Modeling of 10 GeV-1 TeV laser-plasma accelerators using Lorentz booster simulations

by

J.-L. Vay. C.G.R. Geddes. E. Esarey, C.B.Schroeder, W.P. Leemans. E. Cormier-Michel, D.P. Grote

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

Tech-X Corporation, Boulder CO

Lawrence Livermore National Laboratory, Livermore CA

December 2011

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 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 of 10 GeV-1 TeV laser-plasma accelerators using Lorentz boosted simulations

J.-L. Vay,[1, a)] C. G. R. Geddes,[1] E. Esarey,[1] C. B. Schroeder,[1] W. P. Leemans,[1] E. Cormier-Michel,[2] and D. P. Grote[3]

[1)] *Lawrence Berkeley National Laboratory, Berkeley, CA, USA*

[2)] *Tech-X Corporation, Boulder, CO, USA*

[3)] *Lawrence Livermore National Laboratory, Livermore, CA, USA*

Modeling of laser-plasma wakefield accelerators in an optimal frame of reference [J.-L. Vay, Phys. Rev. Lett. 98 130405 (2007)] allows direct and efficient full-scale modeling of deeply depleted and beam loaded laser-plasma stages of 10 GeV-1 TeV (parameters not computationally accessible otherwise). This verifies the scaling of plasma accelerators to very high energies and accurately models the laser evolution and the accelerated electron beam transverse dynamics and energy spread. 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. Agreement at the percentage level is demonstrated between simulations using different frames of reference for a 0.1 GeV class stage. Obtaining these speedups and levels of accuracy was permitted by solutions for handling data input (in particular particle and laser beams injection) and output in a relativistically boosted frame of reference, as well as mitigation of a high-frequency instability that otherwise limits effectiveness.

PACS numbers: 03.30.+p, 52.38.Kd, 29.20.Ej, 52.65.Rr

[a)] Electronic mail: jlvay@lbl.gov

I. INTRODUCTION

Laser plasma accelerators (LPAs) offer order of magnitude increase in accelerating gradient over standard radio-frequency accelerators (which are limited by electrical breakdown), thus holding the promise of much shorter particle accelerators[1,2]. High quality electron beams of energy up-to 1 GeV have been produced in just a few centimeters[3], with 10 GeV stages being planned as modules of a high energy collider[4,5].

As a 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 (plasma wave) with a periodic structure following the laser. The alternate concentration of positive and negative charges in the wake creates very intense electric fields. An electron (or positron) beam injected with the right phase can be accelerated by those fields to high energy in a much shorter distance than is possible in conventional particle accelerators. The efficiency and quality of the acceleration is governed by several factors which require precise three-dimensional shaping of the plasma column, as well as the laser and particle beams, and understanding of their evolution.

Computer simulations have had a profound impact on the design and understanding of past and present LPA experiments[6–9], with accurate modeling of wake formation, electron self-trapping and acceleration requiring fully kinetic methods (usually Particle-In-Cell) using large computational resources due to the wide range of space and time scales involved[7,9]. Future LPA experiments include those that will be carried out using the BELLA (Berkeley lab laser accelerator) facility at LBNL (Lawrence Berkeley National Laboratory), which will use a 40 J, 1 PW laser system to research the production of 10 GeV electron beams in a meter-length plasma[10]. Simulations of parameters relevant to such a 10 GeV stage demand as many as 5000 processor hours for a one-dimensional simulation on a NERSC supercomputer[11]. Various reduced models have been developed to allow multidimensional simulations at manageable computational costs: fluid approximation[12–14], quasistatic approximation[12,15–18], laser envelope models[12,14,16,17,19], and scaled parameters[20,21]. However, the various approximations that they require result in a narrower range of applicability. As a result, even using several models concurrently does not usually provide a complete description. For example, scaled simulations of 10 GeV LPA stages do not capture correctly some essential transverse physics, e.g. the laser and beam betatron motion, which can lead to inaccurate beam emittance (a measure of the beam quality). An envelope description using a reduce wave operator can capture these effects correctly at full scale for the early propagation through the plasma but can fail as the laser spectrum broadens due to energy depletion as it propagates further in the plasma[13,19,22]. However, capturing depletion accurately is essential to the design of efficient stages, in order to optimize the transfer of energy from the laser to the wake and particle bunch.

An alternative approach allows for orders of magnitude speedup of simulations, whether at full or reduced scale, via the proper choice of a reference frame moving near the speed of light in the direction of the laser[23]. It does so without alteration to the fundamental equations of particle motion or electrodynamics, provided that high-frequency light emitted counter to the direction of propagation of the beam can be neglected. This approach exploits the properties of space and time dilation and contraction associated with the Lorentz transformation. It was shown[23] that the ratio of longest to shortest space and time scales of a system of two or more components crossing at relativistic velocities is not invariant under such a transformation (a laser crossing a plasma is just such a relativistic crossing). Since for simulations based on formulations from first principles, the number of computer operations (e.g., time steps) is proportional to the ratio of the longest to shortest time scale of interest, it follows that such simulations will eventually have different computer runtimes, yet equivalent accuracy, depending solely upon the choice of frame of reference.

The procedure appears straightforward: identify the frame of reference which will minimize the range of space and/or time scales and perform the calculation in this frame. However, several practical complications arise. Most importantly, while the fundamental equations of electrodynamics and particle motion are written in a covariant form, the numerical algorithms that are derived from them may not retain this property, and calculations in frames moving at different velocities may not be successfully conducted with the use of the exact same algorithms. For example, it was shown[24] that calculating the propagation of ultra-relativistic charged particle beams in an accelerator using standard Particle-In-Cell techniques leads to large numerical errors, which were fixed by developing a new particle pusher. The modeling of a LPA stage in a boosted frame involves the fully electromagnetic modeling of a plasma propagating at near the speed of light, for which Numerical Cerenkov[25,26] is a potential issue. Second, the input and output data are usually known from, or compared to, experimental data. Thus, calculating in a frame other than the laboratory entails transformations of the data between the calculation frame and the laboratory frame. Third, electromagnetic calculations that include wave propagation will include waves propagating forward and backward in any direction. For a frame of reference moving in the direction of the accelerated beam (or equivalently the wake of the laser), waves emitted by the plasma in the forward direction expand while the ones emitted in the backward direction contract, following the properties of the Lorentz transformation. If one is to resolve both forward and backward propagating waves emitted from the plasma, there is no gain in selecting a frame different from the laboratory frame. However, the physics of interest for a laser wakefield is the laser driving the wake, the wake, and the accelerated beam. Backscatter is weak in the short-pulse regime, and does not interact as strongly with the beam as do the forward propagating waves which stay in phase for a long period. It is thus often assumed that the backward propagating waves can be neglected in the modeling of LPA stages. The accuracy of this assumption has been demonstrated by comparison between explicit codes which include both forward and backward waves and envelope or quasistatic codes which neglect backward waves[7,21,27].

After the idea and basic scaling for performing simulations of LPAs in a Lorentz boosted frame were published[23], there have been several reports of the application of the technique to various regimes of LPA[9,11,14,28–34]. Speedups varying between several and a few thousands were reported with various levels of accuracy in agreement between simulations performed in a Lorentz boosted frames and in a laboratory frame. High-frequency instabilities were reported to develop in 2D or 3D calculations, that were limiting the velocity of the boosted frame and thus the attainable speedup[31,32,35].

We presented elsewhere[36] numerical techniques that were implemented in the Particle-In-Cell code Warp[37] for mitigating the short wavelength instability, including a solver with tunable coefficients. A detailed study of the application of these techniques to the simulations of scaled LPA stages also revealed that choosing a frame near the frame of the wakefield as the reference frame allows for more aggressive application of filtering or damping for mitigating short wavelength instabilities, than is possible in laboratory frame simulations[36]. We showed that this is due to hyperbolic rotation of the laser oscillations in space time, which is another beneficial consequence of the Lorentz transformation when transforming the laser from the laboratory to a boosted frame, in particular for frames near the frame of the wakefield[38].

In the present paper, we present accurate modeling of 10 GeV-1 TeV LPA stages with beam loading relevant to laser driven collider designs and stages for upcoming lasers[5,20,21], verifying the scaling of efficient, deeply depleted LPAs to very high energies[4]. This is enabled by controlling an instability that develops with high-boost frames by using methods that we developed and presented elsewhere[36,38], allowing 2D and 3D simulations of 100 GeV and 1 TeV class LPA stages in the wakefield frame, thus achieving the maximum theoretical speedups of over 10^5 and 10^6, respectively. Accuracy of the method is demonstrated at the percentage level. This method is used for for the numerical exploration at full scale of the performance of a 10 GeV stage with a 40 J laser, taking accurately into account laser depletion and spectrum broadening, as well as the accelerated electron beam energy spread, and the transverse dynamics of both the laser and the electron beam.

The theoretical speedup expected for performing the modeling of a LPA stage in a boosted frame is derived in Section 2. Section 3 summarizes the issues that have limited speedups in previous work and solutions. Accurate modeling of full scale and scaled 10 GeV class stages is demonstrated in Section 4, and the method is used to simulate stages in the 100 GeV-1 TeV range in Section 5. The evolution of the laser spectrum with respect to the frame boost is given in Appendix A and the consequences on the choice of the optimal boost are discussed. Enabling techniques that were implemented in Warp for input and output of data in a boosted frame are described in Appendix B.

II. THEORETICAL SPEEDUP DEPENDENCY WITH THE FRAME BOOST

The obtainable speedup is derived as an extension of the formula that was derived earlier[23], taking in addition into account the group velocity of the laser as it traverses the plasma. In our previous work[23], the laser was assumed to propagate at the velocity of light in vacuum during the entire process, which is a good approximation when the relativistic factor of the frame boost γ is small compared to the relativistic factor of the laser wake γ_w in the plasma. The expression is generalized here to higher values of γ, for which the actual group velocity of the laser in the plasma must be taken into account. We shall show that for a 10 GeV class LPA stage, the maximum attainable speedup is above four orders of magnitude.

Assuming that the simulation box is a fixed number of plasma periods long, which implies the use (which is standard) of a moving window following the wake and accelerated beam, the speedup is given by the ratio of the time taken by the laser pulse and the plasma to cross each other, divided by the shortest time scale of interest, that is the laser period. To first order, the wake velocity v_w is set by the 1D group velocity of the laser driver, which in the linear (low intensity) limit, is given by[2]:

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

where $\omega_p = \sqrt{(n_e e^2)/(\epsilon_0 m_e)}$ is the plasma frequency, $\omega = 2\pi c/\lambda$ is the laser frequency, n_e is the plasma density, λ is the laser wavelength in vacuum, ϵ_0 is the permittivity of vacuum, c is the speed of light in vacuum, and e and m_e are respectively the charge and mass of the electron.

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

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

where $\lambda_p \approx 2\pi c/\omega_p$ is the wake wavelength, L is the plasma length, v_w and $v_p = \beta_p c$ are respectively the velocity of the wake and of the plasma relative to the frame of reference, and η is an adjustable parameter for taking into account the fraction of the wake which exited the plasma at the end of the simulation. For a beam injected into the n^{th} bucket, η would be set to $n - 1/2$. If positrons were considered, they would be injected half a wake period ahead of the location of the electrons injection position for a given period, and one would have $\eta = n - 1$. The numerical cost R_t scales as the ratio of the total time to the shortest timescale of interest, which is the inverse of the laser frequency, and is thus given by

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

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

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

In a frame moving at βc, the quantities become

$$\lambda_p^* = \lambda_p / \left[\gamma \left(1-\beta_w \beta\right)\right] \tag{5}$$

$$L^* = L/\gamma \tag{6}$$

$$\lambda^* = \gamma \left(1+\beta\right) \lambda \tag{7}$$

$$\beta_w^* = \left(\beta_w - \beta\right) / \left(1-\beta_w \beta\right) \tag{8}$$

$$v_p^* = -\beta c \tag{9}$$

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

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

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

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

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

We note that assuming that $\beta_w \approx 1$ (which is a valid approximation for most practical cases of interest) and that $\gamma << \gamma_w$, this expression is consistent with the expression derived earlier[23] for the LPA case which states that $R_t^* = \alpha R_t / \left(1+\beta\right)$ with $\alpha = \left(1-\beta+l/L\right)/\left(1+l/L\right)$, where l is the laser length which is generally proportional to $\eta\lambda_p$, and $S = R_t/R_T^*$. However, higher values of γ are of interest for maximum speedup, as shown below.

