


UNIVERSIDADE DE LISBOA
INSTITUTO SUPERIOR TÉCNICO



**Circumventing dephasing and
pump-depletion for long and stable
single-stage plasma acceleration**
from flying-focus laser pulses to proton-bunch drivers

Chiara Badiali

Supervisor: Doctor Jorge Miguel Ramos Domingues Ferreira Vieira

Co-Supervisor: Doctor Thales Marques Corrêa da Silva

Thesis approved in public session to obtain the PhD Degree in

Physics

Jury final classification:

Pass with Distinction and Honour



UNIVERSIDADE DE LISBOA
INSTITUTO SUPERIOR TÉCNICO

**Circumventing dephasing and pump-depletion for long and
stable single-stage plasma acceleration
from flying-focus laser pulses to proton-bunch drivers**

Chiara Badiali

Supervisor: Doctor Jorge Miguel Ramos Domingues Ferreira Vieira

Co-Supervisor: Doctor Thales Marques Corrêa da Silva

Thesis approved in public session to obtain the PhD Degree in

Physics

Jury final classification:

**Pass with Distinction and Honour
Jury**

Chairperson: Doctor Luís Miguel de Oliveira e Silva, Instituto Superior Técnico, Universidade de Lisboa

Members of the Committee:

Doctor Robert Bingham, Faculty of Science, University of Strathclyde, Glasgow, Reino Unido

Doctor Luís Miguel de Oliveira e Silva, Instituto Superior Técnico, Universidade de Lisboa

Doctor Massimo Ferrario, INFN – Laboratori Nazionali di Frascati, Itália

Doctor Jorge Miguel Ramos Domingues Ferreira Vieira, Instituto Superior Técnico, Universidade de Lisboa

Funding Institutions

FCT Fundação para a Ciência e a Tecnologia

Resumo

A aceleração por campos de esteira em plasma oferece gradientes aceleradores várias ordens de grandeza superiores aos da tecnologia convencional; contudo, a sua aplicação prática é limitada pelo esgotamento da energia do feixe excitador da esteira e pela perda de sincronização entre o feixe excitador e o feixe acelerado. Esta tese explora duas vias paralelas para alcançar uma aceleração estável num único estágio.

Primeiro, utilizamos feixes excitadores que são lasers espaço-temporais com velocidade controlada para gerar campos de esteira com velocidade de fase ajustável. Estes pulsos de forma invariantes permitem uma aceleração prolongada sem esgotamento da energia do feixe. Concebemos um esquema de alinhamento da fase aceleradora para partículas de vida curta, sub-relativísticas, acoplando feixes subluminais a um perfil de densidade de plasma personalizado. Modelos teóricos e simulações partícula-na-célula demonstram a aceleração sustentada de partículas várias ordens de grandeza mais pesadas do que os elétrons, mantendo simultaneamente elevada captura (atingindo até 100%) e preservação da emitância transversal.

Uma alternativa viável para mitigar o esgotamento da energia do feixe excitador consiste em utilizar drivers de feixes de alta energia, conceito subjacente à experiência AWAKE no CERN. Aqui, um feixe de prótons longo e altamente energético acelera elétrons ao longo de dezenas de metros, recorrendo à instabilidade de auto-modulação para formar micro-feixes que excitam, de forma ressonante, um campo acelerador estável. Nesta tese, investigamos a instabilidade transversal de hosing, que é disruptiva e compete com o processo de auto-modulação. Estudamos a coevolução destas duas instabilidades através de dados experimentais e de simulações partícula-na-célula quase-estáticas.

Em conjunto, estes resultados fornecem soluções robustas para acelerar partículas de vida curta e estabilizar feixes excitadores de alta energia. Por isso, abrem caminho para a próxima geração de fontes de partículas baseadas em aceleração em plasma.

Palavras-chave: aceleração baseada em plasma; flying focus; aceleração de muões; instabilidade no plasma; simulações particle-in-cell (PIC)

Abstract

Plasma wakefield acceleration offers accelerating gradients orders of magnitude beyond conventional technology, opening a path to compact high-energy particle sources. Realizing this promise in a single accelerating stage, however, is constrained by two fundamental processes: the depletion of the driver as it transfers energy to the plasma, and the loss of synchronism between the driver and the accelerated beam. This thesis develops two complementary strategies to circumvent these limits and achieve long, stable acceleration in a single stage.

First, we employ velocity-controlled spatio-temporal laser drivers to obtain wakefields with controllable phase velocity and apply them to the acceleration of short-lived particle beams. These shape-invariant pulses enable extended, depletion-free acceleration. We design a phase-locking scheme for decaying, sub-relativistic particles by coupling subluminal drivers with plasma density tapering. Theoretical models and particle-in-cell simulations demonstrate the sustained acceleration of particles orders of magnitude heavier than electrons, while still maintaining high capture (reaching 100% in the optimal beam-loading regime) and preservation of transverse emittance.

Another viable alternative to attenuate the role of pump depletion is to use proton bunches as drivers, the concept behind the AWAKE experiment at CERN. Here, a long, highly energetic proton beam accelerates electrons over tens of meters, relying on the self-modulation instability to form micro-bunches that resonantly excite a stable accelerating field. In this thesis, we investigate the disrupting transverse hosing instability, which competes with the self-modulation process. We study the co-evolution of these two instabilities through experimental data and quasi-static particle-in-cell simulations.

Collectively, these results provide robust solutions for accelerating short-lived particles and stabilizing high-energy drivers. As a consequence, this work paves the way for the next generation of plasma-based particle sources.

Keywords: plasma-based acceleration; muon acceleration; flying focus; plasma instability; particle-in-cell simulations;

ACKNOWLEDGMENTS

I would like to start by expressing my deepest gratitude to my supervisors, Thales Silva and Jorge Vieira, for their guidance, support, and encouragement throughout these years. Their scientific insights, patience, and dedication have shaped not only this work but also my understanding of research and mentorship. I hope in the future I will be able to carry forward their commitment to forming students and to creating an environment where people can grow, think freely, and reinforce their love for science.

I also sincerely thank Patric Muggli for his support, perspective, and encouragement during this journey. His experience and scientific vision have been a valuable source of inspiration. I also thank Marija Vranic for her mentorship over the years.

I would also like to thank all the incredible people I met in GoLP. I think over the years we have made GoLP a place of support, friendship, and shared memories, and I am deeply grateful to have been part of it. I will always remember the mandatory breaks at 4pm, the lunch together, the Fridays playing pool or the karaoke nights at Júlio's, and all the moments that made these years lighter. We should be proud of what we built together.

I would also like to thank my office mates for the everyday company, mutual respect, and for making the daily routine of these years much more enjoyable. I thank Rafael for the many discussions throughout this work, from which we often left with more doubts than we had before starting, and for sharing the *custody* of our lasers. I am sure they will be in very good hands from now on. I would also like to thank Mariana, not only as a colleague but also as a friend, for the conversations, the support, and for being someone with whom I could always vent and feel understood. I would also like to thank Sofiia, my colleague and aerial silks companion, for sharing with me a much-needed space outside physics. I would like to thank André for his support during the final stretch of this thesis, especially on the days when stress made me less than easy to be around. Thank you for taking care of me when I really needed it, for always finding a way to make me laugh, and for making even the longest days pass in a second.

I would also like to thank my colleagues in AWAKE for welcoming me during the experimental run and for sharing such an incredible experience with me. In particular, I would like to thank Carolina, Nikita, and Arthur for patiently guiding me through my first experimental data analysis.

Gostaria também de agradecer à Cláudia e à Susana por todo o apoio ao longo destes anos. O trabalho que fazem é essencial e muitas vezes invisível, mas faz uma diferença enorme no

nosso dia a dia. Quero também agradecer à Patrícia Simões pelo apoio nas reuniões NDA e pelo acompanhamento ao longo deste percurso. A sua orientação, disponibilidade e atenção foram muito importantes para mim, e sou muito grata por ter podido contar com esse apoio.

Quero também agradecer à minha colega de casa e amiga, Fabia, que ao longo destes anos se tornou praticamente uma irmã. Mesmo quando não percebia bem o que eu estava a fazer em física, esteve sempre lá para me ouvir desabafar, para acolher as minhas reclamações depois dos dias mais difíceis, ou simplesmente para ver uma série comigo e ajudar-me a desligar um pouco de tudo.

Vorrei ringraziare Elena, Lorenzo e Federica per il loro sostegno, che sono sempre riusciti a farmi sentire sempre, anche se lontani. Vorrei ringraziare anche Arianna per aver condiviso con me gioie e dolori (e tatami) in questi anni. In fondo, poche cose uniscono quanto una sofferenza condivisa, soprattutto quando si riesce a riderci sopra. Infine, vorrei ringraziare la mia famiglia per il loro supporto, per aver creduto in me anche quando gli ho detto che mi sarei trasferita in un paese che conoscevamo a malapena, e per essermi sempre stati accanto. Anche quando non capivano del tutto cosa stessi facendo, non hanno mai smesso di fidarsi di me e di incoraggiarmi.

This work was supported by FCT I.P. under Project 2024.07895.CPCA.A3 – DOI: <https://doi.org/10.54499/2024.07895.CPCA.A3> – at MareNostrum 5 supercomputer, jointly funded by EuroHPC JU, Portugal, Turkey and Spain. This work was supported by the HE EuPRAXIA-PP under grant agreement No. 101079773. The support of the Portuguese Science Foundation, FCT, through Grant No. PRT/BD/152270/2021, DOI: 10.54499/PRT/BD/152270/2021, is also gratefully acknowledged.

CONTENTS

1	Introduction	3
1.1	State of the Art in Plasma Acceleration	4
1.2	Anatomy of a Plasma Wakefield Accelerator	5
1.3	Motivation and Scope of the Thesis	6
1.4	Methods	7
1.4.1	Particle-in-cell simulations with OSIRIS	7
1.4.2	Quasi-static simulations with QuickPIC	9
1.4.3	Thesis Outline	10
1.5	List of Contributions	11
2	Simulation Design for Velocity Controlled Drivers	13
2.1	Spectral construction and numerical implementation of spatio-temporal pulses	14
2.1.1	Spectral design of spatio-temporal pulses	15
2.1.2	Discrete implementation in OSIRIS	16
2.1.3	Accuracy constraints of the spectral initialization	17
2.2	Simulation Design and Physical Validation	18
2.2.1	Optimization of spatio-temporal drivers for wakefield excitation	19
2.2.2	Simulation domain size and wall injection for spatio-temporal pulses	23
2.3	Conclusions	27
2.4	Simulation setup details	28
3	Decaying Particles Acceleration	31
3.1	Surviving Particles Fraction in an Accelerator	32
3.2	Plasma Wakefield Acceleration with Subluminal Drivers	33
3.2.1	Muon Acceleration via Standard Laser Wakefields	34
3.2.2	Subluminal-driver plasma wakefield acceleration	35
3.2.3	Subluminal Driver with Different Velocities	38

3.3	Tailored plasma density profile for extended phase-lock	40
3.4	Model validation and versatility	41
3.4.1	Acceleration of Different Particles	41
3.4.2	Acceleration using Secondary Acceleration Buckets	42
3.4.3	Acceleration from Different Initial Velocities	44
3.4.4	Estimate of the end of phase-locked acceleration	45
3.5	Conclusions	46
4	Beam Quality and Experimental Feasibility	49
4.1	Beam Loading Effects	50
4.2	Transverse Dynamics and Beam Quality Preservation	52
4.3	Study on Transverse Trapping in Presence of a Tapered Density Profile	54
4.4	Experimental Feasibility of our Tapered Density Profile	56
4.5	Conclusions	58
5	Co-living of the Self-Modulation and Hosing Instabilities in AWAKE	61
5.1	Physical background: SMI and HI in AWAKE	62
5.1.1	Self-modulation instability (SMI)	63
5.1.2	Hosing instability (HI)	64
5.2	The Run 2b of AWAKE experiment	65
5.2.1	Electron bunch seeding of self-modulation and hosing	65
5.2.2	AWAKE Run 2B diagnostics used in this work	67
5.2.3	Experimental results of Hosing Instability Evolution	69
5.3	Simulations Using the Quasi-static Code QuickPIC	71
5.3.1	Simulation Setup	72
5.3.2	Centroid growth versus effective plasma length in simulation	73
5.3.3	Dependence on the seeding-beam misalignment	74
5.4	Conclusions and Future Directions	77
6	Conclusions	79
A	Machine-learning surrogates in PIC codes	81
A.1	Workflow and initial Compton Scattering proof of concept	81
A.2	Neural-network sampling of Bethe-Heitler pair production in PIC codes	82

LIST OF FIGURES

1.1	Schematic of a plasma wakefield behind a driver, showing accelerating buckets indexed by n_w and a witness beam injected into an accelerating phase. The wake propagates at the phase velocity v_ϕ , and the characteristic longitudinal scale is set by the plasma wavelength λ_p	5
1.2	Schematic of a single iteration of the particle-in-cell (PIC) algorithm. The index p labels individual macroparticles, while the indices i, j denote grid locations. The highlighted blocks correspond to the particle integration step and the electromagnetic field update, respectively.	8
2.1	Spectral design of ST pulses with controllable focal velocities. The desired spectral components (red curve) are defined by the intersection of the vacuum light cone (green surface) with the plane imposed by the spatio-temporal coupling (purple plane). (a) In the superluminal regime ($v_f = 1.5c$), the intersection yields an open hyperbolic trajectory, characteristic of X-waves with infinite spectral extent. (b) In the subluminal regime ($v_f = 0.5c$), the coupling plane cuts through the time-like region of the cone, confining the allowed spectral range to a closed elliptical loop.	15
2.2	Two examples of spatio-temporal pulses obtained with the discrete implementation in OSIRIS. In (a), we have an example of the shape of a superluminal pulse with $v_f = 1.5c$. In (b), we show an example of the shape of a subluminal pulse with $v_f = 0.5c$. The blue-red color-scale represents the electric field, and the black line corresponds to the pulse's full longitudinal envelope $f_{ }(z)$ in configuration space.	17
2.3	Comparison of the injected pulse electric field for three choices of the longitudinal and transverse repetition lengths used in the spectral reconstruction. (a) $k_0 T_{\text{rep},z} = 700$ and $k_0 T_{\text{rep},(x,y)} = 1000$, for which the reconstructed pulse is smooth and free of aliasing artefacts. (b) $k_0 T_{\text{rep},z} = 200$ and $k_0 T_{\text{rep},(x,y)} = 100$, where the longitudinal repetition length is too small, producing spurious replicas along the propagation axis. (c) $k_0 T_{\text{rep},z} = 700$ and $k_0 T_{\text{rep},(x,y)} = 200$, where the transverse repetition length is insufficient, and replicas enter the simulation domain, leading to distorted, unphysical structures.	19

- 2.4 Wakefield excitation for different choices of v_f and w_0 . Panels (a) and (b) show the case $v_f = 0.8c$: for (a) $k_0 w_0 = 10$ the driver is too long and overlaps with most of the first plasma period, whereas for (b) $k_0 w_0 = 7.1$ the driver length is better matched and a clearer wake structure is obtained. The same trend is more pronounced in panels (c) and (d) for $v_f = 0.6c$, comparing (c) $k_0 w_0 = 10$ with the tighter-focus case (d) $k_0 w_0 = 4.3$. The greyscale shows the electron density, and the blue-red colour scale shows the laser electric field. All cases use $\omega_0/\omega_p = 10$ and $a_0 = 2.2$ 21
- 2.5 Waterfall plots of the accelerating electric-field amplitude for different focal velocities: (a) $v_f = 0.8c$, (b) $v_f = 0.99c$, and (c) $v_f = 1.4c$ 22
- 2.6 Wakefield excitation for a fixed driver with $v_f = 0.96c$ and $k_0 w_0 = 17$, for different values of a_0 . The upper panels show the electron density response and the on-axis laser electric-field envelope, while the lower panels show the on-axis accelerating field. (a) $a_0 = 0.22$, (b) $a_0 = 0.44$, (c) $a_0 = 1.3$, and (d) $a_0 = 2.2$ 22
- 2.7 Schematic of a subluminal ST pulse showing the full laser envelope (black) and the localized high-intensity region (grey line; blue-red shading). Because the envelope propagates at c while the focus advances at $v_f < c$, the intensity peak slips backward within the envelope between (a) t_1 and (b) $t_2 > t_1$. Once the peak reaches the trailing edge of the envelope, the driver can no longer sustain an efficient wake. 24
- 2.8 Reduced-domain simulations of a subluminal ST driver with and without wall injection. (a) schematic of the injected field and boundary locations. (b,d) early-time snapshots in the co-moving coordinate ξ without and with wall injection, respectively. (c,e) late-time snapshots showing wake degradation without injection and sustained excitation with injection. (f) peak intensity in vacuum versus time, confirming that wall injection preserves the intended propagation-invariant behavior in the reduced domain. Greyscale shows the electron density, and the blue-red scale shows the laser electric field. 26
- 3.1 Time evolution of the laser-driven wakefield acceleration of a monoenergetic muon beam with initial velocity of $0.9c$. Panels (a–c) show snapshots at (a) $t = 32.03 \omega_p^{-1}$, (b) $t = 62.06 \omega_p^{-1}$, and (c) $116.12 \omega_p^{-1}$. The grayscale colormap represents the plasma density, in green we have the muon beam density, and the red line marks the longitudinal accelerating electric field on-axis. Panels (d–f) show the evolution of the beam's maximum longitudinal momentum p_1 as a function of propagation coordinate x_1 , in units of c/ω_p . A pattern of repeated energy peaks is observed in correspondence with the peak intensity of the wakefield electric field, with a maximum momentum of $p_{1,\max} \approx 2.21 m_\mu c$ reached in panel (e). For reference, a background density of $n_0 = 10^{18} \text{ cm}^{-3}$ corresponds to $1/\omega_p \simeq 17.7 \text{ fs}$ and $c/\omega_p \simeq 5.3 \mu\text{m}$ 36

- 3.2 Time evolution of wakefield acceleration of a monoenergetic muon beam with initial velocity of $0.9c$ using a spatio-temporal laser driver with a focal velocity of $v_f = 0.9c$. Panels (a–c) show snapshots at (a) $t = 41.97 \omega_p^{-1}$, (b) $t = 159.88 \omega_p^{-1}$, and (c) $261.80 \omega_p^{-1}$. The grayscale colormap represents the plasma density, in green we have the muon beam density, and the red line marks the longitudinal accelerating electric field. Panels (d–f) show the evolution of the beam's maximum longitudinal momentum p_1 as a function of propagation coordinate x_1 , in units of c/ω_p . A pattern of repeated energy peaks is observed in correspondence of the peak intensity of the wakefield electric field, with a maximum momentum of $p_{1,\max} \approx 2.3 m_\mu c$ reached in panel (e). For reference, a background density of $n_0 = 10^{18} \text{ cm}^{-3}$ corresponds to $1/\omega_p \simeq 17.7 \text{ fs}$ and $c/\omega_p \simeq 5.3 \mu\text{m}$ 38
- 3.3 Maximum energy gain for a beam of muons (initial velocity $v_0 = 0.9c$, corresponding to normalized momentum $p_1 = 2.06 m_\mu c$), accelerated by spatiotemporal laser pulses with different focal velocities (v_f). The plot compares the evolution of the longitudinal momentum (p_1) for three cases: $v_f = 0.91c$ (green), $v_f = 0.92c$ (pink), and $v_f = 0.93c$ (blue). The case with $v_f = 0.91c$ achieves the highest final momentum ($p_1^{\max} \approx 3.1 m_\mu c$), as this intermediate focal velocity balances initial trapping with delayed phase slippage, maximizing the effective acceleration length. For higher driver velocities ($v_f = 0.92c$ and $0.93c$), the particles dephase more rapidly and cannot stay in phase with the accelerating field, resulting in lower maximum energy gains. 39
- 3.4 (a) Tailored plasma density profiles $n(z)/n_0$ computed for different particles using the phase-locking model described in Section 3.3. Each profile maintains synchronism between the accelerating wake and particles of varying mass: pion mass m_π (solid), muon mass m_μ (dashed), and reduced muon masses $m_\mu/2$ (dash-dotted) and $m_\mu/4$ (dotted). (b) Maximum Lorentz factor $\gamma_{\max}$ achieved as a function of particle mass. The black curve shows theoretical predictions based on the tailored profiles, while red squares represent OSIRIS simulation results. All cases assume initial injection at $\gamma_i = 2.9$ (gray dashed line). The agreement confirms the predictive accuracy and versatility of the model. . . . 42
- 3.5 (a) Wakefield structure generated by a spatiotemporal laser pulse propagating in plasma with a focal velocity $v_f = 0.97c$. The grayscale colormap shows the plasma density response to the laser's propagation n/n_0 , while the blue-orange and red-blue colormaps correspond to the transverse field component $E_x - B_y$ and the longitudinal electric field E_L , respectively. The laser pulse enters from the right and excites a train of accelerating buckets, with the first seven labeled. The black line shows the on-axis accelerating field E_z . (b) Ideal density profile for the acceleration of a monoenergetic particle beam with mass $m_\mu/4$ for the first three accelerating buckets. 43

- 3.6 Maximum Lorentz factor $\gamma_{\max}$ reached as a function of the accelerating bucket used for the acceleration n_w . Theoretical predictions (black squares) are compared with OSIRIS simulations (red circles), showing excellent agreement. All cases correspond to initial injection at $\gamma_i = 2.29$ (gray dashed line). 43
- 3.7 Maximum velocity evolution v/c of three beams starting at $v_0/c = 0.5$, $v_0/c = 0.7$, and $v_0/c = 0.9$. Each beam is driven with focus velocities of $0.8c$, $0.9c$, and $0.97c$, respectively. The dashed lines represent the theoretical velocity evolution, and the solid markers are from OSIRIS simulations. 44
- 3.8 Estimate of the maximum acceleration time. The longitudinal separation $\Delta x_1(t) = x(t) - x_f(t)$ is computed by comparing the witness trajectory $x(t)$ (reconstructed via Eq.(3.15)) with the prescribed driver peak $x_f(t) = x_f(0) + v_f t$. The overtake time t_{over} (red marker) denotes the zero-crossing where the witness outruns the driver, marking the termination of the phase-locked accelerating stage. Calculations assume a reference density $n_0 = 4 \times 10^{18} \text{ cm}^{-3}$ and a laser frequency ratio $\omega_0/\omega_p = 20$ 45
- 4.1 Longitudinal electric field (E_z) lineouts for different beam loading regimes characterized by the parameter Λ , evaluated for a short witness beam ($\sigma_z \simeq \lambda_p/2$, $\sigma_r = 0.3 \mu\text{m}$) at initial velocity $0.9c$ in a plasma of density $n_0 = 4 \cdot 10^{18} \text{ cm}^{-3}$ ($k_p = 0.377 \mu\text{m}^{-1}$). The black line ($\Lambda = 1.3 \cdot 10^{-5}$, $n_b = 10^{-3}n_0$) represents the unloaded limit where the witness-beam charge is negligible. The red line ($\Lambda = 0.13$, $n_b = 10 n_0$) corresponds to a moderate beam-loading scenario where the wakefield is modified by the presence of the bunch. The blue line ($\Lambda = 1.3$, $n_b = 100 n_0$) illustrates the overloaded regime, where excessive beam charge significantly distorts the accelerating structure and degrades the gradient. . . . 50
- 4.2 Evolution of the normalized transverse emittance ($\bar{\epsilon}_n$) of the trapped beam core as a function of propagation distance z . The unloaded case ($\Lambda \approx 10^{-5}$, black) and the optimal beam loading case ($\Lambda = 0.13$, red) show comparable emittance preservation, demonstrating that the flattened accelerating fields do not degrade transverse quality, but instead slightly improve it. Conversely, the overloaded regime ($\Lambda = 1.3$, blue) exhibits rapid emittance growth driven by severe wakefield distortion. 53
- 4.3 Transverse stability analysis of different sub-relativistic particles injected into the tailored density profile. (a-d) Evolution of the theoretical plasma density $n(z)/n_0$ (log scale) required for longitudinal phase-locking for electrons, muons, pions, and kaons, respectively. (e-h) Computed transverse trajectories for the corresponding particle species. All particles are injected with a radial offset of $r_{b0} = 1.9 \mu\text{m}$ into the wakefield bubble with an initial radius of $r_{b0} = 2.0 \mu\text{m}$. The grey shaded regions represent the plasma background evolution ($r_b(z) = r_{b0} \cdot [n(z)/n_0]^{-1/2}$). Despite the proximity to the injection wake's border, the particles remain confined within the accelerating structure throughout the entire propagation. 55

- 4.4 Validation of the experimental feasibility using a piecewise-linear density profile. (a) Comparison between the exact theoretical profile ('ideal density (theory)', blue) and a surrogate profile constructed from linear segments ('approximated density', orange dashed). The bottom panel shows the relative density error. (b) Evolution of the maximum beam velocity $\beta_{\max}$ obtained from OSIRIS simulations employing the two density profiles shown in (a). The results using the approximated profile (orange) closely track the ideal simulation (blue) and the ODE prediction (black dashed). The negligible deviation $\Delta\beta$ confirms that the discretization of the density gradient does not impact the final energy gain. 57
- 4.5 Experimental feasibility using a two-jet gas profile. (a) Comparison between the ideal tapered density $n_{\text{ODE}}(z)$ (blue) and an implementable profile (orange, dashed) constructed from two Gaussian gas jets (parameters in Table 4.2). The top panel shows $n(z)$ on a logarithmic scale; the bottom panel shows the relative percentage error, defined as $100 \times (n_{\text{jets}} - n_{\text{ODE}})/n_{\text{ODE}}$. (b) Comparison of the energy gain between the theoretical ODE prediction (dashed black), the PIC simulation using the ideal density (blue), and the PIC simulation using the gas jet approximation (orange). The two-jet surrogate closely reproduces the required profile and maintains the expected acceleration performance. 58
- 5.1 Illustrative comparison between a self-modulation-dominated case (a) and a hosing-dominated case (b) for the propagation of a long particle beam in plasma. The grey scale shows the normalized proton-bunch density n/n_0 (arb. units), while the red-blue scale shows the transverse wakefield component in the plane of each panel (W_x in (a) and W_y in (b); red positive, blue negative). The green curve indicates the slice-by-slice transverse centroid. In (a), the density and transverse-field structure remain symmetric with respect to the beam axis, and the centroid stays close to zero. In (b), a clear centroid displacement develops together with a transverse-field and density pattern that breaks axial symmetry, characteristic of hosing growth. 63
- 5.2 Experimental setup of seeding self-modulation and hosing instability. Sketch of the set-up, not to scale. 66
- 5.3 AWAKE Run 2b experimental setup and diagnostics. 68

- 5.4 Plunger-scan visualisation of the seeded transverse centroid dynamics in AWAKE Run 2b. Each panel shows the mean streak-camera image of the proton bunch at the plasma exit for a given plasma length (labels in metres), after shot-to-shot alignment using the timing marker (not shown) and restricting to the proton-beam region of interest. The vertical axis ζ is the longitudinal coordinate along the bunch, obtained from the time resolved sweep of the streak camera, and the horizontal axis y is the transverse coordinate at the screen in the plane of the imposed seed misalignment. The red markers overlay the transverse centroid extracted slice-by-slice along the time-resolved axis using a centre-of-mass estimator within a length-dependent window. The progressive increase of the centroid excursion with plasma length provides a direct experimental handle on the growth of the hosing-like transverse displacement when the instability is seeded by an intentionally misaligned electron bunch. 70
- 5.5 Centroid displacement extracted from the plunger scan streak camera data, in the seeded misalignment configuration, as a function of the effective plasma length L . Black circles show the mean of the slice resolved centroid envelope $|y_c(\zeta)|$ averaged over the full longitudinal extent of the proton bunch. Grey squares show the root mean square of the slice centroid $y_c(\zeta)$ over the same range, which is less sensitive to the sign of the displacement. The red curve is a one parameter fit of the convective hosing envelope of Eq. (5.4), $|y_c| \propto \exp(N_h)/\sqrt{N_h}$ with $N_h = AL^{2/3}$, to the mean envelope points. 71
- 5.6 Simulated proton-bunch density projections in the (y, z) plane for a seeded-misalignment configuration at different effective plasma lengths (panel labels). Greyscale shows the particle-density histogram, and the red curve is the reconstructed transverse centroid $y_c(z)$ obtained from a centre-of-mass estimate in thin longitudinal bins. Top row: distributions at the end of the plasma (i.e. at the selected effective length). Bottom row: the same particles ballistically propagated through vacuum until the streak camera position. 74
- 5.7 Centroid-displacement metric extracted from the simulation data plunger-scan (seeded-misalignment configuration) as a function of effective plasma length L . Black circles are evaluated at the plasma exit, grey squares after free streaming to the streak camera position. Each point is the mean of the slice resolved centroid envelope $|y_c(\zeta)|$ along the bunch, and the bar is its standard deviation along ζ . The red and blue dashed curves are fits of the convective hosing envelope of Eq. (5.4) to the plasma exit and streak camera points respectively, with fitted growth coefficients A reported in the legend. 75

5.8	QuickPIC simulation of the transverse centroid evolution in a seeded-hosing configuration, where the hosing seed is provided by injecting an electron bunch with an initial transverse offset Δy relative to the proton-bunch axis. The plot shows the proton-bunch centroid $y_c(z)$ as a function of propagation coordinate z , for the set of offsets indicated in the legend. All coordinates are expressed in plasma units (c/ω_p), and all parameters other than the seed offset are kept fixed across the scan.	75
5.9	QuickPIC simulations of the seeded-hosing response to varying electron-bunch transverse offset Δy relative to the proton-bunch axis. (a) Transverse centroid evolution $y_c(z)$ of the proton bunch as a function of propagation coordinate z , for the set of offsets indicated in the legend. The maximum centroid misalignment reached during propagation is extracted for each offset value. (b) Black curve, transverse lineout of the wakefield component $W_y \equiv E_y - cB_x$ along the y axis at the longitudinal position of the proton bunch axis of symmetry and at propagation coordinate $z = 3$ m (left axis), showing the dipole structure of the seed driven wakefield in the plane of the imposed misalignment. Coloured stars, maximum centroid displacement y_c from panel (a) for the offsets indicated in the legend, plotted at the corresponding Δy position along the horizontal axis (right axis). All coordinates are expressed in plasma units (c/ω_p), and all parameters other than the seed offset are kept fixed across the scan.	76
A.1	ML-PIC workflow	82
A.2	Bethe-Heitler NN sampling validation	83

LIST OF TABLES

3.1	Estimated survival fraction S (and corresponding loss fraction $1 - S$) for μ^- , π^- , and K^- accelerated from the initial to the final Lorentz factor (γ_i, γ_f) in a plasma accelerator and in a conventional RF linac. The plasma and RF cases assume constant accelerating gradients of 0.2 GV/cm and 20 MV/m , respectively.	33
3.2	Simulation set-up for muon acceleration using standard laser-driven wakefield acceleration.	35
3.3	Simulation set-up for muon acceleration using standard spatio-temporal pulses driven wakefield acceleration.	37
4.1	Trapping and accelerated charge fractions of the witness beam for the three beam-loading regimes of Fig. 4.1. The beam-loading parameter is defined as $\Lambda = (n_b/n_0)(k_p\sigma_r)^2$. Q_0 denotes the initial bunch charge and Q_{trap} the absolute trapped charge. The Trapping column reports the charge fraction confined within the accelerating bucket; the $\geq 0.97c$, $\geq 0.98c$, and $\geq 0.99c$ columns report the charge fraction reaching at least the indicated longitudinal velocity. All fractions are normalized to Q_0	51
4.2	Parameters of the two Gaussian gas jets used to approximate the ideal tapered density profile. Lengths are derived using a reference density $n_0 = 5.0 \times 10^{17} \text{ cm}^{-3}$ (corresponding to a plasma skin depth $c/\omega_p \simeq 7.515 \text{ }\mu\text{m}$). Each jet is described by the function $n_j(z) = b + A \exp\left[-4 \ln 2 \left((z - z_0)/\text{FWHM}\right)^2\right]$. The last column reports the normalized on-axis peak density.	58
5.1	Experimental Parameters of AWAKE Run 2B for seeded hosing studies.	67
5.2	Baseline parameters for the AWAKE SMI-HI simulations performed with QuickPIC.	72

To the ones who taught me patience,
and to those who helped me keep it.

