Journal of Computational Physics 230 (2011) 5908–5929

Contents lists available at ScienceDirect

Journal of Computational Physics

journal homepage: www.elsevier.com/locate/jcp

Numerical methods for instability mitigation in the modeling of laser wakefield accelerators in a Lorentz-boosted frame

J.-L. Vay [a,*], C.G.R. Geddes [a], E. Cormier-Michel [b], D.P. Grote [c]

[a] Lawrence Berkeley National Lab., 1 Cyclotron Road, Berkeley, CA 94720, USA
[b] Tech-X Corporation, 5621 Arapahoe Ave., Suite A, Boulder, CO, 80303, USA
[c] Lawrence Livermore National Lab., L-637, Livermore, CA 94550, USA

ARTICLE INFO

Article history:
Received 18 November 2010
Received in revised form 1 April 2011
Accepted 5 April 2011
Available online 12 April 2011

Keywords:
Laser wakefield acceleration
Particle-in-cell
Plasma simulation
Special relativity
Boosted frame
Numerical instability

ABSTRACT

Modeling of laser-plasma wakefield accelerators in an optimal frame of reference [1] has been shown to produce orders of magnitude speed-up of calculations from first principles. Obtaining these speedups required mitigation of a high-frequency instability that otherwise limits effectiveness. In this paper, methods are presented which mitigated the observed instability, including an electromagnetic solver with tunable coefficients, its extension to accommodate Perfectly Matched Layers and Friedman's damping algorithms, as well as an efficient large bandwidth digital filter. It is observed that choosing the frame of the wake as the frame of reference allows for higher levels of filtering or damping than is possible in other frames for the same accuracy. Detailed testing also revealed the existence of a singular time step at which the instability level is minimized, independently of numerical dispersion. A combination of the techniques presented in this paper prove to be very efficient at controlling the instability, allowing for efficient direct modeling of 10 GeV class laser plasma accelerator stages. The methods developed in this paper may have broader application, to other Lorentz-boosted simulations and Particle-In-Cell simulations in general.

© 2011 Elsevier Inc. All rights reserved.

1. Introduction

Laser driven plasma waves offer orders of magnitude increases in accelerating gradient over standard accelerating structures [2] (which are limited by electrical breakdown), thus holding the promise of much shorter particle accelerators [3]. In a laser plasma accelerator (LPA), a laser propagates through a plasma, displacing electrons while ions remain essentially static, thus creating a pocket of positive charges that the displaced electrons rush to fill. The resulting coherent periodic motion of the electrons oscillating around their original position creates a wake with periodic structure following the laser. The alternate concentration of positive and negative charges in the wake creates very intense electric fields. An electron (or positron) beam injected with the right phase can be accelerated by those fields to high energy in a much shorter distance than is possible in conventional particle accelerators. High quality electron beams of energy up-to 1 GeV have been produced in just a few centimeters [4–7], with 10 GeV stages being planned as modules of a high energy collider [8]. The efficiency and quality of the acceleration is governed by several factors which require precise three-dimensional shaping of the plasma column, as well as the laser and particle beams, and understanding of their evolution.

Recently, an approach has been demonstrated for speeding up by orders of magnitude the modeling of LPA from first principle. The speedup is provided by a dramatic reduction of the number of time steps, thanks to properties of special relativity

* Corresponding author. Tel.: +1 510 486 4934; fax: +1 510 486 5392.
E-mail address: jlvay@lbl.gov (J.-L. Vay).

0021-9991/$ - see front matter © 2011 Elsevier Inc. All rights reserved.
doi:10.1016/j.jcp.2011.04.003

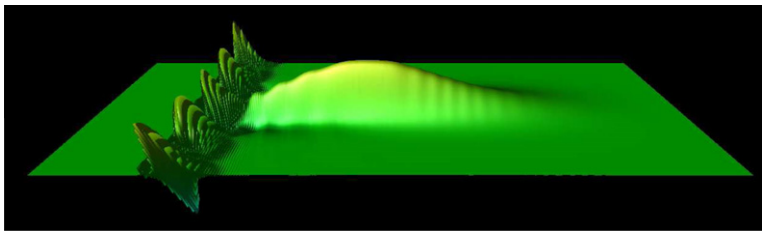

Fig. 1. Surface plot of the longitudinal field from a 2–1/2D simulation of a full scale 10 GeV LPA in a boosted frame at $\gamma = 130$ (elevation/color are proportional to the magnitude of the electric field). The laser is propagating from left to right and the plasma from right to left. The plasma wake (physical feature) is the yellow hump near the center, and as explained in [34], in this frame (near the frame of the wake, $\gamma \approx 132$), the laser manifests itself as a fast temporal beating at long wavelength. A fast growing short wavelength instability is visible developing at the front (left side) of the plasma. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

that apply when choosing an optimal relativistically boosted frame of reference [1]. Early successes [9–17] have been limited to modest relativistic boosts (and thus speedups) that fall below the optimal value [18] due to the development of a short wavelength instability at the front of the plasma (see Fig. 1), as reported in [14,15,18]. The presence and growth rate of the instability was observed to be very sensitive to the resolution (slower growth rate at higher resolution), choice of field solver, and to the amount of damping of high frequencies and smoothing of short wavelengths. The instability is always propagating at some angle from the longitudinal axis, and is observed in 2D and 3D runs but was never observed in any of the 1D runs performed by the authors.

Due to spatial and time discretization of the Maxwell equations, numerical light waves may travel faster or slower on the computational grid than the actual speed of light in vacuum c, with the magnitude of the effect being larger at short wavelength, where discretization errors are the largest. When the numerical speed is lower than c, it is possible for fast macro-particles to travel faster than the wave modes, leading to numerical Cerenkov effects that may result in instabilities [19–23]. The effect was studied analytically and numerically in detail for one-dimensional systems in [21,22]. When modeling an LPA setup in a relativistically boosted frame, the background plasma is traveling near the speed of light and it has been conjectured [15] that the observed instability might be caused by numerical Cerenkov. We investigate in this paper whether the instability that is observed in boosted frame simulations of LPA is indeed of numerical Cerenkov type and if the cures aimed at mitigating numerical Cerenkov are effective.

In this paper, we present in Section 2 numerical techniques that were implemented in the Particle-In-Cell code Warp [24] for mitigating the short wavelength instability, including an efficient numerical filter and a solver for advancing the electromagnetic fields with tunable coefficients. The new techniques are applied to the modeling of full scale and downscaled 10 GeV class LPA stages in Section 3, with special emphasis on the characteristics, and methods for control, of the short wavelength instability. Sensitivity of the instability to numerical parameters is analyzed in Section 4. In [25], we show that the techniques and analysis developed in the present paper allow for the first time calculations at the maximal theoretical frame boost, hence maximal speedup, for LPA stages in the 10 GeV–1 TeV range.

Although this is not developed in this paper, the techniques that are presented below may have a broader range of applications, to other boosted frame simulations of free electron lasers [26,27] or coherent synchrotron radiation [28], as well as to Particle-In-Cell (PIC) simulations in general, where an electromagnetic solver with tunable numerical dispersion and/or efficient filtering may be of interest.

2. Novel numerical techniques for the mitigation of short wavelength instabilities

Several solutions have been proposed for mitigating numerical Cerenkov effects and short wavelength instabilities in general: smoothing the current deposited by the macro-particles [19,21], damping the electromagnetic field [23,29,30], solving the Maxwell equations in Fourier space [20], or using a field solver with a larger stencil to provide lower numerical dispersion [23]. Several of these techniques have been implemented in Warp. All the simulations presented in this paper employed cubic splines for current deposition and electromagnetic force gathering between the macro-particles and the grid [31], whose beneficial effects on standard LPA PIC simulations have been demonstrated in [32]. In addition, efficient filtering of the deposited current (and gathered electromagnetic fields, eventually) was implemented, as well as a Maxwell solver with tunable coefficients, and a damping scheme, which are described in this section. The effects from each on LPA simulations is shown in Section 3. The use of Fourier based Maxwell solvers is not considered in this paper.

2.1. Wideband lowpass digital filtering

It is common practice to apply digital filtering to the charge or current density in Particle-In-Cell simulations, for smoothing purposes [33]. As explained in [34] and as verified in Section 3, LPA simulations performed in the frame of the wake allow for wideband filtering, from the Nyquist wavelength up-to 10s of cells, as the physics that needs to be resolved is at longer

wavelength. Here, we characterize effects of such filters and show that enhanced efficiency is obtained by using "strided" filtering.

A commonly used filter in PIC simulations is the three points filter

$$\phi_j^f = \alpha\phi_j + (1-\alpha)\frac{\phi_{j-1}+\phi_{j+1}}{2} \tag{1}$$

where ϕ^f is the filtered quantity. This filter is called a binomial filter when $\alpha = 0.5$. Assuming $\phi = e^{ikx}$ and $\phi^f = g(\alpha,k)e^{ikx}$, where g is the filter gain, which is a function of the filtering coefficient α and the wavenumber k, we find from (1) that

$$g(\alpha,k) = \alpha + (1-\alpha)\cos(k\Delta x) \tag{2}$$

$$\approx 1 - (1-\alpha)\frac{(k\Delta x)^2}{2} + O(k^4) \tag{3}$$

For n successive applications of filters of coefficients $\alpha_1,\ldots,\alpha_n$, the total attenuation G is given by

$$G = \prod_{i=1}^{n} g(\alpha_i,k) \tag{4}$$

$$\approx 1 - \left(n - \sum_{i=1}^{n}\alpha_i\right)\frac{(k\Delta x)^2}{2} + O(k^4) \tag{5}$$

A sharper cutoff in k space is provided by using $\alpha_n = n - \sum_{i=1}^{n-1}\alpha_i$ so that $G \approx 1 + O(k^4)$. Such a step is called a *compensation* step [33]. For the bilinear filter ($\alpha = 1/2$), the compensation factor is $\alpha_c = 2 - 1/2 = 3/2$. For a succession of n applications of the bilinear factor, it is $\alpha_c = n/2 + 1$. The gain versus wavelength is plotted in Fig. 2 for the bilinear filter without compensation ($G = g(1/2,k)$), with compensation ($G = g(1/2,k)\cdot g(3/2,k)$), and for n-pass bilinear filters with compensation ($G = g(1/2,k)^n \cdot g(n/2+1,k)$) for $n = 4, 20, 50$ and 80.

The bilinear filter provides complete suppression of the signal at the grid Nyquist wavelength (twice the grid cell size). Suppression of the signal at integers of the Nyquist wavelength can be obtained by using a stride s in the filter

$$\phi_j^f = \alpha\phi_j + (1-\alpha)\frac{\phi_{j-s}+\phi_{j+s}}{2} \tag{6}$$

for which the gain is given by

$$g(s,\alpha,k) = \alpha + (1-\alpha)\cos(sk\Delta x) \tag{7}$$

$$\approx 1 - (1-\alpha)\frac{(sk\Delta x)^2}{2} + O(k^4) \tag{8}$$

The gain is plotted in Fig. 3 (top) for four passes bilinear filters with compensation ($G = g(s,1/2,k)^4 \cdot g(s,3,k)$) for strides $s = 1$ to 4. For a given stride, the gain is given by the gain of the bilinear filter shifted in k space, with the pole $g = 0$ shifted from $\lambda = 2/\Delta x$ to $\lambda = 2s/\Delta x$, with additional poles, as given by

$$sk\Delta x = \arccos\left(\frac{\alpha}{\alpha-1}\right) \pmod{2\pi} \tag{9}$$

The resulting filter is pass band between the poles, but since the poles are spread at different integer values in k space, a wide band low pass filter can be constructed by combining filters using different strides. Examples are given in Fig. 3 (bottom) for combinations of the filters with stride 1 to 4.

The combined filters with strides 2, 3 and 4 have nearly equivalent fall-offs in gain (in linear scale) to the 20, 50 and 80 passes of the bilinear filter (see Fig. 4). Yet, the filters with stride need respectively only 10, 15 and 15 passes of a three-point filter, compared to respectively 21, 51 and 81 passes for an equivalent n-pass bilinear filter, yielding gains of respectively 2.1, 3.4 and 5.4 in number of operations in favor of the filters with stride.

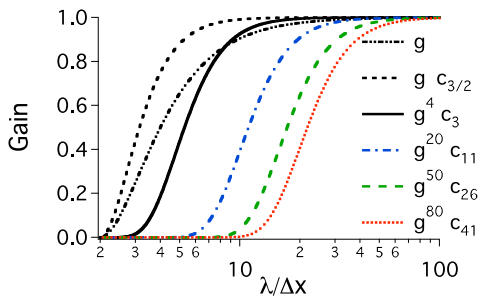

Fig. 2. Gain versus wavelength for the bilinear filter without compensation ($g = g(1/2,k)$), with compensation ($g \cdot c_{3/2} = g(1/2,k)\cdot g(3/2,k)$), and n-pass bilinear filters with compensation ($g^n \cdot c_{\alpha_c} = g(1/2,k)^n \cdot g(\alpha_c,k)$) for $n = 4, 20, 50$ and 80.

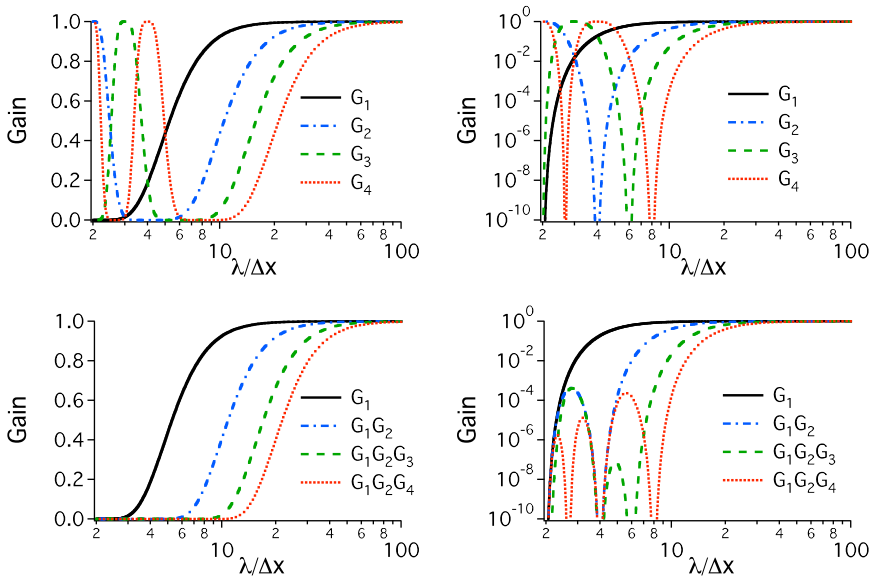

Fig. 3. (Top) Gain for four passes bilinear filters with compensation ($G_s = g(s,1/2,k)^4 \cdot g(s,3,k)$) for strides s = 1 to 4 linear with (left) linear ordinate (right) logarithmic ordinate; (bottom) gain for four low pass filters combining the G_1 to G_4 filters with (left) linear ordinate (right) logarithmic ordinate.

