

Unphysical kinetic effects in particle-in-cell modeling of laser wakefield acceleratorsEstelle Cormier-Michel,^{1,2} B. A. Shadwick,³ C. G. R. Geddes,¹ E. Esarey,^{1,2} C. B. Schroeder,¹ and W. P. Leemans^{1,2}¹*Lawrence Berkeley National Laboratory, Berkeley, California 94720, USA*²*Department of Physics, University of Nevada, Reno, Nevada 89557, USA*³*Department of Physics and Astronomy, University of Nebraska, Lincoln, Nebraska 68588-0111, USA*

(Received 10 December 2007; revised manuscript received 29 April 2008; published 23 July 2008)

Unphysical heating and macroparticle trapping that arise in the numerical modeling of laser wakefield accelerators using particle-in-cell codes are investigated. A dark current free laser wakefield accelerator stage, in which no trapping of background plasma electrons into the plasma wave should occur, and a highly nonlinear cavitated wake with self-trapping, are modeled. Numerical errors can lead to errors in the macroparticle orbits in both phase and momentum. These errors grow as a function of distance behind the drive laser and can be large enough to result in unphysical trapping in the plasma wake. The resulting numerical heating in intense short-pulse laser-plasma interactions grows much faster and to a higher level than the known numerical grid heating of an initially warm plasma in an undriven system. The amount of heating, at least in the region immediately behind the laser pulse, can, in general, be decreased by decreasing the grid size, increasing the number of particles per cell, or using smoother interpolation methods. The effect of numerical heating on macroparticle trapping is less severe in a highly nonlinear cavitated wake, since trapping occurs in the first plasma wave period immediately behind the laser pulse.

DOI: [10.1103/PhysRevE.78.016404](https://doi.org/10.1103/PhysRevE.78.016404)

PACS number(s): 52.38.Kd, 52.50.Jm, 41.75.Jv

I. INTRODUCTION

Plasma-based accelerators [1] are capable of supporting large amplitude plasma waves with electric fields up to hundreds of GV/m, approximately three orders of magnitude beyond those obtainable with conventional accelerators. Laser-plasma accelerator experiments, in which the source of the electrons was self-trapping from the background plasma, have been successful in generating highly energetic electrons. For example, near-monoenergetic electron bunches have been produced in laser-plasma accelerator experiments in the 100 MeV range [2–4] as well as the 1 GeV range [5]. In these experiments, a relativistically intense laser pulse is injected into a plasma with an initial laser pulse length on the order of or extending over a few plasma wavelengths. As the laser pulse propagates through the plasma, it undergoes self-modulation and self-steepening, leading to a highly nonlinear cavitated wake [6] (often referred to as the bubble [7] or blowout regime [8]). Simulations show that electrons are self-trapped from the background plasma near the back of the cavitated wake region and accelerated to high energies. Narrow energy spread electron beams were produced through control of the interaction length such that the acceleration occurred over a dephasing length [3,9].

To further improve the electron bunch quality and stability, a variety of laser-triggered injection methods have been proposed [10–14], and controlled injection via colliding laser pulses [15] and by control of plasma density gradients [16] has been achieved experimentally. The next generation of plasma accelerator experiments is likely to use a two-stage approach [17]. The first stage would be a relatively low energy injector, wherein the accelerated electron bunch is produced through self-trapping or controlled injection. The second, accelerating, stage must be such that additional plasma electrons are not self-trapped; in accelerator terminology, this stage would be “dark-current free.”

In a cold plasma, self-trapping of plasma electrons will not occur in a one-dimensional plasma wave for amplitudes below the cold, relativistic wave-breaking field [18] given by $E_{\text{WB}} = \sqrt{2}(\gamma_\phi - 1)^{1/2} E_0$, where $\gamma_\phi = 1/(1 - \beta_\phi^2)$, $v_\phi = c\beta_\phi$ is the plasma wave phase velocity (approximately the group velocity of the driver), $E_0 = c m \omega_p / e$ is the cold wave-breaking field, $\omega_p = ck_p = (4\pi n_0 e^2 / m)^{1/2}$ is the plasma frequency, and n_0 is the ambient electron plasma density. In a warm plasma, the maximum plasma wave amplitude is reduced to a value $E_{\text{max}} < E_{\text{WB}}$ [19]. Even for electric field amplitudes E_z less than the warm wave-breaking field it is possible for the hot electrons on the tail of the distribution to become trapped [20,21]. However, for plasma temperatures on the order of 10 eV, typical in laser wakefield accelerator experiments [22,23], the amount of self-trapped charge is expected to be very small for values of E_z on the order of E_0 . For example, the trapping fraction [20,21] is on the order of 10^{-6} for a 100 eV plasma and vanishingly small for 10 eV with $E_z = 2.25E_0$ and $\gamma_\phi = 10$. Furthermore, for an intense, ultrashort laser pulse (with length on the order of the plasma wavelength) interacting with a plasma below self-trapping, there are no significant physical heating mechanisms [24,25], in addition to compressional temperature oscillation in the case of a warm plasma, since the effect of collisions and the generation of slow phase velocity waves (e.g., Raman back-scattering) are negligible.

The particle-in-cell (PIC) method [26–28] is a widely used computational tool to study laser plasma interactions and, in particular, PIC codes have been used extensively to model laser-plasma accelerator experiments. This includes the study of self-trapping and electron acceleration. In a particle-grid approach such as PIC, finite-sized, charged macroparticles (typically representing 10^3 – 10^4 real electrons) interact with electromagnetic fields defined on a grid. The unavoidable discretization of the physical model and the necessarily small number of macroparticles relative to the




FIG. 1. Temperature evolution as a function of time for a 1D thermal plasma, initially at 0.1 eV and with no driver, with linear interpolation (solid curve), quadratic interpolation (dotted curve), and cubic interpolation (dash-dotted curve). The fast growing part (visible only with the linear interpolation) is grid heating (saturating at $k_g \lambda_D \sim 1$) and the slowly growing part is due to macroparticle scattering. See the Appendix for a discussion of quadratic and cubic interpolation.

number of physical electrons give rise to unphysical heating [27,28]. In addition, since the macroparticle positions are not restricted to mesh points, some form of interpolation is necessary to evaluate the force, resulting in trajectory errors. These numerical errors will alter phase-space and can mimic physical processes leading to incorrect interpretation of computational results. This will be of particular importance when attempting to model detailed kinetic effects, such as trapping of the background plasma electrons and generation of dark current in a plasma-based accelerator.

To illustrate the conventionally understood heating mechanisms in a PIC code, Fig. 1 shows simulations carried out using a modified version of the PSC code [29] (discussed below) in one-dimension (1D) for a constant density, quasineutral electron plasma in the absence of any laser or electron beam driver. The plasma is initially loaded uniformly in space with a Maxwellian momentum distribution having a temperature of 0.1 eV, 400 macroparticles per cell, and 240 cells per plasma wavelength ($\lambda_p = 2\pi c / \omega_p = 8 \mu\text{m}$ for $n_0 = 1.7 \times 10^{19} \text{ cm}^{-3}$), in a $130 \mu\text{m}$ spatial domain with periodic boundary conditions. The evolution of the mean squared momentum spread of the plasma, normalized to $m^2 c^2$, as a function of normalized time $\omega_p t$, is shown in Fig. 1. Two numerical heating mechanisms can be distinguished in Fig. 1: scattering [30] and grid heating [31]. Scattering, with a continuous slow growth rate due to the finite number of macroparticles (visible at early and late times in Fig. 1), depends mainly on the number of macroparticles per cell and on the interpolation method (see the Appendix for details on interpolation methods). Grid heating, with a fast growth rate (seen in the central portion of Fig. 1 with linear interpolation), saturates at $k_g \lambda_D \sim 1$ [28,31], where $k_g = 2\pi / \Delta z$ is the grid wave number, Δz the grid size, $\lambda_D = (k_B T / 4\pi n_0 e^2)^{1/2}$ the Debye length, T the electron plasma temperature, and k_B the Boltzmann constant. For the parameters used in Fig. 1, grid heating saturates at $k_B T \approx 9$ eV. Significant numerical heating occurs only after many plasma periods. The use of smoother interpolation functions greatly reduces the self-

heating rates, as is evident by the dotted (quadratic interpolation) and dash-dotted (cubic interpolation) curves in Fig. 1. Furthermore, when the macroparticles are loaded uniformly with zero temperature, no heating occurs in the absence of a driver. (A cold plasma at rest is a fixed point of the PIC algorithm.)

While both heating mechanisms illustrated in Fig. 1 are unphysical, they have distinct origins. Scattering is a “discrete particle effect” related to approximating the continuous phase-space density by individual macroparticles. As macroparticles drift from one cell to another in simulation of an unforced thermal plasma, there will be fluctuations in the number of macroparticles per cell. Since the average number of macroparticles per cell is rather small, even fluctuations of ± 1 macroparticle can lead to a large potential. These large, approximately random, fluctuations in potential lead to a localized electric field that acts on the macroparticles. In effect, the macroparticles scatter off fluctuations in the potential. Grid heating is a kinetic instability resulting from the grid aliasing of high-frequency modes (not resolved by the grid) to low frequencies [31]. The modification of the usual warm-plasma dispersion relation by aliasing effects results in roots with positive imaginary parts. Since this numerical instability arises from the same basic physical processes (but triggered by grid aliasing) as true plasma modes, it can both mimic real plasma phenomena and act to drive other, physical plasma responses. Common to both mechanisms is the creation of particle energy; likewise both mechanisms result in dynamical evolution of a state that physically is an equilibrium, but not a fixed point of the PIC algorithm.

In this paper we discuss numerical effects that can lead to artificial heating and unphysical trapping in PIC simulations of laser wakefield accelerators. It is shown that in the presence of a laser driver, a plasma that is initially loaded uniformly and cold (no initial momentum) heats rapidly, over just a few plasma periods, to temperatures greatly exceeding $k_g \lambda_D \sim 1$. This is in contrast with the well-known numerical heating mechanisms shown in Fig. 1, for which no heating occurs when a plasma is loaded cold. Numerical heating in the laser-driven case can lead to effective plasma temperatures high enough for macroparticles to become trapped and accelerated in the wake. In other words, trapping seen in PIC simulations can be, depending on the parameter regime, entirely due to artifacts in the PIC algorithm and not physical processes in the plasma.

