

Anisotropic Diffusion in Mesh-Free Numerical Magnetohydrodynamics

Philip F. Hopkins^{1*}

¹*TAPIR & The Walter Burke Institute for Theoretical Physics, Mailcode 350-17, California Institute of Technology, Pasadena, CA 91125, USA*

Submitted to MNRAS, January, 2016

ABSTRACT

We extend recently-developed mesh-free Lagrangian methods for numerical magnetohydrodynamics (MHD) to arbitrary anisotropic diffusion equations, including: passive scalar diffusion, Spitzer-Braginskii conduction and viscosity, cosmic ray diffusion/streaming, anisotropic radiation transport, non-ideal MHD (Ohmic resistivity, ambipolar diffusion, the Hall effect), and turbulent “eddy diffusion.” We study these as implemented in the code GIZMO for both new meshless finite-volume Godunov schemes (MFM/MFV). We show the MFM/MFV methods are accurate and stable even with noisy fields and irregular particle arrangements, and recover the correct behavior even in arbitrarily anisotropic cases. They are competitive with state-of-the-art AMR/moving-mesh methods, and can correctly treat anisotropic diffusion-driven instabilities (e.g. the MTI and HBI, Hall MRI). We also develop a new scheme for stabilizing anisotropic tensor-valued fluxes with high-order gradient estimators and non-linear flux limiters, which is trivially generalized to AMR/moving-mesh codes. We also present applications of some of these improvements for SPH, in the form of a new integral-Godunov SPH formulation that adopts a moving-least squares gradient estimator and introduces a flux-limited Riemann problem between particles.

Key words: methods: numerical — hydrodynamics — MHD — instabilities — conduction — diffusion

1 INTRODUCTION

Diffusion operators are ubiquitous in physical systems, manifest in e.g. conduction, viscosity, cosmic ray diffusion/streaming, photon diffusion (in optically thick media), sub-grid “turbulent eddy” diffusion, passive scalar (e.g. metal) diffusion, Ohmic resistivity, ambipolar diffusion, the Hall effect, and more. Most of these are fundamentally anisotropic. Magnetic fields, fluid flow, or radiation flux can all break isotropy and introduce a (local) preferred direction. The effects are large – for the same initial conditions but different magnetic field configurations, effects like viscosity or conduction might dominate the hydrodynamics, or be completely suppressed. Moreover, anisotropy can introduce qualitatively new behaviors and physical instabilities (e.g. Balbus 2000; Quataert 2008).

Clearly, it is important to capture these phenomena accurately in numerical simulations. In regular, non-moving mesh methods (including fixed-grid and adaptive mesh-refinement [AMR]), the properties of diffusion equations are reasonably well-understood, but even in these cases it is highly non-trivial to formulate anisotropic operators in a numerically stable fashion (see, e.g. Sharma & Hammett 2007). But for some classes of problems, these methods are sub-optimal. They tend to produce excessive diffusion when a fluid is moving rapidly relative to the grid, especially across contact discontinuities (Tasker & Bryan 2008; Springel 2010; Hopkins 2015); have difficulty coupling to N-body gravity methods and handling self-gravitating hydrostatic equilibrium (Müller & Steinmetz 1995; LeVeque 1998; Zingale et al. 2002); introduce low-order errors around refinement boundaries (O’Shea et al. 2005; Heitmann et al. 2008); and feature inherently preferred directions which can introduce systematic errors even at high resolution if physical anisotropies are not aligned with the grid axes (e.g. Peery & Imlay 1988; Hahn et al. 2010; Hopkins 2015; Hopkins & Raives 2016).

Mesh-free methods can avoid these sources of error, and so are popular for many problems in astrophysics. However, the

most popular mesh-free method, smoothed-particle hydrodynamics (SPH) has its own difficulties. For example, SPH requires the use of artificial diffusion operators to maintain numerical stability, which can produce greater numerical viscosity and over-smoothing of shocks and discontinuities compared to Riemann-based shock solvers (Cullen & Dehnen 2010); it tends to suppress fluid-mixing instabilities (Morris 1996; Ritchie & Thomas 2001; Agertz et al. 2007); and suffers from zeroth-order errors (“E0” errors; Morris 1996; Dilts 1999; Read et al. 2010) which produce systematic errors that do not converge with resolution alone. Although there have been dramatic improvements in all of these (see Hopkins 2013; Rosswog 2014; Hu et al. 2014), the E0 errors cannot be completely eliminated from SPH without making the method numerically unstable (Price 2012b).

Recently, however, Lanson & Vila (2008a,b); Gaburov & Nittadori (2011); Hopkins (2015) have developed a class of new, mesh-free Lagrangian finite-volume methods which are both high-order consistent and fully conservative. Similar to moving-mesh codes (Springel 2010; Duffell & MacFadyen 2011; Gaburov et al. 2012), these new methods appear to capture many of the advantages of AMR and SPH, while avoiding the disadvantages above (although, of course, they feature their own sources of error such as enhanced “grid noise” and volume “partition noise”; see Hopkins 2015). In Hopkins 2015 (hereafter Paper I) and Hopkins & Raives (2016) (Paper II), these are developed for magnetohydrodynamics (MHD) in the multi-method magnetohydrodynamics+gravity code GIZMO,¹ built on the N-body gravity and domain decomposition algorithms from GADGET-3 (Springel 2005).

In this paper, therefore, we extend these new methods to include arbitrary anisotropic diffusion operators. We consider a systematic comparison of tests in order to determine the degree of anisotropy that can be reliably treated. In the Appendices, we show

¹ A public version of the code is available at <http://www.tapir.caltech.edu/~phopkins/Site/GIZMO.html>

* E-mail: phopkins@caltech.edu

how some parts of the new methods here can also be applied to SPH.

2 GENERAL DIFFUSION OPERATORS IN MESHLESS FINITE-VOLUME GODUNOV METHODS

The general diffusion equation in conservative form is:

$$\frac{d\mathbf{U}}{dt} = -\nabla \cdot \mathbf{F}_{\text{diff}} \quad (1)$$

$$\mathbf{F}_{\text{diff}} = -\mathbf{K} \cdot (\nabla \otimes \mathbf{q}) \quad (2)$$

where d/dt is a Lagrangian derivative, and \mathbf{F}_{diff} a diffusive flux given by a linear combination of the tensor field \mathbf{K} and the gradients of the field \mathbf{q} . We denote the inner product “ \cdot ” and outer product “ \otimes ”.

In the mesh-free finite-volume or finite-mass (MFV or MFM, respectively) formulations of the Lagrangian methods in Gaburov & Nitadori (2011) and Paper I, we begin from a general conservation equation of the form $DU/Dt = -\nabla \cdot \mathbf{F}$ and use this to derive a second-order accurate Godunov-type numerical expression of the hydrodynamic equations. It is therefore trivial to apply the same here, giving

$$\frac{d}{dt}(V\mathbf{U})_i = -\sum_j \mathbf{F}_{\text{diff},ij}^* \cdot \mathbf{A}_{ij} \quad (3)$$

where \mathbf{A}_{ij} is the “effective face area” (defined in Paper I; it depends on the inter-element spacing and kernel shape, and reduces to the geometric faces of a Voronoi tessellation in the limit of a delta-function kernel). The same would obtain for Cartesian grid methods and moving-mesh codes, with \mathbf{A}_{ij} the usual inter-face area. Here $\mathbf{F}_{\text{diff},ij}^*$ is the *interface* value of the flux.

The solution then follows our usual MHD method: we calculate the coefficients, perform a reconstruction step to estimate quantities “at the face,” replace the flux with the solution to a Riemann problem (RP), and use this to update the conserved quantities. We will explain this in more detail below.

Note that Eq. 3 is manifestly antisymmetric, so conserved quantities ($V\mathbf{U}$) are manifestly conserved as desired.

Additionally, we note that the MFM and MFV methods from Paper I differ only at second order in how they handle advection between particles. In diffusion-only problems, they are manifestly identical. Therefore, we do not show both (although we have confirmed their identical results), but will, for simplicity, adopt MFM as our reference method.

2.1 Gradient Estimation & Reconstruction

We require gradients for all quantities; to do this, we adopt the standard gradient estimator in GIZMO, a moving, second-order accurate and consistent least-squares estimator. For a scalar f , this is

$$(\nabla f)_i^\alpha = \sum_j (f_j - f_i) \left(\mathbf{W}_i^{-1} \right)^{\alpha\beta} (\mathbf{x}_j - \mathbf{x}_i)^\beta \omega_j(\mathbf{x}_i) \quad (4)$$

$$\mathbf{W}_i^{\alpha\beta} \equiv \sum_j (\mathbf{x}_j - \mathbf{x}_i)^\alpha (\mathbf{x}_j - \mathbf{x}_i)^\beta \omega_j(\mathbf{x}_i) \quad (5)$$

here we assume an Einstein summation convention over the indices β corresponding to the spatial dimensions, and $\omega_j(\mathbf{x}_i)$ is an (arbitrary) weight function defined in Paper I. This estimator is second-order accurate for an *arbitrary* mesh/particle configuration, minimizes the (weighted) least-squares deviation $\sum_j \omega_j |f_i + \nabla f_i \cdot (\mathbf{x}_j - \mathbf{x}_i) - f_j|^2$, and has been applied in a wide range of different numerical methods (see e.g. Oñate et al. 1996; Kuhnert 2003; Maron & Howes 2003; Luo et al. 2008; Lanson & Vila 2008a). Note that

Eq. 4 applies to a scalar field f . For a general tensor \mathbf{q} , we apply Eq. 4 separately to every component $f = q_{\alpha\beta\gamma\dots}$ to determine all partial derivatives needed to construct $\nabla \otimes \mathbf{q}$ and \mathbf{F}_{diff} . The gradients are slope-limited as described in Paper I, such that they do not create new local extrema within the interacting kernel.

In the reconstruction step we must extrapolate the values from i and j to the left and right sides of the face. For hydrodynamics and MHD, we perform a linear reconstruction of all the MHD quantities based on their slope-limited gradients (see Paper I). For additional quantities needed for diffusion, our default approach is a first-order reconstruction, e.g. $(\nabla \otimes \mathbf{q})_{R,L} = (\nabla \otimes \mathbf{q})_{i,j}$. This is easy to implement, and most stable, but comes at the cost of greater numerical diffusion. We have therefore also considered limited tests using the “double linear” reconstruction of Muñoz et al. (2013). This amounts to treating $\nabla \otimes \mathbf{q}$ like any other primitive variable. In a first loop, we calculate $\nabla \otimes \mathbf{q}$ with our standard method; in a second loop, we calculate $\nabla \otimes (\nabla \otimes \mathbf{q})$, and use this to linearly reconstruct $(\nabla \otimes \mathbf{q})_{R,L}$ (or any component of it) like any other primitive variable. This gives a simple, second-order Taylor-series representation which implicitly includes an appropriately larger stencil (since the second pass sums over the values $(\nabla \otimes \mathbf{q})_j$, themselves constructed from the \mathbf{q}_k in all neighbors of j). To ensure smoothness, the reconstructed gradients in the double-linear method are slope-limited with respect to the particle-centered gradients using the same limiter we adopt for all the usual MHD quantities (see Paper I).

2.2 The Riemann Problem

We treat diffusion in operator-split fashion from the pure-MHD RP. We have experimented with several Riemann solvers for the diffusion RP, and find the best compromise between numerical diffusion, stability, and flexibility using the Harten et al. (1983) (HLL) solver in the Lagrangian frame. In this case:

$$\mathbf{F}_{\text{diff},ij}^* = \frac{\lambda^+ \mathbf{F}_{\text{diff},L} - \lambda^- \mathbf{F}_{\text{diff},R} + \alpha \lambda^+ \lambda^- (\mathbf{U}_R - \mathbf{U}_L)}{\lambda^+ - \lambda^-} \quad (6)$$

where the maximum/minimum wavespeeds λ^+/λ^- are determined appropriate to the problem. Here $\mathbf{U}_R, \mathbf{U}_L$ are the appropriate right/left states reconstructed at the face following Paper I (extrapolated from the particle-centered values to the face with the derivatives of \mathbf{U} , with a MINMOD slope limiter).

Consider the case where we are using our second-order reconstruction in a Lagrangian frame. We adopt the wavespeed estimate from Paper I, $\lambda^+ = \text{MAX}(v_L, v_R) + c_{\text{fast}}$, $\lambda^- = \text{MIN}(v_L, v_R) - c_{\text{fast}}$. Here c_{fast} is a fastest wavespeed determined by the 1D RP. Then, because in our default GIZMO method our frame is moving with $v = (v_L + v_R)/2$, in the Lagrangian frame we simply obtain $\lambda^+ = -\lambda^- = |v_R - v_L|/2 + c_{\text{fast}}$, in which case the HLL solution reduces to the Global Lax Friedrich (GLF) function: $\mathbf{F}_{\text{diff},ij}^* = (\mathbf{F}_{\text{diff},L} + \mathbf{F}_{\text{diff},R})/2 - \alpha (|v_R - v_L|/2 + c_{\text{fast}}) (\mathbf{U}_R - \mathbf{U}_L)/2$.

Although using the GLF flux in *non*-Lagrangian methods tends to be excessively diffusive, here it gives nearly indistinguishable results from using the HLL solution with another wavespeed estimate (e.g. Roe wavespeeds or the exact eigenvalues of the Jacobian $\partial \mathbf{F}/\partial \mathbf{U}$). This is because the first-order term owing to frame motion is automatically accounted for by the Lagrangian nature of the method. Even if we use a double-linear reconstruction, we see only a tiny (percent-level) increase in diffusion using this particular form for the RP solution.

Using the above with $\alpha = 1$ gives an effective numerical diffusivity $\kappa_n \sim c_s |\mathbf{x}|_{ij}$, which can easily exceed the physical $\kappa_p \sim \|\mathbf{K}\|$ by large factors at low resolution. We prevent this by replacing

Eq. 6 with the flux-limited equation:

$$\mathbf{F}_{\text{HLL}} = \text{MINMOD}((1 + \psi)\mathbf{F}_2, \mathbf{F}_2 + \mathbf{F}_U) \quad (7)$$

$$\mathbf{F}_2 = \frac{\lambda^+ \mathbf{F}_{\text{diff},L} - \lambda^- \mathbf{F}_{\text{diff},R}}{\lambda^+ - \lambda^-} \quad (8)$$

$$\mathbf{F}_U \equiv \alpha \left[\frac{\lambda^+ \lambda^- (\mathbf{U}_R - \mathbf{U}_L)}{\lambda^+ - \lambda^-} \right] \quad (9)$$

$$\alpha \equiv \left(\frac{\|\mathbf{K}^* \cdot (\nabla \otimes \mathbf{q})^*\|}{\|\mathbf{K}^*\| \|(\nabla \otimes \mathbf{q})^*\|} \right) \left(\frac{0.2 + r}{0.2 + r + r^2} \right) \quad (10)$$

$$r \equiv \left(\frac{|v_R - v_L|}{2} + c_{\text{fast}}^* \right) \frac{\|\mathbf{x}_j - \mathbf{x}_i\| \|\mathbf{q}^*\|}{\|\mathbf{K}^*\| \|\mathbf{U}^*\|} \quad (11)$$

where above \mathbf{F}_2 is the flux we would obtain in the simplest 2nd-order accurate (but numerically unstable) formulation. \mathbf{F}_U is the diffusive flux from the Riemann problem, appropriately limited by the function α . Here $\|\mathbf{x}\|$ refers to the Frobenius norm of the tensor \mathbf{x} , and f^* refers to the interface value of the primitive variable f ; where possible we adopt these from the appropriate solution to the pure-MHD RP, otherwise we approximate $f^* = (f_L + f_R)/2$.

The first term in α (in \mathbf{K} and $\nabla \otimes \mathbf{q}$) limits \mathbf{F}_U for the anisotropic case, vanishing as F_{diff} does where there is full anisotropic suppression (even at low resolution). Note that the form of $\mathbf{K} \cdot (\nabla \otimes \mathbf{q})$ in Eq. 10 follows directly from the form of $\mathbf{F}_{\text{diff}} = \mathbf{K} \cdot (\nabla \otimes \mathbf{q})$. For the cases where F_{diff} is represented by a different linear combination of gradients, the term in Eq. 10 should be modified accordingly.

