

Lawrence Berkeley National Laboratory

LBL Publications

Title

Particle-in-cell Simulations of Relativistic Magnetic Reconnection with Advanced Maxwell Solver Algorithms

Permalink

<https://escholarship.org/uc/item/0mx6s893>

Journal

The Astrophysical Journal, 952(1)

ISSN

0004-637X

Authors

Klion, Hannah
Jambunathan, Revathi
Rowan, Michael E
et al.

Publication Date

2023-07-01

DOI

10.3847/1538-4357/acd75b

Copyright Information

This work is made available under the terms of a Creative Commons Attribution License, available at <https://creativecommons.org/licenses/by/4.0/>

Peer reviewed



Particle-in-cell Simulations of Relativistic Magnetic Reconnection with Advanced Maxwell Solver Algorithms

Hannah Klion¹ , Revathi Jambunathan¹ , Michael E. Rowan² , Eloise Yang¹ , Donald Willcox¹ , Jean-Luc Vay¹ ,
Remi Lehe¹ , Andrew Myers¹ , Axel Huebl¹ , and Weiqun Zhang¹

¹ Lawrence Berkeley National Laboratory, 1 Cyclotron Road, Berkeley, CA 94720, USA; klion@lbl.gov

² Advanced Micro Devices, Inc., Santa Clara, CA, USA

Received 2023 April 20; revised 2023 May 17; accepted 2023 May 18; published 2023 July 13

Abstract

Relativistic magnetic reconnection is a nonideal plasma process that is a source of nonthermal particle acceleration in many high-energy astrophysical systems. Particle-in-cell (PIC) methods are commonly used for simulating reconnection from first principles. While much progress has been made in understanding the physics of reconnection, especially in 2D, the adoption of advanced algorithms and numerical techniques for efficiently modeling such systems has been limited. With the GPU-accelerated PIC code WarpX, we explore the accuracy and potential performance benefits of two advanced Maxwell solver algorithms: a nonstandard finite-difference scheme (CKC) and an ultrahigh-order pseudo-spectral method (PSATD). We find that, for the relativistic reconnection problem, CKC and PSATD qualitatively and quantitatively match the standard Yee-grid finite-difference method. CKC and PSATD both admit a time step that is 40% longer than that of Yee, resulting in a $\sim 40\%$ faster time to solution for CKC, but no performance benefit for PSATD when using a current deposition scheme that satisfies Gauss's law. Relaxing this constraint maintains accuracy and yields a 30% speedup. Unlike Yee and CKC, PSATD is numerically stable at any time step, allowing for a larger time step than with the finite-difference methods. We found that increasing the time step 2.4–3 times over the standard Yee step still yields accurate results, but it only translates to modest performance improvements over CKC, due to the current deposition scheme used with PSATD. Further optimization of this scheme will likely improve the effective performance of PSATD.

Unified Astronomy Thesaurus concepts: [Computational methods \(1965\)](#); [GPU computing \(1969\)](#); [Plasma astrophysics \(1261\)](#)

1. Introduction

High-energy radiation is observed from various astrophysical systems, such as pulsar wind nebulae, and from jets in active galactic nuclei (Giannios 2010), X-ray binaries (Tetarenko et al. 2017), and gamma-ray bursts (Piran 2004; Kumar & Zhang 2015). In particular, pulsars produce high-energy gamma-ray flares that evolve too rapidly to be explained by conventional particle acceleration theory (Abdo et al. 2011; Tavani et al. 2011). Magnetic reconnection is often invoked to explain the rapid nonthermal particle acceleration and emission in these systems (Lytikov & Uzdensky 2003; Cerutti et al. 2012; McKinney & Uzdensky 2012; Nalewajko et al. 2015; Petropoulou et al. 2019; Philippov et al. 2019).

During magnetic reconnection, magnetic field energy is converted to particle kinetic energy in the form of both bulk motion and plasma heating. For the highly magnetized astrophysical systems, particle acceleration is caused by relativistic magnetic reconnection, where the unconnected upstream plasma has a magnetic field energy density many times its enthalpy density. The plasma gains relativistic bulk and thermal velocities, and the particle energy distributions develop nonthermal high-energy power-law tails. This population of energized particles are thought to be a source of high-energy emission.

Particle-in-cell (PIC) is a well-established method for studying nonthermal plasma acceleration from first principles (Birdsall & Langdon 1991). Each species in the plasma is modeled with computational particles, which generate currents as they move in the domain. The current is deposited on a spatial grid and is a source term in the Maxwell equations. An electromagnetic field-solve step (also referred to as a Maxwell solve) calculates the electric and magnetic fields. The electromagnetic forces are interpolated to the particle positions, and their positions and velocities are advanced accordingly in time. The PIC method therefore fully captures kinetic particle acceleration, as well as the feedback of the accelerated plasma onto the fields.

A number of PIC studies have investigated relativistic reconnection for collisionless electron–positron (Zenitani & Hoshino 2001, 2007; Cerutti et al. 2012; Sironi & Spitkovsky 2014; Nalewajko et al. 2015; Werner et al. 2016) and electron–ion (Melzani et al. 2014) plasma in two-dimensions. Long thin current sheets become unstable to the tearing-mode instability, resulting in the formation and mergers of chains of trapped plasma, called plasmoids. Simulations have also indicated that reconnection progresses at a rate of approximately 0.1 in such systems, normalized to the reconnecting magnetic field and the Alfvén velocity (Guo et al. 2015; Cassak et al. 2017; Werner et al. 2018). Cerutti et al. (2014) investigated the dispersion relations of the tearing mode (in 2D) and the drift-kink mode that develops in 3D. The fastest-growing tearing mode in simulations has been shown to agree well with analytical expectations (Zenitani & Hoshino 2007). Recent work with PIC simulations has focused on the mechanisms underpinning

the onset of reconnection and phases of particle energization (Guo et al. 2019; Hakobyan et al. 2021; Sironi 2022). The particle energy spectra that result from reconnection show hard power laws that extend to high energies (Guo et al. 2015; Werner et al. 2016, 2018; Petropoulou et al. 2019; Hakobyan et al. 2021). These spectra can then be combined with radiation models to predict observational signatures of reconnection in astrophysical systems (Cerutti et al. 2013; Nalewajko et al. 2018).

While substantial progress has been made in understanding particle acceleration in 2D, and more recently in 3D (e.g., Guo et al. 2015; Zhang et al. 2021; Schoeffler et al. 2023), algorithmic and computational innovation in PIC simulations of such systems has been limited. Virtually all studies employ a finite-difference time-domain Maxwell solver with a staggered Yee grid (Yee 1966; sometimes called FDTD), and only a few have explored the advantages of GPU acceleration for astrophysical PIC simulations (Bussmann et al. 2013; Chien et al. 2020; Xiong et al. 2023). The Yee approach is second-order in both space and time. For certain plasma systems, the numerical dispersion inherent to the method can lead to significant errors. With the Yee solver using a time step at the Courant limit, the numerical dispersion error is maximal along the axes and zero along the principal diagonals of the cells. To obtain solutions with less dispersion, we need alternate solvers, eventually based on higher-order methods. Cole and Kärkkäinen proposed a nonstandard finite-difference approach to mitigate the effects of numerical dispersion along the principal axes when using the time step at the Courant limit (Cole 1997, 2002; Kärkkäinen et al. 2006), which Cowan et al. (2013) extended to noncubic cells. This combination is known as the Cole–Kärkkäinen–Cowan (CKC) scheme. While numerical dispersion can be suppressed with CKC along the main axes, it remains at other angles.

Higher-order methods, including Fourier-based spectral methods, can be used to reduce dispersion even further. Pseudo-Spectral Analytical Time Domain (PSATD; Haber et al. 1973; Vay et al. 2013) is one such method, which enables arbitrary-order accuracy that can be set at runtime. Because they are finite-difference schemes, Yee and CKC are only numerically stable when the time step is below a value set by the Courant limit (see Section 2.3); on the other hand, PSATD is based on analytical integration in Fourier space and has no such constraint. In this paper, we compare the performance and accuracy of the two nonstandard approaches, CKC and PSATD, for relativistic reconnection with the widely adopted Yee scheme. While PSATD does not have a Courant *stability* limit on the time step with regard to the Maxwell solve, an overly large time step may still reduce the accuracy of the simulation (as particles traveling close to the speed of light may travel over a cell size in a single time step). We therefore explore the performance and accuracy of PSATD with time steps above the light travel time across a cell.

The PIC algorithm captures reconnection physics accurately from first principles, but it can be computationally expensive, especially when performing high-resolution simulations with higher-order particle shape factors. Graphics processing units (GPUs) can offer remarkable acceleration over conventional CPU architectures for a number of scientific applications, including PIC (Bussmann et al. 2013; Germaschewski et al. 2016; Chien et al. 2020; Vay et al. 2020; Myers et al. 2021). We use the GPU-accelerated electromagnetic PIC code, WarpX

(Vay et al. 2020; Myers et al. 2021). It has excellent full-machine scaling at leadership-class computing facilities, including Summit and Perlmutter (NVIDIA GPUs) and the world’s first reported exascale machine, Frontier (AMD GPUs; Fedeli et al. 2022). The code is built on the AMReX (Zhang et al. 2019) framework, which supports MPI+X parallelism, where MPI enables inter-rank communications, and X corresponds to an interface such as OpenMP, CUDA, HIP, or SYCL for parallel programming on multicore CPUs or GPUs. PIConGPU (Huebl 2019), VPIC 2.0 (Bird et al. 2022), and the Plasma Simulation code (PSC; Germaschewski et al. 2016) also employ similar strategies that enable performance portability and allow scaling to multiple GPU nodes. Nonrelativistic magnetic reconnection has been used as a comparison case to validate multiple GPU-accelerated PIC codes, including PSC and sputniPIC (Chien et al. 2020), which can make use of a single node with multiple GPUs, and a CUDA Fortran single-GPU PIC code (Xiong et al. 2023).

In this paper, we perform first-of-their-kind 2D GPU simulations of relativistic reconnection with the advanced Maxwell solvers CKC and PSATD. Because these have never before been used for relativistic reconnection, we validate our results by comparing against simulations that use the conventional Yee solver, which has been well studied for 2D systems. In particular, we focus on the evolution of the current sheet structures, the particle-field energy balance, particle energy spectrum, and reconnection rate. We investigate the accuracy-based constraints on PSATD time steps by parameterizing the time step relative to the standard Courant limit of the finite-difference simulations. As with the simulations with different solvers, we compare our results with different time steps to the baseline Yee simulations. For the same cell size, the Courant limit for CKC admits a longer time step than Yee. PSATD is unconditionally stable with limitations on accuracy that may be imposed by other time-integration algorithms in the PIC simulation. Both CKC and PSATD may allow for faster simulations, so we compare the time to solution for these advanced solvers and assess the performance gains of a large time step with PSATD while holding the cell size constant.

The rest of the paper is organized as follows: in Section 2, we describe our simulations, including details about the initial configuration of the current sheets, the perturbation to trigger reconnection, and the algorithmic and numerical parameters. In Section 3, we discuss the accuracy and performance results from using different Maxwell solvers. In Section 4, we present the results from increasing the time step past the Courant limit with PSATD. We summarize and discuss future directions for our work in Section 5.