Simulations are carried out using a modified version of the plasma simulation code (PSC) [29], which implements the standard (momentum-conserving) PIC algorithm described in Ref. [28]. Differential equations are solved using a leap-frog finite differencing method (i.e., with staggered mesh and time steps) which is second order accurate in grid size and time step. The Lorentz force is applied by using the Boris algorithm [32]. Deposition of particle current on the grid is obtained using a local charge conserving scheme as described in Ref. [33]. A description of the interpolation algorithms is presented in the Appendix. These standard algorithms are used in other PIC codes frequently used to model laser-plasma accelerators (e.g., Refs. [34–36]) and, therefore, the results of this paper will apply to simulations done with any code based on these standard algorithms.

In the absence of collisions, the PIC algorithm is attempting to solve the coupled Vlasov-Maxwell equations. In Secs. II and III of this work we compare numerical solutions, via the PIC algorithm, to known properties of the exact solutions of the underlying physical model (the Vlasov-Maxwell equations). In particular, we consider an initial plasma phase-space distribution with a delta function momentum distribution (an initially cold plasma) and a wake amplitude below wave breaking. This is an exact solution of the Vlasov-Maxwell equations, i.e., a δ function momentum distribution remains a δ function throughout the plasma evolution [25]. For these initial conditions the Vlasov-Maxwell equations yield the cold fluid model. Therefore, for a short laser pulse interacting with an initially cold, collisionless, underdense plasma below wave breaking, there are no mechanisms for heating and any growth in the plasma momentum spread must then be due to numerical artifacts in the PIC algorithm. In Sec. II, details of phase space for a 1D case in the standard laser wakefield accelerator (LWFA) regime, relevant to a dark current free accelerator stage (well below wavebreaking), are presented. Section III explores this standard LWFA regime in two-dimensions (2D). Section IV presents 2D simulations indicating the impact of these effects on self-trapping in a highly nonlinear case (the cavitating, blowout or bubble regime), which is a regime relevant to recent experiments. A conclusion is presented in Sec. V.

II. ONE-DIMENSIONAL STANDARD LWFA REGIME

In this section, the artificial heating and trapping of macroparticles, and the details of phase space, are studied for a standard LWFA in 1D. The peak amplitude of the linear polarized laser field is $a_0=2$, where a_0 is related to the peak laser intensity I_0 and laser wavelength $\lambda_0=2\pi c/\omega_0$ by $a_0^2 \approx 7.32 \times 10^{-19} (\lambda_0 [\mu\text{m}])^2 I_0 [\text{W}/\text{cm}^2]$, i.e., $I_0=8.5 \times 10^{18} \text{ W}/\text{cm}^2$ for $\lambda_0=0.8 \mu\text{m}$ and $a_0=2$. The initial normalized laser intensity profile is of the form $a_0^2 \exp(-2z^2/L^2)$, with $k_p L=2$, and z is the coordinate along the direction of laser propagation. Here $\omega_0/\omega_p=10$ is considered, which implies a plasma wavelength of $\lambda_p=8 \mu\text{m}$ ($n_0=1.7 \times 10^{19} \text{ cm}^{-3}$) and a laser pulse length of $L=2.5 \mu\text{m}$ [10 fs full width at half maximum (FWHM) laser intensity duration]. The 1D simulation box is $130 \mu\text{m}$ long, and the laser is launched from the boundary of the simulation box into a uniform plasma with a $l=105 \mu\text{m}$ cosine ramp of the form $n(z)=0.5[1-\cos(\pi z/l)]n_0$ for $0 < z < l$, and with $n(z)=n_0$ for $z \geq l$. A window moving at the speed of light is used to follow the laser until $t=31.5\lambda_p/c$, which is a propagation distance of approximately twice the window size. The macroparticles are loaded uniformly at rest, using either $N_{\text{PPC}}=100$ or $N_{\text{PPC}}=400$, where N_{PPC} is the number of macroparticles per cell. Resolution is varied to study numerical heating.

The normalized longitudinal wakefield E_z/E_0 as a function of z is shown in Fig. 2 at $t=31.5\lambda_p/c$ for different resolutions. Note that convergence of the field occurs relatively quickly: there is little difference between the curves for $\Delta z < \lambda_0/36$. For $\Delta z=\lambda_0/24$, the wakefield is not yet converged and is approximately 3.5% lower in amplitude than with




FIG. 2. (Color online) Longitudinal electric field E_z/E_0 for $N_{\text{PPC}}=400$ and the resolutions: (a) $\Delta z=\lambda_0/60$ [magenta (light gray)] and $\Delta z=\lambda_0/36$ [blue (dark gray) dashed curve] and (b) $\Delta z=\lambda_0/60$ [magenta (light gray)] and $\Delta z=\lambda_0/24$ [blue (dark gray) dashed curve].

$\Delta z=\lambda_0/60$. As we will see below, kinetic quantities such as momentum spread converge far more slowly.

No particle trapping should occur in this regime, since the plasma is initially cold and the wakefield is well below the cold relativistic wave-breaking field $E_z < E_0[2(\gamma_\phi - 1)]^{1/2}$. For a cold plasma, the particle orbits are identical to the cold fluid characteristics and thus trapping can only occur at wave breaking, at which point the plasma density has singularities at the peak compression locations. The evolution of the plasma temperature has previously been studied using a warm plasma model [24,25], which predicts that an initially cold collisionless plasma remains cold in this regime, i.e., a δ function momentum distribution remains a δ function. Thus it is typically expected that the plasma should keep its initial temperature $T=T_0=0$. Note that in this regime, there are no physical heating mechanisms in the model, since there are no collisions and there are no slow phase velocity plasma waves. Thus, a converged PIC simulation with an initially cold plasma should match the cold fluid result.

The PIC simulations, however, show macroparticles trapped in the wake, as seen in Fig. 3. Figure 3 shows the macroparticle phase space (momentum versus position) at $t=31.5\lambda_p/c$ for different resolutions and $N_{\text{PPC}}=400$. One signature of physical trapping in a warm 1D plasma would be the appearance of trapped particles in the first plasma wave period immediately behind the laser pulse [20]. This is not the case in Fig. 3, in which no trapping is present for several plasma periods behind the laser pulse.

Figure 4 shows the evolution of the longitudinal momentum spread $\sigma_{u_z}^2 = \langle (u_z - \langle u_z \rangle)^2 \rangle$, proportional to the effective plasma temperature $\sigma_{u_z}^2 \approx k_B T / mc^2$ if $\sigma_{u_z}^2 \ll 1$, where $u_z = p_z / mc$ is the longitudinal normalized momentum of the macroparticles and the angular brackets denote an averaging over particles, for (a) different number of macroparticles per cell and (b) different longitudinal resolutions. Figure 4 was obtained by finding the mean momentum in each cell, which is then linearly interpolated at the particle positions to calcu-




FIG. 3. Macroparticle phase space at $ct=31.5\lambda_p$, with $\omega_0/\omega_p=10$, $a_0=2$, and $k_p L=2$, using $N_{PPC}=400$, and (a) $\Delta z=\lambda_0/24$, (b) $\Delta z=\lambda_0/30$, (c) $\Delta z=\lambda_0/36$, (d) $\Delta z=\lambda_0/48$, (e) $\Delta z=\lambda_0/60$.

late the variance in each cell. Warm fluid theory [19,24,25] indicates that in this regime, the physical plasma temperature in the wake oscillates according to $T=(n/n_0\gamma_f)^2 T_0$, where T_0 is the initial temperature and γ_f is the Lorentz factor of the fluid oscillation. There should be no growth in the temperature in the wake. For an initially warm plasma, the temperature oscillates at the plasma period, and an initially cold plasma ($T_0=0$) should remain cold. This is in contrast to the numerical heating shown in Fig. 4, which shows a periodic oscillation at the plasma period superimposed on a rapid growth of the momentum spread as a function of distance behind the laser pulse. In the first several plasma wave periods, the macroparticle momentum spread is too low to allow trapping. Trapping is only present in the later trailing plasma wave buckets, as seen in Fig. 3, because the momentum spread is increasing as a function of distance behind the laser driver. Increasing the number of macroparticles per cell [Fig. 4(a)] has a small effect on the momentum spread. As the resolution increases [Fig. 4(b)], however, the momentum spread decreases, which leads to macroparticles trapped later in the wake, as seen in Fig. 3. The exception is for the lowest resolution, Fig. 3(a), where fewer particles are trapped. Here




FIG. 4. (Color online) Longitudinal momentum spread for (a) $\Delta z=\lambda_0/36$ and with $N_{PPC}=400$ (black curve) and $N_{PPC}=100$ [red (gray) curve]; (b) $N_{PPC}=400$ and with $\Delta z=\lambda_0/36$ (black curve) and $\Delta z=\lambda_0/60$ [red (gray) curve].

the wakefield is not yet converged, i.e., at $\Delta z=\lambda_0/24$ the field is 3.5% lower in amplitude than at $\Delta z=\lambda_0/60$, and the lower wake amplitude implies a high threshold temperature (momentum spread) for trapping and hence less trapped particles [20]. In comparison to the undriven self-heating case shown in Fig. 1, the growth in the momentum spread is much more rapid (~ 5 plasma periods against ~ 100 in the non-driven case), and reaches values two orders of magnitude higher.

To understand the development of the numerical momentum spread, the phase space (momentum versus position) of the macroparticles is plotted in Fig. 5 for $\Delta z=\lambda_0/36$ and $N_{PPC}=400$ with detailed plots of the first (A), third (B), and fifth (C) buckets after the laser pulse. As a function of distance behind the driver, phase space develops an increasingly complex structure. Figure 6 shows the phase space for (a) $\Delta z=\lambda_0/36$ and $N_{PPC}=400$, (b) $\Delta z=\lambda_0/36$ and $N_{PPC}=100$, and (c) $\Delta z=\lambda_0/60$ and $N_{PPC}=400$. The insets show a magnification of the phase space at the first (A) and fifth (B) buckets after the laser pulse. As shown in Figs. 6(a)–6(c), the phase-space structure is dependent on the resolution and number of macroparticles per cell. At a resolution of $\Delta z=\lambda_0/36$ the longitudinal electric field is accurately represented. Increasing the resolution leads to little change in the wakefield, but results in significant changes in the macroparticle phase space.

