

POSITRON ACCELERATION IN A PLASMA COLUMN

Dissertation

Zur Erlangung des Doktorgrades an der Fakultät für Mathematik, Informatik und
Naturwissenschaften Fachbereich Physik der Universität Hamburg

vorgelegt von

NIELS SEVERIN DIEDERICHS

Februar 2023

BETREUER:

Dr. Jens Osterhoff
Prof. Dr. Carl B. Schroeder
Prof. Dr. Wolfgang Hillert

GUTACHTER DER DISSERTATION:

Dr. Jens Osterhoff
Prof. Dr. Wolfgang Hillert
Prof. Dr. Jorge Vieira

ZUSAMMENSETZUNG DER PRÜFUNGSKOMMISSION:

Dr. Jens Osterhoff
Prof. Dr. Carl B. Schroeder
Prof. Dr. Wolfgang Hillert
Prof. Dr. Gudrid Moortgat-Pick
Prof. Dr. Daniela Pfannkuche

VORSITZENDE DER PRÜFUNGSKOMMISSION:

Prof. Dr. Daniela Pfannkuche

DATUM DER DISPUTATION:

21.03.2023

VORSITZENDER FACH-PROMOTIONSAUSSCHUSS PHYSIK:

Prof. Dr. Günter H. W. Sigl

LEITER DES FACHBEREICHS PHYSIK:

Prof. Dr. Wolfgang J. Parak

DEKAN DER FAKULTÄT MIN:

Prof. Dr.-Ing. Norbert Ritter

Niels Severin Diederichs: *Positron acceleration in a plasma column*, Dissertation - Zur Erlangung des Doktorgrades an der Fakultät für Mathematik, Informatik und Naturwissenschaften Fachbereich Physik der Universität Hamburg, © Februar 2023

You should sit in meditation for twenty minutes every day – unless
you're too busy; then you should sit for an hour.

— Zen Proverb

Enough of this miserable, whining life. Stop monkeying around!

— Marcus Aurelius, *Meditations*, 9.37

This task is appointed for you, Frodo; and that if you do not find a
way, no one will.

— Elrond, *Lord of the Rings* by J. R. R. Tolkien

ABSTRACT

The particle physics community has expressed significant interest in a 10 TeV-class electron-positron collider to advance our understanding of matter. However, the costs associated with such a collider using conventional radio-frequency technology seem to be prohibitive. Thus, significant research and development is needed on the accelerator side to reduce both construction and power consumption costs.

Plasma-based accelerators, due to their extremely high accelerating gradients, are a potential solution to reduce construction costs. Although significant progress has been made in plasma-based electron acceleration in recent years, positron acceleration still poses a notorious challenge and has been identified as a critical point in the realization of a plasma-based electron-positron collider. Additionally, plasma-based positron acceleration, while preserving the required high beam quality for a collider, is challenging even conceptually and difficult to model with currently available simulation tools.

Positron acceleration in a plasma column is a promising new concept owing to its high accelerating gradient and the ability to preserve positron beam emittance, a fundamental requirement for a collider. In this scheme, a relativistic electron beam travels along the axis of a plasma column, driving a plasma wake in the blowout regime. Due to the finite radius of the plasma, the restoring force acting on the plasma electrons is reduced and the plasma electrons return in a long, high-density electron filament to the axis some distance behind the driver. This electron filament generates high-gradient accelerating and focusing fields for positrons. However, questions regarding stability and beam quality remain.

In this thesis, positron acceleration in a plasma column is put to the test. Beyond proof of concept, the acceleration of high-quality positron beams and its resilience to realistic conditions, such as misalignment is demonstrated. Realistic plasma profiles and temperature effects are also considered and found to be beneficial to the scheme. The results show that positron acceleration in a plasma column is indeed a viable concept for a possible collider application, although additional research is needed to optimize efficiency. Furthermore, the GPU-accelerated, quasi-static Particle-in-Cell code HiPACE++ has been developed. The new code allows for modeling plasma-based positron acceleration scenarios in full 3D with previously unattainable precision while drastically reducing computational costs, enabling the studies presented. An outlook for the next possible steps towards a plasma-based collider is given.

ZUSAMMENFASSUNG

Die Gemeinschaft der Teilchenphysik hat großes Interesse an einem Hochenergie-Elektron-Positron-Kollider geäußert, um das Verständnis von Materie zu erweitern. Die damit verbundenen Kosten scheinen jedoch unerschwinglich zu sein. Um sowohl Bau- als auch Stromverbrauchskosten zu reduzieren, ist daher erhebliche Forschung und Entwicklung der Beschleunigertechnologie erforderlich.

Plasmabasierte Beschleuniger sind aufgrund ihrer extrem hohen Beschleunigungsfelder ein vielversprechender Ansatz, um die Baukosten zu reduzieren. Obwohl in den letzten Jahren erhebliche Fortschritte bei der plasmabasierten Elektronenbeschleunigung erzielt wurden, stellt die Positronenbeschleunigung immer noch eine ungelöste Herausforderung dar und wurde als entscheidendes Problem bei der Realisierung eines plasmabasierten Elektron-Positron-Kolliders identifiziert. Darüber hinaus ist plasmabasierte Positronenbeschleunigung, unter Berücksichtigung der erforderlichen hohen Strahlqualität für einen Kollider, selbst konzeptionell herausfordernd und schwer mit den derzeit verfügbaren Simulationswerkzeugen zu modellieren.

Die Positronenbeschleunigung in einem Plasmazyylinder ist ein neues, vielversprechendes Konzept aufgrund der hohen Beschleunigungsfelder und der Erhaltung der Emittanz bei der Beschleunigung, welches ein fundamentales Kriterium für einen Kollider darstellt. Im Schema durchläuft ein relativistischer Elektronenstrahl einen Plasmazyylinder auf der Achse und treibt eine Plasmawelle im Blowout-Regime. Durch den endlichen Radius der Plasmas verringert sich die Rückstellkraft, die auf die Plasmateilchen wirkt, und durch die verzögerte Rückkehr formen die Plasmateilchen ein langes Filament mit hoher Elektronendichte auf der Achse. Dieses Elektronenfilament erzeugt Beschleunigungsfelder und Fokussierfelder für Positronen mit hohen Feldgradienten. Trotzdem blieben einige Fragen bezüglich der Stabilität und der Strahlqualität noch ungeklärt.

In dieser Arbeit wird die Positronenbeschleunigung in einem Plasmazyylinder auf die Probe gestellt. Über den Beweis des Konzepts hinaus wird eine Beschleunigung von Positronen mit hoher Strahlqualität und ihre Widerstandsfähigkeit gegen realistische Bedingungen wie Fehlstellungen demonstriert. Realistische Plasmaprofile und Temperatureinflüsse werden ebenfalls betrachtet und als vorteilhaft für das Konzept befunden. Die Ergebnisse zeigen, dass die Positronenbeschleunigung in einem Plasmazyylinder in der Tat ein praktikables Konzept für eine mögliche Anwendung in einem Kollider ist, auch wenn noch weitere Forschung zur Optimierung der Effizienz erforderlich ist. Darüber hinaus wurde der GPU-beschleunigte, quasi-statische Particle-in-Cell-Code HiPACE++ entwickelt. Der neue Code ermöglicht

die Modellierung plasmabasierter Positronenbeschleunigungsszenarien in voller 3D-Darstellung mit bisher unerreichter Präzision bei gleichzeitiger drastischer Reduzierung der Rechenkosten, was die vorgestellten Studien erst ermöglichte. Es wird ein Ausblick auf die nächsten möglichen Schritte hin zu einem plasmabasierten Kollider gegeben.

LIST OF PUBLICATIONS

This cumulative thesis builds upon four first-authored core publications. The publications and the respective description of my contribution can be found in Appendix A.

1. S. Diederichs, C. Benedetti, A. Huebl, R. Lehe, A. Myers, A. Sinn, J.-L. Vay, W. Zhang, and M. Thévenet,
HiPACE++: a portable, 3D quasi-static Particle-in-Cell code
Computer Physics Communications **278**, 108421 (2022) [44].
2. S. Diederichs, C. Benedetti, E. Esarey, J. Osterhoff, and C. B. Schroeder,
High-quality positron acceleration in beam-driven plasma accelerators
Phys. Rev. Accel. Beams **23**, 121301 (2020) [41].
3. S. Diederichs, C. Benedetti, E. Esarey, M. Thévenet, J. Osterhoff, and C. B. Schroeder,
Stable electron beam propagation in a plasma column
Physics of Plasmas **29**, 043101 (2022) [42].
4. S. Diederichs, C. Benedetti, E. Esarey, M. Thévenet, J. Osterhoff, and C. B. Schroeder,
Self-stabilizing positron acceleration in a plasma column
Phys. Rev. Accel. Beams **25**, 091304 (2022) [45].

The presented work builds upon and extends the following publication, which has already been presented in my master's thesis and thus is not included in this thesis.

5. S. Diederichs, T. Mehrling, C. Benedetti, C. B. Schroeder, E. Esarey, A. Knetsch, and J. Osterhoff
Positron transport and acceleration in beam-driven plasma wakefield accelerators using plasma columns
Phys. Rev. Accel. Beams **22**, 081301 (2019) [46].

The following work was not finalized before the completion of this thesis and is therefore not included.

6. S. Diederichs, C. Benedetti, E. Esarey, M. Thévenet, A. Sinn, J. Osterhoff, and C. B. Schroeder
Temperature effects in plasma-based positron acceleration schemes using electron filaments
Physics of Plasmas **30**, 073104 (2023) [43].

The following publications have been contributed to, but are not further discussed in this thesis.

7. R. D’Arcy, J. Chappell, J. Beinortaite, S. Diederichs, G. Boyle, B. Foster, M. J. Garland, P. Gonzalez Caminal, C. A. Lindstrøm, G. Loisch, S. Schreiber, S. Schröder, R. J. Shalloo, M. Thévenet, S. Wesch, M. Wing, and J. Osterhoff
Recovery time of a plasma-wakefield accelerator
Nature **603**, 58–62 (2022) [49].
8. S. Schröder, C. A. Lindstrøm, S. Bohlen, G. Boyle, R. D’Arcy, S. Diederichs, M. J. Garland, P. Gonzalez Caminal, A. Knetsch, V. Libov, P. Niknejadi, K. Pöder, L. Schaper, B. Schmidt, B. Sheeran, G. Tauscher, S. Wesch, J. Zemella, M. Zeng, and J. Osterhoff
High-resolution sampling of beam-driven plasma wakefields
Nature Communications **11**, 5984 (2020) [135].
9. A. Knetsch, B. Sheeran, L. Boulton, P. Niknejadi, K. Pöder, L. Schaper, M. Zeng, S. Bohlen, G. Boyle, T. Brümmer, J. Chappell, R. D’Arcy, S. Diederichs, B. Foster, M. J. Garland, P. Gonzalez Caminal, B. Hidding, V. Libov, C. A. Lindstrøm, A. Martinez de la Ossa, M. Meisel, T. Parikh, B. Schmidt, S. Schröder, G. Tauscher, S. Wesch, P. Winkler, J. C. Wood, and J. Osterhoff
Stable witness-beam formation in a beam-driven plasma cathode
Phys. Rev. Accel. Beams **24**, 101302 (2021) [78].

CONTENTS

PREFACE	1
I Particle accelerators for high energy physics	
1 PARTICLE ACCELERATION	7
2 PLASMA-BASED PARTICLE ACCELERATION	13
3 OVERVIEW OF PLASMA-BASED POSITRON ACCELERATION SCHEMES	21
II Numerical methods	
4 MODELLING OF PLASMA-BASED ACCELERATORS	29
4.1 The quasi-static particle-in-cell method	29
4.2 GPU-accelerated high-performance computing	32
4.3 HiPACE++: a portable, 3D quasi-static Particle-in-Cell code	34
III Positron acceleration in a plasma column	
5 CONCEPT	39
6 HIGH-QUALITY POSITRON ACCELERATION	49
6.1 Optimal beam loading in plasma-based accelerators	49
6.2 The Slicing Advanced Loading Algorithm for Minimizing Energy spread (SALAME)	52
6.3 Low-energy-spread positron acceleration	55
7 STABILITY ANALYSIS	61
7.1 Transverse beam instabilities	61
7.2 Electron drive beam stability	63
7.3 Positron witness beam stability	67
SUMMARY AND CONCLUSION	73
IV Appendix	
A AUTHOR CONTRIBUTION AND PUBLICATIONS	77
HiPACE++: a portable, 3D quasi-static Particle-in-Cell code	77
High-quality positron acceleration in beam-driven plasma accelerators	87
Stable electron beam propagation in a plasma column	95
Self-stabilizing positron acceleration in a plasma column	105
LIST OF FIGURES	117
LISTINGS	121
ACRONYMS	121
BIBLIOGRAPHY	123

PREFACE

Particle accelerators, and especially particle colliders, have been the driving tools to shape our understanding of matter and the fundamental forces of nature. In a collider, particle bunches are accelerated to high kinetic energies and brought to head-on collisions. Following Einstein's famous relation between energy E and mass m of particles via $m = E/c^2$, where c is the speed of light in vacuum, the energy of the colliding particles can be transformed into new, possibly heavier, particles during the collision.

Increasingly powerful particle accelerators have revealed more and more new particles, which have been embedded into the so-called Standard Model of particle physics (SM) [176]. From this theoretical framework, the existence of further particles was deduced and predicted. Recently, almost 50 years after its prediction [51, 64], the last predicted particle by the SM, the Higgs boson, was confirmed experimentally at the Large Hadron Collider at CERN [1, 33].

Yet, many important questions remain unsolved. The asymmetry of the prevalence of matter over antimatter, the mass of neutrinos, the unification of the three fundamental forces described by the SM with gravitation, and the astrophysical observations that could be explained by so-called dark matter and dark energy are strong indications that the SM is incomplete. Since there are currently no clear predictions for further particles in the SM, the next steps to uncover the laws of nature are not straightforward. Two approaches are pursued: first, a high-precision measurement of the not yet fully understood Higgs boson in a so-called *Higgs factory* and, second, colliding particles at unprecedented high energies to possibly reveal unknown physics phenomena in a so-called *discovery machine*.

The particle physics communities have gathered to discuss the next steps for the upcoming decades in the European Strategy for Particle Physics (ESPP) [61] in Europe, and the Snowmass process [27] in the United States. Both the ESPP and the Snowmass process agree that the precision measurement of the Higgs boson in a dedicated electron-positron collider has the highest priority as a next facility. Electrons and positrons are beneficial, because they are elementary particles and therefore their initial conditions at the collision can be precisely determined. Knowing the initial conditions allows for a high-precision measurement of all the properties of the Higgs boson, namely mass, width, quantum numbers, and coupling to other particles [24, 113]. There are two mature proposals for a Higgs factory, the International Linear Collider (ILC) [12] and the Compact Linear Collider (CLIC) [4, 5]. More proposals at an earlier stage can be found in the Snowmass

report of the collider implementation task force [132]. Furthermore, the report also includes a contemporary review of proposals for discovery machines that aim to reach unprecedented energies beyond 10 TeV. However, the associated costs and environmental impact of these machines due to construction and power consumption are immense and might exceed the limit of affordability [132]. Therefore, the ESPP and the Snowmass process concluded that research and development must be prioritized to drastically reduce the cost and power consumption of these big machines [61, 132].

As a consequence, the ESPP has developed a community-driven Accelerator Research & Development Roadmap [3] for the next decades. In the roadmap, the concept of plasma-based accelerators was identified as one of the five key areas of accelerator R&D owing to their extreme accelerating gradients [13] that could potentially minimize the construction costs.

In a plasma-based accelerator, a particle beam [35] or laser driver [157] excites a wakefield in a plasma that can be used to accelerate particles with gradients orders of magnitudes larger than those achievable in radio-frequency (RF) cavities. Despite not yet being ready for application, plasma wakefield accelerators have made formidable progress in the past decades, though, mostly for the acceleration of electrons. Plasma-based positron acceleration still poses a notorious challenge, and has been identified as a crucial point in the realization of a plasma-based collider [3, 144]. So far, the Facility for Accelerator science and Experimental Test beams (FACET) at Stanford Linear Accelerator Center (SLAC) has been the only plasma-based accelerator to demonstrate positron acceleration. Due to the limited availability of positron beams, experimental progress is scarce. In addition to that, high-quality positron acceleration in a plasma is challenging, even conceptually.

Since positron acceleration is fundamental to the realization of a plasma-based electron-positron collider, the community has responded with various new plasma-based positron acceleration concepts. However, many of them have not been studied sufficiently. Often, a new scheme is proposed and the basic acceleration is demonstrated, but many important aspects are not yet shown that are vital for the realization of a plasma-based collider, like acceleration of high-quality beams, robustness against transverse instabilities, or efficiency. Experimental test facilities for positron acceleration require significant investments, thus, without convincing concepts, they are not funded. Therefore, it is a necessity to bring positron acceleration schemes conceptually to maturity such that experiments can be planned, funded, and conducted.

Due to the lack of experiments, plasma-based positron acceleration is studied predominantly via computer simulations. On top of that, even the simulation of plasma-based positron acceleration for collider

Radio-frequency-based accelerators are typically referred to as conventional accelerators

applications is challenging and computationally expensive. For this reason, the development of adequate modelling tools is an integral part of the progress.

In this work, the promising concept of positron acceleration in a plasma column [46] is put to the test. This thesis is intended as a comprehensive guide of the scheme, laying the foundation for the experimental feasibility. The acceleration of high-quality positron beams and the robustness against transverse instabilities is demonstrated via simulations. Additionally, a new simulation code is presented that allows for modelling challenging plasma-based accelerator scenarios including various positron acceleration schemes at previously unattainable resolution.

Part I

PARTICLE ACCELERATORS FOR HIGH ENERGY
PHYSICS

PARTICLE ACCELERATION

A particle collider is designed to study rare physics processes that can occur when high-energy particles collide head-on. To reach the desired energies, the particles must be accelerated. Due to rarity of the events of interest, as many particles as possible must be brought to collision. Therefore, these machines do not accelerate single particles, but instead so-called particle beams or particle bunches, which consist of up to $\sim 10^{11}$ particles with a combined charge on the nanocoulomb-level. In general, particle colliders operate at the frontier of what is technologically possible and are highly complex machines. Covering their respective challenges in detail exceeds the scope of this introduction, which focuses only on the most important accelerator parameters of a collider, so that the proposed plasma-based positron acceleration scheme can be discussed in that context. A more detailed introduction to colliders and their challenges can be found in e.g., Shiltsev and Zimmermann [147] and a proper introduction into accelerator physics in general can be found in handbooks, e.g., Chao et al. [32] or Reiser [129].

Charged particles are relatively straightforward to accelerate via the electromagnetic force. The equation of motion for a charged particle with the position \mathbf{x} , velocity $\mathbf{v} = d\mathbf{x}/dt$, mass m , and charge q is given by the Lorentz force \mathbf{F} [95]:

$$\mathbf{F} = \frac{d\mathbf{p}}{dt} = q\mathbf{E} + q(\mathbf{v} \times \mathbf{B}), \quad (1.1)$$

where $\mathbf{p} = \gamma m\mathbf{v}$ is the relativistic momentum of the particle, $\gamma = (1 - (\mathbf{v}^2/c^2))^{-1/2}$ the Lorentz factor, c is the speed of light, and \mathbf{E} and \mathbf{B} are the electric and magnetic fields, respectively. Notably, the first term $q\mathbf{E}$ is fully responsible for the acceleration, because the second term $q(\mathbf{v} \times \mathbf{B})$ is always perpendicular to the velocity \mathbf{v} due to the cross product. Thus, the electric field \mathbf{E} can be used for acceleration and deflection, while the magnetic field \mathbf{B} can be used for deflection only. In conventional accelerators, magnetic fields are typically used to guide relativistic charged particle beams, since moderate magnetic fields of 1 T generate an equivalent force as large electric fields of 300 MV/m, which are, unlike 1 T magnetic fields, non-trivial to achieve with current technology. In plasma-based accelerators both the electric and magnetic fields significantly contribute to the focusing of the charged particles [53].

The total energy gained by the particle ΔE_p is given by the force \mathbf{F} integrated over the path s . For convention, the longitudinal coordinate z is chosen to be parallel to the path of the particle. Then, only the

Symbols written in bold denote three-dimensional vectors

longitudinal electric field E_z contributes to the acceleration and ΔE_p is given by

$$\Delta E_p = q \int_0^L E_z(s) ds, \quad (1.2)$$

where L is the total length of the accelerator. Since the charge q for a given particle is fixed, the energy gain can only be increased via two ways: either the accelerator length L or the longitudinal electric field E_z must be increased. Conventional accelerators rely on RF to generate the electric fields and are currently limited to field gradients of ~ 100 MV/m due to material breakdown [28]. Thus, the easiest solution is to increase the path within the accelerating field. This can be achieved by coupling many accelerator modules behind each other – a linear accelerator. Another option is to re-use the same accelerator module many times by guiding the accelerated particles in a circle – a circular particle accelerator.

LINEAR AND CIRCULAR ACCELERATORS

In a linear accelerator, many accelerator modules are coupled behind each other and each of these accelerator stages are used only once. In a circular accelerator, the particles are guided via magnetic fields to follow a closed orbit, and an accelerator module can be used many times. Re-using the accelerator modules allows for reaching high energies with few accelerator modules. The currently most powerful particle accelerator is the Linear Hadron Collider (LHC) at CERN, which accelerates protons up to 7 TeV in a ring with a circumference of 27 km. Still, the maximum energy in circular accelerators is limited for two reasons.

The first reason is the available magnetic field strength. The required magnetic field strength B in a dipole to keep a charged particle with charge q and momentum p on a circular path with a radius of curvature of ρ_c is given by

$$\frac{1}{\rho_c} = \frac{qB}{p}. \quad (1.3)$$

For example, to keep the 7 TeV protons on the circular track in the LHC with a radius of curvature of $\rho_c \simeq 2.8$ km, superconducting dipole magnets with 8.3 T field strength are required, which was at the edge of what was achievable at that time. Nowadays, 11 T have been demonstrated [23] and 16 T are envisioned [161], enabling smaller radii of curvature or higher beam energies.

The second limiting factor of a circular accelerator is synchrotron radiation. When relativistic particles are accelerated, they emit synchrotron radiation. While this effect is negligible for longitudinal acceleration, it can be significant for transverse acceleration, i.e., when

the particles are guided via magnetic fields. The radiated power P_S is proportional to the fourth power of the energy of the particle E_p and inversely proportional to the mass of the particle squared and the bending radius squared:

$$P_S \propto \frac{E_p^4}{m^2 \rho_c^2}. \quad (1.4)$$

Especially for light particles like electrons and positrons, synchrotron radiation can be the limiting factor for reaching high energies. It was found that the associated costs with a circular collider scales with E_p^2 [130], while it scales only with E_p for a linear collider. As a consequence, reaching high energies with light particles is more cost-effective in a linear accelerator, and the turning point for an RF-based accelerator was determined to be a beam energy of roughly 100 GeV [129]. Since this work investigates a plasma-based positron acceleration scheme, the introduction focuses on linear colliders. High beam energies are necessary to generate heavy, exotic particles – however, the chance of generating them is determined by the luminosity, which is discussed in the next section.

A more detailed explanation of synchrotron radiation can be found in e.g., Reiser [129]

BEAM REQUIREMENTS FOR A LINEAR PARTICLE COLLIDER

Besides reaching high energies to enable certain physics processes, a particle collider requires a high event rate, so that the desired processes can be studied with sufficient statistics and signal can be distinguished from noise. The event rate dN/dt in a collider is given by

$$dN/dt = \mathcal{L} \sigma_I, \quad (1.5)$$

where \mathcal{L} is the luminosity and σ_I the cross section of the event. Since the event cross section is a property of the physical process, it is here not considered as a free parameter in the accelerator, although polarization of electrons and positrons can lead to significantly higher cross sections [112]. The luminosity is the pivotal property of the collider that defines the achievable event rate and is desired to be maximized.

The luminosity \mathcal{L} for two colliding Gaussian bunches of equal sizes is defined as [144]

$$\mathcal{L} = H_D \frac{N^2 f}{4\pi \sigma_x^* \sigma_y^*} \quad (1.6)$$

with N being the number of particles per bunch, f the frequency of the bunch collisions, σ_x^* and σ_y^* being the root mean square (RMS) bunch sizes in the horizontal and vertical plane, respectively, and H_D a geometric factor. The geometric factor H_D takes into account the focusing effect of the electron bunch on the positron bunch and vice-versa,

The luminosity is discussed in more detail in Herr and Muratori [63]

which reduce the beam sizes and thus increase the luminosity [144]. While for RF-based accelerators it can be up to 2, it is expected to be $H_D \approx 1$ due to the ultrashort bunches in plasma-based accelerators [140]. The asterisks indicate the parameters at the interaction point. Thus, to achieve a high luminosity, a high bunch charge, a high collision frequency, and a small bunch size at the interaction point are needed. Obtaining a small bunch size is strongly coupled to the beam quality, which is elucidated in the next paragraphs.

The bunch sizes at the interaction point can be written as

$$\sigma_{x,y}^* = \sqrt{\frac{\epsilon_{x,y} \beta_{x,y}^*}{\gamma}} \quad (1.7)$$

with the normalized emittance $\epsilon_{x,y}$, and the beta function $\beta_{x,y}^*$ in the horizontal and vertical plane. The beta function describes the envelope of the beam width and depends on the focusing system of the accelerator. At the interaction point, the final focusing system squeezes the beam as tightly as possible to maximize the luminosity. In the final focusing system, the minimum beta function at the interaction point is limited by the focusing strength of the magnets and the Oide effect [121]. Thus, the emittance is the crucial parameter that determines the beam size and must be minimized to achieve a high luminosity.

There are different definitions of the emittance in the literature. In this work, following the convention by Reiser [129], the emittance is defined as a measure of transverse beam quality. Furthermore, the normalized RMS emittance is used, because it is preserved if the acceleration is adiabatic and allows for a better description of the beam quality throughout the accelerator. Then, the emittance in x (and equivalently in y) is given by

$$\epsilon_x = \sqrt{\langle x^2 \rangle \langle u_x^2 \rangle - \langle x u_x \rangle^2} \quad (1.8)$$

where $u_x = p_x/(mc)$ is the normalized momentum in x and $\langle \cdot \rangle$ denotes the average over the beam particle distribution. Neglecting the effect by an energy spread, the emittance is preserved during the beam transport using linear focusing fields, while non-linear focusing fields can deteriorate the beam quality. In general, the emittance can increase due to various effects, but it requires significant effort to reduce the emittance again, e.g., radiation cooling via synchrotron radiation. A detailed description on sources of emittance growth can be found in Reiser [129], Ch. 6 for conventional accelerators and in Lindström and Thévenet [89] for plasma-based accelerators.

Besides the emittance, another parameter linked to the beam quality influences the spot size at the collision point, and that is the energy spread of the beam. Similarly to the bending in the dipole in Eq. 1.3, the focusing force in a quadrupole also depends on the particle momentum p and thus on the particle energy. Consequently, particles

with higher energies are focused less strongly than particles with lower energies, which leads to a smearing of the focus, an effect called *chromaticity*. Although this can be corrected via sextupoles (see, e.g., Shiltsev and Zimmermann [147] and references therein for further information), it still sets a stringent limit on the energy spread acceptance of the final focusing system. For example, the RMS energy spread design value for CLIC is 0.35 % [4]. Designing a beamline with a larger momentum acceptance requires further research and development.

In summary, to achieve a high luminosity it is of utmost importance to accelerate *high-charge, low-emittance, and low-energy-spread* particle bunches at a high frequency. More advanced effects that impact the luminosity including crossing angle, hour-glass effect, or misalignment are described in the literature, e.g., in Herr and Muratori [63].

ACCELERATION MECHANISM REQUIREMENTS FOR A PARTICLE COLLIDER

Despite general requirements to the beam quality, a particle collider must fulfill a few more constraints that impacts the possible acceleration schemes.

Although the effect of misalignment of the colliding beams on the luminosity is not discussed here in detail, stability is an important issue in a collider. The underlying particle acceleration scheme must be intrinsically stable, as otherwise small misalignment, which is inevitable in any accelerator, can build up and ultimately lead to beam loss or severe luminosity degradation. In fact, transverse instabilities can be the limiting factor for the achievable beam current in a collider.

Further considerations for particle colliders are driven by their sheer size and associated costs. Two important cost factors are the construction costs and the electrical power consumption. The estimated total cost for a Higgs factory is in the range of multiple billions to a few tens of billions of dollars, while the cost for a discovery machine above 10 TeV is estimated to be tens of billions of dollars. These are quite restrictive costs considering that the total annual budget of the world's particle physics research is only the order of 3 billion dollars [148]. Therefore, the cost reduction must be one of the primary goals of accelerator R&D [3, 132].

To reduce the construction costs for linear colliders, it is important to reach high acceleration gradients to limit the accelerator length. In fact, to achieve above 10 TeV center-of-mass energy, a paradigm shift in the accelerator technology is necessary, as otherwise the construction costs are prohibitive. Plasma-based accelerators are a promising candidate to overcome this limitation due to their ultra high field gradients of tens to hundreds of GV/m, orders of magnitude above the breakdown limit of RF-cavities.

To reduce the electrical power consumption, there are different approaches. First, one has to make sure that the particles are accelerated as efficiently as possible. A decisive figure of merit is the wall-plug efficiency, which is defined by how efficiently the power from the electric grid can be transferred to the particle bunch. To achieve a high wall-plug efficiency, all subsequent energy transfers within the acceleration process, e.g., the generation of the electric fields and the extraction of the energy by the accelerated bunch, must be highly efficient, as otherwise the wall-plug efficiency plummets and the power consumption gets prohibitively high. Another approach to reduce the power consumption is energy recovery. This can in principle be realized by recycling the energy of the accelerated bunch after the collision, as proposed in several energy recovery colliders [15, 90, 158]. Still, this is ongoing research and currently no mature design ready for implementation in a collider exists. To compare different collider schemes including their power consumption, the luminosity per power \mathcal{L}/P is used as a figure of merit. This is especially important when comparing circular and linear colliders since circular colliders emit synchrotron radiation, but the bunches collide many times, while in linear colliders there are no radiation losses, but the beams collide only once.

In conclusion, a particle collider at the energy frontier is the most challenging machine from the viewpoint of accelerator physics and still requires significant research and development. Strict requirements in beam quality, beam stability, efficiency, and accelerating gradient are needed to achieve relevant luminosities of interesting particle physics events at a reasonable cost. Any collider-relevant particle accelerator scheme must fulfill these requirements and its performance can be measured by these key parameters. As stated before, significant accelerator R&D is needed to reduce the costs of any future high-energy collider. This work focuses on a plasma-based positron acceleration scheme, therefore, the basic principles of plasma-based accelerators are introduced in the next chapter.

In a plasma-based accelerator, a particle or laser beam excites a plasma wake via the space charge field of the particle beam or the ponderomotive force of the laser and thereby transfers energy to the plasma. Subsequently, a so-called witness bunch can extract energy from the plasma wake, leading to its acceleration and increased energy. This chapter gives a brief introduction into the basic concepts of plasma-based accelerators required to understand the main studies of this work. A thorough review is given in Esarey, Schroeder, and Lee-mans [53] for laser-driven and in Hogan [67] for beam-driven plasma wakefield accelerators, respectively.

Fundamental properties of plasma

Plasma is the fourth fundamental state of matter after solid, liquid and gas. By subsequently increasing the energy imparted, a solid turns into a liquid and a liquid turns into a gas. Then, when the energy imparted is on the order of the binding energy of the electrons of the gas (e.g., 13.6 eV for hydrogen), the gas can be ionized into a plasma. Strictly speaking, not every ionized gas is considered a plasma, since any gas always has some finite degree of ionization. Therefore, a more rigorous definition is needed. This work follows the textbook definition from Chen [34]: "*A plasma is a quasi-neutral gas of charged and neutral particles, which exhibits collective behavior*". The additional requirements of *quasi-neutrality* and *collective behavior* are explained in the next paragraphs.

In a plasma, local charge perturbations are shielded by freely moving charges on the spatial scale of the characteristic shielding length, the so-called Debye length λ_D . From distances larger than the Debye length, the plasma appears quasi-neutral. Due to the mass asymmetry between the light electrons and the heavy ions, the shielding happens mostly due to the mobile electrons, while the ions form a static background. Under the assumption of an immobile ion background, the Debye length is given by

$$\lambda_D = \sqrt{\frac{\epsilon_0 k_b T_e}{n_0 e^2}}, \quad (2.1)$$

with ϵ_0 , k_b , and e being the vacuum permittivity, the Boltzmann constant, and the elementary charge, respectively. T_e denotes the electron temperature and n_0 the unperturbed plasma electron density. For the plasma shielding to work, two more conditions need to be satisfied.

First, there must be a sufficient amount of particles for the shielding to be a statistically valid concept. The plasma parameter N_D describes the number of particles within a sphere with the radius of a Debye length:

$$N_D = \frac{4}{3}\pi n_e \lambda_D^3 \approx 1.72 \times 10^{12} \sqrt{\frac{(k_b T_e)^3}{n_0}} \quad (k_b T_e \text{ in eV}), \quad (2.2)$$

and it must be a large number, which already hints at the required collective behavior. Second, the spatial dimension of the plasma L must be much larger than the Debye length, otherwise no quasi-neutrality can be obtained within the plasma.

Small perturbations of the plasma electron density lead to oscillations of the electrons around the charge equilibrium due to the restoring force of the ion background. From the laws of electrostatics, the angular oscillation frequency of the plasma electrons ω_p can be calculated [162] and is given by

$$\omega_p = \sqrt{\frac{n_0 e^2}{m_e \epsilon_0}} \quad (2.3)$$

where m_e denotes the mass of the electrons. The plasma frequency ω_p gives the characteristic time scale on which a plasma collectively reacts to perturbations.

For the plasma to behave differently than a neutral or weakly ionized gas, the so-called collective behavior is required. Collective behavior means that the dynamics in the plasma are dominated by the electromagnetic forces which many charged particles take part in, and not by single particle interactions like collisions. Therefore, it is required that the mean time between collisions τ is long compared to the characteristic time scale of the plasma defined by the plasma frequency, namely, that $\omega_p \tau \gg 1$

In summary, the definition of plasma used in this work is given by the following three conditions:

1. $\lambda_D \ll L$, the Debye length must be significantly smaller than the spatial dimension of the plasma to allow for Debye shielding and thus, quasi-neutrality.
2. $N_D \gg 1$, the plasma parameter must be much larger than one for the Debye shielding to work, as many particles are needed for it to be a statistically valid concept.
3. $\omega_p \tau \gg 1$, the time scale of single particle collisions must be larger than that of the collective, electromagnetic plasma response.

PLASMA ACCELERATION

The potential of plasmas as an accelerator medium becomes apparent when the electric field amplitude of a plasma wake is examined. Assuming a plasma oscillation in one dimension, the electric field can be obtained from Gauss's law $\nabla \cdot \mathbf{E} = \rho/\epsilon_0$, with ρ being the charge density. If electrons and ions are fully separated on the scale of a plasma skin depth $k_p^{-1} = c/\omega_p$, then $\rho = n_0e$ and one obtains an estimate for the maximum electric field amplitude by

$$E_{\max} = \frac{1}{k_p} \frac{n_0e}{\epsilon_0} \quad (2.4)$$

which is referred to as the cold, non-relativistic wave-breaking limit amplitude E_0 [53] and commonly written in the form of, by substituting $k_p^{-1} = c/\omega_p$,

$$E_0 = \frac{\omega_p m_e c}{e}. \quad (2.5)$$

The cold, non-relativistic wave-breaking limit amplitude can be approximated by $E_0(\text{V/m}) \approx 100\sqrt{n_0(\text{cm}^{-3})}$. Thus, plasma densities of 10^{16} - 10^{18} cm^{-3} can sustain ultra-high electric fields of 10-100 GV/m, orders of magnitude above the breakdown threshold of RF cavities.

To utilize these fields in plasmas for a high-gradient accelerator, two elements are needed: a drive beam, which induces a charge separation, and a witness beam, which is accelerated in the arising fields. The driver can either be a high-power laser beam [157] that relies on the ponderomotive force to achieve charge separation [79] or a relativistic charged particle bunch [35, 134] that separates the charge due to its space-charge fields. The driver traverses the plasma with close-to speed of light, pushing (laser or electron drivers) or pulling (positron or proton drivers) the light electrons, thereby inducing a co-propagating plasma wake that sustains the aforementioned ultra-high electromagnetic fields, called *wakefields*. From the Lorentz force, the fields \mathbf{W} acting on an ultra-relativistic particle beam ($v_z \approx c$) reduce to

$$\mathbf{W} = \begin{pmatrix} E_x - cB_y \\ E_y + cB_x \\ E_z \end{pmatrix}. \quad (2.6)$$

Therefore, E_z is referred to as the longitudinal wakefield, and $E_x - cB_y$ and $E_y + cB_x$ are referred to as the transverse wakefields, respectively. In cylindrically symmetric geometries, the transverse and longitudinal wakefields are connected by the Panofsky-Wenzel theorem [124] via

$$\nabla_{\perp} W_z = \partial_{\zeta} W_{\perp}. \quad (2.7)$$

with the co-moving, longitudinal variable $\zeta = z - ct$, the time t , and $W_{\perp} = (W_x, W_y)$.

To drive a plasma wake effectively, the driver must excite the wake resonantly, i.e., the length of driver must be on the same scale as the plasma wavelength $\lambda_p = 2\pi/k_p$. For a plasma density of $n_0 = 10^{18}\text{cm}^{-3}$, this corresponds to a beam length of $\sim 30\ \mu\text{m}$ or, as typically provided for lasers, a pulse duration of $\sim 100\ \text{fs}$. Longer beams need to undergo self-modulation transforming into multiple short bunches on the length scale of the plasma wavelength before they can resonantly drive a wake in the plasma, see Sec III.D in Esarey, Schroeder, and Leemans [53] for self-modulated laser bunches or Caldwell et al. [29] for long, self-modulated proton bunches.

Linear wakefields

In the linear regime, the driver-induced plasma density perturbation is much smaller than the ambient plasma density. This is the case for particle beams with a peak density n_b much smaller than the plasma density n_0 , i.e., $n_b/n_0 \ll 1$, and for laser drive beams with a laser strength parameter a_0 much smaller than one, $a_0 \ll 1$. In the linear regime, the plasma wake can be calculated fully analytically [59, 134] from the cold fluid equations, namely the continuity equation and the fluid momentum equation, and the Poisson equation. Then, the trailing wakefields behind the driver are of sinusoidal shape, as shown in Fig 2.1, where the perturbation of the electron plasma density $\delta n_e/n_0$ (top plot), the accelerating longitudinal wakefield E_z (center plot) and the focusing transverse wakefield $E_x - cB_y$ (bottom plot) are shown, respectively. Although the focusing and accelerating phase is shifted, they still overlap for both electrons and positrons for quarter the period. Thus, simultaneous acceleration and focusing of both is possible. However, the accelerating gradient and the possible accelerated charge is limited in the linear regime. Increasing the driver strength to $n_b/n_0 \lesssim 1$ for beam drivers or $a_0 \lesssim 1$ for laser drivers increases the accelerating gradient and allows for higher witness beam charge, while the field structure is still close-to sinusoidal, which is called the quasi-linear regime. If the driver strength is further increased, the wake becomes fully non-linear and a kinetic description of the plasma is required.

Blowout wakefields

In the non-linear regime, one case is of special interest. In the so-called blowout or bubble regime [126, 131], the space charge force of a narrow drive bunch ($n_b/n_0 \gg 1$) or the ponderomotive force of a tightly focused laser pulse ($a_0 \gg 1$) is strong enough to fully expel the plasma electrons, forming a cavity filled with only the ions. The




Figure 2.1: Plasma wake and corresponding wakefields in the linear regime. The electron density perturbation normalized to the background plasma density, the accelerating field and the focusing field normalized to the cold, non-relativistic wave-breaking field E_0 are shown in the x - ζ -plane in (a)–(c), respectively.

ions attract the electrons and so they return in some distance behind the driver to the propagation axis of the driver. Notably, almost all electrons return in the same longitudinal position, forming a large density spike.

The plasma wake and the corresponding wakefields in the blowout regime are shown for an exemplary electron drive beam in Fig. 2.2. The field structure in the ion cavity consists of a highly non-linear accelerating wakefield and a linear transverse focusing wakefield for electrons (if the motion of the background ions can be neglected and the transverse ion profile is uniform), making the blowout regime attractive for high-gradient, emittance-preserving electron acceleration. Furthermore, the on-axis transverse focusing field has no longitudinal dependence within the blowout cavity. Therefore, as a consequence from the Panofsky-Wenzel theorem (see Eq. 2.7), the accelerating field is uniform transversely, which allows for uniform acceleration at a fixed longitudinal position in the wake. Note that the accelerating field has a longitudinal dependence and therefore uniform acceleration of a complete bunch requires *beam loading*, an effect described in more detail in Ch. 6.

The blowout regime has demonstrated high-gradient electron acceleration experimentally. Furthermore, significant progress has been achieved in the terms of beam quality of the accelerated bunches in the last years. Since early milestones like the energy doubling achieved for fractions of a 42 GeV beam in a 85 cm long plasma target [22], the wakefield has been precisely measured [135], and low-energy-spread acceleration for both beam-driven [88] and laser-driven [77] plasma accelerators in the blowout regime have been demonstrated.

Despite its successes in electron acceleration, the blowout is not usable for positron acceleration – the ion cavity is defocusing for positrons and only the narrow electron cusp at the back of the blowout is focusing for positrons, see Fig. 2.2 (c). Using a positron instead of an electron drive beam does not transfer to a blowout regime usable for positron acceleration due to the mass asymmetry between the light, mobile plasma electrons and the heavy, rigid plasma ions. Instead, the mobile electrons are pulled towards the axis, forming a electron density spike on axis [37, 82]. The resulting focusing fields are non-linear and significant emittance growth is expected for a Gaussian positron bunch. Thus, unlike the blowout regime for electrons, there is currently no clear solution for plasma-based positron acceleration that fulfills all the requirements of collider, namely efficient, low-emittance, low-energy-spread, high-gradient, stable positron acceleration. There are many approaches though, and the most popular are introduced in the next chapter.

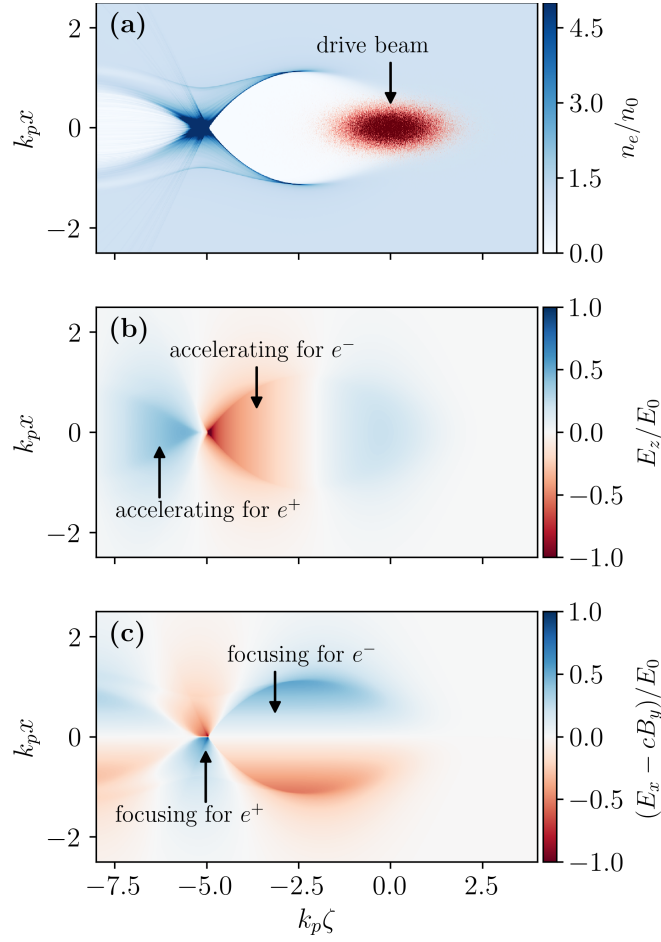


Figure 2.2: Plasma wake and corresponding wakefields in the blowout regime. The electron density normalized to the background plasma density, the accelerating field and the focusing field normalized to the cold, non-relativistic wave-breaking field E_0 are shown in the x - ζ -plane in (a)–(c), respectively.

Efficiency in plasma accelerators

As discussed in Ch. 1, the wall-plug efficiency is of paramount importance for a linear collider as otherwise the power consumption costs are prohibitive. The wall-plug efficiency is defined by the energy transfer ratio from the electrical grid to the accelerated beams. In a plasma accelerator, the wall-plug efficiency is the product of the wall-plug-to-driver efficiency, the driver-to-plasma efficiency, and the plasma-to-witness efficiency. The wall-plug-to-driver efficiency is estimated to be up to 55 % for particle beam drivers in the CLIC design [4]. Although it is typically below 1 % for widespread titanium-sapphire laser systems, fiber lasers may achieve up to 60 % [55, 125]. Here, the wall-plug-to-driver efficiency is excluded from the discussion of various plasma-based positron acceleration schemes, as it may impact all beam-driven schemes similarly or significantly altered in the case of upcoming laser drivers. Furthermore, plasma-based positron schemes rarely provide the total driver-to-witness efficiency, since that requires modelling the full accelerator, but instead only the *instantaneous* energy-transfer efficiency from the drive beam to the witness beam. Note that the total driver-to-witness efficiency is always smaller than the instantaneous energy transfer efficiency, since it requires full driver depletion for both to be equal, which is difficult to achieve. Still, the instantaneous transfer efficiency is a useful estimate. Therefore, throughout the rest of the thesis, the term efficiency refers to the instantaneous energy-transfer efficiency and is used as a figure of merit for the applicability in a linear collider.

OVERVIEW OF PLASMA-BASED POSITRON ACCELERATION SCHEMES

Plasma-based positron acceleration is challenging, due to the asymmetric plasma response and the strict beam requirements of the, currently, only known application – a high-energy electron-positron collider. As highlighted in Chapter 1, any linear-collider-relevant acceleration scheme must fulfill the key requirements of high efficiency, high gradient, low energy spread, low emittance, and stability. Thus, simultaneously accelerating and quality-preserving focusing fields for positrons in a plasma are needed. To obtain a positron-accelerating and -focusing wakefield structure, an area with a large excess of plasma electrons must be generated. This is not trivial due to the high mobility of the plasma electrons, i.e., a positron witness bunch with a reasonable charge affects the plasma electron distribution itself significantly. Still, there are many different approaches to generate positron-accelerating and -focusing wakefields. At the current status of the research, it is not clear which scheme is the most promising for a future linear collider. The most important schemes and their respective challenges are introduced in this chapter.

The effect of the witness bunch on the wake is discussed in detail in Ch. 6

Linear wakefields

The linear regime of plasma accelerators (see Ch. 2) requires a drive beam with a peak density of $n_{b,d}/n_0 \ll 1$ or a laser driver with normalized peak laser amplitude $a_0 \ll 1$. Due to the symmetric plasma response, it is suitable for both electron and positron acceleration. The linear regime for positron acceleration was first demonstrated at FFTB at SLAC [21]. Recently, the regime was discussed in depth in Hue et al. [70]. Although the linear regime is suitable for positron acceleration and focusing in general, there is a trade-off between high beam quality and charge that can be accelerated.

The problem lies in the requirement for the linear regime, namely $n_{b,w}/n_0 \ll 1$, where $n_{b,w}$ denotes the density of the positron witness bunch and n_0 the background plasma density, respectively. The beam size reduces during the acceleration, thereby increasing the peak density, leading to the witness beam driving a non-linear wake itself. Even worse, the background electrons may be sufficiently perturbed so that the positron beam is defocused. To prevent the beam collapse, a large witness beam emittance or weak focusing is required. For FACET-like positron beam parameters, the required emittance in the linear regime was reported to be 4000 mm mrad [9], which is unacceptably

high for collider applications. Using a smaller emittance requires to decrease the charge in the positron bunch. Alternatively, weak focusing can be achieved by transversely tailoring the driver, reducing the transverse wakefield and thus, the required emittance. Still, the decreasing beam size due to the acceleration eventually leads to a non-linear wake driven by the witness bunch. Notably, a positron witness bunch that causes a nonlinear plasma response only suffers limited emittance growth, if it is quasi-matched to the non-linear focusing fields [70] and no defocusing occurs. However, this does not resolve the issue that in general only a small charge can be accelerated in linear wakes while maintaining reasonable beam quality.

The quasi-linear regime to mildly non-linear regime requires a drive beam with $n_{b,d}/n_0 \lesssim 1 - 2$ and enables higher accelerating gradients while maintaining a similar focusing structure. Although higher accelerating gradients can be achieved, the regime is generally plagued by the same issues as the linear regime.

Non-linear, positron-beam-driven regime

For $n_b/n_0 \gg 1$ the plasma electrons overshoot the propagation axis and form a wake similar to the electron blowout

A tightly focused, strong electron drive beam drives a wake in the blowout regime. In contrast, a positron drive beam with a density of $n_b/n_0 > 1$ attracts the mobile electrons, forming an electron density filament on axis, which enables positron acceleration and focusing. The regime was first described by Lee et al. [82] and has been named *flow-in* regime. The focusing fields in this regime are highly non-linear in the transverse direction and vary along the bunch longitudinally. The first experiments in 2003 used a single bunch, where the head of the bunch acts as a driver and the tail of the bunch as a witness beam. In that setup, the longitudinally varying focusing force was first observed experimentally in Hogan et al. [66] and a significant emittance growth was predicted, which was later experimentally confirmed by Muggli et al. [117]. The acceleration of the bunch tail by 5 GeV in a 1.3 m long plasma was demonstrated in Corde et al. [37]. Later, similar results were obtained in a two beam setup, with a distinct drive and witness positron beam [47].

Due to the strongly non-linear focusing fields and the resulting emittance growth, the positron-beam-driven, non-linear regime does not seem suitable for a collider.

Elongating the electron spike at the back of a blowout

Another way to generate a large excess of plasma electrons is utilizing the electron cusp at the back of the blowout regime. In the blowout regime (see Ch. 2), all plasma electrons are expelled and return after some distance behind the driver to the driver's propagation axis. This leads to a small area with a very high plasma electron density that

generates focusing and accelerating fields for positrons. The region is longitudinally and transversely narrow and does not provide sufficient space for a positron bunch with reasonable charge and length to be accelerated. By *elongating* this electron density spike at the back of the blowout, a longer, high-density electron filament can be created. This filament provides sufficient space for the transport and acceleration of positron beams with reasonable charge and length.

Utilizing the electron cusp at the back of the blowout was first studied via simulations by Lotov [98] for beam-driven plasma accelerators. In that work, a positron witness bunch is placed in the electron cusp at the back of the blowout, attracting more plasma electrons and thereby leading to an elongated electron sheath. Recently, this approach was further investigated and extended by a theoretical model to enable high quality positron acceleration [187] in this setup. Thereby, the acceleration of a positron beam with \sim GV/m accelerating gradient, several μ m emittance, few percent RMS energy spread, and tens of percent transfer efficiency was demonstrated via simulations. Despite the progress, the main limitation of this approach [98] has not been addressed: Due to the steep, non-linear accelerating gradient along the longitudinal axis, any small positional jitter of the witness beam drastically affects the outcome in terms of energy gain and energy spread, rendering a stable experimental implementation of this approach challenging.

The approach of utilizing the electron cusp at the back of the wake was also investigated for laser-driven plasma accelerators via simulations [92]. There, a laser drives a wake close to the blowout regime and is guided in a parabolic plasma channel. By carefully choosing the laser parameters and by loading the wake with a positron witness bunch, the electron cusp at the back of the blowout is elongated. Additionally, the positron witness bunch attracts more plasma electrons and thereby strengthens the electron filament. Positron acceleration on the order of hundreds of MeV-level was demonstrated. However, so far the acceleration was only demonstrated with a large energy spread. The approach requires significant optimization to become feasible for a collider application.

Another approach proposes to use an electron witness beam in the back of the blowout placed shortly before the electron cusp to decelerate the plasma electrons in the sheath [173]. Additionally, a warm plasma was assumed. The combination of a plasma temperature and the deceleration of the sheath electrons via an electron witness bunch leads to an elongation of the electron cusp, generating an extended electron filament, suitable for positron transport and acceleration. Despite achieving a high accelerating gradient of tens of GV/m, the demonstrated energy spread in the simulation was on the order of ten percent for a high-charge positron bunch. Thereby it should be noted

that the results were not optimized and the energy spread could be drastically reduced by shaping the positron beam current profile.

Last, using an electron drive beam and a finite radius plasma column modifies the trajectories of the plasma sheath electrons, leading to an elongated electron filament [46]. Again, the elongated electron filament allows for simultaneous focusing and accelerating of a positron bunch. The concept of positron acceleration in a plasma column is studied in depth in this thesis and is explained in detail in Ch. 5. The concept was extended by colleagues [128] by using a two column structure, where a laser (instead of simply generating the column in front of the drive beam) is placed behind the drive beam. Then, the drive beam itself ionizes a column due to its space charge fields and the laser increases the column radius by further ionizing the surrounding gas. In the studied regime, a ring-shaped field structure focusing and accelerating for positrons is generated, but the emittance is growing significantly in this setup.

Hollow drivers

Another way to generate a high-density electron filament suitable for positron acceleration is by means of a donut-shaped driver, i.e., a hollow core electron drive beam [70, 72] or Laguerre-Gaussian laser driver [170, 171]. For these hollow core drivers, the plasma electrons within the hollow core radius are pushed towards the propagation axis, leading to a high-density, on-axis plasma electron filament. Although the setup enables simultaneous acceleration and focusing of positron beams, it is challenging to preserve the shape of the driver (both particle beam and laser) over a distance much longer than the characteristic diffraction length. The hollow core electron drive beam is susceptible to the Weibel instability [175]. It arises from noise in the current distribution of the beam, leading to small fields that enhance the noise in the currents, leading to beam filamentation. A transverse (in a ring-shaped beam, azimuthal) momentum spread within the beam can mitigate the instability [156]. However, the required momentum spread to damp the instability in a ring-shaped beam would likely be large and make this approach impractical.

Hollow and quasi-hollow channels

In the previously presented concepts, the accelerating and focusing fields are generated by the means of an excess electron density. If one neglects the need for a focusing field, hollow core plasma channels are a promising solution [36, 82, 142]. In a hollow plasma channel, electron beams [188, 189], positron beams [58, 82], or proton beams [181, 182] can be used to excite a longitudinal wakefield that is accelerating for positrons. For a symmetric bunch that travels on the axis of the hollow

core channel, no focusing field is present. Due to the absence of any focusing field, in principle ultra-low emittance positron beams can be accelerated. Therefore, the hollow core plasma accelerator concept has spurred significant theoretical interest [76, 137, 139] and has been proposed as a basis for a plasma-based linear collider [143].

Not too long ago, the hollow core plasma accelerator has been demonstrated experimentally [58]. Despite its promising properties, the concept is prone to the beam breakup instability [142]. Any offset with respect to the hollow channel axis induces a deflecting force, enhances the offset and leads to beam loss, which was measured and confirmed experimentally [87]. Thus, the stability of both drive and witness beam is of concern in a hollow core plasma accelerator. Although external focusing with quadrupole magnets and a linear energy chirp have been proposed as a possible solution [137, 139, 181], a detailed study on the stability including external focusing is yet missing.

Recently, using an asymmetric electron drive beam in a hollow core plasma channel was proposed [188]. The asymmetric beam induces a quadrupole moment in the focusing field, which leads to focusing in one plane and defocusing in the other. Due to the defocusing force, the drive beam is split and hits the hollow core channel wall in a controlled fashion, providing improved stability against the initial offsets. Another interesting finding of that study is that, in case of a sufficiently high-charged electron drive beam, the electrons of the hollow core plasma cylinder flow into the hollow channel and provide an on-axis electron filament, similar to what can be obtained from elongating the electron spike at the back of the bubble. In that configuration, both accelerating and focusing fields are provided for the positron witness bunch. An energy extraction efficiency from wake to witness bunch of 30 % was demonstrated via simulations, while maintaining a 1 % energy spread. The positron beam emittance used in this example was > 50 mm mrad, too large for a collider. Thus, this scheme needs to be demonstrated at lower emittances to be considered for a collider design.

