



Fully Kinetic Shearing-box Simulations of Magnetorotational Turbulence in 2D and 3D. I. Pair Plasmas

Fabio Bacchini¹ , Lev Arzamasskiy² , Vladimir Zhdankin^{3,4} , Gregory R. Werner¹ , Mitchell C. Begelman^{5,6} , and Dmitri A. Uzdensky¹

¹ Center for Integrated Plasma Studies, Department of Physics, University of Colorado, 390 UCB, Boulder, CO 80309-0390, USA; fabio.bacchini@colorado.edu

² School of Natural Sciences, Institute for Advanced Study, Princeton, NJ 08544, USA

³ Department of Astrophysical Sciences, Princeton University, 4 Ivy Lane, Princeton, NJ 08544, USA

⁴ Center for Computational Astrophysics, Flatiron Institute, 162 Fifth Avenue, New York, NY 10010, USA

⁵ JILA, University of Colorado and National Institute of Standards and Technology, 440 UCB, Boulder, CO 80309-0440, USA

⁶ Department of Astrophysical and Planetary Sciences, University of Colorado, 391 UCB, Boulder, CO 80309-0391, USA

Received 2022 June 14; revised 2022 August 9; accepted 2022 August 16; published 2022 October 14

Abstract

The magnetorotational instability (MRI) is a fundamental mechanism determining the macroscopic dynamics of astrophysical accretion disks. In collisionless accretion flows around supermassive black holes, MRI-driven plasma turbulence cascading to microscopic (i.e., kinetic) scales can result in enhanced angular-momentum transport and redistribution, nonthermal particle acceleration, and a two-temperature state where electrons and ions are heated unequally. However, this microscopic physics cannot be captured with standard magnetohydrodynamic (MHD) approaches typically employed to study the MRI. In this work, we explore the nonlinear development of MRI turbulence in a pair plasma, employing fully kinetic particle-in-cell (PIC) simulations in two and three dimensions. First, we thoroughly study the axisymmetric MRI with 2D simulations, explaining how and why the 2D geometry produces results that differ substantially from 3D MHD expectations. We then perform the largest (to date) 3D simulations, for which we employ a novel shearing-box approach, demonstrating that 3D PIC models can reproduce the mesoscale (i.e., MHD) MRI dynamics in sufficiently large runs. With our fully kinetic simulations, we are able to describe the nonthermal particle acceleration and angular-momentum transport driven by the collisionless MRI. Since these microscopic processes ultimately lead to the emission of potentially measurable radiation in accreting plasmas, our work is of prime importance to understand current and future observations from first principles, beyond the limitations imposed by fluid (MHD) models. While in this first study we focus on pair plasmas for simplicity, our results represent an essential step toward designing more realistic electron–ion simulations, on which we will focus in future work.

Unified Astronomy Thesaurus concepts: Plasma astrophysics (1261); Accretion (14); Plasma physics (2089)

1. Introduction

Astrophysical accretion disks are ubiquitous in the universe, often surrounding massive compact objects such as black holes and neutron stars. Plasmas in accretion disks rotate around the central object with a differential-rotation profile determined by the object’s gravity; gravity also drives accretion of gas onto the central object, thus requiring a mechanism of angular-momentum transport and redistribution throughout the disk. Such a mechanism has not been completely understood so far. The magnetorotational instability (MRI; Velikhov 1959; Chandrasekhar 1960; Balbus & Hawley 1991, 1998) has been proposed as a possible pathway to promote angular-momentum transport in accretion disks: the MRI causes a dramatic amplification of any pre-existing seed magnetic field, however small, driving and sustaining turbulence at macroscopic scales. The development of the MRI thus creates the conditions to connect the plasma dynamics at the smallest and largest length scales in accretion disks. Although the development of strong turbulence is an inevitable consequence of the MRI, it remains to be understood whether, and under which conditions, this

turbulence can also be responsible for efficient angular-momentum transport.

Several studies have been conducted on the linear MRI theory (Balbus & Hawley 1991; Goodman & Xu 1994; Latter et al. 2009; Pessah & Goodman 2009), but the full nonlinear development of this instability cannot be treated analytically. Decades of research efforts have employed magnetohydrodynamic (MHD) simulations to study the complex, turbulent state that arises during the nonlinear stage of the MRI. In particular, local models of accretion disks—where only a small sector of the disk is simulated—have attracted widespread recognition, owing to the possibility of studying the MRI development separately from its embedding astrophysical environment. From the very beginning, such simulations have been conducted with the so-called *shearing-box* approach (Hill 1878; Hawley et al. 1995), where a small simulation box with locally Cartesian coordinates is taken as representative of the large-scale behavior of the whole disk. In the shearing-box paradigm without density stratification, the vertical (z , parallel to the rotation axis) and toroidal (φ , the direction of rotation; locally y) coordinates are periodic, and the radial (r , locally x) coordinate employs *shearing-periodic* conditions that model the background differential rotation of the disk. Several MHD codes implementing the shearing box have been developed (e.g., Gressel & Ziegler 2007; Stone & Gardiner 2010) and applied (e.g., Matsumoto & Tajima 1995; Stone et al. 1996;



Original content from this work may be used under the terms of the [Creative Commons Attribution 4.0 licence](https://creativecommons.org/licenses/by/4.0/). Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.

Sano & Inutsuka 2001; Pastorello et al. 2013; Hirai et al. 2018) to study the MRI. These works have found that, when the initial conditions involve a weak, purely vertical magnetic field, the system initially evolves to a state of so-called “channel flows” where strong radial and toroidal magnetic fields develop. The magnetic-field polarity and flow bulk velocity change sign at the channel interfaces, which become susceptible to the development of parasitic instabilities (e.g., tearing, Kelvin–Helmholtz, or drift-kink modes among others; Goodman & Xu 1994; Pessah & Goodman 2009) and magnetic reconnection. These secondary modes feed off the primary instability, eventually destroying the macroscopic channels and driving a turbulent state.

While MHD studies have provided fundamental insight into the physics of the MRI, plasmas in astrophysical accretion disks should in many cases be treated with kinetic approaches. In particular for radiatively inefficient accretion flows (RIAFs) around supermassive black holes (SMBHs) accreting at low rates, it is expected that plasma particles (electrons and ions) can travel long distances between binary collisions, i.e., their mean free path is comparable to the macroscopic scales in these environments. As a consequence, electrons and ions in RIAFs can be thermally decoupled at short distances from the central object; both species can experience strong nonthermal acceleration, but energy partition can be unequal. In particular, electron acceleration to high energies can result in the emission of most of the outgoing observable disk radiation. Such a collisionless, two-temperature, nonthermal, radiative state cannot be fully captured with the standard MHD simulations typically employed for accretion-disk modeling; it is, however, imperative to investigate plasmas in these conditions with appropriate methods to correctly interpret current and future observations, such as the first direct images of the accretion flow around M87* and SgrA* (EHT Collaboration 2019, 2022). So far, no self-consistent (i.e., first-principles) model of collisionless MRI-driven turbulence, including reconnection, angular-momentum transport, plasma heating, and particle acceleration has been produced.

The scarcity of accurate theoretical models has profound impact on our understanding of accretion physics. For example, recent observational campaigns are affected by heavy uncertainties due to a lack of a reliable theory for the electron/ion energy partition in accretion disks. One of the most widely employed models of plasma heating in general-relativistic magnetohydrodynamic (GRMHD) simulations used to interpret Event Horizon Telescope (EHT) images (e.g., Ressler et al. 2015; Chael et al. 2018; EHT Collaboration 2019) is based on the nonrelativistic gyrokinetic (GK) approximation (e.g., Schekochihin et al. 2009), in which the magnetic moment of particles is assumed to be conserved. This assumption reduces the dimensionality of the problem and makes it much less computationally challenging, allowing accurate calculations of energy partition between different species in turbulence (Howes 2010; Kawazura et al. 2019), including MRI turbulence (e.g., Kawazura et al. 2021, in the reduced-MHD limit). However, the GK assumptions prevent one from studying nonadiabatic heating mechanisms (in which the magnetic moment is not conserved). Additionally, GK heating models do not take into account deviations of the particle distribution function from thermodynamic equilibrium (e.g., pressure anisotropy). Such deviations can significantly modify the rate of Landau damping (Kunz et al. 2015, 2018), which is

the main dissipation mechanism in GK. In addition, in realistic astrophysical plasmas, the magnetic moment of particles is not expected to be conserved on global dynamical timescales, allowing for a number of heating mechanisms not included in GK. For example, if high-frequency fluctuations are present, turbulence can be dissipated due to cyclotron resonance (Kennel & Engelmann 1966; Isenberg 2004; Isenberg & Vasquez 2007; Arzamasskiy et al. 2019; Squire et al. 2022); alternatively, if electric-field fluctuations at scales comparable to particle gyroradii have sufficiently large amplitudes, particles can start experiencing a random walk in velocity space and extract energy from these fluctuations (“stochastic heating”; McChesney et al. 1987; Chandran et al. 2010; Hoppock et al. 2018; Cerri et al. 2021). It is unclear whether these two mechanisms are viable in the context of accretion flows, because both are expected to become less and less important at large scale separation (i.e., the ratio between global scales and the scales of particle gyromotion). Additionally, if plasmas significantly deviate from local thermodynamic equilibrium, various kinetic microinstabilities can be excited, such as firehose (Rosenbluth 1956; Chandrasekhar et al. 1958; Parker 1958; Vedenov & Sagdeev 1958; Kunz et al. 2014a; Riquelme et al. 2015), mirror (Shapiro & Shevchenko 1964; Kunz et al. 2014a; Riquelme et al. 2015), ion-cyclotron (Sagdeev & Shafranov 1960; Ley et al. 2019), and whistler (Sagdeev & Shafranov 1960; Riquelme et al. 2016). These instabilities create small-scale fluctuations that scatter particles and introduce an effective viscosity in the plasma, which becomes the dominant dissipation mechanism of the cascade (Arzamasskiy et al. 2022). The partition of energy due to this “gyroviscosity” has been studied with “kinetic-MHD” methods (similar to the Braginskii-MHD approach) by Sharma et al. (2006), who concluded that the ion-to-electron heating ratio depends on which instabilities are excited and how pressure anisotropy is mediated by them, making the fully kinetic modeling of these processes crucial for understanding plasma heating.

Importantly, plasma energization can also occur in the presence of current sheets, due to the tearing instability (leading to magnetic reconnection; see, e.g., Zweibel & Yamada 2009) and the drift-kink instability (DKI, Pritchett et al. 1996; Zenitani & Hoshino 2007; Hoshino 2020). Both these instabilities can operate in turbulent plasmas, where myriad small-scale current sheets can form (e.g., Zhdankin et al. 2013, 2015, 2018a; Mallet et al. 2017; Loureiro & Boldyrev 2020). These mechanisms can rapidly convert built-up magnetic energy to plasma energy, accelerating a significant fraction of particles to very high (nonthermal) energies (as shown by kinetic simulations, e.g., Daly et al. 2008; Guo et al. 2014; Sironi & Spitkovsky 2014; Werner et al. 2016; Zhdankin et al. 2017, 2019; Comisso & Sironi 2018, 2019; Werner et al. 2018; Wong et al. 2020). In particular, the tearing instability tends to disrupt thin current sheets; concurrently, the DKI can cause a current sheet to kink (ripple) and become distorted, dissipating magnetic energy as the sheet turbulently folds over on itself (Werner & Uzdensky 2021). These processes play an important role in dissipating magnetic fields generated by the MRI, and require a kinetic treatment to be captured self-consistently. Recently, prescriptions for electron/ion heating in magnetic reconnection obtained from kinetic simulations (e.g., Ball et al. 2018; Werner et al. 2018) have been employed in GRMHD as a way to include the missing microphysics leading

to plasma heating and subsequent emission (Dexter et al. 2021; Hankla et al. 2022; Scepi et al. 2022).

Studying the MRI with kinetic approaches is important beyond the electron/ion thermal- and nonthermal-energization problem. A primary question is how the aforementioned angular-momentum transport is realized in collisionless disks; turbulence could, in principle, supply an effective “collisionless” viscosity and resistivity determined by electromagnetic, bulk, and anisotropic-pressure stresses. However, the latter is not included in standard MHD models, which typically assume an isotropic pressure. As mentioned above, building on kinetic theory (Quataert et al. 2002), Sharma et al. (2006) conducted simulations with a more sophisticated “kinetic” (i.e., Braginskii) MHD model that evolves the parallel and perpendicular pressure independently, limiting the allowed pressure anisotropy with an ad hoc prescription. These simulations have shown that pressure anisotropy can enhance the effective viscosity, but a first-principles explanation for the observed dynamics is still lacking.

Turbulent angular-momentum transport and particle acceleration in collisionless plasmas can be self-consistently modeled from first principles with particle-in-cell (PIC) simulations. Riquelme et al. (2012) and Hoshino (2013) presented the first attempts to conduct shearing-box PIC simulations in 2D. These pioneering works focused on pair plasmas (for simplicity), and found that reconnection and turbulence during the nonlinear MRI stage can produce substantial nonthermal particle acceleration and momentum transport. Hoshino (2015) carried out the first-ever 3D shearing-box PIC simulation for a pair plasma and found similar results during the nonlinear MRI stage. These findings were corroborated by additional large-scale 2D pair-plasma simulations by Inchingolo et al. (2018), who pointed out the role of drift-kink instabilities in driving the turbulent state. All these results align well with 3D simulations conducted by Kunz et al. (2016) with a nonrelativistic hybrid approach (although these employed a different initial magnetic-field geometry), where ions are fully kinetic and electrons are treated as a charge-neutralizing massless fluid.

Although they have provided new, fundamental insight, the aforementioned PIC studies of the MRI so far presented in the literature have been rather inconclusive in determining the exact dependence of simulations on the chosen physical and numerical parameters. The main issue impeding a comprehensive understanding of the collisionless MRI is the presence of very different scales of length and time that have to be included in PIC simulations for these to achieve the converged, mesoscale dynamics that is expected from MHD numerical experiments. In nonrelativistic astrophysical plasmas, the kinetic (microscopic) length scales for each particle species j subjected to a magnetic field $\mathbf{B}$ are the Debye length $\lambda_{D,j} = \sqrt{kT_j/(4\pi n_j q_j^2)}$, the skin depth $c/\omega_{p,j}$ (with the plasma frequency $\omega_{p,j} = \sqrt{4\pi n_j q_j^2/m_j}$), and the Larmor radius $\rho_{C,j} = m_j c u_{\perp,j}/(q_j B)$ (where $u_{\perp}$ represents the particle velocity perpendicular to $\mathbf{B}$); here, T and n are the plasma temperature and number density, and q and m are the particle charge and mass. For an accretion disk subjected to the MRI, the macroscopic length scales at a specific distance r_0 from the central object are the most-unstable MRI wavelength $\lambda_{\text{MRI}} = 2\pi v_A/\Omega_0$ and the disk scale-height $H = v_{\text{th}}/\Omega_0$, where $v_A = B/\sqrt{4\pi \sum_j n_j m_j}$ is the (nonrelativistic) Alfvén speed, v_{th} is




Figure 1. Representation of the hierarchy from microscopic (kinetic) to macroscopic (fluid) length scales that are relevant for kinetic shearing-box simulations of the MRI.

the plasma thermal speed, and Ω_0 is the disk rotational frequency at r_0 . The relative ordering of microscopic (kinetic) and macroscopic (fluid) spatial scales for cold plasmas is depicted in Figure 1. The separation between macroscopic and microscopic scales in real astrophysical systems is typically enormous, and including a realistic separation in PIC simulations translates into prohibitive computational costs. To face the problem, PIC methods typically employ a reduced (i.e., unrealistic) separation between fluid and kinetic scales. Unrealistic physical parameters can, however, have a substantial effect on simulation results. Studying how a certain choice of parameters translates into a specific simulated result is a key element in interpreting many kinetic studies presented in the literature, but it has so far been underexplored. In addition, the reduced dimensionality of 2D simulations can dramatically alter the MRI evolution (Sano & Stone 2002; Masada & Sano 2008), adding to the uncertainty of the PIC results presented in literature.

In this work, we aim at filling the existing knowledge gap in the physics of the collisionless MRI by carrying out a thorough exploration of PIC simulations in two and three dimensions. We assess the effect of all relevant physical parameters on kinetic simulations, and most importantly, we conduct large-scale, fully kinetic 3D simulations of unprecedented size. With these runs, we can analyze the phenomenology of MRI-driven turbulence, angular-momentum transport, and particle acceleration self-consistently, exploring a wide parameter space. In this first study, we focus on pair plasmas for simplicity, with the aim of approaching the more relevant electron–ion case in a follow-up work.

This paper is organized as follows: in Section 2, we present the relevant equations for shearing-box PIC simulations and the methods employed to carry out our numerical study. In Section 3, we discuss the choice of simulation parameters and the implications for the numerical cost of the simulations we run. In Section 4, we present a large parameter-space exploration of the 2D axisymmetric pair-plasma MRI, discussing its phenomenology, the numerical convergence of simulations, and the mechanisms of angular-momentum transport and particle acceleration. In Section 5, we present the first large-scale, 3D, fully kinetic simulations of the MRI conducted with a newly developed “orbital-advection” shearing-box method. Finally, in Section 6, we discuss our results and present our main conclusions.

2. Kinetic Shearing-box Equations

In this section, we present the theoretical models employed for the numerical simulations presented in Section 4 and onward. First, we briefly describe the particle and field equations in a rotating frame (Section 2.1) and the transformation to shearing coordinates (Section 2.2). Then, we present a new formulation of the governing equations that simplifies the existing approaches and is appropriate for 3D simulations

(Section 2.3). We will call this new paradigm *kinetic orbital advection*.

2.1. Equations in a Rotating Frame

We consider a small Cartesian sector of an accretion disk with coordinates (r, φ, z) , i.e., the radial, toroidal, and vertical (parallel to the central object's axis of rotation) directions. We apply a boost to a frame of reference that moves with rotational velocity $\mathbf{v}_0 = \boldsymbol{\Omega}_0 \times \mathbf{r}$, where $\boldsymbol{\Omega}_0 = (0, 0, \Omega_0)$ is the rotational frequency. The local box assumes a coordinate system (x, y, z) , where $r_0 + x$ gives the radial position by means of a (small) increment x with respect to a reference radius r_0 ; y is the azimuthal (toroidal) direction, ignoring effects of curvature; and z is still the vertical direction. The rotational velocity is assumed to be nonrelativistic, i.e., $v_0^2 \ll c^2$. The disk is assumed to be differentially rotating, implying that an observer in the corotating frame, at some radius r_0 with velocity $\mathbf{v}_0(r_0)$, measures a background shear in the toroidal direction. Across short distances, the corresponding background velocity profile can be linearized as $\mathbf{v}_s(x) = -s\Omega_0 x \hat{\mathbf{e}}_y$, with $s = 3/2$ for a Keplerian disk (note that the equations we write below remain applicable for arbitrary s).

Maxwell's equations for the electromagnetic fields in the corotating frame become (e.g., Schiff 1939; Hoshino 2013⁷)

$$\partial_t \mathbf{B} = -c \nabla \times \mathbf{E}, \quad (1)$$

$$\partial_t \left(\mathbf{E} - \frac{\mathbf{v}_0}{c} \times \mathbf{B} \right) = c \nabla \times \mathbf{B} - 4\pi \mathbf{J} - \nabla \times \left(\mathbf{v}_0 \times \left(\mathbf{E} - \frac{\mathbf{v}_0}{c} \times \mathbf{B} \right) \right), \quad (2)$$

$$\nabla \cdot \mathbf{B} = 0, \quad (3)$$

$$\nabla \cdot \left(\mathbf{E} - \frac{\mathbf{v}_0}{c} \times \mathbf{B} \right) = 4\pi \rho, \quad (4)$$

for the electric and magnetic fields $\mathbf{E}$ and $\mathbf{B}$, evolving in time according to the source terms $\mathbf{J}$ and ρ , i.e., current and charge density. We note that, in this frame, there exist four vector fields, $\mathbf{E}$, $\mathbf{E}^*$, $\mathbf{B}$, and $\mathbf{B}^*$, obeying the relations (Arendt 1998)

$$\mathbf{E} = \mathbf{E}^* + \frac{\mathbf{v}_0}{c} \times \mathbf{B}, \quad (5)$$

$$\mathbf{B}^* = \mathbf{B} - \frac{\mathbf{v}_0}{c} \times \mathbf{E}^*, \quad (6)$$

and one could write different (but equally valid) versions of Maxwell's equations in terms of any set of such fields. Equations (1)–(4) employ $\mathbf{E}$ and $\mathbf{B}$, as this choice makes the resulting equations more suitable for numerical integration.

The relativistic particle equations of motion in the rotating frame are

$$\frac{d\mathbf{x}}{dt} = \frac{\mathbf{u}}{\gamma}, \quad (7)$$

$$\frac{d\mathbf{u}}{dt} = \frac{q}{m} \left(\mathbf{E} + \frac{\mathbf{u}}{c\gamma} \times \mathbf{B} \right) + 2\mathbf{u} \times \boldsymbol{\Omega}_0 + 2\gamma s \Omega_0^2 x \hat{\mathbf{e}}_x, \quad (8)$$

for a particle with position $\mathbf{x}$ and spatial part of the four-velocity $\mathbf{u} = (u_x, u_y, u_z)$ (with the Lorentz factor

$\gamma = \sqrt{1 + u^2/c^2}$). On the right-hand side of Equation (8), the Coriolis, centrifugal, and radial gravitational forces (simplified to the small-box approximation) add to the usual electromagnetic Lorentz force acting on particles. Equations (1)–(8) will be referred to as the *kinetic shearing-box* (KSB) system from now on.

Because of the background shearing profile, peculiar boundary conditions are necessary in shearing-box simulations. Maxwell's Equations (1)–(4) require a special treatment at the x -boundaries, reflecting the fact that the shear advects and distorts field lines and fluid bulk motions. These *shearing-periodic boundary conditions* are

$$\mathbf{E}(x_i, y, z) = \mathbf{E}(x_o, y - \Delta y_s, z) + \mathbf{v}_s(L_x) \times \mathbf{B}(x_o, y - \Delta y_s, z)/c, \quad (9)$$

$$\mathbf{B}(x_i, y, z) = \mathbf{B}(x_o, y - \Delta y_s, z), \quad (10)$$

$$\mathbf{J}(x_i, y, z) = \mathbf{J}(x_o, y - \Delta y_s, z) + \mathbf{v}_s(L_x) \rho(x_o, y - \Delta y_s, z), \quad (11)$$

$$\rho(x_i, y, z) = \rho(x_o, y - \Delta y_s, z), \quad (12)$$

where x_i and x_o represent the position of the inner and outer x -boundaries and $L_x = x_o - x_i$ is the box size along x . The difference in y -coordinate $\Delta y_s = s\Omega_0 L_x t$ increases with time, and at the *periodic times* $t = N L_y / (s\Omega_0 L_x)$ with $N \in \mathbb{N}$, the simulation box is strictly periodic again along x . The y - and z -directions always remain strictly periodic.

Finally, note that, in particle simulations, special boundary conditions must be applied to particles crossing x -boundaries. First, the velocity of each particle must be modified by adding or subtracting (upon crossing the outer or inner boundary, respectively) the global velocity offset $s\Omega_0 L_x$. Second, the position of each crossing particle must be shifted by a distance Δy_s along the positive or negative y -direction (for particles crossing the outer or inner boundary, respectively).

The numerical implementation of the KSB system presents significant complications, due to the shearing-periodic boundary conditions. These difficulties mainly arise in computer codes employing MPI parallelization: because corresponding points along the inner and outer x -boundaries shift in position as time passes, the communication between processors across these boundaries is substantially more complex than in the case of standard periodic boundary conditions. Additionally, the values of $\mathbf{E}$ and $\mathbf{J}$ must be corrected (with $\mathbf{B}$ and ρ , respectively) prior to communication, which requires additional interpolation because these quantities are typically not defined at the same spatial locations.

2.2. Equations in Shearing Coordinates

A simpler alternative approach for PIC simulations has been proposed by Riquelme et al. (2012) in the form of a transformation to the so-called *shearing coordinates* (originally developed and employed for MHD calculations; Goldreich & Lynden-Bell 1965; Goldreich & Tremaine 1978; Narayan et al. 1987; Kaisig 1989). The transformation consists of a boost along y to a reference frame moving with velocity $\mathbf{v}_s(x)$, thus obtaining the cancellation of velocity offsets at the x -boundaries.⁸ Assuming a 2D (x – z) geometry, Maxwell's

⁷ Note that there is a sign typo in Hoshino (2013) in the first $\mathbf{v}_0$ term on the right-hand side of Equation (2).

⁸ Note that all terms containing $\mathbf{v}_0$ in the KSB system are neglected prior to boosting the reference frame, under the assumption that $v_0^2 \ll c^2$.

equations can then be rewritten and simplified to obtain

$$\partial_t \mathbf{B} = -c \nabla \times \mathbf{E} - s \Omega_0 B_x \hat{\mathbf{e}}_y, \quad (13)$$

$$\partial_t \mathbf{E} = c \nabla \times \mathbf{B} - s \Omega_0 E_x \hat{\mathbf{e}}_y - 4\pi \mathbf{J}, \quad (14)$$

$$\nabla \cdot \mathbf{B} = 0, \quad (15)$$

$$\nabla \cdot \mathbf{E} = 4\pi \rho, \quad (16)$$

which only differ from the standard Maxwell's equations by two source terms that drive the shearing of electric- and magnetic-field lines along y (see Riquelme et al. 2012 for further details). Correspondingly, the particle equations of motion are rewritten by subtracting the background shearing velocity, such that

$$\frac{d\mathbf{x}}{dt} = \frac{\mathbf{u}}{\gamma}, \quad (17)$$

$$\frac{d\mathbf{u}}{dt} = \frac{q}{m} \left(\mathbf{E} + \frac{\mathbf{u}}{c\gamma} \times \mathbf{B} \right) + 2\mathbf{u} \times \Omega_0 + s \Omega_0 u_x \hat{\mathbf{e}}_y. \quad (18)$$

We will refer to Equations (13)–(18) as the *kinetic shearing box in shearing coordinates* (KSB-SC) system. Note that, because of the boost, all quantities (electromagnetic fields, currents, velocities, etc.) measured in this reference frame are not equivalent to those measured in the standard corotating frame employed for the KSB system. In particular, particle velocities here do not include the background shearing-velocity part, implying that the current density measured in this frame lacks the corresponding contribution from the background shearing motion.

The KSB-SC equations above are adequate for 2D axisymmetric simulations such as those conducted by Riquelme et al. (2012) and Inchingolo et al. (2018) (as well as in Section 4 of this work). These equations are advantageous for such simulations also because they are far less complicated to solve numerically than those of the KSB system, since the x -boundaries are here strictly periodic. However, the full 3D version of the KSB-SC equations includes several additional terms that explicitly depend on the y -coordinate (Riquelme et al. 2012); the numerical implementation of these terms is substantially more complicated than that of the original KSB equations. For this reason, the KSB-SC system has so far only been employed for 2D simulations.