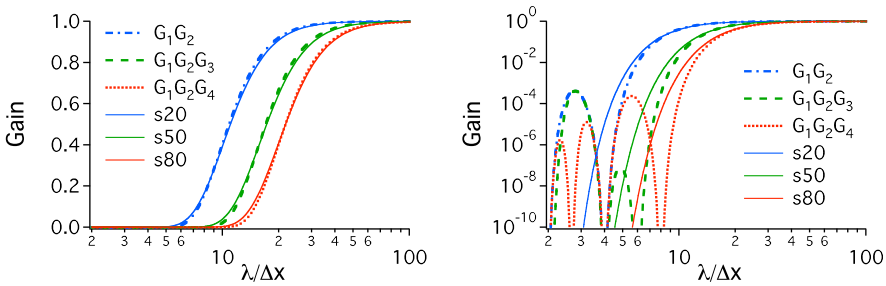

Fig. 4. Comparison between filters with stride and filter s20–80 with (left) linear ordinate (right) logarithmic ordinate.

Filters with larger stencils can be constructed that are equivalent to the n-pass three points filters that were considered here (with or without stride). However, their efficient implementation may be challenging due to the size of the stencil which may eventually exceed the grid length in a given dimension when using domain decomposition on parallel computers. Hence, the use of stride offers a versatile solution for efficient filtering of short wavelength noise without the complications associated with the usage of very large stencils.

2.2. Tunable solver

The electromagnetic solver in conventional Particle-In-Cell methods has numerical dispersion [33]. In order to evaluate the effect of numerical dispersion on boosted frame LPA simulations, an electromagnetic solver with tunable numerical dispersion was implemented in Warp and is described below.

In [35,36], Cole introduced an implementation of the source-free Maxwell's wave equations for narrow-band applications based on non-standard finite-differences (NSFD). In [37], Karkkainen et al. adapted it for wideband applications. At the Courant limit for the time step and for a given set of parameters, the stencil proposed in [37] has no numerical dispersion along the principal axes, provided that the cell size is the same along each dimension (i.e. cubic cells in 3D). The solver from [37] was modified to be consistent with the Particle-In-Cell methodology and implemented in the code Warp, with the ability given to the user of setting the solver adjustable coefficients, providing tunability of the numerical properties of the solver to better fit the requirements of a particular application. In the present case, this solver allows the evaluation of numerical dispersive effects on the instability that is observed in boosted frame simulations of LPAs.

The "Cole–Karkkainnen" (or CK) solver [37] uses a non-standard finite difference formulation (based on extended stencils) of the Maxwell–Ampere equation. For implementation into a Particle-In-Cell code, the formulation must introduce the source term into Cole–Karkkainen's source free formulation in a consistent manner. However, modifying the NSFD formu-

lation of the Maxwell–Ampere equation so that it includes the source term in a way that is consistent with the current deposition scheme is challenging. To circumvent this problem, Warp implementation departs from Karkkainen's by applying the enlarged stencil on the Maxwell–Faraday equations, which does not contain any source term but is formally equivalent to the source-free Maxwell–Ampere equation. Consequently, in Warp's implementation, the discretized Maxwell–Ampere equation is the same as in the Yee scheme, and the discretized Maxwell's equations read:

$$D_t\mathbf{B} = -\nabla^* \times \mathbf{E} \tag{10}$$

$$D_t\mathbf{E} = c^2\nabla \times \mathbf{B} - \frac{\mathbf{J}}{\epsilon_0} \tag{11}$$

$$\nabla \cdot \mathbf{E} = \frac{\rho}{\epsilon_0} \tag{12}$$

$$\nabla^* \cdot \mathbf{B} = 0 \tag{13}$$

where ϵ_0 is the permittivity of vacuum, and Eqs. 12 and 13 are not being solved explicitly but verified via appropriate initial conditions and current deposition procedure. The differential operators are defined as

$$\nabla = D_x\hat{\mathbf{x}} + D_y\hat{\mathbf{y}} + D_z\hat{\mathbf{z}} \tag{14}$$

$$\nabla^* = D_x^*\hat{\mathbf{x}} + D_y^*\hat{\mathbf{y}} + D_z^*\hat{\mathbf{z}}, \tag{15}$$

the finite differences and sums operators being respectively

$$D_tG|_{i,j,k}^n = \frac{G|_{i,j,k}^{n+1/2} - G|_{i,j,k}^{n-1/2}}{\Delta t} \tag{16}$$

$$D_xG|_{i,j,k}^n = \frac{G|_{i+1/2,j,k}^n - G|_{i-1/2,j,k}^n}{\Delta x} \tag{17}$$

$$D_x^* = \left(\alpha + \beta S_x^1 + \gamma S_x^2\right)D_x \tag{18}$$

and

$$S_x^1G|_{i,j,k}^n = G|_{i,j+1/2,k}^n + G|_{i,j-1/2,k}^n + G|_{i,j,k+1/2}^n + G|_{i,j,k-1/2}^n \tag{19}$$

$$S_x^2G|_{i,j,k}^n = G|_{i,j+1/2,k+1/2}^n + G|_{i,j-1/2,k+1/2}^n + G|_{i,j+1/2,k-1/2}^n + G|_{i,j-1/2,k-1/2}^n \tag{20}$$

The quantity G is a sample vector component, Δt and Δx are respectively the time step and the grid cell size along x, while α, β and γ are constant scalars verifying $\alpha + 4\beta + 4\gamma = 1$. As with the Yee algorithm, the quantities with half-integer are located between the nodes (electric field components) or in the center of the cell faces (magnetic field components). The operators along y and z, i.e. $D_y, D_z, D_y^*, D_z^*, S_y^1, S_z^1, S_y^2$, and S_z^2, are obtained by circular permutation of the indices.

In 2D, assuming the plane (x, z), the enlarged finite operators simplify to

$$D_x^* = \left(\alpha + \beta S_x^1\right)D_x \tag{21}$$

$$S_x^1G|_{i,j,k}^n = G|_{i,j+1/2,k}^n + G|_{i,j-1/2,k}^n. \tag{22}$$

An extension of this algorithm for non-cubic cells provided by Cowan in [38] is not considered in this paper. However, all considerations given here for the solver implemented in Warp apply readily to the solver developed by Cowan.

2.2.1. Numerical dispersion

The dispersion relation of the solver is given by

$$\left(\frac{\sin\frac{\omega\Delta t}{2}}{c\Delta t}\right)^2 = C_x\left(\frac{\sin\frac{k_x\Delta x}{2}}{\Delta x}\right)^2 + C_y\left(\frac{\sin\frac{k_y\Delta y}{2}}{\Delta y}\right)^2 + C_z\left(\frac{\sin\frac{k_z\Delta z}{2}}{\Delta z}\right)^2 \tag{23}$$

with

$$C_x = \alpha + 2\beta(c_y + c_z) + 4\gamma c_y c_z \tag{24}$$

$$C_y = \alpha + 2\beta(c_z + c_x) + 4\gamma c_z c_x \tag{25}$$

$$C_z = \alpha + 2\beta(c_x + c_y) + 4\gamma c_x c_y \tag{26}$$

and

$$c_x = \cos(k_x\Delta x) \tag{27}$$

$$c_y = \cos(k_y\Delta y) \tag{28}$$

$$c_z = \cos(k_z\Delta z) \tag{29}$$

The Courant–Friedrichs–Lewy condition (CFL) is given by imposing

$$\left(\sin\frac{\omega\Delta t}{2}\right)^2 \leqslant 1 \tag{30}$$

Table 1
List of possible coefficients for the CK solver and their corresponding CFL time step limit, assuming cubic cells ($\Delta x = \Delta y = \Delta z$).

	Yee	CK	CK2	CK3	CK4	CK5
β'	0	−1	−1/2	0	−1/2	−9/10
γ'	0	1	1/2	−1	0	9/10
α	1	7/12	19/24	11/12	3/4	5/8
β	0	1/12	1/24	1/24	1/16	3/40
γ	0	1/48	1/96	−1/48	0	3/160
$c\Delta t_{CFL}/\Delta x$	$1/\sqrt{3}$	1	$1/\sqrt{2}$	$1/\sqrt{2}$	$\sqrt{2}/\sqrt{3}$	$\sqrt{5}/\sqrt{6}$

We make the ansatz that the most unstable modes propagate at the Nyquist wavelength along the cell (3D) diagonal, the cell faces (2D) diagonals or the main axes. For the cell diagonal, setting $k_x\Delta x = k_y\Delta y = k_z\Delta z = \pi$ gives $\sin(k_x\Delta x/2) = \sin(k_y\Delta y/2) = \sin(k_z\Delta z/2) = 1$ and $c_x = c_y = c_z = -1$, so that $C_x = C_y = C_z = \alpha - 4\beta + 4\gamma$. Repeating for the main axes and 2D diagonals, and solving for (23) and (30), the CFL condition reads

$$c\Delta t_c \leqslant \min\left[\Delta x, \Delta y, \Delta z, 1/\sqrt{(\alpha - 4\gamma)\max[\kappa_x + \kappa_y, \kappa_x + \kappa_z, \kappa_y + \kappa_z]}, 1/\sqrt{(\alpha - 4\beta + 4\gamma)(\kappa_x + \kappa_y + \kappa_z)}\right] \tag{31}$$

where $\kappa_x = 1/\Delta x^2,\ \kappa_y = 1/\Delta y^2$ and $\kappa_z = 1/\Delta z^2$.

Assuming cubic cells ($\Delta x = \Delta y = \Delta z$), the coefficients given in [37] ($\alpha = 7/12,\ \beta = 1/12$ and $\gamma = 1/48$) allow $c\Delta t = \Delta x$, and thus no dispersion along the principal axes.

It is of interest to note that (23) can be rewritten

$$\left(\frac{\sin\frac{\omega\Delta t}{2}}{c\Delta t}\right)^2 = \left(s_x^2 + s_y^2 + s_z^2\right) + \beta'\left(s_x^2 s_y^2 + s_x^2 s_z^2 + s_y^2 s_z^2\right) + \gamma'\left(s_x^2 s_y^2 s_z^2\right) \tag{32}$$

with $s_x = \sin(k_x\Delta x/2),\ s_y = \sin(k_y\Delta y/2),\ s_z = \sin(k_z\Delta z/2),\ \beta' = -8\beta - 16\gamma$ and $\gamma' = 48\gamma$, for which the coefficients from [37] ($\alpha = 7/12,\ \beta = 1/12$ and $\gamma = 1/48$) correspond to the nice values $\beta' = -1$ and $\gamma' = 1$.

Sets of possible coefficients and the corresponding CFL condition, assuming cubic cells, are given in Table 1. The numerical dispersion using those coefficients are plotted in Fig. 5 along the principal axes and diagonals for cubic cells ($\Delta x = \Delta y = \Delta z$) and contrasted with the one of the Yee solver (all taken at each solver's CFL time step limit). At the CFL limit, the Yee algorithm offers no numerical dispersion along the 3D diagonal, but relatively large numerical dispersion at the Nyquist frequency along the main axes. Conversely, the Cole–Karkkainen solver (CK) offers no numerical dispersion along the main axes but significant dispersion along the diagonals. The CK solver also allows larger time steps than the Yee solver by almost a factor of two in 3D. The solver labeled CK2 offers numerical dispersion that is intermediate between the Yee solver and the CK solver along the main axes and the 3D diagonal, but slightly degraded along the 2D diagonal. Conversely, while solver CK3 also offers intermediate numerical dispersion along the main axes and the 3D diagonal, it offers no numerical dispersion along the 2D diagonal. Solver CK4 improves slightly the numerical dispersion along the main axes over CK2 and CK3 at the expense of the dispersion along the diagonals. Finally, CK5 offers the highest level of isotropy. The CFL time steps of solvers CK2, 3, 4 and 5 are intermediate between the Yee and the CK CFL time steps. This provides solvers with a range of numerical dispersion among which some may be more favorable with regard to the minimization of errors or mitigation of numerical instabilities for a given application.

To reduce numerical dispersion to its lowest level, it is desirable to operate the CK solver as close as possible to the CFL limit $c\Delta t = \Delta x$. However, an instability (other than numerical Cerenkov) arises at the Nyquist frequency in such a case. The analysis is given in 1D in Appendix A, as well as its mitigation using digital filtering. Since for the CK solver, the CFL limit is independent of dimensionality, the analysis and mitigation apply readily to 2D and 3D simulations.

For absorption of outgoing waves at the computational box boundaries, the extension of the solver to a Perfectly Matched Layer [39] is described below.

2.2.2. Current deposition and Gauss' Law

In most applications, it is essential to prevent accumulations of errors to the discretized Gauss' Law. This is accomplished by providing a method for depositing the current from the particles to the grid which is compatible with the discretized Gauss' Law, or by providing a mechanism for "divergence cleaning" [33,40–42]. For the former, schemes which allow a deposition of the current that is exact when combined with the Yee solver is given in [43] for linear form factors and in [44] for higher order form factors. Since the discretized Gauss' Law and Maxwell–Faraday equation are the same in our implementation as in the Yee solver, charge conservation is readily verified using the current deposition procedures from [43,44], and this was verified numerically. Hence, in this case, as for standard PIC simulations using the Yee solver, divergence cleaning is not necessary provided that the discretized Gauss' Law is satisfied initially.

