

RESEARCH ARTICLE | JANUARY 28 2025

Application of mesh refinement to relativistic magnetic reconnection

Revathi Jambunathan ; Henry Jones ; Lizzette Corrales ; Hannah Klion ; Michael E. Rowan ; Andrew Myers ; Weiqun Zhang ; Jean-Luc Vay 



Phys. Plasmas 32, 013905 (2025)

<https://doi.org/10.1063/5.0233583>

 CHORUS



Articles You May Be Interested In

Laser-plasma interactions with a Fourier-Bessel particle-in-cell method

Phys. Plasmas (March 2016)

Accurate modeling of plasma acceleration with arbitrary order pseudo-spectral particle-in-cell methods

Phys. Plasmas (March 2017)

Stable discrete representation of relativistically drifting plasmas

Phys. Plasmas (October 2016)



Physics of Plasmas

Special Topics Open
for Submissions

[Learn More](#)

Application of mesh refinement to relativistic magnetic reconnection

Cite as: Phys. Plasmas **32**, 013905 (2025); doi: [10.1063/5.0233583](https://doi.org/10.1063/5.0233583)

Submitted: 15 August 2024 · Accepted: 27 December 2024 ·

Published Online: 28 January 2025



View Online



Export Citation



CrossMark

Revathi Jambunathan,^{1,a)} Henry Jones,¹ Lizzette Corrales,^{1,b)} Hannah Klion,¹ Michael E. Rowan,²
Andrew Myers,¹ Weiqun Zhang,¹ and Jean-Luc Vay¹

AFFILIATIONS

¹Lawrence Berkeley National Laboratory, 1 Cyclotron Road, Berkeley, California 94720, USA

²Advanced Micro Devices, Inc., Santa Clara, California 95054, USA

^{a)} Author to whom correspondence should be addressed: rjambunathan@lbl.gov

^{b)} Present address: Cornell University, Ithaca, New York 14850, USA.

ABSTRACT

During relativistic magnetic reconnection, antiparallel magnetic fields undergo a rapid change in topology, releasing a large amount of energy in the form of non-thermal particle acceleration. This work explores the application of mesh refinement to 2D reconnection simulations to efficiently model the inherent disparity in length-scales. We have systematically investigated the effects of mesh refinement and determined necessary modifications to the algorithm required to mitigate non-physical artifacts at the coarse–fine interface. We have used the ultrahigh-order pseudo-spectral analytical time-domain Maxwell solver to analyze how its use can mitigate the numerical dispersion that occurs with the finite-difference time-domain (or “Yee”) method. Absorbing layers are introduced at the coarse–fine interface to eliminate spurious effects that occur with mesh refinement. We also study how damping the electromagnetic fields and current density in the absorbing layer can help prevent the non-physical accumulation of charge and current density at the coarse–fine interface. Using a mesh refinement ratio of 8 for two-dimensional magnetic reconnection simulations, we obtained good agreement with the high-resolution baseline simulation, using only 36% of the macroparticles and 71% of the node-hours needed for the baseline. The methods presented here are especially applicable to 3D systems where higher memory savings are expected than in 2D, enabling comprehensive, computationally efficient 3D reconnection studies in the future.

© 2025 Author(s). All article content, except where otherwise noted, is licensed under a Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>). <https://doi.org/10.1063/5.0233583>

I. INTRODUCTION

Magnetic reconnection is a fundamental process where the topology of magnetic fields rapidly rearranges (they break and reconnect) converting energy stored in the stressed regions of strong magnetic fields to non-thermal particle energy. This process is often invoked to explain particle energization leading to high-energy emissions in a wide range of plasma systems. These systems include solar flares, extreme astrophysical systems such as pulsars, active galactic nuclei, gamma-ray bursts, and black hole jets as well as laboratory astrophysics and even magnetic fusion devices.¹ Studying the plasma kinetic effects underpinning particle energization is critical to understanding high-energy emissions from astrophysical systems. Therefore, a first-principles approach is required to capture the complex interaction of charged particles with the electromagnetic fields in these systems.

We use a fully kinetic, explicit, electromagnetic particle-in-cell (PIC) approach^{2–4} to study relativistic reconnection physics. Significant

work has been done in the past two decades investigating the kinetic effects that are important to 2D relativistic reconnection in collisionless pair plasmas^{5–10} as well as electron-ion plasmas.¹¹ These studies showed that the high-aspect ratio current sheets become unstable to the tearing mode instability, leading to the formation of trapped plasma islands, called plasmoids, that undergo merging and secondary reconnection during the non-linear phase. Detailed investigations on the mechanisms that drive the onset of reconnection and phases of particle energization have also been performed.^{12–14} The particle energy spectra due to reconnection show hard power laws that extend to high energies.^{10,13,15–17} These spectra can then be combined with radiation models to predict observational signatures of reconnection in astrophysical systems.^{18,19} For extreme astrophysical processes, additional quantum electrodynamics (QED) processes, such as pair production and synchrotron radiation, are important. PIC simulation studies including these effects have been reported recently^{20–22} for 2D systems.

While substantial work has been done in 2D, similar detailed investigations of 3D systems are lacking due to the computational intensity of the PIC method. Nevertheless, a few 3D studies performed recently reveal that in addition to the tearing-mode, drift-kink instabilities dominate the evolution of the current sheet in the out-of-plane dimension,^{8,16,23,24} significantly complicating the picture and allowing particles to escape plasmoids. Escaped particles can reenter reconnecting regions and become multiply energized.²⁵ Cerutti *et al.* investigated the dispersion relations of the tearing mode (in 2D) and the drift-kink mode (in 3D). Numerical investigations performed by Werner and Uzdensky²⁴ showed that in systems with a large guide-field (i.e., out-of-plane magnetic field component), 3D instabilities are suppressed and 2D simulations can be used as a proxy to study these systems. However, with small guide-fields, 3D effects become important and cannot be accurately represented by 2D simulations. More detailed investigations of 3D systems with radiation physics and QED effects have not yet been conducted, especially with large magnetization, mainly due to the computational expense of the PIC method.

Development and application of advanced numerical algorithms can improve computational efficiency and thereby enable detailed studies of 3D reconnection systems. In Klion *et al.*,²⁶ we used the ultrahigh-order pseudo-spectral analytical time-domain (PSATD) Maxwell solver and found that for 2D uniform grid simulations with same spatial resolution, it can also accurately capture reconnection similar to the widely used finite-difference time-domain (FDTD) Yee solver. By contrast, it is not restricted by a Courant–Friedrichs–Lewy (CFL) stability criterion in theory, and our results showed excellent agreement up to CFL = 1.6 (i.e., $c\Delta t/\Delta x = 1.6$, where c is the speed of light, and Δt and Δx are respectively the time step and the mesh size of the simulation in each direction). The simulations were performed on graphics processing units (GPUs) using the exascale-capable WarpX code, and the reconnection rate as well as particle acceleration obtained from our simulations agreed well with the results from the literature.

Most PIC simulations reported in the literature have used a uniform grid. Note that the grid resolution for the PIC method must resolve the local skin depth in order to accurately capture the plasma kinetic effects. Magnetic reconnection involves disparate length scales wherein the plasma density is often higher in the current sheet (by a factor of 5 at minimum) than in the upstream regions. In order to capture these kinetic effects, the grid resolution in the current sheet must be less than the local skin depth. However, uniform resolution grids resolve even the region upstream of the current sheet, where the plasma density is lower and the corresponding skin depth is larger. Note that the PIC method also requires at-least tens of particles per cell to obtain statistically accurate description of the non-thermal particle acceleration. Mesh refinement (MR) is therefore a natural choice to alleviate the memory requirement and improve computational efficiency. However, it has not been explored for relativistic reconnection to the best of our knowledge.

In this paper, we apply mesh refinement (MR) to relativistic magnetic reconnection using WarpX to investigate the impact of using different resolutions for the high-density current sheet and low-density upstream regions. The static MR strategy we leverage was first developed by^{27,28} and has been previously applied to study particle accelerators and laser plasma interactions.^{30–32} However, it has not yet been applied to systems such as relativistic magnetic reconnection, which

pose a unique set of challenges due to the high current density and large flux of particles crossing the coarse–fine interface (from upstream toward the x-points in the current sheet). While the implementation of the MR method is generalized to 3D in the code, the study of the method is performed in 2D in this work, as it enables the exploration of the key issues and their mitigation more effectively without lack of generality (i.e., the issues and mitigations identified in 2D extend readily to 3D). Previously validated, high-resolution, uniform grid 2D simulations²⁶ serve as the baseline to compare with the MR simulations. To study the accuracy, we compare current sheet evolution, energy conservation and conversion, and particle spectra. We also demonstrate the advantage of using a spectral method, such as PSATD, instead of the widely used FDTD method, thanks to its ultra-low numerical dispersion.

The rest of the paper is organized as follows: Sec. II briefly describes the idealized Harris-like current sheet setup and the MR method that we used for the relativistic reconnection simulations along with our choices of numerical parameters. In Sec. III, we first present results obtained from coarsening the spatial resolution of uniform grid simulations, then from applying static MR patches surrounding the current sheet. The current sheet evolution, energy conversion, and particle energization are compared with the uniform resolution simulations, and speedups are given. Section IV discusses the effect of solver and parameter choices for the MR reconnection simulations that were discussed in this work. Finally, conclusions are given in Sec. V together with suggestions of improvements of the method discussed here that could be included as part of future work.

II. SIMULATION SETUP AND MESH REFINEMENT STRATEGY

A. Harris-like current sheets

The simulations shown in this paper are of two-dimensional, pair-plasma, relativistic magnetic reconnection, starting from Harris-like current sheets³² on a periodic domain. Since the background magnetic field changes sign at the current sheets, two sheets are needed to ensure periodicity of the magnetic field direction. This section summarizes the initial configuration; detailed description of the derivations of spatial profiles for all values discussed are given in our prior work, Klion *et al.*²⁶ Unless otherwise indicated, the simulations in this paper use the same initial configuration and simulation parameters as used previously. Table I summarizes relevant simulation parameters in scaled units and corresponding values in SI units is provided in Table IV in Appendix A. Code and input files to replicate our simulations and results are available online.³³

The upstream, unreconnected, magnetic field is $\mathbf{B} = \pm B_0 \hat{z}$. The upstream magnetic field strength B_0 sets the inverse upstream electron gyrofrequency, $\omega_c^{-1} \equiv m_e/(eB_0)$, which is our main time-scale. Here, m_e is the electron mass and e is the elementary charge. Our base unit of length is the Larmor radius, $\rho_c = c\omega_c^{-1}$, where c is the speed of light. The 2D simulation domain extends from $-L_x$ to L_x in x and $-L_z$ to L_z in z , with periodic boundary conditions in both directions. We set $L_x = 2508 \rho_c$ and $L_z = 1254 \rho_c$. The top and bottom current sheets are centered at $x = \pm x_c \equiv \pm L_x/2$ with half-width $\delta = 12.15 \rho_c$.

The spatial distribution of number density, $n(x)$ and bulk velocity $\beta(x)$ is as follows:

TABLE I. Physical parameters and symbols common to all of our simulations. Quantities marked with * are freely chosen; others are derived.

| Parameter | Symbol | Value |
|-----------------------------------|-------------|----------------|
| Background (cold) magnetization* | σ | 30 |
| Background temperature* | θ_b | 0.15 |
| Current sheet half-width* | δ | $12.15 \rho_c$ |
| Current sheet skin depth | λ_e | $2.45 \rho_c$ |
| Current sheet overdensity factor* | n_d/n_b | 5 |
| Current sheet velocity | β_0 | $0.22 c$ |
| Current sheet temperature | θ_d | 1.57 |
| Domain half-width (x)* | L_x | $2508 \rho_c$ |
| Domain half-width (z)* | L_z | $1254 \rho_c$ |

$$n(x) = n_b + (n_d - n_b) \left(\operatorname{sech} \frac{x + x_c}{\delta} + \operatorname{sech} \frac{x - x_c}{\delta} \right), \quad (1)$$

$$\boldsymbol{\beta}(x) = \beta(x) \hat{y}, \quad (2)$$

$$= \beta_0 \left(\operatorname{sech} \frac{x + x_c}{\delta} - \operatorname{sech} \frac{x - x_c}{\delta} \right) \hat{y}. \quad (3)$$

Positrons and electrons are initialized with bulk velocities $\boldsymbol{\beta}(x)$ and $-\boldsymbol{\beta}(x)$, respectively. The bulk velocity at the center of the current sheet is β_0 , and \hat{y} is the unit vector parallel to the out-of-plane y axis. Both species are initialized with density $n(x)$. The number density in the upstream region is n_b and n_d in the current sheet. Solving for Ampère's law gives the expression for the initial spatial variation of magnetic field, given by

$$B_z(x) = -\frac{2B_0}{\left(\frac{\pi}{2} + \frac{n_d}{n_b} - 1\right)} \times \left(\arctan \tanh \frac{x + x_c}{2\delta} - \arctan \tanh \frac{x - x_c}{2\delta} - \frac{\pi}{4} + \frac{1}{2} \left[\frac{n_d}{n_b} - 1 \right] \times \left[\tanh \frac{x + x_c}{\delta} - \tanh \frac{x - x_c}{\delta} - 1 \right] \right), \quad (4)$$

where

$$B_0 = 2\mu_0 e n_b \beta_0 c \delta \left(\frac{\pi}{2} + \frac{n_d}{n_b} - 1 \right). \quad (5)$$

The upstream “cold” magnetization, $\sigma \equiv B_0^2 / (\mu_0 n_b m_e c^2) = 30$. At the center of the current sheets, the magnetic field decreases to zero, and therefore, the magnetic pressure is also negligible. To keep the system in pressure equilibrium, the gas pressure in the current sheet must compensate for this loss of magnetic pressure. To accomplish this, the current sheet plasma is both denser and hotter than the upstream plasma. The temperature profile is given as

$$\theta(x) = \frac{\sigma (4 + \eta) / \eta - [B_z(x) / B_0]^2}{4 [n(x) / n_b] \sqrt{1 - \beta(x)^2}}. \quad (6)$$

The current sheet number density per species is chosen to be $n_d = 5n_b$, where n_b is the upstream number density per species. This sets the dimensionless temperature at the center of the current sheets

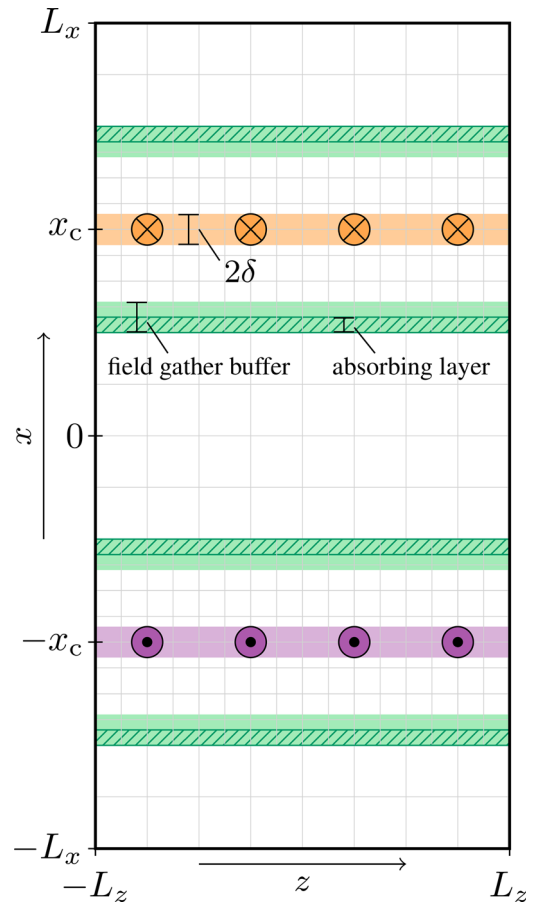


FIG. 1. Initial configuration for two-dimensional relativistic magnetic reconnection with mesh refinement. The initial current sheets (orange and purple) have half-widths of δ and are located at $x = \pm x_c$. The difference in grid density demonstrates the location of the two refinement patches. Along the x axis boundaries of each refinement patch, there is an absorbing layer (dark green hatched) and a larger field gather buffer (FGB) region (green highlight). The grid lines and other features on this schematic are not shown to scale.

