# Influence of proton bunch and plasma parameters on the AWAKE experiment 

## Mariana Azevedo Trocado Moreira

Thesis to obtain the Master of Science Degree in

## Engineering Physics

Supervisor: Prof. Jorge Miguel Ramos Domingues Ferreira Vieira

## Examination Committee

Chairperson: Prof. João Pedro Saraiva Bizarro
Supervisor: Prof. Jorge Miguel Ramos Domingues Ferreira Vieira
Members of the Committee: Prof. Ricardo Parreira de Azambuja Fonseca Prof. Nelson Manuel Carreira Lopes

## Acknowledgments

I bumped into my thesis subject quite serendipitously. I had heard about the AWAKE project about two years before from a friend who was working in it, but I had no idea that anyone at IST was involved. When she told me that there would be a collaboration meeting at my university, I quickly did some research and found out about this Jorge Vieira person (thank you, Janet).

Before I was finished realizing what a great opportunity had fallen into my hands, I soon discovered that my lucky break extended to the character of my advisor. I couldn't have hoped for a more patient, available and demanding-to-the-right-measure supervisor, and it truly was a pleasure to work with him.

But it didn't end there. At GoLP I found a welcoming, professional atmosphere, a desk to sit at, and some company for lunch (alright, and also some friends). I truly am thankful for the healthy doses of assorted buffoonery, alongside the serious lessons I learned about being a researcher.

So perhaps it is only fair that the last months of working on this thesis were fraught with some personal adversity, for those who believe in a finite amount of good fortune. In any case I would like to express a heartfelt thank you to the people who stood closer to the storm.

After all this (and to my surprise) I still have the courage to say that I hope the work I encounter in the future is as challenging and growth-inducing as this was. Go figure.

## Resumo

Aceleradores baseados em plasma são fortes candidatos a uma geração de aceleradores de partículas compactos altamente antecipada. Uma experiência no CERN (chamada AWAKE) vai testar o conceito de wakefields conduzidos por protões em plasma, em que os bunches de protões longos se sujeitam à instabilidade de auto-modelação (SMI). A eficácia da experiência depende fortemente do desenvolvimento previsível e bem-sucedido desta instabilidade, que fragmenta o bunch inicial em pequenos (da ordem do comprimento de onda de plasma) feixes. Utilizando simulações numéricas particle-in-cell, este trabalho investiga os efeitos de variações inevitáveis dos parâmetros iniciais da experiência de evento para evento nos parâmetros de saída da mesma, concluindo que os efeitos observados são da ordem das variações iniciais e que a injeção determinística de eletrões é, portanto, possível. Uma versão hipotética da experiência com antiprotões no lugar de protões também é estudada, em que os campos alcançam amplitudes mais altas através de efeitos não-lineares que não são completamente compreendidos. Por último, uma ferramente poderosa é desenvolvida para investigar a fase não-linear dos wakefields conduzidos por feixes auto-modulados (depois da saturação da SMI), para a qual não existe teoria e onde a amplitude dos campos cai rapidamente. Esta ferramenta consiste num programa paralelo que calcula as contribuições individuais de feixes auto-modelados individuais para o wakefield total, e leva à conclusão de que a queda espacial é devida à interferência incoerente dos wakes individuais, enquanto que a temporal é devida à perda de carga condutora, as razões para a qual ainda não são compreendidas completamente.

Keywords: aceleradores baseados em plasma, instabilidade de auto-modulação, teoria de wakefield, aceleração em plasma por protões, simulação particle-in-cell, computação paralela


#### 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 acceleration, particle-in-cell simulation, parallel computing

## Contents

List of Tables ..... ix
List of Figures ..... xi
Acronyms ..... xvii
1 Introduction ..... 1
1.1 Plasma-based accelerators ..... 3
1.2 Proton-driven PWFA ..... 4
1.3 The AWAKE experiment ..... 5
1.4 State of the art ..... 7
1.5 Objectives and outline ..... 8
1.6 Methods ..... 8
1.6.1 Some notes about the PIC method and OSIRIS ..... 8
2 Sensitivity analysis ..... 11
2.1 Simulation and parameters ..... 11
2.2 Results ..... 12
2.2.1 Direct results ..... 12
2.2.2 Indirect results ..... 15
2.3 Conclusion ..... 22
3 Parameter scans of the AWAKE simulation ..... 23
3.1 Simulation and parameters ..... 23
3.2 Results ..... 26
3.2.1 Properties of the wake ..... 26
3.2.2 Behavior of accelerated electrons ..... 28
Analysis of the parameter scans ..... 30
General observations ..... 32
3.3 Conclusion ..... 34
4 Single beamlet linear wakefield solver ..... 35
4.1 Description of the problem and requirements ..... 36
4.2 Algorithms and implementation ..... 38
4.2.1 Stages 1 through 5: Beamlet isolation ..... 38
Stage 1: First try at setting marks ..... 38
Stage 2: Identifying groups of marks ..... 39
Stage 3: Correcting the boundaries ..... 39
Stage 4: Filling out missing marks ..... 39
Stage 5: Smoothing borders ..... 40
4.2.2 Stage 6: Wakefield computation ..... 40
4.2.3 Implementation ..... 43
4.3 Benchmarking of the solving algorithm ..... 44
4.4 Applications in the AWAKE simulations ..... 44
4.4.1 General wakefield ..... 45
4.4.2 Separation of individual wakes ..... 46
4.5 Conclusion ..... 49
5 Antiprotons as wakefield drivers ..... 51
5.1 Simulation ..... 51
5.2 Results ..... 52
5.2.1 Primary outputs ..... 52
5.2.2 Deeper analysis of charge preservation ..... 52
5.2.3 Deeper analysis of bunch-plasma energy transfer ..... 56
5.2.4 Applying the wakefield solver to the antiproton simulations ..... 60
5.3 Conclusion ..... 62
6 Summary and future work ..... 63
Bibliography ..... 65
A Description of the diagnostics developed for this work ..... 69
A. 1 Peak and average axial electric field ..... 69
A. 2 Concatenation of axial electric field line-outs ..... 69
A. 3 Total charge of particle species ..... 71
A. 4 Maximum witness electron energy ..... 72
A. 5 Local phase shift of the wake ..... 72
A. 6 Hypothetical energy gain and minimum injection point of witness electrons ..... 72
A. 7 Energy contained in focusing and defocusing regions of $E_{r}$ ..... 73
A. 8 Charge density in the longitudinal and transverse force plane ..... 74

## List of Tables

1.1 Plasma and beam parameters of the AWAKE experiment. . . . . . . . . . . . . . . . . . . . 7
1.2 Normalization of physical quantities in OSIRIS (a hat denotes the normalized quantity). . 9
2.1 General parameters of the baseline simulation for the sensitivity analysis.. . . . . . . . . 13
2.2 Summary of the runs performed for the sensitivity analysis with each respective param-
eter variation. . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 13
2.3 Regions of the simulation window and simulation times chosen to compare the results from the sensitivity analysis. 13
3.1 General simulation parameters of the standard simulation for the parameter scans. The altered parameters (w.r.t. Table[2.1) are shown in boldface. . . . . . . . . . . . . . . . . . 25
3.2 Experiment parameters chosen for the analysis, their respective variations, and the con-
5.1 General parameters of the baseline simulation for the study of antiprotons as drivers. . . 51

## List of Figures

1.1 A chart of the progress of particle colliders in terms of the center-of-mass energy. ..... 1
1.2 Snapshot of a numerical simulation of an intense laser pulse (in rainbow colors) prop-agating in plasma (in grey shadowing) superimposed with a graph of the local electricfield.2
1.3 Snapshots from a three-dimensional PIC simulation of the initial beam and after the SMI has developed. ..... 4
1.4 A schematic image of the AWAKE experiment site and the surrounding accelerator com- plex at CERN. ..... 6
1.5 Demonstrative arrangement of the proton and laser beams in AWAKE. ..... 6
1.6 Image of the longitudinal profile of a proton bunch traveling through plasma with distinct signs of modulation. ..... 7
1.7 Diagram of the PIC loop implemented in OSIRIS. The $p$ index denotes a macroparticle, while the indices $i j$ refer to the grid coordinates. The orange and green boxes representthe particle pusher and the field solver, respectively.9
2.1 Plot of the density profile from Eq. (2.1) for $r=0$. ..... 13
2.2 Three exemplary comparisons of the results from the sensitivity analysis runs in the "head" region ..... 14
2.3 Comparison of the distribution of the radial electric field $E_{2}$ in the "middle" region for variations of the resolution in $x_{2}$ at three simulation times. ..... 14
2.4 Comparison of the distribution of the radial electric field $E_{2}$ in the "middle" region for variations of the resolution in $x_{1}$ at three simulation times. ..... 16
2.5 Comparison of the distribution of the axial electric field $E_{1}$ in the "tail" region for an increased number of particles per cell of the proton species at $t_{1}$ and $t_{2}$. ..... 16
2.6 Comparison of the distribution of the axial electric field $E_{1}$ in the "tail" region for varia- tions of the resolution in $x_{1}$ at $t_{2}$ and $t_{3}$. ..... 17
2.7 Comparison of the distribution of the radial electric field $E_{2}$ in the "tail" region for vari- ations of the resolution in $x_{1}$ at $t_{2}$ and $t_{3}$. ..... 17
2.8 Comparison of the distribution of the proton charge density $\rho_{p^{+}}$in the "middle" region for variations of the resolution in $x_{1}$ at $t_{2}$ and $t_{3}$. ..... 17
2.9 Peak axial electric field $E_{1}$ versus the propagation distance for all the sensitivity analysis runs. ..... 18
2.10 Average axial electric field $E_{1}$ versus the propagation distance for all the sensitivity anal- ysis runs. ..... 18
2.11 Comparison of the concatenated line-outs of $E_{1}$ close to the axis over the propagation distance ..... 19
2.12 Total charge of the proton species below a radius of one plasma skin depth over the propagation distance and for all sensitivity analysis runs. ..... 20
2.13 Absolute values of the total charge of the witness beam species over the propagation distance and for all the sensitivity analysis runs (amplified scale). ..... 21
2.14 Maximum energy reached by a witness electron versus propagation distance for all sen- sitivity analysis runs. ..... 21
3.1 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 rect-angle), based on simulation data. The horizontal coordinate is the distance w.r.t. thehead of the beam.25
3.2 Average amplitude of $E_{1}$ versus the propagation distance for all parameter scans. ..... 26
3.3 Local phase shift of $E_{1}$ close to the axis versus the propagation distance $z$ and the po-sition along the beam w.r.t. to its head $\xi-\xi_{h}$, for the standard simulation. The positionalong the beam where all the line-outs were taken for comparison is indicated by thelight gray dashed line.27
3.4 Local phase shift of the axial component of the wake 11.72 cm behind the head of the beam versus the propagation distance for all parameter scans. A highlight of a portion of the graph is included on the right. ..... 28
3.5 Concatenated line-outs of $E_{z}$ close to the axis and example of an electron trajectory(the superimposed yellow line) calculated by the diagnostic for injected electrons. Thelabels $z_{\mathrm{inj}}, \xi_{f}$ and $\xi_{i}$ denote the minimum injection point and the final and initial beamcoordinates of the electron, respectively.29
3.6 Energy gain for injected electrons with $\gamma_{f}=2000$ versus their final position along the ..... 31
3.7 Minimum injection point for electrons with $\gamma_{f}=2000$ versus their final position along the beam w.r.t. the head $\xi_{f}-\xi_{h}$ for all parameter scans. This figure is the line-out marked by the white dashed line in Fig. 3.9.| ..... 32
3.8 Required initial Lorentzfactor of an injected electron according to its final position along the beam w.r.t. the head $\xi_{f}-\xi_{h}$ and its final energy in terms of a Lorentz factor $\gamma_{f}$ for two different regions of the beam (note the logarithmic scale of the data). These results are from the standard simulation. ..... 33
3.9 Deduced minimum injection point for electrons according to their final position along the beam w.r.t. the head $\xi_{f}-\xi_{h}$ and their final energy in terms of a Lorentz factor $\gamma_{f}$ for two different regions of the beam (results from the standard run). The white dashed linerepresents the line-outs taken for the analysis of the parameter scans.33
4.1 Typical evolution of the average amplitude of $E_{z}$ along time or, equivalently, propagation distance (data taken from an OSIRIS simulation). The linear and nonlinear stages of theSMI are indicated on the graph.35
4.2 Example of the axial electric field $E_{z}$ along the beam $(\xi)$ for a fixed time. The gray data is taken from an OSIRIS simulation, while the red lines are a smoothed envelope of thisdata.36
4.3 Sample of charge density data for a portion of the simulation window just before the head of the beam (on the right), where $x_{1}$ and $x_{2}$ are the OSIRIS coordinates correspondingto $\xi$ and $r$, respectively. The white dashed line represents the line-out shown in Fig.|4.4.|. 38
4.4 Line-out of the charge density data in the $x_{1}$ direction at $x_{2}=2$ in the same region as Fig. 4.3. ..... 38
4.5 Images exemplifying all five stages required for beamlet isolation in the same sample region. Images a) and b) include the beam density distribution in the background, while
the colored points denote marks set by the algorithm. . . . . . . . . . . . . . . . . . . . . 40
4.6 Symbolic representation of the method by which the wakefield is calculated for every cell in a column, saving redundant operations.
4.7 Pseudo-code for the algorithm that determines the contribution to the global wakefield for an entire column (fixed $\xi_{0}$ ) and for a certain beamlet, based on the C language.
4.8 Schematic data grid for a given beamlet (represented by the shaded area on the right) to demonstrate the distribution of columns by different processes, where the unit amount of columns ( $c_{\text {unit }}$ ) is 2, the cell length of the grid $(L)$ is 15 and the number of processes $\left(n_{p}\right)$ is 5 .
4.9 Column chart of two example distributions using 27 processes, for two different beamlets which require grids of two different lengths.43
4.10 Diagram showing the structure of the program and the different files composing it. Diamondshaped objects with a red outline represent external libraries. Stadium-shaped objects

| stand for header files and rectangular objects for source files, whereby a green outline |
| :--- | :--- |
| signifies that the files are not original. The direction of the arrows means "includes".. . . 43 |

4.11 Comparison between the axial wakefield obtained from an OSIRIS simulation (top) and from the program (bottom) for the benchmarking test.44

4.12 Comparison of the average amplitude of the axial electric field according to the simula
tion's results and the wakefield solver's results.
4.13 Comparison of the axial wake in the tail region at two different times and for the simulation and the wakefield solver. The line-outs of the data close to the axis are superimposed (at $\mathrm{z} \approx 10.65 \mathrm{~m}$ the simulation line-out is reproduced in gray in the graph for the wakefield solver for comparison).45
4.14 Comparison of the axial wake in the head region for the simulation and for the wakefield solver.
4.15 Wake caused by the 8th beamlet (top row) and wake caused by the first 7 beamlets (bottom row) for two propagation distances (left and right). The position of the 8th beamlet corresponds to about 96\% of the simulation window. The line-outs of the field at the axis are superimposed. The line-outs for the 8th beamlet (top row) are reproduced in the bottom row (gray) for comparison, though their amplitudes are not to scale (compare the different color scales for both rows, on the right).
4.16 Wake caused by the 90th beamlet (top row) and wake caused by the first 89 beamlets (bottom row) for two propagation distances (left and right). The position of the 90th beamlet corresponds to about 60\% of the simulation window. The line-outs of the field at the axis are superimposed. The line-outs for the 90th beamlet (top row) are reproduced in the bottom row (gray) for comparison, though their amplitudes are again not to scale. 48
4.17 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. . 48

[^0]5.1 Peak and average amplitudes of $E_{1}$ (left and right graph, respectively) versus the prop- agation distance for the proton and antiproton runs. ..... 53
5.2 Total charge of the driver species below a radius of one plasma skin depth (left), andtotal charge of the witness electron species in the entire simulation window (right), bothversus the propagation distance and for both runs.53
5.3 Maximum witness electron energy for both runs versus the propagation distance. ..... 53
5.4 Energies contained in focusing and defocusing regions of $E_{2}$ (w.r.t. the charge of thedriver particles) for both runs (left) and ratio between these energies for both runs (right). 54
5.5 Evolution of $E_{2}-B_{3}$ (the unsigned transverse force, in red and blue) and the absolutevalue of the driver bunch charge density $|\rho|$ (in green) between 2.56 and 4.69 m for aproton and an antiproton driver. Note that negative (blue) areas of $E_{2}-B_{3}$ are focusingfor protons while positive (red) areas are focusing for antiprotons.55
5.6 Charge density of antiprotons in the transverse/longitudinal force plane at $z \approx 1.28 \mathrm{~m}$ (example). ..... 57
5.7 Proportion of initial driver charge occupying each of the four quadrants in the force planeversus the propagation distance and for three different regions of the beam (designatedfor every row). The left and right columns show the results for the simulations withprotons and antiprotons, respectively.57
5.8 Smoothed line-outs of the axial electric field close to the axis for both runs at around ..... $\square$
6.39 m of propagation distance, at the head of the beam (top) and in the middle of the beam (bottom). ..... 58
5.9 Proportion of initial driver charge in accelerating and decelerating regions for both runs. ..... 59
5.10 Proportion of current driver charge in accelerating and decelerating regions, weighted by the axial field $E_{z}$ acting on it and unweighted. The left and right graphs correspond to the results from the simulations with protons and antiprotons, respectively. ..... 60
5.11 Comparison of the average amplitude of the axial electric field according to simulation results (lines) and the wakefield solver's results (triangles), for antiprotons and protonsas drivers.61
5.12 Comparison of the axial wake in the middle region with an antiproton driver as deter- mined by the simulation (top) and the wakefield solver (bottom). ..... 61
A. 1 Example of the product from the diagnostic described in this section for a segment of the simulation window, overlaid with three slopes at different points with different meanings for the wake's phase velocity $v_{\phi}$. ..... 70
A. 2 Schematic illustration of the quantities required to calculate the wake's phase velocity using the concatenated line-outs of $E_{1}$. ..... 70
A. 3 Schematic representation of the force plane $W_{r} / W_{z}$ and meaning of each quadrant for protons (purple) and antiprotons (teal). ..... 75

## Acronyms

```
AWAKE Advanced Proton Driven Plasma Wakefield Experiment. 5, 8, 11, 12, 22, 24, 27, 30, 32,34,
    44,45,49,51,62,64,72,73
CERN European Organization for Nuclear Research. 4,51
CLIC Compact Linear Collider.4
CPA chirped-pulse amplification. 3
GoLP Group of Lasers and Plasmas. 8,43
HDF Hierarchical Data Format. 36,43
IDL Interactive Data Language. 43, 63,69
ILC International Linear Collider. 4,5
IPFN Institute for Plasmas and Nuclear Fusion. 8
IST Instituto Superior Técnico. 8,43
LHC Large Hadron Collider. 1,4,5
LWFA Laser Wakefield Acceleration. 2-5
MPI Message Passing Interface.36,43
PBWA Plasma Beat-Wave Acceleration. 3,5
PIC particle-in-cell. 5, 7, 9, 30, 33, 34, 44, 47, 64,73
PWFA Plasma Wakefield Acceleration. 3, 5, 8, 51,
RF radiofrequency. 1, 2,4,32
rms root mean square. 3, 5, 7, 11, 24, 26, 28,52
SMI self-modulation instability. 5, 7, 8, 19, 23, 28, 32, 33, 35, 36, 44, 45, 47, 49, 54, 56, 59, 63, 64
SPS Super Proton Synchrotron. 5,7
UCLA University of California, Los Angeles. 8
```


## Chapter 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. Simultaneous technological development has allowed for the construction of particle accelerators that could approximately keep up with an exponential growth in the collision energy, until two decades ago. The Large Hadron Collider (LHC), despite its budget of $€ 7.5$ billion and a circumference of 27 km , was the first accelerator to escape this trend (see Fig. 1.1).

The current situation is that future particle accelerators will require increasingly high budgets to go beyond the energy frontier. The main reason for this is closely linked to the use of radiofrequency (RF) cavities to accelerate particles. The magnitude of the electric fields in these cavities is limited by the dielectric breakdown of the material that constitutes them. In order to achieve higher collision energies, one then needs acceleration distances longer than those available today.

In order to face this challenge, in 1979 the concept of a laser-driven plasma accelerator [2] was first proposed. According to the authors, a relativistic and intense laser driver would excite a largeamplitude plasma wave with a phase velocity close to $c$ (the speed of light), and relativistic particles could then be accelerated by these waves in case they were phase-locked with the wave. The wave is excited because the propagating laser exerts a ponderomotive force (which results from the radiation


Figure 1.1: A chart of the progress of particle colliders in terms of the center-of-mass energy (adapted from [1]). The square symbols indicate hadron colliders and the circular symbols indicate electron-positron colliders.


Figure 1.2: Snapshot of a numerical simulation of an intense laser pulse (in rainbow colors) propagating in plasma (in gray shadowing) superimposed with a graph of the local electric field (in red) [6].
pressure of the pulse and is proportional to the gradient of the field's intensity) on the plasma particles, pushing them away from regions with higher intensity to regions with lower intensity. It is the electrons in the plasma that are mostly pushed away from the laser, since the ions are much heavier and can be considered immobile. For this reason, the space left behind the blown electrons becomes positively charged and they are drawn back towards the laser axis, whereby they will tend to overshoot the axis at high enough speeds. Thus, as the laser pulse progresses in the plasma, periodic regions with positive and negative space charges are left behind. Due to this charge separation, a longitudinal oscillating electric field develops behind the pulse.

By injecting electrons at the right phase of the electric field and provided that the wakefield amplitude or their initial velocity satisfy the trapping condition, they become trapped and are accelerated in the wake of the laser. This concept came to be known as Laser Wakefield Acceleration (LWFA). A simulation of the laser pulse propagating in plasma can be seen in Fig. 1.2. Relativistic electrons become trapped in the left end of the ion channel (the white region).

Since plasmas are already ionized, they can sustain much higher electric fields than any other medium. For waves to propagate through the plasma, the maximum field that can be supported is of the order of the nonrelativistic wavebreaking field [3]:

$$
\begin{equation*}
E_{0}=\frac{c m_{e} \omega_{p}}{e} \approx 0.96 \sqrt{n_{0}\left[\mathrm{~cm}^{-3}\right]}[\mathrm{V} / \mathrm{cm}] \tag{1.1}
\end{equation*}
$$

where $\omega_{p}$ 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 $m_{e}$ ( $\varepsilon_{0}$ is the vacuum permittivity). For a typical plasma density of $10^{18} \mathrm{~cm}^{-3}, E_{0} \sim 100 \mathrm{GeV}$. The highest acceleration gradient known to be supported by any kind of RF cavity is of the order of $200 \mathrm{MV} / \mathrm{m}$ [4], but they usually operate at much smaller gradients, close to tens of MV/m (e.g. [5]). Hence, plasma can support accelerating fields that can be more than three orders of magnitude higher than conventional RFcavities. Plasma-based accelerators could thus provide the way towards more compact, affordable and powerful particle accelerators. Despite this tremendous potential, there are several open questions (mostly related to beam quality and efficiency) which are subject to intense research efforts today.

