

Modeling laser wakefield accelerators in a Lorentz boosted frame

by

J.-L. Vay, C. G. R. Geddes, E. Cormier-Michel, D. P. Grote

Accelerator and Fusion Research Division
Ernest Orlando Lawrence Berkeley National Laboratory
Berkeley, California 94720

And

Tech-X Corp., Boulder, CO, USA

and

LLNL, Livermore, CA, USA

June 2010

This work was supported by the Director, Office of Science, Office of Fusion Energy Sciences, of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231.

This document was prepared as an account of work sponsored by the United States Government. While this document is believed to contain correct information, neither the United States Government nor any agency thereof, nor The Regents of the University of California, nor any of their employees, makes any warranty, express or implied, or assumes any legal responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by its trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof, or The Regents of the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof or The Regents of the University of California.

This work was supported by the Director, Office of Science, Office of Fusion Energy Sciences, of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231.

Modeling 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*LBNL, Berkeley, CA, USA*

^b*Tech-X Corp., Boulder, CO, USA*

^c*LLNL, Livermore, CA, USA*

Abstract

Modeling of laser-plasma wakefield accelerators in an optimal frame of reference [1] is shown to produce orders of magnitude speed-up of calculations from first principles. Obtaining these speedups requires mitigation of a high-frequency instability that otherwise limits effectiveness in addition to solutions for handling data input and output in a relativistically boosted frame of reference. The observed high-frequency instability is mitigated using methods 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 shown 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 for the same accuracy. Detailed testing also revealed serendipitously the existence of a singular time step at which the instability level is minimized, independently of numerical dispersion, thus indicating that the observed instability may not be due primarily to Numerical Cerenkov as has been conjectured. The techniques developed for Cerenkov mitigation prove nonetheless to be very efficient at controlling the instability. Using these techniques, agreement

at the percentage level is demonstrated between simulations using different frames of reference, with speedups reaching two orders of magnitude for a 0.1 GeV class stages. The method then allows direct and efficient full-scale modeling of deeply depleted laser-plasma stages of 10 GeV-1 TeV for the first time, verifying the scaling of plasma accelerators to very high energies. Over 4, 5 and 6 orders of magnitude speedup is achieved for the modeling of 10 GeV, 100 GeV and 1 TeV class stages, respectively.

Keywords: laser wakefield acceleration, particle-in-cell, plasma simulation, special relativity, frame of reference, boosted frame

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]. High quality electron beams of energy up-to 1 GeV have been produced in just a few centimeters [4, 5, 6, 7], with 10 GeV stages being planned as modules of a high energy collider [8].

As the laser propagates through a plasma, it displaces electrons while ions remain essentially static, 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

16 conventional particle accelerators. The efficiency and quality of the accelera-
17 tion is governed by several factors which require precise three-dimensional
18 shaping of the plasma column, as well as the laser and particle beams, and
19 understanding of their evolution.

20 Computer simulations have had a profound impact on the design and
21 understanding of past and present experiments [9], with accurate modeling
22 of the wake formation and beam acceleration requiring fully kinetic meth-
23 ods (usually Particle-In-Cell) with large computational resources due to the
24 wide range of space and time scales involved [10, 11]. For example, modeling
25 10 GeV stages for the LOASIS (LBNL) BELLA proposal [12] demanded as
26 many as 5,000 processor hours for a one-dimension simulation on a NERSC
27 supercomputer [13]. Various reduced models have been developed to allow
28 multidimensional simulations at manageable computational costs: fluid ap-
29 proximation [14], quasistatic approximation [15, 16, 17], laser envelope mod-
30 els [16], scaled parameters [18, 19]. However, the various approximations
31 that they require result in a narrower range of applicability. As a result,
32 even using several models concurrently does not usually provide a complete
33 description. For example, scaled simulations of 10 GeV LPA stages do not
34 capture correctly some essential transverse physics, e.g. the laser and beam
35 betatron motion, which can lead to inaccurate beam emittance (a measure
36 of the beam quality). An envelope description can capture these effects cor-
37 rectly at full scale for the early propagation through the plasma but can
38 fail as the laser spectrum broadens due to energy depletion as it propagates
39 further in the plasma.

40 An alternative approach allows for orders of magnitude speedup of simu-

41 lations, whether at full or reduced scale, via the proper choice of a reference
42 frame moving near the speed of light in the direction of the laser [1]. It
43 does so without alteration to the fundamental equations of particle motion
44 or electrodynamics, provided that the high-frequency part of the light emit-
45 ted counter to the direction of propagation of the beam can be neglected.
46 This approach exploits the properties of space and time dilation and con-
47 traction associated with the Lorentz transformation. As shown in [1], the
48 ratio of longest to shortest space and time scales of a system of two or more
49 components crossing at relativistic velocities is not invariant under such a
50 transformation (a laser crossing a plasma is just such a relativistic crossing).
51 Since the number of computer operations (e.g., time steps), for simulations
52 based on formulations from first principles, is proportional to the ratio of
53 the longest to shortest time scale of interest, it follows that such simulations
54 will eventually have different computer runtimes, yet equivalent accuracy,
55 depending solely upon the choice of frame of reference.

56 The procedure appears straightforward: identify the frame of reference
57 which will minimize the range of space and/or time scales and perform the
58 calculation in this frame. However, several practical complications arise.
59 First, the input and output data are usually known from, or compared to,
60 experimental data. Thus, calculating in a frame other than the laboratory
61 entails transformations of the data between the calculation frame and the lab-
62 oratory frame. Second, while the fundamental equations of electrodynamics
63 and particle motion are written in a covariant form, the numerical algorithms
64 that are derived from them may not retain this property, and calculations in
65 frames moving at different velocities may not be successfully conducted with

66 the use of the exact same algorithms. For example, it was shown in [20] that
67 calculating the propagation of ultra-relativistic charged particle beams in an
68 accelerator using standard Particle-In-Cell techniques lead to large numerical
69 errors, which were fixed by developing a new particle pusher. The modeling
70 of a laser plasma accelerator (LPA) stage in a boosted frame involves the
71 fully electromagnetic modeling of a plasma propagating at near the speed of
72 light, for which Numerical Cerenkov [21, 22] is a potential issue. Third, elec-
73 tromagnetic calculations that include wave propagation will include waves
74 propagating forward and backward in any direction. For a frame of reference
75 moving in the direction of the accelerated beam (or equivalently the wake of
76 the laser), waves emitted by the plasma in the forward direction expand while
77 the ones emitted in the backward direction contract, following the proper-
78 ties of the Lorentz transformation. If one is to resolve both forward and
79 backward propagating waves emitted from the plasma, there is no gain in
80 selecting a frame different from the laboratory frame. However, the physics
81 of interest for a laser wakefield is the laser driving the wake, the wake, and
82 the accelerated beam. Backscatter is weak in the short-pulse regime, and
83 does not interact as strongly with the beam as do the forward propagating
84 waves which stay in phase for a long period. It is thus often assumed that
85 the backward propagating waves can be neglected in the modeling of LPA
86 stages. The accuracy of this assumption is shown by comparison between
87 explicit codes which include both forward and backward waves and envelope
88 or quasistatic codes which neglect backward waves [10, 19, 23].

89 After the idea and basic scaling for performing simulations of LPA in a
90 Lorentz boosted frame were published in [1], there have been several reports

91 of the application of the technique to various regimes of LPA [13, 24, 25,
92 26, 11, 27, 28]. Speedups varying between several to a few thousands were
93 reported with various levels of accuracy in agreement between simulations
94 performed in a Lorentz boosted frames and in a laboratory frame. High-
95 frequency instabilities were reported to develop in 2D or 3D calculations,
96 that were limiting the velocity of the boosted frame and thus the attainable
97 speedup [29, 27, 28].

98 In this paper, we present numerical techniques that were implemented in
99 the Particle-In-Cell code Warp [30] for mitigating the numerical Cerenkov
100 instability, including a solver with tunable coefficients, and show that these
101 techniques are effective for suppressing the high frequency instability ob-
102 served in boosted frame simulations. A detailed study of the application
103 of these techniques to the simulations of downscaled LPA stages reveals
104 that choosing the frame of the wakefield as the reference frame allows for
105 more aggressive application of the standard techniques mitigating numerical
106 Cerenkov, than is possible in laboratory frame simulations. It is shown that
107 the instability that develops with high-boost frames is well controlled, allow-
108 ing for the first time 2D and 3D simulations of LPA in the wakefield frame,
109 for 100 GeV and 1 TeV class stages, achieving the maximum theoretical
110 speedups of over 10^5 and 10^6 respectively.

111 This paper is organized as follows. The theoretical speedup expected for
112 performing the modeling of a LPA stage in a boosted frame is derived in
113 Section 2. Section 3 addresses the issue of input and output of data in a
114 boosted frame. High frequency instability issues and remedies are presented
115 in Section 4. These techniques enable accurate modeling of 0.1 GeV-1 TeV

116 LPA stages. Stage modeling results are presented in section 5, and observed
117 speedup is contrasted to the theoretical speedup of section 2.

118 **2. Theoretical speedup dependency with the frame boost**

119 The obtainable speedup is derived as an extension of the formula that
120 was derived in [1], taking in addition into account the group velocity of the
121 laser as it traverses the plasma. In [1], the laser was assumed to propagate
122 at the velocity of light in vacuum during the entire process, which is a good
123 approximation when the relativistic factor of the frame boost γ is small
124 compared to the relativistic factor of the laser wake γ_w in the plasma. The
125 expression is generalized here to higher values of γ , for which the actual group
126 velocity of the wake in the plasma must be taken into account. We shall show
127 that for a 10 GeV class LPA stage, the maximum attainable speedup is above
128 four orders of magnitude.

129 Assuming that the simulation box is a fixed number of plasma periods
130 long, which implies the use (which is standard) of a moving window following
131 the wake and accelerated beam, the speedup is given by the ratio of the time
132 taken by the laser pulse and the plasma to cross each other, divided by
133 the shortest time scale of interest, that is the laser period. Assuming for
134 simplicity that the wake propagates at the group velocity of plane waves in
135 a uniform plasma of density n_e , the group velocity of the wake is given by

$$v_w/c = \beta_w = \left(1 + \frac{\omega_p^2}{\omega^2}\right)^{-1/2} \quad (1)$$

136 where $\omega_p = \sqrt{(n_e e^2)(\epsilon_0 m_e)}$ is the plasma frequency, $\omega = 2\pi c/\lambda$ is the
137 laser frequency, ϵ_0 is the permittivity of vacuum, c is the speed of light in

138 vacuum, and e and m_e are respectively the charge and mass of the electron.

139 In the simulations presented herein, the runs are stopped when the last
 140 electron beam macro-particle exits the plasma, and a measure of the total
 141 time of the simulation is given by

$$T = \frac{L + \eta\lambda_p}{v_w - v_p} \quad (2)$$

142 where $\lambda_p \approx 2\pi c/\omega_p$ is the wake wavelength, L is the plasma length, v_w and
 143 $v_p = \beta_p c$ are respectively the velocity of the wake and of the plasma relative
 144 to the frame of reference, and η is an adjustable parameter for taking into
 145 account the fraction of the wake which exited the plasma at the end of the
 146 simulation. The numerical cost R_t scales as the ratio of the total time to the
 147 shortest timescale of interest, which is the inverse of the laser frequency, and
 148 is thus given by

$$R_t = \frac{Tc}{\lambda} = \frac{(L + \eta\lambda_p)}{(\beta_w - \beta_p)\lambda} \quad (3)$$

149 In the laboratory, $v_p = 0$ and the expression simplifies to

$$R_{lab} = \frac{Tc}{\lambda} = \frac{(L + \eta\lambda_p)}{\beta_w\lambda} \quad (4)$$

150 In a frame moving at βc , the quantities become

$$\lambda_p^* = \lambda_p / [\gamma(1 - \beta_w\beta)] \quad (5)$$

$$L^* = L/\gamma \quad (6)$$

$$\lambda^* = \gamma(1 + \beta)\lambda \quad (7)$$

$$\beta_w^* = (\beta_w - \beta) / (1 - \beta_w\beta) \quad (8)$$

$$v_p^* = -\beta c \quad (9)$$

$$T^* = \frac{L^* + \eta\lambda_p^*}{v_w^* - v_p^*} \quad (10)$$

$$R_t^* = \frac{T^* c}{\lambda^*} = \frac{(L^* + \eta\lambda_p^*)}{(\beta_w^* + \beta) \lambda^*} \quad (11)$$

151 where $\gamma = 1/\sqrt{1 - \beta^2}$.

152 The expected speedup from performing the simulation in a boosted frame
153 is given by the ratio of R_{lab} and R_t^*

$$S = \frac{R_{lab}}{R_t^*} = \frac{(1 + \beta)(L + \eta\lambda_p)}{(1 - \beta\beta_w)L + \eta\lambda_p} \quad (12)$$

154 Assuming that $\gamma \ll \gamma_w$, and that $\beta_w \approx 1$ (which is a valid approximation
155 for most practical cases of interest), this expression is consistent with the ex-
156 pression derived in [1] for the LPA case which states that $R_t^* = \alpha R_t / (1 + \beta)$
157 with $\alpha = (1 - \beta + l/L) / (1 + l/L)$, where l is the laser length which is gen-
158 erally proportional to $\eta\lambda_p$, and $S = R_t / R_t^*$.

159 The linear theory predicts that for the intense lasers ($a \gtrsim 1$) typically used
160 for acceleration, the laser depletes its energy over approximately the same
161 length $L_d = \lambda_p^3 / 2\lambda^2$ over which the particles dephase from the wake [2].
162 Acceleration is compromised beyond L_d and in practice, the plasma length is
163 proportional to the dephasing length, i.e. $L = \xi L_d$. In most cases, $\gamma_w^2 \gg 1$,
164 which allows the approximations $\beta_w \approx 1 - \lambda^2 / 2\lambda_p^2$, and $L = \xi \lambda_p^3 / 2\lambda^2 \approx$
165 $\xi \gamma_w^2 \lambda_p / 2 \gg \eta\lambda_p$, so that Eq.(12) becomes

$$S = (1 + \beta)^2 \gamma^2 \frac{\xi \gamma_w^2}{\xi \gamma_w^2 + (1 + \beta) \gamma^2 (\xi \beta / 2 + 2\eta)} \quad (13)$$

166 For low values of γ , i.e. when $\gamma \ll \gamma_w$, Eq.(13) reduces to

$$S_{\gamma \ll \gamma_w} = (1 + \beta)^2 \gamma^2 \quad (14)$$

167 Conversely, if $\gamma \rightarrow \infty$, Eq.(13) becomes

$$S_{\gamma \rightarrow \infty} = \frac{4}{1 + 4\eta/\xi} \gamma_w^2 \quad (15)$$

168 Finally, in the frame of the wake, i.e. when $\gamma = \gamma_w$, assuming that $\beta_w \approx 1$,
 169 Eq.(13) gives

$$S_{\gamma=\gamma_w} \approx \frac{2}{1 + 2\eta/\xi} \gamma_w^2 \quad (16)$$

170 Since η and ξ are of order unity, and the practical regimes of most interest
 171 satisfy $\gamma_w^2 \gg 1$, the speedup that is obtained by using the frame of the wake
 172 will be near the maximum obtainable value given by Eq.(15).

173 Note that without the use of a moving window, the relativistic effects that
 174 are at play in the time domain would also be at play in the spatial domain,
 175 as shown in [1], and the γ^2 scaling would transform to γ^4 . In the frame
 176 of the wake, there is no need of the moving window, thus simplifying the
 177 procedure, while in a frame traveling faster than the wake in the laboratory,
 178 a moving window propagating in the backward direction is needed. However,
 179 the scaling shows that there would be very little gain in doing the latter.

180 *2.1. Estimated speedup for 0.1-100 GeV stages*

181 Formula (13) is used to estimate the speedup for the calculations of 100
 182 MeV, 1 GeV, 10 GeV and 100 GeV class stages, assuming a laser wavelength
 183 $\lambda = 0.8\mu m$. Using parameters and scaling laws from [18], the corresponding
 184 initial plasma densities n_e are respectively $10^{19}cc$, $10^{18}cc$, $10^{17}cc$ and $10^{16}cc$,
 185 while the plasma lengths L are 1.5 mm, 4.74 cm, 1.5 m, and 47.4 m, with
 186 $\xi \approx 1.63$. For these values, the wake wavelengths λ_p are respectively $10.6\mu m$,

187 $33.4\mu m$, $106.\mu m$, $334.\mu m$, and relativistic factors γ_w are 13.2, 41.7, 132 and
 188 417. In the simulations presented in this paper, the beam is injected near the
 189 end of the wake period (first "bucket"). In first approximation, the beam
 190 has propagated through about half a wake period to reach full acceleration,
 191 and we set $\eta \approx 0.5$. For a beam injected into the n^{th} bucket, η would be set
 192 to $n - 1/2$. If positrons were considered, they would be injected half a wake
 193 period ahead of the location of the electrons injection position for a given
 194 period, and one would have $\eta = n - 1$. For the parameters considered here,
 195 $L \approx \lambda_p/\gamma_w^2$, and (15) gives $S_{\gamma \rightarrow \infty} \approx 2\gamma_w^2$.

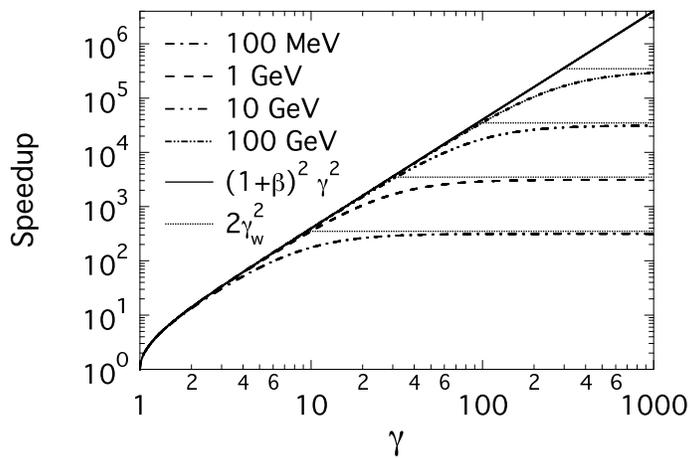


Figure 1: Speedup versus relativistic factor of the boosted frame from Eq.(13) for 100 MeV - 100 GeV LPA class stages.

196 The speedup versus the relativistic factor of the boosted frame γ is plotted
 197 in Fig. 1. As expected, for low values of γ , the speedup scales as (14), and
 198 asymptotes to a value slightly lower than $2\gamma_w^2$ for large values of γ . It is of
 199 interest to note that the qualitative behavior is identical to the one obtained
 200 in [1] (see Fig. 1 and accompanying analysis) in the analysis of the crossing

201 of two rigid identical beams, confirming the generality of the generic analysis
202 presented in [1]. For a 100 GeV class stage, the maximum estimated speedup
203 is as large as 300,000.

204 **3. Input and output to and from a boosted frame simulation**

205 This section describes the procedures that have been implemented in the
206 Particle-In-Cell framework Warp [30] to handle the input and output of data
207 between the frame of calculation and the laboratory frame. Simultaneity of
208 events between two frames is valid only for a plane that is perpendicular
209 to the relative motion of the frame. As a result, the input/output processes
210 involve the input of data (particles or fields) through a plane, as well as output
211 through a series of planes, all of which are perpendicular to the direction of
212 the relative velocity between the frame of calculation and the other frame of
213 choice.

