

Overcoming timestep limitations in boosted-frame particle-in-cell simulations of plasma-based acceleration

Olga Shapoval ^{*}, Remi Lehe [†], Maxence Thévenet [‡], Edoardo Zoni , Yinjian Zhao , and Jean-Luc Vay [‡]
Lawrence Berkeley National Laboratory, Berkeley, California 94720, USA



(Received 20 May 2021; accepted 28 October 2021; published 29 November 2021)

Explicit electromagnetic particle-in-cell (PIC) codes are typically limited by the Courant-Friedrichs-Lewy (CFL) condition, which implies that the timestep multiplied by the speed of light must be smaller than the smallest cell size. In the case of boosted-frame PIC simulations of plasma-based acceleration, this limitation can be a major hindrance, as the cells are often very elongated along the longitudinal direction and the timestep is thus limited by the small, transverse cell size. This entails many small-timestep PIC iterations and can limit the potential speed-up of the boosted-frame technique. Here, by using a CFL-free analytical spectral solver, and by mitigating additional numerical instabilities that arise at large timestep, we show that it is possible to overcome traditional limitations on the timestep and thereby realize the full potential of the boosted-frame technique over a much wider range of parameters.

DOI: [10.1103/PhysRevE.104.055311](https://doi.org/10.1103/PhysRevE.104.055311)

I. INTRODUCTION

Particle-in-cell (PIC) simulations [1,2] are key to the development of plasma-based accelerators and of their potential future applications [3]. However, these simulations can typically be very computationally expensive. One way to reduce their computational cost is to use the boosted-frame technique [4], whereby the simulation is performed in a Lorentz frame moving relativistically in the same direction as the beam or laser driver. The boosted-frame technique is nowadays routinely used in simulations of plasma-based accelerators, and can speed up simulations by several orders of magnitude. To a large extent, this was made possible by the development of a number of algorithms that mitigate the numerical Cherenkov instability (NCI) [5–20]—a numerical instability that would otherwise rapidly grow in the boosted frame and irremediably corrupt the simulated physics.

The remarkable speedup afforded by the boosted-frame technique is due largely to the possibility of increasing the timestep in the boosted frame, and thereby reducing the number of PIC iterations to be performed compared to a corresponding laboratory-frame simulation. For example, in a typical laboratory-frame simulation of laser-wakefield acceleration, both the longitudinal cell size Δz_{lab} and timestep Δt_{lab} are constrained to resolve the small laser oscillations at wavelength λ_{lab} : $\Delta z_{\text{lab}} \ll \lambda_{\text{lab}}$ for a laser propagating along z , with $c\Delta t_{\text{lab}} \leq \Delta z_{\text{lab}} \ll \lambda_{\text{lab}}$ (while the transverse cell size is usually much larger: $\Delta x, \Delta y \gg \Delta z_{\text{lab}}$). By contrast, in a Lorentz boosted frame drifting along z at relativistic velocity with a Lorentz factor $\gamma_b \gg 1$, the laser oscillations are dilated by a factor of approximately $2\gamma_b$ ($\lambda \approx 2\gamma_b\lambda_{\text{lab}}$), which

greatly relaxes the constraints on the longitudinal cell size and timestep: $\Delta z \ll 2\gamma_b\lambda_{\text{lab}}$, with $c\Delta t \leq \Delta z \ll 2\gamma_b\lambda_{\text{lab}}$ (where the quantities Δz and Δt denote the longitudinal cell size and timestep *in the boosted frame*).

However, for large γ_b , as the constraints imposed by the laser are relaxed, the timestep often becomes constrained instead by the transverse cell size: $c\Delta t \leq \Delta x, \Delta y$. (Note that the transverse cell size is left unchanged in the boosted-frame simulation as compared to the corresponding laboratory-frame simulation, since transverse physical length scales are unchanged by the Lorentz transform). In the case of finite-difference time-domain (FDTD) PIC algorithms, this constraint on the timestep is due to the Courant-Friedrichs-Lewy (CFL) condition [21,22]. Similarly, the pseudo-spectral time-domain (PSTD) PIC algorithm [23] also has a CFL condition. As a consequence of the CFL condition, the timestep of the boosted-frame simulation is relatively small and limits the potential computational speedup, even though the physics at stake does not necessarily require such a high temporal resolution.

However, unlike FDTD and PSTD PIC algorithms, pseudo-spectral analytical time-domain (PSATD) PIC algorithms [24,25], which integrate analytically Maxwell's equations over one timestep in Fourier space, do not have a similar CFL condition. It follows that boosted-frame PIC simulations that use the PSATD Maxwell solver could use a larger timestep, as it is thus not explicitly constrained by the transverse resolution. However, it turns out that PSATD boosted-frame simulations are empirically unstable for $c\Delta t > \Delta x, \Delta y$. More specifically, the Galilean PSATD algorithm [16–18], which does efficiently mitigate the NCI for $c\Delta t < \Delta x, \Delta y$, does not seem to suppress the NCI anymore for $c\Delta t > \Delta x, \Delta y$.

This paper examines the nature of this resurgent NCI and shows that this instability can be strongly mitigated with a new algorithm, referred to as the *averaged Galilean PSATD*, whereby a key feature of the PSATD algorithm is

^{*}oshapoval@lbl.gov

[†]rlehe@lbl.gov

[‡]Present address: DESY, 22607 Hamburg, Germany.

exploited to analytically average the electromagnetic fields in time before gathering them onto the macroparticles. Hence, with this new algorithm, simulations can run with a large timestep ($c\Delta t \leq \Delta z \gg \Delta x, \Delta y$) and exhibit the corresponding computational speedup, while preserving the integrity of the simulated physics. While this development was motivated here with the example of laser-wakefield acceleration, it is generally applicable to any simulation where the physics imposes a high transverse spatial resolution but does not impose such strong constraints on the timestep, so that it would be advantageous to use a large timestep compared to the cell size. For instance, this also includes the simulations of low-emittance pencil-like beams [26], in which the space charge requires a high transverse resolution, but has a relatively slow time evolution.

The remainder of the paper is structured as follows. We first examine in more detail the NCI that occurs for large timesteps in the case of the standard Galilean PSATD algorithm. Based on this analysis, we introduce the averaged Galilean PSATD algorithm and describe the corresponding modified PIC loop. We then demonstrate the stability of this new algorithm with large timesteps, first for a uniform plasma, and then for 2D simulations of laser-wakefield acceleration (LWFA) and 3D simulations of plasma wakefield acceleration (PWFA).

II. LIMITATIONS OF THE STANDARD GALILEAN PSATD ALGORITHM FOR LARGE Timesteps

As mentioned in the Introduction, boosted-frame simulations with the Galilean PSATD algorithm are typically unstable when using a large timestep $c\Delta t > \Delta x, \Delta y$. Here, we illustrate this by examining the theoretical NCI growth rate of the Galilean PSATD algorithm for two-dimensional simulations of a uniform plasma drifting at a relativistic velocity $\mathbf{v}_0 = v_0 \mathbf{u}_z$ (where \mathbf{u}_z is the unit vector along the z axis). As a reminder, the Galilean PSATD algorithm solves the Maxwell equations on a moving grid, which drifts at a velocity $\mathbf{v}_{\text{gal}} = v_{\text{gal}} \mathbf{u}_z$. This algorithm was shown to suppress the NCI when $v_{\text{gal}} = v_0$ [16].

In this section, we in fact consider two cases: that of a matching velocity $v_{\text{gal}} = v_0$ and that of a slightly detuned velocity $v_{\text{gal}} = 0.99 v_0$. Conceptually, these two cases represent—at a simplified level—different areas of the simulation box, in the case of a realistic LWFA simulation. More specifically, the matched case ($v_{\text{gal}} = v_0$) represents the background, quiescent plasma, far from the driver and the wakefield, since the Galilean velocity is typically chosen to match its velocity (i.e., $v_{\text{gal}} = -\sqrt{1 - 1/\gamma_b^2} c$). However, the case of the detuned velocity represents the perturbed plasma around the laser driver and in the wakefield, where the local velocity is different than that of the background plasma, and thus different than the Galilean velocity.

In both of these cases, we choose $\Delta z \gg \Delta x$. This is typical for boosted-frame simulations with a large γ_b , since the longitudinally Lorentz-dilated driver and wakefield relax the requirement on the longitudinal resolution. We then further consider two cases: that of a small timestep $c\Delta t = \Delta x$ and that of a large timestep $c\Delta t = \Delta z$. Note that the latter case would not be allowed by the CFL condition of an FDTD algorithm.

Figure 1 displays the theoretical NCI growth rate for the four possible combinations (i.e., small and large timestep, matched and detuned Galilean velocity). The growth rates are obtained by solving the theoretical dispersion relation, namely, Eq. (19) in Ref. [16]. To guide the interpretation of this figure, we also plot the position of well-known NCI resonant modes [8], which are caused by temporal and spatial aliasing. For the Galilean PSATD algorithm, the equation of these aliased resonant modes is given by

$$k_{x,\text{res}} = \sqrt{\left(k_z \frac{v_0}{c} + m_z \frac{2\pi}{\Delta z} \frac{(v_0 - v_{\text{gal}})}{c} - \frac{2\pi n}{c\Delta t} \right)^2 - k_z^2}, \quad (1)$$

for any $m_z, n \in \mathbb{Z}$, where m_z is the spatial alias index and n is the temporal alias index [18]. As one can observe, if $v_{\text{gal}} \approx v_0$, the term proportional to m_z almost cancels and the position of these lines mainly depends on the time aliasing n .