For intense lasers ($a \sim 1$) typically used for acceleration, the energy gain is limited by dephasing[39], which occurs over a scale length $L_d \sim \lambda_p^3/2\lambda^2$. Acceleration is compromised beyond L_d and in practice, the plasma length is proportional to the dephasing length, i.e. $L = \xi L_d$. In most cases, $\gamma_w^2 >> 1$, which allows the approximations $\beta_w \approx 1 - \lambda^2/2\lambda_p^2$, and $L = \xi\lambda_p^3/2\lambda^2 \approx \xi\gamma_w^2\lambda_p/2 >> \eta\lambda_p$, so that Eq.(12) becomes

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

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

$$S_{\gamma<<\gamma_w} = \left(1+\beta\right)^2 \gamma^2 \tag{14}$$

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

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

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

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

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

Note that without the use of a moving window, the relativistic effects that are at play in the time domain would also be at play in the spatial domain[23], and the γ^2 scaling would transform to γ^4. In the frame of the wake, there is no need of the moving window, thus simplifying the procedure, while in a frame traveling faster than the wake in the laboratory, a moving window propagating in the backward direction would be needed. However, the scaling shows that there would be very little gain in doing the latter. Furthermore, analysis presented elsewhere[36,38] and below show that choosing a frame near the frame of the wake is optimum for mitigation of a high frequency instability. This point is refined by a detailed analysis of the laser spectrum on axis in Appendix A which shows that for heavily depleted lasers where the spectrum red-shifts during propagations, the optimal γ might be at a slightly lower value, but this does not greatly affect speedup.

A. Estimated speedup for 0.1-100 GeV stages

Formula (13) is used to estimate the speedup for the calculations of 100 MeV to 1 TeV class stages, assuming a laser wavelength $\lambda = 0.8\mu m$. Parameters for the 100 MeV stage are given in Table 1 below, and parameters for higher energies are derived using scaling laws from[20]. The initial plasma densities n_e for the 100 MeV, 1 GeV, 10 GeV, 100 GeV and 1 TeV stages are respectively 10^{19} cm^{-3}, 10^{18} cm^{-3}, 10^{17} cm^{-3}, 10^{16} cm^{-3} and 10^{15} cm^{-3}, while the plasma lengths L are 1.5 mm, 4.74 cm, 1.5 m, 47.4 m and 1.5 km if choosing $\xi \approx 1.63$. For these values, the wake wavelengths λ_p are respectively 10.6μm, 33.4μm, 106.μm, 334.μm, 1.06mm, and relativistic factors γ_w are 13.2, 41.7, 132, 417 and 1320. In the simulations presented in this paper, the beam is injected near the end of the wake period (first "bucket"). The beam has propagated through about half a wake period to reach full acceleration (due to dephasing), and we set $\eta \approx 0.5$. For the parameters considered here, $L \approx \lambda_p/\gamma_w^2$, and (15) gives $S_{\gamma\to\infty} \approx 2\gamma_w^2$.

The speedup versus the relativistic factor of the boosted frame γ is plotted in Fig. 1-a. As expected, for low values of γ, the speedup scales as (14), and asymptotes to a value slightly lower than $2\gamma_w^2$ for large values of γ. Calculations using the frame of the wake ($\gamma = \gamma_w$) attain nearly the maximum speedup. It is of interest to note that the qualitative behavior is identical to the one obtained in our earlier work[23] (see Fig. 1 and accompanying analysis) in the analysis of the crossing of two rigid identical beams, confirming the generality of the generic analysis presented previously[23]. For 10 GeV, 100 GeV and 1 TeV class stages, the maximum estimated speedups are as large as 3×10^4, 3×10^5 and 3×10^6 respectively. Estimated computational time without boost scales as λ_p^6 (λ_p^3 volume $\times$ λ_p^3 long)$\approx E^3$ (where E is the stage energy) making them harder to model. Fortunately, the boost provides more computational gain for the higher energy stages, making them accessible.

III. NUMERICAL ISSUES IN PAST BOOSTED FRAME SIMULATIONS AND OBSERVED SPEEDUPS

Several numerical limits can restrict the boost performance. Here we review limits in past simulations and their impact on performance (a short wavelength instability, laser initialization, statistics), and present methods for circumventing these limits.

A violent numerical instability developing at the front of the plasma column for $\gamma \gtrsim 100$ in 2D and $\gamma \gtrsim 50$ in 3D was reported[31,32,35] using the Particle-In-Cell codes Osiris[40], Vorpal[41] and Warp[37]. The presence and growth rate of the instability is observed to be very sensitive to the resolution (slower growth rate at higher resolution), and to the amount of damping of high frequencies and filtering of short wavelengths[36]. The instability is always propagating at an angle from the longitudinal axis, and is observed in 2D and 3D runs but was never observed in any of the 1D runs. When modeling an LPA setup in a relativistically boosted frame, the background plasma is traveling near the speed of light and it has been conjectured[32] that the observed instability might be caused by numerical Cerenkov effect. The instability was studied in detail with Warp and effective mitigation was demonstrated on 10 GeV class LPA stages using newly developed algorithms and results[36].

Secondly, boosted frame simulations may require larger simulation boxes in the transverse dimension if the entire laser is to be initialized at $t = 0$, as is common practice for standard laboratory frame simulations[11,30,32]. The Rayleigh length of the laser is contracted by γ in the boosted frame, while the laser duration increases by $\gamma(1+\beta)$, implying an increase of the entire laser spot size by $\gamma^2(1+\beta)$[32]. If the laser is to be initialized entirely in the simulation box at $t=0$, then the simulation box transverse surface increases as $\gamma^4(1+\beta)^2$. Although the cost of the simulation does not scale linearly with the simulation box transverse surface, as most of it is used only for laser initialization and does not contain macro-particles, the scaling is so unfavorable that gains of γ^2 provided by the reduction of time steps can be overtaken by the γ^4 additional costs in grid size, thus limiting the usefulness of the method to low values of γ boost. Diagrams of the laser emission procedures used for boosted frame simulations with the Osiris, Vorpal and Warp are given in Fig. 2. Osiris initializes the entire laser at once and is thus subject to the abovementioned limitations. To circumvent those, Vorpal emits the laser from all but one faces of the simulation box[11] using total field/scattered field technique[42], while Warp emits via a moving planar antenna as described in Appendix B.

Thirdly, for a given number of plasma macro-particles per cell, the total number of macro-particles in the entire plasma column goes down as $1/\gamma^2$ where γ is the relativistic factor of the Lorentz boost[32]. However, simulations of self-injection regimes require a sufficient number of macro-particles in the plasma column so that adequate statistics ensues in the number of trapped macro-electrons, imposing a ceiling in the value of γ that can be used. For a typical scheme, a $\gamma_{max} \simeq 50$ was derived[32] using purely statistical arguments assuming the usage of macro-particles of equal weights. This limit might be relaxed by using varying macro-particle weights such that regions with high probability of trapping (as determined from the accumulated knowledge of previous work) are populated with a higher density of macro-particles of smaller weights. This is already practiced in ordinary runs (i.e. without boosted-frame) for minimizing the computational cost while maximizing the statistics within "dynamically interesting" regions[43]. For instance it is found[6,44,45] that in the bubble regime, self-injected particles are initially located within a relatively narrow ring region along the laser axis whose radius is of the order of the laser waist. Previous simulations can be utilized to determine exactly the radius and thickness of the ring region. This issue does not affect the modeling of stages with external injection that will be considered in this paper.

Observed speedups from simulations using the Particle-In-Cell codes Osiris, Vorpal and Warp are plotted for 0.1 GeV to 1 TeV stages in Fig. 1-b and contrasted to the theoretical speedups from Eq. (13). All three codes were using the same standard Particle-In-Cell method[46]. They all successfully performed 2D and/or 3D calculations with boosts at γ in the range of 20-70, reaching speedups over three orders of magnitude (projected for Osiris assuming no computational cost from laser injection). Without the use of special techniques to mitigate the short wavelength instability, none of the codes could perform successfully 2D or 3D simulations for γ boost values over 100. With the use of the special techniques described elsewhere[36] and in Appendix B, Warp simulations were successfully performed using γ boost as high as 1,300 in 2D and 400 in 3D for 1 TeV and 100 GeV class stages respectively.

It is important to note that observed speedups were obtained from simulations of different setups and thus do not offer a direct comparison of the merits of the different codes with regards to boosted frame simulations: Osiris simulations were of trapped self-injection stages, while Vorpal and Warp simulations were of external injection stages with beam loading. Furthermore, while Vorpal and Warp simulations used special procedures to launch the laser that minimize the transverse grid size, Osiris' did not and used transverse grid sizes that were notably larger (as described above). This made Osiris runs in boosted frames substantially more costly, which does not show in the speedups reported by Osiris as this effect was not factored in. However, it is also important to recognize that the most important limiting factor was the high frequency instability (observed in 2D and 3D), which seems to have affected boosted frame simulations of LPA equally, independently of the code used or the simulated LPA setup, legitimizing the comparison in this respect.

The numerical techniques that were developed and implemented in Warp[36] [as described in Appendix B] are used in the next sections to demonstrate stability and convergence of the boosted frame method up to the 100 GeV-1 TeV range.

IV. MODELING 10 GEV CLASS LASER PLASMA ACCELERATION STAGES

This section presents the modeling of deeply depleted and beam loaded 10 GeV LPA stages at full scale in 2-1/2D and 3D using the new numerical techniques that we implemented in Warp in[36] and in Appendix B, which has not been done fully self-consistently without the Lorentz boost method.

It has been shown that many parameters of high energy LPA stages can be accurately simulated at reduced cost by simulating stages of lower energy gain, with higher density and shorter acceleration distance, by scaling the physical quantities relative to the plasma wavelength, and this has been applied to design of 10 GeV LPA stages[20,21]. The number of oscillations of a mismatched laser pulse in the plasma channel however depends on stage energy and does not scale, though this effect is minimized for a channel guided stage[20,21]. The number of betatron oscillations of the trapped electron bunch will also depend on the stage energy, and may affect quantities like the emittance of the beam. For these reasons, and to prove validity of scaled designs of other parameters, it is necessary to perform full scale simulations, which is only possible by using reduced models (e.g., Ref.[19]) or simulations in the boosted frame.

The basic prescription for scaling a LPA simulation to lower plasma density and higher electron energy gain can be briefly summarized as follows (for additional details see Refs[20,21]). First, a fully resolved simulation (i.e., sufficient number of grid points per laser wavelength λ) is performed at a relatively low value of $\lambda_p/\lambda = k/k_p$ (i.e., at a relatively high plasma density for a fixed laser wavelength). Next, the simulation results are scaled to a higher value of k/k_p (i.e., lower density) by keeping the normalized laser and beam parameters fixed (constant a_0, k_pL, $k_p\sigma$, $k_p\sigma_z$, $k_p\sigma_r$, and n_b/n_0, where a_0 is the laser normalized vector potential, k_p is the plasma wavenumber, n_b is the electron beam density, L and σ are the longitudinal and transverse sizes of the laser, and σ_z and σ_r are the longitudinal and transverse sizes of the beam), since these normalized parameters determine the structure of the accelerating and focusing plasma wakefields. The acceleration length L_d (e.g., the length for the electron beam to reach maximum energy) and the electron energy gain γ_{max} scale as $L_dk_p^3/k^2 =$ constant and $\gamma_{max}k_p^2/k^2 =$ constant, since the dephasing and depletion lengths scale as $L_d \sim k^2/k_p^3$ and the accelerating field scales as $E_z \sim k_p$. As noted above, some physically relevant quantities do not remain constant when scaled to higher values of k/k_p, such as the trapping threshold for particles in the wake[47] as well as the normalized Rayleigh length, $k_pZ_R = kk_p\sigma^2/2$, which determines for example the number of oscillations of a mismatched laser pulse in a plasma channel. For this reason, in general, fully resolved simulations at the correct value of k/k_p are still desirable.

For benchmarking purposes, scaled simulations[20] are performed, first at a density of $n_e = 10^{19}$ cm^{-3}, using various values of the boosted frame relativistic factor γ to show the accuracy and convergence of the technique. These scaled simulations were shown to efficiently accelerate both electrons and positrons with low energy spread, and predicted acceleration of hundreds of pC to 10 GeV energies using a 40 J laser. The accuracy of the boosted frame technique is evaluated by modeling scaled stages[20,21] at 0.1 GeV, which allows for a detailed comparison of simulations using a reference frame ranging from the laboratory frame to the frame of the wake. Excellent agreement is obtained on wakefield histories on axis, beam average energy and transverse RMS size histories, and momentum spread at peak energy, with speedup over a hundred, in agreement with the theoretical estimates from Section 2. The boosted frame technique is then applied in the next section to provide full scale simulation of high efficiency quasilinear LPA stages at higher energy, verifying the scaling laws in the 10 GeV-1 TeV range.