CHAPTER 1

INTRODUCTION

Recent progress in plasma-based accelerators is driven by their potential to deliver compact, sustainable, and high-performance particle or radiation sources. Although conventional radio-frequency (RF) technology has enabled significant scientific advancements, its accelerating gradients are limited to approximately 100 MV/m due to material breakdown [1]. Consequently, achieving higher energies requires infrastructure at the kilometer scale, imposing considerable spatial and energy demands [2]. In contrast, plasma accelerators employ an already ionized medium, which is not subject to material breakdown, and can sustain electric fields several orders of magnitude stronger. When excited by intense lasers or relativistic particle bunches, plasma waves can support gradients exceeding tens of GV/m [3–5]. However, translating these extreme fields into stable acceleration over long distances remains a major scientific challenge.

In laser wakefield acceleration (LWFA), an ultrashort ($\tau \lesssim 100$ fs), high-intensity ($I \gtrsim 10^{18}$ W/cm²) laser pulse drives the wake through the ponderomotive expulsion of plasma electrons [3]. This approach has successfully demonstrated multi-GeV electron energy gains in centimeter-scale plasmas [6, 7]. However, the maximum attainable energy in a single stage is intrinsically limited by two factors: depletion of the laser pulse energy and accelerated particle beam dephasing, where the accelerated particles outrun the wake. Addressing these limits is critical for future applications and motivates the development of advanced schemes such as plasma waveguides [8] or, as explored in this work, the use of shape-invariant structured light pulses [9, 10].

Alternatively, in beam-driven plasma wakefield acceleration (PWFA), a relativistic charged particle bunch excites the wake through its intense space-charge fields [4]. As in the LWFA case, the process is limited by the finite energy stored in the driver, which is depleted as it transfers energy to the witness. A particularly compelling solution is to use high-energy proton bunches as drivers. Since stored proton bunch energies can be orders of magnitude larger than those available in single laser pulses or electron bunches, proton-driven PWFA offers a path toward single-stage, multi-meter acceleration lengths [11]. The central challenge is that the available proton bunches (such as the SPS beam at CERN) are typically much longer than the plasma wavelength. Consequently, the wake must be generated through the Self-

Modulation Instability (SMI), a process that micro-bunches the driver at λ_p , but is inherently coupled to other destructive transverse instabilities such as hosing.

This thesis addresses these complementary challenges by exploring novel driver-control mechanisms designed to maximize the length of a single-stage acceleration. Specifically, it investigates the use of spatio-temporally structured laser pulses to overcome dephasing limits and analyzes the competitive dynamics between self-modulation and hosing to ensure stable proton-driven acceleration. Together, these studies aim to extend the reach of plasma-based schemes toward the length scales required for high-energy physics applications.

1.1 STATE OF THE ART IN PLASMA ACCELERATION

Over the last two decades, plasma accelerators have progressed from proof-of-principle demonstrations to experiments targeting application-prone beam quality. In beam-driven PWFA, landmark energy gains were demonstrated at SLAC with meter-scale plasmas [12], and high-efficiency transfer to a distinct trailing bunch was later achieved through strong beam loading [13]. More recently, experiments have shown that careful control of the current profile can simultaneously preserve low energy spread and achieve high energy-transfer efficiency [14], while dedicated studies have addressed constraints on repetition rate through direct measurements of plasma recovery dynamics [15]. A key beam-quality milestone is the demonstration of emittance preservation in a high-gradient, high-efficiency PWFA, establishing that plasma stages can accelerate particle beams without degrading their transverse quality [16].

In the LWFA domain, multi-GeV electron beams have been produced using guided, high-power laser pulses in plasma channels [6, 7]. Although LWFAs have historically faced challenges related to shot-to-shot stability and longitudinal phase-space control, recent results increasingly emphasize reproducibility, spectral control, and transport into conventional beam-lines. These advances are also supported by facility-scale initiatives designed to deliver user-ready plasma-accelerated beams [17].

A particularly strong indicator of maturity is the growing set of demonstrations connecting plasma accelerators to photon-science applications. Compact plasma-based beams have driven free-electron lasing in different configurations and wavelength ranges, including beam-driven PWFA-based FEL operation [18] and LWFA-based lasing in the extreme ultraviolet [19], as well as seeded FEL operation enabled by controlled synchronization and improved longitudinal coherence [20]. Very recently, FACET-II experiments have reported simultaneous energy and brightness enhancement in a meter-scale PWFA, producing injected bunches exceeding 20 GeV with sub-percent energy spread and mm-mrad-class normalized emittance, pointing toward plasma stages that act not only as accelerators but also as beam-quality transformers [21]. In parallel, large collaborative efforts such as EuPRAXIA are defining a roadmap toward dedicated European plasma-accelerator infrastructures delivering multi-GeV, high-brightness electron beams for applications [17].

1.2 ANATOMY OF A PLASMA WAKEFIELD ACCELERATOR

A plasma wakefield accelerator consists of three essential elements: the *plasma* medium that supports the wave, a *driver* that excites a plasma wave within it, and a *witness beam* that is injected into a suitable phase of that wave and gains energy from its longitudinal electric field [3–5]. The driver can be either an intense laser pulse (LWFA) or a relativistic charged-particle bunch (PWFA). In both cases, the driver perturbs the plasma electron density, launching a co-moving oscillation whose electromagnetic fields can sustain accelerating gradients far beyond those achievable with conventional radio-frequency structures.




FIGURE 1.1: Schematic of a plasma wakefield behind a driver, showing accelerating buckets indexed by n_w and a witness beam injected into an accelerating phase. The wake propagates at the phase velocity v_ϕ , and the characteristic longitudinal scale is set by the plasma wavelength λ_p .

For a plasma of electron density n_0 , the plasma frequency and plasma wavelength are

$$\omega_p = \sqrt{\frac{n_0 e^2}{\epsilon_0 m_e}}, \quad \lambda_p = \frac{2\pi c}{\omega_p}, \quad (1.1)$$

which set the natural temporal and longitudinal spatial periodicity of the wake [5]. Throughout this thesis, the longitudinal direction z denotes the propagation axis of the driver and wake (see Fig. 1.1), and a cylindrical geometry is assumed with radial coordinate $r = \sqrt{x^2 + y^2}$ measuring the transverse distance from the propagation axis. The on-axis longitudinal electric field $E_z(z, t)$ varies periodically along the wake, producing alternating accelerating and decelerating phase regions for a given charge species within each wake bucket.

Behind the driver, the wake is naturally segmented into successive accelerating phases, commonly referred to as *wake buckets*. Figure 1.1 illustrates this idea schematically: the first accelerating bucket immediately behind the driver is labeled $n_w = 1$, while *secondary accelerating buckets* further back correspond to $n_w = 2, 3, \dots$. The longitudinal separation between adjacent buckets is set by the plasma wavelength λ_p . The ability to inject and accelerate a witness in either the first or a secondary bucket is an important degree of freedom in both simulations and experiments, and will also be employed in this PhD thesis.

The driver travels in the $+z$ direction and excites a wake that propagates at phase velocity v_ϕ (Fig. 1.1). In conventional LWFA, the wake phase velocity is typically close to the laser group velocity v_g in plasma [5]. In PWFA, v_ϕ is generally close to the driver-bunch velocity. In this thesis, we also consider *spatio-temporal (ST) laser drivers* for which the velocity of the intensity peak (the “focal velocity”) can be prescribed (Chapter 2). In that case, the wake phase velocity can be engineered to be subluminal and to vary along the propagation path, enabling synchronism with initially slow, heavy particles.

In addition to the longitudinal accelerating field, plasma wakes feature strong transverse fields that can focus or defocus the witness. A common quantity is the transverse (Lorentz) force $F_\perp = q(E_r - v_b B_\theta)$, which determines whether a given phase provides net focusing for a relativistic witness [5]. In the nonlinear “blowout” (or “bubble”) regime, the driver expels plasma electrons, leaving an ion cavity that provides a nearly linear focusing force for negatively charged relativistic witnesses, while E_z can remain approximately uniform over a significant portion of the bucket [22]. Although the present simplified schematic (Fig. 1.1) emphasizes the longitudinal bucket structure, the transverse focusing properties are equally central for stable trapping and emittance preservation.

To avoid ambiguity, we will use the following terms consistently: (i) *driver*: the laser pulse or particle bunch that excites the wake; (ii) *wakefield*: the co-moving plasma wave and associated fields; (iii) *witness beam*: the accelerated bunch injected into the wake; (iv) *accelerating bucket/wake* n_w : the n_w -th accelerating phase behind the driver (with $n_w = 1$ the first bucket); (v) *phase velocity* v_ϕ : the propagation velocity of the wake structure; (vi) *plasma wavelength* λ_p : sets the longitudinal scale of the buckets.

1.3 MOTIVATION AND SCOPE OF THE THESIS

A central limitation of plasma wakefield accelerators is that the useful acceleration distance is not set by the plasma itself, but by the evolution of the driver and the synchronization between the witness beam and the wake. In laser wakefield acceleration (LWFA), energy is continuously transferred from the laser to the plasma and the accelerated beam, leading to pump depletion. Similarly, relativistic electrons may outrun the accelerating phase of the wake (dephasing), further capping the maximum energy gain sustainable in a single stage [5]. Analogously, in beam-driven PWFA, the driver energy is depleted as it excites the wake, determining the maximum transfer efficiency [13, 23].

The conventional route to overcoming these limits is staging: extracting the beam from one plasma module and injecting it into a subsequent stage powered by a fresh driver [5]. While conceptually straightforward, staging imposes stringent requirements on beam transport, inter-stage synchronization, and alignment. Furthermore, capturing and refocusing the beam between stages without degrading its emittance remains a strong technological barrier [24]. These challenges motivate the search for advanced schemes that can extend the effective acceleration distance within a single stage, thereby relaxing or eliminating the need for complex staging structures.

One of the most compelling frontiers for such advanced acceleration schemes is the produc-

tion of a high-quality decaying particle beam. A highly relevant application is, for example, the acceleration of muons [25]. Muon colliders promise to provide clean lepton-lepton interactions at multi-TeV energies, enabling precision studies of the Higgs boson and potential discoveries of new physics [26–29]. Additionally, neutrino factories based on muon storage rings could revolutionize the study of neutrino oscillations and CP violation [30, 31]. However, the efficient acceleration of muon beams is challenging due to the particle’s finite lifetime ($\tau_\mu \approx 2.2 \mu\text{s}$ at rest).

This finite lifetime necessitates rapid capture and acceleration techniques to boost muons to relativistic velocities before they decay. Standard LWFA schemes, which are optimized for ultra-relativistic electrons ($v \approx c$), are not well suited for this task because the wake phase velocity is too high for initially slow muons. To address this, this thesis develops a theoretical scheme employing structured laser pulse drivers with focal velocity control. By manipulating the laser’s focal velocity, we can generate a wakefield that travels at a subluminal velocity, synchronized with the accelerating muons, thereby preventing dephasing and enabling efficient capture and acceleration over extended distances. The proposed acceleration method is designed to be versatile and applicable to particles with different masses, initial velocities, and laser powers; for example, it can also accelerate pions and kaons (with respective lifetimes of $\tau_\pi = 26 \text{ ns}$ and $\tau_k = 12 \text{ ns}$), which are otherwise impossible to accelerate with standard RF cavities. Complementing the laser-driven approach, the second part of this thesis is motivated by the possibility of employing high-energy proton bunches to overcome the pump-depletion limit. Unlike lasers or electron beams, proton drivers can carry kilojoules of stored energy, offering a direct route to single-stage acceleration over multi-meter distances. However, this energy comes in the form of long bunches that must undergo the Self-Modulation Instability (SMI) to resonantly drive a wake. In this regime, the challenge shifts from maintaining driver velocity to maintaining driver stability. Specifically, the same forces that modulate the beam can also amplify transverse centroid oscillations, the Hosing Instability, potentially destroying the beam before acceleration can occur. Therefore, a study on the control of this instability is necessary if long acceleration stages are to be achieved.

These two complementary research lines converge on a unified objective: identifying physically grounded routes to stable, single-stage plasma acceleration. Whether through the spatio-temporal engineering of laser drivers to accelerate short-lived muons, or the mitigation of transverse instabilities in high-energy proton beams, this thesis seeks to extend the interaction length of plasma accelerators toward the scales required for next-generation particle physics.

1.4 METHODS

1.4.1 PARTICLE-IN-CELL SIMULATIONS WITH OSIRIS

The majority of the simulation results presented in this thesis are obtained using the OSIRIS code, a fully relativistic, electromagnetic particle-in-cell (PIC) framework widely used for the study of laser-plasma and beam-plasma interactions [32]. OSIRIS was originally developed to address problems in plasma-based acceleration, relativistic laser-matter interaction, and high-

energy-density physics, and has since evolved into a versatile platform for kinetic plasma simulations in one, two, and three dimensions.

The PIC method provides a first-principles numerical solution of the Klimontovich–Maxwell system by representing the plasma as an ensemble of finite-sized macroparticles and solving Maxwell’s equations on a spatial grid. Macroparticles are advanced in time according to the Lorentz force, while their charge and current densities are deposited onto the grid to update the electromagnetic fields. This self-consistent cycle captures both collective plasma dynamics and kinetic effects, making PIC particularly well suited for highly nonlinear regimes such as plasma wakefield acceleration, beam instabilities, and relativistic laser propagation.




FIGURE 1.2: Schematic of a single iteration of the particle-in-cell (PIC) algorithm. The index p labels individual macroparticles, while the indices i, j denote grid locations. The highlighted blocks correspond to the particle integration step and the electromagnetic field update, respectively.

Figure 1.2 illustrates the standard PIC cycle. At each time step, the electromagnetic fields defined on the spatial grid are interpolated to the particle positions using a weighting scheme. The particles are then advanced according to the relativistic Lorentz force, after which their charge and current densities are deposited back onto the grid. Maxwell’s equations are finally integrated on the grid to update the electromagnetic fields, completing a self-consistent iteration of the algorithm.

OSIRIS employs a finite-difference time-domain (FDTD) solver for Maxwell’s equations and supports both Cartesian and cylindrical geometries, including a quasi-cylindrical spectral decomposition in azimuthal modes [33, 34]. This formulation allows for efficient simulation of effectively three-dimensional phenomena while retaining much of the physics of full 3D sim-

ulations at reduced computational cost. The code includes advanced features such as moving simulation windows, absorbing boundary conditions, field and particle diagnostics, and support for external particle and field injection.

Unlike direct particle-particle methods such as molecular dynamics, whose cost scales as $\mathcal{O}(N^2)$ with the number of particles N , the PIC approach scales as $\mathcal{O}(N)$ by using a grid to mediate particle interactions through the electromagnetic fields. In practice, the computational cost is dominated by the number of macroparticles rather than the number of grid cells, and the continued growth of high-performance computing has made increasingly realistic simulations feasible, although the cost can still become prohibitive for long propagation distances or large transverse domains. The ability to capture the complete spatio-temporal evolution of both fields and particles makes PIC indispensable for studying strongly nonlinear plasma processes, where reduced models or envelope approximations may fail. In this work, we employ the OSIRIS code to investigate wakefield excitation, beam-driven instabilities, and the interaction of structured laser pulses with plasma, where full electromagnetic and kinetic fidelity is required.

1.4.2 QUASI-STATIC SIMULATIONS WITH QUICKPIC

In addition to full PIC simulations, a subset of the simulations presented in this thesis was performed using the quasi-static QUICKPIC code [35, 36], which is specifically designed for the efficient modeling of beam-driven plasma wakefield acceleration involving highly relativistic particle beams.

QUICKPIC is based on the quasi-static approximation (QSA), which exploits the strong separation of timescales between the ultra-relativistic driver (electron or proton beam) and the plasma response. In this framework, the driver is assumed to evolve slowly compared to the plasma oscillation period, allowing the electromagnetic fields and plasma particles to be solved in a co-moving frame as a sequence of steady-state problems. The fast plasma dynamics is computed as a function of the longitudinal coordinate $\xi = z - ct$, while the slow evolution of the driver is advanced over larger propagation steps.

This approximation reduces the original 3D time-dependent problem to a set of 2D transverse-field solutions, resulting in orders-of-magnitude reductions in computational cost compared to full PIC. As a consequence, QUICKPIC enables simulations of meter-scale plasma stages and realistic experimental parameters that would be challenging with explicit PIC methods.

The main limitation of the quasi-static approach is that it cannot capture processes that involve rapid temporal evolution of the driver or strong back-action on timescales comparable to the plasma period. Effects such as laser evolution, self-injection, or strong beam loading with rapid phase-space distortion fall outside the strict validity of the approximation. Nevertheless, for highly relativistic particle drivers with slowly varying envelopes, QUICKPIC provides an accurate and computationally efficient description of wakefield formation, beam focusing, and energy gain.

In this thesis, QUICKPIC is primarily used for long-distance propagation studies in the context of the AWAKE experiment presented in Chapter 5.

1.4.3 THESIS OUTLINE

This thesis is structured around two complementary strategies to extend the single-stage interaction length of plasma accelerators: the use of velocity-controlled laser drivers to overcome dephasing, and the use of high-energy proton drivers to overcome pump depletion. The structure is as follows:

Chapter 2: Spatio-Temporal Control of Laser Drivers

This chapter establishes the theoretical and numerical foundations for the first part of the thesis. It introduces the physics of spatio-temporally (ST) structured laser pulses, specifically focusing on the "flying focus" technique to decouple the peak intensity velocity from the group velocity. We present the simulation framework used throughout this work and detail the numerical construction of these complex drivers. Furthermore, the chapter provides a systematic parameter study, offering design guidelines for selecting optimal pulse configurations to maximize wake excitation and control the effective phase velocity of the generated wakes.

Chapter 3: Acceleration of Short-Lived Particles

Building on the tools developed in Chapter 2, this chapter applies velocity control to a critical challenge: the acceleration of short-lived, sub-relativistic particles such as muons, pions, and kaons. We demonstrate that standard wakefield configurations are insufficient due to severe phase slippage and introduce a novel scheme combining subluminal ST drivers with tailored plasma density profiles. The method is first introduced theoretically and then tested in self-consistent simulations. This approach allows for sustained phase-locking, enabling the efficient capture and acceleration of these particles before they decay. We also show that the method is versatile for different particle masses, initial velocities, and the acceleration wake used.

Chapter 4: Beam Quality and Stability Analysis

Building on the acceleration method presented in Chapter 3, Chapter 4 performs a rigorous analysis of the resulting beam quality. Here, we investigate the multidimensional dynamics of the acceleration process, focusing on trapping efficiency, emittance preservation, and energy spread. The chapter assesses the trade-offs between capture efficiency and beam quality, providing a comprehensive study of the transverse dynamics of the accelerated beam. This Chapter also includes a discussion of the experimental feasibility of the acceleration beam, focusing in particular on the peculiar density profile required by the theory.

Chapter 5: Proton-Driven Wakefield Acceleration

The focus shifts in this chapter to the second route for length extension: exploiting the kilojoule-level energy stored in proton bunches to drive wakefields over meter scales. Centered on the physics of the AWAKE experiment at CERN, this chapter examines its competition with the parasitic Hosing Instability and the Self-Modulation Instability. Specifically, we analyze the use of controlled electron-beam misalignment as a seed for these instabilities in the context of AWAKE Run 2B. The study connects Particle-in-Cell (PIC) simulations directly to experimentally accessible observables with the final goal of improving stability in proton-driven acceleration.

Chapter 6: Conclusion and Outlook

Finally, this chapter summarizes the main findings of the thesis. A final discussion of the possible impact of these results is also included.

1.5 LIST OF CONTRIBUTIONS

The work developed during this PhD has resulted in the following publications.

- [1] **C. Badiali**, R. Almeida, B. Malaca, R. Fonseca, T. Silva, and J. Vieira. “Tailored laser wakefield acceleration for decaying particles”. *Physical Review Research* (Mar. 2026). DOI: 10.1103/1pp8-qx35.
- [2] **C. Badiali**, R. Almeida, T. Silva, and J. Vieira. “Simulation Design for Velocity-Controlled Spatio-Temporal Drivers in Laser Wakefield Acceleration”. 2026. Submitted to *Plasma Physics and Controlled Fusion*. ARXIV: 2603.28473 [physics.plasm-ph].
- [3] **C. Badiali**, P. J. Bilbao, F. Cruz, and L. O. Silva. “Machine-learning-based models in particle-in-cell codes for advanced physics extensions”. *Journal of Plasma Physics* **88.6** (2022), p. 895880602. DOI: 10.1017/S0022377822001180.
- [4] **C. Badiali**, P. J. Bilbao, and F. Cruz. “A machine learning based approach to Compton scattering models for particle-in-cell codes”. In: *47th EPS Conference on Plasma Physics*. 2021, pp. 377–380. [proceedings].
- [5] Ó. Amaro, **C. Badiali**, and B. Martinez. “Neural network sampling of Bethe-Heitler process in particle-in-cell codes”. *Journal of Computational Physics* **553** (2026), p. 114707. DOI: 10.1016/j.jcp.2026.114707.
- [6] N. Z. van Gils *et al.* and the **AWAKE Collaboration**. “Seeding of Self-Modulation using Truncated Seed Bunches as a Path to High Gradient Acceleration”. 2026. ARXIV: 2603.29472 [physics.acc-ph].
- [7] D. Alesini, M. P. Anania, F. Anelli, . . . , **C. Badiali**, *et al.* *EuPRAXIA@SPARC_LAB Technical Design Report*. Tech. Rep. INFN–2026-01-LNF, INFN Laboratori Nazionali di Frascati, Feb. 2026. DOI: 10.15161/oar.it/k57qz-5qk09.
- [8] A. Clairembaud *et al.* and the **AWAKE Collaboration**. “Measurement of the Saturation Length of the Self-Modulation Instability”. 2026. ARXIV: 2602.16624 [physics.plasm-ph].
- [9] F. Zimmermann, S. Moebius, S. Kraus, G. Franchetti, **C. Badiali**, A. Pukhov, Y. Bellouard, Y. Yang, R. Dadashi Motlagh, D. Schulte, *et al.* “Advanced accelerator concepts for dark sector searches and fast muon acceleration”. *JACoW IPAC 2024, MOPR17* (2024). DOI: 10.18429/JACoW-IPAC2024-MOPR17.

The work developed during this PhD has been presented as invited talks at the following international conferences and workshops.

- [T1] “Acceleration of Decaying Particles in Laser-Plasma Accelerators Using Structured Light”. 51st EPS Conference on Plasma Physics (EPS 2025), Vilnius, Lithuania, July 2025.
- [T2] “Controlled Acceleration of Unstable Particles in Laser-Plasma Wakefields Using Structured Light”. Laser-Plasma Accelerators Workshop 2025 (LPAW 2025), Ischia Island, Naples, Italy, April 2025.
- [T3] “Simulations of a plasma wakefield accelerator for pions and muons”. I.FAST Workshop on GHz Rate and Rapid Muon Acceleration for Particle Physics, Bern, Switzerland, December 2023.

CHAPTER 2

SIMULATION DESIGN FOR VELOCITY CONTROLLED DRIVERS

Following the physical motivation introduced in Chapter 1, this chapter focuses on the numerical and modeling aspects required to study velocity-controlled spatio-temporal drivers in plasma wakefield acceleration. Laser-plasma accelerators have emerged as a promising route towards compact sources of relativistic electron and ion beams [11, 37–39] and secondary radiation [40–43], but their performance is often limited by dephasing and pump depletion [44, 45].

A central challenge is therefore to control how the laser intensity is distributed in space and time as it propagates through plasma. Spatio-temporal (ST) pulse shaping, and in particular the *ideal flying focus* concept, addresses this challenge by engineering laser pulses whose intensity peak travels at a prescribed velocity that is decoupled from the group velocity of the pulse envelope [10, 46–48]. This control is achieved through chromatic focusing combined with a tailored spectral phase, allowing the focal spot to propagate with a chosen focal velocity v_f (including subluminal and superluminal values) while maintaining an approximately constant peak intensity over distances that can greatly exceed the Rayleigh range of a conventional focus. This added control has triggered strong interest in ST pulses across nonlinear optics and plasma physics [49, 50].

Flying-focus and related ST structures have been proposed and, in some cases, demonstrated for driving ionization waves with prescribed velocities [51], forming extended plasma channels [52, 53], enabling dephasing-free laser wakefield acceleration, and tailoring wakefield structures for high-efficiency electron acceleration [54–57]. More recently, variants such as discrete and transverse flying-focus schemes have been introduced to enhance energy gain per stage or to accelerate ions, and structured beams have been used to directly drive wakefields in plasmas [58–60]. Overall, ST pulses with a prescribed focal velocity provide a flexible way to match the driver to the plasma response, rather than requiring the plasma to adapt to the constraints of a conventional laser pulse. We focus on the subluminal regime because it is directly motivated by our studies of phase-locking wakefields for initially subrelativistic, short-lived particles using $v_f < c$ drivers, which will be presented in Chapters 3 and 4. Nev-

ertheless, several of the considerations discussed here, including the spectral construction and practical simulation-design guidelines, are also relevant to other operating regimes, such as $v_f \simeq c$ and $v_f > c$ schemes aimed at mitigating dephasing in laser-wakefield acceleration.

This Chapter addresses the physical and numerical requirements for modeling subluminal flying-focus-type ST pulses in particle-in-cell simulations, with emphasis on practical use within OSIRIS [32]. Building on the Maxwell-consistent spectral decomposition of [61], we organize the Chapter around three main objectives: (i) presenting the spectral construction and its discrete k -space representation used for numerical initialization; (ii) quantifying how the ST pulse parameters determine the conditions for resonant wakefield excitation; and (iii) characterizing the geometric and numerical constraints that make these drivers computationally demanding, including focus-envelope slippage, transverse expansion, and the replica fields introduced by spectral discretizations.

The Chapter is organized as follows. Section 2.1 introduces the Maxwell-consistent spectral framework used to model velocity-controlled ST drivers and its numerical realization in OSIRIS. Section 2.1.1 defines flying-focus pulses in Fourier space and distinguishes the subluminal and superluminal regimes through the geometry of the light-cone intersection. Section 2.1.2 describes how the continuous spectral weight is mapped onto a discrete set of modes for field initialization in PIC codes. Section 2.1.3 analyzes the artifacts introduced by discrete k -space sampling and derives practical criteria to mitigate them.

Section 2.2 focuses on wakefield-relevant parameter choices and on the computational constraints that arise for subluminal drivers. Section 2.2.1 discusses how the spatio-temporal geometry constrains the effective driver length and transverse scale, and provides scaling guidelines for near-resonant wakefield excitation. Section 2.2.2 discusses the domain-size requirements imposed by envelope-focus slippage and the large angular content of velocity-controlled ST pulses, derives practical box-sizing estimates, and shows how wall injection enables long-distance simulations in reduced transverse domains. Finally, Sec. 2.3 summarizes the main conclusions and practical recommendations for the reader.

2.1 SPECTRAL CONSTRUCTION AND NUMERICAL IMPLEMENTATION OF SPATIO-TEMPORAL PULSES

Particle-in-cell codes commonly rely on paraxial beam models and the slowly varying envelope approximation to initialize laser pulses. Although these approaches are adequate for narrowband Gaussian drivers, they are not suitable for velocity-controlled space-time (ST) pulses. The prescribed focal trajectory of ST pulses requires a broadband spectral content, large-angle propagation components, and a fixed correlation between spatial and temporal frequencies. To address this limitation, we adopt a generalized, Maxwell-consistent initialization in which the electromagnetic field is expressed as a superposition of vacuum solutions of Maxwell's equations, thereby enabling accurate modeling of ST pulses. This section presents the mathematical formulation of this approach (Sec. 2.1.1), summarizes its implementation in OSIRIS (Sec. 2.1.2), and discusses one of the issues arising from this method together with




FIGURE 2.1: Spectral design of ST pulses with controllable focal velocities. The desired spectral components (red curve) are defined by the intersection of the vacuum light cone (green surface) with the plane imposed by the spatio-temporal coupling (purple plane). (a) In the superluminal regime ($v_f = 1.5c$), the intersection yields an open hyperbolic trajectory, characteristic of X-waves with infinite spectral extent. (b) In the subluminal regime ($v_f = 0.5c$), the coupling plane cuts through the time-like region of the cone, confining the allowed spectral range to a closed elliptical loop.

a strategy to circumvent it, namely the artificial periodic repetition of the fields induced by representing a continuous spectrum on a discrete k -space grid (Sec. 2.1.3).

2.1.1 SPECTRAL DESIGN OF SPATIO-TEMPORAL PULSES

In vacuum, any electromagnetic wave packet can be represented as a superposition of basis solutions that satisfy the vacuum dispersion relation

$$\omega(\mathbf{k}) = c|\mathbf{k}|. \quad (2.1)$$

where c is the speed of light in vacuum, ω is the angular frequency of the wave, and $\mathbf{k} = (k_x, k_y, k_z)$ is the wavevector. In Cartesian geometry, the electric field can be written as [61]

$$\mathbf{E}(\mathbf{x}, t) = \Re \left\{ E_0 \int_{\mathbb{R}^3} f_k(\mathbf{k}) \hat{\mathbf{e}}_E(\mathbf{k}) \exp [i (\mathbf{k} \cdot \mathbf{x} - \omega(\mathbf{k})t)] d^3\mathbf{k} \right\}, \quad (2.2)$$

where z is defined as the main propagation direction, and $\hat{\mathbf{e}}_E(\mathbf{k})$ is a unit polarization vector satisfying $\mathbf{k} \cdot \hat{\mathbf{e}}_E = 0$. A corresponding expression for the magnetic field can be written. Equation (2.2), together with the analogous expression for the magnetic field, guarantees that each spectral component is an exact solution of Maxwell's equations in vacuum [62, 63]. The complex spectral weight $f_k(\mathbf{k})$ determines the amplitude and phase of each mode and therefore fully specifies the pulse, while the overall maximum amplitude of the electric field is set by E_0 .

We now use Eq. (2.2) to model ST pulses. Figure 2.1 illustrates how a focal spot moving at a prescribed velocity v_f can be generated. The allowed spectral components must be restricted to the intersection of the vacuum light cone, Eq. (2.1), shown in green in Fig. 2.1, with a spectral plane defined by the desired trajectory,

$$\omega = v_f k_z + |\mathbf{k}_0| c (1 - v_f/c), \quad (2.3)$$

shown in purple in Fig. 2.1. The resulting intersection is highlighted by the red curve in Fig. 2.1. Here, $(\omega_0, \mathbf{k}_0)$ defines the specific spectral plane under consideration, which passes through the point $\omega(\mathbf{k}) = \omega_0$ when $\mathbf{k} = (0, 0, k_0)$. In the $(k_z, \omega/c)$ plane, this curve is open in the superluminal regime ($v_f > c$), as shown in Fig. 2.1 (a), and becomes a closed loop in the subluminal regime ($v_f < c$), as shown in Fig. 2.1 (b). In addition, we restrict the spectrum to components with $k_z > 0$ in order to describe forward-propagating pulses.