### 1.1 Plasma-based accelerators

Plasma-based accelerators can be driven by an intense laser pulse driver (LWFA, mentioned above) or by a particle bunch driver. There has been a number of experimental proof-of-principle demonstrations of both concepts [3].

Ideally, the LWFA uses a single short ultra-intense laser pulse to excite a plasma wave, where "short" means a pulse duration smaller than $\lambda_{p} / c$, with the plasma wavelength $\lambda_{p}$ given by $\lambda_{p}=2 \pi c / \omega_{p}$. Since lasers with both requirements (short and very intense) were not available when experiments started, an alternative configuration was suggested where two long laser pulses (of lower intensity) produced a beat frequency that would resonantly excite the plasma wave. This configuration is called Plasma Beat-Wave Acceleration (PBWA) [7]. This mechanism can be compared to a harmonic oscillator that is driven to oscillate at its natural frequency, and which will therefore oscillate to ever higher amplitudes. In PBWA the driver of the oscillations is the beat between the two laser pulses, which must coincide with the plasma's natural frequency $\omega_{p}$ (the laser frequencies $\omega_{1}$ and $\omega_{2}$ therefore have to fulfill the condition $\omega_{2}-\omega_{1} \simeq \omega_{p}$ ). The amplitude of the wake then grows along the driver.

Similar beating patterns can be obtained even when only one frequency mode is used. There are several mechanisms that can lead to this behavior. In purely one-dimensional scenarios, for instance, the presence of an initial laser with some $\omega_{0}$ and a small seed (in the form of a plasma density perturbation, for example) is enough to set off an instability, called RamAN forward scattering [8], that transfers energy from the laser to daughter waves with $\omega=\omega_{0} \pm \omega_{p}$. This ultimately creates a beating pattern equivalent to the mechanism behind PBWA and the generation of high-amplitude plasma waves as well as particle acceleration have in fact been observed through Raman forward scattering [9].

In three-dimensional scenarios, however, the dominant mechanism to drive large amplitude plasma waves is the laser self-modulation instability [10], where the laser must be long compared with $\lambda_{p}$. In order to preserve the high intensity of the laser, this can be achieved by decreasing $\lambda_{p}$ through higher plasma densities, since $\lambda_{p} \propto 1 / \sqrt{n_{0}}$. As seen before, the ponderomotive force exerted by the leading edge of the laser pulse pushes electrons away and thus produces inhomogeneities in the plasma density. This means that the local plasma frequency $\omega_{p}$ and consequently the local refractive index will vary, with the effect of periodically focusing and defocusing the rest of the laser pulse, and gradually fragmenting it into shorter pulses with widths of the order of $\lambda_{p}$. This train of pulses in turn resonantly excites a plasma wave (as in PBWA), which amplifies the periodic focusing and defocusing. This mechanism thus develops into an instability that saturates when the maximum electric field amplitude is reached, and it is designated as self-modulated LWFA [11].

Eventually, though, the use of a single short laser pulse became possible with the advent of chirpedpulse amplification (CPA) [12]. CPA enabled the amplification of ultra-short light pulses to peak intensities of the order of petawatts, which was impossible before due to the detrimental effects of the high intensity on the gain medium of the amplifier. The current record for laser-driven plasma wakefield acceleration (or LWFA) is for electrons with energies up to 4.2 GeV in 9 cm of plasma with a density of $7 \cdot 10^{17} \mathrm{~cm}^{-3}$ [13], which corresponds to an acceleration gradient of approximately $470 \mathrm{GV} / \mathrm{m}$. This beam of electrons was furthermore produced with a $6 \%$ root mean square (rms) energy spread and contained 6 pC of charge.

The strong suit of LWFA schemes is the creation of very high acceleration gradients over small distances (of the order of millimeters) [14]. However, the scaling of such schemes to longer distances (and thus higher final energies) is not straightforward, since short high-intensity laser pulses carry a limited amount of energy. This was the initial motivation behind Plasma Wakefield Acceleration (PWFA) [15] (which uses a particle beam to drive the plasma wake), since particle beams can have higher energies than laser pulses and thus promise higher energy gains. The traditional concept relies on a relativistic
electron beam to play the role of the driver, and it has now been demonstrated experimentally several times [16-21] (one experiment used positrons as driver particles).

Similarly to the laser pulse in LWFA, the electron driver beam repels electrons in the plasma as it propagates, which creates areas devoid of electrons and populated by ions behind the beam. These areas attract the repelled electrons back towards the axis, and so on. If the driver has a positive charge, such as a positron beam, the electrons in the plasma are pulled towards the passing beam instead of pushed away. Since the beam is already gone by the time the electrons reach the axis though, they overshoot it and a region of positive ions is once again formed.

In the linear regime, where the ratio of beam density to plasma density is smaller than one, or $n_{b} / n_{p}<1$, and for a longitudinally symmetric driving beam, the amount of energy that an accelerated electron can gain from the wakefield is limited to twice the maximum beam energy [22]. This is expressed by the transformer ratio

$$
\begin{equation*}
R_{t}=\frac{E_{\text {max } x, \text { witness }}}{E_{\text {max }, \text { beam }}} \leq 2 \tag{1.2}
\end{equation*}
$$

where "witness" refers to the accelerated particle. This limit can be increased both for asymmetric driving beams [23] and with a symmetric beam if one is in the nonlinear (or "blowout") regime [24], where $n_{b} / n_{p}>1$.

### 1.2 Proton-driven PWFA

Though the teraelectronvolt (TeV) regime in particle colliders is already a reality for hadrons, with the LLC , this scale has not yet been reached for leptons. A TeV lepton collider is the next aspiration for particle physicists, to complement and extend the knowledge obtained with the $\boxed{L H C}$. A circular synchrotron scheme like the one used to accelerate protons at the European Organization for Nuclear Research (CERN) is, however, less efficient for lepton acceleration, since the radiation particles emit when forced on a curved trajectory is inversely proportional to their mass and leptons thus lose much of their energy.

A TeV lepton accelerator would then have to be linear, which is the setup pursued for the International Linear Collider (ILC) and for the Compact Linear Collider (CLIC), two prospective particle accelerators currently in planning (the ILC is represented in Fig. 1.1 as "LC"). The acceleration technology considered for both projects is based on RF cavities. The ILC has a projected length of 31 km and would only be expected to reach collision energies of the order of TeV at a second stage of the collider's lifetime. CLIC could have a total length of around 40 km and would be operated at the TeV


Figure 1.3: Snapshots from a three-dimensional PIC simulation of the initial beam (left) and after the SMI has developed (right) [6].
range from the beginning. The cost estimate for the ILC, for instance, totals around €7.7 billion [25], so these projects will come at a high price.

Precisely with this dilemma in mind, the idea to use a proton beam instead of an electron beam as the driver in PWFA]was first advanced in 2009 [26]. Since the maximum energy that can be transferred to a witness electron is of the order of the driving beam energy, and since the proton bunches at the LHC carry more energy than any other driver currently available, the authors suggested that these proton bunches could drive the wakefield. The initial concept assumed that the LHC proton bunch could be compressed from its original length of 7.55 cm to around $100 \mu \mathrm{~m}$ (which is close to $\lambda_{p}$ ). This would be necessary because, the aim being to generate large-amplitude plasma waves, the maximum electric field that can be achieved in a linear regime is given by [26]

$$
\begin{equation*}
E_{\max }=240[\mathrm{MV} / \mathrm{m}] \frac{N}{4 \cdot 10^{10}}\left(\frac{0.6}{\sigma_{z}[\mathrm{~mm}]}\right)^{2} \propto \sigma_{z}^{-2} \tag{1.3}
\end{equation*}
$$

where $N$ is the number of particles in the driving bunch and $\sigma_{z}$ is therms length of the bunch (assuming a Gaussian profile).

Taking a proton energy of 1 TeV in a bunch with $10^{11}$ protons, a plasma density of about $10^{14} \mathrm{~cm}^{-3}$ and the number of electrons in the witness bunch to be around $10^{10}$ (which corresponds to an accelerated charge of approximately 1.6 nC ), the authors showed through numerical particle-in-cell (PIC) simulations that witness electrons initially with 1 GeV could be accelerated up to 0.62 TeV after 450 m .

In reality, a procedure to compress long, highly energetic proton bunches to the micron scale has not materialized. An alternative solution to this obstacle was proposed soon after, which was inspired by self-modulated LWFA [27]. This idea consists of letting the long proton beam undergo the selfmodulation instability (SMI) whereby the bunch progressively fragments into much smaller "beamlets" (with a length of the order of $\lambda_{p}$ ) due to periodic transversely focusing and defocusing regions (see Fig. 1.3). This instability eventually saturates. Through this mechanism the long proton beam is selfconsistently transformed to a format that can resonantly excite the plasma wave, similarly to PBWA,

There is another instability that competes with the SMI which occurs when the beam is not perfectly axisymmetric as it propagates in plasma: the hosing instability. This instability tends to disrupt the beam and is thus undesirable for the purpose of PWFA. However, it was soon suggested how to suppress this instability [27,28], by strongly seeding the SMI and by operating in a regime where the waves are linear, i.e. where the density fluctuations of the plasma wave are small compared to the background plasma density ( $\delta n / n_{0} \ll 1$ ). This last aspect keeps the hosing instability from dispersing the proton bunch even after the SMI has saturated. The seed for the SMI could be a very steep edge leading the density profile, for example by "severing" the beam in half, as suggested in [27, 29].

### 1.3 The AWAKE experiment

With the basic feasibility of the proposed proton-driven plasma wakefield accelerator established, an experiment has been launched at CERN to demonstrate the concept. The Advanced Proton Driven Plasma Wakefield Experiment (AWAKE) [31] (see [32]) is located at the former site of the CERN Neutrinos to Gran Sasso (CNGS) experiment (see Fig. 1.4) and had its first short data-taking run at the end of 2016. The experiment receives 400 GeV protons from the Super Proton Synchrotron (SPS), the last booster before the LHC, and directs them to a 10-meter-long plasma cell filled with rubidium vapor.

In order to seed the SMI, the approach of the sharp density profile edge was taken. However, it is currently not possible to shape high-energy proton bunches [33], so an ingenious setup was devised to both solve this problem and deal with the question of how to create the plasma (which has to be


Figure 1.4: A schematic image of the AWAKE experiment site and the surrounding accelerator complex at CERN (adapted from [30]).
extremely homogeneous). A short laser pulse is timed so that it travels with the center of the proton bunch, and as both enter the plasma cell the laser ionizes the rubidium gas, thus forming uniform plasma behind it. As such, the first half of the proton bunch propagates freely in the rubidium gas, while the second half propagates in the plasma generated by the laser pulse. In other words, the plasma only reacts to the second half of the proton bunch and thus a sharp rise in the density profile (see Fig. 1.5).

Once the proton bunch becomes fully self-modulated (after 4 m ), a 15 MeV electron beam will be injected into the plasma. The electron beam cannot be injected earlier because the phase velocity of the wake varies significantly while the instability is growing [34]. This would cause the beam electrons to slip into defocusing and decelerating phases of the wake, thus breaking up the electron beam. After saturation the phase velocity of the wake approaches the proton beam velocity [35], and the witness electrons can be injected at this point.

The parameters of the plasma and of both particle beams for the AWAKE experiment are listed in Table 1.1. As scientific goals, the AWAKE collaboration has chiefly proposed the study of the physics of self-modulation, the measurement of the accelerating wakefields through witness electrons for different injection and plasma conditions, the study of injection dynamics and of the production of GeV electrons with this accelerator, and the development of longer scalable plasma cells as well as setups for the production of plasma-accelerated protons [33].


Figure 1.5: Demonstrative arrangement of the proton and laser beams in AWAKE [6].

Table 1.1: Plasma and beam parameters of the AWAKE experiment [36].

| Parameter | Value |
| :--- | :---: |
| Plasma density $n_{0}$ | $7 \cdot 10^{14} \mathrm{~cm}^{-3}$ |
| Plasma length $L$ | 10 m |
| Proton bunch population $N_{b}$ | $3 \cdot 10^{11}$ |
| Proton bunch length $\sigma_{z b}$ | 12 cm |
| Proton bunch radius $\sigma_{r b}$ | $200 \mu \mathrm{~m}$ |
| Proton energy $W_{b}$ | 400 GeV |
| Proton bunch normalized emittance $\epsilon_{b n}$ | $3.5 \mu \mathrm{~m}$ |
| Electron bunch radius $\sigma_{r e}$ | $250 \mu \mathrm{~m}$ |
| Electron energy $W_{e}$ | 15 MeV |

### 1.4 State of the art

The SMI was thoroughly studied both analytically and in numerical PIC simulations for certain beam and plasma parameters [37-39]. In the real experiment, however, the proton beam parameters will inevitably vary from one SPS proton bunch to another. Similarly, the injected electron bunches and the plasma will not be identical from shot to shot. These variations can have detrimental effects on the AWAKE experiment and specifically on the development of the SMI. There has been research that investigated the role of shot-to-shot fluctuations of several experiment parameters. For instance, possible plasma inhomogeneities and their effect on the SMI and the witness electrons are well studied since they constituted a design constraint from the beginning (see for example [39]). There is also literature on the relationship between beam noise and the SMI seed, since the latter must be large enough to suppress the hosing instability [40], with the conclusion that noise from SPS proton bunches is very low and that there is therefore no significant constraint on the seed. The influence of proton beam emittance on the amplitude of the wakefield has also been analyzed [41].

Fluctuations of the characteristics of the proton bunches delivered by the SPS are naturally expected, and their role both on the development of the SMI and on the final state of the witness beam must be understood. An investigation into fixed variations of $\pm 5 \%$ in three proton bunch parameters (bunch population $N_{b}$, rms bunch length $\sigma_{z b}$ and rms bunch radius $\sigma_{r b}$ ) was begun by N. Savard [42]. This work will be expanded here by exploring the physics associated with these results in detail.

Besides the theoretical advances, the AWAKEcollaboration has reported the observation of strong modulation of high-energy proton bunches in December 2016, immediately after its first data-taking period [43] (see Fig. 1.6). The injection of the witness electron beam has not been tested yet.


Figure 1.6: Image of the longitudinal profile of a proton bunch traveling through plasma with distinct signs of modulation (from [43]).

### 1.5 Objectives and outline

Bearing in mind that AWAKErelies on a beam-plasma instability to succeed, it is vital to understand how sensitive the SMI is to variations of the initial parameters, as mentioned before. The goal of this thesis is precisely to contribute to this understanding. Showing that proton-driven PWFA is robust against initial perturbations despite relying on an instability to develop would be an interesting and important result for AWAKE.

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 in this work.

The objectives for this thesis and the chapters where they will be pursued are thus following:

- Chapter2: To establish a sound foundation for any conclusions that may be drawn in the following chapters by guaranteeing the reliability of the PICsimulations used throughout this thesis.
- Chapter3 To understand the influence of small perturbations to the initial configuration of the AWAKE experiment (by varying the properties of the driving proton bunch and the radius of the plasma channel) on the excited accelerating gradients and final energies attained by accelerated electrons.
- Chapter 4 To develop an advanced tool with the ability to provide some answers regarding the behavior of the wakefield after saturation of the SMI which is not fully understood at the moment.
- Chapter 5. To explore a hypothetically different regime of the AWAKE experiment where the driver particles are substituted by antiprotons, which will add to the current knowledge about the asymmetry between opposite charges when driving a wakefield.


### 1.6 Methods

The thesis work was conducted at the Group of Lasers and Plasmas (GoLP) of the Institute for Plasmas and Nuclear Fusion (IPFN) at Instituto Superior Técnico (IST). The PIC simulations were performed with the code OSIRIS [44].

### 1.6.1 Some notes about the PIC method and OSIRIS

OSIRIS is a relativistic, massively parallel PICcode developed by the Osiris Consortium, which is composed of IST and the University of California, Los Angeles (UCLA). The code can perform simulations of up to three dimensions and includes especially developed visualization and post-processing tools collectively designated as VisXD.

The first applications of the particle-in-cell method appeared during the 1950's and were related to the movement of electrons in vacuum tubes, though the now ubiquitous use of the method for plasma physics was introduced by Dawson and Buneman, among others [45, p. 3]. The general idea behind the PIC method consists of tracking individual particles on top of a stationary grid and computing all the physical quantities originated or influenced by those particles (charge density, current, electric and magnetic fields) at the grid points. This information is then used to interpolate the forces felt by each particle and push it across the grid accordingly, where the quantities at the grid points are once again updated. This rough sequence constitutes the computation loop in PICsimulations.


Figure 1.7: Diagram of the PIC loop implemented in OSIRIS. The $p$ index denotes a macroparticle, while the indices $i j$ refer to the grid coordinates. The orange and green boxes represent the particle pusher and the field solver, respectively.

Due to the high particle densities involved in typical plasma simulations, OSIRIS (and most PIC codes) does not represent each real particle individually, but instead distributes "macroparticles" over the grid, which represent an averaged ensemble of real particles.

The loop loosely described above is represented schematically in Fig. 1.7 as it is implemented in OSIRIS. The part of the iteration that solves the equations of motion for each macroparticle and updates its position (represented by the orange box in Fig. 1.7) is commonly called the particle pusher, which is based on the Boris scheme in OSIRIS. The field counterpart to this step (the green box in Fig. 1.7) is a finite difference solver for the electric and magnetic fields in space and time based on the Yee scheme [46].

The specification of the grid in a numerical simulation is not completely arbitrary. On one hand, the resolution of the grid must of course be adapted to the characteristic length of the expected phenomena (the grid must be able to resolve them). On the other hand, the cell dimensions are tied to the maximum time step that guarantees stability of the code through a Courant-Friedrichs-Lewy condition. In the case of electromagnetic leapfrog schemes [45, p. 355] and thus OSIRIS, this condition can be expressed in

Table 1.2: Normalization of physical quantities in OSIRIS (a hat denotes the normalized quantity).

| Quantity | Definition |
| :--- | :--- |
| Length $x$ | $\hat{x}=x k_{p}$ |
| Time $t$ | $\hat{t}=t \omega_{p}$ |
| Charge $q$ | $\hat{q}=\frac{q}{e}$ |
| Density $n$ | $\hat{n}=\frac{n}{n_{0}}$ |
| Velocity $v$ | $\hat{v}=\frac{v}{c}$ |
| Momentum $p$ | $\hat{p}=\frac{p}{m_{e} c}$ |
| Energy $W$ | $\hat{W}=\frac{W}{m_{e} c^{2}}$ |
| Electric/magnetic field $E / B$ | $[\hat{E} / \hat{B}]=\frac{[E / B]}{E_{0}}=[E / B] \frac{e}{c_{e} \omega_{p}}$ |

two dimensions as (in OSIRIS normalized units, see below):

$$
\begin{equation*}
\Delta t \leq \frac{1}{\sqrt{\frac{1}{\Delta x_{1}^{2}}+\frac{1}{\Delta x_{2}^{2}}}} \tag{1.4}
\end{equation*}
$$

where $\Delta t$ is the time step, and $\Delta x_{1}$ and $\Delta x_{2}$ are the cell lengths in both directions, given by $\Delta x_{i}=L_{i} / n_{i}$.
The computational task of following the paths of all the particles in a plasma (even if they are aggregated into macroparticles) is of course very demanding. This is why OSIRIS was designed to support parallel computing, where multiple computations are carried out simultaneously on different processing units. It has been tested successfully in some of the largest supercomputers in the world, using more than $10^{6}$ computing cores [47].

All the quantities in OSIRIS are normalized to plasma parameters or physical constants, which facilitates computation and generalizes the results. Some of these normalizations are summarized in Table 1.2, where $k_{p}$ is the plasma wavenumber, defined by $k_{p}=\omega_{p} / c, \gamma$ is the LoRENTZ factor given by $\gamma=1 / \sqrt{1-\beta^{2}}$ and $\beta$ is the ratio $\beta=v / c$, and $E_{0}$ is the nonrelativistic wavebreaking field defined in Eq. (1.1). The electric and the magnetic field have the same units since OSIRIS uses the Gaussian unit system. The independent variable which allows the conversion of OSIRIS quantities into physical ones is the plasma density $n_{0}$ (from which $\omega_{p}$ and $k_{p}$ can be calculated).

## Chapter 2

## Sensitivity analysis

The following chapter aims to assure the reliability of the results from a baseline simulation of the AWAKE experiment, so that further results from similiar simulations in this work can be trusted as well.

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) were already set, through a combination of reasonable estimation and computational compromise, and the purpose of the analysis was rather to establish whether the results were already safely within an area of convergence using these parameters. Even if this proves not to be the case, the analysis will have exposed weak points of the model and when they are most evident, which is in itself an important result.

