



Local Simulations of MRI turbulence with Meshless Methods

Hongping Deng¹ , Lucio Mayer¹, Henrik Latter², Philip F. Hopkins³ , and Xue-Ning Bai⁴

¹ Center for Theoretical Astrophysics and Cosmology, Institute for Computational Science, University of Zurich, Winterthurerstrasse 190, 8057 Zurich, Switzerland
hpdeng@physik.uzh.ch

² Department of Applied Mathematics and Theoretical Physics, University of Cambridge, Centre for Mathematical Sciences, Wilberforce Road, Cambridge CB3 0WA, UK

³ TAPIR, California Institute of Technology, Pasadena, CA 91125, USA

⁴ Institute for Advanced Study and Tsinghua Center for Astrophysics, Tsinghua University, Beijing 100084, People's Republic of China
 Received 2018 October 31; revised 2019 February 13; accepted 2019 February 20; published 2019 April 1

Abstract

The magneto-rotational instability (MRI) is one of the most important processes in sufficiently ionized astrophysical disks. Grid-based simulations, especially those using the local shearing box approximation, provide a powerful tool to study the nonlinear turbulence the MRI produces. On the other hand, meshless methods have been widely used in cosmology, galactic dynamics, and planet formation, but have not been fully deployed on the MRI problem. We present local unstratified and vertically stratified MRI simulations with two meshless MHD schemes: a recent implementation of smoothed-particle magnetohydrodynamics (SPH MHD), and a meshless finite-mass (MFM) MHD scheme with constrained gradient divergence cleaning, as implemented in the GIZMO code. Concerning variants of the SPH hydro force formulation, we consider both the “vanilla” SPH and the PSPH variant included in GIZMO. We find, as expected, that the numerical noise inherent in these schemes significantly affects turbulence. Furthermore, a high-order kernel, free of the pairing instability, is necessary. Both schemes adequately simulate MRI turbulence in unstratified shearing boxes with net vertical flux. The turbulence, however, dies out in zero-net-flux unstratified boxes, probably due to excessive numerical dissipation. In zero-net-flux vertically stratified simulations, MFM can reproduce the MRI dynamo and its characteristic butterfly diagram for several tens of orbits before ultimately decaying. In contrast, extremely strong toroidal fields, as opposed to sustained turbulence, develop in equivalent simulations using SPH MHD. The latter unphysical state is likely caused by a combination of excessive artificial viscosity, numerical resistivity, and the relatively large residual errors in the divergence of the magnetic field.

Key words: accretion, accretion disks – magnetohydrodynamics (MHD) – methods: numerical – turbulence

1. Introduction

The turbulence instigated by the magneto-rotational instability (MRI) can transport angular momentum outward, thus enabling accretion in several sources—such as dwarf novae, low-mass X-ray binaries, and Active Galactic Nuclei (AGNs). Numerical MHD simulations are necessary to study this highly nonlinear problem. Simulations of MRI range from local shearing box simulations, both unstratified (e.g., Hawley et al. 1995, 1996; Sano et al. 2004; Simon & Hawley 2009) and stratified (e.g., Brandenburg et al. 1995; Stone et al. 1996; Miller & Stone 2000; Davis et al. 2010; Simon et al. 2011), to global simulations (e.g., Armitage 1998; Hawley 2000; Steinacker & Papaloizou 2002; Fromang & Nelson 2006; Parkin & Bicknell 2013; Zhu & Stone 2018). Three-dimensional simulations carried out with different grid-based codes, such as ZEUS (Hawley et al. 1995), Pencil (Brandenburg & Dobler 2002), RAMSES (Teyssier 2002; Fromang & Nelson 2006), ATHENA (Stone et al. 2008), and the spectral code Snoopy (Lesur & Longaretti 2007) report similar statistics for the turbulence.

Local MRI simulations are especially challenging because the saturated state appears to depend on the small-scale diffusion, be it physical or numerical. For instance, zero-net-flux simulations in unstratified boxes do not converge with increasing resolution, as the turbulent motions reach their peak amplitude near the smallest resolvable scales. However, if physical sources of diffusivity are incorporated and resolved, turbulence can die out when the magnetic Prandtl number is too small (Fromang & Papaloizou 2007; Fromang et al. 2007). The

latter dissipation is also sensitive to the vertical aspect ratio of the computational domain (L_z/L_x) (Shi et al. 2015). On the other hand, in net vertical flux simulations, angular momentum transport depends on the magnetic Prandtl number yet again—at least, when the latter takes values of order unity (Meheut et al. 2015). Vertically stratified shearing box simulations without a net flux also suffer convergence problems (Bodo et al. 2014; Ryan et al. 2017). Adding a net vertical flux, however, can radically change the character of MRI turbulence; for example, magnetic winds may be launched. Some of the properties of these winds also suffer from non-convergence (Bai & Stone 2013; Fromang et al. 2013; Lesur et al. 2013). Finally, non-ideal MHD effects can suppress or radically alter the nature and strength of turbulence (e.g., Fleming et al. 2000; Turner et al. 2007; Bai & Stone 2011; Bai 2014; Lesur et al. 2014; Simon et al. 2015). Such effects are still under investigation.

Currently there are very few published studies of the MRI undertaken with mesh-free methods (see, e.g., Gaburov & Nitadori 2011; Pakmor & Springel 2013; Hopkins & Raives 2015), and no systematic exploration of MRI properties over the various standard flow and magnetic field configurations routinely examined with grid-based codes. This is despite the frequent use of smoothed-particle magnetohydrodynamics (SPH MHD) (Springel 2010a; Price 2012), particularly to probe galaxy, star, and planet formation (Price & Bate 2007, 2008, 2009; Dolag & Stasyszyn 2009; Dobbs et al. 2016; Forgan et al. 2016; Steinwandel et al. 2018). The neglect is

perhaps connected to several numerical problems specific to meshless methods, which we shall now discuss briefly.

In codes without meshes or with arbitrary mesh geometries, minimizing the divergence of magnetic fields (Tricco & Price 2012; Hopkins 2016a) is a major challenge. Because proper minimization of the divergence is hard to achieve, small “magnetic monopoles” can arise, leading to spurious magnetic field reconfiguration, reconnection, and artificial dissipation in neighboring domains. On fixed, rectilinear, regular, non-moving grids, the *Constrained Transport* (CT) scheme (Evans & Hawley 1988) can maintain zero divergence to machine precision. Until recently, CT schemes had only been implemented for regular, non-moving meshes, but recently Mocz et al. (2014, 2016) successfully generalized the CT method to moving meshes that adopt a Voronoi tessellation as their volume partition, e.g., those in AREPO (Springel 2010b). However, most Lagrangian or quasi-Lagrangian methods, including moving-mesh as well as particle-based methods or mesh-free finite-volume methods, use “divergence cleaning” schemes to keep $\nabla \cdot \mathbf{B}$ minimal (Powell et al. 1999; Dedner et al. 2002). Tricco & Price (2012) developed improved divergence-cleaning implementations in SPH (adapting the hyperbolic cleaning scheme from Dedner et al. (2002)), and showed that this could successfully reproduce some standard MHD tests (e.g., the Orszag-Tang vortex). However, in nonlinear MRI simulations—in fact, in any regime of MHD turbulence—effective divergence cleaning is especially difficult due to the complex, multiscale field geometry. Consequently, the latter methods are not guaranteed to work satisfactorily.

SPH also suffers from known numerical dissipation sourced by various terms, including the E0 error (Read et al. 2010), pairing instability (Dehnen & Aly 2012; Rosswog 2015), and incorrectly triggered artificial viscosity (Deng et al. 2017). It is well-known that this additional numerical dissipation impedes SPH’s capability to model subsonic turbulence, even *without* magnetic fields (Bauer & Springel 2012; Hopkins 2015; Deng et al. 2019).

The Lagrangian meshless finite-volume (MFV) method was developed two decades ago (e.g., Vila 1999; Hietel et al. 2000) and has been significantly improved since (e.g., Lanson & Vila 2008a, 2008b). Recently, it has stimulated growing interest in the astrophysical community (e.g., Gaburov & Nitadori 2011; Hubber et al. 2017). Hopkins (2015) generalized the method in Gaburov & Nitadori (2011) to other mesh-free finite-volume Godunov schemes, including the closely related “meshless finite-mass” (MFM) method. These methods, similarly to moving mesh methods, attempt to combine advantages of grid-based and particle-based codes. In particular, they can describe subsonic hydrodynamical turbulence reasonably well (with a quality comparable to regular mesh grid-codes; see Hopkins 2015), though at greater computational cost. They also avoid advection problems in complex flow geometries that are better modeled in the Lagrangian frame—respecting, for example, Galilean invariance. Finally, they can naturally extend to self-gravitating flows by exploiting accurate state-of-the-art gravity solvers, such as treecodes, which have native implementations in particle codes. Meshless methods have been generalized to MHD (Hopkins 2016a). In subsequent work, Hopkins (2016b) developed a constrained-gradient (CG) divergence cleaning scheme that can maintain a much smaller $\nabla \cdot \mathbf{B}$ (by ~ 2 orders of magnitude) compared to

hyperbolic divergence cleaning. These methods, as implemented in the public code GIZMO,⁵ have already been used to simulate the MRI in two-dimensional unstratified shearing sheets (Hopkins & Raives 2015), and these tests have demonstrated that it recovers the correct linear growth rates and behaves similarly to well-tested grid codes (e.g., ATHENA). However, how these methods perform in three dimensions, in stratified configurations, and/or during non-linear saturation, remain untested.

In this paper, we have carried out MRI simulations in both unstratified and vertically stratified shearing boxes, with both SPH and MFM MHD implementations (as they are implemented in the multimethod GIZMO code), in order to explore the numerical requirements for these methods to treat the MRI in the nonlinear regime. We focus on MFM, as opposed to MFV or more general moving-mesh schemes (several of which are also implemented in GIZMO and can, in principle, use the same CG divergence “cleaning” method), because MFM is designed, like SPH, to precisely conserve the mass of fluid elements (i.e., there is identically zero advection), so the method is “purely” Lagrangian. This is perhaps the most challenging case for our purpose, because hybrid moving-mesh or MFV-type methods, in which the grid moves but mass fluxes are also allowed, effectively act as a smoothing of grid motion, thus interpolating between the “pure Lagrangian” (constant mesh-motion) and the “pure Eulerian” (fixed-grid) representation of a fluid.