[evaluating Eq. (6) at $x = x_c$] to be $\theta_d = 1.57$. The electron–positron plasma is initialized at the start of the simulation by sampling momenta from a Maxwell–Jüttner distribution at the local temperature and with the local bulk velocity.³⁴ Details of the derivation of these expressions are given in Klion *et al.*²⁶ A schematic of this configuration is shown in Fig. 1.

To initiate reconnection, a one percent sinusoidal perturbation is applied to the vector potential \mathbf{A} , which reduces the magnetic pressure at $z = 0$ just above and below the current sheets. The magnetic field at initialization is the curl of this perturbed vector potential, so $\nabla \cdot \mathbf{B} = 0$ at the start of the simulation. The functional form of this perturbation is given in our previous work.²⁶ The numerical methods used in this work preserve this property.

B. Mesh refinement method

The mesh refinement method implemented in WarpX is briefly described here; more details can be found in Vay *et al.*²⁸ The




FIG. 2. Schematic to illustrate the static mesh refinement algorithm, with a parent grid at level 0 (L_0) and a refined region at level 1 (L_1). The refined region involves three patches, namely, a fine patch, a coarse patch, and an auxiliary patch. The fine and auxiliary patches have L_1 resolution while the coarse patch has the same resolution as the level below it, i.e., the parent grid. Maxwell’s equations are solved on the fine and coarse patches of L_1 , and these regions are terminated by absorbing layers indicated by the green bands surrounding these patches.

terminology is introduced in Fig. 2 using a simple example with one level of refinement. The coarsest level, denoted L_0 , is called the parent grid. The refinement region built on top of the parent grid is referred to as level 1 (L_1). The overlapping area is delimited by the dotted lines on the parent grid in Fig. 2. The core principle of this method relies on the linearity of Maxwell’s equations to separate the coarse and fine resolution solves of Maxwell’s equations on every level. The total electromagnetic fields on an “auxiliary” patch is constructed using substitution that corrects the high-resolution solutions on the fine level with the long-range effects captured on the coarse level. The particles gather field from the “auxiliary” patch on the finest level that corresponds to their position. That is, if their position is within the fine patch, then particles gather fields from the “auxiliary” patch on L_1 , else they will gather fields from the parent grid. In order to achieve this, every refinement region has three separate grids, referred to as fine patch, coarse patch, and auxiliary patch. The fine and auxiliary patches have the same resolution as L_1 , and the grid resolution for the coarse patch on L_1 is that of the level below it, in this case the parent grid, L_0 . Note that the parent grid does not require an auxiliary patch and the electromagnetic solution on the parent grid is the final solution for the parent grid level, L_0 . The particles deposit their current on the fine patch of the level that corresponds to their positions, i.e., if the particle position overlaps with the refinement patch, then the current is deposited on the fine patch of L_1 . The current density on the fine patch of L_1 is then interpolated to the corresponding coarse patch on the same level (L_1), and copied from coarse patch of L_1 to the parent grid L_0 , in the region overlapping with the refinement patch (and delimited by the dotted line). Maxwell’s equations are solved independently on the fine and coarse patches of L_1 and on the parent grid, L_0 . The fine and coarse patches on level, L_1 , are terminated with absorbing boundary conditions, shown by green bands in Fig. 2. For the parent grid, the physical boundary condition at the edge of the domain is applied. Note that in the plasma accelerator simulations that use Warp/WarpX,^{28,35} a perfect matching layer (PML) was used to damp (or absorb) the electromagnetic signals leaving the fine and coarse patches of L_1 .^{27,36} This treatment at the coarse–fine boundary works well for applications like particle accelerators, where the plasma in the coarse–fine interface is nearly vacuum. However, for applications such as

reconnection where the plasma is dense at the coarse–fine interface and the current density can be large, and numerical artifacts were found to build over time during reconnection with PML. In this work, we have implemented a new absorbing layer feature that is applied inside the refinement patch, as shown by the hashed region in Fig. 1 to prevent numerical artifacts that were observed with PML. The effect of using this new feature is discussed later in Sec. IV B.

After the Maxwell solve, the full electromagnetic solution on the auxiliary patch of L_1 is obtained by the following substitution:^{28,35}

$$F_1^a(p) = F_1^f(p) + I^a[F_0(q) - F_1^c(q)](p), \quad (7)$$

where F is the field and p and q are grid points at the fine and coarse resolutions, respectively. The subscript denotes the level of the field, and the superscripts, f , c , and a refer to the fine, coarse, and auxiliary patch, respectively. $I[\cdot](p)$ is the operator that interpolates fields from the coarse resolution grid points (q) to the grid points on the auxiliary patch of the fine resolution (p) of level L_1 . For the interpolation, the solution from the coarse patch of L_1 (F_1^c) is first subtracted from the solution of the underlying region from the parent grid (F_0). This solution described by $F_0(q) - F_1^c(q)$ at grid points (q) on the coarse level is then interpolated to the higher-resolution auxiliary patch grid points (p), to which the solution from the fine patch, $F_1^f(p)$ is also added. Such a substitution followed by interpolation given by $I^a[F_0(q) - F_1^c(q)](p)$ and addition with the fine-resolution solution ($F_1^f(p)$), ensures that the fine-resolution solutions (terminated at the coarse/fine boundary) capture the short-wavelength solution and are also corrected to include the long-range interactions captured by the parent grid.²⁸ The region s , delimited by a red line on the auxiliary patch of L_1 , is offset from the coarse–fine boundary to ensure that particles do not gather the fine-level solution from the “auxiliary” patch of L_1 close to the edge of the patch, in order to avoid spurious effects that occur when the particles leave or enter the refinement patch.^{28,37} This buffer region where particles do not gather fields close to the patch edge is referred to as the field gather buffer region, shown in green in Fig. 1.

Simulations of relativistic magnetic reconnection with mesh refinement were initialized as follows. The magnetic field for the Harris-sheet setup, along with the perturbation, is initialized only on the parent grid. As previously mentioned, an absorbing layer is used instead of PMLs to terminate the patches. We also deposit the current and damp it with the same damping profile in the absorbing region at the edges of the fine patch, consistent with the electromagnetic fields.³¹ The buffer gather region, starting from the edge of the absorbing layer, was chosen to be of similar physical width as the absorbing layer on the fine patch and is indicated by the hashed region in Fig. 1. The choice of these parameters for the MR simulations performed in this work is provided in Sec. III.

III. RECONNECTION SIMULATIONS WITH MESH REFINEMENT

A. Effect of mesh refinement on current density and reconnection evolution

1. Effect of coarsening uniform grid resolution

We first perform a baseline uniform grid simulation for the two-dimensional Harris-sheet setup described in Sec. II A, with a grid size of 4096×2048 cells, which has a resolution of two cells per current-

sheet skin depth ($\Delta x = \Delta z = \lambda_e/2$). Note that the upstream Larmor radius is also resolved with nearly one cell (0.8 cells), and the upstream and current sheet Debye lengths are resolved by 1.25 and 1.8 cells, respectively. Three additional uniform grid simulations are performed by coarsening the baseline simulation by factors of 2, 4, and 8, i.e., with grid size, 2048×1024 , 1024×512 , and 512×256 , respectively. These cases are named *baseline*, *coarse2*, *coarse4*, and *coarse8*, respectively. To differentiate the effect of the macroparticle resolution and spatial resolution of the grid, we fix the initial macroparticle resolution per unit area to be the same in these simulations, i.e., we initialize the simulations with 64, 256, 1024, and 4096 macroparticles per cell (ppc). Note that the physical plasma density at initialization is the same in all

these simulations (this is achieved by the choice of macroparticle weight). All the simulations are performed using the PSATD Maxwell solver, with a CFL=0.95, Esirkepov deposition, and cloud-in-cell interpolation.

A comparison of the evolution of the top current sheet obtained from the *baseline*, *coarse4*, and *coarse8* simulations is shown in Fig. 3. The current sheet in the *baseline* simulation evolves similarly to the simulation performed in our previous work.²⁶ Soon after the current sheets are initialized, and a 1% perturbation is applied, the magnetic pressure drops above and below the current sheet, causing it to collapse. The system continues to evolve, producing regions of trapped plasma called plasmoids. These plasmoids move outward along the

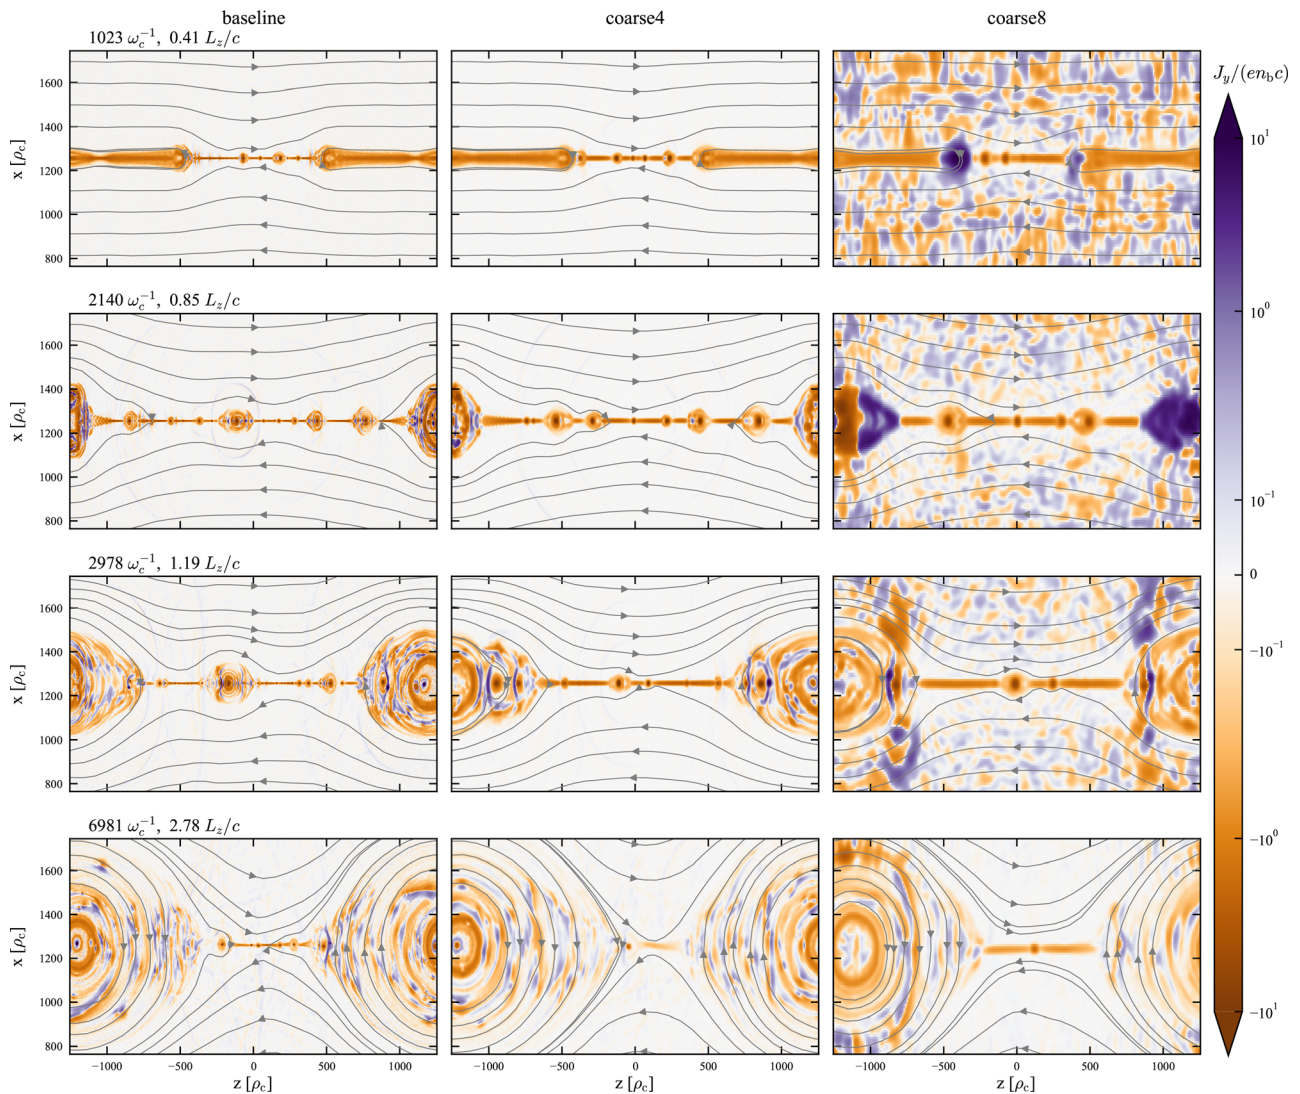


FIG. 3. Comparison of temporal evolution of normalized out-of-plane current density, $j_y / (en_b c)$, obtained from three 2D uniform grid simulations with baseline, coarse4, and coarse8 resolutions. The magnetic field lines are shown by the gray lines. The *baseline* and *coarse4* simulations show qualitatively similar behavior, though the initial current sheet is somewhat underresolved in *coarse4*. The *coarse8* simulation is strongly affected by numerical instabilities (or numerical heating) due to underresolution of the necessary plasma length scales.

TABLE II. MR simulations and parameters.

| MR case | <i>coarse2RR2</i> | <i>coarse4RR4</i> | <i>coarse8RR8</i> |
|-------------------------------------|-------------------|-------------------|-------------------|
| Parent [N_x, N_z] | [2048, 1024] | [1024, 512] | [512, 256] |
| RR | 2 | 4 | 8 |
| Absorbing layer ^a | 20 | 40 | 80 |
| FGB ^b layer ^a | 48 | 96 | 192 |

^aThe width of the layers are set by the number of fine-patch cells.

^bField gather buffer (FGB).

current sheet, and merge, forming larger plasmoids (as seen at 2094 and 3025 ω_c^{-1}), finally leaving a single plasmoid at the end of reconnection (at $\sim 7000 \omega_c^{-1}$). The *coarse4* simulation qualitatively shows a similar evolution; however, due to the $4 \times$ coarser cell size, the thin current sheet is not captured as well as in the *baseline* simulation. On the other hand, the *coarse8* simulation is severely under-resolved with just 0.25 cells per current sheet skin depth, 0.153 cells per upstream Debye length, and 0.2225 cells per current sheet Debye length, resulting in a numerical instability (caused by numerical heating) that appears prominently in the rightmost column of Fig. 3. We note that *coarse2* had similar evolution to *baseline*, and so have omitted it from Fig. 3 for brevity.