As can be seen in Figs. 1(a) and 1(b), in the matched-velocity case ($v_{\text{gal}} = v_0$), the Galilean PSATD algorithm suppresses the NCI, both for $c\Delta t = \Delta x$ (upper left panel) and $c\Delta t \gg \Delta x$ (upper right panel). By contrast, in the detuned case ($v_{\text{gal}} \neq v_0$, lower panels in Fig. 1), the NCI has a more noticeable growth rate. This growth rate is relatively small for $c\Delta t = \Delta x$, but is much larger for $c\Delta t \gg \Delta x$. In practice, this implies that the Galilean algorithm is relatively robust to velocity perturbations (e.g., in the wakefield) in the case of a small timestep ($c\Delta t \leq \Delta x$), but it is much less robust to those perturbations in the case of a large timestep ($c\Delta t \gg \Delta x$). This explains the empirical observation, mentioned in the Introduction, that boosted-frame simulations of LWFA are typically unstable with large timestep.

Furthermore, in the case $v_{\text{gal}} \neq v_0$ and $c\Delta t \gg \Delta x$ [Figs. 1(c) and 1(d)], we see that the large NCI growth rate is concentrated near time-aliased resonances. Thus, the NCI arises here primarily from the resonant interaction of particles with electromagnetic modes that are not resolved in time. More specifically, the electromagnetic modes are in principle oscillating in time (as predicted by the analytical formulas used in the derivation of the PSATD algorithm), and in most cases their net effect on particles averages to zero. Yet because these fields are sampled with a (large) discrete timestep—which can be comparable to the period of their oscillations—particles may in certain cases see an almost constant value (because of aliasing) instead of an oscillating one, and thus experience a lasting, resonant effect.

III. AVERAGED GALILEAN PSATD ALGORITHM

The above analysis suggests a natural remedy: when pushing the particles, instead of using the value of the fields sampled at a specific time $t = n\Delta t$ for some integer n , the particles should instead be pushed with the fields *averaged in time* between $t = (n - 1/2)\Delta t$ and $t = (n + 1/2)\Delta t$. Averaging the fields over one timestep will barely affect the physics, provided that it is well-resolved in time with the chosen timestep. However, this average will damp the under-resolved modes that are spuriously resonant in Fig. 1. Since this is a *temporal* average, not a *spatial* one, it will not affect the above-mentioned fine spatial details that typically impose

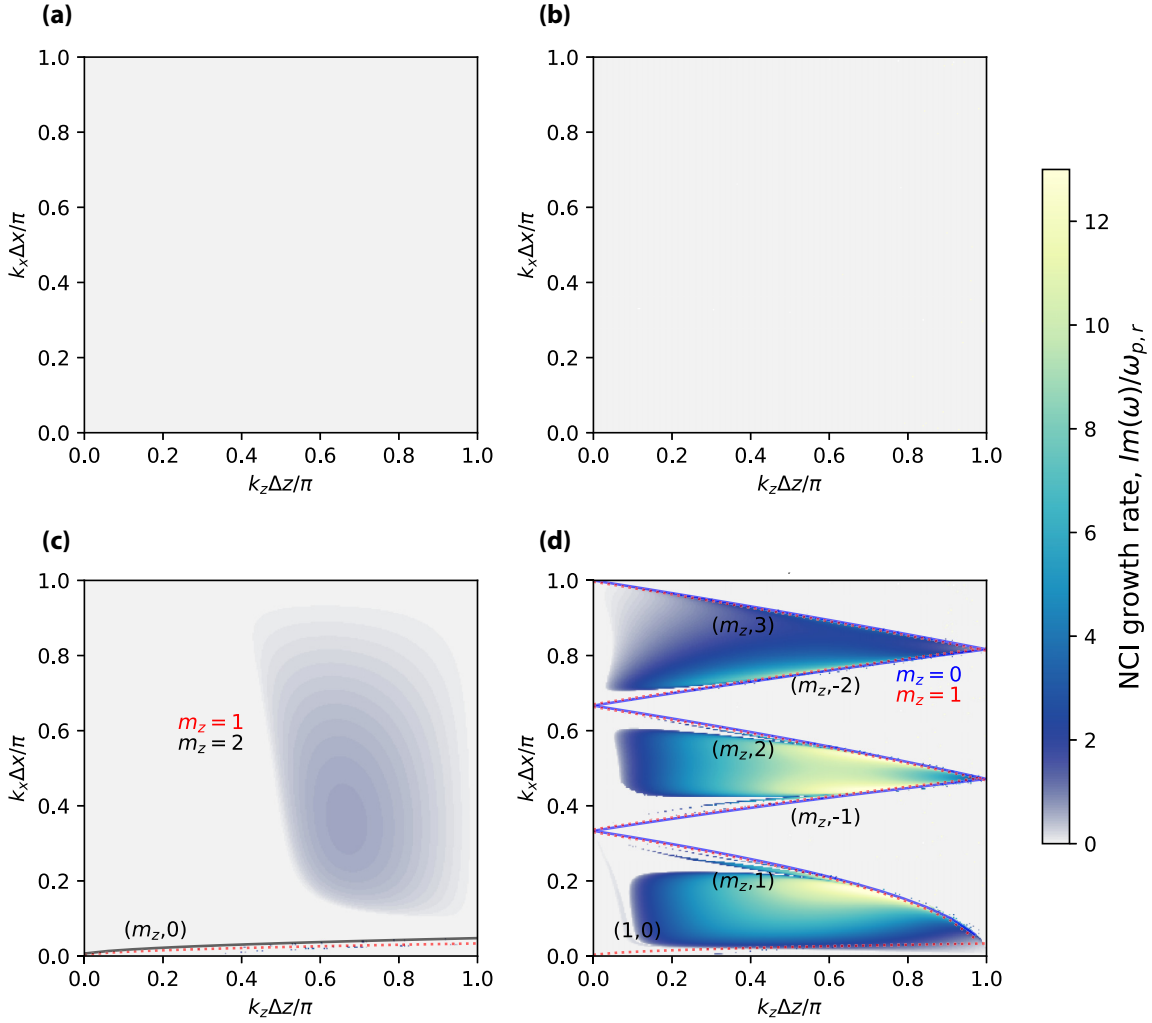


FIG. 1. NCI growth rate of the Galilean PSATD scheme. Normalized NCI growth rate $\text{Im}(\omega)/\omega_{p,r}$ in spectral (k_x, k_z) space, calculated from the analytical stability analysis for different Galilean velocities $v_{\text{gal}} = v_0$ (a), (b) and $v_{\text{gal}} = 0.99 v_0$ (c), (d), and for different timesteps $c\Delta t = \Delta x$ (left) and $c\Delta t = 6\Delta x$ (right). The solid and dotted lines correspond to well-known aliased NCI resonant modes, with alias number (m_z, n) , as given by Eq. (1). In this simulation, a uniform plasma drifts at a velocity $v_0 = c(1 - 1/\gamma_b^2)^{1/2}$ with $\gamma_b = 130$, and the transverse and longitudinal cell sizes are $\Delta x = 6.4 \times 10^{-2} k_{p,r}^{-1}$ and $\Delta z = 6\Delta x$, respectively (where $k_{p,r}^2 = n_0 e^2 / (m_e \epsilon_0 c^2 \gamma_0)$, and where n_0 is the plasma density).

a high transverse resolution $\Delta x, \Delta y$ (for example, the space-charge field of a low-emittance beam), as long as they vary slowly in time (in comparison to the timestep used in the simulation).

We note that with most Maxwell solvers (for example, the FDTD and PSTD algorithms), the time evolution of the electromagnetic fields *within* one timestep is in general not known. However, this evolution is indeed known in the case of the PSATD algorithm. More specifically, as part of the derivation of the PSATD algorithm [16,24,25], the time evolution of the \mathbf{E} and \mathbf{B} fields in Fourier space is calculated analytically. Here, we propose to average this analytical expression over one timestep in Fourier space [see Eqs. (A9) and (A10) in Appendix A], and then to transform these averaged fields $\langle \mathbf{E} \rangle$ and $\langle \mathbf{B} \rangle$ to real space, where they are gathered onto the macroparticles and then discarded. (However, the *unaveraged* \mathbf{E} and \mathbf{B} fields are still kept in memory, and are updated by the standard Galilean PSATD equations [16,17] at each PIC

iteration). The corresponding modified PIC loop is illustrated in Fig. 2 and described in more detail in Appendix B.

In the rest of this article, we refer to this new scheme as the *averaged Galilean PSATD algorithm*, since it combines the Galilean PSATD scheme [16–18] and the temporal average of the fields over one timestep. By construction, the averaged Galilean PSATD algorithm inherits the main advantages of the Galilean PSATD scheme: it has a low amount of spurious numerical dispersion (for high-order spatial derivatives [27,28]) and does not have a CFL limit. In addition, as shown in the next sections, the averaged Galilean PSATD algorithm efficiently mitigates the NCI for large timesteps.

