnetic component of the SPH force (4.33), whilst $\mathbf{f}$ is the total force on the particle. It is also useful to simply monitor the evolution in the maximum, minimum and average of $|\nabla \cdot \mathbf{B}|$ with time as well as the conserved quantities given in §4.2.2.

Visualisation

In order to make a direct comparison of our results with those of grid-based MHD codes, we interpolate the results from the particles to an array of pixels using the SPH kernel. That is, for a contour or rendered plot of a scalar quantity ϕ we interpolate to the pixels using

$$\phi(x,y) = \sum_b m_b \frac{\phi_b}{\rho_b} W(x-x_b, y-y_b, h_b) \tag{5.36}$$

where W is the cubic spline kernel used in the calculations (§3.2.5) and the summation is over contributing particles. Note that in practise this is quite simple to implement, as it involves only one loop over the particles, during which the contributions from the current particle to all pixels within a smoothing radius ($2h$) are calculated. For a vector quantity a similar interpolation can be performed for each component, however since in this case the interpolation is usually to a coarser grid, it is simpler just to bin the particles into grid cells and take the average of the vector components in each cell.

5.3.2 $\nabla \cdot \mathbf{B}$ advection

The first problem we examine is a simple test used by Dedner et al. (2002) in which a non-zero magnetic divergence is introduced into the simulation as an initial condition. This is a particularly good test for comparing various divergence cleaning procedures. The initial conditions are a uniform density distribution ($\rho = 1$) in the domain $-0.5 < x < 1.5, -0.5 < y < 1.5$ with a constant initial velocity field $\mathbf{v} = [1,1]$. The initial gas pressure is $P = 6$ with $\gamma = 5/3$ and the magnetic field has a constant component perpendicular to the plane $B_z = 1/\sqrt{4\pi}$. The divergence is introduced as a peak in the $x-$component of the field in the form

$$B_x = 4096(r^2)^4 - 128(r^2)^2 + 1 \qquad\qquad r^2 = x^2 + y^2 \tag{5.37}$$

the contours of which are shown in the left column of Figure 5.1. The particles are arranged on a cubic lattice for simplicity and periodic boundary conditions are enforced using ghost particles. Since the density is uniform throughout the simulation the results are insensitive to whether (4.20) or (4.22) is used and also to the anticlumping term since the simulation is not unstable to negative stress. The artificial dissipation terms are turned off for this problem in order to isolate the effects of the divergence cleaning procedures.

The results of this test using four different divergence cleaning techniques are shown in Figure 5.1. The plots show contours of the $x-$component of the magnetic field as it evolves in each case (30 contours are plotted, evenly spaced between the minimum and maximum of B_x over all the simulations). The results using the consistent formulation of $\nabla \cdot \mathbf{B}$ terms discussed in the previous chapter and in §5.2.1 are shown in the top row. In this case the divergence error is passively advected by the flow and both the field and the divergence error remain unchanged (relative to the flow) at $t = 1$, demonstrating that the formalism is indeed consistent in the presence of magnetic monopoles. In order to compare these results with a conservative formulation of the MHD equations, we have performed a simulation using an SPH

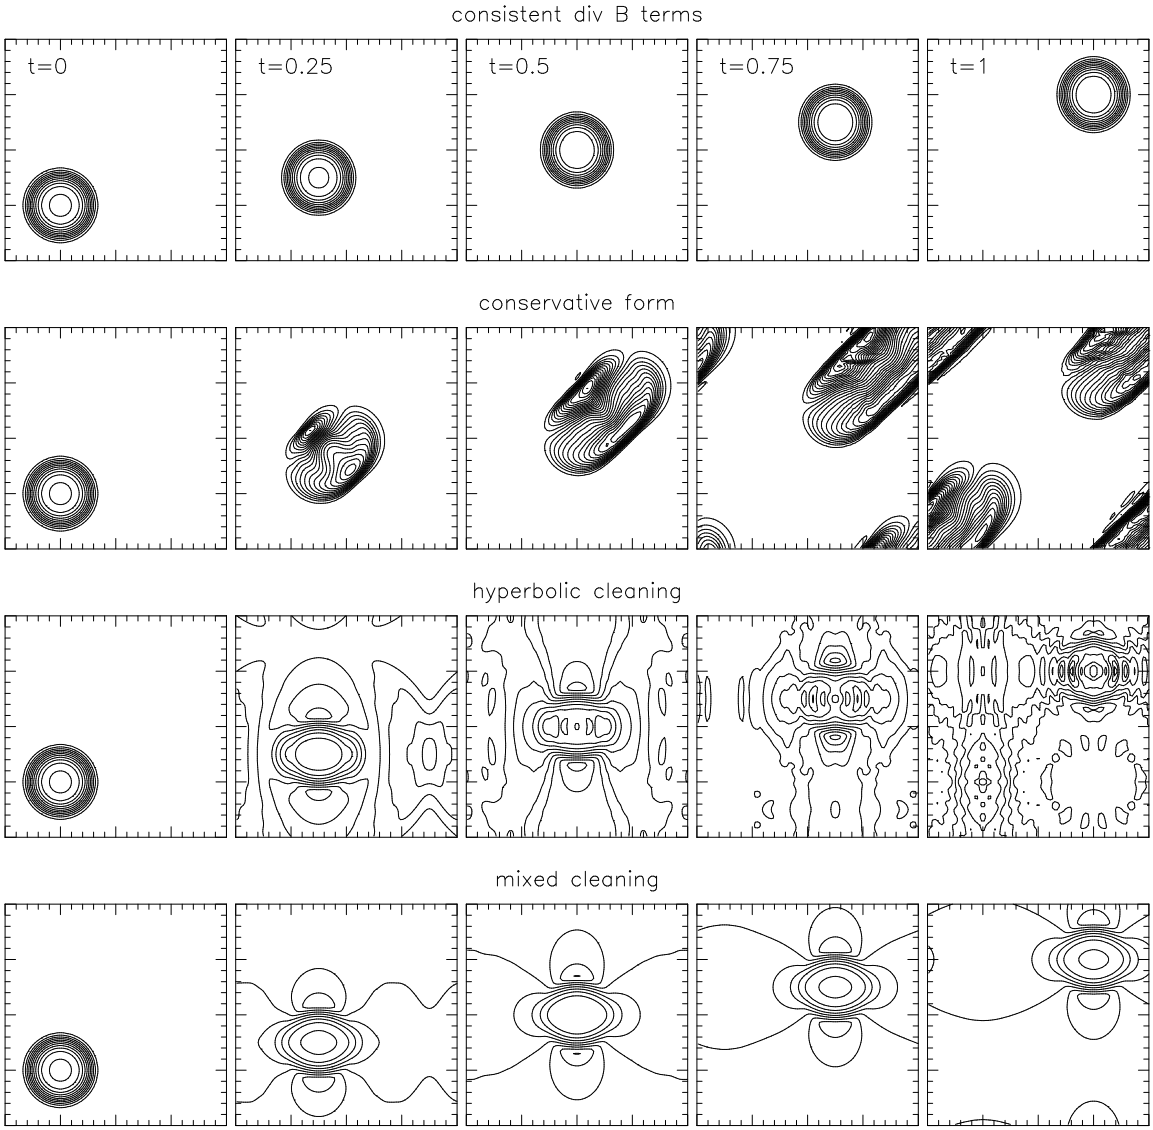

Figure 5.1: Results of the $\nabla \cdot \mathbf{B}$ advection problem. An initially non-zero divergence is setup as a peak in the $x-$component of the magnetic field (leftmost figures), with a velocity field $\mathbf{v}(x,y) = [1,1]$ and periodic boundaries. The plots show contours in B_x at various times throughout the simulation for various divergence cleaning procedures. The consistent treatment of $\nabla \cdot \mathbf{B}$ terms (top row) is clearly seen to advect the divergence without change, which is an improvement over a conservative formulation of the MHD equations in which the divergence is smeared throughout the simulation volume (second row). With the use of hyperbolic cleaning in addition to the consistent $\nabla \cdot \mathbf{B}$ terms, the divergence error is spread rapidly (middle row), whilst with a mixed hyperbolic/parabolic cleaning (fourth row) this error is also diffused away, resulting in a divergence-free field configuration (compare the bottom row with the results using the projection method in Figure 5.3).

induction equation of the form

$$\frac{d}{dt}\left(\frac{B_a^i}{\rho_a}\right) = \sum_b m_b \left[\frac{B_a^j}{\rho_a^2}(\mathrm{v}_b^i - \mathrm{v}_a^i) + \frac{\mathrm{v}_a^i}{\rho_a^2}(B_b^j - B_a^j)\right] \frac{\partial W_{ab}}{\partial x_a^j} \tag{5.38}$$

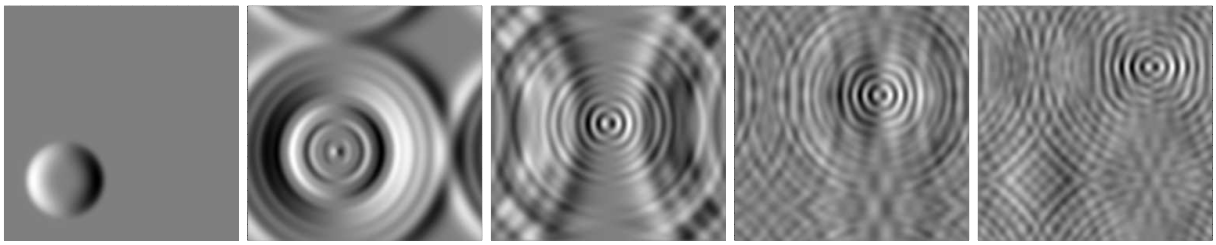

Figure 5.2: Divergence of the magnetic field in the $\nabla \cdot \mathbf{B}$ advection problem at the times shown in Figure 5.1 using the hyperbolic divergence cleaning discussed in §5.2.3. The divergence error is rapidly spread in a wavelike manner throughout the simulation volume, although in the absence of diffusion the overall error does not decrease in magnitude. Periodic boundary conditions are used, resulting in the interference patterns seen at later times.

