Thermal Trapping in the Blowout Regime of the Plasma Wakefield Accelerator

Scott Kravitz
Department of Physics and Astronomy, University of California, Los Angeles, California 90095
(Dated: September 29, 2010)

The nonlinear, blowout regime of the plasma wakefield accelerator has been the subject of considerable interest due to its potential for use as a next generation accelerator in high energy physics experiments. Much of the analytical work and simulation in this regime has been restricted to scenarios of cold background plasma. Significant theoretical results concerning plasma thermal effects have been obtained only in one dimension. This paper addresses the phenomenon of background electron trapping, in particular by thermal means, and its effects on wakefield development. In order to examine a key source of initial temperature, we also discuss plasma electron heating through the decay of the nonlinear wake waves themselves. Through 2D particle-in-cell simulations, we find that some thermal trapping occurs even at modest plasma temperatures, creating Landau damping of the wakefield that increases rapidly with temperature up to ∼500 eV, and then less strongly for higher temperatures. In addition, the peak electric field produced in the plasma scales linearly with the driving beam charge over a wide range before degrading suddenly once a charge threshold has been reached. These results suggest that temperature is a significant factor in a plasma wakefield accelerator's effectiveness, and that parameters such as the charge of the driving beam and the injection time of the witness beam should be carefully chosen to minimize negative thermal effects.

I. INTRODUCTION

The concept of the plasma wakefield accelerator (PWFA) has been around since 1985 [1] and was first experimentally tested in 1988 [2]. In such an accelerator, an electron beam (the driving beam) is sent through background plasma consisting of electrons and positively-charged ions. The plasma electrons in the vicinity of the driving beam are deflected away, while the higher-mass ions remain relatively unperturbed. The resulting fluctuations in the plasma give rise to an electric field which has regions of alternating sign trailing behind the beam in the transverse direction. In a typical, linear case, the field is a wave of wavelength λ_p (the plasma wavelength) which is inversely proportional to the square root of the plasma density. This wakefield can then be used to accelerate and focus a second, witness electron beam for as long as it remains in the correct region of the wake. This method of acceleration allows the energy from many electrons (in the driving beam) to be transferred to relatively few electrons in the witness beam, thus achieving very high energies over short distances. Recent experiments have created accelerating gradients on the order of 50 GV/m which are sustained over a distance of just under a meter, leading to an energy gain of 44 GeV/m in some of the accelerated electrons [3]. To enable plasma accelerators to compete with large, traditional accelerators such as the LHC, even higher accelerating gradients are necessary. To this end, new experimental methods and considerations are being explored.

Of particular interest is the blowout regime of the PWFA; all simulations presented here were blowout scenarios. This regime is characterized by a driving electron beam which is denser than the background plasma, causing all electrons in the vicinity of the beam to be completely expelled4. In contrast to the linear regime (called so because of the linear restoring force experienced by the plasma electrons) which has sinusoidal plasma waves, this regime exhibits wakefields with a sharp spike near the region where the plasma electrons return to the beam axis. This spike is capable of being much larger than the amplitude of the wave in the linear case. Analytical work done in the case of a driving beam with infinitesimal length suggests that this spike is roughly linear in the normalized beam charge $\tilde{Q}$, defined by $\tilde{Q} = \frac{N_b k_p^3}{n_0}$, where N_b is the number of beam electrons, k_p is the plasma skin depth, and n_0 is the unperturbed plasma electron density [4, 5]. With the proper parameters, wakefields on the order of TV/m can be achieved in this regime.

However, thermal trapping has the potential to undermine attempts at realizing such high fields. Trapping exists when background electrons have sufficient velocities in the transverse direction to be accelerated by the wakefield over large distances (trapped in the accelerating region), thus behaving similarly to witness beam electrons. This takes energy away from the wake, resulting in a smaller accelerating gradient for the witness beam. The effect is particularly significant when the energy gradient of the wakefield is strong enough to accelerate electrons to near the speed of light over a short distance. When the source of the requisite initial velocities is thermal in nature, this is referred to as thermal trapping. In addition, the decay of the wake itself creates plasma heating, which can induce trapping. When the blown-out electrons return to the beam axis, they crash into each other, transferring their directed kinetic energy into heat in what is referred to as wave breaking [7]. This heating signifies both a loss of wakefield energy and a potential source of trapping, and is thus undesirable.

