SHARP: A Spatially Higher-order, Relativistic Particle-in-cell Code

Mohamad Shalaby^{1,2,3}, Avery E. Broderick^{1,2}, Philip Chang⁴, Christoph Pfrommer^{5,6}, Astrid Lamberts⁷, and Ewald Puchwein⁸

¹Department of Physics and Astronomy, University of Waterloo, 200 University Avenue West, Waterloo, ON, N2L 3G1, Canada; mshalaby@live.ca

²Perimeter Institute for Theoretical Physics, 31 Caroline Street North, Waterloo, ON, N2L 2Y5, Canada

³Department of Physics, Faculty of Science, Cairo University, Giza 12613, Egypt

⁴Department of Physics, University of Wisconsin-Milwaukee, 1900 E. Kenwood Boulevard, Milwaukee, WI 53211, USA

⁵Leibniz-Institut für Astrophysik Potsdam (AIP), An der Sternwarte 16, D-14482 Potsdam, Germany

⁶Heidelberg Institute for Theoretical Studies, Schloss-Wolfsbrunnengasse 35, D-69118 Heidelberg, Germany

⁷Theoretical Astrophysics, California Institute of Technology, Pasadena, CA 91125, USA

⁸Institute of Astronomy and Kavli Institute for Cosmology, University of Cambridge, Madingley Road, Cambridge, CB3 0HA, UK

Received 2017 February 15; revised 2017 April 5; accepted 2017 April 10; published 2017 May 23

Abstract

Numerical heating in particle-in-cell (PIC) codes currently precludes the accurate simulation of cold, relativistic plasma over long periods, severely limiting their applications in astrophysical environments. We present a spatially higher-order accurate relativistic PIC algorithm in one spatial dimension, which conserves charge and momentum exactly. We utilize the smoothness implied by the usage of higher-order interpolation functions to achieve a spatially higher-order accurate algorithm (up to the fifth order). We validate our algorithm against several test problems—thermal stability of stationary plasma, stability of linear plasma waves, and two-stream instability in the relativistic and non-relativistic regimes. Comparing our simulations to exact solutions of the dispersion relations, we demonstrate that SHARP can quantitatively reproduce important kinetic features of the linear regime. Our simulations have a superior ability to control energy non-conservation and avoid numerical heating in comparison to common second-order schemes. We provide a natural definition for convergence of a general PIC algorithm: the complement of physical modes captured by the simulation, i.e., those that lie above the Poisson noise, must grow commensurately with the resolution. This implies that it is necessary to simultaneously increase the number of particles per cell and decrease the cell size. We demonstrate that traditional ways for testing for convergence fail, leading to plateauing of the energy error. This new PIC code enables us to faithfully study the long-term evolution of plasma problems that require absolute control of the energy and momentum conservation.

Key words: instabilities – methods: numerical – plasmas – relativistic processes

1. Introduction

The PIC method is a very powerful numerical tool to study the evolution of plasmas, it is used to model plasmas ranging from laboratory experiments to astrophysical environments. First proposed in one spatial dimension (1D) by Buneman (1959) and Dawson (1962), the general idea of this algorithm is straightforward: it follows the trajectory of particles with N -body methods, while solving Maxwell’s equation on a Eulerian grid. The communication between grid points and particles is achieved through interpolation. The general loop (described in Figure 1) consists of first interpolating particle positions and velocities to a spatial grid to solve for the resulting charge and current density. Maxwell’s equations are then solved on the grid with these source terms to find the self-consistent electromagnetic (electrostatic in 1D) fields. Fields are then interpolated back to the particle positions to calculate the Lorentz force, and hence, acceleration, to forward evolve the particles in time using a so-called pusher. This reduces the number of computational operations from $\sim O(N^2)$ (such as in the case for N -body methods) to $\sim O(N)$, where N is the number of particles in the simulation. This also results in eliminating all wave modes in the electromagnetic fields on scales smaller than the cell size on the grid upon which they are computed.

A major test of the accuracy and fidelity of different PIC schemes is their ability to preserve conserved quantities such as energy, momentum, and charge. Often, this is required to accurately study subdominant, relativistic populations that typically arise in astrophysical contexts. Examples include

nonthermal particle populations accelerated at shocks and reconnection events (Spitkovsky 2008; Lyutikov et al. 2016), propagation of cosmic rays (Riquelme & Spitkovsky 2009), interaction of accretion disks and coronae (Miller & Stone 1999), and TeV blazar driven beam instabilities (Broderick et al. 2012). In the latter, this problem is especially severe, with the beams being both numerically and energetically subdominant while being highly relativistic; even a small degree of heating in the background can impact or overwhelm the evolution of beam plasma instabilities.

Direct interpolation of particle data to construct charge and current densities on the grid, in general, leads to a violation of charge conservation. However, modern algorithms typically use charge conserving methods to perform such step while maintaining the charge conservation (e.g., Eastwood 1991; Villasenor & Buneman 1992; Esirkepov 2001; Umeda et al. 2003). Energy and momentum conservation on the other hand appear to be mutually exclusive (see, e.g., Brackbill 2016). Due to their importance, several schemes have been developed to ameliorate their non-conservation using different underlying methodologies.

Recently introduced implicit methods (Chen et al. 2011; Lapenta & Markidis 2011; Markidis & Lapenta 2011) can in theory preserve total energy exactly (though in practice may not), while violating momentum conservation. For computationally simpler explicit, momentum conserving schemes, energy conservation is improved by filtering the deposited grid moments (charge and current densities), as is done in the TRISTAN-MP code (Buneman et al. 1993; Spitkovsky 2005). However, filtering,

Figure 1. Schematic representation for the general loop in the PIC method: starting from the three o'clock position and moving clockwise, the macro-particles' position and velocity data (x, v) are deposited onto a physical grid to construct charge and current densities $(\rho_k$ and $J_k)$ at control points (k) of the grid (Section 2.3). These are used to solve Maxwell's equations, which yield the self-consistently computed (electric and magnetic) fields at these control points of the physical grid (Section 2.4). The updated fields are, then, back-interpolated on the macro-particles to construct the Lorentz force on macro-particles (Section 2.5), which is then used to evolve them via a particle pusher (Section 2.6).

when used with a momentum conserving scheme, leads to a violation of momentum conservation and non-vanishing self-forces (see Appendix A.3). In lieu of filtering, energy conservation is also improved by using higher-order interpolation functions, which is the approach that we will adopt below for both forward- (from particles to fields) and back-interpolation (from fields to particles) steps. This has the added advantage that momentum conservation is maintained.

In this paper, we describe an implementation of the PIC algorithm in 1D that uses high-order spline functions (up to the fifth order) for the forward- and back-interpolation steps of the algorithm. We couple this with an exact Poisson solver and a second-order symplectic integrator, i.e., leap frog, to produce a second-order accurate code called SHARP-1D. SHARP-1D displays superior energy-conservation properties while conserving the momentum exactly. The smoothness coming from the usage of high-order interpolation functions is utilized to construct an up to fifth-order spatially accurate algorithm.⁹ Despite the high spatial order accuracy, our code remains second-order accurate because it is limited by the accuracy of the particle pusher. This will be addressed in future work.

The paper is organized as follows. In Section 2, we describe the basic equations of the PIC method, our choice of discretization of the equations, and discuss sources for the numerical error. After discussing the order of accuracy of the solution, we describe our choice of code units and the implementation of SHARP-1D. In Section 3, we discuss the conservation properties of different PIC algorithms. In Section 4, we demonstrate the different capabilities of SHARP-1D by

validating it against several test problems: thermal stability of plasma, plasma oscillation frequency and linear-Landau damping of standing plasma waves, and two-stream instabilities in both relativistic and non-relativistic regimes. We compare several of the results of SHARP-1D to the results of TRISTAN-MP simulations in Section 5. Finally, we study the convergence properties for our algorithm in Section 6 and discuss the performance of SHARP-1D in Section 7. We conclude in Section 8.

2. The PIC Method

The evolution of the particles that comprise a plasma is described by the Boltzmann and Maxwell's equations. In the absence of collisions, the particles are described by the Vlasov equation, which in one spatial dimension is

$$\partial_t f_s(x, u, t) + \frac{u}{\gamma} \partial_x f_s(x, u, t) + \frac{q_s E(x, t)}{m_s} \partial_u f_s(x, u, t) = 0, \quad (1)$$

where s denotes a particle species, characterized by its charge q_s and mass m_s , $u = \gamma v$ is the spatial component of the four-velocity, $\gamma = \sqrt{1 + (u/c)^2}$ is the Lorentz factor, $E(x, t)$ is the electric field, and $f_s(x, u, t)$ is the phase-space distribution functions of particles of species s .

In one dimension, Maxwell's equations imply that the magnetic field is constant and along the direction of the particle motion. Thus, it impacts the evolution of neither the particle nor the electric field, and we take it to be zero henceforth. Therefore, Maxwell's equations reduce to

$$\partial_x E(x, t) - \frac{\rho(x, t)}{\epsilon_0} = 0 \quad \text{and} \quad \partial_t E(x, t) + \frac{j(x, t)}{\epsilon_0} = 0, \quad (2)$$

where ρ and j are the charge and current densities, respectively.¹⁰ This set of equations is closed by the following equations for ρ and j .

$$\begin{aligned} \rho(x, t) &= \sum_s q_s \int f_s(x, u, t) du \\ j(x, t) &= \sum_s q_s \int \frac{u}{\gamma} f_s(x, u, t) du. \end{aligned} \quad (3)$$

2.1. Smoothing Phase-space Distribution Function

The distribution function for point-like particles is given by the Klimontovich distribution function:

$$f_s^K(x, u) = \sum_{i_s}^{N_s^p} \delta(x - x_{i_s}) \delta(u - u_{i_s}), \quad (4)$$

where N_s^p is the number of physical (point-like) particles and $\delta(x)$ is the Dirac delta function.

Direct usage of this singular distribution function is impractical for two reasons. First, the number of particles to be simulated is too large to be tractable, and thus, second, this distribution function would result in an overwhelming shot

⁹ Due to the usage of lower order interpolation functions, current algorithms perform the back-interpolation step with second-order spatial accuracy (e.g., Lapenta & Markidis 2011; Haugbølle et al. 2013; Brackbill 2016; Lapenta 2016).

¹⁰ Here, we use SI units, and ϵ_0 is the permittivity of the vacuum. To convert to CGS units, ϵ_0 needs to be replaced by $1/4\pi$. Ultimately, we convert to a system of numerical units, obviating the distinction between these (Section 2.8).

noise due to the finite number of particles used in practice (Lipatov 2002). Both can be mitigated in simulations by using a smoothed approximation for the distribution function. Thus, we approximate the distribution function by

$$f_s(x, u) = w \sum_{i_s}^{N_s} S(x - x_{i_s}) \delta(u - u_{i_s}) \approx \int dx' f_s^K(x', u) S(x', x), \quad (5)$$

where N_s is the number of macro-particles (defined below), and $w = N_s^p / N_s$ is the number of physical particles that a macro-particle represents.

The macro-particles (also called computational particles¹¹) have a “shape” $S(x, x_{i_s})$, i.e., a smoothed localized charge distribution that has w physical particles centered at x_{i_s} , with

$$\int S(x, x_{i_s}) dx = 1. \quad (6)$$

The charge and the mass for these computational particles are $Q_s = wq_s$ and $M_s = wm_s$, respectively.

The macro-particles have the same plasma frequency ω_p as the physical particles they represent in simulations: they have same charge-to-mass ratio and hence $\omega_p^2 = q_s^2 N_s^p / (Vm_s \epsilon_0) = Q_s^2 N_s / (VM_s \epsilon_0)$, where V is the volume. Both macro-particles and physical particles also have the same normalized temperature $\theta_s = k_B T_s / M_s c^2 = k_B T_s^p / m_s c^2$, where T_s and T_s^p are the effective temperatures of the macro-particles and physical particles, respectively. This arises from the assumption that the macro-particles are monolithic, and thus have the same velocity distribution as the underlying physical particles, and thus

$$\frac{k_B T_s}{M_s c^2} = \frac{\int du (\gamma - 1) f_s}{\int du f_s} = \frac{\int du (\gamma - 1) f_s^K}{\int du f_s^K} = \frac{k_B T_s^p}{m_s c^2}. \quad (7)$$

Inserting the phase-space distribution function in Equation (5) into the Vlasov–Maxwell system (Equations (1) and (2)), we obtain from the first two moments the following equations of motion for the macro-particle of species s :

$$\frac{dx_{i_s}}{dt} = \frac{u_{i_s}}{\gamma_{i_s}}, \quad \frac{du_{i_s}}{dt} = \frac{Q_s}{M_s} E_{i_s}, \quad (8)$$

$$E_{i_s} \equiv \int E(x) S(x, x_{i_s}) dx, \quad (9)$$

where $E(x)$ is the solution of Maxwell’s Equations (2) with moments given by

$$\rho(x, t) = \sum_s Q_s \sum_{i_s}^{N_s} S(x, x_{i_s}), \quad (10)$$

$$j(x, t) = \sum_s Q_s \sum_{i_s}^{N_s} v_{i_s} S(x, x_{i_s}), \quad (11)$$

where $v_{i_s} = u_{i_s} / \gamma_{i_s}$. Below we explain how implicit discretization of such a system of equations is achieved in our code.

¹¹ Both macro-particles and particles will be interchangeably used to mean the same thing: particles used in simulations. When we refer to physical particles, it will be explicitly specified.

2.2. Spatial Grid

For a system of macro-particles in a periodic box (line) of length L , we divide our domain into N_c cells each of size $\Delta x = L / N_c$. Assuming $k \in \{0, 1, \dots, N_c - 1\}$, we define the k th cell (c_k) centered at $x_{k+1/2} \equiv x_k + \Delta x / 2$ as $x \in [x_k, x_{k+1})$, where $x_k \equiv \Delta x k$.

We adopt spline functions extending over a number of grid cells as the shape function of the macro-particles. Therefore, the distribution of physical particles inside these macro-particles is symmetric around their center and extends over a number of computational cells depending on the order m of the spline functions used. For instance, $m = 1$ is a top-hat distribution (shape) given by

$$S(x, x_i) \rightarrow S^1\left(\frac{|x - x_i|}{\Delta x}\right) = \frac{1}{\Delta x} \begin{cases} 1, & \text{If } |x - x_i| < \Delta x / 2, \\ 0, & \text{otherwise.} \end{cases} \quad (12)$$

We also define, $t^n = n \times \Delta t$, $E^n \equiv E(t^n)$, i.e., superscript n denotes the n th time step for the particular quantity. We choose the time step Δt such that $c\Delta t \leq \Delta x$, to obey the Courant–Friedrichs–Lewy (CFL) stability condition in 1D (Courant et al. 1967).

2.3. Charge and Current Deposition

To obtain a discrete set of equations that governs the evolution of such macro-particles, we begin by integrating the first equation in (2) over the k th cell c_k :

$$E_{k+1}^n - E_k^n = \int_{c_k} \frac{\rho^n}{\epsilon_0} dx = \frac{\Delta x}{\epsilon_0} \rho_{k+1/2}^n. \quad (13)$$

Here, $\rho_{k+1/2}^n$ is the average charge density inside c_k at $t = t^n$:

$$\begin{aligned} \rho_{k+1/2}^n &= \int_{c_k} \rho^n \frac{dx}{\Delta x} = \sum_s Q_s \sum_{i_s} \int_{x_k}^{x_{k+1}} S^m\left(\frac{x - x_{i_s}^n}{\Delta x}\right) \frac{dx}{\Delta x} \\ &= \sum_s \frac{Q_s}{\Delta x} \sum_{i_s} W^m\left(\frac{x_{k+1/2} - x_{i_s}^n}{\Delta x}\right), \end{aligned} \quad (14)$$

where

$$W^m\left(\frac{x_k - x_{i_s}^n}{\Delta x}\right) = \int_{x_k - \Delta x / 2}^{x_k + \Delta x / 2} S^m\left(\frac{x - x_{i_s}^n}{\Delta x}\right) dx \quad (15)$$

defines the weighed contribution of a macro-particle at x_{i_s} to the average charge density of the k th cell.

The explicit forms for both shape S^m and weight W^m functions that we use in our code are given in Appendix B. In Figure 2, we plot different weight functions.

2.4. Solving Maxwell’s Equations

Equation (13) gives the change in the electric field at cell edges exactly. However, a complete solution also requires the boundary condition, $E_{k=0}$. Therefore, we rewrite Equation (13) as

$$E_k^n = E_0^n + \frac{\Delta x}{\epsilon_0} \sum_{j=0}^{k-1} \rho_{j+1/2}^n. \quad (16)$$

We see that E_k^n inherits the error of E_0^n . To find E_0^n , we first find the sum of cell-edges field $E_{\text{tot}}^n \equiv \sum_k E_k^n$. The second equation

Figure 2. Different weight functions implemented in the code. The relation between shape and weight functions is defined in Equation (15). Explicit forms for both shape and weight functions are given in Appendix B.

in (2) can be re-written (for $m > 0$) as

$$\begin{aligned} \partial_t E_{\text{tot}} &= - \sum_k \frac{j_k}{\epsilon_0} = - \frac{1}{\epsilon_0} \sum_s Q_s \sum_{i_s} v_{i_s} \sum_k S^m(x_k, x_{i_s}) \\ &= - \sum_s \frac{Q_s}{\Delta x \epsilon_0} \sum_{i_s} v_{i_s} = \frac{-j_{\text{tot}}}{\epsilon_0}. \end{aligned} \quad (17)$$

The first equality uses $\sum_k \Delta x S^m(x_k, x_i) = \sum_k W^{m-1}(x_k, x_i) = 1$, which is a property of spline functions implemented in our code (see Table 4). Note that j_{tot} can be calculated exactly at each time-step using the macro-particles' velocities. If the plasma macro-particles have a total current, then $\partial_t E \neq 0$, i.e., setting $\partial_t E$ to zero will artificially add a constant electric field on the grid or equivalently a counter current.

Equation (17) is approximated to second, third, and fourth orders of accuracy, respectively, as

$$E_{\text{tot}}^n = E_{\text{tot}}^{n-1} - \frac{\Delta t}{\epsilon_0} j_{\text{tot}}^{n-\frac{1}{2}} + O(\Delta t^3) \quad (18)$$

$$\begin{aligned} E_{\text{tot}}^n &= \frac{21}{23} E_{\text{tot}}^{n-1} + \frac{3}{23} E_{\text{tot}}^{n-2} - \frac{1}{23} E_{\text{tot}}^{n-3} \\ &\quad - \frac{24}{23} \frac{\Delta t}{\epsilon_0} j_{\text{tot}}^{n-\frac{1}{2}} + O(\Delta t^4) \end{aligned} \quad (19)$$

$$\begin{aligned} E_{\text{tot}}^n &= \frac{17}{22} E_{\text{tot}}^{n-1} + \frac{9}{22} E_{\text{tot}}^{n-2} - \frac{5}{22} E_{\text{tot}}^{n-3} \\ &\quad + \frac{1}{22} E_{\text{tot}}^{n-4} - \frac{12}{11} \frac{\Delta t}{\epsilon_0} j_{\text{tot}}^{n-\frac{1}{2}} + O(\Delta t^5). \end{aligned} \quad (20)$$

Proceeding in a similar way to generate higher-order accurate, asymmetric estimates for E_{tot}^n results in numerically unstable approximations. Thus, we stop at the fourth-order accurate method given in Equation (20).