### 2.1 Simulation and parameters

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 the same initial density profile as the proton bunch, and 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 2.1 in both OSIRIS units and physical units for a background plasma density of $3.5 \cdot 10^{14} \mathrm{~cm}^{-3}$. The resolutions which follow from the chosen lengths and cell numbers, through $\Delta x_{i}=L_{i} / n_{i}$, translate into about 75 and 333 cells per plasma wavelength in the axial and radial direction, respectively ( $\Delta x_{1} \approx 0.083 k_{p}^{-1}$ and $\Delta x_{1} \approx 0.019 k_{p}^{-1}$ ). The initial profile of the proton bunch $n_{b}$ was implemented according to:

$$
\begin{equation*}
n_{b}(\xi, r)=\frac{n_{b 0}}{2}\left[1+\cos \left(\sqrt{\frac{\pi}{2 \sigma_{z b}^{2}}}\left(\xi-\xi_{h}\right)\right)\right] \exp \left(-\frac{r^{2}}{2 \sigma_{r b}^{2}}\right), \quad \xi_{0} \leq \xi \leq \xi_{h} \tag{2.1}
\end{equation*}
$$

where $\xi$ is the beam co-moving coordinate given by $\xi=z-c t, n_{b 0}$ is the density at the peak of the profile (assumed here as $n_{b 0}=4 \cdot 10^{12} \mathrm{~cm}^{-3}$ ), $\sigma_{z b}$ is the rms bunch length (taken as $\sigma_{z b} \approx 17.8 \mathrm{~cm}$ ), $\sigma_{r b}$ the rms bunch width (taken as $\sigma_{r b} \approx 282.8 \mu \mathrm{~m}$ ), $\xi_{0}$ is the position where the function crosses the $\xi$ axis (in this simulation $\xi_{0}=1 k_{p}^{-1}$ ), and $\xi_{h}$ is the position of the beam head, where the ionizing laser pulse is co-propagating in the real experiment (in this simulation $\xi_{h}=1480 k_{p}^{-1}$ ), though it is not modelled
here. A plot of this density profile at the axis is shown in Fig. 2.1
The background plasma begins after the length of one simulation window, and does not cover its entire height, but is rather limited by a boundary $r_{p}(z)$ that grows linearly along the simulation. This linear growth is designed to emulate the way the plasma is formed in the AWAKE experiment by the ionizing laser, with a decreasing ionization radius as the pulse propagates through the plasma cell. As implemented in these simulations, the radii at the entrance and at the exit of the plasma cell (i.e. at $z_{0}=1492 k_{p}^{-1}$ and $z_{f}=51243.2 k_{p}^{-1}$, respectively) are $r_{p}\left(z_{0}\right)=212.1 \mu \mathrm{~m}$ and $r_{p}\left(z_{f}\right)=141.4 \mu \mathrm{~m}$.

Though the physical parameters chosen for these simulations do not correspond entirely to the AWAKEnominal parameters, the sensitivity analysis performed here will still be valid, as all the following simulations in this work closely resemble the baseline used here. The numerical parameters chosen for the analysis were the resolution in the $x_{1}$ direction, the resolution in the $x_{2}$ direction and the number of numerical particles per cell of the proton beam species, henceforth signified by $\mathrm{ppc}_{p^{+}}$.

The number of particles per cell is given in the input file as a vector with the size of the dimensions of the simulation, in this case two (in the standard simulation $\mathrm{ppc}_{p^{+}}=(2,2)$ ). The elements of the vector specify the number of macroparticles per cell and per corresponding direction, so the total number of particles in each cell is the product of the vector components (macroparticles are the numerical representation of an aggregate of real particles, as mentioned in section 1.6.1). This parameter was tested by increasing the number of particles in each direction by the same factor.

The resolutions were varied by both increasing and decreasing the number of cells in each direction. A change in resolution furthermore implies a change in the time step through the Courant-FriedrichsLewy condition from Eq. 1.4 . A summary of the simulations performed and their respective parameter variations is presented in Table 2.2 .

### 2.2 Results

OSIRIS offers an optional set of diagnostics as output from the simulations besides the raw macroparticle data, which includes the distributions of all three spatial components of the electric and magnetic fields and of the charge density of each particle species, as well as spatial and temporal averages of these quantities. The most relevant quantities for displaying the phenomena expected in AWAKE are the axial and the radial electric field components $E_{z}$ and $E_{r}$ (called $E_{1}$ and $E_{2}$ in OSIRIS), respectively, and the charge density of the proton beam $\rho_{p^{+}}$, so the distributions of these quantities (which can be considered direct simulation results as opposed to post-processed ones) were used for comparing the runs described above.

In addition to the spatial distributions of these quantities, however, further diagnostics were developed to evaluate the results from the simulations throughout this entire work. These diagnostics are described in detail in Appendix A, though in this particular chapter only the ones described in sections A. 1 through A. 4 were used. These diagnostics produce indirect, or post-processed, results from the simulation data.

### 2.2.1 Direct results

A direct comparison between the resulting distributions of the axial and radial electric fields $E_{1}$ and $E_{2}$ and of the proton charge density $\rho_{p^{+}}$was performed at three different regions along the beam (always with a width of $\Delta x_{1}=45 k_{p}^{-1}$ ) and at three simulation times, specified in Table 2.3. The comparison times were chosen due to the fact that they are common multiples of the different time steps, with the exception of the run with decreased resolution in $x_{1}$. In this case the closest simulation time was taken for comparison, which is denoted by a prime in Table 2.3 ( $t^{\prime}$ and $z^{\prime}$ ).

Table 2.1: General parameters of the baseline simulation for the sensitivity analysis.

| Parameter | Normalized units | Physical units <br> $\left(n_{0}=3.5 \cdot 10^{14} \mathrm{~cm}^{-3}\right)$ |
| :--- | :---: | :---: |
| Length of simulation box, $L_{1}$ | $1492 k_{p}^{-1}$ | 42.41 cm |
| Height of simulation box, $L_{2}$ | $8 k_{p}^{-1}$ | 2.28 mm |
| Number of cells in $x_{1}$ direction, $n_{1}$ | 18000 | - |
| Number of cells in $x_{2}$ direction, $n_{2}$ | 425 | - |
| Time step | $0.012 \omega_{p}^{-1}$ | $3.41 \mu \mathrm{~m}(\mathrm{at} c)$ |
| Propagation distance | $53000 k_{p}^{-1}$ | 15.06 m |



Figure 2.1: Plot of the density profile from Eq. (2.1) for $r=0$.

Table 2.2: Summary of the runs performed for the sensitivity analysis with each respective parameter variation.

| Parameter | Factor | Altered quantities |
| :---: | :---: | :---: |
| $x_{1}$ resolution | 1.5 (increase) | $\begin{gathered} n_{1}=27000 \\ \Delta t=0.017 \end{gathered}$ |
| $x_{2}$ resolution | 1.5 (increase) | $\begin{gathered} n_{2}=638 \\ \Delta t=0.012 \end{gathered}$ |
| $x_{1}$ resolution | $1 / \sqrt{2}$ (decrease) | $\begin{aligned} & n_{1}=12728 \\ & \Delta t=0.013 \end{aligned}$ |
| $x_{2}$ resolution | 1/ $\sqrt{2}$ (decrease) | $\begin{gathered} n_{2}=301 \\ \Delta t=0.018 \end{gathered}$ |
| particles per cell | $1.5 \times 1.5=2.25$ (increase) | $\mathrm{ppc}_{p^{+}}=(3,3)$ |

Table 2.3: Regions of the simulation window and simulation times chosen to compare the results from the sensitivity analysis.

| Regions |  | Simulation times |  |
| :---: | :---: | :---: | :---: |
| head | $x_{1}=[1435,1480] k_{p}^{-1}=[40.79,42.07] \mathrm{cm}$ | $t_{1}$ | $\begin{array}{ll} t=3179.99 \omega_{p}^{-1}, & z=0.91 \mathrm{~m} \\ \left(t^{\prime}=3444.99 \omega_{p}^{-1},\right. & \left.z^{\prime}=0.96 \mathrm{~m}\right) \end{array}$ |
| middle | $x_{1}=[721,766] k_{p}^{-1}=[20.49,21.78] \mathrm{cm}$ | $t_{2}$ | $\begin{array}{ll} t=28619.89 \omega_{p}^{-1}, & z=8.13 \mathrm{~m} \\ \left(t^{\prime}=28708.22 \omega_{p}^{-1},\right. & \left.z^{\prime}=8.16 \mathrm{~m}\right) \end{array}$ |
| tail | $x_{1}=[21,66] k_{p}^{-1}=[0.59,1.88] \mathrm{cm}$ | $t_{3}$ | $\begin{gathered} t=50879.81 \omega_{p}^{-1}, \quad z=14.47 \mathrm{~m} \\ \left(t^{\prime}=50526.48 \omega_{p}^{-1}, \quad z^{\prime}=14.37 \mathrm{~m}\right) \end{gathered}$ |



Figure 2.2: Three exemplary comparisons of the results from the sensitivity analysis runs in the "head" region: the radial electric field $E_{2}$ at $t_{1}$ for a decreased $x_{1}$ resolution (left), the axial electric field $E_{1}$ at $t_{2}$ for an increased $x_{2}$ resolution (middle), and the proton charge density $\rho_{p}+$ at $t_{3}$ for an increased number of particles per cell of the proton species (right).


Figure 2.3: Comparison of the distribution of the radial electric field $E_{2}$ in the "middle" region for variations of the resolution in $x_{2}$ at three simulation times.

Considering the distributions of the three quantities across the three probing times, it is noticeable that barely any differences exist in the "head" region of the beam in general. Fig. 2.2 exemplifies this fact through the comparisons of $E_{1}, E_{2}$ and $\rho_{p^{+}}$at $t_{1}, t_{2}$ or $t_{3}$ for three different parameter variations. 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.

At the beginning of the simulation the wakefield is still growing and its amplitude can be comparable to the amplitude 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. This is exemplified in Figs. 2.3 and 2.4 by the radial electric field $E_{2}$ in the middle region of the beam and for different resolutions of $x_{2}$ and $x_{1}$, respectively. In both figures some differences are perceptible at $t_{1}$. As the fields grow, however, these differences virtually disappear for variations in the number of particles per cell and in $x_{2}$ resolution, though only the latter is shown here, again in Fig. 2.3. The sensitivity of the results towards variations in $x_{1}$ resolution is higher, as can be inferred from the different shapes of the fields in Fig. 2.4 though the results of the baseline run are quite close to the ones from the run with increased resolution (especially for later simulation times).

As might be expected, the tail of the beam is even more noisy and its field amplitudes lower than the middle at early times. The run with an increased number of particles per cell displays the most pronounced differences in such regions, where there is more noise. As shown in Fig. 2.5. this parameter change effectively annihilates all noise even in the "tail" region. After $t_{1}$ the noise mostly disappears and so do the differences between these runs.

Regarding the resolution in $x_{1}$, the axial field $E_{1}$ in the tail region (see Fig. 2.6) displays a similar behavior to the field $E_{2}$ in the middle region: the baseline fields agree reasonably well with the increased-resolution fields (especially at $t_{3}$ ), once again suggesting a healthy degree of convergence towards correct results despite the higher sensitivity to this parameter. Nonetheless, the amplitude of the field at $t_{2}$ differs visibly between these two runs.

The effects on the radial electric field $E_{2}$ in the same region are more significant. In Fig. 2.7. most prominently at $t_{2}$, differences both in amplitude and field shape are blatantly observable. Keeping with the assumption that greater resolution leads to more accurate results, the run with an increase in $x_{1}$ resolution seems to suggest that the radial electric field is in reality slightly more powerful and, at $t_{2}$, of a different shape, though towards the end of the simulation the results converge moderately well. In general the sensitivity of the fields with respect to the resolution in $x_{1}$ seems to decrease for later simulation times (at $t_{3}$ ).

The incomplete convergence towards correct results w.r.t. (with respect to) the resolution in the $x_{1}$ direction does not manifest itself only in the field distributions, but also in the distribution of the proton charge density $\rho_{p^{+}}$. Fig. 2.8 shows the sizable differences at $t_{2}$, and in this case they do not seem to converge as well as the fields at $t_{3}$.

### 2.2.2 Indirect results

The peak and average axial electric fields along the simulation (or propagation distance) were calculated using the diagnostic described in section A.1, whereby the first few rows of data from the axis were excluded from the analysis (a distance of 5 cells in the $x_{2}$ direction, except for the runs with increased and decreased resolution in $x_{2}$, where the distance was 8 and 4 cells, respectively). This is done in


Figure 2.4: Comparison of the distribution of the radial electric field $E_{2}$ in the "middle" region for variations of the resolution in $x_{1}$ at three simulation times.


Figure 2.5: Comparison of the distribution of the axial electric field $E_{1}$ in the "tail" region for an increased number of particles per cell of the proton species at $t_{1}$ and $t_{2}$.


Figure 2.6: Comparison of the distribution of the axial electric field $E_{1}$ in the "tail" region for variations of the resolution in $x_{1}$ at $t_{2}$ and $t_{3}$.
$E_{2}$, tail


Figure 2.7: Comparison of the distribution of the radial electric field $E_{2}$ in the "tail" region for variations of the resolution in $x_{1}$ at $t_{2}$ and $t_{3}$.


Figure 2.8: Comparison of the distribution of the proton charge density $\rho_{p^{+}}$in the "middle" region for variations of the resolution in $x_{1}$ at $t_{2}$ and $t_{3}$.


Figure 2.9: Peak axial electric field $E_{1}$ versus the propagation distance for all the sensitivity analysis runs.
order to exclude noisy data and unphysically high values that appear close to the axis boundary.
The result for the peak electric field can be seen in Fig. 2.9 , where the previously observed high sensitivity to the resolution in $x_{1}$ is confirmed (the largest differences with respect to the baseline results happen for variations of this parameter). There is an appreciable distance between the values of the baseline simulation and the run with increased $x_{1}$ resolution, especially around the maxima of the curves.

The average electric field, shown in Fig. 2.10 shows a similar trend. It seems that the "true" (or more accurate) values of the axial field are slightly lower before the peak (up to 9\%) and their maximum is slightly shifted forwards by around $60 \mathrm{~cm}(6 \%)$. This means that the baseline simulation, a close form of which will be used throughout the rest of this work, delivers slightly overestimated values for the amplitude of $E_{1}$ up to the peak, and underestimated values after the curve begins to fall again (at most by approximately 8\%).

The diagnostic elucidated in section A.2 joins the line-outs of the axial field $E_{1}$ taken at every file dump, so that the behavior of the phase velocity of the wake w.r.t. the simulation window, which is moving at $c$, can be visualized (the result of the diagnostic is shown in Fig. 2.11 . In this case the line-outs composing the concatenated matrix were again taken at a distance of 5 cells from the axis

Average $I E_{1} \mid$


Figure 2.10: Average axial electric field $E_{1}$ versus the propagation distance for all the sensitivity analysis runs.


Figure 2.11: Comparison of the concatenated line-outs of $E_{1}$ close to the axis over the propagation distance in the "middle" and "tail" regions (top and bottom row, respectively) and for a decreased and an increased $x_{1}$ resolution (middle and right column, respectively).
for most simulations, except for the runs with decreased and increased resolution in $x_{2}$, where the distance was of 4 and 8 cells, respectively (for the same reasons as stated above).

The phase velocity $v_{\phi}$ can be estimated by following the trajectory of a fixed phase (i.e. a position along the field wave where the value is constant through time, for example where $\left.E_{1}(\xi)=0\right)$ and measuring its local slope. Imagining that the $z$ axis becomes the horizontal axis (rotating the graph $90^{\circ}$ clockwise), a positive slope will indicate that $v_{\phi}<c$, a negative slope that $v_{\phi}>c$, and a zero-slope that $v_{\phi}=c$ (this is explored in more detail in section A.2).

The results yielded by this diagnostic are practically indistinguishable for most of the sensitivity analysis runs. This is especially true at the head of the beam, a region which was unlikely to show any differences in any case, as was noted before. On the other hand, perceptible differences arise once again for variations of the $x_{1}$ resolution.

Upon close scrutiny of the results in the "middle" region, in the top row of Fig. 2.11, the trends of the phase velocity at 15 m (the end of the simulation) prove quite different for the three runs. In the baseline simulation the phase velocity appears to be very slightly superluminal and increasing (an observation that is supported by a clearer version of this trend in the decreased-resolution run), while in the increased resolution run it appears to be stagnating, and is not superluminal at any point during the simulation. Future statements about the superluminosity of $v_{\phi}$ in the middle region at the end of the simulation should therefore be avoided.

In the zone of saturation of the SMI however, the phase velocities of the baseline and the increasedresolution run are extremely close. Between 3.5 and 7 m the LORENTz factor corresponding to the phase velocity coincides to four decimal places ( $\gamma_{\phi} \approx 1.0008$ ).

In the "tail" region of the window the differences are more subtle between the baseline and the increased-resolution version, where there are no doubts that the phase velocity is superluminal and increasing around 10 m . Between 10 and 13 m , however, the average $v_{\phi}$ is lower (though very slightly, with a difference of $0.2 \%$ ) in the run with increased resolution, though the final trend points to a faster growth than the one observed in the baseline run (as evidenced by the curved slope as opposed to the


Figure 2.12: Total charge of the proton species below a radius of one plasma skin depth over the propagation distance and for all sensitivity analysis runs.
almost linear slope in the baseline case).
The total charges of both the proton species and the witness beam species (the diagnostic for which is described in section A.3) were also used to compare the results from the sensitivity analysis.

As the SMI develops, some portions of the proton bunch will be defocused from the axis and eventually leave the simulation box. Since it takes some time for the defocused protons to reach the border of the box and leave it, the total charge calculated for the whole box would only indicate some loss after a certain propagation distance, so only a portion of the protons below a certain radius were taken into account in the calculation of the total charge. The value chosen for this radial threshold was one plasma skin depth, or the inverse of the plasma wavenumber $k_{p}$. For a plasma density of $3.5 \cdot 10^{14} \mathrm{~cm}^{-3}$, $k_{p}^{-1}$ is around 0.3 mm . In OSIRIS units this radius corresponds to $x_{2}=1 k_{p}^{-1}$ (see Fig. 2.8 for reference, for example). In the case of the protons the charge density data was taken as input.

The result, shown in Fig. 2.12, suggests that the total proton charge is underestimated in the baseline run (especially between 6 and 15 m , reaching a discrepancy of $10 \%$ ), though the different curves converge at the end of the simulation. There is a lower parameter sensitivity to the $x_{2}$ resolution, though a divergence is still visible.

The total charge of the witness beam was summed over the entire simulation window using the raw macroparticle data. Though Fig. 2.13 would seem to contradict this at first glance, here the results from all the different runs are in very good agreement, to within $3 \%$ (note the scale). A graph spanning the entire range of the electron charge data would display a sudden drop from around 2.7 pC to the values seen in Fig. 2.13 before 1.5 m , and an almost indiscernible and flat superposition of the different lines from then on.

A graph of the maximum electron energy in the witness beam along the simulation is shown in Fig. 2.14 , as described in section A.4. Bearing in mind that this diagnostic takes the energy values from one single macroparticle at each dump, and is thus not very representative of phenomena on a larger scale, there is moderate agreement between the curves of the different runs. At around 7 m all the sensitivity analysis runs seem to indicate a higher value than the baseline run, so it seems that this simulation is not capturing the more realistic energy value. Other than this, no trend of convergence can be identified clearly.


Figure 2.13: Absolute values of the total charge of the witness beam species over the propagation distance and for all the sensitivity analysis runs (amplified scale).


Figure 2.14: Maximum energy reached by a witness electron versus propagation distance for all sensitivity analysis runs.

### 2.3 Conclusion

The most obvious conclusion to be extracted from this parameter sensitivity analysis is that all the simulation results (both direct and indirect) were most sensitive to the resolution in the $x_{1}$ direction. The direct outputs of the simulations evinced no mentionable differences w.r.t. the other parameters and a reasonable degree of convergence with the results from the runs with increased $x_{1}$ resolution (which improves for later simulation times).

The indirect results displayed greater sensitivities to both resolutions (in $x_{1}$ and $x_{2}$ ), and the values did not converge as well (occasionally at all) to the refined parameters. Nonetheless, the situations where these deviations exist were identified and the magnitude of the deviations determined, which at least enables the estimation of a margin of error when drawing conclusions about the quantities affected.

It should be mentioned that the latter result is particularly relevant when deducing absolute values from the simulations or their post-processed results, of course. However, most objectives set out for this thesis are based on relative results between different runs, where the findings from this chapter are of no great consequence.

Having established a certain level of reliability and exposed the weak points regarding simulations of AWAKE it is now safe to proceed to the main objectives of this work.

## Chapter 3

## Parameter scans of the AWAKE simulation

This chapter pursues one of the main goals of the present thesis, which is to determine how expectable variations of the initial parameters of the AWAKE experiment can affect the experiment's outputs, most notably the wake characteristics and the final energy of hypothetical accelerated electrons. Shot-toshot 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.

For this to happen it is critical that the wakefield phase does not change much from shot-to-shot. The reason why the wakefield phase matters can be elucidated by drawing an analogy between the RF cavities in a conventional circular accelerator and the wakefield troughs (more precisely, half-troughs) in a plasma accelerator. In conventional accelerators the electric fields in RF cavities oscillate with a precise and well-defined frequency and the positions and dimensions of the cavities themselves are known, which means that it is possible to know where and when it is best to insert particles (in socalled buckets). In a plasma accelerator, however, the "cavities" (field troughs) may change positions, which corresponds to a variation in the wake's phase velocity or, equivalently, a phase shift.

Besides guaranteeing that electrons can always be injected at the same position and reach the desired final state despite expectable fluctuations, then, it will be important to assure the stability of the wakefield's amplitude as well, for beam-quality purposes. The amplitude determines the energy gain of the accelerated particles, so if it varies considerably between shots the accelerated beam will also have different energies between shots, which is generally not desired from a particle accelerator. Beam-quality is nevertheless of lower priority for the AWAKE experiment, whose purpose is first and foremost to prove a concept.

### 3.1 Simulation and parameters

The standard simulation used for the comparison of the parameter scans is slightly different than the baseline one described in section 2.1 and in Table 2.1. 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. Additionally, the
plasma radius is kept constant in these runs, i.e. $r_{p}(z)=r_{p}$ ( $150 \mu \mathrm{~m}$ in the standard run).
The values for the beam dimensions used in these simulations are closer to AWAKEs nominal parameters, with the rms length $\sigma_{z b}=12.6 \mathrm{~cm}$ and width $\sigma_{r b}=200 \mu \mathrm{~m}$. The proton beam profile is described by the same equation as before, that is,

$$
\begin{equation*}
n_{b}(\xi, r)=\frac{n_{b 0}}{2}\left[1+\cos \left(\sqrt{\frac{\pi}{2 \sigma_{z b}^{2}}}\left(\xi-\xi_{h}\right)\right)\right] \exp \left(-\frac{r^{2}}{2 \sigma_{r b}^{2}}\right), \quad \xi_{0} \leq \xi \leq \xi_{h} \tag{3.1}
\end{equation*}
$$

though the positions of the beam head and the function's zero are shifted: here $\xi_{h}=1651 k_{p}^{-1}$ and $\xi_{0}=80.2 k_{p}^{-1}$. The peak density $n_{b 0}$ also has a different value here, calculated from the bunch dimensions mentioned above and the fixed bunch population $N_{b}=1.5 \cdot 10^{11}$ through the equation

$$
\begin{equation*}
n_{b 0}=\frac{N_{b}}{(2 \pi)^{2 / 3} \sigma_{r b}^{2} \sigma_{z b}}, \tag{3.2}
\end{equation*}
$$

giving $n_{b 0} \approx 1.89 \cdot 10^{12} \mathrm{~cm}^{-3}$. In all other respects the standard simulation here is identical to the baseline one described in the last chapter. The differing simulation parameters are displayed in bold in Table 3.1 , where the physical values of the remaining parameters for a plasma density of $7 \cdot 10^{14} \mathrm{~cm}^{-3}$ have also been included.

