Modeling Laser Wakefield Accelerators in a Lorentz Boosted Frame

J.-L. Vay*, C. G. R. Geddes*, C. Benedetti*, D. L. Bruhwiler†, E. Cormier-Michel†, B. M. Cowan†, J. R. Cary† and D. P. Grote**

*Lawrence Berkeley National Laboratory, CA, USA
†Tech-X Corporation, CO, USA
**Lawrence Livermore National Laboratory, CA, USA

Abstract. Modeling of laser-plasma wakefield accelerators in an optimal frame of reference has been shown to produce up to three orders of magnitude speed-up in calculations from first principles of stages in the 100 MeV-10 GeV energy range. Maximum obtainable speedups calculated using linear theory predict that higher speedups are attainable, in the range of 4-6 orders of magnitude for stages in the energy range of 10 GeV-1 TeV respectively. Practical limitations have been reported and discussed which have prevented reaching these speedups so far, including a violent high frequency numerical instability. The limitations are briefly reviewed and discussed in this paper, as well as their mitigation. It is also reported that the high frequency numerical instability can be controlled effectively using novel numerical techniques that have been implemented in the Particle-In-Cell code Warp, and that 5 and 6 orders of magnitude speedups were demonstrated on 100 GeV and 1 TeV stages respectively, verifying the scaling of plasma accelerators to very high energies, and providing highly efficient tools for the detailed designs of experiments on new lasers such as BELLA.

Keywords: Laser Plasma Acceleration, Particle-In-Cell, Boosted Frame
PACS: 52.38.Kd, 41.75.Jv, 52.65.Rr

INTRODUCTION

Modeling of laser-plasma wakefield accelerators in an optimal frame of reference [1] has been shown to produce up-to three orders of magnitude speed-up in calculations from first principles of stages in the 100 MeV-10 GeV energy range [2-10]. Maximum obtainable speedups calculated using linear theory predict that higher speedups are attainable, in the range of 4-6 orders of magnitude for stages in the energy range of 10 GeV-1 TeV respectively. Practical limitations have been reported and discussed which have prevented reaching these speedups [2, 5, 8], including a decrease in the laser Rayleigh length which results in an increase in the laser spot size, a decrease in macro-particle statistics in the plasma column affecting the self-injection regime, and a violent high frequency numerical instability. The limitations are reviewed and discussed in this paper, as well as their mitigation. It is also reported that the high frequency numerical instability can be controlled effectively using novel numerical techniques that have been implemented in the Particle-In-Cell code Warp [11], and that 5 and 6 orders of magnitude speedups were demonstrated on 100 GeV and 1 TeV stages respectively.

THEORETICAL SCALING OF THE SPEEDUP

The theoretical scaling of the speedup of a Lorentz boosted frame simulation, in comparison to an equivalent simulation using the laboratory frame, is given in this section. It is assumed that the simulation box is a fixed number of plasma periods long, which implies the use (which is common) of a moving window following the wake and accelerated beam. The speedup is then 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 [1]. Assuming that (a) the wake propagates at the group velocity $v_w = \beta_w c$ of plane waves in a uniform plasma, (b) that simulations stop when the last electron beam macro-particle exits the plasma, and (c) that $\gamma_w = 1/\sqrt{1-\beta_w^2} \gg 1$, it is shown in [12] that the speedup is given by

CP1299, *Advanced Accelerator Concepts: 14th Workshop,* edited by S. H. Gold and G. S. Nusinovich
2010 American Institute of Physics 978-0-7354-0853-1/10/$30.00