The intersection of Eqs. (2.1) and (2.3) defines the curve $\kappa = 0$, where

$$\kappa = |\mathbf{k}| - k_0 - (k_z - k_0). \quad (2.4)$$

To construct a physical pulse with finite energy, we allow the spectrum to have a finite bandwidth around $\kappa = 0$. For simplicity, we further assume a separable form for the spectral weight,

$$f_k(\mathbf{k}) = f_{\parallel}(\kappa) f_{\perp}(\mathbf{k}_{\perp}), \quad (2.5)$$

where $\mathbf{k}_{\perp} = (k_x, k_y)$ denotes the transverse wavevector. Among the possible choices for $f_{\perp}$, we adopt a Gaussian transverse spectrum,

$$f_{\perp}(\mathbf{k}_{\perp}) = \frac{w_0^2}{4\pi} \exp\left(-\frac{k_{\perp}^2 w_0^2}{4}\right), \quad (2.6)$$

which corresponds to a beam waist w_0 in real space.

The longitudinal spectral function $f_{\parallel}(\kappa)$ is obtained by prescribing a desired temporal envelope and computing its Fourier transform with respect to the variable κ . In this work, we employ a flattened longitudinal profile (e.g., a squared-sine envelope with a rapid rise and a long flat-top region) to produce a nearly constant on-axis intensity over an extended propagation distance. Figure 2.2 shows representative examples of superluminal and subluminal space-time pulses constructed using this procedure, together with the temporal envelope employed.

2.1.2 DISCRETE IMPLEMENTATION IN OSIRIS

We now briefly describe the OSIRIS implementation of the method introduced in Sec. 2.1.1, originally developed by R. Almeida [61]. The algorithm constructs the electromagnetic fields by evaluating the superposition in Eqs. (2.2), with the continuous spectral weight f_k mapped onto a discrete grid in k -space. This sampling implies a periodic reconstruction in space-time, which sets the replica periodicities discussed in Sec. 2.1.3. The fields are obtained by summing the resulting set of exact vacuum modes (for example, plane waves in Cartesian geometry or Bessel modes in cylindrical geometry) over the entire spatial domain at $t = 0$.

The discretized implementation of the spectral weight introduces two main challenges. The first is that initialization is performed in Fourier space and, upon transformation to configuration space, the simulation domain must be sufficiently large to contain the entire pulse. One possible way to address this issue is to continuously inject waves from the simulation boundaries, as will be discussed in detail in Sec. 2.2.2.




FIGURE 2.2: Two examples of spatio-temporal pulses obtained with the discrete implementation in OSIRIS. In (a), we have an example of the shape of a superluminal pulse with $v_f = 1.5c$. In (b), we show an example of the shape of a subluminal pulse with $v_f = 0.5c$. The blue-red color-scale represents the electric field, and the black line corresponds to the pulse's full longitudinal envelope $f_{\parallel}(z)$ in configuration space.

The second challenge is that discretizing k -space naturally introduces periodicity in configuration space, which may manifest in the simulation domain if not properly controlled. This issue is addressed in the following section.

2.1.3 ACCURACY CONSTRAINTS OF THE SPECTRAL INITIALIZATION

As discussed in Sec. 2.1.2, the numerical implementation relies on sampling the continuous spectral weight $f_k(\mathbf{k})$ on a finite and discretized set of modes with spacing Δk_i , where $i = (x, y, z)$. As in any discrete Fourier-type representation, this sampling leads to a periodic reconstruction of the electromagnetic fields in configuration space, with repetition lengths given by

$$T_{\text{rep},i} = \frac{2\pi}{\Delta k_i}. \quad (2.7)$$

These periodic replicas are numerical artifacts and may overlap with the physical pulse inside the simulation domain if not properly controlled.

To avoid unphysical effects in the longitudinal direction, the repetition length must exceed the simulation window size. This requires

$$T_{\text{rep},z} > L_{\text{box},z}, \quad (2.8)$$

so that no longitudinal replica enters the computational domain during the time interval of interest.

In the transverse directions, the repetition length must be sufficiently large to prevent overlap between the physical pulse and its periodic replicas throughout propagation. The quantity $T_{\text{rep},(x,y)}$ represents the distance in configuration space between adjacent replicas of the pulse introduced by the discrete sampling in k -space. To avoid unphysical interference, this distance must exceed the maximum transverse extent of the pulse within the simulation domain.

For subluminal focal velocities $v_f < c$, the transverse evolution of the pulse can be estimated using the paraxial Gaussian-beam model [61, 64],

$$w(\xi) = w_0 \sqrt{1 + \left(\frac{2\xi}{k_0 w_0^2 (1 - v_f/c)} \right)^2}, \quad (2.9)$$

where $\xi = z - v_f t$ is the co-moving coordinate.

The transverse evolution of the pulse is described by the waist function $w(\xi)$, where $\xi = z - v_f t$ is the co-moving coordinate. At the focus, $w(0) = w_0$, while away from the focus, the pulse diffracts and its transverse size increases according to Eq. (2.9). Since the simulation domain is finite, the pulse only occupies a limited longitudinal interval ξ , determined by the boundaries of the computational box and the location of the focal point. The largest transverse size attained by the pulse inside the simulation window is therefore

$$w_{\max} = w(\xi_{\max}), \quad (2.10)$$

which corresponds to the maximum beam radius that appears anywhere in the simulation.

Focusing on the transverse repetition length $T_{\text{rep},(x,y)}$, artificial replicas of the pulse appear at transverse positions $x = \pm T_{\text{rep},x}$ and $y = \pm T_{\text{rep},y}$. To avoid unphysical overlap between the physical pulse and its replicas, the distance between neighboring replicas must exceed the transverse size of the pulse throughout its evolution. Since the pulse occupies approximately the region $|x| \lesssim w_{\max}$, a sufficient geometric condition to prevent overlap is

$$T_{\text{rep},(x,y)} \gtrsim 2w_{\max}. \quad (2.11)$$

This criterion ensures that, even at its widest point, the pulse remains spatially isolated from its periodic replicas. If this condition is violated, the replicas overlap with the main pulse, resulting in artificial interference patterns and transverse field modulations.

The impact of these constraints is illustrated in Fig. 2.3 for a laser pulse with parameters $a_0 \equiv eE_0/\omega_0 m_e c = 3$, $k_0 w_0 = 10$, and $v_f = 0.96c$, where E_0 is the laser electric field maximum amplitude, e is the elementary charge, and m_e is the electron mass. Figure 2.3(a) shows an optimized configuration in which both longitudinal and transverse repetition lengths satisfy the above criteria, resulting in a smooth pulse free from replica overlap. In contrast, Fig. 2.3(b) shows the onset of artificial longitudinal modulations when $T_{\text{rep},z}$ is too small, while Fig. 2.3(c) illustrates the analogous transverse distortion that arises when $T_{\text{rep},(x,y)}$ is insufficient compared with the pulse waist.

In summary, an artifact-free implementation can be obtained by first computing the maximum transverse size $w_{\max}$ of the pulse inside the simulation window using Eq. (2.9). The transverse repetition lengths should then be chosen such that $T_{\text{rep},(x,y)} \geq 2w_{\max}$, while the longitudinal repetition length must satisfy $T_{\text{rep},z} > L_{\text{box}}$.

2.2 SIMULATION DESIGN AND PHYSICAL VALIDATION

The spectral geometry of the spatio-temporal pulses introduced in Sec. 2.1 couples the driver's transverse scale to the longitudinal extent of its high-intensity region. This differs from con-




FIGURE 2.3: Comparison of the injected pulse electric field for three choices of the longitudinal and transverse repetition lengths used in the spectral reconstruction. (a) $k_0 T_{\text{rep},z} = 700$ and $k_0 T_{\text{rep},(x,y)} = 1000$, for which the reconstructed pulse is smooth and free of aliasing artefacts. (b) $k_0 T_{\text{rep},z} = 200$ and $k_0 T_{\text{rep},(x,y)} = 100$, where the longitudinal repetition length is too small, producing spurious replicas along the propagation axis. (c) $k_0 T_{\text{rep},z} = 700$ and $k_0 T_{\text{rep},(x,y)} = 200$, where the transverse repetition length is insufficient, and replicas enter the simulation domain, leading to distorted, unphysical structures.

ventional Gaussian drivers, for which the spot size w_0 and pulse duration τ are usually treated as independent design parameters. As a result, resonant wakefield excitation does not follow the standard scaling laws directly and instead requires a dedicated exploration of the driver parameter space $(\omega_0, w_0, v_f, a_0)$.

Moreover, because the prescribed focal velocity v_f is decoupled from the envelope group velocity, long-distance propagation introduces practical numerical constraints. The relative drift between the moving focus and the finite envelope sets a survival length, while the broad spectral support required for $v_f \neq c$ leads to substantial transverse spreading. Together, these effects can make full-domain initialization increasingly expensive in the subluminal regime.

This section addresses both the physical and numerical requirements. Section 2.2.1 analyses how v_f and the driver geometry constrain the resonant conditions for wake excitation and derives practical scaling guidelines for parameter selection. Section 2.2.2 then introduces and benchmarks continuous wall injection, showing how it alleviates the domain-size constraints and establishes the criteria needed for efficient and stable particle-in-cell modeling of these spatially extended drivers.

2.2.1 OPTIMIZATION OF SPATIO-TEMPORAL DRIVERS FOR WAKEFIELD EXCITATION

Efficient wakefield generation relies on matching the laser-pulse dimensions to the characteristic plasma scales. Transversely, the energy should be concentrated within a region comparable to the plasma skin depth, $k_p^{-1} = c/\omega_p$, where $\omega_p = (n_0 e^2 / \epsilon_0 m_e)^{1/2}$ is the electron plasma frequency, n_0 is the plasma density, and ϵ_0 is the vacuum permittivity. Longitudinally, resonant excitation typically requires the driver length to be of the order of half the wake wave-

length [65]. Two points are important in the present context. First, for the spatio-temporal pulses considered in this work, the transverse and longitudinal profiles are intrinsically coupled and cannot be tuned independently. Second, when the wake is analyzed in the co-moving coordinate $\xi = z - v_f t$, the apparent wake period in ξ depends on the focal velocity v_f . In particular, a plasma oscillation at ω_p corresponds to the spatial period

$$\lambda_{\text{wake}}(v_f) = \frac{2\pi v_f}{\omega_p}, \quad (2.12)$$

so the longitudinal matching condition in ξ is $L_{\text{pulse}} \sim \lambda_{\text{wake}}/2$.

The corresponding field envelope, valid in the paraxial limit ($k_0 w_0 \gg 1$), can be written as

$$E(r, \xi) \propto \frac{w_0}{w(\xi)} \exp\left(-\frac{r^2}{w^2(\xi)}\right), \quad (2.13)$$

where $w(\xi)$ is given by Eq. (2.9). In the subluminal regime ($|v_f| < c$), Eq. (2.9) has the same functional form as the diffraction of a vacuum Gaussian beam, but with the propagation coordinate replaced by the co-moving variable ξ . This motivates defining an effective Rayleigh length in ξ ,

$$\xi_R \equiv \frac{k_0 w_0^2}{2} \left(1 - \frac{v_f}{c}\right), \quad (2.14)$$

so that $w(\xi) = w_0 \sqrt{1 + (\xi/\xi_R)^2}$. Within this paraxial model, the on-axis intensity scales as

$$I(0, \xi) \propto \left[1 + \left(\frac{\xi}{\xi_R}\right)^2\right]^{-1}, \quad (2.15)$$

which is Lorentzian in ξ with full width at half maximum

$$L_{\text{FWHM}} = 2\xi_R = k_0 w_0^2 \left(1 - \frac{v_f}{c}\right). \quad (2.16)$$

For a standard Gaussian beam, the Rayleigh length characterizes the spatial decay of the peak intensity with propagation distance. Here, by contrast, ξ_R sets the extent of the high-intensity region in the co-moving coordinate. Since a Lorentzian profile has long tails, we characterize the effective driver length by the $1/e^2$ intensity width (e is the Euler number),

$$L_{\text{pulse}} \equiv L_{1/e^2} = \sqrt{e^2 - 1} L_{\text{FWHM}}. \quad (2.17)$$

Imposing the longitudinal matching condition $L_{\text{pulse}} \sim \lambda_{\text{wake}}/2$ then yields the estimate

$$k_0 w_0 \approx \left[\frac{\pi}{\sqrt{e^2 - 1}} \frac{\omega_0}{\omega_p} \frac{v_f/c}{1 - v_f/c} \right]^{1/2}. \quad (2.18)$$

The validity of the paraxial expression in Eq. (2.18) is strictly bounded by the condition $k_0 w_0 \gg 1$. Several of the regimes considered in this manuscript are characterised by a broad angular spectra that lie outside this strict domain of validity. Equation (2.18) should therefore




FIGURE 2.4: Wakefield excitation for different choices of v_f and w_0 . Panels (a) and (b) show the case $v_f = 0.8c$: for (a) $k_0 w_0 = 10$ the driver is too long and overlaps with most of the first plasma period, whereas for (b) $k_0 w_0 = 7.1$ the driver length is better matched and a clearer wake structure is obtained. The same trend is more pronounced in panels (c) and (d) for $v_f = 0.6c$, comparing (c) $k_0 w_0 = 10$ with the tighter-focus case (d) $k_0 w_0 = 4.3$. The greyscale shows the electron density, and the blue-red colour scale shows the laser electric field. All cases use $\omega_0/\omega_p = 10$ and $a_0 = 2.2$.

be regarded as a practical design estimate, capturing the correct scaling of w_0 with v_f and ω_0/ω_p rather than as a quantitatively exact prediction.

This longitudinal estimate should be complemented by the usual transverse matching condition for wakefield excitation. At the pulse center ($\xi = 0$), the transverse profile reduces to a Gaussian with spot size w_0 , so efficient coupling in conventional wakefield theory typically requires $k_p w_0 \sim \mathcal{O}(1)$. Since both the longitudinal resonance condition and the transverse matching condition constrain the same parameter w_0 , it is not, in general, possible to satisfy both simultaneously. Efficient operation therefore requires working in a regime where the two estimates are compatible, or prioritizing one constraint and accepting a compromise in the other.

Figure 2.4 illustrates how the interplay between w_0 and v_f can produce either weakly resonant or near-optimal wakefields for fixed parameters $\omega_0/\omega_p = 10$ and $a_0 = 2.2$. Figures 2.4(a) and (b) correspond to $v_f = 0.8c$. In panel (a), the choice $k_0 w_0 = 10$ is not optimal, whereas in panel (b) we use the value predicted by Eq. (2.18), $k_0 w_0 \approx 7.1$. The latter case exhibits a markedly cleaner plasma response (shown in grey) to the ST driver. Similarly, Figs. 2.4 (c) and (d) correspond to $v_f = 0.6c$. For $k_0 w_0 = 10$ [panel (c)], the driver is noticeably longer than the plasma response, while using the predicted optimal value $k_0 w_0 \approx 4.3$ [panel (d)] restores a well-defined wake structure even at this lower focal velocity.




FIGURE 2.5: Waterfall plots of the accelerating electric-field amplitude for different focal velocities: (a) $v_f = 0.8c$, (b) $v_f = 0.99c$, and (c) $v_f = 1.4c$.

However, this optimization strategy may introduce trade-offs. Experimentally, generating tightly focused broadband wave packets requires large-aperture optics. In addition, a smaller spot size produces a narrower transverse wake structure, which reduces the acceptance for a witness beam. Despite these constraints, the wakefield remains stable over extended propagation distances when the appropriate spatio-temporal coupling is employed, as illustrated in Fig. 2.5, which shows waterfall plots of the longitudinal electric field for spatio-temporal drivers with different focal velocities, together with the corresponding wake evolution that closely follows the imposed focal velocity.




FIGURE 2.6: Wakefield excitation for a fixed driver with $v_f = 0.96c$ and $k_0 w_0 = 17$, for different values of a_0 . The upper panels show the electron density response and the on-axis laser electric-field envelope, while the lower panels show the on-axis accelerating field. (a) $a_0 = 0.22$, (b) $a_0 = 0.44$, (c) $a_0 = 1.3$, and (d) $a_0 = 2.2$.

Once the spatio-temporal geometry is optimized for resonance, the wakefield amplitude

can be further increased by raising the driver strength a_0 . As illustrated in Fig. 2.6, increasing from (a) $a_0 = 0.22$, (b) $a_0 = 0.44$, (c) $a_0 = 1.3$, and (d) $a_0 = 2.2$ leads to progressively larger accelerating fields and to changes in the wake structure, going from the linear to the non-linear regime. The wakefield patterns obtained, particularly in panel (d), are qualitatively similar to those reported in recent experiments demonstrating wakefield generation driven by structured light [66]. We note that, in the strongly driven cases shown in panels (c) and (d), the on-axis laser envelope develops a modulation behind the focus with a period close to the plasma wavelength. This is a signature of self-modulation of the driver by the wake. Unlike a short Gaussian pulse, ST drivers carry significant intensity well beyond the focus, so the laser tails are intense enough to interact nonlinearly with the plasma such that the resulting modulation of the plasma density imprints a corresponding modulation on the laser envelope. Although ST pulses are constructed to be shape-invariant in vacuum, their structure can therefore be modified by the laser-plasma interaction in the nonlinear regime.

In summary, the paraxial model provides a practical framework for selecting spatio-temporal pulse parameters that yield near-resonant and stable wakefield excitation. Because the longitudinal and transverse properties of the driver are intrinsically coupled, resonance cannot be enforced by tuning the pulse duration alone. In the subluminal regime, reducing the focal velocity v_f increases the effective longitudinal extent of the high-intensity region in the co-moving coordinate, which can violate the resonance condition. Maintaining resonant excitation, therefore, requires compensating for this elongation by reducing the spot size w_0 according to the scaling in Eq. (2.18). This provides a simple design strategy to identify suitable (w_0, v_f) combinations before performing full numerical optimization.

2.2.2 SIMULATION DOMAIN SIZE AND WALL INJECTION FOR SPATIO-TEMPORAL PULSES

In conventional laser-wakefield simulations, the driver is typically a short pulse whose longitudinal extent is small compared with its diffraction length. A moving window that follows the driver at (approximately) the speed of light can therefore contain both the laser and the plasma response within a compact computational domain.

Velocity-controlled ST pulses differ in two key ways, with direct consequences for the required box size. First, the driver is not a short wave packet: the laser energy is distributed over a comparatively long envelope, while the high-intensity peak is formed by a constructive-interference condition that propagates through this envelope at a prescribed focal velocity v_f . In a standard moving window that advances at c (and thus tracks the envelope), the envelope remains approximately stationary, whereas the intensity peak drifts through it at the relative speed $|c - v_f|$ in subluminal cases. This slippage is illustrated in Fig. 2.7, where the focus slips from the leading side of the envelope towards the trailing side as propagation proceeds. In the superluminal regime, the focus instead slips from the trailing side towards the leading side.

This relative motion implies that the high-intensity peak has a finite lifespan within a given envelope. We quantify this effect by defining a survival length L_{surv} , namely the maximum distance over which the focus remains contained inside an envelope of longitudinal extent




FIGURE 2.7: Schematic of a subluminal ST pulse showing the full laser envelope (black) and the localized high-intensity region (grey line; blue-red shading). Because the envelope propagates at c while the focus advances at $v_f < c$, the intensity peak slips backward within the envelope between (a) t_1 and (b) $t_2 > t_1$. Once the peak reaches the trailing edge of the envelope, the driver can no longer sustain an efficient wake.

L. For a subluminal pulse with the focus initially near the leading edge, the survival time is $\Delta t_{\text{surv}} \simeq L/|c - v_f|$, which gives

$$L_{\text{surv}} \simeq \frac{v_f}{|c - v_f|} L. \quad (2.19)$$

This expression is valid in both the subluminal and superluminal regimes, and diverges as $v_f \rightarrow c$, recovering the conventional moving-window picture. In wakefield-driven acceleration, L_{surv} provides an upper bound on the propagation distance over which the driver can maintain a stable wake structure purely from this geometric constraint. In practice, achieving long interaction lengths at focal velocities far from c therefore requires a correspondingly long envelope (and thus a larger longitudinal domain).

Constructing a pulse with a prescribed focal velocity $v_f \neq c$ requires a broad spectral support with significant transverse wavevector content. In the subluminal regime, this manifests as a large effective divergence of the envelope, so that the pulse occupies a wider transverse extent than a conventional Gaussian driver of comparable w_0 . Consequently, avoiding artificial truncation generally demands a larger transverse domain.

These considerations must be reconciled with the numerical constraints discussed in Sec. 2.1.3. The longitudinal repetition length introduced by spectral discretizations may need to be increased as the envelope length and interaction distance grow.

The simulation domain requirements depend on the chosen ST pulse parameters and can be estimated from simple geometric considerations. For subluminal pulses with a Gaussian transverse profile, the envelope expands transversely as it propagates. For a pulse with spot size w_0 and focal velocity v_f , an effective divergence half-angle can be estimated as

$$\tan \theta_{\text{div}} \approx \frac{2}{k_0 w_0 (1 - v_f/c)}. \quad (2.20)$$

Figure 2.8 shows an example with $v_f = 0.97c$, $k_0 w_0 = 14$, and $\omega_0/\omega_p = 10$. Using Eq. (2.20), this corresponds to $\theta_{\text{div}} \approx 67^\circ$, indicated by the dashed black lines in Fig. 2.8 (a). For a Gaussian transverse profile, this aperture contains approximately 86

In a quasi-3D geometry, where the transverse coordinate ranges from $r \in [0, L_\perp]$, a practical guideline is

$$L_\perp \gtrsim 4w_0 + z_{\text{prop}} \tan \theta_{\text{div}}, \quad (2.21)$$

so that the relevant initial envelope extent (taken here as $4w_0$) and its subsequent expansion remain contained within the computational domain. For a propagation throughout the full survival time, a conservative choice is $z_{\text{prop}} \sim c \Delta t_{\text{surv}}$, which typically implies a large transverse domain.

Fortunately, the transverse-domain requirements can be relaxed by using a boundary-injection approach. Instead of initializing the full field distribution throughout the simulation domain at $t = 0$, OSIRIS can inject the driver continuously through the transverse boundaries by evaluating the superposition in Eq. (2.2) at each time step. In this scheme, the angular spectrum required to form the ST pulse is supplied dynamically from the walls, so the simulation does not need to contain the full laterally expanding envelope at all times. This allows the transverse domain to be reduced substantially while preserving the correct field structure in the interaction region and maintaining consistency with the fields obtained from full-domain initialization.

Figure 2.8 (a) illustrates that the transverse domain can be reduced substantially by focusing the computational region on the portion of the pulse that is most relevant for the laser-plasma interaction, i.e. where the intensity is sufficiently high. Figures 2.8 (b) and (c) show simulations performed in this smaller box without boundary injection. While the early-time response is captured, the wakefield at later times [panel (c)] becomes markedly weaker because the pulse energy is no longer represented correctly: the laterally diverging components required to sustain the moving focus are truncated by the boundaries.

Figures 2.8 (d) and (e) show the corresponding simulations using wall injection. At early times, the results are indistinguishable from those without wall injection (compare panels (b) and (d)). At later times, wall injection sustains a wakefield amplitude comparable to that obtained initially, although deviations from the early-time structure (panel (d)) emerge due to laser-plasma interaction effects. In particular, the plasma modifies the field components that would otherwise converge from the transverse “wings” to form the interference pattern that defines the moving focus. This highlights an important point: although ST pulses can be designed to be shape invariant in vacuum, their structure can be altered by interaction with plasma.

To corroborate that the wall-injection algorithm preserves the intended propagation-invariant intensity profile, Fig. 2.8 (f) shows the evolution of the peak laser intensity in vacuum, with and without wall injection. Without injection, the peak intensity decays as the pulse footprint expands beyond the reduced domain, whereas with wall injection it remains approximately constant. This demonstrates that, in a vacuum, wall injection maintains the desired field structure while enabling substantially smaller transverse domains.

To quantify the practical benefit of wall injection in a realistic PIC setup, we benchmark the wall-injection algorithm against the conventional approach in which the full pulse is initialized throughout the domain at $t = 0$. Because wall injection evaluates the analytic field at the transverse boundaries throughout the run, it introduces an approximately constant computa-




FIGURE 2.8: Reduced-domain simulations of a subluminal ST driver with and without wall injection. (a) schematic of the injected field and boundary locations. (b,d) early-time snapshots in the co-moving coordinate ξ without and with wall injection, respectively. (c,e) late-time snapshots showing wake degradation without injection and sustained excitation with injection. (f) peak intensity in vacuum versus time, confirming that wall injection preserves the intended propagation-invariant behavior in the reduced domain. Greyscale shows the electron density, and the blue-red scale shows the laser electric field.

tional overhead per time step, increasing the time per iteration by a factor of ~ 1.5 – 2 . This overhead is typically more than compensated by the reduction in domain size enabled by wall injection: without injection, containing the laterally expanding field can require a transverse domain that is ~ 10 – 100 times larger, depending on the pulse parameters and the interaction length. Since the overall computational workload scales with the number of degrees of freedom advanced in time, reducing the domain size directly reduces the total computational resources required. In quasi-3D, the cost scales approximately with the number of grid cells, while in full 3D it scales as $N_{\perp}^2$ in the transverse resolution. Consequently, wall injection can reduce both the memory footprint and the overall computational cost, often by one to two orders of magnitude relative to otherwise equivalent simulations.

As a practical guide, simulations of subluminal ST drivers should be sized using the survival time $\Delta t_{\text{surv}} \simeq L/(c - v_f)$ and the effective divergence $\tan \theta_{\text{div}} \approx 2/[k_0 w_0 (1 - v_f/c)]$. In

quasi-3D geometry, choose the radial extent to satisfy $L_{\perp} \gtrsim 4w_0 + z_{\text{prop}} \tan \theta_{\text{div}}$, with z_{prop} set by the intended interaction length (a conservative choice is $z_{\text{prop}} \sim c, \Delta t_{\text{surv}}$ when propagating over the full survival time). If these requirements make the transverse domain impractical, enable wall injection, which supplies the angular spectrum at the boundaries and allows the interaction region to be simulated in a much smaller box.

2.3 CONCLUSIONS

This Chapter presented a practical workflow for modeling velocity-controlled ST laser drivers in particle-in-cell simulations. We relied on a Maxwell-consistent spectral construction that defines ST pulses as a superposition of exact vacuum solutions, ensuring the correct spatio-temporal correlations. We then described how to map this continuous formulation onto a discrete k -space grid suitable for numerical initialization.

A critical numerical consequence of discretizing the spectrum is the appearance of real-space replicas, whose periodicity is determined by the k -space sampling. We explicitly addressed this issue by providing simple, geometry-based criteria for selecting the longitudinal and transverse repetition lengths. These criteria ensure that replica pulses do not enter the interaction region during the simulation, making the spectral method reliable in practical wakefield setups.

We have shown that the ST spectral geometry intrinsically couples the driver's transverse scale and the longitudinal extent of its high-intensity region. As a result, wakefield optimization cannot be performed by varying the pulse duration independently of the spot size, as is typically done for conventional Gaussian drivers. In the subluminal regime, decreasing v_f elongates the effective high-intensity region in the co-moving coordinate; we have demonstrated that near-resonant excitation can be recovered by reducing w_0 according to the scaling in Eq. (2.18), and we have validated this design rule with PIC simulations across representative parameter scans.

We have addressed the numerical constraints that arise when $v_f \neq c$. First, the slip between the envelope (advancing at c) and the moving focus introduces a survival length, limiting how long a stable wake can be sustained within a finite envelope. Second, the broad angular content required for subluminal pulses leads to substantial transverse expansion. These geometric effects make full-domain initialization increasingly expensive for long subluminal runs. We have therefore employed a continuous wall-injection scheme that supplies the required self-consistent fields dynamically through the transverse boundaries. We have shown that wall injection preserves the intended propagation-invariant behavior in vacuum and sustains the driven wake in reduced transverse domains. We have also highlighted that, although ST drivers are constructed to be shape-invariant in vacuum, their structure can be modified by the laser-plasma interaction in the nonlinear regime: in particular, the long envelope characteristic of ST pulses keeps the tails at sufficient intensity to drive a nonlinear plasma response, leading to self-modulation of the laser by the wake. This is a generic feature of velocity-controlled ST drivers in wakefield-relevant regimes and should be taken into account when designing experiments or simulations with these pulses. Although wall injection introduces a modest per-step

overhead, the reduction in domain size typically yields a net decrease in overall computational cost.

Collectively, these results provide a concise recipe for the robust PIC modeling of velocity-controlled spatio-temporal drivers: (i) specify the driver through a Maxwell-consistent spectral decomposition; (ii) choose the k -space discretizations so that the associated repetition lengths exceed the relevant longitudinal and transverse window; (iii) select (w_0, v_f) using the resonance scaling in Eq. (2.18) to obtain strong and stable wake excitation; and (iv) for long subluminal propagation, use wall injection to avoid representing the full laterally expanding envelope inside the computational domain. The methodology established here forms the numerical foundation for the phase-locking acceleration scheme developed in Chapter 3 and the beam-quality analysis of Chapter 4.

2.4 SIMULATION SETUP DETAILS

All simulations are performed with the OSIRIS azimuthal-mode decomposition algorithm in $(r, z-\varphi)$ geometry (quasi-3D) [33, 34]. In this approach, the computational grid is two-dimensional in (r, z) , while the dependence on the azimuthal angle φ is represented analytically by expanding the fields and currents into a finite number of Fourier modes. In other words, there is no mesh in φ : non-axisymmetric structure is captured through the retained azimuthal harmonics, and the accuracy is controlled by the number of modes included. This provides a practical compromise between fully three-dimensional modeling and the lower cost of axisymmetric simulations. We retain the cylindrically symmetric mode $m = 0$ and the first non-axisymmetric mode $m = 1$, which is sufficient to capture the laser polarization and the main wake structure at reduced computational cost. The plasma is initialized with a uniform density n_0 , preceded by a short vacuum-to-plasma density ramp of length $L_{\text{ramp}} = 4c/\omega_p$. Plasma electrons are sampled with 32 particles per cell. A moving window propagating at the speed of light is employed in all simulations presented in this thesis's Chapter. For the parameter scans shown in Figs. 2.5 and 2.6, the simulation domain is $L_z \times L_r = 100c/\omega_p \times 20c/\omega_p$ with $N_z \times N_r = 10000 \times 400$, corresponding to $\Delta z = 0.01c/\omega_p$ and $\Delta r = 0.05c/\omega_p$. Absorbing boundary conditions are applied at the transverse edge of the domain using perfectly matched layers. The laser driver is polarised along E_y . Its longitudinal envelope is a flat-top profile with squared-sine ramps on both sides (rise/fall length $5c/\omega_p$ and flat-top length $90c/\omega_p$).

Take-home messages

- **We define Maxwell-consistent ST drivers in Fourier space.** The pulses are constructed as a superposition of exact vacuum solutions to ensure the correct spatio-temporal correlations.
- **Discretizing the spectrum introduces periodic replicas.** The finite k -space grid imposes longitudinal and transverse repetition lengths. To avoid replicas, ensure $T_{\text{rep},z} > L_{\text{box}}$ and $T_{\text{rep},\perp} \gtrsim 2w_{\text{max}}$, where w_{max} is the largest transverse size of the pulse inside the simulation window.
- **For ST pulses, longitudinal and transverse scales are coupled.** In ST pulses the spot size and the effective high-intensity length cannot be tuned independently, so optimization must be performed in the joint $(\omega_0, w_0, v_f, a_0)$ space.
- **Resonance conditions and pulse dimensions.** To maximize wakefield excitation, the driver dimensions must match the plasma response. Transversely, the spot size w_0 is typically chosen such that $k_p w_0 \sim \mathcal{O}(1)$. Longitudinally, the effective pulse length must match half the wake wavelength $\lambda_{\text{wake}} \approx 2\pi v_f / \omega_p$. In the subluminal regime, satisfying this longitudinal resonance imposes a specific scaling on the spot size w_0 :

$$k_0 w_0 \approx \left[\frac{\pi}{\sqrt{e^2 - 1}} \frac{\omega_0}{\omega_p} \frac{v_f/c}{1 - v_f/c} \right]^{1/2}$$

This relation implies that as the focal velocity v_f decreases, the spot size must be reduced (tighter focusing) to prevent the effective driver length from becoming too long for the wake.