2. Simulation Setup

2.1. Harris-like Sheets

The simulations are initialized with two pair-plasma Harris-like current sheets (Harris 1962) in equilibrium. This section describes the configuration, which is summarized in the diagram in Figure 1. Full details of the derivation are in the Appendix, and summaries of important parameters can be found in Table 1 (scaled units) and Table 4 (SI units). Input files to replicate the simulations and analysis are available online.³

³ [10.5281/zenodo.7847375](https://doi.org/10.5281/zenodo.7847375)

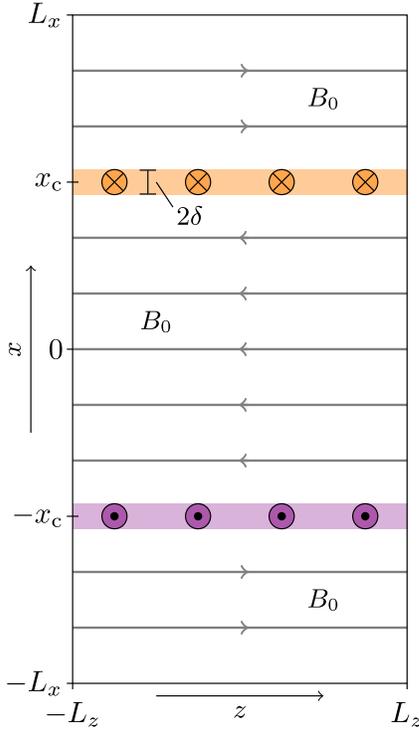


Figure 1. Diagram of the equilibrium double Harris-like sheet, before the application of the perturbation. The current sheets are centered at $x = \pm x_c$ and have half-widths of δ in the x -direction. They extend infinitely in the y - z plane. The current at $x = x_c$ ($x = -x_c$) is into (out of) the plane. The upstream magnetic field has magnitude B_0 . See Equations (1), (3), and (6) for expressions for the number densities, particle velocities, and magnetic fields.

Table 1

Physical Parameters and Symbols Common to All of Our Simulations

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_c	$2.45 \rho_c$
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	$2195 \rho_c$
Domain half-width (z)*	L_z	$1058 \rho_c$

Note. Quantities marked with * are freely chosen; others are derived.

Far from the current sheets, in the upstream, the magnetic field is $\mathbf{B} = \pm B_0 \hat{z}$. Its sign changes at the current sheets. Our principal unit of time will be the inverse upstream electron gyrofrequency, $\omega_c^{-1} \equiv m_e/(eB_0)$, where c is the speed of light, e is the elementary charge, and m_e is the electron mass. Our base unit of length will be $\rho_c = c\omega_c^{-1}$, which is the nominal relativistic Larmor radius. The current sheets are chosen to have half-width $\delta = 12.15 \rho_c$, and they extend in the y - z plane. They are centered at $x = \pm x_c \equiv \pm L_x/2$, where L_x is the half-width of the domain in the x -direction, which spans the interval $[-L_x, L_x]$.

We establish spatial distributions of number density $n(x)$ and bulk velocity $\beta(x)$:

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

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

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

Each species has number density $n(x)$. Positrons have velocity $\beta(x)$, and electrons have velocity $-\beta(x)$. The number density in the upstream region is n_b , and that in the current sheet is n_d . The bulk velocity at the center of the current sheet is β_0 , and \hat{y} is the unit vector parallel to the y -axis. From Ampere's Law, $\nabla \times \mathbf{B} = \mu_0 \mathbf{J}$:

$$-\frac{\partial}{\partial x} B_z = 2\mu_0 e n(x) \beta(x) c, \quad (4)$$

where \mathbf{B} is the magnetic field, \mathbf{J} is the current density, and μ_0 is the vacuum permeability. The factor of 2 in front of J_y is due to there being two species, positrons and electrons, that contribute to the total current density. Except where otherwise indicated, all quantities are given in the observer frame. We solve Equation (4) for $B_z(x)$ with the boundary conditions

$$B_z(\infty) = -B_z(0) = B_0. \quad (5)$$

This gives

$$B_z(x) = -\frac{2B_0}{\left(\frac{\pi}{2} + \frac{n_d}{n_b} - 1\right)} \times \left(\arctan \tanh \frac{x + x_c}{2\delta} - \arctan \tanh \frac{x - x_c}{2\delta} - \frac{\pi}{4} + \frac{1}{2} \left[\frac{n_d}{n_b} - 1 \right] \left[\tanh \frac{x + x_c}{\delta} - \tanh \frac{x - x_c}{\delta} - 1 \right] \right), \quad (6)$$

where

$$B_0 = 2\mu_0 e n_b \beta_0 c \delta \left(\frac{\pi}{2} + \frac{n_d}{n_b} - 1 \right). \quad (7)$$

We choose the upstream ‘‘cold’’ magnetization $\sigma \equiv B_0^2/(\mu_0 n_b m_e c^2) = 30$. This is somewhat different from the typical relativistic ‘‘hot’’ magnetization $\sigma_h = B_0^2/(\mu_0 h)$, with h the relativistic enthalpy density. The temperature in the upstream will be very mildly relativistic, with $h \lesssim 1.5 n_b m_e c^2$, so the hot magnetization is around 20, and $\sigma \sim \sigma_h$. Based on either definition of magnetization, reconnection will proceed in the highly relativistic regime.

We choose the current sheet overdensity to be a factor of five, such that $n_d = 5n_b$. The skin depth in the current sheet is, by definition,

$$\lambda_e \equiv \frac{c}{\omega_p} = c \sqrt{\frac{m_e \epsilon_0}{n_d e^2}}, \quad (8)$$

where ω_p is the plasma frequency in the current sheet. From the quantities fixed thus far, $\lambda_e = 2.45 \rho_c$. Combining Equation (7) with expressions for σ and λ_e gives an expression for the

velocity at the center of the current sheet:

$$\beta_0 = \frac{1}{\left(\frac{\pi}{2} + \frac{n_d}{n_b} - 1\right)} \frac{\lambda_e \sqrt{\frac{n_d}{n_b}} \sigma}{2\delta} = 0.22 c. \quad (9)$$

We calculate the temperature profile from pressure balance in the x -direction in the observer's (unprimed) frame:

$$P_{\text{gas}}(x) + P_{\text{mag}}(x) = C, \quad (10)$$

where C is a constant. The gas pressure $P_{\text{gas}} = T^{xx}$ where $T^{\mu\nu}$ is the stress-energy tensor. In the fluid (primed) frame, $T'^{xx} = 2n'\theta mc^2$, where $\theta = k_B T / (m_e c^2)$ is the dimensionless temperature, and n' is the number density in the fluid frame. The fluid bulk velocity is in the y -direction, so $T'^{xx} = T^{xx}$. The observer-frame number density is $n = \gamma n'$, where $\gamma = 1/\sqrt{1 - \beta^2}$. Substituting,

$$P_{\text{gas}} = T^{xx} = 2nmc^2\theta\gamma^{-1} \quad (11)$$

$$= 2nmc^2\theta\sqrt{1 - \beta^2}. \quad (12)$$

The magnetic pressure,

$$P_{\text{mag}} = \frac{B^2}{2\mu_0}, \quad (13)$$

does not need to be transformed, given that \mathbf{B} is already in the observer's frame.

We now have

$$2nmc^2\theta\sqrt{1 - \beta^2} + \frac{B^2}{2\mu_0} = C. \quad (14)$$

Far from the current sheets,

$$n = n_b \quad (15)$$

$$\theta = \theta_b = \sigma/\eta \quad (16)$$

$$\beta = 0 \quad (17)$$

$$\mathbf{B} = \mathbf{B}_0, \quad (18)$$

where we have chosen $\eta \equiv 200$, so $\theta_b = 0.15$. Solving for $\theta(x)$:

$$\theta(x) = \frac{\sigma(4 + \eta)/\eta - [B_z(x)/B_0]^2}{4[n(x)/n_b]\sqrt{1 - \beta(x)^2}}. \quad (19)$$

The temperature at the center of the current sheet is calculated by evaluating Equation (19) at $x = \pm x_c$, giving $\theta_d = 1.57$.

Electron–positron pairs are initialized at the start of the simulation; they are arranged such that they are uniformly spaced, and their momenta are initialized by sampling from a Maxwell–Jüttner distribution at the local temperature θ (Zenitani 2015).

We note that the setup described here is slightly nonstandard. We have used spatially varying distributions of number density (Equation (1)), bulk velocity (Equation (3)), and temperature (Equation (19)) to represent both the upstream plasma and current sheets. More commonly, the current sheet plasma (with fixed hot temperature θ_d and drift velocity β_0) is overlaid on a domain filled with the upstream plasma (temperature θ_b , number density n_b). Consequently, at initialization, the plasma at the center of the current sheets will have a two-temperature distribution. The number density of the drifting plasma is the only quantity that varies spatially. Ultimately, the resulting current distributions are very similar, as are the induced

magnetic fields, and both configurations are in equilibrium. The primary difference is that the values of θ_d and β_0 used here differ from those of Werner et al. (2018), whose Harris sheets may otherwise appear identical to ours.

2.2. Perturbation

We add a perturbation to the equilibrium configuration in order to control the location and number of x -points. We model our perturbation after that of Werner et al. (2018), who add a one percent sinusoidal perturbation to the vector potential, which decreases the magnetic pressure at the z -axis just outside of each of the current sheets:

$$A_y = \int B_z(x) dx \times \left[1 - 0.01 \cos^{51}\left(\frac{\pi}{L_z} z\right) \cos^2\left(\frac{\pi(x - x_c)}{L_x}\right) \right]. \quad (20)$$

We set the constant of integration equal to zero. In most studies that use this perturbation, the cosine in z is not raised to any power. We find that instead raising it to the 51st power narrows the region being perturbed and thus reduces the number of initial X -points in our configuration. This helps to ensure a more consistent comparison between the simulations. However, after the initial phases of reconnection, the form of the perturbation does not affect the results. The exact choice of power is arbitrary, though it must be odd in order to maintain the sign of the cosine term.

We initialize the fields in the simulation with $\mathbf{B} = \nabla \times (A_y \hat{\mathbf{y}})$. The simulations will therefore all start with $\nabla \cdot \mathbf{B} = 0$, which will be conserved by all numerical methods used here.

2.3. Particle-in-cell Simulations

Simulations of relativistic electron–positron pair-plasma reconnection were performed using the electromagnetic PIC code WarpX (Vay et al. 2020; Myers et al. 2021). The two-dimensional domain is discretized with a uniform grid with a cell size of $\Delta x = \Delta z = \lambda_e/4$. Both the skin depth in the current sheet and upstream Larmor radius $\rho_c = 1.63\Delta x$ will be resolved. Near the end of the simulation, the expectation is that energized particles will have Larmor radii of $\rho_{c,f} = \sigma\rho_c$. Because the grid resolves the smaller initial Larmor scale, it will continue to resolve it throughout the simulation as it grows.

The simulations were initialized with 64 particles per species per cell, evenly spaced in both directions. At this particle density, the overall results are converged, and statistical noise is not significant. The grid extends to $[-L_x, L_x] \times [-L_z, L_z]$, with $L_x = 2195 \rho_c = 73.2 \rho_{c,f}$ and $L_z = 1058 \rho_c = 35.3 \rho_{c,f}$. The simulations have a resolution of (7168×3456) cells, and a total of 3.1 billion particles. Periodic boundary conditions are applied on all domain edges. The full domain is $\sim 70 \rho_{c,f}$ in the smaller dimension, situating our simulations in the “large-domain” regime, where particle acceleration is expected to be primarily limited by the plasma properties rather than the box size (Werner et al. 2016).