As the phase-space structure develops increasingly more complexity behind the driver, the macroparticle orbits are displaced (both in momentum and in phase) from the ideal cold fluid orbit. Since the electric field of the wake is well




FIG. 5. Macroparticle phase space for $\Delta z = \lambda_0/36$ and $N_{\text{PPC}} = 400$. The boxes A, B, and C show detail of phase space at the first, third, and fifth bucket after the laser pulse.

resolved at these resolutions, so is the separatrix (determined by the electric field structure) that separates trapped from untrapped orbits in phase space. Eventually, the macroparticle orbits are sufficiently displaced from the cold fluid orbits such that some macroparticles become trapped in the wakefield. Figure 7 shows the macroparticle phase space for $\Delta z = \lambda_0/36$ and $N_{\text{PPC}} = 400$. On top of the macroparticle phase space, the cold fluid orbit is shown as well as the separatrix between trapped and untrapped orbits. The cold fluid orbit and the separatrix are obtained using the relationship $u_z = \beta_\varphi \gamma_\varphi^2 (H + \Phi) \pm \gamma_\varphi [\gamma_\varphi^2 (H + \Phi)^2 - 1]^{1/2}$, where $\gamma_\varphi = \omega_0/\omega_p$ is the phase velocity of the wake $\beta_\varphi = \sqrt{1 - 1/\gamma_\varphi^2}$ and H is the Hamiltonian with $H = 1$ for the cold fluid orbit and $H = 1/\gamma_\varphi - \Phi_{\text{min}}$ for the separatrix [37]. The potential Φ is the potential obtained from the macroparticle momentum through the relation $\gamma - u_z = 1 + \Phi$. As the spread around the cold fluid orbit increases with distance behind the laser pulse, some macroparticles cross the separatrix and are trapped and accelerated in the wake.

Increasing the resolution decreases the phase-space orbit displacements, which results in less trapping, as shown in Figs. 6(a) and 6(c). The effects of smoother interpolation functions and current smoothing (described in the Appendix) are shown in Fig. 8, which plots the macroparticle phase space for $\Delta z = \lambda_0/36$ and $N_{\text{PPC}} = 400$ for (a) linear interpolation, (b) quadratic interpolation [Eq. (A7)], (c) cubic interpolation [Eq. (A8)], and (d) linear interpolation with a four-pass (1, 2, 1) filter plus compensator [38] on the current. Figure 9 shows the longitudinal momentum spread evolution for the same cases. Using quadratic or cubic interpolation reduces the orbit displacements in phase space and reduces the amount of trapping, although the difference between quadratic and cubic interpolation is less than between linear and quadratic interpolation. The current smoothing algorithm reduces the phase-space structure early, but later in the wake behind the driver the phase-space structure contains large orbit displacements, as is evident by comparing position A

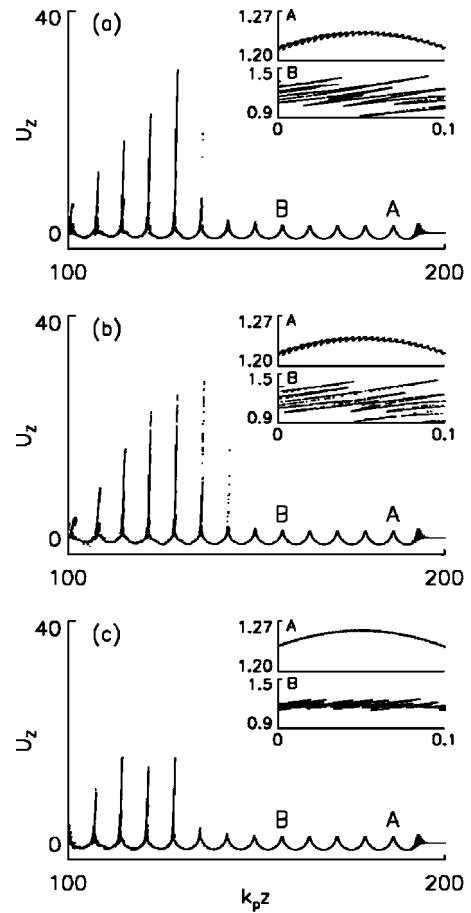


FIG. 6. Macroparticle phase space at $ct = 31.5\lambda_p$, with $\omega_0/\omega_p = 10$, $a_0 = 2$, and $k_p L = 2$, using (a) $\Delta z = \lambda_0/36$ and $N_{\text{PPC}} = 400$, (b) $\Delta z = \lambda_0/36$ and $N_{\text{PPC}} = 100$, and (c) $\Delta z = \lambda_0/60$ and $N_{\text{PPC}} = 400$. The insets A, B show a magnification of the phase space at the first and fifth buckets after the laser pulse.

with position B in Fig. 8(d). Consequently, current smoothing does not reduce the artificial trapping for these parameters. That quadratic interpolation is more effective in reducing the growth of these phase-space structures suggests that errors are mainly due to accumulation of interpolation errors of the fields onto the macro-particle locations (which is not




FIG. 7. (Color online) Macroparticle phase space of the seventh and eighth bucket behind the laser pulse (black) for $\Delta z = \lambda_0/36$ and $N_{\text{PPC}} = 400$, with the cold fluid orbit [green (gray) dashed curve] and the separatrix [magenta (gray) dash-dotted curve] representing the limit between untrapped and trapped orbits.

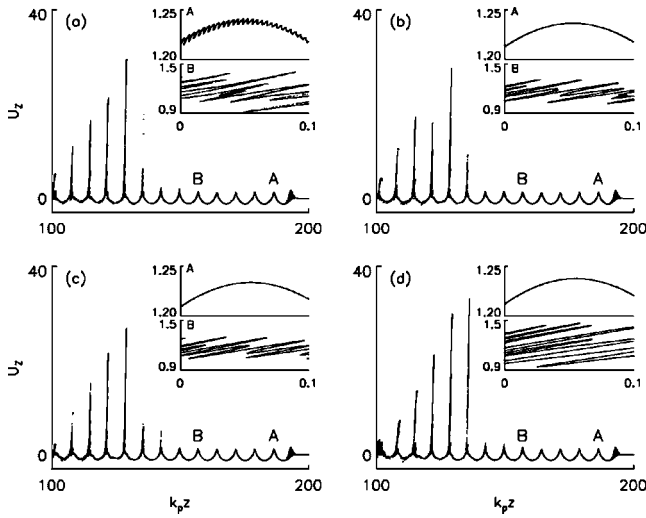


FIG. 8. Macroparticle phase space at $ct=31.5\lambda_p$, with $\omega_0/\omega_p = 10$, $a_0=2$, $k_p L=2$, $\Delta z=\lambda_0/36$, and $N_{PPC}=400$, and using (a) linear interpolation, (b) quadratic interpolation, (c) cubic interpolation, and (d) linear interpolation with a (1, 2, 1) filter with compensator applied to the current. The insets A and B show a magnification of the phase space at the first and fifth buckets after the laser pulse.

affected by the smoothing of the current), rather than high frequencies introduced in depositing the current on the grid.

The initial evolution of the phase-space structures (as seen in Fig. 5) is primarily due to spatial correlations in the interpolation error, which is different from the heating mechanisms shown in Fig. 1. Only once the orbit displacements grow to a point such that, after having been “coarse grained” to the grid, the resulting current distribution corresponds to a broadened phase-space distribution, can the collective instability responsible for grid heating turn on. In addition, the momentum distribution function exhibits multiple peaks at a given point in space which can lead to additional heating via a multiple beam instability [39]. Linear interpolation, which forces zero interpolation error at the grid points, leads to rather large relative errors and results in nearby macroparticles experiencing significantly different fields. Since these errors are correlated with the position within the cell, subsequent time steps tend to reinforce the relative error. As the resolution is increased, for a given interpolation method, the absolute magnitude of the difference in interpolation error across the cell is reduced, leading to neighboring orbits having smaller relative displacements, slowing the growth of the structures (see Fig. 9). Similarly, by forcing greater continuity (at the cost of absolute accuracy, see discussion in the Appendix), smoother interpolation methods result in an intercell approximation for the fields that varies much less over the cell than for the linear case.

Figure 10 shows the dependence of the temperature evolution as a function of laser intensity ($a_0=1, 2$, and 3) for $\Delta z=\lambda_0/36$ and using 400 macroparticles per cell. As the laser intensity becomes larger the plasma heats more rapidly and to a higher temperature. Figure 11 shows details of phase space for $a_0=3$ using linear interpolation and $\Delta z=\lambda_0/36$. As in the previous case, the longitudinal electric field is below the wave-breaking limit ($E_z/E_0=1.8 < E_{WB}/E_0=4.24$), so no




FIG. 9. (Color online) Longitudinal momentum spread $\sigma_{u_z}^2$ for $\Delta z=\lambda_0/36$ and $N_{PPC}=400$ with (a) linear interpolation (black curve), quadratic interpolation [magenta (light gray) curve], cubic interpolation [blue (dark gray) dashed curve], and (b) linear interpolation (black curve), quadratic interpolation [magenta (light gray) curve], and linear interpolation with a (1, 2, 1) filter with compensator [blue (dark gray) dashed curve].

trapping is expected. Figure 11 shows some significant phase-space structure development already in the first bucket after the laser field (inset) and trapping occurs in the third bucket whereas for $a_0=2$, at the same resolution, trapping appears in the eighth bucket after the laser pulse.

Figure 12 shows the amount of trapped charge in the wake per m^2 (the number of macroparticles normalized by $n_0\Delta z/N_{PPC}$, where n_0 is the plasma density in m^{-3}) on a log scale versus number of cells per laser wavelength for various N_{PPC} , for the $a_0=2$ case and using linear interpolation. Here, only trapped macroparticles with $u_z \geq u_g \approx \omega_0/\omega_p=10$ are considered, where u_g is the approximate laser group and wake phase velocity. The amount of trapped charge does not always follow a consistent behavior, but the general trend is that trapping decreases when increasing the resolution and




FIG. 10. (Color online) Longitudinal momentum spread for $a_0=3$ (black), $a_0=2$ [magenta (light gray)], and $a_0=1$ [blue (dark gray) dotted curve], using $\Delta z=\lambda_0/36$ and $N_{PPC}=400$.