- **Envelope-focus slippage sets a survival length.** For $v_f < c$, the intensity peak drifts through the envelope at the relative speed $c - v_f$. We quantify this effect by defining a survival length L_{surv} , namely the maximum distance over which the focus remains contained inside an envelope of longitudinal extent L as

$$L_{\text{surv}} \simeq \frac{v_f}{c - v_f} L.$$

- **Wall injection makes long subluminal runs practical.** Subluminal pulses require broad angular content and therefore large transverse domains. Continuous wall injection supplies the required spectrum at the boundaries and can strongly reduce the computational cost, in particular for highly non-paraxial pulses.

CHAPTER 3

DECAYING PARTICLES ACCELERATION

The particle physics community is actively considering muon colliders as one possible next step for high-energy physics studies [67]. Muons, like electrons, are fundamental particles, meaning their total energy is available in collisions, unlike hadrons, whose substructure reduces the effective collision energy [68–70]. Nonetheless, due to their finite lifetime ($\tau_\mu = 2.2\,\mu\text{s}$ in the proper frame), muons must be collected, cooled, and rapidly accelerated before a significant number of them decay [71]. This poses a major challenge for conventional accelerators, which rely on RF cavities with limited accelerating gradients.

Similarly, other short-lived particles such as pions and kaons, whose lifetimes are even shorter ($\tau_{\pi^\pm} = 26\,\text{ns}$ and $\tau_{K^\pm} = 12\,\text{ns}$), cannot be accelerated efficiently using standard RF technology [72]. Plasma-based acceleration, with its ultra-high gradients on the order of $\sim\text{GV/cm}$, offers a promising pathway to quickly boost the energy of such rapidly decaying particles, potentially overcoming the limitations imposed by their short lifetimes.

This Chapter presents an acceleration method that addresses the challenge of accelerating short-lived, sub-relativistic particles using laser-driven plasma-wakefield acceleration [73]. We begin in Sec. 3.1 by quantifying the survival fraction of particles during acceleration using both conventional radio-frequency (RF) cavities and plasma-based accelerators, highlighting the impact of acceleration gradient and particle lifetime.

Following this motivation, Sec. 3.2 investigates the use of velocity-controlled subluminal drivers as a first-order solution to the trapping problem. We start in Sec. 3.2.1 by testing the application of standard Laser Wakefield Acceleration (LWFA) to muons, confirming that conventional luminal wakes fail to trap sub-relativistic particles. This motivates the introduction of the subluminal-driver concept in Sec. 3.2.2, where we show that while matching the driver velocity to the particle allows for initial trapping, it fundamentally limits the maximum attainable energy due to dephasing where the accelerating beam overtakes the wake. Section 3.2.3 then explores the optimization of this simple scheme by varying the focal velocity of the driver for the maximum energy gain.

To further push the maximum energy gain in a single stage, Sec. 3.3 presents the phase-locking framework at the heart of this chapter: a general theory for phase-locking via tailored plasma density profiles in combination with spatio-temporal pulses with subluminal

focal velocity. We derive the exact density gradients required to synchronize the wakefield accelerating region with the accelerated particle beam. We then demonstrate the universality of this model in Sec. 3.4.1 by calculating the specific density profiles required for different species, including muons, pions, and kaons. The versatility of the scheme is further expanded in Sec. 3.4.2, which shows that phase-locking can be achieved even if the beam is injected into secondary accelerating buckets rather than the first wake period. Finally, Sec. 3.4.3 validates the robustness of the scheme against variations in the initial particle velocity, proving its feasibility for realistic, broad-spectrum particle sources.

3.1 SURVIVING PARTICLES FRACTION IN AN ACCELERATOR

Before proceeding to test the feasibility of employing plasma-based accelerators to promptly extend the decay time of unstable particles, we can estimate the survival fraction during an acceleration stage as a function of the accelerating gradient. Starting from the decay law in the laboratory frame [74, 75],

$$\frac{dN}{dt} = -\frac{N}{\gamma \tau_p}, \quad (3.1)$$

where N is the number of surviving particles, γ is the Lorentz factor, and τ_p is the particle's proper lifetime (rest-frame mean lifetime). Assuming a constant accelerating field E_{acc} and $\beta \simeq 1$, where $\beta = v/c$ is the particle's velocity normalized by the speed of light, during most of the acceleration stage the energy gain satisfies

$$\frac{d\gamma}{dt} \simeq \frac{q E_{\text{acc}}}{m_b c}, \quad (3.2)$$

which integrates to the fraction of surviving particles S [74, 75]

$$S \simeq \left(\frac{\gamma_i}{\gamma_f} \right)^{\nu_p}, \quad \nu_p \equiv \frac{m_b c}{q E_{\text{acc}} \tau_p}, \quad (3.3)$$

where the Lorentz factors $\gamma_{i,f} = 1 + K_{i,f}/(m_b c^2)$ can be computed from the initial and final kinetic energies K_i and K_f , m_b is the rest mass of the particles, τ_p the proper lifetime, q the charge of the particle and c the speed of light.

As a representative case, we examine the acceleration of muon (μ^-), pion (π^-), and kaon (K^-) beams from an initial kinetic energy of $K_i = 200 \text{ MeV}$ to a final energy of $K_f = 1 \text{ GeV}$. Two scenarios are considered: one employing a plasma accelerator with an accelerating gradient of $E_{\text{acc}} = 20 \text{ GV/m}$, and another using a conventional RF linac operating at $E_{\text{acc}} = 20 \text{ MV/m}$. The corresponding survival fractions S and loss fractions $1 - S$ for each particle species are listed in Table 3.1. The results clearly highlight the strong advantage of plasma-based accelerators, whose high gradients enable the prompt acceleration of short-lived particles, thereby mitigating decay losses.

Important Note — The compact expression $S = (\gamma_i/\gamma_f)^{\nu_p}$ assumes $\beta \simeq 1$ during the entire acceleration, i.e. that the particles are ultra-relativistic. For heavier particles at lower energies

TABLE 3.1: Estimated survival fraction S (and corresponding loss fraction $1 - S$) for μ^- , π^- , and K^- accelerated from the initial to the final Lorentz factor (γ_i, γ_f) in a plasma accelerator and in a conventional RF linac. The plasma and RF cases assume constant accelerating gradients of 0.2 GV/cm and 20 MV/m, respectively.

Particle	γ_i	γ_f	Accelerator	S	Loss ($1 - S$)
μ^-	2.9	10.5	Plasma (0.2 GV/cm)	99.999 %	0.001 %
μ^-	2.9	10.5	RF Linac (20 MV/m)	98.951 %	1.049 %
π^-	2.4	8.2	Plasma (0.2 GV/cm)	99.888 %	0.112 %
π^-	2.4	8.2	RF Linac (20 MV/m)	32.640 %	67.360 %
K^-	1.4	3.0	Plasma (0.2 GV/cm)	99.404 %	0.596 %
K^-	1.4	3.0	RF Linac (20 MV/m)	0.252 %	99.748 %

this assumption becomes less accurate. Without this approximation, the survival fraction reads

$$S = \exp \left[-\frac{m_b c}{q E_{\text{acc}} \tau_p} \int_{\gamma_i}^{\gamma_f} \frac{d\gamma}{\beta(\gamma) \gamma} \right], \quad \beta(\gamma) = \sqrt{1 - \frac{1}{\gamma^2}}. \quad (3.4)$$

The extra factor $1/\beta(\gamma)$ increases the decay probability for particles with lower initial γ , since they spend more time in the accelerator. This effect is negligible for muons and pions at the energies considered here, but can slightly increase the predicted loss for kaons.

3.2 PLASMA WAKEFIELD ACCELERATION WITH SUBLUMINAL DRIVERS

The estimate in Sec. 3.1 shows that ultra-high accelerating gradients can strongly mitigate decay losses by reducing the time spent in the accelerator. However, high gradient alone is not sufficient: to realize this benefit for muons, pions, and kaons, the particles must be promptly trapped into an accelerating bucket and remain phase-synchronous with the accelerating field long enough to reach the relativistic regime. This requirement is non-trivial because realistic secondary beams are typically produced at kinetic energies of a few hundred MeV, i.e. with $\beta < 1$, so their velocity evolves appreciably during acceleration. In addition, their large inertia, orders of magnitude greater than that of electrons, implies that the initial increase in β is slow compared to electrons, causing the velocity to increase slowly and making them particularly sensitive to phase slippage.

In conventional laser-driven wakefield acceleration, the wake phase velocity is ($v_\phi \simeq c$), so a sub-relativistic heavy witness quickly slips with respect to the accelerating phase and dephases before gaining substantial energy. In this Section, we quantify this limitation and then explore an alternative route enabled by velocity-controlled spatio-temporal (ST) laser drivers, which utilize a ‘flying focus’ to decouple the intensity peak’s velocity from the group velocity in the medium (as detailed in Chapter 2). By tuning the focal velocity of the driver, we

can reduce the phase velocity of the wakefield to match the initial witness velocity and enable prompt capture. We first demonstrate the dephasing/trapping limitation of standard LWFA for muons in PIC simulations (Sec. 3.2.1). We then introduce the subluminal-driver concept and derive the corresponding dephasing constraint in a uniform plasma (Sec. 3.2.2). This represents a fundamental shift from the conventional dephasing regime: whereas in standard LWFA the wake outruns the witness, here it is the witness that eventually outruns the wake, imposing a new upper limit on the achievable energy gain.

Finally, we examine how the energy gain depends on the chosen focal velocity (Sec. 3.2.3), identifying the trade-off between initial trapping and the onset of overtake. These results motivate the need for dynamic phase-velocity control during acceleration, which we address in Sec. 3.3 using tailored plasma-density profiles to maintain extended phase-locking and achieve fully relativistic energies in a single acceleration stage.

3.2.1 MUON ACCELERATION VIA STANDARD LASER WAKEFIELDS

While the ultra-high gradients of plasma wakefields are ideal for minimizing decay, standard laser-driven (LWFA) schemes fail to efficiently trap sub-relativistic secondaries. In a conventional LWFA setup, the wakefield phase velocity is intrinsically tied to the speed of the driver, so for a laser driver this corresponds to $v_\phi \approx c$. Because realistic secondary sources typically produce muons with $v_\mu \leq 0.95c$, the resulting velocity mismatch with a luminal wake makes sustained acceleration highly inefficient.

To illustrate this limitation, we performed a particle-in-cell (PIC) simulation using the code OSIRIS. We modeled the interaction of a muon bunch, initialized with $v_\mu = 0.9c$, with a standard laser-driver wakefield. The simulation was performed in a quasi-3D cylindrical geometry, with parameters summarized in Table 3.2. A laser pulse driver propagates through a parabolic plasma channel, while a muon beam is injected ahead of the driver. The plasma profile features a linear density ramp for $z < 4c/\omega_p$ and transverse parabolic confinement for guided laser propagation. We observe an initial energy gain for the muon bunch, followed by rapid dephasing and loss of trapping: the quasi-luminal wake ($v_\phi \approx c$) outruns the beam. The large muon inertia ($m_\mu/m_e \simeq 206.7$) limits the rate of velocity increase, so that the bunch exits the accelerating region before gaining sufficient momentum to remain phase-locked with the wakefield (Figure 3.1).

In particular, at early times, Figure 3.1 (a, d) shows clearly how the bunch sits in the accelerating phase of the wake, where the electric field is negative. By $t = 60.06 [1/\omega_p]$ (Fig. 3.1 (b,e)), the particles' longitudinal momentum reaches its overall maximum $p_{1,\max} \approx 2.21 m_\mu c$. However, since the driver's velocity is much faster than the witness beam's velocity, the continuous phase slippage caused by the velocity mismatch pushes the bunch out of the acceleration bucket. In Figure 3.1 (c,f), we show how the bunch is no longer trapped in the first accelerating bucket.

These limitations motivate the use of *subluminal drivers* (Section 3.2.2) that generate wakes that match the initial muon velocity, allowing better trapping of particles with large inertia and non-relativistic initial velocities.

TABLE 3.2: Simulation set-up for muon acceleration using standard laser-driven wakefield acceleration.

	Parameter	Value
Simulation	Algorithm	Quasi-3D
	Grid ($z \times r$) [c/ω_p]	12000×400
	Domain ($z \times r$) [c/ω_p]	$[-120, 15] \times [0, 10]$
	Time step	$7 \times 10^{-3} \omega_p^{-1}$
Laser	a_0	2.0
	ω_0/ω_p	20
	w_0	$2.9 c/\omega_p$
Plasma	n_0/n_c	$(\omega_0/\omega_p)^{-2}$
	Parabolic Plasma Profile	$n(r) = n_0 \left(1 - \frac{r^2}{w^2}\right)$
Muons	m_μ/m_e	206.7
	γ_μ	2.06
	n_μ/n_0	10^{-3}
	$\sigma_z \times \sigma_r$	$4 \times 0.5 c/\omega_p$

3.2.2 SUBLUMINAL-DRIVER PLASMA WAKEFIELD ACCELERATION

As discussed in the previous subsection, using a conventional laser driver for plasma wakefield acceleration presents a key limitation when accelerating particles like muons, pions and kaons: the phase velocity of the wakefield is too high for efficient initial trapping. These particles, typically injected with subrelativistic velocities, fail to remain in phase with the accelerating field long enough to gain significant energy.

For this reason, we investigate the use of a subluminal driver. Although such drivers can successfully trap sub-relativistic beams, their effectiveness is ultimately constrained by phase slippage between the witness and the wake (dephasing), which limits the effective interaction length [76]. We consider a driver moving with a subluminal focal velocity $v_f < c$ and a mono-energetic witness beam moving at a velocity $v_b = v_f$ at injection. For a wake moving at v_f , the spatial plasma period in the laboratory frame scales as $\lambda_p^f = 2\pi v_f/\omega_p = (v_f/c)\lambda_p$. In other words, if $\lambda_p = 2\pi c/\omega_p$ is the nominal plasma wavelength, then for a wake moving at v_f the relevant half-wavelength (i.e. the approximate longitudinal extent of one accelerating bucket) becomes

$$\ell_w \simeq \frac{\lambda_p}{2} \frac{v_f}{c}. \quad (3.5)$$

The wake phase velocity v_ϕ in the nonlinear regime can be approximated by the laser group velocity with a small nonlinear correction, $v_\phi \simeq v_g \left(1 - \frac{3\omega_p^2}{4\omega_0^2}\right)$, where we adopt an effective group-velocity scaling for a driver propagating with focal velocity v_f , namely [77] $v_g \simeq$




FIGURE 3.1: Time evolution of the laser-driven wakefield acceleration of a monoenergetic muon beam with initial velocity of $0.9c$. Panels (a–c) show snapshots at (a) $t = 32.03 \omega_p^{-1}$, (b) $t = 62.06 \omega_p^{-1}$, and (c) $116.12 \omega_p^{-1}$. The grayscale colormap represents the plasma density, in green we have the muon beam density, and the red line marks the longitudinal accelerating electric field on-axis. Panels (d–f) show the evolution of the beam's maximum longitudinal momentum p_1 as a function of propagation coordinate x_1 , in units of c/ω_p . A pattern of repeated energy peaks is observed in correspondence with the peak intensity of the wakefield electric field, with a maximum momentum of $p_{1,\text{max}} \approx 2.21 m_\mu c$ reached in panel (e). For reference, a background density of $n_0 = 10^{18} \text{ cm}^{-3}$ corresponds to $1/\omega_p \simeq 17.7 \text{ fs}$ and $c/\omega_p \simeq 5.3 \mu\text{m}$.

$v_f \sqrt{1 - \frac{\omega_p^2}{\omega_0^2}}$. Then, following [77], the dephasing length, $L_d \simeq \frac{v_b}{|v_b - v_\phi|} \ell_w$, now evaluates to

$$L_d \simeq \frac{2}{3} \lambda_p \left(\frac{\omega_0}{\omega_p} \right)^2 \frac{v_f}{c}, \quad (3.6)$$

where we assumed $v_b = v_f$. The result is reduced to the standard $L_d \simeq \frac{2}{3} \lambda_p \left(\frac{\omega_0}{\omega_p} \right)^2$ when $v_f \sim c$. Consequently, the maximum energy gain for a mono-energetic witness co-propagating over the phase-slippage length is

$$\frac{\Delta W_{\text{max}}}{(mc^2)} = L_d \langle E_{\text{accel}} \rangle \simeq \frac{2}{3} \lambda_p \left(\frac{\omega_0}{\omega_p} \right)^2 \frac{v_f}{c} \langle E_{\text{accel}} \rangle. \quad (3.7)$$

So, if, for example, we want to accelerate particles using a laser with $\omega_0/\omega_p = 20$, a laser wavelength of $\lambda = 800 \text{ nm}$, and an average accelerating field of $\langle E_{\text{accel}} \rangle = 50 \text{ GV/m}$, our energy gain will be only around tens of MeV.

We performed a quasi-3D OSIRIS simulation to validate this estimate. The parameters

TABLE 3.3: Simulation set-up for muon acceleration using standard spatio-temporal pulses driven wake-field acceleration.

	Parameter	Value
Simulation	Algorithm	Quasi-3D (cylindrical)
Grid	$(z \times r) [c/\omega_p]$	27000×400
	Domain $(z \times r) [c/\omega_p]$	$[-120, 15] \times [0, 20]$
	Cell size $(\Delta z \times \Delta r) [c/\omega_p]$	0.005×0.05
Muons	m_μ/m_e	206.7
	Initial γ_μ	2.06
	Peak density	$n_\mu/n_0 = 10^{-3}$
	Size $(\sigma_z \times \sigma_r) [c/\omega_p]$	1.0×0.5
ST laser driver	a_0	2
	ω_0/ω_p	20
	Focal velocity	$v_f = 0.9c$
	Transverse profile	Gaussian, $w_0 = 1 c/\omega_p$

for the simulation are summarized in Table 3.3 Figure 3.2 illustrates a simulation of the acceleration of a monoenergetic muon beam with initial velocity $v_b = 0.9c$ driven by a spatio-temporal laser pulse with focal velocity $v_f = 0.9c$. Panels (a–c) show snapshots of the plasma and beam dynamics at different times. The grayscale colormap represents the plasma density, the green regions show the muon beam density, and the red curve the accelerating electric field. Panels (d–f) present the beam’s longitudinal momentum p_1 as a function of the propagation coordinate x_1 again at different times. The muon bunch experiences repeated energy peaks aligned with the structure of the wakefield, reaching a maximum momentum of $p_{1,\max} \approx 2.3 m_\mu c$ in panel (e). For reference, a background density of $n_0 = 10^{18} \text{ cm}^{-3}$ corresponds to $1/\omega_p \simeq 17.7 \text{ fs}$ and $c/\omega_p \simeq 5.3 \mu\text{m}$.

While this configuration enables initial trapping, it is still ultimately limited by dephasing. As the beam gains energy, its velocity $v_b(t)$ eventually exceeds the phase velocity of the wake ($v_\phi \approx v_f$), causing the bunch to outrun the accelerating region. The beam moves faster than the driver, eventually exiting the high-field region and imposing a limitation on the achievable energy gain in uniform plasmas using subluminal drivers. Thus, we are still not satisfying the condition of bringing our particles from subrelativistic to relativistic velocities to mitigate decay losses (Sec. 3.1).

We can conclude that employing a subluminal driver improves the acceleration outcome compared with the standard laser-driven case discussed in Sec. 3.2.1, enabling the witness to reach higher energies and velocities before phase slippage terminates efficient acceleration. Nevertheless, a fixed subluminal phase velocity still imposes an intrinsic limit on the attainable energy gain: as the witness accelerates and approaches (or exceeds) the wake phase velocity, it inevitably slips out of the optimal accelerating phase. These results motivate a more advanced strategy in which the wake phase velocity is dynamically controlled during propagation, as




FIGURE 3.2: Time evolution of wakefield acceleration of a monoenergetic muon beam with initial velocity of $0.9c$ using a spatio-temporal laser driver with a focal velocity of $v_f = 0.9c$. Panels (a–c) show snapshots at (a) $t = 41.97 \omega_p^{-1}$, (b) $t = 159.88 \omega_p^{-1}$, and (c) $261.80 \omega_p^{-1}$. The grayscale colormap represents the plasma density, in green we have the muon beam density, and the red line marks the longitudinal accelerating electric field. Panels (d–f) show the evolution of the beam's maximum longitudinal momentum p_1 as a function of propagation coordinate x_1 , in units of c/ω_p . A pattern of repeated energy peaks is observed in correspondence of the peak intensity of the wakefield electric field, with a maximum momentum of $p_{1,\max} \approx 2.3 m_\mu c$ reached in panel (e). For reference, a background density of $n_0 = 10^{18} \text{ cm}^{-3}$ corresponds to $1/\omega_p \simeq 17.7 \text{ fs}$ and $c/\omega_p \simeq 5.3 \mu\text{m}$.

will be shown in Sec. 3.3. In the next section we will briefly discuss how this result can be improved by varying the driver's focal velocity.

3.2.3 SUBLUMINAL DRIVER WITH DIFFERENT VELOCITIES

In the previous section, we demonstrated that using a subluminal driver with a focal velocity matched to the initial muon velocity enables efficient trapping and acceleration of sub-relativistic muon beams, but the energy gain remains limited by rapid dephasing as the particles catch up and overtake the wakefield.

We now extend this analysis by considering drivers with focal velocities slightly above the initial beam velocity ($v_f > v_0$). Figure 3.3 compares the maximum longitudinal momentum reached by muon beams starting at $v_0 = 0.9c$ ($p_1 = 2.06 m_\mu c$) for three different focal velocities: $v_f = 0.91c$, $0.92c$, and $0.93c$. The simulation set-up is identical to the one presented in Sec. 3.2.2 apart from the focal velocity of the driver. The results show that a modest increase




FIGURE 3.3: Maximum energy gain for a beam of muons (initial velocity $v_0 = 0.9c$, corresponding to normalized momentum $p_1 = 2.06 m_\mu c$), accelerated by spatiotemporal laser pulses with different focal velocities (v_f). The plot compares the evolution of the longitudinal momentum (p_1) for three cases: $v_f = 0.91c$ (green), $v_f = 0.92c$ (pink), and $v_f = 0.93c$ (blue). The case with $v_f = 0.91c$ achieves the highest final momentum ($p_1^{\max} \approx 3.1 m_\mu c$), as this intermediate focal velocity balances initial trapping with delayed phase slippage, maximizing the effective acceleration length. For higher driver velocities ($v_f = 0.92c$ and $0.93c$), the particles dephase more rapidly and cannot stay in phase with the accelerating field, resulting in lower maximum energy gains.

in the driver velocity can significantly enhance the achievable energy gain. In particular, the case with $v_f = 0.91c$ yields the highest final momentum ($p_1^{\max} \approx 3.1 m_\mu c$), as the wake remains synchronous with the beam for a longer distance, allowing continued acceleration. For higher velocities ($v_f = 0.92c$ and $0.93c$), the wake outruns the beam, leading to premature dephasing and reduced energy gain. This indicates a critical trade-off in driver selection. While setting a focal velocity v_f slightly above the beam's initial velocity delays the point at which the particles overtake the wake—thereby extending the interaction length—an excessively large v_f causes the wake to outrun the particles entirely, leading to a loss of trapping. Effectively, one must balance the delay of dephasing against the risk of immediate de-trapping.

However, even in the optimal case ($v_f = 0.91c$), the maximum velocity obtained by the muons is only about $0.95c$, which remains significantly below the fully relativistic ($v \approx c$) regime we want to achieve. This limitation highlights the need for more advanced strategies to further extend the acceleration length and increase the final particle energy. In the next section, we present an approach based on a dynamically tailored plasma density profile, which enables phase-locking between the beam and the wake and allows the particles to reach higher energies.

3.3 TAILORED PLASMA DENSITY PROFILE FOR EXTENDED PHASE-LOCK

As introduced in the previous section, effective acceleration of subluminal, heavy particles to relativistic velocities requires not only initial trapping but also a significant extension of the dephasing length. To achieve this, we employ a tailored plasma-density profile that dynamically controls the phase velocity of the accelerating wake, thereby maintaining a continuous phase-lock with the particles as they accelerate. By coupling the evolution of the particle's longitudinal momentum with the variation of the plasma frequency caused by the change in the density profile, we derive the exact density profile needed to keep the beam within the accelerating region of the wake. The physical mechanism underlying this control can be understood directly from Eq. 3.12. The phase velocity of the accelerating bucket depends not only on the focal velocity v_f of the driver, but also on the local density gradient dn/dz . A decreasing density profile shifts v_w downward, slowing the wake to match a still sub-relativistic witness, while a subsequent up-ramp raises v_w again as the particle approaches relativistic velocities. The density profile therefore acts as a continuous tuning knob for the wake phase velocity, allowing it to follow the evolving particle velocity throughout the entire acceleration stage.

To derive the conditions required for phase-locking, we begin by expressing the relevant physical quantities and equations governing particle acceleration in a varying-density plasma. The ideal plasma-density profile that preserves phase synchronism between the accelerated particles and the longitudinal wakefield can be obtained by requiring the instantaneous particle velocity $v(z)$ to coincide with the phase velocity $v_w(z)$ of the n_w -th accelerating bucket throughout the interaction length.

The on-axis accelerating field is proportional to the square root of the plasma density, and it can be written as

$$E_z(z) = \alpha_0 \frac{m_e \omega_{p0} c}{e} \sqrt{\frac{n(z)}{n_0}}, \quad (3.8)$$

where α_0 is the normalized field amplitude. Here n_0 and ω_{p0} denote the reference density and plasma frequency respectively at $z = 0$. The longitudinal momentum of a test particle of charge q and mass m_b evolves according to

$$\frac{dp_z}{dt} = m_b \frac{d[\gamma(z)v(z)]}{dt} = q E_z(z). \quad (3.9)$$

The spatio-temporal driver's peak intensity that is responsible for the wakefield excitation moves at v_f . Using the propagation distance of the driver peak $z = v_f t$ as an independent variable, we replace the time derivative by a space derivative, obtaining

$$\frac{d}{dz} [\gamma(z)v(z)] = \alpha_0 \frac{|q|}{e} \frac{m_e}{m_b} \frac{\omega_{p0} c}{v_f} \sqrt{\frac{n(z)}{n_0}}. \quad (3.10)$$

Equation (3.10) constitutes the first element of the coupled system that will allow us to determine the ideal density profile for sustained phase synchronism.

The second equation can be derived by noting that each accelerating bucket has an approximate size of $\sim \lambda_p$, the axial position of the center of the n_w -th accelerating cavity with respect to the driver peak is

$$z_w(t) = v_f t - n_w \lambda_p [z(t)], \quad \lambda_p(z) = \frac{2\pi v_f}{\omega_p(z)}, \quad (3.11)$$

with $\omega_p(z) = \sqrt{n(z)e^2/(m_e \epsilon_0)}$. Differentiation with respect to time gives the phase velocity of the wake

$$\begin{aligned} v_w &= \frac{dz_w}{dt} = v_f - n_w v_f \frac{d\lambda_p}{dz} \\ &= v_f - \frac{2\pi n_w v_f^2}{\omega_{p0}} \frac{d}{dz} \sqrt{\frac{n_0}{n(z)}}. \end{aligned} \quad (3.12)$$

To obtain the ideal density profile that keeps our particles phase-locked to the ideal acceleration position, we impose $v_w = v$ when solving Eqs. (3.10) and (3.12). It is important to note that this model has a high degree of generality: the same approach can be adapted to a wide range of driver configurations, particle species, and initial conditions, making it a powerful tool for designing advanced plasma-based acceleration schemes.

3.4 MODEL VALIDATION AND VERSATILITY

The theoretical framework developed in Sec. 3.3 provides a general recipe for phase-locked acceleration of massive, sub-relativistic particles using tailored plasma-density profiles. A key advantage of this model lies in its universality: it can be directly applied to particles with different masses and charges by adjusting the relevant parameters in the governing equations. To evaluate its predictive accuracy and demonstrate this versatility, we perform a systematic investigation of how the acceleration dynamics is influenced by intrinsic particle properties, the selected accelerating bucket (n_w), and the initial injection velocity. Specifically, we compare theoretical predictions for energy gain against OSIRIS simulations for different particle masses (Sec. 3.4.1), injection into secondary accelerating buckets (Sec. 3.4.2), and varying initial velocities (Sec. 3.4.3).

3.4.1 ACCELERATION OF DIFFERENT PARTICLES

To demonstrate the model's adaptability to the acceleration of different particles, we first compute the tailored density profiles required for species with different rest masses while maintaining a constant initial Lorentz factor of $\gamma_i = 2.9$. We consider particles with masses corresponding to the pion ($m_\pi \approx 139.6 \text{ MeV}/c^2$), the muon ($m_\mu \approx 105.7 \text{ MeV}/c^2$), and two hypothetical 'reduced-mass' muons ($m_\mu/2$ and $m_\mu/4$).

The resulting profiles, shown in Fig. 3.4(a), reveal that heavier species require a more gradual density transition to accommodate their slower velocity evolution. By contrast, the lighter reduced-mass cases demand steeper gradients as they approach the relativistic regime more rapidly.




FIGURE 3.4: (a) Tailored plasma density profiles $n(z)/n_0$ computed for different particles using the phase-locking model described in Section 3.3. Each profile maintains synchronism between the accelerating wake and particles of varying mass: pion mass m_π (solid), muon mass m_μ (dashed), and reduced muon masses $m_\mu/2$ (dash-dotted) and $m_\mu/4$ (dotted). (b) Maximum Lorentz factor $\gamma_{\max}$ achieved as a function of particle mass. The black curve shows theoretical predictions based on the tailored profiles, while red squares represent OSIRIS simulation results. All cases assume initial injection at $\gamma_i = 2.9$ (gray dashed line). The agreement confirms the predictive accuracy and versatility of the model.

Figure 3.4(b) compares the maximum Lorentz factor $\gamma_{\max}$ predicted by the model with values obtained from OSIRIS simulations. The excellent agreement confirms the predictive accuracy of the framework across a wide range of particle masses. Furthermore, the clear scaling observed suggests that the model is not limited to the species tested here, but can be reliably extrapolated to even heavier particles, provided the density profile is adjusted to match their specific inertia.

3.4.2 ACCELERATION USING SECONDARY ACCELERATION BUCKETS

The phase-locking model introduced in the previous sections also provides valuable insight into the role of the accelerating bucket used for injection. In a typical wakefield structure, multiple accelerating buckets are formed behind the driver, each with slightly different field strengths and spatial characteristics. Although the first bucket, right behind the laser, ($n_w = 1$) is often the most commonly used due to its strongest fields, the model predicts that efficient acceleration is still possible in secondary buckets, provided phase synchronism is maintained.

Figure 3.5(a) shows the wakefield structure generated by a spatiotemporal laser pulse propagating through plasma with a focal velocity of $v_f = 0.97c$. A sequence of accelerating buckets is excited behind the laser, with the first seven labeled. The on-axis accelerating field E_z (black line) is shown alongside the transverse and longitudinal field components. Particles injected into any of these buckets can be phase-locked and accelerated by tailoring the plasma density profile accordingly.

Figure 3.5(b) shows examples of ideal-fit density profiles for the first accelerating buckets for the acceleration of a monoenergetic beam of particles with $m_\mu/4$ mass and initial velocity of $v_{\mu/4} = 0.9c$. As illustrated in Fig. 3.5(b), phase-locking in secondary buckets requires less extreme density reductions compared to the first bucket. This implies less degradation of the longitudinal accelerating field during the acceleration process.




FIGURE 3.5: (a) Wakefield structure generated by a spatiotemporal laser pulse propagating in plasma with a focal velocity $v_f = 0.97c$. The grayscale colormap shows the plasma density response to the laser's propagation n/n_0 , while the blue-orange and red-blue colormaps correspond to the transverse field component $E_x - B_y$ and the longitudinal electric field E_L , respectively. The laser pulse enters from the right and excites a train of accelerating buckets, with the first seven labeled. The black line shows the on-axis accelerating field E_z . (b) Ideal density profile for the acceleration of a monoenergetic particle beam with mass $m_\mu/4$ for the first three accelerating buckets.

