Influence of proton bunch and plasma parameters on the AWAKE experiment

Mariana Moreira
mariana.t.moreira@tecnico.ulisboa.pt

Instituto Superior Técnico, Lisboa, Portugal

November 2017

Abstract

Plasma-based accelerators are strong candidates for a highly anticipated generation of compact particle accelerators. A proton-driven plasma wakefield experiment (the AWAKE experiment) at CERN will test the underlying concept using long proton bunches that undergo the self-modulation instability (SMI). The effectiveness of the experiment hinges on the successful and predictable development of this instability, which fragments the initial proton bunch into smaller beamlets with lengths of the order of the plasma wavelength.

Using numerical particle-in-cell simulations, this work investigates the effects of inevitable event-to-event variations of the initial parameters of the experiment on its outputs, concluding that the differences in the observed effects are always of the same order as the initial variations and deterministic injection of electrons is thus possible. A hypothetical version of the experiment where the driver particles are substituted by antiprotons is also studied, which is able to achieve higher wakefield amplitudes through not-fully-understood nonlinear effects. Lastly, a powerful tool is developed to investigate the nonlinear phase of self-modulated particle-driven wakefields (after saturation of the SMI), for which no theory exists and where the wakefield amplitude drops steeply. This tool consists of a parallel program that calculates the individual contributions of single beamlets to the overall wakefield, and leads to the conclusion that the spatial decline of the amplitude (along the beam) is due to incoherent interference of single wakes, while the temporal decline (along the propagation distance) is due to the loss of driver charge, the reasons for which are not yet understood in depth.
Keywords: plasma-based accelerators, self-modulation instability, wakefield theory, proton-driven plasma wakefield acceleration, particle-in-cell simulation, parallel computing

1. Introduction

The rapid development of particle physics in the last seventy years has created a demand for ever higher collision energies in particle accelerators, where theories can ultimately be tested. Cutting-edge technology has pushed the energy frontier again and again and has consistently been able to provide the collision energies sought by the scientific community. Nevertheless, current technology is reaching a limit, as the impressive budgets for both existing and projected colliders can attest: e.g. €7.5 billion for the Large Hadron Collider (LHC) and €7.7 billion for the International Linear Collider (ILC) [1], respectively.

Current accelerator schemes rely on radiofrequency (RF) cavities to accelerate particles, by letting electromagnetic waves resonate inside them. The magnitude of the electric fields in these cavities, however, is limited by the dielectric breakdown of the material that constitutes them. A limit on the acceleration gradient (which is of the order of 100 MV/m for RF cavities) translates into longer acceleration distances if higher energies are to be reached.

This is why the concept of a plasma-based accelerator is so interesting in the pursuit of more powerful particle acceleration. Plasmas are already ionized, and can therefore sustain much higher electric fields than any other medium. The maximum field that still allows the propagation of waves inside a plasma is of the order of the non-relativistic wavebreaking field [2]:

$$E_0 = \frac{cm_e\omega_p}{e} \approx 0.96\sqrt{n_0\,[\mathrm{cm}^{-3}]}\ [\mathrm{V/cm}]\ , \quad (1)$$

where c is the speed of light and ω_p is the plasma frequency given by $\omega_p^2 = n_0 e^2/\varepsilon_0 m_e$, which depends on the plasma density n_0, the elementary charge e and the electron mass m_e (ε_0 is the vacuum permittivity). For a typical plasma density of 10^{18} cm^{-3}, $E_0 \sim 100$ GV/m, an acceleration gradi-

ent that is three orders of magnitude higher than conventional RF cavities.

The first concept for a plasma-based accelerator was proposed in 1979, and consisted of a relativistic and intense laser pulse which would excite a large-amplitude plasma wave with a phase velocity close to c [3]. Relativistic particles could then be accelerated by this wave, as long as they became phase-locked with it. The unavailability of intense and short laser pulses at the time stimulated the invention of other schemes reliant on long pulses of lower intensity, such as plasma beat-wave acceleration (PBWA) [4]. In this case the wakefield is excited resonantly by the beat frequency between two laser pulses (which must fulfill the condition $\omega_2 - \omega_1 \simeq \omega_p$).

Plasma wakefields can also be driven by a relativistic particle bunch, however. This scheme is called simply plasma wakefield acceleration (PWFA) [5]. Short, ultra-intense laser pulses carry a limited amount of energy, and as long as staging methods are still in their infancy, PWFA is the most promising scheme to achieve high particle energies in a single stage, since particle bunches can hold more energy than laser pulses.

The wakefield in a plasma-based accelerator acts as a transformer between the driver and the witness particles, transferring energy from one to the other. In order to maximize the final energy of the witness particles, then, the energy of the driver must be as high as possible. The most energetic particle bunches on Earth are produced at the European Organization for Nuclear Research (CERN), which is what led to the inception of the Advanced Proton Driven Plasma Wakefield Experiment (AWAKE) [6, 7].

This experiment intends to demonstrate the concept of PWFA using 400-gigaelectronvolt proton bunches supplied by the Super Proton Synchrotron (SPS) at CERN to accelerate injected electrons. The SPS proton bunches are not ideal since their considerable length (12 cm) could not drive a wakefield effectively (the length of the driver should be of the order of the plasma wavelength λ_p, given by $\lambda_p = 2\pi c/\omega_p$). On the other hand, it is impossible to compress such highly energetic bunches to lengths of the order of 100 µm, ie. the approximate plasma wavelength λ_p that this scheme would require [6].

The solution to this issue consists of letting the long proton bunch undergo an instability that was first observed in long laser pulses propagating in plasma: the self-modulation instability (SMI) [8]. Under its action, the bunch progressively fragments into much smaller "beamlets" (with a length of the order of λ_p) due to periodic transversely focusing and defocusing regions. This instability eventually saturates and the initial proton bunch is self-consistently transformed to a format that can resonantly excite the wakefield.