2.2.3. Perfectly Matched Layer

The absorption of outgoing waves is typically required in LPA and many other simulations as the computational box is smaller than the experiment and unphysical reflections must be avoided at the simulation box boundaries. The split form

of Perfectly Matched Layer (PML) [45] framework applies readily to Eqs. (10) and (11). The equations on the component along z of the magnetic field are given by

$$(D_t + \sigma_x)B_{zx} = -D_x^* E_y \tag{33}$$
$$(D_t + \sigma_y)B_{zy} = D_y^* E_x \tag{34}$$
$$(D_t + \sigma_x)E_y = -c^2 D_x(B_{zx} + B_{zy}) \tag{35}$$
$$(D_t + \sigma_y)E_x = c^2 D_y(B_{zx} + B_{zy}) \tag{36}$$

where σ_x and σ_y are the absorbing layer coefficients along x and y, respectively. The equations for the other components of the magnetic field and for the electric field are obtained similarly, applying the standard difference operator on the spatial

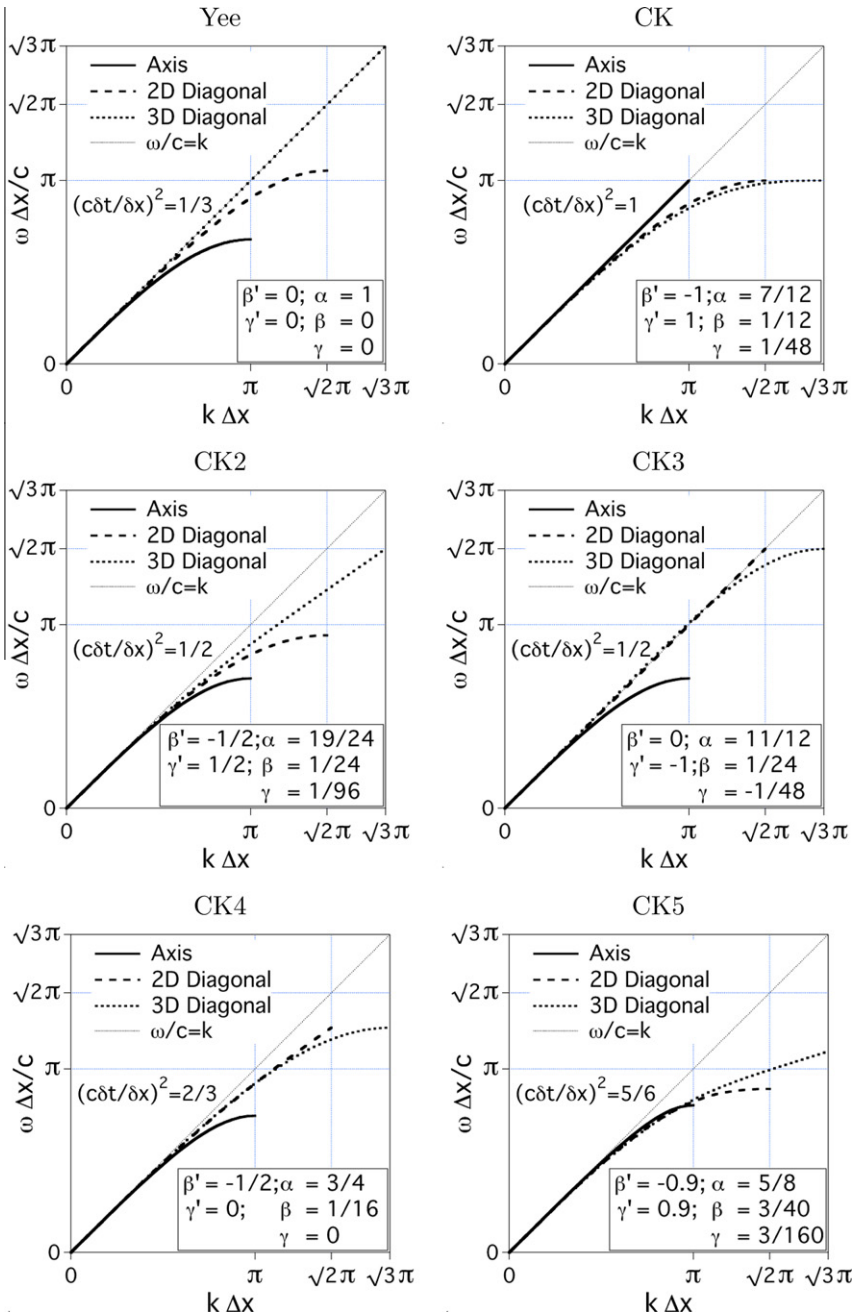

Fig. 5. Numerical dispersion along the principal axis and diagonals for cubic cells ($\Delta x = \Delta y = \Delta z$) at the Courant limit for the solver with adjustable numerical dispersion using the parameters from Table 1.

derivatives of the magnetic field and the enlarged difference operator on the spatial derivatives of the electric field. The formula to update the fields is obtained by solving the finite-difference equations or by integrating over one time step, giving

$$B_{zx}|_{i+1/2,j+1/2,k}^{n+1/2} = \xi_x B_{zx}|_{i+1/2,j+1/2,k}^{n-1/2} - \frac{1-\xi_x}{\sigma_x} D_x^* E_y|_{i+1/2,j+1/2,k}^{n} \tag{37}$$

$$B_{zy}|_{i+1/2,j+1/2,k}^{n+1/2} = \xi_y B_{zy}|_{i+1/2,j+1/2,k}^{n-1/2} + \frac{1-\xi_y}{\sigma_y} D_y^* E_x|_{i+1/2,j+1/2,k}^{n} \tag{38}$$

$$E_y|_{i,j+1/2,k}^{n+1} = \xi_x E_y|_{i,j+1/2,k}^{n} - c^2 \frac{1-\xi_x}{\sigma_x} D_x (B_{zx} + B_{zy})|_{i,j+1/2,k}^{n+1/2} \tag{39}$$

$$E_x|_{i+1/2,j,k}^{n+1} = \xi_y E_x|_{i+1/2,j,k}^{n} + c^2 \frac{1-\xi_y}{\sigma_y} D_y (B_{zx} + B_{zy})|_{i+1/2,j,k}^{n+1/2} \tag{40}$$

where $\xi = (1 - \sigma\Delta t/2)/(1 + \sigma\Delta t/2)$ via direct solve, or $\xi = e^{-\sigma\Delta t}$ via time integration (note that in our tests, both implementations gave nearly identical results).

The PML using the stencil given by (40) was tested and compared to the standard Yee implementation in 2D and 3D. Fig. 6 shows snapshots from 2D simulations of the reflected residue from a PML layer of a pulse with amplitude given by the Harris function $(10 - 15 * \cos(2\pi ct/L) + 6 * \cos(4\pi ct/L) - cos(6\pi ct/L))/32$ where t is time, c is the speed of light and $L = 50\Delta x$ is the pulse length in cell size units. A grid of 400×400 cells was used with $\Delta x = \Delta y$. The absorbing layer was 8 cells deep and the dependency of the PML coefficients with the index position i in the layer was $\sigma_i = \sigma_m(i\Delta x/\Delta)^n$ with $\sigma_m = 4/\Delta x, \Delta = 5\Delta x$ and $n = 2$. The alternative prescription for the coefficients given in [46,47], which reads $\sigma_i^* = (\xi_{i+1/2} - 1/\xi_i)/\Delta x$ with $\xi_i = e^{-\sigma_i \Delta t}$ and $\sigma_i = \sigma_m(i\Delta x/\Delta)^n$, was also tested.

For the generic test case that has been considered, the new implementation exhibited a very low residue of reflections from the PML layer, which are qualitatively and quantitatively very similar to the residue obtained with a standard PML implementation. In agreement with results from [46,47], the use of the modified coefficients σ^* led to an order of magnitude improvement over the use of the standard coefficients.

The 3D tests gave similar absorption efficiency between the Yee and the new solver implementations of the PML, for all the CK solver coefficients given in Table 1.

It was shown in [46,47] that the efficiency of the layer can be improved further for the standard PML by augmenting the equations with additional terms. However, a similar extension may not be readily available when using the Cole–Karkkainen stencil and is not considered here.

2.3. Friedman adjustable damping

The tunable damping scheme developed by Friedman [30] was shown to be the most potent practical method for mitigating the numerical Cerenkov instability in [23], among the selected methods that were considered. It is readily applicable to the solver presented above and has been implemented in Warp by modifying (10) to

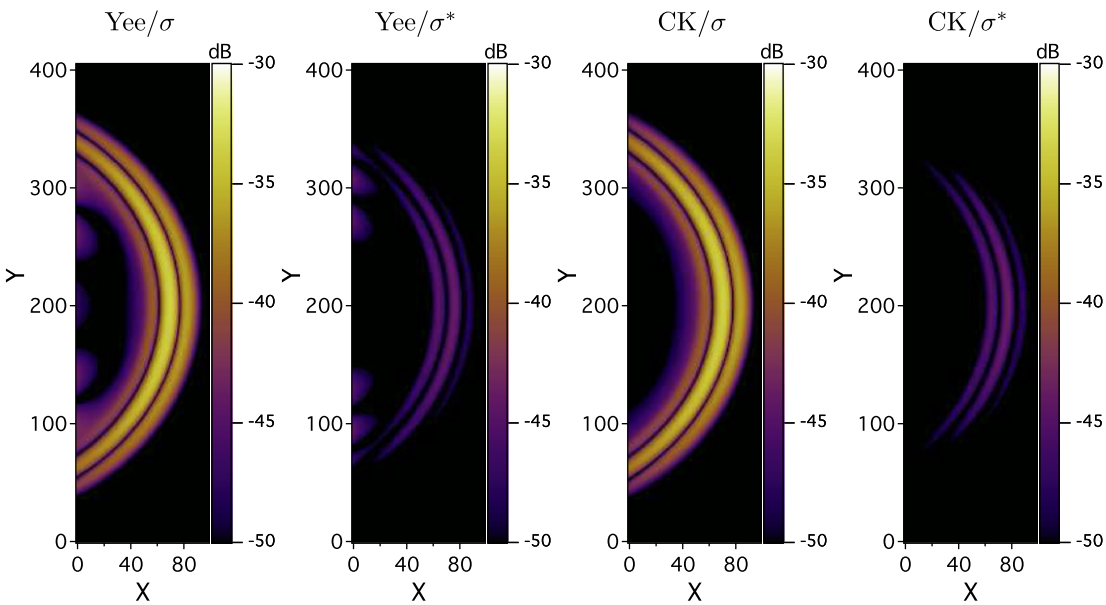

Fig. 6. Reflected signal (in dB) from a PML layer using the Yee or the Cole–Karkkainen solver, and the absorbing coefficient σ or σ^*. Each simulation was run with the time step set at the Courant limit.

$$\mathbf{B}^{n+3/2} = \mathbf{B}^{n+1/2} - \Delta t \nabla^* \times \left[\left(1 + \frac{\theta}{4}\right) \mathbf{E}^{n+1} - \frac{1}{2}\mathbf{E}^n + \left(\frac{1}{2} - \frac{\theta}{4}\right) \overline{\mathbf{E}}^{n-1} \right] \tag{41}$$

with

$$\overline{\mathbf{E}}^{n-1} = \left(1 - \frac{\theta}{2}\right) \mathbf{E}^n + \frac{\theta}{2} \overline{\mathbf{E}}^{n-2} \tag{42}$$

where $0 \leqslant \theta \leqslant 1$ is the damping factor. The numerical dispersion becomes

$$\left(\frac{\sin \frac{\omega \Delta t}{2}}{c \Delta t} \right)^2 = F \Omega^2 \tag{43}$$

where

$$F = 1 - \frac{2\theta \sin^2(\omega \Delta t/2)}{2e^{-i\omega\Delta t} - \theta} \tag{44}$$

and

$$\Omega^2 = \left[C_x \left(\frac{\sin \frac{k_x \Delta x}{2}}{\Delta x} \right)^2 + C_y \left(\frac{\sin \frac{k_y \Delta y}{2}}{\Delta y} \right)^2 + C_z \left(\frac{\sin \frac{k_z \Delta z}{2}}{\Delta z} \right)^2 \right] \tag{45}$$

The CFL is given by

$$c\Delta t_c^* = c\Delta t_c \sqrt{\frac{2+\theta}{2+3\theta}} \tag{46}$$

where Δt_c is the critical time step of the numerical scheme without damping ($\theta = 0$), as given by (31).

The numerical dispersion of the Cole–Karkkainen–Friedman (CKF) solver (using the coefficients from the CK solver in Table 1) is plotted in Fig. 7 along the principal axis and diagonals for cubic cells ($\Delta x = \Delta y = \Delta z$) and contrasted with the one of

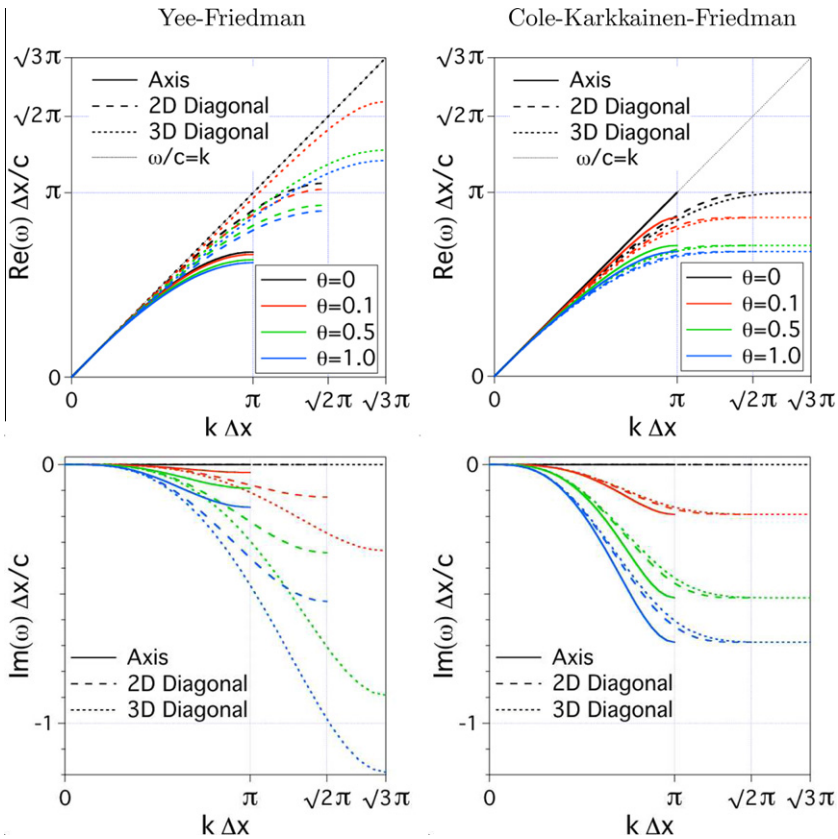

Fig. 7. Numerical dispersion along the principal axis and diagonals for cubic cells ($\Delta x = \Delta y = \Delta z$) at the Courant limit for: (left) the Yee–Friedman solver; (right) the Cole–Karkkainen–Friedman solver. The real part (phase) and the imaginary part (amplitude) are plotted respectively in the top and bottom rows.

the Yee–Friedman (YF) solver (both taken at the Courant time step limit). The amount of phase error rises with the value of the damping parameter θ (partly due to the slightly more constraining limit on the critical time step). However, it was shown in [23] that the amount of damping provided by the YF solver was sufficient to counteract the slight degradation of numerical dispersion with raising θ, reducing the numerical Cerenkov effects to an acceptable level for the problem that was considered. The damping is very isotropic with the CKF solver but not with the YF one. The YF implementation offers a higher level of damping of the shortest wavelengths along the 3D diagonals, while the CKF offers more damping along the axes, and the amount of damping along the 2D diagonals are similar. In summary, the YF implementation delivers respectively the highest/lowest level of damping in the direction of lowest/highest numerical dispersion, while the CKF implementation delivers a proportionally higher level of damping than the YF implementation along the direction of highest numerical dispersion. Thus it may be expected that the CKF implementation will be more efficient in reducing numerical Cerenkov effects.