214 *3.1. Input*

215 *3.1.1. Particles*

216 Particles are launched through a plane using a technique which applies
217 to many calculations in a boosted frame, including LPA, and is illustrated
218 using the case of a positively charged particle beam propagating through a
219 background of cold electrons in an assumed continuous transverse focusing
220 system, leading to a growing transverse instability [1]. In the laboratory
221 frame, the electron background is initially at rest and a moving window is
222 used to follow the beam progression. Traditionally, the beam macropar-
223 ticles are initialized all at once in the window, while background electron

224 macroparticles are created continuously in front of the beam on a plane that
 225 is perpendicular to the beam velocity. In a frame moving at some fraction
 226 of the beam velocity in the laboratory frame, the beam initial conditions at
 227 a given time in the calculation frame are generally unknown and one must
 228 initialize the beam differently. However, it can be taken advantage of the
 229 fact that the beam initial conditions are often known for a given plane in the
 230 laboratory, either directly, or via simple calculation or projection from the
 231 conditions at a given time. Given the position and velocity $\{x, y, z, v_x, v_y, v_z\}$
 232 for each beam macroparticle at time $t = 0$ for a beam moving at the aver-
 233 age velocity $v_b = \beta_b c$ (where c is the speed of light) in the laboratory, and
 234 using the standard synchronization ($z = z' = 0$ at $t = t' = 0$) between the
 235 laboratory and the calculation frames, the procedure for transforming the
 236 beam quantities for injection in a boosted frame moving at velocity βc in the
 237 laboratory is as follows (the superscript ' relates to quantities known in the
 238 boosted frame while the superscript * relates to quantities that are know at
 239 a given longitudinal position z^* but different times of arrival):

- 240 1. project positions at $z^* = 0$ assuming ballistic propagation

$$t^* = (z - \bar{z})/v_z \quad (17)$$

$$x^* = x - v_x t^* \quad (18)$$

$$y^* = y - v_y t^* \quad (19)$$

$$z^* = 0 \quad (20)$$

241 the velocity components being left unchanged,

- 242 2. apply Lorentz transformation from laboratory frame to boosted frame

$$t'^* = -\gamma t^* \quad (21)$$

$$x'^* = x^* \quad (22)$$

$$y'^* = y^* \quad (23)$$

$$z'^* = \gamma\beta ct^* \quad (24)$$

$$v_x'^* = \frac{v_x^*}{\gamma(1 - \beta\beta_b)} \quad (25)$$

$$v_y'^* = \frac{v_y^*}{\gamma(1 - \beta\beta_b)} \quad (26)$$

$$v_z'^* = \frac{v_z^* - \beta c}{1 - \beta\beta_b} \quad (27)$$

243 where $\gamma = 1/\sqrt{1 - \beta^2}$. With the knowledge of the time at which each
 244 beam macroparticle crosses the plane into consideration, one can inject
 245 each beam macroparticle in the simulation at the appropriate location
 246 and time.

- 247 3. synchronize macroparticles in boosted frame, obtaining their positions
 248 at a fixed $t' (= 0)$ which is before any particle is injected

$$z' = z'^* - \bar{v}_z'^* t'^* \quad (28)$$

249 This additional step is needed for setting the electrostatic or electro-
 250 magnetic fields at the plane of injection. In a Particle-In-Cell code, the
 251 three-dimensional fields are calculated by solving the Maxwell equa-
 252 tions (or static approximation like Poisson, Darwin or other [20]) on
 253 a grid on which the source term is obtained from the macroparticles
 254 distribution. This requires generation of a three-dimensional represen-
 255 tation of the beam distribution of macroparticles at a given time before
 256 they cross the injection plane at z'^* . This is accomplished by expand-
 257 ing the beam distribution longitudinally such that all macroparticles
 258 (so far known at different times of arrival at the injection plane) are

259 synchronized to the same time in the boosted frame. To keep the beam
 260 shape constant, the particles are "frozen" until they cross that plane:
 261 the three velocity components and the two position components per-
 262 pendicular to the boosted frame velocity are fixed, while the remaining
 263 position component is advanced at the average beam velocity. As parti-
 264 cles cross the plane of injection, they become regular "active" particles
 265 with full 6-D dynamics.

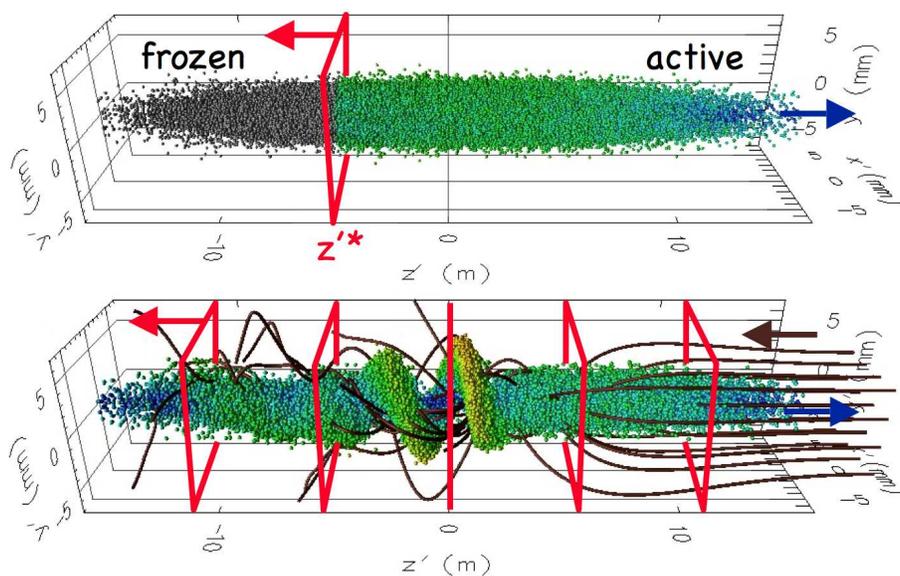


Figure 2: (top) Snapshot of a particle beam "frozen" (grey spheres) and "active" (colored spheres) macroparticles traversing the injection plane (red rectangle). (bottom) Snapshot of the beam macroparticles (colored spheres) passing through the background of electrons (dark brown streamlines) and the diagnostic stations (red rectangles). The electrons, the injection plane and the diagnostic stations are fixed in the laboratory plane, and are thus counterpropagating to the beam in a boosted frame.

266 Figure 2 (top) shows a snapshot of a beam that has passed partly through
 267 the injection plane. As the frozen beam macroparticles pass through the

268 injection plane (which moves opposite to the beam in the boosted frame),
 269 they are converted to “active” macroparticles. The charge or current density
 270 is accumulated from the active and the frozen particles, thus ensuring that
 271 the fields at the plane of injection are consistent.

272 3.1.2. Laser

273 Similarly to the particle beam, the laser is injected as an electric field
 274 E_{\perp} through a plane perpendicular to the axis of propagation of the laser (by
 275 default z), using the formula

$$E_{\perp}(x, y, t) = E_0 f(x, y, t) \sin[\omega t + \phi(x, y, \omega)] \quad (29)$$

276 where E_0 is the amplitude of the laser electric field, $f(x, y, t)$ is the laser
 277 envelope, ω is the laser frequency, $\phi(x, y, \omega)$ is a phase function to account
 278 for focusing, defocusing or injection at an angle, and t is time. By default,
 279 the laser envelope is a three dimensional gaussian of the form

$$f(x, y, t) = e^{-(x^2/2\sigma_x^2 + y^2/2\sigma_y^2 + c^2 t^2/2\sigma_z^2)} \quad (30)$$

280 where σ_x , σ_y and σ_z are the dimensions of the laser pulse; or it can be defined
 281 arbitrarily by the user at runtime. If $\phi(x, y, \omega) = 1$, the laser is injected at
 282 a waist and parallel to the axis z .

283 If, for convenience, the injection plane is moving at constant velocity $\beta_s c$,
 284 the formula is modified to take the Doppler effect on frequency and amplitude
 285 into account and becomes

$$E_{\perp}(x, y, t) = (1 - \beta_s) E_0 f(x, y, t) \sin[(1 - \beta_s) \omega t + \phi(x, y, \omega)]. \quad (31)$$

286 The laser field is added to the transverse electric field components as

$$E_x(x, y, z) = E_x(x, y, z) + \zeta(z, t) \frac{c \delta t}{\delta z} \cos(\psi) E_{\perp} \quad (32)$$

$$E_y(x, y, z) = E_y(x, y, z) + \zeta(z, t) \frac{c\delta t}{\delta z} \sin(\psi) E_\perp \quad (33)$$

287 where ψ accounts for polarization and $\zeta(z, t)$ is an interpolation function.
 288 The latter is necessary since $\beta_s c\delta t \neq \delta z$ in general, and the plane of injec-
 289 tion may not coincide with a grid plane; or to provide smoothing of high
 290 frequencies if needed (see discussion in section 4.2 below).

291 The injection of a laser of frequency ω is considered for a simulation using
 292 a boosted frame moving at βc with respect to the laboratory. Assuming that
 293 the laser is injected at a plane that is fixed in the laboratory, and thus moving
 294 at $\beta_s = -\beta$ in the boosted frame, the injection in the boosted frame is given
 295 by

$$\begin{aligned} E_\perp(x', y', t') &= (1 - \beta_s) E'_0 f(x', y', t') \sin[(1 - \beta_s)\omega t' + \phi(x', y', \omega)] \\ &= (E_0/\gamma) f(x', y', t') \sin[\omega t'/\gamma + \phi(x', y', \omega)] \end{aligned} \quad (34)$$

296 since $E'_0/E_0 = \omega'/\omega = 1/(1 + \beta)\gamma$.

297 The technique implemented in Warp presents several advantage over other
 298 procedures that have been proposed elsewhere [13, 28]. In [28], the laser
 299 beam is initialized entirely in the computational box, leading to larger boxes
 300 transversely in a boosted frame, as the Rayleigh length of the laser shortens
 301 and the overall laser pulse radius rises, eventually offsetting the benefits of
 302 the boosted frame. The transverse broadening of the box is avoided in [13]
 303 at the cost of a more complicated injection scheme, requiring to launch the
 304 laser from all but one faces of the simulation box. The method presented here
 305 avoids the caveat of the broadening and retains simplicity with a standard
 306 injection technique through one plane.

307 3.2. Output

308 Some quantities, e.g. charge, are Lorentz invariant, while others, like di-
309 mensions perpendicular to the boost velocity, are the same in the laboratory
310 frame. Those quantities are thus readily available from standard diagnostics
311 in the boosted frame calculations. Quantities which do not fall in this cat-
312 egory are recorded at a number of regularly spaced “stations”, immobile in
313 the laboratory frame, at a succession of time intervals to record data history,
314 or averaged over time. A visual example is given on Fig. 2 (bottom). Since
315 the space-time locations of the diagnostic grids in the laboratory frame gen-
316 erally do not coincide with the space-time positions of the macroparticles and
317 grid nodes used for the calculation in a boosted frame, some interpolation is
318 performed at runtime during the data collection process. As a complement
319 or an alternative, selected particle or field quantities are dumped at regu-
320 lar interval for post-processing. The choice of the methods depends on the
321 requirements of the diagnostics and particular implementations.

322 4. High frequency instability and Numerical Cerenkov

323 As reported in [27] and [28], for simulations using a boosted frame at
324 $\gamma \geq 10 - 20$ (depending on parameters), a fast growing short wavelength
325 instability was observed developing at the front of the plasma (see Fig. 3).
326 The presence and growth rate of the instability was observed to be very sen-
327 sitive to the resolution (slower growth rate at higher resolution), choice of
328 field solver, and to the amount of damping of high frequencies and smoothing
329 of short wavelengths. The instability is always propagating at some angle
330 from the longitudinal axis, and is observed in 2D and 3D runs but was never

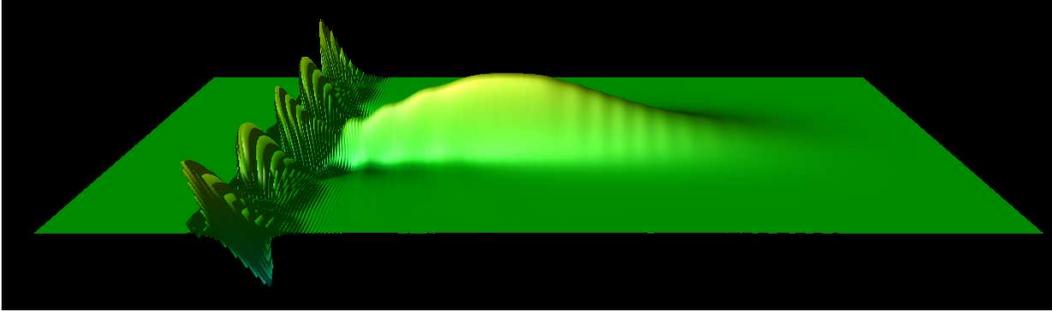


Figure 3: Snapshot of a surface plot of the longitudinal field from a 2-1/2D simulation of a full scale 10GeV LPA in a boosted frame at $\gamma = 130$ (elevation is proportional to the magnitude of the electric field). The laser is propagating from left to right and the plasma from right to left. A fast growing short wavelength instability is developing at the front of the plasma.

331 observed in any of the 1D runs performed by the authors. When modeling
 332 an LPA setup in a relativistically boosted frame, the background plasma is
 333 traveling near the speed of light and it has been conjectured [28] that he
 334 observed instability might be caused by numerical Cerenkov. We investigate
 335 in this paper whether the instability that is observed in boosted frame sim-
 336 ulations of LPA is indeed of numerical Cerenkov type and if the cures aimed
 337 at mitigating numerical Cerenkov are effective.

338 Due to spatial and time discretization of the Maxwell equations, numerical
 339 light waves may travel faster or slower on the computational grid than the
 340 actual speed of light in vacuum c , with the magnitude of the effect being
 341 larger at short wavelength, where discretization errors are the largest. When
 342 the numerical speed is lower than c , it is possible for fast macro-particles
 343 to travel faster than the wave modes, leading to numerical Cerenkov effects
 344 that may result in instabilities [21, 22, 31, 32, 33]. The effect was studied

345 analytically and numerically in detail for one-dimensional systems in [31,
346 32]. Several solutions were proposed: smoothing the current deposited by
347 the macro-particles [21, 31], damping the electromagnetic field [33, 34, 35],
348 solving the Maxwell equations in Fourier space [22], or using a field solver
349 with a larger stencil to provide lower numerical dispersion [33].

350 Several of the abovementioned techniques to mitigate numerical Cerenkov
351 and high frequency errors have been implemented in Warp. All the simula-
352 tions presented in this paper employed cubic splines for current deposition
353 and electromagnetic force gathering between the macro-particles and the
354 grid [36], whose beneficial effects on standard LPA PIC simulations have
355 been demonstrated in [37]. In addition, a Maxwell solver with tunable coef-
356 ficients was implemented, as well as a damping scheme, and filtering of the
357 deposited current and gathered electromagnetic fields, which are described
358 in this section. The use of Fourier based Maxwell solvers is not considered
359 in this paper.

360 4.1. Wideband lowpass digital filtering

361 It is common practice to apply digital filtering to the charge or current
 362 density in Particle-In-Cell simulations, for smoothing or compensation pur-
 363 pose [46]. The most commonly used filter is the three points filter

$$\phi_j^f = \alpha\phi_j + (1 - \alpha) \frac{\phi_{j-1} + \phi_{j+1}}{2} \quad (36)$$

364 where ϕ^f is the filtered quantity. This filter is called a binomial filter when
 365 $\alpha = 0.5$. Assuming $\phi = e^{jkx}$ and $\phi^f = g(\alpha, k) e^{jkx}$, where g is the filter gain,
 366 which is function of the filtering coefficient α and the wavenumber k , we find
 367 from (36) that

$$g(\alpha, k) = \alpha + (1 - \alpha) \cos(k\delta x) \quad (37)$$

$$\approx 1 - (1 - \alpha) \frac{(k\delta x)^2}{2} + O(k^4) \quad (38)$$

368 For n successive applications of filters of coefficients $\alpha_1 \dots \alpha_n$, the total atten-
 369 uation G is given by

$$G = \prod_{i=1}^n g(\alpha_i, k) \quad (39)$$

$$\approx 1 - \left(n - \sum_{i=1}^n \alpha_i \right) \frac{(k\delta x)^2}{2} + O(k^4) \quad (40)$$

370 If $\alpha_n = n - \sum_{i=1}^{n-1} \alpha_i$ then $G \approx 1 + O(k^4)$, providing a sharper cutoff in k
 371 space. Such step is called a *compensation* step [46]. For the bilinear filter
 372 ($\alpha = 1/2$), the compensation factor is $\alpha_c = 2 - 1/2 = 3/2$. For a succession
 373 of n applications of the bilinear factor, it is $\alpha_c = n/2 + 1$. The gain versus
 374 wavelength is plotted in Fig. 4 for the bilinear filter without compensation

375 ($G = g(1/2, k)$), with compensation ($G = g(1/2, k) \cdot g(3/2, k)$), and four n-
 376 pass bilinear filters with compensation ($G = g(1/2, k)^n \cdot g(3/2, k)$) for $n = 4,$
 20, 50 and 80.

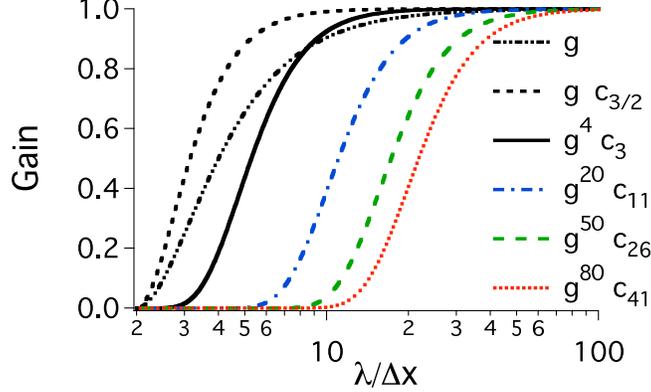


Figure 4: Gain versus wavelength for the bilinear filter without compensation ($g = g(1/2, k)$), with compensation ($g \cdot c_{3/2} = g(1/2, k) \cdot g(3/2, k)$), and n-pass bilinear filters with compensation ($g^n \cdot c_{\alpha_c} = g(1/2, k)^n \cdot g(\alpha_c, k)$) for $n = 4, 20, 50$ and 80.

377

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

$$\phi_j^f = \alpha \phi_j + (1 - \alpha) \frac{\phi_{j-s} + \phi_{j+s}}{2} \quad (41)$$

382 for which the gain is given by

$$g(s, \alpha, k) = \alpha + (1 - \alpha) \cos(sk\delta x) \quad (42)$$

$$\approx 1 - (1 - \alpha) \frac{(sk\delta x)^2}{2} + O(k^4) \quad (43)$$

383 The gain is plotted in Fig. 5 (top) for four passes bilinear filters with
 384 compensation ($G = g(s, 1/2, k)^4 \cdot g(s, 3/2, k)$) for strides $s=1$ to 4. For a

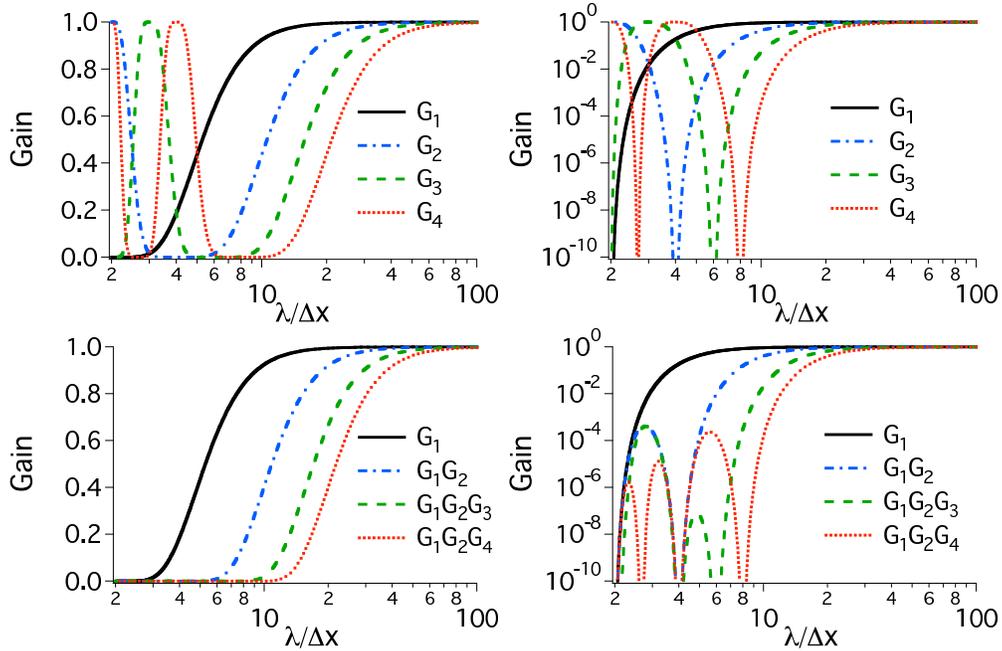


Figure 5: (top) gain for four passes bilinear filters with compensation ($G_s = g(s, 1/2, k)^4 \cdot g(s, 3/2, k)$) for strides $s=1$ to 4 linear with (left) linear ordinate (right) logarithmic ordinate; (bottom) gain for four low pass filters combining the G_1 to G_4 filters with (left) linear ordinate (right) logarithmic ordinate.

