Time-Resolved Phase-Space Characterisation of Plasma-Wakefield-Accelerated Electrons at FLASHForward

Dissertation

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

vorgelegt von
Pau González Caminal

Hamburg
2022

Gutachter der Dissertation:	Prof. Dr. Wolfgang Hillert Dr. Richard D'Arcy
Zusammensetzung der Prüfungskommission:	Prof. Dr. Daniela Pfannkuche Prof. Dr. Wolfgang Hillert Prof. Dr. Gudrid Moortgat-Pick Dr. Jens Osterhoff Dr. Richard D'Arcy
Vorsitzende/r der Prüfungskommission:	Prof. Dr. Daniela Pfannkuche
Datum der Disputation:	06.07.2022
Vorsitzender Fach-Promotionsausschuss PHYSIK:	Prof. Dr. Wolfgang J. Parak
Leiter des Fachbereichs PHYSIK:	Prof. Dr. Günter H. W. Sig
Dekan der Fakultät MIN:	Prof. Dr. Heinrich Graener

Eidesstattliche Erklärung

Hiermit versichere ich an Eides statt, die vorliegende Dissertationsschrift selbst verfasst und keine anderen als die angegebenen Hilfsmittel und Quellen benutzt zu haben.

Die eingereichte schriftliche Fassung entspricht der auf dem elektronischen Speichermedium.

Die Dissertation wurde in der vorgelegten oder einer ähnlichen Form nicht schon einmal in einem früheren Promotionsverfahren angenommen oder als ungenügend beurteilt.

Hamburg, den 03.05.2022

Unterschrift der Doktorandin / des Doktoranden

Abstract

A plasma can sustain electric fields orders of magnitude larger than those attainable with the conventional radio-frequency (RF) technology typically used in particle accelerators, which are limited to $\sim$ 100 MV/m due to electrical breakdowns occurring at the metallic boundary of the accelerating structures. In a particle-beam-driven plasma-wakefield accelerator (PWFA), a charge-density wake sustaining field gradients in excess of GV/m is driven by the passage of a relativistic high-intensity particle bunch through a plasma. By harnessing the gradients of the wake, particles trailing behind the wakefield-driving bunch can be accelerated to GeV energies over meter distances, thus enabling a drastic reduction of the size of accelerator componenents and, consequently, potentially reducing the costs of future accelerator facilities. Despite this promise, however, for PWFA to be a viable technology, the quality of the accelerated bunches must match that achieved by RF-based state-of-the-art FEL linacs and particle colliders. Even though theoretical predictions suggest that PWFA schemes are capable of producing electric-field profiles with properties sufficient to preserve the longitudinal-phase-space structure of the accelerating beam, direct experimental demonstration has not yet been achieved.

In the work presented in this thesis the diagnostic capabilities of a novel X-band transverse deflection structure (TDS)—featuring femtosecond resolution and a variable polarisation of the streaking field—are exploited to investigate two mechanisms enabling the preservation of the energy spread of electron beams accelerated in a nonlinear plasma wake: optimal beam loading to preserve the correlated energy spread and a fully evacuated ion column to preserve the uncorrelated energy spread. By directly observing the longitudinal phase space of 1-GeV bunches accelerated 44 MeV in a nonlinear plasma wake, experiments performed at the FLASHForward facility (DESY, Hamburg) demonstrate that the longitudinal accelerating gradients are transversely homogeneous to within 0.8 % (1.5 %) at an interval of confidence of 68 % (95 %) and show variable amounts of beam loading depending on the exact shape of the current profile of the driver-trailing-bunch pair. The results presented in this work experimentally demonstrate the predicted suitability of PWFA for future applications requiring the preservation of high longitudinal beam quality. Furthermore, a reconstruction of the beam-plasma interaction in a particle-in-cell code has been accomplished, which illustrates the extreme sensitivity of the PWFA acceleration process to the phase-space distribution of the incoming beams. These achievements suggest that, while PWFA is capable of producing the desired field geometries, an improved control over the production of driver-trailing-bunch pairs will be required to demonstrate stable and quality-preserving acceleration at higher energy gains.

Kurzfassung

Teilchenstrahl-getriebene Plasma-Kielfeld-Beschleuniger (englisch plasma-wakefield accelerator, PWFA) können Beschleunigungsgradienten von GV/m erreichen. Damit übertreffen sie die derzeit standardmäßig verwendete Hochfrequenztechnologie um mehrere Größenordnungen. Die Hochfrequenztechnologie besteht aus metallischen Beschleunigerstrukturen, deren beschleunigende Felder wegen elektrischen Überschlags auf ~ 100 MV/m limitiert ist. Plasmabeschleuniger sind damit ein vielversprechender Ansatz um die Größe und folglich auch die Kosten zukünftiger Beschleunigeranlagen drastisch zu verringern. Durchläuft ein auf wenige Mikrometer fokussiertes, ultra-relativistisches, geladenes Teilchenbündel Plasma, wird eine Dichtemodulation der Plasmaelektronen im Nachlauf des Teilchenbündels angeregt. In den so entstandenen elektrischen Feldern kann ein weiteres Teilchenbündel eingesetzt werden und in wenigen Metern auf Energien von GeV beschleunigt werden. Um Plasmabeschleunigung als zukunftsweisende Beschleunigertechnologie zu etablieren, muss eine Strahlqualität ähnlich der in herkömmlichen Beschleunigern möglich sein. Theoretische Vorhersagen deuten darauf hin, dass Plasmabeschleuniger in der Lage sind die Strahlqualität während des Beschleunigungsprozesses zu erhalten. Der aktuelle Forschungsschwerpunkt ist es dies auch experimentell nachzuweisen.

In dieser Arbeit wird das Potenzial einer transversal ablenkenden X-Band Struktur (englisch transverse deflection structure, TDS) als unterstützende Diagnostik für einen Plasmabeschleuniger ausgelotet. Diese neu entwickelte Diagnostik erlaubt es den longitudinalen Phasenraum eines Elektronenbündels in bisher unerreichtem Detail zu vermessen und zeichnet sich insbesondere durch die Femtosekunden-Auflösung als auch die variable Polarisation des ablkenkenden Feldes aus. Mit dieser neuartigen Diagnostik konnten zwei fundamentale Mechanismen zur Erhaltung der Qualität des beschleunigten Teilchenbündels in einem Plasmabeschleuniger untersucht werden. Zunächst wird der Mechanismus des optimalen Beam Loadings zur Erhaltung der Energiebreite beobachtet. Außerdem wird die Homogenität der erzeugten Gradienten im Kielfeld untersucht. In der FLASHForward-Anlage (DESY, Hamburg) wurde der resultierende longitudinale Phasenraum eines in Plasma um 44 MeV beschleunigten 1-GeV-Teilchenbündels direkt beobachtet. Dabei wurde eine transversale Homogeneität des beschleunigenden Feldes von $0,8$ % ($1,5$ %) mit einem Konfidenzintervall von 68 % (95 %) gemessen. Die direkte Abhängigkeit des Beam Loadings von der genauen Form des Stromprofils der Teilchenbündel konnte außerdem direkt beobachtet werden. Die in dieser Arbeit vorgestellten Ergebnisse demonstrieren experimentell die theoretisch prognostizierte Eignung von PWFA für zukünftige Anwendungen, die den Erhalt einer hohen longitudinalen Phasenraum-Qualität des Strahls erfordern. Die Modellierung der

Strahl-Plasma-Wechselwirkung in einem Particle-in-Cell-Code macht die Empfindlichkeit des PWFA-Beschleunigungsprozesses gegenüber der Phasenraumverteilung der einfallenden Teilchenbündels deutlich. Die hier gewonnenen Erkenntnisse zeigen auf, dass PWFA zwar die gewünschten Feldgeometrie liefern kann, aber eine verbesserte Kontrolle über die Produktion von Bündel-Paaren erforderlich sein wird, um eine stabile und qualitätserhaltende Beschleunigung bei höheren Energiegewinnen zu demonstrieren.

Contents

Introduction

High-energy particle accelerators are powerful machines that enable the investigation of the fundamental structure of matter. Since the discovery of X-rays [Röntgen, 1898], this part of the electromagnetic-radiation spectrum—with wavelengths ranging from a few nm to a few pm—has played a central role in fundamental and applied research, as well as in medicine. The discovery of the DNA structure [Crick and Watson, 1953; Franklin and Gossling, 1953] or that of the ribosome [Ramakrishnan, 2010], for example, represent well known historical milestones. After the first observation of the radiation emitted from electrons in a synchrotron [Elder et al., 1947], referred to as synchrotron radiation (or synchrotron light), large accelerator facilities have become essential tools for the production of increasingly brighter and shorter X-ray photon pulses. While initially being used as a by-product at electron storage rings designed and built for nuclear and sub-nuclear physics, their successful application to experiments in atomic, molecular, and solid state physics motivated the construction of new facilities exclusively dedicated and optimised to serve as light sources [Schneider, 2010]. During the first decade of the 2000s, the advent of single-pass X-ray free-electron lasers (FEL) [Madey, 1971] based on a linear accelerator (linac) and electron-bunch compressors [Dohlus, Limberg and Emma, 2005], enabled the production of high-intensity, coherent photon pulses with lengths of only a few to 100 fs [Pellegrini, Marinelli and Reiche, 2016]. These new machines, referred to as fourth-generation synchrotron-light sources, have opened the door to the exploration of ultra-fast dynamics of complex molecules and condensed matter in their natural length and time scales, i.e. ~ 1 Å and ~ 1 fs [Altarelli et al., 2007; Schroer et al., 2019].

In the field of high-energy particle physics, the collision of ultra-relativistic particles with either fixed targets or other counter-propagating particles provides a deep insight into the constituents of matter and their interactions. Since the discovery of the electron [Thomson, 1897] and the proton [Rutherford, 1911] at the turn of the 20th century, the constant development of accelerator technologies has decisively contributed to a deeper understanding of the inner workings of nature by enabling the detection of a host of elementary and composite particles—the lastest being the Higgs boson [Higgs, 1964], produced for first time in 2012 at the Large Hadron Collider (LHC) [Evans and Bryant, 2008] in proton-proton collisions at energies of 7–8 TeV [Aad et al., 2012; Chatrchyan et al., 2012]. This milestone constitutes the most recent confirmation of a prediction based on the Standard Model (SM) of particle physics [Weinberg, 1967], a quantum-field theory that currently provides the most accurate description of three of the four fundamental forces of nature—the electromagnetic, weak, and strong interactions, omitting gravity.

While the SM exhibits an overwhelming agreement with a broad range of experimental observations [e.g., Odom et al., 2006], its incompleteness is exposed by a number of unanswered questions—among others: why the scale of the Higgs-boson mass is so much different from naive quantum-mechanical expectations (the "naturalness" or "hierarchy" problem) [Arkani-Hamed, Dimopoulos and Dvali, 1998; Randall, 1999]; what is the composition of dark matter [Bertone, Hooper and Silk, 2005]; what is the reason of the baryon (matter-antimatter) asymmetry [Canetti, Drewes and Shaposhnikov, 2012]; or whether the interactions described by the SM can be unified with general relativity, which is expected to be observable at energies near the Planck scale ($\sim 1.2 \cdot 10^{16}$ TeV). In spite of widespread hope, besides the Higgs-boson, the experiments performed so far at the LHC have not provided evidence of new particles that could pave the way towards an alternative theory beyond the SM [Rappoccio, 2019]. As a result, the consensus among the scientific community is that if new physics exists, it is either above the current collision-energy frontier or below the current precision frontier—i.e., indistinguishable from SM backgrounds. The disturbing uncertainty of this scenario aggravates the per se controverted cost-benefit analysis of a future multi-TeV particle collider, which is strongly constrained by the highest accelerating electric-field strengths of ~ 100 MV/m achievable with conventional radio-frequency (RF) technology before material breakdown occurs [Aicheler et al., 2012]. Therefore, the quest for higher energies necessitates the exploration of novel approaches in accelerator technology that enable higher acceleration gradients, thus enabling a reduction of the size and costs of next generation TeV-particle sources [Schroeder et al., 2010; Adli et al., 2013]. Simultaneously, such technological developments could have a far-reaching impact in society by providing affordable and compact high-energy particle accelerators for a broad range of applications—e.g., industry, medicine, and photon science.

Plasma-wakefield acceleration offers a solution to this challenge through two inherent properties. Firstly, the charge-density wakes excited in a plasma medium are capable of sustaining electric fields in excess of GV/m within a spatial extent on the order of $\mathcal{O}(100\ \mu\text{m})$, thus providing a miniaturised high-gradient accelerating structure that outperforms current RF-based technology by orders of magnitude. Secondly, if the excitation is strong enough to blow out all the plasma electrons in its wake and form a cavitation, the geometry of both the longitudinal and the transverse wakefields is expected to be compatible with an aberration-free acceleration [Rosenzweig, Breizman et al., 1991], which is a fundamental requirement for mantaining the quality of the accelerated beams in order to match that achieved with state-of-the-art RF accelerators in terms of energy spread and emittance [Di Mitri and Cornacchia, 2014; Benedikt, Schulte and Zimmermann, 2015].

Such plasma waves can be excited either by intense laser pulses [Tajima and Dawson, 1979] or by ultra-relativistic particle beams [P. Chen et al., 1985; Ruth et al., 1985], which is referred to as laser-wakefield acceleration (LWFA) and beam-driven plasma-wakefield acceleration (PWFA), respectively. PWFA has the advantage that the ultra-relativistic drive beam—and consequently the wakefield structure that it produces—can propagate through the plasma over large distances at (or very close to) the speed of light $v_d \approx c$, thus automatically fulfilling the synchronicity condition required for acceleration. This is generally

not possible in the case of a laser pulse, the group velocity of which is lower than the speed of light $v_g < c$ [Esarey, Schroeder and Leemans, 2009]. Additionally, ultra-relativistic particle beams can be produced at kHz-to-MHz repetition rates and relatively high wall-plug efficiencies of about 60 % [Aicheler et al., 2012], whereas current laser systems based on chirped-pulse amplification [Strickland and Mourou, 1985] typically work at $\sim$ 1 Hz and have efficiencies below 0.1 % [Hooker et al., 2014]. By harnessing the large gradients of a PWFA, particles trailing behind the driving bunch can be accelerated to GeV energies over meter distances. Such a PWFA configuration with externally injected trailing bunches constitutes a kind of energy booster that transfers some of the driver energy to the trailing bunch, and is therefore of special interest for future high-energy particle-accelerator applications [Colby and Len, 2017; Cros et al., 2019]. The work of this thesis is devoted to PWFA studies.

Due to the small dimensions of the blown-out plasma-wake structure and the huge strength of its associated electric fields, the overall acceleration process is extremely sensitive to the phase-space distribution of both the driver and the trailing bunch as well as to relative misalignments between the two. For example, the driver bunch must have longitudinal and transverse sizes on the order or below those of the wakefield structure and its peak current must be large enough to fully expell the plasma electrons as it propagates [Rosenzweig, Breizman et al., 1991]. Furthermore, the transverse size of the trailing bunch must be matched to the strong radially focusing field of the ion volume left behind the driver [Muggli et al., 2004; Mehrling, Grebenyuk et al., 2012] and its current profile must be carefully shaped in order to flatten the otherwise steep accelerating field and maximise the energy transfer from the wake to the bunch—an effect known as beam loading [Tzoufras et al., 2008]. Additionally, also due to the strong focusing fields, centroid offsets in the plane perpendicular to the direction of propagation trigger oscillations that either destabilise the acceleration process [Whittum et al., 1991] or lead to a dilution of the phase space of the trailing bunch—with consequent degradation of the beam quality [Lindstrøm, Adli, Pfingstner et al., 2016].

Over the last two decades, experimentation with PWFA has seen rapid progress. This has been facilitated by the developement of laser-driven RF photocathode guns [Dowell et al., 2003; Stephan et al., 2010] capable of delivering electron bunches with a very high phase-space density. After acceleration to ultra-relativistic energies and longitudinal compression, these bunches represent excellent drivers of plasma wakes. Pioneering experiments were performed between 1998 and 2015 at the FFTB[1] and FACET[2] facilities. By using a single electron bunch, Blumenfeld et al., 2007 demonstrated high accelerating gradients and high energy gain of a small fraction of particles at the back of the wakefield structure, and Clayton et al., 2016 provided the first experimental observation of the potential suitability of PWFA for quality-preserving acceleration through indirect measurements of the longitudinal and transverse field structure of a blown-out plasma wake. In more advanced double-bunch experiments, M. Litos, Adli, An et al., 2014 showed a high instantaneous energy-transfer

[1] Final Focus Test Beam (FFTB) facility [Joshi et al., 2002] at the Stanford Linear Accelerator Center (SLAC).

[2] Facilities for Accelerator Science and Experimental Test (FACET) [Hogan, T. O. Raubenheimer et al., 2010] at SLAC.

efficiency of $\sim$ 30 % between a driver and a trailing bunch. However, the energy spread of the latter increased by 100 % and about $\sim$ 75 % of its charge was lost during the beam-plasma interaction, clearly indicating that a number of issues remained to be addressed in order for PWFA to be ready for applications. One of the facilities that picked up the baton of PWFA research was the Future Oriented Wakefield Acceleration Research and Developement facility at FLASH[3] (FLASHForward) at DESY, where the work presented in this thesis has been conducted.

FLASHForward [D'Arcy, Aschikhin et al., 2019], which started operation in 2017, benefits from the FEL-quality electron beams delivered by the FLASH linac [Schreiber and Faatz, 2015]. After appropriate manipulation of their longitudinal phase space, the FLASH beams are bisected to produce driver-trailing-bunch pairs that are subsequently sent through a plasma produced in a windowless gas cell with lengths of order cm. One of the most recent milestones achieved at the facility was the demonstration of energy-spread preservation and high energy-transfer efficiency by strongly beam loading the wakefield with tailored-current-profile bunches [Lindstrøm, Garland et al., 2021]. Building upon this success, the scientific goal of this thesis is to further investigate two aspects that enable the preservation of both the projected and slice energy spread of the accelerated bunch:

- optimal beam loading for projected energy spread preservation
- transverse homogeneity of the (longitudinal) accelerating field produced in the fully evacuated ion column left behind the driver for slice energy spread preservation.

To that end, the diagnostic capabilites of a novel transverse deflection RF structure (TDS) [Altenmueller, Larsen and Loew, 1964] operating in the X-band ($\sim$ 12 GHz) frequency range and featuring a variable polarisation of the streaking field are exploited [Grudiev, 2016; Craievich, Bopp et al., 2020; Marchetti, Grudiev et al., 2021].

The new TDS, dubbed the PolariX-TDS, imprints the longitudinal information of the bunch into a transverse coordinate by means of time-varying transverse fields excited in a disk-loaded RF waveguide structure—i.e., it imprints a time-dependent kick that either yaws or pitches the bunch. Through an appropriate transport of the beam to a screen station located downstream of the device, its current profile together with its sliced beam parameters in either x or y can be directly measured. Additionally, in combination with a dipole magnet, the longitudinal phase space of the bunch—i.e., time versus energy—becomes accessible in a single-shot basis. In the course of this work, the first PolariX-TDS prototype has been installed and commissioned in a dedicated beamline located downstream of the FLASHForward experimental chamber, enabling the diagnostic of the phase space of electron bunches with femtosecond-scale time resolution.

[3] Free-electron Laser in Hamburg (FLASH) [Schreiber and Faatz, 2015] at DESY.

Thesis overview

The ultimate scientific goal of this thesis is to gain a better insight into the wakefield geometries produced in a PWFA by means of a time-resolved phase-space characterisation of electron bunches accelerated in a blown-out plasma wake. Such a characterisation unveils important aspects of the acceleration process that need to be addressed for PWFA to produce an aberration-free acceleration. To present these studies in an organised manner, the thesis is divided in the following chapters:

Chapter 1 – Reviews the fundamental concepts relevant to the experimental work presented in this thesis, including: 1) the theory of beam dynamics in charged-particle transport systems in linear approximation; 2) different methods for measuring the transverse emittance of ultra-relativistic particle beams and a novel method for matching small beta functions using centroid jitter and two beam position monitors developed at FLASHForward [Lindstrøm, D'Arcy et al., 2020]; 3) a detailed account of the operation principles and diagnostic capabilities of a TDS; 4) key aspects of plasmas and PWFA schemes with external injection of electron beams.

Chapter 2 – Introduces the FLASH linac, the FLASHForward experimental facility and the new PolariX-TDS diagnostics beamline, including several design aspects of the latter that have beed addressed in the framework of this study.

Chapter 3 – Gives an account of the RF design of the PolariX-TDS system conceived by Grudiev, 2016 and the hardware commissioning of some of its critical RF-components performed by several technical groups at DESY. It further reports on the beam-based commissioning of the PolariX-TDS conducted in the course of this work, including two applications for beam characterisation and optimisation.

Chapter 4 – Outlines the details of the optimisation and characterisation of FLASH beams for plasma acceleration in the context of the experimental campaign conducted to address the main scientific goal of this thesis. The chapter provides a broad perspective of the complete procedure required to set up the experiment and introduces some of the distinct techniques used therein.

Chapter 5 – Presents a detailed characterisation of electron bunches accelerated in a blown-out beam-loaded plasma wake and discusses key aspects of the beam-plasma interaction related to the preservation of both the slice and the projected energy spread of the accelerated beams. Finally, it analyses the acceleration process by means of a reconstruction of the beam-plasma interaction in the particle-in-cell code HiPACE++.

Conclusion – Summarises the work executed during the different stages of the present study and discusses the implications of the obtained results. To conclude, it outlines possible strategies to improve the outcome of the performed experiments with a view towards PWFA at higher energy gains.

1. Fundamental concepts

In this chapter, the fundamental concepts relevant to the experimental work presented in this thesis are provided. Section 1.1 lays out the theoretical background that enables the description of the dynamics of relativistic electron beams within a nonperiodic *transport system*—also referred to as a *single-pass transfer beamline* or simply *beamline*. In Section 1.2, different methods to measure the transverse emittance of the beams are discussed, which constitute one of the basic tools to experimentally assess the beam dynamics along the beamline. Section 1.3 delves into the longitudinal-diagnostic capabilities of RF transverse deflection structures (TDS) that are at the core of the experimental studies presented in Chapters 3–5. Finally, the underlying physics of beam-driven plasma-wakefield acceleration (PWFA) is presented in Section 1.4 together with the requirements for applications and the experimental progress made during the last decades.

1.1. Beam dynamics in charged-particle transport systems

The motion of a charged particle in an electromagnetic field is described by the Lorentz equation:

$$\boldsymbol{F} = q(\boldsymbol{\mathcal{E}} + \boldsymbol{v} \times \boldsymbol{\mathcal{B}}), \tag{1.1.1}$$

where $\boldsymbol{\mathcal{E}}$ and $\boldsymbol{\mathcal{B}}$ are the electric and magnetic fields, q and $\boldsymbol{v}$ are the charge and velocity of the particle, and $\boldsymbol{F}$ is the resulting force exerted by the fields on the particle. This equation implies that only the electric field can be used to increase the kinetic energy of the particle, whereas the magnetic field, due to the cross product $(\boldsymbol{v} \times \boldsymbol{\mathcal{B}}) \perp \boldsymbol{v}$, can only deflect its path. Moreover, since the effect induced by the magnetic field scales with the particle velocity, the deflection at relativistic velocities is most efficiently achieved by magnetic fields[1]. As a result, in conventional accelerators, electric fields are used to accelerate the particles—mainly by either containing electromagnetic RF fields in resonant cavities or by guiding them through loaded RF waveguides [Wangler, 2008]—while static magnetic fields are used to guide and focus the particles along the accelerator—the latter being the subject of this section. Before delving into the equations of motion and their solutions, it is convenient to introduce some

[1] As an example, the electric field required to produce the same effect as a moderate magnetic field strength of $\mathcal{B} = 1$ T, corresponds to $\mathcal{E} = c\mathcal{B} \sim 3 \cdot 10^8$ V/m, which is far beyond the limits of conventional technology—although, as will be seen in Section 1.4, it is within reach in a plasma wake.

concepts that set up the framework of the discussion.

Relativistic energy and momentum. The energy and momentum of a relativistic particle can be expressed as:

$$E = \gamma_{rel} m_0 c^2, \tag{1.1.2a}$$

$$\boldsymbol{p} = \gamma_{rel} m_0 \boldsymbol{v}, \tag{1.1.2b}$$

where $m_0 \approx 0.511$ MeV is the rest mass of the electron and $\gamma_{rel} = 1/\sqrt{1-\beta_{rel}^2}$ is the *relativistic Lorentz factor*, in which $\beta_{rel} = v/c$ is the electron velocity scaled by the speed of light—the so-called *relativistic beta*.

General equation of motion and conserved quantities. In a static magnetic field the Lorentz force is perpendicular to the particle velocity, and consequently the rate of change of the energy vanishes automatically:

$$\frac{dE}{dt} = \boldsymbol{v} \cdot \boldsymbol{F} = q\boldsymbol{v} \cdot (\boldsymbol{v} \times \boldsymbol{\mathcal{B}}) = 0, \tag{1.1.3}$$

which implies that both the energy and the momentum of the particle are constants of the motion. The general equation of motion can be expressed as [Rossbach and Schmüser, 1994]:

$$\frac{d}{dt}\left(m_0 \gamma_{rel} \boldsymbol{v}\right) = \boldsymbol{F} = q\boldsymbol{v} \times \mathcal{B}, \tag{1.1.4}$$

For an arbitrary magnetic field this equation cannot be solved explicitly, and several constraints and approximations are required in order to obtain an appropriate mathematical description of the particle dynamics.

Beam rigidity. Equation 1.1.1 indicates that magnetic fields perpendicular to the direction of propagation of the particles $\boldsymbol{v} \perp \boldsymbol{\mathcal{B}}$ are the most efficient in achieving its deflection. In this configuration, for a relativistic electron with charge $-e$ propagating with velocity v through a dipole field with strength $\mathcal{B}$, the balance between the Lorentz force $-ev\mathcal{B}$ and the centrifugal force $\gamma_{rel} m_0 v^2/\rho$ leads to the expression:

$$\mathcal{B}\rho = \frac{p}{e}, \tag{1.1.5}$$

where ρ is the radius of curvature of the trajectory—i.e., the bending radius. The relation between these parameters as expressed in Equation 1.1.5 is typically referred to as *beam*

rigidity [Wiedemann, 2007] and quantifies the amount of deflection experienced by a particle in a dipole field—particles with a large momentum are more rigid and therefore less bent than particles with a small momentum, which are less rigid and therefore more easily bent.

Design orbit and magnetic midplane. In the example described above, the electron describes a curved path in a plane perpendicular to the dipole field. The major constraint imposed to the static magnetic fields designed to guide and focus the particles along the beamline is that their corresponding scalar potential $\boldsymbol{\mathcal{B}} = -\boldsymbol{\nabla}\phi$ must be an odd function in one of the axis of the transverse plane x–y. Conventionally, y is chosen and therefore: $\phi(x,y,z) = -\phi(x,-y,z)$. The symmetry plane x–z is designated as the *magnetic midplane* [Brown, 1982] or *median plane* [Courant and Snyder, 1958]. Once this constraint is fulfiled, the magnetic field $\boldsymbol{\mathcal{B}}$ on the midplane is always normal to the plane, and the trajectory of a *reference particle* with design momentum p_0 lies entirely within the magnetic midplane. This trajectory is subsequently referred to as the *reference trajectory* or *design orbit*.

Curved coordinate system. The design orbit of a beamline can become rather complex and solving the equations of motion in the coordinate system of the laboratory frame brings little reward. Therefore a new right-handed set of coordinates is introduced that transforms the laboratory frame into the reference-particle frame $(X,Y,Z) \rightarrow (x,y,s)$. This is illustrated in Figure 1.1. For a given particle in the new coordinate system, s represents the distance along the design orbit (i.e., $s = vt$), x its horizontal displacement perpendicular to $\hat{\boldsymbol{s}}$, and y its vertical displacement. The plane x–s typically corresponds to the magnetic midplane—i.e., a dipole magnet bends the particles in the horizontal plane—and $y = Y$.

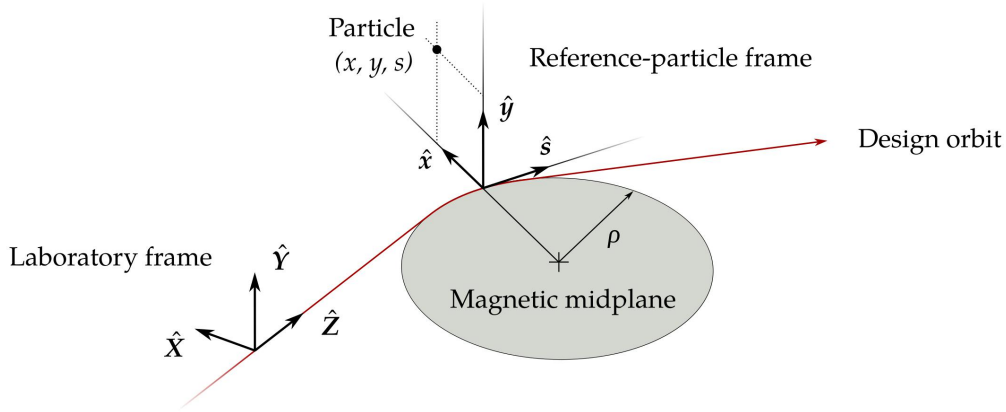

Figure 1.1.: Curved coordinate system in the reference-particle frame (reproduced from [Brown, 1982]).

Electron beams, electron bunches and paraxial approximation. A distribution of particles forms a *beam* if its mean longitudinal momentum largely exceeds its transverse momentum $\langle p_s \rangle \gg \langle p_{x/y} \rangle$ and the momentum spread in the transverse plane is much

smaller than the longitudinal momentum $\sigma_{p_{x/y}} \ll \langle p_s \rangle$ [Wiedemann, 2007]—so that the trajectories of all particles form a small angle respect to the design orbit and always lie close to it. In analogy to geometric optics, beams that fulfil the second condition are referred to as *paraxial beams* and the approximation exploiting it is called *paraxial approximation*. Due to the alternating nature of the accelerating fields used in conventional accelerators, the beams are produced as particle *bunches*: i.e., their longitudinal extent is on the same order of magnitude as their transverse extent $\sigma_s \sim \sigma_{x/y}$ [Reiser, 2008]. Additionally, for the applications of interest in this thesis, the beams typically have a small longitudinal-momentum spread $\sigma_{p_s} \ll \langle p_s \rangle$ and fulfil the condition $\sigma_{p_s} / \langle p_s \rangle \lesssim 0.1$ %.

Dynamical variables. In the curved coordinate system of the reference particle, the state of a particle at a position s along the design orbit is given by the dynamic variables: $\boldsymbol{X} = (x, x', y, y', \xi, \delta)$, where x', y', ξ and δ are defined as:

$$x' = \frac{dx}{ds} = \frac{p_x}{p_s} \approx \frac{p_x}{p_0} \tag{1.1.6a}$$

$$y' = \frac{dy}{ds} = \frac{p_y}{p_s} \approx \frac{p_y}{p_0} \tag{1.1.6b}$$

$$\xi = s - s_0 \tag{1.1.6c}$$

$$\delta = \frac{p_s - p_0}{p_0}, \tag{1.1.6d}$$

and s_0 and p_0 are the longitudinal position and momentum of the reference particle, correspondingly. In the transverse plane, the new variables x' and y' represent the *slope* (or angle) described between the particle trajectory and the design orbit. The pairs of variables x–x' and y–y' are typically referred to as *transverse phase space*—sometimes also as *trace space* [Floettmann, 2003]—and the pair of variables ξ–δ as *longitudinal phase space*.

Magnet building blocks. The magnets used in accelerators are engineered to produce a specific effect on the particles. This is achieved by means of basic field configurations that correspond to the first terms of the multipole expansion of an arbitrary magnetic field. In particular, for a simple 1D case along the x axis, the multipoles of the vertical component of a magnetic field $\mathcal{B}_y$ can be obtained by means of a Taylor expansion around $x = 0$:

$$\frac{e}{p}\mathcal{B}_y(x) = \frac{e\mathcal{B}_y}{p}\bigg|_{x=0} + \frac{e}{p}\frac{\partial \mathcal{B}_y}{\partial x}\bigg|_{x=0} x + \frac{e}{p}\frac{1}{2!}\frac{\partial^2 \mathcal{B}_y}{\partial x^2}\bigg|_{x=0} x^2 + \ldots \tag{1.1.7a}$$

$$= \frac{1}{\rho} + k_1 x + k_2 x^2 + \ldots \tag{1.1.7b}$$

dipole quadrupole sextupole

where the magnetic field $\mathcal{B}_y$ has been normalised by the beam rigidity, and the normalized quadrupole k_1 and sextupole k_2 gradients have been introduced. A constant magnetic field (dipole term) leads to a bending of the particle trajectory (cf. Equation 1.1.5), whereas a positive gradient (quadrupole term) results in a linear focusing of the particles towards the axis $x = 0$. The magnetic-field components of the first three normal multipoles in the transverse plane x–y in cartesian coordinates are summarised in Table 1.1 [Wiedemann, 2007]. Their corresponding lines of flux density and the force that they would exert on electrons travelling into the magnet—i.e., into the page—are schematically depicted in Figure 1.2.

Multipole term	$\hat{\mathcal{B}}_x$	$\hat{\mathcal{B}}_y$
Dipole		1
Quadrupole	y	x
Sextupole	$2xy$	$x^2 - y^2$

Table 1.1.: Magnetic-field components of the first three normal multipoles in the transverse plane x–y in cartesian coordinates [Wiedemann, 2007].

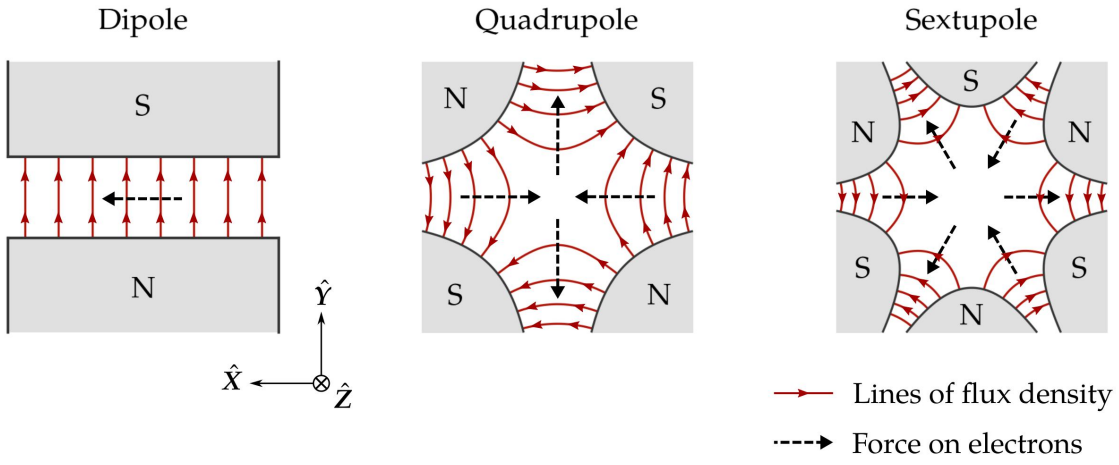

Figure 1.2.: Lines of flux density of the first three multipole terms. The force that they would exert on an electron travelling into the magnet—i.e., into the page—are indicated with dark-dashed arrows.

1.1.1. Single-particle dynamics in linear approximation

Matrix formalism

To obtain an appropriate mathematical representation of the particle dynamics in an arbitrary magnetic field with midplane symmetry, the equations of motion (cf. Eq. 1.1.4) need to be solved in the curved coordinate system of the reference particle. On account of the paraxial approximation and the condition $\delta \ll 1$, the equations of motion are expanded about the reference trajectory in terms of a Taylor series on the dynamic variables and the

multipole fields. When only the first-order terms are retained, the *Hill's equations* are obtained [Wiedemann, 2007]:

$$x'' + \left(\frac{1}{\rho^2(s)} - k_1(s)\right) x = \frac{\delta}{\rho(s)} \tag{1.1.8a}$$

$$y'' + k_1(s) y = 0, \tag{1.1.8b}$$

which make clear that the first-order description of the particle dynamics is fully determined by the magnetic dipole and quadrupole terms. For particles without momentum deviation $\delta = 0$, these are homogeneous second-order differential equations known from the theory of the classical harmonic oscillator, and their principal solutions are:

$$\left.\begin{aligned} C_u(s) &= \cos\left(\sqrt{K_u}s\right) \\ S_u(s) &= \tfrac{1}{\sqrt{K_u}} \sin\left(\sqrt{K_u}s\right) \end{aligned}\right\} \quad \text{for} \quad K_u > 0 \tag{1.1.9a}$$

$$\left.\begin{aligned} C_u(s) &= \cosh\left(\sqrt{K_u}s\right) \\ S_u(s) &= \tfrac{1}{\sqrt{K_u}} \sinh\left(\sqrt{K_u}s\right) \end{aligned}\right\} \quad \text{for} \quad K_u < 0, \tag{1.1.9b}$$

where u represents either x or y, $K_x = \dfrac{1}{\rho^2(s)} - k_1(s)$ and $K_y = k_1(s)$.

For particles with momentum deviation $\delta \neq 0$, Equation 1.1.8a becomes an inhomogeneous second-order differential equation, whose complete solution is expressed as a superposition of the principal solutions of the homogeneous case, a particular solution $D_x(s)$ of the inhomogeneous case and the initial conditions of the particle:

$$\begin{aligned} x(s) &= C_x(s)x_0 + S_x(s)x_0' + D_x(s)\delta_0 \\ x'(s) &= C_x'(s)x_0 + S_x'(s)x_0' + D_x'(s)\delta_0 \end{aligned} \tag{1.1.10a}$$

$$\begin{aligned} y(s) &= C_y(s)y_0 + S_y(s)y_0' \\ y'(s) &= C_y'(s)y_0 + S_y'(s)y_0'. \end{aligned} \tag{1.1.10b}$$

The function $D_x(s)$ is referred to as *first-order dispersion* and reflects the fact that particles with a finite momentum deviation under the influence of a dipole field are transversely dispersed around the design orbit. The first-order dispersion fulfils the equation:

$$D_x''(s) + \left(\frac{1}{\rho^2(s)} - k_1(s)\right) D_x(s) = \frac{1}{\rho(s)}, \tag{1.1.11}$$

and is solved via a Green's function integral containing the driving term $1/\rho(s)$ and the

characteristic solutions of the homogeneous equations are [Brown, 1982]:

$$\begin{aligned} D_x(s) &= S_x(s)\int_0^s \frac{C_x(\tau)}{\rho(\tau)}d\tau - C_x(s)\int_0^s \frac{S_x(\tau)}{\rho(\tau)}d\tau \\ D'_x(s) &= S'_x(s)\int_0^s \frac{C_x(\tau)}{\rho(\tau)}d\tau - C'_x(s)\int_0^s \frac{S_x(\tau)}{\rho(\tau)}d\tau, \end{aligned} \tag{1.1.12}$$

where the initial conditions $D_x(0) = 0$ and $D'_x(0) = 0$ have been assumed for simplicity.

The first-order path length is obtained by expanding the differential path length in the curved coordinate system $dr = \left(dx^2 + dy^2 + (1 + x/\rho)^2 ds^2\right)^{1/2}$ and retaining only the first-order term [Brown, 1982]:

$$\begin{aligned} \xi &= \int_0^s \frac{x(\tau)}{\rho(\tau)}d\tau + \text{higher-order terms} \\ &\approx x_0\int_0^s \frac{C_x(\tau)}{\rho(\tau)}d\tau + x'_0\int_0^s \frac{S_x(\tau)}{\rho(\tau)}d\tau + \xi_0 + \delta_0\int_0^s \frac{D_x(\tau)}{\rho(\tau)}d\tau. \end{aligned} \tag{1.1.13}$$

The solutions obtained in Equations 1.1.10 and 1.1.13, which transfer the initial particle state $\boldsymbol{X_0}$ at location s_0 to its final state $\boldsymbol{X}$ at location s, can be expressed in matrix form as $\boldsymbol{X} = \boldsymbol{R} \cdot \boldsymbol{X_0}$:

$$\begin{pmatrix} x \\ x' \\ y \\ y' \\ \xi \\ \delta \end{pmatrix} = \begin{pmatrix} C_x(s) & S_x(s) & 0 & 0 & 0 & D_x(s) \\ C'_x(s) & S'_x(s) & 0 & 0 & 0 & D'_x(s) \\ 0 & 0 & C_y(s) & S_y(s) & 0 & 0 \\ 0 & 0 & C'_y(s) & S'_y(s) & 0 & 0 \\ R_{51} & R_{52} & 0 & 0 & 1 & R_{56} \\ 0 & 0 & 0 & 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} x_0 \\ x'_0 \\ y_0 \\ y'_0 \\ \xi_0 \\ \delta_0 \end{pmatrix}, \tag{1.1.14}$$

where the matrix elements R_{51}, R_{52} and R_{56} correspond to the three first-order path-length integrals on the r.h.s. of Equation 1.1.13, by order of appearance. The absence of couplings between x–x' and y–y' is the result of the midplane-symmetry condition imposed on the magnetic fields, and the presence of couplings between x–x' and ξ–δ is the result of the dipole magnets acting on the x–s plane—typically referred to as *dispersive plane*.

The matrix shown in Eq. 1.1.14 corresponds to that of a combined function magnet, containing dipole and quadrupole components simultaneously. In practice, as mentioned before, separate-function magnets with constant fields are used, and $\rho(s)$ and $k_1(s)$ are piecewise-constant functions of the longitudinal coordinate s. To a good approximation, the *hard-edge model* can be used to characterise each element—i.e. the particles are assumed to transition abruptly from a field-free region outside the magnet into the full-strength region inside it. A more accurate description, however, requires the fringing fields extending

towards the outside of the edges to be considered, which result in additional gradients that lead to slight focusing effects. The explicit form of the matrices for the first-order-imaging magnets (i.e. dipoles and quadrupoles) in the hard-edge model is summarised in Appendix A.

The matrices describing the first-order transport of a particle through a beamline element are symplectic [Weyl, 1946] and have the following properties [Meyer, Hall and Offin, 2009]:

1. A matrix is symplectic if and only if:

$$\boldsymbol{M}^T \boldsymbol{J} \boldsymbol{M} = \boldsymbol{J} \tag{1.1.15}$$

2. The inverse of symplectic matrix is given by:

$$\boldsymbol{M}^{-1} = -\boldsymbol{J} \boldsymbol{M}^T \boldsymbol{J} \tag{1.1.16}$$

3. Its determinant is positive and equals 1:

$$\det(\boldsymbol{M}) = +1 \tag{1.1.17}$$

where:

$$\boldsymbol{J}^{2\times 2} = \begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix}, \quad \text{and} \quad \boldsymbol{J} \equiv \boldsymbol{J}_{2n} \equiv \underbrace{\boldsymbol{J}^{2\times 2} \oplus ... \oplus \boldsymbol{J}^{2\times 2}}_{n\times}. \tag{1.1.18}$$

As a consequence of these properties, the transport of an initial particle state $\boldsymbol{X_0}$ through the entire beamline can be calculated as a multiplication of matrices of individual beamline elements $\boldsymbol{R_i}$ $(i = 0, 1, ..., n)$:

$$\boldsymbol{X} = \boldsymbol{R_n} \cdot \boldsymbol{R_{n-1}} \cdot ... \cdot \boldsymbol{R_0} \cdot \boldsymbol{X_0}. \tag{1.1.19}$$

In the work presented in this thesis, calculations involving the magnetic lattice are restricted to the first-order formalism. However, the higher-order formalism is briefly referenced for completeness. Second-order terms appear in the Taylor-expanded equations of motion as a dyadic product of the dynamic variables (i.e., x_0^2, $x_0 x_0'$, $x_0 y_0$, ..., δ^2) and translate into additional driving terms that perturb the oscillatory nature of the motion. In a similar manner as exemplified with the first-order dispersion, the solutions of the resulting inhomogeneous differential equations are obtained by means of Green's function integrals that provide the couplings between the dynamical variables [Brown, 1982]. Ultimately, for each beamline element, a tensor representing the second-order aberrations of the system is added to the first-order-imaging matrices:

$$x_i = \sum_{j=1}^{6} R_{ij} x_j(0) + \sum_{j=1}^{6} \sum_{k=1}^{6} T_{ijk} x_j(0) x_k(0), \tag{1.1.20}$$

where x_i ($i = 1, ..., 6$) correspond to the dynamic variables, and $x_i(0)$ is the initial state of the particle. Explicit expressions for all the non-vanishing elements of the tensor T_{ijk} can be found in references [Brown, 1982; Iselin, 1994].

Courant-Snyder parameters

In the following, the equation of motion of a particle with design (or nominal) momentum p_0 in a focusing field is considered:

$$x'' + k_1(s)x = 0, \tag{1.1.21}$$

where $k_1(s)$ is a piecewise constant function of s. The analysis here is restricted to the horizontal motion in the x–x' plane—which is assumed to be completely decoupled from the vertical and the longitudinal planes—but applies to the y–y' plane as well. Equation 1.1.21 can be solved by introducing the trial function [Wille, 2000]:

$$x(s) = A\sqrt{\beta(s)} \cos\left(\psi(s) + \psi_0\right), \tag{1.1.22}$$

where A and ψ_0 are constants. By substituting the trial function (Eq. 1.1.22) into the Hill's equation (Eq. 1.1.21), the relation between the *beta function* $\beta(s)$ and the *phase advance* $\psi(s)$ is obtained:

$$\psi(s) = \int_0^s \frac{d\tau}{\beta(\tau)}. \tag{1.1.23}$$

If $k_1(s)$ is constant (and positive) for all s, the particle will undergo simple harmonic motion with constant frequency and amplitude. In general, however, changes in focusing strength k_1 translate into amplitude and phase modulations and the envelope of the motion is given by $A\sqrt{\beta(s)}$.

The first derivative of the trajectory $x(s)$ is:

$$x'(s) = \frac{A}{\sqrt{\beta(s)}} \left[\frac{\beta'(s)}{2} \cos\left(\psi(s) + \psi_0\right) - \sin\left(\psi(s) + \psi_0\right) \right]. \tag{1.1.24}$$

With the introduction of the new parameters:

$$\alpha(s) = -\frac{\beta'(s)}{2} \quad \text{and} \quad \gamma(s) = \frac{1+\alpha^2(s)}{\beta(s)}, \tag{1.1.25}$$

the combination[2] of Eqs. 1.1.22 and 1.1.24 results in the so-called *Courant-Snyder invariant* A [Courant and Snyder, 1958]:

$$A^2 = \gamma(s)x^2(s) + 2\alpha(s)x(s)x'(s) + \beta(s)x'^2(s), \tag{1.1.26}$$

which is the equation of an ellipse in the trace space (x, x') centered at $(0, 0)$, whose shape is characterised by the *Courant-Snyder* (or *Twiss*) *parameters* α, β and γ. On the one hand, for a constant $A = A_1$ at a fixed position $s = s_1$ along the beamline, a family of trace space points $\boldsymbol{X}(s_1, A_1) = (x, x')$ characterised by Eqs. 1.1.22 and 1.1.24 with initial phases $\psi_0 \in [0, 2\pi]$ define an ellipse that encloses an area of πA_1^2. On the other hand, at a different position $s = s_2$ the ellipse will change its shape, but its area will remain the same. Therefore, the Courant-Snyder invariant A is a constant of the motion of the particles as they move along the beamline. The invariant A^2 is sometimes expressed in terms of the so-called *single-particle emittance* ϵ_x, and in some literature is also referred to as the *action* J_x [Wolski, 2014]. The

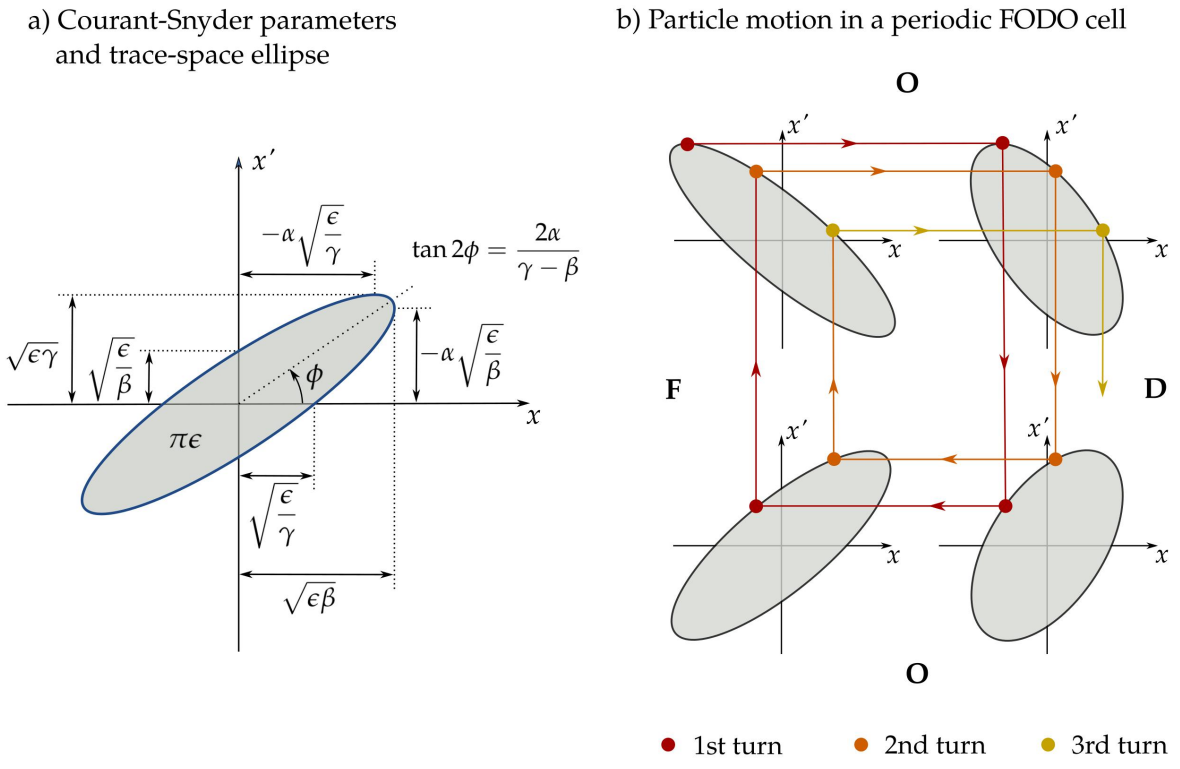

Figure 1.3.: a) Relation between the Courant-Snyder parameters and the trace-space ellipse. b) Circulation of a particle over the trace-space ellipse along its transport in a periodic FODO cell.

[2] In practice this requires to eliminate the terms depending on the phase ψ, i.e., expressing the cos and sin in terms of x, x', $\beta(s)$, $\alpha(s)$, $\gamma(s)$, and using the general relation $\cos^2 + \sin^2 = 1$.

relation between the Courant-Snyder parameters and the shape of the trace-space ellipse is represented in Figure 1.3 a).

For a closed circular machine the magnetic lattice is periodic, which leads to a periodic solution of the beta function with the periodicity of the lattice[3]. In this case, the Courant-Snyder parameters at each location s are fully determined by the lattice. Figure 1.3 b) illustrates the circulation of a particle over the trace-space ellipse and the change of shape of the ellipse for a periodic cell composed by a pair of focusing and defocusing quadrupoles with drifts in between—which is referred to as a FODO cell, where F stands for focusing, D for defocusing, and O for the drifts. Notice that in a drift the beamsize changes but the divergence is constant, whereas in a quadrupole the beamsize is approximately constant and the divergence is changed.

In contrast to a circular machine, in a single-pass transfer beamline the beta function is not uniquely determined by the transfer matrix, but depends on initial conditions that have to be specified in an adequate way. Since β and α are related to the beam size and divergence, the Courant-Snyder parameters are derived from the particle distribution in the phase space at the entrance of the beamline. This is discussed in more detail in Section 1.1.2.

By introducing the *beta matrix* $\boldsymbol{B}$:

$$\boldsymbol{B} = \begin{pmatrix} \beta & -\alpha \\ -\alpha & \gamma \end{pmatrix}, \tag{1.1.27}$$

which according to the definition of the Courant-Snyder parameters (cf. Equation 1.1.25) fulfils the condition $\det(\boldsymbol{B}) = 1$, the equation of the trace-space ellipse (cf. Equation 1.1.26) can be expressed in matrix form as:

$$A^2 = \boldsymbol{X}^T \cdot \boldsymbol{B}^{-1} \cdot \boldsymbol{X} = \begin{pmatrix} x & x' \end{pmatrix} \cdot \begin{pmatrix} \beta & \alpha \\ \alpha & \gamma \end{pmatrix} \cdot \begin{pmatrix} x \\ x' \end{pmatrix}. \tag{1.1.28}$$

Given a transfer matrix $\boldsymbol{R_x}(s_1, s_0)$ that transports the state of a particle from position s_0 to s_1 and using the condition $\boldsymbol{R_x}^{-1}\boldsymbol{R_x} = 1$ (cf. 1.1.15 and 1.1.16), the beta matrix is transported along the beamline according to the expression [Wille, 2000]:

$$\boldsymbol{B}(s_1) = \boldsymbol{R_x}(s_1, s_0) \cdot \boldsymbol{B}(s_0) \cdot \boldsymbol{R_x}^T(s_1, s_0), \tag{1.1.29}$$

which can be alternatively written as:

[3] As a matter of fact, the formal derivation of the beta function makes use of the Floquet theorem, which requires the lattice to fulfil periodicity conditions [Courant and Snyder, 1958; Rossbach and Schmüser, 1994]. However, since the focus of this thesis is on single-pass transfer beamlines, a simplified approach that yields the same results has been followed.

$$\begin{pmatrix}\beta_1\\ \alpha_1\\ \gamma_1\end{pmatrix} = \begin{pmatrix} R_{11}^2 & -2R_{11}R_{12} & R_{12}^2 \\ -R_{11}R_{21} & R_{11}R_{22}+R_{12}R_{21} & -R_{22}R_{12} \\ R_{21}^2 & -2R_{22}R_{21} & R_{22}^2 \end{pmatrix} \cdot \begin{pmatrix}\beta_0\\ \alpha_0\\ \gamma_0\end{pmatrix} \tag{1.1.30}$$

By using Eqs. 1.1.22 and 1.1.24, the transport matrix between two points along the beamline s_0 and s_1 can be expressed in terms of the Courant-Snyder parameters [Wille, 2000]:

$$\boldsymbol{R_x}(s_1, s_0) = \begin{pmatrix} \sqrt{\dfrac{\beta_1}{\beta_0}}\left(\cos\Delta\psi + \alpha_0 \sin\Delta\psi\right) & \sqrt{\beta_1\beta_0}\sin\Delta\psi \\ \dfrac{(\alpha_0-\alpha_1)\cos\Delta\psi - (1+\alpha_0\alpha_1)\sin\Delta\psi}{\sqrt{\beta_1\beta_0}} & \sqrt{\dfrac{\beta_0}{\beta_1}}(\cos\Delta\psi - \alpha_1\sin\Delta\psi) \end{pmatrix}, \tag{1.1.31}$$

where $\Delta\psi = \psi(s_1) - \psi(s_0)$ is the *phase advance* between the two locations. For a phase advance of $\Delta\psi = 90$ deg, a transverse offset at the initial location s_0 translates into an angular kick at the final location s_1 and vice versa.

1.1.2. Dynamics of particle beams

In the previous section, the mathematical tools used to describe the motion of individual particles have been presented. In practice, however, the interest is in the dynamics of beams composed by a large number of particles—e.g., 1 nC $\approx 6 \cdot 10^9$ electrons—and a statistical description in terms of averaged quantities is required, which is the subject of this section.

Liouville's theorem

The statistical description of the beam dynamics relies on Liouville's theorem, which states that the six dimensional phase-space density along any particle trajectory is a conserved quantity [Wiedemann, 2007]. The statement is valid for conservative systems whose motion is described by the Hamilton's equations—i.e., whose dynamic variables form a set of proper canonically conjugated variables [Goldstein, Poole and Safko, 2001]—and implies that the volume in phase space occupied by the beam is constant along the beam line. This imposes a constraint on how the shape of the particle distribution evolves in phase space during transport and, therefore, on the dynamic evolution of their averaged quantities.

Beam correlations and beam matrix

The statistical description of an ensemble of particles with a given phase-space density distribution $\Phi(x_1, ..., x_6)$ is substantiated by means of the covariance between the dynamic variables x_i ($i = 1, ..., 6$):

$$\langle x_i x_j \rangle = \frac{\int_{-\infty}^{\infty} (x_i - \langle x_i \rangle)(x_j - \langle x_j \rangle)\Phi(x_1, ..., x_6)dV}{\int_{-\infty}^{\infty} \Phi(x_1, ..., x_6)dV}, \tag{1.1.32}$$

where $\langle x_i \rangle$ is the first moment of the distribution on the variable x_i:

$$\langle x_i \rangle = \frac{\int_{-\infty}^{\infty} x_i \Phi(x_1, ..., x_6)dV}{\int_{-\infty}^{\infty} \Phi(x_1, ..., x_6)dV}, \tag{1.1.33}$$

$dV = \Pi_i dx_i$ is an infinitessimal element of volume and $\int_{-\infty}^{\infty}$ represents integration over all the variables x_i. For $i = j$ the covariance equals the second central moment (or mean square value) and is a measure of the spread of the distribution in x_i. The projected beamsize and divergence in the transverse phase space are given by the root mean square (rms) values:

$$\sigma_{x,rms} = \sqrt{\langle x^2 \rangle} \quad \text{and} \quad \sigma_{x',rms} = \sqrt{\langle x'^2 \rangle}, \tag{1.1.34}$$

and equally in the y–y' plane. The equivalents in the longitudinal phase space ξ–δ are referred to as *rms bunch length* and *projected rms energy spread*:

$$\sigma_{\xi,rms} = \sqrt{\langle \xi^2 \rangle} \quad \text{and} \quad \sigma_{\delta,rms} = \sqrt{\langle \delta^2 \rangle}. \tag{1.1.35}$$

The correlation between the transverse coordinates $x(s)$ or $x'(s)$—and similarly for y and y'—and the energy δ leads to the statistical definition of *beam dispersion* η:

$$\eta_x(s) = \frac{\langle x(s)\delta \rangle}{\langle \delta^2 \rangle} \quad \text{and} \quad \eta'_x(s) = \frac{\langle x'(s)\delta \rangle}{\langle \delta^2 \rangle}, \tag{1.1.36}$$

which is uniquely defined at each position s along the beamline—since so is $x(s)$. It is worth mentioning here that in single-pass transfer beamlines the lattice dispersion is defined between two locations s_0 and s along the beamline (cf. Eq. 1.1.12), and none of the dispersion sources located upstream of s_0 contribute to the dispersion functions $D_x(s, s_0)$ and $D'_x(s, s_0)$ evaluated at s [Prat, 2009]. Therefore, in general, the beam dispersion and the lattice dispersion do not need to be equivalent unless the lattice dispersion is measured from the beginning of the beamline, where the electron bunch is created. However, even in that case the two quantities might differ from each other, since the lattice dispersion only depends on the beamline geometry and does not take into account eventual collective effects (cf. Section

1.1.3) that can imprint additional correlations on the particle distribution.

The correlation between the transverse coordinates $x(s)$ or $x'(s)$ and the longitudinal coordinate $\xi(s)$ leads to the statistical definition of *beam tilt* μ:

$$\mu_x = \frac{\langle x\xi \rangle}{\langle \xi^2 \rangle} \qquad \text{and} \qquad \mu'_x = \frac{\langle x'\xi \rangle}{\langle \xi^2 \rangle}, \tag{1.1.37}$$

where the dependence of all variables on s is assumed, but has not been written for the sake of clarity. Equivalent definitions can be introduced for the vertical plane μ_y and μ'_y.

Similarly, in the longitudinal phase space, the correlation between ξ and δ leads to the statistical definition of *chirp* h:

$$h = \frac{\langle \xi\delta \rangle}{\langle \xi^2 \rangle}. \tag{1.1.38}$$

When $h \neq 0$, the projected energy spread can be easily dominated by the correlation, which accounts for the use of *correlated energy spread* as an alternative term to projected energy spread.

The matrix containing the covariance between all the dynamical variables is know as the *beam* or *covariance matrix*:

$$\boldsymbol{\Sigma} = \mathrm{Cov}(\boldsymbol{X}, \boldsymbol{X}) = \begin{pmatrix} \langle x^2 \rangle & \langle xx' \rangle & \langle xy \rangle & \langle xy' \rangle & \langle x\xi \rangle & \langle x\delta \rangle \\ \langle x'x \rangle & \langle x'^2 \rangle & \langle x'y \rangle & \langle x'y' \rangle & \langle x'\xi \rangle & \langle x'\delta \rangle \\ \langle yx \rangle & \langle yx' \rangle & \langle y^2 \rangle & \langle yy' \rangle & \langle y\xi \rangle & \langle y\delta \rangle \\ \langle y'x \rangle & \langle y'x' \rangle & \langle y'y \rangle & \langle y'^2 \rangle & \langle y'\xi \rangle & \langle y'\delta \rangle \\ \langle \xi x \rangle & \langle \xi x' \rangle & \langle \xi y \rangle & \langle \xi y' \rangle & \langle \xi^2 \rangle & \langle \xi\delta \rangle \\ \langle \delta x \rangle & \langle \delta x' \rangle & \langle \delta y \rangle & \langle \delta y' \rangle & \langle \delta\xi \rangle & \langle \delta^2 \rangle \end{pmatrix}. \tag{1.1.39}$$

It can be shown that the rotation in phase space of the rms values of any distribution is equivalent to the rotation of a hyper-ellipsoid [Floettmann, 2003]. Thus, the transport of the beam matrix from s_0 to s_1 along the beamline is achieved similarly to the beta matrix introduced in Equation 1.1.29:

$$\boldsymbol{\Sigma}(s_1) = \boldsymbol{R}(s_1, s_0) \cdot \boldsymbol{\Sigma}(s_0) \cdot \boldsymbol{R}^{\boldsymbol{T}}(s_1, s_0), \tag{1.1.40}$$

which is referred to as *beam-transport equation*.

Sliced beam parameters

In the context of this thesis, it is of special interest to characterise the particle distribution of thin slices at different locations ξ_i along the bunch, which leads to the definition of *sliced beam parameters*. In this case, the integration in Equations 1.1.32 and 1.1.33 is not performed over the whole phase-space volume V, but the longitudinal variable ξ is restricted to a region $\xi \in [\xi_i - \Delta\xi/2, \xi_i + \Delta\xi/2]$, where $\Delta\xi$ is the slice width (or size). In particular, the first moments of the transverse variables x, x', y and y' within a slice along ξ are referred to as *slice centroids*, and as *slice mean energy* when the variable is δ. Correspondingly, the second moments are referred to as *slice beam sizes* for x and y, *slice divergences* for x' and y', and *slice energy spread* for δ. On account of the considerations about the effect of the chirp on the projected energy spread discussed above, the slice energy spread is also referred to as *uncorrelated energy spread*.

When the beam correlations are not very intricate, it can be convenient to expand the sliced beam centroids and the sliced mean energy in a Taylor series around $\xi = 0$:

$$\langle x_i \rangle(\xi) = \langle x_i \rangle(0) + \frac{d\langle x_i \rangle(\xi)}{d\xi}\Big|_{\xi=0} \xi + \frac{d^2\langle x_i \rangle(\xi)}{d\xi^2}\Big|_{\xi=0} \xi^2 + \mathcal{O}(\xi^3) + \ldots \tag{1.1.41}$$

tilt / chirp — curvature — 3^{rd} order derivative

where $\langle x_i \rangle$ represents the sliced first moments of any of the transverse variables or the energy $x_i \in \{x, x', y, y', \delta\}$, and $\langle x_i \rangle(0)$ is the corresponding first moment of the projected particle distribution. As indicated in Equation 1.1.41, the linear term corresponds to either the tilt or the chirp defined previously and the second order term is referred to as curvature.

Beam emittance

The volume in phase space occupied by a particle distribution is characterised by the square root of the determinant of the covariance matrix and is referred to as the *6D rms emittance*:

$$\epsilon_{rms}^{6D} = \sqrt{\det(\Sigma)}, \tag{1.1.42}$$

which in virtue of the Liouville's theorem is a conserved quantity. Rather than the complete volume occupied by the particles, the rms emittance provides an appropriate scaling parameter that represents the high-density core of the distribution.

In the case of a completely decoupled beamline—i.e., all magnetic fields have midplane symmetry and no dispersion sources are present—the first-order motion of the particles does not show any coupling between the three spatial dimensions and the covariance matrix of the beam is block diagonal:

$$\boldsymbol{\Sigma} = \begin{pmatrix} \langle x^2 \rangle & \langle xx' \rangle & 0 & 0 & 0 & 0 \\ \langle x'x \rangle & \langle x'^2 \rangle & 0 & 0 & 0 & 0 \\ 0 & 0 & \langle y^2 \rangle & \langle yy' \rangle & 0 & 0 \\ 0 & 0 & \langle y'y \rangle & \langle y'^2 \rangle & 0 & 0 \\ 0 & 0 & 0 & 0 & \langle \xi^2 \rangle & \langle \xi\delta \rangle \\ 0 & 0 & 0 & 0 & \langle \delta\xi \rangle & \langle \delta^2 \rangle \end{pmatrix}, \tag{1.1.43}$$

from which it results that the emittance in each plane is simply:

$$\epsilon_{x,rms} = \sqrt{\langle x^2 \rangle \langle x'^2 \rangle - \langle xx' \rangle^2} \tag{1.1.44a}$$

$$\epsilon_{y,rms} = \sqrt{\langle y^2 \rangle \langle y'^2 \rangle - \langle yy' \rangle^2} \tag{1.1.44b}$$

$$\epsilon_{\xi,rms} = \sqrt{\langle \xi^2 \rangle \langle \delta^2 \rangle - \langle \xi\delta \rangle^2}. \tag{1.1.44c}$$

These are referred to as *transverse rms emittance* for the x–x' and the y–y' planes, and *longitudinal rms emittance* for the ξ–δ plane. In this case, the 6D rms emittance is exactly the product of the three projected emittances $\epsilon_{rms}^{6D} = \epsilon_{x,rms} \cdot \epsilon_{y,rms} \cdot \epsilon_{\xi,rms}$.

If sources of dispersion are present in, for instance, the x–x' plane and the beam has a finite energy spread $\langle \delta^2 \rangle \neq 0$, the trace space correlates with the longitudinal phase space:

$$\boldsymbol{\Sigma} = \begin{pmatrix} \langle x^2 \rangle & \langle xx' \rangle & 0 & 0 & \langle x\xi \rangle & \langle x\delta \rangle \\ \langle x'x \rangle & \langle x'^2 \rangle & 0 & 0 & \langle x'\xi \rangle & \langle x'\delta \rangle \\ 0 & 0 & \langle y^2 \rangle & \langle yy' \rangle & 0 & 0 \\ 0 & 0 & \langle y'y \rangle & \langle y'^2 \rangle & 0 & 0 \\ \langle \xi x \rangle & \langle \xi x' \rangle & 0 & 0 & \langle \xi^2 \rangle & \langle \xi\delta \rangle \\ \langle \delta x \rangle & \langle \delta x' \rangle & 0 & 0 & \langle \delta\xi \rangle & \langle \delta^2 \rangle \end{pmatrix}. \tag{1.1.45}$$

In this situation, the projection of the 6D rms emittance into the trace space becomes larger than the transverse rms emittance obtained in the absence of dispersion—even though the total phase-space volume is the same. This effect has to be taken into account during the experimental determination of the emittance by making sure that the dispersion along the beamline section used in the measurement is cancelled (cf. Section 1.2).

Eventually, couplings between the x–x' and the y–y' planes exist—e.g. due to the presence of skew quadrupoles or to misalignments between different beamline elements. In that case, the transverse 4D covariance matrix can be generally expressed as:

$$\boldsymbol{\Sigma}_{xy} = \begin{pmatrix} \langle x^2 \rangle & \langle xx' \rangle & \langle xy \rangle & \langle xy' \rangle \\ \langle x'x \rangle & \langle x'^2 \rangle & \langle x'y \rangle & \langle x'y' \rangle \\ \langle yx \rangle & \langle yx' \rangle & \langle y^2 \rangle & \langle yy' \rangle \\ \langle y'x \rangle & \langle y'x' \rangle & \langle y'y \rangle & \langle y'^2 \rangle \end{pmatrix} . \tag{1.1.46}$$

By means of a symplectic transformation it is possible to diagonalise this matrix, which yields the *intrinsic* emittances ϵ_1 and ϵ_2 [Woodley and Emma, 2000]:

$$\overline{\boldsymbol{\Sigma}}_{xy} = \overline{\boldsymbol{R}} \cdot \boldsymbol{\Sigma}_{xy} \cdot \overline{\boldsymbol{R}}^T = \begin{pmatrix} \epsilon_1 & 0 & 0 & 0 \\ 0 & \epsilon_1 & 0 & 0 \\ 0 & 0 & \epsilon_2 & 0 \\ 0 & 0 & 0 & \epsilon_2 \end{pmatrix} , \tag{1.1.47}$$

which can be shown to fulfil the condition [Buon, 1993]:

$$\epsilon_{rms}^{4D} = \epsilon_1 \cdot \epsilon_2 \leq \epsilon_{x,rms} \cdot \epsilon_{y,rms}. \tag{1.1.48}$$

The transverse rms emittance in trace space is called *geometric emittance* due to the geometric nature of the definition of the trace-space variables. When the beam is accelerated—i.e. the longitudinal momentum p_s is increased while the transverse momenta p_x and p_y are kept constant—the slopes are decreased (cf. Equation 1.1.6), leading to a reduction of the geometric emittance that is referred to as *adiabatic damping*. In that case, the slopes x' and y' are not proper canonical conjugates of the coordinates x and y, and, as stated by Liouville's theorem, the geometric emittance is not a conserved quantity anymore. It is therefore convenient to compute the transverse rms emittance with the canonical momenta p_x and p_y and, in addition, normalise it by $m_0 c$ to make it independent of the particle species. This leads to the definition of *normalised emittance*:

$$\begin{aligned} \epsilon_{x,n} &= \frac{1}{m_0 c} \sqrt{\langle x^2 \rangle \langle p_x^2 \rangle - \langle x p_x \rangle^2} = \frac{p_0}{m_0 c} \sqrt{\langle x^2 \rangle \langle x'^2 \rangle - \langle xx' \rangle^2} \\ &= \beta_{rel} \gamma_{rel} \epsilon_{x,rms} \approx \gamma_{rel} \epsilon_{x,rms} \ , \end{aligned} \tag{1.1.49}$$

where in the last step it has been considered that $\beta_{rel} \approx 1$. The use of normalised values allows to compare the emittace at different stages of acceleration and to compare emittances of different particle species. Similarly, the longitudinal rms emittance can be normalised so that its value is invariant under acceleration:

$$\epsilon_{\xi,n} = \frac{1}{m_0 c}\sqrt{\langle \xi^2\rangle\langle p_s^2\rangle - \langle \xi p_s\rangle^2} \approx \sqrt{\langle \xi^2\rangle\langle \gamma_s^2\rangle - \langle \xi\gamma_s\rangle^2}. \tag{1.1.50}$$

Transverse dynamics

The 2D beam matrix in trace space can be related to the Courant-Snyder parameters by calculating the covariance of an ensemble of particles according to their coordinates x and slopes x' (cf. Eqs. 1.1.22 and 1.1.24) [Wolski, 2014]. Assuming that the single-particle emittance ϵ_x and Courant-Snyder parameters β_x, α_x and γ_x are not correlated to the initial phase advance ψ_0 and that the initial phases of different particles are homogeneously distributed between 0 and 2π, the following expressions are obtained:

$$\langle x^2\rangle = \langle \epsilon_x \beta_x \cos^2\psi_x\rangle = \epsilon_{x,rms}\ \beta_x \tag{1.1.51a}$$

$$\langle xx'\rangle = \langle -\epsilon_x\alpha_x\cos^2\psi_x - \epsilon_x\cos\psi_x\sin\psi_x\rangle = -\epsilon_{x,rms}\ \alpha_x \tag{1.1.51b}$$

$$\begin{aligned}\langle x'^2\rangle &= \langle \frac{\epsilon_x}{\beta_x}\left(\alpha_x^2\cos^2\psi_x + \sin^2\psi_x - 2\alpha_x\cos\psi_{x,0}\sin\psi_x\right)\rangle \\ &= \epsilon_{x,rms}\frac{1+\alpha_x^2}{\beta_x} = \epsilon_{x,rms}\ \gamma_x,\end{aligned} \tag{1.1.51c}$$

where the subscript *rms* has been omitted from the Courant-Snyder parameters for simplicity. Equations 1.1.51 state that the statistical description of the transverse beam parameters can be expressed in terms of the transverse rms emittance and the Courant-Snyder parameters of an equivalent single particle. This can be represented in terms of the beta matrix (cf. Eq. 1.1.27) as:

$$\boldsymbol{\Sigma_x} = \begin{pmatrix}\langle x^2\rangle & \langle xx'\rangle \\ \langle x'x\rangle & \langle x'^2\rangle\end{pmatrix} = \epsilon_{x,rms}\cdot\boldsymbol{B} = \epsilon_{x,rms}\begin{pmatrix}\beta_x & -\alpha_x \\ -\alpha_x & \gamma_x\end{pmatrix}, \tag{1.1.52}$$

which, in account of the relation $\det(\boldsymbol{B}) = 1$, is consistent with the definition of the transverse rms emittance introduced in Equation 1.1.44. The validity of Eq. 1.1.52 is subject to the conditions for which the Courant-Snyder parameters are well defined—viz. monochromatic beams without couplings between x and y. When this is the case, the beamsize is simply:

$$\sigma_{x,\beta} = \sqrt{\epsilon_{x,rms}\beta_x} \tag{1.1.53}$$

and is referred to as the *betatron* (or *natural*) *beamsize*. For beams with a finite energy spread $\sigma_\delta \neq 0$ in a dispersive section, the beamsize is transported from s_0 to s according to Eq. 1.1.40:

$$\langle x_s^2 \rangle = \underbrace{R_{11}^2 \langle x_0^2 \rangle + 2R_{11}R_{12}\langle x_0 x_0' \rangle + R_{12}^2 \langle x_0'^2 \rangle}_{\sigma_{x,\beta}^2} + R_{16}^2 \langle \delta_0^2 \rangle = \sigma_{x,\beta}^2 + D_x^2 \sigma_\delta^2 \,, \tag{1.1.54}$$

from which it becomes clear that for $D_x \neq 0$ a dispersive term proportional to the energy spread is added to the betatron beamsize.

Since the normalised quadrupole strength depends on the particle momentum, particles with different energies are focused differently according to the expression:

$$k_1 = \frac{e\, g}{p_0(1+\delta)}, \tag{1.1.55}$$

where e is the electron charge, g is the gradient of the magnetic field (cf. Eq. 1.1.7), p_0 is the nominal momentum, and δ the energy deviation of the particle considered. Consequently, the trace-space ellipse associated to each energy evolves at a particular rate as the beam travels along the beamline, which can lead to an emittance growth. These *chromatic effects* are of second-order and can be corrected by means of sextupoles located in a dispersive section [Brown, 1982].

For drift-quadrupole beamlines in which chromatic effects cannot be compensated, it is useful to estimate their impact on the transverse rms emittance. To that end, the *chromatic amplitude* or *W-function* is introduced [Montague, 1979]:

$$W = \sqrt{\left(\frac{\partial \alpha}{\partial \delta} - \frac{\alpha}{\beta}\frac{\partial \beta}{\partial \delta}\right)^2 + \left(\frac{1}{\beta}\frac{\partial \beta}{\partial \delta}\right)^2}, \tag{1.1.56}$$

which quantifies the chromatic error of the beam. With this figure of merit, the relative projected-rms-emittance growth can be shown to be [Lindstrøm and Adli, 2016]:

$$\frac{\Delta \epsilon}{\epsilon_0} = \frac{1}{2} W^2 \sigma_\delta^2 + \mathcal{O}(\sigma_\delta^4), \tag{1.1.57}$$

where ϵ_0 is the initial transverse rms emittance and σ_δ the relative energy spread of the beam. In conventional accelerator beamlines, the chromatic amplitude can become large in final focusing sections in which the beamsize is easily demagnified several orders of magnitude—at FLASHForward, for instance, the beta function is focused from ~ 100 m to ~ 0.1 m at the entrance of the plasma cell (cf. Section 2.2). On the other hand, for highly divergent beams with large energy spreads as those produced by plasma-based accelerators, chromatic effects might be a concern when capturing the beam for further transport and diagnostics.

Beam matching and mismatch parameter

As mentioned in Section 1.1.1, the beta function in a single-pass transfer beamline is not uniquely determined by the lattice. Therefore, the relation between the beam matrix and the Courant-Snyder parameters (cf. Eq. 1.1.52) has to be determined experimentally—which is the subject of Section 1.2. Since the experimental beam parameters can differ from those assumed by the design optic, the strength of several quadrupoles might need to be modified to adjust the measured Courant-Snyder parameters to the design values at a selected location of the beamline. This process is referred to as *beam matching* and is critical to guarantee the proper operation of the accelerator—e.g., when establishing precise beam parameters at the waist in a final focusing section. Therefore, sections dedicated to measure the beam parameters and to match the beam to the design optics are typically foreseen at strategic locations along the beamline.

To quantify the amount of mismatch between the measured and the design Courant-Snyder parameters, the so-called *beta mismatch paramter* (or simply *mismatch parameter*) is introduced [Sands, 1991; Minty and Zimmermann, 2003]:

$$\mathrm{B_{mag}} = \frac{1}{2}\left(\beta\gamma_D - 2\alpha\alpha_D + \gamma\beta_D\right), \tag{1.1.58}$$

where the parameters with no subscript are those measured, and the ones with the subscript D are the design values. A perfectly matched beam fulfils $\mathrm{B_{mag}} = 1$. Otherwise $\mathrm{B_{mag}} > 1$.

Longitudinal dynamics

The statistical definition of chirp h was introduced with Eq. 1.1.38, which led to the nomenclature of correlated and uncorrelated energy spread as alternative terms for the projected and the slice energy spread, respectively. In principle, as will be shortly seen, the correlated energy spread can be modified by means of longitudinal-phase-space manipulations without affecting the longitudinal rms emittance. The uncorrelated energy spread $\sigma_{\delta,i}$, in contrast, is an intrinsic property of the bunch that does affect the emittance. At FLASH, its minimum value is given by the technical specifications of the photo-cathode RF gun (cf. Section 2.1) and is increased, among others, due to space-charge forces and coherent synchrotron radiation (cf. Section 1.1.3).

An energy chirp is produced by adjusting the time-of-arrival of a bunch relative to the phase of an accelerating RF field. In this situation, the energy gain at a longitudinal position ξ relative to the reference particle is given by [Wangler, 2008]:

$$\Delta E = eV\cos\left(\phi_{RF} + k_{RF}\xi\right), \tag{1.1.59}$$

where e is the electron charge, V is the RF voltage amplitude, ϕ_{RF} is the phase of the RF field relative to the reference particle, and $k_{RF} = 2\pi/\lambda_{RF}$ is the RF wave number. If the bunch length is much smaller than the RF wavelength $\sigma_\xi \ll \lambda_{RF}$, the correlation can be approximated by the slope of the field h_{RF} at the location of the bunch, and the single-particle energy is expressed as:

$$\delta = h_{RF}\xi + \delta_i, \tag{1.1.60}$$

where δ_i is the energy deviation associated to the uncorrelated (intrinsic) energy spread. The main purpose of imprinting a chirp is to trigger longitudinal dynamics in dedicated dispersive sections in order to gain control over the length of the bunch—and therefore over its current [Dohlus, Limberg and Emma, 2005].

Figure 1.4 shows the conceptual design of a widely used bunch compressor referred to as *C-chicane*. The dipole arrangement fulfils three conditions: 1) the trajectory of the outcoming bunch is colinear with that of the incoming bunch; 2) the dispersion is zero before and after the chicane $D_x = 0$, but it is nonzero inbetween; 3) the first-order longitudinal dispersion accrued along the compressor is positive $R_{56} > 0$. As the bunch travels through the compressor, particles with higher momentum (blue dots) follow a shorter path than those with lower momentum (red dots). Due to the initial negative chirp $h < 0$, the tail of the bunch catches up with the head while the head is pulled back towards the tail, leading to an overall compression. The evolution of the longitudinal phase space is shown in the three upper plots (A to C), where it can be seen that the projected energy spread is kept constant while the bunch length is progressively reduced, which translates into a dramatic increase in the peak current. The mathematical formulation of this process is briefly described in the following.

It has been shown in Section 1.1.1 that the differential path length ξ of a particle with a finite energy deviation $\delta \neq 0$ is modulated in the presence of dispersion (cf. Eqs. 1.1.13 and 1.1.20):

$$\xi_s = \xi_0 + R_{56}\delta + T_{566}\delta^2 + \mathcal{O}(\delta^3), \tag{1.1.61}$$

where R_{56} and T_{566} are the first and second-order longitudinal dispersion, correspondingly. The effect that this has on a bunch with a given chirp h can be evaluated by substituting Eq. 1.1.60 into Eq. 1.1.61 and taking the second central moment of the distribution. When retaining only terms to first order in δ, the following expressions are obtained:

$$\xi_s = (1 + R_{56}h)\xi_0 + R_{56}\delta_i \quad \longrightarrow \quad \sigma^2_{\xi,s} = (1 + R_{56}h)^2\sigma^2_{\xi,0} + R^2_{56}\sigma^2_{\delta,i}\,. \tag{1.1.62}$$

The factor $1 + R_{56}h$ determines the longitudinal dynamics of the system, leading to either a compression or a stretching of the bunch. If the condition $R_{56}h = -1$ is fulfilled, maximal

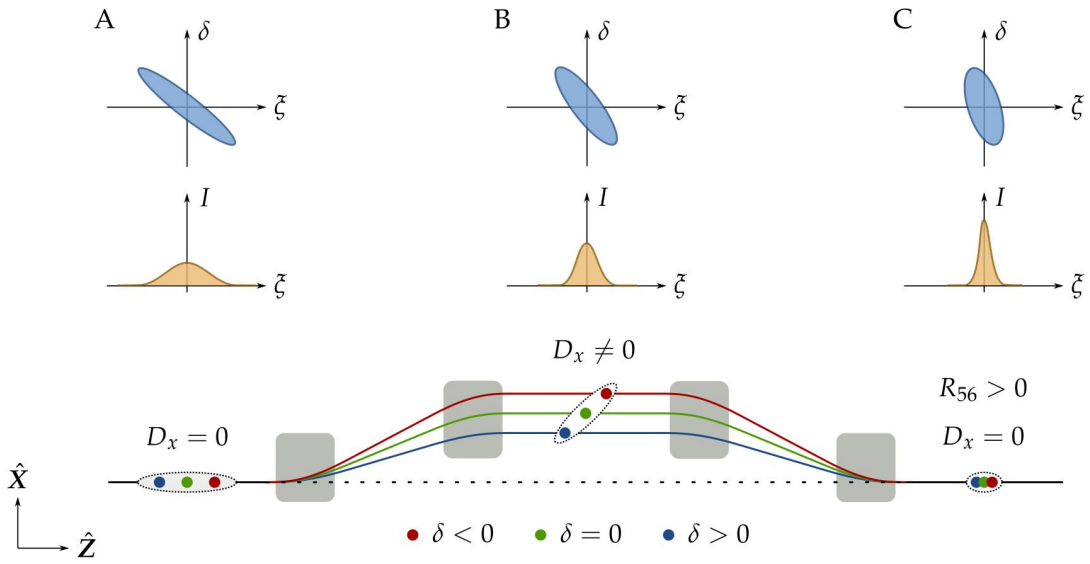

Figure 1.4.: Conceptual design of a C-chicane magnetic bunch compressor. Top: shape of the longitudinal phase space before (A) in the middle (B) and at the end (C) of the compressor. Middle: current profiles corresponding to the locations A, B and C. Bottom: trajectory of particles with different energies along the C-chicane. Scheme adapted from Wesch, 2012.

compression is achieved and the resulting bunch length is limited by the instrinsic uncorrelated energy spread $\sigma_{\xi,s} = R_{56}\sigma_{\delta,i}$. Otherwise, when $\sigma_{\xi,s} \gg R_{56}\sigma_{\delta,i}$, the overall compression is well approximated by the so-called *compression factor*:

$$C = \frac{\sigma_{\xi,0}}{\sigma_{\xi,s}} = \frac{1}{1 + R_{56}h}. \tag{1.1.63}$$

In the first acceleration stages before compression, the bunch can be long enough as to sample a large range of RF phases and the imprinted chirp cannot be approximated by the slope of the RF field anymore. These non-linearities have an impact on the current profile that is detrimental to the success of both the SASE process at FELs [Dohlus, Limberg and Emma, 2005] and the double-bunch generation in PWFA experiments with externally injected beams (cf. Section 4.2). In order to linearise the longitudinal phase space, a higher-harmonic RF system is typically used in combination with the one working at the fundamental RF frequency (cf. Section 2.1).