3. Application to the modeling of laser wakefield acceleration

The methods presented in Section 2 were applied to the modeling of full scale and downscaled 10 GeV LPA stages in 2–1/2D and 3D. Downscaled simulations are presented first, as they allowed for fast evaluation of the effects of potential remediations due to quick run times and because the instability was present at modest levels which did not generally disrupt the simulation. The modeling of full scale stages, which allow for higher values of γ for the reference frame due to the high γ of the plasma wake structure, is more prone to the high frequency instability and are presented to strongly test these techniques and show effectiveness in modeling high energy stages.

3.1. Scaled 10 GeV stages

It has been shown that many parameters of high energy LPA stages can be accurately simulated at reduced cost by simulating stages of lower energy gain, with higher density and shorter acceleration distance, by scaling the physical quantities relative to the plasma wavelength, and this has been applied to design of 10 GeV LPA stages [48,49]. The parameters were chosen here to be close (though not identical) to the case where $k_pL = 2$ in [48] where k_p is the plasma wavenumber and L is the laser pulse length. In the cases considered in this paper, the beam is composed of test particles only, with the goal of testing the fidelity of the algorithm in modeling laser propagation and wake generation. These simulations are scaled replicas of 10 GeV stages that would operate at actual densities of 10^{17} cm^{-3} [48,49]. The main physical and numerical parameters of the simulations are given in Table 2. Unless noted otherwise, in all the simulations presented herein, the field is gathered from the grid onto the particles directly from the Yee mesh locations, i.e. using the "energy conserving" procedure (see [33], chapter 10). The approximate relativistic factor of the wake that is formed by the laser traveling in the plasma is given, according to linear theory, by $\gamma_w = 2\pi c/\lambda\omega_p$ where $\omega_p = \sqrt{n_e e^2/\epsilon_0 m_e}$ is the electron plasma frequency. For the given parameters, $\gamma_w \approx 13.2$. Thus, Warp simulations were performed using reference frames at $\gamma = 1$ (laboratory frame) and frames moving at up to 13.

In this paper, the effect of the instability, filtering or damping on the physical quantities of interest (laser, wake and accelerated beam evolutions), are captured by examination of the beam energy history, which was chosen since it integrates the

Table 2
List of parameters for scaled 10 GeV class LPA stage simulation.

Beam radius	R_b	82.5 nm
Beam length	L_b	85. nm
Beam transverse profile		$\exp\left(-r^2/8R_b^2\right)$
Beam longitudinal profile		$\exp\left(-z^2/2L_b^2\right)$
Laser wavelength	λ	0.8 μm
Laser length (FWHM)	L	10.08 μm
Normalized vector potential	a_0	1
Laser longitudinal profile		$\sin(\pi z/L)$
Plasma density on axis	n_e	10^{19} cm^{-3}
Plasma longitudinal profile		Flat
Plasma length	L	1.5 mm
Plasma entrance ramp profile		Half sinus
Plasma entrance ramp length		4 μm
Number of cells in x	N_x	75
Number of cells in z	N_z	860 ($\gamma = 13$)-1691 ($\gamma = 1$)
Cell size in x	Δx	0.65 μm
Cell size in z	Δz	$\lambda/32$
Time step	Δt	At CFL limit
Particle deposition order		Cubic
# of plasma particles/cell		1 macro-e$^-$+1 macro-p$^+$

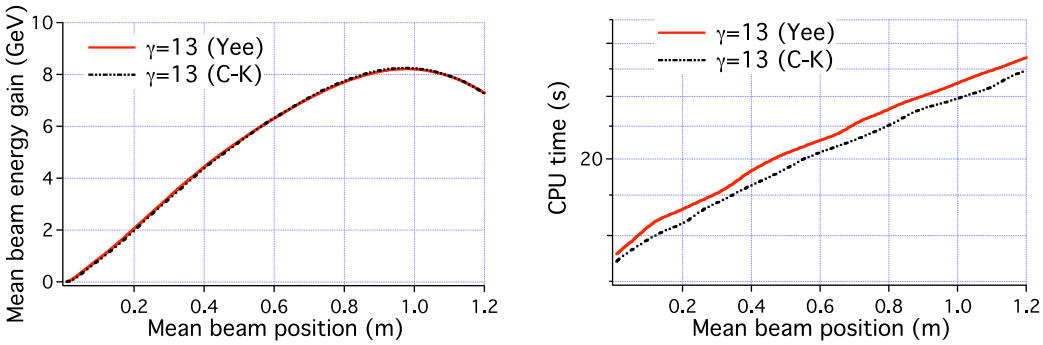

Fig. 8. (Left) Average scaled beam energy gain and (right) CPU time, versus longitudinal position in the laboratory frame from simulations in boosted frames at $\gamma = 13$, using the Cole–Karkkainen solver with filter S (1) (red curve is reference from calculation with Yee solver and filter S(1)). (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

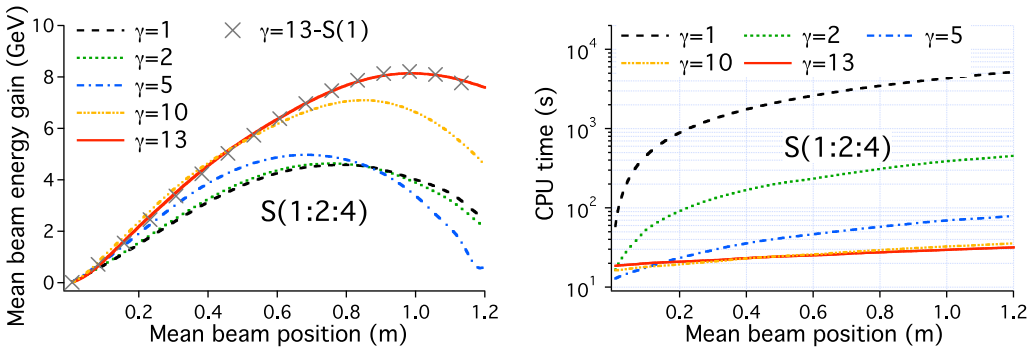

Fig. 9. (Left) Average scaled beam energy gain and (right) CPU time, versus longitudinal position in the laboratory frame from simulations in the laboratory frame ($\gamma = 1$) and boosted frames at $\gamma = 2, 5, 10$ and 13, using the Yee solver with digital filter S(1:2:4) (gray cross is reference from run with filter S(1)).

other effects, as the beam energy is affected by the wake evolution which is itself affected by the laser evolution. Details of the convergence of the laser and wake evolutions, as well as of the beam transverse and longitudinal properties, are analyzed in [25] and in a paper to be published elsewhere.

The results from a simulation using the Cole–Karkkainen solver is contrasted to a simulation using the standard Yee solver in Fig. 8. The two calculations gave identical results in this case, validating our implementation of the CK solver. Despite the more expensive stencil, the run with the CK solver was almost 40% faster, due to a time step larger by $\sqrt{2}$. The two calculations, as well as those presented in the rest of this paper, were performed with the time step set at the CFL limit, so as to minimize numerical dispersion errors from the electromagnetic solver.

Simulations were then conducted using the Yee solver with digital filter S(1:2:4) as shown in Fig. 9, and using the Yee-Friedman solver as shown in Fig. 10. Results are contrasted to a run at $\gamma = 13$ with little filtering and no damping (crosses in Figs. 9 and 10) that was in turn benchmarked against a run in the laboratory frame [25], and serves as the reference. None of these simulations were disrupted by short wavelength instability. The results hence show what mitigation techniques are allowable by comparing physical quantities and benchmarking to reference run. Smoothing with the wideband filter S(1:2:4) did not produce significant degradations for the calculation in the wake frame ($\gamma = 13$) but did otherwise. Similar to filtering, damping aggressively did not degrade the result in the range $10 \leqslant \gamma \leqslant 13$ but did significantly in the range $1 \leqslant \gamma \leqslant 5$. Comparing the timings with those of Fig. 13 (middle-left) from [25] shows that the smoothing and the damping added less than a factor of two of total runtime to the simulations. These results lead to several observations:

1. while the grid dimensions and number of cells were chosen such that square cells were obtained for $\gamma = 13$, meaning a larger dispersion in the longitudinal direction with the Yee solver than with the Cole–Karkkainen solver, both gave the same result. This is significant since for simulations of LPA in the laboratory frame reported in the literature, the need to have nearly perfect numerical dispersion in the longitudinal direction imposes a constraint on the cell aspect ratio and thus on resolution [50,51]. This constraint is removed when simulating in the frame of the wake ($\gamma = 13 \approx \gamma_w$),
2. damping of high frequencies with the Yee–Friedman solver or wideband smoothing of short wavelength has a negligible effect on accuracy for simulations in the frame of the wake, but degrades the accuracy very significantly for slower moving reference frames. The dependency of the effect of damping and smoothing with γ boost has two causes. First, simulations with a boost $\gamma \approx \gamma_w$ require fewer time steps than simulations using a lower value of γ. Thus, for a given value of

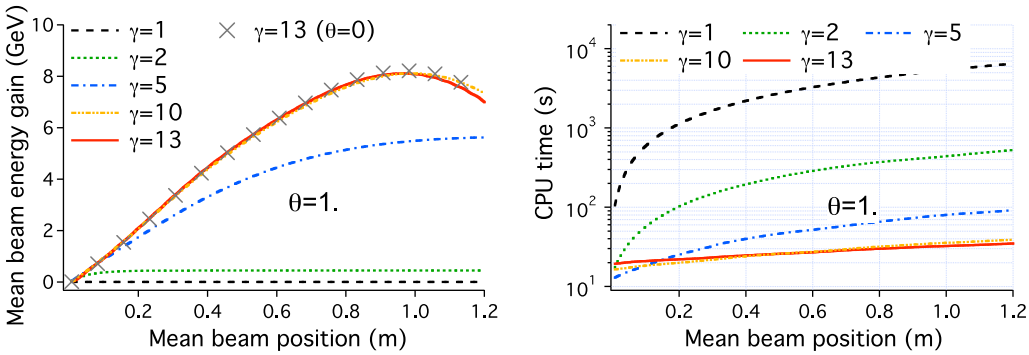

Fig. 10. (Left) Average scaled beam energy gain and (right) CPU time, versus longitudinal position in the laboratory frame from simulations in the laboratory frame ($\gamma = 1$) and boosted frames at $\gamma = 2, 5, 10$ and 13, using the Yee–Friedman solver with $\theta = 1$ (grey cross is reference from run with no damping).

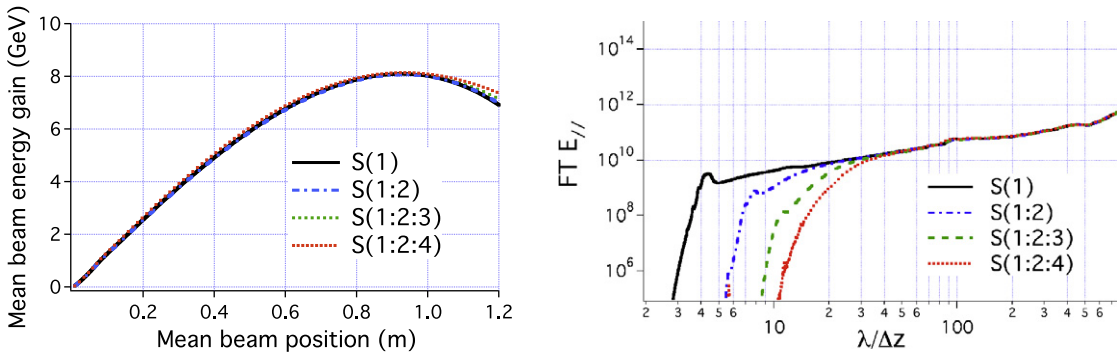

Fig. 11. (Left) Average beam energy gain versus longitudinal position (in the laboratory frame), (right) average (over transverse direction) of the absolute value of the Fourier Transform of the longitudinal electric field at t = 40 ps, from 2D-1/2 simulations of a full scale 10 GeV LPA in a boosted frame at $\gamma = 130$, using the Yee solver and various digital filter kernels. Square cells ($\Delta x = \Delta z = 6.5\mu$m) and the CFL time step ($c\Delta t/\Delta z = 1/\sqrt{2}$) were used.

the damping coefficient θ, the integrated amount of damping will be lower for the simulations with $\gamma \approx \gamma_w$. Second, as explained in [25,34], a large fraction of the short wavelength content that is present in the simulations in the laboratory frame is transformed into time oscillations in simulations in the wake frame. Hence, filtering short wavelengths has less effect on the physics when calculating in the wake frame than when calculating in the laboratory frame,

3. the computational cost of using even the most aggressive damping or smoothing is low, especially considering that the simulations presented here were using only two plasma macro-particles per cell.

In summary, calculating in a boosted frame near the frame following the wake ($\gamma \approx \gamma_w$) relaxes the constraint on the numerical dispersion in the direction of propagation of the laser (which is essential in simulations in the laboratory frame), and allows for more aggressive damping of high frequencies and smoothing of short wavelengths than is possible in standard laboratory frame calculations, while maintaining accuracy in resolution of the physical quantities of interest (wake accelerating field).

3.2. Full scale 10 GeV class stages

Full scale simulations are required for studying laser focusing and beam emittance evolution, as these quantities are not reproduced correctly in downscaled simulations. However, as noted in [9], full scale simulations using the laboratory frame of 10 GeV stages at plasma densities of 10^{17} cm^{-3} are not practical on present computers in 2D and 3D. At this density, the wake relativistic factor $\gamma_w \approx 132$, and 2–1/2D and 3D simulations were done in boosted frames up to $\gamma = 130$ to evaluate which numerical techniques are required to allow such simulations.