The second term in α ensures $\alpha \rightarrow 1$ at small r and $\alpha \rightarrow 1/r$ at large r (the functional form is motivated by Rosdahl & Teyssier 2015); together with the MINMOD application this prevents the numerical diffusivity from exceeding physical values by more than some tolerance parameter ψ . As usual this parameter represents some tradeoff: increasing ψ gives smoother solutions at the expense of numerical diffusivity. We find our qualitative results are robust for all $0.05 < \psi < 1$, and use $\psi = 0.1$ as our default.

Finally, we also compute a “direct” flux based on pairwise direct-difference gradients, and use this to restrict $\mathbf{F}_{\text{diff},ij}^*$ via:

$$(\mathbf{F} \cdot \mathbf{A})_{\text{direct}} \equiv - \left[\frac{\mathbf{K}_i + \mathbf{K}_j}{2} \cdot (\nabla \otimes \mathbf{q})_{\text{dir}} \right] \cdot \|\mathbf{A}_{ij}\| \frac{\mathbf{x}_j - \mathbf{x}_i}{\|\mathbf{x}_j - \mathbf{x}_i\|} \quad (12)$$

$$(\nabla \otimes \mathbf{q})_{\text{dir}} \equiv \frac{(\mathbf{x}_j - \mathbf{x}_i) \otimes (\mathbf{q}_j - \mathbf{q}_i)}{\|\mathbf{x}_j - \mathbf{x}_i\|^2} \quad (13)$$

$$\mathbf{F}_{\text{diff},ij}^* = \begin{cases} \mathbf{0} & (\text{SIGN}[(\mathbf{F} \cdot \mathbf{A})_{\text{direct}}] \neq \text{SIGN}[\mathbf{F}_{\text{HLL}} \cdot \mathbf{A}_{ij}] \\ & \text{and } \|(\mathbf{F} \cdot \mathbf{A})_{\text{direct}}\| > \epsilon \|\mathbf{F}_{\text{HLL}} \cdot \mathbf{A}_{ij}\| \\ \mathbf{F}_{\text{HLL}} & \text{otherwise} \end{cases} \quad (14)$$

with the tolerance parameter $0 \leq \epsilon \leq 1$. We find stable (albeit slightly more noisy) results in all our problems using values as large as $\epsilon = 2$, in fact, and larger values do give improved performance at low resolution on multi-dimensional tests (e.g. the diffusing ring), but we adopt $\epsilon = 1/2$ as our default here.

We note that this is not the only way to stabilize the anisotropic diffusion equations. For example, Sharma & Hammett (2007) propose an elegant slope-limiting method; however, it is not obvious how to extend this to unstructured meshes. Recently, Kannan et al. (2015) implemented the method of Gao & Wu (2013) for moving-mesh codes. The advantage of their method is that it is extremum-preserving and generalizes relatively easily to implicit integrators. However, we choose to explore alternatives for three reasons. First, the Gao & Wu (2013) method implicitly uses the lower-order “direct” gradients, as opposed to our (in principle arbitrarily high-order) matrix-based gradient estimator, neces-

sarily making the method lower-order. Moreover in flows where there is a clear mean-field gradient but large resolution-scale noise, matrix-based estimators are significantly more robust (see García-Senz et al. 2012; Mocz et al. 2014; Pakmor et al. 2016). Second, our method here allows for non-linear flux-limiting terms (e.g. α above) which allow us to limit numerical diffusion to physical values even at arbitrarily low resolution (potentially important in multi-physics problems where the diffusion may, in some places, dominate only on un-resolved scales). And third, our method here, unlike most in the literature, allows for any (arbitrarily complex) tensor \mathbf{q} , and/or linear combinations of \mathbf{K} and the elements of $\nabla \otimes \mathbf{q}$ in the fluxes (relevant for e.g. radiation transport, Braginskii viscosity, and the Hall effect).

2.3 Timestepping

In addition to the usual timestep limiters (e.g., the CFL condition, gravitational acceleration-based limiters) which always apply, ensuring numerical stability in explicit methods for diffusion equations requires an additional timestep criterion:

$$\Delta t \leq \frac{1}{\|\mathbf{K}\|} \left[\frac{\|\nabla \otimes \mathbf{q}\|}{\|\mathbf{q}\|} + \frac{1}{\Delta x} \right]^{-2} \frac{\|\mathbf{U}\|}{\|\mathbf{q}\|} \quad (15)$$

For the simplest diffusion example, $\partial U / \partial t = \kappa \partial^2 U / \partial x^2$ and an appropriately slope-limited gradient, this reduces to the common expression $\Delta t \leq \Delta x^2 / \kappa$. We also require that the inter-particle flux between a particle and any neighbors of the conserved $V\mathbf{U}$ cannot exceed half the minimum of the $|V\mathbf{U}|$ in the pair, in a single timestep, but find this criterion is always satisfied if the above (more strict) criterion is as well.

3 EXAMPLES IMPLEMENTED

We implement this method in the code GIZMO. A number of specific physical cases have been implemented; our focus here is not on the microphysics but on the numerical methods. Still, it is useful to list the relevant examples, both to explicitly see how they correspond to the general formulation above, and to give examples of the different forms of anisotropy that are physically expected.

3.0.1 Isotropic Passive-Scalar Diffusion

For this simple case, we have

$$\mathbf{K} = \kappa \mathbf{I} \quad (16)$$

$$\mathbf{q} = \mathbf{U} = n_{\text{scalar}} \quad (17)$$

where κ is the (arbitrary) diffusion coefficient, \mathbf{I} the identity matrix, and n_{scalar} the scalar density.

3.0.2 Anisotropic Thermal Conduction

Isotropic and anisotropic thermal conduction are represented by:

$$\mathbf{K} = \kappa_{\perp} \mathbf{I} + \kappa_{\parallel} \hat{\mathbf{B}} \otimes \hat{\mathbf{B}} \quad (18)$$

$$\mathbf{q} = T \quad , \quad \mathbf{U} = \rho u \quad (19)$$

where T is the gas temperature, ρ the gas density and u the internal energy per unit mass, κ_{\perp} the isotropic diffusion coefficient, κ_{\parallel} the anisotropic (parallel) diffusion coefficient, and $\hat{\mathbf{B}} \equiv \mathbf{B}/|\mathbf{B}|$ the direction of the local magnetic field. Here we allow either arbitrarily specified κ , or default to the Spitzer conductivity (for κ_{\parallel} , with $\kappa_{\perp} = 0$) with a limiter for cases where the implied heat transport rate exceeds the free electron streaming rate.

4 Hopkins

3.0.3 Anisotropic Cosmic Ray Diffusion & Streaming

Cosmic rays diffuse along magnetic field lines. Moreover, as shown by Uhlig et al. (2012), cosmic-ray streaming can be represented numerically via a diffusion operator. Therefore we have, for the simple case of a single-species cosmic-ray population

$$\mathbf{K} = (\kappa_{\text{diff}} + \kappa_{\text{stream}}) \hat{\mathbf{B}} \otimes \hat{\mathbf{B}} \quad (20)$$

$$\mathbf{q} = P_{\text{cr}} \quad , \quad \mathbf{U} = e_{\text{cr}} \quad (21)$$

where P_{cr} is the cosmic ray pressure (proportional for a single-species model to the CR number density), $e_{\text{cr}} \approx P_{\text{cr}}/(\gamma_{\text{cr}} - 1) = \rho u_{\text{cr}}$ is the cosmic ray energy density, and κ_{diff} and κ_{stream} are the appropriate “effective” coefficients for diffusion and streaming processes, respectively.

3.0.4 Turbulent Eddy Diffusion Models

In popular “sub-grid” models for turbulence, the effect of unresolved eddies is treated as a diffusion process. Following the Smagorinsky (1963) model, we model this for e.g. diffusion of a scalar field via

$$\mathbf{K} = \rho (C \Delta x)^2 \|\mathbf{S}\| \mathbf{I} \quad (22)$$

$$\mathbf{q} = u_{\text{scalar}} = \frac{n_{\text{scalar}}}{\rho} \quad , \quad \mathbf{U} = n_{\text{scalar}} \quad (23)$$

where Δx is the resolution scale (for our meshless methods, defined to be equal to the rms inter-element spacing inside the kernel), $C \sim 0.1$ is a constant calibrated to numerical simulations in Smagorinsky (1963), and \mathbf{S} is the symmetric shear tensor.

3.0.5 Anisotropic Radiation Transport in the Diffusion Limit

In the optically thick limit, the radiative transfer moment equations are commonly expressed as a diffusion equation. It is convenient to represent this in the form

$$\mathbf{F}_{\text{diff}} = -\mathbf{K} \cdot (\nabla \cdot \mathbf{q}) \quad (24)$$

$$\mathbf{K} = \frac{\lambda c}{\kappa_{\nu} \rho} \mathbf{I} \quad (25)$$

$$\mathbf{q} = n_{\nu} \mathbb{D}_{\nu} \quad , \quad \mathbf{U} = n_{\nu} \quad (26)$$

where c is the speed of light, κ_{ν} the opacity at frequency ν , n_{ν} the photon number (or energy) density, λ an optional “radiative flux limiter,” and \mathbb{D}_{ν} the dimensionless Eddington tensor. Note the $\nabla \cdot \mathbf{q}$ instead of $\nabla \otimes \mathbf{q}$ has no effect on our methodology, it simply re-orders the gradient terms in \mathbf{F}_{diff} . Various numerical methods (e.g. flux-limited diffusion and “variable Eddington tensor” methods) rely on this description.

3.0.6 Anisotropic Viscosity

For a viscous fluid, we have

$$\mathbf{q} = \mathbf{v} \quad , \quad \mathbf{U} = \rho \mathbf{v} \quad (27)$$

where \mathbf{v} is the velocity and ρ the density. In this case the form of \mathbf{F}_{diff} depends on the viscous parameterization. In the case of a magnetized fluid, the Braginskii viscosity can be written

$$\mathbf{F}_{\text{diff}} = -\mathbf{K} \left[\hat{\mathbf{K}} : (\nabla \otimes \mathbf{q}) \right] \quad (28)$$

$$\mathbf{K} = 3\nu_{\parallel} \left(\hat{\mathbf{B}} \otimes \hat{\mathbf{B}} - \frac{1}{3} \mathbf{I} \right) \quad (29)$$

$$\hat{\mathbf{K}} = \frac{\mathbf{K}}{3\nu_{\parallel}} = \hat{\mathbf{B}} \otimes \hat{\mathbf{B}} - \frac{1}{3} \mathbf{I} \quad (30)$$

while for the un-magnetized case, it is common practice to decompose the viscosity into shear (η) and bulk (ζ) terms, following

$$\mathbf{F}_{\text{diff}} = \Pi_{\eta} + \Pi_{\zeta} \quad (31)$$

$$\Pi_{\eta} = -\mathbf{K}_{\eta} \left[\nabla \otimes \mathbf{q} + (\nabla \otimes \mathbf{q})^T - \frac{2}{3} (\nabla \cdot \mathbf{q}) \right] \quad , \quad \mathbf{K}_{\eta} = \eta \mathbf{I} \quad (32)$$

$$\Pi_{\zeta} = -\mathbf{K}_{\zeta} (\nabla \cdot \mathbf{q}) \quad , \quad \mathbf{K}_{\zeta} = \zeta \mathbf{I} \quad (33)$$

Here “:” denotes the double-dot-product. As above, the particular arrangement of \mathbf{K} and $\hat{\mathbf{K}}$ within the double-dot-product (for the magnetized case) or decomposition of $\nabla \otimes \mathbf{v}$ into shear/bulk terms have no effect on our methodology; they simply re-order the gradients of \mathbf{q} within \mathbf{F}_{diff} . As with conduction we allow either freely-specified η , ζ , ν_{\parallel} , or can calculate the coefficients according to Spitzer-Braginskii theory with the appropriate limiters. We also add the corresponding viscous term to the energy equation.

3.0.7 Non-Ideal MHD

Astrophysical non-ideal MHD effects are typically parameterized as Ohmic dissipation (controlled by the resistivity η_o), the Hall effect (η_H) and ambipolar diffusion (η_A). All appear as diffusion operators in the induction equation; if we operator-split the ideal MHD term (already solved in GIZMO), we have

$$\frac{d\mathbf{B}}{dt} = -\nabla \times \left[\eta_o \mathbf{J} + \eta_H (\mathbf{J} \times \hat{\mathbf{B}}) - \eta_A (\mathbf{J} \times \hat{\mathbf{B}}) \times \hat{\mathbf{B}} \right] \quad (34)$$

where $\mathbf{J} = \nabla \times \mathbf{B}$. We can, with some elaborate algebra, write this as an equation in $\nabla \cdot \mathbf{F}_{\text{diff}}$, but it is easier to cast this directly into the alternative Godunov form:

$$\frac{d}{dt} (V \mathbf{U})_i = - \sum_j \mathbf{A}_{ij} \times \mathbf{F}_{\text{diff},ij}^* \quad (35)$$

$$\mathbf{F}_{\text{diff}} = -[\mathbf{K}_O + \mathbf{K}_H + \mathbf{K}_A] \cdot (\nabla \times \mathbf{q}) = -\mathbf{K} \cdot (\nabla \times \mathbf{q}) \quad (36)$$

$$\mathbf{K}_O = \eta_o \mathbf{I} \quad (37)$$

$$\mathbf{K}_A = \eta_A (\mathbf{I} - \hat{\mathbf{B}} \otimes \hat{\mathbf{B}}) \quad (38)$$

$$\mathbf{K}_H = \eta_H \begin{bmatrix} 0 & \hat{B}_z & -\hat{B}_y \\ -\hat{B}_z & 0 & \hat{B}_x \\ \hat{B}_y & -\hat{B}_x & 0 \end{bmatrix} \quad (39)$$

$$\mathbf{q} = \mathbf{B} \quad , \quad \mathbf{U} = \mathbf{B} \quad (40)$$

where $\hat{B}_x, \hat{B}_y, \hat{B}_z$ are the components of $\hat{\mathbf{B}}$ and the coefficients are calculated from the plasma properties. The appropriate magnetic terms are added to the energy equation.

4 NUMERICAL TESTS

4.1 Diffusing Sheet

For a first test in Fig. 1, we study a simple discontinuity in one dimension, $\mathbf{q} = \mathbf{q}_L$ for $x < 0$, $\mathbf{q} = \mathbf{q}_R$ for $x > 0$, with all other properties fixed. To focus on the diffusion equations, we turn off all MHD forces (for now). We also simulate this in a 3D box (with an arbitrarily large periodic domain), because although it is essentially a 1D problem, sensitivity to particle arrangement makes accurate solution much more challenging in 2D or 3D. Finally, we also add noise (which is critical to test for numerical instability): we add a uniform random value to q between $\pm 0.05 q_L$, for all particles.




Figure 1. Diffusing sheet test (§ 4.1). A 3D box ($N = 64$ elements across plotted x -range) is initialized with a step-function discontinuity ($x = 0$) in diffusing quantity \mathbf{q} . The solution is scale-free in units shown, and independent of the physics (we confirm identical behavior for implementations of passive-scale diffusion, conduction, cosmic rays, turbulent eddy diffusion, radiation transport in the diffusion limit, viscosity, and Ohmic resistivity; § 3). We consider isotropic and anisotropic cases with diffusion tensor $\mathbf{K} = K(\hat{\mathbf{B}} \otimes \hat{\mathbf{B}})$, with parallel, partially-aligned, and perpendicular fields \mathbf{B} . In each case we compare our mesh-free finite-element methods (MFM, our MFV method is manifestly identical on these problems; § 2) to exact solutions.