The choice of electromagnetic field solver sets the time step. In this paper, three choices are explored: the standard finite-difference time-domain method on a Yee grid (here called “Yee”; Yee 1966), a nonstandard finite-difference

Cole–Kärkkäinen solver with Cowan coefficients (CKC; Cole 1997, 2002; Kärkkäinen et al. 2006; Cowan et al. 2013), and a pseudo-spectral analytical time domain (PSATD) method (Haber et al. 1973; Vay et al. 2013) with a 16th-order stencil (Vincenti & Vay 2016). On our uniform 2D grid with square cells, Yee has a maximum stable time step Δt_C set by the Courant–Friedrichs–Lewy (CFL) condition: $\Delta t_C = \Delta x / (c\sqrt{2})$. CKC admits a larger time step $\Delta t_C = \Delta x / c$. Unlike the finite-difference schemes, PSATD does not have a theoretical limit on time step, because it is unconditionally stable. In that case, by default, WarpX sets the time step to be the smallest cell light-crossing time $\Delta t_C = \Delta x / c$ as a Courant-like time step. For the remainder of the paper, we define the “CFL factor,” simply denoted “CFL,” as a multiplying factor to the CFL time step limit Δt_C of a given Maxwell solver, such that the time step of a simulation is given by $\Delta t = \text{CFL} \times \Delta t_C$.

Two sets of simulations were performed. The first set compares the three electromagnetic field solvers with a time step $\Delta t = 0.95\Delta t_C$ (CFL factor of 0.95). The second studies the effect of time steps that exceed Δt_C with PSATD. While there is no formal stability limit on the time step, the accuracy of our simulations can be expected to deteriorate for CFL factors larger than some value to be determined. Using a systematic study of the evolution of the simulations for a range of CFL factors, we will determine if and in what cases such a practical time step ceiling exists.

By default, each Maxwell solver is paired with a current deposition algorithm that guarantees charge conservation. The Esirkepov deposition scheme (Esirkepov 2001) is charge-conserving when used with either Yee or CKC (Vay et al. 2011). However, Gauss’s law is not satisfied when Esirkepov deposition is combined with high-order PSATD. Vay et al. (2013) developed a deposition scheme (hereafter referred to as “Vay”) that preserves $\nabla \cdot \mathbf{E} = \rho / \epsilon_0$ when used with PSATD. Here, ρ is the charge density, so we also refer to algorithmic combinations that satisfy this equation as “charge-conserving.” We exclusively use Esirkepov with Yee and CKC. When not otherwise specified, PSATD is coupled with the Vay deposition. These three combinations form our main comparison. When measuring the time to solution, we find it instructive to decouple the performance differences from the Maxwell solvers and the current deposition schemes. To do so, we compare the main three simulations against a PSATD simulation with Esirkepov deposition. It happens that, for this problem, PSATD+Esirkepov produces a correct result despite violating Gauss’s law. Both deposition schemes use cubic splines for the particles, and once on the grid, currents are smoothed with a single-pass bilinear filter.

In all cases, the field gather operation also uses cubic splines. We use a relativistic second-order Boris push to advance the particle positions (Boris 1970).

3. Results: Maxwell Solvers

3.1. Energy Conversion and Particle Acceleration

The PIC simulations of reconnecting Harris-like sheets (described in Section 2) were conducted until reconnection has completed and magnetic energy is no longer being converted to particle kinetic energy, at around $t = 8000\text{--}9000 \omega_c^{-1} \approx 4 \times (2L_z/c)$. The qualitative current evolution of the Yee simulation is captured in Figure 2. The

other simulations evolve similarly. Shortly after the start of the simulation, both the top and bottom current sheets (left and right columns in Figure 2) collapse due to the perturbation (Equation (20)) that lowers the magnetic pressure just above and below the current sheet (second row in Figure 2). Quasi-circular structures of trapped plasma and current form, called plasmoids. They inherit their average current from the current sheet where they form. Several magnetic X -points form between the plasmoids in the current sheets. The plasmoids move outward along the current sheet, and occasionally they merge, as highlighted in the blue box in the fourth row of Figure 2. This merger creates a new current sheet antiparallel to the bulk current in the plasmoids, and that extends perpendicularly from the original. This is the site of “secondary reconnection.” At $t \approx 8000 \omega_c^{-1}$, both primary and secondary reconnection have ended, leaving a single plasmoid at $z \approx \pm L_z$ (bottom row in Figure 2).

In all three of the Maxwell solvers studied, magnetic reconnection proceeds at approximately the same rate and with the same structures. Figure 3 compares the plasma structures around the upper current sheet at $t = 1470 \omega_c^{-1}$ for simulations that use the Yee, CKC, and PSATD solvers. Several small plasmoids have formed in each current sheet, with a single larger one forming at the edge of the domain. The exact positions and sizes of the smaller plasmoids differ between the solvers, but the current sheet fragments in a similar way in all cases.

Energy conversion and particle acceleration also proceed similarly with all three solvers. When reconnection saturates at around $t = 6000 \omega_c^{-1}$, about 40% of the energy in electromagnetic fields has been converted to particle kinetic energy (Figure 4). This includes energy associated with both thermal and bulk motion. Energy conversion proceeds nearly identically for the first $1800 \omega_c^{-1}$, after which there are slight differences between the numerical methods. This initial interval of identical evolution appears to be one of linear growth of the tearing-mode instability. The amplitudes of the fastest-growing spatial Fourier modes are approximately exponential in this interval. At around $t = 1800 \omega_c^{-1}$, the exponential growth stops, suggesting the beginning of a nonlinear phase of evolution. This nonlinearity coincides with the appearance of small but noticeable differences between the energy conversions when using the different solvers, suggesting that this divergence is a consequence of nonlinear evolution. PSATD+Esirkepov shows a similar result.

All simulations conserve energy to within one part in 2000, which is well within an acceptable level of energy nonconservation. The CKC simulation loses 4×10^{-4} of the initial total energy, slightly more than the 3.5×10^{-4} and 3×10^{-4} lost by the Yee and PSATD simulations, respectively.

The three methods and PSATD+Esirkepov also produce quantitatively similar particle acceleration. The highest particle γ at the start of the simulation is around 30, while by the end of reconnection at $t = 8500 \omega_c^{-1}$, the fastest particles have $\gamma = 500$, over ten times higher. The majority of the energy in the domain is in particles with $\gamma \lesssim \sigma$. Prior work has found that $dN/d\gamma \propto \gamma^{-\alpha}$ with $\alpha \approx 1\text{--}3$ (Guo et al. 2015; Werner et al. 2016, 2018). This corresponds to power laws with indices between 0 and -2 for the Lorentz factor distribution. Our energy distributions roughly follow this slope (Figure 5).

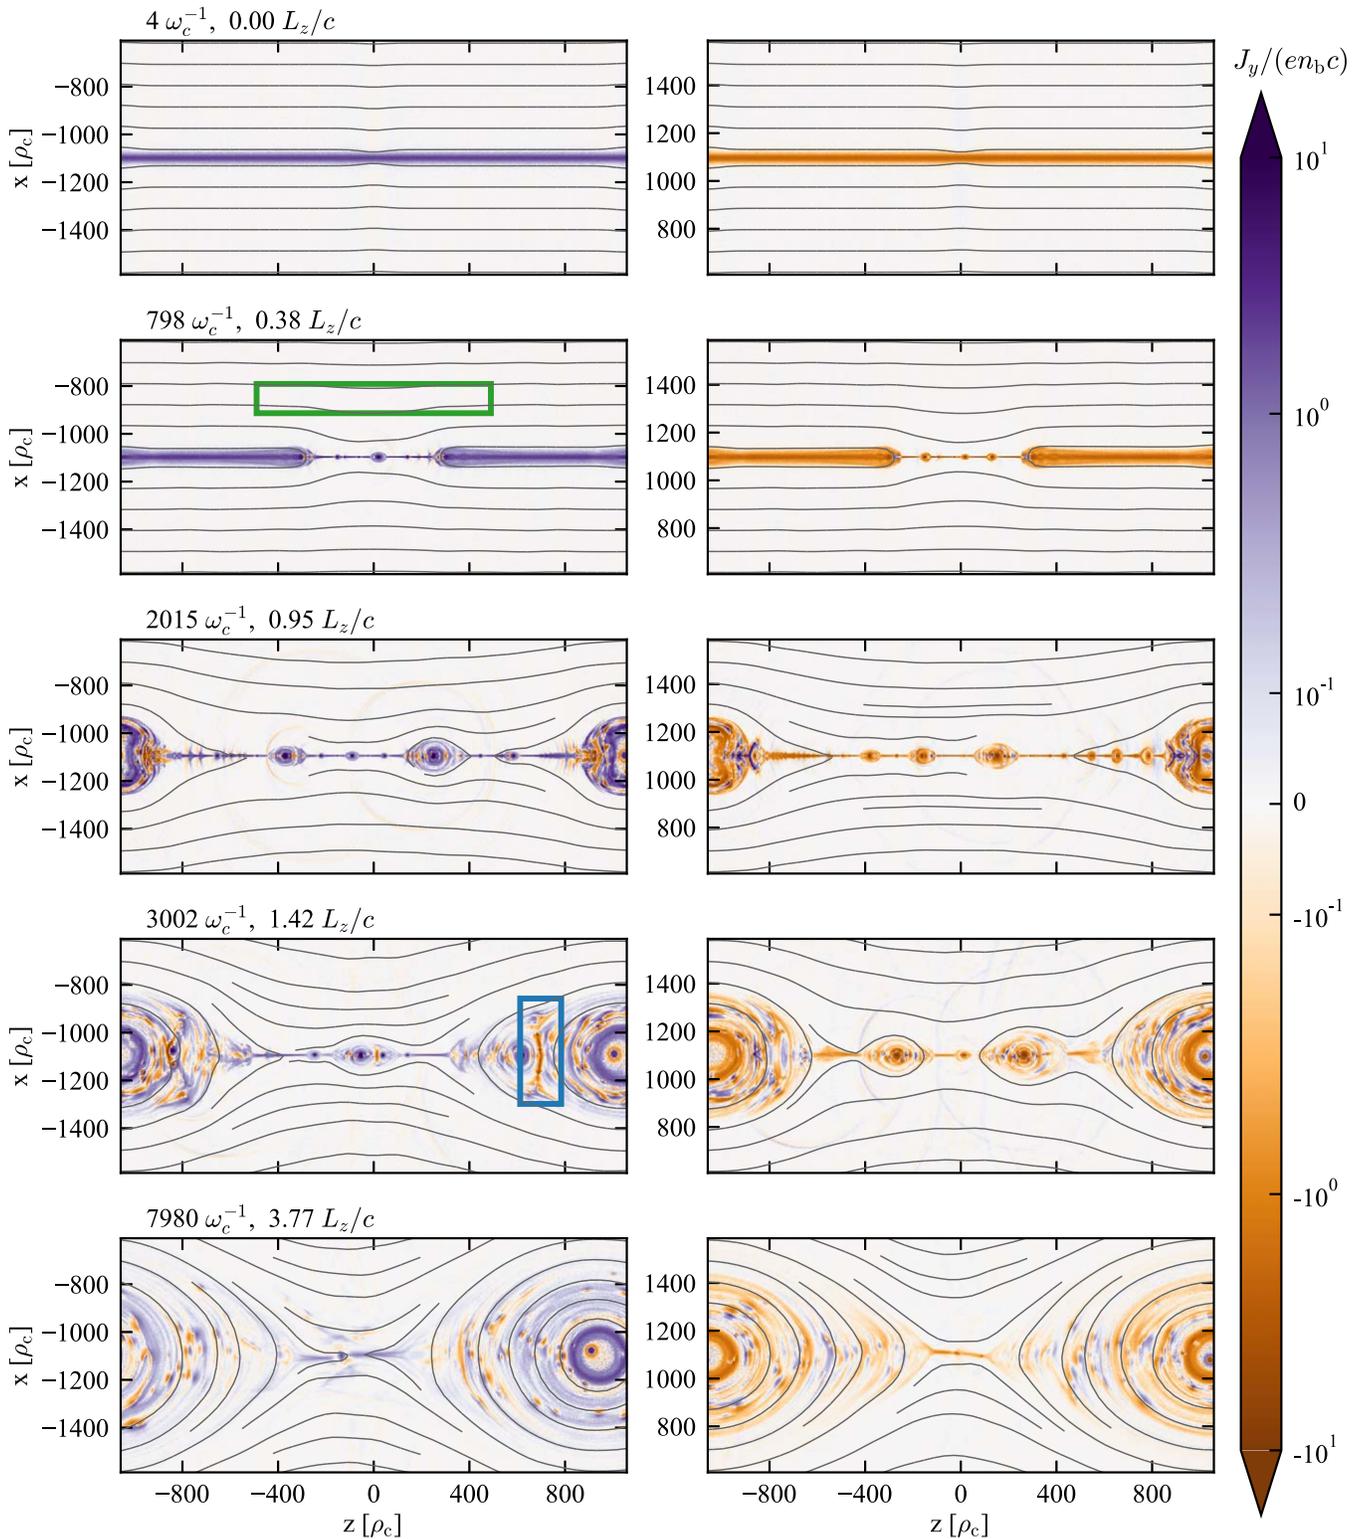


Figure 2. Time evolution of top (left column) and bottom (right column) current and in-plane magnetic field (black lines) in the Yee simulation. The small perturbation to the magnetic fields at $z = 0$ at the initial current sheets leads the current sheet to collapse at that point, thinning it out and causing reconnection to start. The current sheet fragments into plasmoids, which move away from the center and merge with one another, causing secondary reconnection (e.g., region marked with blue box). At the end, there is a single large plasmoid, and reconnection ends. The green box shows the region used to calculate average inflow velocity in Section 3.2.