Significant theoretical results concerning the effects of background plasma temperature on wakefield development have been restricted to one-dimensional analysis, with a particular focus on laser wakefield accelerators over PWFAs [8, 9]. By examining the effects of thermal

trapping on wakefield development and the wave breaking mechanism of heating through simulations in 2D, we can better design experiments in the blowout regime of the PWFA to minimize negative thermal effects.

II. SIMULATIONS

In order to simulate a PWFA in the blowout regime, the 2D particle-in-cell (PIC) code OOPIC was used. In the simulations run, a driving electron beam was sent through a cylinder filled with underdense plasma. The system was treated as cylindrically-symmetric, providing information about the particles' motion in the longitudinal (z) and radial (r) directions, but not in the azimuthal (ϕ) direction. The simulation length was fixed at 13.3 μm, while its radius was 2.22 μm. Cell size was fixed at $dz = dr = 61.7$ nm, with a time step of $dt = 8.43 \times 10^{-17}$ s, satisfying the Courant condition that light cannot pass through more than one cell in a single time step. Each boundary of the simulation (left, right, and radial) is treated as a conductor. The plasma is modeled as ionized hydrogen atoms and their electrons, and has an unperturbed density of $n_0 = 7.0 \times 10^{25} m^{-3}$ with 16 macroparticles per cell.

What made these simulations most unique was the driving beam, modeled after those produced by the Linac Coherent Light Source (LCLS) at Stanford. It was modeled as a bigaussian distribution, with dimensions of $r_{rms} = 0.197\mu$m and $z_{rms} = 0.6\mu$m (corresponding to a duration of 2 fs). The short duration of the beam enabled a larger proportion of the beam's energy to be transferred to the wakefield rather than to the back of the beam itself through self-wake effects. In addition, the mean energy of beam electrons was 14.3 GeV, with a minimal rms energy spread of 0.1%.

The first series of diagnostics was intended to give a measure of wakefield decay due to thermal trapping for a range of beam charges. It involved measuring the amplitude of the accelerating wakefield, both the wave immediately following the beam (first wake) and the wave following that (second wake), just before the driving beam begins to exit the simulation ($t = 40$ fs) at different initial plasma temperatures, ranging from 0 eV to 5000 eV. Though the LCLS parameters give a beam charge of 2 pC ($\tilde{Q} = 6.97$), simulations were run for beam charges in the range of $2 \leq \tilde{Q} \leq 50$.

Boundary conditions are often critical to the outcome of PIC simulations, so the same diagnostic was run a second time with an additional 4 μm (one plasma wavelength) of vacuum preceding the plasma. This was used to distance the plasma from the effects of the left conducting boundary, which would tend to attract plasma electrons close to it, thus decreasing the number of electrons with sufficient longitudinal velocity to be trapped. The time at which the wakefields were measured was altered to account for the added distance through which the driving beam traveled. By comparing the results of this series of diagnostics to those of the first, the effects of the conducting boundary could be determined.

Additional simulations were run to determine the effects of wavebreaking on the plasma temperature. The first of these used a driving beam of normalized charge $\tilde{Q} = 7$ which traveled through a perfectly cold plasma ($T = 0$ eV). At various times throughout the simulation, the temperatures of the plasma electrons (radial, longitudinal, and isotropic) were taken and averaged across the entire simulation, to get a measure of the total energy of the wake transferred to heating. The second set of diagnostics measured the plasma electron temperatures well after the driving beam had left ($t = 0.32$ ps) for values of $\tilde{Q}$ ranging from 2 to 50, again with zero initial temperature. Both simulations were run without an initial vacuum region.