FIG. 11. Macroparticle phase space at $ct=31.5\lambda_p$ with $\omega_0/\omega_p=10$, $a_0=3$, and $k_p L=2$, using $\Delta z=\lambda_0/36$ and $N_{\text{PPC}}=400$. The inset A shows a magnification of the phase space at the first bucket after the laser pulse.

increasing the number of macroparticles per cell. To reach convergence both resolution and number of macroparticles per cell must be increased at the same time [25]. Departures from the general trend may be due to the fact that the momentum spread tends to reach similar levels near the back of the simulation window, even though the momentum spread immediately after the driver can be lower for finer resolution. It should be noted that, as the phase structure is highly “filamented” (as seen in Fig. 5) and far from a Maxwellian distribution, the macroparticle momentum distribution is not a thermal distribution and the fraction of trapped particles is several orders of magnitude higher than predicted in Ref. [20], which assumed an initial thermal momentum distribution and the correct physical response of the plasma [i.e., $T \approx T_0(n/n_0\gamma_f)^2$].

For comparison to the uniform cold load used throughout this paper, it is insightful to compare to a case in which the macroparticles are loaded uniformly on the spatial grid, but with random momenta selected from an initial momentum distribution that is Maxwellian with a 10 eV initial temperature. The results are shown in Fig. 13 for $a_0=2$, $k_p L=2$, $\lambda_p=8 \mu\text{m}$, $\lambda_0=0.8 \mu\text{m}$, $\Delta z=\lambda_0/60$, and $N_{\text{PPC}}=400$ [identical to those in Fig. 6(c), but with a warm initial state]. For the initial temperature of 10 eV, the Debye length $\lambda_{D0} \approx \lambda_p/1420 = \lambda_0/142$ while, $k_g \lambda_{D0} \approx 2.7$. Hence, the Debye




FIG. 12. Amount of trapped charge (number of macroparticles with $u_z \geq 10$ multiplied by $n_0 \Delta z / N_{\text{PPC}}$ in m^{-2}) as a function of longitudinal resolution for $N_{\text{PPC}}=400$ (\circ), $N_{\text{PPC}}=100$ ($+$), and $N_{\text{PPC}}=20$ (Δ), using linear interpolation.




FIG. 13. (Color online) Longitudinal momentum spread at $ct=31.5\lambda_p$ for $\Delta z=\lambda_0/60$ and $N_{\text{PPC}}=400$ with an initial temperature of 10 eV ($\sigma_{u_z}^2 \sim 2 \times 10^{-5}$) (black), compared to the momentum spread with the same numerical parameters but with an initially cold plasma [magenta (gray)].

length is resolved and self-heating should be minimal (see Fig. 1). Figure 13 shows that the plasma temperature oscillates around $T=10$ eV at the plasma period [24,25], until contribution from the numerical heating dominates the momentum spread, which then grows at the same rate and to the same value as in the cold case, greatly exceeding $\lambda_D = \Delta z / 2\pi$.

Often in the study of laser-plasma accelerators a technique known as the “moving window” is used. In this case, the simulation box is typically only a few plasma periods long and moves at the speed of light in the direction of laser propagation (in this way, the simulation box is approximately comoving with the laser pulse). When using a moving window the simulation box is constantly moving into undisturbed plasma. This limits the extent to which the plasma is heated by numerical artifacts because the “residence time” for plasma in the simulation is limited to the box length over c . This approach is especially useful for cases in which the phenomena of interest occurs immediately behind the driver, such as in the blowout regime discussed in Sec. IV. However, some specific cases may require longer windows, e.g., colliding pulse injection, where a counterpropagating pulse is launched from the right side of the box, or simulation of down-ramp injection, where the trapping may occur a several plasma periods behind the laser pulse [16].

III. TWO-DIMENSIONAL STANDARD LWFA REGIME

In this section, the effects of unphysical heating in PIC simulations (using PSC) are studied in 2D for a dark current free stage of a standard LWFA. The initial normalized peak amplitude of the linearly polarized laser pulse is $a_0=1.15$ and the initial normalized laser intensity profile is of the form $a_0^2 \exp(-2x^2/r_0^2 - 2z^2/L^2)$, with $k_p L=2$ and $k_p r_0=8$, where z is the longitudinal coordinate and x is the transverse coordinate. Also, ω_0/ω_p is 10 with $\lambda_0=0.8 \mu\text{m}$, which implies $\lambda_p=8 \mu\text{m}$ ($n_0=1.7 \times 10^{19} \text{cm}^{-3}$), the laser pulse length is $L=2.5 \mu\text{m}$ (10 fs FWHM laser intensity duration), and the laser spot size is $r_0=10.2 \mu\text{m}$. The laser propagates in a parabolic plasma channel matched to its spot size, i.e., the initial plasma density profile [1] is $n=n_0+\Delta n_c x^2/r_0^2$ with $\Delta n_c [\text{cm}^{-3}]=1.13 \times 10^{20}/r_0 [\mu\text{m}]$. The 2D simulation box is $65 \mu\text{m}$ in the z direction (with a moving window) and




FIG. 14. Schematic of the laser polarizations: (a) Transverse electric field of the laser polarized in the simulation plane (*p* polarized). (b) Transverse electric field of the laser polarized out of the simulation plane (*s* polarized).

153 μm in the x direction. Figure 14 shows a schematic of the laser polarization with respect to the 2D simulation box that is oriented in the x - z plane. Polarization in (*p*-polarized) and out (*s*-polarized) of the x - z plane is considered, with the laser propagating in the z direction. The macroparticles are loaded uniformly and cold (i.e., no initial momentum), using four macroparticles per cell.

The longitudinal electric field of the wake is shown in Fig. 15 as a function of grid size, which indicates that both E_z and the ponderomotive force driving E_z are well resolved for $\Delta z \leq \lambda_0/24$, leading to a peak value of approximately $0.44E_0$ at early time. Compared to $\Delta z = \lambda_0/48$, the peak field is 2% lower for $\Delta z = \lambda_0/24$ and 13% lower for $\Delta z = \lambda_0/12$. At $\omega_p t = 218$ the laser has evolved and $E_z = 0.8E_0$. Since wake front curvature effects can lower the wave breaking and self-trapping threshold in 2D [40], a lower peak laser intensity ($a_0 = 1.15$) was chosen for these 2D studies compared to the 1D studies ($a_0 = 2$) in Sec. II. The value of $a_0 = 1.15$ was also chosen because it well illustrates the difference in trapping between in and out of plane polarization (see Fig. 16). Numerical solutions with a cold fluid model [41] do not show a singular plasma density for these parameters, which is an indication that there should be no trapping in this example. A δ function momentum distribution is an exact solution of Vlasov and, therefore, an initially cold, collisionless plasma must remain cold in this regime [25]. Thus, in this regime, the PIC simulation with an initially cold plasma should con-




FIG. 15. (Color online) Longitudinal electric field on axis at $\omega_p t = 55$ for (a) $\Delta x = \lambda_0/2.5$ and $\Delta z = \lambda_0/48$ [magenta (light gray) curve] and $\Delta x = \lambda_0/2.5$ and $\Delta z = \lambda_0/12$ [blue (dark gray) dashed curve], (b) $\Delta x = \lambda_0/2.5$ and $\Delta z = \lambda_0/48$ [magenta (light gray) curve] and $\Delta x = \lambda_0/2.5$ and $\Delta z = \lambda_0/24$ [blue (dark gray) dashed curve].




FIG. 16. Longitudinal phase-space of the electrons at $\omega_p t = 218$, with the laser (a) *p* polarized, (b) *s* polarized. The resolution is $\Delta x = \lambda_0/2.5 = r_0/32$ and $\Delta z = \lambda_0/24$.

verge to the cold fluid result with no self-trapping. However, the PIC simulations show macroparticles trapped in the wake, as seen in Fig. 16. Here, trapped macroparticles are defined to be those with normalized longitudinal momenta $u_z = p_z/mc > 1$, that exceed the peak momentum for cold fluid oscillation in this case. The trapped charge, which depends on the numerical parameters (resolution, number of macroparticles per cell), is caused by numerical heating of the plasma. As with the 1D case, numerical errors lead to displacements of the macroparticle orbits in both momentum and phase (relative to the physical cold fluid orbits). These errors increase with distance behind the laser pulse, and trapping can occur when the errors become sufficiently large.

A. Polarization dependence and macroparticle push

The interpolation of the electromagnetic fields [Eqs. (A2) and (A3)] at the macroparticle positions leads to errors in the evaluation of the force [Eq. (A4)] acting on the macroparticles, which leads to errors in the resulting momenta of the macroparticles. These errors are more significant when the macroparticles experience forces with short periods on the simulation grid, e.g., when the macroparticles experience fast oscillations in the laser field, which occurs when the laser is polarized in the 2D simulation plane [Fig. 14(a)].

This is illustrated in Fig. 17, which shows selected macroparticle orbits for *p* and *s* polarization with $\Delta x = \lambda_0/2.5$ and $\Delta z = \lambda_0/24$, with linear interpolation. Two rows of particles




FIG. 17. (Color online) Macroparticle orbits for $\Delta x = \lambda_0/2.5$ and $\Delta z = \lambda_0/24$, with linear interpolation. (a) The laser is p polarized. (b) The laser is s polarized. The blue and red are for different initial longitudinal positions of the macroparticles.

are shown, each with fixed initial longitudinal position and each consisting of a set of 32 particles evenly spaced in the transverse direction. For the case of a s -polarized laser, the trajectories of the macroparticles correspond to a nearly laminar flow, which is the expected result for an initially cold plasma below the wave-breaking limit. However, when the laser is p polarized, the field interpolation errors caused by the under-resolved grid in the fast ($\omega_0/\omega_p \gg 1$) laser oscillations lead to errors in the momentum advance resulting in significant perturbation of the macroparticle orbits near the end of the laser pulse.