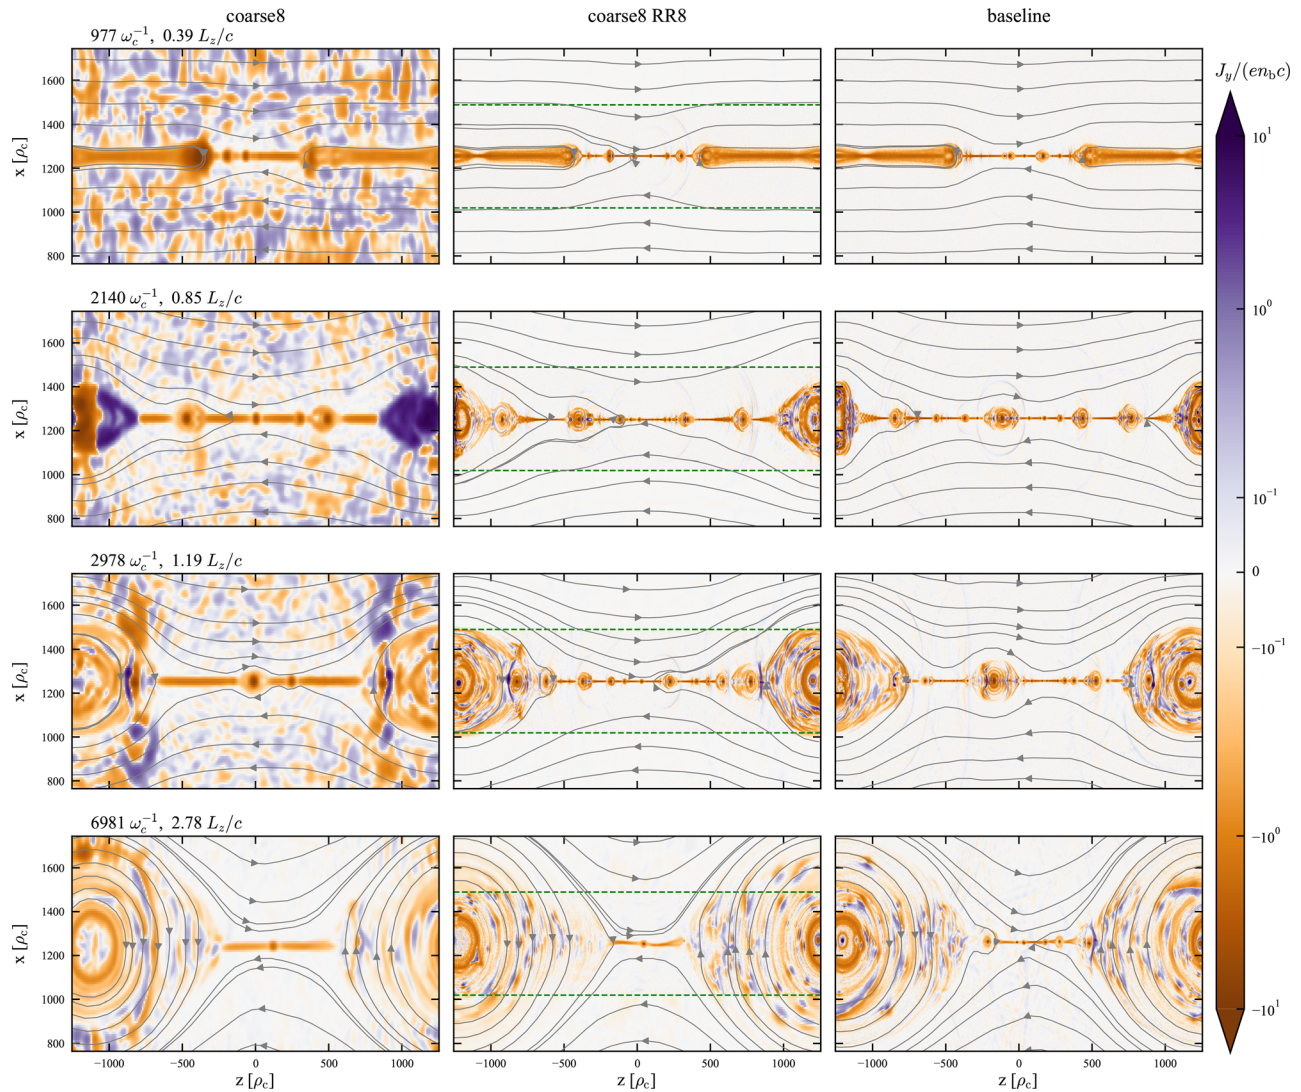


FIG. 4. Comparison of temporal evolution of the normalized out-of-plane current density, $j_y / (en_b c)$, obtained from three 2D grid simulations, namely, uniform *coarse8*, mesh refinement case, *coarse8RR8*, and the *baseline* uniform grid simulation. The magnetic field lines are shown by the gray lines. The addition of mesh refinement greatly reduces the effects of low resolution in the *coarse8* simulation. The qualitative evolution of *coarse8RR8* matches that of our high-resolution *baseline* case. Note that the fields are shown on the valid regions of L_1 and L_0 from which the particles gather electromagnetic fields for the top current sheet. The valid region of the mesh refinement patch on L_1 , surrounding the top current sheet, spans $1040 < x/\rho_c < 1520$, and is indicated by the dashed green lines. (For the bottom current sheet, the valid region of the mesh refinement patch on L_1 spans from $-1520 < x/\rho_c < -1040$). Both current sheets evolve similarly and for clarity we zoom in and show only the top current sheet.

2. Application of mesh refinement and its effect on current density

Next, we study how the application of mesh refinement affects the results for the coarse resolution simulations. We perform three MR simulations, with parent grids that have the same resolution as the *coarse2*, *coarse4*, and *coarse8* simulations, i.e., (2048×1024) (1024×512) (512×256) grid-sizes. For each of these cases, static refinement patches are applied around both current sheets (as illustrated in Fig. 1) with a refinement ratio (RR) such that the resolution of the fine patch is the same as the resolution in the *baseline* case. Refinement ratio (RR) is the ratio of the cell size on the parent grid to the refined patch. The RR for the refinement patches applied to parent grids with resolution same as *coarse2*, *coarse4*, and *coarse8* is set to $RR = 2, 4,$ and $8,$ respectively, and these MR cases are named *coarse2RR2*, *coarse4RR4*, and *coarse8RR8*, respectively. Note that these patches are static, and therefore, the refinement patch size is chosen such that it can capture reconnection physics until the end of




FIG. 5. Comparison of energy conversion (top) and relative energy conservation (bottom) obtained from the uniform grid 2D *baseline*, *coarse2*, *coarse4*, and *coarse8* simulations. In the top panel, the magnetic field energy and particle energy are normalized by the total initial energy and shown by solid and dashed lines, respectively. The *baseline* result shown in black is averaged over five repeated simulations along with the standard deviation denoted by error bars.

reconnection. From our previous work, we learned that the size of largest plasmoid extends up to $\sim 800\rho_c$. Therefore, the static mesh refinement patches are initialized to be $800\rho_c$, resulting in 37% of the domain being refined. The time step in the mesh refinement simulations is set by the $CFL = 0.95$, based on the cell size at the finest resolution (i.e., the time step is the same as the *baseline* case). For all the mesh refinement simulations, the initial number of macroparticles per species per unit area is set to be the same as the *baseline* case, i.e., 256, 1024, and 4096 macroparticles per parent cell (64 macroparticles per fine patch cell) for the *coarse2RR2*, *coarse4RR4*, and *coarse8RR8* cases, respectively. As mentioned previously in Sec. II B, the refinement patch contains an absorbing boundary layer that extends into the fine patch from the coarse-fine boundary, and a field-gather buffer region, within which particles gather fields from the parent grid to avoid numerical artifacts. For all the MR simulations, the absorbing layer was set to be 10 parent grid cells wide starting from the coarse-fine interface, corresponding to 20, 40, and 80 fine-patch cells for the *coarse2RR2*, *coarse4RR4*, and *coarse8RR8* simulations, respectively. Similarly, the width of the field gather buffer region required to avoid numerical artifacts for the chosen refinement ratio are 48, 96, and 192 fine-patch cells from the coarse-fine interface (the green region in Fig. 1) for the *coarse2RR2*, *coarse4RR4*, and *coarse8RR8* simulations, respectively. The numerical parameters used for the mesh refinement simulations are summarized in Table II.

In Fig. 4, we compare the evolution of the out-of-plane current density, j_y , obtained from *coarse8*, *coarse8RR8*, and *baseline* simulations. It can be seen that the numerical instability observed for the *coarse8* simulation is mitigated when using mesh refinement in the *coarse8RR8* case, because the mesh-refined region resolves the initial current sheet skin depth. Compared to the *baseline* simulation, it can be seen that *coarse8RR8* can capture the current sheet, the formation of plasmoids, merging of plasmoids leading to secondary reconnection (seen at $t = 2978\omega_c^{-1}$), finally forming a single plasmoid at late times when reconnection has quenched. Note that, while we do not expect the evolution of plasmoids to be exactly the same as the *baseline* case, these results confirm that even with a high refinement ratio of 8, the simulations are able to capture reconnection characteristics. The evolution of current density was qualitatively similar for the *coarse2RR2* and *coarse4RR4* simulations. We chose to highlight *coarse8RR8* because the uniform grid simulation without mesh refinement, i.e., *coarse8* simulation exhibited instability, hence providing a more challenging test for mesh refinement. We also performed simulations with the refined patch confined to the center of the current sheet, such that the coarse-fine interface cuts across the current sheet, and found that the mesh refinement algorithm captures reconnection characteristics well for these cases as well. Results from these simulations are presented in Appendix C.

B. Effect of mesh refinement on energy conservation and conversion

A comparison of the energy transfer from magnetic field to particle kinetic energy (thermal and bulk acceleration) obtained from the *baseline*, *coarse2*, *coarse4*, and *coarse8* simulations is shown in the top panel of Fig. 5. The current sheet evolution includes a linear regime, when the current sheet breaks and forms small regions of trapped plasma, and during this time, the particle energy increases exponentially. At around $t \sim 1800\omega_c^{-1}$, transition to the non-linear regime

begins where plasmoids merge to form larger plasmoids also causing secondary reconnection. This continues until reconnection ceases by $t \sim 7000\omega_c^{-1}$ when the magnetic field energy and particle energy reach quasi steady-state. Since energy transfer can exhibit some small differences in the non-linear regime beginning at $t = 1800\omega_c^{-1}$, five simulations were performed with the *baseline* numerical parameters. The averaged magnetic field (solid) and particle energy (dashed) obtained from these simulations are shown in black for the *baseline* case in Fig. 5, along with the standard deviation (error bars). Energy conversion obtained from *coarse2* and *coarse4* simulations proceeds identically with the *baseline* case in the linear regime, where the tearing mode instability dominates, i.e., until $t = 1800\omega_c^{-1}$. Beyond this linear regime, there are some differences in energy transfer but within the standard deviation of the *baseline* for the *coarse2*, and a few standard deviations of the *baseline* for the *coarse4* case in the mid-reconnection phase ($1800\omega_c^{-1} < t < 5000\omega_c^{-1}$), beyond which it is within the standard deviation until the end of reconnection. On the other hand, the *coarse8* simulation does not capture the energy transfer process accurately since the skin depth for this case is highly under-resolved, as was also observed from the current density evolution shown previously in Fig. 3. In the bottom panel of Fig. 5, we compare the relative energy conservation and it can be seen that the *baseline* and *coarse2* energies are well-conserved and *coarse4* deviates from *baseline* when the non-linear reconnection regime begins ($t = 1800\omega_c^{-1}$); however, it is still within a relative difference of 4×10^{-3} . The *coarse8* simulation did not conserve energy, consistent with the increase in magnetic field energy and particle energy observed in the top panel. We also compared momentum conservation for the uniform grid cases and found that the x , y , and z total momenta are conserved with $\pm 1 \times 10^{-4}$ from the respective initial total momenta (further details in Fig. 12 in Appendix B).

In Fig. 6, we compare the energy conversion obtained from the 2D mesh refinement simulations, *coarse2RR2*, *coarse4RR4*, and *coarse8RR8* with the magnetic field and particle energy evolution averaged from five *baseline* simulations. As seen from the top panel, the energy conversion proceeds identically for all the cases in the linear regime (until $t = 1800\omega_c^{-1}$) and then within the standard deviation from the baseline simulation. In the bottom panel, we compare energy conservation (relative energy with respect to initial total energy) from the MR simulations and the baseline. By construction, the MR method is not energy-conserving, because we damp the electromagnetic fields and the current density in the absorbing layer adjacent to the coarse-fine interface in the fine patch (i.e., level 1 grid). Even then, the energy is conserved within 1% for the *coarse8RR8* simulation with the largest refinement ratio of 8. We investigated the initial bump at around $t = 200\omega_c^{-1}$ observed for the *coarse8RR8*, and *coarse4RR4* cases and found that it is caused by a very small signal that crosses the edges of the field-gather buffer region in the fine patch. However, later on, this small bump in the signal does not affect the reconnection physics or energy conversion processes as apparent from the time evolution of energy transfer from the top panel. We also found that the x , y , and z total momenta are conserved with $\pm 1 \times 10^{-4}$ of the initial total momentum for the MR simulations (Fig. 13 in the Appendix C). Thus, mesh refinement does not significantly alter momentum conservation compared to the highly refined *baseline* case.

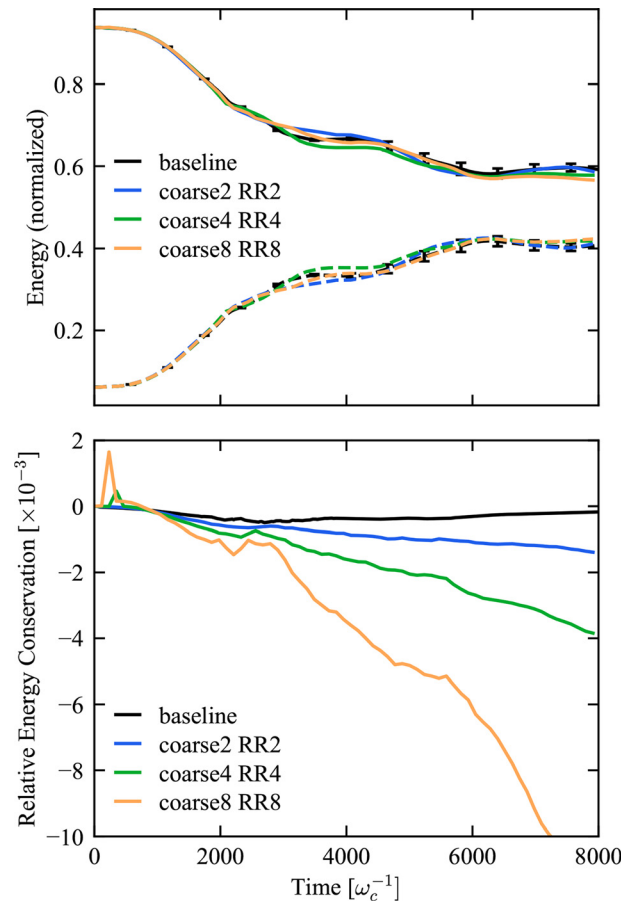


FIG. 6. Comparison of energy conversion (top) and relative energy conservation (bottom) obtained from the uniform grid 2D *baseline*, *coarse2RR2*, *coarse4RR4*, and *coarse8RR8* simulations. In the top panel, the magnetic field energy and particle energy are normalized by the total initial energy and shown by solid and dashed lines, respectively. The *baseline* result shown in black is averaged over five repeated simulations along with the standard deviation denoted by error bars.

C. Effect of grid resolution and mesh refinement on particle acceleration and reconnection rate

The evolution of particle acceleration from the mesh refinement simulations, *coarse2RR2*, *coarse4RR4*, and *coarse8RR8*, is compared with the *baseline* and *coarse8* uniform grid simulations, in Fig. 7. Similar to our previous work,²⁶ the highest particle γ at the start of reconnection is 30, and at the end of reconnection, $t = 7000\omega_c^{-1}$, the highest particle γ for our baseline simulation increased by an order of magnitude to 500. Majority of the particles have $\gamma \leq \sigma$, where $\sigma = 30$ is the magnetization used for our 2D relativistic reconnection simulations, which is consistent with previous results in the literature.^{10,15,16} It can be seen that compared to the *baseline* simulation, the particle spectra obtained from the *coarse8* uniform grid simulation are subject to numerical heating and do not exhibit the expected power law. This is consistent with the energy increase observed in Fig. 5 and inability to capture current sheet in Fig. 3. However, for the parent grid with the same resolution, when including a mesh refinement patch as in




FIG. 7. Comparison of time evolution of particle spectra for the *baseline* (solid black), and MR simulations, *coarse2RR2* (solid blue), *coarse4RR4* (solid green), and *coarse8RR8* (solid orange) simulations. The solid red line is obtained from the uniform grid *coarse8* simulation.

coarse8RR8 simulation, the particle spectra and thus particle acceleration are captured accurately. Note that *coarse2RR2* and *coarse4RR4* also show the same quantitative evolution of particle spectra. The uniform grid *coarse2* and *coarse4* simulations are not shown for brevity, but they quantitatively agree with the *baseline* simulations.