Another approach was proposed by Silva et al. [150], where a warm, quasi-hollow channel and an electron drive beam are used to create a wake with both positron accelerating and focusing fields. Notably, these warm, quasi-hollow channels arise as an aftereffect of the blowout regime. The longitudinally varying focusing fields provide high robustness against transverse instabilities. The simulation-based study demonstrated the stable acceleration of positron beams of 100 pC charge, several mm mrad emittance, and a final energy spread of several percent. Again, it should be noted that the results were not optimized and tailoring the positron beam current profile, as explained in Ch. 6, could drastically improve the performance.

This is the same effect that provides stability in the plasma column, which is explained in greater detail in Ch. 7.

Other concepts

Since the development of a practical positron acceleration scheme is of high priority within the plasma accelerator community, significant R&D is ongoing and the brief description above cannot include all proposed positron acceleration concepts. A few more proposals include the utilization of a long, self-modulated proton bunch [100], coaxial plasma channels [127], or direct laser acceleration [102]. These concepts are not further discussed in this work since they are not mature enough to be considered for a collider application. The concept of using a self-modulated proton bunch has a low efficiency and the emittance of the positron bunch has not yet been analyzed [100], coaxial plasma channels require ring-shaped positron bunches [127], where the emittance preservation is unclear, and direct laser acceleration has only been demonstrated for positron beams with a charge of 17 fC [102], too little for achieving reasonable luminosities.

Conclusion

The experimental progress on plasma-based positron acceleration has been scarce in the last decade owing to the limited availability of beam facilities generating positron beams. Yet, using simulations, many promising approaches have been developed, although reaching collider-relevant positron beam quality is challenging even conceptually. Many schemes have yet to be thoroughly examined and lack optimization for high-quality beams, or have not addressed critical issues such as stability. At this stage, a fair assessment of the schemes that would identify the ideal one for a plasma-based collider is not possible, since further optimization can drastically affect the outcome in terms of beam quality or efficiency. It is beyond the scope of this work to optimize each of the presented schemes. This work focuses on one of the schemes, positron acceleration in a plasma column, optimizes it in terms of beam quality, and demonstrates the intrinsic transverse stability. This is an important step to assess the feasibility of positron acceleration in a plasma column for a possible collider.

Part II

NUMERICAL METHODS

The access to experimental positron wakefield acceleration facilities is limited. So far, plasma-based positron acceleration were performed only at FACET [37, 47, 58, 87] at SLAC [179], although new concepts [155, 186] could make GeV-level, low-emittance positron beams more accessible. Until new experiments emerge, numerical modeling is simulations are the fundamental approach for assessing and evaluating plasma-based positron acceleration schemes. As stressed before, linear colliders strictly require low-emittance beams, high beam energies, and long propagation distances. The non-linear positron acceleration schemes often rely on fine regions with high electron density. These requirements impose challenges on the simulations, as these fine electron-excess areas and the sub-micrometer scale beams must be resolved. In fact, the physical studies presented in this work were not attainable with previously existing plasma-accelerator simulation codes. Therefore, the novel, quasi-static, 3D particle-in-cell code HiPACE++ was developed and is introduced in this chapter. It combines the latest high-performance computing software practices with an optimized quasi-static particle-in-cell (PIC) algorithm and reaches orders of magnitude speed-up on GPU-accelerated supercomputers over its CPU-based predecessor. The quasi-static PIC method and modern High-Performance Computing (HPC) practices forming the basis for HiPACE++ are outlined in the following sections.

4.1 THE QUASI-STATIC PARTICLE-IN-CELL METHOD

The PIC method [20, 65] is a fundamental tool for modeling and understanding kinetic plasma physics, including plasma-based accelerators. The evolution of a plasma in the collisionless regime can be described by a Maxwell-Vlasov system. The PIC method solves an approximation of this system numerically. Thereby, the plasma is modelled kinetically by so-called macro-particles, each representing an ensemble of physical particles, while the fields are represented on a spatial grid. To solve the discretized Maxwell equations on the spatial grid, the currents and densities must be known. Likewise, advancing particles with the Lorentz force requires the electromagnetic fields at the particles position. In PIC, this is solved by interpolating particle quantities from and to grid quantities.

The core algorithm of the PIC method consists of

A detailed, contemporary review of the PIC method is given by Vay and Lehe [168] or Arber et al. [11].

1. The particle push: the positions and momenta of the macro-particles are advanced by a time step Δt using the Newton-Lorentz equations of motion.
2. The charge/current deposition: the charge and current densities of the macro-particles are interpolated on the spatial grid.
3. The field solver: The electromagnetic fields are calculated on the spatial grid based on the charge and current densities.
4. The field gather: The electromagnetic fields are interpolated to the macro-particles for updating the force terms for the next particle push.

The PIC algorithm includes collective plasma effects, however, single particle effects (e.g., incoherent radiation emission) are absent from the Maxwell-Vlasov formulation. Still, additional physics such as ionization, collisions, or quantum effects can be added to the routine (see e.g., [11]).

The PIC routine is generally considered robust, although the time step Δt is typically restricted by a Courant-Friedrichs-Lewy condition [38] for finite-difference [180] or pseudo-spectral time-domain field solvers [91]. A significant speedup can be achieved by performing the simulation in a Lorentz-boosted frame [165] or by utilizing the quasi-static approximation [153, 154, 160]. The speedup in quasi-static PIC scales with increasing beam energies, therefore it is especially useful for simulating high-energy beams as needed in a collider. Thus, it is often the method of choice for simulating plasma-based positron acceleration and also used throughout the simulation studies in this work.

Quasi-static PIC

The time and length scales in a plasma-based accelerator span many orders of magnitude: On one hand, the plasma wavelength λ_p is on the order of tens of μm and the laser wavelength (in case of a laser driver) is on the order of $\sim \mu\text{m}$. On the other hand, the length of the plasma target is on the order of centimeters to meters, and multiple plasma target stages in a collider are on the order of kilometers. Thus, simulating a collider with a time step resolving the plasma wavelength takes on the order of $\mathcal{O}(10^9)$ time steps, which is beyond feasibility. A relativistic particle beam evolves on the order of its betatron wavelength $\lambda_\beta \simeq \sqrt{2\gamma}\lambda_p$ [52]. Similarly, a laser envelope evolves transversely on the order of its Rayleigh-length $z_R \simeq \pi w_0^2/\lambda$ [115], with w_0 being the beam waist and λ the laser wavelength, respectively. Longitudinally, a laser envelope evolves on the order of its pulse dispersion time [153, 154]. In addition, the spectral phase content of a laser-pulse in a wake-field can evolve more quickly and influence the acceleration process.

Still, the evolution of the envelope for a typical laser used for plasma acceleration is on a much longer distance than the plasma wavelength. The disparity of time scales between the evolution of the plasma and the laser envelope or particle beam gives rise to a separate treatment. This was formulated as the quasi-static approximation [153, 154, 160], where the laser or particle beams are assumed to be frozen during the propagation through a plasma slice. As a consequence, only the betatron motion of the beams or the evolution of the laser envelope must be resolved, leading to a drastically increased time step, especially for collider-relevant energies. For example, for an electron beam with an energy of 10 GeV, the corresponding Lorentz factor is $\gamma \approx 20000$, thus the betatron wavelength is approximately $\sqrt{2\gamma} = 200$ times larger than the plasma wavelength, allowing for a 200 times larger time step in a quasi-static simulation. Thus, the simulation requires 200 times fewer time steps to simulate the same accelerator length.

The separate treatment of the plasma and the beams is apparent in the speed-of-light coordinate, using the co-moving variable $\zeta = z - ct$. While the fast plasma response is calculated in ζ , the beams are assumed to be rigid. Then, when the fields for the given beams are calculated, the beams are advanced in time t with a large time step Δt while the plasma is assumed to be static during the beam advance.

The quasi-static PIC algorithm shown in Fig. 4.1 is outlined as follows:

1. The beam deposits its current and charge densities
2. An unperturbed slice of plasma is initialized ahead of the bunches.
3. In a loop over the longitudinal grid points, the plasma slice is pushed backwards through the beams along ζ . Thereby, at each longitudinal grid point, the plasma particles deposit their currents, the 2D fields on the slice are solved, and the plasma particles are pushed to the next grid point. This corresponds to the core PIC algorithm described above in 2D, whereby the particles are pushed along the longitudinal variable ζ instead of the time t .
4. When the full fields are calculated, the beams are advanced by $\Delta t \gg \Delta\zeta/c$, which is the reason for the large speedup.

Besides the larger time step, this approach has another advantage over full electromagnetic PIC. By solving the 3D problem as n_ζ 2D problems, with n_ζ being the number of longitudinal grid points, the memory requirements are much lower, as only a 2D plasma slice (instead of a full 3D plasma) needs to be stored, and only few 2D slices (instead of a 3D grid) are simultaneously needed for the computation of the 3D fields. This will be important for the porting to Graphics Processing Unit (GPU)s, as explained in section 4.3.




Figure 4.1: Schematic of the quasi-static PIC algorithm. The 3D simulation domain is calculated slice-by-slice in a loop over the longitudinal grid points from the head of the box to its tail. Only the beam particles, a 2D slice of plasma particles, and a few 2D slices of fields are required to determine the wake in the 3D simulation domain. Source: [168]

The advantage of the increased time step in quasi-static PIC comes at the cost of reduced physics. Due to separate treatment of plasma and beam particles, it is not possible to self-consistently model self-injection, where background plasma electrons are trapped and accelerated to relativistic velocities, i.e., the transition from plasma particles to beam particles. This drawback can be circumvented by modeling the injection process with a fully electromagnetic PIC code and then transferring to a quasi-static PIC code for the acceleration process. Also, modeling strong longitudinal plasma density gradients is challenging for quasi-static PIC, but latest research is aiming at overcoming this limitation [163].

The quasi-static approximation has first been utilized in the 2D PIC code WAKE by Mora and Antonsen [114] and since then has been successfully implemented in various codes [19, 68, 86, 96, 106, 172]. Still, even with quasi-static PIC codes, modeling low-emittance beams and resolving fine electron structures in 3D, e.g., as needed for the stability analysis of positron acceleration schemes, is extremely computationally demanding and only feasible if modern supercomputers can be utilized.

4.2 GPU-ACCELERATED HIGH-PERFORMANCE COMPUTING

Supercomputers have significantly evolved in the last years, breaking the exascale barrier with a performance above 10^{18} IEEE 754 Double Precision (64-bit) floating point operations per second (FLOPS) [93]. They are required in all natural sciences to model challenging tasks like earthquakes [56], the infectivity of the SARS-CoV-2 spike pro-

tein [30], or, in the realm of plasma physics, modelling a Tokamak fusion reactor [178]. Modelling a plasma-based collider in full 3D is also tremendously challenging and requires efficient use of modern supercomputers. Recent efforts have managed to model the first three stages of a plasma-based collider using the electromagnetic PIC code WarpX [167] on the supercomputer Summit [169]. However, up to hundreds of stages are needed for a full collider and the later stages are more challenging to model due to the decreasing beam width with increasing beam energy. Thus, much more computing power is needed to model a full collider. Even worse, stability tolerances must be assessed, requiring many of these yet unfeasible simulations.

In the last decade, the supercomputer landscape has shifted to accelerated computing, in particular using GPUs as accelerating devices. Currently, 15 of the 20 supercomputers listed in the TOP500 [110] are GPU-based with 11 using NVIDIA and 4 using AMD GPUs, respectively. As nowadays the carbon footprint has gained increased importance, it is noteworthy that 19 out of top 20 of the Green500, which ranks the TOP500 supercomputers by power-efficiency, are GPU-based. Thus, to be able to use the most powerful and most power-efficient supercomputers of today, simulation software must be able to run on GPUs.

GPU-accelerated computing

Modern HPC GPU architectures provide tens of TFLOPS in double-precision performance, tens of GB fast-access GPU memory, fast atomic operations, and a relatively slow transfer between the host (CPU) and the device (GPU) memories. Different vendors use different GPU architectures and thus GPUs from different vendors require different programming models (CUDA for NVIDIA, ROCm for AMD, OpenCL for Intel, Metal for Apple). As it is impractical to re-write simulation software for each programming model of the corresponding vendor, it has become increasingly important to write portable code, e.g., by using portability layers like Kokkos [50], RAJA [16], or Alpaka [184]. A performance portability layer provides an abstraction for both parallel execution of code and data management, such that a single source code can be executed efficiently on different platforms, including different GPUs. However, a portability layer alone does not guarantee good performance on accelerators such as GPUs, because PIC simulations on accelerators are typically memory-bandwidth-limited. Therefore, to gain optimal performance on GPUs, an algorithm must be optimized to minimize memory transfer between the host and device.

So far, the speedup and power-efficiencies provided by GPUs has only been utilized in electromagnetic PIC codes [25, 26, 84, 119]. Here, the first GPU-capable, quasi-static, 3D PIC code HiPACE++ is presented,

which combines the large time step of quasi-static codes with the possibility to run on the latest compute architectures.

Corresponding publication

S. Diederichs, C. Benedetti, A. Huebl, R. Lehe, A. Myers, A. Sinn, J.-L. Vay, W. Zhang, and M. Thévenet,
HiPACE++: a portable, 3D quasi-static Particle-in-Cell code
Computer Physics Communications **278**, 108421 (2022) [44].

4.3 HIPACE++: A PORTABLE, 3D QUASI-STATIC PARTICLE-IN-CELL CODE

In HiPACE++, the quasi-static PIC algorithm is restructured for optimal GPU performance by reducing the host-device transfers and a novel parallelization scheme, which provides near optimal scaling up to hundreds of GPUs, making the code a prime candidate to utilize GPU-based supercomputers. On modern HPC GPUs, it outperforms CPU implementations by orders of magnitude and allows for modeling previously unattainable wakefield accelerator setups at unprecedented resolution.

In the following, the GPU-optimized quasi-static PIC algorithm is introduced.

GPU-optimized quasi-static PIC

In the quasi-static PIC algorithm outlined in Sec. 4.1, the beam operations are detached from the plasma operations: First, the plasma response to the beam current is calculated slice-by-slice, including the plasma particle deposition and push, and the field solve. Finally, once the fields are computed on the full 3D domain, the field is interpolated to the beam particles and they are pushed. Thus, pushing the beam particles implies that the fields on the 3D domain must be known and therefore allocated in memory, which can be on the order of hundreds of GB of memory, exceeding the memory available in current state-of-the-art GPUs. In the GPU-optimized quasi-static PIC algorithm, the 3D memory allocation is circumvented by integrating the beam operations in the loop over slices, so that it reads as follows.

1. Gather fields and push plasma particles backwards from $\zeta + \Delta\zeta$ to ζ ;
2. Deposit plasma currents and densities;
3. Deposit beam currents and densities;
4. Solve fields on slice for given currents and densities;

5. Gather fields and push beam particles located in slice ζ from t to $t + \Delta t$.

Since the fields are interpolated per slice to the beam particles, there is no requirement to allocate the memory of the 3D simulation domain, which is a key feature to harness the full computing potential of a GPU. It drastically reduces the needed memory of the algorithm, which now only consists of the beam particles, a 2D slice of plasma particles, and a few 2D slices. These typically fit in the memory of a single GPU. Therefore, all quantities required for the computation are directly allocated in the GPU memory, minimizing the slow transfer between the host and the GPU. Host-device transfers are only needed for input/output (I/O) and the longitudinal parallelization. Additionally, since all required data is directly allocated in the GPU memory, single-GPU fast Fourier transforms (FFTs) can be used, which are significantly faster than single- or multi-CPU FFTs.

Furthermore, simulations can be accelerated by using multiple GPUs in parallel. Since the fields and the plasma response are calculated successively in a loop over the longitudinal slices, a quasi-static PIC code can only be parallelized longitudinally via pipelining algorithms [54, 152]. In these pipelining algorithms, each GPU calculates a different time step in a staggered way. Since host-device memory transfers are expensive for GPUs, it is important to minimize both the number of memory transfers between the GPUs and their size. A temporal domain decomposition was developed, such that the amount of data to be transferred decreases with the number of used GPUs. This novel longitudinal parallelization allows for near-optimal scaling to hundreds of state-of-the-art GPUs, allowing to fully harness the computing power of modern supercomputers.

Besides the the advantage in raw computing power, GPUs are more energy-efficient than CPUs. Exactly determining the gain in energy-efficiency on a supercomputer is complex. Therefore, the following simple estimate is used to determine the gain in energy-efficiency of HiPACE++ when running on the GPU-equipped supercomputer JUWELS Booster instead of running on the CPU-only supercomputer JUWELS. The energy used per simulation is estimated as the product of the run time on GPU and CPU, the fraction of number of compute nodes used, and the total power consumption of JUWELS for CPUs and the JUWELS Booster for GPUs, respectively. Using the example from the GPU-versus-CPU performance benchmark in the corresponding publication, running HiPACE++ on the GPUs of the JUWELS Booster is estimated to be $\sim 10\times$ more energy-efficient than using the CPUs on JUWELS.

More details on the GPU-optimized pipeline can be found in Sec. 5.2 of the corresponding publication.

Software implementation practices

HiPACE++ is built on top of two C++ libraries: first, AMReX [185], which provides the data structures, Message Passing Interface communications [109], and a performance-portability layer, and second, the openPMD-api [71], which provides the I/O handling. The portability layer allows the code to run on CPU, on modern NVIDIA GPUs and on ROCm-capable AMD GPUs. In support of open science, the code fulfills modern scientific coding practices [166] and has an open repository [159], is well-documented, ensures robustness due to continuous integration testing, and can be installed using modern package managers such as Spack [57].

CONCLUSION

Particle-in-cell simulations are an invaluable tool for modelling plasma-based accelerators, especially for positron acceleration schemes. Three-dimensional simulations of positron beams with collider-relevant parameters, i.e., low emittance and a small transverse beam size, are extremely challenging. In this chapter, the novel code HiPACE++ was introduced, which combines the advantages of the large time steps in a quasi-static PIC code with an optimized algorithm to utilize the immense computing power of modern GPUs. HiPACE++ achieves orders of magnitude speed-up on modern GPUs over CPU implementations and scales efficiently to hundreds of GPUs. Thereby, it enables modelling previously unattainable scenarios or parameter scans of previously demanding simulation settings, such as challenging positron acceleration schemes or the proton-driven plasma accelerator AWAKE [62].

With this new, powerful tool at hand, the numerical basis for the following physical studies is founded.

Key findings of the corresponding publication

- The first GPU-capable, quasi-static 3D particle-in-cell code HiPACE++ is developed.
- A modified quasi-static PIC loop enables orders of magnitude speedup on modern GPUs over a Central Processing Unit (CPU) implementation.
- Utilizing modern GPUs is estimated to be $\sim 10\times$ more energy-efficient, reducing the carbon footprint of the simulations.
- A novel longitudinal parallelization allows for near-optimal scaling to hundreds of cutting-edge GPUs.

Part III

POSITRON ACCELERATION IN A PLASMA
COLUMN

CONCEPT

Recently, a novel concept for positron acceleration in a beam driven plasma accelerator has been proposed in Diederichs et al. [46]: positron acceleration in a plasma column. The concept relies on an electron drive beam and a finite-width plasma column as an accelerating medium and enables high field gradients and emittance-preserving positron acceleration. However, other requirements for a collider, i.e., low-energy-spread positron acceleration and robustness to transverse instabilities, were not considered in the original publication.

In this part, the scheme is studied in depth: low-energy-spread and low-emittance positron acceleration is demonstrated, the stability of the drive and witness beams are shown, and the effect of initial misalignment on the witness beam quality is evaluated. Non-ideal column profiles are examined and temperature effects are considered. But first, the basic principle of the concept and the field structure and its properties are introduced.

Setup of positron acceleration in a plasma column

The concept of positron acceleration in a plasma column relies on an electron drive beam propagating on the column axis and driving a blowout wake with a blowout radius R_b larger than the column radius R_p . In the case of a blowout wake in a homogeneous plasma, the transverse focusing field increases linearly in r up to the bubble radius [101], before it decays, close-to exponentially, in the plasma electron sheath [183]. In contrast, in the case of a plasma column with a radius R_p smaller than the blowout radius R_b , the transverse wakefield increases linearly only up to the column radius, $\propto r$ for $0 < r \leq R_p$, and then falls off as $\propto r^{-1}$ for $R_p < r \leq R_b$, before it quickly decays in the plasma sheath for $r > R_b$. Due to the lack of ions outside of the column radius, the transverse wakefield at the location of the electron sheath is reduced in comparison to the homogeneous plasma. This reduction of the focusing field leads to a modification and, most importantly, to a spread of the electron trajectories in the plasma sheath. As a consequence, the longitudinal position where the electrons cross the propagation axis of the drive beam strongly depends on their initial transverse position in the plasma column. Furthermore, the electrons return in a wide-spread area behind the blowout bubble, creating an elongated, high-density electron filament. This electron filament in turn provides strong focusing and accelerating fields for positrons. The scheme is visualized in Fig. 5.1.

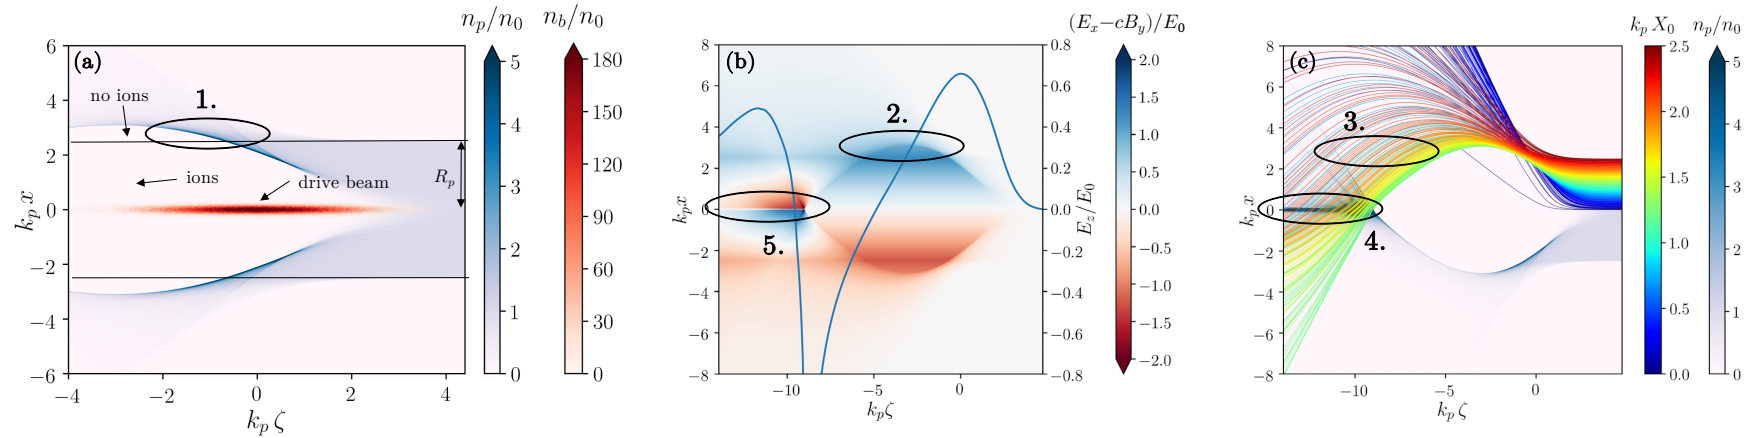


Figure 5.1: Schematic of positron acceleration in a plasma column. (a) Electron density (blue) and drive beam density (red), (b) focusing wakefield (2D plot) and on-axis accelerating field (blue line), and (c) electron trajectories (rainbow-colored lines). **1.** An electron drive beam expels the plasma electrons beyond the column radius R_p . **2.** The lack of ions beyond the column radius reduces the focusing field in that region. **3.** The reduced focusing field leads to a spread of electron trajectories. **4.** The electrons return on axis in a long, high-density electron filament. **5.** The filament leads to an elongated region with accelerating and focusing fields for positrons.

Field structures in a plasma column

The transverse focusing field $(E_x - cB_y)/E_0$ and the accelerating field E_z/E_0 in the region suitable for positron acceleration determine the evolution of the bunch and its parameters, i.e., the emittance, energy gain, and energy spread.

As an example, a setup is considered that consists of an electron drive beam with a Gaussian density profile with RMS sizes of $\sigma_x = \sigma_y = 0.1 k_p^{-1}$ in the transverse and $\sigma_z = 1.41 k_p^{-1}$ in the longitudinal direction, and a beam current of $I_b/I_A = 1$, with $I_A = 4\pi\epsilon_0 m_e c^3/e \approx 17$ kA being the Alfvén current. As a target, a plasma column with a radius of $R_p = 2.5 k_p^{-1}$ is used. Many plasma physics phenomena scale with the background plasma density n_0 , therefore more general results can be obtained by normalizing the length scales to plasma skin depth k_p^{-1} and the fields to the cold, non-relativistic wave-breaking field E_0 . The limits of the normalization are discussed in the last section of this chapter.

For the example setup, the resulting accelerating field E_z/E_0 in the x - ζ plane is shown in Fig. 5.2. The inset shows its transverse dependence at three different longitudinal positions within the region suitable for positron acceleration, namely $k_p\zeta = -12.5$ (dashed line), $k_p\zeta = -11.5$ (solid line), and $k_p\zeta = -10.5$ (dotted line). As one can see, the accelerating field decreases for increasing distance from the propagation axis. Unlike in the transversely uniform accelerating field of blowout in a homogeneous plasma for electrons, a positron beam with a finite width obtains a per-slice uncorrelated energy spread in this type of wake. The longitudinal dependence of the accelerating field is shown in the on-axis field in Fig. 5.1 (b). In the presented case, the peak of the accelerating field for the positrons is roughly $\approx 75\%$ of the peak decelerating field of the drive bunch, enabling accelerating gradients on the order of 10 GV/m for plasma densities on the order of $n_0 \sim 10^{17} \text{ cm}^{-3}$. The field is non-uniform along the propagation axis, such that a beam acquires a correlated energy spread, or chirp, during the acceleration process. This can be prevented by using tailored positron current profiles that flatten the accelerating field, as discussed in Sec. 6.3.

The focusing wakefield $(E_x - cB_y)/E_0$ along the x - ζ plane is shown in Fig. 5.3. The inset shows the transverse dependence at the same three longitudinal positions $k_p\zeta = -12.5, -11.5, -10.5$. The focusing field has a step-like jump at the column axis, before it decays almost linearly for increasing distances from the axis. Although an initially Gaussian witness beam cannot be perfectly matched to these non-linear fields, a quasi-matched optimal spot size can be determined such that the emittance growth is minimized during the acceleration process. To




Figure 5.2: Two-dimensional (ζ, x) map of the accelerating wakefield, E_z/E_0 . Positrons can be accelerated in the region $-14 \lesssim k_p \zeta \lesssim -10$. Inset: transverse dependence of accelerating field in the positron accelerating region at three different longitudinal locations denoted by the dashed ($k_p \zeta = -12.5$), solid ($k_p \zeta = -11.5$), and dotted ($k_p \zeta = -10.5$) lines. The accelerating field falls off for increasing distance from the propagation axis. The transverse gradient of the field decreases further behind the driver.




Figure 5.3: Two-dimensional (ζ, x) map of the focusing wakefield, $(E_x - cB_y)/E_0$. Positrons can be focused in the region $-14 \lesssim k_p \zeta \lesssim -9$, which widely overlaps with the positron accelerating region. Inset: transverse dependence of focusing field at three different longitudinal locations denoted by the dashed ($k_p \zeta = -12.5$), solid ($k_p \zeta = -11.5$), and dotted ($k_p \zeta = -10.5$) lines. The focusing field decays almost linearly for increasing distances from the propagation axis. The field decreases further behind the driver.

allow for analytical estimations, the focusing field is assumed to be step-like in the transverse direction x , namely

$$(E_x - cB_y)/E_0 = -\alpha \text{sgn}(x), \quad (5.1)$$

where α is the field strength and $\text{sgn}(x)$ is the sign function. Then, from a simplified model, the optimal RMS size σ_x that minimizes the emittance growth in such a wake was found to be [46]

$$\sigma_x^3 \simeq 1.72 \frac{\epsilon_x^2}{k_p \alpha \gamma}, \quad (5.2)$$

where ϵ_x is the transverse emittance and γ the Lorentz factor of the beam, respectively. For a single beam slice with an initial emittance of $\epsilon_x = 0.1 k_p^{-1}$, the estimated emittance growth at saturation for a quasi-matched beam was found to be $\simeq 2\%$. In PIC simulations, the emittance growth for a single slice was found to be $\approx 3\%$, which is in reasonable agreement with the model in the view of the approximations.

In conclusion, the field structures in the presented setup are suitable for high-gradient positron acceleration. Despite the non-linear focusing fields, the emittance growth for an initially Gaussian beam is small, and saturates as soon as the beam is matched to these non-linear focusing fields.

Realistic plasma column targets

The experimental realization of positron acceleration in a plasma column strongly depends on the availability of the required plasma sources. Therefore, it is important to determine how sensitive the scheme is to the column radius or shape. In the original publication, the column radius was scanned numerically, and the optimal column radius was found to scale as [46]

$$k_p R_p \approx 2 \sqrt[3]{2I_b/I_A} \quad (5.3)$$

for drive beam currents I_b between $1 \leq I_b/I_A \leq 10$. Additionally, a reasonable flexibility on the column radius was found, i.e., a tolerance of a few percent. So far, only idealized columns have been considered and the scheme must be robust against realistic profiles obtainable in the experiment. Therefore, realistic profiles are considered below.

Plasma columns can be generated by optical field ionization (OFI) or via beam field ionization of the drive beam itself. The latter has the advantage that the generated column is inherently aligned with the drive beam. However, since the beam-field-induced ionization rate is strongly coupled to the drive beam, the usable parameter space is limited. Additionally, head erosion is a concern for a beam propagating in a neutral gas [8], further restricting a stable column

generation. Therefore, OFI-generated plasma columns by a laser seem to be preferable.

Plasma columns have been generated as targets for beam-driven plasma accelerators at FACET [60], but have attracted great attention as laser waveguides in the context of laser-driven wakefield acceleration recently [6, 48, 111, 145, 146, 149]. The radial density profiles obtained in experiments [111, 145, 146] differ notably from a perfect column. Therefore, the applicability of the positron acceleration scheme in these realistic profiles must be demonstrated. After hydrodynamic expansion [146], the OFI-generated radial density profiles can be approximated by a parabolic channel multiplied with an exponential decay, namely

$$n_e(r) = n_{e,0} \left(1 + \frac{r^2}{w_M^2} \right) \exp \left[- \left(\frac{r}{d} \right)^s \right] \quad (5.4)$$

with r being the radius, $n_{e,0}$ the on-axis density, w_M the characteristic channel radius, d the decay constant, and s the exponent, respectively. This description allows for different, column-like density profiles. Various density profiles and the resulting on-axis accelerating wakefield (using the same drive beam as before) are shown in Fig. 5.4 (a) and (b), respectively. Despite the vastly different shapes, the accelerating wakefield and the focusing wakefield are very similar to those in the simple column. Note that each of these shapes has been optimized and has a tolerance of a few percent for achieving fields suitable for positron acceleration.. Interestingly, the ratio of the peak positron-accelerating and the peak drive-beam-decelerating field in case of the dashed orange line is notably higher than that of the simple column. Thus, tailoring the transverse density profile, as it is possible in OFI-generated density profiles, can increase the efficiency of the scheme.

Temperature effects

Besides the smooth density profiles obtained from OFI using circular polarized lasers, these realistic profiles have a temperature of roughly tens of eV for common laser technologies such as titanium-sapphire laser systems. The thermal electron pressure broadens the electron peak at the back of the blowout bubble [138, 141] and could potentially disrupt the positron accelerating wakefield structure. The effect of a temperature on the wakefields is presented in Fig. 5.5.

The on-axis accelerating field at plasma temperatures of 5, 15, 25, and 50 eV for the setup introduced in the beginning of this chapter is displayed in Fig. 5.5 (a). The accelerating field is hardly affected by a temperature and even at 50 eV the positron-accelerating part of the wakefield is not visibly changed. In contrast, the focusing field is significantly altered, as shown along the transverse coordinate in the

Temperature effects for positron acceleration in a plasma column are discussed in detail in Diederichs et al. [43]

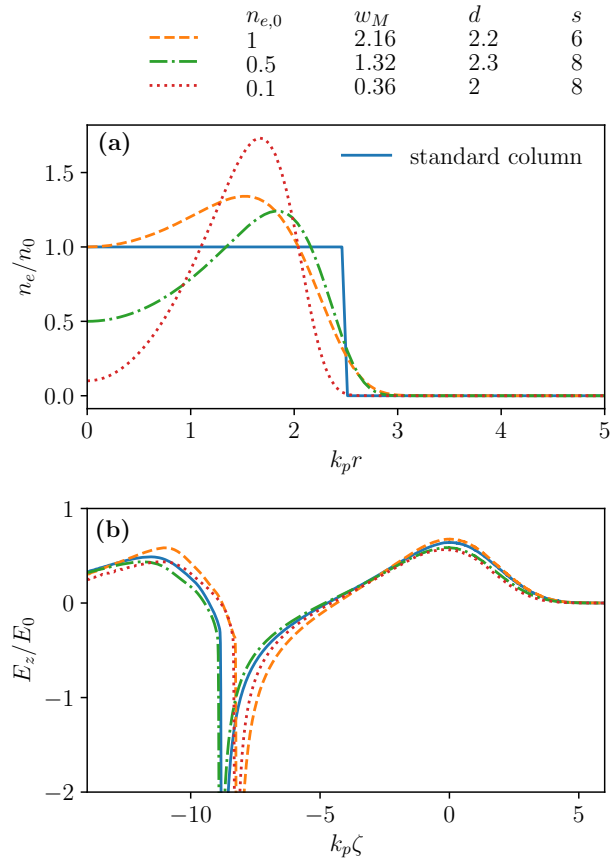


Figure 5.4: (a) Various OFI-like, radial density profiles obtained from Eq. 5.4 and the given parameters. (b) The respective on-axis accelerating wakefield E_z along the co-moving variable ζ .




Figure 5.5: (a) Longitudinal wakefield, E_z/E_0 , versus co-moving variable, $k_p \zeta$ and (b) transverse wakefield amplitude, $(E_x - cB_y)/E_0$, versus transverse position, $k_p x$ for various plasma temperatures. While a temperature of up to 50 eV has almost no effect on the accelerating field, it reduces the transverse wakefield amplitude and broadens its linear region around the zero crossing.

region of a positron bunch at a longitudinal position of $k_p\zeta = -10.5$ in Fig. 5.5 (b). The step-like behavior of the focusing field presented in the original publication [46] is smeared out and the focusing field has a linear part on the column axis that widens with increasing temperature. This effect has also been reported in a warm, elongated bubble regime [173]. The step-like behavior only occurs for an infinite density spike, as it is only possible in a cold plasma [40]. With a temperature, the thermal electron pressure limits the on-axis density spike and increases the width of the on-axis electron filament, leading to the smeared-out focusing field. Note that these linear focusing fields could in principle be used for emittance-preserving positron acceleration. Thus, a temperature, as it is always prevalent in realistic plasma columns, does not significantly change the accelerating field but modifies the focusing field in a favorable way by decreasing its transverse non-linearity. As a consequence, a temperature reduces the emittance growth of the positron beam [43]. On a side note, a temperature also alleviates numerical convergence in quasi-static PIC codes because they typically struggle with divergence of the on-axis density spike in cold plasmas [43, 73].

Density limits

So far, the physics of positron acceleration in a plasma column have been discussed using normalized units, because the results scale with the plasma density. However, the approach cannot be scaled to arbitrary densities. In fact, if the density is too high, since the wakefield amplitude scales as $E \sim E_0 \propto n_0^{1/2}$, the wakefield within the blowout radius (which exceeds the column radius) can ionize the neutral gas outside of the column, causing the column to expand and perturb the positron accelerating and focusing fields. In case of a plasma column consisting of ions that are not fully stripped of all their electrons, e.g., He^+ , the wakefield or the space charge field of the drive beam can ionize the plasma in the column to a higher ionization level. The plasma electrons generated by this additional ionization can disrupt the wakefield structure and deteriorate the positron accelerating wakefields. Therefore, for given drive beam parameters and depending on the gas used, the plasma density must be below some maximum value where additional ionization occurs.

For the sample drive beam parameters used throughout this section, the maximum density is determined by scanning the background plasma density with PIC simulations. The ionization in HiPACE++ is modelled via tunneling ionization based on the Ammosov-Delone-Krainov (ADK) model [7]. The on-axis accelerating fields for various densities using a helium plasma ionized to the first level and a hydrogen plasma are shown in Fig. 5.6 (a) and (b), respectively. The maximum densities in a helium plasma and hydrogen plasma for




Figure 5.6: Longitudinal wakefield, E_z/E_0 , versus co-moving variable, $k_p\zeta$ for various plasma densities in (a) helium and (b) hydrogen. For higher densities, additional ionization occurs and disrupts the accelerating wakefield.

the given drive beam are $\simeq 1 \times 10^{17} \text{ cm}^{-3}$ and $\simeq 5 \times 10^{16} \text{ cm}^{-3}$, respectively. The focusing field is also not affected at these densities. In principle, higher densities can be reached in a helium plasma ionized to the second level or lowering the drive beam current. The required laser intensity to ionize helium to the second level is on the level of $\sim 10^{16} \text{ W/cm}^2$, which entails new challenges, as a resonant laser pulse in that regime can already drive a small wake itself [74], thus, it is not further discussed here.

Particle colliders have strict requirements on the beam quality. As discussed in Ch. 1, besides a low emittance and high bunch charge, the energy spread among the beam particles is crucial to achieve a high luminosity. Due to chromatic effects, a low energy spread is needed to tightly focus the beam at the interaction point to decrease the beam spot size and increase the luminosity. The energy acceptance of the final focusing system determines the maximum tolerable energy spread. As an example, the RMS energy spread design value for the CLIC is 0.35% [4], which is a substantial challenge for plasma-based accelerators.

Typically, one differentiates between the energy spread within a beam slice (the uncorrelated energy spread) and the energy spread between beam slices (the correlated energy spread, often called *chirp*). Uncorrelated (slice) energy spread is caused by non-uniformity in the accelerating field within each slice (see inset of Fig. 5.2). Correlated energy spread occurs when different slices see different accelerating fields. The latter can be compensated for by an effect called *beam loading*, which works as follows.

High-charge particle bunches generate strong electromagnetic fields, which affect the fields within the accelerating structure. The resulting net field is the field that ultimately accelerates the bunch. By shaping the current distribution of the to-be-accelerated bunch, the accelerating field can be flattened [103], so that uniform acceleration is possible, leading to low energy spreads. While the beam-loading effect is a simple superimposition in the case of a linear wake [75], it is more challenging in case of high-charge witness bunches, as used in the non-linear regime. Then, the fields of a witness bunch act on the plasma, changing the plasma dynamics and thus significantly affecting the wakefields.

6.1 OPTIMAL BEAM LOADING IN PLASMA-BASED ACCELERATORS

Optimal beam loading in the blowout regime

The blowout regime is most interesting for high-gradient electron acceleration in a uniform plasma if background ion motion can be neglected. The uniform ion background provides a linear focusing field that is constant for all slices of the witness bunch. Additionally, the accelerating field is radially constant within the blowout cavity. Thus, the energy spread is only induced by the longitudinal variation

of the accelerating field. Due to the simple wake structure, analytical estimates are possible, although a description of the field structure itself is challenging.

A seminal work by Tzoufras et al. [164] derived a theoretical description of the effect of the witness bunch on the resulting wakefield based on the model by Lu et al. [101] (which has been recently extended to increased accuracy [39]). In the ultra-relativistic limit and assuming a maximal blowout radius much larger than the plasma skin depth, the ideal bunch profile to flatten the accelerating field in the blowout regime was found to be trapezoidal. This is shown in Fig. 6.1, which shows an optimally beam-loaded wake in the blowout regime. The witness beam profile was numerically calculated, see Sec. 6.2. Another important finding from the theory is that the efficiency in the optimally-loaded blowout regime is constant, i.e., that the product of the witness beam charge and the (flattened) accelerating field $Q_w \times E_{acc}$ is constant.

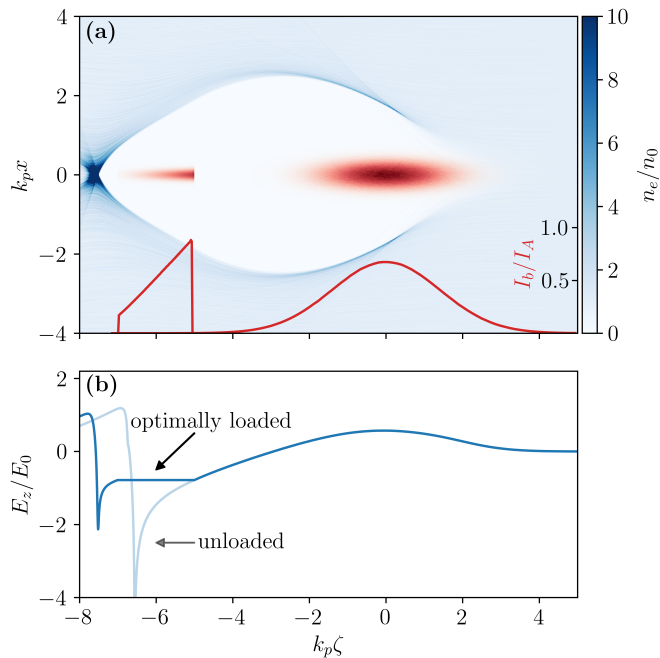


Figure 6.1: Optimally beam-loaded wake. (a) Two-dimensional (ζ, x) map of the plasma electron density n_e/n_0 (blue) and the beam distributions (red). The drive and witness beam current profiles I_b/I_A are denoted by the red line. The trapezoidal witness beam current flattens the accelerating wakefield E_z/E_0 along the co-moving variable ζ as shown in (b). Thus, the wake is optimally loaded and the witness beam does not obtain a correlated energy spread under the assumption that the wakefield stays constant during acceleration.

Still, optimal beam loading is an open topic of research, as the theory is relatively restrictive and did not include a description of the driver. A recent effort used the Broyden–Fletcher–Goldfarb–Shanno

algorithm and a simulation parameter scan to optimally load the blowout wake with Gaussian witness bunches [174]. However, the algorithm is intricate and does not reveal the optimal witness current shape. This limitation is resolved in Sec. 6.2, where an algorithm that numerically calculates the beam current to flatten an arbitrary wakefield is presented.

Close-to-optimal beam loading in the blowout regime has recently been achieved experimentally by Lindström et al. [88], preserving per-mill level energy spread in a beam-driven plasma accelerator.

Note that laser beams propagate in a plasma at a group velocity slower than the speed of light. Thus, the ultra-relativistic witness bunch catches up with the laser driver and outruns the accelerating phase, a phenomenon called dephasing. Consequently, it is more difficult to achieve uniform average acceleration along the plasma stage in a laser-driven plasma accelerator. Flattening the average wakefield in a laser-driven plasma accelerator has been studied in Schroeder et al. [136] and demonstrated experimentally in Kirchen et al. [77], achieving energy spreads as low as 1.2%.

Optimal beam loading for positron acceleration

Although in the linear regime optimal beam loading can be solved analytically for both a uniform plasma [70] or hollow core plasma channels [139, 189], the non-linear positron acceleration schemes [37, 46, 72, 150, 188] are considered more relevant due to the higher accelerating gradients and positron bunch charge. However, so far no general analytic theory was found for these non-linear regimes due to their complex structure. Plasma-based positron acceleration requires areas with a large excess of plasma electrons. Due to the mass asymmetry of the plasma, the electrons are much more mobile. Thus, unlike the almost static ions in the blowout regime to an electron witness bunch, the plasma electrons react much stronger to a high-charge positron bunch, resulting in complex, non-linear behavior that is difficult to describe analytically. Recently, an analytic description was found for the beam loading of a positron bunch in the electron cusp at the back of the blowout [187], but this is not generally applicable for other positron acceleration schemes.

Another challenge of these non-linear electron structures is the radial dependence of the accelerating field. Thus, optimal beam loading in these structures does not only involve flattening the field longitudinally, but also in the transverse directions, which has not been demonstrated yet.

So far, beam loading in these non-linear positron schemes has been optimized by trial-and-error, mostly using Gaussian-shaped witness bunches. This typically results in energy-spreads on the percent-level, too large for a collider application. In the next section, an algorithm is

presented that utilizes the quasi-static approximation to numerically calculate the witness beam current profile that flattens the average accelerating field in an *arbitrary* wakefield – optimal for non-linear positron acceleration schemes.

Corresponding publication

S. Diederichs, C. Benedetti, E. Esarey, J. Osterhoff, and C. B. Schroeder,
High-quality positron acceleration in beam-driven plasma accelerators
Phys. Rev. Accel. Beams **23**, 121301 (2020) [41]

6.2 THE SLICING ADVANCED LOADING ALGORITHM FOR MINIMIZING ENERGY SPREAD (SALAME)

*A similar approach
has been used by
Lotov [97, 98]*

The following algorithm utilizes the slice-by-slice calculation of the wakefields in quasi-static PIC codes to recursively construct, starting from a predetermined position, the optimal current profile of a witness bunch that flattens the accelerating field along the bunch. The algorithm is introduced in the context of positron acceleration, but it equivalently works to generate electron witness bunches.

A radially symmetric bunch is considered with a profile of the form

$$n_b(\zeta, r) = g_{\parallel}(\zeta)g_{\perp}(\zeta, r) \quad (6.1)$$

where $g_{\parallel}(\zeta)$ and $g_{\perp}(\zeta, r)$ denote the longitudinal and transverse slice-dependent density profiles, respectively. Thereby, it is required that $\int g_{\perp}(\zeta, r)rdr = \int g_{\perp}(\zeta = \zeta_{head}, r)rdr$ for any ζ , where ζ_{head} is the location of the bunch head, such that the bunch current density profile only depends on $g_{\parallel}(\zeta)$. Although the algorithm works for arbitrary transverse density profiles, first only radially Gaussian and longitudinally uniform profiles g_{\perp} are considered for simplicity, i.e.,

$$g_{\perp}(r) = \exp\left(-\frac{r^2}{2\sigma_r^2}\right) \quad (6.2)$$

with σ_r being the radial RMS bunch size.

The algorithm iterates from the location of the bunch head backwards over every longitudinal grid point and determines the optimal beam current via $g_{\parallel}(\zeta)$ that flattens the average accelerating field $\langle E_z \rangle$ in that particular beam slice. The accelerating field E_z is averaged over

the transverse beam distribution to take into account the transverse dependence of E_z via

$$\begin{aligned}\langle E_z \rangle &= \frac{\int_0^\infty E_z(r) g_\perp(r) r dr}{\int_0^\infty g_\perp(r) r dr} \\ &= \frac{\int_0^\infty E_z(r) \exp[-r^2/(2\sigma_r^2)] r dr}{\int_0^\infty \exp[-r^2/(2\sigma_r^2)] r dr}.\end{aligned}\quad (6.3)$$

Note that, in case of a transversely uniform accelerating field, e.g., as in a blowout wake, the average simplifies to the on-axis accelerating field.

The algorithm is incorporated in the quasi-static PIC algorithm as follows: First, the quasi-static plasma response, which is calculated as a loop over the longitudinal grid points, is evaluated up to the slice where the witness bunch head is supposed to start. There, the average accelerating field at the head $\langle E_{z,\text{head}} \rangle$ is determined. Then, at each slice with the longitudinal index i and the respective averaged accelerating gradient $\langle E_{z,i} \rangle$, the following procedure, which is visualized in Fig. 6.2, is performed to find the optimal beam current:

1. The initial condition $|\langle E_{z,i} \rangle| > |\langle E_{z,\text{head}} \rangle|$ is evaluated. If it is not satisfied, no further beam loading is possible and the algorithm terminates. The absolute value of the average field is used so that the algorithm works to generate both electron and positron witness beams.
2. To determine the optimal beam-loading current with a bisection method, beam density values that are lower ($g_{\parallel,\text{min}}$) and higher ($g_{\parallel,\text{max}}$) than the optimal one are required. Due to the initial condition, $g_{\parallel,\text{min}} = 0$ is a valid first value. $g_{\parallel,\text{max}}$ is determined by generating a beam slice with g_{\parallel} and evaluating $\langle E_{z,i} \rangle$ by solving the quasi-static equations. Iteratively, g_{\parallel} is increased (e.g. by a factor of 10 per iteration) until the wake is overloaded, i.e., $|\langle E_{z,i} \rangle| < |\langle E_{z,\text{head}} \rangle|$. Then, $g_{\parallel,\text{max}} = g_{\parallel}$ and $g_{\parallel,\text{min}}$ is set to the last g_{\parallel} that was not overloading the wake.
3. A modified bisection method is used to determine the optimal beam current from $g_{\parallel,\text{min}}$ and $g_{\parallel,\text{max}}$. The new value of the beam density is given by

$$g_{\parallel} = w_g g_{\parallel,\text{min}} + (1 - w_g) g_{\parallel,\text{max}}, \quad (6.4)$$

where the weight w_g is given by

$$w_g = \frac{|\langle E_{z,\text{head}} \rangle| - |\langle E_{z,\text{min}} \rangle|}{|\langle E_{z,\text{max}} \rangle| - |\langle E_{z,\text{min}} \rangle|}, \quad (6.5)$$

with $\langle E_{z,\text{min}} \rangle$ and $\langle E_{z,\text{max}} \rangle$ being the averaged fields on the slice with a beam density of $g_{\parallel,\text{min}}$ and $g_{\parallel,\text{max}}$, respectively. Using this

With the explicit field solver in HiPACE++ the bisection is not needed, since the current to flatten E_z can be calculated directly, see [151].

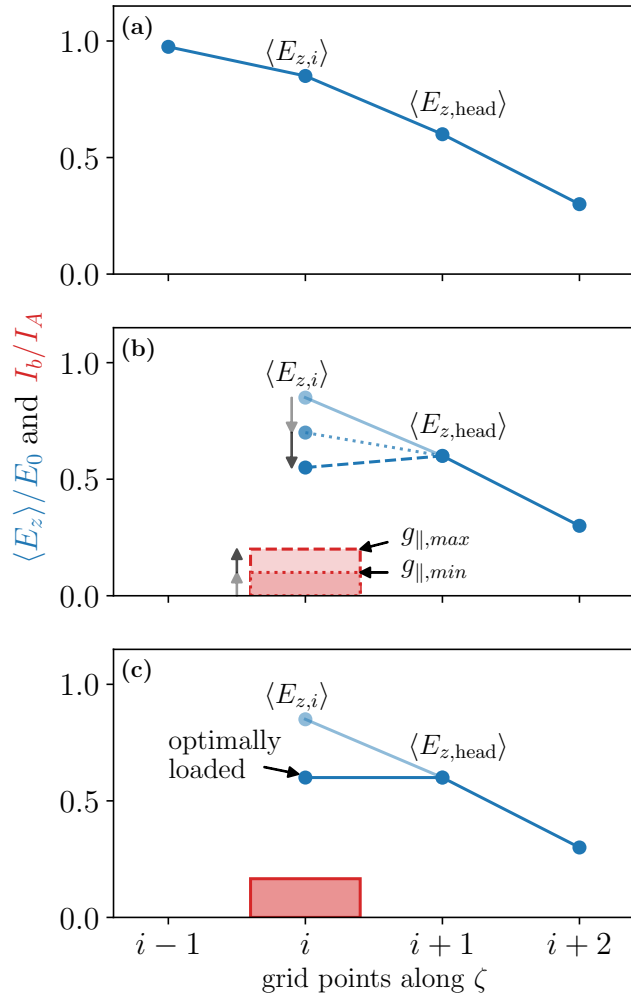


Figure 6.2: Visualization of the optimal beam loading algorithm : (a) Ensure that the initial condition $|\langle E_{z,i} \rangle| > |\langle E_{z,\text{head}} \rangle|$ is satisfied, otherwise terminate. (b) By step-wisely increasing the current (red bars) on grid point i , the $\langle E_{z,i} \rangle$ (blue line) decreases until the wake is overloaded (i.e. $|\langle E_{z,i} \rangle| < |\langle E_{z,\text{head}} \rangle|$). The overloading current determines $g_{\parallel,\text{max}}$, the last underloading current $g_{\parallel,\text{min}}$. (c) A modified bisection method using $g_{\parallel,\text{min}}$ and $g_{\parallel,\text{max}}$ is used to determine the optimally loading current on grid point i . (a)-(c) are repeated in a loop over the longitudinal grid points, until the initial condition is violated and the wake cannot be loaded any further

modified bisection method has shown to converge with far fewer iterations than the standard bisection method with $w_g = 0.5$. The bisection method terminates and moves to the next slice, if $\langle E_{z,i} \rangle$ converges to $\langle E_{z,\text{head}} \rangle$ within a predetermined tolerance, otherwise $g_{\parallel,\text{min}}$ and $g_{\parallel,\text{max}}$ are updated and a new bisection is performed.

The bunch current profile obtained from this procedure flattens the average accelerating field, preventing a correlated energy spread under the assumption that the bunch size does not change during acceleration. However, this assumption is generally not true. If the size of a beam slice is not matched to the transverse focusing field, it will evolve until it is matched (and suffer some emittance growth). Additionally, during the acceleration, the matched spot size reduces adiabatically with increasing particle energy. In a radially non-uniform accelerating field, a changed spot size samples different parts of the accelerating field, leading to a correlated energy spread. To fully eliminate the uncorrelated energy spread, these effects need to be taken into account, which is briefly summarized in next paragraph and explained in more detail in the corresponding article.

The mismatch can be mitigated by introducing a slice-dependent bunch size $\sigma_r(\zeta)$ and numerical matching of the bunch size to the focusing field per slice. This slice-by-slice matching also minimizes the emittance growth [18]. To mitigate the effect of the adiabatic spot size reduction due to acceleration, the scaling of the matched spot size with the particle energy must be known. Then, if one anticipates the final energy of the bunch at the end of the accelerator stage, an average transverse spot size during the acceleration can be calculated and used to weight the averaged accelerating field in Eq. 6.3. Note that these advanced techniques are not needed in a simple blowout wake for electron acceleration, because then the accelerating field does not depend on the transverse extent of the beam, provided that the beam is fully contained in the ion cavity, but only for a radially non-uniform accelerating field as encountered in various positron acceleration schemes.

6.3 LOW-ENERGY-SPREAD POSITRON ACCELERATION

The positron-accelerating wake in a plasma column generated by an electron drive beam is highly non-linear. No analytic description of the optimal current profile to flatten the field has been found so far, therefore it is a prime candidate to be optimized using the algorithm presented in Sec. 6.2. Here, the optimal beamloading in the plasma column configuration is studied by varying the starting position of the witness bunch head and using the SALAME to generate an optimally loaded wake. The current profiles (red lines) for the optimally loaded




Figure 6.3: Current profiles I_b/I_A (red) for an optimally loaded wake and the corresponding averaged accelerating field $\langle E_z \rangle / E_0$ (blue) along the co-moving variable ζ . The charge values for the different witness beam profiles are given assuming a background plasma density of $n_0 = 5 \times 10^{17} \text{cm}^{-3}$.

wake at different starting positions and the corresponding average accelerating field $\langle E_z \rangle / E_0$ (blue lines) are shown in Fig. 6.3.

The optimal current profile is close to Gaussian, with a little shift to the head of the bunch. Interestingly, placing the bunch head further in the front of the wake, leading to a lower acceleration gradient, does not necessarily increase the witness bunch charge Q_w . Thus, the efficiency is not constant and does depend on the starting position of the bunch, which is in contrast to the blowout regime for electrons, as demonstrated in Tzoufras et al. [164]. Defining the instantaneous driver-to-witness efficiency as

$$\eta = \frac{Q_w \times E_w^+}{Q_d \times E_d^-}, \quad (6.6)$$

where Q_d is the drive beam charge and E_w^+ and E_d^- are the average energy gain and loss rates per particle of the witness and drive beam, respectively. The highest efficiency in this setting was found to be $\eta \approx 3\%$, for a witness bunch head position of $-10.8 \lesssim k_p \zeta_{head} \lesssim -10.6$. This is less than the $\eta = 4.8\%$ achieved with a simple Gaussian-shaped witness beam current in Diederichs et al. [46], however, at the cost of a higher energy spread.