3.2. Reconnection Rate

The dimensionless reconnection rate

$$\beta = -\frac{1}{v_A B_0 L_x} \frac{d\Phi}{dt} \quad (21)$$

parameterizes much of the linear theory of reconnection (e.g., Werner & Uzdensky 2021). The Alfvén velocity is v_A , and Φ is the unreconnected flux. Directly measuring the amount of unreconnected flux is difficult, so we instead use the approximation $\beta \approx v_{\text{in}}/v_{\text{out}}$, where v_{in} is the inflow velocity

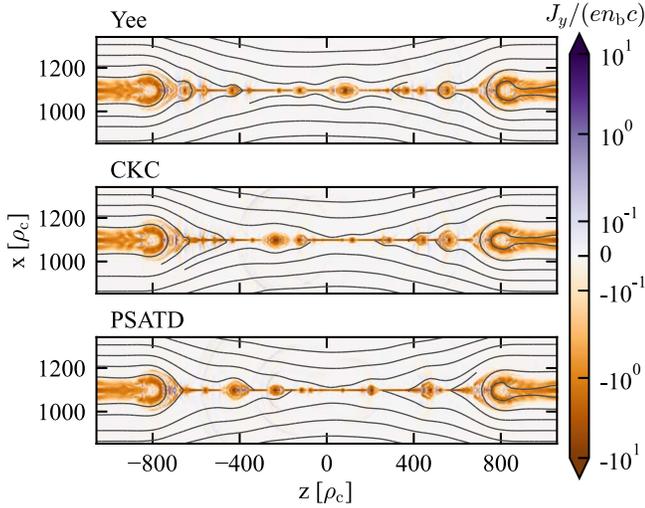


Figure 3. Comparison between current sheet and magnetic field structure in simulations using different Maxwell solvers. All snapshots are shown at $t = 1470 \omega_c^{-1}$. At this phase, all simulations show that the current sheet has fragmented into several small plasmoids, and a single large plasmoid is forming around the edge of the domain. There are small differences between the current and magnetic field structures that are due to the nonlinearity of reconnection. The energy conversion and particle acceleration are similar to one another (see Figures 4 and 5).

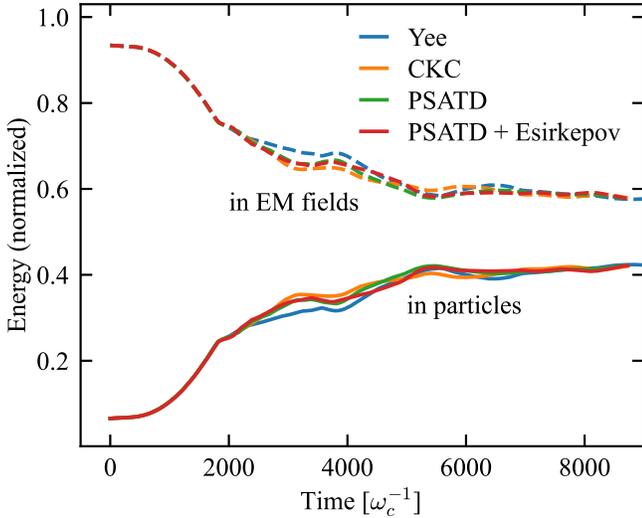


Figure 4. Evolution of energy balance between electromagnetic fields (dashed lines) and particles (solid lines) during reconnection simulations. The y-axis is normalized to the total energy at the start of the simulation. Results from different Maxwell solvers are shown in different colors (Yee: blue, CKC: orange, PSATD+Vay current deposition: green, PSATD+Esirkepov current deposition: red). The initial $\sim 1500 \omega_c^{-1}$ of evolution is nearly identical between all solvers, and after that all simulations show similar evolution, ending when about 40% of the field energy has been converted to particle kinetic energy.

into the reconnection layer, and v_{out} is the terminal exhaust velocity downstream (Cassak et al. 2017; Liu et al. 2022). We calculate v_{in} by averaging $|v_x|$ within a region of size $x_R \times z_R = 245 \rho_c \times 980 \rho_c = 0.11 L_x \times 0.92 L_z$ centered on $(x, z) = (-x_c + x_R, 0)$. This region is marked by the green rectangle in Figure 2. The measured inflow velocity is relatively insensitive to the choice of x_R and z_R . The measurement is also symmetric across the current sheet, i.e., moving the box to $x = -x_c - x_R$ does not change the

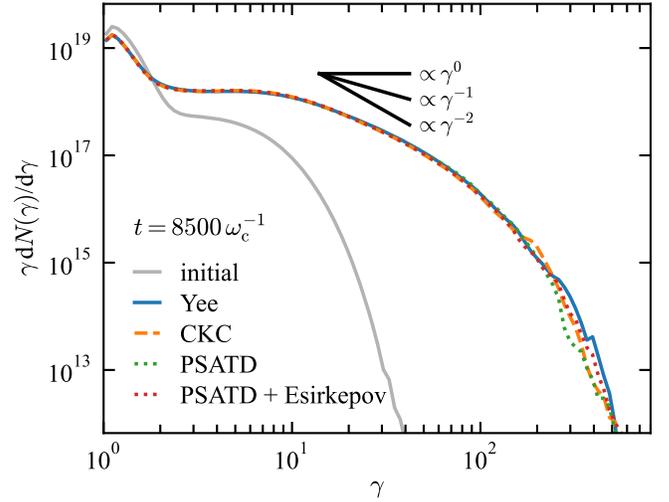


Figure 5. Distribution of particle Lorentz factor γ weighted by particle energy at the start of the simulation (gray line) and once reconnection has ended ($t = 8500 \omega_c^{-1}$). We compare final distributions for the different Maxwell solvers: Yee (solid blue), CKC (dashed orange), PSATD with Vay current deposition (dotted green), and PSATD with Esirkepov current deposition (dotted red). All three of the Maxwell solvers show very similar particle acceleration, as does the PSATD+Esirkepov combination. The majority of the new particle kinetic energy has gone into particles with $\gamma \sim \sigma$. The spectral slope varies between 0 and -2 (overplotted in black), which roughly matches the ranges found in prior work (Guo et al. 2015; Werner et al. 2016, 2018).

measurement. The measurement is also similar on the other current sheet. For the sake of simplicity, we therefore only show reconnection rates on the $+x$ side of the lower current sheet. The v_{in} average only includes cells in the upstream, i.e., where more than 70% of the plasma originated on the same side of the current sheet. Because reconnection mixes plasma across the current sheet, this excludes plasmoids and other reconnection exhaust. The value of this threshold does not strongly affect our results.

We measure the outflow velocity, v_{out} , by taking the median of the 10 highest cell-averaged z -velocities within $\delta = 12 \rho_c$ of the center of the current sheet. By $t = 1000 \omega_c^{-1}$, this approaches the expected value of $v_{\Lambda} = c\sigma/\sqrt{\sigma^2 + 1} \approx c$.

This estimates of the reconnection rate for Yee, PSATD, and CKC simulations are shown in Figure 6. The ratio $v_{\text{in}}/v_{\text{out}}$ grows nearly identically in all simulations from 0 to 0.2 at a time of $1200 \omega_c^{-1}$. For the following $1200 \omega_c^{-1}$, the estimate of the rate holds relatively constant between 0.15 and 0.2 before dropping at around $t = 2500 \omega_c^{-1}$. This coincides with $2L_z/c$, the light-crossing time across the current sheet. After this point, the steady-state assumption under which $\beta \approx v_{\text{in}}/v_{\text{out}}$ breaks down as the reconnection fronts interfere with one another, due to the periodic boundary condition.

The measurement of a rate between 0.15 and 0.2 matches the expectation from prior work. Typically, β for nonrelativistic reconnection is observed to be around 0.1, with higher rates for relativistic reconnection (Blackman & Field 1994; Guo et al. 2015; Comisso & Bhattacharjee 2016; Liu et al. 2017).

3.3. Performance

Current sheet structures (Figure 3), energy conversion (Figure 4), particle acceleration (Figure 5), and reconnection rate (Figure 6) are all similar between the Yee, PSATD, and

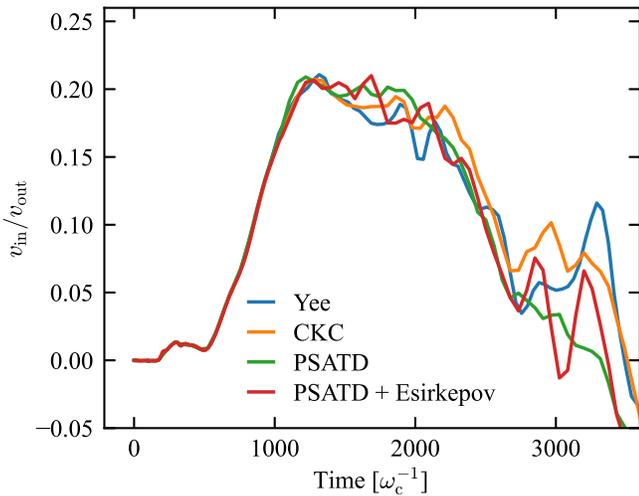


Figure 6. Estimated dimensionless reconnection rate v_{in}/v_{out} for simulations with Yee (blue), CKC (orange), PSATD with Vay current deposition (green), and PSATD+Esirkepov (red). The rate evolves similarly in all cases, peaking at 0.2 and remaining between 0.15 and 0.2 from $t = 1000 \omega_c^{-1}$ to $2000 \omega_c^{-1}$. This matches the expected “universal” reconnection rate of 0.1 (Comisso & Bhattacharjee 2016; Liu et al. 2017).

CKC Maxwell solvers. This suggests that reconnection proceeds similarly in all of our simulations, demonstrating that the results produced by the PSATD and CKC solvers for the reconnection problem are comparable to those from the more conventional Yee solver. The PSATD+Esirkepov simulations also produce comparable results, despite the violation of Gauss’s law. Adding this simulation to our performance comparison allows us to decouple the performance effects of the Maxwell solvers and current deposition schemes.