2.3. A New Approach: Kinetic Orbital Advection

Here, we propose a new approach to kinetic shearing-box simulations in 3D. This is based on a combination of existing techniques that have been employed in previous MHD and PIC studies of the MRI. Our method is similar to the implementation of the KSB system as presented by Hoshino (2013, 2015), but simplifies the required boundary conditions and employs equations similar to those of the KSB-SC system, while remaining fully applicable in 3D. The construction of our method follows the so-called “orbital advection” approach that is widely applied in MHD shearing-box implementations (e.g., Gressel & Ziegler 2007; Stone & Gardiner 2010 and references therein). As we will show in the following, the basic idea is that electric fields and particle momenta are replaced with corresponding quantities boosted to a frame moving with $\mathbf{v}_s$, but the coordinates of the simulation box are kept in the lab frame. We note that a similar strategy was employed in

nonrelativistic hybrid simulations by Kunz et al. (2014a, 2014b, 2016).

To construct our set of equations, we start from the KSB system. First, as done in all previous works, we neglect the terms involving $\mathbf{v}_0$ in Maxwell's equations (Riquelme et al. 2012; Hoshino 2013, 2015; Inchingolo et al. 2018). Then, assuming $v_s^2 \ll c^2$, we apply a Galilean transformation,

$$\mathbf{E}' \simeq \mathbf{E} + \frac{\mathbf{v}_s}{c} \times \mathbf{B}, \quad (19)$$

where $\mathbf{E}'$ is the electric field in the comoving frame. Substituting for $\mathbf{E}'$ in Maxwell's equations, we get

$$\partial_t \mathbf{B} = -c \nabla \times \left(\mathbf{E}' - \frac{\mathbf{v}_s}{c} \times \mathbf{B} \right), \quad (20)$$

$$\partial_t \left(\mathbf{E}' - \frac{\mathbf{v}_s}{c} \times \mathbf{B} \right) = c \nabla \times \mathbf{B} - 4\pi \mathbf{J}, \quad (21)$$

$$\nabla \cdot \mathbf{B} = 0, \quad (22)$$

$$\nabla \cdot \left(\mathbf{E}' - \frac{\mathbf{v}_s}{c} \times \mathbf{B} \right) = 4\pi \rho. \quad (23)$$

Further, substituting Equation (20) into (21), we obtain

$$\begin{aligned} \partial_t \mathbf{E}' &= c \nabla \times \mathbf{B} - 4\pi \mathbf{J} \\ &\quad - \mathbf{v}_s \times (\nabla \times \mathbf{E}') + \frac{\mathbf{v}_s}{c} \times \left(\nabla \times \left(\frac{\mathbf{v}_s}{c} \times \mathbf{B} \right) \right), \end{aligned} \quad (24)$$

where, for consistency with the assumption $v_s^2 \ll c^2$, we will ignore the last term on the right-hand side. Note that, having removed the bulk-motion electric field $\mathbf{v}_s \times \mathbf{B}/c$ from $\mathbf{E}$, one would expect $E' \ll E$; however, E' can still be large due to nonideal effects—e.g., at reconnection sites—hence we retain the term $\mathbf{v}_s \times (\nabla \times \mathbf{E}')$ in Equation (24). A clear advantage of a formulation involving $\mathbf{E}'$ as the electric field is that the boundary condition along radial boundaries becomes

$$\mathbf{E}'(x_i, y, z) = \mathbf{E}'(x_0, y - \Delta y_s, z), \quad (25)$$

which, in contrast with Equation (9) for the KSB system, does not require correction of the electric field at the radial boundaries with the corresponding magnetic field. This simplifies the numerical implementation of boundary conditions significantly.

Before rewriting Maxwell's equations into a form that is suitable for a numerical treatment, it is convenient (for reasons that will become apparent later) to substitute the current, $\mathbf{J}$, with the comoving current $\mathbf{J}'$. By applying another Galilean transformation, we get

$$\mathbf{J}' \simeq \mathbf{J} - \rho \mathbf{v}_s. \quad (26)$$

Note that using $\mathbf{J}'$ as a source term for Maxwell's equations modifies the boundary condition (11), which becomes

$$\mathbf{J}'(x_i, y, z) = \mathbf{J}'(x_0, y - \Delta y_s, z) \quad (27)$$

and is manifestly simpler to implement numerically, since it requires no correction of the current at the boundaries. We can now employ Equations (22) and (23) to obtain the following

evolution equations for $\mathbf{E}'$ and $\mathbf{B}$,

$$\begin{cases} \partial_t B_x = -c\partial_y E'_z + c\partial_z E'_y - v_{s,y}\partial_y B_x \\ \partial_t B_y = c\partial_x E'_z - c\partial_z E'_x - v_{s,y}\partial_y B_y - s\Omega_0 B_x, \\ \partial_t B_z = -c\partial_x E'_y + c\partial_y E'_x - v_{s,y}\partial_y B_z \end{cases} \quad (28)$$

$$\begin{cases} \partial_t E'_x = c\partial_y B_z - c\partial_z B_y + v_{s,y}\partial_y E'_x - v_{s,y}\partial_x E'_y - 4\pi J'_x \\ \partial_t E'_y = -c\partial_x B_z + c\partial_z B_x \\ \quad - v_{s,y}\partial_x E'_x - v_{s,y}\partial_y E'_y - v_{s,y}\partial_z E'_z - 4\pi J'_y \\ \partial_t E'_z = c\partial_x B_y - c\partial_y B_x + v_{s,y}\partial_y E'_z - v_{s,y}\partial_z E'_y - 4\pi J'_z \end{cases} \quad (29)$$

where we have neglected all terms of order Ω_0^2 . This form of Maxwell's equations is simple to handle numerically, since it can be written as

$$\partial_t \mathbf{B} = -v_{s,y}\partial_y \mathbf{B} + (\text{other terms}), \quad (30)$$

$$\partial_t \mathbf{E}' = \pm v_{s,y}\partial_y \mathbf{E}' + (\text{other terms}), \quad (31)$$

where the “ $\pm$ ” in the equation for $\mathbf{E}'$ is a “ $-$ ” for the y -component and a “ $+$ ” for the other components. These equations present a manifestly advective term in the y -direction, which can then be treated with any of the available schemes for numerical advection (see Appendix A for further details).

In our approach, we choose to keep the particle position in the lab frame, but to solve for the particle momenta $\mathbf{u}'$ in the comoving frame. In the nonrelativistic limit, this simply reduces to separating out the shearing velocity $\mathbf{v}_s$ from velocity fluctuations superimposed to this background. This choice serves three purposes: first, we avoid the representation of a constant background by a finite amount of computational particles, greatly reducing computational noise (as also noted by Kunz et al. 2014b); second, by evolving the particle velocity in the comoving frame, we can directly collect the current $\mathbf{J}'$, justifying the use of the comoving current as a source term in Maxwell's equations; third, we avoid the need to modify the y -velocity of particles crossing x -boundaries, thus simplifying the application of boundary conditions.

To derive our particle equations of motion, we start from Equation (8); we then boost $\mathbf{u}$ to the frame moving with $\mathbf{v}_s$,

$$\begin{cases} u'_x = u_x \\ u'_y = \gamma_s u_y - \gamma_s v_{s,y} \gamma, \\ u'_z = u_z \end{cases} \quad (32)$$

where $\gamma_s = 1/\sqrt{1 - v_s^2/c^2}$. If we again assume $v_s^2 \ll c^2$, this essentially becomes a Galilean transformation, and we have $\gamma_s \simeq 1$ and $u'_y = u_y - \gamma' v_{s,y}$, with $\gamma' \simeq \gamma$. We then substitute into the equation of motion (8) to obtain

$$\begin{aligned} \frac{d\mathbf{u}}{dt} &\simeq \frac{d\mathbf{u}'}{dt} + \gamma' \frac{d\mathbf{v}_s}{dt} + \mathbf{v}_s \frac{d\gamma'}{dt} \\ &= \frac{q}{m} \left(\mathbf{E} + \frac{\mathbf{u}'}{c\gamma'} \times \mathbf{B} + \frac{\mathbf{v}_s}{c} \times \mathbf{B} \right) \\ &\quad + 2(\mathbf{u}' \times \boldsymbol{\Omega}_0 + \gamma' \mathbf{v}_s \times \boldsymbol{\Omega}_0 + \gamma' s\Omega_0^2 x \hat{\mathbf{e}}_x), \end{aligned} \quad (33)$$

and for the particle position we have

$$\frac{d\mathbf{x}}{dt} = \frac{\mathbf{u}'}{\gamma'} + \mathbf{v}_s. \quad (34)$$

The terms containing γ' on the left-hand side of Equation (33) can be expanded: the first term is

$$\gamma' \frac{d\mathbf{v}_s}{dt} = -\gamma' s\Omega_0 \frac{d\mathbf{x}}{dt} \hat{\mathbf{e}}_y = -s\Omega_0 u'_x \hat{\mathbf{e}}_y, \quad (35)$$

and the second term can be treated in the following way: if $\gamma' \simeq \gamma$,

$$\frac{d\gamma'}{dt} \simeq \frac{d\gamma}{dt} = \frac{\mathbf{u}}{\gamma} \cdot \frac{d\mathbf{u}}{dt} \simeq (\mathbf{v}' + \mathbf{v}_s) \cdot \frac{d\mathbf{u}}{dt}, \quad (36)$$

where $\mathbf{v}' = \mathbf{u}'/\gamma'$. Now using Equation (8) together with $\mathbf{E}' = \mathbf{E} + (\mathbf{v}_s/c) \times \mathbf{B}$, we get

$$\begin{aligned} \frac{d\gamma'}{dt} &\simeq (\mathbf{v}' + \mathbf{v}_s) \cdot \frac{q}{m} \left(\mathbf{E}' + \frac{\mathbf{v}'}{c} \times \mathbf{B} \right) \\ &\quad + (\mathbf{v}' + \mathbf{v}_s) \cdot (2\mathbf{u} \times \boldsymbol{\Omega}_0 + 2\gamma s\Omega_0^2 x \hat{\mathbf{e}}_x). \end{aligned} \quad (37)$$

Multiplying by $\mathbf{v}_s$ and neglecting all terms of order Ω_0^2 (for consistency with the assumption $\gamma_s \simeq 1$), we thus have

$$\mathbf{v}_s \frac{d\gamma'}{dt} \simeq \frac{q}{m} (\mathbf{v}' \cdot \mathbf{E}') \mathbf{v}_s. \quad (38)$$

Finally, on the right-hand side of Equation (33), we can expand the term

$$\gamma' \mathbf{v}_s \times \boldsymbol{\Omega}_0 = -\gamma' s\Omega_0^2 x \hat{\mathbf{e}}_x, \quad (39)$$

which cancels exactly with the rotational-gravitational term in Equation (33). The final form of our momentum equation is therefore

$$\begin{aligned} \frac{d\mathbf{u}'}{dt} &= \frac{q}{m} \left(\mathbf{E}' + \frac{\mathbf{u}'}{c\gamma'} \times \mathbf{B} \right) \\ &\quad + 2\mathbf{u}' \times \boldsymbol{\Omega}_0 + s\Omega_0 u'_x \hat{\mathbf{e}}_y - \frac{q}{m} \left(\frac{\mathbf{u}'}{\gamma'} \cdot \mathbf{E}' \right) \mathbf{v}_s, \end{aligned} \quad (40)$$

which is exactly the same as Equation (18) from the KSB-SC system plus the additional term⁹ $(q/m)(\mathbf{u}' \cdot \mathbf{E}') \mathbf{v}_s/\gamma'$. Note that the correspondence with the KSB-SC momentum equation is not a coincidence: in both cases, the evolution of the particle velocity is calculated by separating out $\mathbf{v}_s$.

To summarize our choice of equations, we have (dropping all primes for simplicity)

$$\begin{cases} \partial_t B_x = -c\partial_y E_z + c\partial_z E_y - v_{s,y}\partial_y B_x \\ \partial_t B_y = c\partial_x E_z - c\partial_z E_x - v_{s,y}\partial_y B_y - s\Omega_0 B_x, \\ \partial_t B_z = -c\partial_x E_y + c\partial_y E_x - v_{s,y}\partial_y B_z \end{cases} \quad (41)$$

⁹ This additional term is not a peculiarity of our formulation; it is simply neglected in Riquelme et al. (2012) and Inchingolo et al. (2018). This term represents additional inertial forces proportional to the energy gain of a particle due to work exerted by electric fields (Arendt 1998).

$$\begin{cases} \partial_t E_x = c \partial_y B_z - c \partial_z B_y + v_{s,y} \partial_y E_x - v_{s,y} \partial_x E_y - 4\pi J_x \\ \partial_t E_y = -c \partial_x B_z + c \partial_z B_x \\ \quad - v_{s,y} \partial_x E_x - v_{s,y} \partial_y E_y - v_{s,y} \partial_z E_z - 4\pi J_y \\ \partial_t E_z = c \partial_x B_y - c \partial_y B_x + v_{s,y} \partial_y E_z - v_{s,y} \partial_z E_y - 4\pi J_z \end{cases}, \quad (42)$$

$$\nabla \cdot \mathbf{B} = 0, \quad (43)$$

$$\nabla \cdot \left(\mathbf{E} - \frac{\mathbf{v}_s}{c} \times \mathbf{B} \right) = 4\pi\rho, \quad (44)$$

for the electromagnetic fields, and

$$\begin{aligned} \frac{d\mathbf{u}}{dt} &= \frac{q}{m} \left(\mathbf{E} + \frac{\mathbf{u}}{c\gamma} \times \mathbf{B} \right) \\ &+ 2\mathbf{u} \times \boldsymbol{\Omega}_0 + s\Omega_0 u_x \hat{\mathbf{e}}_y - \frac{q}{m} \left(\frac{\mathbf{u}}{\gamma} \cdot \mathbf{E} \right) \mathbf{v}_s, \end{aligned} \quad (45)$$

$$\frac{d\mathbf{x}}{dt} = \frac{\mathbf{u}}{\gamma} + \mathbf{v}_s, \quad (46)$$

for the particle motion. Equations (41)–(46) constitute the paradigm for the *kinetic shearing box with orbital advection* to be employed in our 3D simulations, and will be referred to as the KSB-OA system from now on. This set of equations requires simpler boundary conditions than the KSB case and can be directly applied to 3D studies, contrary to the KSB-SC framework. The numerical approach to solve these equations is described in detail in Appendix A.

Finally, we note that Equation (46) can result in superluminal particle motion, especially when $u/\gamma \sim c$, because the velocity addition with $\mathbf{v}_s$ is nonrelativistic. To avoid this issue, we can apply the proper relativistic boost back to the observer's frame, to ensure that the velocity on the right-hand side is always subluminal:

$$\frac{dx}{dt} = \frac{u_x}{[\gamma + u_y v_{s,y}] \gamma_s}, \quad (47)$$

$$\frac{dy}{dt} = \frac{u_y + \gamma v_{s,y}}{\gamma + u_y v_{s,y}}, \quad (48)$$

$$\frac{dz}{dt} = \frac{u_z}{[\gamma + u_y v_{s,y}] \gamma_s}. \quad (49)$$

The numerical solution of these equations is slightly more complicated, and requires a simple nonlinear iteration (see Appendix A.2 for details). In large-scale simulations, however, this difference in computational cost may be significant. We thus resort to employing Equations (47)–(49) only when superluminal motion is detected during a run. Note, however, that these equations of motion are not entirely consistent with the derivation of the momentum Equation (45) given above, where $d\mathbf{x}/dt = \mathbf{u}/\gamma + \mathbf{v}_s$ was assumed. We have verified that this discrepancy does not impact our results significantly.

We emphasize that, in deriving the equations above, we have implicitly assumed that average bulk flows remain essentially nonrelativistic throughout the development of the MRI and the subsequent nonlinear turbulent stage. This assumption is manifest in the adoption of Galilean transformations for the boost of quantities to the comoving frame; the requirement of maintaining subrelativistic flows also applies to other methods discussed in the previous sections. In all our runs presented in this work, we have indeed verified that this assumption is respected.

3. Choice of Physical Parameters and Computational Considerations

In Section 1, we have discussed the scale hierarchy that characterizes multiscale systems such as collisionless accretion disks. In realistic scenarios, the various length scales represented in Figure 1 typically differ by orders of magnitude; in simulations, we are bound to employ a drastically reduced scale separation, due to computational constraints. Here, we illustrate how the choice of physical parameters influences the computational cost of simulations, as well as the limitations that affect such numerical experiments.

The simulations presented in the next sections always start from a weak vertical magnetic field $B_{z,0}$ in a box filled with uniform, thermal pair plasma (with particle mass $m = m_i = m_e$, number density $n_0 = n_{i,0} = n_{e,0}$) with dimensionless temperature $\theta_0 = kT_0/(mc^2)$ (with $T_0 = T_{i,0} = T_{e,0}$) and zero bulk speed (measured in the frame moving with $\mathbf{v}_s$). In addition to θ_0 , the other basic dimensionless parameters we must choose in our simulations are the initial Alfvén speed (normalized over the speed of light, and calculated with the nonrelativistic expression) $v_{A,0}/c = B_{z,0}/\sqrt{4\pi n_0 m_e (1 + m_i/m_e) c^2}$ and the cyclotron-to-rotational frequency ratio $\omega_c/\Omega_0 = qB_{z,0}/(mc\Omega_0)$. Setting these quantities also automatically produces the value of the total (counting both species) thermal-to-magnetic pressure ratio $\beta_0 = 2n_0 kT/[B_{z,0}^2/(8\pi)] = 2\theta_0/(v_{A,0}^2/c^2)$. To choose the free simulation parameters, we need to consider (i) the underlying assumptions of our model, i.e., a nonrelativistic shearing box (where volume-averaged bulk flows always remain subrelativistic), and (ii) the astrophysical environments we target, i.e., accretion disks around SMBHs. To respect the nonrelativistic shearing-box assumptions, we need $v_{A,0}/c \ll 1$; for conditions appropriate to astrophysical accretion disks, we need $\theta_0 < 1$ and $\omega_c/\Omega_0 \gg 1$. Setting these free parameters produces an ordering of microscopic and macroscopic scales; to appropriately model accretion around SMBHs, we should in principle respect the ordering depicted in Figure 1, but we will show that this is not always straightforward.

In terms of the free parameters θ_0 , $v_{A,0}/c$, and ω_c/Ω_0 , the length scales shown in Figure 1, normalized to the total (including both species) plasma skin depth c/ω_p , are given by

$$\frac{\lambda_D}{c/\omega_p} = \sqrt{\theta_0}, \quad (50)$$

$$\frac{\rho_C}{c/\omega_p} = \frac{\sqrt{\theta_0}}{v_{A,0}/c}, \quad (51)$$

$$\frac{\lambda_{\text{MRI}}}{c/\omega_p} = 2\pi \frac{\omega_c}{\Omega_0}, \quad (52)$$

$$\frac{H}{c/\omega_p} = \frac{\omega_c}{\Omega_0} \frac{\sqrt{\theta_0}}{v_{A,0}/c}, \quad (53)$$

and $r_0/(c/\omega_p) \sim H/(c/\omega_p)$ if we assume a thick accretion disk, i.e., setting the pressure scale-height H automatically sets the distance r_0 from the central object.¹⁰ Imposing the aforementioned typical values $\theta_0 < 1$, $v_{A,0}/c \ll 1$ thus produces the nonrelativistic ordering of microscopic scales $\lambda_D < c/\omega_p < \rho_C$.

¹⁰ This last assumption indicates an underlying discrepancy in the application of the unstratified shearing-box model to thick accretion disks, where vertical stratification would in principle be required.

Next, we need to choose the separation between the kinetic scales and the macroscopic scales λ_{MRI} and H . These are related by

$$\frac{\lambda_{\text{MRI}}}{\rho_C} = 2\pi \frac{\omega_C}{\Omega_0} \frac{v_{A,0}/c}{\sqrt{\theta_0}}, \quad (54)$$

$$\frac{H}{\rho_C} = \frac{\omega_C}{\Omega_0}, \quad (55)$$

with

$$\frac{H}{\lambda_{\text{MRI}}} = \frac{\sqrt{\theta_0}}{2\pi v_{A,0}/c} = \frac{1}{2\pi\beta_0}. \quad (56)$$

We thus see that it is easy to impose $\lambda_{\text{MRI}}/\rho_C \gg 1$, $H/\rho_C \gg 1$ by carefully choosing $\omega_C/\Omega_0 \gg 1$, but the separation H/λ_{MRI} depends instead on the previous choice of θ_0 and $v_{A,0}/c$. Hence, care is needed in the choice of plasma temperature and Alfvén speed to ensure $\rho_C < \lambda_{\text{MRI}} < H$ while also keeping the correct microscopic ordering. We note that the free dimensionless parameters also determine the macroscopic-to-microscopic temporal-scale separation,

$$\frac{P_0}{\omega_p^{-1}} = 2\pi \frac{\omega_C/\Omega_0}{v_{A,0}/c}, \quad (57)$$

$$\frac{P_0}{\omega_C^{-1}} = 2\pi \frac{\omega_C}{\Omega_0}, \quad (58)$$

where $P_0 = 2\pi/\Omega_0$ is the orbital period, and the ordering $\omega_p^{-1} < \omega_C^{-1} < P_0$ must be respected.

When choosing parameters for numerical simulations, a realistic scale separation and correct scale ordering are generally hard to achieve. The computational cost of a typical shearing-box simulation directly derives from the relations expressed above, with the addition that the vertical box size L_z/λ_{MRI} and the final simulation time t_{end}/P_0 enter as additional macroscopic scales that should respect the ordering $\lambda_{\text{MRI}} < L_z < H$ and $P_0 < t_{\text{end}}$. As we will show in Section 4, physically valid simulations indeed require box sizes of at least a few λ_{MRI} in each spatial direction, as well as $t_{\text{end}} \gtrsim 10P_0$. In this work, we will always employ a spatial resolution $\Delta x \simeq c/\omega_p$; while this choice does not typically resolve the Debye length λ_D at the beginning of the simulation, in all cases the plasma is heated quickly enough to increase the effective Debye length to scales larger than the grid spacing well before the nonlinear MRI stage. We have also checked that our results do not change significantly with higher resolution. The choice of temporal resolution immediately follows as $\Delta t \simeq r_{\text{CFL}}\omega_p^{-1}$ (where $r_{\text{CFL}} \simeq 2^{-1/2}$ and $r_{\text{CFL}} \simeq 2^{-1/3}$ are the 2D and 3D Courant–Friedrichs–Lewy factors, respectively).

For 2D simulations, we will employ the KSB-SC paradigm and typically allow for $L_z > H$ in order to avoid excessive computing costs (see below) when conducting the parameter-space scan presented in Section 4.3. We have verified that the qualitative behavior of the results is the same when employing parameters that instead respect the constraint $L_z \leq H$. The computational cost for our 2D runs can thus be estimated using

$$N_g \equiv \frac{L}{\Delta x} \simeq \frac{L}{c/\omega_p} = \frac{L}{\lambda_{\text{MRI}}} \frac{\lambda_{\text{MRI}}}{\rho_C} \frac{\rho_C}{c/\omega_p} \propto \frac{L}{\lambda_{\text{MRI}}} \frac{\omega_C}{\Omega_0} \quad (59)$$

for the number of required grid points N_g along a generic spatial direction, and

$$N_t \simeq \frac{t_{\text{end}}}{r_{\text{CFL}}\omega_p^{-1}} = \frac{t_{\text{end}}}{P_0} \frac{P_0}{r_{\text{CFL}}\omega_p^{-1}} \propto \frac{t_{\text{end}}}{P_0} \frac{\omega_C/\Omega_0}{v_{A,0}/c} \quad (60)$$

for the number of time steps N_t needed to reach a final time t_{end} . We can observe that, in terms of the free parameters, the required grid size increases linearly (in each direction, i.e., as N_g^2 in 2D runs) with ω_C/Ω_0 . In 2D runs, the typical choice for the free parameters is $\omega_C/\Omega_0 = 10\text{--}30$ (but we have run simulations with much larger values), $v_{A,0}/c = 0.01$, $\theta_0 = 1/32$ (i.e., $\beta_0 = 624$ for both species), which translate into $\lambda_{\text{MRI}}/\Delta x \simeq 60\text{--}180$ and $P_0/\Delta t \simeq 8000\text{--}25,000$. With these values, the microscopic scales are rather well-separated, $\rho_C/(c/\omega_p) \simeq 17.8$; however, the macroscopic scales are much less distinguishable, i.e., $H/\lambda_{\text{MRI}} \simeq 2.8$. Furthermore, the separation between macroscopic and microscopic scales is large, but still far from realistic, i.e., $\lambda_{\text{MRI}}/\rho_C \simeq 10.6$. For a domain size $L = 8\lambda_{\text{MRI}}$ in each dimension and a final simulation time $t_{\text{end}} = 15P_0$, these parameters imply evolving a grid of size up to $\sim 1500 \times 1500$ cells for $\sim 400,000$ time steps. Achieving better scale separation is in principle possible, e.g., by increasing ω_C/Ω_0 ; however, the required grid size and simulation duration can become prohibitive well before reaching realistic values (e.g., an estimated $\omega_C/\Omega_0 \sim 10^7$ around M87*).

With the insight gained from our 2D campaign, we will run 3D simulations employing the KSB-OA formulation and demand $L_z = H$, for complete consistency with the spatial scale ordering. We will then vary L_x/L_z and L_y/L_z independently, taking care to ensure that the velocity offset $v_s(L_x) = s\Omega_0 L_x < c$ at x -boundaries.¹¹ The required grid size in each direction is thus

$$N_g \simeq \frac{L}{c/\omega_p} = \frac{L}{H} \frac{H}{\lambda_{\text{MRI}}} \frac{\lambda_{\text{MRI}}}{\rho_C} \frac{\rho_C}{c/\omega_p} \propto \frac{L}{H} \frac{\omega_C}{\Omega_0} \frac{\sqrt{\theta_0}}{v_{A,0}/c}, \quad (61)$$

while the number of time steps is computed as in the 2D case. In general, $\sqrt{\theta_0}/(v_{A,0}/c) > 1$, and therefore it is clear that 3D runs will be more expensive than 2D simulations, not simply due to the dimensionality but also in terms of the effect of the free parameters. For example, a large-box 3D simulation with $\omega_C/\Omega_0 = 15$, $v_{A,0}/c \simeq 0.007$, $\theta_0 = 1/128$ in a box of size $L_z = 2\lambda_{\text{MRI}}$, $L_y = 2L_x = 8\lambda_{\text{MRI}}$, run until $t_{\text{end}} = 15P_0$, implies evolving a grid of $\sim 256 \times 1024 \times 128$ cells for $\sim 250,000$ time steps.

Finally, we note that the reasoning outlined above only relates to the *initial conditions* of each simulation; in practice, as the system evolves, macroscopic and microscopic scales tend to become better separated as the magnetic field is amplified and ρ_C is reduced. The MRI dynamics also helps to better resolve the kinetic scales (e.g., the skin depth), because these will become larger as plasma is heated during the system evolution.

¹¹ This is not necessary in 2D: as mentioned by Riquelme et al. (2012), in the shearing-coordinate frame, the global profile of the shearing velocity is $v_s/c = -\tanh(s\Omega_0 x/c)\hat{e}_y$, such that $v_s < c$ regardless of the simulation size. This does not apply to our 3D simulations, which do not employ shearing coordinates.