As explained in Sec. 6.2, an evolving beam in a radially non-uniform accelerating field still obtains a correlated energy spread. Via slice-by-slice matching to the non-uniform focusing field along the bunch, and incorporating the spot size reduction due to the bunch acceleration (it scales as $\sigma_{r,matched} \propto \sqrt[3]{1/\gamma}$ according to Eq. 5.2 in Sec. 5), the cor-




Figure 6.4: Mean energy of each slice vs. longitudinal position in the positron bunch after acceleration. The blue line refers to algorithm flattening $\langle E_z \rangle$ with a longitudinal uniform σ_r . The red line and the green line refer to additionally applying slice-by-slice matching and averaging of the bunch spot size over the acceleration distance, respectively. That way, the correlated energy spread is reduced to the noise level.

related energy spread can be fully mitigated. As shown in Fig. 6.4, the simple field flattening algorithm (blue line) still exhibits a correlated energy spread of $\Delta E = 100$ MeV after acceleration from 1 GeV to ≈ 5.5 GeV. The slice-by-slice matching (red line) reduces the correlated energy spread to $\Delta E = 60$ MeV, and finally, after taking into account the spot size reduction due to the energy gain, the correlated energy spread is suppressed to the noise level. Note that these additional corrections do not significantly change the current profiles and charge values.

Reducing the uncorrelated energy spread

Including the corrections, the SALAME is able to fully suppress a correlated energy spread. Due to the transverse non-uniformity of the accelerating field (see Fig. 5.2) and the finite width of the beam, an uncorrelated energy spread cannot be prevented, it can only be reduced since, so far, no beam-loading strategy to flatten E_z transversely has been identified.

In the following, the dependence of the uncorrelated energy spread on the witness beam emittance is investigated. A matched witness bunch with a smaller emittance has a smaller transverse width and samples a smaller transverse domain of E_z , and, thus, obtains a smaller uncorrelated energy spread. The radial dependence of the accelerating field in the vicinity of the axis can approximately be described as $E_z(r) = E_{z,0} - \beta r$, with $E_{z,0}$, β , and r being the on-axis accelerating field, the transverse gradient of E_z , and the radius, respectively. Then, from

geometric considerations, one can expect that the relative uncorrelated energy spread of a beam slice at saturation scales as $\sigma_\gamma/\gamma \sim \beta\sigma_r/E_{z,0}$, and so $\sigma_\gamma/\gamma \rightarrow 0$ in the limit of a small bunch.

Unfortunately, it is not possible to model much smaller emittance beams with the currently available numerical tools, as they require extreme transverse resolutions and tremendous computational resources. Therefore, a reduced model is employed to provide some insight into the scaling to smaller emittances. Since the correlated energy spread can be fully eliminated, only a single slice of the bunch is considered in the reduced model. The slice of the peak current of the positron beam is chosen, which was found to reasonably represent the total energy-spread of the bunch. In the reduced model, test particles are advanced with a second-order-accurate particle pusher in the static fields extracted from the PIC simulation. The reduced model shows reasonable agreement in terms of energy spread and emittance with the PIC simulation for an initial emittance of $\epsilon_x = 0.05 k_p^{-1}$, which corresponds to $\epsilon_x = 0.38 \mu\text{m}$ at a background density of $n_0 = 5 \times 10^{17} \text{cm}^{-3}$. Then, using the reduced model, the emittance can be reduced to previously numerically inaccessible values. The relative uncorrelated energy spread for $0.38 \mu\text{m}$ (blue line for reduced model, dashed-black line for PIC simulation), $0.19 \mu\text{m}$ (red line), and $0.08 \mu\text{m}$ (green line) along a propagation of 15 cm are depicted in Fig. 6.5.

More details on the reduced model can be found in Sec. III C of the corresponding publication.

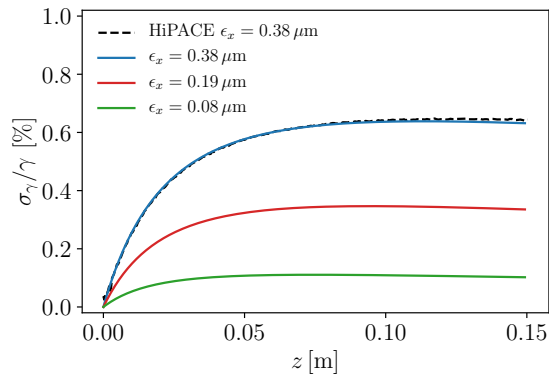


Figure 6.5: Relative slice energy spread vs. acceleration distance. Advancing test particles in an approximated step function yields similar energy-spread in comparison with the HiPACE simulation. The results indicate that emittances smaller than $0.1 \mu\text{m}$ induce an energy spread below 0.1%.

The results indicate that for emittances below $0.1 \mu\text{m}$, the relative uncorrelated energy spread is below 0.1%, sufficient for a collider. Sub-per-mill energy spreads have previously not been anticipated for plasma-based positron acceleration and this indicates an important step to collider-relevant parameters. It should be noted that the obtained results assume a static wakefield and do not capture the effect

of a reduced positron beam size on the wake itself. Eventually, one could expect a disruption of the positron accelerating wake structure for an extremely intense positron bunch. However, beams smaller than the electron filament can currently not be modelled even with current numerical tools, and will be subject of future studies.

Key findings of the corresponding publication

- A novel algorithm to numerically determine the optimal beam current to flatten the accelerating field is introduced.
- Tailoring the positron current profile allows to fully suppress the correlated energy spread.
- Decreasing the witness beam emittance decreases the uncorrelated energy spread that is obtained during acceleration.
- Plasma columns potentially enable sub-micron emittances and sub-per-mill energy spread positron acceleration.

In the previous chapter, positron acceleration in a plasma column has shown promising results in terms of beam quality, making it subject of interest for collider applications. Nonetheless, these promising results hold only true if, and only if, the system is stable towards small asymmetries, which are inevitable in an actual accelerator. For example, the infamous instability of the hollow core plasma channel (see Ch. 3) has made its application challenging in practise. Since positron acceleration in a plasma column relies on a cylindrically symmetric wakefield structure, one might expect it to be prone to similar beam breakup instabilities. In this chapter, the stability of positron acceleration is studied in terms of drive and witness beam misalignments, and the induced effect on the witness beam quality. The chapter is organized as follows: In Sec. 7.1, the most prevalent beam instabilities for plasma-based accelerators are introduced. In Sec. 7.2 the stable propagation of the electron drive beam through a plasma column is demonstrated. Finally, in Sec. 7.3, the witness beam stability and the deteriorating effects of drive and witness beam misalignment are evaluated.

7.1 TRANSVERSE BEAM INSTABILITIES

Beam breakup instabilities pose a major challenge for linear colliders [31, 80, 123]. Small misalignments between the beam and the accelerator can amplify exponentially, blow up the emittance, and thus, decrease the luminosity, or in the worst case, lead to total beam loss. In conventional accelerators, these instabilities limit the charge that can be safely accelerated and hence, determine the achievable efficiency. In plasma-based accelerators, the extreme transverse field gradients enhance the rapid growth of these instabilities. Therefore, a study came to the conclusion that the maximum charge in plasma-based colliders in the blowout regime is severely limited [81], however, the study overlooked important damping mechanisms [104].

In general, beam breakup instabilities can be mitigated by introducing a betatron frequency spread along the beam, which is often referred to as the Balakin-Novokhatsky-Smirnov (BNS) damping mechanism, named after the proponents [14]. Although BNS damping was originally proposed for conventional accelerators, the same underlying physics are nowadays also applied in plasma-based accelerators in various regimes [83, 104, 107, 122, 150].

The BNS damping mechanism has been recently reviewed again by Novokhatski [120]

In the following, the most relevant transverse instability in plasma accelerators, the hosing instability, is introduced.

Hosing instability in plasma-based accelerators

The so-called *hosing instability*, first identified by Whittum et al. [177], arises when initially misaligned beams couple coherently to their induced plasma response. For example, a beam that is initially misaligned with respect to its propagation axis, e.g., by a tilt, induces a shift of the plasma wake centroid that in turn acts on the beam and, in the worst case, exponentially amplifies the initial tilt, leading to beam breakup.

A detailed review of the hosing instability in plasma-based accelerators is given in Mehrling et al. [108]

The hosing instability is particularly prevalent in the blowout regime due to its longitudinally constant and transversely linear focusing fields, which enables a coherent coupling for a mono-energetic bunch. It was initially found to grow exponentially [69, 177], threatening applicability of the blowout regime. Later, the relativistic mass change of the electrons in the plasma sheath [69] and the collective behavior of the sheath [105] were found to reduce the coupling. An energy chirp (inducing a BNS-damping-like betatron frequency spread) was found to saturate the hosing instability [107]. However, the required energy spread is on the order of few percents, which may be too much for the witness beam in a collider application. Therefore, this mechanism can only be utilized to stabilize the drive beam propagation. Additionally, large drive beams, for which the head does not witness a full blowout wake, and hence, a nonlinear focusing force that varies along the head, were reported to be more stable against hosing [122]. Again, this mechanism is inapplicable to the witness beam in the blowout wake.

For the witness beam stabilization, another mechanism must be utilized. Intense witness beams, as required for a linear collider, provide space charge fields that modify the background ion distribution, leading to a longitudinally-varying non-linear focusing field. Thus, a betatron spread along the beam is induced that effectively damps the oscillation [104]. Note that, these non-linear focusing fields induce emittance growth [10], although the emittance growth can be mitigated by matching the beams to these non-linear fields [17, 18]. Alternatively to the blowout regime, other regimes, like the quasi-linear regime, naturally provide longitudinally-varying focusing fields that induce a betatron frequency spread that suppress the hosing instability [83].

In conclusion, the hosing instability is a concern for plasma-based accelerators, although many BNS-like mechanisms have been identified to efficiently suppress it. Analogously to conventional accelerators, a trade-off between the stabilizing effects of ion motion and energy spread versus the ion-motion-induced emittance growth and the energy spread acceptance in the collider must be found. The trade-off determines the possible efficiency in plasma-based accelerators and

needs to be studied in the future, as the aforementioned study [81] neglected important damping effects and thus, misses the key parameter trade-off. This optimization study is out of the scope of this thesis.

7.2 ELECTRON DRIVE BEAM STABILITY

Since the positron accelerating wakefield structure in a plasma column is based on cylindrical symmetry, the drive beam stability is of utmost importance. Fortunately, unlike the hollow core channel, the drive beam itself is subject to the transverse wakefield in the plasma column, which provides focusing and, as shown later, stability. The plasma column (compared to a homogeneous plasma) introduces a new symmetry axis though, therefore it is important to determine the robustness of the scheme against misalignment between the beams and the plasma column.

Previous work has reported on the long-range attraction of an electron beam towards a neutral plasma column [2] and on the deflection of an electron beam at the plasma-gas boundary [99, 116, 118]. Furthermore, a severe misalignment between the column and the electron beam was found to seed the hosing instability [8]. In case of an isotropic acceleration medium, e.g., a homogeneous plasma, the hosing instability is seeded only by asymmetries of the beam itself. In a plasma column, both the beam asymmetries and misalignment between the beam and the column axis can seed the hosing instability. To assess the robustness against the hosing instability of an electron beam in the plasma column versus a homogeneous plasma, first asymmetric beams, i.e., tilted beams, are considered, before the column-specific robustness against misalignment with respect to the column center is evaluated. The results are discussed in depth in the publication below and the core findings are presented in the following.

Corresponding publication

S. Diederichs, C. Benedetti, E. Esarey, M. Thévenet, J. Osterhoff, and C. B. Schroeder,

Stable electron beam propagation in a plasma column

Physics of Plasmas **29**, 043101 (2022) [42].

Tilted drive beams

Modelling the hosing instability has significantly contributed to understanding the relevant physics. Advanced analytical models have identified the relativistic mass change of the plasma electrons [69] as well as the collective behavior in the blowout sheath [105] as mitigation mechanisms to the hosing instability. Recent efforts [105] describe the plasma response based on an analytic model of the blowout wake [183],

including the collective behavior of the plasma sheath. In this work, the model is extended to capture finite-radius plasma columns, as required for the proposed positron acceleration scheme. Details on the extended model are presented in [Sec. II A](#) and in the [Appendix](#) of the corresponding publication. The model, confirmed by [PIC](#) simulations, demonstrates that a tilted beam traversing a plasma column is slightly more robust against the hosing instability than it is traversing a homogeneous plasma. The model shows that the increased robustness is mostly caused by the increased plasma sheath thickness in the column, which arises because of the larger spread of the electron trajectories due to the reduced focusing field in the area of the sheath. The increased plasma sheath thickness decreases the coupling of the beam centroid to the plasma centroid, leading to the observed increase of stability. However, the plasma column has an additional degree of freedom, namely that the propagation axis of the beam must be aligned with the center of the drive beam, which is discussed in the next subsection.

Drive beams misaligned with respect to the column axis

Similarly to small beam asymmetries, the whole beam centroid can be displaced with respect to the column axis, e.g., due to machine jitters. For a transversely displaced beam, a restoring force is acting on the beam in a plasma column, unlike the deflecting force in the hollow core channel. When all plasma electrons of the column are expelled beyond the column radius, more background ions are on the opposing site of the displaced beam. The transverse wakefields are mostly caused by the ion distribution, therefore the wake centroid X_p is then not aligned with the beam centroid X_b , but instead shifted towards the column center. Thus, the beam is pulled towards the column axis. The wake centroid X_p depends on the co-moving variable ζ , since the ion distribution depends on the blowout radius, as illustrated in [Fig. 7.1](#). The ζ -dependence of the wake centroid causes the head and the tail to experience different restoring forces, potentially seeding the hosing instability.

The evolution of a transversely displaced drive beam is studied via [PIC](#) simulations. The previously mentioned analytical model is not applicable, because, as explained in the publication, one of the key assumptions of the model is not met for fully displaced beams. The simulations reveal that, indeed, the hosing instability is seeded due to the different restoring force from head to tail. To mitigate the induced oscillation, most of the [BNS](#)-like damping mechanisms discussed in [Sec. 7.1](#) are applicable. The damping via an ion-motion-induced, longitudinally varying focusing fields for a helium plasma is illustrated in [Fig. 7.2](#), where the tail of the drive beam, which is most susceptible to the hosing instability, is plotted against the

Mitigating the hosing instability due to large beams [122] is not reasonable here, as the hosing is only induced at the center and tail of the beam.

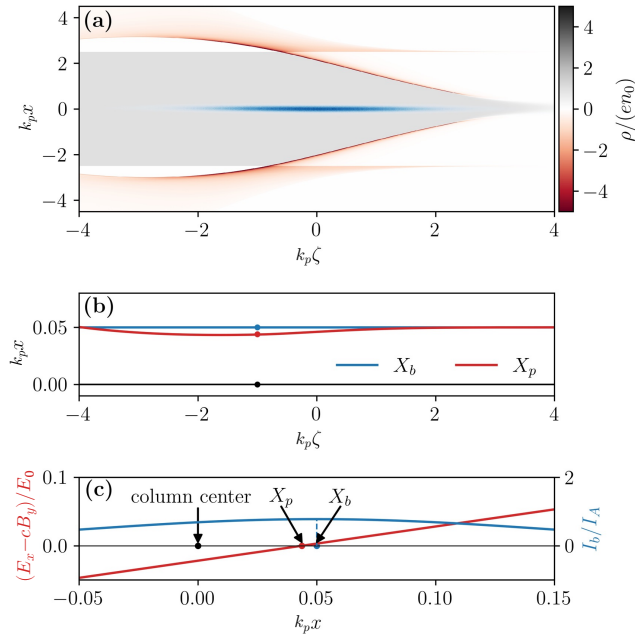


Figure 7.1: Characterization of the wake structure of a beam with a transverse offset with respect to the column axis. (a) Snapshot of the plasma charge density in the x - ζ -plane for a beam (shown in blue) with an initial transverse offset of $X_b = 0.05 k_p^{-1}$ with respect to the plasma column center. (b) The resulting ζ -dependent wake centroid X_p (red line) that lays between the beam centroid X_b (blue line) and the center of the column (black line), leading to an attraction of the beam towards the center of the column. (c) Lineout of the beam current profile (blue line) and focusing wakefield (red line) along x at $\zeta = -1 k_p^{-1}$.

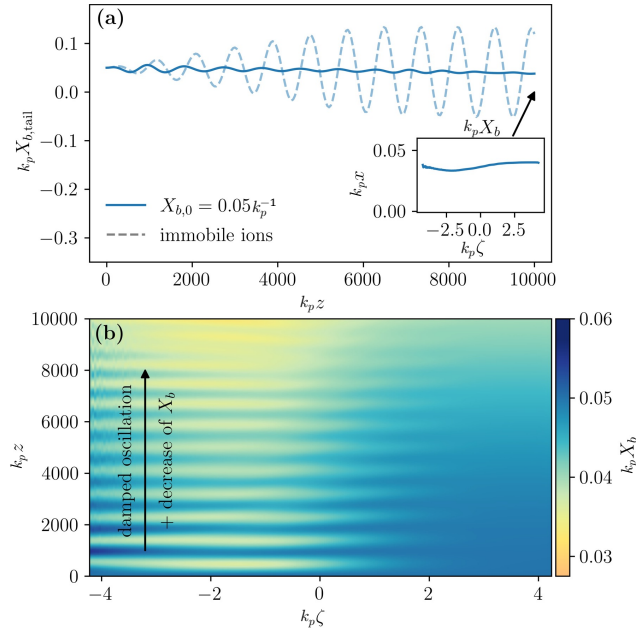


Figure 7.2: Evolution (a) of the beam centroid at the tail of the beam $X_{b,tail}$ along the propagation distance for an initial offset of $X_{b,0} = \sigma_x = 0.05 k_p^{-1}$ assuming an immobile ion background (dashed line) and with mobile ions in a Helium plasma (solid line). The inset shows the final slice-dependent beam centroid. Waterfall plot (b) with the evolution of the beam centroid for $X_{b,0} = 0.05 k_p^{-1}$ in Helium vs longitudinal coordinate (horizontal axis) and propagation distance (vertical axis). The beam centroid decreases via a damped oscillation towards the column center.

propagation distance. In the case of immobile ions (dashed line) an oscillation is induced, which saturates at an amplitude on the order of the initial offset. Assuming a helium plasma (solid line), the oscillation is damped, leaving dominantly a drift towards the column axis. The final ζ -dependent bunch centroid is shown in the inset of Fig. 7.2 (a). In agreement with the predicted wake centroid X_p in Fig. 7.1, the tail of the beam has drifted closer to the column axis than the head. The ζ -dependent temporal evolution of the full bunch is shown in Fig. 7.2 (b). While the tail and the center of the bunch undergo oscillations, the head slowly drifts towards the axis. The oscillations are damped, leaving only a slow drift towards the column axis.

Besides transversely displaced beams, pointing jitters also deteriorate the alignment between beam and column axis. Pointing jitters are tested via PIC simulations using a beam with an initial transverse momentum (results shown in Sec. IV of the corresponding publication). Thereby, it is shown that the column hardly affects the beam for small pointing jitters and that the beam is bent towards the column axis for large jitters, in agreement with previous studies [99, 116, 118]. Stable propagation requires initial pointing angles of $X'_{b,0} \ll R_p/L_p$ with R_p and L_p being the column radius and the plasma length, respectively.

Consequently, for beam parameters similar to FACET-II and assuming a plasma column with density $n_0 = 1 \times 10^{17} \text{cm}^{-3}$ and a radius of $R_p = 42 \mu\text{m}$, a pointing jitter below $\lesssim 10 \mu\text{rad}$ is required to allow for stable beam propagation and thus stable wakefield structures usable for positron acceleration.

In conclusion, the propagation of an electron beam in a plasma column is stable. Although initial misalignment induces hosing, it can be effectively mitigated by ion motion or an energy spread. Furthermore, the intrinsic stability allows for alignment of the beam via active feedback loops [8], since the emitted betatron radiation by an oscillating bunch scales with the amplitude of the oscillation, which in turn scales with initial offset. Thus, the emitted betatron radiation scales with the initial offset, allowing for alignment of the electron beam and the plasma column by minimizing the emitted betatron radiation.

Key findings of the corresponding publication

- The electron drive beam is intrinsically attracted to the column center.
- Initial misalignment of the drive beam and the column axis induces the hosing instability.
- The hosing instability is efficiently damped by BNS-like mechanisms in form of ion motion or energy spread.
- The overall drive beam propagation through a plasma column is stable.

7.3 POSITRON WITNESS BEAM STABILITY

After the drive beam stability has been shown in the previous section, the witness beam stability is investigated. Currently, there is no analytic model that describes the highly non-linear field structure in the region suitable for positron acceleration. Therefore, the witness beam stability is investigated via 3D PIC simulation using HiPACE++ and a reduced model in this section. The findings are described in detail in the publication below and the key findings are presented in the following.

Corresponding publication

S. Diederichs, C. Benedetti, E. Esarey, M. Thévenet, J. Osterhoff, and C. B. Schroeder,
Self-stabilizing positron acceleration in a plasma column
Phys. Rev. Accel. Beams **25**, 091304 (2022) [45].

In the simulations, the same drive beam and plasma parameters are used as before. To preclude the hosing suppression due to a slice-dependent energy chirp [107], a witness beam profile that optimally loads the wake (see Sec. 6.2) is used. Since the current profiles obtained in Sec. 6.3 are non-trivial, a linear combination of two Gaussian current profiles is used for simplicity that captures the salient features of the profiles presented in Sec. 6.3.

The full simulation settings are provided in Sec. III of the corresponding publication

The evolution of the witness beam centroid X_w is studied for various initial misalignments of the witness beam centroid $X_{w,0}$, drive beam centroid $X_{d,0}$, or both. The results are shown in Fig. 7.3, where the difference between the witness beam centroid and the wakefield centroid $\langle W \rangle$ is plotted. The wakefield centroid is subtracted to take into account that the wakefield centroid is shifted for a displaced drive beam. As one can see, the witness beam quickly converges to the wakefield centroid in a strongly damped oscillation in all presented cases.

The oscillation is damped by two effects: first, the longitudinal varying focusing field induces a varying betatron frequency (see Sec. 7.1), which has also been observed in other positron acceleration schemes [150]. Second, the transverse non-linearity of the focusing field causes phase mixing within a single slice, ultimately damping the witness beam centroid motion.

Damping via transverse non-linearity of the focusing field

The scaling of the damping due to the transverse non-linearity of the focusing field is investigated via a simplified 1D model. Thereby, the focusing field is again assumed to be step-like

Details on the simplified model are presented in the Appendix A of the corresponding publication

$$\frac{E_x - cB_y}{E_0} = -\alpha \operatorname{sgn}(x), \quad (7.1)$$

where x is the transverse coordinate, α that the strength of the wakefield, and $\operatorname{sgn}(x)$ the sign function. Furthermore, the model neglects drive beam evolution (i.e., the focusing field is assumed to be constant), head-to-tail effects (i.e., a single slice within the beam is considered), and assumes the particles are not significantly accelerated within a betatron period. Finally, the initial centroid displacement is assumed to be small compared to the transverse RMS beam size, i.e., $X_{w,0} \ll \sigma_{x,w}$. Under these assumptions, the evolution of the beam centroid X_w is found to be a damped oscillation with a damping length S_{damp} scaling as

$$k_p S_{\text{damp}} \propto \sqrt{\frac{k_p \sigma_{x,w} \gamma}{\alpha}} \quad (7.2)$$

where γ is the Lorentz factor of the beam. Notably, and confirmed by the simulations, the damping length does not depend on the initial

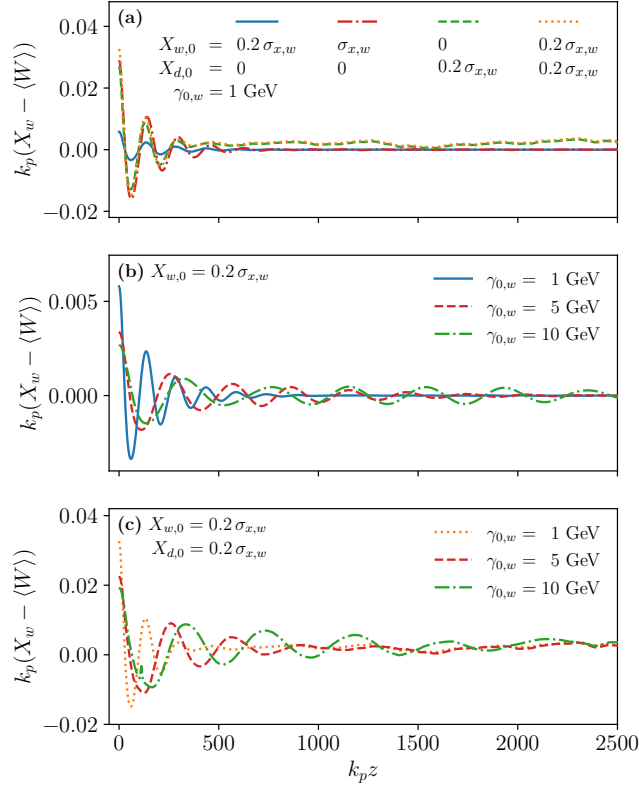


Figure 7.3: (a) Difference between witness beam centroid X_w and focusing wake centroid $\langle W \rangle$ for different initial witness and/or drive beam offsets as a function of the propagation distance. The witness beam centroid quickly converges to the wake centroid within a few damped oscillations, demonstrating the stability of the scheme versus initial offsets. The cases of a misaligned driver (b) and both a misaligned driver and misaligned witness beam (c) are also shown for higher witness energies of 5 and 10 GeV, respectively. Despite an increased damping length, the evolution is still stable.

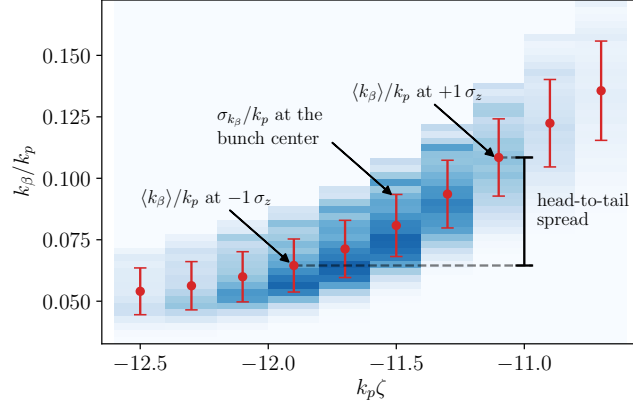


Figure 7.4: Distribution of betatron wavenumbers k_β/k_p along the co-moving variable ζ . The rms wavenumber spread per slice σ_{k_β} is depicted by the red bars. The head-to-tail wavenumber spread, defined as the difference between the mean wavenumber at the head of the bunch (located at $+1\sigma_z$) and the tail (located at $-1\sigma_z$) is depicted by the black bar. The head-to-tail spread is roughly four times larger than the average of the intra-slice spread.

offset, at least for small displacements. The scaling with γ , α , and $\sigma_{x,w}$ is in good agreement with the simulations. Note that the limit of $\sigma_{x,w} \rightarrow 0$ violates the underlying assumption of the model $X_{w,0} \ll \sigma_{x,w}$ and is therefore not described by the model. In fact, in the limit of an infinitesimal beam size, the per-slice betatron spread becomes zero (similarly to the uncorrelated energy spread in Sec. 6.3), thus the damping length becomes infinite.

Damping via longitudinally varying focusing field

The damping caused by the longitudinally varying focusing field is harder to quantify, since it strongly depends on the witness bunch parameters, i.e., the beam loading and bunch length. To quantify, whether the transverse non-linearity or the longitudinal variance of the focusing field is dominating the damping process, the betatron wavenumber distribution along the bunch can be extracted numerically. Using the example with an initial witness beam centroid offset of $X_{w,0} = 0.2\sigma_{x,w}$, the betatron wavenumber spread from the head to the tail of the bunch is found to be four times larger than the betatron spread within a single slice, as shown in Fig. 7.4. Thus, for this example, the longitudinally varying focusing field is the dominating effect in the damping process.

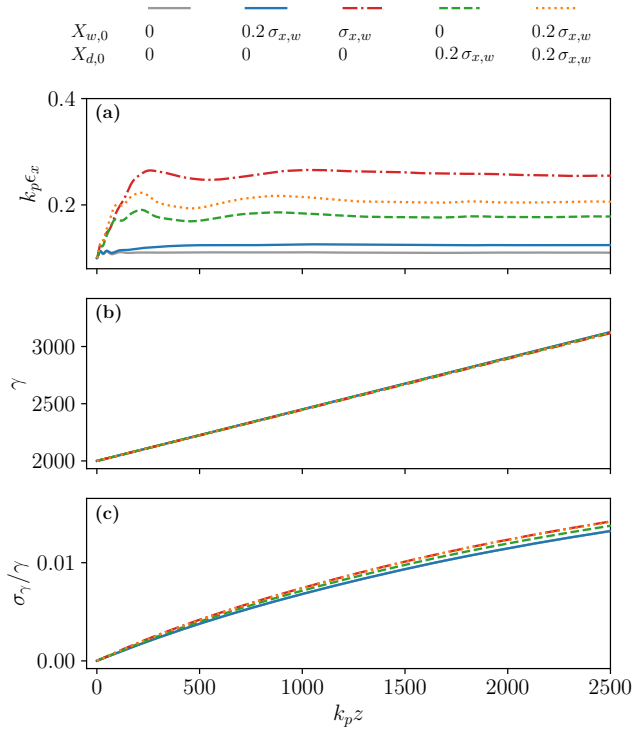


Figure 7.5: (a)–(c) Emittance, energy gain, and relative energy spread as a function of the propagation distance. Depending on the offset, emittance growth can be observed, which saturates as soon as the beam is aligned with the wake centroid. Only marginal differences in energy gain and energy spread are observed.

Effect of misalignment on the positron beam quality

Having discussed the stability, now the effect of initial misalignments on the positron beam quality is addressed. For the cases presented in Fig. 7.3, the effect of the misalignment on the emittance, energy, and relative energy spread of the witness bunch are shown in Fig. 7.5 (a)–(c). While a small witness beam offset (solid blue line) increases the emittance by only a few percent in comparison to the aligned case (solid grey line), a large witness beam offset (dash-dotted red line), a drive beam offset (dashed green line), or a drive and witness beam offset (dotted orange line) significantly deteriorate the emittance. Still, the emittance growth saturates as soon as the oscillation observed in Fig. 7.3 is damped and beam centroid converges to the wake centroid. This is even the case for an evolving wake centroid, as it occurs for an misaligned drive beam. The energy gain and relative energy spread are hardly affected by initial misalignments. Additionally, the strong focusing fields fully capture the positron bunch and the bunch charge is conserved in all studied cases.

Since the damping of the oscillation is based on phase mixing, a certain emittance growth is to be expected in the process. To gain

A full derivation is given in the Appendix B of the corresponding publication

further understanding of the scaling with the initial offset, the theory for emittance growth at saturation presented in Diederichs et al. [46] is extended to capture initial misalignments. Then, for an initial displacement of the witness beam $\Delta_x = X_{w,0}/\sigma_{x,w}$, the emittance growth and saturation is given by

$$\frac{\sqrt{\langle x^2 \rangle \langle u_x^2 \rangle}}{\sigma_x \sigma_{u_x}} \approx \left\{ \frac{8}{45} \left[\left(1 + \frac{4}{\pi} \right) (1 + \Delta_x^2) + \sqrt{\frac{2}{\pi}} \eta^{-1} (2 + 3\Delta_x^2) + \frac{5}{2} \sqrt{\frac{2}{\pi}} \eta \left(1 + \frac{\Delta_x^2}{2} \right) + \frac{3}{4} \eta^2 \right] \right\}^{1/2}, \quad (7.3)$$

with

$$\eta = \frac{\sigma_{u_x}^2}{k_p \sigma_x \gamma \alpha}. \quad (7.4)$$

While the emittance growth predicted by the theoretical model shows reasonable agreement with the case of a small witness beam offset, it significantly underestimates the emittance growth for larger initial misalignments. The difference arises because the model does not capture the beam-loading effect on the focusing field. A misaligned, high-charge witness beam asymmetrically modifies the focusing field, leading to an increased emittance growth that is not covered by the model. Therefore, the model provides a reasonable scaling only for small initial offsets $\Delta_x \ll 1$. Additionally, since some of the assumptions made in the model are quite rigorous (flat beam, neglect of head-to-tail effects, neglect of acceleration) the model should be used for qualitative purposes only.

In conclusion, the longitudinally varying and transversely non-linear focusing fields provide stability since initial offsets are quickly damped due to phase mixing. This mechanism comes at a cost of emittance growth. Therefore, tight tolerances in terms of beam alignment must be met in the context of a linear collider, to limit the emittance growth and subsequent luminosity decrease.

Key findings of the corresponding publication

- Initial misalignment of the drive beam, witness beam, or both induces the hosing instability to the witness bunch.
- The longitudinally varying and transversely nonlinear focusing fields efficiently damp the oscillation of the hosing instability.
- An approximate scaling for the damping length and the induced emittance growth is given.
- Overall, positron acceleration in a plasma column is stable. As expected, misalignments between the beams and the plasma column affect the witness beam quality.

CONCLUSION AND OUTLOOK

Conclusion

The particle physics community has demonstrated interest towards a lepton collider with a center-of-mass energy above 10 TeV, which cannot be met by conventional RF-based accelerators at a reasonable cost. Thus, an energy frontier machine of this kind requires a paradigm shift in the accelerator technology. Plasma accelerators are a promising candidate due to their 10-100s of GV/m field gradients that could reduce the construction costs drastically. Although plasma acceleration of electrons has made rapid progress in the last decade, the acceleration of positrons is still challenging, even conceptually. Different schemes have been proposed and high-gradient positron acceleration has been demonstrated experimentally, but the results were plagued with low beam quality or transverse instabilities. On top of that, even modelling positron acceleration in the collider-relevant regime is challenging, since this requires resolving small beams and can be numerically extremely expensive.

This dissertation has served the field of plasma-based positron acceleration in two ways:

First, the portable, quasi-static, 3D particle-in-cell code HiPACE++ was developed. By adopting the quasi-static PIC algorithm and utilizing modern HPC programming standards, HiPACE++ is able to fully harness the immense compute power of modern GPUs, enabling orders of magnitude speedup in comparison with equivalent CPU-based codes. As a consequence, previously unattainable beam parameters can be modelled with HiPACE++ at reasonable computational cost. Since the code is open-source, the whole community benefits from the development.

And second, the concept of positron acceleration in a plasma column has been significantly advanced. Low-emittance, low-energy-spread positron acceleration has been demonstrated. To obtain the low energy spread, an algorithm was introduced that calculates the witness bunch profile that optimally loads an arbitrary plasma wake. The algorithm is beneficial to all positron and electron acceleration schemes and has already been adopted by others [187]. Meeting the strict energy spread requirements of a collider is an important step to its realization. The transverse stability of positron acceleration in a plasma column has been demonstrated. Both the drive and the witness beam are inherently stable to small transverse misalignment.

The positron beam quality is affected by an initial misalignment, therefore precise alignment of the beams with the plasma column is

crucial to the realization of the scheme. Notably, the stability of the positron bunch arises due to the transversely non-linear and longitudinally non-uniform focusing field structure, which has also been observed in other positron acceleration schemes. Furthermore, the concept of positron acceleration in a plasma column has been demonstrated in realistic setups. Smooth transverse plasma profiles and temperature effects have been studied. Both can be beneficial to the acceleration (smooth plasma profiles) or focusing (temperature effects) of positron bunches. The density limit for certain beam parameters was evaluated and is in reasonable range for the experimental realization.

Outlook

Despite the promising progress this work has shown for positron acceleration in a plasma column, the concept requires further optimization. The efficiency must be increased to tens of percent to be applicable in a collider, as otherwise the power consumption renders it prohibitively expensive. As shown in this work, the efficiency can in principle be increased by using a smooth transverse plasma profile, which has to be studied in detail. Furthermore, shaping the drive beam current can significantly increase the efficiency [85, 94, 133], which has also been utilized in other positron acceleration concepts to achieve tens of percent energy transfer efficiency [187].

After the optimization of the efficiency, positron acceleration in a plasma column needs to be compared and evaluated against the other positron acceleration schemes. Currently, a fair comparison is difficult to make, as the publications use vastly different witness beam parameters, i.e., different emittances, which consequently leads to different efficiencies and energy spreads. Therefore, in the spirit of the [ESPP Accelerator R&D Roadmap \[13\]](#), the community should agree on positron beam emittance as a reference value. Then, the beam loading algorithm presented in this work can be utilized to minimize the energy spread and a proper comparison between plasma-based positron acceleration schemes can be made.

Finally, true progress towards a plasma-based collider can only be achieved by experiments. This simulation- and theory-based work serves as a stepping stone towards an experimental test facility. Based on this work, positron acceleration in a plasma column has been selected as one of the positron acceleration schemes to be tested at [FACET-II](#) at [SLAC](#). Although there are many other open questions regarding a plasma-based collider, such as staging, efficiency, or a sufficient repetition rate, the experimental realization of high-quality positron acceleration in a plasma accelerator will mark an important milestone.

Part IV

APPENDIX



AUTHOR CONTRIBUTION AND PUBLICATIONS


My contributions to the core publications of this thesis are listed in chronological order below:

1. **HiPACE++: a portable, 3D quasi-static Particle-in-Cell code** [44]
I proposed the idea of porting HiPACE to GPUs and the novel, GPU-optimized parallelization scheme. I implemented large fractions of the code, conducted most of the underlying PIC simulations and the corresponding data analysis, and mainly wrote the manuscript.
2. **High-quality positron acceleration in beam-driven plasma accelerators** [41]
I implemented the SALAME algorithm in HiPACE and adjusted it to the nonlinear fields to fully suppress the correlated energy spread. I performed all underlying simulations and the respective data analysis, and mainly wrote the manuscript.
3. **Stable electron beam propagation in a plasma column** [42]
I extended the analytic hosing model to capture finite-radius plasma columns, and conducted the numerical comparison. I performed all underlying simulations and the respective data analysis, and mainly wrote the manuscript.
4. **Self-stabilizing positron acceleration in a plasma column** [45]
I derived the analytic model to describe the positron beam emittance growth at saturation with an initial offset. I performed all underlying PIC simulations and the respective data analysis, and mainly wrote the manuscript.



Contents lists available at ScienceDirect

Computer Physics Communications

www.elsevier.com/locate/cpc


HiPACE++: A portable, 3D quasi-static particle-in-cell code ☆,☆☆

S. Diederichs^{a,b,c,*}, C. Benedetti^b, A. Huebl^b, R. Lehe^b, A. Myers^b, A. Sinn^a, J.-L. Vay^b, W. Zhang^b, M. Thévenet^a^a Deutsches Elektronen-Synchrotron DESY, Notkestr. 85, 22607 Hamburg, Germany^b Lawrence Berkeley National Laboratory, 1 Cyclotron Rd, Berkeley, CA 94720, USA^c University of Hamburg, Institute of Experimental Physics, Luruper Chaussee 149, 22607 Hamburg, Germany

ARTICLE INFO

Article history:

Received 21 September 2021
 Received in revised form 29 April 2022
 Accepted 17 May 2022
 Available online 24 May 2022

Keywords:

Particle-in-cell
 Plasma acceleration
 GPU computing
 Quasi-static approximation

ABSTRACT

Modeling plasma accelerators is a computationally challenging task and the quasi-static particle-in-cell algorithm is a method of choice in a wide range of situations. In this work, we present the first performance-portable, quasi-static, three-dimensional particle-in-cell code HiPACE++. By decomposing all the computation of a 3D domain in successive 2D transverse operations and choosing appropriate memory management, HiPACE++ demonstrates orders-of-magnitude speedups on modern scientific GPUs over CPU-only implementations. The 2D transverse operations are performed on a single GPU, avoiding time-consuming communications. The longitudinal parallelization is done through temporal domain decomposition, enabling near-optimal strong scaling from 1 to 512 GPUs. HiPACE++ is a modular, open-source code enabling efficient modeling of plasma accelerators from laptops to state-of-the-art supercomputers.

Program summary

Program Title: HiPACE++

CPC Library link to program files: <https://doi.org/10.17632/zh3rc7hvrml.1>Developer's repository link: [HiPACE++](#) [GitHub repository](#)

Licensing provisions: BSD 3-clause

Programming language: C++

Nature of problem: Modeling plasma accelerators is a computationally challenging task requiring nanometer-scale resolutions over meter-scale propagation distances. The quasi-static particle-in-cell method enables high-fidelity simulations of this strongly non-linear process, but these simulations can be very expensive.

Solution method: The quasi-static particle-in-cell algorithm is modified to enable efficient utilization of accelerated hardware, in particular with GPU computing, reducing the cost of simulations by orders of magnitude. A novel longitudinal parallelization enables excellent strong scaling of this method up to hundreds of GPUs.

Reference: [10.5281/zenodo.5358483](https://doi.org/10.5281/zenodo.5358483)

© 2022 The Author(s). Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

1. Introduction

Plasma accelerators [1,2] enable the acceleration of charged particles over short distances due to their multi-GeV/m field gradients. Although great progress in terms of beam quality and stability has recently been achieved [3–6], significant advance is still required to make plasma-accelerator-driven applications feasible. The Particle-in-Cell (PIC) method [7,8] is a reliable tool to simulate plasma acceleration, and PIC simulations play a major role in understanding, exploring and improving plasma accelerators [9–11].

* The review of this paper was arranged by Prof. David W. Walker.

☆☆ This paper and its associated computer program are available via the Computer Physics Communications homepage on ScienceDirect (<http://www.sciencedirect.com/science/journal/00104655>).

* Corresponding author at: Deutsches Elektronen-Synchrotron DESY, Notkestr. 85, 22607 Hamburg, Germany.

E-mail addresses: severin.diederichs@desy.de (S. Diederichs), maxence.thevenet@desy.de (M. Thévenet).

<https://doi.org/10.1016/j.cpc.2022.108421>

0010-4655/© 2022 The Author(s). Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

Simulation of a multi-GeV plasma-based accelerator typically requires modeling sub-micron-scale structures propagating over meter-scale distances, hence full electromagnetic PIC simulations require millions of time steps due to the Courant-Friedrichs-Lewy (CFL) condition [12], which makes them unpractical. Several methods were developed to circumvent this limitation and enable larger time steps, including running PIC in a Lorentz-boosted frame [13] or using a quasi-static approximation [14–17], both of which have proved performant for modeling of high-energy plasma accelerator stages [18–26].

Besides algorithmic improvements, further speedup can be accomplished from hardware improvement. Accelerated computing is growing in popularity in the supercomputer landscape [27], and in particular using GPUs (Graphics Processing Units) as accelerators enabled significant speedup in High-Performance Computing (HPC) applications including PIC [28–31].

The heterogeneity of processor architectures in HPC makes it difficult to maintain a portable codebase but, following modern HPC practices, this challenge can be efficiently addressed with performance-portability layers [32–34].

In this article, we present the portable, three-dimensional, open-source, quasi-static PIC code HiPACE++¹ [35]. HiPACE++ is written in C++ and is built on top of the AMReX [36] framework, which provides field data structure, Message Passing Interface (MPI) communications, and a performance-portability layer. In particular, the quasi-static PIC algorithm is adapted to accelerated computing, and HiPACE++ demonstrates orders-of-magnitude speedup over CPU implementations as well as near-optimal scaling up to hundreds of cutting-edge GPUs. These performances enable realistic simulations of $1024 \times 1024 \times 1024$ cells for 1000 time steps in less than two minutes on modern GPU-accelerated supercomputers. HiPACE++ is a new software, combining the algorithm from the legacy C code HiPACE [22] with a portability layer and specific modifications enabling GPU computing.

The article is organized as follows: Section 2 summarizes the well-known quasi-static PIC algorithm. The GPU-porting strategy is introduced in Sec. 3. Correctness of the code is demonstrated in Sec. 4. Sec. 5 presents performance results and a novel parallelization strategy improving scalability on accelerated platforms. Additional code features are highlighted in Sec. 6.

2. The quasi-static particle-in-cell algorithm

In a plasma accelerator, a driver perturbs the plasma electrons (the ions, heavier, can generally be considered immobile) and drives an electron plasma wave. While the driver can be a laser pulse or a particle beam, we hereafter focus on the case of a particle beam (beam-driven wakefield acceleration) for simplicity, as this is what is currently implemented in HiPACE++. In the wake of the driver, a witness beam of charged particles can be accelerated with a high field gradient. In most conditions (with the notable exception of witness beam self-injection), the driver and witness beams evolve on a time scale much longer than the plasma response [37]. The quasi-static approximation (QSA) treats the beams as rigid when computing the plasma response at a given beam location, hence decoupling the beam and plasma evolutions. Under this approximation, the Maxwell equations take the form of Poisson equations, and the scheme is not subject to a CFL condition. Then, the time step is determined by the smallest betatron period of the beams (the betatron period is the characteristic time scale over that a beam evolves), making it possible to use time steps orders of magnitude larger than those in conventional electromagnetic PIC [22]. The algorithm has two main parts: first, from the




Fig. 1. Snapshot of the quasi-static PIC algorithm. The 3D simulation domain is calculated slice-by-slice in a loop over the longitudinal grid points from the head of the box to its tail. Only the beam particles, a 2D slice of plasma particles, and a few 2D slices of fields are required to determine the wake in the 3D simulation domain.

distributions of the beams, compute the plasma response (computationally expensive). Second, from the plasma fields, advance the beams by one time step (computationally cheap).

For a given distribution of the beams, the plasma response is computed in the co-moving frame defined by $\zeta = z - ct$, with c being the speed of light in vacuum (the beams propagate in the $+z$ direction). A slice of unperturbed plasma is initialized ahead of the beams and pushed backwards along the ζ coordinate. At each longitudinal position, the wakefields are calculated as a 2D problem in the transverse plane. The 3D problem is then solved as n_ζ 2D transverse problems (called *slices*), with n_ζ being the number of longitudinal grid points in the simulation domain. Fig. 1 illustrates the algorithm. In the standard algorithm, the fields are calculated in the whole 3D domain before the beams are advanced by one time step. In this work, we propose to integrate the beam advance in the loop over slices, hence pushing beam particles slice by slice. With this change, all parts of the simulation are done per slice, which is a crucial condition for our performance-portability strategy, in particular on GPUs.

From Maxwell's equations and the QSA, the following field equations can be derived [21]. The wake potential $\psi = \phi - c A_z$, with ϕ and \mathbf{A} being the scalar and vector potential respectively, is obtained from

$$\nabla_\perp^2 \psi = -\frac{1}{\epsilon_0} \left(\rho - \frac{1}{c} J_z \right), \quad (1)$$

where ϵ_0 is the vacuum permittivity and ρ and \mathbf{J} are the total (beams + plasma) charge and current densities, respectively. The transverse wakefields $E_x - c B_y$ and $E_y + c B_x$ are calculated from the transverse derivatives of ψ :

$$E_x - c B_y = -\frac{\partial}{\partial x} \psi, \quad (2a)$$

$$E_y + c B_x = -\frac{\partial}{\partial y} \psi. \quad (2b)$$

The longitudinal field E_z is obtained from

$$\nabla_\perp^2 E_z = \frac{1}{\epsilon_0 c} \nabla_\perp \cdot \mathbf{J}_\perp. \quad (3)$$

The components of the magnetic field are given by

$$\nabla_\perp^2 B_x = \mu_0 (-\partial_y J_z + \partial_\zeta J_y), \quad (4a)$$

$$\nabla_\perp^2 B_y = \mu_0 (\partial_x J_z - \partial_\zeta J_x), \quad (4b)$$

and

¹ <https://github.com/Hi-PACE/hipace>.

$$\nabla_{\perp}^2 B_z = \mu_0(\partial_y J_x - \partial_x J_y), \quad (5)$$

where μ_0 denotes the vacuum permeability. All quantities except the longitudinal derivatives $\partial_{\zeta} J_x$ and $\partial_{\zeta} J_y$ are directly accessible after the current deposition. These derivatives can be obtained with a predictor-corrector loop [20,22] or by explicit integration [38,39]. Both options are available in HiPACE++, hereafter referred to as predictor-corrector or explicit method respectively, and their implementations are described in Sec. 3.2.

In the QSA PIC algorithm presented here, the fields at slice ζ are advanced in space (from $\zeta + \Delta\zeta$ to ζ) and the beam particles are advanced in time (from t to $t + \Delta t$) together, in the following sequence:

1. Gather fields and push plasma particles backwards from $\zeta + \Delta\zeta$ to ζ ;
2. Deposit plasma currents and densities;
3. Deposit beam currents and densities;
4. Solve equation (1) for ψ to calculate $E_x - cB_y$ and $E_y + cB_x$ with equation (2);
5. Solve equation (3) for E_z ;
6. Solve equation (5) for B_z ;
7. Solve equations (4) for $B_{x/y}$.
8. Gather fields and push beam particles located in slice ζ from t to $t + \Delta t$.

In the standard QSA PIC algorithm (see [20,22]), the beam operations (current deposition, field gather and particle push) are separated from the loop over slices: at the end of the loop over slices, once the fields are computed on the whole 3D domain, the 3D fields are used to advance the beam. Here, we integrate the beam operations into the 2D slice routine, which removes the necessity to allocate the memory of the 3D simulation domain. This modification is a key requirement to harness the full compute potential of a GPU, as explained in the next section. Additionally, the 3D field arrays are never used for computation and therefore do not need to be allocated, which allows for fitting high-resolution simulations on a single GPU. This enables production-quality simulations even with modest GPU resources.

3. Porting quasi-static PIC to GPU

3.1. Porting strategy

HiPACE++ is written considering modern GPU architectures with tens-of-GB global memory, relatively slow transfers between host (CPU) and device (GPU) memories, and fast atomic operations. As illustrated in Fig. 1, the data that needs to be allocated for computation is modest, and only consists in the beam particles, a 2D slice of plasma particles, and a 2D slice of grid quantities. In a vast majority of practical cases, these quantities fit in the global memory of a single GPU. Thus, these are directly allocated in the GPU memory, reducing the need for host-device transfers during computation to its minimum. Host-device transfers are only used for I/O and communication during the longitudinal parallelization, although both can be in principle circumvented by using buffering methods combined with optimized transfers, such as using NVIDIA GPUDirect. Additionally, keeping all required data directly on the GPU makes it possible to use single-GPU Fast Fourier Transforms (FFTs) that are considerably faster than single- or multi-CPU FFTs, which accounts for a significant fraction of the observed speedup.

A practical consequence of this strategy is that the full 3D domain is not needed for computation and thus never allocated on GPU, leading to much reduced memory utilization. The GPU memory limits the transverse resolution, though high (2048 × 2048 grid

points) up to extreme (8192 × 8192 grid points) resolutions are achievable with small GPUs (8 GB GPU-memory) and state-of-the-art GPUs (80 GB GPU-memory), respectively.

This has another important consequence: Since the 3D domain is not allocated on the GPU, the fields on a slice overwrite the previous values and are not known at the end of the loop over slices. For this reason, the beam operations (field gather, particle push and current deposition) must be performed per slice within the loop over slices. To that end, beam particles are sorted per slice at the beginning of each time step. Although this results in many small and inefficient kernels (a slice of beam particles contains ~1000 particles for typical simulation parameters), the beam operations take a negligible amount of time overall. Finally, field data in the full 3D domain can be stored for the purpose of diagnostics. In that case, the required data is stacked in host (CPU) memory until the last slice is computed, and then flushed to disk.

In summary, the porting strategy combines two elements: (i) fit the 2D transverse problem in GPU memory, so it can be solved without excessive communications and (ii) reduce host-device transfers to a minimum. Space is saved on device memory by not allocating 3D field arrays, thus enabling the computation of large domains on a single GPU.

3.2. Implementation

As part of the main loop over slices, the most time-consuming functions in the quasi-static PIC method are the field solver, the plasma particle pusher, and the plasma current deposition. Performance-portability is achieved via the AMReX framework [36], and the same methods as in Ref. [30] are applied: low-level loops (over all particles or over all grid points in a slice) are written in an abstract form (through a function `amrex::ParallelFor`), which is compiled into a vectorized loop (for CPU) or a GPU kernel depending on the target platform, enabling portability of a unique source code. The GPU implementation exploits fast single-GPU FFTs as well as fast atomic operations on modern GPUs (in particular for the current deposition, the most expensive particle operation). The implementation of the core functions is described in the next paragraphs.

For the **plasma particle push**, HiPACE++ uses a fifth-order Adams-Bashforth particle pusher, as described in [22]. The transverse beam position x_{\perp} , the transverse normalized momentum $u_{\perp} = \frac{1}{Mc}(p_x, p_y)$, and the normalized plasma wake potential $\Psi_p = \frac{e}{m_e c^2} \psi_p$ of each plasma particle (with M being the mass of the pushed plasma particle, m_e the mass of an electron, and e the elementary charge) are updated as follows:

$$\partial_{\zeta} x_{\perp} = -\frac{u_{\perp}}{1 + \Psi_p}, \quad (6)$$

$$\partial_{\zeta} u_{\perp} = -\frac{q}{M} \left[\frac{\gamma_p}{1 + \Psi_p} \begin{pmatrix} E_x - cB_y \\ E_y + cB_x \end{pmatrix} + \begin{pmatrix} cB_y \\ -cB_x \end{pmatrix} + \frac{cB_z}{1 + \Psi_p} \begin{pmatrix} u_y \\ -u_x \end{pmatrix} \right], \quad (7)$$

$$\partial_{\zeta} \Psi_p = -\frac{qm_e}{Me} \left[\frac{1}{1 + \Psi_p} \begin{pmatrix} u_x \\ u_y \end{pmatrix} \cdot \begin{pmatrix} E_x - cB_y \\ E_y + cB_x \end{pmatrix} - E_z \right], \quad (8)$$

with q being the charge of the particle and γ_p the Lorentz factor given by

$$\gamma_p = \frac{1 + u_{\perp}^2 + (1 + \Psi_p)^2}{2(1 + \Psi_p)}. \quad (9)$$

The beam particles are advanced by a second-order symplectic integrator. The field gather and particle push are embarrassingly parallel operations well-suited to GPU computing.

Due to the handling per slice, the **current deposition** is so far limited to zeroth order longitudinally for both plasma and beam particles, while orders 0-3 are available in the transverse direction. On GPU, the otherwise expensive current deposition is performed using fast atomic operations to global device memory.

As can be seen in Eqs. (2), (3) and (5), most fields are computed by **solving a transverse Poisson equation** and applying finite-difference operators. The Poisson equation with Dirichlet boundary conditions is solved by means of fast Poisson solvers [40], which are based on a Discrete Sine Transform (DST) of the first type. The DST is provided by the FFTW [41] library on CPU and by a custom implementation using FFTs [42] on GPU. The FFTs on GPU are provided by vendor libraries. The capability to run 2D FFTs on a single GPU instead of parallel FFTs on many CPUs is critical to provide good performance, considering that parallel FFTs require large amount of communications. Special care is needed to compute $B_{x/y}$, because of the longitudinal derivatives $\partial_\zeta J_x$ and $\partial_\zeta J_y$ in Eq. (4). The $B_{x/y}$ field solver is usually the most expensive part of the 3D QSA PIC method.

Two options are implemented for the $B_{x/y}$ field solver algorithm. The first option is a **predictor-corrector field solver**, as implemented in the legacy code HiPACE. The longitudinal derivatives $\partial_\zeta J_{x/y}$ are evaluated on slice ζ from the previously-computed slice $\zeta + \Delta\zeta$ and slice $\zeta - \Delta\zeta$ still to be computed. An initial guess is made for $B_{x/y}$, with which particles are pushed from slice ζ to $\zeta - \Delta\zeta$ where their current is deposited. The current on slice $\zeta - \Delta\zeta$ is used to calculate $B_{x/y}$ at ζ , and the procedure is repeated until a convergence criterion is reached or a maximum number of iterations is attained. Each iteration involves all PIC operations for plasma particles as well as several Poisson solves.