The following experiment parameters were varied for this study:

- rms proton bunch length $\sigma_{z b}( \pm 5 \%)$,
- rms proton bunch width $\sigma_{r b}$ ( $\pm 5 \%$ ),
- proton bunch population $N_{b}( \pm 5 \%)$,
- plasma radius $r_{p}$ ( $\pm 5 \%$ ),
- rmstiming jitter of the proton bunch with respect to the ionizing laser pulse $\Delta t$ ( $\pm 15 \mathrm{ps}$ ).

Equation (3.2) 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. The bunch population (or the driver charge, equivalently), however, is a parameter that is to be analyzed independently. Hence, in order to avoid mixing variations of two parameters, the density $n_{b 0}$ was altered as well in runs with variations of the bunch dimensions so as to conserve the charge and the bunch population $N_{b}$. On the other hand, the timing jitter $\Delta t$ is in practice a phase shift of the cosine in Eq. (3.1) with respect to the position of the beam head $\xi_{h}$, thus encapsulating either more or less charge depending on whether the maximum of the cosine is moved to the right or left of $\xi_{h}$. This shift will be denoted by an additional term $\Delta \xi$ in the cosine argument from Eq. (3.1):

$$
\begin{equation*}
\cos \left(\sqrt{\frac{\pi}{2 \sigma_{z b}^{2}}}\left(\xi-\xi_{h}\right)\right) \quad \longrightarrow \quad \cos \left(\sqrt{\frac{\pi}{2 \sigma_{z b}^{2}}}\left(\xi-\xi_{h}-\Delta \xi\right)\right) \tag{3.3}
\end{equation*}
$$

Besides a slight change of the bunch's profile, the timing jitter thus concomitantly translates into an increased or decreased amount of charge driving the wakefield, which is equivalent to a change of the overall bunch population $N_{b}$.

The concrete values for each parameter variation are listed in Table 3.2. As an example of what the accelerating mechanism at AWAKE looks like, Fig. 3.1 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. 3.1 , 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 $\left(\lambda_{p} / 4\right)$ long. The highlighted area in Fig. 3.1, in green, represents one of these regions.

Table 3.1: General simulation parameters of the standard simulation for the parameter scans. The altered parameters (w.r.t. Table 2.1) are shown in boldface.

| Parameter | Normalized units | Physical units <br> $\left(n_{0}=7 \cdot 10^{14} \mathrm{~cm}^{-3}\right)$ |
| :--- | :---: | :---: |
| Length of simulation box, $L_{1}$ | $\mathbf{1 6 6 3} \boldsymbol{k}_{p}^{-1}$ | 33.43 cm |
| Height of simulation box, $L_{2}$ | $8 k_{p}^{-1}$ | 1.61 mm |
| Number of cells in $x_{1}$ direction, $n_{1}$ | $\mathbf{2 0 0 6 3}$ | - |
| Number of cells in $x_{2}$ direction, $n_{2}$ | 425 | - |
| Time step | $0.012 \omega_{p}^{-1}$ | $2.41 \mu \mathrm{~m} \mathrm{(at} c)$ |
| Propagation distance | $53000 k_{p}^{-1}$ | 10.65 m |

Table 3.2: Experiment parameters chosen for the analysis, their respective variations, and the consequently altered simulation inputs w.r.t. the standard simulation.

| Parameter | Variation | Altered simulation inputs |  |  |
| :---: | :---: | :--- | :--- | :--- |
| $\sigma_{z b}$ | $+5 \%$ | $\sigma_{z b}=13.23 \mathrm{~cm}$ | $n_{b 0} \approx 1.80 \cdot 10^{12} \mathrm{~cm}^{-3}$ | $\xi_{0}=1.0 k_{p}^{-1}$ |
|  | $-5 \%$ | $\sigma_{z b}=11.97 \mathrm{~cm}$ | $n_{b 0} \approx 1.99 \cdot 10^{12} \mathrm{~cm}^{-3}$ | $\xi_{0}=157.85 \mathrm{k}_{p}^{-1}$ |
|  | $+5 \%$ | $\sigma_{r b}=210 \mu \mathrm{~m}$ | $n_{b 0} \approx 1.71 \cdot 10^{12} \mathrm{~cm}^{-3}$ |  |
|  | $-5 \%$ | $\sigma_{r b}=190 \mu \mathrm{~m}$ | $n_{b 0} \approx 2.09 \cdot 10^{12} \mathrm{~cm}^{-3}$ |  |
| $N_{b}$ | $+5 \%$ | $n_{b 0} \approx 1.98 \cdot 10^{12} \mathrm{~cm}^{-3}$ |  |  |
|  | $-5 \%$ | $n_{b 0} \approx 1.80 \cdot 10^{12} \mathrm{~cm}^{-3}$ |  |  |
| $r_{p}$ | $+5 \%$ | $r_{p}=157.5 \mu \mathrm{~m}$ |  |  |
|  | $-5 \%$ | $r_{p}=142.5 \mu \mathrm{~m}$ |  |  |
| $\Delta t$ | +15 ps | $\Delta \xi \approx-22.39 k_{p}^{-1}$ | $\xi_{0}=57.81 k_{p}^{-1}$ |  |
|  | -15 ps | $\Delta \xi \approx 22.39 k_{p}^{-1}$ | $\xi_{0}=102.59 k_{p}^{-1}$ |  |



Figure 3.1: 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 w.r.t. the head of the beam.

### 3.2 Results

Further post-processing diagnostics were developed to evaluate the parameter scans specifically. Besides the averaging of the electric field from section A. 1 , two further post-processing tools were employed: one which determines the local phase shift of the wake w.r.t. its head (described in Appendix A section A.5), and one which tracks the axial movement and energy gain of externally injected electrons (in section A.6).

### 3.2.1 Properties of the wake

The effects of the initial parameter variations on the amplitude and on the phase of the wakefield will be presented here. As stated in the introduction to this chapter, the stabilities of the amplitude and phase in the face of natural driver bunch fluctuations are a necessary condition for a reliable version of a plasma accelerator.

The results for the average amplitude of the axial electric field $E_{1}$ are shown for all the runs in Fig. 3.2. The axial field reaches a maximum at roughly 5 m , where the effects from the parameter variations overall are also maximized. As the average amplitude decreases, these effects seem to follow suit and converge towards fixed variations at the end of the propagation distance. The plasma radius $r_{p}$ is a noteworthy parameter here since it hardly influences the results of this diagnostic throughout the entire length of the simulation.

Two groups of experiment parameters can be identified according to their qualitative effects on the axial field. Variations of the rms beam dimensions $\sigma_{z b}$ and $\sigma_{r b}$ cause inversely proportional variations of $<\left|E_{1}\right|>$, while the bunch population $N_{b}$, the substantially equivalent timing jitter $\Delta t$ and the plasma radius $r_{p}$ (at a much smaller scale) are proportional to $<\left|E_{1}\right|>$. Linear wakefield theory predicts that $E_{z} \propto n_{b 0}$ (see Eq. (4.1) on p. 37 , where $n_{b 0}$ 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. (3.2) $n_{b 0} \propto N_{b} \sigma_{r b}^{-2} \sigma_{z b}^{-1}$.

It should be noted that, although $\sigma_{r b}$ should have a larger effect on the axial field $E_{1}$ than $\sigma_{z b}$, this is only observed until the maximum of the average axial field. Before the peak, the rms bunch width $\sigma_{r b}$ has the greatest influence on $<\left|E_{1}\right|>$, reaching variations of $\pm 10 \%$ w.r.t. the standard run at 5 m . At the peak these effects begin to decrease until they are comparable with those of the $\sigma_{z b}$ variations.


Figure 3.2: Average amplitude of $E_{1}$ versus the propagation distance for all parameter scans.

At the same time the effects due to changes of $N_{b}$ become dominant, with a maximum of around $\pm 8 \%$ at 5.3 m .

The timing jitter of the proton bunch is varied using absolute values ( $\pm 15 \mathrm{ps}$ ), which renders the correlation between variation and effect less directly perceivable than for the remaining parameters. Excluding possible effects due to the different shape of the proton bunch density profile, the $\Delta t$ variations correspond to differences of $\pm 2.85 \%$ in the amount of charge driving the wake (or, equivalently, in the bunch population). In the case of this diagnostic, the maximum observed effect is of roughly $\pm 3 \%$ at 5.5 m , which provides some notion about the influence of this parameter.

At the end of the plasma cell, or at 10.33 m (since the plasma medium only begins after a 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 bunch population $N_{b}$. 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.

The behavior of the wake's phase is even more important than the wakefield amplitude. If the phase varies too much from shot to shot there can be no deterministic injection of electrons, which would render a particle accelerator based on the AWAKEconcept nonviable.

The approach to determine the local phase shift of the wake w.r.t. its head along time consists of taking a line-out of $E_{z}$ close to the axis at each dump, dividing the line-out into small regions (each encapsulating a few wavelengths), and fitting the function

$$
\begin{equation*}
f(\xi)=A \sin \left[k_{p}\left(\xi-\xi_{h}\right)+\phi\right] \tag{3.4}
\end{equation*}
$$

where $A$ and $\phi$ are the fitting parameters, to the data in each small region (or fitting window). The value of the parameter $\phi$ will then indicate how much that particular region has shifted w.r.t. the sinusoidal wake beginning at the head of the beam $\left(\xi_{h}\right)$. Further details about this diagnostic are supplied in section A. 5 of Appendix A

In this case, a number of 2.5 plasma wavelengths $\lambda_{p}$ per fitting window was chosen, generating the phase shift distribution shown in Fig. 3.3. The results are for the standard run and the beam coordinate


Figure 3.3: Local phase shift of $E_{1}$ close to the axis versus the propagation distance $z$ and the position along the beam w.r.t. to its head $\xi-\xi_{h}$, for the standard simulation. The position along the beam where all the line-outs were taken for comparison is indicated by the light gray dashed line.


Figure 3.4: Local phase shift of the axial component of the wake 11.72 cm behind the head of the beam versus the propagation distance for all parameter scans. A highlight of a portion of the graph is included on the right.
has been represented in reference to the position of the laser pulse or beam head, $\xi_{h}$. The comparison between the results of the parameter scans was done for a line-out of the phase shift along $z$ at a fixed position along the beam which is approximately 11.7 cm behind the head (see Fig. 3.3 , where the line-out is represented by a light gray dashed line).

The line-outs from all the runs are shown in Fig. 3.4 . There is a kind of plateau around 3.5 m which, judging by its movement along the beam between 2 and 4 m of propagation (see Fig. 3.3), is linked to the saturation of the SMI. Before and after this plateau region no clear effects from the parameter variations are discernible, either because the variations are negligibly small (before the plateau) or because their magnitudes are within the level of noise in the data (after the plateau).

At the plateau, however, the same tendencies as before can be identified: in this case the bunch dimensions $\sigma_{z b}$ and $\sigma_{r b}$ are proportional to the phase shift, and the bunch population $N_{b}$ and the timing jitter $\Delta t$ are inversely proportional, although these tendencies seem to be reversed after the plateau (this is seen more clearly in line-outs taken further behind the head, which are not shown here). The plasma radius $r_{p}$ once again seems to exert very little effect on the characteristics of $E_{z}$.

The largest variations found in this sample are for therms bunch width $\sigma_{r b}$, which would make sense if the phase shift was assumed to be $\phi \propto E_{z}^{-1} \propto \sigma_{r b}^{2}$ at least before saturation. The phase is around $5 \%$ larger at 3.5 m and $5 \%$ smaller at 3.2 m for the $\sigma_{r b}$ runs.

As mentioned before, at the end of the simulation hardly any differences can be detected, possibly due to the amount of noise. In any case, the relative magnitude of these effects must be below the initial percentage variations of the parameters. This means that deterministic injection is indeed possible for AWAKEdespite initial parameter variations, as long as it is performed after saturation of the SMI Overall the properties of the wake are robust against shot-to-shot fluctuations of the driving proton bunch, and they meet the requirements for successful proof-of-principle experiments at AWAKE.

### 3.2.2 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. This tool is based on the output of another diagnostic developed in this work: the concatenated line-outs of $E_{z}$ close to the axis (see section A.2, where this
diagnostic is described, or Fig. 2.11 on p. 19).
The purpose of this diagnostic is to know, for a certain specified final energy and beam coordinate, where an electron would have to be injected along the propagation distance so as to attain that final energy and without being decelerated at any point, while additionally taking into account the possible dephasing of the electron w.r.t. the wake (longitudinally). One of the desired outputs is the energy gain obtained by the electron.

This is accomplished by "positioning" an electron at $z=10.33 \mathrm{~m}$ (the end of the plasma cell) and at a certain $\xi_{f}$ (some final position along the wake) with a certain final energy $\gamma_{f}$, and propagating it backwards in time through a leapfrog-type relativistic particle pusher until the electron meets a cell where the electric field is positive (i.e., where it would lose energy). This cell then determines the minimum injection point $z_{\mathrm{inj}}$, since before this propagation distance the electron would have to lose energy. An alternative condition to find $z_{\text {inj }}$ was imposed due to the currently planned electron injection energy for AWAKE, which is 15 MeV [36]. This corresponds to a minimum Lorentz factor threshold $\gamma_{\text {min }}$ of 29. The energy gain $\Delta \gamma$ is then obtained from the difference between the imposed final energy $\gamma_{f}$ and the initial energy $\gamma_{0}$ arrived at by the particle pusher, i.e. $\Delta \gamma=\gamma_{f}-\gamma_{0}$. The diagnostic is described in more detail in section A. 6 of Appendix A.

An example of a propagated electron using this diagnostic is shown in Fig. 3.5. Its trajectory w.r.t. the concatenated $E_{z}$ line-outs is shown in yellow, whereby the algorithm begins at the filled circle and ends at the unfilled one. Note that the electron trajectories can only run backwards or stay constant in $\xi$, since the electrons would otherwise be traveling faster than the speed of light (both the simulation window and the beam coordinate $\xi$ move at $c$ ). The significant variables $\xi_{f}$ (input), $z_{\mathrm{inj}}$ and $\xi_{i}$ (outputs) are also indicated for this case.

Needless to say, the results from this diagnostic only apply to the wake close to the axis, since this constitutes the data being used as input, and the electron is only pushed along one dimension (axially). A further caveat is that the accuracy of the results is of course heavily dependent on the number of file dumps from the simulation, because this determines the size of the time step covered by each iteration of the particle pusher. The larger this time step is, the more uncertainty there is as to the exact position $z_{\text {inj }}$ where the electron would have encountered a decelerating region, and consequently as to the total


Figure 3.5: Concatenated line-outs of $E_{z}$ close to the axis and example of an electron trajectory (the superimposed yellow line) calculated by the diagnostic for injected electrons. The labels $z_{\mathrm{inj}}, \xi_{f}$ and $\xi_{i}$ denote the minimum injection point and the final and initial beam coordinates of the electron, respectively.
energy gain.
In this case the chosen number of file dumps in the simulations corresponds to a time step $d t$ of around 0.11 ns or a distance $d x($ at $c)$ of 3.55 cm . Since the point of this diagnostic is to track the energy gain of an electron in a wakefield while taking the dephasing between both into account, the time step of the particle pusher must be able to resolve the detuning time $t_{d}$, which is defined as the time required for an electron to phase-slip half a wave period w.r.t. the wakefield, i.e. for an electron to enter a region where it would lose energy. In order to make a conservative estimate of $t_{d}$ for the present case, it will be assumed that the velocity of the electron is approximately $c$, and the lowest phase velocity of the wake registered in the concatenated $E_{1}$ line-out data for the entire simulation window will be chosen. This velocity can be approximately calculated from this data through the expression

$$
\begin{equation*}
\frac{v_{\phi}}{c}=\frac{\Delta \xi}{\Delta z}+1 \tag{3.5}
\end{equation*}
$$

where $\Delta \xi / \Delta z$ is a slope measured for a constant phase in the data exemplified in Fig. 3.5(more details about this and the derivation of Eq. 3.5 can be found in section A.2). The detuning time is given by [3]:

$$
\begin{equation*}
t_{d}=\frac{\pi}{\omega_{p}\left(\frac{c}{v_{\phi}}-1\right)}, \tag{3.6}
\end{equation*}
$$

and the corresponding detuning length by $L_{d}=c t_{d}$. Plugging all the values into these equations, the conservative detuning time is found to be approximately 3.89 ns , and the detuning length around 1.17 m . A comparison of the values for $d t$ and $d x$ with $t_{d}$ and $L_{d}$, respectively, leads to the conclusion that the number of dumps is sufficient for the algorithm to deliver correct results.

The reader may reasonably wonder why the study of injected electrons was not based on the PIC simulations that are constantly mentioned in this work. The answer is that this diagnostic is a compromise between less accurate, but more wide-ranging results and the opposite. The simulations of the AWAKE experiment mentioned in this chapter require around 30,000 computing hours, so the investigation of the behavior of injected electrons through [IC simulations for every single parameter variation would become prohibitively lengthy. This diagnostic instead offers fast, indicative results.

The characteristics of injected electrons were studied for eleven evenly spaced values of the final Lorentz factor between 1800 and 3800 (corresponding to final energies of 900 MeV and 1.9 GeV , respectively), and for two different regions of the final position along the beam at roughly 12 cm and 23 cm behind the head, both with a length of $\lambda_{p} / 2$, which is the expected length of the accelerating regions of the wake. A virtual electron was positioned at every cell in the $x_{1}$ direction (the resolution in $x_{1}$ determines the maximum number of electron trajectories that can be followed). Studying the injected electrons directly, besides the wake properties, is equally crucial to guarantee the possibility of deterministic injection for AWAKE, since any blind spots regarding their concrete behavior in the wake are thereby covered.

## Analysis of the parameter scans

For the comparison of the parameter scan results, the region around 12 cm behind $\xi_{h}$ and the final Lorentz factor $\gamma_{f}$ of 2000 (corresponding to 1 GeV , see white dashed line in Fig. 3.9) were chosen. The quantities being compared are the energy gain in terms of a Lorentz factor $\Delta \gamma$ and the minimum injection point $z_{\text {inj }}$.

Fig. 3.6 shows the results for the energy gain within a region with a length of $\lambda_{p} / 2$ (corresponding to the accelerating half-period of the wake). Two immediately prominent features of the graph demand some clarification: the flat top for all the simulations and the seemingly inconsistent trough of the


Figure 3.6: Energy gain for injected electrons with $\gamma_{f}=2000$ versus their final position along the beam w.r.t. the head $\xi_{f}-\xi_{h}$ for all parameter scans.
$\sigma_{z b}+5 \%$ run on the right.
The energy gain plateau (where $\Delta \gamma \sim \gamma_{f}$ ) is the region where the particles had low initial energies $\gamma_{0}$, or, in other words, where they had most available distance to be accelerated (this is also the region with the lowest injection points, compare with Fig. 3.7. Though $\Delta \gamma$ appears perfectly constant here, the values actually oscillate between around 1940 and 1970, which is consistent both with the initial energy threshold $\gamma_{\min }=29$ and with the variation of $z_{\mathrm{inj}}$ in this region, seen in Fig. 3.7. The length of this plateau gradually decreases for higher final energies, since more particles along $\xi_{f}$ require higher initial energies (if $\gamma_{0}$ is higher, $\Delta \gamma$ is no longer of the order of $\gamma_{f}$ but smaller).

The intuitively out-of-place drop of $\Delta \gamma$ to zero in the $\sigma_{z b}+5 \%$ run has a more prosaic explanation (this aberration can also be seen in the $z_{\text {inj }}$ results as a peak, in Fig. 3.7). The electric field data from which this diagnostic departs is considerably noisy, as Fig. 3.5 can attest. This noise is especially critical in the transitions from a positive to a negative electric field or vice-versa, due to the no-energyloss condition of the algorithm. This condition means that as soon as the algorithm encounters a cell with a positive value for the electric field, no matter how small, the particle trajectory is stopped. The anomalous data point at around -12.00 cm is due to a "parasitic" cell at $z=10.33 \mathrm{~cm}$ close to the electric field transition with a positive value which should in fact have a negative value, and this is why the algorithm assigned a zero energy gain to the electron in question.

Proceeding now to the analysis of the remaining data in Fig. 3.6, and neglecting the data point mentioned in the last paragraph, one can observe that the decreased bunch dimensions $\sigma_{z b}$ and $\sigma_{r b}$ and the decreased timing jitter $\Delta t$ seem to shorten the accelerating region by around $2.6 \%$. Most relevant for the witness electrons, however, is the right half of the accelerating half-period, since it is also focusing, and this region is actually $\lambda_{p} / 4$ long, where the difference becomes around $5 \%$. The fact that the timing jitter follows the same tendency here as the bunch dimensions is unexpected, due to its similarity with the bunch population parameter, as mentioned before. Interestingly, the opposite variation of these three parameters, nor any other parameter, have the effect of lengthening the region of energy gain on the right side.

The constant value for the energy gain registered in most of the region suffers hardly any effects from the parameter variations. The left-hand limit of the graph is too irregular to allow for any conclusions, though this half is in any case less interesting since it is defocusing.

$$
\gamma_{\mathrm{f}}=2000
$$



Figure 3.7: Minimum injection point for electrons with $\gamma_{f}=2000$ versus their final position along the beam w.r.t. the head $\xi_{f}-\xi_{h}$ for all parameter scans. This figure is the line-out marked by the white dashed line in Fig. 3.9

The right end of the graph of the minimum injection point, in Fig. 3.7, reflects almost directly the same interval in the graph of the energy gain, which is not the case for the remaining region. 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. Electron injection and acceleration can thus be accomplished reliably in AWAKE experiments.

## General observations

Having performed a quantitative analysis of the robustness of injected electron parameters against expectable driving bunch fluctuations, some general observations about the results from the diagnostic described above are also in order. The AWAKE experiment is currently exploring the physics of the SMI which means that only the proton beam and the ionizing laser pulse are required. The next two phases of the experiment will, however, introduce the injection and acceleration of electrons to multiple GeV , first as a long bunch from which only a fraction of the electrons are captured, and later as very short bunch ( $\sigma_{z e}<\lambda_{p}$ ) that will be injected into a precise accelerating and focusing phase of the wakefield. For the immediately next phase an electron beam injector (in the form of an RFgun) with the capacity for energies of $10-20 \mathrm{MeV}$ will be installed, whereby the electrons are supposed to be injected at $4-5 \mathrm{~m}$ of propagation distance and around $\sigma_{z b}(12 \mathrm{~cm})$ behind the laser pulse [36].