Table 1: Plasma and beam parameters of the AWAKE experiment [10].

Parameter	Value
Plasma density n_0	$7{\cdot}10^{14}$ cm^{-3}
Plasma length L	10 m
Proton bunch population N_b	$3{\cdot}10^{11}$
Proton bunch length σ_{zb}	12 cm
Proton bunch radius σ_{rb}	200 µm
Proton energy W_b	400 GeV
Proton bunch normalized emittance ϵ_{bn}	3.5 µm
Electron bunch radius σ_{re}	250 µm
Electron energy W_e	15 MeV

This instability must be seeded, however, to prevent the competing hosing instability from disrupting the proton beam [8, 9]. In AWAKE the seed consists of an ionizing laser pulse that copropagates with the center of the proton bunch in a 10-meter-long plasma cell filled with rubidium gas, thereby simultaneously fulfilling the function of creating plasma.

In the future, electrons will be injected with around 15 MeV at 4 m of propagation distance, after the SMI has saturated and the bunch has become fully self-modulated [10]. The main parameters of the AWAKE experiment are listed in Table 1.

As an example of what the accelerating mechanism at AWAKE looks like, Fig. 1 depicts the self-modulated proton bunch and a line-out of the axial wakefield E_z (close to the axis). Electrons are accelerated in regions where E_z is negative, i.e. every half-wavelength. However, and though this is not shown in Fig. 1, the radial electric field E_r always trails the axial field with a phase difference of $\pi/2$, so the regions where electrons are both accelerated and focused are actually a fourth of a wavelength ($\lambda_p/4$) long. The highlighted area in Fig. 1, in green, represents one of these regions.

Bearing in mind that AWAKE relies on a beam-plasma instability to succeed, it is vital to understand whether the experiment's outputs are robust against inevitable and realistic variations of the initial parameters. This work will build on previous efforts in this direction [11]. On the other hand, there are some open questions concerning both the physics of self-modulated plasma wakefield accelerators and the AWAKE experiment specifically, which will also be addressed here. These studies will rely on particle-in-cell simulations performed with the relativistic, massively parallel code OSIRIS [12].

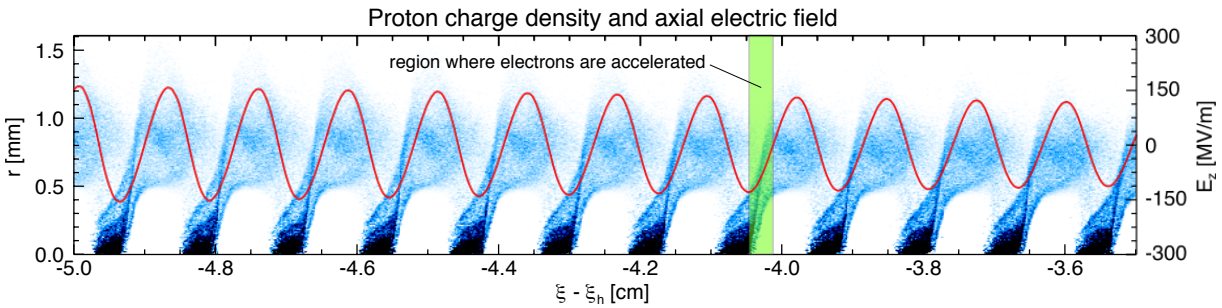

Figure 1: Representation of the self-modulated proton bunch (in blue), the axial electric field E_z close to the axis (red line) and a region where electrons can be accelerated (green rectangle), based on simulation data. The horizontal coordinate is the distance with respect to the head of the beam.

2. Sensitivity Analysis

The purpose of a sensitivity analysis is to ascertain the robustness of a numerical model. This is usually done by refining the model (or certain parameters of the model) and guaranteeing the convergence of the results for a certain level of refinement. In this case the parameters of the model (i.e. the simulation) are already set, through a combination of reasonable estimation and computational compromise, and the purpose of the analysis is to establish whether the results are already safely within an area of convergence. This will lay a foundation of reliability upon which the results from similar simulations throughout this work can rest.

The AWAKE experiment is simulated in two-dimensional cylindrical coordinates , where the OSIRIS coordinates x_1 and x_2 correspond to the usual z and r coordinates, respectively. The simulation window moves in the x_1 direction at the speed of light, thus accompanying the proton bunch and its evolution. There are three particle species in the simulation: the background plasma electrons, the proton beam and a witness electron beam. The latter is implemented with an extremely low density so that this species does not affect the main phenomena.

The general parameters of the simulation are listed in Table 2. Note that all quantities in OSIRIS are normalized to combinations of n_0, ω_p, c, m_e, e and E_0 (from Eq. 1), where the background plasma density n_0 is an independent variable that allows their conversion to physical units. Taking n_0 = 3.5·10^{14} cm^{-3}, L_1 in Table 2 corresponds to 42.41 cm, L_2 to 2.28 mm and the total propagation distance to 15.06 m. The physical parameters chosen for the simulations in this section do not correspond entirely to the AWAKE nominal parameters, but the sensitivity analysis performed here will still be valid.

The initial profile of the proton bunch n_b was implemented according to:

Table 2: General parameters of the baseline simulation for the sensitivity analysis.