4.1.1 Isotropic Case

In the isotropic case, after the noise has been (quickly) damped away,² this has the trivial analytic solution

$$q(x, t) = \frac{q_R + q_L}{2} + \frac{q_R - q_L}{2} \text{Erf} \left[\frac{x}{\sqrt{4\kappa t}} \right] \quad (41)$$

The solution is completely self-similar so the absolute values of all quantities are irrelevant; the only numerically important quantity is the number of resolution elements over which the contact discontinuity has been diffused $\sim \sqrt{4\kappa t}/\Delta x$ (as we evolve the solution, it becomes better-resolved). We therefore plot results at fixed $\sqrt{4\kappa t}/\langle \Delta x \rangle$.

Our results converge well to the analytic solution with both methods even at low resolution ($\sqrt{4\kappa t}/\langle \Delta x \rangle \gtrsim 4$). Because the initial noise is grid-scale, it should be damped away on a small timescale $\sqrt{4\kappa t}/\langle \Delta x \rangle \sim 1$; we confirm this.

Note that the solution here is independent of the physics; we have explicitly verified that our implementations of passive-scalar diffusion ($n \propto q(x, t)$ above), conduction ($T \propto q$ with $\mathbf{K} = \kappa_{\perp} \mathbf{I}$), cosmic rays ($e_{\text{cr}} \propto q$, with $\mathbf{K} = \kappa_{\text{diff}} \mathbf{I}$), eddy diffusion ($n \propto q$), isotropic radiation transport ($n_{\nu} \propto q$ with $\mathbb{D}_{\nu} = \mathbf{I}/3$), viscosity ($\mathbf{B} = 0$ with $v_x = v_z = 0$, $v_y \propto q$, and $\zeta = 0$), and Ohmic resistivity ($B_y \propto q$) all give *identical* results on this test (given the same diffusivity), as they should.

² We have considered both a perfect step-function initial condition, and (alternatively) initializing the profile with the analytic solution corresponding to time t_0 such that $\sqrt{4\kappa t_0}/\langle \Delta x \rangle \approx 2$. These are indistinguishable after a short time ($\sqrt{4\kappa t_0}/\langle \Delta x \rangle \gtrsim 4$).

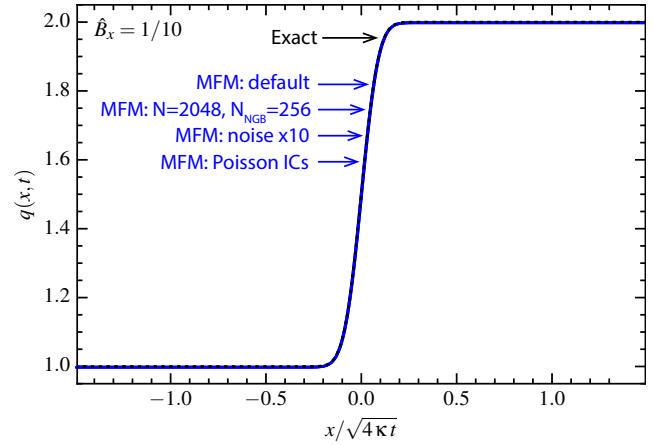


Figure 2. Anisotropic sheet (as Fig. 1) with varied ICs (§ 4.1.3). We compare (1) A “default” case with resolution $N = 64$ particles across the plotted range in the x -direction (≈ 4 particles across the range of diffusion at this time, in MFM), and effective neighbor number $N_{\text{NGB}} = 32$ within the kernel partition function. Particles are initially in a cubic lattice, with noise in \mathbf{q} equal to 2.5% of the jump value. (2) High-resolution ($N = 2048$) and neighbor number ($N_{\text{NGB}} = 256$, with a quintic kernel function). (3) Initial noise = 25% of jump value. (4) Initial particles laid down randomly (according to a Poisson distribution). Our MFM results are indistinguishable from each other and the exact solution in each case.

4.1.2 Anisotropic Case

Next we consider the anisotropic case. Here we take $\mathbf{K} = \kappa \hat{\mathbf{B}} \otimes \hat{\mathbf{B}}$ with κ and $\hat{\mathbf{B}}$ constant and $|\hat{\mathbf{B}}| = 1$. In this case, the solution is identical to the isotropic case but with $\kappa \rightarrow \kappa |\hat{\mathbf{B}} \cdot \hat{\mathbf{x}}|^2 = \kappa \hat{B}_x^2$, so it is entirely specified by the absolute value of the projection of $\hat{\mathbf{B}}$ in the gradient ($\hat{\mathbf{x}}$) direction, $|\hat{B}_x|$.

We have considered ~ 100 values between $-1 \leq \hat{B}_x \leq 1$ to check for pathological behavior; these are summarized with three representative cases shown here: $|\hat{B}_x| = 0$ (perpendicular fields, which should completely suppress diffusion), $|\hat{B}_x| = 1$ (parallel fields; the solution should be identical to the isotropic case), and $|\hat{B}_x| = 1/\sqrt{2}$. For all scalar diffusion cases (passive scalars, cosmic rays, eddy diffusion, conduction) this gives identical results. For our radiation diffusion, we make the system anisotropic by instead taking $\mathbb{D} = \hat{n} \otimes \hat{n}$, with $\hat{n} = \hat{\mathbf{B}}$ constant. We confirm this produces exactly identical solutions to the cases above.

Our MFM method is able to handle all three cases accurately; we confirm that there is zero diffusivity for the perpendicular case (up to machine precision, if our initial \mathbf{q} depends on x alone) and that the parallel case exactly matches the isotropic case. Convergence is again good even at low resolution $\sqrt{4\kappa t}/\Delta x \gtrsim 4$, and (we show below) the results are insensitive to noise or particle order.

4.1.3 Dependence on Particle Order, Resolution, Noise, and Neighbor Number

Fig. 2 considers variations of the anisotropic sheet. In order to test whether our methods are sensitive to the local arrangement of particles, we have considered (in both 2D and 3D tests) an initial particle distribution following (1) a regular square lattice, (2) uniformly randomly-distributed particles over the volume (a Poisson distribution), (3) a glass (generated from the random distribution), and (4) a densest sphere packing. In 2D we have also considered a regular triangular and hexagonal grid. For our MFM/MFV method, the results are almost indistinguishable (after the initial noise is damped) in every one of these cases (even the “worst-case” Poisson distribution, shown in Fig. 2), clearly demonstrating that the method does not depend sensitively on particle order.

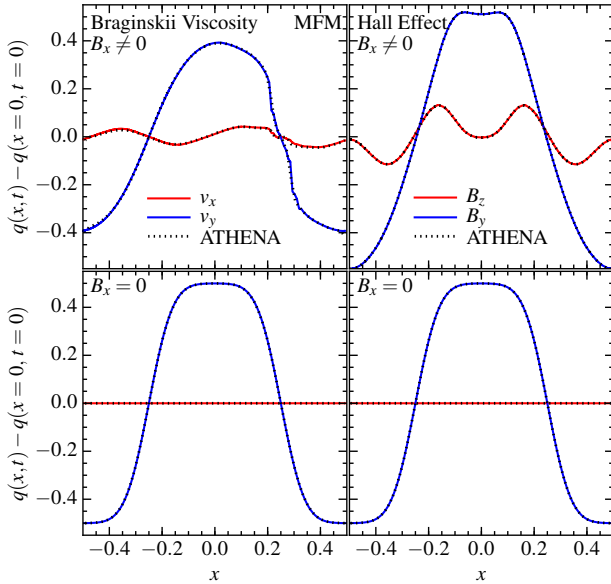


Figure 3. Anisotropic diffusing sheet problem, as Fig. 1 ($N = 64$), but here focusing on two special cases where the solution is not the same as in Fig. 1 (see § 4.1.4). There is no analytic solution for the cases shown, so, we compare a converged solution from ATHENA. *Left:* Braginskii viscosity (§ 3.0.6). The initial discontinuity is in v_y only, but this form of viscosity generates $v_x \neq 0$ at later times. We consider two cases: $\hat{B} = (1, 1, 0)/\sqrt{2}$ (*upper*; this should produce non-zero diffusion) and $\hat{B} = (0, 1, 0)$ (*lower*; this should have vanishing diffusion). *Right:* Hall effect (§ 3.0.7). The initial discontinuity is in B_y ($B_x = \text{constant}$, B_z vanishes). We consider a case which should diffuse ($B_x \neq 0$; *upper*) and one which should be suppressed ($B_x = 0$; *lower*). MFM and ATHENA agree well in all cases.

We consider resolution tests, varying both the absolute resolution and number of kernel neighbors. Our MFM/MFV results are well-converged even with just $\sim 4 - 8$ resolution elements across the jump (i.e. $(\sqrt{4\kappa t})/\Delta x \gtrsim 4$). The convergence rate of our method is also independent of neighbor number, as shown in Paper I and Paper II (going to larger neighbor number simply trades a reduction in noise for increased numerical diffusivity).

We have also tested for sensitivity to initial noise, varying the seed noise level from 0 – 25% of the jump level. In all cases our results are stable and do not qualitatively change with respect to this. Likewise, the steepness of the initial discontinuity has no effect on our results (whether we begin with the full solution at a resolved scale, or a perfectly steep discontinuity across a single particle).

4.1.4 Braginskii Viscosity and the Hall Effect

Two cases in § 3 are more complicated, even in a simple diffusing sheet setup. These are Braginskii viscosity (§ 3.0.6) and the Hall effect (§ 3.0.7). In both, \mathbf{q} is a vector (\mathbf{v} or \mathbf{B} , respectively), and even if we initialize a gradient only in one element of that vector and set the other elements to vanish everywhere, the form of the diffusion operator leads to growth of other components of \mathbf{q} . Mathematically, the anisotropic part of the diffusion operator in most cases is a projection operator; here, it also includes a rotation operator. This leads to different non-linear solutions and makes it more challenging to achieve stability.

Therefore we consider these cases specifically in Fig. 3. Since the non-linear solutions do not have analytic forms even for this simple test, we compare to a converged, high-resolution ($N = 2048$ across the plotted domain), one-dimensional solution from the well-tested grid-based code ATHENA (Stone et al. 2008). We run ATHENA in its most accurate constrained-transport, PPM-CTU

(highest-order) mode. For simplicity, we do not disable the other hydrodynamic forces, but these are not dominant. To avoid certain boundary condition effects in ATHENA we initialize the test problem in slightly different fashion: we take

$$q = 1.5 - 0.5 \left(1 + \text{Erf} \left[\frac{x-0.25}{0.01} \right] - \text{Erf} \left[\frac{x+0.25}{0.01} \right] \right) \quad (42)$$

For the Braginskii problem, we take initial $\mathbf{v} = (0, q, 0)$ and consider both $\hat{B} = (1, 0, 0)$ (which should produce zero diffusion) and $\hat{B} = (1, 1, 0)/\sqrt{2}$ (which should diffuse); we take $\rho = u = 1$, and $|\mathbf{B}| = 10^{-6}$ to be small so it has no effect on the dynamics except to control the anisotropy. For the Hall problem, we take vanishing initial velocities and $\mathbf{B} = 10^{-6}(0, q, 0)$ (no diffusion) $\mathbf{B} = 10^{-6}(1, q, 0)$ (diffusion). Our MFM simulations are run in 3D boxes with resolution $N = 128$ across the plotted domain.

Fig. 3 shows these develop non-zero v_x and B_z , despite these vector components initially vanishing. In both cases our MFM/MFV methods produce solutions in excellent agreement with ATHENA, even well into non-linear evolution. Both cases with diffusion, and cases with full anisotropic suppression, are captured.

4.2 Diffusion Across a Moving, Rotated, Shearing Contact Discontinuity

We next consider a more challenging variation of the diffusion sheet, illustrated in Fig. 4. We **(1)** re-enable our normal MHD physics. **(2)** Insert a contact discontinuity at $x = 0$ (the location of the jump in \mathbf{q}) with a density jump of a factor of 2, so the diffusion (of some passive scalar) is across the discontinuity. This implies a different particle arrangement (since our particles are equal-mass) across $x = 0$. **(3)** Make the discontinuity shearing, with $v_y = x/2$. This means the particle geometry around the diffusing interface is constantly being re-arranged. **(4)** Uniformly boost the system by $\mathbf{v}_{\text{boost}} = (10, 3, 2)c_s$. **(5)** Rotate the entire system by $+35^\circ$.

None of these changes the physical solution for the diffusion. However they are all, in principle, numerically challenging. Because our methods are fully Lagrangian, our solutions are trivially invariant to the boost **(4)** and rotation **(5)** operations. This is not the case, however, for Eulerian methods. Our MFM method is also invariant to **(1)** and **(2)**, i.e. the evolution of a stable contact discontinuity produces vanishing fluxes, and the gradient estimator is explicitly insensitive to particle arrangement within the kernel for linear gradients (this is not always the case in other Lagrangian methods such as SPH, where solutions around contact discontinuities can depend on particle arrangement; see references in § 1).

Fig. 4 shows that, despite these complications, our MFM solution still agrees very well with the exact result. There is some noise around $x = 0$, introduced by the shear **(3)**, particularly around the contact discontinuity, but it is not visible on the scale plotted. This constant re-arrangement of the particles (hence constant re-arrangement of the effective faces and implicit mesh) introduces “grid noise” (for detailed discussion of this noise term see Duffell & MacFadyen 2014; Hopkins 2015; Mocz et al. 2015; Pakmor et al. 2016).

A detailed comparison with grid-based methods is outside the scope of this paper; however, this problem is very challenging for such codes. A moving contact discontinuity, especially one not aligned with the grid, produces large numerical diffusivity at low resolution (see Hopkins 2015; Springel 2010). Using ATHENA to run a version of this problem with anisotropic conduction, we find the numerical diffusivity exceeds the physical diffusion even in the isotropic case unless we use more than ~ 512 elements across the

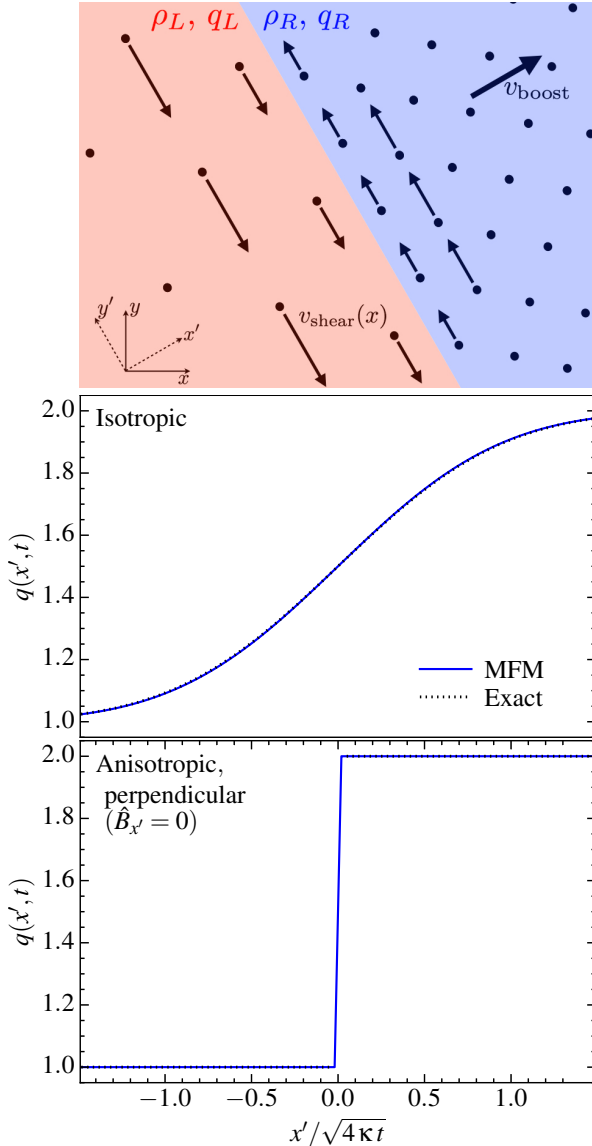


Figure 4. Diffusion across a super-sonically moving, rotated, shearing contact discontinuity (§ 4.2), as Fig. 1, with full MHD physics enabled. *Top:* Illustration of the initial problem setup. We measure properties along the x' axis, perpendicular to the initial discontinuity in density and \mathbf{q} . We consider an isotropic case (*middle*) and an anisotropic, perpendicular (fully-suppressed; *bottom*) case. The contact discontinuity means the particle arrangement is different on either side of $x' = 0$, and the shear field means that the arrangement of neighbor particles is constantly changing. These do not destroy our solution.

plotted range (compared to 64 here); even higher resolution is required for the anisotropic case.