To find E_0 from E_{tot} , we multiply Equation (13) by the index k , and then sum over all cells. The left-hand side is given by

$$\begin{aligned} \sum_{k=0}^{N_c-1} k (E_{k+1}^n - E_k^n) &= \sum_{k=1}^{N_c} (k-1) E_k^n - \sum_{k=1}^{N_c-1} k E_k^n \\ &= N_c E_{N_c}^n - E_{\text{tot}}^n. \end{aligned} \quad (21)$$

Then, using the periodicity of E_k , i.e., $E_{N_c} = E_0$, we can write

$$E_0^n = \frac{E_{\text{tot}}^n}{N_c} + \frac{\Delta x}{\epsilon_0 N_c} \sum_{k=0}^{N_c-1} k \rho_{k+1/2}^n. \quad (22)$$

Thus, for a given set of particle data $\{x_{i_s}^n, v_{i_s}^{n-1/2}\}$, we are able to find the electric field at the edges of the cells, E_k^n . Importantly, the only source of error is the error introduced in finding E_{tot}^n .

The error, for fixed $\Delta t/\Delta x$ (motivated by the CFL condition), is of the order of $O(\Delta x^4)$ if we use Equation (18) to update E_{tot}^n . The error drops to $O(\Delta x^5)$ or $O(\Delta x^6)$ with Equations (19) or (20), respectively. It is also important here to note that an $O(\Delta x^k)$ error in the force would introduce an $O(\Delta x^{k+1})$ error in the updated data of the macro-particles.

2.5. Back-interpolation: Force on Macro-particles

Having determined the electric field using Equations (13), (14), and (18) or (19) or (20) to different orders of accuracy from the macro-particle positions and velocities, we now calculate the force on the individual macro-particles. The effective electric field that acts on the macro-particle (Equation (9)) can be determined from the electric field at each cell face, E_k^n as follows

$$\begin{aligned} E_{i_s}^n &= \sum_k \int_{c_k} E^n(x) S^m(x, x_{i_s}^n) dx \\ &= \sum_k \frac{E_{k+1}^n + E_k^n}{2} \int_{x_k}^{x_{k+1}} S^m(x, x_{i_s}^n) dx + O(\Delta x^2) \\ &= \sum_k \frac{E_{k+1}^n + E_k^n}{2} W^m[(x_{k+1/2} - x_{i_s}^n)/\Delta x] + O(\Delta x^2). \end{aligned} \quad (23)$$

In Appendix A.2, we show that the approximation in Equation (23) leads to exact momentum conservation and vanishing self-forces. Another possible approximation of $E_{i_s}^n$, that has the same order of accuracy is given by

$$\begin{aligned} E_{i_s}^n &= \sum_k E_k^n \int_{x_{k-1/2}}^{x_{k+1/2}} S^m(x, x_{i_s}^n) dx + O(\Delta x^2) \\ &= \sum_k E_k^n W^m[(x_k - x_{i_s}^n)/\Delta x] + O(\Delta x^2). \end{aligned} \quad (24)$$

The approximation in Equation (24) is used, for instance, in Haugbølle et al. (2013). It generally leads to a violation of momentum conservation and unphysical self-forces as shown in Appendix A.1. The order of accuracy for the back-interpolation step is typically second order (e.g., Lapenta & Markidis 2011; Haugbølle et al. 2013; Brackbill 2016; Lapenta 2016) because of the use of lower order interpolation functions.

Using higher-order interpolation functions implies a smoother representation of the phase-space distribution function. As a consequence, we can assume a smoother representation of the electric fields.¹² Therefore, since we implemented up to fifth-order spline interpolation function, this allows constructing up to *fifth-order* accurate back-interpolation. To derive a higher-order accurate method, we approximate the integration

¹² An order m interpolation function means that f is $m+1$ times spatially differentiable.

$I_k = \int_{c_k} E^n(x) S(x, x_{i_s}^n) dx$ using Simpson's rule:

$$\begin{aligned} I_k &= \int_{x_k}^{x_{k+1}} E^n(x) S^m(x, x_{i_s}^n) dx \\ &= \frac{\Delta x}{6} \left[E_k^n S_{k,i_s}^m + 4E_{k+\frac{1}{2}}^n S_{k+\frac{1}{2},i_s}^m + E_{k+1}^n S_{k+1,i_s}^m \right] \\ &\quad + O(\Delta x^5), \end{aligned} \quad (25)$$

where we define $S_{k,i_s}^m \equiv S^m(x_k, x_{i_s}^n)$.

To utilize Equation (25), we need to approximate $E_{k+\frac{1}{2}}^n$ in terms of cell-edge fields E_k^n as follows

$$\begin{aligned} E_{k+\frac{1}{2}} &= \frac{-E_{k+2} + 9(E_k + E_{k+1}) - E_{k-1}}{16} + O(\Delta x^4) \quad (26) \\ &= \frac{3(E_{k+3} + E_{k-2}) + 150(E_{k+1} + E_k) - 25(E_{k+2} + E_{k-1})}{256} \\ &\quad + O(\Delta x^6). \end{aligned} \quad (27)$$

Using Equation (25) and since $\Delta x S^m \in [0, 1]$ (see Table 4), $O(\Delta x^4)$ error in $E_{k+\frac{1}{2}}$ implies $O(\Delta x^4)$ error order in E_{i_s} and $O(\Delta x^6)$ error in $E_{k+\frac{1}{2}}$ implies $O(\Delta x^5)$ error order in E_{i_s} (because of the error order in Equation (25)). Therefore, the approximate electric field $E_{i_s}^n = \sum_k I_k$ on a macro-particle at x_{i_s} , using the periodicity of E_k , can be expressed as

$$\begin{aligned} E_{i_s}^n &= \sum_k \Delta x S_{k+\frac{1}{2},i_s}^m \left[\frac{9(E_k^n + E_{k+1}^n) - (E_{k+2}^n + E_{k-1}^n)}{24} \right] \\ &\quad + \sum_k \frac{E_k^n}{3} [\Delta x S_{k,i_s}^m] + O(\Delta x^4) \quad (28) \\ &= \sum_k \frac{E_k^n}{3} \left[\Delta x S_{k+\frac{1}{2},i_s}^m \right] + \sum_k \Delta x S_{k+\frac{1}{2},i_s}^m \\ &\quad \times \left[\frac{3(E_{k+3}^n + E_{k-2}^n) + 150(E_{k+1}^n + E_k^n) - 25(E_{k+2}^n + E_{k-1}^n)}{384} \right] \\ &\quad + O(\Delta x^5). \end{aligned} \quad (29)$$

To summarize, here we showed how we can find the forces on the individual macro-particles for a given discretized field on a grid. The force error can be of the order of $O(\Delta x^2)$, $O(\Delta x^4)$, or $O(\Delta x^5)$ if we use Equations (23), (28), or (29), respectively. Given the numerical error in finding E_k^n , so far we have shown how, for a given set of particle data $\{x_{i_s}^n, v_{i_s}^{n-1/2}\}$, we can find the forces on such particles $F_{i_s}^n$. The error can be of the order of $O(\Delta x^2)$, $O(\Delta x^4)$, or $O(\Delta x^5)$, by employing consecutively higher-order equations in finding E_0^n .

2.6. Pusher: Particle Update

To push the individual particles, we use a leapfrog scheme to discretize the equations of motions for the particles,

$$u_{i_s}^{n+1/2} = u_{i_s}^{n-1/2} + \Delta t \frac{Q_s}{M_s} E_{i_s}^n + O(\Delta t^3), \quad (30)$$

$$x_{i_s}^{n+1} = x_{i_s}^n + \Delta t v_{i_s}^{n+1/2} + O(\Delta t^3). \quad (31)$$

The code assumes that the initial positions of the particles are provided at $t = 0$, but the initial momenta are given at $t = -\Delta t/2$. It also assumes that at $t = -\Delta t$ the sum of electric field at cell edges was zero, i.e., $E_{\text{tot}}^{n=-1} = 0$.

2.7. Error Sources in Our Algorithm

Here, we summarize the different sources of error in the algorithm presented above.

1. The use of discretized equations to compute E_{tot} , yielding up to $O(\Delta t^5)$ accurate schemes (Equations (18)–(20)).
2. The interpolation of the field from the grid to the macro-particle to calculate the force F_{i_s} . This can be done at $O(\Delta x)$ in such a fashion that total momentum is conserved exactly, or at higher accuracy, $O(\Delta x^3)$ and $O(\Delta x^4)$ at the cost of (slightly) violating momentum conservation.
3. The updating of the particle positions, which is currently performed at $O(\Delta t^3)$.

For fixed $\Delta t/\Delta x$, as implied by the CFL condition, the above imply that our code is fundamentally second-order accurate, limited by the particle pusher. That is, despite improving energy and momentum conservation, the order of the interpolation function does not set the order of accuracy of the overall scheme. Implementing a higher-order symplectic integrator within our scheme would improve the convergence order; we leave this point for future work. Nevertheless, as we will show in Sections 4–6, the improved spatial order produces substantial practical enhancements in the code performance.

2.8. Normalized Equations

Using the fiducial units (c, n_0, q_0, m_0) , we define the following scales

$$\begin{aligned} t_0 &= \sqrt{m_0 \epsilon_0 / (q_0^2 n_0)}, \quad \mathbb{E}_0 = \sqrt{n_0 m_0 c^2 / \epsilon_0}, \\ \rho_0 &= q_0 n_0, \quad j_0 = \rho_0 c, \quad x_0 = c t_0. \end{aligned} \quad (32)$$

We then define our dimensionless variables as $\bar{t} = t/t_0$, $\bar{dt} = \Delta t/t_0$, $\bar{x} = x/x_0$, $\bar{h} = \Delta x/x_0$, $\bar{u} = u/c$, $\bar{Q}_s = Q_s/q_0$, $\bar{M}_s = M_s/m_0$, $\bar{\rho} = \rho/\rho_0$, $\bar{j} = j/j_0$, and $\bar{E} = E/\mathbb{E}_0$. We also identify t_0 as the timescale of the plasma frequency of the entire plasma that includes contributions from all species. This defines n_0 as follows

$$\omega_0^2 \equiv t_0^{-2} = \frac{q_0^2 n_0}{m_0 \epsilon_0} = \omega_p^2 = \sum_s \frac{Q_s^2 n_s}{M_s \epsilon_0} \rightarrow n_0 = \sum_s \frac{\bar{Q}_s^2 n_s}{\bar{M}_s}, \quad (33)$$

where n_s is the number density of the macro-particles of species s and

$$n_0 \Delta x = \sum_s \frac{\bar{Q}_s^2 n_s \Delta x}{\bar{M}_s} = \frac{1}{N_c} \sum_s \frac{\bar{Q}_s^2 N_s}{\bar{M}_s}. \quad (34)$$

In the case where all species have the same mass and charge, $n_0 \Delta x = (\sum_s N_s)/N_c = N_t/N_c$. In such a case, we chose our fiducial units so that $\bar{Q}_s^2 = 1$, $\bar{M}_s = 1$, therefore, $n_0 = \sum_s n_s$.

In terms of the dimensionless variables defined above, the equations that are solved by the code can be written as follows

$$\bar{\rho}_{k+1/2}^n = \frac{N_c}{\sum_s (\bar{Q}_s^2 N_s) / \bar{M}_s} \sum_s \bar{Q}_s \sum_{i_s} W^m[(\bar{x}_{k+1/2} - \bar{x}_{i_s}^n)/h], \quad (35)$$

$$\bar{E}_{k+1}^n = \bar{E}_k^n + h \bar{\rho}_{k+1/2}^n, \quad (36)$$

$$\bar{E}_{i_s}^n = \sum_k \frac{\bar{E}_{k+1}^n + \bar{E}_k^n}{2} W^m[(\bar{x}_{k+1/2} - \bar{x}_{i_s}^n)/h], \quad (37)$$

$$\bar{u}_{i_s}^{n+1/2} = \bar{u}_{i_s}^{n-1/2} + dt \frac{\bar{Q}_s}{\bar{M}_s} \bar{E}_{i_s}^n, \quad (38)$$

$$\bar{x}_{i_s}^{n+1} = \bar{x}_{i_s}^n + dt \bar{v}_{i_s}^{n+1/2}.$$

The general loop in the code is then

$$\left\{ \begin{array}{l} \bar{x}_{i_s}^n \\ \bar{v}_{i_s}^{n-1/2} \end{array} \right\} \xrightarrow{(35)} \bar{\rho}_{k+1/2}^n \xrightarrow{(36)} \bar{E}_{\text{tot}}^{n-1} \xrightarrow{(37)} \bar{E}_k^n \xrightarrow{(38)} \left\{ \begin{array}{l} \bar{x}_{i_s}^{n+1} \\ \bar{v}_{i_s}^{n+1/2} \end{array} \right\}.$$

A schematic representation of this loop is shown in Figure 1.

2.9. Implementation

SHARP-1D is implemented in C++, and is massively parallelized using MPI. The parallelization is done by distributing macro-particles on different processors, while reserving the first processor (rank = 0) to compute the electric field on the grid and to manage outputs.

3. Conserved Quantities in PIC

As we mentioned in the introduction, a major test of the accuracy and fidelity of PIC schemes is their ability to preserve conserved quantities such as energy, momentum, and charge. Below we discuss such conservation laws in PIC schemes and how well they are respected when different methodologies are used. We begin by showing that our algorithm is charge conserving. Then, we compare the energy and momentum conservation properties when implicit and explicit techniques are employed.¹³ We also show that the usage of higher-order interpolation leads to a decrease in the aliasing, which improves energy conservation while maintaining exact momentum conservation.

3.1. Charge Conservation

In traditional implementations of PIC, the direct interpolation from particles to grid points, in order to calculate the grid charge and current densities, leads to violation of the continuity equation on the grid. The reason for that is calculating the current density requires the knowledge of both particles positions and velocities at the half-time step and approximating the particles positions at half-time step leads to an error in the calculated current density when particles cross cell-boundaries.

Recently, several methods were proposed where the calculation of the current density on the grid from the particles is done such that the continuity equation is satisfied on the grid at all times (Eastwood 1991; Villasenor & Buneman 1992; Esirkepov 2001; Umeda et al. 2003). In the presented algorithm, we locally obey the discretized continuity equation at all times (i.e., we use a charge conserving scheme). The discretized current density¹⁴ coincides with the current density proposed in Esirkepov (2001).

¹³ Typically discretization in PIC is done using either explicit or implicit schemes. In explicit schemes (such as the algorithm presented here), particle data are used first to calculate the fields on the grid and then particle advancement in time is carried out using these fields. On the other hand, when implicit discretization is employed, the equations for fields on the grid and evolution equations of particles have to be solved simultaneously in order to evolve forward in time.

¹⁴ In fact, our presented algorithm does not require to calculate the current density for solving Poisson's equation.

To see this, we integrate the continuity equation $\partial_t \rho(x, t) + \partial_x J_x(x, t) = 0$ over a cell of size Δx ,

$$\partial_t \rho_{k+1/2}(t) + \frac{J_{k+1} - J_k}{\Delta x} = 0. \quad (39)$$

Therefore, we can write

$$\frac{\rho_{k+1/2}^{n+1} - \rho_{k+1/2}^n}{\Delta t} + \frac{J_{k+1}^{n+1/2} - J_k^{n+1/2}}{\Delta x} + O(\Delta t^3) = 0. \quad (40)$$

Using the second equation in (2), the current density at cell edges can be expressed as follows

$$J_k^{n+1/2} \equiv \sum_s Q_s \sum_{i_s} v_{i_s}^{n+1/2} S(x_k, x_{i_s}^{n+1/2}) \quad (41)$$

$$= -\frac{\epsilon_0}{\Delta t} [E_k^{n+1} - E_k^n] + O(\Delta t^3). \quad (42)$$

Then using Equations (42) and (13), we can write

$$\begin{aligned} J_{k+1}^{n+1/2} - J_k^{n+1/2} &= -\frac{\epsilon_0}{\Delta t} [(E_{k+1}^{n+1} - E_k^{n+1}) \\ &\quad - (E_{k+1}^n - E_k^n)] + O(\Delta t^3). \\ &= -\frac{\Delta x}{\Delta t} [\rho_{k+1/2}^{n+1} - \rho_{k+1/2}^n] + O(\Delta t^3). \end{aligned} \quad (43)$$

Therefore, our scheme obeys exactly the second-order accurate continuity equation.

3.2. Energy and Momentum Conservation

In general, for PIC schemes, energy and momentum conservation appear to be mutually exclusive (see, e.g., Brackbill 2016). Momentum non-conservation typically comes from non-vanishing self-forces and errors in the interaction forces—an example of such a case is shown in Appendix A.1. These non-physical forces can produce macroscopic non-physical instabilities (Langdon 1973). Energy non-conservation can also produce dramatic changes in the evolution of the plasmas. Since such an error has a secular (unbounded) growth, energy non-conservation imposes a serious limitation on the ability to study the nonlinear phenomena, which occurs on long timescales (compared to ω_p^{-1}).

For instance, Lapenta & Markidis (2011) demonstrated that in the two-stream instability, the errors in the energy are disproportionately distributed to the fast particles. In particular, they demonstrated that even though the per-particle violation in energy conservation is small, the disproportionate distribution of energy non-conservation leads to errors in the distribution of particles, especially at the high energy end. This is important for particle acceleration in relativistic situations (Lapenta & Markidis 2011). Also, results that probe the long-term behavior of particle distribution functions starting from linearly unstable conditions such as tenuous beam instabilities or particle acceleration are subject to these issues (e.g., Sironi & Giannios 2014; Ardaneh et al. 2015; Park et al. 2015).

Traditional explicit algorithms lead to numerical increase in the total energy (numerical heating) while conserving the total momentum exactly (Birdsall & Langdon 1991; Hockney & Eastwood 1988). For explicit schemes, an energy conserving algorithm was developed in Lewis (1970). However, the total energy is conserved only in the limit of $\Delta t \rightarrow 0$. In practice, there will be numerical heating of the plasma because of the finite timestep.

On the other hand, traditional implicit algorithms tend to decrease the total energy numerically (numerical cooling), while violating the total momentum conservation (Brackbill & Forslund 1985). Recently, implicit algorithms that, in principle, conserve the total energy exactly, while still violating momentum conservation, were introduced: for non-relativistic/classical plasmas by Markidis & Lapenta (2011) and Chen et al. (2011) and relativistic plasmas by Lapenta & Markidis (2011). However, in practice, these algorithms use the Jacobian-Free Newton Krylov (JFNK) method to solve the full implicit system, introducing an error that depends on the accuracy of the Newton or Picard iteration. This leads to violation in energy conservation that can be controlled by increasing the accuracy of such methods.