As shown in Figure 3.6, the maximum Lorentz factor $\gamma_{\max}$ increases with the bucket index n_w . Even though secondary buckets have a lower peak accelerating field compared to the first one, they require a less pronounced initial density reduction to maintain phase locking. As a result, the accelerating field is better preserved throughout the stage, leading to a higher final energy gain. The electric field is overall then more conserved. This means that the theoretical predictions (black squares) are in excellent agreement with OSIRIS simulations (red circles), confirming the model's ability to accurately capture bucket-dependent dynamics. All cases correspond to particles injected at $\gamma_i = 2.29$.




FIGURE 3.6: Maximum Lorentz factor $\gamma_{\max}$ reached as a function of the accelerating bucket used for the acceleration n_w . Theoretical predictions (black squares) are compared with OSIRIS simulations (red circles), showing excellent agreement. All cases correspond to initial injection at $\gamma_i = 2.29$ (gray dashed line).

These results demonstrate the flexibility of the phase-locking model, which can be applied to optimize acceleration not only across particle species and initial energies, but also across

different regions of the wakefield structure. This capability is particularly relevant for experimental configurations where injection into the first bucket is not feasible, for example due to the presence of strong plasma electrons self-injection in the first accelerating bucket.

3.4.3 ACCELERATION FROM DIFFERENT INITIAL VELOCITIES

A key advantage of the phase-locking acceleration framework is its flexibility with respect to the initial velocity of the injected particles. In realistic scenarios, the particles we are interested in for this study (muons, pions and kaons) are produced with a broad range of initial velocities, depending on the production mechanism and the characteristics of the source. The most fruitful production mechanisms though are the ones where the particles are produced at subrelativistic velocities [78, 79]. Therefore, it is crucial for any practical acceleration scheme to be adaptable to varying initial conditions.

To investigate this aspect, we analyze the acceleration of three distinct beams with mass $m_{\mu/4}$, each injected with a different initial velocity ($v_0/c = 0.5, 0.7, 0.9$), and matched to a tailored plasma density profile driven by a spatiotemporal laser pulse with a focal velocity set just above the respective injection velocity ($v_f = 0.8c, 0.9c, 0.97c$). Figure 3.7 shows the evolution of the normalized velocity v/c for each case, comparing the predictions of the phase-locking model (dashed lines) to fully kinetic OSIRIS simulations (solid markers). Dotted horizontal lines indicate the initial velocities. To obtain stable wakefield excitation while varying the focal velocity of the driver we followed the guidelines presented in Chapter 2.




FIGURE 3.7: Maximum velocity evolution v/c of three beams starting at $v_0/c = 0.5$, $v_0/c = 0.7$, and $v_0/c = 0.9$. Each beam is driven with focus velocities of $0.8c$, $0.9c$, and $0.97c$, respectively. The dashed lines represent the theoretical velocity evolution, and the solid markers are from OSIRIS simulations.

The excellent agreement between the theoretical curves and simulation results demonstrates the robustness of the model over a broad range of injection energies. Even for low initial velocities, the tailored density profile ensures sustained phase-lock and efficient energy transfer from the wakefield to the beam, with the accelerated particles rapidly gaining speed and approaching relativistic regimes. These results confirm that the presented phase-locking strategy is not limited to a single injection condition, but can be flexibly adapted to accommodate the natural spread of velocities produced in realistic sources. This versatility further highlights the applicability of the approach for advanced plasma-based accelerators targeting short-lived particles across a range of experimental conditions.

3.4.4 ESTIMATE OF THE END OF PHASE-LOCKED ACCELERATION

The density-tailoring strategy derived in Sec. 3.3 is designed to maintain phase-locking between the accelerating bucket and a heavy, initially sub-relativistic witness. However, because the spatio-temporal driver peak propagates at a prescribed subluminal focal velocity $v_f < c$, the acceleration distance is fundamentally finite. As the witness reaches the ultra-relativistic regime, its velocity $v(t) \rightarrow c$, causing it to eventually catch and overtake the driver peak. This "overtake limit" defines the maximum achievable energy in a single stage for a given v_f .

A critical objective of this framework is to predict the maximum effective acceleration time, which directly determines the peak energy gain. A convenient way to define the end of the useful acceleration window is by monitoring the longitudinal separation between the witness and the driver peak:

$$\Delta x_1(t) \equiv x(t) - x_f(t), \quad (3.13)$$

where $x(t)$ denotes a representative beam position (e.g., the witness centroid) and $x_f(t)$ is the position of the driver intensity peak. For spatio-temporal pulses, the peak moves at a constant focal velocity v_f , leading to:

$$x_f(t) = x_f(0) + v_f t. \quad (3.14)$$

The overtake time t_{over} is defined as the moment this separation vanishes ($\Delta x_1 = 0$). Predicting t_{over} allows for the precise selection of v_f to maximize the interaction length and, consequently, the final energy of the beam. The maximum acceleration time t_{over} can be extracted theoretically by solving the reduced model derived in Sec. 3.3. For a driver with a constant focal velocity v_f , we first determine the witness velocity evolution $\beta(t)$ by integrating the reduced ordinary differential equation. The witness trajectory is then reconstructed as

$$x(t) = x(0) + c \int_0^t \beta(t') dt', \quad (3.15)$$

By solving for the time when the witness position $x(t)$ matches the driver peak position $x_f(t)$, we can determine the exact duration of the effective acceleration window. This analytical




FIGURE 3.8: Estimate of the maximum acceleration time. The longitudinal separation $\Delta x_1(t) = x(t) - x_f(t)$ is computed by comparing the witness trajectory $x(t)$ (reconstructed via Eq.(3.15)) with the prescribed driver peak $x_f(t) = x_f(0) + v_f t$. The overtake time t_{over} (red marker) denotes the zero-crossing where the witness outruns the driver, marking the termination of the phase-locked accelerating stage. Calculations assume a reference density $n_0 = 4 \times 10^{18} \text{ cm}^{-3}$ and a laser frequency ratio $\omega_0/\omega_p = 20$.

approach enables us to optimize the initial offset and driver velocity v_f to ensure the beam remains within the accelerating region for the longest possible duration.

Figure 3.8 illustrates the numerical reconstruction of the witness and driver trajectories to determine the overtake time for a representative case with $v_f = 0.97c$ and $\beta_0 = 0.9$. In this configuration, an initial longitudinal offset of $x_f(0) - x(0) = 8.0 \mu\text{m}$ is prescribed to define the starting phase within the accelerating bucket. As the witness accelerates, the gap $\Delta x_1(t)$ narrows until the trajectories intersect at t_{over} (red marker). At this point, the witness physically overtakes the driver peak, exiting the accelerating phase and entering the region of phase-deceleration. This crossover marks the absolute physical limit of energy gain for a single-stage subluminal driver in a uniform plasma.

3.5 CONCLUSIONS

In this Chapter, we addressed how to bring short-lived, sub-relativistic heavy particles into the ultra-relativistic regime within a single, compact accelerating stage, fast enough to outpace their decay. Starting from a quantitative analysis of in-flight losses (Sec. 3.1), we showed that the GV/cm-class gradients accessible to plasma wakefields can reduce decay losses by orders of magnitude with respect to conventional RF linacs (Table 3.1). This advantage, however, can only be realized if the accelerating structure is matched to the kinematics of the witness beam from injection up to the relativistic limit.

We then demonstrated, both analytically and through self-consistent OSIRIS particle-in-cell simulations, that conventional laser-wakefield acceleration cannot meet this requirement efficiently. A quasi-luminal wake ($v_\phi \simeq c$) outruns any sub-relativistic, heavy witness on a timescale much shorter than that needed to gain substantial energy, leading to rapid dephasing and detrapping (Sec. 3.2.1). This limitation is a direct consequence of the velocity mismatch between a luminal wake and a witness with large inertia.

To circumvent this mismatch, we introduced velocity-controlled spatiotemporal laser drivers. By exploiting their flying focus to set the wake phase velocity through the focal velocity v_f , we showed that a witness can be promptly trapped when $v_f \simeq v_b$. Yet, in a uniform plasma, this strategy is itself fundamentally limited by a dephasing process: as the witness accelerates, it eventually overtakes the wake (Sec. 3.2). The associated dephasing length, $L_d \simeq (2/3) \lambda_p (\omega_0/\omega_p)^2 (v_f/c)$, and the corresponding maximum energy gain set a limit that cannot be overcome by simply tuning v_f slightly above the injection velocity, as further illustrated in Sec. 3.2.3.

To further improve our acceleration, we proposed a method to calculate the ideal plasma density profile to keep the particle beam phase-locked (Sec. 3.3). By coupling the longitudinal momentum equation to the local plasma frequency, we derived the exact tailored density profile $n(z)$ that enforces the synchronism condition $v_w(z) = v(z)$ throughout the entire interaction. The density profile thus acts as a continuous tuning mechanism to control the wake phase velocity: it slows the wake when the witness is sub-relativistic and speeds it up again as the witness approaches c . Crucially, the framework is formulated in terms that explicitly depend on the witness mass, charge, injection energy, and the chosen wake bucket, making it

directly transferable across different physical scenarios.

The model has been systematically validated through a series of OSIRIS simulations. First, we showed that the predicted density profiles drive sustained acceleration of particles with masses spanning from $m_\mu/4$ up to the pion mass, with the maximum Lorentz factor following the analytic prediction across nearly an order of magnitude in m_b (Sec. 3.4.1). Second, we demonstrated that phase-locking is not restricted to the first wake bucket: injection into secondary buckets requires milder density gradients and, as a consequence, preserves the accelerating field more effectively, leading to higher final energies (Sec. 3.4.2). This feature is particularly relevant for experimental scenarios in which the first bucket is contaminated by self-injected plasma electrons. Third, we verified the robustness of the scheme against a broad range of injection velocities: witnesses launched at $v_0/c = 0.5, 0.7, 0.9$ all reach near-luminal velocities while closely following the analytic predictions (Sec. 3.4.3). Together, these results establish the phase-locking framework as a versatile and predictive design tool for plasma-based acceleration of decaying particles.

Finally, we identified the absolute physical limit of this single-stage scheme. Even under perfect phase-locking, a witness driven to $v \rightarrow c$ must eventually overtake a driver moving at a finite focal velocity $v_f < c$. By comparing the analytic witness trajectory with the prescribed driver trajectory, we obtained an explicit estimate of the overtake time t_{over} and, equivalently, of the maximum useful acceleration length (Sec. 3.4.4). This sets a clear design target: v_f and the initial witness–driver offset must be jointly optimized so as to maximize the integrated energy gain while keeping the witness within the accelerating phase up to t_{over} .

Taken together, the results of this chapter provide a complete, predictive, theoretical, and numerical framework for the single-stage acceleration of decaying, sub-relativistic particles. The combination of velocity-controlled spatio-temporal drivers with tailored plasma density profiles enables efficient trapping, sustained phase synchronism, and final energies that are not accessible with either ingredient alone. Several questions remain open on the path to a full experimental implementation, most notably the realization of the required density profiles in the laboratory, the systematic inclusion of beam-loading and transverse emittance preservation in the optimized regime, and the coupling of the scheme to realistic muon, pion, and kaon sources. These aspects, together with possible extensions to multi-stage architectures and to alternative driver concepts, are natural avenues for future work and motivate the complementary proton-driven approach that will be discussed in the following chapters.

Take-home messages

- **High accelerating gradients are necessary to avoid decay losses.** For an average accelerating gradient E_{acc} the surviving fraction is

$$S \simeq \left(\frac{\gamma_i}{\gamma_f} \right)^{v_p}, \quad v_p \equiv \frac{m_b c}{q E_{\text{acc}} \tau_p}.$$

- **Standard LWFA is too fast to trap sub-relativistic particles.** With a quasi-luminal wake ($v_\phi \simeq c$), a beam injected with $\beta_0 < 1$ rapidly slips out of the accelerating phase before gaining substantial energy.
- **Subluminal ST drivers trap particles, but are very limited by dephasing.** For a wake moving at the focal velocity $v_f < c$, the accelerating-bucket length and dephasing length scale as

$$\ell_w \simeq \frac{\lambda_p}{2} \frac{v_f}{c}, \quad L_d \simeq \frac{2}{3} \lambda_p \left(\frac{\omega_0}{\omega_p} \right)^2 \frac{v_f}{c} \quad (v_b = v_f \text{ at injection}),$$

and acceleration terminates when the witness reaches $v_b > v_f$ and overtakes the wake.

- **Single-stage energy gain requires phase-locking via a tailored density profile.** Use $n(z)$ to make the wake velocity follow the particle velocity by solving the coupled system

$$\frac{d}{dz}[\gamma(z)v(z)] = \alpha_0 \frac{|q|}{e} \frac{m_e}{m_b} \frac{\omega_{p0} c}{v_f} \sqrt{\frac{n(z)}{n_0}}, \quad v_w(z) = v_f - \frac{2\pi n_w v_f^2}{\omega_{p0}} \frac{d}{dz} \sqrt{\frac{n_0}{n(z)}},$$

with the *locking condition* $v_w(z) = v(z)$. The model was tested in PIC simulations across different β_0 (sub- to near-relativistic velocities), injection buckets n_w , and different particles' masses. The model proved to be *versatile* and predictive.

- **The time the witness overtaking the driver defines the maximum acceleration time.** Acceleration cannot extend beyond the time when the witness reaches the driver intensity maximum:

$$x(t) = x_f(t), \quad x(t) = x(0) + c \int_0^t \beta(t') dt', \quad x_f(t) = x_f(0) + \int_0^t v_f(t') dt',$$

(for constant v_f : $x(0) - x_f(0) + c \int_0^t \beta dt' - v_f t = 0$). This sets the *maximum* acceleration time even under perfect phase-locking.

CHAPTER 4

BEAM QUALITY AND EXPERIMENTAL FEASIBILITY

In the previous chapter, we established a robust framework for accelerating short-lived, sub-relativistic particles using ST drivers. We demonstrated that by dynamically tailoring the plasma density profile to match the accelerating beam's inertia, it is possible to maintain longitudinal phase-locking and achieve relativistic energies before significant decay occurs. However, the phase-locking mechanism was validated using low-density witness beams to isolate single-particle dynamics from collective effects. Although this established the theoretical validity of the subluminal trapping concept, for realistic applications, the accelerator must handle significant charge that loads the wakefield. Therefore, this chapter shifts the focus from longitudinal synchronism to the critical issues of beam quality preservation and beam loading and beam quality preservation at high charge.

We begin by investigating the limits of beam loading in our subluminal acceleration scheme (Sec . 4.1). When a witness bunch with a substantial charge is introduced, it drives its own wakefield that superimposes on the driver field and modifies the wake structure. We characterize this effect and identify the optimal loading regime, which flattens the on-axis accelerating field and minimizes the energy spread. We then examine the transverse dynamics of the witness beam. (Sec. 4.2). The varying density $n(z)$ required for longitudinal phase locking inevitably modulates the ion channel focusing force along the propagation direction. We analytically demonstrate that because short-lived particles possess a mass much larger than that of an electron ($m_p \gg m_e$), their betatron wavelength λ_β is lengthened. This makes the beam significantly less sensitive to density variations. Furthermore, PIC simulations confirm that the normalized emittance $\bar{\epsilon}_n$ is preserved under optimal loading, whereas overloading triggers wake disruption and irreversible emittance growth.

A dedicated transverse trapping analysis for muons, pions, and kaons (Sec . 4.3) confirms that these species remain stably confined within the accelerating bubble throughout the entire interaction, owing to their large inertia. Finally, we discuss the experimental feasibility of tailored density profiles (Sec . 4.4) and propose two practical realizations, validated through PIC simulations.

4.1 BEAM LOADING EFFECTS

The introduction of a witness bunch with a substantial charge modifies the wakefield structure, a phenomenon known as beam loading [80]. This effect arises because the bunch induces its own wakefield, which superimposes destructively on the accelerating field generated by the driver, thereby flattening it. Ideal beam loading occurs when all bunch particles accelerate with the same electric field, thereby minimizing the growth of the beam energy spread and emittance. This is also the condition that maximizes the efficiency of the accelerator [22, 81]. The magnitude of the beam loading effect in the nonlinear blowout regime is governed by the dimensionless parameter [82, 83]

$$\Lambda = \frac{n_b}{n_0} (k_p \sigma_r)^2, \quad (4.1)$$

where n_b is the peak bunch density, n_0 the background plasma density, k_p the plasma wavenumber, and σ_r the transverse size of the beam. To explore the impact of beam loading in our subluminal acceleration scheme, we then conducted a series of simulations varying the witness beam density to achieve different values of Λ .




FIGURE 4.1: Longitudinal electric field (E_z) lineouts for different beam loading regimes characterized by the parameter Λ , evaluated for a short witness beam ($\sigma_z \simeq \lambda_p/2$, $\sigma_r = 0.3 \mu\text{m}$) at initial velocity $0.9c$ in a plasma of density $n_0 = 4 \cdot 10^{18} \text{ cm}^{-3}$ ($k_p = 0.377 \mu\text{m}^{-1}$). The black line ($\Lambda = 1.3 \cdot 10^{-5}$, $n_b = 10^{-3}n_0$) represents the unloaded limit where the witness-beam charge is negligible. The red line ($\Lambda = 0.13$, $n_b = 10n_0$) corresponds to a moderate beam-loading scenario where the wakefield is modified by the presence of the bunch. The blue line ($\Lambda = 1.3$, $n_b = 100n_0$) illustrates the overloaded regime, where excessive beam charge significantly distorts the accelerating structure and degrades the gradient.

Figure 4.1 illustrates the modification of the wakefield structure in our subluminal driver configuration for three distinct loading regimes. We repeat the simulation of Fig. 3.7 in the case of the initial velocity of $0.9c$ using a short witness beam of length $\sigma_z \simeq \lambda_p/2$ and transverse size $\sigma_r = 0.3 \mu\text{m}$, driven by a flying-focus pulse with velocity $v_f = 0.97c$, scanning the beam density over three orders of magnitude ($n_b = 10^{-3}n_0, 10n_0, 100n_0$) at fixed plasma density $n_0 = 4 \cdot 10^{18} \text{ cm}^{-3}$. The black line ($\Lambda \approx 10^{-5}$) shows the unloaded wake, characterized by a sharp spike in the accelerating field E_z . At intermediate loading (red line, $\Lambda = 0.13$),

TABLE 4.1: Trapping and accelerated charge fractions of the witness beam for the three beam-loading regimes of Fig. 4.1. The beam-loading parameter is defined as $\Lambda = (n_b/n_0)(k_p\sigma_r)^2$. Q_0 denotes the initial bunch charge and Q_{trap} the absolute trapped charge. The **Trapping** column reports the charge fraction confined within the accelerating bucket; the $\geq 0.97c$, $\geq 0.98c$, and $\geq 0.99c$ columns report the charge fraction reaching at least the indicated longitudinal velocity. All fractions are normalized to Q_0 .

(n_b/n_0)	Λ	Q_0	Trapping	Q_{trap}	$\geq 0.97c$	$\geq 0.98c$	$\geq 0.99c$
10^{-3}	1.3×10^{-5}	1.1 fC	88%	0.97 fC	87%	76%	32%
10	0.13	11 pC	100%	11 pC	100%	76%	50%
100	1.3	110 pC	50%	55 pC	45%	20%	14%

the superposition of the bunch's self-field locally flattens the gradient, producing a region of nearly uniform E_z across the bunch length. Finally, the blue line ($\Lambda = 1.3$) demonstrates the overloaded regime, where the excessive charge creates a retarding field that distorts the wake structure and degrades the accelerating gradient. The local flattening at $\Lambda = 0.13$ is particularly favorable for our scheme: by equalizing the field experienced by particles across the bunch, it enlarges the longitudinal phase-space volume in which sub-relativistic, heavy particles can be simultaneously phase-locked. As we quantify below in Table 4.1, this translates into a substantial improvement in trapping efficiency, beyond the standard benefit of reduced energy spread. It is important to note that, in this subluminal configuration, beam loading also perturbs the front of the wake. This effect is a unique feature of a subluminal wakefield and is inaccessible in standard quasi-luminal wakes, where the witness cannot influence the field ahead without violating causality.

Having established the qualitative impact of beam loading on the wakefield structure, we now quantify the acceleration efficiency and trapping performance for realistic witness bunch densities. In particular, we evaluate the fraction of the injected bunch that is successfully trapped throughout the acceleration process. Table 3.1 summarizes the trapping and acceleration statistics for the same three representative beam loading regimes presented in Section 4.1, parameterized by the normalized parameter Λ . The results are derived from quasi-3D OSIRIS simulations that track the phase-space evolution of the witness bunch. In the unloaded limit ($\Lambda \approx 10^{-5}$, $Q_0 \simeq 1.1$ fC), trapping is high (88%), but the acceleration quality is limited by the steep longitudinal field gradient, which induces emittance growth. In contrast, the optimal beam loading regime ($\Lambda = 0.13$, $Q_0 = 11$ pC) achieves 100% trapping, corresponding to 11 pC of fully captured witness charge. In this regime, the witness beam's self-field flattens the accelerating gradient (as seen in Fig. 4.1), creating a larger effective volume of phase-space where particles can be stably accelerated. As a result, 5.5 pC of the bunch reaches relativistic velocities $\geq 0.99c$.

In the overloaded regime ($\Lambda = 1.3$, $Q_0 = 110$ pC), the excessive charge creates a strong retarding field that effectively cancels the driver's wakefield, reducing the depth of the accelerating region. Trapping drops to 50% (55 pC), and only 15 pC reaches $\geq 0.99c$. This confirms that beam loading can improve performance up to a critical charge limit $\Lambda \sim 1$, beyond which the acceleration mechanism collapses.

4.2 TRANSVERSE DYNAMICS AND BEAM QUALITY PRESERVATION

While the primary goal of the tapered density profile is to maintain longitudinal phase locking, the variation of the plasma density $n(z)$ inevitably affects the transverse dynamics of the witness beam. To ensure that the accelerated beam remains suitable for applications (e.g., tomography or colliders), we must monitor its transverse quality, quantified by the normalized emittance ϵ_n .

In the nonlinear blowout regime, plasma electrons are expelled from the optical axis, leaving a uniform ion column. This ion channel creates a radial focusing force acting on the witness beam, whose magnitude depends on the local plasma density $n(z)$. For a witness particle of charge q and longitudinal velocity v_z , the radial Lorentz force at a distance r from the axis reads $F_r = q(E_r - v_z B_\theta)$. In the ideal blowout, the bubble interior contains only static ions, so the azimuthal magnetic field is negligible compared to the radial electric field ($B_\theta \ll E_r/c$), and the radial force reduces to the purely electrostatic ion-column expression for any v_z . For a singly-charged negative witness ($q = -e$), valid for the species considered in this work (μ^- , π^- , K^- , e^-), this gives

$$F_r = -\frac{m_e \omega_p^2(z)}{2} r = -\frac{e^2 n(z)}{2\epsilon_0} r, \quad (4.2)$$

where $\omega_{p0} = \sqrt{n_0 e^2 / (m_e \epsilon_0)}$ is the plasma frequency at the reference density n_0 , and $\omega_p(z) = \omega_{p0} \sqrt{n(z)/n_0}$ is its local value at position z along the tapered profile. The dynamics of a witness particle with mass m_p (e.g., a muon) and relativistic Lorentz factor $\gamma(z)$ are governed by the relativistic equation of motion:

$$\frac{d}{dt}(\gamma m_p \dot{r}) = F_r. \quad (4.3)$$

Assuming the adiabatic approximation where the variation of γ is slow compared to the oscillation period ($\dot{\gamma}/\gamma \ll \omega_\beta$), the equation reduces to that of a harmonic oscillator $\ddot{r} + \omega_\beta^2 r = 0$, with a characteristic betatron frequency ω_β [84–86]

$$\omega_\beta(z) = \sqrt{\frac{-F_r}{m_p \gamma r}} = \frac{\omega_{p0} \sqrt{n(z)/n_0}}{\sqrt{2\gamma(z)}} \sqrt{\frac{m_e}{m_p}}, \quad (4.4)$$

consequently, the spatial period of these oscillations, the betatron wavelength $\lambda_\beta \approx 2\pi c/\omega_\beta$, evolves as

$$\lambda_\beta(z) \propto \sqrt{\frac{\gamma(z)}{n(z)}} \sqrt{\frac{m_p}{m_e}}. \quad (4.5)$$

The theoretical framework highlights a specific challenge for our scheme. The injection phase combines a low initial velocity ($\beta \ll 1$) with a rapid density down-ramp. For an electron beam under the same conditions, this combination would pose a serious challenge. Setting $m_p = m_e$ in Eq. 4.5, the betatron wavelength at low γ is only a few micrometers, comparable to, or shorter than, the local density-gradient length set by the steep down-ramp. The

matched envelope, $\sigma_r \propto \sqrt{\gamma/k_\beta}$, is therefore forced to track focusing variations that occur within a single betatron period: the slow-evolution condition that protects the matched envelope ($|\dot{k}_\beta/k_\beta^2| \ll 1$) is violated, leading to envelope mismatch, emittance growth, and ultimately particle loss.

However, the large rest mass of the witness particles introduces a decisive advantage. In this context, the large mass ($m_p \gg m_e$) significantly increases the betatron wavelength, effectively counteracting the destabilizing effects of the low initial velocity (e.g. muons $\lambda_{\beta,\mu} \approx 14\lambda_{\beta,e}$). This increased *inertial stiffness* implies that the beam responds much more slowly to changes in the focusing force. Consequently, the heavy-particle beam remains robust against rapid density variations in the tapered profile, mitigating the risk of defocusing and preserving transverse beam quality even during the critical injection phase.

To validate the qualitative hypothesis on beam quality conservation just presented, we analyzed the evolution of the average normalized transverse emittance. This quantity is defined as $\bar{\epsilon}_n = (\epsilon_{n,x} + \epsilon_{n,y})/2$, where $\epsilon_{n,x}$ and $\epsilon_{n,y}$ represent the normalized root-mean-square (RMS) emittances in the horizontal and vertical planes, respectively. For the x-plane, this is given by:

$$\epsilon_{n,x} = \sqrt{\langle x^2 \rangle \langle p_x^2 \rangle - \langle x \cdot p_x \rangle^2} \quad (4.6)$$

where $\langle \cdot \rangle$ denotes the statistical average over the particle distribution of the trapped beam, x is the transverse position, and p_x is the transverse momentum normalized to $m_p c$. An analogous expression applies to the y-plane. Figure 4.2 presents the results from PIC simulations, comparing the three beam loading presented in Table 3.1.




FIGURE 4.2: Evolution of the normalized transverse emittance ($\bar{\epsilon}_n$) of the trapped beam core as a function of propagation distance z . The unloaded case ($\Lambda \approx 10^{-5}$, black) and the optimal beam loading case ($\Lambda = 0.13$, red) show comparable emittance preservation, demonstrating that the flattened accelerating fields do not degrade transverse quality, but instead slightly improve it. Conversely, the overloaded regime ($\Lambda = 1.3$, blue) exhibits rapid emittance growth driven by severe wakefield distortion.

The evolution of the emittance reveals two distinct phases. In the initial phase ($z < 0.3$ mm), corresponding to the density down-ramp, emittance increases in both the unloaded and

optimal cases. This behavior is expected in plasma-based injection schemes, where the beam envelope must adapt to the rapidly varying focusing forces [86]. However, crucially for our scheme, this growth is strictly contained.

The beam quality evolution was studied in the same three different beam loading regimes investigated in Table 3.1 to show the role of beam loading in beam quality preservation. The simulations show that for the optimal loading case ($\Lambda = 0.13$, red line), the emittance remains comparable to the unloaded baseline ($\Lambda \approx 10^{-5}$, black line). This confirms that the optimal loading regime achieves complete trapping while preserving the linearity of the transverse focusing forces, as reflected in the emittance evolution shown in Fig. 4.2.

In contrast, the overloaded case ($\Lambda = 1.3$, blue line) exhibits rapid, continuous emittance degradation. In this regime, the excessive beam charge severely distorts the wakefield structure, disrupting the linear focusing fields and causing irreversible loss of beam quality.

4.3 STUDY ON TRANSVERSE TRAPPING IN PRESENCE OF A TAPERED DENSITY PROFILE

In this section, we derive the equation of radial motion for a generic massive particle and solve it numerically to study the transverse trapping of particles injected at different radial positions.

Consider a witness particle of mass m_p (e.g., a muon, pion, or kaon) being accelerated by a longitudinal electric field E_z while subject to a linear transverse focusing force F_r , typical of the blowout regime. The equations of motion are [87, 88]:

$$\frac{d\vec{p}}{dt} = qE_z\hat{e}_z - \frac{m_e\omega_p^2(z)}{2}r\hat{e}_r \quad (4.7)$$

where q is the particle charge, and $\omega_p(z) = \omega_{p0}\sqrt{n(z)/n_0}$ is the local electron plasma frequency. Projecting onto the longitudinal (z) and radial (r) axes, and assuming the paraxial approximation ($v_r \ll v_z \approx v$), we obtain a coupled system. Under the paraxial and ultra-relativistic approximations ($v_r \ll v_z \rightarrow c$), the longitudinal energy gain is governed by

$$\frac{d\gamma}{dt} \approx \frac{qE_z(z)}{m_pc}. \quad (4.8)$$

For the radial direction, expanding the time derivative of the momentum $p_r = \gamma m_p \dot{r}$:

$$\frac{d}{dt}(\gamma m_p \dot{r}) = m_p(\gamma \ddot{r} + \dot{\gamma} \dot{r}) = F_r \quad (4.9)$$

Substituting the focusing force expression $F_r = -\frac{m_e}{2}\omega_p^2(z)r$, we arrive at the differential equation describing the transverse oscillations for any particle species:

$$\ddot{r} + \frac{\dot{\gamma}}{\gamma}\dot{r} + \frac{m_e}{m_p}\frac{\omega_p^2(z)}{2\gamma}r = 0 \quad (4.10)$$

Equation (4.10) represents a harmonic oscillator with a time-dependent frequency and a ‘friction’ term $\frac{\dot{\gamma}}{\gamma}\dot{r}$ arising from relativistic mass increase (adiabatic damping). Crucially, the stability depends explicitly on the mass ratio m_e/m_p . Heavier particles experience a reduced effective restoring force (scaled by m_e/m_p), which leads to longer betatron wavelengths and increased robustness against density variations compared to electrons.




FIGURE 4.3: Transverse stability analysis of different sub-relativistic particles injected into the tailored density profile. (a-d) Evolution of the theoretical plasma density $n(z)/n_0$ (log scale) required for longitudinal phase-locking for electrons, muons, pions, and kaons, respectively. (e-h) Computed transverse trajectories for the corresponding particle species. All particles are injected with a radial offset of $r_{b0} = 1.9 \mu\text{m}$ into the wakefield bubble with an initial radius of $r_{b0} = 2.0 \mu\text{m}$. The grey shaded regions represent the plasma background evolution ($r_b(z) = r_{b0} \cdot [n(z)/n_0]^{-1/2}$). Despite the proximity to the injection wake's border, the particles remain confined within the accelerating structure throughout the entire propagation.

To investigate the transverse trapping capabilities of our tapered density profile, we numerically solved Equation (4.10) for different particle species (electrons, muons, pions, and kaons). For each particle, we computed the specific tailored density profile $n(z)$ required to maintain perfect phase-locking for each specific mass. All particles are injected with a radial offset $r_{b0} = 1.9 \mu\text{m}$ into a wakefield bubble of initial (when $n(z=0)/n_0 = 1$) radius $R_b(0) = 2.0 \mu\text{m}$.