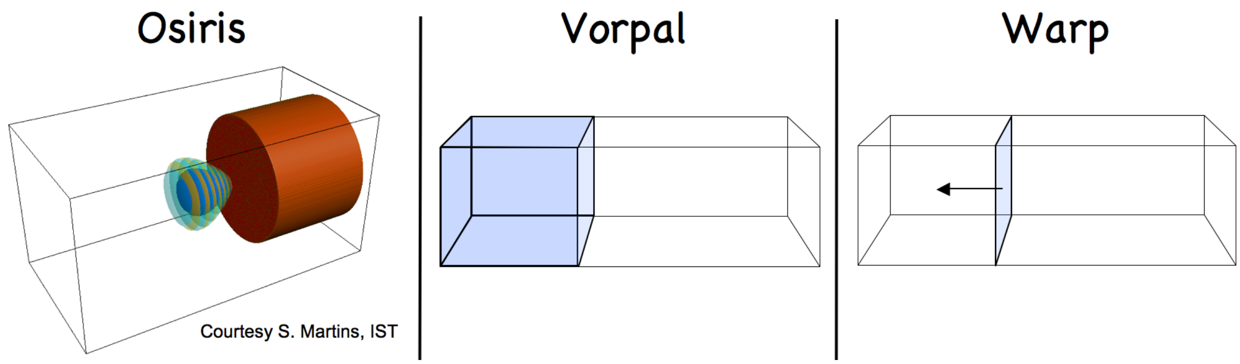

FIGURE 1. Diagrams of laser emission procedures in the Particle-In-Cell codes Osiris (left), Vorpal (middle) and Warp (right) for Lorentz boosted frame simulations. Osiris initializes the entire laser at once. Vorpal emits the laser from all but one faces (blue) of the simulation box. Warp emits through a moving plane (blue). For all three diagrams, the laser propagates from left to right.

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

where βc is the velocity of the boost, $\gamma = 1/\sqrt{1-\beta^2}$, ξ is the ratio between the plasma length and the dephasing length and η is the fraction of the wake which exited the plasma at the end of the simulation (see [12] for details).

For low values of γ, i.e. when $\gamma \ll \gamma_w$, Eq.(1) reduces to $S_{\gamma\ll\gamma_w} = (1+\beta)^2\gamma^2$. Conversely, if $\gamma \to \infty$, Eq.(1) becomes $S_{\gamma\to\infty} = \frac{4}{1+4\eta/\xi}\gamma_w^2$. Finally, in the frame of the wake, i.e. when $\gamma = \gamma_w$, assuming that $\beta_w \approx 1$, Eq.(1) gives $S_{\gamma=\gamma_w} \approx \frac{2}{1+2\eta/\xi}\gamma_w^2$. Since η and ξ are of order unity, and the practical regimes of most interest satisfy $\gamma_w^2 \gg 1$, the speedup that is obtained by using the frame of the wake will be near the maximum obtainable value given by $S_{\gamma\to\infty}$ and is over 4-6 orders of magnitude for stages in the 10 GeV-1 TeV energy range respectively.

NUMERICAL ISSUES IN PAST BOOSTED FRAME SIMULATIONS

Several numerical limits can restrict the boost performance. Here we review limits in past simulations and their impact on performance: laser initialization, statistics and a short wavelength instability.

As discussed in [2, 5, 8], 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. 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 $\gamma^2(1+\beta)$ [8]. 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. Diagrams of the laser emission procedures used for boosted frame simulations with the codes Osiris [13], Vorpal [14] and Warp [11] are given in Fig. 1. 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 [2] using total field/scattered field technique [15], while Warp emits via a moving planar antenna as described in [12].

As discussed in [8], 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. 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 in [8] 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