One major source of energy non-conservation is the coupling between the wave modes resolved by the grid and their aliases. Aliases are wave modes that differ by an integer number of $2\pi/\Delta x$. The reason for such coupling is that continuous particle data (which support wave modes that include aliases of wave modes resolved by the grid) are used in the construction of phase-space moments at a discrete set control of points on the physical grid. Therefore, for momentum conserving schemes, we improve the energy conservation by decreasing the coupling of the wave modes resolved by the grid with their aliases.

One way to decrease the effect of aliasing is done by filtering the deposited grid moments (charge and current densities), this is used, for instance, in TRISTAN (Buneman et al. 1993) and its parallel version TRISTAN-MP (Spitkovsky 2005). Filtering, however, when used with a momentum conserving scheme results in violating the momentum conservation and non-vanishing self-forces (an example for such case is presented in Appendix A.3).

Alternatively, energy conservation is naturally improved when using higher-order interpolation functions, as employed here. In Fourier space, using higher-order interpolation function is qualitatively equivalent to low-pass filtering. However, for instance, filtering when used with W^0 produces larger energy errors in comparison to the errors produced when first-order interpolation, W^1 , is used (cf. Section 8.7 of Birdsall & Langdon 1991). That is, higher-order shape functions considerably decrease the impact of aliasing, resulting in an improved energy conservation for the same reasons as filtering, while simultaneously maintaining higher accuracy evolution of the underlying system. This includes, potentially, exact momentum conservation. As shown explicitly in Appendix C, the discretization of the plasma into macro-particles with shape functions of order m produces a spectral smearing with a width that scales as $(k\Delta x)^{-m}$, producing a corresponding an exponential decrease in aliasing for the resolved modes with shape-function order.

4. Validation

We begin the assessment of the numerical algorithm presented in Section 2 with the physical validation of simulation results against known results. This represents the first in two critical numerical tests, the second being convergence, which is treated in Section 6.

Specifically, we present comparisons with the analytical or semi-analytical results on the following test problems:

1. thermal stability of a uniform plasma (Section 4.1),

2. standing plasma waves—Plasma oscillations and linear Landau damping (Section 4.2), and
3. two-stream instabilities—non-relativistic and relativistic (Section 4.3).

Because validation essentially consists of quantitative comparisons with known results, we are currently limited to phenomena in the linear regime. This appears to be a widespread difficulty within plasma simulations. Nevertheless, weak validation in the nonlinear regime can be found when we discuss code comparisons in Section 5. Here, we assume that all plasma species have the same mass, and, up to a sign, the same charge. Thus, we take as our fiducial units, $q_0^2 = Q_s^2$ and $m_0 = M_s$, which implies that $n_0 = \sum_s n_s$.

4.1. Thermal Stability of Plasma

In the absence of microscopic radiative processes (e.g., Compton scattering, bremsstrahlung), a uniform, thermal plasma should not evolve. In practice, even a uniform thermal plasma will numerically heat. This arises mainly as a consequence of aliasing as explained in Section 3.2.

There are two temperature scales that are often relevant for the numerical heating of plasma simulated with variants of the PIC algorithm. The first is the numerical Debye temperature,

$$\theta_D \equiv h^2, \quad (44)$$

which is the temperature at which the Debye length is equal to square of the cell size in code units. Note that this is a purely numerical quantity that defines those temperatures below which the discretization of Maxwell's equations no longer resolves the Debye length. Typically, for $\theta < \theta_D$ second-order accurate codes will exhibit virulent numerical heating until $\theta \approx \theta_D$ (see, e.g., Birdsall & Maron 1980). Thus, θ_D presents a key numerical limitation on the classes of plasmas that have been simulated to date. However, as shown in Section 3.2, implementing higher-order spatial interpolation decreases the effect of aliasing and hence considerably decrease such heating as shown in Figure 3.

The second is the Poisson temperature, set by the Poisson fluctuations in the reconstruction of the particle distribution. This is set by equating the average potential and kinetic energies of randomly distributed particles (Appendix D.3), though in the cold and hot limits (i.e., non-relativistic and relativistic velocity dispersion limits, respectively) this reduces to

$$\theta_P = \theta_D \frac{N_c^2}{12N} \left[1 - \frac{6f_m}{N_c} \right] \times \begin{cases} 2, & \theta_P \ll 1, \\ 1, & \theta_P \gg 1. \end{cases} \quad (45)$$

Temperatures below θ_P are not well defined numerically. Note that the ordering of θ_P and θ_D is not fixed, though we will consider cases when $\theta_P < \theta_D$ exclusively.

Here, we study the heating due to different approximations in PIC algorithms. To this end, we start all of our simulations with temperatures higher than θ_P and study the evolution of the plasma temperatures. We perform a series of simulations for two populations of negatively and positively charged macro-particles in a periodic box of length $\bar{L} = 5$ and cell size of $h = 0.1$. The total number of macro-particles $N_t = 2 \times 10^5$. Therefore, $\theta_P \approx 1.04 \times 10^{-5}$ and $\theta_D = 10^{-2}$.

Each simulation is started at a different initial temperature: $\theta = 10^{-1}, 10^{-2}, 10^{-3}, 10^{-4}$. In all simulations presented here,

Figure 3. Impact of numerical heating on the temperature (top) and energy error (bottom) evolution in simulations of a stationary thermal plasma for different interpolation functions: first-order interpolation W^1 (left), third-order W^3 (middle), and fifth-order W^5 (right). Here, $\theta = k_B T / m_0 c^2$ is the normalized temperature, $\Delta \mathcal{E}$ is the energy change (error) in the total energy and $\mathcal{E}_{\text{th}}$ is the initial thermal energy of plasma. Therefore, $\Delta \mathcal{E} / \mathcal{E}_{\text{th}}$ measures the fractional energy error with respect to the initial thermal energy of plasma. For each interpolation order we perform simulations, at fixed cell size $h=0.1$, with initial temperatures of $\theta = 10^{-1}$ (solid-red curves), $\theta = 10^{-2}$ (dotted-green curves), $\theta = 10^{-3}$ (dashed-cyan curves), and $\theta = 10^{-4}$ (dotted-dashed blue curves). The top panels show the long-term (up to $4 \times 10^6 \omega_p^{-1}$) evolution of different temperatures for different interpolation orders, while the bottom panels show the evolution of the fractional energy error of the plasma. The dashed black line in the top panel shows the Debye temperature θ_D . The purple lines in the top panels indicate θ_p that is given by Equation (45): temperatures below θ_p are not well defined numerically.

we use second-order accurate back-interpolation, i.e., Equation (23).

Since for all simulations here, $\theta \ll 1$, we start with an initial distribution function that is given by

$$f(x, v, t = 0) = \frac{N_t}{L} \sqrt{\frac{\theta}{2\pi c^2}} e^{-(v/c)^2/2\theta}. \quad (46)$$

For the various initial temperatures, we show in Figure 3 the evolution of the plasma temperatures (top panel) and that of the fractional energy error of the plasma (bottom panel). We show results when first-order interpolation W^1 (left), third-order W^3 (middle), and fifth-order W^5 (right) are used in deposition and back-interpolation steps.

As expected, for the first-order interpolation (the scheme that is most commonly employed in existing PIC codes), uncontrolled heating is observed for all $\theta \leq \theta_D$. This heating subsides when $\theta \approx 3\theta_D$, requiring between one and two cells per Debye length. Hence, first order PIC algorithms face severe computational requirements to resolve cold plasmas.

However, using higher-order spatial interpolation significantly reduces the temperature at which uncontrolled numerical heating occurs. By fifth order, temperatures four orders of magnitude smaller than θ_D , and only an order of magnitude larger than θ_p , can be resolved for millions of plasma timescales. That is, high-order spatial interpolation extends the range of temperatures and timescales that can be simulated.

The marked improvement of SHARP-1D is a direct result of the corresponding improvement in energy conservation. The bottom panels of Figure 3 show the evolution of the growth in the energy of each simulation; in all cases, the unphysical heating can be fully attributed to the failure to conserve energy.

However, the fractional energy non-conservation is improved by nearly three orders of magnitude for each simulated temperature as the spatial interpolation order is increased from W^1 to W^5 . The net result of higher spatial order is, therefore, the ability to run simulations orders of magnitude longer with orders of magnitude lower resolutions.

4.2. Stability of Standing Linear Plasma Waves

We now turn to the stability and evolution of plasmas with linear perturbations, specifically, standing waves. Key validation tests are the reproduction of oscillation frequencies, dispersion relations, and linear Landau damping rates. We begin with a discussion of the anticipated values followed by quantitative comparisons of standing wave evolution.

4.2.1. Linear Dispersion Relations and Growth Rates

As shown in Appendix E, the linear dispersion relation for thermal plasmas are conveniently expressed in terms of a handful of dimensionless quantities:

$$\hat{\omega} = \frac{\omega}{\omega_p}, \quad (47)$$

$$\hat{k} = \frac{kv_{\text{th}}}{\omega_p} = \frac{k}{k_D}, \quad \text{and} \quad (48)$$

$$v_p = \frac{\hat{\omega}}{\hat{k}} = \frac{\omega}{kv_{\text{th}}}, \quad (49)$$

where $v_{\text{th}} = \theta^{1/2}c$ is the thermal velocity dispersion and k_D is the wavenumber associated with the Debye length. The linear

dispersion relation for a non-relativistic thermal population of uniformly distributed electrons with a fixed positively charged background, i.e., infinitely massive ions, is given by (see Appendix E for more details)

$$\hat{k}^2 + 1 = \sqrt{\frac{\pi}{2}} [\text{Erfi}(v_p/\sqrt{2}) - i] v_p e^{-\frac{v_p^2}{2}}, \quad (50)$$

where Erfi is the complex error function defined as $\text{Erfi}(v_p) = -i \text{Erf}(iv_p)$. For a given $\hat{k}$, the roots $\hat{\omega}_j = \omega_j/\omega_p$ can be found by solving Equation (50) numerically; the real and imaginary values for ω_j yield the oscillating frequency and the growing/damping rate of the mode with wavenumber k , respectively, at a given thermal velocity v_{th} .

Approximate expressions for the roots of Equation (50) are often obtained for the limit $\hat{k} \ll 1$. The most common of these, and the standard expression found in most textbooks (e.g., Boyd & Sanderson 2003, and referred to here as “Standard”) is

$$\hat{\omega}_i = -\frac{1}{2} \sqrt{\frac{\pi}{2}} \frac{1}{\hat{k}^3} e^{-\frac{1}{2\hat{k}^2} - \frac{3}{2}} \quad \text{and} \quad \hat{\omega}_r = 1 + \frac{3}{2} \hat{k}^2, \quad (51)$$

where ω_r and ω_i are the real and imaginary components of ω , respectively. A more accurate expression, derived by a higher-order approximation to Equation (50) is given by McKinstrie et al. (1999; referred to as “Extended”).

$$\begin{aligned} \hat{\omega}_i &= -\frac{1}{2} \sqrt{\frac{\pi}{2}} \left(\frac{1}{\hat{k}^3} - 6\hat{k} \right) e^{-\frac{1}{2\hat{k}^2} - \frac{3}{2} - 3\hat{k}^2 - 12\hat{k}^4} \\ \hat{\omega}_r &= 1 + \frac{3}{2} \hat{k}^2 + \frac{15}{8} \hat{k}^4 + \frac{147}{16} \hat{k}^6. \end{aligned} \quad (52)$$

Both of these approximates are shown in comparison to the full numerical solution (“Numerical”) in Figure 4. It is immediately evident that the regime of applicability of both the Standard and Extended approximations is limited to $\hat{k} < 0.25$, with the numerical consequence that neither are quantitatively accurate in the rapid damping regime, where numerical validation experiments can most easily be performed.

For convenience, we provide below a numerical fitting formula based on the formulation of McKinstrie et al. (1999) to the full numerical solution for $\hat{k} \in [0, 0.6]$ that is good to 4% throughout and better than 0.5% above $\hat{k} = 0.3$

$$\begin{aligned} \hat{\omega}_i &= -\frac{1}{2} \sqrt{\frac{\pi}{2}} \left(\frac{1}{\hat{k}^3} - 6\hat{k} - 40.7173\hat{k}^3 \right. \\ &\quad \left. - 3900.23\hat{k}^5 - 2462.25\hat{k}^7 - 274.99\hat{k}^9 \right) \\ &\quad \times \exp \left(-\frac{1}{2\hat{k}^2} - \frac{3}{2} - 3\hat{k}^2 - 12\hat{k}^4 - 575.516\hat{k}^6 \right. \\ &\quad \left. + 3790.16\hat{k}^8 - 8827.54\hat{k}^{10} + 7266.87\hat{k}^{12} \right), \\ \hat{\omega}_r &= 1 + \frac{3}{2} \hat{k}^2 + \frac{15}{8} \hat{k}^4 + \frac{147}{16} \hat{k}^6 \\ &\quad + 736.437\hat{k}^8 - 14729.3\hat{k}^{10} + 105429\hat{k}^{12} \\ &\quad - 370151\hat{k}^{14} + 645538\hat{k}^{16} - 448190\hat{k}^{18}. \end{aligned} \quad (53)$$

In the following subsections, we report a series of simulations to test the code in the regimes $\omega_i \sim 0$ (undamped modes) and $\omega_i \neq 0$ (damped modes). We use a fixed neutralizing

Figure 4. Stability of standing plasma waves in the linear regime. The black curves show the numerical solution to the linear dispersion relation, Equation (50). Red and blue curves are different analytical approximate solutions for the linear dispersion relation (50). Purple data points show results from different simulations in Section 4.2.

background and negatively charged plasma macro-particles whose initial distribution function is given by

$$f(x, v, t = 0) = \frac{e^{-v^2/2\theta}}{\sqrt{2\pi\theta}} [1 + \alpha \cos(kx)]. \quad (54)$$

In all simulations, we use $\theta = 10^{-3}$, $\alpha = 10^{-2}$, and fifth-order spatial interpolation, i.e., W^5 . The damping rates and oscillation frequencies in different simulations of this section are also shown in Figure 4.

4.2.2. Plasma Oscillations

At small $\hat{k}$, the linear Landau damping rate is vanishingly small, i.e., $\omega_i \approx 0$. As a result, a linear perturbation should oscillate providing quantitative tests in the form of the oscillation frequency and evolution of the mode amplitude. We initialize the simulation with an excited mode such that $\hat{k} = 0.01$. The theoretical predictions (numerical solution of Equation (50)) are $\omega_r/\omega_p = 1.00015$ and $\omega_i = -4.7 \times 10^{-2167}$. The rest of our simulation parameters are given in Table 1.

Figure 5 shows the fitted values for the electric field of the initially excited oscillation modes. The oscillation frequency of the initially excited mode is found to be within 0.09% of the theoretically predicted oscillation frequency (this is measured by fitting the oscillation of such a mode in the simulation over about 159 oscillation periods). The oscillation frequency and amplitude are still in excellent agreement with the theoretical prediction until the end of the simulation time $t = 10^3 \omega_p^{-1}$.

In Figure 6, we show the evolution of the averaged energy in the excited modes (over five plasma periods). We find that most

Table 1
Standing Plasma Wave Simulation Parameters and Results

$\hat{k}$	N_t	$\bar{\lambda}^a$	$\bar{\lambda}/h^b$	$\bar{L}/\bar{\lambda}$	ω_r/ω_p^c	$\omega_r^{\text{sim}}/\omega_p^d$	R_r^e	ω_i/ω_p^c	$\omega_i^{\text{sim}}/\omega_p^d$	R_i^e	$\omega_p t_{\min}^f$	$\omega_p t_{\max}^f$
0.01	10^9	19.869	1761	2	1.00015	$0.99925 \pm 3 \times 10^{-7}$	0.9991	...	...	...	0	10^3
0.35	5×10^8	0.568	67	12	1.22095	$1.21989 \pm 2 \times 10^{-5}$	0.9991	-0.03432	$-0.0344 \pm 4 \times 10^{-5}$	1.0036	1.7	30.84
0.40	5×10^8	0.497	62	15	1.28506	$1.28145 \pm 4 \times 10^{-5}$	0.9972	-0.06613	$-0.0687 \pm 8 \times 10^{-5}$	1.0389	1.7	16.80
0.45	5×10^8	0.442	68	19	1.35025	$1.34617 \pm 1 \times 10^{-5}$	0.9969	-0.10629	$-0.1066 \pm 2 \times 10^{-5}$	1.00326	1.6	16.65
0.50	5×10^8	0.397	68	19	1.41566	$1.40786 \pm 2 \times 10^{-5}$	0.9945	-0.15336	$-0.1546 \pm 6 \times 10^{-5}$	1.0081	1.6	8.77

Notes.

^a Wavelength of the initially excited mode in code units.

^b Number of cells used to resolve the initially excited wavelength.

^c Numerical solution of Equation (50).

^d Oscillation frequencies and damping rates found by fitting the changes in simulations. The error in the fitting parameters corresponds to a 99% confidence level.

^e $R_r = \omega_r^{\text{sim}}/\omega_r$ and $R_i = \omega_i^{\text{sim}}/\omega_i$, where ω_r^{sim} and ω_i^{sim} are the oscillation frequency and damping rate obtained from fitting simulation results, respectively. An example for such a fit is shown in Figure 5 for the case of $\hat{k} = 0.01$ and in Figure 7 for the case of $\hat{k} = 0.45$.

^f To fit our simulation results, we specify a time range $[t_{\min}, t_{\max}]$ in the simulation over which we carry out the fitting.

of the power stays in the excited mode for the whole simulation period. Coupled with the degree of energy conservation during the simulation, this implies that no more than 0.1% of the initial energy in the mode leaks into other degrees of freedom of the plasma (e.g., other modes or heating of the plasma).

4.2.3. Linear Landau Damping Rates

For larger $\hat{k}$, corresponding to comparatively larger wave-numbers, the damping rates become large. By $\hat{k} = 0.35$, the wave should damp by one e -fold in about six wave oscillation periods. Again, this provides a number of quantitative tests of SHARP-1D: oscillation frequencies and damping rates. Thus, here we report on simulations at large $\hat{k}$; the values of parameters in these simulations and their results are summarized in Table 1. The damping rates and oscillation frequencies in different simulations are also shown in Figure 4.

In all simulations, the oscillation frequencies and damping rates are found to be within 0.5% and 0.8% of the theoretical predictions of linear theory, respectively. The two columns R_r and R_i of Table 1 report the ratio between the oscillation frequency and damping rate in simulations to their theoretically predicted values, obtained by numerically solving Equation (50). The damping rate and the oscillation frequency of the simulation are obtained by fitting the evolution of the Fourier component of the electric field that corresponds to the initially excited wave mode. In Figure 7, we show an example (for $\hat{k} = 0.45$ simulation) of the fitting carried out to find oscillation frequencies and damping rates of different simulations.