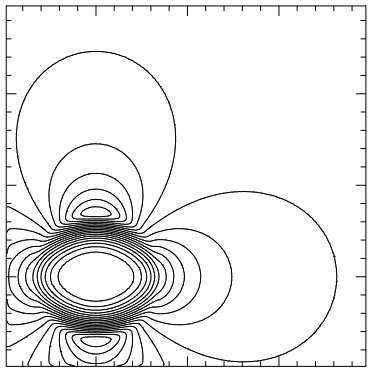

Figure 5.3: Divergence cleaning using the approximate projection method described in §5.2.2. The plot shows 30 contours of B_x in the $\nabla \cdot \mathbf{B}$ advection problem after a single projection step at $t = 0$. The results may be compared to those shown in Figure 5.1. The projected magnetic field adopts an essentially divergence-free configuration in a single step.

which is an SPH form of the conservative (in a volume sense) induction equation

$$\frac{d}{dt}\left(\frac{\mathbf{B}}{\rho}\right) = \left(\frac{\mathbf{B}}{\rho}\cdot\nabla\right)\mathbf{v} + \mathbf{v}\left(\frac{\nabla\cdot\mathbf{B}}{\rho}\right). \tag{5.39}$$

The results using this formalism are shown in the second row of Figure 5.1. The peak in B_x is distorted by the flow and the divergence error is smeared throughout the simulation.

The third row in Figure 5.1 shows the results using the divergence correction discussed in §5.2.3 using only the hyperbolic term in (5.27)(ie. with $\sigma = 0$) in conjunction with the usual monopole formulation of the induction equation (4.22). The divergence error is spread rapidly in a wavelike manner by the constraint equation (5.27)(this is graphically illustrated in Figure 5.2 which shows the propagation of the divergence error in this simulation). However, the magnitude does not decrease in this case. Using the mixed hyperbolic/parabolic cleaning with a small amount of diffusion (using the parabolic term in (5.27), in this case with $\sigma = 0.1$), this error is rapidly diffused away, resulting in a divergence-free field configuration (Figure 5.1, bottom row). For comparison, the results of a single projection step at $t = 0$ are shown in Figure 5.3, showing the divergence-free configuration adopted by the field. The projection step is calculated as described in §5.2.2.

The time evolution of various quantities throughout these simulations are shown in Figure 5.4. The

left panels show the evolution of the maximum (top) and average (bottom) of $|\nabla \cdot \mathbf{B}|$. In conservative form (solid line) the maximum divergence varies slightly and initially becomes larger than the initial value. The bottom panel shows that the average value in this case steadily increases over time, due to the smearing effect of the divergence propagation (5.4). The consistent formulation of $\nabla \cdot \mathbf{B}$ terms (dashed line) maintains a steady value of both the maximum and average, as observed in Figure 5.1. With hyperbolic cleaning (dot-dashed) the maximum divergence error is quickly reduced (although increases at later times as the divergence waves cross the periodic domain and interact) whilst the average climbs as the divergence error is spread throughout the domain. Using the mixed hyperbolic/parabolic cleaning as described above (dotted line), both the maximum and average divergence is swiftly reduced. For comparison, results using the projection method where a projection step is taken every 10 timesteps are also plotted (dashed-dot-dashed). Note however that the boundary conditions are assumed to be open for this problem which means that the periodicity is not accounted for. At early times this is a valid assumption as the source term for the Poisson equation (ie. $\nabla \cdot \mathbf{B}$) is non-zero in only a finite region of the simulation volume. However as the divergence is spread by the cleaning this assumption breaks down and a fully periodic treatment should be used.

The magnitude of the volume integral of the flux (4.17) and of the cross helicity (4.16) are shown in the right hand panels of Figure 5.4. Although (as discussed in §4.2.2) the conservation of the volume integral of the flux is not particularly important physically, this plot demonstrates that this quantity is conserved more accurately using a conservative formulation of the induction equation than when using the monopole-consistent formulation [4]. However, the opposite is true in the conservation of cross-helicity (which measures the preservation of the flux-frozenness condition, c.f. §4.2.2). With any kind of divergence cleaning, the flux integral is conserved to a much higher degree of accuracy and the same is true for the cross-helicity except in the case of the projection method. The projection method does not conserve the cross-helicity invariant since the divergence cleaning is done without any knowledge of the velocity field. In the hyperbolic/parabolic cleaning the induction equation is still explicitly evolved and therefore the flux-frozenness condition is still maintained.

Finally the effect of varying the strength of the parabolic (diffusion) term in (5.27) is examined. In Figure 5.5 the time evolution of the maximum of $|\nabla \cdot \mathbf{B}|$ over the particles is shown, varying the diffusion parameter σ. A small amount of diffusion is necessary to remove the divergence error, however as σ is increased the cleaning becomes less effective as the slow-acting parabolic effects dominate. The fastest reduction in $\nabla \cdot \mathbf{B}$ is obtained using $\sigma \sim 0.1 - 0.2$, giving a good balance between the slow-acting diffusion and the spreading produced by the hyperbolic term.

[4] Note that using the conservative induction equation in the form (5.38) does not exactly conserve the volume integral of the flux (4.17) since the gradient terms are not symmetric between the particle pairs. A formalism which does conserve this integral is simple to construct based on (5.39). For example

$$\frac{d}{dt}\left(\frac{B_a^i}{\rho_a}\right) = \sum_b m_b \left(\frac{B_a^j}{\rho_a^2} \mathrm{v}_b^i + \frac{B_b^j}{\rho_b^2} \mathrm{v}_a^i\right) \frac{\partial W_{ab}}{\partial x_a^j} \tag{5.40}$$

explicitly conserves the integral (4.17) since

$$\sum_a m_a \frac{d}{dt}\left(\frac{\mathbf{B}_a}{\rho_a}\right) = 0, \tag{5.41}$$

although the interpolation provided by the terms in (5.40) is not a particularly good one (c.f. §3.2.2).

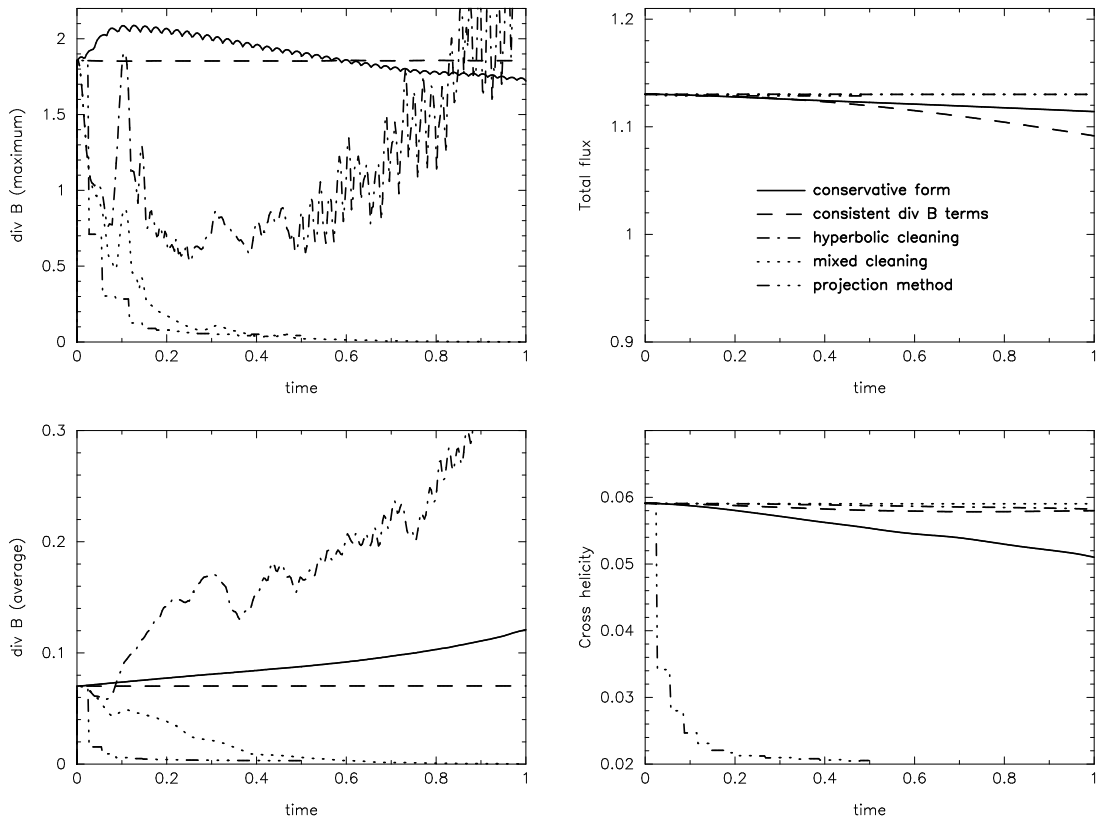

Figure 5.4: Time evolution of various quantities in the $\nabla \cdot \mathbf{B}$ advection test. The left hand panels show the maximum (top) and average (bottom) values of $|\nabla \cdot \mathbf{B}|$ over the particles. With a conservative formulation of the induction equation the divergence error increases with time (solid line) whereas the errors are conserved using a formulation which is consistent in the presence of magnetic monopoles (dashed line). With hyperbolic cleaning (dot-dashed) the maximum is quickly reduced although the average increases, however with the parabolic term included the error is rapidly diffused away (dotted line), giving results comparable to the projection method (dashed-dot-dashed). The right hand panels show the conservation of the volume integral of the flux (top) and the cross-helicity invariant (bottom), which in all cases is improved by the divergence cleaning except in the case of the projection method which does not conserve the cross-helicity.

5.3.3 Circularly polarized Alfvén wave

This test is described by Tóth (2000) where it is used to test a variety of multidimensional MHD schemes in grid based codes. The test involves a circularly polarized Alfvén wave propagating in a two dimensional domain. The advantage of using a circularly (as opposed to linearly) polarized wave is that it turns out to be an exact, non-linear solution to the MHD equations, which means that the solution after one period should exactly match the initial conditions, without the effects of nonlinear steepening (as observed, for example, in the magnetosonic wave tests described in §4.6.4). This also means that the wave can be setup with a much larger amplitude than would be used for purely linear waves.

In Tóth (2000), the wave is setup to propagate at an angle $\theta = 30^\circ$ with respect to the $x-$axis. In SPH the orientation of the wave vector with respect to the co-ordinates is not particularly important because there is no spatial grid. However, we have retained the rotated configuration as firstly it ensures that there are no spurious effects resulting from the initial arrangement of the particles and secondly enables a fair comparison with the results shown in Tóth (2000). The particles are setup on a hexagonal close packed

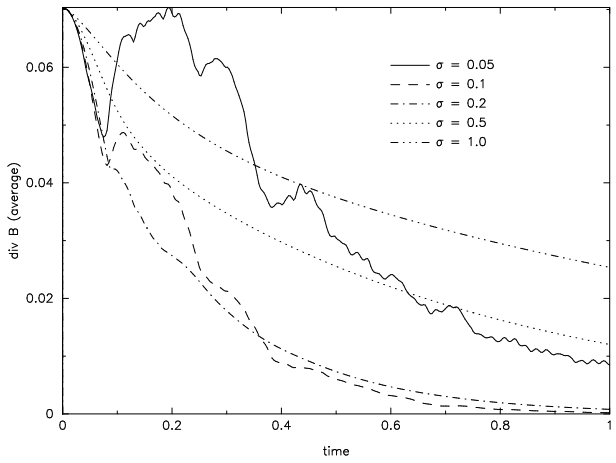

Figure 5.5: Time evolution of the average value of $\nabla \cdot \mathbf{B}$ in the divergence advection problem, varying the diffusion parameter σ. A small amount of diffusion is necessary to remove the divergence error, however as σ is increased the reduction in the divergence lessens as the slow acting diffusion dominates over the rapid spreading produced by the hyperbolic term. The fastest reduction is obtained using $\sigma \sim 0.1 - 0.2$.