The effect of the field interpolation errors at the macroparticle position was estimated by introducing analytically specified laser fields (no plasma waves), and examining test particle orbits in the transverse plane, the laser being p polarized. Here, the particle push is the same as in the PIC algorithm, without field update since the laser field is specified analytically. Figure 18(a) shows two test macroparticles initially loaded on axis (same transverse position) and in the same cell but separated longitudinally by $\Delta z/2$. Behind the laser pulse, the two particles in Fig. 18(a) have different final momenta, which gives rise to the trajectory crossing seen in Fig. 17(a). Figures 18(b) and 18(c) show the effect of the different resolutions on the particle trajectory error. It is found that the main role is played by the longitudinal resolution, and that a resolution of at least $\Delta z \leq \lambda_0/50$ is necessary to minimize the macroparticle trajectory errors.

The particle push error [Fig. 18(a)] will also depend on the laser parameters. The errors will become more important




FIG. 18. (Color online) Orbit of (a) two test particles in prescribed fields initially on axis and in the same cell using linear interpolation with $\Delta x = \Delta z = \lambda_0/24$, (b) two test particles at the same position with $\Delta x = \lambda_0/24$ and $\Delta z = \lambda_0/12$ [magenta (light gray) curve] and $\Delta z = \lambda_0/48$ [blue (dark gray) dashed curve], (c) two test particles at the same position with $\Delta z = \lambda_0/48$ and $\Delta x = \lambda_0/2.5$ [magenta (light gray) curve] and $\Delta x = \lambda_0/24$ [blue (dark gray) dashed curve], and (d) two test particles initially on axis and in the same cell using quadratic interpolation and the resolution $\Delta x = \Delta z = \lambda_0/24$.

if the laser field is more intense and as the particles go through more laser oscillations (longer pulse) and accumulate interpolation errors.

Using smoother interpolation functions significantly decreases the trajectory error. Figure 18(d) shows that with quadratic interpolation the trajectory error inside the laser pulse is suppressed, which indicates that the error mainly depends on how well the field is interpolated at the macroparticle positions. Macroparticle orbits, in the self-consistent case, with the laser p polarized, obtained using quadratic [see Eq. (A7)] interpolation is shown in Fig. 19. The trajectory error inside the laser pulse is greatly reduced and macropar-




FIG. 19. (Color online) Macroparticle orbits for $\Delta x = \lambda_0/2.5$ and $\Delta z = \lambda_0/24$, using the quadratic interpolation function.




FIG. 20. (Color online) Macroparticle orbits, for $\Delta x = \lambda_0/2.5$ and $\Delta z = \lambda_0/24$, using linear interpolation and current smoothing with compensator: (a) p polarization and (b) s polarization.

ticles follow a laminar flow as expected. At later time however, there still remains small macroparticle orbit separation, owing to the numerical heating which is further suppressed when using cubic interpolation [see Eq.(A8)].

Shown in Fig. 20 are macroparticle trajectories computed using linear interpolation [Eq.(A6)] with the addition of smoothing using a (1, 2, 1) filter with compensator (described in the Appendix) on the current densities. When the laser is s polarized heating is reduced [Fig. 20(b)] compared to Fig. 17(b)]. Significant trajectory errors remain when the laser is p polarized [Fig. 20(a)] because the interpolation error is not reduced by smoothing. This suggests that smoother interpolation functions are more efficient than current smoothing for obtaining accurate trajectories.

B. Plasma momentum spread

The transverse and longitudinal momentum spread $\sigma_u^2 = \langle (u_{x,z} - \langle u_{x,z} \rangle)^2 \rangle$ evaluated on axis, is shown in Figs. 21–25 for a variety of 2D cases with different resolutions and interpolation methods. For clarity the momentum spread has been smoothed over 12 cells which does not change the trend. The momentum spread is proportional to the plasma temperature $\sigma_u^2 \approx k_B T / mc^2$ assuming $\sigma_u^2 \ll 1$ (i.e., typical plasma temperature of 10 eV corresponds to $\sigma_u^2 \approx 2 \times 10^{-5}$). The general trend, as previously, is that the macroparticles, which are initially loaded uniformly and cold, develop errors in momentum and phase as they are pushed through the laser field.




FIG. 21. (Color online) Transverse momentum spread (smoothed over 12 cells) with the resolution $\Delta x = \lambda_0/2.5$ and $\Delta z = \lambda_0/24$, the laser being p polarized (black curve) or s polarized [red (gray) curve], with (a) linear interpolation and (b) quadratic interpolation. Note that the binning of the macroparticles to calculate the momentum spread can lead to artificially high values in the fast oscillations of the laser (first peak).

These orbit errors continue to grow in the wake behind the driver toward the end of the simulation box, leading to an increase in the momentum spread. When the orbit errors become large, macroparticles are trapped in the wake. Cases run with larger longitudinal simulation boxes indicate that the momentum spread does saturate sufficiently far behind the driver, but the behavior is likely to be strongly affected by the large number of trapped macroparticles and the resulting damping (beam loading) of the wake.

Figure 21 shows the transverse momentum spread when the laser is p polarized (black curve) and s polarized (red curve) for (a) linear interpolation and (b) quadratic interpolation. When the laser is p polarized, it triggers a higher transverse momentum spread immediately behind the laser, due to the particle push errors. On the contrary, the longitudinal momentum spread will not depend on the laser polarization. The discrepancy between p and s polarization is reduced when using smoother interpolation [Fig. 21(b)] as the macroparticles experience smaller push errors inside the laser.

Figures 22 and 23 show the longitudinal and transverse momentum spread as a function of resolution with linear and quadratic interpolation, respectively, when the laser is p polarized. The general trend is that higher resolution reduces the unphysical momentum spread. Both the longitudinal and transverse momentum spreads are a function of both Δx and Δz ; the momentum spread is driven by the largest grid size




FIG. 22. (Color online) Longitudinal (a) and transverse (b) momentum spread (smoothed over 12 cells) with $N_{\text{PPC}}=4$ and using linear interpolation; for the resolution $\Delta x=\lambda_0/2.5$ and $\Delta z=\lambda_0/24$ (black curve), $\Delta x=\lambda_0/24$ and $\Delta z=\lambda_0/24$ [blue (dark gray) dashed curve—changing transverse resolution compare to the black curve], and $\Delta x=\lambda_0/2.5$ and $\Delta z=\lambda_0/48$ [magenta (light gray) curve—changing longitudinal resolution compare to the black curve].

(the transverse resolution Δx in this case), especially when a smoother interpolation function is used. Note that the observed macroparticle trapping is located at the back of the box, where the momentum spread is higher. This implies that both the longitudinal and transverse resolutions need to be of the same order to reduce the numerical heating of the plasma and thus to reduce the resulting spurious trapping.

Figure 24 shows the effects of interpolation functions and number of macroparticles per cell N_{PPC} on the longitudinal momentum spread, using square cells $\Delta x=\Delta z=\lambda_0/24$: (a) linear interpolation with $N_{\text{PPC}}=4$, $N_{\text{PPC}}=16$, and $N_{\text{PPC}}=64$; (b) quadratic interpolation with $N_{\text{PPC}}=4$, $N_{\text{PPC}}=16$, and $N_{\text{PPC}}=64$; and (c) cubic interpolation with $N_{\text{PPC}}=4$, $N_{\text{PPC}}=16$, and $N_{\text{PPC}}=25$. A direct comparison of the effects of interpolation methods is shown in Fig. 25, which shows the longitudinal momentum spread using linear, quadratic, and cubic interpolation schemes with $\Delta x=\lambda_0/24$, $\Delta z=\lambda_0/24$, and $N_{\text{PPC}}=4$. The general trend is that smoother interpolation functions and larger N_{PPC} reduce the unphysical momentum spread.

The growth in the plasma momentum spread is a measure of how much the momenta of the macroparticles are departing from the cold fluid orbits. As discussed above, in addition to displacements in momentum, the macroparticles also develop displacements in phase relative to the cold fluid orbits. When these orbit displacements become sufficiently large, macroparticles can become trapped in the wake. The




FIG. 23. (Color online) Longitudinal (a) and transverse (b) momentum spread (smoothed over 12 cells) with $N_{\text{PPC}}=4$ and using quadratic interpolation; for the resolution $\Delta x=\lambda_0/2.5$ and $\Delta z=\lambda_0/24$ (black curve), $\Delta x=\lambda_0/24$ and $\Delta z=\lambda_0/24$ [blue (dark gray) dashed curve—changing transverse resolution compare to the black curve], and $\Delta x=\lambda_0/2.5$ and $\Delta z=\lambda_0/48$ [magenta (light gray) curve—changing longitudinal resolution compare to the black curve].

number of trapped macroparticles (defined as those with momenta $p_z/mc > 1$, i.e., momenta larger than the maximum of the cold fluid orbit) for a variety of cases is summarized in Figs. 26 and 27, where the number of trapped macroparticles has been normalized by $\Delta x \times \Delta z \times n_0 / N_{\text{PPC}}$, yielding the charge per meter.

Figure 26 shows the amount of trapped charge when using linear interpolation, as a function of resolution and laser polarization. Even though the longitudinal electric field is well converged at $\Delta x=\lambda_0/2.5$ and $\Delta z=\lambda_0/24$ (see Fig. 15), the amount of trapped charge is highly dependent on the numerical parameters. At low resolution there is a large difference depending on whether the laser is p polarized or s polarized. Note that the trend is not consistent and can reverse as the resolution increases. In fact, the polarization dependence vanishes when the resolution is sufficiently high in both transverse and longitudinal directions (for $\Delta x < \lambda_0/24$ and $\Delta z < \lambda_0/24$ in this example).

Figure 27 shows the evolution of trapped charge as a function of resolution for different interpolation functions, when using square cells ($\Delta x=\Delta z$). With linear interpolation, the results seem to converge to a finite value of trapped charge; however, using more macroparticles per cell or smoother interpolation functions decreases again the trapping which indicates that the amount of trapped charge is not converged. The lower trapped charge for cubic interpolation




FIG. 24. (Color online) Longitudinal momentum spread (smoothed over 12 cells) using (a) linear interpolation with $N_{\text{PPC}}=4$ (black curve), $N_{\text{PPC}}=16$ [magenta (light gray) curve], and $N_{\text{PPC}}=64$ [blue (dark gray) dashed curve]; (b) quadratic interpolation with $N_{\text{PPC}}=4$ (black curve), $N_{\text{PPC}}=16$ [magenta (light gray) curve], and $N_{\text{PPC}}=64$ [blue (dark gray) dashed curve]; and (c) cubic interpolation with $N_{\text{PPC}}=4$ (black curve), $N_{\text{PPC}}=16$ [magenta (light gray) curve], and $N_{\text{PPC}}=25$ [blue (dark gray) dashed curve].

and $\Delta x = \Delta z = \lambda_0/18$ may be due to a lower wakefield amplitude which is not yet converged at that resolution. When the resolution is high enough (i.e., Δx and $\Delta z \lesssim \lambda_0/36$), increasing the number of macroparticles or using smoother interpolation can reduce the momentum spread and the amount of unphysical macroparticle trapping. It should be noted that both number of macroparticles per cell and resolution need to be increased to reach convergence [25].