A sudden drop in the mode energy at the first period can be observed in Figure 7. This is a characteristic feature of linear Landau simulations, i.e., when damping rate is comparable to the plasma frequency. Such a feature is also present when other simulation methods are used to simulate the evolution of such modes in a thermal plasma (e.g., Besse & Sonnendrücker 2003; Rossmannith & Seal 2011).

4.3. Two-stream Instability

We now consider the quantitative accuracy with which SHARP-1D can reproduce a dynamical instability—in 1D the primary example is the two-stream instability. This provides an

opportunity to also assess the relativistic performance of the code, through the simulation of relativistic beams. Thus, we will consider two limiting regimes: non-relativistic ($v_b \ll c$) and relativistic ($u_b \gg c$), where v_b is the speed of the streams and $u_b = v_b/\sqrt{1 - (v_b/c)^2}$ is the specific momentum.

As with linear Landau damping, we will begin with a general discussion of the anticipated instability properties and then move onto quantitative comparisons.

4.3.1. Instability Growth Rates

In the non-relativistic regime, i.e., the streams are moving with non-relativistic speeds and have non-relativistic temperatures, the linear dispersion relation for two counter streams in such a limit is given by (see Appendix E for more details)

$$\begin{aligned} \hat{k}^2 + 1 = & \sqrt{\frac{\pi}{8}} \left[(v_p + z_b) \left(\text{Erfi} \left[\frac{v_p + z_b}{\sqrt{2}} \right] - i \right) e^{-v_p z_b} \right. \\ & \left. + (v_p - z_b) \left(\text{Erfi} \left[\frac{v_p - z_b}{\sqrt{2}} \right] - i \right) e^{v_p z_b} \right] e^{-(v_p^2 + z_b^2)/2}, \end{aligned} \quad (55)$$

where $z_b = \bar{v}_b/\sqrt{\theta}$, $\bar{v}_b = v_b/c$, $\hat{k} = kc\sqrt{\theta}/\omega_p$, and $v_p = \omega/kc\sqrt{\theta}$. In general, this must be numerically solved to obtain the mode frequencies. The solution generally consists of two oscillatory modes, a growing mode, and several damping modes. We present some of these solutions in Table 2. In Figure 8, we show the numerical solutions for the maximum growth rates (and the mode growing with such a rate) as a function of the stream speed and its temperature.

In the relativistic limit, the beam velocity distribution exhibits a narrow peak very close to the speed of light c , and is thus well described by the cold-plasma limit, i.e., $\theta = 0$. Within this limit, the linear dispersion relation for two relativistic counter streaming $e^+ - e^-$ populations (with speed v_b) is given by

$$1 = \frac{\omega_p^2/2}{\gamma_b^3(kv_b - \omega)^2} + \frac{\omega_p^2/2}{\gamma_b^3(kv_b + \omega)^2}. \quad (56)$$

Figure 5. Fitting the oscillation frequency for plasma oscillation simulations. The left panel shows that the simulation data (red curve) is in excellent agreement with the fit. It continues to excellently fit the simulation until its end at $10^3 \omega_p^{-1}$ (right).

Figure 6. Evolution of the averaged (over five plasma periods) energy in the initially excited mode normalized to averaged energy in the first five plasma periods of such mode. The inset shows that, after evolving the simulation to $10^3 \omega_p^{-1}$, the level of variation on energy carried by the mode is about 0.1% of the initial energy in that mode.

In this case, it is possible to obtain analytic solutions:

$$\frac{\omega}{\omega_p} = \pm \sqrt{\frac{1 + 2\hat{k}_b^2 \gamma_b^3 \pm \sqrt{8\hat{k}_b^2 \gamma_b^3 + 1}}{2\gamma_b^3}}, \quad (57)$$

where $\hat{k}_b = kv_b/\omega_p$. When $\gamma_b^3 \hat{k}_b^2 < 1$, these again correspond to two oscillating modes, a growing mode and a damping mode. The positive imaginary root is maximized at the wavenumber

$$\frac{k_m c}{\omega_p} = \sqrt{\frac{3/8}{\bar{v}_b^2 \gamma_b^3}} \quad (58)$$

at which the growth rate is

$$\frac{\Gamma_m}{\omega_p} = \frac{1}{2\sqrt{2}\gamma_b^3}, \quad (59)$$

where $\Gamma_m = \Im(\omega_m)$ is the growth rate of that wavenumber.

The anticipated growth rates for both the relativistic and non-relativistic regimes are shown in Figure 9. At low beam velocities the finite temperatures of the beams suppress the growth rates appreciably relative to the cold-plasma limit,

Figure 7. Evolution of standing linear plasma mode in the regime of high linear Landau damping rate ($\hat{k} = 0.45$). The figure shows the fitting of the Fourier component of the grid electric field that corresponds to the initially excited mode. The noise level (blue line) corresponds to the estimated Fourier component of the electric field in Equation (107), i.e., the noise due to the finite number of macro-particles distributed uniformly on a periodic grid.

highlighting the importance of numerically solving the dispersion relation in the non-relativistic regime. In both cases (non-relativistic and relativistic streams) investigated below, we use a fifth-order interpolation function W^5 .

4.3.2. Non-relativistic Two-stream Simulations

The non-relativistic two-stream simulations are initialized with an initial co-moving temperature of the streams at $\theta = 10^{-4}$. We use $\Delta x \sim 0.001c/\omega_p$ and $L/\lambda_m = 20$, where $\lambda_m = 2\pi/k_m$ is the fastest growing wavelength. We start with a uniform distribution of macro-particles. Therefore, the initial distribution function is given by

$$f(x, v, t = 0) = \frac{1}{\sqrt{2\pi}\theta} \left[\frac{e^{-(\bar{v}-\bar{v}_b)^2/2\theta}}{2} + \frac{e^{-(\bar{v}+\bar{v}_b)^2/2\theta}}{2} \right]. \quad (60)$$

Other simulation parameters and the theoretical prediction are given in Table 2.

Generally, these simulations exhibit excellent quantitative agreement with the results of the linear theory. This is clearly evident in Figure 10, which shows the evolution of the amplitude of the most rapidly growing mode (based on linear

Table 2
Non-relativistic Two-stream Simulation Parameters

v_b/c	N_c^a	N_{pc}^b	Γ_m/ω_p^c	$k_m c/\omega_p^d$	$\hat{\omega}_m^1$ and $\hat{\omega}_m^{2e}$	$\hat{\omega}_m^{3e}$	$\hat{\omega}_m^{4e}$	$\hat{\omega}_m^{5e}$
0.02	5046	900	0.168555	24.9	$\pm 1.40842 - 0.0161625i$	$-0.485447i$	$-0.915823i$	$-1.32416i$
0.025	5324	900	0.243771	23.6	$\pm 1.47658 - 0.0123756i$	$-0.342545i$	$-0.673817i$	$-0.987125i$
0.03	6013	899.85	0.284945	20.9	$\pm 1.48041 - 0.00479742i$	$-0.249662i$	$-0.48542i$	$-0.715823i$
0.035	6904	900	0.307734	18.2	$\pm 1.46288 - 0.00103408i$	$-0.213058i$	$-0.356584i$	$-0.524922i$
0.04	7903	900	0.321003	15.9	$\pm 1.44312 - 0.00011898i$	$-0.247524i$	$-0.391793i$	$-0.527878i$
0.05	9974	899.91	0.334645	12.6	$\pm 1.41697 - 3.246 \times 10^{-7}i$	$-0.340762i$	$-0.388706i$	$-0.494053i$
0.06	12083	900	0.34111	10.4	$\pm 1.40108 - 1.294 \times 10^{-10}i$	$-0.341415i$	$-0.467936i$	$-0.496439i$
0.08	16110	900	0.346906	7.8	$\pm 1.39159 - 3.816 \times 10^{-19}i$	$-0.346906i$	$-0.396172i$	$-0.524619i$
0.09	18212	900	0.348379	6.9	$\pm 1.38628 - 1.148 \times 10^{-24}i$	$-0.348379i$	$-0.548489i$	$-0.556374i$

Notes.

^a Number of cells.

^b Number of computation particles per cell.

^c Theoretical predictions for the maximum growth rate.

^d Fastest growing wave mode predicted theoretically.

^e $\hat{\omega} = \omega/\omega_p$: other solutions of Equation (55) at the fastest growing wave mode k_m .

theory) in the non-relativistic simulations as a function of linear growth times. These should be compared to the solid black line, which shows the expected exponential growth, i.e., $e^{\Gamma_m t}$; the correspondence lasts over 4–8 e -folding times, i.e., 2–3.5 orders of magnitude, ending when the instability saturates non-linearly.

The initial oscillations correspond to other solutions of the dispersion relation, i.e., they are fully described by the linear analysis of the two-stream instability: in addition to the ultimately dominant exponentially growing mode, the linear dispersion relation admits other damping and oscillatory modes.

The inset in Figure 10 shows a reconstruction of the $v_b/c = 0.02$ simulation, which shows prominent oscillations at the beginning, in terms of the full complement of plasma modes shown in Table 2. Thus, the simulation quantitatively reproduces all of the anticipated linear features. A similar exercise is possible with the remaining non-relativistic simulations as well.

4.3.3. Relativistic Two-stream Simulations

We also performed a series of simulations for streams moving with relativistic speeds $u_b = \gamma_b v_b$. In all simulations, the initial co-moving temperature of streams is $\theta = 3 \times 10^{-3}$, $\Delta x \sim 0.05c/\omega_p$, and $L/\lambda_m = 10$, where $\lambda_m = 2\pi/k_m$ is the fastest growing wavelength. We start with uniformly distributed macro-particles, i.e., the instability here also grows from numerical noise. The initial distribution function is given by a combination of Maxwell–Jüttner distributions:

$$f(x, v, t = 0) = \frac{[e^{\bar{u}\bar{u}_b/\theta} + e^{-\bar{u}\bar{u}_b/\theta}]e^{-\gamma_b\gamma/\theta}}{4K_1(1/\theta)}, \quad (61)$$

where K_1 is the Bessel function of the first kind, and θ is the temperature in the co-moving frame of each beam. All other simulations parameters, along with theoretical predictions, are shown in Table 3.

Figure 11 shows again an excellent quantitative agreement between the growth rates of the fastest growing Fourier component of the electric field in different simulations and the theoretical predictions (solid-black curve). As in the non-relativistic case, the simulation exhibits exponential growth with the anticipated growth rate over three to four e -folding

Figure 8. Numerical solutions to the non-relativistic two-stream instability dispersion relation (Equation (55)). The blue curve shows the dependence of maximum growth rate on the stream speed (and temperature). The red curve shows the dependence of maximally growing wave mode on the stream speed (and temperature).

timescales. As for the non-relativistic case, the initial oscillations can be identified with the oscillatory, non-growing modes (described by and below Equation (57)). A fit for $u_b/c = 4.0$ that includes the oscillatory components is shown as an inset inside Figure 11.

5. Comparison with TRISTAN-MP

TRISTAN-MP is a publicly available PIC code to study plasma physics relevant for astrophysical problems (e.g., Spitkovsky 2008; Riquelme & Spitkovsky 2009; Philippov et al. 2015). Here, we compare some of the results we obtain from SHARP-1D with those obtained using TRISTAN-MP. In all cases, the same initial data are used. Generally, we find a substantial improvement in the ability to conserve energy and avoid numerical heating in SHARP simulations, and a good agreement in short timescale phenomena for which energy non-conservation is not substantial.

In our test problems, the performance of SHARP compares very favorably to TRISTAN-MP, typically running roughly an order of magnitude faster. We caution, however, that this may not be an entirely fair comparison since TRISTAN-MP, as a 3D code, may not be optimized for 1D problems.

Figure 9. Theoretical predictions for the two-stream instability growth rates in both relativistic and non-relativistic regimes. The black data points correspond to our simulated growth rates. Here $u_b = \gamma_b v_b$.

5.1. Numerical Particulars of TRISTAN-MP

In addition to the initial conditions, TRISTAN-MP has a number of specific numerical parameters that impact its performance. While we are unable to perform an exhaustive analysis of each, we did explore the result of varying a handful of these. TRISTAN-MP is a 3D and 2D PIC code. Here, we use the 2D version with one or two cells in one of the spatial dimensions to run it in an effective 1D setup, which enables a fair comparison to SHARP-1D.

5.1.1. Filtering

TRISTAN-MP provides the ability to low-pass filter the deposited grid moments, e.g., current densities on the grid, damping high-frequency noise prior to using them to solve Maxwell's equations. This reduces the coupling between the wave modes resolved on the grid with their aliases leading to improvements in the momentum and energy conservation of the algorithm. In TRISTAN-MP, filtering is accomplished with a three-cell stencil that generates a weighted average between the charge current density in a given cell with its neighbors. This operation may be repeated as many times as desired, smoothing the moments on progressively larger scales.

It is not a priori clear how many passes of the three-cell filtering operation are optimal in a given problem and we experimented with a number of different choices for the comparison problems presented here. We find that after a small number of filtering passes, typically three to five, the qualitative improvement is only moderate for $\theta \lesssim \theta_D$ and negligible for $\theta > \theta_D$.

This modest improvement comes with the additional computational cost, set by the addition of a substantial number of transverse grid cells required by many filtering passes. Thus, when comparing the numerical heating in SHARP-1D (when W^S is used), we use four filtering passes with a 2D simulation box for TRISTAN-MP that is only two cells wide in the x -direction. All other comparisons employ only three filtering passes or less with a 2D simulation box for TRISTAN-MP that is only one cell wide in the x -direction.

5.1.2. Electromagnetic Mode Speeds

To suppress the numerical Cerenkov instability in TRISTAN-MP it is possible to independently set the ratio of the propagation speed of transverse electromagnetic modes to the speed of light.

Figure 10. Two-stream instability simulation results in the non-relativistic regime. Different curves show the growth of the maximally growing mode (predicted by the theory) with time in units of the growth rate of this mode (predicted also theoretically) for streams with different speeds. The solid-black line is a line with slope = 1. An excellent agreement between simulation results and theoretical predictions is therefore evident: the growth rate found in the simulation and the rates calculated by numerically solving the linear dispersion relation for such plasma, Equation (55), are quantitatively similar. The inset shows a reconstruction of the $v_b/c = 0.020$ simulation in terms of the full complement of plasma modes shown in Table 2. Note that Γ_m is different for different beam velocities v_b (see Table 2 and Figure 9). Since the instability grows from the noise in all simulations, we shifted the time so that linear-phase instability growth starts at the same time for all simulations.

This is implemented explicitly via an additional numerical coefficient in Maxwell's equations. Typically, this is set near unity, e.g., 1.025. However, in the 1D electrostatic case, the numerical Cerenkov instability does not exist, and we have verified that this factor does not qualitatively change any of the results from the TRISTAN-MP simulations.

5.2. Thermal Stability and Energy Conservation

In Figure 12, we compare the evolution of the temperature and energy error in a pair of simulations described in Section 4.1. These are chosen such that in one case the Debye length is resolved by the grid cell (red curves) and when it is not resolved (blue curves). In both cases, TRISTAN-MP (dashed lines) exhibits a significantly larger violation of energy conservation, differing only in the timescale over which this occurs. When the Debye length is resolved, the numerical heating occurs more slowly, becoming untenable only after 3.5×10^6 plasma timescales. On the other hand, when the Debye length is not resolved, the numerical heating dominates the initial thermal energy almost instantly. In both cases, by $4 \times 10^6 \omega_p^{-1}$, both simulations have generated similar relative degrees of numerical heating, i.e., the ratio of the energy errors to the original thermal energy of the plasma.

In comparison, SHARP-1D (solid lines) reduces the numerical heating rate drastically. When the Debye length is resolved (high temperature) the fractional errors are fixed near 10^{-5} throughout the simulation. Lower temperature plasmas exhibit similar absolute heating rates, and therefore the relative heating for cold plasmas appears larger. However, even when the Debye length is unresolved by an order of magnitude, the plasma continues to be

Table 3
Relativistic Two-stream Simulation Parameters

u_b/c	v_b/c	N_c^a	N_{pc}^b	Γ_m/ω_p^c	$k_m c/\omega_p^d$	$\hat{\omega}_m^1$ and $\hat{\omega}_m^2^e$	$\hat{\omega}_m^3^e$
1	0.707	2441	491.6	0.2102	0.5147	± 0.8142	$-0.2102i$
2	0.894	6138	488.76	0.1057	0.2047	± 0.4095	$-0.1057i$
3	0.948	10948	493.24	0.0629	0.1148	± 0.2435	$-0.0629i$
4	0.970	16668	497.96	0.0422	0.0754	± 0.1636	$-0.0422i$

Notes.

^a Number of cells.

^b Number of computation particles per cell.

^c Theoretical predictions for the maximum growth rate.

^d Fastest growing wave mode predicted theoretically.

^e $\hat{\omega} = \omega/\omega_p$: other solutions given by Equation (57) at the fastest growing wave mode k_m .

Figure 11. Two-stream instability simulation results in the relativistic regime. Different curves show the growth of the maximally growing mode (predicted by the theory) with time in units of the growth rate of this mode (predicted also theoretically) for streams with different speeds. The solid-black line is a line with slope = 1. An excellent agreement between simulation results and theoretical predictions is therefore evident: the growth rates found in the simulations and the rates calculated by solving the linear dispersion relation in the cold-limit Equation (59) are quantitatively similar. The inset here shows a fit for $u_b/c = 4.0$ that includes the oscillatory components given in Table 3. Note that for different stream velocities, Γ_m is different (see Table 3 and Figure 9). Since the instability here grows from the noise in all simulations, we shifted the time so that linear instability growth starts at the same time for all simulations.

well modeled. The origin of the improvement in the numerical heating is the improvement in the order of interpolation.

5.3. Stability of Standing Linear Plasma Waves

Here, we compare the evolution of a standing plasma wave, where the linear Landau damping can be ignored, i.e., the ability of both TRISTAN-MP and SHARP-1D to maintain a small amplitude oscillating wave mode.

The simulation setup is similar to that in Section 4.2: a fixed uniform background of ions with thermal electrons in a box with size $L = 40c/\omega_p$. The electrons are initially uniformly distributed. The initially excited mode is added through a velocity perturbation: we first initialize electron velocities using Equation (46) with $\theta = 10^{-3}$, then add a position dependent velocity perturbation to individual particles' velocity by adding $\beta \cos(2\pi x/\lambda)$ to their velocities, where the initially excited wavelength $\lambda = 20c/\omega_p$, with $\beta = 0.01\lambda/2\pi$, i.e., after about 0.25 of a plasma period this will introduce a density perturbation

with an amplitude of 0.01.¹⁵ The cell size is $\Delta x = 0.01c/\omega_p$ and since $kv_{th}/\omega_p = 0.009934$, linear Landau damping can be ignored. In all simulations described here, we fix the number of electrons per cell to $N_{pc} = 1250$. We also note that in all simulations the Debye length is well resolved, i.e., $\lambda_D = 3.162\Delta x$.