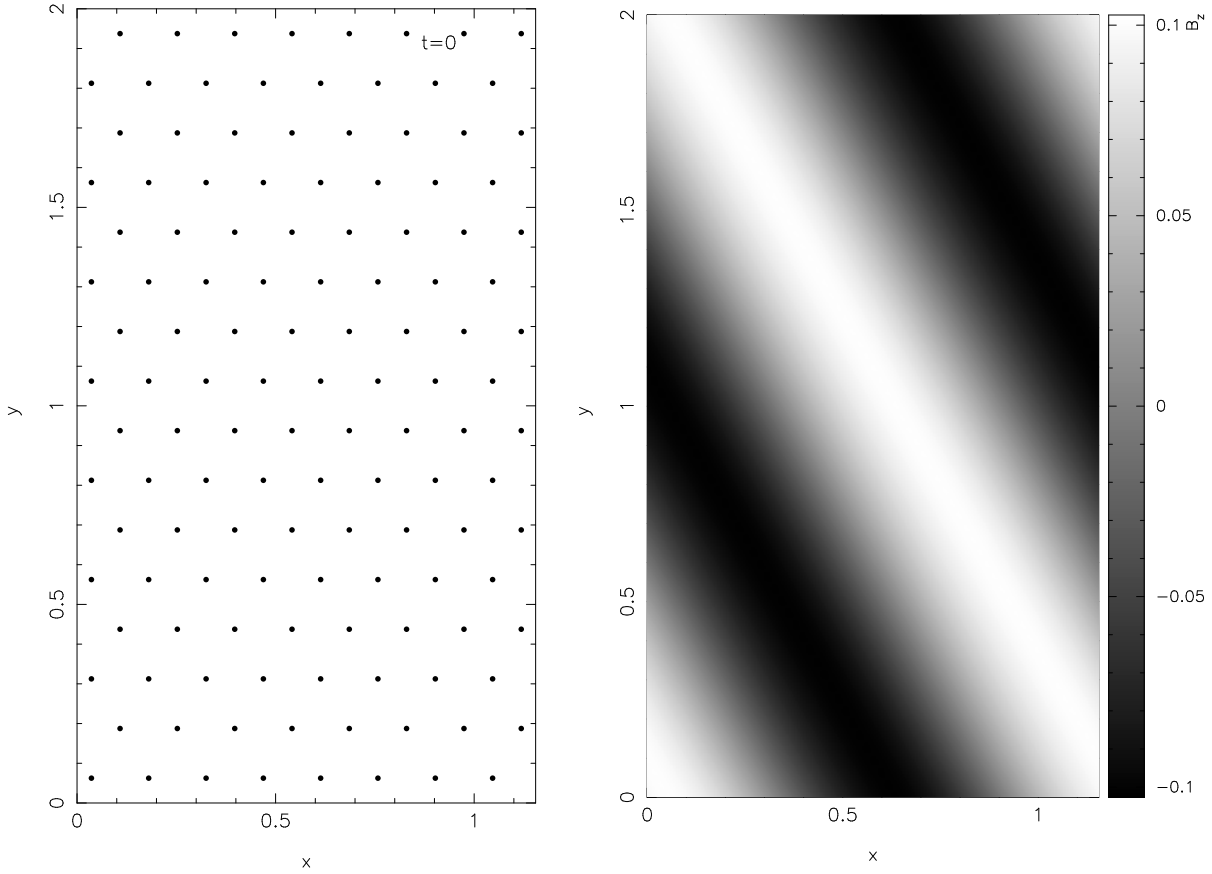

Figure 5.6: Circularly polarized Alfvén wave test. The left figure shows the particle setup in the lowest resolution run. On the right the vertical component of the magnetic field is plotted as a rendered image from the 32×64 particle run at $t = 5$, showing the propagation of the wave with respect to the domain and the particle setup.

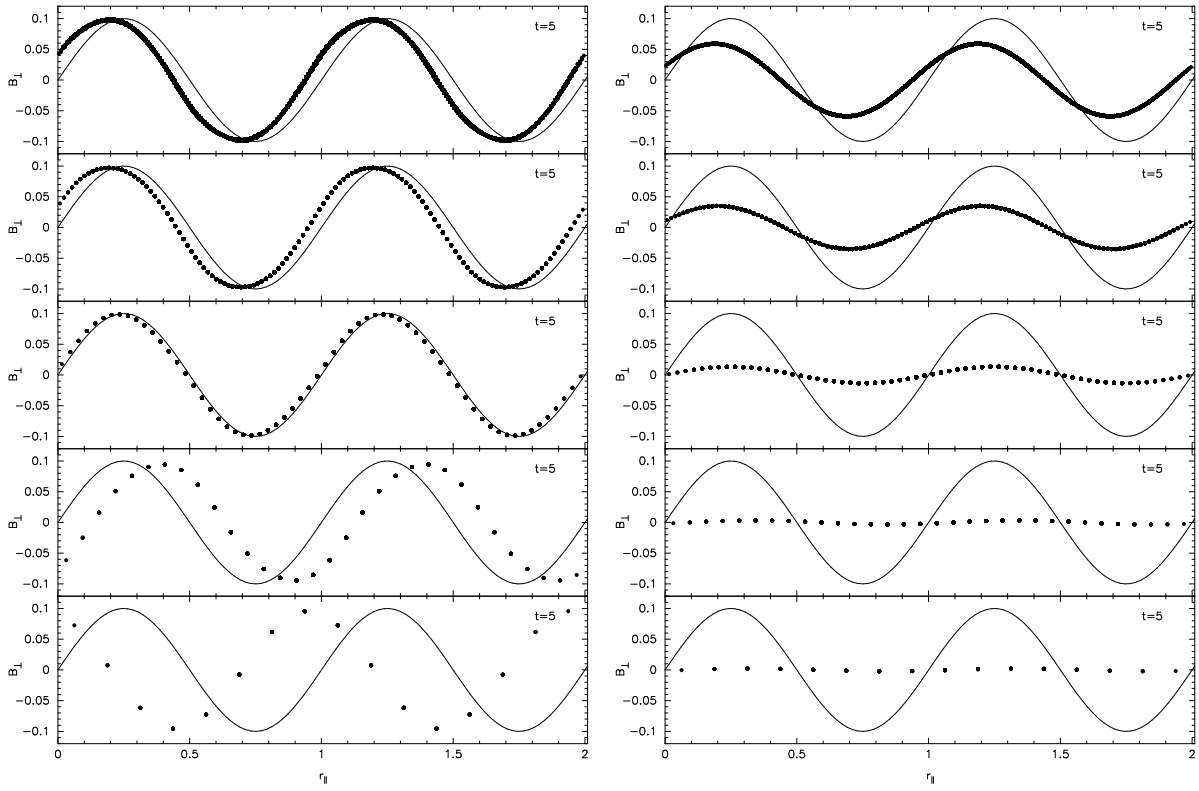

Figure 5.7: Results of the circularly polarized Alfvén wave test at $t = 5$ (corresponding to 5 wave periods). The plots show the perpendicular component of the magnetic field vector $B_\perp = B_y \cos\theta - B_x \sin\theta$ for all of the particles, projected against a vector parallel to the direction of wave propagation $r_\parallel = x\cos\theta + y\sin\theta$ (where $\theta = 30^\circ$ in this case). The SPMHD results are shown at five different resolutions which are, from bottom to top, 8×16, 16×32, 32×64, 64×128 and 128×256. Initial conditions are indicated by the solid line. The numerical results should match these initial conditions at the time shown. The left panel shows the results in the absence of dissipative terms and demonstrates that the SPMHD algorithm contains very little intrinsic numerical dissipation even at low resolutions, although there is a small phase error present even in the converged higher resolution runs. The right hand panel shows the results using the dissipative terms as required in the shock tube problems. In this case the wave amplitude is damped by the artificial resistivity term and exhibits somewhat slow convergence.

lattice (ie. such that particles are equispaced) in a rectangular domain $0 < x < 1/\cos\theta; 0 < y < 1/\sin\theta$. This positioning of the boundaries means that periodic boundary conditions can be used, although some care is required to ensure the continuity of the lattice across the boundaries. This is achieved by stretching the lattice slightly in the $y-$direction to ensure that the boundaries lie at exactly half the spacing of the rows in the lattice. The particle setup at the lowest resolution is shown in the left hand side of Figure 5.6.

The wave is setup with a unit wavelength along the direction of propagation (ie. in this case along the line at an angle of 30° with respect to the x-axis). The initial conditions are $\rho = 1$, $P = 0.1$, $\mathrm{v}_\parallel = 0$, $B_\parallel = 1$, $\mathrm{v}_\perp = B_\perp = 0.1\sin(2\pi r_\parallel)$ and $v_z = B_z = 0.1\cos(2\pi r_\parallel)$ with $\gamma = 5/3$ (where $r_\parallel = x\cos\theta + y\sin\theta$). The $x-$ and $y-$ components of the magnetic field are therefore given by $B_x = B_\parallel\cos\theta - B_\perp\sin\theta$ and $B_y = B_\parallel\sin\theta + B_\perp\cos\theta$ (and similarly for the velocity). Conversely, $B_\parallel = B_y\sin\theta + B_x\cos\theta$ and $B_\perp = B_y\cos\theta - B_x\sin\theta$. Note that this setup means that $\nabla\cdot\mathbf{B} = 0$ holds as a combination of the $\partial B_x/\partial x$ and $\partial B_y/\partial y$ terms, rather than both components being zero individually. The vertical component of the magnetic field after 5 periods is plotted as a rendered image in the right hand side of Figure 5.6, showing

the direction of wave propagation with respect to the domain and the particle setup.