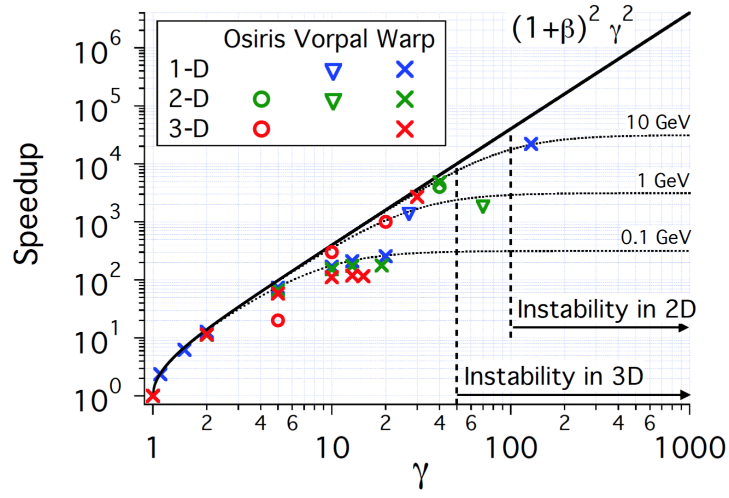

FIGURE 2. Theoretical speedups (lines) for 0.1 GeV, 1 GeV and 10 GeV stages and observed speedups from simulations using the code Osiris (circles), Vorpal (triangles) and Warp (crosses). Osiris reported speedups courtesy of S. F. Martins, IST, Portugal, and W. B. Mori, UCLA, USA.

density of macro-particles of smaller weights. This is already practiced in ordinary runs (i.e. without boosted-frame) for example in Ref. [16] for minimizing the computational cost while maximizing the statistics within "dynamically interesting" regions. For instance it is well known [17, 18] 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.

A violent numerical instability developing at the front of the plasma column for $\gamma \gtrsim 100$ in 2-D and $\gamma \gtrsim 50$ in 3-D was reported using all three codes [7, 8, 19]. The presence and growth rate of the instability was observed to be very sensitive to the resolution (slower growth rate at higher resolution), and to the amount of damping of high frequencies and smoothing of short wavelengths. The instability is always propagating at some angle from the longitudinal axis, and is observed in 2-D and 3-D runs but was never observed in any of the 1-D 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 [8] that the observed instability might be caused by numerical Cerenkov effect. The instability has been studied in details with Warp and mitigations have been developed using new algorithms, as described in a following section.

SUMMARY OF OBTAINED SPEEDUP WITHOUT MITIGATION OF INSTABILITY

The theoretical speedups from Eq. (1) for 0.1 GeV to 10 GeV stages are plotted in Fig. 2, as well as observed speedups from simulations using the Particle-In-Cell codes Osiris, Vorpal and Warp. All three codes were using the same standard Particle-In-Cell method [20] and were simulating laser plasma acceleration stages in the range of 0.1-10 GeV. They all successfully performed 2-D and/or 3-D calculations with boosts at gamma in the range of 20-70, reaching speedups over three orders of magnitude (projected for Osiris assuming no computational cost from laser injection). Only in 1-D could Warp successfully perform a simulation with a boost at $\gamma > 100$, approaching the theoretical maximum speedup of 3×10^4 speedup.

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, Vorpal simulations were of external injection stages with beam loading, and Warp simulations were of external injection stages without 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.