Parameter	Normalized units
Length of simulation box, L_1	1492 k_p^{-1}
Height of simulation box, L_2	8 k_p^{-1}
Number of cells in x_1 direction, n_1	18000
Number of cells in x_2 direction, n_2	425
Time step	0.012 ω_p^{-1}
Propagation distance	53000 k_p^{-1}

$$n_b(\xi, r) = \frac{n_{b0}}{2}\left[1+\cos\left(\sqrt{\frac{\pi}{2\sigma_{zb}^2}}\,(\xi-\xi_h)\right)\right] e^{-\frac{r^2}{2\sigma_{rb}^2}}, \quad (2)$$

for $\xi_0 \leq \xi \leq \xi_h$, where ξ is the beam co-moving coordinate given by $\xi = z - ct$, n_{b0} is the density at the peak of the profile (assumed here as n_{b0} = 4·10^{12} cm^{-3}), σ_{zb} is the rms bunch length (taken as $\sigma_{zb} \approx$ 17.8 cm), σ_{rb} is the rms bunch width (taken as $\sigma_{rb} \approx$ 282.8 µm), ξ_0 is the position where the function crosses the ξ axis (in this simulation ξ_0 = 1 k_p^{-1}) , and ξ_h is the position of the beam head, where the ionizing laser pulse is co-propagating in the real experiment (in this simulation ξ_h = 1480 k_p^{-1}).

The numerical parameters chosen for the analysis were the resolutions in the x_1 and x_2 direction (increase by 1.5 and decrease by $1/\sqrt{2}$ with respect to the baseline) and the number of numerical particles per cell of the proton beam species, henceforth signified by ppc$_{p^+}$ (increase by 1.5^2 = 2.25). These numerical particles, also called macroparticles, represent an averaged ensemble of real particles.

The output from OSIRIS simulations includes the distribution of the fields and charge densities in the simulation region. A direct comparison of this data for the axial and radial electric fields E_z and E_r and for the proton charge density ρ_{p^+} was performed at three different regions (each with a width of 45 k_p^{-1}) along the beam: the head (cap-

turing the front of the proton bunch), the middle (about halfway through the simulation window) and the tail (close to the left boundary of the window). In addition, three different simulation times t_1, t_2 and t_3 were probed for this comparison, corresponding to around 1, 8 and 14 m of propagation distance, respectively.

From this comparison it was noticed that barely any differences exist in the head region of the beam for all three quantities (E_z, E_r and ρ_{p^+}) and across all three probing times. The head of the beam is the location along the simulation window where there is less numerical noise (even as the simulation advances in time), probably because it is composed of "fresh" (undisturbed) plasma entering the window from the right, while numerical error accumulates along the beam and causes noisier data towards the tail region. This lack of numerical noise may be the reason why the results here are so robust against simulation parameter variations.

In the remaining regions (middle and tail) differences could be perceived. At the beginning of the simulation the wakefield is still growing and its amplitude can be comparable to that of small fields associated with beam density fluctuations caused by the initial bunch temperature. This means that at this stage the electric fields are both noisier (except around the head of the beam) and more sensible to variations of the simulation parameters. As might be expected, this is especially true for the tail of the beam.

The run with an increased number of particles per cell displays the most pronounced differences in such regions (where there is most noise). The higher number of macroparticles effectively annihilates all noise even in the tail region, though after t_1 the noise mostly disappears and so do the differences between these runs.

It was also found that there is generally most sensitivity towards the resolution in the x_1 direction, and that most results are not fully convergent with those of this run, both in terms of the configuration of the fields/charge density and their magnitudes.

An overview of the evolution of the wakefields can be obtained by calculating the average amplitude of E_z in the entire simulation window for each simulation time. The results produced by this diagnostic are shown in Fig. 2, where the highest sensitivity can once again be observed for changes of the resolution in the x_1 direction. The results from the baseline simulation have clearly not converged to those of the refined resolution in x_1. This means that the baseline simulation delivers slightly overestimated values (by around 8%) for the amplitude of E_z up to the peak (which is

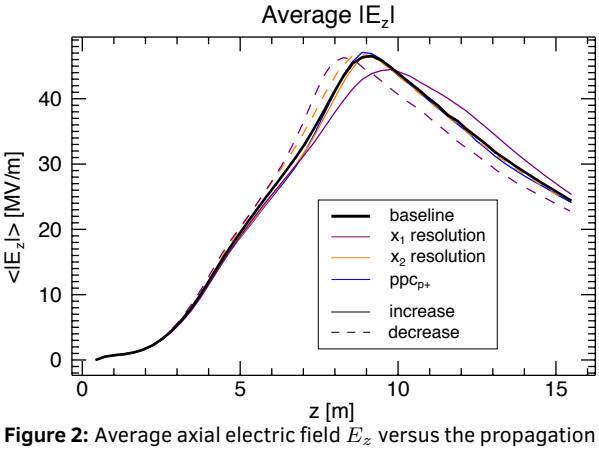

Figure 2: Average axial electric field E_z versus the propagation distance for all the sensitivity analysis runs.

itself anticipated in time by roughly 6%), and underestimated values after the curve begins to fall again (at most by 8%).

The total proton charge at each time was also computed and compared for the different parameter variations. Though the different curves converge at the end of the simulation, between 6 and 15 m the total proton charge is underestimated in the baseline run by up to 10%.

There is in fact a sizable discrepancy between the results from the baseline simulation and those with a finer resolution in x_1, but the magnitude of these deviations has been determined, which at least enables the estimation of a margin of error when drawing conclusions about the quantities affected. It should be mentioned that these findings are particularly relevant when deducing *absolute* values from the simulations or their post-processed results. However, most objectives set out for this work are based on *relative* results between different runs, where the incomplete convergence found in this section is not critical.

3. Parameter scans of the AWAKE simulation

The objective of this section is to determine how expectable variations of the initial parameters of the AWAKE experiment can affect its outputs, most notably the wake characteristics and the final energy of hypothetical accelerated electrons. Shot-to-shot fluctuations of the proton bunch driver parameters of the order of a few percent are to be expected.

The motivation for this study is two-sided. On one hand, it is interesting to understand how the instability at stake here, the SMI, reacts to variations of the driver bunch that triggers it, which is a more fundamental question and not necessarily tied to AWAKE. On the other hand, it is crucial for AWAKE to show that, despite its reliance on an instability, electrons can be injected deterministically into this accelerator.