Stability analysis for a uniform plasma drifting at relativistic velocity

To analyze the stability of the new averaged Galilean PSATD algorithm, we consider again the case of a

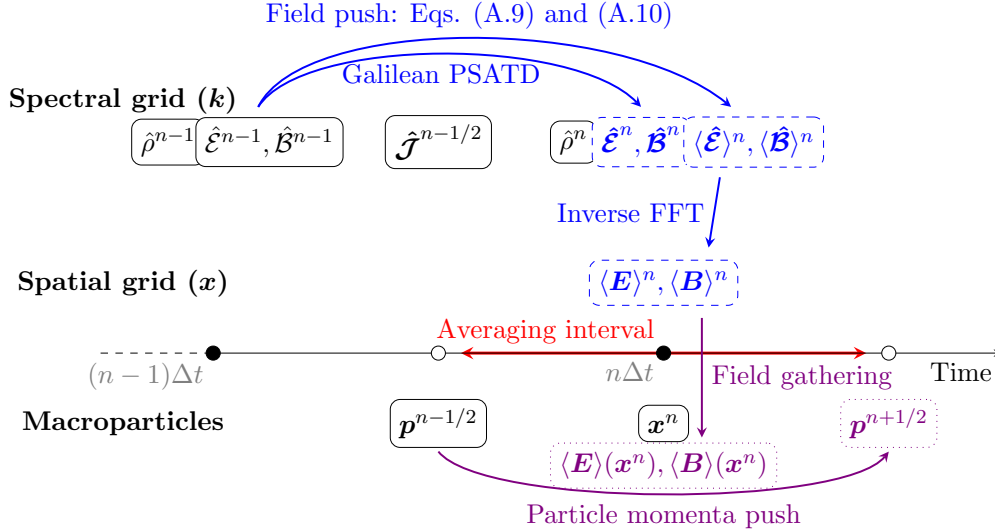


FIG. 2. Illustration of the field push and particle momenta push, in the averaged Galilean PSATD algorithm. The quantities represented in solid frames are the ones that are known just before the field and particle push. Note that this includes the deposited charge and current $\hat{\rho}$ and $\hat{\mathcal{J}}$. The quantities in dashed and dotted frames are the ones that are being computed during the field and particle push. As part of the field push, the regular fields $\hat{\mathcal{E}}$ and $\hat{\mathcal{B}}$ are updated, and the averaged fields $\langle \mathcal{E} \rangle$ and $\langle \mathcal{B} \rangle$ are calculated and transformed to the spatial grid. As part of the particle push, the averaged fields $\langle E \rangle$ and $\langle B \rangle$ are gathered onto the macroparticles, to update the macroparticles' momenta. The rest of the PIC cycle (e.g., charge and current deposition, particle position push) is not shown here but is identical to the standard Galilean PSATD [16,17]. See the appendices for more details on the PIC loop and exact definitions of the notation.

two-dimensional uniform plasma drifting with a relativistic velocity. We derived the theoretical dispersion equation for this system, by using a similar method as for the standard Galilean PSATD algorithm [16]—while taking into account the additional average in time. The full derivation of this theoretical dispersion equation is given in the Appendix C. By solving this dispersion equation numerically, we can extract the NCI growth rate $\text{Im}(\omega)/\omega_{p,r}$ as a function of \mathbf{k} . In addition, we also performed actual PIC simulations for the same system. We used the PIC code WarpX [29,30], in which we implemented the averaged Galilean PSATD algorithm, and we then extracted the NCI growth rate of the NCI in postprocessing.

The growth rates extracted from both the WarpX simulations and the theoretical dispersion equation are shown in Fig. 3—both for standard Galilean PSATD (left panels) and for the averaged Galilean PSATD (right panels). For this case, we used the same parameters as for the lower-right panel of Fig. 1, i.e., $\gamma_b = 130$, $\Delta x = 6.4 \times 10^{-2} k_{p,r}^{-1}$, a large timestep $c\Delta t = \Delta z = 6\Delta x$, and a detuned velocity $v_{\text{gal}} = 0.99 v_0$. (Recall from the previous sections that the case of a detuned velocity is the one for which using a large timestep presents a major issue.)

As can be seen in Fig. 3, the theoretical predictions (upper panels) and simulation results (lower panels) are in good agreement, which confirms that the theoretical dispersion equation correctly captures the nature of the instability. (Note that growth rate measured from simulations is typically noisy, which limits the comparison). More importantly, both the theoretical predictions and simulation results show that averaged Galilean PSATD (right panels) strongly reduces the growth of the instability compared to the standard Galilean PSATD (left panels). This confirms that averaging the fields in time inhibits spurious resonances with under-resolved electromagnetic

modes, and thereby enables stable simulations with large timesteps.

IV. APPLICATION TO SIMULATIONS OF PLASMA WAKEFIELD IN A LORENTZ-BOOSTED FRAME

A. 2D LWFA

This section illustrates that the stability properties observed in the case of a uniform plasma also apply to realistic simulations of plasma-based acceleration. To this end, we first perform two-dimensional (2D) simulations of a laser-wakefield accelerator. In these simulations, an x -polarized Gaussian laser pulse with amplitude $a_0 = 1$, duration $\tau = 20$ fs and waist $w_0 = 15 \mu\text{m}$ propagates in a matched parabolic plasma channel with a background density of $1.0 \times 10^{18} \text{cm}^{-3}$. The simulation runs in a Lorentz-boosted frame ($\gamma_b = 30$) with a nodal PSATD solver with finite order 16 [18,27,28]. The longitudinal resolution (in the boosted frame) is set to $\Delta z = 2\gamma_b \lambda_{\text{lab}}/32 = 1.52 \mu\text{m}$, while the transverse resolution is $\Delta x = 0.15 \mu\text{m}$, so that $\Delta z = 10\Delta x$. We run the simulation with the standard and averaged Galilean PSATD, and with a small timestep ($c\Delta t = \Delta x$) as well as large timesteps ($c\Delta t = 5\Delta x$ and $c\Delta t = 10\Delta x$). Figure 4 displays snapshots of the longitudinal electric field E_z and of the longitudinal current density J_z in the boosted frame, for these different cases. In these colormaps, the rapid oscillations of E_z for $z > 0$ mm correspond to the longitudinal component of the laser field, which undergoes significant nonlinear evolution and red-shifting, while the slow oscillations of E_z and J_z for $z < 0$ mm correspond to the plasma wakefield.

As expected, the standard Galilean PSATD is stable for a small timestep [Fig. 4(a)], but unstable for large timesteps [Figs. 4(c) and 4(e)]. More specifically, in panels Figs. 4(c) and 4(e), spurious oscillations rapidly grow in the wakefield

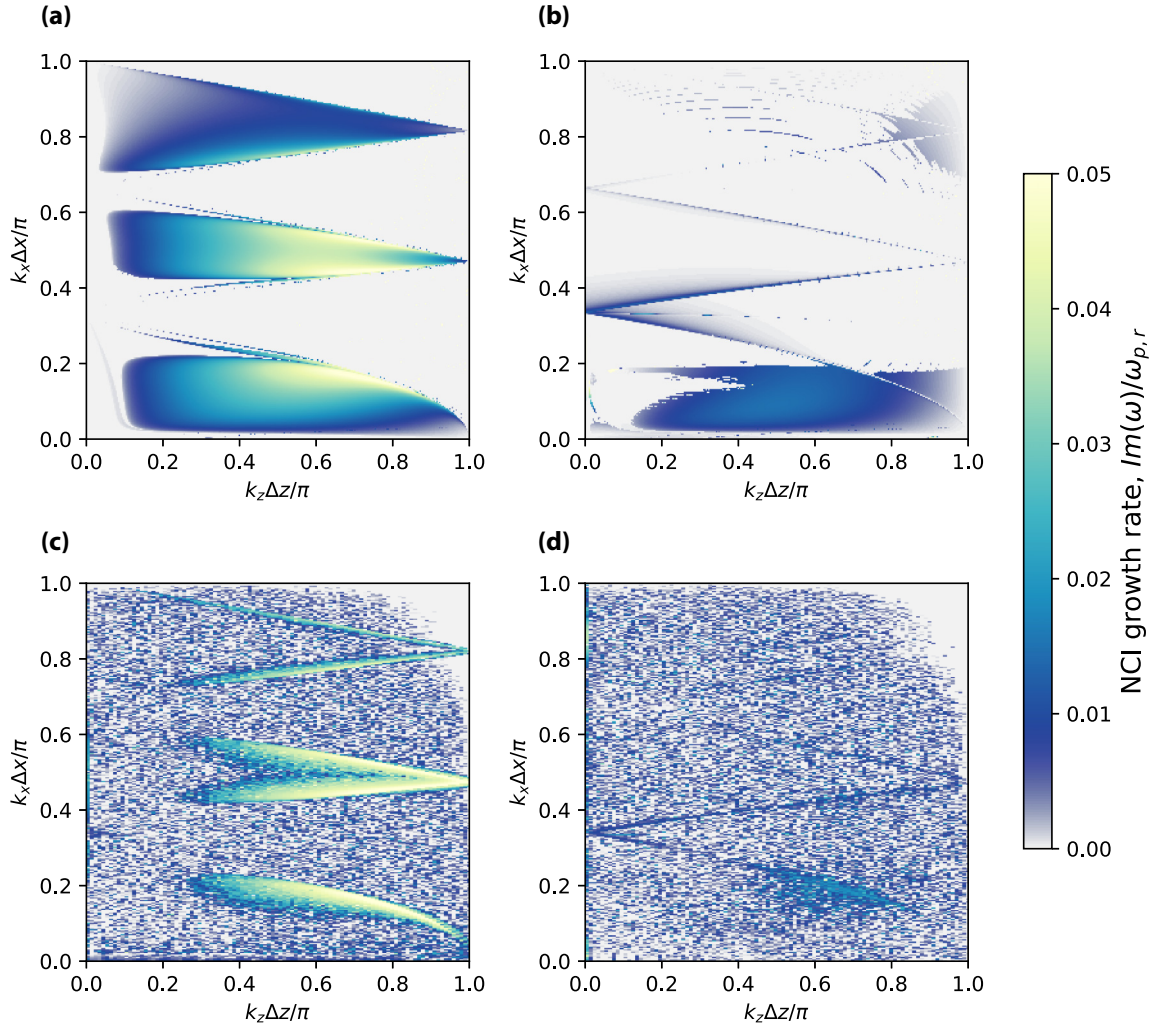


FIG. 3. NCI growth rate: Galilean PSATD vs. averaged Galilean PSATD schemes. Normalized NCI growth rate $\text{Im}(\omega)/\omega_{p,r}$ in spectral (k_x, k_z) space, calculated from the analytical stability analysis (a), (b) and from WarpX simulation results (c), (d), obtained using the Galilean PSATD (a), (c) and averaged Galilean PSATD (b), (d) schemes, at infinite spectral order, with large timestep $c\Delta t = \Delta z = 6\Delta x$ and slightly detuned Galilean velocity $v_{\text{gal}} = 0.99 v_0$.

and severely disrupt its structure. We also note that the results of the averaged Galilean PSATD with a small timestep [Fig. 4(b)] are almost indistinguishable from those of the standard Galilean PSATD [Fig. 4(a)]—thereby confirming that averaging the fields in time preserves the essential physics. More importantly, for large timestep [Figs. 4(d) and 4(f)], the averaged Galilean PSATD achieves stability while preserving the overall structure of the wakefield. Indeed for $c\Delta t = 5\Delta x$ [Fig. 4(d)] the E_z field is still almost indistinguishable from that of Fig. 4(a). For $c\Delta t = 10\Delta x$ [Fig. 4(f)], small differences become noticeable, especially in the red-shifted laser oscillations—although they hardly affect the structure of the accelerating wakefield. This may indicate that this value of Δt starts to reach the limit for which the simulation is not well-resolved in time anymore.