The second option for the $B_{x/y}$ field solver is an **explicit field solver** using analytic integration, as done in Refs. [38,39]. A 2D non-homogeneous Helmholtz-like equation must be solved (see equation (19) in Ref. [39]), for which HiPACE++ uses the GPU-capable multigrid solver provided by AMReX. The multigrid solver is an expensive operation relying on an iterative solver, but does not require multiple iterations of PIC operations.

4. Correctness

The reference setup used throughout this article consists in a typical beam-driven wakefield acceleration simulation containing a driver beam and witness beam with Gaussian distributions with rms sizes $k_p\sigma_{\perp,d} = 0.3$, $k_p\sigma_{\zeta,d} = 1.41$ and $k_p\sigma_{\perp,w} = 0.1$, $k_p\sigma_{\zeta,w} = 0.2$, where $k_p = \omega_p/c$ is the plasma wavenumber, $\omega_p = \sqrt{n_0 e^2 / (m_e \epsilon_0)}$ is the plasma frequency, and n_0 the ambient plasma density (subscripts d and w stand for driver and witness, respectively).

The driver beam is located at the origin and has a peak density of $n_{b,d}/n_0 = 10$. The witness beam is centered around longitudinal position $k_p\zeta_{0,witness} = -5$ and has a peak density of $n_{b,w}/n_0 = 100$. The electron plasma is modeled with 4 particles per cell, and the background ions are assumed to be immobile. The simulation domain in x , y , and ζ is, in units of k_p^{-1} , $(-8, 8)$, $(-8, 8)$, and $(-7, 5)$ and uses $1024 \times 1024 \times 1024$ grid points. All simulation parameters and the used software are listed in the Appendix.

Fig. 2 shows a comparison between HiPACE++, the legacy code HiPACE, and the full GPU-capable 3D electromagnetic PIC code WarpX [25]. The accelerating field E_z/E_0 , where $E_0 = cm_e\omega_p/e$ is commonly referred to as the cold non-relativistic wave breaking limit [37], shows excellent agreement between these three codes. For HiPACE++, both the predictor-corrector and the explicit solver were used, and both demonstrate again very good agreement, as

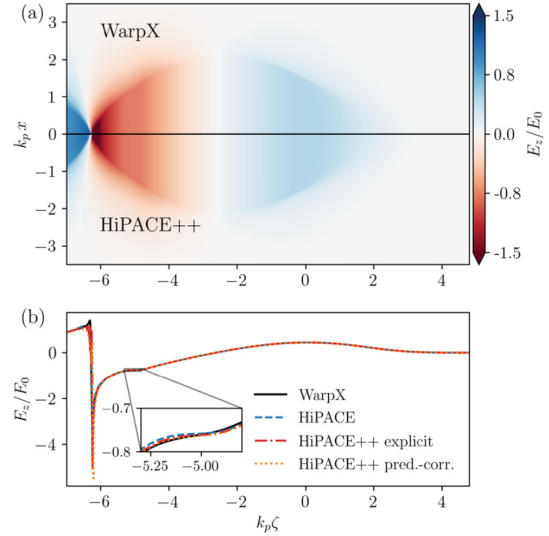


Fig. 2. (a) x - ζ snapshot of the electric field in a beam-driven wakefield acceleration simulation using WarpX (top) and HiPACE++ (bottom). (b) Lineout of the accelerating field from WarpX, HiPACE, and HiPACE++. The inset shows a zoom on the witness beam region, where flattening of the accelerating field due to beam loading is visible. Both the predictor-corrector (pred.-corr.) and the explicit field solvers are shown. (For interpretation of the colors in the figure(s), the reader is referred to the web version of this article.)

shown in Fig. 2(b). For WarpX, the rigid beams propagated in a uniform plasma long enough for the wake to reach a steady state. Minor differences in the witness beam region and in the spike at the back of the bubble can be attributed to different physical models, numerical methods, and initialization.

5. Performance and parallelization

In this section, the performance of HiPACE++ is evaluated. Although QSA PIC codes are considered fast due to their large time steps, full 3D simulations remain computationally expensive. To reduce the runtime, the QSA PIC loop can be parallelized in two independent ways: First, for so-called *transverse parallelization*, the computation of a 2D slice can be performed by multiple processing units via transverse domain decomposition. Second, in the *longitudinal parallelization*, the domain is decomposed longitudinally, and different processing units compute different parts of this domain. Due to intrinsic dependencies of the QSA PIC method (at a given time step, the head of the domain *must* be computed before the tail), this longitudinal parallelization takes the form of a pipeline [43,44], where different ranks compute different time steps.

Transversely, the computation of individual slices in HiPACE++ is performed on 1 GPU (when running on GPU, see Sec. 5.1) or using multiple OpenMP threads (when running on CPU). Longitudinally, the code is parallelized with MPI through a novel pipeline algorithm, see Sec. 5.2.

Unless stated otherwise, all simulations in this section ran on the JUWELS Booster, where each node is equipped with 2 AMD EPYC 7402 processors with 24 cores each and 4 NVIDIA A100 GPUs (40 GB, NVLink3) per node.

5.1. Single-GPU performance

As discussed in Sec. 3 and illustrated in Fig. 1, the amount of data that must be allocated for a 3D domain is relatively modest and consists of beam particles and 2D slices of plasma particles




Fig. 3. Performance comparison between GPU and CPU, for the same setup as Sec. 4 for a single time step with medium ($512 \times 512 \times 1024$ cells, blue lines) and high ($2048 \times 2048 \times 1024$ cells, red lines) resolutions, with predictor-corrector field solver) on the JUWELS Booster. Simulations on CPU used HiPACE (MPI-parallel, dashed lines) and HiPACE++ (OpenMP-parallel, dotted line). Simulations on GPU used HiPACE++. The high resolution run with HiPACE does not fit on less than 16 nodes on CPU.

and fields on the grid. Due to the relatively small size of beam data, the amount of data virtually depends only on the transverse number of cells. For example, the total allocated data on the global memory of a NVIDIA A100 GPU with the explicit solver (respectively predictor-corrector method) with 2 million beam particles (accounting for ~ 230 MB) and 1 particle per cell for the plasma electrons is 2.0 GB (resp. 2.0 GB) for a problem of 128×128 cells transversely, 2.9 GB (resp. 2.6 GB) for a 1024×1024 problem and 19.2 GB (resp. 12.9 GB) for a 4096×4096 problem size (for details on all simulation parameters see the Appendix). Therefore, most practical problems fit on a single GPU, and the performance of HiPACE++ on a single NVIDIA A100 GPU is detailed below.

The benefit of fitting the problem on a single GPU is clearly demonstrated in Fig. 3. Typical CPU implementations of the 3D QSA PIC method [22,39,21] accelerate the calculation by decomposing the domain transversely, resulting in large amounts of communications (in particular in the Poisson solver) that dominate the runtime and cause non-ideal scaling. The CPU runs used only the 48 CPU cores on the nodes of the JUWELS Booster. The GPU runs also used the 4 GPUs. The CPU runs were parallelized in the transverse direction only. Longitudinal parallelization is an orthogonal problem, and is done in the exact same way on CPU and GPU (see Sec. 5.2). On GPU, the simulations at medium and high resolutions take 3.6 sec and 22.9 sec and cost 2.5×10^{-4} node-hours and 1.6×10^{-3} node-hours, respectively. For the same simulations using 1024 cores on CPU, HiPACE requires 17.5 sec and 556.1 sec for a cost of 0.10 node-hours and 3.3 node-hours. At medium resolution, the run on 1 (1024) CPU cores was $145 \times (4.7 \times)$ slower and cost $12 \times (630 \times)$ more node-hours than on 1 GPU. At high resolution, the run on 16 (1024) CPU cores was $197 \times (24 \times)$ slower and cost $261 \times (2050 \times)$ more node-hours than on 1 GPU. The number of node-hours was calculated as $[\text{number of CPU cores}]/48$ for CPU runs, and $[\text{number of GPUs}]/4$ for GPU runs, as each node has 48 CPU cores and 4 GPUs.

For CPU computing with no hardware accelerator, shared-memory parallelization with OpenMP is implemented to enable transverse parallelization when running on CPU only. In that case, tiling is implemented for plasma particle operations (field gather, particle push and current deposition), and the threaded version of FFTW can be called. As shown by the dotted line in Fig. 3, the transverse OpenMP parallelization of HiPACE++ gives a similar scaling as the pure MPI transverse parallelization of the legacy code HiPACE up to 16 threads (running on 16 cores of the 24-core JUWELS Booster CPUs). We attribute the performance improvement of HiPACE++ over HiPACE to better memory handling, but




Fig. 4. Runtime for different transverse resolutions on (a) NVIDIA A100 GPUs and (b) a single Graphics Compute Die (GCD) of an AMD Instinct MI250X. Left bars: using the predictor-corrector loop. Right bars: using the explicit field solver. The runtimes of 1024×1024 for AMD Instinct MI250X only, 2048×2048 and 4096×4096 transverse grid points are plotted on a separate y-axis to improve readability of the figure.

detailed profiling of the legacy code HiPACE is out of scope of this article.

For further insight into the performance of HiPACE++, we ran the reference setup presented in Sec. 4 with increasing transverse resolution, keeping all other parameters constant (for more details see the Appendix). This scan uses 1024 longitudinal grid points, and performance data is given for both the predictor-corrector loop and the explicit field solver. The predictor-corrector loop used up to 5 iterations, which typically yields a comparable level of convergence between the two solvers in standard beam-driven plasma accelerator scenarios. We observed that the explicit solver converges faster than the predictor-corrector loop in challenging simulation settings, such as large transverse box sizes or abrupt beam current spikes.

The most time-consuming functions of the two solvers on an NVIDIA A100 are shown in Fig. 4 (a). In both cases a vast majority of the time is spent in solving for B_x/y . While both the fast Poisson solver and particle operations dominate the predictor-corrector solver at different resolutions, the multigrid solver is always the most expensive operation for the explicit solver. As a reminder, each iteration in the predictor-corrector loop involves all PIC operations for the plasma particles (field gather, particle push, current deposition and field solve) repeated up to 5 times per slice. Note that this study is not a comparison of the two field solvers, as they have different convergence properties, but rather a performance analysis of each solver separately.

The performance portability of HiPACE++ on ROCm-capable AMD GPUs is demonstrated by running the transverse scaling on a single Graphics Compute Die (GCD) of an AMD Instinct MI250X, shown in Fig. 4 (b). The scan was performed on the early-access test system Crusher at the Oak Ridge Leadership Facility, which is equipped with a 64-core AMD EPYC 7A53 “Optimized 3rd Gen EPYC” CPU and four AMD Instinct MI250X. Each MI250X contains two GCDs, which can be viewed as two separate GPUs from a program-

ming perspective. The run time differs roughly by $0.7 - 1.6\times$ (resp. $0.9 - 3.7\times$) for the predictor-corrector solver (resp. explicit solver) in comparison with the NVIDIA A100. Note that these results were obtained on a test platform to demonstrate the portability and significant performance improvements are to be expected, for example by optimizations of rocFFT, which currently suffers from performance penalties for FFTs with grid sizes that are not a power of two.

5.2. Longitudinal parallelization via temporal domain decomposition

As presented in Sec. 2, the computation of the fields in the full domain at a given time step relies on a loop from the head-most slice to the tail, and consequently cannot be parallelized longitudinally by standard domain decomposition. When computing multiple time steps, longitudinal parallelization can be achieved via pipelining algorithms [43,44], which were first realized in the form of a spatial decomposition [43]. Because of the combination of (i) faster computation of each slice and (ii) using a single rank per slice, the standard spatial decomposition demonstrates poor scaling with our GPU implementation. We hereafter present a temporal domain decomposition, more suitable to the GPU implementation in HiPACE++. This implementation has some similarities with the streaming pipeline presented in Ref. [44]. Both pipelines are summarized below, assuming the problem is decomposed longitudinally in as many sub-domains (*boxes*) as the number of ranks n_{ranks} , and runs for nt time steps.

In the spatial decomposition, each rank gets assigned one sub-domain, which it consecutively calculates for every time step. After the rank has calculated its sub-domain for time step t , it passes the plasma particles and slices required for computation via MPI to the next rank downstream. It then receives the plasma particles and slices at $t + \Delta t$ from the next rank upstream. A rank keeps its associated beam particles, unless they slip backwards out of the sub-domain due to longitudinal velocities smaller than the speed of light. The communication caused by beam particle slippage is usually negligible.

The algorithm (in pseudo-code) reads:

```
# Rank r computes box b for all time steps
for t in 0:nt-1:
    Receive last slice from box b+1 at time t
    Compute box b at time t
    Send last slice to box b-1 at time t
```

where a slice consists in field data and plasma particle data. This pipeline is represented on the left of Fig. 5. The number of scalars communicated per time step and per rank reads $N_s = n_x n_y (S_{cell} + n_{ppc} S_{plasma})$ where n_x (n_y) is the number of cells in the transverse direction x (y), S_{cell} is the number of scalars communicated per cell (in HiPACE++, $S_{cell} = 6$ for J_x and J_y of the previous slice, and $B_{x/y}$ of the two previous slices) and S_{plasma} is the number of scalars communicated per plasma particle (in HiPACE++, $S_{plasma} = 35$ due to fifth-order Adams-Bashforth pusher). Here, n_{ppc} is the number of plasma particles per cell. Each rank always communicates a full slice, so the amount of data communicated does not scale with n_{ranks} .

In the temporal decomposition, each rank computes the full domain for the subset of time steps t for which $t \equiv r \pmod{nt}$ (where r is the current rank). At each time step, the assigned rank computes the full domain in a loop over the boxes, from head to tail. After each box is calculated, the beam particles within that box are sent to the next rank downstream, which calculates the next time step for this box. Then, the rank receives the beam particles of the next box from the rank upstream and continues its

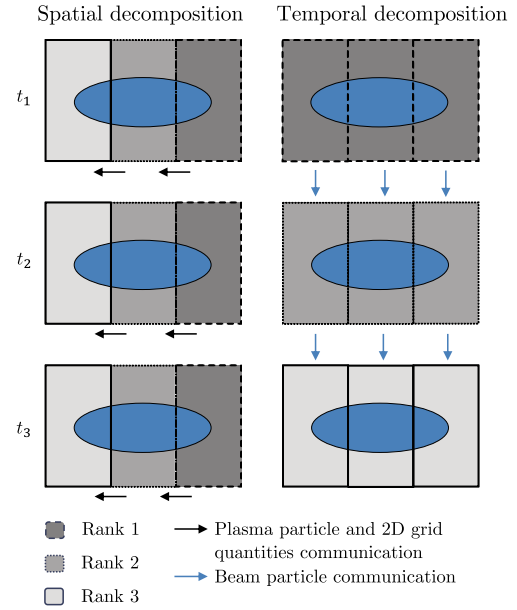


Fig. 5. Left: spatial domain decomposition. Each rank calculates a fixed sub-domain for all time steps. Plasma particles and 2D field slices need to be communicated. Right: temporal domain decomposition. Each rank calculates the full domain for a sub-set of time steps. The beam particles of a sub-domain need to be communicated.

calculation from head to tail. The rank keeps the plasma particles and 2D grid quantities and resets them at each new time step in the first box.

The algorithm (in pseudo-code) reads:

```
# Rank r computes all boxes for time step t
for b in nb-1:0:
    Receive beam from box b at time t-1
    Compute box b at time t
    Send beam to box b at time t+1
```

This pipeline is represented on the right of Fig. 5. The number of scalars communicated per time step and per rank reads $N_t = n_{beam,r} \times S_{beam}$ where $n_{beam,r}$ and S_{beam} denote the number of beam particles on that rank and number of scalars communicated per beam particle, respectively. As can be seen in Fig. 5, $n_{beam,r}$ scales with the number of ranks (i.e., the number of sub-domains), so this pipeline should perform better for strong scaling.

Scalability is the key advantage of the temporal decomposition: in the spatial decomposition, the amount of data to send/receive is constant (one slice of fields and plasma particles) while, in the temporal decomposition, it scales with the number of ranks (only the beam particles within the sub-domain are communicated).

As an example, let us consider a typical problem with $n_x = n_y = 1024$, $n_{ppc} = 1$, and $n_{beam,total} = 2 \times 10^6$. Even in the most favorable case (exchanging as few scalars as possible), $S_{plasma} = 7$ for position, momentum and particle weight (HiPACE++ uses 35), $S_{beam} = 7$ and $S_{cell} = 6$ to calculate the initial guess, the amount of data (assuming IEEE 754 double precision) exchanged per rank and per time step is 110MB for the spatial decomposition. Although this is usually not the case, we assume a load-balanced beam particle distribution across ranks for simplicity, so that $n_{beam,r} = n_{beam,total}/n_{ranks}$. The temporal decomposition exchanges roughly $110\text{MB}/n_{ranks}$ per rank and per time step. For $n_{ranks} = 256$ the temporal decomposition exchanges roughly two




Fig. 6. Strong scaling for two different problems with a) $1024 \times 1024 \times 1024$ cells with 4 plasma particles per cell for 1000 time steps and b) $2048 \times 2048 \times 2048$ cells with 1 plasma particle per cell for 2048 time steps. Both settings used two beams with 10^6 beam particles each. The final run time is given for the maximum number of ranks used. In b) the scaling starts at 4 GPUs due to time limit restriction on the supercomputer. The problems are parallelized in the longitudinal direction only.

orders of magnitude less data than the spatial decomposition, and is hence expected to show better scalability.

The performance of the temporal decomposition pipeline is assessed via a strong scaling of two different setups. One setup is the reference simulation setup from Sec. 4 with $n_{steps} = 1000$ time steps and the other uses $2048 \times 2048 \times 2048$ cells and 2048 time steps (for more details see the Appendix). The efficiency η is given by $\eta(n_{ranks}) = t(1)/[n_{ranks}t(n_{ranks})]$ where $t(i)$ is the run time on i ranks. Due to the filling and emptying of the pipeline, the ideal efficiency for both pipelines is not identically 1 but rather given by

$$\eta_{ideal}(n_{ranks}) = \left(1 + \frac{n_{ranks} - 1}{n_{steps}}\right)^{-1}. \quad (10)$$

An efficiency of 1 is obtained in the limit of $n_{steps} \gg n_{ranks}$. The results are shown in Fig. 6. The temporal domain decomposition (red lines) shows an efficiency close to the ideal pipeline scaling (black dashed line). The spatial decomposition (blue lines) suffers from efficiency degradation above 8 ranks. Both scalings were performed on the JUWELS Booster and the reference setup was also run on Summit (red dotted line), which is equipped with 6 NVIDIA V100 GPUs per node. The maximum number of ranks is chosen so that only 4 slices remain per sub-domain, which was the case at 256 ranks ($= 256$ GPUs) for the reference setup and 512 ranks for the higher-resolution case.

Note, that the temporal domain decomposition outperforms the spatial decomposition even though it is at a disadvantage: due to performance enhancements unrelated to the parallelization, the absolute run time is reduced, causing the communications to take up a larger fraction of the total run time.

6. Software practice and additional features

HiPACE++ is a versatile, open-source, 3D, quasi-static PIC software with an object-oriented design to invite the integration of new numerical methods or physics packages. HiPACE++ uses the cross-platform build system CMake and can be installed, as well as its dependencies, with software package managers, such as Spack [45].

HiPACE++ complies with the openPMD standard [46] and uses the openPMD-api [47] for I/O, allowing for interoperability and simple benchmarking with other codes. Both HDF5 [48] and ADIOS2 [49] file formats are supported (a feature inherited from




Fig. 7. (a) Evolution of the emittance during propagation over 3000 time steps of the witness beam presented in Sec. 4 with an initial transverse offset of the witness bunch centroid of $x_b = \sigma_x$, for the two field solvers, in normalized units. The error due to single precision is computed for each field solver with respect to the double-precision simulation; (b) Same for a simulation running in SI units, where $k_p^{-1} = 10 \mu\text{m}$.

the openPMD-api), and the capability to read an external beam from file at the openPMD format is available.

Two unit systems, SI units and normalized units, are available as a runtime parameter. In normalized units, all lengths are re-scaled to the plasma skin depth k_p^{-1} , the fields to the cold, non-relativistic wave breaking limit E_0 , and all densities to the background plasma density n_0 . All operations are performed in the unit system chosen by the user. An advanced parser makes it possible to write the input file in a unit system and run the simulation in the other one, allowing to use the advantages of both unit systems (numerical accuracy, interoperability with other codes, convenience for multi-physics implementations, etc.) in a flexible manner.

The code can be compiled in either double (C++ double) or single (C++ float) precision, a feature inherited from AMReX. The effect of the precision on the simulation accuracy is investigated by comparing the evolution of the emittance of the witness beam of the reference setup with an initial emittance of $\epsilon_{x,0} = 0$ when an initial transverse offset of the bunch centroid $x_b = \sigma_x$ is present in Fig. 7. For both predictor-corrector and explicit field solvers, the error attributed to using single precision remains well below 1% (2%) in normalized units (SI units) after 3000 time steps. As expected, the difference between single and double precision is higher for SI units than for normalized units, although both remain on the percent level.

Table 1 shows the runtime in single and double precisions for the two solvers on two different architectures, a cutting-edge HPC GPU (NVIDIA A100) and a typical consumer-grade (“gaming”) GPU (NVIDIA RTX2070), easily available on a laptop. As anticipated, double-precision calculations are much faster on the HPC GPU than on the gaming GPU. However, with the capability to run high-resolution production simulations on a gaming GPU with comparable accuracy and performance as on an HPC GPU in single precision, HiPACE++ provides useful scalability from laptops to the largest supercomputers.

Table 1

Runtime for 1 time step of the simulation presented in Fig. 7 on NVIDIA A100 and NVIDIA RTX2070 GPUs. P-C stands for the predictor-corrector loop.

GPU	Solver	T_{double}	T_{single}	speedup
A100	P-C	14.9 s	11.4 s	1.3×
A100	explicit	8.0 s	5.8 s	1.4×
RTX2070	P-C	96.9 s	33.7 s	2.9×
RTX2070	explicit	52.4 s	17.0 s	3.1×

Though in active development, HiPACE++ features numerous capabilities useful for production simulations including multiple beams and plasma with different species and profiles (driver and witness beam, ion motion, etc.), the possibilities to load external beams and apply external fields as well as specialized mesh refinement capabilities [50]. Field ionization using the ADK-model [51] is available in SI units and could readily be extended to normalized units. HiPACE++ proposes different field solvers, and uses modern software practices to make it a user-friendly and stable code (continuous integration, open development repository, extensive documentation and comments). Planned upgrades include a laser envelope model [52,17,53], full mesh refinement, and the support of more GPU architectures from other providers. The code is fully operational on CPU, on modern NVIDIA GPUs, and modern, ROCm-capable AMD GPUs.

7. Conclusion

This paper presented the open-source, performance-portable, 3D quasi-static PIC code HiPACE++. The main adjustments required to port the quasi-static PIC loop to accelerated computing consist in (i) ensuring that all operations, including the beam operations, are performed within the loop over slices so little data needs to be stored on device memory and (ii) proposing a different longitudinal parallelization (pipeline) with which the amount of data communicated per rank scales down with the number of ranks, enabling excellent scalability for production simulations up to hundreds of GPUs.

Focusing on runtime rather than scalability, HiPACE++ is not MPI-parallel transversally: each slice is computed on a single GPU, enabling orders-of-magnitude speedup with respect to CPU implementations. Transverse parallelization will be considered if it provides significant speedup without impacting the code clarity. Benchmarks show excellent agreement with legacy code HiPACE and full electromagnetic PIC code WarpX.

HiPACE++ is built on top of cutting-edge libraries (in particular AMReX and openPMD-api) to harness top performance-portability and encourage open science, while improving sustainability. It enables production simulations of plasma acceleration from laptops to supercomputers.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

The authors gratefully acknowledge helpful discussions with T.J. Mehrling, C.B. Schroeder, J. Osterhoff, T. Wetzell, and B. Diederichs. We gratefully acknowledge the Gauss Centre for Supercomputing e.V. (www.gauss-centre.eu) for funding this project by providing computing time through the John von Neumann Institute for Computing (NIC) on the GCS Supercomputer JUWELS Booster at Jülich Supercomputing Centre (JSC). We acknowledge

Table A.2

Simulation parameters of the presented studies in the respective figure. P-C stands for predictor-corrector solver, n_{ppc} refers to the number of plasma particles per cell, the *beam* column indicates whether the beam was initialized as a random beam with a fixed number of particles and a fixed weight or by a beam with fixed particles per cell and a variable weight, and n_{steps} denotes the number of time steps. The maximum number of iterations in the predictor-corrector solver is given between brackets in the Solver column.

Fig.	Solver	$n_x/y \times n_z$	n_{ppc}	Beam	n_{steps}
2	P-C (5)	1024 × 1024	4	fixed ppc	1
	explicit	1024 × 1024	4	fixed ppc	1
3	P-C (1)	512 × 1024	1	random	1
		2048 × 1024	1	random	1
4	P-C (5)	$2^n \times 1024$	1	random	1
	explicit	$2^n \times 1024$	1	random	1
6	P-C (1)	1024 × 1024	4	random	1000
	P-C (1)	2048 × 2048	1	random	2048
7	P-C (5)	1024 × 1024	1	random	3000

the Funding by the Helmholtz Matter and Technologies Accelerator Research and Development Program. This research was also supported by the Exascale Computing Project (No. 17-SC-20-SC), a collaborative effort of the U.S. Department of Energy's Office of Science and National Nuclear Security Administration. We acknowledge the support of the Director, Office of Science, Office of High Energy Physics, of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231. This research used resources of the Oak Ridge Leadership Computing Facility at the Oak Ridge National Laboratory, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC05-00OR22725.

Appendix A. Simulation parameters

In all simulations, second-order current deposition was used in the transverse direction. The reference setup consists of a drive and a witness beam. Both beams are Gaussian with rms sizes $k_p \sigma_{\perp,d} = 0.3$, $k_p \sigma_{\xi,d} = 1.41$ and $k_p \sigma_{\perp,w} = 0.1$, $k_p \sigma_{\xi,w} = 0.2$. The peak densities are $n_{b,d}/n_0 = 10$ and $n_{b,w}/n_0 = 100$. The drive beam has an initial energy of 10 GeV and 0.1% rms energy spread, the witness beam has an initial energy of 1 GeV and no initial energy spread. The drive beam is located at the origin, the witness beam is centered around longitudinal position $k_p \zeta_{0,w} = -5$. The beams are initialized at waist, either with a fixed number of particles per cell with a variable weight or with random positions and fixed weights. The beams initialized by a variable weight use 1 particle per cell. The randomly initialized beams use 10^6 fixed weight particles per beam. For all simulations, the domain is, in units of k_p^{-1} , $(-8, 8)$, $(-8, 8)$, and $(-7, 5)$ in x , y , and z . The varying simulation parameters for the presented studies are listed in Table A.2. The time step in all simulations is $dt = 6 \omega_p^{-1}$.

For the HiPACE++ and WarpX simulations on the JUWELS Booster, we used GCC 9.3.0, CUDA 11.0, OpenMPI 4.1.0rc1, CMake 3.18.0, and FFTW 3.3.8, except for the strong scaling using the temporal domain decomposition on the JUWELS Booster, which used GCC 10.3.0, CUDA 11.3, and OpenMPI 4.1.1. The GPU runs were compiled with nvcc 11.0.221 using the following flags:

```
-O3 -gencode=arch=compute_80,code=sm_80
-gencode=arch=compute_80,code=compute_80
-maxrregcount=255 --use_fast_math.
```

The HiPACE++ CPU runs were compiled using the following flags: `-O3 -DNDEBUG -pthread -fopenmp -Werror=return-type.`

The legacy code HiPACE was compiled with ICC 19.1.2.254 using the flags: `-std=c99 -march=native -O3 -Os.`

On the laptop, we used GCC 8.4.0, CUDA 11.0, MPI 3.1, and CMake 3.20.3. The code was compiled with nvcc 11.0.194 using the following flags: `-O3 -DNDEBUG --xpt-relaxed-constexpr`

```

--expt-extended-lambda -Xcudafe
--diag_suppress=esa_on_defaulted_function_ignored
--maxrregcount=255 -Xcudafe --display_error_number
--Wext-lambda-captures-this --use_fast_math
-Xcompiler -pthread.

```

Throughout the studies, we used AMReX v21.05 to v22.04 and HiPACE++ from commit 3f2f4e15a607 to v22.04, except for the simulation using spatial domain decomposition, which was conducted on commit 11523c24c0f7c73ce3fe8d3424ede54565f58d50.

References

- [1] T. Tajima, J.M. Dawson, *Phys. Rev. Lett.* 43 (1979) 267–270, <https://doi.org/10.1103/PhysRevLett.43.267>.
- [2] P. Chen, J.M. Dawson, R.W. Huff, T. Katsouleas, *Phys. Rev. Lett.* 54 (1985) 693–696, <https://doi.org/10.1103/PhysRevLett.54.693>.
- [3] C.A. Lindström, J.M. Garland, S. Schröder, L. Boulton, G. Boyle, J. Chappell, R. D'Arcy, P. Gonzalez, A. Knetsch, V. Libov, G. Loisch, A. Martinez de la Ossa, P. Niknejadi, K. Pöder, L. Schaper, B. Schmidt, B. Sheeran, S. Wesch, J. Wood, J. Osterhoff, *Phys. Rev. Lett.* 126 (2021) 014801, <https://doi.org/10.1103/PhysRevLett.126.014801>.
- [4] J. Couperus, R. Pausch, A. Köhler, O. Zarini, J. Krämer, M. Garten, A. Huebl, R. Gebhardt, U. Helbig, S. Bock, et al., *Nat. Commun.* 8 (1) (2017) 1–7, <https://doi.org/10.1038/s41467-017-00592-7>.
- [5] M. Kirchen, S. Jalas, P. Messner, P. Winkler, T. Eichner, L. Hübner, T. Hülsenbusch, L. Jeppé, T. Parikh, M. Schnepf, A.R. Maier, *Phys. Rev. Lett.* 126 (2021) 174801, <https://doi.org/10.1103/PhysRevLett.126.174801>.
- [6] W. Wang, K. Feng, L. Ke, C. Yu, Y. Xu, R. Qi, Y. Chen, Z. Qin, Z. Zhang, M. Fang, et al., *Nature* 595 (7868) (2021) 516–520, <https://doi.org/10.1038/s41586-021-03678-x>.
- [7] R. Hockney, J. Eastwood, *Computer Simulation Using Particles*, Advanced Book Program: Addison-Wesley, McGraw-Hill, 1981, <https://doi.org/10.1201/9780367806934>.
- [8] C. Birdsall, A. Langdon, *Plasma Physics via Computer Simulation*, The Adam Hilger Series on Plasma Physics, McGraw-Hill, 1985, <https://doi.org/10.1201/9781315275048>.
- [9] R. Zgadzaj, T. Silva, V. Khudiyakov, A. Sosedkin, J. Allen, S. Gessner, Z. Li, M. Litos, J. Vieira, K. Lotov, et al., *Nat. Commun.* 11 (1) (2020) 1–11, <https://doi.org/10.1038/s41467-020-18490-w>.
- [10] A. Caldwell, K. Lotov, A. Pukhov, F. Simon, *Nat. Phys.* 5 (5) (2009) 363–367, <https://doi.org/10.1038/nphys1248>.
- [11] S. Schröder, C. Lindström, S. Bohlen, G. Boyle, R. D'Arcy, S. Diederichs, M. Garland, P. Gonzalez, A. Knetsch, V. Libov, et al., *Nat. Commun.* 11 (1) (2020) 1–6, <https://doi.org/10.1038/s41467-020-19811-9>.
- [12] R. Courant, K. Friedrichs, H. Lewy, *Math. Ann.* 100 (1) (1928) 32–74, <https://doi.org/10.1007/BF01448839>.
- [13] J.-L. Vay, *Phys. Rev. Lett.* 98 (2007) 130405, <https://doi.org/10.1103/PhysRevLett.98.130405>.
- [14] P. Sprangle, E. Esarey, A. Ting, *Phys. Rev. A* 41 (1990) 4463–4469, <https://doi.org/10.1103/PhysRevA.41.4463>.
- [15] P. Sprangle, E. Esarey, A. Ting, *Phys. Rev. Lett.* 64 (1990) 2011–2014, <https://doi.org/10.1103/PhysRevLett.64.2011>.
- [16] P. Mora, T.M. Antonsen Jr., *Phys. Rev. E* 53 (1996) R2068–R2071, <https://doi.org/10.1103/PhysRevE.53.R2068>.
- [17] P. Mora, T.M. Antonsen Jr., *Phys. Plasmas* 4 (1) (1997) 217–229, <https://doi.org/10.1063/1.872134>.
- [18] K.V. Lotov, *Phys. Plasmas* 5 (3) (1998) 785–791, <https://doi.org/10.1063/1.872765>.
- [19] K.V. Lotov, *Phys. Rev. Spec. Top., Accel. Beams* 6 (2003) 061301, <https://doi.org/10.1103/PhysRevSTAB.6.061301>.
- [20] C. Huang, V. Decyk, C. Ren, M. Zhou, W. Lu, W. Mori, J. Cooley, T.M. Antonsen Jr., T. Katsouleas, *J. Comput. Phys.* 217 (2) (2006) 658–679, <https://doi.org/10.1016/j.jcp.2006.01.039>.
- [21] W. An, V.K. Decyk, W.B. Mori, T.M. Antonsen Jr., *J. Comput. Phys.* 250 (2013) 165–177, <https://doi.org/10.1016/j.jcp.2013.05.020>.
- [22] T.J. Mehrling, C. Benedetti, C.B. Schroeder, J. Osterhoff, *Plasma Phys. Control. Fusion* 56 (8) (2014) 084012, <https://doi.org/10.1088/0741-3335/56/8/084012>.
- [23] C. Benedetti, C.B. Schroeder, E. Esarey, C.G.R. Geddes, W.P. Leemans, *AIP Conf. Proc.* 1812 (1) (2017) 050005, <https://doi.org/10.1063/1.4975866>.
- [24] A. Pukhov, *CERN Yellow Rep.* 1 (2016) 181, <https://doi.org/10.5170/CERN-2016-001.181>.
- [25] J.-L. Vay, A. Huebl, A. Almgren, L. Amorim, J. Bell, L. Fedeli, L. Ge, K. Gott, D. Grote, M. Hogan, et al., *Phys. Plasmas* 28 (2) (2021) 023105, <https://doi.org/10.1063/5.0028512>.
- [26] P. Yu, X. Xu, V.K. Decyk, W. An, J. Vieira, F.S. Tsung, R.A. Fonseca, W. Lu, L.O. Silva, W.B. Mori, *J. Comput. Phys.* 266 (2014) 124–138, <https://doi.org/10.1016/j.jcp.2014.02.016>.
- [27] H. Meuer, E. Strohmaier, J. Dongarra, H. Simon, Top 500 supercomputers, <http://www.top500.org/lists/2021/11/>, 2021.
- [28] H. Bureau, R. Widera, W. Höning, G. Juckeland, A. Debus, T. Kluge, U. Schramm, T.E. Cowan, R. Sauerbrey, M. Bussmann, *IEEE Trans. Plasma Sci.* 38 (10) (2010) 2831–2839, <https://doi.org/10.1109/TPS.2010.2064310>.
- [29] M. Bussmann, H. Bureau, T.E. Cowan, A. Debus, A. Huebl, G. Juckeland, T. Kluge, W.E. Nagel, R. Pausch, F. Schmitt, U. Schramm, J. Schuchart, R. Widera, in: *Proceedings of the International Conference on High Performance Computing, Networking, Storage and Analysis*, SC '13, ACM, New York, NY, USA, 2013, 5, <https://doi.org/10.1145/2503210.2504564>.
- [30] A. Myers, A. Almgren, L. Amorim, J. Bell, L. Fedeli, L. Ge, K. Gott, D. Grote, M. Hogan, A. Huebl, R. Jambunathan, R. Lehe, C. Ng, M. Rowan, O. Shapoval, M. Thévenet, J.-L. Vay, H. Vincenti, E. Yang, N. Zaïm, W. Zhang, Y. Zhao, E. Zoni, *Parallel Comput.* 108 (2021) 102833, <https://doi.org/10.1016/j.parco.2021.102833>.
- [31] R. Bird, N. Tan, S.V. Luedtke, S. Harrell, M. Tauber, B. Albright, *IEEE Trans. Parallel Distrib. Syst.* 33 (04) (2022) 952–963, <https://doi.org/10.1109/TPDS.2021.3084795>.
- [32] H.C. Edwards, C.R. Trott, D. Sunderland, in: *Domain-Specific Languages and High-Level Frameworks for High-Performance Computing*, J. Parallel Distrib. Comput. 74 (12) (2014) 3202–3216, <https://doi.org/10.1016/j.jpdc.2014.07.003>.
- [33] D.A. Beckingsale, J. Burmark, R. Hornung, H. Jones, W. Killian, A.J. Kunen, O. Pearce, P. Robinson, B.S. Ryujiin, T.R. Scogland, in: *2019 IEEE/ACM International Workshop on Performance, Portability and Productivity in HPC (P3HPC)*, 2019, pp. 71–81, <https://doi.org/10.1109/P3HPC49587.2019.00012>.
- [34] E. Zenker, B. Worpitz, R. Widera, A. Huebl, G. Juckeland, A. Knüpfner, W.E. Nagel, M. Bussmann, in: *2016 IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW)*, IEEE, 2016, pp. 631–640, <https://doi.org/10.1109/IPDPSW.2016.50>.
- [35] M. Thévenet, S. Diederichs, A. Huebl, A. Myers, A. Sinn, R. Lehe, J.-L. Vay, W. Zhang, *Hi-pace/hipace*: v21.09, <https://doi.org/10.5281/zenodo.5358484>, Sep. 2021.
- [36] W. Zhang, A. Almgren, V. Beckner, J. Bell, J. Blaschke, C. Chan, M. Day, B. Friesen, K. Gott, D. Graves, M. Katz, A. Myers, T. Nguyen, A. Nonaka, M. Rosso, S. Williams, M. Zingale, *J. Open Sour. Softw.* 4 (37) (2019) 1370, <https://doi.org/10.21105/joss.01370>.
- [37] E. Esarey, C.B. Schroeder, W.P. Leemans, *Rev. Mod. Phys.* 81 (2009) 1229–1285, <https://doi.org/10.1103/RevModPhys.81.1229>.
- [38] T. Wang, V. Khudik, B. Breizman, G. Shvets, *Phys. Plasmas* 24 (10) (2017) 103117, <https://doi.org/10.1063/1.4999629>.
- [39] T. Wang, V. Khudik, G. Shvets, *WAND-PIC: a three-dimensional quasi-static particle-in-cell code with parallel multigrid solver and without predictor-corrector*, <https://doi.org/10.48550/arXiv.2012.00881>, 2020.
- [40] C. Van Loan, *Computational Frameworks for the Fast Fourier Transform*, *Frontiers in Applied Mathematics, Society for Industrial and Applied Mathematics*, 1992, <https://doi.org/10.1137/1.9781611970999>.
- [41] M. Frigo, S. Johnson, *Proc. IEEE* 93 (2) (2005) 216–231, <https://doi.org/10.1109/JPROC.2004.840301>.
- [42] J.W. Cooley, P. Lewis, P. Welch, *J. Sound Vib.* 12 (3) (1970) 315–337, [https://doi.org/10.1016/0022-460X\(70\)90075-1](https://doi.org/10.1016/0022-460X(70)90075-1).
- [43] B. Feng, C. Huang, V. Decyk, W. Mori, P. Muggii, T. Katsouleas, *J. Comput. Phys.* 228 (15) (2009) 5340–5348, <https://doi.org/10.1016/j.jcp.2009.04.019>.
- [44] A. Sosedkin, K.V. Lotov, *Nucl. Instrum. Methods Phys. Res., Sect. A, Accel. Spectrom. Detect. Assoc. Equip.* 829 (2016) 350–352, <https://doi.org/10.1016/j.nima.2015.12.032>.
- [45] T. Gambelin, M. LeGendre, M.R. Collette, G.L. Lee, A. Moody, B.R. de Supinski, S. Futral, in: *SC'15: Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, IEEE, 2015, pp. 1–12, <https://doi.org/10.1145/2807591.2807623>.
- [46] A. Huebl, R. Lehe, J.-L. Vay, D.P. Grote, I.F. Sbalzarini, S. Kuschel, D. Sagan, C. Mayes, F. Perez, F. Koller, M. Bussmann, *openPMD: a meta data standard for particle and mesh based data*, <https://doi.org/10.5281/zenodo.591699>, 2015.
- [47] A. Huebl, F. Poeschel, F. Koller, J. Gu, *openPMD-api: C++ & Python API for Scientific I/O with openPMD*, <https://doi.org/10.14278/rodare.27>, 2018.
- [48] The HDF Group, *Hierarchical data format version 5*, <http://www.hdfgroup.org>, 2000–2019.
- [49] W.F. Godoy, N. Podhorski, R. Wang, C. Atkins, G. Eisenhauer, J. Gu, P. Davis, J. Choi, K. Geraschewski, K. Huck, et al., *SoftwareX* 12 (2020) 100561, <https://doi.org/10.1016/j.softx.2020.100561>.
- [50] T.J. Mehrling, C. Benedetti, C.B. Schroeder, E. Esarey, in: *2018 IEEE Advanced Accelerator Concepts Workshop (AAC)*, IEEE, 2018, <https://doi.org/10.1109/AAC.2018.8659404>.
- [51] M.V. Ammosov, N.B. Delone, V.P. Krainov, in: *J.A. Alcock (Ed.), High Intensity Laser Processes*, in: *International Society for Optics and Photonics*, vol. 0664, SPIE, 1986, pp. 138–141, <https://doi.org/10.1117/12.938695>.
- [52] P. Sprangle, E. Esarey, J. Krall, G. Joyce, *Phys. Rev. Lett.* 69 (15) (1992) 2200, <https://doi.org/10.1103/PhysRevLett.69.2200>.
- [53] C. Benedetti, C.B. Schroeder, C.G.R. Geddes, E. Esarey, W.P. Leemans, *Plasma Phys. Control. Fusion* 60 (1) (2017) 014002, <https://doi.org/10.1088/1361-6587/aa8977>.

High-quality positron acceleration in beam-driven plasma acceleratorsS. Diederichs^{1,2,3}, C. Benedetti², E. Esarey², J. Osterhoff³, and C. B. Schroeder^{2,4}¹*University of Hamburg, Institute of Experimental Physics, D-22761 Hamburg, Germany*²*Lawrence Berkeley National Laboratory, Berkeley, California 94720, USA*³*Deutsches Elektronen-Synchrotron DESY, D-22607 Hamburg, Germany*⁴*Department of Nuclear Engineering, University of California, Berkeley, California 94720, USA*

(Received 27 August 2020; accepted 18 November 2020; published 3 December 2020)

Acceleration of positron beams in plasma-based accelerators is a highly challenging task. To realize a plasma-based linear collider, acceleration of a positron bunch with high-efficiency is required, while maintaining both a low emittance and a subpercent-level energy spread. Recently, a plasma-based positron acceleration scheme was proposed in which a wake suitable for the acceleration and transport of positrons is produced in a plasma column by means of an electron drive beam [Diederichs *et al.*, *Phys. Rev. Accel. Beams* **22**, 081301 (2019)]. In this article, we present a study of beam loading for a positron beam in this type of wake. We demonstrate via particle-in-cell simulations that acceleration of high-quality positron beams is possible, and we discuss a possible path to achieve collider-relevant parameters.

DOI: [10.1103/PhysRevAccelBeams.23.121301](https://doi.org/10.1103/PhysRevAccelBeams.23.121301)**I. INTRODUCTION**

Plasma-based particle accelerators potentially enable compact linear electron-positron colliders due to their large acceleration gradients [1]. In a plasma wakefield accelerator (PWFA), an ultrarelativistic, high-charge density particle beam expels all plasma electrons from its propagation axis and an ion cavity is formed [2,3]. The cavity, also referred to as bubble or blowout, features a region with a large, longitudinally accelerating gradient and a transversely linear restoring force for relativistic electrons. Whereas high-energy gain, high-efficiency [4,5], and stably beam-loaded [6] electron acceleration has been demonstrated experimentally in PWFAs, stable and quality preserving positron acceleration remains a challenge. Identifying a positron acceleration scheme that fulfills the requirements imposed by a particle collider, namely the stable and efficient acceleration of high-charge positron bunches, while maintaining both a low emittance and a low energy spread, has been an outstanding challenge, and previously proposed positron acceleration concepts were not able to meet all the necessary requirements. For instance, utilizing hollow core electron drive beams showed only a per-mille-level driver-to-witness energy conversion efficiency [7]. PWFAs driven by a positron beam have been investigated in Ref. [8]. While this scheme demonstrated high-efficiency acceleration of the positron witness beam,

the nonlinear nature of the transverse focusing fields, and their variation as the drive beam evolves renders the preservation of the witness beam emittance challenging. Hollow core plasma channels have been proposed as potential plasma target candidates for positron acceleration [9,10]. However, owing to the lack of any focusing field for the beam in a hollow channel, this scheme suffers from severe beam breakup instability [9,11].

In a recent article, a novel method for positron acceleration was proposed that uses an electron beam as driver and a plasma column as the acceleration medium [12]. For a plasma column with a column radius smaller than the blowout radius, the transverse wakefields are altered, resulting in an elongation of the background plasma trajectories returning toward the axis. This creates a long, high-density electron filament, leading to the formation of a wake phase region which is suitable for acceleration and transport of positron beams. Despite the nonlinear nature of the transverse wakefields, it was shown that quasimatched propagation of positron beams was possible. Due to the non-uniformity of the accelerating field created in these structures, the energy-spread was found to be at the percent-level, which is too high for application in a plasma-based linear collider. Another study has investigated beam loading of simple Gaussian beams in these plasma structures [13]. Despite achieving higher efficiency than the one reported in Ref. [12], the emittance was not preserved.

In this article, we investigate beam loading of a positron bunch in the nonlinear wake formed in a plasma column with the goal of minimizing the energy spread of the bunch, while maintaining both a low emittance and a high charge. Beam loading has been first described for linear wakes in

Published by the American Physical Society under the terms of the Creative Commons Attribution 4.0 International license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

Ref. [14]. In the nonlinear blowout regime, beam loading of electron beams was studied in Ref. [15], where an analytical expression for the longitudinal witness beam current profile that eliminates the energy spread was obtained. Owing to the different nature of the wakefield structure, this type of analytic result is not valid in the case of the nonlinear positron accelerating fields considered in this study. Here, beam loading is studied by means of a numerical algorithm that reconstructs, slice-by-slice and self-consistently, the longitudinal current profile of an optimal witness beam which flattens the accelerating fields within the bunch. We further discuss the transport of the positron witness bunch and its optimization with the goal of minimizing the energy spread and preserving the emittance, both crucial parameters for the employment of this acceleration scheme in a future plasma-based linear collider.

Lastly, we assess a possible path to achieve collider-relevant parameters.

II. NONLINEAR WAKEFIELDS FOR POSITRON ACCELERATION

The generation of positron beam focusing and accelerating wakes using plasma columns was first described in [12]. Using an electron drive beam and a plasma column with a radius smaller than the blowout radius leads to the formation of a wide longitudinal electron filament behind the blowout bubble. This elongated region of high electron density provides accelerating and focusing fields for positron beams. This is illustrated in Figs. 1 and 2, which show two-dimensional maps of the accelerating field E_z/E_0 and focusing field $(E_x - B_y)/E_0$, respectively. The fields are normalized to the cold, nonrelativistic wave-breaking




FIG. 1. Two-dimensional (ζ, x) map of the accelerating wakefield, E_z/E_0 . Positrons can be accelerated in the region $-14 \lesssim k_p \zeta \lesssim -10$. Inset: transverse dependence of accelerating field in the positron accelerating region at three different longitudinal locations denoted by the dashed ($k_p \zeta = -12.5$), solid ($k_p \zeta = -11.5$), and dotted ($k_p \zeta = -10.5$) lines. The accelerating field falls off for increasing distance from the propagation axis. The gradient of the transverse field decreases further behind the driver.




FIG. 2. Two-dimensional (ζ, x) map of the focusing wakefield, $(E_x - B_y)/E_0$. Positrons can be focused in the region $-14 \lesssim k_p \zeta \lesssim -9$, which widely overlaps with the positron accelerating region. Inset: transverse dependence of focusing field at three different longitudinal locations denoted by the dashed ($k_p \zeta = -12.5$), solid ($k_p \zeta = -11.5$), and dotted ($k_p \zeta = -10.5$) lines. The focusing field decays almost linearly for increasing distances from the propagation axis. The field decreases further behind the driver.

limit $E_0 = \omega_p mc/e$, where c denotes the speed of light in vacuum, $\omega_p = (4\pi n_0 e^2/m)^{1/2}$ the plasma frequency, n_0 the background plasma density, and e and m the electron charge and mass, respectively. In this example, we consider a plasma column with a radius $k_p R_p = 2.5$ and a Gaussian electron drive beam with sizes $k_p \sigma_{x,y}^{(d)} = 0.1$, $k_p \sigma_z^{(d)} = \sqrt{2}$, and peak current $I_b^{(d)}/I_A = 1$, where $I_A = mc^3/e \simeq 17$ kA is the Alfvén current. The modeling was performed using the quasistatic particle-in-cell (PIC) code HiPACE [16]. To reduce the high computational cost of the modeling imposed by the required numerical resolution, the wakefields were computed using an axisymmetric cylindrical solver based on the one implemented in the quasistatic version of the code INF&RNO [17], while the particles are advanced in full 3D. Denoting by $k_p = \omega_p/c$ the plasma wave number, the dimensions of the computational domain are $12 \times 12 \times 20 k_p^{-3}$ in the coordinates $x \times y \times \zeta$, where x , and y are the transverse coordinates, and $\zeta = z - ct$ is the longitudinal co-moving coordinate, with z and t being the longitudinal coordinate and the time, respectively. The resolution is $0.0056 \times 0.0056 \times 0.0075 k_p^{-3}$. The background electron plasma was modeled with 25 constant weight particles per cell. The drive beam was sampled with 10^6 constant-weight particles.

The positron focusing and accelerating phase is located between $-14 \lesssim k_p \zeta \lesssim -9$. The accelerating field has its peak at $k_p \zeta \approx -11.5$. Unlike in the blowout regime case, E_z has a transverse dependence. The inset of Fig. 1 shows E_z/E_0 along the transverse coordinate x at three different longitudinal locations denoted by the dashed ($k_p \zeta = -12.5$), solid ($k_p \zeta = -11.5$), and dotted ($k_p \zeta = -10.5$) lines in Fig. 1.

In all three locations, $E_z(x)$ has an on-axis maximum and decays for increasing distances from the propagation axis. Notably, the transverse gradient of the accelerating field is smaller further behind the driver. The nonuniformity of E_z will lead to a ζ -dependent uncorrelated slice energy spread since particles that remain closer to the axis will experience a larger accelerating gradient compared to the ones further off axis. This effect will be investigated more thoroughly in Sec. III C.

The transverse behavior of the focusing field, $(E_x - B_y)/E_0$, is depicted in the inset of Fig. 2, where we show transverse lineouts of the focusing wakefields for the same three longitudinal locations used in Fig. 1. We see that the transverse wakefield decays almost linearly for increasing distances from the propagation axis. The field decrease is smaller further behind the driver. As shown in Ref. [12], the field becomes almost a step-function when sufficiently loaded by a positron bunch.

III. SELF-CONSISTENT BEAM LOADING TO MINIMIZE THE ENERGY-SPREAD

In many beam-driven plasma wakefield accelerator applications, both the driver and the witness beams are usually highly relativistic and evolve on a much longer time scale than the background plasma. In this case the quasistatic approximation [18], which allows treatment of the plasma and the relativistic beams in a separate manner, can be used. In the quasistatic approximation, the wakefields generated by a given beam are determined by initializing a slice of unperturbed plasma ahead of the beam and then follow its evolution as the slice is pushed through the beam from head to tail along the negative ζ direction (here ζ can be interpreted as a fast “time” that parametrizes plasma-related quantities), while the beam is assumed to be frozen. This implies that to calculate the fields at some longitudinal position ζ , only the information upstream of this point is required.

We used this feature of the quasistatic solution to design an algorithm that recursively constructs, slice-by-slice and starting from the head, the optimal current profile of a witness bunch such that the accelerating field along the bunch is constant and equal to a set value. This leads to a reduced energy spread of the accelerated particles. The algorithm is described in detail in the Appendix. We considered a (radially symmetric) bunch initially described as

$$n_b(\zeta, r) = g_{\parallel}(\zeta)g_{\perp}(\zeta, r), \quad (1)$$

where $g_{\parallel}(\zeta)$ and $g_{\perp}(\zeta, r)$ denote the longitudinal and transverse density profiles, respectively. We require that, for any ζ , $\int g_{\parallel}(\zeta, r)dr = \int g_{\perp}(\zeta = \zeta_{\text{head}}, r)dr$, where ζ_{head} is the location of the bunch head, so that the bunch current density profile only depends on $g_{\parallel}(\zeta)$. For simplicity, we first consider bunches with transverse profiles that are radially

Gaussian and longitudinally uniform, i.e., $g_{\perp}(\zeta, r) = \exp[-r^2/(2\sigma_r^2)]$, where σ_r is the (longitudinally constant) rms bunch size. At every longitudinal location ζ (bunch slice), the algorithm performs an iterative search for the optimal bunch current, determined via $g_{\parallel}(\zeta)$, that flattens the accelerating field in that particular slice. The procedure is repeated recursively for all the slices going from the head to the tail of the bunch. Note that, besides a constant accelerating field along the bunch, other field configurations yielding an energy chirp during acceleration are possible. In order for a solution to be found, the positron bunch has to be located in a phase of the wake where $\partial_{\zeta}E_z < 0$. To take into account the fact that, in general, E_z varies in the transverse plane across the beam, the figure of merit considered by the algorithm is a transversally weighted accelerating field $\langle E_z \rangle$, defined as

$$\begin{aligned} \langle E_z \rangle &= \frac{\int_0^{\infty} E_z(r)g_{\perp}(r)dr}{\int_0^{\infty} g_{\perp}(r)dr} \\ &= \frac{\int_0^{\infty} E_z(r) \exp[-r^2/(2\sigma_r^2)]dr}{\int_0^{\infty} \exp[-r^2/(2\sigma_r^2)]dr}. \end{aligned} \quad (2)$$

In case of a transversally uniform accelerating field, e.g., as in the blowout regime, the averaged accelerating field simply reduces to the on-axis accelerating field.

A. Optimization of the witness bunch position

The choice of the location of the witness bunch head, ζ_{head} , sets the amplitude of the accelerating gradient and determines the shape of bunch current profile. In the following, we study the effect of different witness head positions for a bunch in the wake described in Sec. II. To fulfill the requirement that the bunch head has to be located in a wake phase such that $\partial_{\zeta}E_z < 0$, and to achieve a reasonable acceleration gradient, we chose $-11.5 \lesssim k_p\zeta_{\text{head}} \lesssim -10$. Also, we consider a witness bunch an emittance such that $k_p\epsilon_x = 0.05$, and a bunch size $k_p\sigma_r = 0.0163$. Numerical results for the current profiles and their corresponding loaded averaged accelerating fields, $\langle E_z \rangle$, for four values of the witness bunch head position are depicted in Fig. 3. Interestingly, placing the bunch head in a more forward position in the wake, corresponding to a lower accelerating gradient, does not necessarily increase the charge of the witness bunch. This can be seen in Table I, where we show the witness charge, Q_w , as a function of the bunch head position. Values of the charge have been computed assuming a background density of $n_0 = 5 \times 10^{17} \text{ cm}^{-3}$. For this density the charge of the drive beam is $Q_d = 1.5 \text{ nC}$.