We have performed this test at five different resolutions: 8×16, 16×32, 32×64, 64×128 and 128×256 particles. In each case the number of particles in the y-direction is determined by the hexagonal lattice arrangement. The results are shown in Figure 5.7 after 5 wave periods (corresponding to $t=5$). The plots show the perpendicular component of the magnetic field $B_\perp$ plotted against $r_\parallel$ for all of the particles in the simulation, with the results from the bottom to top panels shown in order of increasing resolution. In each case the initial conditions are indicated by the solid line which is identical to the exact solution at the time shown. The left hand side of Figure 5.7 shows the results in the absence of dissipative terms (that is with the artificial viscosity, resistivity and thermal conductivity turned off). In this case the amplitude agrees very well with the exact solution even at the lowest resolutions. This demonstrates that SPH has a very low intrinsic numerical dissipation (compare for example with the damping of the wave at lower resolutions in the plots shown in Tóth 2000). However there is a small phase error which remains even in the highest resolution run. This is similar to the phase error observed in the one dimensional sound wave tests presented in §3.7.2 and in the one dimensional magnetosonic waves tests in §4.6.4. In these cases the phase error was found to be essentially removed by accounting for the variable smoothing length terms (§3.3.4,§4.3.6). The results shown in Figure 5.7 incorporate the variable smoothing length terms, however in this case the phase error is not completely removed (although is still an improvement over the results using simple averages of the smoothing lengths or kernel gradients). The right hand side of Figure 5.7 shows the results of this test using the dissipative terms as required in the shock tube problems. In this case the wave is severely damped and convergence of the amplitude towards the exact solution is quite slow. The damping is largely caused by the uniform application of artificial resistivity (ie. using $\alpha_B=1$ everywhere) resulting in a somewhat large dissipation even in the absence of shocks. Substantially improved results could be obtained using the resistivity switch discussed in §4.5.2, however for the shock tube problems it was found that use of such a switch could result in too little dissipation at rotational discontinuities in the absence of a shear viscosity term. The divergence error remains very small $\left[(\nabla\cdot\mathbf{B})_{max}\sim10^{-3}\right]$ in all of the simulations shown.

5.3.4 2.5D shock tube

The next two tests are simply two dimensional versions of the one dimensional shock tube tests described in §4.6.3 and demonstrate the effects of divergence errors in the shock capturing scheme. In two dimensions we setup the particles on a cubic lattice in the $x-$direction in the domain $x=[-0.5-\mathrm{v}_{x(L)}t_{max},0.5-\mathrm{v}_{x(R)}t_{max}]$, where $\mathrm{v}_{x(L)}$ and $\mathrm{v}_{x(R)}$ are the initial velocities assigned to the left and right states. This means that at the time t_{max} the particles are contained in the domain $x=[-0.5,0.5]$. The domain has a width of 4 particle spacings in the $y-$direction for computational efficiency. Boundary conditions are implemented by fixing the particle properties in two buffer regions at the edges of the $x-$domain, in which particles are evolved with a fixed velocity but copy their properties $(\rho,P,\mathbf{B})$ from the nearest 'active' particle. Periodic boundary conditions are used in the $y-$direction, implemented using ghost particles. The exact position of the $y-$boundary is chosen to ensure periodicity of the lattice arrangement, ie. at half the spacing of the initial rows of particles in the y-direction. The initial shock is setup as a discontinuity in the fluid quantities at $x=0$ to which no smoothing is applied.

The first shock test is the adiabatic shock tube problem involving seven different discontinuities given

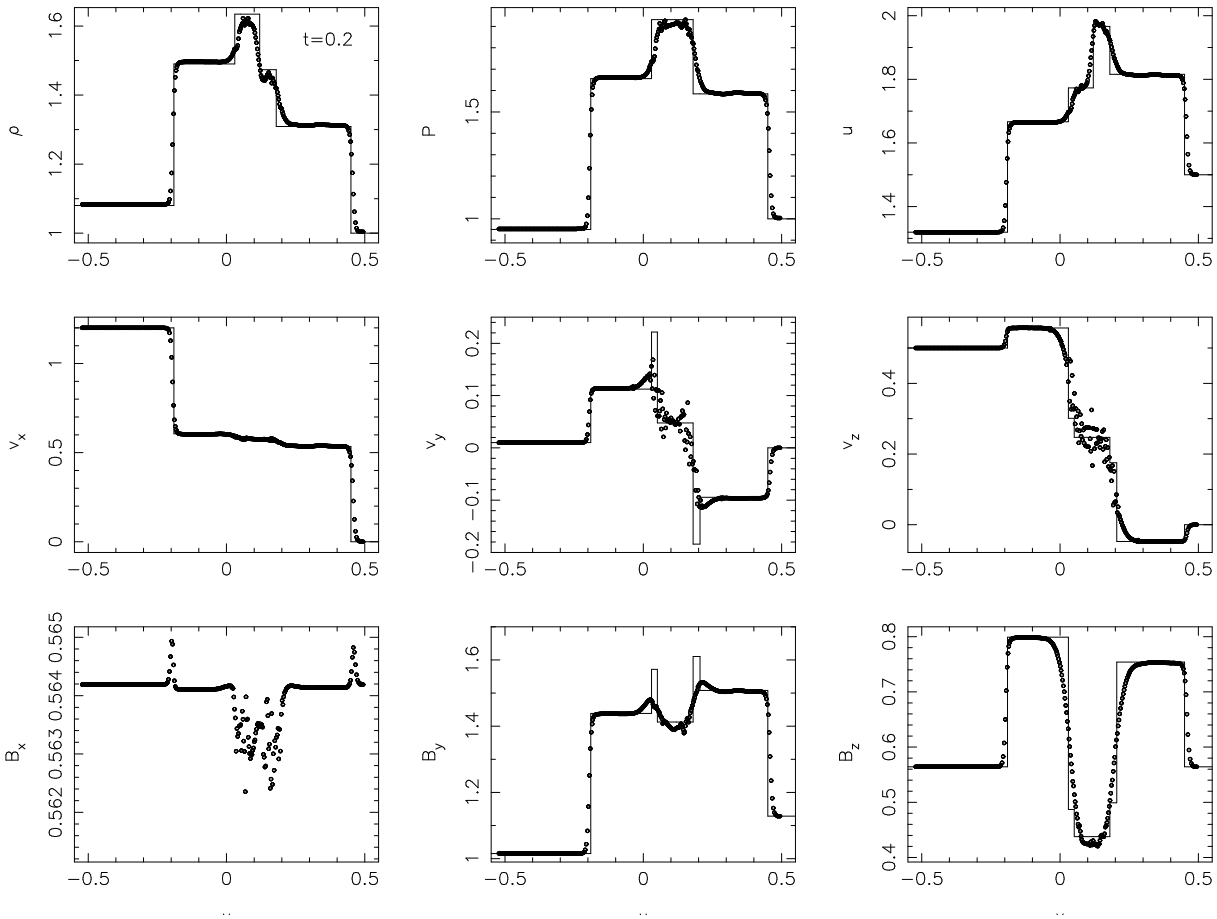

Figure 5.8: Results of the 2.5D shock tube test using an initial smoothing length of $h = 1.2(m/\rho)^{1/2}$ and the dissipative terms as implemented for the one dimensional shock tube problems. In two dimensions at this value of smoothing length small oscillations in the transverse velocity components appear primarily as a result of the non-zero magnetic divergence.

in §4.6.3 (Figure 4.15). Strictly this is a '$2\frac{1}{2}$' dimensional problem since the transverse velocity and magnetic field also have components in the $z-$direction. Conditions to the left of the discontinuity (the left state) are given by $(\rho, P, \mathrm{v}_x, \mathrm{v}_y, \mathrm{v}_z, B_y, B_z) = [1.08, 0.95, 1.2, 0.01, 0.5, 3.6/(4\pi)^{1/2}, 2/(4\pi)^{1/2}]$ whilst to the right (the right state) the conditions are $(\rho, P, \mathrm{v}_x, \mathrm{v}_y, \mathrm{v}_z, B_y, B_z) = [1, 1, 0, 0, 0, 4/(4\pi)^{1/2}, 2/(4\pi)^{1/2}]$ with $B_x = 2/(4\pi)^{1/2}$ everywhere and $\gamma = 5/3$. The problem has been studied by in one dimension by many authors (e.g. Ryu and Jones, 1995; Balsara, 1998) and in two dimensions by Tóth (2000) and Dedner et al. (2002).

The problem is computed using 310×4 particles[5] which corresponds to particle being uniformly spaced on a cubic lattice with separation 0.004, although results are similar using a hexagonal close packed lattice arrangement. Note that this resolution is less than half of that used in the one dimensional case (§4.6.3) but is comparable to, if slightly lower than, the resolutions used in Tóth (2000). The small density difference between the left and right states in this problem is setup by changing the particle masses. The solution using an initial smoothing length of $h = 1.2(m/\rho)^{1/2}$ is shown in Figure 5.8 at $t_{max} = 0.2$ and may be compared with the exact solution taken from Ryu and Jones (1995) (solid line)

[5]Note that this is the number of particles in the domain $-0.5 < x < 0.5$ at $t_{max} = 0.2$ and that the resolution in this domain is correspondingly lower at earlier times due to the inflow boundary condition.

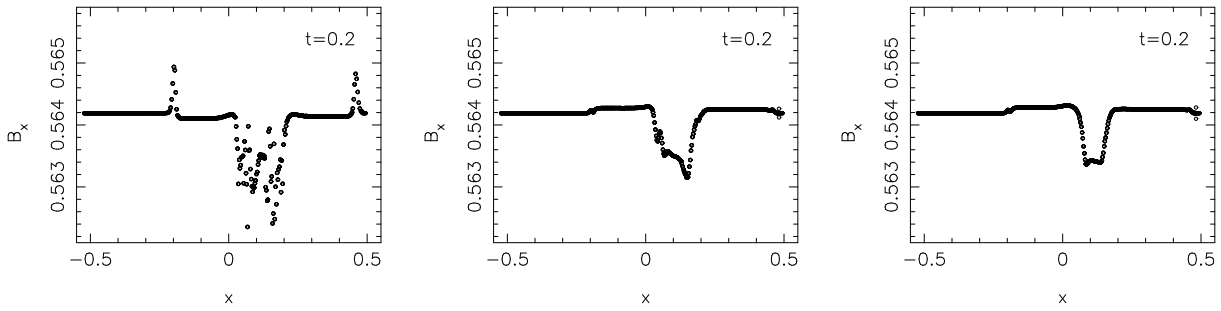

Figure 5.9: The parallel component of the magnetic field in the 2.5D shock tube problem using the dissipative terms as implemented for the 1D problems (left), using the total magnetic energy (centre) and using the total magnetic and kinetic energies (right). Using the total magnetic energy in the dissipative terms means that jumps in the parallel field components are smoothed in addition to the jumps in transverse field. Using the total kinetic energy smooths jumps in the transverse (as well as parallel) velocity components, however this explicitly adds an undesirable shear component to the artificial viscosity term. Details of these formalisms are given in §4.5.