385 given stride, the gain is given by the gain of the bilinear filter shifted in
 386 k space, with the pole $g = 0$ shifted from $\lambda = 2/\delta x$ to $\lambda = 2s/\delta x$, with
 387 additional poles, as given by

$$sk\delta x = \arccos\left(\frac{\alpha}{\alpha - 1}\right) \pmod{2\pi} \quad (44)$$

388 The resulting filter is pass band between the poles, but since the poles are
 389 spread at different integer values in k space, a wide band low pass filter can
 390 be constructed by combining filters at different strides. Examples are given
 391 in Fig. 5 (bottom) for combinations of the filter with stride 1 to 4.

392 The combined filters with strides 2, 3 and 4 have nearly equivalent fall-

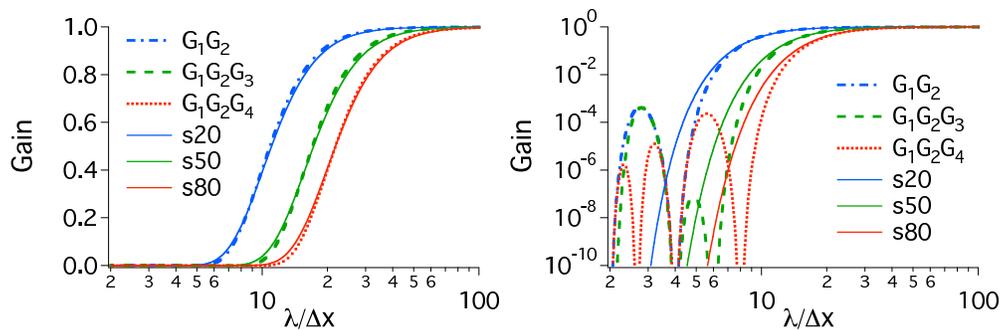


Figure 6: Comparison between filters with stride and filter s20-80 with (left) linear ordinate (right) logarithmic ordinate.

393 offs in gain (in linear scale) to the 20, 50 and 80 passes of the bilinear filter
 394 (see Fig. 6). Yet, the filters with stride need respectively 10, 15 and 15
 395 passes of a three-point filter while the n-pass bilinear filter need respectively
 396 21, 51 and 81 passes, giving gains of respectively 2.1, 3.4 and 5.4 in number
 397 of operations in favor of the filters with stride.

398 *4.2. Tunable solver*

399 In [38] and [39], Cole introduced an implementation of the source-free
400 Maxwell's wave equations for narrow-band applications based on non-standard
401 finite-differences (NSFD). In [40], Karkkainen et al adapted it for wideband
402 applications. At the Courant limit for the time step and for a given set of
403 parameters, the stencil proposed in [40] has no numerical dispersion along the
404 principal axes, provided that the cell size is the same along each dimension
405 (i.e. cubic cells in 3D). The solver from [40] was modified to be consistent
406 with the Particle-In-Cell methodology and implemented in the code Warp,
407 with the ability given to the user of setting the solver adjustable coefficients,
408 providing tunability of the numerical properties of the solver to better fit the
409 requirements of a particular application.

410 The "Cole-Karkkainen" 's solver [40] uses a non-standard finite difference
411 formulation (extended stencil) of the Maxwell-Ampere equation. For imple-
412 mentation into a Particle-In-Cell code, the formulation must introduce the
413 source term into Cole-Karkkainen's source free formulation in a consistent
414 manner. However, modifying the NSFD formulation of the Maxwell-Ampere
415 equation so that it includes the source term in a way that is consistent with
416 the current deposition scheme is challenging. To circumvent this problem,
417 Warp implementation departs from Karkkainen's by applying the enlarged
418 stencil on the Maxwell-Faraday equations, which does not contain any source
419 term but is formally equivalent to the source-free Maxwell-Ampere equation.
420 Consequently, in Warp's implementation, the discretized Maxwell-Ampere
421 equation is the same as in the Yee scheme, and the discretized Maxwell's

422 equations read:

$$\Delta_t \mathbf{B} = -\nabla^* \times \mathbf{E} \quad (45)$$

$$\Delta_t \mathbf{E} = c^2 \nabla \times \mathbf{B} - \frac{\mathbf{J}}{\epsilon_0} \quad (46)$$

$$\nabla \cdot \mathbf{E} = \frac{\rho}{\epsilon_0} \quad (47)$$

$$\nabla^* \cdot \mathbf{B} = 0 \quad (48)$$

423 where ϵ_0 is the permittivity of vacuum, and Eq. 47 and 48 not being solved
 424 explicitly but verified via appropriate initial conditions and current deposi-
 425 tion procedure. The differential operators are defined as

$$\nabla = \Delta_x \hat{\mathbf{x}} + \Delta_y \hat{\mathbf{y}} + \Delta_z \hat{\mathbf{z}} \quad (49)$$

$$\nabla^* = \Delta_x^* \hat{\mathbf{x}} + \Delta_y^* \hat{\mathbf{y}} + \Delta_z^* \hat{\mathbf{z}}, \quad (50)$$

426 the finite differences and sums operators being

$$\Delta_t G|_{i,j,k}^n = \frac{G|_{i,j,k}^{n+1/2} - G|_{i,j,k}^{n-1/2}}{\delta t} \quad (51)$$

$$\Delta_x G|_{i,j,k}^n = \frac{G|_{i+1/2,j,k}^n - G|_{i-1/2,j,k}^n}{\delta x} \quad (52)$$

$$\Delta_x^* = (\alpha + \beta S_x^1 + \gamma S_x^2) \Delta_x \quad (53)$$

427 with

$$\begin{aligned} S_x^1 G|_{i,j,k}^n &= G|_{i,j+1/2,k}^n + G|_{i,j-1/2,k}^n \\ &+ G|_{i,j,k+1/2}^n + G|_{i,j,k-1/2}^n \end{aligned} \quad (54)$$

$$\begin{aligned} S_x^2 G|_{i,j,k}^n &= G|_{i,j+1/2,k+1/2}^n + G|_{i,j-1/2,k+1/2}^n \\ &+ G|_{i,j+1/2,k-1/2}^n + G|_{i,j-1/2,k-1/2}^n \end{aligned} \quad (55)$$

428 The quantity G is a sample vector component, δt and δx are respectively
 429 the time step and the grid cell size along x , while α , β and γ are constant

430 scalars verifying $\alpha + 4\beta + 4\gamma = 1$. The operators along y and z , i.e. Δ_y ,
 431 Δ_z , Δ_y^* , Δ_z^* , S_y^1 , S_z^1 , S_y^2 , and S_z^2 , are obtained by circular permutation of the
 432 indices.

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

$$\Delta_x^* = (\alpha + \beta S_x^1) \Delta_x \quad (56)$$

$$S_x^1 G_{i,j,k}^n = G_{i,j+1/2,k}^n + G_{i,j-1/2,k}^n. \quad (57)$$

434 An extension of this algorithm for non-cubic cells provided by Cowan in
 435 [42] is not considered in this paper. However, all considerations given here
 436 for the solver implemented in Warp apply readily to the solver developed by
 437 Cowan.

438

439 4.2.1. Numerical dispersion

440 The dispersion relation of the solver is given by

$$\left(\frac{\sin \frac{\omega \delta t}{2}}{c \delta t} \right)^2 = C_x \left(\frac{\sin \frac{k_x \delta x}{2}}{\delta x} \right)^2 + C_y \left(\frac{\sin \frac{k_y \delta y}{2}}{\delta y} \right)^2 + C_z \left(\frac{\sin \frac{k_z \delta z}{2}}{\delta z} \right)^2 \quad (58)$$

441 with

$$C_x = \alpha + 2\beta(c_y + c_z) + 4\gamma c_y c_z \quad (59)$$

$$C_y = \alpha + 2\beta(c_z + c_x) + 4\gamma c_z c_x \quad (60)$$

$$C_z = \alpha + 2\beta(c_x + c_y) + 4\gamma c_x c_y \quad (61)$$

442 and

$$c_x = \cos(k_x \delta x) \quad (62)$$

$$c_y = \cos(k_y \delta y) \quad (63)$$

$$c_z = \cos(k_z \delta z) \quad (64)$$

443 The Courant-Friedrichs-Lewy condition (CFL) is given by

$$\begin{aligned}
c\delta t_c \leq & \min[\delta x, \delta y, \delta z, \\
& 1/\sqrt{(\alpha - 4\gamma) \max[\kappa_x + \kappa_y, \kappa_x + \kappa_z, \kappa_y + \kappa_z]}, \\
& 1/\sqrt{(\alpha - 4\beta + 4\gamma) (\kappa_x + \kappa_y + \kappa_z)}] \tag{65}
\end{aligned}$$

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

445 Assuming cubic cells ($\delta x = \delta y = \delta z$), the coefficients given in [40] ($\alpha =$
446 $7/12$, $\beta = 1/12$ and $\gamma = 1/48$) allow $c\delta t = \delta x$, and thus no dispersion along
447 the principal axes.

448 It is of interest to note that (58) can be rewritten

$$\left(\frac{\sin \frac{\omega\delta t}{2}}{c\delta t} \right)^2 = (s_x^2 + s_y^2 + s_z^2) + \beta' (s_x^2 s_y^2 + s_x^2 s_z^2 + s_y^2 s_z^2) + \gamma' (s_x^2 s_y^2 s_z^2) \tag{66}$$

449 with $s_x = \sin(k_x \delta x/2)$, $s_y = \sin(k_y \delta y/2)$, $s_z = \sin(k_z \delta z/2)$, $\beta' = -8\beta - 16\gamma$
450 and $\gamma' = 48\gamma$, for which the coefficients from [40] take the nice values $\beta' = -1$
451 and $\gamma' = 1$.

452 Sets of possible coefficients and the corresponding CFL condition, assum-
453 ing cubic cells, are given in Table 1. The numerical dispersion using those
454 coefficients are plotted in figure 7 along the principal axes and diagonals for
455 cubic cells ($\delta x = \delta y = \delta z$) and contrasted with the one of the Yee solver (all
456 taken at each solver's CFL time step limit). At the CFL limit, the Yee al-
457 gorithm offers no numerical dispersion along the 3D diagonal, but relatively
458 large numerical dispersion at the Nyquist frequency along the main axes.
459 Conversely, the Cole-Karkkainen solver (CK) offers no numerical dispersion
460 along the main axes but significant dispersion along the diagonals. The CK

	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
$c\delta t/\delta x$	$1/\sqrt{3}$	1	$1/\sqrt{2}$	$1/\sqrt{2}$	$\sqrt{2}/\sqrt{3}$	$\sqrt{5}/\sqrt{6}$

Table 1: List of coefficients

461 solver also allows larger time steps than the Yee solver by almost a factor
462 of two in 3D. The solver labeled "CK2" offers numerical dispersion that is
463 intermediate between the Yee solver and the CK solver along the main axes
464 and the 3D diagonal, but slightly degraded along the 2D diagonal. Con-
465 versely, while solver CK3 also offers intermediate numerical dispersion along
466 the main axes and the 3D diagonal, it offers no numerical dispersion along the
467 2D diagonal. Solver CK4 improves slightly the numerical dispersion along
468 the main axes over CK2 and CK3 at the expense of the dispersion along the
469 diagonals. Finally, CK5 offers the highest level of isotropy. The CFL time
470 steps of solvers CK2, 3, 4 and 5 are intermediate between the Yee and the CK
471 CFL time steps. This provides solvers with a range of numerical dispersion
472 among which some may be more favorable with regard to the mitigation of
473 numerical instabilities for a given application.

474 To reduce numerical dispersion to its lowest level, it is desirable to operate
475 the CK solver as close as possible to the CFL limit $c\delta t = \delta x$. However, an
476 instability (other than numerical Cerenkov) arises at the Nyquist frequency

477 in such a case. The analysis is given in 1D in Appendix I, as well as its
478 mitigation using digital filtering. Since for the CK solver, the CFL limit is
479 independent of dimensionality, the analysis and mitigation apply readily to
480 2D and 3D simulations.

481 For absorption of outgoing waves at the computational box boundaries,
482 the extension of the solver to a Perfectly Matched Layer [43] is given in
483 Appendix II.

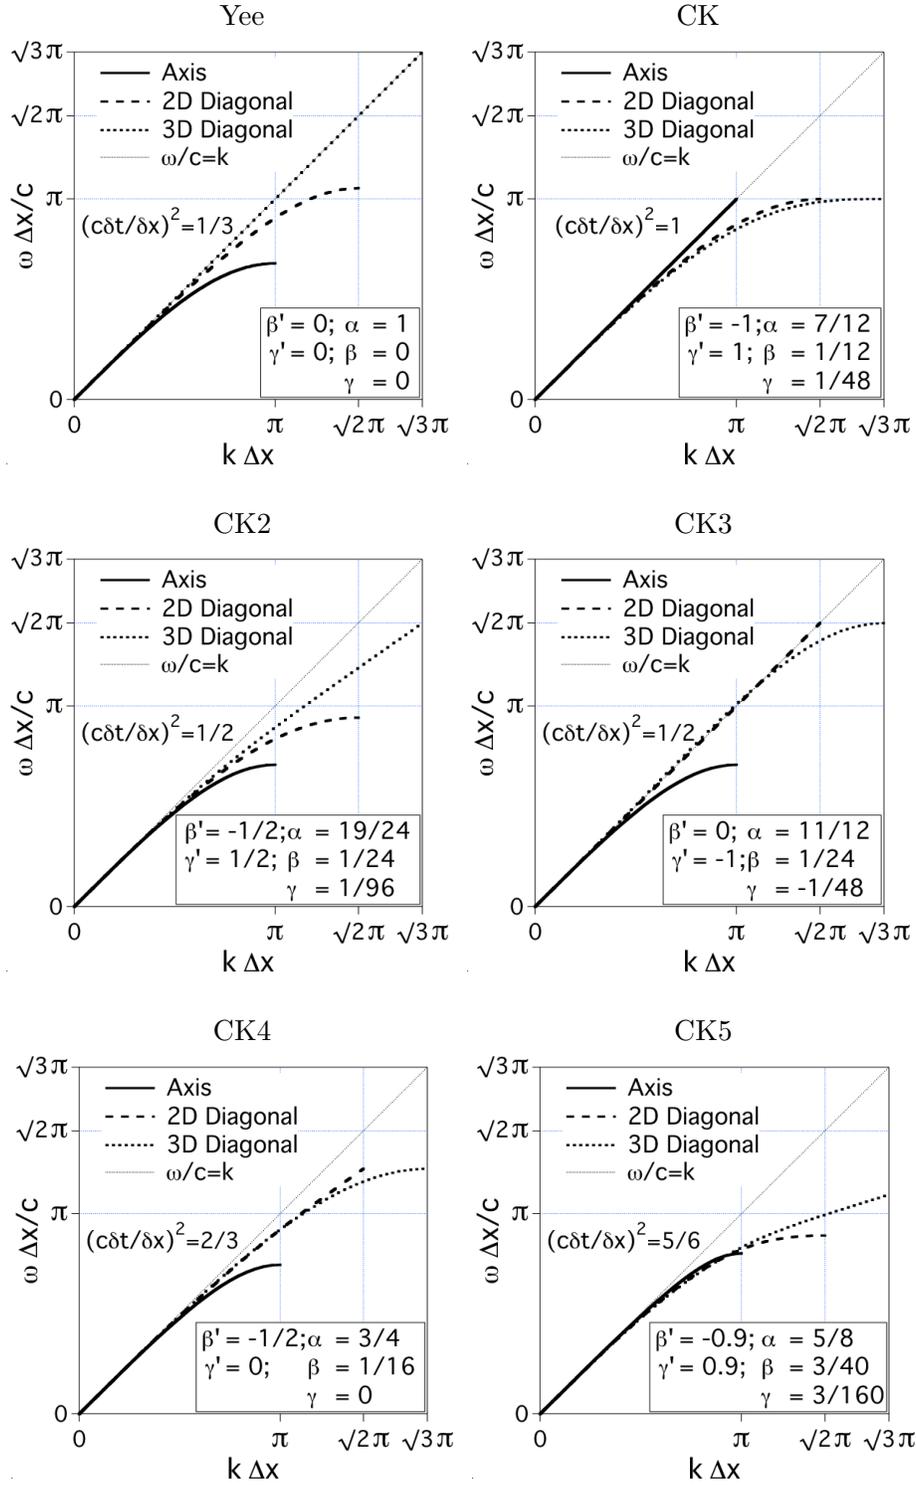


Figure 7: Numerical dispersion along the principal axis and diagonals for cubic cells ($\delta x = \delta y = \delta z$) at the Courant limit for the solver with adjustable numerical dispersion using the parameters from Table 1.

484 *4.2.2. Current deposition and Gauss' Law*

485 In most applications, it is essential to prevent accumulations of errors to
 486 the discretized Gauss' Law. This is accomplished by providing a method
 487 for depositing the current from the particles to the grid which is compatible
 488 with the discretized Gauss' Law, or by providing a mechanism for "divergence
 489 cleaning" [46, 47, 48, 49]. For the former, schemes which allow a deposition
 490 of the current that is exact when combined with the Yee solver is given in
 491 [50] for linear form factors and in [51] for higher order form factors. Since the
 492 discretized Gauss' Law and Maxwell-Faraday equation are the same in our
 493 implementation as in the Yee solver, charge conservation is readily verified
 494 using the current deposition procedures from [50] and [51], and this was
 495 verified numerically. Hence divergence cleaning is not necessary.