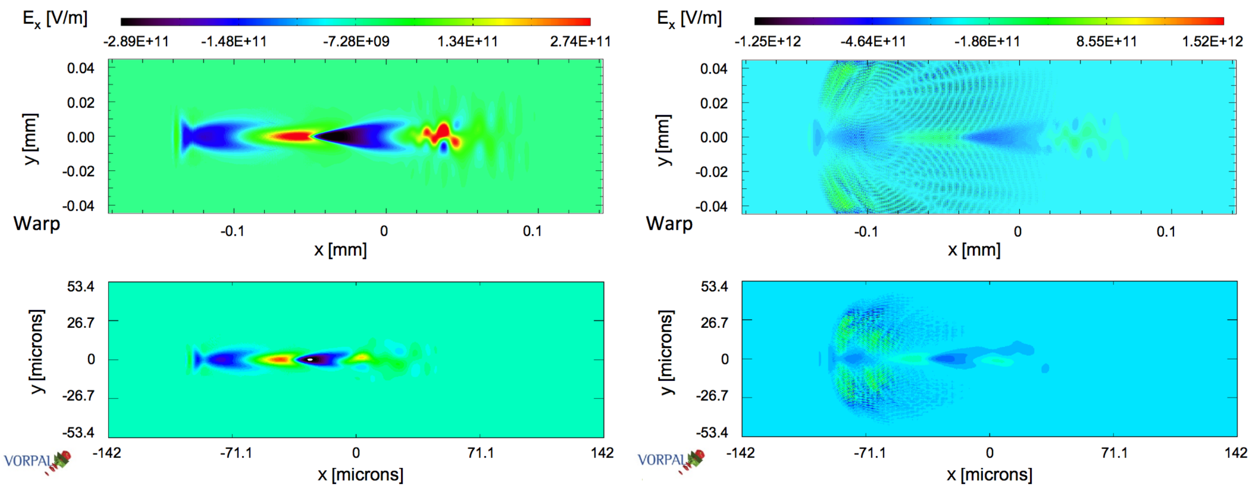

FIGURE 3. Colored contour plots of the longitudinal electric field from 2-1/2D simulations of laser plasma acceleration stages in a boosted frame of $\gamma = 5.6$ with the "magical" time step $c\delta t \approx \delta x/\sqrt{2}$ (left) and another time step $c\delta t \approx \delta x/\sqrt{1.6}$ (right) from the codes Warp (top) and Vorpal (bottom). The two codes used similar but not identical parameters.

CONTROL OF NUMERICAL INSTABILITY AND NOVEL ALGORITHMS

Several algorithms to increase boost performance have been implemented in Warp, aiming at controlling high-frequency numerical instabilities and have been applied to the observed instability in boost: (a) wide band filtering of current density and/or gathered electromagnetic field, (b) tunable damping of electromagnetic field, (c) electromagnetic solver with enlarged stencil with tunable numerical dispersion. The filtering is based on the use of combined bilinear (3 points) digital filters with different strides, so as to adjust the band in the frequency domain, providing efficient filtering over a wide band. The electromagnetic field damping is based on the scheme with adjustable damping parameter from Friedman [21], which was flagged in [22] as being the most potent practical method for mitigating the numerical Cerenkov instability among the selected methods that were considered. The electromagnetic solver with enlarged stencil is based on the solver proposed by Cole and Karkkainen in [23, 24], adapted to the Particle-In-Cell technique in a way that ensures compatibility with standard "exact" current deposition schemes [25, 26]. Perfectly Matched Layers [27] for efficient absorption of outgoing waves and the Friedman adjustable damping extension were also implemented with the enlarged stencil. Details of the implementations are given in [12].

An extensive set of testing was performed on the simulation of a downscaled 100 MeV, [28, 29] and full scale 10 GeV LPA stages, revealing 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. Filtering of the current density (and eventually gathered fields) was shown to be more effective at preserving stability and accuracy than damping of the electromagnetic field. The testing also revealed that there exists a singular time step for which the level of instability is minimal, independently of other numerical parameters, especially the numerical dispersion of the solver. The existence of such a time step was confirmed by D. Bruhwiler using Vorpal. Figure 3 displays the contour plots of the longitudinal electric field from 2-1/2D simulations of laser plasma acceleration stages in a boosted frame of $\gamma = 5.6$ using the "magical" time step $c\delta t \approx \delta x/\sqrt{2}$ or a different time step $c\delta t \approx \delta x/\sqrt{1.6}$ from the codes Warp and Vorpal, both using the Yee solver, no damping of the field nor filtering of the current density nor gathered fields. All snapshots have been taken at the same physical time but the physical parameters of the Warp and Vorpal simulations, although being close, were not identical, explaining the observed differences. Nonetheless, both codes exhibited similar levels of instability at a given physical time which was greatly reduced (or delayed) if using the special time step.

The extensive testing and its analysis indicate that the observed instability may not be caused by numerical Cerenkov effects. Analysis of the nature of the instability is underway, but regardless of cause, the methods presented mitigate it effectively. The tunability of the field solver is key in providing stability in 3D at the singular time step, which is not attainable by the standard Yee solver [30]. Details of the testing and analysis can be found in [12].