We will explore both the traditional “density-energy” formulation of SPH (named hereafter “TSPH”) (Springel 2005) and the more recently developed “pressure-energy” formulation (“PSPH”) (Hopkins 2012; Saitoh & Makino 2013). The SPH MHD used here represents the state-of-the-art implementation described by Price (2012), with the advanced artificial viscosity/resistivity switches developed in Cullen & Dehnen (2010) and Tricco & Price (2013), and divergence cleaning following (Tricco & Price 2012). In unstratified shearing box simulations, no significant density contrast is present, and we expect TSPH and PSPH to perform similarly. Therefore, we did not run TSPH and PSPH comparisons for this particular setup. In the GIZMO MFM runs, we adopt the CG divergence cleaning of Hopkins (2016b).

We note that the effect of the initial noise in MFV, which appears to depend on the regularity of the initial particle distribution (Gaburov & Nitadori 2011), is poorly understood. In general, the dependence of MRI properties on the numerical setup of the initial condition should be expected because MRI is extremely sensitive to numerical dissipation. Due to this added complexity of the initial condition design (although, in principle, MFV-type methods should be less prone to the effect of particle discretization noise during slope limiting as well as in the divergence cleaning step), we defer their scrutiny in the context of MRI to future work.

We start, in Section 2, with a discussion of our shearing box implementation and the role of the smoothing kernel function. We also tested the resolution needed for accurate MRI eigenmode growth. In Section 3, we present unstratified shearing box simulations with and without net vertical flux, and in short and tall boxes. Stratified shearing box simulations are described in Section 4, where we compare different simulation setups and the two methods (MFM and SPH). A

⁵ The public version of the code, containing all the algorithms used here, is available at <http://www.tapir.caltech.edu/~phopkins/Site/GIZMO.html>.

discussion and conclusion follow in Sections 5 and 6, respectively.

2. The Shearing Box Approximation

The shearing box is a local model of the equations of motion widely used in MRI simulations to achieve high resolution (Goldreich & Lynden-Bell 1965; Hawley et al. 1995; Latter & Papaloizou 2017). One considers a small patch of a disk centered at a radius R and rotating at the angular velocity $\Omega(R)$. In the corotating frame, one installs a Cartesian geometry at the box center, using x and y to represent the radial and azimuthal directions, respectively. In compressible ideal MHD, the governing equations are

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

$$\begin{aligned} \frac{\partial \mathbf{v}}{\partial t} + \mathbf{v} \cdot \nabla \mathbf{v} = & -\frac{1}{\rho} \nabla \left(P + \frac{B^2}{8\pi} \right) + \frac{(\mathbf{B} \cdot \nabla) \mathbf{B}}{4\pi\rho} \\ & + 2q\Omega^2 x \hat{\mathbf{x}} - \Omega^2 z \hat{\mathbf{z}} - 2\Omega \times \mathbf{v}, \end{aligned} \quad (2)$$

$$\frac{\partial \mathbf{B}}{\partial t} = \nabla \times (\mathbf{v} \times \mathbf{B}), \quad (3)$$

$$\frac{\partial \rho u}{\partial t} + \nabla \cdot (\rho u \mathbf{v}) = -P \nabla \cdot \mathbf{v}, \quad (4)$$

where $\hat{\mathbf{x}}$ and $\hat{\mathbf{z}}$ are the unit vectors in the x and z directions, and ρ , u , P , c_s , and $\mathbf{v}$ represent the density, specific internal energy, gas pressure, sound speed, and velocity, respectively. The tidal acceleration $2q\Omega^2 x \hat{\mathbf{x}}$ in Equation (2) comes from the expansion of the effective potential (gravitational plus centrifugal). The constant $q \equiv -d \ln \Omega / d \ln R$, and for a Keplerian disk, $q = 1.5$. The vertical component of the star's gravity is represented by $-\Omega^2 z \hat{\mathbf{z}}$; if included, this results in a vertical density stratification with a scale height of $H = c_s/\Omega$, where c_s is the initial sound speed. In simplified models examining motions confined near the disk midplane, this term can be dropped. The ratio between the gas pressure and magnetic energy, $\beta \equiv P/(B^2/8\pi)$, is a dimensionless measure of the magnetic field strength.

We assume an ideal gas equation of state (EOS),

$$P = \rho u (\gamma - 1). \quad (5)$$

We choose $\gamma = 5/3$, except when we set $\gamma = 1.001$ to mimic an isothermal EOS. In particular, we have one stratified simulation with $\gamma = 1.001$ to show how such a soft EOS exacerbates long-term numerical dissipation.

2.1. Shearing Box Boundary Conditions

The computation domain is a rectangular prism with sides of length L_x , L_y , and L_z . In unstratified boxes, the domain is periodic in y and z , and shear periodic in x . These boundary conditions can be expressed mathematically for a fluid variable f as

$$f(x, y, z) = f(x + L_x, (y - q\Omega L_x t) \bmod L_y, z), \quad (6)$$

$$f(x, y, z) = f(x, y + L_y, z), \quad (7)$$

$$f(x, y, z) = f(x, y, z + L_z). \quad (8)$$

They apply to the thermodynamic variables and the perturbed fluid velocity, where the background is $\mathbf{v}_0 = -q\Omega x \hat{\mathbf{y}}$. The

implementation of the shearing periodic boundary conditions in Lagrangian codes is relatively easy because we do not need to extrapolate fluid quantities to ghost zones, as in grid codes. When a fluid element (“particle”) moves across the radial boundary, it reappears at the other radial boundary with a velocity offset added to its azimuthal velocity.

In vertically stratified simulations, we apply an outflow boundary condition in the z direction by removing any element whose smoothing length is larger than $1.2H$. This yields a density floor of about 0.0002 in code units (see below).

2.2. Equilibrium Tests and the Kernel Function

The shearing box admits the following simple equilibrium: $\mathbf{v} = -\mathbf{v}_0$, $\rho = \text{constant}$. To test whether the code properly describes this state in addition to the shearing periodic boundary conditions, we conduct an MFM simulation using this as an initial condition. We employ a cubic box of one disk scale height per side. In the calculation, we choose units so that $\Omega = 1$, $c_s = 1$, and $\rho = 1$. At a resolution of $48 \times 48 \times 48$ elements with the Wendland C4 kernel (200 neighbors, $N_{\text{ngb}} = 200$), the equilibrium can be maintained to machine precision for the duration of the simulation ($\sim 200 \Omega^{-1}$) with no sign of breakdown.

We next reran the simulation using the cubic spline kernel (55 neighbors) and found that the radial velocity becomes *non-zero* and the perfect lattice breaks into a glass configuration. The velocity errors are a few percent of the sound speed. In this case, elements form pairs as shown in Figure 1. In SPH, this pairing (or clumping) instability (Springel 2010a; Price 2012) happens with any kernel whose Fourier transform is negative for some wave vectors, at a sufficiently large neighbor number (Dehnen & Aly 2012). Thus, it would appear that MFM also suffers a similar instability if “too many elements” are included in the kernel radius of compact support (i.e., if one does not, as one should, use higher-order kernels with a larger number of neighbors).⁶ The Wendland C4 kernel does not suffer from these issues at this “enclosed neighbor number” (see e.g., Dehnen & Aly 2012). It also helps to keep elements well-ordered, which is crucial for accurate gradient estimation in any unstructured method (Rosswog 2015). MRI turbulence is generally subsonic (except in the disk corona of stratified box), so we always use the Wendland C4 kernel to minimize numerical noise/dissipation (except when noted).

2.3. Resolution

Ideally, we would like to estimate: the shortest length scale adequately resolved in a simulation's h ; the number of resolution elements required to furnish that level of accuracy; and ultimately, the numerical resources required by a given code to achieve this.

Unfortunately, it is not straightforward to determine the shortest resolvable length scale, even for grid codes (which possess an unambiguous grid spacing δz). Particle codes exhibit several length scales that might be used as references in this

⁶ The interpretation of the pairing instability in MFM differs slightly from that in SPH—the kernel function is used in MFM to define the volume partition between neighboring resolution elements. If one takes a low-order kernel—say, the cubic spline—and forces its radius of compact support to enclose too many elements (such that the mean inter-element separation is much smaller than the kernel function width), the effective faces between elements essentially “overlap” into a single face (which becomes ill-defined).




Figure 1. Resolution element (“particle” or “mesh-generating point”) locations at $t = 8.4 \Omega^{-1}$ in a *steady state* MFM run with the cubic spline kernel (top) and the Wendland C4 kernel (bottom). Elements form pairs in the simulation using the cubic spline kernel while the Wendland C4 kernel simulation maintains nice element order (initially, cubic lattices are sheared)

context. Perhaps the most obvious choice is the mean particle separation $\bar{\delta}$. Of course, this is rather crude because particle codes are intrinsically adaptive, with denser regions resolved better than less dense regions. However, even if we set aside this concern, we must emphasize that resolution is not so much determined by the size of $\bar{\delta}$ as it is by the kernel volume, because this is where the interpolation occurs (which is instrumental in defining a volume element in MFM). We may then be tempted to use, instead of $\bar{\delta}$, the kernel support radius H_1 , because the particles within a kernel together form a “resolution unit”. More accurate estimates exist, however, as we now discuss.

In SPH, an alternative definition for the resolution scale is $h = 2\sigma$, where σ is the standard deviation of the weighting kernel function $W(\mathbf{x}, H_1)$, as defined in Dehnen & Aly (2012):

$$\sigma^2 = \frac{1}{3} \int dx x^2 W(\mathbf{x}, H_1). \quad (9)$$

We henceforth denote this length scale as h_{SPH} .

In MFM, for a well-chosen kernel (i.e., H_1 is within a factor of a few rms inter-particle separation, such that faces are well-defined), an alternative estimate of h is the face-area weighted inter-neighbor separation (where the face areas are themselves determined by the volume partition from the kernel function; see Hopkins 2015). Note that, consequently, h so defined is the scale upon which the Riemann problem is solved, different from both H_1 and $\bar{\delta}$. For the Wendland C4 kernel parameters adopted here, this yields $h \approx 0.45H_1$, so we will use this value throughout. We denote this by h_{MFM} .