Table 1
List of 2D Simulations

Box Size ($L_x \times L_z$)	ω_C/Ω_0	β_0	$\lambda_{\text{MRI}}/\rho_C$
$2 \times 2\lambda_{\text{MRI}}^2$	6	624	2.13
$2 \times 2\lambda_{\text{MRI}}^2$	30	624	10.67
$2 \times 2\lambda_{\text{MRI}}^2$	120	624	42.65
$2 \times 2\lambda_{\text{MRI}}^2$	240	624	85.30
$4 \times 4\lambda_{\text{MRI}}^2$	30	624	10.67
$8 \times 8\lambda_{\text{MRI}}^2$	6	624	2.13
$8 \times 8\lambda_{\text{MRI}}^2$	15	624	5.33
$8 \times 8\lambda_{\text{MRI}}^2$	30	0.1	842.30
$8 \times 8\lambda_{\text{MRI}}^2$	30	1	266.36
$8 \times 8\lambda_{\text{MRI}}^2$	30	5	119.12
$8 \times 8\lambda_{\text{MRI}}^2$	30	78	30.16
$8 \times 8\lambda_{\text{MRI}}^2$	30	624	10.67
$8 \times 8\lambda_{\text{MRI}}^2$	30	2,496	5.33
$8 \times 8\lambda_{\text{MRI}}^2$	60	624	21.33
$16 \times 4\lambda_{\text{MRI}}^2$	30	624	10.67
$16 \times 16\lambda_{\text{MRI}}^2$	30	624	10.67
$32 \times 8\lambda_{\text{MRI}}^2$	30	624	10.67
$32 \times 32\lambda_{\text{MRI}}^2$	30	624	10.67

Note. In all cases, $v_{A,0}/c = 0.01$; the numerical resolution is such that $\Delta x \simeq c/\omega_p$, and we employ 25 particles per cell per species.

4. Two-dimensional Simulations of the MRI in Collisionless Pair Plasmas

In this section, we present the results of several PIC simulations of the collisionless MRI in 2D. The parameters for each simulation are specified in Table 1. In all runs, we solve the KSB-SC system with the relativistic PIC code ZELTRON (Cerutti et al. 2013). Our initial setup consists of a 2D (in the xz poloidal plane) box filled with thermal plasma and a weak vertical magnetic field determined by our choice of initial parameters ω_C/Ω_0 , $v_{A,0}$, and β_0 (or equivalently, θ_0). We always employ 25 particles per cell per species, and adopt a resolution such that $\Delta x \simeq c/\omega_p$. Our largest 2D simulations employ computational grids with 4480^2 cells, for a total of over one billion particles. Finally, note that, in our runs, we do not perturb the initial conditions, such that the MRI grows from random, base-level fluctuations (i.e., numerical noise).

In this section, we first briefly describe the typical stages of the evolution of the collisionless MRI in two-dimensions (Section 4.1). We then provide a physical explanation for all the features observed in this evolution (Section 4.2). We subsequently present an exploration of the physical-parameter space characterizing kinetic shearing-box simulations of the MRI (Section 4.3). Then, we discuss the characteristics of MRI-driven turbulence (Section 4.4), particle energization (Section 4.5), angular-momentum transport (Section 4.6), and the zero net-flux case (Section 4.7). From these runs, we identify a suitable set of parameters to be employed in the much more expensive 3D simulations presented in Section 5.

4.1. Typical Evolution of the 2D Collisionless MRI

Figure 2 describes the evolution of a typical two-dimensional, PIC MRI simulation of limited box size. The top-right panel shows evolution of the change (with respect to the initial value) in magnetic energy for each component of $\mathbf{B}$ (normalized to the total energy at $t = 0$). This run employs the

parameters $\omega_C/\Omega_0 = 120$, $\beta_0 = 624$, $v_{A,0}/c = 0.01$ in a “small” box of size $2 \times 2\lambda_{\text{MRI}}^2$. The evolution proceeds through several stages, with the beginning of each phase marked by vertical black lines in Figure 2. These phases qualitatively characterize all 2D simulations with finite, net vertical magnetic flux, independently of the system size and parameters chosen for the run (aside from a less clear distinction between phases in larger systems—see the explanation below).

The magnetic energy grows from a pre-instability state (phase 1), which lasts for $\sim 2P_0$, to the linear MRI growth (phase 2), where the magnetic-field strength increases exponentially by several orders of magnitude. During the linear stage, the initially vertical magnetic field acquires radial and toroidal components, resulting in bent magnetic field lines (in the xz -plane) that self-organize into “channel” flows. This is visible in Figure 2 (top left), which shows that both the in-plane magnetic field $B_{\text{pol}} = \sqrt{B_x^2 + B_z^2}$ and the out-of-plane B_y present channel structures. These channels are of vertical size $\sim \lambda_{\text{MRI}}$, and the radial and toroidal fields are spatially uniform inside each channel; both field components reverse sign at the channel interfaces. The channels represent an exact nonlinear, equilibrium solution of the MHD equations (Goodman & Xu 1994); this state is, however, unstable to linear perturbations, which can lead to further evolution through secondary (parasitic) instabilities developing on top of the primary MRI (Goodman & Xu 1994; Latter et al. 2009; Pessah & Goodman 2009). These include tearing modes, which we see developing in our simulation (Figure 2): throughout the linear stage, the magnetic-field strength in each channel increases, and the thickness of the interfaces where $\mathbf{B}$ changes sign decreases; current sheets that form at these interfaces thin out until, at $t \simeq 3P_0$, we observe an abrupt drop in magnetic-field energy. This moment marks the onset of magnetic reconnection, mediated by the tearing instability in the thinning current sheets, and the beginning of phase 3 of the evolution (Figure 2, center left).

Phase 3 is the most interesting stage of the evolution. This phase is in principle representative of the typical state of turbulent accretion disks, where particles (electrons and ions) can experience acceleration. For electrons, this acceleration may become strong enough to reach radiative regimes (e.g., Ryan et al. 2017; Dexter et al. 2021). The nonlinear stage in phase 3 is characterized by the competition between the MRI, which actively strengthens the magnetic field in the channels, and magnetic reconnection, which depletes magnetic energy in the current sheets, thereby energizing the plasma. In addition to the reconnection process, substantial kinking of the current sheets can in principle disrupt the channels and bring the system to a saturated, fully turbulent state, which should remain sustained throughout phase 3. However, here we do not observe the development of such a dynamics. In fact, turbulence does not generally develop in small-box runs, where channels persist without being disrupted throughout several reconnection events (Figure 2, bottom left), marked by repeated drops in the magnetic-energy evolution during phase 3. The continuous action of the MRI during this phase is indicated by subsequent stages of magnetic-energy growth, which do, however, proceed with decreasing growth rates. This magnetic-field amplification is particularly strong along the (out-of-plane) y -direction; the in-plane components show a much smaller increase.




Figure 2. Evolution of a typical PIC simulation of the pair-plasma MRI in a small 2D box ($2 \times 2\lambda_{\text{MRI}}^2$). Left column: spatial distribution of the out-of-plane magnetic field B_y and of the in-plane field strength $B_{\text{pol}} = \sqrt{B_x^2 + B_z^2}$ at subsequent times during the simulation. Right column: evolution of the change in magnetic energy in all three components of $\mathbf{B}$ (top) and evolution of the volume-averaged $\beta_{\parallel} = p_{\parallel}/[B^2/(8\pi)]$ and $\beta_{\perp} = p_{\perp}/[B^2/(8\pi)]$ (bottom). The beginnings of different phases characterizing the system evolution (described in the text) are marked with vertical black lines.

Finally, at $t \simeq 7P_0$, we observe that the increase in magnetic energy is halted, and the system rapidly relaxes to a quiescent state that remains unchanged endlessly (phase 4). This results from the MRI growth rate decreasing and the unstable MRI modes migrating to longer and longer wavelengths. Eventually, these modes grow to length scales larger than the box size, and the MRI stops. Previous works have shown that magnetic fields then reorganize into flux tubes that are stable to axisymmetric perturbations and impede further MRI growth (Riquelme et al. 2012; Hoshino 2013); in our runs, we observe the same behavior during phase 4 (not shown in the left panels of Figure 2). In realistic environments, the MRI is not expected to be suppressed over such short timescales; as we discuss in the next sections, this effect results from the typical (unrealistic) choice of parameters employed in PIC runs.

As mentioned above, the nonlinear stage of the MRI (phase 3) is expected to achieve a fully turbulent state, which is not observed in any small-box run. A regime of developed turbulence can be attained only by increasing the domain size. Figure 3 shows the typical evolution of a large-box PIC MRI simulation in 2D, with parameters $\omega_c/\Omega_0 = 30$, $\beta_0 = 624$, $v_{A,0}/c = 0.01$, and a box size $32 \times 32\lambda_{\text{MRI}}^2$, i.e., 16 times larger than the previous run. The left panels show that the (much more numerous) initial channels formed during the linear stage break up into smaller structures: this breakup corresponds to the beginning of the nonlinear state (i.e., phase 3 in the evolution). The magnetic-energy evolution (top right) shows that, within 30 orbits, the system is slowly reaching the end of phase 3 (notice the difference in timescales between this and the previous run); the energy in B_y is still increasing, indicating

that the MRI has not been completely halted yet. However, over sufficiently long timescales, the behavior of the MRI in the large- and small-box cases is similar: in our simulations, we have observed that arbitrarily large systems eventually evolve to the same final quiescent state where the MRI has stopped, with large-scale magnetic loops that persist indefinitely. The remarkable difference between the two cases is thus represented by the presence of turbulence, which only manifests when employing large domain sizes. Notice also how the presence of turbulence in the large-box case masks subsequent cycles of MRI growth/reconnection, which can instead be clearly identified in the small-box run.

The two simulations discussed in this section are representative of the state of the art of fully kinetic, two-dimensional MRI simulations in pair plasmas with initial vertical field. In fact, both simulations employ much larger scale separation—in terms of ω_c/Ω_0 and/or box size—than any previous work. Still, several questions on the physical evolution of such a system remain open: it is thus far unclear what exactly causes the subsequent MRI growth-reconnection cycles during phase 3 in the small-box case (and the lack of such a clear distinction between cycles in the large-box case); it is also not apparent how the choice of physical parameters drives the transition to a turbulent state and what causes the decay of turbulence over time; moreover, it remains unexplained why, even in the large-box case, the B_y energy grows much larger than the other components, without reaching a clear saturated state such as that observed in typical MHD simulations (e.g., Pastorello et al. 2013). In the next section, we provide a physical explanation for all these observations.




Figure 3. As in Figure 2, but for a large 2D box ($32 \times 32 \lambda_{\text{MRI}}^2$).

4.2. Physics of the 2D Collisionless MRI

In principle, the desired outcome of any MRI simulation is a saturated, turbulent nonlinear state that lasts for long times. However, no 2D PIC simulation of the collisionless MRI with nonzero net flux presented in literature (including this work) shows this behavior. While the linear MRI stage is generally well-captured (for reasonable simulation parameters), from the onset of magnetic reconnection, the overall behavior of 2D PIC runs seems to diverge from 3D MHD expectations. Figures 2 and 3 in particular show that the magnetic energy increases steadily, instead of remaining roughly at the level reached at the end of the linear stage. Furthermore, Figure 3 points out that, in large systems, a quasi-constant, saturated level can in fact be maintained at least for the energy in B_x and B_z (i.e., the in-plane components), but not for B_y . We explain this issue, together with the overall 2D MRI dynamics in PIC, with the following reasoning.

At the end of the linear stage, to maintain a saturated state, the magnetic-field amplification by the MRI should be approximately balanced by dissipative mechanisms. Strong magnetic fields induced by the MRI can in principle be rapidly destroyed via reconnection, generally resulting in turbulence and small-scale magnetic structures, which can then seed the MRI again in a continuous cycle. In 3D MHD simulations with sufficiently large box size, this is realized without issues (see Latter et al. 2009 for a discussion on MHD simulations with varying box size). However, 2D PIC simulations have reduced dimensionality, and also allow for the presence of pressure anisotropy. The latter is typically not included in standard MHD (with exceptions, e.g., Sharma et al. 2006). Both the 2D geometry and the development of pressure anisotropy play an important role in axisymmetric PIC MRI simulations.

Figures 2 and 3 (bottom right panels) show the evolution of the volume-averaged $\beta_{\parallel} = p_{\parallel}/[B^2/(8\pi)]$ and $\beta_{\perp} = p_{\perp}/[B^2/(8\pi)]$, where $p_{\parallel}$ and $p_{\perp}$ are the parallel and perpendicular (to the magnetic field) thermal plasma pressure. We observe that non-negligible pressure anisotropy $p_{\perp} > p_{\parallel}$ develops during the MRI evolution, starting from the linear phase: a strong increase in B is accompanied by an increase in perpendicular particle energy, due to the invariance of the magnetic moment $\mu = mu_{\perp}^2/(2B)$. This anisotropy alone can drive three distinct mechanisms: (i) an increase in the growth rate of tearing modes (e.g., Chen & Palmadesso 1984); (ii) the excitation of mirror modes (e.g., Kunz et al. 2014a), which can also couple with tearing modes, further increasing the tearing growth rate (Alt & Kunz 2019; Winarto & Kunz 2022); and (iii) a decrease in the growth rate of the MRI (Quataert et al. 2002). During the linear stage, the first two mechanisms cause the rapid onset of magnetic reconnection: the enhancement of tearing modes occurs concurrently with the natural thinning of current sheets that form at the interface of channel flows. At the same time, the MRI growth rate is reduced both by pressure anisotropy (i.e., the third mechanism above) and by the fact that, by this time, v_A has increased substantially and the most-unstable MRI wavelength $\lambda_{\text{MRI}} = 2\pi v_A/\Omega_0$ is much larger (possibly larger than the box size) than in the beginning of the simulation. As a result, as reconnection is initiated, the MRI drastically (but temporarily) slows down.

We note that the excitation of the mirror instability mentioned above should in principle limit the growth of pressure anisotropy, because in realistic systems, mirror modes can act on timescales comparable to those of the MRI. This represents a crucial point in typical PIC MRI modeling: the usual choice of physical parameters is such that, in simulations, mirror modes grow unrealistically slowly. This can be simply

estimated by demanding that the mirror-instability growth rate be comparable to that of the fastest-growing MRI mode, i.e., $\omega_C \Lambda_m^2 \sim 3\Omega_0/4$, where $\Lambda_m \equiv p_\perp/p_\parallel - 1 - 1/\beta_\perp$ (Kunz et al. 2014a). During the linear stage, we measure a typical peak value $\Lambda_m \sim 10^{-1}$, implying that $\omega_C/\Omega_0 \sim 10^2$ would be required for fast mirror modes; the measured typical value of Λ_m is even lower during the nonlinear phase, achieving $\Lambda_m \leq 10^{-2}$, which requires $\omega_C/\Omega_0 \sim 10^4$ to satisfy the condition above. As discussed in Section 3, this regime simply cannot be achieved realistically in simulations. As a consequence, the MRI dynamics in PIC runs will inevitably result in the development of larger pressure anisotropy than expected with realistic physical parameters.¹²

A large pressure anisotropy, combined with the migration of unstable wavelengths to larger scales, could in principle suppress the MRI from proceeding further already at the end of the linear stage. Fortunately, magnetic reconnection can at least partially counter this process, both by decreasing the field strength and by decreasing pressure anisotropy via parallel particle acceleration at reconnection sites. This is exactly what we observe for both the small- and the large-box cases right after the onset of reconnection (Figures 2 and 3, bottom right). In all our runs, we consistently find that reconnection events push the anisotropy in the direction opposite to that induced by the MRI, at times achieving $p_\parallel \simeq p_\perp$. We also note that an additional (but probably largely unimportant) mechanism of anisotropy reduction may stem from the conservation of the invariant $p_\parallel B^2/\rho^3$ (e.g., Sharma et al. 2006 and references therein), which further increases $p_\parallel$ when magnetic energy is depleted during reconnection (Hoshino 2013).

As reconnection dissipates magnetic energy and reduces pressure anisotropy, the system returns to conditions that allow for the MRI to develop. At the same time, tearing modes are progressively quenched due to the thickening of current sheets at channel interfaces. This temporarily halts the reconnection process. The MRI is therefore free to start again, commencing a new exponential-growth stage (as can be observed during phase 3 of the evolution in Figure 2, top right). However, recall that, by this time, v_A is much larger than in the initial state, and the most-unstable MRI modes may no longer fit inside the simulation box. The MRI growth rate is then appreciably smaller than the asymptotic value $3\Omega_0/4$. This new MRI phase thus amplifies magnetic fields more slowly, but the overall dynamics is similar to that of phase 2: $p_\perp > p_\parallel$ anisotropy is again realized due to μ conservation; current sheets become thinner at the channel interfaces; magnetic reconnection eventually disrupts the sheets, and so on. During each cycle, the MRI growth rate decreases as v_A increases further; in addition, while reconnection tends to re-establish isotropy, in our 2D runs we observe that the $p_\perp > p_\parallel$ anisotropy accumulates across subsequent cycles, increasing over time. Reconnection and (unrealistically slow) mirror modes thus fail to keep the MRI-driven anisotropy limited. At each subsequent cycle, the accumulated effect of the increase in λ_{MRI} and in the pressure anisotropy pushes the MRI to smaller and smaller growth rates. Eventually, all unstable MRI wavelengths grow

beyond the box size, and the process cannot remain self-sustaining. The MRI completely stops and reconnection dissipates magnetic energy at the remaining current sheets. The system then reaches the quiescent state described in this and previous works (Riquelme et al. 2012; Hoshino 2013).

The reasoning outlined above explains why we observe the MRI slowing down and eventually stopping altogether in 2D simulations of any size: the lack of saturation in magnetic energy (especially in B_y) and unregulated pressure anisotropy, over time, cause the migration of MRI wavelengths to scales larger than the box size. This could be mitigated if magnetic-field saturation were achieved after the linear stage, and if pressure anisotropy could be efficiently limited. This is better verified in larger boxes: indeed, even though the end state of the evolution is the same, qualitative differences between the nonlinear evolution of the MRI arise between large and small simulations. The differences are due to three main reasons: (i) In small boxes, a large number of tearing modes are suppressed due to the limited separation between current-sheet length ($\sim L_x$) and thickness ($\sim c/\omega_p$), which slows reconnection and allows for magnetic fields to grow stronger via the MRI. (ii) Large boxes also allow for the formation of a larger number of channels along z , and for the excitation of the DKI in long current sheets, which is instead suppressed in small simulations (as noted by Inchingolo et al. 2018 as well as in Section 4.3.4 of this work). The DKI in larger boxes significantly contributes to breaking up and mixing the more numerous channel flows, inducing a turbulent state. (iii) Relatedly, the lack of developed turbulence in small boxes results in less efficient reduction of pressure anisotropy via turbulent magnetic-field dissipation and pitch-angle particle scattering. Magnetic fields and pressure anisotropy thus grow much faster when the MRI evolves in smaller boxes, resulting in a more rapid migration of unstable modes to wavelengths larger than the box size. This migration is slower in larger boxes, which can also fit larger wavelengths; as a consequence, the nonlinear turbulent stage is sustained for far longer.

The arguments above still do not explain why the MRI, in all 2D simulations, fails to saturate at an approximately constant level during the nonlinear stage, but rather induces a continuous growth in the magnetic energy (specifically in B_y ; see Figure 3, top right). This occurs irrespective of the box size, and represents a major discrepancy with 3D MHD simulations. This issue was not addressed in previous works, but it has in fact a clear explanation: in 2D axisymmetric simulations, the out-of-plane, toroidal magnetic field B_y cannot be subjected to magnetic reconnection. This is a fundamental problem for maintaining a saturated nonlinear state;¹³ without reconnection along y , the amplification of B_y via the MRI is practically unimpeded. The only mechanism counteracting this effect is represented by advection, mixing, and reconfiguration of opposite-sign B_y regions; but as observed in Figure 3 (left panels), even for large system sizes, the long-term evolution of the system inevitably consists of macroscopic B_y structures emerging from the turbulent state as the MRI is halted.

The reasoning outlined above leads us to conclude that, in 2D, achieving a minimally sustained turbulent state is in fact possible, but only in simulations with sufficiently large boxes (in accordance with Inchingolo et al. 2018). However, our

¹² It is important to emphasize that, with our typical choice of parameters, mirror instabilities will not be completely absent in the simulations we perform. Characteristic mirror “stripes” are indeed visible inside the channels during the linear stage (e.g., Figure 2, top left). Our (necessarily) unrealistic scale separation simply reduces the importance of these instabilities with respect to the overall dynamics. This must be taken into account when interpreting the physical results of our simulations.

¹³ Especially because the MRI is particularly efficient at amplifying magnetic fields in the y -direction, as can be inferred by simply observing the KSB-SC system—the only first-order (in Ω_0) source term appears in the equation for B_y .

analysis also shows that the ultimate fate of any 2D simulation is to reach a nonlinear stage that does not truly saturate, but rather features a strong, unimpeded increase in B_y , eventually resulting in the suppression of the MRI as unstable modes migrate to longer and longer wavelengths. We note that the absence of saturation for B_y was previously observed by Riquelme et al. (2012), who reported that 2D MHD simulations also display a similar behavior. In Section 5, we will show that the issues outlined above can be ameliorated in 3D simulations, where reconnection along y is allowed.

4.3. Parameter-space Exploration

Having established the general features of 2D kinetic MRI simulations, we now present a detailed parameter-space scan where we exhaustively explore the effect of four specific parameters. As illustrated in the previous sections, we expect 2D simulations to be fundamentally different from 3D runs; our exploration is nevertheless instructive, because many features of the 2D dynamics are still present in 3D (see Section 5). In particular, 2D results help us choose a reasonable set of parameters to be employed in 3D runs, which are much more expensive and are necessarily more limited in the parameter space that can be explored. The parameters we explore are:

1. The separation between macroscopic and microscopic temporal scales, parameterized by the ratio between the initial cyclotron frequency and the rotation frequency, ω_C/Ω_0 ; previous works have explored the range $\omega_C/\Omega_0 \simeq 10$ –20 (Riquelme et al. 2012; Hoshino 2013, 2015; Inchingolo et al. 2018). Here, we will employ a much larger range of values.
2. The separation between physical domain size and (macro-) instability scales in terms of the box size (e.g., in the vertical direction, $L=L_z$) compared to the initial most-unstable MRI wavelength λ_{MRI} . Previous 2D works have advocated the need to employ values $L/\lambda_{\text{MRI}} > 8$ to reach convergence in the results of kinetic MRI simulations (Inchingolo et al. 2018).
3. The initial plasma temperature, or equivalently, the initial thermal-to-magnetic pressure ratio $\beta_0 = 8\pi p/B_0^2$ (with $p = p_e + p_i$ the total pressure). This was previously investigated in the range $\beta_0 \simeq 100$ –6000 (Hoshino 2013).
4. The aspect ratio L_x/L_z of the simulation box. In particular, we will explore the case of boxes elongated along x . This point was not explored in previous kinetic-MRI works.

In this section, we will focus on assessing convergence in the results based on the magnetic-energy amplification. In light of the discussion presented in Section 4.1, we opt for excluding the B_y component from the convergence analysis, because the lack of reconnection along y in 2D simulations fundamentally affects the evolution of magnetic fields in the direction perpendicular to the simulation plane. The analysis of other important physical processes (such as turbulence, nonthermal particle acceleration, and angular momentum transport) will be presented in Section 4.4 and onward.

4.3.1. Effect of the Temporal-scale Separation (ω_C/Ω_0)

We begin our investigation by evaluating the impact of the ω_C/Ω_0 ratio. At first, we will focus on a limited box size; although small-box runs differ substantially from the large-box

case (see Section 4.2), they are still useful to understand the effect of scale separation when only MRI and reconnection are present, and fully developed turbulence is excluded. Moreover, the small-box case allows us to explore a much wider parameter space, since larger ω_C/Ω_0 values imply larger $\lambda_{\text{MRI}}/(c/\omega_p)$ ratios and therefore more resolution needed per λ_{MRI} (see Section 3). Limited box sizes will be the necessary condition for affordable 3D runs and it is therefore interesting to analyze this case in 2D before performing 3D simulations. Later, we will consider larger boxes in 2D resulting in a more realistic system evolution.

For this first analysis, our box has size $2 \times 2\lambda_{\text{MRI}}^2$. We run four simulations with $\omega_C/\Omega_0 = 6, 30, 120, 240$, keeping the initial $\beta_0 = 624$ (including both electrons and positrons; the corresponding temperature is $\theta_0 = 1/32$ for each species) and the initial Alfvén speed $v_{A,0}/c = 0.01$. The large range of frequency ratios we explore allows us to assess the effect of the separation between fluid (of size λ_{MRI}) and kinetic (of size ρ_C) scales when these differ substantially. In particular, our choices of parameters result in $\lambda_{\text{MRI}}/\rho_C \simeq 2, 10, 40$, and 80 for the four runs, respectively.

In Figure 4 (top left), we show the evolution of the change in magnetic energy (excluding the energy in B_y as mentioned in the previous section), normalized by the initial total energy of the system, for the four runs until $t = 10P_0$. All simulations present the common features outlined in Section 4.1, i.e., four typical phases of evolution. Each phase develops similarly as ω_C/Ω_0 increases, but fundamental differences are present. In particular, the linear phase is well-captured only for $\omega_C/\Omega_0 \geq 30$. As we increase ω_C/Ω_0 from 6 to 240, we observe that the saturation level of the magnetic energy (during the nonlinear stage) increases by roughly one order of magnitude. Larger ω_C/Ω_0 values also correspond to an increase in the duration of the nonlinear stage (before the flattening of the curves). The cases $\omega_C/\Omega_0 = 120$ and $\omega_C/\Omega_0 = 240$ exhibit very small differences in the results; the nonlinear stage has approximately the same duration, and the saturation level is approximately the same (at least in the early nonlinear stage, before the MRI slows down). This indicates that, with a fixed box size, the results are well-converged¹⁴ with respect to an increase in ω_C/Ω_0 , even though the latter is still orders of magnitude away from any realistic values (e.g., $\omega_C/\Omega_0 \sim 10^7$ around M87*).

The small-box case is not representative of a realistic MRI evolution, as it lacks important dynamics such as fully developed turbulence. For this reason, we run another series of simulations, employing a larger box of size $8 \times 8\lambda_{\text{MRI}}^2$ and the values $\omega_C/\Omega_0 = 6, 15, 30$, and 60 (all other parameters are the same as for the small-box case). The results are shown in Figure 4 (top right; note the longer timescales on the x -axis). Qualitatively, this medium-box case presents the same trend displayed by the previous set of runs: the linear phase is well-captured for $\omega_C/\Omega_0 \geq 15$, and the nonlinear saturation level and nonlinear stage duration increase with ω_C/Ω_0 . Additionally, this case displays a full breakup of channels and a turbulent behavior throughout the nonlinear stage. We observe that the magnetic-energy saturation is roughly the same for $\omega_C/\Omega_0 \geq 15$; during the nonlinear stage, the results appear well-converged between $\omega_C/\Omega_0 = 30$ and $\omega_C/\Omega_0 = 60$, and they only diverge at late times (when the MRI is slowing down). We

¹⁴ Recall that, in these plots, we have excluded the B_y component from the total magnetic energy.