A. Scaled 10 GeV class stages

The main physical and numerical parameters of the simulations are given in Table I. They were chosen to be close (though not identical) to a case reported elsewhere[20] with $k_pL = 2$ where L is the laser pulse length as defined in Table I, the main differences being a sinusoidal versus gaussian laser longitudinal profile and a laser spot size larger by $\sqrt{2}$. These simulations are for a fully resolved 100 MeV stage at a density of 10^{19} cm^{-3}, which can be scaled to describe a 10 GeV stage at a density of 10^{17} cm^{-3}, thereby allowing short run times to permit effective benchmarking between the algorithms[20,21]. These runs were done using the standard Yee solver with no damping, and with the 4-pass stride-1 filter plus compensation[36]. No signs of detrimental numerical instabilities were observed at the resolutions reported here with these settings in 2-1/2D or 3D.

For the given parameters, the wake relativistic factor $\gamma_w \approx 13.2$. Thus, Warp simulations were performed using reference frames moving between $\gamma = 1$ (laboratory frame) and 13. For a boosted frame associated with a value of γ approaching γ_w in the laboratory, the wake is expected to travel at low velocity, and the physics to appear somewhat different from that observed in the laboratory frame, in accordance to the properties of the Lorentz transformation[38]. Figures 3 and 4 show surface renderings of the transverse and longitudinal electric fields respectively, as the beam enters its early stage of acceleration by the plasma wake, from a calculation in the laboratory frame and another in the frame at $\gamma = 13$. The two snapshots offer strikingly different views of the same physical processes: in the laboratory

TABLE I. List of parameters for a LPA stage simulation at 100 MeV

plasma density on axis	n_e	10^{19} cm^{-3}
plasma longitudinal profile		flat
plasma length	L_p	1.5 mm
plasma entrance ramp profile		half sine
plasma entrance ramp length		20 μm
laser profile		$a_0 \exp\left(-r^2/2\sigma^2\right) \sin\left(\pi z/3L\right)$
normalized vector potential	a_0	1
laser wavelength	λ	0.8 μm
laser spot size (RMS)	σ	8.91 μm
laser length (HWHM)	L	3.36 μm
normalized laser spot size	$k_p\sigma$	5.3
normalized laser length	k_pL	2
beam profile		$n_{b0} \exp\left(-r^2/2\sigma_r^2 - z^2/2\sigma_z^2\right)$
beam transverse size (RMS)	σ_r	165 nm
beam length (RMS)	σ_z	85 nm
normalized beam spot size	$k_p\sigma_r$	0.1
normalized beam length	$k_p\sigma_z$	0.05
beam transverse emittance	ϵ	73.5 nm.mrad
beam total charge (3D)	Q	6.42 pC
beam initial energy	E_0	$1.5\gamma_w m_e c^2$
injection distance after laser max		$0.7\lambda_p$
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	$0.65\mu m$
cell size in z	δz	$\lambda/64$
time step	δt	at CFL limit
particle deposition order		cubic
# of plasma particles/cell		1 macro-e$^-$+1 macro-p$^+$

frame, the wake is fully formed before the beam undergoes any significant acceleration and the imprint of the laser is clearly visible ahead of the wake; in the boosted frame calculation, the beam is accelerated as the plasma wake develops, and the laser imprint is not visible on the snapshot. Close examination reveals that the short spatial variations which make the laser imprint at the front of the wake are transformed into time variations in the boosted frame of $\gamma = 13$. This effect is due to hyperbolic rotation in Minkowski space of the laser propagation in plasma, as explained in more detail elsewhere[38]. The imprint of the beam loading is clearly visible on the plot of the longitudinal electric field (wake) in the laboratory frame (top plot of Figure 4).

Histories of the perpendicular and longitudinal electric fields recorded at a number of stations at fixed locations in the laboratory offer direct comparison between the simulations in the laboratory frame ($\gamma = 1$) and boosted frames at $\gamma = 2$, 5, 10 and 13. Figure 5 and 6 show respectively the transverse and longitudinal electric fields collected at the positions $z = 0.3$ mm and $z = 1.05$ mm (in the laboratory frame) on axis ($x = y = 0$). The agreement is excellent and confirms that despite the apparent differences from snapshots taken from simulations in different reference frames, the same physics was recovered. The effect of beam loading is visible in Figure 6 at $t \approx 1.15$ ps and $t \approx 3.61$ ps, confirming that the amplitude and phase of beam loading was correctly recovered in all frames. This is further confirmed by the plot of the average scaled beam energy gain and transverse RMS size as a function of position in the laboratory frame, and of relative longitudinal momentum dispersion at peak energy (Fig. 7). These show that the correct laser evolution and electron beam energy, momentum spread and transverse dynamics were modeled in all frames. The small differences observed in the mean beam energy histories and on the longitudinal momentum spread are due to a lack of convergence at the resolution that was chosen, and we have verified that convergence was improving with increasing resolution. The beam was launched with the same phase in the 2-1/2D and the 3D simulations, resulting in lower energy gain in 3D, due to proportionally larger laser depletion effects in 3D than in 2-1/2D.

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

All the simulations presented so far in this section were using the Yee solver[36], for which the Courant condition is given by $c\delta t < \left(1/\delta x^2 + 1/\delta z^2\right)^{-1/2}$ in 2D and $c\delta t < \left(1/\delta x^2 + 1/\delta y^2 + 1/\delta z^2\right)^{-1/2}$ in 3D where δt is the time step and δx, δy and δz are the computational grid cell sizes in x, y and z. As γ was varied, the transverse resolution was kept constant, while the longitudinal resolution was kept at a constant fraction of the incident laser wavelength $\delta z = \zeta\lambda$, such that in a boosted frame, $\delta z^* = \zeta\lambda^* = \zeta\left(1+\beta\right)\gamma\lambda$. As a result, the speedup becomes, when using the Yee solver

$$S_{yee2D} = S\frac{\delta z\sqrt{1/\delta x^2 + 1/\delta z^2}}{\delta_z^*\sqrt{1/\delta x^2 + 1/\delta z^{*2}}} \tag{17}$$

in 2D and

$$S_{yee3D} = S\frac{\delta z\sqrt{1/\delta x^2 + 1/\delta y^2 + 1/\delta z^2}}{\delta_z^*\sqrt{1/\delta x^2 + 1/\delta y^2 + 1/\delta z^{*2}}} \tag{18}$$

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

The speedup versus relativistic factor of the reference frame is plotted in Fig. 8, from (13), (17) and (18), and contrasted with measured speedups from 1D, 2-1/2D and 3D Warp simulations, confirming the scaling obtained analytically.

This subsection demonstrated accurate modeling of the evolution of the laser and the electron beam energy, momentum spread and transverse dynamics with agreement at the percentage level between simulations using various reference frames. The scaling of the speedup was also confirmed, validating our understanding of the boosted frame method scaling with γ boost.

B. Full scale 10 GeV class stages

The boosted frame technique was next applied to the modeling of 10 GeV stages at full scale (i.e. at plasma density $n_e = 10^{17}$ cm^{-3}, with parameters scaled from Table 1). As noted elsewhere[11], full scale simulations using the laboratory frame of 10 GeV stages at plasma densities of 10^{17} cm^{-3} are not practical on present computers in 2D and 3D. At this density, the wake relativistic factor $\gamma_w \approx 132$, and 2-1/2D and 3D simulations were done in boosted frames up to $\gamma = 130$, realizing the maximum theoretical speedup. This section demonstrates accurate modeling of the particle beam acceleration and transverse dynamics evolution of full scale beam loaded 10 GeV stages in 2D and 3D for boosted frames up to $\gamma = 130$.

Fig. 9 and Fig. 10 show the average beam energy gain and transverse RMS size versus longitudinal position from respectively 2-1/2D and 3D simulations in boosted frames at $\gamma = 30$ to 130 in 2-1/2D and at $\gamma = 60$ to 130 in 3D (runs at $\gamma = 1$ are impractical and were not performed). All runs gave the same beam energy history within a few percent. The average energy gain peaks around 10 GeV in 2-1/2D and 8 GeV in 3D, in agreement with the scaled simulations (see Fig. 7). The abovementioned short wavelength instability that occurs at high values of γ boost is described elsewhere[36] and has been mitigated in the 3D simulations using $\gamma \geq 120$ using a novel electromagnetic solver and time step ($c\delta t/\delta z = 1/\sqrt{2}$) for which the instability growth rate is greatly reduced, completed by smoothing of short wavelengths[36]. The small discrepancy between the results of the runs at various γ is due to lack of convergence and difficulty in attaining identical initial conditions (see Appendix B) at the resolution that was chosen (32 grid cells per laser wavelength in vacuum). Preliminary scans with varying resolution (not shown here) show that agreement improves with higher resolution and suggest that boosted frame simulations may converge faster than laboratory frame simulations.

The boosted frame technique was applied to the direct simulation of a 10 GeV stage ($n_e = 10^{17}$ cm^{-3}) in which the accelerated charge was sufficiently high so that the effects of beam loading were readily evident. These parameters are relevant to experiments that will be carried out on new lasers, such as the BELLA facility at LBNL[10], and to LPA stages that can serve as the basis for high energy collider modules[4,5]. In this simulation, a laser pulse with intensity $a_0 = 1.414$, wavelength $\lambda = 0.8$ microns and RMS duration of $L/c = 40$ fs (i.e. $k_pL = 1/\sqrt{2}$), where L is the longitudinal RMS size of the gaussian laser pulse profile $a(r, z) = a_0 \exp\left(-r^2/2\sigma^2 - z^2/2L^2\right)$ was focused to a RMS gaussian transverse spot size of 63 microns (i.e. $k_p\sigma = 3.75$) at the channel entrance. The plasma channel had an on-axis density $n_0 = 10^{17}$ cm^{-3}, a length of 0.75 m with a parabolic channel (factor=0.6) and a longitudinal taper[20,48] of the form $n(x) = n_0(1.32x + 1)$. An electron beam with a gaussian profile and 17 microns size (i.e. $k_p\sigma_r = 1$) and 8.5 microns length (i.e. $k_p\sigma_z = 0.5$) was externally injected at a distance of $1.53\lambda_p$ behind the location of maximum laser intensity with an initial energy of 100 MeV and an initial emittance of 60 mm-mrad. The calculation was in 2-1/2D, and the beam charge density corresponded to a total charge of ~ 52 pC in 3D. The large input emittance was chosen to maximize the beam radius for efficient beam loading and for emittance matching to the wakes focusing fields[20]. Figure 11 shows the density wake excited by an intense laser pulse and the externally injected electron beam accelerated by the wake. The color coding indicates the energy reached by the electrons. The depression in the density wake is due to self-consistent beam loading of the injected electron bunch. The histories of electron beam mean energy, emittance and RMS size are given in Figure 12, while the longitudinal momentum distribution is given in Figure 13. At the exit of the structure, electrons with energy of up to 11 GeV were observed. The time projected energy spread and normalized emittance when exiting the plasma channel were 15% and 61 mm-mrad, respectively. The slice energy spread and emittance of a slice at 9.5 GeV were 1% and 54 mm-mrad. Whereas these values are larger than acceptable for collider and light source applications, it has been shown that lower emittance bunches can be accelerated by using high order laser modes to control the transverse focusing forces[49] and lower energy spread by controlling beam loading[20,50]. Future work will aim at optimizing the phase space properties of the bunch including optimization of taper[48] and use of higher order laser modes to minimize emittance.

The present work demonstrates the ability to simulate at full scale a 10 GeV stage that exhibits significant laser depletion and beam loading. It confirms that the electron beam acceleration and energy gain is very well predicted by scaled simulations, and shows that emittance conservation is obtained through good matching, which is only accurately accounted for at full scale.

V. FULL SCALE 100 GEV-1 TEV STAGES

The numerical techniques that we developed[36] and successfully applied to the modeling of 10 GeV class stages in the preceding section are applied in this section to the modeling of stages in the 0.1 GeV-1 TeV range in 2-1/2D and in the 0.1-100 GeV range in 3D, showing scaling of LPAs to high energies. The plasma density n_e scales inversely to the energy gain, from 10^{19} cm^{-3} down to 10^{15} cm^{-3} in the 0.1 GeV-1 TeV range. These simulations used the parameters given in Table I scaled appropriately to higher energies[20], and used the high speed of the boosted simulations to allow fast turnaround improvement of the stage design presented in[20,21].

The average beam energy gain history is plotted in Fig. 14, scaling the 0.1-100 GeV runs to the 1 TeV run in 2-1/2D, and the 0.1-10 GeV runs to the 100 GeV run in 3D. The differences at 10^{19} cm^{-3} of the scaled beam energy gain history can be attributed to the effects from having only a few laser oscillations per pulse.