III. RESULTS AND DISCUSSION

Plots of the first and second wakefield amplitudes as a function of initial plasma temperature and normalized beam charge in the no-vacuum case are shown below in Fig. 1. The results are shown for $\tilde{Q}$ up to 50 for completeness, although for $\tilde{Q} > 20$, the wakefield began to exhibit behavior that differed radically from that of a more standard PWFA experiment. In particular, the waves did not appear to show periodicity, and in many cases identification of a first and second wake was not straightforward. Nevertheless, in the primary region of interest ($2 \leq \tilde{Q} \leq 20$), both plots indicate a sharp decrease in wakefield amplitude (in the range of 10-20%) with temperature up to approximately 500 eV, with continued but more gradual decrease at higher temperatures (up to 30-40% at 5 keV). This supports the notion that increases in the plasma temperature can significantly affect the wakefield amplitude due to increased thermal trapping. It is worth noting that the first wake degradation decreases with increasing $\tilde{Q}$ in the range of $4 \leq \tilde{Q} \leq 30$, e.g. the amplitude of the wake at $T = 5$ keV is a larger fraction of the amplitude at $T = 0$ for $\tilde{Q} = 20$ than at $\tilde{Q} = 10$. In addition, at a fixed initial temperature, the amplitude of the first wake increases linearly with increasing $\tilde{Q}$ up to $\tilde{Q} = 30$, with a sudden increase at $\tilde{Q} = 40$ followed by a dramatic decrease at $\tilde{Q} = 50$. This initial linear relationship matches theoretical expectations for the nonlinear regime. The spike may have been so dramatic because of the wave amplitude's sensitivity to the time it was measured at for such large $\tilde{Q}$, although its replication at several different temperatures below 500 eV suggests this is not the case. As was observed in simulations (and can be seen by comparing the two plots in Fig. 1), the decrease in the first wakefield was accompanied by an increase in the second wakefield.

The data for the second wakefield are similar, though somewhat less well-behaved. For $\tilde{Q} \leq 20$, the same strong degradation occurs up to 500 eV, with weaker degradation up to 5 keV. Additionally, the wake am-

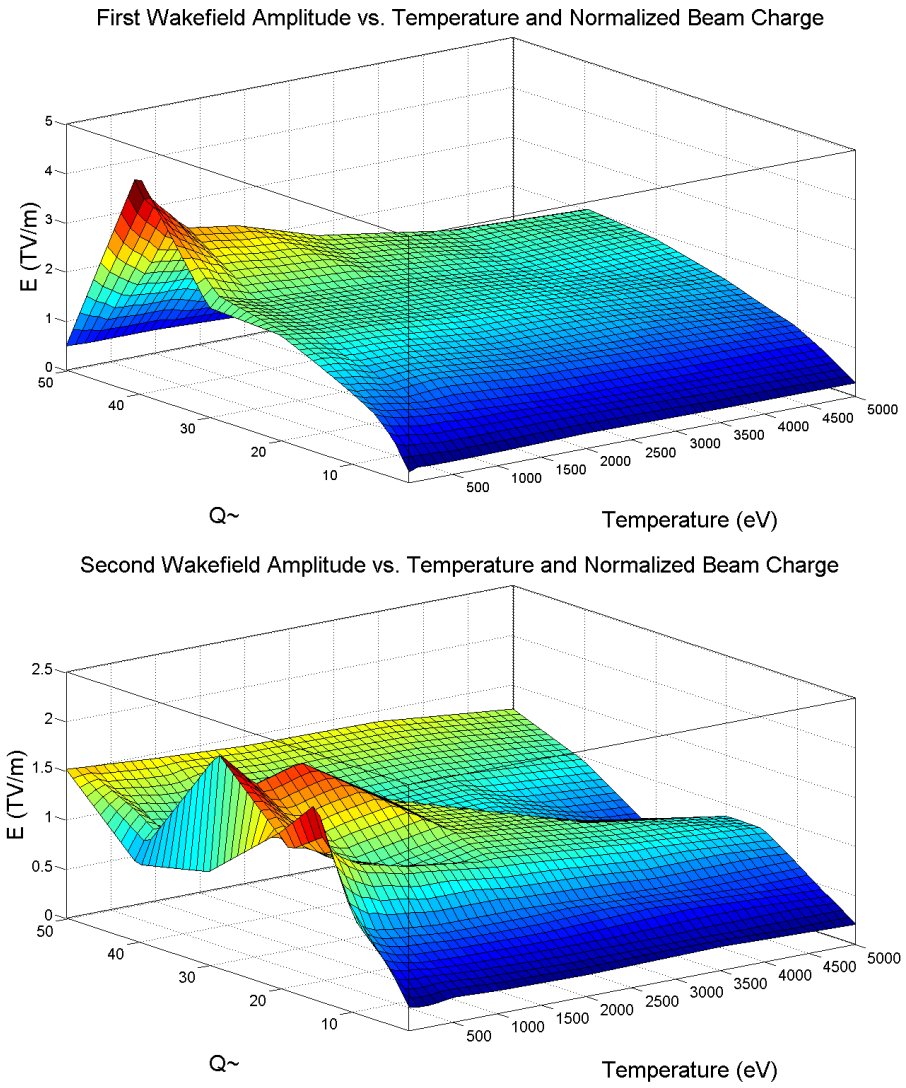

FIG. 1: Wakefield amplitude data for the no-vacuum case.

plitude scales linearly with $\tilde{Q}$ in this region. However, at larger beam charges, the wake amplitude is significantly smaller for low temperatures (< 100 eV) and is positively correlated with temperature in this temperature range, though it resumes its expected decrease for larger temperatures. One possible explanation for this is that, in this temperature region, larger temperatures facilitate the transfer of energy from the first wake to the second, before thermal trapping dominates and decays both wakes.