In these reconnection simulations, the majority of the runtime ($\sim 60\%$) is spent in the current deposition routine, the performance of which is unaffected by the Maxwell solver. Consequently, the walltime per time step is approximately the same between the Yee, CKC, and PSATD+Esirkepov runs. CKC and PSATD permit a time step that is 40% longer than in Yee, reducing the time to solution for CKC and PSATD+Esirkepov by 40% and 30%, respectively, over the baseline Yee (see Table 2). PSATD+Esirkepov’s steps take 10% longer than Yee(+Esirkepov) or CKC(+Esirkepov), indicating that the PSATD field-solve itself has only a small impact on the computational performance. However, PSATD(+Vay) has a 50% longer time per step, due to differences in the current deposition kernels. Thus, despite the longer time step permitted by the PSATD field solver, there is only a slight net increase in the time to solution. The implementation of Esirkepov deposition in WarpX is highly optimized, so it is possible that similar optimization in Vay deposition could make PSATD+Vay a more advantageous combination. This is an area for future work.

4. Results: Large Time Steps with PSATD

A particular advantage of the PSATD method over either CKC or Yee is that it is not subject to a Courant stability criterion. Consequently, we are able to further increase the time step in the PSATD simulations by raising the CFL factor above 1. If other algorithmic choices are kept the same, then the time per time step is unlikely to change, reducing the time to

Table 2
Performance Comparison between Solvers

Algorithmic Options	Time Step (ω_c^{-1})	Walltime per Step (s)	Walltime to $1470 \omega_c^{-1}$ (s)
Yee	0.411	0.077	274.6
CKC	0.581	0.077	193.5
PSATD (+ Vay)	0.581	0.115	290.0
PSATD+Esirkepov	0.581	0.083	209.9

Note. The walltime per time step is similar for Yee, CKC, and PSATD+Esirkepov, but 50% more expensive for PSATD(+Vay). Because PSATD and CKC allow longer time steps, CKC reaches a solution 40% faster, and PSATD(+Vay) is slightly slower than Yee. PSATD+Esirkepov performs comparably to CKC, reflecting that walltime per step is largely governed by the deposition scheme. These numbers are from simulations run to a final time of $t = 1470 \omega_c^{-1}$ with dynamic load balancing and no I/O on 21 OLCF Summit nodes (126 GPUs).

solution. In this section, we study a sequence of simulations that are identical except for their CFL factors and therefore time steps. The CFL factors studied range from a baseline of 0.95–2.2. We refer to simulations in this sequence as CX.X where “X.X” is the CFL factor. C0.95 is the simulation labeled as “PSATD” in Section 3. As discussed in Section 2, we exclusively use the Vay current deposition scheme for the simulations in this sequence.

The statistical properties of the reconnecting plasma tend to remain the same early in the simulations and for lower values of the CFL factor. For values of the CFL factor $\lesssim 1.65$, the energy conversion from magnetic fields to particles proceeds almost identically (Figure 7). We see the same behavior as with the baseline runs (Figure 4), where during the initial $\sim 1500 \omega_c^{-1}$, about 20% of the field energy is transferred to the particles. Following that phase, the energy conversion proceeds more slowly, saturating at a final distribution where about 40% of the energy is in particles and 60% remains in the fields.

C1.7 matches the simulations with a lower CFL factor through the main period of interest, until reconnection saturates at around $7000 \omega_c^{-1}$. In the last $1000 \omega_c^{-1}$, it shows a slight increase in both electromagnetic field and particle energy, indicating that energy is not conserved. This is reflected in the relative energy conservation (see Figure 8). Nonconservation starts off small but increases sharply at around $t = 7000 \omega_c^{-1}$, reaching 1% by the end of the simulation at $t = 9000 \omega_c^{-1}$.

The runs that all appear the same in the energy balance plot (CFL factor ≤ 1.65 , Figure 7) show greater degrees of energy conservation. For runs with a CFL factor $\lesssim 1.6$, nonconservation is comparable to the baseline, C0.95. C1.65 also shows low levels of nonconservation until the very end of the simulation, when it starts to increase. Left to run further, we expect that nonconservation would continue to increase like in the simulations with higher CFL factors. For the duration of the simulation, C1.65 is within a reasonable tolerance, only violating energy conservation by less than 1 part in 500.

For simulations with CFL factors ≥ 1.8 , significant energy nonconservation develops during the last part of reconnection, much earlier than in C1.7. Errors in the energy have built up significantly by $4000 \omega_c^{-1}$ for C1.8, and by $2000 \omega_c^{-1}$ for both C2.0 and C2.2 (Figure 8). These are first apparent in the energy conservation plot, but continue to grow until they show in the plot of energy conversion (orange line in Figure 7). By

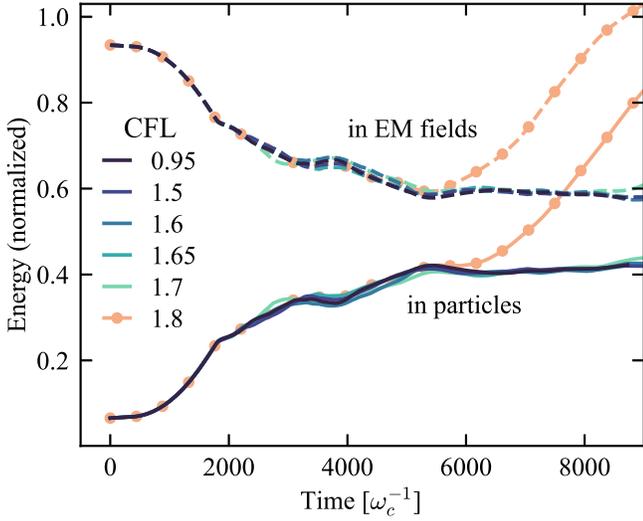


Figure 7. Energy balance between electromagnetic field (dashed lines) and particle kinetic energy (solid lines) in PSATD simulations of magnetic reconnection with a CFL factor greater than 1. CFL factors less than 1.7 produce results that are nearly identical to the benchmark C0.95. For larger values (i.e., C1.8 in orange and C2.0 and C2.2, not shown), we see a qualitative increase in both particle and field energy at progressively earlier times.

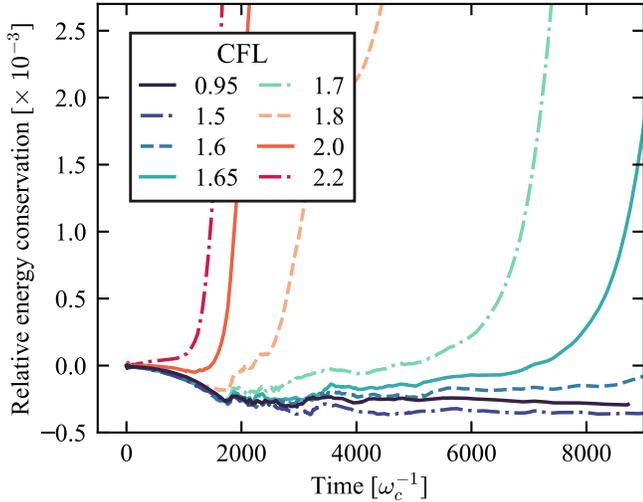


Figure 8. Relative energy conservation throughout reconnection simulations with CFL factors greater than one. The baseline C0.95 is shown in solid, dark purple. Up to C1.6 (dashed blue), energy nonconservation stays relatively small and comparable to C0.95 (less than one part in 2000). This suggests that a CFL factor $\lesssim 1.6$ is sufficient to capture the necessary physics throughout the time interval of interest ($t \lesssim 9000 \omega_c^{-1}$). For most of the evolution, C1.65 (solid light blue) also keeps a low degree of nonconservation, but larger errors in energy appear at around $t = 7000 \omega_c^{-1}$. By the end of the simulation, its nonconservation is still less than 1%, but growing rapidly. As we further increase the CFL factor, the rapid increase in energy errors moves to progressively earlier times, with the runaway occurring before $2000 \omega_c^{-1}$ for C2.2 (red dashed-dotted line).

$t = 6000 \omega_c^{-1}$ the field and particle energies are visibly different from the baseline. C2.0 and C2.2 also show earlier runaway growth in field and particle energy, but they are omitted from Figure 7 for clarity.

The particle energy spectra further demonstrate agreement between the simulations with a CFL factor $\lesssim 1.7$ throughout the majority of the simulation. Figure 9 compares the initial particle energy distributions (gray lines) with those at $t = 2000 \omega_c^{-1}$ (mid-reconnection) and $t = 8500 \omega_c^{-1}$ (after

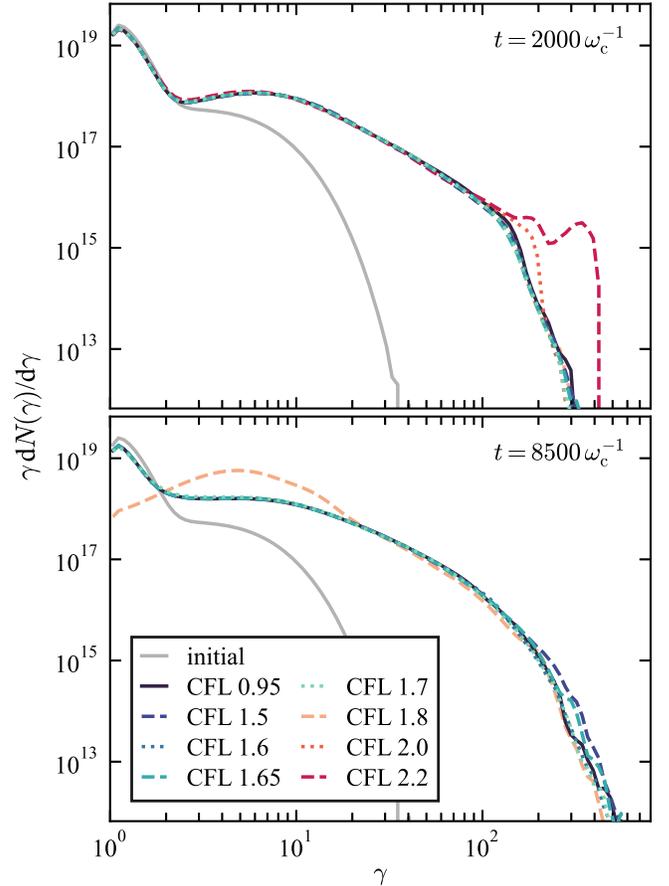


Figure 9. Distribution of particle energy as a function of Lorentz factor (γ) mid-reconnection (top panel; $t = 2000 \omega_c^{-1}$) and near the end of simulations (bottom panel; $t = 8500 \omega_c^{-1}$). The baseline C0.95 appears in solid dark purple. Larger values of the CFL factor are shown in dashed and dotted lines. For the values of the CFL factor ≤ 1.7 , the distributions agree closely throughout the simulations, especially at lower energies ($\lesssim 100$). At high energies, there are small differences, which cannot be distinguished from statistical noise. As with energy conservation (Figure 8), the simulations with the highest CFLs diverge the earliest; for C2.0 and C2.2 (orange dotted and red dashed), we see disagreement with the baseline simulations at $t = 2000 \omega_c^{-1}$, while C1.8 (dashed orange) is still in agreement. By the end of the simulation, C1.8 shows strong disagreement with the baseline, peaking at $\gamma \sim 4$, rather than close to 1.

reconnection, almost at the end of the simulation). At the earlier time, all of the runs we study show close agreement with the baseline C0.95 out to $\gamma \sim 100$. Those runs with a CFL factor less than 2 also agree closely with one another at the highest energies; there are few, if any, particles with Lorentz factors in excess of 300. C2.0 shows a slight overabundance of particles with $\gamma \sim 100$, which grows as the simulation progresses. C2.2 shows an even larger overabundance of high-energy particles, reaching a maximum Lorentz factor of 400, 30% higher than the maximum achieved by C1.8 at this time. This also coincides with the beginning of the energy nonconservation seen in Figure 8. The overabundance of high-energy particles may directly cause the initial nonconservation, but shortly thereafter we also see artificial heating in the cold upstream. In C0.95 and the runs with lower (≤ 1.7) CFL factors, the cold upstream plasma appears as a peak at $\gamma \sim 1$, and it remains largely unchanged even at the end of the simulation. However, in C1.8 and above, this peak broadens and moves out to $\gamma \sim 4$. This is apparent near the end of C1.8 (bottom panel, Figure 9). This also occurs in C2.0, C2.2,