In Figure 13, we show the long-term evolution of the energy in this isolated wave mode. In both the SHARP-1D (solid curves) and TRISTAN-MP (dashed curves) simulations, the square of the amplitude ($|\tilde{E}_k|^2$) exhibits very small, long-timescale oscillations. In the SHARP-1D simulations, these are confined to within 0.8% of the initial value over the entire simulation. In contrast, the TRISTAN-MP simulations also exhibit a secular growth in the mode amplitude, leading to an approximate energy increase of 10% by $8 \times 10^4 \omega_p^{-1}$.

This behavior is independent of the interpolation order of SHARP-1D or number of smoothing filters employed in TRISTAN-MP. Even when employing W^1 with SHARP-1D, the mode amplitude continues to execute only small oscillations about the fixed value, accurately reproducing the expectation for the linear evolution of the mode. Because the Debye length is well resolved in this case, filtering improves the energy conservation only slightly in TRISTAN-MP, again, making little difference to the mode evolution.

The origin of the unphysical growth in the mode in the TRISTAN-MP simulations is unclear. The heating of the background is insufficient to appreciably Landau damp or excite the mode. We have run additional simulations with SHARP-1D employing a momentum non-conserving scheme in SHARP-1D with W^1 , i.e., using Equation (24) (which is the back-interpolation scheme used in TRISTAN-MP), finding similar results. That is, excluding both the differences in order and the back-interpolation of the fields as the source of the secular growth in the mode energy. The only untested distinction remaining between SHARP-1D and TRISTAN-MP is the way in which the electric fields are updated.

Next, we compare the evolution of a shorter wavelength mode ($kv_{th}/\omega_p = 0.45$) that has a high linear Landau damping rate. The simulation's setup is exactly the same as in Section 4.2.3, but with a lower number of macro-particles

¹⁵ Unlike the simulations in Section 4.2, here we initialize a perturbation in the velocity, which is more easily done with the native initialization routines of TRISTAN-MP. We have verified that there are no significant differences when the mode is initialized as a density perturbation, and hence do so when we later compare a Landau damped wave.

Figure 12. Comparison of the numerical heating (top) and energy non-conservation (bottom) when TRISTAN-MP is used (four filters, dashed curves) and when SHARP-1D (fifth order, solid curves) is used. Here, $\theta = k_B T / m_0 c^2$ is the normalized temperature, $\Delta \mathcal{E}$ is the energy change/error in the total energy, and $\mathcal{E}_{th}$ is the initial thermal energy, i.e., excluding rest mass energy of macro-particles, therefore $\Delta \mathcal{E} / \mathcal{E}_{th}$ measures the fractional error with respect to the initial thermal energy of plasma. We compare the evolution of the plasmas with initial normalized temperatures of $\theta_i = 10^{-1}$ (red curves) and $\theta_i = 10^{-4}$ (blue curves) for the two codes. The dashed black line in the top panel shows the Debye temperature θ_D .

($N_t = 5 \times 10^7$). The timescale for damping is much smaller than the time needed for the energy non-conservation to affect the evolution of such modes. Thus, the result from both TRISTAN-MP (dashed blue curve) and SHARP-1D (solid red line) match exactly as shown in Figure 14. As noted before, the initial drop in the wave amplitude is also present when other simulation methods are used (e.g., Besse & Sonnendrücker 2003; Rossmanith & Seal 2011).

5.4. Two-stream Instability

We now compare the performance on dynamical instabilities, i.e., the two-stream instability as described in Section 4.3, for non-relativistic and relativistic streams. Here, again the timescale on which such instability grows, is much shorter than the timescale needed for non-energy conservation to affect the evolution. Therefore, we find the same linear evolution in both codes.

Figure 15 shows a comparison for two of the non-relativistic stream velocities reported in Section 4.3.2. As before, time is measured in e -foldings of the most unstable mode. In Figure 16, comparisons for two relativistic stream simulations reported in Section 4.3.3 are shown. Again the same linear evolution of the instability is found in both codes.

Figure 13. Evolution of the averaged energy in the initially excited mode (blue), energy in all other modes resolved by the grid (red), and averaged energy error (black). Note that all energies, including the energy error, are normalized by the initial average energy in the excited mode, which is 16.87% of the initial thermal energy in the plasma, i.e., excluding the rest mass energy. Results from SHARP-1D are shown as solid lines, while results from TRISTAN-MP are shown as dashed lines. The averaging is done over 37 plasma periods, while the normalization is done with respect to the initial average energy in the excited mode, i.e., in the first 37 plasma periods. Top (bottom) panel shows a comparison of the SHARP simulations employing W^1 (W^5) to the TRISTAN-MP simulations with no (three) filtering passes.

6. Convergence

At the end of Section 2.2, we discussed the accuracy and different errors introduced in our numerical scheme. We saw that the dominant error is of the order of $O(h^3)$, arising from the order of particle pusher. Here, we assess the convergence of SHARP-1D and demonstrate that the numerical error decreases as expected. In particular, we develop a general criterion for convergence studies of PIC simulations by requiring that the ratio of the energy in the shortest wave mode to the energy in the Poisson noise of simulation to be, at least, fixed.

We then present a test case where the error in the total energy of plasmas is used as the measure of error: such a test shows that the definition for convergence motivated above leads to a decrease in the error at the expected rate and typical methods to test for convergence fail: a slower decrease in the error is observed as resolutions increase leading to a plateau in the error, where increasing the resolutions no longer lead to a decrease in the error.

6.1. “Resolution” in PIC Algorithms

In general, for a PIC-type algorithm, three notions of resolution are relevant for simulations.

Figure 14. Comparison of the evolution of a linear plasma wave in the regime where the linear Landau damping rate is high ($k = 0.45$). Here, we compare the results of SHARP-1D with W^5 (solid curves) to TRISTAN-MP with three filtering passes (dashed curves). Since, the number of particles here is lower than what was used in Section 4.2.3, there is a higher level of noise leading to a slightly faster damping rate than seen in Figure 7.

1. Spatial resolution of the grid, i.e., h . This also determines the temporal resolution.
2. Momentum resolution, set by the number of particles used to construct the charge and current density at each cell, i.e., the number of particles per cell, N_{pc} .
3. Spectral resolution, set by the size of the “spectral-cell,” which, for each spatial-dimension, is given by of $2\pi/L$, where L is the box-size.

The third is rarely discussed in PIC simulations and arises when a physical phenomena (in linear or nonlinear regimes) has a narrow spectral support. In such cases, higher resolution simulations will require increasing the three types of resolutions simultaneously. Here we will focus on the first two, leaving a complete discussion of the third for future work.

Generally, it is necessary to increase all relevant resolutions simultaneously to study algorithmic convergence. As we will see below, this requires increasing the number of particles per cell, rather than fixing it as is typically done. A similar requirement was found for smooth-particle hydrodynamics (SPH) simulations, where the convergence also requires increasing the number of fluid-particles within the smoothing volume of each particle to study convergence (Zhu et al. 2015).

6.2. Definition of Convergence—Equivalent Simulation

The notions of both spatial resolution (i.e., h) and momentum resolution (i.e., N_{pc}) place different constraints on the range of underlying wave modes that can effectively be simulated. Thus, some care must be taken to ensure that as these resolutions are increased simultaneously, the underlying wave complement of the physical system resolves ever smaller scales.

The discrete nature of the macro-particles places a floor on the amplitude of a mode that can be effectively resolved.¹⁶ The average potential energy of the particle distribution, or Poisson

Figure 15. Comparison of the non-relativistic two-streams instability. Here, we compare the results of SHARP-1D with W^5 (solid curves) to TRISTAN-MP with three filtering passes (dashed curves).

Figure 16. Comparison of the relativistic two-streams instability. Here, we compare the results of SHARP-1D with W^5 (solid curves) to TRISTAN-MP with three filtering passes (dashed curves).

noise, is

$$\mathcal{E}_{\text{noise}}^m = \frac{L}{12\epsilon_0} \left[1 - \frac{6f_m}{N_c} \right] q_0^2 N_p, \quad (62)$$

where f_m is a coefficient that depends on the spatial order of the algorithm. Any mode with an energy less than $\mathcal{E}_{\text{noise}}^m$ is effectively unresolvable.

For comparison, we compute the energy in a single plasma mode that can be resolved on the grid.¹⁷ For a single mode in plasma with charge density given by

$$\rho(x, t) = -A \left(\frac{q_0 N_p}{L} \right) \cos(\omega t) \cos(2\pi x / \lambda), \quad (63)$$

the electric field is given by

$$\begin{aligned} E(x, t) &= E_0(t) + \int_0^x \frac{dx' \rho(x', t)}{\epsilon_0} \\ &= E_0(t) - \left(\frac{q_0 N_p}{L} \frac{\lambda}{2\pi} \right) \frac{A}{\epsilon_0} \cos(\omega t) \sin(2\pi x / \lambda), \end{aligned} \quad (64)$$

¹⁶ This is what sets the thermal floor delineated by θ_p in Equation (45).

¹⁷ The wavelength λ is resolved by the grid, if λ/L is an integer.

where A is the amplitude of the initial perturbation for a mode with wavelength λ . Hence, the total electric field energy is given by

$$\begin{aligned}\mathcal{E}(t) &= \frac{\epsilon_0}{2} \int_0^L E(x, t)^2 dx \\ &= \frac{\epsilon_0 L}{2} E_0^2(t) + \frac{1}{2\epsilon_0} \left(\frac{q_0 N_p}{L} \frac{\lambda}{2\pi} \right)^2 A^2 \cos^2(\omega t) \frac{L}{2}.\end{aligned}\quad (65)$$

Averaging over a full-period and assuming that $\langle E_0^2 \rangle = 0$, the averaged potential energy in the wave mode is then

$$\langle \mathcal{E} \rangle = \frac{1}{2\epsilon_0} \left(\frac{q_0 N_p}{L} \frac{\lambda}{2\pi} \right)^2 A^2 \frac{L}{4}.\quad (66)$$

Whether or not a mode can be resolved is then determined by the ratio of $\langle \mathcal{E} \rangle$ to $\mathcal{E}_{\text{noise}}^m$,

$$r = \frac{\langle \mathcal{E} \rangle}{\mathcal{E}_{\text{noise}}^m} = \frac{3}{2} \frac{A^2 N_p}{(2\pi L/\lambda)^2} \left[1 - \frac{6f_m}{N_c} \right]^{-1} \approx \frac{3A^2}{8\pi^2} \frac{\lambda^2 N_p}{L^2}.\quad (67)$$

When $r > 1$, the mode is resolved on the grid, while when $r < 1$ it is dominated by the Poisson noise in the simulation and rapidly randomized. Because $r \propto \lambda^2$ this statement is also a function of wavelength, with the smallest wavelength modes being the most marginal. That is, the Poisson noise limit, $\mathcal{E}_{\text{noise}}^m$, sets a minimum mode wavelength, $\lambda_{\min}$ the simulation can resolve, independent of the spatial resolution of the grid. Improving the fidelity of the simulation requires, therefore, concurrent increases in spatial resolution (e.g., h), momentum resolution (e.g., N_{pc}), and spectral resolution (e.g., L).

Explicitly, requiring that modes on the smallest spatial scales are resolved, i.e., $\lambda_{\min} \propto h$, then translates into the requirement that hN_{pc}/L is, at least, fixed.¹⁸ When spectral resolution is not important, this requires that if the spatial resolution increases by a factor η then N_{pc} must grow by a similar factor, i.e.,

$$h \rightarrow h/\eta \quad \text{and} \quad N_{\text{pc}} \rightarrow \eta N_{\text{pc}}.\quad (68)$$

This is unsurprising—convergence requires simultaneous and equal increases in the spatial and momentum resolutions. It does mean, however, that convergence studies are numerically demanding, as they scale as η^3 , even in 1D. This is, of course, exactly the factor one would find in a purely Eulerian scheme for solving the 1D Boltzmann equation, which is similar to 2D hydrodynamics. It does make clear, however, that the inherent randomness of the particle description does nothing to improve the convergence characteristics.

6.3. SHARP-1D Convergence—An Explicit Example

We now provide an explicit example of convergence testing, as described in the previous section, using SHARP-1D. To illustrate both the convergence of SHARP-1D under this definition, and equally importantly, the lack of convergence under separate definitions often employed, we do this for an extreme range of η , extending over two orders of magnitude.

¹⁸ Such a requirement (a fixed ratio of the energy in the shortest wavelength mode to the Poisson energy) implies that the ratio of the energy in a specific mode to the energy in the Poisson noise of simulation increases as η^2 with improving resolution.

We begin with a fiducial simulation, which defines $\eta = 1$. This is comprised of a population of electrons, with a total number of macro-particles of $N_p = 8950$, and a fixed neutralizing background. We use a box with normalized length $\bar{L} = 39.96175$ and the initial normalized temperature for electrons $\theta_{\text{ini}} = 10^{-3}$. We start with a single excited mode with amplitude $A = 10^{-2}$ and wavelength of $\bar{\lambda} = \bar{L}/2$. Therefore, $\hat{k} = 2\pi\sqrt{\theta}/\bar{\lambda} = 0.0099441$, i.e., the linear perturbation should oscillate without damping during the entire simulation time, $T = 100\omega_p^{-1}$.

For our fiducial simulation, we set the cell size, Δx , such that $1/h = c/(\Delta x \omega_p) = 8.958$, i.e., $N_c = 358$ and $N_{\text{pc}} = N_p/N_c = 25$. For all simulations in this section, we use fifth-order interpolation (W^5). Note that in this simulation the box is sufficiently large to spectrally resolve all relevant features of the dispersion relation, and thus we do not consider it further here.

The accuracy measure we employ is the normalized maximum error over the duration of the simulation:

$$\epsilon_m \equiv \frac{\max(\Delta \mathcal{E})}{\mathcal{E}_{\text{th}}},\quad (69)$$

where $\Delta \mathcal{E}$ is the energy change in the total energy, $\mathcal{E}_{\text{th}}$ is the initial thermal energy, i.e., excluding rest mass energy. Using various definitions of the energy error, i.e., average error, result in qualitatively identical results. Note that this is not the only accuracy measure we might use; others include the amplitude or phase of the wave, or the ability to reproduce other known solutions. It does have the property that it is fundamentally well understood (energy is conserved), not explicitly conserved by the code (like momentum), and universally defined.

Figure 17 shows the impact of increasing the various relevant notions of resolution independently and together. Increasing either the spatial or momentum resolution independently leads to a plateau in ϵ_m . The location and magnitude of this plateau depends upon the value of the non-converging resolution, i.e., when converging in spatial resolution, N_{pc} , or when converging in momentum resolution, h . This is qualitatively distinct from the case when both resolutions are increased simultaneously, for which no plateau is evident over two orders of magnitude.

Quantitatively, using our definition of convergence, SHARP-1D converges as $\eta^{-2.8} \propto h^{2.8}$. This is very similar to the anticipated h^3 , with the implication that the algorithm performance is well understood. Because this is ultimately set by the currently second-order symplectic integration employed in the particle pusher, implementing higher-order spatial interpolation will not improve this convergence rate. However, the value of the higher-order spatial interpolation appears in the magnitude of the energy error (i.e., its coefficient). This is clearly evident in Figure 18, which shows the error in the $\eta = 1$ simulation for different interpolation orders. Nevertheless, achieving the full benefits of the higher-order spatial implementations will require implementing an appropriate particle pusher.

7. Performance of SHARP-1D

To quantify the increase in the computational cost of using higher-order interpolation functions, we ran a simulation using 8950 macro-particles on five processors for all implemented orders. In Figure 19, we show the relative

Figure 17. Effect of increasing different resolutions starting with our fiducial simulation, on the maximum of the normalized energy error $\epsilon_m \equiv \max(\Delta\mathcal{E})/\mathcal{E}_{\text{th}}$, where $\Delta\mathcal{E}$ is the energy change in the total energy, $\mathcal{E}_{\text{th}}$ is the initial thermal energy, i.e., excluding rest mass energy. The figures show the effect of increasing the spatial resolution, h , while fixing the momentum resolution, N_{pc} , (right), the effect of increasing the momentum resolution while fixing the spatial resolution (middle), and the effect of simultaneously increasing both momentum and spatial resolution (left). The red-star result, which is the same simulation for all plots here, corresponds to our fiducial simulation ($N_{\text{pc}} = 25$ and $1/h = 8.958$). We define $\eta \equiv N_{\text{pc}}/25 = 1/(8.958 h)$.

Figure 18. Effect of using higher-order interpolation functions on the normalized maximum energy error ϵ_m . Red, blue, green, and black lines are, respectively, the results after running the simulation up to 10^3 , 10^4 , 10^5 , and $10^6 \omega_p^{-1}$. This shows the importance of using higher-order interpolation functions in controlling the energy non-conservation, the effect is specially important for long time simulations. Here, the maximum of the normalized energy error $\epsilon_m \equiv \max(\Delta\mathcal{E})/\mathcal{E}_{\text{th}}$, where $\Delta\mathcal{E}$ is the energy change in the total energy, $\mathcal{E}_{\text{th}}$ is the initial thermal energy, i.e., excluding rest mass energy.

increase in computational time for both deposition and back-interpolation steps after running each simulation up to $t \omega_p = 10^3$ (red), 10^5 (green), and 10^6 (black). The computational cost per update of using W^5 is 2.28 times larger than W^1 . We have verified that this is independent of the number of macro-particles and the number of processors.

The advantages of using higher-order interpolation are problem dependent. However, if we use the error in the total energy as a measure of accuracy, we can attempt to quantify the difference by computing the relative computational cost of simulations with different interpolation orders holding the level of accuracy fixed. Figure 18 shows that for a simulation that runs until time $t = 10^6 \omega_p^{-1}$, the energy error, ϵ_m , is smaller by a factor of 10^3 when W^5 is used instead of W^1 .

To achieve a similar accuracy using W^1 , i.e., decreasing ϵ_m by a factor of 10^3 , η needs to be increased by $10^{3/2.8} \sim 11.8$ (where we employed the scaling in the right-hand panel of Figure 17). Consequently, both N_{pc} and N_c each have to increase by 11.8, increasing the number of steps by the same factor. Therefore, using W^1 , the computational cost increases

Figure 19. Increase in computational cost when higher-order interpolation functions are used in SHARP-1D, at a fixed number of macro-particles ($N_p = 8950$) and number of cells. The normalized time is the computation time used in both deposition and back-interpolation steps when simulations are run up to $t \omega_p = 10^3$ (red), 10^5 (green) and 10^6 (black), normalized to the corresponding time when W^1 is used. Here, we used five processors in all simulations. We have verified that the relative increase in the computational cost, we found here, is independent of the number of macro-particles and the number of processors.

by a factor of $11.8^3 \sim 1640$, ~ 730 times that required by W^5 . That is, to achieve the same level of accuracy, a simulation that uses W^5 is about 730 faster than a simulation that uses W^1 with improved resolutions.¹⁹

SHARP-1D exhibits a near linear strong scaling, i.e., for fixed problem size, with the number of processors employed, N_{pr} . By varying N_{pr} between 20 and 300, we find

$$t \sim N_{\text{pr}}^{-0.96}. \quad (70)$$

8. Conclusion

In this paper, we present a self-consistent discretization for the governing equations of plasma made of macro-particles in 1D (i.e., the Vlasov–Poisson equations) implemented in the SHARP-1D code. It employs a self-consistent force on such

¹⁹ For reference, on Intel Xeon 3.47 GHZ CPUs, the computational time to evolve 8950 macro-particles on five processors for 22396417 steps is 1732 seconds per processor when using W^1 .

macro-particles that is accurate up to fifth order and provides an essential step toward higher-order accurate PIC schemes. The overall accuracy of the algorithm is, however, limited by the accuracy of the particle pusher which is still a second order accurate symplectic method (leap-frog).