The required initial energies $\gamma_{0}$ obtained from the diagnostic (limited by a minimum threshold of 29) are shown in Fig. 3.8 for two regions of the beam, where the right-hand one is approximately the currently planned position for electron injection (i.e. 12 cm behind the head of the beam). The mostly purple "blotches" in Fig. 3.8 correspond to $\gamma_{0} \sim 50$, or initial energies of the order of 25 MeV , which means that these blotches roughly cover what is accessible to AWAKEwith the current electron injection conditions.

On one hand, for a fixed final energy (e.g. $\gamma_{f}=2500$ ), the length of the section where $\gamma_{0} \sim 50$ is


Figure 3.8: Required initial Lorentz factor of an injected electron according to its final position along the beam w.r.t. the head $\xi_{f}-\xi_{h}$ and its final energy in terms of a LORENTz factor $\gamma_{f}$ for two different regions of the beam (note the logarithmic scale of the data). These results are from the standard simulation.
enough to reach that $\gamma_{f}$ shrinks for the beam region further ahead (right-hand region), which suggests that the left-hand region ( 23 cm behind the head) would be a more efficient position to inject the electrons. On the other hand, the purple "accessibility blotch" stretches all the way up to higher final energies in the left-hand region, unlike the other region. This may mean that there is an obstacle to the multi-GeV energy range with current conditions, and these results again indicate that the region around 23 cm behind the head of the beam would be a better place to inject the electrons.

This discrepancy has likely less to do with the diagnostic used here than with the reason behind the choice of the injection position by the AWAKE collaboration. A number of publications related to the experiment state that the maximum amplitude of the wakefield along the beam after SMI saturation is located around $\sigma_{z b}$ behind the laser pulse, citing [27] (where this is not conveyed explicitly). However, the simulations performed for the present work produce conflicting results. A comparison between the concatenated $E_{1}$ line-outs in the two regions considered here (not shown) reveals that the amplitude of $E_{1}$ is clearly higher in the region around -23 cm . If the location for injection was chosen based on the simulations presented in [27], it should be noted that, despite being highly reliable 3D PICssimulations, some of the parameters assumed therein differ significantly from the current nominal parameters for AWAKE (e.g. a proton beam energy of 24 instead of 200 GeV , or a peak beam density ratio $n_{b 0} / n_{0}$ of 0.04 instead of around 0.003 , where $n_{0}$ is the plasma density).

The minimum injection points $z_{\mathrm{inj}}$ determined by the electron-pushing diagnostic are also presented in Fig. 3.9 for the same regions as above. This data generally implies higher propagation distances for


Figure 3.9: Deduced minimum injection point for electrons according to their final position along the beam w.r.t. the head $\xi_{f}-\xi_{h}$ and their final energy in terms of a Lorentz factor $\gamma_{f}$ for two different regions of the beam (results from the standard run). The white dashed line represents the line-outs taken for the analysis of the parameter scans.
injection than the range planned for the next phase of AWAKE, which is $4-5 \mathrm{~m}$. For a final energy of $1 \mathrm{GeV}\left(\gamma_{f} \approx 2000\right)$, for example, the range for injection in the right-hand region is about 6-7 m . However, it could be argued that the conditions used in this algorithm to obtain the values of $z_{\text {inj }}$ are too rigid, more prominently the condition that the electron does not lose any energy. An electron in a decelerating region may be able to be refocused and slip back into an accelerating region (note also that the conservative estimate of the detuning length discussed above is much smaller than 10 m ).

The only way to gain certainty about some contradictory results presented here and the behavior of injected electrons in general would be to consider the full-fledged 3D (or symmetrically 2D) movement of electrons under the influence of both longitudinal and transverse forces, which was not attainable within the time frame of this work but could constitute an item for future work.

### 3.3 Conclusion

A positive outcome for the work presented in this chapter would be the absence of irregularities, since ideally the outputs of the AWAKE experiment should not be susceptible to natural oscillations of the initial parameters. This chapter has shown that, despite depending completely on an instability, and as far as the tools developed for this work can attest, the outputs from the AWAKE experiment are robust against expectable shot-to-shot proton bunch fluctuations.

In a broad manner, the largest effects on the wakefield properties were always observed somewhere along the propagation distance instead of at the end, and they rarely exceeded the magnitude of the original variation in relative terms. At 10.33 m the effects (if existent) always amounted to less than $5 \%$, proportionally. This does not mean, of course, that the larger differences along $z$ do not have any consequences on other outputs, such as the final energy of the witness electrons. The tool developed to investigate this aspect, however, showed that the consequences for the injection of electrons are of a similarly small scale as before, where conclusions could be drawn. The regions where this was not possible luckily coincide with the defocusing half of the accelerating $\lambda_{p} / 2$, and are therefore less important.

It was also noted that the results from the simulations performed for this chapter imply that there is a more efficient region along the beam to inject electrons than the one currently envisaged for AWAKE, though this would require a deeper investigation. The results from the diagnostic that studies the behavior of injected electrons should in any case be confirmed by PIC simulations that include the electrons in the future.

This work leads to the conclusion that electron acceleration in the AWAKE experiment is viable at least to the level of proof-of-principle experiments, but also possibly beyond.

## Chapter 4

## Single beamlet linear wakefield solver

The physics behind self-modulated particle-driven plasma waves is quite complex. 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 [34]. What is known about the phase after saturation of the instability, however, derives solely from numerical simulations.

During this nonlinear stage the amplitude of the wakefield drops significantly (see Fig. 4.1) and, despite numerical studies, it is not fully understood why this happens. Such studies have revealed, for example, that the introduction of a small step in the longitudinal density profile of the plasma can dramatically counteract the decline of the wakefield's amplitude [48]. They have also found that the decline must be connected with a lengthening of the plasma wavelength $\lambda_{p}$ and the consequent rephasing of the self-modulated beamlets. The exact mechanisms that support these observations are still under discussion, though.

One of the possibilities currently under consideration as an explanation for the nonlinear behavior of the self-modulated beam is related to the way the wake driven by each beamlet interferes with those of the previous beamlets. The wakefield amplitude can remain constant or increase if these single-beamlet-driven wakes are all coherent. 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. At the same time, however, any beamlet with a length different than zero (i.e. any profile that is not a DIRAC delta function) introduces a small phase shift at the beginning of its own wakefield, which means that even a train of rigid and perfectly equidistant beamlets displays the same nonlinear behavior of decreasing amplitude [49].

To complicate matters further, this problem is spatiotemporal. The behavior of the wakefield must


Figure 4.1: Typical evolution of the average amplitude of $E_{z}$ along time or, equivalently, propagation distance (data taken from an OSIRIS simulation). The linear and nonlinear stages of the SMI are indicated on the graph.


Figure 4.2: Example of the axial electric field $E_{z}$ along the beam ( $\xi$ ) for a fixed time. The gray data is taken from an OSIRIS simulation, while the red lines are a smoothed envelope of this data.
be understood both along the beam at a fixed time (spatially, as shown in Fig. 4.2) and along a certain propagation distance for a fixed beam position (temporally). The solution will probably consist of an interplay between both phenomena.

Understanding theoretically why the SMI converges towards the described behavior is an essential first stage to mitigating it. This is the motivation behind the development of the computer program presented in this chapter, which will use the charge density data from OSIRIS simulations after selfmodulation has occurred, calculate each beamlet's wakefield and sum all the contributions to an overall $E_{1}$ distribution. This method will for the first time separate the contributions from single beamlets and hopefully provide some insight about their interaction with the overall wakefield in the nonlinear stage.

### 4.1 Description of the problem and requirements

In a more detailed manner, this program will have to perform the following tasks: read the charge density data from the OSIRIS output files, which use the Hierarchical Data Format (HDF); identify and isolate the self-modulated beamlets in the simulation window; calculate the corresponding axial electric field for each beamlet through a discretized version of linear wakefield theory and the simulation data; sum all the contributions; and lastly save the result to the HDFformat so as to take full advantage of the available visualization routines. In addition, due to the large amount of computation required to calculate the wakefield, this program will be implemented with support for parallel computation, using the Message Passing Interface (MPI)

These tasks can be condensed into the following key challenges.

- Data input/output: The beam density data must be retrieved from simulation files in HDF5 format, and the result of the program must be saved to a format that is readable by the OSIRIS visualization tool VisXD (see section 1.6.1).
- Beamlet isolation: From the spatial distribution of charge density obtained in simulations, the beamlets that form after self-modulation must be identified and separable from the rest of the data, so that they can be passed to the calculation of the respective wakefield.
- Wakefield computation: The wakefield equation must be implemented correctly and as efficiently as possible, as well as a method to incorporate parallelization into its computation.

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 [3]

$$
\begin{equation*}
E_{z}(\xi, r)=4 \pi e k_{p}^{2} \int_{\infty}^{\xi} d \xi^{\prime} \int_{0}^{\infty} d r^{\prime} r^{\prime} \cos \left[k_{p}\left(\xi-\xi^{\prime}\right)\right] I_{0}\left(k_{p} r_{<}\right) K_{0}\left(k_{p} r_{>}\right) n_{b}\left(\xi^{\prime}, r^{\prime}\right), \tag{4.1}
\end{equation*}
$$

where $I_{0}$ and $K_{0}$ are the zeroth-order modified Bessel functions, and $r_{<,>}$stands for the smallest/largest of $r$ and $r^{\prime}$.

As a first step, Eq. (4.1) must be discretized. Both infinities in the integrals represent the regions beyond the beam in question (in the case of $\xi$ ) and above the beam (in the case of $r$ ). In practice, however, after each beamlet there will be a finite position $\xi_{f}$ after which all density values are zero. The same is true in the transverse direction, for some $r_{f}$. This means that the limits of integration may be replaced by $\xi_{f}$ and $r_{f}$. Hence, the electric field $E_{z}$ for a point $\left(\xi_{0}, r_{0}\right)$ can be rewritten as:

$$
\begin{align*}
E_{z}\left(\xi_{0}, r_{0}\right)=-4 \pi e k_{p}^{2} \int_{\xi_{0}}^{\xi_{f}} d \xi \cos \left[k_{p}\left(\xi_{0}-\xi\right)\right]\left[K_{0}\left(k_{p} r_{0}\right) \int_{0}^{r_{0}}\right. & d r r n_{b}(\xi, r) I_{0}\left(k_{p} r\right) \\
& \left.+I_{0}\left(k_{p} r_{0}\right) \int_{r_{0}}^{r_{f}} d r r n_{b}(\xi, r) K_{0}\left(k_{p} r\right)\right] . \tag{4.2}
\end{align*}
$$

Substituting the normalized quantities for the electric field, the density and all spatial variables as listed in Table 1.2 (and dispensing with the hats), and discretizing the integrals, Eq. (4.2) becomes

$$
\begin{align*}
E_{z}\left(\xi_{0}, r_{0}\right)=-\sum_{\xi_{i}=\xi_{0}}^{\xi_{f}} \Delta \xi_{i} \cos \left[k_{p}\left(\xi_{0}-\xi_{i}\right)\right]\left[K_{0}\left(k_{p} r_{0}\right)\right. & \sum_{r_{i}=0}^{r_{0}} \Delta r_{i} r_{i} n_{b}\left(\xi_{i}, r_{i}\right) I_{0}\left(k_{p} r_{i}\right) \\
& \left.+I_{0}\left(k_{p} r_{0}\right) \sum_{r_{i}=r_{0}}^{r_{f}} \Delta r_{i} r_{i} n_{b}\left(\xi_{i}, r_{i}\right) K_{0}\left(k_{p} r_{i}\right)\right] \tag{4.3}
\end{align*}
$$

where all quantities are normalized and $\Delta \xi$ and $\Delta r$ are the cell dimensions in the axial and radial directions, respectively. Since these lengths are constant, they can be taken out of the sums. Furthermore, moving the sum in $\xi$ inside each sum in $r$ will greatly reduce the number of computations per grid point. The final expression to be implemented in the code reads

$$
\begin{align*}
E_{z}\left(\xi_{0}, r_{0}\right)=-\Delta \xi \Delta r\left[K_{0}\left(r_{0}\right) \sum_{0}^{r_{0}} r_{i} I_{0}\left(r_{i}\right) \sum_{\xi_{0}}^{\xi_{f}} \cos \left(\xi_{0}-\xi_{i}\right) n_{b}\left(\xi_{i}, r_{i}\right)\right.
\end{aligned} \quad \begin{aligned}
& \left.\quad+I_{0}\left(r_{0}\right) \sum_{r_{0}}^{r_{f}} r_{i} K_{0}\left(r_{i}\right) \sum_{\xi_{0}}^{\xi_{f}} \cos \left(\xi_{0}-\xi_{i}\right) n_{b}\left(\xi_{i}, r_{i}\right)\right] .
\end{align*}
$$

This equation requires the beam density profile $n_{b}\left(\xi_{i}, r_{i}\right)$ as input, which prompts the issue of processing the simulation data into separate beamlet regions. The density distribution will typically be a data grid with a larger amount of cells in the axial direction than in the radial direction. An example of such a set of data can be seen in Fig. 4.3 for a region close to the head of the beam (from this point on all quantities are normalized unless specified otherwise, which means that $\xi$ and $r$ are equivalent to the OSIRIS variables $x_{1}$ and $x_{2}$, respectively). Since the beamlets are separated along $\xi$, the algorithm to identify them must work its way along the axial direction, taking the line-outs of the data at each $r$ as input. Such a line-out is also exemplified in Fig. 4.4 for the same $\xi$ interval as Fig. 4.3

After the beamlets have been isolated, their respective wakefields must be calculated separately at every point within and behind the beamlet. If $\xi_{f}$ is the delimiting position before the next beamlet, then the computation area for each beamlet (and thus each wake) streches from the beginning of the


Figure 4.3: Sample of charge density data for a portion of the simulation window just before the head of the beam (on the right), where $x_{1}$ and $x_{2}$ are the OSIRIS coordinates corresponding to $\xi$ and $r$, respectively. The white dashed line represents the line-out shown in Fig. 4.4
simulation box to $\xi_{f}$. This means that the closer a beamlet is to the head (the end of the simulation box), the larger its grid and the number of computations involved.

### 4.2 Algorithms and implementation

This program has been divided into six sequential stages: five of them are dedicated to the identification and isolation of the beamlets, and the sixth one to computing the wakefield. Parallelization was implemented in only some of the first five algorithms (stages 1,4 and 5 ), since the others would involve too much communication to be efficient, so each of these five stages will merely earn a loose description. The sixth stage, being the crucial and most parallelized part of this program, will be explored in more detail.

### 4.2.1 Stages 1 through 5: Beamlet isolation

## Stage 1: First try at setting marks

In this first step towards beam identification, the charge density data are scanned along $\xi$ and for every $r$, i.e. the algorithm is applied to each line-out similar to the one in Fig. 4.4 at every $r$. The different rows (line-outs) of data are distributed among the available processes and smoothed along $\xi$.

Finding a beamlet implies finding a gap, which is the case when a minimum interval for which the density values are below a certain threshold has been traced. After an interval has qualified as a gap, the algorithm looks for the next point where the density value is again above the threshold, and places a mark exactly between this point and the point where the gap began.

The result can be seen in Fig. 4.5a), for a small region. Since the data is still quite noisy and irregular after the smooth, there are clearly missed spots where a mark should have been. The marks are stored


Figure 4.4: Line-out of the charge density data in the $x_{1}$ direction at $x_{2}=2$ in the same region as Fig. 4.3
in a binary grid (where a mark is represented by 1 , and no mark by 0 ) the same size of the original density grid.

## Stage 2: Identifying groups of marks

The goal of the first five stages of the program is to deliver clear, full-fledged boundaries between beamlets, not a set of scattered marks. This stage is pivotal for the achievement of that objective, and is the most complex of all five beamlet-isolating algorithms.

The general idea is to scan the grid where the marks are stored from below ( $r=0$ ) and from the head of the beam $\left(\xi=\xi_{\max }\right)$. When a mark is found along the row that is currently being probed at some $r_{0}$, the algorithm enters another loop to look for marks above that one, starting at $r=r_{0}+1$ (if the coordinates are now understood as grid indexes) and within a certain tolerance around the position $\xi_{0}$ where the mark was found. This process is repeated in the vertical direction while $\xi_{0}$ is adapted to positions of the new marks that are found. If the amount of marks found this way surpasses a certain threshold (in order to avoid misidentified beamlets), then their coordinates are stored as a set and the marks are copied to a new grid. If the algorithm reaches the maximum radius without meeting the mark threshold, the marks found until then are disregarded. In both cases all the marks found in the loop are erased from the original grid, so that the algorithm does not encounter them again.

As the coordinates for all the marks in each set are stored in a matrix, these sets, henceforth called boundaries, are ordered according to the position of their first mark. The first boundary from the right of the grid is boundary number one, for example. The identification of different boundaries is shown in Fig. 4.5b) in the form of different colors for different boundaries.

## Stage 3: Correcting the boundaries

The previous stage cannot produce a perfect output and it is noticeable that some sets of marks identified as two boundaries should in fact be one boundary. This is problematic since the boundaries are to be supplemented with missing marks later on, and they must be correctly identified for this.

By comparing the positions of the first marks of every two adjacent sets, and if the difference between them is below a certain value, the two boundaries are merged. Fig. 4.5k) shows a corrected version of the previous stage.

## Stage 4: Filling out missing marks

Now that all boundaries are found and well-defined, one can proceed to completing them. In general the new marks are determined through a linear interpolation between the existent marks. The boundaries are distributed among the available processes.

If the first mark from the bottom is not located at $r=0$, then there is no interpolation and the space in between is filled with marks at the same $\xi$ as that first mark. For all the marks between this first mark and the last mark, a linear interpolation is made between every consecutive pair, and the new marks are positioned accordingly. Finally, if the last mark is not at $r=r_{\text {max }}$, a linear regression is performed using a certain number of the last preceding original marks, and the space between the last mark and $r_{\text {max }}$ is filled with new marks according to the result of the regression.

The same sample as before is shown in Fig. 4.5d) after undergoing this stage.


Figure 4.5: Images exemplifying all five stages required for beamlet isolation in the same sample region. Images a) and b) include the beam density distribution in the background, while the colored points denote marks set by the algorithm.

## Stage 5: Smoothing borders

As a final step, the axial positions of the marks are smoothed with a certain width, so that the boundaries are more continuous. This task is once again divided among processes. The final result after all these stages (for the same sample) is depicted in Fig. 4.5 e). The final output of this part of the program is a matrix with the coordinates of all marks in each boundary.

### 4.2.2 Stage 6: Wakefield computation

Each sum in Eq. (4.4) requires a loop in the program implementation. The sum in $\xi$, given by

$$
\begin{equation*}
\sum_{\xi_{0}}^{\xi_{f}} \cos \left(\xi_{0}-\xi_{i}\right) n_{b}\left(\xi_{i}, r_{i}\right) \tag{4.5}
\end{equation*}
$$

can be computed only once for each radius and becomes an expression that depends on $r_{i}$, e.g. $C\left(r_{i}\right)$. Then there are two sums for the two sections of $r$ (below and above $r_{0}$ ). An initial approach could


Figure 4.6: Symbolic representation of the method by which the wakefield is calculated for every cell in a column, saving redundant operations.
consist of doing the calculation $r_{i} I_{0}\left(r_{i}\right) C\left(r_{i}\right)$ or $r_{i} K_{0}\left(r_{i}\right) C\left(r_{i}\right)$ at every iteration of these loops in order to determine the sum. However, since this would imply performing the same computation many times, a method has been devised to avoid this redundancy.

The expressions contained in each of the two sums in $r$ in Eq. (4.4) can be regarded as vectors, whose elements are summed within a certain range according to the current $r_{0}$. Adopting the following definitions,

$$
\begin{equation*}
C\left[r_{i}\right]=\sum_{\xi_{0}}^{\xi_{f}} \cos \left(\xi_{0}-\xi_{i}\right) n_{b}\left(\xi_{i}, r_{i}\right), \quad V 1\left[r_{i}\right]=r_{i} I_{0}\left(r_{i}\right) C\left[r_{i}\right], \quad V 2\left[r_{i}\right]=r_{i} K_{0}\left(r_{i}\right) C\left[r_{i}\right], \tag{4.6}
\end{equation*}
$$

and noticing that $\sum_{r_{0}}^{r_{f}} V 2\left[r_{i}\right]=\sum_{0}^{r_{f}} V 2\left[r_{i}\right]-\sum_{0}^{r_{0}-1} V 2\left[r_{i}\right]$, the value of the contribution to the electric field at $r_{0}=0$ and for any $\xi_{0}$ can be written as

$$
\begin{equation*}
E_{z}\left(\xi_{0}, r_{0}=0\right)=\Delta \xi \Delta r\left[K_{0}(0) \operatorname{sum} 1+I_{0}(0) \text { sum } 2\right] \tag{4.7}
\end{equation*}
$$

where sum $1=0$ and sum $2=\sum_{0}^{r_{f}} V 2\left[r_{i}\right]$ (the sum of all elements in $V 2$ ). All consecutive values of $E_{z}$ in this column (i.e. for a fixed $\xi_{0}$ and increasing $r_{0}$ ) can be easily calculated by performing an addition/subtraction of the corresponding element of the vectors $V 1$ and $V 2$ to sum 1 or sum2, respec-

```
for (int r = 0; r <= rmax; r++) {
    C = 0;
    for (int xi = xiO; xi < xif; xi++)
        C += cos(xi0 - xi) * density[xi][r];
    V1[r] = r * IO(r) * C;
    V2[r] = r * KO(r) * C;
    }
    sum1 = 0;
    sum2 = 0;
    for (int r = 0; r < rf; r++) sum2 += V2[r];
    for (int r0 = 0; r0 < rf; r0++) {
        sum1 += V1[r0];
        sum2 -= V2[r0];
    Ez[xiO][r0] += dxi * dr * [KO(r0) * sum1 + IO(r0) * sum2];
}
```

Figure 4.7: Pseudo-code for the algorithm that determines the contribution to the global wakefield for an entire column (fixed $\xi_{0}$ ) and for a certain beamlet, based on the C language.