Transverse-longitudinal coupling

For particle beams with an energy chirp in a dispersive section, the correlation between the longitudinal coordinate and the energy translates into a transverse beam tilt in the dispersive plane, thus leading to a transverse-longitudinal coupling. This can be clearly observed in the schematic representation of a bunch compressor shown in Figure 1.4, in which the bunch in the middle of the C-chicane (location B) develops a tilt in the x–s plane. In particular,

assuming a perfectly linear chirp $\delta(\xi) = h \cdot \xi$, it can be seen that a beam with a given first-order dispersion $x(\delta) = \eta_x \cdot \delta$ acquires a tilt $x(\xi) = \mu_x \cdot \xi$ given by the expression:

$$\mu_x = \eta_x \cdot h. \tag{1.1.64}$$

As will be discussed in Chapter 2, at FLASHForward this coupling is used to split the bunch transversely in a dispersive section in order to produce the double-bunch structure required to perform beam-driven plasma-wakefield experiments with externally injected beams. Furthermore, the coupling allows to manipulate the first- and second-order terms of the transverse slice centroids in order to compensate phase-space distortions originated along the beamline due, among others, to collective effects. The technique, which was proposed in [Guetg, Beutner et al., 2015], consists of adjusting the focusing strengths of quadrupoles and/or sextupoles located in a dispersive section to introduce a beam tilt and/or curvature of the same magnitude but opposite sign with respect to the observed distortions. A practical example of that is discussed in Section 3.4.

1.1.3. Collective effects

The mathematical treatment presented in the two previous sections has only considered external fields acting on either a single particle or a distribution of non-interacting particles. However, there are several situations in which the interaction of the beam with itself has to be taken into account. The fields associated to these interactions are called *self-fields* and the resulting effects are referred to as *collective effects*. As will be seen in Chapter 3 and 4, these effects have a large impact on the phase space of the bunch and constrain the operational range of PWFA experiments performed at FLASHForward. Since a thorough study of these issues is out of the scope of this thesis, only a brief description of the most relevant effects is provided in the following.

Transverse space charge. Space-charge forces arise from the Coulomb field of the charge distribution that is seen by the beam itself. In the transverse plane, a relativistic particle travelling in the z direction not only experiences an electric field $\mathcal{E}_r$ but also an additional azimuthal magnetic field $\mathcal{B}_\phi = (\beta_{rel}/c)\mathcal{E}_r$ that contributes to the Lorentz force as a focusing term. The resulting effect can be expressed in cylindrical coordinates as [Ferrario, Migliorati and Palumbo, 2014]:

$$F_r^{SC} = q(\mathcal{E}_r + v_z \times \mathcal{B}_\phi) = q(1 - \beta_{rel}^2)\mathcal{E}_r = \frac{q\mathcal{E}_r}{\gamma_{rel}^2}. \tag{1.1.65}$$

Due to the factor $1/\gamma_{rel}^2$, at low energies the repulsive force can lead to a radial blowup of the beam and it is therefore convenient to compensate for it with radially symmetric focusing fields as those produced by a solenoid (cf. Section 2.1 and 4.1). At higher energies,

in contrast, the focusing magnetic field completely compensates the electric repulsion and transverse-space-charge effects are negligible.

Longitudinal space charge. Longitudinal-space-charge (LSC) forces result from local density variations and lead to longitudinal-phase-space distortions that can deteriorate the beam quality. An expression that captures the main features of these effects is obtained by considering a uniform transverse charge distribution with radius a and an arbitrary non-uniform longitudinal density $\lambda(\xi)$, travelling through a pipe of radius b [Ferrario, Migliorati and Palumbo, 2014]:

$$F_{\xi}^{SC}(r,\xi) = \frac{q}{4\pi\epsilon_0}\frac{1}{\gamma_{rel}^2}\left[1 - \frac{r^2}{a^2} + 2\ln\frac{b}{a}\right]\frac{\partial\lambda(\xi)}{\partial\xi}. \tag{1.1.66}$$

On the one hand, similarly to the case of transverse space charge, the dependence on the beam energy behaves as $1/\gamma_{rel}^2$—i.e., low-energy beams are strongly affected by these forces. On the other hand, the charge-density gradient determines the sign of the force, so that particles found in regions with a positive (negative) gradient experience an energy gain (loss). Initial charge density fluctuations can be driven, among others, by laser-intensity modulations at the photo-cathode and translate into energy changes through LSC forces. The resulting energy-correlated charge-density modulations are subsequently amplified when the bunch travels through a magnetic chicane, leading to the so-called *microbunching instabilities* [Saldin, Schneidmiller and Yurkov, 2004].

Coherent synchrotron radiation. Synchrotron radiation is emitted when the trajectory of a particle is bent, which in an accelerator beamline typically occurs in dipole magnets. If the bunch length is much shorter than the radiation wavelength, the radiated fields of each single electron are in phase and interfere fully coherently. Therefore, the radiation becomes fully coherent and the total radiated power P_f increses by a factor N (the number of particles) compared with the incoherent case P_{in} [Dohlus, Limberg and Emma, 2005]:

$$P_f = NP_{in} = N^2\frac{e^2c}{6\pi\epsilon_0}\frac{\gamma_{rel}^4}{\rho}, \tag{1.1.67}$$

where ρ is the bending radius of the dipole magnet. On account of the large number of electrons typically contained in a bunch ($N \approx \mathcal{O}(10^9)$), the coherent radiation is several orders of magnitude larger than the incoherent one. For small bending radii the radiation emitted from the tail can overtake the bunch and interact with the head, because the electromagnetic field travels on a shorter straight path [Dohlus and Limberg, 1997]. Since the radiated power can vary along the bunch—depending, for instance, on the current profile—particles at different longitudinal positions can experience a different energy gain (or loss). CSR effects can also result in an increase of both the slice energy spread and the slice transverse rms emittance, as well as to transverse centroid offsets along the bunch

that translate into an increase of the projected emittance. The impact that these distortions have on the electron beams used for PWFA experiments at FLASHForward is discussed in Chapter 4.

1.2. Methods for measuring the transverse emittance

In this section a general mathematical procedure based on the least-squares method is presented, that allows to retrieve the 2D transverse rms emittances in x and y from a set of beam size measurements [Minty and Zimmermann, 2003; Loehl, 2005]. After that, several experimental realisations of the method are described, each of them adjusting to particular boundary conditions.

For the rest of the section the following assumptions apply:

1. The particles are relativistic $\beta_{rel} \approx 1$.
2. The nominal momentum p_0 is known beforehand.
3. No couplings between x and y exist.
4. The dispersion is zero along the beamline section used in the measurement $D_x = 0$ and $D'_x = 0$.
5. Chromatic aberrations are negligible.

These conditions imply that the relation between the covariance matrix and the Courant-Snyder parameters is well defined and fulfils Equation 1.1.52. The discussion is restricted to the x–x' plane to ease the notation, but is also valid for the y–y' plane.

1.2.1. General approach

Given an initial beam distribution at a fixed location s_0—henceforth referred to as *reference plane*—and a set of n different transfer matrices $\boldsymbol{R}^{(i)}$ $(i = 1, ..., n)$, the beam size resulting from all transformations can be calculated by means of the usual beam-matrix-transport equation $\boldsymbol{\Sigma} = \boldsymbol{R} \cdot \boldsymbol{\Sigma_0} \cdot \boldsymbol{R}^T$ according to the following expression:

$$\underbrace{\begin{pmatrix} \langle x^2_{(1)} \rangle \\ \langle x^2_{(2)} \rangle \\ \vdots \\ \langle x^2_{(n)} \rangle \end{pmatrix}}_{\boldsymbol{B}} = \underbrace{\begin{pmatrix} {R^{(1)}_{11}}^2 & 2R^{(1)}_{11}R^{(1)}_{12} & {R^{(1)}_{12}}^2 \\ {R^{(2)}_{11}}^2 & 2R^{(2)}_{11}R^{(2)}_{12} & {R^{(2)}_{12}}^2 \\ \vdots & \vdots & \vdots \\ {R^{(n)}_{11}}^2 & 2R^{(n)}_{11}R^{(n)}_{12} & {R^{(n)}_{12}}^2 \end{pmatrix}}_{\boldsymbol{A}} \cdot \underbrace{\begin{pmatrix} \langle x_0^2 \rangle \\ \langle x_0 x'_0 \rangle \\ \langle x'^2_0 \rangle \end{pmatrix}}_{\boldsymbol{o}}, \tag{1.2.1}$$

where $\boldsymbol{o}$ contains the initial beam matrix elements, $\boldsymbol{A}$ the elements of the transfer matrices, and $\boldsymbol{B}$ the squared beam sizes of the n measurements. Since the beam matrix is defined by three unknowns, at least three beam size measurements ($n = 3$) are required to reconstruct the trace space. Typically, however, a larger number of measurements is preferred in order to be less sensitive to experimental errors.

The matrices $\boldsymbol{R}^{(i)}$ in Eq. 1.2.1 represent the beam transport along either 1) a fixed section between the reference plane s_0 and the location s of a single screen with different transfer-optics settings, or 2) multiple sections between s_0 and s_i (where s_i are the locations of several screen stations) with fixed quadrupole settings. For optimal accuracy, the difference between the betatron phase-advance associated to consecutive transfer matrices should be $\Delta\psi \approx \pi/n$. This can be understood if the measurement is thought of as a kind of a tomography in which a 2D object is reconstructed from 1D projections into one axis obtained by "rotating" the object in phase space. Therefore: 1) a reflection respect the projection axis is redundant, so that only a rotation by 180° is required; 2) the information contained in the phase-space reconstruction is maximised when the projections differ from each other by the largest amount.

The problem of determining the elements of the beam matrix $\boldsymbol{o}$ in Eq. 1.2.1 can be generally solved with the linear-least-squares method by minimising the function:

$$\chi^2 = \sum_{l=1}^{n} \frac{1}{\sigma_l^2} \left[B_l - \sum_{i=1}^{3} A_{li} o_i \right]^2 = \sum_{l=1}^{n} \left[b_l - \sum_{i=1}^{3} a_{li} o_i \right]^2 , \tag{1.2.2}$$

where σ_l is the error associated to the squared beam size $\langle x_{(l)}^2 \rangle$ and the weighted parameters have been introduced:

$$b_l = \frac{B_l}{\sigma_l} \qquad \text{and} \qquad a_{li} = \frac{A_{li}}{\sigma_l}. \tag{1.2.3}$$

The estimation of the beam matrix parameters $\boldsymbol{o}$ and its covariance matrix $\boldsymbol{\Sigma_o}$ can be shown to be [Minty and Zimmermann, 2003; Loehl, 2005]:

$$\boldsymbol{o} = \left(\boldsymbol{a}^T \boldsymbol{a} \right)^{-1} \boldsymbol{a}^T \boldsymbol{b} \qquad \text{and} \qquad \boldsymbol{\Sigma_o} = \left(\boldsymbol{a}^T \boldsymbol{a} \right)^{-1} , \tag{1.2.4}$$

from which the geometric transverse rms emittance and the Courant-Snyder parameters at the location s_0 are directly obtained:

$$\epsilon_{x,rms} = \sqrt{o_1 o_3 - o_2^2} \tag{1.2.5a}$$

$$\beta_{x,0} = \frac{o_1}{\epsilon_{x,rms}} \tag{1.2.5b}$$

$$\alpha_{x,0} = -\frac{o_2}{\epsilon_{x,rms}} \tag{1.2.5c}$$

Each of the Equations 1.2.5 is a function of the beam-matrix parameters $f = f(o_1, o_2, o_3)$ and its associated error is:

$$\sigma_f^2 = (\boldsymbol{\nabla}_o f)^T \, \boldsymbol{\Sigma}_o \, (\boldsymbol{\nabla}_o f), \tag{1.2.6}$$

where $\boldsymbol{\nabla}_o f$ is the vector formed by the partial derivatives $\boldsymbol{\nabla}_o f = \left(\frac{\partial f}{\partial o_1}, \frac{\partial f}{\partial o_2}, \frac{\partial f}{\partial o_3} \right)^T$.

1.2.2. Experimental realisation

In this section, four different experimental realisations of the general method described above are discussed—each of them adjusting to particular boundary conditions—which in one way or another are relevant to the work presented in this thesis. Additionally, a novel method developed at FLASHForward referred to as *2-BPM-tomography* is introduced, which, even if it does not provide information about the beam emittance, it allows to characterise the beam parameters with a dramatic improvement of time economy compared with the other methods.

Multiple-screen method. In the multiple-screen method the beam size is measured at several screen stations located downstream of the reference plane and the lattice optic is fixed to fulfil the measurement constraints. As an example, the DBC2 section located after the second bunch compressor at FLASH (cf. Section 2.1) is considered in the following, which is shown in Figure 1.5. In this case, four screens (A to D) are separated by three FODO cells that, according to the design optic, achieve a periodic evolution of the beta function and a progressive increase of the phase advance of about 45° ($\pi/4$) per cell. This is accomplished for x and y simultaneously, which allows to retrieve the emittance in both planes from the same measurement. Between the reference plane and the first screen (A) six quadrupoles are available to match the beam to the design optic. Due to the low number of available screens, the accuracy of the reconstructed parameters is severely compromised when the beam is mismatched, and several measurement-matching iterations are typically required to progressively reduce the errors until $\mathrm{B}_{\mathrm{mag}} \approx 1$.

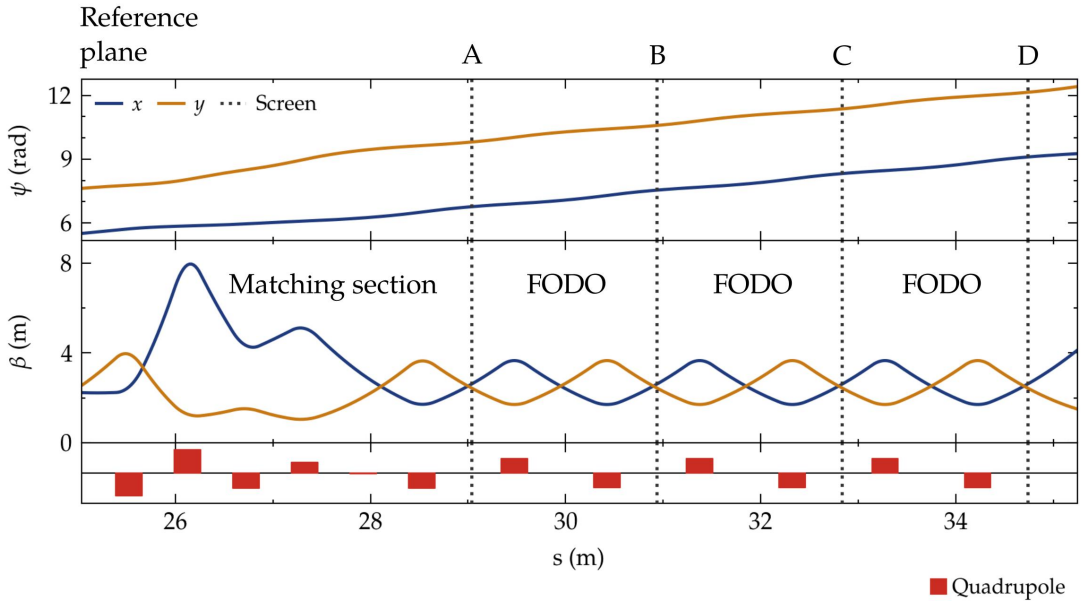

Figure 1.5.: Design lattice optics of multiple-screen emittance measurements at DBC2 (FLASH). Top: phase advance. Bottom: beta functions. The reference plane is located at the top-left, whereas the screens are installed at locations A, B, C and D.

Single-quadrupole scan. The simplest realisation of an emittance measurement is achieved by scanning the strength of a single quadrupole located upstream of a screen monitor. In this case the transfer matrix of the system is the product of a quadrupole $\boldsymbol{Q}$ and a drift $\boldsymbol{D}$ matrix:

$$\boldsymbol{R_x} = \boldsymbol{D_x} \cdot \boldsymbol{Q_x} = \begin{pmatrix} 1 & l \\ 0 & 1 \end{pmatrix} \cdot \begin{pmatrix} 1 & 0 \\ k_1 & 1 \end{pmatrix} = \begin{pmatrix} k_1 l + 1 & l \\ k_1 & 1 \end{pmatrix}, \tag{1.2.7}$$

where the thin-lens approximation has been used for the quadrupole (cf. Appendix A). After transport of the beam matrix $\boldsymbol{\Sigma_x^{scr}} = \boldsymbol{R_x} \cdot \boldsymbol{\Sigma_x^0} \cdot \boldsymbol{R_x^T}$, the beam size at the screen is:

$$\Sigma_{11}^{scr} = \left(\Sigma_{11}^0 + 2\Sigma_{12}^0 l + \Sigma_{22}^0 l^2\right) \; + \; 2\left(\Sigma_{11}^0 l + \Sigma_{12}^0 l^2\right) \cdot k_1 \; + \; \Sigma_{11}^0 l^2 \cdot k_1^2, \tag{1.2.8}$$

which is a quadratic function of k_1. Therefore, the retrieval of the beam matrix is reduced to fitting a parabola to the squared beam size for an appropriate range of quadrupole strengths—which requires the beam size to go through a waist in order to optimise the resolution. Since the quadrupole field is focusing in one plane and defocusing in the other, it is generally not possible to scan x and y simultaneously. Nevertheless, this can be achieved if the lattice geometry and the lattice optic fulfil particular boundary conditions, which enable the so-called *symmetric single-quadrupole scan* [Prat, 2014].

Multiple-quadrupole scan. In case several constraints have to be fulfilled during the measurement, a single quadrupole does not provide enough degrees of freedom and the use of multiple quadrupoles is required. This situation is encountered, for instance, when the beam size at the screen has to be kept inside a given range of values to adjust to the operational requirements of the optical system. In these circumstances, at least four quadrupoles are required in order to measure the emittance in x and y simultaneously—i.e., four variables (quadrupoles) for each of the four constraints imposed: σ_x, σ_y, ψ_x and ψ_y. Slice-emittance measurements constitute another multiple-constrain problem, in which the beam size is kept at fixed values in both planes while the phase advance is scanned over 180° in only one of them. An example of the latter is discussed in Section 3.4.

Object-plane scan. In the previous three methods, a fixed reference plane exists at which the beam parameters are reconstructed. The approach followed with the object-plane scan, instead, is to set the reference plane at different positions s_i along a natural beam waist and image the beam-size evolution by setting the imaging system to first-order point-to-point imaging with a given magnification power. In this way, the reference plane becomes the *object plane* and the imaging screen the *imaging plane*. First-order point-to-point imaging is achieved when the transport matrix fulfills the condition [Brown, 1982]:

$$\boldsymbol{R_x}(s_{scr}, s_i) = \begin{pmatrix} R_{11} & 0 \\ R_{21} & R_{22} \end{pmatrix}, \tag{1.2.9}$$

where R_{11} corresponds to the magnification of the system. That being the case, the beam size at the imaging plane—i.e., at the screen—is simply $\Sigma_{11}^{scr} = R_{11}^2 \Sigma_{11}^{s_i}$ (cf. Equation 1.1.40). Since multiple constraints need to be fulfilled simultaneously—viz. point-to-point imaging condition and magnification power—this technique constitutes a particular type of multiple-quadrupole scan. When the beam is strongly focused at a fixed location—as is the case at the interaction point in beam-driven plasma-wakefield-acceleration experiments—the phase advance rapidly accrues values close to 180 deg, which is required to achieve enough accuracy in the measurement. For a beta function at the waist of $\beta = 10$ mm, for instance, the phase advance accures $\Delta\psi = 170$ deg in a range of $\Delta s \approx 23$ cm around the waist. In general, the relation between the longitudinal extent Δs covered by the beam along a given phase advance $\Delta\psi$ around the waist can be expressed as $\Delta s = 2\beta \tan(\Delta\psi/2)$, where β is the beta function at the waist.

Alternatively to the linear-least-squares method, the data can be evaluated by considering that, according the transport equation $\boldsymbol{\Sigma} = \boldsymbol{R} \cdot \boldsymbol{\Sigma}_0 \cdot \boldsymbol{R}^T$, the beam size in a drift evolves as:

$$\Sigma_{11}(s_i) = \Sigma_{11}^0 + 2\Sigma_{12}^0 \cdot s_i + \Sigma_{22}^0 \cdot s_i^2, \tag{1.2.10}$$

where s_i are the different object-plane locations around the waist. On the other hand, according to the point-to-point imaging condition, the beam size at the screen is related

to the beam size at the location of the object plane by the expression $\Sigma_{11}^{scr} = \Sigma_{11}(s_i) \cdot {R_{11}^{(i)}}^2$. Similar to the case of a single-quadrupole scan, the beam matrix can be retrieved by fitting Eq. 1.2.10 to a parabola, from which the waist location s_w and the beam size at the waist σ_w are directly obtained:

$$s_w = -\frac{\Sigma_{12}^0}{\Sigma_{22}^0} \quad \text{and} \quad \sigma_w = \sqrt{\Sigma_{11}^0 - \frac{\left(\Sigma_{12}^0\right)^2}{\Sigma_{22}^0}}. \tag{1.2.11}$$

This method is of special interest in final-focusing sections because it provides direct information about the waist size and location. Multiple examples of that can be found in Chapters 4 and 5.

2-BPM-tomography. In PWFA experiments, matching the beam parameters at the entrance of the plasma cell requires a high degree of accuracy and several emittance measurements followed by matching iterations are typically needed until the target values are achieved. Since changing the current of a magnet is a relatively slow operation, the quadrupole scans discussed so far turn the matching routines into a time-consuming process. To overcome this problem, a novel method has been devised at FLASHForward that allows to match small β functions using centroid jitter and two beam position monitors (BPM) located upstream and downstream of the waist [Lindstrøm, D'Arcy et al., 2020]. Since this method is non-invasive, the online assessment of the beam parameters is enabled, which leads to a dramatic speed-up of the matching routines.

Similarly to any particle distribution, the beam-centroid jitter occupies a finite area in trace-space x–x' that can be used to define jitter-related Courant-Snyder parameters. The central assumption underpinning this technique is that the Courant-Snyder parameters of the jitter are similar to those of the beam. This is expected to be the case in long FODO-like transmission lines, in which different sources of jitter—e.g. magnets and accelerating cavities—are distributed accross a range of phase advances and none of them dominates exceedingly over the others. In that case, the jitter-phase-space ellipse gradually expands while it rotates to acquire a shape similar to the beam-phase-space ellipse. Even if they do not coincide exactly along the whole beamline, their evolution will be particularly similar in regions where the beam is strongly focused, which is the relevant observation for the application described here. It should be noted that this technique is not suitable for measuring the beam emittance, but this is also not required for matching.

Considering that the two BPMs are only separated by a drift[4], the jitter trace space is simply given by the centroids x and the slopes x':

[4]The measurement can also be generalized to nonballistic orbits (i.e., with magnets between the BPMs), as is relevant to, for instance, plasma accelerators with strong permanent quadrupoles close to the plasma entrance [Lindstrøm, D'Arcy et al., 2020].

$$x' = \frac{x_2 - x_1}{\Delta s}, \tag{1.2.12}$$

where x_1 and x_2 are the centroid coordinates at the upstream and downstream BPM correspondingly, and Δs is the distance between the two. By means of the covariance matrix in trace space $\boldsymbol{\Sigma}_x^{jitter} = \mathrm{Cov}(x, x')$ the evolution of the associated Courant-Snyder parameters can be described exactly as in Eq. 1.2.10, from which the waist location and the β function at the waist are directly obtained (cf. Eqs. 1.2.11). Experimental applications of this technique are discussed in Section 4.2.3.

1.3. Longitudinal diagnostics of electron bunches

In the context of this thesis, *longitudinal diagnostics* is referred to diagnostic devices that are able to resolve the longitudinal structure of an electron bunch—such as its current profile—rather than those that only aim to inform on its first moments. During the last two decades, one of the main drivers of technological development for the longitudinal diagnostic of electron bunches has been the advent and rapid progress of compact XUV and X-ray free-electron lasers (FELs), which are currently able to produce electron beams with lengths of only few femtoseconds and even below [Ganter et al., 2019]. Indeed, it is the availability of such sources of short, high-quality electron bunches that has opened the door to electron-beam-driven PWFA experiments (cf. Section 1.4.4). Therefore, to a large extent, longitudinal-diagnostic devices already available at FEL facilities are per se well suited to the requirements of PWFA experiments.

Electro-optical sampling (EOS) makes use of the birefringence of an electro-optical crystal to sample the Coulomb field of relativistic electron bunches and is able to retrieve the current profile noninvasively with femtosecond time resolution in a single-shot basis at a repetition rate of ~ 1 kHz [Steffen et al., 2009]. Alternatively, the Coulomb field of the bunch can be exploited to trigger the emission of coherent radiation, the spectrum of which contains information about the current profile of the bunch [Lai, Happek and Sievers, 1994]. Different realisations of this frequency-domain technique are possible depending on the type of radiation utilised, including: coherent diffraction radiation (CDR) [Lockmann et al., 2020], coherent transition radiation (CTR) [Wesch et al., 2011; B. Schmidt et al., 2020], Smith-Purcell radiation [Andrews et al., 2014], synchrotron radiation (CSR) and others [Andonian et al., 2009; Curcio et al., 2020]. However, even if the different types of diagnostics mentioned so far are able to achieve time resolutions on the femtosecond scale, they only provide information about the current profile of the bunch. Transverse-deflection structures (TDS), in contrast, overcome this limitation and enable the retrieval of the sliced beam parameters.

The TDS imprints the longitudinal information of the bunch into a tranverse coordinate by means of time-varying transverse fields excited in a dedicated RF structure. Through an appropiate transport of the beam to a screen station located downstream of the device, its

current profile together with its sliced transverse dimension in either x or y can be directly measured. Additionally, this technique can be combined with conventional phase-space manipulations to retrieve further beam parameters as a function of time: by scanning the phase advance in a multi-shot measurement, the slice emittance in the plane perpendicular to the transverse deflection can be retrieved; in combination with a dispersive section the longitudinal phase space of the bunch becomes accessible in a single-shot basis.

Transverse-deflection structures were invented in the 60s in the context of particle-physics research [Altenmueller, Larsen and Loew, 1964] and gained popularity at accelerator-based light sources in the 2000s [Emma, Frisch and Krejcik, 2000]. Since then, they have been regularly used at FELs to diagnose the longitudinal-phase-space of electron bunches [Alesini et al., 2006; Röhrs et al., 2009; Craievich, Ischebeck et al., 2013] and technological developments have allowed to extend their operation to the X-band RF frequency range [Ding et al., 2011; Krejcik et al., 2013; Dolgashev et al., 2014], leading to a compact hardware design and an improvement of the time resolution down to the few-femtosecond level [Behrens, F.-J. Decker et al., 2014]. In the following, the principles of operation and the diagnostics capabilities common to all transverse-deflection structures are described.

1.3.1. Principles of operation of a TDS

Panofsky-Wenzel theorem

According to the Lorenz equation $\boldsymbol{F} = e(\boldsymbol{\mathcal{E}} + \boldsymbol{v} \times \boldsymbol{\mathcal{B}})$, in order to imprint a transverse momentum to the particles, the RF structure must contain either transverse electric fields, transverse magnetic fields or both of them, that are perpendicular to the direction of the beam propagation. An analysis of the resonant modes present in a circular cylindrical cavity—which is referred to as *pillbox* cavity—helps to understand which are the fundamental electromagnetic field configurations that can be used to this end.

In a pillbox cavity, there exist two types of azimuthally symmetric solutions for the 3D electric- and magnetic-field components: transverse magnetic modes[5] TM_{mnp} with $\mathcal{B}_z = 0$, and transverse electric modes TE_{mnp} with $\mathcal{E}_z = 0$ [Wangler, 2008]. From the analytical expressions of the TM and TE modes it can be seen that: 1) only the dipole modes with $m = 1$ have nonzero transverse fields at $r = 0$, and 2) most of these modes have both electric- and magnetic-field components in the tranverse plane. The question whether these fields add constructively or cancel out each other is answered by the Panosfky-Wenzel theorem [Panofsky and Wenzel, 1956], which provides a general expression for the transverse momentum gained by a bunch of relativistic particles passing through a cavity parallel to its axis excited either in a TM or a TE mode:

[5]The subscript m ($m = 0, 1, 2, ...$) is the number of full period variations in θ of the field components; the subscript n ($n = 1, 2, 3, ...$) is the number of zeros of the axial field component in the radial direction; and the subscript p ($p = 0, 1, 2, ...$) is the number of half period varitions in z of the fields [Wangler, 2008].

$$\Delta \boldsymbol{p}_{\perp} = \frac{je}{\omega_{RF}} \int_0^d \nabla_{\perp} \boldsymbol{\mathcal{E}}_z dz, \tag{1.3.1}$$

where ω_{RF} is the angular frequency of the field, and d is the length of the cavity. This expression has two important consequences. Firstly, it excludes TE modes from being used to deflect the particles, since they fulfil the constraint $\mathcal{E}_z = 0$. In the case of a TE mode, what actually happens is that the deflecting impulse of the magnetic field always cancels out that of the electric field, whereas in a TM mode the effect of the first multiplies that of the latter. Therefore only TM modes are appropriate to achieve transverse deflection. In second place, the requirement $\nabla_{\perp} \boldsymbol{\mathcal{E}}_z \neq 0$ implies the existence of a transverse gradient of the longitudinal field that induces an energy spread for beams with finite beam size—i.e., there is a trade-off for being able to deflect the particles. The fundamental mode typically used in a TDS shares key properties with the first transverse mangetic dipole mode TM_{110} of a pillbox cavity, which is schematically depicted in Figure 1.6.

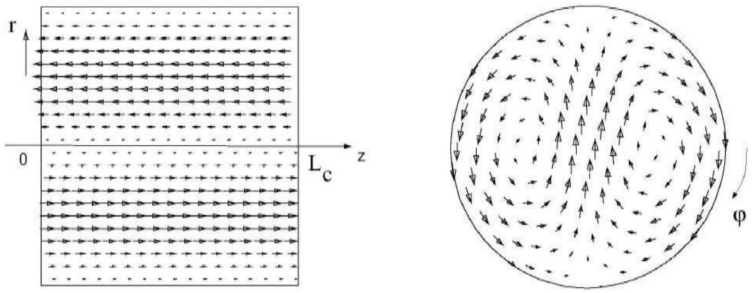

Figure 1.6.: Electric and magnetic fields of the TM_{110} mode in a pillox cavity. Figure reproduced from Röhrs, 2008.

Disk-loaded waveguide structures

In order to accrue a sufficiently large transverse momentum, a series of pillbox cavities coupled through radial apertures need to be arranged in a periodic structure. There are two ways in which such a structure can be conceived: 1) as a periodic array of coupled resonant cavities containing a standing wave (SW), or 2) as a waveguide loaded with a periodic arrangement of disks that reduce the phase velocity of a travelling wave (TW)—which in a uniform waveguide is always larger than the speed of light [Wangler, 2008]. A SW is obtained by simply closing both ends of the structure, whereas for TW operation the power is coupled at one of the outermost cells and is dumped into an RF load at the other one after travelling through the structure. One of the main differences between the two types of structures is the time required for the RF field to reach the steady state, which is referred to as *filling time* t_f. For a TW the filling time equals the structure length devided by the group velocity $t_f = L/v_g$ and is much faster than for a SW, which requires the field to slowly build up as the waves are reflected at the end-walls. Therefore, for short relativistic electron bunches as

those of interest in this thesis, TW structures are much more power efficient, since they can be operated with very short RF pulses and reach much higher peak fields compared to SW structures [Gerigk, 2012].

Hybrid modes

On account of the Floquet theorem and the boundary conditions imposed on a periodic structure, the propagating-wave solution can be expanded in a Fourier series in terms of the so-called *space harmonics*, each with its wavenumber $k_n = k_0 + 2\pi n/d$, where d is the length of a single cell [Wangler, 2008]. For a given RF frequency ω_{RF}, the phase velocity of the fundamental wave $\beta_0 = \omega_{RF}/ck_0$—which has the largest Fourier amplitude and is therefore the mode at which the structure is most often operated—is determined by the cell geometry, which is designed to fulfil the synchronocity condition and optimise the performance of the structure (cf. Section 3.1). The fundamental space harmonic in a TDS is a linear combination of the TM_{11} and TE_{11} modes and is referred to as *hybrid mode* HEM [Hahn, 1963; Paramonov, 2013]

Beam dynamics within a TDS

From the electromagnetic fields in a TDS excited in the fundamental HEM mode, it can be shown that the Lorentz force in cartesian coordinates experienced by a particle with charge e amid its passage through the device is [Altenmueller, Larsen and Loew, 1964]:

$$F_x = 0, \tag{1.3.2a}$$

$$F_y = e\mathcal{E}_0 \sin(\phi_{RF} + k_{RF}\xi), \tag{1.3.2b}$$

$$F_z = e\mathcal{E}_0 \cos(\phi_{RF} + k_{RF}\xi) k_{RF} y, \tag{1.3.2c}$$

where $\mathcal{E}_0$ is the amplitude of a travelling wave $\mathcal{E} = \mathcal{E}_0 e^{i(k_{RF}\xi - \omega_{RF}t)}$, ω_{RF} and k_{RF} are the angular frequency and the wavenumber correspondingly (with $\omega_{RF}/k_{RF} = c$), and ϕ_{RF} is the RF phase at the bunch center ($\xi = 0$) relative to the zero crossing of the field. For practical RF frequencies[6] and typical bunch lengths on the order of ~ 100 fs, the relation holds $\xi \ll \lambda_{RF}$, which justifies to approximate the Lorenz force by a Taylor series expansion around $k_{RF}\xi \approx 0$. By doing so, after integration from the start of the TDS to a location s inside the structure and assuming that the total momentum p_0 remains constant, the following expression for the accrued vertical kick is obtained:

[6] For the PolariX-TDS, for instance, with a frequency around 12 GHz and a wavelength of $\lambda_{RF} \approx 25$ mm, even long bunches of $\sigma_\xi \approx 1$ ps (0.3 mm) only see a small fraction of the field $\sigma_\xi/\lambda_{RF} \approx 1$ %.

$$y'(s) = y'_0 + \int_0^{p_y(s)} \frac{dp_y}{p_0} = y'_0 + \int_0^s \frac{1}{p_0} F_y \frac{ds}{c} = y'_0 + \frac{F_y}{p_0 c} s$$
$$\approx y'_0 + \frac{e\mathcal{E}_0 \left(k_{RF}\xi \cos\phi_{RF} + \sin\phi_{RF}\right)}{p_0 c} s. \tag{1.3.3}$$

As devised for applications in longitudinal beam diagnostics [Emma, Frisch and Krejcik, 2000], the TDS is operated at the zero crossings of the field—i.e., ϕ_{RF} being either 0 or π—so that the deflection is linearly dependent on the longitudinal particle coordinate ξ. After integration over the whole TDS length L, equation 1.3.3 becomes:

$$y'_{s=L} = y'_0 + \frac{e\mathcal{E}_0 L k_{RF}}{p_0 c}\xi \,, \tag{1.3.4}$$

from which the total displacement, $y(s) = y_0 + \int_0^L y'(s)ds$, can also be derived:

$$y_{s=L} = y_0 + y'_0 L + \frac{e\mathcal{E}_0 L^2 k_{RF}}{2p_0 c}\xi \tag{1.3.5}$$

Regarding the longitudinal force F_z, a particle travelling off-axis ($y \neq 0$) will experience a change in momentum that can be shown to be:

$$c\Delta p_z = \int_0^L F_z ds \approx e\mathcal{E}_0 \left(\cos\phi_{RF} - k_{RF}\xi \sin\phi_{RF}\right) k_{RF} \int_0^L y(s)ds$$
$$= e\mathcal{E}_0 k_{RF} \left(L y_0 + \frac{L^2}{2} y'_0 + \frac{e\mathcal{E}_0 L^3 k_{RF}}{6p_0 c}\xi \right) \tag{1.3.6}$$

On the other hand, the transverse deflecting force also imparts a change in momentum, which can be expressed as:

$$c\Delta p_y = F_y y_{s=L} \approx e\mathcal{E}_0 \left(\sin\phi_{RF} + k_{RF}\xi \cos\phi_{RF}\right) y_{s=L}$$
$$= e\mathcal{E}_0 k_{RF} \left(y_0\xi + L y'_0 \xi + \frac{e\mathcal{E}_0 L^2 k_{RF}}{2p_0 c}\xi^2 \right) \tag{1.3.7}$$

By comparing Equations 1.3.7 and 1.3.6 it can be seen that:

$$\frac{\Delta p_y}{\Delta p_z} \approx \frac{\xi}{L} \ll 1, \tag{1.3.8}$$

where in the last inequalty it has been considered that typically ξ is on the order of

$\mathcal{O}(100\ \mu\text{m})$, whereas $L \approx 1$ m. Therefore, to a very good approximation, the change in longitudinal momentum dominates the overall change of the particle momentum, so that $c\Delta p \approx c\Delta p_z$. With this assumption, and introducing the usual definition $\delta = (p - p_0)/p_0$, the equations that describe the dynamics of a charged particle through a TDS are obtained:

$$y_{s=L} = y_0 + Ly_0' + \frac{\mathcal{K}L}{2}\xi \tag{1.3.9a}$$

$$y_{s=L}' = y_0' + \mathcal{K}\xi \tag{1.3.9b}$$

$$\delta_{s=L} = \delta_0 + \mathcal{K}y_0 + \frac{\mathcal{K}L}{2}y_0' + \frac{\mathcal{K}^2L}{6}\xi, \tag{1.3.9c}$$

where the *integrated kick*:

$$\mathcal{K} = \frac{eV_0k_{RF}}{p_0c} \tag{1.3.10}$$

and the *integrated voltage*:

$$V_0 = \mathcal{E}_0L, \tag{1.3.11}$$

have been introduced. Equations 1.3.9 allow to construct the transport matrix for the TDS. In particular, for a streaking field in the vertical direction:

$$\boldsymbol{M}_{\mathbf{TDS}} = \begin{pmatrix} 1 & L & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & L & \frac{\mathcal{K}L}{2} & 0 \\ 0 & 0 & 0 & 1 & \mathcal{K} & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & \mathcal{K} & \frac{\mathcal{K}L}{2} & \frac{\mathcal{K}^2L}{6} & 1 \end{pmatrix}. \tag{1.3.12}$$

An approximation for $L \to 0$ yields the simplified transfer matrix:

$$\boldsymbol{M}_{\mathbf{TDS}}^{\text{thin}} = \begin{pmatrix} 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & \mathcal{K} & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & \mathcal{K} & 0 & 0 & 1 \end{pmatrix}, \tag{1.3.13}$$

which is referred to as *thin-TDS matrix*. The conditions in which the thin-TDS approximation is valid depend on the transport of the beam downstream of the TDS and will become clear in the next section, once the longitudinal resolution is defined.

1.3.2. Diagnostics capabilities of a TDS

Bunch length and longitudinal resolution

The transport of the beam matrix $\boldsymbol{\Sigma}$ (cf. Section 1.1.2) from the TDS to a screen—or to a position s downstream of the TDS in general—is described by the expression (cf. Eq. 1.1.40):

$$\boldsymbol{\Sigma_s} = (\boldsymbol{R_{TDS\to s}} \cdot \boldsymbol{M_{TDS}}) \cdot \boldsymbol{\Sigma_0} \cdot (\boldsymbol{R_{TDS\to s}} \cdot \boldsymbol{M_{TDS}})^T \tag{1.3.14}$$

where $\boldsymbol{\Sigma_0}$ is the beam matrix at the TDS entrance, $\boldsymbol{M_{TDS}}$ is the TDS matrix (cf. Equation 1.3.12), and $\boldsymbol{R_{TDS\to s}}$ is the transfer matrix from the TDS exit downstream to the screen, whose elements will be represented simply as R_{11}, R_{12}, ... , R_{66}. In the following, the thin TDS approximation is used.

According to Equation 1.3.14, for a particle distribution without couplings in the plane y–ξ, the beam size measured at the screen can be expressed as:

$$\langle y_s^2 \rangle = \underbrace{R_{33}^2\langle y_0^2\rangle + 2R_{33}R_{34}\langle y_0 y_0'\rangle + R_{34}^2\langle y_0'^2\rangle}_{\langle y_{s,\beta}^2\rangle \,=\, \sigma_{y,\beta}^2} + \underbrace{\mathcal{K}^2 R_{34}^2}_{\mathcal{K}R_{34} \,=\, S}\langle \xi_0^2\rangle \tag{1.3.15a}$$

$$\sigma_{y,s}^2 = \sigma_{y,\beta}^2 + S^2\sigma_\xi^2, \tag{1.3.15b}$$

where the brace gathers the terms corresponding to the natural beam size due to the betatron motion of the particles $\sigma_{y,\beta}$, and the new *shear parameter* $S = \mathcal{K}R_{34}$ has been introduced. Equations 1.3.15 show that the total beam size at the screen $\sigma_{y,s}$ is the convolution of the streaked beam size $S\sigma_\xi$ with the natural beam size $\sigma_{y,\beta}$. Therefore, in order to resolve the longitudinal structure of the beam at the screen, the streaked beam size has to dominate over the natural one. Rearranging terms in Equation 1.3.15b, this can be reformulated in the following way:

$$\sigma_{y,s} = S\sqrt{\left(\frac{\sigma_{y,\beta}}{S}\right)^2 + \sigma_\xi^2} = S\sqrt{R_\xi^2 + \sigma_\xi^2} \tag{1.3.16}$$

where the *longitudinal resolution* has been defined:

$$R_\xi := \frac{\sigma_{s,\beta}}{|S|}. \tag{1.3.17}$$

Since the shear parameter is proportional to the frequency of the RF field $S = R_{34}\mathcal{K} \propto f_{RF}$

(cf. Eq. 1.3.10), the time resolution depends on the frequency as $R_{\tilde{\xi}} \propto f_{RF}^{-1}$. Therefore, compared to the S-band RF frequency range ($\sim$ 3 GHz), the resolution of an X-band TDS is improved by a factor of 4, leading to a reduction of the required length of the structure.

Lattice optic for optimal longitudinal resolution

If the matrix element R_{34} contained in the shear parameter $S = \mathcal{K}R_{34}$ is expressed in terms of the Courant-Snyder parameters (cf. Eq. 1.1.31), the longitudinal resolution can be shown to be:

$$R_{\tilde{\xi}} := \frac{\sigma_{y,\beta}}{|S|} = \frac{\sqrt{\beta_y(s)\epsilon_y}}{|\mathcal{K}|\sqrt{\beta_y(s)\beta_y(s_0)}|\sin\Delta\psi_y(s,s_0)|} = \frac{1}{|\mathcal{K}|}\sqrt{\frac{\epsilon_y}{\beta_y(s_0)}}\frac{1}{|\sin\Delta\psi_y(s,s_0)|}, \tag{1.3.18}$$

where ϵ_y is the geometric emittance of the beam, $\beta_y(s_0)$ the beta function at the location of the TDS, and $\Delta\psi_y(s,s_0) = \psi_y(s) - \psi_y(s_0)$ the phase advance from the TDS to the screen. Equation 1.3.18 indicates that to maximise the resolution, the phase advance should be close to $\Delta\psi_y(s,s_0) \approx 90$ deg and the beta function at the TDS $\beta_y(s_0)$ should be as large as possible. The maximum achievable integrated kick $\mathcal{K}$ is given by the technical specifications of the RF system, and the emittance ϵ_y by the quality of the beam delivered. The beta function at the screen $\beta_y(s)$ does not influence the longitudinal resolution and can be set according to the technical specifications of the diagnostic: the lower bound is given by the resolution of the optical system and the upper bound by its field-of-view. In the end, in most of the cases the condition $\beta_y(s) \ll \beta_y(s_0)$ is fulfilled, which can be used to better understand the nature of an optimal TDS lattice optic, as discussed in the following.

The transfer matrix elements relevant to the beam size at the screen are R_{33} and R_{34}, which for a phase advance of $\Delta\psi_y(s,s_0) \approx 90$ deg are simplified to the expressions:

$$R_{33} = \sqrt{\frac{\beta_y(s)}{\beta_y(s_0)}}\alpha_y(s_0) \qquad \text{, and} \qquad R_{34} = \sqrt{\beta_y(s)\beta_y(s_0)}, \tag{1.3.19}$$

where $\alpha_y(s_0)$ is the alpha function at the location of the TDS. From the discussion about the beta function at the screen, it derives that generally $\beta_y(s)/\beta_y(s_0) \ll 1$. Additionally, at the location of the TDS, an alpha function close to zero $\alpha_y(s_0) \approx 0$ is tipically chosen to avoid changes in beam size along the device. These two conditions guarantee that to a good approximation $R_{33} \approx 0$ and hence the resulting matrix is that of a first-order parallel-to-point imaging system in the vertical plane [Brown, 1982]:

$$\boldsymbol{R_y} = \begin{pmatrix} 0 & R_{34} \\ R_{43} & R_{44} \end{pmatrix}, \tag{1.3.20}$$

in which information about the original transverse location of the particles is removed and only their kicks are imaged with a magnification of R_{34}. This is consistent with the discussion about the beam dynamics within a TDS, where it has been shown that the bunch length information is carried by the ξ-dependent kick y' imprinted by the TDS.

Calibration of the shear parameter

The transport of a particle vector $\boldsymbol{X}$ from the TDS to the screen is described by the expression $\boldsymbol{X_s} = \boldsymbol{R_{TDS\to s}} \cdot \boldsymbol{M_{TDS}} \cdot \boldsymbol{X_0}$. According to that, when considering a particle beam, the centroid in the streaking direction measured at the screen is obtained:

$$\langle y_s \rangle = \underbrace{R_{33}\langle y_0 \rangle + R_{34}\langle y_0' \rangle}_{\langle y_{s,\beta} \rangle} + S\langle \xi_0 \rangle, \tag{1.3.21}$$

where the brace gathers the terms corresponding to the natural betatron motion of the transverse centroids $\langle y_{s,\beta} \rangle$, and $\langle \xi_0 \rangle$ is the longitudinal location of the beam centroid relative to the zero crossing of the streaking field. For small displacements around the zero crossing much smaller than the RF wavelength $\Delta\xi \ll \lambda_{RF}$, the RF phase varies linearly with the longitudinal coordinate $\Delta\xi \approx \Delta\phi_{RF}/k_{RF}$, and the beam centroid measured at the screen can be expressed as:

$$\langle y_s \rangle = \langle y_{s,\beta} \rangle + \frac{S}{k_{RF}} \cdot \phi_{RF}. \tag{1.3.22}$$

The shear parameter S is calculated by performing measurements of the beam centroid at different RF phases around the zero crossing and applying a linear regression to the data according to Eq. 1.3.22. Once the shear parameter has been calibrated, the longitudinal resolution R_ξ is obtained by measuring the natural (unstreaked) beamsize of the bunch at the screen $\sigma_{s,\beta}$ and using Equation 1.3.17.

Current profile and two-point-tomography

Considering a particle beam, the transport of its slice centroids from the TDS to the screen is obtained by means of the usual particle-transport equation $\boldsymbol{X_s} = \boldsymbol{R_{TDS\to s}} \cdot \boldsymbol{M_{TDS}} \cdot \boldsymbol{X_0}$ and taking the first moments of thin slices in ξ:

$$\langle y_s \rangle_\xi = \langle y_{s,\beta} \rangle_\xi + S\xi, \tag{1.3.23}$$

where the subscript ξ identifies the slice at a particular longitudinal coordinate of the beam and $\langle y_{s,\beta} \rangle_\xi$ corresponds to the betatron-motion term containing the beam correlations.

Assuming that the longitudinal resolution achieved with the TDS is large enough that the minimum resolvable slice is practically infinitessimal compared to the total length of the bunch, the beam can be considered to be a single filament in phase space with an associated 1D charge-density distribution varying along the longitudinal coordinate $\lambda(\xi)$. In that case, Eq. 1.3.23 can be expressed as a mapping function μ from the longitudinal beam coordinate ξ to the screen coordinate y:

$$\mu: \quad \xi \mapsto y(\xi) = y_\beta(\xi) + S\xi, \tag{1.3.24}$$

where the subscript s has been dropped, since y always refers to the screen—whereas ξ always refers to the longitudinal coordinate of the beam at the TDS. If the beam does not contain correlations in the ξ–y and ξ–y' planes, the betatron motion term equals zero $y_\beta(\xi) = 0$ and the vertical bunch coordinate at the screen can be simply mapped back to the longitudinal coordinate of the beam $\xi(y) = y/S$. Thus, the original charge density $\lambda(\xi)$ is directly related to the projected charge density measured at the screen $\lambda_y(y)$ by means of the expression:

$$\lambda_y(y) = \int \lambda(\xi)\ \delta(y - S\xi)\ d\xi = \frac{\lambda(\xi(y))}{|S|} \qquad \longrightarrow \qquad \lambda(\xi) = \lambda_y(y(\xi))|S|, \tag{1.3.25}$$

where $\delta(y - S\xi)$ is the δ-function evaluated at $y = S\xi$ and the generalised scaling property[7] of the δ-function has been used in the integration.

In general, however, since the beam does exhibit correlations in both planes y–ξ and y'–ξ, the mapping from the TDS to the screen $\mu(\xi)$ (cf. Eq. 1.3.24) becomes a complicated function of ξ modulated by the unknown term $y_\beta(\xi)$ and the measured current profile cannot be simply mapped back to the real one. Moreover, this function will affect each zero crossing differently, because the sign of the streak S is inverted when the RF phase is shifted 180 deg, whereas the betatron-motion term does not change.

By combining the measurements performed at the two zero crossings, the distortions associated to the bunch correlations $y_\beta(\xi)$ can be compensated and the real longitudinal current profile can be reconstructed. Inspired by similar problems encountered in the longitudinal-phase-space reconstruction based on the RF zero-phasing technique [Bane, 1990; Huang et al., 2010], this method was proposed by Paul Emma in [Emma, Frisch and Krejcik, 2000] and was first privately noted down in detail by Henrik Loos in [Loos, n.d.]. The method, commonly referred to as *two-phase averaging* or *two-point-tomography*, is illustrated in Figure 1.7 and is described in the following.

[7] The generalised scaling property of the δ-function in integral form states that:

$$\int_{-\infty}^{\infty} f(x)\ \delta(g(x))\ dx = \frac{f(x_0)}{|g'(x_0)|}, \tag{1.3.26}$$

where x_0 are the zeros of $g(x)$—i.e., x_0 is the inverse function of $g(x)$.

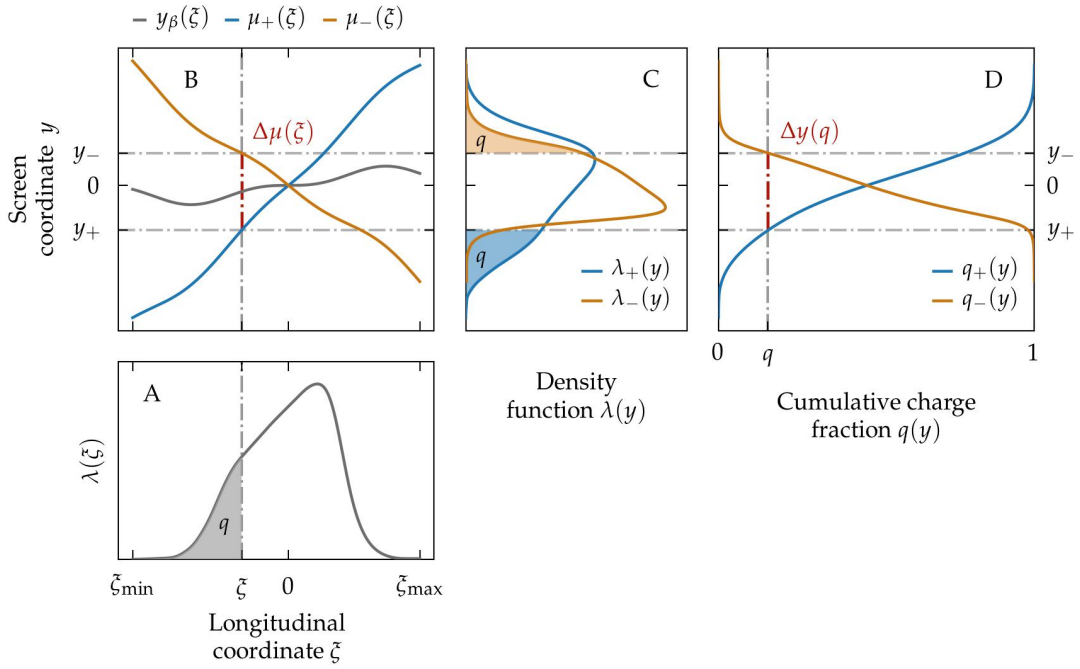

Figure 1.7.: Schematic representation of the two-point-tomography technique. A) original (longitudinal) charge density distribution. B) intrinsic bunch correlations $y_\beta(\xi)$ (grey line) and mapping functions $\mu_\pm(\xi)$ for positive (blue line) and negative (orange line) shear parameters. C) resulting charge density distributions projected at the screen $\lambda_\pm(y)$. D) cumulative charge fractions as a function of the screen coordinate $q_\pm(y)$.

The original charge density $\lambda(\xi)$ is shown in Fig. 1.7A and the bunch-correlation term $y_\beta(\xi)$ is represented by the grey line in Fig. 1.7B. Depending on the sign of the streak—i.e., on the zero crossing—two different mapping functions μ are obtained (blue and orange lines in Figure 1.7B):

$$\mu_\pm(\xi) = y_\beta(\xi) + S_\pm\xi, \tag{1.3.27}$$

which produce two different charge-density functions at the screen (Figure 1.7C):

$$\lambda_\pm(y) = \int \lambda(\xi)\ \delta(y - \mu_\pm(\xi))\ d\xi, \tag{1.3.28}$$

where $\delta(y - \mu_\pm(\xi))$ is the δ-function evaluated at $y = \mu_\pm(\xi)$. For given charge-density functions at the screen $\lambda_\pm$, their cumulative charge fractions are defined as:

$$q_+(y) = \int_{-\infty}^{y} \lambda_+(y')dy' \tag{1.3.29a}$$

$$q_-(y) = \int_{y}^{\infty} \lambda_-(y')dy', \tag{1.3.29b}$$

where the different integration limits reflect the fact that the two zero crossings result in inverted mappings from ξ to y—i.e., the particles need to be summed up starting from the same end of the bunch. This is shown in Figure 1.7D, where the functions $q_{\pm}(y)$ evaluated at $y_{\pm}$ correspond to the shaded areas below the functions $\lambda_{\pm}(y)$ in Fig. 1.7C.

There are two fundamental conditions for the applicability of the reconstruction method: 1) the streak S is large enough that the function μ is injective and, therefore, invertible:

$$\mu^{-1} : \; y \mapsto \xi(y); \tag{1.3.30}$$

2) the original charge density profile is always positive along the extent of the particle distribution—i.e. no double-bunch structure is allowed:

$$\lambda(\xi) > 0 \,, \quad \text{for} \;\; \xi_{\min} < \xi < \xi_{\max}. \tag{1.3.31}$$

When these two conditions are fulfilled, the cumulative charge fractions are strictly monotonic and can be inverted $y_{\pm}(q) = q_{\pm}(y)^{-1}$. The reconstruction is based on the fact that the value of the difference $\Delta y(q) = y_{+}(q) - y_{-}(q)$ for a given q is identical to the value of the difference of the streaked correlation functions $\mu_{\pm}$ at a particular ξ (cf. vertical dash-dotted red lines in Figures 1.7B and 1.7D):

$$\Delta\mu(\xi) = \mu_{+}(\xi) - \mu_{-}(\xi) = (S_{+} - S_{-})\xi. \tag{1.3.32}$$

This allows to uniquely correlate q with ξ by setting $\Delta y(q) = \Delta\mu(\xi)$ and solving for ξ:

$$\xi(q) = \frac{y_{+}(q) - y_{-}(q)}{S_{+} - S_{-}}. \tag{1.3.33}$$

By inverting the function $\xi(q)$, the fractional charge is obtained $q(\xi) = \xi(q)^{-1}$ as a function of ξ (shaded grey area in Fig. 1.7A). Finally, the original charge density is the derivative of that function:

$$\lambda(\xi) = \frac{dq(\xi)}{d\xi}. \tag{1.3.34}$$

In a similar way, the (unstreaked) beam correlation $y_{\beta}(q)$ as a function of q can be retrieved by eliminating the $S_{\pm}\xi$ terms in Eq. 1.3.27 and setting $\mu_{\pm} = y_{\pm}(q)$:

$$y_{\beta}(q) = \frac{S_{+}y_{-}(q) - S_{-}y_{+}(q)}{S_{+} - S_{-}}, \tag{1.3.35}$$

which, on account of the relation $q = q(\xi)$ previously obtained, can be expressed as a

function of ξ. Since now the mapping functions $\mu_\pm(\xi)$ become explicit functions of ξ, it is possible to map any slice beam parameter $f_\pm(y)$ measured at the screen at either of the two zero crossings to the original longitudinal coordinate of the bunch by setting $y = \mu_\pm(\xi)$:

$$f(\xi) = f_\pm(\mu_\pm(\xi)). \tag{1.3.36}$$

Longitudinal phase space

In a diagnostics beamline, a TDS is typically combined with a dipole magnet spectrometer operating in the plane perpendicular to the streak. In this way, the longitudinal bunch coordinate ξ and its energy δ are mapped into the vertical y_s and horizontal x_s screen coordinates respectively, and the longitudinal phase space density of the bunch can be retrieved from the measured images. In the following, the (independent) operation of a dipole spectrometer is briefly summarised.

For an uncoupled beamline, the transfer matrix of a dispersive section operating in the x–s plane has the following form:

$$\boldsymbol{R}_{\boldsymbol{D_x}} = \begin{pmatrix} R_{11} & R_{12} & 0 & 0 & 0 & R_{16} \\ R_{21} & R_{22} & 0 & 0 & 0 & R_{26} \\ 0 & 0 & R_{33} & R_{34} & 0 & 0 \\ 0 & 0 & R_{43} & R_{44} & 0 & 0 \\ R_{51} & R_{52} & 0 & 0 & 1 & R_{56} \\ 0 & 0 & 0 & 0 & 0 & 1 \end{pmatrix}, \tag{1.3.37}$$

where the elements R_{16} and R_{26} correspond to the dispersion functions D_x and $D_{x'}$ (cf. Eq. 1.1.14). According to the transport of the beam matrix through such a dispersive section downstream to a position s where the screen is located, $\boldsymbol{\Sigma_s} = \boldsymbol{R}_{\boldsymbol{D_x}} \cdot \boldsymbol{\Sigma_0} \cdot \boldsymbol{R}_{\boldsymbol{D_x}}^T$, the measured beam size can be expressed as:

$$\langle x_s^2 \rangle = \underbrace{R_{11}^2 \langle x_0^2 \rangle + 2R_{11}R_{12}\langle x_0 x_0' \rangle + R_{12}^2 \langle x_0'^2 \rangle}_{\langle x_{s,\beta} \rangle \,=\, \sigma_{x,\beta}^2} \;+\; R_{16}^2 \langle \delta_0^2 \rangle \tag{1.3.38a}$$

$$\sigma_x^2 \;=\; \sigma_{x,\beta}^2 \;+\; D_x^2 \sigma_E^2, \tag{1.3.38b}$$

where the curly bracket gathers the terms corresponding to the natural beam size due to the betatron motion of the particles. Equations 1.3.38 show that the total beam size at the screen σ_x is the convolution between the natural beamsize $\sigma_{x,\beta}$ of the bunch and the dispersive term due to the energy spread $D_x\sigma_E$. Reasoning along the same lines as with the derivation of the longitudinal resolution (cf. Eq. 1.3.17 and preceeding discussion), in order to be able to

measure the energy spectrum of the bunch, the dispersed beam size has to dominate over the natural beam size, which leads to the definition of *energy resolution*:

$$R_\delta := \frac{\sigma_{x,\beta}}{|D_x|}. \tag{1.3.39}$$

In contrast to the longitudinal resolution, the energy resolution R_δ cannot be directly measured, because the dispersive elements in the beamline cannot be simply removed and the natural betatron beamsize $\sigma_{x,\beta}$ at the screen is not directly accessible. Therefore it has to be derived from knowledge about both the beam parameters and the lattice optic. The beam parameters must be measured upstream of the dispersive section and a good degree of matching to the lattice optic is required to be able to estimate their value at the screen. The dispersion at the screen is typically measured by scanning the dipole-spectrometer current.

Experimentally, the screen resolution R_{scr} has to be taken into account and the beam size σ_M effectively measured at the dispersive screen can be expressed as:

$$\begin{aligned} \sigma_M^2 &= D_x^2\sigma_E^2 + \sigma_{x,\beta}^2 + R_{scr}^2 \\ &= D_x^2\left(\sigma_E^2 + R_\delta^2\right) + R_{scr}^2. \end{aligned} \tag{1.3.40}$$

When the TDS is switched on to access the longitudinal phase space of the bunch, the transverse gradient of the longitudinal field $\boldsymbol{\nabla}_\perp \boldsymbol{\mathcal{E}}_z \neq 0$ associated to the transverse-deflection dipole mode—cf. Panofsky-Wenzel theorem, Eq. 1.3.1—induces several effects that must be taken into account. These effects and some of their most important consequences are discussed in the following.

Induced energy spread

The energy spread induced by a TDS can be calculated just by propagating the beam matrix through the TDS according to Eq. 1.3.14—using in this case the complete TDS matrix, since the thin approximation is not enough to account for all effects. Assuming an uncoupled beamline, this results in the following expression:

$$\begin{aligned} \langle\delta^2\rangle = \mathcal{K}^2\langle y_0^2\rangle + \mathcal{K}^2 L\langle y_0 y_0'\rangle + \left(\frac{\mathcal{K}L}{2}\right)^2\langle y_0'^2\rangle \\ + \left(\frac{\mathcal{K}^2 L}{6}\right)^2\langle\xi_0^2\rangle + \frac{\mathcal{K}^2 L}{3}\langle\delta_0\xi_0\rangle + \langle\delta_0^2\rangle. \end{aligned} \tag{1.3.41}$$

Taking the limit to an infinitesimal thin slice in ξ, so that $\langle \xi_0^2 \rangle \to 0$ and $\langle \delta_0 \xi_0 \rangle \to 0$, one becomes the following expression for the induced slice energy spread σ_{IES}:

$$\begin{aligned} \langle \delta^2 \rangle_\xi - \langle \delta_0^2 \rangle_\xi &= \mathcal{K}^2 \left(\langle y_0^2 \rangle_\xi + L \langle y_0 y_0' \rangle_\xi + \frac{L^2}{4} \langle y_0'^2 \rangle_\xi \right) \\ &= \mathcal{K}^2 \epsilon_y(\xi) \left(\beta_y(s_0, \xi) - L\alpha_y(s_0, \xi) + \frac{L^2}{4} \gamma_y(s_0, \xi) \right) \\ &= \sigma_{IES}^2(\xi) \end{aligned} \tag{1.3.42}$$

which is independent of the zero crossing at which the TDS is operated because of the term $\mathcal{K}^2$. Therefore, when measuring the slice energy spread of the bunch at the screen, the TDS-induced energy spread must be added into Eq. 1.3.40:

$$\sigma_M^2 = D_x^2 \left(\sigma_E^2 + R_\delta^2 + \sigma_{IES}^2 \right) + R_{scr}^2. \tag{1.3.43}$$

Fundamental resolution limit

In order to achieve a good longitudinal resolution, the beta function $\beta_y(s_0)$ at the location of the TDS typically has large values and dominates over the other two Courant-Snyder parameters $\alpha_y(s_0)$ and $\gamma_y(s_0)$. Hence, in general, to a good approximation the induced energy spread is given by:

$$\sigma_{IES} \approx \mathcal{K} \sqrt{\epsilon_y \beta_y(s_0)} = \mathcal{K} \sigma_y(s_0). \tag{1.3.44}$$

Multiplying this equation by the longitudinal resolution, the following expression results:

$$R_\xi \cdot \sigma_{IES} \approx \frac{\epsilon_y}{|\sin \Delta\psi_y(s, s_0)|}, \tag{1.3.45}$$

which asserts that there is a fundamental limit to the resolvable area of the longitudinal phase space, given by the geometric emittance and the phase advance from the TDS to the screen [Behrens and Gerth, 2009]. If the phase advance fulfils the condition $\Delta\psi_y(s, s_0) \approx 90$ deg, the expression can be simplified to $R_\xi \cdot \sigma_{IES} \approx \epsilon_y$.

In the following, a numerical estimation of the energy spread induced by the PolariX-TDS at FLASHForward is described. The TDS operates at a frequency of $f_{RF} = 11.988$ GHz and it has a wavenumber of $k_{RF} \approx 251$ m^{-1}. If the RF system delivers an effective voltage of $V_0 = 10$ MV, an electron beam with an energy of $E = 1$ GeV experiences an integrated kick of $\mathcal{K} \approx 2.51$ m^{-1}. Assuming a normalized emittance of $\epsilon_{n,y} = 1$ mm mrad (being the geometric emittance $\epsilon_y = 5.11 \cdot 10^{-10}$ m rad) and a beta function at the TDS of $\beta_y = 100$ m,

the induced energy spread has a value of $\sigma_{IES} = 0.57 \cdot 10^{-3}$, which sets a lower bound to the measurable energy spreads. The dispersion at the location of the longitudinal-phase-space screen is typically about $D_x = 0.7$ m. For beta functions in the dispersive plane of $\beta_x \lesssim 10$ m, energy resolutions of $R_\delta \lesssim 0.1 \cdot 10^{-3}$ are achieved, which are enough to resolve the induced energy spread. According to Equation 1.3.45, the longitudinal resolution achievable for such a beam would be $R_\xi \approx 0.9$ µm (3 fs).

Induced energy chirp

In the absence of bunch correlations in the y–ξ plane, the effect of the TDS on the beam matrix term correlating the longitudinal coordinate to the energy is obtained from the usual beam-transport equation $\boldsymbol{\Sigma_s} = \boldsymbol{M_{TDS}} \cdot \boldsymbol{\Sigma_0} \cdot \boldsymbol{M_{TDS}^T}$ and can be expressed as:

$$\langle \xi_s \delta_s \rangle = \frac{\mathcal{K}^2 L}{6} \langle \xi_0^2 \rangle + \langle \xi_0 \delta_0 \rangle, \tag{1.3.46}$$

which, on account of the chirp definition introduced in Eq. 1.1.38 and assuming that the bunch length is not affected along its passage through the TDS so that $\langle \xi_0^2 \rangle = \langle \xi_s^2 \rangle$, can be rewritten as:

$$\frac{\langle \xi_s \delta_s \rangle}{\langle \xi_s^2 \rangle} = \frac{\mathcal{K}^2 L}{6} + \frac{\langle \xi_0 \delta_0 \rangle}{\langle \xi_0^2 \rangle} \quad \longrightarrow \quad h_s = h_{TDS} + h_0 \,, \tag{1.3.47}$$

where the *induced energy chirp* $h_{TDS} = \mathcal{K}^2 L/6$ has been defined, which is independent of the zero crossing at which the TDS is operated due to the factor $\mathcal{K}^2$.

The statistical definition of the chirp used in Eq. 1.3.47, however, does not capture the complexity that the longitudinal phase space might exhibit—due, for instance, to non-linearities in the imprinted chirp or to collective effects—and it is more convenient to analyse the sliced mean energy of the bunch after the TDS. Taking the first moments of thin slices along ξ and considering correlations in the y–ξ and y'–ξ planes, the following expression is obtained from the usual particle-transport equation $\boldsymbol{X_s} = \boldsymbol{M_{TDS}} \cdot \boldsymbol{X_0}$:

$$\langle \delta_s \rangle_\xi = \langle \delta_0 \rangle_\xi + \mathcal{K} \langle y_0 \rangle_\xi + \frac{\mathcal{K} L}{2} \langle y_0' \rangle_\xi + \frac{\mathcal{K}^2 L}{6} \xi. \tag{1.3.48}$$

Equation 1.3.48 shows that, if bunch correlations in the y–ξ and y'–ξ planes exist, the slice centroids in the y–y' plane are amplified by the streaking field and change the slice mean energy of the bunch along its passage through the TDS. Additionally, the effect inverts its sign with the change of zero crossing, whereas the induced energy chirp h_{TDS}—i.e., last term on the r.h.s. of Eq. 1.3.48—does not. This TDS-induced changes of the sliced-mean-energy profile are similar to those affecting the current profile discussed previously. A method to retrieve the real sliced mean energy of the bunch based on the application, among others, of

Eq. 1.3.36 is described in Chapter 3.

Slice emittance

The methods to measure the projected emittance discussed in Section 1.2 can be applied simultaneously to the TDS operation in order to retrieve the *slice emittance* of the bunch. In this case, the transfer matrices have to fulfil the constriants for the emittance measurement between the reference plane and the screen in one direction, while fulfilling the ones for the TDS measurement between the TDS and the screen in the other. Since the new PolariX-TDS allows to adjust the polarisation of the streaking field, the slice emittance in both x and y can be measured. Examples of this using either the multiple-quadrupole-scan or the object-plane-scan approaches are disscused in Sections 3.4 and 4.4.2, respectively.

1.4. Beam-driven plasma-wakefield acceleration (PWFA)

This section provides the basic concepts of plasma-based acceleration required to understand the work presented in this thesis. Section 1.4.1 gives an account of the general definition of plasmas, some of their distinctive features, their mathematical description and their production. The properties of the wakefields that are excited when a relativistic electron bunch is focused into a plasma are analysed in Section 1.4.2 with regard to their suitability for acceleration purposes. Important aspects about the external injection of electron beams into a blown out plasma wake are discussed in Section 1.4.3. Finally the requirements that particle-beam-driven plasma-wakefield accelerators (PWFA) have to fulfil to rival the state-of-the-art RF-based technology are described together with the experimental progress made during the last decades.

1.4.1. Definition, properties and description of plasmas

A plasma can be defined as a *quasineutral* gas of charged (and neutral) particles which exhibits *collective* behaviour [F. F. Chen, 1984]. This means that an ionised gas per se is not a plasma, for it is not guaranteed that the two properties highlighted in the definition always apply. In the following paragraphs the meaning of collective behaviour and quasineutrality is elucidated and some of the key parameters characterising the plasma properties are briefly discussed.

Debye length. The high mobility of the unbound charges of a plasma naturally tend to shield the effect of an electric potential immersed in it, thus isolating the rest of the plasma from its influence. Once the plasma electrons have rearranged around the perturbation and have shielded its effect, the electrostatic potential is not described by the

usual Coulomb potential in vacuum, but it decays asymptotically according to the expression: $\Phi \sim q/r \exp(-r/\lambda_D)$, where λ_D is the characteristic decay length and is referred to as *Debye length*. The Debye length is defined as:

$$\frac{1}{\lambda_D^2} = \frac{1}{\lambda_{D,e}^2} + \sum_i \frac{1}{\lambda_{D,i}^2}, \tag{1.4.1}$$

where:

$$\lambda_{D,e} = \sqrt{\frac{\epsilon_0 k_B T_e}{e^2 n_e}} \quad , \qquad \lambda_{D,i} = \sqrt{\frac{\epsilon_0 k_B T_i}{q_i^2 n_i}} \tag{1.4.2}$$

are the Debye lengths of the plasma electrons and ions correspondingly, with ϵ_0 being the vacuum permitivity, k_B the Boltzmann's constant, e the electron charge, q_i the net charge of the ions of species i, and $T_{e,i}$ and $n_{e,i}$ the temperature and particle density of electrons and ions correspondingly. The condition of quasineutrality requires the spatial extent of the plasma to be much larger than the Debye length $L \gg \lambda_D$, so that the neutrality of the overall system is preserved regardless of local perturbations resulting from the Debye shielding. For large plasma temperatures, the strong thermal motion can allow the particles to escape the potential well, thus counteracting the shielding effect. When such thermal effects are negligible, the system is referred to as a *cold plasma*, which is the case of interest for applications to plasma-based acceleration. In the framework of this thesis, the plasma ions can be considered to be stationary compared to the much lighter plasma electrons, which rearrange their positions comparatively instantaneously to create net negative or positive space charges. It is therefore assumed that $\lambda_D \approx \lambda_{D,e}$.

Plasma parameter. A quantity related to the Debye length is the *plasma parameter*, which accounts for the number of particles contained within the Debye sphere:

$$\Lambda = \frac{4\pi}{3} n_e \lambda_D^3, \tag{1.4.3}$$

where n_e is electron-plasma density and λ_D its Debye length. If the Debye sphere is sparsely populated ($\Lambda \ll 1$) the motion of the particles is dominated by individual scattering events. In these conditions, collective behaviour is prevented and, strictly speaking, such a system cannot be considered a plasma. For the shielding to be effective—and, therefore, quasineutrality to be guaranteed—the Debye sphere must be dense ($\Lambda \gg 1$) so that collective Coulomb interactions dominate over binary particle collisions. Such a collisionless plasma is referred to as an *ideal plasma*.

Plasma frequency. On a length scale shorter than λ_D, Debye shielding does not occur and a slab of plasma electrons displaced from its equilibrium position will experience a restoring

force produced by the (immobile) positive ion background and the excess of negative charges at the displacement position, thus initiating an oscillatory motion. Considering this motion in 1D, the frequency of the oscillations can be shown to be [F. F. Chen, 1984]:

$$\omega_p = \sqrt{\frac{n_e e^2}{\epsilon_0 m_e}}, \tag{1.4.4}$$

which is referred to as the *plasma frequency* and defines the time scale on which collective effects take place. This is exactly the time scale required for the plasma to recover quasineutrality—i.e., to achieve an effective Debye shielding. A third condition for an ionised gas to be considered a plasma is that the plasma frequency must be much larger than the collision rate between electrons and ions (or neutral atoms). If τ is the mean time between collisions, this can be expressed as:

$$\omega_p \tau \gg 1. \tag{1.4.5}$$

Skin depth. The dispersion relation for transverse electromagnetic waves propagating in a plasma is given by the expression [F. F. Chen, 1984]:

$$\omega^2 = \omega_p^2 + c^2 k^2, \tag{1.4.6}$$

where ω and k are the angular frequency and the wavenumber of the incident wave, respectively, and ω_p is the plasma frequency defined previously. Considering a wave with a fixed frequency ω, if the plasma density is progressively increased, the wavenumber will be reduced according to Eq. 1.4.6 until it becomes zero. This cuttof condition occurs when $\omega = \omega_p$, at a *critical density* given by the expression $n_c = \omega^2 m_e / 4\pi e^2$. For larger densities, k becomes imaginary and the electromagnetic wave, which cannot propagate through the plasma, is attenuated. Plasmas with densities above and below the critical density are called *over-dense* and *under-dense plasmas*, respectively. Since k^{-1} characterises the length over which the shielding of low frequency radiation takes place, it is typically referred to as skin depth—in analogy to the skin effect observed at the inner surface of RF-accelerator cavities/structures exposed to electromagnetic fields [Wangler, 2008]. For oscillations at the plasma frequency, the *plasma skin depth* is obtained:

$$k_p^{-1} = \frac{\lambda_p}{2\pi} = \frac{c}{\omega_p}, \tag{1.4.7}$$

where the *plasma wavelength* λ_p has been introduced. Typically, in plasma-based accelerators, the plasma skin depth k_p^{-1} is used to characterise the length scale of the acceleration process. For an electron-plasma density of $n_e = 10^{16}$ cm^{-3}, for instance, $k_p^{-1} \approx 50$ µm.

Wave-breaking limit. The ultimate goal of a plasma-based accelerator is to excite a large plasma wake producing fields that can be used to accelerate particles. A useful estimate of the electric fields achievable in such a wake is obtained by considering a 1D plane wave—as in the case considered previously for the derivation of the plasma frequency—producing complete charge separation. Mathematically, this is derived by solving Gauss's law $\nabla \cdot E = \rho/\epsilon_0$, where the charge density is given by the electron-plasma density $\rho = n_e\, e$, and reduces to the expression:

$$E_0 = \frac{m_e c \omega_p}{e} = \sqrt{\frac{m_e c^2 n_e}{\epsilon_0}}, \tag{1.4.8}$$

which is referred to as the *(cold nonrelativistic) wave-breaking limit* [Dawson, 1959], since plasma waves with larger amplitudes than those achieved at this point will tend to destroy the oscillations. For a plasma density of $n_e = 10^{16}$ cm^{-3} a wave-breaking limit of $E_0 \approx 10$ GV/m is obtained, which illustrates the large field gradients that can be potentially used for particle acceleration.

Mathematical description of plasmas. There are different approaches to mathematically describe the behaviour of a plasma, which are briefly summarised below.

- **Microscopic picture.** The most elementary approach is to consider the motion of each of the particles contained in the plasma, including their collisions and long-range Coulomb interactions. This is achieved by means of the Klimontovich equation coupled with Maxwell's equations [Mehrling, 2014]. While this approach provides an exact classical description of the system, it is computationally very costly and of little practical use—although it serves as a starting point for the kinetic picture.
- **Kinetic picture.** Since a plasma is dominated by collective effects, it is more useful to consider the evolution of the statistical averages of the particle distribution, which is the approach followed by the kinetic description of plasmas. In this case, the evolution of the distribution is governed by the Vlasov equation, which, together with Maxwell's equations, forms a closed set of equations that can be solved numerically [Mehrling, 2014].
- **Macroscopic (fluid) picture.** A third, more simplified approach is to treat the plasma as a fluid by considering the evolution of its macroscopic variables—i.e., its density and velocity—which are then coupled to Maxwell's equations. This approach is particularly valid for a cold plasma, when the thermal motion of electrons can be neglected. Additionally, in the time scale of the acceleration processes of interest for plasma-based acceleration, the ions can be considered to form a static background for most experimental cases. The last two approximations constitute the basic assumptions of the *cold electron-fluid model*, which provides an accurate description of longitudinal plasma waves in the linear and mildly nonlinear regimes (cf. Section 1.4.2). The equations describing the system are the continuity equation for the electron-plasma density and

the fluid momentum equation [Mehrling, 2014], which when coupled to Maxwell's equations form a closed system of equations that, in particular cases, can be solved analytically.

Plasma sources. Plasmas are formed via ionisation of a gas by means of two basic mechanisms: a) collisional ionisation, by which fast energetic particles collide with the atoms; b) photoionisation by electromagnetic radiation. Photoionisation can be achieved in different ways, among others, with the strong fields of an intense laser pulse—which can operate in different regimes, leading to single-photon, multi-photon, tunneling or above-threshold ionisation [Krausz and Ivanov, 2009]—or with the Coulomb fields of a relativistic particle beam [O'Connell et al., 2006]. On the other hand, collisional ionisation can be achieved, for instance, by triggering an electrical discharge in a gas capillary contained between two electrodes sustaining a high electric potential. After the discharge is initiated, the free electrons are quickly accelerated and drive an avalanche-ionization process that heats (and eventually fully ionises) the plasma [Spence and Hooker, 2000]. The latter is the mechanism used in the work presented in this thesis (cf. Section 2.2).

1.4.2. Excitation of nonlinear plasma wakes

In this section the excitation of charge-density wakes, achieved when a relativistic electron bunch is focused into a plasma, is analysed. The goal is to gain a basic insight into the different beam-plasma-interaction regimes and the properties of the resulting wakefields in order to identify which of them is the best suited for particle acceleration—which, as will be seen, corresponds to the highly nonlinear (or blowout) regime. Henceforth, the local electron density will be represented simply by $n \equiv n_e$ and the ambient unperturbed density by n_0.

Cold electron-fluid model in the quasi-static approximation (QSA). Under certain approximations, the cold electron-fluid model (cf. Section 1.4.1) provides a simple means of understanding the basic physical phenomena resulting from the beam-plasma interaction. Firstly, the particle beams considered have a transverse extent much larger than the plasma skin depth $k_p\sigma_r \gg 1$, and, consequently, only the longitudinal direction needs to be taken into account—i.e., the model is restricted to 1D plasma waves, although the three momentum components of the fluid are retained. Similar to the reference-particle frame used in Section 1.1, the formulation of the physics of plasma waves is carried through in the *co-moving frame*, in which the longitudinal coordinates become $\xi = z - ct$ and $\tau = t$. The second approximation consists of considering that the characteristic evolution of the particle beam—which is represented by the inverse of the betatron frequency $\tau_b \sim \omega_\beta^{-1} \approx \sqrt{2\gamma_{rel}}\omega_p^{-1}$ [Esarey, Catravas and Leemans, 2001], where γ_{rel} is the relativistic gamma of the beam and ω_p is the plasma frequency—is much slower than the time scale of the plasma. Therefore, the particle beam can be considered as frozen during the characteristic response time of the plasma, and the time variation of the electron-fluid quantities of interest and the electromagnetic

fields is dominated by the dependency in ξ—i.e., $\partial t \approx -c^{-1}\partial\xi$. This is referred to as the *quasi-static approximation* [Sprangle, Esarey and Ting, 1990]. The equations governing the cold electron-fluid model are the continuity equation and the fluid-momentum equation, which combined can be shown to result in a second-order ordinary differential equation describing the structure of the *wakefield potential* $\psi(\xi)$ in an initially unperturbed plasma under the influence of a relativistic particle beam [Mehrling, 2014]:

$$k_p^{-2}\frac{\partial^2\psi(\xi)}{\partial\xi^2} = \frac{1}{2(1+\psi(\xi))^2} - \frac{n_b(\xi)}{n_0} - \frac{1}{2}, \tag{1.4.9}$$

where $n_b(\xi)$ is the density distribution of the beam. The fluid quantities are related to the wakefield potential by the following equations:

$$\text{electron-fluid (plasma) density:} \qquad \frac{n}{n_0} = \frac{1+(1+\psi)^2}{2(1+\psi)^2} \tag{1.4.10a}$$

$$\text{normalised electron-fluid momentum:} \qquad u_z = \frac{1-(1+\psi)^2}{2(1+\psi)^2} \tag{1.4.10b}$$

$$\text{relativistic electron-fluid Lorentz factor:} \qquad \gamma_{rel} = \frac{1+(1+\psi)^2}{2(1+\psi)}, \tag{1.4.10c}$$

where the normalised momentum is $\boldsymbol{u} = \boldsymbol{p}/mc$. A general relationship between the wakefield potential ψ and the longitudinal and transverse electromagnetic fields in cylindrical coordinates is given by the expressions [Mehrling, 2014]:

$$\frac{E_z}{E_0} = -k_p^{-1}\frac{\partial\psi}{\partial\xi} \tag{1.4.11a}$$

$$\frac{E_r - cB_\theta}{E_0} = -k_p^{-1}\frac{\partial\psi}{\partial r}, \tag{1.4.11b}$$

where E_0 is the wave-breaking field (cf. Section 1.4.1). The last two equations are valid even beyond the cold electron-fluid model. Equations 1.4.9 and 1.4.10 accurately describe longitudinal plasma waves driven by a broad particle beam ($k_p^{-1}\sigma_r \gg 1$) in the cold electron-fluid model in the QSA in two different regimes:

a) the ***linear regime***, when $\dfrac{n_b}{n_0} \ll 1$ and $\psi \ll 1$

b) the ***nonlinear regime***, when $\dfrac{n_b}{n_0} \approx 1$ and $\psi \approx 1$,

both of which being extendable to 3D if the transverse electron-beam distribution and its

influence on the plasma dynamics need to be taken into account. The model, however, fails for very dense beams that are able to transversely expell (blow out) all the plasma electrons, thus creating a cavitated volume behind the drive beam. This gives rise to a third regime that requires a 3D kinetic description (cf. Section 1.4.1):

c) the ***highly nonlinear*** or ***blowout regime***, when $\frac{n_b}{n_0} \gg 1$ and $\psi \gg 1$.

Linear regime. When $\psi \ll 1$ the following approximation can be applied to Eq. 1.4.9 $(1+\psi)^{-2} \approx 1 - 2\psi$, thus resulting in the expression:

$$\left(\frac{\partial^2}{\partial \xi^2} + k_p^2\right) \psi(\xi) = -k_p^2 \frac{n_b(\xi)}{n_0}, \tag{1.4.12}$$

which is the equation of a forced harmonic oscillator. It can be shown that for a driver with a Gaussian profile:

$$n_b(\xi) = N_b \exp\left(-\frac{(\xi - \xi_c)^2}{2\sigma_\xi^2}\right) \exp\left(-\frac{r^2}{2\sigma_r^2}\right), \tag{1.4.13}$$

the amplitude of the excited wakefield is maximum for $k_p \sigma_\xi = 1$ [Mehrling, 2014], in which case the longitudinal and transverse fields reduce to:

$$\frac{E_z}{E_0} = N_b \sqrt{\frac{2\pi}{\exp(1)}} \cos\left(k_p(\xi - \xi_c)\right) \exp\left(-\frac{r^2}{2\sigma_r^2}\right), \tag{1.4.14a}$$

$$\frac{E_z - cB_\theta}{E_0} = -N_b \sqrt{\frac{2\pi}{\exp(1)}} \sin\left(k_p(\xi - \xi_c)\right) \frac{k_p r}{k_p^2 \sigma_r^2} \exp\left(-\frac{r^2}{2\sigma_r^2}\right). \tag{1.4.14b}$$

These are sinusoidal wakefields propagating with the plasma wavenumber for which the accelerating and focusing parts have a 90 deg phase shift, implying that only half of the accelerating phases are also focusing—and, therefore, well suited for accelerating particles. Notice that the longitudinal accelerating field retains the radial dependence of the drive bunch, which would induce an increase of the uncorrelated energy spread when accelerating beams with a finite transverse size. Similarly, the transverse focusing field is not linear with the radial coordinate and, additionally, its longitudinal variation is sinusoidal, which would induce strong aberrations in the transverse properties of the beam. Due to the large impact of these field geometries on the quality of the accelerated beams, the linear regime is not well suited for particle acceleration.