SHARP-1D conserves momentum exactly, and despite its second-order accuracy, when higher-order interpolation functions are used, better energy conservation and lower numerical heating is evident. SHARP-1D simulations of a thermal plasma, whose Debye length is 10 times smaller than the cell size and which employ spatial interpolation accurate to fifth order only have an energy error, which is better than 1% of the initial thermal energy. Moreover, it shows a negligible numerical heating over a very long time (up to millions of inverse plasma frequencies, see Figure 3).

We present a validation of SHARP-1D against some test problems: the thermal stability of plasmas, the stability of linear plasma modes, and the two-stream instability in the relativistic and non-relativistic regimes.

To perform such validation tests, we determine the correct modes of thermal plasmas (oscillation frequencies and damping and growing rates) by solving the corresponding linear dispersion relations numerically. This is done for thermal plasmas that are both stationary and counter streaming. For convenience, we provide a fit to the oscillation frequencies and the damping rates in the linear regime of thermal plasmas up to $\hat{k} = kv_{th}/\omega_p = 0.6$.

In all test problems, SHARP-1D demonstrates the ability to reproduce kinetic effects of the linear regime both qualitatively and quantitatively. This includes reproducing the correct oscillation frequencies and damping rates for different modes of the thermal plasma, and also reproducing all oscillating, growing, and damping modes in counter streaming plasmas in the relativistic and non-relativistic regimes.

Results from SHARP-1D in both relativistic and non-relativistic regimes are contrasted with results from TRISTAN-MP. A substantial improvement in the ability to conserve energy and control numerical heating is shown when SHARP-1D is used.

Importantly, the improved performance due to higher spatial order does not come at the cost of increased execution time; to achieve the same level of accuracy, we have shown that, for SHARP-1D, a simulation with W^5 is almost three orders of magnitude faster than a simulation with W^1 and improved resolutions.

Finally, we develop a general criterion for convergence studies of PIC simulations by requiring the ratio of the energy in the shortest wave mode to the energy in the Poisson noise of the simulation to be at least fixed.

An example study for such a convergence test is presented, where the decrease in the energy error for plasma, as different relevant resolutions are increased, is used as a measure for convergence. Both the number of particles per cell and the spatial resolution of the grid are crucial resolution elements: increasing only one relevant resolution results not only in a slower decrease of the error, but also in a plateau where the error does not decrease any longer as such resolution increases. Faster decrease of the error without any plateau is achieved when all relevant resolutions are increased simultaneously for such plasma.

This new PIC code provides a new avenue that enables the faithful study of the long-term evolution of plasma problems

(in one dimension) that require absolute control of the energy and momentum conservation. Those include, e.g., the oblique instability driven by the highly anisotropic TeV pair beams that emerge from TeV gamma-rays that propagate from blazars to us or interactions of relativistic plasma components with a non-relativistic background plasma over long timescales (Broderick et al. 2012).

We would like to thank Anatoly Spitkovsky for providing access to TRISTAN-MP, substantial guidance in its use, and a number of constructive comments on the manuscript. M.S. and A.E.B. receive financial support from the Perimeter Institute for Theoretical Physics and the Natural Sciences and Engineering Research Council of Canada through a Discovery Grant. Research at Perimeter Institute is supported by the Government of Canada through Industry Canada and by the Province of Ontario through the Ministry of Research and Innovation. P.C. gratefully acknowledges support from the NASA ATP program through NASA grant NNX13AH43G, and the NSF through grant AST-1255469. C.P. gratefully acknowledges support by the European Research Council through ERC-CoG grant CRAGSMAN-646955 and by the Klaus Tschira Foundation. E.P. gratefully acknowledges support by the Kavli Foundation. Support for A.L. was provided by an Alfred P. Sloan Research Fellowship, NASA ATP Grant NNX14AH35G, and NSF Collaborative Research Grant #1411920 and CAREER grant #1455342.

Appendix A Momentum Conservation

For the purpose of comparison, we start by calculating the correct interaction force for point particles in 1D, by using

$$\begin{aligned} q_0 \mathbb{E}_0 &= m_0 \epsilon_0 \omega_0 \quad \text{and} \quad h N_{cp} = \frac{\omega_0 L / c}{\sum_s (\bar{Q}_s^2 N_s) / M_s} = \frac{\omega_0 q_0^2}{m_0 c \epsilon_0 \omega_0} \\ &= \frac{q_0^2}{\epsilon_0 q_0 \mathbb{E}_0} \implies \frac{1}{\epsilon_0 q_0 \mathbb{E}_0} = \frac{h N_{cp}}{q_0^2}. \end{aligned} \quad (71)$$

Therefore, the correct interaction term in 1D for point particles ($m = 0$) (in code units) is given by

$$\begin{aligned} \bar{F}_{int} &= \frac{Q_1 Q_2}{2 \epsilon_0} \frac{1}{q_0 \mathbb{E}_0} \begin{cases} +1 & x_2 > x_1 \\ -1 & x_2 < x_1 \end{cases} \\ &= \frac{\bar{Q}_1 \bar{Q}_2 h}{2} N_{cp} \begin{cases} +1 & x_2 > x_1 \\ -1 & x_2 < x_1 \end{cases}. \end{aligned} \quad (72)$$

On the other hand, the force on a macro-particle, with charge q_α and centered at $x_\alpha \in [0, L]$ on a periodic box, is given by $\bar{F}_\alpha \equiv F_\alpha / (q_0 \mathbb{E}_0) = \bar{Q}_\alpha \bar{E}_\alpha$, where

$$\begin{aligned} \bar{E}_\alpha &= \int_0^L \bar{E}(x) S^m[(x - x_\alpha) / \Delta x] dx \\ &= \sum_{k=0}^{N_c-1} \int_{x_{k-1/2}}^{x_{k+1/2}} \bar{E}(x) S^m dx \\ &= \sum_{k=0}^{N_c-1} \int_{x_k}^{x_{k+1}} \bar{E}(x) S^m dx. \end{aligned} \quad (73)$$

By defining

$$N_{cp} \equiv \frac{N_c}{\sum_s (\bar{Q}_s^2 N_s) / \bar{M}_s} = \frac{1}{n_0 \Delta x} \quad \text{and} \quad N_{pc} \equiv \frac{1}{N_{cp}}$$

and $W_{k,i}^m \equiv W^m[(\bar{x}_k - \bar{x}_i)/h]$.

(74)

The exact equations for the fields on the grid edges are given by Equations (35), (36), these can be written as

$$\bar{E}_{k+1} = \bar{E}_k + h \bar{\rho}_{k+\frac{1}{2}} \quad \text{and} \quad \bar{\rho}_{k+\frac{1}{2}} = N_{cp} \sum_s \bar{Q}_s \sum_{i_s} W_{k+\frac{1}{2},i_s}^m.$$
(75)

The solution of Equation (75) can be expressed as

$$\begin{aligned} \bar{E}_k &= \bar{E}_0 + h \sum_{j=0}^{k-1} \bar{\rho}_{j+\frac{1}{2}} = \bar{E}_0 - h \sum_{j=k}^{N_c-1} \bar{\rho}_{j+\frac{1}{2}} \\ &= \bar{E}_0 + \frac{h}{2} \sum_{j=0}^{k-1} \bar{\rho}_{j+\frac{1}{2}} - \frac{h}{2} \sum_{j=k}^{N_c-1} \bar{\rho}_{j+\frac{1}{2}} \\ &= \bar{E}_0 - \frac{h}{2} \bar{\rho}_{k+\frac{1}{2}} + \frac{h}{2} \sum_{j=0}^{N_c-1} A_{jk} \bar{\rho}_{j+\frac{1}{2}}, \end{aligned}$$
(76)

where A_{jk} is anti-symmetric matrix given by

$$A_{jk} = \begin{cases} +1 & j < k, \\ 0 & j = k, \\ -1 & j > k. \end{cases}$$
(77)

Two possible, second-order accurate, approximations for Equation (73) can be written as follows

$$F_{is} \approx \begin{cases} \bar{Q}_s \sum_{k=0}^{N_c-1} \bar{E}_k W_{k,i_s}^m, \\ \frac{\bar{Q}_s}{2} \sum_{k=0}^{N_c-1} [\bar{E}_k + \bar{E}_{k+1}] W_{k+\frac{1}{2},i_s}^m = \frac{\bar{Q}_s}{2} \sum_{k=0}^{N_c-1} \bar{E}_k \left[W_{k+\frac{1}{2},i_s}^m + W_{k-\frac{1}{2},i_s}^m \right]. \end{cases}$$
(78,79)

A.1. Non-momentum Conserving Second-order Scheme: Approximation (78)

If we use (76) and net-charge neutrality, i.e., $\sum_k \bar{\rho}_k = \sum_k \bar{\rho}_{k+\frac{1}{2}} = 0$, then the total force on all macro-particles is given by

$$\begin{aligned} \bar{F}_{Net} &= \sum_s \sum_{i_s} \bar{F}_{is} = \sum_k \bar{E}_k \sum_s \bar{Q}_s \sum_{i_s} W_{k,i_s}^m = N_{pc} \sum_k \bar{E}_k \bar{\rho}_k \\ &= -N_{pc} \frac{h}{2} \sum_k \bar{\rho}_k \bar{\rho}_{k+1/2} + N_{pc} \frac{h}{2} \sum_{j,k} A_{jk} \bar{\rho}_k \bar{\rho}_{j+\frac{1}{2}}. \end{aligned}$$
(80)

Therefore, the net-force on the system does not depend on the choice of E_0 (because of charge neutrality), and since using higher interpolation functions makes the variation in the interpolated density smoother, it decreases both terms in Equation (80), i.e., using higher-order interpolation improves the momentum conservation.

A.1.1. Origin of Momentum Non-conservation: Self-forces and Wrong Interaction Forces

To see the origin of momentum non-conservation, we examine the interpolated force in the case in which there are only two macro-particles. Using $\bar{Q}_1 + \bar{Q}_2 = 0$, $\sum_k W_{k,i_s}^m = 1$, and

$$\bar{\rho}_{j+\frac{1}{2}} = N_{cp} \left[\bar{Q}_1 W_{j+\frac{1}{2},1}^m + \bar{Q}_2 W_{j+\frac{1}{2},2}^m \right],$$

the force on macro-particle at x_1 is given by

$$\begin{aligned} \bar{F}_1 &= \bar{Q}_1 \sum_{k=0}^{N_c-1} \left(\bar{E}_0 - \frac{h}{2} \bar{\rho}_{k+\frac{1}{2}} + \frac{h}{2} \sum_{j=0}^{N_c-1} A_{jk} \bar{\rho}_{j+\frac{1}{2}} \right) W_{k,1}^m \\ &= \bar{Q}_1 E_0 - \frac{\bar{Q}_1 h}{2} N_{cp} \sum_k \left[\bar{Q}_1 W_{k+\frac{1}{2},1}^m + \bar{Q}_2 W_{k+\frac{1}{2},2}^m \right] W_{k,1}^m \\ &\quad + \frac{\bar{Q}_1 h}{2} N_{cp} \sum_{j,k} A_{jk} \left[\bar{Q}_1 W_{k+\frac{1}{2},1}^m + \bar{Q}_2 W_{k+\frac{1}{2},2}^m \right] W_{k,1}^m \\ &= \bar{Q}_1 E_0 + \frac{\bar{Q}_1^2 h}{2} N_{cp} \left[\sum_{j,k} A_{jk} W_{j+\frac{1}{2},1}^m W_{k,1}^m - \sum_k W_{k+\frac{1}{2},1}^m W_{k,1}^m \right] \\ &\quad + \frac{\bar{Q}_2 \bar{Q}_1 h}{2} N_{cp} \left[\sum_{j,k} A_{jk} W_{j+\frac{1}{2},2}^m W_{k,1}^m - \sum_k W_{k+\frac{1}{2},2}^m W_{k,1}^m \right]. \end{aligned}$$
(81)

Therefore, the non-vanishing self-force, $\bar{F}_{self}$, and the numerical interaction force, $\bar{F}_{int}$, are given by

$$\begin{aligned} \bar{F}_{self} &= \frac{\bar{Q}_1^2 h}{2} N_{cp} \left[\sum_{j,k} A_{jk} W_{j+\frac{1}{2},1}^m W_{k,1}^m - \sum_k W_{k+\frac{1}{2},1}^m W_{k,1}^m \right] \neq 0 \\ \text{and } \bar{F}_{int} &= \frac{\bar{Q}_2 \bar{Q}_1 h}{2} N_{cp} \\ &\quad \times \left[\sum_{j,k} A_{jk} W_{j+\frac{1}{2},2}^m W_{k,1}^m - \sum_k W_{k+\frac{1}{2},2}^m W_{k,1}^m \right]. \end{aligned}$$
(82)

Therefore, the self-force ($\sim h/N_{pc} = L/N_p$) here vanishes only in the limit of infinite number of macro-particles $N_p \rightarrow \infty$.

A.2. Momentum Conserving Second-order Scheme: Approximation (79)

If we use (76) and net-charge neutrality, i.e., $\sum_k \bar{\rho}_k = \sum_k \bar{\rho}_{k+1/2} = 0$, and

$$\begin{aligned} \bar{E}_k + \bar{E}_{k+1} &= 2\bar{E}_0 + h \sum_{j=0}^{k-1} \bar{\rho}_{j+\frac{1}{2}} - h \sum_{j=k+1}^{N_c-1} \bar{\rho}_{j+\frac{1}{2}} \\ &= 2\bar{E}_0 + h \sum_{j=0}^{N_c-1} A_{jk} \bar{\rho}_{j+\frac{1}{2}}. \end{aligned}$$
(83)

The net-force is, then, given by

$$\begin{aligned}\bar{F}_{Net} &= \sum_s \sum_{i_s} \bar{F}_{i_s} = \sum_k \frac{\bar{E}_k + \bar{E}_{k+1}}{2} \sum_s \bar{Q}_s \sum_{i_s} W_{k+\frac{1}{2}, i_s}^m \\ &= N_{pc} \sum_k \frac{\bar{E}_k + \bar{E}_{k+1}}{2} \bar{\rho}_{k+\frac{1}{2}} = \frac{h N_{pc}}{2} \sum_{j,k} A_{jk} \bar{\rho}_{j+\frac{1}{2}} \bar{\rho}_{k+\frac{1}{2}} = 0. \quad (84)\end{aligned}$$

Therefore, the net-force on the system, again, does not depend on the choice of E_0 (because of charge neutrality), and it is always exactly zero, therefore, the momentum is exactly conserved. Hence we call this a momentum conserving scheme.

A.2.1. Vanishing of the Self-forces Exactly

To see how self-forces vanish, we examine the interpolated force in case in which there are only two macro-particles. Using $\bar{Q}_1 + \bar{Q}_2 = 0$, $\sum_k W_{k,i_s}^m = 1$ and

$$\bar{\rho}_{j+\frac{1}{2}} = N_{cp} \left[\bar{Q}_1 W_{j+\frac{1}{2},1}^m + \bar{Q}_2 W_{j+\frac{1}{2},2}^m \right],$$

the force on macro-particle at x_1 is given by

$$\begin{aligned}\bar{F}_1 &= \bar{Q}_1 \sum_{k=0}^{N_c-1} \frac{\bar{E}_k + \bar{E}_{k+1}}{2} W_{k+\frac{1}{2},1}^m \\ &= \bar{Q}_1 \sum_{k=0}^{N_c-1} \left[E_0 + \frac{h}{2} \sum_{j=0}^{N_c-1} A_{jk} \bar{\rho}_{j+\frac{1}{2}} \right] W_{k+\frac{1}{2},1}^m \\ &= \bar{Q}_1 \bar{E}_0 + \frac{\bar{Q}_1 h}{2} N_{cp} \sum_{j,k} A_{jk} \left[\bar{Q}_1 W_{j+\frac{1}{2},1}^m + \bar{Q}_2 W_{j+\frac{1}{2},2}^m \right] W_{k+\frac{1}{2},1}^m \\ &= \bar{Q}_1 \bar{E}_0 + \frac{\bar{Q}_1 \bar{Q}_2 h}{2} N_{cp} \sum_{j,k} A_{jk} W_{j+\frac{1}{2},2}^m W_{k+\frac{1}{2},1}^m.\end{aligned} \quad (85)$$

Therefore, the vanishing self-force, $\bar{F}^{\text{self}}$, and the numerical interaction force, $\bar{F}^{\text{int}}$, are given by

$$\bar{F}^{\text{self}} = \frac{\bar{Q}_1^2 h}{2} N_{cp} \sum_{j,k} A_{jk} W_{j+\frac{1}{2},1}^m W_{k+\frac{1}{2},1}^m = 0$$

$$\text{and} \quad \bar{F}^{\text{int}} = \frac{\bar{Q}_1 \bar{Q}_2 h}{2} N_{cp} \sum_{j,k} A_{jk} W_{j+\frac{1}{2},2}^m W_{k+\frac{1}{2},1}^m. \quad (86)$$

We can see from comparing Equations (72) and (86) that the numerically calculated interaction force is shorter in range, and that the usage of higher-order interpolation (larger macro-particles) results in smoothing the interaction forces.

A.3. Filtering for Momentum Conserving Scheme: Non-vanishing Self-forces

Here, we show that filtering of the grid charge densities, e.g., 1D equivalence of filtering implemented in TRISTAN-MP, lead to a non-vanishing self-forces and a violation of momentum conservation.

After one-filter, the filtered charge densities $\bar{\rho}_k^f$ are given by

$$\bar{\rho}_{k+\frac{1}{2}}^f = \frac{\bar{\rho}_{k-\frac{1}{2}} + 2\bar{\rho}_{k+\frac{1}{2}} + \bar{\rho}_{k+\frac{3}{2}}}{4}. \quad (87)$$

Therefore, for a momentum conserving scheme, e.g., Equation (79), if we used (83) and replaced $\rho_{k+1/2}$ by

$\rho_{k+1/2}^f$, the net force is given by

$$\begin{aligned}\bar{F}_{Net} &= \sum_k \frac{\bar{E}_k + \bar{E}_{k+1}}{2} \sum_s \bar{Q}_s \sum_{i_s} W_{k+\frac{1}{2}, i_s}^m \\ &= \frac{N_{pc} h}{2} \sum_{k,j} A_{jk} \bar{\rho}_{k+\frac{1}{2}}^f \bar{\rho}_{k+\frac{1}{2}} \\ &= \frac{N_{pc} h}{8} \sum_{k,j} A_{jk} \left(\bar{\rho}_{k-\frac{1}{2}} + \bar{\rho}_{k+\frac{3}{2}} \right) \bar{\rho}_{k+\frac{1}{2}} \neq 0. \quad (88)\end{aligned}$$

Which shows that filtering the grid deposited moment (grid charge density) leads to violation of momentum conservation in otherwise a momentum conserving scheme.