Figure 4.8: Schematic data grid for a given beamlet (represented by the shaded area on the right) to demonstrate the distribution of columns by different processes, where the unit amount of columns ( $c_{\text {unit }}$ ) is 2 , the cell length of the grid $(L)$ is 15 and the number of processes $\left(n_{p}\right)$ is 5 .
tively. This logic is represented in Fig. 4.6. The C-based pseudo-code for the computation of the $E_{z}$ contribution for a fixed $\xi_{0}$ is shown in Fig. 4.7

This method works recursively inside each column, and so it would be inefficient to split columns between processes due to the communication overhead that this would incur. Consequently, in designing the parallelization of this stage, this calculation was defined as the indivisible work unit that any single process can complete.

The parallelization of the sixth stage is accomplished through the distribution of columns (or sections of $\xi$ ) within the current beamlet grid to different processes, whereby the beamlets are handled successively. As mentioned before, the position of a column in the grid will determine the workload required for that column. A column at $\xi_{0}=1$, for example, will require a higher amount of calculations than a column at $\xi_{0}=10$ (this can be deduced from the loop to determine $C$ in Fig. 4.7). This fact recommends the division of the beamlet grid in regions of linearly increasing size instead of a uniform distribution, where the largest region will be rightmost and the smallest region (at least one column) leftmost in the grid, as exemplified in Fig. 4.8. Attempting this kind of distribution for an arbitrary number of processes is not trivial.

Since there is the initial constraint that there can be at most one process per column (or, conversely, at least one column per process), if the number of available processes is greater than the number of cells in the $\xi$ direction, then the excessive processes will stay idle and the columns will be uniformly distributed (one column per process). This situation is in any case unlikely to be encountered, since the typical number of cells in the axial direction is of the order of tens of thousands.

Ideally, the number of processes would allow a perfectly linear distribution of the number of columns that simultaneously fulfills the condition

$$
\begin{gather*}
c_{\text {unit }}+2 c_{\text {unit }}+3 c_{\text {unit }}+\cdots+n_{p} c_{\text {unit }}=L \quad \Leftrightarrow \quad c_{\text {unit }} \sum_{i=1}^{n_{p}} i=L \\
\Leftrightarrow \quad c_{\text {unit }} \frac{1}{2} n_{p}\left(n_{p}+1\right)=L \tag{4.8}
\end{gather*}
$$

where $c_{\text {unit }}$ is a unit amount of columns, $n_{p}$ is the number of processes and $L$ is the length of the beamlet grid in cells (these quantities are illustrated in Fig.4.8). The algorithm for process distribution will calculate a tentative $c_{\text {unit }}$ by rearranging Eq. (4.8) and will round the result down. If this number is zero, there is no possible arrangement that can begin from a linear distribution (even with some
remainder of columns that can be distributed afterwards), but the algorithm attemps to redistribute excess processes in smaller regions to the larger regions. Such a case can be seen in Fig. 4.9 in the red data. The blue data exemplifies a case where a linear distribution could be applied.

The information regarding which process handles which columns is stored into an array, which will be useful when using the MPI functions MPI_Scatterv and MPI_Gatherv. These functions will serve to distribute these sections of the total $E_{z}$ matrix to the processes and recover them after each process has performed its calculation.


Figure 4.9: Column chart of two example distributions using 27 processes, for two different beamlets which require grids of two different lengths.

### 4.2.3 Implementation

This program was written in the C programming language, requiring the external libraries MPI, HDF5 and GSL (GNU Scientific Library). The format chosen to save the final data was "zdf", a custom-built file format developed at GoLP. The source files for its implementation were supplied and are not a product of this project [50]. Before writing anything in C, all the algorithms were devised and tested (serially) in Interactive Data Language (IDL), which supports array programming and is thus convenient. An organizational overview of the final program is shown in Fig. 4.10 .

The program was compiled on a 1800-core cluster at IST, and tested using up to the full cluster capacity without any issues.


Figure 4.10: Diagram showing the structure of the program and the different files composing it. Diamond-shaped objects with a red outline represent external libraries. Stadium-shaped objects stand for header files and rectangular objects for source files, whereby a green outline signifies that the files are not original. The direction of the arrows means "includes".

### 4.3 Benchmarking of the solving algorithm

In order to prove the accuracy of the algorithm, a simple benchmark was devised to test the program. This benchmark consists of a single short electron bunch propagating in plasma in the linear regime of wakefield excitation, with an axially Gaussian and radially cosinusoidal profile, i.e. a profile of the form

$$
\begin{equation*}
n_{b}(\xi, r)=n_{b 0} \exp \left(-\frac{\xi^{2}}{2 \sigma_{z b}^{2}}\right) \frac{1}{2}\left[1+\cos \left(\sqrt{\frac{\pi}{2 \sigma_{r b}^{2}} r}\right)\right] \tag{4.9}
\end{equation*}
$$

where the ratio $n_{b 0} / n_{0}=0.1$ was chosen. The test was performed by running an OSIRIS simulation of this setup (in 2D cylindrical coordinates), using the charge density data of the electron beam at the end of the simulation as input for the program, and comparing the wakefield calculated by the program with the one from the PIC simulation (also at the end of the simulation). This test obviously applies exclusively to the solving algorithm, since there is no beamlet isolation involved.

The resulting fields are shown in Fig. 4.11. The agreement seems nearly perfect, though there is a difference of $5 \%$ in the maximum amplitude registered in the entire simulation window. This may be due to the cells close to the axial boundary in the simulation (at $r=0$ ), since, as mentioned before, the values there tend to be noisier and assume improbably high values. This is also supported by the fact that, overall, the amplitudes are in excellent agreement.


Figure 4.11: Comparison between the axial wakefield obtained from an OSIRIS simulation (top) and from the program (bottom) for the benchmarking test.

### 4.4 Applications in the AWAKE simulations

Having assured the reliability of the code, it can now be applied to a typical simulation of the AWAKE experiment so as to investigate the issues mentioned in the introduction to this chapter, namely what causes the observed wakefield behavior in the nonlinear stage of the SMI Two applications of the program on the AWAKE simulation and their results are presented in the following, where the spatial and temporal characteristics of the problem are both addressed.


Figure 4.12: Comparison of the average amplitude of the axial electric field according to the simulation's results and the wakefield solver's results.

### 4.4.1 General wakefield

The program was used to determine the axial wake from the density data of the AWAKE simulation at three different times (equivalently, propagation distances or dumps), chosen after the linear phase of the SMI. at dump $30(z \approx 6.39 \mathrm{~m}), 40(z \approx 8.52 \mathrm{~m})$ and $50(z \approx 10.65 \mathrm{~m})$.

A comparison of the results from the program between the different simulation times yields no obvious differences (such as field configuration) except for the generally decreasing amplitude (see purple data points in Fig. 4.12). Using the diagnostic for the average amplitude of the electric field (see section A.1) to condense this information and compare it with the results from the simulation, the graph in Fig. 4.12 was obtained. Since the program was only used at three different dumps, there are only three data points for these results.

A direct comparison between the configurations of the simulation and solver fields for the three different times reveals that the simulation field converges towards the solver field for later times, as


Figure 4.13: Comparison of the axial wake in the tail region at two different times and for the simulation and the wakefield solver. The line-outs of the data close to the axis are superimposed (at $z \approx 10.65 \mathrm{~m}$ the simulation line-out is reproduced in gray in the graph for the wakefield solver for comparison).


Figure 4.14: Comparison of the axial wake in the head region for the simulation and for the wakefield solver.

Fig. 4.12 already subtly hints at. This is demonstrated in Fig. 4.13, where the tai ${ }^{1}$ region of the simulation window (where the differences are strongest) is sampled for the first and third dumps where results were produced by the wakefield solver. Despite the phase shift between both fields at 10.65 m , the initially noise-riddled tail region of the simulation approximates the solver's field very well in the final dump.

By contrast, there is not much evolution in the head region, as it stays identical to the solver's configuration through all three times, even w.r.t. the phase. This is shown in Fig. 4.14 , where a comparison is shown solely for the first dump where the solver was used (when the differences should be largest). 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 solver results in other regions is mostly due to numerical noise (barely present in the wakefield solver's results).

### 4.4.2 Separation of individual wakes

The objective is now to test the theory regarding the interference between the wakes caused by each beamlet using the program. The wakefield solver can be modified so as to stop after a given number of beamlets and to produce two files instead of the overall wakefield file: one with the wake originated exclusively by the last beamlet, and one with the overall wake generated by the remaining beamlets before the last one. This allows the direct observation of how the field caused by a consecutive beamlet adds to the field already present due to all the previous beamlets, thereby offering further insight into the wake incoherence hypothesis.

This modified version of the code was applied to the last two dumps mentioned in the last section ( $z \approx 8.52$ and 10.65 m ), and for three different beamlet numbers (i.e. the number of beamlets after which the program stops): 8,90 and 170 . The beamlets are counted from the head of the beam, and these numbers are given according to the isolation of the beamlets in the last dump ( $z \approx 10.65 \mathrm{~m}$ ).

The beamlet isolation algorithm of the program naturally delivers different results for different density data. At the moment (the algorithm could be arbitrarily refined) it tends to amalgamate not yet

[^1]fully modulated regions of the bunch into a single beamlet. This happens more frequently at the beam head and for earlier times, since the SMI has not saturated entirely and the beamlets are not wholly separated in this region. The algorithm also tends to miss or misidentify beamlets in the tail region for later times, since, though the particle bunch is fully modulated, charge is gradually lost and the density peaks become smaller and more diffuse here. These aspects will, however, not change the conclusions obtained here, since it was controlled whether the same beamlets were being compared for both simulation times.

The results are shown in Fig.s 4.15-4.17. Noting that the OSIRIS variables $x_{1}$ and $t$ (implicitly normalized) correspond to $\hat{z}$ and $\hat{t}$ in a more careful physical notation, the horizontal axis is displayed as the normalized co-moving coordinate $\hat{\xi}$, since

$$
\begin{equation*}
\hat{z}-\hat{t}=z k_{p}-t \omega_{p}=k_{p}(z-c t)=k_{p} \xi=\hat{\xi} \tag{4.10}
\end{equation*}
$$

where the relationship $c=\omega_{p} / k_{p}$ was used and the hats denote normalized quantities. The three beamlets represent the different regions of the wake (head, middle and tail). Note that the normalized length of the simulation window is $1492 k_{p}^{-1}$.

Though this cannot be seen clearly from the figures, 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 justifies the high amplitude of its wake. This may also explain why its wake is almost imperceptibly out of phase with the existing wake, as Fig. 4.15 shows. By the ninetieth beamlet, however, there is an unambiguous phase-shift of about a fourth of a period, or $\lambda_{p} / 4$ (see Fig. 4.16). 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. 4.17. 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 $\lambda_{p} / 2$ (the shift of the existing wake is negligible in all three cases). Some mechanism is effectively "pushing" beamlets forwards in the tail region, as well as their respective wakes. 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 that other protons were refocused on the right side of the original beamlet, so that it appears as though the beamlet has advanced. This mechanism, which is related to charge recapture, is nonetheless only observed at the tail of the beam, where the relative amplitudes of the individual wakes are extremely small, and so it is not held as a significant factor behind the amplitude decline along $z$ (temporally).

These observations 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 $\xi$ ). The fact that the results from the linear wakefield solver overlap with the ones from PICsimulations 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 $\xi$.

There are two sides to the phenomenology of the amplitude decrease along the propagation distance $z$, where one of them is a negligible effect in this case but will be mentioned for the sake of completeness.

If each beamlet's wake was propagating at exactly the same speed as the overall wake (and if the driver charge did not decline), the general wake resulting from the interference between all the

Addition of wake from beamlet nr. 8


Figure 4.15: Wake caused by the 8th beamlet (top row) and wake caused by the first 7 beamlets (bottom row) for two propagation distances (left and right). The position of the 8th beamlet corresponds to about $96 \%$ of the simulation window. The line-outs of the field at the axis are superimposed. The line-outs for the 8th beamlet (top row) are reproduced in the bottom row (gray) for comparison, though their amplitudes are not to scale (compare the different color scales for both rows, on the right).

Addition of wake from beamlet nr. 90


Figure 4.16: Wake caused by the 90th beamlet (top row) and wake caused by the first 89 beamlets (bottom row) for two propagation distances (left and right). The position of the 90th beamlet corresponds to about $60 \%$ of the simulation window. The line-outs of the field at the axis are superimposed. The line-outs for the 90th beamlet (top row) are reproduced in the bottom row (gray) for comparison, though their amplitudes are again not to scale.


Figure 4.17: 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.


Figure 4.18: 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.
individual wakes would stay constant in time, which means that the wake amplitude could vary along the beam but the average amplitude of the entire wake would stay constant along $z$. In fact there exist small variations in the relative phases between each individual wake and the overall wake, as Fig. 4.17 demonstrates through the shift of the position of beamlet number 170 (discussed above). The interference pattern thus changes from instant to instant and more regions of the overall wakefield are thereby gradually canceled out. Such a shift is, however, only observable at the tail of the beam and, as stated already, this mechanism cannot have any meaningful weight in the explanation for the temporal decline of the wakefield amplitude. Charge recapture and the effects related to it, though, are still an interesting aspect of the nonlinear stage of the SMI which could be studied further and become significant in other conditions.

What is truly decisive for the amplitude decrease in $z$ is simply the fact that driver charge is gradually lost. Based on the simulation results, the average amplitude of $E_{z}$ over $z$ is plotted alongside the total proton charge in Fig. 4.18 , whereby only the charge below $r=3 k_{p}^{-1}$, where the fields start vanishing, was accounted for. This 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 $z$ ).

### 4.5 Conclusion

This chapter began with the description of a parallel program to calculate the axial electric field excited by a train of modulated beamlets in a way that takes their individual contributions into account, and proceeded to the presentation of results obtained from this program when applied to a simulation of AWAKE. It was shown that the spatial decline of the wakefield amplitude (along the beam) in the nonlinear stage of the SMI is due to incoherent interference between the wakes caused by each single beamlet, and the temporal one (along the propagation distance) due to driver charge loss. It is, however, still not clear why this is so and how these effects could be counteracted. The insight acquired in this chapter, though, will surely prove a sound starting point for further investigations.

## Chapter 5

## Antiprotons as wakefield drivers

The AWAKE experiment is unusual and different from previous experimental efforts in PWFA for many different reasons, but the most relevant one is the type of driver particle. As mentioned in Chapter 1 PWFA was initially proposed with electrons in mind as the wakefield drivers [15], and most experiments use this setup. Protons differ from electrons not only in their mass but also in their charge. It has been found that the positive counterparts of electrons, positrons, seem to be less efficient at driving a wakefield [51], which begs the question whether the negative counterparts of protons would in turn excite the wakefield at AWAKEmore effectively. This question is all the more worth exploring given the fact that negative ion beams may be available at CERN in the future [52].

Though linear wakefield theory predicts a perfectly symmetric behavior for driver particles of opposite charges, the asymmetry observed in reality is not unexpected. With the exception of electronpositron plasmas, there is always an underlying asymmetry in plasma, which consists of the higher inertia of positive particles versus negative ones. Since the plasma wave caused by a particle driver (the wake) is mostly carried by electrons, it makes sense that a negative driver would result in a different wake than a positive driver (which might retard the motion of electrons somewhat in comparison with a repulsive interaction).

This chapter thus explores a hypothetical substitution of the particles in the AWAKE driver bunch by their antiparticles, i.e. antiprotons, with the goal of understanding not only the differences between both drivers, but also their efficiencies in the specific context of a self-modulated accelerator.

### 5.1 Simulation

With the exception of the bunch dimensions ( $\sigma_{z b}=12.6 \mathrm{~cm}$ and $\sigma_{r b}=200 \mu \mathrm{~m}$ ) and the plasma density ( $n_{0}=7 \cdot 10^{14} \mathrm{~cm}^{-3}$ ), the baseline simulation used for this study is exactly the same as the one described in section 2.1 . The physical values corresponding to the simulation parameters for the plasma density

Table 5.1: General parameters of the baseline simulation for the study of antiprotons as drivers.

| Parameter | Normalized units | Physical units <br> $\left(n_{0}=7 \cdot 10^{14} \mathrm{~cm}^{-3}\right)$ |
| :--- | :---: | :---: |
| Length of simulation box, $L_{1}$ | $1492 k_{p}^{-1}$ | 29.99 cm |
| Height of simulation box, $L_{2}$ | $8 k_{p}^{-1}$ | 1.61 mm |
| Time step | $0.012 \omega_{p}^{-1}$ | $2.41 \mu \mathrm{~m}(\mathrm{at} \mathrm{c})$ |
| Propagation distance | $53000 k_{p}^{-1}$ | 10.65 m |

chosen here are listed in Table 5.1. The input for the simulation with the antiproton driver needed no modification other than a minus sign in the normalized charge-mass product.

### 5.2 Results

A large amount of diagnostics was applied to and/or developed for the run with antiprotons, because the results from some of the more straightforward diagnostics mentioned in other chapters raised deeper questions which required new post-processing tools. The first analyses were based on the diagnostic for the peak and average amplitude of $E_{z}$ (described in section A.1), for the total charge of a particle species (section A.3) and for the maximum witness electron energy (section A.4), which is based on the raw particle data for the witness electron species included in the simulation.

### 5.2.1 Primary outputs

The evolutions of the peak and average axial electric field for both simulations (with a proton and an antiproton driver) are shown in Fig. 5.1. There are some differences along the propagation distance, especially for the average field, where there is a notable difference in the smoothness of both maxima. The most striking differences, however, are observed at the end of the plasma cell ( $z \approx 10.30 \mathrm{~m}$ ), where the peak and average $E_{1}$ amplitudes are almost 2 and 1.5 times as high, respectively, for the run with antiprotons.

The total charge measured for both the driver species and the witness electrons shows even more impressive results for the antiproton run, as Fig. 5.2 demonstrates. The rms bunch radius $\sigma_{r b}$ is $200 \mu \mathrm{~m}$, so the amount of driver charge that is effectively driving the wakefield is found below a radius of one plasma skin depth $k_{p}^{-1}$, which corresponds to $201 \mu \mathrm{~m}$. The absolute amount of antiproton charge within this radius is almost 2.8 times higher than the proton charge at the end of the simulation, while almost 4.6 times more witness electrons remain inside the simulation window by the end of the plasma cell.

At this point 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 $\xi$ and $r$ in Eq. (4.1)). For this reason, one would expect the increase factors in peak 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 2 and 2.8 , respectively). This means that, although antiprotons are able to excite wakefields more efficiently than protons, the amount of charge available in the antiproton case would suggest that even higher wakefields could be achieved. The reasons for this discrepancy will be examined later.

Since the maximum electric field $E_{z}$ is higher for an antiproton driver, one could also expect the maximum energy reached by witness electrons to increase. This quantity, though it does not contain any information about the accelerated electrons' energy spectrum, provides an inkling of the potential for acceleration. As expected, the antiproton simulation again delivers better results in this respect (see Fig. 5.3). After 10 m of plasma the highest energy achieved by a witness electron is almost 1.7 times higher for an antiproton driver than for the standard setup, which is consistent with the 1.8 increase factor of the peak amplitude of $E_{z}$ at the same propagation distance.

### 5.2.2 Deeper analysis of charge preservation

The differences observed for the quantities considered above raise obvious questions, such as why the antiproton driver is able to conserve so much of its charge, and why the amplitude of the wakefield


Figure 5.1: Peak and average amplitudes of $E_{1}$ (left and right graph, respectively) versus the propagation distance for the proton and antiproton runs.


Figure 5.2: Total charge of the driver species below a radius of one plasma skin depth (left), and total charge of the witness electron species in the entire simulation window (right), both versus the propagation distance and for both runs.


Figure 5.3: Maximum witness electron energy for both runs versus the propagation distance.


Figure 5.4: Energies contained in focusing and defocusing regions of $E_{2}$ (w.r.t. the charge of the driver particles) for both runs (left) and ratio between these energies for both runs (right).
is not as high as this amount of charge would imply. In order to find answers to these questions, the post-processing of the simulation results must go further than the outputs presented above.

A first attempt at understanding why there is more charge left to excite the wake for an antiproton driver was directed at the field energy contained in focusing and defocusing (w.r.t. the charge of the driver) areas of the radial electric field $E_{r}$, which contributes most to the radial forces acting on the bunch particles (the contribution from the azimuthal component of the magnetic field $B_{\varphi}$ will be considered later). The diagnostic developed for this purpose (see section A.7) also calculates the ratio between both energies.

Besides showing in which case there is more energy available for focusing, this analysis will also supply 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.

The energies stored in $E_{r}$ (named $E_{2}$ in OSIRIS), shown on the left in Fig. 5.4 , display the typical SMI pattern of exponential growth [34] followed by a peak and a saturation phase. Just as in the evolution of the average amplitude of $E_{z}$ (Fig. 5.1 right), the peak is lower and happens later for antiprotons, while the decrease associated with the saturation of the SMI after the peak is slower than in the proton case. 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.

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, which may help to explain the disparity in absolute driver charge retained along the propagation distance. The ratio of focusing energy to defocusing energy, seen on the right in Fig. 5.4 , demonstrates both this fact and the fact that there is a larger difference between the energies in the antiproton run. This imbalance indicates that the wake generated by antiprotons is more nonlinear than the one generated by protons.

The concrete way through which particles are focused in both cases can be understood by observing the transverse forces acting on the driver bunch and the bunch density at different times. These quantities are superimposed in Fig. 5.5 for the middle region of the beam and for four different propagation distances. The transverse force acting on a relativistic particle beam propagating along the $z$ direction (with $v \sim v_{z}$ ) is given by

$$
\begin{equation*}
\mathbf{W}_{\perp}=q\left(\mathbf{E}_{\perp}+\mathbf{e}_{z} \times \mathbf{B}_{\perp}\right) \quad \Rightarrow \quad W_{r}=q\left(E_{r}-B_{\varphi}\right) \quad W_{\varphi}=q\left(E_{\varphi}+B_{r}\right) \tag{5.1}
\end{equation*}
$$

Transverse force and charge density in the middle region
protons
antiprotons


Figure 5.5: Evolution of $E_{2}-B_{3}$ (the unsigned transverse force, in red and blue) and the absolute value of the driver bunch charge density $|\rho|$ (in green) between 2.56 and 4.69 m for a proton and an antiproton driver. Note that negative (blue) areas of $E_{2}-B_{3}$ are focusing for protons while positive (red) areas are focusing for antiprotons.
in cylindrical coordinates and in the Gaussian system of units, where $\mathbf{e}_{z}$ is the unit vector in the $z$ direction. The $\varphi$ component of the transverse force can be neglected since cylindrical symmetry is assumed. Implementing the OSIRIS normalizations and variable names (more details in section A.8), the normalized transverse (radial) force is then simply

$$
\begin{equation*}
\hat{W}_{r}=\hat{q}\left(\hat{E}_{r}-\hat{B}_{\varphi}\right)= \pm\left(E_{2}-B_{3}\right) \tag{5.2}
\end{equation*}
$$

where the plus sign refers to protons and the minus sign to antiprotons. The quantity represented in Fig. 5.5 is in fact $\left|\hat{W}_{r}\right|$, so that the focusing areas have negative values for protons and positive values for antiprotons.