Table 3: Differing simulation parameters with respect to the ones in Table 2, used for the parameter scans.

Parameter	Normalized units
Length of simulation box, L_1	1663 k_p^{-1}
Number of cells in x_1 direction, n_1	20063

The standard simulation used for comparing the parameter scans is slightly different than the one described in the previous section and in Table 2. The simulation window is longer in order to accommodate parameter variations of the proton bunch length, and the number of cells in the x_1 direction is also increased so that the resolution in this direction remains equivalent.

The values for the beam dimensions used in these simulations are closer to AWAKE's nominal parameters, with the rms length σ_{zb} = 12.6 cm and width σ_{rb} = 200 µm. The proton beam profile is described by the same equation as before (Eq. 2), though the positions of the beam head and the function's zero are shifted: here ξ_h = 1651 k_p^{-1} and ξ_0 = 80.2 k_p^{-1}. The peak density n_{b0} also has a different value here, calculated from the bunch dimensions mentioned above and the fixed bunch population N_b = 1.5·10^{11} through the equation

$$n_{b0} = \frac{N_b}{(2\pi)^{2/3}\,\sigma_{rb}^2\,\sigma_{zb}}\ , \tag{3}$$

giving $n_{b0} \approx$ 1.89·10^{12} cm^{-3}. In addition, the ionization radius r_p is set as 150 µm (this is the height up to which there is plasma in the simulation window). In all other respects the standard simulation here is identical to the baseline one from the last section. The differing simulation parameters are displayed in Table 3. For a background plasma density of 7·10^{14} cm^{-3}, L_1 corresponds to 33.43 cm, L_2 to 1.61 mm and the total propagation distance to 10.65 m.

The following experiment parameters were varied for this study: the rms proton bunch length σ_{zb} ($\pm$5%) and width σ_{rb} ($\pm$5%), the proton bunch population N_b ($\pm$5%), the plasma radius r_p ($\pm$5%), and the rms timing jitter of the proton bunch with respect to the ionizing laser pulse Δt ($\pm$15 ps).

Equation (3) means that, if the density is kept constant, any variations of the proton bunch's dimensions lead to a simultaneous change of the bunch's population (a parameter we wish to analyze independently as well). Hence, in order to avoid mixing variations of two parameters, the density n_{b0} was altered in runs with variations of bunch dimensions so as to conserve N_b. On the other hand, the timing jitter Δt is in practice a phase shift of the cosine in Eq. (2) with respect to the position of the beam head ξ_h, thus encapsulating either more or less charge depending on whether the maximum of the cosine is moved to the right or left of ξ_h. This shift will be denoted by an additional term $\Delta\xi$ in the cosine argument from Eq. (2):

$$\cos\left(\sqrt{\frac{\pi}{2\sigma_{zb}^2}}\,(\xi - \xi_h - \Delta\xi)\right)\ . \tag{4}$$

Ignoring possible effects due to the different shape of the bunch's profile, the absolute values of variation of Δt ($\pm$15 ps) correspond to differences of $\pm$2.85% in the amount of charge driving the wake.

3.1. Properties of the wake

The results for the average amplitude of the axial electric field E_z are shown for all the variations in Fig. 3. Two groups of experiment parameters can be identified according to their qualitative effects on E_z. Variations of the rms beam dimensions σ_{zb} and σ_{rb} cause inversely proportional variations of the average $|E_z|$, while the bunch population N_b, the substantially equivalent timing jitter Δt and the plasma radius r_p (at a much smaller scale) are proportional to it. Linear wakefield theory predicts that $E_z \propto n_{b0}$ (see Eq. (5) on p. 7, where n_{b0} is implicit in the density profile $n_b(\xi, r)$), and the results from this diagnostic seem to confirm this relationship, since according to Eq. (3) $n_{b0} \propto N_b\,\sigma_{rb}^{-2}\,\sigma_{zb}^{-1}$.

Although σ_{rb} should have a larger effect on E_z than σ_{zb}, this is only observed until the maximum of the average axial field, where variations of $\pm$10% w.r.t. (with respect to) the standard run are reached at 5 m. At the peak these variations begin to decrease and the effects of σ_{zb} variations become dominant, with a maximum of around $\pm$8% at 5.3 m.

At the end of the plasma cell, or at 10.33 m (since the plasma medium only begins after a

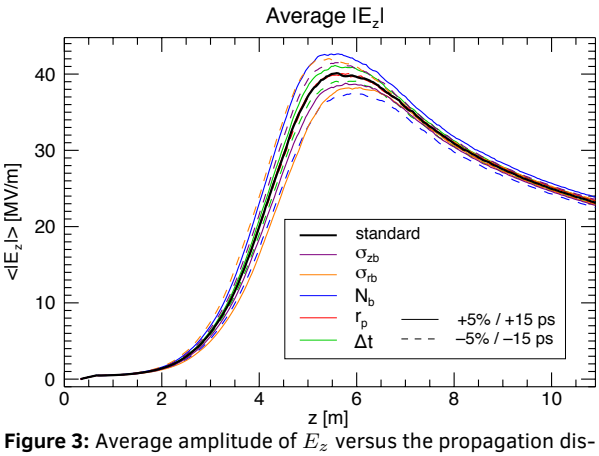

Figure 3: Average amplitude of E_z versus the propagation distance for all parameter scans.

length of one simulation window, or approximately 33 cm), the results from all the runs have converged to different values and the largest variation w.r.t. the standard run is of around $\pm 3\%$ for the rms bunch length σ_{zb}. This result is perfectly satisfactory for the proof-of-principle experiments to be conducted at AWAKE, though the larger amplitude variations around the middle of the propagation distance could be problematic for the production of high-quality electron beams in the future.