496 *4.3. Friedman adjustable damping*

497 The tunable damping scheme developed by Friedman [35] was shown to
 498 be the most potent practical method for mitigating the numerical Cerenkov
 499 instability in [33], among the selected methods that were considered. It is
 500 readily applicable to the solver presented above by modifying (45) to

$$\mathbf{B}^{n+3/2} = \mathbf{B}^{n+1/2} - \delta t \nabla^* \times \left[\left(1 + \frac{\theta}{4}\right) \mathbf{E}^{n+1} - \frac{1}{2} \mathbf{E}^n + \left(\frac{1}{2} - \frac{\theta}{4}\right) \bar{\mathbf{E}}^{n-1} \right] \quad (67)$$

501 with

$$\bar{\mathbf{E}}^{n-1} = \left(1 - \frac{\theta}{2}\right) \mathbf{E}^n + \frac{\theta}{2} \bar{\mathbf{E}}^{n-2} \quad (68)$$

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

$$\left(\frac{\sin \frac{\omega \delta t}{2}}{c \delta t} \right)^2 = F \Omega^2 \quad (69)$$

503 where

$$F = 1 - \frac{2\theta \sin^2(\omega\delta t/2)}{2e^{-i\omega\delta t} - \theta} \quad (70)$$

504 and

$$\Omega^2 = \left[C_x \left(\frac{\sin \frac{k_x \delta x}{2}}{\delta x} \right)^2 + C_y \left(\frac{\sin \frac{k_y \delta y}{2}}{\delta y} \right)^2 + C_z \left(\frac{\sin \frac{k_z \delta z}{2}}{\delta z} \right)^2 \right] \quad (71)$$

505 The CFL is given by

$$c\delta t_c^* = c\delta t_c \sqrt{\frac{2 + \theta}{2 + 3\theta}} \quad (72)$$

506 where δt_c is the critical time step of the numerical scheme without damping
507 ($\theta = 0$), as given by (65).

508 The numerical dispersion of the Cole-Karkkainen-Friedman (CKF) solver
509 (using the coefficients from the CK solver in Table 1) is plotted in figure 8
510 along the principal axis and diagonals for cubic cells ($\delta x = \delta y = \delta z$) and
511 contrasted with the one of the Yee-Friedman (YF) solver (both taken at the
512 Courant time step limit). The amount of phase error rises with the value
513 of the damping parameter θ (partly due to the slightly more constraining
514 limit on the critical time step). However, it was shown in [33] that the
515 amount of damping provided by the YF solver was sufficient to counteract
516 the slight degradation of numerical dispersion with raising θ , reducing the
517 numerical Cerenkov effects to an acceptable level for the problem that was
518 considered. The damping is very isotropic with the CKF solver but not with
519 the YF one. The YF implementation offers a higher level of damping of
520 the shortest wavelengths along the 3D diagonals, while the CKF offers more
521 damping along the axes, and the amount of damping along the 2D diagonals
522 are similar. In summary, the YF implementation delivers respectively the

523 highest/lowest level of damping in the direction of lowest/highest numerical
524 dispersion, while the CKF implementation delivers a proportionally higher
525 level of dispersion than the YF implementation along the direction of highest
526 numerical dispersion. Thus it may be expected that the CKF implementation
527 will be more efficient in reducing numerical Cerenkov effects.

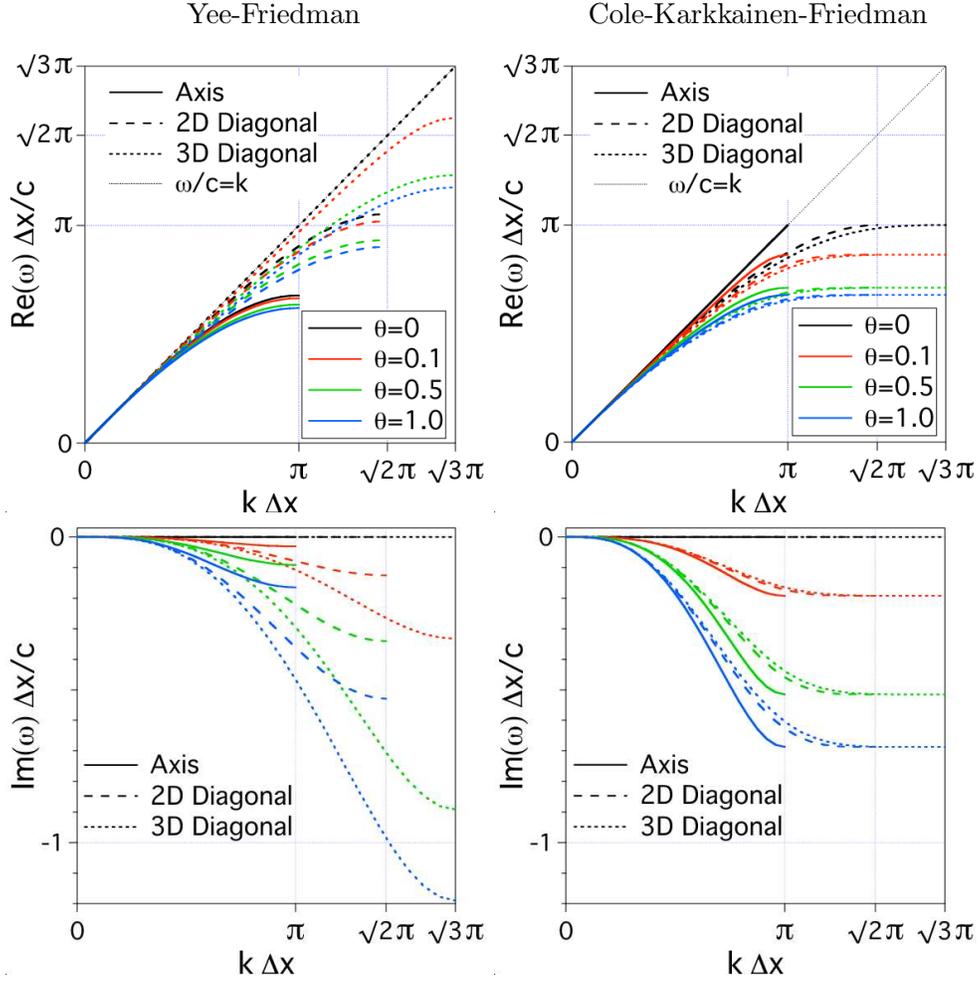


Figure 8: Numerical dispersion along the principal axis and diagonals for cubic cells ($\delta x = \delta y = \delta z$) at the Courant limit for: (left) the Yee-Friedman solver; (right) the Cole-Karkkainen-Friedman solver. The real part (phase) and the imaginary part (amplitude) are plotted respectively in the top and bottom rows.

528 5. Application to the modeling of laser wakefield acceleration

529 This section presents applications of the methods to the modeling of 10
530 GeV LPA stages at full scale in 2-1/2D and 3D, which has not been done
531 fully self-consistently with other methods. It has been shown that many
532 parameters of high energy LPA stages can be accurately simulated at reduced
533 cost by simulating stages of lower energy gain, with higher density and shorter
534 acceleration distance, by scaling the physical quantities relative to the plasma
535 wavelength, and this has been applied to design of 10 GeV LPA stages [18,
536 19]. The number of oscillations of a mismatched laser pulse in the plasma
537 channel however depends on stage energy and does not scale, though this
538 effect is minimized for a channel guided stage as considered in [18, 19]. The
539 number of betatron oscillations of the trapped electron bunch will also depend
540 on the stage energy, and may affect quantities like the emittance of the beam.
541 For these reasons, and to prove validity of scaled designs of other parameters,
542 it is necessary to perform full scale simulations, which is only possible by
543 using reduced models or simulations in the boosted frame.

544 As a benchmarking exercise, we first perform scaled simulations similar
545 to the ones performed in [18], at a density of $n_e = 10^{19} \text{ cm}^{-3}$, using various
546 values of the boosted frame relativistic factor γ to show the accuracy and con-
547 vergence of the technique. These stages were shown to efficiently accelerate
548 both electrons and positrons with low energy spread, and the scaled simula-
549 tions predicted acceleration of hundreds of pC to 10 GeV energies using a 40
550 J laser. The accuracy of the technique is evaluated by modeling scaled stages
551 [18, 19] at 0.1 GeV, which allows for a detailed comparison of simulations
552 using a reference frame ranging from the laboratory frame to the frame of the

553 wake. Excellent agreement is obtained on wakefield histories on axis, beam
554 average energy history and momentum spread at peak energy, with speedup
555 over a hundred, in agreement with the theoretical estimates from Section
556 2. The downscaled simulations are also used for an in-depth exploration of
557 the effects of the methods presented in Sections 3 and 4, and evaluation of
558 which techniques are required to permit maximum γ boost while maintaining
559 high accuracy. We then apply the boosted frame technique to provide full
560 scale simulation of high efficiency quasilinear LPA stages at higher energy,
561 verifying the scaling laws in the 10 GeV-1 TeV range.

562 *5.1. Scaled 10 GeV stages*

563 The parameters were chosen to be close (though not identical) to the
564 case where $k_p L = 2$ in [18] where k_p is the plasma wavenumber and L is
565 the laser pulse length. In the cases considered in this paper, the beam is
566 composed of test particles only, with the goal of testing the fidelity of the
567 algorithm in modeling laser propagation and wake generation. The results
568 from simulations of LPA in a boosted frame where beam loading is present
569 will be presented elsewhere. These simulations are scaled replicas of 10 GeV
570 stages that would operate at actual densities of 10^{17} cm⁻³ [18, 19] and allow
571 short run times to permit effective benchmarking between the algorithms.
572 The main physical and numerical parameters of the simulation are given in
573 Table 2. Unless noted otherwise, in all the simulations presented herein, the
574 field is gathered from the grid onto the particles directly from the Yee mesh
575 locations, i.e. using the 'energy conserving' procedure (see [46], chapter 10).

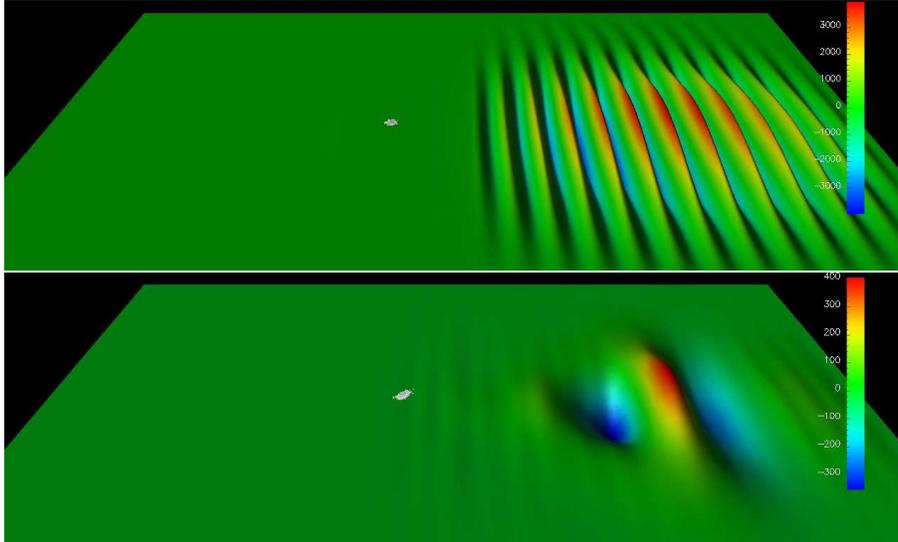


Figure 9: Colored surface rendering of the transverse electric field from a 2-1/2D Warp simulation of a laser wakefield acceleration stage in the laboratory frame (top) and a boosted frame at $\gamma = 13$ (bottom), with the beam (white) in its early phase of acceleration. The laser and the beam are propagating from left to right.

576 *5.1.1. Using standard numerical techniques*

577 These runs were done using the standard Yee solver with no damping, and
 578 with the 4-pass stride-1 filter plus compensation, similarly to the simulations
 579 reported in [18]. No signs of detrimental numerical instabilities were observed
 580 at the resolutions reported here with these settings.

581 The approximate relativistic factor of the wake that is formed by the
 582 laser traveling in the plasma is given, according to linear theory, by $\gamma_w =$
 583 $2\pi c/\lambda\omega_p$ where $\omega_p = \sqrt{n_e e^2/\epsilon_0 m_e}$ is the electron plasma frequency. For the
 584 given parameters, $\gamma_w \approx 13.2$. Thus, Warp simulations were performed using
 585 reference frames moving between $\gamma = 1$ (laboratory frame) and 13. For a
 586 boosted frame associated with a value of γ approaching γ_w in the laboratory,

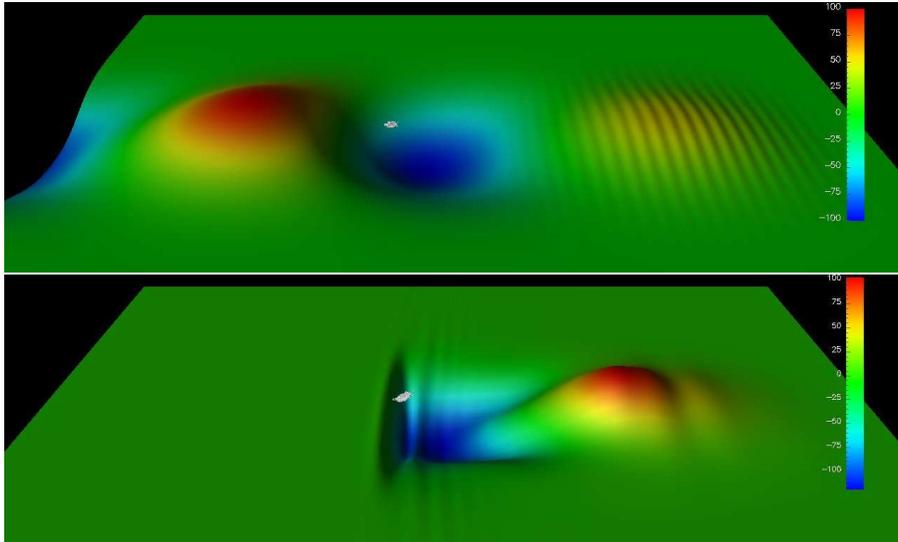


Figure 10: Colored surface rendering of the longitudinal electric field from a 2-1/2D Warp simulation of a laser wakefield acceleration stage in the laboratory frame (top) and a boosted frame at $\gamma = 13$ (bottom), with the beam (white) in its early phase of acceleration. The laser and the beam are propagating from left to right.

587 the wake is expected to travel at low velocity in this boosted frame, and the
 588 physics to appear somewhat different from the one observed in the laboratory
 589 frame, in accordance to the properties of the Lorentz transformation. Figure
 590 9 and 10 show surface renderings of the transverse and longitudinal electric
 591 fields respectively, as the beam enters its early stage of acceleration by the
 592 plasma wake, from a calculation in the laboratory frame and another in the
 593 frame at $\gamma = 13$. The two snapshots offer strikingly different views of the
 594 same physical processes: in the laboratory frame, the wake is fully formed
 595 before the beam undergoes any significant acceleration and the imprint of
 596 the laser is clearly visible ahead of the wake; while in the boosted frame
 597 calculation, the beam is accelerated as the plasma wake develops, and the

598 laser imprint is not visible on the snapshot. Close examination reveals that
 599 the short spatial variations which make the laser imprint in front of the wake
 600 are transformed into time variations in the boosted frame of $\gamma = 13$.

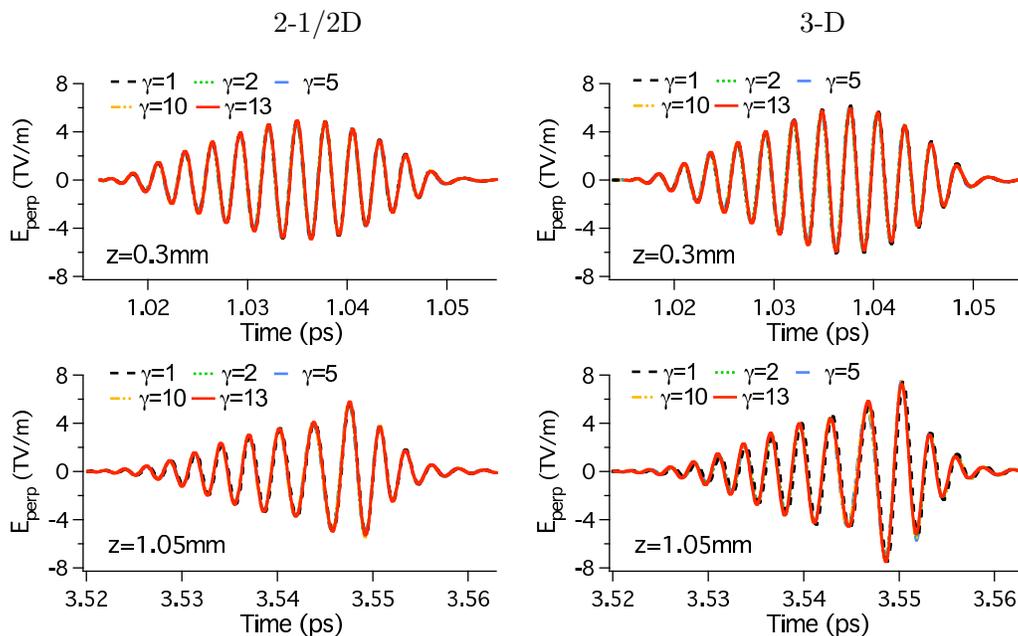


Figure 11: History of transverse electric field at the position $x = y = 0$, $z = 0.3$ mm and $z = 1.05$ mm (in the laboratory frame) from simulations in the laboratory frame ($\gamma = 1$) and boosted frames at $\gamma = 2, 5, 10$ and 13 .

601 Histories of the perpendicular and longitudinal electric fields recorded at a
 602 number of stations at fixed locations in the laboratory offer direct comparison
 603 between the simulations in the laboratory frame ($\gamma = 1$) and boosted frames
 604 at $\gamma = 2, 5, 10$ and 13 . Figure 11 and 12 show respectively the transverse
 605 and longitudinal electric fields collected at the positions $z = 0.3$ mm and
 606 $z = 1.05$ mm (in the laboratory frame) on axis ($x = y = 0$). The agreement
 607 is excellent and confirms that despite the apparent differences from snap-

608 shots taken from simulations in different reference frames, the same physics
609 was recovered. This is further confirmed by the plot of the average scaled
610 beam energy gain as a function of position in the laboratory frame, and of
611 relative longitudinal momentum dispersion at peak energy (Fig. 13). The
612 small differences observed on the mean beam energy histories and on the lon-
613 gitudinal momentum spread are attributed to a lack of convergence at the
614 resolution that was chosen. The beam was launched with the same phase in
615 the 2-1/2D and the 3D simulations, resulting in lower energy gain in 3D, due
616 to proportionally larger laser depletion effects in 3D than in 2-1/2D.

617 The CPU time recorded as a function of the average beam position in
618 the laboratory frame (Fig. 13-middle) indicates that the simulation in the
619 frame of $\gamma = 13$ took ≈ 25 s in 2-1/2D and ≈ 150 s in 3D versus $\approx 5,000$
620 s in 2-1/2D and $\approx 20,000$ s in 3D in the laboratory frame, demonstrating
621 speedups of ≈ 200 in 2-1/2D and ≈ 130 in 3D, between calculations in a
622 boosted frame at $\gamma = 13$ and the laboratory frame.

623 All the simulations presented so far in this section were using the Yee
624 solver, for which the Courant condition is given by $c\delta t < (1/\delta x^2 + 1/\delta z^2)^{-1/2}$
625 in 2D and $c\delta t < (1/\delta x^2 + 1/\delta y^2 + 1/\delta z^2)^{-1/2}$ in 3D where δt is the time step
626 and δx , δy and δz are the computational grid cell sizes in x , y and z . As γ
627 was varied, the transverse resolution was kept constant, while the longitudinal
628 resolution was kept at a constant fraction of the incident laser wavelength
629 $\delta z = \zeta\lambda$, such that in a boosted frame, $\delta z^* = \zeta\lambda^* = \zeta(1 + \beta)\gamma\lambda$. As a result,
630 the speedup becomes, when using the Yee solver

$$S_{yee2D} = S \frac{\delta z \sqrt{1/\delta x^2 + 1/\delta z^2}}{\delta_z^* \sqrt{1/\delta x^2 + 1/\delta z^{*2}}} \quad (73)$$

631 in 2D and

$$S_{yee3D} = S \frac{\delta z \sqrt{1/\delta x^2 + 1/\delta y^2 + 1/\delta z^2}}{\delta_z^* \sqrt{1/\delta x^2 + 1/\delta y^2 + 1/\delta z^{*2}}} \quad (74)$$

632 in 3D where S is given by Eq. (13).