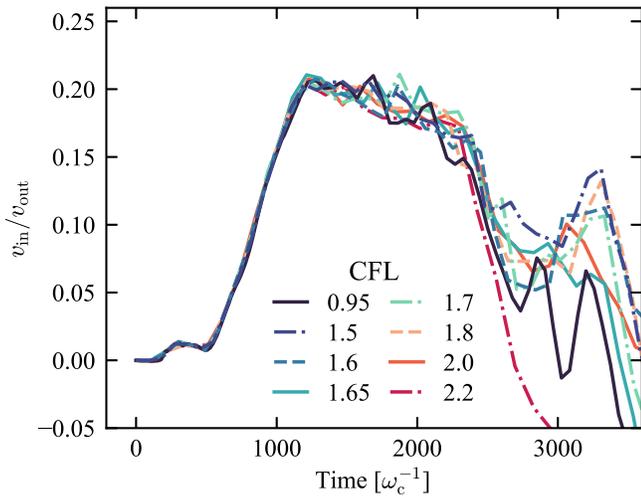


Figure 10. Dimensionless reconnection rate estimated as the ratio of inflow to outflow velocities in the current sheet (see Section 3.2). The rate appears to be mostly insensitive to the CFL factor. This is largely because our estimate of the reconnection rate is only valid for the first $2000 \omega_c^{-1}$ of the simulations, when energy nonconservation is minimal (Figure 8). C2.2 is the only run that shows significant deviation in the first $2500 \omega_c^{-1}$ of evolution, which coincides with a runaway in its energy nonconservation.

beginning shortly after the time of the top panel; we omit their distributions in the later panel for the sake of clarity.

At the end of the simulations, the runs with CFL factors $\lesssim 1.7$ have particle energy distributions that agree almost exactly for $\gamma < 100$. For the highest-energy particles, there are minor differences, comparable to the spread seen in the different electromagnetic solvers (see Figure 5). These highest-energy particles are also the rarest—there are ~ 5000 times fewer particles at $\gamma = 2$ as there are at $\gamma = 100$, so the higher energies are more subject to statistical noise.

The nonconservation seen for CFL factors ≥ 1.65 in Figure 8 appears to be at least partially driven by numerical heating in the upstream region. The exact cause is uncertain, though we expect it is a direct consequence of the large CFL factor, rather than the under-resolution of a physical timescale. Coincidentally, C1.65 has a time step just over $1 \omega_c^{-1}$, meaning that the cyclotron motion of upstream particles cannot be captured. However, this heating does not occur if we obtain that same time step by, for example, doubling the cell size in each direction and reducing the CFL factor below 1. The effect is driven by the upstream plasma. Simulations of a domain with no current sheet, filled only with upstream ($\theta = 0.15$, $\sigma = 30$) plasma, show the same numerical heating.

The dimensionless reconnection rate is remarkably similar for all runs except for C2.2 (Figure 10). As discussed in Section 3.2, we are primarily interested in the value of the estimated reconnection rate for $t < 2000 \omega_c^{-1}$. Within that time interval, the simulations all show a nearly identical rise to 0.2, followed by a slight decline to 0.15 over the following $1000 \omega_c^{-1}$. C2.2 largely follows this pattern, but declines much more quickly than the other runs at $t = 2000 \omega_c^{-1}$. This highlights the robustness of the reconnection rate and indicates that it likely should not be used to diagnose whether the simulations are producing correct results. At $t = 2000 \omega_c^{-1}$, C2.0 and C2.2 both show qualitative disagreement with the baseline models in the high-energy particle spectra (Figure 9).

Table 3

Performance of Simulations with CFL Factors between 0.95 and 2.2

CFL (ω_c^{-1})	Time Step Per Step (s)	Walltime $1470 \omega_c^{-1}$ (s)	Walltime to
0.95	0.582	0.115	290.0
1.5	0.919	0.126	200.9
1.6	0.978	0.127	190.9
1.65	1.010	0.128	186.4
1.7	1.041	0.128	180.9
1.8	1.102	0.129	172.5
2.0	1.225	0.131	157.2
2.2	1.347	0.131	143.0

Note. These numbers are from otherwise-identical simulations run to a final time of $t = 1470 \omega_c^{-1}$ with dynamic load balancing and no I/O on 21 OLCF Summit nodes (126 GPUs).

However, they agree with baseline models in the reconnection rate measurement.

Of the diagnostics discussed here, energy conservation is the most sensitive to numerical problems that arise due to a high CFL factor. Because our simulations are closed systems, the total energy should remain the same throughout the evolution, providing a straightforward ground truth comparison. When increasing the CFL factor, we find that the particle acceleration and reconnection rate obtained from those simulations do not deviate from the baseline until after energy nonconservation begins to rise rapidly. Even C2.2 agrees with C0.95 for the first $1500 \omega_c^{-1}$ of the simulation, before its nonconservation starts to increase. While particle energy spectra may match the baseline even after energy nonconservation begins to run away (e.g., C1.7 at $t = 8500 \omega_c^{-1}$, bottom panel of Figure 9), we would not be able to verify that particle acceleration was still correct in the absence of a known baseline. When increasing the CFL factor, we suggest using a runaway in energy nonconservation as a heuristic for closed systems to determine when a simulation result is unreliable.

In Table 3, we summarize the computational performance of simulations with CFL factors greater than 1. Again, we compare against the baseline C0.95 with the PSATD Maxwell solver. We find that, between C0.95 and C2.2, the walltime to simulate a fixed physical time interval decreases by half. Over this range of CFL factors, the physical time step increases by a factor of 2.3, and the walltime per time step increases by about 15%. Most of this walltime difference is because current deposition is slower (on a per step basis) at higher CFL factors. When particles move farther per step, they are more likely to move between cells, and this likely reduces the cache performance of the deposition routines.

The speedup obtained for a particular application is limited by how high one can increase the CFL factor while maintaining a reliable result throughout the time interval of interest. We expect that the interval over which energy is conserved for a given CFL factor will vary based on the problem setup. For the configuration described here, a CFL factor of 2.2 may be adequate if we were only interested in the onset of reconnection. In that case, the time to solution will be 50% of what it is in C0.95. In this study, we were interested in following reconnection from its onset until it saturates. For that, we needed a CFL factor of at most 1.7, and would have terminated the simulation at around $6500 \omega_c^{-1}$. At that CFL

factor, we reach the solution 1.6 times faster than in the baseline C0.95.

5. Conclusions

We have performed first-of-their-kind particle-in-cell simulations of relativistic reconnection with a 16th-order pseudo-spectral Maxwell solver (PSATD) and a time step that exceeds the conventional CFL limit. We find that PSATD and the nonstandard finite-difference scheme CKC qualitatively and quantitatively produce the same results as the standard second-order finite-difference scheme Yee. We have verified that all three schemes produce the same qualitative plasmoid evolution, particle-field energy balance, particle acceleration, and reconnection rate (Figures 3–6). The particle energy distribution has a power-law tail with a spectral slope α of $dN(\gamma)/d\gamma \propto \gamma^{-\alpha}$. We measure α between 1 and 3, as expected (Guo et al. 2015; Werner et al. 2016, 2018). The reconnection rate is between 0.15 and 0.2 for most of the linear phase of the simulation, consistent with expectations for relativistic reconnection.

We also compare the performance of the solvers and measure efficiency in terms of time to solution. For the same CFL factor, CKC and PSATD allow for a time step that is longer than that of Yee by a factor of $\sqrt{2}$. The walltime taken per step, though, depends more on the current deposition scheme than on the Maxwell solver. The walltime per step is the same in CKC as in Yee (which both use Esirkepov deposition), giving a $\sim 40\%$ speedup in time to solution. High-order PSATD is only charge-conserving when used with Vay deposition, which in its current WarpX implementation is not as optimized as Esirkepov and is therefore more computationally expensive. Thus, we decouple the comparison of the solvers and current deposition methods by performing an additional PSATD+Esirkepov simulation. In doing so, we verify that PSATD itself is not the primary cause of the more expensive time step. Despite not being charge-conserving, PSATD+Esirkepov gives the correct answer for this problem, and has a walltime per step only 10% higher than CKC or Yee. In conjunction with the $\sqrt{2}$ larger time step allowed by PSATD, this produces a net 30% reduction in time to solution. A time step in the charge-conserving PSATD+Vay takes 50% longer than in Yee, so it has a slightly longer time to solution than the most commonly used charge-conserving Yee algorithm.

Unlike either of the finite-difference schemes, PSATD is numerically stable at any time step, even one greater than the light travel time across a cell. We explore the accuracy and performance of CFL factors >1 in the relativistic reconnection problem. We find that the timescale of interest sets the maximum allowable time step parameterized by the CFL factor. Factors ≤ 1.65 conserve energy comparably well to the baseline PSATD case until $t = 8000 \omega_c^{-1}$, well past the end of reconnection. In that interval, they show good agreement in both particle acceleration (Figure 5) and reconnection rate (Figure 10). A slightly higher CFL factor of 1.7 still shows good agreement in energy distribution (Figure 7) and particle acceleration, but it begins to show signs of runaway errors in energy at $t \approx 6500 \omega_c^{-1}$, right after reconnection saturates. This suggests that, for this particular problem and computational setup, the CFL factor could reach 1.7 without compromising the accuracy of the simulation results during reconnection. As we progressively increase the time step (up to a CFL factor of 2.2), runaway energy nonconservation begins earlier and

earlier. In the interval where energy is conserved, the other diagnostics such as particle energization and reconnection rate agree with the baseline case, suggesting that energy non-conservation is a good diagnostic of when the simulation results are reliable for closed systems with periodic boundary conditions.

The walltime per time step only increases modestly with CFL factor, about 15% from the baseline C0.95 to C2.2. The increases in the time step outweigh the increases in walltime per step, reducing the walltime per physical time by a factor of about two between C0.95 and C2.2, if all other algorithmic choices are unchanged. Comparing against CKC+Esirkepov, the most efficient of the CFL = 0.95 simulations, we obtain a 25% reduction in time to solution by adopting a CFL factor of 2.2, and a $<10\%$ reduction by using the CFL factor of 1.7 that we have determined to be suitable for our scientific purposes. If a CFL factor below 1.7 were necessary, CKC+Esirkepov would be the most efficient option, due to the faster current deposition. Future work will include further optimization of the Vay deposition routine, which would shift the trade-offs to favor the PSATD+Vay combination in more situations. While PSATD+Esirkepov was accurate in these simulations, we do not recommend its use without an additional current correction or divergence-cleaning operation to guarantee charge conservation.

One of PSATD’s strengths is that it reduces numerical dispersion that may appear when using Yee or CKC solvers. While an exploration into the mitigating effect of the PSATD solver on numerical dispersion in relativistic plasmas is outside the scope of this paper, it is indeed a known effect that can contaminate simulations of astrophysical jets, shocks, and magnetic reconnection (Godfrey 1974; Melzani et al. 2013; Godfrey & Vay 2014; Ikeya & Matsumoto 2015; Li et al. 2017; Nishikawa et al. 2021; Tomita et al. 2022). PSATD can therefore provide new opportunities to study highly relativistic plasmas, a numerically challenging regime that characterizes a number of astrophysical systems.