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

This section demonstrated the scaling of highly depleted beam loaded stages up to 1 TeV in 2D and 100 GeV in 3D, providing greater credibility for evaluation of various LPA based collider options[4].

VI. CONCLUSIONS

Calculations using an optimal boosted frame of 10 GeV, 100 GeV and 1 TeV class stages including beam loading were presented, with speedups over 4, 5 and 6 orders of magnitude respectively over what would be required by "standard" laboratory frame calculations, which are impractical for such stages due to computational requirements. Our previous theoretical speedup estimate[23] was extended to high boost values, while complications associated with the handling of input and output data between a boosted frame and the laboratory frame were discussed. Practical solutions were implemented, including a technique for injecting the laser that is simpler and more efficient than methods proposed previously.

The boosted frame Particle-In-Cell technique has been shown to accurately model the laser evolution and resolve the wavelength shifting and broadening (as described in Appendix A) that occurs as the laser depletes, offering advantages over other models (for example envelope, quasistatic) while providing the speed required for direct simulation of 10 GeV and beyond laser plasma accelerators to accurately model laser and beam transverse oscillations. It has been shown to also model accurately the electron beam acceleration, longitudinal and transverse dynamics. The results are within a few percent of those from 'standard' laboratory frame simulations, which is within acceptable range for the design of proof-of-principle experiments. The boosted frame technique is being applied to the direct simulation of 10 GeV beam loaded stages, which is relevant to experiments on new lasers (e.g., the

BELLA facility at LBNL), as well as next generation controlled laser plasma accelerator stages and collider modules.

VII. ACKNOWLEDGMENTS

We are thankful to C. Benedetti, D. L. Bruhwiler, J. R. Cary, B. Cowan, A. Friedman, C. Huang, S. F. Martins, W. B. Mori, and B. A. Shadwick for insightful discussions. Work supported by US-DOE Contracts DE-AC02-05CH11231 and DE-AC52-07NA27344, and US-DOE SciDAC program ComPASS. Used resources of NERSC, supported by US-DOE Contract DE-AC02-05CH11231.

Appendix A: Laser spectrum on axis and optimal frame

In this section, we discuss in more detail how the choice of the optimal frame for smoothing is determined by the laser spectrum. It was shown previously[38] that choosing γ boost near γ of the wakefield is a possible option. We extend the discussion to consider depletion of the laser and show that in this case a lower value of γ boost might be desirable.

The spectrum history of the laser field on axis is given in Fig. 15 for selected values of γ between 1 and 135. The history is given up to the time of the electron beam peak energy. In the laboratory frame at $\gamma = 1$, the initial (t=0) spectral content is very localized in a narrow band around the laboratory frame vacuum laser wavelength λ_0, spreading and redshifting as the laser propagates and depletes its energy into the wake[2]. Although it is not visible in the spectrum, the laser waves propagate in the positive direction in the laboratory frame.

At higher values of γ frame, the initial spectral content of the laser shifts to longer wavelengths relative to the boosted frame vacuum laser wavelength λ_0'. As the frame approaches the wake frame $\gamma_w \approx 132$, the initial spectrum is displaced toward very long wavelengths (standing waves), because the frame is moving near the laser group velocity. At later times, the high γ frame spectra show content repopulating progressively shorter wavelengths. This corresponds to the redshifting observed in the lab frame; with the calculation frame matching the initial laser group velocity, the redshifted light which propagates slower now slips backward in the moving frame. As γ frame rises, eventually all waves propagate in the negative direction for $\gamma \geq \gamma_w$.

Mitigation of the short wavelength instability necessitates higher amounts of smoothing at higher γ, and smoothing is most effective (and has minimal effect on simulation physics) when spectral content is confined to long wavelengths[36]. This occurs for $\gamma \approx \gamma_w$ initially, and Fig. 15 indicates how for strongly depleted stages the optimum γ may be adjusted slightly below γ_w, in order to maximize the wavelengths of the average spectral content over the propagation length rather than only at the start. The plasma column also usually exhibits a parabolic transverse profile so as to provide transverse focusing of the laser, which also slightly reduces group velocity[2]. These corrections to the optimal gamma are relatively small and thanks to the weak dependency of the speedup with γ near γ_w (cf Fig. 1), simulations with γ approaching γ_w offer speedups that are very near the maximum attainable, thus offering in practice the maximum benefit of the boosted frame technique while maintaining the highest level of accuracy. Other effects such as tapering of the plasma density may further decrease the optimal γ for smoothing. A large value of γ boost can still be used for high energy stages, thus achieving orders of magnitude speedups in practice.

At higher resolution, the instability level is reduced[36], and so is the amount of smoothing that is necessary to control it. Furthermore, the instability spectrum is confined to a very narrow band located near the Nyquist cutoff of the simulation grid[36], and thus separates further from the spectrum with physical content of interest as resolution increases. Hence high resolution simulations may use γ boost near γ wake and achieve maximum speedup even for runs using a tapered plasma. Simulations have been conducted at up to three times the base resolution, where use of $\gamma = \gamma_w$ is possible even including the above effects. The boosted frame speedup allows such high resolutions which may be important for evaluation of future low emittance stage concepts.

Appendix B: Input and output to and from boosted frame simulations in Warp

This section describes the procedures that have been implemented in the Particle-In-Cell framework Warp[37] to handle the input and output of data between the frame of calculation and the laboratory frame. Simultaneity of events between two frames is valid only for a plane that is perpendicular to the relative motion of the frame. As a result, the input/output processes involve the input of data (particles or fields) through a plane, as well as output through a series of planes, all of which are perpendicular to the direction of the relative velocity between the frame of calculation and the other frame of choice.

1. Input

a. Particles

Particles are launched through a plane using a technique which applies to many calculations in a boosted frame, including LPA, and is illustrated using the case of a positively charged particle beam propagating through a background of cold electrons in an assumed continuous transverse focusing system, leading to a growing transverse instability[23]. In the laboratory frame, the electron background is initially at rest and a moving window is used to follow the beam progression. Traditionally, the beam macroparticles are initialized all at once in the window, while background electron macroparticles are created continuously in front of the beam on a plane that is perpendicular to the beam velocity. In a frame moving at some fraction of the beam velocity in the laboratory frame, the beam initial conditions at a given time in the calculation frame are generally unknown and one must initialize the beam differently. However, it can be taken advantage of that the beam initial conditions are often known for a given plane in the laboratory, either directly, or via simple calculation or projection from the conditions at a given time. Given the position and velocity $\{x, y, z, v_x, v_y, v_z\}$ for each beam macroparticle at time $t = 0$ for a beam moving at the average velocity $v_b = \beta_b c$ (where c is the speed of light) in the laboratory, and using the standard synchronization ($z = z' = 0$ at $t = t' = 0$) between the laboratory and the calculation frames, the procedure for transforming the beam quantities for injection in a boosted frame moving at velocity βc in the laboratory is as follows (the superscript $'$ relates to quantities known in the boosted frame while the superscript * relates to quantities that are know at a given longitudinal position z^* but different times of arrival):

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

$$t^* = (z - \bar{z})/v_z \tag{B1}$$

$$x^* = x - v_x t^* \tag{B2}$$

$$y^* = y - v_y t^* \tag{B3}$$

$$z^* = 0 \tag{B4}$$

the velocity components being left unchanged,

2. apply Lorentz transformation from laboratory frame to boosted frame

$$t'^* = -\gamma t^* \tag{B5}$$

$$x'^* = x^* \tag{B6}$$

$$y'^* = y^* \tag{B7}$$

$$z'^* = \gamma\beta c t^* \tag{B8}$$

$$v'^*_x = \frac{v^*_x}{\gamma\left(1-\beta\beta_b\right)} \tag{B9}$$

$$v'^*_y = \frac{v^*_y}{\gamma\left(1-\beta\beta_b\right)} \tag{B10}$$

$$v'^*_z = \frac{v^*_z - \beta c}{1-\beta\beta_b} \tag{B11}$$

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

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

$$z' = z'^* - \bar{v}'^*_z t'^* \tag{B12}$$

This additional step is needed for setting the electrostatic or electromagnetic fields at the plane of injection. In a Particle-In-Cell code, the three-dimensional fields are calculated by solving the Maxwell equations (or static approximation like Poisson, Darwin or other[24]) on a grid on which the source term is obtained from the macroparticles distribution. This requires generation of a three-dimensional representation of the beam distribution of macroparticles at a given time before they cross the injection plane at z'^*. This is accomplished by expanding the beam distribution longitudinally such that all macroparticles (so far known at different times of arrival at the injection plane) are synchronized to the same time in the boosted frame. To keep the beam shape constant, the particles are "frozen" until they cross that plane: the three velocity components and the two position components perpendicular to the boosted frame velocity are kept constant, while the remaining position component is advanced at the average beam velocity. As particles cross the plane of injection, they become regular "active" particles with full 6-D dynamics.

Figure 16 (top) shows a snapshot of a beam that has passed partly through the injection plane. As the frozen beam macroparticles pass through the injection plane (which moves opposite to the beam in the boosted frame), they are converted to "active" macroparticles. The charge or current density is accumulated from the active and the frozen particles, thus ensuring that the fields at the plane of injection are consistent.

b. Laser

Similarly to the particle beam, the laser is injected through a plane perpendicular to the axis of propagation of the laser (by default z). The electric field $E_\perp$ that is to be emitted is given by the formula

$$E_\perp(x,y,t) = E_0 f(x,y,t)\sin\left[\omega t + \phi(x,y,\omega)\right] \tag{B13}$$

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

$$f(x,y,t) = e^{-\left(x^2/2\sigma_x^2 + y^2/2\sigma_y^2 + c^2t^2/2\sigma_z^2\right)} \tag{B14}$$

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

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

$$\begin{aligned} E_\perp(x,y,t) = &\left(1-\beta_s\right) E_0 f(x,y,t) \\ &\times \sin\left[\left(1-\beta_s\right)\omega t + \phi(x,y,\omega)\right]. \end{aligned} \tag{B15}$$

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

$$\begin{aligned} E_{\perp}\left(x',y',t'\right) &= \left(1-\beta_s\right) E'_0 f\left(x',y',t'\right) \\ &\times \sin\left[\left(1-\beta_s\right)\omega' t' + \phi\left(x',y',\omega'\right)\right] \qquad \text{(B16)} \\ &= \left(E_0/\gamma\right) f\left(x',y',t'\right) \\ &\times \sin\left[\omega t'/\gamma + \phi\left(x',y',\omega'\right)\right] \qquad \text{(B17)} \end{aligned}$$

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

The electric field is then converted into currents that get injected via a 2D array of macro-particles, with one positive and one dual negative macro-particle for each array cell in the plane of injection, whose weights and motion are governed by $E_{\perp}\left(x',y',t'\right)$. Injecting using this dual array of macroparticles offers the advantages of automatically including the longitudinal component which arise from emitting into a boosted frame, and to automatically verify the discrete Gauss' law thanks to using the Esirkepov current deposition scheme[51].

As discussed in section III, the technique implemented in Warp presents several advantages over other procedures that have been proposed elsewhere[11,32]. The method presented here avoids the caveat of the broadening of the transverse size of the laser while retaining simplicity and versatility by injecting through one plane rather than several faces of the box.

2. Output

Some quantities, e.g. charge or dimensions perpendicular to the boost velocity, are Lorentz invariant. Those quantities are thus readily available from standard diagnostics in the boosted frame calculations. Quantities which do not fall in this category are recorded at a number of regularly spaced "stations", immobile in the laboratory frame, at a succession of time intervals to record data history, or averaged over time. A visual example is given on Fig. 16 (bottom). Since the space-time locations of the diagnostic grids in the laboratory frame generally do not coincide with the space-time positions of the macroparticles and grid nodes used for the calculation in a boosted frame, some interpolation is performed at runtime during the data collection process. As a complement or an alternative, selected particle or field quantities are dumped at regular interval for post-processing. The choice of the methods depends on the requirements of the diagnostics and particular implementations.

REFERENCES

[1]T. TAJIMA and J. DAWSON, PHYSICAL REVIEW LETTERS **43**, 267 (1979), ISSN 0031-9007.

[2]E. Esarey, C. B. Schroeder, and W. P. Leemans, Rev. Mod. Phys. **81**, 1229 (2009), ISSN 0034-6861.

[3]W. P. Leemans, B. Nagler, A. J. Gonsalves, C. Toth, K. Nakamura, C. G. R. Geddes, E. Esarey, C. B. Schroeder, and S. M. Hooker, NATURE PHYSICS **2**, 696 (2006), ISSN 1745-2473.

[4]C. B. Schroeder, E. Esarey, C. G. R. Geddes, C. Benedetti, and W. P. Leemans, Phys. Rev. Spec. Topics Accel. Beams. **13**, 101301 (2010).