Nonlinear regime. The wakefield potential becomes $\psi \approx 1$ for intense driver beams $n_b/n_0 \approx 1$, thus invalidating the linear approximation discussed above. In this case, the full expression of Eq. 1.4.9 must be used. Additionally, for drivers with a finite transverse extent, the transverse force exerted on the plasma electrons becomes important, which typically requires the use of a 3D fluid model to capture the plasma dynamics accurately. Figure 1.8 shows the transition from the linear to the nonlinear regimes for three increasing driver-beam intensities $n_b/n_0 = 0.01$, $n_b/n_0 = 0.2$ and $n_b/n_0 = 0.6$—computed from numerical solutions of Eq. 1.4.9. As can be observed, for moderately intense beams (B) the fields still retain their sinusoidal features. However, the electron-fluid density is slightly peaked, with the peak being shifted towards longer distances behind the driver bunch, and the wakefield potential acquires a parabolic shape that results in an asymmetric accelerating field. For even higher intensities (C), these features are intensified and the final accelerating field develops a saw-tooth shape. Due to this field asymmetry, the overlap between the focusing and the accelerating fields increases, which implies that the phase range suitable for particle acceleration asymptotically approaches the full span of accelerating phases. The progressive elongation of the wakefield wavelength in the nonlinear regime is related to the fact that the electrons start to become relativistic and gain (relativistic) mass, thus reducing their plasma

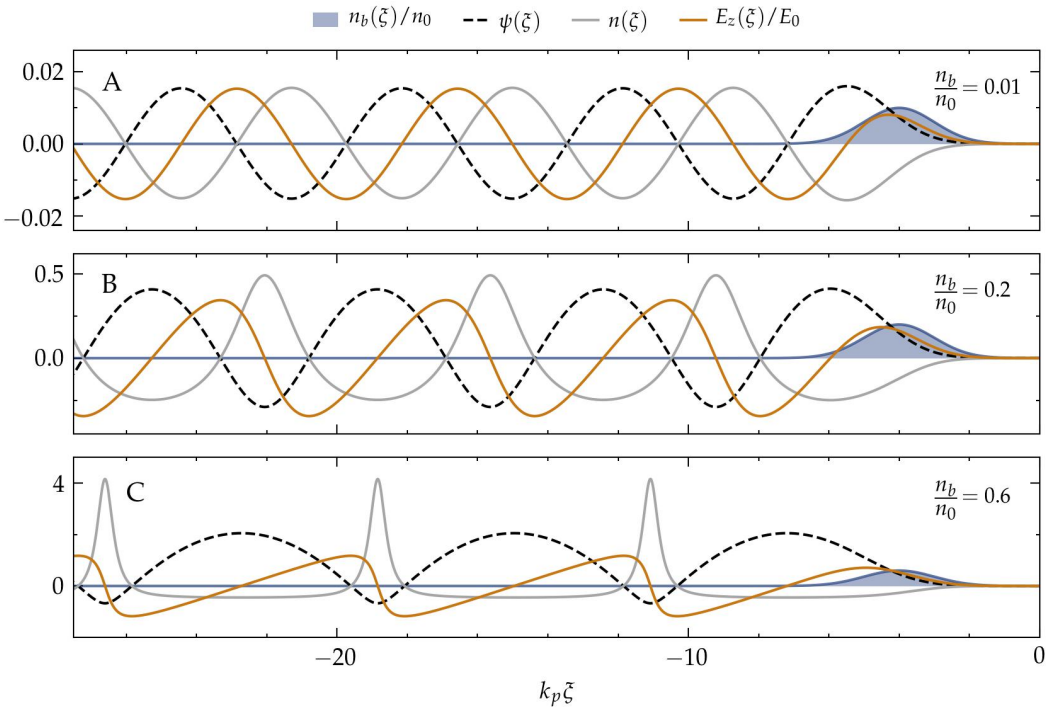

Figure 1.8.: Excitation of plasmas waves in the 1D cold electron-fluid model and the quasi-static approximation (QSA). Normalized wakefield potential, density variations and longitudinal fields generated by a Gaussian electron beam with $k_p\sigma_\xi = 1$, and $n_b/n_0 = 0.01$ (a), $n_b/n_0 = 0.2$ (b), and $n_b/n_0 = 0.6$ (c). Results obtained from numerical solutions of Eq. 1.4.9. The transition from linear to nonlinear wakefields, which feature an elongated wavelength and nonsinusoidal shape, can be clearly seen. Adapted from Ferran Pousa, 2019; numerical-integration framework provided by Ángel Ferran Pousa.

frequency. In this case the cold relativistic wave-breaking field can be shown to be [Esarey, Schroeder and Leemans, 2009]:

$$E_{0,rel} = \sqrt{2(\gamma_{rel} - 1)}\, E_0 \,, \tag{1.4.15}$$

where E_0 is the cold nonrelativistic wave-breaking field (cf. Eq. 1.4.8) and γ_{rel} is the relativistic electron-fluid Lorentz factor (cf. Eq. 1.4.10c).

Highly nonlinear or blowout regime

For short ($k_p\sigma_\xi < 1$) and intense ($n_b/n_0 \gg 1$) driver beams with a transverse beam size on the order and below the skin depth ($k_p\sigma_r \lesssim 1$), the plasma electrons are radially expelled (or blown out) from the region near to the propagation axis and are subsequently attracted by the static-ion volume left behind the driver, overshooting the beam-propagation axis after the bunch has propagated over a half of the plasma wavelength $\lambda_p/2$. In this way a cavitation that co-propagates behind the driver is formed, which is delimited by a thin and dense sheath of plasma electrons that cicumvent the blown-out volume and is surrounded by a weakly interacting neutral plasma. As mentioned before, in the blowout regime the fluid model is not valid and a 3D (or 2D if the system is cylindrically symmetric) kinetic model based on the Vlasov-Maxwell equations is required, which in general needs to be solved numerically by means of particle-in-cell (PIC) codes [Dawson, 1983]. Nonetheless, the characteristic cavity-like geometry formed by the sheath of plasma electrons enables a phenomenological approach to be applied. By considering a bi-Gaussian electron driver, the electric current associated with a sheath of width Δ_s with a radial position depending on the longitudinal coordinate $r_b = r_b(\xi)$ and uniform density $n_s = r_b^2/((r_b + \Delta_s)^2 - r_v^2)$ for $r_b \leq r \leq r_b + \Delta_s$, and assuming quasi-static conditions, a model developed by Lu et al., 2006 provides the following wakefield potential:

$$\psi(r,\xi) = \frac{r_b^2(\xi)}{4}(1+\beta) - \frac{r^2}{4} \,, \quad \text{for} \quad r < r_b \,, \tag{1.4.16}$$

where $\beta = \frac{(1+\alpha)^2 \ln(1+\alpha)^2}{(1+\alpha)^2 - 1)} - 1$ and $\alpha = \Delta_s/r_b$. For short and sufficiently intense drive beams such that the maximum blowout radius is $R_b k_p \gtrsim 4$ and $\beta \ll 1$—which is referred to as the *ultra-relativistic blowout regime*—the trajectory of the innermost electrons of the sheath can be expressed as:

$$r_b \frac{d^2 r_b}{d\xi^2} + 2\left(\frac{dr_b}{d\xi}\right)^2 + 1 = \frac{4\lambda(\xi)}{r_b^2} \,, \tag{1.4.17}$$

where $\lambda(\xi)$ is the current profile of the drive beam. According to Eqs. 1.4.11, the resulting

wakefields are:

$$\frac{E_z}{E_0} = \frac{k_p r_b}{2} \frac{dr_b}{d\xi} \tag{1.4.18a}$$

$$\frac{E_r - cB_\theta}{E_0} = \frac{k_p r}{2}. \tag{1.4.18b}$$

In the ultra-relativistic blowout regime, the trajectory of the sheath electrons $r_b(\xi)$ can be approximately described by the equation of a circle, and the slope of the longitudinal (accelerating) field fulfils [Lu et al., 2006]:

$$\frac{1}{E_0} \frac{\partial E_z}{\partial \xi} \approx \frac{k_p}{2}. \tag{1.4.19}$$

Two important properties can be readily observed from Eqs. 1.4.18:

1) the longitudinal field E_z at a location ξ is transversely homogeneous and is only determined by the radius r_b and the slope $dr_b/d\xi$ of the sheath of plasma electrons;
2) the transverse focusing field $E_r - cB_\theta$ increases linearly with the radial position r inside the cavitated volume and is independent of the longitudinal coordinate.

These two properties make the wakefields in the blowout regime especially well suited for particle acceleration, since they simultaneouly provide a transversely homogeneous longitudinal acceleration and an aberration-free transverse focusing that is constant along the whole range of accelerating phases. In contrast to the linear and mildly nonlinear regimes, these field geometries enable the preservation of the quality of the accelerated electron bunches, which, as will be discussed in Section 1.4.4, is one of the principal requirements for the successful application of PWFA to particle acceleration. Therefore, the experiments performed at FLASHForward operate in the blowout regime and the focus in the next sections will be on this regime.

To illustrate the formation of a plasma blowout in a range of parameters similar to those of the experiments discussed in this thesis, simulations with the code Wake-T [Ferran Pousa, Assmann and Alberto Martinez de la Ossa, 2019] are performed. A bi-Gaussian electron drive bunch with a normalised emittance of $\epsilon_{n,r} = 4$ µm, an rms length of $\sigma_\xi \approx 25$ µm ($k_p \sigma_\xi = 0.5$) and a peak current of $I_b = 1.5$ kA—corresponding to a charge of $Q \approx 350$ pC—is focused into an unperturbed plasma with an electron density of $n_0 = 10^{16}$ cm^{-3} and a longitudinal extent of $L = 5$ cm. Two different transverse focal sizes are simulated:

a) $\beta_r = 10$ mm ($k_p \sigma_r \approx 0.1$), which results in a peak particle-beam density of $n_b \approx 2.5 \cdot 10^{17}$ cm^{-3} and $n_b/n_0 \approx 25.0$;
b) $\beta_r = 0.5$ m ($k_p \sigma_r \approx 0.6$), with $n_b \approx 5.0 \cdot 10^{15}$ cm^{-3} and $n_b/n_0 \approx 0.5$, representing

a situation in which an optimal focusing into the plasma is prevented due to, e.g., focusing errors.

The simulations are performed in the QSA, assuming cylindrical symmetry and using 4 plasma macro-particles per radial cell. The results after propagation of the beam over the full plasma length are shown in Figure 1.9, on the left for the strongly focused beam (case a) and on the right for the weakly focused beam (case b). The top plots (A and D) show the plasma density (grey colorscale) and the driver particle-beam distribution (coloured colorscale), together with the current profile and an ellipse on top of the beam indicating an area $A = \pi\sigma_\xi\sigma_r$ occupied by the initial particle distribution. The longitudinal and radial wakefields are shown in the middle (B and E) and bottom (C and F) plots, with a lineout of the field at $k_p r = 0$ and $k_p\xi \approx -3.5$, respectively. As can be observed, even in the case of a strongly focused driver (A), the bubble radius is slightly below $k_p R_b \lesssim 1$, indicating that the blowout does not develop in the ultra-relativistic regime. This also implies that the slope of the longitudinal field E_z (B) for accelerating phases is lower than $\partial E_z/\partial\xi \cdot E_0^{-1} < k_p/2$. For the weakly focused driver (D), the bubble radius is decreased, since a smaller fraction of

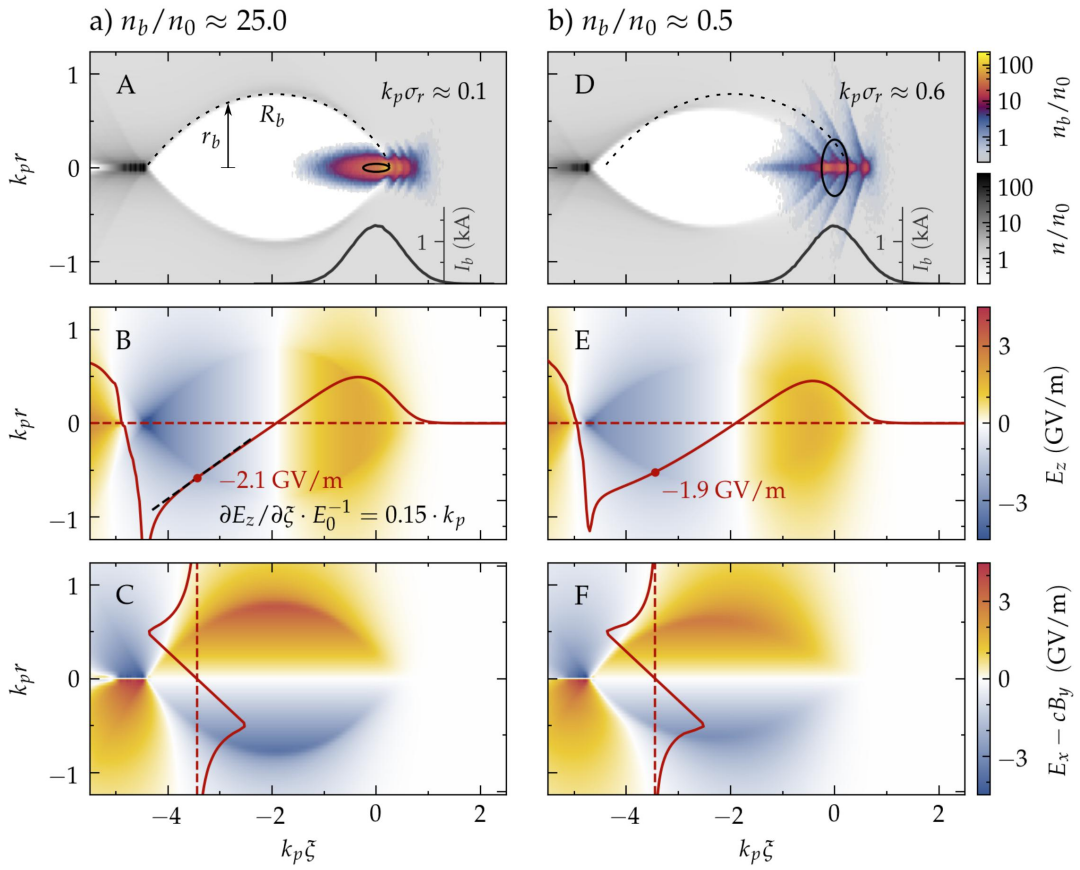

Figure 1.9.: Wakefields in the blowout regime. Tracking simulations showing the blowout achieved with a strongly focused (left) and a weakly focused (right) beam into a plasma with density $n_0 = 10^{16}$ cm^{-3} after a propagation length of $L = 5$ cm. See text for details. Simulations performed with WAKE-T [Ferran Pousa, Assmann and Alberto Martinez de la Ossa, 2019].

the driver-beam particles participate in the formation of the blowout. As a consequence, the longitudinal field (E) decreases accordingly and its derivative starts to exhibit a nonlinear behaviour—i.e. the slope is not constant. The elongation of the longitudinal bubble size is related to the descrease in E_z amplitude, since the sheath electrons at the rear part of the cavity experience a smaller pull and are less accelerated towards the direction of propagation—i.e., are less relativistic than in the other case. A relevant observation is that in both cases the transverse focusing field depends linearly with the radius r and has the same gradient $K_r/E_0 = k_p/2$—i.e., it suffices that the volume behind the driver is fully cavitated for the background ions to produce a focusing field. As a final remark, it should be pointed out that once the driver is focused to enable a large fraction of driver-beam particles to contribute to the formation of the blowout, a further reduction of its transverse beamsize does not change the shape of the cavity and the resulting wakefields—i.e., only an increase in current can increase the strength of the wakefields. This is actually implicit in the differential equation of the blowout radius (cf. Eq. 1.4.17) in the ultra-relativistic regime, in which only the current profile appears. Furthermore, as the wakefields develope, the transverse focusing fields lead to a sort of pinching of the driver.

1.4.3. External injection in the blowout regime

There are two basic mechanisms used to inject electrons into the accelerating phase of a wakefield structure: 1) *internal injection*, by which electrons from the plasma background are trapped in the cavity; 2) *external injection*, which relies on using a pre-accelerated (relativistic) electron bunch from an external source. The work presented in this thesis deals with the latter. In the acceleration process, some of the energy of the drive bunch is transferred to the externally injected (trailing) bunch by means of the excited wakefields, so that the plasma medium can be regarded as a kind of energy transformer. As such, a key aspect of the acceleration process is the efficiency with which this energy transfer is accomplished. However, since the beam parameters of the accelerated bunch must fulfil the requirements for the envisaged applications (cf. Section 1.4.4), other aspects of the beam-plasma interaction must be taken into account. In the next paragraphs some of these aspects are briefly discussed.

Beam loading

While the blowout regime provides large accelerating gradients, their pronounced dependence on the longitudinal coordinate $\partial E_z/\partial\xi$ can imprint extreme energy chirps that dramatically increase the projected energy spread of the accelerated bunch. The following example illustrates the point. As discussed in Section 1.4.2, in the ultra-relativistic blowout regime the slope of the accelerating gradient is $\partial E_z/\partial\xi = E_0 k_p/2$ (cf. Eq. 1.4.19), which for a plasma density of $n_0 = 10^{16}$ cm^{-3} has a value of $\partial E_z/\partial\xi \approx 90$ MV/m/µm. If a bunch with energy $E = 1$ GeV and a length $\sigma_\xi = 10$ µm is accelerated at a mean gradient of $\langle E_z \rangle = 2.0$ GV/m

over a plasma length of $L_p = 0.5$ m—so that the bunch doubles its energy—the energy difference between the head and the tail will be $\Delta E \approx 450$ MeV, which at 2 GeV corresponds to a relative energy of $\Delta\delta \approx 20$ %. It is therefore obvious that a substantial degree of control over the (longitudinal) flatness of the accelerating field is required to produce bunches with a low energy spread. The basic mechanism to achieve this flattening is *beam loading*, by which the Coulomb fields produced by the trailing bunch itself interact with the plasma-wake structure to modify the accelerating wakefield.

In the blowout regime, the conditions that the trailing bunch must fulfil in order to optimally flatten E_z are described by an analytical model developed by Tzoufras et al., 2008, which is in turn based on the phenomenological model of the blowout regime from Lu et al., 2006 (cf. Section 1.4.2). When a trailing-bunch is located at the rear of the bubble, the plasma-sheath electrons feel the repelling force from the additional charge in addition to the force from the ion channel. This modifies the trajectory of the sheath electrons, which, as described by Eq. 1.4.18a, changes the resulting accelerating field. By inferring the shape of the sheath that results in a constant field, Tzoufras et al., 2008 established that the current profile of the trailing bunch $\lambda(\xi)$ should be trapezoidal:

$$\lambda(\xi) = \sqrt{E_s^4 + \left(\frac{R_b}{2}\right)^4} - E_s(\xi - \xi_s), \tag{1.4.20}$$

where ξ_s is the location of the bunch head, E_s the unperturbed field at ξ_s, and R_b is the maximum blowout radius. This is highlighted in Figure 1.10, which shows simulations performed with the tracking code Wake-T, without (left) and with (right) the load of a trailing bunch.

According to the current profile given by Eq. 1.4.20, the total amount of charge that can be coupled into the back of the bubble is given by the expression [Tzoufras et al., 2008]:

$$Q_s E_s = \frac{\pi R_b^4}{16} \tag{1.4.21}$$

which illustrates the trade-off between the number of particles that can be accelerated and the accelerating gradient. Since Eq. 1.4.21 expresses the energy absorbed per unit length, it derives that for optimal-beam-loading conditions the efficiency of the energy transfer from the wake to the trailing bunch does not depend on the accelerating field E_s.

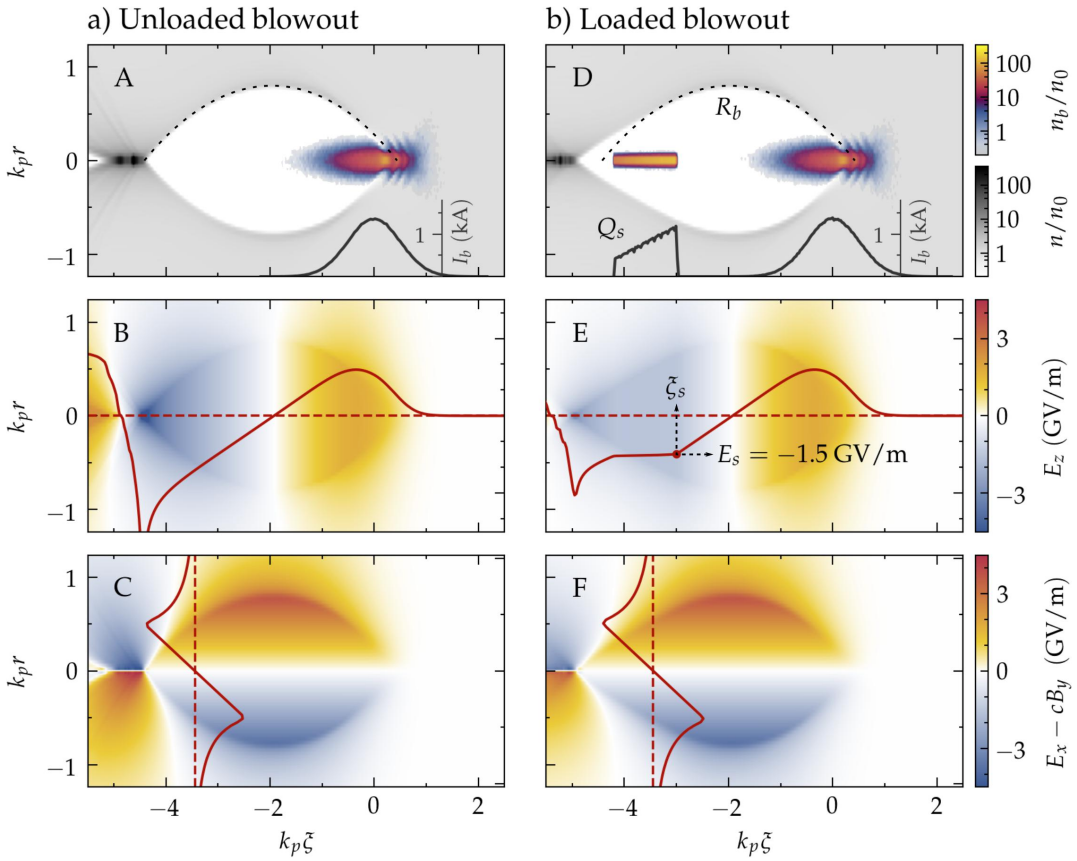

Figure 1.10.: Optimal beam loading in the blowout regime. Tracking simulations showing a blowout of an unloaded (left) and a loaded (right) plasma cavity, together with their corresponding accelerating E_z (middle) and focusing $E_r - cB_\theta$ fields (bottom). The background plasma density is $n_0 = 10^{16}$ cm^{-3}. A drive bunch with an energy of $E = 1$ GeV, a normalised emittance of $\epsilon_{n,r} = 4$ µm and a charge of $Q_d \approx 330$ pC is focused into the plasma with a beta of $\beta_r = 10$ mm ($k_p \sigma_r \approx 0.1$), producing a relative particle density of $n_b/n_0 \approx 25$. As described in [Tzoufras et al., 2008], the trapezoidal current profile of the trailing bunch flattens the accelerating field, lowering the amplitude of the unloaded field to the value E_s behind the location where the trailing bunch arrives ξ_s. The charge of the trailing bunch is $Q_s \approx 185$ pC. The simulations are performed with the tracking code WAKE-T [Ferran Pousa, Assmann and Alberto Martinez de la Ossa, 2019] with 4 particles per cell, using the QSA and assuming cylindrical symmetry.

Energy-transfer efficiency and transformer ratio

The *energy-transfer efficiency* from the wake to the accelerated particles is defined as the ratio between the energy gained by the trailing bunch W_{gain} and the energy deposited into the wake W_{dep}—i.e., the energy lost by the driver:

$$\eta_{\mathrm{wake}\to w} = \frac{W_{\mathrm{gain}}}{W_{\mathrm{dep}}}. \tag{1.4.22}$$

Assuming a propagation length through the plasma of L_p and an idealised case with optimal

beam loading and a constant decelerating field in the region of the driver, the energy gain of the trailing bunch is $W_{\text{gain}} = -Q_w E_z^w L_p$ and the energy loss of the driver $W_{\text{dep}} = Q_d E_z^d L_p$. In this case, the energy-transfer efficiency can be expressed as:

$$\eta_{\text{wake}\to w} = -\frac{Q_w E_z^w}{Q_d E_z^d} = \frac{Q_w}{Q_d} \cdot T_R, \tag{1.4.23}$$

where the *transformer ratio* $T_R = -E_z^w / E_z^d$ has been introduced [Ruth et al., 1985]. Since $Q_i E_z^i$ expresses the rate of energy gain or loss per unit length, Eq. 1.4.23 is sometimes referred to as the *instantaneous* energy-transfer efficiency. Equation 1.4.23 shows that increasing the transformer ratio is one element required to maximise the energy-transfer efficiency [Loisch et al., 2018; Roussel et al., 2020].

As seen in the previous section (cf. Eq. 1.4.21), for an optimally beam loaded wake the energy extracted by the (trapezoidal) trailing bunch is determined by the blowout radius $Q_w E_z^w \propto R_b^4$ (cf. Eq. 1.4.21). Since acceleration is possible as long as all driver particles have an energy > 0, the maximum amount of energy that can be deposited into the wake is reached if all driver particles experience the same decelerating field. Similar to the beam loading mechanism described previously, this is achieved when the driver bunch has a trapezoidal current profile that flattens the decelerating field [Lotov, 2005; Tzoufras et al., 2008; Roussel et al., 2020].

The overall efficiency η_{tot} of the PWFA must take into account the fraction of driver energy deposited into the wakefield $\eta_{d\to\text{wake}}$:

$$\eta_{d\to\text{wake}} = \frac{Q_d E_z^d L_p}{W_d} \quad \longrightarrow \quad \eta_{\text{tot}} = \eta_{d\to\text{wake}} \cdot \eta_{\text{wake}\to w}, \tag{1.4.24}$$

where W_d is the total energy of the driver before plasma interaction. Therefore, to maximise the overall efficiency, the energy of the driver must be fully depleted $\eta_{d\to\text{wake}} \approx 1$. Full energy depletion is ultimately limited at the point where the driver starts to become nonrelativistic and 1) the peak current is not enough to sustain the wake or 2) the synchronicity between the wake and the trailing bunch is lost. Considering that for an ultra-relativistic driver the energy at which these issues start to be a concern is negligible compared to its initial energy, the depletion length can be approximated by the expression:

$$L_{\text{dep}} \approx \frac{W_d}{Q_d E_z^d} \,. \tag{1.4.25}$$

Beta matching and emittance growth

In Section 1.4.2 it has been stated that the transverse focusing wakefields produced in the blowout regime exhibit a linear dependence on the radial coordinate r. According to Eq. 1.4.18, in such a cylindrically symmetric focusing channel, the focusing strength k_1 experienced by a relativistic electron with momentum $p_0 \approx \gamma_{rel} m_e c$ can be expressed as (cf. Eq. 1.1.7):

$$k_1 = \frac{e}{p_0} \cdot \frac{n_0 e}{2c\epsilon_0} = \frac{k_p^2}{2\gamma_{rel}} \,, \tag{1.4.26}$$

and the beam envelope is described by the evolution of the beta function (cf. Section 1.1.1):

$$\beta(s) = \beta_0 \cos^2\left(\sqrt{k_1} s\right) + \frac{1}{\beta_0 k_1} \sin^2\left(\sqrt{k_1} s\right), \tag{1.4.27}$$

where β_0 is the beta function at the entrance of the plasma channel and s the propagation length. Equation 1.4.27 suggests that the beam envelope will oscillate between β_0 and $1/k_1\beta_0$—i.e., the beam ellipse will 'rotate' in transverse phase space. For beams with a finite energy spread, these oscillations lead to a progressive smearing out of the phase space in a process referred to as *betatron decoherence*, which results in a growth of the projected emittance. In order to mitigate these effects, the tendency of the beam to diverge due to its emittance must be balanced by the attractive force of the ion focusing, which is referred to as the *matching* condition and is fulfilled when:

$$\beta_m = \frac{1}{\sqrt{k_1}} = \frac{\sqrt{2\gamma_{rel}}}{k_p} \tag{1.4.28a}$$

$$\alpha_m = 0. \tag{1.4.28b}$$

For an injected bunch with an energy of 1 GeV and a plasma density of $n_0 = 10^{16}$ cm^{-3}, the matched beta has a value of $\beta_m = 3.3$ mm, which clearly indicates the challenging nature of the lattice optic required to focus the beams into the plasma[8]. A progressive, adiabatic matching of the beta function is possible by precisely shaping the plasma-density profile at the entrance of the plasma cell [Marsh et al., 2005; Floettmann, 2014]. However, this is difficult to achieve experimentally and is still a subject of active research [Xu et al., 2016; Frederico, Hogan and T. Raubenheimer, 2017; M. D. Litos et al., 2019].

If the beam is not matched, the emittance saturates after full decoherence at a value given by the expression [Mehrling, Grebenyuk et al., 2012]:

[8] As a comparison, the nominal beta function at the interaction point at the LHC is $\beta^* = 55$ cm [Benedikt, Collier et al., 2004].

$$\epsilon_{n,\mathrm{sat}} = \frac{\epsilon_{n,\mathrm{init}}}{2}\left(\frac{1+\alpha_0}{\beta_0}\beta_m + \frac{\beta_0}{\beta_m}\right) = \epsilon_{n,\mathrm{init}}\ \mathrm{B}_{\mathrm{mag}}, \tag{1.4.29}$$

where $\epsilon_{n,\mathrm{init}}$ is the initial beam emittance and $\mathrm{B}_{\mathrm{mag}}$ is the mismatch parameter introduced in Eq. 1.1.58—considering that the matched parameters are the 'design' ones ($\beta_m = \beta_D$). The *decoherence* (or *saturation*) *length* is related to the matched beta and the energy spread of the bunch by the expression [Lindstrøm, 2019a]:

$$L_{dc} = \frac{\beta_m}{\sigma_\delta}, \tag{1.4.30}$$

which for a plasma density of $n_0 = 10^{16}$ cm^{-3} and an FEL-quality bunch with an energy spread of $\sigma_\delta \lesssim 0.1$ % has a value of $L_{dc} \gtrsim 3.3$ m. In the work presented in this thesis the plasma has a longitudinal extent of $L = 0.05$ m and is operated at densities around $n_0 \lesssim 10^{16}$ cm^{-3}. Therefore, the emittance growth induced by beta mismatch should be negligible.

Transverse misalignments

If the trailing bunch enters the plasma with an offset in the transverse plane—i.e., $\langle x \rangle \neq 0$ and/or $\langle x' \rangle \neq 0$ and equally for the vertical plane—it will undergo transverse oscillations around the axis of the wakefield structure as it propagates. For beams with a finite energy spread, this leads to betatron decoherence and, therefore, to an emittance growth [Lindstrøm, Adli, Pfingstner et al., 2016]. The decoherence length, however, is on the same order than that given in Eq. 1.4.30, and is not a concern for the experiments presented in this thesis. Nonetheless, if the amplitude of the oscillations is large enough, the trailing bunch can sample the outer regions of the wakefield structure in which the accelerating field is not transversely constant and the focusing field is not linear anymore, thus leading to a degradation of the beam quality in terms of both energy spread and emittance and to an eventual loss of charge into the plasma background. Additionally, as a result of the interaction between the (oscillating) trailing bunch and the plasma-sheath electrons, the shape of the wakefield structure can be distorted, which can have an impact on the quality of the accelerating and focusing fields and on the overall stability of the acceleration process. These issues will be investigated in Chapter 5.

Regarding the driver bunch, the presence of sliced transverse centroid offsets can seed an exponential amplification of transverse oscillations during propagation in the plasma, ultimately leading to a beam breakup. This effect, which is known as the *hosing* (or *electron-hose*) *instability* [Whittum et al., 1991], severely constraints the length over which the wakefield structure can be sustained, thus preventing the driver-energy depletion and critically limiting the applicability of PWFA. Strategies to mitigate the instability have been proposed, including the use of a drive beam with an energy chirp—either introduced prior to the interaction or

developed during propagation—or tapered plasma profiles [Mehrling, Fonseca et al., 2017], and focusing the driver beam into the plasma to a transverse size comparable to the blowout radius [A. Martinez de la Ossa, Mehrling and Osterhoff, 2018].

1.4.4. Requirements for applications and experimental progress

As mentioned in the previous section, PWFA schemes with external injection of particle beams can be considered as a kind of energy booster (or energy transformer). Therefore, its natural application is preeminently on high-energy particle colliders, a research area in which the demand for new compact and cost-effiecient accelerator technologies is more pressing than ever (cf. Introduction). Since the amount of energy that can be transferred to the trailing bunch is limited by the energy contained within the driver, the ultimate goal of building a TeV-energy particle accelerator must rely on the concatenation of multiple PWFA stages. Currently, however, there are still significant scientific unknowns preventing PWFA from being scaled to ultra-high energies. As such the near term goal is to demonstrate the operation of a single PWFA high-efficiency, quality preserving stage with large energy gain [Albert et al., 2021]. The issues that need to be addressed to that end go beyond the physics of the acceleration process itself and include, among others, the development of appropriate plasma sources with, e.g., tunable density profiles [Xu et al., 2016] or new diagnostic techniques [Downer et al., 2018]. In the following the focus is on the physics of the acceleration process.

At the current stage of research, the challenge is to ensure that a single PWFA stage fulfils multiple requirements simultaneously [Lindstrøm, Beinortaite et al., 2021]:

- a high beam-loaded acceleration gradient significantly larger than those produced in RF cavities
- beam-quality preservation, including energy spread—to which the work presented in this thesis is devoted—and emittance;
- a high instantaneous energy-transfer efficiency of ~ 50 %;
- trailing-bunch charge-coupling efficiency of 100 %;
- depletion of ~ 50 % of the driver-beam energy;
- overall stability of the acceleration process.

As discussed in the previous section, in order to achieve the first four goals a high degree of control over the parameters of the incoming bunches is mandatory—viz. precise shaping of their current profile, appropriate focusing into the plasma and mitigation of transverse misalignments between the driver and the trailing bunch. The depletion of a large fraction of the driver-beam energy requires the acceleration process to be sustained over a sufficient propagation length while optimising the energy transfer from the driver to the wake—which, again, requires an appropriate shaping ($k_p\sigma_\xi < 1$) and focusing ($k_p\sigma_r < 1$) of the driver bunch. In this regard, the presence of sliced transverse centroid offsets in the

driver can seed the hosing instability and, even if the propagation length is not enough to cause the beam breakup, the overall stability of the acceleration process might be severely compromised. Therefore, a precise control over the sliced transverse beam parameters is also required. In the blowout regime, a variation of the distance between the driver and the trailing bunch $\Delta\xi_{d-w}$ leads to an energy-gain variation on the order of $\Delta\delta_w \approx k_p\Delta\xi_{d-w}/2$ (cf. Eq. 1.4.19). Therefore, in order to keep the energy jitter of a PWFA stage below $\Delta\delta_w < 1$ % at a plasma density of $n_0 = 10^{16}$ cm^{-3}, the timing between the two bunches must be better than $\Delta\xi_{d-w} \lesssim 1.0$ µm (3.5 fs). This illustrates the level of accuracy that any strategy adopted for the production of the double-bunch structure must fulfil.

Due to the small transverse size of the accelerated beams exiting the wakefield structure, their capture and transport is prone to be affected by chromatic effects that can lead to an emittance growth (cf. Section 1.1.2, Eq. 1.1.57). Since this issue is just the reverse case of focusing the beam into the plasma, such effects can be mitigated by using tapered plasma-density profiles, thus enabling an adiabatic release of the accelerated bunch. An alternative approach is to use a plasma lens, i.e., a charged-particle optics device that provides strong focusing in both planes simultaneously [Lindstrøm, Adli, Boyle et al., 2018].

Experimental progress. While the experimentation with PWFA schemes was started by Rosenzweig, Cline et al., 1988 in the late 80s at the Advanced Accelerator Test Facility (Argonne National Laboratory), it has been during the past two decades that this field of research has undergone its most rapid progress. This has been facilitated by the developement of laser-driven RF photocathode guns [Dowell et al., 2003; Stephan et al., 2010] that are able to deliver electron bunches with a very high phase-space density. After acceleration to ultra-relativistic energies and longitudinal compression, these bunches constitute excellent drivers of plasma wakes. Pioneering experiments were performed between 1998 and 2015 at the Final Focus Test Beam (FFTB) facility [Joshi et al., 2002] and at the Facilities for Accelerator Science and Experimental Test (FACET) [Hogan, T. O. Raubenheimer et al., 2010], both at the Stanford Linear Accelerator Center (SLAC). By using a single electron bunch, Blumenfeld et al., 2007 demonstrated high accelerating gradients and high energy gain of a small fraction of particles at the back of the wakefield structure, and Clayton et al., 2016 provided a first experimental validation of the potential suitability of PWFA for quality-preserving acceleration through indirect measurements of the longitudinal and transverse field structure of a blown-out plasma wake. In more advanced double-bunch experiments, M. Litos, Adli, An et al., 2014 showed a high instantaneous energy-transfer efficiency of $\sim$ 30 % between a driver and a trailing bunch. However, the energy spread of the latter increased by 100 % and about $\sim$ 75 % of its charge was lost during the beam-plasma interaction, clearly indicating that a number of issues remained to be addressed in order for PWFA to be ready for applications. One of the facilities that picked up the baton of PWFA research was the Future Oriented Wakefield Acceleration Research and Developement facility at FLASH[9] (FLASHForward) at DESY. Since its first operation in 2017, four important milestones have been achieved: the

[9]Free-electron Laser in Hamburg (FLASH) [Schreiber and Faatz, 2015] at DESY.

experimental demonstration of a tunable plasma-based energy dechirper [D'Arcy, Wesch et al., 2019]; a novel measurement technique that enables sampling (longitudinal) accelerating wakefields of order GV/m with femtosecond time resolution [Schröder, Lindstrøm et al., 2020]; energy-spread preservation and high energy-transfer efficiency in a plasma-wakefield accelerator [Lindstrøm, Garland et al., 2021]; study of the ion motion in a plasma-wakefield accelerator and determination of its recovery time [D'Arcy, Chappell et al., 2022]. Table 1.2 lists some of the most important experimental milestones of PWFA with electron beams worldwide.

Year	Facility	Experimental results
1988	ANL-AATF[1]	First experimental observation of plasma wakefield acceleration in the linear regime by Rosenzweig, Cline et al., 1988.
1990	KEK[2]	Acceleration in the linear regime in a plasma wakefield excited by bunch trains by Nakajima et al., 1990.
2000	ANL-AATF	First observation of plasma wakefield acceleration in the blowout regime by Barov et al., 2000.
2004	SLAC-FFTB[3]	Meter-scale plasma-wakefield accelerator driven by a matched electron beam by Muggli et al., 2004.
2005	SLAC-FFTB	2.7 GeV energy gain of electrons accelerated in a 10-cm-long plasma by Hogan, Barnes et al., 2005.
2007	SLAC-FFTB	Energy doubling of 42 GeV electrons in a metre-scale plasma wakefield accelerator by Blumenfeld et al., 2007.
2008	BNL-ATF[4]	First driver-trailing-bunch pair experiments. High-gradient plasma-wakefield acceleration with two subpicosecond electron bunches by Kallos et al., 2008.
2014	SLAC-FACET[5]	High-efficiency acceleration of an electron beam in a plasma wakefield accelerator by M. Litos, Adli, An et al., 2014.
2016	SLAC-FACET	9 GeV energy gain in a beam-driven plasma wakefield accelerator by M. Litos, Adli, Allen et al., 2016. Self-mapping the longitudinal field structure of a nonlinear plasma accelerator cavity by Clayton et al., 2016.
2018	DESY-PITZ[6]	Observation of the self-modulation instability via time-resolved measurements by Gross et al., 2018. Observation of high transformer ratio plasma wakefield acceleration by Loisch et al., 2018.
2019	DESY-FLASHForward[7]	Tunable plasma-based energy dechirper by D'Arcy, Wesch et al., 2019.
2021	DESY-FLASHForward	Energy-spread preservation and high energy-transfer efficiency in a plasma-wakefield accelerator by Lindstrøm, Garland et al., 2021.
2022	DESY-FLASHForward	Recovery time of a plasma-wakefield accelerator by D'Arcy, Chappell et al., 2022.

Table 1.2.: An abridged historical overview of experimental progress in PWFA with electron beams. [1]Argonne National Laboratory, Advanced Accelerator Test Facility, Chicago, IL, USA; [2]Kō-Enerugī butsurigaku Kenkyūsho (High-energy Physics Research Institute), Tsukuba, Ibaraki, Japan; [3]Stanford Linear Accelerator Center, Final Focus Test Beam, Stanford, CA, USA; [4]Brookhaven National Laboratory, Accelerator Test Facility, Brookhaven, NY, USA; [5]Stanford Linear Accelerator Center, Facility for Advanced Accelerator Experimental Tests, Stanford, CA, USA; [6]Deutsches Elektronen Synchrotron, Photoinjektor-Teststand in Zeuthen, Zeuthen, Brandenburg, Germany; [7]Future Oriented Wakefield Acceleration Research and Developement facility at FLASH, DESY, Hamburg, Germany. Table adapted from [Lindstrøm, 2019a].

2. Experimental facility

The FLASHForward experimental facility [D'Arcy, Aschikhin et al., 2019] is a test-bed for beam-driven plasma-wakefield (PWFA) research and developement, with a view towards both photon-science (e.g. FELs) and high-energy-physics applications (e.g. linear colliders). The facility benefits from the FEL-quality electron bunches provided by the FLASH linac to drive a wakefield in a plasma produced inside a windowless gas cell with lengths of order cm. The experimentation is broken down into three flagship experiments:

X-1 an **internal-injection** experiment, which uses the plasma itself as a cathode in order to generate ultrashort, low-emittance beams by trapping and accelerating electrons from the ambient plasma background [Knetsch et al., 2021];

X-2 an **external-injection** experiment, accelerating existing bunches from FLASH with high efficiency while preserving their beam quality [Libov et al., 2018; Lindstrøm, Beinortaite et al., 2021; Lindstrøm, Garland et al., 2021];

X-3 a **high-repetition-rate** experiment, dedicated to high-average-power PWFA studies enabled by the MHz-repetition-rate bunch-train structure delivered by the FLASH linac [D'Arcy, Aschikhin et al., 2019; D'Arcy, Chappell et al., 2022].

Additionally, the scientific programme of FLASHForward incorporates a diverse range of ancillary experiments, including, among others, the characterisation and development of plasma cells [Garland et al., 2021] and the development of active plasma lenses [Meisel, 2018; Lindstrøm, Adli, Boyle et al., 2018; Röckemann et al., 2018].

The installation of the FLASHForward beamline started in summer 2016 and the first beam-based commissioning campaign took place in August 2017. In January 2018 the construction of the beamline was completed with the installation of the experimental chamber for plasma experimentation and a short post-plasma diagnostics section with a broadband dipole spectrometer for energy-spectrum measurements. The first signatures of a beam-driven plasma wakefield were observed in June 2018 and external injection was achieved shortly thereafter, in August 2018. The beamline was upgraded in summer 2019 to incorporate the diagnostics capabilities of the PolariX-TDS system. In summer 2020 a new dipole spectrometer with high spatial resolution was added to the diagnostics section located immediately after the plasma chamber.

The goal of this chapter is to describe the core constituents of the experimental facility and provide an overview of both its basic operation and its capabilities with emphasis on

the X-2 experiment. In Section 2.1 the focus is on the FLASH superconducting RF linac, which provides electron bunches to its three beamlines—FLASH1 and FLASH2, both photon-science user facilities, and FLASHForward. The first operational FLASHForward beamline installed before the 2019 upgrade is described next (cf. Section 2.2), including its beamline design and its diagnostic capabilities. The specifics of the PolariX-TDS beamline design and its commissioning are addressed separately in Section 2.3. Finally the key aspects of the distinctive vacuum design that enables the operation of windowless gas cells is described.

2.1. FLASH

Originally conceived as a test stand for a future linear collider—the TeV Superconducting Linear Accelerator (TESLA)—with an integrated x-ray free-electron laser [Richard et al., 2001], the **F**ree-electron **Las**er in **H**amburg (FLASH) is a photon-science user facility located at DESY producing SASE-FEL radiation with wavelengths ranging from the vacuum ultraviolet (VUV) to the soft x-ray (XUV) regime. In many respects it has served as a prototype for the European XFEL [Altarelli et al., 2007] and to some extent as a test bed for the International Linear Collider Project (ILC) [Brau, Okada and Walker, 2007]. FLASH demonstrated SASE operation at a wavelength of 32 nm in 2005 [Ayvazyan et al., 2006], at 13 nm in 2006 [Rossbach, 2006; Ackermann et al., 2007] and at 6.5 nm (design value) in 2007 [Schreiber, Faatz and Honkavaara, 2008]. An extensive historical review of the scientific progress driven by FLASH over the last two decades—in both accelerator research and development (ARD) and photon science—can be found in Rossbach, Schneider and Wurth, 2019.

One of the distinctive features of the superconducting RF (SCRF) technology used at FLASH is that, due to the low resistive losses generated at the walls of the acceleration cavities, long RF-pulses are sustainable, thus enabling the production of long bursts of electron and photon pulses. Therefore, the FLASH facility combines the extreme peak brightness of an FEL with a very high average brightness [Rossbach, 2016]. Since the experiments in photon-science rarely exhaust the large number of pulses available in a burst, a second (parallel) beamline equipped with variable gap undulators [Schöps et al., 2014], FLASH2, was constructed between late 2011 and early 2014 [Vogt, Faatz et al., 2015; Rönsch-Schulenburg, Faatz et al., 2017], practically doubling the beamtime accessible to the users. Parallel operation is achieved by extracting a fraction of the bunches from the burst by means of a kicker-septum system and directing it to the second beamline (cf. Section 2.1.2). The FLASHForward beamline is attached to the FLASH linac after the FLASH2 extraction arc. This is schematically represented in Figure 2.1 together with the main components of the FLASH linac, which are described in the following section.

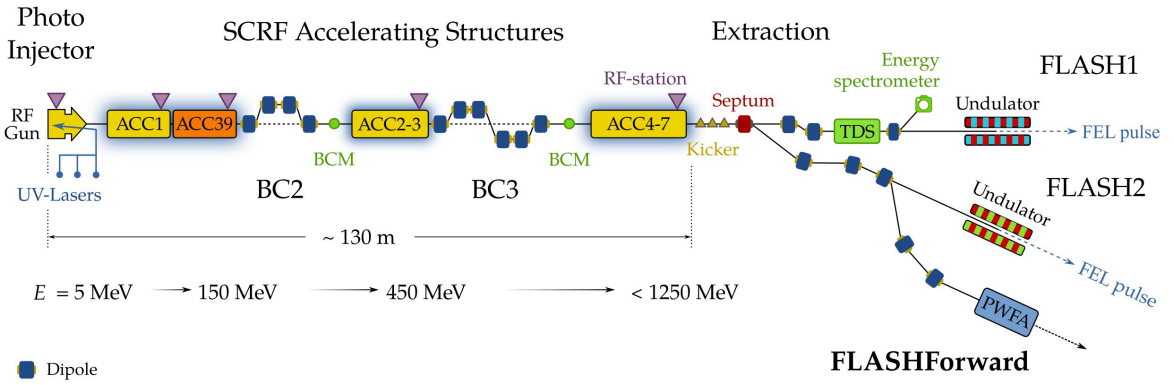

Figure 2.1.: Schematic layout of the FLASH facility with its linac and its three beamlines: FLASH1 and FLASH2—both photon-science user facilities equipped with undulators—and FLASHForward. The linac, with a total length of ∼ 130 m, consists of a normal-conducting photo-injector and seven superconducting RF (SCRF) modules separated by two bunch compressors (BC2 and BC3) at the indicated beam energies. A third harmonic module (ACC39) linearises the longitudinal phase space for effective compression, which can be monitored with a TDS and an energy spectrometer located upstream of the FLASH1 undulators, and two BCMs located downstream of BC2 and BC3. After acceleration up to an energy of $E < 1250$ MeV the electrons are directed to either FLASH1 or FLASH2 by means of a kicker-septum system. Extraction from the FLASH2 beamline to FLASHForward is achieved with a static dipole magnet.

2.1.1. FLASH linac

Photo injector. At FLASH the electrons are extracted from a Cs_2Te photocathode [Schreiber, Lederer et al., 2017] by using one of three different UV[1] lasers: two of similar pulse length producing pulse durations of 4.5 ps and 6.5 ps and injecting charges up to 2 nC, and a third one producing pulses with durations between 0.8 ps and 1.6 ps and injecting charges between 20 pC and 120 pC for short-pulse-mode operation [Rönsch-Schulenburg, Hass et al., 2013; Plath et al., 2013]. The electrons are emitted directly into a normal-conducting RF gun working at 1.3 GHz (L-band) that provides acceleration to an energy of 5.6 MeV. The rapid acceleration together with the focusing field produced by a built-in solenoid, contribute to the reduction of space-charge-induced emittance growth [Schreiber and Faatz, 2015].

SCRF modules and bunch compression. After the gun, further acceleration to a maximum (minimum) energy of 1.25 GeV (0.35 GeV) is provided by up to seven SCRF modules operating also at 1.3 GHz, which are bath-cooled by superfluid helium to 2 K and are designed to reach an accelerating gradient of 25 MV/m. A third-harmonic structure (ACC39) operating in decelerating mode [Edwards, Behrens and Harms, 2010; Vogel et al., 2010] is located after the first SCRF module (ACC1) in order to linearize the longitudinal phase space of the bunch for efficient longitudinal compression. This is composed of two magnetic chicanes working at nominal beam energies of 150 GeV (BC2, C-chicane) and 450 MeV (BC3, S-chicane), with an R_{56} of 180 mm and 43 mm, respectively [Limberg et al.,

[1] Ultra Violet

1996; Stulle, 2004; Schreiber and Faatz, 2015]. This setup delivers electron bunches with energies within a range of 0.35–1.25 GeV, bunch lengths of 50–5000 ps with peak currents of up to 2.5 kA, transverse normalised emittance between 0.7 and $\sim$ 2 µm, an uncorrelated energy spread below 0.1% and a tunable longitudinal phase space [Schreiber and Faatz, 2015]—all key features exploited at FLASHForward to enable the production of a double-bunch structure suitable for external-injection plasma-wakefield experimentation.

Time structure. The photo injector is able to deliver pulse trains at repetition rates of either 1, 5 or 10 Hz, each with a maximum length of 800 µs—limited by the RF-power system. These pulse trains are typically referred to as *macropulses*. The pulse spacing within a macropulse can be adjusted between 0.33 and 25 µs according to a discrete range of operating laser frequencies [Schreiber and Faatz, 2015]. The described pulse structure is schematically represented in Figure 2.2a.

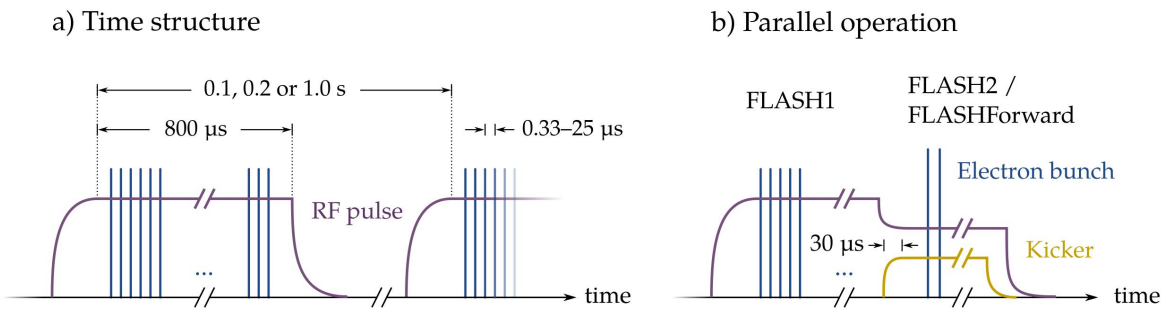

Figure 2.2.: Bunch train patterns delivered by the FLASH linac. a) RF pulses (purple) with a maximum flat-top length of 800 µs are produced at repetition rates of 0.1, 0.2 or 1.0Hz. Each macropulse can be filled with electron bunches (blue vertical lines) with spacings of 0.33–25 µs. b) Possible bunch pattern for parallel operation. The macropulse is devided into two bunch-train parts with different destinations and adjustable RF-amplitude and RF-phase for optimised bunch compression. Customised electron-bunch charge and spacing is achieved by operating each bunch-train part with a different UV laser. The kicker provides a minimum separation of 30 µs.

Parallel operation. In order to achieve the parallel operation of FLASH2, three kickers providing a fast separation of 30 µs are used in combination with a Lambertson septum deflecting in the horizontal plane (cf. Figure 2.1) to extract a fraction of the electron bunches contained in a macropulse after acceleration, thus dividing a bunch train into two consecutive (bunch-train) parts or sub-trains. The bunch parameters of each bunch-train part can be independently adjusted to fulfil the experimental requirements of each beamline. Control over the charge is achieved by operating the photo-cathode with independent UV lasers, which also enables the customisation of the number of bunches and the intra-bunch separation. Additionally, the RF-amplitude and phase of each RF-station can be modified separately to generate different longitudinal compressions. The resulting time structure of a bunch train with two parts corresponding to parallel operation is exemplified in Figure 2.2b. With respect to the FEL-pulse, the variable gap undulators of FLASH2 allow the tuning of the wavelength

within a factor of ~ 4 while using a fixed bunch energy dictated by the FLASH1-beamline requirements [Schreiber and Faatz, 2015].

Electron-beam diagnostics. Various diagnostics are integrated along the FLASH linac, including toroidal current transformers to measure the bunch charge, beam-loss monitors (BLM) to measure the radiation levels induced by charge losses, beam-position monitors (BPM) to measure the transverse centroid of the bunch, and beam arrival-time monitors (BAM) [Viti et al., 2017] to measure the timing of the bunches with respect to a reference signal provided by the synchronisation system. A FODO section located after the first bunch compressor BC2 with 4 screen stations is used to measure the transverse emittances and match the beam envelope to the lattice optics (cf. Section 1.2) [Loehl, 2005]. Longitudinal beam diagnostics relevant to the experiments presented in this thesis are the bunch-compression monitors (BCM) and a TDS (cf. Section 1.3).

- **BCM:** the BCM uses the total intensity of coherent diffraction radiation (CDR) emitted by an electron bunch to estimate its length [Wesch, 2012]. Two BCMs are installed downstream of BC2 and BC3 (green circles in Fig. 2.1) and are routinely used to provide a feedback signal for the automated stabilisation of the bunch compression.
- **LOLA-TDS:** the TDS, which is installed upstream of the FLASH1 undulators (cf. Fig. 2.1), is a LOLA-type disc-loaded RF waveguide structure [Altenmueller, Larsen and Loew, 1964] operating at a frequency of 2.856 GHz and streaking the bunch in the vertical direction. Under optimal operating conditions, it is able to achieve time resolutions of ~ 8 fs [Behrens, Gerasimova et al., 2012]. In combination with a dispersive section equipped with an energy spectrometer, the LOLA-TDS enables the measurement of the longitudinal phase space of the bunch. This feature is routinely used at FLASHForward in the first stages of the experimental setup to fine tune the linac parameters (cf. Section 4.1).

Synchronisation and beam stabilisation. The synchronisation framework at FLASH is based on a pulsed-laser optical system using stabilised fiber links for the timing distribution and optical cross-correlators for the synchronization of external Ti:sapphire lasers [Schulz, 2011]. A master laser oscillator (MLO) is phase-locked to an oven-controlled (quartz) crystal oscillator (OCXO) operating at 9.0277775 MHz—corresponding to the reference used in the old RF-based synchronisation scheme [Lorbeer, 2006]. Three BAMs installed before and after BC2 and after BC3 are used to measure the timing of the electron bunch with respect to the MLO, which, in combination with an active feedback system, is able to achieve a beam arrival-time jitter of ~ 30 fs [Schulz et al., 2015]. A dedicated low-level RF (LLRF) system stabilizes and flattens the amplitude and phase of the accelerating fields produced in the SCRF modules, resulting in a rms RF-amplitude stability of better than 0.005 %, and a RF-phase stability better than 0.01 deg (at 1.3 GHz) [C. Schmidt et al., 2014]. The BCMs located after BC2 and BC3 are used to stabilise the bunch compression by means of a slow

RF-feedback that controls the RF-stations of the acceleration modules. In a similar way, the beam orbit measured by the BPMs is stabilised by means of an automated control system that actuates appropriate steerer magnets according to the lattice optics, thus compensating eventual slow machine drifts.

2.1.2. FLASH2 extraction arc

With a length of approximately 30 m, the FLASH2 extraction arc [Scholz, 2013] achieves a total deflection angle of 12 deg and a vertical offset of $\sim$ 36 mm with respect to the linac. The kicker-septum system introduces both vertical and horizontal dispersion. The horizontal dispersion is closed after the last two horizontal bends, whereas the vertical dispersion is closed by means of a vertical mini-dogleg comprising two adjacent dipoles located shortly after the last two horizontal bends. The longitudinal dispersion R_{56} at the end of the arc is zero, implying that the bunch compression generated in the linac is preserved after the extraction. This is achieved by means of a horizontal bend with a negative angle—typically referred to as a "reverse bend"—located roughly in the middle of the arc. In order to mitigate CSR effects, the horizontal beta function has a small waist at the septum, at the last horizontal bends, and a few meters before the end of the arc, corresponding to the location of the first FLASHForward extraction dipole (cf. Section 2.2). However, due to the large bending radius of both the septum and the last horizontal bends, these effects are in general non-negligle and, depending on the beam current, might still have a large impact on the beam quality (cf. Chapter 4).

2.2. FLASHForward

2.2.1. Design

The FLASHForward beamline is installed in the FLASH2 experimental area (tunnel) and consists of five sections: a short extraction (FLFEXTR) with a horizontal bending angle of 8.0 deg; a longer dispersive section enabling further bunch compression (FLFCOMP), providing 4 m transverse separation with respect to the FLASH2 beamline and featuring a collimator device for the production of a double-bunch structure; a straight section for beam matching and final focusing into the plasma (FLFMAFF); a central plasma-interaction chamber; a post-plasma diagnostic section (FLFDIAG) equipped with two dipole spectrometers. The conceptual design is schematically depicted in Figure 2.3—in which the distances between the elements are roughly to scale. In the following, key aspects about each of these sections are provided.

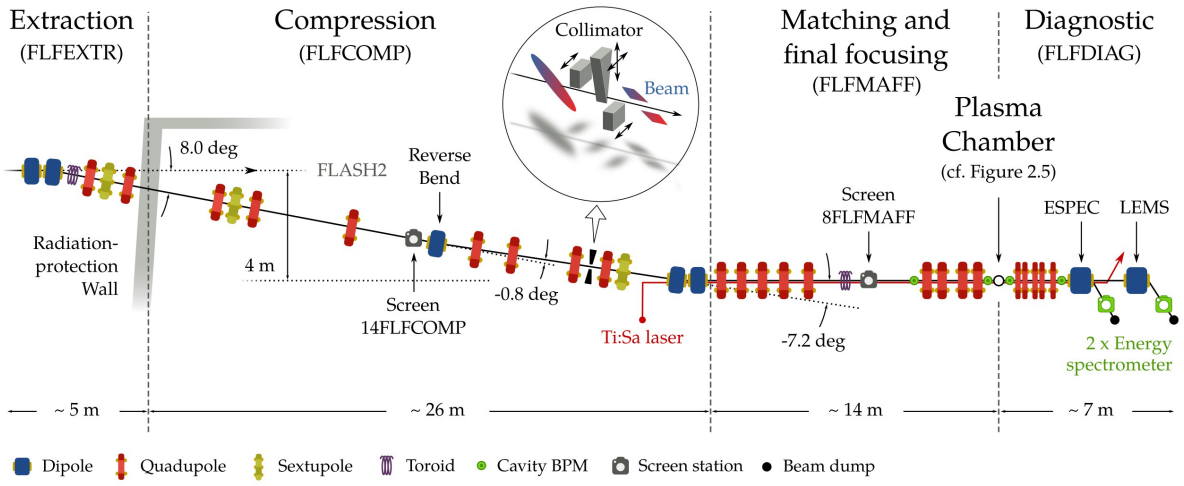

Figure 2.3.: Layout of the FLASHForward beamline with its five sections: a short extraction (FLFEXTR) with a horizontal bending angle of 8.0 deg; a longer dispersive section enabling a further bunch compression (FLFCOMP), providing 4 m transverse separation with respect to the FLASH2 beamline and featuring a collimator device for the production of a double-bunch structure; a straight section for beam matching and final focusing to the plasma (FLFMAFF); a central plasma-interaction chamber; a post-plasma diagnostic section (FLFDIAG) equipped with two dipole spectrometers.

Extraction - FLFEXTR. The extraction is composed of two adjacent horizontal dipole magnets followed by two quadrupoles and a sextupole. One of the major constraints imposed on the design of this section is the presence of a 2-m-thick radiation-safety wall that allows access to the FLASH2 tunnel independent of the FLASH1 tunnel. On the one hand, the dipoles must provide a large bending angle that achieves enough transverse separation before the wall to accomodate the additional magnets required for proper beam transport. On the other hand, the CSR effects generated by the large beam currents through bend sections must be kept at tolerable levels to avoid an excessive degradation of the beam quality. A total bending angle of 8 deg was considered to be a good compromise between these two constraints [Behrens, 2015].

Compression - FLFCOMP. The compression section was designed to fulfil multiple requirements simultaneously [Behrens, 2015]:

1. achieve a transverse offset of 4 m with respect to the FLASH2 beamline, which, according to the bending angle of the extraction section, requires the beam to propagate a total length of $\sim$ 26 m;
2. accomplish an achromatic beam translation—i.e., closed first- and second-order transverse and angular dispersions, and mitigation of chromatic effects—and no net deflection at the end of the section, which is achieved by the last two bends in combination with an appropriate lattice optic composed of first- and second-order imaging elements—i.e., quadrupole and sextupole magnets;
3. mitigate CSR effects by adjusting the lattice optic to form a small waist at the last two

bends;
4. provide a tunable R_{56} to increase control over the longitudinal compression independent of FLASH linac settings by means of a reverse bend;
5. enable a precise energy collimation by positioning the collimator device at a location of large dispersion and a small beta function—i.e., a location with a good energy resolution $R_\delta = \sigma_{x,\beta}/|D_x|$.

As mentioned in Chapter 1, the proper operation of the collimator relies on the linearisation of the (negative) energy chirp of the bunch to induce a transverse beam tilt in the presence of dispersion. This is schematically represented in Figure 2.3. The collimator device consists of a three-component mask: two horizontally adjustable blocks to remove the low- (head) and high-energy (tail) parts of the bunch, and a vertically and horizontally adjustable wedge-shaped collimator to split it into a driver and trailing bunch pair [Schröder, Ludwig et al., 2020]. Apart from the production of a double-bunch structure, such a device enables further diagnostics capabilities, most prominently a high-resolution wakefield-sampling technique [Schröder, Lindstrøm et al., 2020] and, in combination with the two-BPM-tomography technique (cf. Section 1.2), a sliced-beam-parameters analysis. An example of the latter is discussed in Section 4.2.3. The titanium-sapphire (Ti:Sa) laser used for X-1 experiments is coupled into the beamline at the end of the compression section, which enables colinear propagation of the electron and laser beams (cf. Figure 2.3). The design optic of FLFCOMP is optimised together with that of FLFEXTR (cf. Figure 2.6).

Matching and final focusing - FLFMAFF. The straight section before the plasma-interaction chamber is divided into two subsections. The first part includes five quadrupoles and a screen station, enabling the measurement of the transverse rms emittances and match the beam to an appropriate lattice optic compatible with a small focus at the interaction point. The second part contains four quadrupoles that complete the final-focusing optics. In order to match a 1-GeV beam to the transverse focusing fields of a wakefield cavity excited in a plasma with an electron density of $n_0 = 1.1 \cdot 10^{15}$ cm^{-3}, for instance, the beta function at the focal point should be $\beta_{x/y} = 10$ mm. A possible lattice optic producing such a small focus is shown in Figure 2.6, where the beta functions at the final-focusing quadrupoles quickly rise to values above $\beta \approx 300$ m after the hard focus at the plasma cell. The quadrupole strengths, on the other hand, reach values of up to $k_1 \approx 10$ m^{-2}, which, in combination with the large beta functions and for beams with a finite energy spread, introduces chromatic effects. Figure 2.4 shows the phase-space ellipses of beam slices with energy deviations of $\Delta\delta = \pm 0.25$ %, normalised to the (design) Courant-Snyder parameters of the nominal-energy beam $\beta_D = 10$ mm and $\alpha_D = 0$. In the horizontal plane, slices with higher (lower) energies exhibit a mismatch of $B_{mag} = 1.13$ ($B_{mag} = 1.12$), which translates into a shift of the waist location of $\Delta s = +5$ mm ($\Delta s = -5$ mm)—i.e., the tail and head of the bunch are focused downstream and upstream of the design focal point, respectively. This behaviour is more pronounced in the vertical plane, since the beta function is larger at the last two quadrupoles, which are the ones with larger focusing strengths. As a result, the waist location is shifted $\Delta s = +13$ mm and $\Delta s = -13$ mm for higher (tail) and lower (head) energies, respectively.

Such chromatic effects might compromise the quality of the plasma-acceleration process and must be considered when optimising the final-focusing lattice optics. An example of this optimisation process is discussed in detail in Chapter 4.

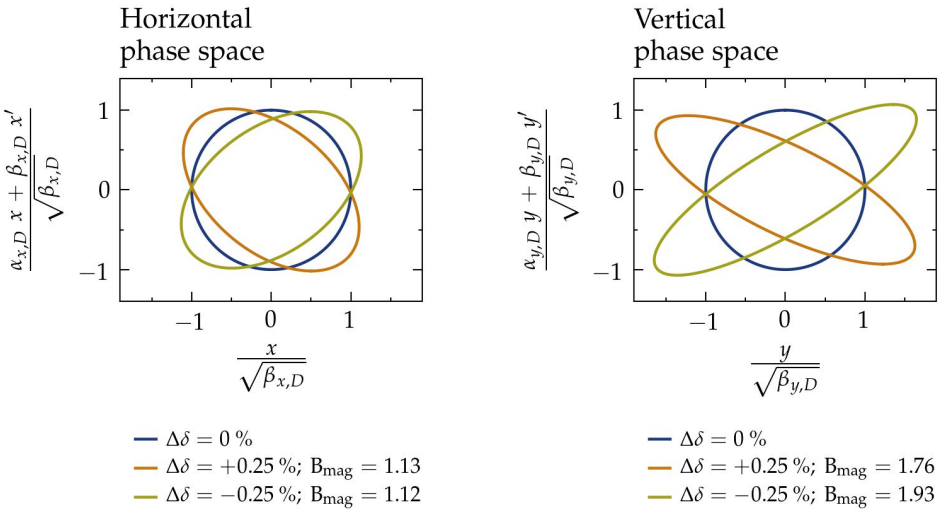

Figure 2.4.: Chromatic effects at a focal point with design Courant-Snyder parameters $\beta_D = 10$ mm and $\alpha_D = 0$, corresponding to a beam matched to an electron-plasma density of $n_0 = 1.1 \cdot 10^{15}$ cm^{-3}. The degree of mismatch is represented by the phase-space ellipses of beam slices with energy deviations of $\Delta\delta = \pm 0.25$ % normalised to the (design) Courant-Snyder parameters at a nominal-energy of $E = 1$ GeV.

Plasma chamber and gas cell. The beam-plasma-interaction chamber is the centerpiece of the beamline (cf. Figure 2.5a). Its bespoke design consists of two stacked cylindrical vacuum chambers with diameters of 500 mm. The upper one is directly connected to the electron beamline, which necessitates a series a differential pumping stages along the beam pipes attached to it to progressively adapt the vacuum levels to the UHV of the accelerator (cf. Section 2.4). It contains the gas cell and various diagnostics for spatial alignment of the laser and electron beams mounted on a baseplate, and incorporates multiple side ports to provide additional access for transverse laser pulses used in internal-injection (X-1) experimentation as well as clear lines of sight for optical diagnostic systems. The bottom chamber is kept at medium-vacuum (MV) conditions at roughly $\sim 10^{-3}$ mbar and contains a hexapod positioning system mechanically connected to the baseplate of the upper chamber by means of a feedthrough-bellow system, providing precisely controlled rotational and translational movement of the plasma cell and diagnostics.

The gas cell, schematically depicted in Figure 2.5b, consists of a stack of two sapphire slabs held together by a PEEK[2] holder in which two gas inlets connected by a thin channel with a diameter of $\mathcal{O}(1\text{ mm})$ are milled. Both ends of the channel are directly coupled to the background vacuum of the plasma chamber. When equal gas pressure is continually

[2]Polyether ether ketone.

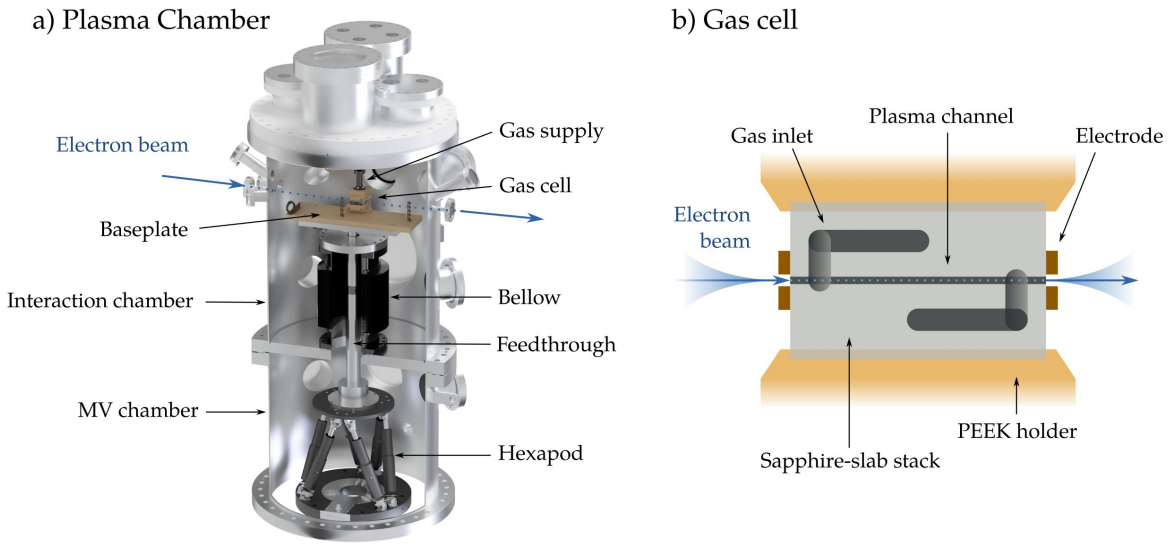

Figure 2.5.: Plasma chamber and gas cell. a) Side-view of the plasma chamber with its two stacked cylindrical vacuum volumes. The upper volume (interaction chamber) is directly connected to the electron beamline and contains a baseplate with the gas cell and various diagnostics for spatial alignment of the laser and electron beams. The lower volume (MV chamber) contains the hexapod connected to the baseplate by means of a feedthrough-bellow system. (Source: Schröder, 2021; CAD model adapted by Sarah Schöder.) b) The gas cell consists of a stack of two milled sapphire slabs held together by a PEEK holder. Two electrodes at each side of the central channel initiate a discharge triggered by means of a high-voltage thyratron front-end, thus promoting the formation of a plasma capillary.

applied to the inlets, a constant neutral gas density profile exists in the central channel [Garland et al., 2021]. Two electrodes attached at each side of the channel are used to initiate a discharge triggered by means of a high-voltage thyratron front-end, thus promoting the formation of a plasma capillary inside the channel, through which the electron beam is subsequently sent. Such a windowless gas-cell design has the advantage of avoiding the degradation of the beam quality that would be otherwise induced by particle scattering through a material interface. In addition to providing a hard-surface environment for the confinement of the plasma, the sapphire material allows the transmission of light emitted during the recombination of the plasma, thus facilitating the application of spectroscopic techniques for plasma diagnostics [Garland et al., 2021]. Different cell geometries with channel lengths of 33, 50 and 195 mm are available, which provide some flexibility to adjust the duration of the beam-plasma interaction. The cells can be operated with various gas species up to plasma-electron densities of $\mathcal{O}(10^{17}\ \mathrm{cm}^{-3})$. The temporal evolution of the density at the center of the channel follows an approximate exponential decay, such that the plasma density value at the time of interaction may be selected by delaying the trigger signal of the discharge with respect to the arrival time of the bunch. The longitudinal density profile, on the other hand depends on the operating conditions of the cell, its geometry, and the relative time of interaction after the discharge, and its shape can range from a flat-top with upstream and downstream density ramps to Gaussian-like profiles [Garland et al., 2021].

Diagnostics - FLFDIAG. The diagnostic section immediately downstream of the plasma chamber accomodates five high-strength quadrupoles operating as a triplet to capture the high-divergence beams coming out of the plasma capillary, followed by a broadband energy spectrometer (ESPEC) to image both the driver and trailing bunch simulataneously and a second narrowband energy spectrometer with higher spatial resolution for low-emittance measurements (LEMS). Details about the diagnostic capabilities of each of the two spectrometers are described in Section 2.2.2. Figure 2.6 shows a possible lattice optic to capture the high-divergence beam coming out of the interaction point. As can be observed, the nature of the beam evolution is the reverse of that seen in the final-focusing section, with beta functions above 300 m and large focusing strengths. In this case, however, the quadrupole triplet not only has to capture the beam, but it also needs to fulfil the requirements imposed by the measurement to be perfomed by either of the two energy spectrometers. The configuration

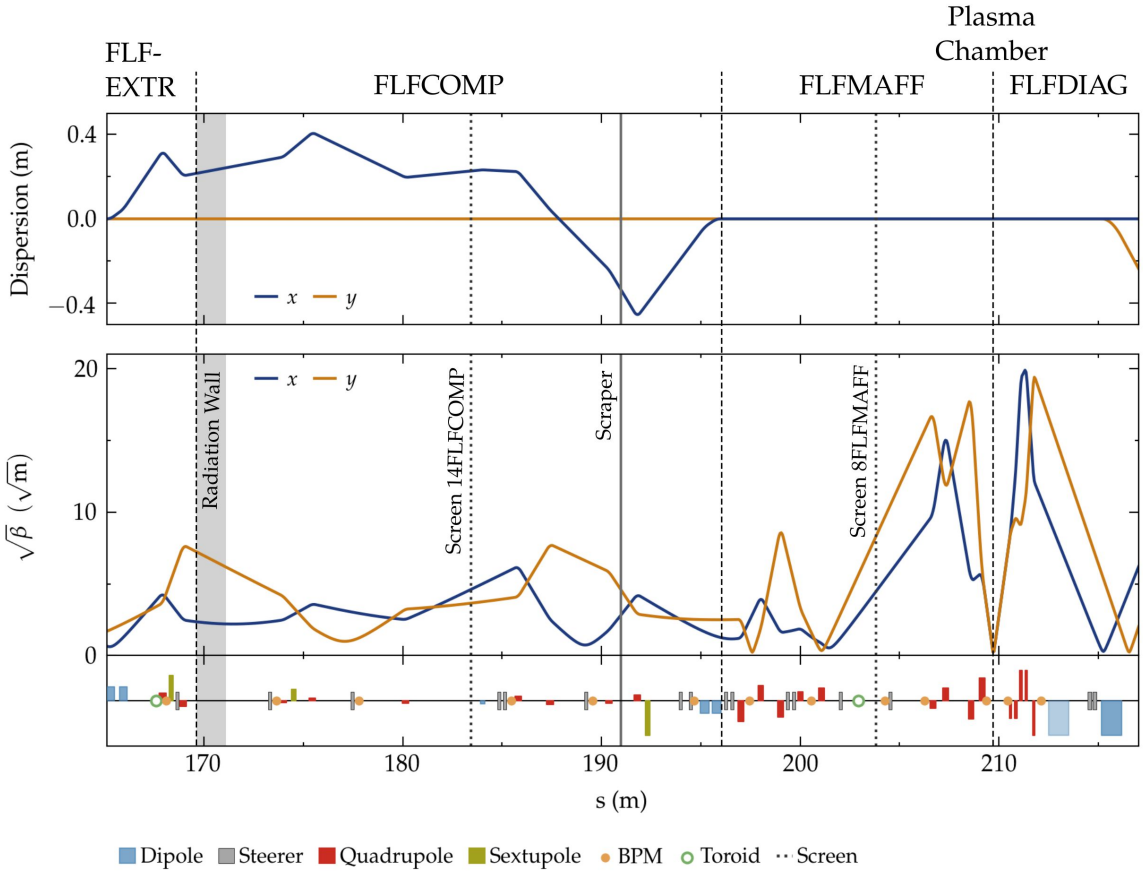

Figure 2.6.: Dispersion (upper plot) and beta functions (bottom plot) corresponding to the design lattice optic of the FLFEXTR and FLFCOMP sections and a possible optic at the FLFMAFF and FLFDIAG sections for PWFA experiments and diagnostics. The horizontal dispersion created at the extraction dipole is closed at the end of the compression section (FLFCOMP). The scraper is installed at $s \approx 191$ m in a location with large dispersion and a moderate beta function for the efficient production of a double-bunch structure. The lattice in the FLFMAFF section is used to focus the beam to beta functions on the order of $\beta_{x/y} \approx \mathcal{O}(10\text{ mm})$ at the interaction point (plasma chamber). The beta functions around the focal point can easily reach values above $\beta_{x/y} \approx 300$ m. The FLFDIAG section is used to capture the beam and achieve a proper transport to the spectrometers.

shown in Figure 2.6 would be appropriate for the measurement of the energy spectrum at the LEMS, while achieving a good resolution of the transverse beam size.

2.2.2. Electron-beam diagnostics

The beamline is equipped with regular electron-beam diagnostics, including toroidal current transformers, beam-loss monitors (BLM), beam-position monitors (BPM), screen stations, and two dedicated magnet dipole energy spectrometers. Relevant aspects of some of these diagnostics are described in the following paragraphs.

BPMs. Different types of BPMs are integrated along the beamline: stripline, button and cavity BPMs, all of them enabling the diagnostic of electron bunches with charges down to around $\lesssim 100$ pC. Stripline [Baboi et al., 2006] and button BPMs [Treyer et al., 2013] provide spatial resolutions of 10 µm and ~ 5 µm, respectively. Cavity BPMs [Lipka, 2009; Lipka et al., 2014], in contrast, can achieve resolutions in the sub-µm range and are therefore the type used around the plasma chamber, where a high-precision monitoring of the bunch pointing jitter is required. Additionally, this setup enables the application of the 2-BPM-tomography technique (cf. Section 1.2), for which a sub-µm resolution is a prerequisite for the reliable estimation of the pointing jitter at the location of the waist [Lindstrøm, D'Arcy et al., 2020].

Screen stations. The screen stations installed at FLASHForward were originally developed for the Europeran-XFEL [Wiebers et al., 2013] and use inorganic scintillator materials in order to avoid coherent effects in the emission of optical transition radiation (OTR) triggered by high-brightness beams in conventional OTR monitors. The scincillator materials used are Cerium-doped Lutetium Yttrium Orthosilicate (LYSO:Ce) and Cerium-doped Gadolinium Aluminium Gallium Garnet (GAGG:Ce). A technical drawing of the core components of the screen station is shown in Figure 2.7. A movable screen holder accomodating up to two different screens and a test chart for calibration purposes is positioned to intercept the beam perpendicularly to the direction of propagation. The emitted light is observed at 45 deg and is imaged out of vacuum into a CCD camera by means of an optical system exploiting the Scheimpflug principle to minimize depth-of-focus effects. The camera used is a Basler avA2300-25gm with a pixel size of 5 µm $\times$ 5 µm and a total of 2330 $\times$ 1750 pixel. The field of view (FOV) can be manually adjusted to obtain a magnification power of either 1:1 or 1:2. The nominal spatial resolution of the system is $R_{\mathrm{scr}} = 10$ µm and can be reduced further down to $\mathcal{O}(1\ \mu\mathrm{m})$ by applying deconvolution algorithms and appropriate single-particle resolution functions [Kube et al., 2015]. Two screen stations with a magnification power of 1:1 are located at the FLFCOMP and at the FLFMAFF sections for transverse beam profile diagnostics.

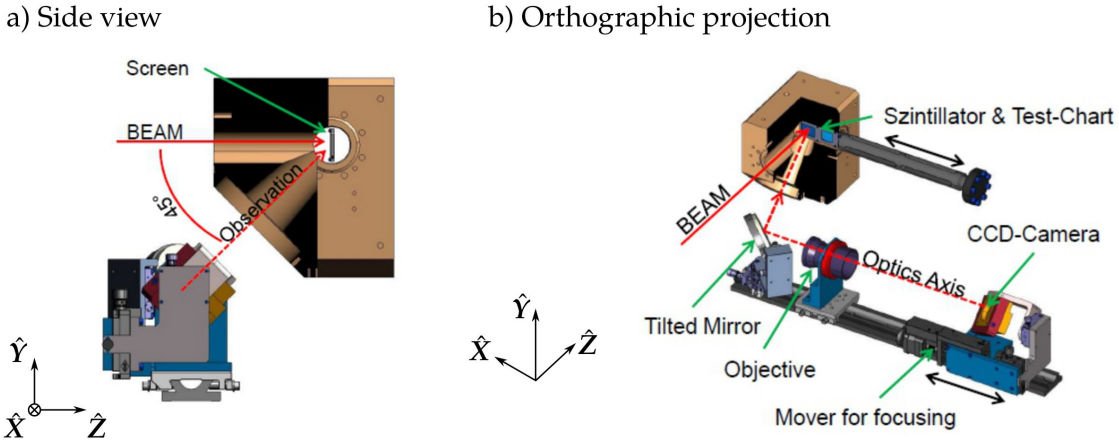

Figure 2.7.: Technical drawing of the screen stations used at FLASHForward. A movable screen holder is positioned to intercept the beam perpendicularly to the direction of propagation. The emitted light is observed at 45 deg and is imaged out of vaccum into a CCD camera by means of an optical system exploiting the Scheimpflug principle to minimize depth-of-focus effects. The movable imaging system can be adjusted along the optical axis to produce magnification powers of 1:1 or 1:2. Image reproduced from Wiebers et al., 2013.

Broadband energy spectrometer - ESPEC. As a result of the beam-plasma interaction, a substantial fraction of the driver-bunch energy is transferred to the trailing bunch, eventually depleting the energy of the driver (cf. Section 1.4). In order to measure the large energy spreads associated to this process, a broad-band 1 m-long dipole spectrometer dispersing in the vertical plane is installed roughly 3 m downstream of the interaction point. A large 1 mm-thick stainless-steel vacuum chamber is embedded between the pole shoes with two scintillating screens (Kodak LANEX Fine) attached to the outer side of its downstream and bottom walls. The screens are imaged by means of a periscopic system and 6 separate CCD cameras (model Basler acA2040-35gm, with a pixel size 3.45 µm × 3.45 µm and a total of 2048 × 1536 pixel). The spatial resolution of the system is limited by scattering in the chamber wall to about 50 µm [Lindstrøm, 2019b]. With maximum dispersions on the order of $D_y \approx 1$ m, the best energy resolutions achievable are about $R_\delta \approx 0.005$ %. However, the beam transport through the high-strength quadrupole triplet is subject to chromatic effects, which prevents the resolution of the entire spectrum of beams with large energy spreads. For such beams, a precise measurement of the full spectrum requires the imaging energy of the quadrupoles to be progressively scanned and to subsequently splice the measured spectra with, e.g., a Cauchy-Lorentz distribution. For beams with a small waist around the interaction point, the transverse emittance can be measured in the nondispersive plane by means of an object-plane scan (cf. Section 1.2). However, the limited demagnification power of the lattice and the poor spatial resolution of the screen constrain the range of measurable emittances to values above $\gtrsim 10$ µm [Lindstrøm, 2019b]. This limitation is overcome with the LEMS spectrometer, which contains a different screen setup and is installed approximately 7 m downstream of the interaction point.