2.4. MRI Quality Factor

In grid-code MRI simulations, the number of cells per fastest growing mode’s wavelength is often used as a resolution metric (Hawley et al. 2011; Parkin & Bicknell 2013). It can potentially tell us where and when in a simulation the MRI is adequately or inadequately resolved.

We may define a quality parameter (Noble et al. 2010; Hawley et al. 2011) via

$$Q_z = \frac{\lambda_{\text{MRI}}}{h} = \frac{2\pi V_{az}}{\Omega h}, \quad (10)$$

where V_{az} is the z component of the Alfvén velocity and h is the (vertical) resolution length. In the presence of a net vertical flux, λ_{MRI} is close to—but not exactly—the fastest growing linear mode’s wavelength, $\lambda_{\text{fastest}} = \sqrt{16/15} \lambda_{\text{MRI}}$. Perturbations with wavelengths smaller than $\lambda_{\text{MRI}}/\sqrt{3}$ are stable for the same configuration. The value of Q_z can be measured and averaged over the disk body during the saturated state. It can also be defined with respect to other coordinate directions, most often in the y direction. Regions where the plasma beta is high are of lower quality factor, and there may be insufficient resolution in these regions.

We must state from the outset that a quality factor, so defined, is a rather crude measure of how well the turbulence is resolved. First, it is based on the linear theory of the net-flux MRI setup and hence may not be generally applicable during the ensuing nonlinear saturation; certainly, its relevance for zero-net flux simulations is debatable. Second, it only describes whether the input scale of the turbulence is resolved; it has nothing to say about the ensuing turbulent cascade on smaller scales. If $Q_z \gtrsim 1$, then there is no inertial range to consider. Furthermore, in that case, the input and dissipative scales are directly adjacent: strictly, there is no real turbulence, but rather a monoscale chaotic flow. Nonetheless, vertically stratified shearing box simulations indicate that $Q_z > 10$, $Q_y > 20$ ensures the convergence with resolution of some large-scale average flow quantities (Hawley et al. 2011).

In our small number of grid-based simulations, we substitute $h = \delta z$ for the grid spacing. In our particle simulations, we use both $\bar{\delta}$ and h_{SPH} or h_{MFM} for comparison. We use Q_{sepr} to denote the quality factor when $\bar{\delta}$ is used for h , and Q^* denotes the quality factor when either h_{SPH} or h_{MFM} is used.

For MFM, the two quality factors are related by

$$\frac{Q^*}{Q_{\text{sepr}}} = \left(\frac{4\pi}{3N_{\text{ngb}}} \right)^{\frac{1}{3}} \kappa, \quad (11)$$

where $\kappa \equiv H_1/h_{\text{MFM}}$ and is related to the type of kernel used in the simulation. The cubic spline kernel can achieve good

density estimation with a small number of neighbors N_{ngb} (around 42), where we define

$$N_{\text{ngb}} = \frac{4\pi}{3} H_1^3 n(x_i) = \frac{4\pi}{3} H_1^3 (\rho_i/m_i), \quad (12)$$

$$n(x_i) = (\bar{\delta})^{-3}, \quad (13)$$

where m_i is the mass of the i th particle. The cubic spline yields Q^* and Q_{sepr} that are relatively close (the latter only a factor 0.9 of the other). However, the Wendland C4 kernel ($N_{\text{ngb}} = 200$) gives $Q^* = 0.6Q_{\text{sepr}}$.

2.5. Channel Flow Growth

The linear MRI eigenmodes in a net vertical flux are called channel flows. Being nonlinear solutions, these eigenmodes will grow to large and nonlinear amplitudes before being destroyed, in the first instance, by parasitic instabilities (Goodman & Xu 1994). A robust turbulence then ensues, sometimes exhibiting the recurrent generation and destruction of the channels. In this section, we measure the growth rates of the simulated channel flow in SPH and MFM, and compare them to the theoretical value. In particular, we calculate the error as a function of resolution (in fact, by Q), to see whether certain scalings can be discerned. For comparison, we also obtain error estimates for the finite volume Godunov code, ATHENA (Stone & Gardiner 2010), with second-order reconstruction and either the Roe and HLLD solvers. Our primary aim here is to compare the performance of the codes on this linear problem, not with respect to computational effort directly (which we address in detail in Section 5.1), but with respect to h . Is a plausible estimate of the minimum resolvable length scale a useful diagnostic for measuring the accuracy of MRI growth, or are other things going on? If the importance of h is overwhelming, we may reasonably expect all three codes to perform similarly on this test.

To begin, we initialize a box of size $H \times H \times H$, threaded by uniform vertical background fields of magnitude B_0 . We then set $\beta = 84$, such that the fastest growing channel mode just fits into the box, and $\gamma = 1.001$ so the gas is effectively isothermal. The initial amplitude of the channel mode is $0.001 B_0$, and the theoretical growth rate of the fastest channel mode is $0.75 \Omega^{-1}$. The simulations are run for $8 \Omega^{-1}$, so the channel mode is smaller than $0.4 B_0$ at the end of the simulation. We calculate the growth rate using the magnitude of the magnetic field at two consecutive snapshots taken every $0.5 \Omega^{-1}$. The growth rate relative error is defined as

$$\max\{(s_i - 0.75)/0.75\} \quad (14)$$

where s_i is the i th measure of the growth rate.

In Figure 2, we plot the growth rate error as a function of Q_{sepr} , i.e., the (mean) number of resolution elements per λ_{MRI} , be they grid cells (in ATHENA) or the average particle spacing (in SPH and MFM). In MFM, the Wendland C4 kernel is employed in order to avoid pair instability. As is clear, and to be expected, the errors decline with increasing resolution. They are less than 1% when $Q_{\text{sepr}} > 32$ in MFM. TSPH captures the MRI better than MFM in the low-resolution simulations, but it converges more slowly. It is known that SPH has zeroth-order errors that only vanish when both Q_{sepr} and N_{ngb} approach infinity (N_{ngb} is fixed here) (Read et al. 2010; Zhu et al. 2015).




Figure 2. Growth rate errors in different MHD schemes vs. the quality factor Q (number of resolution lengths h per MRI wavelength) for various codes and different definitions of h . The solid red line indicates the performance of the ATHENA code for three different resolutions; here, $h = \delta z$. The solid blue line represents MFM and the solid orange curve TSPH, both using the mean particle separation $\bar{\delta}$ for h ; i.e., $Q = Q_{\text{sep}}$ and equals the mean number of particles per MRI wavelength. The dashed lines represent the performance of MFM and TSPH when $Q = Q^*$ (see Section 2.4).

ATHENA performs better than SPH or MFM, if we are to directly compare number of grid points against number of particles. However, as explained in the previous two subsections, even if such a comparison may reflect relative numerical effort, it does not do justice to the particle codes; for them, the mean particle spacing is not equal to the shortest resolvable length scales. In Figure 2, we also plot the relative growth rate error for both SPH and MFM versus Q^* , the number of resolution lengths h per MRI wavelength (less than the number of particles). These dashed curves now sit on top of or adjacent to each other. If we are to consider resolution this way, the codes actually perform comparably; ATHENA and MFM (which share the same Riemann solver) agree exceptionally well. In comparison, SPH possesses a different and adverse scaling, indicative of its sensitivity to not only the number of particles but also to the number of near neighbors. We will see this poorer performance again later in nonlinear tests of MRI saturation.

Finally, a few words dealing with the cubic spline kernel in MFM, which yields a larger effective quality parameter Q^* for smaller N_{ngb} (see Section 2.3). Indeed, it outperforms the Wendland C4 kernel when the resolution is low. However, when $Q_{\text{sepr}} > 20$, the growth rate errors increase due to the pairing instability (which itself is resolution-dependent; see Dehnen & Aly 2012). This numerical noise (see Section 2.2) can dominate over the weak channel mode in the early stage. The channel modes eventually outcompete this noise, but the errors in gradient estimation lead to extra dissipation (see Section 4.2) that is hard to quantify.

3. Unstratified Shearing Box Simulations

We summarize all the simulations we undertook in both unstratified and stratified boxes in Table 1, which includes key parameters, physical and numerical configurations, comments, and references. Further details can be found in the referenced subsections.

Table 1
Simulations, Results, and Comments

Simulations	Initial fields	Box size	Resolution	MHD-methods	EOS	Sections/Figures	Reference
Unstratified	NZ	$H \times 6.28H \times H$	$64 \times 360 \times 64$	MFM	adiabatic	Section 3.1/Figure 3	1
		$H \times 4H \times H$	48 elements per H	MFM/PSPH			
	ZNZ	$H \times \pi H \times H$	$64 \times 200 \times 64$	MFM	isothermal	Section 3.2/Figure 4	2
		$H \times 4H \times 4H$	48/64 elements per H	MFM	adiabatic(c)	Section 3.2.1/Figure 5	3
Stratified	B_y	$\sqrt{2}H \times 4\sqrt{2} \times 24H$	1.5 M elements	TSPH/PSPH	adiabatic	Section 4.2/Figure 7	4
				MFM	ad/iso	Section 4.3/Figure 9	
			3 M elements	MFM	adiabatic(c)	Section 4.4/Figure 11	

Note. The following abbreviations have been used: NZ—Net vertical flux; ZNZ—Zero net vertical flux; MFM—meshless finite mass method with constrained gradient divergence cleaning; TSPH—Density-energy (traditional) formulation of SPH; PSPH—Pressure-energy formulation of SPH. Here, we always use the Wendland C4 kernel, except one experiment run with the quartic spline kernel in Figure 9. In stratified shearing box simulations, elements with a smoothing length larger than $1.2H$ are clipped, resulting in a density floor of ~ 0.0002 . Both SPH MHDs employ the Cullen & Dehnen artificial viscosity switch (Cullen & Dehnen 2010), the hyperbolic divergence cleaning of Tricco & Price (2012), and the artificial resistivity of Tricco & Price (2013). Furthermore, adiabatic runs use $\gamma = 5/3$, isothermal runs $\gamma = 1.001$, and adiabatic(c) runs apply an ad hoc cooling (see Equation (19)). We expect very fast turbulence decay due to numerical dissipation using isothermal EOS (see Figure 9 and Section 4.3). We always try to use an adiabatic EOS to minimize long-term numerical dissipation, except when we want to enable direct comparison with previous studies.

References. 1. Hawley et al. (1996); 2. Fromang & Papaloizou (2007); 3. Shi et al. (2015); 4. Davis et al. (2010).