1. Simulations with fewer macroparticles and their effect on particle acceleration

Particle-in-cell simulation runtimes are dominated by the total number of particles, especially for the relativistic reconnection simulations presented in this work. To isolate the effect of spatial resolution from macroparticle resolution, the total number of particles was kept the same for all cases presented so far (as discussed in Sec. II). For the *baseline* and *coarse8* uniform grid simulations, there were (8×8) and (64×64) macroparticles per cell, respectively, at initialization, i.e., same number of macroparticles per unit area. For the *coarse8RR8* simulation also, the parent grid was initialized with (64×64) particles per cell everywhere, such that the cells in the refined patch had (8×8) macroparticles per species at initialization. Results are presented from three additional simulations performed with the same grid as the *coarse8RR8* case but with fewer total number of macroparticles. For these simulations, the number of initial macroparticles was sequentially reduced only in the coarse region from (64×64) to (32×32) , (16×16) , and (8×8) macroparticles per coarse cell per species. The number of macroparticles in the fine patch cells was maintained at (8×8) for all these simulations, similar to the *baseline* and previously discussed MR simulations. An additional uniform grid simulation was performed named *baseline-w/8 × 8MR-part*, where the resolution of the uniform grid was the same as *baseline*, but the initial particle distribution was same as MR simulation, *coarse8RR8* with (8×8) macroparticles per coarse cell per species, and 8×8 macroparticles per fine cell per species (i.e., the total macroparticles were same as the *coarse8RR8* simulation with (8×8) particles in fine and coarse patch). This initial particle distribution in the *baseline-w/8 × 8MR-part* simulation results in 8×8 macroparticles per uniform grid cell per species in the regions corresponding to the fine patch of MR simulations, and 1×1 macroparticles per uniform grid cell corresponding to the coarse regions of the *coarse8RR8* simulation.

We compare the particle acceleration obtained from these *coarse8RR8* simulations initialized with (64×64) , (32×32) , (16×16) , and (8×8) macroparticles in the coarse regions, with the uniform grid *baseline*, *baseline w/8 × 8MR-part*, and *coarse8* simulations in Fig. 8. It can be seen that at the start of the non-linear regime, at $t \sim 2000\omega_c^{-1}$, the particle acceleration obtained from all the *coarse8RR8* simulations with different initial macroparticle resolutions agree well with the *baseline* simulation. This agreement is also observed mid-reconnection at $t \sim 3000\omega_c^{-1}$. At the end of reconnection, the particle spectra in the *coarse8RR8* simulations with few particles agree very well with the *baseline* simulation for $\gamma > 5$. As shown in the zoomed-in inset, some minor deviation appears from the *baseline* simulation for the simulations with fewer particles. These minor differences for the low γ region are likely due to the low macroparticle resolution in the coarse and fine upstream regions at the end of reconnection. During reconnection, particles from the upstream flow toward the current sheet and become trapped in plasmoids. As a result, at the end of reconnection, the upstream region in the fine patch has less than five macroparticles per cell in some parts, contributing to differences in the low- γ spectra. The distribution of particle count as reconnection progresses is shown in Fig. 17 (discussed in Appendix D). In Appendix D, we split the energy spectra into contributions from the coarse region, labeled as upstream, and the fine patch regions. We find that spectral contributions from the coarse region are also affected by low-particle resolution, even when the number of macroparticles per cell is much greater

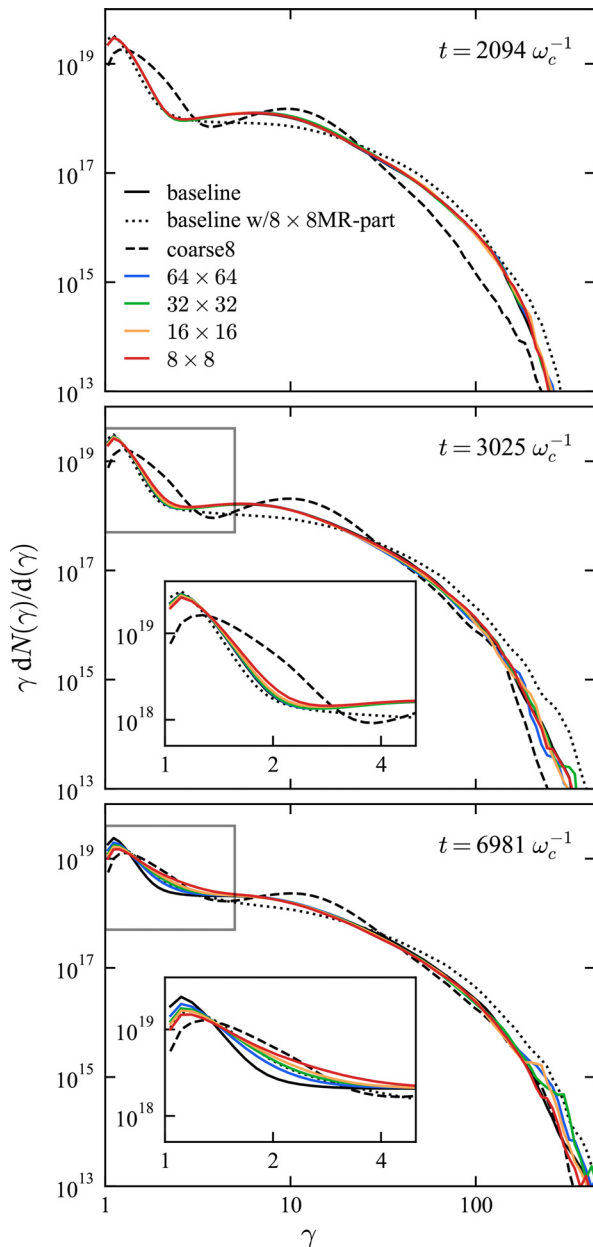


FIG. 8. Comparison of time evolution of particle spectra for the *baseline* (solid black), *coarse8* (dashed black), and *coarse8RR8* simulations initialized with (64×64) (solid blue) (32×32) (solid green) (16×16) (solid orange), and (8×8) (solid red) macroparticles in the coarse region cells, and (8×8) macroparticles per cell in the fine-patch regions (similar to *baseline*). Additionally, particle spectra obtained from *baseline-w/8 x 8MR-part* simulation (dotted black) with similar initial particle distribution as the *coarse8RR8* case with (8×8) macroparticles in the coarse region cells is also compared.

than 5. Despite this underresolution in both the coarse and fine regions, the low-energy spectral signature is minimally affected. There is no effect, however, on the high-energy portion of the power law, which is a critical signature of reconnection. The high-energy spectrum

is captured with high fidelity even in the *coarse8RR8* simulation initialized with 8×8 particles per cell in the coarse region.

In contrast, the particle spectra obtained from the *baseline w/8 x 8MR-part* simulation show that more particles are accelerated compared to the *baseline*, even though the grid resolves the current sheet skin-depth and upstream and current sheet Debye lengths. Qualitatively, the plasmoid formation was found to be different in this case compared to the *baseline* and *coarse8RR8* simulations with different macroparticles (shown in Fig. 18 in Appendix E). This is consistent with higher particle acceleration observed in the particle spectra. This further supports the efficiency of mesh refinement, where, with the same initial particle distribution, the particle spectra can be captured with MR, such that initially, particle statistics per cell is well maintained. The efficiency can be further improved by applying particle splitting or merging, which will be performed as future work.

2. Effect of mesh refinement on reconnection rate

Accuracy of the mesh refinement simulations is also investigated by comparing the reconnection rate with the highly refined *baseline* uniform grid simulation. The dimensionless reconnection rate, β , is given by²⁴

$$\beta = -\frac{1}{v_A B_0 L_x} \frac{d\Phi}{dt}, \quad (8)$$

where v_A is the Alfvén velocity, B_0 is the upstream magnetic field, L_x is the size of the domain, and Φ is the unreconnected flux. Since directly measuring unreconnected flux is difficult, we use the approximation, $\beta \equiv v_{in}/v_{out}$, where v_{in} is the inflow velocity into the reconnection current sheet, and v_{out} is the terminal exhaust velocity downstream.^{38,46} Similar to our previous work,²⁶ we calculate v_{in} by averaging $|v_x|$ within a region of size $x_R \times z_R = 122 \rho_c \times 490 \rho_c = 0.098 L_x \times 0.195 L_z$ centered on $(x, z) = (-x_c + x_R, 0)$. We found that the measured inflow velocity is relatively insensitive to the choice of x_R and z_R , and is also symmetric across the current sheet, and insensitive to the choice of the top and bottom current sheets. To compare the reconnection rate from the *baseline* and MR simulations, we choose the $+x$ side of the bottom current sheet. The outflow velocity, v_{out} , is measured by taking the median of the 10 highest cell-averaged z -velocities within $\delta = 12 \rho_c$ of the center of the current sheet. The outflow velocity approaches the expected Alfvén velocity, $v_A = c\sigma/\sqrt{\sigma^2 + \sigma} \approx c$.

The estimated reconnection rate for the *baseline*, *coarse2RR2*, *coarse4RR4*, and *coarse8RR8* simulations are compared in Fig. 9. Note that these simulations were performed with same number of macroparticles in the domain. We also compare the reconnection rate obtained from the *coarse8RR8* simulation with 8×8 particles per coarse and fine cell. It can be seen that the growth of the ratio of v_{in}/v_{out} from 0 to approximately 0.26 occurs within $t \sim 1000\omega_c^{-1}$. This growth is captured well by the mesh refinement simulations performed with same number of macroparticles, and also by the *coarse8RR8* simulation performed with fewer particles (8×8 macroparticles in coarse and fine cells initially). From $t = 1000\omega_c^{-1}$ to $2500\omega_c^{-1}$, reconnection rate estimates obtained from all simulations remain steady between 0.15 and 0.25, before decaying to < 0.1 . The time $2500\omega_c^{-1}$ approximately corresponds to the non-linear current sheet evolution. Thus, after this time, the linearity assumption underlying the estimate $\beta \approx v_{in}/v_{out}$ breaks down and is no longer an accurate estimate of the reconnection rate. From the reconnection rate

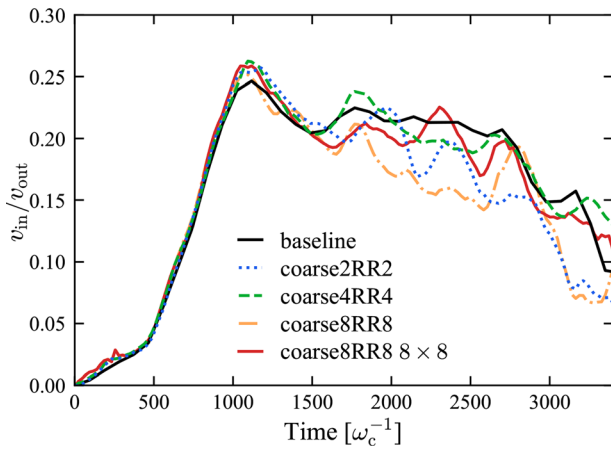


FIG. 9. Comparison of time evolution of reconnection rate, $\beta \equiv v_{in}/v_{out}$, obtained for *baseline* (solid black), *coarse2RR2* (dotted blue), *coarse4RR4* (dashed green), *coarse8RR8* (dashed-dotted orange), and *coarse8RR8* with 8×8 macroparticles per coarse and fine cell per species at initialization (solid red). Reconnection rate evolves similarly in all cases, peaking at 0.26 at $t \sim 1000\omega_c^{-1}$, remaining between 0.17 and 0.23 from $t = 1000\omega_c^{-1}$ to $2500\omega_c^{-1}$. After this time, reconnection becomes non-linear and the assumptions made to estimate the reconnection rate are no longer valid. Results from using MR thus agree well with the *baseline* uniform grid simulations.

TABLE III. Timing comparison of *baseline* and MR simulation.

| Case | <i>Baseline</i> | <i>coarse8RR8</i> (64ppc ^a) |
|----------------------|-----------------|---|
| No. of GPU nodes | 4 | 2 |
| No. of particles (M) | 1073.7 | 413.3 |
| Total wall time (s) | 560 | 793 |
| Total node hours | 0.62 | 0.44 |
| Performance increase | ... | 1.4 |

^a 8×8 particles per cell (ppc) per species in both coarse and fine cells at initialization.

comparisons made in Fig. 9, it can be seen that the mesh refinement method used in this work captures reconnection rate as well as the high-resolution *baseline* uniform grid simulations.

D. Performance comparison with uniform grid

Since the results obtained from the *coarse8RR8* simulation initialized with 64 particles per cell in the coarse and fine cells agree well with the *baseline* simulation, we compare the wall time and node-hours used between the uniform grid *baseline* simulation and the *coarse8RR8*. Since the memory footprint for the *coarse8RR8* simulations with eight times coarse particle resolution in the coarse patch is reduced by a factor of 4, the latter simulation fits on just two nodes of the OLCF Summit supercomputer while the *baseline* simulation required four nodes. We used 4 GPUs per node so that the grids can be equally divided between the nodes, instead of using all the 6GPUs available per Summit node. In Table III, we compare the total runtime and node hours used for the simulation. For both simulations, we did not include diagnostics or I/O and performed the simulations up to the end of reconnection (7000 timesteps). The *baseline* simulation

with 1073×10^6 particles required 560 s with four Summit nodes, i.e., 0.62 node-hours. With the *coarse8RR8* simulation with 413.3×10^6 particles, we obtained a performance increase by a factor of 1.4 in terms of node-hours used. In the uniform grid and MR simulations, the most time-consuming step of the PIC loop is current-deposition, which is directly proportional to the number of particles. Note that the number of particles decreased by a factor of 2.6 in the MR simulation and we also decreased the number of nodes by a factor of 2. Thus, a speed-up of $1.4 \times$ is expected if we consider that the main cost of the PIC loop is from particle deposition. In the MR simulations, in addition to current deposition which takes 30% of the total time, nearly 14% of the total time is consumed by the routine that sorts particles into field gather buffers. As mentioned before, particles in the fine patch that are in the field gather buffer region gather fields from the parent level. Sorting particles ensures that particles that gather from Level 1 and Level 0 are contiguous in memory. However, detailed investigations on optimization of this routine were not conducted in this work, and targeted optimizations may improve performance in the future. In addition, the time required for communications doubled since the mesh refinement algorithm involves communication of field data from parent grid to coarse patch of Level 1, which is interpolated to the auxiliary patch. The domain decomposition grids on Levels 0 and 1 are different, and as a result, the communication of field data across levels may be computationally expensive and add to the time required for the MR simulation to complete.

Preliminary simulations for 3D reconnection were also performed to determine the performance improvement from using a static MR for 1000 timesteps. 3D uniform grid simulations were performed with the same resolution as the 2D *coarse2* case, resulting in a $(2048 \times 1024 \times 1024)$ grid, since it would be computationally intensive to have the same resolution as the 2D *baseline* case. The 3D uniform grid simulation was initialized with eight particles per cell, resulting in 34.36×10^9 macroparticles. For 1000 timesteps, the wall clock time was 119.8 s using 512 GPU nodes, i.e., 17.04 node hours. A 3D mesh refinement simulation was also performed with parent grid $4 \times$ coarser than the 3D uniform grid simulation with parent grid size of $(512 \times 256 \times 256)$, similar to the resolution of the 2D *coarse8* case. Static mesh refinement patches were initialized surrounding the two current sheets, such that, 25% of the domain was refined with a refinement ratio of 4 in each direction, and the coarse and fine cells were initialized with 8 particles per cell. The total macroparticles as a result were reduced by a factor of 3.8, requiring eight times fewer nodes and 282 s wall clock time, i.e., 5.01 node hours. Thus, with a refinement ratio of 4, and refinement patch that covered 25% of the domain, the total performance improved by a factor of 3.37 for the 3D simulation. We also extended the MR patch to extend 37% of the domain, and the performance increased by a factor of 1.4. Note that further investigations need to be performed for the full 3D reconnection, where the particle distribution will become more unbalanced as reconnection proceeds. Nevertheless, these preliminary results already indicate that 3D reconnection with MR will enable higher resolutions, which might be computationally expensive or even infeasible with a uniform grid.