Fig. 11 shows the average beam energy gain versus longitudinal position and the averaged Fourier Transform of the longitudinal electric field taken at t = 40 ps, from 2D-1/2 simulations of a full scale 10 GeV LPA in a boosted frame at $\gamma = 130$, using the Yee solver and various smoothing kernels. All runs gave the same beam energy history within a few percent, and no sign of instability is detected in the Fourier transform plot of the longitudinal electric field.

In 3D, all simulations at $\gamma = 130$ using the Yee solver (using cubic cells and a time step at the CFL limit) developed the instability and loss of the beam, regardless of the amount of filtering or damping that has been tried (not shown). The failure of the 3D simulations using the Yee solver motivated use of the Cole–Karkkainen–Friedman (CKF) solver which offers no dispersion along the principal axes with the time step at the CFL limit, so as to evaluate whether it reduces the level of instability. Various levels of filtering and damping were also applied. Data from 3D simulations using the CKF solver and various smoothing kernels are plotted in Fig. 12. Stability is attained when using a sufficient level of filtering. Damping is detrimental to stability at low levels ($\theta = 0.1$) but is beneficial at a higher level ($\theta = 0.5$) as shown by the need for more (less) filtering to allow propagation of the beam to peak energy with low (high) damping relative to the base case. Overall, filtering more strongly affects the instability than damping, and damping was not investigated further.

Next, simulations using the solver coefficients CK2-5 from Table 1 were performed, with the time step set at their respective CFL limit. The best results were obtained using solvers CK2 and CK3, while CK4 and CK5 did not offer substantial improvement over the CK solver. The results from the runs using CK2 and CK3 were nearly identical and hence only those

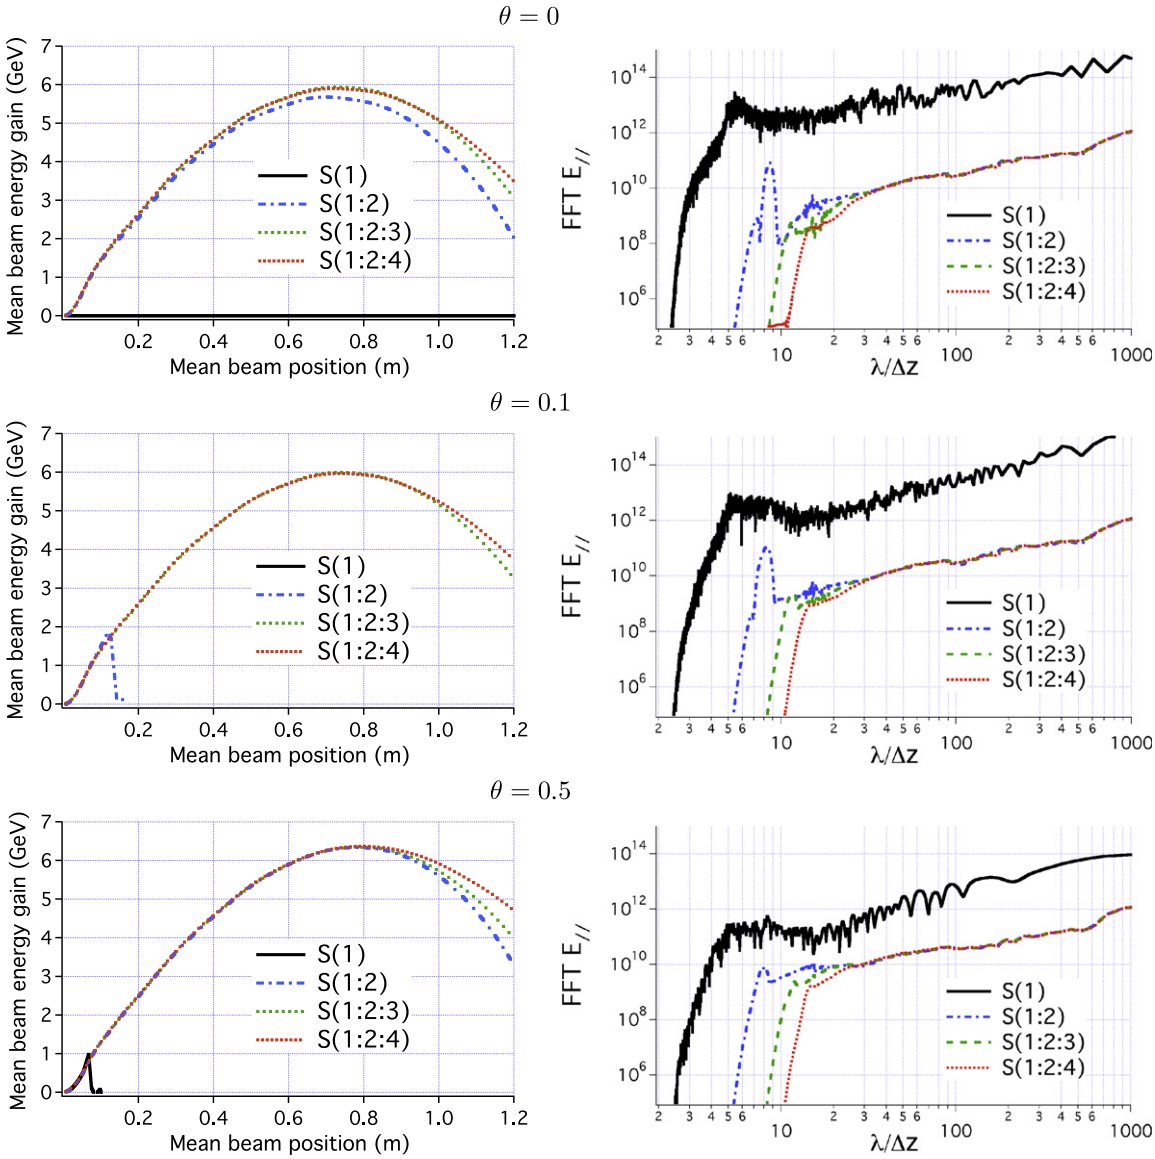

Fig. 12. (Left) Average beam energy gain versus longitudinal position (in the laboratory frame); (right) average (over transverse direction) of the absolute value of the Fourier Transform of the longitudinal electric field at $t = 40$ ps, from 3D simulations of a full scale 10 GeV LPA in a boosted frame at $\gamma = 130$, using the Cole–Karkkainen–Friedman solver and various smoothing kernels, with (top) no numerical damping ($\theta = 0$), (middle) damping with $\theta = 0.1$ and (bottom) $\theta = 0.5$.

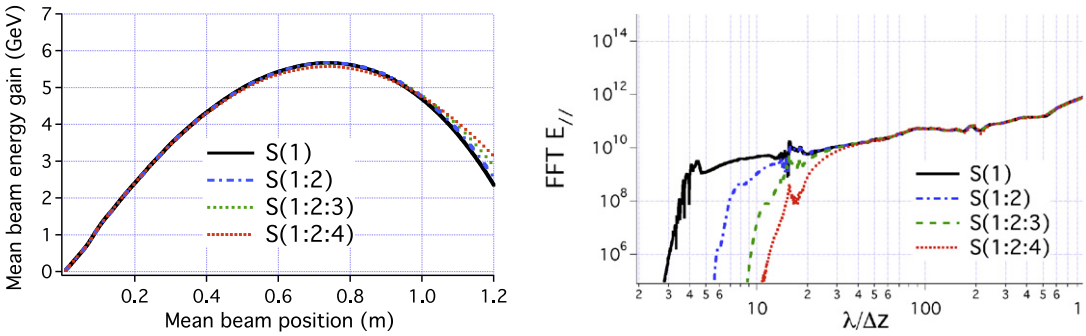

Fig. 13. (Left) Average beam energy gain versus longitudinal position (in the laboratory frame), (right) average (over transverse direction) of the absolute value of the Fourier Transform of the longitudinal electric field at $t = 40$ ps, from 3D simulations of a full scale 10 GeV LPA in a boosted frame at $\gamma = 130$, using the CK2 solver and various digital filter kernels. The low level activity that is present around $\lambda/\Delta z = 16$ is from the ponderomotive force at half the wavelength of the laser in vacuum and is physical.

from CK2 are reported in Fig. 13, which show very consistent beam energy gain histories, and no sign of instability in the Fourier Transform plot of the longitudinal electric field at $t = 40$ ps.

In summary, the full scale 10 GeV class simulations using the frame of the wake performed in this subsection show that:

1. 2–1/2D simulations using the Yee solver at the CFL limit (with square cells) were essentially free of instability,
2. 3D simulations using the Yee solver developed strong instabilities that could not be mitigated using filtering or damping, even at high levels,
3. 3D simulations using the CK solver developed moderately strong instabilities that could be mitigated using moderate to high levels of damping and/or filtering, the latter being the most effective,
4. 3D simulations using the CK2 (or CK3) solver were essentially free of instability.

In conclusion of this section, stable direct modeling of full scale 10 GeV class stages is demonstrated using the boosted frame approach in 2–1/2D and 3D using respectively the Yee and the CK2 (or CK3) solvers. Stable modeling in 3D using the CK solver is possible if using a sufficient level of filtering, which is shown to be adequate if using the frame of the wake as frame of reference.

4. Effects of numerical parameters on the observed instability

In the previous section, it was demonstrated that the usage of the new tunable solver and moderate filtering could be successfully combined for mitigating the short wavelength instability, allowing for efficient direct simulations of full scale 10 GeV class stages. In this section, we analyze in detail the sensitivity of the instability to numerical parameters.

The Fourier transform of the longitudinal electric field averaged over the whole domain at $t = 40$ ps, from 3D simulations using the Yee solver, is given in Fig. 14 (left). It is contrasted to the same data taken from 2–1/2D simulations (right). Both simulations used the same time step at the 3D CFL limit $c\Delta t = \Delta z/\sqrt{3}$. The similarity of the two plots indicates that the degradation of the numerical dispersion that resulted from going from the 2D to the 3D CFL limit is the cause of the failure of the 3D runs using the Yee solver. It follows that the instability that is observed in 3D can be studied in 2–1/2D, provided that the time step is set at the 3D CFL limit.

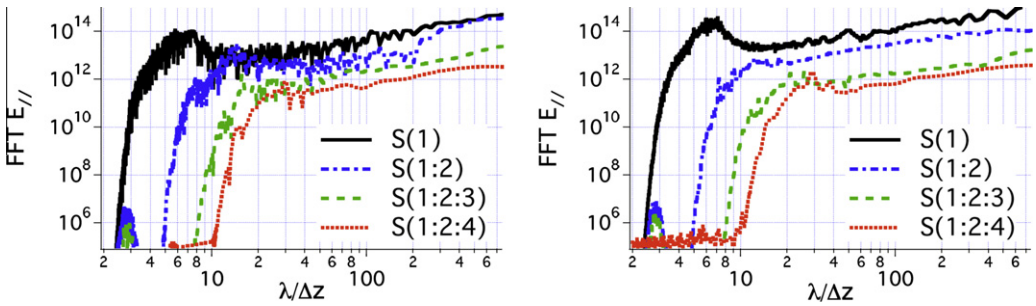

Fig. 14. Average (over transverse direction) of the absolute value of the Fourier Transform of the longitudinal electric field at $t = 40$ ps, from (left) 3D and (right) 2D-1/2 simulations of a full scale 10 GeV LPA in a boosted frame at $\gamma = 130$, using the Yee solver and various smoothing kernels. The same time step at the 3D CFL limit $c\Delta t = \Delta z/\sqrt{3}$ was used for both simulations.

4.1. Effects of spatial resolution

Snapshots of the longitudinal electric field at the front of the plasma taken at t = 12.5 ps, and their corresponding Fourier transform, are given in Fig. 15, from 2–1/2D simulations using the Yee solver with the time step at the 3D CFL limit $c\Delta t = \Delta z/\sqrt{3}$. Three resolutions were considered: (a) $\Delta x = \Delta z = 13$ μm, (b) $\Delta x = \Delta z = 6.5$ μm, and (c) $\Delta x = \Delta z = 3.25$ μm. The amplitude of the instability is roughly inversely proportional to the resolution. For this configuration, the instability exhibits two primary modes at various relative levels, both at a fixed number of grid cells in the longitudinal direction, and at a fixed absolute length in the transverse direction. This indicates that the transverse part of the modes is governed by the physical geometry of the problem while the longitudinal part is governed by numerical resolution. Because the instability always appears near the same value in λ/dz, filtering to remove the instability always requires a filter of similar width,