Figure 4.3(a-d) shows the evolution of the plasma density $n(z)/n_0$ (log scale) required for the longitudinal phase-locking of each species. Figure 4.3(e-h) presents the corresponding transverse trajectories. For the electron case [Fig. 4.3(a,e)], the rapid relativistic acceleration dictates a steep density gradient to maintain synchronism, which results in a fast contraction of the wakefield bubble radius. Nevertheless, the particle remains successfully confined within the accelerating structure. Conversely, heavier particles like muons, pions, and kaons [Fig. 4.3(b-d, f-h)] accelerate more slowly due to their inertia. This requires a much more gradual density variation, creating a slowly evolving accelerating channel that remains wide and stable over significantly longer distances.

These results highlight an important feature of our acceleration scheme. The density profile required for phase locking is fundamentally different across species. Because electrons have a significantly lower mass than muons, pions, and kaons, they experience rapid relativistic acceleration and reach the velocity of the wake almost instantaneously. Thus, in this case, the ideal density profile is steep and monotonically rising. In contrast, heavier particles like muons, pions, and kaons accelerate more slowly due to their inertia, requiring a pronounced initial density down-ramp to maintain phase-locking (Fig. 4.3, panels a–d). A down-ramp weakens the ion channel focusing force, making this the most demanding regime for transverse confinement. The fact that all three heavy species remain stably confined throughout this down-ramp, as shown in panels (f–h), is a direct consequence of their large inertia, which prevents the beam from responding to the rapid variation in focusing strength.

4.4 EXPERIMENTAL FEASIBILITY OF OUR TAPERED DENSITY PROFILE

Controlled density variations have been successfully employed in plasma wakefield acceleration experiments since the earliest demonstrations of the technique. At first, sharp density transitions (down-ramps) were introduced to facilitate electron injection [89–91], formed by utilizing the natural edges of gas jets or by placing a thin blade at the flow boundary to create a localized shock [92]. Subsequently, longer-scale up-ramps were implemented to extend the dephasing length and enhance energy gain [93–95], typically achieved through staged gas jets with independent backing pressures or tilted nozzle geometries [96, 97]. These established techniques demonstrate that both abrupt and gradual density variations can be reliably implemented in laboratory settings.

For acceleration stages limited to the millimeter range, a piecewise-linear approximation offers a viable option for experimentally realizing the tailored plasma density profile presented in Sec. 3.3. From an experimental perspective, generating a multi-segment density profile with precise piecewise-linear transitions remains an open challenge. A potential pathway to realize such structures involves the use of custom-machined gas cells or variable-geometry nozzles, which provide longitudinal control over the gas flow [94, 98]. While advanced plasma wakefield experiments successfully employ these techniques to produce individual, controlled up-ramps and down-ramps [90, 95], extending them to the complex tailoring required for this scheme would demand dedicated target development and careful experimental characterization.

In the previous Sections, we derived an ideal density profile $n(z)$ capable of maintaining phase synchronism throughout the acceleration process. In this section, we will briefly discuss the feasibility of the tailored density profiles presented in this thesis. We consider two primary experimental approaches for density tailoring: supersonic gas jets and segmented flow apertures.

To validate this approach and test the acceleration scheme’s tolerance to non-ideal experimental conditions, we approximated the continuous theoretical curve using a discretized se-




FIGURE 4.4: Validation of the experimental feasibility using a piecewise-linear density profile. (a) Comparison between the exact theoretical profile (‘ideal density (theory)’, blue) and a surrogate profile constructed from linear segments (‘approximated density’, orange dashed). The bottom panel shows the relative density error. (b) Evolution of the maximum beam velocity β_{max} obtained from OSIRIS simulations employing the two density profiles shown in (a). The results using the approximated profile (orange) closely track the ideal simulation (blue) and the ODE prediction (black dashed). The negligible deviation $\Delta\beta$ confirms that the discretization of the density gradient does not impact the final energy gain.

quence of five linear ramps. Figure 4.4 presents a direct comparison between the ideal physics case and this discrete approximation. Figure 4.4(a) shows the density profiles, where the relative error introduced by the linear segments remains bounded below 10%. Crucially, Figure 4.4(b) demonstrates that this discretization does not compromise acceleration efficiency.

We performed full OSIRIS particle-in-cell (PIC) simulations for both density profiles. The evolution of the maximum beam velocity β_{max} for the approximated profile (orange solid line) closely tracks both the PIC simulation of the ideal density (blue solid line) and the theoretical prediction derived from the ODE model (black dashed line). The velocity mismatch, defined as $\Delta\beta = \beta_{\text{approx}} - \beta_{\text{theory}}$, remains negligible (< 0.01), confirming the robustness of the phase-locking mechanism against localized density deviations and manufacturing imperfections.

For higher energy gains, heavier particle masses, or very low initial velocities, the acceleration length extends to the centimeter scale. A profile spanning several centimeters and exhibiting a non-monotonic shape cannot be produced by a single nozzle. A practical solution consists of combining two partially overlapping supersonic gas jets, whose superposition reproduces both the down-ramp and the subsequent up-ramp of the theoretical profile. Figure 4.5 shows the comparison between the ideal ODE profile and the two-jet approximation, constructed using the parameters listed in Table 4.2. The top panel shows that the superposition of the two Gaussian jets closely tracks the theoretical curve across the full stage, correctly capturing the entrance ramp, the low-density plateau, and the exit ramp. The bottom panel confirms that the pointwise density error remains within 25% throughout the interaction length. As shown in Figure 4.5(a), the total density resulting from the superposition of two

TABLE 4.2: Parameters of the two Gaussian gas jets used to approximate the ideal tapered density profile. Lengths are derived using a reference density $n_0 = 5.0 \times 10^{17} \text{ cm}^{-3}$ (corresponding to a plasma skin depth $c/\omega_p \simeq 7.515 \text{ } \mu\text{m}$). Each jet is described by the function $n_j(z) = b + A \exp[-4 \ln 2 ((z - z_0)/\text{FWHM})^2]$. The last column reports the normalized on-axis peak density.

Jet	Center z_0 [mm]	FWHM [mm]	σ [mm]	Pedestal b [cm^{-3}]	Peak A [cm^{-3}]	A/n_0
1	-0.98	2.48	1.05	2.0×10^{16}	1.0×10^{18}	2.00
2	21.49	4.67	1.98	2.0×10^{16}	1.8×10^{18}	3.60

Gaussian jets mimics the theoretical profile, capturing both the down-ramp and up-ramp transitions with a mismatch between the theoretical and approximated curve always below 25%. Full kinetic validation of this cm-scale configuration was not performed due to computational




FIGURE 4.5: Experimental feasibility using a two-jet gas profile. (a) Comparison between the ideal tapered density $n_{\text{ODE}}(z)$ (blue) and an implementable profile (orange, dashed) constructed from two Gaussian gas jets (parameters in Table 4.2). The top panel shows $n(z)$ on a logarithmic scale; the bottom panel shows the relative percentage error, defined as $100 \times (n_{\text{jets}} - n_{\text{ODE}})/n_{\text{ODE}}$. (b) Comparison of the energy gain between the theoretical ODE prediction (dashed black), the PIC simulation using the ideal density (blue), and the PIC simulation using the gas jet approximation (orange). The two-jet surrogate closely reproduces the required profile and maintains the expected acceleration performance.

constraints. However, we tested the first 4 mm of acceleration.

The results, shown in Figure 4.5(b), confirm that the acceleration performance of the two-jet profile closely follows the ODE prediction (presented in Sec. 3.3). This agreement validates the effectiveness of the two-jet approximation in reproducing the key features of the ideal density profile and achieving the expected acceleration performance.

4.5 CONCLUSIONS

In this chapter, we expanded upon the purely longitudinal, single-particle framework established in Chapter 3 to address the critical physical and experimental challenges required for

a functional accelerator: managing high-charge beam loading, preserving transverse beam quality, and implementing the tailored density profiles in the laboratory. First, we demonstrated that our subluminal phase-locking scheme successfully supports significant witness bunch charges. By characterizing the wakefield under various beam-loading conditions, we identified an optimal loading regime ($\Lambda \simeq 0.13$). In this regime, the superposition of the bunch's self-field flattens the accelerating gradient, effectively enlarging the stable phase-space volume. This flattening allows for 100% charge trapping without disrupting the underlying wakefield structure, whereas moving into an overloaded regime ($\Lambda \sim 1$) severely degrades both the gradient and the capture efficiency.

Crucially, we established that the transverse quality of the accelerated beam is preserved under these optimal conditions. The tailored density down-ramp required to maintain longitudinal phase-locking inherently weakens the ion channel focusing force, posing a severe risk of envelope mismatch and emittance growth for standard electron beams. However, analytical modeling and fully kinetic simulations revealed that heavier sub-relativistic particles (muons, pions, and kaons) possess an intrinsic robustness against these rapid density variations. Their large inertia substantially lengthens their betatron wavelength, acting as a stabilizing anchor that prevents the beam from responding destructively to the shifting focusing forces. As a result, the normalized transverse emittance remains strictly contained, validating the scheme's ability to deliver high-quality beams.

Finally, we addressed the experimental feasibility of generating the required non-monotonic density profiles. Recognizing that perfect continuous curves are theoretical idealizations, we tested the robustness of the phase-locking mechanism against discrete approximations. For millimeter-scale stages, a piecewise-linear surrogate profile maintained the relative density error below 10% and yielded a negligible velocity deviation ($\Delta\beta < 0.01$) compared to the ideal case. For longer centimeter-scale stages, the superposition of two Gaussian gas jets successfully reproduced the essential down-ramp and up-ramp transitions. While constructing these exact multi-segment geometries using custom-machined gas cells or variable-aperture nozzles remains a demanding experimental challenge, our results prove that the underlying physics is highly resilient to localized density errors, provided the global gradient trend is maintained.

Collectively, these findings confirm that the combination of subluminal spatio-temporal drivers and tailored density profiles is not merely a theoretical curiosity, but a robust, load-tolerant mechanism capable of delivering high-quality bunches of short-lived particles.

Take-home messages

- **Beam loading controls energy spread and maximum charge we can accelerate.** In blowout, the loading strength is controlled by

$$\Lambda = \frac{n_b}{n_0} (k_p \sigma_r)^2.$$

In our subluminal stage, an intermediate regime ($\Lambda \simeq 0.13$) flattens E_z and maximizes performance (100% trapping; 50% reaching $\beta \geq 0.99c$), while overloading ($\Lambda \sim 1$) collapses the structure (50% trapping; 14% reaching $\beta \geq 0.99c$).

- **The tapered density modifies focusing, but heavy particles are intrinsically robust.** The ion-channel focusing force scales with the local density,

$$F_r = -\frac{m_e \omega_p^2(z)}{2} r, \quad \omega_p(z) = \omega_{p0} \sqrt{\frac{n(z)}{n_0}},$$

giving a betatron frequency

$$\omega_\beta(z) = \frac{\omega_{p0}}{\sqrt{2\gamma(z)}} \sqrt{\frac{n(z)}{n_0}} \sqrt{\frac{m_e}{m_p}}, \quad \lambda_\beta \propto \sqrt{\frac{\gamma(z)}{n(z)}} \sqrt{\frac{m_p}{m_e}}.$$

Thus, for $\mu/\pi/K$ the large $m_p \gg m_e$ lengthens λ_β and suppresses sensitivity to rapid focusing changes in the down-ramp.

- **Beam quality is preserved in the optimal-loading regime.** We monitor transverse quality through the normalized RMS emittance

$$\epsilon_{n,x} = \sqrt{\langle x^2 \rangle \langle p_x^2 \rangle - \langle x p_x \rangle^2}, \quad \bar{\epsilon}_n = \frac{1}{2} (\epsilon_{n,x} + \epsilon_{n,y}).$$

PIC shows $\bar{\epsilon}_n$ remains comparable between the unloaded case and $\Lambda \simeq 0.13$, whereas overloading ($\Lambda \sim 1$) induces strong emittance growth via wake distortion.

- **Transverse confinement in the taper follows a universal ODE with explicit mass scaling.** Under the paraxial approximation the radial motion is

$$\ddot{r} + \frac{\dot{\gamma}}{\gamma} \dot{r} + \frac{m_e}{m_p} \frac{\omega_p^2(z)}{2\gamma} r = 0,$$

and solving this with the species-specific phase-locking $n(z)$ confirms stable confinement for electrons and even more robust trajectories for $\mu/\pi/K$. The same down-ramp that enables phase control also increases transverse acceptance through $R_b \propto n^{-1/2}$ during injection.

- **The required density tailoring is experimentally doable.** A piecewise-linear surrogate keeps the density error bounded (below $\sim 10\%$) and reproduces the acceleration (PIC: $\Delta\beta < 0.01$ versus the ideal profile). For cm-scale stages, an implementable two-jet surrogate can be built by superposition of two gas jets.

CHAPTER 5

CO-LIVING OF THE SELF-MODULATION AND HOSING INSTABILITIES IN AWAKE

Plasma-based accelerators offer the potential to sustain accelerating gradients orders of magnitude beyond those achievable with conventional RF technology. However, the achievable acceleration distance, and thus the maximum energy gain per single acceleration stage, is fundamentally constrained by the finite energy stored within the driver. In particle-driven wakefield acceleration, this limit is dictated by driver depletion. As the drive beam excites the wakefield, its kinetic energy is progressively depleted, transforming into both the field and the plasma's kinetic energy.

The AWAKE experiment at CERN addresses this limitation by employing a high-energy *proton* bunch as the wakefield driver [99]. Proton beams currently available at CERN possess sufficient stored energy to theoretically sustain GV/m-class wakefields over *hundreds of meters* in a single plasma stage, thereby reducing the necessity for multi-stage acceleration. However, high-energy proton bunches, such as those from the Super Proton Synchrotron (SPS) at CERN, which is used at AWAKE, have longitudinal sizes ($\sigma_z \approx 6 - 12$ cm) that far exceed the plasma wavelength at densities required for high-gradient acceleration. To achieve longitudinal resonance ($\lambda_p \approx \sigma_z$) with a centimeter-scale bunch, the plasma density must be drastically reduced. This, however, yields gradients that accelerate so slowly that they offer no advantage over standard RF cavities. To circumvent this, we rely on the Self-Modulation Instability (SMI) [100] to modulate the proton beam into a train of micro-bunches. This mechanism enables the beam to resonantly drive the wakefield while remaining in the high-density, high-gradient regime.

The SMI is a transverse instability that micro-bunches the long proton beam at the plasma frequency. The resulting micro-bunches act cooperatively to drive a robust wakefield capable of accelerating witness electrons. However, while SMI is fundamental to the AWAKE concept, it inevitably co-exists with a competing mode: the Hosing Instability (HI) [101]. HI manifests itself as a slice-dependent transverse oscillation of the beam centroid. In contrast to SMI, which modulates the beam envelope to excite the wake, HI enhances beam centroid displacements that can lead to the breakup of the proton bunch. If uncontrolled, hosing can severely degrade

the quality of the wakefield, reduce the overlap between the drive and witness beams, and ultimately destroy the acceleration process. Scaling proton-driven wakefield acceleration requires a precise understanding of how these two competing instabilities co-evolve. Sustaining high-gradient wakefields over 10-100 meter-scale plasmas relies on saturating the SMI while simultaneously suppressing the beam-disrupting transverse oscillations characteristic of the HI. To study this interplay experimentally and better understand the development of the HI, in the AWAKE experiment, the injected electron seed bunch can be intentionally misaligned, thereby providing a seed to observe the HI growth under controlled initial conditions.

Section 5.1 summarizes the essential theoretical background on the transverse instabilities relevant to proton-driven wakefield acceleration. Section 5.1.1 introduces the self-modulation instability (SMI) and the distinction between noise-driven and seeded growth, while Sec. 5.1.2 reviews the hosing instability (HI) and the convective envelope scaling that we will later use as a reference for the evolution of centroid observables. In Sec. 5.2, we then introduce the AWAKE experimental set-up and the Run 2b program, with Sec. 5.2.1 discussing how controlled seeding sets the transverse asymmetry and how SMI and HI can co-evolve in seeded configurations, Sec. 5.2.2 describing the diagnostics and analysis observables used in this work (in particular time-resolved streak-camera centroid tracking and effective-length scans), and Sec. 5.2.3 presenting the plunger-scan streak-camera measurements of the seeded transverse centroid evolution, including the centroid-displacement observable used for quantitative comparison. Section 5.3 then models the same seeded-misalignment scenario with the quasi-static PIC code QUICKPIC. In particular, Sec. 5.3.1 summarizes the simulation set-up, Sec. 5.3.2 reproduces the plunger-scan analysis in the simulations, and Sec. 5.3.3 examines the dependence of the centroid growth on the electron-seed misalignment and relates it to the transverse wakefield structure. Section 5.4 outlines possible future extensions of this work. The chapter closes with a set of take-home messages summarizing the main conclusions.

5.1 PHYSICAL BACKGROUND: SMI AND HI IN AWAKE

This section provides an overview of the physical mechanisms underlying the self-modulation instability (SMI) and the hosing instability (HI) in the context of the AWAKE experiment. Figure 5.1 illustrates the fundamental differences between these two regimes for a long particle beam propagating in plasma. For each case, the plot shows the normalized proton-bunch density (grey scale) together with the transverse wakefield component in the plane of the panel, $W_i = E_i - cB_j$ (with $i, j = x, y$ and $i \neq j$ represented in red-blue scale). For the positively charged proton beam, the red and blue regions across the axis correspond to alternating focusing and defocusing phases. The green curve indicates the transverse centroid slice-by-slice as a function of the longitudinal coordinate. In a self-modulation-dominated scenario (Fig. 5.1(a)), the proton-bunch density remains symmetric with respect to the beam axis, keeping the centroid close to zero. In contrast, in a case dominated by hosing (Fig. 5.1(b)), a clear centroid oscillation develops alongside a transverse field and a density pattern that visibly breaks the axial symmetry, driving the characteristic growth of the hosing.

We summarize the key theoretical concepts and characteristic signatures of both instabili-

ties, setting the stage for the subsequent simulation and experimental analysis. In particular, we highlight the differences between seeded and noise-driven scenarios, the role of initial conditions (e.g. transverse offsets), and the expected interplay between longitudinal modulation and transverse centroid dynamics.




FIGURE 5.1: Illustrative comparison between a self-modulation-dominated case (a) and a hosing-dominated case (b) for the propagation of a long particle beam in plasma. The grey scale shows the normalized proton-bunch density n/n_0 (arb. units), while the red–blue scale shows the transverse wake-field component in the plane of each panel (W_x in (a) and W_y in (b); red positive, blue negative). The green curve indicates the slice-by-slice transverse centroid. In (a), the density and transverse-field structure remain symmetric with respect to the beam axis, and the centroid stays close to zero. In (b), a clear centroid displacement develops together with a transverse-field and density pattern that breaks axial symmetry, characteristic of hosing growth.

5.1.1 SELF-MODULATION INSTABILITY (SMI)

The fundamental mechanism driving SMI is the feedback established between transverse wakefields and the local bunch density. Initial perturbations, whether arising from noise or an external seed, induce transverse wakefields that alternate between focusing and defocusing phases with a periodicity set by the plasma wavelength, so that successive longitudinal slices of the beam experience opposite transverse forces. This transverse motion modulates the local charge density, which in turn reinforces the wakefield at the plasma frequency ω_p [100]. As the beam propagates, this positive feedback drives a longitudinal density modulation at ω_p , transforming the bunch into a train of microbunches separated by $\lambda_p = 2\pi v_\phi / \omega_p \simeq 2\pi c / \omega_p$ that resonantly excite the plasma wave. For AWAKE, this microbunch train is a prerequisite that allows a long proton driver to sustain large wakefields over extended distances.

A critical distinction exists between *unseeded* (driven by random fluctuations) and *seeded* self-modulation. In the unseeded regime, the instability grows from intrinsic fluctuations within the bunch and plasma, resulting in significant shot-to-shot variability in both the phase and amplitude of the wakefields. Conversely, seeded SMI utilizes an externally controlled perturbation to impose a well-defined initial condition, thereby fixing the modulation phase and ensuring reproducibility [102].

The simulations and analyses presented in this work rely on seeded configurations, treating seed parameters (such as offset, size, and timing) as control variables that govern both

modulation growth and the beam's transverse stability. For a weakly modulated beam slice at co-moving position ζ (with $\zeta \equiv k_p \zeta$ and $\hat{z} \equiv k_p L$), the density modulation amplitude scales as

$$\delta n_b / n_{b0} \propto \frac{\exp(N_{sm})}{\sqrt{N_{sm}}}, \quad N_{sm} = \frac{3\sqrt{3}}{4} \left(\mu \hat{k} \beta^2 |\zeta| \hat{z}^2 \right)^{1/3}, \quad (5.1)$$

where $\mu = 2I_1(r_0)K_1(r_0)$ accounts for finite beam radius and return-current effects ($r_0 = k_p r_b$) and $\hat{k}\beta = (k_b/k_p)(2\gamma)^{-1/2}$ [103]. Using $(k_b/k_p)^2 = (n_{b0}/n_0)(m_e/M_b)$, Eq. (5.1) makes explicit the scaling (up to $\mathcal{O}(1)$ profile factors)

$$N_{sm} \propto \left[(n_{b0}/n_0) \gamma^{-1} k_p^3 |\zeta| L^2 \right]^{1/3}. \quad (5.2)$$

5.1.2 HOSING INSTABILITY (HI)

In a standard beam-plasma interaction, the Hosing Instability (HI) can arise from random fluctuations in the beam centroid or plasma density, resulting in unpredictable transverse oscillations of the bunch centroid. However, in seeded SMI configurations, the interplay between longitudinal modulation and transverse dynamics changes significantly. While introducing a seed to trigger SMI tends to control the wake phase, any asymmetry in the seed as for example, a transverse offset between the seed and the proton bunch, may introduce a preferred transverse direction. If, say, the seed is injected with a vertical offset Δy relative to the proton axis (with $\Delta x = 0$), the seeded wakefield is itself vertically displaced, and centroid oscillations of the proton bunch develop predominantly in the y - z plane, while the orthogonal x - z plane remains comparatively symmetric. Moreover, the resulting behavior may not always resemble a clean periodic centroid oscillation along the bunch: depending on the configuration, the centroid evolution can instead appear as a progressive, monotonic drift of successive longitudinal slices away from the beam axis.

In the overdense, long-beam regime relevant to proton-driven PWEA, the transverse hosing instability is convective. The centroid displacement does not grow as a simple exponential with a constant rate, but instead follows an asymptotic envelope [103]. The growth rate of this envelope is set by the competition between the focusing action of the plasma wakefields, which the beam resists through its own transverse stiffness, and the self-consistent coupling between beam centroid motion and wake centroid motion. The beam's transverse stiffness in this regime is characterized by the betatron wavenumber

$$k_b = \sqrt{\frac{n_{b0} e^2}{\epsilon_0 m_b c^2}}, \quad (5.3)$$

which sets the natural transverse oscillation scale of the beam particles inside the plasma wake.

For a weakly tilted beam slice at co-moving position ζ , an analogous analysis yields a centroid amplitude that grows as

$$|x_c| \propto \frac{\exp(N_h)}{\sqrt{N_h}}, \quad N_h = \frac{3\sqrt{3}}{4} \left(\mu \hat{k} \beta^2 |\zeta| \hat{z}^2 \right)^{1/3}, \quad (5.4)$$

with the same definitions of μ , $\hat{k}\beta$, ζ , and $\hat{z}$ as in Sec. 5.1.1 [103]. Equation (5.4) has the same functional form as the SMI exponent N_{sm} of Eq. (5.1) and therefore obeys the same scaling,

$$N_h \propto \left[(n_{b0}/n_0) \gamma^{-1} k_p^3 |\zeta| L^2 \right]^{1/3}, \quad (5.5)$$

up to $\mathcal{O}(1)$ numerical prefactors. The coincidence of the two exponents is a central result of the coupled-mode analysis [103]: in the overdense regime, SMI and hosing grow at comparable rates. The qualitative distinction between the two instabilities therefore lies in the transverse mode structure they excite—a symmetric density modulation for SMI versus an antisymmetric centroid displacement for hosing.

5.2 THE RUN 2B OF AWAKE EXPERIMENT

In AWAKE Run 1, 400 GeV SPS proton bunches propagated through a 10 m rubidium plasma column. Triggered by a relativistic ionization front, the proton bunch self-modulated into microbunches at the plasma period, and externally injected electrons were accelerated to the multi-GeV level. These results established seeded self-modulation (SM) as a viable method to transform a long proton bunch into an efficient resonant wakefield driver.

The AWAKE plasma source is a rubidium vapor cell of nominal length $L_p = 10$ m, heated to a uniform temperature so that the rubidium vapor density, and hence the plasma density once the vapor is ionized, is set by the saturated vapor pressure and is uniform along the cell to better than 0.2% [104, 105]. Rubidium is chosen for its low first ionization potential, which allows full single ionization of the vapor by a co-propagating short laser pulse of moderate intensity, and for its relatively heavy ion mass, which limits ion motion on the timescale of the proton bunch transit [106]. The same laser pulse provides the relativistic ionization front that defines the start of the plasma column seen by the trailing proton bunch.

AWAKE Run 2 is designed to move from proof-of-principle to a scalable accelerator concept: (i) produce a reproducible, saturated self-modulated proton driver (the self-modulator), (ii) demonstrate controlled capture/acceleration of an externally injected electron bunch while preserving beam quality, and (iii) prepare the technology for longer plasmas by separating self-modulation from acceleration and by developing plasma sources and injection systems scalable beyond the 10 m class. In parallel, Run 2 enables dedicated studies of the physics of seeded self-modulation and its interplay with transverse instabilities, which are crucial for understanding the limitations and optimization of the AWAKE concept. In particular, the use of upgraded plasma sources with controlled density profiles and the ability to vary the effective plasma length provide unique opportunities to probe the growth of modulation and hosing signatures as a function of propagation distance.

5.2.1 ELECTRON BUNCH SEEDING OF SELF-MODULATION AND HOSING

As discussed in the previous sections, the long SPS proton driver used in the AWAKE experiment relies on SMI to transform into a resonant microbunch train [11, 99, 100]. In AWAKE

Run 2, this process can be seeded by a short electron bunch [102]. Figure 5.2 summarizes the conceptual layout, highlighting the distinction between (a) seeded SMI and (b) seeded hosing. In both scenarios, the interaction is initiated by a short electron bunch injected ahead of the long proton driver at the plasma entrance. The beam and plasma parameters utilized in this work correspond to the long-bunch, overdense regime, a parameter space where both SMI and hosing naturally co-develop [103, 107]




FIGURE 5.2: Experimental setup of seeding self-modulation and hosing instability. Sketch of the set-up, not to scale.

As the electron seed bunch propagates through the preformed plasma ahead of the proton driver, it excites an initial wakefield with a well defined phase. This wakefield acts as a controlled initial condition for the SMI feedback loop described in Sec. 5.1.1, replacing the random fluctuations that would otherwise set the modulation phase shot to shot. The result is a reproducible microbunch train of periodicity λ_p , as illustrated in Fig. 5.2 (a), and ultimately the controlled electron injection and acceleration demonstrated in Run 1 [11, 102].

Compared to the ionisation front seeding used in Run 1, an electron bunch seed has the additional advantage that it can in principle enforce a more uniform onset of modulation along the driver, mitigating the phase slippage associated with an initially unmodulated bunch head in a preformed plasma. In AWAKE, hosing can be deliberately seeded by introducing a controlled transverse offset between the seed electron bunch trajectory and the proton-bunch axis. The proton slices then experience a coherent ξ -dependent transverse kick that initiates centroid motion predominantly in the plane of the imposed misalignment, as sketched in Fig. 5.2 (b).

Run 2 explicitly exploits the electron bunch as a *tunable* seed: by varying the seed charge, timing, and transverse alignment relative to the proton bunch, one can control the initial wakefield amplitude, the onset of seeded self-modulation, and the strength of the initial transverse asymmetry responsible for seeded hosing [102, 108, 109]. This thesis focuses on the seeded regime where the electron bunch provides reproducible initial conditions for SMI, while an intentional transverse misalignment is used to excite hosing in a controlled manner. This approach enables quantitative comparisons between diagnostics and simulations based on experimentally accessible observables described in Sec. 5.2.

A key novelty of the Run 2b campaign analysed in this chapter is the ability to vary the effective plasma length L using the plunger system, while keeping the proton bunch, the seed electron bunch, and the rubidium vapour conditions fixed. Previous AWAKE measurements

gave access to the fully developed hosing instability at the end of the plasma column. The plunger scan extends this to a series of shorter effective lengths, so that the development of the instability can be followed along the propagation of the beam in plasma rather than observed only in its saturated state.

The parameters of the experimental set-up investigated in this Chapter have been summarized in Tab. 5.1.

TABLE 5.1: Experimental Parameters of AWAKE Run 2B for seeded hosing studies.

Group	Parameter	Value
Proton bunch		
Population	N_p	2×10^{11}
Charge (from $N_p e$)	Q_p	17.1 nC
Energy	E_p	400 GeV
Bunch length (rms)	$\sigma_{z,p}$	4.35 cm
Transverse size (rms)	$\sigma_{\perp,p}$	0.2 mm
Seeding electron bunch		
Population	N_e	10^9
Charge (from $N_e e$)	Q_e	≈ 0.16 nC
Energy	E_e	19 MeV
Bunch length (rms)	$\sigma_{z,e}$	0.580 mm
Transverse size (rms)	$\sigma_{\perp,e}$	0.470 mm
Longitudinal distance $t = 0$	Δx_1	22 cm
Plasma and diagnostics		
Plasma density	n_0	$2 \times 10^{14} \text{ cm}^{-3}$
Plasma length	L_p	10.3 m
Plasma radius	r_p	~ 1.5 mm
Distance to detector	Δx_D	3.5 m

5.2.2 AWAKE RUN 2B DIAGNOSTICS USED IN THIS WORK

A summary of the components of the AWAKE experiment is shown in Fig. 5.3. To experimentally study wakefield-driven dynamics, this thesis relies primarily on two key diagnostic: time-resolved streak-camera imaging and the ability to vary the effective plasma length [109]. The streak camera system images optical transition radiation emitted by the proton bunch as it crosses a screen located downstream of the plasma. The incoming light is swept across the detector, so that arrival time is mapped onto a spatial coordinate. The result is a single shot, time resolved image in which one axis encodes the longitudinal position along the bunch, $z \equiv ct$, and the other a transverse coordinate at the screen.




FIGURE 5.3: AWAKE Run 2b experimental setup and diagnostics.

In Run 2b, two streak cameras image the same screen along orthogonal transverse axes, so that one camera produces an (x, t) projection and the other a (y, t) projection of the proton bunch on a shot by shot basis. This dual view capability is essential for hosing studies. A seeded transverse offset breaks the axial symmetry preferentially in one plane, and the centroid excursion is expected to develop predominantly along the misalignment direction while the orthogonal plane remains comparatively symmetric. Having both projections available allows us to verify this directly, by checking that the centroid grows in the seeded plane and not in the perpendicular one, rather than inferring it from a single view. While Run 2b retains time-integrated imaging of defocused protons to offer a complementary handle on transverse wakefield strength via downstream halo distributions [110, 111], the streak camera remains the most direct diagnostic of the longitudinal development of the proton bunch [112].

Furthermore, a crucial capability added during Run 2b is the use of plungers in the plasma source to vary the effective plasma length under identical upstream conditions [109]. Plungers are movable rods that can be inserted radially into the rubidium vapour column at chosen longitudinal positions, locally disrupting the vapour density and thereby preventing plasma formation downstream of the insertion point. By selecting the position of the plunger, the effective interaction length L can be set on a shot to shot basis, while the proton bunch, the seeding electron bunch, and the upstream plasma conditions are kept fixed. The resulting plunger scans allow the evolution of the beam to be probed as a function of propagation distance in plasma, rather than relying on a single fixed interaction length.