For comparison, plots of the first and second wakefield amplitudes in the scenario with an initial vacuum region are shown in Fig. 2. The first is qualitatively very similar to its no-vacuum counterpart, exhibiting the most degradation in the first 500 eV, with continued damping at higher temperatures. In addition, the same pattern of decreased degradation with increasing beam charge up to $\tilde{Q} = 30$ is apparent, though less dramatically so. The most noticeable difference between this and the case without vacuum is that the spike in amplitude at $\tilde{Q} = 40$ present before is now much less pronounced. With this notable exception, however, the amplitudes remained relatively unperturbed, being on the average slightly larger than in the no-vacuum case. This suggests that the conducting boundary condition of the first simulation did not significantly affect the level of thermal trapping in the simulation.

The second wakefield, on the other hand, was considerably more well-behaved with the addition of initial vacuum. It followed the same general pattern as the first wakefield in both scenarios, both in terms of the temperature and beam charge dependence. The main disparities were the lack of a sudden drop in amplitude at $\tilde{Q} = 50$ (in this aspect, it more closely resembled the second wakefield in the no-vacuum case) and a positive correlation with temperature between 500 and 2000 eV for

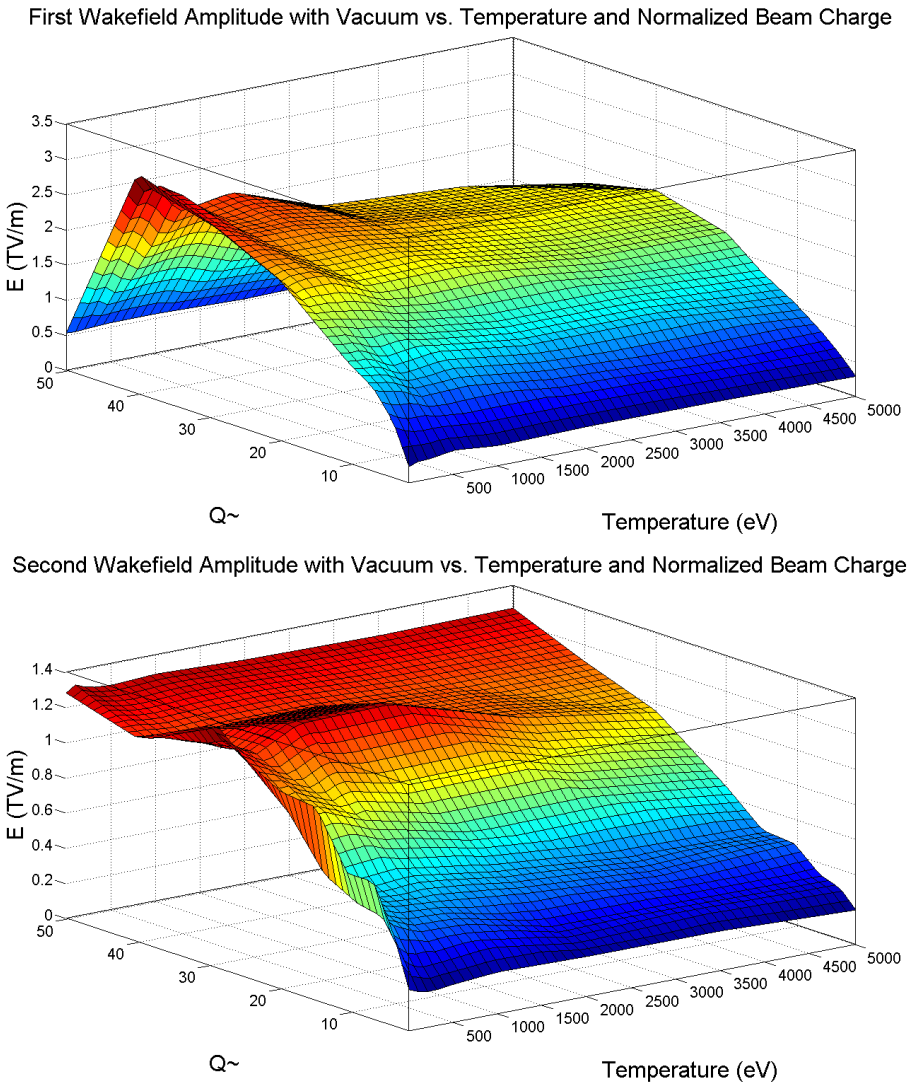

FIG. 2: Wakefield amplitude data for the case with an initial vacuum region.

the $\tilde{Q} = 30$ case. For lower $\tilde{Q}$ values (< 20), where the second wakefield in both scenarios had a similar temperature dependence, the amplitude was consistently smaller (by about 20%) than in the no-vacuum scenario. These data seem to indicate that the existence of an initial vacuum region serves to mitigate unpredictable effects arising from the conducting boundary conditions, while increasing the overall level of thermal trapping moderately.