IV. EFFECT OF MESH REFINEMENT PARAMETERS

A. Effect of PSATD vs FDTD

Uniform grid simulation results obtained from the widely used FDTD Yee method were compared to those obtained using the

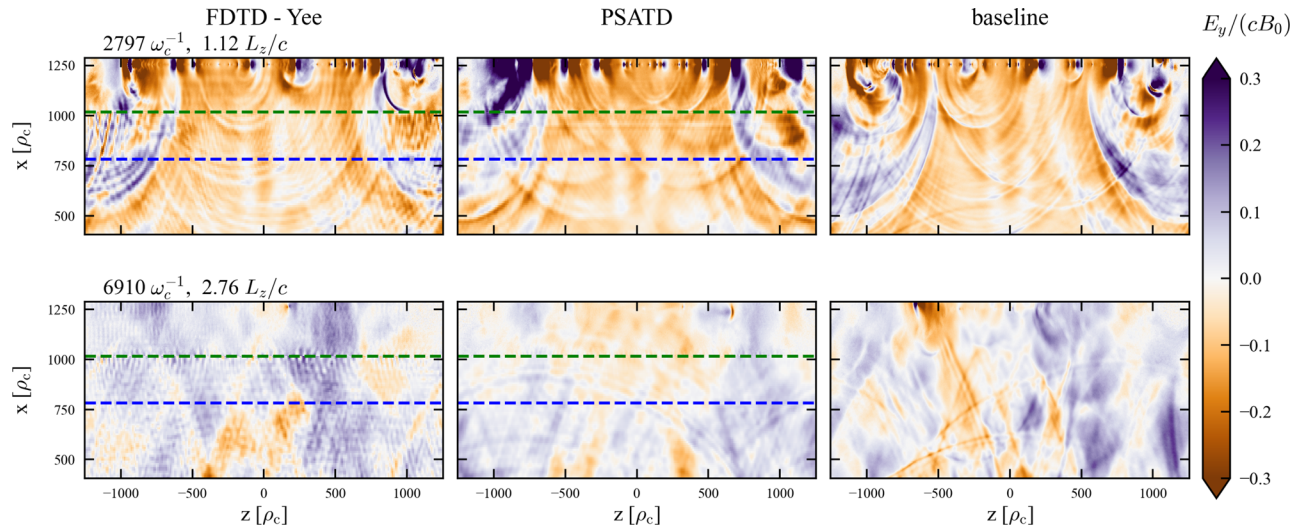


FIG. 10. Comparison of spatial variation of the normalized out-of-plane electric field, $E_y/(cB_0)$ for mesh refinement simulations performed with same grid configuration as *coarse8RR8* with the Yee solver (left), the PSATD solver (middle), with the high resolution uniform grid baseline case (right). The FIG is zoomed in near the coarse-fine interface, depicted by the dotted blue line, and the edge of the field-gather buffer region, depicted by the dotted green line.

PSATD Maxwell solvers in Ref. 27 where it was found that both solvers capture the evolution of reconnection identically for the numerical parameters that were considered. In this work, we performed MR simulations for the *coarse8RR8* test case with the two solvers. Figure 10 shows a comparison of the normalized out-of-plane electric field for the FDTD Yee simulation, PSATD, and *baseline* uniform grid simulation. The fine-coarse interface is shown by the dashed-blue lines for the MR simulation results. The dashed green line indicates the edge of the field gather buffer region in the fine patch, where particles in the fine-patch region between the green and blue lines gather fields from the level below (parent grid in this case). Note that the fields in these plots are from the levels that the particles gather from, and therefore, the coarse-grid result is shown in the region below the dashed-green line, since particles gather from the parent grid in those regions. At mid-reconnection ($t \sim 3000\omega_c^{-1}$), the E_y solution obtained from the Yee simulation started to develop spurious structures near the coarse and fine patch interface, and by the end of reconnection, at $t \sim 7000\omega_c^{-1}$, these structures are present everywhere in the domain, with larger wavelength in the coarse-grid compared to the fine-patch. A similar structure was seen previously for plasma-accelerator simulations with the Yee scheme and the mismatch of numerical dispersion at the coarse and fine grid.²⁹ Similar to previously studied numerical dispersion in the reconnection simulations we present here, the electromagnetic waves propagate at different speeds on the fine and coarse grid, which is further aggravated by the single time step chosen to solve Maxwell's equations on every level in our simulations. Subcycling in the fine-patch with time step ratio on the coarse and fine cells such that the corresponding CFL is the same on each grid and close to unity was found to improve the results in previously performed accelerator simulations.²⁹ Another solution is to use an ultrahigh-order PSATD solver. Compared to the Yee solver, the near-dispersionless PSATD method does not develop these short-wavelength structures and the normalized electric field qualitatively compares well with the *baseline*

uniform grid simulations. Thus, all the MR simulations reported in Sec. III used the PSATD solver (with order 16).

B. Field damping method and parameters in absorbing layer

As mentioned previously in Sec. II B, Maxwell's equations are solved on the fine and coarse patches of each level. At the edges of the fine patch, the electromagnetic fields are terminated by an absorbing layer. In this work, we implemented an absorbing layer to damp the fields and compared its effect to the simulations that use a perfectly matched layer (PML), which is the default for mesh refinement patches in WarpX. In the PML method, the fields are split into normal and tangential components, and the normal components are damped. Since the PML method was tailored to absorb electromagnetic waves in vacuum, this method works well with particle accelerator simulations where typically, the plasma density near the coarse-fine interface is very small.^{31,36} However, for applications such as magnetic reconnection, where the plasma density and current density at the coarse-fine interface is high, it was found empirically that damping all the components, as is done with an absorbing layer, performs better. For damping fields in the absorbing layer, the following non-physical conductivity is used:³⁶

$$\sigma_{x,i} = \sigma_{\max} \left(\frac{i\Delta x}{\delta} \right)^{N_{\text{exp}}}, \quad i = 0, \dots, N_{\text{layer}}, \quad (9)$$

where $\sigma_{x,i}$ is the non-physical conductivity in the i th cell of the absorbing (or PML) layer, N_{layer} is the number of cells in the absorbing (or PML) layer, Δx is the size of the cell, δ is the size of the damping layer given by $N_{\text{layer}} * \Delta x$, N_{exp} sets the profile for the function (we use quadratic, $N_{\text{exp}} = 2$ or cubic $N_{\text{exp}} = 3$), and $\sigma_{\max} = \frac{\kappa_{ds}c}{\Delta x}$, where c is the speed of light and κ_{ds} is the damping

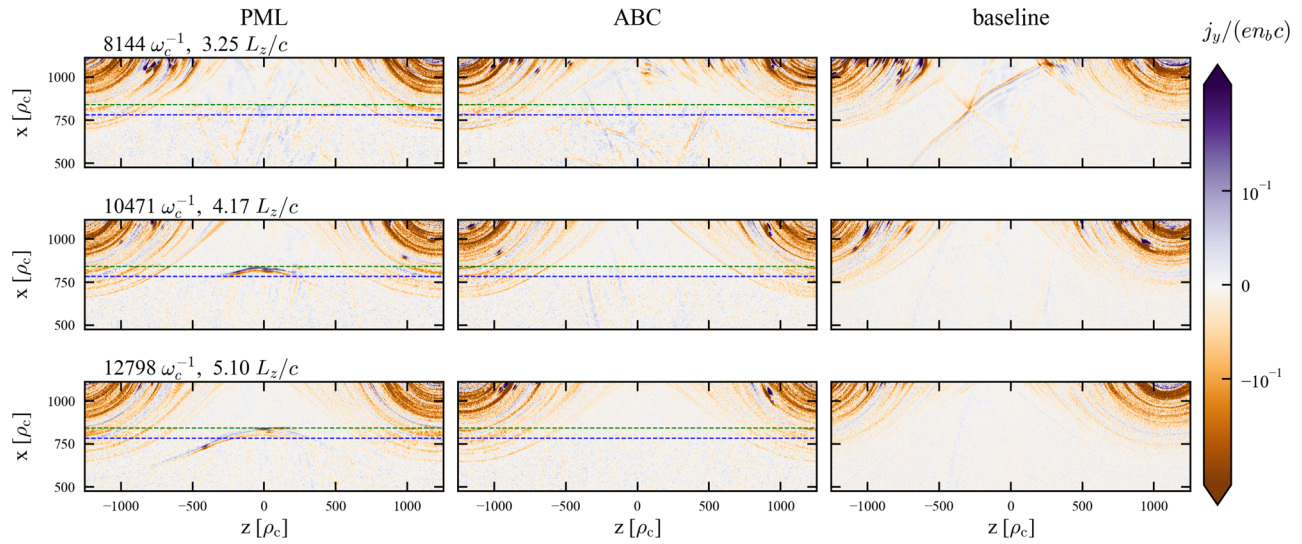


FIG. 11. Comparison of spatial variation of normalized out-of-plane current density, $j_y / (en_b c)$ for mesh refinement simulations performed with *coarse2RR2* grid configuration with the PML (left), the absorbing layer (middle), and the uniform grid baseline (PSATD) case (right). The figure is zoomed-in near the coarse-fine interface, depicted by the dotted blue line, and the edge of the field-gather buffer region, depicted by the dotted green line.

strength. We build on the two-step PML formulation described in Shapoval *et al.*,³⁶ and instead of damping only the normal components, we also damp the tangential components. Note that for both the PML and absorbing layer, the same conductivity profile is used to damp the fields in the absorbing layers at the edges of the refinement patches. The main difference between the two treatments, is that, in the PML, the fields are split into tangential and normal components, and only the tangential components are damped, while with the absorbing layer implemented in this work, all the components are damped equally.

We compare the effect of using an absorbing layer or a PML for the *coarse2RR2* case and compare with the *baseline* simulation. In addition to damping the fields, the current density from the macroparticles is also deposited in the absorbing layer and damped using the same damping profile as for the fields. For the comparison, a cubic profile was used for the conductivity ($N_{\text{exp}} = 3$), with a damping strength, $\kappa_{ds} = 4$. The absorbing layer is 20 cells in the fine-patch, and the field gather buffer region is 28 cells for both simulations. The spatial variation of the out-of-plane current density, j_y , is shown in Fig. 11. We observed that until the end of reconnection, the solution with the PML method compared well with the absorbing layer and *baseline* simulation. However, when we performed these simulations for a few more light-crossing times, beyond $t > 8000\omega_c^{-1}$, we observed an accumulation of non-physical current density near the PML region at the coarse-fine interface, as seen prominently at $t = 10471\omega_c^{-1}$, which continues to grow at $t = 12800\omega_c^{-1}$. While the PML method did not significantly affect the solution until the end of reconnection, we decided to investigate and found that damping all the components with the absorbing layer method mitigates these numerical artifacts at the coarse-fine boundary as can be seen from the absorbing layer solution in Fig. 11. We found that the choice of the damping profile (quadratic or cubic) or the damping strength (varied from 4 to 30) does not significantly affect the solution.

V. CONCLUSIONS AND FUTURE WORK

In this work, we studied the application of static mesh refinement to first principles 2D particle-in-cell simulations of relativistic magnetic reconnection. Uniform grid simulations were performed first by sequentially coarsening the highest resolution uniform grid by factors of 2, 4, and 8. To distinguish the effect of macroparticle resolution and grid resolution, all the uniform grid simulations were initialized with the same total number of macroparticles (increasing macroparticles per cell by (2×2) , (4×4) , and (8×8) compared to the baseline uniform grid simulation for simulations with grid coarsening factors 2, 4, and 8, respectively). The eight times coarser grid simulation did not resolve the current sheet skin depth and therefore did not accurately model magnetic reconnection, as expected. Applying a static MR fine-patch with a refinement ratio of 8 on top of the $8 \times$ coarser parent grid led to improved resolution of the current sheet. The MR simulation was able to capture the evolution of the current sheet during reconnection, energy conversion, energy conservation, and particle spectra accurately as indicated by the excellent agreement with the uniform grid baseline simulations. The number of macroparticles was then reduced such that, at initialization, the number of macroparticles for the MR simulation with refinement ratio of 8, was 8×8 in the coarse and fine cells. These simulations also modeled the particle spectra and the power law accurately at all times. However, it was found that, at the end of reconnection, by which time particles from the upstream are trapped in one large plasmoid, the number of macroparticles in the coarse upstream cells is not sufficient to capture the low-energy spectra ($\gamma = 2$) as accurately, and led to small deviation from the baseline solution. However, beyond $\gamma > 5$, which is relevant regime for particle acceleration, the spectra compare well with the uniform resolution baseline case.

The FDTD simulations that employed the Yee solver displayed spurious short-wavelength structures attributed to the large numerical dispersion occurring on the coarse parent grid and coarse patch. This is due

to the single time step set by a CFL ~ 1 on the refined patch, leading to an effective CFL of 0.125 on the coarse patch for a refinement ratio of 8. On the other hand, the ultrahigh-order PSATD solver is less susceptible to numerical dispersion and showed good agreement with the high-resolution uniform grid baseline results. A new absorbing layer was introduced to reduce the numerical artifacts at the coarse–fine interface that were observed with the PML method long after reconnection quenched.

Based on results from previous studies,²⁶ a refinement patch was chosen that covers at least 80% of the largest plasmoid size expected at the end of reconnection. This resulted in nearly 37% of the 2D domain being refined, reducing the number of macroparticles by one half for a refinement ratio of 8. A $1.4\times$ performance increase was observed in terms of node-hours used, compared to the high-resolution 2D uniform grid simulation. Manual performance optimization of the mesh refinement algorithm for magnetic reconnection, not explored here, should also provide additional performance improvements. Preliminary 3D uniform grid and two MR simulations were performed for 1000 timesteps to compare the performance improvement with refined regions covering 25% and 37.5% of the domain. The number of macroparticles required decreased by a factor of 0.25 and 0.5, improving the performance (in terms of node hours) by a factor of 3.4 and 1.46, respectively, reducing the number of GPU nodes required by a factor of 8 and 4, respectively. Thus larger memory savings and performance increases can be expected when using mesh refinement in 3D. Additionally, subcycling with time step ratio between fine and coarse levels being the same as the refinement ratio may also allow for increase efficiency. Detailed investigations are deferred to future work.

The MR strategies presented in this work have implications beyond the 2D relativistic reconnection application demonstrated here. The strategies presented here will also benefit non-relativistic magnetic reconnection and other high-energy systems with large disparities in length-scales and with high plasma currents crossing the coarse–fine interface. Additionally, preliminary 3D simulations show promising performance improvement, requiring fewer GPU nodes than the uniform grid counterpart. This suggests that 3D simulations with higher resolution in the current sheet are now possible due to reduced memory requirement compared to the uniform grid counterpart. This is especially significant when using radiative cooling, where cooling rates have been artificially decreased or turned off in regions where the local density becomes large, and the skin depth is not resolved by the restrictive uniform grid. It will therefore open a new window to study 3D effects.

The MR strategies applied to reconnection in this work lay the groundwork for future improvements. The MR simulations presented here used a static mesh, and this meant using a large region for refinement even at the start of the simulation when the current sheet thickness, or region requiring refinement, is much smaller. Future work will extend this method to include adaptive refinement as the current sheet evolves to form plasmoids and as the plasmoids merge growing in size, i.e., as the region requiring refinement evolves. The advantage of adaptive mesh refinement is that it would allow for a small refinement patch initially when the current sheet is thin, and as the current sheet evolves to form plasmoids, the refinement patch would adapt according to the local skin-depth. While static mesh refinement itself is quite challenging, dynamic mesh refinement adds to the challenge due to an evolving coarse–fine interface. This requires a detailed analysis of load-balancing and evolving damping layers. However, in

principle, the methods used for the static MR simulations discussed in this paper can be applied to adaptive MR, but with more engineering to enable adaptivity. Similarly, the simulations used a fixed number of macroparticles with weights that were set at initialization. However, it may be more efficient to split particles when they cross from the coarse to fine region and merge particles that transition to the coarse region. Some studies have performed particle splitting and merging,³⁹ and these will be explored in our future work. The MR algorithm presented in this work can readily be used to perform 3D simulations. While only single-level MR is presented in this work, the code and methods presented here are also capable of performing multiple levels of refinement. As mentioned previously, in high-energy astrophysical systems, radiative effects are important. With the MR method presented in this work, one can resolve the local skin depth due to higher densities caused by the cooling, without having to refine the other regions. Due to the reduced memory requirement, the mesh refinement approach presented in this work will render 3D simulations more tractable as they can be performed more efficiently.