However, a significant interpretation caveat is that streak-camera images are recorded on

a screen located downstream of the plasma after a finite drift in vacuum. Consequently, the measured transverse distribution is not necessarily identical to the distribution at the plasma exit because the drift maps angular components into real space. This vacuum drift effect must be accounted for when comparing experimental data to simulations, which naturally provide particle distributions at the plasma exit [113], a topic discussed in further detail in Sec. 5.3.

5.2.3 EXPERIMENTAL RESULTS OF HOSING INSTABILITY EVOLUTION

As mentioned in Sec. 5.2, a key capability introduced in Run 2b is to vary the effective plasma interaction length using movable plungers, without altering the upstream beam configuration. This allows to repeat the same seeded configuration at different lengths L and measure how an experimentally accessible observable evolves with propagation distance in plasma. In the seeded-misalignment configuration considered here, the electron bunch is intentionally injected with a transverse offset with respect to the proton-bunch axis, providing a controlled seed for a hosing, as described in Sec. 5.2.1.

Figure 5.4 shows one of the plunger scan to study hosing growth measured with the streak-camera system in the AWAKE Run 2b. In this case, an electron beam was misaligned by 0.56 mm with respect to the proton bunch axis (for other parameters see Tab. 5.1). To construct this dataset, multiple (typically 10) single-shot streak-camera images were recorded at each plasma length L , and averaged. The images focus on the transverse region of interest containing the proton bunch. The vertical streak axis provides a time-resolved projection of the beam density. Throughout this chapter, we express this axis as a longitudinal coordinate along the bunch, $z \equiv ct$ (equivalently, a co-moving coordinate ξ up to a constant offset), while the other axis corresponds to the transverse coordinate in the measurement plane (here denoted y).

Superimposed on each panel of Fig. 5.4 is the centroid trajectory of the reconstructed slice, $y_c(z)$ (red line). This metric is derived using a slice-wise centre-of-mass estimator applied to thin longitudinal bins along the streak axis. As the proton bunch propagates in plasma, the slices in the defocusing phase of the wakefield are pushed outward and form a diffuse, low intensity halo of protons at large transverse distance from the bunch core. A centre of mass computed over the full transverse extent of the streak camera image would be biased by this halo, because even a small fraction of charge at large $|y|$ contributes a large lever arm to the average. To avoid this bias, the centroid is computed only within a transverse window around the bunch core, whose width is increased with the effective plasma length L to follow the broadening of the core as the instability develops.

Two qualitative trends are evident from Fig. 5.4. At short effective lengths, the centroid remains close to the reference axis over most of the bunch, consistent with an early stage of transverse asymmetry growth. As L increases, the centroid develops a progressively larger slice-to-slice excursion, indicating the growth of a hosing-like transverse displacement in the seeded configuration. Because the streak images are recorded on a screen downstream of the plasma, the observed distributions can include the effect of transport in vacuum. Consequently, the measured centroid excursion should be interpreted as an experimentally accessible proxy for transverse wake-driven dynamics rather than as a direct measurement of wakefields (Sec. 5.2.2). In Sec. 5.3, we will further explore this point showing that even the




FIGURE 5.4: Plunger-scan visualisation of the seeded transverse centroid dynamics in AWAKE Run 2b. Each panel shows the mean streak-camera image of the proton bunch at the plasma exit for a given plasma length (labels in metres), after shot-to-shot alignment using the timing marker (not shown) and restricting to the proton-beam region of interest. The vertical axis ζ is the longitudinal coordinate along the bunch, obtained from the time resolved sweep of the streak camera, and the horizontal axis y is the transverse coordinate at the screen in the plane of the imposed seed misalignment. The red markers overlay the transverse centroid extracted slice-by-slice along the time-resolved axis using a centre-of-mass estimator within a length-dependent window. The progressive increase of the centroid excursion with plasma length provides a direct experimental handle on the growth of the hosing-like transverse displacement when the instability is seeded by an intentionally misaligned electron bunch.

free-streamed particles still can give us hints on the instability development.

To compare the plunger scan measurements with the convective hosing envelope derived in Sec. 5.1.2, we proceed as follows. For each effective plasma length L , we form a (y, z) density projection of the proton bunch and reconstruct the slice centroid trajectory $y_c(z)$ via a centre-of-mass estimate in thin longitudinal bins (red traces in Fig. 5.6). In Fig. 5.5, we can see that as L increases, the centroid trajectory develops a progressively larger slice-to-slice misplacement, in good alignment with growth of a hosing-like transverse asymmetry in a seeded configuration as self-modulation develops.

Two caveats should be kept in mind when reading Fig. 5.5. First, the convective envelope of Eq. (5.5) is obtained from a linearised analysis of the coupled centroid and envelope equations in the overdense long beam regime, in the asymptotic limit of large N_h [103]. Both ends of the plasma length range probed by the plunger scan sit at the edge of this regime. At the smallest values of L the argument N_h is not yet large enough for the asymptotic envelope to be accurate, and at the largest values of L the streak images of Fig. 5.4 already show a fully developed microbunch train, indicating that the self modulation has saturated and that the linear assumption underlying the coupled mode analysis no longer strictly holds. The fit in Fig. 5.5 should therefore be read as a consistency check on the functional form of the convective envelope rather than as a quantitative validation of the linear growth rate.

Second, the streak camera sits at a fixed position $\Delta x_D = 3.5$ m downstream of the nominal plasma exit, so when the plunger reduces L the vacuum drift between the plasma exit and the screen increases by the same amount. During this drift, the transverse momentum

acquired inside the plasma is converted into additional spatial displacement, and the centroid amplitude at the screen is in general larger than the amplitude at the plasma exit by a factor that depends on L . This effect cannot be removed from the experimental data alone, since the diagnostic only sees the centroid after the drift. Both points are addressed quantitatively in Sec. 5.3, where the QuickPIC simulations give access to the centroid both at the plasma exit and after ballistic transport to the diagnostic plane.




FIGURE 5.5: Centroid displacement extracted from the plunger scan streak camera data, in the seeded misalignment configuration, as a function of the effective plasma length L . Black circles show the mean of the slice resolved centroid envelope $|y_c(\zeta)|$ averaged over the full longitudinal extent of the proton bunch. Grey squares show the root mean square of the slice centroid $y_c(\zeta)$ over the same range, which is less sensitive to the sign of the displacement. The red curve is a one parameter fit of the convective hosing envelope of Eq. (5.4), $|y_c| \propto \exp(N_h)/\sqrt{N_h}$ with $N_h = AL^{2/3}$, to the mean envelope points.

5.3 SIMULATIONS USING THE QUASI-STATIC CODE QUICKPIC

In this section, we use the quasi-static PIC code QuickPIC [36] to model the AWAKE configuration in which the HI is intentionally seeded by injecting an electron bunch with a transverse offset with respect to the proton driver. Carrying out such studies with a standard fully electromagnetic PIC code (e.g. OSIRIS) is computationally demanding in the AWAKE parameter regime: the driver is a long proton bunch propagating through metre-scale plasma, and in the seeded-misalignment scenario considered here the electron bunch can be displaced longitudinally by up to 22 mm from the proton beam, which further increases the longitudinal extent required for the simulation domain. A commonly adopted workaround is to replace the proton driver with lighter particles to accelerate the instability's growth and shorten the simulated propagation distance; however, the resulting HI dynamics are not strictly equivalent to the proton-driven case.

Here, we aim to quantify the HI growth for a proton driver, to isolate the role of the vacuum propagation up to the streak camera, which the experiment cannot do, and to follow

the centroid evolution in a regime where the linear assumptions underlying Eq. (5.5) can be checked directly. For these reasons, we adopt the quasi-static approach of QuickPIC introduced in Sec. 1.4.2, which exploits the separation between the fast plasma response and the slow beam evolution to enable three-dimensional simulations over the relevant propagation lengths at a much lower cost than a fully electromagnetic PIC code. Since HI corresponds to a symmetry-breaking transverse mode, its onset and subsequent evolution cannot be captured in two-dimensional (axisymmetric) simulations and require fully three-dimensional modeling.

5.3.1 SIMULATION SETUP

We now summarize the simulation setup to reproduce the experimental results presented in Sec. 5.2. Table 5.2 summarises the baseline parameters used in the QuickPIC simulations of the seeded SMI-HI configuration.

At the entrance of the plasma ($t = 0$), the electron bunch is injected with a longitudinal delay $\Delta z = 16$ cm with respect to the proton bunch and with a transverse offset $\Delta z = 0.56$ mm,

TABLE 5.2: Baseline parameters for the AWAKE SMI-HI simulations performed with QuickPIC.

Group	Parameter	Value
Proton driver		
Population	N_p	1.07×10^{11}
Energy	E_p	400 GeV
RMS length	$\sigma_{z,p}$	4.35 cm
RMS radius	$\sigma_{\perp,p}$	0.2 mm
Seeding electron bunch		
Population	N_e	10^9
Energy	E_e	19 MeV
RMS length	$\sigma_{z,e}$	0.580 mm
RMS radius	$\sigma_{\perp,e}$	0.470 mm
Longitudinal delay at $t = 0$	Δx_1	16 cm
Transverse offset at $t = 0$	Δx_2	0.56 mm
Plasma and diagnostics		
Plasma density	n_0	10^{14} cm^{-3}
Plasma radius	r_p	22 mm
Plasma length (max.)	L_p	10.3 m
Distance to detector	Δx_D	3.5 m
Numerical set-up		
Moving window		on
Simulation box [mm]		$z \in [0, 226], x \in [-22, 22], y \in [-22, 22]$

which seeds the symmetry-breaking hosing mode in the plane of misalignment. In the experiment, the real longitudinal delay is of 22 cm. A longitudinal-offset scan was performed to assess the impact of enlarging the simulation domain. Within the range of offsets explored, no significant change in the wakefield structure was observed that would justify the increased computational cost associated with a larger transverse box.

The plasma is modeled as a uniform column, and we consider propagation lengths up to $L_p = 10.3$ m to match the longest effective plasma lengths of the plunger scans. To enable a direct comparison with the streak-camera measurements at the plasma exit, the simulated proton phase space is additionally transported through a vacuum drift of $\Delta x_D = 3.5$ m (free streaming) to the detector plane. All simulations are performed in a moving window.

5.3.2 CENTROID GROWTH VERSUS EFFECTIVE PLASMA LENGTH IN SIMULATION

To facilitate a rigorous comparison between the experimental plunger-scan measurements and our numerical model, we applied a post-processing workflow to the simulation data that mirrors the experimental analysis. Specifically, for each effective plasma length L , we generated a (y, z) (plane of the misalignment) density projection of the proton bunch and reconstructed the slice centroid trajectory $y_c(z)$ using a centre-of-mass estimator within narrow longitudinal bins (indicated by the red traces in Fig. 5.6). Qualitatively, as L increases, the centroid trajectory exhibits progressively larger slice-to-slice displacements, capturing the characteristic growth of a hosing-like transverse asymmetry in the seeded configuration.

As with the streak-camera measurements, the interpretation of plunger-scan data must account for the beam transport downstream of the plasma. Because the diagnostics are located at a fixed position, the transverse structure observed at the screen differs from the distribution at the plasma exit: each particle's position at the screen is determined by both its initial transverse position and its transverse momentum at the exit, since these latter map into additional displacements during the drift. Consequently, when benchmarking against numerical models, it is essential to distinguish between (i) the distribution at the end of the interaction length and (ii) the propagated distribution at the measurement screen. In the comparisons that follow, we address this by explicitly modeling the ballistic transport to the diagnostic plane to ensure a rigorous one-to-one comparison.

A critical aspect of the plunger-scan geometry is that the plunger scan is the fixed position of the diagnostic plane. Consequently, any reduction in effective plasma length results in a corresponding increase in the vacuum drift distance. To faithfully model this, we propagate the simulated particles ballistically from the plasma exit to the common observation plane at $L = 10.3$ m via the transformation $y_{\text{obs}} = y_{\text{exit}} + (p_y/p_z)\Delta L$ (see bottom row of Fig. 5.6). This free-streaming step is critical, as the drift converts transverse momentum into spatial displacement, effectively amplifying the apparent modulation and centroid excursion recorded by the streak camera.

Furthermore, to compare plunger-scan measurements to theory and to simulations, it is useful to reduce the centroid trajectory $x_c(\xi)$ to a single scalar observable for each plasma




FIGURE 5.6: Simulated proton-bunch density projections in the (y, z) plane for a seeded-misalignment configuration at different effective plasma lengths (panel labels). Greyscale shows the particle-density histogram, and the red curve is the reconstructed transverse centroid $y_c(z)$ obtained from a centre-of-mass estimate in thin longitudinal bins. Top row: distributions at the end of the plasma (i.e. at the selected effective length). Bottom row: the same particles ballistically propagated through vacuum until the streak camera position.

length. From the slice centroids extracted as in the bottom row of Fig. 5.6, we define a centroid-displacement metric $(\Delta y_c)_L$ as the typical centroid deviation from the global bunch centre. Figure 5.5 shows the centroid displacement metric $(\Delta y_c)_L$ as a function of the effective plasma length L , evaluated both at the plasma exit (black circles) and after free streaming the simulated particles through the residual vacuum drift to the streak camera position (grey squares). The red curve is the convective r

5.3.3 DEPENDENCE ON THE SEEDING-BEAM MISALIGNMENT

To assess how the initial misalignment of the external seed influences the transverse dynamics, we performed a parameter scan using QuickPIC, varying the electron-bunch offset Δy while maintaining all other beam and plasma parameters constant. For each configuration, we tracked the proton-bunch transverse centroid y_c as a function of the propagation coordinate z .

Figure 5.8 compares the resulting centroid trajectories for offsets spanning three orders of magnitude. For very small offsets, the centroid excursion remains weak, implying the




FIGURE 5.7: Centroid-displacement metric extracted from the simulation data plunger-scan (seeded-misalignment configuration) as a function of effective plasma length L . Black circles are evaluated at the plasma exit, grey squares after free streaming to the streak camera position. Each point is the mean of the slice resolved centroid envelope $|y_c(\zeta)|$ along the bunch, and the bar is its standard deviation along ζ . The red and blue dashed curves are fits of the convective hosing envelope of Eq. (5.4) to the plasma exit and streak camera points respectively, with fitted growth coefficients A reported in the legend.




FIGURE 5.8: QuickPIC simulation of the transverse centroid evolution in a seeded-hosing configuration, where the hosing seed is provided by injecting an electron bunch with an initial transverse offset Δy relative to the proton-bunch axis. The plot shows the proton-bunch centroid $y_c(z)$ as a function of propagation coordinate z , for the set of offsets indicated in the legend. All coordinates are expressed in plasma units (c/ω_p), and all parameters other than the seed offset are kept fixed across the scan.

seed is insufficient to drive a coherent dipole component of the wakefield before intrinsic self-modulation dominates. Intermediate offsets, however, yield the largest coherent oscillations, indicating efficient coupling of the misaligned seed to the transverse dipole wake. Finally, for very large offsets, the centroid response decreases again, suggesting a reduced ef-




FIGURE 5.9: QuickPIC simulations of the seeded-hosing response to varying electron-bunch transverse offset Δy relative to the proton-bunch axis. **(a)** Transverse centroid evolution $y_c(z)$ of the proton bunch as a function of propagation coordinate z , for the set of offsets indicated in the legend. The maximum centroid misalignment reached during propagation is extracted for each offset value. **(b)** Black curve, transverse lineout of the wakefield component $W_y \equiv E_y - cB_x$ along the y axis at the longitudinal position of the proton bunch axis of symmetry and at propagation coordinate $z = 3$ m (left axis), showing the dipole structure of the seed driven wakefield in the plane of the imposed misalignment. Coloured stars, maximum centroid displacement y_c from panel (a) for the offsets indicated in the legend, plotted at the corresponding Δy position along the horizontal axis (right axis). All coordinates are expressed in plasma units (c/ω_p), and all parameters other than the seed offset are kept fixed across the scan.

fective coupling as the seed perturbs the wake too far off-axis relative to the proton bunch core. These trends align with theoretical expectations: the seeded dipole component must be strong enough to dominate early dynamics, but sufficiently overlapped with the driver to exert a net transverse force [103].

To gain further physical insight, we examine how the proton beam has evolved after 3 m of propagation in plasma and compare the centroid locations across the offset scan (Fig. 5.9 (a)). We then overlay the transverse lineout of the wakefield component $W_y \equiv E_y - cB_x$, taken at the slice where the centroid is evaluated, together with the centroid positions extracted for

the different Δy cases (Fig. 5.9 (b)). W_y is the transverse force per unit charge that a proton experiences in the misalignment plane, so its sign and amplitude as a function of y encode the dipole structure of the seeded wakefield. For each seed offset Δy , the maximum centroid displacement y_c reached during propagation falls on the value of W_y at that position, supporting the interpretation that the observed non monotonic dependence on Δy reflects how efficiently the imposed seed establishes the dipole wake experienced by the proton bunch. In other words, the misalignment scan provides an indirect measure of transverse wake symmetry breaking. Intermediate offsets maximise the effective dipole drive, whereas very small or very large offsets reduce the coherent transverse forcing and therefore the centroid excursion.

5.4 CONCLUSIONS AND FUTURE DIRECTIONS

The results presented in this chapter establish a controlled framework, both experimental and numerical, to study the seeded co-evolution of self-modulation and hosing in AWAKE. Several natural extensions follow from this work, spanning theoretical modeling, experimental diagnostics, and simulation campaigns.

A first natural direction concerns the theoretical description of the centroid envelope. The convective envelope of Eq. (5.5) describes the centroid growth at the plasma exit, but the streak camera measures the bunch after a vacuum drift Δx_D that itself depends on the effective plasma length L . The QuickPIC simulations of Sec. 5.3 reveal a systematic difference between the growth coefficients fitted at the plasma exit ($A = 0.76 \text{ m}^{-2/3}$) and at the screen ($A = 0.56 \text{ m}^{-2/3}$), highlighting the role of ballistic transport. A natural next step is to extend the linear analysis of Ref. [103] to include free-streaming after the plasma, deriving a closed-form prediction for the centroid envelope at a fixed downstream plane as a function of L . Such a model would allow the streak-camera data to be used as a calibrated probe of the in-plasma centroid, rather than as a proxy.

On the experimental side, the plunger-scan analysis of Sec. 5.2.3 relies on the streak camera as the primary diagnostic of seeded HI growth. The Run 2b setup (Fig. 5.3) also includes the halo camera, which images defocused protons and carries information about the transverse wakefield amplitude rather than the centroid [110, 114]. Cross-correlating halo radius and streak centroid across the same plunger scan would provide an independent handle on the SMI–HI co-evolution and a direct test of the coupled-mode prediction that both instabilities grow with comparable rates.

Together, these possible directions would turn the controlled seeded-misalignment configuration introduced in this chapter into a quantitative platform for studying the co-evolution of SMI and HI in AWAKE. This work could bridge the linear-coupled-mode predictions, the Run 2b diagnostics, and the requirements for stable single-stage proton-driven acceleration.

Take-home messages

- **For resonant wakefield AWAKE we need self-modulation instability.** With $\sigma_{z,p} \gg \lambda_p$, the proton driver must self-modulate into microbunches at λ_p to drive a resonant wake. Electron-bunch seeding fixes the modulation phase and improves shot-to-shot reproducibility.
- **A controlled transverse seed makes HI develop in the misalignment plane.** A transverse offset between the seed electron bunch and the proton axis introduces a dipole wake component and seeds a hosing-like centroid evolution while SMI develops.
- **In the overdense long-bunch regime, HI is convective.** The centroid amplitude follows the convective envelope

$$|x_c| \propto \frac{\exp(N_h)}{\sqrt{N_h}}, \quad N_h = \frac{\sqrt[3]{3}}{4} \left(\mu \hat{k}_\beta^2 |\zeta| \hat{z}^2 \right)^{1/3},$$

implying (up to $\mathcal{O}(1)$ profile factors)

$$N_h \propto \left[(n_{b0}/n_0) \gamma^{-1} k_p^3 |\zeta| L^2 \right]^{1/3}.$$

This is the reference scaling used to interpret the measured/simulated centroid growth with effective plasma length.

- **A single scalar observable enables a clean simulation-experiment comparison.** From streak-camera (or simulated) (y, z) projections, reconstruct the slice centroid $y_c(z)$ with a centre-of-mass estimator in thin longitudinal bins, and reduce it to a length-dependent metric $\langle |\Delta y_c| \rangle_L$ for direct comparison to the convective model.
- **Simulations of plunger scans must account for vacuum propagation.** Because the diagnostic screen is at a fixed location, the downstream vacuum drift inherently varies with the effective plasma length L . As demonstrated by comparing the beam profiles at the plasma exit to those at the detector plane, this ballistic transport is critical. It converts transverse momentum into spatial displacement, significantly amplifying the measured centroid excursion. The trend of the instability growth is still conserved.

CHAPTER 6

CONCLUSIONS

Plasma wakefield acceleration offers ultra-high gradients and the prospect of compact accelerators, yet translating this promise into robust, single-stage energy gain is limited by two recurring constraints: driver evolution (specifically depletion and self-consistent reshaping) and the loss of synchronism between the accelerating phase and the witness beam. This thesis addresses these limitations along two complementary directions. First, we explored driver engineering in laser-driven wakefield acceleration using velocity-controlled spatio-temporal (ST) pulses, in which the intensity peak propagates at a prescribed focal velocity. Second, we studied long-distance acceleration with high-energy proton drivers in the context of the AWAKE experiment, where wakefield formation relies on self-modulation and is intrinsically coupled to transverse instabilities such as hosing.

To enable the study of velocity-controlled spatio-temporal (ST) pulses, we first established a robust numerical framework detailed in Chapter 2. Because subluminal ST drivers require broad spectral support, large-angle propagation components, and exhibit an intrinsic envelope-focus slippage, conventional paraxial envelope models are insufficient. To address this, we utilized a Maxwell-consistent spectral initialization in the particle-in-cell code OSIRIS, identifying explicit geometric criteria to mitigate unphysical spatial replicas induced by discrete Fourier sampling. Furthermore, we derived analytical scaling laws to match the pulse dimensions to the plasma response, ensuring resonant wakefield excitation. Finally, to overcome the severe computational constraints imposed by the substantial transverse expansion of the pulse and its focus-envelope slippage, we implemented a continuous wall-injection scheme. This approach dynamically supplies the required angular spectrum at the transverse boundaries, preserving the correct interaction physics while substantially reducing the necessary transverse domain size and overall computational cost.

Building upon this numerical foundation, a central outcome of this work is that controlling the wake's phase velocity provides a practical route to accelerate particles that are not initially ultra-relativistic. We have shown that flying focus pulses combined with carefully tailored density profiles can accelerate initially subrelativistic decaying particles (*e.g.*, muons, pions, and kaons) to relativistic energies over remarkably short distances. We have developed a model to calculate the ideal plasma profile by matching the wake's velocity to the beam's

evolving speed. Our results were validated through numerical simulations, resulting in an excellent agreement in predicting the beam energy gain. Because the scheme operates with few-joule drivers, exploring plasma-based acceleration of decaying particles is within reach and represents an application of plasma acceleration that might not have an analog in standard RF accelerators. This approach may serve as a rapid booster stage for decaying particles in muon-collider designs or as a compact source of high-energy pions, kaons, and muons for neutrino physics [72], opening a range of scientific and societal applications. The scheme's versatility could enable its application to improve electron acceleration in plasma wakefields and to isotope separation.

Beyond longitudinal phase control, we assess the preservation of beam quality and the experimental feasibility of the required density profiles. We demonstrate that optimized beam loading is a critical ingredient for simultaneously maximizing energy-transfer efficiency and maintaining beam quality: by injecting an optimal witness charge, the accelerating field can be flattened, minimizing energy spread without compromising wake stability. Furthermore, we analyze transverse dynamics within the tapered density to ensure that the aggressive gradients required for phase-locking do not compromise transverse confinement. Finally, we bridge theory and experiment by showing that the ideal theoretical density profiles can be effectively approximated using realistic multi-jet gas configurations, reproducing the essential phase-locking behavior with minimal deviation.

The second major contribution of this thesis is a quantitative study of the co-development of the self-modulation instability and transverse hosing in proton-driven wakefield accelerators. In the AWAKE experiment, the long proton bunch must microbunch at the plasma wavelength to resonantly excite a wake; however, this process is intrinsically sensitive to symmetry-breaking perturbations. Using Run 2B diagnostics with a seeded misalignment, we establish a robust centroid-based observable extracted from time-resolved images to characterize the growth of transverse motion during plunger scans (varying plasma length). Crucially, we demonstrate that accurately benchmarking simulations against experiment requires accounting for post-plasma ballistic drift, in which free-streaming converts transverse momentum into amplified spatial displacements at the detector. Complementary quasi-static QuickPIC simulations reproduce the experimental conditions, confirming that SMI and hosing can coexist under controlled seeding, with centroid dynamics strictly governed by the plane and amplitude of the initial perturbation.

Taken together, the results of this thesis emphasize that extending single-stage plasma acceleration is not achieved by a single fix, but rather by achieving a better control of the *evolution* of the driver-wake-witness system. In the laser-driven case, this control is achieved through spatio-temporal engineering of the driver and density tailoring of the plasma to actively manage phase dynamics. In the proton-driven case, it is achieved by seeding and diagnosing the instabilities that enable wake formation, while quantifying and constraining the transverse dynamics that can disrupt stable acceleration. These two threads converge on a common theme: *phase-space control*—both longitudinal and transverse—is a decisive ingredient that turns high gradient into usable, high-quality energy gain.

APPENDIX A

MACHINE-LEARNING SURROGATES IN PIC CODES

This appendix briefly summarises a work brought on in parallel to the main PhD focus of thesis that resulted in two publications [115, 116]. Both works address a common computational challenge in particle-in-cell (PIC) simulations: the evaluation and sampling of probabilistic micro-physics (e.g., collisional or QED processes) inside Monte-Carlo (MC) modules. The central idea is to use compact neural-network (NN) surrogates to replace expensive analytical expressions or memory-heavy look-up tables, while maintaining the accuracy needed for production-scale simulations. The first work [115] demonstrates the workflow to integrate NNs in run-time into PIC codes, using the Compton Scattering cross section as a proof of concept. The second work [116] applies this approach to the Bethe-Heitler pair production process, showing that NN surrogates can achieve comparable accuracy to traditional methods while significantly reducing memory requirements. These results suggest that NN surrogates could be a powerful tool for enhancing the efficiency and practicality of MC modules in PIC simulations, enabling more complex and realistic modeling of plasma processes.

A.1 WORKFLOW AND INITIAL COMPTON SCATTERING PROOF OF CONCEPT

In this work. [115], we developed and demonstrated for the first time a workflow to integrate NNs in run-time into PIC codes. The workflow is designed to be modular and flexible, allowing for the training of NN models to be performed in Python, whereas the inference is performed at run time inside the PIC code, which is typically written in Fortran. The motivation is that ML development (data handling, training, validation) is most efficient in Python, whereas large-scale PIC simulations rely on highly optimised Fortran kernels. To bridge these ecosystems, the workflow exports the trained NN parameters and loads them at run time in the PIC code.

As a proof of concept, the NN surrogate was employed to replace the Compton Scattering




FIGURE A.1: Schematic workflow to integrate neural-network surrogates in a Fortran PIC code: (i) data-set generation (from existing PIC routines), (ii) training in a Python/Keras environment, (iii) export of NN parameters, and (iv) run-time import and inference inside the PIC simulation.

cross section calculation. Training data were generated from controlled PIC runs by logging the analytical probability along with the relevant particle inputs, and the NN was trained to reproduce this target over the range of interest. The resulting NN-based approach accurately reproduced the photon spectra obtained with the conventional algorithm and showed a slightly faster run-time.

A.2 NEURAL-NETWORK SAMPLING OF BETHE-HEITLER PAIR PRODUCTION IN PIC CODES

In this second work [116], a continuation of the one presented in App. A.1, we applied neural networks to improve the practicality of Monte-Carlo implementations of the Bethe-Heitler electron-positron pair-production process in PIC codes. Standard MC modules often rely on pre-computed tables for total and differential cross-sections. However, these tables become rapidly impractical as models grow in complexity and dimensionality. The approach replaces (a) the total Bethe-Heitler cross-section and (b) the inverse cumulative distribution function used to sample the emitted positron energy by compact dense neural networks.




FIGURE A.2: Example validation of NN-based sampling of Bethe-Heitler pair production in a production-like simulation: final positron spectrum comparing the standard look-up-table approach and the neural-network approach. The agreement illustrates that NN sampling can reproduce the reference spectra while using a much smaller model memory wise.

The NN models were trained over a wide parameter range relevant to laser-plasma applications, covering multiple atomic numbers and photon energies. A key result is that the NN surrogates can reach an accuracy comparable to the table approach (Fig. A.2) while reducing the memory footprint by around two orders of magnitude, possibly facilitating the introduction of more sophisticated cross-section models in future MC modules. The implementation was validated in an idealised benchmark (mono-energetic photons incident on a solid target) and in a representative production-like laser-foil interaction simulation, showing that positron spectra obtained with NN sampling closely match the table-based reference.

BIBLIOGRAPHY

- [1] E. I. Simakov, V. A. Dolgashev, and S. G. Tantawi. “Advances in high gradient normal conducting accelerator structures”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 907 (2018), pp. 221–230.
- [2] S. Gourlay, T. Raubenheimer, and V. Shiltsev. “Challenges of future accelerators for particle physics research”. In: *Frontiers in Physics* 10 (2022), p. 920520.
- [3] T. Tajima and J. M. Dawson. “Laser Electron Accelerator”. In: *Physical Review Letters* 43.4 (1979), pp. 267–270. DOI: 10.1103/PhysRevLett.43.267.
- [4] 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”. In: *Physical Review Letters* 54 (1985), pp. 693–696.
- [5] E. Esarey, C. B. Schroeder, and W. P. Leemans. “Physics of laser-driven plasma-based electron accelerators”. In: *Reviews of Modern Physics* 81 (2009), pp. 1229–1285.
- [6] W. P. Leemans, A. J. Gonsalves, H.-S. Mao, K. Nakamura, C. Benedetti, C. B. Schroeder, C. Tóth, J. Daniels, D. E. Mittelberger, S. S. Bulanov, J.-L. Vay, C. G. R. Geddes, and E. Esarey. “Multi-GeV Electron Beams from Capillary-Discharge-Guided Subpetawatt Laser Pulses in the Self-Trapping Regime”. In: *Physical Review Letters* 113.24 (2014), p. 245002. DOI: 10.1103/PhysRevLett.113.245002.
- [7] A. J. Gonsalves, K. Nakamura, J. Daniels, C. Benedetti, C. Pieronek, T. C. H. de Raadt, S. Steinke, J. H. Bin, S. S. Bulanov, J. van Tilborg, C. G. R. Geddes, C. B. Schroeder, C. Tóth, E. Esarey, and W. P. Leemans. “Petawatt Laser Guiding and Electron Beam Acceleration to 8 GeV in a Laser-Heated Capillary Discharge Waveguide”. In: *Physical Review Letters* 122.8 (2019), p. 084801. DOI: 10.1103/PhysRevLett.122.084801.
- [8] H. Milchberg, T. Clark, C. Durfee III, T. Antonsen, and P. Mora. “Development and applications of a plasma waveguide for intense laser pulses”. In: *Physics of Plasmas* 3.5 (1996), pp. 2149–2155.
- [9] H. E. Kondakci and A. F. Abouraddy. “Optical space-time wave packets having arbitrary group velocities in free space”. In: *Nature communications* 10.1 (2019), p. 929.
- [10] D. H. Froula, D. Turnbull, A. S. Davies, T. J. Kessler, D. Haberberger, J. P. Palastro, S.-W. Bahk, I. A. Begishev, R. Boni, S. Bucht, et al. “Spatiotemporal control of laser intensity”. In: *Nature photonics* 12.5 (2018), pp. 262–265.