The driver-to-beam efficiency, η , can be calculated from the charge of the witness beam, its energy gain rate, E_w^+ , the charge of the drive beam, and its energy loss rate, E_d^- , via




FIG. 3. Current profiles for an optimally loaded wake I_b/I_A (red), and corresponding lineouts of the transversally averaged accelerating field, $\langle E_z \rangle / E_0$ (blue). The current profiles and their corresponding accelerating gradient is given for four values of the witness bunch head position, $k_p \zeta_{\text{head}} = -11.0$ (dotted dashed), -10.8 (solid), -10.6 (dashed), -10.4 (dotted). Their relative averaged accelerating gradients within the beam are (in the same order as the starting position) $0.462, 0.438, 0.403, 0.367$ with respect to the maximum averaged accelerating gradient of the unloaded wake $\langle E_{\text{max}} \rangle / E_0 = 0.49$. The same line style marks the current profile and its corresponding field lineout. The presented current profiles optimally load the wake, resulting in a flattened $\langle E_z \rangle / E_0$.

$$\eta = \frac{Q_w E_w^+}{Q_d E_d^-}. \quad (3)$$

For the chosen density the driver energy loss rate is $E_d^- = 34$ GeV/m. Values of the energy gain for the witness bunch and the efficiency as a function of the witness head position are given in Table I.

The results show that the efficiency peaks around the position $-10.8 \lesssim k_p \zeta_{\text{head}} \lesssim -10.6$. As shown in Sec. II, the accelerating field is transversely flatter for more negative head positions, therefore the case $k_p \zeta_{\text{head}} = -10.8$ is preferable since the choice of this witness position will result in a smaller energy-spread, while maintaining close to maximum efficiency. We recall that for this witness position the charge of the bunch is 52 pC and the efficiency $\eta \approx 3\%$. This is less than what was achieved with a simple Gaussian density profile in [12], which featured a witness bunch charge of $Q_w = 84$ pC and an efficiency of $\eta \approx 4.8\%$.

TABLE I. Charge and energy gain of the witness bunch, and driver-to-witness efficiency as a function of the witness head position. Values are computed assuming a background plasma density $n_0 = 5 \times 10^{17}$ cm $^{-3}$. The driver parameters are the same as in Sec. II, yielding $Q_d = 1.5$ nC and $E_d^- = 34$ GeV/m.

$k_p \zeta_{\text{head}}$	Q_w [pC]	E_w^+ [GeV/m]	η [%]
-10.4	54	25.0	2.7
-10.6	57	27.4	3.1
-10.8	52	29.8	3.0
-11.0	36	31.4	2.2

However, energy-spread minimization was not taken into consideration in that study, which lead to an energy-spread on the few-percent-level.

Choosing bunch head positions that are closer to the driver, i.e., $k_p \zeta_{\text{head}} \geq -10.2$, yields complex (e.g., multi-peaked) bunch current profiles. In this case, the positron beam significantly alters the background plasma electron trajectories, resulting in the formation of a second on-axis electron density peak behind the blowout region. This, in principle, allows for the loading of a second positron beam or an increase of the length of the first. This can be seen in Fig. 3. In fact, for $k_p \zeta_{\text{head}} = -10.4$ (dotted line) we see that $\langle E_z \rangle$ has a local maximum behind the bunch which is higher than the value within the bunch, allowing for further beam loading. We did not investigate further such forward starting positions because we consider the resulting complex bunch structures difficult to realize experimentally.

B. Minimizing the correlated energy-spread

Using the weighted accelerating field $\langle E_z \rangle$ from Eq. (2) as the figure of merit in the proposed algorithm yields a bunch current profile that eliminates the correlated energy-spread only under the assumption that the bunch size does not change during acceleration. However, this assumption is generally not true. First, if the spot size is not matched to the focusing field at some position along the bunch due to, e.g., the slice-dependent nature of the transverse wakefields, it will evolve until it is matched. Second, due to the acceleration, the matched spot size adiabatically decreases with increased particle energy. Both effects must be taken into account in order to eliminate the correlated energy spread entirely. Eliminating the mismatch requires performing a slice-by-slice matching of the beam, i.e., introducing a slice-dependent bunch size, $\sigma_r(\zeta)$. Note that this also leads to a ζ -dependence of g_{\perp} . In our algorithm, calculation of the self-consistent, slice-dependent bunch size can be done numerically while the optimal bunch is generated. As a desirable side effect, the slice-by-slice matching also minimizes the emittance growth [19]. To take into account the change of $\sigma_r(\zeta)$ due to acceleration, the averaged spot size over the acceleration distance should be used in calculating $\langle E_z \rangle$ in Eq. (2). We recall that for the here considered steplike wakes with a field strength of α the matching condition for a given emittance ϵ_x is $\sigma_{r,\text{matched}} \approx 1.72 \epsilon_x^2 / (k_p \alpha \gamma)$ and so the matched spot size is expected to scale with the energy as $\sigma_{r,\text{matched}} \propto \sqrt[3]{1/\gamma}$, where γ is the bunch relativistic factor [12].

The averaged bunch size over the acceleration distance can then be estimated as

$$\begin{aligned} \bar{\sigma}_r(\zeta) &= \frac{\sigma_r(\zeta)}{\gamma_{\text{final}} - \gamma_{\text{init}}} \int_{\gamma_{\text{init}}}^{\gamma_{\text{final}}} \sqrt[3]{\frac{\gamma_{\text{init}}}{\gamma}} d\gamma \\ &= \frac{3}{2} \sigma_r(\zeta) \gamma_{\text{init}}^{1/3} \frac{\gamma_{\text{final}}^{2/3} - \gamma_{\text{init}}^{2/3}}{\gamma_{\text{final}} - \gamma_{\text{init}}}, \end{aligned} \quad (4)$$

where γ_{init} and γ_{final} refer to the initial and final bunch energy, respectively. Note that to calculate $\bar{\sigma}_r(\zeta)$, the final beam energy γ_{final} is required. We also notice that the inclusion of the slice-by-slice matching and of the energy-averaged bunch size when computing the optimal bunch profiles do not significantly alter the current profiles and charges discussed in Sec. III A. Changes to the optimal beam loading algorithm including slice-by-slice matching and averaged spot size are described in the Appendix.

The efficacy of the slice-by-slice matching and inclusion of the average spot size is demonstrated in Fig. 4, where we show the mean energy of each slice for a positron witness bunch that accelerates from 1 GeV to ≈ 5.5 GeV in a distance of 15 cm. The blue line refers to algorithm flattening $\langle E_z \rangle$ with a longitudinal uniform σ_r . The red line and the green line refer to additionally applying slice-by-slice matching and averaging of the bunch spot size over the acceleration distance, respectively. In this example, the location of the bunch head was $k_p \zeta_{\text{head}} = -10.8$ and the bunch had an initial emittance such that $k_p \epsilon_x = 0.05$. All the other parameters were as before (see Sec. II). In order to mitigate the computational cost, these results were obtained with a frozen field approximation (i.e., the particles of the witness bunch are pushed in a nonevolving wakefield). This approach has shown both reasonable agreement of the energy spread and the emittance of the witness bunch with full quasistatic PIC simulations. The agreement is facilitated by the slice-by-slice matching, which mitigates the witness beam evolution. We see that the bunch obtained without slice-by-slice matching and using the initial longitudinal uniform σ_r (blue line in Fig. 4) shows a range of mean energy variation of $\Delta E \approx 100$ MeV. Using the slice-by-slice matching (red line) reduces the amplitude of the variation to $\Delta E \approx 60$ MeV. Finally, by using the energy-averaged bunch size $\bar{\sigma}_r(\zeta)$ in the calculation of the optimal




FIG. 4. Mean energy of each slice vs longitudinal position in the positron bunch after acceleration. The blue line refers to algorithm flattening $\langle E_z \rangle$ with a longitudinal uniform σ_r . The red line and the green line refer to additionally applying slice-by-slice matching and averaging of the bunch spot size over the acceleration distance, respectively. That way, the correlated energy spread can be reduced to the noise level.

bunch (green line) the correlated energy spread is essentially removed ($\Delta E \approx 3$ MeV).

C. Minimizing the uncorrelated energy-spread

Whereas the correlated energy-spread can be completely eliminated, the uncorrelated energy-spread can only be reduced, as it arises from the transverse nonuniformity of E_z . We did not identify a strategy to reduce the transverse gradient of E_z by loading the wake with a positron bunch. However, we explored two possible solutions to minimize the impact of such gradient and reduce the uncorrelated energy spread. First, one can position the witness bunch in a region of the wake where E_z is transversally as flat as possible, and second, one can use a transversally smaller witness beam.

As described in Sec. II, E_z flattens transversally further behind the driver (i.e., for more negative ζ). Therefore, it is favorable to choose the starting position of the bunch furthest behind the driver, which still has a reasonable efficiency. According to this criterion and the results from Sec. III A, the optimal starting position is $k_p \zeta_{\text{head}} = -10.8$.

In the following, we study the dependence of the uncorrelated energy-spread on the witness bunch emittance. A matched bunch with a smaller emittance will have a smaller transverse extent and will sample a smaller domain of E_z and, hence, it will acquire a smaller uncorrelated energy spread. For a flat beam, and assuming that in the vicinity of the axis the accelerating field can be modeled as $E_z(x) = E_{z,0} - \beta|x|$, where β describes the transverse gradient of E_z (see inset of Fig. 1), then, from geometric considerations, we expect the relative slice (uncorrelated) energy spread at saturation to scale as $\sigma_y/\gamma \sim \beta\sigma_r/E_{z,0}$, and so $\sigma_y/\gamma \rightarrow 0$ in the limit of a small bunch. We note that it is not possible with our current numerical tools to model collider-relevant low-emittance witness beams, owing to the required high resolution and associated computational costs. To overcome this limitation, we use a reduced model to assess the scaling of the energy-spread for these conditions. Since the correlated energy-spread can be eliminated with the procedure discussed in the previous section, we consider a single slice of the beam in the reduced model. We chose the slice of the peak current of the positron bunch, which we have found to reasonably represent the total energy-spread of the bunch. Using the previous example with a starting position of $k_p \zeta_{\text{head}} = -10.8$, the peak of the current is located at $k_p \zeta_{\text{peak}} = -11.45$. We reuse the simulation, which included the slice-by-slice matching and the averaging over the acceleration distance. Assuming the same density of $n_0 = 5 \times 10^{17} \text{ cm}^{-3}$ as for the efficiency consideration, the emittance of the beam is $\epsilon_x = 0.05 k_p^{-1} = 0.38 \mu\text{m}$. In the reduced model, we generate test particles, which we advance with a second-order-accurate particle pusher in the radial fields provided by the simulation. High-resolution

simulations with the cylindrically symmetric PIC code INF&RNO indicate that the focusing field converges toward a step function [12]. Likewise, we model the focusing field in the reduced model with a piecewise constant function, $(E_x - B_y)/E_0 = -\alpha \text{sign}(x)$, where $\alpha = 0.6$ for our example. We have found the model in reasonable agreement with HiPACE simulations in terms of energy-spread, emittance, and bunch size evolution. This is shown for the energy-spread in Fig. 5. The black dashed line and the blue solid line describe the energy-spread at an emittance of $\epsilon_x = 0.38 \mu\text{m}$ obtained from the HiPACE simulation and the reduced model, respectively. The energy-spread of the peak-current slice of the beam is $\approx 0.65\%$ for both the simulation (dashed line) and the reduced model (blue line). The final energy-spread and the emittance growth of the whole bunch in the PIC simulation are $\approx 0.7\%$ and $\approx 2\%$ (both not shown in Fig. 5), respectively. Under the assumption that a smaller emittance beam with the same charge does not significantly change the wake structure, we can decrease the beam emittance in the reduced model to previously numerically inaccessible values. The results are shown in Fig. 5. The results of the reduced model indicate that for emittances well below $0.1 \mu\text{m}$, we can achieve energy spreads below 0.1% . The red and green line denote the energy spreads for initial emittances of $\epsilon_x = 0.19 \mu\text{m}$ and $\epsilon_x = 0.08 \mu\text{m}$, respectively. Their corresponding final energy spreads are 0.3% and 0.1% . This indicates the path to possible collider-relevant parameters. However, this model does not capture the change of the wake structure due to a reduced witness bunch spot size. Eventually, when the on-axis density of the positron bunch exceeds the density of the background electrons, we expect a significant disruption of the positron accelerating wake structure. Additionally, a finite initial background plasma temperature can smooth the piecewise constant focusing field, possibly affecting the results presented here. These effects

will be the topic of further research and require extensive development of simulation tools to enable detailed studies.

IV. CONCLUSION

High-quality positron acceleration with sub-percent-level energy spread is possible in beam driven plasma wakefield accelerators. Utilizing an electron drive beam and a narrow plasma column allows for high-charge, and low-emittance positron beams. By shaping the longitudinal density profile of a transversally Gaussian witness beam, the energy-spread can be controlled and kept at the subpercent-level. Thereby, correlated energy spread can be completely eliminated. The uncorrelated energy spread scales with the transverse beam spot size. Our results indicate that using collider-relevant beam emittances might yield energy spreads as low as 0.1% . Further research will aim to strengthen this result. Additionally, the efficiency might be increased by proper shaping the drive beam [20,21], by optimizing the transverse plasma profile [12] or by using the here proposed technique to generate longitudinally chirped bunches. Extending these results to higher efficiencies will pave the path to a plasma-based collider.

ACKNOWLEDGMENTS

We acknowledge the support of the Director, Office of Science, Office of High Energy Physics, of the U.S. Department of Energy, and its Science User Facility National Energy Research Scientific Computing Center (NERSC) under Contract No. DE-AC02-05CH11231. We gratefully acknowledge the Gauss Centre for Supercomputing e.V. [22] for funding this project by providing computing time through the John von Neumann Institute for Computing (NIC) on the GCS Supercomputer JUWELS at Jülich Supercomputing Centre (JSC). We acknowledge the Funding by the Helmholtz Matter and Technologies Accelerator Research and Development Program.

APPENDIX: ALGORITHM FOR DETERMINING THE BUNCH PROFILE WITH OPTIMIZED BEAM LOADING

The algorithm used in this study calculates, by exploiting the quasistatic approximation, the longitudinal current profile of a witness bunch that maintains the average accelerating gradient over the full bunch length. The average accelerating gradient is set at the bunch head, $\langle E_{z,\text{head}} \rangle$. The bunch is constructed recursively by stacking infinitesimal longitudinal slices of charge, one after the other, starting from the head and going toward the tail of the bunch. For each slice, the calculation of the optimal current is done using an optimized bisection procedure.

The steps of the algorithm are as follows. First, for any generic longitudinal slice i ($i = 0$ represents the bunch head, slices are counted starting from the head), the




FIG. 5. Relative slice energy spread vs acceleration distance. Advancing test particles in an approximated step function yields similar energy-spread in comparison with the HiPACE simulation. The results indicate that emittances smaller than $0.1 \mu\text{m}$ induce an energy spread below 0.1% .

algorithm computes the weighted accelerating field right behind the current slice assuming zero charge in the slice. We denote this quantity by $\langle E_{z,i} \rangle$. We then check if $|\langle E_{z,i} \rangle| > |\langle E_{z,\text{head}} \rangle|$. The absolute value is used so the algorithm works for both electron and positron witness bunches. If this condition is not fulfilled then no further beam loading is possible and the recursive procedure terminates (i.e., the bunch tail is reached). On the other hand, if the condition is satisfied, then beam loading is possible and the algorithm initializes the optimized bisection procedure to determine the current in the slice. We recall that the current is set via the g_{\parallel} function in Eq. (1). In order for the bisection procedure to converge, values of the current lower ($g_{\parallel,\text{min}}$) and higher ($g_{\parallel,\text{max}}$) than the optimal one need to be determined. Since we know that with no charge in the i th slice we have $|\langle E_{z,i} \rangle| > |\langle E_{z,\text{head}} \rangle|$, then we can set $g_{\parallel,\text{min}} = 0$. Determining $g_{\parallel,\text{max}}$ requires a trial and error procedure where, starting from, e.g., $g_{\parallel,\text{max}} = 1.2g_{\parallel,i-1}$, the value of the current in the slice is progressively increased in a geometric way (i.e., typically multiplying the current by a factor 10) until overloading of the wake is reached, i.e., until the condition $|\langle E_{z,i} \rangle| < |\langle E_{z,\text{head}} \rangle|$ is satisfied. Note that every time the value of the current in the slice is changed, a solution of the quasistatic field equations for the slice is required in order to determine the current value of the weighted accelerating field behind the slice. Once $g_{\parallel,\text{min}}$ and $g_{\parallel,\text{max}}$ are known, the optimized bisection procedure begins. A new value of the current is computed according to

$$g_{\parallel} = w_g g_{\parallel,\text{min}} + (1 - w_g) g_{\parallel,\text{max}}, \quad (\text{A1})$$

where $w_g = (|\langle E_{z,\text{head}} \rangle| - |\langle E_{z,\text{min}} \rangle|) / (|\langle E_{z,\text{max}} \rangle| - |\langle E_{z,\text{min}} \rangle|)$, and where $\langle E_{z,\text{min}} \rangle$ and $\langle E_{z,\text{max}} \rangle$ are the averaged field values behind the slice correspond to $g_{\parallel,\text{max}}$ and $g_{\parallel,\text{min}}$, respectively. The bisection procedure terminates, and the algorithm advances to the next slice ($i + 1$), if the averaged field computed with g_{\parallel} converges to $\langle E_{z,\text{head}} \rangle$ within a predetermined tolerance, otherwise $g_{\parallel,\text{min}}$ and $g_{\parallel,\text{max}}$ are updated and a new optimized bisection is performed. We note that by using Eq. (A1) instead of the classical bisection procedure, i.e., $g_{\parallel} = 0.5(g_{\parallel,\text{min}} + g_{\parallel,\text{max}})$, the number of iterations required to reach convergence is significantly reduced.

Algorithm modifications for slice-by-slice matching and average bunch size

Incorporating the slice-by-slice matching procedure into the algorithm requires the following modification. At each slice i , the matched spot size $\sigma_{r,\text{matched}}$ needs to be determined. This is done by exploiting a fixed-point method. We generate a Gaussian test particle distribution with some rms size, $\sigma_r(i)$. As an initial guess, the spot size from the previous slice, $\sigma_r(i - 1)$, is used. Then, the test

particles are evolved in time without acceleration in the focusing field given by $(E_x - B_y)(i - 1)$ by using a second order accurate particle pusher until the second order spatial moment of the distribution has saturated. The value of the moment is used to set a new value for $\sigma_r(i)$, and the whole process is repeated until the sequence of values of $\sigma_r(i)$ has converged. Note that the focusing field of the slice $i - 1$ is used to compute $\sigma_r(i)$ under the assumption that the longitudinal resolution is high enough that the focusing field changes only marginally between two adjacent slices. To further take into account the spot size reduction due to acceleration of the particles, the averaged matched spot size $\bar{\sigma}_{r,\text{matched}}$ can be calculated via Eq. (4). Finally, $\sigma_{r,\text{matched}}$ or $\bar{\sigma}_{r,\text{matched}}$ can be used to calculate the average accelerating field $\langle E_z \rangle$ by Eq. (2). It should be noted that $\bar{\sigma}_{r,\text{matched}}$ is only used to calculate $\langle E_z \rangle$, the bunch is still generated with a spot size of $\sigma_{r,\text{matched}}$.

- [1] C. B. Schroeder, E. Esarey, C. G. R. Geddes, C. Benedetti, and W. P. Leemans, Physics considerations for laser-plasma linear colliders, *Phys. Rev. Accel. Beams* **13**, 101301 (2010).
- [2] J. B. Rosenzweig, B. Breizman, T. Katsouleas, and J. J. Su, Acceleration and focusing of electrons in two-dimensional nonlinear plasma wake fields, *Phys. Rev. A* **44**, R6189 (1991).
- [3] A. Pukhov and J. Meyer-ter Vehn, Laser wake field acceleration: The highly non-linear broken-wave regime, *Appl. Phys. B* **74**, 355 (2002).
- [4] 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, E. Oz, R. H. Siemann, D. Walz, and M. Zhou, Energy doubling of 42 GeV electrons in a metre-scale plasma wakefield accelerator, *Nature (London)* **445**, 741 (2007).
- [5] M. Litos *et al.*, High-efficiency acceleration of an electron beam in a plasma wakefield accelerator, *Nature (London)* **515**, 92 (2014).
- [6] C. Lindström, J. M. Garland, S. Schröder, L. Boulton, G. Boyle, J. Chappell, R. D'Arcy, P. Gonzalez, A. Knetsch, V. Libov, G. Loisch, A. Martinez de la Ossa, P. Niknejadi, K. Pöder, L. Schaper, B. Schmidt, B. Sheeran, S. Wesch, J. Wood, and J. Osterhoff, Energy-Spread Preservation and High Efficiency in a Plasma-Wakefield Accelerator (to be published).
- [7] N. Jain, T. M. Antonsen, Jr., and J. P. Palastro, Positron Acceleration by Plasma Wakefields Driven by a Hollow Electron Beam, *Phys. Rev. Lett.* **115**, 195001 (2015).
- [8] S. Corde, E. Adli, J. Allen, W. An, C. Clarke, C. Clayton, J. Delahaye, J. Frederico, S. Gessner, S. Green *et al.*, Multi-gigaelectronvolt acceleration of positrons in a self-loaded plasma wakefield, *Nature (London)* **524**, 442 (2015).
- [9] C. B. Schroeder, D. H. Whittum, and J. S. Wurtele, Multimode Analysis of the Hollow Plasma Channel Wakefield Accelerator, *Phys. Rev. Lett.* **82**, 1177 (1999).

- [10] S. Gessner, E. Adli, J. M. Allen, W. An, C. I. Clarke, C. E. Clayton, S. Corde, J. Delahaye, J. Frederico, S. Z. Green *et al.*, Demonstration of a positron beam-driven hollow channel plasma wakefield accelerator, *Nat. Commun.* **7**, 11785 (2016).
- [11] C. A. Lindström *et al.*, Measurement of Transverse Wakefields Induced by a Misaligned Positron Bunch in a Hollow Channel Plasma Accelerator, *Phys. Rev. Lett.* **120**, 124802 (2018).
- [12] S. Diederichs, T. J. Mehrling, C. Benedetti, C. B. Schroeder, A. Knetsch, E. Esarey, and J. Osterhoff, Positron transport and acceleration in beam-driven plasma wakefield accelerators using plasma columns, *Phys. Rev. Accel. Beams* **22**, 081301 (2019).
- [13] S. Zhou, W. Lu, W. Mori, and W. An, Positron acceleration in transversely tailored plasmas, *FACET-II Science Workshop* (2019) https://conf.slac.stanford.edu/facet-2-2019/sites/facet-2-2019.conf.slac.stanford.edu/files/basic-page-docs/positron%20acceleration%20in%20transversely%20tailored%20plasmas_ShiyuZhou.pdf.
- [14] T. Katsouleas, S. Wilks, P. Chen, J. M. Dawson, and J. J. Su, Beam loading in plasma accelerators, *Part. Accel.* **22**, 81 (1987), <https://s3.cern.ch/inspire-prod-files-3/30e3b8e467eb7298ce19d6bbc95b457f>.
- [15] 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, *Phys. Rev. Lett.* **101**, 145002 (2008).
- [16] T. Mehrling, C. Benedetti, C. B. Schroeder, and J. Osterhoff, HiPACE: A quasi-static particle-in-cell code, *Plasma Phys. Controlled Fusion* **56**, 084012 (2014).
- [17] C. Benedetti, C. B. Schroeder, E. Esarey, C. G. R. Geddes, and W. P. Leemans, Efficient modeling of laser-plasma accelerator staging experiments using INF&RNO, *AIP Conf. Proc.* **1812**, 050005 (2017).
- [18] P. Sprangle, E. Esarey, and A. Ting, Nonlinear interaction of intense laser pulses in plasmas, *Phys. Rev. A* **41**, 4463 (1990).
- [19] C. Benedetti, C. B. Schroeder, E. Esarey, and W. P. Leemans, Emittance preservation in plasma-based accelerators with ion motion, *Phys. Rev. Accel. Beams* **20**, 111301 (2017).
- [20] P. Chen, J. J. Su, J. M. Dawson, K. L. F. Bane, and P. B. Wilson, Energy Transfer in the Plasma Wake-Field Accelerator, *Phys. Rev. Lett.* **56**, 1252 (1986).
- [21] G. Loisch, G. Asova, P. Boonpornprasert, R. Brinkmann, Y. Chen, J. Engel, J. Good, M. Gross, F. Grüner, H. Huck *et al.*, Observation of High Transformer Ratio Plasma Wakefield Acceleration, *Phys. Rev. Lett.* **121**, 064801 (2018).
- [22] www.gauss-centre.eu

Stable electron beam propagation in a plasma column

Cite as: Phys. Plasmas **29**, 043101 (2022); doi: 10.1063/5.0087807

Submitted: 9 February 2022 · Accepted: 16 March 2022 ·

Published Online: 4 April 2022



S. Diederichs,^{1,2,3,a)} C. Benedetti,² E. Esarey,² M. Thévenet,¹ J. Osterhoff,¹ and C. B. Schroeder^{2,4}

AFFILIATIONS

¹Deutsches Elektronen-Synchrotron DESY, Notkestr. 85, 22607 Hamburg, Germany

²Lawrence Berkeley National Laboratory, Berkeley, California 94720, USA

³University of Hamburg, Institute of Experimental Physics, Luruper Chaussee 149, 22607 Hamburg, Germany

⁴Department of Nuclear Engineering, University of California, Berkeley, California 94720, USA

^{a)}Author to whom correspondence should be addressed: severin.diederichs@desy.de

ABSTRACT

The stability of plasma-based accelerators against transverse misalignments and asymmetries of the drive beam is crucial for their applicability. Without stabilizing mechanisms, even small initial offsets of the drive beam centroid can couple coherently to the plasma wake, grow, and ultimately lead to emittance degradation or beam loss for a trailing witness beam. In this work, we demonstrate the intrinsic stability of a beam propagating in a plasma column. This result is relevant in the context of plasma-based positron acceleration, where a wakefield suitable for the transport and acceleration of a positron witness beam is generated in a plasma column by means of an electron drive beam. The stable propagation of the drive beam is a necessary condition for the experimental implementation of this scheme. The differences and similarities of stabilizing mechanisms in a plasma column compared to a homogeneous plasma are identified via theory and particle-in-cell simulations. Experimental tolerances are given, demonstrating the experimental feasibility of the scheme.

© 2022 Author(s). All article content, except where otherwise noted, is licensed under a Creative Commons Attribution (CC BY) license (<http://creativecommons.org/licenses/by/4.0/>). <https://doi.org/10.1063/5.0087807>

I. INTRODUCTION

Plasma-based accelerators provide extreme field gradients and enable compact accelerator facilities, potentially drastically reducing their costs. In a plasma-based accelerator, a high-intensity laser pulse¹ or an ultra-relativistic, high-density particle bunch² drives a plasma wake, which can be utilized to accelerate a trailing witness bunch. In the so-called blowout regime,^{3,4} the driver is strong enough to expel all plasma electrons along the propagation axis, leaving an ion cavity with accelerating and focusing fields on the order of tens to hundreds of GV/m. While high-energy-gain,^{5,6} high-efficiency,⁷ and low-energy-spread⁸ electron acceleration were demonstrated experimentally in this regime, the blowout regime is not applicable to positron acceleration because the ion cavity defocuses positron beams. Application of plasma accelerator technology to an electron-positron collider requires the development of plasma-based methods for positron acceleration. To date, stable, high-quality positron beam acceleration in plasma remains a challenge.

Several positron acceleration concepts have been proposed, including utilizing positron drive beams,⁹ hollow-core drive beams¹⁰ or lasers pulses,¹¹ or the back of the blowout wake.¹² Unfortunately,

none of these concepts support low emittance, low energy spread, and reasonable efficiency simultaneously. Although hollow core plasma channels are a promising candidate,^{13,14} they suffer from intrinsic beam-breakup instability due to the absence of focusing fields for the drive beam.^{13,15} Using asymmetric drive beams provides stability in at least one direction,¹⁶ but only positron beams with large beam emittances ($>50 \mu\text{m rad}$) have been accelerated in simulations, which is too large for collider applications. The wake generated in a thin, warm, quasi-hollow plasma channel provides accelerating fields for positrons while being robust against instabilities.¹⁷ However, this scheme has been investigated via simulations for a positron beam with several $\mu\text{m rad}$ emittance and several percent relative energy spread only, i.e., a beam quality insufficient for collider-relevant applications.

Recently, a concept has been proposed that utilizes an electron drive beam and a plasma column to generate positron-accelerating and -focusing wakefield structures that show sub- $\mu\text{m rad}$ emittance and sub-percent energy spread positron acceleration in simulations.^{18,19} Since the scheme relies on a cylindrically symmetric wakefield structure, the drive beam stability is of utmost importance. Previous work has reported on the attraction of an electron beam

toward a column of neutral plasma²⁰ and on the deflection of an electron beam at the plasma-gas boundary.^{21–23} Furthermore, a severe misalignment between the column and the electron beam was found to seed the hosing instability.²⁴

In this article, we analyze the drive beam stability in depth via theory and particle-in-cell (PIC) simulations. We compare the susceptibility of tilted beams toward the hosing instability in a plasma column and in a homogeneous plasma. We find that the drive beam is self-stabilizing toward tilts and transverse offsets with respect to the center of the plasma column. Finally, we quantify the expected stability of a beam in a plasma column in the context of a potential future experimental realization at FACET-II.²⁵

The paper is organized as follows: In Sec. II, we derive an analytical model for the hosing instability in plasma columns and validate it with PIC simulations. Results obtained for the column are compared with that obtained in a homogeneous plasma. In Sec. III, we investigate drive beams with a transverse displacement with respect to the plasma column center. In Sec. IV, we evaluate the stability for a possible experimental implementation. Section V concludes this work.

II. HOSING FOR A BEAM WITH A TILT IN THE CENTROID

The transverse stability of plasma wakefield accelerators is a long-standing problem that arises due to the coupling between the beam centroid, X_b , and the plasma wake centroid, X_p , created by the beam space-charge. For example, an initial displacement of the beam centroid with respect to the wake centroid affects the evolution of the wake centroid itself, which in turn resonantly affects the beam centroid at all phases behind and so forth. The amplification of an initial transverse beam displacement had originally been predicted to grow exponentially in time and along the beam,^{26,27} threatening the application of plasma wakefield accelerators. Later, various mechanisms similar to the Balakin–Novokhatsky–Smirnov (BNS) damping in conventional accelerators²⁸ were identified, predicting saturation of the hosing instability. These include energy spread of the drive beam,²⁹ non-uniform focusing fields due to ion motion,^{30,31} or longitudinally varying focusing fields in the case of large beams sizes.³²

Modeling the hosing instability has significantly increased the understanding of the relevant physics. In a homogeneous plasma, tilted beams are often used to seed and study the hosing instability. In the following, we further develop and expand the most recent hosing model presented in Ref. 33 to include the effect of a finite-radius plasma column. An accurate description of the plasma response is critical because, as was shown in previous works, the relativistic mass gain/loss^{27,33} and the collective behavior of the plasma sheath have been found to mitigate the hosing instability.³³ We use the adapted model to study the stability of a tilted beam against hosing in a column and compare it with the homogeneous case, thereby gaining understanding of the physics relevant in the column case.

We consider a setup where an electron beam propagating in a plasma column drives a wake in the blowout regime, as illustrated in Fig. 1(a), where we plot the plasma charge density ρ (red-gray color-scale) and the beam density (blue color-scale). Within Secs. II and III, all length scales are normalized to the plasma wave number $k_p = \omega_p/c$, all times to the plasma frequency $\omega_p = \sqrt{4\pi n_0 e^2/m_e}$, densities to the ambient plasma density n_0 , charges to the elementary charge e , masses to the electron mass m_e , fields to the cold,

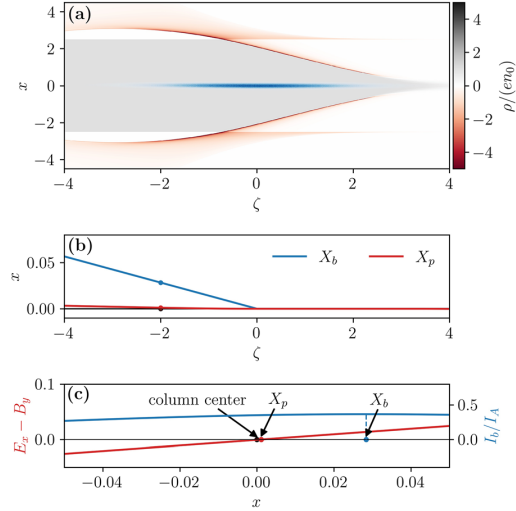


FIG. 1. Characterization of the wake structure of a tilted drive beam in a plasma column. (a) Snapshot of the plasma charge density in the x - ζ -plane for a drive beam (shown in blue) with an initially tilted beam centroid. (b) The corresponding ζ -dependent wake centroid X_p (red line) and beam centroid X_b (blue line). Over a longer propagation distance, the small induced offset of X_p by the beam offset X_b can build up and lead to strong hosing. (c) Line-out of the beam current profile (blue line) and focusing wakefield (red line) along x at $\zeta = -2$.

non-relativistic wave-breaking field $E_0 = m_e c \omega_p / e$, and potentials to $m_e c^2 / e$. The beam centroid is tilted transversely, as depicted in Fig. 1(b), which shows the beam centroid X_b (blue line) and the plasma wake centroid X_p (red line) along the co-moving variable $\zeta = z - ct$ (z is the longitudinal variable, c is the speed of light in vacuum, and t is the time). Here, the wake centroid $X_p(\zeta)$ is defined as the x -position of the zero-crossing of the focusing field $E_x - B_y$. A transverse line-out at $\zeta = -2$ is presented in Fig. 1(c), where the blue line denotes the beam current profile and the red line a lineout of the transverse wakefield, respectively.

A. Analytical model

The mathematical model presented in Ref. 33 describes the hosing by means of a set of coupled partial differential equations for the evolution of the ζ -dependent plasma wake centroid and the beam centroid, namely,

$$\frac{\partial^2 X_p}{\partial \zeta^2} + C_d(\zeta) \frac{\partial X_p}{\partial \zeta} + \frac{C_p(\zeta)}{2} X_p = \frac{C_b(\zeta)}{2} X_b, \quad (1a)$$

$$\frac{\partial^2 X_b}{\partial t^2} + \lambda(\zeta, t) \frac{\partial X_b}{\partial t} + \Omega_\beta^2(\zeta, t) X_b = \Omega_\beta^2(\zeta, t) X_p. \quad (1b)$$

Both are driven, damped harmonic oscillator-like equations for $X_p(\zeta, t)$ along $\zeta = z - ct$ and $X_b(\zeta, t)$ along t , respectively.

The coefficients of the beam centroid equation $\lambda(\zeta, t)$ and $\Omega_\beta^2(\zeta, t)$ were introduced in Ref. 29. The term $\lambda(\zeta, t)$ represents the change in the amplitude of the beam centroid oscillation due to relativistic mass gain or loss of the beam electrons, while $\Omega_\beta^2(\zeta, t)$ accounts

for the effect of the change in the betatron frequency for a changing beam energy.

The coefficients for the plasma wake centroid C_d , C_p , and C_b were introduced in Ref. 33: $C_p/2$ is the square of the undamped oscillation wavenumber of the system, C_b describes the coupling of the wake centroid to the beam centroid, and C_d describes the damping/amplification of the wake centroid oscillation due to the relativistic mass gain/loss of the plasma electrons in the sheath. In the case of a narrow beam in the blowout regime, where the overlap between the beam current density and the electron sheath is negligible, the three wake centroid coefficients are³³

$$C_p = -\frac{1}{\langle 1 + \Psi \rangle_r} \left[\left\langle \left(1 + \frac{J_{z,p}}{\tilde{n} - J_{z,p}} \right) \frac{\partial_r \Psi}{r} \right\rangle_r + \left\langle \frac{\partial_z A_r}{r} \right\rangle_r + \left\langle \frac{\partial_r A_{z,p}}{r} \right\rangle_r \right], \quad (2)$$

$$C_b = \frac{1}{\langle 1 + \Psi \rangle_r} \left\langle \frac{\partial_r A_{z,b}}{r} \right\rangle_r, \quad (3)$$

$$C_d = \frac{\langle \partial_z \Psi \rangle_r}{\langle 1 + \Psi \rangle_r}, \quad (4)$$

with the radial moments being defined as

$$\langle \Phi(r) \rangle_r = \frac{1}{N} \int_0^\infty \tilde{n}(r) \Phi(r) r dr, \quad (5)$$

and where the normalization factor N is given by

$$N = \int_0^\infty \tilde{n}(r) r dr. \quad (6)$$

Here, r denotes the radius; Ψ is the pseudo-wakefield potential; A_r is the radial component of the vector potential; $A_{z,b}$ and $A_{z,p}$ are the longitudinal components of the vector potential associated with the beam and plasma, respectively; $J_{z,p}$ represents the longitudinal plasma current density; and $\tilde{n}(r)$ the plasma density profile of the electron sheath. Thus, if the electromagnetic potentials, the plasma current density, and the plasma density profile are known, the coefficients can be calculated. Several analytical models in the blowout regime^{34–37} calculate the electromagnetic potentials for a given beam current profile and parameterization of the electron sheath. The model in Ref. 35 assumes an exponentially decaying current and electron density profile in the sheath and has shown good agreement with simulations in the analysis of hosing in a homogeneous plasma.³³ Namely, the electron density in the sheath is assumed to be

$$\tilde{n} = \tilde{n}_s \exp\left(-\frac{r - R_b}{\Delta_\rho}\right) \Theta(r - R_b), \quad (7)$$

where R_b is the blowout radius, Δ_ρ is the sheath thickness, \tilde{n}_s is the peak sheath electron density, and $\Theta(x)$ is the Heaviside step function. We adapt the model in Ref. 35 for plasma columns by the following assumptions:

1. For a blowout radius R_b smaller than the column radius R_p , the model remains unchanged;
2. For a blowout radius R_b larger than the column radius R_p , we assume the ions extend to R_p and that there is vacuum/non-

ionized gas between R_p and R_b . This reduces the peak sheath electron density, since less electrons contribute to it.

The assumptions are illustrated in Fig. 2, which shows the focusing field (blue lines) for the column (solid lines) and a homogeneous plasma (dashed lines) and the respective source for the plasma potential $S = -(\rho - J_{z,p})$ (black lines), with ρ being the normalized plasma charge density. The lack of ions in the column between R_p and R_b reduces the focusing field at R_b as well as the peak of the source. The electron sheath is spread over a larger radial extent in case of the column, which is consistently observed in simulations. Using the stated assumptions of the model, the potentials A_r , A_z , and Ψ for both the plasma column and the homogeneous plasma are derived in the Appendix. Using these potentials and assuming that the thickness of the plasma electron sheath Δ_ρ is much smaller than the blowout radius ($\Delta_\rho/R_b \ll 1$) and the thickness of the plasma electron current Δ_J spreads over a plasma depth ($\Delta_J \simeq 1$, see the Appendix), we obtain the following coefficients in the column for $R_b < R_p$ (which are equal to that in the homogeneous case³³):

$$C_p = \frac{1 - R_b^2}{4} - \frac{\Delta_\rho}{4R_b} \left[1 + \Lambda - \frac{R_b^2}{4} (1 - R_b^2 - 2R_b R_b' + R_b^2) \right] + \mathcal{O}\left(\frac{\Delta_\rho^2}{R_b^2}\right), \quad (8)$$

$$C_b = \frac{\Lambda}{R_b^2} \left[1 - \Delta_\rho \left(\frac{2}{R_b} + \frac{R_b}{4} \right) \right] + \mathcal{O}\left(\frac{\Delta_\rho^2}{R_b^2}\right), \quad (9)$$

$$C_d = \frac{R_b'}{4} \left[R_b - \frac{\Delta_\rho}{2} \left(\frac{R_b^2}{2} - 1 \right) \right] + \mathcal{O}\left(\frac{\Delta_\rho^2}{R_b^2}\right), \quad (10)$$

with $R_b'(\zeta) = \partial_\zeta R_b(\zeta)$, $\Lambda(\zeta) = 4I_b/I_A$ the integrated beam current, and where $I_A = m_e c^3 / e \simeq 17$ kA is the Alfvén current. Note that the coefficient C_p was rewritten in comparison with Eq. (36) in Ref. 33 to highlight the differences of the coefficient in the plasma column.

In the plasma column with $R_b > R_p$, we obtain the new coefficients

$$C_p = \kappa \frac{1 - R_b^2}{4} - \frac{\kappa \Delta_\rho}{4R_b} \times \left[1 + \Lambda - 4R_b^2 - \frac{R_b^2}{4} (1 - R_b^2 - 2R_b R_b' + R_b^2) \right] + \mathcal{O}\left(\frac{\Delta_\rho^2}{R_b^2}\right), \quad (11)$$




FIG. 2. Source term $S = -(\rho - J_{z,p})$ (black lines) as assumed in the model and resulting focusing field (blue lines) in a plasma column (solid lines) and a homogeneous plasma (dashed lines). In a plasma column with a radius R_p smaller than the blowout radius R_b , the focusing field is reduced at the blowout radius.

$$C_b = \frac{\Lambda}{R_b^2} \left[1 - \Delta_p \left(\frac{2}{R_b} + \frac{\kappa R_b}{4} \right) \right] + \mathcal{O} \left(\frac{\Delta_p^2}{R_b^2} \right), \quad (12)$$

$$C_d = \frac{\kappa R_b}{4} \left[R_b - \frac{\Delta_p}{2} \left(\frac{R_p^2}{2} + 3 \right) \right] + \mathcal{O} \left(\frac{\Delta_p^2}{R_b^2} \right), \quad (13)$$

with $\kappa = R_p^2/R_b^2 < 1$. A qualitative comparison of the stability in the column versus the homogeneous plasma based on the coefficients alone is difficult, because the plasma column consistently features a larger sheath thickness Δ_p than the homogeneous plasma, which can be attributed to the larger spread in the electron trajectories in the sheath due to the reduced focusing field.¹⁸

B. Comparison with simulation results

For a comparison of the transverse beam stability in a plasma column and in a homogeneous plasma, we employ three-dimensional particle-in-cell (PIC) simulations performed using the quasi-static PIC code HiPACE++.³⁸

We consider beam parameters that have been used to demonstrate high-quality positron acceleration^{18,19} and introduce a tilt in the bunch distribution to seed the hosing instability. The beam has a bi-Gaussian distribution with $\sigma_{x,y} = 0.1$, $\sigma_z = 1.41$, $I_b/I_A = 1$, it has a mean energy of $\gamma_0 = 20\,000$, and the emittance is such that the beam is matched in the focusing field of a blowout wake. The initial centroid of the beam is linearly tilted for $\zeta < 0$ according to $X_{b,0} = 0.014 \times \zeta \Theta(-\zeta)$. The beam is sampled with 10^7 macro-particles. The computational domain is $(-16, 16) \times (-16, 16) \times (-6, 6)$ in $x \times y \times \zeta$ in case of the plasma column and $(-8, 8) \times (-8, 8) \times (-6, 6)$ in case of a homogeneous plasma. In both cases, the mesh resolution is $0.0078 \times 0.0078 \times 0.0012$. The plasma column has a radius of $R_p = 2.5$ and is modeled by 25 macro-particles per cell. The homogeneous plasma is modeled by 4 macro-particles per cell. The numerical parameters were chosen to ensure convergence, which is reached for different parameters in the column and the homogeneous plasma. The ion background is assumed immobile. The beam is advanced in both cases by a constant time step of $\Delta t = 10$.

First, the effect of the beam tilt on the wake centroid is compared in the column case and in the homogeneous plasma case by integrating Eq. (1a) backwards in ζ from the start of the tilt to the tail of the beam. To calculate the coefficients C_p , C_d , and C_b in Eq. (1a), $R_b(\zeta)$ is extracted from the first time step of a PIC simulation using a beam without a tilt. The slope of the blowout radius $R_b'(\zeta)$ is calculated from the extracted $R_b(\zeta)$. The parameterization of the electron sheath using Eq. (7) is in reasonable agreement with the simulation for $\Delta_{p,h} = 0.25$ and $\Delta_{p,c} = 0.5$ for the homogeneous plasma and plasma column, respectively. We compare the predicted wake centroid from the model (solid lines) with wake centroid from the simulations (dashed and dashed-dotted lines) in Fig. 3. The wake centroid is extracted by the interpolation to the zero crossing of the transverse wakefield $E_x - B_y$. The model shows reasonable agreement with the simulations. Both the model and the PIC simulation show that the wake centroid displacement in the column is smaller than in the homogeneous plasma. This indicates that a beam propagating in a column is less prone to the hosing instability than in a homogeneous plasma due to the larger sheath thickness in the column case.

To confirm these findings, we solve the coupled equations Eqs. (1a) and (1b). During the numerical integration of Eq. (1b) in t , at




FIG. 3. Wake centroid as a function of ζ for the beam shown in Fig. 2. In the model, the sheath thicknesses are assumed to be $\Delta_{p,h} = 0.25$ and $\Delta_{p,c} = 0.5$ for the homogeneous case and the column case, respectively. The wake centroid deviation due to the beam tilt is smaller for the plasma column than in a homogeneous plasma, indicating a higher stability in the column. This feature is observed in simulations and is well captured by the model.

each time step, we integrate Eq. (1a) in ζ to calculate the wake centroid offset. To calculate $\lambda(\zeta, t)$ (cf. Ref. 29) in Eq. (1b), we additionally extract the accelerating field E_z from the PIC simulation with the aligned beam. We evaluate the beam centroid at the position of the tail $X_{b,tail}$, which is most susceptible to the hosing instability, and compare it with the PIC simulation results. We define the tail in the simulation as $X_{b,tail} \equiv X_b(-4.26 < \zeta < -4.20)$, while we can directly evaluate it in the model at $\zeta = -4.23$.

The model (black lines) shows good agreement with the simulations (blue lines for homogeneous plasma and red lines for plasma column), as shown in Fig. 4. As it can be seen in Figs. 4(a) and 4(b), the envelope of the oscillation in the homogeneous plasma has a higher slope than in the column and reaches its maximum after $\approx 40 \omega_{p,0}^{-1}$, while it takes $\approx 50 \omega_{p,0}^{-1}$ in the column. Here, $\omega_{p,0}$ denotes the initial (i.e., at injection) betatron frequency of the beam. Both cases saturate at roughly two times the initial offset. The presence of an initial uncorrelated relative energy spread detunes the oscillation,²⁹ suppressing the hosing. For instance, an initial energy spread of $\sigma_y/\gamma_0 = 3\%$ damps the oscillation in both the homogeneous plasma and the column case, as shown in Figs. 4(c) and 4(d). Here, the oscillation envelope is damped quicker in the column, which we attribute to the reduced coupling (see Fig. 3). Besides reducing the hosing at the tail, an uncorrelated energy spread also prevents the emergence of a wave-like pattern in the beam profile, which can be seen in the video of the full evolution of the bunches in the supplementary material.

In summary, we find that beam propagation in the column case is more robust against hosing seeded by a tilt in the beam centroid compared to the homogeneous plasma case. Our model demonstrates that this is mostly due to the increased sheath thickness in the plasma column, leading to a decreased coupling of the wakefield centroid X_p to the beam centroid X_b . However, the column has an additional constraint in comparison with the homogeneous plasma, namely, that the propagation axis of the beam must be aligned with the center of the column. The stability for a transverse displacement of the beam with respect to the column axis is discussed in Sec. III.

III. HOSING FOR A BEAM NOT ALIGNED WITH THE COLUMN AXIS

The transverse displacement of the electron beam with respect to the plasma column axis is a problem that may arise in experiments due to machine jitters. A transversely misaligned beam that drives a




FIG. 4. Comparison between the evolution of $X_{b,tail}$ in the model and the simulation in (a) a homogeneous plasma and (b) the column without an initial energy spread. The oscillation in the plasma column rises less steeply than in the homogeneous plasma but saturates at the same value of roughly two times the initial offset. (c) and (d) Show the results for an initial uncorrelated energy spread of $\sigma_\gamma/\gamma_0 = 3\%$. The detuning of the oscillation due to the uncorrelated, relative energy spread has a greater effect in the plasma column and leads to oscillations with an amplitude slightly smaller than that in the homogeneous plasma.

wake in the blowout regime experiences a restoring force toward the column center. In fact, when all plasma electrons of the column are expelled, the majority of the background ions are located on the opposing side of the column with respect to the location of the displaced beam. As the transverse wakefields are mostly caused by the ion distribution, the wake centroid X_p is not aligned with the beam centroid X_b but shifted toward the column center. As a consequence, X_p depends on the co-moving variable ζ , since the ion distribution depends on the blowout radius, as illustrated in Fig. 5. Therefore, the head and the tail of the beam are experiencing different restoring forces, which potentially seed the hosing instability.

The analytical model derived in Sec. II is not applicable in this case. In the solution of the coupled centroid equations, the integration has to begin at the start of the beam displacement, which is at the head of the beam. However, in this location, the assumption that the electron sheath is much smaller than the blowout radius ($\Delta_\rho \ll R_b$) is not fulfilled anymore, rendering the model inaccurate. Therefore, PIC simulations were used to investigate the stability of a transversely displaced beam.

The parameters of the beam considered for this study are the ones introduced in Sec. II B, except for the uncorrelated energy spread that has been set to zero, and a beam size of $\sigma_x = 0.05$. The initial transverse beam offset with respect to the column is set equal to one root-mean-square (rms) size of the beams, namely, $X_{b,0} = \sigma_x = 0.05$.

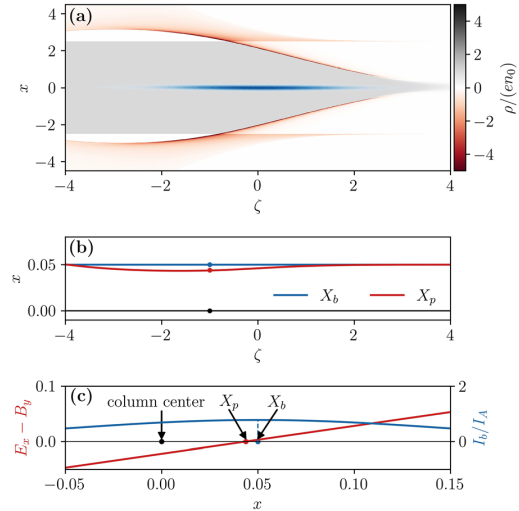


FIG. 5. Characterization of the wake structure of a beam with a transverse offset with respect to the column axis. (a) Snapshot of the plasma charge density in the $x-\zeta$ -plane for a beam (shown in blue) with an initial transverse offset of $X_b = 0.05$ with respect to the plasma column center. (b) The resulting ζ -dependent wake centroid X_p (red line) that lays between the beam centroid X_b (blue line) and the center of the column (black line), leading to an attraction of the beam toward the center of the column. (c) Lineout of the beam current profile (blue line) and focusing wakefield (red line) along x at $\zeta = -1$.

The beam is advanced in the plasma column for 2000 time steps with an adaptive time step, resolving each energy-dependent betatron period of the beam with 81 temporal steps. The evolution of the bunch centroid at the tail of the beam, $X_{b,tail}$, is shown in Fig. 6(a). As indicated in Fig. 5, the tail of the beam drifts toward the column axis and hosing is seeded. Using an immobile ion background (dashed line), the induced hosing grows initially and saturates at an amplitude on the order of the initial offset. When ion motion is enabled (assuming a Helium plasma, solid line), the oscillation is damped, leaving dominantly a drift toward the column axis. The final ζ -dependent bunch centroid is shown in the inset of Fig. 6(a). As predicted by the wake centroid X_p in Fig. 5, the tail of the beam has drifted closer to the column axis than the head. The ζ -dependent temporal evolution of the full bunch is shown in Fig. 6(b). We see that while the tail and the center of the bunch undergo oscillations, the head slowly drifts toward the axis. The oscillations are damped, leaving only a slow drift toward the column axis.

The damping of the oscillation is caused by the detuning effect associated with the motion of background Helium ions.^{39–41} The originally uniform background ion distribution is perturbed by the strong space charge fields of the bunch, leading to ζ -dependent, non-linear focusing fields. Any longitudinal dependence of the focusing fields causes decoherence of the centroid oscillations,^{31,52,42} which is similar to the BNS damping mechanism²⁸ in conventional accelerators. While the longitudinal variation of the focusing fields for small, high-density drive beams is caused by ion motion, larger, less-intense drive beams can be subject to longitudinally varying focusing fields due to the fact

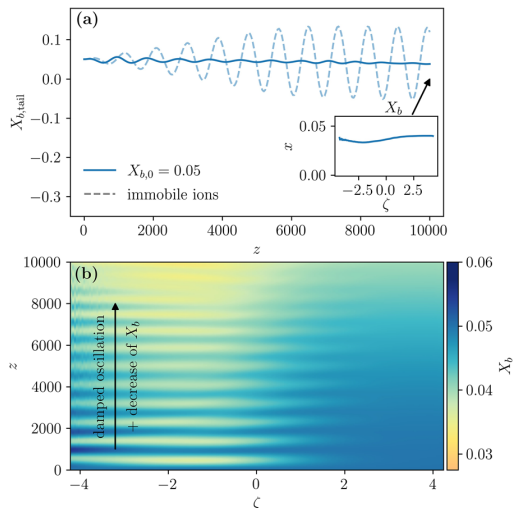


FIG. 6. Evolution (a) of the beam centroid at the tail of the beam $X_{b,tail}$ along the propagation distance for an initial offset of $X_{b,0} = \sigma_x = 0.05$ assuming an immobile ion background (dashed line) and with mobile ions in a Helium plasma (solid line). The inset shows the final slice-dependent beam centroid. Waterfall plot (b) with the evolution of the beam centroid for $X_{b,0} = 0.05$ in Helium vs longitudinal coordinate (horizontal axis) and propagation distance (vertical axis). The beam centroid decreases via a damped oscillation toward the column center.

that their head is not experiencing a full blowout wake, and this can suppress the hosing instability, as discussed in Ref. 32.

In summary, an initially transversely displaced electron beam in a plasma column is attracted by the column center, but seeds the hosing instability by inducing a longitudinal dependence in the wake centroid X_p . However, multiple damping mechanisms cause decoherence of the oscillation, such as ion motion, energy spread, or the head not being in the blowout wake. These processes mitigate the hosing instability and lead to an overall stable propagation of the beam in the plasma column.

IV. ASSESSMENT OF STABILITY FOR FACET-II BEAM PARAMETERS

The demonstrated stability against tilts and transverse offsets makes plasma columns a promising candidate for positron acceleration. Currently, FACET-II²⁵ is the only beam-driven plasma wakefield accelerator facility that plans to accelerate positrons. Using electron drive beams similar to those available at FACET-II, we test both transverse offsets as well as pointing jitters to evaluate tolerances in a possible experiment. Using SI units in this section, the beam is assumed to be bi-Gaussian with $\sigma_{x,y} = 10 \mu\text{m}$, $\sigma_z = 20 \mu\text{m}$, has a charge of $Q_b = 3 \text{ nC}$, a mean energy of $\gamma_0 = 20\,000$, an initial emittance of $\epsilon_0 = 10 \mu\text{m rad}$, and an uncorrelated relative energy spread of $\sigma_\gamma/\gamma_0 = 3\%$. We consider a Helium plasma column with a density of $1 \times 10^{17} \text{ cm}^{-3}$ and with a radius of $R_p = 2.5 k_p^{-1} = 42 \mu\text{m}$. Helium is used owing to its high ionization threshold: for the given beam parameters, other gases would be ionized outside of the column by the driven wakefield (which exceeds the radius of the column) causing the

column to expand and deteriorate the positron accelerating and focusing fields.

Three different offsets are tested in PIC simulations with the same numerical parameters as in Sec. II B. The evolution of the full beam centroid (a) and the centroid measured at the tail (b) are shown in Fig. 7. Here, the tail is defined as $X_{b,tail} = X_b(-62 \mu\text{m} < \zeta < -58 \mu\text{m})$. The blue line corresponds to a small initial offset, $X_{b,0} = 1 \mu\text{m}$, the red line to an offset equal to a rms size of the beam, $X_{b,0} = \sigma_x = 10 \mu\text{m}$, and the green line to an offset equal to half the column radius, $X_{b,0} = 0.5 R_p = 21 \mu\text{m}$. Simulations show that in all cases, an oscillation is induced, which, as expected, is more pronounced at the tail of the beam. However, even for an initial misalignment of $0.5 R_p$ (green line), which significantly perturbs the wakefield, the drive beam does not breakup, and the seeded hosing is mitigated. We find that, in this case, the mitigation is not caused by ion motion, but rather by the uncorrelated energy spread developing as the beam propagates in the plasma. Although the beam is slightly pulled toward the axis, the drift is not sufficient to compensate and cancel the initial offset. For the presented cases, the beam energy is partially depleted at the end of the propagation distance.