ACKNOWLEDGMENTS

This work was partially supported by the U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research, Exascale Computing Project (Grant No. 17-SC-20-SC) under Contract No. DE-AC02-05CH11231, and Simulation and Analysis of Reacting Flow, FWP Grant No. FP00011940, funded by the Applied Mathematics Program in ASCR. This research used resources of the Oak Ridge Leadership Computing Facility, which is a DOE Office of Science User Facility supported under Contract No. DE-AC05-00OR22725 with the SummitPLUS award 2023-2024. This research also used resources of the National Energy Research Scientific Computing Center (NERSC), a U.S. Department of Energy Office of Science User Facility located at Lawrence Berkeley National Laboratory, operated under Contract No. DE-AC02-05CH11231 using NERSC award ASCR-ERCAP mp111 for 2023 and 2024. This research used the open-source particle-in-cell code WarpX <https://github.com/ECP-WarpX/WarpX>, primarily funded by the U.S. DOE Exascale Computing Project. Primary WarpX contributors are with LBNL, LLNL, CEA-LIDYL, SLAC, DESY, CERN, and Modern Electron. We acknowledge all WarpX contributors and also thank Remi Lehe for the insightful discussions on mesh refinement methods used in this work.

Software: WarpX,^{30,40} AMReX,^{41,42} matplotlib,⁴³ numpy,⁴⁴ scipy,⁴⁵ yt.⁴⁶

AUTHOR DECLARATIONS

Conflict of Interest

The authors have no conflicts to disclose.

Author Contributions

Revathi Jambunathan: Conceptualization (equal); Investigation (equal); Software (equal); Supervision (equal); Validation (equal); Visualization (equal); Writing – original draft (equal). **Henry Jones:** Formal analysis (equal); Validation (equal); Visualization (equal); Writing – review & editing (equal). **Lizzette Corrales:** Investigation (equal); Validation (equal); Visualization (equal). **Hannah Klion:** Conceptualization (equal); Formal analysis (equal); Investigation (equal); Methodology (equal); Validation (supporting); Visualization

(equal); Writing – original draft (supporting); Writing – review & editing (equal). **Michael Rowan:** Conceptualization (equal); Investigation (equal); Validation (supporting); Writing – original draft (supporting); Writing – review & editing (equal). **Andrew Myers:** Methodology (supporting); Software (supporting); Writing – review & editing (equal). **Weiqun Zhang:** Methodology (supporting); Software (supporting); Writing – review & editing (equal). **Jean-Luc Vay:** Conceptualization (equal); Formal analysis (equal); Investigation (equal); Methodology (equal); Software (equal); Writing – review & editing (equal).

DATA AVAILABILITY

The data that support the findings of this study are openly available in <https://doi.org/10.5281/zenodo.13324091>, Ref. 34.

APPENDIX A: HARRIS-LIKE CURRENT SHEET SETUP

Physical parameters and symbols common to all of our simulations (Table IV).

TABLE IV. Physical parameters and symbols common to all of our simulations.

| Parameter | Symbol | Value |
|------------------------------|-------------|---------------------------------------|
| Background Larmor radius | ρ_c | 4.1×10^{-3} m |
| Background Larmor frequency | ω_c | 7.3×10^{-10} s ⁻¹ |
| Skin depth | λ_e | 0.01 m |
| Current sheet half-width | δ | 0.05 m |
| Background magnetization | σ | 30 |
| Background magnetic field | B_0 | 0.42 T |
| Current sheet number density | n_d | 2.8×10^{17} m ⁻³ |
| Background number density | n_b | 5.6×10^{16} m ⁻³ |
| Current sheet velocity | β_0 | 0.22 c |
| Background temperature | θ_b | 0.15 |
| Current sheet temperature | θ_d | 1.57 |
| Domain half-width (x) | L_x | 10.24 m |
| Domain half-width (z) | L_z | 5.12 m |

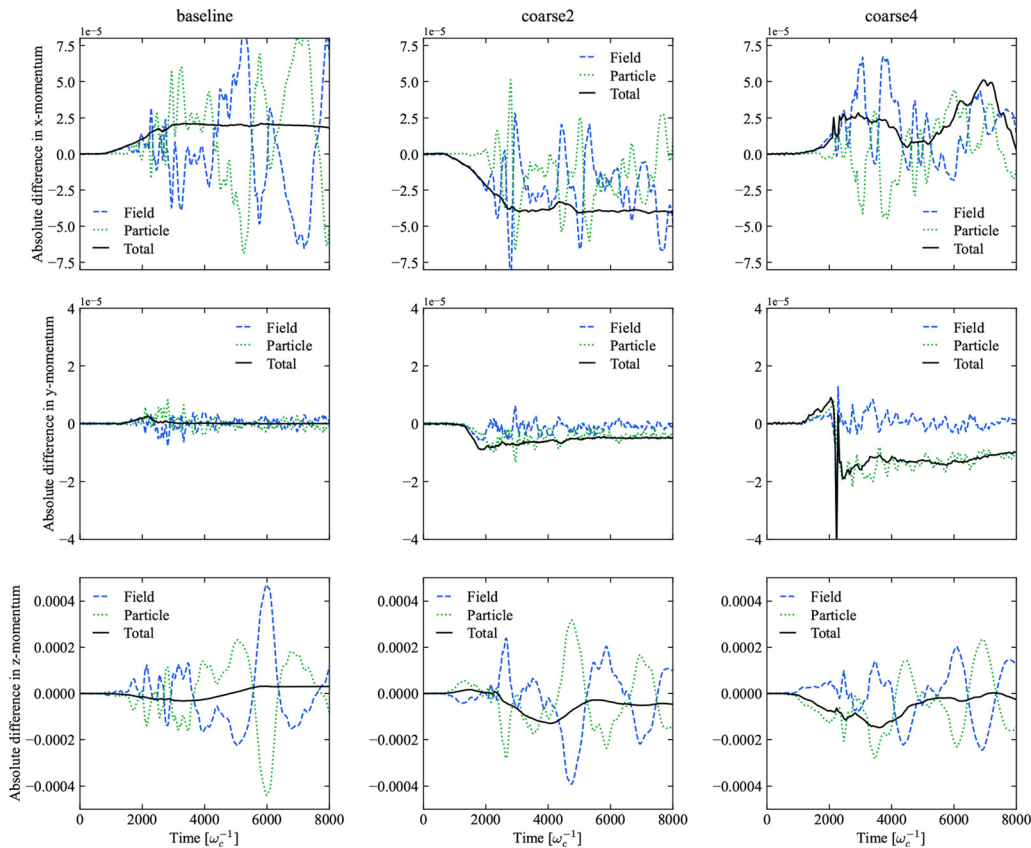


FIG. 12. Comparison of conservation of the x-, y-, and z- components of the total momentum obtained from *baseline*, *coarse2*, and *coarse4* uniform grid simulations. The field momentum is shown by the blue dashed lines, particle momentum denoted by the green dotted lines, and the total momentum is shown by the solid black line.

12 February 2025 20:07:42

APPENDIX B: COMPARISON OF MOMENTUM CONSERVATION

Similar to energy conservation comparisons shown previously in Figs. 5 and 6, a comparison of conservation of x -, y -, and z - components of momentum for the uniform grid simulations, *baseline*, *coarse2*, and *coarse4* is shown in Fig. 12. The absolute difference in field, particle (including both species), and total momentum in the x , y , and z directions (integrated over the whole domain) with their respective values at initialization are shown. Note that the initial field momentum in all directions is zero by construction for the Harris-like current sheet setup. The particle momentum theoretically would cancel out, since inside the current sheet the out-of-plane velocity for the electrons and positrons are opposite, resulting in the initial current sheet. However, the x , y , and z momenta of just electrons summed over the top half of the domain, which includes just the top current sheet is -4×10^{-7} , 5.7×10^{-5} , and -7×10^{-8} (kg m s^{-1}), respectively, at initialization. When summed over the entire domain, due to statistical noise, the momenta do not cancel out

exactly and the total initial momenta (including both electrons and positrons) in the x , y , and z simulations are on the order of 10^{-7} for the *baseline* simulation and approximately in the same range for all the other simulations compared here. [Adding the magnitude of the out-of-plane momentum for both species over the whole domain results in 2.28×10^{-4} (kg m s^{-1})].

It can be seen that for *baseline* simulation, the x -, y -, and z - components of total momentum (which is the sum of the field and particle x - momentum) is conserved within $\pm 2 \times 10^{-5}$, $\pm 1 \times 10^{-5}$, and $\pm 8 \times 10^{-5}$, respectively. As we coarsen the uniform grid by a factor of 2, and 4, in the *coarse2*, and *coarse4* simulations, the absolute difference in the x - total momentum increases to $\pm 4.2 \times 10^{-5}$ and $\pm 5 \times 10^{-5}$, respectively, the absolute difference in the y - total momentum increases to $\pm 1 \times 10^{-4}$, and $\pm 4 \times 10^{-5}$, respectively, and finally, the absolute difference in the z - total momentum increases to $\pm 1.6 \times 10^{-4}$ and $\pm 1.65 \times 10^{-4}$, respectively. These deviations from the initial momentum are within acceptable tolerances for the explicit electromagnetic PIC method used in this work. The momentum is not conserved for the *coarse8* case, similar to the energy non-conservation due to numerical heating.

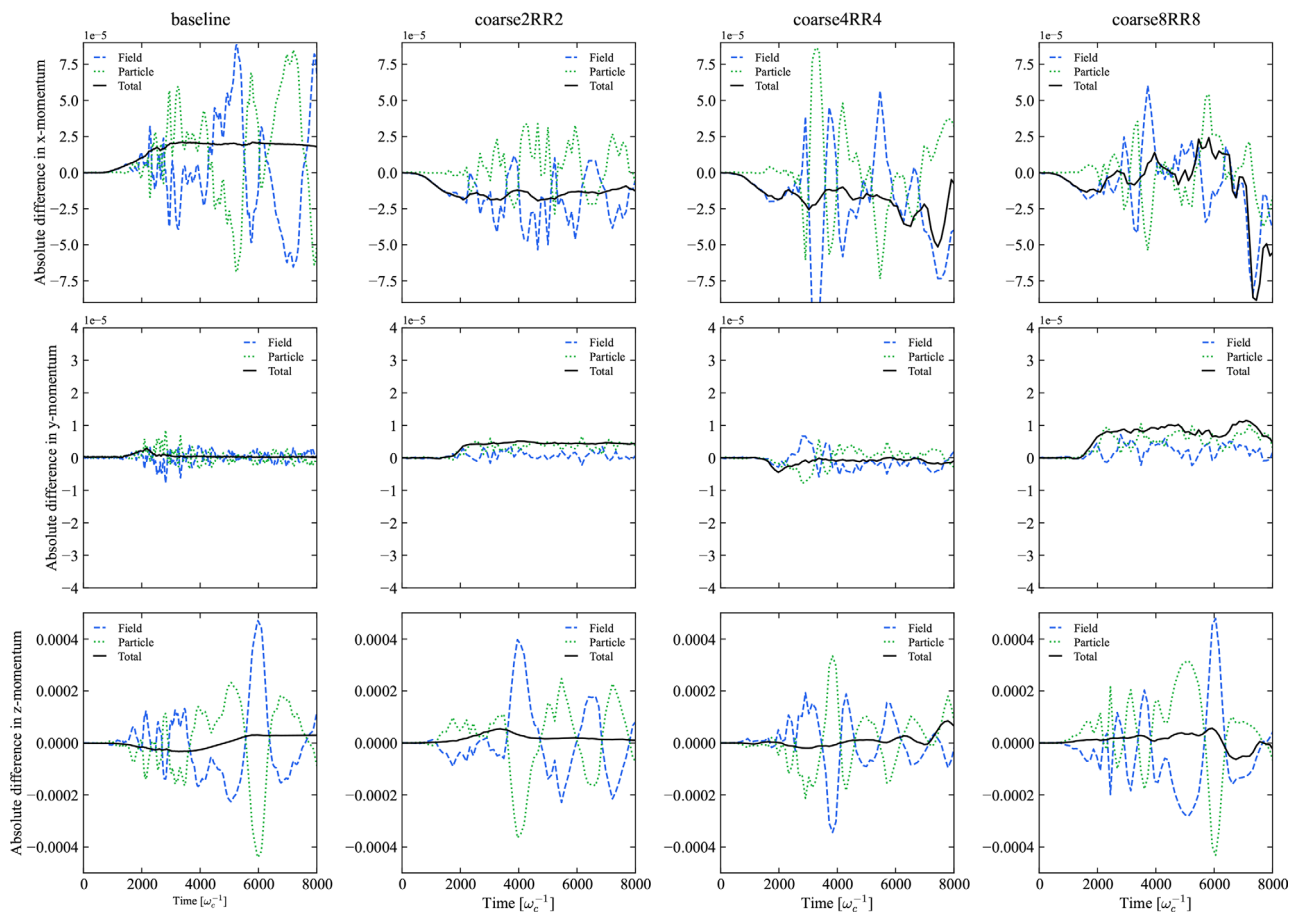


FIG. 13. Comparison of conservation of the x -, y -, and z - components of the total momentum obtained from *baseline*, *coarse2RR2*, *coarse4RR*, and *coarse8RR8* simulations. The field momentum is shown by the blue dashed lines, particle momentum denoted by the green dotted lines, and the total momentum is shown by the solid black line.

12 February 2025 20:07:42

To study the effect of mesh refinement on momentum conservation, the x , y , and z momenta (including field, particle, and total) obtained from the MR simulations, *coarse2RR2*, *coarse4RR4*, and *coarse8RR8* are compared with the high-resolution uniform grid *baseline* simulation. Similar to the uniform grid simulations, the difference in x , y , and z components of the total momentum is within $\pm 4 \times 10^{-4}$. Similar to the energy conservation studies, we do not expect momentum to be conserved with MR since we damp the solution with absorbing layers. Even then, the total momentum in the x , y , and z directions are within $\pm 1 \times 10^{-4}$ of the initial total momentum.

APPENDIX C: EFFECT OF REFINEMENT PATCH SHAPE

In addition to the mesh refinement simulations tabulated in Table II and discussed in Sec. III A, we also studied the effect of patch shape. The MR simulations tabulated in Table II and discussed in Sec. III A involved refinement patches that extended throughout the current sheet up to the edges of the domain boundary in z . We also performed these simulations with a rectangular patch for both the top and the bottom current sheets, where the refinement patch is terminated in the z direction. The valid region of the refinement box extended from $-548.6 < z/\rho_c < 548.6$ for both current sheets. We performed three MR simulations, namely,