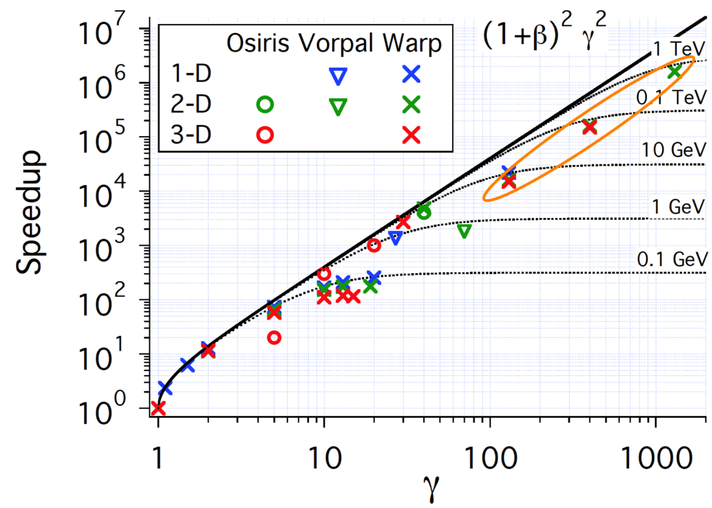

FIGURE 4. Theoretical speedups (lines) for 0.1 GeV to 1 TeV stages and observed speedups from simulations using the code Osiris (circles), Vorpal (triangles) and Warp (crosses), updated with latest Warp results.

APPLICATION TO THE MODELING OF 100 GEV-1 TEV STAGES

Using the novel numerical techniques developed in Warp, simulations of stages in the range of 0.1 GeV-1 TeV were performed in 2-1/2D and in the range of 0.1-100 GeV in 3-D, using Lorentz boosted frame with γ as high as $1,300$, verifying the scaling of plasma accelerators into the 1 TeV range. Using Eq. (1), the speedup of the full scale 100 GeV class run, which used a boosted frame of $\gamma = 400$ as frame of reference, is estimated 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. With the same analysis, the speedup of the 2-1/2D 1 TeV stage is estimated to be over a million. The observed speedups updated with the newest Warp runs are summarized in Fig. 4. With the increased performance now accessible for higher boost values, it becomes possible to use high performance computing to perform high resolution simulations of low emittance beams with beam loading, plasma ramp, etc.

CONCLUSION

Using novel numerical techniques that were implemented in the Particle-In-Cell code Warp to control a high-frequency numerical instability, it was shown that modeling of laser-plasma wakefield accelerators in an optimal frame of reference can produce up to six orders of magnitude speed-up in calculations from first principles of stages in the 1 TeV range, which is three orders of magnitude higher than the higher speedup reported previously for a 1-10 GeV range stage. For a 10 GeV stage, the speedup is over four orders of magnitude, which is one order of magnitude larger than the maximum speedup previously reported. The new results offered a verification of the scaling of plasma accelerators to the 1 TeV range and provide highly efficient tools for the detailed designs of experiments on new lasers such as BELLA [31].

ACKNOWLEDGMENTS

Work supported by US-DOE Contracts DE-AC02-05CH11231 and DE-AC52-07NA27344, and US-DOE SciDAC program ComPASS, including partial support from the ComPASS SciDAC project, grant No. DE-FC02-07ER41499, and from Tech-X Corporation. Used resources of NERSC and LBNL cluster Lawrencium, supported by US-DOE Contract DE-AC02-05CH11231. The authors are grateful to S. F. Martins and W. B. Mori for providing input data.