The numerical and algorithmic innovations leveraged in this study can be used to enable larger and more efficient simulations of astrophysical and laboratory plasmas (Ji et al. 2022). Our simulations are also some of the first GPU-accelerated astrophysical PIC simulations. As a first step, we have only focused on two-dimensional systems without additional kinetic and radiative physics. A third spatial dimension is dynamically important in reconnection because it makes the current sheet susceptible to an additional instability, called “drift-kink,” which can suppress particle acceleration (e.g., Cerutti et al. 2014; Sironi & Spitkovsky 2014; Guo et al. 2015; Werner & Uzdensky 2017). CKC and PSATD may be especially efficient in 3D because their time steps are larger than Yee’s by a factor of $\sqrt{3}$, rather than $\sqrt{2}$ in 2D. In many astrophysical reconnection environments, synchrotron emission and pair production play an important role. With WarpX, we will be able to take advantage of GPU-accelerated exascale computing resources to perform 3D simulations that include this radiative physics.

Acknowledgments

We thank Luca Comisso for the insightful discussions on magnetic reconnection and reconnection rate. This work was partially supported by the U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research,

Exascale Computing Project (17-SC-20-SC) under contract DE-AC02-05CH11231, and Simulation and Analysis of Reacting Flow, FWP #FP00011940, funded by the Applied Mathematics Program in ASCR. This research used the open-source particle-in-cell code WarpX <https://github.com/ECP-WarpX/WarpX>, primarily funded by the US DOE Exascale Computing Project. The primary WarpX contributors are affiliated with LBNL, LLNL, CEA-LIDYL, SLAC, DESY, CERN, and Modern Electron. We acknowledge all WarpX contributors. This work used resources of the Oak Ridge Leadership Computing Facility at the Oak Ridge National Laboratory, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC05-00OR22725. 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 awards DDR-ERCAP m3881 for 2022 and ASCR-ERCAP mp111 for 2022 and 2023.

Software: WarpX (Fedeli et al. 2022; Vay et al. 2023), AMReX (Zhang et al. 2019; the AMReX Development Team et al. 2023), matplotlib (Hunter 2007), numpy (Harris et al. 2020), scipy (Virtanen et al. 2020), yt (Turk et al. 2011).

Appendix Magnetic Field Configuration

A.1. Equilibrium

In all of our simulations, we fix $\rho_c = 4.1 \times 10^{-3}$ m, which gives $B = 0.42$ T and fixes the other dimensional values in Table 4. These are used in the WarpX simulations, which require SI units. This choice of length scale is largely arbitrary. In this nonradiative regime, the magnetization determines most of the physics; the rest of the results will scale accordingly.

From the relationship between ρ_c and λ_e and the definition of skin depth, we can calculate number density in the current sheet n_d :

$$n_d = \frac{m\epsilon_0 c^2}{e^2 \lambda_e^2} = 2.8 \times 10^{17} \text{ m}^{-3}, \quad (\text{A1})$$

and in turn, the background number density $n_b = 5.6 \times 10^{16} \text{ m}^{-3}$. At this density, a magnetization of 30 requires the background magnetic field $B_0 = 0.42$ T.

After integrating Ampere's Law (Equation (4)),

$$\begin{aligned} B_z(x) = & -4\mu_0 en_b c \beta_0 \delta \left[\arctan\left(\tanh\left(\frac{x+x_c}{2\delta}\right)\right) \right. \\ & \left. - \arctan\left(\tanh\left(\frac{x-x_c}{2\delta}\right)\right) \right] \\ & + \frac{1}{2} \left(\frac{n_d}{n_b} - 1\right) \left(\tanh\left(\frac{x+x_c}{\delta}\right) - \tanh\left(\frac{x-x_c}{\delta}\right) \right) + C. \end{aligned} \quad (\text{A2})$$

To solve for the constant of integration C , we apply the condition from Equation (5):

$$B_z(\infty) = -B_z(0) = B_0. \quad (\text{A3})$$

Table 4

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	8.96 m
Domain half-width (z)	L_z	4.32 m

First, taking the limit as $x \rightarrow \infty$,

$$\begin{aligned} \lim_{x \rightarrow \infty} B_z(x) = & -4\mu_0 en_b \beta_0 c \delta (\arctan(1) - \arctan(1)) \\ & + \frac{1}{2} \left(\frac{n_d}{n_b} - 1\right) (1 - 1) + C \\ = & C, \end{aligned} \quad (\text{A4})$$

because $\lim_{x \rightarrow \infty} \tanh(x) = 1$. To double precision, $\tanh(x) = 1$ if $x \geq 20$, so the following is true if $x_c/\delta \geq 40$, which is the case in our simulations:

$$\begin{aligned} B_z(0) = & -4\mu_0 en_b \beta_0 c \delta (\arctan(1) - \arctan(-1)) \\ & + \frac{1}{2} \left(\frac{n_d}{n_b} - 1\right) (1 - (-1)) + C \end{aligned} \quad (\text{A5})$$

$$= -4\mu_0 en_b \beta_0 c \delta \left(\frac{\pi}{2} + \frac{n_d}{n_b} - 1\right) + C. \quad (\text{A6})$$

Combining the conditions on B_z as $x \rightarrow \infty$ and at $x=0$ (Equation (5)):

$$B_0 = \lim_{x \rightarrow \infty} B_z(x) = C = -B_z(0) \quad (\text{A7})$$

$$C = 4\mu_0 en_b \beta_0 c \delta \left(\frac{\pi}{2} + \frac{n_d}{n_b} - 1\right) - C \quad (\text{A8})$$

$$C = 2\mu_0 en_b \beta_0 c \delta \left(\frac{\pi}{2} + \frac{n_d}{n_b} - 1\right). \quad (\text{A9})$$

This yields the expression for magnetic field in Equations (6) and (7):

$$\begin{aligned} B_z(x) = & -\frac{2B_0}{\left(\frac{\pi}{2} + \frac{n_d}{n_b} - 1\right)} \\ & \times \left(\arctan\left(\tanh\left(\frac{x+x_c}{2\delta}\right)\right) - \arctan\left(\tanh\left(\frac{x-x_c}{2\delta}\right)\right) - \frac{\pi}{4} \right. \\ & \left. + \frac{1}{2} \left(\frac{n_d}{n_b} - 1\right) \left(\tanh\left(\frac{x+x_c}{\delta}\right) - \tanh\left(\frac{x-x_c}{\delta}\right) - 1 \right) \right), \end{aligned} \quad (\text{A10})$$

with

$$B_0 = 2\mu_0 en_b \beta_0 c \delta \left(\frac{\pi}{2} + \frac{n_d}{n_b} - 1 \right). \quad (\text{A11})$$

The full expressions for the equilibrium fields and plasma properties are

$$\begin{aligned} \frac{B(x)}{B_0} &= \frac{2}{\frac{\pi}{2} + \frac{n_d}{n_b} - 1} \\ &\times \left[\arctan \left(\tanh \left(\frac{x - x_c}{2\delta} \right) \right) - \arctan \left(\tanh \left(\frac{x + x_c}{2\delta} \right) \right) + \frac{\pi}{4} \right. \\ &\left. + \frac{1}{2} \left(\frac{n_d}{n_b} - 1 \right) \left(\tanh \left(\frac{x - x_c}{\delta} \right) - \tanh \left(\frac{x + x_c}{\delta} \right) + 1 \right) \right] \end{aligned} \quad (\text{A12})$$

$$\frac{n(x)}{n_b} = 1 + \left(\frac{n_d}{n_b} - 1 \right) \left(\operatorname{sech} \left(\frac{x + x_c}{\delta} \right) + \operatorname{sech} \left(\frac{x - x_c}{\delta} \right) \right) \quad (\text{A13})$$

$$\beta(x) = \frac{\lambda_e \sqrt{\sigma}}{2\delta \left(\frac{\pi}{2} + \frac{n_d}{n_b} - 1 \right)} \sqrt{\frac{n_d}{n_b}} \left(\operatorname{sech} \left(\frac{x + x_c}{\delta} \right) - \operatorname{sech} \left(\frac{x - x_c}{\delta} \right) \right) \quad (\text{A14})$$

$$\theta(x) = \frac{\sigma (4 + \eta) / \eta - [B_z(x) / B_0]^2}{4 [n(x) / n_b] \sqrt{1 - \beta(x)^2}}. \quad (\text{A15})$$

A.2. Magnetic Field Perturbation

The magnetic field perturbation is based on the the vector potential (Equation (20)), which is an integral of $B_z(x)$ (Equation (6)). We split $B(x)$ into three types of terms, based on functional form: constants, arctan(tanh) terms, and tanh terms. Constants are straightforward to integrate. The integrals of the other two types of terms are

$$\int \tanh \left(\frac{x \pm x_c}{\delta} \right) dx \equiv \zeta_{\pm}(x) = \delta \log \cosh \left(\frac{x \pm x_c}{\delta} \right) \quad (\text{A16})$$

$$\begin{aligned} \int \arctan \tanh \left(\frac{x \pm x_c}{2\delta} \right) dx &\equiv \xi_{\pm}(x) \\ &= 2\delta \left(\frac{x \pm x_c}{2\delta} \right) \left[\arctan \left(\exp \left(-\frac{x \pm x_c}{\delta} \right) \right) + \arctan \left(\tanh \left(\frac{x \pm x_c}{2\delta} \right) \right) \right. \\ &\left. + \frac{\delta i}{2} \left[\operatorname{Li}_2 \left(-i \exp \left(-\frac{x \pm x_c}{\delta} \right) \right) - \operatorname{Li}_2 \left(i \exp \left(-\frac{x \pm x_c}{\delta} \right) \right) \right] \right]. \end{aligned} \quad (\text{A17})$$

The value of $\int B_z(x) dx$ is

$$\begin{aligned} \int B_z(x) dx &= \frac{2B_0}{\frac{\pi}{2} + \frac{n_d}{n_b} - 1} \\ &\times \left[\xi_{-}(x) - \xi_{+}(x) + \frac{\pi}{4} x + \frac{1}{2} \left(\frac{n_d}{n_b} - 1 \right) (\zeta_{-}(x) - \zeta_{+}(x) + x) \right] \end{aligned} \quad (\text{A18})$$

$$\begin{aligned} &= \frac{2B_0}{\frac{\pi}{2} + \frac{n_d}{n_b} - 1} \left[\frac{\pi}{4} x + (x - x_c) \left(\arctan \exp \left(-\frac{x - x_c}{\delta} \right) + \arctan \tanh \left(\frac{x - x_c}{2\delta} \right) \right) \right. \\ &- (x + x_c) \left(\arctan \exp \left(-\frac{x + x_c}{\delta} \right) + \arctan \tanh \left(\frac{x + x_c}{2\delta} \right) \right) \\ &+ \frac{\delta i}{2} \left(\operatorname{Li}_2 \left(-i \exp \left(-\frac{x - x_c}{\delta} \right) \right) - \operatorname{Li}_2 \left(i \exp \left(-\frac{x - x_c}{\delta} \right) \right) \right) \\ &- \frac{\delta i}{2} \left(\operatorname{Li}_2 \left(-i \exp \left(-\frac{x + x_c}{\delta} \right) \right) - \operatorname{Li}_2 \left(i \exp \left(-\frac{x + x_c}{\delta} \right) \right) \right) \\ &\left. + \frac{1}{2} \left(\frac{n_d}{n_b} - 1 \right) \left(\delta \log \cosh \left(\frac{x - x_c}{\delta} \right) - \delta \log \cosh \left(\frac{x + x_c}{\delta} \right) + x \right) \right] \end{aligned} \quad (\text{A19})$$

$$\begin{aligned} &= \frac{2B_0}{\frac{\pi}{2} + \frac{n_d}{n_b} - 1} \left[\frac{\pi}{4} (x - 2x_c) + \frac{1}{2} \left(\frac{n_d}{n_b} - 1 \right) \left(\delta \log \cosh \left(\frac{x - x_c}{\delta} \right) - \delta \log \cosh \left(\frac{x + x_c}{\delta} \right) + x \right) \right. \\ &+ \frac{\delta i}{2} \left(\operatorname{Li}_2 \left(-i \exp \left(-\frac{x - x_c}{\delta} \right) \right) - \operatorname{Li}_2 \left(i \exp \left(-\frac{x - x_c}{\delta} \right) \right) \right) \\ &\left. - \frac{\delta i}{2} \left(\operatorname{Li}_2 \left(-i \exp \left(-\frac{x + x_c}{\delta} \right) \right) - \operatorname{Li}_2 \left(i \exp \left(-\frac{x + x_c}{\delta} \right) \right) \right) \right]. \end{aligned} \quad (\text{A20})$$

Equation (A20) results from applying the simplification

$$\arctan(e^{-2y}) + \arctan(\tanh(y)) = \frac{\pi}{4}. \quad (\text{A21})$$

For the purposes of symmetry and preserving the periodic boundary conditions, we have set the constant of integration equal to zero.