Figure 4. Evolution of the change in the in-plane magnetic energy for our parameter-space exploration of the 2D pair-plasma MRI. Top left and top right: varying ω_C/Ω_0 at a fixed β_0 for a small- and a medium-size simulation box, respectively. Bottom left: varying the box size at a fixed ω_C/Ω_0 and β_0 . Bottom right: varying β_0 at a fixed ω_C/Ω_0 for a medium-size box.

conclude that, when fixing the box to larger sizes, the results converge more quickly, as ω_C/Ω_0 increases, than in the small-box case.

This analysis shows that, when keeping the box size fixed, solely increasing the temporal-scale separation produces a nonlinear MRI evolution that converges for “large enough” ω_C/Ω_0 , both for small box sizes ($\sim 2 \times 2\lambda_{\text{MRI}}^2$) and for larger boxes ($\gtrsim 8 \times 8\lambda_{\text{MRI}}^2$). However, the converged state is qualitatively different in the small-box and in the large-box case, and the value of ω_C/Ω_0 at which convergence is reached also depends on the box size. These considerations are likely to carry over to 3D simulations, which will be explored in Section 5.

4.3.2. Effect of the Box Size (L/λ_{MRI})

In this section, we explore the effect of an increasing box size L/λ_{MRI} . We run five simulations, progressively enlarging the simulation domain from $2 \times 2\lambda_{\text{MRI}}^2$ up to $32 \times 32\lambda_{\text{MRI}}^2$. In all cases, we fix $\omega_C/\Omega_0 = 30$, $\beta_0 = 624$, and $v_{A,0}/c = 0.01$. Figure 4 (bottom-left panel) reports the change in magnetic energy for all runs until $t = 30P_0$. We observe that larger

domains tend to increase the duration of the nonlinear stage; for a domain size $\gtrsim 16 \times 16\lambda_{\text{MRI}}^2$, the system is not yet completely quiescent at 30 orbits. In addition, the saturation level is roughly the same for system sizes of at least $8 \times 8\lambda_{\text{MRI}}^2$, which suggests that convergence has been reached. This observation appears consistent with the conclusions presented in Inchingolo et al. (2018), but there are fundamental differences between their considerations and the ones presented here.

First, our runs employ a larger $\omega_C/\Omega_0 = 30$ ($\omega_C/\Omega_0 = 10$ was chosen for the simulations in Inchingolo et al. 2018). This difference may appear unimportant, but we have shown in Figure 4 (top right) that $\omega_C/\Omega_0 < 15$ does not reliably produce converged magnetic-field amplification for a box size of $8 \times 8\lambda_{\text{MRI}}^2$. Second, Inchingolo et al. (2018) claim that, for a box size of at least $8 \times 8\lambda_{\text{MRI}}^2$, the *total* magnetic energy (including the out-of-plane B_y) converges; we do not observe such convergence in our results simply by increasing the box size. Rather, we see that the *in-plane* magnetic-field strength reaches convergence for sufficiently large boxes. Adding the out-of-plane field to this analysis does not show convergence even for very large values of ω_C/Ω_0 . Third, Inchingolo et al.

(2018) argue that an indicator of convergence is represented by a volume-averaged $v_A/c < 1$ (using the nonrelativistic expression for v_A), which they observe for box sizes of at least $8 \times 8 \lambda_{\text{MRI}}^2$. However, their analysis is limited to rather short timescales; in all our 2D numerical experiments (regardless of the box size), we invariably find that, on average, $v_A/c > 1$ over sufficiently long times, due to a lack of saturation in B_y (as we have argued in Section 4.2). This specific issue can be mitigated in 3D simulations, as we will show in Section 5.1.1.

From the analyses carried out here and in the previous section, we conclude that a combination of “large enough” frequency ratio and system size is necessary for two reasons, namely (i) to reach convergence (at least in terms of the in-plane magnetic-field amplification), i.e., to obtain results that become insensitive to further increase in both ω_C/Ω_0 and L/λ_{MRI} ; and (ii) to include all the relevant physics (such as turbulence) with sufficient scale separation. Solely varying one of these parameters does not ensure that the results will be converged and also physically valid. From our results, it appears that $\omega_C/\Omega_0 = 30$ and a system size of $8 \times 8 \lambda_{\text{MRI}}^2$ suffice in respecting this criterion while attaining satisfactory convergence in the results. However, we will show in Section 4.5 that extra care must be taken when choosing simulation parameters if one considers their effect on particle energization. Finally, we note that the aforementioned ω_C/Ω_0 and box size may still impose excessive computational costs in 3D runs. However, we will show in Section 5 that the different nature of the 3D MRI dynamics may allow for less stringent conditions on the physical parameters.

4.3.3. Effect of the Initial Plasma- β

In this section, we assess the effect of β_0 on the system evolution. This was previously investigated by Hoshino (2013), who explored the range $\beta_0 \simeq 100$ –6000; here, we run six simulations with $\beta_0 = 0.1, 1, 5, 78, 624, 2, \text{ and } 496$, respectively. For consistency with the parameters employed in the previous sections, we choose $\omega_C/\Omega_0 = 30$, $v_{A,0}/c = 0.01$, and a box size of $8 \times 8 \lambda_{\text{MRI}}^2$. Fixing these parameters implies that we are bound to change β_0 by changing the initial plasma temperature θ_0 correspondingly (see Section 3); to keep the initial temperature sufficiently far from relativistic values, we take a maximum $\beta_0 = 2496$ (corresponding to $\theta_0 = 1/8$).

The results of the six runs are shown in Figure 4 (bottom right). The magnetic-field strength at saturation appears to decrease steadily as β_0 increases; however, this is simply an effect of our choice of normalization over the initial total energy. In fact, we observe that the total (i.e., between initial state and the saturated nonlinear stage) magnetic-field amplification B^2/B_0^2 is roughly the same in all these simulations;¹⁵ because the energy of the system is dominated by the thermal energy, the ratio between magnetic energy and total initial energy will decrease with increasing θ_0 (and therefore β_0).

To better evaluate convergence in this case, we thus consider the quantity $B_{\text{pol}}/B_{\text{pol},0}$. This is shown in Figure 5 (left panel), which demonstrates that the amplification of the in-plane magnetic field is well-converged across all simulations and does not depend on the initial β_0 . In the same figure (right panel), we also plot the value of the volume-averaged

β in time. Regardless of β_0 , all simulations reach $\langle \beta \rangle \lesssim 1$ at the beginning of the nonlinear stage, followed by a gradual increase until $\langle \beta \rangle \sim 1$ –10 at $t \simeq 20P_0$, marking the end of the turbulent stage and the emergence of quiescent magnetic loops. The fact that $\langle \beta \rangle$ at saturation does not depend on β_0 and is roughly consistent (around 0.1–1) across simulations, together with the fact that the magnetic-field amplification is also comparable, implies that the particle energy at saturation is approximately the same for all simulations. This signifies that the particle-energy gain (i.e., measured as θ/θ_0 at saturation) is smaller if the initial θ_0 is larger. This also explains the behavior of the $\beta_0 = 0.1$ case, in which $\langle \beta \rangle$ increases between the initial state and the saturated value, instead of decreasing: in this case, θ/θ_0 is larger than B^2/B_0^2 and thus $\beta/\beta_0 > 1$ at saturation. This dynamics can be understood via the following argument: because the magnetic energy at saturation is the same for all cases, the amount of energy available for dissipation via reconnection (and thus convertible into kinetic energy) is the same. The net increase in particle energy will thus be equal across all cases, but the ratio between final and initial energy will be larger if θ_0 (and hence β_0) is smaller.

These results show that no substantial differences arise for values $\beta_0 \simeq 0.1$ –2500, at least in terms of global quantities such as the magnetic-field amplification and the plasma- β at saturation. The system evolution appears very similar, particularly during the nonlinear phase, in which the value of β may influence mechanisms such as turbulent particle energization. Our results are in good agreement with those presented in Hoshino (2013), with the caveat that, in their case, $v_{A,0}$ was also varied alongside θ_0 to change the initial β_0 . We have also explored the case $\beta_0 \sim 0.1$, while Hoshino (2013) employed a minimum value $\beta_0 \sim 96$. Our small- β_0 case is particularly interesting, as it produces the same results found for much larger β_0 cases; in literature, it is instead generally thought that a small β_0 should completely impede the development of the MRI. However, this only applies when the vertical box size is comparable to or smaller than H : indeed, choosing $H/L_z = 1$ and varying θ_0 (and hence β_0) while keeping ω_C/Ω_0 and $v_{A,0}$ fixed could produce a situation in which the MRI cannot develop because the vertical box size becomes smaller than λ_{MRI} . In our case, however, we have ignored this constraint and taken $L_z = 8 \lambda_{\text{MRI}}$ (which can result in $L_z \gg H$, e.g., in our $\beta_0 = 0.1$ case). Thus, in all our simulations, the vertical box size can still fit channel modes and the MRI can develop regardless of the initial β_0 . Here, we choose to violate $H/L_z \geq 1$ because, when this condition is imposed (e.g., in our 3D simulations, Section 5), the chosen value of β_0 results in additional constraints on length and timescales (see Section 3). We have, however, verified that our results do not qualitatively change when respecting the $H/L_z \geq 1$ condition.

4.3.4. Effect of the Box Aspect Ratio

Finally, we analyze the effect of the box aspect ratio (L_x/L_z) in 2D MRI simulations, which has not been explored in previous kinetic studies. This analysis is motivated by multiple reasons: first, 3D MHD simulations have shown that domains with aspect ratio $L_x/L_z \neq 1$ may produce substantially different results in terms, e.g., of the efficiency of angular-momentum transport (e.g., Bodo et al. 2008; Latter et al. 2009 and references therein); second, as we (and previous works) have shown, turbulence can more easily develop in large-box

¹⁵ This observation agrees with the results presented in the previous section: the magnetic-field amplification depends on ω_C/Ω_0 , which does not change during this scan.




Figure 5. Comparison of 2D pair-plasma simulations of the MRI with variable β_0 in a medium-size box and fixed ω_c/Ω_0 . Left: in-plane magnetic-energy variation normalized over the initial field strength. Right: corresponding evolution of the volume-averaged β in time.




Figure 6. Comparison of 2D pair-plasma simulations of the MRI with variable box aspect ratio. Top: spatial distribution of the in-plane field strength at $t = 4.2P_0$ for a simulation with box size $4 \times 4 \lambda_{\text{MRI}}^2$ (left) and $16 \times 4 \lambda_{\text{MRI}}^2$ (right), with $\omega_c/\Omega_0 = 30$, $\beta_0 = 624$. The elongated-box case shows turbulent structures that do not develop in the corresponding square-box case. Bottom left: evolution of the change in magnetic energy for runs with variable box size and aspect ratio at fixed ω_c/Ω_0 and β_0 . Bottom right: evolution of the volume-averaged $\beta_{\parallel}$ and $\beta_{\perp}$ for the square-box case $4 \times 4 \lambda_{\text{MRI}}^2$ and for the elongated-box case $16 \times 4 \lambda_{\text{MRI}}^2$.

simulations that fit a larger number of channel modes (via large L_z/λ_{MRI} ratios) and that allow for the DKI to contribute to channel breakup and mixing (via large $L_x/(c/\omega_p)$ ratios). However, if enough channel modes fit in L_z , it may not be necessary to further enlarge the simulation box in both x and z ; the development of the DKI may be achieved solely by enlarging the box along x , given that current sheets develop

only along this direction (in 2D). This is attractive because it allows for simulations with large L_x (resulting in developed turbulence) without the need for a correspondingly large L_z . This is also of great importance for 3D simulations where we will demand that $H/L_z = 1$: as discussed in Section 3, this condition implies that the chosen value of L_z/λ_{MRI} will determine the computational cost (in terms of resolution and

time steps). Keeping a reduced L_z while still including all the relevant physics by enlarging L_x (and L_y , in 3D) can greatly reduce the computational effort required.

To assess the possible advantages of aspect ratios $L_x/L_z > 1$, we compare the square-box cases of size $4 \times 4 \lambda_{\text{MRI}}^2$ and $8 \times 8 \lambda_{\text{MRI}}^2$ with simulations where we choose $L_x/L_z = 4$, resulting in boxes of size $16 \times 4 \lambda_{\text{MRI}}^2$ and $32 \times 8 \lambda_{\text{MRI}}^2$. We fix the simulation parameters to $\omega_C/\Omega_0 = 30$, $\beta_0 = 624$, and $v_{A,0}/c = 0.01$. The results are presented in Figure 6. The top panels show a snapshot of the in-plane field strength at $t = 4.2P_0$ (shortly after the beginning of the nonlinear stage) for a square-box simulation of size $4 \times 4 \lambda_{\text{MRI}}^2$ and for an elongated-box run of size $16 \times 4 \lambda_{\text{MRI}}^2$. Turbulent structures are visible in the distribution of B_{pol} when $L_x/L_z > 1$; such structures do not develop in the square-box case, indicating that a limited x -extent inhibits important dynamics leading to the breakup and mixing of channel flows. This difference clearly appears in the bottom-left panel of the same figure, which shows the evolution of the in-plane magnetic energy for all four runs. While the linear MRI stage is well-captured in all cases, the nonlinear stage proceeds substantially differently for a $4 \times 4 \lambda_{\text{MRI}}^2$ box, with the system displaying the small-box behavior discussed in Section 4.3.1. The evolution of the elongated box of size $16 \times 4 \lambda_{\text{MRI}}^2$, however, is much more similar to the larger-box case $8 \times 8 \lambda_{\text{MRI}}^2$, resulting in a saturated, turbulent nonlinear stage. Similarly, the elongated case $32 \times 8 \lambda_{\text{MRI}}^2$ shows a much longer nonlinear stage (which lasts until the end of the run) than the corresponding square-box case.

As a further confirmation of the presence of additional channel breakup and mixing dynamics, in Figure 6 (bottom-right panels) we plot the evolution of the volume-averaged $\beta_{\parallel}$ and $\beta_{\perp}$ for the cases $4 \times 4 \lambda_{\text{MRI}}^2$ and $16 \times 4 \lambda_{\text{MRI}}^2$. In Section 4.2, we argued that the duration and saturation level of the nonlinear stage are heavily influenced by the efficiency of reconnection to limit magnetic-field amplification and pressure anisotropy; here, we observe that the magnetic energy saturates at lower values, and that the pressure anisotropy is indeed better quenched in the elongated-box case, which allows for stronger reconnection and more developed turbulence. The migration of MRI modes toward long wavelengths is therefore slower in elongated boxes. The square-box case, instead, prevents fast DKI modes and very rapidly builds strong magnetic fields and large pressure anisotropy right after the first onset of reconnection. These results prove that simulations in a box of limited vertical extent can still reproduce all the MRI physics of interest (including a saturated turbulent state) simply by increasing the domain's spatial extent in the direction(s) normal to z . For 3D simulations, this will be particularly important when choosing the optimal box geometry for physically meaningful simulations that are also computationally affordable.

4.4. Spectra of 2D MRI-driven Turbulence

In this section, we briefly analyze the power spectrum of the magnetic-field component in the simulation plane (B_{pol}) as well as in the out-of-plane direction (B_y). This analysis provides a clear indication of turbulent activity during the system evolution, and allows us to infer whether MRI-driven turbulence develops at all.

In Figure 7, we plot the isotropic power spectrum for a simulation run in a domain of size $8 \times 8 \lambda_{\text{MRI}}^2$ with $\omega_C/\Omega_0 = 60$, $\beta_0 = 624$, and $v_{A,0}/c = 0.01$. We choose this




Figure 7. Isotropic power spectrum of the in-plane (red lines) and out-of-plane (blue lines) magnetic field for a 2D pair-plasma MRI simulation with parameters $\omega_C/\Omega_0 = 60$, $\beta_0 = 624$, and box size $8 \times 8 \lambda_{\text{MRI}}^2$. The spectra during the linear stage (dashed lines) are peaked around the most-unstable MRI wavenumber k_{MRI} . At late times during the nonlinear stage (solid lines), the spectra feature characteristic slopes indicating an inertial range, and a clear spectral break in the vicinity of the average Larmor-radius wavenumber k_p .

particular set of parameters, among the various cases explored in the previous sections, as it represents a good compromise between box size and separation between macroscopic and microscopic scales. In the plot, the dashed colored lines represent the spectrum of B_{pol} (red) and B_y (blue) measured during the linear MRI phase. Both spectra are strongly peaked around the most-unstable wavenumber k_{MRI} (marked by a vertical dotted line), indicating active MRI growth at the fastest rate. A second, less pronounced peak is visible at this time in both spectra around $kc/\omega_p \simeq 0.1$; this is compatible with the most-unstable wavenumber of mirror modes, suggesting that mirror activity is present at this time in the system evolution (as also discussed in Section 4.2 and as noted by Inchingolo et al. 2018). We repeat the measurement of magnetic-field spectra at a late time $t \simeq 10P_0$ (solid lines), when the system has entered the nonlinear stage. These spectra feature characteristic slopes that identify an inertial range: at length scales larger than the typical Larmor radius, i.e., $k < k_p \equiv 1/\rho_C$, the B_{pol} spectrum has a shallow slope that could be comparable with $k^{-5/3}$, which hints at a cascade process akin to the typical strong-turbulence MHD picture (e.g., Goldreich & Sridhar 1995). At $k \simeq k_p/2$, this spectrum changes to a slope close to k^{-3} or slightly steeper, indicating the transition from Alfvénic turbulence to a kinetic range. The B_y spectrum is instead characterized by this $\propto k^{-3}$ slope (or steeper) at all scales, in agreement with previous 2D studies (Inchingolo et al. 2018). Power in both spectra falls rapidly at large k , as the cascade reaches sub-Larmor scales, where strong magnetic energy dissipation is expected to occur. We have verified that numerical dissipation at the grid scale is minimal; we elaborate upon this in Appendix C.

The measured spectra suggest that the system has indeed reached a state of well-developed turbulence during the nonlinear stage. In general, we measure the same spectral slopes in the majority of the simulations explored in our parameter-space scan. In a few cases, however, the results deviate from this expectation: in particular, we find that (square) system sizes $\leq 4 \times 4 \lambda_{\text{MRI}}^2$ and scale separations $\omega_C/\Omega_0 \leq 15$ do not produce spectra with such clear power-

law behavior, hinting at a turbulent cascade that is not fully developed. In the previous sections, we have argued that such cases lack the action of certain physical mechanisms, such as drift-kink instabilities, that promote the transition to a turbulent state. Our analysis supports this statement, as well as our claim from Section 4.3.4 that an elongated simulation box facilitates the development of turbulence by giving the additional instabilities sufficient room to grow: indeed, we find that, e.g., the case with box size $16 \times 4\lambda_{\text{MRI}}^2$ produces magnetic-field spectra with characteristic turbulent slopes, while the $4 \times 4\lambda_{\text{MRI}}^2$ case does not.

Even though our results are in good agreement with the only other (to date) spectral analysis of the 2D, fully kinetic MRI presented in literature (Inchingolo et al. 2018), a few additional observations are worth reporting here. First, the spectra we show in Figure 7 are not necessarily representative of the whole nonlinear stage. Measuring the spectra at other times reveals additional features, e.g., a progressive shift of the spectral break in B_{pol} toward smaller wavenumbers. This is expected in numerical simulations of turbulence due to the continuous plasma heating that pushes ρ_C closer to the macroscopic scales (Zhdankin et al. 2018a; see Section 4.5), but the scale separation we include in our runs is rather limited from the start, which allows for tracking this phenomenon only over relatively short timescales. Additionally, as we have illustrated in Section 4.2, our 2D simulations are characterized by the absence of reconnection along y , which probably alters the development of the B_y spectrum with respect to the 3D case. This issue, combined with the fact that the MRI continuously slows down (and then stops; see Section 4.1) in our simulations, produces late-time power spectra that are no longer representative of MRI turbulence and differ substantially from those shown here. Finally, we note that, over long times, we observe a progressive pile-up of energy at large k values (not shown in Figure 7), which we attribute to numerical noise accumulating on short wavelengths; this complicates the long-term spectral analysis.

Our results also differ from the typical spectra seen in 3D MHD and hybrid simulations. For example, Kunz et al. (2016) and Walker et al. (2016) measured a $\propto k^{-2}$ spectrum for B_y in the inertial range, which they attribute to sharp reversals of B_y over short length scales. It is entirely possible that the difference between our 2D results and 3D MHD and hybrid results originates from our fully kinetic, pair-plasma setup. Moreover, the simulations conducted in Kunz et al. (2016) also differ in the background magnetic-field configuration, which has zero net flux. Our initial setup includes a nonzero net-flux magnetic field, which changes the MRI evolution drastically (see Section 4.7). This may produce additional differences in the turbulent spectra. In Section 5.2, we will discuss spectra for our 3D simulations.

4.5. Particle Energization in 2D

We now turn our attention to the question of particle energization in the 2D collisionless MRI. With our PIC approach, we can measure nonthermal effects, which are not included in MHD models. Our pair-plasma simulations are not representative of the electron-ion plasma expected to make up typical RIAFs, but we can still infer qualitative features in the mechanisms driving particle energization during the MRI development that could potentially carry over to the case of large m_i/m_e .

Prior to conducting any analysis, we emphasize a few key points raised in earlier sections. First, it is now clear that, in all of our 2D simulations (regardless of system size and parameters chosen), the MRI eventually stops and the system reaches a quiet steady state. We should bear in mind that such a “quiet” state consisting of magnetic loops is not physically realistic, and is a product of the 2D geometry. Second, in the ideal situation where the MRI has evolved into a fully turbulent state, we do not expect the particle energy to saturate in time at a well-defined value. This turbulent state can be regarded as analogous to driven-turbulence simulations (e.g., Zhdankin et al. 2018b) where energy is constantly injected into the system. In our case, the injection mechanism is the MRI on large scales, and in addition, heating and dynamo processes self-consistently regulate the energy injection in our simulations. Third, the rate of particle energy gain will obviously depend on the physical conditions under which the MRI develops. We have already shown that the system evolution is drastically different for small/large boxes and for small/large ω_C/Ω_0 (see Section 4.1). We will need to consider such differences to attribute the measured energization to the correct mechanism. Finally, we remark that, due to the fundamentally different physical processes at play (e.g., lack of reconnection along y in 2D), we expect the results of 3D simulations to potentially differ from those presented here. These differences will be discussed in Section 5.3.

Figure 8 shows the evolution of the average Lorentz factor (top row) in three simulations that are representative of the parameter-space scan presented in Section 4.3: the small-box/large-frequency ratio case $2 \times 2\lambda_{\text{MRI}}^2$, $\omega_C/\Omega_0 = 240$ (left column); the medium-box/medium-frequency ratio case $8 \times 8\lambda_{\text{MRI}}^2$, $\omega_C/\Omega_0 = 60$ (central column); and the large-box/small-frequency ratio¹⁶ case $32 \times 32\lambda_{\text{MRI}}^2$, $\omega_C/\Omega_0 = 30$ (right column). In all cases, $\beta_0 = 624$ and $v_{A,0}/c = 0.01$. In each case, the energy evolution is shown until shortly after the moment when the MRI wavelength grows beyond the box size. This moment depends on the simulation parameters, and it differs in each of the three cases. For each simulation, we show the evolution of electron and positron energies separately, as well as their average and difference.

The first feature we observe is a differential energization between electrons and positrons, which becomes substantial at late times. Although surprising (compared to standard PIC simulations), this is a genuine physical phenomenon in the shearing-box paradigm. Our setup involves a simple (and by itself, symmetric) electromagnetic configuration, i.e., a uniform vertical field; but differently from standard PIC simulations, the addition of a unidirectional (along y) background rotational profile results in a symmetry-breaking effect. Charged particles in this situation ought to behave differently depending on the sign of their charge, which determines how their cyclotron gyromotion combines with the background rotation. Analytically, we can observe that the shearing-box equations include extra forces acting on particles in addition to the usual Lorentz force. These forces cause drifts in the particle motion, which depend on the sign of the charge and of the mean magnetic field. We have found that this additional dynamics plays a fundamental role in the energization process, with electrons consistently gaining more energy than positrons (by a factor ~ 3) in our

¹⁶ Note that the “small” ω_C/Ω_0 employed here is still approximately twice as large as the largest value employed in previous works on the pair-plasma MRI.




Figure 8. Particle energization in 2D pair-plasma MRI simulations. Top row: average Lorentz factor for electrons and positrons for simulations with a small box/large ω_C/Ω_0 (left), medium box/medium ω_C/Ω_0 (center), large box/small ω_C/Ω_0 (right). In each case, the difference between $\langle \gamma \rangle$ of the two species is shown as a dashed green line, and the end of the nonlinear stage is marked with a vertical dashed–dotted line (notice the difference in timescales between different runs). Bottom row: evolution of energy distributions in time (including both species) for the three runs. In each case, the distributions are shown until the same final time as in the top row. The dotted lines indicate the average and cutoff Lorentz factors $\gamma \sim qB\lambda_{\text{MRI}}/(mc^2)$ (with $B = B_c$ or $B = B_c$; see text) characterizing the quiescent end state of each run.

largest simulations, for our specific choice of an initial $\mathbf{B}$ aligned with the object’s axis of rotation. This phenomenon is exactly reversed when reversing the orientation of the vertical field: we have indeed verified that simulations in which $B_{z,0}$ is antiparallel to the rotation axis result in positrons gaining more energy than electrons. The differential-heating mechanism reported here has not been discussed in previous works, but it is in fact a crucial point to consider in future electron–ion simulations, where the heating ratio between species is a primary analysis target, because it relates directly to observations. In the following, we will focus on analyzing particle energization on average (including both species); we discuss the causes and the importance of this differential heating in detail in Appendix B.

Comparing the three cases shown in Figure 8, we observe that the average γ at the end of each run consistently reaches values $\gtrsim 10^2$; however, the energization proceeds very differently for different simulation parameters, as is evident from the slope of each curve over time. In the small-box case (left column in Figure 8), the heating rate is evidently much larger during brief, “bursty” periods that correspond to large-scale reconnection events, starting at $t \simeq 3.5P_0$, which marks the onset of reconnection after the linear stage. Each event is followed by a “quiet” phase where particles gain energy steadily at a much slower rate. During each quiet phase, the energization rate is very similar, and the transition to and from reconnection is visible. Conversely, the large-box case (right column) shows a smooth increase in the kinetic energy, without the energization “bursts” observed in the small-box case. The heating rate here progressively decreases as time passes; the

energization proceeds much more slowly (notice the difference in timescales between different cases) and around $t \simeq 30P_0$, when the MRI is slowing down and halting, the average Lorentz factor reaches a plateau. The medium-box case (central column) shows features that appear as a combination of the other two cases: the energy increases very rapidly at a few moments in time, with smooth phases of slower heating in between.