[5]W. Leemans and E. Esarey, PHYSICS TODAY **62**, 44 (2009), ISSN 0031-9228.

[6]F. Tsung, W. Lu, M. Tzoufras, W. Mori, C. Joshi, J. Vieira, L. Silva, and R. Fonseca, PHYSICS OF PLASMAS **13**, 056708 (2006), ISSN 1070-664X.

[7]C. Geddes, D. Bruhwiler, J. Cary, W. Mori, J.-L. Vay, S. Martins, T. Katsouleas, E. Cormier-Michel, W. Fawley, C. Huang, et al., in *Journal of Physics: Conference Series* (2008), vol. 125, p. 012002 (11 pp.), ISSN 1742-6596.

[8]C. Geddes, *et al.*, in *SciDAC Review 13* (2009), p. 13.

[9]C. Huang, W. An, V. Decyk, W. Lu, W. Mori, F. Tsung, M. Tzoufras, S. Morshed, T. Antonsen, B. Feng, et al., Journal of Physics: Conference Series **180**, 012005 (11 pp.) (2009), ISSN 1742-6596.

[10]W. P. Leemans, R. Duarte, E. Esarey, S. Fournier, C. G. R. Geddes, D. Lockhart, C. B. Schroeder, C. Toth, J.-L. Vay, and S. Zimmermann, AIP Conference Proceedings **1299**, 3 (2010).

[11]D. Bruhwiler, J. Cary, B. Cowan, K. Paul, C. Geddes, P. Mullowney, P. Messmer, E. Esarey, E. Cormier-Michel, W. Leemans, et al., in *AIP Conference Proceedings* (2009), vol. 1086, pp. 29–37, ISSN 0094-243X.

[12]J. KRALL, A. TING, E. ESAREY, and P. SPRANGLE, Physical Review E **48**, 2157 (1993).

[13]B. A. Shadwick, C. B. Schroeder, and E. Esarey, PHYSICS OF PLASMAS **16**, 056704 (2009), ISSN 1070-664X.

[14] C. Benedetti, C. B. Schroeder, E. Esarey, C. G. R. Geddes, and W. P. Leemans, AIP Conference Proceedings **1299**, 250 (2010).

[15] P. SPRANGLE, E. ESAREY, and A. TING, PHYSICAL REVIEW LETTERS **64**, 2011 (1990), ISSN 0031-9007.

[16] T. M. ANTONSEN and P. MORA, Physical Review Letters **69**, 2204 (1992).

[17] C. Huang, V. K. Decyk, C. Ren, M. Zhou, W. Lu, W. B. Mori, J. H. Cooley, T. M. Antonsen, Jr., and T. Katsouleas, JOURNAL OF COMPUTATIONAL PHYSICS **217**, 658 (2006), ISSN 0021-9991.

[18] B. Feng, C. Huang, V. Decyk, W. Mori, P. Muggli, and T. Katsouleas, Journal of Computational Physics **228**, 5340 (2009), ISSN 0021-9991.

[19] B. M. Cowan, D. L. Bruhwiler, E. Cormier-Michel, E. Esarey, C. G. Geddes, P. Messmer, and K. M. Paul, Journal of Computational Physics **230**, 61 (2011), ISSN 0021-9991.

[20] E. Cormier-Michel, C. Geddes, E. Esarey, C. Schroeder, D. Bruhwiler, K. Paul, B. Cowan, and W. Leemans, in *AIP Conference Proceedings* (2009), vol. 1086, pp. 297–302, ISSN 0094-243X.

[21] C. G. R. Geddes, *et al.*, in *Proc. Particle Accelerator Conference* (Vancouver, Canada, 2009), wE6RFP075.

[22] S. V. BULANOV, I. N. INOVENKOV, V. I. KIRSANOV, N. M. NAUMOVA, and A. S. SAKHAROV, Physics of Fluids B-plasma Physics **4**, 1935 (1992).

[23] J.-L. Vay, Physical Review Letters **98**, 130405/1 (2007), ISSN 0031-9007.

[24] J. L. Vay, Physics of Plasmas **15**, 056701 (2008).

[25] J. BORIS and R. LEE, JOURNAL OF COMPUTATIONAL PHYSICS **12**, 131 (1973), ISSN 0021-9991.

[26] I. Haber, R. Lee, H. Klein, and J. Boris, in *Proc. Sixth Conf. Num. Sim. Plasmas* (Berkeley, CA, 1973), pp. 46–48.

[27] B. Cowan, D. Bruhwiler, E. Cormier-Michel, E. Esarey, C. Geddes, P. Messmer, and K. Paul, in *AIP Conference Proceedings* (2009), vol. 1086, pp. 309–14, ISSN 0094-243X.

[28] J.-L. Vay, *et al.*, in *Proc. Particle Accelerator Conference* (Vancouver, Canada, 2009), tU1PBI04.

[29] S. F. Martins, *et al.*, in *Proc. Particle Accelerator Conference* (Vancouver, Canada, 2009), tH4GBC05.

[30] J.-L. Vay, D. Bruhwiler, C. Geddes, W. Fawley, S. Martins, J. Cary, E. Cormier-Michel, B. Cowan, R. Fonseca, M. Furman, et al., Journal of Physics: Conference Series **180**, 012006 (5 pp.) (2009), ISSN 1742-6596.

[31] J.-L. Vay, *et al.*, arXiv:0909.5603 (2009).

[32] S. F. Martins, R. A. Fonseca, L. O. Silva, W. Lu, and W. B. Mori, COMPUTER PHYSICS COMMUNICATIONS **181**, 869 (2010), ISSN 0010-4655.

[33] S. F. Martins, R. A. Fonseca, W. Lu, W. B. Mori, and L. O. Silva, NATURE PHYSICS **6**, 311 (2010), ISSN 1745-2473.

[34] S. F. Martins, R. A. Fonseca, J. Vieira, L. O. Silva, W. Lu, and W. B. Mori, PHYSICS OF PLASMAS **17**, 056705 (2010), ISSN 1070-664X.

[35] D. L. Bruhwiler (2008), private Communication.

[36] J. L. Vay, C. G. R. Geddes, E. Cormier-Michel, and D. P. Grote, Journal of Computational Physics **230**, 5908 (2011).

[37] D. Grote, A. Friedman, J.-L. Vay, and I. Haber, in *AIP Conference Proceedings* (2005), 749, pp. 55–8, ISSN 0094-243X.

[38] J. Vay, C. G. R. Geddes, E. Cormier-Michel, and D. P. Grote, Physics of Plasmas **18**, 030701 (2011).

[39] C. B. Schroeder, C. Benedetti, E. Esarey, and W. P. Leemans, Physical Review Letters **106**, 135002 (2011).

[40] R. A. Fonseca, *et al.*, Lec. Notes in Comp. Sci. **2329**, 342 (2002).

[41] C. Nieter and J. R. Cary, J. Comput. Phys. **196**, 448 (2004).

[42] Taflove and Hagness, *Computational Electrodynamics: The Finite-Difference Time-Domain Method* (MA: Artech House, Norwood, 2000), 2nd ed.

[43] C. Benedetti, *et al.*, Nuclear Inst. and Methods in Physics Research A **608**, 94 (2009).

[44] C. G. R. Geddes, *Plasma Channel Guided Laser Wakefield Accelerator* (Ph.D. dissertation, University of California, Berkeley, Berkeley CA, 2005).

[45] S. Kalmykov, *et al.*, Phys. Rev. Lett. **103**, 135004 (2009).

[46] C. Birdsall and A. Langdon, *Plasma physics via computer simulation* (Adam-Hilger, 1991), ISBN 0 07 005371 5.

[47] C. B. Schroeder, E. Esarey, B. A. Shadwick, and W. P. Leemans, Physics of Plasmas **13**, 033103 (2006).

[48] W. Rittershofer, C. B. Schroeder, E. Esarey, F. J. Gruner, and W. P. Leemans, Physics of Plasmas **17**, 063104 (2010).

[49] E. Cormier-Michel, E. Esarey, C. G. R. Geddes, C. B. Schroeder, K. Paul, P. J. Mullowney, J. R. Cary, and W. P. Leemans, Physical Review Special Topics-accelerators and Beams **14**, 031303 (2011).

[50] M. Tzoufras, W. Lu, F. S. Tsung, C. Huang, W. B. Mori, T. Katsouleas, J. Vieira, R. A. Fonseca, and L. O. Silva, Physical Review Letters **101**, 145002 (2008).

[51] T. Esirkepov, COMPUTER PHYSICS COMMUNICATIONS **135**, 144 (2001), ISSN 0010-4655.

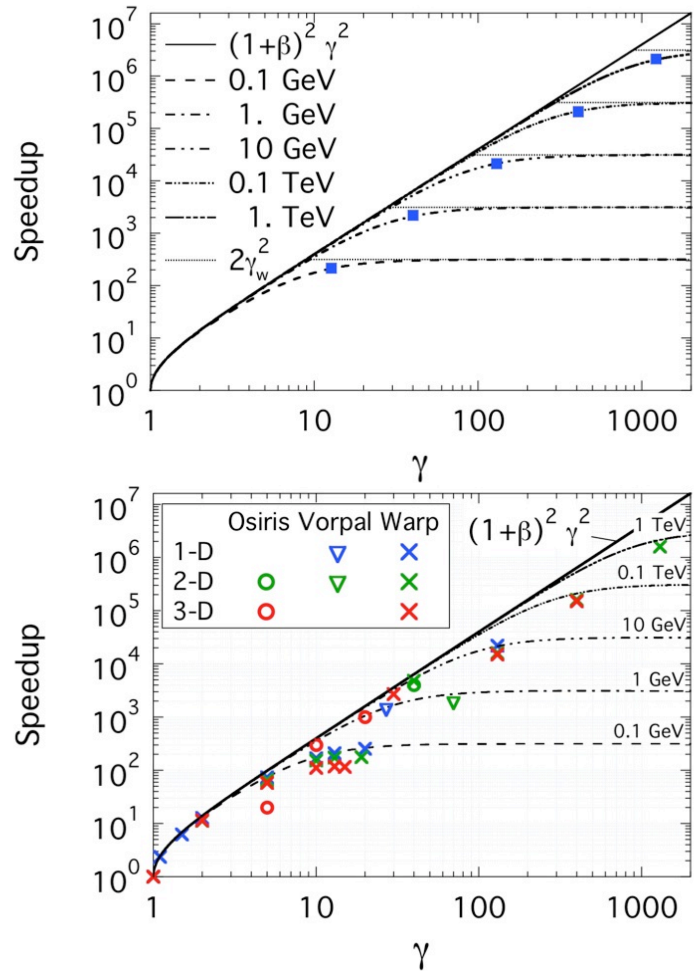

FIG. 1. (top) Theoretical speedup from Eq.(13) versus relativistic factor of the boosted frame for 0.1 GeV - 1 TeV LPA class stages (squares indicate speedup obtained using the frame of the wake $\gamma = \gamma_w$); (bottom) observed speedups from simulations using the code Osiris (circles), Vorpal (triangles) and Warp (crosses) and theoretical speedups (lines) for 0.1 GeV to 1 TeV stages. Vorpal reported speedups courtesy of D. L. Bruhwiler, Tech-X Corp., USA. Osiris reported speedups courtesy of S. F. Martins, IST, Portugal, and W. B. Mori, UCLA, USA.

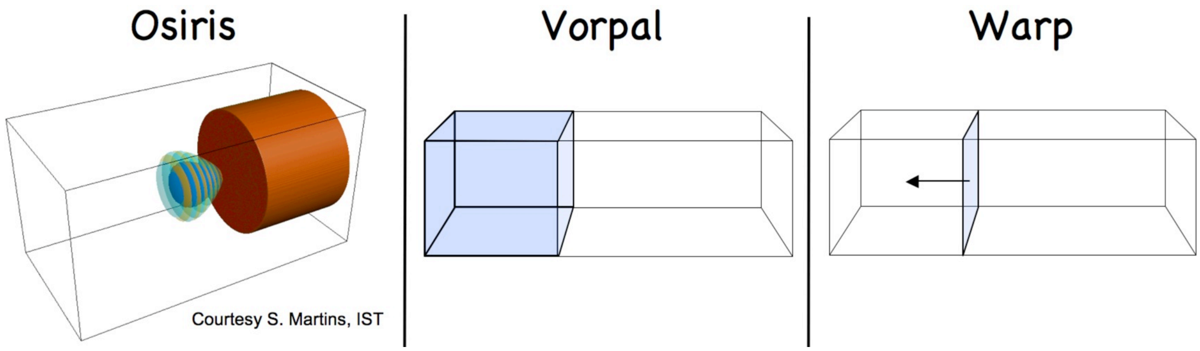

FIG. 2. (color online) Diagrams of laser emission procedures in the Particle-In-Cell codes Osiris (left), Vorpal (middle) and Warp (right) for Lorentz boosted frame simulations. Osiris initializes the entire laser at once. Vorpal emits the laser from all but one faces (blue) of the simulation box. Warp emits through a moving plane (blue). For all three diagrams, the laser propagates from left to right. Reprinted with permission from J.-L. Vay et al., *Proc. 14th Workshop Advanced Accelerator Concepts*. Copyright 2010 American Institute of Physics.