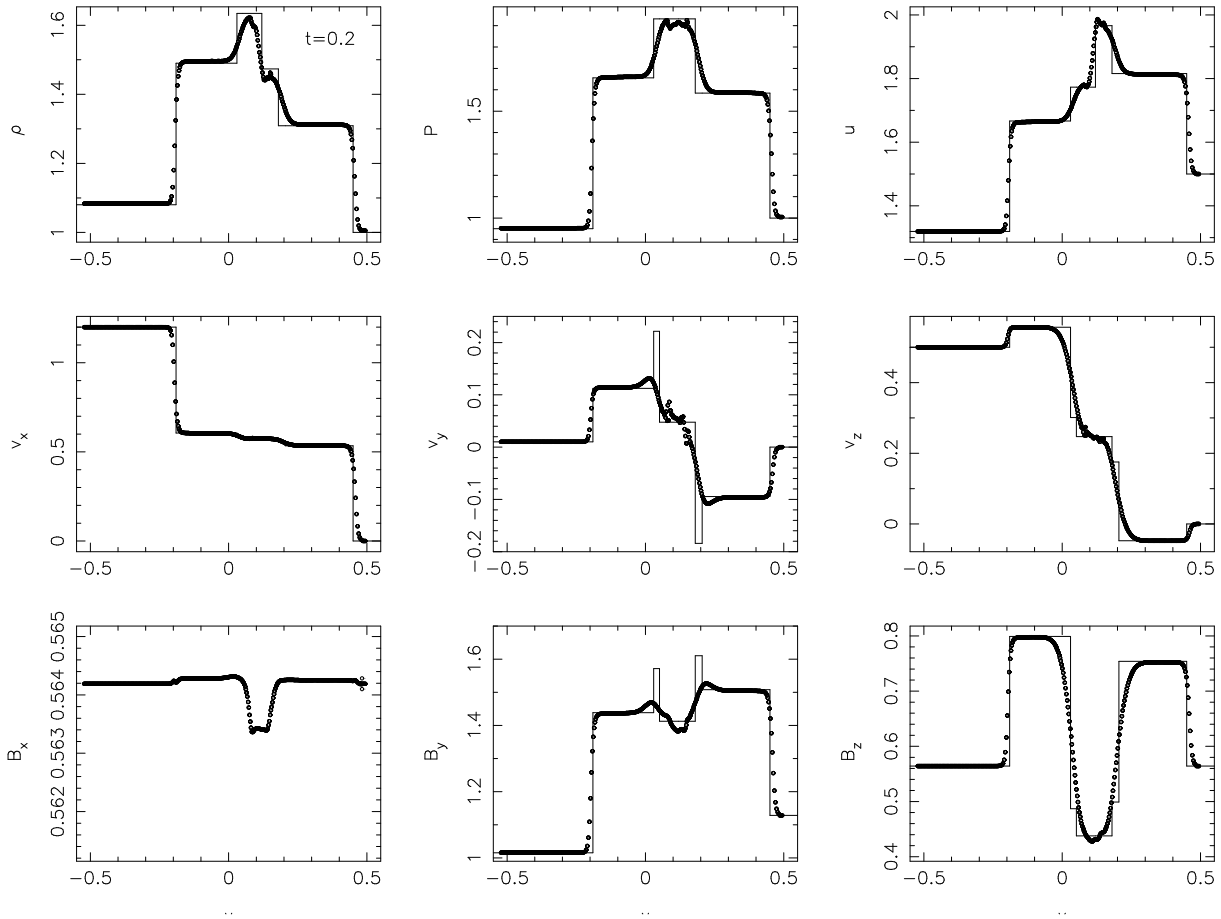

Figure 5.10: Results of the 2.5D shock tube test using an initial smoothing length of $h = 1.2(m/\rho)^{1/2}$ and using the total magnetic and kinetic energies in the dissipative terms as described in §4.5. The oscillations in the transverse velocity components observed in Figure 5.8 are damped in this case by the presence of an additional shear term in the artificial viscosity.

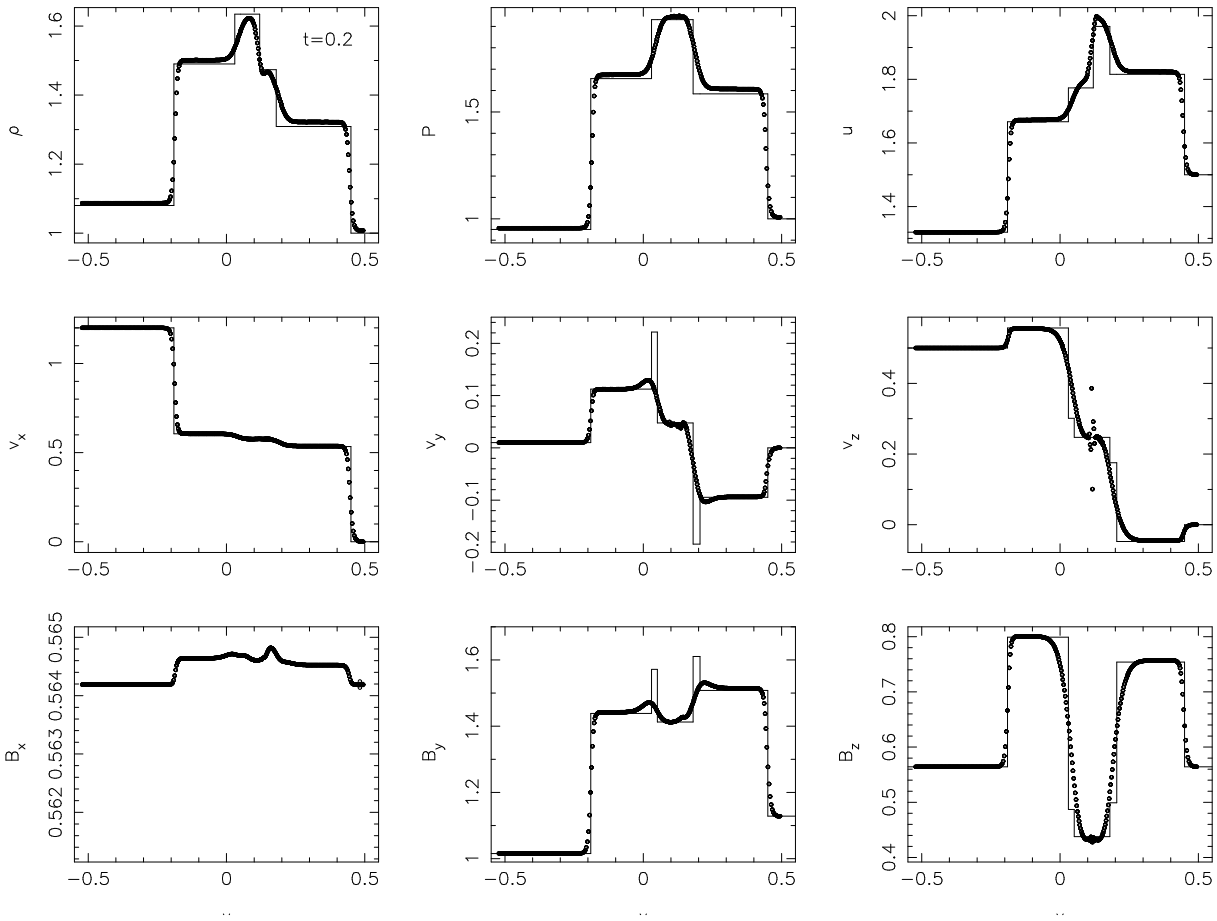

Figure 5.11: Results of the 2.5D shock tube test using a slightly larger initial smoothing length of $h = 1.5(m/\rho)^{1/2}$ and the total magnetic energy in the artificial resistivity term but using the usual artificial viscosity term. The results are a substantial improvement on those presented in Figure 5.8 for a very modest increase in the number of neighbours.

and with the one dimensional results shown in Figure 4.15. In the two dimensional case the transverse velocity components exhibit small oscillations near the contact discontinuity. It should be noted first of all that these oscillations are quite small and do not appear to affect the dynamics significantly (mainly because the jumps in the transverse velocity components are an order of magnitude less than the jump in v_x). However, the oscillations appear to result from a combination of three factors: the unsmoothed initial conditions, the fact that we do not explicitly apply any smoothing to the transverse velocity components and the effects of the small jumps in the $x-$component of the magnetic field.

To remove these oscillations two approaches can be taken: The first approach is to modify the artificial viscosity terms slightly in order to smooth the transverse velocity profiles. The dissipative terms used in order to capture shocks were discussed at length in §3.5, §4.5 and in the one dimensional shock tube tests described in §4.6.3. In the one dimensional case the dissipation terms for MHD (comprising an artificial viscosity, artificial thermal conductivity and artificial resistivity) were derived assuming that jumps would only occur in components of the magnetic field transverse to the line joining the particles that jumps in velocity would only occur parallel to this line. Neither of these assumptions strictly hold in the shock tube problem shown in Figure 5.8 since the transverse velocity components clearly jump and there is also a small jump in the parallel field component due to the divergence errors. A reformulation of the dissipative

terms relaxing both of these assumptions was presented in §4.5.1, deriving the artificial viscosity and artificial resistivity terms from jumps in the total kinetic and magnetic energies respectively in the total energy equation. The effects of using these formulations on the profile of the parallel component of the magnetic field are shown in Figure 5.9. From the centre panel we see that using the total magnetic energy formulation for the artificial resistivity has clear advantages in preventing oscillations in the parallel component of the field at shock fronts. Using the total kinetic energy version of the artificial viscosity (in order to smooth out jumps in the transverse velocity) effectively adds an explicit shear component to the viscosity term. In §4.5.1 it was noted that discontinuities in the transverse velocity components can only occur at corresponding jumps in the magnetic field and therefore that such discontinuities are already smoothed somewhat by the application of artificial resistivity there. For this reason the total kinetic energy formalism was *not* used in one dimension. The results using this formalism for the two dimensional problem are shown in Figure 5.10 in which we see that the oscillations are quite effectively damped. In this case the shear viscosity term has been applied in a minimal way by using the usual artificial viscosity switch (§3.5.2) which responds to $(-)\nabla \cdot \mathbf{v}$ (although since the jumps in transverse velocity are small even the minimum level of $\alpha = 0.1$ away from the shocks is sufficient to damp the oscillations seen). Adding an explicit shear viscosity is, however, highly undesirable since it increases the spurious transport of angular momentum caused by the artificial viscosity term.