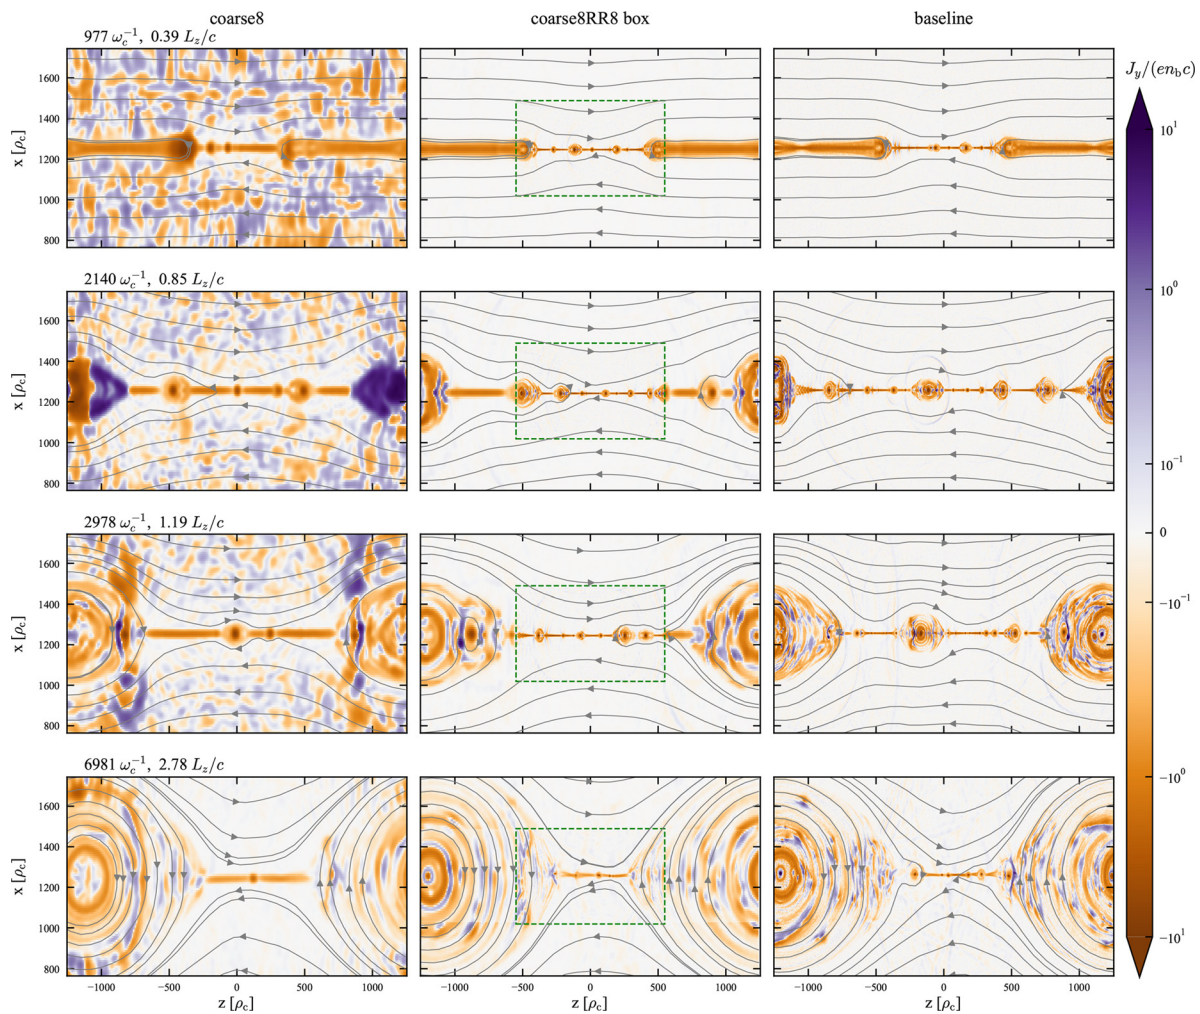


FIG. 14. Comparison of temporal evolution of the normalized out-of-plane current density, $j_y/(en_b c)$, obtained from three 2D grid simulations: uniform *coarse8*, mesh refinement case *coarse8RR8-box*, and the *baseline* uniform grid simulation. The magnetic field lines are shown by the gray lines. The addition of mesh refinement greatly reduces the effects of low resolution in the *coarse8* simulation. The qualitative evolution of *coarse8RR8-box* matches that of our high-resolution *baseline* case. The plotted fields include the valid regions of L_1 and L_0 from which the particles gather electromagnetic fields. The valid region of the mesh refinement patch on L_1 , surrounding the top current sheet, spans $1040 < x/\rho_c < 1520$ and $-548.6 < z/\rho_c < 548.6$, and is indicated by the dashed green lines. [For the bottom current sheet, the valid region of the mesh refinement patch on L_1 spans from $-1520 < x/\rho_c < -1040$ and $(-548.6 < z/\rho_c < 548.6)$]. Both current sheets evolve similarly, and for clarity we zoom in and show only the top current sheet.

coarse2RR2-box, *coarse4RR4-box*, and *coarse8RR8-box*, with the same MR parameters as in Table II, and the same dimensions for the parent grid and refinement ratios, but with a rectangular refinement patch. In Fig. 14, we compare the evolution of the out-of-plane current density, j_y , obtained from the uniform grid *coarse8*, mesh refined *coarse8RR8-box*, and the uniform grid *baseline* simulations. Similar to the observations made in Fig. 4, we can see that the numerical instability observed for the *coarse8* simulation is

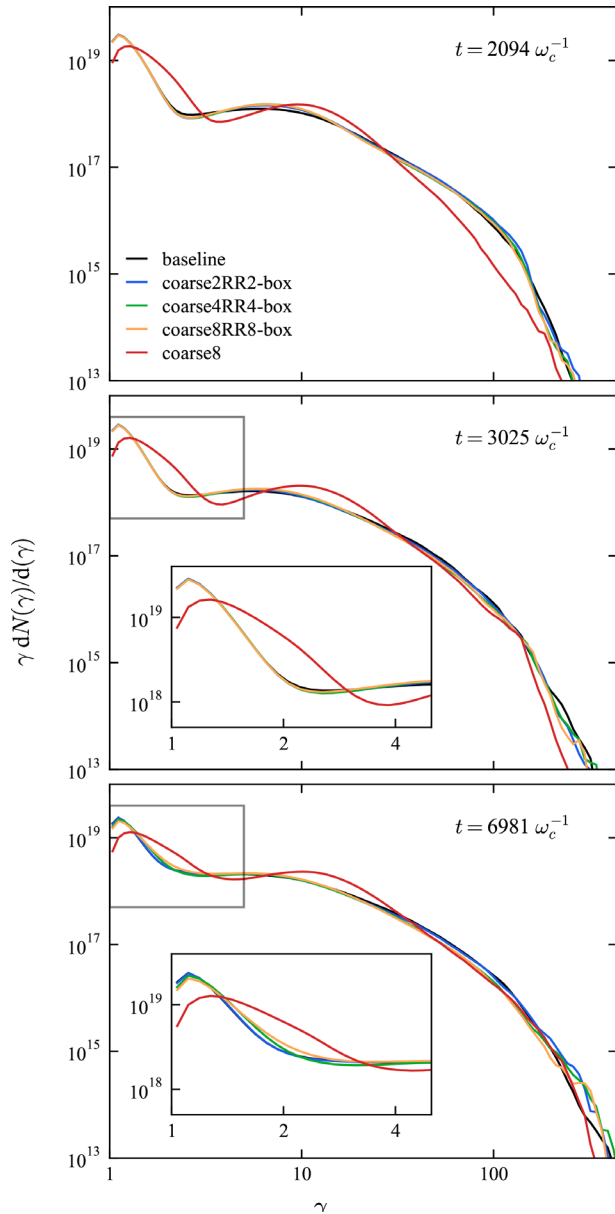


FIG. 15. Comparison of time evolution of particle spectra for the *baseline* (solid black), and MR simulations, *coarse2RR2-box* (solid blue), *coarse4RR4-box* (solid green), and *coarse8RR8-box* (solid orange). The solid red line is obtained from the uniform grid *coarse8* simulation.

mitigated when using the rectangular mesh refinement patch (indicated by the dashed green line) in *coarse8RR8-box*. The evolution of the current sheet, i.e., formation of plasmoids and their merging is captured well by the *coarse8RR8-box* case, similar to the *baseline*. As previously discussed in Sec. III A, we do not expect the evolution of the current sheet to be exactly the same as the *baseline* with the high refinement ratio of 8, where the small wavelength features are not captured by the coarse parent grid. However, the simulations are able to capture the overall reconnection characteristics well. With a rectangular patch in *coarse8RR8-box*, the coarse–fine interface occurs in the upstream region and across the current sheet. These results demonstrate that the MR algorithm discussed in Sec. II B can capture reconnection characteristics while preventing numerical artifacts at the coarse–fine interface. Similar results we obtained for *coarse2RR2-box* and *coarse4RR4-box* simulations, and we highlight only *coarse8RR8-box* as that has the highest aspect ratio of 8, which is more challenging for mesh refinement.

The evolution of particle acceleration from mesh refinement simulations *coarse2RR2-box*, *coarse4RR4-box*, and *coarse8RR8-box* is compared with *baseline* and *coarse8* uniform grid simulations in Fig. 15. Similar to the observations in Fig. 7, adding a mesh refinement patch in *coarse8RR8-box* simulation allows for capturing particle acceleration, similar to *baseline*, and alleviates the numerical heating observed for *coarse8* simulation. Note that *coarse2RR2-box* and *coarse4RR4-box* also agree well with the *baseline* simulation.

APPENDIX D: PARTICLE SPECTRA CONTRIBUTIONS FROM UPSTREAM AND REFINEMENT PATCH REGIONS, AND NUMBER OF PARTICLES PER CELL

Minor differences were observed in the particle spectra (shown in Fig. 8) obtained from *coarse8RR8* simulations initialized with (64×64) (32×32) (16×16) , and (8×8) macroparticles in the coarse region cells, and (8×8) macroparticles per cell in the fine-patch regions. The energy spectra were split into contributions from particles in the refinement patch and from particles outside of the patch, called “upstream.” The spectra from these regions along with the total spectra (which is the sum of split spectra) is shown in Fig. 16 focusing in the $\gamma < 5$ region. It can be seen that at $t \sim 2000\omega_c^{-1}$, the patch (dashed), upstream (dotted), and total spectra (solid) from all simulations agree well. As time progresses, at $t = 6981\omega_c^{-1}$, it can be seen that the difference in the total spectra at $\gamma = 2.5$ is largely due to the particle spectra in the upstream region, where the number of macroparticles is lower at initialization and is further decreased during reconnection as upstream particles are pulled toward the current sheet region covered by the refinement patch. In Fig. 17, we show a histogram of particles per cell for all the cells in the domain, as a fraction of the total number of cells. It can be seen that at $t = 840\omega_c^{-1}$, all simulations have more than four particles per cell everywhere, except for a small fraction of cells in the simulations initialized with (16×16) and (8×8) macroparticles in the coarse region cells. As time progresses, the particles from the upstream coarse region are drawn to the X-points in the current sheets, resulting in reduced particles per cell. For the simulations initialized with (64×64) and (32×32) in the coarse region, for most of the simulation there are more than four particles per cell, and a very small fraction (< 0.01) of cells have less than four




FIG. 16. Time evolution of particle spectra (in the $\gamma < 5$ range) obtained from the *coarse8RR8* simulations initialized with (64×64) (32×32) (16×16) , and (8×8) macroparticles in the coarse region cells, and (8×8) macroparticles per cell in the fine-patch regions. The total spectra (solid) is the sum of spectra from particles in the refinement patch (dotted) and particles outside of the refinement patch called upstream (dashed). This plot focuses on the box inset region shown previously in Fig. 8.

macroparticles per cell. All the cells with less than 30 macroparticles per cell are from the fine patch, which is 8 times more refined than the parent grid. For the simulation initialized with (16×16) macroparticles per cell, until $t = 2560\omega_c^{-1}$, most cells have more than five particles per cell, and a small fraction < 0.05 have less than four macroparticles per cell. Near the end of reconnection, at $t = 6000\omega_c^{-1}$, the (16×16) case has 5% of total cells with one or less macroparticles per cell. For the simulation initialization with (8×8) macroparticles in the coarse region, at the end of reconnection, nearly 18% of the cells have zero and one macroparticles per cell. Note that these simulations also have a significant fraction of




FIG. 17. Particle count distribution showing fraction of total cells with a given number of particles per cell, shown at times throughout reconnection in the *coarse8RR8* simulations initialized with (64×64) , (32×32) , (16×16) , and (8×8) macroparticles in the coarse region cells, and (8×8) macroparticles per cell in the fine-patch regions.

cells with more than 100 macroparticles per cell at the end of reconnection, where particles are trapped in plasmoids. With a linear interpolation for deposition and gathering, the current density for these times is shown in Fig. 18 (middle column). Even with noisy current density and few macroparticles per cell in the upstream regions, the differences in the spectra for $\gamma < 5$ are minor and as mentioned previously, and the non-thermal particle acceleration is captured well by all the simulations as seen from the good agreement in the high γ regions of the spectra in Fig. 8.

APPENDIX E: EFFECT OF MESH REFINEMENT AND PARTICLE RESOLUTION ON CURRENT DENSITY VARIATION

To delineate the effect of grid resolution and particle resolution, uniform grid simulations, *baseline*, *coarse2*, *coarse4*, and *coarse8*, and MR simulations, *coarse2RR2*, *coarse4RR4*, and *coarse8RR8* were performed with same total macroparticles. All MR simulations were found to agree well with the highly refined *baseline* simulation. To improve efficiency as enabled by MR, *coarse8RR8*

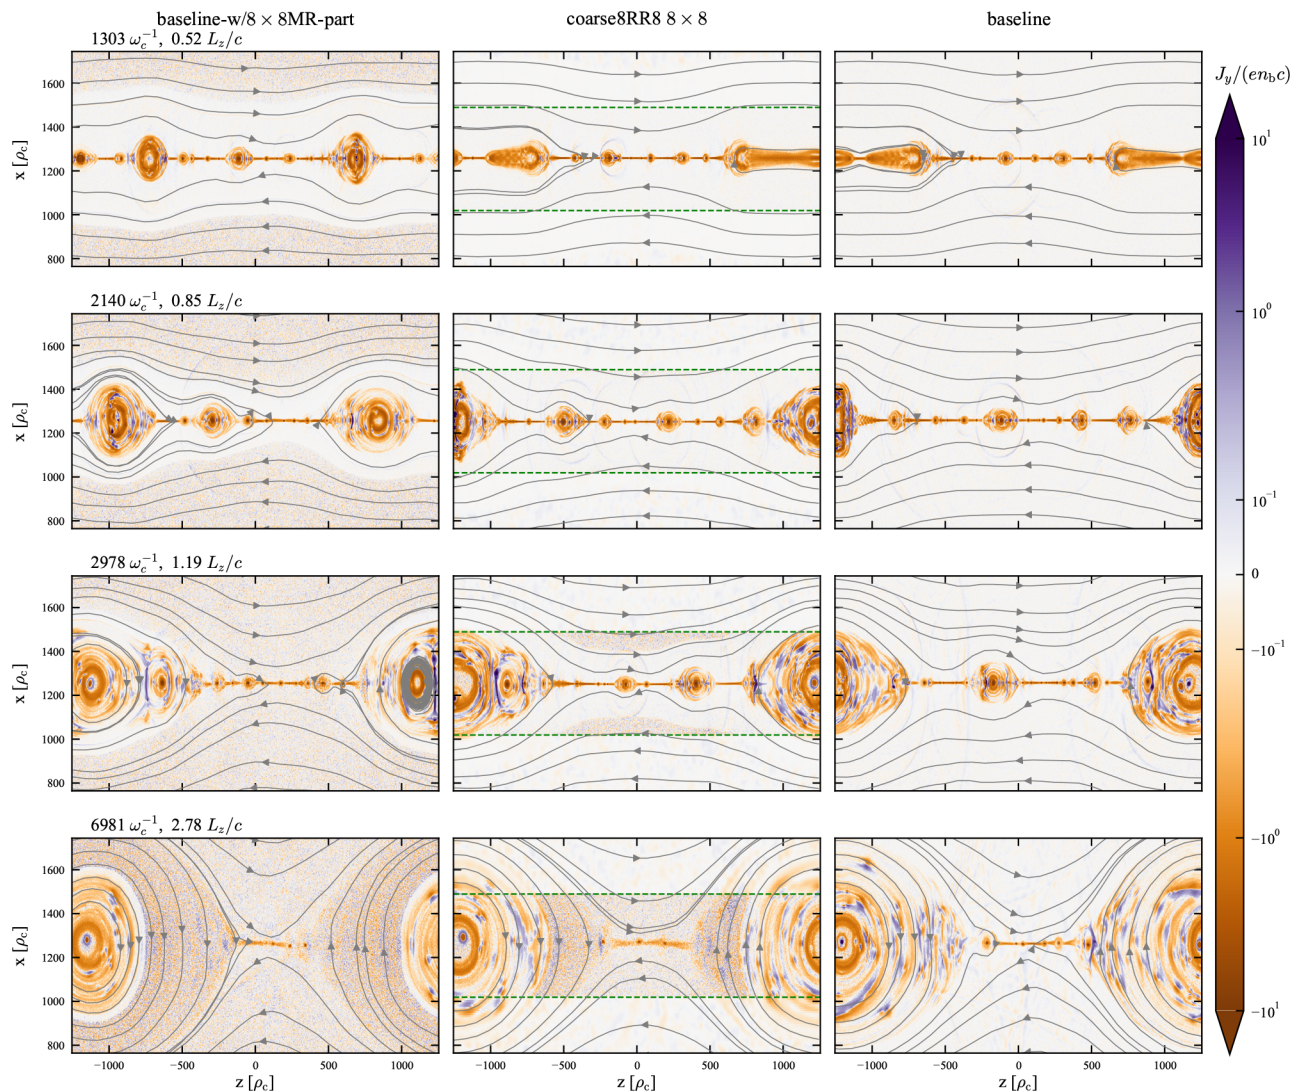


FIG. 18. Comparison of temporal evolution of the normalized out-of-plane current density, $j_y / (en_b c)$, obtained from three 2D grid simulations, namely, uniform *baseline-w/8 x 8MR-part* case, mesh refinement case, *coarse8RR8* with 8×8 macroparticles per coarse and fine cell per species, and the *baseline* uniform grid simulation. The magnetic field lines are shown by the gray lines. Note that for the MR simulation, the fields are shown on the valid regions of L_1 and L_0 from which the particles gather electromagnetic fields for the top current sheet. The valid region of the mesh refinement patch on L_1 , surrounding the top current sheet, spans $1040 < x / \rho_c < 1520$ and is indicated by the dashed green lines. (For the bottom current sheet, the valid region of the mesh refinement patch on L_1 spans from $-1520 < x / \rho_c < -1040$). Both current sheets evolve similarly and for clarity we zoom in and show only the top current sheet.

simulations were initialized with (64×64) (32×32) (16×16) , and (8×8) macroparticles in the coarse region cells, and (8×8) macroparticles per cell in the fine-patch regions. The particle spectra and current density evolution from these cases were comparable with the *baseline* simulation. To further assess the need for mesh refinement, uniform grid simulation, *baseline-w/8* \times *8MR-part*, was performed with grid resolution same as *baseline*, and initial macroparticle distribution same as the *coarse8RR8* simulation with 8×8 macroparticles per coarse and fine cell per species. Even though the total number of macroparticles and their initial distributions were the same, the current sheet evolution and plasmoid formation from the *coarse8RR8* MR simulation compared better with the high-resolution *baseline* simulation than the *baseline-w/8* \times *8MR-part* as shown in Fig. 18. The current density evolution obtained from the *baseline-w/8* \times *8MR-part* simulations shows faster rate of plasmoid formation as seen from the snapshots at $t = 977\omega_c^{-1}$ and $t = 2140\omega_c^{-1}$. This is also consistent with the higher particle acceleration observed from the particle spectra obtained from this simulation, shown previously in Fig. 8. Thus, mesh refinement allows for improved particle statistics by using coarse cells in the upstream region, and fine cells in regions where the plasma skin-depth is high. This results in efficient simulations with fewer total macroparticles while maintaining accuracy. The use of particle splitting and merging in the future will further improve the noise observed in the fine patch of the *coarse8RR8* simulation, further reducing noise with improved particle statistics per cell.