633 The speedup versus relativistic factor of the reference frame is plotted in
634 Fig. 14, from (13), (73) and (74), and contrasted with measured speedups
635 from 1D, 2-1/2D and 3D Warp simulations, confirming the scaling obtained
636 analytically.

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

beam radius	R_b	82.5 nm
beam length	L_b	85. nm
beam transverse profile		$\exp(-r^2/8R_b^2)$
beam longitudinal profile		$\exp(-z^2/2L_b^2)$
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^+

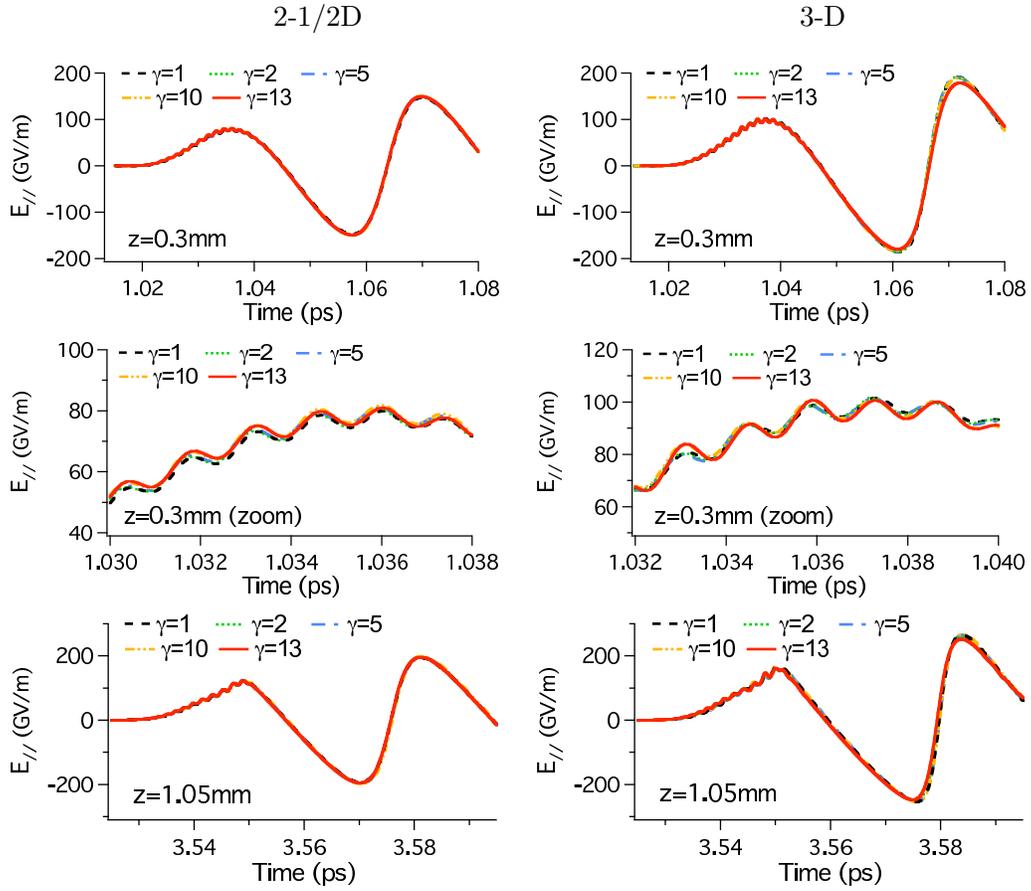


Figure 12: History of longitudinal electric field at the position $x = y = 0$, $z = 0.3$ mm and $z = 1.05$ mm (in the laboratory frame) from simulations in the laboratory frame ($\gamma = 1$) and boosted frames at $\gamma = 2, 5, 10$ and 13 .

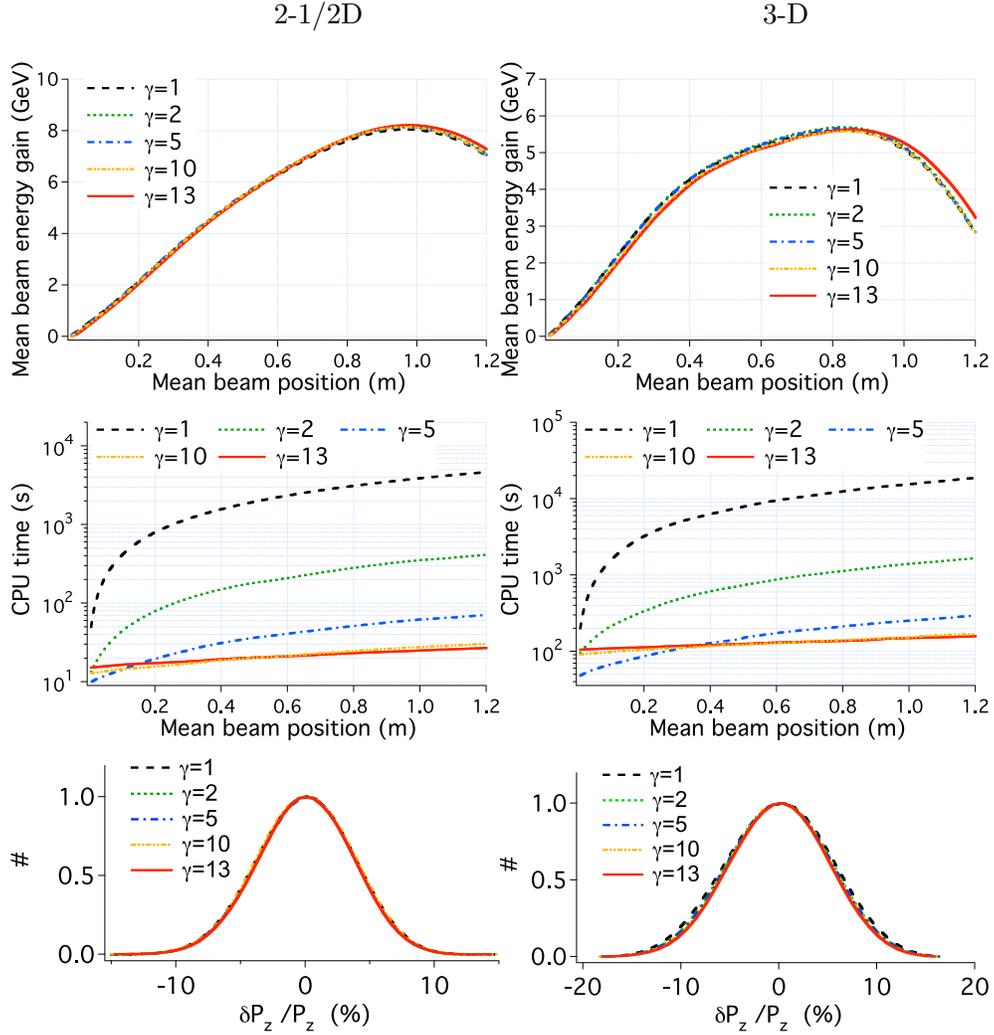


Figure 13: (top) Average scaled beam energy gain and (middle) CPU time, versus longitudinal position in the laboratory frame from simulations; (bottom) distribution of relative longitudinal momentum dispersion at peak energy, in the laboratory frame ($\gamma = 1$) and boosted frames at $\gamma = 2, 5, 10$ and 13 .

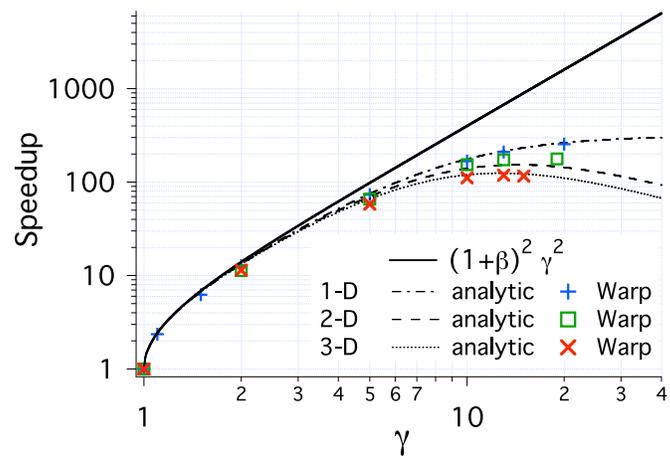


Figure 14: Speedup versus relativistic factor of the boosted frame from Eq. (13), (73), (74), and Warp simulations.

637 *5.1.2. Effect of filtering, solver with adjustable dispersion and damping*

638 The modeling of full scale stages, which allow for higher values of γ for
639 the reference frame, is more prone to the high frequency instability that was
640 mentioned in a previous section, as we will show below. In anticipation of
641 the application of the method presented above to mitigate the instability,
642 simulations of the scaled stage were conducted using the Yee solver with
643 digital filter S(1:2:4) as described above (Fig. 15), the Cole-Karkkainen solver
644 (Fig. 16) or the Yee-Friedman solver (Fig. 17).

645 Smoothing with the wideband filter S(1:2:4) did not produce significant
646 degradations for the calculation in the wake frame ($\gamma = 13$) but did otherwise.
647 The calculations with the Yee solver and the Cole-Karkkainen solver gave
648 identical results, validating our implementation of the CK solver. Despite the
649 more expensive stencil, the run with the CK solver was almost 40% faster,
650 due to a time step larger by $\sqrt{2}$. Similarly to filtering, damping aggressively
651 did not degrade the result in the range $10 \leq \gamma \leq 13$ but did significantly in
652 the range $1 \leq \gamma \leq 5$. Comparing the timings with those of Fig. 13 (middle-
653 left) shows that the smoothing and the damping added less than a factor of
654 two of total runtime to the simulations.

655 Those results lead to several observations: (i) while the grid dimensions
656 and number of cells were chosen such that square cells were obtained for
657 $\gamma = 13$, meaning a larger dispersion in the longitudinal direction with the Yee
658 solver than with the Cole-Karkkainen solver, both gave the same result. This
659 is significant since for simulations of LPA in the laboratory frame reported
660 in the literature, the need to have nearly perfect numerical dispersion in the
661 longitudinal direction imposes a constraint on the cell aspect ratio and thus

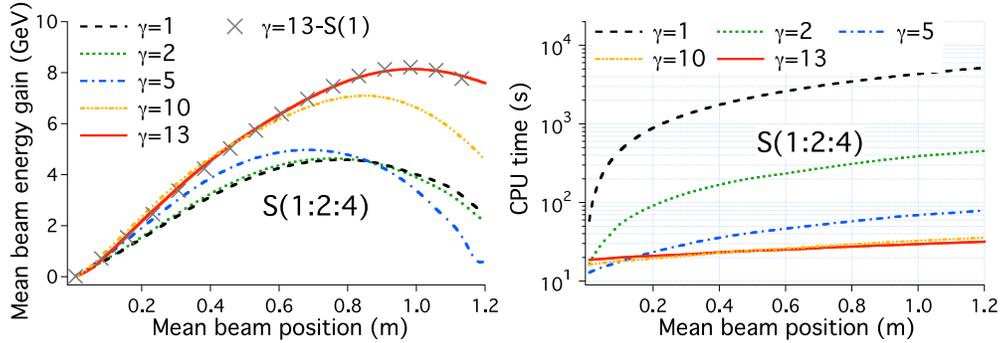


Figure 15: (left) Average scaled beam energy gain and (right) CPU time, versus longitudinal position in the laboratory frame from simulations in the laboratory frame ($\gamma = 1$) and boosted frames at $\gamma = 2, 5, 10$ and 13 , using the Yee solver with digital filter S(1:2:4) (grey cross is reference from run with filter S(1)).

662 on resolution [44, 45]. This constraint is removed when simulating in the
663 frame of the wake ($\gamma = 13 \approx \gamma_w$); (ii) damping of high frequencies with
664 the Yee-Friedman solver or wideband smoothing of short wavelength have a
665 negligible effect on accuracy for simulations in the frame of the wake, but
666 degrade the accuracy very significantly for slower moving reference frames.
667 The dependency of the effect of damping and smoothing with γ boost has
668 two causes. First, simulations with a boost $\gamma \approx \gamma_w$ require fewer time steps
669 than simulations using a lower value of γ . Thus, for a given value of the
670 damping coefficient θ , the integrated amount of damping will be lower for
671 the simulations with $\gamma \approx \gamma_w$. Second, as mentioned above in the discussion
672 of the surface renderings shown in Fig. 9 and 10, a large fraction of the short
673 wavelength content that is present in the simulations in the laboratory frame
674 is transformed into time oscillations in simulations in the wake frame. Hence,
675 filtering short wavelength has less effect on the physics when calculating in

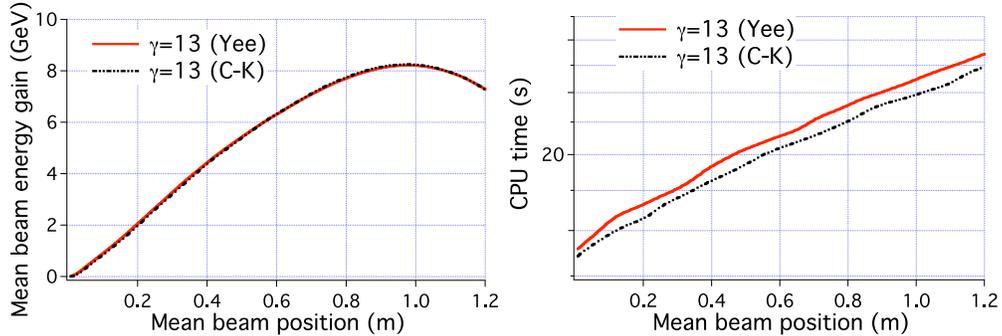


Figure 16: (left) Average scaled beam energy gain and (right) CPU time, versus longitudinal position in the laboratory frame from simulations in the laboratory frame ($\gamma = 1$) and boosted frames at $\gamma = 2, 5, 10$ and 13 , using the Cole-Karkkainen solver with filter $S(1)$ (red curve is reference from calculation with Yee solver and filter $S(1)$).

676 the wake frame than when calculating in the laboratory frame; (iii) the cost
 677 of using even the most aggressive damping or smoothing is low, especially
 678 considering that the simulations presented here were using only two plasma
 679 macro-particles per cell.

680 In summary, calculating in a boosted frame near the frame following the
 681 wake ($\gamma \approx \gamma_w$) relaxes the constraint on the numerical dispersion in the
 682 direction of propagation of the laser (which is essential in simulations in the
 683 laboratory frame), and allows for more aggressive damping of high frequencies
 684 and smoothing of short wavelengths than is possible in standard laboratory
 685 frame calculations.

686 *5.2. Full scale 10 GeV class stages*

687 As noted in [13], full scale simulations using the laboratory frame of 10
 688 GeV stages at plasma densities of 10^{17} cm^{-3} are not practical on present
 689 computers in 2D and 3D. At this density, the wake relativistic factor $\gamma_w \approx$

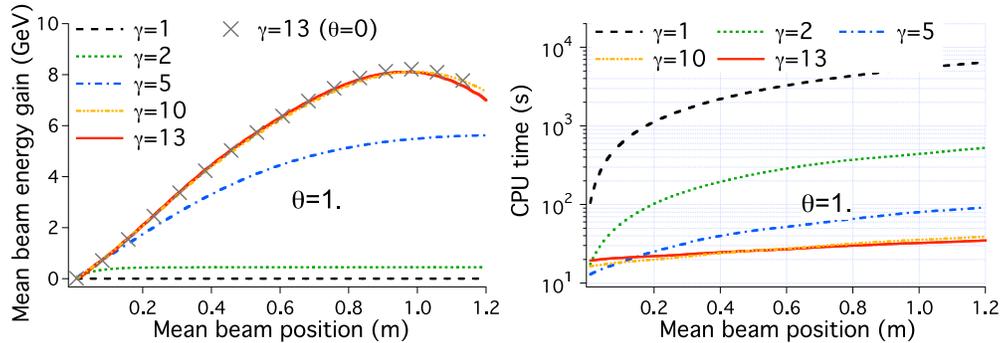


Figure 17: (left) Average scaled beam energy gain and (right) CPU time, versus longitudinal position in the laboratory frame from simulations in the laboratory frame ($\gamma = 1$) and boosted frames at $\gamma = 2, 5, 10$ and 13 , using the Yee-Friedman solver with $\theta = 1$ (grey cross is reference from run with no damping).

690 132, and 2-1/2D and 3D simulations were done in boosted frames up to
 691 $\gamma = 130$.

692 5.2.1. Simulations in 2-1/2D

693 Fig. 18 shows the average beam energy gain versus longitudinal position
 694 and the averaged Fourier Transform of the longitudinal electric field taken
 695 at $t=40$ ps, from 2D-1/2 simulations of a full scale 10GeV LPA in a boosted
 696 frame at $\gamma = 130$, using the Yee solver and various smoothing kernels. Fig.
 697 19 shows the average beam energy gain versus longitudinal position from
 698 simulations in boosted frames at $\gamma = 30, 60$ and 130 . All runs gave the
 699 same beam energy history within a few percents, and no sign of instability
 700 is detected in the Fourier transform plot of the longitudinal electric field.
 701 The average energy gain peaks around 8 GeV, in agreement with the scaled
 702 simulations (see Fig. 13).

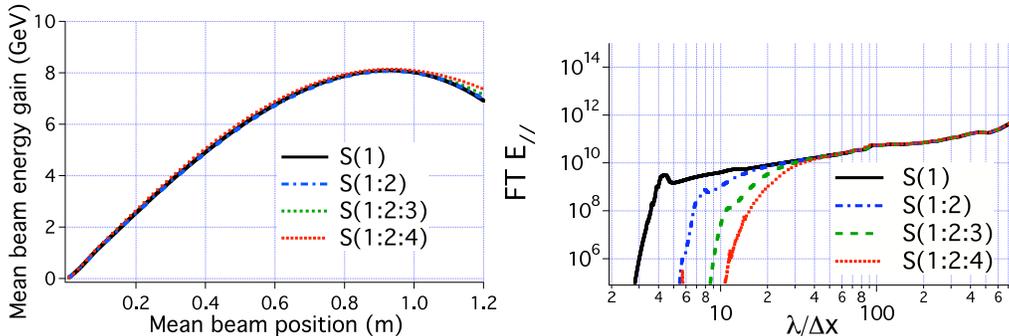


Figure 18: (left) Average beam energy gain versus longitudinal position (in the laboratory frame), (right) Fourier Transform of the longitudinal electric field at $t=40$ ps, averaged over whole domain, from 2D-1/2 simulations of a full scale 10GeV LPA in a boosted frame at $\gamma = 130$, using the Yee 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.

703 5.2.2. Simulations in 3D

704 In 3D, all simulations at $\gamma = 130$ using the Yee solver (using cubic
 705 cells and a time step at the CFL limit) developed the instability and loss
 706 of the beam, regardless of the amount of filtering or damping that has been
 707 tried. The failure of the 3D simulations using the Yee solver motivated use of
 708 the Cole-Karkkainen-Friedman (CKF) solver, with various levels of filtering
 709 and damping. Data from 3D simulations using the CKF solver and various
 710 smoothing kernels are plotted in Fig. 20. Stability is attained when using a
 711 sufficient level of filtering. Damping is detrimental to stability at low levels
 712 ($\theta = 0.1$) but is beneficial at a higher level ($\theta = 0.5$).