The images in Fig. 5.5 take place during the linear stage (growth) of the SMI, and a charge recapturing mechanism can be discerned for antiprotons. At 2.56 m the transverse force is beginning to modulate the driver bunch, whereby a second row of modulated particles, further from the axis, is formed in both cases. The crucial difference between them is the configuration of the fields close to the axis: for protons they tilt backwards in $\xi$, while for antiprotons they tilt slightly forward up to a certain height. Bearing in mind that the simulation window moves at $c$, particles in the second row can only "slide" back towards the axis by either moving vertically or falling behind w.r.t. $\xi$ in a focusing area. The shape of the fields in the antiproton case provides an especially large window of opportunity for this to happen, though some charge is also recaptured in the proton case due to a deceleration of the wakefield (moving backwards and thus providing a vertical path towards the axis), which is common to both drivers.

These steps are demonstrated in the snapshots corresponding to 3.41 and 3.84 m in Fig. 5.5. At 3.84 m the larger part of each proton beamlet finds itself in a defocusing zone, causing a significant loss of charge in the ensuing distance. By contrast, smaller portions of the antiproton beamlets are in a defocusing area at this time, and some charge is still in the process of being recaptured. At 4.69 m , after this process is completed, there is evidently more charge in the antiproton-driven case.

There is one aspect of the evolution of the total driver charge below one $k_{p}^{-1}$ shown in Fig. 5.2 (lefthand graph) which, though unremarked before, can now be justified in light of the charge recapturing mechanism. Around 3 m of propagation distance, the total proton charge in Fig. 5.2 briefly surpasses that of the antiprotons. This is because that data only includes the charge below $x_{2}=1 k_{p}^{-1}$, and between 3 and 3.5 m most of the second row of modulated charge (which is outside the range of the diagnostic) has spread out for the proton case, but not for the antiproton one. This means that there is temporarily more proton charge below one $k_{p}^{-1}$, though this situation is promptly reversed when the second antiproton row begins "falling" towards the axis (around 3.4 m ).

These results have shown that an antiproton-driven wakefield is able to maintain more driver charge than a proton-driven one because there is more energy available for focusing, on one hand, and because of the radial field configuration on the other hand, which enables an especially effective charge recapturing mechanism. In addition, there are some indications that the antiproton-driven wakefield is more nonlinear than the proton-driven one.

### 5.2.3 Deeper analysis of bunch-plasma energy transfer

Addressing now the question pertaining to the antiproton wakefield amplitude, which should be higher according to the amount of charge available, a more extensive overview of the movement of driver charge in longitudinal and transverse fields is desired to test the hypothesis that too many antiprotons are being accelerated by the wakefield instead of transferring their energy to it.

A further diagnostic (described in section A.8) combs through the data of each simulation dump


Figure 5.6: Charge density of antiprotons in the transverse/longitudinal force plane at $z \approx 1.28 \mathrm{~m}$ (example).


Figure 5.7: Proportion of initial driver charge occupying each of the four quadrants in the force plane versus the propagation distance and for three different regions of the beam (designated for every row). The left and right columns show the results for the simulations with protons and antiprotons, respectively.


Figure 5.8: Smoothed line-outs of the axial electric field close to the axis for both runs at around 6.39 m of propagation distance, at the head of the beam (top) and in the middle of the beam (bottom).
and deposits driver charge on a grid spanned by the longitudinal force $W_{z}$ and the transverse force $W_{r}$, depending on the position of that charge and the force values registered at that position. The longitudinal and transverse forces are given by $W_{z}=q E_{z}$ and $W_{r}=q\left(E_{r}-B_{\varphi}\right)$. Neglecting the charge factor, $W_{z}$ and $W_{r}$ simply correspond to the quantities $E_{1}$ and $E_{2}-B_{3}$ in OSIRIS, respectively.

The four quadrants of the force plane correspond to the four possible combinations of accelerating/decelerating with focusing/defocusing fields (see Fig. A. 3 in Appendix A where this is shown schematically). A concrete example of the output from this diagnostic is presented in Fig. 5.6. After the charge density on the force plane has been determined for every simulation time, this data can be condensed further to the proportion of the total initial charge occupying each quadrant at each time, providing a temporal overview of the movement of driver particles.

This diagnostic was applied to different sections of the beam, which correspond to the "head", "middle" and "tail" regions listed in Table 2.3, Fig. 5.7 shows six different graphs for these three regions and for protons and antiprotons. Each graph depicts the proportion of initial charge in each of the four quadrants of the $W_{r} / W_{z}$ plane (this means, for example, that adding the four curves at any given distance $z$ results in the total charge left relative to the initial total charge).

There are several observations that can be surmised from Fig. 5.7, some of which are expected by now. This includes the fact that there is in general much more charge left at the end of the propagation distance in the antiproton case (especially in the middle region and excepting the head of the beam). Another somewhat expectable fact is that, after 5 m , there is less antiproton charge in defocusing fields than proton charge (note the purple and blue curves in the middle and the tail of the beam).

The first glaring insight provided by this diagnostic is the fact that protons and antiprotons behave almost exactly the same way at the head of the beam through the entire simulation. In other words, the wakes are consistent with the linear regime of wakefield theory (where opposite charges should be equivalent) exclusively in this region. Accordingly, there is almost no difference in the total charge in this region. This suggests that the causes for all the previously observed disparities between a proton and an antiproton driver are of a nonlinear character, since the wake behaves linearly at the head of the beam and gradually becomes nonlinear along it, as Fig. 5.8 demonstrates. For a fixed
time the $E_{z}$ field at the head of the beam is almost perfectly symmetrical for the opposite charges of protons and antiprotons, as predicted by linear theory (though a bit of dephasing can already be observed at the left end of the graph). By the middle of the beam, the fields are phase-shifted from their symmetrical position and display clear signs of wave steepening, traits which are distinctive of nonlinear behavior [3]. In addition, the amplitude of the field in the antiproton case has reached a higher value than the proton case.

The graphs in Fig. 5.7 yield one further piece of information regarding the middle and tail regions. Before 4 m symmetric peaks and valleys of two pairs of curves can be observed in all four graphs (in both regions and for protons and antiprotons). The pairs of curves that seem to be correlated are always the same (purple/orange and green/blue), and the two curves in each pair merely oscillate between focusing and defocusing regions while maintaining the same $W_{z}$ (either accelerating or decelerating). This suggests that the proportion of particles giving up or gaining energy barely changes at the beginning of the growth of the SMI, with the particles being periodically focused and defocused while under the influence of the same accelerating or decelerating force.

This observation is confirmed when the results of Fig. 5.7 are condensed further so as to obtain the proportion of charge that is either in accelerating or decelerating areas, without discriminating between focusing and defocusing ones. This data (for the entire beam) is displayed in Fig. 5.9 , where a phase at the beginning of the propagation distance can be observed during which the proportions change very slowly (this phase is shorter for antiprotons). With this kind of graph it is much easier to estimate how much energy the driver is transferring to the wake and vice-versa. It is of course desirable that the driving bunch transfers as much energy as possible to the wake, i.e. that it is mostly in decelerating regions, so that injected particles can in turn receive that energy through acceleration.

As evidenced in Fig. 5.9, protons transfer more energy to the wakefield than antiprotons around 4 m , but then most of the proton charge is rapidly lost. During the saturation phase of the SMI ( $6-10 \mathrm{~m}$ ) the antiproton driver transfers much more energy to the wake. However, this proportion of 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.


Figure 5.9: Proportion of initial driver charge in accelerating and decelerating regions for both runs.


Figure 5.10: Proportion of current driver charge in accelerating and decelerating regions, weighted by the axial field $E_{z}$ acting on it and unweighted. The left and right graphs correspond to the results from the simulations with protons and antiprotons, respectively.

Nonetheless, it could be argued that the diagnostic discussed above is crude in the sense that it simply adds the charge in accelerating or decelerating fields without containing any information about the magnitude of the force acting on each increment of charge. It could be, for example, that antiprotons tend to gather in lower-amplitude regions of the decelerating field for some reason, and are thus giving up even less energy than indicated by Fig. 5.9. In order to test this hypothesis and verify whether it would provide an explanation for the low antiproton wake amplitude, the previous diagnostic was extended to include the weight of the longitudinal force (in practice $E_{z}$ ) acting on each cell of the $W_{r} / W_{z}$ grid as its charge is added to the respective proportion calculation. Contrary to all the previous "proportion" diagnostics, where the value at each time (or distance $z$ ) was relative to the initial total charge, here the proportion is always relative to the current total charge. The results are shown in Fig. 5.10 , where the unweighted version was included. The difference between normalization to an initial charge versus the current charge can be visualized by comparing the dash-dotted lines for any species in Fig. 5.10 with the corresponding curve in Fig. 5.9 , since they are equivalent.

The $E_{z}$-weighted results do indeed differ from the previous ones. Nevertheless, the difference between the two portions of decelerating charge for protons and antiprotons ( $72 \%$ and $66 \%$, respectively) is still too small to account for the lower-than-expected wake amplitude with an antiproton driver, as both percentages were merely offset by 11 points.

### 5.2.4 Applying the wakefield solver to the antiproton simulations

It was initially deemed interesting to use the program presented in the last chapter (single beamlet linear wakefield solver) on the results of the antiproton run. It was hoped that the program might provide some answers regarding the wakefield amplitude of the antiproton case and its nonlinearity.

The same three dumps as in Chapter 4 were chosen for the application of the wakefield solver, which correspond to the propagation distances $6.39,8.52$ and 10.65 m . In stark contrast with the results shown in the last chapter, however, the electric fields produced by the wakefield solver are in obvious disagreement with the ones from the antiproton simulation, as evinced by Fig. 5.11. The proton simulation underlying section 4.4 is the same as the one mentioned throughout this chapter, so the data discussed in that section is reproduced in Fig. 5.11 for comparison. Whereas the wakefield solver was able to reproduce the axial field determined by the proton simulation, the average amplitudes


Figure 5.11: Comparison of the average amplitude of the axial electric field according to simulation results (lines) and the wakefield solver's results (triangles), for antiprotons and protons as drivers.


Figure 5.12: Comparison of the axial wake in the middle region with an antiproton driver as determined by the simulation (top) and the wakefield solver (bottom).
registered by the wakefield solver in the antiproton case differ at least by a factor of two. Interestingly, these amplitudes are also lower than the ones extracted from the program for the proton case, bearing in mind that linear wakefield theory is perfectly symmetric for opposite charges.

The differences between simulation and wakefield solver for the antiproton case are not simply a matter of amplitude, too. They also consist of diverging field configurations, as demonstrated for the middle region in Fig. 5.12 (the head region is quite similar for all three dumps, though).

As observed above, the wakefields generated by an antiproton driver display a higher level of nonlinearity than the ones created by protons. This, together with the fact that there is generally more antiproton charge available, and that some regions are therefore probably locally nonlinear (whenever $n_{b} / n_{0}>1$ ), leads to the assumption that linear wakefield theory breaks down for the antiproton case, thus rendering the wakefield solver useless as a tool for analysis.

### 5.3 Conclusion

In summary, the investigation of an antiproton-driven version of AWAKE has revealed that antiprotons truly are more efficient as wakefield drivers, as previous research on electron- and positron-driven wakefields suggests. However, antiprotons are also not as efficient as one could expect given the amount of driver charge available.

The larger amount of charge retained in the antiproton simulation seems to be explained mainly by the stronger focusing fields associated with higher nonlinearity of the wake. An especially favorable charge recapturing mechanism further contributes to the increased driver charge when compared to protons. The precise causes for these two effects are nevertheless still not understood.

The fact that a lower proportion of the antiproton charge left towards the end of the simulation is in decelerating fields (and thus feeding energy to the wake) in comparison with the proton case was initially considered as the reason behind the lower-than-expected amplitude of the antiproton-driven wake. However, a more detailed computation of the proportions, which included the weight of the field $E_{z}$, confirmed that the difference between both cases is too small to fully explain the observed amplitude for antiprotons.

Moreover, it was concluded that all these dissimilarities between the proton- and antiproton-driven cases are caused by nonlinear effects and can only be observed behind the head of the beam.

## Chapter 6

## Summary and future work

The work presented here can be understood as being composed of three main areas of inquiry, each one respectively seeking to address the following central questions.

- How do small variations of AWAKEs initial parameters affect the outputs of the experiment?
- Why does the amplitude of the wakefield decrease after saturation of the SMI and what tools can be devised to find out?
- How would the hypothetical substitution of the driver protons by antiprotons change the AWAKE experiment?

In order to answer these questions both linear wakefield theory and the physics of the SMI were studied in considerable depth, but most information was extracted from the extensive simulation data through a mentionable amount of original diagnostics (consisting of dozens of IDL scripts, one MATLAB program and one modestly complex parallel C program).

A sensitivity analysis as the one exposed in Chapter 2 should be present in any rigorous investigation that is based on numerical simulations, as it can reveal possible weak points of the simulation results or simply attest to their reliability. It was indeed found that the simulation results generally displayed a high degree of sensitivity to the resolution in the $x_{1}$ direction, and a lack of convergence towards the results using a higher resolution was observed especially for post-processed quantities. However, it was noted that this is only problematic when deducing absolute values for these quantities (e.g. peak amplitude of the electric field or its location) from the simulations.

The parameter scans in Chapter 3 constitute the necessary task of verifying whether a future phase of the AWAKE experiment where electrons will be injected will be successful, by attempting to anticipate any issues that might arise from realistic variations of the initial parameters. It was observed that the largest effects occurred somewhere along the propagation distance and dwindled towards the end of the simulation. The relative effects were consistently seen to be of the order of the initial parameter variation (the largest recorded factor between both was 2 ). In addition, the diagnostic designed to trace the movement and energy gain of electrons in the axial field in one dimension led to the suggestion of a more efficient location along the beam to inject them in the next phase of AWAKE. It was nevertheless concluded that electron acceleration in the AWAKE experiment is viable for proof-of-principle experiments and possibly beyond, for high-quality electron beams.

The future work concerning this chapter undoubtedly envisages a higher-dimensional version of the electron-tracking diagnostic, where the transverse forces will be considered and the electrons can be traced both along $\xi$ and $r$. This will do the complex configuration of the longitudinal and transverse fields justice and potentially reveal more information regarding the constraints for electron injection
and acceleration. This type of diagnostic should furthermore be benchmarked against PIC simulation results. It would also be interesting to develop a theory for the parameter fluctuations based on linear wakefield theory, and compare the predictions with the results obtained here.

The development of a parallel program which determines the overall wakefield by adding the contributions of single self-modulated beamlets was delineated in Chapter 4. This included the description of the algorithms that automatically identify and isolate the beamlets, and of how the computation of the wakefield was parallelized. After applying the program to a simulation of AWAKE, it was found that the spatial decline of the wakefield amplitude (along the beam) in the nonlinear stage of the SMI 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.

The clarification of the precise causes for wake incoherence and charge loss should be pursued in the future, as well as possible methods to avoid these effects. There is also some room for improvement of the program, for example by refining the beamlet isolation algorithm.

Chapter 5 set out to discover whether the currently known different efficiencies of electrons and positrons as wakefield drivers apply to protons and antiprotons as well. This was shown to be true with exceeding clarity (the negatively charged version of a proton is indeed more efficient), though the wakefield amplitude is lower than might be expected considering the increased driver charge available for antiprotons. It was found 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, the effects that cause the different wakefields for protons and antiprotons are of a nonlinear character, and are hence not observed at the front of the beam, which behaves linearly.

Despite identifying the mechanisms that lead to a more efficient wakefield excitation, it is still not fundamentally understood why an antiproton-driven wakefield behaves more nonlinearly than a protondriven one. Both this point and the gap between available driver charge and observed amplitude should be investigated further, for example by running simulations where particles are tracked (in order to see their exact movement in the wakefield).

The answers provided in this work for the three central questions mentioned above will hopefully have contributed to the knowledge of self-modulated beam-driven plasma waves and helped pave the way towards the successful demonstration of proton-driven plasma wakefield acceleration at AWAKE.

## Bibliography

[1] J. Nash, "Current and future developments in accelerator facilities," presented at the Joint meeting of the High Energy Particle Physics and Astro-Particle Physics groups of the Institute of Physics, London, 2010.
[2] T. Tajima and J. M. Dawson, "Laser electron accelerator," Physical Review Letters, vol. 43, no. 4, pp. 267-270, 1979.
[3] E. Esarey, P. Sprangle, J. Krall, and A. Ting, "Overview of plasma-based accelerator concepts," IEEE Transactions on Plasma Science, vol. 24, no. 2, pp. 252-288, 1996.
[4] N. A. Solyak, "Gradient limitations in room temperature and superconducting acceleration structures," AIP Conf. Proc., vol. 1086, pp. 365-372, 2009.
[5] J. Arthur, P. Anfinrud, P. Audebert, K. Bane, I. Ben-Zvi, et al., "Linac Coherent Light Source (LCLS) Conceptual Design Report," SLAC National Accelerator Laboratory," Design Report, 2002.
[6] J. Vieira, private communication.
[7] Y. Kitagawa, T. Matsumoto, T. Minamihata, K. Sawai, K. Matsuo, et al., "Beat-wave excitation of plasma wave and observation of accelerated electrons," Physical Review Letters, vol. 68, pp. 4851, 1992.
[8] W. B. Mori, C. D. Decker, D. E. Hinkel, and T. Katsouleas, "Raman forward scattering of short-pulse high-intensity lasers," Physical Review Letters, vol. 72, pp. 1482-1485, 1994.
[9] A. Modena, Z. Najmudin, A. E. Dangor, C. E. Clayton, K. A. Marsh, et al., "Observation of Raman forward scattering and electron acceleration in the relativistic regime," IEEE Transactions on Plasma Science, vol. 24, no. 2, pp. 289-295, 1996.
[10] N. E. Andreev, L. Gorbunov, V. Kirsanov, A. A. Pogosova, and R. R. Ramazashvili, "Resonant excitation of wakefields by a laser pulse in a plasma," JETP Letters, vol. 55, pp. 571-576, 1992.
[11] K. Nakajima, D. Fisher, T. Kawakubo, H. Nakanishi, A. Ogata, et al., "Observation of ultrahigh gradient electron acceleration by a self-modulated intense short laser pulse," Physical Review Letters, vol. 74, pp. 4428-4431, 1995.
[12] D. Strickland and G. Mourou, "Compression of amplified chirped optical pulses," Optics Communications, vol. 55, no. 6, pp. 447-449, 1985.
[13] W. P. Leemans, A. J. Gonsalves, H.-S. Mao, K. Nakamura, C. Benedetti, et al., "Multi-GeV electron beams from capillary-discharge-guided subpetawatt laser pulses in the self-trapping regime," Physical Review Letters, vol. 113, p. 245002, 2014.
[14] R. Bingham, J. T. Mendonça, and P. K. Shukla, "Plasma based charged-particle accelerators," Plasma Physics and Controlled Fusion, vol. 46, no. 1, p. R1, 2004.
[15] P. Chen, J. M. Dawson, R. W. Huff, and T. Katsouleas, "Acceleration of electrons by the interaction of a bunched electron beam with a plasma," Physical Review Letters, vol. 54, pp. 693-696, 1985.
[16] J. B. Rosenzweig, D. B. Cline, B. Cole, H. Figueroa, W. Gai, et al., "Experimental observation of plasma wake-field acceleration," Physical Review Letters, vol. 61, pp. 98-101, 1988.
[17] N. Barov, M. E. Conde, W. Gai, and J. B. Rosenzweig, "Propagation of short electron pulses in a plasma channel," Physical Review Letters, vol. 80, pp. 81-84, 1998.
[18] V. Yakimenko, I. V. Pogorelsky, I. V. Pavlishin, I. Ben-Zvi, K. Kusche, et al., "Cohesive acceleration and focusing of relativistic electrons in overdense plasma," Physical Review Letters, vol. 91, p. 014802, 2003.
[19] I. Blumenfeld, C. E. Clayton, F. J. Decker, M. J. Hogan, C. Huang, et al., "Energy doubling of 42 GeV electrons in a metre-scale plasma wakefield accelerator," Nature, vol. 445, pp. 741-744, 2007.
[20] M. Litos, E. Adli, W. An, C. I. Clarke, C. E. Clayton, et al., "High-efficiency acceleration of an electron beam in a plasma wakefield accelerator," Nature, vol. 515, no. 7525, pp. 92-95, 2014.
[21] S. Corde, E. Adli, J. M. Allen, W. An, C. I. Clarke, et al., "Multi-gigaelectronvolt acceleration of positrons in a self-loaded plasma wakefield," Nature, vol. 524, no. 7566, pp. 442-445, 2015.
[22] R. Ruth, A. Chao, P. Morton, and P. Wilson, "A plasma wake field accelerator," Particle Accelerators, vol. 17, pp. 171-189, 1985.
[23] P. Chen, J. J. Su, and J. M. Dawson, "Energy transfer in the plasma wake field accelerator," Physical Review Letters, vol. 56, no. 12, pp. 1252-1255, 1986.
[24] J. Rosenzweig, "Nonlinear plasma dynamics in the plasma wake-field accelerator," Physical Review Letters, vol. 58, no. 6, pp. 555-558, 1987.
[25] N. Patel, "Accelerator physics: The plasma revolution," Nature, vol. 449, no. 7159, pp. 133-135, 2007.
[26] A. Caldwell, K. V. Lotov, A. Pukhov, and F. Simon, "Proton-driven plasma-wakefield acceleration," Nature Physics, vol. 5, pp. 363-367, 2009.
[27] N. Kumar, A. Pukhov, and K. V. Lotov, "Self-modulation instability of a long proton bunch in plasmas," Physical Review Letters, vol. 104, p. 255003, 2010.
[28] J. Vieira, W. B. Mori, and P. Muggli, "Hosing instability suppression in self-modulated plasma wakefields," Physical Review Letters, vol. 112, p. 205001, 2014.
[29] Y. Fang, J. Vieira, L. D. Amorim, W. Mori, and P. Muggli, "The effect of plasma radius and profile on the development of self-modulation instability of electron bunches," Physics of Plasmas, vol. 21, p. 056703, 2014.
[30] "Underground structures at CERN." [Online]. Available: https://proj-cngs.web.cern.ch/proj-cngs/ Download/UndergroundStructures/SITUA\%20170603m.jpg
[31] R. Assmann, R. Bingham, T. Bohl, C. Bracco, B. Buttenschön, et al., "Proton-driven plasma wakefield acceleration: a path to the future of high-energy particle physics," Plasma Physics and Controlled Fusion, vol. 56, p. 084013, 2014.
[32] AWAKE website. [Online]. Available: http://awake.web.cern.ch/awake/
[33] P. Muggli, A. Caldwell, O. Reimann, E. Oz, R. Tarkeshian, et al., "Physics of the AWAKE project," in Proceedings of IPAC2013, Shanghai, China, 2013, pp. 1179-1181.
[34] C. B. Schroeder, C. Benedetti, E. Esarey, F. J. Grüner, and W. P. Leemans, "Growth and phase velocity of self-modulated beam-driven plasma waves," Physical Review Letters, vol. 107, p. 145002, 2011.
[35] A. Pukhov, N. Kumar, T. Tückmantel, A. Upadhyay, K. V. Lotov, et al., "Phase velocity and particle injection in a self-modulated proton-driven plasma wakefield accelerator," Physical Review Letters, vol. 107, p. 145003, 2011.
[36] P. Muggli, E. Adli, A. Alexandrova, L. Amorim, R. Apsimon, et al., "AWAKE readiness for the study of the seeded self-modulation of a 400 GeV proton bunch," 2017, unpublished.
[37] K. V. Lotov and A. Sosedkin, "Simulation of self-modulating particle beams in plasma wakefield accelerators," in Proceedings of IPAC2013, Shanghai, China, 2013, pp. 1238-1240.
[38] K. V. Lotov, A. Pukhov, and A. Caldwell, "Effect of plasma inhomogeneity on plasma wakefield acceleration driven by long bunches," Physics of Plasmas, vol. 20, no. 1, p. 013102, 2013.
[39] K. V. Lotov, A. P. Sosedkin, and A. V. Petrenko, "Long-term evolution of broken wakefields in finiteradius plasmas," Physical Review Letters, vol. 112, p. 194801, 2014.
[40] K. V. Lotov, G. Z. Lotova, V. I. Lotov, A. Upadhyay, T. Tückmantel, et al., "Natural noise and external wakefield seeding in a proton-driven plasma accelerator," Physical Review Accelerators and Beams, vol. 16, p. 041301, 2013.
[41] K. V. Lotov, "Effect of beam emittance on self-modulation of long beams in plasma wakefield accelerators," Physics of Plasmas, vol. 22, no. 12, p. 123107, 2015.
[42] N. Savard, J. Vieira, and P. Muggli, "Effect of proton bunch parameter variation on AWAKE," in Proceedings of PAC, submitted, 2016.
[43] A. Caldwell, E. Gschwendtner, and P. Muggli, "Milestone reached by the AWAKE collaboration." [Online]. Available: http://awake.web.cern.ch/awake/images/Milestone1.pdf
[44] R. A. Fonseca, L. O. Silva, F. S. Tsung, V. K. Decyk, W. Lu, et al., "OSIRIS: A three-dimensional, fully relativistic particle in cell code for modeling plasma based accelerators," in Computational Science — ICCS 2002: International Conference Amsterdam, The Netherlands, April 21-24, 2002 Proceedings, Part III. Springer Berlin Heidelberg, 2002, pp. 342-351.
[45] C. K. Birdsall and A. B. Langdon, Plasma Physics via Computer Simulation, 3rd ed., P. Stott and H. Wilhelmsson, Eds. Bristol and Philadelphia: Institute of Physics Publishing, 1998.
[46] K. Yee, "Numerical solution of initial boundary value problems involving maxwell's equations in isotropic media," IEEE Transactions on Antennas and Propagation, vol. 14, no. 3, pp. 302-307, May 1966.
[47] R. A. Fonseca, J. Vieira, F. Fiuza, A. Davidson, F. S. Tsung, et al., "Exploiting multi-scale parallelism for large scale numerical modelling of laser wakefield accelerators," Plasma Physics and Controlled Fusion, vol. 55, no. 12, p. 124011, 2013.
[48] K. V. Lotov, "Physics of beam self-modulation in plasma wakefield accelerators," Physics of Plasmas, vol. 22, no. 10, p. 103110, 2015.
[49] K. V. Lotov, "Excitation of two-dimensional plasma wakefields by trains of equidistant particle bunches," Physics of Plasmas, vol. 20, no. 8, p. 083119, 2013.
[50] R. A. Fonseca, private communication.
[51] S. Lee, T. Katsouleas, R. G. Hemker, E. S. Dodd, and W. B. Mori, "Plasma-wakefield acceleration of a positron beam," Physical Review E, vol. 64, p. 045501, 2001.
[52] J. Lettry, D. Aguglia, J. Alessi, P. Andersson, S. Bertolo, et al., "Linac4 H-ion sources," Review of Scientific Instruments, vol. 87, no. 2, p. 02B139, 2016.
[53] R. W. Hockney and J. W. Eastwood, Computer Simulation Using Particles, 5th ed. Bristol and Philadelphia: Institute of Physics Publishing, 1999.