4.3 Sinusoidal Temperature Distribution

Fig. 5 considers a one-dimensional test problem in which a scalar q follows a sinusoidal distribution from Arth et al. (2014). This tests the same physics as the later stages of the diffusing sheet, but is much “easier” (since it is 1D and there is never a steep gradient). No other MHD physics are active, and we take a periodic box of unit length $L = 1$, with unit density and sound speed and $\gamma = 5/3$, with the physical solution

$$q(t, x) = \frac{3}{2} + \sin(2\pi x) \exp\left(-4\pi^2 K \hat{B}_x^2 t\right) \quad (43)$$




Figure 5. Diffusion of a sinusoidal pulse ($\mathbf{q} \propto \sin(2\pi x)$; § 4.3), in a 1D periodic box with $0 < x < 1$ with $N = 64$ elements across the x -range. We show anisotropic cases with perpendicular (full suppression) and parallel (no suppression) fields. In the perpendicular case, the initial pulse does not evolve; in the parallel case, the amplitude of the pulse decays as $\exp(-\pi^2 4 K t)$ (we show results at $4 K t = 1/5$).

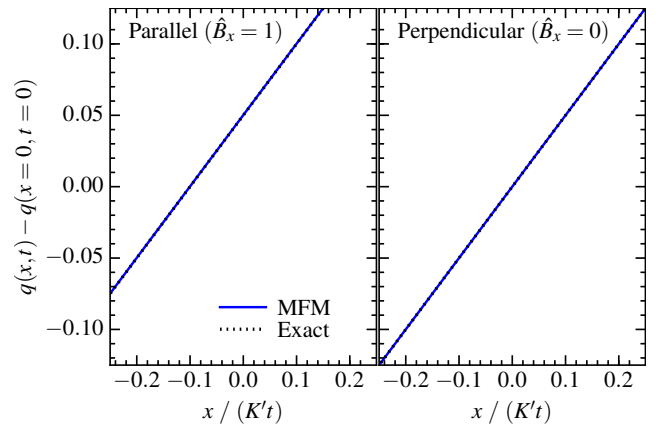


Figure 6. Diffusion with a variable diffusivity \mathbf{K} and diffused quantity \mathbf{q} (§ 4.4). Both vary linearly in \hat{x} . The box is 3D with $N = 64$ elements across the x -range, but the results are converged (up to tested $N = 4096$). We consider anisotropic cases with parallel (*left*; no suppression) and perpendicular (*right*; full suppression) fields.

initialized at $t = 0$. We have confirmed (as expected) that all the conclusions from our diffusing sheet tests apply in this test.

4.4 Variable Diffusivity

We previously took \mathbf{K} to be constant. Here, consider the case $q(x, t = 0) = q_0 + q'x$, $K(x, t) = K_0 + K'x$; this produces the analytic solution $q(x, t) = q(x, t = 0) + q'K't$. In the simple anisotropic case with $\mathbf{K} = K\hat{B} \otimes \hat{B}$ this becomes $q(x, t) = q(x, t = 0) + |\hat{B}_x|^2 q'K't$. We consider this in the same 3D box setup as before (here with large enough distance in the x -direction so that the boundary conditions do not enter the considered domain).

The results are shown in Fig. 6. Our methods perform well in both the anisotropic case and (more trivially, therefore not shown), the isotropic case.

4.5 Gaussian Pulse

We now consider a multi-dimensional problem – the diffusion of a quantity injected as a δ -function instantaneously into a homogeneous background. In a periodic box of unit size, we initialize a 3D, spherically-symmetric Gaussian for \mathbf{q} centered on the origin, where the diffused quantity is treated as a passive scalar with all other background properties constant (so no other MHD effects appear).

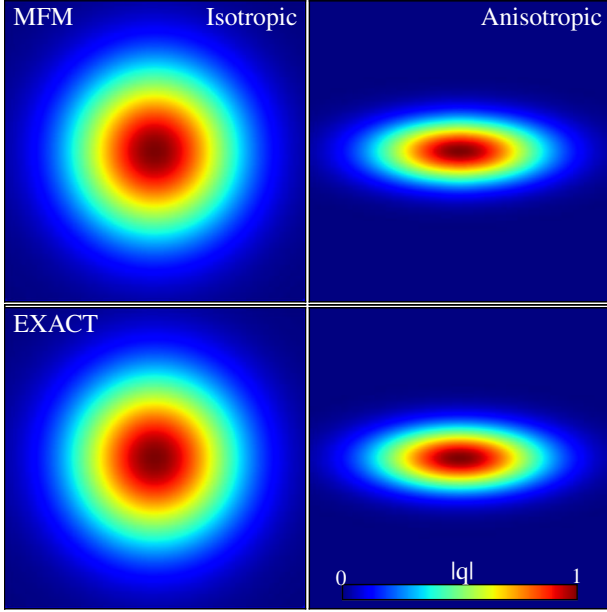


Figure 7. Three-dimensional Gaussian pulse test (§ 4.5). A Gaussian distribution of the diffused quantity \mathbf{q} is injected at the center with initial spherically symmetric width $\epsilon = 0.05$. We plot a slice through the $x-y$ plane, at time $t = 0.8$ for diffusivity $K = 0.01$, in a box of unit size with 128^2 elements. Color denotes $|\mathbf{q}|$, scaled linearly, in arbitrary units (so we simply set $\text{MAX}[|\mathbf{q}|] = 1$). *Left:* Isotropic diffusion. The Gaussian simply expands; all methods accurately capture the exact solution (*bottom*). *Right:* Anisotropic diffusion with $\mathbf{K} = K \hat{\mathbf{B}} \otimes \hat{\mathbf{B}}$ and $\hat{\mathbf{B}} = \hat{x}$. The distribution should expand only in the x -direction. MFM captures this, with minor artifacts from the initial particle arrangement which converge away at higher resolution.

In the isotropic case, this evolves as:

$$q(\mathbf{x}, t) = \frac{q_0 (2\pi)^{-3/2}}{(\epsilon^2 + 2\kappa t)^{3/2}} \exp\left[-\frac{1}{2} \left(\frac{x^2 + y^2 + z^2}{\epsilon^2 + 2\kappa t} \right)\right] \quad (44)$$

where q_0 is an arbitrary normalization, κ is the diffusivity, t the time since the problem was initialized, and ϵ defines the initial width of the distribution ($\epsilon \rightarrow 0$ becomes a δ -function; larger ϵ correspond to starting from an already-evolved solution). We take $\epsilon = 0.05$, comparable to our inter-particle spacing, so that there is a well-defined gradient in our initial condition.

In the anisotropic case, if we assume $\mathbf{K} = K \hat{\mathbf{B}} \otimes \hat{\mathbf{B}}$ with constant (in space and time) $\hat{\mathbf{B}}$, we can always define our axes so that $\hat{\mathbf{B}} = \hat{x}$; then this evolves as:

$$q(\mathbf{x}, t) = \frac{q_0 (2\pi)^{-3/2}}{\epsilon^2 (\epsilon^2 + 2\kappa t)^{1/2}} \exp\left[-\frac{1}{2} \left(\frac{x^2}{\epsilon^2 + 2\kappa t} + \frac{y^2 + z^2}{\epsilon^2} \right)\right]$$

i.e. q diffuses normally along $\hat{\mathbf{B}}$, and not perpendicular.

Fig. 7 shows the results. The isotropic case is easily recovered accurately (even at lower resolution than shown). MFM is also able to recover the anisotropic case; however at low resolution (a 64^3 box) the agreement is not perfect, as some artifacts from the grid structure (here particles were initially laid in a Cartesian grid) are present. These are invisible by-eye if we go to $> 256^3$ resolution, although still measurable in the L1 error norm. In multi-dimensional problems such as this, the perpendicular width of structures must, in general, be a few particles across before complete anisotropy can be fully captured. This is required so that a reliable gradient in the relevant direction can be determined (similar to the requirement in grid-based codes).




Figure 8. The diffusing ring test problem (§ 4.6; style as Fig. 7). A “hot spot” (high- $|\mathbf{q}|$) within a small radial annulus is initialized with pure azimuthal fields; this should diffuse into a ring without additional radial diffusion. We compare MFM at two different resolutions (labeled), and both early (*top*) and late (*bottom*) times. For early times there is an approximate analytic solution, shown; for late times there is no analytic solution, but the high- $|\mathbf{q}|$ material should remain confined to the same radial annulus as at early times, gradually diffusing around the ring until it is isothermal. In MFM/MFV we see good agreement with the expected behaviors at both times, even at low resolution. There is some numerical radial diffusion, but this gradually converges away as we go towards higher resolution.




Figure 9. Convergence in the diffusing ring test (Fig. 8), with our MFM method. We plot the L1 error norm in q averaged over the domain as a function of the number of elements on a side N_{1D} , at a fixed time $t = 0.2$. Convergence is close to the ideal linear scaling (dotted red).

4.6 Diffusing Ring

A more challenging version of this problem is diffusion with azimuthal anisotropy, following Parrish & Stone (2005); Sharma & Hammett (2007, 2011). In a periodic box of unit size centered on the origin and cylindrical (r, ϕ, z) coordinates, $\rho = u = 1$, $\mathbf{v} = 0$, and purely azimuthal magnetic fields $\mathbf{B} = B_0 \hat{\phi}$, we initialize $q(t=0) = q_0 + q_1 (\exp[-(1/2) [(r-r_0)^2 / \delta r_0^2 + \phi^2 / \delta \phi_0^2]])$. Here q_0 and q_1 are an arbitrary background and normalization (we take $q_0 = 10^{-10}$ and $q_1 = 1$), $\delta r_0 = 0.05$ and $r_0 = 0.3$ define a Gaussian ring at radius r_0 of width δr , and $\delta \phi_0 = 0.5$ is an initial Gaussian spread about $\phi = 0$ in the $\hat{\phi}$ direction. We assume $\mathbf{K} = K \hat{\mathbf{B}} \otimes \hat{\mathbf{B}}$.

Because the fields are purely azimuthal, the quantity should diffuse in the purely azimuthal direction, “around the ring,” rather than in the radial direction. At early times, this has an exact solution of the same functional form: $q(t > 0) = q_0 + q_1(t) (\exp[-(1/2) [(r-r_0)^2 / \delta r_0^2 + \phi^2 / \delta \phi^2]])$ where $\delta \phi^2 =$




Figure 10. Diffusing ring test for radiative diffusion (§ 4.6.2; style as Fig. 7), where the diffused quantity $\mathbf{q} = q\mathbb{D}_\nu$ is an anisotropic tensor, with $\mathbb{D}_\nu = \hat{\phi} \otimes \hat{\phi}$ (i.e. a purely azimuthal Eddington tensor). We initialize a small ($r = 0.1$), thin ($\delta r = 0.01$), azimuthally symmetric ring (in a box of unit length); the ring should expand with speed $\approx K/r$ and remain thin. The resolution here is 64^2 , and we show the result at early ($t = 0.1$; *top*) and late ($t = 4$; *bottom*) times. MFM captures the physically correct behavior.

$\delta\phi^2(r, t) = \delta\phi_0^2 + 2Kr^{-2}t$ (with normalization $q_1(r, t) = q_1(t = 0)(\delta\phi_0/\delta\phi)$). This assumes we can neglect the periodic boundary conditions around the ring – i.e. is valid for $\delta\phi \ll \pi$ (hence early times). At late times, the diffusion from both directions self-intersects on the opposite side of the ring, and there is no simple exact solution. Eventually, though, as $t \rightarrow \infty$, the system becomes isothermal within each azimuthal annulus. This is a challenging problem even in high-order grid codes (see Parrish & Stone 2005).

Fig. 8 compares the results at early and late times, at two different resolution levels. As before, MFM is able to capture the azimuthal anisotropy. Even at low resolution (64^2), there are only weak grid artifacts, but these and the amount of perpendicular diffusion improve at higher resolution. At late times, for comparison, in the fixed-grid code ATHENA on a Cartesian mesh (where the preferred direction of the grid is not the azimuthal direction), it requires going to $\sim 256^2$ resolution before the diffusion properly “wraps” into a ring at all (see e.g. Sharma & Hammett 2011); our 128^2 case resembles a $\sim 512^2$ case with ATHENA. And we note that we have not aligned the particles with the anisotropy (the particles are in a regular triangular lattice).

Note that we have tested both the 2D ring version of this problem and the 3D version, where the ring becomes a cylinder elongated in the \hat{z} direction and the box is periodic in that direction. The results are very similar in both cases.

4.6.1 Convergence

Note that our diffusing ring setup is slightly different from that in Sharma & Hammett (2007, 2011) and Kannan et al. (2015), who used this problem to measure the convergence properties of their method. We have therefore also compared initial conditions which match their choice: we initialize a step-function discontinuity with $q = q_0$ and $q = q_1$ inside or outside (respectively) of an annulus $0.25 < r < 0.35$ and $-\pi/6 < \phi < \pi/6$. The qualitative results with all methods are identical to those shown in Fig. 8.

Fig. 9 quantifies the convergence (in 2D tests using this initial condition) by measuring the L1 norm of q at time $t = 0.2$ relative to a high-resolution solution (2048^2) interpolated to the particle positions. We find a convergence rate $L1 \propto N^{-0.9}$, close to the ideal $\propto N^{-1}$. This is competitive with and in some cases superior to the implementations studied in fixed-grid and moving-mesh codes in Sharma & Hammett (2007, 2011); Kannan et al. (2015). However we caution that we are comparing to our own high-resolution solution, not an exact solution (because none exists); so systematic errors which may converge more slowly do not appear.

© 0000 RAS, MNRAS 000, 000–000

4.6.2 Tensor Diffusion: The Radiative Diffusion Case

Interestingly, if we consider the radiative diffusion version of this problem, the behavior is qualitative different. Take $\mathbb{D}_\nu = \hat{n} \otimes \hat{n}$ and $\hat{n} = \hat{\phi}$ (define $K \equiv \lambda c/(\kappa_\nu \rho)$); because the anisotropy is *inside* the first gradient in the diffusion equation (i.e. we have $\mathbf{F} = K\nabla \cdot [(\hat{\phi} \otimes \hat{\phi})|\mathbf{q}|]$ as opposed to $\mathbf{F} = K(\hat{\phi} \otimes \hat{\phi}) \cdot \nabla|\mathbf{q}|$), the solution is distinct. For an azimuthally symmetric $\mathbf{U} = n_\nu = n_\nu(r)$ and $\hat{n} = \hat{\phi}$, the diffusion equation for scalar q reduces to $\partial q/\partial t = -(K/r)\partial q/\partial r$, i.e. the ring expands radially with a speed $= K/r$. Fig. 10 shows the results of this test. In our MFM/MFV methods, the ring expands as expected, and the correct qualitative behavior is captured even at extremely low resolution ($\sim 32^2$).

4.7 Anisotropic Diffusion-Driven Instabilities: MTI & HBI

We now consider two instabilities specific to anisotropically conducting plasmas: the magneto-thermal instability (MTI) and heat-flux driven buoyancy instability (HBI) (Balbus 2000; Quataert 2008). These have been studied in various astrophysical contexts as drivers of turbulence, convection, and mechanisms to enhance or suppress conduction; our specific problem setup is motivated by the studies in Parrish & Stone (2005); Parrish et al. (2008); Parrish & Quataert (2008); McCourt et al. (2011); Kannan et al. (2015).

Here, we are not interested in the physics of the instabilities themselves, but they are useful numerical tests for several reasons. (1) They require accurate coupling of the anisotropic conduction to the magneto-hydrodynamics of the flow (not guaranteed in operator-split methods). (2) They are specific to anisotropic conduction and are suppressed with isotropic conduction, so directly test whether isotropic numerical diffusion can overwhelm physical diffusion. (3) They test the ability of the anisotropic conduction operator to recover small-amplitude seed perturbations. (4) They lead to non-linear, sub-sonic turbulence, which is particularly challenging for mesh-free methods (historically, SPH) to treat accurately (see e.g. Bauer & Springel 2012; Price 2012a), and require that our operator preserves anisotropy even in a turbulent flow.