- [11] E. Adli, A. Ahuja, O. Apsimon, et al. "Acceleration of electrons in the plasma wakefield of a proton bunch". In: *Nature* 561.7723 (2018), pp. 363–367. DOI: 10.1038/s41586-018-0485-4.
- [12] I. Blumenfeld, C. E. Clayton, F.-J. Decker, M. J. Hogan, C. Huang, R. Ischebeck, R. Iverson, C. Joshi, T. Katsouleas, N. Kirby, W. Lu, K. A. Marsh, W. B. Mori, P. Muggli, R. H. Siemann, and D. Walz. "Energy doubling of 42 GeV electrons in a metre-scale plasma wakefield accelerator". In: *Nature* 445.7129 (2007), pp. 741–744. DOI: 10.1038/nature05538.
- [13] M. Litos, E. Adli, W. An, C. I. Clarke, C. E. Clayton, S. Corde, J.-P. Delahaye, R. J. England, A. S. Fisher, J. Frederico, S. Gessner, S. Z. Green, M. J. Hogan, C. Joshi, W. Lu, K. A. Marsh, W. B. Mori, P. Muggli, N. Vafaei-Najafabadi, et al. "High-efficiency acceleration of an electron beam in a plasma wakefield accelerator". In: *Nature* 515.7525 (2014), pp. 92–95. DOI: 10.1038/nature13882.
- [14] C. A. Lindstrøm et al. "Energy-Spread Preservation and High Efficiency in a Plasma-Wakefield Accelerator". In: *Physical Review Letters* 126.1 (2021), p. 014801. DOI: 10.1103/PhysRevLett.126.014801.
- [15] R. D'Arcy et al. "Recovery time of a plasma-wakefield accelerator". In: *Nature* 603.7899 (2022), pp. 58–62. DOI: 10.1038/s41586-021-04348-8.
- [16] C. A. Lindstrøm et al. "Emittance preservation in a plasma-wakefield accelerator". In: *Nature Communications* 15 (2024), p. 6097. DOI: 10.1038/s41467-024-50320-1.
- [17] R. W. Assmann, M. K. Weikum, T. Akhter, D. Alesini, A. S. Alexandrova, M. P. Anania, N. E. Andreev, I. Andriyash, M. Artioli, A. Aschikhin, et al. "EuPRAXIA Conceptual Design Report". In: *The European Physical Journal Special Topics* 229.24 (2020), pp. 3675–4284. DOI: 10.1140/epjst/e2020-000127-8.
- [18] R. Pompili, D. Alesini, M. P. Anania, et al. "Free-electron lasing with compact beam-driven plasma wakefield accelerator". In: *Nature* 605.7911 (2022), pp. 659–662. DOI: 10.1038/s41586-022-04589-1.
- [19] X.-N. Wang, X.-F. Lan, Y.-S. Huang, H. Zhang, and T.-P. Yu. "Rapid Muon Acceleration in Beam-Driven Plasma Wakefield". In: *Phys. Plasmas* 28 (2021), p. 093101.
- [20] M. Labat et al. "Seeded free-electron laser driven by a compact laser plasma accelerator". In: *Nature Photonics* 17 (2023), pp. 150–156. DOI: 10.1038/s41566-022-01104-w.
- [21] Y. Zhang et al. "Plasma-wakefield accelerator simultaneously boosts electron beam energy and brightness". In: *Nature Communications* (2025). DOI: 10.1038/s41467-025-65742-8.
- [22] W. Lu. *Nonlinear plasma wakefield theory and optimum scaling for laser wakefield acceleration in the blowout regime*. PhD Thesis, University of California, Los Angeles. 2006.
- [23] M. Tzoufras, W. Lu, F. S. Tsung, C. Huang, W. B. Mori, T. Katsouleas, J. Vieira, R. A. Fonseca, and L. O. Silva. "Beam Loading in the Nonlinear Regime of Plasma-Based Acceleration". In: *Physical Review Letters* 101.14 (2008), p. 145002. DOI: 10.1103/PhysRevLett.101.145002.

- [24] S. Steinke, J. van Tilborg, C. Benedetti, et al. "Multistage coupling of laser-wakefield accelerators with compact beam transport". In: *Nature* 530 (2016), pp. 190–193. DOI: 10.1038/nature16525.
- [25] F. Zimmermann, S. Moebius, S. Kraus, G. Franchetti, C. Badiali, A. Pukhov, Y. Belouard, Y. Yang, R. D. Motlagh, D. Schulte, V. Shiltsev, R. Ischebeck, M. Krasny, U. Niedermayer, P. Crivelli, and T. Latychevskaia. "Advanced accelerator concepts for dark sector searches and fast muon acceleration". English. In: *Proc. IPAC'24* (Nashville, TN). IPAC'24 - 15th International Particle Accelerator Conference 15. JACoW Publishing, Geneva, Switzerland, May 2024, pp. 489–492. ISBN: 978-3-95450-247-9. DOI: 10.18429/JACoW-IPAC2024-MOPR17. URL: <https://jacow.org/ipac2024/doi/JACoW-IPAC2024-MOPR17.html>.
- [26] K. R. Long, D. Lucchesi, M. A. Palmer, N. Pastrone, D. Schulte, and V. Shiltsev. "Muon colliders to expand frontiers of particle physics". In: *Nature Physics* 17.3 (2021), pp. 289–292.
- [27] R. B. Palmer. "Muon Colliders". In: *Reviews of Accelerator Science and Technology* 7 (2014), pp. 137–159.
- [28] J.-P. Delahaye, M. Diemoz, K. Long, M. A. Palmer, N. Pastrone, D. Schulte, and V. Shiltsev. "Muon Colliders". In: *arXiv* (2019). eprint: 1901.06150 (physics.acc-ph).
- [29] D. Stratakis, D. Bowring, J. Johnstone, M. Chung, K. Yonehara, J. K. S. T. Hassan, A. K. Jain, R. P. Johnson, D. Neuffer, P. Snopok, and M. A. Palmer. "Accelerator performance analysis of the Fermilab Muon Campus". In: *Phys. Rev. Accel. Beams* 20 (2017), p. 111003.
- [30] A. Cervera, A. Donini, M. B. Gavela, J. J. G. Cadenas, P. Hernandez, O. Mena, and S. Rigolin. "Golden measurements at a neutrino factory". In: *Nucl. Phys. B* 579 (2000), pp. 17–55.
- [31] M. Apollonio, A. Blondel, et al. "Oscillation Physics with a Neutrino Factory". In: *arXiv* (2002). eprint: hep-ph/0210192.
- [32] R. A. Fonseca et al. "OSIRIS: A Three-Dimensional, Fully Relativistic Particle in Cell Code for Modeling Plasma Based Accelerators". In: *Computational Science — ICCS 2002*. Ed. by P. M. A. Sloot, A. G. Hoekstra, C. J. K. Tan, and J. J. Dongarra. Berlin, Heidelberg: Springer Berlin Heidelberg, 2002, p. 342. ISBN: 978-3-540-47789-1.
- [33] A. F. Lifschitz, X. Davoine, E. Lefebvre, J. Faure, C. Rechatin, and V. Malka. "Particle-in-cell modelling of laser-plasma interaction using Fourier decomposition". In: *Journal of Computational Physics* 228.5 (2009), pp. 1803–1814.
- [34] A. Davidson, A. Tableman, W. An, F. S. Tsung, W. Lu, J. Vieira, R. A. Fonseca, L. O. Silva, and W. B. Mori. "Implementation of a hybrid particle code with a PIC description in r-z and a gridless description in ϕ into OSIRIS". In: *Journal of Computational Physics* 281 (2015), pp. 1063–1077.
- [35] C. Huang, V. K. Decyk, C. Ren, M. Zhou, W. Lu, W. B. Mori, J. H. Cooley, and T. M. Antonsen. "QuickPIC: A highly efficient particle-in-cell code for modeling wakefield acceleration in plasmas". In: *Journal of Computational Physics* 217 (2006), pp. 658–679.

- [36] W. An, V. K. Decyk, W. B. Mori, and T. M. Antonsen. “An improved iteration scheme for the quasi-static particle-in-cell algorithm”. In: *Journal of Computational Physics* 250 (2013), pp. 165–177.
- [37] T. Tajima. “High energy laser plasma accelerators”. In: *Laser and Particle Beams* 3.4 (1985), pp. 351–413.
- [38] E. Esarey, C. B. Schroeder, and W. P. Leemans. “Physics of laser-driven plasma-based electron accelerators”. In: *Rev. Mod. Phys.* 81 (3 Aug. 2009), pp. 1229–1285. DOI: 10.1103/RevModPhys.81.1229. URL: <https://link.aps.org/doi/10.1103/RevModPhys.81.1229>.
- [39] V. Horný, P. G. Bleotu, D. Ursescu, V. Malka, and P. Tomassini. “Efficient laser wakefield accelerator in pump depletion dominated bubble regime”. In: *Phys. Rev. E* 110 (3 Sept. 2024), p. 035202. DOI: 10.1103/PhysRevE.110.035202. URL: <https://link.aps.org/doi/10.1103/PhysRevE.110.035202>.
- [40] H. Daido, M. Nishiuchi, and A. S. Pirozhkov. “Review of laser-driven ion sources and their applications”. In: *Reports on progress in physics* 75.5 (2012), p. 056401.
- [41] S. Corde, K. Ta Phuoc, G. Lambert, R. Fitour, V. Malka, A. Rousse, A. Beck, and E. Lefebvre. “Femtosecond x rays from laser-plasma accelerators”. In: *Rev. Mod. Phys.* 85 (1 Jan. 2013), pp. 1–48. DOI: 10.1103/RevModPhys.85.1. URL: <https://link.aps.org/doi/10.1103/RevModPhys.85.1>.
- [42] M. Galletti, R. Assmann, M. Couprie, M. Ferrario, L. Giannessi, A. Irman, R. Pompili, and W. Wang. “Prospects for free-electron lasers powered by plasma-wakefield-accelerated beams”. In: *Nature Photonics* 18.8 (2024), pp. 780–791.
- [43] Y. Glinec, J. Faure, L. L. Dain, S. Darbon, T. Hosokai, J. J. Santos, E. Lefebvre, J. P. Rousseau, F. Burgy, B. Mercier, and V. Malka. “High-Resolution γ -Ray Radiography Produced by a Laser-Plasma Driven Electron Source”. In: *Phys. Rev. Lett.* 94 (2 Jan. 2005), p. 025003. DOI: 10.1103/PhysRevLett.94.025003. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.94.025003>.
- [44] T. Tajima, X. Yan, and T. Ebisuzaki. “Wakefield acceleration”. In: *Reviews of Modern Plasma Physics* 4.1 (2020), p. 7.
- [45] A. Debus, R. Pausch, A. Huebl, K. Steiniger, R. Wiedera, T. E. Cowan, U. Schramm, and M. Bussmann. “Circumventing the Dephasing and Depletion Limits of Laser-Wakefield Acceleration”. In: *Phys. Rev. X* 9 (3 Sept. 2019), p. 031044. DOI: 10.1103/PhysRevX.9.031044. URL: <https://link.aps.org/doi/10.1103/PhysRevX.9.031044>.
- [46] A. Sainte-Marie, O. Gobert, and F. Quere. “Controlling the velocity of ultrashort light pulses in vacuum through spatio-temporal couplings”. In: *Optica* 4.10 (2017), pp. 1298–1304.
- [47] D. Froula, J. Palastro, D. Turnbull, A. Davies, L. Nguyen, A. Howard, D. Ramsey, P. Franke, S.-W. Bahk, I. Begishev, et al. “Flying focus: Spatial and temporal control of intensity for laser-based applications”. In: *Physics of Plasmas* 26.3 (2019).

- [48] M. Yessenov, L. A. Hall, K. L. Schepler, and A. F. Abouraddy. "Space-time wave packets". In: *Advances in Optics and Photonics* 14.3 (2022), pp. 455–570.
- [49] A. Howard, D. Turnbull, A. Davies, P. Franke, D. Froula, and J. Palastro. "Photon acceleration in a flying focus". In: *Physical review letters* 123.12 (2019), p. 124801.
- [50] P. Franke, D. Ramsey, T. T. Simpson, D. Turnbull, D. Froula, and J. Palastro. "Optical shock-enhanced self-photon acceleration". In: *Physical Review A* 104.4 (2021), p. 043520.
- [51] C. Caizergues, S. Smartsev, V. Malka, and C. Thaury. "Phase-locked laser-wakefield electron acceleration". In: *Nature Photonics* 14 (2020), pp. 475–479. DOI: 10.1038/s41566-020-0657-2.
- [52] J. Palastro, D. Turnbull, S.-W. Bahk, R. Follett, J. Shaw, D. Haberberger, J. Bromage, and D. Froula. "Ionization waves of arbitrary velocity driven by a flying focus". In: *Physical Review A* 97.3 (2018), p. 033835.
- [53] D. Turnbull, P. Franke, J. Katz, J. Palastro, I. Begishev, R. Boni, J. Bromage, A. Milder, J. Shaw, and D. Froula. "Ionization waves of arbitrary velocity". In: *Physical Review Letters* 120.22 (2018), p. 225001.
- [54] J. P. Palastro, J. L. Shaw, P. Franke, D. Ramsey, T. T. Simpson, and D. H. Froula. "Dephasingless Laser Wakefield Acceleration". In: *Phys. Rev. Lett.* 124 (13 Mar. 2020), p. 134802. DOI: 10.1103/PhysRevLett.124.134802. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.124.134802>.
- [55] M. Ambat, J. Shaw, J. Pigeon, K. Miller, T. Simpson, D. Froula, and J. P. Palastro. "Programmable-trajectory ultrafast flying focus pulses". In: *Optics Express* 31.19 (2023), pp. 31354–31368.
- [56] K. G. Miller, J. R. Pierce, M. V. Ambat, J. L. Shaw, K. Weichman, W. B. Mori, D. H. Froula, and J. P. Palastro. "Dephasingless laser wakefield acceleration in the bubble regime". In: *Scientific Reports* 13.1 (2023), p. 21306.
- [57] A. Ghasemi, S. Mirzanejhad, and T. Mohsenpour. "Flying focus laser wake field acceleration by donut shape pulse". In: *Applied Physics B* 130.6 (2024), p. 105.
- [58] Z. Gong, S. Cao, J. P. Palastro, and M. R. Edwards. "Laser Wakefield Acceleration of Ions with a Transverse Flying Focus". In: *Phys. Rev. Lett.* 133 (2024), p. 265002. DOI: 10.1103/PhysRevLett.133.265002.
- [59] J. R. Pierce, K. G. Miller, F. Li, J. P. Palastro, and W. B. Mori. "Laser wakefield acceleration driven by a discrete flying focus". In: *Physical Review Accelerators and Beams* 28.10 (2025), p. 101303.
- [60] A. Liberman, S. Smartsev, A. Golovanov, S. Tata, I. A. Andriyash, S. Benracassa, E. Y. Levine, E. Kroupp, R. Lahaye, C. Thaury, et al. "Electron acceleration in a wakefield generated by structured light". In: *Laser Acceleration of Electrons, Protons, and Ions VIII*. Vol. 13534. SPIE. 2025, pp. 19–25.

- [61] R. R. de Almeida. “Arbitrarily non-paraxial electromagnetic wave-packets in particle-in-cell codes”. Maxwell-consistent non-paraxial laser injection enabling prescribed focal velocity (flying focus) in PIC. Master’s thesis. Lisbon, Portugal: Instituto Superior Técnico, Universidade de Lisboa, 2023. URL: https://scholar.tecnico.ulisboa.pt/records/Bu-3rpCZiOLBFJbyLb73T_KtviVfTKzV94eB.
- [62] C. Sheppard and S. Saghafi. “Electromagnetic Gaussian beams beyond the paraxial approximation”. In: *Journal of the Optical Society of America A* 16.6 (1999), pp. 1381–1386.
- [63] E. Heyman and L. B. Felsen. “Gaussian beam and pulsed-beam dynamics: complex-source and complex-spectrum formulations within and beyond paraxial asymptotics”. In: *Journal of the Optical Society of America A* 18.7 (2001), pp. 1588–1611.
- [64] D. Ramsey, A. Di Piazza, M. Formanek, P. Franke, D. H. Froula, B. Malaca, W. B. Mori, J. R. Pierce, T. T. Simpson, J. Vieira, M. Vranic, K. Weichman, and J. P. Palastro. “Exact solutions for the electromagnetic fields of a flying focus”. In: *Phys. Rev. A* 107 (1 Jan. 2023), p. 013513. DOI: 10.1103/PhysRevA.107.013513. URL: <https://link.aps.org/doi/10.1103/PhysRevA.107.013513>.
- [65] V. Malka. *CAS Proceedings of the 2019 course on High Gradient Wakefield Accelerators, Sesimbra, Portugal, 11–22 March 2019*. Tech. rep. 2020.
- [66] A. Liberman, A. Golovanov, S. Smartsev, S. Tata, I. A. Andriyash, S. Benracassa, E. Y. Levine, Y. Wan, E. Kroupp, and V. Malka. “Direct observation of a wakefield generated with structured light”. In: *Nature Communications* 16.1 (Dec. 2025), p. 10957. ISSN: 2041-1723. DOI: 10.1038/s41467-025-66056-5. URL: <https://doi.org/10.1038/s41467-025-66056-5>.
- [67] K. R. Long, D. Lucchesi, M. A. Palmer, N. Pastrone, D. Schulte, and V. Shiltsev. “Muon colliders to expand frontiers of particle physics”. In: *Nature Physics* 17.3 (2021), pp. 289–292.
- [68] A. Costantini, F. De Lillo, F. Maltoni, L. Mantani, O. Mattelaer, R. Ruiz, and X. Zhao. “Vector boson fusion at multi-TeV muon colliders”. In: *Journal of High Energy Physics* 2020.9 (2020), pp. 1–59.
- [69] M. Chiesa, F. Maltoni, L. Mantani, B. Mele, F. Piccinini, and X. Zhao. “Measuring the quartic Higgs self-coupling at a multi-TeV muon collider”. In: *Journal of High Energy Physics* 2020.9 (2020), pp. 1–29.
- [70] F. Zimmermann. “Future colliders for particle physics—“Big and small””. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 909 (2018), pp. 33–37.
- [71] M. Collaboration. “Demonstration of cooling by the muon ionization cooling experiment”. In: *Nature* 578.7793 (2020), pp. 53–59.
- [72] A. Apyan, M. W. Krasny, and W. Placzek. “Gamma Factory high-intensity muon and positron source: Exploratory studies”. In: *Phys. Rev. Accel. Beams* 26 (8 Aug. 2023), p. 083401.

- [73] C. Badiali, R. Almeida, B. Malaca, R. Fonseca, T. Silva, and J. Vieira. “Tailored laser wakefield acceleration for decaying particles”. In: *Phys. Rev. Res.* 8 (2 Apr. 2026), p. 023077. DOI: 10.1103/1pp8-qx35. URL: <https://link.aps.org/doi/10.1103/1pp8-qx35>.
- [74] T. Goldman and R. R. Silbar. “Pion LINAC as an Energy-Tagged ν Source”. In: *Phys. Rev. ST Accel. Beams* 11 (2008), p. 124701.
- [75] R. J. Noble. “Particle Production and Survival in Muon Acceleration”. In: *Proc. 1992 APS DPF Meeting*. FERMILAB-CONF-92-165. 1992.
- [76] J. D. Sadler, C. Arran, H. Li, and K. A. Flippo. “Overcoming the dephasing limit in multiple-pulse laser wakefield acceleration”. In: *Physical Review Accelerators and Beams* 23.2 (2020), p. 021303.
- [77] W. Lu, M. Tzoufras, C. Joshi, F. S. Tsung, W. B. Mori, J. Vieira, R. A. Fonseca, and L. O. Silva. “Generating multi-GeV electron bunches using single stage laser wakefield acceleration in a 3D nonlinear regime”. In: *Phys. Rev. ST Accel. Beams* 10 (6 June 2007), p. 061301.
- [78] M. Boscolo, J.-P. Delahaye, and M. Palmer. “The future prospects of muon colliders and neutrino factories”. In: *Reviews of Accelerator Science and Technology* 10.01 (2019), pp. 189–214.
- [79] C. Accettura, D. Adams, R. Agarwal, C. Ahdida, C. Aimè, N. Amapane, D. Amorim, P. Andreetto, F. Anulli, R. Appleby, et al. “Towards a muon collider”. In: *The European Physical Journal C* 83.9 (2023), p. 864.
- [80] S. W. T. Katsouleas and J. Su. “Beam loading efficiency in plasma accelerators”. In: *Part. Accel* 22 (1987), pp. 81–99.
- [81] M. Tzoufras, W. Lu, F. S. Tsung, C. Huang, W. B. Mori, T. Katsouleas, J. Vieira, R. Fonseca, and L. Silva. “Beam loading by electrons in nonlinear plasma wakes”. In: *Physics of Plasmas* 16.5 (2009).
- [82] C. A. Lindstrøm, J. Beinortaitė, J. Björklund Svensson, L. Boulton, J. Chappell, S. Diederichs, B. Foster, J. Garland, P. Gonzalez Caminal, G. Loisch, et al. “Emittance preservation in a plasma-wakefield accelerator”. In: *Nature Communications* 15.1 (2024), p. 6097.
- [83] R. Pompili, D. Alesini, M. P. Anania, M. Behtouei, M. Bellaveglia, A. Biagioni, F. Bisesto, M. Cesarini, E. Chiadroni, A. Cianchi, et al. “Energy spread minimization in a beam-driven plasma wakefield accelerator”. In: *Nature Physics* 17.4 (2021), pp. 499–503.
- [84] D. W. Kerst. “Historical development of the betatron”. In: *Nature* 157.3978 (1946), pp. 90–95.
- [85] S. Corde, K. T. Phuoc, G. Lambert, R. Fitour, V. Malka, A. Rousse, A. Beck, and E. Lefebvre. “Femtosecond x-rays from laser-plasma accelerators”. In: *Rev. Mod. Phys.* 85 (2013), pp. 1–48.
- [86] P. San Miguel Claveria, E. Adli, L. Amorim, W. An, C. Clayton, S. Corde, S. Gessner, M. Hogan, C. Joshi, O. Kononenko, et al. “Betatron radiation and emittance growth in plasma wakefield accelerators”. In: *Philosophical Transactions of the Royal Society A* 377.2151 (2019), p. 20180173.

- [87] Y. Glinec, J. Faure, A. Lifschitz, J. Vieira, R. A. Fonseca, L. Silva, and V. Malka. "Direct observation of betatron oscillations in a laser-plasma electron accelerator". In: *Europhysics Letters* 81.6 (2008), p. 64001.
- [88] J. Vieira, C.-K. Huang, W. B. Mori, and L. O. Silva. "Polarized beam conditioning in plasma based acceleration". In: *Phys. Rev. ST Accel. Beams* 14 (7 July 2011), p. 071303. DOI: 10.1103/PhysRevSTAB.14.071303. URL: <https://link.aps.org/doi/10.1103/PhysRevSTAB.14.071303>.
- [89] C. G. R. Geddes, K. Nakamura, G. R. Plateau, C. Toth, E. Cormier-Michel, E. Esarey, C. B. Schroeder, J. R. Cary, and W. P. Leemans. "Plasma-Density-Gradient Injection of Low Absolute-Momentum-Spread Electron Bunches". In: *Phys. Rev. Lett.* 100 (21 May 2008), p. 215004. DOI: 10.1103/PhysRevLett.100.215004. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.100.215004>.
- [90] M. Hansson, B. Aurand, X. Davoine, H. Ekerfelt, K. Svensson, A. Persson, C.-G. Wahlström, and O. Lundh. "Down-ramp injection and independently controlled acceleration of electrons in a tailored laser wakefield accelerator". In: *Phys. Rev. ST Accel. Beams* 18 (7 July 2015), p. 071303. DOI: 10.1103/PhysRevSTAB.18.071303. URL: <https://link.aps.org/doi/10.1103/PhysRevSTAB.18.071303>.
- [91] D. Ullmann, P. Scherkl, A. Knetsch, T. Heinemann, A. Sutherland, A. F. Habib, O. S. Karger, A. Beaton, G. G. Manahan, A. Deng, G. Andonian, M. D. Litos, B. D. O'Shea, J. R. Cary, M. J. Hogan, V. Yakimenko, J. B. Rosenzweig, and B. Hidding. "All-optical density downramp injection in electron-driven plasma wakefield accelerators". In: *Phys. Rev. Res.* 3 (4 Dec. 2021), p. 043163. DOI: 10.1103/PhysRevResearch.3.043163. URL: <https://link.aps.org/doi/10.1103/PhysRevResearch.3.043163>.
- [92] A. Buck, J. Wenz, J. Xu, K. Khrennikov, K. Schmid, M. Heigoldt, J. M. Mikhailova, M. Geissler, B. Shen, F. Krausz, S. Karsch, and L. Veisz. "Shock-Front Injector for High-Quality Laser-Plasma Acceleration". In: *Phys. Rev. Lett.* 110 (18 May 2013), p. 185006. DOI: 10.1103/PhysRevLett.110.185006. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.110.185006>.
- [93] A. Pukhov and I. Kostyukov. "Control of laser-wakefield acceleration by the plasma-density profile". In: *Phys. Rev. E* 77 (2 Feb. 2008), p. 025401. DOI: 10.1103/PhysRevE.77.025401. URL: <https://link.aps.org/doi/10.1103/PhysRevE.77.025401>.
- [94] C. Wang, H. Tang, and X. Cheng. "Tailoring plasma density for laser plasma accelerators". In: *Journal of Physics D: Applied Physics* 42.23 (2009), p. 235206.
- [95] C. Aniculaesei, V. B. Pathak, H. T. Kim, K. H. Oh, B. J. Yoo, E. Brunetti, Y. H. Jang, C. I. Hojbota, J. H. Shin, J. H. Jeon, et al. "Electron energy increase in a laser wakefield accelerator using up-ramp plasma density profiles". In: *Scientific Reports* 9.1 (2019), p. 11249.
- [96] H. T. Kim, K. H. Pae, H. J. Cha, I. J. Kim, T. J. Yu, J. H. Sung, S. K. Lee, T. M. Jeong, and J. Lee. "Enhancement of Electron Energy to the Multi-GeV Regime by a Dual-Stage Laser-Wakefield Accelerator Pumped by Petawatt Laser Pulses". In: *Phys. Rev. Lett.* 111 (16 Oct. 2013), p. 165002. DOI: 10.1103/PhysRevLett.111.165002. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.111.165002>.

- [97] E. Guillaume, A. Döpp, C. Thaury, K. Ta Phuoc, A. Lifschitz, G. Grittani, J.-P. Goddet, A. Tafzi, S. W. Chou, L. Veisz, and V. Malka. “Electron Rephasing in a Laser-Wakefield Accelerator”. In: *Phys. Rev. Lett.* 115 (15 Oct. 2015), p. 155002. DOI: 10.1103/PhysRevLett.115.155002. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.115.155002>.
- [98] G. Costa, M. Anania, S. Arjmand, A. Biagioni, M. Del Franco, M. Del Giorno, M. Galletti, M. Ferrario, D. Pellegrini, R. Pompili, et al. “Characterisation and optimisation of targets for plasma wakefield acceleration at SPARC_LAB”. In: *Plasma Physics and Controlled Fusion* 64.4 (2022), p. 044012.
- [99] A. Caldwell, K. Lotov, A. Pukhov, and F. Simon. “Proton-driven plasma-wakefield acceleration”. In: *Nature Physics* 5 (2009), pp. 363–367. DOI: 10.1038/nphys1248.
- [100] N. Kumar, A. Pukhov, and K. V. Lotov. “Self-Modulation Instability of a Long Proton Bunch in Plasmas”. In: *Phys. Rev. Lett.* 104 (2010), p. 255003. DOI: 10.1103/PhysRevLett.104.255003. arXiv: 1003.5816.
- [101] J. Vieira, W. Mori, and P. Muggli. “Hosing instability suppression in self-modulated plasma wakefields”. In: *Physical Review Letters* 112.20 (2014), p. 205001.
- [102] L. e. a. Verra. “Controlled Growth of the Self-Modulation of a Relativistic Proton Bunch in Plasma”. In: *Phys. Rev. Lett.* 129 (2 July 2022), p. 024802. DOI: 10.1103/PhysRevLett.129.024802. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.129.024802>.
- [103] C. Schroeder, C. Benedetti, E. Esarey, F. Grüner, and W. Leemans. “Coupled beam hose and self-modulation instabilities in overdense plasma”. In: *Physical Review E—Statistical, Nonlinear, and Soft Matter Physics* 86.2 (2012), p. 026402.
- [104] G. Plyushchev, R. Kersevan, A. Petrenko, and P. Muggli. “A rubidium vapor source for a plasma source for AWAKE”. In: *Journal of Physics D: Applied Physics* 51.2 (2018), p. 025203.
- [105] E. Öz and P. Muggli. “A novel Rb vapor plasma source for plasma wakefield accelerators”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 740 (2014), pp. 197–202.
- [106] A. Caldwell and K. Lotov. “Plasma wakefield acceleration with a modulated proton bunch”. In: *Nuclear Instruments and Methods in Physics Research Section A* 829 (2016), pp. 3–16.
- [107] J. Vieira, Y. Fang, W. Mori, L. Silva, and P. Muggli. “Transverse self-modulation of ultra-relativistic lepton beams in the plasma wakefield accelerator”. In: *Physics of Plasmas* 19.6 (2012).
- [108] V. Bencini, S. Doeberl, P. Muggli, J. Farmer, A. Pukhov, E. Gschwendtner, A. Caldwell, G. Zevi Della Porta, M. Wing, E. Granados, et al. *AWAKE status report 2022*. Tech. rep. 2022.
- [109] C. AWAKE. *AWAKE Status Report 2025*. Tech. rep. 2025.

- [110] M. Turner, A. Petrenko, B. Biskup, S. Burger, E. Gschwendtner, K. V. Lotov, S. Mazzoni, and H. Vincke. “Indirect self-modulation instability measurement concept for the AWAKE proton beam”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 829 (2016), pp. 314–317. DOI: 10.1016/j.nima.2016.01.060. arXiv: 1511.06537 [physics.acc-ph]. URL: <https://doi.org/10.1016/j.nima.2016.01.060>.
- [111] M. Turner, E. Gschwendtner, and P. Muggli. “A method to determine the maximum radius of defocused protons after self-modulation in AWAKE”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 909 (2018). EAAC 2017 NIM A proceedings, pp. 123–125. DOI: 10.1016/j.nima.2018.02.007. arXiv: 1712.00315 [physics.acc-ph]. URL: <https://doi.org/10.1016/j.nima.2018.02.007>.
- [112] S. S. Wyler. “Control of Proton Beam Self-Modulation for AWAKE via Initial Beam Parameters”. CERN-THESIS-2022-107. MA thesis. École Polytechnique Fédérale de Lausanne (EPFL), 2022. URL: <https://cds.cern.ch/record/2824591>.
- [113] M. Turner and P. Muggli. *Understanding time-resolved images of AWAKE proton bunches*. 2024. DOI: 10.48550/arXiv.2409.18771. arXiv: 2409.18771 [physics.plasm-ph]. URL: <https://arxiv.org/abs/2409.18771>.
- [114] M. Turner, E. Adli, A. Ahuja, O. Apsimon, R. Apsimon, A.-M. Bachmann, M. Barros Marin, D. Barrientos, F. Batsch, J. Batkiewicz, et al. “Experimental observation of plasma wakefield growth driven by the seeded self-modulation of a proton bunch”. In: *Physical review letters* 122.5 (2019), p. 054801.
- [115] C. Badiali, P. J. Bilbao, F. Cruz, and L. O. Silva. “Machine-learning-based models in particle-in-cell codes for advanced physics extensions”. In: *Journal of Plasma Physics* 88.6 (2022), p. 895880602. DOI: 10.1017/S0022377822001180.
- [116] Ó. Amaro, C. Badiali, and B. Martinez. “Neural network sampling of Bethe-Heitler process in particle-in-cell codes”. In: *Journal of Computational Physics* 553 (2026), p. 114707.