We also note that, both for $c\Delta t = 5\Delta x$ [Fig. 4(d)] and $c\Delta t = 10\Delta x$ [Fig. 4(f)], small transverse oscillations become noticeable in J_z for $z < -1.0$ mm. These oscillations may be due to the remaining nonzero growth rate of the averaged Galilean algorithm [see the growth rates represented

in Figs. 3(b) and 3(d)]. However, their magnitude is small enough that they do not lead to a modulation of the electric field, hence they do not affect the dynamics. Again, this represents a clear improvement compared to the standard Galilean algorithm [Figs. 4(c) and 4(e)].

B. 3D LWFA

To show that the advantageous stability properties of the averaged Galilean PSATD generalize to 3D, we ran a 3D LWFA simulation using the same physical and numerical parameters as in Sec. IV A. We note that, in this case, we had to use domain decomposition to be able to run the simulation. This setup was run with the standard Galilean PSATD and the averaged Galilean PSATD, both with a large timestep $c\Delta t = 10\Delta z$. Figure 5 displays corresponding snapshots. As can be seen, the standard Galilean PSATD algorithm is unstable for this large timestep, while the averaged Galilean PSATD algorithm remains stable. This is again consistent with the

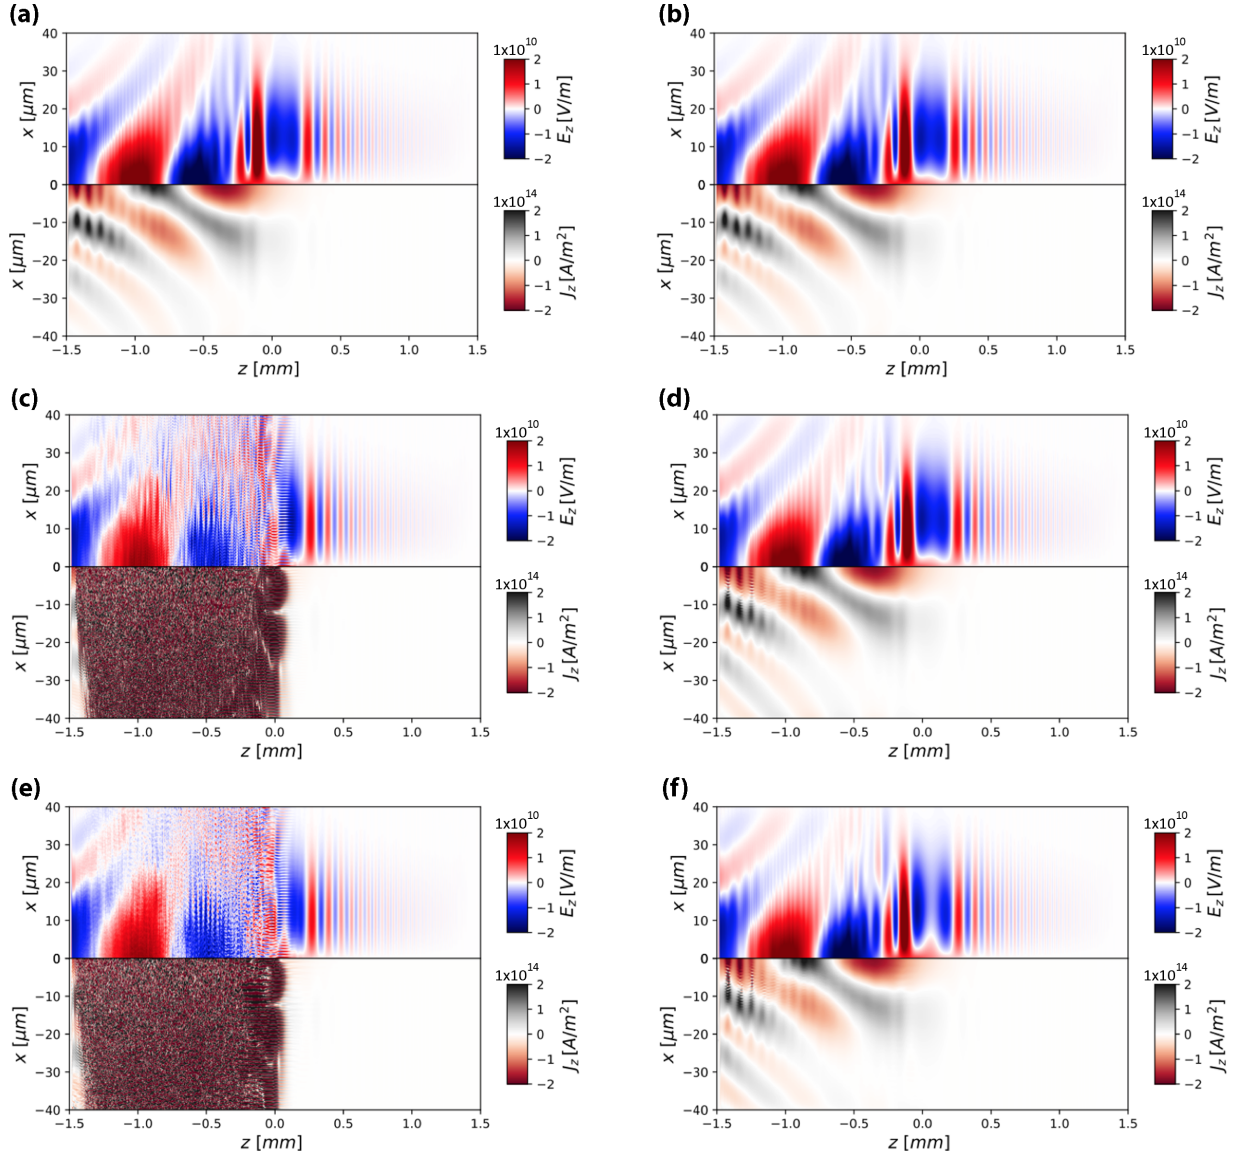


FIG. 4. Instability mitigation in a 2D laser-wakefield simulation with large timestep. The upper half of each subplot shows the E_z field seen by the macroparticles (i.e., the regular E_z field for the standard Galilean PSATD, and the averaged (E_z) field for the averaged Galilean PSATD). The lower half of each subplot shows the longitudinal current density J_z of the plasma. All the fields are shown in the boosted frame. The different subplots correspond to the standard Galilean PSATD (a), (c), (e) and averaged Galilean PSATD (b), (d), (f), with $c\Delta t = 0.1\Delta z = \Delta x$ (a), (b), $c\Delta t = 0.5\Delta z = 5\Delta x$ (c), (d), and $c\Delta t = \Delta z = 10\Delta x$ (e), (f).

2D results and confirms the benefits of the proposed averaged algorithm.

C. 3D PWFA

We now go beyond the stability analysis, and show that the algorithm can speed up a full-scale plasma-wakefield simulation while preserving the physical results of interest. Here we focus on a 3D simulation of beam-driven wakefield acceleration. In this simulation, a 1 nC Gaussian electron beam propagates in a plasma with a background density of $1.0 \times 10^{17} \text{ cm}^{-3}$, and experiences a typical evolution whereby the head of the beam erodes while the tail of the beam performs betatron oscillations in the generated wakefield. The electron beam initially has a mean Lorentz factor $\gamma = 2000$,

with a relative RMS spread $\Delta\gamma/\gamma = 0.01$, and a transverse and longitudinal RMS size of 5 and 20 μm , respectively. The simulation is run in a Lorentz-boosted frame ($\gamma_b = 5.6$), with 512^3 cells of size $\Delta x = \Delta y = 0.78 \mu\text{m}$, $\Delta z = 5 \mu\text{m} = 6.4 \Delta x$ (in the boosted frame), and a nodal PSATD solver with finite order 16 [18,27,28]. To verify again that the averaged Galilean PSATD algorithm preserves the physics of interest, we run the simulation both with the standard Galilean PSATD algorithm and a small timestep $c\Delta t = \Delta x$ (fiducial case) and with the averaged Galilean PSATD algorithm and a large timestep $c\Delta t = \Delta z$. In both cases, we ran the WarpX code on the Summit supercomputer, using 24 GPUs with domain decomposition along z .