Low-emittance measurement station - LEMS. The LEMS design is based on the same type of dipole magnet used in the ESPEC and also disperses the beam in the vertical plane. Since its purpose is to measure the beam parameters of either noninteracted bunches or plasma-accelerated trailing bunches, the station only needs to cover a small energy range on the order of $\Delta\delta \approx 10$ %, which greatly simplifies the design of the vacuum system to accomodate the optical-imaging system. An XFEL-type screen station with an in-vacuum GAGG scintillator is installed on a short beam pipe at an angle of approximately -10 deg with respect to the undeflected-beam trajectory and $\sim$ 1 m downstream of the dipole. Due to the longer distance from the interaction point, the magnetic lattice can achieve demagnification powers of up to $R_{11} \approx -12$, which, in combination with a measured screen resolution of $R_{\mathrm{scr}} = 7$ µm, results in an emittance resolution of 0.1 µm [Lindstrøm, Beinortaite et al., 2021]. In the vertical plane, a nominal dispersion of approximately $|D_y| \approx 250$ mm provides an energy resolution of $R_\delta \approx 0.003$ %.

2.3. PolariX-TDS beamline

The motivation to extend the diagnostic capabilities of FLASHForward with an X-band TDS is manifold. On the one hand, collective effects arising along the FLASH2 extraction and the FLASHForward extraction and compression sections can result in: 1) distortions of the linearity of the energy chirp of the bunch; 2) transverse-slice centroid offsets. The presence of a TDS at the end of the FLASHForward beamline enables the precise assessment of these effects in order to take the appropriate actions to mitigate them (cf. Section 4.2). In this sense, the ability of the PolariX-TDS system to adjust the polarisation of the streaking field represents a great advantage, since the large transverse gradients present in a plasma wakefield require the mitigation of residual transverse slice centroids in both planes (cf. Section 3.4). On the other hand, the visualisation of the longitudinal phase space of the beam downstream of the collimator device is required to precisely characterise and optimise its performance—which is critical for the success of external-injection (X-2) experimentation (cf. Section 4.2). Finally, a TDS operating in the X-band range located after the interaction chamber enables the diagnostic of plasma-accelerated electron bunches with femtosecond-scale time resolution, thus providing direct information about both the transverse and the longitudinal structure of the plasma wakefields (cf. Chapter 5).

In this section, key aspects of the PolariX-TDS beamline design are discussed, which are relevant to the diagnostics capabilities of the device. Details about the RF-hardware of the PolariX-TDS and its dedicated commissioning are presented in Chapter 3. A short summary of the commissioning of the beamline components concludes the section.

2.3.1. Design

The major constraint imposed on the beamline design is dictated by the location of the RF-power station of the PolariX-TDS system, which must be close enough to the structure to minimise the power losses produced along the waveguide network (cf. Section 3.1). As part of the PolariX-TDS collaboration, both FLASHForward and FLASH2 were interested in the acquisition of a TDS, which, in order to reduce financial costs, were foreseen to share the RF-power system. On account of that and due to the limited space availability in the pre-existing FLASH2 beamline, the TDS had to be located approximately $\sim$ 25 m downstream of the plasma chamber.

The magnetic lattice around the TDS is organised as schematically depicted in Figure 2.8. Four quadrupoles separated by $\sim$ 5 m-long gaps are located upstream of the structure in order to allow for the future installation of three undulators with a view towards the generation of FEL-radiation. These four quadrupoles together with the high-strength triplet of the FLFDIAG section provide enough flexibility to adjust the lattice optic to the requirements of the different types of measurements performed with the TDS—viz. slice transverse beam parameters and slice emittance in x and y, and longitudinal-phase-space measurements. The slice transverse beam parameters are measured at a screen station located at the end of a short straight section downstream of the structure comprising two quadrupoles. After that, a dispersive section is dedicated to the diagnostic of the longitudinal phase space, which includes a dipole introducing a horizontal deflection of -5 deg, a quadrupole to add flexibility to the simultaneous optimisation of the longitudinal and energy resolutions, and a second screen station. A provisional beam dump is located at the end of the beamline. The longitudinal and energy resolutions achievable by means of this lattice arrangement are provided in the following paragraphs.

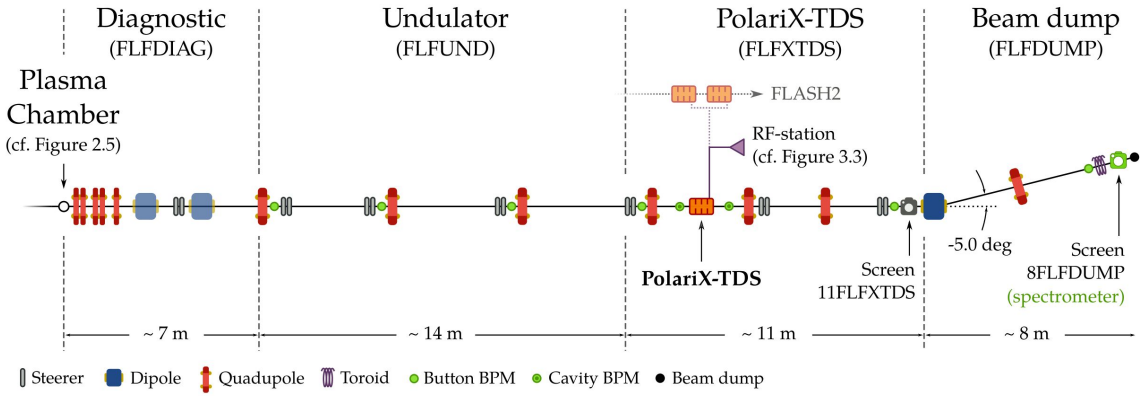

Figure 2.8.: Layout of the PolariX-TDS diagnostics beamline consisting of three sections: FLFUND, FLFXTDS and FLFDUMP. The TDS is located $\sim$ 25 m downstream of the plasma chamber, near the two structures installed at FLASH2 and the shared RF-power station. Four quadrupoles upstream of the structure allow the future installation of three undulators with a view towards the generation of FEL-radiation. The straight section downstream of the TDS is dedicated to the measurement of the transverse sliced beam parameters. A dispersive section with an additional quadrupole is used for the diagnostic of the longitudinal phase space.

Longitudinal resolution

As discussed in Section 1.3.2, the longitudinal resolution is optimal when the beam is sent collimated through the TDS and has a large beta function compatible with a well behaved optic along the beamline—i.e., not too large as to cause transmission problems or to produce chromatic effects. Considering the iris radius of the structure $a = 4$ mm (cf. Table 3.1) and a 1 GeV-beam with a normalised emittance of $\epsilon_{y,n} = 1$ µm, a beta function of $\beta_{y,\mathrm{TDS}} = 100$ m is found to be a good compromise, resulting in 99.7 % of the particles (i.e., $6\sigma_{y,rms}$) being contained in $\sim$ 17 % of the TDS aperture. Such a beta function is easily achieved with the quadrupoles upstream of the TDS. On the other hand, the phase advance between the TDS and the screen is required to be $\Delta\psi_{y,\mathrm{TDS}\to\mathrm{scr}} = 90$ deg, which is accomplished with the quadrupoles downstream of the structure. Since there are (at least) two quadrupoles between the TDS and the screen, an additional degree of freedom is available to adjust the beta function in the streaking plane at either of the two screens in order to: 1) appropriately resolve the unstreaked beam size; 2) fit the streaked beam size in the FOV of the screen while simultaneouly allowing an RF-phase scan to calibrate the shear parameter (cf. Section 2.3.2). A possible lattice optic for each of the screen stations and a symmetric unfocused beam at the plasma chamber with $\beta_x = \beta_y = 10$ m and $\alpha_x = \alpha_y = 0.0$ is shown in Figure 2.9. Lattice optics suitable for the diagnotic of beams with a small waist at the interaction point are discussed in Chapters 4 and 5. Notice that to adjust the horizontal beta function at the 11FLFXTDS screen, the quadrupoles upstream of the TDS must be used, since those downstream of the TDS are used to optimise the longitudinal resolution. In the optic shown in Figure 2.9, a beta

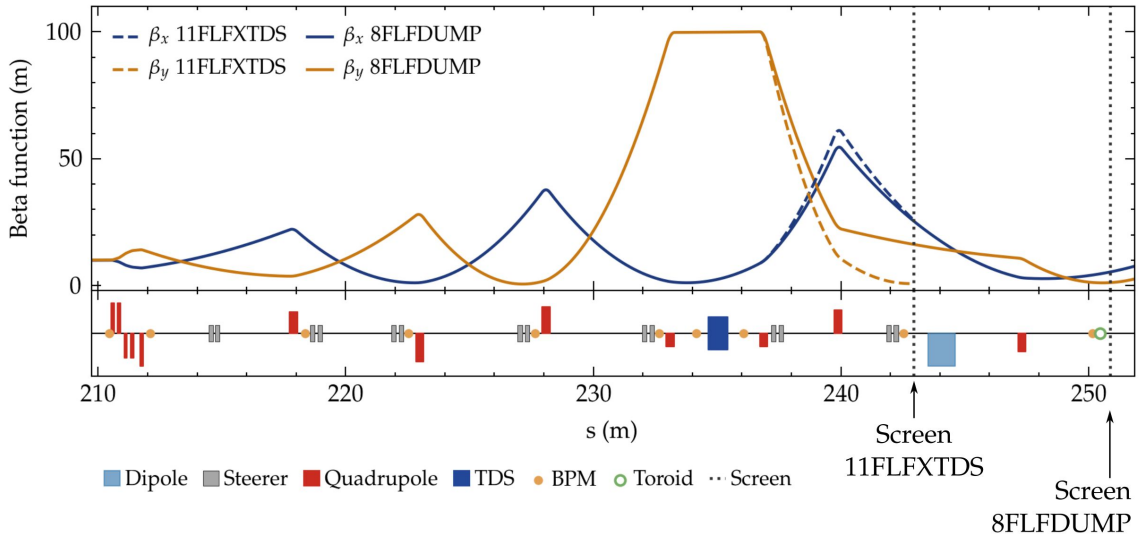

Figure 2.9.: Design lattice optic of the PolariX-TDS diagnostics beamline for a symmetric unfocused beam at the plasma chamber with $\beta_x = \beta_y = 10$ m and $\alpha_x = \alpha_y = 0.0$. The beta function in the streaking plane at the location of the TDS is typically $\beta_{y,\mathrm{TDS}} \approx 100$ m. The quadrupoles downstream of the structure are adjusted to achieve a phase advance between the TDS and the screen of $\Delta\psi_{y,\mathrm{TDS}\to\mathrm{scr}} = 90$ deg at either of the two screen stations to optimise the longitudinal resolution.

of $\beta_{x,\mathrm{scr}} \approx 25$ m is achieved, which provides good resolution of the beam size as required for the measurement of the sliced beam parameters. At the 8FLFDUMP screen, as mentioned before, the presence of an additional quadrupole grants more flexibility to simultaneously adjust the optics in the horizontal plane to achieve a good energy resolution.

The remaining terms contributing to the longitudinal resolution not related to the lattice optic are the integrated TDS voltage V_{TDS}, the beam energy E and the operating RF frequency $f_{RF} = 11.9888$ GHz. The integrated TDS voltage depends on the available RF-power and the power-to-voltage constant of the structure, which has a nominal value of $D = 5.2\ \mathrm{MV}/\sqrt{\mathrm{MW}}$ (cf. Table 3.1). As will be discussed in Section 3.1, the klystron is able to deliver up to $P = 6$ MW, and this value will be increased to $P = 24$ MW with the future installation of an RF-pulse compressor. The resolutions achievable accross the full RF-power range and different beam parameters are shown in Figure 2.10. The best resolution for a 1 GeV-beam with a normalised emittance of $\epsilon_{y,n} = 1$ µm provided by the system used in this thesis (klystron only) is below $R_\xi \lesssim 3$ fs. In practice, however, due to the losses produced in the RF-power transport from the klystron to the structure, this value is increased slightly above $R_\xi \gtrsim 3$ fs.

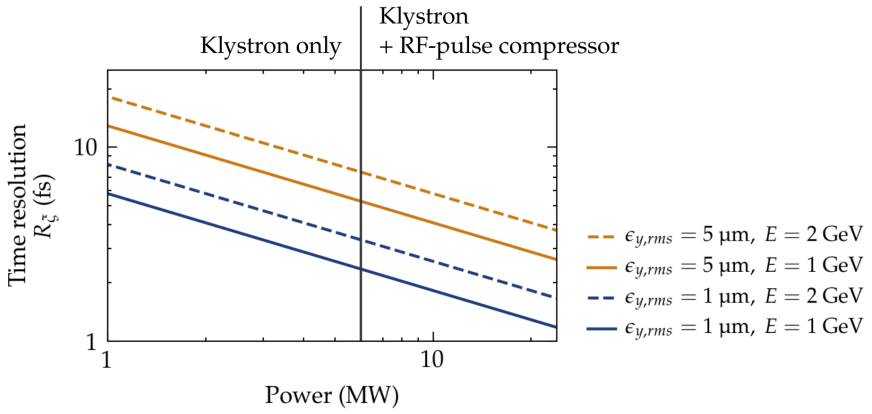

Figure 2.10.: Time resolution achievable at FLASHForward accross the full RF-power range and different beam parameters. The RF-power system used in the work presented in this thesis consists of only a klystron and is able to deliver a maximum of $P = 6$ MW, with which best-case longitudinal resolutions of $R_\xi \approx 3$ fs are achieved. The future installation of an RF-pulse compresseor will reduce the resolution further to slightly above $R_\xi \approx 1.0$ fs.

Energy resolution

The dispersions produced at the 8FLFDUMP screen are in the range $|D_x| = 0.7$–1.0 m, depending on the exact focusing strength of the quadrupole located between the dipole and the screen, which typically has negative values not lower than $k_1 \approx -1.0\ \mathrm{m}^{-2}$. The achievable energy resolutions in the most unfavourable case of having $D_x = 0.7$ m are shown in the left plot of Figure 2.11 for a range of beta functions at the screen and different beam parameters. As can be observed, for moderate betas of $\beta_{x,\mathrm{scr}} \approx 10$ m the resolution

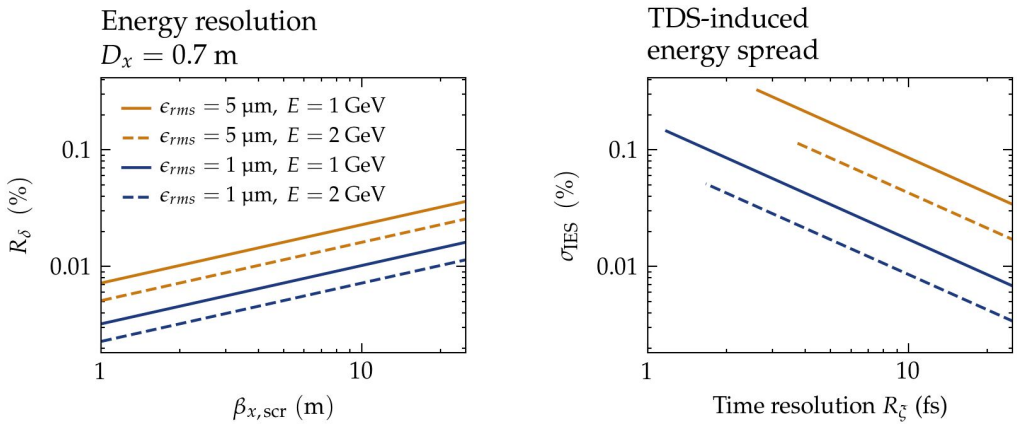

Figure 2.11.: Left plot: energy resolution achievable at the 8FLFDUMP screen with a dispersion of $D_x = 0.7$ m and different beam parameters. Moderate betas of $\beta_{x,\mathrm{scr}} \approx 10$ m produce resolutions around or well below $R_\delta \lesssim 0.01$ %, which enable to resolve the sub-per-mille uncorrelated energy spreads produced by the FLASH linac. Right plot: TDS-induced energy spread associated to the longitudinal resolutions shown in Figure 2.10. The induced energy spread is typically above the energy resolution threshold and can thus be well resolved.

is around or well below $R_\delta \lesssim 0.01$ % for all beam parameters, which is enough to resolve the sub-per-mille uncorrelated energy spreads produced by the FLASH linac (cf. Section 2.1). When performing longitudinal-phase-space measurements, however, the slice energy spread is broadened due to TDS-induced effects (cf. Section 1.3), which is typically the actual factor limiting the measurement. To illustrate this, the TDS-induced energy spread associated with the longitudinal resolutions discussed in the previous section is shown on the right plot of Figure 2.11, where it can be seen that for time resolutions below $R_{\tilde{\xi}} \lesssim 10$ fs the TDS effects exceed the energy resolutions achieved with the spectrometer—and thus can be well resolved. Notice that for the best time resolutions achievable the induced energy spread can be larger than the natural uncorrelated energy spread of the bunch.

Electron-beam diagnostics

The beamline is equipped with regular beam diagnostics, including a toroid located ≈ 0.5 m upstream of the screen 8FLFDUMP, plus several BLMs and BPMs. The BPMs around the PolariX-TDS are of the cavity type to allow a precise monitoring of the beam trajectory through the structure and avoid eventual transmission problems. XFEL-type screen stations (cf. Section 2.2) are installed donwstream of the PolariX-TDS, both including LYSO:Ce and GAGG:Ce scintillators. In order to guarantee their adequacy for bunch-length and energy-spectrum measurements over the whole range of experimental conditions expected at FLASHForward, a magnification ratio of 1:1 and 1:2 is chosen for the first (11FLFXTDS) and second (8FLFDUMP) station, corresponding to a field-of-view (FOV) of 9 mm × 12 mm and 18 mm × 24 mm, respectively. The motivation for these choices is described in the following section.

2.3.2. Jitter and tolerance studies

The accuracy of the bunch-length measurements performed with the TDS relies on the stability of the shear parameter S and the ability to precisely determine its magnitude. In the same way, the measurement of the energy spectrum relies on the calibration of the dipole spectrometer. These calibration routines can be compromised by spurious variations of the bunch centroid at the screen. Therefore an analysis of the different sources of jitter and their impact on the operation of the PolariX-TDS diagnostics beamline is required. At FLASHForward, the sources of jitter[3] can be divided into three categories:

1. RF stability of the PolariX-TDS system, which include the *phase* $\bar{\sigma}_\phi$ and the *amplitude* $\bar{\sigma}_A$ jitter;
2. linac stability, including the *arrival-time* $\bar{\sigma}_t$, the *pointing* $\bar{\sigma}_r$ and the *energy* jitter $\bar{\sigma}_\delta$;
3. stability of the plasma-acceleration process, which is reflected mainly on the *energy* jitter of the accelerated witness bunch $\bar{\sigma}_{\delta,\mathrm{PWFA}}$.

In practice, not all these sources can be completely disentangled from each other and their combined effect has to be considered. In this section, the aforementioned sources of jitter are defined and the relation between them is examined. After that, their impact on both the calibration of the shear parameter and the calibration of the dipole spectrometer is evaluated and appropriate parameters for the optics system of the two screen stations of the diagnostics beamline are determined accordingly.

RF-phase and linac arrival-time jitter. The phase jitter $\bar{\sigma}_\phi$ refers to spurious variations of the RF phase of the TDS system relative to the arrival time of the bunch. For a given shear parameter S, this translates into a pointing jitter at the screen:

$$\bar{\sigma}_{\langle y\rangle,\phi} = \frac{|S|}{k_{RF}}\bar{\sigma}_\phi. \tag{2.3.1}$$

Similarly, the linac arrival-time jitter $\bar{\sigma}_t$ results in centroid variations at the screen:

$$\bar{\sigma}_{\langle y\rangle,t} = |S|\, c\, \bar{\sigma}_t, \tag{2.3.2}$$

where c is the speed of light. Since the RF-phase and the linac arrival-time are inextricabaly linked by definition, their combined effect has to be considered, which leads to the expression:

$$\bar{\sigma}_{\langle y\rangle,\phi+t} \;=\; \frac{|S|}{k_{RF}}\sqrt{\bar{\sigma}_\phi^2 + (k_{RF}\, c\, \bar{\sigma}_t)^2} \;=\; \frac{|S|}{k_{RF}}\;\bar{\sigma}_{\phi,\mathrm{eff}}\;, \tag{2.3.3}$$

[3]The jitter parameters are indicated with a bar $\bar{\sigma}$ to differentiate from the conventional notation used for the second moments of the beam parameters.

where the *effective RF-phase* jitter has been defined. At FLASH, the arrival-time jitter is estimated to be around $\bar{\sigma}_t \approx 30$ fs [Schreiber and Faatz, 2015], which corresponds to ≈ 0.13 deg at 11.988 GHz. The effective RF-phase jitter of the PolairX-TDS system is estimated to be $\bar{\sigma}_{\phi,\mathrm{eff}} = 0.25$ deg, slightly above the arrival-time jitter of FLASH.

RF-amplitude jitter. The amplitude jitter $\bar{\sigma}_A$ reflects eventual variations of the RF-signal amplitude that directly translate into variations of the absolute value of the shear parameter:

$$\bar{\sigma}_S = |S|\ \bar{\sigma}_A\ , \tag{2.3.4}$$

which for the PolariX-TDS system is estimated to be $\bar{\sigma}_A \approx 1$ % [Reukauff, 2022]. When the TDS is operated at phases around the zero crossing $\phi_{RF} \neq 0$ the amplitude jitter translates into a pointing jitter:

$$\bar{\sigma}_{\langle y\rangle,A} = \frac{\phi_{RF}}{k_{RF}}\ \bar{\sigma}_S = \frac{\phi_{RF}}{k_{RF}}|S|\ \bar{\sigma}_A. \tag{2.3.5}$$

Assuming an amplitude jitter of $\bar{\sigma}_A \approx 1$ %, however, a comparison between this equation and the effective RF-phase jitter (cf. Eq. 2.3.3) reveals that for small RF phases $\phi_{RF} \ll 1$ the impact on the pointing is largely dominated by the phase jitter and the contribution from the RF-amplitude jitter becomes negligible.

Energy jitter. In a dispersive section operating in the x–s plane, mean-energy variations of the bunch translate into a horizontal centroid jitter:

$$\bar{\sigma}_{\langle x\rangle,\delta} = |D_x|\ \bar{\sigma}_\delta \tag{2.3.6}$$

The estimated linac-energy jitter at FLASH is $\bar{\sigma}_\delta \approx 0.1$ %, which, for a typical dispersion at the 8FLFDUMP screen of $|D_x| = 0.7$ m, represents a pointing jitter of $\bar{\sigma}_{\langle x\rangle,\delta} = 0.7$ mm. On the other hand, for comparatively large energy variations of $\bar{\sigma}_{\delta,\mathrm{PWFA}} \approx 5$ % expected from an unoptimised beam-plasma interaction in PWFA experiments, a pointing jitter of $\bar{\sigma}_{\langle x\rangle,\delta} = 35$ mm would be observed at the screen.

Linac-pointing jitter. The pointing jitter at FLASH is estimated to be $\bar{\sigma}_r \approx 50$ µm [Schreiber and Faatz, 2015]. However, as discussed in Section 1.2.2, its phase space progressively resembles that of the beam envelope as the bunch travels along the beamline and its actual value is expected to deviate from its nominal value depending on the lattice optic [Lindstrøm, D'Arcy et al., 2020]. At locations where the bunch is strongly focused to $\beta \approx 10$ mm, for instance, the jitter can be as low as $\bar{\sigma}_r \approx 1$ µm (cf. Figure 4.7), whereas for

moderately large beta functions of $\beta \approx 100$ m it can increase up to around $\bar{\sigma}_r \approx 150$ µm. For small shear-parameter values around $S \approx 10$, a pointing jitter of $\bar{\sigma}_r \approx 50$ µm is equivalent to 0.07 deg, and for $S \approx 40$ it goes down to below < 0.02 deg. In the dispersive plane, assuming a dispersion of $|D_x| = 0.7$ m, a pointing jitter of $\bar{\sigma}_r \approx 50$ µm would be equivalent to an energy jitter of less than < 0.01 % and generally its effect can be neglected.

Tolerance-calculation method

The amount of jitter that can be tolerated in a calibration routine depends on the FOV of the screen station, the beam parameters and either the operating conditions of the PolariX-TDS or the dispersion introduced by the dipole spectrometer. The approach followed here is to calculate the effective FOV available to perform a scan after substracting the projected (or dispersed) beamsize σ and the effect of the jitter $\bar{\sigma}$:

$$\mathrm{FOV}_{\mathrm{eff}} = \mathrm{FOV} - 6\sigma - 2\bar{\sigma}. \tag{2.3.7}$$

The effective field of view can be expressed in terms of the available RF phase $\Delta\phi_{RF}$ and the available energy range $\Delta\delta$ for the streaking plane and the dispersive plane, respectively:

$$\Delta\phi_{RF} = \mathrm{FOV}\frac{k_{RF}}{S} - 6\, k_{RF}\, \sigma_\xi - 2\, \bar{\sigma}_{\phi,\mathrm{eff}} \tag{2.3.8a}$$

$$\Delta\delta = \mathrm{FOV}\frac{1}{D_x} - 6\, \sigma_\delta - 2\, \bar{\sigma}_\delta, \tag{2.3.8b}$$

where in the first equation σ_ξ is the bunch length and in the second equation σ_δ and $\bar{\sigma}_\delta$ are the energy spread and the energy jitter of the bunch, respectively. Once the available scan range is obtained, N measurements are simulated by varying the scan parameter in 10 equally spaced steps and generating N_s normally distributed bunch centroids according to the magnitude of the jitter. The centroids that fall outside of the FOV are discarded. The figures of merit used to characterise the jitter tolerance are: 1) the deviation between the calibrated parameter and the real value used to simulate the measurement; 2) the relative error of the calibrated parameter, i.e., the error of the (simulated) measurement divided by the absolute value of the calibrated parameter—e.g., $\sigma_S/|S|$.

Shear-parameter calibration

As discussed above, when performing bunch-length measurements, the effective RF-phase variations $\bar{\sigma}_{\phi,\mathrm{eff}}$ are the principal source of jitter affecting the pointing stability at the screen. Therefore, only this source—with an estimated value of $\bar{\sigma}_{\phi,\mathrm{eff}} = 0.25$ deg—will be taken into account in the following. Regarding the bunch parameters, three different lengths are

considered: $\sigma_\xi = 3$ µm (10 fs), $\sigma_\xi = 30$ µm (100 fs), and $\sigma_\xi = 90$ µm (300 fs). The jitter tolerance is analysed for different shear parameters between $S = 10$ and $S = 75$, which, for the typical TDS-measurement optics settings (cf. Section 2.3.1) and a 1-GeV-beam with a normalised emittance of $\epsilon_{y,n} = 1$ µm, result in longitudinal resolutions of $R_\xi \approx 3$ µm (10 fs) and $R_\xi \approx 0.45$ µm (1.5 fs), respectively. For each shear parameter, $N = 10^4$ calibration measurements with $N_s = 10$ samples per scan step are simulated.

The results for a screen-optic system with a FOV ratio of 1:1—corresponding to a vertical extent of 12 mm—are shown in Figure 2.12. The left axis indicates the available scan range in degree and the right the relative error of the calibrated shear parameter. In all cases, the calculated shear parameter coincides with that used to simulate the measurement to within 5 % and it is therefore not shown in the plot. As expected, the available scan range is inversely proportional to the shear parameter $\Delta\phi_{RF} \propto S^{-1}$ and decreases progressively for longer bunches, whereas the calibration error increases as the scan range is reduced—i.e., as the jitter starts to dominate over the bunch length. For the longest bunch ($\sigma_\xi = 90$ µm), only shear parameters below $S < 20$ are operational. However, even in that case, longitudinal resolutions of $R_\xi \lesssim 3$ µm (10 fs) are achievable, which are still excellent for such long bunches. For the shortest bunch $\sigma_\xi = 3$ µm, an effective RF-phase jitter of $\bar{\sigma}_\phi = 0.25$ deg results in a tolerable calibration error of ≈ 8 % at the highest shear parameter achievable in the future with the RF compressor. A FOV ratio of 1:1 also provides enough optical-resolution power to measure the unstreaked beam size, which for a beta function at the screen of $\beta_y = 2$ m and a 1-GeV-bunch with $\epsilon_{y,n} = 1$ µm, would result in $\sigma_y \approx 30$ µm, equivalent to 6 pixels. A screen-optic system with a FOV ratio of 1:2, in contrast, would tolerate higher levels of jitter at the cost of reducing the optical resolution.

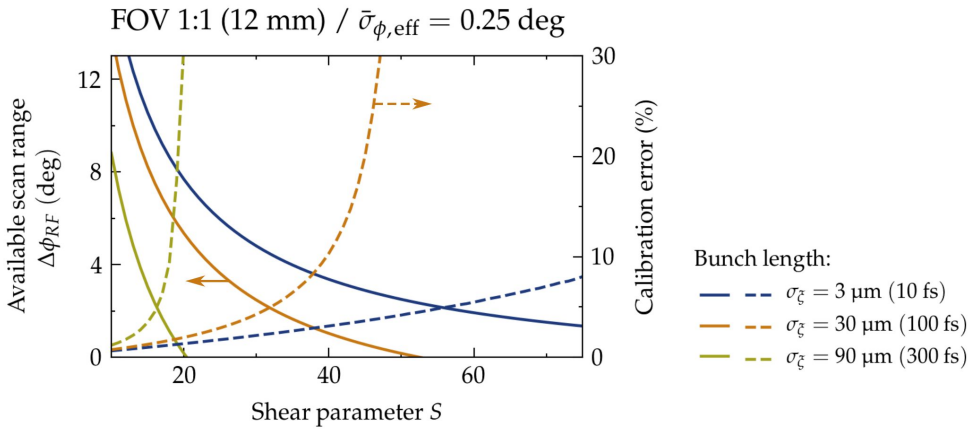

Figure 2.12.: Tolerance of transverse beam centroid jitter associated to an effective RF-phase jitter of $\bar{\sigma}_{\phi,\mathrm{eff}} = 0.25$ deg for a screen station with a FOV ratio of 1:1. The available scan range (left axis) is reduced as the shear parameter increases, which translates into a dramatic increase of the measurement errors (right axis) as the scan range approaches zero. The constraints observed for longer bunches can be avoided by operating the system with a reduced shear parameter while still producing a sufficient time resolution (see text).

Energy calibration

The most critical scenario to be considered at the screen 8FLFDUMP is that of an accelerated trailing bunch with a high mean-energy jitter $\bar{\sigma}_\delta$. The dispersion at that screen can vary slightly depending on the strength of the quadrupole Q4FLFDUMP located between the dipole and the screen, but its value in normal operating conditions is $D_x \approx 0.7$ m. Three different energy spreads are considered in the following: $\sigma_\delta = 0.05$ % and $\sigma_\delta = 0.1$ %, corresponding roughly to optimally and suboptimally beam-loaded trailing bunches, repectively, and $\sigma_\delta = 0.25$ %, corresponding to a full (unscraped) bunch with a negative chirp or to a strongly overloaded/underloaded trailing bunch. The tolerance is analysed for a range of energy-jitter values between $\bar{\sigma}_\delta = 0.1$ %—i.e., the nominal energy jitter of the FLASH linac—and $\bar{\sigma}_\delta = 1$ %. For each energy jitter, $N = 10^4$ calibration measurements with 10 scan steps and $N_s = 20$ samples per scan step are simulated.

The results for a screen-optic system with a FOV ratio of 1:2—corresponding to a horizontal extent of 18 mm—are shown in Figure 2.13. As expected, the available scan range (top plot) decreases linearly with increasing energy-jitter values $\Delta\delta \propto -\bar{\sigma}_\delta$ and increasing energy spreads $\Delta\delta \propto -\sigma_\delta$. The reference grey-dashed line in the plot indicates where the

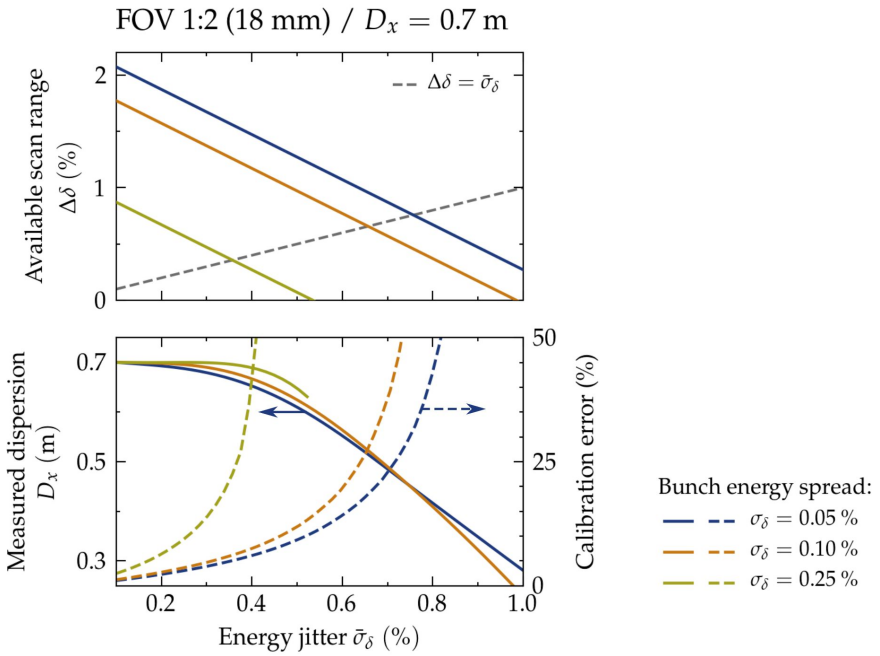

Figure 2.13.: Tolerance of transverse beam centroid jitter associated to the range of mean energy jitters expected at FLASHForward for a screen station with a FOV ratio of 1:2. The available scan range (top plot) is reduced for higher jitter amplitudes. The calibration error (right axis in the bottom plot) increases dramatically as the available scan range equals the energy spread of the bunch. The measured dispersion (left axis in the bottom plot) strongly underestimates the actual lattice dispersion above a mean energy jitter of about $\bar{\sigma}_\delta \approx 0.4$ %, which corresponds to the maximum tolerable amount of jitter.

available range equals the energy jitter $\Delta\delta = \bar{\sigma}_\delta$. The bottom plot shows the dispersion value obtained in the calibration (left axis) and the relative error of the measurement (right axis). The nominal energy jitter of FLASH—which is expected to apply in the absence of beam-plasma interaction—is tolerable in all energy-spread scenarios. However, as the energy jitter increases and approaches the value of the available scan range, the relative error increases exponentially to values above > 25 %, which occurs at progressively higher energy jitters for lower bunch energy spreads. On the other hand, the measured dispersion is gradually underestimated and reaches values below $D_x < 0.6$ m (≈ 15 % relative error) for energy-jitter values above $\bar{\sigma}_\delta \approx 0.5$ %. Taking into account these two observations, the level of jitter tolerable with a FOV ratio of 1:2 would be found within a range of $0.3\ \% < \bar{\sigma}_\delta < 0.5\ \%$, depending on the energy spread of the bunch. A FOV ratio of 1:1, in contrast, would severely compromise the operation of the spectrometer.

2.3.3. Beamline commissioning

In order to achieve an accurate level of control over the beam transport, a precise knowledge of the beamline and the operation of its basic components is required. A good level of agreement between the theoretical and the experimental models of the magnetic lattice, for instance, is of special importance when performing emittance measurements, in which the transport matrices determine the outcome of the measurement. In the context of this thesis, the characterisation of the beam emittance plays a central role, and therefore, special care was taken to validate the proper operation of the beamline in a dedicated beam-based commissioning campaign. The goal of this section is not to provide a detailed analysis of the measurements performed but to shortly describe the steps followed and to report on the major outcomes of the campaign. Prior to the beam-based commissioning, the polarity of all the magnets was checked in situ with a Hall sensor.

BPM calibration. The most basic assessment of the beam transport consists of characterising the beam orbit along the beamline, enabled by the BPMs. Therefore, the first step of the commissioning campaign was to calibrate all the BPMs of the PolariX-TDS beamline. This was performed by introducing a known kick with a steerer magnet located upstream of the BPM generating a sufficient phase advance as to clearly observe the resulting transverse offsets at the monitor. The quadrupole magnets between each chosen pair of steerer magnet and BPM must be degaussed and switched off, thus ensuring that there is only a drift space between them. Since the accuracy of the BPM readings is reduced for large transverse beam sizes, an appropriate lattice optic producing beta functions below $\beta_x = \beta_y \approx 20$ m along the whole beamline was used. For the same reason, a beam with a small energy spread was produced at the linac in order to obtain a small transverse beam size at the dispersive section (FLFDUMP). Apart from obtaining an accurate calibration of the BPMs, the measurement enabled the validation of both the location and the type of each of the steerer magnets installed, which is a mandatory requirement to perform the orbit-response measurements

discussed next.

Orbit-response measurements (ORM). Similar to the procedure followed for the BPM calibration, an ORM is performed by introducing a range of known angular kicks with a steerer magnet and measuring their transport along the beamline. In this case, however, all the quadrupoles are switched on. The resulting dependence of the transverse beam centroid offset measured at the BPMs is used to determine the transport-matrix elements R_{12} (for the horizontal kick) and R_{34} (for the vertical kick), which are subsequently compared to those of the theoretical model. As a result of the measurement, a small error was detected in the calibration of the power supply of the first quadrupole upstream of the PolariX-TDS, which was corrected by the MKK[4] technical group at DESY. Apart from that, a good agreement between the theoretical model and the experiment was observed, thus validating the knowledge about the positions of the magnets, their magnetic fields, and their current calibrations.

TDS-screen-based ORM. For the straight lattice section downstream of the PolariX-TDS a method similar to the ORM was employed, which used the TDS as a steering element and the screen 11FLFXTDS as a monitor of the transverse centroid offset. This setup has the advantage of conferring a higher level of precision to the measurement due to the larger range of offsets measurable at the screen (9 mm $\times$ 12 mm) compared to that typically available within the linear range of a BPM ($\approx$ 1 mm). The measurement enabled the identification of a large error in the focusing strength of one of the quadrupoles, which was related to an underestimation of the current delivered by the power supply by more than 30 %.

2.4. Differential-pumping system

As mentioned in Section 2.2, the continuous-flow design of the gas cell used at FLASH-Forward requires a vacuum-system configuration that allows direct connection of the high pressure region around the plasma chamber to the ultra-high-vacuum (UHV) conditions mandatory for the safe operation of the FLASH facility—i.e., $p \lesssim 10^{-8}$ mbar at the transmission beamline and $p \lesssim 10^{-9}$ mbar at the linac or special diagnostics equipment as, for instance, the PolariX-TDS. A careful design of the vacuum system around the TDS was therefore of the utmost importance to the safe operation of the system.

A standard approach to efficiently decouple two vacuum regions with a large pressure difference is to connect them by means of one or several differential-pumping stages, which consist of an appropriate arrangement of vacuum pumps connected by small apertures that greatly reduce the amount of gas flowing towards the direction of lower pressures. While small apertures increase the efficiency of the system, their dimensions are constrained by the

[4]Maschine Kraft Kühlung Klima.

operational requirements of the beamline—e.g., the lattice optic and the transverse size of the electron beam. Regarding the pumping speed needed at each stage, the gas load to be evacuated is determined by the pressure obtained at the previous stage and the conductance of the pipe connecting the two stages. Once UHV conditions are achieved, the system is limited by the conductance of the pipes and an increase of pumping speed does not translate into an improvement of the vacuum levels along the beamline [Kersevan, 2007].

At FLASHForward, the gas cell is operated with a gas load of up to $Q_{\text{in}} = 20\ \text{mbar} \cdot \text{l/s}$ [Dale and Schaper, 2017]. The principal requirement of the vacuum system is that UHV conditions must be achieved at two locations: 1) just upstream of the FLFMAFF section; 2) at the PolariX-TDS. A large amount of the gas load is already evacuated with two pumps attached to the plasma chamber and the remaining part is progressively removed by three differential-pumping stages installed at each side of the chamber—i.e., six stages in total. The minimum diameter of the apertures allowed is mainly constrained by the 25 TW Ti:Sapphire laser used for internal-injection (X-1) experiments, which propagates colinear to the electron beam and requires progressively larger apertures with diameters of up to ≈ 40 mm to avoid clipping the transverse profile far away from the waist. Immediately upstream and downstream of the plasma chamber, small apertures with diameters of 10 mm are allowed, since both the electron and the laser beams are close to their focus and have transverse sizes on the order of $\mathcal{O}(10\ \mu\text{m})$.

The dimensioning of the system for the plasma chamber and the three upstream differential-pumping stages was conducted by an external company [Hilbert, 2017]. Following a similar approach, the dimensioning of the three downstream differential-pumping stages and the sections immediately upstream and downstream of the PolariX-TDS was performed at FLASHForward. The vacuum design of the PolariX-TDS system itself—which involves a high level of complexity—was carried out by the MVS[5] technical group at DESY. Taking into account the constraints described above, analytical formulae were used to characterise the system and determine the best configuration of beam-pipe diameters and pumping speeds. Subsequently, the whole system was modeled with the simulation program Molflow+ [Kersevan and Ady, 2019], which meant the presence of beaming effects that reduce the efficiency of the system could be discarded [Kersevan, 2007]. All calculations and simulations were performed considering operation with H_2, which is the gas species most difficult to evacute. A thermal outgassing rate of $q \approx 10^{-10}\ \text{mbar} \cdot \text{l/s/cm}^2$ was assumed for all pipes. Figure 2.14 shows the pressure profile obtained along the whole beamline and the location and type of the different vacuum pumps used. The pressure at the plasma chamber is kept at $p = 0.1$ mbar by means of two turbo-molecular pumps (TMP) with an effective pumping speed of $S \approx 200\ \text{l/s}$. Each differential-pumping stage is equipped with a TMP with $S \approx 500\ \text{l/s}$ and achieves a pressure reduction between two and three orders of magnitude. Two ion-getter pumps (IGP) with $S \approx 70\ \text{l/s}$ are attached at each side of the PolariX-TDS with a high-conductance (large diameter) pipe to facilitate the evacuation of residual gas coming from the structure. The conductance-limited section downstream of the TDS is kept

[5]Maschinen Vakuum Systeme

at UHV conditions by two additional IGPs.

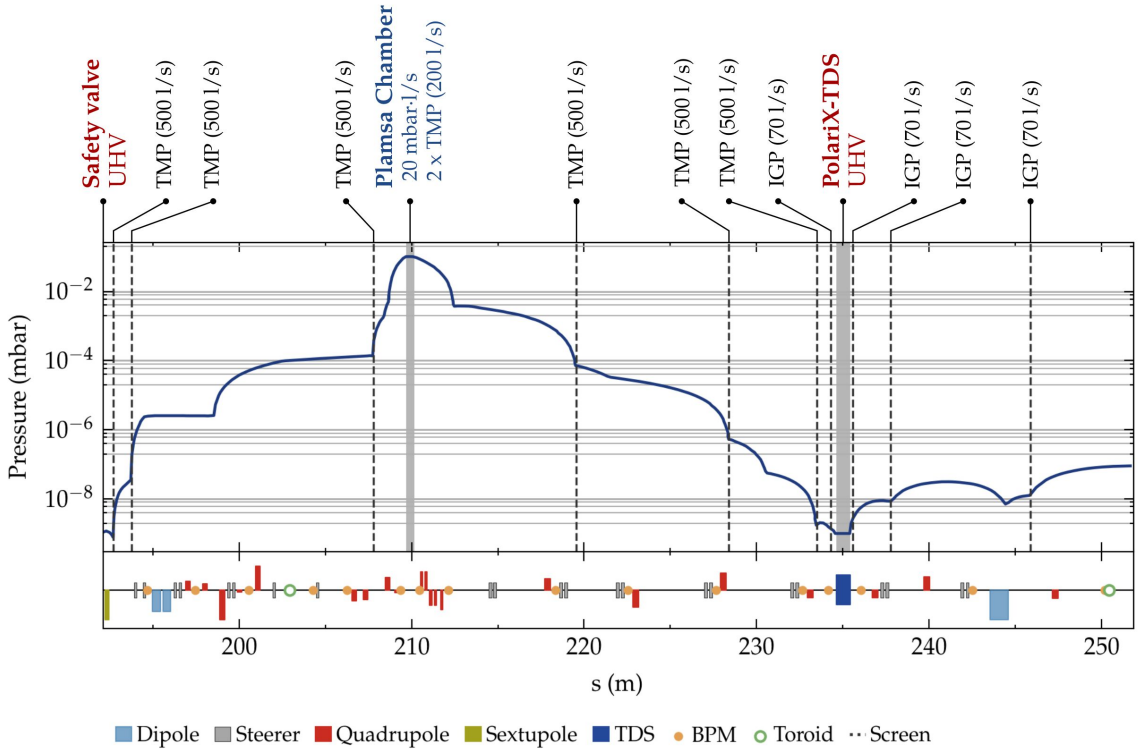

Figure 2.14.: Vacuum-pressure profile achieved at FLASHForward with its dedicated differential-pumping system. The gas cell is operated with a gas load of up to $Q_{\text{in}} = 20$ mbar $\cdot$ l/s, which must be completely absorbed shortly upstream of the FLFMAFF section (Safety valve) and at the location of the PolariX-TDS. The pressure at the plasma chamber is kept at $p = 0.1$ mbar by means of two turbo-molecular pumps (TMP) with an effective pumping speed of $S \approx 200$ l/s. Each differential-pumping stage is equipped with a TMP with $S \approx 500$ l/s and achieves a pressure reduction between two and three orders of magnitude. The section around and downstream of the PolariX-TDS is kept at UHV conditions by four IGPs with $S \approx 70$ l/s.

3. PolariX-TDS commissioning

In this chapter, the RF design of the PolariX-TDS system, the hardware commissioning of some of its critical RF-components, and its beam-based commissioning are described, including several applications for beam characterisation and optimisation. Before that, the circumstances that enabled the development and production of the new device—which led to the creation of the PolariX-TDS collaboration in 2017—are briefly described.

PolariX-TDS Collaboration. In 2016, in the framework of a study for a future compact linear collider (CLIC) based on high-gradient accelerating RF structures [Aicheler et al., 2012], Alexej Grudiev published a report with the design of compact high-power RF components with frequencies in the X-band range [Grudiev, 2016]. Some of these components, which will be described in the following sections, are at the core of the novel TDS design. In parallel, the Paul-Scherrer-Institut (PSI, Switzerland) developed the so-called *high-precision tuning-free assembly procedure* [Ellenberger et al., 2013; Zennaro et al., 2014], which was required to produce the C-band accelerating structures for the SwissFEL [Prat et al., 2020] in order to avoid any RF-tuning after fabrication. These procedures are the second critical element needed to guarantee the proper operation of the new device. Finally, there was a demand for new longitudinal beam diagnostics coming from several research facilities: FLASHForward, FLASH2 [Christie, Roensch-Schulenburg and Vogt, 2019] and the ARES-SINBAD facility [Dorda et al., 2016] at DESY, and the ATHOS beamline [Ganter et al., 2019] at SwissFEL at the PSI. All these facilities require either femtosecond or sub-femtosecond resolution and can largely benefit from the extended capabilities enabled by a TDS with variable polarisation. Altogether this is the background that motivated the creation of an international collaboration between CERN, PSI and DESY in 2017 [Marchetti, Assmann et al., 2017], with the aim of developing, implementing, and commission the new device—which later on was given the name *PolariX-TDS*. The first prototype produced was installed at FLASHForward in 2019 and its successful commissioning, discussed in this section, represents an important milestone of the PolariX-TDS collaboration.

3.1. RF design

The conceptual design of the PolariX-TDS system is schematically depicted in Figure 3.1. Initially proposed by Alexej Grudiev in Ref. [Grudiev, 2016], it was further developed by him in Ref. [Grudiev, 2017]. The operation foreseen by Grudiev works as follows: an RF

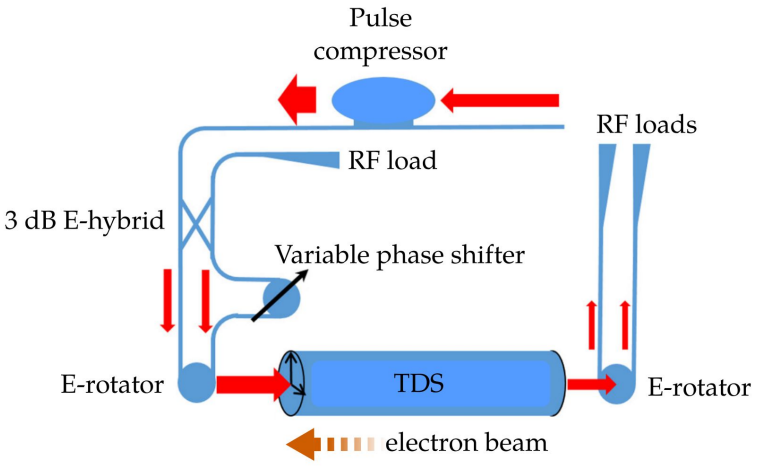

Figure 3.1.: Conceptual RF design of the PolariX-TDS system (figure reproduced from [Craievich, Bopp et al., 2020]).

pulse of about ~ 1 µs is first compressed[1] to achieve high-power signal over a shorter time length and is afterwards split into two signals by a 3 dB splitter. One of the two split arms contains a variable phase shifter that can modify the relative phase between the two signals, which are then combined by an E-rotator to produce a backward-travelling wave with a linearly polarised dipole mode (cf. Section 1.3.1). At the output of the structure, the two signals are disentagled by a second E-rotator and are directed into two separate RF loads. In the following, key aspects of the geometry and operation of the aforementioned components are summarised. After that, the complete RF hardware system is described, including the RF high-power source and the low-level RF (LLRF) system.

3.1.1. RF components

3 dB E-hybrid. All RF components described in [Grudiev, 2016] use the 3 dB E-hybrid as a point of departure. The device is shown in Figure 3.2 a). Its function is to split a single input-power signal into two output-power signals, each having half the amplitude of the input—i.e., each attenuated by 3 dB. In this case, in contrast to other 3 dB hybrid geometries, both output signals flow parallel to each other in the forward direction and are 90 deg out of phase. To achieve this, four ports are required, one of them terminated with an RF load. The device is called an E-hybrid because the RF-power splitting takes place in the plane tangential to the electric field.

E-rotator. The E-rotator is obtained by short-circuiting the two output ports of a 3 dB E-hybrid and attaching a circular waveguide at its center as shown in Figure 3.2 b). In this configuration, the resulting device has two ports: one input corresponding to one of the

[1]The compression is actually optional and might not be required if the incoming RF pulse already has enough power.

a) 3-dB E-hybrid

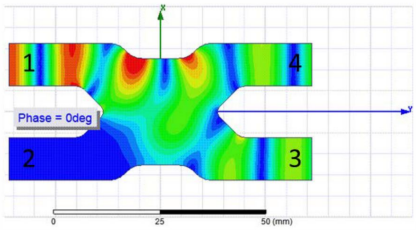

b) circular TE_{11} rotating mode launcher (E-rotator)

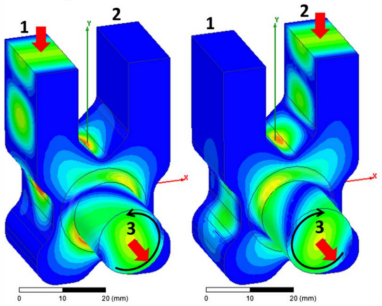

c) variable polarisation circular waveguide TE_{11} mode launcher

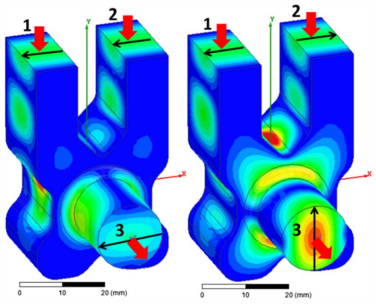

d) variable phase shifter

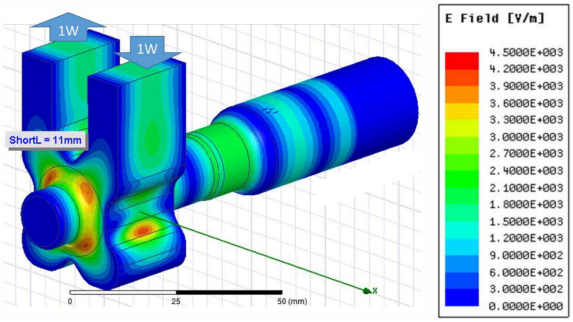

Figure 3.2.: Compact high-power RF components at the X-band used in the PolariX-TDS system (figures reproduced from [Grudiev, 2016]).

E-hybrid ports and one output corresponding to the circular waveguide—the remaining E-hybrid port being terminated with an RF load. At the circular output, two degenerate TE_{11} modes with the same amplitude but with 90 deg phase difference can propagate, thus producing a circularly polarised TE_{11} mode. Depending on which input is used to feed the power, the launched mode at the output rotates either counterclockwise $\boldsymbol{\mathcal{E}}_L$ or clockwise $\boldsymbol{\mathcal{E}}_R$. In cartesian coordinates this can be expressed as:

$$\boldsymbol{\mathcal{E}}_L = \mathcal{E}_L \cos(\omega_{RF} t + \phi_{L,0})\hat{\boldsymbol{x}} \; + \; \mathcal{E}_L \sin(\omega_{RF} t + \phi_{L,0})\hat{\boldsymbol{y}} \tag{3.1.1a}$$

$$\boldsymbol{\mathcal{E}}_R = \mathcal{E}_R \cos(\omega_{RF} t + \phi_{R,0})\hat{\boldsymbol{x}} \; - \; \mathcal{E}_R \sin(\omega_{RF} t + \phi_{R,0})\hat{\boldsymbol{y}}. \tag{3.1.1b}$$

If instead, two signals with the same power $P_L = P_R$ are fed simultaneously into the two E-hybrid ports, two rotating modes (with $\mathcal{E}_L = \mathcal{E}_R = \mathcal{E}_0$) are launched at the output and their overlap results in a linearly polarised TE_{11} dipole mode:

$$\boldsymbol{\mathcal{E}}_\perp = \boldsymbol{\mathcal{E}}_L + \boldsymbol{\mathcal{E}}_R = 2\mathcal{E}_0 \cos(\omega_{RF} t)\hat{\boldsymbol{x}}\,, \tag{3.1.2}$$

where the initial conditions $\phi_{L,0} = \phi_{R,0} = 0$ have been assumed for simplicity. As shown in Figure 3.2 c), depending on the relative phase between the two inputs, the direction of the

dipole mode can be selected: if they are in phase ($\phi_{L,0} = \phi_{R,0}$) the polarisation is horizontal, and if they are 180 deg out of phase ($\phi_{L,0} = \phi_{R,0} + 180$) it is vertical. The phase modulation is achieved with a variable phase shifter, which is described in the following.

Variable phase shifter. By feeding power into one of the inputs of the E-rotator and short-circuiting the circular waveguide, the launched circularly polarised wave is reflected back into the device. Since the rotation direction is inverted after reflection, the signal propagates to the second input port, which in this configuration acts as an output. Adjusting the effective length of the circular waveguide—i.e., the position of the short circuit, which in practice is substantiated by means of a movable piston—the phase of the outgoing signal can be selected.

TDS. The TDS itself is a constant-impedance circular disk-loaded waveguide structure operating in a TM_{110}-like mode (cf. Section 1.3.1) in the backward-travelling-wave regime—i.e., the power (group velocity) flows in the direction opposite to the beam propagation (phase velocity)—and is fabricated from brazed stacks of machined copper cylinders. The main difference with respect to a typical TDS design, is that the cells need to be azymuthally symmetric in order to guarantee the degeneracy of the two TM_{110}-like modes, which must coexist to produce a circularly polarised wave. The design operating frequency of the structure—which is directly related to the radius b of the cell—is $f_0 = 11.9952$ GHz. To adapt to the requirements of each facility, however, the frequency needs be fine tuned by adjusting the operating temperature T according to the equation:

$$f = f_0 - \alpha_{\text{Cu}}\, f_0\, (T - T_0), \tag{3.1.3}$$

where $\alpha_{\text{Cu}} = 1.8 \cdot 10^{-5}\ \text{K}^{-1}$ is the copper expansion coefficient and $T_0 = 30\ ^\circ\text{C}$ is the design operating temperature. The remaining parameters of the cell geometry and the number of cells are determined to optimise the performance of the structure [Craievich, Bopp et al., 2020], which can be caracterised by a figure of merit called the *transverse shunt impedance* $R_\perp$, representing the conversion efficiency from power to particle deflection:

$$R_\perp = \frac{V_0^2}{P_{\text{in}}}, \tag{3.1.4}$$

where $V_0 = \mathcal{E}_0 L$ (cf. Eq. 1.3.11) is the integrated deflecting voltage delivered to the beam and P_{in} is the power fed into the structure. An equivalent parameter widely used in the RF literature is the *power-to-voltage* constant D:

$$D = \frac{V_0}{\sqrt{P_{\text{in}}}}. \tag{3.1.5}$$

Other aspects of the TDS performance which had to be taken into account in the design are, among others, the maximum field gradients sustainable at the inner surface of the structure before electrical breakdown occurs, the aperture available for electron- and photon-beam transmission, and the HOM[2] longitudinal and transverse wake potentials excited by the electron beam, which might affect the beam quality. In order to couple the RF power into the structure without reflections, an additional matching cell with a slightly different geometry is introduced between the E-rotator and the TDS. Two different structure designs with different total lengths are devised to allow for more flexibility in case of eventual spatial constraints. Table 3.1 summarises key parameters of the two structures operating with and without RF-pulse compression. The short version was fabricated as a first prototype, which is the one installed at FLASHForward used in all measurements presented in this thesis.

Cell Parameter	Symbol	Units	Value	
Frequency†	f_0	GHz	11.9952 (11.9888)	
Temperature†	T_0	°C	30 (62)	
Iris radius	a	mm	4	
Cell-to-cell phase advance	$\delta\phi_0$	deg	120	
Shunt impedance	$R'_\perp$	MΩ	50	
TDS parameter			Short	Long
Number of cells			96	120
Filling time	t_f	ns	104.5	129.5
Attenuation	τ_0	dB	−5.21	−6.48
Active length	L	mm	800	1000
Total length‡	L_{tot}	mm	960	1160
TDS alone				
Shunt impedance	$R_\perp$	MΩ	27.3	37.5
Power-to-voltage	D	$\mathrm{MV/MW^{0.5}}$	5.2	6.1
TDS + XBOC				
Shunt impedance	$R_\perp$	MΩ	142	178
Power-to-voltage	D	$\mathrm{MV/MW^{0.5}}$	11.9	13.3

Table 3.1.: Summary of key design RF parameters of the structure of the PolariX-TDS system. Data reproduced from [Craievich, Bopp et al., 2020]. †The quantities in parenthesis correspond to the operating frequency and temperature at FLASHForward. ‡The total length includes the E-rotator and the coupling cell at each side of the structure.

[2]Higher order modes.

3.1.2. Complete RF-hardware system

The complete RF-hardware system installed at FLASHForward in summer 2019 is schematically depicted in Figure 3.3. In contrast to the conceptual design shown in Figure 3.1, this one includes the high-power RF source and the LLRF system, whereas the RF-pulse compressor is excluded, since it was not installed until early 2021. All components at the high RF-power side are kept under ultra-high vacuum (UHV) conditions at pressures well below $p < 10^{-8}$ mbar and are temperature stabilised by means of a water-cooling system.

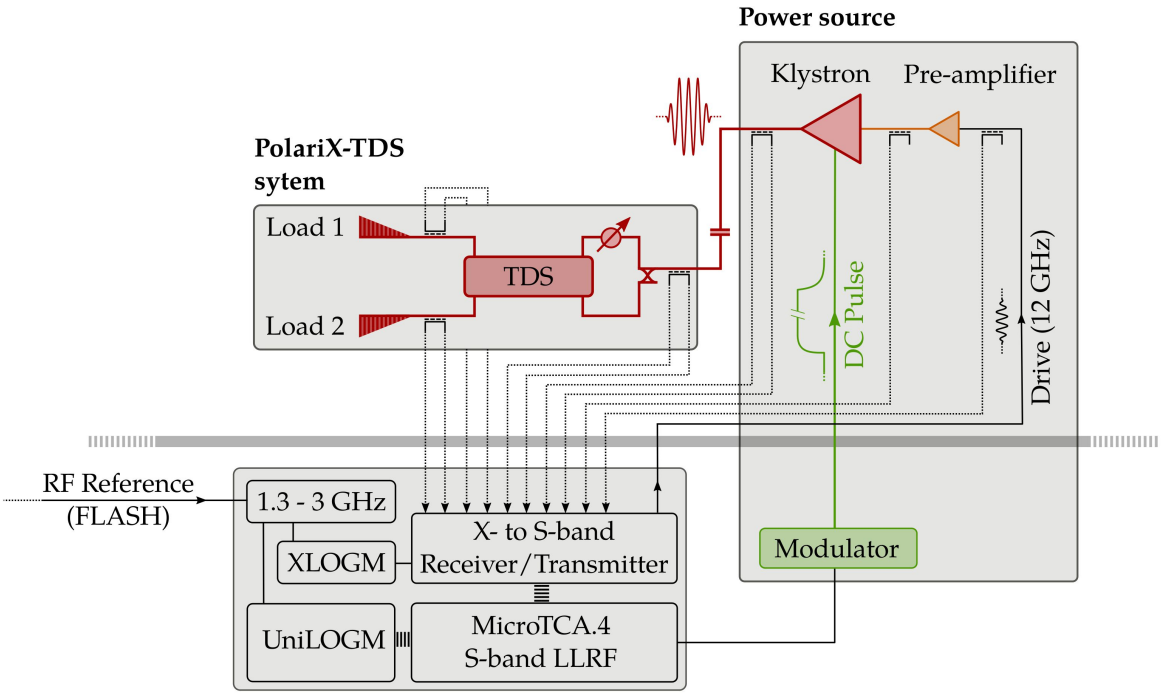

Figure 3.3.: Schematic representation of the complete PolariX-TDS RF-hardware system.

Power source. The klystron (6 MW Toshiba E37113A) operates as an amplifier that converts some of the DC energy input carried by a high-voltage ($\sim$ 160 kV), high-current ($\sim$ 100 A) pulse delivered by a modulator (Ampegon Type-μ M-Class) into RF energy with a frequency of 12 GHz given by a pre-amplified drive signal delivered by the LLRF system. Thus, the drive signal is amplified to a high power of up to 6 MW and is subsequently directed from the klystron output to the TDS through a waveguide network. For reasons of machine safety, an RF window decouples the UHV of the klystron network from that of the structure and the components attached directly to it, thus preventing an accidental venting of the machine vacuum through the high power RF system.

LLRF system. Several directional waveguide couplers are used to pick up a small fraction (≈ -60 dB) of the power at different locations accross the system. At low power—before and after the pre-amplifier—only the forwarded signal is measured, whereas at high

power the reflected signal is also provided in order to: 1) monitor the actual power directed towards the structure; 2) detect an eventual malfunction of the system that could translate into an increase of the reflected signal and damage the klystron. Since the PolariX-TDS is the first X-band RF structure installed at DESY, a dedicated LLRF system has been developed by members of the MSK[3] group to extend the already existing systems operating at 1.3 GHz (FLASH, EuXFEL) and 3 GHz (SINBAD, REGAE) to the 12 GHz required by the new device [Reukauff, 2022]. The 12 GHz signals coming from the directional couplers are converted into the 3 GHz domain at which the MicroTCA.4 S-band LLRF module operates, enabling the integration of the hardware system into the DOOCS framework of FLASH (cf. Section 2.1).

3.2. RF-hardware commissioning at DESY

The references [Romano del Pozo et al., 2019; Craievich, Bopp et al., 2020] report on the manufacture, the low- and high-power tests, and the high-power pre-conditioning of individual components of the PolariX-TDS system performed at CERN and at the PSI. In this section, specific aspects of the hardware commissioning performed at DESY, which are relevant to the beam-based commissioning, are briefly summarised.

RF-component performance tests

Before installation in the experimental area, the performance of all individual waveguide components is inspected by the MIN[4] group. To that end, two important parameters are used:

- *reflection-attenuation coefficient* τ_{att}, which is defined as the ratio between the fraction of power P_{refl} reflected by an impedance discontinuity back to the source and the full power forwarded into it P_{forw} expressed in dB:

$$\tau_{\text{att}} = 10 \cdot \log_{10}\left(\frac{P_{\text{refl}}}{P_{\text{forw}}}\right) \tag{3.2.1}$$

- *insertion loss* α_I, which is defined as the loss of RF power in a transmission line:

$$\alpha_I = 10 \cdot \log_{10}\left(\frac{P_{\text{in}}}{P_{\text{out}}}\right), \tag{3.2.2}$$

  where the P_{in} is the input power and P_{out} is the output power.

[3] Maschine Strahlkontrolle
[4] Maschine: Injektoren und Injektion

In all cases, reflection-attenuation coefficients below −26 dB are obtained—i.e., less than 0.25 % of the forwarded power is reflected back to the source. Regarding the 3-dB E-hybrid, measurements performed with ideal loads connected to its outputs confirm that the input signal is effectively split into two equal parts. Special care is taken to verify the proper operation of the phase shifter and to establish safe limits for the insertion range of the piston to avoid eventual problems with its mechanical actuation. The results of the measurement (cf. Figure 3.9) demonstrate insertion losses below 0.1 dB and reflection-attenuation coefficients below −26 dB within the established actuator limits; in good agreement with the measurements performed at CERN [Romano del Pozo et al., 2019].

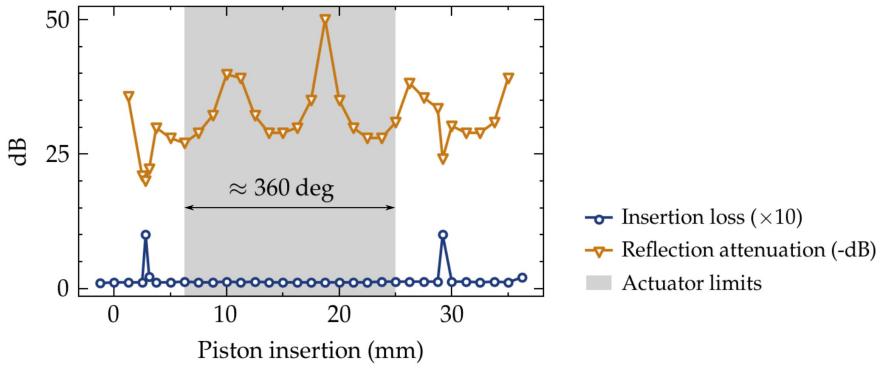

Figure 3.4.: Measurement of the insertion loss and reflection-attenuation coefficient of the phase shifter. Measurement performed by the MIN group.

High-RF-power conditioning

Once the complete hardware system is installed in the experimental area and the inner volume of the waveguide network is evacuated to fulfil UHV conditions, the high-RF-power conditioning of the system is performed to enable the TDS operation at the required field gradients. The conditioning procedure consists of progressively increasing the power fed into the system to promote the desorption of material and/or the field emission of electrons from the inner metallic surfaces in a controlled way, thus avoiding major discharge events that could irreversibly damage the RF components. This is effectively equivalent to removing inpurities adsorbed in the surface and polishing residual micro-protrusions/micro-tips.

RF-power calibration

The RF-power calibration of the low-level RF (LLRF) signals is performed by the MSK group according to the following procedure. A signal with a known power level coming from an external RF-signal generator is injected at the directional couplers distributed along the waveguide network. The signal sampled by the LLRF system is then used to determine a calibration factor that introduces an appropriate amplitude scaling to compensate for

the differences between the reference and the measured signals. This calibration includes all losses along the measurement chain—i.e., cable losses, conversion gain of the receiver, additional attenuators, etc. Afterwards the coupling factor of the directional waveguide coupler was included in the calibration.

Since the described procedure is performed at low power, the LLRF calibration is subsequently verified by operating the klystron at high power and comparing the measured (calibrated) signal with the nominal values provided in the technical specifications. The results of the measurement are shown in Figure 3.5, in which the klystron output is plotted versus the pre-amplified input signal. As can be observed, the maximum (nominal) achievable output power of the klystron is 6 MW, whereas the measured signal indicates values of almost 6.5 MW. The klystron power is overestimated by almost 8 %, while the drive power from the preamplifier is underestimated by $\sim$ 10 %. This mismatch is not fully understood, but could be related either to deviations from the nominal coupling factors of the directional couplers after installation in the experimental area or to nonlinearities in the dedicated X-to-S-band receiver. The calibration of the remaining directional couplers cannot be verified, since the actual power transported to locations downstream of the klystron is not precisely known. However, on account of the results discussed above, a default calibration error of 0.5 dB is assigned to all LLRF power readings.

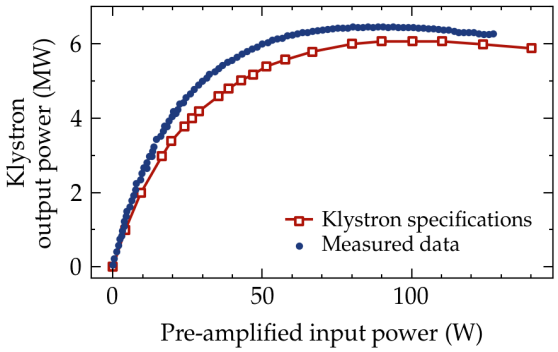

Figure 3.5.: Comparison between the calibrated LLRF signals and the klystron specifications. Measurement performed by the MSK group.

3.3. Beam-based commissioning

The main goal of the beam-based commissioning is to validate key RF parameters of the PolariX-TDS system, including: 1) the balance between the power fed into the two arms of the E-rotator; 2) the power-to-voltage constant of the structure; 3) the stability of the RF system. In the following, the characterisation of each of these parameters is described separately.

For the commissioning campaign, a bunch with a charge of $Q \approx 100$ pC, an energy of $\langle E \rangle = 750$ MeV, and a moderately short rms length of $\sigma_{\xi,rms} \approx 135$ fs (40 µm) is used. The

beamline-setup routines consist of: 1) successively minimising the dispersion at the end of the FLASH2 extraction and at the end of the FLASHForward extraction and compression sections; 2) measuring the projected emittance at the FLFMAFF section; 3) matching the beam to symmetric optics at the start of the FLFDIAG section, with $\beta_x = \beta_y = 10$ m and $\alpha_x = \alpha_y = 0$, compatible with slice-emittance measurements in either of the two transverse planes x–x' and y–y', which is required for the diagnostics applications discussed in Section 3.4. The lattice optic downstream of that point is adjusted to the requirements of each measurement separately.

3.3.1. Power imbalance

During the early operation tests with the PolariX-TDS system, an unexpected kick perpendicular to the streaking direction was observed. In this section, this behaviour—which is related to the principles of operation of the E-rotator—is analysed. After that, the experimental determination of its magnitude is presented.

As discussed in Section 3.1, the nominal operation of the E-rotator requires feeding two equal-power signals into each of the two input ports of the device $P_L = P_R$. If the last condition is not fulfilled, the amplitudes of the launched circular waves are not balanced $\mathcal{E}_L \neq \mathcal{E}_R$ and their overlap results in an elliptically polarised mode $\mathcal{E}_E$:

$$\begin{aligned} \boldsymbol{\mathcal{E}}_E &= (\mathcal{E}_L - \mathcal{E}_R)\cos(\omega_{RF}t)\hat{\boldsymbol{x}} \; + \; (\mathcal{E}_L + \mathcal{E}_R)\sin(\omega_{RF}t)\hat{\boldsymbol{y}} \\ &= \mathcal{E}_\perp \cos(\omega_{RF}t)\hat{\boldsymbol{x}} \; + \mathcal{E}_\parallel \sin(\omega_{RF}t)\hat{\boldsymbol{y}}, \end{aligned} \tag{3.3.1}$$

where it has been assumed that the left and right waves are 180 deg out of phase. The elliptical mode of Eq. 3.3.1 is the combination of two linearly polarised modes with different amplitudes $\mathcal{E}_\parallel \neq \mathcal{E}_\perp$ and 90 deg off phase respect to each other. In these conditions, if the TDS is operated at the zero crossing of the streaking mode $\mathcal{E}_\parallel$, the deflecting mode $\mathcal{E}_\perp$ is on-crest and introduces a net kick to the bunch perpendicular to the streaking direction.

In order to measure the amplitude of each of the two components of Eq. 3.3.1—and therefore the power imbalance between the two E-rotator inputs—a 360 deg RF-phase scan is performed. Since the TDS is required to operate over the full range of RF phases, the equations describing the particle dynamics differ from those presented in Section 1.3.1, which were derived under the assumption that the bunch is always located around the zero crossing $k_{RF}\xi \approx 0$. The expressions for the bunch centroids in y and y' at the TDS exit

compatible with any operating RF phase are:

$$y = y_0 + y'_0 L + \frac{e\mathcal{E}_{\parallel} L \sin\phi_{RF}}{p_0 c} \frac{L}{2} \tag{3.3.2a}$$

$$y' = y'_0 + \frac{e\mathcal{E}_{\parallel} L \sin\phi_{RF}}{p_0 c}, \tag{3.3.2b}$$

where L is the length of the structure, ϕ_{RF} is the RF phase relative to the zero crossing, and $p_0 c$ is the bunch energy. After transport to the screen, the following expression results:

$$y_s = \underbrace{R_{33} y_0 \; + \; (L R_{33} + R_{34})\, y'_0}_{y_{s,\beta}} \; + \; \left(R_{33} \frac{L}{2} + R_{34} \right) \frac{e\mathcal{E}_{\parallel} L \sin\phi_{RF}}{p_0 c} \tag{3.3.3a}$$

$$= \; y_{s,\beta} \; + \; \left(R_{33} \frac{L}{2} + R_{34} \right) \frac{e}{p_0 c} V_{\parallel} \sin\phi_{RF}, \tag{3.3.3b}$$

where the brace in the first equation gathers the terms corresponding to the natural betatron motion $y_{s,\beta}$ and the effective voltage $V_{\parallel} = \mathcal{E}_{\parallel} L$ has been introduced in the second equation. Equivalent equations are obtained for the x–x' plane by using the corresponding transport-matrix terms (R_{11} and R_{12}), and substituting $\mathcal{E}_{\parallel} \sin\phi_{RF}$ with $\mathcal{E}_{\perp} \cos\phi_{RF}$. Rearranging terms, the effective voltage corresponding to the streaking $V_{\parallel}$ and the deflecting $V_{\perp}$ modes for a given RF phase ϕ_{RF} can be expressed in terms of the bunch centroids x_s and y_s measured at the screen:

$$V_x = V_{\perp} \cos\phi_{RF} = \frac{p_0 c}{e} \cdot \frac{1}{R_{11} \frac{L}{2} + R_{12}} \left(x_s - x_{s,\beta} \right) \tag{3.3.4a}$$

$$V_y = V_{\parallel} \sin\phi_{RF} = \frac{p_0 c}{e} \cdot \frac{1}{R_{33} \frac{L}{2} + R_{34}} \left(y_s - y_{s,\beta} \right). \tag{3.3.4b}$$

To perform the measurement, the TDS is operated with a vertically polarised streak—i.e., 180 deg phase difference between the left and right E-rotator inputs—whereas the beam parameters at the TDS are adjusted as if the streak were in the horizontal plane, with $\beta_{x,TDS} \approx 50$ m and $\beta_{y,TDS} \approx 10$ m. The lattice optic between the TDS and the 11FLFXTDS screen is optimised accordingly to amplify the "unwanted" kick with $R_{12} \approx 11$ and $R_{34} \approx 2.5$ m. The klystron power is subsequently adjusted to be able to monitor the bunch centroid at the screen over the full range of RF phases, with an obtained value of $P_k \approx 200$ kW. The results are shown in Figure 3.6. On the left, the elliptical path described by the bunch centroids at the screen is plotted. Since the polarisation of the streaking term $V_{\parallel}$ slightly deviates from the vertical axis—with an angle of $\theta = 1.53$ rad (87.7 deg)—the voltages V_x

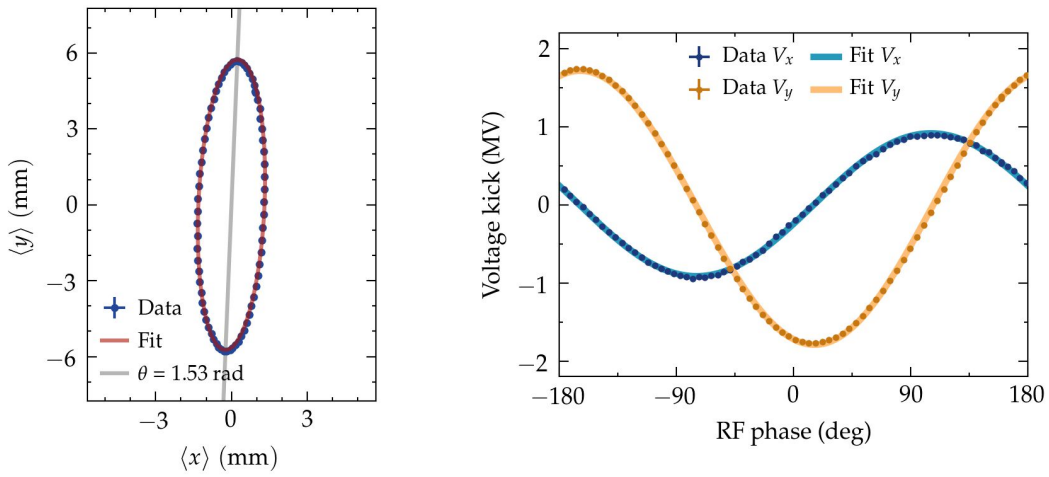

Figure 3.6.: 360 deg RF-phase scan performed to calculate the power imbalance between the left and right input ports of the E-rotator.

and V_y (cf. Eq. 3.3.4) are calculated from the projection of the bunch centroids onto each of the normal axes of the ellipse. On the right, the values obtained are plotted versus the RF phase ϕ_{RF}, together with sinusoidal functions fitted to the data. Notice that in the plot the amplitude of the deflecting term $V_\perp$ is multiplied by a factor 10 to better appreciate the 90 deg offset between the two signals.

From Eq. 3.3.1, the effective voltage amplitude coming from the left and right E-rotator inputs can be calculated from the streaking and deflecting terms obtained from the scan:

$$V_L = \frac{V_\parallel + V_\perp}{2} = 0.92 \text{ MV} \tag{3.3.5a}$$

$$V_R = \frac{V_\parallel - V_\perp}{2} = 0.83 \text{ MV}, \tag{3.3.5b}$$

which indicate that the left input gets 55 % of the total power fed into the system, whereas the right input only gets 45 %—or stated differently, the right input gets 19 % less power than the left input.

A possible reason for this imbalance could be related to differences between the signal attenuation along each of the two arms of the split waveguide network connecting the 3-dB E-hybrid and the E-rotator. Assuming that the 3-dB E-hybrid operates as expected and distributes the incoming power into two equal parts, this would imply that the right arm introduces an additional attenuation of $20 \cdot \log_{10}(V_R/V_L) = -0.91$ dB with respect to the left arm (i.e., 19 % power loss). However, this is unlikely to be the case, since all waveguide segments are qualified within the standard tolerances (cf. Section 3.2). It is also unconceivable that the phase shifter is responsible for this difference, since the measurements performed at both CERN [Romano del Pozo et al., 2019] and DESY (cf. Section 3.2) show insertion-loss

values of $\lesssim 0.1$ dB and reflection-attenuation values of $\lesssim -26$ dB.

A more plausible explanation for the observed effects is related to the operation of the 3-dB E-hybrid. If the impedances at the two 3-dB-splitter outputs are not properly matched, the power can be unevenly distributed. Indeed, according to RF simulations of a conventional coax power splitter performed by the MIN group, a moderate mismatch factor of ~ 1.25 can account for a -0.9 dB difference between the two splitter outputs, corresponding to a realistic reflection-attenuation coefficient of -19.4 dB—i.e., 1.15 % of reflected power in one of the outputs. This is indeed confirmed experimentally by the power readings obtained from the directional coupler located before the splitter, which result in a reflection attenuation of $10 \cdot \log_{10}(P_{\text{refl}}/P_{\text{forw}}) = -19.5$ dB, where P_{refl} and P_{forw} are the reflected and forward power, respectively. This consideration is relevant to the power-to-voltage calculation discussed in the next section, since it means that, apart from the expected waveguide attenuation, the entire power fed into the 3-dB splitter reaches the structure.

3.3.2. Power-to-voltage constant D

In order to calculate the power-to-voltage constant of the TDS (cf. Section 3.1), the shear parameter S is measured at both zero crossings and at different RF-power values between 20 % and 100 % of the available klystron power. For all measurements, the quadrupoles between the TDS and the screen 11FLFXTDS are switched off to minimise eventual errors in the transport matrix. The beta function in the streaking plane is set to $\beta_{y,TDS} \approx 25$ m at the TDS and to $\beta_{y,scr} \approx 5$ m at the screen.

Assuming a waveguide attenuation of -0.1 dB/m and according to the waveguide-network geometry and the power imbalance previously measured, the power effectively fed into the structure represents 94 % of the power measured at the directional coupler located right upstream of the structure (cf. Appendix B). On the other hand, the integrated voltage V_0 is related to the measured shear parameter S by the expression:

$$V_0 = \frac{p_0 c}{e} \frac{1}{k_{RF}} \frac{S}{R_{33}\frac{L}{2} + R_{34}} \tag{3.3.6}$$

where $p_0 c = 750$ MeV is the beam energy, e the electron charge, $k_{RF} \approx 251$ m^{-1} the wavenumber of the structure, $R_{33} = 1.0$ and $R_{34} = 8.3$ m the transport-matrix terms from the TDS to the screen, and $L = 0.8$ m is the effective length of the structure. Notice that the thick-TDS matrix (cf. Eq. 1.3.12) has been used in Equation 3.3.6, since the transport-matrix element R_{33} is non-negligible compared to R_{34}.

The result of the measurement is shown in Figure 3.7, with an obtained power-to-voltage constant of $D = 4.6 \pm 0.1$ MV$/\sqrt{\text{MW}}$. This represents a considerable difference of almost 12 % respect to the nominal value of $D = 5.2$ MV$/\sqrt{\text{MW}}$, which could be due to several reasons. Most likely is that the LLRF signal sampled at the directional coupler used in the

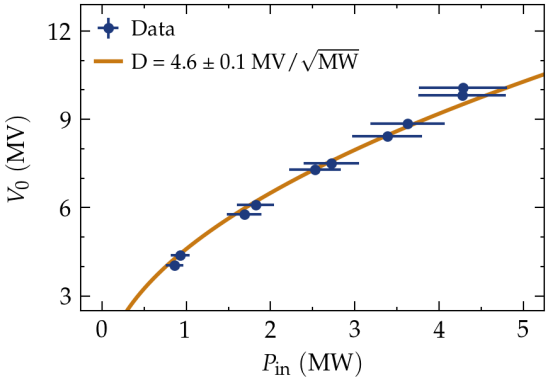

Figure 3.7.: Measurement of the power-to-voltage constant D.

measurement has a large associated error of (at least) -0.5 dB. Another explanation could be that the estimation of the waveguide losses discussed in Appendix B might not be sufficiently accurate. This is due to the fact that the calculation of how cascading imperfect parts affect the whole system is extremely involved, even though the RF-components used in the setup were individually qualified. Additionally, the definition of the power-to-voltage constant relates the power fed into the structure to the achieved particle deflection in the streaking plane. However, due to the power imbalance, a fraction of the input power is converted into a transverse (unwanted) kick, effectively decreasing the shunt impedance of the structure. On account of these considerations, it can be concluded that a physical model of the complete RF system around the TDS involves a high level of complexity, and that the uncertainties associated with the actual performance of the system as a whole prevent the calculation of the real power-to-voltage constant of the structure with the desired level of accuracy.

3.3.3. RF-system stability

In Section 2.3.2, different sources of jitter have been analysed and the relation between them has been emphasised. This section focuses on the experimental determination of the two jitter sources directly related to the PolariX-TDS system: the effective RF-phase jitter $\bar{\sigma}_{\phi,\text{eff}}$ and the RF-amplitude jitter $\bar{\sigma}_A$.

Effective RF-phase jitter $\bar{\sigma}_{\phi,\text{eff}}$

In order to estimate the effective RF-phase jitter of the PolariX-TDS system, the data from the power-to-voltage measurements discussed in the previous section is used. For each measured shear parameter S, the average of the centroid jitter sampled at each of the RF-phase scan steps is calculated together with the spread of the sample means—i.e., $\sigma_{\bar{X}} = \sigma/\sqrt{N}$, where N is the number of RF-phase steps. The computed values are subsequently fitted by Equation 2.3.3:

$$\bar{\sigma}_{\langle y\rangle,\phi+t} = \frac{|S|}{k_{RF}}\sqrt{\bar{\sigma}_{\phi}^2 + (k_{RF}\,c\,\bar{\sigma}_t)^2} = \frac{|S|}{k_{RF}}\,\bar{\sigma}_{\phi,\mathrm{eff}}\,,$$

where the effective RF-phase jitter $\bar{\sigma}_{\phi,\mathrm{eff}}$ represents the combined effect of the TDS RF-phase jitter $\bar{\sigma}_{\phi}$ and the linac arrival-time jitter $\bar{\sigma}_t$. The result of the measurement is shown in Figure 3.8. The linac-pointing jitter has a value of $\bar{\sigma}_r \approx 20$ µm, well below the TDS-induced jitter observed at the screen, and can therefore be ignored. The obtained effective RF-phase jitter is $\bar{\sigma}_{\phi,\mathrm{eff}} = 0.096 \pm 0.003$ deg (22.2 ± 0.7 fs), which represents an excellent RF stability that outperforms the expected FLASH-linac arrival-time jitter ($\bar{\sigma}_t \approx 50$ fs) by more than a factor 2.