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

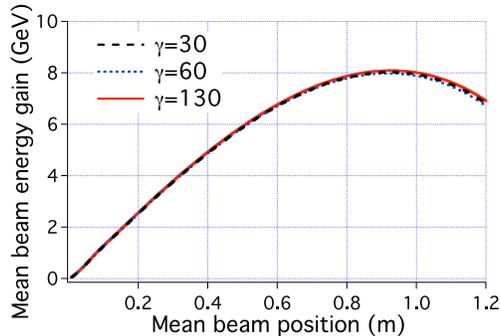


Figure 19: Average beam energy gain versus longitudinal position (in the laboratory frame) from 2D-1/2 simulations of a full scale 10GeV LPA in a boosted frame at $\gamma = 30, 60$ and 130, using the Yee solver.

717 runs using CK2 and CK3 were nearly identical and hence only those from
 718 CK2 are reported in Fig. 21, which show very consistent beam energy gain
 719 histories, and no sign of instability in the Fourier Transform plot of the
 720 longitudinal electric field at $t=40$ ps (closer inspection revealed that when
 721 using the lowest level of filtering S(1), a mild instability was developing but
 722 it was not affecting the average beam energy gain history). As shown on Fig.
 723 22, those results are in good agreement with runs at $\gamma = 30$ and 60 using the
 724 Yee solver, and with the maximum energy gain predicted around 5.7 GeV by
 725 the scaled simulations shown on Fig. 13 (top-right).

726 In summary, the full scale 6-7 GeV simulations using the frame of the
 727 wake performed in this subsection show: (i) 2-1/2D simulations using the
 728 Yee solver at the CFL limit (with square cells) were free of instability; (ii)
 729 3D simulations using the CK solver developed moderately strong instabilities
 730 that were mitigated using moderate to high levels of damping and/or filtering,
 731 the latter being the most effective; (iii) 3D simulations using the CK2 (or

732 CK3) solver developed very mild instabilities that were mitigated with a low
733 level of filtering.

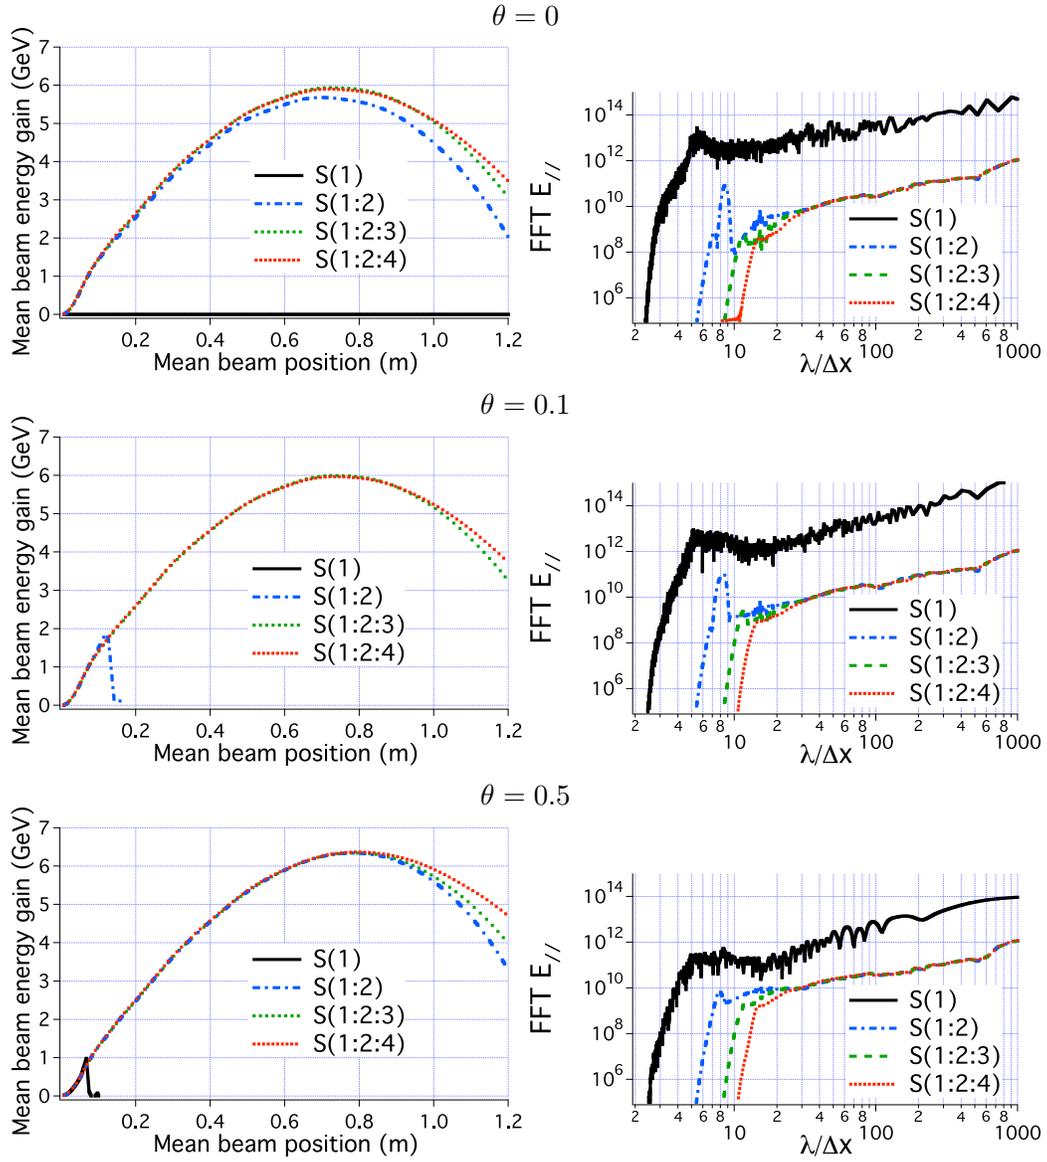


Figure 20: (left) Average beam energy gain versus longitudinal position (in the laboratory frame); (right) Fourier Transform of the longitudinal electric field at $t=40$ ps, averaged over plane on axis perpendicular to laser polarization, from 3D simulations of a full scale 10GeV LPA in a boosted frame 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) $\theta = 0.5$.

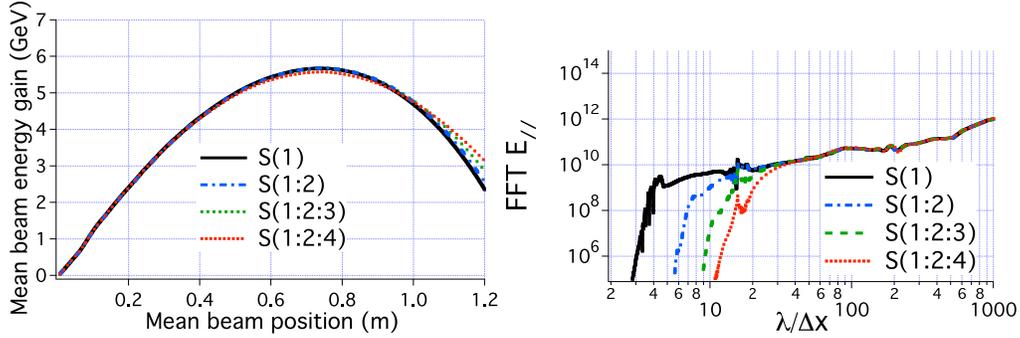


Figure 21: (left) Average beam energy gain versus longitudinal position (in the laboratory frame), (right) Fourier Transform of the longitudinal electric field at $t=40$ ps, averaged over whole domain, from 3D simulations of a full scale 10GeV LPA in a boosted frame at $\gamma = 130$, using the CK2 solver and various digital filter kernels.

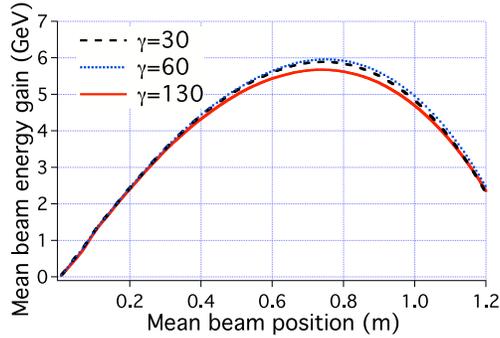


Figure 22: Average beam energy gain versus longitudinal position (in the laboratory frame) from 3D simulations of a full scale 10GeV LPA in a boosted frame at $\gamma = 30, 60$ and 130 , using the Yee solver ($\gamma = 30$ and 60) and the CK2 solver ($\gamma = 130$), with digital filter S(1) and with the time step set by $c\delta t/\delta z = 1/\sqrt{2}$ for stability (see discussion below) .

734 5.3. Effects of numerical parameters on the observed instability

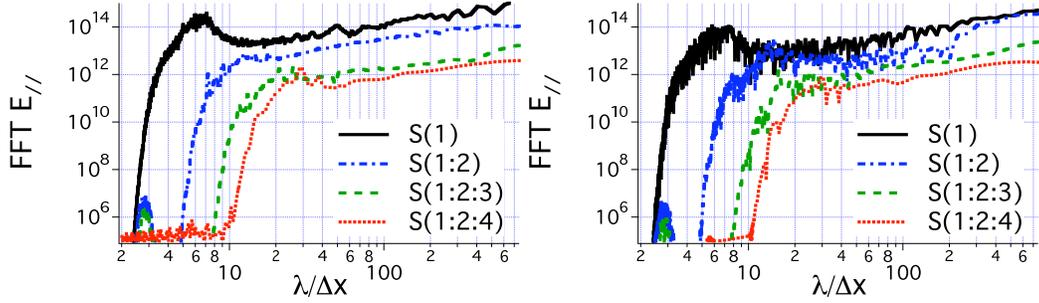


Figure 23: Fourier Transform of the longitudinal electric field at $t=40$ ps, averaged over plane on axis perpendicular to laser polarization, from (left) 3D and (right) 2D-1/2 simulations of a full scale 10GeV LPA in a boosted frame at $\gamma = 130$, using the Yee solver and various smoothing kernels. The same time step at the 3D CFL limit $c\delta t = \delta x/\sqrt{3}$ was used for both simulations.

735 The Fourier transform of the longitudinal electric field averaged over the
 736 whole domain at $t=40$ ps, from 3D simulations using the Yee solver, is given
 737 in Fig. 23 (left). It is contrasted to the same data taken from 2-1/2D
 738 simulations (right). Both simulations used the same time step at the 3D
 739 CFL limit $c\delta t = \delta z/\sqrt{3}$. The similarity of the two plots indicates that the
 740 degradation of the numerical dispersion that resulted from going from the
 741 2D to the 3D CFL limit is the cause of the failure of the 3D runs using the
 742 Yee solver. Taking advantage of this observation, we study in this section
 743 the instability arising in 2-1/2D simulations using a time step at the 3D CFL
 744 limit.

745 5.3.1. Effects of spatial resolution

746 Snapshots of the longitudinal electric field at the front of the plasma taken
 747 at $t = 12.5$ ps, and their corresponding Fourier transform, are given in Fig.

748 24, from 2-1/2D simulations using the Yee solver with the time step at the
749 3D CFL limit $c\delta t = \delta z/\sqrt{3}$. Three resolutions were considered: (a) $\delta_x =$
750 $\delta_z = 13\mu m$, (b) $\delta_x = \delta_z = 6.5\mu m$, and (c) $\delta_x = \delta_z = 3.25\mu m$. The amplitude
751 of the instability is roughly inversely proportional to the resolution. For this
752 configuration, the instability exhibits two primary modes at various relative
753 levels, both at a fixed number of grid cells in the longitudinal direction, but
754 at a fixed absolute length in the transverse direction. This indicates that
755 the transverse part of the modes is governed by the physical geometry of the
756 problem while the longitudinal part is governed by numerical resolution.

757 Results from 2-1/2D simulation using the CK solver at the 3D CFL limit
758 $c\delta t/\delta z = 1/\sqrt{3}$ at the resolution $\delta_x = \delta_z = 6.5\mu m$ are given in Fig. 25.
759 The same two modes that were observed in the plots from the equivalent
760 simulation using the Yee solver (see Fig. 24-middle), are present, and the
761 overall amplitude of the instability is similar. These similarities on the details
762 of the instability between the Yee and CK solvers indicate that the differences
763 in numerical dispersion of the solvers do not constitute a key factor affecting
764 the instability.

765 5.3.2. *Effects of time step*

766 It is striking that all the solvers that lead to the lowest levels of instability
767 had the same CFL time step $c\delta t_{CFL} = \delta z/\sqrt{2}$. For checking whether this is
768 coincidental, simulations were performed using the CK solver, scanning the
769 time step between $c\delta t/\delta z = 0.5$ and $c\delta t/\delta z = 1$. The Fourier Transform of
770 the longitudinal field averaged over the entire domain taken at $t = 40$ ps,
771 is given in Fig. 26, exhibiting a sharp reduction of the instability level in a
772 narrow band around $c\delta t = \delta z/\sqrt{2}$. Since the numerical dispersion degrades

773 in all directions when the time step diminishes, this indicates that the value
774 of the time step value is of more importance than the numerical dispersion
775 of the solver being used.

776 Simulations using the Yee or the CK solver with the singular time step
777 $c\delta t = \delta z/\sqrt{2}$ were performed and produced levels of instabilities that were
778 much reduced (and delayed) compared to the 3D CFL time step (not shown).
779 The snapshot of the electric field and its Fourier Transform taken at $t = 49$
780 ps are given in Fig. 27. The Fourier spectrum is very similar in each case,
781 although the instability is slightly stronger with the CK solver than with the
782 Yee solver. In both cases, the instability is easily removed by using the S(1:2)
783 filter (see Fig. 28).

784 As mentioned in the previous section, the solvers CK, CK4 and CK5,
785 which all have a CFL time step above the singular time step $c\delta t = \delta z/\sqrt{2}$,
786 produced significant levels of instability when run at their CFL limit. It was
787 verified that using those solvers in 3D at the time step $c\delta t = \delta z/\sqrt{2}$ resulted
788 in greatly reduced levels of instability. It was also observed that running
789 the Yee solver using non-cubic cells, e.g. with lower resolution transversely
790 such as $\delta x = 2\delta z$ at $\gamma = 130$, or $\delta x = 2.6\delta z$ at $\gamma = 50$, produced the same
791 pattern: a significant instability was present when using the CFL time step
792 and was greatly reduced by using $c\delta t = \delta z/\sqrt{2}$. Hence for the suppression of
793 the instability, the choice of the solver seems to depend solely on whether
794 its CFL condition allows stability at the special time step $c\delta t = \delta z/\sqrt{2}$ for a
795 given grid cell aspect ratio, but not significantly on its numerical dispersion
796 nor on the value of the grid cell aspect ratio.

797 *5.3.3. Effects of field gathering procedure*

798 The scan of time step was repeated using the 'momentum conserving'
799 procedure [46] , in which the field values are interpolated at the grid nodes
800 before being gathered onto the particles. The result is given in Fig. 29. With
801 the momentum conserving procedure, the level of instability is consistently
802 high and independent of the time step. Since the numerical dispersion of the
803 solver varies substantially with the time step, this result supports the con-
804 clusion that the instability may not be of numerical Cerenkov nature. The
805 identification of the nature of the instability and the explanation of the sin-
806 gular time step $c\delta t_S$ call for a multidimensional (no instability was observed
807 in 1D regardless of the field gathering method) analysis of the discretized
808 Vlasov algorithm that was employed, which is left for future work.

809 The results that were obtained lead to the following conclusions: (i) the
810 time step $c\delta t_S = \delta z/\sqrt{2}$ consistently produces the lowest levels of instability,
811 regardless of dimensionality (2D vs 3D), the field solver being used, reso-
812 lution, aspect ratio of cells (within the range of the finite number of cases
813 that were experimented); (ii) the main advantage of the tunable field solver
814 resides in allowing access to the singular time step $c\delta t_S$ rather than pro-
815 viding improved numerical dispersion, which consequently do not appear to
816 be a primary driver of the instability; (iii) the instability is not completely
817 removed at $c\delta t_S$ and filtering is still needed, albeit at lower levels; (iv) the
818 field gathering procedure is key, as the existence of a singular time step at
819 which the instability is greatly reduced is observed using an 'energy conserv-
820 ing' procedure, but not using a 'momentum conserving' procedure. These
821 results indicate that the instability that is being observed may not be a type

822 of numerical Cerenkov instability, as originally conjectured.

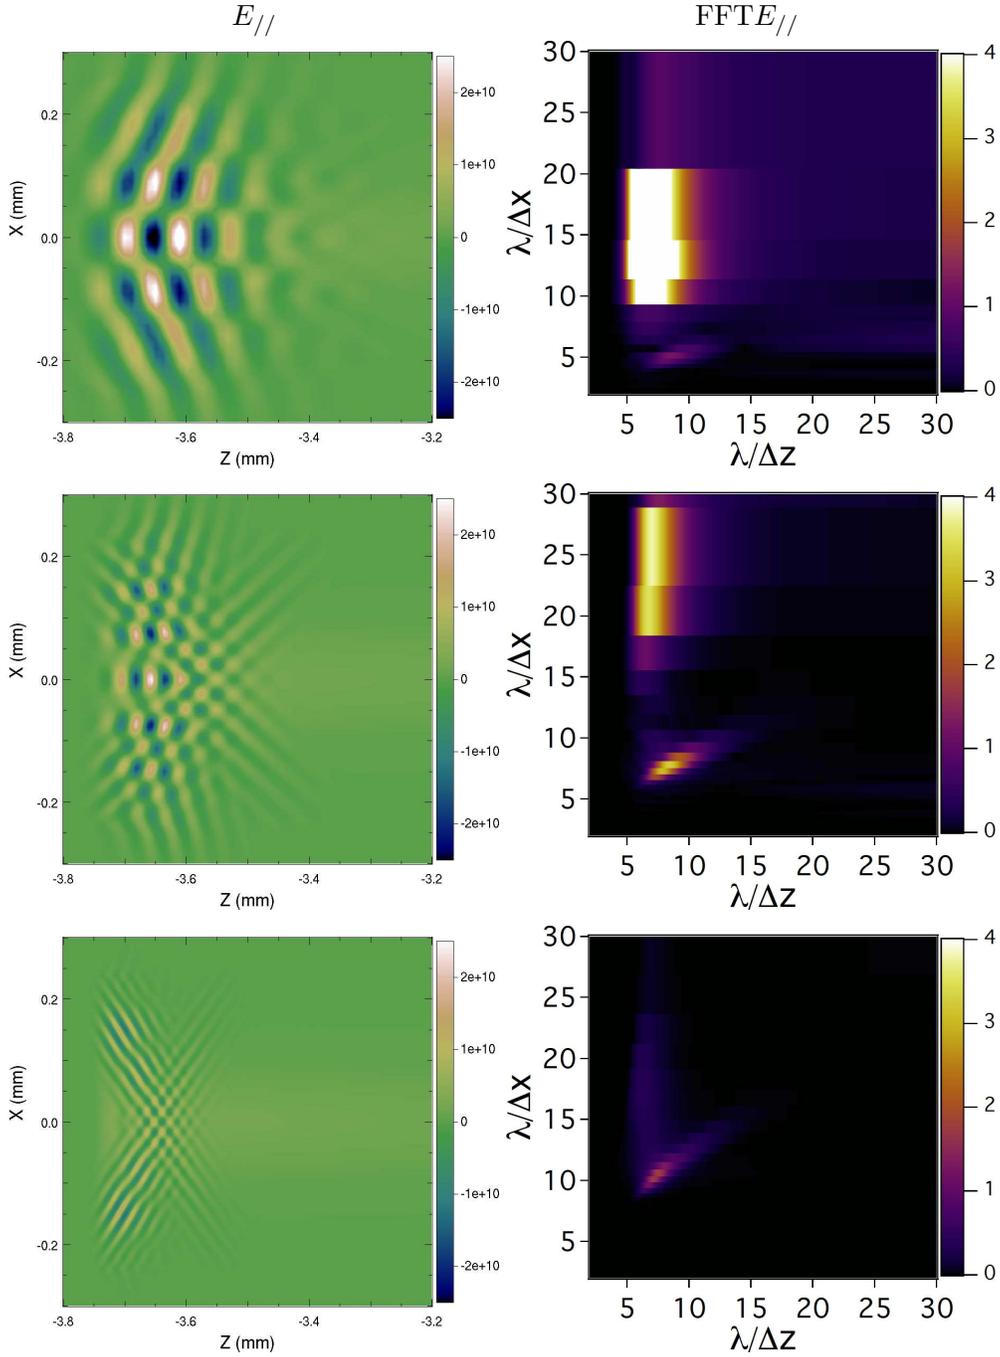


Figure 24: (left) Snapshot of the longitudinal electric field ($E_{//}$) at the front of the plasma at $t = 12.5$ ps; (right) Fourier Transform of the longitudinal electric field, from 2-1/2D simulations of a full scale 10GeV LPA in a boosted frame at $\gamma = 130$, using the Yee solver, for (top) $\delta x = \delta z = 13 \mu\text{m}$; (middle) $\delta x = \delta z = 6.5 \mu\text{m}$; (bottom) $\delta x = \delta z = 3.25 \mu\text{m}$. The time step at the 3D CFL limit $c\delta t = \delta z/\sqrt{3}$ was used for all three simulations.