3.1. Net-vertical-flux Simulations

We first ran an unstratified shearing box simulation with net vertical flux similar to the fiducial model of Hawley et al. (1995). This is the simplest 3D MRI setup. For such a configuration, we are able to reproduce the main features of previous grid-based simulations using high-resolution MFM simulations. The box is of size $H \times 6.28H \times H$ and threaded by vertical fields with $\beta = 400$. We used a resolution of $64 \times 360 \times 64$ elements, which corresponds to 28 elements per λ_{MRI} , and hence roughly 17 h_{MFM} per λ_{MRI} (or roughly 9 H_1 per λ_{MRI}). In comparison, Hawley et al. (1995) employed about 16 grid points per MRI wavelength. The box size also affects the simulated turbulence: the smaller box tends to have stronger outbursts in the turbulent state (Bodo et al. 2008; Lesaffre et al. 2009). We set two other runs in a box of $H \times 4H \times H$ with $\beta = 330$ ($H = 2\lambda_{\text{MRI}}$), and with either PSPH and MFM using 48 elements per H . We added random velocity perturbations (5% of the sound speed) to the shear flow at initialization. The simulations are run for 11 orbits with $\gamma = 5/3$.

To characterize the saturated turbulence, we plot several density-weighted, averaged quantities in Figure 3. We take the arithmetic average of fluid variables at all the MFM fluid elements so that the average is naturally density-weighted because of the adaptive nature of GIZMO. The density fluctuations are small in unstratified turbulence, so the density-weighted average should be close to the volume average in previous studies. This should be the case also for stratified turbulence, another situation in which we will apply this averaging method (see Section 4), because the stress is almost independent of the density when $|z| < 2\sqrt{2}H$ (Simon et al. 2011). In the $L_y = 6.28H$ simulation (red curves), both the magnetic energy and stresses are in agreement with the results of Hawley et al. (1995), which were obtained with a Eulerian code. The ratio of the Maxwell stress to the magnetic pressure,

$$\alpha_M = \frac{\langle -2B_x B_y \rangle}{\langle B^2 \rangle}, \quad (15)$$

is about 0.5, similar to the aforementioned previous adiabatic calculations. The agreement between the older ZEUS runs and MFM may have been expected because the resolution is similar in the two (about 16 resolution lengths, either δz or h_{MFM} , per MRI wavelength). It should be noted, however, that the ZEUS code is more diffusive than ATHENA—our point of comparison in the linear growth tests.

The simulations in smaller boxes, $L_y = 4H$, show stronger bursts in stresses and magnetic energy because fewer active (non-axisymmetric) modes can fit in the box, leading to an artificial truncation of the participating modes in the nonlinear dynamics. As a consequence, the system is nearer to criticality, and single-channel modes can intermittently dominate.

The internal energy increases due to the turbulent dissipation. This is most significant in the two $L_y = 4H$ simulations, as they exhibit the strongest bursts from the channel flows: because these flows achieve large amplitudes, a great deal of energy is dissipated into heat when they break down. The PSPH simulation, in particular, is some four times “hotter” than the large-box MFM simulation. The PSPH run also undergoes a much greater increase of internal energy compared to the equivalent MFM run (see Figure 3, upper right panel) due to stronger channel activity near the beginning of the run (signaled by the very large initial spike in the various diagnostics shown in Figure 3). The dominance of channels in the PSPH run early on suggests that the system is closer to marginal stability than in MFM; this is probably due to the additional numerical diffusivity in PSPH. Note that the plasma β increases substantially as the gas is heated up, and the boxes will ultimately approach the incompressible zero-net-flux regime. This explains why, in general, the bursts become less powerful as the simulations continue. Finally, for the net-flux unstratified setup, we ran only PSPH and not TSPH; because there are no steep density gradients, we expect no significant difference due to the actual formulation of the SPH hydro force.

3.2. Zero-net-vertical-flux Simulations in a “Standard” Box

We run a standard zero-net-flux unstratified box simulation with MFM with no explicit physical dissipation terms (Stone et al. 1996; Fromang & Papaloizou 2007). We initialize a box




Figure 3. From top left to bottom right corner, the time evolution of the averaged magnetic energy, kinetic energy, thermal energy, Maxwell stress, α_M , and Reynolds stress in the unstratified vertical flux simulations are shown (see text for the explanation of how average quantities are computed). Time is given in orbits. Here, P_0 is the initial pressure. The MFM simulation with $L_y = 6.28H$ (red curves) gives results close to those of Hawley et al. (1996). The PSPH simulation with $L_y = 4H$ (black curves) has larger internal energy than the two MFM simulations. We note that the increasing internal energy leads to a larger plasma β and smaller outbursts.

of size $H \times \pi H \times H$ with $64 \times 200 \times 64$ elements, threaded by magnetic fields,

$$\mathbf{B} = B_0 \hat{z} \sin(2\pi x/H). \quad (16)$$

The field strength B_0 is chosen so that the volume-averaged β equals 400. We use an isothermal EOS ($\gamma = 1.001$) to align with Fromang & Papaloizou (2007). MRI turbulence is sensitive to the nature of both physical and numerical dissipation: without physical viscosity and resistivity, Fromang & Papaloizou (2007) found that zero-net-flux MRI turbulence was driven to smaller scales as resolution increased with no sign of convergence. Fromang et al. (2007) showed that the saturated state depended on the magnetic Prandtl number when a source of diffusivity—physical or numerical—is present; if this was too low, turbulence would die after some period of time. With these problematic results in mind, we will now assess how well a zero-net-flux MRI setup can be modeled by a meshless code.

We plot the averaged magnetic energy, Maxwell stress, and α_M in Figure 4. In contrast to MRI runs with grid codes, the magnetic fields and magnetic stress rapidly decay after an initial burst. There is no sustained turbulence, as in Fromang & Papaloizou (2007), nor is there some period of MRI turbulence before decay, as in Fromang et al. (2007). It is true that MFM smooths fluid variables within a kernel, so the resolution (with respect to h_{MFM}) is roughly half that of the standard simulation with $64 \times 200 \times 64$ grid cells in Fromang et al. (2007) (see Section 2.3). A simulation with >128 particles per scale height is prohibitively expensive with MFM (see Section 5.1).

However, even an “ideal MHD” run undertaken with low resolution in a grid code can sustain MRI turbulence (Stone et al. 1996). Therefore, our result is disappointing.

One way to interpret this is to consider the relative sizes of the numerical resistivity and viscosity. At low resolutions, MFM should have a moderate numerical viscosity, ν (see Appendix A), and relatively large numerical resistivity, η (see Appendix B). As a consequence, the effective numerical Prandtl number $P_m = \nu/\eta$ must be small (less than 1). It is perhaps not surprising, then, that the turbulence decays (Fromang et al. 2007). However, the fact that it decays so abruptly might point to a simpler cause: the numerical resistivity is very large and simply prohibits turbulence of any kind past the initial spike. Indeed, in some respects, our MFM run resembles the Fleming et al. (2000) run with a magnetic Reynolds number of 13,000, which abruptly dies off after an initial burst.

3.2.1. Zero-net-vertical-flux Simulations in a Tall Box

It has been shown that, when the numerical domain is reshaped such that it exhibits a large vertical aspect ratio ($L_z/L_x \geq 2.5$), a new, more vigorous and cyclical MRI dynamo emerges (Shi et al. 2015) (see also Lesur & Ogilvie 2008). Importantly, its saturated stress is independent of resolution. To test the effect of a taller box, we redo the simulations in Section 3.2 in a box of size $H \times 4H \times 4H$. We present two simulations with 48 and 64 elements per scale height; their details are shown in Figure 5. We set $\gamma = 5/3$ and add a cooling term in order to keep the internal energy roughly




Figure 4. Time evolution of averaged magnetic energy, Maxwell stress, and α_M in the zero-net-flux MFM simulation. Time is given in orbits. The magnetic field decays quickly and the Maxwell stress becomes nearly zero after about 20 orbits (see text).




Figure 5. Time evolution of α_M and averaged magnetic energy in the zero-net-vertical-flux tall-box MFM simulation are shown in the upper panels. In the lower panel, the temporal evolution of the averaged horizontal magnetic field of the Nz64 simulation is presented. The simulation (Nz64) with 64 elements per scale height shows a sign of convergence comparable to the fast decay of magnetic fields in the short box zero-net-flux simulation in Figure 4. The pattern of the averaged azimuthal field is also similar to that of Shi et al. (2015). However, the magnetic fields eventually decays.

constant (see Equation (19)). The numerical dissipation is EOS-related and we expect very fast decay with the isothermal EOS (see Section 4.3). This does mean, however, that we cannot make direct quantitative comparison with Shi et al. (2015).

Initially, our simulations exhibit turbulence as shown in Figure 5. Moreover, they reproduce the averaged toroidal field patterns produced by Shi et al. (2015). However, while the turbulence is sustained for much longer than in the standard

box, the activity ultimately dies out after some 30–40 orbits. During the turbulent phase, $\alpha_M \sim 0.44$ in the higher-resolution tall box simulation, but the averaged magnetic energy $\langle B^2/8\pi \rangle/P_0 \sim 0.01$ is much lower than the values (>0.1) obtained with the ATHENA code (Shi et al. 2015). As a result, the stress ($\sim 0.006 P_0$) is also much smaller.

In tall box simulations, when the magnetic Prandtl number $P_m \geq 4$, the saturated stress is independent of P_m while the

turbulence vanishes for $P_m = 1$ with even 128 cells per scale height (Shi et al. 2015). While P_m effects might be at play in our simulations, it should be noted that our 64 particle simulation is quite low-resolution (elements are further smoothed within a kernel). Worse resolution leads to more rapid decay. It is likely that both meshless simulations possess too great a numerical resistivity to support a sustained MRI dynamo.

To summarize, the MFM simulations—at least with the current implementation of the method—do not allow us to sustain the MRI in unstratified zero-net-flux simulations, either in short or tall boxes. This is probably the result of either too large a numerical resistivity, or more generally, too low an effective P_m , at least at the resolutions we were able to access (see Section 5.1). It will be particularly interesting to explore both (a) higher-resolution MFM simulations (where the resistivity should be lower and P_m larger), and (b) simulations using other, closely related schemes that are not completely “fixed-mass” but instead closer to moving-mesh schemes (e.g., the MFV scheme or arbitrarily shearing-mesh schemes with the CG divergence cleaning). We did not run zero-net-flux simulations with SPH, due to evidence suggesting that it has an intrinsically larger numerical viscosity that would cause it to yield higher Prandtl numbers, quite irrespective of the specific implementation of artificial viscosity (Bauer & Springel 2012; Deng et al. 2017, 2019), and because it produces substantially larger magnetic field divergence with available divergence cleaning methods (see Section 4.1).