The top panels of Fig. 6 display colormaps of the wakefield in the laboratory frame, which were reconstructed on-the-fly

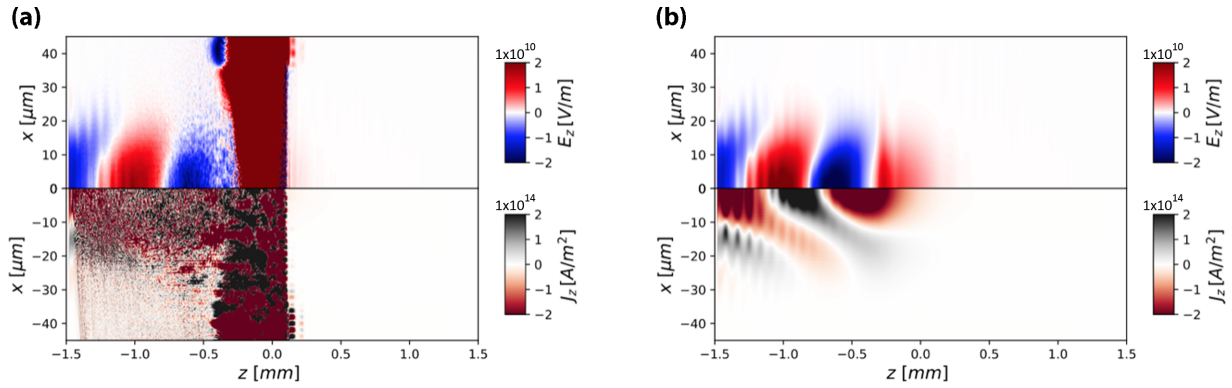


FIG. 5. Instability mitigation in a 3D laser-wakefield simulation with large timestep. The upper half of each subplot shows the E_z field seen by the macroparticles, i.e., the regular E_z field for the standard Galilean PSATD (a), and the averaged $\langle E_z \rangle$ field for the averaged Galilean PSATD (b). The lower half of each subplot shows the longitudinal current density J_z of the plasma. All the fields are shown in the boosted frame with $c\Delta t = \Delta z = 10\Delta x$.

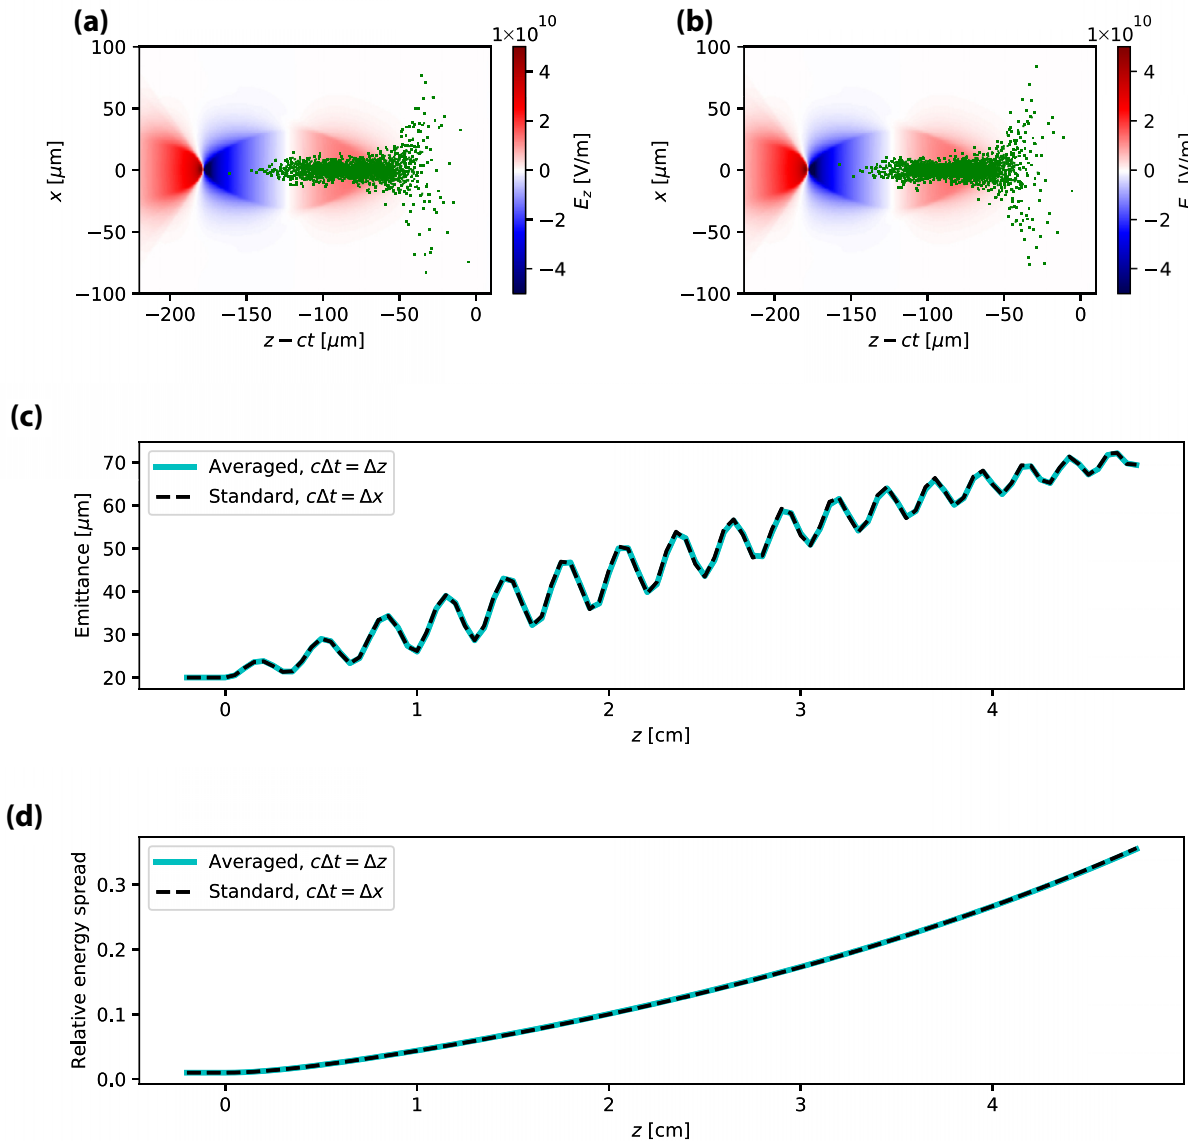


FIG. 6. Algorithm comparison for a 3D plasma-wakefield simulation. (a), (b) Snapshot of the E_z field ($\langle E_z \rangle$) in the case of the averaged Galilean PSATD in the laboratory frame, for the small-timestep ($c\Delta t = \Delta x$) standard Galilean PSATD simulation (a) and large-timestep ($c\Delta t = \Delta z$) averaged Galilean PSATD simulation (b) with fixed $\Delta z = 6.4\Delta x$. The green dots are representative random samples of the macroparticles in the beam driver. (c), (d) Evolution of the emittance and relative energy spread of the beam, in the laboratory frame.

during the boosted-frame simulation. Again, the simulation with the large timestep and the averaged Galilean PSATD is stable [Fig. 6(b)] and the simulated wakefield is almost indistinguishable from that produced from the fiducial small-timestep simulation [Fig. 6(a)]. In addition, Figs. 6(c) and 6(d) show the evolution of the emittance and relative energy spread of the driver beam in the laboratory frame, as it undergoes head erosion and betatron oscillation. This is obtained from laboratory-frame particle data that is reconstructed on-the-fly during the boosted-frame simulation. As can be seen, the evolution of these beam quantities show excellent agreement between the fiducial small-timestep simulation and the large-timestep with the averaged Galilean PSATD.

Thus, in the above example, the averaged Galilean PSATD allowed stable simulations to be run with a large timestep while preserving the integrity of the physics at stake. We note that, as a consequence of the large timestep, the simulation using the averaged Galilean PSATD exhibited a $5\times$ overall speed-up compared to the small-timestep standard Galilean PSATD simulation, on the Summit supercomputer.

V. DISCUSSION AND CONCLUSION

In this paper, we proposed a modified PIC algorithm that enables stable boosted-frame simulations of plasma-based acceleration with a large timestep $c\Delta t \gg \Delta x$, where Δx represents the smallest cell size. This was achieved by using a CFL-free Galilean PSATD solver and by averaging the \mathbf{E} and \mathbf{B} fields in time, to inhibit spurious resonances with under-resolved, aliased electromagnetic modes. We demonstrated this novel scheme in realistic 2D and 3D plasma-wakefield simulations.

We note that the proposed algorithm could certainly be further refined and improved upon in the future. For instance, although the proposed algorithm strongly reduces the NCI growth rate for large timesteps, it does not completely eliminate it. As a consequence, the NCI at large timesteps could still be an issue for certain sets of parameters.

In conclusion, this work demonstrates that it is possible to run boosted-frame simulations with a much larger timestep than the traditional CFL limit, while still accurately capturing the physics. This new development enables potential speedups of an order of magnitude or more, opening up a new area of investigation within the field of first-principles, Particle-In-Cell modeling of plasma-wakefield particle accelerators, whereby the simulation timestep is chosen much more freely than before.

Although the present work is focused on simulations of plasma accelerators and on a particular method (the averaged Galilean PSATD), it could have a wider impact. For instance, even though the algorithm proposed here builds specifically upon the PSATD framework, the central idea (namely, averaging the fields in time) is fairly general and could thus also guide the future development of similar solutions for FDTD-based methods. In addition, beyond the plasma accelerator community, this work may be of interest to the modeling of advanced light sources concepts, coherent synchrotron radiation in particle accelerators, astrophysical shocks or beam-plasma instabilities of astrophysical relativistic jets, which can also utilize the boosted-frame PIC method

for accurate modeling from first principles. We also envision that the method that is used in this paper can be employed to overcome similar timestep limitations in PIC simulations that do not employ the boosted-frame technique, with impact to a much wider range of applications.