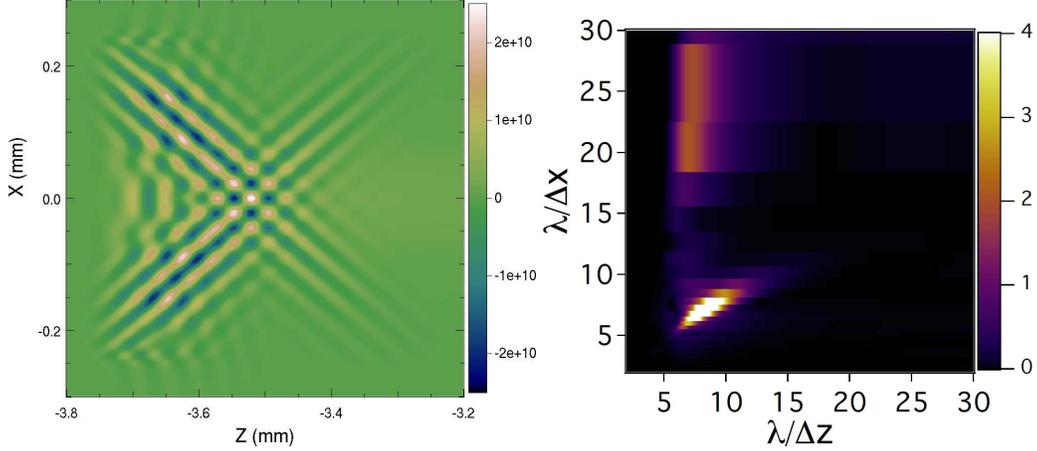


Figure 25: (left) Snapshot of the longitudinal electric field ($E_{//}$) at the front of the plasma at $t = 12.5$ ps; (right) Fourier Transform of the longitudinal electric field, from 2-1/2D simulations of a full scale 10GeV LPA in a boosted frame at $\gamma = 130$, with the CK solver, using $\delta x = \delta z = 6.5 \mu\text{m}$, and the time step at the 3D CFL limit $c\delta t = \delta x/\sqrt{3}$.

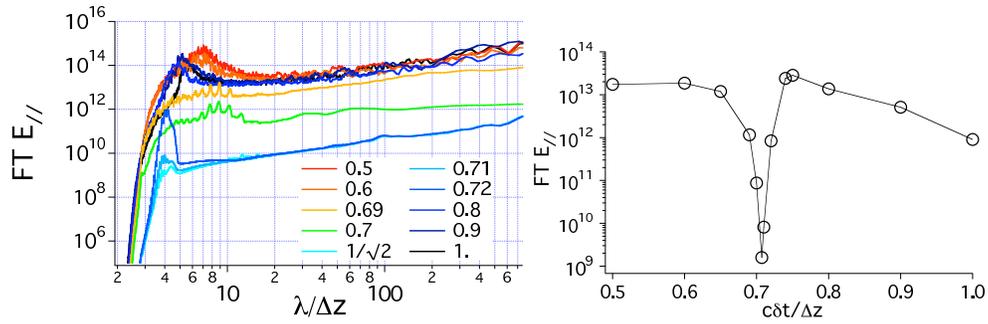


Figure 26: Fourier Transform of the longitudinal electric field at $t=40$ ps, averaged over the whole domain, from 2-1/2D simulations of a full scale 10GeV LPA in a boosted frame at $\gamma = 130$, using the CK solver, for time steps between $c\delta t/\delta z = 0.5$ and $c\delta t/\delta z = 1$, versus $\lambda/\delta z$ (left) and at $\lambda/\delta z = 4$ (right).

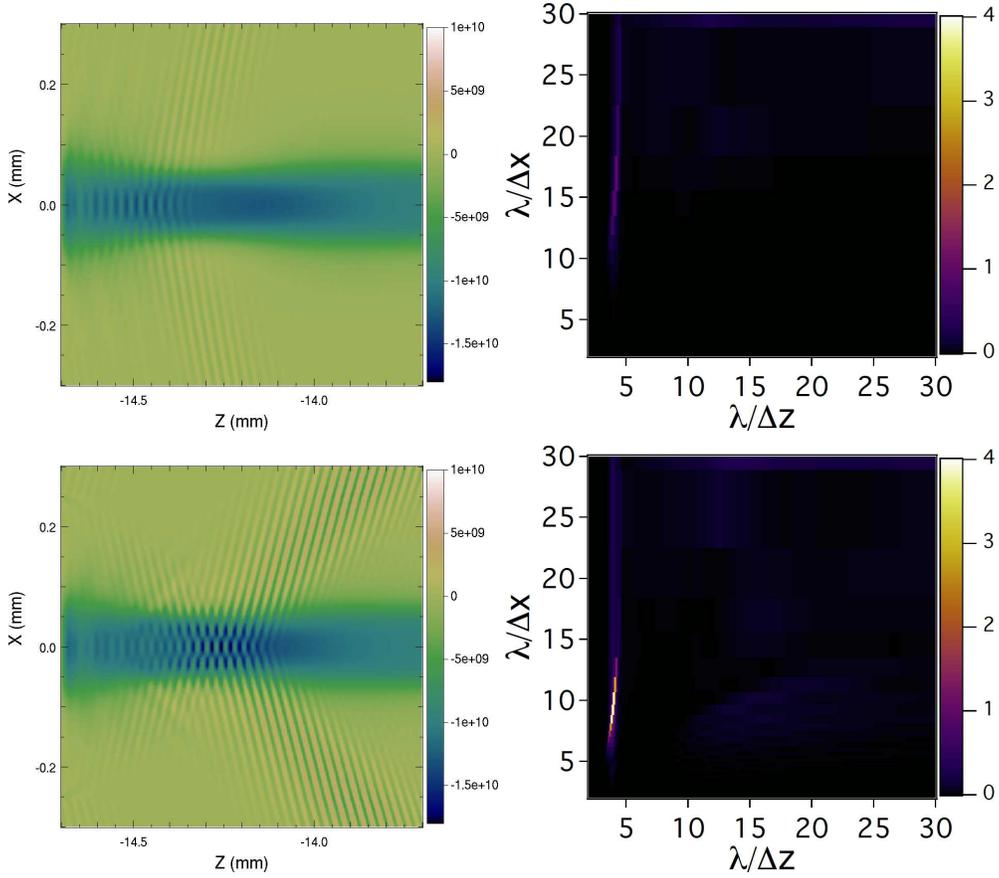


Figure 27: (left) Snapshot of the longitudinal electric field ($E_{//}$) at the front of the plasma at $t = 49$ ps; (right) Fourier Transform of the longitudinal electric field, from 2-1/2D simulations of a full scale 10GeV LPA in a boosted frame at $\gamma = 130$, using $\delta x = \delta z = 6.5\mu m$, 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.

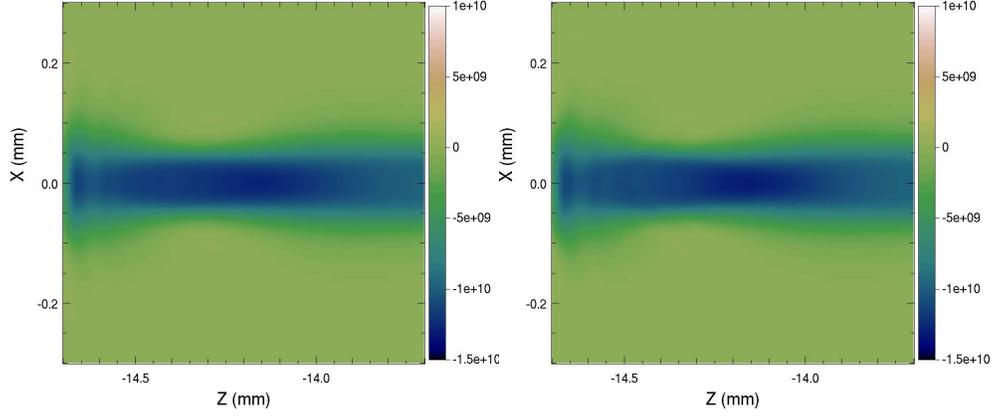


Figure 28: Snapshot of the longitudinal electric field ($E_{//}$) at the front of the plasma at $t = 49$ ps from 2-1/2D simulations of a full scale 10GeV LPA in a boosted frame at $\gamma = 130$, using $\delta_x = \delta_z = 6.5\mu m$, and the time step at the 2D CFL limit $c\delta t = \delta z/\sqrt{2}$, 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. 27. The remaining feature is the wake.

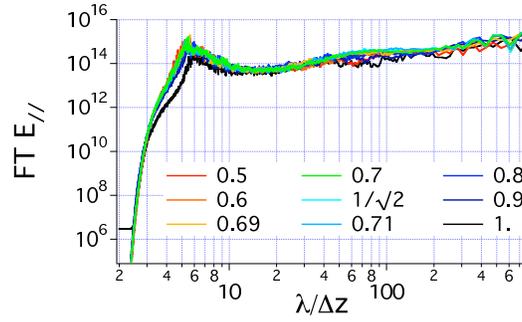


Figure 29: Fourier Transform of the longitudinal electric field at $t=40$ ps, averaged over the whole domain, from 2-1/2D simulations of a full scale 10GeV LPA in a boosted frame at $\gamma = 130$, using the CK solver, for time steps between $c\delta t/\delta z = 0.5$ and $c\delta t/\delta z = 1$, using a 'momentum conserving' field gathering scheme.

823 5.4. Full scale 100 GeV - 1 TeV class stages

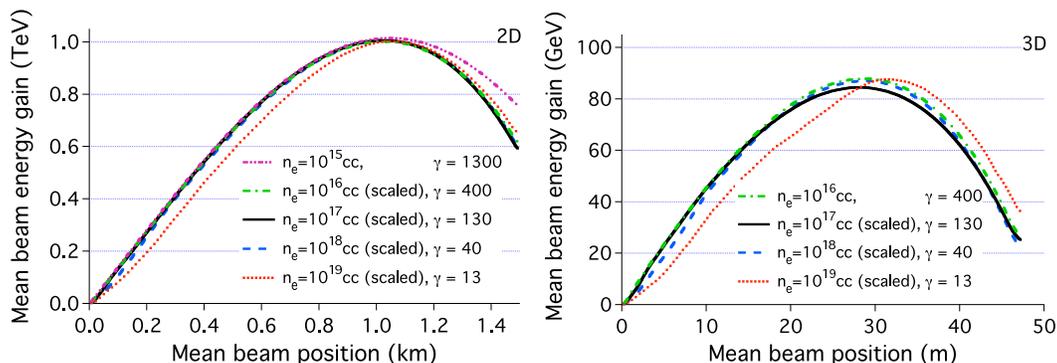


Figure 30: Average beam energy gain versus longitudinal position (in the laboratory frame) for simulations at $n_e = 10^{19} \text{ cc}$ down to 10^{15} cc , using frames of reference between $\gamma = 13$ and $\gamma = 1300$, in 2-1/2D (left) and 3D (right).

824 Using the knowledge acquired from the 10 GeV class study, simulations
825 of stages in the range of 0.1 GeV-1 TeV were performed in 2-1/2D and in
826 the range of 0.1-100 GeV in 3D. The plasma density n_e scales inversely to
827 the energy gain, from 10^{19} cc down to 10^{15} cc in the 0.1 GeV-1 TeV range.
828 These simulations used the parameters given in Table 2 scaled appropriately,
829 and used the high speed of the boosted simulations to allow fast-turnaround
830 improvement of the stage design [18, 19]. Scaled energy gain was increased
831 by adjusting the phase of the beam injection behind the laser by $\sim 12\%$
832 in 3D and 7% in 2D, with respect to the results presented in the preceding
833 section. The 5% level difference between the 2D and 3D beam phases is
834 likely due to small differences in wake structure, laser depletion, and the small
835 number of betatron oscillations of the laser. To minimize beam loss, the beam
836 dimensions were reduced by a factor of 3 in each dimension. Simulations
837 showing performance of this design in 2-1/2D were performed using the Yee

838 solver with filter S(1) for the 0.1-10 GeV runs, S(1:2) for the 100 GeV and
839 S(1:2:3) for the 1 TeV ones. The 3D simulations were performed using the
840 CK2 solver with filter S(1) for the 0.1-1 GeV runs, and S(1:2) for the 10-100
841 GeV ones. The average beam energy gain history is plotted in Fig. 30, scaling
842 the 0.1-100 GeV runs to the 1 TeV range in 2-1/2D, and the 0.1-10 GeV runs
843 to the 100 GeV range in 3D. The results exhibit an excellent agreement on
844 the peak scaled beam energy gain between 0.1-100 GeV runs, and on the
845 scaled beam energy gain histories between the 1-100 GeV runs. A higher
846 level of smoothing was needed for the 1TeV case, explaining the deviation
847 past 1 km. This deviation is of little importance in practice, where one is
848 mostly interested in the beam evolution up-to the peak energy point. The
849 differences at 10^{19} on the scaled beam energy gain history can be attributed
850 to the effects from having only a few laser oscillations per pulse.

851 Using (13), the speedup of the full scale 100 GeV class run, which used
852 a boosted frame of $\gamma = 400$ as frame of reference, is estimated to be over
853 100,000, as compared to a run using the laboratory frame. Assuming the
854 use of a few thousands of CPUs, a simulation that would require several
855 decades to complete using standard PIC techniques in the laboratory frame,
856 was completed in four hours using 2016 CPUs of the Cray system at NERSC.
857 With the same analysis, the speedup of the 2-1/2D 1 TeV stage is estimated
858 to be over a million.

859 **6. Conclusion and outlook**

860 The technique proposed in [1] was applied successfully to speedup by
861 orders of magnitude calculations of laser-plasma accelerators from first prin-
862 ciples. The theoretical speedup estimate from [1] was improved, while com-
863 plications associated with the handling of input and output data between a
864 boosted frame and the laboratory frame were discussed. Practical solutions
865 were presented, including a technique for injecting the laser that is simpler
866 and more efficient than methods proposed previously.

867 Control of an instability that was limiting the speedup of such calcula-
868 tions in previous work is demonstrated, via the use of a field solver with
869 tunable coefficients and digital filtering. The tunable solver was shown to be
870 compatible with existing "exact" current deposition techniques for conserva-
871 tion of Gauss Law, and accommodates Perfectly Matched Layers for efficient
872 absorption of outgoing waves.

873 Extensive testing of the methods presented for numerical Cerenkov miti-
874 gation reveals that choosing the frame of the wake as the frame of reference
875 allows for higher levels of filtering and damping than is possible in other
876 frames with the same accuracy. It also revealed that there exists a singu-
877 lar time step for which the level of instability is minimal, independently of
878 other numerical parameters, especially the numerical dispersion of the solver.
879 This indicates that the observed instability may not be caused by numerical
880 Cerenkov effects. Analysis of the nature of the instability is underway, but
881 regardless of cause, the methods presented mitigate it effectively. The tun-
882 ability of the field solver is key in providing stability in 3D at the singular
883 time step, which is not attainable by the standard Yee solver.

884 The use of those techniques permitted the first calculations in the opti-
 885 mal frame of 10 GeV, 100 GeV and 1 TeV class stages, with speedups over
 886 4, 5 and 6 orders of magnitude respectively over what would be required
 887 by "standard" laboratory frame calculations, which are impractical for such
 888 stages due to computational requirements.

889 These results show that the technique can be applied to the modeling of 10
 890 GeV stages, and future work will include the effects of beam loading, plasma
 891 density ramps, as well as particle trapping in the near future. Future work
 892 on the numerical methods include a comprehensive analysis of the instability
 893 and the existence of a singular time step under certain conditions, as well as
 894 the local application of filtering, smoothing and/or mesh refinement [57, 58]
 895 around the front of the plasma, where the instability develops. The latter is
 896 expected to provide mitigation of the instability while preserving accuracy
 897 in the core of the simulation.

898 7. Appendix I: One dimensional analysis of the CK solver

899 Although the most interesting applications of the CK solver require two or
 900 three dimensions, analysis of the method in one dimension reveals a potential
 901 issue when $c\delta t = \delta x$. In one dimension (choosing x), Equations (45)-(46)
 902 reduce to

$$B_y|_{i+1/2}^{n+1/2} = B_y|_{i+1/2}^{n-1/2} + \frac{\delta t}{\delta x} (E_z|_{i+1}^n - E_z|_i^n) \quad (75)$$

$$E_z|_i^{n+1} = E_z|_i^n + \frac{c^2\delta t}{\delta x} (B_y|_{i+1/2}^{n+1/2} - B_y|_{i-1/2}^{n+1/2}) - \frac{J_i^n}{\epsilon_0} \quad (76)$$

903 Due to uniform time discretization and linearity, the response of the sys-
 904 tem (75)-(76) to arbitrary distributions and evolutions of sources (i.e. macro-

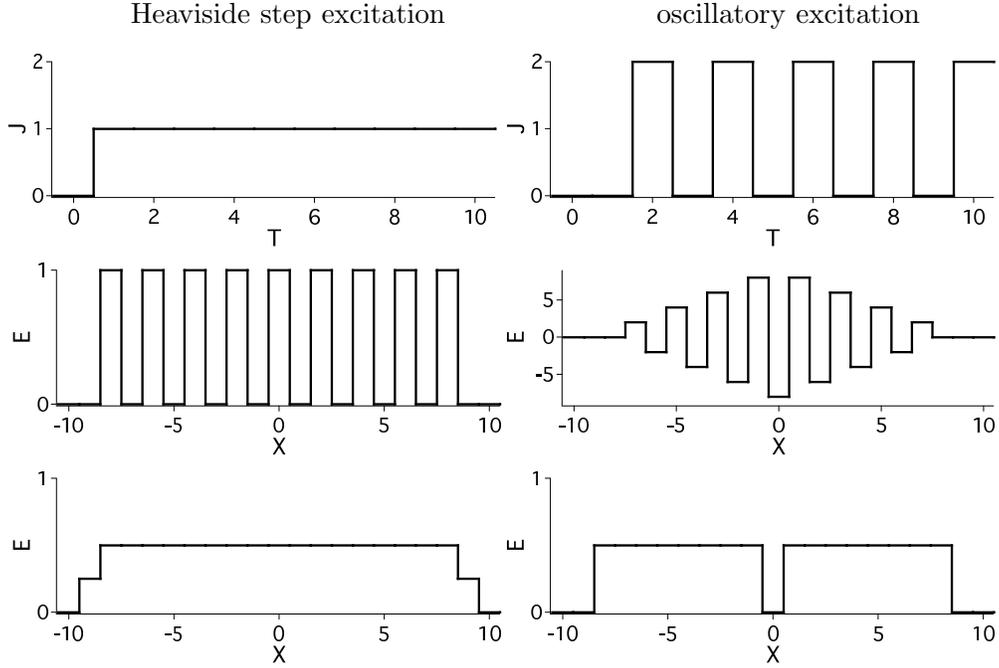


Figure 31: (top) time history (in time steps) of the current source for (left) a Heaviside step (right) a heaviside step modulated by a sinusoidal oscillation at the Nyquist frequency; (middle) response of the system of equations (75)-(76) via a snapshot of the electric field after 10 time steps, without filtering of the source term; (bottom) response of the system of equations (75)-(76) with application of bilinear digital filter of the source term in space. A time step of $c\delta t = \delta x$ was used in all runs and scaled constants $c = \epsilon_0 = 1$ were assumed.