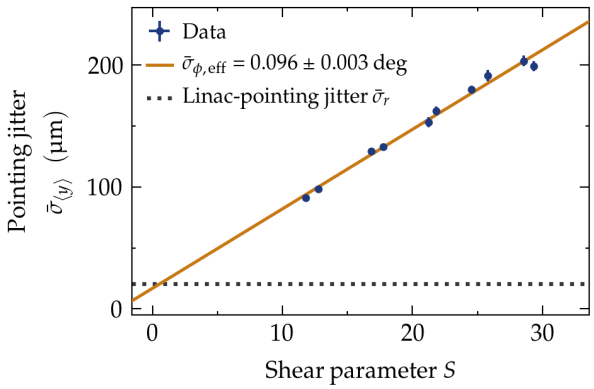

Figure 3.8.: Effective RF-phase jitter measurement.

RF-amplitude jitter $\bar{\sigma}_A$

In Section 2.3.2, the effect of the RF-amplitude jitter $\bar{\sigma}_A$ on the bunch centroid is demonstrated to be negligible compared to the effective RF-phase jitter $\bar{\sigma}_{\phi,\mathrm{eff}}$ when operating the TDS around the zero crossing. Its effect, on the other hand, translates into beam-size fluctuations at the screen:

$$\bar{\sigma}_{\langle y^2\rangle,A} = |S|\;\sigma_{\tilde{\xi}}\;\bar{\sigma}_A, \tag{3.3.7}$$

where $\sigma_{\tilde{\xi}}$ is the rms length of the bunch. If the bunches delivered by the linac had a perfectly stable (constant) length, the RF-amplitude jitter could be estimated by observing the beam-size variations at different values of the shear parameter. However, the bunch-compression schemes typically used at linacs (cf. Section 1.1.2) are affected by sources of jitter related, among others, to the RF stability of the accelerating modules. Therefore, eventual bunch-length variations $\bar{\sigma}_{\tilde{\xi}}$ must be taken into account in the computation of the beam-size variations

at the screen:

$$\bar{\sigma}_{\langle y^2\rangle, A+\xi} = |S| \sqrt{\left(\sigma_\xi \, \bar{\sigma}_A\right)^2 + \bar{\sigma}_\xi^2} = |S| \, \sigma_\xi \sqrt{\bar{\sigma}_A^2 + \bar{\sigma}_{\xi,\mathrm{rel}}^2} \tag{3.3.8a}$$

$$= |S| \, \sigma_\xi \, \bar{\sigma}_{A,\mathrm{eff}} \,, \tag{3.3.8b}$$

where the relative bunch-length jitter $\bar{\sigma}_{\xi,\mathrm{rel}} = \bar{\sigma}_\xi / \sigma_\xi$ and the effective RF-amplitude jitter $\bar{\sigma}_{A,\mathrm{eff}}$ have been introduced. Equations 3.3.8 reveal that the RF-amplitude jitter cannot be disentangled from the linac-induced bunch-length jitter. Nonetheless, if the linac-induced bunch-length variations can be estimated, their effect can be substracted from the effective RF-amplitude jitter, thus enabling the retrieval of the PolariX-TDS RF-amplitude jitter.

The effective RF-amplitude jitter $\bar{\sigma}_{A,\mathrm{eff}}$ is calculated from the power-to-voltage measurements discussed previously and the linac-induced bunch-length variations are measured with the bunch-compressor monitor BCM-4DBC3.2 located downstream of the second bunch compressor in the FLASH linac. The bunch length estimated from each of the two diagnostics is shown in Figure 3.9 and is plotted versus the shear parameter S only to identify the measurements—i.e., no correlation is expected between the shear parameter and the estimated bunch length. The absolute values obtained from the two diagnostics agree within the error bars (upper plots). However, the resolution achieved with the BCM is considerably worse than that achieved with the PolariX-TDS, and produce relative errors around $\bar{\sigma}_{\xi,\mathrm{rel}} = 17$ % (bottom-right plot). Therefore, the calculated effective RF-amplitude jitter $\bar{\sigma}_{A,\mathrm{eff}} = 3.6 \pm 0.4$ % (bottom-left plot) represents the best estimate of an upper bound of the RF-amplitude jitter $\bar{\sigma}_A$. Nonetheless, it is not unreasonable to expect that a considerable fraction of this value is associated with linac instabilities and, consequently, that the actual

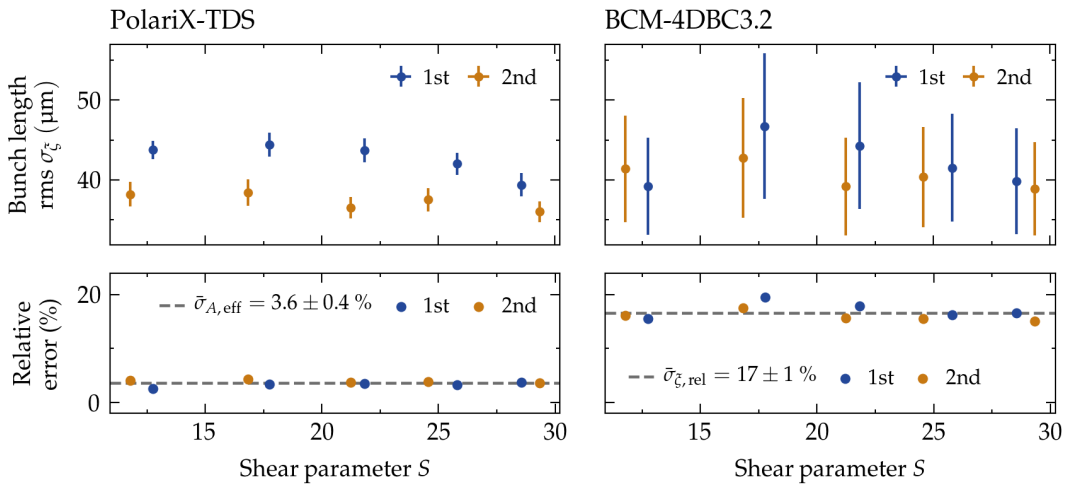

Figure 3.9.: RF-amplitude jitter $\bar{\sigma}_A$ calculation. Left: bunch length measurements performed with the PolariX-TDS. Right: bunch length measurements performed with the BCM-4DBC3.2.

RF-amplitude jitter of the PolariX-TDS system is very close to the predicted value of 1 %.

An alternative approach to estimate the RF-amplitude jitter $\bar{\sigma}_A$ would be to evaluate its effect on the pointing jitter at the screen for large RF phases. Far away from the zero crossing, the bunch centroid measured at the screen evolves with the RF phase as (cf. Equations 3.3.2):

$$\langle y_s \rangle = \langle y_{s,\beta} \rangle + \frac{S}{k_{RF}} \sin \phi_{RF}. \tag{3.3.9}$$

Considering both the RF-amplitude and the effective RF-phase jitter, the pointing jitter measured at the screen is obtained by propagating errors:

$$\bar{\sigma}_{\langle y \rangle, A+\phi} = \frac{S}{k_{RF}} \sqrt{\left(\sin \phi_{RF}\ \bar{\sigma}_A\right)^2 + \left(\cos \phi_{RF}\ \bar{\sigma}_{\phi,\mathrm{eff}}\right)^2}. \tag{3.3.10}$$

By operating the TDS on-crest (i.e., $\phi_{RF} = 90$ deg) the effect of the RF-phase jitter vanishes and the RF-amplitude jitter remains as the unique term contributing to the centroid pointing jitter:

$$\bar{\sigma}_{\langle y \rangle, A}(\phi_{RF} = 90) = \frac{S}{k_{RF}} \bar{\sigma}_A\ , \tag{3.3.11}$$

thus providing a way to measure the RF-amplitude stability of the system. Due to time constraints during the commissioning campaign, however, this approach could not be realised experimentally.

3.4. Beam-diagnostic applications

In this section two diagnostics applications that exploit the capability of the PolariX-TDS system to vary the polarisation of the streaking field are described: a time-resolved phase-space reconstruction and a dispersion-based beam-tilt correction. The first one is of general interest in a broad spectrum of beam-dynamics investigations. In the context of PWFA, in particular, the phase-space reconstruction can be used as an input for particle-in-cell (PIC) simulations, thus enabling the reproduction of the plasma-acceleration process and providing a better insight into the physics of the beam-plasma interaction (cf. Section 5.4). The second application enables the minimisation of the sliced transverse-centroid offsets of the beam and constitutes a valuable technique for the optimisation of electron bunches for PWFA experiments (cf. Section 4.3).

3.4.1. Time-resolved phase-space reconstruction

The reconstruction method presented here is based on the independent measurement of the longitudinal phase space and the slice emittances in x and y, and enables the retrieval of the time-resolved decoupled transverse phase-space of a particle distribution together with its longitudinal phase space. This means, that for each longitudinal slice the first moments of all dynamic variables and the following terms of the beam matrix are obtained:

$$\langle \boldsymbol{X} \rangle_\xi = \begin{pmatrix} \langle x \rangle_\xi \\ \langle x' \rangle_\xi \\ \langle y \rangle_\xi \\ \langle y' \rangle_\xi \\ \langle \xi \rangle \\ \langle \delta \rangle_\xi \end{pmatrix} \quad , \quad \boldsymbol{\Sigma}_\xi = \begin{pmatrix} \langle x^2 \rangle_\xi & \langle xx' \rangle_\xi & 0 & 0 & 0 & 0 \\ \langle x'x \rangle_\xi & \langle x'^2 \rangle_\xi & 0 & 0 & 0 & 0 \\ 0 & 0 & \langle y^2 \rangle_\xi & \langle yy' \rangle_\xi & 0 & 0 \\ 0 & 0 & \langle y'y \rangle_\xi & \langle y'^2 \rangle_\xi & 0 & 0 \\ 0 & 0 & 0 & 0 & \langle \xi^2 \rangle & 0 \\ 0 & 0 & 0 & 0 & 0 & \langle \delta^2 \rangle_\xi \end{pmatrix}, \tag{3.4.1}$$

where the subindex ξ identifies the slice and the longtudinal resolution is $R_\xi = \sqrt{\langle \xi^2 \rangle}$.

For this application, a bunch with a charge of $Q \approx 250$ pC, an energy of $\langle E \rangle = 750$ MeV, and an rms length of $\sigma_{\xi,rms} \approx 280$ fs (85 µm) with a negative energy chirp is used. The beamline is set up in the same way as described in Section 3.3, matching the beam to symmetric optics at the start of the FLFDIAG section, with $\beta_x = \beta_y = 10$ m and $\alpha_x = \alpha_y = 0$, which eases the performance of slice-emittance measurements in both transverse planes x–x' and y–y'—i.e., a single set of quadrupole-scan parameters can be used for both planes by just reversing the sign of the quadrupole strengths $k_1^i \to -k_1^i$. The lattice optic downstream of that point is adjusted to the requirements of each measurement separately. The different slow-RF feedback systems of the linac (cf. Section 2.1) are activated before the measurements, including the compression, the charge, and the orbit feedback.

Longitudinal-phase-space measurements

For the longitudinal-phase-space measurements, the beam parameters at the TDS in the streaking plane are set to $\beta_{y,TDS} = 100$ m and $\alpha_{y,TDS} = 0$. The lattice between the TDS and the screen is set to parallel-to-point imaging (cf. Section 1.3), with $R_{33} \approx 0$ and $R_{34} = 14.6$ m. Operating the TDS at its maximum power, with $V_0 \approx 10$ MV, and according to the normalised emittance $\epsilon_{y,n} = 1.46$ µm measured at the FLFMAFF section, the longitudinal resolution should be approximately $R_\xi \approx 3.1$ fs. In the dispersive plane, the design beta at the screen is $\beta_{x,\mathrm{scr}} \approx 9$ m, which, according to the normalised emittance $\epsilon_{x,n} = 2.14$ µm measured at the FLFMAFF section and a dispersion at the screen of $|D_x| \approx 0.777$ m, results in an estimated energy resolution of $R_\delta = 1.5 \cdot 10^{-4}$.

The first step of the measurement consists of performing a time calibration at each of the two zero crossings by means of an RF-phase scan and an energy calibration by means of a dipole-current scan. The results are summarised in Table 3.2. As can be observed, the shear parameters measured at the two zero crossings differ from each other by ≈ 10 %, which does not seem to be related to changes in RF power, since the RF-coupler readings indicate that the power is constant to within 1 % during all measurements. To clarify the causes of this behaviour, additional measurements would be required—which the strict time constraints of the commissioning campaign unfortunately did not allow. The differences in the shear parameter are also reflected in the longitudinal resolutions obtained and, on the other hand, are larger than the value estimated above, indicating a non-negligible impact of beam correlations at the screen and/or a mismatch between the beam and the lattice optic. The unstreaked beam size at the dispersive screen is shown in Figure 3.10. The top plot shows an analysis of the time resolution versus energy, with the left and right axes corresponding to the 1st and 2nd zero crossings, which are scaled according to the measured shear parameter S. The resolution has a peak around $\delta \approx -0.1$ % and decreases towards higher (right) and lower (left) energies. Since the bunch has an approximately linear (negative) chirp, this is expected to result in an effective variation of the resolution along the bunch. The mean resolution is indicated with a dotted-grey line, which corresponds to the projected beam size in the vertical axis (middle-left plot).

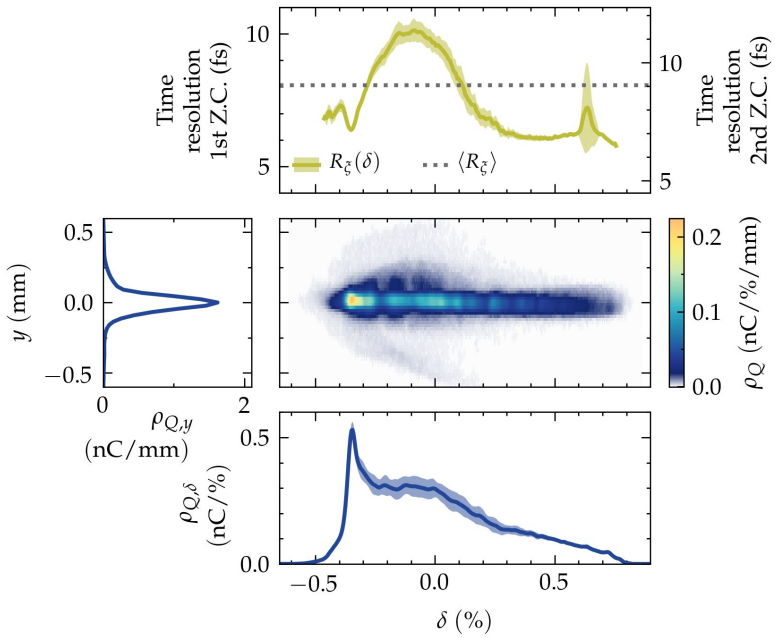

Figure 3.10.: Time resolution versus energy computed from the (dispersed) unstreaked beam size measured at the 8FLFDUMP screeen.

Parameter	Symbol	Units	Zero crossing	
			1st	2nd
Time calibration	t_{cal}	fs/pixel	-1.50	1.69
Shear parameter	S		-50.48 ± 0.02	44.96 ± 0.02
Time resolution	R_{ξ}	fs	8.06	9.05
Energy calibration	δ_{cal}	10^{-3}/pixel	-0.042	
Dispersion	D_x	m	-0.777	
Energy resolution†	R_{δ}	10^{-4}	1.5	

Table 3.2.: Summary of time and energy calibrations of the longitudinal-phase-space measurement performed at the 8FLFDUMP screen. †The energy resolution is estimated from the design optics used in the measurement.

The measured longitudinal phase space is shown in Figure 3.11. The two upper-right plots contain sample images measured at the 1st and 2nd zero crossing. The averaged projected energy $\rho_{\delta}(\delta)$ is plotted in the two left plots, and the averaged sliced energy spread $\sigma_{\delta}(\xi)$ and the current profile $I(\xi)$ in the two bottom-right plots. The differences between the two zero crossings are most prominent in the energy axis, whereas they are almost negligible in the longitudinal axis. The slice energy spread has a peak around the core of the bunch, which results from CSR effects (cf. Section 1.1.3) that arise in the bending magnets in the linac and in the FLASH2 extraction and FLASHForward extraction-compression sections. Owing to the good time resolution achieved, the effect of microbunching instabilities (cf. Section 1.1.3) is clearly seen in each of the sampled images. The large impact of TDS-induced effects on the energy axis observed in the measurements is shortly analysed below.

According to the particle-transport equation $\boldsymbol{X} = \boldsymbol{M}_{TDS} \cdot \boldsymbol{X}_0$, the sliced mean energy at the TDS exit can be expressed as:

$$\delta_{\pm}(\xi) = \delta_0(\xi) + \mathcal{K}_{\pm}\, y_0(\xi) + \frac{\mathcal{K}_{\pm} L}{2}\, y_0'(\xi) + \frac{\mathcal{K}_{\pm}^2 L}{6}\xi \,. \tag{3.4.2}$$

Thus, due to the large integrated kicks $\mathcal{K}$ generated with the PolariX-TDS system, the resulting energy profile $\delta_{\pm}(\xi)$ strongly depends on the zero crossing at which the TDS is operated—whenever the beam correlations $y_0(\xi)$ and $y_0'(\xi)$ are not negligible. In the longitudinal axis, in contrast, the large shear parameter S implies that the betatron-motion term at the screen $\sigma_{y,\beta}$ is negligible compared to the streaked beam size $S\sigma_{\xi}$ and, to a large extent, the measured current profile is independent of the zero crossing:

$$\sigma_y^2 = \sigma_{y,\beta}^2 + S^2\sigma_{\xi}^2 \quad \xrightarrow{\sigma_{y,\beta} \ll S\sigma_{\xi}} \quad \sigma_y \approx S\sigma_{\xi} \,. \tag{3.4.3}$$

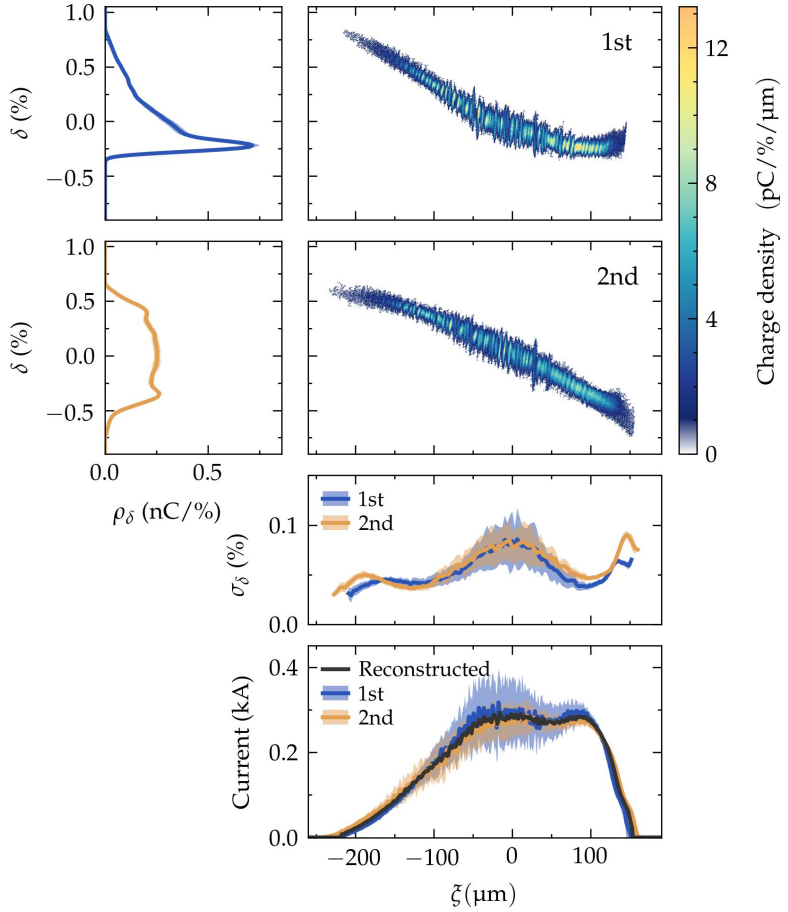

Figure 3.11.: Longitudinal-phase-space measurement performed at the 8FLFDUMP screen. See text for details.

In the following, a method to retrieve the true slice mean energy $\delta_0(\xi)$ of the bunch from the measurements performed at the two zero crossings is described.

Longitudinal-phase-space reconstruction method

To measure the energy experimentally, the observable used is the x coordinate of the particle distribution at the screen. According to the particle-transport equation $\boldsymbol{X} = \boldsymbol{R}_{TDS\to scr} \cdot \boldsymbol{M}_{TDS} \cdot \boldsymbol{X}_0$, the sliced centroids in x at the screen can be expressed as:

$$\begin{aligned} x(\xi) &= \underbrace{R_{11}\, x_0(\xi) \;+\; (LR_{11} + R_{12})\, x_0'(\xi)}_{x_\beta(\xi)} \;+\; R_{16}\Bigg(\underbrace{\delta_0(\xi) + \mathcal{K}\, y_0(\xi) + \frac{\mathcal{K}L}{2}\, y_0'(\xi) + \frac{\mathcal{K}^2 L}{6}\xi}_{\delta(\xi)\quad \text{(cf. Equation 1.3.48)}}\Bigg) \\ &= x_\beta(\xi) \;+\; D_x(s, s_0)\, \delta(\xi) \end{aligned} \tag{3.4.4}$$

where the first brace gathers the terms corresponing to the natural betatron motion of the centroids, and the second the true mean energy of the slice plus the changes induced by the TDS. Equation 3.4.4 reveals that the slice centroids at the screen $x(\xi)$ become modulated by the betatron-motion term $x_\beta(\xi)$ and, consequently, that the horizontal coordinate cannot be directly mapped back to energy $\delta(\xi) \neq x(\xi)/D_x(s, s_0)$. However, if the bunch correlations are negligible $x_\beta(\xi) \approx 0$ or the dispersive term largely dominates over the betatron-motion term, the slice mean energy measured at each of the zero crossings can be approximated by the expression:

$$\delta_\pm(\xi) \approx \frac{x_\pm(\xi)}{D_x(s, s_0)}, \tag{3.4.5}$$

and the measured slice mean energy can be related to the beam parameters at the TDS location by means of Equation 3.4.2. Notice that Eq. 3.4.2 is expressed as a function of ξ, which implies that the vertical screen coordinate y of the measurement has been mapped back to the internal bunch coordinate by means of the two-point-tomography method (cf. Equation 1.3.36). Using Eq. 3.4.2, the expression $\mathcal{K}_+\delta_-(\xi) - \mathcal{K}_-\delta_+(\xi)$ can be computed, which provides the true slice mean energy of the bunch:

$$\delta_0(\xi) = \frac{\delta_+(\xi) + \delta_-(\xi)}{2} - \frac{|\mathcal{K}|^2 L}{6}\xi, \tag{3.4.6}$$

where it has been assumed that typically $|\mathcal{K}_+| \approx |\mathcal{K}_-| = |\mathcal{K}|$. This expression will be used to reconstruct the longitudinal phase space later on.

By reasoning along the same lines, the bunch correlations in the full y–y' plane can be additionally retrieved. The term $y_0'(\xi)$ can be directly obtained from the two-point-tomography, since according to the parallel-to-point imaging optics typically used in time-resolved measurements, only the kicks at the location of the TDS are imaged at the screen—i.e., $R_{33} \approx 0$ and $R_{34} \gg 1$. Therefore, to a very good approximation, the betatron term at the screen retrieved with the two-point-tomography $y_\beta(\xi)$ contains information about beam correlations in the y'–ξ plane at the TDS:

$$y_\beta(\xi) = R_{34}\left(\frac{R_{33}}{R_{34}}\, y_0(\xi) + y_0'(\xi)\right) \quad \xrightarrow{\dfrac{R_{33}}{R_{34}} \ll 1} \quad y_0'(\xi) \approx \frac{y_\beta(\xi)}{R_{34}}. \tag{3.4.7}$$

On the other hand, the term $y_0(\xi)$ can be retrieved from the measurements at the two zero crossings by substracting the slice mean energy for the negative streak $\delta_-(\xi)$ from that for the positive streak $\delta_+(\xi)$. When doing so, after rearranging terms, the following expression results:

$$y_0(\xi) = \frac{\delta_+(\xi) - \delta_-(\xi)}{\mathcal{K}_+ - \mathcal{K}_-} - \frac{L}{2} y_0'(\xi) - \frac{L}{6} \frac{(\mathcal{K}_+^2 - \mathcal{K}_-^2)}{\mathcal{K}_+ - \mathcal{K}_-} \cdot \xi \tag{3.4.8}$$

$$\approx \frac{\delta_+(\xi) - \delta_-(\xi)}{2|\mathcal{K}|} - \frac{L}{2} y_0'(\xi), \tag{3.4.9}$$

where again it has been assumed that typically $|\mathcal{K}_+| \approx |\mathcal{K}_-| = |\mathcal{K}|$.

Slice emittance in x and y

The reference plane used in the slice-emittance measurements corresponds to the start of the FLFDIAG section—in which, according to the beamline setup, the Courant-Snyder parameters are symmetric, with $\beta_x = \beta_y = 10$ m and $\alpha_x = \alpha_y = 0.0$. Furthermore, the design optic for the measurement of the horizontal phase space x–x' is optimised to: 1) keep the Courant-Snyder parameters in the streaking (vertical) plane at the location of the TDS at $\beta_{y,TDS} = 100$ m and $\alpha_{y,TDS} = 0$; 2) set the lattice between the TDS and the screen 11FLFXTDS to parallel-to-point imaging, with $R_{33} \approx 0$ and $R_{34} = 7.6$ m. These two conditions produce a beta at the screen of $\beta_{y,\mathrm{scr}} \approx 0.5$ m. Assuming an integrated TDS voltage of $V_0 \approx 10$ MV, a shear parameter of $S = 27$ is obtained, which, according to the normalised emittance $\epsilon_{y,n} = 1.46$ µm measured at the FLFMAFF section, results in a longitudinal resolution of $R_\xi = 0.9$ µm (3.0 fs). In the horizontal axis, on the other hand, the phase advance between the reference plane and the screen is progressively scanned between 0 deg and 180 deg, whereas the beta function at the screen is fixed to $\beta_{x,\mathrm{scr}} = 25$ m. The key parameters of the scan are shown in Figure 3.12, including the beta functions at the screen β_x and β_y, the phase advance in x between the reference plane and the screen $\Delta\mu_{x,\mathrm{ref}\to\mathrm{scr}}$, the phase advance in y between the TDS and the screen $\Delta\mu_{y,\mathrm{TDS}\to\mathrm{scr}}$, and the quadrupole strengths of the magnets used—excluding those between the TDS and the screen, which are always kept at the same values. The evolution of the beta function in x and y for all scan steps is shown in Figure 3.13. Since the beam optic at the reference plane is symmetric, when measuring the slice emittance in the vertical plane y–y' the quadrupole strengths only need to be reversed and the polarisation of the streak rotated by 90 deg.

The shear parameter S is calibrated at both zero crossings only once before each slice-emittance scan. The results are summarised in Table 3.3. Since the FOV of the screen-optic system is smaller in the horizontal axis, the RF power is reduced when streaking the bunch horizontally, which is reflected by the lower shear-parameter values obtained for the vertical slice-emittance. Similar to what has been observed in the longitudinal-phase-space screen, the shear parameters at the two zero crossings differ by ~ 10 % and 5 % for the horizontal and vertical slice-emittance measurements, respectively, whereas the RF power in all measurements performed at the 1st zero crossing is only 5 % lower than at the 2nd zero crossing—which should result in integrated voltage differences of only ~ 2.5 %.

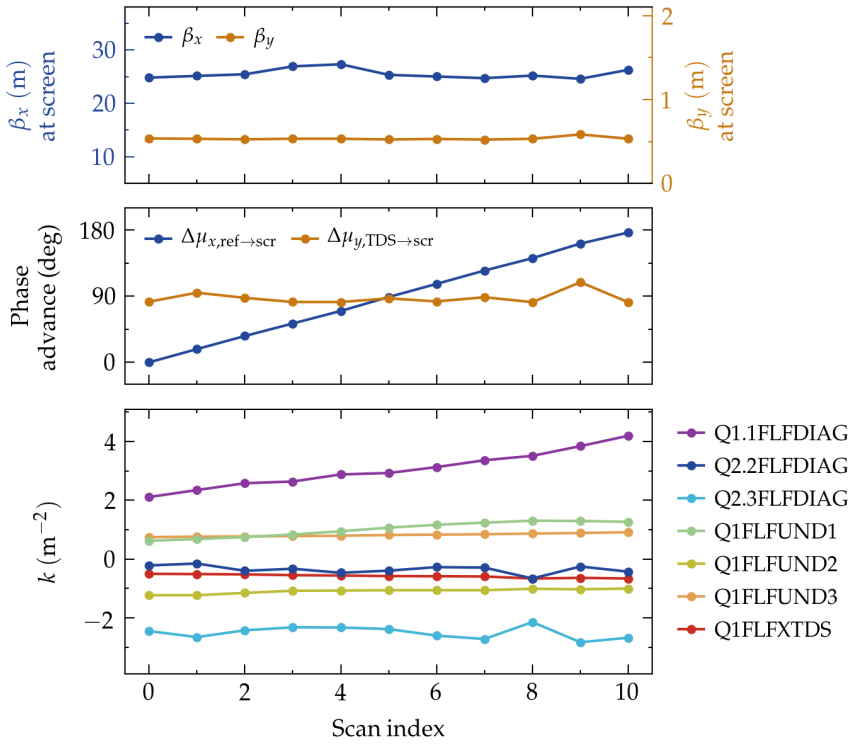

Figure 3.12.: Key design parameters of the slice-emittance measurement. Top plot: beta functions at the screen in x and y. Middle plot: phase advance in x between the reference plane and the measurement screen $\Delta\mu_{x,\mathrm{ref}\to\mathrm{scr}}$ scanning 180 deg; phase advance in y between the TDS and the screen $\Delta\mu_{y,\mathrm{TDS}\to\mathrm{scr}}$, kept constant at 90 deg. Bottom plot: quadrupole strengths of the magnets used in the scan.

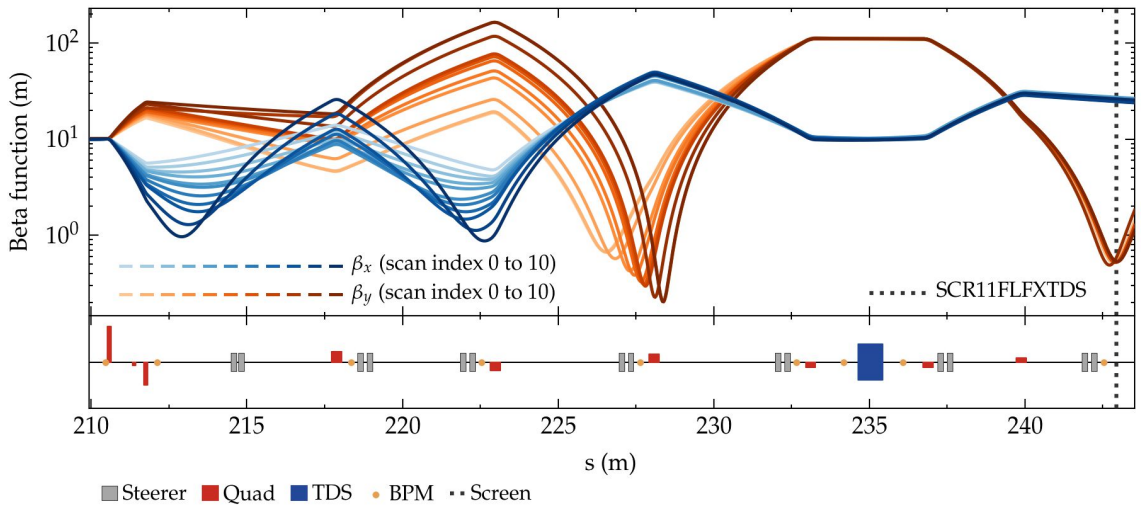

Figure 3.13.: Lattice optic for the slice-emittance measurement in the horizontal plane. For the measurement in the vertical plane the transverse axes are swapped.

Parameter	Symbol	Units	**Horizontal emittance**	
			1st Z.C.	2nd Z.C.
Time calibration	t_{cal}	fs/pixel	-0.77	0.69
Shear parameter	S		-24.02 ± 0.01	26.75 ± 0.01
Time resolution	$R_{\tilde{\zeta}}$	fs	9.3	8.3
			Vertical emittance	
			1st Z.C.	2nd Z.C.
Time calibration	t_{cal}	fs/pixel	-0.89	0.86
Shear parameter	S		-20.41 ± 0.01	21.49 ± 0.01
Time resolution	$R_{\tilde{\zeta}}$	fs	10.7	10.4

Table 3.3.: Summary of the shear-parameter calibration performed for the slice-emittance measurements in the horizontal and vertical planes.

Figure 3.14 shows sample images taken from the horizontal slice-emittance measurement. As can be observed, the sliced beam size decreases towards the middle of the scan, hinting at a potential mismatch between the beam and the lattice optic. Additionally, the bunch length at the 1st zero crossing seems to be larger around the 6th and 7th scan indices. If these differences were caused by bunch correlations, the bunch length at the 2nd zero crossing should be correspondingly shorter. However, this is not observed, which indicates that the real bunch length is indeed varying along the scan. This could be due to several reasons, the most likely of which are: slow drifts of the compression settings, variable charge generated by the photocathode even though the slow-RF feedback systems are activated, or transmission problems.

In order to exclude the presence of slow drifts in the SCRF modules that affect the compression of the bunch at the linac, the bunch length measured at the BCM-4DBC3.2 located downstream of the 2nd bunch compressor (cf. Section 2.1) is analysed. Figure 3.15 shows the length measured at the BCM (grey line) over the complete time span of the slice-emittance measurement, together with the rms length reconstructed from the PolariX-TDS measurements by means of the two-point-tomography (red circles connected with solid lines). The time range of each scan step is indicated with blue and orange vertical bands for the 1st and 2nd zero crossing, respectively. Up to the 4th scan index, the BCM values overestimate those of the TDS by roughly ~ 10 %—most probably due to differences in time resolution between the two diagnostics—but the trend exhibited is similar in both of them. Between steps 5 and 8, however, the TDS values describe a bump that is not observed in the BCM, thus suggesting that the bunch-length variations originate downstream of the linac. The vertical-slice-emittance measurements exhibit a similar behaviour, with bunch-length variations that go up to 10 % around the 6th scan index and charge losses of ~ 20 % towards the initial and final steps of the scan.

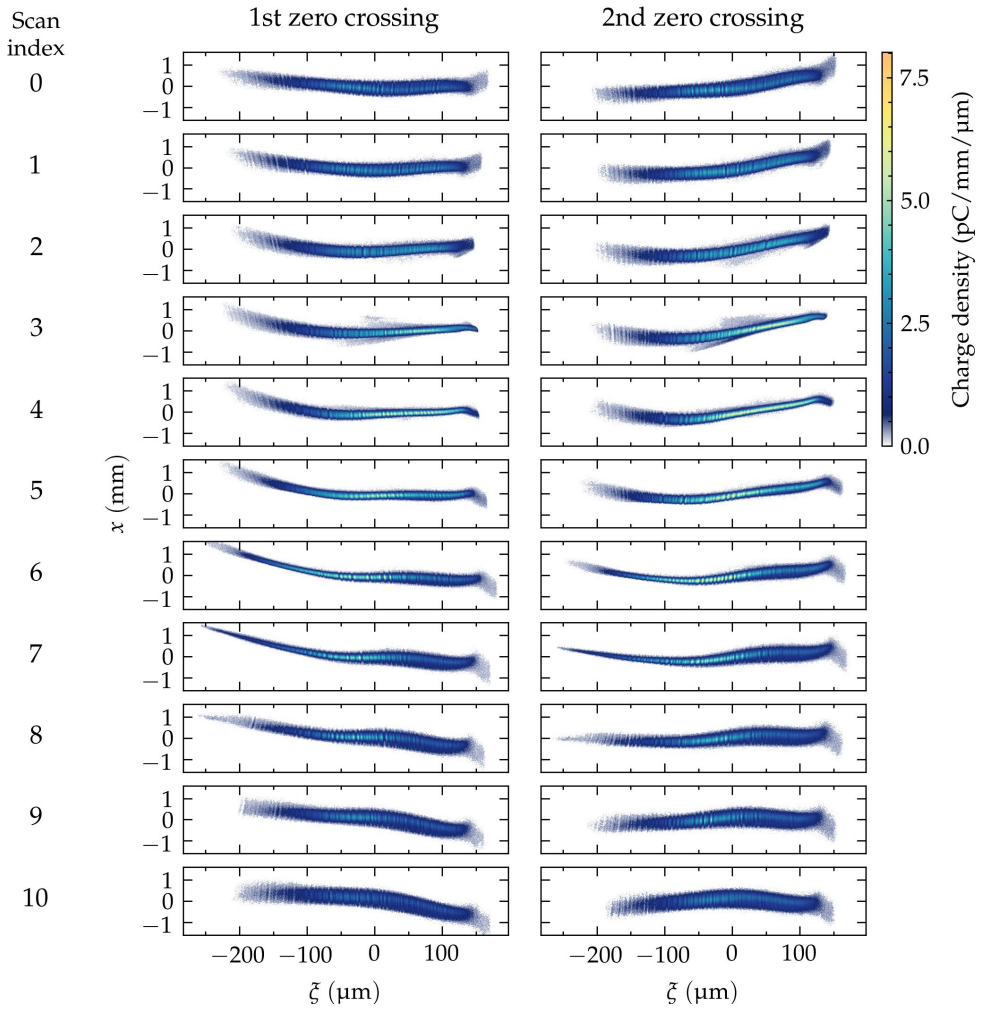

Figure 3.14.: Sample images for each of the scan indices of the horizontal slice-emittance measurement for the 1st (left) and 2nd (right) zero crossing.

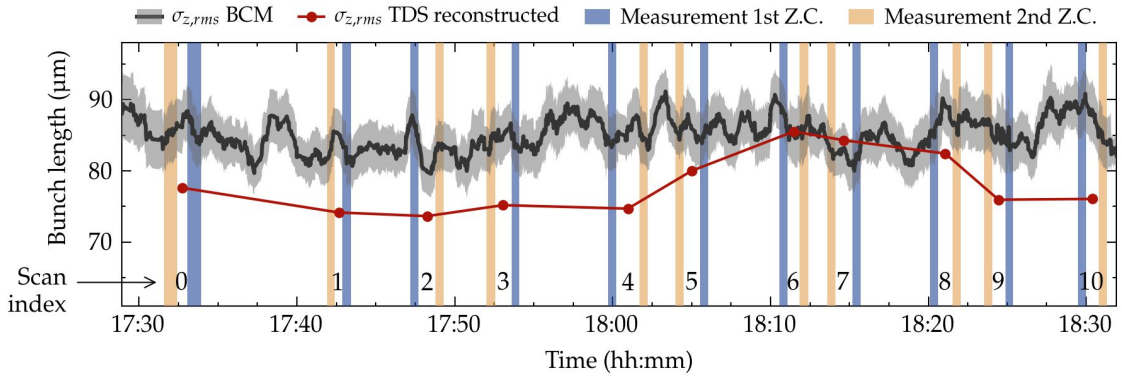

Figure 3.15.: Comparison between the bunch length measured with the PolariX-TDS (red lines) and with the BCM-4DBC3.2 located downstream of the second bunch compressor at the FLASH linac.

To investigate transmission problems in the FLASHForward beamline, the rms bunch length reconstructed by means of the two-point tomography is compared to the charge obtained from the CCD-camera counts. This is shown in Figure 3.16. The relative length variation (left axis) exhibits fluctuations of up to ~ 10 % and has its peak around the 6th scan index. On the other hand, the relative charge (right axis) reveals that the measurement is substantially affected by transmission problems that translate into a progressive charge loss slightly below and above 10 % towards the initial and final steps of the scan, respectively. The fact that the charge loss does not correlate with the bunch-length variations does not preclude that a defective transmission indeed leads to the observed variations. The small aperture of the PolariX-TDS (with a diameter of 8 mm) and the large beam size at its location ($\beta_{y,\mathrm{TDS}} \approx 100$ m) in combination with complex bunch correlations between the transverse and longutudinal planes (y–ξ and y'–ξ) and eventual transverse misalignments, could lead to a loss of charge from different regions of the bunch depending on the quadrupole-scan settings. Due to time constraints during the experimental campaign, however, the performance of a new iteration of slice-emittance measurements with an optimised beam transmission could not be performed. Nevertheless, this experience suggests that, whenever a high longitudinal resolution R_ξ is not required, the beta function of the beam at the location of the TDS should be reduced to moderate values of, e.g., $\beta_{y,\mathrm{TDS}} \approx 50$ m.

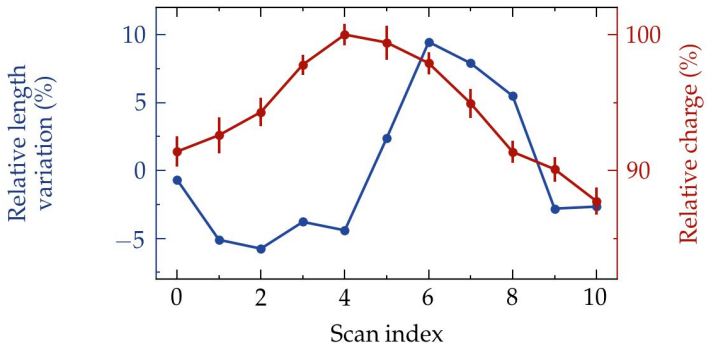

Figure 3.16.: Reconstructed rms bunch length (left axis) and charge measured at the screen computed from the CCD counts (right axis) along the slice-emittance measurements.

The calculation of the slice emittance is performed according to the procedure described in Section 1.2, i.e., by analysing the beam size in slices of width equal to the longitudinal resolution. The results are shown in Figure 3.17, including the normalised slice emittance $\epsilon_{x/y,n}$ and the beta $\beta_{x/y}$ and alpha $\alpha_{x/y}$ functions at the reconstruction plane and the resulting mismatch parameter $\mathrm{B}_{\mathrm{mag},x/y}$. The design Courant-Snyder parameters $\beta_x = \beta_y = 10$ m and $\alpha_x = \alpha_y = 0.0$, corresponding to a perfect match of $\mathrm{B}_{\mathrm{mag}} = 1.0$, are indicated with grey-dashed lines. The current profile is computed as the average of the reconstructed current profiles computed for each scan step. On average, the horizontal emittance $\epsilon_{x,n} = 0.85 \pm 0.03$ µm is roughly 3 times larger than the vertical emittance $\epsilon_{y,n} = 0.34 \pm 0.01$ µm, which can be associated to CSR effects (cf. Section 1.1.3) that arise in the horizontal dipoles along the beamline, since both the bunch compression and the extraction from the linac take place in the horizontal plane. The slice emittance increase observed in both planes at the head

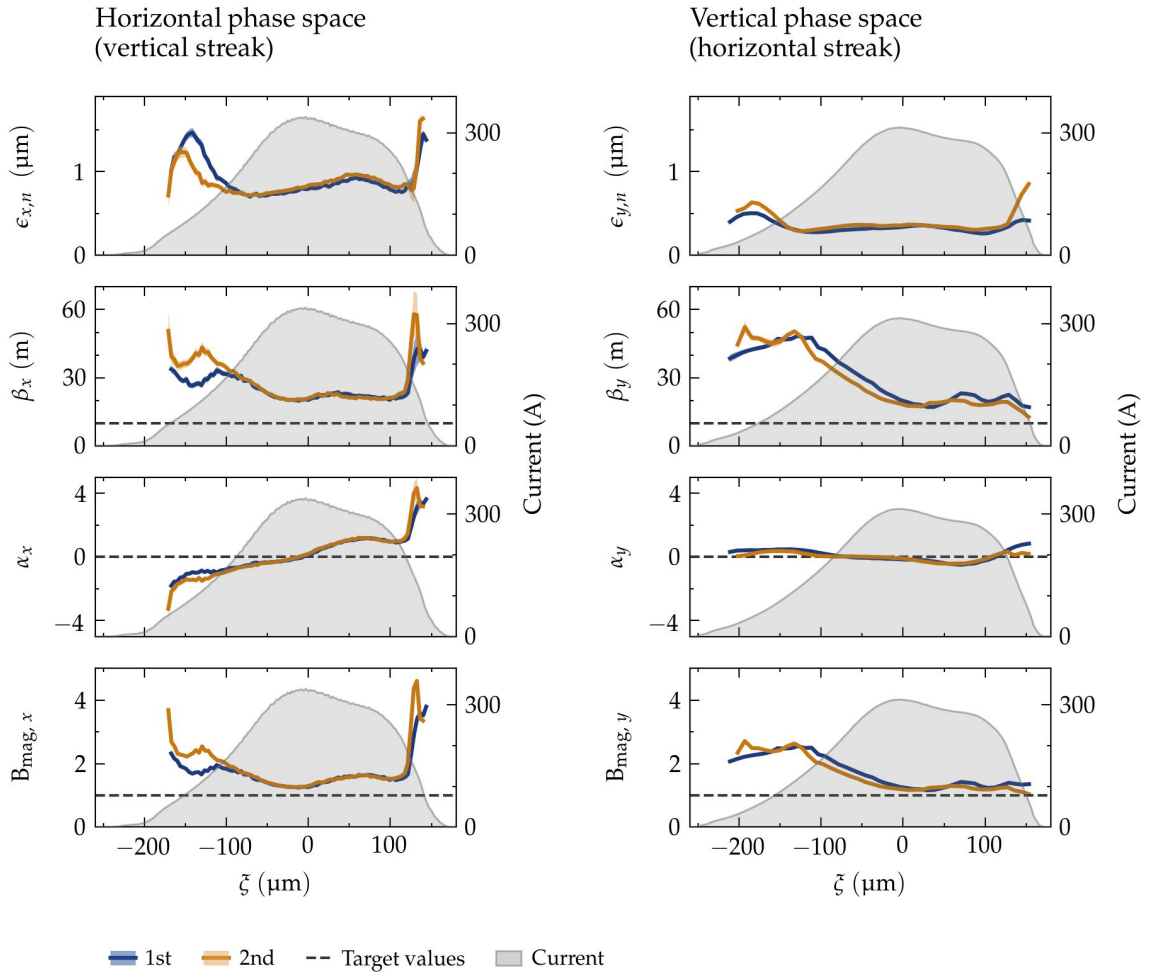

Figure 3.17.: Slice emittance and slice Courant-Snyder parameters in the horizontal (left) and vertical (right) planes.

and tail of the bunch is most probably the result of an overfocusing produced by the solenoid field of the RF gun due to the decrease of charge density—and, consequently, transverse-space-charge forces—at the front and rear parts of the bunch [Röhrs, 2008]. However, this interpretation might be compromised by measurement errors, since the signal-to-noise ratio of the images taken at the screen decreases at the regions of low charge density. The slight increase of the horizontal emittance around $\xi \approx 50$ µm could be caused either by CSR effects or by TSC forces changing along the bunch according to the charge density. The beta functions in both planes are consistently higher than the design value, clearly indicating that the bunch is not well matched to the lattice optic. Similarly, the alphas in the horizontal plane reveal that the bunch head converges whereas the tail diverges—i.e., the head is focused downstream of the reference plane, whereas the tail is focused upstream. In the vertical plane, on the other hand, the alpha has values very close to the design parameter $\alpha_y \approx 0$. These effects are reflected in the mismatch parameter which is larger than one in both planes $B_{mag,x/y} > 1.0$.

Figure 3.18 shows the sliced centroids in the transverse planes x–x' and y–y' computed by averaging the measurements performed at the 1st and 2nd zero crossings. The shaded ellipses represent the area occupied by the projected rms emittance—i.e., $A = \pi \cdot \epsilon_{rms}$. Both the horizontal and the vertical phase space exhibit relatively small divergences on the order of 10 µrad, whereas the centroids are on the same order of magnitude as the bunch length $\sim$ 100 µm. In the horizontal plane, the transverse displacements at the front of the bunch are especially pronounced.

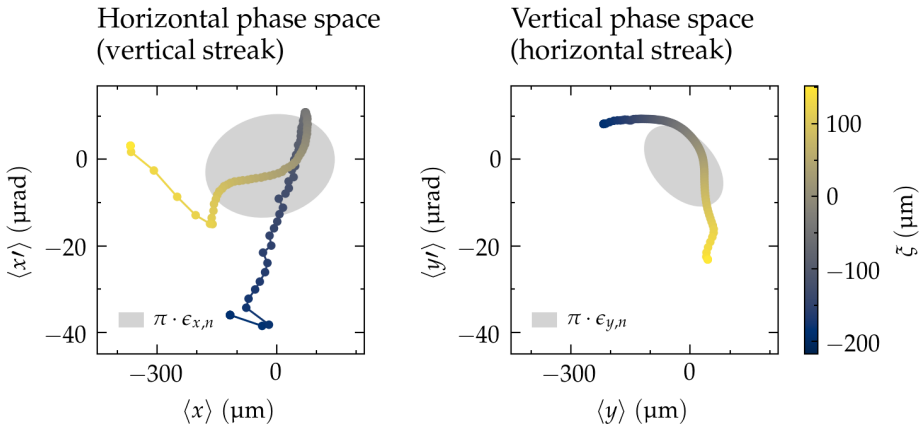

Figure 3.18.: Transverse-slice centroid offsets in the horizontal (left) and the vertical (right) planes. The plots show the average between the two zero crossings.

Validation of the phase-space reconstruction using particle-tracking simulations

In order to validate the phase-space reconstruction method presented in the previous sections, particle-tracking simulations were carried out with the code Ocelot [Agapov et al., 2014]. After taking into account the TDS-induced effects—i.e., induced energy spread and induced energy chirp—the bunch parameters retrieved from the longitudinal-phase-space and the slice-emittance measurements are used to reproduce the 6D particle-density distribution at the reference plane—i.e., at the start of the FLFDIAG section. The simulated bunch is subsequently used to recreate the longitudinal-phase-space measurements at the 8FLFDUMP screen at both zero crossings, and the results are compared to the real measurements. The comparison is shown in Figure 3.19, with the real experiment on the left and the simulation on the right. The good agreement between the two validates the reconstruction method presented in this section. Additionally, it supports the argument—already anticipated in Section 1.3—that the large differences between the two zero crossings in the energy axis are a TDS-induced effect that results in an amplification of the internal beam correlations in the streaking plane y–ξ and y'–ξ.

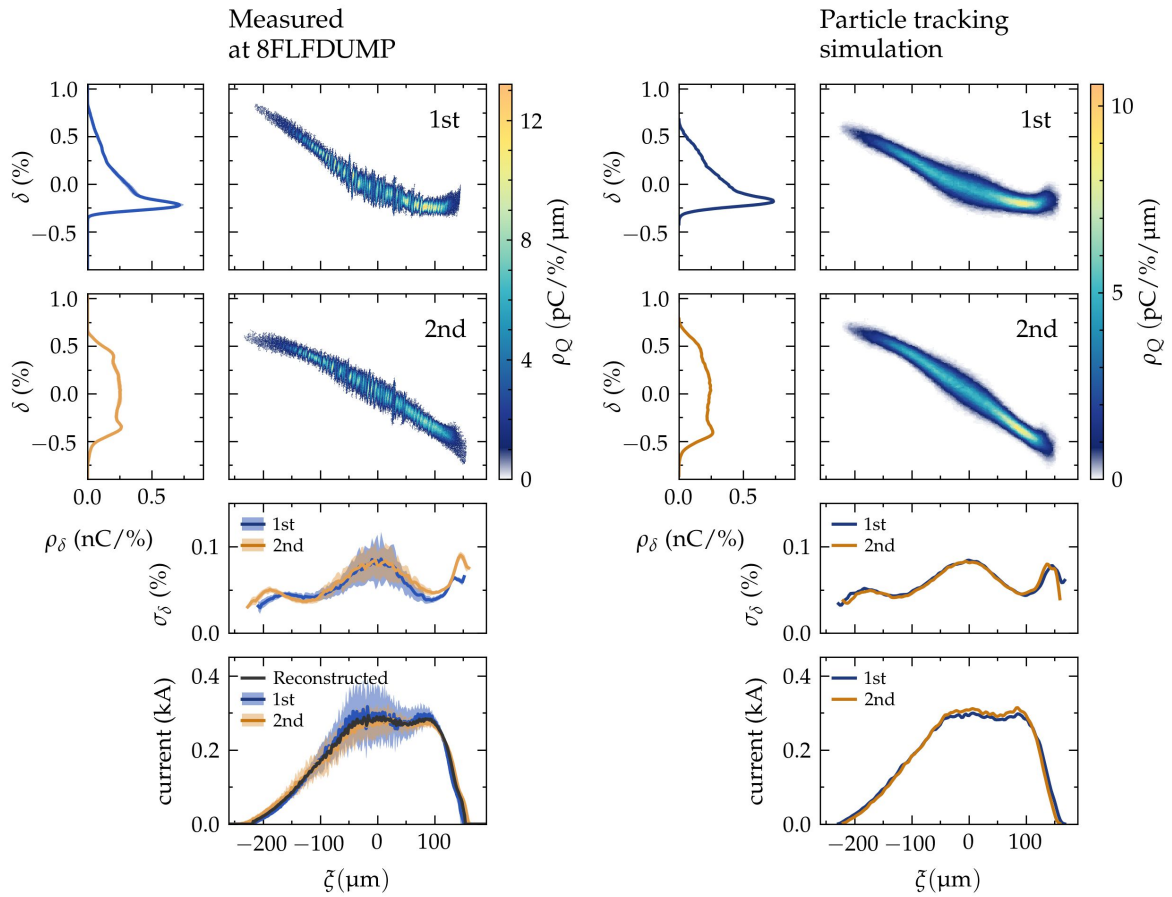

Figure 3.19.: Comparison between experimental and simulation longitudinal-phase-space measurements performed at the screen 8FLFDUMP.

3.4.2. Dispersion-based beam-tilt correction

In beam-driven plasma-wakefield (PWFA) experiments performed in the blowout regime (cf. Section 1.4), the presence of transverse-slice centroid offsets leads to instabilities that can severely compromise the acceleration process. In the case of a driver bunch, these offsets seed the hosing instability (cf. Section 1.4.3), by which the tranverse-centroid displacements are amplified exponentially during propagation in the plasma, ultimately leading to a beam breakup. Furthermore, in external-injection schemes, transverse misalignments betwen the driver and the trailing bunch cause the latter to undergo transverse oscillations. Depending on the amplitude of these oscillations, the interaction of the beam with the walls of the plasma cavity can lead to charge losses and to a deformation of the wakefield structure, eventually reducing the quality of the accelerated bunch (cf. Sections 1.4.3 and Chapter 5). Moreover, for large propagation lengths in plasma these oscillations progressively smear out the transverse phase space of the bunch, thus leading to an emittance growth (cf. Section 1.4.3).

There are multiple sources of slice centroid offsets, including, amongst others, CSR effects, coupler kicks in the RF-acceleration modules, transverse wake-fields and spurious dispersion (cf. Section 1.1). In the following, the strategy followed at FLASHForward to correct these transverse offsets is outlined, taking advantage of: 1) the coupling between the transverse and the longitudinal beam dynamics that arise in dispersive sections when the bunch has a linear energy chirp; and 2) the diagnostics capabilities of the PolariX-TDS. As mentioned in Section 1.1.2, the method was proposed by Guetg, Beutner et al., 2015 and examples of its application can be found in [Guetg, F. J. Decker et al., 2016; Vogt and Zemella, 2021]. Since at FLASHForward a purely empirical approach of the method is followed, a mathematical model—which can be found in the aforementioned references—is not discussed here.

Figure 3.20 shows the design lattice optics along the FLASH2 extraction (FL2EXTR) and the FLASHForward extraction (FLFEXTR) and compression (FLFCOMP) sections. Since the extraction from the linac is performed with vertical kickers and a horizontal Lambertson septum (cf. Section 2.1), both horizontal and vertical dispersions are produced at the FLASH2 side. Appropriate magnets to correct vertical offsets are the quadrupole Q14FL2EXTR and, to a lesser degree, the dipole D16FL2EXTR. This is because the vertical dispersion D_y in this region has the highest values, whereas the horizontal dispersion has already been closed.

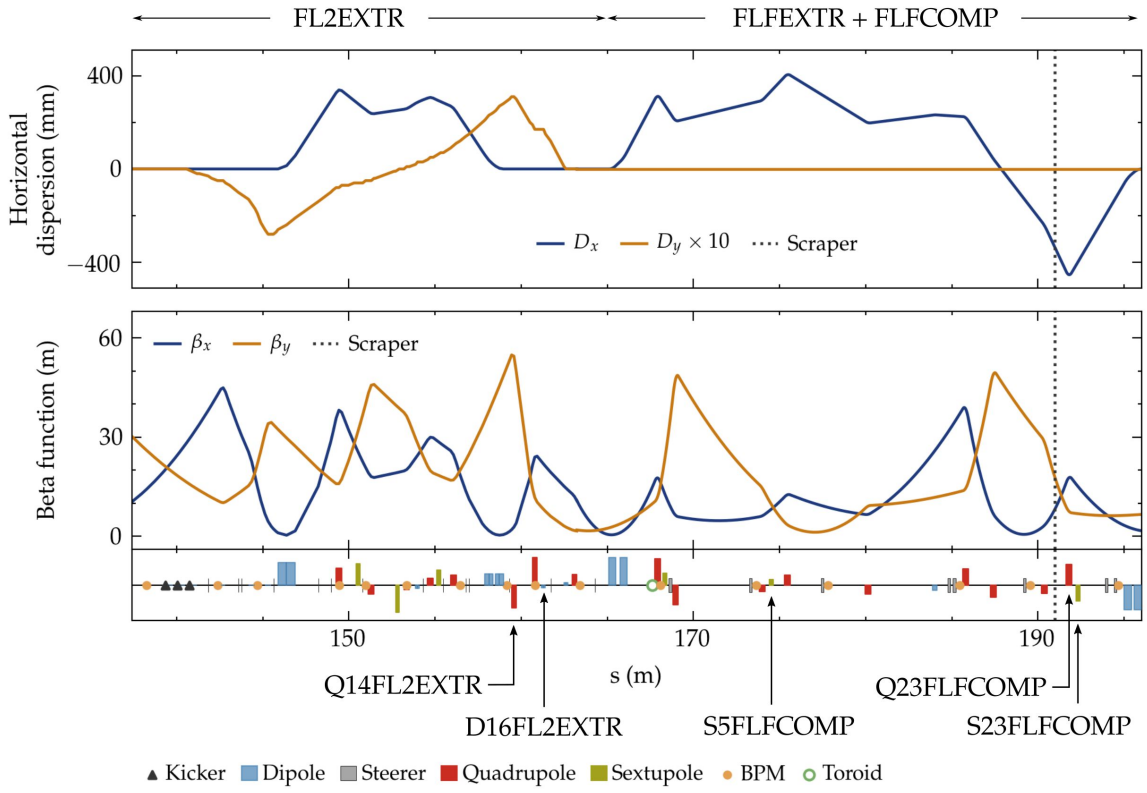

Figure 3.20.: Design lattice optic of the FLASH2 extraction (FL2EXTR) and the FLASHForward extraction (FLFEXTR) and compression (FLFCOMP) sections.

The fact that the vertical beta function β_y at the location of the quadrupole is relatively large, however, limits the flexibility to change its focusing strength, which might otherwise excessively distort the beam envelope. On the other hand, the horizontal beta function β_x goes through a minimum, which implies that the effect of the changes on the horizontal plane is, to a large extent, negligible.

To compensate the horizontal slice centroid offsets, additional considerations must be taken into account. As mentioned in Sections 1.1.2 and 2.2, at FLASHForward, the transverse-longitudinal coupling that an energy-chirped beam develops in a dispersive section is used to split the bunch delivered by the linac into a double-bunch structure by means of a wedge collimator located at $s \approx 191$ m (grey-dotted lines in Figure 3.20). One of the first steps in the procedure followed to set up the FLASHForward beamline requires the fine tuning of the bunch compression at the linac and the transport of the bunch in order to optimise the production of the double-bunch structure. Therefore, any change introduced upstream of the collimator after the optimisation, will be detrimental to the purposes of the experiment. On account of that, appropriate magnets to correct horizontal offsets are the quadrupole Q23FLFCOMP and the sextupole S23FLFCOMP, located right after the collimator and before the end of the FLFCOMP section. In this narrow region, the horizontal dispersion is close to its largest values and both the horizontal and the vertical beta functions are moderate. When the considerations about the longitudinal-phase-space optimisation are not required, additional magnets can be used. The sextupole S5FLFCOMP, for instance, is a good candidate to manipulate the curvature of the bunch. This is indeed the magnet used in a measurement performed to test the dispersion-based beam-tilt correction routines discussed in this section, which is described in the following.

The measurement focuses on the manipulation of the tilt and curvature of the bunch in the horizontal plane. The bunch is monitored at the screen 11FLFXTDS located at the straight section after the PolariX-TDS, which is adjusted to streak in the vertical plane. The focusing strength of the quadrupole Q23FLFCOMP is used to control the tilt, whereas that of the sextupole S5FLFCOMP to control the curvature, and their values are varied simulateously to perform a 2D scan. Figure 3.21 shows sample images for selected scan indices. As can be observed, for each sextupole setting, the variation of the quadrupole current translates into a change of the tilt from negative (bottom) to positive (top) values. Conversely, for each quadrupole setting, the variation of the sextupole current changes the curvature of the bunch from convex (left) to concave (right). This is most clearly appreciated for vertical tilts—i.e. for quadrupole currents around $I \approx 56.07$ A.

It must be noticed that a measurement as the one described above does not provide direct information about the dynamics of the particles. This means, that the "straightening" of the bunch monitored at the screen only guarantees that the tilt and curvature has been corrected at the location of the measurement—i.e. at the screen. In order to make sure that the transverse-slice slopes $\langle x' \rangle_\xi$ and $\langle y' \rangle_\xi$ have been corrected as well, the measurements must be performed at different phase advances. This is effectively equivalent to performing a manual slice-emittance scan.

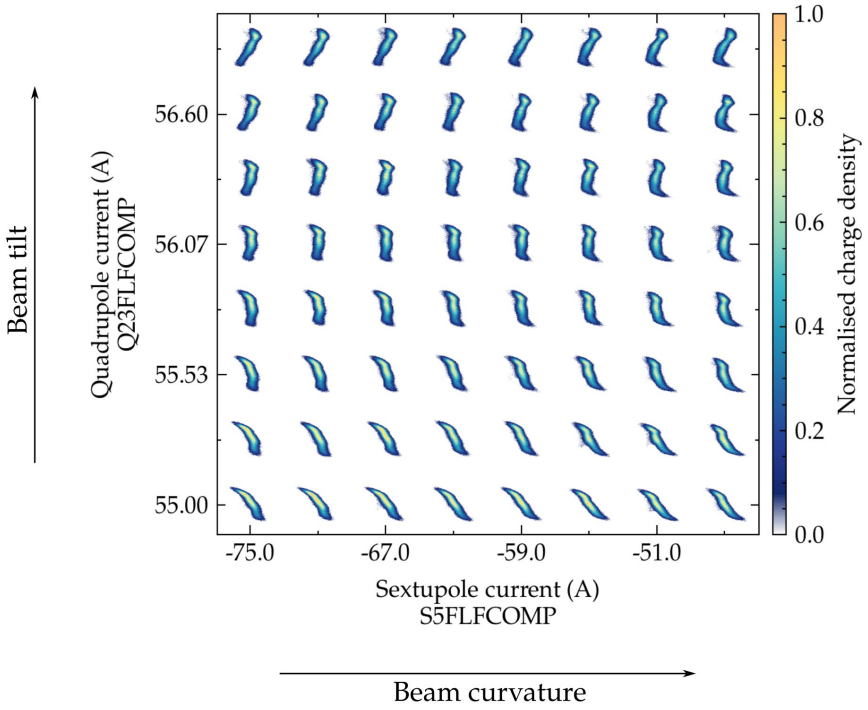

Figure 3.21.: Tilt-curvature 2D scan performed with Q23FLFCOMP and S5FLFCOMP.

4. Optimisation and characterisation of FLASH beams for plasma acceleration

In the following, the optimisation and characterisation of FLASH electron beams for plasma acceleration will be addressed in the context of the X-2 core experiment at FLASHForward [Libov et al., 2018; Lindstrøm, Beinortaite et al., 2021]. As explained in Chapter 2 the main goal of this experiment is to demonstrate quality preservation of an externally injected beam after its acceleration in a beam-driven plasma wakefield. One of the most recent achievements of X-2 has been to demonstrate energy-spread preservation and high energy-transfer efficiency by strongly beam loading the wakefield with tailored-current-profile bunched beams [Lindstrøm, Garland et al., 2021]. On the grounds of this success and with the powerful diagnostic capabilities of the PolariX-TDS, an ensuing experimental campaign was planned to further investigate two aspects that lead to quality preservation: optimal beam loading for chirp compensation and uniformity of the accelerating gradients inside the volume of the wakefield cavity.

The current chapter deals with the first half of this campaign, which comprises the preparation of the experiment and the establishment of a proper working point. The characterisation of the accelerated beams and a discussion of the ultimate goals pursued with the campaign are described in Chapter 5. As the title suggests, the aim here is twofold: 1) to provide a broad perspective of the complete procedure followed to set up the experiment and to introduce some of the distinct techniques used therein; 2) to characterise the attained working point to such a level of accuracy as to lay a solid ground for the discussion of the next chapter.

With few exceptions, all the data presented herein was taken during a dedicated ARD[1] beam time that took place between the 28th of October and the 1st of November 2020. Framed in the chronological sequence of events and with the two aforementioned aims in mind, the present chapter is structured as follows: Section 4.1 gives an overview of the linac and the beamline setup routines and elaborates on the resulting beam parameters; Section 4.2 describes three preliminary steps critical to the success of the X-2 experiment: a) the fine tuning of the longitudinal phase space of the electron bunch; b) its coarse straightening in both transverse planes (x–ξ and y–ξ); c) the strong-focus optics setup required to match the beam to a given plasma density. The subsequent optimisation of the interaction between the beam and the plasma is addressed in Section 4.3, concluding a coherent sequence of three

[1] Accelerator Research and Development

shift blocks extending over 24 hours. In the last part of this section, the specificities of the plasma interaction required for the actual PolariX-TDS campaign are addressed. After that, in Section 4.4, the working point of the PolariX-TDS campaign is characterised.

4.1. Linac and beamline setup

The general goal of the linac and the beamline setup is to deliver electron beams compatible with the basic experimental requirement of X-2: the generation of a driver-witness double-bunch structure with precisely shaped current profiles. As discussed in Section 2.2, the double-bunch structure is generated in the first dispersive section of the FLASHForward beamline and relies upon 1) the linearity of the longitudinal phase space of the beam and 2) particular constraints imposed on its Courant-Snyder parameters at the location of the three collimators. Additionally, the first (driver) bunch has to be capable of driving a nonlinear perturbation in the plasma, while the second (trailing) bunch capable of strongly beam loading the wakefield driven by the first. Taking these constraints into account and according to the operational parameter space available at FLASHForward (cf. Section 2.2), the beams delivered should have an energy of ~ 1 GeV, a linear chirp with a correlated rms energy spread ~ 0.25 %, a high peak current $\gtrsim 1$ kA with a linearly ramped and roughly triangular current profile, and must be matched to the design lattice optics. Typically, the full charge required to fulfil these conditions is ~ 1 nC.

Linac setup

After some preliminary routines to adjust the components of the photo cathode and the phases of the different RF modules, the focus of the linac setup is on matching the beam coming out of the photo injector to the lattice optics. For that, the compression settings must be such as to minimise the energy spread of the beam or, alternatively, correspond to the case in which the beam is *on-crest*—i.e., located at the maximum or crest of the RF signal, that incidentally coincides with an RF phase of $\phi_{RF} = 0$ deg, otherwise the measurement of the Courant-Snyder parameters would be spoiled by chromatic effects.

Firstly, the current of the main solenoid in the RF-Gun is adjusted to compensate for the space-charge forces that dramatically degrade the beam at these low energies (cf. Sections 1.1.3 and 2.1). After that, once the orbit has been tuned to avoid or compensate for kicks—arising from the beam propagating off-center through the quadrupoles or the SCRF modules, but also from RF couplers even when the beam is centered [Prat, 2009]—and the 1st and 2nd order lattice dispersion have been closed, an emittance measurement of the multiple-screen type (cf. Section 1.2.2) is performed downstream of the 1st bunch compressor, in a section called DBC2. It is worth mentioning here that this section is the only one in the linac that accommodates the accrual of the required phase advance to perform such a measurement. Therefore, no expense should be spared in the number of iterations needed

until the mismatch amplitude is close enough to $B_{mag} = 1.0$ (cf. Section 1.2).

At this point, the beam is sent to the LOLA-TDS (cf. Section 2.1) for its longitudinal phase space to be measured. An appropriate optic to resolve both energy and time has to be loaded and the transport of the beam has to be optimised down to the screen station located at the dispersive section of the LOLA-TDS beamline.

Figure 4.1 shows a longitudinal-phase-space measurement of the beam best suited to the requirements of the October 2020 ARD beam time campaign. The time resolution of the measurement is $R_\xi = 18$ µm (60 fs). The differences between the current profiles measured at the 1st and the 2nd zero crossings result from correlations in the streaking plane y–ξ at the location of the TDS (cf. Section 1.3.2). The reconstructed current profile has an appropriate shape, with a slowly and linearly ramped current at the back to enable the production of trapezoidal profiles after collimation and a peak current of ~ 1.25 kA. The correlation between energy deviation δ and longitudinal position ξ should be perfectly linear along the whole bunch, but the figure clearly shows two different slopes at the back and at the front.

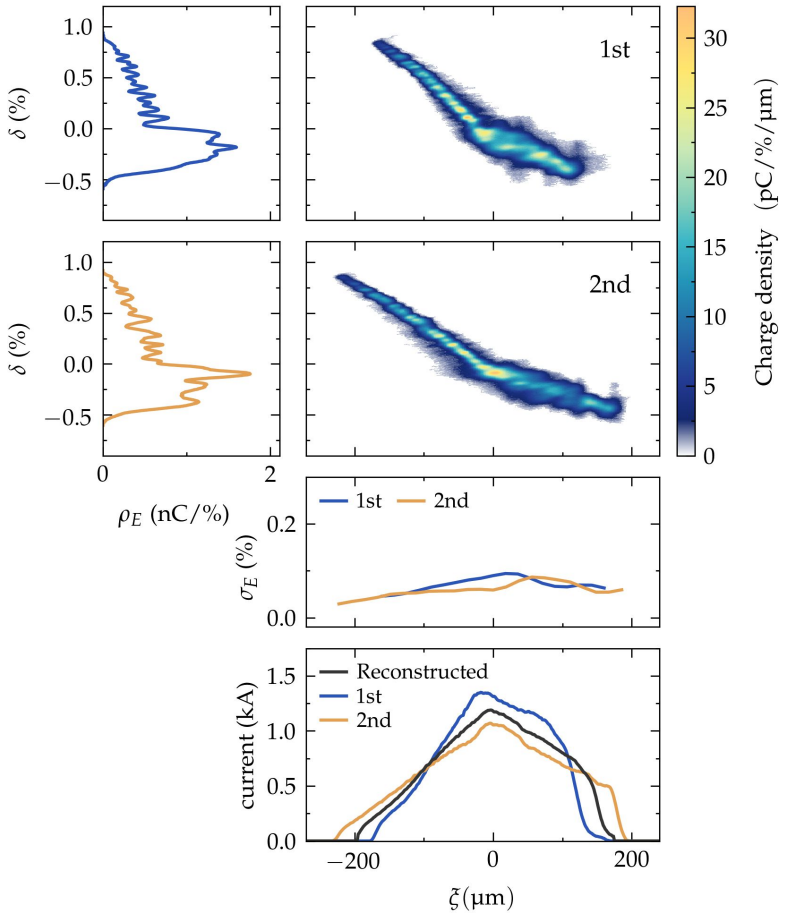

Figure 4.1.: Single-shot longitudinal-phase-space measurement with the LOLA-TDS performed at the SDUMP screen after the linac setup.

Parameter	Symbol	Units	Value	Measurement device/location
Charge	Q	pC	896 ± 1	Toroid 5DBC3
Mean energy	$\langle E \rangle$	MeV	1050	
Normalised emittance in x	$\epsilon_{n,x}$	mm mrad	0.84 ± 0.01	DBC2
Normalised emittance in y	$\epsilon_{n,y}$	mm mrad	0.76 ± 0.02	DBC2
Uncorrelated energy spread RMS	$\sigma_{u,E}$	%	0.07	LOLA-TDS
Correlated energy spread RMS	σ_E	%	0.32	LOLA-TDS
Bunch length RMS	σ_ξ	µm	83.3	LOLA-TDS

Table 4.1.: Key beam parameters measured during the linac setup, including the normalised emittance measured at DBC2 before compression and the energy spread and bunch length measured at the LOLA-TDS after compression.

This is the result of CSR effects (cf. Section 1.1.3) induced in the two bunch compressors, by which the beam progressively loses energy from the tail towards the center and gains energy from the center towards the head. These effects are also responsible for the slight increase in slice energy spread at the center and the head of the bunch. As discussed in Section 1.1.2, these changes in the longitudinal phase space are associated to centroid offsets in the dispersive plane x–s of the compressors and can develop as the beam propagates. This issue will be addressed later on, in Section 4.2.2. The beam also shows some high-frequency density modulations both in energy and space, which are the result of microbunching instabilities (cf. Section 1.1.3). The key parameters of the LOLA-TDS measurement, together with the projected emittance measured in DBC2 are summarised in Table 4.1.

Beamline setup

The next goal is to transport the beam to the straight section of the FLASHForward beamline with highest fidelity and throughput. The two basic routines to achieve this are precise orbit tuning and dispersion closure. Regarding the FLASH2 extraction, a description of the elements introducing dispersion and those foreseen to close it can be found in Section 2.1. Ideally, before switching to the FLASHForward extraction section, a new emittance measurement would be performed to make sure that the beam is still matched to the lattice optics. However, due to the large energy spread, such a measurement yields little reward.

High quality transport of the beam through the FLASHForward extraction and compression sections requires a certain degree of complexity. On the one hand, eventual misalignment between different hardware components mean that the center of the BPMs do not equate to an ideal orbit. On the other hand, the substantial energy spread in combination with the lattice dipersion results in large beams that compromise the accuracy of the BPM readings. Additionally, the steering of the beam through the differential pumping stage located at the end of the FLFCOMP section—with a thin pipe of 10 mm diameter—can quickly lead to

losses. Nevertheless, it is of course possible to overcome all these problems and proceed with the dispersion minimisation at the beginning of the FLFMAFF section, for which the same magnets as those used to straighten the beam can be used (cf. Section 3.3).

Once the beam is at the start of FLFMAFF, an emittance measurement of the multiple-quadrupole type (cf. Section 1.2) is performed at the 8FLFMAFF screen. As explained in Section 2.2, this requires the use of the head and tail energy collimators to reduce the charge of the beam and avoid the saturation of the scintillator screen, and to reduce the large energy spread that would otherwise spoil the measurement. By doing so, what is effectively measured is a relatively thin slice around the core of the bunch that is considered to well represent its overall beam parameters. The result of the measurement is then used to match the beam to the appropriate lattice optics that transport the beam through the PolariX-TDS and enable longitudinal-phase-space measurements at the 8FLFDUMP screen. This concludes the linac and beamline setup routines and constitutes the point of departure for tuning the beam parameters at the FLASHForward beamline with a higher lever of accuracy.

4.2. Beam preparation

The preparation of the beam for plasma interaction is performed in three consecutive steps with increasing degrees of refinement: longitudinal-phase-space optimisation, beam-tilt correction, and final-focus setup. Each of these steps will be addressed separately in the following subsections.

4.2.1. Longitudinal-phase-space optimisation

It has already been shown that collective effects, which degrade the beam quality, are triggered at several locations in the linac, and it is expected that, during the passage through the FLASH2 and FLASHForward extraction and compression sections, similar effects will develop. Therefore, it is important to monitor the beam at the end of the FLASHForward beamline, so that the compression settings can be re-adjusted by taking these additional effects into account. Besides that, the only way to directly assess the performance of the three collimators is to observe the beam downstream of their location.

With these considerations in mind, the first goal of the beam preparation is to optimise the longitudinal phase space of the beam by carrying out three interrelated routines in an iterative way: fine adjustment of the compression settings; re-assessment of the beam fitness for the double-bunch generation; and evaluation of the proper operation of each of the three collimators. This procedure is performed by setting new parameters manually and getting online feedback of their outcome. The key aspects of the final working point achieved are discussed below.

Figure 4.2 shows a longitudinal-phase-space measurement of the full bunch performed with the PolariX-TDS at the 8FLFDUMP screen. The time resolution is $R_\xi = 2.3\,\mu\text{m}$ (7.7 fs), an improvement of a factor ~ 7.5 compared to the LOLA-TDS measurement—clearly apparent with the increased definition of the bunch substructure resulting from microbunching instabilities. According to the lattice optic used and assuming, for instance, a normalised emittance of $\epsilon_{n,x} \sim 1\,\text{mm mrad}$, the energy resolution would be $R_\delta = 1.2 \cdot 10^{-4}$ and the induced energy spread $\sigma_{IES} = 4.7 \cdot 10^{-4}$, which means that the lowest measurable slice energy spread is limited by TDS-induced effects. The difference between the current profiles in the 1st and the 2nd zero crossings is almost negligible—not because correlations in the y–ξ plane do not exist, but because the impact of these correlations on the measured current profile is reduced compared to the LOLA-TDS case. This is actually a direct consequence of the increased time resolution, since the streaked beam size at the screen largely dominates over the natural beam size and hence over its correlations. On the other hand, the larger kick $\mathcal{K}$ achieved with the PolariX-TDS implies that the same correlations at the location of the TDS have a larger impact on the measured energy (cf. Equation 1.3.48), which explains the noticeable differences between the two zero crossings concerning the overall projected

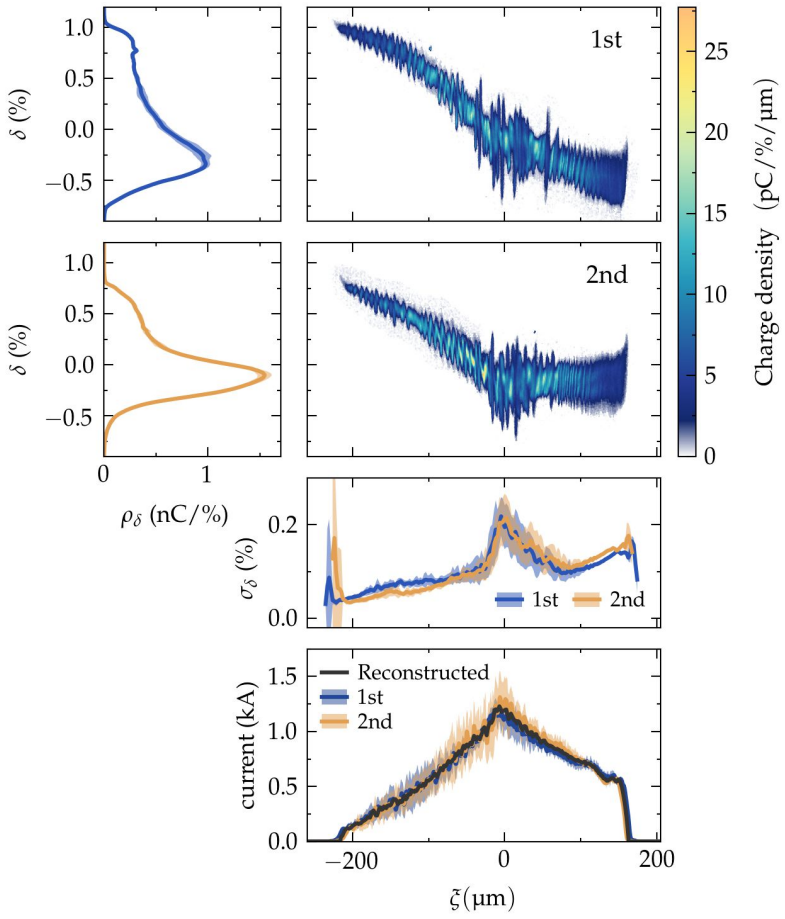

Figure 4.2.: PolariX-TDS longitudinal-phase-space measurement of unscraped bunches during the beam preparation. See text for details.

energy and the sliced mean energy of the bunch. Regarding the bunch correlations in the x–ξ plane, as soon as the centroid offsets are on the order of $\langle x \rangle \sim 100$ µm and $\langle x' \rangle \sim 0.5$ mrad, their impact on the measured sliced mean energy is expected to result in relative errors on the 1 % level (cf. Equation 3.4.4) and can therefore be ignored.

The influence of collective effects along the FLASH2 and FLASHForward extraction and compression sections is also clearly visible. As a result of CSR, the slice energy spread around the center of the bunch shows an increase from ~ 0.1 % to ~ 0.2 %, and, at the front part of the bunch, it increases steadily up to ~ 0.15 %. At the same time, also as a consequence of CSR, the energy of the center of the bunch decreases, while the front part has slightly increased. Despite these effects, the beam retains the features targeted during the setup—viz. linearly ramped current profile, peak current $\gtrsim 1$ kA and roughly linear chirp—and it is well suited to generate a double-bunch structure.

Figure 4.3 shows selected steps of scraper scans performed with each of the three collimators, corresponding to approximately the final positions used in the PolariX-TDS

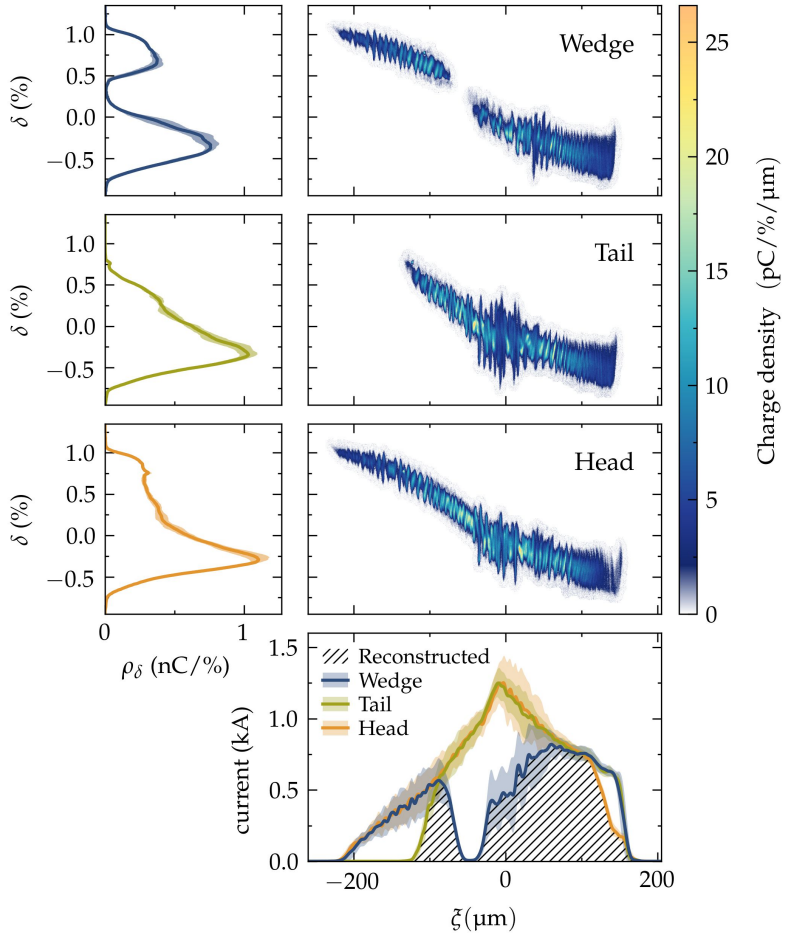

Figure 4.3.: Scraper scans performed at the 8FLFDUMP screen. See text for details.

campaign. The wedge scraper is able of completely removing the charge of a thin longitudinal slice at a location roughly ~ 50 µm behind the center of the bunch, clearly dividing it into two separate parts. The steepness of the cut is more pronounced at higher energies (with a fast rise of ~ 20 µm) than at lower energies (with its effect extending over a length of ~ 100 µm towards the head). This is due to both the higher energy spread around the center of the bunch and to the flatter chirp over the front half of the bunch. The operation of the tail scraper shows the same features as the higher energy part of the wedge scraper, with a short decay of ~ 20 µm. Moreover, the production of an almost perfectly trapezoidal current profile is possible by just shifting the tail scraper further back to higher energies. The operation of the head scraper is again affected by the large slice energy spread and the flatter chirp at the front of the bunch, but enables the modulation of the up-ramp shape of the driver, which has an impact on the transformer ratio and, consequently, on the overall efficiency of the acceleration process (cf. Section 1.4).