IV. TWO-DIMENSIONAL HIGHLY NONLINEAR REGIME

In this section, PIC simulations are presented of a case illustrating numerical effects on self-trapped bunch properties in a highly nonlinear 2D wake. In this regime the electrons are completely expelled from the region near the first plasma period (referred to as the cavitated [6], bubble [7], or blowout [8] regime). This regime is similar to that of current experiments producing high quality electron bunches [2–4], with typical plasma temperatures of $T \sim 10$ eV [22,23]. The




FIG. 25. (Color online) Longitudinal momentum spread using linear (black curve), quadratic [magenta (light gray) curve], and cubic [blue (dark gray) dashed curve] interpolation algorithms with $\Delta x = \lambda_0/24$, $\Delta z = \lambda_0/24$, and $N_{\text{PPC}}=4$.

initial normalized laser intensity profile is of the form $a_0^2(r_0/r_s)\exp(-2x^2/r_s^2 - 2z^2/L^2)$, with $r_s^2/r_0^2 = 1 + (z - Z_f)^2/Z_R^2$, Z_f is the vacuum focal position, $Z_R = \pi r_0^2/\lambda_0$ the Rayleigh length, $r_0 = w_0 = 5.5 \mu\text{m}$ the spot size at focus, $L = 5 \mu\text{m}$, $\lambda_0 = 0.8 \mu\text{m}$. Also, we take $a_0 = 5$, which gives a peak laser intensity of $I = 5.4 \times 10^{19} \text{W/cm}^2$. The laser is focused at the top of a $200 \mu\text{m}$ cosine ramp in a preformed plasma channel matched (in the low intensity limit) to the laser spot size, with a density on axis of $n_0 = 5 \times 10^{18} \text{cm}^{-3}$ ($\omega_0/\omega_p \approx 19$). The simulation box, comoving with the laser pulse at the speed of light, is $50 \mu\text{m}$ wide and $80 \mu\text{m}$ long, and the macroparticles are loaded uniformly and cold (no initial momentum) with $N_{\text{PPC}}=4$.

Figure 28 shows the density profile after 1 mm of propagation ($\omega_p t = 534$), for (a) $\Delta x = \lambda_0/2.4 = r_0/16.5$ and $\Delta z = \lambda_0/31.4$ and (b) $\Delta x = \lambda_0/4.8 = r_0/33$ and $\Delta z = \lambda_0/62.8$. A cavitating, bubble, region is seen directly after the laser pulse and electron bunches are trapped and accelerated inside the bubble. In Fig. 28(a) two electron bunches (spatially separated) are trapped in the wake. However, when the resolution is increased [Fig. 28(b)] only the first electron bunch remains, indicating that the second electron bunch is due to artificial trapping in the wake due to the low resolution.

Figure 29 shows orbits of selected macroparticles when the laser is s polarized [Fig. 29(a)] and p polarized [Fig.

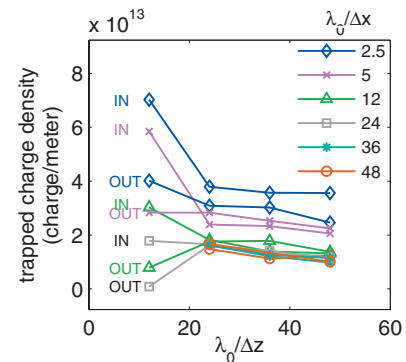


FIG. 26. (Color) Amount of trapped charged ($p_z/m_e c > 1$) multiplied by $\Delta x \times \Delta z \times n_0/N_{\text{PPC}}$ (in charge/meter) as a function of longitudinal resolution with linear interpolation and $N_{\text{PPC}}=4$, for different transverse resolutions ($\Delta x = \lambda_0/2.5 - \lambda_0/48$), with p polarization (in) and s polarization (out).




FIG. 27. Amount of trapped charge ($p_z/m_e c > 1$) versus resolution, with $\Delta x = \Delta z$, for different interpolation algorithms: linear interpolation (square), quadratic interpolation (circle), and cubic interpolation (triangle). Full and empty points are for 4 and 16 macroparticles per cell, respectively. The laser is p polarized.

29(b)] using linear interpolation, and when the laser is p polarized using quadratic interpolation [Fig. 29(c)]. The three pictures are very similar at early $\omega_p t$ when the bubble is formed, indicating that the particle push error in the laser field observed in Sec. III [Figs. 18 and 17(a)] is not as important in this case as the macroparticles are rapidly expelled transversely. However, the orbits seem to have a different behavior at later times, indicating that the macroparticles are subject to different values of numerical errors. For these numerical parameters the hydrodynamics (e.g., wake formation) is well resolved, but the kinetic effects (e.g., numerical heating and trapping) are not converged.




FIG. 28. Electron density at $\omega_p t = 534$ for (a) $\Delta x = \lambda_0/2.4 = r_0/16.5$ and $\Delta z = \lambda_0/31.4$ and (b) $\Delta x = \lambda_0/4.8 = r_0/33$ and $\Delta z = \lambda_0/62.8$, and $N_{ppc} = 4$




FIG. 29. Orbits of selected macroparticles for $\Delta x = \lambda_0/2.4$ and $\Delta z = \lambda_0/31.4$, using linear interpolation and with (a) s polarization, (b) p polarization, and (c) using quadratic interpolation with p polarization.

Figure 30 shows macroparticle phase space for $\Delta x = \lambda_0/4.8$, $\Delta z = \lambda_0/48$, and s polarization with different interpolation functions, illustrating the effects of interpolation errors on the trapped bunch. The properties of the trapped bunch vary only slightly with interpolation scheme, in Figs. 30(a) and 30(b): the charge decreases by 10%, the mean energy increases by 1%, and the energy spread varies from 3% (r.m.s.) to 2.6%. Changing from quadratic to cubic interpolation [Fig. 30(c)] the properties of the first trapped bunch vary even less, compared to quadratic interpolation: 0.6 and 0.1% fluctuation in charge and mean energy, respectively, and no significant change in energy spread. Results are similar for a p polarized laser ($\sim 1\%$ variation for the mean energy and $\sim 10\%$ variation for the charge).

In this case ($\Delta x = \lambda_0/4.8$ and $\Delta z = \lambda_0/48$), the on-axis longitudinal momentum spread, behind the cavitated region ($510 < k_{pz} < 515$), is $\sigma_{u_z}^2 \sim 10^{-4}$ for linear interpolation and $\sim 10^{-6}$ for cubic interpolation. Near the transverse edge of the bubble ($k_{px} = 3.3$ and $519 < k_{pz} < 524$) the longitudinal momentum spread fluctuates between 10^{-4} and 10^{-1} , on a period much smaller than the plasma period, the mean value being $\sigma_{u_z}^2 = 0.4 \times 10^{-2}$. These values, near the edge of the bubble, are found not to change with resolution or interpola-




FIG. 30. Macroparticle phase space for the resolution $\Delta x = \lambda_0/4.8$ and $\Delta z = \lambda_0/48$ for (a) linear interpolation, (b) quadratic interpolation, and (c) cubic interpolation. The laser is s polarized.

tion method and may be dominated by the trajectory crossing next to the blowout region (see Fig. 29).

When increasing the transverse resolution, using linear interpolation and $\Delta z = \lambda_0/31$, more charge is trapped in both the first and the second bunch whereas the mean energies decrease, which is consistent with beam loading. This results in a 20% difference in charge and 10% difference in mean energy for the first trapped bunch between $\Delta x = \lambda_0/4.8$ and $\Delta x = \lambda_0/18$, and a factor of 2 difference in energy spread. The change in charge and energy spread is of the per cent level between $\Delta x = \lambda_0/18$ and $\Delta x = \lambda_0/31.4$, but the mean energy still varies by 10% between these two cases although this can be explained by a decrease of the laser group velocity when increasing the transverse resolution [42]. This indicates that the transverse resolution must be sufficiently high to well resolve the density and field structures, and to accurately model the trapping physics.

In the case of the highly nonlinear regime discussed in this section (where we expect to be significantly above the trapping threshold for physically relevant plasma temperatures), results are less sensitive to resolution than would be the case for less nonlinear cases closer to threshold (e.g., cases presented in Sec. III where the trapped charge can vary by a factor of 4 with resolution and by almost an order of magnitude with different interpolation functions) although one still typically finds artificial trapping due to numerical heating (as noted in Fig. 28). Note also that the bunch trapped in the cavity is located at a short distance behind the laser field. As shown in the previous section, the artificial

momentum spread of the plasma grows as a function of distance behind the laser pulse, and artificial kinetic effects will be more important at the back of the simulation box. Also, in the highly nonlinear case, the macroparticle orbits are not as sensitive to the laser polarization (in or out of the simulation plane), since the macroparticles are rapidly expelled transversely from the laser field.

V. CONCLUSION

In this work, the effects of numerical heating on the simulation of laser wakefield accelerators using PIC codes have been investigated. Three cases were examined: a dark current free LWFA stage in 1D, a dark current free LWFA stage in 2D, and a highly nonlinear wake with a self-trapped electron bunch in the 2D blowout regime. The parameters used in these cases are typical of those in present (blowout regime with self-trapping [2–4]) and future (dark current free accelerating stage with external injection [17]) experiments. In general, in order to properly model kinetic phenomena, such as self-trapping of electrons, very high numerical resolution is required in comparison to that required to model fluid phenomena, such as the structure of the electric field of the wake, which depends on the bulk properties (e.g., moments) of the macroparticle distribution. Numerical errors in PIC codes can lead to artificial plasma heating, as measured by the momentum spread of the macroparticles. This heating grows as a function of the distance behind the laser pulse and can result in nonphysical trapping of macroparticles in the wake, thereby resulting in an artificial dark current or artificially high levels of charge in the trapped bunch with lower mean energy.