Regardless of the specific processes contributing to particle energization, analyzing the energy distribution functions in each case reveals several similarities. Figure 8 (bottom row) shows the energy distributions (including both species) at several different times, until the moment the MRI is halted in each of the three cases (again, notice the difference in timescales). In all cases, at the onset of magnetic reconnection (marking the transition to the nonlinear stage), the distribution abruptly jumps to much larger average energies and a clear nonthermal component arises. The nonthermal population appears to follow a power-law distribution with index ~ 2.5 – 3 (both slopes are shown in the plot to guide the eye), but this value changes substantially as time passes. As the system evolves, the spectra become shallower and the main part of the distribution resembles a Maxwellian; in addition, in all cases, we observe the progressive formation of a second peak in the distribution at high energies, and a stabilization of the cutoff γ at a value that grows larger with larger ω_C/Ω_0 .

That the energy distributions continue to evolve as long as the MRI is active is unsurprising, given that the system includes a free energy source. However, the shift from a power-law distribution to a double-peaked shape is peculiar.

Measuring the peak and cutoff γ across our large parameter space reveals a common feature: these values correspond almost exactly to the Lorentz factors of particles with gyroradii of size $\sim \lambda_{\text{MRI}}$, when measuring the magnetic-field strength at the center and at the edge (which we indicate with B_c and B_e , respectively) of the large-scale, stable magnetic loops that form at the end of the system evolution (the corresponding values of γ are indicated in Figure 8 by black dotted lines). These loops originate from plasmoids created via reconnection, well before the end of the nonlinear stage when the MRI starts to slow down. Particles are gathered inside the loops as these continue to merge; when the MRI is completely halted, these loops have grown to size $\sim \lambda_{\text{MRI}}$ and contain almost all particles inside. The energy distribution evolves to reflect this state: particles trapped inside the loops, where $B \sim B_c$, tend to increase their energy such that $mc^2\gamma/(qB_c) \sim \lambda_{\text{MRI}}$ (e.g., Hillas 1984), which creates the peaked distribution we observe. The maximum particle energy will then be limited by the strong magnetic field at the loop edge such that $mc^2\gamma/(qB_e) \sim \lambda_{\text{MRI}}$, with $B_e \sim 10B_c$. This creates the energy cutoff that we measure in our runs.

As discussed in earlier sections, the main difference between the MRI evolution in the three cases is the presence of sustained turbulence in the nonlinear stage. In the small-box case, turbulence cannot develop and the energization is mediated by intermittent reconnection events, separated by “quiet” periods where magnetic fields are amplified in coherent channel structures and particles can increase their perpendicular momentum adiabatically via μ conservation. In large boxes, there is no clear distinction between subsequent (discrete) reconnection events—turbulence results from the breakup and mixing of channel flows, and more chaotic processes come into play. To verify quantitatively that turbulent heating is the main driver of particle energization in the large-box case, we follow an approach similar to that described by Zhdankin et al. (2018a). Assuming Alfvénic turbulence, the energy injection from electromagnetic fields to particles can be written as

$$\frac{dK_{\text{int}}}{dt} \sim \frac{\delta B_{\text{RMS}}^2 v_A}{8\pi \ell_{\text{inj}}}, \quad (62)$$

where $K_{\text{int}} \equiv 2nmc^2(\langle\gamma\rangle - 1)$ (counting both particle species) is the average internal (kinetic) energy and δB_{RMS} , v_A , and ℓ_{inj} are the characteristic root-mean-square magnetic energy, the corresponding Alfvén speed, and the injection length scale of the turbulent cascade. In our case, the relevant magnetic field and Alfvén speed are those corresponding to the late-time, saturated state achieved in the xz -plane, because the out-of-plane B_y only acts as a guide field and does not participate in the cascade process (due to the lack of reconnection along y). For the injection scale, we select $\ell_{\text{inj}} = \lambda_{\text{MRI}}$: during the turbulent nonlinear stage, in all cases, we measure the largest magnetic-field structures to be of size comparable to the *initial* MRI wavelength (despite the system evolution pushing the MRI to much larger scales); this is an indication that the system retains some memory of the initial conditions. In Figure 9, we show this measurement of energy injection (integrated in time) for the three simulations presented above. We observe that the large-box case reaches, at late times, a saturated state that agrees very well with our simple estimate. This appears to indicate that the turbulent-cascade process at play is well-described by Alfvénic turbulence. For the small- and medium-




Figure 9. Estimate for injection of magnetic energy and conversion into kinetic energy assuming Alfvénic turbulence, for the three simulations presented in Section 4.5. In each case, the end of the nonlinear phase is marked with a black cross.

box cases, the physics involved in the nonlinear stage is substantially different, and our estimate fails to accurately describe the energy-injection process. In 3D, we expect the physics of this injection to include additional dynamics along y , which allows B_y to participate in the cascade process. The absence of this mechanism in 2D is a potential source of discrepancy between 2D and 3D results.

4.6. Stresses and Angular-momentum Transport in 2D

In this section, we analyze the effective viscosity that originates from the MRI nonlinear dynamics in 2D pair-plasma simulations. We define an effective-viscosity parameter (Shakura & Sunyaev 1973),

$$\alpha \equiv \frac{\langle R_{xy} + M_{xy} + A_{xy} \rangle}{\langle p \rangle}, \quad (63)$$

where $\langle \dots \rangle$ denotes a volume average. Here, $R_{xy} = mnU_x U_y / \Gamma^2$ is the Reynolds stress, with $\Gamma = \sqrt{1 + U^2/c^2}$ and $\mathbf{U} = (U_x, U_y, U_z)$ the spatial part of the bulk four-velocity of the plasma in the frame moving with velocity $\mathbf{v}_s$; $M_{xy} = -B_x B_y / (4\pi)$ is the Maxwell stress; and $A_{xy} = -(p_\perp - p_\parallel) B_x B_y / B^2$ is the anisotropic stress (e.g., Sharma et al. 2006). Our PIC simulations, where the plasma can be relativistically hot, require special care when measuring α and its components. Given a distribution function $f(\mathbf{x}, \mathbf{u}, t)$, we first compute the bulk velocity $\mathbf{U} = (1/n) \int \mathbf{u} f(\mathbf{x}, \mathbf{u}, t) d^3\mathbf{u}$ for each particle species separately. Then, we apply a Lorentz boost to obtain $\tilde{\mathbf{u}}$, the spatial part of the four-velocity of each particle in the frame moving with $\mathbf{U}$. We subsequently define a pressure tensor for each particle species as

$$\mathbf{P} = m \int \frac{\tilde{\mathbf{u}} \tilde{\mathbf{u}}}{\tilde{\gamma}} f(\mathbf{x}, \tilde{\mathbf{u}}, t) d^3\tilde{\mathbf{u}}, \quad (64)$$

where $\tilde{\gamma} = \sqrt{1 + \tilde{u}^2/c^2}$. The isotropic, parallel, and perpendicular pressure are then $p = \text{tr}(\mathbf{P})/3$, $p_\parallel = \mathbf{P} : \hat{\mathbf{b}}\hat{\mathbf{b}}$, and




Figure 10. Left column: evolution of the α -parameter and of the volume-averaged Maxwell, Reynolds, and anisotropic stresses for a small box/large ω_c/Ω_0 (top), medium box/medium ω_c/Ω_0 (center), and large box/small ω_c/Ω_0 (bottom) simulation of the 2D pair-plasma MRI. Right column: evolution of the volume-averaged pressure anisotropy and $\beta_\parallel, \beta_\perp$ for the same simulations.

$p_\perp = (1/2)\mathbf{P}:(\mathbb{I} - \hat{\mathbf{b}}\hat{\mathbf{b}})$, respectively, where $\hat{\mathbf{b}} = \tilde{\mathbf{B}}/\tilde{B}$ and $\tilde{\mathbf{B}}$ is the magnetic field in the frame moving with $\mathbf{U}$. R_{xy} and A_{xy} are computed for each species separately, and then summed across species.

In Figure 10, we show the time evolution of the average α and its components (left column) as well as the volume-averaged $\langle\beta_\parallel\rangle$, $\langle\beta_\perp\rangle$, and the pressure anisotropy $\langle(p_\perp - p_\parallel)/p\rangle$ (right column) for the small-box/high-frequency ratio case $2 \times 2\lambda_{\text{MRI}}^2$, $\omega_c/\Omega_0 = 240$ (top row), the medium-box/medium-frequency ratio $8 \times 8\lambda_{\text{MRI}}^2$, $\omega_c/\Omega_0 = 60$ (middle row), and the large-box/small-frequency ratio $32 \times 32\lambda_{\text{MRI}}^2$, $\omega_c/\Omega_0 = 30$ (bottom row). In all cases, $\beta_0 = 624$ and $v_{A,0}/c = 0.01$. All simulations are characterized by a rapid increase in all components of α during the linear stage, with a peak at $\alpha \simeq 1$ right before the onset of reconnection. At this point, different parameters determine strong differences in the results: during the nonlinear stage, we observe that the volume-averaged M_{xy} can grow and remain close to unity at early times in the small- and medium-box cases; conversely, in the large-box case, M_{xy} rapidly settles to ~ 0.1 after the onset of reconnection and slowly decreases to ~ 0.01 . In all cases, the anisotropic stress A_{xy} achieves the highest (volume-averaged) value ~ 0.1 at the onset of reconnection, and then firmly settles around 0.02–0.05 for the whole duration of the nonlinear phase. The Reynolds stress is consistently the least important component of α and remains much smaller than the other two during the nonlinear phase. Overall, the α -parameter closely follows the evolution of M_{xy} in the small- and medium-box cases, attaining values 0.1–1; in the

large-box case, $\alpha \simeq 0.01$ –0.1 and the contribution of A_{xy} is more significant, providing most of the angular-momentum transport in the late nonlinear stage when the magnetic field is dissipating via reconnection and turbulence. We note, however, that in the large-box case, ω_c/Ω_0 is substantially smaller than in the small-box case (respectively $\omega_c/\Omega_0 = 30$ and $\omega_c/\Omega_0 = 240$); because the scale separation affects the development of mirror modes (see Section 4.2), pressure anisotropy may be exaggerated in the small frequency-ratio case. Larger values of ω_c/Ω_0 would be needed to verify the trend of the anisotropic stress in the large-box case, but this involves computational costs beyond our current possibilities. We will investigate the matter more in detail in future work.

These results align well with the discussion provided in the previous sections. The MRI in the large-box case behaves qualitatively differently from the small-box case: the system develops strong turbulence during the nonlinear stage, promoting the continuous, more efficient dissipation of magnetic fields and faster pressure isotropization. In small boxes, this dissipation is primarily mediated by episodic, large-scale reconnection events, with phases of strong magnetic-field amplification driving pressure anisotropy in between. This suggests that, in a large box, M_{xy} and A_{xy} will be reduced in comparison with a small-box case. The right-hand panels of Figure 10 confirm this reasoning: pressure anisotropy is much more pronounced at the beginning of the nonlinear stage in smaller boxes, while it remains rather limited in the large-box case.




Figure 11. Evolution of the in-plane magnetic-energy variation for several 2D pair-plasma simulations of the MRI with zero net magnetic flux. Left: simulations with varying box size and fixed ω_C/Ω_0 , β_0 . Right: simulations with varying ω_C/Ω_0 and fixed box size and β_0 .

If α is invoked as a coefficient of turbulent angular-momentum transport, then the large-box case where turbulence can develop is more relevant to provide a first-principles measure of the effective turbulent viscosity in a collisionless plasma. In this case, our 2D numerical experiments indicate that the Maxwell stress provides the primary contribution to α , with the anisotropic stress achieving smaller values $A_{xy}/M_{xy} \simeq 0.2$. This result differs from the findings of Kunz et al. (2016), who reported $A_{xy}/M_{xy} \simeq 1$; however, we stress again that 2D runs have the intrinsic limitation of lacking reconnection processes in the dominant magnetic-field component B_y . In 3D, we may then expect large differences with the results presented here: dissipating B_y via reconnection may imply that the Maxwell stress could achieve much smaller values than those measured in 2D, possibly increasing the relative contribution of pressure anisotropy in determining α . On the other hand, in the previous sections we have argued that reconnection along y may help reduce pressure anisotropy by increasing $p_{\parallel}$ and hence decreasing A_{xy} . We will discuss the dynamics of angular-momentum transport in 3D in Section 5.4.

4.7. Net Flux versus Zero Net Flux

We conclude our investigation of the 2D collisionless-MRI dynamics by briefly considering the case of zero net flux. In contrast with the simulations presented in the previous sections, here the initial magnetic field is set to $\mathbf{B}(t=0) = -B_{z,0} \sin(2\pi x/L_x) \hat{e}_z$ such that no net flux is present along z . This case was briefly discussed in Riquelme et al. (2012) in 2D and was the focus of the only (to date) 3D hybrid study of the MRI (Kunz et al. 2016).

In Figure 11, we plot the evolution of the in-plane magnetic energy for a few simulations with exactly the same parameters employed in Section 4.3 (aside from the initial magnetic-field configuration). In particular, we compare B_{pol} in several cases: first, by keeping a fixed $\omega_C/\Omega_0 = 30$, $\beta_0 = 624$, and $v_{A,0}/c = 0.01$ (the latter two measured with the maximum value of B), in larger and larger simulation boxes $L_z/\lambda_{\text{MRI}} = L_x/\lambda_{\text{MRI}} = 4, 8, 16$, and 32 (left panel); then, by fixing the box size $L_z/\lambda_{\text{MRI}} = L_x/\lambda_{\text{MRI}} = 8$, $\beta_0 = 624$, and $v_{A,0}/c = 0.01$, and increasing the frequency ratio $\omega_C/\Omega_0 = 6$,

15, 30, and 60 (right panel). From this comparison, we observe that the general statements made for the nonzero net-flux case hold here as well: the duration of the turbulent, saturated nonlinear stage that follows the linear growth of the MRI increases both with box size and with ω_C/Ω_0 . Over time, the turbulence rapidly decays as magnetic fields grow and pressure anisotropy develops (not shown), pushing the MRI wavelength to scales larger than the box size. However, the zero net-flux case is also characterized by striking differences, with the most prominent being a much more efficient breakup and mixing of channel flows. This was also observed by Riquelme et al. (2012), who attributed this behavior to the fact that the instability develops nonuniformly in different regions of the domain; in particular, the most-unstable MRI wavelength $\lambda_{\text{MRI}} = 2\pi v_A/\Omega_0$ is shorter where B_z is weaker and longer where B_z is stronger, producing a spatial inhomogeneity in the dominant instability wavevector. Channels forming inside regions of weaker/stronger B_z are therefore spatially nonuniform, and tend to be destroyed much more quickly. Hence, differently from the net-flux case, the breakup of channels formed during the linear stage is not solely mediated by the onset of tearing modes (and concurrently of kink instabilities): the dynamics does not need to “wait” for the magnetic field in the channels to grow strong enough (and for current sheets to thin out enough) that reconnection can occur. Instead, the intrinsic inhomogeneity of the background field mediates the mixing and redistribution of magnetic energy. As a consequence, the saturation level reached in all simulations with zero net flux is much smaller than in the corresponding nonzero flux cases.

Figure 11 also shows important differences in how these simulations tend to converge: while it appears that progressively better convergence is achieved for increasing ω_C/Ω_0 at a fixed box size, enlarging the box size does not result in converging results. Larger boxes seemingly produce a steady increase in the magnetic-field amplification, in contrast with the clear convergence observed for the net-flux case. We also note that, because the magnetic field here reverses its direction through the box, we observe no sign of the differential heating of electrons versus positrons that we measured in the net-flux case (see Section 4.5). Charges with opposite sign are indeed

Table 2
List of 3D Simulations

Box Size ($L_x \times L_y \times L_z$)	ω_c/Ω_0	β_0	$\lambda_{\text{MRI}}/\rho_C$
$2 \times 2 \times 2\lambda_{\text{MRI}}^3$	15	312	7.46
$4 \times 4 \times 2\lambda_{\text{MRI}}^3$	15	312	7.46
$4 \times 8 \times 2\lambda_{\text{MRI}}^3$	15	312	7.46

Note. In all cases, $v_{A,0}/c \simeq 0.007$; the numerical resolution is such that $\Delta x \simeq c/\omega_p$, and we employ 27 particles per cell per species.

equally likely to experience differential gyromotion and drift toward current sheets, overall resulting in an almost perfectly balanced energization among electrons and positrons.

In this work, we focus on nonzero net-flux simulations primarily because, in any astrophysical accretion disk, the complete absence of vertical magnetic flux is not expected; realistically, any portion of the disk will be threaded by some magnetic field carrying a net flux, however small. Additionally, the behavior of zero net-flux simulations is fundamentally different from the expected development of a “pure” MRI: the former situation is more akin to a dynamo problem, and does not feature the same qualitative evolution of the nonzero net-flux case. Our results show that the case of zero net flux differs enough from the case with nonzero magnetic flux to deserve a thorough, dedicated exploration, which we will pursue in future work.

5. Three-dimensional Simulations of the MRI in Collisionless Pair Plasmas

In this section, we present 3D PIC simulations conducted with ZELTRON, where we have implemented the newly developed KSB-OA model (see Section 2.3). In general, we expect 3D runs to produce results more physically realistic (at least in terms of the magnetic-field evolution and subsequent development of turbulence) than corresponding 2D simulations, for the reasons discussed in Section 4.2. We will therefore select simulation parameters that strictly enforce three physical constraints: (i) a sufficiently subluminal velocity offset at the x -boundaries, i.e., $v_s(L_x) = s\Omega_0 L_x < c$; (ii) a vertical extent of the simulation box that does not exceed the disk pressure scale height, i.e., $L_z/H \leq 1$; and (iii) a sufficiently large separation between macroscopic and microscopic scales (at least in the beginning of each run), in terms of $\lambda_{\text{MRI}}/\rho_C$.

Due to the much higher computational cost of 3D runs, in this first study we will explore a limited parameter range inspired by the results of 2D simulations. We will show that 3D runs of relatively small size already provide fundamental new insight into the physics of the collisionless MRI. The 3D simulations discussed in this work are summarized in Table 2; all runs are initialized analogously to the 2D case, with a weak vertical magnetic field determined from the initial simulation parameters. We always employ 27 particles per cell per species, and a numerical grid such that $\Delta x \simeq c/\omega_p$. In future work, we will conduct larger simulations with better scale separation to achieve more realistic regimes.

In the following sections, we will first describe the physics of small- and large-box 3D PIC simulations of the pair-plasma MRI, and how these relate to corresponding 2D simulations. We will then analyze the 3D turbulence, particle energization, and angular-momentum transport that characterize the 3D case.

5.1. Physics of the 3D Collisionless MRI

Figures 12 and 13 show the results for two representative 3D simulations of the collisionless MRI in pair plasmas, employing a small ($2 \times 2 \times 2\lambda_{\text{MRI}}^3$) and a large ($4 \times 8 \times 2\lambda_{\text{MRI}}^3$) box size, respectively. In these figures, we show subsequent snapshots of the spatial distribution of B_x and B_y (left panels), of the evolution of the change in magnetic energy (top-right panel), and of the volume-averaged parallel and perpendicular plasma- β (bottom-right panel). These simulations are initialized with parameters such that $L_z = H$, $\lambda_{\text{MRI}}/\rho_C \simeq 7.5$, and $v_s(L_x)/c \simeq 0.125$ (in the small-box case) and $v_s(L_x)/c \simeq 0.25$ (in the large-box case). In both cases, imposing these conditions results in a frequency ratio $\omega_c/\Omega_0 = 15$ and a plasma temperature $\theta_0 = 1/128$ such that $\beta_0 = 312$ (including both species), with an initial Alfvén speed $v_{A,0}/c \simeq 0.007$.

For the small-box case ($2 \times 2 \times 2\lambda_{\text{MRI}}^3$) shown in Figure 12, we observe that the evolution of the magnetic energy (top-right panel) is initially qualitatively similar to those of 2D simulations of comparable L_x , L_z (see the 2D case of Figure 2): from the initial state, the field strength grows exponentially until saturation is reached. The nonlinear stage, however, shows qualitative differences: we observe that the magnetic energy in all components (dominated by B_y , which does not saturate in 2D) remains approximately around the saturation level for ~ 2 orbital periods after the end of the linear stage, before starting to slowly decrease. This behavior is due to the 3D geometry and is illustrated by subsequent snapshots of the B_x and B_y distribution in the domain, shown in Figure 12 (left panels). Around $t \simeq 3.7P_0$ (top-left panels) we observe that channel flows form, during the linear stage, in the xz -plane over length scales of λ_{MRI} , similarly to the 2D dynamics. However, in 3D, these channels also span the entire y -extent of the simulation domain. At the channel interfaces, reconnection starts dissipating magnetic energy when current sheets become sufficiently thin and tearing modes are excited (around $t \simeq 4.2P_0$, center-left panels). In 3D, reconnection can take place also along y , thus dissipating energy stored in B_y (and not just in B_x and B_z as in 2D). The channel interfaces are indeed heavily perturbed, and magnetic reconnection can act along all spatial directions (e.g., around $t \simeq 5.4P_0$, bottom-left panels). The saturated energy level of the nonlinear stage we measure between $t \simeq 4P_0$ and $t \simeq 6P_0$ is maintained due to the balance achieved between injection (via the MRI) and dissipation (via reconnection) of energy into the system.

Starting from $t \simeq 6P_0$, the system slowly relaxes toward an end state where the MRI slows down and then stops. This final state (not shown in the left panels of Figure 12) is a two-channel configuration with lower B_y energy. From $t \simeq 8P_0$ onward, we observe that the energy in B_x and B_z oscillates around approximately the same values; the spatial distribution of the magnetic field tends to assume a “shear-wave” configuration with coherent, large-scale structures aligned with the shear flow appearing cyclically. No reconnection occurs at this point; the system appears to maintain this state endlessly. The dynamics that characterizes this state is different from that of the early nonlinear stage, and it is doubtful whether it can even be associated with the MRI at all. Qualitatively, it appears that MRI modes have migrated to wavelengths longer than the box size, and that the 3D geometry is allowing a quasi-steady state that consists of periodic oscillations. This dynamics is of no relevance for this work, but it is interesting to note the




Figure 12. Evolution of the pair-plasma MRI in a small 3D box ($2 \times 2 \times 2\lambda_{\text{MRI}}^3$). Left: spatial distribution of the radial (B_x) and toroidal (B_y) magnetic field at subsequent times during the simulation. Right: evolution of the change in magnetic energy in all three components of $\mathbf{B}$ (top), and evolution of the volume-averaged $\beta_{\parallel}$ and $\beta_{\perp}$ (bottom).

difference between the 3D and 2D end states, with the latter consisting of stable loops in the xz -plane.

The evolution of the volume-averaged $\beta_{\parallel}$ and $\beta_{\perp}$, shown in Figure 12 (bottom right), closely follows that of the magnetic field: in the early nonlinear stage (between $t \simeq 4P_0$ and $t \simeq 6P_0$), we measure an increase $\langle \beta_{\parallel} \rangle \simeq \langle \beta_{\perp} \rangle \sim 0.2$ –10, with very brief periods where $\langle \beta_{\perp} \rangle > \langle \beta_{\parallel} \rangle$ and subsequent “bursty” reconnection events re-establishing isotropy periodically. At late times, pressure anisotropy accumulates and pushes the MRI growth rate to smaller and smaller values.

The small 3D run discussed above achieves magnetic-energy saturation in the nonlinear stage, thus already improving over 2D runs of much larger size (in x and z). However, in such a small simulation, we do not obtain a developed turbulent state, consistently with the results of the sole 3D fully kinetic study previously presented (Hoshino 2015). For this reason, we now focus on the larger simulation shown in Figure 13 with box size $4 \times 8 \times 2\lambda_{\text{MRI}}^3$. This case is chosen as representative of a situation where the numerical domain size is rather limited in z (in our case, by the requirement $L_z \leq H$)—hence a small number of channel modes are expected to grow during the linear stage—but the larger extent in x and y now allows for additional instability mechanisms that cannot be included in the previous $2 \times 2 \times 2\lambda_{\text{MRI}}^3$ case. For 2D simulations, in Section 4.3.4 we have argued that boxes elongated in x promote the development of turbulence by allowing for drift-kink modes that can disrupt the channels; here, we observe exactly this effect in 3D. Elongated channels are here susceptible to more numerous DKI modes than in the small-box case. These long-wavelength modes also cause ripples in

the current sheets of larger amplitude, up to the point where current sheets at the interfaces of different channels interact with each other. This eventually causes a loss of coherence in the channel structures and the complete disruption of the sheets. This is not observed in the small-box case, where ripples along the current sheets are of much smaller amplitude and do not cause the breakup of channels.

The top-right panel of Figure 13 shows that the system evolution now proceeds through a linear stage where the magnetic field is amplified roughly 10 times less than in the previous cubic-box case. The onset of magnetic reconnection results in a nonlinear stage that is qualitatively different from that of the small-box case: repeated cycles of period $\sim 2P_0$ are observed where the magnetic energy in all components grows and is rapidly dissipated before growing again. Each growth phase corresponds to the formation of macroscopic channels, and each dissipation phase results in the complete breakup of said channels, which lose coherence and decay into “turbulent” structures (whether this can truly be called a turbulent state will be discussed in Section 5.2). This behavior is visible in the spatial distribution of B_x and B_y at subsequent times in Figure 13 (left panels).

Reconnection in this elongated-box simulation is evidently much more active than in the cubic-box case. We observe that the macroscopic current sheets violently kink in both x and y , eventually resulting in the emergence of the turbulent state. The bottom-right panel of Figure 13 shows that the stronger reconnection dynamics also heavily impacts the evolution of β : because of the smaller magnetic-field amplification (with respect to the cubic-box case), the nonlinear stage is




Figure 13. As in Figure 12, but for a box size of $4 \times 8 \times 2\lambda_{\text{MRI}}^3$.

characterized by a volume-averaged $\langle\beta\rangle \sim 10$, which can grow up to ~ 100 during the dissipation phase of the channel-breakup cycles. Pressure anisotropy is well-regulated, during the early nonlinear stage, by subsequent reconnection events, which periodically re-establish isotropy; in the late nonlinear stage, a finite (tens of percent), unquenched pressure anisotropy is retained, marking the transition to the same end state described above for the small-box simulation.

The numerical experiment reported in Figure 13, however limited in box size and scale separation, provides the first, striking evidence that large 3D PIC simulations of the collisionless MRI possess qualitative features analogous to corresponding MHD simulations. In contrast to the 2D case, where important physical processes are impeded by the reduced dimensionality, here we observe an MRI behavior akin to that exhibited by 3D MHD runs, with the cyclic creation (via the MRI) and destruction (via reconnection) of channel flows. In the next subsections, we discuss more quantitative comparisons with 2D simulations as well as the effect of choosing different aspect ratios in 3D runs.

5.1.1. 2D versus 3D MRI Dynamics

Figure 14 shows the evolution of the volume-averaged Alfvén speed (left panel), calculated with the nonrelativistic expression, and of the volume-averaged pressure anisotropy ($p_{\perp} - p_{\parallel}$)/ p (right panel) for the two runs presented in the previous section (with box size $2 \times 2 \times 2\lambda_{\text{MRI}}^3$ and $4 \times 8 \times 2\lambda_{\text{MRI}}^3$, respectively), as well as for an analogous run with box size $4 \times 4 \times 2\lambda_{\text{MRI}}^3$. In the same plots, these 3D runs are compared with a small-box 2D simulation (with box size $2 \times 2 \times 2\lambda_{\text{MRI}}^2$) employing the same physical parameters. We

observe that the 3D simulations display a qualitatively different behavior overall.