4. Stratified Shearing Box Simulations

In this section, we undertake simulations in the vertically stratified shearing box, in which the (leading-order) vertical acceleration from the central star’s gravity is incorporated. We initialize the simulations with a Gaussian density profile with uniform temperature

$$\rho(z) = \rho_0 \exp\left(-\frac{z^2}{2H^2}\right), \quad (17)$$

where $H = c_s/\Omega$ is a factor different from $\sqrt{2}c_s/\Omega$ in some previous work (Davis et al. 2010; Simon et al. 2011). Note that c_s is the initial sound speed; in adiabatic runs, the sound speed (and hence scale height) will change. We adopt units such that $H = 1$, $\Omega = 1$, $c_s^2/\gamma = 1$ and use the adiabatic EOS with $\gamma = 5/3$. The density profile is sampled using the Monte Carlo method and then relaxed to a glassy configuration. The density errors in the disk body ($-3H < z < 3H$) are below the 1% level. We initialize azimuthal magnetic fields with $\beta = 25$ in a box of size $\sqrt{2}H \times 4\sqrt{2}H \times 24H$ (the box is extremely tall, but no element has $z > 6$ in our simulations; see Figure 12). Outflow boundary condition are applied, but there is, in fact, no significant outflow and few elements are clipped. Random velocity perturbations $\sim 0.01c_s$ are added to seed the instability.

In our fiducial model, we use 1.5M elements in total, leading to $\delta \sim 0.03$ at the disk midplane. However, due to the adaptive feature of our method, the resolution is lower further away from the midplane. This helps to save some computational resources, because high resolution is not needed in the MRI-stable disk corona with strong fields (see Figure 12 and also Miller & Stone 2000). However, the nearly zero-flux MRI turbulence in

the disk body still requires high resolution and is computationally demanding.

4.1. Divergence Cleaning of Magnetic Fields

Both SPH and MFM are unable to strictly maintain *exactly* solenoidal magnetic fields naturally, and thus must employ cleaning schemes to keep their divergences minimal. We try to quantify the efficacy of this procedure in this section before showing our main results.

We define the dimensionless divergence of magnetic fields as

$$\text{divB} = \frac{h|\nabla \cdot \mathbf{B}|}{B}. \quad (18)$$

In our unstratified box simulations of Section 3, the divB diagnostic is smaller than 10^{-3} at the location of most fluid elements in MFM. Divergence control in the stratified shearing box MRI is more challenging, however. We run the fiducial model to compare the level of non-zero divergence in MFM with the CG cleaning (Hopkins & Raives 2015; Hopkins 2016a) and TSPH with the hyperbolic divergence cleaning (Tricco & Price 2012). In Figure 6, the hyperbolic cleaning keeps $\text{divB} \sim 0.1$ in TSPH while the CG cleaning keeps divB two orders of magnitude lower in MFM. Large divB only occurs at the vertical boundaries and in the weak field regions in MFM; the vertical boundaries are poorly resolved because MFM fluid elements, which are built from particles, are fewer—but there, the divergence should have negligible influence on the turbulent disk body because the correlation length of magnetic fields is smaller than H (Davis et al. 2010; Bai & Stone 2013). We also stress that the maps shown in Figure 6 are quite representative of the differences between SPH and MFM in our tests.

Summarizing, the CG cleaning method in MFM significantly outperforms its competitors here, and we shall see how this is important in the following subsection. As a word of caution, we note that SPH methods that are recast at least partially in a finite-volume formulation, such as Godunov-SPH (Inutsuka 2002), might be amenable to implementations of the GC cleaning method. It would be interesting to explore the latter avenue in order to find out how much better an SPH method can perform once it is equipped with a superior divergence cleaning scheme.

4.2. Unphysical Behavior in SPH Simulations

We present three SPH simulations, two of which are run with the Wendland C4 kernel and adiabatic EOS but using different SPH formulations. In Figure 7, we plot α_M and the scaled magnetic energy versus time, in addition to spacetime diagrams of the horizontally averaged toroidal field. As is clear, the TSPH and PSPH simulations provide similar results. At first, the MRI grows and expels the initial azimuthal fields to the disk corona, where strong fields accumulate and are amplified (at about $\sim 60 \Omega^{-1}$). Domains dominated by magnetic energy propagate from the corona to the disk midplane, and the entire box ultimately is dominated by strong, growing azimuthal fields ($\beta \sim 1$). Note that α_M is negligible from some 10 to 20 orbits, indicating that the MRI is quenched. At the end of the simulation, the magnetic fields are at equipartition with the gas pressure and almost entirely azimuthal, with no turbulent activity. Simultaneously, the disk expands vertically as it




Figure 6. Typical $\text{Log}_{10}(\text{divB})$ value in TSPH (PSPH shows similar results) and MFM. MFM with CG divergence cleaning maintains divB two orders of magnitudes smaller than the hyperbolic divergence cleaning scheme used in TSPH. In MFM, most elements have $\text{divB} < 0.01$. Large divB only occurs in the weak field ($\beta > 1000$) regions.

becomes magnetically supported. Obviously, these simulations bear little resemblance to previous grid-code-stratified shearing box simulations, which report robust subsonic turbulence in the disk body (Davis et al. 2010; Simon et al. 2011). To test whether the strong fields persist in a more numerically dissipative setup, we reran the TSPH simulation with the quartic spline kernel (which has a noisier element distribution) and an isothermal EOS (see Figure 9). We find that this TSPH variant actually damps the turbulence faster and reaches the $\beta \sim 1$ state earlier.

We should note that SPH growing strong toroidal fields is not a result unique to our simulation setup or code. Dobbs et al. (2016), using the SPHNG SPH code to simulate global galactic disk models, also reported unaccountable growth of magnetic fields. Likewise, similar behavior has also been seen in MHD-SPH simulations of disk formation in tidal disruption events with the code PHANTOM (C. Bonnerot 2019, private communication). Both of these codes also implement the hyperbolic cleaning method from Tricco & Price (2012) in SPH. Stasyszyn & Elstner (2015) present a detailed study that discusses the likely numerical issues: they consider 3D global simulations of a differentially rotating disk with an initially pure-toroidal field, designed such that the system is stable and should exhibit no field growth. Using more accurate (CT or vector-potential-based) schemes, they recover this solution. However, using SPH with similar hyperbolic divergence cleaning, they show that discretization error produces small radial field components, which couples to the rotational shear and amplifies this—and in turn, the toroidal field—exponentially. They specifically show that the form of the SPH MHD induction equation leads (in essentially any internally consistent, SPH-based cleaning scheme) to the divergence-cleaning *amplifying* the vertical field, instead of *damping* the radial field, in order to locally restore $\nabla \cdot \mathbf{B} = 0$.

This demonstrates a few key ingredients that interact here: the particularly virulent form of this instability in SPH requires shear/differential rotation (either in global disk simulations or shearing boxes), non-zero radial, azimuthal, and vertical field components where there is a vertical gradient present that can offset the radial gradient (hence the 3D stratified simulations), and relatively large $\nabla \cdot \mathbf{B}$ errors (note these are large here, with $\text{divB} \sim 0.01\text{--}0.1$).

We should also note that it is possible to construct divergence-cleaning schemes, such as the one in Tricco & Price (2012), that are total-energy conserving. In highly idealized test problems, this will serve to limit the nonlinear magnitude of any erroneous magnetic field amplification. However, in a shearing box or global thin disk simulation, there is an essentially infinite source of energy from shear, so this does not “rescue” the simulations from excessive numerical dissipation.

More generally, it is well-known that, without *any* divergence-cleaning, the $\nabla \cdot \mathbf{B}$ errors are violently numerically unstable: magnetic monopoles grow explosively and the amplitude of $\mathbf{B}$ is correspondingly rapidly amplified. It is also well-established that this artificial, explosive field growth can occur even with divergence-cleaning, if the cleaning is not sufficiently accurate, or if it acts “too slowly” to respond to the growth rate. For example, Mocz et al. (2016) showed that using just the (considerably less-sophisticated) Powell et al. (1999) type of divergence cleaning, even in ordered meshes, produces large artificial magnetic field growth (on essentially the Courant timescale) and much larger magnetic field strength in idealized tests, as compared to CT methods.

Regarding the damping of turbulence; we should, of course, note that SPH requires artificial viscosity and resistivity to capture MHD shocks (Cullen & Dehnen 2010; Tricco & Price 2013, and references therein). It is well-known that SPH tends to overdamp subsonic turbulence due to imperfectly




Figure 7. Evolution of the magnetic fields in the SPH MHD stratified shearing box simulations. The upper panels show the time evolution of the magnetic energy and α_M as noted. The lower panels show the time evolution of the horizontally averaged azimuthal magnetic fields. Strong toroidal fields grow through shear amplification of radial fields. Secondary instability cannot develop efficiently in the low-resolution disk corona, and the strong toroidal fields spread gradually to the disk midplane. The stratified box is eventually filled with strong toroidal fields ($\beta \sim 1$, stable to MRI) and the disk expands vertically. The PSPH and TSPH simulation are almost identical because their artificial viscosity damps subsonic turbulence similarly (Bauer & Springel 2012; Hopkins 2015). We add a more dissipative TSPH simulation, with the quartic spline kernel and isothermal EOS (see Figure 9). The TSPH-variant simulation does not dissipate the strong toroidal fields; rather, it grows the fields even quicker due to larger numerical noise.

triggered artificial viscosity (Bauer & Springel 2012; Hopkins 2015; Deng et al. 2019). The GIZMO code applies an artificial viscosity switch similar to that described in Cullen & Dehnen (2010) with $\alpha_{\min} = 0.05$ and $\alpha_{\max} = 2$ (see Appendix F2 of Hopkins 2015 for details) to suppress unwanted artificial viscosity. This switch works most efficiently in regions away from shocks and may not be effective in regions with large velocity derivatives (Deng et al. 2017). In our SPH simulations, relatively large artificial viscosity with $\alpha_{\text{sph}} > 0.2$ is still triggered (see Figure 8). Artificial viscosity certainly helps the turbulence dissipate. With damped velocity fluctuations, the MRI and its parasitic modes (Latter et al. 2009; Pessah & Goodman 2009) cannot grow efficiently. We therefore also explore what happens if we revert to the more dissipative artificial viscosity in GADGET2 (Springel 2005) and restart the TSPH simulation from $t = 50 \Omega^{-1}$ (see Figure 12); the turbulence does decay faster (as expected) and the strong toroidal fields develop more quickly. Thus, as expected, the MRI turbulence damping is due significantly to the artificial viscosity.