ACKNOWLEDGMENTS

The authors acknowledge the discussions and code contributions from the whole WarpX development team. We are also thankful to Andrew Myers for valuable feedback on the manuscript. This research was supported by the Exascale Computing Project (Grant No. 17-SC-20-SC), a collaborative effort of the U.S Department of Energy Office of Science and the National Nuclear Security Administration. This work was performed in part under the auspices of the U.S. Department of Energy by Lawrence Berkeley National Laboratory under Contract No. DE-AC02-05CH11231. This research used resources of the Oak Ridge Leadership Computing Facility at the Oak Ridge National Laboratory, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC05-00OR22725.

APPENDIX A: DERIVATION OF THE AVERAGED FIELDS

Throughout the paper, we use the notation $\langle \hat{\mathcal{F}} \rangle^n$ to refer to the averaged field within the time interval $[(n - \frac{1}{2})\Delta t, (n + \frac{1}{2})\Delta t]$, defined as

$$\langle \hat{\mathcal{F}} \rangle^n = \frac{1}{\Delta t} \int_{(n-1/2)\Delta t}^{(n+1/2)\Delta t} \hat{\mathcal{F}}(\mathbf{k}, \tau) d\tau, \quad (\text{A1})$$

where $\hat{\mathcal{F}} = \hat{\mathcal{F}}(\mathbf{k}, t) = \int \mathbf{F}(\mathbf{x}, t) e^{-i\mathbf{k}\cdot\mathbf{x}} d^3\mathbf{x}$ refers to the Fourier transform of the $\mathbf{F}(\mathbf{x}, t)$ field.

In the Galilean coordinates drifting at \mathbf{v}_{gal} , the p -order discretized Maxwell equations transformed to Fourier space read [18]

$$\left(\frac{\partial}{\partial t} - i[\mathbf{k}] \cdot \mathbf{v}_{\text{gal}} \right)^2 \hat{\mathcal{B}} + c^2 [k]^2 \hat{\mathcal{B}} = \frac{1}{\epsilon_0} i[\mathbf{k}] \times \hat{\mathcal{J}}, \quad (\text{A2})$$

$$\begin{aligned} & \left(\frac{\partial}{\partial t} - i[\mathbf{k}] \cdot \mathbf{v}_{\text{gal}} \right)^2 \hat{\mathcal{E}} + c^2 [k]^2 \hat{\mathcal{E}} \\ & = -\frac{c^2}{\epsilon_0} \hat{\rho} i[\mathbf{k}] - \frac{1}{\epsilon_0} \left(\frac{\partial}{\partial t} - i[\mathbf{k}] \cdot \mathbf{v}_{\text{gal}} \right) \hat{\mathcal{J}}, \end{aligned} \quad (\text{A3})$$

where $[k] = \sqrt{[\mathbf{k}]^2} = \sqrt{[k_x]^2 + [k_y]^2 + [k_z]^2}$ and where $[k_u]$ with $u = \{x, y, z\}$ is the Fourier transform of the p -order discretized stencil $\hat{\nabla}_u$ [i.e., such that a p -order Taylor expansion yields $\hat{\nabla}_u F = \partial_u F + O(\Delta u^p)$] [27,28].

As explained in Ref. [16], these equations can be integrated analytically under the assumption that the current $\hat{\mathcal{J}}$ is constant over one timestep, and that the fields $\hat{\mathcal{E}}, \hat{\mathcal{B}}, \hat{\mathcal{J}}$, and $\hat{\rho}$ satisfy the conservation equations. More specifically, assuming that the \mathbf{E} and \mathbf{B} fields are known at $t = (n - 1)\Delta t$, and under the assumption that $\hat{\mathcal{J}}(\mathbf{k}, t)$ is constant and equal to $\hat{\mathcal{J}}^{n-1/2}(\mathbf{k})$ over the time interval $[(n - 1)\Delta t, n\Delta t]$, we can obtain the expressions of $\hat{\mathcal{E}}(\mathbf{k}, t), \hat{\mathcal{B}}(\mathbf{k}, t)$ as a function

of $\hat{\mathcal{E}}^{n-1}$, $\hat{\mathcal{B}}^{n-1}$, $\hat{\mathcal{J}}^{n-1/2}$, $\hat{\rho}^{n-1}$, $\hat{\rho}^n$:

$$\begin{aligned} \hat{\mathcal{E}}(\mathbf{k}, t) = & \left\{ \hat{\mathcal{E}}^{n-1} - \frac{\alpha_1}{c^2[k]^2(1-\nu^2)} - \frac{\beta_1}{c^2[k]^2} \right\} \cos\{c[k][t - (n-1)\Delta t]\} e^{i\nu c[k][t - (n-1)\Delta t]} \\ & + \frac{\alpha_1}{c^2[k]^2(1-\nu^2)} + \frac{\beta_1}{c^2[k]^2} e^{i\nu c[k][t - (n-1)\Delta t]} \\ & + \frac{1}{c[k]} \left[c^2 i[\mathbf{k}] \times \hat{\mathcal{B}}^{n-1} - \frac{1}{\varepsilon_0} \hat{\mathcal{J}}^{n-1/2} + i\nu \frac{\alpha_1}{c[k](1-\nu^2)} \right] \sin\{c[k][t - (n-1)\Delta t]\} e^{i\nu c[k][t - (n-1)\Delta t}, \end{aligned} \quad (\text{A4})$$

$$\begin{aligned} \hat{\mathcal{B}}(\mathbf{k}, t) = & \left\{ \hat{\mathcal{B}}^{n-1} - \frac{\alpha_2}{c^2[k]^2(1-\nu^2)} \right\} \cos\{c[k][t - (n-1)\Delta t]\} e^{i\nu c[k][t - (n-1)\Delta t]} + \frac{\alpha_2}{c^2[k]^2(1-\nu^2)} \\ & + \frac{1}{c[k]} \left\{ -i[\mathbf{k}] \times \hat{\mathcal{E}}^{n-1} + i\nu \frac{\alpha_2}{c[k](1-\nu^2)} \right\} \sin\{c[k][t - (n-1)\Delta t]\} e^{i\nu c[k][t - (n-1)\Delta t}, \end{aligned} \quad (\text{A5})$$

where $\nu = [\mathbf{k}] \cdot \mathbf{v}_{\text{gal}}/c[k]$ and

$$\alpha_1 = \frac{i\nu c[k]}{\varepsilon_0} \hat{\mathcal{J}}^{n-1/2} - \frac{c^2 \hat{\rho}^n - \hat{\rho}^{n-1} e^{i\nu c[k]\Delta t}}{\varepsilon_0 (1 - e^{i\nu c[k]\Delta t})} i[\mathbf{k}], \quad (\text{A6})$$

$$\beta_1 = \frac{c^2 \hat{\rho}^n - \hat{\rho}^{n-1}}{\varepsilon_0 (1 - e^{i\nu c[k]\Delta t})} i[\mathbf{k}], \quad (\text{A7})$$

$$\alpha_2 = \frac{1}{\varepsilon_0} i[\mathbf{k}] \times \hat{\mathcal{J}}^{n-1/2}. \quad (\text{A8})$$

Strictly speaking, these expressions of $\hat{\mathcal{E}}(\mathbf{k}, t)$, $\hat{\mathcal{B}}(\mathbf{k}, t)$ in Eqs. (A4) and (A5) are only valid for t in the interval $[(n-1)\Delta t, n\Delta t]$ [because of the assumption $\hat{\mathcal{J}}(\mathbf{k}, t) = \hat{\mathcal{J}}^{n-1/2}(\mathbf{k})$]. However, we assume that they are also approximately valid on the interval $[(n-1/2)\Delta t, (n+1/2)\Delta t]$ over which the fields are averaged (see Fig. 2). This is valid if $\hat{\mathcal{J}}$ varies slowly from one timestep to the next—i.e., if the plasma response is well-resolved in time.

Under these assumptions, we average Eqs. (A4) and (A5) in time as defined in Eq. (A1), and obtain

$$\langle \hat{\mathcal{E}} \rangle^n = \Psi_1 \hat{\mathcal{E}}^{n-1} - ic\Psi_2 \frac{[\mathbf{k}]}{[k]} \times \hat{\mathcal{B}}^{n-1} + \left(\frac{i\nu c[k]}{\varepsilon_0} A_1 + \frac{\Psi_2}{c[k]\varepsilon_0} \right) \hat{\mathcal{J}}^{n-1/2} + C_\rho (A_2 - A_1) \hat{\rho}^n \frac{[\mathbf{k}]}{[k]} + C_\rho (\theta^2 A_1 - A_2) \hat{\rho}^{n-1} \frac{[\mathbf{k}]}{[k]}, \quad (\text{A9})$$

$$\langle \hat{\mathcal{B}} \rangle^n = \Psi_1 \hat{\mathcal{B}}^{n-1} + \frac{i}{c} \Psi_2 \frac{[\mathbf{k}]}{[k]} \times \hat{\mathcal{E}}^{n-1} + \frac{i[\mathbf{k}]}{\varepsilon_0} A_1 \frac{[\mathbf{k}]}{[k]} \times \hat{\mathcal{J}}^{n-1/2}. \quad (\text{A10})$$

Here again, $\nu = [\mathbf{k}] \cdot \mathbf{v}_{\text{gal}}/c[k]$, $\theta = e^{i[\mathbf{k}] \cdot \mathbf{v}_{\text{gal}} \Delta t/2}$, and the other coefficients are given by

$$C_\rho = \frac{ic^2[k]}{\varepsilon_0(1-\theta^2)}, \quad (\text{A11a})$$

$$A_1 = \frac{\Psi_1 - 1 + i\nu\Psi_2}{c^2[k]^2(\nu^2 - 1)}, \quad (\text{A11b})$$

$$A_2 = \frac{\Psi_3 - \Psi_1}{c^2[k]^2}, \quad (\text{A11c})$$

$$\Psi_1 = \theta \frac{(S_1 + i\nu C_1) - \theta^2(S_3 + i\nu C_3)}{c[k]\Delta t(\nu^2 - 1)}, \quad (\text{A11d})$$

$$\Psi_2 = \theta \frac{(C_1 - i\nu S_1) - \theta^2(C_3 - i\nu S_3)}{c[k]\Delta t(\nu^2 - 1)}, \quad (\text{A11e})$$

$$\Psi_3 = \frac{i\theta(1-\theta^2)}{c[k]\Delta t\nu}, \quad (\text{A11f})$$

with $C_m = \cos(m c[k] \Delta t/2)$ and $S_m = \sin(m c[k] \Delta t/2)$ for $m = 1, 2, 3$.