The second approach is to simply increase the number of neighbours slightly for each particle to give a more accurate interpolation. The results using an initial smoothing length of $h = 1.5(m/\rho)^{1/2}$ are shown in Figure 5.11 using the total magnetic energy formulation of the artificial resistivity but retaining the usual artificial viscosity formulation. In this case the jump in the parallel field component is much lower and the oscillations in the transverse velocity components do not appear, although there is a small glitch at the contact discontinuity similar to that observed in the one dimensional case (§4.6.3). Increasing the smoothing length from $h = 1.2(m/\rho)^{1/2}$ to $h = 1.5(m/\rho)^{1/2}$ corresponds to an increase in the number of neighbours from ≈ 20 to ≈ 28 on a uniform cubic lattice in two dimensions. This quite a small increase in computational expense for a substantial gain in accuracy (and stability). It therefore seems much more desirable to increase the smoothing length slightly for multidimensional problems rather than to explicit add a shear viscosity term.

Finally, although this problem is not unstable to the clumping instability we have also investigated the effects of various instability correction methods on the shock profile. In particular use of the anticlumping term (§4.4.1) was found to produce additional noise in the shock profile. Using either the Morris formalism for the anisotropic force (§4.4.2) or subtracting the constant component of the magnetic field (§4.4.4) both give results very similar to those shown in Figures 5.8-5.11.

5.3.5 Two dimensional shock tube

The second shock tube test is used by both Tóth (2000) and Dedner et al. (2002) in two dimensions to compare the results of various divergence cleaning schemes, although the one dimensional version of this test has been used by many authors (e.g Dai and Woodward, 1994; Ryu and Jones, 1995). The results of the one dimensional test using the SPMHD algorithm were presented in §4.6.3 (Figures 4.18 and 4.19). Although this is a purely two dimensional test we present it after the 2.5D shock tube since it presents a much more challenging problem with regards to the non-zero divergence of the magnetic field due to the

stronger shocks.

The particle setup is as described in the previous section, except that the initial left state is given by $(\rho, P, \mathrm{v}_x, \mathrm{v}_y, B_y) = [1, 20, 10, 0, 5/(4\pi)^{1/2}]$ and the right state is $(\rho, P, \mathrm{v}_x, \mathrm{v}_y, B_y) = [1, 1, -10, 0, 5/(4\pi)^{1/2}]$ with $B_x = 5.0/(4\pi)^{1/2}$ and $\gamma = 5/3$. The boundaries are correspondingly adjusted in the $x-$direction to allow the particles to fill the domain $-0.5 < x < 0.5$ at $t_{max} = 0.08$. Particles are arranged initially on a cubic lattice with particle spacing 0.004, giving 660 particles in the $x-$direction and a total particle number of $660 \times 4 = 2640$. As in the previous test, the results using an initial smoothing length of $h = 1.2(m/\rho)^{1/2}$ exhibit significant oscillations in the transverse velocity (v_y). In this case the oscillations are substantially worse because the jump in the parallel field component is much larger. Hence we have performed this test using $h = 1.5(m/\rho)^{1/2}$. However, even in this case the oscillations remain present and so we have also added the shear viscosity term, using (4.92) with $\alpha = 1$ everywhere (that is, not using the viscosity switch). The results using these settings are shown in Figure 5.12 and may be compared with the exact solution taken from Dai and Woodward (1994) (solid line) and with the one dimensional results shown in Figure 4.18. Even in this case some oscillations are visible in the v_y profile, corresponding exactly with a spike in $\nabla \cdot \mathbf{B}$. In the $h = 1.2(m/\rho)^{1/2}$ case this spike is much larger [$(\nabla \cdot \mathbf{B})_{max} \sim 40$], causing significantly more disruption to the velocity profile. Thus despite the various tweaks we have attempted for this test, the oscillations appear to be primarily caused by the divergence errors generated at the shocks.

The effects of increasing the number of neighbours and changing the strength of the dissipation terms may be summarised as follows: Increasing the number of neighbours reduces the jumps in the parallel field component (for example with $h = 1.2(m/\rho)^{1/2}$ the jump is given by $\Delta B_x = [B_{x(max)} - B_{x(min)}]/B_{x0} \approx$ 18% whilst for $h = 1.5(m/\rho)^{1/2}$ we have $\Delta B_x \approx 3\%$ and for $h = 1.6(m/\rho)^{1/2}$ this reduces further still to $\Delta B_x \approx 1\%$). On the other hand, adding dissipation at the jumps in parallel field means that although such jumps may be present, the discontinuities (causing strong divergence errors) are smoothed. The effect of adding the shear viscosity term is to increase the dissipation at these discontinuities, thus reducing to some extent the associated spike in the magnetic divergence.

In Tóth (2000) the results of this test were presented using the source term approach of Powell et al. (1999) (discussed in §5.2.1), showing similar jumps in the parallel magnetic field component which were unchanged even in the converged numerical results. The fact that the jumps in parallel field reduce with an increasing number of neighbours indicates that the SPMHD algorithm converges to the exact solution in the limit of $h \to \infty$ and $N \to \infty$ where N is the number of particles. Tóth (2000) attributes the errors in the parallel field components in the Powell method to the non-conservative source terms in the induction equation. We have also performed this simulation using the 'conservative' induction equation (5.38), however we find that the jumps in B_x are not changed significantly by including the $\mathbf{v}\nabla \cdot \mathbf{B}$ term (although contain substantially more numerical noise). We attribute this difference to the fact that we use a non-conservative[6] formulation only in the induction equation, unlike in the Powell method where non-conservative forms are also used in the momentum and energy equations.

The shock tube tests presented above have been computed without using any form of divergence cleaning (other than the consistent formulation of the MHD equations in the presence of magnetic monopoles discussed in §5.2.1). Thus a way of eliminating both the jumps in parallel field and the

[6] where 'non-conservative' means that the volume integral of the flux (4.17) is not conserved exactly.

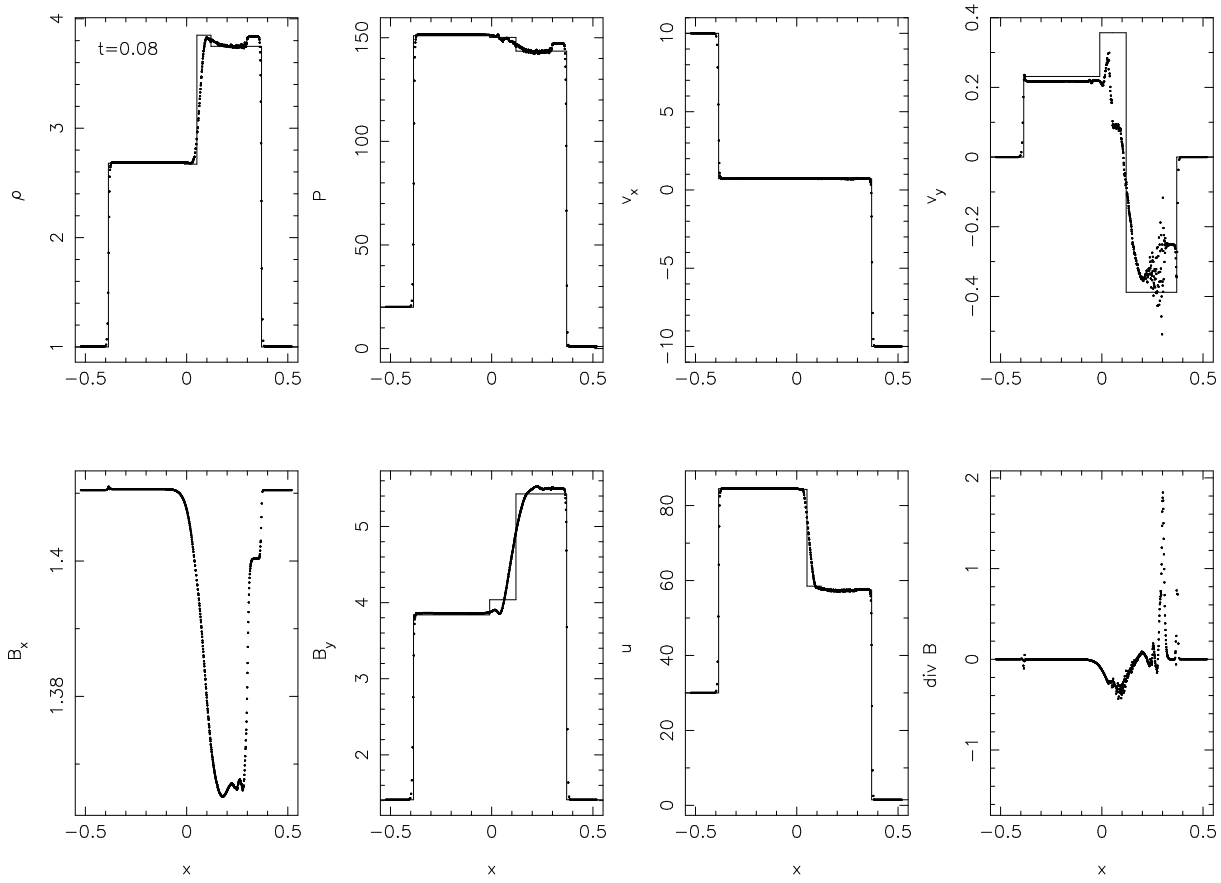

Figure 5.12: Results of the two dimensional shock tube test at $t = 0.08$ using $h = 1.5(m/\rho)^{1/2}$ and the shear viscosity term. The results may be compared with the one dimensional results shown in Figure 4.18 and the exact solution given by the solid line. In this stronger shock tube problem the jumps in the parallel field can cause significant oscillations in the transverse velocity components due to the non-zero divergence terms. Increasing the number of neighbours acts to reduce the jumps in the B_x component of the magnetic field, whilst increasing the dissipation spreads these discontinuities such that the resulting divergence errors are lower.

resulting oscillations in the transverse velocity components is to clean up the divergence error. Using the hyperbolic/parabolic cleaning discussed in §5.2.3 is not particularly effective for this problem, since the divergence errors are propagated away from their source at the fastest wave speed which is similar to the rate at which they are created by the shocks. Thus the diffusion introduced by the parabolic term does not have time to eliminate the divergence error before oscillations in the velocity components are produced. This is illustrated in Figure 5.13 which shows the results using this type of cleaning with $\sigma = 0.1$ on the parabolic term (c.f. §5.3.2). The divergence errors are reduced by a factor of ≈ 2 compared to the results shown in Figure 5.12. In order to eliminate the divergence errors from problems such as this one where divergence errors are created rapidly it would be better to use the projection method (§5.2.2). The projection method is somewhat complicated to implement in this case, however, because of the periodic boundary conditions (although this would not be the case using an iterative scheme as discussed in §5.2.2). An alternative would be to use the 'superfast' hyperbolic cleaning discussed in §5.2.3.