REFERENCES

- ¹H. Ji, W. Daughton, J. Jara-Almonte, A. Le, A. Stanier, and J. Yoo, "Magnetic reconnection in the era of exascale computing and multiscale experiments," *Nat. Rev. Phys.* **4**, 263–282 (2022).
- ²C. K. Birdsall and A. B. Langdon, *Plasma Physics via Computer Simulation* (CRC Press, 2004).
- ³T. Arber, K. Bennett, C. Brady, A. Lawrence-Douglas, M. Ramsay, N. Sircombe, P. Gillies, R. Evans, H. Schmitz, A. Bell *et al.*, "Contemporary particle-in-cell approach to laser-plasma modelling," *Plasma Phys. Controlled Fusion* **57**, 113001 (2015).
- ⁴A. Taflove and S. C. Hagness, *Computational Electrodynamics*, Vol. 28 (Artech House Publishers Norwood, MA, 2000).
- ⁵S. Zenitani and M. Hoshino, "The generation of nonthermal particles in the relativistic magnetic reconnection of pair plasmas," *Astrophys. J. Lett.* **562**, L63–L66 (2001).
- ⁶S. Zenitani and M. Hoshino, "Particle acceleration and magnetic dissipation in relativistic current sheet of pair plasmas," *Astrophys. J.* **670**, 702–726 (2007).
- ⁷B. Cerutti, D. A. Uzdensky, and M. C. Begelman, "Extreme particle acceleration in magnetic reconnection layers: Application to the gamma-ray flares in the Crab Nebula," *Astrophys. J.* **746**, 148 (2012).
- ⁸L. Sironi and A. Spitkovsky, "Relativistic reconnection: An efficient source of non-thermal particles," *Astrophys. J.* **783**, L21 (2014).
- ⁹K. Nalewajko, D. A. Uzdensky, B. Cerutti, G. R. Werner, and M. C. Begelman, "On the distribution of particle acceleration sites in plasmoid-dominated relativistic magnetic reconnection," *Astrophys. J.* **815**, 101 (2015).
- ¹⁰G. R. Werner, D. A. Uzdensky, B. Cerutti, K. Nalewajko, and M. C. Begelman, "The extent of power-law energy spectra in collisionless relativistic magnetic reconnection in pair plasmas," *Astrophys. J. Lett.* **816**, L8 (2016).
- ¹¹M. Melzani, R. Walder, D. Folini, C. Winisdoerffer, and J. M. Favre, "Relativistic magnetic reconnection in collisionless ion-electron plasmas explored with particle-in-cell simulations," *Astron. Astrophys.* **570**, A111 (2014).
- ¹²F. Guo, X. Li, W. Daughton, P. Kilian, H. Li, Y.-H. Liu, W. Yan, and D. Ma, "Determining the dominant acceleration mechanism during relativistic magnetic reconnection in large-scale systems," *Astrophys. J. Lett.* **879**, L23 (2019).
- ¹³H. Hakobyan, M. Petropoulou, A. Spitkovsky, and L. Sironi, "Secondary energization in compressing plasmoids during magnetic reconnection," *Astrophys. J.* **912**, 48 (2021).
- ¹⁴L. Sironi, "Nonideal fields solve the injection problem in relativistic reconnection," *Phys. Rev. Lett.* **128**, 145102 (2022).
- ¹⁵G. R. Werner, D. A. Uzdensky, M. C. Begelman, B. Cerutti, and K. Nalewajko, "Non-thermal particle acceleration in collisionless relativistic electron-proton reconnection," *Mon. Not. R. Astron. Soc.* **473**, 4840–4861 (2018).
- ¹⁶F. Guo, Y.-H. Liu, W. Daughton, and H. Li, "Particle acceleration and plasma dynamics during magnetic reconnection in the magnetically dominated regime," *Astrophys. J.* **806**, 167 (2015).
- ¹⁷M. Petropoulou, L. Sironi, A. Spitkovsky, and D. Giannios, "Relativistic magnetic reconnection in electron-positron-proton plasmas: implications for jets of active galactic nuclei," *Astrophys. J.* **880**, 37 (2019).
- ¹⁸B. Cerutti, G. R. Werner, D. A. Uzdensky, and M. C. Begelman, "Simulations of particle acceleration beyond the classical synchrotron burnoff limit in magnetic reconnection: An explanation of the crab flares," *Astrophys. J.* **770**, 147 (2013).
- ¹⁹K. Nalewajko, Y. Yuan, and M. Chruslińska, "Kinetic simulations of relativistic magnetic reconnection with synchrotron and inverse Compton cooling," *J. Plasma Phys.* **84**, 755840301 (2018).
- ²⁰H. Hakobyan, A. Philippov, and A. Spitkovsky, "Effects of synchrotron cooling and pair production on collisionless relativistic reconnection," *Astrophys. J.* **877**, 53 (2019).
- ²¹K. Schoeffler, T. Grismayer, D. Uzdensky, R. Fonseca, and L. Silva, "Bright gamma-ray flares powered by magnetic reconnection in QED-strength magnetic fields," *Astrophys. J.* **870**, 49 (2019).
- ²²J. M. Mehlhaff, G. R. Werner, D. A. Uzdensky, and M. C. Begelman, "Kinetic beaming in radiative relativistic magnetic reconnection: A mechanism for rapid gamma-ray flares in jets," *Mon. Not. R. Astron. Soc.* **498**, 799–820 (2020).
- ²³B. Cerutti, G. R. Werner, D. A. Uzdensky, and M. C. Begelman, "Three-dimensional relativistic pair plasma reconnection with radiative feedback in the crab nebula," *Astrophys. J.* **782**, 104 (2014).
- ²⁴G. R. Werner and D. A. Uzdensky, "Reconnection and particle acceleration in three-dimensional current sheet evolution in moderately magnetized astrophysical pair plasma," *J. Plasma Phys.* **87**, 905870613 (2021).
- ²⁵L. Sironi and A. M. Beloborodov, "Kinetic simulations of radiative magnetic reconnection in the coronae of accreting black holes," *Astrophys. J.* **899**, 52 (2020).
- ²⁶H. Klion, R. Jambunathan, M. E. Rowan, E. Yang, D. Willcox, J.-L. Vay, R. Lehe, A. Myers, A. Huebl, and W. Zhang, "Particle-in-cell simulations of relativistic magnetic reconnection with advanced maxwell solver algorithms," *Astrophys. J.* **952**, 8 (2023).
- ²⁷J.-L. Vay, J.-C. Adam, and A. Héron, "Asymmetric pml for the absorption of waves. application to mesh refinement in electromagnetic particle-in-cell plasma simulations," *Comput. Phys. Commun.* **164**, 171–177 (2004).
- ²⁸J.-L. Vay, D. P. Grote, R. H. Cohen, and A. Friedman, "Novel methods in the particle-in-cell accelerator code-framework warp," *Comput. Sci. Disc.* **5**, 014019 (2012).
- ²⁹J.-L. Vay, A. Almgren, J. Bell, R. Lehe, A. Myers, J. Park, O. Shapoval, M. Thévenet, W. Zhang, D. P. Grote, M. Hogan, L. Ge, and C. Ng, "Toward plasma wakefield simulations at exascale," in *2018 IEEE Advanced Accelerator Concepts Workshop (AAC)* (IEEE, 2018), pp. 1–5.
- ³⁰L. Fedeli, A. Huebl, F. Boillod-Cerneux, T. Clark, K. Gott, C. Hillairet, S. Jaure, A. Leblanc, R. Lehe, A. Myers, C. Piechurski, M. Sato, N. Zaim, W. Zhang, J. Vay, and H. Vincenti, "Pushing the frontier in the design of laser-based electron accelerators with groundbreaking mesh-refined particle-in-cell simulations on exascale-class supercomputers," in *SC22: International Conference for High Performance Computing, Networking, Storage and Analysis* (IEEE Computer Society, Los Alamitos, CA, USA, 2022), pp. 1–12.
- ³¹R. Lehe, A. Brelly, L. Giacomet, R. Jambunathan, and J.-L. Vay, "Absorption of charged particles in perfectly matched layers by optimal damping of the deposited current," *Phys. Rev. E* **106**, 045306 (2022).
- ³²E. G. Harris, "On a plasma sheath separating regions of oppositely directed magnetic field," *Nuovo Cimento* **23**, 115–121 (1962).

- ³³R. Jambunathan, H. Jones, L. Corrales, H. Klion, M. Rowan, A. Myers, W. Zhang, and J.-L. Vay (2024). “Application of mesh refinement to relativistic magnetic reconnection (Version v1),” Zenodo. <https://doi.org/10.5281/zenodo.13324091>
- ³⁴S. Zenitani, “Loading relativistic Maxwell distributions in particle simulations,” *Phys. Plasmas* **22**, 042116 (2015).
- ³⁵J.-L. Vay, A. Almgren, J. Bell, L. Ge, D. Grote, M. Hogan, O. Kononenko, R. Lehe, A. Myers, C. Ng, J. Park, R. Ryne, O. Shapoval, M. Thevenet, and W. Zhang, “Warp-x: A new exascale computing platform for beam-plasma simulations,” in *3rd European Advanced Accelerator Concepts workshop (EAAC2017)* [Nucl. Instrum. Methods Phys. Res. Sect. A **909**, 476–479 (2018)].
- ³⁶O. Shapoval, J.-L. Vay, and H. Vincenti, “Two-step perfectly matched layer for arbitrary-order pseudo-spectral analytical time-domain methods,” *Comput. Phys. Commun.* **235**, 102–110 (2019).
- ³⁷A. Friedman, R. H. Cohen, D. P. Grote, S. M. Lund, W. M. Sharp, J.-L. Vay, I. Haber, and R. A. Kishek, “Computational methods in the warp code framework for kinetic simulations of particle beams and plasmas,” *IEEE Trans. Plasma Sci.* **42**, 1321–1334 (2014).
- ³⁸P. A. Cassak, Y. H. Liu, and M. A. Shay, “A review of the 0.1 reconnection rate problem,” *J. Plasma Phys.* **83**, 715830501 (2017).
- ³⁹Y.-H. Liu, P. Cassak, X. Li, M. Hesse, S.-C. Lin, and K. Genetreti, “First-principles theory of the rate of magnetic reconnection in magnetospheric and solar plasmas,” *Commun. Phys.* **5**, 97 (2022).
- ⁴⁰Q. Dong, B. Wang, X. Duan, L. Yan, K. Liu, W. Luo, F. Shao, and T. Yu, “A dynamical particle merging and splitting algorithm for particle-in-cell simulations,” *Comput. Phys. Commun.* **294**, 108913 (2024).
- ⁴¹J.-L. Vay, A. Almgren, L. D. Amorim, I. Andriyash, D. Belkin, “D. Bizzozero, A. Blely, S. E. Clark, L. Fedeli, M. Garten, L. Ge, K. Gott, C. Harrison, A. Huebl, L. Giacometti, R. E. Groenewald, D. Grote, J. Gu, R. Jambunathan, H. Klion, P. Kumar, M. Thévenet, G. Richardson, O. Shapoval, R. Lehe, B. Loring, P. Miller, A. Myers, E. Rheaume, M. E. Rowan, R. T. Sandberg, P. Scherpelz, E. Yang, W. Zhang, Y. Zhao, K. Z. Zhu, E. Zoni, and N. Zaim (2023). “Ecp-warp-x/warp-x: 23.04,” Zenodo. <https://doi.org/10.5281/zenodo.7793695>
- ⁴²W. Zhang, A. Almgren, V. Beckner, J. Bell, J. Blaschke, C. Chan, M. Day, B. Friesen, K. Gott, D. Graves *et al.*, “AMReX: A framework for block-structured adaptive mesh refinement,” *J. Open Source* **4**, 1370 (2019).
- ⁴³A. Almgren, V. Beckner, J. Blaschke, C. Chan, M. Day, B. Friesen, K. Gott, D. Graves, A. Huebl, M. Katz, A. Myers, T. Nguyen, A. Nonaka, M. Rosso, J. Sexton, S. Williams, W. Zhang, M. Zingale, and the AMReX Development Team (2023). “Amrex-codes/amrex: Amrex 23.04,” Zenodo. <https://doi.org/10.5281/zenodo.7790123>
- ⁴⁴J. D. Hunter, “Matplotlib: A 2d graphics environment,” *Comput. Sci. Eng.* **9**, 90–95 (2007).
- ⁴⁵C. R. Harris, K. J. Millman, S. J. van der Walt, R. Gommers, P. Virtanen, D. Cournapeau, E. Wieser, J. Taylor, S. Berg, N. J. Smith, R. Kern, M. Picus, S. Hoyer, M. H. van Kerkwijk, M. Brett, A. Haldane, J. Fernández del Río, M. Wiebe, P. Peterson, P. Gérard-Marchant, K. Sheppard, T. Reddy, W. Weckesser, H. Abbasi, C. Gohlke, and T. E. Oliphant, “Array programming with NumPy,” *Nature* **585**, 357–362 (2020).
- ⁴⁶P. Virtanen, R. Gommers, T. E. Oliphant, M. Haberland, T. Reddy, D. Cournapeau, E. Burovski, P. Peterson, W. Weckesser, J. Bright, S. J. van der Walt, M. Brett, J. Wilson, K. J. Millman, N. Mayorov, A. R. J. Nelson, E. Jones, R. Kern, E. Larson, C. J. Carey, Í. Polat, Y. Feng, E. W. Moore, J. VanderPlas, D. Laxalde, J. Perktold, R. Cimrman, I. Henriksen, E. A. Quintero, C. R. Harris, A. M. Archibald, A. H. Ribeiro, F. Pedregosa, P. van Mulbregt, and SciPy 1.0 Contributors, “SciPy 1.0: Fundamental algorithms for scientific computing in Python,” *Nat. Methods* **17**, 261–272 (2020).
- ⁴⁷M. J. Turk, B. D. Smith, J. S. Oishi, S. Skory, S. W. Skillman, T. Abel, and M. L. Norman, “yt: A multi-code analysis toolkit for astrophysical simulation data,” *Astrophys. J. Suppl. Ser.* **192**, 9 (2011).