4.7.1 High-Resolution Tests

Following Parrish & Stone (2005); Parrish et al. (2008); Parrish & Quataert (2008); McCourt et al. (2011); Kannan et al. (2015), we initialize a 2D box in $x-z$ coordinates with an analytic constant gravitational acceleration $\mathbf{g} = -\hat{z}$, size $L = 1/10$, resolution 256^2 , polytropic $\gamma = 5/3$, and conductivity $K = 0.01$ ($\mathbf{K} = K\mathbf{I}$ for the isotropic case, $\mathbf{K} = K\hat{B} \otimes \hat{B}$ for the anisotropic case). For the MTI, we initialize $u = (3/2)(1-z/H)$ with $H = 3$, and $\mathbf{B} = 10^{-11}\hat{x}$, with a seed velocity perturbation³ $\mathbf{v} = \mathcal{M}_0 c_s \sin(4\pi x/L)\hat{z}$ with $\mathcal{M}_0 = 10^{-2}$. For the HBI, $u = (3/2)(1+z/H)$ with $H = 2$, and $\mathbf{B} = 10^{-11}\hat{z}$ and $\mathbf{v} = \mathcal{M}_0 c_s [\sin(3\pi z/L)\hat{x} + \sin(4\pi x/L)\hat{z}]$. The density $\rho(z)$ is initialized so that the initial pressure gradient balances gravity with $\langle \rho \rangle = 1$ in the box. The domain is divided vertically into three equal sub-domains of length $L/3$; the top and

³ In our high-resolution tests, we use a seed velocity perturbation with magnitude $\mathcal{M}_0 = 10^{-2}$ instead of $\mathcal{M}_0 = 10^{-4}$ as in Parrish & Stone (2005); Parrish et al. (2008); Parrish & Quataert (2008). This was chosen for computational convenience. Because we use an explicit integration method, the timestep for high resolution in a small spatial domain becomes very short ($\sim 10^{-5}$). Using the smaller seed velocity requires approximately an order-of-magnitude longer runtimes to develop non-linear behavior, so we adopt the larger seed for high-resolution tests. However we explicitly show in § 4.7.2 below that our MFM method can accurately capture very small seed velocities even at much lower resolution $\sim 32^2$.




Figure 11. Magneto-thermal instability (MTI) and heat-flux driven buoyancy instability (HBI) tests (§ 4.7), with our MFM method. *Top:* MTI: The initial condition (*left*) is a 2D box of length $L = 0.1$, with 256^2 resolution elements, with a vertically stratified atmosphere with $dT/dz < 0$ (temperature decreases upwards), in equilibrium with a constant vertical gravitational field $\mathbf{g} = -\hat{z}$ and conductivity $\kappa = 0.01$. Magnetic field lines are shown at different times t ; these are initialized with trace values aligned along \hat{x} . Color denotes $|\mathbf{B}|$, increasing linearly from the minimum to maximum values at each time (cyan to purple to magenta, respectively; see Fig. 13 for quantitative values). With no diffusion, or isotropic diffusion ($\mathbf{K} = \kappa \mathbf{I}$; *center*), the system is stable and the field configuration should be preserved. With anisotropic diffusion ($\mathbf{K} = \kappa \hat{B} \otimes \hat{B}$; *right*), the system is unstable and develops convection. This re-orient the field to be initially near vertical (at e.g. the time shown), then breaks up into turbulence and the field becomes isotropic. The characteristic timescale is the buoyancy time ~ 1.7 in these units. *Bottom:* HBI. The resolution and conductivity are the same but the initial atmosphere now has $dT/dz > 0$, and initial $\hat{B} = \hat{z}$. Again with no diffusion or isotropic diffusion, the system is stable. With anisotropic diffusion, the HBI amplifies transverse motions that re-orient the field to be perpendicular ($\hat{B} \rightarrow \hat{x}$). MFM recovers all of the expected behaviors for both instabilities with anisotropic diffusion, and correctly suppresses them with isotropic diffusion.

bottom layers have isotropic conduction (i.e. are buoyantly neutral) while the central layer has anisotropic conduction (this follows the previous studies and reduces the sensitivity to boundary conditions). The boundaries are periodic in \hat{x} ; and constant-temperature reflecting boundaries in \hat{z} . Note that a sharp, reflecting and conducting “wall” is particularly challenging to implement in mesh-free methods. We treat the reflecting boundaries as follows: a layer (3 particles deep) of boundary particles with fixed positions and temperatures is placed at $z < 0$ and $z > L$. For every interaction between a boundary particles and normal particle (gradient calculation, the hydrodynamic operations, etc), the boundary particle is assumed to have all specific properties matched to the normal particle, except the temperature (fixed to the IC value), velocity and magnetic field (which follow the usual reflection rules) and density (adapted to give equal pressure, given the different temperatures).

Both of these ICs represent a stably stratified atmosphere. With no conduction, or isotropic conduction, the system should remain in equilibrium and the seed perturbations should damp. With anisotropic conduction, provided a large enough diffusivity (as chosen here) such that the system is approximately isother-

mal along field lines, the instabilities should grow and eventually re-orient the field lines. In the MTI (vertical temperature profile $dT/dz < 0$), the small vertical velocity grows into large, non-linear convective cells which carry the field lines and re-orient the field in the vertical direction, until the cells cross the domain and the fixed-temperature boundary conditions produce sustained turbulence. In the HBI ($dT/dz > 0$), the initial mixed perturbation generates growing separation/compression of field lines, which leads to horizontal motions that try to stretch the field lines horizontally, until the instability saturates when the field lines are horizontal and suppress further convection. In both the HBI and MTI, the characteristic timescale is the buoyancy time $|g \partial \ln T / \partial z|^{-1/2}$ ($\sim (1.7, 1.4)$ in code units for our MTI and HBI setups, respectively).

Fig. 11 shows the results of these tests at a few buoyancy times, using our MFM method. Recall the relatively large initial perturbation here means non-linear behavior can develop in just a couple buoyancy times, and we see that occur in the anisotropic case. The qualitative behavior of the non-linear HBI and MTI compares well to that seen in Parrish & Stone (2005); Parrish &

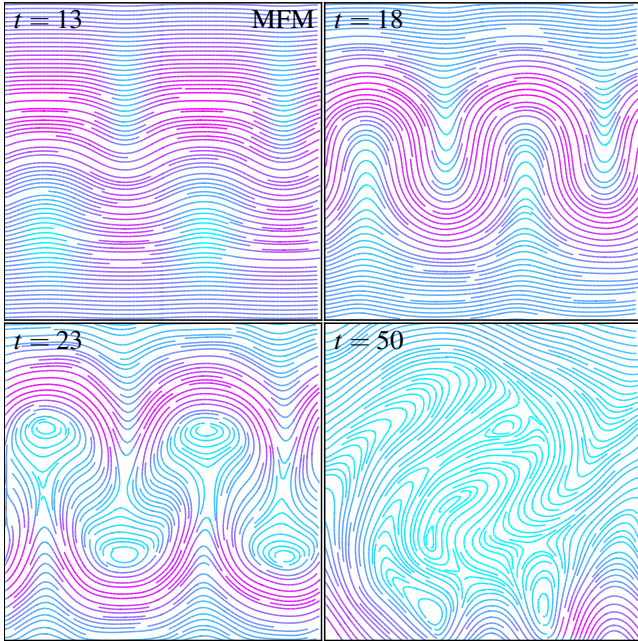


Figure 12. Development of the MTI in an MFM simulation at low (32^2) resolution, now with a box of side-length $L = 1$. Field lines are shown as Fig. 11. Initially small seed velocity perturbations grow in the vertical direction and become the convective cells. Once the cells reach the box boundaries ($t \gtrsim 25$, the convection is sustained by the constant-temperature, reflecting boundary conditions and produces sustained turbulence that isotropizes the magnetic field. Again, MFM can capture all the important behaviors even at very low resolution, especially for the larger box studied here which produces larger Mach numbers (here ~ 0.03) compared to Fig. 11. The buoyancy time for this setup is ~ 1.7 .

Quataert (2008); McCourt et al. (2011). In both MTI and HBI, the cases with isotropic diffusion show (correctly) no evidence of instability, and the initial perturbations are eventually fully damped.

4.7.2 Low-Resolution Tests

To explore late-time evolution and resolution effects, we now consider an initial condition which is identical to our high-resolution tests, but with lower resolution 32^2 and larger box size $L = 1$ (these both allow larger timesteps), smaller seed velocity $\mathcal{M}_0 = 10^{-4}$, and larger conductivity $K = 0.1$ (to preserve the desired limit where the conduction is sufficiently fast).

Fig. 12 shows the time evolution of the MTI in one of these runs; we obtain similar (good) behavior in our low-resolution MTI runs. In both cases, even at 32^2 resolution, the instabilities are (correctly) completely suppressed with isotropic diffusion. Quantitatively, the results from these runs are shown in Fig. 13.

In MFM, we see the MTI grow (despite the very low resolution), going non-linear after $\sim 10 - 20$ buoyancy times (similar to high-resolution cases with the same seed perturbation amplitude; Parrish & Stone 2005; McCourt et al. 2011; Kannan et al. 2015), at which point the convective plumes re-order the magnetic field from horizontal to vertical. At late times, the plumes break up into sustained, non-linear convection, which isotropizes the field (although this is known to be sensitive to the boundary conditions). The fluctuations around isotropy owe to the ongoing turbulence (and are smaller at higher resolution). The saturated Mach numbers $\sim 0.01 - 1$ are larger than those in the smaller $L = 0.1$ box, as expected based on the scaling seen in McCourt et al. 2011. Both qualitative and quantitative evolution compare favorably to




Figure 13. Time evolution of our low-resolution MTI (left) and HBI (right) simulations from Fig. 12 (§ 4.7.2). *Top:* rms Mach number (volume-averaged) in the center of the computational domain, as a function of time. *Bottom:* Mean squared z -component of the magnetic field orientation, $\hat{B}_z^2 = (\hat{B} \cdot \hat{z})^2$. In MFM, the instabilities develops and velocities increase from their seed values rapidly around $t \sim 10 - 20 t_{\text{buoyancy}}$. The magnetic fields are rapidly re-oriented at the same time. In the MTI, fields go from horizontal to vertical, then the sustained convection with steady-state Mach numbers $\sim 0.03 - 0.1$ (given the large box-size $L = 1$) isotropizes the fields (red dotted line represents isotropic fields). The fluctuations about isotropy owe to the low resolution. In the HBI, fields go from vertical to isotropic, but the large Mach numbers (~ 0.02), low resolution, and boundary conditions here cause an overshoot that sustains fluctuations around isotropy until the velocities decay to $\lesssim 0.01 c_s$, at which point the instability rapidly completes the horizontal re-orientation of the field.

higher-resolution studies in fixed-grid and moving-mesh schemes (Mccourt et al. 2011; Kannan et al. 2015).

For the HBI, the instability goes non-linear on a similar timescale to the MTI, but the Mach numbers slowly decay as the field re-aligns. Interestingly, for the $L = 1$ box, there is an intermediate period of isotropic fields, before the instability completes their horizontal re-alignment; this does not appear in the $L = 0.1$ box, so we suspect it owes to the larger turbulent “overshoot” through the unstable zone induces by the first phase of the instability, which must be damped before the re-alignment can complete.

4.8 Hall MRI

As discussed in § 4.1.4, the Hall effect presents unique numerical challenges; therefore, we now consider the magneto-rotational instability (MRI) with a Hall term and Ohmic resistivity (“Hall MRI”). This again involves the effects of anisotropic diffusion on plasma instabilities and turbulence. The effects of the Hall term are especially important in the context of proto-stellar and protoplanetary disks (for reviews, see Balbus 2003; Wardle 2007), and have been studied in detail using Eulerian methods (Flock et al. 2011; Simon et al. 2011; Bai 2011; Simon et al. 2015).

The MRI itself is astrophysically interesting in a wide range of contexts involving magnetized disks, and is numerically particularly interesting because it has historically proven challenging for SPH methods to correctly capture its growth (see e.g. Rosswog & Price 2007; Price & Bate 2008; Dolag & Stasyszyn 2009). In Paper II and Hopkins (2016), we consider extensive studies of the

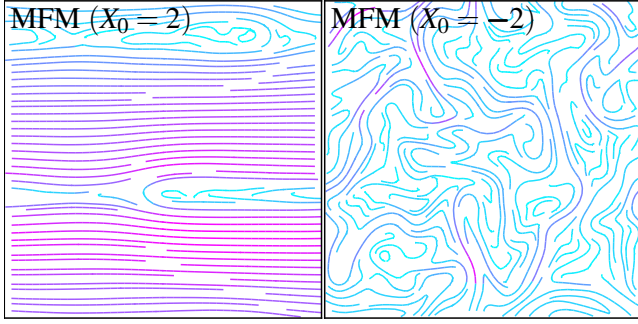


Figure 14. Hall MRI tests (§ 4.8). We compare the field lines (as Fig. 11) well into non-linear evolution. The initial condition is a 2D $R-z$ shearing box of length $L = H$ (the pressure scale length) with 128^2 resolution, a constant field $B_0 \hat{z}$, trace pressure/velocity perturbations, and explicit Ohmic resistivity (with magnetic Reynolds number $\text{Re}_{M,0} = 10 \times 1/\eta_0$) and Hall effects (with Hall parameter $X_0 \propto 1/\eta_H$). Opposite signs of X_0 correspond to the identical setups with opposite signs of \mathbf{B} ; this has no effects unless the Hall term is present. Results are shown at $t \approx 6t_{\text{orbital}}$ ($t_{\text{orbital}} = 2\pi/\Omega$). *Left:* For $X_0 \geq 0$, modes should grow quickly. In the non-linear state, the fastest-growing modes become bigger than the box, leading to the horizontal channel modes seen; these increase $|\mathbf{B}|$ without limit. *Right:* For $-4 \lesssim X_0 \lesssim -2$, growth is slower and smaller modes grow fastest, leading to near-isotropic MRI turbulence in the saturated state as opposed to channel modes.

MRI in ideal MHD ($X_0 = 0$, $\text{Re}_{M,0} \rightarrow \infty$) in GIZMO. We showed that our MFM and MFV methods recover the correct linear growth rates and non-linear properties in good agreement with well-tested higher-order Eulerian codes such as ATHENA. We also showed that at least some SPH schemes were capable of doing the same (albeit with greater noise), if a larger neighbor number is used.

Here, we follow Sano & Stone (2002) and consider a simplified test problem. We adopt a 2D axisymmetric shearing box from Guan & Gammie (2008); this is locally Cartesian with one coordinate (x) representing the radial direction and the other coordinate (z) the vertical direction, representing a small, azimuthally symmetric “ring” about the midplane of a disk in a Keplerian potential. Details of the boundary conditions and gravitational terms, as implemented in GIZMO, are in Paper II (this is the same box used for the MRI simulations therein). The box has initially constant density $\rho_0 = 1$, side-length $L = 1 = H$ (where $H \equiv (2/\gamma)^{1/2} c_{s,0}/\Omega$ is the scale-height), orbital frequency $\Omega = 1$, mean pressure $P_0 = c_{s,0}^2/\gamma = \Omega^2/2$ ($\gamma = 5/3$), spatially uncorrelated random pressure and velocity fluctuations with uniform distribution and $|\delta P|/P_0 = |\delta \mathbf{v}|/c_{s,0} \leq 0.5 \times 10^{-2}$, and uniform vertical field $\mathbf{B} = B_0 \hat{z}$. The MRI is then characterized by three numbers, the plasma beta $\beta_0 \equiv P_0/(v_{A,0}^2 \rho_0) = 3200$ ($v_{A,0} \equiv |B_0|/(4\pi\rho_0)^{1/2}$); the magnetic Reynolds number $\text{Re}_{M,0} \equiv v_{A,0}^2/(\eta_0\Omega) = 10$, which determines the Ohmic resistivity η_0 ; and the Hall parameter $X_0 \equiv cB_0\Omega/(2\pi en_{e,0}v_{A,0}^2)$, which determines $\eta_H = |X_0|v_{A,0}^2(2\Omega)^{-1}(|\mathbf{B}|/|B_0|)(\rho_0/\rho)$ (we assume the free electron fraction is constant).