Notably, the induced oscillation scales with the initial offset. Although offsets by a large fraction of the column radius significantly distort the wakefield behind the bunch, the intrinsic stability of the bunch is an important finding. In combination with the scaling of the oscillation with the initial offset, the intrinsic stability allows for the alignment of the beam via active feedback loops.²⁴ The emitted betatron radiation by the oscillating bunch scales with the oscillation amplitude, and, thus, with the initial offset, allowing for alignment of the electron beam and the plasma column by minimizing the emitted betatron radiation.

We also tested the stability of the beam against pointing jitters, modeled by adding an initial normalized transverse momentum u_x to the beam particles. It is expected that the maximum pointing angle $X'_b = u_x/u_z$ to achieve stable beam propagation should be much smaller than the ratio of the column radius R_p and the plasma length L_p , i.e., $X'_b \ll R_p/L_p$. Here, u_z is the mean normalized longitudinal




FIG. 7. Evolution of the (a) full beam centroid and (b) beam centroid measured at the tail of the beam for different offsets. Despite offsets of large fractions of the column radius, the overall propagation in a plasma column is stable.




FIG. 8. Evolution of the (a) full beam centroid and (b) beam centroid measured at the tail of the beam for different offsets. The column affects the beam propagation only for large pointing jitters.

momentum of the beam. A plasma length of $L_p = 42$ cm is assumed, such that $R_p/L_p = 10^{-4}$. Three different pointing angles are used to study the pointing jitter. The evolution of the full beam centroids (a) and the centroids measured at the tail (b) are shown in Fig. 8. The blue line corresponds to an initial pointing angle of $X'_{b,0} = 0.01 R_p/L_p = 1 \mu\text{rad}$, the red line to an initial pointing angle of $X'_{b,0} = 0.1 R_p/L_p = 10 \mu\text{rad}$, and the green line to an initial pointing angle of $X'_{b,0} = R_p/L_p = 100 \mu\text{rad}$. The dashed lines indicate the vacuum propagation, which was calculated analytically assuming ballistic propagation in vacuum.

For the small and medium pointing angles, the propagation of the centroids hardly differs from that in vacuum. For the large pointing angle of $X'_{b,0} = R_p/L_p$, the beam centroid significantly deviates from the column axis. In contrast to the vacuum propagation, propagation in the plasma column results in the centroid trajectory being bent toward the column center. This affects only the center and tail of the beam, while the head still propagates almost unchanged away from the column axis, as shown in the videos in the supplementary material and in agreement with previous studies.^{21–23} Therefore, the initial assumption to achieve stable beam propagation that the pointing jitter should be much smaller than R_p/L_p was verified. In the context of the presented parameters, a pointing jitter of $\lesssim 10 \mu\text{rad}$ is required to allow for stable wakefield structures usable for positron acceleration.

In summary, beam parameters similar to those available at FACET-II predict stable beam propagation for even large offsets and moderate pointing jitters. Large offsets or significant pointing jitters deteriorate the wakefield structure and, therefore, a positional alignment on the μm -level and an angular alignment on the $\lesssim 10 \mu\text{rad}$ -level are required for positron-relevant applications. The implications of a misaligned drive beam on the quality of a potential trailing positron beam will be addressed in detail in a forthcoming work.⁴³

V. CONCLUSION

In this work, we have extended an existing model to describe the hosing instability to plasma columns and have validated the model by

means of PIC simulations. We found that plasma columns are in fact more stable against hosing of tilted beams than homogeneous plasmas because of the increased sheath thickness in the column. However, the propagation axis of the beam needs to be aligned with the central axis of the column. A transversely misaligned drive beam is attracted toward the column center with longitudinally varying forces. Although hosing is seeded, it is efficiently damped by various BNS-like mechanisms, including ion motion³¹ or energy spread of the beam.²⁹ Finally, we have evaluated the tolerance for misalignment of an electron beam for the experimental realization at FACET-II. We found that even for large offsets, the beam is stable. The scaling of the oscillation with the initial offset allows for alignment via active feedback loops. The demonstrated overall stability of an electron beam in a plasma column is an important step toward the experimental realization of positron acceleration in a plasma column.

SUPPLEMENTARY MATERIAL

See the supplementary material for videos of the full evolution of the tilted bunches from Sec. II.

ACKNOWLEDGMENTS

We gratefully acknowledge helpful discussions with T.J. Mehrling, A. Martinez de la Ossa, and S. Gessner. This work was supported by the Director, Office of Science, Office of High Energy Physics, of the U.S. Department of Energy, under Contract No. DE-AC02-05CH11231 and used the computational facilities at the National Energy Research Scientific Computing Center (NERSC). We gratefully acknowledge the Gauss Centre for Supercomputing e.V. (www.gauss-centre.eu) for funding this project by providing computing time through the John von Neumann Institute for Computing (NIC) on the GCS Supercomputer JUWELS at Jülich Supercomputing Centre (JSC). This research was supported in part through the Maxwell computational resources operated at Deutsches Elektronen-Synchrotron DESY, Hamburg, Germany. We acknowledge the Funding by the Helmholtz Matter and Technologies Accelerator Research and Development Program.

AUTHOR DECLARATIONS

Conflict of Interest

The authors have no conflicts to disclose.

DATA AVAILABILITY

The data that support the findings of this study are available from the corresponding author upon reasonable request.

APPENDIX: DERIVATION OF THE POTENTIALS IN A PLASMA COLUMN

In the quasi-static approximation, Maxwell's equations can be written as³⁵

$$-\nabla_{\perp}^2 \begin{bmatrix} \mathbf{A} \\ \phi \end{bmatrix} = \begin{bmatrix} \mathbf{J} \\ \rho \end{bmatrix}, \quad (\text{A1})$$

with the Lorentz gauge condition being

$$\nabla_{\perp} \cdot \mathbf{A}_{\perp} = \frac{\partial \psi}{\partial \zeta}, \quad (\text{A2})$$

and the wake potential $\psi = \phi - A_z$, where ϕ and \mathbf{A} are the scalar and vector potentials, respectively. Assuming cylindrical symmetry ($\mathbf{A}_{\perp} = A_r \hat{r}$),

$$\frac{1}{r} \frac{\partial}{\partial r} \left(r \frac{\partial}{\partial r} \right) \left[\frac{\psi}{A_z} \right] = - \left[\frac{\rho - J_z}{J_z} \right]. \quad (\text{A3})$$

Thus, for a given parameterization of ρ and J_z , we can use Eq. (A3) and (A2) to calculate ψ , A_z , and A_r . Following Ref. 35, we assume: First, the beam expels all electrons within the blowout radius, such that ρ is solely determined by the ions inside the ion cavity. Note that $\rho = 0$ for $R_p < r < R_b$. Second, the plasma electron sheath density and current can be parameterized by an exponential function and the source term for Eq. (A3) is $S(\zeta, r) \equiv -(\rho - J_z) \approx -\rho$, which is consistently observed in PIC simulations. Using the exponential parameterization of the sheath, $S(\zeta, r)$ in the homogeneous plasma and in the column, with $R_b < R_p$, is

$$S(\zeta, r) = \begin{cases} -1, & r < R_b(\zeta), \\ S_0(\zeta) e^{-[r-R_b(\zeta)]/\Delta_p}, & r \geq R_b(\zeta), \end{cases} \quad (\text{A4})$$

with $S_0(\zeta)$ being the peak and Δ_p the thickness of the electron sheath, respectively. In the case of the plasma column with $R_b \geq R_p$, the source term is

$$S(\zeta, r) = \begin{cases} -1, & r < R_p < R_b(\zeta), \\ 0, & R_p < r < R_b(\zeta), \\ S_0(\zeta) e^{-[r-R_b(\zeta)]/\Delta_p}, & R_p < R_b(\zeta) < r. \end{cases} \quad (\text{A5})$$

The peak of the source terms $S_0(\zeta)$ is calculated from Eq. (12) in Ref. 35 and is in the case of a plasma column given by

$$S_0(\zeta) = \begin{cases} \frac{R_b^2(\zeta)}{2\Delta_p [R_b(\zeta) + \Delta_p]}, & R_b(\zeta) < R_p, \\ \frac{R_p^2}{2\Delta_p [R_b(\zeta) + \Delta_p]}, & R_b(\zeta) \geq R_p. \end{cases} \quad (\text{A6})$$

The case $R_b(\zeta) < R_p$ is equal to $S_0(\zeta)$ in a homogeneous plasma. The plasma current density $J_{z,p}$ is parameterized by

$$J_{z,p}(\zeta, r) = \begin{cases} 0, & r < R_b(\zeta), \\ J_s(\zeta) e^{-[r-R_b(\zeta)]/\Delta_r}, & r \geq R_b(\zeta), \end{cases} \quad (\text{A7})$$

where $J_s(\zeta)$ is the peak current density and Δ_r is the thickness of the sheath current density, respectively. $J_s(\zeta)$ is calculated by Eq. (16) in Ref. 35 via

$$J_s(\zeta) = \frac{\Lambda(\zeta)/2 - \int_0^{\infty} r dr d^2\Psi/d\zeta^2}{\Delta_J (R_b + \Delta_J)}, \quad (\text{A8})$$

where $\Lambda(\zeta) = 4I_b/I_A$ is the normalized beam current.

Using the source term, we can calculate ψ from (A3) by integration and assuming that $\lim_{r \rightarrow \infty} \psi(r) = 0$. For the case

$R_p \geq R_b(\zeta)$, which is equivalent to a homogeneous plasma, we obtain the same result as Ref. 35,

$$\psi(\zeta, r) = \frac{R_b^2 - r^2}{4} + \frac{\Delta_p R_b^2}{2(R_b + \Delta_p)} \left[1 + e^{R_b/\Delta_p} E_1 \left(\frac{R_b}{\Delta_p} \right) \right] \quad (\text{A9})$$

for $r < R_b(\zeta)$ and

$$\psi(\zeta, r) = \frac{\Delta_p R_b^2}{2(R_b + \Delta_p)} e^{-(r-R_b)/\Delta_p} \left[1 + e^{r/\Delta_p} E_1 \left(\frac{r}{\Delta_p} \right) \right] \quad (\text{A10})$$

for $r \geq R_b(\zeta)$, with E_1 being the exponential integral function. For the case $R_p < R_b(\zeta)$, we obtain the new results,

$$\psi(\zeta, r) = \frac{R_p^2 - r^2}{4} + \frac{R_p^2}{2} \log \left(\frac{R_b}{R_p} \right) + \frac{\Delta_p R_p^2}{2(R_b + \Delta_p)} \left[1 + e^{R_b/\Delta_p} E_1 \left(\frac{R_b}{\Delta_p} \right) \right]$$

for $r < R_p < R_b(\zeta)$,

$$\psi(\zeta, r) = \frac{\Delta_p R_p^2}{2(R_b + \Delta_p)} \left[1 + e^{R_b/\Delta_p} E_1 \left(\frac{R_b}{\Delta_p} \right) \right] + \frac{R_p^2}{2} \left[\log \left(\frac{R_b}{R_p} \right) - \log \left(\frac{r}{R_p} \right) \right]$$

for $R_p < r < R_b(\zeta)$, and

$$\psi(\zeta, r) = \frac{\Delta_p R_p^2}{2(R_b + \Delta_p)} e^{-(r-R_b)/\Delta_p} \left[1 + e^{r/\Delta_p} E_1 \left(\frac{r}{\Delta_p} \right) \right] \quad (\text{A11})$$

for $R_p < R_b(\zeta) < r$.

For the longitudinal component of the vector potential A_z , we get the following results both in the plasma column as in the homogeneous plasma. Within the blowout radius $r < R_b(\zeta)$,

$$A_z(\zeta, r) = \frac{\Lambda(\zeta)}{2} \ln r, \quad (\text{A12})$$

and outside of the blowout radius $r \geq R_b(\zeta)$,

$$A_z(\zeta, r) = \frac{\Lambda(\zeta)}{2} \ln r - J_s(\zeta) \Delta_J \left\{ R_b \ln \left(\frac{r}{R_b} \right) + \Delta_J \left[-1 + \ln \left(\frac{r}{R_b} \right) + e^{-(r-R_b)/\Delta_r} + e^{R_b/\Delta_r} \left(E_1 \left(\frac{r}{\Delta_J} \right) - E_1 \left(\frac{R_b}{\Delta_J} \right) \right) \right] \right\}. \quad (\text{A13})$$

Note that although A_z has the same form in the column and in the homogeneous case, they still differ, because the peak current J_s depends on ψ , which is different in the column and homogeneous case [see Eq. (A8)].

The radial component of the vector potential A_r in the homogeneous plasma yields within the blowout radius $r < R_b(\zeta)$,

$$A_r(\zeta, r) = \frac{R_b}{4(R_b + \Delta_p)^2} \frac{\partial R_b}{\partial \zeta} (R_b^2 + 2\Delta_p R_b + 2\Delta_p^2) \times r \left[1 + e^{R_b/\Delta_p} E_1 \left(\frac{R_b}{\Delta_p} \right) \right], \quad (\text{A14})$$

and outside of the blowout radius $r \geq R_b(\zeta)$, we get

$$A_r(\zeta, r) = \frac{R_b}{4(R_b + \Delta_\rho)^2} \frac{\partial R_b}{\partial \zeta} (R_b^2 + 2\Delta_\rho R_b + 2\Delta_\rho^2) \times \frac{1}{r} \left[R_b^2 + 3\Delta_\rho R_b + 3\Delta_\rho^2 + r^2 e^{R_b/\Delta_\rho} E_1 \left(\frac{r}{\Delta_\rho} \right) - 3\Delta_\rho (\Delta_\rho + r) e^{-(r-R_b)/\Delta_\rho} \right]. \quad (\text{A15})$$

Note that there are sign differences between Eqs. (A14) and (A15) and the corresponding Eqs. (21) and (22) in Ref. 35, which do not originate from the choice of co-moving variable. In the column case with $r, R_p < R_b(\zeta)$, we obtain the new results

$$A_r(\zeta, r) = \frac{R_b R_p^2}{4(R_b + \Delta_\rho)^2} \frac{\partial R_b}{\partial \zeta} r \left[1 + e^{R_b/\Delta_\rho} E_1 \left(\frac{R_b}{\Delta_\rho} \right) \right], \quad (\text{A16})$$

and for $R_p < R_b(\zeta) < r$,

$$A_r(\zeta, r) = \frac{R_b R_p^2}{4(R_b + \Delta_\rho)^2} \frac{\partial R_b}{\partial \zeta} \times \frac{1}{r} \left[R_b^2 + 3\Delta_\rho R_b + 3\Delta_\rho^2 + r^2 e^{R_b/\Delta_\rho} E_1 \left(\frac{r}{\Delta_\rho} \right) - 3\Delta_\rho (\Delta_\rho + r) e^{-(r-R_b)/\Delta_\rho} \right]. \quad (\text{A17})$$


REFERENCES

- ¹E. Esarey, C. B. Schroeder, and W. P. Leemans, "Physics of laser-driven plasma-based electron accelerators," *Rev. Mod. Phys.* **81**, 1229–1285 (2009).
- ²P. Chen, J. M. Dawson, R. W. Huff, and T. Katsouleas, "Acceleration of electrons by the interaction of a bunched electron beam with a plasma," *Phys. Rev. Lett.* **54**, 693–696 (1985).
- ³J. B. Rosenzweig, B. Breizman, T. Katsouleas, and J. J. Su, "Acceleration and focusing of electrons in two-dimensional nonlinear plasma wake fields," *Phys. Rev. A* **44**, R6189–R6192 (1991).
- ⁴K. V. Lotov, "Blowout regimes of plasma wakefield acceleration," *Phys. Rev. E* **69**, 046405 (2004).
- ⁵I. 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, E. Oz, R. H. Siemann, D. Walz, and M. Zhou, "Energy doubling of 42 GeV electrons in a metre-scale plasma wakefield accelerator," *Nature* **445**, 741–744 (2007).
- ⁶A. 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, K. Swanson, L. Fan-Chiang, G. Bagdasarov, N. Bobrova, V. Gasilov, G. Korn, P. Sasorov, and W. P. Leemans, "Petawatt laser guiding and electron beam acceleration to 8 GeV in a laser-heated capillary discharge waveguide," *Phys. Rev. Lett.* **122**, 084801 (2019).
- ⁷M. 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, D. Walz, G. White, Z. Wu, V. Yakimenko, and G. Yocky, "High-efficiency acceleration of an electron beam in a plasma wakefield accelerator," *Nature* **515**, 92–95 (2014).
- ⁸C. A. Lindström, J. M. Garland, S. Schröder, L. Boulton, G. Boyle, J. Chappell, R. D'Arcy, P. Gonzalez, A. Knetsch, V. Libov, G. Loisch, A. Martinez de la Ossa, P. Niknejadi, K. Pöder, L. Schaper, B. Schmidt, B. Sheeran, S. Wesch, J. Wood, and J. Osterhoff, "Energy-spread preservation and high efficiency in a plasma-wakefield accelerator," *Phys. Rev. Lett.* **126**, 014801 (2021).
- ⁹S. Corde, E. Adli, J. Allen, W. An, C. Clarke, C. Clayton, J. Delahaye, J. Frederico, S. Gessner, S. Green *et al.*, "Multi-gigaelectronvolt acceleration of positrons in a self-loaded plasma wakefield," *Nature* **524**, 442 (2015).
- ¹⁰N. Jain, T. M. Antonsen, Jr., and J. P. Palastro, "Positron acceleration by plasma wakefields driven by a hollow electron beam," *Phys. Rev. Lett.* **115**, 195001 (2015).
- ¹¹J. Vieira and J. T. Mendonça, "Nonlinear laser driven donut wakefields for positron and electron acceleration," *Phys. Rev. Lett.* **112**, 215001 (2014).
- ¹²K. V. Lotov, "Acceleration of positrons by electron beam-driven wakefields in a plasma," *Phys. Plasmas* **14**, 023101 (2007).
- ¹³C. B. Schroeder, D. H. Whittum, and J. S. Wurtele, "Multimode analysis of the hollow plasma channel wakefield accelerator," *Phys. Rev. Lett.* **82**, 1177–1180 (1999).
- ¹⁴S. Gessner, E. Adli, J. M. Allen, W. An, C. I. Clarke, C. E. Clayton, S. Corde, J. Delahaye, J. Frederico, S. Z. Green *et al.*, "Demonstration of a positron beam-driven hollow channel plasma wakefield accelerator," *Nat. Commun.* **7**, 11785 (2016).
- ¹⁵C. A. Lindström, E. Adli, J. M. Allen, W. An, C. Beekman, C. I. Clarke, C. E. Clayton, S. Corde, A. Doche, J. Frederico, S. J. Gessner, S. Z. Green, M. J. Hogan, C. Joshi, M. Litos, W. Lu, K. A. Marsh, W. B. Mori, B. D. O'Shea, N. Vafaei-Najafabadi, and V. Yakimenko, "Measurement of transverse wakefields induced by a misaligned positron bunch in a hollow channel plasma accelerator," *Phys. Rev. Lett.* **120**, 124802 (2018).
- ¹⁶S. Zhou, J. Hua, W. An, W. B. Mori, C. Joshi, J. Gao, and W. Lu, "High efficiency uniform wakefield acceleration of a positron beam using stable asymmetric mode in a hollow channel plasma," *Phys. Rev. Lett.* **127**, 174801 (2021).
- ¹⁷T. Silva, L. D. Amorim, M. C. Downer, M. J. Hogan, V. Yakimenko, R. Zgadzaj, and J. Vieira, "Stable positron acceleration in thin, warm, hollow plasma channels," *Phys. Rev. Lett.* **127**, 104801 (2021).
- ¹⁸S. Diederichs, T. J. Mehrling, C. Benedetti, C. B. Schroeder, A. Knetsch, E. Esarey, and J. Osterhoff, "Positron transport and acceleration in beam-driven plasma wakefield accelerators using plasma columns," *Phys. Rev. Accel. Beams* **22**, 081301 (2019).
- ¹⁹S. Diederichs, C. Benedetti, E. Esarey, J. Osterhoff, and C. B. Schroeder, "High-quality positron acceleration in beam-driven plasma accelerators," *Phys. Rev. Accel. Beams* **23**, 121301 (2020).
- ²⁰E. Adli, C. A. Lindström, J. Allen, C. Clarke, J. Frederico, S. Gessner, S. Green, M. Hogan, M. Litos, B. O'Shea *et al.*, "Long-range attraction of an ultrarelativistic electron beam by a column of neutral plasma," *New J. Phys.* **18**, 103013 (2016).
- ²¹P. Muggli, S. Lee, T. Katsouleas, R. Assmann, F.-J. Decker, M. J. Hogan, R. Iverson, P. Raimondi, R. H. Siemann, D. Walz *et al.*, "Refraction of a particle beam," *Nature* **411**, 43–43 (2001).
- ²²P. Muggli, S. Lee, T. Katsouleas, R. Assmann, F. J. Decker, M. J. Hogan, R. Iverson, P. Raimondi, R. H. Siemann, D. Walz, B. Blue, C. E. Clayton, E. Dodd, R. A. Fonseca, R. Hemker, C. Joshi, K. A. Marsh, W. B. Mori, and S. Wang, "Collective refraction of a beam of electrons at a plasma-gas interface," *Phys. Rev. ST Accel. Beams* **4**, 091301 (2001).
- ²³K. Lotov, "Force exerted on particle bunch propagating near plasma-vacuum boundary," *Plasma Phys. Controlled Fusion* **62**, 085002 (2020).
- ²⁴W. An, M. Zhou, N. Vafaei-Najafabadi, K. A. Marsh, C. E. Clayton, C. Joshi, W. B. Mori, W. Lu, E. Adli, S. Corde, M. Litos, S. Li, S. Gessner, J. Frederico, M. J. Hogan, D. Walz, J. England, J. P. Delahaye, and P. Muggli, "Strategies for mitigating the ionization-induced beam head erosion problem in an electron-beam-driven plasma wakefield accelerator," *Phys. Rev. ST Accel. Beams* **16**, 101301 (2013).
- ²⁵V. Yakimenko, L. Alsberg, E. Bong, G. Bouchard, C. Clarke, C. Emma, S. Green, C. Hast, M. J. Hogan, J. Seabury, N. Lipkowitz, B. O'Shea, D. Storey, G. White, and G. Yocky, "FACET-II facility for advanced accelerator experimental tests," *Phys. Rev. Accel. Beams* **22**, 101301 (2019).
- ²⁶D. H. Whittum, W. M. Sharp, S. S. Yu, M. Lampe, and G. Joyce, "Electron-hose instability in the ion-focused regime," *Phys. Rev. Lett.* **67**, 991–994 (1991).
- ²⁷C. Huang, W. Lu, M. Zhou, C. E. Clayton, C. Joshi, W. B. Mori, P. Muggli, S. Deng, E. Oz, T. Katsouleas, M. J. Hogan, I. Blumenfeld, F. J. Decker,

- R. Ischebeck, R. H. Iverson, N. A. Kirby, and D. Walz, "Hosing instability in the blow-out regime for plasma-wakefield acceleration," *Phys. Rev. Lett.* **99**, 255001 (2007).
- ²⁸V. E. Balakin, A. V. Novokhatsky, and V. P. Smirnov, "VLEPP: Transverse beam dynamics," in *Proceedings of the 12th International Conference on High-Energy Accelerators, HEACC 1983: Fermilab, Batavia, August 11–16, 1983* (1983), Vol. C830811, pp. 119–120.
- ²⁹T. J. Mehrling, R. A. Fonseca, A. Martinez de la Ossa, and J. Vieira, "Mitigation of the hose instability in plasma-wakefield accelerators," *Phys. Rev. Lett.* **118**, 174801 (2017).
- ³⁰W. An, W. Lu, C. Huang, X. Xu, M. J. Hogan, C. Joshi, and W. B. Mori, "Ion motion induced emittance growth of matched electron beams in plasma wakefields," *Phys. Rev. Lett.* **118**, 244801 (2017).
- ³¹T. J. Mehrling, C. Benedetti, C. B. Schroeder, E. Esarey, and W. P. Leemans, "Suppression of beam hosing in plasma accelerators with ion motion," *Phys. Rev. Lett.* **121**, 264802 (2018).
- ³²A. Martinez de la Ossa, T. J. Mehrling, and J. Osterhoff, "Intrinsic stabilization of the drive beam in plasma-wakefield accelerators," *Phys. Rev. Lett.* **121**, 064803 (2018).
- ³³T. J. Mehrling, C. Benedetti, C. B. Schroeder, A. Martinez de la Ossa, J. Osterhoff, E. Esarey, and W. P. Leemans, "Accurate modeling of the hose instability in plasma wakefield accelerators," *Phys. Plasmas* **25**, 056703 (2018).
- ³⁴W. Lu, C. Huang, M. Zhou, W. B. Mori, and T. Katsouleas, "Nonlinear theory for relativistic plasma wakefields in the blowout regime," *Phys. Rev. Lett.* **96**, 165002 (2006).
- ³⁵S. A. Yi, V. Khudik, C. Siemon, and G. Shvets, "Analytic model of electromagnetic fields around a plasma bubble in the blow-out regime," *Phys. Plasmas* **20**, 013108 (2013).
- ³⁶A. A. Golovanov, I. Y. Kostyukov, J. Thomas, and A. Pukhov, "Analytic model for electromagnetic fields in the bubble regime of plasma wakefield in non-uniform plasmas," *Phys. Plasmas* **24**, 103104 (2017).
- ³⁷T. N. Dalichaouch, X. L. Xu, A. Tableman, F. Li, F. S. Tsung, and W. B. Mori, "A multi-sheath model for highly nonlinear plasma wakefields," *Phys. Plasmas* **28**, 063103 (2021).
- ³⁸S. Diederichs, C. Benedetti, A. Huebl, R. Lehe, A. Myers, A. Sinn, J.-L. Vay, W. Zhang, and M. Thévenet, "HiPACE++: a portable, 3D quasi-static particle-in-cell code," *Comput. Phys.* (in press); [arXiv:2109.10277](https://arxiv.org/abs/2109.10277).
- ³⁹J. B. Rosenzweig, A. M. Cook, A. Scott, M. C. Thompson, and R. B. Yoder, "Effects of ion motion in intense beam-driven plasma wakefield accelerators," *Phys. Rev. Lett.* **95**, 195002 (2005).
- ⁴⁰C. Benedetti, C. B. Schroeder, E. Esarey, and W. P. Leemans, "Emittance preservation in plasma-based accelerators with ion motion," *Phys. Rev. Accel. Beams* **20**, 111301 (2017).
- ⁴¹C. Benedetti, T. J. Mehrling, C. B. Schroeder, C. G. R. Geddes, and E. Esarey, "Adiabatic matching of particle bunches in a plasma-based accelerator in the presence of ion motion," *Phys. Plasmas* **28**, 053102 (2021).
- ⁴²R. Lehe, C. B. Schroeder, J.-L. Vay, E. Esarey, and W. P. Leemans, "Saturation of the hosing instability in quasilinear plasma accelerators," *Phys. Rev. Lett.* **119**, 244801 (2017).
- ⁴³S. Diederichs, C. Benedetti, M. Thévenet, E. Esarey, and C. B. Schroeder, "Self-stabilizing positron acceleration in a plasma column," (submitted).

Self-stabilizing positron acceleration in a plasma columnS. Diederichs^{1,2,3,*}, C. Benedetti², M. Thévenet¹, E. Esarey²,
J. Osterhoff¹ and C. B. Schroeder^{2,4}¹Deutsches Elektronen-Synchrotron DESY, Notkestraße 85, 22607 Hamburg, Germany²Lawrence Berkeley National Laboratory, 1 Cyclotron Road, Berkeley, California 94720, USA³Institute of Experimental Physics, University of Hamburg,

Luruper Chaussee 149, 22607 Hamburg, Germany

⁴Department of Nuclear Engineering, University of California, Berkeley, California 94720, USA (Received 22 October 2021; revised 23 June 2022; accepted 29 August 2022; published 28 September 2022)

Plasma accelerators sustain extreme field gradients and potentially enable future compact linear colliders. Although tremendous progress has been achieved in accelerating electron beams in a plasma accelerator, positron acceleration with collider-relevant parameters is challenging. A recently proposed positron acceleration scheme relying on the wake generated by an electron drive beam in a plasma column has been shown to be able to accelerate positron witness beams with low emittance and low energy spread. However, since this scheme relies on cylindrical symmetry, it is possibly prone to transverse instabilities that could lead, ultimately, to beam breakup. In this article, we show that the witness beam itself is subject to various damping mechanisms and, therefore, this positron acceleration scheme is inherently stable toward the misalignment of the drive and witness beams. This enables stable, high-quality plasma-based positron acceleration.

DOI: 10.1103/PhysRevAccelBeams.25.091304

I. INTRODUCTION

The ability to build the next, TeV-class linear particle collider is severely constrained due to its high construction costs and power consumption. Providing extreme accelerating gradients, plasma-based accelerators potentially enable compact linear colliders, promising drastic cost reductions [1,2]. In a plasma-based accelerator, either an intense laser pulse [3] or an ultrarelativistic, high-charge particle bunch [4] drives a plasma wake, which can sustain accelerating gradients on the order of tens to hundreds of GV/m. Although high-energy-gain [5,6], high-efficiency [7], and low-energy-spread [8] electron acceleration was demonstrated experimentally, positron acceleration with collider-relevant parameters is significantly more challenging.

Several concepts have been proposed, including utilizing positron drive beams [9], use of hollow-core drive beams [10] or laser pulses [11], a combination of particle and laser drivers [12], or use of the rear portion of a blowout bubble wake [13]. Unfortunately, these concepts lack low emittance, low energy spread, or reasonable efficiency.

*severin.diederichs@desy.de

Published by the American Physical Society under the terms of the *Creative Commons Attribution 4.0 International license*. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

Hollow core plasma channels have been a promising candidate [14,15], but they suffer from instabilities due to the absence of focusing fields for the drive beam [14,16]. Using asymmetric drive beams provides stability in at least one transverse direction [17], but only positron beams with large beam emittances ($>50 \mu\text{m rad}$) could be accelerated, which is too large for collider applications. It was found that the wake generated in a thin, warm, quasi-hollow plasma channel provides accelerating fields for positrons while being robust against instabilities [18]; however, this scheme was demonstrated for a positron beam with several $\mu\text{m rad}$ emittance and several percent relative energy spread only, i.e., beam quality too poor for a collider.

In Ref. [19], a concept was proposed that utilizes an electron drive beam and a plasma column to generate positron-accelerating and focusing wakefield structures and has shown sub- $\mu\text{m rad}$ emittance and sub-percent energy spread positron acceleration [19,20]. Since the scheme relies on cylindrical symmetry, one might expect it to be prone to beam breakup instabilities similar to the ones affecting the hollow core plasma channel. In a recent study [21], the electron driver was found to propagate stably in a plasma column for initial misalignments smaller than the beam size. Still, neither the effect of a misaligned drive beam on the witness beam nor the stability of a misaligned witness beam itself has been investigated so far.

In this article, we demonstrate by means of theory and particle-in-cell (PIC) simulations the stability of a positron

witness beam in the plasma column configuration in the case of a misaligned drive beam or when the witness beam itself is misaligned, and we discuss the corresponding witness beam quality deterioration in presence of such asymmetries. We show that the witness beam is subject to various damping mechanisms and, therefore, this positron acceleration scheme is inherently stable with respect to the misalignment of both the drive and witness beams. These results pave the path for a stable acceleration of low-emittance, low-energy-spread positron beams, a critical step toward the realization of a plasma-based electron-positron collider.

The article is organized as follows: In Sec. II, we recapitulate positron acceleration in a plasma column. In Sec. III, the stability of the positron beam in the presence of initial misalignments is investigated. In Sec. IV, the effect of initial misalignments on the positron beam quality (e.g., deterioration of emittance and energy spread) is determined. Sec. V concludes this work.

II. POSITRON ACCELERATION IN A PLASMA COLUMN

As was first discussed in Ref. [19], if an electron beam drives a wake in a plasma column with a column radius smaller than the blowout radius, the transverse wakefields are reduced outside the column due to the lack of ions. The reduced focusing fields induce a spread in the background plasma electron trajectories moving near the boundary of the blowout wake, which return to the propagation axis over an elongated longitudinal region, forming a high-density electron filament at some distance behind the drive beam. The filament creates a wakefield suitable for acceleration and transport of a trailing witness positron bunch, as shown in the left column (a)–(c) of Fig. 1, where we plot the plasma charge density (a), normalized to en_0 , where e is the elementary charge and n_0 is the ambient plasma density, respectively, and the accelerating (b) and focusing (c) fields, normalized to the cold, nonrelativistic wave-breaking field $E_0 = m_e c^2 k_p / e$, in the x - ζ -plane. Hereby, $k_p = (4\pi n_0 e^2 / m_e c^2)^{1/2}$ is the plasma wave number, e is the electron charge, m_e is the electron mass, c is the speed of light in vacuum, x is the transverse coordinate, $\zeta = z - ct$ is the longitudinal comoving variable, with z and t being the longitudinal coordinate and the time, respectively. We use a bi-Gaussian electron drive beam with a peak current $I_d/I_A = 1$, and root-mean-square (rms) sizes $\sigma_{z,d} = 1.41 k_p^{-1}$ (longitudinal) and $\sigma_{x,d} = \sigma_{y,d} = 0.05 k_p^{-1}$ (transverse), with $I_A = m_e c^3 / e \simeq 17$ kA being the Alfvén current. The plasma column has a radius $R_p = 2.5 k_p^{-1}$. For simplicity, a column with a transverse steplike profile is used, however, the scheme also works for realistic, smooth profiles like the ones obtainable via optical field ionization [22–26].

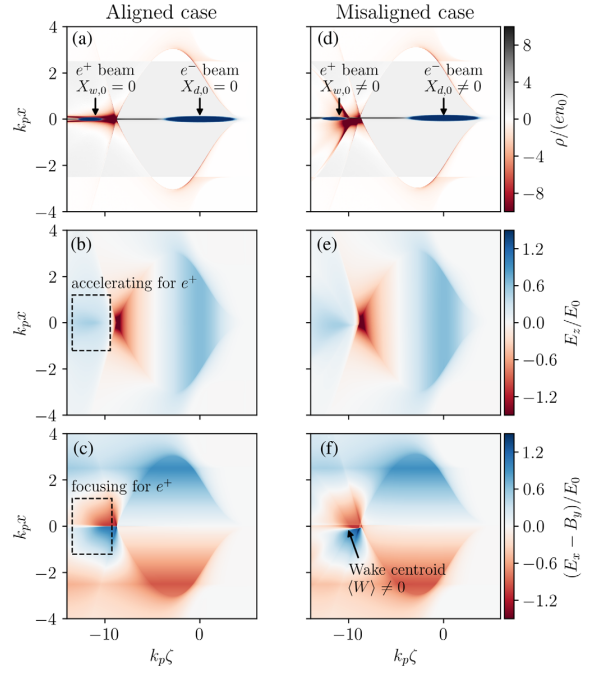


FIG. 1. Wakefield generated in a plasma column by an aligned (left column) and misaligned (right column) drive beam. The normalized plasma charge density, the accelerating field, and the focusing field in the x - ζ -plane are shown for the case with aligned drive and witness beams (a)–(c) and for misaligned drive and witness beams with offsets of $X_{d,0} = X_{w,0} = 0.05 k_p^{-1}$ (d)–(f). As shown in (f), an offset of the drive beam of $X_{d,0} > 0$ results in a wake centroid of $\langle W \rangle < 0$.

In the presented configuration, a helium plasma ionized to the first level is assumed within the column radius and neutral helium gas outside of the column radius. Helium is optimal owing to its high ionization threshold. The usable density range for the plasma column scheme is limited by the wakefield-induced ionization at the boundary of the column. Since the wakefield amplitude scales as $E \sim E_0 \propto n_0^{1/2}$, if the density is too high, the wakefield within the blowout radius (which exceeds the column radius) ionizes the neutral gas outside of the column, causing the column to expand and perturb the positron accelerating and focusing fields. Additionally, the space charge fields of the drive beam can further ionize the plasma in the column to a higher ionization level and disturb the wakefield. For the above drive beam parameters, the maximum densities in a helium plasma and hydrogen plasma are $\simeq 1 \times 10^{17} \text{ cm}^{-3}$ and $\simeq 5 \times 10^{16} \text{ cm}^{-3}$, respectively.

A transverse displacement of the drive beam centroid by one rms size, $X_{d,0} = \sigma_{x,d}$, modifies the electron trajectories asymmetrically, perturbing the wake centroid $\langle W \rangle$ at the back of the blowout. The effect of a misaligned drive beam

is shown in the right column (d)–(f) of Fig. 1. The stable propagation of a misaligned electron drive beam has been shown in a recent study [21]. There, the transversely displaced drive beam was found to be attracted to the center of the plasma column. Although an initial misalignment was found to seed the hosing instability [27], the induced oscillation is damped by various well-known mechanisms such as ion motion [28], energy spread [29], and others [30,31]. Thus, transversely displaced drive beams undergo damped oscillations and drift toward the plasma column axis. For the drive beam parameters presented in this work, and assuming a background density of $1 \times 10^{17} \text{ cm}^{-3}$, a positional alignment at the μm level and an angular alignment on the $\lesssim 10 \mu\text{rad}$ -level are required for stable drive beam propagation in meter-scale plasma columns.

The effect of a misaligned drive beam on the positron witness bunch is discussed in the next sections.

III. STABILITY OF THE WITNESS BEAM

Based on the premise that the drive beam is stable, we now investigate the witness beam stability and the witness beam quality degradation in the presence of a misaligned drive or witness beam using 3D PIC simulations with the quasistatic code HiPACE++ [32]. In the simulations, the computational domain is $(-16, 16) \times (-16, 16) \times (-14, 6) k_p^{-3}$ in $x \times y \times \zeta$, where x and y are the transverse coordinates. The mesh resolution is $0.0078 \times 0.0078 \times 0.001 k_p^{-3}$. The same drive beam and helium plasma column parameters, as in Sec. II, are used. The helium plasma is modeled by electrons and ions with 16 macroparticles per cell each. The drive beam has an initial energy of 5 GeV, no energy spread, and a normalized emittance of $\epsilon_0 = 0.18 k_p^{-1}$, which is matched to the focusing field provided by the background ions in the column. The drive beam is sampled with 10^7 macroparticles. The simulation is propagated for 1000 time steps with a constant time step of $\Delta_t = 5(c k_p)^{-1}$.

The witness beam has a transversely Gaussian and a longitudinally tailored profile. To preclude hosing suppression due to the presence of a slice-dependent energy chirp along the positron beam [29] (the energy spread required to suppress the instability is several percent, and this might be incompatible with collider applications), we use a witness beam current profile that optimally loads the wake, minimizing the development of a correlated energy spread. As shown in Ref. [20], the current profile that flattens the accelerating wakefield by optimal beamloading is nontrivial for the nonlinear positron-accelerating wake. To approximately flatten the wakefield, we consider a current profile that captures the salient features of the one described in Ref. [20]. This is obtained by linearly combining two Gaussian distributions centered in $\zeta_{0,w1} = -11.57 k_p^{-1}$ and $\zeta_{0,w2} = -11.3 k_p^{-1}$, and with a length of

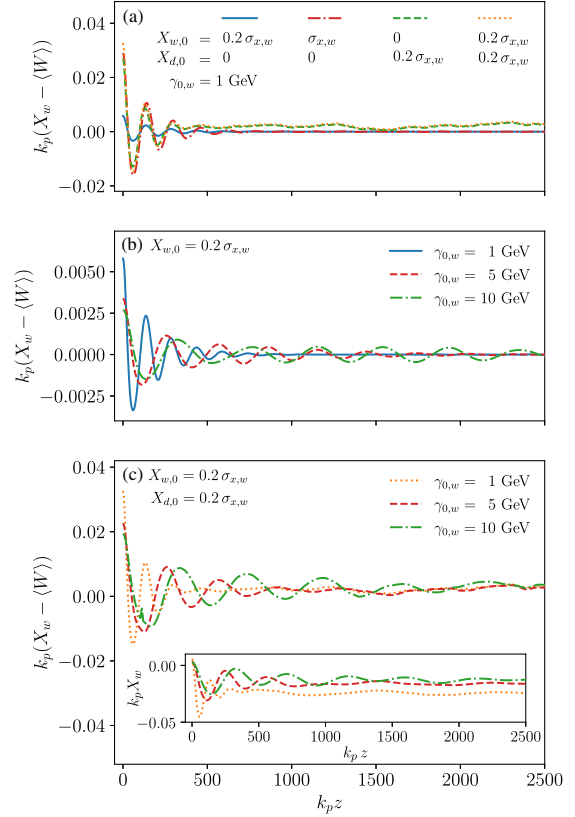


FIG. 2. (a) Difference between witness beam centroid X_w and focusing wake centroid $\langle W \rangle$ for different initial witness and/or drive beam offsets as a function of the propagation distance. The witness beam centroid quickly converges to the wake centroid within a few damped oscillations, demonstrating the stability of the scheme versus initial offsets. The cases of a misaligned driver (b) and both a misaligned driver and misaligned witness beam (c) are also shown for higher witness energies of 5 and 10 GeV, respectively. Inset in (c): X_w along the propagation axis. Despite an increased damping length, the evolution is still stable, i.e., X_w converges to $\langle W \rangle$.

$\sigma_{z,w1} = 0.5 k_p^{-1}$ and $\sigma_{z,w2} = 0.2 k_p^{-1}$, respectively. The corresponding peak densities are $n_{w1}/n_0 = 260$ and $n_{w2}/n_0 = 44.6$. The transverse size of the witness bunch is $\sigma_{x,w} = \sigma_{y,w} = 0.029 k_p^{-1}$, its initial energy is 1 GeV, it has no initial energy spread, and the initial normalized emittance is $\epsilon_x = \epsilon_y = 0.1 k_p^{-1}$ (corresponding to $0.75 \mu\text{mrad}$ for $n_0 = 5 \times 10^{17} \text{ cm}^{-3}$). The witness bunch is sampled with 1.25×10^8 macroparticles.

After presenting the basic numerical setting, we now analyze the evolution of the witness beam in the case of various misalignments of the witness beam centroid, the drive beam centroid, or both with respect to the plasma column axis. The evolution of the witness bunch centroid X_w is shown in Fig. 2(a), where we plot the difference

between the witness beam centroid X_w and the wakefield centroid, which we define as the zero crossing of the transverse focusing field, at the phase location of the witness bunch (W). In the next paragraphs, different beam configurations are reviewed in detail.

First, the setup of an *on-axis drive beam* and a *misaligned witness beam* is considered. In this configuration, a small initial transverse witness beam offset ($X_{w,0} = 0.2\sigma_{x,w}$, solid blue line) and a larger witness beam offset ($X_{w,0} = \sigma_{x,w}$, dash-dotted red line) are tested. For both offsets, the witness beam centroid quickly converges to the wake centroid via a strongly damped oscillation. For the small initial transverse beam offset, we also test higher initial witness beam energies, namely 5 GeV (dashed red line) and 10 GeV (dash-dotted green line) as shown in Fig. 2(b). To keep the witness beams quasimatched [19], their transverse rms sizes are reduced $\sigma_{x,w} = \sigma_{y,w} = 0.017 k_p^{-1}$ and $\sigma_{x,w} = \sigma_{y,w} = 0.0135 k_p^{-1}$ for 5 GeV and 10 GeV, respectively. The initial offset is kept constant relative to the reduced beam size with $X_{w,0} = 0.2\sigma_{x,w}$. As shown in Fig. 2(b), a higher witness beam energy results in a longer damping length.

Second, the setup of a *misaligned drive beam* and an *on-axis witness beam* is considered. For initial displacements of $X_{d,0} = 0.2\sigma_{x,w}$ and $X_{w,0} = 0$ [dashed green line in Fig. 2(a)], the witness beam centroid again quickly converges to the wake centroid in a damped oscillation. Small deviations of the witness beam centroid from the wake centroid are caused by small asymmetries of the wakefield around the zero crossing so that the bunch equilibrium distribution in that wake is itself slightly asymmetric, implying a small offset of the centroid. Notably, the witness beam centroid stays aligned with the wake centroid, even if the wake centroid is changing due to the evolution of the drive beam.

Finally, the same behavior is also observed in the most challenging setup of a *misaligned drive beam* and a *misaligned witness beam*, where both the drive and the witness beams are misaligned by $X_{d,0} = 0.2\sigma_{x,w}$ and $X_{w,0} = 0.2\sigma_{x,w}$, respectively [dotted orange line in Fig. 2(a)]. Hereby, we chose the most challenging case, where both beams are displaced in the same direction, which increases the offset between X_w and $\langle W \rangle$, see Fig. 1. Due to its lower energy, the witness beam adjusts to the (slowly) evolving wake centroid driven by the 5-GeV drive beam. Thus, we also tested quasimatched, higher witness beam energies of 5 GeV [dashed red line in Fig. 2(c)] and 10 GeV [dash-dotted green line in Fig. 2(c)] to verify the stability even in cases where the witness beam is evolving at a slower rate compared to the drive beam. Despite the increased damping length due to the higher witness beam energy, the witness is still able to follow the wake centroid. The inset in (c) shows the corresponding witness beam centroids X_w along the propagation distance without the subtraction of the wake centroid. In the presented cases, the witness beam centroids are displaced from the column axis. Since the attraction of

the drive beam toward the column center is rather slow (see Ref. [21]), the initial drive beam offset is not fully corrected during the propagation. The resulting nonzero wake centroid leads to a displacement of the witness beam centroids.

We also considered cases where the drive and witness beams are offset in orthogonal directions (results not shown), but we did not observe any detrimental coupling between the motion in the two planes.

In all presented cases, the oscillation of the positron witness beam centroid is quickly damped and the witness centroid converges to the wake centroid. The damping is caused by two effects that result in high stability of the witness beam toward transverse displacements: first, the oscillation is damped due to the longitudinally varying focusing fields (see Fig. 1), which has also been observed in other positron acceleration schemes [18]. This effect is similar to the Balakin-Novokhatsky-Smirnov (BNS) damping mechanism [33] in conventional accelerators and has been discussed in the context of quasilinear wakefields in Ref. [31]. Second, the nonlinear transverse wakefields cause phase mixing within a single slice of the beam that ultimately leads to a damping of the witness beam centroid motion.

The scaling of the damping due to the nonlinear transverse wakefields can be obtained from a simplified theoretical model, described in Appendix A. In the simplified model (see also Ref. [19]), we consider a flat beam (i.e., $\sigma_x \gg \sigma_y$) in a steplike confining wakefield, namely $(E_x - B_y)/E_0 = -\alpha \text{sgn}(x)$, where x is the transverse coordinate, α is the strength of the wakefield, and $\text{sgn}(x)$ is the sign function. The evolution of the drive beam is neglected (i.e., the confining wake is constant), we neglect head-to-tail effects (i.e., we consider damping in a slice at a given longitudinal location along the beam), and we assume the acceleration to be slow (i.e., the particle energy changes significantly over distances longer than the characteristic betatron period). Finally, we assume the centroid displacement is small compared to the characteristic beam size, namely $X_{w,0} \ll \sigma_{x,w}$. With these assumptions, the evolution of the beam centroid X_w is found to be a damped oscillation with a damping length S_{damp} scaling as

$$k_p S_{\text{damp}} \propto \sqrt{\frac{k_p \sigma_{x,w} \gamma}{\alpha}}. \quad (1)$$

Notably, and confirmed by the simulations, the damping length does not depend on the initial offset (at least for small displacements). The scaling of the damping length with γ , α , and $\sigma_{x,w}$ is in good agreement with the simulation results.

Since the damping is caused by phase mixing, the damping length is linked to the spread of the betatron frequencies within the bunch. In the case of a Gaussian betatron wave number distribution, the damping length is proportional to $S_{\text{damp,Gauss}} \propto 1/\sigma_{k_\beta}$, where σ_{k_β} is the rms

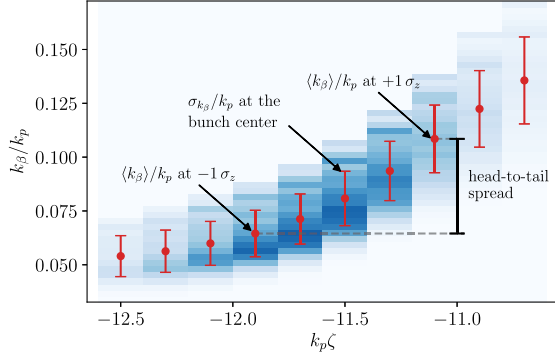


FIG. 3. Distribution of betatron wavenumbers k_β/k_p along the comoving variable ζ . The rms wave number spread per slice σ_{k_β} is depicted by the red bars. The head-to-tail wave number spread, defined as the difference between the mean wave number at the head of the bunch (located at $+1\sigma_z$) and the tail (located at $-1\sigma_z$), is depicted by the black bar. The head-to-tail spread is larger than the average of the intraslice spread.

spread of betatron wave numbers. Although the wave number distribution within the bunch is not Gaussian in this case, their distribution can be used to quantify which mechanism dominates the damping process between (i) BNS damping due to the longitudinally varying focusing strength resulting in a ζ -dependent betatron period for particles in different bunch slices, and (ii) nonlinear focusing wakefield resulting in an amplitude-dependent betatron period for the particles in a beam slice. Therefore, we examine the distribution of betatron frequencies within the witness bunch. Using the example with the misaligned witness beam with an initial offset of $X_{w,0} = 0.2\sigma_{x,w}$, the betatron frequencies are extracted from 1D test particles simulations by employing a second-order particle pusher and using the initial focusing field of the PIC simulation along x at various slices in the beam. Note that it is possible to obtain the betatron frequencies directly from the PIC simulation, but it is impractical due to the acceleration of the particles and the required data intensity as a short output period is needed for reasonable accuracy.

The distribution of the wave numbers of the test particles, k_β/k_p , along the bunch is shown in Fig. 3. The red dots and bars denote the per-slice average $\langle k_\beta \rangle$ and the per-slice rms spread σ_{k_β} of the betatron wave number, respectively. We define the head and the tail of the bunch as the slices at $\zeta = \sigma_z$ and $\zeta = -\sigma_z$. Then, the head-to-tail wave number spread is obtained as the difference between the mean betatron wave numbers at the head and at the tail and is visualized in the plot as a black bar. As shown, the head-to-tail wave number spread, which is caused by the longitudinal variation of the focusing field, is approximately 4 times larger than the average intraslice wave number spread caused by the transverse nonlinearity of the focusing field.

Thus, the variance of the focusing field along the bunch dominates the damping of the positron bunch. Note that the longitudinal variance strongly depends on the beamloading of the witness bunch.

IV. EFFECT OF MISALIGNMENT ON WITNESS BEAM QUALITY

Having discussed the positron stability, we now address the effect of an initial misalignment on the beam quality. For the cases described in Fig. 2(a), the effect of the misalignment on the emittance, energy gain, and relative energy spread of the positron witness beam are shown in Figs. 4(a)–4(c). For a small witness beam offset (solid blue line), the emittance grows only by a few percent in comparison to the aligned case (solid gray line). A large witness beam offset (dash-dotted red line), a drive beam offset (dotted green line), or both (dashed orange line) degrade the emittance due to the larger initial difference between wake centroid and beam centroid. Nevertheless, the emittance growth saturates as soon as the beam centroid converges to the wake centroid. This is also the case for the misaligned drive beam, where the wake centroid is evolving. The energy and the relative energy spread are not sensitive to initial misalignments of the beam centroids, and

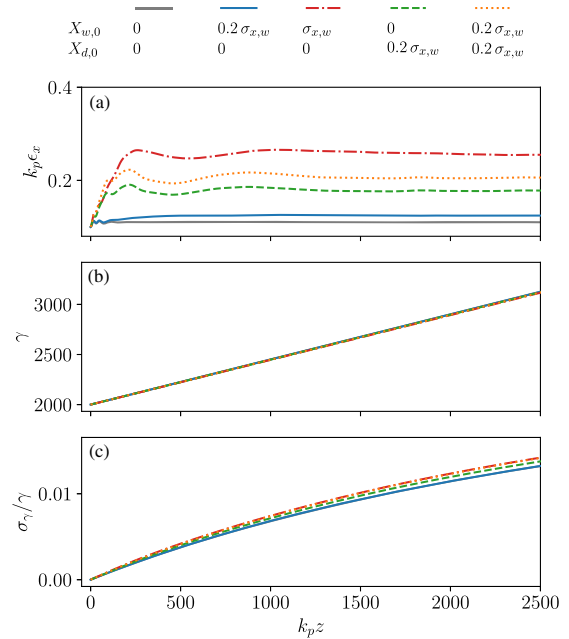


FIG. 4. (a)–(c) Emittance, energy gain, and relative energy spread as a function of the propagation distance. Depending on the offset, emittance growth can be observed, which saturates as soon as the beam is aligned with the wake centroid. Only marginal differences in energy gain and energy spread are observed.

results differ only marginally from the aligned case. Notably, no positrons are lost during the stabilization process in all studied cases.

To gain a further understanding of the effect of an initial offset of the witness beam centroid on the emittance, we use the model for the emittance growth at saturation presented in Ref. [19]. In the model, the same general assumptions as in the previous section are made. Then, assuming a small relative initial offset of the witness beam $\Delta_x = X_{w,0}/\sigma_x \lesssim 1$, the emittance growth at saturation is given by

$$\frac{\sqrt{\langle x^2 \rangle \langle u_x^2 \rangle}}{\sigma_x \sigma_{u_x}} \approx \left\{ \frac{8}{45} \left[\left(1 + \frac{4}{\pi} \right) (1 + \Delta_x^2) + \sqrt{\frac{2}{\pi}} \eta^{-1} (2 + 3\Delta_x^2) + \frac{5}{2} \sqrt{\frac{2}{\pi}} \eta \left(1 + \frac{\Delta_x^2}{2} \right) + \frac{3}{4} \eta^2 \right] \right\}^{1/2}, \quad (2)$$

with

$$\eta = \frac{\sigma_{u_x}^2}{k_p \sigma_x \gamma \alpha}. \quad (3)$$

A detailed derivation of Eq. (2) is presented in Appendix B.

We now compare the emittance growth obtained from the model with results from the PIC simulations. The emittance growth at saturation is estimated using Eq. (2) with an energy of $\gamma = 2000$, an initial transverse beam size of $\sigma_x = 0.029 k_p^{-1}$, an initial emittance of $\epsilon_0 = 0.1 k_p^{-1}$, and a field strength of $\alpha = 0.5$, which is extracted from the PIC simulation with aligned beams with respect to the column axis. For the aligned case with $\Delta_x = 0$, the small offset case with $\Delta_x = 0.2$, and the large offset case with $\Delta_x = 1$, the predicted emittance growth is 6%, 8%, and 50%, respectively. While the emittance growth predicted by the simplified model is in reasonable agreement with the simulation in the aligned and small offset cases, it is significantly underestimated in the large offset case. A reason for the underestimation is that the beamloading effect on the transverse wakefield of a high-charge, strongly misaligned witness beam is not captured by the model. In the case of a large offset, the transverse focusing field is asymmetrically altered, thus the assumption of a simple steplike wake in the model is not fulfilled anymore. Therefore, the model provides a reasonable scaling only for small offsets. Additionally, some of the assumptions (flat beam, neglect of head-to-tail effects) made in the model are not fulfilled in the full PIC simulation, so the model should be used for qualitative purposes only.

In summary, the nonlinearity of the focusing field transversely provides stability as initial offsets are quickly damped due to the phase mixing of the particles. However, this mechanism comes at a cost, since the phase mixing increases the witness beam size and the emittance.

V. CONCLUSION

In this paper, positron acceleration in a plasma column is shown to be inherently stable: in case of misalignments between the drive and witness beams with respect to the plasma column, the witness beam centroid is attracted to the centroid of the focusing wake and the witness beam is not susceptible to the hosing instability due to a longitudinally varying and transversely nonlinear focusing field in the region of the positron bunch. The initial misalignment is corrected via a damped oscillation, which affects the witness beam quality. A scaling for the damping length of the oscillation and the induced emittance growth is presented. By analysis of the betatron wave numbers distribution within the bunch, the variation of the focusing field along the bunch is identified to be the dominating effect in the damping. To achieve high-quality positron acceleration, the alignment of drive and witness beam with respect to the plasma column is essential, but, in principle, achievable due to the intrinsic stability, which allows for the implementation of active feedback loops. Plasma columns enable stable, low-emittance, low-energy-spread positron acceleration as required for a linear collider. Initial experiments of this configuration may be performed at beam test facilities, such as FACET-II [34].