To respect the validity of the nonrelativistic shearing-box paradigm, we have verified that bulk flows always remain subrelativistic in all our runs. However, as was also found in earlier work (Riquelme et al. 2012; Inchingolo et al. 2018), here we observe that the average v_{A}^{NR} in the 2D run vastly exceeds the speed of light during the nonlinear stage, violating the underlying model assumptions. For 2D, Inchingolo et al. (2018) attributed this behavior to the limited size of the simulation box, showing that, for larger domains $\gtrsim 8 \times 8\lambda_{\text{MRI}}^2$, the constraint $v_{\text{A}}^{\text{NR}}/c < 1$ was fulfilled over short times. However, our 2D campaign (Section 4) shows that, in all cases (irrespective of the box size), the system inevitably evolves such that $v_{\text{A}}^{\text{NR}}/c > 1$ over sufficiently long times. This occurs because the out-of-plane B_y fails to saturate in 2D, due to a lack of dissipation processes (i.e., reconnection) along y (see Section 4.2). The violation of the $v_{\text{A}}^{\text{NR}}/c < 1$ condition breaks the underlying nonrelativistic shearing-box assumption, and therefore puts into question the physical validity of the 2D results. In addition, the lack of saturation in 2D heavily contributes to the migration of MRI modes to scales larger than the box size. Our 3D simulations shown in Figure 14 (left panel) display a different behavior, with the magnetic-field amplification (and therefore $v_{\text{A}}^{\text{NR}}/c$) reaching much smaller values. Even a limited-size, cubic-box run already results in $v_{\text{A}}^{\text{NR}}/c \sim 1$ (i.e., ~ 10 times smaller than in the corresponding 2D case), and larger and larger 3D boxes progressively improve this result. Our largest two cases, with box size $4 \times 4 \times 2\lambda_{\text{MRI}}^3$ and $4 \times 8 \times 2\lambda_{\text{MRI}}^3$, firmly maintain $v_{\text{A}}^{\text{NR}}/c \lesssim 1$ for the whole duration of the nonlinear stage.

The fact that, in 3D, MRI saturates at lower magnetic-field amplitudes implies that unstable MRI wavelengths do not




Figure 14. Left: evolution of the volume-averaged Alfvén speed (calculated with the nonrelativistic expression) for several 3D PIC simulations of the pair-plasma MRI, compared with a 2D simulation using the same initial physical parameters. In the 3D runs, the constraint $v_A^{\text{NR}}/c < 1$ is much better respected. Right: evolution of the volume-averaged pressure anisotropy, which influences the MRI growth rate, in the same runs. The 3D runs develop smaller anisotropy over longer times.

migrate to larger length scales as quickly as in 2D. Thanks to the more efficient reconnection during the nonlinear stage, in 3D the MRI is not aggressively suppressed by the unbound growth of magnetic fields. This supports our statement on the significance of 3D simulations: capturing the MRI and reconnection dynamics along the toroidal direction is essential to model a physically meaningful evolution of the instability, which can then maintain a saturated state for longer times.

Two-dimensional simulations also display a qualitatively different behavior in the development of pressure anisotropy: Figure 14 (right panel) shows the evolution of the volume-averaged $\langle (p_\perp - p_\parallel)/p \rangle$ in the same 2D and 3D simulations. In 2D, the anisotropy rapidly increases during the nonlinear stage, and progressively accumulates as the MRI continues to amplify B_y without the dissipative effect of reconnection along y , up to a maximum $\langle (p_\perp - p_\parallel)/p \rangle \simeq 1$. In all 3D runs, this anisotropy is limited to smaller values (up to a maximum $\langle (p_\perp - p_\parallel)/p \rangle \simeq 0.5$) for the whole duration of the nonlinear phase, mildly increasing toward the end of the run. This is verified both in the cubic-box case and in larger boxes, indicating that the additional reconnection along y plays a major role in constraining the development of pressure anisotropy. This anisotropy eventually contributes to pushing MRI wavelengths to larger scales and to halting the MRI development, but over longer timescales than in corresponding 2D cases.

5.1.2. Effect of the Aspect Ratio

The 2:4:1 aspect ratio employed for the simulation presented in Figure 13 has been chosen by carefully analyzing a number of runs with varying L_x , L_y , and L_z . As illustrated in Section 3, the physical domain size in z (i.e., L_z/λ_{MRI}) heavily impacts the simulation cost, when demanding that the macroscopic scale ordering $H/L_z \geq 1$ be respected. In practice, choosing L_z/λ_{MRI} and $\lambda_{\text{MRI}}/\rho c$ determines the required number of grid cells per λ_{MRI} and the number of time steps per P_0 . Then, the aspect ratios L_x/L_z and L_y/L_z are free parameters, each impacting the overall simulation cost linearly (i.e., quadratically when combined). Ideally, we are free to choose a small vertical extent, e.g., $L_z/\lambda_{\text{MRI}} = 2$, and then elongate the box in x and y to better include the reconnection physics (as shown in

Sections 4.3.4 and 5.1). However, the x -extent of the box cannot be increased indefinitely, because we demand that the condition $v_s(L_x)/c < 1$ be satisfied at the radial boundaries. Our tests indicated that numerical artifacts appear at the boundary when $L_x/L_z \geq 4$ (corresponding to $v_s(L_x)/c \geq 0.5$). Hence, we choose $L_x/L_z = 2$ in our 3D simulations. Elongating the box in y is less problematic, since no constraints related to the nonrelativistic shearing-box model exist in the toroidal direction.¹⁷ We have chosen $L_y/L_z = 4$ in our largest run.

In earlier sections, we have argued that boxes with larger L_x/L_z produce better-developed turbulence; Figure 14 indeed demonstrates that, as we enlarge a 1:1:1 box to 2:2:1 and then to 2:4:1, we can observe the emergence of the characteristic cycles of creation and destruction of channel flows. Reconnection acts more efficiently when $L_x/L_z > 1$, resulting in dynamics more similar to that expected from 3D MHD numerical experiments. It remains to be assessed whether the effect of the aspect ratio changes when macroscopic and microscopic length scales are better separated, and whether even larger L_y/L_z produce significant changes in the results. Due to the exceptional computing costs involved for such simulations, we defer this analysis to future work.

5.2. Spectra of 3D MRI-driven Turbulence

In Figure 15, we show the isotropic power spectrum of the poloidal ($B_{\text{pol}} = \sqrt{B_x^2 + B_z^2}$) and toroidal (B_y) magnetic-field components for the large-box simulation discussed in Section 5.1 (with parameters $\omega_C/\Omega_0 = 15$, $\beta_0 = 312$, and box size $4 \times 8 \times 2 \lambda_{\text{MRI}}^3$; see Figure 13). This run is characterized by an “episodic” behavior during the nonlinear stage, where cycles of channel creation/destruction can be clearly identified. The magnetic-field strength and configuration change drastically throughout each cycle; for this reason, in Figure 15, we focus on the cycle occurring toward the end of the nonlinear stage (between $t \simeq 9P_0$ and $t \simeq 11P_0$; see Figure 13), and time-average the power spectra throughout the cycle.

¹⁷ Note, however, that curvature effects have been neglected when constructing the shearing-box equations; L_y/L_z can then be larger than L_x/L_z , but must remain limited to a certain extent.




Figure 15. Isotropic power spectrum of the poloidal (red line) and toroidal (blue line) magnetic field for the 3D pair-plasma MRI simulation of Figure 13 ($\omega_C/\Omega_0 = 15$, $\beta_0 = 312$, and box size $4 \times 8 \times 2\lambda_{\text{MRI}}^3$). The spectra are averaged over one channel disruption/creation cycle during the nonlinear stage (between $t \simeq 9P_0$ and $t \simeq 11P_0$), showing a persistent peak around the most-unstable MRI wavenumber k_{MRI} . The spectra feature characteristic slopes indicating an inertial range, and a spectral break in the vicinity of the average Larmor-radius wavenumber k_ρ measured at $t \simeq 10P_0$.

The results show a persistent peak in the spectra around k_{MRI} , the wavenumber of the most-unstable MRI mode calculated with the initial simulation parameters. This peak is indicative of the fact that the MRI is active through the cycle, although magnetic-field dissipation via reconnection is clearly dominant during the dissipation phase, when channels are being disrupted and magnetic energy is converted into heat. Such dissipation results in a turbulent cascade, which can be identified in the spectra via the presence of an inertial range with characteristic slopes. At moderate wavenumbers, the poloidal-field spectrum features a shallow slope that could be comparable to $k^{-5/3}$, although the spectra are quite noisy (especially in the inertial range) and a precise measurement of the spectral slopes is hard to obtain. Regardless of the exact power-law index, we clearly observe that a spectral break occurs around $k \simeq k_\rho$, with k_ρ being the wavenumber corresponding to the average Larmor radius measured at $t \simeq 10P_0$ (i.e., at the midpoint in time through the cycle). In the sub-Larmor range ($k > k_\rho$), the B_{pol} spectrum steepens to a slope $\propto k^{-3.2}$; in Figure 15, a slope $\propto k^{-3}$ is also indicated to guide the eye. The toroidal-field spectrum similarly features a shallow slope consistent with k^{-2} at $k < k_\rho$ (although this part of the spectrum is also rather noisy), and a spectral break and steepening to $\propto k^{-3.2}$ in the kinetic range.

These 3D spectra present similarities as well as substantial differences with those shown in Figure 7 for the 2D case. Like in 2D, for the poloidal field, we observe an inertial range with a $\propto k^{-5/3}$ slope, indicative of a turbulent MHD-type cascade driven by energy injection at large scales; similarly to 2D, for B_{pol} , we measure a spectral break and a steepening in the sub-Larmor range, although in 3D this part of the spectra features a slightly steeper $\propto k^{-3.2}$ slope. Especially at kinetic scales, 2D and 3D spectra show substantial differences: the B_y spectrum in 3D presents a spectral break not observed in 2D, and the kinetic range is here much more extended. In addition, we pointed out that, in 2D, the spectra continuously evolve, and over time, the measured characteristic slopes do not persist. We believe that

this is an artifact of the reduced 2D dimensionality; in 3D, we have verified that the spectra shown in Figure 15 consistently appear over each channel creation/disruption cycle in the nonlinear stage as long as the MRI is active.

Our 3D results are in good agreement, at least in the inertial range ($k < k_\rho$), with the 3D hybrid simulations of Kunz et al. (2016) as well as with 3D MHD results (e.g., Walker et al. 2016). Slight differences with these works arise in the kinetic range ($k > k_\rho$), where we measure a $\propto k^{-3.2}$ slope instead of the $\propto k^{-3}$ reported by Kunz et al. (2016). However, their result was related to the ion sub-Larmor scales in a hybrid simulation, while here we are considering fully kinetic pair plasmas. Future 3D PIC runs with electron-ion plasma will clarify whether this discrepancy is related to our pair-plasma assumption.

5.3. Particle Energization in 3D

In this section, we briefly analyze particle energization in the 3D pair-plasma MRI. We focus on our largest simulation presented in Figure 13 ($\omega_C/\Omega_0 = 15$, $\beta_0 = 312$, and box size $4 \times 8 \times 2\lambda_{\text{MRI}}^3$) and measure the evolution of the mean Lorentz factor for electrons and positrons separately, as well as the average value for both species. In addition, we compute the energy distribution function at subsequent times during the simulation. The results are shown in Figure 16.

Much like in the 2D case (Section 4.5), we observe that the particle energy constantly increases throughout the run, with the average Lorentz factor reaching values ~ 10 – 40 at the end of the simulation (left panel). Also similarly to our 2D simulations, we measure significant differential heating between the two particle species, with electrons reaching energies ~ 4 times larger than positrons on average. The energy distribution functions measured for both species together (right panel) show that, right after the onset of reconnection at $t \simeq 4.7P_0$, there is a significant nonthermal component in the energy spectra; this nonthermal tail follows a power law that persists until the end of the run as the plasma heats, while an additional high-energy peak progressively develops around $\gamma \simeq 100$ (indicated by a dashed line in Figure 16). At the end of the simulation ($t \simeq 15P_0$), particles have accumulated around the peak energy, eroding the power-law tail of the spectrum.

These 3D results show both remarkable similarities and qualitative differences with our 2D runs. First, the differential heating already found in 2D simulations arises in 3D essentially in the same way, with electrons gaining more energy than positrons; the mechanism driving this effect is the same as in 2D simulations (which we elaborate on in Appendix B). The plasma heating proceeds, on average, very similarly to our large-box 2D runs (see Figure 8), where developed Alfvénic turbulence drives a rather smooth increase in particle kinetic energy. In 3D, the “episodic” nature of the channel creation/disruption cycles manifests itself in the evolution of $\langle \gamma \rangle$, where we observe periodic (small) increases corresponding to the repeated onset of reconnection events. The particle energy distribution also evolves similarly to the 2D case, at least in the beginning of the nonlinear stage; like in 2D, the slope of the power-law tail is ~ 2.5 – 3 , persisting until the end of the run. A qualitative difference with 2D simulations lies in the end state reached in 3D runs: the high-energy peak shown in Figure 16 here precisely corresponds to the Lorentz factor of particles with $\rho_C \simeq L_z$ (calculated with the mean magnetic-field strength at saturation). In 2D, we observed that a high-energy peak develops due to the “quiet-loop” end state that emerges due to




Figure 16. Left: evolution of the average Lorentz factor for electrons, positrons, and for the two species combined in the 3D pair-plasma MRI simulation of Figure 13 ($\omega_c/\Omega_0 = 15$, $\beta_0 = 312$, and box size $4 \times 8 \times 2\lambda_{\text{MRI}}^3$). As in 2D, electrons heat significantly more than positrons. Right: evolution of the particle energy distribution during the nonlinear stage ($t \simeq 4.7\text{--}15P_0$) of the same simulation. Right after the onset of reconnection at $t \simeq 4.7P_0$, the distribution has developed a nonthermal tail following a power law (the slopes of power laws with indices 2.5 and 3 are shown to guide the eye). The nonthermal population persists until the end of the run, while a separated high-energy peak develops around $\gamma \sim 100$ (dotted line), which corresponds to the maximum attainable energy, due to the system size in the z -direction.

the reduced dimensionality; here, it appears that the system is simply reaching a state where particles accumulate around the highest attainable energy according to the simulation-box size. This evolution is very similar to that of forced-turbulence kinetic simulations, where particles accumulate, at late times, around the maximum energy determined by the box size (Zhdankin et al. 2018b). In our runs, the accumulation of particles around the maximum energy and the shift of the distribution to larger average energies progressively reduce the extent of the power law. The development of the high-energy peak also relates to (but is not a direct consequence of) the differential heating discussed above, with electrons accumulating faster around the highest attainable energy.

Both the differential heating and the emergence of the high-energy peak can be ameliorated with larger boxes and/or better scale separation, which imply higher computational costs; for example, we have conducted preliminary runs with larger domain sizes, confirming that in such cases the high-energy peak shifts to higher γ , proportionally to the extent of the smallest side (usually L_z) of the simulation domain. Systematically studying the effect of box size and scale separation on particle energization in 3D involves high computational costs, and we thus defer it to future work.

5.4. Stresses and Angular-momentum Transport

To conclude our first study of 3D MRI simulations, we consider the evolution of the α -parameter and of the viscous stresses. Figure 17 (top left) shows the measured values of α and of the volume-averaged M_{xy} , R_{xy} , and A_{xy} for the simulation presented in Figure 13 ($\omega_c/\Omega_0 = 15$, $\beta_0 = 312$, and box size $4 \times 8 \times 2\lambda_{\text{MRI}}^3$). We clearly observe that, for this run, the anisotropic stress dominates the overall angular-momentum transport; this differs substantially from 2D cases, where $A_{xy} > M_{xy}$ was verified only at late times (when the MRI is halting and pressure anisotropy accumulates). The top-right, bottom-left, and bottom-right panels of Figure 17 show, respectively, the evolution of α and of the volume-averaged M_{xy} and A_{xy} for the 2D and 3D cases already compared in

Section 5.1.1. In all these cases, $\omega_c/\Omega_0 = 15$ and $\beta_0 = 312$, while the box size and dimensionality are varied to compare the 2D and 3D dynamics as well as the 3D small- and large-box evolution. We observe that, during the active nonlinear phase, α is smaller for 3D runs than in the 2D case; a further decrease in α occurs from small to large box sizes in 3D runs. Comparing the most important stress components, we notice that M_{xy} follows the same decreasing trend observed for α , while A_{xy} remains essentially unchanged across all runs. It is therefore evident that the smaller value of $\alpha \simeq 0.1$ achieved in the large-box 3D case (whereas $\alpha \geq 1$ in the active nonlinear phase of the 2D run) is solely related to a decrease in the Maxwell stress.

The explanation for this behavior is straightforward: as we have argued in the previous sections, large 3D boxes drive much more efficient reconnection dynamics, which in turn results in a magnetic-field saturation at lower values than both small 3D simulations and arbitrarily large 2D runs. The Maxwell stress, on average, becomes correspondingly smaller and so does α . Since A_{xy} remains somewhat constant across all these cases, our results could in principle suggest that the main driving mechanism for angular-momentum transport in large 3D simulations is represented by pressure anisotropy (this result was also reported by Hoshino 2015). However, we believe that the large value of A_{xy} observed here may simply be an artifact of the reduced scale separation employed in our 3D numerical experiments. In particular, we measure $\langle A_{xy} \rangle / \langle p \rangle \simeq 0.1$ in these runs, which is nearly twice as large as the value found in 2D simulations with large boxes or large scale separation (compare Figure 17). In such cases, pressure anisotropy is more efficiently limited by developed turbulence and/or faster mirror modes; with the limited scale separation and box size employed in our 3D runs, it is not unreasonable to think that the measured anisotropy is thus exaggerated. Three-dimensional Braginskii-MHD and hybrid simulations (Sharma et al. 2006; Kunz et al. 2016) have indeed shown that $A_{xy} \sim M_{xy}$ is expected in fluid modeling, but that anisotropic momentum transport does not typically dominate the overall value of α by




Figure 17. Top left: evolution of α and of the volume-averaged Maxwell, Reynolds, and anisotropic stress in the 3D pair-plasma MRI simulation of Figure 13 ($\omega_c/\Omega_0 = 15$, $\beta_0 = 312$, and box size $4 \times 8 \times 2\lambda_{\text{MRI}}^3$). Top right: comparison of α for 2D and 3D simulations employing the same ω_c/Ω_0 and β_0 and different box sizes. Bottom left and right: comparison of the volume-averaged Maxwell and anisotropic stress for the same 2D and 3D simulations.

the large margin we observe here. To assess the precise role and trend of anisotropic stresses in 3D fully kinetic simulations, we will conduct large-box studies with better scale separation in future work.

6. Discussion and Conclusions

In this work, we have presented a comprehensive exploration of the magnetorotational instability (MRI) in collisionless pair plasmas via fully kinetic particle-in-cell simulations. With a shearing-box setup implemented in our relativistic PIC code ZELTRON (Cerutti et al. 2013), we have carried out a vast array of 2D runs, exploring an unprecedentedly large parameter space. In particular, we have conducted very large-scale 2D simulations with macroscopic-to-microscopic temporal-scale separation up to ~ 10 times larger than previous works and with system size up to twice as large as the largest simulations presented in literature. In addition, we have carried out large-scale 3D PIC simulations of the MRI in pair plasmas, achieving for the first time a global mesoscale dynamics akin to that observed in MHD works. To study the axisymmetric MRI with PIC, we have resorted to the established 2D shearing-coordinate approach employed in previous publications (Riquelme et al. 2012; Inchingolo et al. 2018); for 3D nonaxisymmetric simulations, we have developed and applied a novel “orbital-advection” formulation of the shearing box that simplifies pre-existing methods (Hoshino 2013, 2015) and is less complicated to implement numerically.

An important result of our study is the demonstration that 2D (axisymmetric) fully kinetic simulations of the MRI are qualitatively different from corresponding 3D kinetic runs. We have pointed out for the first time such differences, and provided a physical explanation for the behavior of the kinetic MRI in 2D. We found that axisymmetric PIC simulations are intrinsically limited in the physical processes they can model, and such a limitation has important consequences for the overall MRI evolution. In particular, the lack of reconnection along the out-of-plane (toroidal) direction results in a lack of saturation of the strongest magnetic-field component, B_y , with

detrimental consequences for the development and sustainment of turbulence. However, we also found that many important lessons learned from 2D runs will carry over to the 3D case: for example, with our extensive parameter scans in 2D, we have determined the minimal values of key physical parameters (in terms of the dimensionless scale separation, system size, etc.) that need to be reached for meaningful studies. Thus, 2D simulations give us guidance (starting point) for designing 3D runs. Differences and similarities between the 2D and 3D cases are illustrated in detail below.

By conducting large-scale 3D simulations of the MRI, for the first time we have demonstrated that fully kinetic simulations can reproduce the mesoscale dynamics typically observed in MHD: multidimensional, macroscopic channel flows cyclically form and get disrupted by parasitic instabilities, particularly via drift-kink modes. The channel disruption results in developed turbulence, and several channel creation/growth/disruption cycles can be observed during the nonlinear stage in large-scale 3D runs. This result also aligns with previous 3D hybrid simulations (Kunz et al. 2016). Conversely, channels developing in small-scale runs are not efficiently disrupted and no turbulence develops, in agreement with the sole 3D fully kinetic study presented in literature so far (Hoshino 2015), which employed a simulation box of limited size.

In addition to studying the general MRI dynamics, we have explored in detail several physical processes that develop concurrently with the linear and nonlinear stages of the main instability. Our results can be summarized as follows:

1. We have extensively described the typical nonlinear evolution of the MRI in fully kinetic shearing-box simulations, explaining each phase of the evolution and the differences between the small- and large-box cases, as well as between the 2D and 3D cases. We have explained how small and large simulations differ in producing developed turbulent states, and how the choice of physical parameters impacts the overall evolution and duration of the nonlinear MRI. We have also identified

the role of pressure anisotropy in these simulations, observing that a large $\Delta p = p_\perp - p_\parallel$ can develop and accumulate throughout the nonlinear stage of the instability; this anisotropy is likely exaggerated by our choice of modest scale separation (but still much larger than that employed in previous studies) that necessarily characterizes PIC simulations. A limited scale separation can result in inefficient mirror modes, which grow much more slowly than the MRI, allowing for persistently large pressure anisotropy. This anisotropy ultimately participates in pushing MRI modes to wavelengths larger than the box size, halting the MRI dynamics over long times.

2. By exploring a large parameter space, we have assessed the numerical convergence of key quantities (e.g., the magnetic-field amplification) in 2D simulations when the physical parameters are varied. We have studied the effect of the separation between the macroscopic (Ω_0^{-1}) and microscopic (ω_C^{-1}) temporal scales, ω_C/Ω_0 ; of the separation between box size and MRI scales L/λ_{MRI} ; of the initial plasma- β ; and of the box aspect ratio. Increasing ω_C/Ω_0 with a fixed box size generally results in a longer duration of the nonlinear stage and in an increase of the magnetic energy at saturation. We have shown that, at some large ω_C/Ω_0 , the results (in particular, the saturated magnetic energy) eventually converge; however, this converged nonlinear state is qualitatively different for small and large boxes, and the precise value of ω_C/Ω_0 resulting in convergence depends on the box size. Similarly, fixing ω_C/Ω_0 and only increasing the box size produces results that converge for sufficiently large box sizes; in large boxes, a sustained-turbulence state can develop, whereas in small simulations it does not. This partly agrees with previous 2D work by Inchingolo et al. (2018), who focused their analysis on increasing box sizes while keeping ω_C/Ω_0 fixed. They concluded that convergence in the magnetic-field amplification is attained, in 2D, for box sizes of at least $8 \times 8 \lambda_{\text{MRI}}^2$, and that large-box simulations can maintain a volume-averaged $v_A/c < 1$ (calculated with the nonrelativistic expression) throughout the system's evolution, respecting the underlying assumptions of the nonrelativistic shearing box. We have instead demonstrated that 2D simulations, over sufficiently long times, invariably develop an average $v_A/c > 1$, owing to the absence of reconnection in the y -direction. This does not occur, however, in our 3D runs of sufficiently large size; we conclude that 3D simulations are of key importance to obtain physically valid results that respect all the underlying assumptions of the shearing-box model. Considering the effect of the initial plasma- β , we found that the results are practically unchanged in the range $\beta_0 \simeq 0.1$ –2500: independent of the initial β_0 , all simulations end up with plasma beta more or less within the same universal final range, of order 0.1–1 for 2D runs. This result holds as long as the unstable MRI modes can fit in the simulation box, i.e., under the constraint $L_z > \lambda_{\text{MRI}}$. Finally, we have observed that the geometry of the simulation box can impact the development of the nonlinear MRI stage dramatically: boxes with aspect ratio $L_x/L_z > 1$ result in developed turbulence much more easily than corresponding boxes with $L_x/L_z = 1$, both in 2D and 3D. We have observed that this is caused by

additional drift-kink modes that are impeded in small simulation domains, and that can promote channel-disruption events in elongated boxes instead.

3. We have verified that sustained turbulence develops, in our simulations, during the nonlinear MRI evolution. The isotropic power spectra of the poloidal ($B_{\text{pol}} = \sqrt{B_x^2 + B_z^2}$) and toroidal (B_y) magnetic field during this phase show the presence of an inertial range with characteristic power laws, indicative of turbulent activity. Both in 2D and 3D, the B_{pol} spectrum features a shallow slope roughly consistent with $\propto k^{-5/3}$ at large and intermediate spatial scales, and a spectral break in the vicinity of the average Larmor-radius wavenumber $k_\rho = \rho_C^{-1}$. In the 3D case, at length scales below ρ_C , the B_{pol} spectrum steepens to a $\propto k^{-3.2}$ slope and a clear kinetic range is present. The B_y spectrum appears to follow a $\propto k^{-3}$ slope at all scales in 2D; in 3D, B_y instead follows a shallower slope compatible with $\propto k^{-2}$ in the inertial range, and a $k^{-3.2}$ slope in the kinetic range. Our 2D results are consistent with previous 2D studies (Inchingolo et al. 2018); our 3D results for B_y , in the inertial range, are in agreement with MHD and hybrid-kinetic simulations (Kunz et al. 2016; Walker et al. 2016). However, the latter comparison is complicated by the difference in the underlying models (our fully kinetic PIC versus MHD or hybrid methods). Considering works focusing on pair plasmas, a similarity can be drawn between our model and that employed in Zhdarkin et al. (2017, 2018a), since both cases consider a forced-turbulence system evolution. However, in those works, a $\propto k^{-4}$ slope (or steeper) was found for $k > k_\rho$, which differs from our shallower $k^{-3.2}$ result in the kinetic range. This measurement is also hard to compare against analytic expectations, as no previous studies (to the best of our knowledge) have focused on the specific case of pair-plasma MRI-driven turbulence we consider. For example, Loureiro & Boldyrev (2018) considered a low- β pair plasma in a tearing-mediated cascade, finding a $\propto k^{-3}$ slope in the kinetic range; our conditions, however, are those of a high-beta plasma, and we have no basis to claim that our cascade is tearing-mediated in nature. It is also interesting to note that solar-wind observations commonly report a spectral slope $\propto k^{-7/3} - k^{-8/3}$ (not extremely far from our $\propto k^{-3.2}$ result) at sub-ion-Larmor scales, consistent with gyrokinetic calculations (e.g., Schekochihin et al. 2009). Whether a similarity can be drawn between this and our case will be verified with future electron-ion simulations.
4. Concerning particle energization in 2D and 3D, we found that energy injection in large runs is akin to that brought about by Alfvénic turbulence. In these simulations, turbulence is well-developed and particle heating proceeds smoothly, with the system achieving a steady-state balance between turbulent magnetic-energy dissipation and particle energization. In contrast, when turbulence is not well-developed (e.g., in small boxes), particle energization occurs mostly during short “bursts” corresponding to large-scale reconnection events. Moreover, for the first time, we have reported that substantial differential heating can occur between electrons and positrons in MRI simulations. The physical mechanism behind this phenomenon is represented by additional drift forces

(related to the background differential rotation) that affect opposite charges differently, causing a symmetry break in the gyromotion of electrons and positrons in a uniform magnetic field. This effect is unrealistically large in our simulations, due to our choice of (necessarily) limited scale separation ω_C/Ω_0 . We have elaborated on how tuning the values of physical parameters can ameliorate this issue and its implications for the interpretation of the results.