Figs. 14-15 show the resulting evolution of the magnetic fields, for $X_0 = +2, 0, -2, -5$. For $X_0 = 0$ (no Hall term), we simply have the MRI with explicit resistivity. In 2D shearing boxes, the fastest-growing modes for $\text{Re}_{M,0} = 10 \gg 1$ should have growth rates only slightly smaller than the 0.75Ω expected in the ideal limit; once non-linear, they should form an inverse cascade until horizontal channel modes appear which grow without limit. For $X_0 > 0$, the behavior should be essentially identical, with slightly faster mode growth at higher X_0 . When $X_0 < -4$, the system is fully stable against the MRI, and the initial perturbations should decay.




Figure 15. Growth of the volume-averaged magnetic pressure (relative to the initial thermal pressure P_0) in the mean-field Hall MRI tests. We show results for different initial Hall parameters $X_0 = +2, 0, -2, -5$ (as labeled), at low resolution (32^2 ; thin lines) and intermediate resolution (128^2 ; thick lines). We also show the expected analytic maximum linear growth rate ($|\mathbf{B}|^2 \propto \exp(1.5t\Omega)$; dotted black line), which should be close to the simulated growth rate for $X_0 \geq 0$. The growth for $X_0 \geq 0$ is similar to the ideal MHD case (as expected), with growth rates in good agreement with the analytic expectation for 128^2 resolution, and late-time formation of channel modes (Fig. 14). Growth rates are suppressed at very low resolution, as expected from the MRI studies in Paper II. For $X_0 = -2$, linear growth is suppressed and the magnetic field strength saturates when the system saturates in MHD turbulence (also as expected). For $X_0 = -5$, the system should be stable against MRI growth; we confirm this and that the initial seed noise damps at the expected rate (decay rate $\sim \Omega$). All the results here are in good agreement with well-tested Eulerian codes (see e.g. Sano & Stone 2002, Fig. 5).

At intermediate $-4 < X_0 < 0$, the MRI should grow but with a suppressed maximum growth rate (for finite Re_M). Within this range, as X_0 becomes more negative (larger $|X_0|$), smaller-scale modes grow faster, until for $X_0 \lesssim -2$ there is no critical scale at all; because of this, the system can saturate for many orbital periods with steady-state MRI turbulence (as the growing small-scale, high- k modes prevent the formation of low- k channel modes). We confirm all these behaviors here (compare Fig. 5 in Sano & Stone 2002). As expected and shown in detail in Paper II, the linear growth rates are suppressed at very low resolution ($\sim 32^2$), but rapidly approach the analytic solution at higher resolution. Moreover, we note that we have re-run every problem in Sano & Stone (2002) with our MFM method (varying β_0 , $\text{Re}_{M,0}$, X_0 , and the box size; considering zero net-field cases; and their whistler wave test problem) and confirm identical qualitative behavior.

Note that, for $-4 < X_0 < 0$, given our setup, the instantaneous Hall parameter $X \propto |\mathbf{B}|^{-1}$ and magnetic Reynolds number $\text{Re}_M \propto |\mathbf{B}|^2$, so as modes grow non-linear ($|\mathbf{B}|$ increases), the system moves closer to the ideal MHD limit. Because our problem has finite $\text{Re}_{M,0} = 10$, if we begin with $-2.1 < X_0 < 0$, there is a fastest-growing mode, with relatively large wavelengths ($\lambda_{\text{max}} \approx 0.03H, 0.09H$ for $X_0 = -2, -1$). During the turbulent phase, these modes can increase $|\mathbf{B}|$ non-linearly, which in turn increase the fastest-growing mode growth rate and wavelength, and suppress the growth of the smaller-scale modes. This can eventually trigger a runaway inverse cascade that produces the channel modes seen for $X_0 \geq 0$. We find that, if λ_{max} is well-resolved, this occurs eventually if $X_0 \geq -2$ (for $\text{Re}_{M,0} = 10$), although it in some cases requires ~ 100 orbital times. For $X_0 < -2.1$, we confirm the turbulence remains essentially indefinitely.

4.9 Other Multi-Physics Problems

We have also vetted our algorithms in other non-linear problems where diffusion is not necessarily the primary physics. For example, we have considered a Sedov-Taylor blastwave with thermal conduction, and a shocktube with physical viscosity following (Sijacki & Springel 2006). We find good agreement between our MFM method and ATHENA; small differences are dominated by the known numerical differences in their solution of the normal MHD equations (see Paper I-Paper II), not the sub-dominant diffusion terms. Therefore we do not consider these problems good tests of the diffusion treatment in itself (and do not study them further here). But they do serve as a validation that the diffusion operators here behave properly when coupled to additional dynamics.

A methods paper specifically devoted to the implementation of radiation transport is in preparation, where we compare the flux-limited diffusion, optically-thin variable Eddington tensor (OTVET; Gnedin & Abel 2001), and M1 moment closure approximations and consider several dynamical test problems designed to study the radiation-hydrodynamics of ionizing photons (Khatami et al., in prep). A similar detailed study of the cosmic ray implementation, including cosmic ray pressure effects and the role of cosmic ray streaming (as well as diffusion) is also in preparation (Chan et al., in prep).

As a “stress test” of our implementations, we have also run full cosmological simulations of galaxy formation using the Feedback in Realistic Environments (FIRE) models (Hopkins et al. 2014; Faucher-Giguere et al. 2015; Ma et al. 2015, 2016; Chan et al. 2015; Muratov et al. 2015; Oñorbe et al. 2015), with GIZMO in its MFM mode. These simulations include gas, stars, supermassive black holes, and dark matter, self-gravity, cosmological integrations, cooling physics and gas chemistry, star formation, and feedback from stars in the forms of photo-heating, radiation pressure, stellar winds, and supernovae. We have considered cases with MHD and using the Spitzer-Braginskii coefficients for anisotropic conduction and viscosity, and Smagorinski eddy diffusion for metals. The results are presented in Su et al. (2016). For our purposes here, since there is no known “right” answer for such simulations, we consider these only to be useful tests of numerical stability under extreme conditions. Critically, we see no evidence for numerical instability or unphysical features owing to the addition of anisotropic diffusion operators in these simulations.

5 DISCUSSION

We present numerical discretizations of general anisotropic tensor diffusion operators for Lagrangian hydrodynamics methods, specifically for both recently-developed meshless finite-mass or finite-volume (MFM/MFV) Godunov schemes. We implement these in the multi-method code GIZMO, with the specific implementations relevant for passive scalar diffusion, non-ideal MHD (Ohmic resistivity, the Hall effect, and ambipolar diffusion), sub-grid “turbulent eddy diffusion” models, anisotropic conduction and viscosity (shear/bulk or Braginskii), cosmic ray diffusion and streaming, and anisotropic radiation diffusion (with a variable Eddington tensor).

We consider a variety of test problems. In all cases, our finite-element MFM/MFV schemes can produce accurate solutions, even at low resolution, and are numerically stable. The schemes are also manifestly conservative. They are able to recover correctly the anisotropic cases, up to and including complete suppression with perpendicular magnetic fields. We show this is true regardless of the local particle arrangement/disorder (even in “worst case” sce-

narios where the particle arrangement is totally random), and regardless of the “neighbor number” in the spline. For some cases of great astrophysical interest, e.g. diffusion across a moving contact discontinuity which is not aligned with the grid, these methods (by virtue of being Lagrangian and mesh-free) may exhibit substantially reduced numerical diffusion compared to non-moving, grid-based codes (e.g. AMR methods) at the same resolution. The MFM/MFV methods are able to capture subtle instabilities driven by anisotropic diffusion (e.g. the magneto-thermal and heat-flux driven buoyancy instabilities, and Hall MRI).

The particular form we adopt for the flux-limited Riemann problem is non-trivial, and should be useful for other explicit anisotropic diffusion methods. Stabilizing these methods without introducing excessive numerical diffusion is challenging, especially in irregular/unstructured meshes or mesh-free configurations. The method we propose has the advantage that it trivially generalizes to arbitrarily high-order (and complicated) gradient estimators and slope-limiters, as well as higher-order reconstruction of the gradients at the faces (we simply replace the left and right states in Eqs. 7-11 with their appropriate values); it also admits arbitrarily complex tensors for both the diffusivity and diffused quantities (as opposed to many methods which are specific to scalar diffusion). And the flux computation is pair-wise and negligible in cost compared to the MHD Riemann problem solution. It therefore is of interest not just to mesh-free methods but also moving-mesh and AMR methods.

At least some aspects of our new method here can also be applied to SPH. Because GIZMO is a multi-method code which includes an optional SPH solver, we present an application of this in Appendix A. There we note that some of the low-order errors in other SPH-based anisotropic diffusion formulations can be resolved by application of the higher-order matrix gradient estimators here, but this in turn requires application of a flux-limiter akin to what we use here, to stabilize the method. We show that the resulting “integral-Godunov” SPH behaves well in at least a subset of our tests.

The major dis-advantage of our method is that it does not trivially generalize for implicit solvers (see Appendix B). This is a subject that merits investigation in future work, since implicit methods can often provide a large speed boost to certain types of problems. However, the methods here are amenable to significant acceleration via super-timestepping, which can provide a comparable speedup as described in Appendix C.

ACKNOWLEDGMENTS

We thank Eliot Quataert for several helpful suggestions of interesting test problems. Support for PFH was provided by the Gordon and Betty Moore Foundation through Grant #776 to the Caltech Moore Center for Theoretical Cosmology and Physics, an Alfred P. Sloan Research Fellowship, NASA ATP Grant NNX14AH35G, and NSF Collaborative Research Grant #1411920. Numerical calculations were run on the Caltech compute cluster “Zwicky” (NSF MRI award #PHY-0960291) and allocation TG-AST130039 granted by the Extreme Science and Engineering Discovery Environment (XSEDE) supported by the NSF.

REFERENCES

- Agertz, O., et al. 2007, MNRAS, 380, 963
- Alexiades, V., Amiez, G., & Gremaud, P. A. 1996, Communications in Numerical Methods in Engineering, 12, 31

- Arth, A., Dolag, K., Beck, A. M., Petkova, M., & Lesch, H. 2014, MNRAS, submitted, arXiv:1412.6533
- Bai, X.-N. 2011, ApJ, 739, 50
- Balbus, S. A. 2000, ApJ, 534, 420
- . 2003, ARA&A, 41, 555
- Bauer, A., & Springel, V. 2012, MNRAS, 423, 3102
- Brookshaw, L. 1985, Proceedings of the Astronomical Society of Australia, 6, 207
- Chan, T. K., Kereš, D., Oñorbe, J., Hopkins, P. F., Muratov, A. L., Faucher-Giguère, C.-A., & Quataert, E. 2015, MNRAS, 454, 2981
- Choi, E., Kim, J., & Wiita, P. J. 2009, ApJS, 181, 413
- Cullen, L., & Dehnen, W. 2010, MNRAS, 408, 669
- Dilts, G. A. 1999, International Journal for Numerical Methods in Engineering, 44, 1115
- Dolag, K., & Stasyszyn, F. 2009, MNRAS, 398, 1678
- Duffell, P. C., & MacFadyen, A. I. 2011, ApJS, 197, 15
- . 2014, MNRAS, in press, arxiv:1407.7300
- Español, P., & Revenga, M. 2003, Phys. Rev. E, 67, 026705
- Faucher-Giguere, C.-A., Hopkins, P. F., Keres, D., Muratov, A. L., Quataert, E., & Murray, N. 2015, MNRAS, 449, 987
- Flock, M., Dzyurkevich, N., Klahr, H., Turner, N. J., & Henning, T. 2011, ApJ, 735, 122
- Gaburov, E., Johansen, A., & Levin, Y. 2012, ApJ, 758, 103
- Gaburov, E., & Nitadori, K. 2011, MNRAS, 414, 129
- Gao, Z., & Wu, J. 2013, Journal of Computational Physics, 250, 308
- García-Senz, D., Cabezón, R. M., & Escartín, J. A. 2012, A&A, 538, A9
- Gnedin, N. Y., & Abel, T. 2001, New Astronomy, 6, 437
- Guan, X., & Gammie, C. F. 2008, ApJS, 174, 145
- Gurski, K. F., & O'Sullivan, S. 2010, in American Institute of Physics Conference Series, Vol. 1281, American Institute of Physics Conference Series, ed. T. E. Simos, G. Psihoyios, & C. Tsitouras, 761–764
- Hahn, O., Teyssier, R., & Carollo, C. M. 2010, MNRAS, 405, 274
- Harten, A., Lax, P., & van Leer. 1983, SIAM Reviews, 25, 35
- Heitmann, K., et al. 2008, Computational Science and Discovery, 1, 015003
- Hopkins, P. F. 2013, MNRAS, 428, 2840
- . 2015, MNRAS, 450, 53
- . 2016, MNRAS, 462, 576
- Hopkins, P. F., Keres, D., Onorbe, J., Faucher-Giguere, C.-A., Quataert, E., Murray, N., & Bullock, J. S. 2014, MNRAS, 445, 581
- Hopkins, P. F., & Raives, M. J. 2016, MNRAS, 455, 51
- Hu, C.-Y., Naab, T., Walch, S., Moster, B. P., & Oser, L. 2014, MNRAS(arXiv:1402.1788)
- Kannan, R., Springel, V., Pakmor, R., Marinacci, F., & Vogelsberger, M. 2015, MNRAS, submitted, arXiv:1512.03053
- Kuhnert, J. 2003, in Lecture Notes in Computational Science and Engineering, Vol. 26, Meshfree Methods for Partial Differential Equations, ed. M. Griebel & M. Schweitzer (Springer Berlin Heidelberg), 239–249
- Lanson, N., & Vila, J.-P. 2008a, SIAM J. Numer. Anal., 46, 1912
- . 2008b, SIAM J. Numer. Anal., 46, 1935
- LeVeque, R. J. 1998, Journal of Computational Physics, 146, 346
- Luo, H., Baum, J. D., & Löhner, R. 2008, Journal of Computational Physics, 227, 8875
- Ma, X., Hopkins, P. F., Faucher-Giguère, C.-A., Zolman, N., Muratov, A. L., Kereš, D., & Quataert, E. 2016, MNRAS, 456, 2140
- Ma, X., Kasen, D., Hopkins, P. F., Faucher-Giguère, C.-A., Quataert, E., Kereš, D., & Murray, N. 2015, MNRAS, 453, 960
- Maron, J. L., & Howes, G. G. 2003, ApJ, 595, 564
- McCourt, M., Parrish, I. J., Sharma, P., & Quataert, E. 2011, MNRAS, 413, 1295
- Meyer, C. D., Balsara, D. S., & Aslam, T. D. 2012, MNRAS, 422, 2102
- . 2014, Journal of Computational Physics, 257, 594
- Mocz, P., Vogelsberger, M., Pakmor, R., Genel, S., Springel, V., & Hernquist, L. 2015, MNRAS, 452, 3853
- Mocz, P., Vogelsberger, M., Sijacki, D., Pakmor, R., & Hernquist, L. 2014, MNRAS, 437, 397
- Morris, J. P. 1996, Publications of the Astronomical Society of Australia, 13, 97
- Muñoz, D. J., Springel, V., Marcus, R., Vogelsberger, M., & Hernquist, L. 2013, MNRAS, 428, 254
- Müller, E., & Steinmetz, M. 1995, Computer Physics Communications, 89, 45
- Muratov, A. L., Kereš, D., Faucher-Giguère, C.-A., Hopkins, P. F., Quataert, E., & Murray, N. 2015, MNRAS, 454, 2691
- Oñate, E., Idelsohn, S., Zienkiewicz, O. C., & Taylor, R. L. 1996, International Journal for Numerical Methods in Engineering, 39, 3839
- Oñorbe, J., Boylan-Kolchin, M., Bullock, J. S., Hopkins, P. F., Kereš, D., Faucher-Giguère, C.-A., Quataert, E., & Murray, N. 2015, MNRAS, 454, 2092
- O'Shea, B. W., Nagamine, K., Springel, V., Hernquist, L., & Norman, M. L. 2005, ApJS, 160, 1
- Pakmor, R., Springel, V., Bauer, A., Mocz, P., Muñoz, D. J., Ohlmann, S. T., Schaal, K., & Zhu, C. 2016, MNRAS, 455, 1134
- Parrish, I. J., & Quataert, E. 2008, ApJL, 677, L9
- Parrish, I. J., & Stone, J. M. 2005, ApJ, 633, 334
- Parrish, I. J., Stone, J. M., & Lemaster, N. 2008, ApJ, 688, 905
- Peery, K., & Imlay, S. 1988, Blunt-body flow simulations (American Institute of Aeronautics and Astronautics), 2904+
- Petkova, M., & Springel, V. 2009, MNRAS, 396, 1383
- Price, D. J. 2012a, MNRAS, 420, L33
- . 2012b, Journal of Computational Physics, 231, 759
- Price, D. J., & Bate, M. R. 2008, MNRAS, 385, 1820
- Quataert, E. 2008, ApJ, 673, 758
- Read, J. I., Hayfield, T., & Agertz, O. 2010, MNRAS, 405, 1513
- Ritchie, B. W., & Thomas, P. A. 2001, MNRAS, 323, 743
- Rosdahl, J., & Teyssier, R. 2015, MNRAS, 449, 4380
- Rosswog, S. 2014, arXiv e-print: 1405.6034
- Rosswog, S., & Price, D. 2007, MNRAS, 379, 915
- Sano, T., & Stone, J. M. 2002, ApJ, 570, 314
- Sharma, P., Colella, P., & Martin, D. F. 2010, SIAM J. Sci. Comput., 32, 3564
- Sharma, P., & Hammett, G. W. 2007, Journal of Computational Physics, 227, 123
- . 2011, Journal of Computational Physics, 230, 4899
- Sijacki, D., & Springel, V. 2006, MNRAS, 371, 1025
- Simon, J. B., Hawley, J. F., & Beckwith, K. 2011, ApJ, 730, 94
- Simon, J. B., Lesur, G., Kunz, M. W., & Armitage, P. J. 2015, MNRAS, 454, 1117
- Smagorinsky, J. 1963, Monthly Weather Review, 91, 99
- Springel, V. 2005, MNRAS, 364, 1105
- Springel, V. 2010, MNRAS, 401, 791
- Springel, V., & Hernquist, L. 2002, MNRAS, 333, 649
- Stone, J. M., Gardiner, T. A., Teuben, P., Hawley, J. F., & Simon, J. B. 2008, ApJS, 178, 137
- Su, K.-Y., Hopkins, P. F., Hayward, C. C., Faucher-Giguere, C.-