To see the origin of such violation, we look, as we did before, at the interpolated force in case of having only two macro-particles, using $\bar{Q}_1 + \bar{Q}_2 = 0$, $\sum_k W_{k,i_s}^m = 1$ and

$$\begin{aligned}\bar{\rho}_{j+\frac{1}{2}} &= N_{cp} \left[\bar{Q}_1 W_{j+\frac{1}{2},1}^m + \bar{Q}_2 W_{j+\frac{1}{2},2}^m \right] \\ \Rightarrow \bar{\rho}_{j+\frac{1}{2}}^f &= \frac{N_{cp}}{4} \left[\bar{Q}_1 \left(W_{j-\frac{1}{2},1}^m + 2W_{j+\frac{1}{2},1}^m + W_{j+\frac{3}{2},1}^m \right) \right. \\ &\quad \left. + \bar{Q}_2 \left(W_{j-\frac{1}{2},2}^m + 2W_{j+\frac{1}{2},2}^m + W_{j+\frac{3}{2},2}^m \right) \right],\end{aligned}$$

the force on macro-particle at x_1 is given by

$$\begin{aligned}\bar{F}_1 &= \bar{Q}_1 \sum_{k=0}^{N_c-1} \frac{\bar{E}_k + \bar{E}_{k+1}}{2} W_{k+\frac{1}{2},1}^m \\ &= \bar{Q}_1 \sum_{k=0}^{N_c-1} \left[E_0 + \frac{h}{2} \sum_{j=0}^{N_c-1} A_{jk} \bar{\rho}_{j+\frac{1}{2}}^f \right] W_{k+\frac{1}{2},1}^m \\ &= \bar{Q}_1 \bar{E}_0 + \frac{\bar{Q}_1 h}{4} N_{cp} \sum_{j,k} A_{jk} \left[\bar{Q}_1 W_{j+\frac{1}{2},1}^m + \bar{Q}_2 W_{j+\frac{1}{2},2}^m \right] W_{k+\frac{1}{2},1}^m \\ &\quad + \frac{\bar{Q}_1 h}{8} N_{cp} \sum_{j,k} A_{jk} \left[\bar{Q}_1 \left(W_{j+\frac{3}{2},1}^m + W_{j-\frac{1}{2},1}^m \right) W_{k+\frac{1}{2},1}^m \right. \\ &\quad \left. + \bar{Q}_2 \left(W_{j+\frac{3}{2},2}^m + W_{j-\frac{1}{2},2}^m \right) W_{k+\frac{1}{2},1}^m \right] \\ &= \bar{Q}_1 \bar{E}_0 + \frac{\bar{Q}_1 \bar{Q}_2 h}{4} N_{cp} \left[\sum_{j,k} A_{jk} W_{j+\frac{1}{2},2}^m W_{k+\frac{1}{2},1}^m \right. \\ &\quad \left. + \frac{1}{2} \sum_{j,k} A_{jk} \left(W_{j+\frac{3}{2},2}^m + W_{j-\frac{1}{2},2}^m \right) W_{k+\frac{1}{2},1}^m \right] \\ &\quad + \frac{\bar{Q}_1^2 h}{8} N_{cp} \sum_{j,k} A_{jk} \left(W_{j+\frac{3}{2},1}^m + W_{j-\frac{1}{2},1}^m \right) W_{k+\frac{1}{2},1}^m.\end{aligned} \quad (89)$$

Therefore, filtering leads to wrong interaction forces and a non-vanishing self-force given by

$$\bar{F}_1^{\text{self}} = \frac{\bar{Q}_1^2 h}{8} N_{cp} \sum_{j,k} A_{jk} \left(W_{j+\frac{3}{2},1}^m + W_{j-\frac{1}{2},1}^m \right) W_{k+\frac{1}{2},1}^m \neq 0.$$

Appendix B

Shape and Weight Functions Explicit Form

If we define $y = (x - x_i) / \Delta x = (\bar{x} - \bar{x}_i) / h$, the shape functions, $S^m(y)$, assumed for the macro-particles and the corresponding weight functions, $W^m(y)$, used in the interpolation steps in the code are given in Table 4.

Table 4
Shape and Weight Functions Implemented in the Code

m	$S^m(y)$	$W^m(y)$
0	$\delta(y)$	$\begin{cases} 1 & 0 \leq y \leq \frac{1}{2} \\ 0 & \text{Otherwise} \end{cases}$
1	$\frac{1}{\Delta x} \begin{cases} 1 & 0 \leq y \leq \frac{1}{2} \\ 0 & \text{Otherwise} \end{cases}$	$\begin{cases} 1 - y & 0 \leq y < 1 \\ 0 & \text{Otherwise} \end{cases}$
2	$\frac{1}{\Delta x} \begin{cases} 1 - y & 0 \leq y < 1 \\ 0 & \text{Otherwise} \end{cases}$	$\begin{cases} \frac{3}{4} - y^2 & 0 < y < \frac{1}{2} \\ \frac{1}{8}(3 - 2 y)^2 & \frac{1}{2} \leq y < \frac{3}{2} \\ 0 & \text{Otherwise} \end{cases}$
3	$\frac{1}{\Delta x} \begin{cases} \frac{3}{4} - y^2 & 0 < y < \frac{1}{2} \\ \frac{1}{8}(3 - 2 y)^2 & \frac{1}{2} \leq y < \frac{3}{2} \\ 0 & \text{Otherwise} \end{cases}$	$\begin{cases} \frac{2}{3} - y^2 + y ^3/2 & 0 < y < 1 \\ \frac{1}{6}(2 - y)^3 & 1 \leq y < 2 \\ 0 & \text{Otherwise} \end{cases}$
4	$\frac{1}{\Delta x} \begin{cases} \frac{2}{3} - y^2 + y ^3/2 & 0 < y < 1 \\ \frac{1}{6}(2 - y)^3 & 1 \leq y < 2 \\ 0 & \text{Otherwise} \end{cases}$	$\begin{cases} \frac{115}{192} - \frac{5y^2}{8} + \frac{y^4}{4} & 0 < y < \frac{1}{2} \\ \frac{1}{96}[55 + 20 y - 120y^2 + 80 y ^3 - 16y^4] & \frac{1}{2} \leq y < \frac{3}{2} \\ \frac{1}{384}(5 - 2 y)^4 & \frac{3}{2} \leq y < \frac{5}{2} \\ 0 & \text{Otherwise} \end{cases}$
5	$\frac{1}{\Delta x} \begin{cases} \frac{115}{192} - \frac{5y^2}{8} + \frac{y^4}{4} & 0 < y < \frac{1}{2} \\ \frac{1}{96}[55 + 20 y - 120y^2 + 80 y ^3 - 16y^4] & \frac{1}{2} \leq y < \frac{3}{2} \\ \frac{1}{384}(5 - 2 y)^4 & \frac{3}{2} \leq y < \frac{5}{2} \\ 0 & \text{Otherwise} \end{cases}$	$\begin{cases} \frac{11}{20} - \frac{y^2}{2} + \frac{y^4}{4} - \frac{ y ^5}{12} & 0 \leq y \leq 1 \\ \frac{17}{40} + \frac{5 y }{8} - \frac{7y^2}{4} + \frac{5 y ^3}{4} - \frac{3y^4}{8} + \frac{ y ^5}{24} & 1 < y < 2 \\ \frac{1}{120}(3 - y)^5 & 2 \leq y < 3 \\ 0 & \text{Otherwise} \end{cases}$

Appendix C Aliasing

The Fourier components of grid quantities, in 1D, g_k are such that $g_k = g_{k+pk_g}$, where p is some integer and $k_g = 2\pi/\Delta x$ is the wave mode associated with the cell size Δx on that physical grid. Therefore, for a continuous particles number density $n(x)$, the Fourier component of the grid charge density is given by (Birdsall & Langdon 1991)

$$\tilde{\rho}_k = \sum_{p=-\infty}^{\infty} \tilde{n}(k - pk_g) \tilde{S}(k - pk_g), \quad (90)$$

where $\tilde{S}(k)$ is Fourier transform of our interpolation function and $\tilde{n}(k)$ is Fourier transform of $n(x)$. Therefore, all aliases of k (wave modes that differ from k by integer number of k_g) contribute when grid quantities are calculated. Clearly, this will feedback on the particle quantities, when the grid quantities are used to calculate the force on the particles to evolve them. The strength of the coupling between aliases (the source of this error) depends on how fast $\tilde{S}^m(k)$ falls off for large k , as can be seen in Equation (90).

The Fourier transform of our interpolation functions (spline functions of the order of m , see Table 4) is given by

$$\tilde{S}^m(k) = \left[\frac{\sin(k\Delta x/2)}{k\Delta x/2} \right]^m. \quad (91)$$

Therefore, Using higher-order interpolation functions (larger m) in our code leads to a decrease in the strength of the

couplings between grid wave modes and their aliases), which results in improvements in energy conservation as seen in Section 4.1.

Appendix D Poisson Noise

Here we calculate the noise when a finite number of computational particles are used to represent a uniform distribution function. We calculate the total energy density due to such noise in Appendix D.1 and then find the power spectrum for such noise in Appendix D.2. In Appendix D.3, we calculate the temperature, θ_p , set by the energy in such noise.

D.1. Average Potential Energy from Uniformly Distributed Macro-particles

Using Equation (10) and the first equation in (2), we can write the electric field associated with plasma particles on a periodic box of length L , i.e., $x_{i_s} \in [0, L)$ as follows

$$E(x) - E_0 = \sum_s \frac{Q_s}{\epsilon_0} \sum_{i_s}^{N_s} q^m(x, x_{i_s}). \quad (92)$$

Where, $q^m(x, x_{i_s}) = \int_0^x dx' S^m(x', x_{i_s})$. The periodicity of the box implies that the plasma is neutral. Therefore,

$$E_L - E_0 = 0 = \frac{1}{\epsilon_0} \sum_s Q_s \sum_{i_s}^{N_s} q^m(L, x_{i_s}) = \frac{1}{\epsilon_0} \sum_s Q_s N_s. \quad (93)$$

The spatial averaging of q^m , for uniformly distributed macro-particles, is

$$\begin{aligned}\langle q^m(x, x_{i_s}) \rangle &= \int_0^L \frac{du}{L} q^m(x, u) = \int_0^x dx' \int_0^L \frac{du}{L} S^m(x', u) \\ &= \int_0^x \frac{dx'}{L} (1) = \frac{x}{L}.\end{aligned}\quad (94)$$

For such macro-particles the average of the electric field is zero:

$$\begin{aligned}\langle E(x) - E_0 \rangle &= \sum_s \left\langle \frac{Q_s}{\epsilon_0} \sum_{i_s}^{N_s} q^m(x, x_{i_s}) \right\rangle \\ &= \sum_s \frac{Q_s}{\epsilon_0} N_s \int_0^L \frac{du}{L} q^m(x, u) \\ &= \sum_s \frac{Q_s}{\epsilon_0} N_s \frac{x}{L} = \frac{x}{\epsilon_0 L} \sum_s Q_s N_s = 0.\end{aligned}\quad (95)$$

However, due to the finite number of macro-particles, the average potential energy is non-zero, to calculate such energy we need to calculate

$$\begin{aligned}\langle E^2(x) - E_0^2 \rangle &= \langle (E(x) - E_0)^2 \rangle \\ &= \left\langle \left(\sum_s \frac{Q_s}{\epsilon_0} \sum_{i_s}^{N_s} q^m(x, x_{i_s}) \right)^2 \right\rangle \\ &= \left\langle \sum_s \left(\frac{Q_s}{\epsilon_0} \sum_{i_s}^{N_s} q^m(x, x_{i_s}) \right)^2 \right\rangle + \left\langle \sum_{s \neq s'} \left(\frac{Q_s}{\epsilon_0} \sum_{i_s}^{N_s} q^m(x, x_{i_s}) \right) \left(\frac{Q_{s'}}{\epsilon_0} \sum_{i_{s'}}^{N_{s'}} q^m(x, x_{i_{s'}}) \right) \right\rangle \\ &= \frac{1}{\epsilon_0^2} \sum_s \left\langle \left(\frac{Q_s}{\epsilon_0} \sum_{i_s}^{N_s} q^m(x, x_{i_s}) \right)^2 \right\rangle \\ &+ \frac{1}{\epsilon_0^2} \sum_{s \neq s'} \left\langle \left(\frac{Q_s}{\epsilon_0} \sum_{i_s}^{N_s} q^m(x, x_{i_s}) \right) \left(\frac{Q_{s'}}{\epsilon_0} \sum_{i_{s'}}^{N_{s'}} q^m(x, x_{i_{s'}}) \right) \right\rangle \\ &= \frac{1}{\epsilon_0^2} \sum_s Q_s^2 \left\langle \sum_{i_s}^{N_s} [q^m(x, x_{i_s})]^2 + \sum_{i_s \neq j_s}^{N_s} q^m(x, x_{i_s}) q^m(x, x_{j_s}) \right\rangle \\ &+ \frac{1}{\epsilon_0^2} \frac{x^2}{L^2} \sum_{s \neq s'} Q_s N_s Q_{s'} N_{s'} \\ &= \sum_s \frac{Q_s^2}{\epsilon_0^2} \left[N_s \int_0^L \frac{du}{L} [q^m(x, u)]^2 + N_s(N_s - 1) \right. \\ &\times \left. \left[\int_0^L \frac{du}{L} q^m(x, u) \right]^2 \right] + \frac{1}{\epsilon_0^2} \frac{x^2}{L^2} \sum_{s \neq s'} Q_s N_s Q_{s'} N_{s'} \\ &= \frac{1}{\epsilon_0^2} \frac{x^2}{L^2} \left(\sum_{s \neq s'} Q_s N_s Q_{s'} N_{s'} + \sum_s N_s^2 Q_s^2 \right) \\ &+ \sum_s \frac{Q_s^2 N_s}{\epsilon_0^2} \left[\int_0^L \frac{du}{L} [q^m(x, u)]^2 - \frac{x^2}{L^2} \right] \\ &= \frac{x^2}{\epsilon_0^2 L^2} \left(\sum_s Q_s N_s \right)^2 + \sum_s \frac{Q_s^2 N_s}{\epsilon_0^2} \left[\int_0^L \frac{du}{L} [q^m(x, u)]^2 - \frac{x^2}{L^2} \right] \\ &= \sum_s \frac{Q_s^2 N_s}{\epsilon_0^2} \left[\int_0^L [q^m(x, u)]^2 \frac{du}{L} - \frac{x^2}{L^2} \right].\end{aligned}\quad (96)$$

For the shape functions implemented in SHARP-1D (their explicit forms are given in Appendix B), the integral in (96) is given by

$$\int_0^L \frac{du}{L} [q^m(x, u)]^2 = \frac{x}{L} - \frac{\Delta x}{L} f_m$$

$$\text{and } f_m = \frac{1}{6} \begin{cases} 0 & m = 0 \\ 1 & m = 1 \\ 1.4 & m = 2 \\ 1.70714 & m = 3 \\ 1.96693 & m = 4 \\ 2.19624 & m = 5 \end{cases}.\quad (97)$$

Therefore,

$$\langle E^2(x) - E_0^2 \rangle = \sum_s \frac{N_s Q_s^2}{\epsilon_0^2} \left[\frac{x}{L} - \frac{x^2}{L^2} - \frac{\Delta x}{L} f_m \right].\quad (98)$$

The average electrostatic potential energy due to the finite number of macro-particles is, then, given by

$$\begin{aligned}\mathcal{E}^m &= \frac{\epsilon_0}{2} \int_0^L dx \langle E^2(x) - E_0^2 \rangle = \sum_s \frac{N_s Q_s^2}{2\epsilon_0} \left[\frac{L}{6} - \Delta x f_m \right] \\ &= \frac{L}{12\epsilon_0} \left[1 - \frac{6f_m}{N_c} \right] \sum_s N_s Q_s^2.\end{aligned}\quad (99)$$

Here, $N_c = L/\Delta x$ is the number of macro-cells. If we assume that all plasma species have the same mass, and absolute value of charge, we then make the choice of our fiducial units as, $q_0^2 = Q_s^2$ and $m_0 = M_s$ (that implies $n_0 = \sum_s n_s$). Therefore,

$$\begin{aligned}\langle \bar{\mathcal{E}}^m \rangle &= \frac{\langle \mathcal{E}^m \rangle}{m_0 c^2} = \frac{L^2}{12c^2} \frac{q_0^2 n_0}{\epsilon_0 m_0} \left[1 - \frac{6f_m}{N_c} \right] \\ &= \frac{L^2 \omega_0^2}{12c^2} \left[1 - \frac{6f_m}{N_c} \right] = \frac{\bar{L}^2}{12} \left[1 - \frac{6f_m}{N_c} \right].\end{aligned}\quad (100)$$

Equation (100) shows that using higher-order shape functions decreases the noise coming from the fact that we are using a finite number of macro-particles. The decrease that we gain in the potential energy noise is $f_m \bar{L}^2 / 2N_c$. For a given box size, this improvement is lowered, if we increase the number of cells N_c because it means a decrease in cell size, which also means a decrease in the size of the macro-particles. On the other hand, if we increase the number of cells while keeping the cell size fixed, i.e., by increasing the box-size $\bar{L}$, that improvement due to using higher-order interpolation functions increases.