Table 4.2 summarises the compression settings obtained after the longitudinal-phase-space optimisation routines (black numbers) and indicates the changes applied respect to the ones obtained previously during the linac setup (coloured numbers). The most relevant change consists of a decrease of -0.40 m^{-1} in the chirp induced by ACC23, which is applied to achieve a larger compression after the second bunch compressor (BC3). The longitudinal-phase-space curvature resulting from shifting the beam further off-crest in ACC23 is subsequently linearised by increasing the curvature introduced by ACC1 and ACC39 by an amount of $+20$ m^{-2}. To assess the overall stability of these settings a large number of bunch-length measurements were performed with the BCM-4DBC3.2 located after the second bunch compressor BC3 (cf. Section 2.1). A histogram with the results is shown in Figure 4.4, which reveals large variations of up to about ± 15 % around a mean value of 94.4 µm (310 fs). Most probably the observed compression jitter is related to fluctuations in the power delivered by the ACC39 klystron, since this module is operated at 19.53 MV very close to its maximum value of 20 MV. However, this issue could not be investigated in detail and these fluctuations persisted during the whole campaign.

Parameter	Units	**ACC1**	**ACC39**	**ACC23**	**ACC45**	**ACC67**
Amplitude	MV	162.86 +0.45	19.53 +0.50	325.88 +2.33	306.96	290.00
Phase	deg	11.08 −0.10	8.14 −0.23	21.11 +1.09	0.00	0.00
Sum voltage	MeV	146.09		304.00		
Chirp	m^{-1}	−4.29 +0.03		−7.10 −0.40		
Curvature	m^{-2}	+71.90 +20.0				
Third derivative	m^{-3}	−6000				

Table 4.2.: Summary of the RF-phase and amplitude of the SCRF modules before and after fine tuning the compression settings. Black numbers: parameters obtained after the beam preparation at FLASHForward. Coloured numbers: changes applied with respect to the parameters obtained during the linac setup with the LOLA-TDS.

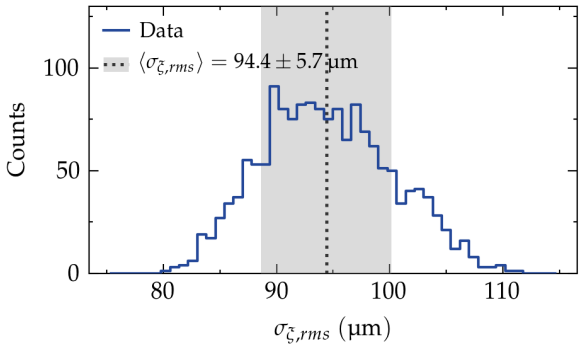

Figure 4.4.: Compression jitter measured at the BCM-4DBC3.2. See text for details.

4.2.2. Beam-tilt correction and R_{56} effects

The sources of beam tilts have been investigated in Section 1.1.2 and the importance of their control has been stressed in Section 1.4. Furthermore, strategies conceived at FLASHForward to overcome this issue have been discussed in detail in Section 3.4, and, with few exceptions regarding the lattice components used to that end, the procedure followed during the October 2020 ARD beam time developed along the same lines. In this section the focus is instead placed on some concomitant effects that the straightening has on the longitudinal beam dynamics: changes in R_{56} and their impact on the bunch length.

Table 4.3 at the end of Section 4.2.3 summarises the magnet currents of specific quadrupoles at different stages of the beam preparation. The quadrupoles are grouped vertically in two categories: those used to correct the tilts and those involved in the strong focusing optics. The first column corresponds to the values used during the longitudinal-phase-space measurements discussed in the previous subsection, and the second column to those achieved after the beam-tilt correction routines performed at the TDS screen 11FLFXTDS. As can be observed, the current of the quadrupole Q14FL2EXTR has been increased by 3.4 %, and that of Q23FLFEXTR decreased by -1.3 %. These quadrupoles have a large impact on the lattice dispersion and can easily modify the R_{56} accrued along the extraction and compression sections, which in turn affects the longitudinal dynamics of the beam. To investigate these effects experimentally, bunch length measurements performed at 8FLFDUMP before the straightening and at 11FLFXTDS after the straightening are correlated to the bunch length measured with the BCM-4DBC3.2 and are afterwards compared to each other. The results of this analysis are discussed in the following paragraph.

The two upper-left plots of Figure 4.5 show the RMS bunch length computed for a large number of images sampled at 8FLFDUMP and 11FLFXTDS. The differences between zero crossings are larger at 11FLFXTDS, since the lower streak achieved at this screen station translates into a larger impact of correlations on the current profile. In order to calculate the correlation between the real bunch length at the screens and the BCM readings, the data is analysed in discrete intervals along the range of BCM values common to all datasets.

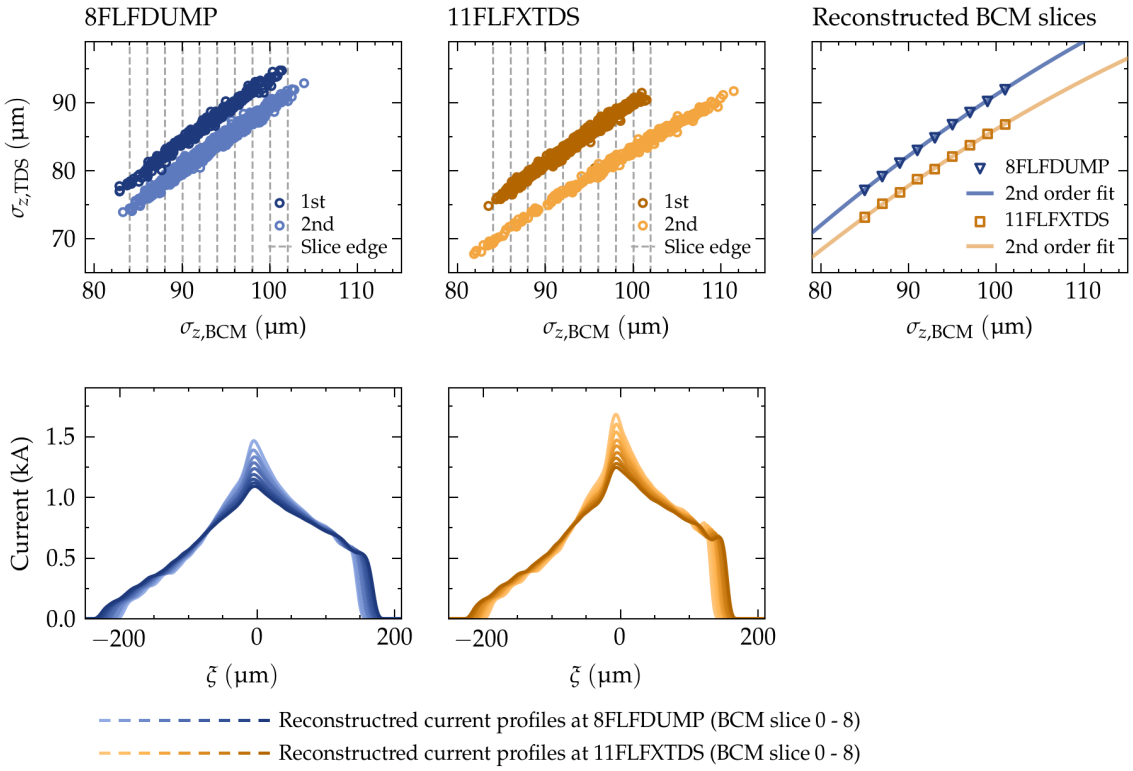

Figure 4.5.: Compression jitter and R56 effects at 8FLFDUMP and 11FLFXTDS. Beam-tilt-correction routines result in an additional compression associated to R56 variations in the FLASHForward extraction (FLFEXTR) and compression (FLFCOMP) sections.

For each interval, the current profiles at each zero crossing are averaged and the 2-point-tomography method is applied. The reconstructed current profiles are shown in the two bottom plots, in which the impact of bunch length variations on the peak current and on the longitudinal extent of the bunch can be observed. In the top-right plot, the RMS bunch length of the reconstructed profiles is correlated to the mean BCM values of each interval, resulting in a compression curve for each screen. The difference between these curves clearly indicates that an additional overall compression has been indeed introduced after the beam-tilt correction routines. Therefore, the analysis suggests that if precise knowledge of the longitudinal structure is required, the beam characterisation should be performed after the execution of all magnet lattice manipulations necessary to set up the experiment. The advantages and shortcomings of this approach are discussed in Section 4.4, which delves into the characterisation of the double-bunch structure.

4.2.3. Final-focus setup

For the beam to blow out the plasma electrons, its 3D particle density must be larger than that of the plasma. Simultaneously, in order to avoid emittance growth, at least the back of the bunch—which is the part that will be accelerated—has to be matched to the radially focusing

forces of the pure ion cavity left behind the driver. These are the two main constraints imposed on the final-focus setup, from which the target values for the beam matching are derived.

For plasma densities on the order of $n_0 \lesssim 10^{16}$ cm^{-3} and a beam with an energy of $E_b \sim 1$ GeV, the matched beta function has a value of $\beta_m \sim 5$ mm (cf. Section 1.4.2). In an example case with normalised emittances of $\epsilon_{x/y} \sim 5$ mm mrad in both planes, a length of $\sigma_\xi \sim 80$ µm (as measured at the PolariX-TDS), and a charge of $Q \sim 900$ pC, the peak density[2] will be $n_{b,\max} \sim 3.5 \cdot 10^{17}$ cm^{-3}. Of course some charge will be lost after collimation and not all slices will be equally demagnified by the focusing quadrupoles, but in a good approximation it can be assumed that such target β_m fulfils the condition $n_b/n_0 > 1$.

However, as the beam is focused symmetrically to such small beta functions, the required beam size in the plane for which the last quadrupole is defocusing grows very fast with increasing distance from the waist, at which point chromatic effects can become a concern. The strategy adopted to overcome this problem is to relax the plasma-matching condition in the plane that would require the largest beta function—this can be freely chosen, since the final focusing section provides enough degrees of freedom. This approach is supported by two considerations: on the one hand, it is expected that the emittance in x will be several times larger than that in y, since beam degradation due to collective effects is most severe in x—i.e., the bending plane of the bunch compressors and extraction sections. Therefore, even if the beta functions are not equal at the focus, the actual beamsize can be made to be roughly symmetric by increasing the beta function in the axis for which the emittance is smaller. This ensures that the whole beam fits into the wakefield structure. On the other hand, the main diagnostic devices located immediately downstream of the plasma cell are two dipole spectrometers designed to operate in the x-δ plane. Hence, for the general purpose of the X-2 experiment, as soon as the beam is matched to the plasma in x and the mismatch in y is moderate, no relevant information is lost.

There is yet another aspect to consider during the final-focus setup: in a radially focusing channel, an appropriate change in focusing strength on a length scale larger than β_m leads to an adiabatic matching of the beam (*c.f.* Section 1.1.2). This is eventually the case if the plasma density profile makes a gentle transition from the open ends of the plasma cell to the background pressure of the chamber [Garland et al., 2021]. However, precisely accounting for these effects in advance is highly complex. As such, the most efficient approach to adjust to the actual experimental conditions is to fine tune the final focus empirically while optimising the interaction of the beam with the plasma (cf. Section 4.3).

[2] For a 3D Gaussian without coupling between the three variables in space, the following formula for the peak particle density is used: $n_{b,\max} = \frac{Q}{e} \frac{1}{(2\pi)^{3/2}\sigma_x\sigma_y\sigma_\xi}$. The alternative formula $n_{b,\max} = \frac{I_{b,\max}}{ec} \frac{1}{2\pi\sigma_x\sigma_y}$ is also customarily used, for which a peak current of $I_{b,\max} \sim 1250$ A as measured with the PolariX-TDS results in the very similar value of $n_{b,\max} \sim 1.7 \cdot 10^{17}$ cm^{-3}.

Figure 4.6 shows the FLFMAFF section with the lattice optic used as a point of departure for the final-focus optimisation routines. The Courant-Snyder parameters at the start of the section were obtained from a quadrupole scan performed at the screen 8FLFMAFF, and were fed into a lattice optimisation routine with target values of $\beta_x \sim 10$ mm and $\beta_y \sim 25$ mm at the location of the cell. The obtained solution matched the target value in x exactly and generated a sufficient value in y to be used as a starting point for further optimisations. In the figure, several quadrupoles are highlighted. These are the ones best suited to control the size and location of the waist. In a good approximation, Q6 modulates the transverse symmetry of the beam, while Q12 and Q13 determine the final size and location of the waist in x and y correspondingly. Since independent control over x and y is not possible, additional quadrupoles can be used to compensate correlation effects.

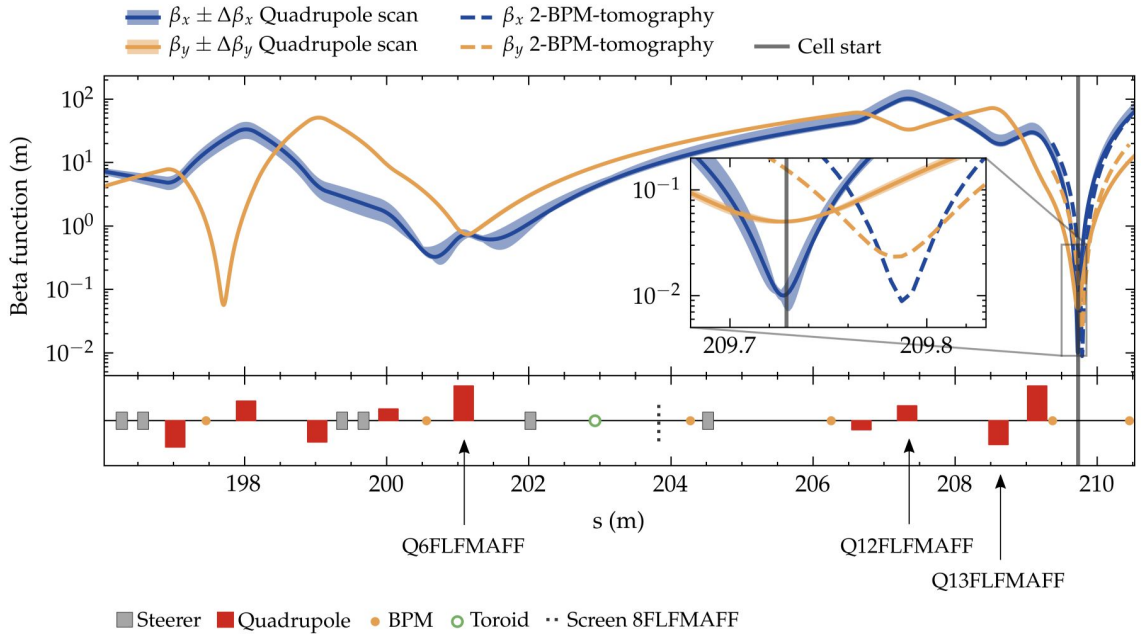

Figure 4.6.: Strong-focus optics at FLFMAFF. Solid lines: lattice optic obtained from optimisation routines fed with the Courant-Snyder parameters measured at the section start. Dashed lines: waist location measured with the two BPMs around the plasma chamber by means of a 2-BPM-tomography. The highlighted quadrupoles are those best suited to adjust the waist size and location.

By sampling the position of the beam centroid with two BPMs located upstream and downstream of the plasma cell, a jitter-based reconstruction of the transverse phase space of the beam can be obtained (cf. Section 1.2.2). The beta functions in x and y reconstructed by means of this technique are represented by dashed lines in Figure 4.6, which reveal a considerable disagreement with the theoretical optic derived from the quadrupole scan. Between the two measurements, the one most sensitive to errors is the quadrupole scan (cf. Section 1.2.2) and it is indeed expected that sources of errors are present—viz. mismatch between the beam and the lattice upstream of the measurement section, and dispersion due to a substantial correlated and uncorrelated energy spread. However, the quadrupole scan is

clearly enough to set up a strong focus at the approximate location of the plasma cell start, which is the condition required for the 2-BPM-tomography to be applicable. From this point on, the new technique—which, apart from being more robust, has the advantage of being non-invasive—is used to fine tune the beam parameters at the waist. At later stages of the optimisation process, the waist location and beam size are double checked with object plane scans performed at the ESPEC and LEMS stations.

After the large changes in focusing strength applied with the new strong-focus optics (cf. third column of Table 4.3), the beam straightness has to be reevaluated and eventually corrected. This is required because the much higher quadrupole field gradients magnify the kicks—and therefore the beam tilts—resulting from slight orbit offsets. Straightforward orbit-steering routines are sufficient to mitigate the largest effects, but the level of accuracy required by the experiment demands a more precise diagnostic and correction of the beam tilts at the exact location of the waist. The technique used to perform this *in-situ* straightening is an extension of the 2-BPM-tomography. Since the beam has an approximately linear chirp, it is possible to produce longitudinal beam slices by using the head and tail collimators located at the COMP section. The 2-BPM-tomography can then be applied to each individual slice, enabling the reconstruction of transverse/longitudinal correlations.

Figure 4.7 shows the final working point achieved after a few iterations of beam-tilt correction manipulations using this technique. The corresponding changes in magnet currents are summarised in the right-most column of Table 4.3. In the two upper-left plots of Figure 4.7, the reconstructed phase space of the centroids at the waist can be seen. The red ellipses enclose an area of $3\pi \cdot \epsilon_n$ (i.e., roughly 92 % of the particles), with estimated normalised emittances of $\epsilon_{n,x} \sim 5$ mm mrad and $\epsilon_{n,y} \sim 1$ mm mrad. The plots clearly show that the beam tilts have been successfully corrected in x and to a large extent also in y. The two upper plots in the center show the absolute trajectory of each slice, in which an overall angle of $\theta_x \lesssim 1$ mrad in x can be appreciated. Very likely, this is the result of slight misalignments between the two BPMs, but, even if the tilt were real, a small overall angle in the trajectory of the beam would not be detrimental to the sucess of the experiment, since the plasma wake is created by the beam itself. More important is the colinearity between different slices, since an angular misalignment between the driver—which determines the symmetry axis of the wakefield structure—and the trailing bunch will lead the latter to undergo transverse oscillations that can compromise the stability of the acceleration. For plasma capillaries longer than the one used in this campaign, angular misalignments in the driver itself can even trigger hosing instabilities, eventually leading to a beam breakup (cf. Section 1.4.2). The measured beta functions around the waist are shown in the two upper-right plots, in which the mean value of all slices at the waist is represented by a horizontal black dashed line at values $\beta_{s_0,x} = 10$ mm and $\beta_{s_0,y} = 25$ mm. Even if the plasma matching condition has been relaxed in y to avoid severe chromatic effects, some signs of mismatch between slices is still present, which translates into the apparent spread of longitudinal waist locations around $s \sim -0.2$ m. As expected, higher (lower) energies are focused further downstream (upstream). In the bottom-right plot the measured charge of each energy slice is shown, which, for a perfectly linear chirp, would correspond to the exact current profile of the beam.

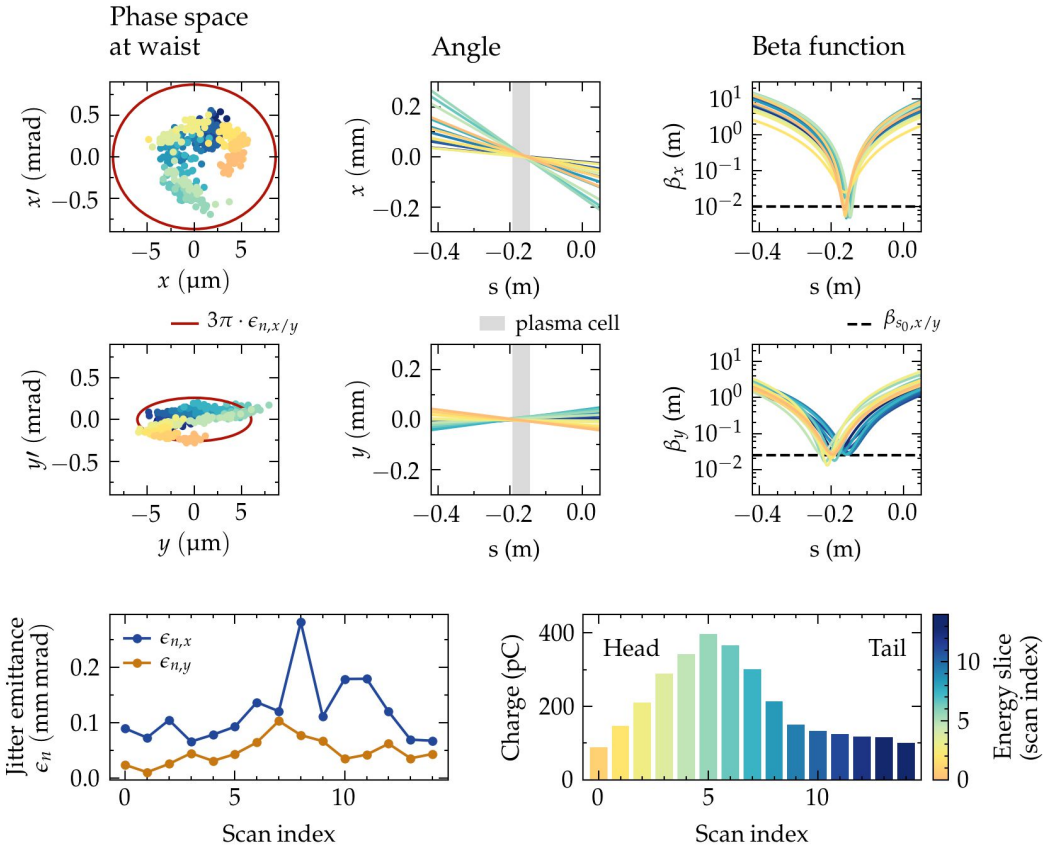

Figure 4.7.: 2-BPM-slice-tomography of the beam at the interaction point. See text for details.

However, as discussed in Section 4.2.1, the longitudinal phase space of the beam is distorted by collective effects and, as can be seen in Figure 4.2, the chirp is larger at the back than at the front part of the bunch. This explains the larger number of slices needed to scan the tail than that needed for the head. The jitter emittance in x and y is shown in the bottom-left plot. It has been discussed in Section 1.2.2 that this figure of merit is not equivalent to the real emittance of the beam, but that the relation between the two parameters is only approximate. Moreover, the relative error of the measurement is proportional to $1/\sqrt{N}$, where N is the number of samples. Since the number of slices analysed is considerable, the samples per slice was limited to $N = 25$, which corresponds to a relative error of at least 20 %—the error slightly increases after the removal of outliers. However, despite the error in the jitter emittance, it can be deduced from the measurement that the overall beam emittance in x is indeed several times larger than that in y.

The *in-situ* straightening concludes the beam-preparation routines, through which an appropriate working point has been established, which enables the plasma-interaction optimisation.

Magnet	Initial settings (A)	Straightening at 11FLFXTDS (A)	Strong focus (A)	Strong focus *in-situ* correction (A)
Q14FL2EXTR	−59.38	−61.38 +3.4 %		−62.38 +1.6 %
Q23FLFCOMP	84.69	83.59 −1.3 %		81.59 −2.4 %
Q1FLFMAFF	−0.75		−188.87 $> +10^3$ %	< 1.0 %
Q2FLFMAFF	−21.19		139.72 +762 %	< 1.0 %
Q3FLFMAFF	−30.94		−152.41 +395 %	< 1.0 %
Q4FLFMAFF	−0.01		83.25 $< -10^3$ %	< 1.0 %
Q6FLFMAFF	42.09		257.17 +515 %	< 1.0 %
Q11FLFMAFF	−49.98		−65.79 +32 %	< 1.0 %
Q12FLFMAFF	18.29		106.65 +502 %	110.14 +3.3 %
Q13FLFMAFF	44.32		−171.29 −495 %	−175.10 +2.2 %
Q14FLFMAFF	−16.97		257.17 $< -10^3$ %	260.73 +1.4 %

Table 4.3.: Magnet currents at different stages of the beam preparation.

4.3. Beam-plasma-interaction optimisation

Direct observation of the interaction between the beam and the plasma provides the most reliable experimental indicators to assess the quality of the acceleration process. In this section, general considerations about these indicators and some experimental techniques derived from them are provided. After that, the optimisation performed in October 2020 to enable time-resolved measurements of plasma-accelerated electrons with the PolariX-TDS is described.

4.3.1. General approach

The optimisation of the beam-plasma interaction is an intrinsically complex problem, because the number of variables involved in the process is quite large and a systematic inspection of the full parameter space is not practicable[3]. Moreover, several merit functions describing different aspects of the acceleration process need to be evaluated simultaneously. Nonetheless, in a first approximation, the available variables can be added consecutively according to their impact on the overall acceleration given the initial experimental conditions. This enables a progressive definition of a bounded region in the operational parameter space in which an optimal solution is expected to be found. Once all experimental variables are included, the

[3] If, for instance, only three variables are to be inspected, a multidimensional scan with 10 steps for each parameter, 25 samples for each step, and a sampling rate of 5 Hz would require 1 h and 20 min—without taking into account the time needed to change settings between scan points.

system reveals its true complexity and an intricate optimisation process is required to find the working point best suited to the requirements of the experiment. In the previous sections, the beam parameters have been independently optimised. The next step is to identify an appropriate range of plasma densities that enables the optimisation of the interaction with the beam.

Plasma-density scan

The usual procedure to identify the range of densities best suited to the experiment is to perform a plasma-density scan, by which the whole bunch is sent through the plasma while the density is progressively increased. The energy spectrum of the interacted beam is then monitored with the dipole spectrometer located immediately after the plasma chamber (ESPEC, cf. Section 2.2), providing information about key features of the wakefield structure. The physics that dominates the process is as follows: the longitudinal size of the bubble depends on the plasma density as $\sim 1/\sqrt{n_0}$ (*cf.* Equation 1.4.4); therefore, as the density increases, the bubble shrinks and the bunch—which has a fixed length—interacts with an increasingly larger range of decelerating and accelerating phases of the wakefield, until it eventually reaches the back of the bubble and starts to see the second wakefield structure. To some extent, and in analogy to the transverse matching discussed in the previous section, the nature of this scan can be regarded as a longitudinal matching of the plasma wakefield to the length scale of the beam.

Figure 4.8 shows such a plasma-density scan performed under experimental conditions very similar to that of the ARD campaign discussed in this chapter. In the top waterfall-like plot, the energy spectrum of the interacted beam is correlated to the discharge-trigger-time delay—i.e., the plasma density. Full images of the bunch at six different points along the scan are shown in the bottom plots. The top-left one corresponds to the non-interacted bunch—i.e., the bunch arrives before plasma generation (right-most samples in the waterfall top plot). The mean energy of the non-interacted beam is represented by a grey vertical line at $\langle E \rangle_{\text{off}} = 1055$ MeV, which is also included in the other five plots for reference. As a result of the negative chirp imprinted in the linac, the head has lower energy, while the tail has higher energy. In image a) the bunch has a size of roughly one half of the structure and experiences an overall deceleration. In image b) part of the tail is already found beyond the zero crossing of the E_z field and starts to experience a slight acceleration. Moreover it can be observed that the lowest energy reached has decreased, indicating a larger field amplitude expected from a higher density. Image c), highlighted in green, corresponds to a working point well suited to experimentation with externally injected beams. This is the case for which the largest amount of charge is accelerated in the most stable way due to the transverse focusing forces of the wakefield. The tail reaches the phase with the highest accelerating field and experiences a maximum energy gain of $\Delta E = 45$ MeV. At this point, for an effective plasma length of 50 mm, the maximum field gradient would be $E_z = 0.9$ GV/m. In image d) the bunch fills the whole length of the wakefield structure and the last particles of the tail

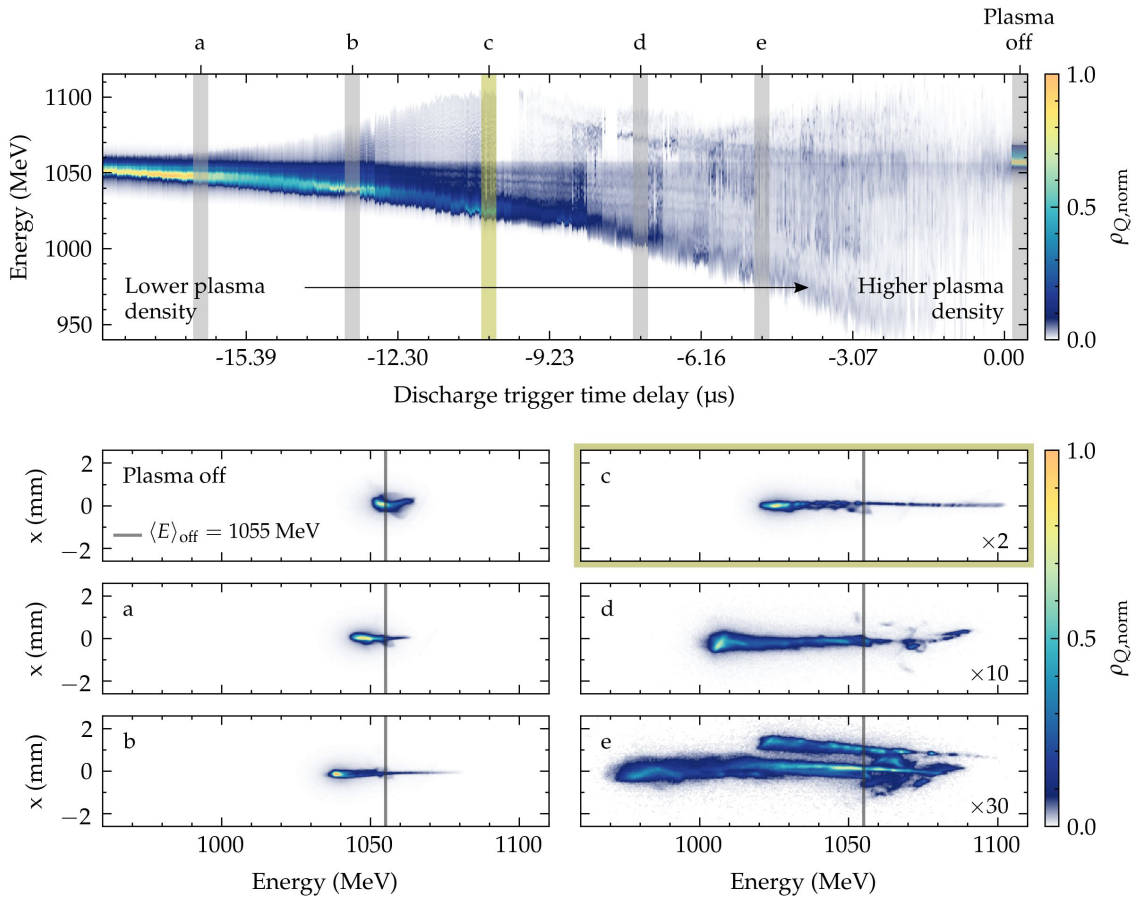

Figure 4.8.: Discharge-trigger-time-delay scan performed to adjust the range of plasma densities suitable for PWFA experiments with externally injected bunches. See text for details.

do not experience any longitudinal acceleration. However, at the very back of the bubble the radial fields are already defocusing and slight centroid offsets translate into a transverse deflection that expels the particles outside of the structure. This constitutes a charge-loss mechanism that has to be addressed when optimising the beam-plasma interaction. Finally, in figure e), a considerable fraction of the tail is already found in the decelerating phase of the second wakefield structure and the particles start to lose energy. In the October 2020 ARD campaign, a working point close to that of image c) was established to proceed with the subsequent optimisation routines.

Wedge-scraper multidimensional scans

The next step in the beam-plasma-interaction optimisation is to insert the wedge scraper to produce a double-bunch structure. With that, two new variables are available: the distance and the charge ratio between the two bunches. These two variables lead to a considerable increase in the complexity of the optimisation procedure, since they simultaneously affect

several properties of the wakefield structure that determine the quality of the acceleration process: the transformer ratio, the amount of beam loading and the energy-transfer efficiency (cf. Section 1.4.2). On the other hand, as seen in the plasma-density scan discussed above, the relation between the length of the bunch and the length of the wakefield structure—which is determined by the plasma density—has a large impact on the acceleration process. Consequently, the optimisation of the wedge-scraper position constitutes in fact a three-dimensional problem that includes the plasma density, which has to be reassessed for each combination of distance and charge ratio between the two bunches.

An approach to solve this problem is the one followed in the experimental campaign that led to the demonstration of energy-spread preservation and high energy-transfer efficiency in early 2021 [Lindstrøm, Garland et al., 2021]. In that campaign, a multidimensional scan including the three variables was performed and appropriate observables were combined to construct a *wakefield-optimisation parameter* Ω with which the results were evaluated. This new figure of merit equals the energy-spread-to-gain ratio σ_δ, divided by the transformer ratio T_R and the energy-transfer efficiency η (i.e., $\Omega = \frac{\sigma_\delta}{T_R \eta}$) and is required due to the fact that the optima of the individual parameters cannot be found simultaneously. As successfully demonstrated in that campaign, the minimisation of this quantity results in the best compromise between the minimisation of the energy-spread-to-gain ratio σ_δ, and the maximisation of the energy-transfer efficiency η and the transformer ratio T_R.

Such a multidimensional scan, however, is extremely time consuming and faster approaches are required at intermediate stages of the optimisation. By manually inspecting the behavior of the system in response to variations of each of the variables, it is possible to establish a hierarchy of variables. The variable with the least impact on the system is then fixed to a specific value and a 2D scan is performed with the two remaining variables. Once an optimal point is found, the variable previously fixed is adjusted separately to refine the solution. Typically, at FLASHForward, a 2D scan of transverse wedge-scraper position and plasma density constitutes a good choice for intermediate optimisations. This is justified by the poor performance of the wedge-scraper at the lower-energy side—which, as seen in Section 4.2.1, results in a considerably unsharp cut—compared with the good performance at the higher-energy side.

Beam matching

Section 4.2.3 describes the procedure to set a beam focus of a specific size at the location of the plasma cell. Typically, the precision with which this is achieved is enough to drive a nonlinear wakefield, resulting in a convenient point of departure to proceed with the optimisation of the beam-plasma interaction. In these conditions, the two routines described above can be applied to improve the interaction regardless of the transverse beamsize of either the driver or the trailing bunch. Once a sufficiently stable working point is reached, the impact of the beamsize of each bunch has to be considered separately. As for the driver, as soon as its charge density is enough to produce a nonlinear perturbation in the plasma,

the development of the wakefield structure is relatively insensitive to its precise transverse beamsize. This is because the transverse gradients keep the beam focused around the axis and the sheath of electrons at the boundary of the cavity become only sensitive to the current profile. Concerning the trailing bunch the situation is quite different. On the one hand, as the plasma-sheath electrons return to the axis, the cavity becomes narrower and the fraction of trailing-bunch electrons effectively captured in the structure can be reduced if its transverse beamsize is too large or if its centroid is transversely offset. On the other hand, if the beta function is not well matched to the transverse focusing forces of the ion column left behind the driver, the quality of the accelerated beam can be degraded—e.g., as a result of emittance growth (cf. Section 1.4.2).

Adequate observables that allow the assessment of the quality of the matching are the amount of charge effectively accelerated or the beamsize and divergence of the accelerated bunch—which can be measured by performing an object plane scan at the ESPEC. In order to match the trailing bunch to the plasma, different strategies can be followed. One option is to scan the location of the capillary along the direction of the beam propagation. Another option is to adjust the final focusing to change the size and location of the beam waist—which eventually involves performing a 2D scan with two of the last four quadrupoles before the plasma chamber.

Beam-tilt correction

Due to the spatial constraints imposed at the back of the wakefield structure, the presence of angular misalignments can lead to a reduction of the amount of charge being coupled into the plasma as well as to an instability of the acceleration process. Its compensation is achieved by means of beam-tilt-correction routines as those described in previous sections, which can be assessed by observation of the amount of charge and the exit angle of the accelerated bunch at the ESPEC.

Tail-scraper scan and charge coupling

Preserving the beam quality requires the full charge of the trailing bunch to be effectively coupled and accelerated in the wakefield[4]. As mentioned before, the charge-coupling can be compromised if the tail of the bunch interacts with the back of the wakefield structure, which can be partially overcome by adjusting the transverse beamsize of the bunch and/or by correcting its angular misalignments. Additionally, at the very back of the structure, the transverse fields are defocusing and acceleration becomes unstable—cf. image d) in Figure 4.8. The extent of the region around the bunch tail affected by these instabilities is increased by plasma-density variations that lead to changes in the size of the wakefield structure. The

[4]In the following, it is assumed that the trailing bunch does not exceed the longitudinal extent of the wakefield structure.

strategy to cope with these charge-loss mechanisms is to scan the tail-scraper position to progressively reduce the length of the trailing bunch until all its particles are found inside the usable and stable volume of the wakefield. For each scraper position, the total charge sent into the plasma is compared to the accelerated one. As long as some charge is lost during the beam-plasma interaction, a reduction of the incoming charge does not affect the amount of accelerated charge, which stays approximately constant. Once the incoming charge equals that of the accelerated bunch, a further decrease affects both the incoming and the accelerated bunches identically. A trail-scraper scan performed during the October 2020 ARD shifts is discussed in Section 4.3.2.

Advanced optimisation approaches

In the optimisation routines described above, the behavior of the beam-plasma interaction is mainly determined by a relatively small number of parameters whose variation, to a large extent, can be unambiguously correlated to a change in the properties of the accelerated bunch. This can be observed at early stages of optimisation, when specific parameters are clearly suboptimal and large changes are required to overcome a minimal threshold of acceleration quality and stability. The situation changes dramatically when a bounded region in the parameter space is found, for which any choice of parameters leads to an acceptable result. At this point, complex operations are required to achieve small improvements, because all parameters tend to entangle each other and a slight modification of a single one have an effect on one or several of the others—and very often their correlation is not obvious. Similar to the strategy followed in [Lindstrøm, Garland et al., 2021], at this advanced stage of optimisation multidimensional scans have to be performed and their evaluation requires the use of appropriate merit functions. A general solution to the problem is not available and each experimental campaign has to be considered individually. In the following section, the optimisation approach followed during the PolariX-TDS campaing in October 2020 is described.

4.3.2. Optimisation for PolariX-TDS measurements

To enable the measurement of plasma-accelerated bunches with the PolariX-TDS, the beam-plasma interaction has to be optimised to achieve specific beam parameters. Low-divergence beams are less prone to cause transmission problems and allow for the use of lattice optics with moderate magnet strengths that reduce chromatic effects. Thus, one of the goals of the optimisation is to minimise the divergence of the accelerated bunch. Even in this case, however, the first quadrupoles after the plasma cell require relatively strong gradients in order to capture the beam. If the accelerated bunch exits the plasma cell with an angle, these quadrupoles introduce kicks that translate into large transverse offsets after beam propagation, which, in turn, can lead to a transmission loss. Therefore, the exit angle of the accelerated bunch has to be minimised as well. Regarding the stability of the plasma

acceleration, the inherently strong gradients produced in plasmas tend to amplify any source of jitter associated to the incoming double-bunch structure. Moreover, due to the large divergence of the beam and the long distance between the plasma cell and the PolariX-TDS, a slight variation in energy gain translates into a considerable mismatch between the beam envelope and the lattice optics, which results in suboptimal TDS measurements. On account of that, a working point with a moderate energy gain is preferred, resulting in weaker plasma wakefields but an increased stability of the acceleration process.

The plasma cell used during the October 2020 ARD campaign consists of a discharge capillary with a length of 50 mm and a diameter of 1.5 mm, milled from two blocks of sapphire. The capillary is filled with Argon at a backing pressure of 10.7 mbar through two gas inlets separated by 34.2 mm, and is discharged by an electrical pulse with a voltage of 25 kV, a current of 500 A and a length of 400 ns. With this setup, the upper bound of achievable densities is on the order of $n_0 \sim 10^{17}$ cm^{-3} and, as discussed in the final-focus setup, the peak particle density of the beam is $n_b \gtrsim 10^{17}$ cm^{-3}. Therefore, it can be assumed that the beam drives a nonlinear wakefield over the whole range of plasma densities, since the condition $n_b/n_0 \gtrsim 1$ always holds.

Prior to optimising the beam-plasma interaction for the PolariX-TDS measurements, a working point for a different experiment was established that resulted in a bunch with $\sim$ 80 pC of charge and an energy gain of $\sim$ 50 MeV. The quality of the incoming beam was not perfectly preserved after acceleration, but a substantial degree of beam-loading was achieved, enabling the minimisation of the energy spread. Using these conditions as a starting point, the first optimisation step consisted of reducing the angle and divergence of the bunch at the plasma exit. This was achieved by scanning the location of the cell along the longitudinal coordinate z and measuring the projected emittance of the accelerated bunch at the ESPEC. The result of the scan is shown in Figure 4.9 (wedge position A). For each scan step, the size and divergence of the incoming beam at the location of the cell is effectively varied. This affects the evolution of the transverse beam envelope inside the

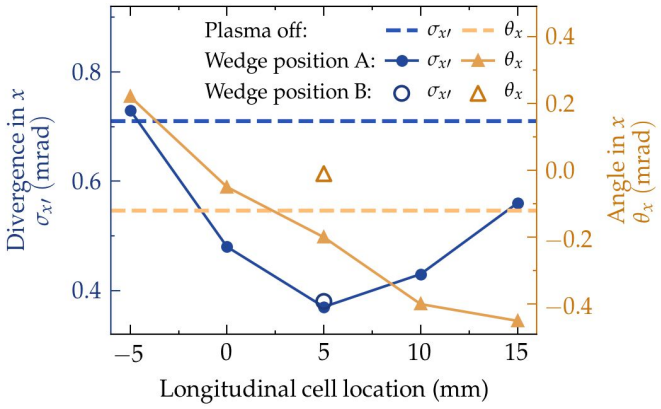

Figure 4.9.: Plasma-cell-location scan performed to optimise the divergence and exit angle of the trailing bunch for PolariX-TDS measurements. See text for details.

plasma and, consequently, the divergence with which the bunch exits the cell (blue-filled circles). At the same time, due to the presence of residual angular misalignments between the driver and trailing bunches, a change in propagation length before plasma interaction translates into a change in the initial transverse coordinate of the trailing bunch relative to the wakefield structure. Therefore, for each scan step, the trajectory of the trailing-bunch centroid inside the plasma is slightly different and the angle with which it exits the cell is modified (orange-filled triangles). The best beam parameters are obtained when the cell is located 5 mm downstream of its reference position (cf. Section 2.2).

In a subsequent optimisation step, the plasma density was slightly decreased to reduce the energy gain and improve the overall stability of the acceleration process. After this correction, a drop in trailing-bunch charge was noticed. To compensate for that the wedge-scraper was pulled out 10000 motor steps, reducing the distance between the driver and trailing bunches and simultaneously increasing their charge. After a slight adjustment of the orbit and the tilt of the incoming beams, the divergence and exit angle of the accelerated bunch were measured again (wedge position B in Figure 4.9). The divergence ($\sigma_{x'} \sim 0.4$ mrad) and angle ($\theta_x \sim 0$ mrad) obtained are considered to be a good working point for the PolariX-TDS measurements. Additionally, a moderate energy spread of ~ 0.1 % observed at the ESPEC, indicates that a substantial degree of beam-loading is still present—which is a requirement for the study goals of the PolariX-TDS campaign.

Finally, to conclude the beam-plasma interaction optimisation, a tail-scraper scan was performed to optimise the charge coupled into the plasma wakefield as described in Section 4.3.1. The results are shown in Figure 4.10. The bottom plot contains the total charge of the incoming double-bunch structure directly measured with a toroid (TOR7FLFMAFF) located upstream of the experimental chamber. The upper and lower bounds corresponding to the two extreme positions of the tail scraper—i.e., completely retracted or completely blocking the trailing bunch—are indicated with the two grey horizontal solid lines. The top plot shows the total CCD counts of the accelerated bunch computed from energy-spectrum images taken at the ESPEC screen. For the range of charges covered in the scan, the CCD counts are expected to correlate linearly with the amount of charge measured at the screen, which was considered to provide enough information to evaluate the degree of charge-coupling. The black vertical dotted line in both plots indicates the tail-scraper position for which 100 % of the charge was deemed to be effectively accelerated, which according to the toroid readings corresponds to ~ 50 pC. In the top plot, at the left side of the dotted line, the large error bars indicate the unstable nature of the charge-loss mechanisms described in Section 4.3.1. At the location of optimal charge-coupling, the accelerated bunch is still affected by charge variations exceeding those of the incoming bunch, which signal a residual instability of the acceleration process. Such instabilities could be an issue in experiments aiming at beam-quality preservation, but should not represent a major drawback to the PolariX-TDS campaign and its goals.

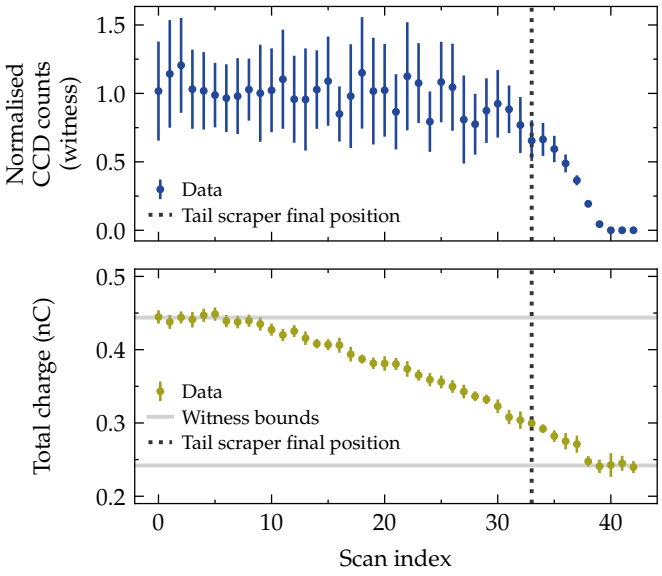

Figure 4.10.: Tail-scraper scan performed to optimise the trailing-bunch charge coupled into the wakefield structure. Top: normalised CCD counts of witness bunches. Bottom: total double-bunch-structure charge.

4.4. Uninteracted double-bunch characterisation

The main focus of this section is the characterisation of the double-bunch structure resulting from the beam-plasma-interaction optimisation discussed so far and comprises two sets of measurements. First, the projected parameters are measured with the two dipole spectrometers installed at the diagnostic section located immediately downstream of the plasma chamber—the ESPEC and the LEMS. After that, time-resolved measurements are carried out with the PolariX-TDS. Even if the focus of this thesis is on the second type of measurements, all the data presented here is required to obtain an appropriate picture of the beam dynamics that enables a basic understanding of the plasma-acceleration process discussed in the next chapter.

4.4.1. Projected beam parameters

The absolute-energy spectrum of the double-bunch structure is measured at the ESPEC. As discussed in Section 2.2, for highly divergent beams with a large energy spread, the imaging system is set up to establish a point-to-point correspondence between the location of the beam waist and the screen, and the focusing energy is subsequently scanned over the range of energies of the beam. Although the total energy spread of the noninteracted double-bunch structure is low enough to be resolved in a single imaging setting, an imaging-energy scan is performed to empirically identify the exact configuration that optimises the resolution, since the overall system is very sensitive to small errors in the magnetic lattice.

Such a scan is shown in Figure 4.11. The focusing energy is varied between 1030 MeV and 1070 MeV in 9 steps and 10 samples per step are taken. On the left, the spectrum of all the samples is plotted. To a good approximation, the best focusing is achieved for both bunches simultaneously at the second-to-last scan step, which is highlighted with a shaded area in the plot. The average spectrum of the highlighted samples is shown on the right plot, in which an excessive overlap between the spectra of driver and trailing bunch hints a problem in the resolution of the measurement. An offline analysis reveals that the object plane used in the point-to-point imaging is erroneously located far away from the expected beam waist, resulting in a disproportionate contribution of the beam divergence that translates into an energy-spread broadening. Additionally, the nominal focusing energies used in the scan (1030–1070 MeV) do not correspond to the actual object plane of the measurement—which has been corrected in labels of the horizontal-axis of the left plot. Despite these problems, the mean energy can still be calculated, since the calibration of the spectrometer does not depend on the quadrupole imaging. The mean energy of each bunch is estimated by fitting a two-Gaussian function to the spectrum, which is displayed in the plot as a dashed (driver) and a dotted (witness) vertical line.

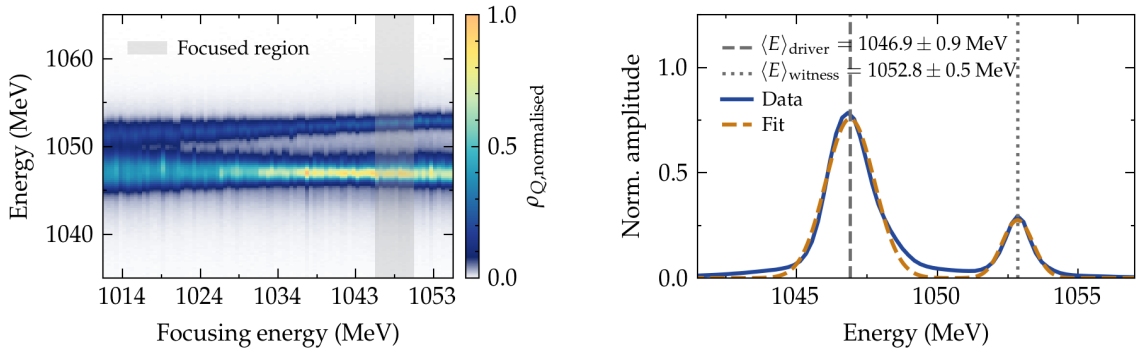

Figure 4.11.: Imaging-energy scan of the non-interacted double-bunch structure at the ESPEC. Left: full-scan energy spectra. Right: averaged energy spectrum at the focused region with two-Gaussian fit.

The projected emittance in x is measured at the LEMS, which provides the required resolution to measure the µm-level beamsizes typically used in PWFA experiments at FLASHForward (cf. Section 2.2). The measurement is performed by setting the quadrupoles to point-to-point imaging and scanning the object plane around the beam waist (cf. Section 1.2.2). The imaging energy is set to the mean value of the double-bunch structure. At the screen, the two bunches are clearly distinguishable, which enables a separate analysis of each. The results of the measurement are shown in Figure 4.12. On the left, the real beamsize at each object plane together with the fitted beamsize evolution is plotted for both the driver (blue) and the tailing bunch (orange). As expected from CSR effects, the driver emittance is larger than that of the trailing bunch, with values of $\epsilon_{x,d} = 7.3 \pm 0.1$ mm mrad and $\epsilon_{x,d} = 4.8 \pm 0.1$ mm mrad, respectively. The waist location of the witness coincides almost perfectly with the entrance of the plasma cell—which is represented by a grey area. Since the driver has a lower energy, its waist is located further upstream. On the right, the centroids and a fit to their drift-like trajectory are shown. Clearly, the paths of driver and trailing

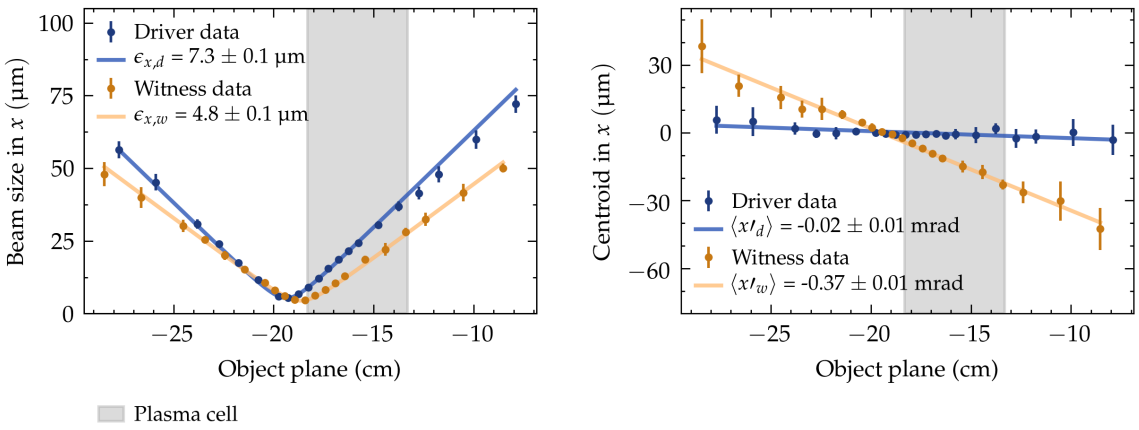

Figure 4.12.: Object-plane scan of the non-interacted double-bunch structure at the LEMS performed to measure the projected emittance of driver and witness bunch. Left: beamsize evolution and calculated projected emittance. Right: centroid trajectories and calculated angle.

bunch are not colinear. Once plasma-interaction is enabled, this will lead the trailing bunch to undergo transverse oscillations around the symmetry axis of the wakefield structure and, consequently, to an eventual emittance growth. However, this effect is expected to be of second order in the short 50-mm-long cell used for this experiment, and therefore its correction is not considered imperative.

Since the energy resolution at the ESPEC is affected by imaging errors, the energy spread is retrieved from the object-plane scan performed at the LEMS. The absolute energy measured at this station is typically affected by angular misalignments (cf. Section 2.2) and only its relative value is considered. The left plot in Figure 4.13 shows the energy spectrum of the bunches, in which a drift in mean energy is clearly observed. This is actually an orbit drift that results from the combined effect of 1) transverse offsets between the beam and the imaging quadrupoles and 2) changes in focusing strength. The right plot shows the RMS energy spread of both bunches computed for all shots, together with their

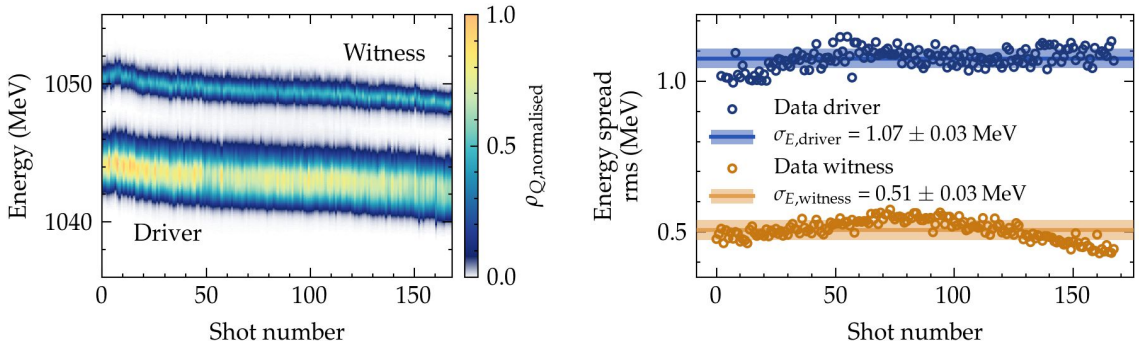

Figure 4.13.: Energy spread of the noninteracted double-bunch structure measured at the LEMS—data taken from the object-plane scan discussed previously. Left: full-scan energy spectra. Right: calculated RMS energy spread of driver and witness bunch.

Parameter	Symbol	Units	**Driver**	**Witness**	Measurement device/location
Charge	Q	pC	238 ± 13	55 ± 6	Tor. 7FLFMAFF
Mean energy	$\langle E \rangle$	MeV	1046.9 ± 0.9	1052.8 ± 0.5	ESPEC
Energy spread, rms	$\sigma_{E,\,rms}$	MeV	1.07 ± 0.04	0.51 ± 0.03	LEMS
Norm. emittance in x	ϵ_x	mm mrad	7.3 ± 0.1	4.8 ± 0.1	LEMS
Beta function x at waist	$\beta_{x,\,s_0}$	mm	8.1 ± 0.2	8.7 ± 0.2	LEMS
Beam size x at waist	$\sigma_{x,\,s_0}$	µm	5.36 ± 0.05	4.48 ± 0.06	LEMS
Divergence x	$\sigma_{x\prime}$	mrad	0.67 ± 0.01	0.52 ± 0.01	LEMS
Angle x	$\langle x\prime \rangle$	mrad	-0.03 ± 0.01	-0.38 ± 0.01	LEMS

Table 4.4.: Projected beam parameters measured at the ESPEC and the LEMS.

mean value and standard deviation. In this case, the orbit drift resulting from transverse misalignments translates into beam-dispersion modulations, which explain the observed changes in energy-spread along the scan. On account of that, it is assumed that an upper bound of the real energy spread of the bunch corresponds to the lowest value obtained in this measurement—which in the case of the trailing bunch is that of the last 10 samples corresponding to the last scan steps, with $\sigma_{E,witness} = 0.45 \pm 0.02$ MeV. This assumption is relevant to the estimation of TDS-induced effects discussed later on.

4.4.2. Time-resolved measurements

This section delves into the longitudinal-phase-space and slice-emittance measurements performed with the PolariX-TDS at the 8FLFDUMP and 11FLFXTDS screen stations. Due to the long distance between the interaction point and the screen stations, chromatic effects are more pronouced than in the ESPEC and the LEMS. For this reason, the longitudinal-phase-space of the bunches is measured first and the resulting sliced energy is taken into account in the subsequent calculation of the slice emittance. The preparation of all measurements—most of them consisting of a quadrupole scan—concentrates on the optimisation of the beam orbit in order to 1) minimise the drifts resulting from variations in quadrupole strength and 2) minimise charge losses. However, due to the inherent difficulty of transporting such divergent beams and to time constraints during the experimental campaign, the desired degree of optimisation could not always be achieved.

Longitudinal phase space

Similar to the projected-energy-spread measurement performed at the ESPEC, the longitudinal-phase-space measurement at the 8FLFDUMP screen is performed by using point-to-point imaging and scanning the focusing energy. This is shown in Figure 4.14, where the spectra of

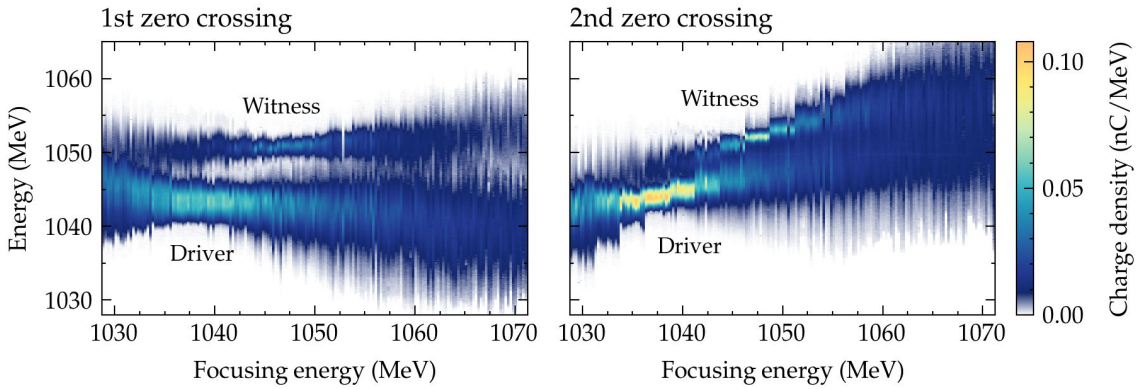

Figure 4.14.: Imaging-energy scan of the non-interacted double-bunch structure at the 8FLFDUMP screen measured with the PolariX-TDS at the 1st (left) and 2nd (right) zero crossings.

the double-bunch structure measured at the 1st (left) and 2nd (right) zero crossings can be seen. Compared to the scan performed at the ESPEC, the effect of an increased chromaticity can be appreciated in the larger separation between the locations where driver and trailing bunch are focused. Despite the effort made to optimise the orbit, the measurements are still affected by drifts, which are different at each zero crossing due to the horizontal kick induced by the PolariX-TDS (cf. Section 3.3). The focusing energies (horizontal axis labels) at which the bunches are properly imaged are not consistent with the true mean energies measured at the ESPEC. This, again, is probably related to a slight misalignment between the object plane and the beam waist. However, for these measurements, only the relative energy is of interest and this issue can be ignored. The projected energy spread is significantly larger at the 1st zero crossing than at the 2nd, which is reflected in the differences in charge density represented by the color scale. In the following, these differences are analysed further from full longitudinal-phase-space images.

To reconstruct the longitudinal phase space from the imaging-energy scan, two images in which either the driver or the trailing bunch is focused need to be combined. The relative position between the two images is adjusted according to the charge-density profile projected into x and y, from which the horizontal and vertical distances between the center of mass of driver and trailing bunch is calculated. In the resulting image, only the ROI around the focused bunch is used. The reconstruction obtained in this way is shown in Figure 4.15, in which the good longitudinal resolution achieved in the measurement, with $R_\xi = 2.1\ \mu\text{m}\ (7.0\ \text{fs})$, provides a considerable level of detail. The most striking differences between the two zero crossings are those seen in the energy plane and are closely related to the longitudinal-phase-space distortions already investigated in Section 3.4. The modulations in the measured projected energy spread can be associated to internal bunch correlations in the y–ξ and $y\prime$–ξ planes, which are amplified when the particles are deflected in the TDS and result in an effective change of energy depending on the longitudinal coordinate ξ (cf. Equation 3.4.4). The differences in slice energy spread, which are also substantial, are not so well understood and could not be further investigated during the campaign. It is possible

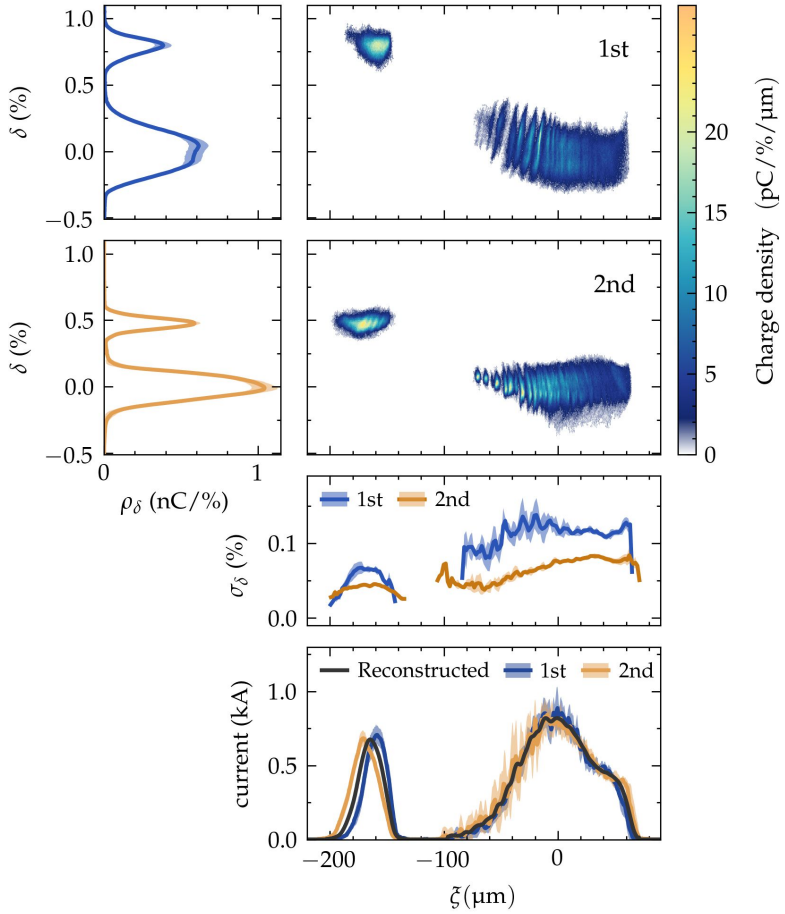

Figure 4.15.: Reconstructed longitudinal phase space measured at the 8FLFDUMP screen.

that they originate from a mismatch between the beam and the lattice optic. As a result of this, the beam envelope is not well collimated at the location of the TDS and shearing effects distort the apparent energy spread at the screen (cf. Equation 5.2.3). Regarding the current profile, the differences between the zero crossings are much less pronounced, as expected from the large streak achieved in the measurement (cf. Section 3.4). The distance between the driver and the trailing bunch in the reconstructed current profile is calculated as the mean value of the distance measured at the two zero crossings, since the 2-point-tomography cannot be applied to the full double-bunch structure (cf. Section 1.3.2).

The goal of the PolariX-TDS campaign is to gain a better insight into the mechanisms of plasma-acceleration through the analysis of the longitudinal phase space of the accelerated bunch, which requires a good understanding of the beam parameters before plasma interaction. To that end, a large number ($\sim$ 150) of trailing-bunch images were taken at each of the zero crossings after removing the driver bunch with the head scraper and adjusting the lattice optics for optimal focusing of the trailing bunch at the screen. This is shown in Figure 4.16. On the left, key beam parameters are plotted versus the full unscraped-bunch RMS length measured at BCM-4DBC3.2. The charge is measured with the toroid TOR-7FLFDUMP,

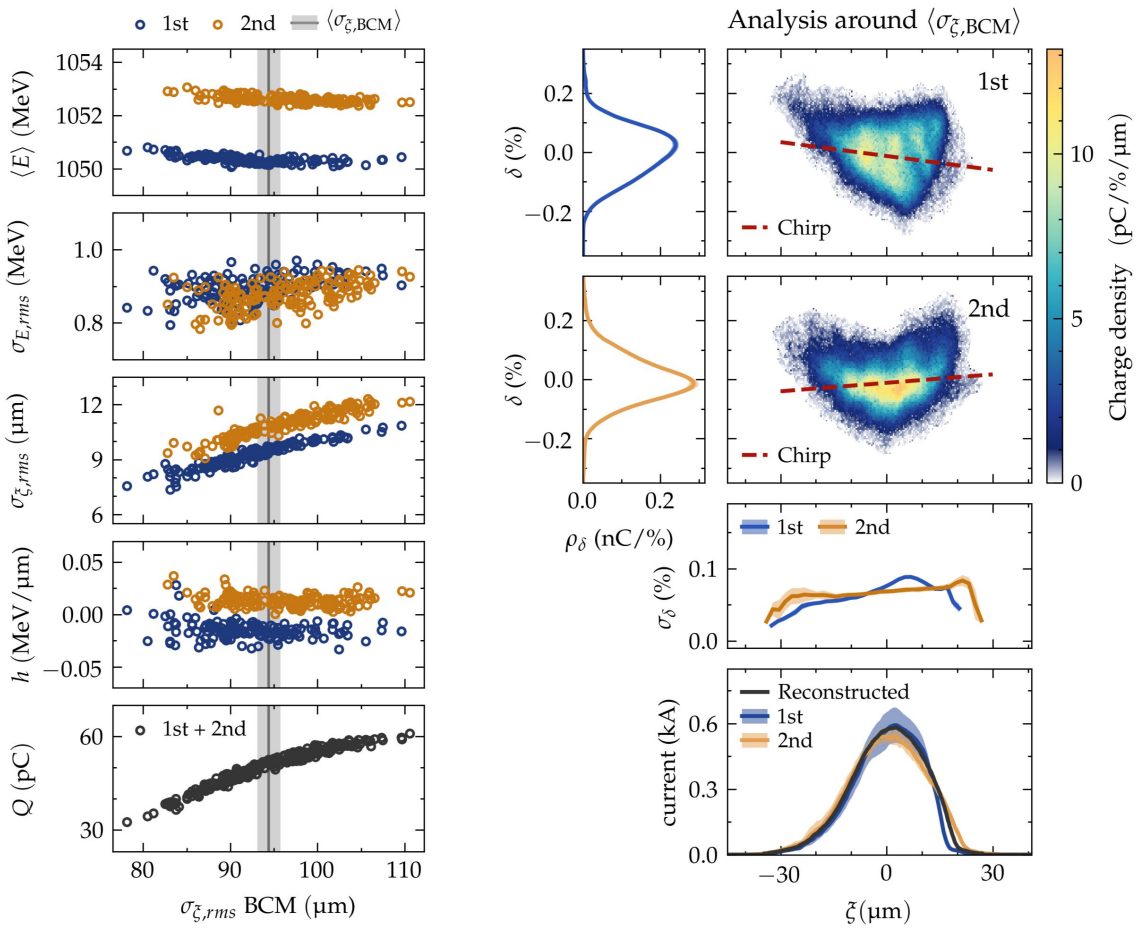

Figure 4.16.: High-statistics longitudinal-phase-space measurements of trailing bunches at the 8FLFDUMP screen. Left: key beam parameters versus unscraped bunch length measured at BCM-4DBC3.2. Right: full longitudinal-phase-space analysis of $\sim$ 20 samples around the mean value of the BCM-4DBC3.2 bunch length.

located $\sim$ 40 cm before the screen. The trailing-bunch RMS length $\sigma_{\xi,rms}$ and its charge Q have a clear correlation with the length of the unscraped bunch, which reveals a strong dependence of the scraping on the linac-induced compression jitter. This is an important finding, since, as has been discussed in Section 4.3, the precise operation of the scraper has an impact on key aspects of the plasma-acceleration process—which will be investigated in Chapter 5. The differences in mean energy $\langle E\rangle$ between the two zero crossings are the result of TDS-induced kicks (cf. Section 3.3) and can be ignored. The right plots show the analysis of $\sim$ 20 images around the mean value of the BCM length and illustrate the differences in projected energy spread and chirp between the two zero crossings. As previously discussed, these longitudinal-phase-space distortions are induced by the TDS and most of them are well understood—with the exception of x–y shearing effects at the screen, potentially related to a lattice-optic mismatch.

The analysis of charge modulations due to compression jitter can be extended to the driver. In Figure 4.17 the trailing-bunch charge from the high-statistics dataset discussed above is added to high-statistics data collected from other measurements performed with the full double-bunch structure. To estimate the charge corresponding to the driver, the trailing-bunch charge is substracted from that of the double bunch. The correlation of the driver charge to the BCM bunch length shows the opposite behaviour than the one for the trailing bunch, so that longer (shorter) unscraped beams produce driver bunches with a lower (higher) charge. Moreover the charge fluctuations in the driver are much larger than in the witness, which is related to the poor performance of both the wedge-scraper at the lower-energy side—i.e., the tail of the driver—and the head-scraper due to CSR effects (cf. Section 4.2.1).

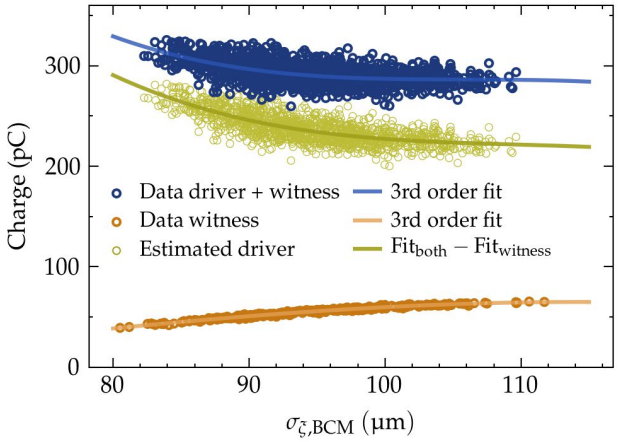

Figure 4.17.: Driver and trailing bunch charge versus unscraped bunch length measured at the BCM-4DBC3.2.

Slice emittance

The slice-emittance measurements at the 11FLFXTDS screen are performed by means of an object-plane scan. The quadrupole triplet that captures the beam is used to adjust the object plane, whereas the remaining quadrupoles are kept at fixed strengths that provide an appropriate optic for TDS measurements, acting as a kind of telescope that requires the beam envolope to be approximately collimated after the triplet. When switching the streaking direction of the PolariX-TDS from y to x, the sign of the focusing stengths k_1 of all quadrupoles is simply reversed.

Figure 4.18 shows the projected beamsize and the charge of both the driver and the witness bunch for each of the two object-plane scans performed. As expected from chromatic effects associated to a strong focusing, the beamsize of each bunch forms a waist at a different location. The larger emittance expected in x is clearly reflected in the larger waist size and steeper beamsize evolution compared to that in y. The charge evolution shown in the bottom plots indicates that the measurements are significantly affected by transmission problems,

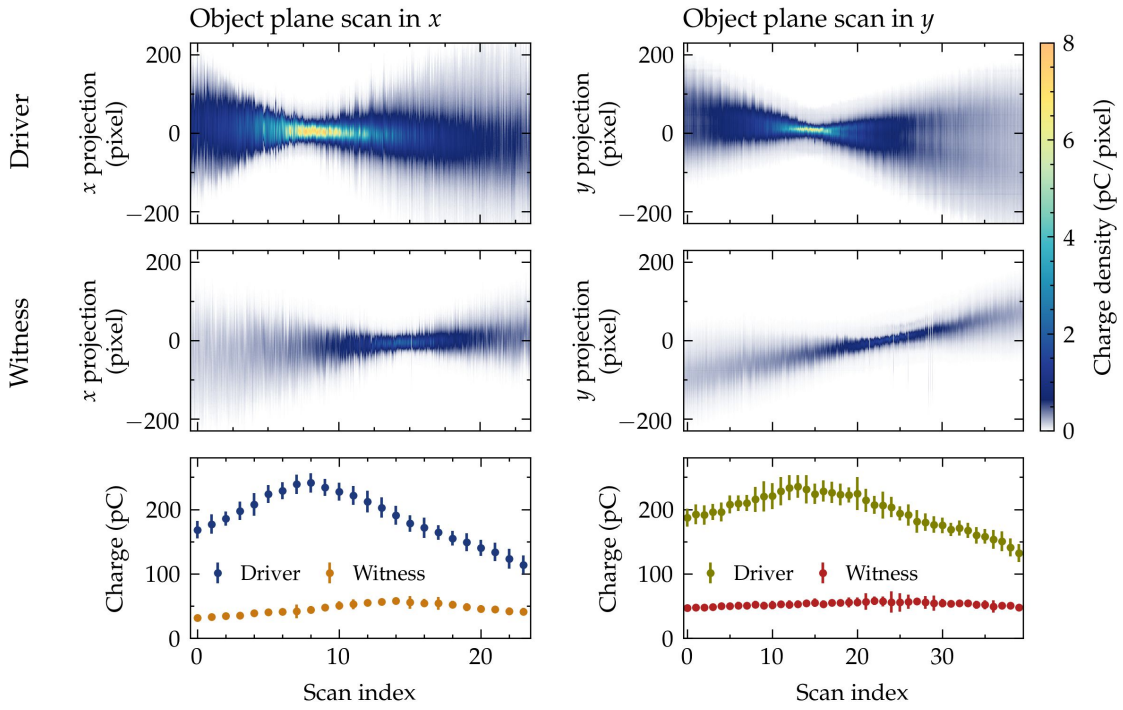

Figure 4.18.: Object-plane scan with double-bunch structure measured at the 11FLFXTDS screen: projected beamsize and charge of driver and witness for vertical (left) and horizontal (right) streaking.

which are clearly correlated with the beamsize evolution. On account of that, the data analysis is restricted to a region around the waist, in which enough points are included to properly define the waist but not more than ~ 25 % of charge is lost. An additional issue affecting the object-plane scan in y unnoticed during the campaign is a faulty operation of the photocathode laser that resulted in the presence of two bunches in the same RF flat-top—this can be observed in the projected beamsize of the driver (top-right plot), which forms two lobes at each side of the waist. This issue represents a serious drawback to the beam characterisation, but since this is the only measurement providing information about the phase space in the y–ξ plane, the analysis is performed by taking the two bunches into consideration with a dedicated beamsize-fitting algorithm.

Prior to the slice-emittance calculation, a projected-emittance analysis is performed, which provides a better insight into the overall waist location and trajectory of driver and trailing bunch relative to the plasma cell. This is shown in Figure 4.19, which contains the calculated beamsizes (left) and centroids (right) at the location of each object plane scanned along the measurement. The differences between the x and y axes are clearly noticeable. The driver emittance in x ($\epsilon_{x,d} = 7.5 \pm 0.6$ mm mrad) is ~ 5 times larger than that in y ($\epsilon_{y,d} = 1.4 \pm 0.1$ mm mrad). In the case of the witness, with an emittance in x ($\epsilon_{x,w} = 3.7 \pm 0.3$ mm mrad) approximately ~ 3 times larger than in y ($\epsilon_{y,w} = 1.2 \pm 0.1$ mm mrad), the difference is less pronounced but also significant. The better degree of matching in x than in y is reflected in both the closer location of the waist to the plasma cell (left) and the

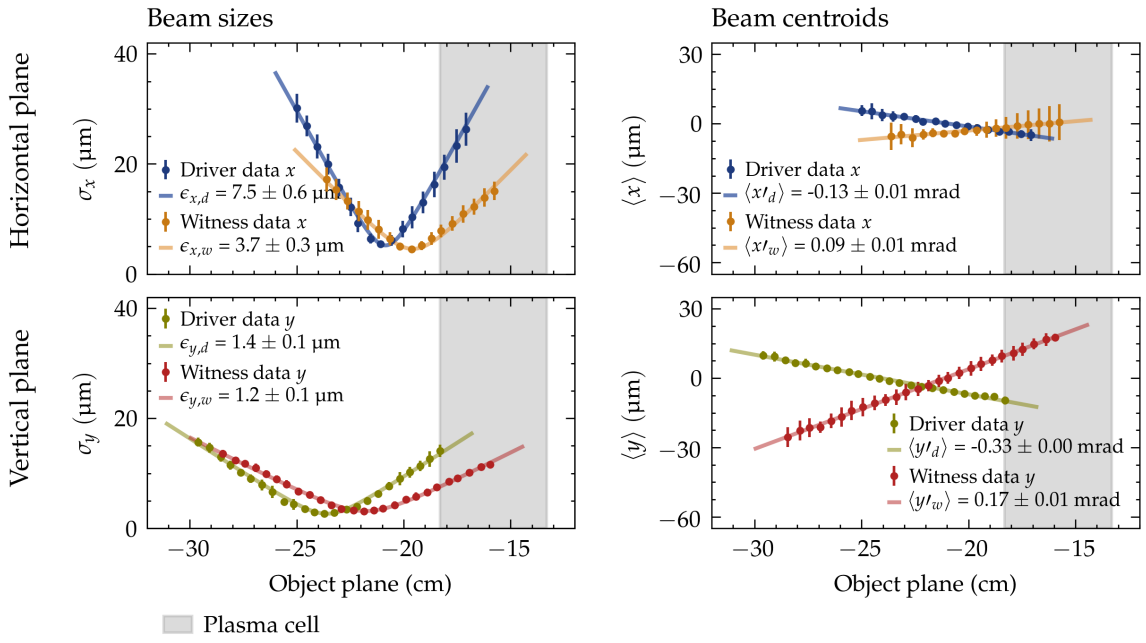

Figure 4.19.: Projected emittance in x and y of the double-bunch structure measured at the 11FLFXTDS screen.

better colinearity between the trajectories of driver and witness (right). A comparison of the results in the horizontal plane (top plots) with those obtained at the LEMS (Figure 4.12) exposes slight differences between the two measurements. In the latter, the waist location of both bunches seems to be closer to the cell, the emittance of the witness bunch is slightly larger, and the bunch trajectories show the opposite behavior. Theses differences are most probably related to the issues affecting the measurement at the 11FLFXTDS screen—i.e., larger chromatic effects and transmission problems—but could also be related to small machine drifts occurring between the two measurements. In any case, the level of agreement between the two of them is considered satisfactory.

Figure 4.20 shows the computed slice-emittance in x and y, for which longitudinal resolutions of $R_{\xi,x} = 1.7$ µm (5.5 fs) and $R_{\xi,y} = 3.0$ µm (10.1 fs) were obtained in each of the measurements, respectively. The reconstruction plane is chosen to be at the location between the waist of the driver and that of the trailing bunch. For each slice, the calculation of the transfer-matrix elements takes into account the relative energy measured at the 8FLFDUMP screen, which is then combined with the absolute mean energy measured at the ESPEC. Regarding the measurement in the horizontal plane several observations can be made. At the front part of the driver, the emittance goes up to values around 12.5 mm mrad, whereas for the rest of the double-bunch structure it has values around and below 5 mm mrad. This can be associated to the larger impact of CSR effects at the front of the bunch. The beta function is considerably well matched along the whole extent of the driver, whereas for the trailing bunch it progressively increases from the head to the tail. As for the alpha function, a general trend going from negative values (divergent) at the head to positive values (convergent)

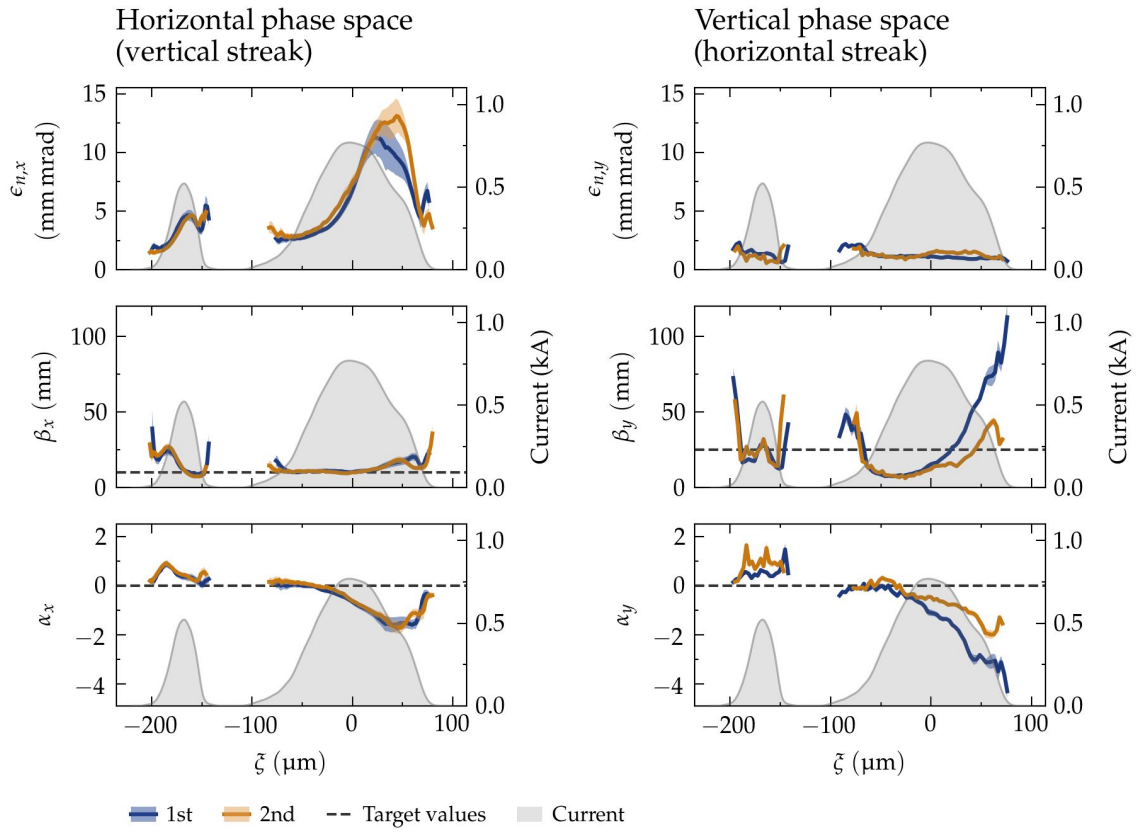

Figure 4.20.: Slice emittance and slice Courant-Snyder parameters in x and y of the double-bunch structure measured at the 11FLFXTDS screen.

at the tail is observed. This overall behavior of the sliced Courant-Snyder parameters is related to differences in mean energy along the bunch, according to which lower energies are focused further upstream than higher energies. The results of the measurement in the vertical plane are manifestly affected by noise associated to the complex analysis required to disentangle the two bunches present at the screen. Nonetheless, some general trends can be observed. The emittance along the whole double-bunch structure is approximately constant, which is expected from the much lower impact of collective effects along the beamline in this plane. Regarding the sliced Courant-Snyder parameters, the trailing bunch seems to be better matched to the target values than the driver, the latter showing an increase in the beta function towards the head and at the same time becoming more divergent.