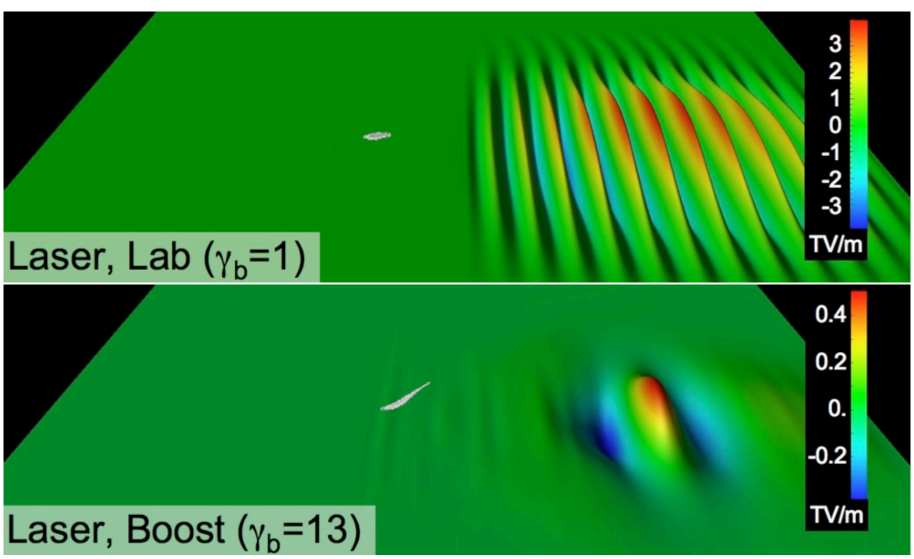

FIG. 3. (color online) Colored surface rendering of the transverse electric field from a 2-1/2D Warp simulation of a laser wakefield acceleration stage in the laboratory frame (top) and a boosted frame at $\gamma = 13$ (bottom), with the beam (white) in its early phase of acceleration. The laser and the beam are propagating from left to right.

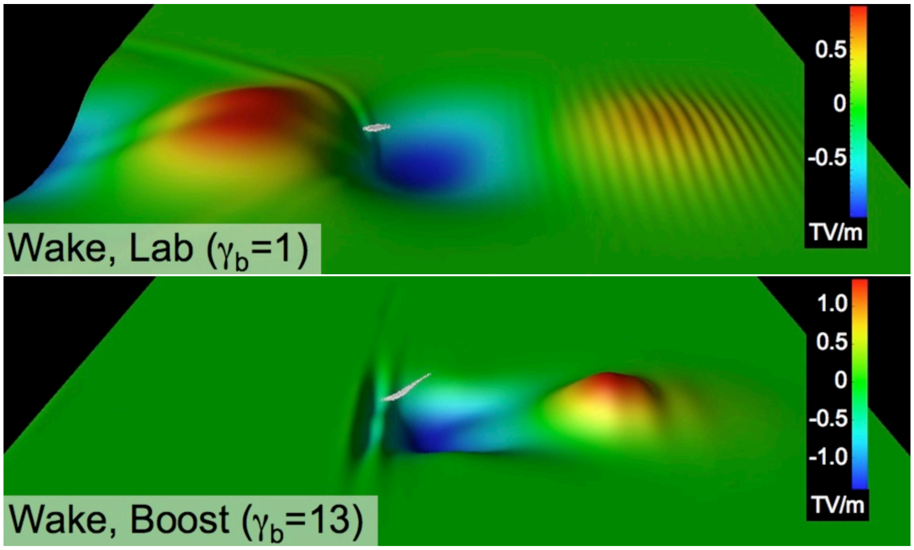

FIG. 4. (color online) Colored surface rendering of the longitudinal electric field from a 2-1/2D Warp simulation of a laser wakefield acceleration stage in the laboratory frame (top) and a boosted frame at $\gamma = 13$ (bottom), with the beam (white) in its early phase of acceleration. The laser and the beam are propagating from left to right.

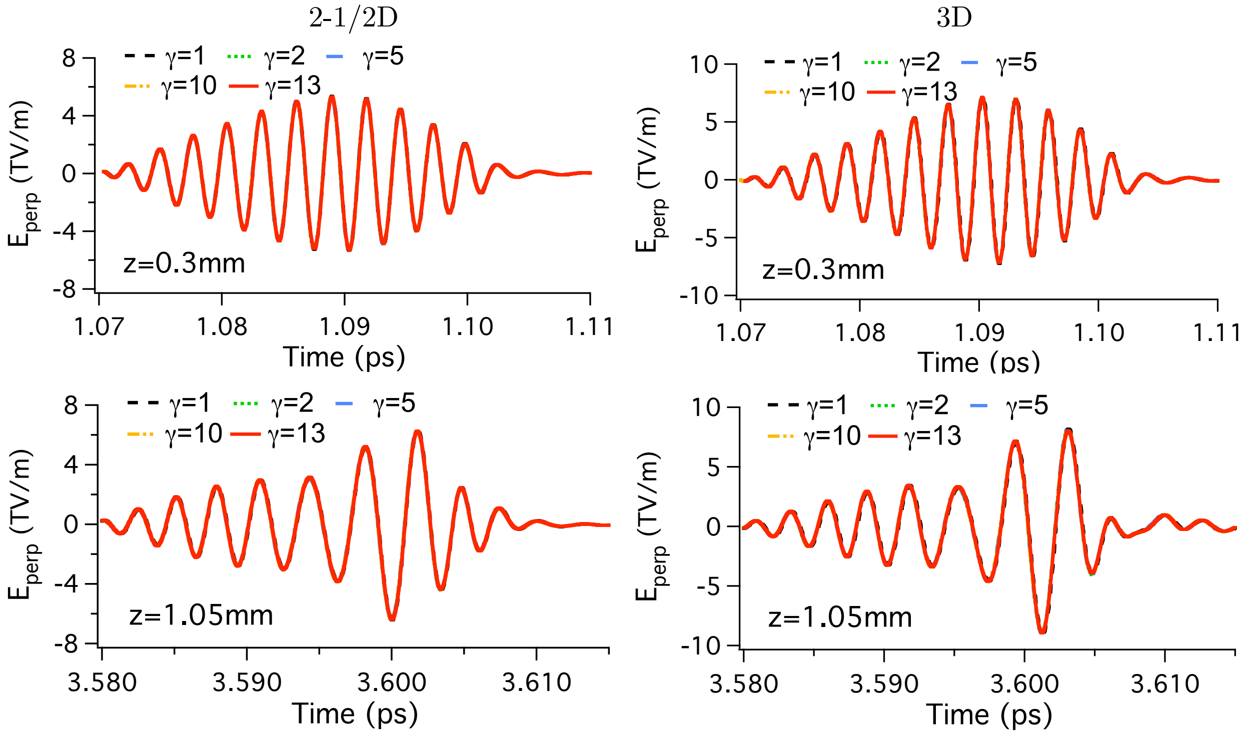

FIG. 5. (color online) History of transverse electric field at the position $x = y = 0$, $z = 0.3$ mm and $z = 1.05$ mm (in the laboratory frame) from simulations in the laboratory frame ($\gamma = 1$) and boosted frames at $\gamma = 2$, 5, 10 and 13.

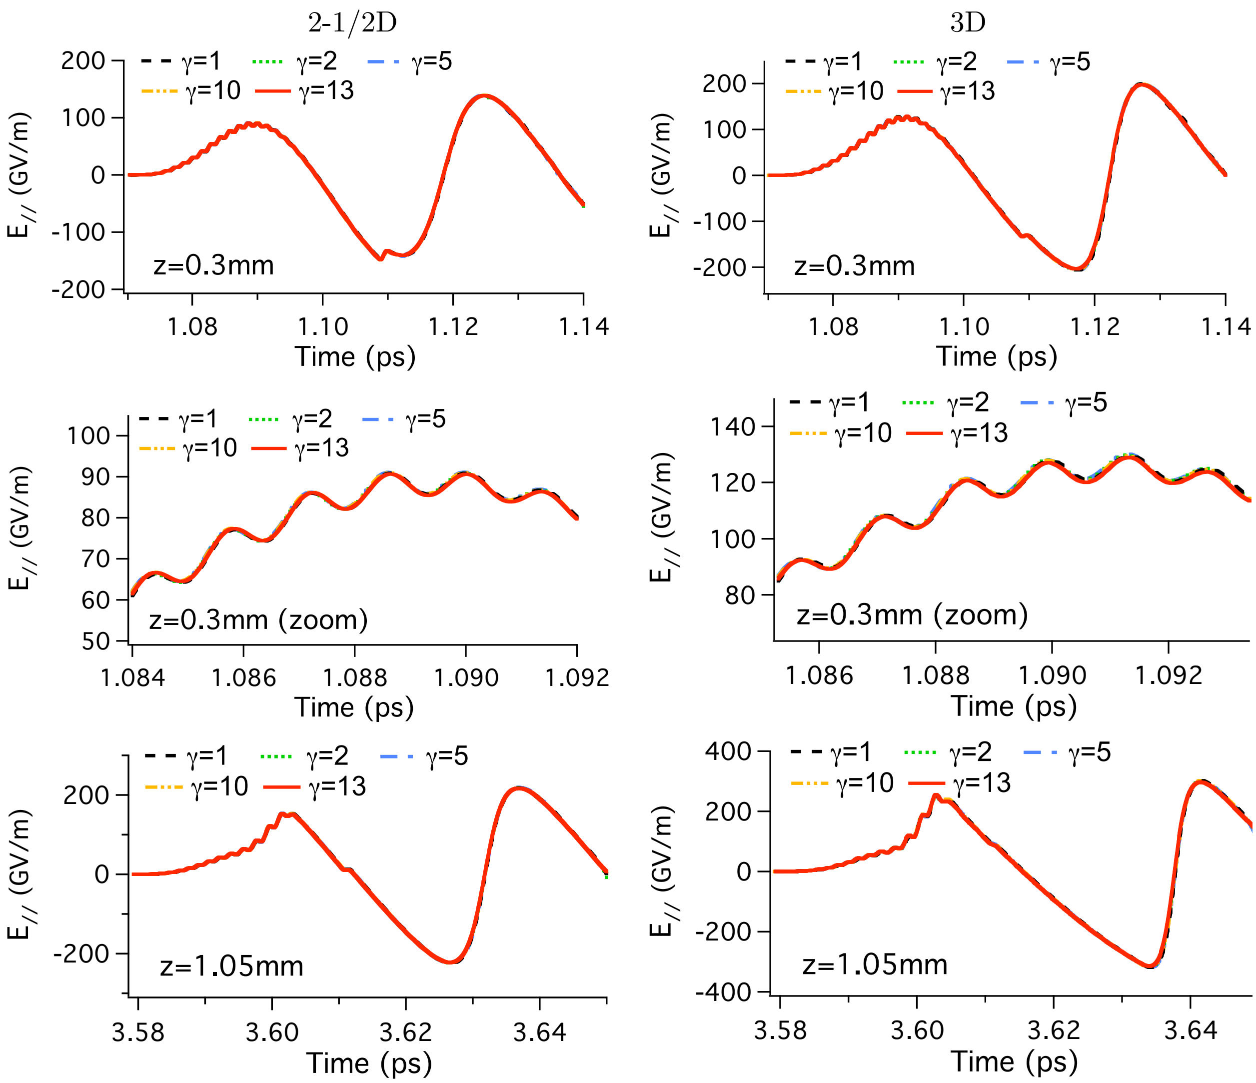

FIG. 6. (color online) History of longitudinal electric field at the position $x = y = 0$, $z = 0.3$ mm and $z = 1.05$ mm (in the laboratory frame) from simulations in the laboratory frame ($\gamma = 1$) and boosted frames at $\gamma = 2$, 5, 10 and 13.