5. The particle energy distribution functions in our 2D and 3D simulations show the presence of substantial nonthermal particle acceleration during the nonlinear MRI stage. A power-law tail with index ~ 2.5 – 3 consistently develops as the MRI transitions to sustained turbulence, both in 2D and in 3D. Our results are in good agreement with previous works; however, here we have demonstrated that in 2D this power-law state is transitory, and that the overall evolution can produce very different energy distributions over time. Inchingolo et al. (2018) carried out a similar analysis but focused on the early nonlinear stage, where a power-law index ~ 2 can indeed be realized. However, such a state is clearly still evolving, and it changes drastically later on. At late times, the power-law part of the distributions progressively disappears, and a high-energy peak develops. This is completely determined by the artificial 2D end state (with characteristic “magnetic loops”), driven by the reduced dimensionality of these runs. Riquelme et al. (2012) described a similar late-time 2D dynamics, but attributed this evolution to the MRI; we emphasize that such a state is in fact not representative of the response of particles to the MRI, because the latter has slowed down (or completely stopped) by the time when the quiescent end state has developed. Conversely, in our 3D simulations, the developed nonthermal features are maintained with the same power-law index until the end of the run. However, the nonthermal tail is progressively eroded by the accumulation of particles around the highest energy attainable (which increases with box size). Hoshino (2015) observed a similar effect in small-box 3D simulations; this result is also consistent with forced-turbulence PIC simulations (Zhdankin et al. 2018b) and points to the need for larger (and more expensive) 3D numerical experiments. Note that the slope of the nonthermal tail we measure in our simulations differs from the results presented by Hoshino (2015); it is also unclear whether the development of the observed nonthermal features in 3D remains unchanged when further increasing the box size and/or scale separation. In future work, we will perform additional simulations with a focus on assessing the impact of physical and numerical parameters on the resulting particle energization.
6. In our 2D and 3D simulations, we observed that viscous stresses (in particular, the dominant xy -component of the stress tensor) can develop during the nonlinear MRI stage. This can lead to efficient angular-momentum transport: an effective (dimensionless) collisionless viscosity $\alpha = (\langle M_{xy} \rangle + \langle A_{xy} \rangle + \langle R_{xy} \rangle) / \langle p \rangle$ (Shakura & Sunyaev 1973) arises both in 2D and 3D runs, mainly due to Maxwell and anisotropic stresses. In 2D, the Maxwell stress $M_{xy} = -B_x B_y / (4\pi)$ is likely exaggerated due to a lack of saturation in B_y caused by the reduced

dimensionality; in 3D large-scale runs, where magnetic fields saturate at lower amplitudes, this stress firmly settles on $\langle M_{xy} \rangle / \langle p \rangle \simeq 0.01$ – 0.1 , consistently with previous hybrid simulations (Kunz et al. 2016). The anisotropic stress $A_{xy} = -(p_\perp - p_\parallel) B_x B_y / B^2$, on the other hand, shows a more complex trend. When keeping the same ω_C/Ω_0 , the average pressure anisotropy $(p_\perp - p_\parallel)$ generally decreases from 2D to 3D and from small to large boxes; however, because the average B^2 decreases as well, $\langle A_{xy} \rangle / \langle p \rangle$ maintains roughly the same average value ~ 0.1 in 2D and 3D simulations. As a result, we measure $A_{xy} > M_{xy}$ in our largest 3D runs, which is counter to the general expectation $A_{xy} \lesssim M_{xy}$. We believe that this enhanced anisotropic angular-momentum transport is related to our (necessarily) limited scale separation and system size; with realistic parameters, efficient mirror modes would rapidly quench pressure anisotropy, likely decreasing A_{xy} . Hoshino (2015) measured a similarly enhanced angular-momentum transport due to pressure anisotropy in small-box 3D simulations, which was likely due to the same mechanisms we describe here.

Although we have addressed several key aspects of the kinetic physics of the MRI, our study still presents several limitations. The most prominent open question is whether 3D PIC simulations could produce results that completely align with hybrid and MHD studies, and how the nonthermal particle acceleration and angular-momentum transport depend on the box size and choice of physical parameters. Here, we have conducted the largest (to date) 3D PIC runs that successfully reproduce the global mesoscale MRI dynamics, but our scale separation and box size remain rather limited. Larger simulations with better separation would be needed (i) to completely characterize the properties of particle energization, e.g., whether large-scale runs, where the box size does not rapidly interfere with the nonthermal part of the spectrum, present the same nonthermal features as those observed here; and (ii) to understand the general trend of viscous stresses when the scale separation is substantially larger, possibly achieving the expected ordering $A_{xy} \lesssim M_{xy}$ in simulations where pressure anisotropy is well regulated. Moreover, the issue of substantial differential heating in 3D runs remains to be addressed, as it can complicate the analysis of particle energization even in pair-plasma simulations.

Despite these limitations, our work paves the way for several future studies focused on the exploration of other properties of the MRI. For example, we have reported here the stark contrast between simulations with and without net magnetic flux; this is a well-known aspect of the MRI dynamics that has been extensively studied with MHD shearing-box models (e.g., Gardiner & Stone 2005). Several works have even pointed out that 3D simulations with and without net magnetic flux behave substantially differently for boxes with different aspect ratios (e.g., Bodo et al. 2008; Walker & Boldyrev 2017); moreover, in ideal MHD, it is generally observed that the α -parameter representing angular-momentum transport tends to zero in the case of zero net magnetic flux, unless explicit dissipation coefficients are included (e.g., Fromang et al. 2007; Fromang & Papaloizou 2007). In collisionless models (including ours), such dissipative effects are naturally included (as shown, e.g., by Kunz et al. 2016), and the phenomenology of angular-momentum transport can be studied from first principles. All in

all, we believe that a systematic PIC study of the zero-net-flux case may be particularly important to understand the exact mechanisms behind the phenomenology of MRI dynamo; we will pursue such a study in the future. Another important analysis to be carried out in future work concerns the initial magnetic-field configuration: here, we have limited ourselves to the purely vertical-field case, but it is well-known that different field geometries (especially when an initial B_y is present) can result in qualitatively different global MRI evolution and turbulence properties (e.g., Gardiner & Stone 2005; Goedbloed & Keppens 2022). This aspect has not been considered in kinetic or hybrid works, but it would be important to study generic field geometries, which are expected in realistic astrophysical scenarios. Finally, we note that studying the kinetic MRI with vertical stratification (as opposed to the case without stratification that we considered) could in principle provide results that are more physically relevant for modeling thick accretion disks (e.g., Hirabayashi & Hoshino 2017) and resolve convergence issues that notably affect MHD studies (e.g., Regev & Umurhan 2008; Bodo et al. 2014).

We conclude by remarking that the next step in our line of work concerns electron–ion ($m_i \gg m_e$) simulations. Studying the large mass ratio case, and especially the energy partition between particle species, is fundamental for the correct interpretation of data from current and future observations targeting RIAFs (EHT Collaboration 2019). In these collisionless astrophysical environments, electrons and ions are decoupled and a two-temperature state can originate from turbulence and/or radiation-reaction dynamics (Arzamasskiy et al. 2019; Zhdankin et al. 2019; Kawazura et al. 2020, 2021; Zhdankin et al. 2021). With our simulations, we will be able to study the process of temperature decoupling (and possibly the dynamics of radiation) during the MRI from first principles. Moreover, information from fully kinetic simulations can be used as input for global fluid simulations of accretion around compact objects employing subgrid models to account for the missing microphysics (Ressler et al. 2015; Chael et al. 2018; Scepi et al. 2022). Given the ever-growing interest in understanding the dynamics of collisionless accretion disks around SMBHs, we believe that this work has the potential to substantially impact our insight into the physics of such space environments.

The authors would like to thank Mario Riquelme, Masahiro Hoshino, Giannandrea Inchingolo, Andrea Mignone, Gianluigi Bodo, Jim Stone, Eliot Quataert, Anatoly Spitkovsky, Matthew Kunz, Elias Most, Sasha Philippov, Nuno Loureiro, and Nicolas Scepi for useful discussions and suggestions throughout the development of this work.

The authors acknowledge support from a NASA ATP grant 80NSSC20K0545. Additional support was provided by National Science Foundation grants AST 1806084 and AST 1903335 to the University of Colorado. F.B. was also partially supported by a Junior Post-doctoral Fellowship (grant number 12ZW220N) from Research Foundation—Flanders (FWO). Support for L.A. was provided by the Institute for Advanced Study. Support for V.Z. at the Flatiron Institute is provided by the Simons Foundation.

This work made use of substantial computational resources provided by several entities. We acknowledge support by the VSC (Flemish Supercomputer Center), funded by the Research

Foundation—Flanders (FWO) and the Flemish Government—department EWI. An award of computer time was provided by the Innovative and Novel Computational Impact on Theory and Experiment (INCITE) program, using resources (namely Theta) of the Argonne Leadership Computing Facility, which is a DOE Office of Science User Facility supported under contract DE-AC02-06CH11357. Additional computer time was provided by the Extreme Science and Engineering Discovery Environment (XSEDE), which is supported by National Science Foundation grant No. ACI-1548562 (Towns et al. 2014), and by the Frontera computing project (Stanzione et al. 2020) at the Texas Advanced Computing Center (TACC, www.tacc.utexas.edu). We acknowledge TACC for providing HPC resources on both Stampede2 and Frontera. Additional resources (namely Pleiades) supporting this work were provided by the NASA High-End Computing (HEC) Program through the NASA Advanced Supercomputing (NAS) Division at Ames Research Center.

Appendix A

Numerical Methods for the Kinetic Shearing Box

A.1. Shearing-coordinate Simulations in 2D

For two-dimensional simulations, in this work we employ the KSB-SC formulation, which is preferable because no complicated boundary conditions are required (the simulation domain is periodic in all directions). Numerically, the implementation of the KSB-SC system of equations in a PIC code requires the solution of modified Maxwell’s equations and a peculiar particle-push step. These modifications are, however, straightforward: for the electromagnetic fields, the additional terms on the right-hand side of Equations (13) and (14) are linear and point-wise, and can therefore be treated explicitly with a standard leapfrog algorithm by interpolating in time and space to maintain second-order accuracy (see Riquelme et al. 2012). The particle-push step is similarly formulated as a modification of a standard Boris scheme that includes here extra (linear) terms. We refer to Section A.2.2 for more details on this modified Boris algorithm.

A.2. Orbital-advection Simulations in 3D

For three-dimensional simulations, in this work we make use of the newly developed KSB-OA framework (Section 2.3). The numerical implementation of our method requires solving modified Maxwell’s equations that generally include implicit terms, as well as a modified Boris push (similar to that required for 2D runs, see Section A.1 above). Our numerical approach also requires shearing-periodic boundary conditions on electromagnetic fields, source terms, and particles.

A.2.1. Solution of Maxwell’s Equations

Our versions of Maxwell’s Equations (20) and (21) involve nonlocal, coupled terms on the right-hand side; this complicates the numerical solution, since a simple leapfrog scheme cannot be applied to explicitly push the electromagnetic fields in time. To avoid costly matrix-inversion operations (see, e.g., Bacchini et al. 2019), in our first numerical implementation, we split the solution in two explicit steps. First, all terms on the right-hand side of Maxwell’s equations that are not related to advection along y are discretized with a central-difference scheme in time and space. The resulting system of discrete

equations (excluding the advective terms) is iterated upon with a simple predictor-corrector step, which is applied three times.¹⁸ At the end of each predictor-corrector iteration, we can include the advection of $\mathbf{E}$ and $\mathbf{B}$ along y in several ways. In our numerical experiments, we have observed that a simple central-difference scheme applied to these terms inevitably results in instabilities developing at the shearing boundaries; we could therefore opt for a more stable advection scheme among the many available in literature. However, as a first approach, we choose an even simpler strategy: to model the advection of field quantities along y , we simply shift each component of $\mathbf{E}$ and $\mathbf{B}$ in space by a distance $v_{s,y}(x)\Delta t$ (with appropriate signs). The shifted arrays are interpolated at the resulting spatial locations, which approximates the advection step. This operation introduces some numerical diffusion, but we observe that this is in fact beneficial, as it suppresses numerical instabilities at the shearing boundaries. Our approach represents a cheap, less sophisticated alternative to more classical schemes; while this leaves ample grounds for improvements, it only requires simple spatial interpolation to model advection, and we observe that this choice produces acceptable results for this first study. In future works, we will implement more robust algorithms to accurately capture the field advection and minimize the numerical diffusion introduced by our approach. Finally, Gauss's law (Equation (23)) is enforced, in our approach, with a Langdon–Marder correction

where $\bar{\gamma}$ is the Lorentz factor defined at an intermediate time between $n + 1/2$ and $n - 1/2$. Next, we recast this equation as

$$\begin{aligned} \mathbf{u}^+ - \mathbf{u}^+ \times \bar{\boldsymbol{\tau}} - \alpha_1 \mathbf{u}_x^+ \hat{\mathbf{e}}_y - \alpha_2 (\mathbf{u}^+ \cdot \bar{\boldsymbol{\epsilon}}) \hat{\mathbf{e}}_y \\ = \mathbf{u}^- + \mathbf{u}^+ \times \bar{\boldsymbol{\tau}} + \alpha_1 \mathbf{u}_x^- \hat{\mathbf{e}}_y + \alpha_2 (\mathbf{u}^- \cdot \bar{\boldsymbol{\epsilon}}) \hat{\mathbf{e}}_y, \end{aligned} \quad (\text{A2})$$

where we have defined $\mathbf{u}^\pm = \mathbf{u}^{n\pm 1/2} \mp \boldsymbol{\epsilon}$, $\bar{\boldsymbol{\tau}} = \boldsymbol{\tau}/\bar{\gamma} + \Delta t \boldsymbol{\Omega}_0$, $\bar{\boldsymbol{\epsilon}} = \boldsymbol{\epsilon}/\bar{\gamma}$, $\alpha_1 = s\Omega_0 \Delta t/2$, $\alpha_2 = s\Omega_0 x^n$, and $\mathbf{u}^+$ is the unknown to be solved for. Assuming that $\bar{\gamma}$ is known (see next paragraph), Equation (A2) can be solved explicitly by recasting it as

$$\mathbf{A} \mathbf{u}^+ = \mathbf{C}, \quad (\text{A3})$$

where the right-hand side is the known term $\mathbf{C} = \mathbf{u}^- + \mathbf{u}^- \times \bar{\boldsymbol{\tau}} + \alpha_1 \mathbf{u}_x^- \hat{\mathbf{e}}_y + \alpha_2 (\mathbf{u}^- \cdot \bar{\boldsymbol{\epsilon}}) \hat{\mathbf{e}}_y$, and

$$\mathbf{A} = \begin{bmatrix} 1 & -\bar{\tau}_z & \bar{\tau}_y \\ \bar{\tau}_z - \alpha_1 - \alpha_2 \bar{\epsilon}_x & 1 - \alpha_2 \bar{\epsilon}_y & -\bar{\tau}_x - \alpha_2 \bar{\epsilon}_z \\ -\bar{\tau}_y & \bar{\tau}_x & 1 \end{bmatrix}. \quad (\text{A4})$$

The solution is then straightforward,

$$\mathbf{u}^+ = \mathbf{A}^{-1} \mathbf{C}, \quad (\text{A5})$$

with the inverse matrix

$$\mathbf{A}^{-1} = \frac{1}{D} \begin{bmatrix} 1 + \bar{\tau}_x^2 + \alpha_2 (\bar{\tau}_x \bar{\epsilon}_z - \bar{\epsilon}_y) & \bar{\tau}_z + \bar{\tau}_x \bar{\tau}_y & -\bar{\tau}_y + \bar{\tau}_x \bar{\tau}_z + \alpha_2 (\bar{\tau}_y \bar{\epsilon}_y + \bar{\tau}_z \bar{\epsilon}_z) \\ -\bar{\tau}_z + \bar{\tau}_x \bar{\tau}_y + \alpha_1 + \alpha_2 (\bar{\tau}_y \bar{\epsilon}_z + \bar{\epsilon}_x) & 1 + \bar{\tau}_y^2 & \bar{\tau}_x + \bar{\tau}_y \bar{\tau}_z - \alpha_1 \bar{\tau}_y + \alpha_2 (-\bar{\tau}_y \bar{\epsilon}_x + \bar{\epsilon}_z) \\ \bar{\tau}_y + \bar{\tau}_x \bar{\tau}_z - \alpha_1 \bar{\tau}_x - \alpha_2 (\bar{\tau}_x \bar{\epsilon}_x + \bar{\tau}_y \bar{\epsilon}_y) & -\bar{\tau}_x + \bar{\tau}_y \bar{\tau}_z & 1 + \bar{\tau}_z^2 - \alpha_1 \bar{\tau}_z - \alpha_2 (\bar{\tau}_z \bar{\epsilon}_x + \bar{\epsilon}_y) \end{bmatrix}, \quad (\text{A6})$$

(Marder 1987; Langdon 1991); the resulting (small) errors on charge conservation are thus bounded in time.

A.2.2. Modified Boris Push and Particle-position Update

To solve the momentum Equation (40), we proceed as follows: first, we define $\boldsymbol{\epsilon} = q\Delta t \mathbf{E}/(2m)$, $\boldsymbol{\tau} = q\Delta t \mathbf{B}/(2mc)$; we then discretize the equation of motion between two consecutive time steps as

$$\begin{aligned} \mathbf{u}^{n+1/2} - \mathbf{u}^{n-1/2} &= 2\boldsymbol{\epsilon} + \frac{\mathbf{u}^{n+1/2} + \mathbf{u}^{n-1/2}}{\bar{\gamma}} \times \boldsymbol{\tau} \\ &+ \Delta t (\mathbf{u}^{n+1/2} + \mathbf{u}^{n-1/2}) \times \boldsymbol{\Omega}_0 \\ &+ \frac{s\Omega_0 \Delta t}{2} (\mathbf{u}_x^{n+1/2} + \mathbf{u}_x^{n-1/2}) \hat{\mathbf{e}}_y \\ &+ s\Omega_0 x^n \left(\frac{\mathbf{u}^{n+1/2} + \mathbf{u}^{n-1/2}}{\bar{\gamma}} \cdot \boldsymbol{\epsilon} \right) \hat{\mathbf{e}}_y, \end{aligned} \quad (\text{A1})$$

where

$$\begin{aligned} D &= 1 + \bar{\tau}^2 - \alpha_1 (\bar{\tau}_z + \bar{\tau}_x \bar{\tau}_y) \\ &- \alpha_2 [\bar{\tau}_y (\bar{\boldsymbol{\tau}} \cdot \bar{\boldsymbol{\epsilon}}) + \bar{\epsilon}_y + \bar{\epsilon}_x \bar{\tau}_z - \bar{\epsilon}_z \bar{\tau}_x]. \end{aligned} \quad (\text{A7})$$

After $\mathbf{u}^+$ has been calculated, the particle velocity at the next time step is simply

$$\mathbf{u}^{n+1/2} = \mathbf{u}^+ + \boldsymbol{\epsilon}. \quad (\text{A8})$$

The definition of $\bar{\gamma}$ is not unique, and it is in fact the peculiarity that generally distinguishes relativistic particle pushers (see, e.g., Ripperda et al. 2018). For simplicity, we can define the mid-point Lorentz factor as

$$\bar{\gamma} = \sqrt{1 + [(u^-)^2 + 2\alpha_1 \mathbf{u}_x^- \mathbf{u}_y^- + 2\alpha_2 (\mathbf{u}^- \cdot \boldsymbol{\epsilon}) \mathbf{u}_y^- / \bar{\gamma}] / c^2}, \quad (\text{A9})$$

which is a choice similar to that used in the standard Boris pusher. This definition is such that, in the absence of electric fields, energy is conserved exactly. We can solve the resulting depressed cubic equation for $\bar{\gamma}$ with any available algorithm, e.g., the one described by Tiruneh (2020): define $X = a_1/3$, $Y = -a_2/2$, where

$$a_1 = -1 - [(u^-)^2 + 2\alpha_1 \mathbf{u}_x^- \mathbf{u}_y^-] / c^2, \quad (\text{A10})$$

$$a_2 = -2\alpha_2 (\mathbf{u}^- \cdot \boldsymbol{\epsilon}) \mathbf{u}_y^- / c^2, \quad (\text{A11})$$

¹⁸ This approach is essentially an iterative Crank–Nicolson scheme. It is well-known that this scheme is not asymptotically stable for an infinite number of iterations (see, e.g., Leiler & Rezzolla 2006); the scheme is, however, stable for three iterations, which is our choice for all simulations. This represents a good compromise between stability and accuracy in the solution.

such that $\bar{\gamma}^3 + a_1\bar{\gamma} + a_2 = 0$. The discriminant $Z = X^3 + Y^2$ indicates that there is only one real root if $Z > 0$,

$$\bar{\gamma} = \sqrt[3]{Y + \sqrt{Z}} + \sqrt[3]{Y - \sqrt{Z}}, \quad (\text{A12})$$

and three real roots (of which only one positive and >1) if $Z < 0$,

$$\bar{\gamma}_j = 2\sqrt{-X} \cos \frac{\phi + 2(j-1)\pi}{3}, \quad (\text{A13})$$

where $j = 1, 2, 3$, and $\phi = \cos^{-1}(Y/\sqrt{-X^3})$.

When the new particle velocity is known, the position push can be carried out in two steps. The position is first updated with the motion with respect to the background, i.e.,

$$\mathbf{x}' = \mathbf{x}^n + \Delta t \frac{\mathbf{u}^{n+1/2}}{\gamma^{n+1/2}}, \quad (\text{A14})$$

after which the current can be collected with the usual Villasenor–Buneman charge deposition (Villasenor & Buneman 1992) using the displacement $\mathbf{x}' - \mathbf{x}^n$. Then, the position is updated with the shearing-velocity part,

$$\mathbf{x}^{n+1} = \mathbf{x}' - s\Omega_0 \frac{\mathbf{x}' + \mathbf{x}^n}{2} \Delta t \hat{\mathbf{e}}_y, \quad (\text{A15})$$

which only modifies the y -coordinate, using the x -position at the midpoint in time (note that $\mathbf{x}^{n+1} \equiv \mathbf{x}'$). Depositing the current using $\mathbf{x}' - \mathbf{x}^n$ instead of $\mathbf{x}^{n+1} - \mathbf{x}^n$ is essential to recover $\mathbf{J}'$ (in this paradigm, the current due to velocity fluctuations with respect to the background) instead of $\mathbf{J}$.

As noted in Section 2.3, if particle velocities are relativistic, the particle position update above may result in superluminal motion, since it is not guaranteed that $(\mathbf{u}/\gamma + \mathbf{v}_s)^2 < c^2$. As an alternative strategy, we can employ the proper relativistic velocity boost to update the particle position with a speed that is ensured to be subluminal. Applying a central-difference scheme to Equations (47)–(49), the position update is then formulated as

$$\begin{aligned} x^{n+1} = x^n + \Delta t \frac{u_x^{n+1/2}}{\gamma^{n+1/2}} \\ + u_y^{n+1/2} v_{s,y}^{n+1/2} / \gamma_s^{n+1/2}, \end{aligned} \quad (\text{A16})$$

$$\begin{aligned} y^{n+1} = y^n + \Delta t \frac{u_y^{n+1/2}}{\gamma^{n+1/2} + u_y^{n+1/2} v_{s,y}^{n+1/2}} \\ + \Delta t \frac{\gamma^{n+1/2} v_{s,y}^{n+1/2}}{\gamma^{n+1/2} + u_y^{n+1/2} v_{s,y}^{n+1/2}}, \end{aligned} \quad (\text{A17})$$

$$z^{n+1} = z^n + \Delta t \frac{u_z^{n+1/2}}{[\gamma^{n+1/2} + u_y^{n+1/2} v_{s,y}^{n+1/2}] \gamma_s^{n+1/2}}, \quad (\text{A18})$$

where $v_{s,y}^{n+1/2} = -s\Omega_0 x^{n+1/2}$, $\gamma_s^{n+1/2} = 1/\sqrt{1 - (v_{s,y}^{n+1/2}/c)^2}$, and $x^{n+1/2} = (x^{n+1} + x^n)/2$. These equations are less trivial to employ, especially given that Equation (A16) is now nonlinearly implicit in the unknown x^{n+1} . Fortunately, the nonlinearity is weak and there are no numerical pathologies on the right-hand side (e.g., the denominator is always >1). The position update can then be carried out in the following steps:

1. Solve the nonlinear Equation (A16) for $x' \equiv x^{n+1}$, obtaining $v_{s,y}^{n+1/2}$, $\gamma_s^{n+1/2}$ in the process. We have found

that three or four fixed-point iterations with x^n as the initial guess typically suffice for converging below an absolute tolerance of 10^{-10} .

2. Partially update the y -coordinate as

$$y' = y^n + \Delta t \frac{u_y^{n+1/2}}{\gamma^{n+1/2} + u_y^{n+1/2} v_{s,y}^{n+1/2}}. \quad (\text{A19})$$

3. Calculate $z' \equiv z^{n+1}$ with Equation (A18).
4. Using $\mathbf{x}'$, deposit the comoving-frame current $\mathbf{J}'$.
5. Complete the update of the y -coordinate with the shearing-velocity part,

$$y^{n+1} = y' + \Delta t \frac{\gamma^{n+1/2} v_{s,y}^{n+1/2}}{\gamma^{n+1/2} + u_y^{n+1/2} v_{s,y}^{n+1/2}}. \quad (\text{A20})$$

Appendix B Differential Heating in Kinetic Shearing-box Simulations of Pair Plasmas

Kinetic shearing-box simulations are characterized by additional drift forces in the particle equations of motion, related to the orbital motion of the accretion disk. These forces are not charge-agnostic—that is, opposite charges of equal mass experience difference forces acting upon their motion. This causes a symmetry break in the dynamics of electrons and positrons; as a result, in all our simulations (both 2D and 3D), we observe that the two species in a pair plasma do not experience the same energization when solving the shearing-box equations. This appears surprising if compared to standard (non-shearing-box) PIC simulations of pair plasmas, where the two species behave essentially in the same way, and the opposite sign in the charge only serves to maintain global quasi-neutrality. Here, we prove that the differential energization we observe is not a numerical artifact, but rather a physical effect that characterizes shearing-box PIC simulations. This statement is also corroborated by several tests—outside of those presented here—performed during the development of this work, including the use of unequal density profiles for electrons and positrons as suggested by Hoshino (2013, 2015) to better respect Gauss’s law for $\mathbf{E}$. We checked that applying such a strategy does not change the resulting differential heating.