APPENDIX B: PIC CYCLE OVERVIEW

Figure 2 gives an overview of a key part of the PIC loop for the averaged Galilean PSATD algorithm. Here we describe the exact PIC loop in more detail. Assuming that we originally know the particle positions and momenta \mathbf{x}^n at $\mathbf{p}^{n-1/2}$ and the fields \mathbf{E}^{n-1} and \mathbf{B}^{n-1} , the loop consists of the following steps:

(1) Deposit the charge and current densities of the particles onto the spatial grid. In particular, we deposit the charge density ρ^n at time $t = n\Delta t$ from the particle positions \mathbf{x}^n and the current density $\mathbf{J}^{n-1/2}$ at time $t = (n - \frac{1}{2})\Delta t$ from the particle positions \mathbf{x}^{n-1} and \mathbf{x}^n and the particle velocities $\mathbf{v}^{n-1/2}$;

(2) Transform all relevant physical quantities from physical space to Fourier space;

(3) Compute the new electromagnetic fields in Fourier space $\hat{\mathcal{E}}^n$ and $\hat{\mathcal{B}}^n$, from the charge and current densities $\hat{\rho}^{n-1}$ (available from the previous PIC iteration), $\hat{\rho}^n$ and $\hat{\mathcal{J}}^{n-1/2}$;

(4) Compute the averaged electromagnetic fields, $\langle \hat{\mathcal{E}} \rangle^n$ and $\langle \hat{\mathcal{B}} \rangle^n$;

(5) Transform all relevant physical quantities from Fourier space back to physical space;

(6) Gather the averaged electromagnetic fields, in physical space, from the spatial grid onto the particles;

(7) Push the particles by updating their positions and momenta based on the current values of the averaged electromagnetic fields, $\langle \mathbf{E} \rangle^n$ and $\langle \mathbf{B} \rangle^n$, whose precise definition is given in the next section. In particular, the momenta are updated from $\mathbf{p}^{n-1/2}$ to $\mathbf{p}^{n+1/2}$ and the positions are then updated from \mathbf{x}^n to \mathbf{x}^{n+1} .

APPENDIX C: DERIVATION OF THE DISPERSION RELATION FOR THE AVERAGED GALILEAN PSATD ALGORITHM

Here we derive the 2D dispersion relation to study the NCI induced by a relativistic plasma flowing through a periodic grid along the z axis with velocity $\mathbf{v}_0 = v_0 \mathbf{u}_z$, where $v_0 = c(1 - 1/\gamma_b^2)^{1/2}$. This is done by combining the discretized Maxwell equations in the Galilean frame and the discretized Vlasov equation, and assuming small perturbations for the electromagnetic fields \mathbf{E} , \mathbf{B} and the distribution function δf . Because the Vlasov equation involves the averaged fields $\langle \mathbf{E} \rangle$, $\langle \mathbf{B} \rangle$, we include their expression as a function of the regular fields \mathbf{E} , \mathbf{B} in the system of equations. Hence, when expressed in spectral space, the different equations of the system are as follows:

(1) the discretized Maxwell equations at n -th timestep in time-symmetrical form [16]:

$$\theta^* c \hat{\mathbf{B}}^n - \theta c \hat{\mathbf{B}}^{n-1} = -t_{ck} \frac{i\mathbf{k} \times (\theta^* \hat{\mathbf{E}}^n + \theta \hat{\mathbf{E}}^{n-1})}{k} + 2\chi_4' \frac{\hat{T}}{\varepsilon_0 c k} \frac{\mathbf{k} \times \hat{\mathcal{J}}_d^{n-1/2}}{k}, \quad (\text{C1})$$

$$\theta^* \hat{\mathbf{E}}^n - \theta \hat{\mathbf{E}}^{n-1} = t_{ck} \frac{i\mathbf{k} \times (\theta^* c \hat{\mathbf{B}}^n + \theta c \hat{\mathbf{B}}^{n-1})}{k} - 2\chi_4 \frac{\hat{T}}{\varepsilon_0 c k} \left[\hat{\mathcal{J}}_d^{n-1/2} - \frac{(\mathbf{k} \cdot \hat{\mathcal{J}}_d^{n-1/2}) \mathbf{k}}{k^2} \right] - \frac{\hat{T} i\mathbf{k}}{\varepsilon_0 k^2} (\theta^* \hat{\rho}^n - \theta \hat{\rho}^{n-1}); \quad (\text{C2})$$

(2) the perturbed Vlasov equation [16]. Note that, here, we replaced the regular fields $\hat{\mathbf{E}}$, $\hat{\mathbf{B}}$ by the averaged fields $\langle \hat{\mathbf{E}} \rangle$, $\langle \hat{\mathbf{B}} \rangle$ to take into account the changes associated with the averaged Galilean PSATD:

$$\delta \hat{f}^{n+1/2}(\mathbf{k}_m, \mathbf{p}) e^{i\mathbf{k}_m(\mathbf{v} - \mathbf{v}_{\text{gal}})\Delta t/2} - \hat{f}^{n-1/2}(\mathbf{k}_m, \mathbf{p}) e^{-i\mathbf{k}_m(\mathbf{v} - \mathbf{v}_{\text{gal}})\Delta t/2} + q\Delta t \hat{S}(\mathbf{k}_m) [\langle \hat{\mathbf{E}}^n(\mathbf{k}) \rangle + \mathbf{v} \times \langle \hat{\mathbf{B}}^n(\mathbf{k}) \rangle] \cdot \frac{\partial f_0}{\partial \mathbf{p}} = 0; \quad (\text{C3})$$

(3) the expression of the averaged field for $t \in [(n-1/2)\Delta t, (n+1/2)\Delta t]$:

$$\langle \hat{\mathbf{B}}^n(\mathbf{k}, t) \rangle = \Psi_1 \hat{\mathbf{B}}^{n-1} + \frac{i}{c} \Psi_2 \frac{\mathbf{k}}{k} \times \hat{\mathbf{E}}^{n-1} + \frac{ik}{\varepsilon_0} A_1 \frac{\mathbf{k}}{k} \times \hat{\mathcal{J}}^{n-1/2}, \quad (\text{C4})$$

$$\langle \hat{\mathbf{E}}^n(\mathbf{k}, t) \rangle = \Psi_1 \hat{\mathbf{E}}^{n-1} - ic\Psi_2 \frac{\mathbf{k}}{k} \times \hat{\mathbf{B}}^{n-1} + \left(\frac{ivck}{\varepsilon_0} A_1 + \frac{\Psi_2}{ck\varepsilon_0} \right) \hat{\mathcal{J}}^{n-1/2} + C_\rho (A_2 - A_1) \hat{\rho}^n \frac{\mathbf{k}}{k} + C_\rho (\theta^2 A_1 - A_2) \hat{\rho}^{n-1} \frac{\mathbf{k}}{k}. \quad (\text{C5})$$

Here, $\hat{T} = \prod_i [1 - \sin(k_i \Delta i/2)]$ represents a one-pass binomial smoother, and $\hat{S}(\mathbf{k}_m)$ is the particle shape factor, with $\mathbf{k}_m = \mathbf{k} + \mathbf{K}_m$ ($\mathbf{K}_m = 2\pi \sum_i u_i m_i / \Delta i$) for $i = \{x, y, z\}$.

As follows from the discrete continuity equation, the corrected current $\hat{\mathcal{J}}^{n-1/2}$ satisfies

$$\hat{\mathcal{J}}^{n-1/2} = \hat{\mathcal{J}}_d^{n-1/2} - \frac{(\mathbf{k} \cdot \hat{\mathcal{J}}_d^{n-1/2}) \mathbf{k}}{k^2} + \frac{(\mathbf{k} \cdot \mathbf{v}_{\text{gal}}) \mathbf{k} \hat{\rho}^n - \hat{\rho}^{n-1} \theta^2}{k^2 (1 - \theta^2)}. \quad (\text{C6})$$