During the growth of the SMI the phase velocity of the wake varies considerably, and continues to change slowly even after the saturation of the instability [13]. This can have the effect of decreasing the dephasing length available for the acceleration of an injected electron, so it is crucial to investigate how the phase of the wakefield evolves and how it reacts to different initial parameters.

The local phase shift of the wake w.r.t. its head was determined through a fit of the function $f(\xi) = A\ \sin\left[k_p\left(\xi - \xi_h\right) + \phi\right]$, where A and ϕ are the fitting parameters, to the electric field (by taking a line-out and dividing it into several fitting windows).

The position along the beam chosen for the comparison of the evolution of the phase shift through time was approximately 11.7 cm behind the head of the beam. The phase shift generally grows along z, with the exception of a plateau around 3.5 m, which is linked to the saturation of the SMI. Before and after this plateau the magnitude of the effects is below the initial parameter variations. At the plateau, however, the largest variations are for the rms bunch width σ_{rb} (around 5% larger at 3.5 m and 5% smaller at 3.2 m).

3.2. Behavior of accelerated electrons

This section will study the behavior of accelerated electrons using a specially devised tool (which is essentially a one-dimensional particle pusher). We wish to determine the injection point along the propagation distance for an electron to achieve a certain final energy.

This is accomplished by "positioning" an electron at z = 10.33 m (the end of the plasma cell) and at a certain ξ_f (some final position along the wake) with a certain final energy γ_f, and propagating it backwards in time through a leapfrog-type relativistic particle pusher. The longitudinal force acting on the electron is surmised from a line-out of E_z close to the axis for that particular simulation time. The time (or propagation distance) when the electron meets a cell where the electric field is positive (i.e., where it would lose energy) will determine the minimum injection point z_{inj} (a minimum threshold for the energy of the electron at injection was also imposed, corresponding to

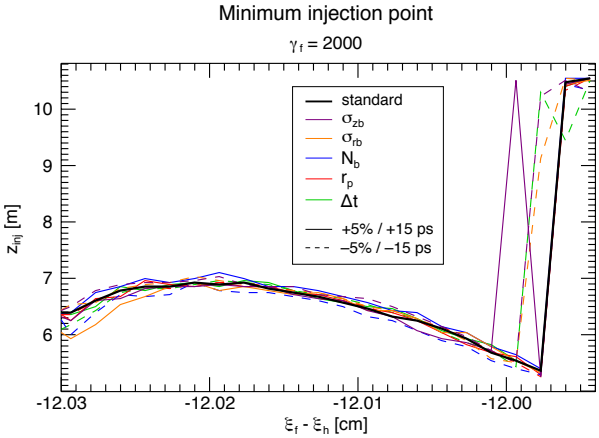

Figure 4: Minimum injection point for electrons with γ_f = 2000 versus their final position along the beam w.r.t. the head $\xi_f - \xi_h$ for all parameter scans.

15 MeV).

The comparison of the resulting injection point data for all the parameter scans was performed for a region located around 12 cm behind the head of the beam and for a final energy of 1 GeV (γ_f = 2000). This data is shown in Fig. 4, whereby the region corresponds to the accelerating and focusing fourth of a plasma wavelength.

It is visible that the decreased bunch dimensions σ_{zb} and σ_{rb} and the decreased timing jitter Δt seem to shorten the accelerating and focusing region by around 5%. Around the middle of the graph, it can be conclusively observed that the injection point is proportional to the bunch population N_b and inversely proportional to the bunch dimensions, though the variations are of only $\pm 2\%$.

These results show that the most important variables underpinning deterministic injection, namely the size and position of an accelerating and focusing wakefield phase and the injection point along z, vary by the same order of a driving bunch's shot-to-shot oscillations.

4. Single beamlet linear wakefield solver

At the moment there is a theoretical apparatus that describes exclusively the linear stage of the SMI, where the long particle bunch becomes fully self-modulated and the axial wakefield E_z grows exponentially [13]. After the saturation of the SMI (the nonlinear stage), however, the amplitude of the wakefield drops significantly (see Figs. 2 and 3) and, despite numerical studies [14, 15], it is not fully understood why this happens.

One of the possibilities currently under consideration as an explanation for the nonlinear phase is related to the incoherent interference between the wakes driven by each beamlet. In order for a single beamlet's wakefield to add constructively to an existent wave, the beamlet must be located in a decelerating and focusing region of the original

wakefield. This problem is in fact spatiotemporal. The behavior of the wakefield must be understood both along the beam at a fixed time (spatially) and along a certain propagation distance for a fixed beam position (temporally).

Understanding the nonlinear phase of the SMI theoretically is the motivation behind the development of a computer program that combines linear wakefield theory with simulation data. This program uses the charge density data from OSIRIS simulations after self-modulation has occurred, identifies and isolates the beamlets, calculates each beamlet's wakefield, and sums all the contributions to an overall E_z distribution. This method will for the first time separate the contributions from single beamlets.

The axial electric field excited in plasma by an axisymmetric particle beam with a given density profile $n_b(\xi, r)$ in cylindrical coordinates is given by the equation [2]

$$E_z(\xi, r) = 4\pi e k_p^2 \int_{\infty}^{\xi} d\xi' \int_0^{\infty} dr' r' \cos\left[k_p(\xi - \xi')\right] I_0(k_p r_<) K_0(k_p r_>) n_b(\xi', r') \,, \tag{5}$$

where I_0 and K_0 are the zeroth-order modified Bessel functions, and $r_{<,>}$ stands for the smallest/largest of r and r'. Both infinities in the integrals represent the regions beyond the beam in question (in the case of ξ) and above the beam (in the case of r). In practice, however, after each beamlet there will be a finite position ξ_f after which all density values are zero. The same is true in the transverse direction, for some r_f. Using this, a discretized version of Eq. 5 was obtained which could be implemented in the program and thus solved numerically. The computation of E_z and other parts of the program were parallelized (distributed over different processes).