Data availability.—The input scripts for the PIC simulations used in Figs. 2 and 4 are available online [35]. The data that support the other figures are available upon reasonable request from the corresponding author.

ACKNOWLEDGMENTS

We acknowledge the Funding from the Helmholtz Matter and Technologies Accelerator Research and Development Program. This work was supported by the Director, Office of Science, Office of High Energy Physics, of the U.S. Department of Energy, under Contract No. DE-AC02-05CH11231 and used the computational facilities at the National Energy Research Scientific Computing Center (NERSC). We gratefully acknowledge the Gauss Centre for Supercomputing e.V. [36] for funding this project by providing computing time through the John von Neumann Institute for Computing (NIC) on the GCS Supercomputer JUWELS at Jülich Supercomputing Centre (JSC). This research was supported in part by the Maxwell computational resources operated at Deutsches Elektronen-Synchrotron DESY, Hamburg, Germany.

APPENDIX A: DAMPING LENGTH OF THE POSITRON WITNESS BEAM CENTROID OSCILLATIONS

Here we derive an expression for the characteristic damping length of the positron witness beam centroid oscillations. In the derivation, the following assumptions are made: (1) The transverse wakefield close to the axis is approximately steplike (see Fig. 4 in Ref. [19]) and is

described by $(E_x - B_y)/E_0 = -\alpha \text{sgn}(x)$, with $\text{sgn}(x)$ being the sign function, and the field strength $\alpha > 0$ a constant, respectively. (2) The drive beam evolution is neglected and the transverse wakefield is constant in time. (3) The beam is considered to be flat, i.e., $\sigma_{w,x} \gg \sigma_{w,y}$ and $\sigma_{u_x} \gg \sigma_{u_y}$, with σ_{u_x} and σ_{u_y} being the transverse rms momentum spreads in x and y direction, respectively. (4) The acceleration process is adiabatic (i.e., the single particle action is conserved). (5) A small centroid offset is considered, namely $X_{w,0} \ll \sigma_{x,w}$.

The damping length depends on the spread in betatron frequencies of the particles in the beam and this, in turn, depends on the beam's initial phase-space distribution. Analytical calculations for the damping length are tractable in the case of an initial phase-space distribution that is perfectly matched (in the case without centroid displacement) in the idealized steplike wakefield (i.e., the unperturbed phase-space distribution is a stationary solution of the Vlasov equation for the system). Note that, as explained in Ref. [19], the Gaussian distribution that is used as the initial condition in all the simulations presented in this work is not a stationary solution of the Vlasov equation. Hence, in our calculations, we will consider an exponential distribution (see below for details on its definition). We verified numerically that the damping length for an exponential distribution is comparable to that of a quasi-matched initial Gaussian distribution with the same rms size (the exponential distribution is a good proxy for a Gaussian distribution). Extension of the results to other types of matched initial distributions is straightforward.

The calculation of the beam centroid evolution requires first determining the centroid motion for a displaced beam with a Kapchinskij-Vladimirskij (KV) [37] phase-space distribution (KV beamlet, note that an undisplaced KV beam is also a stationary solution of the Vlasov equation, see Sec. A 1). Then, the centroid evolution for the whole beam is obtained by considering that a generic matched beam can be decomposed into a sum of KV beamlets (see Sec. A 2).

1. Evolution of the beam centroid of a single KV beamlet

The KV phase-space distribution is defined as [37]

$$f_{\text{KV}}(x, u_x) \propto \delta[H(x, u_x) - H_0], \quad (\text{A1})$$

where δ is the Dirac delta distribution, $H(x, u_x)$ is the single particle Hamiltonian for a particle in a steplike wakefield, and it is given by [19]

$$H(x, u_x) = \frac{u_x^2}{2\gamma} + \alpha k_p |x|, \quad (\text{A2})$$

with x , u_x , and γ being the position, the momentum, and the Lorentz factor of the particle, respectively. Finally, H_0 is a

parameter that depends on the initial condition and sets the characteristic size of the beamlet. We also introduce the parameter X_0 , defined as

$$k_p X_0 = \frac{H_0}{\alpha}. \quad (\text{A3})$$

Since the phase-space distribution depends on the coordinates via the Hamiltonian, then it is, by construction, a stationary solution of the Vlasov equation and so the beamlet is matched (i.e., all the beam moments are constant in time). In a KV beamlet, all the particles satisfy, at all the times, the condition $H(x, u_x) = H_0$ or, equivalently,

$$\frac{u_x^2}{2\gamma} + \alpha k_p |x| = \alpha k_p X_0. \quad (\text{A4})$$

Thus, the particle momentum is a function of the position according to

$$u_x = \pm \sqrt{2\gamma \alpha k_p (X_0 - |x|)}. \quad (\text{A5})$$

We see that $|x| \leq X_0$, and so X_0 represents the maximum coordinate for a particle in a KV beamlet. The betatron period L_β of a single particle in the steplike wake is obtained by integrating the equation of motion $dx/ds = u_x/\gamma$, with $u(x)$ given by Eq. (A5), over a closed orbit, yielding

$$L_{\beta,\text{KV}} = 2 \int_{-X_0}^{X_0} \sqrt{\frac{\gamma}{2\alpha k_p (X_0 - |x|)}} dx = 8 \sqrt{\frac{\gamma X_0}{2\alpha k_p}}. \quad (\text{A6})$$

All the particles in a KV beamlet have the same betatron period that depends on the beamlet size, X_0 . The betatron wave number is then given by $k_{\beta,\text{KV}} = 2\pi/L_{\beta,\text{KV}}$, or

$$k_{\beta,\text{KV}}(X_0) \equiv k_\beta^0 = \frac{\pi}{4} \sqrt{\frac{2\alpha k_p}{\gamma X_0}}. \quad (\text{A7})$$

We now consider a KV beamlet whose centroid is initially displaced by $X_{w,0}$ (in the following, we assume $X_{w,0} > 0$ for illustration purposes). The particles in such beamlet satisfy

$$\frac{u_x^2}{2\gamma} + \alpha k_p |x - X_{w,0}| = \alpha k_p X_0. \quad (\text{A8})$$

Notably, a displaced KV beamlet is not an equilibrium distribution anymore, as shown in Fig. 5, which displays the phase space of the unperturbed KV beam (gray solid line) and that of the displaced KV beamlet (blue-red-black solid line). The particles in the displaced beamlet can be categorized into three subsections: first, the particles with $x < 0$ (solid red line, we denote this set of particles with the symbol C^-) that satisfy

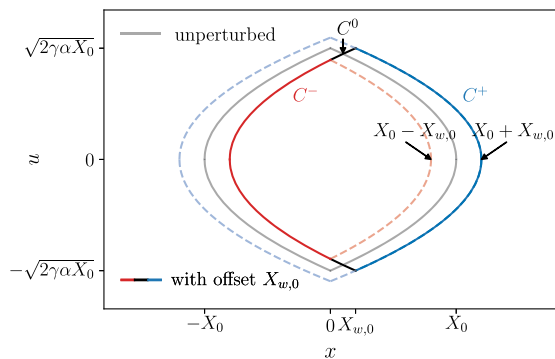


FIG. 5. Phase-space of an unperturbed (solid gray line) KV beamlet and of a KV beamlet with an initial offset $X_{w,0} > 0$ (solid blue-red-black line). The KV beamlet with an initial offset is not an equilibrium distribution anymore: the particles that are initially on the blue (C^+) and red (C^-) arch continue their trajectories on the blue and red dashed lines, respectively. Thus, they are effectively part of KV beamlets with sizes $X_0 + X_{w,0}$ (blue) and $X_0 - X_{w,0}$ (red), respectively. The particles that are initially on the black arches (C^0) fill trajectories in between the red and blue orbits (filamentation).

$$\frac{u_x^2}{2\gamma} + \alpha k_p |x| = \alpha k_p (X_0 - X_{w,0}). \quad (\text{A9})$$

Second, the particles with $x > X_{w,0}$ (solid blue line, C^+) that satisfy

$$\frac{u_x^2}{2\gamma} + \alpha k_p |x| = \alpha k_p (X_0 + X_{w,0}), \quad (\text{A10})$$

and third, the remaining particles with $0 < x < X_{w,0}$ (solid black lines, C^0). When comparing Eqs. (A9) and (A10) with Eq. (A4), one can see that the particles in C^- and C^+ are effectively part of KV beamlets with parameters $X_0 - X_{w,0}$ and $X_0 + X_{w,0}$, respectively. Thus, during evolution, the particles initially in C^- and C^+ remain in their respective perturbed KV orbits with sizes $X_0 \mp X_{w,0}$, as indicated by the dashed lines in the corresponding colors in Fig. 5. The particles in C^0 fill the space area between the red and the blue orbits (filamentation), but the fraction of such particles is small for small centroid displacements. Hence, we see that we can represent a displaced KV beamlet as the sum of two half KV beamlets of different characteristic sizes. From Eq. (A7), the betatron wave numbers for the particles in C^\pm are $k_\beta^\pm = k_{\beta,\text{KV}}(X_0 \pm X_{w,0})$, which reduce, for small offsets $X_{w,0}/X_0 \ll 1$, to

$$k_\beta^\pm = k_\beta^0 \left(1 \mp \frac{1}{2} \frac{X_{w,0}}{X_0} \right). \quad (\text{A11})$$

To calculate the evolution of the centroid of the full, displaced KV beamlet $X_{w,\text{KV}}$, the centroids and the fraction

of particles in C^- , C^+ , and C^0 must be computed. Starting with the latter, we define f^- , f^+ , and f^0 as the fraction of particles in C^- , C^+ , and C^0 , respectively. We have $f^- + f^+ + f^0 = 1$. For small offsets, $X_{w,0}/X_0 \ll 1$, we can neglect the fraction of particles in C^0 , and we modify C^+ by adding the particles in $0 < x < X_{w,0}$ along the dashed blue orbit. In this case, we have $f^- + f^+ = 1$. The fraction in C^\pm is obtained by integrating the corresponding phase-space distribution over the suitable domain. For instance, we have

$$f^- = \frac{1}{N} \int_{-\infty}^0 dx \int_{-\infty}^{\infty} du_x \delta \left[\frac{u_x^2}{2\gamma} + \alpha k_p (|x| - X_0 + X_{w,0}) \right], \quad (\text{A12})$$

with the normalization

$$N = \int_{-\infty}^{\infty} dx \int_{-\infty}^{\infty} du_x \delta \left(\frac{u_x^2}{2\gamma} + \alpha k_p |x| - \alpha k_p X_0 \right). \quad (\text{A13})$$

We obtain $f^- = (1/2) \sqrt{1 - X_{w,0}/X_0}$, and assuming $X_{w,0}/X_0 \ll 1$, we have

$$f^- \simeq \frac{1}{2} \left(1 - \frac{1}{2} \frac{X_{w,0}}{X_0} \right), \quad (\text{A14})$$

and so

$$f^+ = \frac{1}{2} \left(1 + \frac{1}{2} \frac{X_{w,0}}{X_0} \right). \quad (\text{A15})$$

The centroid of the KV beamlet at any time is given by

$$X_{w,\text{KV}}(s) = f^- X^-(s) + f^+ X^+(s), \quad (\text{A16})$$

with $X^\pm(s)$ being the centroids of C^\pm , respectively. Note that since particles in C^\pm will rotate clockwise following the respective orbits with betatron periods k_β^\pm , we can assume that the corresponding centroids are performing a harmonic motion with the same betatron periods, hence

$$X^\pm(s) = X^\pm(0) \cos k_\beta^\pm s. \quad (\text{A17})$$

Here $X_w^\pm(0)$ is the initial C^\pm centroid computed by averaging x over the corresponding initial distributions,

$$X^\pm(0) = \pm \frac{2}{3} (X_0 \pm X_{w,0}). \quad (\text{A18})$$

Inserting Eqs. (A14), (A15), (A17), and (A18) into Eq. (A16), we have

$$X_{w,\text{KV}}(s) = -\frac{X_0}{3} \left(1 - \frac{3X_{w,0}}{2X_0}\right) \cos k_\beta^- s + \frac{X_0}{3} \left(1 + \frac{3X_{w,0}}{2X_0}\right) \cos k_\beta^+ s. \quad (\text{A19})$$

Keeping only first order terms in $X_{w,0}/X_0$, the expression for $X_{w,\text{KV}}(s)$ can be further simplified to

$$X_{w,\text{KV}}(s) = \frac{2}{3} X_0 \sin\left(\frac{k_\beta^0 X_{w,0}}{2X_0} s\right) \sin k_\beta^0 s + X_{w,0} \cos k_\beta^0 s. \quad (\text{A20})$$

We see that, as a consequence of the beating between the betatron frequencies k_β^\pm associated with the particles in C^\pm , the centroid of a perturbed KV beamlet performs transverse oscillations with an amplitude that changes periodically in time. The periodicity of the modulation is determined by $X_{w,0}$.

Finally, we can further simplify Eq. (A20) for early times, $(X_{w,0}/X_0)k_\beta^0 s \ll 1$, to

$$X_{w,\text{KV}}(s) \simeq X_{w,0} \left[\cos k_\beta^0 s + \frac{k_\beta^0 s}{3} \sin k_\beta^0 s \right]. \quad (\text{A21})$$

2. Evolution of the centroid for a beam with an exponential phase-space distribution

We now compute the centroid motion for a displaced beam with an exponential phase-space distribution. The exponential phase-space distribution, including the correct normalization (i.e., such that $\iint f(x, u_x) dx du_x = 1$), is defined as

$$f(x, u_x) = \frac{\alpha}{\sqrt{8\pi\gamma h_0^3}} \exp\left[-\frac{H(x, u_x)}{h_0}\right], \quad (\text{A22})$$

where h_0 is a parameter setting the width of the beam. The rms size of the bunch is given by $\sigma_{\text{exp}} \equiv \sqrt{\langle x^2 \rangle} = \sqrt{2}h_0/\alpha$.

A beam with an exponential phase-space distribution can be decomposed into a sum of different KV beamlets of different sizes. From geometric considerations, we have that the fraction of particles in KV beamlets with sizes between X_0 and $X_0 + dX_0$ is given by

$$dN(X_0) = \frac{2^{7/4}}{\sqrt{\pi}\sigma_{\text{exp}}^{3/2}} \sqrt{X_0} \exp\left(-\sqrt{2}\frac{X_0}{\sigma_{\text{exp}}}\right) dX_0. \quad (\text{A23})$$

The motion of the centroid for the whole beam, X_w , in case of an initial centroid displacement is calculated by superimposing the centroids of different KV beamlets, Eq. (A21), taking into account the weighting given by Eq. (A23), we obtain

$$X_w(s) = \int_0^\infty X_{w,\text{KV}}(s; X_0) dN(X_0), \quad (\text{A24})$$

where we explicitly indicated the dependence of $X_{w,\text{KV}}$ on X_0 . After some algebra, we obtain the following expression for the evolution of the whole beam centroid:

$$X_w(s) = X_{w,0} \frac{4}{\pi} \int_0^\infty \left[\cos(\psi_0 t) + \frac{\psi_0 t}{3} \sin(\psi_0 t) \right] \frac{e^{-t^2}}{t^4} dt, \quad (\text{A25})$$

with

$$\psi_0 = \psi_0(s) = \frac{\pi}{4} \left[\frac{2\sqrt{2}\alpha}{k_p \sigma_{\text{exp}} \gamma} \right]^{\frac{1}{2}} k_p s. \quad (\text{A26})$$

For short propagation distances, $k_p s \ll (4\sqrt{2}k_p \sigma_{\text{exp}} \gamma / \alpha)^{1/2} / \pi$ (i.e., when $\psi_0 \ll 1$), the solution Eq. (A25) can be approximated as

$$\frac{X_w(s)}{X_{w,0}} \simeq 1 - \frac{\psi_0^2}{3} = 1 - \frac{\pi^2 \sqrt{2}\alpha}{24\sigma_{\text{exp}} \gamma} k_p s^2. \quad (\text{A27})$$

For longer propagation distances, $\psi_0(s)$ grows, and for large enough values of ψ_0 , the fast oscillating trigonometric functions in Eq. (A25) will cause the integral to vanish. Hence, Eq. (A25) describes the damping of the centroid oscillations. A numerical study of Eq. (A25) shows that this occurs for $\psi_0 \gtrsim 2\pi$. From this, and using Eq. (A26), we derive the following expression for the characteristic distance over which damping of the centroid oscillations occurs:

$$k_p s \gtrsim k_p S_{\text{damp}} = 4 \left[\frac{\sqrt{2}k_p \sigma_{\text{exp}} \gamma}{\alpha} \right]^{\frac{1}{2}}. \quad (\text{A28})$$

Since the damping length was derived for a beam with an exponential phase-space distribution, while all the simulations in this work use an initial Gaussian phase-space distribution, we verified numerically that the damping length agrees with that of a Gaussian beam to a reasonable level. For the comparison, we assumed that $\sigma_{\text{exp}} = \sigma_{x,w}$, with $\sigma_{x,w}$ being the rms size of the Gaussian witness beam. The evolution of the centroid is obtained with test-particles simulations (using a second-order particle pusher) for beams in a steplike wakefield. Results are shown in Fig. 6, where we compare the centroid evolution from test particles simulations for an exponential beam (red dashed line), and for a Gaussian beam (green dash-dotted line) to the numerical solution of Eq. (A25) (blue solid line). The physical parameters used in this example are $\alpha = 0.5$, $\gamma = 2000$, $\sigma_{x,w} = 0.029 k_p^{-1}$, $X_{w,0} = 0.005 k_p^{-1}$, and an emittance of the Gaussian beam of $\epsilon_x = 0.1 k_p^{-1}$. We find




FIG. 6. Damping in a steplike transverse focusing field during propagation time s . The numerical solution of Eq. (A25) agrees well with the reduced modeling of an exponential beam and with an exponential decay assuming the damping length S_{damp} . The damping of a Gaussian beam in the reduced model takes slightly longer but the physical scaling was found to be the same.

that the solution of Eq. (A25) is in excellent agreement with the simulation results for an exponential beam and in good agreement with that of a Gaussian beam. We also see that the envelope of the damped centroid oscillations can be well approximated by the exponential $\pm \exp(-s/S_{\text{damp}})$, with S_{damp} given by Eq. (A28) (gray dotted lines).

Note that in the full PIC simulations presented in this work, many of the assumptions discussed here are only approximately fulfilled (e.g., round beams are used instead of flat beams, the actual confining wakefield is not an exact steplike function, the drive beam is evolving, etc.) and thus we expect the damping length derived here to be only qualitatively correct.

APPENDIX B: EMITTANCE GROWTH AT SATURATION WITH INITIAL OFFSET

Under the assumptions of the model described in Sec. III, the emittance growth at saturation is calculated similarly to Refs. [19,38]. The Hamiltonian for a particle in a steplike transverse wakefield, $(E_x - B_y)/E_0 = -\alpha \text{sgn}(x)$, is given by

$$H(x, u_x) = \frac{u_x^2}{2\gamma} + \alpha k_p |x|. \quad (\text{B1})$$

The maximum of any particle trajectory x_{max} is obtained by setting $u_x = 0$ for a given initial condition (x_0, u_{x_0}) , yielding

$$x_{\text{max}} = H(x_0, u_{x_0})/\alpha. \quad (\text{B2})$$

Using Eqs. (B1) and (B2), the momentum u_x for any given initial condition is

$$u_x = \pm \sqrt{2\gamma\alpha k_p (x_{\text{max}} - |x|)}, \quad (\text{B3})$$

with $|x| \leq x_{\text{max}}$.

Following Ref. [38], the time average of phase-space moments over a closed particle orbit is given by

$$\overline{x^2}(x_0, u_{x_0}) = \left. \frac{\partial_h \mathcal{P}_2}{\partial_h \mathcal{P}_0} \right|_{h=H(x_0, u_{x_0})=\alpha k_p x_{\text{max}}} \quad (\text{B4a})$$

$$\overline{u_x^2}(x_0, u_{x_0}) = \left. \frac{\gamma \mathcal{P}_0}{\partial_h \mathcal{P}_0} \right|_{h=H(x_0, u_{x_0})=\alpha k_p x_{\text{max}}}, \quad (\text{B4b})$$

with

$$\mathcal{P}_\ell(h) = \int_0^{x_{\text{max}}} x^\ell (h - \alpha k_p |x|)^{1/2} dx. \quad (\text{B5})$$

Evaluating Eq. (B4) with the given Hamiltonian yields

$$\overline{x^2}(x_0, u_{x_0}) = \frac{8}{15} x_{\text{max}}^2 \quad (\text{B6a})$$

$$\overline{u_x^2}(x_0, u_{x_0}) = \frac{2}{3} \gamma \alpha k_p x_{\text{max}}. \quad (\text{B6b})$$

The second-order phase-space moments for the beam are obtained by averaging over the initial Gaussian phase-space distribution. In contrast to Ref. [19], the Gaussian distribution used for the averaging has an initial offset of X_w :

$$\langle x^2 \rangle = \frac{1}{2\pi\sigma_x\sigma_{u_x}} \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \overline{x^2}(x_0, u_{x_0}) \times \exp\left(-\frac{(x_0 - X_w)^2}{2\sigma_x^2} - \frac{u_{x_0}^2}{2\sigma_{u_x}^2}\right) dx_0 du_{x_0} \quad (\text{B7a})$$

$$\langle u_x^2 \rangle = \frac{1}{2\pi\sigma_x\sigma_{u_x}} \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \overline{u_x^2}(x_0, u_{x_0}) \times \exp\left(-\frac{(x_0 - X_w)^2}{2\sigma_x^2} - \frac{u_{x_0}^2}{2\sigma_{u_x}^2}\right) dx_0 du_{x_0}. \quad (\text{B7b})$$

Solving the integrals in Eq. (B7) gives the second-order phase-space moments

$$\langle x^2 \rangle = \frac{8}{15} \left\{ \sigma_x^2 + X_w^2 + \frac{\sigma_{u_x}^2}{\gamma\alpha k_p} \left[\sigma_x e^{-\Delta_x^2/2} \sqrt{\frac{2}{\pi}} + X_w \text{Erf}\left(\frac{\Delta_x}{\sqrt{2}}\right) \right] + \frac{3}{4} \frac{\sigma_{u_x}^4}{\gamma^2 \alpha^2} \right\} \quad (\text{B8a})$$

$$\langle u_x^2 \rangle = \frac{1}{3} \left\{ \sigma_{u_x}^2 + 2\gamma\alpha k_p \left[\sigma_x e^{-\Delta_x^2/2} \sqrt{\frac{2}{\pi}} + X_w \text{Erf}\left(\frac{\Delta_x}{\sqrt{2}}\right) \right] \right\}, \quad (\text{B8b})$$

with $\Delta_x = X_w/\sigma_x$ and the error function $\text{Erf}(x) = \frac{2}{\sqrt{\pi}} \int_0^x e^{-t^2} dt$. Finally, the relative emittance growth is given by the emittance at saturation $\epsilon_f = \sqrt{\langle x^2 \rangle \langle u_x^2 \rangle}$ divided by the initial emittance $\epsilon_i = \sigma_x \sigma_{u_x}$:

$$\frac{\sqrt{\langle x^2 \rangle \langle u_x^2 \rangle}}{\sigma_x \sigma_{u_x}} = \left[\frac{8}{45} \left(1 + \frac{4}{\pi} e^{-\Delta_x} + \Delta_x^2 + 4e^{-\Delta_x/2} \sqrt{\frac{2}{\pi}} \Delta_x \text{Erf} \left(\frac{\Delta_x}{\sqrt{2}} \right) + 2\Delta_x \text{Erf} \left(\frac{\Delta_x}{\sqrt{2}} \right)^2 + 2\eta^{-1} \left\{ [1 + \Delta_x^2] \left[\sqrt{\frac{2}{\pi}} e^{-\Delta_x^2/2} + \Delta_x \text{Erf} \left(\frac{\Delta_x}{\sqrt{2}} \right) \right] \right\} + \frac{5}{2} \eta \left[\sqrt{\frac{2}{\pi}} e^{-\Delta_x^2/2} + \Delta_x \text{Erf} \left(\frac{\Delta_x}{\sqrt{2}} \right) + \frac{3}{4} \eta^2 \right] \right]^{1/2}, \quad (\text{B9})$$

with

$$\eta = \frac{\sigma_{u_x}^2}{k_p \sigma_x \gamma \alpha}. \quad (\text{B10})$$

A second-order Taylor expansion for $\Delta_x \ll 1$ yields

$$\frac{\sqrt{\langle x^2 \rangle \langle u_x^2 \rangle}}{\sigma_x \sigma_{u_x}} \approx \left\{ \frac{8}{45} \left[\left(1 + \frac{4}{\pi} \right) (1 + \Delta_x^2) + \sqrt{\frac{2}{\pi}} \eta^{-1} (2 + 3\Delta_x^2) + \frac{5}{2} \sqrt{\frac{2}{\pi}} \eta \left(1 + \frac{\Delta_x^2}{2} \right) + \frac{3}{4} \eta^2 \right] \right\}^{1/2}. \quad (\text{B11})$$

The Taylor expansion is accurate for $\Delta_x \lesssim 1$ and underestimates the emittance growth at saturation by less than 1% and 5% for $\Delta_x = 0.5$ and $\Delta_x = 1$, respectively.

-
- [1] C. B. Schroeder, E. Esarey, C. G. R. Geddes, C. Benedetti, and W. P. Leemans, Physics considerations for laser-plasma linear colliders, *Phys. Rev. ST Accel. Beams* **13**, 101301 (2010).
 - [2] C. Joshi, S. Corde, and W. Mori, Perspectives on the generation of electron beams from plasma-based accelerators and their near and long term applications, *Phys. Plasmas* **27**, 070602 (2020).
 - [3] E. Esarey, C. B. Schroeder, and W. P. Leemans, Physics of laser-driven plasma-based electron accelerators, *Rev. Mod. Phys.* **81**, 1229 (2009).
 - [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, *Phys. Rev. Lett.* **54**, 693 (1985).
 - [5] 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, E. Oz, R. H. Siemann, D. Walz, and M. Zhou, Energy doubling of 42 GeV electrons in a metre-scale plasma wakefield accelerator, *Nature (London)* **445**, 741 (2007).
 - [6] A. J. Gonsalves *et al.*, Petawatt Laser Guiding and Electron Beam Acceleration to 8 GeV in a Laser-Heated Capillary Discharge Waveguide, *Phys. Rev. Lett.* **122**, 084801 (2019).
 - [7] M. Litos *et al.*, High-efficiency acceleration of an electron beam in a plasma wakefield accelerator, *Nature (London)* **515**, 92 (2014).
 - [8] C. A. Lindström, J. M. Garland, S. Schröder, L. Boulton, G. Boyle, J. Chappell, R. D'Arcy, P. Gonzalez, A. Knetsch, V. Libov, G. Loisch, A. Martinez de la Ossa, P. Niknejadi, K. Pöder, L. Schaper, B. Schmidt, B. Sheeran, S. Wesch, J. Wood, and J. Osterhoff, Energy-Spread Preservation and High Efficiency in a Plasma-Wakefield Accelerator, *Phys. Rev. Lett.* **126**, 014801 (2021).
 - [9] S. Corde, E. Adli, J. Allen, W. An, C. Clarke, C. Clayton, J. Delahaye, J. Frederico, S. Gessner, S. Green *et al.*, Multi-gigaelectronvolt acceleration of positrons in a self-loaded plasma wakefield, *Nature (London)* **524**, 442 (2015).
 - [10] N. Jain, T. M. Antonsen, Jr., and J. P. Palastro, Positron Acceleration by Plasma Wakefields Driven by a Hollow Electron Beam, *Phys. Rev. Lett.* **115**, 195001 (2015).
 - [11] J. Vieira and J. T. Mendonça, Nonlinear Laser Driven Donut Wakefields for Positron and Electron Acceleration, *Phys. Rev. Lett.* **112**, 215001 (2014).
 - [12] L. Reichwein, A. Pukhov, A. Golovanov, and I. Y. Kostyukov, Positron acceleration via laser-augmented blowouts in two-column plasma structures, *Phys. Rev. E* **105**, 055207 (2022).
 - [13] K. V. Lotov, Acceleration of positrons by electron beam-driven wakefields in a plasma, *Phys. Plasmas* **14**, 023101 (2007).
 - [14] C. B. Schroeder, D. H. Whittum, and J. S. Wurtele, Multi-mode Analysis of the Hollow Plasma Channel Wakefield Accelerator, *Phys. Rev. Lett.* **82**, 1177 (1999).
 - [15] S. Gessner, E. Adli, J. M. Allen, W. An, C. I. Clarke, C. E. Clayton, S. Corde, J. Delahaye, J. Frederico, S. Z. Green *et al.*, Demonstration of a positron beam-driven hollow channel plasma wakefield accelerator, *Nat. Commun.* **7**, 11785 (2016).
 - [16] C. A. Lindström *et al.*, Measurement of Transverse Wakefields Induced by a Misaligned Positron Bunch in a Hollow Channel Plasma Accelerator, *Phys. Rev. Lett.* **120**, 124802 (2018).
 - [17] S. Zhou, J. Hua, W. An, W. B. Mori, C. Joshi, J. Gao, and W. Lu, High Efficiency Uniform Wakefield Acceleration of a Positron Beam Using Stable Asymmetric Mode in a Hollow Channel Plasma, *Phys. Rev. Lett.* **127**, 174801 (2021).
 - [18] T. Silva, L. D. Amorim, M. C. Downer, M. J. Hogan, V. Yakimenko, R. Zgadzaj, and J. Vieira, Stable Positron Acceleration in Thin, Warm, Hollow Plasma Channels, *Phys. Rev. Lett.* **127**, 104801 (2021).
 - [19] S. Diederichs, T. J. Mehrling, C. Benedetti, C. B. Schroeder, A. Knetsch, E. Esarey, and J. Osterhoff, Positron transport and acceleration in beam-driven plasma wakefield accelerators using plasma columns, *Phys. Rev. Accel. Beams* **22**, 081301 (2019).

- [20] S. Diederichs, C. Benedetti, E. Esarey, J. Osterhoff, and C. B. Schroeder, High-quality positron acceleration in beam-driven plasma accelerators, *Phys. Rev. Accel. Beams* **23**, 121301 (2020).
- [21] S. Diederichs, C. Benedetti, E. Esarey, M. Thévenet, J. Osterhoff, and C. B. Schroeder, Stable electron beam propagation in a plasma column, *Phys. Plasmas* **29**, 043101 (2022).
- [22] R. J. Shalloo, C. Arran, L. Comer, J. Holloway, J. Jonnerby, R. Walczak, H. M. Milchberg, and S. M. Hooker, Hydrodynamic optical-field-ionized plasma channels, *Phys. Rev. E* **97**, 053203 (2018).
- [23] R. J. Shalloo, C. Arran, A. Picksley, A. von Boetticher, L. Comer, J. Holloway, G. Hine, J. Jonnerby, H. M. Milchberg, C. Thornton, R. Walczak, and S. M. Hooker, Low-density hydrodynamic optical-field-ionized plasma channels generated with an axicon lens, *Phys. Rev. Accel. Beams* **22**, 041302 (2019).
- [24] B. Miao, L. Feder, J. E. Shrock, A. Goffin, and H. M. Milchberg, Optical Guiding in Meter-Scale Plasma Waveguides, *Phys. Rev. Lett.* **125**, 074801 (2020).
- [25] J. E. Shrock, B. Miao, L. Feder, and H. M. Milchberg, Meter-scale plasma waveguides for multi-GeV laser wakefield acceleration, *Phys. Plasmas* **29**, 073101 (2022).
- [26] A. Alejo, J. Cowley, A. Picksley, R. Walczak, and S. M. Hooker, Demonstration of kilohertz operation of hydrodynamic optical-field-ionized plasma channels, *Phys. Rev. Accel. Beams* **25**, 011301 (2022).
- [27] D. H. Whittum, W. M. Sharp, S. S. Yu, M. Lampe, and G. Joyce, Electron-Hose Instability in the Ion-Focused Regime, *Phys. Rev. Lett.* **67**, 991 (1991).
- [28] T. J. Mehrling, C. Benedetti, C. B. Schroeder, E. Esarey, and W. P. Leemans, Suppression of Beam Hosing in Plasma Accelerators with Ion Motion, *Phys. Rev. Lett.* **121**, 264802 (2018).
- [29] T. J. Mehrling, R. A. Fonseca, A. Martinez de la Ossa, and J. Vieira, Mitigation of the Hose Instability in Plasma-Wakefield Accelerators, *Phys. Rev. Lett.* **118**, 174801 (2017).
- [30] A. Martinez de la Ossa, T. J. Mehrling, and J. Osterhoff, Intrinsic Stabilization of the Drive Beam in Plasma-Wakefield Accelerators, *Phys. Rev. Lett.* **121**, 064803 (2018).
- [31] R. Lehe, C. B. Schroeder, J.-L. Vay, E. Esarey, and W. P. Leemans, Intrinsic Stabilization of the Drive Beam in Plasma-Wakefield Accelerators, *Phys. Rev. Lett.* **119**, 244801 (2017).
- [32] S. Diederichs, C. Benedetti, A. Huebl, R. Lehe, A. Myers, A. Sinn, J.-L. Vay, W. Zhang, and M. Thévenet, HiPACE++: A portable, 3D quasi-static particle-in-cell code, *Comput. Phys. Commun.* **278**, 108421 (2022).
- [33] V. E. Balakin, A. V. Novokhatsky, and V. P. Smirnov, VLEPP: Transverse beam dynamics, in *Proceedings of the 12th International Conference on High-Energy Accelerators, HEACC 1983: Fermilab, Batavia, 1983* [Conf. Proc. C830811, 119 (1983)], <https://inspirehep.net/literature/198113>.
- [34] V. Yakimenko, L. Alsberg, E. Bong, G. Bouchard, C. Clarke, C. Emma, S. Green, C. Hast, M. J. Hogan, J. Seabury, N. Lipkowitz, B. O'Shea, D. Storey, G. White, and G. Yocky, FACET-II facility for advanced accelerator experimental tests, *Phys. Rev. Accel. Beams* **22**, 101301 (2019).
- [35] S. Diederichs, C. Benedetti, E. Esarey, M. Thévenet, J. Osterhoff, and C. B. Schroeder, HiPACE++ input scripts for self-stabilizing positron acceleration in a plasma column, [10.5281/zenodo.6671371](https://zenodo.org/record/6671371) (2022).
- [36] www.gauss-centre.eu.
- [37] I. M. Kapchinskij and V. V. Vladimirskij, Limitations of proton beam current in a strong focusing linear accelerator associated with the beam space charge, *Proceedings of the International Conference on High Energy Accelerators and Instrumentation, HEACC* (CERN, Geneva, 1959) p. 274.
- [38] C. Benedetti, C. B. Schroeder, E. Esarey, and W. P. Leemans, Emittance preservation in plasma-based accelerators with ion motion, *Phys. Rev. ST Accel. Beams* **20**, 111301 (2017).

LIST OF FIGURES

- Figure 2.1 Plasma wake and corresponding wakefields in the linear regime. The electron density perturbation normalized to the background plasma density, the accelerating field and the focusing field normalized to the cold, non-relativistic wave-breaking field E_0 are shown in the x - ζ -plane in (a)–(c), respectively. 17
- Figure 2.2 Plasma wake and corresponding wakefields in the blowout regime. The electron density normalized to the background plasma density, the accelerating field and the focusing field normalized to the cold, non-relativistic wave-breaking field E_0 are shown in the x - ζ -plane in (a)–(c), respectively. 19
- Figure 4.1 Schematic of the quasi-static PIC algorithm. The 3D simulation domain is calculated slice-by-slice in a loop over the longitudinal grid points from the head of the box to its tail. Only the beam particles, a 2D slice of plasma particles, and a few 2D slices of fields are required to determine the wake in the 3D simulation domain. Source: [168] 32
- Figure 5.1 Schematic of positron acceleration in a plasma column. 40
- Figure 5.2 Two-dimensional (ζ, x) map of the accelerating wakefield, E_z/E_0 . Positrons can be accelerated in the region $-14 \lesssim k_p\zeta \lesssim -10$. Inset: transverse dependence of accelerating field in the positron accelerating region at three different longitudinal locations denoted by the dashed ($k_p\zeta = -12.5$), solid ($k_p\zeta = -11.5$), and dotted ($k_p\zeta = -10.5$) lines. The accelerating field falls off for increasing distance from the propagation axis. The transverse gradient of the field decreases further behind the driver. 42

- Figure 5.3 Two-dimensional (ζ, x) map of the focusing wakefield, $(E_x - cB_y)/E_0$. Positrons can be focused in the region $-14 \lesssim k_p\zeta \lesssim -9$, which widely overlaps with the positron accelerating region. Inset: transverse dependence of focusing field at three different longitudinal locations denoted by the dashed ($k_p\zeta = -12.5$), solid ($k_p\zeta = -11.5$), and dotted ($k_p\zeta = -10.5$) lines. The focusing field decays almost linearly for increasing distances from the propagation axis. The field decreases further behind the driver. 42
- Figure 5.4 Different OFI-like density profiles and their respective accelerating field 45
- Figure 5.5 Longitudinal and transverse wakefield for various plasma temperatures 46
- Figure 5.6 Longitudinal wakefield, E_z/E_0 , versus co-moving variable, $k_p\zeta$, for various plasma densities 48
- Figure 6.1 Optimally beam-loaded wake. (a) Two-dimensional (ζ, x) map of the plasma electron density n_e/n_0 (blue) and the beam distributions (red). The drive and witness beam current profiles I_b/I_A are denoted by the red line. The trapezoidal witness beam current flattens the accelerating wakefield E_z/E_0 along the co-moving variable ζ as shown in (b). Thus, the wake is optimally loaded and the witness beam does not obtain a correlated energy spread under the assumption that the wakefield stays constant during acceleration. 50
- Figure 6.2 Visualization of the optimal beam loading algorithm : (a) Ensure that the initial condition $|\langle E_{z,i} \rangle| > |\langle E_{z,\text{head}} \rangle|$ is satisfied, otherwise terminate. (b) By step-wisely increasing the current (red bars) on grid point i , the $\langle E_{z,i} \rangle$ (blue line) decreases until the wake is overloaded (i.e. $|\langle E_{z,i} \rangle| < |\langle E_{z,\text{head}} \rangle|$). The overloading current determines $g_{\parallel,\text{max}}$, the last underloading current $g_{\parallel,\text{min}}$. (c) A modified bisection method using $g_{\parallel,\text{min}}$ and $g_{\parallel,\text{max}}$ is used to determine the optimally loading current on grid point i . (a)-(c) are repeated in a loop over the longitudinal grid points, until the initial condition is violated and the wake cannot be loaded any further 54

- Figure 6.3 Current profiles I_b/I_A (red) for an optimally loaded wake and the corresponding averaged accelerating field $\langle E_z \rangle/E_0$ (blue) along the co-moving variable ζ . The charge values for the different witness beam profiles are given assuming a background plasma density of $n_0 = 5 \times 10^{17} \text{cm}^{-3}$. 56
- Figure 6.4 Mean energy of each slice vs. longitudinal position in the positron bunch after acceleration. The blue line refers to algorithm flattening $\langle E_z \rangle$ with a longitudinal uniform σ_r . The red line and the green line refer to additionally applying slice-by-slice matching and averaging of the bunch spot size over the acceleration distance, respectively. That way, the correlated energy spread is reduced to the noise level. 57
- Figure 6.5 Relative slice energy spread vs. acceleration distance. Advancing test particles in an approximated step function yields similar energy-spread in comparison with the HiPACE simulation. The results indicate that emittances smaller than $0.1 \mu\text{m}$ induce an energy spread below 0.1%. 58
- Figure 7.1 Characterization of the wake structure of a beam with a transverse offset with respect to the column axis. (a) Snapshot of the plasma charge density in the x - ζ -plane for a beam (shown in blue) with an initial transverse offset of $X_b = 0.05 k_p^{-1}$ with respect to the plasma column center. (b) The resulting ζ -dependent wake centroid X_p (red line) that lays between the beam centroid X_b (blue line) and the center of the column (black line), leading to an attraction of the beam towards the center of the column. (c) Lineout of the beam current profile (blue line) and focusing wakefield (red line) along x at $\zeta = -1 k_p^{-1}$. 65

Figure 7.2 Evolution (a) of the beam centroid at the tail of the beam $X_{b,\text{tail}}$ along the propagation distance for an initial offset of $X_{b,0} = \sigma_x = 0.05 k_p^{-1}$ assuming an immobile ion background (dashed line) and with mobile ions in a Helium plasma (solid line). The inset shows the final slice-dependent beam centroid. Waterfall plot (b) with the evolution of the beam centroid for $X_{b,0} = 0.05 k_p^{-1}$ in Helium vs longitudinal coordinate (horizontal axis) and propagation distance (vertical axis). The beam centroid decreases via a damped oscillation towards the column center. 66

Figure 7.3 (a) Difference between witness beam centroid X_w and focusing wake centroid $\langle W \rangle$ for different initial witness and/or drive beam offsets as a function of the propagation distance. The witness beam centroid quickly converges to the wake centroid within a few damped oscillations, demonstrating the stability of the scheme versus initial offsets. The cases of a misaligned driver (b) and both a misaligned driver and misaligned witness beam (c) are also shown for higher witness energies of 5 and 10 GeV, respectively. Despite an increased damping length, the evolution is still stable. 69

Figure 7.4 Distribution of betatron wavenumbers k_β/k_p along the co-moving variable ζ . The rms wavenumber spread per slice σ_{k_β} is depicted by the red bars. The head-to-tail wavenumber spread, defined as the difference between the mean wavenumber at the head of the bunch (located at $+1\sigma_z$) and the tail (located at $-1\sigma_z$) is depicted by the black bar. The head-to-tail spread is roughly four times larger than the average of the intra-slice spread. 70

Figure 7.5 (a)–(c) Emittance, energy gain, and relative energy spread as a function of the propagation distance. Depending on the offset, emittance growth can be observed, which saturates as soon as the beam is aligned with the wake centroid. Only marginal differences in energy gain and energy spread are observed. 71

ACRONYMS

ADK	Ammosov-Delone-Krainov
BNS	Balakin-Novokhatsky-Smirnov
CLIC	Compact Linear Collider
CPU	Central Processing Unit
ESPP	European Strategy for Particle Physics
FACET	Facility for Accelerator science and Experimental Test beams
FFT	fast Fourier transform
GPU	Graphics Processing Unit
HPC	High-Performance Computing
ILC	International Linear Collider
I/O	input/output
OFI	optical field ionization
PIC	particle-in-cell
RF	radio-frequency
RMS	root mean square
SLAC	Stanford Linear Accelerator Center
SM	Standard Model of particle physics

BIBLIOGRAPHY

- [1] G. Aad et al. “Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC.” In: *Physics Letters B* 716.1 (2012), pp. 1–29. ISSN: 0370-2693. DOI: [10.1016/j.physletb.2012.08.020](https://doi.org/10.1016/j.physletb.2012.08.020).
- [2] Erik Adli, Carl Andreas Lindstrøm, J Allen, CI Clarke, J Frederico, SJ Gessner, SZ Green, MJ Hogan, MD Litos, B O’Shea, et al. “Long-range attraction of an ultrarelativistic electron beam by a column of neutral plasma.” In: *New Journal of Physics* 18.10 (2016), p. 103013. DOI: [10.1088/1367-2630/18/10/103013](https://doi.org/10.1088/1367-2630/18/10/103013).
- [3] C. Adolphsen et al. *European Strategy for Particle Physics - Accelerator R&D Roadmap. European Strategy for Particle Physics – Accelerator R&D Roadmap*. Vol. 1. CERN Yellow Reports: Monographs. 270 pages, 58 figures. Editor: N. Mounet. LDG chair: D. Newbold. Panel chairs: P. Védérine (HFM), S. Bousson (RF), R. Assmann (plasma), D. Schulte (muon), M. Klein (ERL). Panel editors: B. Baudouy (HFM), L. Bottura (HFM), S. Bousson (RF), G. Burt (RF), R. Assmann (plasma), E. Gschwendtner (plasma), R. Ischebeck (plasma), C. Rogers (muon), D. Schulte (muon), M. Klein (ERL). 2022. DOI: [10.23731/CYRM-2022-001](https://doi.org/10.23731/CYRM-2022-001).
- [4] M Aicheler, P Burrows, M Draper, T Garvey, P Lebrun, K Peach, N Phinney, H Schmickler, D Schulte, and N Toge. *A Multi-TeV Linear Collider Based on CLIC Technology: CLIC Conceptual Design Report*. CERN Yellow Reports: Monographs. Geneva: CERN, 2012. DOI: [10.5170/CERN-2012-007](https://doi.org/10.5170/CERN-2012-007).
- [5] M Aicheler, P.N. Burrows, N. Catalan, R. Corsini, M Draper, J. Osborne, D. Schulte, S. Stapnes, and M.J. Stuart. *CERN Yellow Reports: Monographs, Vol 4 (2018): The Compact Linear Collider (CLIC) – Project Implementation Plan*. en. 2019. DOI: [10.23731/CYRM-2018-004](https://doi.org/10.23731/CYRM-2018-004).
- [6] A. Alejo, J. Cowley, A. Picksley, R. Walczak, and S. M. Hooker. “Demonstration of kilohertz operation of hydrodynamic optical-field-ionized plasma channels.” In: *Phys. Rev. Accel. Beams* 25 (1 Jan. 2022), p. 011301. DOI: [10.1103/PhysRevAccelBeams.25.011301](https://doi.org/10.1103/PhysRevAccelBeams.25.011301).
- [7] Maxim V Ammosov, Nikolai B Delone, and Vladimir P Krainov. “Tunnel Ionization Of Complex Atoms And Atomic Ions In Electromagnetic Field.” In: *High Intensity Laser Processes*. Ed. by John A. Alcock. Vol. 0664. International Society for Optics and Photonics. SPIE, 1986, pp. 138–141. DOI: [10.1117/12.938695](https://doi.org/10.1117/12.938695).

- [8] W. An et al. "Strategies for mitigating the ionization-induced beam head erosion problem in an electron-beam-driven plasma wakefield accelerator." In: *Phys. Rev. ST Accel. Beams* 16 (10 Oct. 2013), p. 101301. DOI: [10.1103/PhysRevSTAB.16.101301](https://doi.org/10.1103/PhysRevSTAB.16.101301).
- [9] Weiming An. "Positron Acceleration in the Electron Driven Plasma Wake Field." In: *ALEGRO Positron Acceleration in Plasma Mini-Workshop*. 2018. URL: https://indico.cern.ch/event/702515/contributions/2884253/attachments/1599455/2535232/Positron_Acceleration_in_the_Electron_Driven_Plasma_Wake_Field_anweiming.pdf.
- [10] Weiming An, Wei Lu, Chengkun Huang, Xinlu Xu, Mark J. Hogan, Chan Joshi, and Warren B. Mori. "Ion Motion Induced Emittance Growth of Matched Electron Beams in Plasma Wakefields." In: *Phys. Rev. Lett.* 118 (24 June 2017), p. 244801. DOI: [10.1103/PhysRevLett.118.244801](https://doi.org/10.1103/PhysRevLett.118.244801).
- [11] TD Arber, Keith Bennett, CS Brady, A Lawrence-Douglas, MG Ramsay, NJ Sircombe, P Gillies, RG Evans, Holger Schmitz, AR Bell, et al. "Contemporary particle-in-cell approach to laser-plasma modelling." In: *Plasma Physics and Controlled Fusion* 57.11 (2015), p. 113001. DOI: [10.1088/0741-3335/57/11/113001](https://doi.org/10.1088/0741-3335/57/11/113001).
- [12] Alexander Aryshev et al. *The International Linear Collider: Report to Snowmass 2021*. 2022. DOI: [10.48550/ARXIV.2203.07622](https://doi.org/10.48550/ARXIV.2203.07622).
- [13] R. Assmann et al. "High-gradient Plasma and Laser Accelerators." In: *European Strategy for Particle Physics - Accelerator R&D Roadmap*. CERN Yellow Reports: Monographs. 2022, pp. 91–143. DOI: [10.23731/CYRM-2022-001](https://doi.org/10.23731/CYRM-2022-001).
- [14] V. E. Balakin, A. V. Novokhatsky, and V. P. Smirnov. "VLEPP: TRANSVERSE BEAM DYNAMICS." In: *Conf. Proc. C830811* (1983), pp. 119–120. URL: <https://inspirehep.net/files/7fed12b8dcaaf3525cc37cf50b56a545>.
- [15] A. Bartnik et al. "CBETA: First Multipass Superconducting Linear Accelerator with Energy Recovery." In: *Phys. Rev. Lett.* 125 (4 2020), p. 044803. DOI: [10.1103/PhysRevLett.125.044803](https://doi.org/10.1103/PhysRevLett.125.044803).
- [16] David A. Beckingsale, Jason Burmark, Rich Hornung, Holger Jones, William Killian, Adam J. Kunen, Olga Pearce, Peter Robinson, Brian S. Ryujin, and Thomas RW Scogland. "RAJA: Portable Performance for Large-Scale Scientific Applications." In: *2019 IEEE/ACM International Workshop on Performance, Portability and Productivity in HPC (P3HPC)*. 2019, pp. 71–81. DOI: [10.1109/P3HPC49587.2019.00012](https://doi.org/10.1109/P3HPC49587.2019.00012).

- [17] C. Benedetti, T. J. Mehrling, C. B. Schroeder, C. G. R. Geddes, and E. Esarey. "Adiabatic matching of particle bunches in a plasma-based accelerator in the presence of ion motion." In: *Physics of Plasmas* 28.5 (2021), p. 053102. DOI: [10.1063/5.0043847](https://doi.org/10.1063/5.0043847).
- [18] C. Benedetti, C. B. Schroeder, E. Esarey, and W. P. Leemans. "Emittance preservation in plasma-based accelerators with ion motion." In: *Phys. Rev. Accel. Beams* 20 (11 Nov. 2017), p. 111301. DOI: [10.1103/PhysRevAccelBeams.20.111301](https://doi.org/10.1103/PhysRevAccelBeams.20.111301).
- [19] C. Benedetti, C. B. Schroeder, C. G. R. Geddes, E. Esarey, and W. P. Leemans. "Efficient modeling of laser-plasma accelerator staging experiments using INF&RNO." In: *AIP Conference Proceedings* 1812.1 (2017), p. 050005. DOI: [10.1063/1.4975866](https://doi.org/10.1063/1.4975866).
- [20] C.K. Birdsall and A.B. Langdon. *Plasma physics via computer simulation*. The Adam Hilger series on plasma physics. McGraw-Hill, 1985. ISBN: 9780070053717. DOI: [10.1201/9781315275048](https://doi.org/10.1201/9781315275048).
- [21] B. E. Blue et al. "Plasma-Wakefield Acceleration of an Intense Positron Beam." In: *Phys. Rev. Lett.* 90 (21 May 2003), p. 214801. DOI: [10.1103/PhysRevLett.90.214801](https://doi.org/10.1103/PhysRevLett.90.214801).
- [22] Ian Blumenfeld et al. "Energy doubling of 42 GeV electrons in a metre-scale plasma wakefield accelerator." In: *Nature* 445.7129 (Feb. 2007), pp. 741–744. DOI: [10.1038/nature05538](https://doi.org/10.1038/nature05538).
- [23] Bernardo Bordini, Luca Bottura, Arnaud Devred, Lucio Fisicarelli, Mikko Karppinen, Gijs de Rijk, Lucio Rossi, Frédéric Savary, and Gerard Willering. "Nb₃Sn 11 T Dipole for the High Luminosity LHC (CERN)." In: *Nb₃Sn Accelerator Magnets: Designs, Technologies and Performance*. Ed. by Daniel Schoerling and Alexander V. Zlobin. Cham: Springer International Publishing, 2019, pp. 223–258. ISBN: 978-3-030-16118-7. DOI: [10.1007/978-3-030-16118-7_9](https://doi.org/10.1007/978-3-030-16118-7_9).
- [24] James E. Brau, Rohini M. Godbole, Francois R. Le Diberder, M. A. Thomson, Harry Weerts, Georg Weiglein, James D. Wells, and Hitoshi Yamamoto. *The Physics Case for an e+e- Linear Collider*. 2012. DOI: [10.48550/ARXIV.1210.0202](https://doi.org/10.48550/ARXIV.1210.0202).
- [25] Heiko Burau, Renée Widera, Wolfgang Hönig, Guido Juckeland, Alexander Debus, Thomas Kluge, Ulrich Schramm, Tomas E Cowan, Roland Sauerbrey, and Michael Bussmann. "PICongPU: a fully relativistic particle-in-cell code for a GPU cluster." In: *IEEE Transactions on Plasma Science* 38.10 (2010), pp. 2831–2839. DOI: [10.1109/TPS.2010.2064310](https://doi.org/10.1109/TPS.2010.2064310).
- [26] M. Bussmann et al. "Radiative Signatures of the Relativistic Kelvin-Helmholtz Instability." In: *Proceedings of the International Conference on High Performance Computing, Networking, Storage*

- and Analysis*. SC '13. ACM, 2013, 5:1–5:12. ISBN: 978-1-4503-2378-9. DOI: [10.1145/2503210.2504564](https://doi.org/10.1145/2503210.2504564).
- [27] Joel Butler, Sekhar Chivukula, Andre de Gouvea, Tao Han, Young-Kee Kim, and Priscilla Cushman. *Snowmass DPF Community Planning Exercise*. 2021. URL: <https://snowmass21.org/>.
- [28] A. D. Cahill, J. B. Rosenzweig, V. A. Dolgashev, S. G. Tantawi, and S. Weathersby. “High gradient experiments with X-band cryogenic copper accelerating cavities.” In: *Phys. Rev. Accel. Beams* 21 (10 Oct. 2018), p. 102002. DOI: [10.1103/PhysRevAccelBeams.21.102002](https://doi.org/10.1103/PhysRevAccelBeams.21.102002).
- [29] Allen Caldwell, Konstantin Lotov, Alexander Pukhov, and Frank Simon. “Proton-driven plasma-wakefield acceleration.” In: *Nat Phys* 5.5 (May 2009), pp. 363–367. ISSN: 1745-2473. DOI: [10.1038/nphys1248](https://doi.org/10.1038/nphys1248).
- [30] Lorenzo Casalino et al. “AI-driven multiscale simulations illuminate mechanisms of SARS-CoV-2 spike dynamics.” In: *The International Journal of High Performance Computing Applications* 35.5 (2021), pp. 432–451. DOI: [10.1177/10943420211006452](https://doi.org/10.1177/10943420211006452).
- [31] Alexander W. Chao, Burton Richter, and Chi-Yuan Yao. “Beam emittance growth caused by transverse deflecting fields in a linear accelerator.” In: *Nuclear Instruments and Methods* 178.1 (1980), pp. 1–8. ISSN: 0029-554X. DOI: [10.1016/0029-554X\(80\)90851-4](https://doi.org/10.1016/0029-554X(80)90851-4).
- [32] Alexander Wu Chao, Karl Hubert Mess, Maury Tigner, and Frank Zimmermann. *Handbook of Accelerator Physics and Engineering*. 2nd. WORLD SCIENTIFIC, 2013. DOI: [10.1142/8543](https://doi.org/10.1142/8543).
- [33] S. Chatrchyan, et al. “Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC.” In: *Physics Letters B* 716.1 (2012), pp. 30–61. ISSN: 0370-2693. DOI: [10.1016/j.physletb.2012.08.021](https://doi.org/10.1016/j.physletb.2012.08.021).
- [34] F.F. Chen. *Introduction to Plasma Physics and Controlled Fusion*. Third. Bd. 1. Springer, 2016. ISBN: 9781441932013. DOI: [10.1007/978-3-319-22309-4](https://doi.org/10.1007/978-3-319-22309-4).
- [35] P. Chen, J. M. Dawson, Robert W. Huff, and T. Katsouleas. “Acceleration of Electrons by the Interaction of a Bunched Electron Beam with a Plasma.” In: *Phys. Rev. Lett.* 54 (7 Feb. 1985), pp. 693–696. DOI: [10.1103/PhysRevLett.54.693](https://doi.org/10.1103/PhysRevLett.54.693).
- [36] T. C. Chiou, T. Katsouleas, C. Decker, W. B. Mori, J. S. Wurtele, G. Shvets, and J. J. Su. “Laser wake-field acceleration and optical guiding in a hollow plasma channel.” In: *Physics of Plasmas* 2.1 (1995), pp. 310–318. DOI: [10.1063/1.871107](https://doi.org/10.1063/1.871107).