As in Ref. [16], we use the following Ansatz for the electromagnetic modes:

$$\hat{\mathbf{E}}^n(\mathbf{k}) = \langle \hat{\mathbf{E}}(\mathbf{k}) \rangle e^{-i(\omega - \mathbf{k} \cdot \mathbf{v}_{\text{gal}})n\Delta t}, \quad (\text{C7a})$$

$$\langle \hat{\mathbf{E}}^n(\mathbf{k}) \rangle = \langle \hat{\mathbf{E}}(\mathbf{k}) \rangle e^{-i(\omega - \mathbf{k} \cdot \mathbf{v}_{\text{gal}})n\Delta t}, \quad (\text{C7b})$$

$$\delta \hat{f}^{n-1/2}(\mathbf{k}_m, \mathbf{p}) = \delta \hat{f}(\mathbf{k}_m, \mathbf{p}) e^{-i(\omega - \mathbf{k} \cdot \mathbf{v}_{\text{gal}})(n-1/2)\Delta t}, \quad (\text{C7c})$$

$$\hat{\mathcal{J}}_d^{n-1/2}(\mathbf{k}) = \hat{\mathcal{J}}_d(\mathbf{k}) e^{-i(\omega - \mathbf{k} \cdot \mathbf{v}_{\text{gal}})(n-1/2)\Delta t}, \quad (\text{C7d})$$

$$\hat{\rho}^n(\mathbf{k}) = \hat{\rho}(\mathbf{k}) e^{-i(\omega - \mathbf{k} \cdot \mathbf{v}_{\text{gal}})n\Delta t}, \quad (\text{C7e})$$

and after some amount of algebra, we derive the following equations for $\hat{\mathcal{J}}_d(\mathbf{k})$ and $\hat{\rho}(\mathbf{k})$ from the Vlasov equation (see a similar derivation in Ref. [16]):

$$\langle \hat{\mathcal{F}}(\mathbf{k}) \rangle = \langle \hat{\mathbf{E}}(\mathbf{k}) \rangle + \mathbf{v}_0 \times \langle \hat{\mathbf{B}}(\mathbf{k}) \rangle - [\mathbf{v}_0 \cdot \langle \hat{\mathbf{E}}(\mathbf{k}) \rangle] \mathbf{v}_0 / c^2, \quad (\text{C8})$$

$$\hat{\mathcal{J}}_d = i \frac{ck\varepsilon_0}{\hat{T}} \left[\xi_1 \langle \hat{\mathcal{F}} \rangle + (\xi_2 \cdot \langle \hat{\mathcal{F}} \rangle) \frac{\mathbf{v}_0}{c} \right], \quad (\text{C9})$$

$$\hat{\rho} = \frac{ik\varepsilon_0}{\hat{T}} (\xi_3 \cdot \langle \hat{\mathcal{F}} \rangle), \quad (\text{C10})$$

In addition, by substituting Eqs. (C6) and (C7) into Eqs. (C1), (C2), and (C4), the problem is reduced to the following set of equations to be solved:

$$s_\omega c \hat{\mathbf{B}} = t_{ck} c_\omega \frac{\mathbf{k} \times \hat{\mathbf{E}}}{k} + \chi_4' \frac{\hat{T}}{\varepsilon_0 c k} \frac{i\mathbf{k} \times \hat{\mathcal{J}}_d}{k}, \quad (\text{C11})$$

$$s_\omega \hat{\mathbf{E}} = -c_\omega t_{ck} \frac{kc\hat{\mathbf{B}}}{k} - is_\omega \frac{\hat{T}\mathbf{k}}{\varepsilon_0 k^2} \hat{\rho} - i\chi_4 \frac{\hat{T}}{\varepsilon_0 c k} \left[\hat{\mathcal{J}}_d - \frac{(\mathbf{k} \cdot \hat{\mathcal{J}}_d) \mathbf{k}}{k^2} \right], \quad (\text{C12})$$

$$\theta^2 c \langle \hat{\mathbf{B}} \rangle = \Psi_1 c \hat{\mathbf{B}} e^{i\omega\Delta t} + i\Psi_2 \frac{\mathbf{k} \times \hat{\mathbf{E}}}{k} e^{i\omega\Delta t} + \frac{ickA_1 \hat{T} \theta}{\varepsilon_0} \frac{\mathbf{k} \times \hat{\mathcal{J}}_d}{k} e^{\frac{i\omega\Delta t}{2}}, \quad (\text{C13})$$

- [5] B. B. Godfrey, Numerical Cherenkov instabilities in electromagnetic particle codes, *J. Comput. Phys.* **15**, 504 (1974).
- [6] B. B. Godfrey, Canonical momenta and numerical instabilities in particle codes, *J. Comput. Phys.* **19**, 58 (1975).
- [7] J.-L. Vay, C. G. R. Geddes, E. Cormier-Michel, and D. P. Grote, Numerical methods for instability mitigation in the modeling of laser wakefield accelerators in a Lorentz-boosted frame, *J. Comput. Phys.* **230**, 5908 (2011).
- [8] B. B. Godfrey and J.-L. Vay, Numerical stability of relativistic beam multidimensional PIC simulations employing the Esirkepov algorithm, *J. Comput. Phys.* **248**, 33 (2013).
- [9] X. Xu, P. Yu, S. F. Martins, F. S. Tsung, V. K. Decyk, J. Vieira, R. A. Fonseca, W. Lu, L. O. Silva, and W. B. Mori, Numerical instability due to relativistic plasma drift in EM-PIC simulations, *Comput. Phys. Commun.* **184**, 2503 (2013).
- [10] B. B. Godfrey, J.-L. Vay, and I. Haber, Numerical stability analysis of the pseudo-spectral analytical time-domain PIC algorithm, *J. Comput. Phys.* **258**, 689 (2014).
- [11] B. B. Godfrey and J.-L. Vay, Suppressing the numerical Cherenkov instability in FDTD PIC codes, *J. Comput. Phys.* **267**, 1 (2014).
- [12] B. B. Godfrey, J.-L. Vay, and I. Haber, Numerical stability improvements for the pseudospectral EM PIC algorithm, *IEEE Trans. Plasma Sci.* **42**, 1339 (2014).
- [13] B. B. Godfrey and J.-L. Vay, Improved numerical Cherenkov instability suppression in the generalized PSTD PIC algorithm, *Comput. Phys. Commun.* **196**, 121 (2015).
- [14] P. Yu, X. Xu, V. K. Decyk, F. Fiuza, J. Vieira, F. S. Tsung, R. A. Fonseca, W. Lu, L. O. Silva, and W. B. Mori, Elimination of the numerical Cherenkov instability for spectral EM-PIC codes, *Comput. Phys. Commun.* **192**, 32 (2015).
- [15] P. Yu, X. Xu, A. Tableman, V. K. Decyk, F. S. Tsung, F. Fiuza, A. Davidson, J. Vieira, R. A. Fonseca, W. Lu, L. O. Silva, and W. B. Mori, Mitigation of numerical Cherenkov radiation and instability using a hybrid finite difference-FFT Maxwell solver and a local charge conserving current deposit, *Comput. Phys. Commun.* **197**, 144 (2015).
- [16] R. Lehe, M. Kirchen, B. B. Godfrey, A. R. Maier, and J.-L. Vay, Elimination of numerical Cherenkov instability in flowing-plasma particle-in-cell simulations by using Galilean coordinates, *Phys. Rev. E* **94**, 053305 (2016).
- [17] M. Kirchen, R. Lehe, B. B. Godfrey, I. Dornmair, S. Jalas, K. Peters, J.-L. Vay, and A. R. Maier, Stable discrete representation of relativistically drifting plasmas, *Phys. Plasmas* **23**, 100704 (2016).
- [18] M. Kirchen, R. Lehe, S. Jalas, O. Shapoval, J.-L. Vay, and A. R. Maier, Scalable spectral solver in Galilean coordinates for eliminating the numerical Cherenkov instability in particle-in-cell simulations of streaming plasmas, *Phys. Rev. E* **102**, 013202 (2020).
- [19] F. Li, P. Yu, X. Xu, F. Fiuza, V. K. Decyk, T. Dalichaouch, A. Davidson, A. Tableman, W. An, F. S. Tsung, R. A. Fonseca, W. Lu, and W. B. Mori, Controlling the numerical Cherenkov instability in PIC simulations using a customized finite difference Maxwell solver and a local FFT based current correction, *Comput. Phys. Commun.* **214**, 6 (2017).
- [20] A. Pukhov, X-dispersionless maxwell solver for plasma-based particle acceleration, *J. Comput. Phys.* **418**, 109622 (2020).
- [21] R. Courant, K. Friedrichs, and H. Lewy, über die partiellen differenzgleichungen der mathematischen physik, *Math. Ann.* **100**, 32 (1928).
- [22] A. Taflov and S. C. Hagness, *Computational Electrodynamics: The Finite-Difference Time-Domain Method*, 3rd ed. (Artech House, Norwood, 2005).
- [23] Q. H. Liu, The pstd algorithm: A time-domain method requiring only two cells per wavelength, *Micro. Optic. Technol. Lett.* **15**, 158 (1997).
- [24] I. Haber, R. Lee, H. Klein, and J. Boris, in *Proceedings of Sixth Conference on Numerical Simulation Plasmas* (Berkeley, CA, 1973), pp. 46–48.
- [25] J.-L. Vay, I. Haber, and B. B. Godfrey, A domain decomposition method for pseudo-spectral electromagnetic simulations of plasmas, *J. Comput. Phys.* **243**, 260 (2013).
- [26] T. J. Mehrling, C. Benedetti, C. B. Schroeder, and E. Esarey, A subgrid algorithm for the efficient modeling of plasma-based accelerators with ion motion using quasi-static particle-in-cell codes, in *Proceedings of the 2018 IEEE Advanced Accelerator Concepts Workshop (AAC)* (IEEE, Piscataway, NJ, 2018), pp. 1–5.
- [27] H. Vincenti and J.-L. Vay, Detailed analysis of the effects of stencil spatial variations with arbitrary high-order finite-difference maxwell solver, *Comput. Phys. Commun.* **200**, 147 (2016).
- [28] S. Jalas, I. Dornmair, R. Lehe, H. Vincenti, J.-L. Vay, M. Kirchen, and A. R. Maier, Accurate modeling of plasma acceleration with arbitrary order pseudo-spectral particle-in-cell methods, *Phys. Plasmas* **24**, 033115 (2017).
- [29] J.-L. Vay, A. Almgren, J. Bell, L. Ge, D. P. Grote, M. Hogan, O. Kononenko, R. Lehe, A. Myers, C. Ng, J. Park, R. Ryne, O. Shapoval, M. Thévenet, and W. Zhang, Warp-x: A new exascale computing platform for beam-plasma simulations, *Nucl. Instrum. Methods Phys. Res. Sect. A* **909**, 476 (2018).
- [30] <http://github.com/ECP-WarpX/WarpX>.