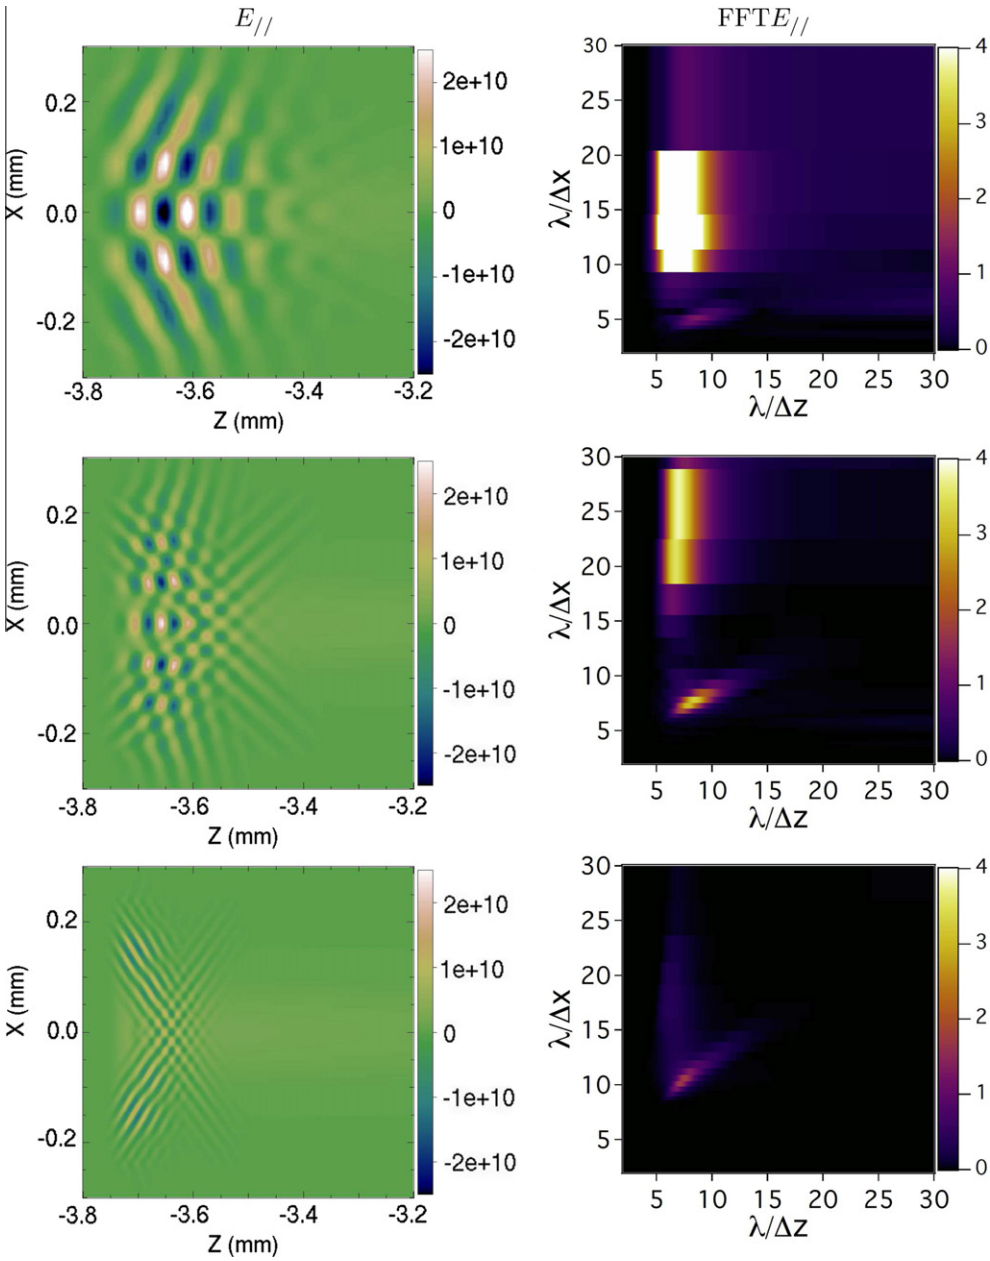

Fig. 15. (left) Snapshot of the longitudinal electric field ($E_{//}$) at the front of the plasma at t = 12.5 ps; (right) Fourier Transform of the longitudinal electric field, from 2–1/2D simulations of a full scale 10 GeV LPA in a boosted frame at $\gamma = 130$, using the Yee solver, for (top) $\Delta x = \Delta z = 13$ μm; (middle) $\Delta x = \Delta z = 6.5$ μm; (bottom) $\Delta x = \Delta z = 3.25$ μm. The time step at the 3D CFL limit $c\Delta t = \Delta z/\sqrt{3}$ was used for all three simulations.

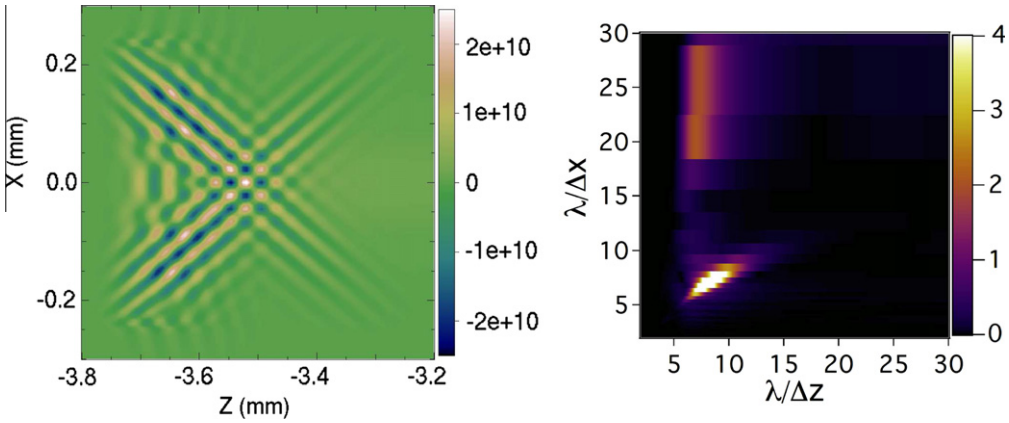

Fig. 16. (left) Snapshot of the longitudinal electric field ($E_{//}$) at the front of the plasma at t = 12.5 ps; (right) Fourier Transform of the longitudinal electric field, from 2–1/2D simulations of a full scale 10 GeV LPA in a boosted frame at $\gamma = 130$, with the CK solver, using $\Delta x = \Delta z = 6.5$ μm, and the time step at the 3D CFL limit $c\Delta t = \Delta z/\sqrt{3}$.

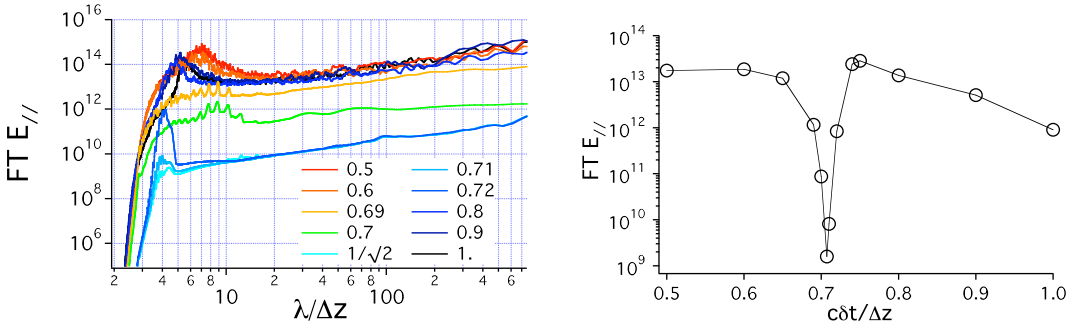

Fig. 17. Average (over transverse direction) of the absolute value of the Fourier Transform of the longitudinal electric field at t = 40 ps, from 2–1/2D simulations of a full scale 10 GeV LPA in a boosted frame at $\gamma = 130$, using the CK solver, for time steps between $c\Delta t/\Delta z = 0.5$ and $c\Delta t/\Delta z = 1$, versus $\lambda/\Delta z$ (left) and at $\lambda/\Delta z = 4$ (right).

and simulation resolution must be chosen high enough to ensure this wavelength is separated from wavelengths of physical interest (laser, plasma wave, etc.).

Results from a 2–1/2D simulation using the CK solver at the 3D CFL limit $c\Delta t/\Delta z = 1/\sqrt{3}$ at the resolution $\Delta x = \Delta z = 6.5$ μm are given in Fig. 16. The same two modes that were observed in the plots from the equivalent simulation using the Yee solver (see Fig. 15-middle), are present, and the overall amplitude of the instability is similar. These similarities in the details of the instability between the Yee and CK solvers indicate that the differences in numerical dispersion of the solvers do not constitute a key factor affecting the instability.

4.2. Effects of time step

It is striking that all the solvers that lead to the lowest levels of instability had the same CFL time step $c\Delta t_{CFL} = \Delta z/\sqrt{2}$. For checking whether this is coincidental, simulations were performed using the CK solver, scanning the time step between $c\Delta t/\Delta z = 0.5$ and $c\Delta t/\Delta z = 1$. The Fourier Transform of the longitudinal field averaged over the entire domain taken at t = 40 ps, is given in Fig. 17, exhibiting a sharp reduction of the instability level in a narrow band around $c\Delta t = \Delta z/\sqrt{2}$. Since the numerical dispersion degrades in all directions when the time step diminishes, this indicates that the value of the time step is of more importance than the numerical dispersion of the solver being used.

Simulations using the Yee or the CK solver with the singular time step $c\Delta t = \Delta z/\sqrt{2}$ were performed and produced levels of instabilities that were much reduced (and delayed) compared to the 3D CFL time step (not shown). The snapshot of the electric field and its Fourier Transform taken at t = 49 ps (versus 12.5 ps for plots in Fig. 15) are given in Fig. 18. The Fourier spectrum is very similar in each case, although the instability is slightly stronger with the CK solver than with the Yee solver. In both cases, the instability is easily removed by using the S(1:2) filter (see Fig. 19).

As mentioned in the previous section, the solvers CK, CK4 and CK5, which all have a CFL time step above the singular time step $c\Delta t = \Delta z/\sqrt{2}$, produced significant levels of instability when run at their CFL limit. It was verified that using those solv-

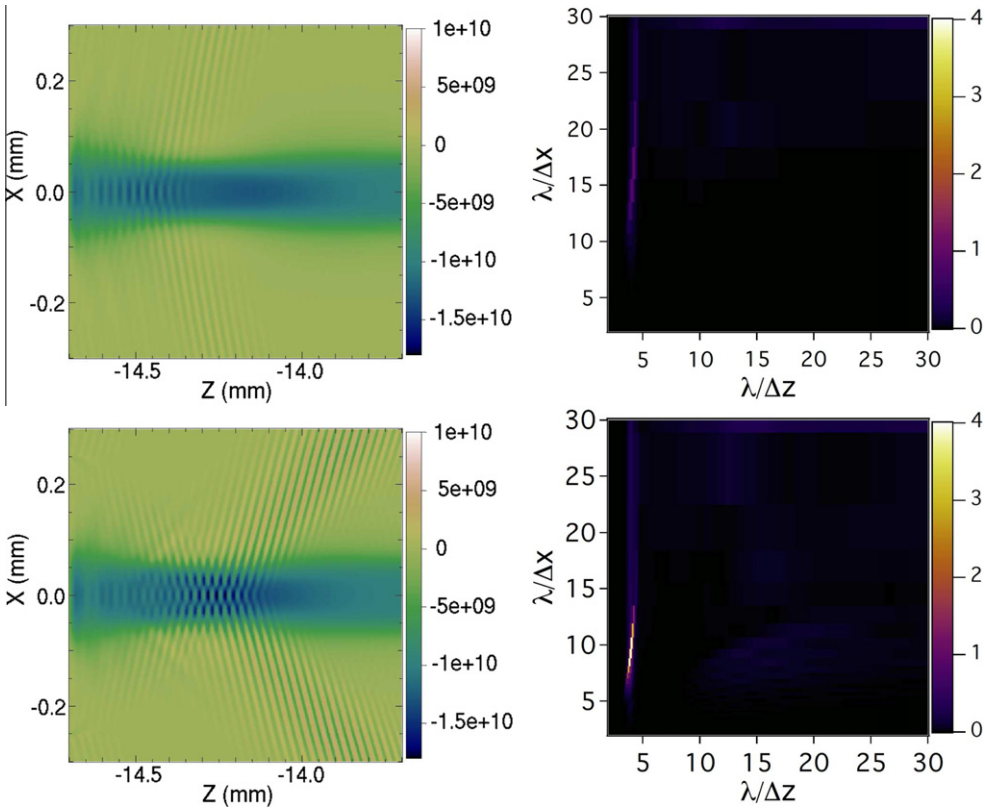

Fig. 18. (left) Snapshot of the longitudinal electric field ($E_{//}$) at the front of the plasma at t = 49 ps; (right) Fourier Transform of the longitudinal electric field, from 2–1/2D simulations of a full scale 10 GeV LPA in a boosted frame at $\gamma = 130$, using $\Delta x = \Delta z = 6.5$ μm, and the time step at the 2D CFL limit $c\Delta t = \Delta z/\sqrt{2}$, for (top) the Yee solver; (bottom) the CK solver. The short wavelength features on the left plots are the unphysical instability. The longer wavelength structure is the wake.

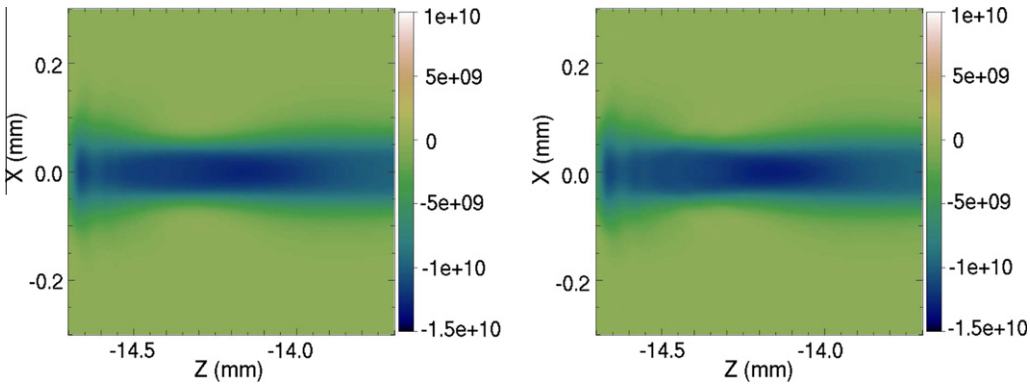

Fig. 19. Snapshot of the longitudinal electric field ($E_{//}$) at the front of the plasma at t = 49 ps from 2–1/2D simulations of a full scale 10 GeV LPA in a boosted frame at $\gamma = 130$, using $\Delta x = \Delta z = 6.5$ μm, and the time step at the 2D CFL limit $c\Delta t = \Delta z/\sqrt{2}$, for (left) the Yee solver; (right) the CK solver. The filter S (1:2) was used to remove the instability that is visible in Fig. 18. The remaining feature is the wake.

ers in 3D at the time step $c\Delta t = \Delta z/\sqrt{2}$ resulted in greatly reduced levels of instability. It was also observed that running the Yee solver using non-cubic cells, e.g. with lower resolution transversely such as $\Delta x = 2 \times \Delta z$ at $\gamma = 130$, or $\Delta x = 2.6 \times \Delta z$ at $\gamma = 50$, produced the same pattern: a significant instability was present when using the CFL time step and was greatly reduced by using $c\Delta t = \Delta z/\sqrt{2}$. Hence for the suppression of the instability, the choice of the solver seems to depend solely on whether its CFL condition allows stability at the special time step $c\Delta t = \Delta z/\sqrt{2}$ for a given grid cell aspect ratio, but not significantly on its numerical dispersion nor on the value of the grid cell aspect ratio.