## Appendix A

## Description of the diagnostics developed for this work

During the course of this work some diagnostics were developed to evaluate the simulation results from OSIRIS, in addition to the already present post-processing capabilities of the VisXD set of tools. In some cases one wished to obtain a very specific quantity which could be derived from the information in the simulation files; in other cases the objective was rather to demonstrate a trend more clearly by condensing some simulation result into a lower-dimensional quantity (taking the average of the distribution in an entire simulation window, for example).

Unless stated otherwise, these small programs (or scripts) were written in the same vectorized programming language used for the OSIRIS visualization tools, called IDL For the sake of rigor and reproducibility, these scripts are described in more detail in this appendix.

## A. 1 Peak and average axial electric field

The purpose of this simple diagnostic is to give an overview of the development of the axial electric field over the propagation distance (i.e. for each recorded time, or dump ${ }^{17}$. Going through all the $E_{1}$ simulation files, the average field for each dump is calculated by summing the absolute values of the field in all cells and dividing the sum by the number of cells. The peak electric field is the maximum absolute value found for the current dump, through an IDL built-in function. Since in this type of simulations (where there is an "axial" boundary) the data close to the axis is often quite noisy and reaches unphysically high values, the data below a given number of cells in the $x_{2}$ direction can optionally be excluded from the analysis. The average and peak electric field values for each dump are then exported to a VisXD-readable format.

## A. 2 Concatenation of axial electric field line-outs

Once again for every dump file from a given simulation, this script retrieves a line-out of the axial electric field along the $x_{1}$ direction radially close to the axis (also with an optional user-specified offset to avoid noisy and unphysical data) and concatenates the line-outs along the propagation distance. The final result is a matrix of $E_{1}$ values where one direction is the beam co-moving coordinate $\xi=z-c t$

[^2]Example of concatenated line-outs of $\mathrm{E}_{1}$


Figure A.1: Example of the product from the diagnostic described in this section for a segment of the simulation window, overlaid with three slopes at different points with different meanings for the wake's phase velocity $v_{\phi}$.
and the other direction is the propagation distance $z$ (or, equivalently, elapsed simulation time or dump number). This diagnostic is useful for perceiving changes of the wake's phase velocity along the propagation distance or changes of the field's local phase along the wake itself.

An example of the data obtained by this diagnostic is shown in Fig. A.1 where the beam coordinate is presented w.r.t. the head of the beam $\xi_{h}$. The phase velocity $v_{\phi}$ can be estimated by following the trajectory of a fixed phase (i.e. a position along the field wave where the value is constant through time, for example where $E_{1}(\xi)=0$ ) and measuring its local slope (see the three different cases illustrated in Fig. A.1). If a certain phase is perfectly vertical in the graph, then the phase velocity at that position $\xi$ is equal to the speed of light, since that phase is perfectly synchronized with the simulation window (which moves at $c$ ). This case is shown in the middle of Fig. A.1 On the other hand, a positive or a negative slope will mean that the phase in question is either overtaking or losing pace with the simulation window, respectively. These two cases are also exemplified in Fig. A. 1 on the left and right.

There is of course a more quantitative way of estimating the phase velocity from the slope measured in a graph such as the one in Fig.A.1. The objective is to find an expression for the phase velocity that depends on the quantities $\Delta z$ and $\Delta \xi$, which can be promptly surmised from these graphs. Firstly, it should be noted that $z$ is simply the distance traveled by the simulation window in a certain time (at $c$ ),


Figure A.2: Schematic illustration of the quantities required to calculate the wake's phase velocity using the concatenated line-outs of $E_{1}$.
so it simultaneously marks the absolute position where the simulation window at a given time begins. Noting also that $v_{\phi}$ is simply the absolute distance traveled by a certain a phase in a certain time, it can be defined as $v_{\phi}=\Delta x / \Delta t$, where $\Delta t=\Delta z / c$. The absolute distance is in turn given by

$$
\begin{equation*}
\Delta x=x_{\max }-x_{\min }=\left[\xi_{\max }+z\left(\xi_{\max }\right)\right]-\left[\xi_{\min }+z\left(\xi_{\min }\right)\right]=\Delta \xi+\Delta z \tag{A.1}
\end{equation*}
$$

where $\Delta \xi=\xi_{\max }-\xi_{\min }$ and $\Delta z=z\left(\xi_{\max }\right)-z\left(\xi_{\min }\right)$. These variables are also defined in Fig. A. 2 Note that $\Delta z$ can assume a negative value. Substituting in the definition of the phase velocity,

$$
\begin{equation*}
v_{\phi}=\frac{\Delta \xi+\Delta z}{\Delta z} c \quad \Leftrightarrow \quad \beta_{\phi}=\frac{v_{\phi}}{c}=\frac{\Delta \xi}{\Delta z}+1 \tag{A.2}
\end{equation*}
$$

an expression for the wake's phase velocity emerges that depends only on $\Delta \xi$ and $\Delta z$. For $\Delta \xi=0$ (a vertical "slope"), for example, $v_{\phi}=c$.

## A. 3 Total charge of particle species

The amount of charge in the driver bunch is directly correlated with the amplitude of the wakefield, so this quantity is worth analyzing over the course of the simulation. This diagnostic can take two simulation outputs as input: either it sums the charges stored in the raw macroparticle data for every recorded simulation time and calibrates them to a real charge value, or it integrates the charge density data. The former option is preferable, since the raw data is generally more accurate. In this way the total charge contained in the simulation box of a particle species can be monitored along the propagation distance.

The total charge of a charge distribution inside a volume in cylindrical coordinates is calculated through

$$
\begin{equation*}
Q=\int \rho d V=\int_{0}^{L} \int_{0}^{2 \pi} \int_{0}^{R} \rho(r, z) r d r d \theta d z=2 \pi \int_{0}^{L} \int_{0}^{R} r \rho(r, z) d r d z \tag{A.3}
\end{equation*}
$$

This equation can be discretized and applied to OSIRIS, becoming

$$
\begin{equation*}
Q=2 \pi e \Delta x_{1} \Delta x_{2} \sum_{i=1}^{n_{1}} \sum_{j=1}^{n_{2}} x_{2}(i) \rho_{p^{+}}(i, j) \tag{A.4}
\end{equation*}
$$

where all quantities are normalized according to OSIRIS (except for the elementary charge $e$ and $Q$, of course). This is the formula applied in the calculation of the total charge when the charge density data is used as input.

As the SMI develops, some portions of the beam will be defocused from the axis and, in the simulation, will eventually leave the box, since the wake has periodic regions of defocusing electric field. Since it takes some time for the defocused protons to reach the border of the simulation box and leave it, the total charge calculated for the whole box would only indicate some loss after a certain propagation distance. For this reason, a radial boundary can be set in the diagnostic, beyond which no charge will be taken into account when summing the macroparticle charges. This allows the observation of beam charge loss from the start of the simulation (and SMI development).

## A. 4 Maximum witness electron energy

This diagnostic combs through the raw macroparticle data of the witness beam species at every file dump and finds the maximum energy attained by a witness electron. The result is a graph of the highest witness beam energy over the propagation distance.

## A. 5 Local phase shift of the wake

This script, written in the programming language MATLAB, produces a distribution of the local phase shift w.r.t. the head of the beam in the axial electric field $E_{1}$ close to the axis versus the propagation distance $z$ (for every recorded file dump) and the position along the beam $\xi$. This is accomplished by subdividing the line-out of $E_{1}$ from each dump into "windows" containing a certain number (given as input) of plasma wavelengths $\lambda_{p}$ (which should coincide with the wavelength of the wake itself) rounded to an integer number of cells, and by fitting the function

$$
\begin{equation*}
f(\xi)=A \sin \left[k_{p}\left(\xi-\xi_{h}\right)+\phi\right] \tag{A.5}
\end{equation*}
$$

where $A$ and $\phi$ are the fitting parameters, to the smoothed data in each window. The amplitude parameter $A$ serves merely to achieve a better fit and for the algorithm to converge, while the value of the second fitting parameter will be appended to a matrix $\phi(z, \xi)$, according to the current dump $(z)$ and window $(\xi)$. The line-out of $E_{1}$ is taken with a distance of five cells from the axis, due to the aforementioned problems involving numerical noise.

## A. 6 Hypothetical energy gain and minimum injection point of witness electrons

The goal of this diagnostic is to determine where along the propagation distance an electron could be injected without losing any energy in the ensuing distance and how much energy it would then gain, depending on its final position along the beam $\xi_{f}$ and its final energy $\gamma_{f}$ (where final means after 10 m of plasma). The results here only apply to the wake close to the axis, since this diagnostic is based on the concatenated line-outs of $E_{1}$ close to the axis, described in section A. 2 .

Though quite simple, this diagnostic can take longitudinal dephasing of the electrons w.r.t. the wake into account, by integrating the 1D relativistic equations of motion for the electrons in the $z$ direction using the leapfrog algorithm. Taking advantage of this algorithm's time-reversal invariance [53, pp. 9496], the electrons begin at their final positions along the beam and are propagated backwards in time until a minimum injection point $z_{\text {inj }}$, before which they would have lost energy. A future injected electron beam will have a minimum energy of 15 MeV at AWAKE [36], so the corresponding minimum threshold $\gamma_{\text {min }}$ of 29 was also imposed on the LORENTZ factor of the electrons to help determine $z_{\mathrm{inj}}$.

The accuracy of the results from this diagnostic is of course heavily dependent on the number of file dumps from the simulation, because this determines the size of the time step covered by each iteration of the particle pusher. The larger this time step is, the more uncertainty there is as to the exact position $z_{\mathrm{inj}}$ where the electron would have encountered a decelerating region, and consequently as to the total energy gain.

This tool produces three outputs as functions of both the final position along the beam $\xi_{f}$ and the final energy $\gamma_{f}$, the ranges and refinement of which can be set arbitrarily (though the maximum refinement of the $\xi_{f}$ axis is limited to the number of cells in the $x_{1}$ direction). The output quantities are
the energy gain $\Delta \gamma$ (i.e. in terms of a LORENTz factor), the minimum injection point $z_{\text {inj }}$ and the implied initial energy $\gamma_{0}$ (at $z_{\text {inj }}$ ).

Some tests were performed to exclude any errors in the code, all of which were passed by the script:

- increasing the final energy of the electrons so much that there is virtually no dephasing between them and the simulation window (moving at $c$ ) and comparing the results with a previous, simpler version of this diagnostic which did not take dephasing into account: the results were identical;
- saving the trajectories of single electrons w.r.t. the moving window (or the co-moving beam coordinate $\xi$ ) in order to check for obviously wrong behavior (such as forward-moving trajectory, which would mean that the electron would be moving at more than $c$ ): no abnormal trajectories were seen and their behavior seemed plausible;
- removing the electric fields, letting electrons with a fixed, relatively low final energy propagate backwards in time through vacuum, and comparing the spatial dephasing resulting from the algorithm with a value calculated manually: the values matched.

Mentioning these tests is especially relevant since the results from this diagnostic were not reproduced in PICsimulations.

## A. 7 Energy contained in focusing and defocusing regions of $E_{r}$

In this diagnostic the energies contained in focusing and defocusing (w.r.t. the charge of the driver) areas of the radial component of the electric field are computed for every simulation dump, through a spatial integration of $E_{r}$ (or $E_{2}$ in OSIRIS). The ratio between these energies (focusing to defocusing) is also calculated and saved.

The energy in the radial electric field is calculated through the spatial integration of the energy density $\eta$ of an electrostatic field, given by $\eta=\frac{1}{2} \varepsilon_{0} E^{2}$, where $\varepsilon_{0}$ is the vacuum permittivity. The energy in $E_{r}$ is then given by (using the volume element in cylindrical coordinates)

$$
\begin{equation*}
W=\int \eta d V=\frac{1}{2} \varepsilon_{0} \int_{0}^{L} \int_{0}^{2 \pi} \int_{0}^{R}\left[E_{r}(r, z)\right]^{2} r d r d \theta d z=\pi \varepsilon_{0} \int_{0}^{L} \int_{0}^{R} r\left[E_{r}(r, z)\right]^{2} d r d z \tag{A.6}
\end{equation*}
$$

For the numerical implementation, the integral can be discretized and the spatial and electric field variables expressed in terms of their normalized counterparts in OSIRIS:

$$
\begin{equation*}
W=\pi \varepsilon_{0} \sum_{i=1}^{n_{1}} \sum_{j=1}^{n_{2}} \frac{x_{2}(i)}{k_{p}}\left[E_{2}(i, j)\right]^{2} E_{0}^{2} \frac{\Delta x_{2}}{k_{p}} \frac{\Delta x_{1}}{k_{p}}=\pi \varepsilon_{0} \frac{E_{0}^{2}}{k_{p}^{3}} \Delta x_{1} \Delta x_{2} \sum_{i=1}^{n_{1}} \sum_{j=1}^{n_{2}} x_{2}(i)\left[E_{2}(i, j)\right]^{2} \tag{A.7}
\end{equation*}
$$

with the electric field normalization factor $E_{0}=m_{e} c \omega_{p} / e$. Depending on the sign of the electric field (w.r.t. the charge of the driver particle), the energy in each cell is either added to the focusing energy $W_{\text {foc }}$ or the defocusing energy $W_{\text {defoc }}$. The ratio is then simply $R=W_{\text {foc }} / W_{\text {defoc }}$.

The plasma species does not cover the entire height of the simulation window in the simulations of AWAKE presented in this work, so a vestigial positive electric field is often observed in the plasma-free strip at the top of the simulation box. In order to avoid skewing the real amount of energy in the radial field inside the plasma, the plasma-free regions are excluded in the diagnostic.

## A. 8 Charge density in the longitudinal and transverse force plane

The insight provided by this diagnostic consists of knowing how much charge is under the influence of which forces at any given point during the simulation. Similarly to the total charge diagnostic, this method can use two data sources from the simulations: either the raw particle data or the charge density data, whereby the raw data is expected to deliver more precise results and the charge density data leads to smoother distributions. For every file dump, the general idea is to find the forces acting on a certain particle (for the raw data) or a certain cell (for the charge density data), and deposit that particle's or that cell's charge into a grid spanned by two ranges of the axial and radial forces (the plane $W_{r} / W_{z}$, see Fig. A.3).

In order to achieve this it is necessary to know the spatial distributions of these forces. The longitudinal force is given by $W_{z}=q E_{z}$, where $q$ is the charge of the driver particles. This equation remains identical in its normalized form, i.e.

$$
\begin{equation*}
W_{z} \frac{1}{m_{e} c \omega_{p}}=q \frac{e}{e} \frac{1}{m_{e} c \omega_{p}} E_{z} \quad \Leftrightarrow \quad \hat{W}_{z}=\hat{q} \hat{E}_{z} \tag{A.8}
\end{equation*}
$$

In OSIRIS the longitudinal force is therefore simply $E_{1}$ (multiplied if necessary by a minus sign). The transverse force acting on a relativistic particle beam propagating along the $z$ direction (with $v \sim v_{z}$ ) is given by (in the Gaussian unit system):

$$
\begin{equation*}
\mathbf{W}_{\perp}=q\left(\mathbf{E}_{\perp}+\mathbf{e}_{z} \times \mathbf{B}_{\perp}\right) \quad \Rightarrow \quad W_{r}=q\left(E_{r}-B_{\varphi}\right) \quad W_{\varphi}=q\left(E_{\varphi}+B_{r}\right) \tag{A.9}
\end{equation*}
$$

where $\mathbf{e}_{z}$ is the unit vector in the $z$ direction.
The third component of the force $W_{\varphi}$ only influences movement around the axis, which is not interesting for this analysis or for cylindrically symmetric simulations. Since the normalized expression for $W_{r}$ is again perfectly equivalent, the distribution of the transverse force in OSIRIS can be determined merely by subtracting $B_{3}$ from $E_{2}$ (and multiplying by a minus sign for antiprotons).

The resulting charge density distribution on the force plane can reveal some information about the movement of the driver particles during the simulation (see an example of this distribution in Fig. 5.6 on $p$.57). However, this information is too detailed to draw clear conclusions, and some condensing diagnostics would be more useful. This is why further tools were developed to postprocess the force plane.

The four quadrants of the force plane correspond to the four possible combinations of accelerating/decelerating with focusing/defocusing fields (specified in Fig. A. 3 for protons and antiprotons). By measuring how much charge is in each quadrant at any given time, it is possible to obtain an overview of the interaction of the driver beam with the wake. If most of the charge is in both decelerating quadrants, for example, then the wake must be absorbing energy from the driver, and vice-versa.

To provide this overview, two scripts were written to analyze the charge distribution on the force plane. One produces the proportion of charge occupying each of the four quadrants and the other produces the proportion of charge occupying accelerating or decelerating fields, both versus the propagation distance. An extension of the latter diagnostic was later devised which weighs the charge with the value of the longitudinal force $W_{z}$ acting on it.


Figure A.3: Schematic representation of the force plane $W_{r} / W_{z}$ and meaning of each quadrant for protons (purple) and antiprotons (teal).


[^0]:    4.18 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.

[^1]:    ${ }^{1}$ Note that the nomenclature "head", "middle" and "tail" in this chapter does not refer to the fixed intervals stated in Table 2.3 in Chapter 2 but to general regions in the simulation window.

[^2]:    ${ }^{1}$ A dump (as in file dump) is an iteration where OSIRIS saves all the simulation outputs to result files, and happens every $n$ iterations (where $n$ is specified by the user). A simulation with $n=10$ that has completed 1000 iterations will result in 100 dumps, for example.