This program was written in the C programming language, where the Message Passing Interface library was used for its parallelization. It has been compiled on a 1920-core cluster at IST, and tested using up to the full cluster capacity (1800 nodes) without any issues.

In order to prove the accuracy of the algorithm, a simple benchmark was devised to test the program. This benchmark consisted of a single short electron bunch propagating in plasma in the linear regime of wakefield excitation, with an axially Gaussian and radially cosinusoidal density profile. The resulting electric field was in excellent agreement with that obtained from a PIC simulation.

4.1. Applications in the AWAKE simulations

The program was used to determine the axial wake from the density data of the AWAKE simulation at three different times chosen after the linear phase of the SMI, corresponding to the propagation distances z ≈ 6.39, 8.52 and 10.65 m.

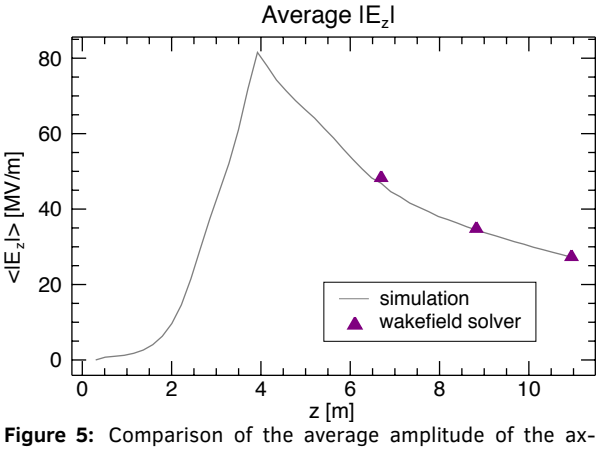

Figure 5: Comparison of the average amplitude of the axial electric field according to the simulation's results and the wakefield solver's results.

There was moderate agreement between the program- and simulation-derived field distributions in the middle and tail regions at 6.39 m, but the results converge for the other two times. The head region is identical at all three times. As mentioned before, the head region is least subject to cumulative noise effects along the simulation window, which may suggest that the disagreement between simulation and program results in other regions is mostly due to numerical noise. The average amplitudes of E_z for the simulation and the wakefield solver are presented in Fig. 5, demonstrating the subtle convergence of both results for later times and their generally good agreement.

The program was also used to test the theory of incoherent interference between the wakes caused by each beamlet. The wakefield solver was modified so as to stop after a given number of beamlets and to produce both the wakefield originated exclusively by the last beamlet, and the overall wakefield generated by the remaining beamlets. This allows the direct observation of how the field caused by a consecutive beamlet adds to the field already present. This modified version of the code was applied to the last two times mentioned above (z ≈ 8.52 and 10.65 m), and for three different beamlet numbers: 8, 90 and 170 (counted from the head of the beam).

The relative amplitudes of the consecutive wake w.r.t. to the exisiting one do not change between both times. The relative amplitude is a measure of how much influence the next beamlet can exert over the present wakefield, and it obviously decreases along the beam. In this case the relative amplitudes are roughly 20%, 0.3% and 0.1% for beamlets number 8, 90 and 170, respectively.

The eighth beamlet is still located very closely to the beam head, which may explain why its wake is

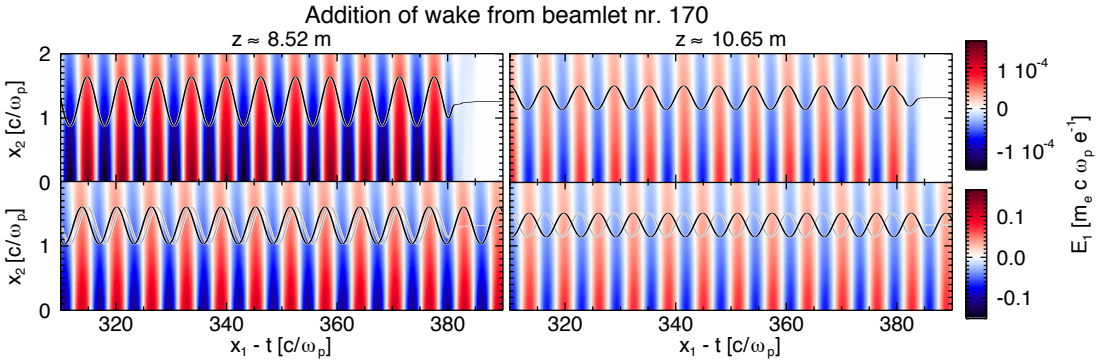

Figure 6: Wake caused by the 170th beamlet (top row) and wake caused by the first 169 beamlets (bottom row) for two propagation distances (left and right). The position of the 170th beamlet corresponds to about 25% of the simulation window. The line-outs of the field at the axis are superimposed. The line-outs for the 170th beamlet (top row) are reproduced in the bottom row (gray) for comparison, though their amplitudes are also not to scale.

almost imperceptibly out of phase with the existing wake, as was observed. By the ninetieth beamlet, however, there is an unambiguous phase-shift of about a fourth of a period, or $\lambda_p/4$. This phase shift remains virtually unchanged through both simulation times.

Around the tail of the beam (beamlet number 170) the phase shift at 8.52 m is the same as the middle region, i.e. $\lambda_p/4$, which can be observed in Fig. 6. Surprisingly, though, between this distance and 10.65 m the position of this beamlet is advanced and the phase-shift thereby evolves to a fully canceling half-period, or $\lambda_p/2$. This obviously does not mean that protons are moving faster than c, but rather that some protons on the left side of the beamlet at 8.52 m were lost to defocusing, and other protons were refocused on the right side, so that it appears as though the beamlet has advanced. This mechanism is nonetheless only observed at the tail of the beam, where the relative amplitudes of the individual wakes are extremely small, and so it cannot be held as a significant factor behind the amplitude decline along z (temporally).