REFERENCES

1. J.-L. Vay, *Phys. Rev. Lett.* **98**, 130405 (2007).
2. D. L. Bruhwiler, *et al.*, "New Developments in the Simulation of Advanced Accelerator Concepts," in *Proc. 13th Advanced Accelerator Concepts Workshop*, Santa Cruz, CA, 2008, p. 29.
3. J.-L. Vay, *et al.*, "Application of the reduction of scale range in a Lorentz boosted frame to the numerical simulation of particle acceleration devices," in *Proc. Particle Accelerator Conference*, Vancouver, Canada, 2009, TU1PBI04.
4. S. F. Martins, R. A. Fonseca, L. O. Silva, W. Lu, and W. B. Mori, "Boosted Frame PIC Simulations of LWFA: Towards the Energy Frontier," in *Proc. Particle Accelerator Conference*, Vancouver, Canada, 2009, TH4GBC05.
5. J.-L. Vay, *et al.*, *J. Phys. Conf. Series* **180**, 12006 (2009).
6. C. Huang, *et al.*, *J. Phys. Conf. Series* **180**, 12005 (2009).
7. J. L. Vay, W. M. Fawley, C. G. Geddes, E. Cormier-Michel, and D. P. Grote, "Speeding up simulations of relativistic systems using an optimal boosted frame," in *Proc. DPF-2009 Conference*, Detroit, MI, 2009, arXiv:0909.5603.
8. S. F. Martins, R. A. Fonseca, L. O. Silva, W. Lu, and W. B. Mori, *Comput. Phys. Comm.* **182**, 869–875 (2010).
9. S. F. Martins, R. A. Fonseca, L. O. Silva, W. Lu, and W. B. Mori, *Phys. Plasmas* **17**, 056705 (2010).
10. S. F. Martins, R. A. Fonseca, W. Lu, W. B. Mori, and L. O. Silva, *Nature Physics* **6**, 311 (2010).
11. D. P. Grote, A. Friedman, J.-L. Vay, and I. Haber, *AIP Conf. Proc.* **749**, 55 (2005).
12. J.-L. Vay, C. G. R. Geddes, E. Cormier-Michel, and D. P. Grote, *arXiv:1009.2727v2* (2010).
13. R. A. Fonseca, *et al.*, *Lec. Notes in Comp. Sci.* **2329**, 342 (2002).
14. C. Nieter, and J. R. Cary, *J. Comput. Phys.* **196**, 448–472 (2004).
15. Taflove, and Hagness, *Computational Electrodynamics: The Finite-Difference Time-Domain Method*, MA: Artech House, Norwood, 2000, 2nd edn.
16. C. Benedetti, *et al.*, *Nuclear Inst. and Methods in Physics Research A* **608**, 94–98 (2009).
17. F. S. Tsung, *et al.*, *Phys. Plasmas* **13**, 056708 (2006).
18. S. Kalmykov, *et al.*, *Phys. Rev. Lett.* **103**, 135004 (2009).
19. D. L. Bruhwiler (2008), *Private Communication*.
20. C. K. Birdsall, and A. B. Langdon, *Plasma Physics Via Computer Simulation*, Adam-Hilger, 1991, ISBN 0-07-005371-5.
21. A. Friedman, *J. Comput. Phys.* **90**, 292 (1990).
22. A. D. Greenwood, K. L. Cartwright, J. W. Luginsland, and E. A. Baca, *J. Comput. Phys.* **201**, 665–684 (2004).
23. J. B. Cole, *IEEE Trans. Antennas Prop.* **50**, 1185–1191 (2002).
24. M. Karkkainen, E. Gjonaj, T. Lau, and T. Weiland, "Low-DispersionWake Field Calculation Tools," in *Proc. of International Computational Accelerator Physics Conference*, Chamonix, France, 2006, pp. 35–40.
25. J. Villasenor, and O. Buneman, *Comput. Phys. Comm.* **69**, 306 (1992).
26. T. Esirkepov, *Comput. Phys. Comm.* **135**, 144–153 (2001).
27. J.-P. Bérenger, *J. Comput. Phys.* **114**, 185 (1994).
28. E. Cormier-Michel, *et al.*, "Scaled Simulations of a 10 GeV Laser Wakefield Accelerator," in *Proc. 13th Advanced Accelerator Concepts Workshop*, Santa Cruz, CA, 2008, p. 297.
29. C. G. R. Geddes, *et al.*, "Scaled simulation design of high quality Laser Wakefield Accelerator stages," in *Proc. Particle Accelerator Conference*, Vancouver, Canada, 2009, WE6RFP075.
30. K. S. Yee, *IEEE Trans. Ant. Prop.* **14**, 302–307 (1966).
31. W. P. Leemans, *et al.* (2010), *These Proceedings*.