The perturbation to the magnetic field is

$$\mathbf{B}' = \nabla \times A_y - \mathbf{B} = -\frac{\partial A_y}{\partial z} \hat{\mathbf{i}} + \frac{\partial A_y}{\partial x} \hat{\mathbf{k}} - \mathbf{B} \quad (\text{A22})$$

$$\begin{aligned} &= -0.01 \times \hat{\mathbf{i}} \left[\frac{51\pi}{L_z} \sin\left(\frac{\pi}{L_z} z\right) \right. \\ &\quad \times \cos^{50}\left(\frac{\pi}{L_z} z\right) \cos^2\left(\frac{\pi(x-x_c)}{L_x}\right) \int B_z(x) dx \left. \right] \\ &\quad + 0.01 \times \hat{\mathbf{k}} \left[\frac{\pi}{L_x} \cos^{51}\left(\frac{\pi}{L_z} z\right) \sin\left(\frac{2\pi(x-x_c)}{L_x}\right) \right. \\ &\quad \times \int B_z(x) dx - \cos^{51}\left(\frac{\pi}{L_z} z\right) \cos^2\left(\frac{\pi(x-x_c)}{L_x}\right) \\ &\quad \times B_z(x) \left. \right]. \end{aligned} \quad (\text{A23})$$

ORCID iDs

Hannah Klion  <https://orcid.org/0000-0003-2095-4293>
 Revathi Jambunathan  <https://orcid.org/0000-0001-9432-2091>
 Michael E. Rowan  <https://orcid.org/0000-0003-2406-1273>
 Eloise Yang  <https://orcid.org/0000-0002-9319-4216>
 Donald Willcox  <https://orcid.org/0000-0003-2300-5165>
 Jean-Luc Vay  <https://orcid.org/0000-0002-0040-799X>
 Remi Lehe  <https://orcid.org/0000-0002-3656-9659>
 Andrew Myers  <https://orcid.org/0000-0001-8427-8330>
 Axel Huebl  <https://orcid.org/0000-0003-1943-7141>
 Weiqun Zhang  <https://orcid.org/0000-0001-8092-1974>

References

- Abdo, A. A., Ackermann, M., Ajello, M., et al. 2011, *Sci*, **331**, 739
 Bird, R., Tan, N., Luedtke, S. V., et al. 2022, *IEEE Trans. Parallel Distrib. Syst.*, **33**, 952
 Birdsall, C. K., & Langdon, A. B. 1991, *Plasma Physics via Computer Simulation* (Boca Raton, FL: CRC Press)
 Blackman, E. G., & Field, G. B. 1994, *PhRvL*, **72**, 494
 Boris, J. P. 1970, in Proc. Fourth Conf. Num. Sim. Plasmas, ed. J. O. Boris & R. Shanny (Washington, DC: Naval Research Laboratory), 3, <https://apps.dtic.mil/sti/citations/ADA023511>
 Bussmann, M., Burau, H., Cowan, T. E., et al. 2013, Proc. Int. Conf. on High Performance Computing, Networking, Storage and Analysis, SC '13 (New York: ACM)
 Cassak, P. A., Liu, Y. H., & Shay, M. A. 2017, *JPhPh*, **83**, 715830501
 Cerutti, B., Uzdensky, D. A., & Begelman, M. C. 2012, *ApJ*, **746**, 148
 Cerutti, B., Werner, G. R., Uzdensky, D. A., & Begelman, M. C. 2013, *ApJ*, **770**, 147
 Cerutti, B., Werner, G. R., Uzdensky, D. A., & Begelman, M. C. 2014, *ApJ*, **782**, 104
 Chien, S. W. D., Nylund, J., Bengtsson, G., et al. 2020, in 2020 IEEE 32nd Int. Symp. Computer Architecture and High Performance Computing (SBAC-PAD) (Los Alamitos, CA: IEEE Computer Society Press), 149
 Cole, J. B. 1997, *ITMTT*, **45**, 991
 Cole, J. B. 2002, *ITAP*, **50**, 1185
 Comisso, L., & Bhattacharjee, A. 2016, *JPhPh*, **82**, 595820601
 Cowan, B. M., Bruhwiler, D. L., Cary, J. R., Cormier-Michel, E., & Geddes, C. G. R. 2013, *PhRvs*, **16**, 041303
 Esirkepov, T. Z. 2001, *CoPhC*, **135**, 144
 Fedeli, L., Huebl, A., Boillod-Cerneux, F., et al. 2022, in SC22: Int. Conf. for High Performance Computing, Networking, Storage and Analysis (Los Alamitos, CA: IEEE Computer Society Press), 1
 Germaschewski, K., Fox, W., Abbott, S., et al. 2016, *JCoPh*, **318**, 305
 Giannios, D. 2010, *MNRAS*, **408**, L46
 Godfrey, B. B. 1974, *JCoPh*, **15**, 504
 Godfrey, B. B., & Vay, J.-L. 2014, *JCoPh*, **267**, 1
 Guo, F., Li, X., Daughton, W., et al. 2019, *ApJL*, **879**, L23
 Guo, F., Liu, Y.-H., Daughton, W., & Li, H. 2015, *ApJ*, **806**, 167
 Haber, I., Lee, R., Klein, H., & Boris, J. 1973, in Proc. Sixth Conf. Num. Sim. Plasmas (Berkeley, CA: Lawrence Berkeley Laboratory), 46
 Hakobyan, H., Petropoulou, M., Spitkovsky, A., & Sironi, L. 2021, *ApJ*, **912**, 48
 Harris, C. R., Millman, K. J., van der Walt, S. J., et al. 2020, *Natur*, **585**, 357
 Harris, E. G. 1962, *NCim*, **23**, 115
 Huebl, A. 2019, *PhD thesis*, Technische Universitaet Dresden
 Hunter, J. D. 2007, *CSE*, **9**, 90
 Ikeya, N., & Matsumoto, Y. 2015, *PASJ*, **67**, 64
 Ji, H., Daughton, W., Jara-Almonte, J., et al. 2022, *NatRP*, **4**, 263
 Kärkkäinen, M., Gjonaj, E., Lau, T., & Weiland, T. 2006, in 9th Int. Computational Accelerator Physics Conf. (Geneva: CERN), 35
 Kumar, P., & Zhang, B. 2015, *PhR*, **561**, 1
 Li, F., Yu, P., Xu, X., et al. 2017, *CoPhC*, **214**, 6
 Liu, Y.-H., Cassak, P., Li, X., et al. 2022, *CmPhy*, **5**, 97
 Liu, Y.-H., Hesse, M., Guo, F., et al. 2017, *PhRvL*, **118**, 085101
 Lyutikov, M., & Uzdensky, D. 2003, *ApJ*, **589**, 893
 McKinney, J. C., & Uzdensky, D. A. 2012, *MNRAS*, **419**, 573
 Melzani, M., Walder, R., Folini, D., Winisdoerffer, C., & Favre, J. M. 2014, *A&A*, **570**, A111
 Melzani, M., Winisdoerffer, C., Walder, R., et al. 2013, *A&A*, **558**, A133
 Myers, A., Almgren, A., Amorim, L., et al. 2021, *ParC*, **108**, 102833
 Nalewajko, K., Uzdensky, D. A., Cerutti, B., Werner, G. R., & Begelman, M. C. 2015, *ApJ*, **815**, 101
 Nalewajko, K., Yuan, Y., & Chruslinska, M. 2018, *JPhPh*, **84**, 755840301
 Nishikawa, K., Duřan, I., Köhn, C., & Mizuno, Y. 2021, *LRCA*, **7**, 1
 Petropoulou, M., Sironi, L., Spitkovsky, A., & Giannios, D. 2019, *ApJ*, **880**, 37
 Philippov, A., Uzdensky, D. A., Spitkovsky, A., & Cerutti, B. 2019, *ApJL*, **876**, L6
 Piran, T. 2004, *RvMP*, **76**, 1143
 Schoeffler, K. M., Grismayer, T., Uzdensky, D., & Silva, L. O. 2023, *MNRAS*, **523**, 3812
 Sironi, L. 2022, *PhRvL*, **128**, 145102
 Sironi, L., & Spitkovsky, A. 2014, *ApJL*, **783**, L21
 Tavani, M., Bulgarelli, A., Vittorini, V., et al. 2011, *Sci*, **331**, 736
 Tetarenko, A. J., Sivakoff, G. R., Miller-Jones, J. C. A., et al. 2017, *MNRAS*, **469**, 3141
 the AMReX Development Team, Almgren, A., Beckner, V., et al. 2023, the AMReX Development Team, AMReX-Codes/amrex: AMReX 23.04, v23.04, Zenodo, doi:10.5281/zenodo.7790123
 Tomita, S., Ohira, Y., Kimura, S. S., Tomida, K., & Toma, K. 2022, *ApJL*, **936**, L9
 Turk, M. J., Smith, B. D., Oishi, J. S., et al. 2011, *ApJS*, **192**, 9
 Vay, J. L., Almgren, A., Amorim, L. D., et al. 2020, *JPhCS*, **1596**, 012059
 Vay, J.-L., Almgren, A., Amorim, L. D., et al. 2023, ECP-WarpX/WarpX: 23.04, v23.04, Zenodo, doi:10.5281/zenodo.7793695
 Vay, J. L., Geddes, C. G. R., Cormier-Michel, E., & Grote, D. P. 2011, *JCoPh*, **230**, 5908
 Vay, J.-L., Haber, I., & Godfrey, B. B. 2013, *JCoPh*, **243**, 260
 Vincenti, H., & Vay, J. L. 2016, *CoPhC*, **200**, 147
 Virtanen, P., Gommers, R., Oliphant, T. E., et al. 2020, *NatMe*, **17**, 261
 Werner, G. R., & Uzdensky, D. A. 2017, *ApJL*, **843**, L27
 Werner, G. R., & Uzdensky, D. A. 2021, *JPhPh*, **87**, 905870613
 Werner, G. R., Uzdensky, D. A., Begelman, M. C., Cerutti, B., & Nalewajko, K. 2018, *MNRAS*, **473**, 4840
 Werner, G. R., Uzdensky, D. A., Cerutti, B., Nalewajko, K., & Begelman, M. C. 2016, *ApJL*, **816**, L8
 Xiong, Q. Y., Huang, S. Y., Yuan, Z. G., et al. 2023, *ApJS*, **264**, 3
 Yee, K. 1966, *ITAP*, **14**, 302
 Zenitani, S. 2015, *PhPI*, **22**, 042116
 Zenitani, S., & Hoshino, M. 2001, *ApJL*, **562**, L63
 Zenitani, S., & Hoshino, M. 2007, *ApJ*, **670**, 702
 Zhang, Q., Guo, F., Daughton, W., Li, H., & Li, X. 2021, *PhRvL*, **127**, 185101
 Zhang, W., Almgren, A., Beckner, V., et al. 2019, *JOSS*, **4**, 1370