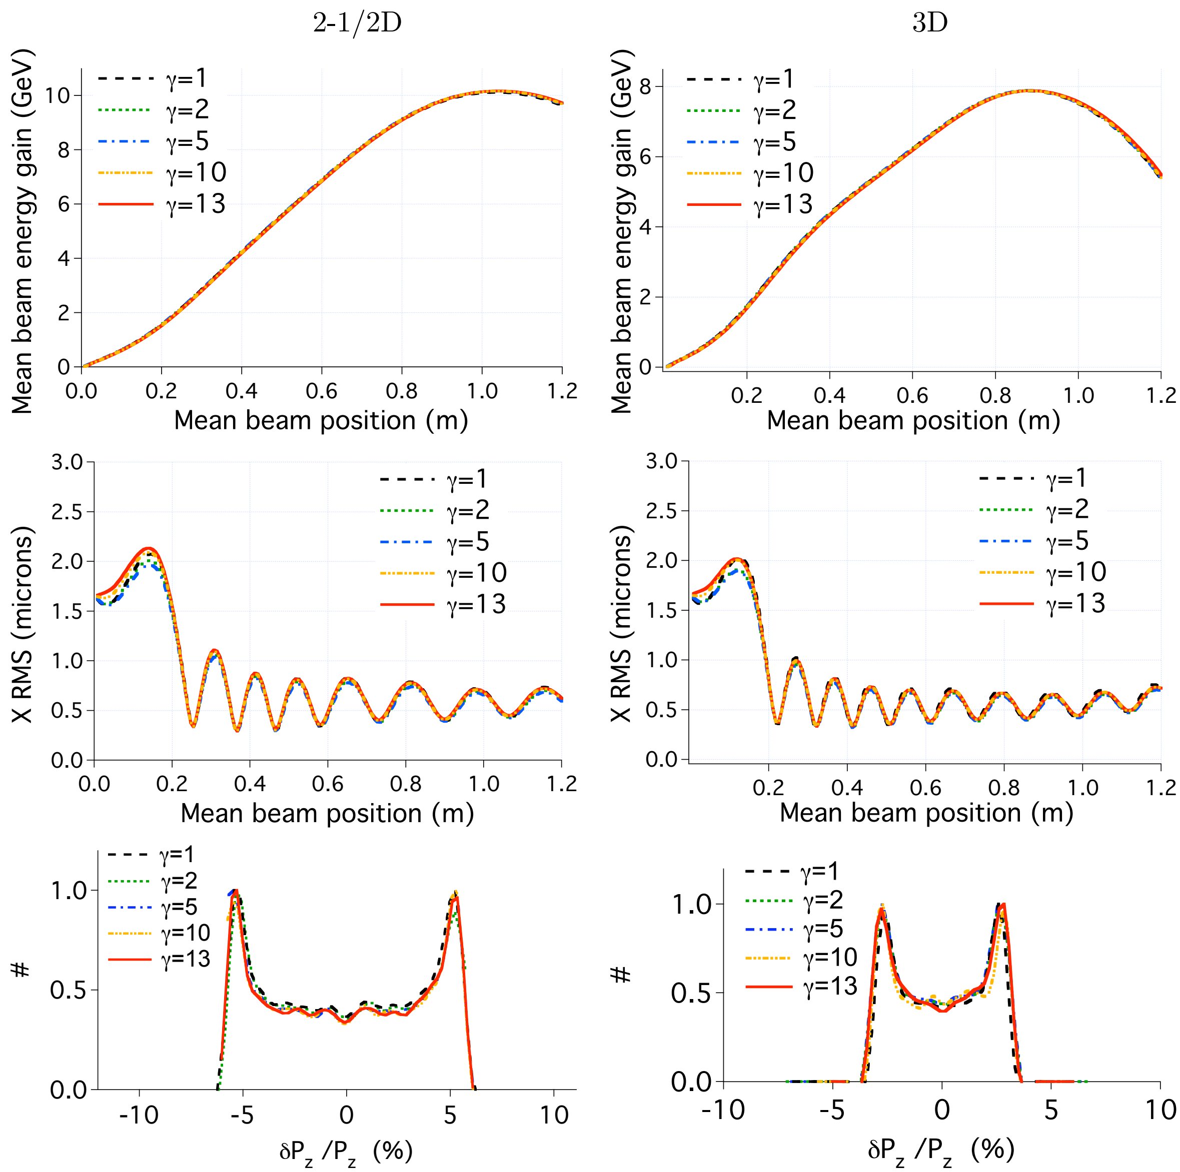

FIG. 7. (color online) Average scaled beam energy gain (top) and beam RMS transverse size (middle) versus longitudinal position in the laboratory frame from simulations; (bottom) distribution of relative longitudinal momentum dispersion at peak energy, in the laboratory frame ($\gamma = 1$) and boosted frames at $\gamma = 2$, 5, 10 and 13.

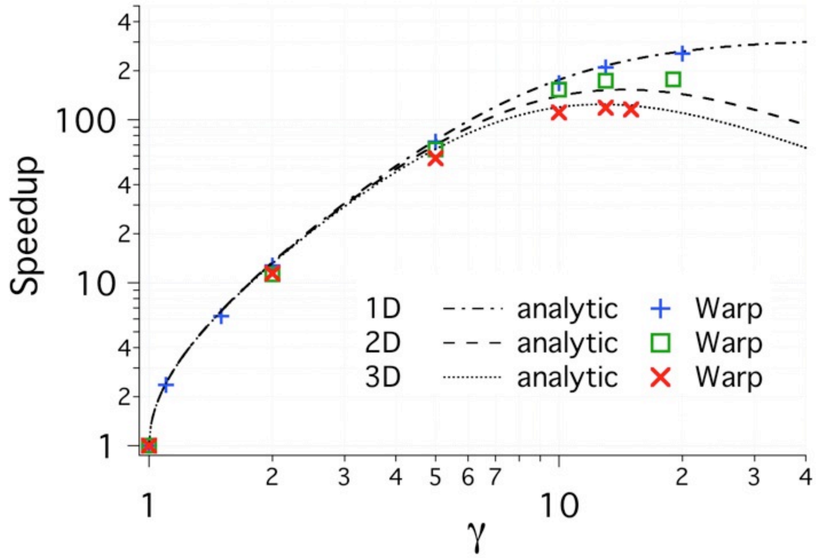

FIG. 8. (color online) Speedup versus relativistic factor of the boosted frame in 1D, 2D and 3D from theoretical estimates (Eq. (13), (17), (18)), and Warp simulations.

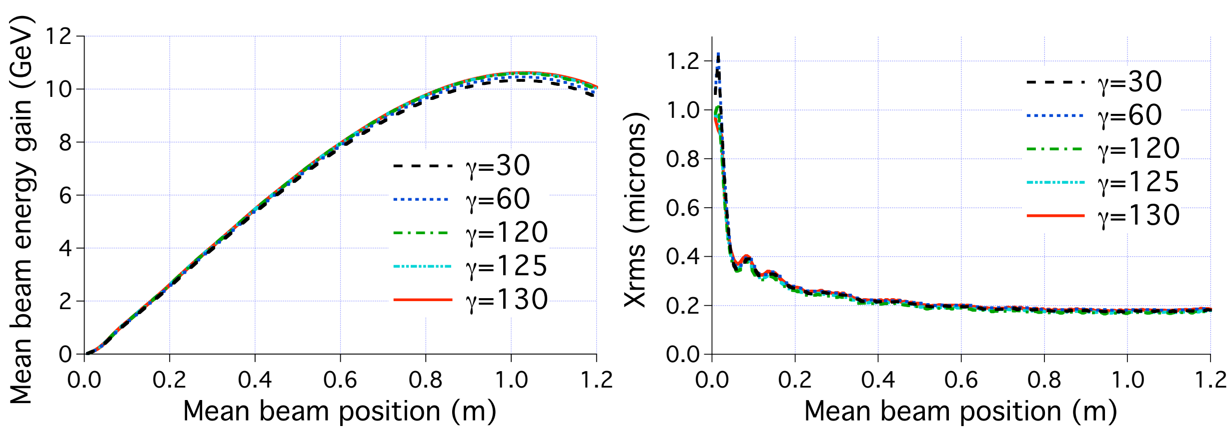

FIG. 9. (color online) Average beam energy gain and transverse size versus longitudinal position (in the laboratory frame) from 2D-1/2 simulations of a full scale 10 GeV LPA in a boosted frame at $\gamma = 30$, 60 and 130, using the Yee solver.

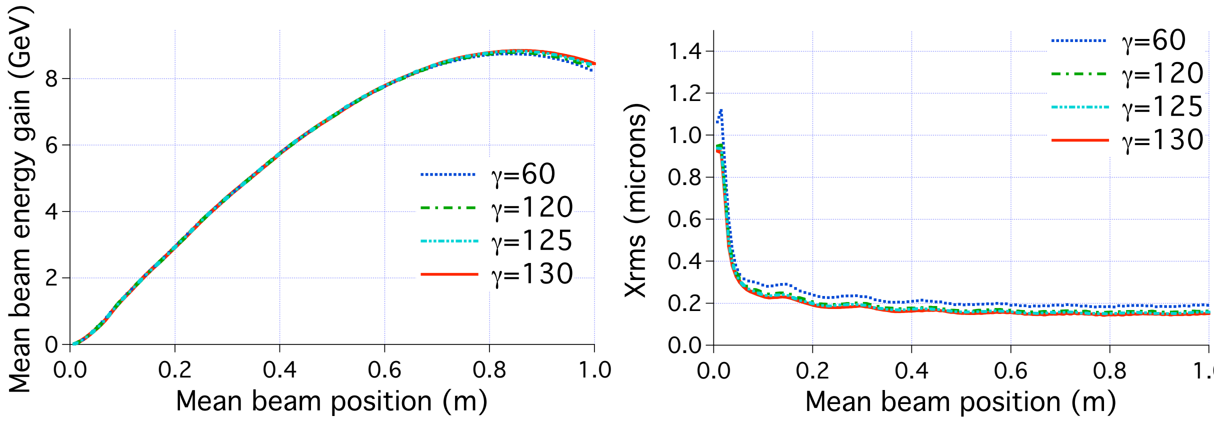

FIG. 10. (color online) Average beam energy gain and transverse size versus longitudinal position (in the laboratory frame) from 3D simulations of a full scale 10 GeV LPA in a boosted frame at $\gamma = 30$, 60, 120, 125 and 130, using the Yee solver ($\gamma = 30$ and 60) and the CK2 solver ($\gamma = 120 - 130$), with digital filter S(1) and with the time step set by $c\delta t/\delta z = 1/\sqrt{2}$ for stability[36].

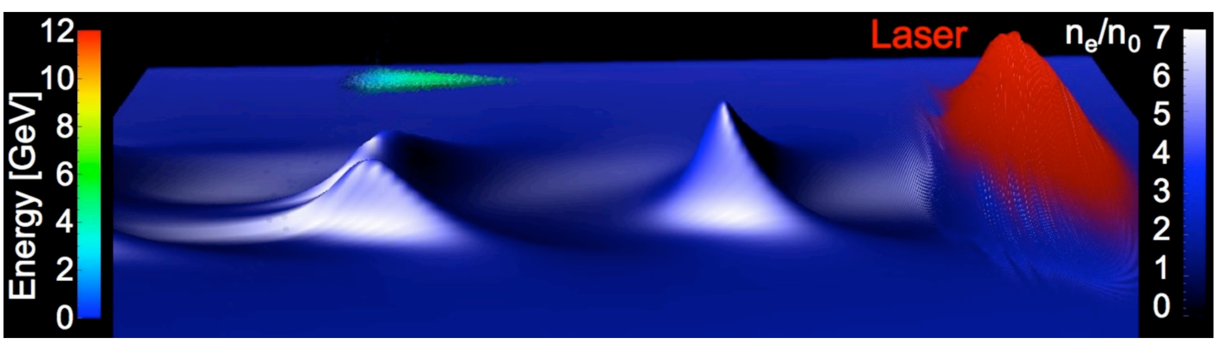

FIG. 11. (color online) Snapshot from a 2-1/2D 10 GeV LPA stage boosted frame simulation as the beam is halfway through acceleration. The image shows an externally injected electron bunch (middle) riding a density wake excited by an intense laser pulse (right), propagating in a 0.75 m long plasma channel.