On the plasma heating side, a plot of the plasma electron temperature as a function of time is presented in Fig. 3. As the driving electron beam propagated through the simulation, two spikes in the temperature occurred, followed by a marked drop soon after the electron beam exited ($t = 60$ fs). After this drop, the temperature rose rapidly and roughly linearly to a peak before steadily decaying. This rapid, linear increase indicates a continual transfer of wakefield energy to heat with continued wavebreaking, while the subsequent decay suggests a steady loss of heat to the surroundings; the lack of further temperature increases after the peak ($t = 0.25$ ps) indicates that by this time plasma oscillations (and hence the wakefield) are minimal and wavebreaking is no longer a significant source of heat. Perhaps more interesting than the shape of the plot are its dimensions; the (isotropic) temperature reached a maximum in excess of 4 keV and remained above 1 keV for much longer (up until $t = 1.5$ ps) than the time scale for the driving beam to travel through the simulation. These data signify that very large heating effects occur in this regime, even well after the driving beam has exited the plasma, so that experiments using multiple driving beams are likely to be strongly hindered by thermal trapping. One potentially redeeming aspect of this large heating is that throughout the wavebreaking process, the longitudinal temperature was significantly smaller, by about a factor of 2, than the radial temperature, which means less thermal trapping.

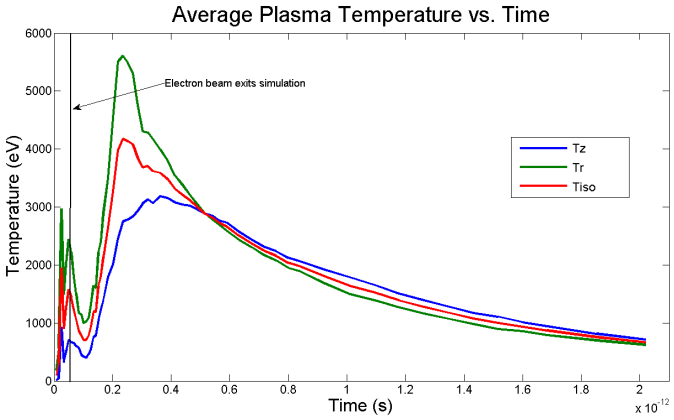

FIG. 3: Plasma electron temperature as a function of time.

To best avoid the negative effects of heating in such an experiment, the witness beam should be timed so as to arrive coincidentally with the drop in temperature (0.11 ps after a driving beam enters, for the simulation parameters) following one of the driving beams as it exits the plasma, perhaps the first.

Lastly, the results of the temperature versus beam charge diagnostic are displayed in Fig. 4. Up through $\tilde{Q}$ = 30, the average plasma electron temperature increases linearly with beam charge. This suggests that the temperature vs. time plot should scale with $\tilde{Q}$ in this range. At larger values, this behavior breaks down, leading to massive heating (80 keV) or, in the $\tilde{Q}$ = 50 case, a complete removal of the plasma electrons from the simulation, making a temperature calculation impossible.