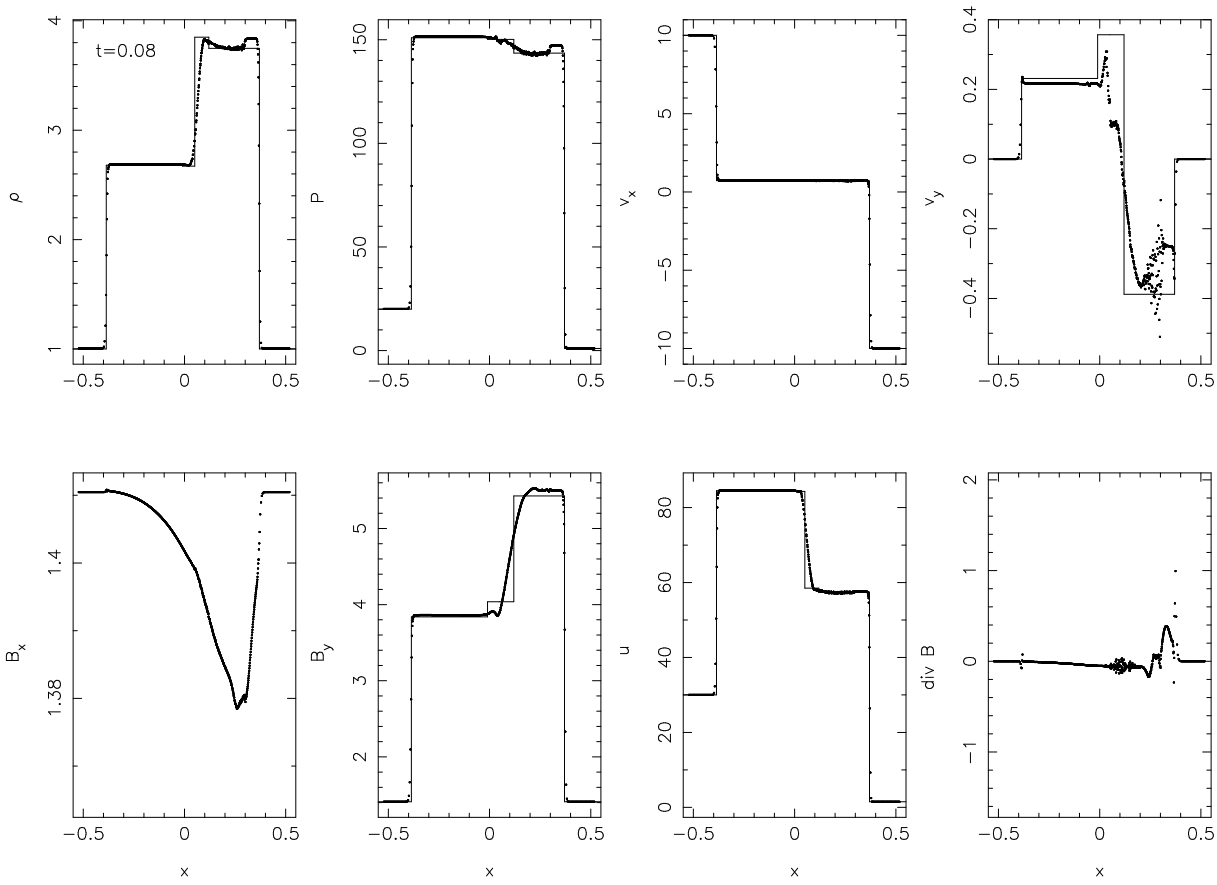

Figure 5.13: Results of the two dimensional shock tube test at $t = 0.08$ computed as in Figure 5.12 but using the hyperbolic/parabolic divergence cleaning (§5.2.3). The exact solution is given by the solid line. The hyperbolic divergence cleaning does not have a large effect on this problem since the divergence errors are propagated at the fastest wave speed which is similar to the rate at which they are generated in the shocks.

5.3.6 Spherical blast waves

Balsara (1998) gives a test involving an adiabatic blast wave propagating in a magnetic medium. Initially the pressure is set to 1000 in a spherical region of radius $r = 0.05$ around the origin in a uniform density box with $P = 1$ elsewhere. The density is initially unity and in the simulation shown we use $\gamma = 1.4$. A constant, uniform field of strength 10G (in code units $B_x = 10/\sqrt{4\pi}$) is setup in the x-direction. We setup this problem using 100×100 particles initially arranged on a cubic lattice in the domain $-0.5 < x < 0.5, -0.5 < y < 0.5$. The results at $t = 0.02$ are shown in Figure 5.14 and may be compared with the numerical solution given in Balsara (1998). The SPMHD results compare very well with the Balsara (1998) solution. In particular the contours of density and pressure show very little scatter, although there are some small effects visible due to the regularity of the initial particle setup.

5.3.7 Orszag-Tang vortex

The final two dimensional test is the compressible Orszag-Tang vortex problem which was first investigated by Orszag and Tang (1979) in order to study incompressible MHD turbulence. The problem was later extended to the compressible case by Dahlburg and Picone (1989) and Picone and Dahlburg (1991).

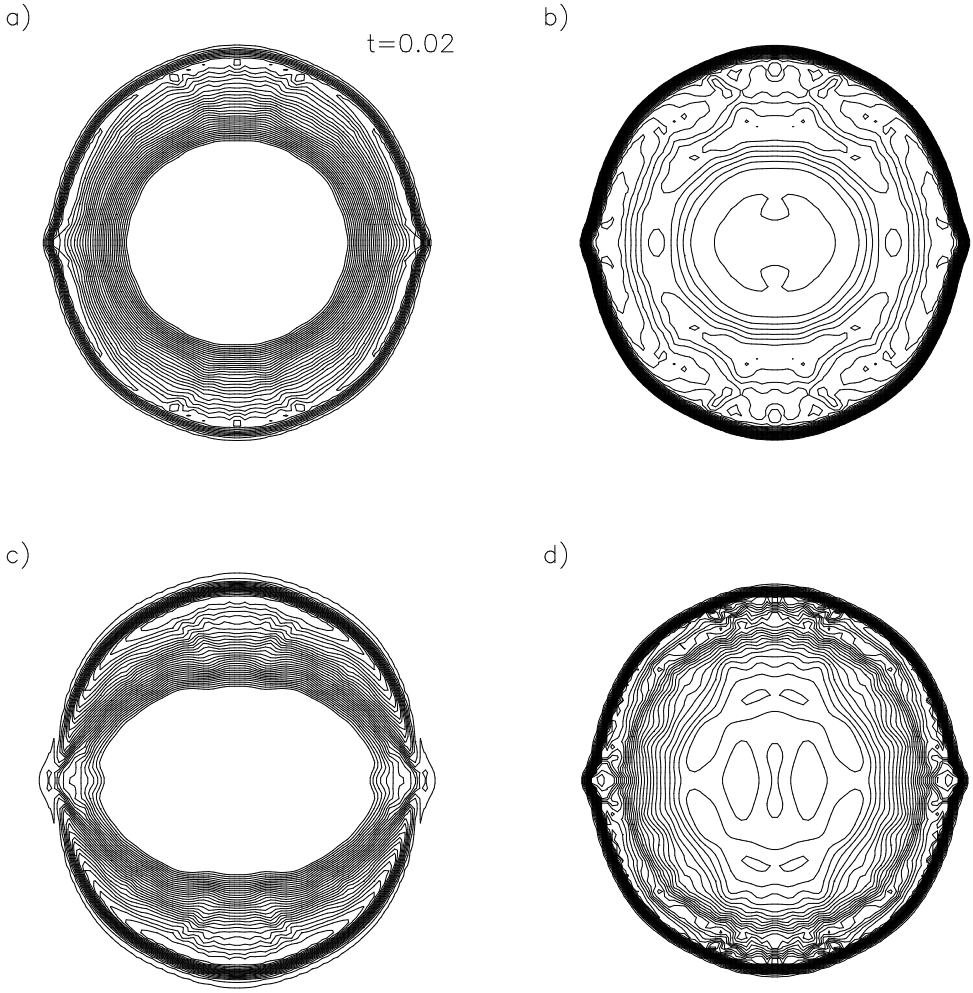

Figure 5.14: Spherical adiabatic MHD blast wave in two dimensions. Plots show: a) logarithm to base 10 of the density; b) logarithm to base 10 of the pressure; c) logarithm to base 10 of the magnetic pressure; d) specific kinetic energy. All plots show 30 contours spaced evenly between the minimum and maximum values of the quantity shown. The results compare extremely well with those shown in Balsara (1998)

More recently it has been widely used as a test problem for multidimensional MHD algorithms (e.g. Ryu et al., 1995; Balsara, 1998; Dai and Woodward, 1998; Londrillo and Del Zanna, 2000; Tóth, 2000).

The setup consists of an initially uniform density, periodic 1×1 box given an initial velocity perturbation $\mathbf{v} = \mathrm{v}_0[-\sin(2\pi y), \sin(2\pi x)]$ where $\mathrm{v}_0 = 1$. The magnetic field is given a doubly periodic geometry $\mathbf{B} = B_0[-\sin(2\pi y), \sin(4\pi x)]$ where $B_0 = 1/\sqrt{4\pi}$. The flow has an initial average Mach number of unity, a ratio of magnetic to thermal pressure of $10/3$ and we use $\gamma = 5/3$. The initial gas state is therefore $P = 5/3B_0^2 = 5/(12\pi)$ and $\rho = \gamma P/\mathrm{v}_0 = 25/(36\pi)$. Note that the choice of length and time scales differs slightly between various implementations in the literature. The setup used above follows that of Ryu et al. (1995) and Londrillo and Del Zanna (2000).