The sliced centroids computed from the object-plane scans are shown in Figure 4.21. The reconstruction plane is the same as for the slice emittance. For each slice ξ_i, the Courant-Snyder parameters are used to calculate the phase-space ellipse corresponding to an area of $\pi\epsilon_n(\xi_i)$, which is plotted on top of each centroid. Slight differences between the 1st (left) and 2nd (right) zero crossings can be noticed, which most probably arise from the kick perpendicular to the streaking plane introduced by the PolariX-TDS—whose sign depends on the zero crossing. Despite this, the measurements performed for each plane remain consistent enough to draw some general conclusions. The overlap in phase space between the driver

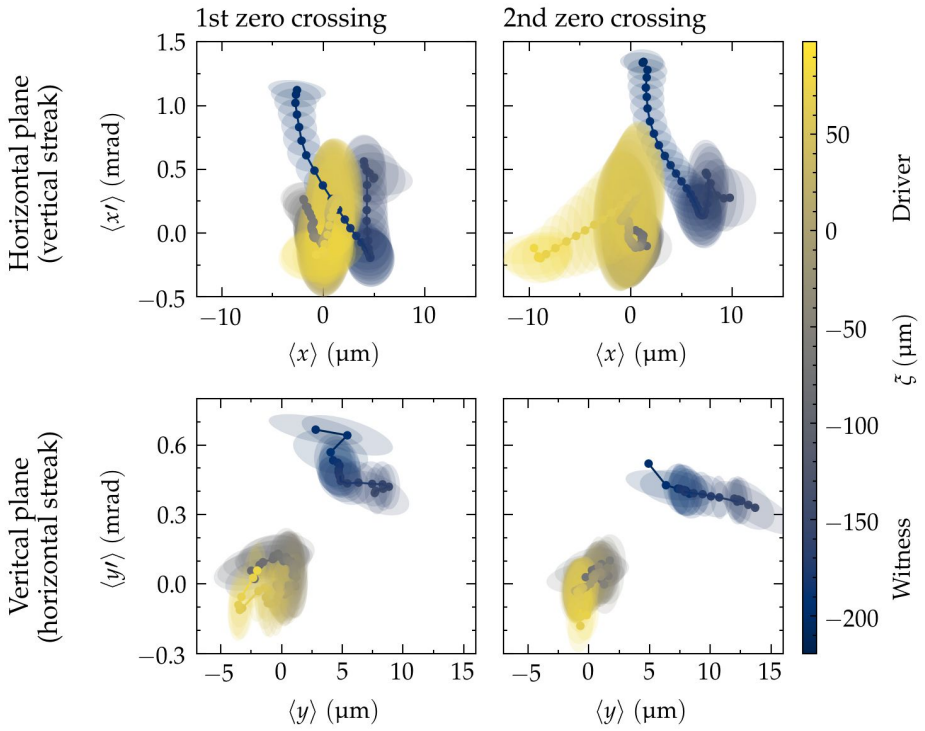

Figure 4.21.: Slice centroids in x and y of the double-bunch structure reconstructed upstream of the gas cell. See text for details.

and the witness bunches is better in the x plane than in the y plane, which is directly related to the fact that the beam-plasma-interaction optimisation is mainly performed by diagnosing the beam parameters in the x axis—i.e., any manipulations affecting the y axis remain mostly unnoticed unless a dedicated sliced 2-BPM-tomography is performed. Nonetheless, since the emittance in y is smaller than in x, in the end the overall phase-space area covered in the two planes is roughly the same. One issue that might compromise the coupling of the witness bunch into the wakefield structure is the relatively large angular misalignment of the tail of the trailing bunch—especially in the x axis. To investigate to what extent this is the case, an analysis of the trailing bunch after plasma interaction is required, which is addressed in the next chapter.

True slice energy spread

With the phase-space information provided by the slice-emittance measurements in x and y it is possible to estimate the TDS-induced slice energy spread and the energy resolution of the longitudinal-phase-space measurements—both of which are required to retrieve the true slice energy spread of the bunch (cf. Section 1.3.2). The first can be calculated by means of Equation 1.3.42 using the Courant-Snyder parameters in the streaking plane at the entrance of the PolariX-TDS and the measured shear parameter. The latter is given by the natural beamsize in the dispersive plane and the lattice dispersion at the location of the screen (cf.

Equation 1.3.39). The screen resolution, which is about ~ 10 µm and results in a contribution of ~ 15 keV to the measured energy spread for typical dispersions of $D_x \sim 0.7$ m^{-1} and ~ 1 GeV beams, will be ignored in the calculations. In the following the central slice of the trailing bunch is considered and the data measured at both zero crossings is treated indistinguishably.

The top plot of Figure 4.22 shows the beta functions in x and y for the lattice optic used in the high-statistics longitudinal-phase-space measurements. The error bands result from propagating the calculated beam-parameter errors along the beamline. To estimate the TDS-induced energy spread, a large number ($N = 2 \cdot 10^5$) of randomly-generated initial Courant-Snyder parameters in the y axis, normally distributed around the measured values, is tracked down to the TDS entrance. With a calculated kick of $\mathcal{K} = 2.6$ m^{-1} and a measured emittance of $\epsilon_y = 1.2 \pm 0.1$ mm mrad, an induced energy spread of $\sigma_{IES} = 0.60 \pm 0.05$ MeV is obtained—where $\mathcal{K} = S/R_{34}$, the measured shear parameter is $S = 65.1 \pm 0.1$, and the

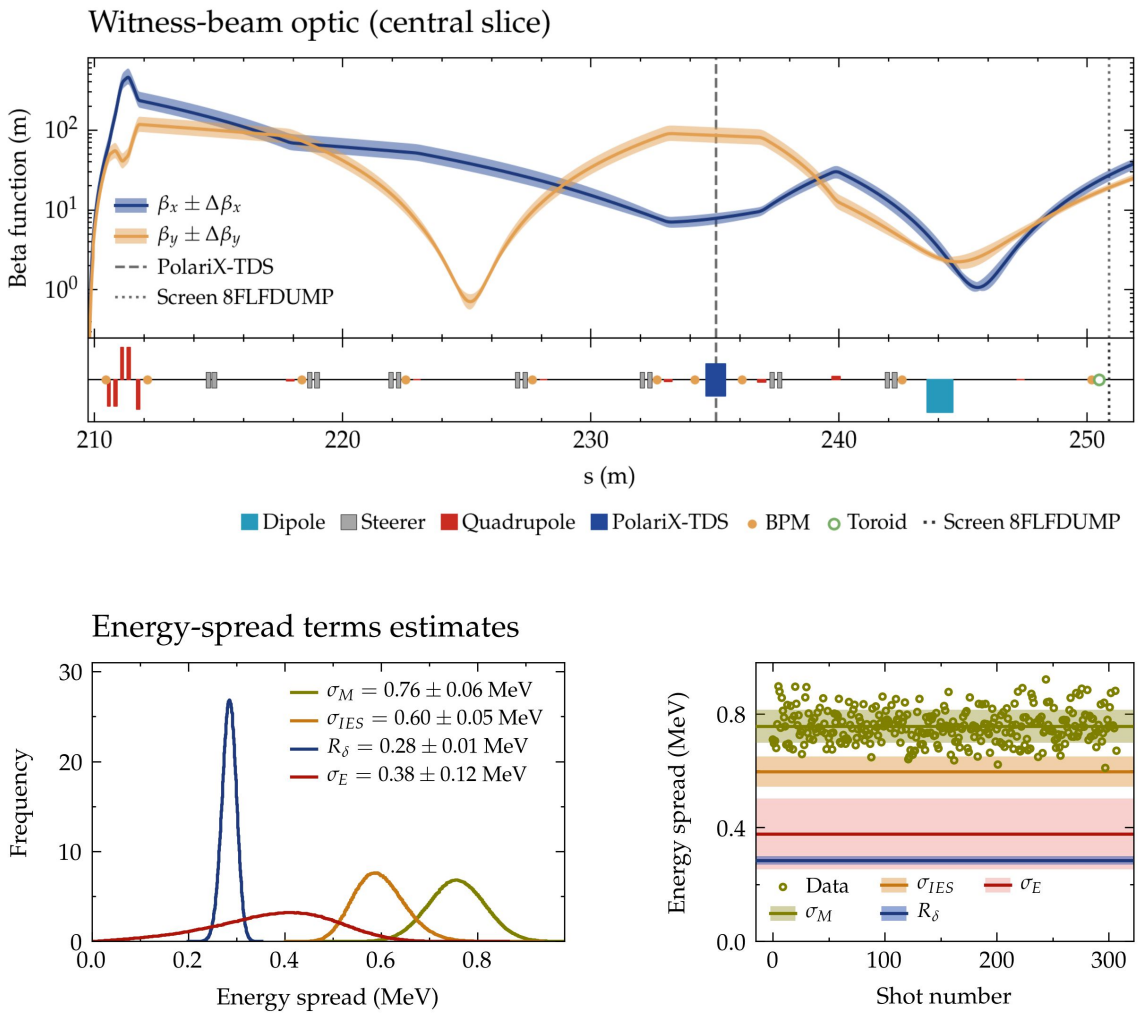

Figure 4.22.: Retrieval of the true energy spread of the central slice of the uninteracted trailing bunch. See text for details.

calculated matrix element between the TDS and the screen is $R_{34} = 25.1$ m. The same method is applied to calculate the energy resolution at the screen from the initial Courant-Snyder parameters in the x plane, the measured emittance $\epsilon_x = 4.1 \pm 0.4$ mm mrad and the dispersion at the screen $|D_x| = 0.591$ m^{-1}, which results in a value of $R_\delta = 0.28 \pm 0.01$ MeV. The probability distributions calculated in this way are shown in the bottom-left plot of Figure 4.22, together with the distribution reproducing the measured energy spread $\sigma_M = 0.76 \pm 0.06$ MeV and the real spread calculated according to the expression:

$$\sigma_E = \sqrt{\sigma_M^2 - \sigma_{IES}^2 - R_\delta^2}, \tag{4.4.1}$$

from which a final value of $\sigma_E = 0.38 \pm 0.12$ MeV is obtained. The same results are shown in the bottom-right plot together with the data points measured at the 8FLFDUMP screen. Due to the errors in the estimated TDS-induced energy spread and energy resolution, the relative error of the true energy spread ($\sigma_{\sigma_E} \approx 32$ %) is substantially larger than that of the measured energy spread ($\sigma_{\sigma_M} \approx 8$ %). Therefore, it can be claimed that the shape of the probability distribution of the true slice energy spread is mainly dominated by measurement errors rather than by shot-to-shot variations of the beam parameters themselves. On the other hand, the estimated real slice energy spread agrees very well with the upper bound of the projected energy spread measured at the LEMS, for which a value of $\sigma_E = 0.45 \pm 0.02$ MeV was obtained, and supports the validity of the method presented here.

5. Characterisation of electron bunches accelerated in a nonlinear beam-loaded plasma-wakefield

In this chapter, the ultimate goal of the PolariX-TDS experimental campaign is addressed, which is to investigate two important aspects required to preserve the quality of electron beams externally injected in a beam-driven plasma wakefield: optimal beam loading for chirp compensation and uniformity of the accelerating gradients in the transverse plane to preserve the uncorrelated energy spread. These investigations are mainly enabled by the powerful capabilities of the PolariX-TDS, which provides an unprecedented insight into the longitudinal-phase-space structure of the accelerated bunches, from which key properties of the field gradients in the plasma can be retrieved. In the experimental conditions of this campaign, however, the measurements are subject to several sources of errors that can be grouped in the following categories: 1) stability of the machine operation; 2) stability of the plasma-acceleration process; 3) inherent limitations of the TDS-diagnostic and its sensitivity to the lattice optic. On account of that, a subordinate but critical goal of the current chapter is to disentangle these sources of errors by means of complementary measurements performed at the remaining diagnostic stations downstream of the plasma chamber—viz. the two dipole spectrometers ESPEC and LEMS and the screen 11FLFXTDS located immediately after the PolariX-TDS—and computer simulations.

Section 5.1 delves into the analysis of all the measurements performed, from which the relation between the experimental conditions of the campaign and the degree and quality of the achieved beam loading is disclosed. In Section 5.2 the focus is on three sources of errors that need to be understood to properly interpret important aspects of the previous analysis: 1) energy jitter resulting from beam-loading modulations; 2) beam dispersion; 3) longitudinal-phase-space distortions induced by the TDS. In Section 5.3 a numerical estimation of the accelerating-gradient uniformity is provided and the accuracy of the result is evaluated in light of the preceding error analysis. Finally, in Section 5.4, particle-in-cell simulations are carried out in order to reproduce the experimental results and, in turn, verify the physical interpretation of the acceleration process discussed in the previous sections.

5.1. Beam-loading modulations and beam characterisation

Similar to the approach followed with the noninteracted double-bunch structure (cf. Section 4.4), the characterisation of the plasma-accelerated bunches requires multiple measurements at each of the diagnostic stations in the post-plasma beamline. On the one hand, the two energy spectrometers—ESPEC and LEMS—provide the most robust and reliable means of absolute energy, projected energy spread and projected emittance in the x–x' plane, and their proximity to the plasma capillary reduces the risk of possible transmission problems. On the other hand, the evaluation of the longitudinal-phase-space measured at the 8FLFDUMP screen relies on the knowledge of the emittance and the Courant-Snyder parameters in x and y along the beamline, which are obtained from slice-emittance measurements performed at the 11FLFXTDS screen.

In the following, all the measurements mentioned above are analysed individually and the obtained results are compared with each other in order to identify the principal sources of errors affecting each of them. The longitudinal phase space is analysed first, since it provides key insights into the acceleration process, thus guiding the analysis of the rest of the measurements.

5.1.1. Longitudinal phase space

Longitudinal-phase-space measurements are performed by point-to-point imaging the beam in both x and y from the output of the capillary to the 8FLFDUMP screen with a nominal energy of $E = 1090$ MeV. In the y plane, additionally, the lattice section between the PolariX-TDS and the screen is set to parallel-to-point imaging and a design beta function of $\beta_y \approx 100$ m at the TDS is assumed. By operating the RF system at its maximum power of 6 MW a time resolution of $R_\xi \approx 8.2$ fs (2.46 µm) is obtained, which corresponds to a rounded-up slice size of 7 pixels. In these conditions, $N = 800$ samples are taken in each of the two zero crossings—which is reduced to $N \approx 700$ after filtering out the shots that are not fully captured within the FOV of the screen.

Figure 5.1 shows selected sample images of the measurements, revealing dramatic differences in the acceleration process: for each zero crossing, the amount of beam loading increases progressively from the top-left (underloaded) to the bottom-right (overloaded) images. As expected from the Gaussian-like current profile of the incoming trailing bunch (cf. Section 4.4), the longitudinal phase space cannot be completely flattened as expected with an optimally shaped trapezoidal profile [Tzoufras et al., 2008]. Therefore, the wakefield is effectively beam loaded only around the core of the bunch while it is mostly unperturbed at its outer regions, where it imprints a clear negative chirp that confers an overall S-shape on the beam. The slice energy spread seems to increase as the wakefield becomes more strongly beam loaded, which could be related to changes in energy resolution. Additionally, there are two remarkable differences between the measurements performed at each zero crossing: 1)

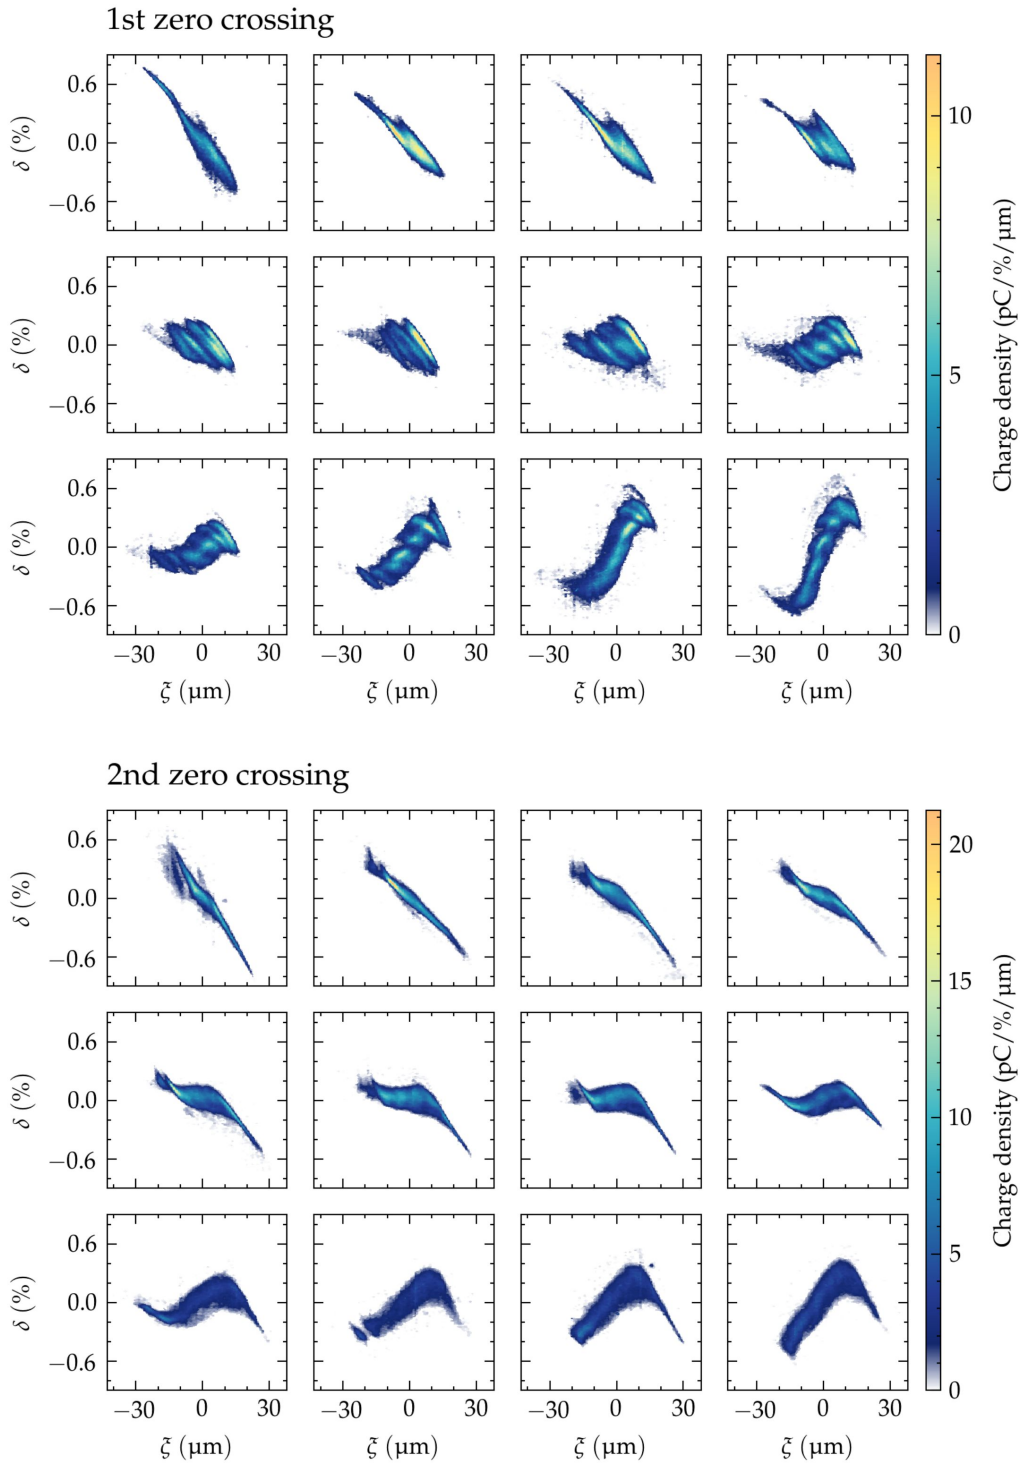

Figure 5.1.: Direct observation of nonlinear chirps of trailing bunches accelerated in a beam-loaded plasma wakefield measured with the PolariX-TDS at the 8FLFDUMP screen at the 1st and 2nd zero crossing.

in the 1st, a kind of shearing seems to distort the phase space by introducing a remarkable negative correlation between the longitudinal coordinate ξ and the energy deviation δ; 2) in

the 2nd zero crossing the head and tail of the bunch are well defined and extend outwards of the core, whereas in the 1st they seem to be folded on top of it, which could be explained by strong y–ξ and y'–ξ correlations along the bunch.

In Chapter 4 the exact shape of the double-bunch structure produced with the scraper was found to depend on the total length of the bunch, which was in turn modulated by compression jitter induced by the SCRF modules at the linac. To investigate to what extent this is the underlying mechanism driving the beam-loading modulations described above, the correlation between the total bunch length measured at the BCM-4DBC3.2 and several key parameters is evaluated in the same way as performed with the uninteracted bunch (cf. Figure 4.16). The projected rms energy spread $\sigma_{E,rms}$ and the chirp h are the two parameters that unambiguously characterise the amount of beam loading. Instead of the conventional statistical definition of chirp (cf. Eq. 1.1.38), however, a new definition is introduced that represents the linear slope of the bunch along one rms sigma $\sigma_{\xi,rms}$ around the core, weighted by the current profile. This is achieved by minimising the function:

$$\chi^2 = \sum_{i=1}^{n} Q_i^2 \left[\delta_i - (\delta_0 + h\xi_i)\right]^2 \ , \quad \text{for} \quad |\xi_i - \langle\xi\rangle| < \frac{\sigma_{\xi,rms}}{2} \ , \tag{5.1.1}$$

where h is the chirp (i.e., the minimiser of the function), Q_i, δ_i and ξ_i are the charge, energy and longitudinal location of each slice, i is the slice index, and n is the total number of slices. Since the beam loading has a significant effect only around the core of the bunch, the new chirp definition represents the amount of beam loading with higher fidelity.

As can be seen on the left hand of Figure 5.2, the energy spread and the chirp do indeed exhibit a clear correlation with the unscraped bunch length: at shorter lengths the chirp is negative and the projected energy spread is large; at lengths around the mean value $\langle\sigma_{\xi,BCM}\rangle$ the chirp approaches $h = 0$ and the energy spread reaches its minimum—both of them close to the uninteracted trailing bunch values (red lines); at larger lengths the energy spread increases again and the chirp becomes positive. The differences between the two zero crossings are due to internal bunch correlations in the streaking plane y–ξ and y'–ξ that modulate the measured chirp and consequently the projected energy spread. These correlations are also responsible for the differences in measured rms accelerated-bunch length $\sigma_{\xi,rms}$. The positive linear correlation between the mean energy $\langle E\rangle$ and the BCM bunch length supports the argument that the compression jitter is the driver of beam-loading modulations, since for shorter (longer) lengths the trailing bunch would arrive at an earlier (later) phase of the wakefield and would experience an overall smaller (larger) field gradient E_z. The difference in mean energy between the two zero crossings can be associated to changes in the beam orbit through the TDS, resulting from the energy-dependent transport of the bunch along the magnetic lattice. An overall charge loss of roughly 40 % can be observed in the bottom plot, which could be due to the combined effect of a suboptimal charge coupling into the wakefield structure and/or transmission problems through the beamline. Compared to the noninteracted bunch (red line) the relative charge loss is more pronounced for BCM lengths above $\langle\sigma_{\xi,BCM}\rangle$. In general, the correlation between all parameters and the BCM

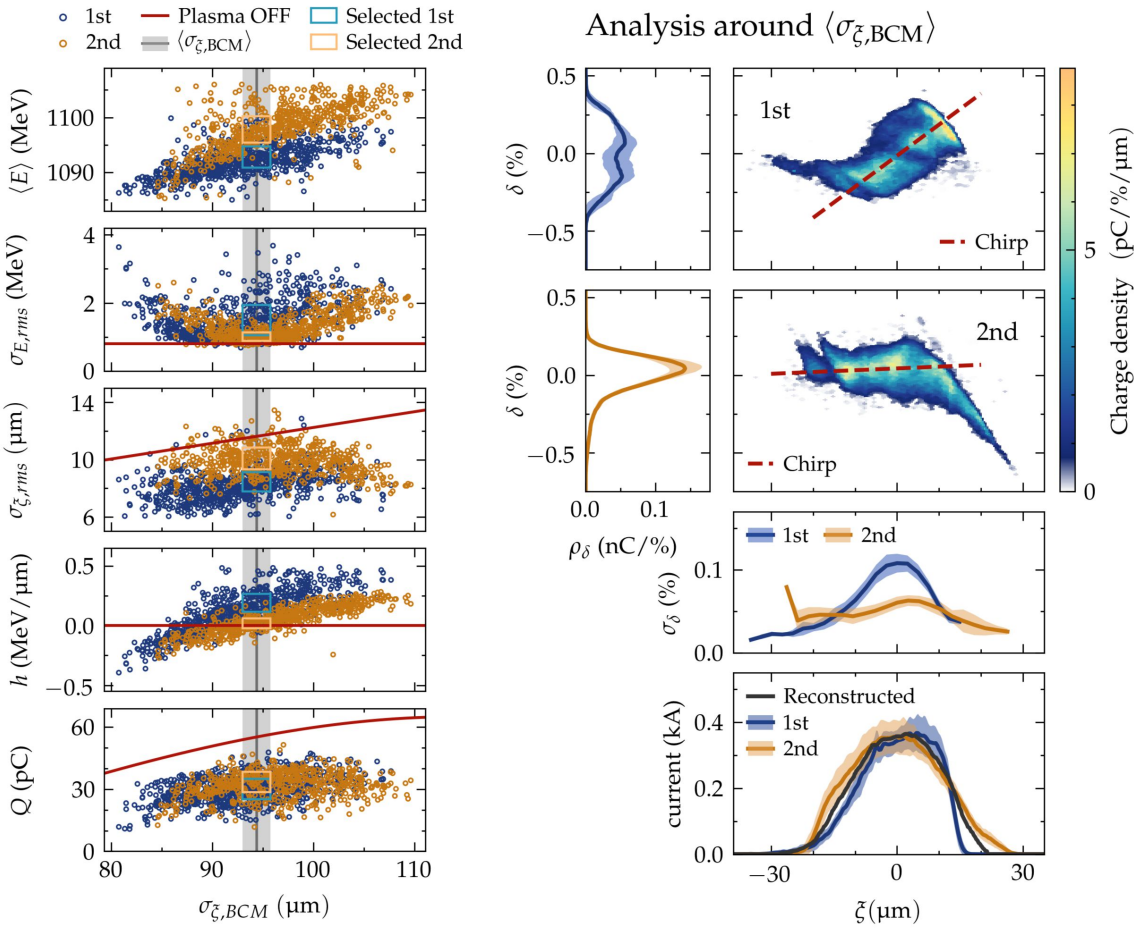

Figure 5.2.: High-statistics data from longitudinal-phase-space measurements. Analysis around the working point $\langle\sigma_{\xi,BCM}\rangle$: for each zero crossing only ~ 20 shots that fulfil the condition $|x - \langle x\rangle| < \sigma_x$ for all parameters simultaneously (i.e., $x \in [\langle E\rangle, \sigma_{E,rms}, \sigma_{\xi,rms}, h, Q]$).

bunch length is smeared out due to additional second-order sources of jitter, among which the most relevant could be expected to be: plasma density variations, BCM-bunch-length measurement errors, and TDS measurement errors.

A detailed longitudinal-phase-space analysis is performed around the mean BCM bunch length $\langle\sigma_{\xi,BCM}\rangle$, which corresponds to the working point targeted during the optimisation routines discussed in the previous chapter. For each zero crossing, $N = 20$ shots are selected that fulfil the condition $|x - \langle x\rangle| < \sigma_x$ for all parameters simultaneously, where $x \in \{\langle E\rangle,\ \sigma_{E,rms},\ \sigma_{\xi,rms},\ h,\ Q\}$. In the left hand of Figure 5.2 the region covered by the selected points is indicated with a light-blue (light-orange) rectangle for the 1st (2nd) zero crossing, and the analysis is shown on the right side of the figure. As already observed in practically all longitudinal-phase-space measurements presented in this thesis, the energy axis is the one that exhibits the largest dependence on internal bunch correlations y–ξ and y'–ξ, which is clearly reflected in the projected energy profile and the calculated chirp at each of the zero crossings. Additionally, the head of the bunch—which is accelerated by

the unperturbed plasma-wakefield—seems to be affected by extreme correlations, since the longitudinal slices in this region overlap on top of each other in the 1st zero crossing, whereas they largely extend ahead of the bunch in the 2nd zero crossing. This issue compromises the applicability of the 2-point-tomographic reconstruction, since the mapping function of the longitudinal coordinate ξ to the screen coordinate y is not injective over the whole length of the bunch (cf. Section 1.3.2). Furthermore, shearing effects that distort the longitudinal phase space are most clearly recognised in the 1st zero crossing and, together with the induced energy chirp, translate into a substantial increase of the slice energy spread. These longitudinal-phase-space distortions will be investigated in Section 5.2.3.

5.1.2. Projected beam parameters

The mean energy is measured at the ESPEC by setting point-to-point imaging from the output of the capillary to the screen and taking a large number of samples ($N = 2000$), the analysis of which is shown in Figure 5.3. On the top left, the projected energy spectra of all shots are sorted by mean energy. According to the longitudinal-phase-space measurements performed at the 8FLFDUMP screen, lower mean energies correspond to underloaded beams—i.e., the bunch has a negative chirp with the head (tail) having lower (higher) energies—whereas higher mean energies correspond to overloaded beams—i.e., the bunch has a positive chirp with the head (tail) having higher (lower) energies. The mean energy distribution (bottom-left plot) describes a slightly asymmetric Gaussian distribution around $\langle E \rangle = 1096.4$ MeV with a standard deviation of $\sigma_{\langle E \rangle} = 3.7$ MeV (0.34 %). On the right, three parameters are plotted versus the unscraped bunch length measured at the BCM-4DBC3.2: the mean energy $\langle E \rangle$ (top right), the rms projected energy spread $\sigma_{E,rms}$ (middle right) and the charge Q (bottom right), which is computed from the CCD counts of the screen camera. Compared to the measurement at the 8FLFDUMP screen, the mean energy and the BCM bunch length correlate unambiguously and the absolute energy values are more reliable due to the negligible effects of angular misalignments of the beam, which also translate into a reduction of the measurement jitter. The projected energy spread describes the same parabolic correlation measured at the 8FLFDUMP screen and it seems to have its minimum slightly below the mean BCM bunch length $\langle \sigma_{\xi,BCM} \rangle$, which would correspond to the best beam loading achievable in the current experimental conditions. However, the resolution of the measurement is limited to around 2 MeV and a final conclusion cannot be drawn in this regard from the data available. The charge, on the other hand, is higher than that of the longitudinal-phase-space measurement, which confirms that a fraction of the charge loss in the latter is due to transmission problems. To characterise the beam parameters of the working point, $N = 50$ samples are taken around the mean BCM bunch length $\langle \sigma_{\xi,BCM} \rangle$ (vertical grey dashed line), which result in a mean energy of $\langle \widetilde{E} \rangle_w = 1097 \pm 2$ MeV and a charge of $\widetilde{Q}_w = 41 \pm 4$ pC. Assuming that no charge is lost between the capillary output and the ESPEC screen, it can be concluded that on average, at the working point, only around 75 % of the 55 pC sent into the plasma is effectively coupled into the wakefield structure. This value decreases further for larger bunch lengths, most probably due to the fact that the

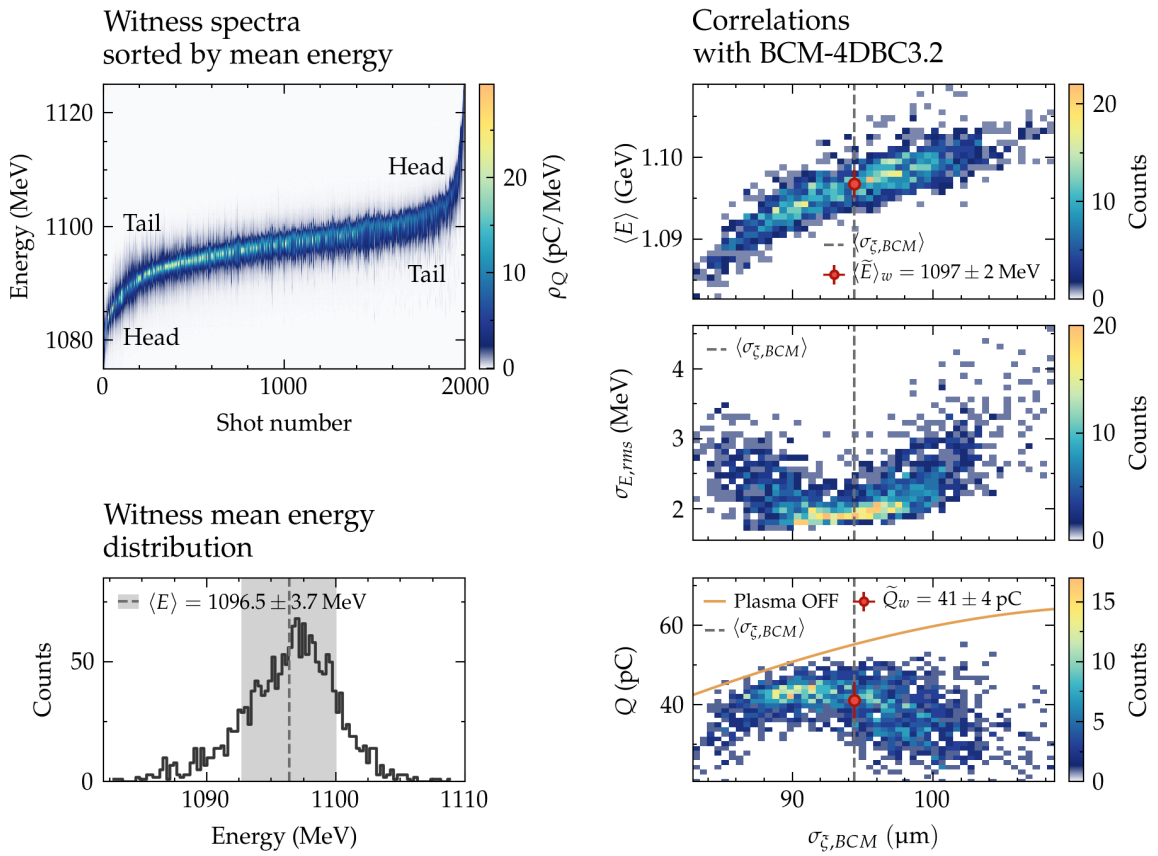

Figure 5.3.: Key projected beam parameters of accelerated bunches measured at the ESPEC, including: energy spectrum (top left), mean energy histogram (bottom left), and correlation of the bunch length $\sigma_{\xi,BCM}$ with mean energy $\langle E \rangle$ (top right), rms projected energy spread $\sigma_{E,rms}$ (middle right), and charge Q (bottom right).

trailing bunch is closer to the back of the bubble.

In order to estimate the true projected energy spread, the data from an emittance measurement performed at the LEMS is analysed, which is shown in Figure 5.4 (energy spread on the left and charge on the right). Even if the energy resolution is largely improved in this station, the parabolic correlation with the BCM bunch length previously observed is substantially distorted. Since there are no additional charge losses compared to the measurement at the ESPEC (cf. Figure 5.3), the energy spread distortions could be associated to beam dispersion modulated by the changes in quadrupole strength at different scan steps. Despite these effects, it is possible to calculate an upper bound of the projected energy spread for $N = 20$ samples around the working point $\langle \sigma_{\xi,BCM} \rangle$ (grey dashed vertical line), for which a value of $\widetilde{\sigma}_{E,rms} = 0.8 \pm 0.2$ MeV is obtained. Compared to the projected energy spread of the noninteracted witness bunch $\sigma_{E,rms} = 0.45 \pm 0.02$ MeV, this represents an increase of ≈ 0.35 MeV that could be mainly associated to the contribution of the head and the tail of the bunch, which are not affected by beam loading and broaden the resulting energy spectrum.

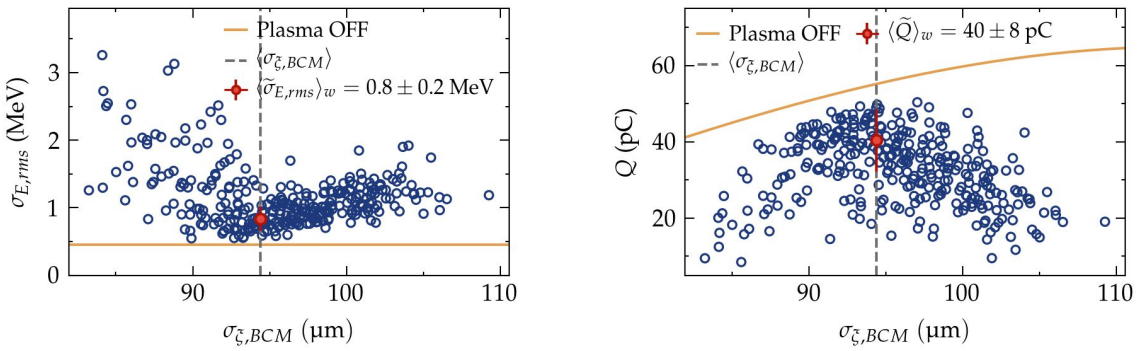

Figure 5.4.: Correlation between bunch length $\sigma_{\xi,BCM}$ and rms projected energy spread $\sigma_{E,rms}$ (left) and charge Q (right) of accelerated bunches measured at the LEMS.

The full energy spectrum of the interacted driver bunch is measured by performing an imaging energy scan at the ESPEC, which is shown on the left plot of Figure 5.5. The reconstructed driver spectrum is computed by averaging the range of focusing energies from 990 to 1050 MeV (grey shaded area) for the $N = 20$ shots closest to the mean BCM bunch length $\langle\sigma_{BCM,rms}\rangle$, whereas the projected energy spread of the accelerated witness is taken from the LEMS analysis discussed before. Both of them are shown on the right plot together with the noninteracted energy spectrum obtained in Section 4.4. The indicated mean energies correspond to the first moment of the energy distribution of both the driver and the trailing bunch.

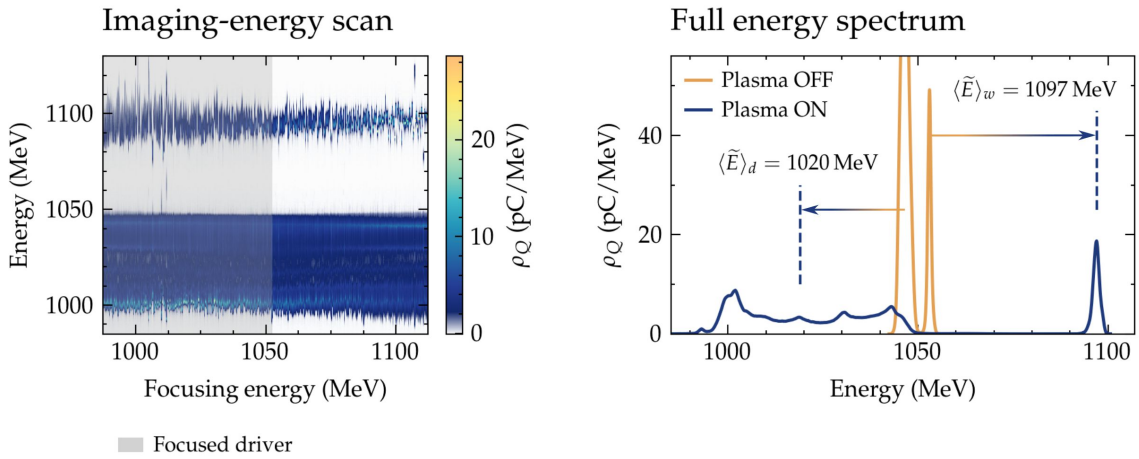

Figure 5.5.: Full energy spectrum of plasma-interacted double-bunch structure. Interacted driver taken from imaging energy scan at ESPEC; interacted witness taken from object plane scan at LEMS. Noninteracted double-bunch structure measured at ESPEC (cf. Section 4.4.1).

The projected emittance in x is measured by means of an object-plane scan performed at the LEMS (cf. Figure 5.6). Compared to the noninteracted bunch, the emittance is reduced from $\epsilon_{x,w} = 4.8 \pm 0.1$ mm mrad (cf. Table 4.4) to $\epsilon_{x,w} = 2.4 \pm 0.2$ mm mrad, which can be directly associated to charge loss from the imperfect coupling of the bunch into the wakefield

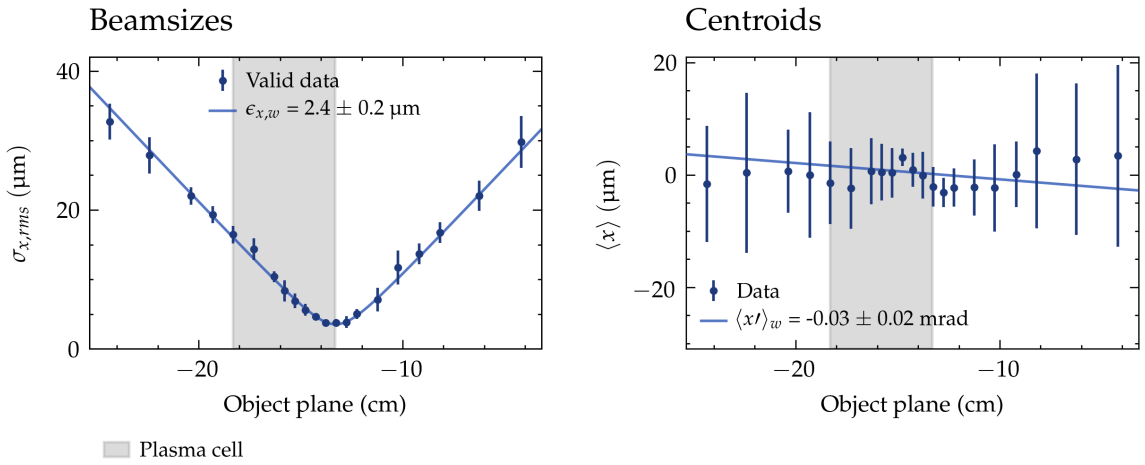

Figure 5.6.: Results of object-plane scan performed at the LEMS to measure the projected emittance in x of accelerated witness bunches. Left: rms beamsize evolution around the waist. Right: transverse beam centroid evolution around the waist.

structure—most probably due to angular misalignment between the driver and the trailing bunch as discussed in the previous chapter (cf. Section 4.4). Otherwise, the beam waist is located almost exactly at the capillary output and the angle is close to the $\langle x' \rangle_w \approx 0.0$ mrad in very good agreement with the value targeted during the optimisation (cf. Section 4.4). At the waist, the measured beta function is $\beta_{x,s0} = 11.5 \pm 0.9$ mm, which is very close to the value measured for the noninteracted trailing bunch at the entrance of the capillary $\beta_{x,s0} = 14.0 \pm 0.3$ mm and indicates an acceptable degree of matching into the focusing ion channel behind the driver.

5.1.3. Slice emittance in x and y

The characterisation of the accelerated trailing bunch is completed with slice-emittance measurements in x and y performed at the 11FLFXTDS screen, for which the same approach as the one followed to characterise the noninteracted double-bunch structure is applied—i.e. the quadrupole triplet that captures the beam is used to scan the object plane around the waist in the direction perpendicular to the TDS streak, while the remaining magnets are set to fixed focusing strengths suitable to transport the beam down to the screen and achieve a good time resolution. For the object-plane scan in the horizontal axis, the RF system is operated at its maximum power and a time resolution of $R_\xi \approx 8.6$ fs (2.56 µm) is achieved, corresponding to a slice size of 13 pixels. For the scan in the vertical axis, the power is reduced to adjust the streaked beam size to the field-of-view of the screen station and a time resolution of $R_\xi \approx 10.2$ fs (3.0 µm) is achieved, corresponding to a slice size of 15 pixels. The magnetic lattice is operated at a nominal momentum of $p = 1096$ MeV/c and for each quadrupole setting $N = 25$ samples are taken.

Figure 5.7 shows the projected particle distribution for each of the two object-plane scans performed at each of the two zero crossings, together with the charge calculated from the CCD-camera pixel counts and the beam size computed by means of a Gaussian fit. Most remarkably, the measurement reveals a drop in charge transmission that increases for the quadrupole settings away from the beam waist. These charge losses are expected to originate at locations along the beamline at which the beta function is too large for the particles with the highest emittance to be effectively transported, and would translate into an overall reduction of the measured beam emittance. To minimize this effect, the data analysis is restricted to the region around the waist for which the beam-size change is large enough to determine the beam divergence and no more than ~ 25 % of the charge is lost.

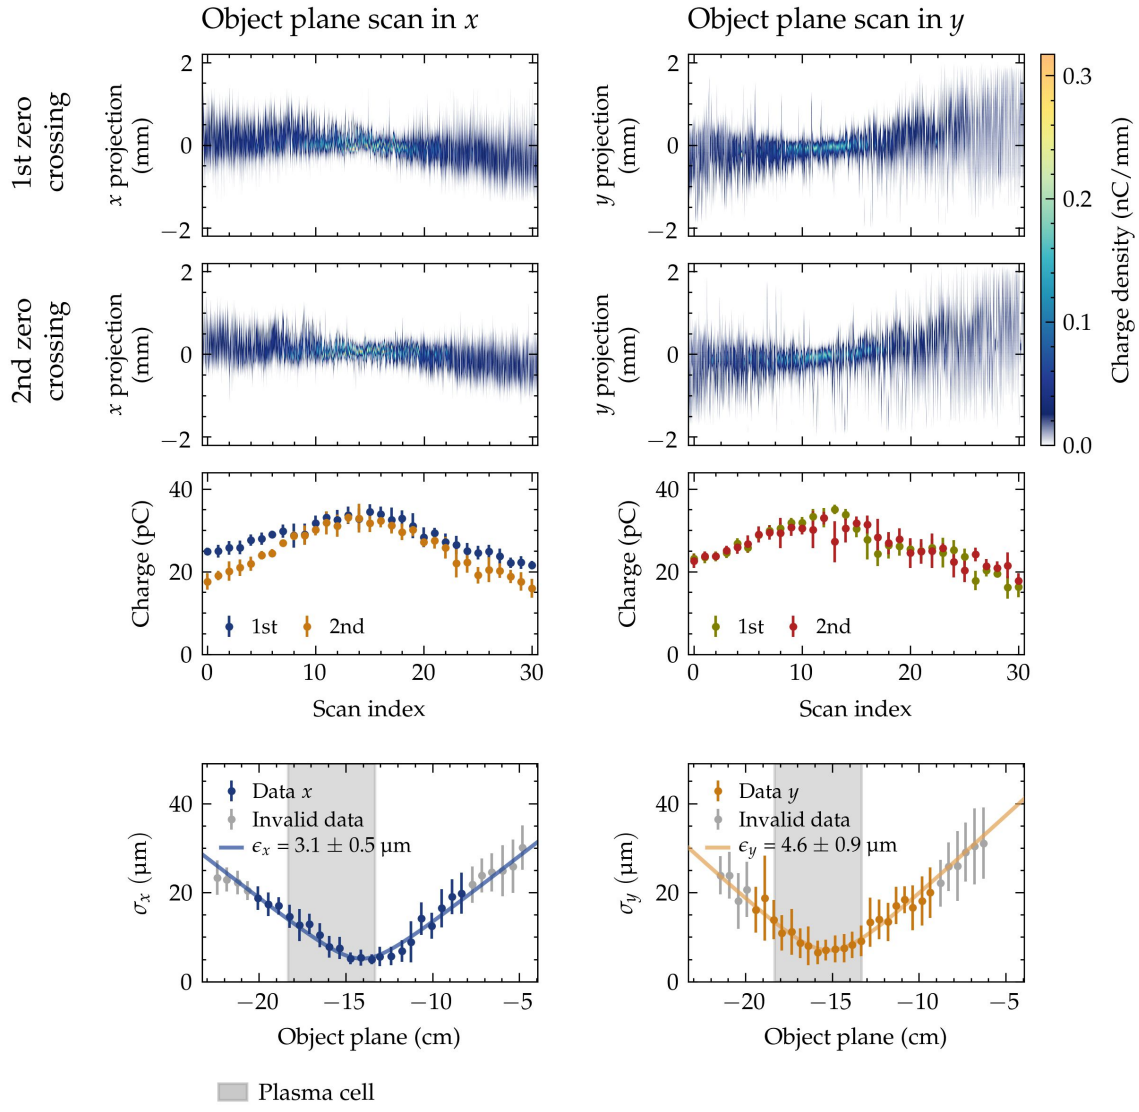

Figure 5.7.: Projected normalised emittance in x and y of accelerated bunches measured at 11FLFXTDS. Top plots: beam projections and measured charges at the screen. Bottom plot: calculated projected beamsizes and normalised emittances.

The result of the analysis is shown in the bottom plot of Figure 5.7, where the excluded scan steps are represented in grey. According to the beam-size evolution, the projected emittance in x is $\epsilon_{x,n} = 3.1 \pm 0.5$ µm, which compared to the measurement performed at the LEMS ($\epsilon_{x,n} = 2.4 \pm 0.2$ µm) exhibits an increase of ~ 30 %—even though a slight decrease would be expected as a result of charge losses even if the bunch were transversely matched into the plasma. In this plane the emittance growth associated with the finite chromatic amplitude W_x at the screen (cf. Equation 1.1.57) has values of $\Delta\epsilon/\epsilon_0 \lesssim 1$ % around the waist and reaches a maximum of ~ 5 % at the outermost scan steps considered in the analysis, and therefore it cannot explain the differences observed. Another source of emittance growth could be the mismatch of the beam to the plasma density at the entrance of the capillary (cf. Section 1.4.2), which according to the Courant-Snyder parameters and the rms energy spread calculated in the previous chapter can reach values up to $\epsilon_{n,sat}/\epsilon_{n,0} \approx 2.1$ for a density of $n_0 \approx 1 \cdot 10^{16}$ cm^{-3}. However, in these conditions, the decoherence length is equivalent to $L_{dc} \approx 9$ m and the effective emittance growth experienced in the 50-mm-long capillary can be considered negligible. On the other hand, the energy jitter resulting from beam-loading modulations does lead to changes in the Courant-Snyder parameters along the beamline that result in an increase of the averaged beam size measured at the screen and could indeed explain the apparent emittance growth. For the projected emittance in y a value of $\epsilon_{y,n} = 4.6 \pm 0.9$ µm is obtained, which represents a dramatic increase of a factor ~ 3.8 compared to the value of the noninteracted trailing bunch ($\epsilon_{y,n} = 1.2 \pm 0.1$ µm, cf. Section 4.4). In this case, the estimated emittance growth resulting from both the finite chromatic amplitude W_y at the screen and the mismatch of the beam to the plasma density has values very similar to that of the x axis and can neither account for the measured increase. The two remaining sources of emittance growth are the energy jitter and dispersion, which will be investigated in more detail in Section 5.2. Table 5.1 summarises the values of the projected emittance measured at the LEMS and at the 11FLFXTDS screen for the noninteracted and the accelerated trailing bunch.

Parameter	Units	**LEMS**		**11FLFXTDS**	
		plasma off	plasma on	plasma off	plasma on
$\epsilon_{x,n}$ (horizontal)	µm	4.8 ± 0.1	2.4 ± 0.2	3.7 ± 0.3	3.1 ± 0.5
$\epsilon_{y,n}$ (vertical)	µm	–	–	1.2 ± 0.1	4.6 ± 0.9

Table 5.1.: Summary of projected normalised emittances measured at the LEMS and the 11FLFXTDS screen.

One of the goals of the campaign is to estimate the accelerating-gradient homogeneity from the comparison of the sliced-energy-spread around the core of the accelerated bunches to that of the noninteracted ones. It is therefore necessary to perform a sliced analysis of the emittance measurements in x and y, for which several systematic errors have to be taken into account in addition to those considered in the projected-emittance case. On the one hand, the amount of energy jitter varies dramatically between the head and the tail of the bunch due to

beam loading modulations. This can be clearly appreciated in the projected energy spectrum measured at the ESPEC (cf. Figure 5.3), from which it becomes obvious that the mean energy of the head covers a range several times larger than that of the tail. On the other hand, as has been stressed in the analysis of the longitudinal-phase-space measurements (cf. Figure 5.2), the head and the tail of the bunch seem to be severely affected by strong correlations in the y–ξ and y'–ξ planes, which seriously compromise the ability to unambiguously assign a particular longitudinal slice to the vertical coordinate y at the screen. With these issues in mind and taking into account that the estimation of the accelerating-gradient homogeneity will be restricted to the core of the bunch—where the wakefield is most effectively beam loaded—a detailed interpretation of the sliced parameters along the whole bunch length is considered superfluous. Nonetheless, the results of a simple sliced analysis that ignores any sort of systematic error are shown in Figure 5.8.

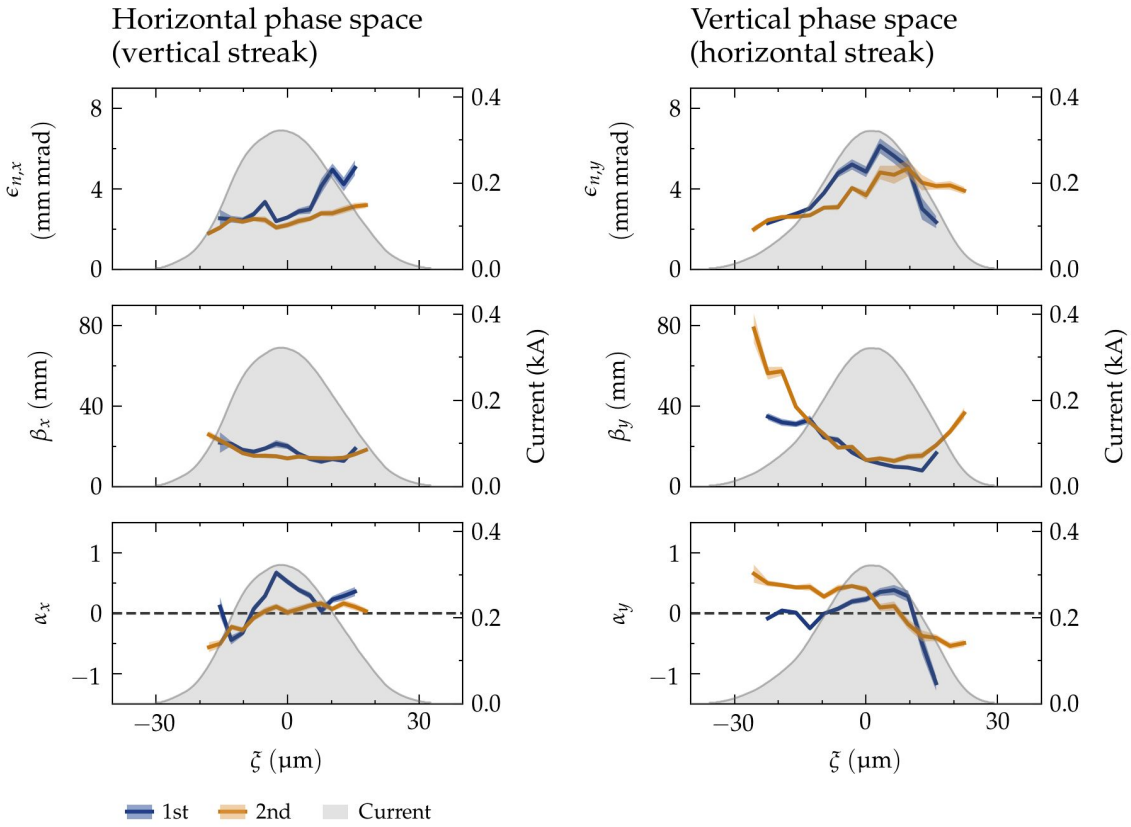

Figure 5.8.: Slice normalised emittance and slice Courant-Snyder parameters in x and y of accelerated trailing bunches measured at the 11FLFXTDS screen.

5.2. Measurement errors

A reliable estimation of the true slice energy spread is of paramount importance for the determination of the accelerating-gradient homogeneity. Therefore, on account of the issues

affecting the longitudinal-phase-space and the emittance measurements discussed in the previous section, an error analysis is carried out in order to evaluate to what extent the observed effects can be compensated.

The accelerated-bunch energy jitter induced by beam-loading modulations will be considered first, which due to the multi-shot nature of the emittance measurements and the limited number of shots sampled therein has to be considered from a statistical point of view. In second place, the effect of beam dispersion on the slice-emittance measurements will be analysed. Finally, the shearing effects distorting the longitudinal-phase-space measurements will be investigated in the framework of the matrix formalism and with the support of particle-tracking simulations performed with the code Ocelot [Agapov et al., 2014].

5.2.1. Energy jitter

In Section 5.1.2, the energy jitter induced by beam-loading modulations has been shown to be $\sigma_{\langle E \rangle} = 0.34$ % around a mean energy of $\langle \widetilde{E} \rangle_w = 1097$ MeV. Due to the high divergence of the beams produced in the PolariX-TDS campaign and the large distance between the plasma source and the screen 11FLFXTDS ($\Delta s \approx 33$ m), such energy variations have a large impact on the measured beam size. In the slice-emittance measurements only $N = 25$ shots per quadrupole setting were sampled, which does not allow to perform an analysis for individual energy bins. Therefore, the impact of the energy jitter on the measured emittance has to be considered statistically. The approach followed here is to simulate emittance measurements for a range of initial β_0 and α_0 parameters and introduce random errors in the nominal energy of the magnetic lattice on a single-shot basis with a normal distribution given by the energy jitter measured at the ESPEC. For each pair of Courant-Snyder parameters β_0 and α_0, 500 emittance measurements are simulated by transporting the beam envelope from the reference plane s_0 to the measurement screen and $N = 25$ samples are taken for each scan step. The beam size is calculated from the beta function at the screen and an emittance of $\epsilon_0 = 2$ µm that is kept fixed for all simulations. Additionally, for each shot a beam-size error of $\sigma = 5$ % is introduced to reproduce realistic measurement conditions.

Figure 5.9 shows the results obtained for the central-slice emittance in y, including the measured beta function β_{meas} (top left), the measured alpha α_{meas} (top right), the measured emittance ϵ_{meas} (bottom left) expressed as the relative error with respect to the emittance used in the simulations $(\epsilon_{meas} - \epsilon_0)/\epsilon_0$, and the relative error of the measured emittance (bottom right) expressed as $\sigma_{\epsilon_{meas}}/\epsilon_{meas}$, where $\sigma_{\epsilon_{meas}}$ is the error derived from the linear-least-squares fit used to determine the emittance. Several isometric curves are plotted in white on top of the colormap as a visual aid to the reader. From the maps obtained through simulation and the Courant-Snyder parameters calculated from the real measurement it is possible to assign a precise coordinate (β_0, α_0) to the true beam paramters at the reference plane s_0, which is indicated as a red dot in each of the plots. According to that, the emittance measured at the screen would be overestimated by 50 %, the measured beta $\beta_{y,meas} = 13 \pm 1$ mm would correspond to a real beta of $\beta_{y,0} = 10 \pm 2$ mm, and the alpha $\alpha_{y,meas} = 0.32 \pm 0.07$ to

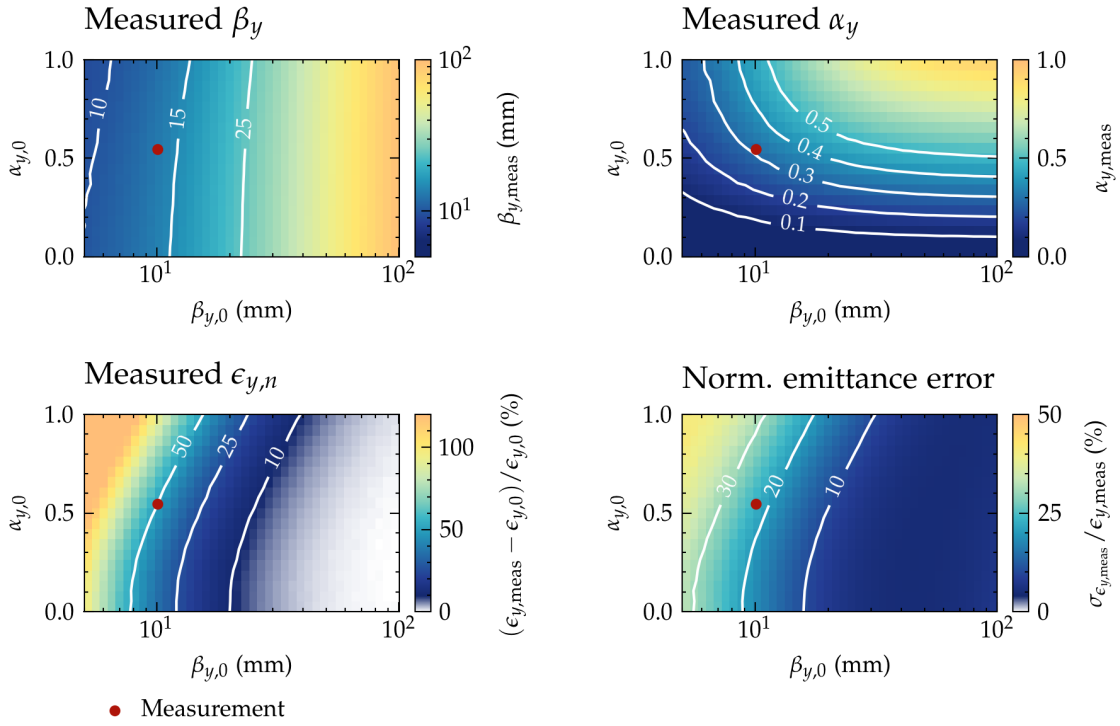

Figure 5.9.: Energy-jitter-induced statistical errors in slice-emittance measurements estimated by means of Monte Carlo simulations.

$\alpha_{y,0} = 0.55 \pm 0.35$.

The same method is applied to the projected emittances measured at the LEMS and at the 11FLFXTDS screen, the results of which are summarised in Table 5.2. As a consequence of the reduced distance from the plasma capillary to the LEMS ($\Delta s \approx 7$ m) the effect of the energy jitter at this station is substantially smaller than that at the 11FLFXTDS screen, and the measured emittance and Courant-Snyder parameters are very close to the estimated real values. At the 11FLFXTDS screen, in contrast, the effect is on the order of 25 % and 30 % for the beta function and the emittance, respectively, whereas in the case of the alpha $\alpha_{x,s0}$ the effect is mainly reflected in the associated error. Altogether, the corrected (real) values in the x–x' plane at the two stations show an excellent agreement, which supports the argument that the energy jitter is the main source of errors in this plane. In the y–y' plane, however, the corrected (real) emittance $\epsilon_{y,n} = 3.8 \pm 0.7$ µm is still a factor ~ 3.2 larger than that of the noninteracted trailing bunch $\epsilon_{y,n} = 1.2 \pm 0.1$ µm, which clearly indicates that the energy jitter alone is not enough the explain the disagreement between the two values. In order to account for these differences, the effect of beam dispersion on the emittance measurement is investigated in the following.

Parameter	Units	**LEMS**		**11FLFXTDS**	
		Measured	Real	Measured	Real
$\epsilon_{x,n}$	µm	2.4 ± 0.2	2.3 ± 0.2	3.1 ± 0.5	2.4 ± 0.4
$\beta_{x,s0}$	mm	12 ± 1	11 ± 1	15 ± 3	12 ± 2
$\alpha_{x,s0}$		0.0 ± 0.1	0.0 ± 0.1	0.0 ± 0.2	0.1 ± 4.9
$\epsilon_{y,n}$	µm			4.6 ± 0.9	3.8 ± 0.7
$\beta_{y,s0}$	mm			17 ± 4	13 ± 2
$\alpha_{y,s0}$				0.0 ± 0.2	0.0 ± 0.1

Table 5.2.: Comparison between measured and jitter-compensated (real) projected emittances for the measurements performed at the LEMS and at the 11FLFXTDS screen.

5.2.2. Beam dispersion

The slice-emittance measurements performed at the 11FLFXTDS screen enable a direct evaluation of the beam dispersion. The most straighforward approach is to calculate the beam tilt measured at the screen for each step of the object-plane scan. The tilt μ is obtained by fitting a linear function to the sliced centroids: $\langle u \rangle_\xi = u_0 + \mu \cdot \xi$, where $\langle u \rangle_\xi$ is the centroid in either x or y at the internal longitudinal coordinate of the bunch ξ. Additionally, in the fitting algorithm the centroids are weighted according to their corresponding charge. The results obtained for the scans in each of the two planes x–x' and y–y' at the first zero crossing are shown in Figure 5.10, which clearly indicate that the bunch tilt in x is kept constant along the whole scan, whereas it transitions from large positive values to large negative values in y. This information allows to derive two conclusions. On the one hand, since the quadrupole triplet used to scan the object plane is located immediately after the plasma chamber, it can be inferred that the observed beam tilt is originated at the beginning of the diagnostics

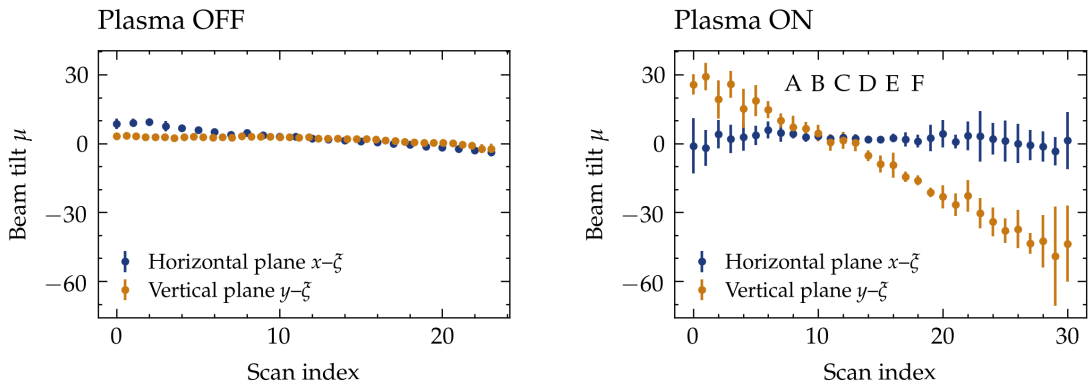

Figure 5.10.: Beam tilt in x and y measured at the 11FLFXTDS screen before (left) and after (right) plasma interaction. The letters A–F in the right plot indicate the object plane locations corresponding to the images shown in Figure 5.12.

beamline—i.e., either in the plasma or between the plasma and the quadrupole triplet. On the other hand, the tilt imprinted on the y–y' plane is a predominantly geometric effect independent of the amount of beam loading, the mean energy or the energy chirp, given that in each scan step—which contains bunches with all possible amounts of beam loading normally distributed around the working point—the relative error is small compared to the range of tilts covered along the scan.

To gain a better insight into the mechanism by which the tilt is imprinted on the bunch, the phase space of the sliced centroids is reconstructed at the output of the plasma capillary, which is shown in Figure 5.11. The results for the y–y' plane (bottom plots) clearly indicate that the trailing bunch exits the plasma capillary with a strong y'–ξ correlation. This correlation could be caused by spatial and angular misalignments between the driver and the trailing bunch at the entrance of the plasma capillary, which would lead different slices to oscillate transversely around the longitudinal axis of the wakefield with different phases and amplitudes. This hypothesis is supported by the analysis of the noninteracted double-bunch structure (cf. Figure 4.19 and 4.21), which indicates that the trailing bunch has an overall angular misalignment in y' with respect to the driver of about $\langle y' \rangle_w - \langle y' \rangle_d \approx 0.4$ mrad. As a result of the transverse dynamics driven by the strong focusing fields inside the plasma, the angular spread among different slices of the trailing bunch increases from $\sigma_{y',w} \lesssim 0.1$ mrad before plasma interaction (cf. Figure 4.21) to $\sigma_{y',w} \approx 0.5$ mrad after acceleration (cf. Figure 5.11).

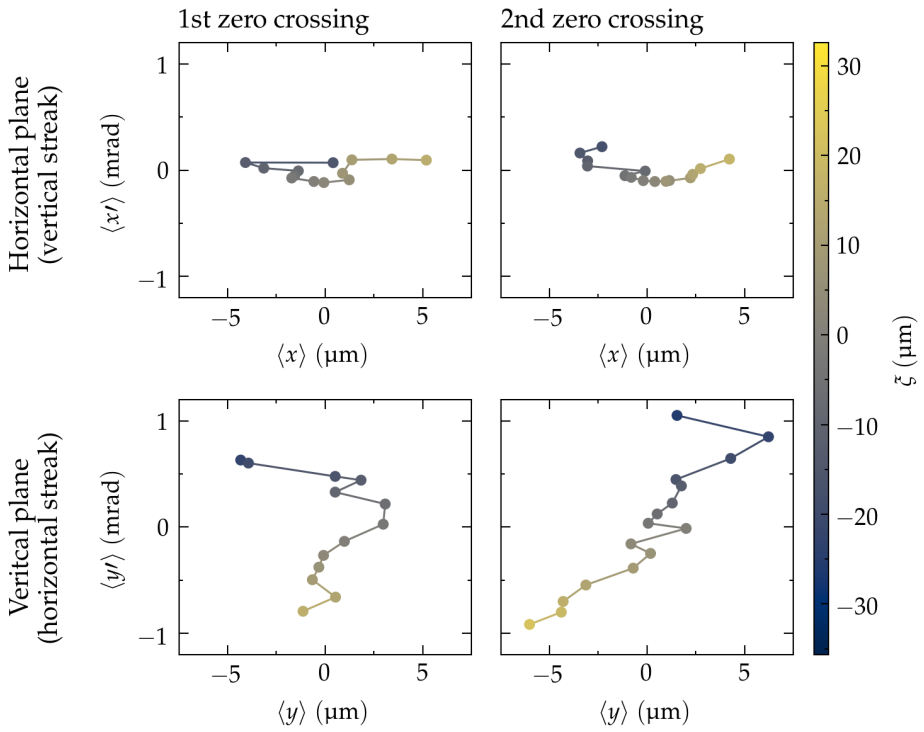

Figure 5.11.: Slice transverse centroid offets of accelerated bunches at the exit of the plasma cell reconstructed from object-plane scans performed at the 11FLFXTDS screen.

Figure 5.12 shows six images from scan steps corresponding to object-plane locations around the waist, which are are labeled with the letters A, B, C, D, E, F and are indicated below the top axis of Figure 5.10. The ability to perform a sliced analysis to a tilted beam allows to partly compensate dispersive effects when calculating the emittance, since for significant tilts as those observed in the images the sliced beam size is smaller than the projected beam size. If the tilt is conceived as a purely transverse displacement of individual slices depending on their longitudinal coordinate, the sliced beam size should not change at all. However, the finite time resolution achieved with the TDS translates into an increase of the sliced beam size measured at the screen. Taking into account that the time resolution changes along the beam if the beam size is not constant over the whole bunch length—which is indeed expected to be the case—the retrieval of the real sliced beam size becomes an extremely complex problem, that falls out of the scope of this thesis. In general, however, it can be claimed that the measured sliced beam size increases with the tilt and, due to the correlation between the beam tilt and the phase advance along the scan, the beam size evolution around the waist is steeper than if no dispersion was present, which leads to an increase of the measured emittance—i.e., the measured slice emittance represents an upper bound of its real value $\epsilon_{\text{real}} \leq \epsilon_{\text{meas}}$. Regarding the Courant-Snyder parameters, on the other hand, the effect of beam dispersion has its largest impact on the reconstructed beta function at the waist, which is smaller than the real one $\beta_{\text{real}} \geq \beta_{\text{meas}}$, whereas the alpha remains practically unaffected $\alpha_{\text{real}} \approx \alpha_{\text{meas}}$.

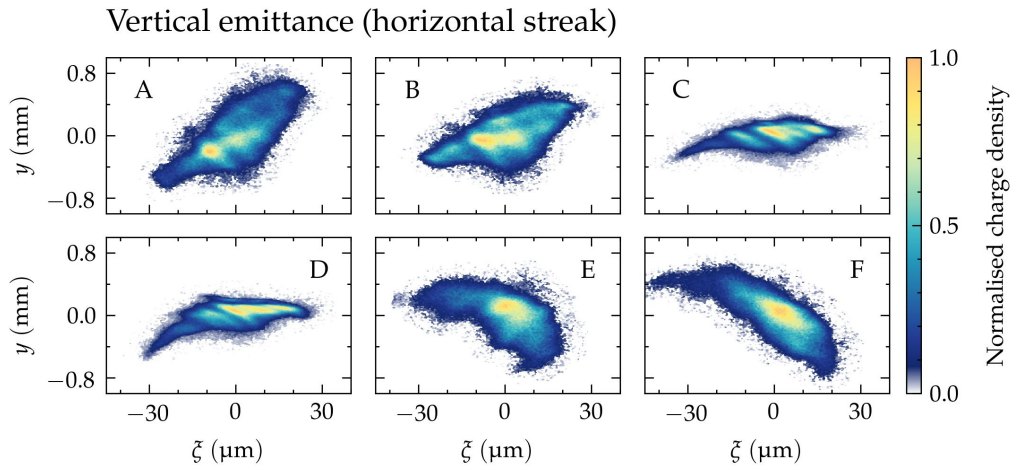

Figure 5.12.: Image samples taken from a slice-emittance measurement in y performed at the 11FLFXTDS screen. The images correspond to object planes around the exit of the plasma cell.

5.2.3. Longitudinal-phase-space shearing effects

In Section 5.1.1 the observation of shearing effects that distort the longitudinal phase space has been described. In the following, these effects are investigated by means of the matrix formalism in the linear approximation (cf. Section 1.1.1) and with the support of particle-tracking simulations performed with Ocelot.

The mathematical analysis focuses on the examination of the correlation between the spatial coordinates x and y at the screen rather than the correlation between the variables that they represent δ and ξ. The beam matrix $\boldsymbol{\Sigma_0}$ is transported from the TDS to the longitudinal-phase-space screen according to the usual formula:

$$\boldsymbol{\Sigma_s} = (\boldsymbol{R_{TDS\to s}} \cdot \boldsymbol{M_{TDS}}) \cdot \boldsymbol{\Sigma_0} \cdot (\boldsymbol{R_{TDS\to s}} \cdot \boldsymbol{M_{TDS}})^T , \tag{5.2.1}$$

where $\boldsymbol{M_{TDS}}$ and $\boldsymbol{R_{TDS\to s}}$ are the transport matrices of the TDS and from the TDS to the screen, respectively (cf. Section 1.3.2). After transport, the term of the beam matrix correlating x and y can be expressed as:

$$\begin{aligned}\langle x_s y_s \rangle = R_{16} \Bigg[& \mathcal{K} R_{33} \langle y_0 \rangle + \mathcal{K} \left(\frac{3L}{2} R_{33} + R_{34} \right) \langle y_0 y_0' \rangle + \frac{\mathcal{K} L}{2} \left(L R_{33} + R_{34} \right) \langle y_0'^2 \rangle \\ & + \frac{\mathcal{K}^3 L}{6} \left(\frac{R_{33} L}{2} + R_{34} \right) \langle \xi_0^2 \rangle + \mathcal{K} \left(\frac{L}{2} R_{33} + R_{34} \right) \langle \xi_0 \delta_0 \rangle \Bigg] .\end{aligned} \tag{5.2.2}$$

If thin slices in ξ are taken, so that $\langle \xi_0^2 \rangle \to 0$ and $\langle \xi_0 \delta_0 \rangle \to 0$, and assuming a lattice optic that fulfils the parallel-to-point imaging condition $R_{33} \ll R_{34}$, Equation 5.2.2 is reduced to:

$$\begin{aligned}\langle x_s y_s \rangle_\xi &= R_{16} \mathcal{K} R_{34} \left(\langle y_0 y_0' \rangle_\xi + \frac{L}{2} \langle y_0'^2 \rangle_\xi \right) \\ &= D_x(s, s_0)\, S\, \epsilon_y(\xi) \left(-\alpha_{y,TDS}(\xi) + \frac{L}{2} \gamma_{y,TDS}(\xi) \right) ,\end{aligned} \tag{5.2.3}$$

where ξ identifies the particles corresponding to a slice at a particular longitudinal bunch coordinate. Eq. 5.2.3 reveals that if the parameters $\alpha_{y,TDS}$ or $\gamma_{y,TDS}$ at the TDS for a given slice ξ are not negligible, a correlation between x and y is imprinted into the slice and is subsequently magnified by the dispersion $D_x(s, s_0)$, resulting in a kind of shearing of the longitudinal phase space measured at the screen that depends on the zero crossing. In general, the typical lattice optic used for TDS diagnostics fulfils the conditions $\alpha_{y,TDS} \approx 0$ and $\gamma_{y,TDS} \approx 1/\beta_{y,TDS} \ll 1$ and this problem should not appear. However, if the beam is not well matched to the design optics, these shearing effects might become a problem. This is exactly the case of the measurements performed with the accelerated trailing bunch, since for highly divergent beams it is difficult to control the beam envelope in order to set $\alpha_{y,TDS} = 0$ at the location of the TDS—even more taking into account that the experiment is affected by different sources of jitter.

In the following the analytical formulation presented above is reproduced by means of particle tracking simulations with Ocelot. In a first step, a realistic particle distribution containing internal beam correlations is generated, for which the measurements of the plasma-

accelerated beam presented in Section 5.1 are used. For the purpose of this investigation the accuracy of the reconstruction plays a minor role and the limitations of the techniques used therein—viz. two-point-tomography, compensation of TDS-induced effects and compensation of measurement errors—will not be discussed. Each longitudinal slice of the reconstructed bunch is tracked separately down to the longitudinal-phase-space screen using two different lattice optics (cf. Figure 5.13): 1) the one corresponding to the real measurements (solid lines), for which the beam has an alpha of $\alpha_{y,TDS} \approx 2.5$ at the TDS, and 2) an optimised one in which the beam is matched to fulfil the condition $\alpha_{y,TDS} = 0.0$ (dashed lines). In the latter, additionally, the time and energy resolution have been improved, but these changes do not have any effect on the shearing of the longitudinal-phase-space since the condition $R_{33} \ll R_{34}$ is fulfilled and the shear parameter S and the dispersion D_x have similar values in both cases—i.e. the shearing effects depend solely on the Courant-Snyder parameters $\alpha_{y,TDS}$ and $\gamma_{y,TDS}$ at the TDS.

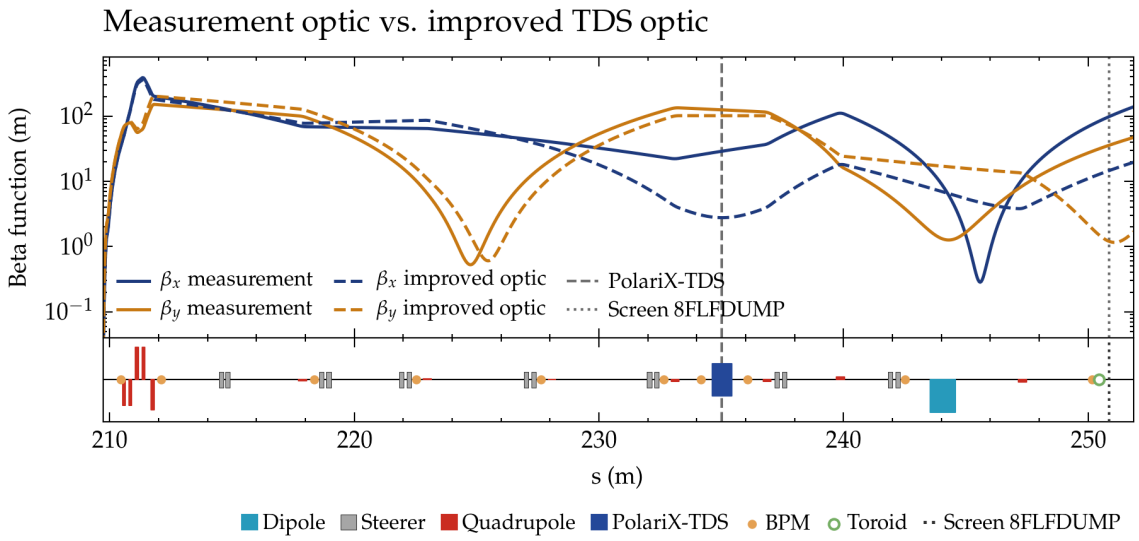

Figure 5.13.: Dependence of longitudinal-phase-space shearing effects on the lattice optic. The lattice optic in the actual experimental conditions is indicated by solid lines, whereas the improved lattice optic matched to the beam parameters is indicated by dashed lines.

The tracked slices at the location of the screen are shown in Figure 5.14 for both the measurement optic (left) and the improved optic (right), and the first (top) and second (bottom) zero crossing. The particles are coloured according to the location of their corresponding slice to better appreciate the shearing effects. With the measurement optic, a clear negative correlation between x and y is observed in all slices. In these conditions, the conventional time-resolved analysis of the bunch at the longitudinal-phase-space screen is severly compromised, because a vertical slice at a location y on the screen does not correspond to a well defined longitudinal bunch coordinate ξ, but contains information about multiple slices extending over a large range of longitudinal bunch locations $\Delta\xi$, which span over an interval that largely exceeds the time resolution $\Delta\xi > R_\xi$. Additionally, the relative location of the slice centroids—which is determined by the internal bunch correlations y–ξ and y'–ξ, and the zero crossing at which the TDS is operated—can lead to a strong modulation of the

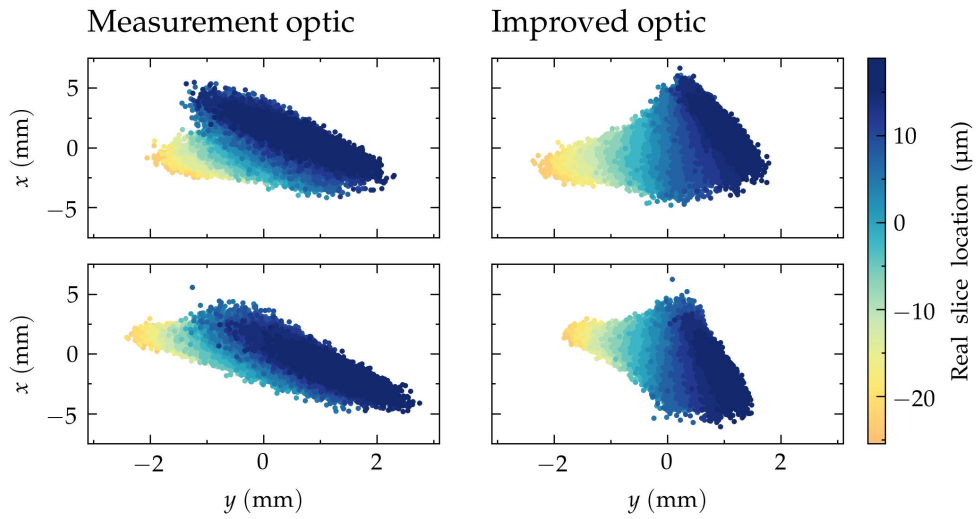

Figure 5.14.: Particle tracking simulations reproducing longitudinal-phase-space shearing effects observed at the 8FLFDUMP screen for the lattice optic used in the actual measurements (left) and an improved lattice optic that would minimise the effects (right).

calculated slice energy spread. With the improved optic, in contrast, the correlation between x and y has been removed for most of the slices and only a residual effect is observed at the head and tail of the bunch due to the fact that the Courant-Snyder parameters vary along the bunch and the condition $\alpha_{y,TDS} = 0$ at the TDS cannot be fulfilled for all the slices simultaneously.

These effects cannot be compensated analytically and must be minimised experimentally by precisely matching the beam to an appropriate lattice optic. On account of the sources of jitter affecting the plasma-accelerated beam, however, the routines required to match the beam are extremely cumbersome and time consuming, and due to the limited availability of beam time during the PolariX-TDS campaign the result of the optimisation was not as optimal as it could have been.

5.3. Accelerating-gradient homogeneity

The experiments of the PolariX-TDS campaign are performed in the blowout regime, in which an ion column is left behind the driver, providing a cylindrically symmetric transverse focusing force that varies linearly with the radius and is constant in ξ over the length of the trailing bunch—assuming that the ion motion in an argon plasma is negligible at the time of arrival of the trailing bunch, which is estimated to be $\langle t \rangle_{witness} - \langle t \rangle_{driver} \approx 0.6$ ps [Rosenzweig, Cook et al., 2005]. According to the Panofsky-Wenzel theorem (cf. Section 1.4.2) the presence of a longitudinally constant transverse-focusing force $\partial F_r / \partial \xi = 0$ implies that the accelerating force is transversely constant $\partial F_z / \partial r = 0$ and, consequently, that the slice energy spread of the bunch is conserved after acceleration. Additionally, if the

longitudinal variation of the accelerating gradient is zero $\partial E_z/\partial\xi = 0$—which corresponds to optimal beam-loading conditions—all forces inside the volume around the trailing bunch are homogeneous, i.e., constant accelerating force accross the whole volume and cylindrically symmetric transverse focusing force varying linearly with radius.

The goal of this section is to provide an estimate of the homogeneity of the accelerating gradient produced in the plasma wakefield while taking into account the limitations of the measurements previously discussed. The figure of merit used to that end is the *field nonuniformity* σ_ζ, which is defined as:

$$\sigma_\zeta = \frac{\sqrt{\sigma_{E,\mathrm{on}}^2 - \sigma_{E,\mathrm{off}}^2}}{\Delta E}, \tag{5.3.1}$$

where $\sigma_{E,\mathrm{on}}$ is the true rms slice energy spread of the accelerated bunch, $\sigma_{E,\mathrm{off}}$ is the true rms slice energy spread of the noninteracted bunch, and ΔE is the energy gain—all terms given in units of absolute energy (e.g. MeV). The field nonuniformity σ_ζ expresses the relative error within which the longitudinal accelerating-gradient can be considered constant over the transverse extent of the trailing bunch. This statement, however, is only rigorously true in the assumption that: 1) the accelerating gradient E_z is constant over the longitudinal extent $\Delta\xi$ of the energy slice considered; 2) the energy spread itself is not affected by any other physical process. On the other hand, a single shot comparison between the slice energy spread before and after acceleration is not possible, and a statistical analysis of ensembles of independent measurements is required. Consequently, both the shot-to-shot stability of the acceleration process and measurement errors must be taken into account when using Equation 5.3.1 as a figure of merit of the actual accelerating-gradient homogeneity.