Studies were presented in Sec. II for a LWFA in 1D with an initially cold plasma. For the parameters in Sec. II (e.g., $a_0 = 2$), the wakefield amplitude ($E_z = 1.1E_0$) is well below the cold, relativistic wakefield amplitude ($E_{WB} = 4.2E_0$). Since the Vlasov equation predicts that a cold plasma should remain cold (a δ function momentum distribution is an exact solution of Vlasov), and since the macroparticles were loaded cold on a uniform grid in the PIC simulation, there should be no physical trapping for this case, i.e., dark current free. The PIC simulations show, however, that the plasma momentum spread increases rapidly as a function of distance behind the drive laser pulse. When compared to conventional numerical self-heating of an initially warm loading of macroparticles in the absence of a drive laser (see Fig. 1), the observed numerical heating in the laser driven case is much more rapid and reaches momentum spreads far exceeding the nondriven, self-heating case. For example, for $\Delta z = \lambda_p/240$ ($\lambda_p = 8 \mu\text{m}$), nondriven numerical self-heating saturates at a level of $k_g \lambda_D \sim 1$ ($T \sim 10$ eV) over a time $\omega_p t \sim 1000$; whereas the laser driven case reaches temperatures in excess of $T \sim 1$ keV over a time $\omega_p t \sim 100$ (see Fig. 1 and 4).

Examination of macroparticle phase space in the 1D PIC simulations (see Figs. 6–8) indicates that phase-space develops a complex structure behind the laser driver. The macroparticle orbits develop errors in both momentum and position relative to the cold fluid orbits. The height of this phase-space structure is a measure of the plasma momentum

spread. These errors, and hence the momentum spread, grow as a function of distance behind the drive laser pulse. Eventually, the orbit errors of some macroparticles are large enough to allow trapping in the wakefield.

In the 1D PIC simulations, the numerical heating is reduced by increasing the resolution, and to a lesser extent by increasing the number of macroparticles per cell. Using smoother interpolation functions is also effective for reducing the numerical heating. Current smoothing is found to be effective in reducing the heating only in the region immediately behind the laser pulse (first few plasma periods). The amount of macroparticle trapping did not always correlate directly with a reduction in artificial momentum spread resulting from using higher resolution or smoother interpolation functions. This is partly due to the fact that although higher resolution or smoother interpolation reduce the momentum spread in the near and intermediate regions behind the laser pulse, the momentum spread near the back of the box, where trapping occurs most readily, reaches similar levels. Furthermore, trapping can reduce the wakefield amplitude via beam loading, which can affect the total charge trapped in the wake and hence the mean energy.

PIC simulations in 2D for a dark current free LWFA stage (with $a_0=1.15$), with parameters typically used in the design of the next generation of experiments, were presented in Sec. III. This included comparisons between in plane versus out-plane polarization. By examining the macroparticle orbits, it is found that in-plane polarization, in which the macroparticles experience strong oscillations in the laser field, leads to larger orbit errors compared to out of plane polarization. This results in a much higher transverse momentum spread; however, the longitudinal momentum spread is comparable. This discrepancy in transverse momentum spread versus polarization is diminished for higher resolution or smoother interpolation functions as the interpolation errors are reduced. Many of the general trends observed in 1D carry over to 2D. Smoother interpolation functions reduce momentum spread. Decreasing both longitudinal (Δz) and transverse (Δx) resolution reduce the momentum spread, however, the momentum spread is driven by the largest grid size (typically Δx). This means that both longitudinal and transverse resolutions need to be of the same order to reduce the numerical plasma heating and the resulting spurious trapping.

The total amount of trapped charge in the 2D wake is not always reduced by increasing the resolution or using smoother interpolation functions. As discussed above, this can be due to the momentum spread rising quickly at the back of the simulation box, leading to trapping, even though the momentum spread may be greatly reduced in the front half of the simulation box. It may also be due to the effects of beam loading reducing the wake amplitude and reducing further trapping. In the case of linear interpolation, trapping seems to saturate when the resolution becomes sufficiently large ($\Delta x, \Delta z \lesssim \lambda_0/36$), i.e., trapping cannot be reduced more by increasing the resolution alone; more macroparticles per cell may be required. However, it is more efficient to use smoother interpolation schemes in terms of computational cost. First, smoother interpolation requires fewer grid points to significantly reduce the truncation errors in the particle push, as shown in Fig. 18. Second, at a given resolution, it is

less expensive to use a smoother scheme than increasing the number of macroparticles to reduce the spurious trapping by the same amount. For example, with the resolution $\Delta x = \Delta z = \lambda_0/24$ the same amount of trapped charge is obtained when using linear interpolation and 64 macroparticles per cell as when using quadratic interpolation and four macroparticles per cell. Using quadratic interpolation instead of linear is only 25% more computationally expensive, which would correspond to increasing the number of macroparticles per cell from 4 to 5 in the linear case.

Lastly, a highly nonlinear wake in the blowout regime ($a_0=5$) with a self-trapped bunch, similar to the regime of recent experiments, was simulated in 2D. Plots of the resulting phase-space of the macroparticles showed that the variations in the trapped bunch properties were at the % level for $\Delta z \sim \lambda_0/31$ and $\Delta x \sim \lambda_0/18$. In this case, macroparticle trapping occurs in the first plasma wave bucket immediately behind the laser pulse where unphysical plasma heating is significantly less important than after a few plasma periods.

In summary, in this paper it has been shown that unphysical kinetic effects in the PIC simulation of a laser wakefield accelerators can result in heating of the plasma to high momentum spread (i.e., temperature) and in artificial self-trapping. Ideally, modeling of kinetic effects in intense, short pulse, laser-plasma interactions using the PIC algorithm would require the initialization of a warm plasma (typically a few tens of eV for photoionized plasmas) and use of sufficiently high resolution such that numerical heating is well below the physical temperature (see Fig. 13). In practice, this is computationally beyond the state of the art for modeling LWFAs at low plasma temperatures in 3D using conventional linearly weighted PIC algorithms. It has been demonstrated, however, that the heating and spurious trapping can be reduced, especially in the region near the back of the laser pulse. This can be accomplished by increasing resolution, number of macroparticles per cell, and using smoother interpolation functions. In particular we have shown the effectiveness of the latter for reducing the growth of momentum spread at reasonable computational cost. These methods will help improve the modeling and the design of future laser wakefield accelerator experiments.

ACKNOWLEDGMENTS

The authors would like to thank Hartmut Ruhl for kindly providing the PIC code PSC and for related discussions. The authors also thank D. Bruhwiler, J. R. Cary, W. B. Mori, and J. L. Vay for many insightful conversations. This work was supported by the University of Nevada, Reno, Grant No. DE-FC52-01NV14050, and by the Director, Office of Science, Office of High Energy Physics, of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231. This work was also supported by a Scientific Discovery through Advanced Computing project, "Advanced Computing for 21st Century Accelerator Science and Technology," and an INCITE grant from the National Energy Research Scientific Computing Center (NERSC), which are supported by the U.S. DOE/SC Office of High Energy Physics and the Office of Advanced Scientific Computing Research.

APPENDIX: MACROPARTICLE WEIGHTING AND SMOOTHING

In the PIC algorithm it is necessary to calculate the charge density (and current density) on a discrete grid from the ungridded particle positions. This is done by an interpolation, or weighting, of the particle location to the numerical grid points. Because the PIC algorithm uses a finite number of macroparticles, spatially localized macroparticles can lead to discontinuities in the numerical plasma density and current. These discontinuities can be quite large when, as is typical practice, the number of macroparticles per cell is small.

The weighting of macroparticle quantities to the spatial grid is not unique; different weighting schemes can be used, which are interpreted as a “shape” for the macroparticles. Specializing to one dimension, the charge density $\rho_j = \rho(X_j)$, at grid point $X_j = j\Delta x$, is obtained from the macroparticle positions x_i by

$$\rho_j = \sum_i q_i S(X_j - x_i), \quad (\text{A1})$$

where q_i is the macroparticle charge, $S(x)$ is the macroparticle shape (weighting) function, and the sum extends over all macroparticles. It is also necessary to determine the force acting on the macroparticle in terms of the gridded fields. To ensure momentum conservation and to avoid possible instabilities associated with self-forces [27], it is necessary to use the same shape function when computing the density and the forces. The force F_i on a macroparticle is obtained by interpolating the electromagnetic fields \mathbf{E} and \mathbf{B} to the macroparticle positions \mathbf{E}_i and \mathbf{B}_i , respectively, by

$$\mathbf{E}_i = \sum_j \mathbf{E}(X_j) S(X_j - x_i), \quad (\text{A2})$$

$$\mathbf{B}_i = \sum_j \mathbf{B}(X_j) S(X_j - x_i), \quad (\text{A3})$$

$$\mathbf{F}_i = q_i(\mathbf{E}_i + \mathbf{v}_i \times \mathbf{B}_i/c), \quad (\text{A4})$$

where \mathbf{v}_i is the macroparticle velocity. For any shape function, the current density at the mesh points is obtained in a way that results in exact charge conservation; see Refs. [33,43].