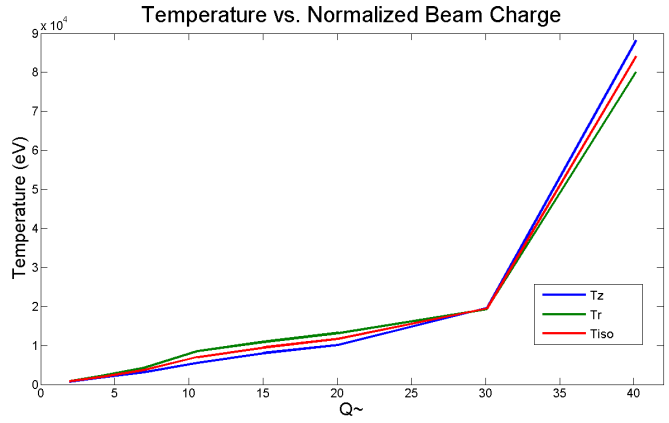

FIG. 4: Plasma electron temperature as a function of beam charge.

IV. CONCLUSIONS

It has been shown that plasma thermal effects can substantially alter the effectiveness of a plasma wakefield accelerator in the blowout regime through thermal trapping of background electrons. Simulations in two dimensions have shown that the largest increases in trapping, with the concomitant degradation of the wakefield, occur for temperatures up to 500 eV, with significant additional degradation continuing through 5 keV. Further simulation data suggest that such degradation becomes a smaller fraction of the initial amplitude with increasing beam charge, but radically different wakefield behavior for the largest charges suggests running experiments at $\tilde{Q} \leq 30$. Effects of conducting boundary conditions on trapping were examined by including a region of vacuum preceding the plasma, with the primary result that some unexpected effects observed in the second wake for the no-vacuum case disappeared. The first wake amplitudes remained relatively unchanged with the addition of vacuum, while the second wake amplitudes decreased moderately, indicating additional trapping due to insulation from the conductor.

Simulations investigating the evolution of plasma electron temperature with time indicate that considerable heating (up to 4 keV for $\tilde{Q} = 7$) occurs after the driving beam exits the plasma, as wavebreaking removes energy from the wakefield, and remains long afterward. In addition, such heating appears to scale linearly with beam charge, potentially placing further limits on the beam charge. Fortunately, more of this heat is transferred to radial motion than transverse motion, reducing the potential amount of future thermal trapping. In addition, a minimum in temperature is obtained soon after the driving beam leaves; taking advantage of this by sending the witness beam through the plasma during this time would alleviate problems arising from thermal trapping, especially in an experiment using multiple driving beams. Taken as a whole, these simulations indicate that initial temperature, beam charge, and timing must be carefully managed to maximize the effectiveness of a PWFA in the blowout regime.

V. ACKNOWLEDGEMENTS

I would like to thank Dr. James Rosenzweig for his help and guidance in completing this project. I would also like to thank Pardis Niknejadi, Brendan O'Shea, Diktys Stratakis, and all of the Particle Beam Physics Laboratory at UCLA for their assistance. Thanks also to Francoise Queval for her direction and coordination of the summer REU program at UCLA. This work was made possible by the generous support of the NSF, grant number NSF-PHY 0850501.

[1] P. Chen, J.M. Dawson, R. W. Huff, and T. Katsouleas, Phys. Rev. Lett. 54, 693 (1985).
[2] J. B. Rosenzweig, D. B. Cline, B. Cole, H. Figueroa, W. Gai, R. Konecny, J. Norem, P. Schoessow, and J. Simpson, Phys. Rev. Lett. 61, 98 (1988).
[3] I. Blumenfeld et al., Nature 445, 741, 15 February 2007.
[4] N. Barov, J. B. Rosenzweig, M. E. Conde, W. Gai, and J. G. Power, Phys. Rev. ST Accel. Beams 3, 011301 (2000).
[5] N. Barov, J. B. Rosenzweig, M. C. Thompson, and R. B. Yoder, Phys. Rev. ST Accel. Beams 7, 061301 (2004).
[6] J. B. Rosenzweig, N. Barov, M. C. Thompson, and R. B. Yoder, Rev. ST Accel. Beams 7, 061302 (2004).
[7] E. Esarey, C. B. Schroeder, and W. P. Leemans, Rev. Mod. Phys. 81, 1229 (2009).
[8] C. B. Schroeder, E. Esarey, B. A. Shadwick, and W. P. Leemans, Physics of Plasmas 13, 033103 (2006).
[9] E. Esarey, C. B. Schroeder, E. Michel, B. A. Shadwick, C. G. R. Geddes, and W. P. Leemans, Physics of Plasmas 14, 056707 (2007).