D.2. Spectrum of the Poisson Noise

To find the spectrum of Poisson noise, we average the Fourier transform of the grid electric fields. Using Equation (92), the Fourier components of the electric field are given by

$$\begin{aligned}\tilde{E}_n &= \int_0^L \frac{dx}{L} \left[E_0 + \sum_s \frac{Q_s}{\epsilon_0} \sum_{i_s}^{N_s} q^m(x, x_{i_s}) \right] e^{-2\pi i n x / L} \\ &= E_0 \delta_{n,0} + \sum_s \frac{Q_s}{\epsilon_0} \sum_{i_s}^{N_s} \int_0^L \frac{dx}{L} [q^m(x, x_{i_s})] e^{-2\pi i n x / L}.\end{aligned}\quad (101)$$

By defining

$$Z(x_{i_s}, n) \equiv \int_0^L \frac{dx}{L} q^m(x, x_{i_s}) e^{-2\pi i n x / L} \Rightarrow Z(x_{i_s}, 0) = \frac{x_{i_s}}{L}, \quad (102)$$

and using the fact that $E(x)$ is a real valued function, the absolute value for such Fourier components are given by

$$\begin{aligned} |\tilde{E}_n|^2 &= \tilde{E}_n \tilde{E}_{-n} = E_0^2 \delta_{n,0} + 2E_0 \delta_{n,0} \sum_s \frac{Q_s}{\epsilon_0} \sum_{i_s}^{N_s} Z(x_{i_s}, 0) \\ &\quad + \left[\sum_s \frac{Q_s}{\epsilon_0} \sum_{i_s}^{N_s} Z(x_{i_s}, n) \right] \left[\sum_s \frac{Q_s}{\epsilon_0} \sum_{i_s}^{N_s} Z(x_{i_s}, -n) \right] \\ &= E_0 \left[E_0 + 2 \sum_s \frac{Q_s}{\epsilon_0} \sum_{i_s}^{N_s} \frac{x_{i_s}}{L} \right] \delta_{n,0} \\ &\quad + \sum_s \frac{Q_s^2}{\epsilon_0^2} \left[\sum_{i_s}^{N_s} Z(x_{i_s}, n) \sum_{j_s}^{N_s} Z(x_{j_s}, -n) \right] \\ &\quad + \sum_{s \neq s'} \frac{Q_s Q_{s'}}{\epsilon_0^2} \left[\sum_{i_s}^{N_s} Z(x_{i_s}, n) \sum_{i_{s'}}^{N_{s'}} Z(x_{i_{s'}}, -n) \right] \\ &= E_0 \left[E_0 + 2 \sum_s \frac{Q_s}{\epsilon_0} \sum_{i_s}^{N_s} \frac{x_{i_s}}{L} \right] \delta_{n,0} \\ &\quad + \sum_s \frac{Q_s^2}{\epsilon_0^2} \left[\sum_{i_s}^{N_s} Z^m(x_{i_s}, n) Z^m(x_{i_s}, -n) \right. \\ &\quad \left. + \sum_{i_s \neq j_s}^{N_s} Z^m(x_{i_s}, n) Z^m(x_{j_s}, -n) \right] \\ &\quad + \sum_{s \neq s'} \frac{Q_s Q_{s'}}{\epsilon_0^2} \left[\sum_{i_s}^{N_s} Z^m(x_{i_s}, n) \sum_{i_{s'}}^{N_{s'}} Z^m(x_{i_{s'}}, -n) \right]. \quad (103) \end{aligned}$$

Averaging over such a periodic box and assuming the macro-particles are uniformly distributed, we can write

$$\begin{aligned} \langle |\tilde{E}_n|^2 \rangle &= E_0 \left[E_0 + 2 \sum_s \frac{Q_s N_s}{2\epsilon_0} \right] \delta_{n,0} \\ &\quad + \sum_s \frac{Q_s^2}{\epsilon_0^2} \left[N_s \int_0^L \frac{du}{L} [Z^m(u, n) Z^m(u, -n)] \right. \\ &\quad \left. + (N_s^2 - N_s) \int_0^L \frac{du}{L} Z^m(u, n) \int_0^L \frac{du}{L} Z^m(u, -n) \right] \\ &\quad + \sum_{s \neq s'} \frac{Q_s N_s Q_{s'} N_{s'}}{\epsilon_0^2} \int_0^L \frac{du}{L} Z^m(u, n) \int_0^L \frac{du}{L} Z^m(u, -n) \\ &= E_0^2 \delta_{n,0} + \left[\sum_{s \neq s'} Q_s N_s Q_{s'} N_{s'} + \sum_s Q_s^2 N_s^2 \right] \\ &\quad \times \int_0^L \frac{du}{L} Z^m(u, n) \int_0^L \frac{du}{L} Z^m(u, -n) \\ &\quad + \sum_s \frac{Q_s^2 N_s}{\epsilon_0^2} \left[\int_0^L \frac{du}{L} [Z^m(u, n) Z^m(u, -n)] \right. \\ &\quad \left. - \int_0^L \frac{du}{L} Z^m(u, n) \int_0^L \frac{du}{L} Z^m(u, -n) \right] \end{aligned}$$

$$\begin{aligned} &= E_0^2 \delta_{n,0} + \sum_s \frac{Q_s^2 N_s}{\epsilon_0^2} \left[\int_0^L \frac{du}{L} [Z^m(u, n) Z^m(u, -n)] \right. \\ &\quad \left. - \int_0^L \frac{du}{L} Z^m(u, n) \int_0^L \frac{du}{L} Z^m(u, -n) \right]. \quad (104) \end{aligned}$$

By using Equations (102), (94)

$$\begin{aligned} \int_0^L \frac{du}{L} Z^m(u, n) &= \int_0^L \frac{du}{L} \int_0^L \frac{dx}{L} q^m(x, u) e^{-2\pi i n x / L} \\ &= \int_0^L \frac{dx}{L} \frac{x}{L} e^{-2\pi i n x / L} = \begin{cases} 1/2, & n = 0, \\ i/(2\pi n), & n \neq 0, \end{cases} \quad (105) \end{aligned}$$

$$\begin{aligned} \int_0^L \frac{du}{L} Z^m(u, n) Z^m(u, -n) &= \begin{cases} 1/3, & n = 0, \\ \frac{1}{(2\pi n)^2} \left[1 + \left(\frac{\sin(\pi n / N_c)}{\pi n / N_c} \right)^{2m} \right], & n \neq 0. \end{cases} \quad (106) \end{aligned}$$

Therefore, the averaged magnitude for the Fourier components can be written as

$$\langle |\tilde{E}_n|^2 \rangle = \begin{cases} E_0^2 + \sum_s \frac{Q_s^2 N_s}{12\epsilon_0^2}, & n = 0, \\ \sum_s \frac{Q_s^2 N_s}{(2\pi n)^2 \epsilon_0^2} \left[\frac{\sin(\pi n / N_c)}{\pi n / N_c} \right]^{2m}, & n \neq 0. \end{cases} \quad (107)$$

If all plasma species have the same mass, and absolute value of charges, such average can be written in code units as

$$\langle |\tilde{E}_n|^2 \rangle = \frac{\langle |\tilde{E}_n|^2 \rangle}{\mathbb{E}_0^2} = \begin{cases} \tilde{E}_0^2 + \frac{\tilde{L}^2}{12N_t}, & n = 0, \\ \frac{1}{(2\pi n)^2} \left[\frac{\sin(\pi n / N_c)}{\pi n / N_c} \right]^{2m} \frac{\tilde{L}^2}{N_t}, & n \neq 0. \end{cases} \quad (108)$$

D.3. Heating Due to Noise

If the energy due to Poisson noise is converted to heat that puts a floor in the temperature PIC scheme can simulate. Here we estimate such a temperature floor θ_p . The energy due to Poisson noise is calculate in Appendix D.1. It is given by

$$\langle \tilde{\mathcal{E}}^m \rangle = \frac{\langle \mathcal{E}^m \rangle}{m_0 c^2} = \frac{\tilde{L}^2}{12} \left[1 - \frac{6f_m}{N_c} \right], \quad (109)$$

where f_m is defined in (97). This noise is due to the finite number of macro-particles used in the simulations. If this energy is converted to thermal energy, it would lead to heating of the plasma up to a temperature θ_p . If the plasmas are at thermal equilibrium, the momentum distribution is given by Maxwell–Jüttner distribution, and hence the temperature of plasmas is related to the kinetic energy, $\tilde{\mathcal{K}}$, as follows

$$\begin{aligned} \langle \tilde{\mathcal{K}} \rangle &= \frac{\langle \mathcal{K} \rangle}{m_0 c^2} = \sum_s \langle (\gamma - 1) \rangle = \sum_s N_s \left[\theta_s + \frac{K_0 [1/\theta_s]}{K_1 [1/\theta_s]} - 1 \right] \\ &= N \left[\theta + \frac{K_0 [1/\theta]}{K_1 [1/\theta]} - 1 \right] = \begin{cases} N\theta/2, & \theta \ll 1, \\ N\theta, & \theta \gg 1. \end{cases} \quad (110) \end{aligned}$$

Here, $\theta_s = k_B T_s / m_0 c^2$ is the normalized temperature of species s with N_s of macro-particles, N is the total number of macro-particles from all species and K_0, K_1 are the Bessel functions of zeroth and first kind respectively. Therefore,

$$\theta_p = \frac{\bar{L}^2}{12N} \left[1 - \frac{6f_m}{N_c} \right] \begin{cases} 2, & \theta_p \ll 1, \\ 1, & \theta_p \gg 1. \end{cases} \quad (111)$$

Hence, if a plasma of macro-particles starts with a temperature below θ_p , the Poisson noise will non-physically heat such plasma.

Appendix E

Dispersion Relation for Non-relativistic Warm Plasma

For non-relativistic ($\gamma^3 \approx 1$) warm plasma, i.e., $0 < \theta \ll 1$, we can write

$$\begin{aligned} f_0(u)du &= f_0(v)dv = \frac{n_0 d(v/c)}{\sqrt{2\pi\theta}} e^{-(v/c)^2/2\theta} \\ &= \frac{n_0 d\bar{v}}{\sqrt{2\pi\theta}} e^{-\bar{v}^2/2\theta} \quad \text{and} \quad \theta = \frac{k_B T}{mc^2}. \end{aligned}$$

If we assume no net current in the plasmas, i.e., the momentum distribution of all species is such that $\sum_s Q_s \int v f_0^s(v) dv = 0$, then the linear dispersion relation of uniformly distributed plasma is given by

$$1 = \sum_s \chi_s(v_p), \quad (112)$$

where $v_p = \hat{\omega}/\hat{k}$, $\hat{\omega} = \omega/\omega_p$, $\hat{k} = kc\sqrt{\theta}/\omega_p$, i.e., $v_p = \omega/kc\sqrt{\theta}$, $\omega_p^2 = \sum_s \omega_{ps}^2$, $\omega_{ps}^2 = Q_s^2 n_s / \epsilon_0 M_s$, and $\bar{v} = v/c$.

If we assume that $\Im(\hat{\omega}) > 0$, i.e., $\Im(\omega) > 0$, then

$$\begin{aligned} \chi_s(v_p) &= \frac{Q_s^2}{\epsilon_0 M_s k^2} \int_{-\infty}^{\infty} \frac{f_0(v)dv}{(v - \omega/k)^2} \\ &= \frac{\omega_{ps}^2}{k^2 c^2} \int_{-\infty}^{\infty} \frac{d\bar{v}}{\sqrt{2\pi\theta}} \frac{e^{-\bar{v}^2/2\theta}}{(\bar{v} - \omega/kc)^2} \\ &= \frac{\omega_{ps}^2 / \omega_p^2}{\hat{\omega}^2} \int_{-\infty}^{\infty} \frac{dz}{\sqrt{2\pi}} \frac{e^{-z^2/2}}{(z - v_p)^2}, \end{aligned} \quad (113)$$

where $z \equiv \bar{v}/\sqrt{\theta}$. Extending the definition of $\chi_s(v_p)$ to the entire complex plane can be done as follows (Brambilla 1998)

$$\begin{aligned} \chi_s(v_p) &= \frac{\hat{\omega}_s^2}{\hat{k}^2} \int_{-\infty}^{\infty} \frac{dz}{\sqrt{2\pi}} \frac{e^{-z^2/2}}{(z - v_p)^2} \\ &\quad - \frac{\hat{\omega}_s^2}{\hat{k}^2} \sqrt{\frac{\pi}{2}} v_p \begin{cases} 0 & \text{if } \Im(v_p) > 0 \\ ie^{-v_p^2/2} & \text{if } \Im(v_p) = 0 \\ 2ie^{-v_p^2/2} & \text{if } \Im(v_p) < 0, \end{cases} \end{aligned} \quad (114)$$

where $\hat{\omega}_s \equiv \omega_{ps}/\omega_p$. Therefore, for $\Im(v_p) \neq 0$, we can then write

$$\begin{aligned} \chi_s(v_p) &= \frac{\hat{\omega}_s^2}{\hat{k}^2} \begin{cases} -1 + \sqrt{\frac{\pi}{2}} v_p [\text{Erfi}(v_p/\sqrt{2}) - i] e^{-\frac{v_p^2}{2}} & \text{if } \Im(v_p) > 0 \\ -1 + \sqrt{\frac{\pi}{2}} v_p [\text{Erfi}(v_p/\sqrt{2}) + i] e^{-\frac{v_p^2}{2}} & \text{if } \Im(v_p) < 0 \end{cases} \\ &\quad - \sqrt{\frac{\pi}{2}} v_p \begin{cases} 0 & \text{if } \Im(v_p) > 0 \\ 2ie^{-v_p^2/2} & \text{if } \Im(v_p) < 0 \end{cases} \\ &= \frac{\hat{\omega}_s^2}{\hat{k}^2} \left[-1 + \sqrt{\frac{\pi}{2}} v_p [\text{Erfi}(v_p/\sqrt{2}) - i] e^{-\frac{v_p^2}{2}} \right], \end{aligned} \quad (115)$$

where Erfi is the complex error function, which is defined as $\text{Erfi}(v_p) = -i \text{Erf}(iv_p)$.

E.1. Standing Linear Plasma Waves

In the case of thermal electrons with fixed neutralizing background ($\hat{\omega}_s = 1$), the dispersion relation is then given by

$$\hat{k}^2 + 1 = \sqrt{\frac{\pi}{2}} v_p [\text{Erfi}(v_p/\sqrt{2}) - i] e^{-\frac{v_p^2}{2}} \quad (116)$$

E.2. Two-stream Instability

In the case of two population of thermal electrons (both have the same number density), propagating in two opposite directions with speed v_b , with fixed neutralizing background, therefore $\hat{\omega}_s^2 = 1/2$, and the linear dispersion relation is then given by ($z_b \equiv v_b/c\sqrt{\theta}$)

$$\begin{aligned} \hat{k}^2 + 1 &= \frac{1}{2} \left[\sqrt{\frac{\pi}{2}} (v_p + z_b) \left[\text{Erfi}\left(\frac{v_p + z_b}{\sqrt{2}}\right) - i \right] e^{-\frac{(v_p + z_b)^2}{2}} \right. \\ &\quad \left. + \sqrt{\frac{\pi}{2}} (v_p - z_b) \left[\text{Erfi}\left(\frac{v_p - z_b}{\sqrt{2}}\right) - i \right] e^{-\frac{(v_p - z_b)^2}{2}} \right] \\ &= \sqrt{\frac{\pi}{8}} \left[(v_p + z_b) \left[\text{Erfi}\left(\frac{v_p + z_b}{\sqrt{2}}\right) - i \right] e^{-v_p z_b} \right. \\ &\quad \left. + (v_p - z_b) \left[\text{Erfi}\left(\frac{v_p - z_b}{\sqrt{2}}\right) - i \right] e^{v_p z_b} \right] \\ &\quad \times e^{-(v_p^2 + z_b^2)/2}. \end{aligned} \quad (117)$$

References

- Ardaneh, K., Cai, D., Nishikawa, K.-I., & Lembége, B. 2015, *ApJ*, **811**, 57
 Besse, N., & Sonnendrücker, E. 2003, *JCoPh*, **191**, 341
 Birdsall, C. K., & Langdon, A. B. 1991, *Plasma Physics via Computer Simulation* (Bristol: Adam Hilger)
 Birdsall, C. K., & Maron, N. 1980, *JCoPh*, **36**, 1
 Boyd, T. J. M., & Sanderson, J. J. 2003, *The Physics of Plasmas* (Cambridge: Cambridge Univ. Press)
 Brackbill, J., & Cohen, B. I. (ed.) 1985, *Multiple Time Scales*, Vol. 3 (Orlando, FL: Academic Press Inc.)
 Brackbill, J. U. 2016, *JCoPh*, **317**, 405
 Brambilla, M. (ed.) 1998, *Kinetic Theory of Plasma Waves: Homogeneous Plasmas* (Oxford: Clarendon Press)
 Broderick, A. E., Chang, P., & Pfrommer, C. 2012, *ApJ*, **752**, 22
 Buneman, O. 1959, *PhRv*, **115**, 503
 Buneman, O., Nishikawa, K.-I., & Neubert, T. 1993, in *Plasma Physics and Controlled Nuclear Fusion (ITC-4)*, ed. H. T. D. Guyenne & J. J. Hunt (ESA), 285
 Chen, G., Chacón, L., & Barnes, D. C. 2011, *JCoPh*, **230**, 7018
 Courant, R., Friedrichs, K., & Lewy, H. 1967, *IBMJ*, **11**, 215
 Dawson, J. 1962, *PhFl*, **5**, 445
 Eastwood, J. W. 1991, *CoPhC*, **64**, 252
 Esirkepov, T. Z. 2001, *CoPhC*, **135**, 144
 Haugbølle, T., Frederiksen, J. T., & Nordlund, A. 2013, *PhPl*, **20**, 062904
 Hockney, R. W., & Eastwood, J. W. 1988, *Computer Simulation Using Particles* (Bristol: Taylor & Francis)
 Langdon, A. B. 1973, *JCoPh*, **12**, 247
 Lapenta, G. 2016, arXiv:1602.06326
 Lapenta, G., & Markidis, S. 2011, *PhPl*, **18**, 072101
 Lewis, H. R. 1970, *JCoPh*, **6**, 136
 Lipatov, A. S. 2002, *The Hybrid Multiscale Simulation Technology: An Introduction with Application to Astrophysical and Laboratory Plasmas* (Berlin: Springer)

- Lyutikov, M., Sironi, L., Komissarov, S., & Porth, O. 2016, arXiv:[1603.05731](#)
- Markidis, S., & Lapenta, G. 2011, [JCoPh](#), **230**, 7037
- McKinstrie, C. J., Giacone, R. E., & Startsev, E. A. 1999, [PhPl](#), **6**, 463
- Miller, K. A., & Stone, J. M. 1999, [BAAS](#), **31**, 1510
- Park, J., Caprioli, D., & Spitkovsky, A. 2015, [PhRvL](#), **114**, 085003
- Philippov, A. A., Spitkovsky, A., & Cerutti, B. 2015, [ApJL](#), **801**, L19
- Riquelme, M. A., & Spitkovsky, A. 2009, [ApJ](#), **694**, 626
- Rossmannith, J. A., & Seal, D. C. 2011, [JCoPh](#), **230**, 6203
- Sironi, L., & Giannios, D. 2014, [ApJ](#), **787**, 49
- Spitkovsky, A. 2005, in AIP Conf. Ser. 801, Astrophysical Sources of High Energy Particles and Radiation, ed. T. Bulik, B. Rudak, & G. Madejski (Melville, NY: AIP), [345](#)
- Spitkovsky, A. 2008, [ApJL](#), **682**, L5
- Umeda, T., Omura, Y., Tominaga, T., & Matsumoto, H. 2003, [CoPhC](#), **156**, 73
- Villasenor, J., & Buneman, O. 1992, [CoPhC](#), **69**, 306
- Zhu, Q., Hernquist, L., & Li, Y. 2015, [ApJ](#), **800**, 6