- A., Keres, D., Ma, X., & Robles, V. H. 2016, MNRAS, in press, arXiv:1607.05274
- Tasker, E. J., & Bryan, G. L. 2008, ApJ, 673, 810
- Tsukamoto, Y., Iwasaki, K., & Inutsuka, S.-i. 2013, MNRAS, 434, 2593
- Tsukamoto, Y., Iwasaki, K., Okuzumi, S., Machida, M. N., & Inutsuka, S. 2015, ApJL, 810, L26
- Uhlig, M., Pfrommer, C., Sharma, M., Nath, B. B., Enßlin, T. A., & Springel, V. 2012, MNRAS, 423, 2374
- Wardle, M. 2007, Astrophysics and Space Science, 311, 35
- Wurster, J., Price, D., & Ayliffe, B. 2014, MNRAS, 444, 1104
- Wurster, J., Price, D. J., & Bate, M. R. 2016, MNRAS, 457, 1037
- Zingale, M., et al. 2002, ApJS, 143, 539

APPENDIX A: APPLICATION TO SPH

As described in Paper I & Paper II, GIZMO is a multi-method code: users can run with our meshless Godunov (MFM or MFV) hydrodynamic methods detailed in the main text, or SPH, if desired.

Isotropic diffusion in SPH is generally well-handled by the hybrid-derivatives formulation of Brookshaw (1985); we confirm this in GIZMO. Anisotropic cases are more difficult. Español & Revenga (2003) derived a second-order accurate generalization of the Brookshaw (1985) method, but Petkova & Springel (2009) and Arth et al. (2014) show that it is (catastrophically) numerically unstable when the degree of anisotropy is large. They both derived the minimum correction/dissipation term necessary to stabilize the equation, giving the hybrid-derivatives SPH equation:

$$\frac{d(V\mathbf{U})_a}{dt} = \frac{m_a}{\rho_a} \sum_b \frac{m_b}{\rho_b} \mathbf{x}_{ab}^T \left[\frac{\mathbf{K}_a + \mathbf{K}_b}{|\mathbf{x}_{ab}^2|} \right] \left[\frac{(\nabla_a W_{ab} + \nabla_b W_{ab})}{2} \otimes (\mathbf{q}_b - \mathbf{q}_a) \right] \quad (\text{A1})$$

where $\mathbf{x}_{ab} = \mathbf{x}_a - \mathbf{x}_b$. Unfortunately, both Petkova & Springel (2009) and Arth et al. (2014) show that this is fundamentally unable to capture significant anisotropy – essentially, Eq. A1 represents a linear combination of ~ 1 part isotropic diffusion for every ~ 5 parts anisotropic diffusion. Thus, the incorrect physical equations are being solved, and the errors (the isotropic part) do not converge with resolution. Even in a simple diffusing slab, this gives systematically incorrect solutions, and effects which depend on anisotropy (the MTI, HBI, diffusing ring) cannot be captured.

A more accurate treatment follows the “gradients of gradients” method as in Sijacki & Springel (2006). Because second-kernel derivatives are unstable and catastrophically noisy, and successive application of the first-order consistent SPH gradient operator yields a non-conservative formulation (Morris 1996; Price 2012b), one applies instead a pair of conjugate first-order and zeroth-order SPH gradient estimators, giving:

$$\frac{d(V\mathbf{U})_a^{\text{SPH}}}{dt} = \sum_b m_a m_b \left(\frac{\mathbf{K}_a \cdot \langle \nabla \otimes \mathbf{q} \rangle_a}{\Omega_a \rho_a^2} \cdot \nabla_a W_{ab} + \frac{\mathbf{K}_b \cdot \langle \nabla \otimes \mathbf{q} \rangle_b}{\Omega_b \rho_b^2} \cdot \nabla_b W_{ab} \right) \quad (\text{A2})$$

where $\Omega_a \sim 1$ is a correction term accounting for variations in the smoothing length h (see Springel & Hernquist 2002), and the SPH first-order consistent derivative operator $\langle \nabla \otimes \mathbf{q} \rangle_a$ is given by

$$\langle \nabla \otimes \mathbf{q} \rangle_{a,\text{SPH}}^{\alpha\beta\dots} = \sum_c \frac{m_c}{\Omega_c \rho_c} (\mathbf{q}_c - \mathbf{q}_a)^{\beta\dots} (\nabla_a W_{ac})^\alpha \quad (\text{A3})$$

Eq. A2 has primarily been used for *isotropic* diffusion (Sijacki

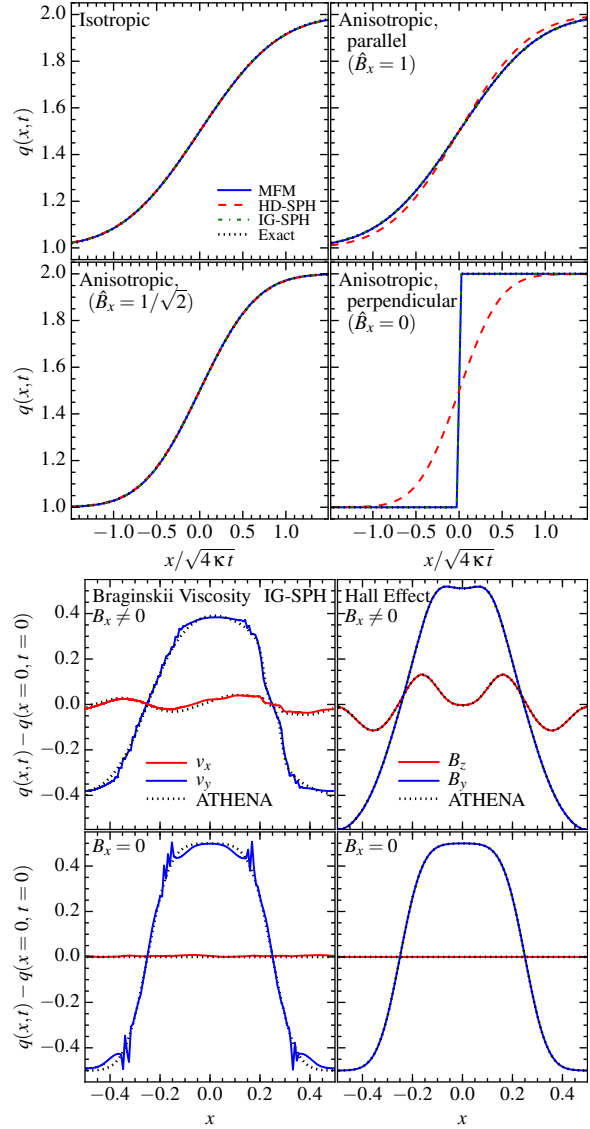


Figure A1. Top: Diffusing sheet test as Fig. 1, but comparing smoothed-particle hydrodynamics (SPH) methods as described in § A. The “hybrid derivatives” HD-SPH method (Eq. A1) is unable to capture proper anisotropy, as is well-known, so we will not consider it further. Our new “integral-Godunov” SPH method, which borrows the higher-order matrix gradient estimators and same flux-limited Riemann solution from the MFM method in the main text, gives very similar results to MFM. Bottom: Braginskii viscosity and Hall effect as Fig. 3, in IG-SPH. The hall effect is captured well; there is some noise in Braginskii viscosity owing to the residual “E0” SPH errors and interaction between the explicit and “artificial” viscosity in SPH.

& Springel 2006; Tsukamoto et al. 2013; Wurster et al. 2014), but recently was applied for anisotropic non-ideal MHD terms in Tsukamoto et al. (2015); Wurster et al. (2016). However, two problems remain. (1) Eq. A3 has a low-order error: for a constant gradient ($q = q_0 + q'x$), the gradient returned by Eq. A3 is $\nabla q = q'(\hat{x} + \epsilon)$ where the error term $\epsilon \sim \mathcal{O}(h^0)$ depends only on the particle arrangement in the kernel ($\epsilon_x = -1 + \sum x_{ca} f_{ca}$, $\epsilon_{y,z} = \sum (y,z)_{ca} f_{ca}$ where $f_{ca} \equiv (x_{ca}/|x_{ca}|)(m_c/\Omega_c \rho_c) \partial W_{ca}/\partial |x_{ca}|$) so does not converge with resolution. (2) Similar, Eq. A2 is zeroth-order, giving error terms $\sim \mathcal{O}(q''h^0) + \mathcal{O}(q'h^{-1})$.

We can potentially improve on this, using our insight from the main text. For (1), replace the standard SPH estimator of $\langle \nabla \otimes \mathbf{q} \rangle_{a,\text{SPH}}$ (Eq. A3) in Eq. A2 with the second-order accu-

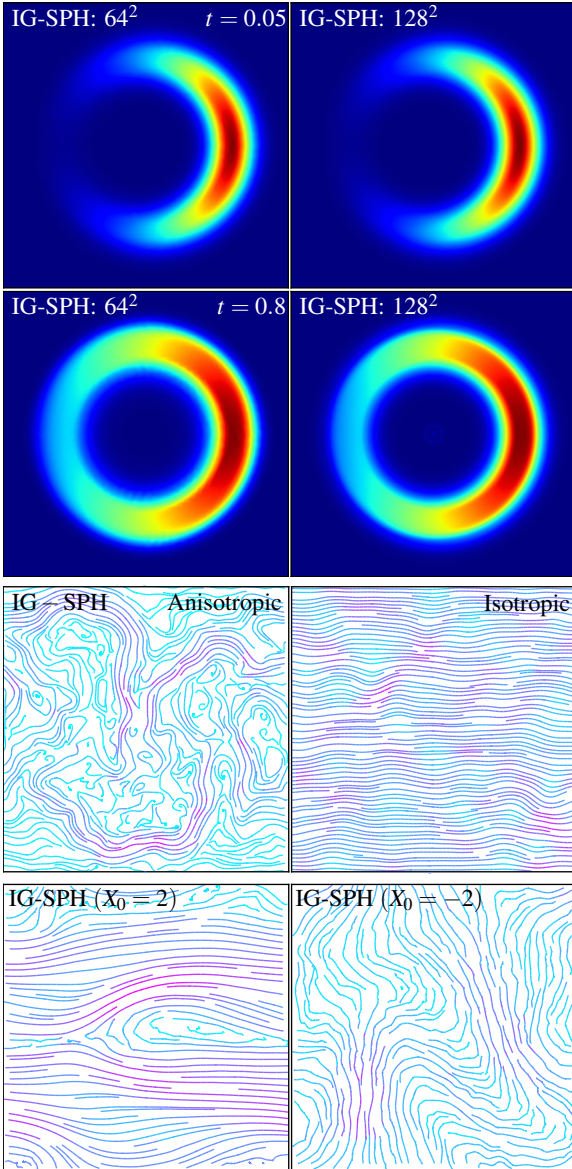


Figure A2. *Top:* Diffusing ring test as Fig. 8. *Middle:* Low-resolution MTI test as Fig. 11-12, for both anisotropic and isotropic diffusion. *Bottom:* Hall MRI test as Fig. 14. IG-SPH performs comparably to our MFM method on the diffusing ring. In the dynamical problems (MTI and Hall MRI), IG-SPH is able to capture some of the qualitative behavior, but known SPH challenges with highly sub-sonic turbulence produce the noise shown (and slower mode growth) and generate some instability even in the isotropic case (which should be stable). These may be fix-able, but are outside the scope of the diffusion operator itself.

rate moving least-squares estimator $(\nabla \otimes \mathbf{q})_a^{\text{matrix}}$ from our MFM method (Eq. 4). This is the same estimator used in so-called “integral” SPH formulations (García-Senz et al. 2012; Rosswog 2014). Unfortunately (as noted by those authors), this alone leads to a numerically unstable method. But we can remedy this by replacing the usual SPH fluxes with our flux-limited Riemann solution. To do this we modify Eq. A2, symmetrizing the density and smooth-

ing length terms as they appear, giving:

$$\frac{d(\mathbf{V}\mathbf{U})_a^{\text{SPH}}}{dt} = - \sum_b \mathbf{A}_{ab}^{\text{SPH}} \cdot \mathbf{F}_{\text{diff}}^{\text{SPH}} \quad (\text{A4})$$

$$\mathbf{A}_{ab}^{\text{SPH}} \equiv \frac{m_a m_b}{\rho_a \rho_b} (\nabla_a W_{ab} + \nabla_b W_{ab}) \quad (\text{A5})$$

$$\mathbf{F}_{\text{diff}}^{\text{SPH}} = -\frac{1}{2} \left(\mathbf{K}_a \cdot (\nabla \otimes \mathbf{q})_a^{\text{matrix}} + \mathbf{K}_b \cdot (\nabla \otimes \mathbf{q})_b^{\text{matrix}} \right) \quad (\text{A6})$$

and then replace $\mathbf{F}_{\text{diff}}^{\text{SPH}} \rightarrow \mathbf{F}_{\text{diff}}^*$ solved exactly as in § 2.2, identifying $\mathbf{F}_{\text{diff},R} = \mathbf{K}_b \cdot (\nabla \otimes \mathbf{q})_b^{\text{matrix}}$ (and likewise L and a). Note that this is now effectively a Godunov-type method, with the “SPH effective face area” $\mathbf{A}_{ab}^{\text{SPH}}$. We therefore refer to this method as “integral-Godunov” (IG) SPH. Finally, note that this does not fully resolve issue (2) above; doing so requires replacing $\mathbf{A}_{ab}^{\text{SPH}}$ with a more consistent face estimator; of course, replacing this with our MFM faces gives *exactly* our MFM method in the main text (the only difference then would be how the other hydro equations are solved).