Artificial resistivity dissipates magnetic fields, and a switch to minimize artificial resistivity away from shocks was

developed by Tricco & Price (2013). Here, we apply this artificial resistivity switch with $\alpha_{B,\min} = 0.005$ and $\alpha_{B,\max} = 0.1$. We choose this conservative $\alpha_{B,\max}$ because the turbulence is subsonic (Hopkins & Raives 2015). We have rerun the TSPH simulation in Figure 7 with $\alpha_{B,\max} = 1$ (suggested by Tricco & Price 2013) and obtained similar results. The numerical resistivity in SPH MHD is evidently different from that of Riemann solvers (see Appendix B).

4.3. A Transient MRI Dynamo in MFM Simulations

Our MFM simulations use the same initial conditions as those of the SPH simulations. We present three simulations. The fiducial model is run with the Wendland C4 kernel and adiabatic EOS. In addition, to test the effect of the kernel function and EOS, we run two simulations with the quartic spline kernel ($N_{\text{ngb}} = 60$) and with an isothermal EOS. In the isothermal run, we solve the energy equation instead of dropping it as done in Stone et al. (2008). To mimic the isothermal EOS, we set $\gamma = 1.001$ so that the thermal energy dominates the total energy and large truncation errors affect the accuracy of magnetic energy calculation. This can cause the fast dissipation of the magnetic fields.




Figure 8. Artificial viscosity α_{sph} parameter in the PSPH run at $50 \Omega^{-1}$ (see also Figure 12). Relatively large artificial viscosity is triggered in the disk body, despite the fact that there are no shocks.

In Figure 9, various quantities are plotted as functions of time. We see here that the MRI grows faster and reaches the magnetic pressure maximum quicker in the simulation with the quartic spline kernel (aqua) as compared to the other two simulations. The quartic spline kernel is more compact and has a larger quality factor, Q^* , than the Wendland C4 kernel (see Section 2.3). Although Dehnen & Aly (2012) show that the quartic spline kernel is superior to the traditional cubic spline kernel, it is still vulnerable to pair instability, which introduces numerical noise in gradient estimation. Because we find that the MRI dies quickly after its initial spike, it would appear that this noise provides a significant degree of numerical dissipation. The isothermal and fiducial adiabatic runs are qualitatively similar: both can sustain MRI turbulence for a period of some 30–40 orbits before dying.

In the first 30 orbits, the fiducial model successfully reproduces the quasiperiodic (~ 10 orbits) butterfly pattern of the averaged azimuthal fields. Note, however, that the butterfly diagram becomes erratic at $\sim 200 \Omega^{-1}$, as seen in other thermal MRI runs (Gressel 2013; Riols & Latter 2018). The saturated $\alpha_M \sim 0.4$ and $\langle B^2/8\pi \rangle/P_0 \sim 0.01$, however, are both in agreement with previous isothermal grid codes’ results (Simon et al. 2011). The box expands vertically during the simulation, due to accretion heating leading to a decrease of resolution in the disk body. In Figure 10, the density at the disk midplane drops to 0.6 at $200 \Omega^{-1}$, which corresponds to a 1.2 times larger mean fluid element separation. The decrease in resolution certainly must affect the sustainability of the turbulence. We next turn to higher-resolution simulations to assess whether our results can improve.

4.4. High-resolution MFM Runs

In order to maintain good resolution over the course of the simulations, thus avoiding expansion resulting from heat transport triggered by turbulence, we add an ad hoc cooling

term as in Noble et al. (2010) and Parkin & Bicknell (2013),

$$\frac{du_{\text{cool}}}{dt} = -\frac{u - u_{\text{init}}}{\tau_{\text{cool}}} \quad (19)$$

where $\tau_{\text{cool}} = 2\pi/\Omega$, and u_{init} is the initial specific internal energy constant. This fast cooling maintains the disk scale height nearly constant, thus preserving the initial resolution across the disk. In addition, we increase the number of elements to three million, which results in $\langle Q_y \rangle \sim 30$, $\langle Q_z \rangle \sim 10$ in the turbulent state.

Various flow properties are plotted in Figure 11. The most important result is that the MRI turbulence is sustained for a longer time (as it should be if the method is converging properly). During this phase, the main flow diagnostics are in good agreement with those of grid code runs: $\alpha_M \sim 0.4$, the averaged magnetic energy is a few percent of the gas pressure, and the Maxwell stress about four times larger than the Reynolds stress (Hawley et al. 2011). The butterfly diagram is reproduced in Figure 11, but the pattern becomes erratic after $300 \Omega^{-1}$. Compared to the MFM fiducial model with 1.5M elements (Figure 9), finer magnetic field structures are captured (see Figure 12) and the butterfly diagram/dynamo is resolved better. We note that even this simple cooling can introduce additional numerical noise at the kernel scale (Rice et al. 2014). However, we cannot afford higher resolution, or to run longer simulations, with this setup (see Section 5.1).

5. Discussion

5.1. Computational Cost and Possible Applications

In addition to the robustness of numerical results, another worthy metric of comparison between codes is their computational cost to carry out a comparable calculation. We run a setup equivalent to our local stratified simulation (the one containing 3M elements) with ATHENA and directly compare the computational costs, using a fixed timestep in both. We found that the MFM simulation is >100 times more computationally expensive than an equivalent run with 32 cells per scale height using ATHENA with the orbital advection method (Masset 2000; Stone & Gardiner 2010) for optimization. The lower computational efficiency of MFM has nothing to do with the hydro solver; rather, it is due to the neighbor “search tree,” which needs to be updated constantly and walked to find neighbors and rebuild the domain (because it allows for arbitrary particle reconfiguration between time-steps). Of course, in simulations where particle order is not dramatically changing and the only forces are local, we could (in principle) save considerable computational expense by simply storing the interacting neighbor lists and rebuilding the domain less often. Furthermore, there is room to significantly improve the neighbor search algorithm on modern massively parallel architectures coupled with accelerators, as is being currently investigated for a range of particle-based codes (Guerrera et al. 2018).

Nevertheless, such a difference in performance is highly problem-dependent. The tree algorithm can be efficiently exploited—and the difference in performance is dramatically mitigated—by the calculation of other physics involving non-local forces. A prime example of the latter is self-gravity. The tree-based gravity solver coupled with MFM, generally speaking, is indeed both faster and more accurate than




Figure 9. Evolution of the magnetic fields in the MFM stratified shearing box simulation with 1.5M elements and different setups. The upper panels are the time evolution of α_M and the magnetic energy as noted. The lower panels show the time evolution of the horizontally averaged azimuthal magnetic fields. In the fiducial run, both the saturated α_M and averaged azimuthal field pattern (butterfly diagram) agree well with previous grid-code simulations (Hawley et al. 2011; Simon et al. 2011) in the 30 early orbits. The fields decay later, partially due to the expansion of the shearing box and consequent decrease of the resolution (see Figure 10). The simulation with the quartic spline kernel cannot reproduce the butterfly diagram, due to numerical noise at the kernel scale, and the magnetic fields decay rapidly. The isothermal stratified shearing box with $\gamma = 1.001$ does not expand vertically, and thus it maintains the resolution. However, the truncation error in the energy equation eventually leads to magnetic field dissipation.




Figure 10. Vertical density profile of the fiducial stratified shearing box MFM simulation at different times. The box expands vertically. At $200 \Omega^{-1}$, the midplane density drops to 0.6 and the mean element separation becomes 1.2 times larger.

traditional gravity solvers coupled with grid-based codes, which is why particle-based methods with tree-based gravity have long been very competitive with grid-based methods in the modeling of self-gravitating protoplanetary disks (e.g., Mayer & Gawryszczak 2008). Furthermore, addressing self-gravitating disks ultimately requires global calculations (Durisen et al. 2007). This is, by itself, a natural regime for mesh-free codes, because one of their major goals is to enable adaptive resolution on global problems, more akin to adaptive mesh refinement (AMR) codes, which have similar computational and memory cost.

The most interesting applications of the mesh-free methods studied here are thus not in idealized MRI setups where accuracy of the MHD calculation over long timescales, absent other physical effects, is the prime objective. Rather, these methods may be more promising for studies of turbulence in magnetized self-gravitating disks, especially the strong dynamo action reported by Riols & Latter (2018). This spiral wave dynamo is vigorous even with large magnetic resistivity and may be responsible for the primordial magnetic field amplification in galaxy formation (Rieder & Teyssier 2016, 2017), a field where adaptive resolution (either with Lagrangian or AMR-type codes)




Figure 11. Evolution of magnetic fields in the high-resolution (3M elements) MFM stratified shearing box simulation. The saturated α_M is ~ 0.4 with $\langle\beta\rangle \sim 100$. The Maxwell stress is roughly four times the Reynolds stress as found in Hawley et al. (1996). All stresses are normalized by P_0 . The butterfly diagram becomes irregular at ~ 50 orbits.

is essentially required. Likewise, the methods described in this paper have considerable potential for applications in other areas of astrophysics where self-gravitating magnetized disks should be relevant, such as the central regions of massive protogalaxies where self-gravitating circumnuclear gas disks could trigger the formation of supermassive black holes (Regan & Haehnelt 2009; Mayer et al. 2010, 2015; Choi et al. 2013) or the outer regions of accretion disks around AGNs (Rafikov 2001). In the case of the protogalactic nuclei, in particular, adaptivity is necessary to capture a wide range of spatial and temporal scales. Furthermore, understanding the interplay between the stabilizing effect of magnetic pressure, turbulence, and gas inflows governed by global self-gravitating modes might be the key to understand whether a monolithic central collapse into a supermassive star (which later will turn into a massive black hole) occurs, as opposed to fragmentation into stars (Latif et al. 2014). The Riols & Latter dynamo action might play an important role in this latter case, in that it might reveal itself to be an important element to understand the process of angular momentum transport—and thus to better evaluate the possibility of a central monolithic collapse.