For the simplest shape function S_0 takes the macroparticle to have uniform density and is the size of a spatial grid cell [28,30,44]

$$S_0(x) = \begin{cases} 1, & \text{for } |x| \leq \frac{1}{2}\Delta x, \\ 0, & \text{otherwise.} \end{cases} \quad (\text{A5})$$

This weighting scheme is known as nearest-grid-point (NGP). NGP is rarely used in practice because it leads to high levels of noise in the gridded quantities. More typically, PIC codes (including PSC) use linear interpolation, also known as cloud-in-cell [45], where the shape function is given by

$$S_1(x) = \begin{cases} 1 - \frac{|x|}{\Delta x}, & \text{for } |x| \leq \Delta x, \\ 0, & \text{otherwise.} \end{cases} \quad (\text{A6})$$

Higher-order interpolation schemes are obtained by an m (m is the order number) times convolution (denoted by an asterisk) of the NGP weighting factor with itself. Hence, the second order, or quadratic, weighting function is piecewise parabolic and is given by $S_2(x) = (S_0 * S_0 * S_0)(x)$, i.e.,

$$S_2(x) = \begin{cases} \frac{3}{4} - \frac{|x|^2}{\Delta x^2}, & \text{for } |x| \leq \frac{1}{2}\Delta x, \\ \frac{1}{2} \left(\frac{3}{2} - \frac{|x|}{\Delta x} \right)^2, & \text{for } \frac{1}{2}\Delta x \leq |x| \leq \frac{3}{2}\Delta x, \\ 0, & \text{otherwise.} \end{cases} \quad (\text{A7})$$

The third order, or cubic, weighting function is piecewise cubic, i.e.,

$$S_3(x) = \begin{cases} \frac{2}{3} - \frac{|x|^2}{\Delta x^2} + \frac{|x|^3}{2\Delta x^3}, & \text{for } |x| \leq \Delta x, \\ \frac{1}{6} \left(2 - \frac{|x|}{\Delta x} \right)^3, & \text{for } \Delta x \leq |x| \leq 2\Delta x, \\ 0, & \text{otherwise.} \end{cases} \quad (\text{A8})$$

Notice that for successive shape functions, the width of the weighting function increases by Δx and the interpolated charge density has its order of continuity increased by 1: S_1 leads to a density that is globally C^0 , while S_2 leads to a density that is globally C^1 , etc. The order of the interpolation is related to the smoothness of the resulting density, not the accuracy. The interpolation error within a cell is smoother with higher order weighting scheme. That is, using higher order (or smoother) weighting functions does not improve the order of the numerical solver, which, in a typical PIC code, is second order in the grid size and time step. In fact, using higher-order functions implies using cells beyond the nearest grid point and therefore lower resolution and smoothing of the results. In the present work, higher-order weighting functions are referred to as smoother interpolation functions (or smoother particle shapes) to avoid confusion with a higher-order accuracy solver. This weighting scheme can be readily generalized to any dimension by using the same shape function in each direction. For example, in three dimensions, we have $S_m(x, y, z) = S_m(x)S_m(y)S_m(z)$.

Another method for suppressing the high frequency components (on the order of k_g), and hence reducing noise, is to smooth the current and/or fields via a filter [38]. In PSC a digital filter can optionally be used to smooth the current. Specifically (see Appendix C of Ref. [28]), the quantity J_j is replaced by

$$J_j = \frac{WJ_{j-1} + J_j + WJ_{j+1}}{1 + 2W}, \quad (\text{A9})$$

with $W=0.5$, i.e., $J_j=(J_{j-1}+2J_j+J_{j+1})/4$, often referred to as a (1, 2, 1) filter. This digital filter is applied four times, each

time the filter is first applied to the longitudinal direction and then to the transverse direction in 2D. Lastly, a filter with $W=-5/14$ (i.e., $J_j=-5J_{j-1}/4+7J_j/2-5J_{j+1}/4$) is applied once to compensate for the roll-off of the low frequencies caused by the previous filtering.

-
- [1] E. Esarey, P. Sprangle, J. Krall, and A. Ting, *IEEE Trans. Plasma Sci.* **24**, 252 (1996).
- [2] C. G. R. Geddes, C. Toth, J. van Tilborg, E. Esarey, C. B. Schroeder, D. Bruhwiler, C. Nieter, J. Cary, and W. P. Leemans, *Nature (London)* **431**, 538 (2004).
- [3] S. P. D. Mangles, C. D. Murphy, Z. Najmudin, A. G. R. Thomas, J. L. Collier, A. E. Dangor, E. J. Divall, P. S. Foster, J. G. Gallacher, C. J. Hooker *et al.*, *Nature (London)* **431**, 535 (2004).
- [4] J. Faure, Y. Glinec, A. Pukhov, S. Kiselev, S. Gordienko, E. Lefebvre, J.-P. Rousseau, F. Burgy, and V. Malka, *Nature (London)* **431**, 541 (2004).
- [5] W. P. Leemans, B. Nagler, A. J. Gonsalves, C. Tóth, K. Nakamura, C. G. R. Geddes, E. Esarey, C. B. Schroeder, and S. M. Hooker, *Nat. Phys.* **2**, 696 (2006).
- [6] P. Mora and T. M. Antonsen, Jr., *Phys. Rev. E* **53**, R2068 (1996).
- [7] A. Pukhov and J. Meyer-ter-Vehn, *Appl. Phys. B: Lasers Opt.* **74**, 355 (2002).
- [8] W. Lu, C. Huang, M. Zhou, M. Tzoufras, F. S. Tsung, W. B. Mori, and T. Katsouleas, *Phys. Plasmas* **13**, 056709 (2006).
- [9] C. G. R. Geddes, C. Tóth, J. van Tilborg, E. Esarey, C. B. Schroeder, D. Bruhwiler, C. Nieter, J. R. Cary, and W. P. Leemans, *Phys. Plasmas* **12**, 056709 (2005).
- [10] D. Umstadter, J. K. Kim, and E. Dodd, *Phys. Rev. Lett.* **76**, 2073 (1996).
- [11] E. Esarey, R. F. Hubbard, W. P. Leemans, A. Ting, and P. Sprangle, *Phys. Rev. Lett.* **79**, 2682 (1997).
- [12] C. B. Schroeder, P. B. Lee, J. S. Wurtele, E. Esarey, and W. P. Leemans, *Phys. Rev. E* **59**, 6037 (1999).
- [13] G. Fubiani, E. Esarey, C. B. Schroeder, and W. P. Leemans, *Phys. Rev. E* **70**, 016402 (2004).
- [14] H. Kotaki, S. Masuda, M. Kando, J. K. Koga, and K. Nakajima, *Phys. Plasmas* **11**, 3296 (2004).
- [15] J. Faure, C. Rechatin, A. Norlin, A. Lifschitz, Y. Glinec, and V. Malka, *Nature (London)* **444**, 737 (2006).
- [16] C. G. R. Geddes, K. Nakamura, G. R. Plateau, C. Tóth, E. Cormier-Michel, E. Esarey, C. B. Schroeder, J. R. Cary, and W. P. Leemans, *Phys. Rev. Lett.* **100**, 215004 (2008).
- [17] A. J. W. Reitsma, W. P. Leemans, E. Esarey, C. B. Schroeder, L. P. J. Kamp, and T. J. Schep, *Phys. Rev. ST Accel. Beams* **5**, 051301 (2002).
- [18] A. I. Akhiezer and R. V. Polovin, *Zh. Eksp. Teor. Fiz.* **30**, 915 (1956); *Sov. Phys. JETP* **3**, 696 (1956).
- [19] C. B. Schroeder, E. Esarey, and B. A. Shadwick, *Phys. Rev. E* **72**, 055401(R) (2005).
- [20] C. B. Schroeder, E. Esarey, B. A. Shadwick, and W. P. Leemans, *Phys. Plasmas* **13**, 033103 (2006).
- [21] E. Esarey, C. B. Schroeder, E. Cormier-Michel, B. A. Shadwick, C. G. R. Geddes, and W. P. Leemans, *Phys. Plasmas* **14**, 056707 (2007).
- [22] C. G. Durfee, J. Lynch, and H. M. Milchberg, *Phys. Rev. E* **51**, 2368 (1995).
- [23] P. Volfbeyn, E. Esarey, and W. Leemans, *Phys. Plasmas* **6**, 2269 (1999).
- [24] B. A. Shadwick, G. M. Tarkenton, and E. H. Esarey, *Phys. Rev. Lett.* **93**, 175002 (2004).
- [25] B. A. Shadwick, G. M. Tarkenton, E. H. Esarey, and C. B. Schroeder, *Phys. Plasmas* **12**, 056710 (2005).
- [26] J. M. Dawson, *Rev. Mod. Phys.* **55**, 403 (1983).
- [27] R. Hockney and J. Eastwood, *Computer Simulation using Particles* (Taylor & Francis, New York, 1988).
- [28] C. K. Birdsall, A. B. Langdon, V. Vehedi, and J. P. Verboncoeur, *Plasma Physics via Computer Simulations* (Adam Hilger, Bristol, 1991).
- [29] H. Ruhl, *Habilitationschrift* (Technische Universität Darmstadt, Darmstadt, 2000).
- [30] R. W. Hockney, *J. Comput. Phys.* **8**, 19 (1971).
- [31] A. B. Langdon, *J. Comput. Phys.* **6**, 247 (1970).
- [32] J. P. Boris, in *Proceedings, Fourth Conference on the Numerical Simulation of Plasma* (Naval Res. Lab., Washington D.C., 1970), pp. 3–67.
- [33] T. Z. Esirkepov, *Comput. Phys. Commun.* **135**, 144 (2001).
- [34] R. A. Fonseca *et al.*, in *ICCS 02: Proceedings of the International Conference on Computational Science—Part III*, edited by P. M. A. Sloot *et al.*, Vol. 2331 of *Lecture Notes In Computer Science* (Springer-Verlag, Heidelberg, 2002), pp. 342–351.
- [35] C. Nieter and J. R. Cary, *J. Comput. Phys.* **196**, 448 (2004).
- [36] J. P. Verboncoeur, A. B. Langdon, and N. T. Gladd, *Comput. Phys. Commun.* **87**, 199 (1995).
- [37] E. Esarey and M. Pilloff, *Phys. Plasmas* **2**, 1432 (1995).
- [38] W. B. Mori (private communication).
- [39] J. M. Dawson, *Phys. Rev.* **118**, 381 (1960).
- [40] S. V. Bulanov, F. Pegoraro, A. M. Pukhov, and A. S. Sakharov, *Phys. Rev. Lett.* **78**, 4205 (1997).
- [41] B. A. Shadwick, G. M. Tarkenton, E. H. Esarey, and W. P. Leemans, *IEEE Trans. Plasma Sci.* **30**, 38 (2002).
- [42] C. G. R. Geddes, Ph.D. thesis, University of California, Berkeley, 2005.
- [43] J. Villasenor and O. Buneman, *Comput. Phys. Commun.* **69**, 306 (1992).
- [44] H. Abe, N. Sakairi, R. Itatani, and H. Okuda, *J. Comput. Phys.* **63**, 247 (1986).
- [45] C. K. Birdsall and D. Fuss, *J. Comput. Phys.* **3**, 494 (1969).