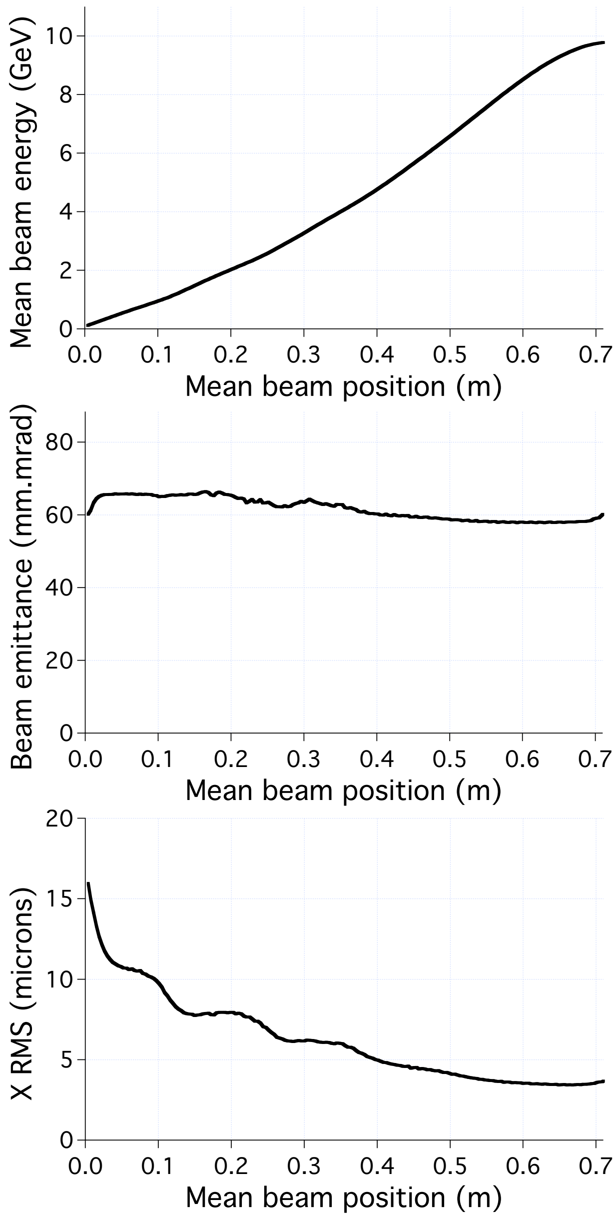

FIG. 12. (top) Average electron beam energy gain, (middle) beam emittance, and (bottom) beam RMS size, versus longitudinal position (in the laboratory frame) from a 2D 10 GeV LPA stage boosted frame simulation.

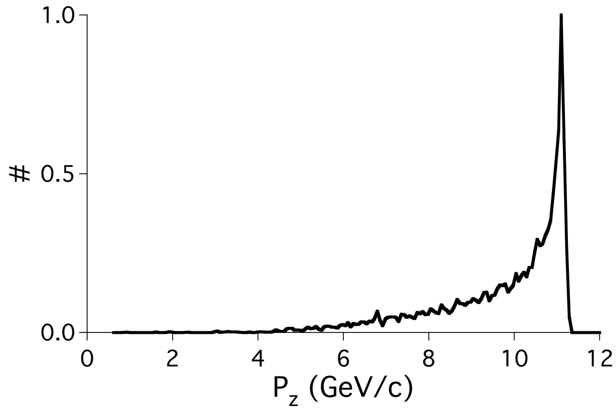

FIG. 13. Longitudinal momentum distribution of the electron beam at maximum energy (z=0.7 m) from a 2D 10 GeV LPA stage boosted frame simulation.

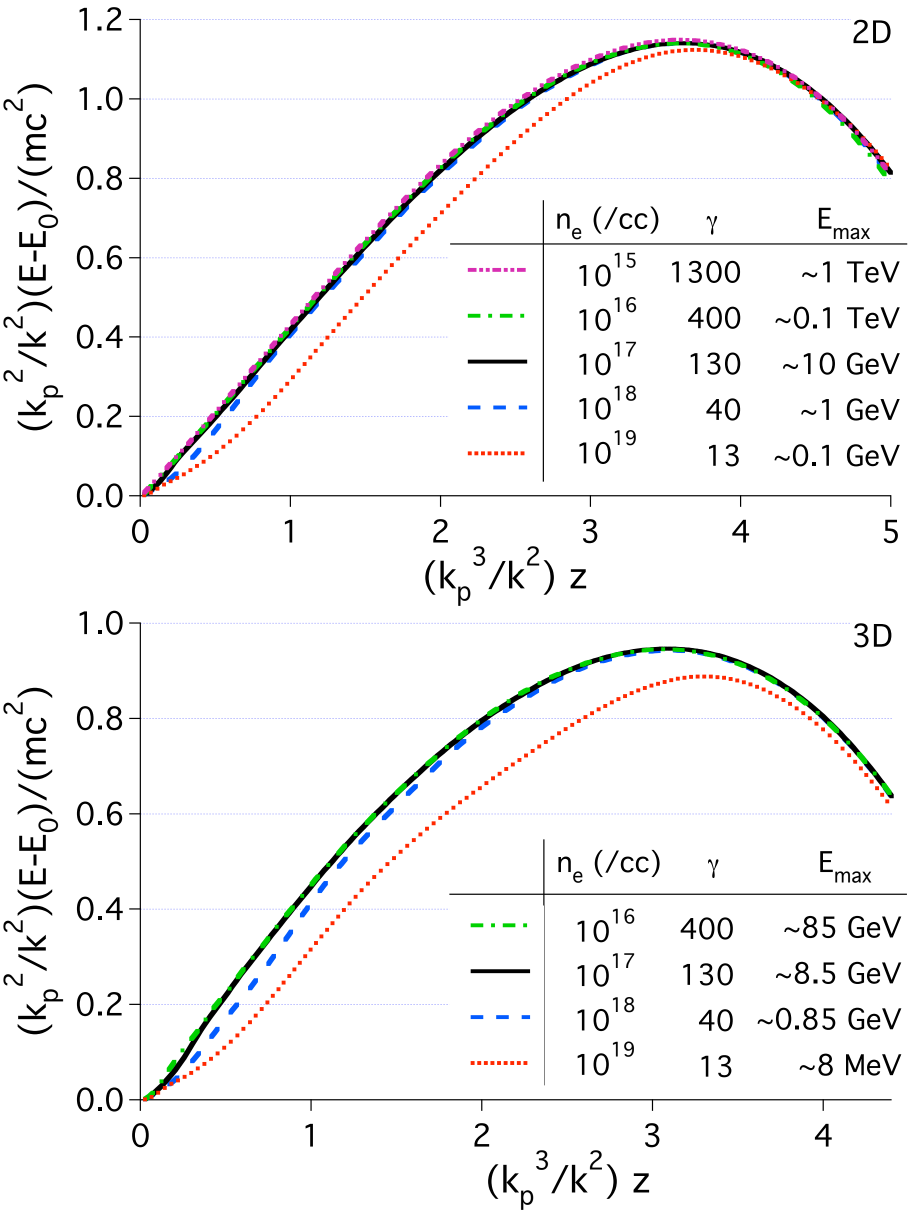

FIG. 14. (color online) Average beam energy gain versus longitudinal position (in the laboratory frame) for simulations at $n_e = 10^{19}$ cm^{-3} down to 10^{15} cm^{-3}, using frames of reference between $\gamma = 13$ and $\gamma = 1300$, in 2-1/2D (top) and 3D (bottom).

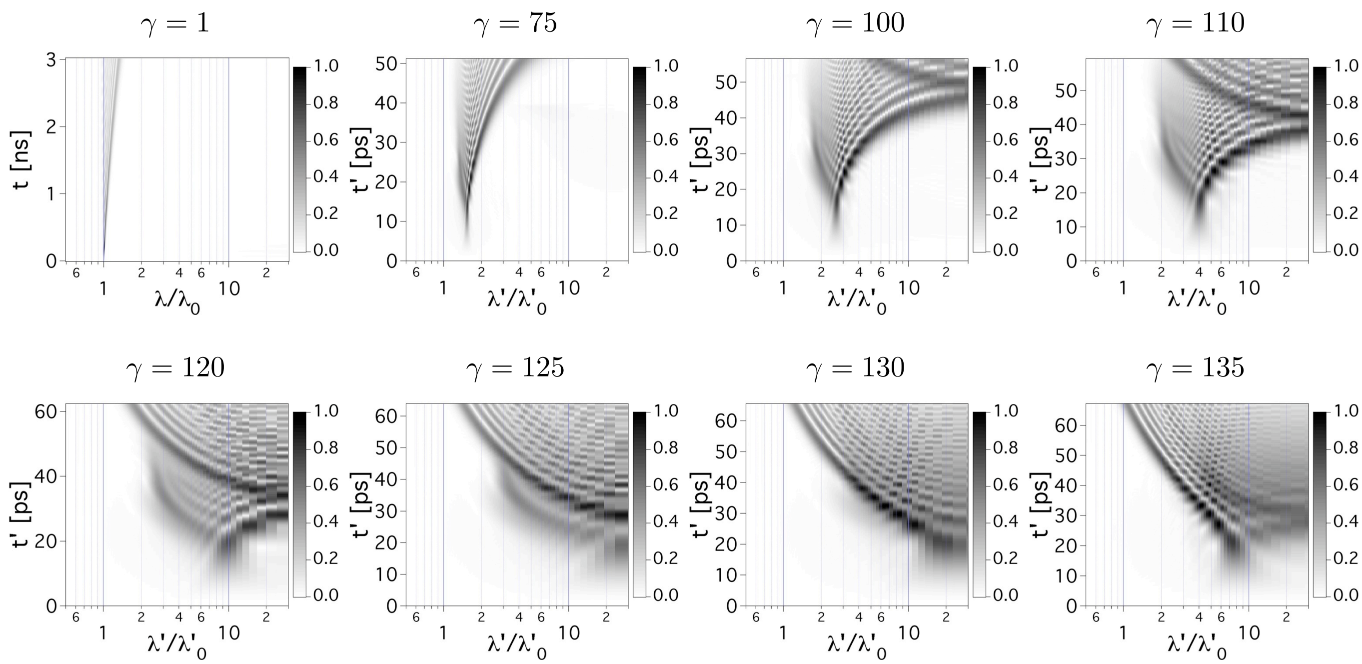

FIG. 15. Spectrum history of the laser field on axis of a 10 GeV stage for selected values of γ between 1 and 135, given up to the time of the electron beam peak energy. The length scale (horizontal axis) is normalized relative to the vacuum laser wavelength as given in each respective frame.

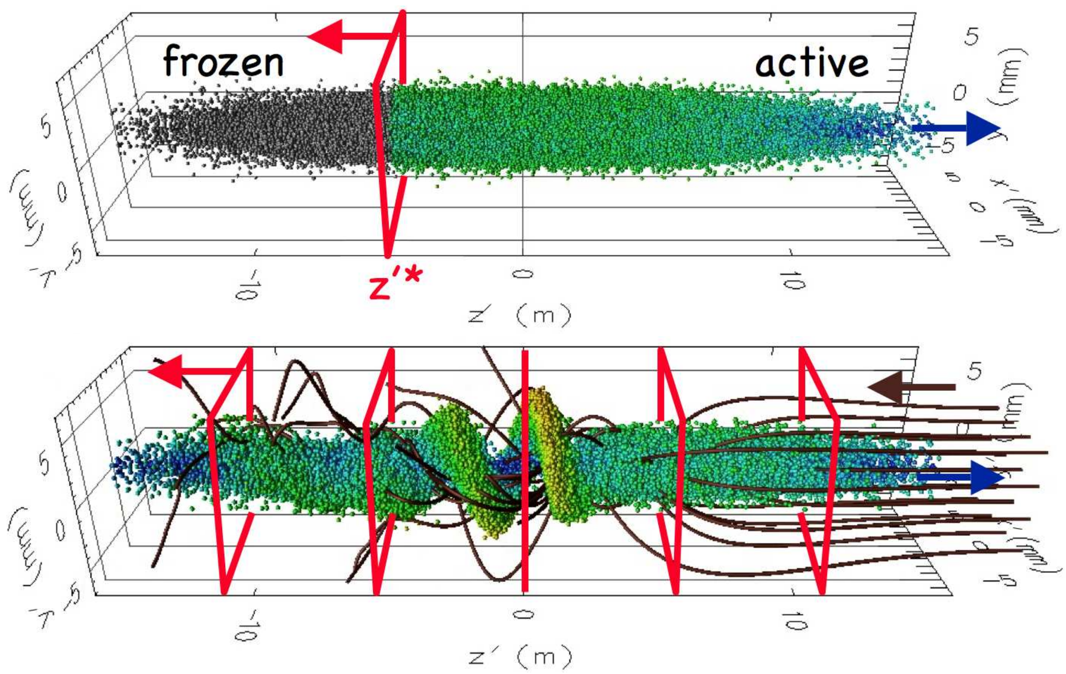

FIG. 16. (color online) (top) Snapshot of a particle beam showing “frozen” (grey spheres) and “active” (colored spheres) macroparticles traversing the injection plane (red rectangle). (bottom) Snapshot of the beam macroparticles (colored spheres) passing through the background of electrons (dark brown 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.