5.2. Other Lagrangian MHD methods

We have restricted our study to just two classes of numerical methods: SPH and MFM (although we did consider a few “variants” of SPH). Furthermore, we have only considered

TSPH and PSPH variants of the SPH method. We should point out that caution is warranted in generalizing any of these results to other Lagrangian methods. Moving meshes or mesh-free finite-volume (MFV)-type methods with divergence-cleaning methods can arbitrarily “smooth” the mesh motion, decreasing the “mesh deformation noise” (McNally et al. 2012; Muñoz et al. 2014) and likely allowing for more accurate divergence cleaning simply because the mesh is deforming less rapidly and less irregularly (so, e.g., smaller gradient errors can be ensured). As noted above, unstaggered CT schemes have now been developed (Mocz et al. 2014, 2016) for certain specific types of moving-mesh schemes, which can maintain $\nabla \cdot \mathbf{B} \approx 0$ at machine precision, so they should perform more similarly to CT-grid schemes here, although the numerical noise/dissipation properties of moving-mesh codes (which determine the MRI damping) are often very different.

Fundamentally distinct SPH MHD methods have also been developed. Although early attempts at implementing SPH MHD based on vector potentials did not allow reconnection (e.g., Rosswog & Price 2007), newer hybrid methods that combine vector potentials with divergence-cleaning in the vector potential space appear to avoid exactly the runaway field amplification discussed here (see Stasyszyn & Elstner 2015). To our knowledge, however, these schemes have not yet been explored in a broader context or used for MRI simulations. Finally, in the last stages of the preparation of this paper, we became aware that a simple variant of an SPH MHD solver based on the




Figure 12. Fully developed MRI turbulence. Snapshots of magnetic field strength for a few stratified runs in Table 1 are labeled beside the panels (note that the field strength is shown in Gauss and $\beta = 1$ corresponds to magnetic field strength of 5 Gauss here). The upper panels are the fiducial MFM simulation containing 1.5M elements, along with its equivalent SPH simulations. The MFM snapshot is taken later than the SPH snapshots because the instability develops early in SPH, due to stronger numeric noise (see Figures 7 and 9). The snapshots are taken roughly when α_M reaches its maximum. Compared to SPH simulations, MFM captures finer magnetic field structures and shows less noise in the fields. In the lower panels, the higher-resolution stratified shearing box (with 3M elements) captures even finer structures.

GDSPH method in the GASOLINE2 code (Wadsley et al. 2017) is currently being tested in local MRI setups similar to those described here (R. Wissing et al 2019, private communication). In the latter, the Lorentz force is smoothed in the same way as the hydro force, possibly helping to reduce numerical dissipation.

6. Conclusions

We have presented the results of a series of MRI simulations with two meshless MHD methods, SPH and MFM, in both vertically unstratified and stratified boxes. Two variants of SPH have been considered: a “vanilla” SPH method based on the density-energy formulation, and PSPH (both as implemented in the GIZMO code). The MRI, especially in its zero-net-flux configuration, is sensitive to numerical or physical dissipation, which makes it challenging for any code to simulate, when physical diffusivities are omitted. This is true for both Eulerian and Lagrangian codes, and the results in the zero-net-flux case will always, to some extent, depend on the

nature of their numerical dissipation. It is perhaps then no surprise that the biggest discrepancies between the codes are observed for this magnetic configuration.

Our main findings can be summarized as:

1. The use of an appropriate kernel function that does not exhibit the pairing instability and allows a relatively large radius of compact support (e.g., Wendland C4) is crucial for maintaining element or mesh-generating-point order and for accurate gradient calculation. In MFM, this is directly akin to using a larger stencil to obtain more accurate, higher-order gradient estimators in traditional regular-grid codes. Although these kernels are less compact than the traditional spline kernels and tend to oversmooth fluid variables in SPH, they help to sustain the turbulence longer.
2. A stiff adiabatic EOS can help to control the noise in solving the energy equation, where the truncation errors can be significant, because the magnetic energy is much smaller than the internal energy.

3. In unstratified shearing boxes with a net vertical field, MFM exhibits a similar error scaling in the linear growth MRI rates as compared to the finite volume Eulerian code ATHENA. Both SPH and MFM can adequately simulate the ensuing turbulence, though the former is more diffusive and thus the MRI is closer to criticality. The consequences of higher diffusivity in SPH include more vigorous channel bursts and very severe heating.
4. In unstratified shearing boxes with zero-net vertical field, SPH and MFM exhibit decaying turbulence at the (relatively low) resolution we are able to simulate here. It is possible this decay is linked to a very low numerical magnetic Prandtl number, but it is more likely that the numerical resistivity in MFM is simply too high for the MRI to be sustained at this resolution.
5. In vertically stratified shearing box simulations, SPH MHD produces radically unphysical behavior: turbulence dies out, but strong toroidal fields continue to grow to equipartition with the gas pressure. This is due to nontrivial coupling of poorly controlled magnetic field divergence, differential rotation/shear, and vertical stratification, at least in the most common SPH form of the induction and divergence-cleaning operators.
6. In vertically stratified shearing boxes, high-resolution MFM simulations produce results comparable to grid codes implementing the CT cleaning method. The classical MRI dynamo, with its characteristic butterfly diagram, is captured for several tens of orbits. Nonetheless, the turbulence ultimately dies out after some 50 orbits, at the relatively low resolution studied as our “baseline” here. Going to higher resolution sustains the dynamo for longer, indicating that the decay is likely due to residual numerical resistivity.

We thank Matthew Bate, Daniel Price, Stephen Rosswog, Jim Stone, James Wadsley, Sijing Shen, and Robert Wissing for useful discussions. We acknowledge support from the Swiss National Science Foundation via the NCCR PlanetS. Support for P.F.H. was provided by an Alfred P. Sloan Research Fellowship, NSF Collaborative Research grant #1715847 and CAREER grant #1455342, and NASA grants NNX15AT06G, JPL 1589742, and 17-ATP17-0214.

Software: GIZMO code (Hopkins 2015), VisIt (Childs et al. 2012).

Appendix A Numerical Viscosity

To test the numerical viscosity, we perturb the background shear flow by adding a radial-dependent azimuthal velocity (a “zonal flow”), i.e., $\delta v = U \sin(2\pi n_x x) \hat{y}$. We adjust the internal energy (for the EOS used here, $\gamma = 5/3$) such that the pressure perturbation is $\delta P = -\frac{\rho U}{\pi n_x} \cos(2\pi n_x x)$ and the initial setup is in equilibrium. Numerical viscosity will cause the perturbation to decay. By drawing an analogy to the Navier–Stokes equations (not necessarily true here), U will decay at a rate $\nu_{\text{num}} k_x^2$, where ν_{num} is the effective numerical viscosity and $k_x = 2\pi n_x$. We fit the decay of $\langle (\delta v_y)^2 \rangle$ to determine ν_{num} , and the decay rate of $\langle (\delta v_y)^2 \rangle$ is shown in Figure 13. Using a shearing box of size $H \times H \times H$ resolved by $32 \times 32 \times 32$ elements, we carry out simulations with $U = 0.1$ and $n_x = [1, 2, 3]$. In Figure 13, the numerical viscosity in MFM is larger than TSPH when the resolution is low. However, MFM outperforms TSPH at 32




Figure 13. Velocity perturbation ($U = 0.1$, $n_x = 1$) decays slightly. In this test, the artificial viscosity is almost zero (in the Cullen & Dehnen, switch $\alpha_{\text{sph}} = \alpha_{\text{min}} = 0.05$) and TSPH has a smaller numerical dissipation when the resolution is low.

elements per wavelength, which is in line with the channel-flow growth-rate test in Section 2.5.

Appendix B Numerical Resistivity

Numerical dissipation can destroy magnetic fields. In a periodic box of size $H \times H \times H$ resolved by $32 \times 32 \times 32$ elements, we initial vertical magnetic fields and adjust the internal energy to set the box in pressure equilibrium. We set $\gamma = 5/3$ here, but tests with an isothermal EOS behave similarly. The fields take the form of $\mathbf{B} = B_0 \hat{z} \sin(2\pi n_x x)$, where $B_0 = \sqrt{8\pi P_0 / \beta}$. In MFM simulations, the field structure decays due to numerical resistivity (see Figure 14). Here, B_0 should decay at a rate of $\eta_{\text{num}} k_x^2$. The decay rate can be determined by fitting the decay of the averaged magnetic energy (which decays twice as fast as B_0). We vary the wavelength (n_x) and field strength (β) to test how strong the numerical dissipation is. In this test, the dimensionless divergence of the magnetic fields is ~ 0.0001 and we believe the dissipation is not caused by the non-zero divergence. However, the magnetic energy does not decay in SPH simulations, even with $n_x = 3$ (see Figure 15). The numerical noise from the artificial viscosity and resistivity break the perfect lattice.




Figure 14. Magnetic energy decays exponentially in the test with $\beta = 100$, $n_x = 3$. We fit the curve to an exponential function to get the decay rate. The decay rate increase fast (faster than a parabola) as the resolution decreases, i.e., n_x increases; it is almost independent of the magnetic field strength in our tested range.




Figure 15. Magnetic energy does not decay in the TSPH test with $\beta = 100$, $n_x = 3$. However, the magnetic field structure becomes noise at $50\sqrt{\gamma}H/c_s$.