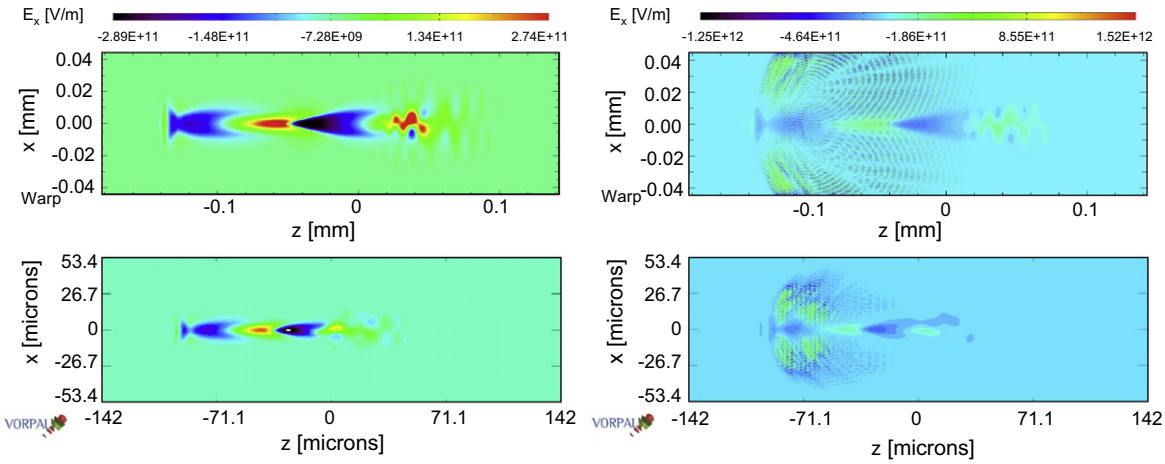

Fig. 20. Colored contour plots of the longitudinal electric field from 2–1/2D simulations of laser plasma acceleration stages in a boosted frame of $\gamma = 5.6$ with the "magical" time step $c\Delta t \approx \Delta z/\sqrt{2}$ (left) and another time step $c\Delta t \approx \Delta z/\sqrt{1.6}$ (right) from the codes Warp (top) and VORPAL (bottom). The two codes used similar but not identical parameters. VORPAL results courtesy of D. L. Bruhwiler from Tech-X Corp.

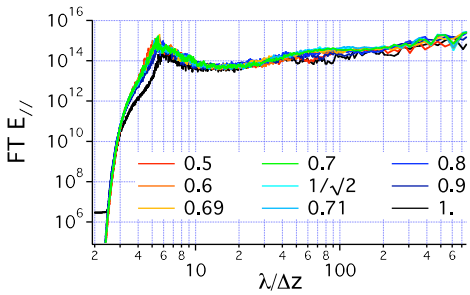

Fig. 21. Average (over transverse direction) of the absolute value of the Fourier Transform of the longitudinal electric field at $t = 40$ ps, from 2–1/2D simulations of a full scale 10 GeV LPA in a boosted frame at $\gamma = 130$, using the CK solver, for time steps between $c\Delta t/\Delta z = 0.5$ and $c\Delta t/\Delta z = 1$, using a 'momentum conserving' field gathering scheme.

The existence of such a time step was confirmed by Bruhwiler using VORPAL [52]. Fig. 20 displays the contour plots of the longitudinal electric field from 2–1/2D simulations of laser plasma acceleration stages in a boosted frame of $\gamma = 5.6$ using the "magical" time step $c\Delta t \approx \Delta z/\sqrt{2}$ or a different time step $c\Delta t \approx \Delta z/\sqrt{1.6}$ from the codes Warp and VORPAL, both using the Yee solver, no damping of the field nor filtering of the current density nor gathered fields. All snapshots have been taken at the same physical time but the physical parameters of the Warp and VORPAL simulations, although being close, were not identical, explaining the observed differences. Nonetheless, both codes exhibited similar levels of instability at a given physical time which was greatly reduced (or delayed) if using the special time step.

The scan of time step was repeated using the "momentum conserving" procedure [33], in which the field values are interpolated at the grid nodes before being gathered onto the particles. The result is given in Fig. 21. With the momentum conserving procedure, the level of instability is consistently high and independent of the time step. Since the numerical dispersion of the solver varies substantially with the time step, these results support the conclusion that the instability may not be of standard numerical Cerenkov nature. The identification of the nature of the instability and the explanation of the singular time step $c\Delta t_S$ call for a multidimensional (no instability was observed in 1D regardless of the field gathering method) analysis of the discretized Vlasov algorithm that was employed, which is left for future work. The results that were obtained in this section lead to the following conclusions:

1. the time step $c\Delta t_S = \Delta z/\sqrt{2}$ consistently produces the lowest levels of instability, regardless of dimensionality (2D vs 3D), the field solver being used, resolution, aspect ratio of cells (within the range of the finite number of cases that were considered),
2. the main advantage of the tunable field solver resides in allowing access to the singular time step $c\Delta t_S$ rather than providing improved numerical dispersion, which consequently do not appear to be a primary driver of the instability,
3. the instability is not completely removed at $c\Delta t_S$ and filtering may still be needed, albeit at lower levels,

4. the field gathering procedure is key, as the existence of a singular time step at which the instability is greatly reduced is observed using an "energy conserving" procedure, but not using a "momentum conserving" procedure.

5. Conclusion

An efficient digital filtering procedure and a solver with tunable numerical dispersion have been implemented in the Particle-In-Cell framework Warp. The tunable solver was shown to be compatible with existing "exact" current deposition techniques for conservation of Gauss' Law, and accommodates Perfectly Matched Layers for efficient absorption of outgoing waves. The new techniques have been used to demonstrate efficient direct modeling of laser plasma accelerator stages in the 10 GeV class range, by providing control of an instability that was limiting the speedup of such calculations in previous work.

Extensive testing confirmed, as was discussed elsewhere [25,34], that choosing the frame of the wake as the frame of reference allows for higher levels of filtering and damping than is possible in other frames with the same accuracy. It also revealed that there exists a singular time step for which the level of instability is minimal, independently of other numerical parameters, especially the numerical dispersion of the solver. This indicates that the observed instability may not be caused by standard numerical Cerenkov effects, as had been conjectured [15]. Analysis of the nature of the instability is underway, but regardless of cause, the methods presented mitigate it effectively. The tunability of the field solver is key in providing stability in 3D at the singular time step if using cubic cells, which is not attainable by the standard Yee solver unless the longitudinal resolution is increased or the transverse resolution relaxed, thus either increasing cost or decreasing accuracy. A broad range of parameters produced identical physical simulation outputs, which are consistent with physics scalings based on scaled runs, giving confidence in accuracy.

The use of these techniques permitted the first calculations in the optimal frame of 10 GeV, 100 GeV and 1 TeV class stages, with speedups over 4, 5 and 6 orders of magnitude respectively [34] over what would be required by "standard" laboratory frame calculations, which are impractical for such stages due to computational requirements. Such simulations are required to support 10 GeV class LPA experiments now being planned, and for future accelerators.

Although this was not developed in this paper, the techniques that were presented may have a broader range of applications, to other boosted frame simulations of free electron lasers [26,27] or coherent synchrotron radiation [28], as well as to PIC simulations in general, where an electromagnetic solver with tunable numerical dispersion and/or efficient filtering may be of interest.

Acknowledgments

We are thankful to D.L. Bruhwiler, J.R. Cary, B. Cowan, E. Esarey, A. Friedman, C. Huang, S.F. Martins, W.B. Mori, B.A. Shadwick, and C.B. Schroeder for insightful discussions, as well as to the VORPAL team from Tech-X for providing plots for benchmarking comparisons. Work supported by US-DOE Contracts DE-AC02-05CH11231 and DE-AC52-07NA27344, and US-DOE SciDAC program ComPASS. Used resources of NERSC, supported by US-DOE Contract DE-AC02-05CH11231.

Appendix A. One dimensional analysis of the CK solver

Analysis of the PIC method using the CK solver in one dimension reveals that the scheme is unstable when $c\Delta t = \Delta z$ if no adequate measure is taken to filter the signal at the Nyquist wavelength as explained here (note that the conclusion and mitigation developed in 1D carry to 2D and 3D). In one dimension (choosing z), Eqs. (10) and (11) reduce to

$$B_y|_{l+1/2}^{n+1/2} = B_y|_{l+1/2}^{n-1/2} - \frac{\Delta t}{\Delta z}\left(E_x|_{l+1}^{n} - E_x|_{l}^{n}\right) \tag{47}$$

$$E_x|_{l}^{n+1} = E_x|_{l}^{n} - \frac{c^2\Delta t}{\Delta z}\left(B_y|_{l+1/2}^{n+1/2} - B_y|_{l-1/2}^{n+1/2}\right) - \frac{J_l^n}{\epsilon_0} \tag{48}$$

The averaging that appears in the stencil of the CK solver applies in the plane perpendicular to the direction of wave propagation and is thus not present in 1D. However, since the 1D system can attain the maximum CFL limit of the CK solver, it is equivalent to solving the CK solver in 2D or 3D for a plane wave propagating along a major axis.

Due to uniform time discretization and linearity, the response of the system (47) and (48) to arbitrary distributions and evolutions of sources (i.e. macro-particles) can be written as the sum of its response to the excitation from a Heaviside function in time, at one location in the grid. Assuming a source term of the form $J|_i^n = H(t)$ where H is the Heaviside function, and setting the time step at the Courant limit $c\Delta t = \Delta z$, the system (47) and (48) produces a spurious "odd–even" oscillation at the Nyquist frequency, as shown in Fig. 22 (second row-left). If a sinusoidal signal oscillating at the Nyquist frequency is added to the source term, the amplitude of the spurious oscillation grows linearly with time, as shown in Fig. 22 (second row-right). The spurious oscillation is effectively suppressed in both cases by the application of a "1–2–1" bilinear digital filter, as shown in Fig. 22 (bottom). Without the filtering the total field energy grows exponentially when the system is excited by an oscillatory source, which is unphysical. These types of filtering are of common use in Particle-In-Cell codes, often

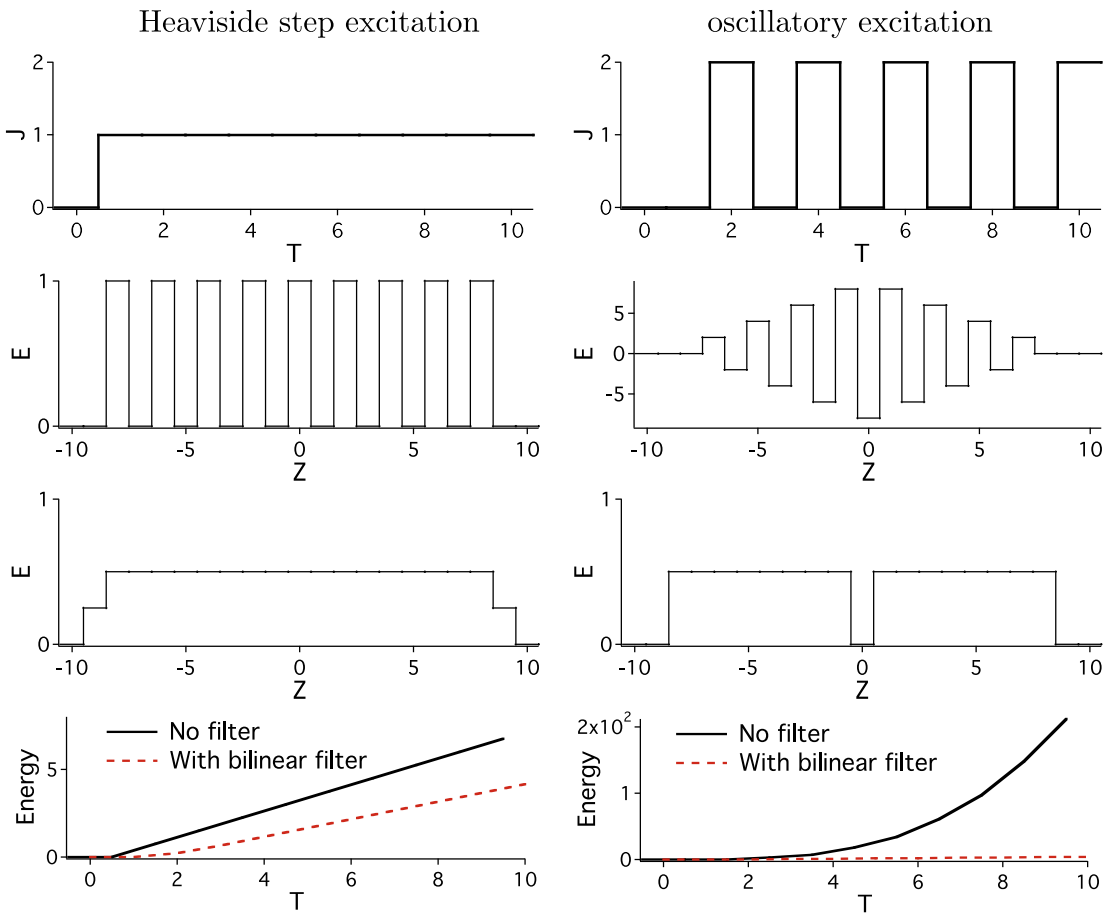

Fig. 22. (Top) Time history (in time steps) of the current source for (left) a Heaviside step (right) a heaviside step modulated by a sinusoidal oscillation at the Nyquist frequency; (second row) response of the system of Eqs. (47) and (48) via a snapshot of the electric field after 10 time steps, without filtering of the source term; (third row) response of the system of Eqs. (47) and (48) with application of bilinear digital filter of the source term in space; (bottom) total field energy history (in time steps). The source term is non-zero at $z = 0$ and zero elsewhere, and the signal has not attained the boundaries ($z = -10$ and $z = 10$) of the simulation grid. A time step of $c\Delta t = \Delta z$ was used in all runs and scaled constants $c = \epsilon_0 = 1$ were assumed.

Table 3
List of parameters for scaled 10 GeV class LPA stage simulation.

Beam length	L_b	85 nm
Beam peak density	n_b	10^{14} cm^{-3}
Beam longitudinal profile		$\exp\left(-z^2/2L_b^2\right)$
Laser wavelength	λ	0.8 μm
Laser length (FWHM)	L	10.08 μm
Normalized vector potential	a_0	1
Laser longitudinal profile		$\sin(\pi z/L)$
Plasma density on axis	n_e	10^{19} cm^{-3}
Plasma longitudinal profile		flat
Plasma length	L	1.5 mm
Plasma entrance ramp profile		half sinus
Plasma entrance ramp length		4 μm
Number of cells	N_z	952
Cell size	Δz	$\lambda/24$
Time step	Δt	$\Delta z/c$
Particle deposition order		cubic
# of plasma particles/cell		10

repeated a prescribed number of times and followed by a compensation stage to avoid excessive damping of long wavelengths [33].