These observations nonetheless prove that the incoherence between consecutive contributions to the overall wakefield by single beamlets must be the dominating cause for the decline of the electric field amplitude along the beam (along ξ). The fact that the results from the linear wakefield solver overlap with the ones from PIC simulations implies that no significant nonlinear phenomena are being overlooked, and that the phenomena exposed by the wakefield solver are the sole cause for the amplitude drop in ξ.

What is truly decisive for the amplitude decrease in z is the fact that driver charge is gradually lost. Based on the simulation results, the average amplitude of E_z is plotted alongside the total proton charge below r = 3 k_p^{-1} in Fig. 7. This radial limit is to avoid tallying driver charge that is still inside the simulation window but is not contributing to wakefield generation. The identical trends of both curves after saturation of the SMI (around 6 m) indicate that driver charge loss is the overwhelming cause for the decreasing wakefield amplitude along time (or z).

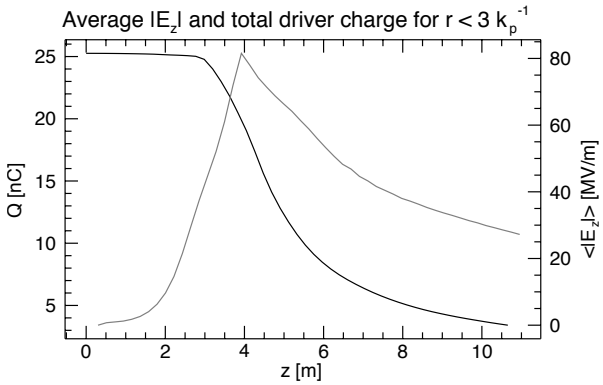

Figure 7: Average amplitude of the axial electric field (gray) and total driver charge below 3 k_p^{-1} or approximately 0.6 mm (black), according to simulation results.

5. Antiprotons as wakefield drivers

As mentioned in section 1, PWFA was initially proposed with electrons in mind as the wakefield drivers [5], and most experiments use this setup. The AWAKE experiment is unusual in comparison with previous efforts in that its wakefield driver is a positive particle. It has been found that the positive counterparts of electrons, positrons, seem to be less efficient at driving a wakefield [16], which begs the question whether the negative counterparts of protons would in turn excite the wakefield at AWAKE more effectively. This section thus explores a hypothetical substitution of the particles in the AWAKE driver bunch by their antiparticles: antiprotons.

With the exception of the bunch dimensions (σ_{zb} = 12.6 cm and σ_{rb} = 200 µm) and the plasma density (n_0 = 7·10^{14} cm^{-3}), the baseline simula-

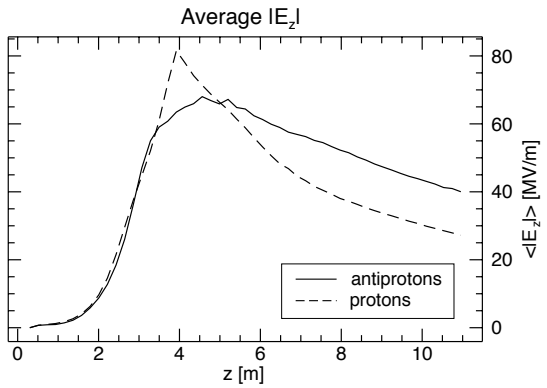

Figure 8: Average amplitude of E_z versus the propagation distance for the proton and antiproton runs.

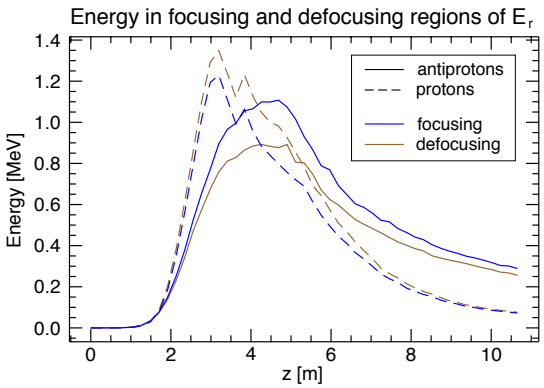

Figure 9: Energies contained in focusing and defocusing regions of E_2 (w.r.t. to the charge of the driver particles) for both runs (left) and ratio between these energies for both runs (right).

tion used for this study is exactly the same as the one described in section 2. For this plasma density, the physical length of the simulation window is 29.99 cm, its height is 1.61 mm and the total propagation distance is 10.65 m.

The simulation with an antiproton driver did produce beneficial results. The average amplitude of E_z, for example, falls less rapidly after saturation and is about 1.5 times higher than that for protons at the end of the plasma cell (see Fig. 8). Perhaps expectedly, the amount of charge driving the wakefield is also higher after the peak, by a factor of around 2.8. In addition, almost 4.6 times more witness electrons remain inside the simulation window by the end of the plasma cell.

A first attempt at understanding why the antiproton driver is able to conserve so much of its charge was directed at the field energy contained in focusing and defocusing areas of the radial electric field E_r, which contributes most to the radial forces acting on the bunch particles. This calculation yielded the graphs displayed in Fig. 9.

In general most of the energy in the proton run is in defocusing fields, while for the antiproton run the majority of it is in focusing fields. The different rates of decline lead to a considerable distance between the focusing energies of both runs after 10 m, where the focusing energy for the antiproton driver is around 3.8 times higher than the one for the proton driver. This may explain the disparity in retained driver charge.

This diagnostic inadvertently supplies some clues regarding the linearity or nonlinearity of the plasma wakes, since a linear regime would imply that the energies associated with focusing and defocusing fields are the same. The larger the imbalance between these energies, the more nonlinear the respective wake will be. This imbalance can be perceived as the distance between the focusing and defocusing curves in Fig. 9 for each driver. The generally larger distance for antiprotons indicates that the wake generated by them is more nonlinear than the one generated by protons.