- [37] Sébastien Corde, E Adli, JM Allen, W An, CI Clarke, CE Clayton, JP Delahaye, J Frederico, S Gessner, SZ Green, et al. "Multi-gigaelectronvolt acceleration of positrons in a self-loaded plasma wakefield." In: *Nature* 524.7566 (2015), p. 442. DOI: [10.1038/nature14890](https://doi.org/10.1038/nature14890).
- [38] R. Courant, K. Friedrichs, and H. Lewy. "Über die partiellen Differenzgleichungen der mathematischen Physik." In: *Mathematische Annalen* 100.1 (1928), pp. 32–74. ISSN: 0025-5831. DOI: [10.1007/BF01448839](https://doi.org/10.1007/BF01448839).
- [39] T. N. Dalichaouch, X. L. Xu, A. Tableman, F. Li, F. S. Tsung, and W. B. Mori. "A multi-sheath model for highly nonlinear plasma wakefields." In: *Physics of Plasmas* 28.6 (2021), p. 063103. DOI: [10.1063/5.0051282](https://doi.org/10.1063/5.0051282).
- [40] J. M. Dawson. "Nonlinear Electron Oscillations in a Cold Plasma." In: *Phys. Rev.* 113 (2 Jan. 1959), pp. 383–387. DOI: [10.1103/PhysRev.113.383](https://doi.org/10.1103/PhysRev.113.383).
- [41] S. Diederichs, C. Benedetti, E. Esarey, J. Osterhoff, and C. B. Schroeder. "High-quality positron acceleration in beam-driven plasma accelerators." In: *Phys. Rev. Accel. Beams* 23 (12 Dec. 2020), p. 121301. DOI: [10.1103/PhysRevAccelBeams.23.121301](https://doi.org/10.1103/PhysRevAccelBeams.23.121301).
- [42] S. Diederichs, C. Benedetti, E. Esarey, M. Thévenet, J. Osterhoff, and C. B. Schroeder. "Stable electron beam propagation in a plasma column." In: *Physics of Plasmas* 29.4 (2022), p. 043101. DOI: [10.1063/5.0087807](https://doi.org/10.1063/5.0087807).
- [43] S. Diederichs, C. Benedetti, E. Esarey, M. Thévenet, A. Sinn, J. Osterhoff, and C. B. Schroeder. "Temperature effects in plasma-based positron acceleration schemes using electron filaments." In: *Physics of Plasmas* 30.7 (July 2023), p. 073104. ISSN: 1070-664X. DOI: [10.1063/5.0155489](https://doi.org/10.1063/5.0155489).
- [44] S. Diederichs, C. Benedetti, A. Huebl, R. Lehe, A. Myers, A. Sinn, J.-L. Vay, W. Zhang, and M. Thévenet. "HiPACE++: A portable, 3D quasi-static particle-in-cell code." In: *Computer Physics Communications* 278 (2022), p. 108421. ISSN: 0010-4655. DOI: [10.1016/j.cpc.2022.108421](https://doi.org/10.1016/j.cpc.2022.108421).
- [45] S. Diederichs, C. Benedetti, M. Thévenet, E. Esarey, J. Osterhoff, and C. B. Schroeder. "Self-stabilizing positron acceleration in a plasma column." In: *Phys. Rev. Accel. Beams* 25 (9 Sept. 2022), p. 091304. DOI: [10.1103/PhysRevAccelBeams.25.091304](https://doi.org/10.1103/PhysRevAccelBeams.25.091304).
- [46] S. Diederichs, T. J. Mehrling, C. Benedetti, C. B. Schroeder, A. Knetsch, E. Esarey, and J. Osterhoff. "Positron transport and acceleration in beam-driven plasma wakefield accelerators using plasma columns." In: *Phys. Rev. Accel. Beams* 22.8 (2019), p. 081301. DOI: [10.1103/PhysRevAccelBeams.22.081301](https://doi.org/10.1103/PhysRevAccelBeams.22.081301).

- [47] A Doche, C Beekman, Sébastien Corde, JM Allen, CI Clarke, J Frederico, SJ Gessner, SZ Green, MJ Hogan, B O’Shea, et al. “Acceleration of a trailing positron bunch in a plasma wakefield accelerator.” In: *Scientific reports* 7.1 (2017), pp. 1–7. DOI: [10.1038/s41598-017-14524-4](https://doi.org/10.1038/s41598-017-14524-4).
- [48] C. G. Durfee and H. M. Milchberg. “Light pipe for high intensity laser pulses.” In: *Phys. Rev. Lett.* 71 (15 Oct. 1993), pp. 2409–2412. DOI: [10.1103/PhysRevLett.71.2409](https://doi.org/10.1103/PhysRevLett.71.2409).
- [49] R D’Arcy, J Chappell, J Beinortaite, S Diederichs, G Boyle, B Foster, MJ Garland, P Gonzalez Caminal, CA Lindstrøm, G Loisch, et al. “Recovery time of a plasma-wakefield accelerator.” In: *Nature* 603.7899 (2022), pp. 58–62. DOI: [10.1038/s41586-021-04348-8](https://doi.org/10.1038/s41586-021-04348-8).
- [50] H. Carter Edwards, Christian R. Trott, and Daniel Sunderland. “Kokkos: Enabling manycore performance portability through polymorphic memory access patterns.” In: *Journal of Parallel and Distributed Computing* 74.12 (2014). Domain-Specific Languages and High-Level Frameworks for High-Performance Computing, pp. 3202–3216. ISSN: 0743-7315. DOI: [10.1016/j.jpdc.2014.07.003](https://doi.org/10.1016/j.jpdc.2014.07.003).
- [51] F. Englert and R. Brout. “Broken Symmetry and the Mass of Gauge Vector Mesons.” In: *Phys. Rev. Lett.* 13 (9 Aug. 1964), pp. 321–323. DOI: [10.1103/PhysRevLett.13.321](https://doi.org/10.1103/PhysRevLett.13.321).
- [52] E. Esarey, P. Catravas, and W. P. Leemans. “Betatron radiation from electron beams in plasma focusing channels.” In: *AIP Conference Proceedings* 569.1 (2001), pp. 473–486. DOI: [10.1063/1.1384377](https://doi.org/10.1063/1.1384377).
- [53] 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](https://doi.org/10.1103/RevModPhys.81.1229).
- [54] B. Feng, C. Huang, V. Decyk, W.B. Mori, P. Muggli, and T. Katsouleas. “Enhancing parallel quasi-static particle-in-cell simulations with a pipelining algorithm.” In: *Journal of Computational Physics* 228.15 (2009), pp. 5340–5348. ISSN: 0021-9991. DOI: [10.1016/j.jcp.2009.04.019](https://doi.org/10.1016/j.jcp.2009.04.019).
- [55] Martin E Fermann, Almantas Galvanauskas, and Gregg Sucha. *Ultrafast lasers: technology and applications*. Vol. 80. CRC Press, 2002. DOI: [10.1201/9780203910207](https://doi.org/10.1201/9780203910207).
- [56] Haohuan Fu et al. “18.9-Pflops Nonlinear Earthquake Simulation on Sunway TaihuLight: Enabling Depiction of 18-Hz and 8-Meter Scenarios.” In: *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analy-*

- sis. SC '17. Association for Computing Machinery, 2017. ISBN: 9781450351140. DOI: [10.1145/3126908.3126910](https://doi.org/10.1145/3126908.3126910).
- [57] Todd Gamblin, Matthew LeGendre, Michael R Collette, Gregory L Lee, Adam Moody, Bronis R de Supinski, and Scott Futral. "The Spack package manager: bringing order to HPC software chaos." In: *SC'15: Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*. IEEE. 2015, pp. 1–12. DOI: [10.1145/2807591.2807623](https://doi.org/10.1145/2807591.2807623).
- [58] Spencer Gessner, Erik Adli, James M Allen, Weiming An, Christine I Clarke, Chris E Clayton, Sebastien Corde, JP Delahaye, Joel Frederico, Selina Z Green, et al. "Demonstration of a positron beam-driven hollow channel plasma wakefield accelerator." In: *Nature communications* 7 (2016), p. 11785. DOI: [10.1038/ncomms11785](https://doi.org/10.1038/ncomms11785).
- [59] L. M. Gorbunov and V. I. Kirsanov. "Excitation of plasma waves by an electromagnetic wave packet." In: *Sov. Phys. JETP* 66 (1987). [*Zh. Eksp. Teor. Fiz.* 93, 509 (1987)], p. 290. URL: http://www.jetp.ras.ru/cgi-bin/dn/e_066_02_0290.pdf.
- [60] SZ Green, E Adli, CI Clarke, Sébastien Corde, SA Edstrom, AS Fisher, J Frederico, JC Frisch, S Gessner, S Gilevich, et al. "Laser ionized preformed plasma at FACET." In: *Plasma Physics and Controlled Fusion* 56.8 (2014), p. 084011. DOI: [10.1088/0741-3335/56/8/084011](https://doi.org/10.1088/0741-3335/56/8/084011).
- [61] European Strategy Group. *2020 Update of the European Strategy for Particle Physics*. Tech. rep. Geneva, 2020. DOI: [10.17181/ESU2020](https://doi.org/10.17181/ESU2020).
- [62] E. Gschwendtner et al. "AWAKE, The Advanced Proton Driven Plasma Wakefield Acceleration Experiment at CERN." In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 829 (2016). 2nd European Advanced Accelerator Concepts Workshop - EAAC 2015, pp. 76–82. ISSN: 0168-9002. DOI: [10.1016/j.nima.2016.02.026](https://doi.org/10.1016/j.nima.2016.02.026).
- [63] Werner Herr and B Muratori. "Concept of luminosity." In: (2006). DOI: [10.5170/CERN-2006-002.361](https://doi.org/10.5170/CERN-2006-002.361).
- [64] Peter W. Higgs. "Broken Symmetries and the Masses of Gauge Bosons." In: *Phys. Rev. Lett.* 13 (16 Oct. 1964), pp. 508–509. DOI: [10.1103/PhysRevLett.13.508](https://doi.org/10.1103/PhysRevLett.13.508).
- [65] R.W. Hockney and J.W. Eastwood. *Computer Simulation Using Particles*. Advanced book program: Addison-Wesley. McGraw-Hill, 1981. ISBN: 9780070291089. DOI: [10.1201/9780367806934](https://doi.org/10.1201/9780367806934).
- [66] M. J. Hogan et al. "Ultrarelativistic-Positron-Beam Transport through Meter-Scale Plasmas." In: *Phys. Rev. Lett.* 90 (20 May 2003), p. 205002. DOI: [10.1103/PhysRevLett.90.205002](https://doi.org/10.1103/PhysRevLett.90.205002).

- [67] Mark J. Hogan. “Electron and Positron Beam–Driven Plasma Acceleration.” In: *Reviews of Accelerator Science and Technology* 09 (2016), pp. 63–83. DOI: [10.1142/S1793626816300036](https://doi.org/10.1142/S1793626816300036).
- [68] C. Huang, V.K. Decyk, C. Ren, M. Zhou, W. Lu, W.B. Mori, J.H. Cooley, T.M. Antonsen Jr., and T. Katsouleas. “QUICKPIC: A highly efficient particle-in-cell code for modeling wakefield acceleration in plasmas.” In: *Journal of Computational Physics* 217.2 (2006), pp. 658–679. ISSN: 0021-9991. DOI: [10.1016/j.jcp.2006.01.039](https://doi.org/10.1016/j.jcp.2006.01.039).
- [69] C. Huang et al. “Hosing Instability in the Blow-Out Regime for Plasma-Wakefield Acceleration.” In: *Phys. Rev. Lett.* 99 (25 Dec. 2007), p. 255001. DOI: [10.1103/PhysRevLett.99.255001](https://doi.org/10.1103/PhysRevLett.99.255001).
- [70] C. S. Hue, G. J. Cao, I. A. Andriyash, A. Knetsch, M. J. Hogan, E. Adli, S. Gessner, and S. Corde. “Efficiency and beam quality for positron acceleration in loaded plasma wakefields.” In: *Phys. Rev. Research* 3 (4 Oct. 2021), p. 043063. DOI: [10.1103/PhysRevResearch.3.043063](https://doi.org/10.1103/PhysRevResearch.3.043063).
- [71] Axel Huebl, Franz Poeschel, Fabian Koller, and Junmin Gu. *openPMD-api: C++ & Python API for Scientific I/O with openPMD*. 2018. DOI: [10.14278/rodare.27](https://doi.org/10.14278/rodare.27).
- [72] Neeraj Jain, TM Antonsen Jr, and JP Palastro. “Positron acceleration by plasma wakefields driven by a hollow electron beam.” In: *Physical Review Letters* 115.19 (Nov. 2015), p. 195001. DOI: [10.1103/PhysRevLett.115.195001](https://doi.org/10.1103/PhysRevLett.115.195001).
- [73] Neeraj Jain, John Palastro, T. M. Antonsen, Warren B. Mori, and Weiming An. “Plasma wakefield acceleration studies using the quasi-static code WAKE.” In: *Physics of Plasmas* 22.2 (2015), p. 023103. DOI: [10.1063/1.4907159](https://doi.org/10.1063/1.4907159).
- [74] S Y Kalmykov, A Englesbe, J Elle, and A Schmitt-Sody. “Single-cycle THz signal accompanying laser wake in photoionized plasmas and plasma channels.” In: *Journal of Physics: Conference Series* 1596.1 (July 2020), p. 012060. DOI: [10.1088/1742-6596/1596/1/012060](https://doi.org/10.1088/1742-6596/1596/1/012060).
- [75] S Wilks T Katsouleas and JDJ Su. “Beam loading efficiency in plasma accelerators.” In: *Part. Accel* 22 (1987), pp. 81–99. URL: https://accelconf.web.cern.ch/p87/PDF/PAC1987_0100.PDF.
- [76] W. D. Kimura, H. M. Milchberg, P. Muggli, X. Li, and W. B. Mori. “Hollow plasma channel for positron plasma wakefield acceleration.” In: *Phys. Rev. ST Accel. Beams* 14 (4 Apr. 2011), p. 041301. DOI: [10.1103/PhysRevSTAB.14.041301](https://doi.org/10.1103/PhysRevSTAB.14.041301).
- [77] Manuel Kirchen et al. “Optimal Beam Loading in a Laser-Plasma Accelerator.” In: *Phys. Rev. Lett.* 126 (17 Apr. 2021), p. 174801. DOI: [10.1103/PhysRevLett.126.174801](https://doi.org/10.1103/PhysRevLett.126.174801).

- [78] A. Knetsch et al. “Stable witness-beam formation in a beam-driven plasma cathode.” In: *Phys. Rev. Accel. Beams* 24 (10 2021), p. 101302. DOI: [10.1103/PhysRevAccelBeams.24.101302](https://doi.org/10.1103/PhysRevAccelBeams.24.101302).
- [79] W. L. Kruer. *The Physics of Laser Plasma Interactions*. Addison-Wesley, Redwood City, 1988. DOI: [10.1201/9781003003243](https://doi.org/10.1201/9781003003243).
- [80] Y. Y. Lau. “Classification of beam breakup instabilities in linear accelerators.” In: *Phys. Rev. Lett.* 63 (11 Sept. 1989), pp. 1141–1144. DOI: [10.1103/PhysRevLett.63.1141](https://doi.org/10.1103/PhysRevLett.63.1141).
- [81] Valeri Lebedev, Alexey Burov, and Sergei Nagaitsev. “Efficiency versus instability in plasma accelerators.” In: *Phys. Rev. Accel. Beams* 20 (12 Dec. 2017), p. 121301. DOI: [10.1103/PhysRevAccelBeams.20.121301](https://doi.org/10.1103/PhysRevAccelBeams.20.121301).
- [82] Seung Lee, Tom Katsouleas, RG Hemker, ES Dodd, and WB Mori. “Plasma-wakefield acceleration of a positron beam.” In: *Phys. Rev. E* 64.4 (Sept. 2001), p. 045501. DOI: [10.1103/PhysRevE.64.045501](https://doi.org/10.1103/PhysRevE.64.045501).
- [83] R. Lehe, C. B. Schroeder, J.-L. Vay, E. Esarey, and W. P. Leemans. “Saturation of the Hosing Instability in Quasilinear Plasma Accelerators.” In: *Phys. Rev. Lett.* 119 (24 Dec. 2017), p. 244801. DOI: [10.1103/PhysRevLett.119.244801](https://doi.org/10.1103/PhysRevLett.119.244801).
- [84] Remi Lehe, Manuel Kirchen, Igor A Andriyash, Brendan B Godfrey, and Jean-Luc Vay. “A spectral, quasi-cylindrical and dispersion-free Particle-In-Cell algorithm.” In: *Computer Physics Communications* 203 (2016), pp. 66–82. DOI: [10.1016/j.cpc.2016.02.007](https://doi.org/10.1016/j.cpc.2016.02.007).
- [85] F. Lemery and P. Piot. “Tailored electron bunches with smooth current profiles for enhanced transformer ratios in beam-driven acceleration.” In: *Phys. Rev. ST Accel. Beams* 18 (8 Aug. 2015), p. 081301. DOI: [10.1103/PhysRevSTAB.18.081301](https://doi.org/10.1103/PhysRevSTAB.18.081301).
- [86] Fei Li, Weiming An, Viktor K. Decyk, Xinlu Xu, Mark J. Hogan, and Warren B. Mori. “A quasi-static particle-in-cell algorithm based on an azimuthal Fourier decomposition for highly efficient simulations of plasma-based acceleration: QPAD.” In: *Computer Physics Communications* 261 (2021), p. 107784. ISSN: 0010-4655. DOI: [10.1016/j.cpc.2020.107784](https://doi.org/10.1016/j.cpc.2020.107784).
- [87] C. A. Lindstrøm et al. “Measurement of Transverse Wakefields Induced by a Misaligned Positron Bunch in a Hollow Channel Plasma Accelerator.” In: *Phys. Rev. Lett.* 120 (12 Mar. 2018), p. 124802. DOI: [10.1103/PhysRevLett.120.124802](https://doi.org/10.1103/PhysRevLett.120.124802).
- [88] C. A. Lindstrøm et al. “Energy-Spread Preservation and High Efficiency in a Plasma-Wakefield Accelerator.” In: *Phys. Rev. Lett.* 126 (1 Jan. 2021), p. 014801. DOI: [10.1103/PhysRevLett.126.014801](https://doi.org/10.1103/PhysRevLett.126.014801).

- [89] C.A. Lindstrøm and M. Thévenet. "Emittance preservation in advanced accelerators." In: *Journal of Instrumentation* 17.05 (May 2022), Po5016. DOI: [10.1088/1748-0221/17/05/p05016](https://doi.org/10.1088/1748-0221/17/05/p05016).
- [90] Vladimir N. Litvinenko, Thomas Roser, and Maria Chamizo-Llatas. "High-energy high-luminosity e+e- collider using energy-recovery linacs." In: *Physics Letters B* 804 (2020), p. 135394. ISSN: 0370-2693. DOI: [10.1016/j.physletb.2020.135394](https://doi.org/10.1016/j.physletb.2020.135394).
- [91] Q. H. Liu. "The PSTD algorithm: A time-domain method requiring only two cells per wavelength." In: *Microwave and Optical Technology Letters* 15.3 (1997), pp. 158–165. DOI: [10.1002/\(SICI\)1098-2760\(19970620\)15:3<158::AID-MOP11>3.0.CO;2-3](https://doi.org/10.1002/(SICI)1098-2760(19970620)15:3<158::AID-MOP11>3.0.CO;2-3).
- [92] Wei-Yuan Liu, Xing-Long Zhu, Min Chen, Su-Ming Weng, Feng He, Zheng-Ming Sheng, and Jie Zhang. "Tail-Wave-Assisted Positron Acceleration in Nonlinear Laser Plasma Wakefields." In: *Phys. Rev. Appl.* 19 (4 2023), p. 044048. DOI: [10.1103/PhysRevApplied.19.044048](https://doi.org/10.1103/PhysRevApplied.19.044048).
- [93] Yong (Alexander) Liu et al. "Closing the "Quantum Supremacy" Gap: Achieving Real-Time Simulation of a Random Quantum Circuit Using a New Sunway Supercomputer." In: *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*. SC '21. Association for Computing Machinery, 2021. ISBN: 9781450384421. DOI: [10.1145/3458817.3487399](https://doi.org/10.1145/3458817.3487399).
- [94] Gregor Loisch, Galina Asova, Prach Boonpornprasert, Reinhard Brinkmann, Ye Chen, Johannes Engel, James Good, Matthias Gross, Florian Grüner, Holger Huck, et al. "Observation of high transformer ratio plasma wakefield acceleration." In: *Physical review letters* 121.6 (2018), p. 064801. DOI: [10.1103/PhysRevLett.121.064801](https://doi.org/10.1103/PhysRevLett.121.064801).
- [95] H. A. Lorentz. "Versuch Einer Theorie der Electricischen und Optischen Erscheinungen in Bewegten Körpern." In: *Collected Papers: Volume V*. Dordrecht: Springer Netherlands, 1937, pp. 1–138. ISBN: 978-94-015-3445-1. DOI: [10.1007/978-94-015-3445-1_1](https://doi.org/10.1007/978-94-015-3445-1_1).
- [96] K. V. Lotov. "Fine wakefield structure in the blowout regime of plasma wakefield accelerators." In: *Phys. Rev. ST Accel. Beams* 6 (6 June 2003), p. 061301. DOI: [10.1103/PhysRevSTAB.6.061301](https://doi.org/10.1103/PhysRevSTAB.6.061301).
- [97] K. V. Lotov. "Efficient operating mode of the plasma wakefield accelerator." In: *Physics of Plasmas (1994-present)* 12.5, 053105 (2005), pp. –. DOI: [10.1063/1.1889444](https://doi.org/10.1063/1.1889444).
- [98] K. V. Lotov. "Acceleration of positrons by electron beam-driven wakefields in a plasma." In: *Physics of Plasmas* 14.2 (2007), p. 023101. DOI: [10.1063/1.2434793](https://doi.org/10.1063/1.2434793).

- [99] K. V. Lotov. “Force exerted on particle bunch propagating near plasma-vacuum boundary.” In: *Plasma Physics and Controlled Fusion* 62.8 (2020), p. 085002. DOI: [10.1088/1361-6587/ab93a4](https://doi.org/10.1088/1361-6587/ab93a4).
- [100] K. V. Lotov and P. V. Tuev. “Plasma wakefield acceleration beyond the dephasing limit with 400 GeV proton driver.” In: *Plasma Physics and Controlled Fusion* 63.12 (Nov. 2021), p. 125027. DOI: [10.1088/1361-6587/ac349a](https://doi.org/10.1088/1361-6587/ac349a).
- [101] W. Lu, C. Huang, M. Zhou, W. B. Mori, and T. Katsouleas. “Nonlinear Theory for Relativistic Plasma Wakefields in the Blowout Regime.” In: *Phys. Rev. Lett.* 96 (16 Apr. 2006), p. 165002. DOI: [10.1103/PhysRevLett.96.165002](https://doi.org/10.1103/PhysRevLett.96.165002).
- [102] Bertrand Martinez, Bernardo Barbosa, and Marija Vranic. “Creation and direct laser acceleration of positrons in a single stage.” In: *Phys. Rev. Accel. Beams* 26 (1 2023), p. 011301. DOI: [10.1103/PhysRevAccelBeams.26.011301](https://doi.org/10.1103/PhysRevAccelBeams.26.011301).
- [103] S van der Meer. *Improving the power efficiency of the plasma wakefield accelerator*. Tech. rep. Geneva: CERN, 1985. URL: <https://cds.cern.ch/record/163918>.
- [104] T. J. Mehrling, C. Benedetti, C. B. Schroeder, E. Esarey, and W. P. Leemans. “Suppression of Beam Hosing in Plasma Accelerators with Ion Motion.” In: *Phys. Rev. Lett.* 121 (26 Dec. 2018), p. 264802. DOI: [10.1103/PhysRevLett.121.264802](https://doi.org/10.1103/PhysRevLett.121.264802).
- [105] T. J. Mehrling, C. Benedetti, C. B. Schroeder, A. Martinez de la Ossa, J. Osterhoff, E. Esarey, and W. P. Leemans. “Accurate modeling of the hose instability in plasma wakefield accelerators.” In: *Physics of Plasmas* 25.5 (2018), p. 056703. DOI: [10.1063/1.5017960](https://doi.org/10.1063/1.5017960).
- [106] T. J. Mehrling, C. Benedetti, C. B. Schroeder, and J. Osterhoff. “HiPACE: a quasi-static particle-in-cell code.” In: *Plasma Physics and Controlled Fusion* 56.8 (2014), p. 084012. DOI: [10.1088/0741-3335/56/8/084012](https://doi.org/10.1088/0741-3335/56/8/084012).
- [107] T. J. Mehrling, R. A. Fonseca, A. Martinez de la Ossa, and J. Vieira. “Mitigation of the Hose Instability in Plasma-Wakefield Accelerators.” In: *Phys. Rev. Lett.* 118 (17 Apr. 2017), p. 174801. DOI: [10.1103/PhysRevLett.118.174801](https://doi.org/10.1103/PhysRevLett.118.174801).
- [108] T. J. Mehrling, R. A. Fonseca, A. Martinez de la Ossa, and J. Vieira. “Mechanisms for the mitigation of the hose instability in plasma-wakefield accelerators.” In: *Phys. Rev. Accel. Beams* 22 (3 Mar. 2019), p. 031302. DOI: [10.1103/PhysRevAccelBeams.22.031302](https://doi.org/10.1103/PhysRevAccelBeams.22.031302).
- [109] Message Passing Interface Forum. *MPI: A Message-Passing Interface Standard - Version 3.0*. Ed. by Message Passing Interface Forum. Sept. 2012. URL: <http://www.mpi-forum.org/docs/mpi-3.0/mpi30-report.pdf>.

- [110] Hans Meuer, Erich Strohmaier, Jack Dongarra, and Horst Simon. *Top 500 Supercomputers*. 2022. URL: <http://www.top500.org/lists/2022/11/>.
- [111] B. Miao, L. Feder, J. E. Shrock, A. Goffin, and H. M. Milchberg. "Optical Guiding in Meter-Scale Plasma Waveguides." In: *Phys. Rev. Lett.* 125 (7 Aug. 2020), p. 074801. DOI: [10.1103/PhysRevLett.125.074801](https://doi.org/10.1103/PhysRevLett.125.074801).
- [112] G. Moortgat-Pick et al. "Polarized positrons and electrons at the linear collider." In: *Physics Reports* 460.4 (2008), pp. 131–243. ISSN: 0370-1573. DOI: [10.1016/j.physrep.2007.12.003](https://doi.org/10.1016/j.physrep.2007.12.003).
- [113] Gudrid Moortgat-Pick, H Baer, M Battaglia, G Belanger, K Fujii, J Kalinowski, S Heinemeyer, Y Kiyono, K Olive, F Simon, et al. "Physics at the e^+e^- linear collider." In: *The European Physical Journal C* 75 (2015), pp. 1–178. DOI: [10.1140/epjc/s10052-015-3511-9](https://doi.org/10.1140/epjc/s10052-015-3511-9).
- [114] P. Mora and T. M. Antonsen. "Electron cavitation and acceleration in the wake of an ultraintense, self-focused laser pulse." In: *Phys. Rev. E* 53 (3 Mar. 1996), R2068–R2071. DOI: [10.1103/PhysRevE.53.R2068](https://doi.org/10.1103/PhysRevE.53.R2068).
- [115] P. Mora and T. M. Antonsen. "Kinetic modeling of intense, short laser pulses propagating in tenuous plasmas." In: *Physics of Plasmas* 4.1 (1997), pp. 217–229. DOI: [10.1063/1.872134](https://doi.org/10.1063/1.872134).
- [116] P. Muggli et al. "Collective refraction of a beam of electrons at a plasma-gas interface." In: *Phys. Rev. ST Accel. Beams* 4 (9 Sept. 2001), p. 091301. DOI: [10.1103/PhysRevSTAB.4.091301](https://doi.org/10.1103/PhysRevSTAB.4.091301).
- [117] P. Muggli et al. "Halo Formation and Emittance Growth of Positron Beams in Plasmas." In: *Phys. Rev. Lett.* 101 (5 July 2008), p. 055001. DOI: [10.1103/PhysRevLett.101.055001](https://doi.org/10.1103/PhysRevLett.101.055001).
- [118] Patric Muggli, Seung Lee, Thomas Katsouleas, Ralph Assmann, Franz-Joseph Decker, Mark J Hogan, Richard Iverson, Pantaleo Raimondi, Robert H Siemann, Dieter Walz, et al. "Refraction of a particle beam." In: *Nature* 411.6833 (2001), pp. 43–43. DOI: [10.1038/35075144](https://doi.org/10.1038/35075144).
- [119] A. Myers et al. "Porting WarpX to GPU-accelerated platforms." In: *Parallel Computing* 108 (2021), p. 102833. ISSN: 0167-8191. DOI: [10.1016/j.parco.2021.102833](https://doi.org/10.1016/j.parco.2021.102833).
- [120] Alexander Novokhatski. "BNS damping." In: *CERN Yellow Reports: Conference Proceedings*. Vol. 9. 2020, pp. 68–68. DOI: [10.23732/CYRCP-2020-009.68](https://doi.org/10.23732/CYRCP-2020-009.68).
- [121] Katsunobu Oide. "Synchrotron-Radiation Limit on the Focusing of Electron Beams." In: *Phys. Rev. Lett.* 61 (15 1988), pp. 1713–1715. DOI: [10.1103/PhysRevLett.61.1713](https://doi.org/10.1103/PhysRevLett.61.1713).

- [122] A. Martinez de la Ossa, T. J. Mehrling, and J. Osterhoff. “Intrinsic Stabilization of the Drive Beam in Plasma-Wakefield Accelerators.” In: *Phys. Rev. Lett.* 121 (6 Aug. 2018), p. 064803. DOI: [10.1103/PhysRevLett.121.064803](https://doi.org/10.1103/PhysRevLett.121.064803).
- [123] W. K. H. Panofsky and M. Bander. “Asymptotic Theory of Beam Break-Up in Linear Accelerators.” In: *Review of Scientific Instruments* 39.2 (1968), pp. 206–212. DOI: [10.1063/1.1683315](https://doi.org/10.1063/1.1683315).
- [124] W. K. H. Panofsky and W. A. Wenzel. “Some Considerations Concerning the Transverse Deflection of Charged Particles in Radio-Frequency Fields.” In: *Review of Scientific Instruments* 27.11 (1956), pp. 967–967. DOI: [10.1063/1.1715427](https://doi.org/10.1063/1.1715427).
- [125] R. Paschotta, J. Nilsson, A.C. Tropper, and D.C. Hanna. “Ytterbium-doped fiber amplifiers.” In: *IEEE Journal of Quantum Electronics* 33.7 (1997), pp. 1049–1056. DOI: [10.1109/3.594865](https://doi.org/10.1109/3.594865).
- [126] A. Pukhov and J. Meyer-ter Vehn. “Laser wake field acceleration: the highly non-linear broken-wave regime.” In: *Applied Physics B* 74.4-5 (2002), pp. 355–361. ISSN: 0946-2171. DOI: [10.1007/s003400200795](https://doi.org/10.1007/s003400200795).
- [127] Alexander Pukhov and John P. Farmer. “Stable Particle Acceleration in Coaxial Plasma Channels.” In: *Phys. Rev. Lett.* 121 (26 Dec. 2018), p. 264801. DOI: [10.1103/PhysRevLett.121.264801](https://doi.org/10.1103/PhysRevLett.121.264801).
- [128] Lars Reichwein, Alexander Pukhov, Anton Golovanov, and Igor Yu. Kostyukov. “Positron acceleration via laser-augmented blowouts in two-column plasma structures.” In: *Phys. Rev. E* 105 (5 2022), p. 055207. DOI: [10.1103/PhysRevE.105.055207](https://doi.org/10.1103/PhysRevE.105.055207).
- [129] M. Reiser. *Theory and Design of Charged Particle Beams*. Wiley series in beam physics and accelerator technology. John Wiley & Sons, New York, 2008. ISBN: 978-3-527-40741-5. DOI: [10.1002/9783527617623](https://doi.org/10.1002/9783527617623).
- [130] B. Richter. “Very high energy electron-positron colliding beams for the study of weak interactions.” In: *Nuclear Instruments and Methods* 136.1 (1976), pp. 47–60. ISSN: 0029-554X. DOI: [0.1016/0029-554X\(76\)90396-7](https://doi.org/10.1016/0029-554X(76)90396-7).
- [131] J. B. Rosenzweig, B. Breizman, T. Katsouleas, and J. J. Su. “Acceleration and focusing of electrons in two-dimensional nonlinear plasma wake fields.” In: *Phys. Rev. A* 44 (10 Nov. 1991), R6189–R6192. DOI: [10.1103/PhysRevA.44.R6189](https://doi.org/10.1103/PhysRevA.44.R6189).
- [132] Thomas Roser et al. *Report of the Snowmass 2021 Collider Implementation Task Force*. 2022. DOI: [10.48550/ARXIV.2208.06030](https://doi.org/10.48550/ARXIV.2208.06030).
- [133] R. Roussel et al. “Single Shot Characterization of High Transformer Ratio Wakefields in Nonlinear Plasma Acceleration.” In: *Phys. Rev. Lett.* 124 (4 Jan. 2020), p. 044802. DOI: [10.1103/PhysRevLett.124.044802](https://doi.org/10.1103/PhysRevLett.124.044802).

- [134] Ronald D Ruth, AW Chao, Perry B Wilson, and PL Morton. "A plasma wake field accelerator." In: *Part. Accel.* 17.SLAC-PUB-3374 (1984), p. 171. URL: <https://cds.cern.ch/record/157249/files/p171.pdf>.
- [135] Sarah Schröder, CA Lindstrøm, S Bohlen, G Boyle, Richard D'Arcy, S Diederichs, MJ Garland, P Gonzalez, A Knetsch, V Libov, et al. "High-resolution sampling of beam-driven plasma wakefields." In: *Nature communications* 11.1 (2020), pp. 1–6. DOI: [10.1038/s41467-020-19811-9](https://doi.org/10.1038/s41467-020-19811-9).
- [136] C. B. Schroeder, C. Benedetti, S. S. Bulanov, D. Terzani, E. Esarey, and C. G. R. Geddes. "Beam dynamics challenges in linear colliders based on laser-plasma accelerators." In: *Journal of Instrumentation* 17.05 (May 2022), P05011. DOI: [10.1088/1748-0221/17/05/p05011](https://doi.org/10.1088/1748-0221/17/05/p05011).
- [137] C. B. Schroeder, C. Benedetti, E. Esarey, and W. P. Leemans. "Beam loading in a laser-plasma accelerator using a near-hollow plasma channel." In: *Physics of Plasmas (1994-present)* 20.12, 123115 (2013), pp. -. DOI: [10.1063/1.4849456](https://doi.org/10.1063/1.4849456).
- [138] C. B. Schroeder and E. Esarey. "Relativistic warm plasma theory of nonlinear laser-driven electron plasma waves." In: *Phys. Rev. E* 81 (5 May 2010), p. 056403. DOI: [10.1103/PhysRevE.81.056403](https://doi.org/10.1103/PhysRevE.81.056403).
- [139] C. B. Schroeder, E. Esarey, C. Benedetti, and W. P. Leemans. "Control of focusing forces and emittances in plasma-based accelerators using near-hollow plasma channels." In: *Physics of Plasmas (1994-present)* 20.8, 080701 (2013), pp. -. DOI: [10.1063/1.4817799](https://doi.org/10.1063/1.4817799).
- [140] C. B. Schroeder, E. Esarey, C. G. R. Geddes, C. Benedetti, and W. P. Leemans. "Physics considerations for laser-plasma linear colliders." In: *Phys. Rev. ST Accel. Beams* 13 (10 Oct. 2010), p. 101301. DOI: [10.1103/PhysRevSTAB.13.101301](https://doi.org/10.1103/PhysRevSTAB.13.101301).
- [141] C. B. Schroeder, E. Esarey, and B. A. Shadwick. "Warm wave breaking of nonlinear plasma waves with arbitrary phase velocities." In: *Phys. Rev. E* 72 (5 Nov. 2005), p. 055401. DOI: [10.1103/PhysRevE.72.055401](https://doi.org/10.1103/PhysRevE.72.055401).
- [142] C. B. Schroeder, D. H. Whittum, and J. S. Wurtele. "Multimode Analysis of the Hollow Plasma Channel Wakefield Accelerator." In: *Phys. Rev. Lett.* 82 (6 Feb. 1999), pp. 1177–1180. DOI: [10.1103/PhysRevLett.82.1177](https://doi.org/10.1103/PhysRevLett.82.1177).
- [143] C.B. Schroeder, C. Benedetti, E. Esarey, and W.P. Leemans. "Laser-plasma-based linear collider using hollow plasma channels." In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 829 (2016). 2nd European Advanced Accelerator

- Concepts Workshop - EAAC 2015, pp. 113–116. ISSN: 0168-9002. DOI: [10.1016/j.nima.2016.03.001](https://doi.org/10.1016/j.nima.2016.03.001).
- [144] Daniel Schulte. “Application of Advanced Accelerator Concepts for Colliders.” In: *Reviews of Accelerator Science and Technology* 09 (2016), pp. 209–233. DOI: [10.1142/S1793626816300103](https://doi.org/10.1142/S1793626816300103).
- [145] R. J. Shalloo et al. “Low-density hydrodynamic optical-field-ionized plasma channels generated with an axicon lens.” In: *Phys. Rev. Accel. Beams* 22 (4 Apr. 2019), p. 041302. DOI: [10.1103/PhysRevAccelBeams.22.041302](https://doi.org/10.1103/PhysRevAccelBeams.22.041302).
- [146] Robert J Shalloo, C Arran, L Corner, J Holloway, J Jonnerby, R Walczak, HM Milchberg, and SM Hooker. “Hydrodynamic optical-field-ionized plasma channels.” In: *Physical Review E* 97:5 (2018), p. 053203. DOI: [10.1103/PhysRevE.97.053203](https://doi.org/10.1103/PhysRevE.97.053203).
- [147] V. Shiltsev and F. Zimmermann. “Modern and future colliders.” In: *Rev. Mod. Phys.* 93 (1 Mar. 2021), p. 015006. DOI: [10.1103/RevModPhys.93.015006](https://doi.org/10.1103/RevModPhys.93.015006).
- [148] Vladimir D Shiltsev. “High-energy particle colliders: past 20 years, next 20 years, and beyond.” In: *Physics-Uspokhi* 55.10 (2012), p. 965. DOI: [10.3367/UFNe.0182.201210d.1033](https://doi.org/10.3367/UFNe.0182.201210d.1033).
- [149] J. E. Shrock, B. Miao, L. Feder, and H. M. Milchberg. “Meter-scale plasma waveguides for multi-GeV laser wakefield acceleration.” In: *Physics of Plasmas* 29:7 (2022), p. 073101. DOI: [10.1063/5.0097214](https://doi.org/10.1063/5.0097214).
- [150] T. Silva, L. D. Amorim, M. C. Downer, M. J. Hogan, V. Yakimenko, R. Zgadzaj, and J. Vieira. “Stable Positron Acceleration in Thin, Warm, Hollow Plasma Channels.” In: *Phys. Rev. Lett.* 127 (10 Sept. 2021), p. 104801. DOI: [10.1103/PhysRevLett.127.104801](https://doi.org/10.1103/PhysRevLett.127.104801).
- [151] Alexander Sinn. *Automatic Beam loading using a near direct SALAME algorithm*. 2022. URL: <https://github.com/Hi-PACE/hipace/pull/814>.
- [152] AP Sosedkin and KV Lotov. “LCODE: a parallel quasistatic code for computationally heavy problems of plasma wakefield acceleration.” In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 829 (2016), pp. 350–352. DOI: [10.1016/j.nima.2015.12.032](https://doi.org/10.1016/j.nima.2015.12.032).
- [153] P. Sprangle, E. Esarey, and A. Ting. “Nonlinear interaction of intense laser pulses in plasmas.” In: *Phys. Rev. A* 41 (8 Apr. 1990), pp. 4463–4469. DOI: [10.1103/PhysRevA.41.4463](https://doi.org/10.1103/PhysRevA.41.4463).
- [154] P. Sprangle, E. Esarey, and A. Ting. “Nonlinear theory of intense laser-plasma interactions.” In: *Phys. Rev. Lett.* 64 (17 Apr. 1990), pp. 2011–2014. DOI: [10.1103/PhysRevLett.64.2011](https://doi.org/10.1103/PhysRevLett.64.2011).

- [155] MJV Streeter, C Colgan, N Cavanagh, E Los, AF Antoine, T Audet, MD Balcazar, L Calvin, J Carderelli, H Ahmed, et al. "Laser Generation of Near-GeV Low Emittance Positron Beams." In: *arXiv preprint arXiv:2205.13850* (2022). DOI: [10.48550/arXiv.2205.13850](https://doi.org/10.48550/arXiv.2205.13850).
- [156] J. J. Su, T. Katsouleas, J. M. Dawson, P. Chen, M. Jones, and R. Keinigs. "Stability of the Driving Bunch in the Plasma Wake-field Accelerator." In: *IEEE Transactions on Plasma Science* 15.2 (1987), pp. 192–198. DOI: [10.1109/TPS.1987.4316684](https://doi.org/10.1109/TPS.1987.4316684).
- [157] T. Tajima and J. M. Dawson. "Laser Electron Accelerator." In: *Phys. Rev. Lett.* 43 (4 July 1979), pp. 267–270. DOI: [10.1103/PhysRevLett.43.267](https://doi.org/10.1103/PhysRevLett.43.267).
- [158] V.I. Telnov. "A high-luminosity superconducting twin e+e- linear collider with energy recovery." In: *Journal of Instrumentation* 16.12 (2021), P12025. DOI: [10.1088/1748-0221/16/12/P12025](https://doi.org/10.1088/1748-0221/16/12/P12025).
- [159] Maxence Thévenet, Severin Diederichs, Axel Huebl, Andrew Myers, Alexander Sinn, Remi Lehe, Jean-Luc Vay, and Weiqun Zhang. *Hi-PACE/hipace: v22.07*. Version v22.07. July 2022. DOI: [10.5281/zenodo.5358483](https://doi.org/10.5281/zenodo.5358483).
- [160] A. Ting, E. Esarey, and P. Sprangle. "Nonlinear wake-field generation and relativistic focusing of intense laser pulses in plasmas." In: *Physics of Fluids B: Plasma Physics* 2.6 (1990), pp. 1390–1394. DOI: [10.1063/1.859561](https://doi.org/10.1063/1.859561).
- [161] Davide Tommasini et al. "Status of the 16 T Dipole Development Program for a Future Hadron Collider." In: *IEEE Transactions on Applied Superconductivity* 28.3 (2018), pp. 1–5. DOI: [10.1109/TASC.2017.2780045](https://doi.org/10.1109/TASC.2017.2780045).
- [162] Lewi Tonks and Irving Langmuir. "Oscillations in Ionized Gases." In: *Phys. Rev.* 33 (2 Feb. 1929), pp. 195–210. DOI: [10.1103/PhysRev.33.195](https://doi.org/10.1103/PhysRev.33.195).
- [163] PV Tuv, RI Spitsyn, and KV Lotov. "Advanced quasistatic approximation." In: *Plasma Physics Reports* 49.2 (2023), pp. 229–238. DOI: [10.1134/S1063780X22601249](https://doi.org/10.1134/S1063780X22601249).
- [164] 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: *Phys. Rev. Lett.* 101 (14 Sept. 2008), p. 145002. DOI: [10.1103/PhysRevLett.101.145002](https://doi.org/10.1103/PhysRevLett.101.145002).
- [165] J.-L. Vay. "Noninvariance of Space- and Time-Scale Ranges under a Lorentz Transformation and the Implications for the Study of Relativistic Interactions." In: *Phys. Rev. Lett.* 98 (13 Mar. 2007), p. 130405. DOI: [10.1103/PhysRevLett.98.130405](https://doi.org/10.1103/PhysRevLett.98.130405).

- [166] J.-L. Vay, A. Huebl, R. Lehe, N.M. Cook, R.J. England, U. Niedermayer, P. Piot, F. Tsung, and D. Winklehner. “Modeling of advanced accelerator concepts.” In: *Journal of Instrumentation* 16.10 (Oct. 2021), T10003. DOI: [10.1088/1748-0221/16/10/t10003](https://doi.org/10.1088/1748-0221/16/10/t10003).
- [167] J.-L. Vay et al. “Modeling of a chain of three plasma accelerator stages with the WarpX electromagnetic PIC code on GPUs.” In: *Physics of Plasmas* 28.2 (2021), p. 023105. DOI: [10.1063/5.0028512](https://doi.org/10.1063/5.0028512).
- [168] Jean-Luc Vay and Rémi Lehe. “Simulations for Plasma and Laser Acceleration.” In: *Reviews of Accelerator Science and Technology* 09 (2016), pp. 165–186. DOI: [10.1142/S1793626816300085](https://doi.org/10.1142/S1793626816300085).
- [169] Sudharshan S. Vazhkudai et al. “The Design, Deployment, and Evaluation of the CORAL Pre-Exascale Systems.” In: *SC18: International Conference for High Performance Computing, Networking, Storage and Analysis*. 2018, pp. 661–672. DOI: [10.1109/SC.2018.00055](https://doi.org/10.1109/SC.2018.00055).
- [170] J. Vieira and J. T. Mendona. “Nonlinear Laser Driven Donut Wakefields for Positron and Electron Acceleration.” In: *Phys. Rev. Lett.* 112 (21 May 2014), p. 215001. DOI: [10.1103/PhysRevLett.112.215001](https://doi.org/10.1103/PhysRevLett.112.215001).
- [171] J. Vieira, J. T. Mendonça, and L. O. Silva. “Positron acceleration in non-linear beam driven plasma wakefields.” In: *AIP Conference Proceedings* 1777.1 (2016), p. 070012. DOI: [10.1063/1.4965655](https://doi.org/10.1063/1.4965655).
- [172] Tianhong Wang, Vladimir Khudik, Jihoon Kim, and Gennady Shvets. “Accelerated three-dimensional quasistatic particle-in-cell code.” In: *Phys. Rev. Accel. Beams* 25 (10 Oct. 2022), p. 104603. DOI: [10.1103/PhysRevAccelBeams.25.104603](https://doi.org/10.1103/PhysRevAccelBeams.25.104603).
- [173] Tianhong Wang, Vladimir Khudik, and Gennady Shvets. *Positron Acceleration in an Elongated Bubble Regime*. 2021. DOI: [10.48550/ARXIV.2110.10290](https://doi.org/10.48550/ARXIV.2110.10290).
- [174] Xiaoning Wang, Jie Gao, Qianqian Su, Jia Wang, Dazhang Li, Ming Zeng, Wei Lu, Warren B Mori, Chan Joshi, and Weiming An. “The optimal beam-loading in two-bunch nonlinear plasma wakefield accelerators.” In: *Plasma Physics and Controlled Fusion* 64.6 (May 2022), p. 065007. DOI: [10.1088/1361-6587/ac6a10](https://doi.org/10.1088/1361-6587/ac6a10).
- [175] Erich S. Weibel. “Spontaneously Growing Transverse Waves in a Plasma Due to an Anisotropic Velocity Distribution.” In: *Phys. Rev. Lett.* 2 (3 Feb. 1959), pp. 83–84. DOI: [10.1103/PhysRevLett.2.83](https://doi.org/10.1103/PhysRevLett.2.83).
- [176] Steven Weinberg. “A Model of Leptons.” In: *Phys. Rev. Lett.* 19 (21 Nov. 1967), pp. 1264–1266. DOI: [10.1103/PhysRevLett.19.1264](https://doi.org/10.1103/PhysRevLett.19.1264).

- [177] David H. Whittum, William M. Sharp, Simon S. Yu, Martin Lampe, and Glenn Joyce. "Electron-hose instability in the ion-focused regime." In: *Phys. Rev. Lett.* 67 (8 Aug. 1991), pp. 991–994. DOI: [10.1103/PhysRevLett.67.991](https://doi.org/10.1103/PhysRevLett.67.991).
- [178] Jianyuan Xiao et al. "Symplectic Structure-Preserving Particle-in-Cell Whole-Volume Simulation of Tokamak Plasmas to 111.3 Trillion Particles and 25.7 Billion Grids." In: *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis. SC '21*. Association for Computing Machinery, 2021. ISBN: 9781450384421. DOI: [10.1145/3458817.3487398](https://doi.org/10.1145/3458817.3487398).
- [179] V. Yakimenko et al. "FACET-II facility for advanced accelerator experimental tests." In: *Phys. Rev. Accel. Beams* 22 (10 Oct. 2019), p. 101301. DOI: [10.1103/PhysRevAccelBeams.22.101301](https://doi.org/10.1103/PhysRevAccelBeams.22.101301).
- [180] K. Yee. "Numerical solution of initial boundary value problems involving maxwell's equations in isotropic media." In: *Antennas and Propagation, IEEE Transactions on* 14.3 (1966), pp. 302–307. ISSN: 0018-926X. DOI: [10.1109/TAP.1966.1138693](https://doi.org/10.1109/TAP.1966.1138693).
- [181] Longqing Yi, Baifei Shen, Liangliang Ji, Konstantin Lotov, Alexander Sosedkin, Wenpeng Wang, Jiancai Xu, Yin Shi, Lingang Zhang, Zhizhan Xu, et al. "Positron acceleration in a hollow plasma channel up to TeV regime." In: *Scientific Reports* 4.1 (2014), pp. 1–5. DOI: [10.1038/srep04171](https://doi.org/10.1038/srep04171).
- [182] Longqing Yi et al. "Scheme for proton-driven plasma-wakefield acceleration of positively charged particles in a hollow plasma channel." In: *Phys. Rev. ST Accel. Beams* 16 (7 July 2013), p. 071301. DOI: [10.1103/PhysRevSTAB.16.071301](https://doi.org/10.1103/PhysRevSTAB.16.071301).
- [183] S. A. Yi, V. Khudik, C. Siemon, and G. Shvets. "Analytic model of electromagnetic fields around a plasma bubble in the blow-out regime." In: *Physics of Plasmas* 20.1, 013108 (2013). DOI: [10.1063/1.4775774](https://doi.org/10.1063/1.4775774).
- [184] Erik Zenker, Benjamin Worpitz, René Widera, Axel Huebl, Guido Juckeland, Andreas Knüpfer, Wolfgang E Nagel, and Michael Bussmann. "Alpaka—An Abstraction Library for Parallel Kernel Acceleration." In: *2016 IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW)*. IEEE, 2016, pp. 631–640. DOI: [10.1109/IPDPSW.2016.50](https://doi.org/10.1109/IPDPSW.2016.50).
- [185] Weiqun Zhang et al. "AMReX: a framework for block-structured adaptive mesh refinement." In: *Journal of Open Source Software* 4.37 (May 2019), p. 1370. DOI: [10.21105/joss.01370](https://doi.org/10.21105/joss.01370).
- [186] Jie Zhao, Yan-Ting Hu, Yu Lu, Hao Zhang, Li-Xiang Hu, Xing-Long Zhu, Zheng-Ming Sheng, Ion Cristian Edmond Turcu, Alexander Pukhov, Fu-Qiu Shao, et al. "All-optical quasi-monoenergetic

- GeV positron bunch generation by twisted laser fields." In: *Communications Physics* 5.1 (2022), pp. 1–10. DOI: [10.1038/s42005-021-00797-9](https://doi.org/10.1038/s42005-021-00797-9).
- [187] Shiyu Zhou, Weiming An, Siqin Ding, Jianfei Hua, Warren B. Mori, Chan Joshi, and Wei Lu. *Positron beam loading and acceleration in the blowout regime of plasma wakefield accelerator*. 2022. DOI: [10.48550/ARXIV.2211.07962](https://doi.org/10.48550/ARXIV.2211.07962).
- [188] Shiyu Zhou, Jianfei Hua, Weiming An, Warren B. Mori, Chan Joshi, Jie Gao, and Wei Lu. "High Efficiency Uniform Wakefield Acceleration of a Positron Beam Using Stable Asymmetric Mode in a Hollow Channel Plasma." In: *Phys. Rev. Lett.* 127 (17 Oct. 2021), p. 174801. DOI: [10.1103/PhysRevLett.127.174801](https://doi.org/10.1103/PhysRevLett.127.174801).
- [189] Shiyu Zhou, Jianfei Hua, Wei Lu, Weiming An, Qianqian Su, Warren B. Mori, and Chan Joshi. "High efficiency uniform positron beam loading in a hollow channel plasma wakefield accelerator." In: *Phys. Rev. Accel. Beams* 25 (9 Sept. 2022), p. 091303. DOI: [10.1103/PhysRevAccelBeams.25.091303](https://doi.org/10.1103/PhysRevAccelBeams.25.091303).

ACKNOWLEDGMENTS

This thesis would not have been possible without the tremendous support I have received during the last years.

I have been truly blessed to be supervised and mentored by *Jens Osterhoff*, *Carl B. Schroeder*, and *Carlo Benedetti*. Thanks to you, I always had an absolute blast doing my PhD and our Thursday night meetings felt more like a chat among friends about physics than actual work. I could not have asked for a better support in any regard. I would like to express my highest gratitude to my group leader *Maxence Thévenet*, who spent many days and late evenings with me working on HiPACE++ during the lockdown. I am thankful for his guidance, profound knowledge, and integrity. I learned a lot.

Wolfgang Hillert has kindly agreed to be both my doctoral supervisor and thesis reviewer. His uncomplicated support of this project, despite the unconventional topic, was greatly appreciated.

I would like to thank *Alexander Sinn*, who quickly became a highly skilled expert in C++ after joining us. His ability to swiftly address and resolve any issues that arose in the code was invaluable. *Axel Hübl*, *Rémi Lehe*, and *Jean-Luc Vay* have been unwavering in their support for HiPACE++ from the very beginning. In the process, Axel has significantly improved my approach to physics simulation software. *Weiqun Zhang* and *Andrew Myers* have been indispensable to this project by providing rapid solutions in AMReX.

The MPA1 team has helped me via various interesting discussions, especially *Alberto Martinez de la Ossa*, *Ángel Ferran Pousa*, and *Mathis Mewes*. Many thanks to the FLASHForward team led by *Richard D'Arcy* for our fruitful collaborations. In particular, I'd like to thank *Carl A. Lindström*, *P. Gonzalez Caminal*, and *Felipe Peña* for their effort as the first real testers of our code. I highly appreciated discussions with Carl and his positive attitude against the challenge of positron acceleration. Special thanks to *Sarah Schröder*, for the very long discussions and raising my awareness of the struggle of the real world of experimental physics.

Eric Esarey has always supported my research stays in Berkeley and promoted my scientific career. His distinctive approach of leading meetings, which involved asking many questions to junior members, has deepened my knowledge and was highly motivating.

I am deeply grateful to *Timon Mehrling*, not only for his excellent supervision I during my master's studies but also for helpful discussions on the theory model, even after he had already left the field. I appreciate the support for our project by *Spencer Gessner* and *Mark Hogan* in the joint quest of accelerating positrons with plasmas. I hope

we can deepen our collaboration with an experimental campaign in the future! I would like to thank *Marlene Turner* for the numerous stimulating discussions and for being an inspiring friend.

Last but certainly not least, I would like to express my highest gratitude to my family for whom I consider myself immensely privileged. I could always rely on their unconditional support.

EIDESSTATTLICHE VERSICHERUNG /
DECLARATION ON OATH

Hiermit versichere ich an Eides statt, die vorliegende Dissertations-
schrift selbst verfasst und keine anderen als die angegebenen Hilfsmittel
und Quellen benutzt zu haben.

Hamburg, den 02. Februar 2023

Niels Severin Diederichs