Figs. A1-A2 show a series of tests of this modified IG-SPH method. First we simply verify the well-known conclusion that the HD-SPH form (Eq. A1) is unable to handle anisotropic diffusion. In contrast our IG-SPH method performs comparably to our MFM method (not surprising, given how similar they are) on simple diffusion sheet problems and e.g. the diffusing pulse and ring tests. We have also compared the “gradients of gradients” SPH method (Eq. A2); if the particle order is ideal this is comparable to our IG-SPH method, but if we seed initial noise in the field \mathbf{q} or particle distribution, then it is significantly more noisy, and it fails to converge with resolution alone (keeping neighbor number fixed) on the diffusing ring test – this is of course expected from our caveats (1)-(2) above. Where IG-SPH shows some issues in Figs. A1-A2: the Braginskii viscosity test, MTI, HBI, and Hall MRI tests, the problems appear to primarily owe to known difficulties with the SPH *hydrodynamics* operators (e.g. artificial viscosity triggering imperfectly, and low-order hydrodynamic errors generating noise in sub-sonic turbulence) – there are various ways to improve these errors but since they are not in the diffusion operators themselves they are outside the scope of this paper.

APPENDIX B: IMPLICIT METHODS

In the main text we solve the diffusion equations explicitly. Qualitatively speaking, we update \mathbf{U} according to

$$(\mathbf{V}\mathbf{U})^{n+1} = (\mathbf{V}\mathbf{U})^n + \Delta t \frac{d(\mathbf{V}\mathbf{U})^{n+1/2}}{dt} \quad (\text{B1})$$

where n refers to the timestep and $n + 1/2$ refers to drifting all quantities to a half-timestep before calculating $d(\mathbf{V}\mathbf{U})/dt$ (see Paper I for details).

The explicit solver has the advantages of computational simplicity, extension to high-order gradient operators/reconstruction methods, allowing non-linear flux limiters, and trivially generalizing to hierarchical adaptive timesteps as we use in our code (akin to “sub-cycling”). However, the disadvantage is that it requires a timestep limit of the form in Eq. 15 (quadratic in the resolution). In some situations this can become extremely small; in these cases, implicit methods are popular. In the implicit case, we take

$$(\mathbf{V}\mathbf{U})^{n+1} = (\mathbf{V}\mathbf{U})^n + \Delta t \frac{d(\mathbf{V}\mathbf{U})^{n+1}}{dt} \quad (\text{B2})$$

Petkova & Springel (2009) show how to cast the isotropic SPH hybrid-derivatives diffusion equation in this form, and then re-cast it as a linear equation of the form $\mathbf{M} \cdot (\mathbf{V}\mathbf{U})^{n+1} = (\mathbf{V}\mathbf{U})^n$ for a

vector (\mathbf{VU}) of elements $(VU)_a$ (representing each resolution element a), and matrix \mathbf{M} . This can then be solved over all particles in a global timestep using a conjugate-gradient inversion. Unfortunately, for our MFM/MFV methods, deriving an implicit method is more challenging. If we simplify by using a first-order reconstruction of $\nabla \otimes \mathbf{q}$ (with corresponding limiters), and (for now) neglect the extra diffusion terms from the Riemann problem, and assume a scalar \mathbf{q} , then Eq. 3 for $d(VU)/dt$ gives us

$$(VU)_a^{n+1} = (VU)_a^n + \frac{\Delta t}{2} \sum_b \left(\mathbf{K}_a \langle \nabla q \rangle_a^{n+1} + \mathbf{K}_b \langle \nabla q \rangle_b^{n+1} \right) \cdot \theta_{ab} \mathbf{A}_{ab} \quad (\text{B3})$$

where θ_{ab} is the limiter function based on comparing the implied flux here to the “direct flux” per Eq. 14. Combining this with Eq. 4 for the gradient estimators, re-arranging and simplifying, we obtain

$$\mathbf{M} \cdot (\mathbf{VU})^{n+1} = (\mathbf{VU})^n \quad (\text{B4})$$

$$M_{ab} \equiv \delta_{ab} + \frac{\Delta t}{2} \frac{\zeta_b}{V_b} \left[\mathbf{K}_b \boldsymbol{\mu}_b^S \cdot \left(\tilde{\mathbf{A}}_{ab} + \tilde{\mathbf{A}}_b^S \delta_{ab} \right) + \sum_c \mathbf{K}_c \boldsymbol{\mu}_{cb} \cdot \left(\tilde{\mathbf{A}}_{ac} - \tilde{\mathbf{A}}_c^S \delta_{ac} \right) \right] \quad (\text{B5})$$

$$\tilde{\mathbf{A}}_{ab} \equiv \theta_{ab} \mathbf{A}_{ab} \quad , \quad \boldsymbol{\mu}_{ab} \equiv \beta_a \mathbf{W}_a^{-1} \mathbf{x}_{ba} \omega_b(\mathbf{x}_a) \quad (\text{B6})$$

$$\tilde{\mathbf{A}}_a^S \equiv \sum_c \tilde{\mathbf{A}}_{ab} \quad , \quad \boldsymbol{\mu}_a^S \equiv \sum_c \boldsymbol{\mu}_{ac} \quad (\text{B7})$$

where δ_{ab} is the Kronecker delta, ω and \mathbf{W}^{-1} are defined in Eq. 4, and β is the slope-limiter for $\langle \nabla \otimes \mathbf{q} \rangle$ (see Paper I). Now, add the diffusive terms from the Riemann problem, $\mathbf{F} \rightarrow \mathbf{F} + \mathbf{F}_{\text{diss}}$ where $\mathbf{F}_{\text{diss}} = \alpha \lambda (\mathbf{U}_R - \mathbf{U}_L) = \alpha_{ab} \lambda_{ab} [u_a - u_b + (\eta_{ab}/2) (\nabla u_a + \nabla u_b) \cdot \mathbf{x}_{ba}]$ if we use a second-order reconstruction of \mathbf{U}_R and \mathbf{U}_L . Here $\lambda = |v_R - v_L|/2 + c_{\text{fast}}$ and η_{ab} is the appropriate limiter for the U_R, U_L reconstruction, both as defined in the text. This gives us

$$\mathbf{M} \rightarrow \mathbf{M} + \mathbf{M}^{\text{diss}} \quad (\text{B8})$$

$$M_{ab}^{\text{diss}} \equiv - \frac{\Delta t}{2V_b} \left[2 \left(\varphi_{ab} - \varphi_b^S \delta_{ab} \right) + \boldsymbol{\mu}_b^S \cdot \left(\mathbf{A}_{ab} + \mathbf{A}_b^S \delta_{ab} \right) - \sum_c \boldsymbol{\mu}_{cb} \cdot \left(\mathbf{A}_{ac} + \mathbf{A}_c^S \delta_{ac} \right) \right] \quad (\text{B9})$$

$$\varphi_{ab} \equiv \alpha_{ab} \lambda_{ab} |\tilde{\mathbf{A}}|_{ab} \quad , \quad \mathbf{A}_{ab} \equiv \eta_{ab} \varphi_{ab} \mathbf{x}_{ba} \quad (\text{B10})$$

$$\varphi_a^S \equiv \sum_c \varphi_{ac} \quad , \quad \mathbf{A}_a^S \equiv \sum_c \mathbf{A}_{ac} \quad (\text{B11})$$

The difference between MFM and IG-SPH simply amounts to the appropriate value for \mathbf{A}_{ab} .

If the limiter functions θ_{ab} , η_{ab} , α_{ab} , and β_a were independent of (\mathbf{VU}) , then we could again perform a simple sparse-matrix inversion to update (\mathbf{VU}) . Obviously even in this limit, the matrix terms are quite complicated: note the extra summations over the index “ c ” that appear, which owe to the higher-order gradient estimator we adopt and cannot be eliminated without lowering the order of the gradient approximation. But the real problem is that the higher-order accuracy and stability of our MFM/MFV methods depends on the limiters being *non-linear* functions of *both* $\mathbf{U}_{a,b}$ and $\nabla \otimes \mathbf{q}_{a,b}$ (itself a sum over $\mathbf{U}_{a,b}$). Of course, non-linear global elliptical equations can be solved, but using an iterative root-finding method to solve for (\mathbf{VU}) , we would have to repeat the summation over \mathbf{U} to determine $\nabla \otimes \mathbf{q}_{a,b}$ and re-compute the limiter functions between each iteration. Combined with a global timestep, this would make the implicit solver much more expensive than sub-cycling our explicit solver (defeating the purpose).

In the limit where the problem is sufficiently smooth, well-resolved, and there is good particle order, $\theta, \eta, \beta \rightarrow 1$ and $\alpha \rightarrow 0$.

Adopting these values and implementing the sparse matrix inversion, we are able to confirm our conclusions for the few test problems that satisfy these criteria. However, with these values fixed, the method is numerically unstable. We can stabilize the method for problems with unresolved gradients (still assuming good particle order) by taking $\theta, \beta, \alpha \rightarrow 1, \eta \rightarrow 0$, but this produces excessive numerical diffusion, especially at low resolution. Unfortunately, we see no obvious way to achieve the combination of accuracy and stability in the text (given the method studied here) without involving non-linear terms in $\nabla \otimes \mathbf{q}$, which are prohibitive for most implicit methods. However, it is possible that a semi-implicit method similar to the one in Sharma & Hammett (2011) could be implemented, where a subset of the flux components which do not require strong limiters are solved implicitly while the others (where the limiters apply) are updated explicitly.

APPENDIX C: SUPER-TIMESTEPPING

Another method to allow larger timesteps in explicit methods for elliptic/parabolic equations is so-called “super-timestepping” (see Alexiades et al. 1996; Gurski & O’Sullivan 2010; Meyer et al. 2012, 2014, and references therein).

A super-timestep Δt_s is composed of N sub-steps δt_j

$$\Delta t_s = \sum_{j=1}^N \delta t_j \quad (\text{C1})$$

The diffusion equations are calculated and updated as usual on each sub-step δt_j , however, these are essentially Runge-Kutta sub-steps in the sense that stability is not guaranteed at after any individual sub-step δt_j , but *only* on the super-step Δt_s .

Alexiades et al. (1996) show that the optimal timesteps which satisfy the necessary stability conditions on Δt_s and simultaneously maximize Δt_s for a given N are given by

$$\delta t_j = \Delta t_{\text{expl}} \left[(1 + \nu) - (1 - \nu) \cos \left(\frac{\pi(2j-1)}{2N} \right) \right]^{-1} \quad (\text{C2})$$

where Δt_{expl} is our usual explicit timestep given by Eq. 15 and $0 < \nu < \lambda_{\text{min}}/\lambda_{\text{max}}$ (where λ_{min} and λ_{max} are the smallest and largest eigenvalues of the diffusion equation). As $\nu \rightarrow 0$, the total time Δt_s covered by the same number N of substeps (and computations) increases, with the limit $\Delta t_s \rightarrow N^2 \Delta t_{\text{expl}}$ as $\nu \rightarrow 0$. Thus this method can provide a speedup of up to a factor N .

This is trivial to implement in explicit methods with a global timestep, and we have done so here. Although all problems shown in the text use the standard, non-optimized explicit timestep Δt_{expl} , we have re-run all of them enforcing a global timestep and using the super-timestepping scheme with varying N and ν . Care is needed, since overly-aggressive timestepping can lead to a loss of accuracy even while maintaining stability (see discussion in § D). In practice, we find $N \sim 5 - 10$ and $\nu \sim 0.04$ give substantial speed-ups (factor $\sim 2.5 - 3$ on the problems here) without any measurable loss of accuracy (consistent with the studies in Alexiades et al. 1996; Choi et al. 2009 and Tsukamoto et al. 2013, who show good accuracy is maintained for $\nu \gtrsim 1/N^{1/2}$).

Moreover, it is straightforward to generalize this to our individual timestepping scheme. In this scheme, particles have independent timesteps but are discretized into hierarchical powers-of-two timebins (see Paper I and Springel 2005); this ensures they remain synchronized. If we take each δt_j and “round down” the nearest (smaller) power-of-two bin, it is straightforward to show that the stability condition in Alexiades et al. (1996) is still satisfied (more easily, in fact). The tradeoff is that the timesteps are no longer quite

“optimal” (not as large as possible). However, for careful choices of N and ν (such as the values above), the difference in the super-step Δt_s (the sum of the δt_j) is only $\sim 10\%$ below optimal.

Note that this is a first-order time integration scheme, which makes it particularly easy to implement. Still greater accuracy may be achieved with higher-order schemes following Meyer et al. (2012, 2014), but we have not experimented with these.

APPENDIX D: TIMESTEP LIMITS WITH SUPER-TIMESTEPPING OR IMPLICIT SOLVERS

With either implicit methods or super-timestepping, although formal stability is guaranteed for large time-steps, accuracy is not. For any test problem, we can simply vary the timestep until a desired accuracy is reached. However for more general problems, it is useful to have a more general timestep constraint.

Regardless of the diffusion problem, MHD must obey the usual CFL condition $\Delta t < C_{\text{CFL}} \Delta x / v_{\text{sig}}$ (where v_{sig} is the standard signal velocity accounting for particle relative motions and their fastest wavespeeds). Additional physics (e.g. gravity, radiation, cooling) come with their own timestep constraints. Note that, for hyperbolic equations, we cannot circumvent these conditions using either of the acceleration schemes above; they are always enforced.

There is no rigorous equivalent condition for the diffusion equation with implicit/super-timestepped solvers, but as pointed out by many authors (e.g. Choi et al. 2009; Gurski & O’Sullivan 2010; Tsukamoto et al. 2013), motivated by hyperbolic problems, we can define an effective diffusive signal speed $v_{\text{diff}} \sim \|\mathbf{K}\| \|\nabla \otimes \mathbf{q}\| / \|\mathbf{U}\|$. Using this to define an analogous Courant condition on the super/implicit timestep:

$$\Delta t_{\text{implicit/super}} < C_{\text{CFL}} \Delta x \frac{\|\mathbf{U}\|}{\|\mathbf{K}\| \|\nabla \otimes \mathbf{q}\|} \quad (\text{D1})$$

we find reasonably good accuracy so long as $C_{\text{CFL}} < 0.5$. For more details we refer to Sharma et al. (2010) who rigorously demonstrate that such a criterion is the correct one for regimes where the diffusion equation behaves advectively (e.g. cosmic ray streaming, or our variable-diffusivity test problem).

Note that when there are un-resolved gradients, this reduces (assuming $\mathbf{U} \sim \mathbf{q}$) to the usual explicit timestep-limiter, $\Delta t < \Delta x^2 / \|\mathbf{K}\|$, as we might expect.⁴ When the gradients are well-resolved, however, larger timesteps are allowed: for super-timestepping, this amounts to a speedup by a factor $\sim L_{\text{grad}} / (N \Delta x)$ where N is the number of sub-steps and $L_{\text{grad}} = \|\mathbf{q}\| / \|\nabla \otimes \mathbf{q}\|$ is the gradient scale length. In practice, for most of the test problems in this paper, this translates to a maximum speedup of a factor of ~ 3 in super-timestepping before a noticeable loss of accuracy appears. Given the large overhead of the implicit methods, for the few cases where we can use them (e.g. “standard” SPH), we actually see relatively little speedup (factor < 2). However, if we considered the same problems, in stages with *resolved* gradients, at much higher resolution, the difference between the simple explicit method used in the text and either implicit or super-timestepped methods should grow accordingly.

⁴ On some problems – for example the diffusing sheet with perpendicular fields, implicit methods remain accurate even with much larger timesteps compared to Eq. D1. However, the only cases we find this is true are ones where the solution is steady-state or equilibrium (in which case arbitrarily long timesteps should not, in principle, be problematic). For all realistic problems we consider which feature actual dynamics, a criterion like Eq. D1 is needed to maintain accuracy.