ORCID iDs

Hongping Deng  <https://orcid.org/0000-0001-6858-1006>
 Philip F. Hopkins  <https://orcid.org/0000-0003-3729-1684>
 Xue-Ning Bai  <https://orcid.org/0000-0001-6906-9549>

References

- Armitage, P. J. 1998, *ApJL*, **501**, L189
- Bai, X.-N. 2014, *ApJ*, **791**, 137
- Bai, X.-N., & Stone, J. M. 2011, *ApJ*, **736**, 144
- Bai, X.-N., & Stone, J. M. 2013, *ApJ*, **767**, 30
- Bauer, A., & Springel, V. 2012, *MNRAS*, **423**, 2558
- Bodo, G., Cattaneo, F., Mignone, A., & Rossi, P. 2014, *ApJL*, **787**, L13
- Bodo, G., Mignone, A., Cattaneo, F., Rossi, P., & Ferrari, A. 2008, *A&A*, **487**, 1
- Brandenburg, A., & Dobler, W. 2002, *CoPhC*, **147**, 471
- Brandenburg, A., Nordlund, A., Stein, R. F., & Torkelsson, U. 1995, *ApJ*, **446**, 741
- Childs, H., Brugger, E., Whitlock, B., et al. 2012, High Performance Visualization—Enabling Extreme-Scale Scientific Insight (Washington, DC: Lawrence Berkeley National Laboratory)
- Choi, J.-H., Shlosman, I., & Begelman, M. C. 2013, *ApJ*, **774**, 149
- Cullen, L., & Dehnen, W. 2010, *MNRAS*, **408**, 669
- Davis, S. W., Stone, J. M., & Pessah, M. E. 2010, *ApJ*, **713**, 52
- Dedner, A., Kemm, F., Kröner, D., et al. 2002, *JCoPh*, **175**, 645
- Dehnen, W., & Aly, H. 2012, *MNRAS*, **425**, 1068
- Deng, H., Mayer, L., & Meru, F. 2017, *ApJ*, **847**, 43
- Deng, H., Reinhardt, C., Benitez, F., Mayer, L., & Stadel, J. 2019, *ApJ*, **870**, 127
- Dobbs, C., Price, D., Pettitt, A., Bate, M. R., & Tricco, T. 2016, *MNRAS*, **461**, 4482
- Dolag, K., & Stasyszyn, F. 2009, *MNRAS*, **398**, 1678
- Durisen, R. H., Boss, A. P., Mayer, L., et al. 2007, in *Protostars and Planets V*, ed. B. Reipurth, D. Jewitt, & K. Keil (Tucson, AZ: Univ. Arizona Press), 607
- Evans, C. R., & Hawley, J. F. 1988, *ApJ*, **332**, 659
- Fleming, T. P., Stone, J. M., & Hawley, J. F. 2000, *ApJ*, **530**, 464
- Forgan, D., Price, D. J., & Bonnell, I. 2016, *MNRAS*, **466**, 3406
- Fromang, S., Latter, H., Lesur, G., & Ogilvie, G. I. 2013, *A&A*, **552**, A71
- Fromang, S., & Nelson, R. P. 2006, *A&A*, **457**, 343
- Fromang, S., & Papaloizou, J. 2007, *A&A*, **476**, 1113
- Fromang, S., Papaloizou, J., Lesur, G., & Heinemann, T. 2007, *A&A*, **1132**, 1123
- Gaburov, E., & Nitadori, K. 2011, *MNRAS*, **414**, 129
- Goldreich, P., & Lynden-Bell, D. 1965, *MNRAS*, **130**, 125
- Goodman, J., & Xu, G. 1994, *ApJ*, **432**, 213
- Gressel, O. 2013, *ApJ*, **770**, 100
- Guerrera, D., Maffia, A., & Burkhardt, H. 2018, *Data in Brief*, **20**, 1148
- Hawley, J. F. 2000, *ApJ*, **528**, 462
- Hawley, J. F., Gammie, C. F., & Balbus, S. a. 1995, *ApJ*, **440**, 742
- Hawley, J. F., Gammie, C. F., & Balbus, S. a. 1996, *ApJ*, **464**, 690
- Hawley, J. F., Guan, X., & Krolik, J. H. 2011, *ApJ*, **738**, 84
- Hietel, D., Steiner, K., & Struckmeier, J. 2000, *Math. Models Methods Appl. Sci.*, **10**, 1363
- Hopkins, P. F. 2012, *MNRAS*, **428**, 2840
- Hopkins, P. F. 2015, *MNRAS*, **450**, 53
- Hopkins, P. F. 2016a, *MNRAS*, **462**, 576
- Hopkins, P. F. 2016b, *MNRAS*, **466**, 3387
- Hopkins, P. F. 2017, arXiv:1712.01294
- Hopkins, P. F., & Raives, M. J. 2015, *MNRAS*, **455**, 51
- Hubber, D. A., Rosotti, G. P., & Booth, R. A. 2017, *MNRAS*, **473**, 1603
- Inutsuka, S.-I. 2002, *JCoPh*, **179**, 238
- Lanson, N., & Vila, J.-P. 2008a, *SJNA*, **46**, 1912
- Lanson, N., & Vila, J.-P. 2008b, *SJNA*, **46**, 1935
- Latif, M., Bovino, S., Van Borm, C., et al. 2014, *MNRAS*, **443**, 1979
- Latter, H. N., Lesaffre, P., & Balbus, S. A. 2009, *MNRAS*, **394**, 715
- Latter, H. N., & Papaloizou, J. 2017, *MNRAS*, **472**, 1432
- Lesaffre, P., Balbus, S. A., & Latter, H. 2009, *MNRAS*, **396**, 779
- Lesur, G., Ferreira, J., & Ogilvie, G. I. 2013, *A&A*, **550**, A61
- Lesur, G., Kunz, M. W., & Fromang, S. 2014, *A&A*, **566**, A56
- Lesur, G., & Longaretti, P. Y. 2007, *MNRAS*, **378**, 1471
- Lesur, G., & Ogilvie, G. 2008, *A&A*, **488**, 451
- Masset, F. 2000, *A&AS*, **141**, 165
- Mayer, L., Fiacconi, D., Bonoli, S., et al. 2015, *ApJ*, **810**, 51
- Mayer, L., & Gawryszczak, A. J. 2008, in *ASP Conf. Ser.* 398, *Extreme Solar Systems*, ed. D. Fischer (San Francisco, CA: ASP), 243
- Mayer, L., Kazantzidis, S., Escala, A., & Callegari, S. 2010, *Natur*, **466**, 1082
- McNally, C. P., Lyra, W., & Passy, J.-C. 2012, *ApJS*, **201**, 18
- Meheut, H., Fromang, S., Lesur, G., Joos, M., & Longaretti, P.-Y. 2015, *A&A*, **579**, A117
- Miller, K. A., & Stone, J. M. 2000, *ApJ*, **534**, 398
- Mocz, P., Pakmor, R., Springel, V., et al. 2016, *MNRAS*, **463**, 477
- Mocz, P., Vogelsberger, M., & Hernquist, L. 2014, *MNRAS*, **442**, 43
- Muñoz, D. J., Kratter, K., Springel, V., & Hernquist, L. 2014, *MNRAS*, **445**, 3475
- Noble, S. C., Krolik, J. H., & Hawley, J. F. 2010, *ApJ*, **711**, 959
- Pakmor, R., & Springel, V. 2013, *MNRAS*, **432**, 176
- Parkin, E., & Bicknell, G. 2013, *ApJ*, **763**, 99
- Pessah, M. E., & Goodman, J. 2009, *ApJL*, **698**, L72
- Powell, K. G., Roe, P. L., Linde, T. J., Gombosi, T. I., & De Zeeuw, D. L. 1999, *JCoPh*, **154**, 284
- Price, D. J. 2012, *JCoPh*, **231**, 759
- Price, D. J., & Bate, M. R. 2007, *MNRAS*, **377**, 77
- Price, D. J., & Bate, M. R. 2008, *MNRAS*, **385**, 1820
- Price, D. J., & Bate, M. R. 2009, *MNRAS*, **398**, 33
- Rafikov, R. R. 2001, *MNRAS*, **323**, 445
- Read, J., Hayfield, T., & Agertz, O. 2010, *MNRAS*, **405**, 1513
- Regan, J. A., & Haehnelt, M. G. 2009, *MNRAS*, **396**, 343
- Rice, W., Paardekooper, S.-J., Forgan, D., & Armitage, P. 2014, *MNRAS*, **438**, 1593
- Rieder, M., & Teyssier, R. 2016, *MNRAS*, **457**, 1722
- Rieder, M., & Teyssier, R. 2017, *MNRAS*, **471**, 2674
- Riols, A., & Latter, H. 2018, *MNRAS*, **474**, 2212
- Rosswog, S. 2015, *MNRAS*, **448**, 3628
- Rosswog, S., & Price, D. 2007, *MNRAS*, **379**, 915
- Ryan, B. R., Gammie, C. F., Fromang, S., & Kestener, P. 2017, *ApJ*, **840**, 6
- Saitoh, T. R., & Makino, J. 2013, *ApJ*, **768**, 44
- Sano, T., Inutsuka, S.-I., Turner, N. J., & Stone, J. M. 2004, *ApJ*, **605**, 321
- Shi, J.-M., Stone, J. M., & Huang, C. X. 2015, *MNRAS*, **456**, 2273
- Simon, J. B., & Hawley, J. F. 2009, *ApJ*, **707**, 833
- 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
- Springel, V. 2005, *MNRAS*, **364**, 1105
- Springel, V. 2010a, *ARA&A*, **48**, 391
- Springel, V. 2010b, *MNRAS*, **401**, 791
- Stasyszyn, F. A., & Elstner, D. 2015, *JCoPh*, **282**, 148
- Steinacker, A., & Papaloizou, J. C. 2002, *ApJ*, **571**, 413
- Steinwandel, U. P., Beck, M. C., Arth, A., et al. 2018, *MNRAS*, **483**, 1008
- Stone, J. M., & Gardiner, T. A. 2010, *ApJS*, **189**, 142
- Stone, J. M., Gardiner, T. A., Teuben, P., Hawley, J. F., & Simon, J. B. 2008, *ApJS*, **178**, 137
- Stone, J. M., Hawley, J. F., Gammie, C. F., & Balbus, S. A. 1996, *ApJ*, **463**, 656
- Teyssier, R. 2002, *A&A*, **385**, 337
- Tricco, T. S., & Price, D. J. 2012, *JCoPh*, **231**, 7214
- Tricco, T. S., & Price, D. J. 2013, *MNRAS*, **436**, 2810
- Turner, N., Sano, T., & Dziourkevitch, N. 2007, *ApJ*, **659**, 729
- Vila, J. 1999, *Math. Models Methods Appl. Sci.*, **9**, 161
- Wadsley, J. W., Keller, B. W., & Quinn, T. R. 2017, *MNRAS*, **471**, 2357
- Zhu, C., Pakmor, R., van Kerkwijk, M. H., & Chang, P. 2015, *ApJL*, **806**, L1
- Zhu, Z., & Stone, J. M. 2018, *ApJ*, **857**, 34