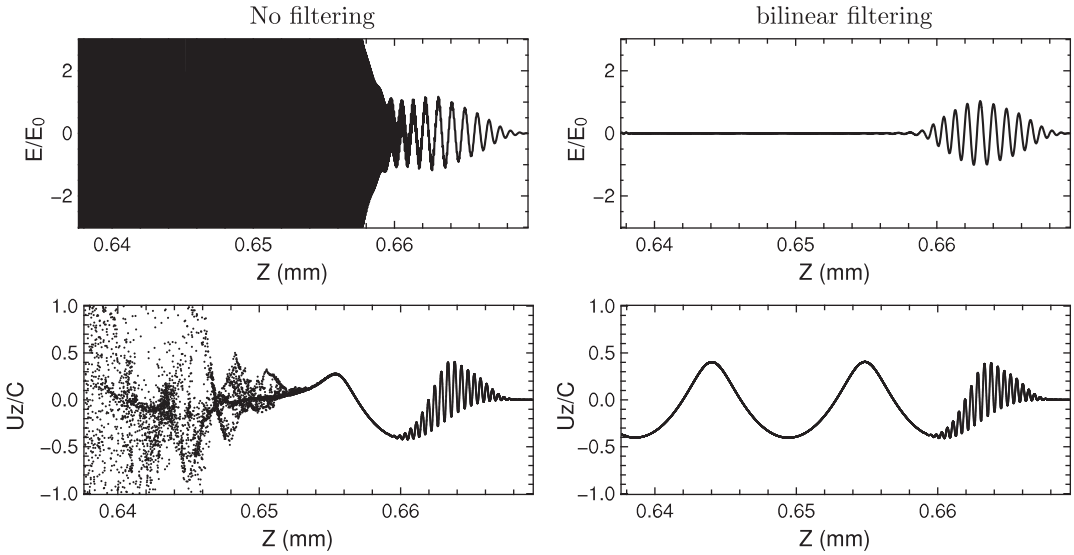

Fig. 23. Snapshots from transverse electric field (normalized to maximum laser amplitude E_0) and plasma electrons longitudinal phase space projection, from a 1D simulation of a laser wakefield acceleration stage the CFL limits ($c\Delta t = \Delta z$) with (left) no filtering of current density; (right) application of a bilinear digital filter to the current density. The simulation used a moving window following the wake and only a small portion of the plasma is present in the plots.

The impact of the spurious oscillations and the effectiveness of the bilinear filtering at suppressing it in actual simulations was tested on a 1D simulation of a scaled wakefield acceleration stage. The physical and numerical parameters of the simulation are given in Table 3. Snapshots of the transverse electric field (aligned with the laser polarization) and the plasma electron phase space, taken once the laser has propagated about half way through the plasma (after ~20,000 time steps) are given in Fig. 23. Without filtering of the current density, an instability develops at the grid Nyquist frequency, severely disrupting the plasma wake, despite the fact that cubic splines were used to deposit current from macro-particles to the grid and gather the electromagnetic field from the grid to the macro-particles. The particles undergo a complicated non-linear response to the instability that is driven by the field. The instability is indeed at the Nyquist frequency but this is not visible on the plot which does not resolve such a small scale. The instability is spuriously pumping some energy into the laser, explaining a larger amplitude of the latter when the instability is present. One application of the bilinear filtering (without compensation) is sufficient to suppress the spurious instability and produce a steady and clean wake.

References

[1] J.-L. Vay, Noninvariance of space- and time-scale ranges under a Lorentz transformation and the implications for the study of relativistic interactions, Physical Review Letters 98 (13) (2007) 130405/1–130405/4.
[2] T. Tajima, J. Dawson, Laser electron-accelerator, Physical Review Letters 43 (4) (1979) 267–270.
[3] E. Esarey, C.B. Schroeder, W.P. Leemans, Physics of laser-driven plasma-based electron accelerators, Reviews of Modern Physics 81 (3) (2009) 1229–1285, doi:10.1103/RevModPhys.81.1229.
[4] C. Geddes, C. Toth, J. van Tilborg, E. Esarey, C. Schroeder, D. Bruhwiler, C. Nieter, J. Cary, W. Leemans, High-quality electron beams from a laser wakefield accelerator using plasma-channel guiding, Nature 431 (7008) (2004) 538–541, doi:10.1038/nature02900.
[5] S. Mangles, C. Murphy, Z. Najmudin, A. Thomas, J. Collier, A. Dangor, E. Divall, P. Foster, J. Gallacher, C. Hooker, D. Jaroszynski, A. Langley, W. Mori, P. Norreys, F. Tsung, R. Viskup, B. Walton, K. Krushelnick, Monoenergetic beams of relativistic electrons from intense laser-plasma interactions, Nature 431 (7008) (2004) 535–538, doi:10.1038/nature02939.
[6] J. Faure, Y. Glinec, A. Pukhov, S. Kiselev, S. Gordienko, E. Lefebvre, J. Rousseau, F. Burgy, V. Malka, A laser-plasma accelerator producing monoenergetic electron beams, Nature 431 (7008) (2004) 541–544, doi:10.1038/nature02963.
[7] W.P. Leemans, B. Nagler, A.J. Gonsalves, C. Toth, K. Nakamura, C.G.R. Geddes, E. Esarey, C.B. Schroeder, S.M. Hooker, Gev electron beams from a centimetre-scale accelerator, Nature Physics 2 (10) (2006) 696–699, doi:10.1038/nphys418.
[8] C. Schroeder, E. Esarey, C. Geddes, C. Toth, W. Leemans, Design considerations for a laser-plasma linear collider, in: AIP Conference Proceedings, vol. 1086, 2009, pp. 208–14.
[9] D. Bruhwiler, J. Cary, B. Cowan, K. Paul, C. Geddes, P. Mullowney, P. Messmer, E. Esarey, E. Cormier-Michel, W. Leemans, J.-L. Vay, New developments in the simulation of advanced accelerator concepts, in: AIP Conference Proceedings, vol. 1086, 2009, pp. 29–37.
[10] J.-L. Vay, et al., Application of the reduction of scale range in a Lorentz boosted frame to the numerical simulation of particle acceleration devices, in: Proceedings of the Particle Accelerator Conference, Vancouver, Canada, 2009, tU1PBI04.
[11] S.F. Martins, et al., Boosted frame pic simulations of lwfa: towards the energy frontier, in: Proceedings of the Particle Accelerator Conference, Vancouver, Canada, 2009, tH4GBC05.
[12] J.-L. Vay, D. Bruhwiler, C. Geddes, W. Fawley, S. Martins, J. Cary, E. Cormier-Michel, B. Cowan, R. Fonseca, M. Furman, W. Lu, W. Mori, L. Silva, Simulating relativistic beam and plasma systems using an optimal boosted frame, Journal of Physics: Conference Series 180 (1) (2009). 5 pp.
[13] C. Huang, W. An, V. Decyk, W. Lu, W. Mori, F. Tsung, M. Tzoufras, S. Morshed, T. Antonsen, B. Feng, T. Katsouleas, R. Fonseca, S. Martins, J. Vieira, L. Silva, E. Esarey, C. Geddes, W. Leemans, E. Cormier-Michel, J.-L. Vay, D. Bruhwiler, B. Cowan, J. Cary, K. Paul, Recent results and future challenges for large scale particle-in-cell simulations of plasma-based accelerator concepts, Journal of Physics: Conference Series 180 (1) (2009) 012005. 11 pp.
[14] J.-L. Vay, et al., Speeding up simulations of relativistic systems using an optimal boosted frame, 2009. <arXiv:0909.5603>.

[15] S.F. Martins, R.A. Fonseca, L.O. Silva, W. Lu, W.B. Mori, Numerical simulations of laser wakefield accelerators in optimal lorentz frames, Computer Physics Communications 181 (5) (2010) 869–875, doi:10.1016/j.cpc.2009.12.023.
[16] S.F. Martins, R.A. Fonseca, W. Lu, W.B. Mori, L.O. Silva, Exploring laser-wakefield-accelerator regimes for near-term lasers using particle-in-cell simulation in Lorentz-boosted frames, Nature Physics 6 (4) (2010) 311–316, doi:10.1038/NPHYS1538.
[17] S.F. Martins, R.A. Fonseca, J. Vieira, L.O. Silva, W. Lu, W.B. Mori, Modeling laser wakefield accelerator experiments with ultrafast particle-in-cell simulations in boosted frames, Physics of Plasmas 17 (5) (2010), doi:10.1063/1.3358139.
[18] J.-L. Vay, et al., Modeling laser wakefield accelerators in a Lorentz boosted frame, in: Proceedings of the 14th Advanced Accelerator Concepts Workshop, Annapolis, MD, 2010.
[19] J. Boris, R. Lee, Nonphysical self forces in some electromagnetic plasma-simulation algorithms, Journal of Computational Physics 12 (1) (1973) 131–136.
[20] I. Haber, R. Lee, H. Klein, J. Boris, in: Proceedings of the Sixth Conference on Numerical Simulation of Plasmas, Berkeley, CA, 1973, pp. 46–48.
[21] B. Godfrey, Numerical Cherenkov instabilities in electromagnetic particle codes, Journal of Computational Physics 15 (4) (1974) 504–521.
[22] B. Godfrey, Canonical momenta and numerical instabilities in particle codes, Journal of Computational Physics 19 (1) (1975) 58–76.
[23] A. Greenwood, K. Cartwright, J. Luginsland, E. Baca, On the elimination of numerical Cerenkov radiation in pic simulations, Journal of Computational Physics 201 (2) (2004) 665–684, doi:10.1016/j.jcp.2004.06.021.
[24] D. Grote, A. Friedman, J.-L. Vay, I. Haber, The warp code: modeling high intensity ion beams, in: AIP Conference Proceedings, vol. 749, 2005, pp. 55–58.
[25] J.-L. Vay, et al., Modeling laser wakefield accelerators in a Lorentz boosted frame, 2010. <arXiv:1009.2727v2>.
[26] W. Fawley, J.-L. Vay, Use of the Lorentz-boosted frame transformation to simulate free-electron laser amplifier physics, in: AIP Conference Proceedings, vol. 1086, 2009, pp. 346–50.
[27] W. Fawley, J.-L. Vay, Full electromagnetic fel simulation via the Lorentz-boosted frame transformation, in: Proceedings of the FEL 2010, paper MOPB01, Sweden, 2010.
[28] W. Fawley, J.-L. Vay, Full electromagnetic simulation of coherent synchrotron radiation via the Lorentz-boosted frame approach, in: Proceedings of the IPAC 2010, paper TUPEC064, Kyoto, Japan, 2010.
[29] B.B. Godfrey, in: Proceedings of the Ninth Conference on Numerical Simulation of Plasmas, 1980.
[30] A. Friedman, A 2nd-order implicit particle mover with adjustable damping, Journal of Computational Physics 90 (2) (1990) 292–312.
[31] H. Abe, N. Sakairi, R. Itatani, H. Okuda, High-order spline interpolations in the particle simulation, Journal of Computational Physics 63 (2) (1986) 247–267.
[32] E. Cormier-Michel, B.A. Shadwick, C.G.R. Geddes, E. Esarey, C.B. Schroeder, W.P. Leemans, Unphysical kinetic effects in particle-in-cell modeling of laser wakefield accelerators, Physical Review E 78 (1, Part 2). doi:10.1103/PhysRevE.78.016404.
[33] C. Birdsall, A. Langdon, Plasma physics via computer simulation, Adam-Hilger, 1991.
[34] J.-L. Vay et al, Effects of hyperbolic rotation in minkowski space on the modeling of plasma accelerators in a Lorentz boosted frame, Physics of Plasmas 18 (2011) 030701/1–4.
[35] J. Cole, A high-accuracy realization of the yee algorithm using non-standard finite differences, IEEE Transactions on Microwave Theory and Techniques 45 (6) (1997) 991–996.
[36] J. Cole, High-accuracy yee algorithm based on nonstandard finite differences: New developments and verifications, IEEE Transactions on Antennas and Propagation 50 (9) (2002) 1185–1191, doi:10.1109/TAP.2002.801268.
[37] M. Karkkainen, E. Gjonaj, T. Lau, T. Weiland, Low-dispersionwake field calculation tools, in: Proceedings of the International Computational Accelerator Physics Conference, Chamonix, France, 2006, pp. 35–40.
[38] B. Cowan, in: Poster of the 10th International Computational Accelerator Physics Conference, San Francisco, CA, 2009.
[39] J. Berenger, Three-dimensional perfectly matched layer for the absorption of electromagnetic waves, Journal of Computational Physics 127 (2) (1996) 363–379.
[40] A.B. Langdon, On enforcing gauss law in electromagnetic particle-in-cell codes, Computer Physics Communications 70 (3) (1992) 447–450.
[41] B. Marder, A method for incorporating Gauss law into electromagnetic pic codes, Journal of Computational Physics 68 (1) (1987) 48–55.
[42] J.-L. Vay, C. Deutsch, Charge compensated ion beam propagation in a reactor sized chamber, Physics of Plasmas 5 (4) (1998) 1190–1197.
[43] J. Villasenor, O. Buneman, Rigorous charge conservation for local electromagnetic-field solvers, Computer Physics Communications 69 (2–3) (1992) 306–316.
[44] T. Esirkepov, Exact charge conservation scheme for particle-in-cell simulation with an arbitrary form-factor, Computer Physics Communications 135 (2) (2001) 144–153.
[45] J. Berenger, A perfectly matched layer for the absorption of electromagnetic-waves, Journal of Computational Physics 114 (2) (1994) 185–200.
[46] J.-L. Vay, A new absorbing layer boundary condition for the wave equation, Journal of Computational Physics 165 (2) (2000) 511–521.
[47] J.-L. Vay, Asymmetric perfectly matched layer for the absorption of waves, Journal of Computational Physics 183 (2) (2002) 367–399, doi:10.1006/jcph.2002.7175.
[48] E. Cormier-Michel, C. Geddes, E. Esarey, C. Schroeder, D. Bruhwiler, K. Paul, B. Cowan, W. Leemans, Scaled simulations of a 10 gev accelerator, in: AIP Conference Proceedings, vol. 1086, 2009, pp. 297–302.
[49] C.G.R. Geddes, et al., Scaled simulation design of high quality laser wakefield accelerator stages, in: Proceedings of the Particle Accelerator Conference, Vancouver, Canada, 2009, wE6RFP075.
[50] F. Tsung, W. Lu, M. Tzoufras, W. Mori, C. Joshi, J. Vieira, L. Silva, R. Fonseca, Simulation of monoenergetic electron generation via laser wakefield accelerators for 5-25 tw lasers, Physics of Plasmas 13 (5) (2006), doi:10.1063/1.2198535.
[51] B. Cowan, Private communication, 2010.
[52] C. Nieter, J.R. Cary, Journal of Computational Physics 196 (2004) 448–472.