The particles are arranged initially on a uniform hexagonal close packed lattice. This distribution means that the particles are isotropically arranged and is the distribution towards which other arrangements naturally settle. However, results are similar using a cubic lattice arrangement. The simulation is performed using 128×146 particles (where the number of particles in the $y-$direction is determined

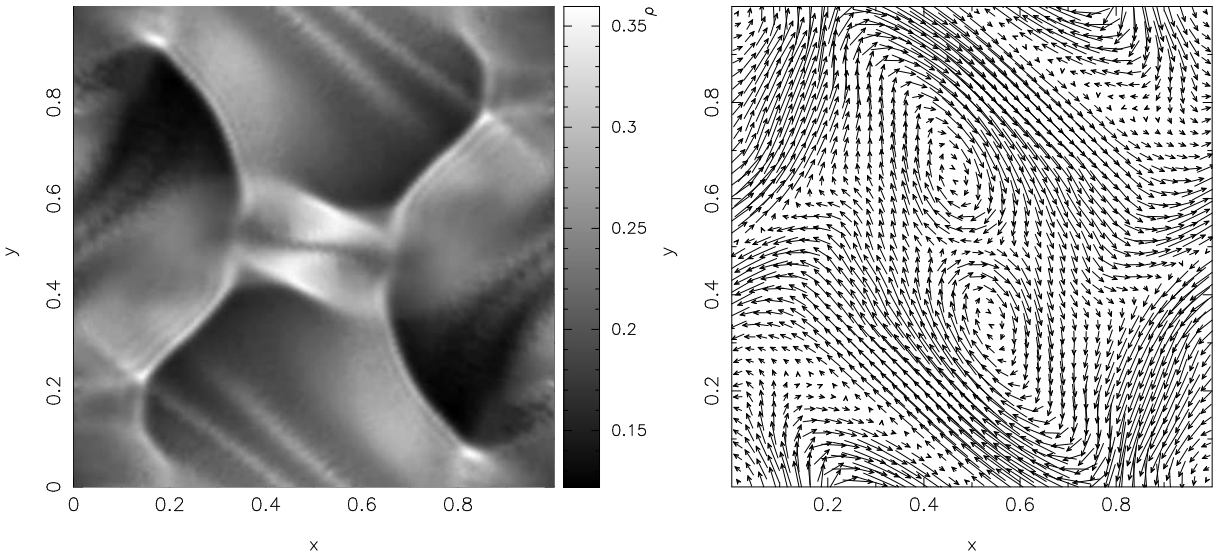

Figure 5.15: Results of the two dimensional Orszag-Tang vortex test, showing the density (left) and magnetic field (right) distribution at $t = 0.5$. The simulation uses 128×146 particles initially arranged on an isotropic hexagonal lattice with periodic boundary conditions. The initial velocity field is a large vortex $\mathbf{v} = [-\sin(2\pi y), \sin(2\pi x)]$ whilst the magnetic field has a doubly periodic geometry $\mathbf{B} = B_0[-\sin(2\pi y), \sin(4\pi x)]$. The SPMHD results are in good qualitative agreement with those presented in (e.g.) Dai and Woodward (1998) and Tóth (2000) although there are some small effects visible in the SPMHD solution due to the distortion of the initial regular lattice arrangement.

by the isotropic lattice arrangement) and the periodic boundary conditions are implemented using ghost particles. Note that this is near the lowest resolution used in Dai and Woodward (1998) (although in SPH the resolution is concentrated preferentially towards regions of high density). The dissipation terms are applied using the artificial viscosity switch and applying the artificial resistivity uniformly. However the artificial thermal conductivity has been turned off for this problem to increase the density resolution. The wall heating effect which the artificial thermal conductivity prevents are discussed in §3.7.3 and are very minor. No shear viscosity term has been used. Simulations of this problem which have been run with or without the variable smoothing length terms, using the Morris formalism for the magnetic force (§4.4.2), evolving either $\mathbf{B}$ or $\mathbf{B}/\rho$ and either the thermal or total energy show essentially no difference in the numerical results.

The results of the density and magnetic field evolution are shown in Figure 5.15 at $t = 0.5$. At this time four shocks are visible which have interacted in the central regions after having crossed the periodic domain. The SPMHD results are in good qualitative agreement with those presented in (e.g.) Dai and Woodward (1994, 1998) and Tóth (2000). In particular the central regions appear to be better resolved than in the 128×128 fixed-grid simulation of Dai and Woodward (1998), although the lower density regions are correspondingly less well resolved. The SPMHD solution shows some small residual effects due to the distortion of the initial regular particle arrangement, noticeable as small ripples behind the shock fronts in Figure 5.15 and a slightly mottled appearance in the low density regions. In Figure 5.15 we have used a smoothing length of $h = 1.5(m/\rho)^{1/2}$ which was found, as in the previous test, to give a substantial improvement in the numerical results over smaller values. In particular the effects from the distortion of the initial lattice are much larger using $h = 1.2(m/\rho)^{1/2}$. With the artificial thermal conductivity term included, the narrow ridges in the density visible near the top and bottom of Figure

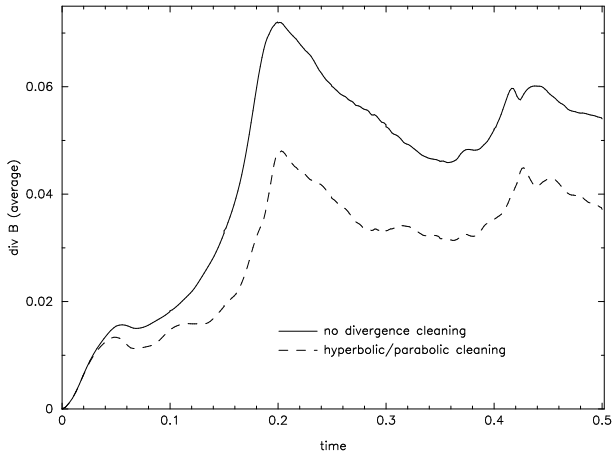

Figure 5.16: Evolution of the average magnetic divergence over the particles in the two dimensional Orszag-Tang vortex problem. Using the hyperbolic divergence cleaning (dashed line) produces only a slight improvement over the results with no divergence cleaning (solid line). The single biggest factor determining the magnitude of the divergence error is the number of neighbours. The results shown are for a smoothing length of $h = 1.5(m/\rho)^{1/2}$.

5.15 are largely smoothed out.

The evolution of the average of the magnetic divergence is shown in Figure 5.16 for two runs with and without divergence cleaning. The results using the hyperbolic/parabolic cleaning with $\sigma = 0.1$ (dashed line) show only a slight improvement ($\sim 30\%$ reduction in the average divergence) over the results with no divergence cleaning (solid line). In fact the single biggest factor which determines the magnitude of the divergence error is the number of neighbouring particles. For example in a simulation using $h = 1.2(m/\rho)^{1/2}$ the divergence errors are approximately twice those shown in Figure 5.16.

5.4 Summary

In this chapter multidimensional aspects of the SPMHD algorithm have been discussed. In particular several methods for maintaining the divergence-free constraint in an SPH context have been presented. Firstly the source term approach of Powell et al. (1999) was outlined and contrasted with the consistent formulation of the MHD (and SPMHD) equations derived in §4.3.2. The major difference between the two approaches is that our approach retains the conservation of momentum and energy whereas the Powell et al. approach does not. The conservation properties of the induction equation were also discussed, in which it was highlighted that using a 'non-conservative' induction equation means that the surface integral of the magnetic flux is conserved, rather than the volume integral. The effect of using the consistent formulation of the MHD equations in the presence of magnetic monopoles (which conserves the surface integral of the flux) is that divergence errors are advected without change by the flow (illustrated in Figure 5.1).

Projection methods for maintaining a divergence free field were discussed in an SPH context in §5.2.2. In particular it was noted that using the Green's function solution to the Poisson equation (as is often used for self-gravity in SPH) provides only an approximate projection. The results using this type of projection on a problem where an initial magnetic divergence was introduced into the simulation were nonetheless very good (§5.3.2). The disadvantages are the substantial computational cost introduced by the solution

of a Poisson equation and for many of the test problems presented in this chapter, the complication introduced by periodic boundary conditions. The potential advantages of using an iterative solution to the Poisson equation were also discussed briefly.

An alternative approach to divergence cleaning suggested recently by Dedner et al. (2002) was discussed in §5.2.3. The method involves adding an additional constraint equation which is coupled to the induction equation for the magnetic field. Chosen appropriately, the effect of this equation is to cause the divergence errors to be propagated in a wave-like manner away from their source (Figure 5.2). Adding a small diffusive term means that the divergence errors are also rapidly reduced to zero. This method is extremely simple to implement and is computationally very inexpensive. The disadvantage is that the error propagation is limited by the timestep condition and, although much faster than using diffusion alone to reduce the divergence, for some problems (for example the shock tube tests given in §5.3.4 and §5.3.5) the cleaning is still not fast enough. However, this method is a substantial improvement over not using any form of divergence cleaning at a negligible additional computational cost.

The various approaches to divergence cleaning were compared in §5.3.2 using a simple test problem in which a non-zero divergence was introduced into the simulation as an initial condition. It was found that using the Dedner et al. (2002) cleaning on this problem could produce results similar to those obtained by taking a projection step every 10 timesteps. It was also noted that the projection method does not conserve the cross-helicity invariant whereas the hyperbolic/parabolic cleaning does.

The SPMHD algorithm was also tested against a variety of multidimensional test problems. A non-linear circularly polarized Alfvén wave was studied in §5.3.3. This test showed that SPMHD has a very low intrinsic numerical dissipation compared to grid based codes, although this property is destroyed by the addition of explicitly dissipative terms for shock-capturing which can cause quite slow convergence on problems where the physical dissipation timescale is of critical importance. Two of the shock tube problems used in the one-dimensional case were examined in two dimensions in §5.3.4 and 5.3.5. For these problems jumps in the component of the magnetic field parallel to the shock front (causing divergence errors) were found to result in oscillations in the transverse velocity profiles. The jumps in the parallel field component were found to decrease as the number of neighbours for each particle was increased, unlike in the Powell et al. method in which the jumps remain unchanged even in the numerically converged results (Tóth, 2000). The corresponding divergence errors produced by these jumps could be reduced by using a form of the dissipative terms derived in §4.5.1 using the total jump in magnetic and kinetic energies. Modifying the artificial viscosity term in this manner results in the addition of an explicit shear viscosity component. It is therefore somewhat undesirable to do so since this can result in excess spurious angular momentum transport elsewhere. A better approach would be to use divergence cleaning to prevent these errors from occurring. However, the hyperbolic cleaning was not found to be particularly effective for this problem because of the restriction to the fastest wave speed and implementation of the projection method is complicated by the periodic boundary conditions. These difficulties are not, however, insurmountable. The single biggest factor in determining the magnitude of the divergence errors in the shock tube tests was found to be the size of the smoothing region (ie. the number of contributing neighbours). It therefore seems advantageous to use a slightly larger number of neighbours for MHD problems (typically $h \gtrsim 1.5(m/\rho)^{1/\nu}$ where ν is the number of spatial dimensions) than might otherwise be used for hydrodynamics.

An initially spherical MHD blast wave test was given in §5.3.6, with good results. Finally the algo-

rithm was tested on the Orszag-Tang vortex problem (§5.3.7) which has been widely used as a benchmark for MHD codes. The SPMHD result were in good qualitative agreement with those presented elsewhere. This test again highlighted the need for a slightly larger number of neighbours, in this case to remove spurious effects related to the initial lattice arrangement and to reduce the magnitude of the divergence errors. The hyperbolic/parabolic divergence cleaning was found to produce only a small ($\sim 30\%$) reduction in the divergence errors, again highlighting the need for some form of sub-timestep cleaning (for example using the projection method).

Unfortunately there was neither the time nor the space in this thesis to benchmark the SPMHD algorithm against the many wonderful exact solutions which can be derived for multidimensional magnetic toy stars.