905 particles) can be written as the sum of its response to the excitation from a
 906 Heaviside function in time, at one location in the grid. Assuming a source
 907 term of the form $J|_i^n = H(t)$ where H is the Heaviside function, and setting
 908 the time step at the Courant limit $c\delta t = \delta x$, the system (75)-(76) produces a
 909 spurious "odd-even" oscillations at the Nyquist frequency, as shown in Fig.
 910 31 (middle-left). If a sinusoidal signal oscillating at the Nyquist frequency is
 911 added to the source term, the amplitude of the spurious oscillation grows lin-

912 early with time, as shown in Fig. 31 (middle-right). The spurious oscillation
 913 is effectively suppressed in both cases by the application of a "1-2-1" bilinear
 914 digital filter, as shown in Fig. 31 (bottom) . These types of filtering are of
 915 common use in Particle-In-Cell codes, often repeated a prescribed number of
 916 times and followed by a compensation stage to avoid excessive damping of
 917 long wavelengths [46].

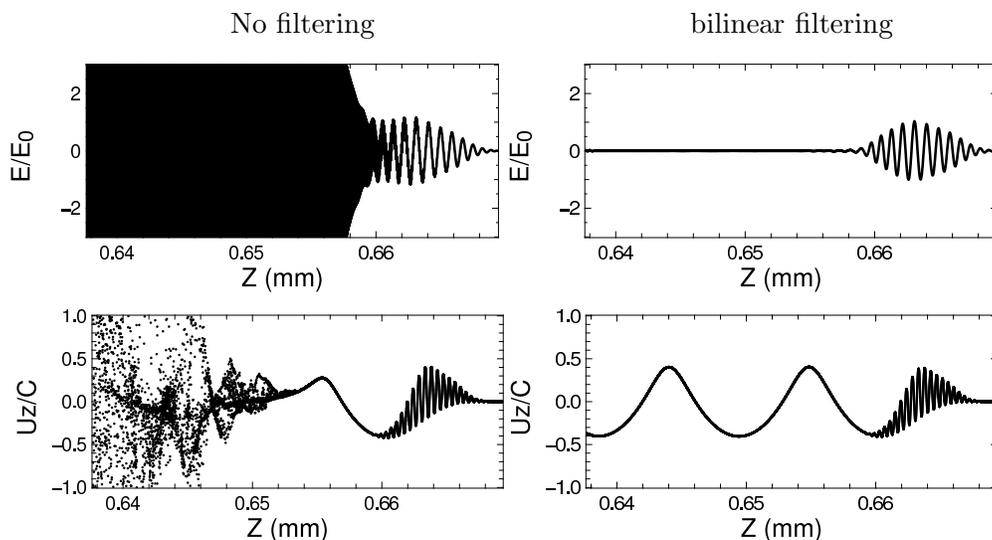


Figure 32: Snapshots from transverse electric field (normalized to maximum laser amplitude E_0) and plasma electrons longitudinal phase space projection, from a 1D simulation of a laser wakefield acceleration stage the CFL limits ($c\delta t = \delta x$) with (left) no filtering of current density; (right) application of a bilinear digital filter to the current density.

918 The impact of the spurious oscillations and the effectiveness of the bi-
 919 linear filtering at suppressing it in actual simulations was tested on a 1D
 920 simulation of a scaled wakefield acceleration stage. The physical and nu-
 921 merical parameters of the simulation are given in table 3. Snapshots of the
 922 transverse electric field (aligned with the laser polarization) and the plasma

923 electron phase space, taken once the laser has propagated about half way
 924 through the plasma (after $\sim 20,000$ time steps) are given in Fig. 32. Without
 925 filtering of the current density, an instability develops at the grid Nyquist
 926 frequency, severely disrupting the plasma wake, despite the fact that cubic
 927 splines were used to deposit current from macro-particles to the grid and
 928 gather the electromagnetic field from the grid to the macro-particles. One
 929 application of the bilinear filtering (without compensation) is sufficient to
 930 suppress the spurious instability and produce a steady and clean wake.

931 8. Appendix II: Perfectly Matched Layer

932 The split form of Perfectly Matched Layer (PML) [52] framework applies
 933 readily to Eqs (45)-(46). The equations on the component along z of the
 934 magnetic field are given by

$$(\Delta_t + \sigma_x) B_{zx} = -\Delta_x^* E_y \quad (77)$$

$$(\Delta_t + \sigma_y) B_{zy} = \Delta_y^* E_x \quad (78)$$

$$(\Delta_t + \sigma_x) E_y = -c^2 \Delta_x (B_{zx} + B_{zy}) \quad (79)$$

$$(\Delta_t + \sigma_y) E_x = c^2 \Delta_y (B_{zx} + B_{zy}) \quad (80)$$

935 where σ_x and σ_y are the absorbing layer coefficients along x and y respectively.
 936 The equations for the other components of the magnetic field and for the
 937 electric field are obtained similarly, applying the standard difference operator
 938 on the spatial derivatives of the electric field and the enlarged difference
 939 operator on the spatial derivatives of the magnetic field. The formula to
 940 update the fields is obtained by solving the finite-difference equations or by

941 integrating over one time step, giving

$$B_{zx}|_{i+1/2,j+1/2,k}^{n+1/2} = \xi_x B_{zx}|_{i+1/2,j+1/2,k}^{n-1/2} - \frac{1-\xi_x}{\sigma_x} \Delta_x^* E_y|_{i+1/2,j+1/2,k}^n \quad (81)$$

$$B_{zy}|_{i+1/2,j+1/2,k}^{n+1/2} = \xi_y B_{zy}|_{i+1/2,j+1/2,k}^{n-1/2} + \frac{1-\xi_y}{\sigma_y} \Delta_y^* E_x|_{i+1/2,j+1/2,k}^n \quad (82)$$

$$E_y|_{i,j+1/2,k}^{n+1} = \xi_x E_y|_{i,j+1/2,k}^n - c^2 \frac{1-\xi_x}{\sigma_x} \Delta_x (B_{zx} + B_{zy})|_{i,j+1/2,k}^{n+1/2} \quad (83)$$

$$E_x|_{i+1/2,j,k}^{n+1} = \xi_y E_x|_{i+1/2,j,k}^n + c^2 \frac{1-\xi_y}{\sigma_y} \Delta_y (B_{zx} + B_{zy})|_{i+1/2,j,k}^{n+1/2} \quad (84)$$

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

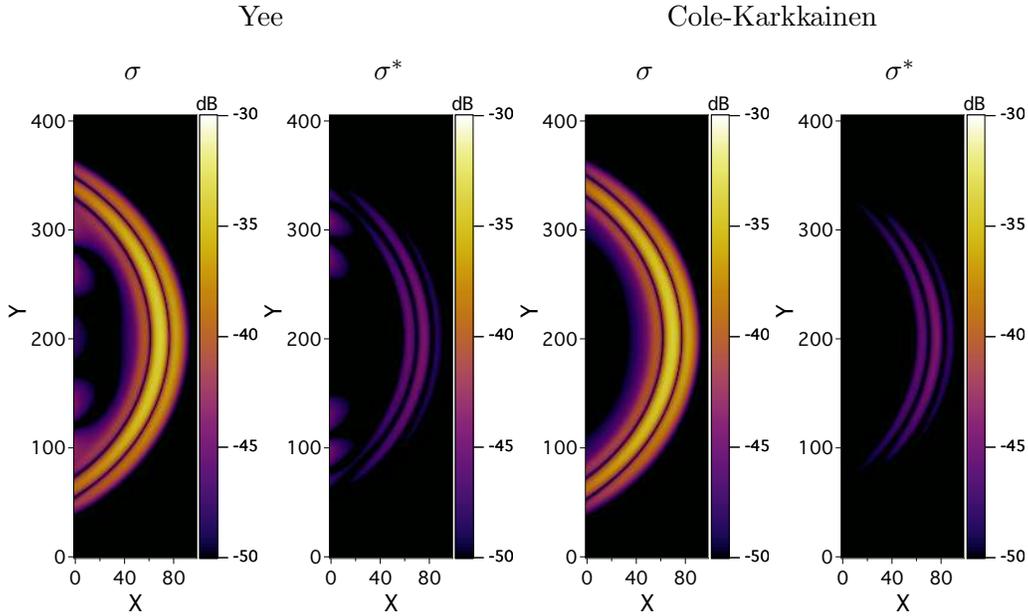


Figure 33: Reflected signal (in dB) from a PML layer using the Yee or the Cole-Karkkainen solver. Each simulation was run for the time step set at the Courant limit.

945 The PML using the stencil given by (84) was tested and compared to
 946 the standard Yee implementation in 2D and 3D. Fig. 33 snapshots from 2D

947 simulations of the reflected residue from a PML layer of a pulse with ampli-
 948 tude given by the Harris function $(10 - 15 * \cos(2\pi ct/L) + 6 * \cos(4\pi ct/L) -$
 949 $\cos(6\pi ct/L))/32$ where t is time, c is the speed of light and $L = 50\delta x$ is the
 950 pulse length in cell size units. A grid of 400x400 cells was used with $\delta x = \delta y$.
 951 The absorbing layer was 8 cells deep and the dependency of the PML co-
 952 efficients with the index position i in the layer was $\sigma_i = \sigma_m (i\delta x/\Delta)^n$ with
 953 $\sigma_m = 4/\delta x$, $\Delta = 5\delta x$ and $n = 2$. The alternative prescription for the coeffi-
 954 cients given in [53, 54], which reads $\sigma_i^* = (\xi_{i+1/2} - 1/\xi_i) / \delta x$ with $\xi_i = e^{-\sigma_i \delta t}$
 955 and $\sigma_i = \sigma_m (i\delta x/\Delta)^n$, was also tested.

956 For the generic test case that has been considered, the new implementa-
 957 tion exhibited a very low residue of reflections from the PML layer, which are
 958 qualitatively and quantitatively very similar to the residue obtained with a
 959 standard PML implementation. In agreement with results from [53, 54], the
 960 use of the modified coefficients σ^* led to an order of magnitude improvement
 961 over the use of the standard coefficients.

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

965 It was shown in [53, 54] that the efficiency of the layer can be improved
 966 further for the standard PML by augmenting the equations with additional
 967 terms. However, a similar extension may not be readily available when using
 968 the Cole-Karkkainen stencil and is not considered here.

969 **9. Acknowledgments**

970 We are thankful to D. L. Bruhwiler, J. R. Cary, B. Cowan, E. Esarey, A.
971 Friedman, C. Huang, S. F. Martins, W. B. Mori, B. A. Shadwick, and C. B.
972 Schroeder for insightful discussions.

- 973 [1] J.-L. Vay, *Phys. Rev. Lett.* **98** (2007) 130405.
- 974 [2] T. Tajima, J. M. Dawson, *Phys. Rev. Lett.*, **43** (1979) 267.
- 975 [3] E. Esarey, *et al.*, *Rev. Modern Phys.* **81**, 252 (2009) 1229.
- 976 [4] C. G. R. Geddes, *et al.*, *Nature* **431**, 538 (2004).
- 977 [5] S. P. D. Mangles, *et al.*, *Nature* **431**, 535 (2004).
- 978 [6] J. Faure, *et al.*, *Nature* **431**, 541 (2004).
- 979 [7] W. P. Leemans, *et al.*, *Nature Physics* **2**, 696 (2006).
- 980 [8] C. B. Schroeder, *et al. Proc. 13th Advanced Accelerator Concepts Work-*
981 *shop*, Santa Cruz, CA (2008) 208.
- 982 [9] C. G. R. Geddes, *et al.*, "Laser Plasma Particle Accelerators: Large
983 Fields for Smaller Facility Sources," *SciDAC Review* **13** (2009) 13.
- 984 [10] C. G. R. Geddes *et al.*, *J. Phys. Conf. Series V* **125** (2008) 12002/1-11.
- 985 [11] C. Huang *et al.*, *J. Phys. Conf. Series* **180** (2009) 12005
- 986 [12] <http://loasis.lbl.gov>

- 987 [13] D. L Bruhwiler *et al.*, *Proc. 13th Advanced Accelerator Concepts Work-*
988 *shop*, Santa Cruz, CA (2008) 29.
- 989 [14] B. A Shadwick, C. B. Schroeder, E. Esarey, *Phys. Plasmas* **16** (2009)
990 056704
- 991 [15] P. Sprangle, E. Esarey, and A. Ting, *Phys. Rev. Letters* **64** (1990) 2011-
992 2014.
- 993 [16] C. Huang *et al.*, *J. of Comput. Phys.* **217** (2006) 658-679.
- 994 [17] B. Feng, C. Huang, V. Decyk, W.B. Mori, P. Muggli, T. Katsouleas, *J.*
995 *Comput. Phys.* **228** (2009) 5340.
- 996 [18] E. Cormier-Michel, *et al. Proc. 13th Advanced Accelerator Concepts*
997 *Workshop*, Santa Cruz, CA (2008) 297.
- 998 [19] C. G. R. Geddes *et al.*, *Proc. Particle Accelerator Conference*, Vancou-
999 *ver*, Canada (2009) WE6RFP075.
- 1000 [20] J.-L. Vay, *Phys. Plasmas*, **15** (2008) 056701.
- 1001 [21] J. P. Boris, *J. Comput. Phys.* **12** (1973) 131-136.
- 1002 [22] I. Haber, R. Lee, H. H. Klein, J. P. Boris, *Proc. Sixth Conf. Num. Sim.*
1003 *Plasmas*, Berkeley, CA (1973) 46-48.
- 1004 [23] B. Cowan, *et al. Proc. 13th Advanced Accelerator Concepts Workshop*,
1005 *Santa Cruz*, CA (2008) 309.
- 1006 [24] J.-L. Vay *et al.*, *Proc. Particle Accelerator Conference*, Vancouver,
1007 *Canada* (2009) TU1PBI04.

- 1008 [25] S. F. Martins, *Proc. Particle Accelerator Conference*, Vancouver,
1009 Canada (2009) TH4GBC05.
- 1010 [26] J.-L. Vay *et al.*, *J. Phys. Conf. Series* **180** (2009) 12006
- 1011 [27] J. -L. Vay, W. M. Fawley, C. G. Geddes, E. Cormier-Michel, D. P. Grote,
1012 arXiv:0909.5603 (Sept. 2009)
- 1013 [28] S. F. Martins, R. A. Fonseca, L. O. Silva, W. Lu, W. B. Mori, *Comput.*
1014 *Phys. Comm.* **182** (2010) 869-875.
- 1015 [29] D. L. Bruhwiler, *Private Communication*.
- 1016 [30] D. P. Grote, A. Friedman, J.-L. Vay, I. Haber, *AIP Conf. Proc.* **749**
1017 (2005) 55.
- 1018 [31] B. B. Godfrey, *J. Comput. Phys.* **15** (1974) 504-521.
- 1019 [32] B. B. Godfrey, *J. Comput. Phys.* **19** (1975) 58-76.
- 1020 [33] A. D. Greenwood, K. L. Cartwright, J. W. Luginsland, E. A. Baca, *J.*
1021 *Comp. Phys.* **201** (2004) 665-684.
- 1022 [34] B. B. Godfrey, *Proc. Ninth Conf. on Num. Sim. of Plasmas* (1980).
- 1023 [35] A. Friedman, *J. Comput. Phys.* **90** (1990) 292.
- 1024 [36] H. Abe, N. Sakairi, R. Itatani, H. Okuda, *J. Comput. Phys.* **63** (1986)
1025 247-267.
- 1026 [37] E. Cormier-Michel, B. A. Shadwick, C. G. R. Geddes, E. Esarey, C. B.
1027 Schroeder, W. P. Leemans, *Phys. Rev. E* **78** (2008) 016404.

- 1028 [38] J. B. Cole, *IEEE Trans. Microw. Theory Tech.*, **45** (1997) 991996.
- 1029 [39] J. B. Cole, *IEEE Trans. Antennas Prop.*, **50** (2002) 11851191.
- 1030 [40] M. Karkkainen, E. Gjonaj, T. Lau, T. Weiland, *Proc. International*
1031 *Computational Accelerator Physics Conference*, Chamonix, France
1032 (2006).
- 1033 [41] K. S. Yee, *IEEE Trans. Ant. Prop.* **14** (1966) 302-307
- 1034 [42] B. Cowan, *Proc. 10th Internat. Comput. Accel. Phys. Conf.*, San Fran-
1035 cisco, CA (2009).
- 1036 [43] J.-P. Bérenger, *J. Comput. Phys.* **114** (1994) 185.
- 1037 [44] F. S. Tsung, *et al.*, *Phys. Plasmas* **13** (2006) 056708.
- 1038 [45] B. Cowan, *Private communication*.
- 1039 [46] C. K. Birdsall and A. B. Langdon, *Plasma Physics Via Computer Sim-*
1040 *ulation* (Adam-Hilger, 1991).
- 1041 [47] A. B. Langdon, *Comput. Phys. Comm.* **70** (1992) 447.
- 1042 [48] B. Marder, *J. Comput. Phys.* **68** (1987) 48.
- 1043 [49] J.-L. Vay, C. Deutsch, *Phys. Plasmas* **5** (1998) 1190.
- 1044 [50] J. Villasenor, O. Buneman, *Comput. Phys. Comm.* **69** (1992) 306.
- 1045 [51] T. Esirkepov, *Comput. Phys. Comm.* **135** (2001) 144-53.
- 1046 [52] J.-P. Berenger, *J. Comput. Phys.* **114** (1994) 185.

- 1047 [53] J.-L. Vay, *J. Comput. Phys.* **165** (2000) 511.
- 1048 [54] J.-L. Vay, *J. Comput. Phys.* **183** (2002) 367.
- 1049 [55] J.-L. Vay, *et al.*, *in preparation*.
- 1050 [56] A. B. Langdon, C. K. Birdsall, *Phys. Fluids* **13** (1970) 2115.
- 1051 [57] J.-L. Vay *et al.*, *Phys. Plasmas* **11** (2004) 2928.
- 1052 [58] J.-L. Vay, J.-C. Adam, A. Héron, *Comput. Phys. Comm.* **164** (2004)
- 1053 171.

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

beam length	L_b	85 nm
beam peak density	n_b	10^{14} cm^{-3}
beam longitudinal profile		$\exp(-z^2/2L_b^2)$
laser wavelength	λ	$0.8 \mu\text{m}$
laser length (FWHM)	L	$10.08 \mu\text{m}$
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		$4 \mu\text{m}$
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