It should be remarked that, according to linear wakefield theory, the amplitude of the axial wakefield E_z should be proportional to the total charge driving the wake (note the integration of the beam density profile $n_b(\xi, r)$ in ξ and r in Eq. (5)). For this reason, one would expect the increase factors in the amplitude of E_z and in total wakefield-driving charge for the antiproton run to be similar, which is not the case (the factors are 1.5 and 2.8, respectively). This means that, although antiprotons are able to excite wakefields more efficiently than protons, the amount of antiproton charge available would suggest that even higher wakefields could be achieved.

In order to investigate why the antiproton driver does not transfer as much energy as expected to the wake, a diagnostic was devised to track the presence of driver charge in accelerating and decelerating fields. The result can be seen in Fig. 10.

During the saturation phase of the SMI (6 – 10 m) the antiproton driver transfers much more energy to the wake. However, this proportion of

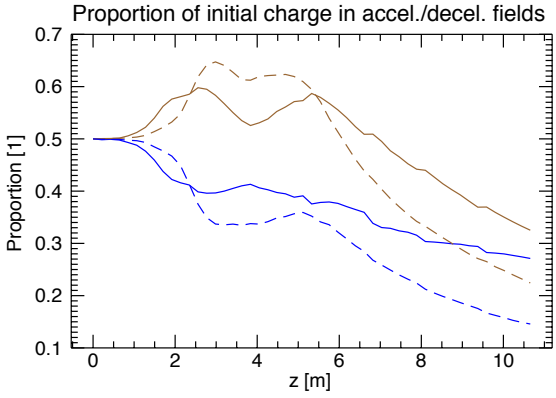

Figure 10: Proportion of initial driver charge in accelerating and decelerating regions for both runs.

charge (the one in decelerating regions) gradually decreases as z grows (note the diminishing distance between the solid brown and blue curves). In fact, after 10 m of plasma, although there is overall less proton charge, a higher portion of it (61%) is contributing energy to the wake as opposed to the portion of antiproton charge left (55%). At first the idea that this difference could explain the underwhelming average amplitude of E_z for antiprotons was entertained, but this difference is far too small to lend this argumentation any credibility.

6. Conclusion

After establishing the reliability of the numerical model being used to simulate the AWAKE experiment, a parameter scan of the initial conditions of the experiment was performed. The relative effects were consistently seen to be of the order of the parameter variation (the largest recorded factor between both was 2), which means that the deterministic injection of electrons is possible. It can thus be concluded that electron acceleration in the AWAKE experiment is viable for proof-of-principle experiments and possibly beyond, for high-quality electron beams.

In section 4 we sought to understand exactly why the amplitude of the wakefield drops during the nonlinear stage of the SMI by developing a parallel program that can solve the linear equation for E_z numerically. It was found that the spatial decline of the wakefield amplitude (along the beam) is due to incoherent interference between the wakes caused by each single beamlet. The temporal decline (along the propagation distance), on the other hand, was found to be correlated with the loss of driver charge.

Finally, we investigated whether antiprotons are more efficient wakefield drivers than protons, as is the case for electrons and positrons. This was shown to be true with exceeding clarity, though the wakefield amplitude is lower than might be expected considering the increased driver charge available for antiprotons. It was observed that a lower portion of antiprotons transfers energy to the wakefield after saturation of the SMI in comparison with the proton case, but this alone cannot account for the discrepancy. Moreover, it was shown that the wakefields driven by antiprotons are more nonlinear than the those driven by protons.

References

[1] N. Patel, Nature **449**, 133 (2007).

[2] E. Esarey, P. Sprangle, J. Krall, and A. Ting, IEEE Transactions on Plasma Science **24**, 252 (1996).

[3] T. Tajima and J. M. Dawson, Physical Review Letters **43**, 267 (1979).

[4] Y. Kitagawa, T. Matsumoto, T. Minamihata, *et al.*, Physical Review Letters **68**, 48 (1992).

[5] P. Chen, J. M. Dawson, R. W. Huff, and T. Katsouleas, Physical Review Letters **54**, 693 (1985).

[6] A. Caldwell, K. V. Lotov, A. Pukhov, and F. Simon, Nature Physics **5**, 363 (2009).

[7] R. Assmann, R. Bingham, T. Bohl, *et al.*, Plasma Physics and Controlled Fusion **56**, 084013 (2014).

[8] N. Kumar, A. Pukhov, and K. V. Lotov, Physical Review Letters **104**, 255003 (2010).

[9] J. Vieira, W. B. Mori, and P. Muggli, Physical Review Letters **112**, 205001 (2014).

[10] P. Muggli, E. Adli, A. Alexandrova, *et al.*, "AWAKE readiness for the study of the seeded self-modulation of a 400 GeV proton bunch," (2017), unpublished.

[11] N. Savard, J. Vieira, and P. Muggli, in *Proceedings of PAC*, submitted, 2016.

[12] R. A. Fonseca, L. O. Silva, F. S. Tsung, *et al.*, in *Computational Science — ICCS 2002: International Conference Amsterdam, The Netherlands, April 21–24, 2002 Proceedings, Part III* (Springer Berlin Heidelberg, 2002) pp. 342–351.

[13] C. B. Schroeder, C. Benedetti, E. Esarey, *et al.*, Physical Review Letters **107**, 145002 (2011).

[14] K. V. Lotov, Physics of Plasmas **22**, 103110 (2015).

[15] K. V. Lotov, Physics of Plasmas **20**, 083119 (2013).

[16] S. Lee, T. Katsouleas, R. G. Hemker, *et al.*, Physical Review E **64**, 045501 (2001).