This effect stems from the kinetic shearing-box paradigm, which involves additional forces in the particle equations of motion. In the KSB-SC system, the momentum Equation (18) contains the force term

$$\mathbf{F}_{\text{SB}} = 2m\mathbf{u} \times \boldsymbol{\Omega}_0 + ms\Omega_0 u_x \hat{\mathbf{e}}_y = \begin{bmatrix} 2m\Omega_0 u_y \\ m(s-2)\Omega_0 u_x \\ 0 \end{bmatrix}, \quad (\text{B1})$$

which is related to Coriolis and centrifugal forces in the comoving frame where all quantities are measured (the same applies to the KSB-OA system). This additional force causes a drift of the particle guiding center, which is absent in typical non-shearing-box kinetic simulations. In the nonrelativistic




Figure 18. Left: trajectories of test particles (an electron and a positron) moving in the xy -plane when a uniform and constant magnetic field $\mathbf{B} = (0, 0, B_0)$ is present. The particle motion is computed by solving the equations employed in the shearing-box paradigm. Right: spatial distribution of the radial and toroidal magnetic field (top row) after the formation of channel flows in a small-box, 2D MRI simulation of size $2 \times 2\lambda_{\text{MRI}}^2$ (with $\omega_C/\Omega_0 = 120$, $\beta_0 = 624$). The corresponding spatial distribution of particle drift velocity along z (bottom row) shows that electrons are pulled toward current sheets, and positrons are pushed away.

limit, this drift has the usual expression

$$\mathbf{v}_{\text{SB}} = \frac{\langle \mathbf{F}_{\text{SB}} \rangle \times \mathbf{B}}{qB^2} c = \frac{mc}{qB^2} \begin{bmatrix} (s-2)\Omega_0 \langle u_x \rangle B_z \\ -2\Omega_0 \langle u_y \rangle B_z \\ 2\Omega_0 \langle u_y \rangle B_y - (s-2)\Omega_0 \langle u_x \rangle B_x \end{bmatrix}, \quad (\text{B2})$$

where $\langle \dots \rangle$ indicates averaging over the particle gyration. Because the sign of the drift depends on the particle charge, in general particles of equal mass but opposite charge will experience opposite drifts in all spatial directions.

We now show that the $\mathbf{v}_{\text{SB}}$ drift is directly related to the dynamics leading to the measured differential energization. First, we note that the initial conditions in our MRI simulations, $\mathbf{B} = (0, 0, B_{z,0})$, lead to a drift motion exclusively in the xy -plane. Such motion is precisely the well-known epicyclic motion of particles in orbit around massive objects (e.g., Balbus & Hawley 1998), which is here modified by the presence of a vertical magnetic field. In Figure 18 (left panel), we show a portion of the trajectory of one electron and one positron (with $m = |q| = 1$) moving according to Equation (18) in the presence of a weak magnetic field $B_{z,0}$ such that $\omega_C/\Omega_0 = 1$, with initial velocity $\mathbf{v} = (0, 0.1, 0)$. We observe that the usual particle gyromotion is here modified by a drift of the guiding center along $\pm y$ depending on the particle charge $q = \mp 1$. Moreover, the gyromotion is in itself different for the two particles, with the electron attaining higher gyrofrequency and smaller gyroradius (and vice versa for the positron).

Although the guiding-center drift is purely in the toroidal direction in this case, this differential gyromotion ought to manifest in our 2D (x - z) simulations as well, given that the particle oscillation along x occurs differently for opposite charges. This difference in particle trajectories implies that charge neutrality cannot be maintained at all times, and a net current must arise.

While the differential gyromotion is present, at least in the beginning, in all our simulations, it is arguably of minor importance in the overall system evolution. Indeed, in Section 4.5 we have shown that significant differential heating primarily manifests during and after the linear stage, when the magnetic-field geometry differs substantially from the initial (pre-instability) phase. We find that relevant differential heating arises due to the channel-flow configuration of B_x and B_y during the later stages: B_z here is heavily subdominant (see Section 4.1), and the main drift component of $\mathbf{v}_{\text{SB}}$ is along z . In Figure 18 (right panels, top row), we plot the spatial distribution of B_x and B_y (normalized over the initial field B_0) for a 2D, small-box ($2 \times 2\lambda_{\text{MRI}}^2$) simulation with parameters $\omega_C/\Omega_0 = 120$, $\beta_0 = 624$ at $t = 3.3P_0$, when channel flows are present and current sheets have formed at the channel interfaces. In the same figure (right panels, bottom row), we also plot the spatial distribution of $v_{\text{SB},z}$ for electrons and positrons taken at the same moment in time. We observe that, in the vicinity of current sheets, both species are drifting along z ; however, the particular magnetic-field configuration causes this drift to be directed *toward* current sheets in the case of electrons, and *away* from them in the case of positrons. In a standard current-sheet (e.g., Harris-like) magnetic




Figure 19. Differential heating in 2D pair-plasma MRI simulations, measured as the difference in kinetic energy between electrons and positrons normalized over the instantaneous total energy of the system. Simulations with small box size and variable ω_C/Ω_0 (left), variable box size and fixed ω_C/Ω_0 (center), and large box size and variable ω_C/Ω_0 (right) are compared.

configuration, particles drift toward current sheets, due to the usual $\mathbf{E} \times \mathbf{B}$ motion, irrespective of their charge. This global behavior is still verified in our case, but the addition of the $\mathbf{v}_{\text{SB}}$ drift now implies that charges of opposite sign drift toward current sheets at different rates.

Given that particles experience most of the energy gain via reconnection in current sheets, this differential drift motion can explain the large difference in particle energy observed between the two species. Electrons, which flow more quickly toward current sheets, are capable of gaining more energy over shorter times; this precisely corresponds to the faster electron energization rates reported in Section 4.5. Eventually, all available magnetic energy is distributed, via reconnection, unequally among the pair-plasma species, with electrons experiencing a larger energy gain.

The differential heating we report here for the first time (to the best of our knowledge) is a genuine physical effect—however, it is arguably of very minimal importance in realistic scenarios. Indeed, it can immediately be observed that the drift speed v_{SB} is inversely proportional to ω_C/Ω_0 (and indeed, for the $\omega_C/\Omega_0 = 120$ case shown in Figure 18, $v_{\text{SB}}/c \ll 1$). In realistic environments where $\omega_C/\Omega_0 \sim 10^7$, this drift would be completely negligible, and therefore no differential heating would be observed. In our simulations (as well as in any other presented in literature), however, we are bound to employ much smaller values of ω_C/Ω_0 , due to otherwise excessive computational costs. Therefore, when measuring particle heating in simulations, this effect must be taken into account, in particular in the context of future electron–ion numerical experiments, where analyzing particle heating is of main interest for observations of accreting compact objects. In such simulations, it will be of major importance to quantify the “spurious” differential heating described here, in order to correctly evaluate the energy partition between particle species.

We conclude by reporting the differential heating measured in our 2D simulation campaign presented in Section 4. This is shown in Figure 19, where the difference in electron and positron total energy $mc^2(\gamma - 1)$ is normalized over the instantaneous total energy of the system, in order to measure the relative importance of differential heating with respect to the global energetics. We observe that larger ω_C/Ω_0 values and larger boxes can ameliorate this issue substantially, keeping the differential heating below 10%. While it is clear that a larger ω_C/Ω_0 implies smaller drift speeds and therefore smaller differences in the rate at which particles flow toward current sheets, justifying the positive effect of larger box sizes is less

straightforward. Qualitatively, we can consider that smaller boxes tend to create persistent large-scale channel flows that are less likely to be destroyed by unstable modes and result in developed turbulence (as illustrated in Section 4.5); this large-scale, ordered distribution of magnetic fields is the most favorable for particles to drift undisturbed toward currents sheets at different rates. In large boxes where turbulence develops, particles are less likely to experience such long-lived, undisturbed motion, being instead continuously scattered across the turbulent structures. Differential heating can still accumulate over long times, but will reach much smaller levels.

Appendix C

Energy Conservation in the Shearing Box

Considering for simplicity the 2D shearing box in shearing coordinates, equations KSB-SC manifestly include source terms affecting the field and particle evolution. These source terms translate into energy sources that are not present in standard (non-shearing-box) PIC simulations. Evaluating energy conservation in the presence of these sources is not straightforward.

To obtain a measure of energy conservation, we compute the energy sources directly at runtime. For the fields, the energy source terms are obtained by dotting Faraday’s and Ampère’s laws with $\mathbf{B}$ and $\mathbf{E}$, respectively:

$$\mathbf{B} \cdot \frac{\partial \mathbf{B}}{\partial t} = \frac{1}{2} \frac{\partial B^2}{\partial t} = (\text{standard energy terms}) - s\Omega_0 B_x B_y, \quad (\text{C1})$$

$$\mathbf{E} \cdot \frac{\partial \mathbf{E}}{\partial t} = \frac{1}{2} \frac{\partial E^2}{\partial t} = (\text{standard energy terms}) - s\Omega_0 E_x E_y, \quad (\text{C2})$$

where in addition to standard energy-exchange terms we now have extra source terms stemming from the shearing-box paradigm. For the particles, we dot the momentum equation with $\mathbf{u}/\gamma$,

$$\frac{\mathbf{u}}{\gamma} \cdot \frac{d\mathbf{u}}{dt} = \frac{d\gamma}{dt} = (\text{standard energy terms}) + s\Omega_0 \frac{u_x u_y}{\gamma}, \quad (\text{C3})$$

which produces one extra energy source. All terms deriving from the shearing-box framework are thus proportional to Ω_0 .

To quantify the importance of the shearing-box sources, we compute each source term in the whole simulation box during a




Figure 20. Energy conservation in a 2D pair-plasma MRI simulation with parameters $\omega_c/\Omega_0 = 30$, $\beta_0 = 624$, and box size $8 \times 8\lambda_{\text{MRI}}^2$. Left: evolution of the total electromagnetic and kinetic energy compared with the modified electromagnetic and kinetic energy, where shearing-box source terms have been subtracted. The modified total energy $\bar{E}_{\text{tot}}$ is a conserved quantity in the absence of numerical errors. Right: evolution during the run of the error on the conserved $\bar{E}_{\text{tot}}$ relative to the instantaneous total energy E_{tot} (the latter includes source terms and is not conserved).

representative 2D pair-plasma MRI simulation with parameters $\omega_c/\Omega_0 = 30$, $\beta_0 = 624$, and box size $8 \times 8\lambda_{\text{MRI}}^2$. The source terms are then integrated in time, to obtain the net energy variation associated with these terms. Subtracting the net contribution of each term from the electromagnetic and particle energy produces a new measure of energy that only accounts for standard energy-exchange channels (i.e., without external energy sources related to the shearing box). In simulations, the total energies including the source terms are given by

$$B^2(t) = \sum_g B_g^2(t) \Delta V_g, \quad (\text{C4})$$

$$E^2(t) = \sum_g E_g^2(t) \Delta V_g, \quad (\text{C5})$$

for the electromagnetic energy of all grid cells g (with cell volume ΔV_g), and

$$K(t) = \sum_p m_p c^2 (\gamma_p(t) - 1), \quad (\text{C6})$$

for the kinetic energy of all particles p . By subtracting the shearing-box terms, we find that the *modified* field energies (i.e., without shearing-box sources) are

$$\frac{1}{2} \bar{B}^2(t) \equiv \frac{1}{2} B^2(t) - \int_0^t dt \sum_g \Delta V_g (-s\Omega_0 B_{x,g}(t) B_{y,g}(t)), \quad (\text{C7})$$

$$\frac{1}{2} \bar{E}^2(t) \equiv \frac{1}{2} E^2(t) - \int_0^t dt \sum_g \Delta V_g (-s\Omega_0 E_{x,g}(t) E_{y,g}(t)), \quad (\text{C8})$$

and the modified particle energy is

$$\bar{K}(t) \equiv K(t) - \int_0^t dt \sum_p s\Omega_0 \frac{u_{x,p}(t) u_{y,p}(t)}{\gamma_p(t)}. \quad (\text{C9})$$

With these definitions, the modified energies are such that

$$\partial_t \bar{E}_{\text{tot}} = \partial_t \left(\frac{\bar{B}^2 + \bar{E}^2}{8\pi} + \bar{K} \right) = 0, \quad (\text{C10})$$

i.e., the sum of the modified electromagnetic and kinetic energy integrated over the simulation box remains constant in time (as in standard PIC simulations), in the absence of numerical errors. Note that, conversely, the total energy including sources $E_{\text{tot}} \equiv (B^2 + E^2)/(8\pi) + K$ is not constant: the sources continuously inject energy into the system.

Figure 20 (left panel) shows the evolution in time of the modified electromagnetic and kinetic energy (integrated over the box) as well as of the sum of the two. The plot shows that the modified electromagnetic energy becomes *negative*: this is because the external energy input for the electromagnetic fields is much larger than the initial energy, and subtracting this energy source results in observing exclusively the standard process of magnetic-energy dissipation, e.g., via reconnection. Because of the external source, more magnetic energy than initially available can be dissipated, and the modified electromagnetic energy can fall below zero. The decrease in electromagnetic energy corresponds almost exactly to an increase in particle energy via dissipation processes. The total (electromagnetic plus kinetic) modified energy is much smaller than its individual components, owing to the fact that the modified electromagnetic energy becomes negative.

The right panel of Figure 20 shows a measure of energy conservation calculated as $|\bar{E}_{\text{tot}}(t) - \bar{E}_{\text{tot}}(0)|/E_{\text{tot}}(t)$. This quantity is such that the numerator expresses the deviation from exact energy conservation in terms of the absolute error on the conserved quantity $\bar{E}_{\text{tot}}(t)$; normalizing over the instantaneous total energy (including sources) serves to provide a measure of how important the numerical errors are compared to the actual amount of energy in the system at a certain time. We see that the error is at most of order 10^{-2} , i.e., numerical energy dissipation is negligible throughout a typical run.

ORCID iDs

Fabio Bacchini <https://orcid.org/0000-0002-7526-8154>
 Lev Arzamasskiy <https://orcid.org/0000-0002-5263-9274>
 Vladimir Zhdankin <https://orcid.org/0000-0001-8363-839X>
 Gregory R. Werner <https://orcid.org/0000-0001-9039-9032>
 Mitchell C. Begelman <https://orcid.org/0000-0003-0936-8488>
 Dmitri A. Uzdensky <https://orcid.org/0000-0001-8792-6698>

References

- Alt, A., & Kunz, M. W. 2019, *JPhI*, **85**, 175850102
- Arendt, P. N., Jr. 1998, arXiv:astro-ph/9801194
- Arzamasskiy, L., Kunz, M. W., Chandran, B. D. G., & Quataert, E. 2019, *ApJ*, **879**, 53
- Arzamasskiy, L., Kunz, M. W., Squire, J., Quataert, E., & Schekochihin, A. A. 2022, arXiv:2207.05189
- Bacchini, F., Amaya, J., & Lapenta, G. 2019, *JPhCS*, **1225**, 012011
- Balbus, S. A., & Hawley, J. F. 1991, *ApJ*, **376**, 214
- Balbus, S. A., & Hawley, J. F. 1998, *RvMP*, **70**, 1
- Ball, D., Sironi, L., & Ozel, F. 2018, *ApJ*, **862**, 80
- 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
- Cerri, S. S., Arzamasskiy, L., & Kunz, M. W. 2021, *ApJ*, **916**, 120
- Cerutti, B., Werner, G. R., Uzdensky, D. A., & Begelman, M. C. 2013, *ApJ*, **770**, 147
- Chael, A., Rowan, M., Narayan, R., Johnson, M., & Sironi, L. 2018, *MNRAS*, **478**, 4
- Chandran, B. D. G., Li, B., Rogers, B. N., Quataert, E., & Germaschewski, K. 2010, *ApJ*, **720**, 503
- Chandrasekhar, S. 1960, *PNAS*, **46**, 253
- Chandrasekhar, S., Kaufman, A. N., & Watson, K. M. 1958, *RSPSA*, **245**, 435
- Chen, J., & Palmadesso, P. 1984, *PhFI*, **27**, 1198
- Comisso, L., & Sironi, L. 2018, *PhRvL*, **121**, 255101
- Comisso, L., & Sironi, L. 2019, *ApJ*, **886**, 122
- Daly, R. A., Djorgovski, S. G., Freeman, K. A., et al. 2008, *ApJ*, **677**, 1
- Dexter, J., Scepi, N., & Begelman, M. C. 2021, *ApJL*, **919**, L20
- EHT Collaboration 2019, *ApJL*, **875**, L5
- EHT Collaboration 2022, *ApJL*, **930**, L16
- Fromang, S., & Papaloizou, J. 2007, *A&A*, **476**, 1113
- Fromang, S., Papaloizou, J., Lesur, G., & Heinemann, T. 2007, *A&A*, **476**, 1123
- Gardiner, T. A., & Stone, J. M. 2005, in AIP Conf. Proc. 784, Magnetic Fields in the Universe: From Laboratory and Stars to Primordial Structures, ed. E. M. de Gouveia Dal Pino, G. Lugones, & A. Lazarian (Melville, NY: AIP), 475
- Goedbloed, H., & Keppens, R. 2022, *ApJS*, **259**, 65
- Goldreich, P., & Lynden-Bell, D. 1965, *MNRAS*, **130**, 125
- Goldreich, P., & Sridhar, S. 1995, *ApJ*, **438**, 763
- Goldreich, P., & Tremaine, S. 1978, *ApJ*, **222**, 850
- Goodman, J., & Xu, G. 1994, *ApJ*, **432**, 213
- Gressel, O., & Ziegler, U. 2007, *CoPhC*, **176**, 652
- Guo, F., Li, H., Daughton, W., & Liu, Y. 2014, *PhRvL*, **113**, 155005
- Hankla, A., Dexter, J., & Scepi, N. 2022, *MNRAS*, **515**, 775
- Hawley, J. F., Gammie, C. F., & Balbus, S. A. 1995, *ApJ*, **440**, 742
- Hill, G. W. 1878, *AmJM*, **1**, 5
- Hillas, A. M. 1984, *ARA&A*, **22**, 425
- Hirabayashi, K., & Hoshino, M. 2017, *ApJ*, **842**, 36
- Hirai, K., Katoh, Y., Terada, N., & Kawai, S. 2018, *ApJ*, **853**, 174
- Hoppeck, I. W., Chandran, B. D. G., Klein, K. G., Mallet, A., & Verscharen, D. 2018, *JPhI*, **84**, 905840615
- Hoshino, M. 2013, *ApJ*, **773**, 118
- Hoshino, M. 2015, *PhRvL*, **114**, 061101
- Hoshino, M. 2020, *ApJ*, **900**, 66
- Howes, G. G. 2010, *MNRAS*, **409**, L104
- Inchingolo, G., Grismayer, T., Loureiro, N. F., Fonseca, R. A., & Silva, L. O. 2018, *ApJ*, **859**, 149
- Isenberg, P. A. 2004, *JGRA*, **109**, A03101
- Isenberg, P. A., & Vasquez, B. J. 2007, *ApJ*, **668**, 546
- Kaisig, M. 1989, *A&A*, **218**, 102
- Kawazura, Y., Barnes, M., & Schekochihin, A. A. 2019, *PNAS*, **771**, 116
- Kawazura, Y., Schekochihin, A. A., Barnes, M., et al. 2020, *PhRvX*, **10**, 041050
- Kawazura, Y., Schekochihin, A. A., Barnes, M., Dorland, W., & Balbus, S. A. 2021, *JPhI*, **88**, 905880311
- Kennel, C. F., & Engelmann, F. 1966, *PhFI*, **9**, 2377
- Kunz, M. W., Abel, I. G., Klein, K. G., & Schekochihin, A. A. 2018, *JPhI*, **84**, 715840201
- Kunz, M. W., Schekochihin, A. A., Chen, C. H. K., Abel, I. G., & Cowley, S. C. 2015, *JPhI*, **81**, 325810501
- Kunz, M. W., Schekochihin, A. A., & Stone, J. M. 2014a, *PhRvL*, **112**, 205003
- Kunz, M. W., Stone, J. M., & Bai, X. 2014b, *JCoPh*, **259**, 154
- Kunz, M. W., Stone, J. M., & Quataert, E. 2016, *PhRvL*, **117**, 235101
- Langdon, A. B. 1991, *CPC*, **70**, 447
- Latter, H. N., Lesaffre, P., & Balbus, S. A. 2009, *MNRAS*, **394**, 715
- Leiler, G., & Rezzolla, L. 2006, *PhRvD*, **73**, 044001
- Ley, F., Riquelme, M., Sironi, L., Verscharen, D., & Sandoval, A. 2019, *ApJ*, **880**, 100
- Loureiro, N. F., & Boldyrev, S. 2018, *ApJL*, **866**, L14
- Loureiro, N. F., & Boldyrev, S. 2020, *ApJ*, **890**, 55
- Mallet, A., Schekochihin, A. A., & Chandran, B. D. G. 2017, *MNRAS*, **468**, 4862
- Marder, B. 1987, *JCoPh*, **68**, 48
- Masada, Y., & Sano, T. 2008, *ApJ*, **689**, 1234
- Matsumoto, R., & Tajima, T. 1995, *ApJ*, **445**, 767
- McChesney, J. M., Stern, R. A., & Bellan, P. M. 1987, *PhRvL*, **59**, 1436
- Narayan, R., Goldreich, P., & Goodman, J. 1987, *MNRAS*, **228**, 1
- Parker, E. N. 1958, *PhRv*, **109**, 1874
- Pastorello, A., Cappellaro, E., Inserra, C., et al. 2013, *ApJ*, **767**, 1
- Pessah, M. E., & Goodman, J. 2009, *ApJL*, **698**, L72
- Pritchett, P. L., Coroniti, F. V., & Decyk, V. K. 1996, *JGR*, **101**, A12
- Quataert, E., Dorland, W., & Hammett, G. W. 2002, *ApJ*, **577**, 524
- Regev, O., & Umurhan, O. M. 2008, *A&A*, **481**, 21
- Ressler, S. M., Tchekhovskoy, A., Quataert, E., Chandra, M., & Gammie, C. F. 2015, *MNRAS*, **454**, 1848
- Ripperda, B., Bacchini, F., Teunissen, J., et al. 2018, *ApJS*, **235**, 21
- Riquelme, M. A., Quataert, E., Sharma, P., & Spitkovsky, A. 2012, *ApJ*, **755**, 50
- Riquelme, M. A., Quataert, E., & Verscharen, D. 2015, *ApJ*, **800**, 27
- Riquelme, M. A., Quataert, E., & Verscharen, D. 2016, *ApJ*, **824**, 123
- Rosenbluth, M. N. 1956, LANL Report LA-2030
- Ryan, B. R., Ressler, S. M., Dolence, J. C., et al. 2017, *ApJL*, **844**, L24
- Sagdeev, R. Z., & Shafranov, V. D. 1960, *JETP*, **181**
- Sano, T., & Inutsuka, S. 2001, *ApJL*, **561**, L179
- Sano, T., & Stone, J. M. 2002, *ApJ*, **577**, 534
- Scepi, N., Dexter, J., & Begelman, M. C. 2022, *MNRAS*, **511**, 3536
- Schekochihin, A. A., Cowley, S. C., Dorland, W., et al. 2009, *ApJS*, **182**, 310
- Schiff, L. I. 1939, *PNAS*, **25**, 391
- Shakura, N. I., & Sunyaev, R. A. 1973, *A&A*, **500**, 33
- Shapiro, V. D., & Shevchenko, V. I. 1964, *JETP*, **18**, 1109
- Sharma, P., Hammett, G. W., Quataert, E., & Stone, J. M. 2006, *ApJ*, **637**, 952
- Sironi, L., & Spitkovsky, A. 2014, *ApJL*, **783**, L21
- Squire, J., Meyrand, R., Kunz, M. W., et al. 2022, *NatAs*, **6**, 715
- Stanzione, D., West, J., Evans, R. T., et al. 2020, Frontera: The Evolution of Leadership Computing at the National Science Foundation (New York, NY: Association for Computing Machinery), 106
- Stone, J. M., & Gardiner, T. A. 2010, *ApJS*, **189**, 142
- Stone, J. M., Hawley, J. F., Gammie, C. F., & Balbus, S. A. 1996, *ApJ*, **463**, 656
- Tirunch, A. T. 2020, arXiv:2002.06976
- Towns, J., Cockerill, T., Dahan, M., et al. 2014, *CSE*, **16**, 62
- Vedenov, A. A., & Sagdeev, R. Z. 1958, *SPhD*, **3**, 278
- Velikhov, E. P. 1959, *JETP*, **36**, 1938
- Villasenor, J., & Buneman, O. 1992, *CPC*, **69**, 306
- Walker, J., & Boldyrev, S. 2017, *MNRAS*, **470**, 2653
- Walker, J., Lesur, G., & Boldyrev, S. 2016, *MNRAS*, **457**, L39
- Werner, G. R., & Uzdensky, D. A. 2021, *JPhI*, **87**, 905870613
- Werner, G. R., Uzdensky, D. A., Begelman, M. C., Cerutti, B., & Nalewajko, K. 2018, *MNRAS*, **473**, 4
- Werner, G. R., Uzdensky, D. A., Cerutti, B., Nalewajko, K., & Begelman, M. C. 2016, *ApJL*, **816**, 1
- Winarto, H. W., & Kunz, M. W. 2022, *JPhI*, **88**, 905880210
- Wong, K., Zhdankin, V., Uzdensky, D. A., Werner, G. R., & Begelman, M. C. 2020, *ApJL*, **893**, L7
- Zenitani, S., & Hoshino, M. 2007, *ApJ*, **670**, 702
- Zhdankin, V., Uzdensky, D. A., & Boldyrev, S. 2015, *PhRvL*, **114**, 065002
- Zhdankin, V., Uzdensky, D. A., & Kunz, M. W. 2021, *ApJ*, **908**, 71
- Zhdankin, V., Uzdensky, D. A., Perez, J. C., & Boldyrev, S. 2013, *ApJ*, **771**, 124
- Zhdankin, V., Uzdensky, D. A., Werner, G. R., & Begelman, M. C. 2018a, *MNRAS*, **474**, 2
- Zhdankin, V., Uzdensky, D. A., Werner, G. R., & Begelman, M. C. 2018b, *ApJL*, **867**, L18
- Zhdankin, V., Uzdensky, D. A., Werner, G. R., & Begelman, M. C. 2019, *PhRvL*, **122**, 055101
- Zhdankin, V., Werner, G. R., Uzdensky, D. A., & Begelman, M. C. 2017, *PhRvL*, **118**, 055103
- Zweibel, E. G., & Yamada, M. 2009, *ARA&A*, **47**, 1