The true slice energy spread of the noninteracted bunch and the energy gain have been calculated in Sections 4.4.2 and 5.1.2, respectively. In the following, the calculation of the true slice energy spread of the accelerated bunch is addressed. After that, the acceleration-gradient homogeneity is calculated and the accuracy and an interpretation of the obtained result is provided.

5.3.1. True slice energy spread

Similar to the case of the noninteracted bunch, the true rms slice energy spread of the accelerated bunch σ_E is calculated according to the formula:

$$\sigma_E = \sqrt{\sigma_M^2 - \sigma_{IES}^2 - R_\delta^2}, \tag{5.3.2}$$

where σ_M is the measured slice energy spread, σ_{IES} is the TDS-induced energy spread, and R_δ is the energy resolution at the longitudinal-phase-space screen. The screen resolution R_{scr} has been excluded from the formula, since in the experimental conditions of the PolariX-

TDS campaign it has a value of $R_{scr} \approx 15$ keV (cf. Section 2.2) and is negligible compared to the other terms. For each of the terms on the r.h.s. of Equation 5.3.2, an ensemble of $N = 2 \cdot 10^6$ random values is generated describing the probability distribution obtained in the measurement, from which the resulting distribution of the true slice energy spread σ_E is subsequently calculated.

On account of the analysis performed in the previous two sections, the accuracy with which the different slice-energy-spread terms can be determined has to be carefully assessed. Three aspects will be considered in the following: 1) the effect of charge loss on the slice energy spread; 2) errors associated to the slice-emittance measurements, which affect the accuracy of the estimated energy resolution and the estimated TDS-induced energy spread; 3) longitudinal-phase-space shearing effects, which introduce distortions in the measured energy spread.

Charge loss. In Section 5.1, it has been seen that the charge loss produced by both the imperfect coupling of the bunch into the wakefield structure and transmission problems leads to a reduction of the beam emittance by ~ 50 %. This charge loss is a purely geometric effect by which the particles with the largest emittance are lost in either the plasma background or the beamline pipes. Since there is no correlation between the single-particle emittance and the particle energy, it is assumed that the slice energy spread is not affected by the charge loss.

Slice-emittance measurement errors. In order to calculate the energy resolution R_δ, the beam parameters retrieved from the slice-emittance measurement in the x–x' plane are required. Since the measurement is only affected by energy jitter—which can be fully compensated analytically—the resolution can be confidently determined to be $R_\delta = 0.60 \pm 0.04$ MeV. It is worth mentioning here that, due to the difficulty to match the beam to the lattice optic, the value obtained is larger than the targeted value of $R_\delta \approx 0.20$ MeV. The measurement in the y–y' plane, in contrast, is not only affected by energy jitter but also by beam dispersion, which cannot be compensated and affects both the measured emittance and the measured Courant-Snyder parameters. In regard to the emittance, as discussed in Section 5.1, no physical process has been identified that could lead to a real emittance growth. Therefore, in order to estimate the TDS-induced energy spread, it is justified to use the slice emittance of the noninteracted trailing bunch—and thus avoid the measurement errors associated to beam dispersion. Regarding the Courant-Snyder parameters, however, the measurements performed with the accelerated bunch constitute the only source of information available. Due to beam dispersion, the calculated beta function at the reconstruction plane—i.e., the capillary output—is known to represent a lower bound of its real value, which becomes an upper bound after propagation to the location of the TDS. Thus, since the beta function makes the largest contribution to the TDS-induced energy spread, the true energy spread obtained with Equation 5.3.2 will represent a lower bound of its real value. On account of these considerations, the induced energy spread is estimated

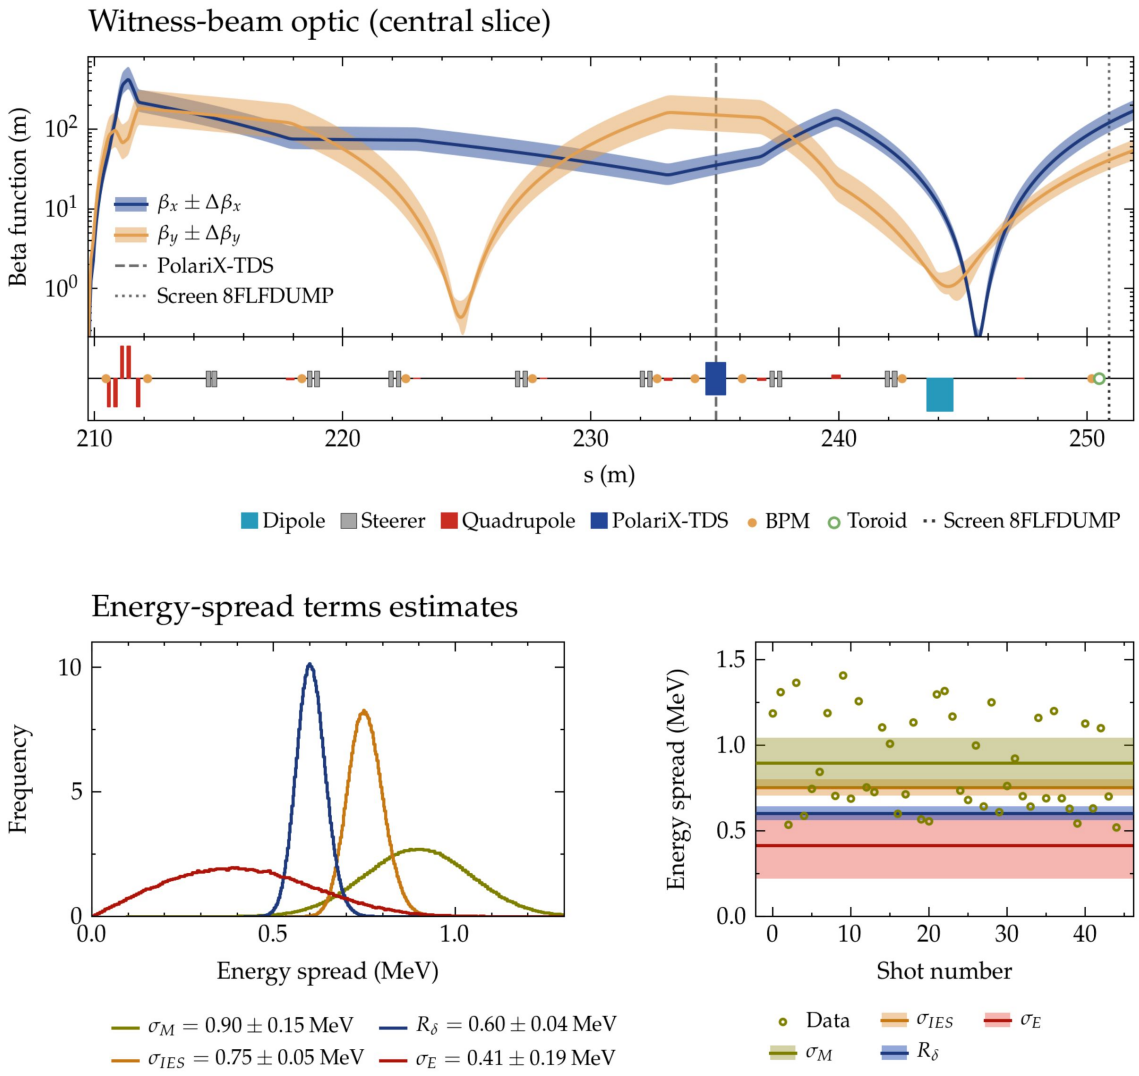

Figure 5.15.: Retrieval of the true energy spread of the central slice of plasma-accelerated trailing bunches. See text for details.

to be $\sigma_{IES} \leq 0.75 \pm 0.05$ MeV. The evolution of the beta function in x and y corresponding to the longitudinal-phase-space measurements of the accelerated trailing bunch are shown in the top plot of Figure 5.15, and the probability distributions of the energy resolution R_δ (blue) and the TDS-induced energy spread σ_{IES} (orange) are shown in the bottom-left plot.

Longitudinal-phase-space shearing effects. The longitudinal-phase-space shearing effects represent the most critical errors affecting the measurement of the slice energy spread, since they can neither be quantified nor compensated analytically. The approach followed in this thesis is to ignore these effects, calculate the energy spread of the central slice by means of a conventional time-resolved analysis using a slice size equivalent to the resolution of the measurement R_ξ, and consider that as soon as the true energy spread resulting from

the application of Equation 5.3.2 is not physically impossible—i.e., that the argument in the square root is not negative—the obtained value will be a rough representation of reality. The analysis of the measured data is restricted to the central slice of the $N \approx 50$ bunches ($N \approx 25$ for each zero crossing) nearest to the working point, for which: 1) the chirp of the accelerated bunch is close to that of the noninteracted bunch $h \approx 0$ and the condition $\partial E_z/\partial\xi \approx 0$ is therefore fulfilled; 2) the calculated energy resolution and TDS-induced energy spread are most accurate—since they depend on slice-emittance measurements that are evaluated in terms of statistical averages around the working point. The computed mean value and standard deviation of the measured samples are $\sigma_M = 0.90 \pm 0.15$ MeV. The obtained standard deviation represents a relative error of $\sigma_{\sigma_E} \approx 17$ %, which is roughly two times larger than that obtained for the noninteracted bunch. This can be associated to: 1) larger differences between the measurements performed at the 1st and 2nd zero crossings due to TDS-induced longitudinal-phase-space distortions; 2) larger instability of the beam parameters due to the acceleration process.

The bottom-left plot of Figure 5.15 shows the probability distribution of all the computed slice energy spread terms: the measured energy spread $\sigma_M = 0.90 \pm 0.15$ MeV (green), the TDS-induced energy spread $\sigma_{IES} = 0.75 \pm 0.05$ MeV (orange), the energy resolution at the screen $R_\delta = 0.60 \pm 0.04$ MeV, and the true slice energy spread $\sigma_E = 0.41 \pm 0.19$ MeV resulting from the application of Equation 5.3.2. The same results are shown in the bottom-right plot together with the data samples selected for the computation. Similar to the case of the noninteracted bunch, the relative error of the true-slice-energy-spread distribution ($\sigma_{\sigma_E} \approx 46$ %) is several times larger than that of the measured one ($\sigma_{\sigma_M} \approx 17$ %), which results from the contribution of the errors associated to the estimated TDS-induced energy spread and energy resolution.

5.3.2. Accelerating-gradient-homogeneity calculation

The accelerating-gradient homogeneity is calculated from the slice-energy-spread of the accelerated $\sigma_{E,\text{on}} = 0.41 \pm 0.19$ MeV and the noninteracted trailing bunch $\sigma_{E,\text{off}} = 0.38 \pm 0.12$ MeV obtained in Sections 5.3.1 and 4.4, respectively, together with the energy gain of the accelerated bunch around the working point $\Delta E = 44 \pm 2$ MeV (cf. Section 5.1.2). If the expectation values of the plasma-on and the plasma-off cases are directly compared, a field nonuniformity of $\sigma_\zeta = 0.35$ % is obtained, which is a remarkable result that implies that the relative energy spread only increases by 3.6 %—from $\sigma_{\delta,\text{off}} \approx 0.036$ % to $\sigma_{\delta,\text{on}} \approx 0.037$ %. As mentioned in the introduction of this section, however, a more rigurous analysis requires to perform a statistical analysis by calculating the field nonuniformity for ensembles of samples reproducing the probability distributions obtained experimentally. This is shown in Figure 5.16, with the true-slice-energy-spread probability distributions on the left and the field nonuniformity calculated for $N = 2 \cdot 10^6$ samples expressed in terms of percentiles on the right. According to the results, the longitudinal accelerating-gradient is transversely homogeneous to within 0.8 % (1.5 %) at an interval of confidence of 68 % (95 %). These

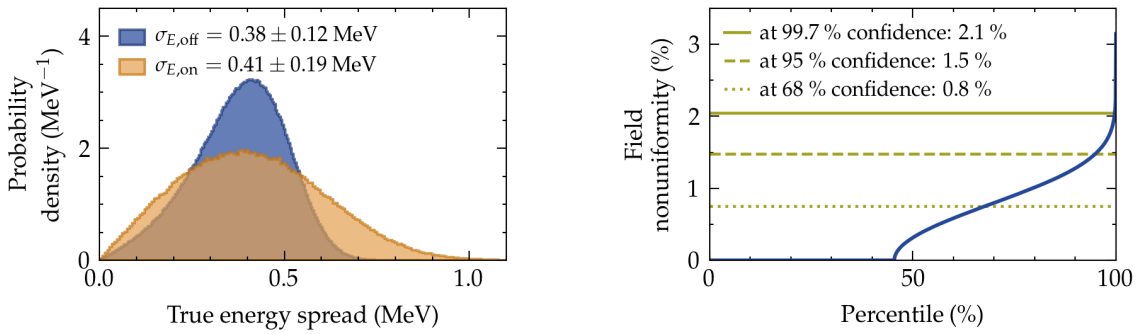

Figure 5.16.: Statistical calculation of the accelerating-gradient homogeneity from time-resolved measurements performed with the PolariX-TDS.

values provide upper bounds based on measurement uncertainties, but, as can be observed, there is also a high chance (at least 46 % of the events) that the field nonuniformity is actually 0 %. In the following, the meaning of the obtained values is shortly discussed.

Considering a single incoming bunch with a slice energy spread equivalent to the mean value of the distribution $\sigma_{E,\mathrm{off}} = 0.38$ MeV and being accelerated by $\Delta E = 44$ MeV, a field nonuniformity of $\sigma_\zeta = 0.8$ % would result in an absolute energy-spread increase of 36 %, which translates into an increase of the relative energy spread from $\sigma_{\delta,\mathrm{off}} \approx 0.036$ % to $\sigma_{\delta,\mathrm{on}} \approx 0.047$ %—i.e., a relative increase of 31 %. Scaling the numbers to the case of doubling the energy gain, the absolute energy spread would increase by a factor of 22, which represents an increase of the relative energy spread from $\sigma_{\delta,\mathrm{off}} \approx 0.036$ % to $\sigma_{\delta,\mathrm{on}} \approx 0.40$ %, leading to an severe degradation of the beam quality. In the experimental conditions of the PolariX-TDS campaign, if the relative slice energy spread had to be preserved, the field nonuniformity should have a value of at most $\sigma_\zeta \approx 0.25$ %. In order to preserve the relative slice energy spread while doubling the energy gain, on the other hand, the value should be decreased down to $\sigma_\zeta \approx 0.06$ %.

As pointed out before, however, due to the statistical nature of the field-nonuniformity calculation, the estimated value is strongly dominated by the shape of the true-slice-energy-spread probability distribution, which is in turn affected by both the stability of the acceleration process and experimental uncertainties. Therefore, the value obtained can be regarded as a figure of merit of the *ability* to retrieve information about the accelerating-gradient homogeneity rather than a measure of the accelerating-gradient homogeneity itself. The following considerations illustrate the point: even if the accelerating gradient was really perfectly homogeneous and the true-slice-energy-spread probability distributions of both the uninteracted and the accelerated bunches were exactly the same, a relative error in the probability distribution of $\sigma_{\sigma_E} \approx 6$ % would only allow to estimate a field nonuniformity of $\sigma_\zeta \approx 0.25$ % at an interval of confidence of 68 %. This could be interpreted as an excellent result that guarantees the preservation of the relative energy spread, but does indeed not reflect the even better quality of the acceleration process.

The fact that experimental uncertainties hinder a precise evaluation of the actual accelerating-gradient homogeneity from Equation 5.3.1, does not preclude the presence of mechanisms that indeed distort the (effective) transverse homogeneity of the accelerating gradient. One possibility would be that the experiment does not take place in the blowout regime, but according to the measured beam parameters this is not expected to be the case. A more plausible source of distortions could be related to spatial and/or angular misalignments between the driver and the trailing bunch. Such misalignments cause the trailing bunch to undergo transverse oscillations inside the wakefield structure, which can lead the beam to sample regions of the accelerating gradient in which the field is strongly nonhomogeneous—as, for instance, near the walls of the structure. Furthermore, the interaction of the oscillating beam with the walls, results in a geometric deformation of the wakefield structure that could translate into a radial dependence of the accelerating-gradient amplitude. A third source of distortions could be that the accelerating gradient is not constant along the longitudinal extent of the slice analysed $\partial E_z/\partial\xi \neq 0$, whose effect on the slice energy spread would be amplified in combination with the geometric deformation of the wakefield structure. Even though a theoretical derivation of the order of magnitude of these effects is out of the scope of this thesis, the field-nonuniformity calculation allows to assert that their combined effect results in distortions of the effective accelerating-gradient homogeneity of at most 0.8 % at an interval of confidence of 68 %. A further consideration of the presence of these effects during the PolariX-TDS campaign will be discussed in the next section with the support of particle-in-cell simulations.

One of the main problems with the calculations presented above is that the true slice energy spread of the accelerated bunch is strongly affected by systematic errors—most prominently associated to longitudinal-phase-space shearing effects—and that the calculated probability distribution has an uncertainty which is difficult to estimate. Therefore, an alternative calculation is described in the following that provides a reliable upper bound to the field nonuniformity.

A measure of the projected energy spread of the accelerated bunch around the working point has been obtained at the LEMS (cf. Figure 5.4). This measurement can be considered to provide an absolute upper bound of the true slice energy spread for two reasons. On the one hand, the head and the tail of the bunch are included in the computation of the rms energy spread, which are known to extend towards lower and higher energies with respect to the mean beam energy, respectively, and therefore broaden the measured energy spectrum. On the other hand, the presence of beam dispersion in the y–y' plane (cf. Section 5.2.2)—which corresponds to the dispersive plane of the LEMS—leads to an overestimated measure of the real projected rms energy spread. Thus, by using the projected energy spread of the accelerated bunch measured at the LEMS instead of the sliced energy spread computed from the longitudinal-phase-space measurements, an upper bound for the field nonuniformity is obtained. The probability distributions used in the new calculation and the field nonuniformity expressed in terms of percentiles are shown in Figure 5.17, according to which the longitudinal accelerating-gradient is transversely homogeneous to within at least 1.8 % with an interval of confidence of 68 %. Even if this value would correspond to

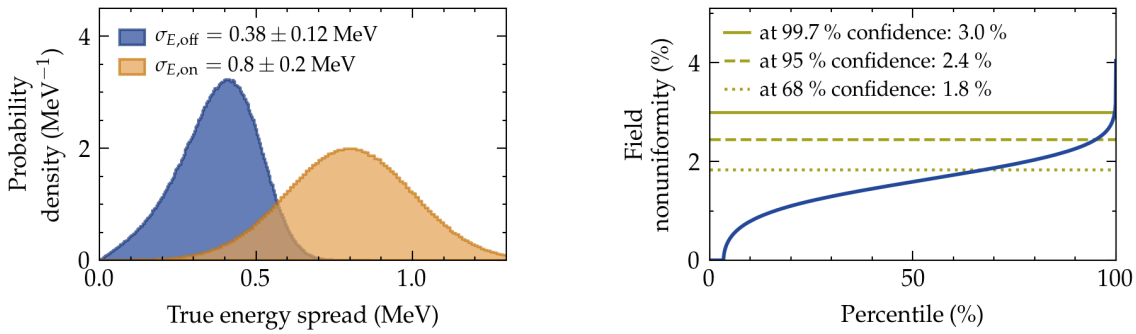

Figure 5.17.: Statistical calculation of the accelerating-gradient homogeneity from projected measurements performed at the LEMS (accelerated trailing bunch) and time-resolved measurements performed with the PolariX-TDS (noninteracted trailing bunch).

an acceleration process with a rather poor performance, the result represents a significant improvement with respect to that published in 2016 by C.E. Clayton et al., who indirectly measured the radially focusing field was approximately constant in the accelerating portion of the wakefield $\partial_\xi F_r \approx 0$ to within ± 3 % and therefore, from the Panofsky-Wenzel theorem, concluded that the accelerating-gradient was also constant $\partial_r F_z \approx 0$ to a similar degree of uncertainty [Clayton et al., 2016].

5.4. Particle-in-cell simulations

In this section, particle-in-cell (PIC) simulations are performed with the code HiPACE++ [Diederichs et al., 2021] to verify the physical interpretation of the acceleration process provided in the previous sections. The simulations require two main inputs: an incoming double-bunch structure and a plasma with definite properties. The phase space of the incoming double-bunch structure has already been thoroughly characterised in Section 4.4, whereas the plasma parameters are only known to be found within an approximate range of values. In the time scale of the acceleration process and from the point of view of the electron bunch, the transverse extent of the plasma is virtually unbounded. Therefore, only the longitudinal density profile needs to be determined, which for the cell geometry used in the experiment has the approximate shape of a flat-top with a tapered transition towards the vacuum background outside the capillary [Garland et al., 2021]. In order to determine the plasma density at the center of the cell, separate measurements are carried out, which are discussed in Section 5.4.1. The remaining parameters—i.e., the extent of the flat-top and the effect of the density taper—are estimated by means of a Bayesian Optimisation (BO) algorithm using an objective function based on the accelerated-bunch parameters measured around the working point. This procedure is presented in Section 5.4.2 together with a discussion on the optimisation outcome. Finally, the beam-loading modulations observed in the experiment are reproduced with an additional set of PIC simulations, which is described in Section 5.4.3.

5.4.1. Plasma-density measurements

In order to determine the plasma density at the center of the capillary, separate measurements are carried out reproducing the experimental conditions of the campaign—viz. capillary filled with Argon at a backing pressure of 10.7 mbar, discharged by an electrical pulse with a voltage of 25 kV, a current of 500 A and a length of 400 ns. The plasma evolution is measured by doping Argon with 3% of Hydrogen and observing the spectral broadening of the H_α emission line at different times after discharge with a dedicated spectrometer [Garland et al., 2021]. Figure 5.18 shows the result of the measurement, for which a range of 6.0 µs was scanned in time steps of $\Delta t = 0.2$ µs, recording densities down to the resolution limit of the spectrometer—which is just below 10^{14} cm^{-3}. The plasma density has its peak at $t = 0.6$ µs right after the end of the main current pulse and starts an approximate exponential decay after that time. A second peak is recognised at $t = 1.2$ µs, which is the result of late current-pulse oscillations passing through the plasma. From this time onwards, the density undergoes slow fluctuations, that can be associated to two main factors: 1) recombination of the plasma; 2) diffusion of hot plasma from the capillary. The plasma density corresponding to the working point of the PWFA experiments is obtained by interpolating the data of the discharge-time-delay scan at the time of arrival of the bunch $t_b = 2.4 \pm 0.2$ µs, which results in a density of $n_0 = (4.2 \pm 1.0) \cdot 10^{15}$ cm^{-3}.

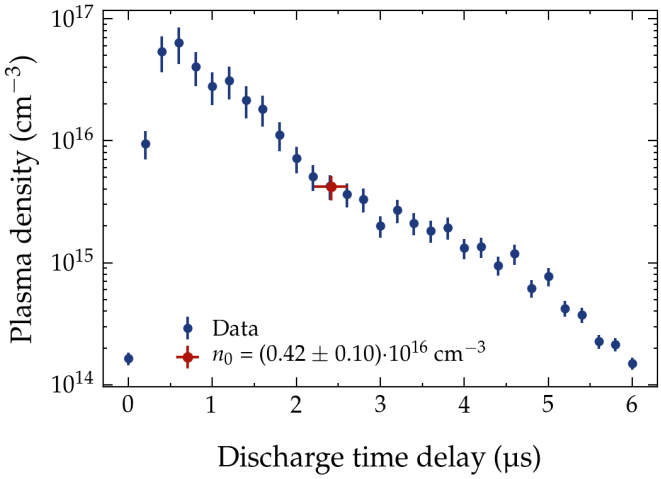

Figure 5.18.: Temporal plasma-density evolution calculated from spectral-line-broadening measurements according to the method described in [Garland et al., 2021].

According to the reconstructed double-bunch structure, if the trailing bunch has to be located within the accelerating phase of the wakefield, the plasma should have a density between $0.5 \cdot 10^{16}$ cm$^{-3} \lesssim n_0 \lesssim 1.5 \cdot 10^{16}$ cm^{-3}. The value obtained from the spectroscopic measurement, however, is slightly below the lower bound of the estimated range. Such a value is unlikely to correspond to the real density at which the experiment takes place, because in that case the trailing bunch would be located right at the beginning of the accelerating phase of the wakefield, which has two consequences: 1) the energy gain would be too low; 2) no margin would be available to produce the observed beam-loading modulations. A

possible reason for this disagreement is the difference between the electrode geometry of the cell used in the spectroscopic measurement and that of the cell used in the experimental campaign, which potentially leads to differences in the dynamics of the plasma-ignition triggered by the current pulse and, in turn, to differences in the temporal evolution of the density. Another explanation—which has been identified by the FLASHForward team during the writing of this dissertation—is related to the fact that the spectroscopic measurements are radially averaged and therefore likely to underestimate the true on-axis density. Since these issues could not be systematically investigated during the course of this work, the result of the measurement cannot be considered conclusive, and the value of the plasma density at the center of the capillary needs to be added as a variable parameter in the Bayesian Optimisation, which is discussed in the next section.

5.4.2. Working-point determination: Bayesian Optimisation

The shape of the longitudinal plasma-density profile together with the particle distribution of the double bunch structure determine the acceleration process and its outcome. In order to reduce the number of parameters that define the system—and therefore the complexity of the optimisation—several approximations are applied. Since the focus of the analysis will be put on beam loading and the modulation of the longitudinal phase space, an eventual change of the transverse beam parameters due to a tapered density transition does not need to be exactly reproduced. This allows to model the full density profile with a perfect flat-top with given effective length L and plasma density n_0. Additionally, to account for residual errors in the reconstructed phase space of the double-bunch structure, a transverse beam-size factor $C_{x/y}$ and a bunch-length compression factor C_ξ are included.

The characterisation of the accelerated trailing bunch around the working point (cf. Section 5.1) provides key aspects of the beam-plasma interaction, that allow to construct an objective function used to evaluate the level of agreement between the experiment and the outcome of the PIC simulations. The chosen parameters are: the amount of charge effectively coupled into the wakefield structure Q_w, the energy gain of the trailing bunch ΔE_w, and the energy loss of the driver bunch ΔE_d. Regarding the energy spread the approach followed is different from that of the other parameters: the experimental value used as a target is that of the true slice energy spread calculated in Section 5.3.1, whereas the simulation value used to compare with is the rms energy spread within the longitudinal extent of one FWHM around the core of the bunch. This enables to effectively minimise the chirp around the core of the bunch in order to find the conditions of optimal beam loading ($h \approx 0$).

The assumption that the experimental working point corresponds to the conditions of optimal beam loading is not strictly justified by the experimental data, which actually suggests that such conditions would be fulfilled for bunch lengths slightly below the mean BCM bunch length $\langle \sigma_{\xi,BCM} \rangle$. However, as discussed in Section 5.1, there are two reasons that prevent to precisely determine the exact beam parameters around the working point: 1) the measurement of the projected parameters is affected by either resolution problems

(at the ESPEC) or beam dispersion (at the LEMS); 2) the retrieval of the true energy chirp of the bunch from the time-resolved measurements performed at the 8FLFDUMP screen is hindered by the impossibility to apply the two-point-tomography to the accelerated bunches due to strong internal bunch correlations in the y–ξ and y'–ξ planes. The approach followed in this thesis is to benchmark the simulation with the PolariX-TDS measurements performed at the 2nd zero crossing, which do fulfill the condition $h \approx 0$ around the working point and whose time-resolved analysis is not as hampered by TDS-induced distortions as those performed at the 1st zero crossing. This approach might introduce slight systematic errors, but it reduces the complexity that the comparison between experiment and simulation would otherwise require.

The objective function is defined as the sum of the squared residuals normalised by given tolerances σ_x (for $x \in \{Q_w, \Delta E_w, \Delta E_d, \sigma_E\}$)—which is similar to the minimisation function used in the linear least-squares method:

$$f = \left(\frac{Q_{w,sim} - Q_{w,exp}}{\sigma_Q}\right)^2 + \left(\frac{\Delta E_{w,sim} - \Delta E_{w,exp}}{\sigma_{\Delta E_w}}\right)^2 + \left(\frac{\Delta E_{d,sim} - \Delta E_{d,exp}}{\sigma_{\Delta E_d}}\right)^2 + \left(\frac{\sigma_{E,sim} - \sigma_{E,exp}}{\sigma_{\sigma_E}}\right)^2, \tag{5.4.1}$$

where σ_Q and $\sigma_{\Delta E_w}$ correspond to the calculated experimental errors, and $\sigma_{\Delta E_d}$ and σ_{σ_E} are adjusted manually. In principle, different algorithms could be used to find the optimal input parameters that minimise Equation 5.4.1. However, due to the high computational cost of evaluating the function f—which implies to run a complete HiPACE++ simulation—a Bayesian Optimisation algorithm has been chosen. In the following, the basic features of BO and the optimisation framework used in this thesis are briefly described.

Optimisation framework

Bayesian Optimisation is a sequential model-based approach to solving the problem of finding a global maximiser (or minimiser) of an unknown objective function f that is noisy and/or expensive to evaluate [Mockus, 1982; Shahriari et al., 2016]. The approach makes use of a surrogate model to approximate f, which is much easier to evaluate and can be used to guide the optimisation process. The model is a probability distribution over all possible objective functions compatible with previous observations, which is often realised by means of a Gaussian Process Regression [Rasmussen and Williams, 2006] and whose mean value is the best estimate of the real objective function f. The mean μ and the standard deviation σ of the probabilistic model are then combined to construct an acquisition function which is used to obtain the best point to evaluate next according to a given trade-off between exploration (large σ) and expectation (large μ). By means of this directed evaluation of f, the number of iterations required to find the maximiser (or minimiser) of f is dramatically reduced, which translates into a more efficient use of computational resources.

The optimisation framework is set up with libEnsemble [Hudson et al., 2021], a Python library for coordinating the evaluation of dynamic ensembles of calculations in parallel. libEnsemble uses a manager process that allocates resources to perform 5 concurrent HiPACE++ simulations with different input parameters, whose outcome is subsequently used to evaluate the objective function f and generate new input parameters based on an underlying BO algorithm implemented with the open source Python library Dragonfly [Kandasamy et al., 2020], which is especially suited for parallel evaluations in synchronous or asynchronous settings. In each HiPACE++ simulation, the double-bunch structure is reproduced by a phase-space distribution of $\sim 4 \cdot 10^6$ constant-weight particles, which evolves in a plasma sampled with 2 particles per cell in a simulation box with dimensions $12 \times 12 \times 9\ k_p^{-3}$ (in $x \times y \times \xi$), resolved by a grid of $512 \times 512 \times 512$ cells. The plasma length is freely adjusted by varying both the time step size and the number of time steps, while always keeping a step size of $\sim 3\ \omega_p^{-1}$.

Results

After $N = 400$ evaluations of the objective function f, the optimisation converged to a minimiser, the parameters of which are summarised in Table 5.3. The energy gain of the witness and the energy loss of the driver are well within the given tolerances. Regarding the witness charge, however, the simulation slightly exceeds the value measured in the experiment. This could be due to either a difficulty to precisely reproduce the dynamics of the charge coupling in the simulation or to an additional charge loss between the capillary

Analysed parameter	Description	Units	**Simulation**	**Experiment**	Tolerance
Q_w	Witness charge	pC	46	41	4
ΔE_w	Witness energy gain	MeV	43	44	2
ΔE_d	Driver energy loss†	MeV	28	27	2
$\sigma_{E,w}$	Witness energy spread‡	MeV	0.6	0.4	0.1
Input parameter					Range
n_0	Plasma density	10^{16} cm^{-3}	0.98		0.5 – 2.0
L	Flat-top length	mm	47.8		25 – 50
$C_{x/y}$	Transverse beam-size factor		0.80		0.5 – 1.0
C_ξ	Bunch-length compression factor		0.96		0.75 – 1.0

Table 5.3.: Summary of analysed and variable (input) parameters for the minimised obtained from the Bayesian Optimisation.

output and the ESPEC not properly identified in the experiment. With respect to the energy spread, the difference observed results from the strategy followed to minimise the chirp—i.e. require the rms energy spread within a longitudinal extent of one FWHM around the core to be as low as the slice energy spread at the center of the bunch observed in the experiment—and a complete agreement was not expected. A closer analysis of the simulation outcome indicates that the chirp of the accelerated trailing bunch is indeed $h \approx 0$, which was the actual target of the optimisation. Altogether, the analysed parameters indicate a satisfactory level of agreement between experiment and simulation.

The input parameters obtained are all within the allowed ranges given in the optimisation. The plasma density obtained is $n_0 = 0.98 \cdot 10^{16}$ cm^{-3} and the effective flat-top length is $L = 47.8$ mm. These values contrast with those obtained in previous X-2 experiments aimed at energy-spread preservation [Lindstrøm, Garland et al., 2021], in which a density of $n_0 = 0.72 \cdot 10^{16}$ cm^{-3} and a flat-top of $L = 34.2$ mm were found to best reproduce the experimental results. These differences could be explained by two considerations: 1) in the PolariX-TDS campaign the capillary is filled with a backing pressure of 10.7 mbar, whereas in the previous experiment the pressure was ~ 20 mbar, thus leading to a different temporal and spatial evolution of the plasma density [Garland et al., 2021]; 2) in the PolariX-TDS campaign the phase space of the double-bunch structure and the charge ratio between the driver and the trailing bunch differ from those of the previous campaign, which is the result of differences in the linac and beam-line setup, the beam preparation routines, and, consequently, a different outcome of the beam-plasma-interaction optimisation. The transverse beam-size factor obtained in the optimisation ($C_{x/y} = 0.80$) could be explained by two different reasons, that do not necessarily exclude each other: 1) an adiabatic change of the transverse beam parameters resulting from the presence of density ramps; 2) experimental errors in the determination of the relative location of the beam and the plasma capillary. Due to the lack of experimental data, however, these issues cannot be investigated further and, since the agreement between the simulation outcome and the experiment is already satisfactory, it is assumed that the simulation model is equivalent to that of the experimental setup. Regarding the bunch-length compression factor, a value of $C_\xi = 0.96$ is obtained, which indicates that the measured current profile adjusts very well to the physics model defined by the remaining simulation parameters.

Figure 5.19 shows a snapshot of the last time step of the simulation corresponding to the working point obtained in the optimisation. Plots **a** and **b** show sections of the plasma across the x–ξ and the y–ξ planes, respectively, in which the plasma density is represented in grey and the beam density in color. The current profile of the double-bunch structure is represented by a dark-grey line on top of the density colormaps with an independent vertical axis. In the x–ξ plane a clear centroid offset in the trailing bunch is observed, which hints the presence of large transverse oscillations during acceleration. These centroid offsets justify the need to locate the trailing bunch at a longitudinal position close to the center of the cavity, where spatial constraints are the least stringent and a sufficient amount of charge coupling into the wakefield structure is achieved. The interaction of the bunch with the plasma-sheath electrons results in a distortion of the wakefield symmetry, which can be clearly observed

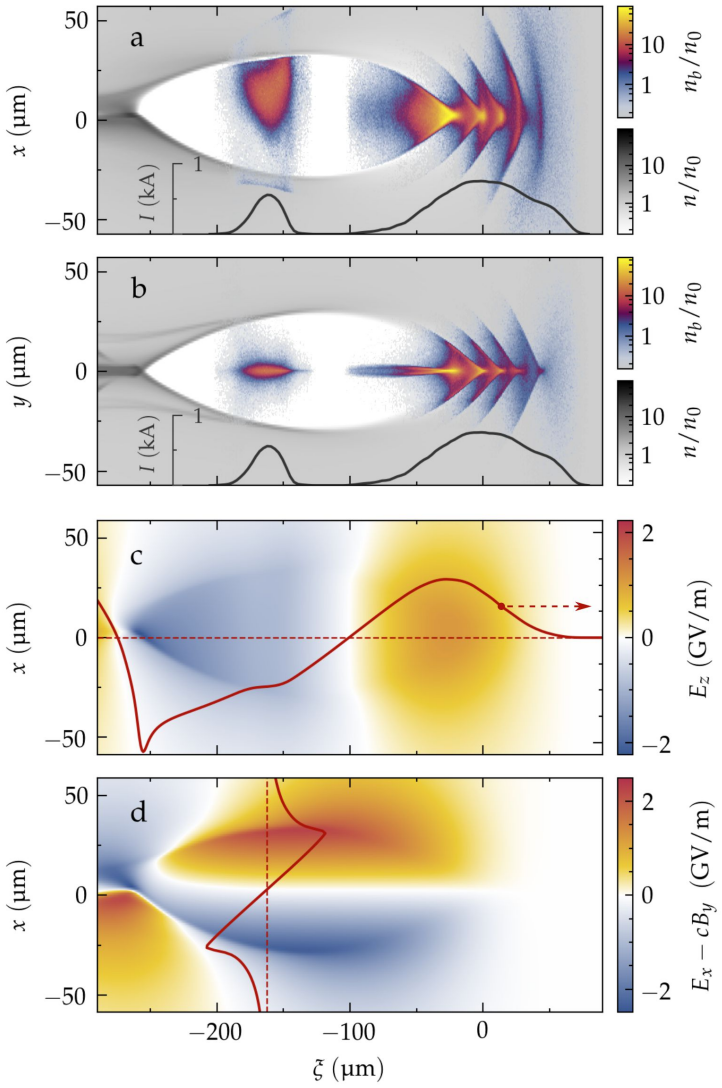

Figure 5.19.: Snapshot of the last time step of a HiPACE++ simulation corresponding to the working point obtained by means of a Bayesian Optimisation. Top plots: horizontal and vertical slices accross the wakefield structure indicating the plasma (grey) and the beam (colour) densities. Bottom plots: horizontal slice accross the wakefield structure indicating the longitudinal accelerating field E_z and the radially focusing field $E_x - cB_y$, including a layout (solid red lines) along the coordinates $x = 0$ µm and $\xi \approx 160$ µm (dashed red lines).

in the colorscale of the cavity boundaries at longitudinal coordinates below −200 µm. In the y–ξ plane, in contrast, the trailing bunch is transversely well centered. Even though this does not preclude the presence of transverse dynamics, an analysis of the beam evolution during acceleration confirms that the trailing bunch stays approximately centered for all time steps, which stays in contradiction with the experimental observation that the bunch exits the plasma capillary with a pronounced correlation in the y'–ξ plane (cf. Section 5.2). After a close inspection of the code used to generate the initial particle distribution of the double-bunch structure, an error was found that translates into a dramatic underestimation

of the centroid offsets in the y–y' plane along the bunch, which explains the disagreement between simulation and experiment. On account of that, the focus of the discussion will be put on the x–ξ plane in the following.

Plot **c** of Figure 5.19 shows a section of the accelerating gradient across the x–ξ plane $E_z = E_z(x, y = 0, \xi)$ (diverging colormap) and a layout at $x = 0$ (red line). As a result of the distortion of the cavity shape, the accelerating gradient around and behind the trailing bunch becomes transversely asymmetric, which can be most clearly recognised in the dark blue spot at the lower part at the back of the bubble ($\xi \approx -250$ µm). The longitudinal layout reveals that a good level of beam loading is achieved, resulting in an accelerating gradient of $E_z \approx -1$ GeV/m that is approximately constant along a fraction of the trailing bunch around the core ($\xi \approx -160$ µm). Plot **d** shows a section of the focusing field across the x–ξ plane (i.e., $E_x - cB_y$) and a layout along the x axis at $\xi \approx -160$ µm, which corresponds roughly to the center of mass of the trailing bunch. As expected in the blowout regime, the layout makes clear that the radially focusing force increases linearly with radius up to the walls of the cavity and decreases rapidly thereafter. In the longitudinal axis, the focusing force increases progressively from the front to the center of the cavity and decreases from the center to the back. The field asymmetry caused by the cavity deformation is especially noticeable towards the back of the structure. The fishbone-like structure observed in the driver results from the combined effect of the geometry of the radially focusing field (plot **d**) and that of the decelerating field (plot **c**) as the wakefield is formed at the front of the cavity, which lead the particles of different slices to oscillate at particular frequencies depending on their distance to the center of the cavity.

The temporal evolution of the interaction between the plasma and the beam in the horizontal plane x–ξ is shown in Figure 5.20. The wakefield structure builds up rapidly and reaches its final dimensions around the third time step, at which point the back of the bubble closes behind the trailing bunch, stabilising at a location slightly below ~ -250 µm. As the driver propagates, the focusing field pinches the beam and the fishbone-like structure builds up progressively. Most remarkably, the figure provides a clear visualisation of the transverse dynamics of the trailing bunch and some of its consequences. The large-amplitude oscillations of the centroid lead the bunch to interact with the plasma-sheath electrons and the fraction of beam electrons that are able to run over the maximum of the focusing field start to oscillate at a lower frequency, eventually detaching from the trailing bunch and being lost in the plasma background outside the cavity. This process can be observed, for instance, at the lower side of the cavity between the time steps $t = 58$–101 ps, and at the upper side between the time steps $t = 101$–159 ps. Additionally, the correlation between the trailing-bunch oscillations and the cavity deformation along the whole acceleration process is clearly noticeable.

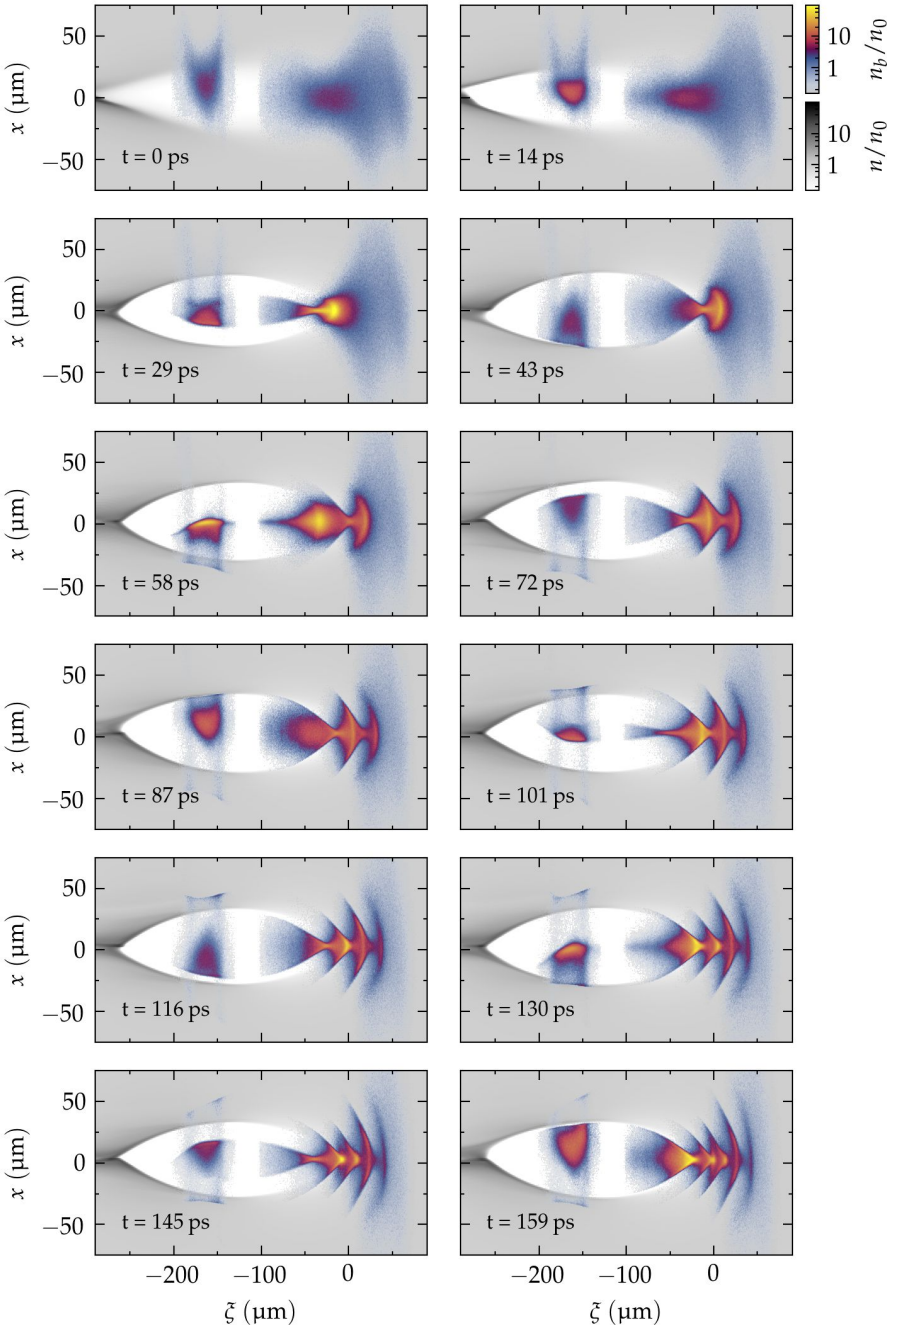

Figure 5.20.: Snapshots of a HiPACE++ simulation of the experimental working point illustrating the time evolution of the beam-plasma interaction. The images show an horizontal slice accross the wakefield structure indicating the plasma (grey) and the beam (colour) densities.

Figure 5.21 provides a better insight into the accelerating gradient in the neighborhood of the trailing bunch. The top-right plot shows the same section of the accelerating gradient across the x–ξ plane as in Figure 5.19c—i.e., a snapshot of the last time step of the simulation—together with the current profile of the double bunch structure. The top-left plot shows transverse layouts of the field $E_z = E_z(x, y = 0, \xi_i)$ at three different longitudinal locations ξ_i

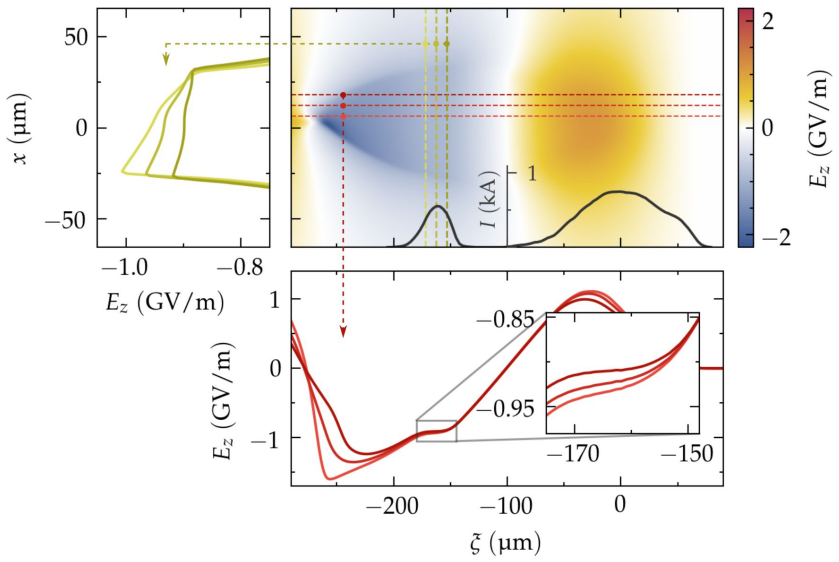

Figure 5.21.: Analysis of the accelerating field E_z homogeneity around the core of the trailing bunch for the last step of the HiPACE++ simulation. Top center: horizontal slice of the accelerating field E_z accross the wakefield structure and current profile of the driver-trailing-bunch pair. Top left: layout of the accelerating field at longitudinal coordinates corresponding to the center of the trailing bunch and $\pm\sigma_{\xi,rms}$. Bottom plot: layout of the accelerating field at horizontal coordinates corresponding to the center of the trailing bunch and $\pm\sigma_{x,rms}$.

along the trailing bunch (olive-green lines), whereas the bottom-right shows longitudinal layouts of the field $E_z = E_z(x_i, y = 0, \xi)$ at three different horizontal locations x_i (red lines). The locations correspond roughly to the center of mass of the bunch, and to the center of mass plus and minus one sigma (i.e., either $\langle\xi\rangle \pm \sigma_\xi$ or $\langle x\rangle \pm \sigma_x$). The three transverse layouts shown in the top-left plot, clearly indicate that the deformation of the cavity strongly affects the uniformity of the accelerating gradient along the x axis, with the slope of the field $\partial E_z/\partial x|_{\xi=\xi_i}$ increasing rapidly from the head to the tail of the trailing bunch. In terms of the longitudinal layouts shown in the bottom-right plot, this effect can be interpreted as a dependence of the amount of beam loading on the transverse coordinate of the beam x. Altogether it is clear that even an infinitesimal longitudinal slice of the bunch experiences an instantaneous accelerating gradient that varies (at least) with the transverse location x of the particles. The instantanous nonuniformity of the accelerating gradient in the transverse plane, however, does not provide direct information about the outcome of the acceleration process over the full plasma length, since these spatial variations depend on dynamic cavity deformations driven by transverse centroid oscillations and, at the same time, the beam electrons oscillate inside the beam envelope. Nonetheless, the total effect of these nonuniformities can be characterised by the slice-energy-spread growth after the beam-plasma interaction, which in turn allows to calculate the (effective) field nonuniformity as defined by Equation 5.3.1. By using the longitudinal resolution obtained in the longitudinal-phase-space measurements ($R_\xi = 2.46$ µm), the rms energy spread of the central slice of the accelerated trailing bunch is $\sigma_{E,\mathrm{on}} = 0.54$ MeV, which results in a (single-shot) effective field nonuniformity of $\sigma_\zeta = 0.9$ %, in good agreement with the experimental calculations of the

(statistical) effective field nonuniformity presented in Section 5.3.

5.4.3. Simulation of beam-loading modulations

The goal of this section is to reproduce the experimental observation of beam-loading modulations around the working point obtained with the Bayesian Optimisation and gain a better insight into the physics of the beam-plasma interaction. To that end, the bunch-length variations characterised in Chapter 4 are used to generate several double-bunch structures, which are used as input of separate HiPACE++ simulations. The parameters of each double-bunch structure are derived from the working-point characterisation described in Section 4.4 by scaling (stretching) the longitudinal coordinate ξ of the particle distribution by a given factor C_ξ between 0.85 and 1.15—which according to the formula $\sigma_{\xi,BCM} = C_\xi \cdot \langle \sigma_{\xi,BCM} \rangle$, where $\langle \sigma_{\xi,BCM} \rangle$ is the mean BCM bunch length, correspond to total bunch lengths of $\sigma_{\xi,BCM} \approx 80$ µm and $\sigma_{\xi,BCM} \approx 109$ µm, respectively—and adjusting the charge ratio between the driver according to the measured charge modulations (cf. Figure 4.17). The transverse phase space is the same for all particle distributions—i.e., that measured in the experiment multiplied by the transverse beam-size factor obtained from the BO. The plasma parameters are fixed to those detailed in Table 5.4.2.

Figure 5.22 shows the layout of the accelerating gradient along the longitudinal coordinate $E_z = E_z(x = 0, y = 0, \xi)$ together with the current profile of the beam for the 13 stretch factors simulated with (top) and without (bottom) trailing bunch.The effect that the driver modulations have on the shape of the accelerating-gradient can be seen in the top plot, in which shorter (longer) bunches with higher (lower) charge drive a stronger (weaker) wakefield. Additionally, for a stronger excitation the slope of the accelerating-gradient $\partial E_z / \partial \xi$ is approximately linear over a wider range of accelerating phases, whereas it decreases progressively towards the back of the bubble when the excitation is weaker. These two effects depend mainly on the peak current, and are therefore determined by the length of the bunch and its charge $I_{max} \propto Q_d \, \sigma_{\xi,d}^{-1}$. Due to the correlation between charge Q and bunch length σ_ξ, the changes in the wakefield strength are particularly pronounced for shorter bunches ($C_\xi < 1.0$), whereas they are almost negligible for longer bunches ($C_\xi > 1.0$). The beam loading produced by the trailing bunch is shown in the bottom plot. The bunch-length variations translate into a modulation of the distance between the driver and the trailing bunch and therefore into a modulation of the location ξ at which beam loading starts. Since the shift in the longitudinal location of the trailing bunch is correlated with its charge, the bunch at ξ_3 has a lower charge than that at ξ_1 (i.e., $Q_w(\xi_3) < Q_w(\xi_2) < Q_w(\xi_1)$. As a result of these two factors, the trailing bunch located at the earliest accelerating-gradient phase ($\xi = \xi_3$) is not capable to produce a sufficient amount of beam loading, whereas that located at the latest phase ($\xi = \xi_1$) even reverses the slope of the accelerating-gradient $\partial E_z / \partial \xi < 0$, resulting in an overloading of the wakefield. At locations around ξ_2 the beam loading is close to the optimal and the slope of the accelerating-gradient is approximately zero at the core of the bunch. Since the mean accelerating-gradient increases as the bunch is located

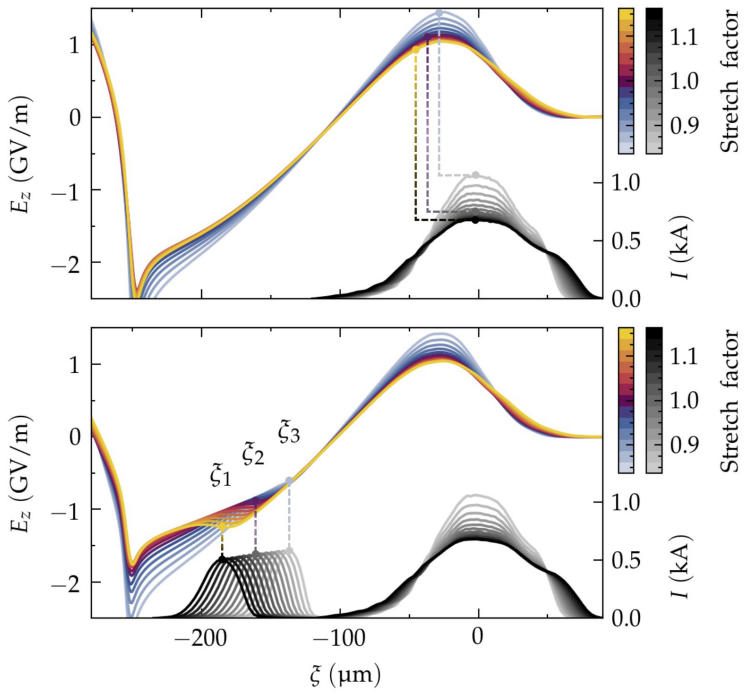

Figure 5.22.: Accelerating-field layout accross $x, y = 0$ (coloured lines, left axis) and beam current (grey lines, right axis) for different beam-compression settings without (top) and with (bottom) beam-loading simulated with HiPACE++.

further behind the wakefield, a larger (smaller) amount of beam loading leads to a higher (lower) energy gain. The overall behavior described above seems to contradict the theory of beam loading in the nonlinear regime developed in reference [Tzoufras et al., 2008], in which the peak current required to beam load the wakefield increases with the strength of the accelerating-gradient—i.e., the current of the trailing bunch at ξ_1 should be larger than that at ξ_3. The disagreement, however, could be explained by two important differences between the experimental conditions of the PolariX-TDS campaign and those assumed in [Tzoufras et al., 2008]: 1) in the latter the acceleration takes place in the strong blowout regime or ultrareletivistic limit, for which $k_p R_b \geq 3$, where k_p is the plasma skin depth and R_b is the blowout radius, whereas in the first $k_p R_b \approx 0.5$; 2) in the experiments presented in this thesis, the shape of the wakefield structure is strongly distorted by the centroid offsets of the trailing bunch, whereas this is not the case in the cited reference.

Figure 5.23 shows a comparison between the longitudinal phase space measured experimentally and that obtained from simulations for three different conditions: underloading (left), best beam-loading conditions (center) and overloading (right). The data from the simulation corresponds to the longitudinal-phase-space ξ–δ obtained at the output of the plasma capillary for the stretch factors $C_\xi = 0.925$, 1.0 and 1.1. The experimental data corresponds to PolariX-TDS measurements performed at the 2nd zero crossing and the shots selected are those that best match two key parameters of the simulation simultaneously: the stretch factor and the chirp h of the accelerated trailing bunch (dotted red line in the plots). The

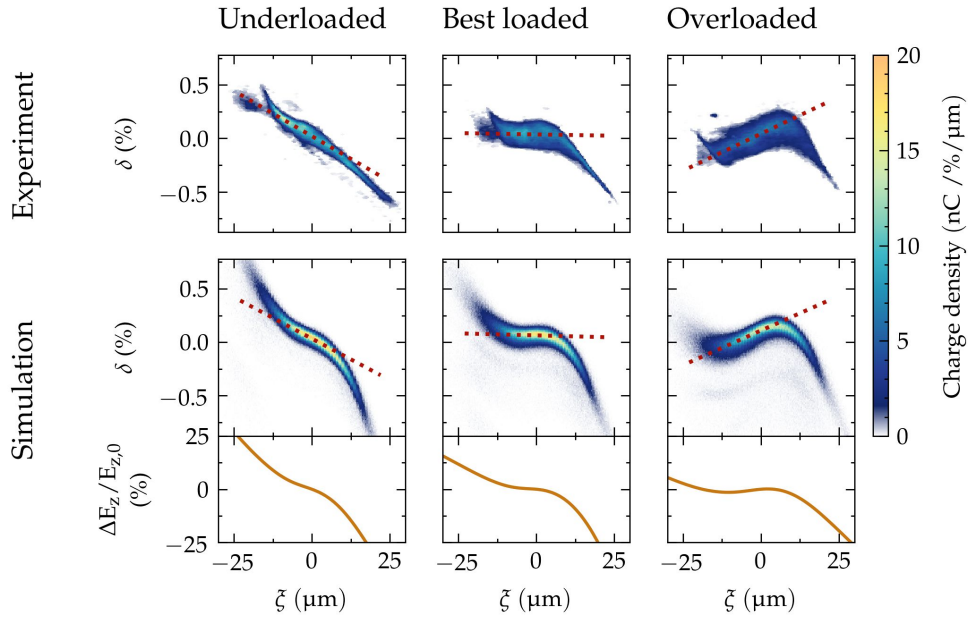

Figure 5.23.: Comparison of the longitudinal phase space of accelerated trailing bunches obtained from measurements at the 8FLDUMP screen at the 2nd zero crossing (top plots) and obtained from HiPACE++ simulations (bottom plots) for three different beam-loading conditions: underloaded (left), best loaded (center) and overloaded (right). The bottom simulation plot shows the relative accelerating-field variation around the core of the bunch.

bottom plots show the relative variation of the accelerating gradient around the core of the bunch for the last time step of the simulation—i.e., $\Delta E_z/E_{z,0} = (Ez - E_{z,0})/E_{z,0}$, where $E_{z,0}$ is the accelerating gradient at the center of mass of the trailing bunch $\langle\xi\rangle_w$. The images of experiment and simulation exhibit a remarkable level of agreement, but there are some differences. On the one hand, the slice energy spread of the experimental images increases progressively with the amount of beam loading. This is mainly due to a deterioration of the energy resolution at the screen for increasing amounts of beam loading through the following mechanism: an increase of beam loading translates into an increase of mean energy, which in turn affects the evolution of the beam envelope along the beamline, resulting in a larger natural beam size at the screen. On the other hand, the slope at the head of the bunch—which is determined by the unperturbed accelerating gradient—is larger in the simulation than in the experimental images. This could be explained by the TDS-induced energy chirp expected from the measurement. However, it is difficult to precisely quantify this effect, since the application of the two-point-tomography on the accelerated bunches—which is required to retrieve the real sliced mean-energy profile—is compromised by systematic errors (cf. Section 5.1.1). Finally, both the head and the tail in the simulation contain more charge than in the experiment, which hints a disagreement in the reproduction of the dynamics of the charge coupling into the cavity and might be the reason why the overall charge in the simulation is overestimated compared to the experiment. Since this issue could not be further investigated in this thesis, the figure of merit used to compare the energy spread between simulation and experiment is the FWHM, which is insensitive to the outermost regions of the energy

distribution.

A further evaluation of the agreement between experiment and simulation is accomplished by comparing the simulation outcome with high-statistics data from the measurements. Four parameters are analysed in the following: the mean energy of the accelerated trailing bunch, its charge, its energy spread (FWHM) and its chirp. The experimental mean energy and charge are taken from the ESPEC, whereas the experimental energy spread and chirp are those measured at the 2nd zero crossing at the 8FLFDUMP screen. The comparison is shown in Figure 5.24. The mean energy obtained from the simulation describes the same trend as in the experiment—i.e., shorter (longer) BCM bunch lengths lead to lower (higher)

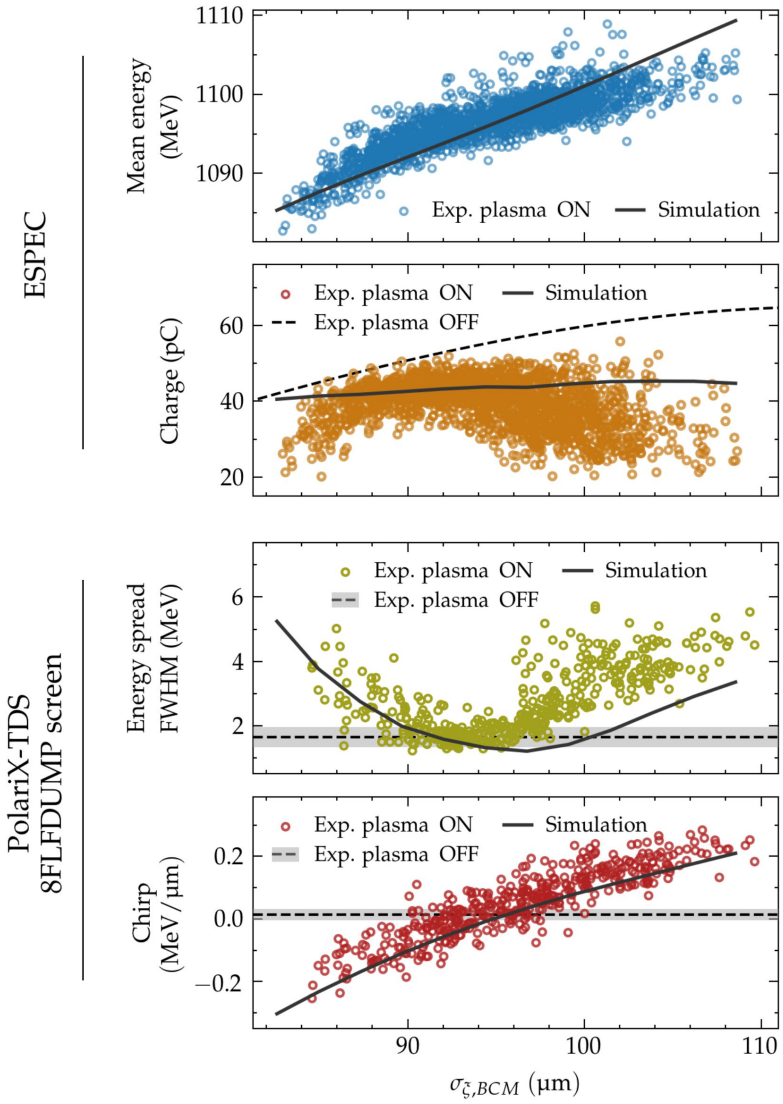

Figure 5.24.: Comparison between experiment and HiPACE++ simulations of key beam parameters of accelerated bunches and their correlation with rms bunch length $\sigma_{\xi,BCM}$. Top plots: mean energy and charge measured at the ESPEC. Bottom plots: FWHM energy spread and chirp measured at the 8FLFDUMP screen at the 2nd zero crossing.

energies—and covers a similar range of energies. Nonetheless, the experimental data can be divided in two regions below and above $\sigma_{\xi,BCM} \approx 90$ µm that show a different correlation with the BCM bunch length, whereas in the simulation the correlation is constant along the whole range of bunch lengths. Regarding the charge, the simulation reflects that an increasing fraction of particles are lost with increasing amounts of beam loading. Nonetheless, the charge loss observed experimentally for the lowest and uppermost bunch lengths is not properly reproduced. This could be a true disagreement between experiment and simulation, but could also reflect experimental limitations associated to the difficulty to: 1) effectively capture and transport the beam; 2) diagnose the actual amount of charge exiting the plasma capillary with a higher level of accuracy. Therefore, it is difficult to derive a conclusion from the available data. The energy spread exhibits a good agreement below the mean BCM bunch length $\langle \sigma_{\xi,BCM} \rangle = 94.4$ µm, but experiment and simulation start to diverge above that value, which could be partly associated to the broadening of the experimental energy spectrum due to the deterioration of the energy resolution at the screen. Regarding the chirp, the level of agreement is excellent.

Altogether, the simulation of beam loading modulations around the working point reproduces the trends observed in the experiment within a similar range of values, and supports the argument that these modulations are triggered by bunch length variations that modify the shape of the double-bunch structure. However, the agreement is not perfect, which could be related to several reasons that are discussed next.

1. Since the target values used to determine the working-point correspond to the experimental data measured at the 2nd zero crossing, the outcome of the Bayesian Optimisation will deviate from the real working point, potentially leading to differences in the observed beam loading modulations. This is actually the result of a more fundamental problem: the difficulty to retrieve the real beam parameters from the time-resolved measurements performed with the PolariX-TDS, which are affected by systematic errors that cannot be compensated analytically. A possible solution would be to include the TDS measurements in the optimisation. However, the simulation of the beam propagation would require, among others, a highly refined model of the longitudinal plasma-density profile in order to properly interface the plasma capillary with the diagnostics beamline and would result in a dramatic increase in the complexity of the simulation framework.
2. Apart from fundamental limitations associated to the time-resolved diagnostics with the PolariX-TDS, there are other beam parameters whose uncertainty could not be properly assessed during the experimental campaign discussed in this thesis—e.g., the amount of charge effectively exiting the plasma capillary or the projected energy spread of the accelerated bunch. These issues can in principle be overcome with both a careful calibration of the available diagnostics and a more extensive set of measurements. However, such procedures are typically subject to time constraints imposed by the beam-time availability, and could not be fully unfolded during the PolariX-TDS campaign.

3. In the physics model used in the target function, the eventual presence of plasma-density ramps has not been explicitly considered and the longitudinal plasma-density profile has been approximated by a perfect flat-top with an effective length that includes the contribution of the ramps to the total energy gain. Depending on the geometry of the ramps, however, this model might not be sufficient to reproduce the experimental results with the desired level of accuracy.
4. It is possible that the target function used in the optimisation does not converge in a unique solution. In that case, different solutions could lead to the same outcome at the working point, but behave differently to bunch-length modulations. This could be compensated by introducing further beam parameters into the target function as, for instance, the observed beam tilts at the output of the capillary. Another option would be to include the beam-loading modulations into the optimisation framework.
5. As mentioned before, the double-bunch structure used in the simulations does not correctly reproduce the phase space of the centroids in the y–y' plane. This could have an impact on the evolution of the wakefield structure and, in turn, on the overall acceleration process simulated with HiPACE++.

In summary, there are multiple aspects of the optimisation framework that can be improved in order to achieve a higher level of agreement between simulation and experiment. On the other hand, the experimental data available has some intrinsic limitations that can only be overcome by performing new measurements. In spite of that, the results presented in this section already provide an insight into critical aspects of the acceleration process that affect both the quality of the beam loading and the homogeneity of the accelerating-gradients produced.

6. Conclusion

Plasma wakes driven by intense, ultra-relativistic electron beams are capable of sustaining electric-field strengths in excess of $\mathcal{O}(\mathrm{GV/m})$ within a spatial extent on the order of $\mathcal{O}(100\ \mu\mathrm{m})$. By harnessing these wakefields, a particle bunch trailing behind the driver can be accelerated to GeV energies over meter distances while, in principle, preserving its quality—i.e., its energy spread and emittance. For this reason, beam-driven plasma-wakefield accelerators (PWFA) are promising candidates to become compact and cost-efficient building blocks of next-generation high-energy particle-accelerator facilities—for instance as the cornerstones of multi-TeV linear colliders and free-electron lasers (FELs). Although tremendous progress has been made during the past two decades, several challanges remain to be tackled before PWFA stages are ready for large-scale application. One of these challenges is to guarantee the preservation of the slice and projected energy spread of the accelerated bunch, which requires an extremely precise control over both the acceleration process and the phase-space distribution of the driver and the trailing bunch.

The work presented in this thesis addresses this issue by exploiting the diagnostic capabilities of a novel X-band transverse deflection structure (TDS) featuring a variable polarisation of the streaking field: the PolariX-TDS. In a recent experiment closely related to this project, energy-spread preservation and high instantaneous energy-transfer efficiency have been demonstrated [Lindstrøm, Garland et al., 2021]. On the grounds of this success, two aspects that lead to quality preservation are further investigated: optimal beam loading for chirp compensation and homogeneity of the longitudinal and transverse electric-field structure within the volume of the blown-out plasma wake. In the following sections, a brief executive summary of the work presented in this thesis is provided and several implications of the obtained results are discussed. The chapter is concluded with some thoughts about possible strategies to improve the outcome of the performed experiments with a view towards PWFA at higher energy gains.

6.1. Executive summary

Beamline design and commissioning. In order to enable the time-resolved phase-space characterisation of electron bunches, a new diagnostics beamline accommodating the PolariX-TDS prototype has been brought into operation. A differential-pumping system that enables the simultaneous operation of the TDS and the windowless plasma target has been designed and commissioned, validating ultra-high vacuum (UHV) conditions with $p \approx 10^{-9}$ mbar at

the boundaries of the RF structure for nominal gas loads of $Q_{\text{in}} = 20\ \text{mbar} \cdot \text{l/s}$ at the plasma chamber. An experimental campaign was planned and conducted to calibrate and verify the proper operation of regular beamline components, including quadrupole magnets, beam steerers, and beam-position monitors (BPMs). Critical problems with the power supplies of several magnets have been detected, investigated, and mitigated. The field-of-view of each of the two screen stations located downstream of the TDS has been determined to enable the longitudinal diagnostic of electron beams in a wide range of experimental conditions subject to different sources of jitter.

PolariX-TDS commissioning. A dedicated beam-based commissioining campaign has been planned and conducted to assess the performance of the PolariX-TDS system. In spite of slight deviations from the nominal operation—which have been systematically investigated—the new device has been shown to fulfil the stability requirements that enable the diagnostic of electron bunches with time resolutions of few femtoseconds. In preparation for the main experimental campaign of this thesis, the capabilities of the TDS have been validated by means of two beam-diagnostics applications: 1) a phase space reconstruction consisting of measurements of the slice-emittance in x and y and the longitudinal-phase-space; 2) a dispersion-based beam-tilt-correction technique that enables the minimisation of the sliced transverse centroid offsets of the bunch—critical to the sucess of PWFA experiments. Regarding the diagnostic of the longitudinal phase space, the energy profile measured at the screen has been shown to be especially sensitive to internal beam correlations in the streaking plane, in agreement with theoretical predictions.

Beam loading and slice-energy-spread preservation. In order to address the ultimate scientific goal of this work, an experimental campaign has been envisaged and carried out. A detailed account of the procedures required to set up the experiment has been provided, including the preparation and optimisation of FLASH beams and the optimisation of the beam-plasma interaction—requiring a high level of complexity. A comprehensive time-resolved phase-space characterisation of the beams before and after interaction with the plasma has been accomplished. The measurements have been found to be affected by several sources of errors that are associated with: 1) the stability of the machine operation; 2) the stability of the plasma-acceleration process; 3) inherent limitations of the TDS-diagnostic and its sensitivity to the beam envelope. By means of a detailed analysis, it has been shown that some of these errors can be partly compensated for, whereas those resulting in longitudinal-phase-space distortions can neither be precisely quantified nor compensated for analytically. Furthermore, the observed limitations of the PolariX-TDS diagnostic have been found to be related to the unoptimised beam-plasma interaction in the vertical plane y–y', which prevents appropriate matching of the highly divergent beam parameters to the lattice optic. In spite of these issues, the measurements have provided an unprecedented insight into the acceleration of electron bunches in a blown-out plasma wake, which has facilitated the main study of the experimental campaign: beam loading and energy-spread preservation.

- **Beam loading.** A direct observation of the longitudinal phase space of electron bunches accelerated in a blown-out beam-loaded plasma wake has been accomplished for first time. The magnitude of beam loading has been shown to correlate with the current profile of the driver-trailing-bunch pair, which in the experimental conditions of this study is affected by bunch-length variations induced in the FLASH linac that result in variations of the double-bunch structure produced by the collimator device. In good agreement with theoretical predictions, the Gaussian-like current profile of the trailing bunch has been shown to flatten the accelerating field around the core of the particle distribution.
- **Slice-energy-spread preservation.** By comparing the slice energy spread of the trailing bunch before and after plasma interaction, it has been demonstrated that the accelerating gradient of the wakefield structure is transversely homogeneous to within 0.8 % (1.5 %) at an interval of confidence of 68 % (95 %). These values represent upper bounds based on measurement uncertainties, but arguments have been provided to presume that the field is perfectly homogeneous in at least 46 % of events. On account of the Panofsky-Wenzel theorem, it has been argued that, simultaneouly, the radially focusing fields are longitudinally constant along the core of the bunch, thus demonstrating that the field geometries of a blown-out beam-loaded plasma wake enable an aberration-free acceleration. At the same time, it has been ascertained that these results can only be considered valid if the observed longitudinal-phase-space distortions are ignored.

Particle-in-cell (PIC) simulations. A tentative reconstruction of the beam-plasma interaction in the PIC code HiPACE++ has been accomplished. While the accuracy of the reconstruction is affected by a number of issues—amongst others, an oversimplification of the physics model—it has been shown that simulation and experiment exhibit a substantial level of agreement, which has revealed some insight into the dynamics of the acceleration process. Most importantly, the oscillations triggered by transverse misalignments between the driver and the trailing bunch have been shown to distort the shape of the plasma-wake structure, which, in turn, has an impact on the homogeneity of the accelerating field that could lead to a degradation of the trailing-bunch quality and a destabilisation of the acceleration process.

6.2. Discussion of results

The main result obtained in this study supports the argument that PWFA schemes are capable of producing aberration-free acceleration by driving a blown-out wake in a plasma and appropriately beam loading the accelerating field. From a theoretical point of view and according to PIC simulations, this is indeed expected to be the case. Furthermore, the experiments provide an unprecedented insight into the time-resolved phase-space structure of plasma-accelerated bunches. This phase-space characterisation demonstrates the extreme

sensitivity of the acceleration process to the exact shape of the driver and the trailing bunch, and suggests that the actual challenge to be addressed for PWFA to fulfil the requirements for applications might be to improve the control over the phase space of the incoming beams. Simultaneously, the experience made in the course of this work indicates that the use of the PolariX-TDS for diagnosing the acceleration process requires an improved control over the plasma-accelerated bunches to enable their appropriate transport to the RF structure. In the following paragraphs these two aspects are briefly discussed.

Optimisation of electron bunches for PWFA. The experiments performed at FLASH-Forward are constrained by a trade-off between peak current and beam quality. While beams with a high peak current are required to drive a strong blow out, they also trigger collective effects during their acceleration and longitudinal compression in the FLASH linac and their transport through the extraction sections upstream of the interaction point. Such collective effects translate into: 1) distortions of the longitudinal phase space of the bunch—viz. a longitudinally-dependent increase of the uncorrelated energy spread and nonlinearities in the sliced mean energy profile; 2) sliced transverse centroid offsets and internal transverse-longitudinal beam correlations. The first affects the operation of the collimator device, which requires the beam to have a linear energy chirp and a small sliced transverse beam size at the location of the collimator. The latter may trigger transverse oscillations during the beam-plasma interaction that destabilises the acceleration process and/or degrades the quality of the trailing bunch.

Due to the transverse-longitudinal couplings in dispersive sections arising from a linear energy chirp, the presence of centroid offsets together with eventual transverse misalignments of magnet-lattice elements leads to a complicated beam evolution along the beamline. Due to this, the optimisation of the beam-plasma interaction becomes an extraordinarily involved process, as different beam parameters that need to be precisely adjusted at the interaction point are correlated to each other in an intricate way. These difficulties have been overcome for short plasma-interaction lengths of about 5 cm by introducing novel experimental methods such as, for instance, the dispersion-based beam-tilt correction or the two-BPM-tomography for matching small beta functions into plasma. An example of the success of such methods is the experiment that motivates the studies presented in this thesis [Lindstrøm, Garland et al., 2021]. However, this might be more challenging when scaling the experiments to larger propagation lengths, which will be required to achieve higher energy gains. Apart from the aforementioned issues related to the presence of transverse centroid offsets, the lack of appropriate diagnostics to assess the beam-plasma interaction in the vertical plane y–y' might become a problem, since the deformation of the accelerating structure induced by trailing bunch oscillations could have a non-negligible impact on the overall acceleration process.

Transport and longitudinal diagnostic of PWFA-accelerated electron bunches. The large integrated kick generated by the PolariX-TDS translates into a magnification of TDS-induced effects predominantly in the energy axis. In the measurements presented in this thesis two of these effects introduce systematic errors that are either difficult or impossible to negotiate:

1) nonlinear terms in the induced energy chirp associated with internal beam correlations in the y–ξ and y'–ξ planes (cf. Eq. 1.3.47)
2) longitudinal-phase-space shearing effects associated with a mismatch of the beam parameters ($\alpha_{y,\mathrm{TDS}} \neq 0$) at the location of the TDS (cf. Eq. 5.2.3).

In order to mitigate the first, the sliced transverse centroid offsets in the streaking plane must be minimised, which is a demanding task due to both the complicated phase-space structure of the incoming beams and the lack of diagnostics required to optimise the beam-plasma interaction in the vertical plane. In order to mitigate the second, the beam envelope must be very well matched to the lattice optic, which is especially difficult for highly divergent beams and also requires an appropriate diagnostic of the beam-plasma interaction in the vertical plane. For these reasons, the longitudinal diagnostic of PWFA-accelerated beams with the PolariX-TDS is only possible for a highly optimised beam-plasma interaction in both the horizontal and the vertical plane, which is currently very challenging.

6.3. Outlook

While the time-resolved phase-space characterisation of electron bunches presented in this work provide a high-level of detail, the reconstruction of the beam-plasma interaction is comparatively less accurate. In order to exploit the full potential of PIC codes and gain a better insight into the acceleration process, several aspects of the simulation framework that have been identified to limit the level of agreement with the experiments must be addressed in the future. Along the same lines, on account of the sensitivity of the acceleration process to the sliced transverse centroid offsets, an in-depth tolerance study by means of PIC simulations could help determine the level of control over the phase-space of the incoming beams required to achieve stable and quality-preserving acceleration at higher energy gains—i.e., over longer propagation lengths through the plasma.

In the experiments discussed in this dissertation—and at FLASHForward in general—the approach adopted to optimise the FLASH beams for PWFA experimentation is purely empirical, i.e., the trade-off between high-peak current and beam quality is determined experimentally to the extent enabled by the available beamline diagnostics. Start-to-end simulations including collective effects could assist the optimisation process by systematically investigating how different machine parameters affect the beam quality. Although it might be challenging to achieve a good level of agreement between simulation and experiment, several upgrades at the FLASH facility are currently under way that could enable accurate

benchmarking procedures. On the one hand, a new and shorter second bunch compressor (BC3, cf. Section 2.1) will grant more control over internal beam correlations and couplings and generates space to accomodate a new section for re-matching the beam from the injector into the linac—which until now was only possible at the section downstream of BC2. On the other hand, the FLASH2 beamline is equipped with two PolariX-TDS for longitudinal diagnostics, thus enabling the assessment of the beam quality after the FLASH2 extraction and before the FLASHForward extraction section.

6.4. Final summary

The PolariX-TDS system was added to the FLASHForward beamline to enable novel measurements within the field of PWFA. Measurements of the longitudinal phase space of electron bunches accelerated in a blown-out, beam-loaded plasma wake were accomplished for first time. These measurements provided the necessary means to establish electric-field structures required for meaningful application of the field to future facilities requiring high levels of longitudinal beam quality. On account of that and the recent milestones achieved at the FLASHForward facility a thriving future for PWFA research and development is expected.

Appendices

Appendix A.

Transfer matrices of individual beamline elements

The derivation of the most common first-order-imaging beamline elements can be obtained from Equations 1.1.8–1.1.14. Below, the resulting matrices are presented.

Drift space. This is the simplest element in a beamline, which represents the propagation of the beam in free space. For a drift of length l:

$$\boldsymbol{R}_{\mathrm{D}} = \begin{pmatrix} 1 & l & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & l & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \end{pmatrix} \tag{A.0.1}$$

Quadrupole magnet. The focusing strength of a pure quadrupole is given by the expression with strength $k_1 = e\,g/p$, where g is the gradient of the magnetic field and p the reference momentum of the particles. For $k_1 > 0$ the matrix is:

$$\boldsymbol{R}_{\mathrm{Q}} = \begin{pmatrix} \cos\phi & \frac{1}{\sqrt{|k_1|}}\sin\phi & 0 & 0 & 0 & 0 \\ -\sqrt{|k_1|}\sin\phi & \cos\phi & 0 & 0 & 0 & 0 \\ 0 & 0 & \cosh\phi & \frac{1}{\sqrt{|k_1|}}\sinh\phi & 0 & 0 \\ 0 & 0 & \sqrt{|k_1|}\sinh\phi & \cosh\phi & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \end{pmatrix} \tag{A.0.2}$$

and for $k_1 < 0$ the two transverse planes are just interchanged:

$$\boldsymbol{R}_{\mathrm{Q}} = \begin{pmatrix} \cosh\phi & \frac{1}{\sqrt{|k_1|}}\sinh\phi & 0 & 0 & 0 & 0 \\ \sqrt{|k_1|}\sinh\phi & \cosh\phi & 0 & 0 & 0 & 0 \\ 0 & 0 & \cos\phi & \frac{1}{\sqrt{|k_1|}}\sin\phi & 0 & 0 \\ 0 & 0 & -\sqrt{|k_1|}\sin\phi & \cos\phi & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \end{pmatrix} \tag{A.0.3}$$

where $\phi = l_{\mathrm{eff}}\sqrt{|k_1|}$, and l_{eff} is the effective length of the quadrupole. For $k_1 > 0$ the quadrupole focuses in x and defocuses in y and vice versa.

In many practical cases, the focal length f of the quadrupole is much larger than its effective length l_{eff}:

$$f = \frac{1}{k_1 l_{\mathrm{eff}}} \gg l_{\mathrm{eff}}\,, \tag{A.0.4}$$

and Eqs. A.0.2 and A.0.3 can be approximated by taking the limit $l_{\mathrm{eff}} \to 0$ while keeping $k_1 l_{\mathrm{eff}} = \text{const.}$:

$$\boldsymbol{R}_{\mathrm{Q,\,thin}} = \begin{pmatrix} 1 & 0 & 0 & 0 & 0 & 0 \\ -\frac{1}{f} & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & \frac{1}{f} & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \end{pmatrix}, \tag{A.0.5}$$

which is referred as the *thin lens approximation*.

Sector dipole magnet. A sector dipole has its end faces perpendicular to the design orbit. For a dipole dispersing in the horizontal plane x–s with a deflection angle α and bending radius ρ, the transfer matrix is:

$$\boldsymbol{R}_{\mathrm{SD}} = \begin{pmatrix} \cos\alpha & \rho\sin\alpha & 0 & 0 & 0 & \rho(1-\cos\alpha) \\ -\frac{1}{\rho}\sin\alpha & \cos\alpha & 0 & 0 & 0 & \sin\alpha \\ 0 & 0 & 1 & \rho\alpha & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 \\ -\sin\alpha & \rho(\cos\alpha - 1) & 0 & 0 & 1 & -\rho(\alpha - \sin\alpha) \\ 0 & 0 & 0 & 0 & 0 & 1 \end{pmatrix}. \tag{A.0.6}$$

Rectangular dipole magnet. In a rectangular dipole magnet, the entrance and exit faces are parallel to each other and, therefore, not perpendicular to the design orbit. If the end faces are symmetrically oriented respect to the reference trajectory, so that they form an angle of $\alpha/2$ that is half of the total bending angle α, its matrix can be described by introducing an additional edge focusing to the sector dipole magnet described above:

$$\boldsymbol{R}_{\mathrm{RD}} = \boldsymbol{F}_{\boldsymbol{\alpha/2}} \cdot \boldsymbol{R}_{\mathrm{SD}} \cdot \boldsymbol{F}_{\boldsymbol{\alpha/2}}, \tag{A.o.7}$$

where $\boldsymbol{F}_{\boldsymbol{\alpha/2}}$ is given by the matrix:

$$\boldsymbol{F} = \begin{pmatrix} 1 & 0 & 0 & 0 & 0 & 0 \\ \frac{\tan \alpha/2}{\rho} & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & -\frac{\tan \alpha/2}{\rho} & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \end{pmatrix}. \tag{A.o.8}$$

Appendix B.

Waveguide-attenuation calculation

In the following, the relation between the power P_{meas} measured at the directional coupler and the power actually fed into the structure $P_L + P_R$ is estimated. Figure B.1 shows a schematic representation the waveguide network of the section of interest. Figure B.2 at the end of this appendix shows a more realistic CAD model.

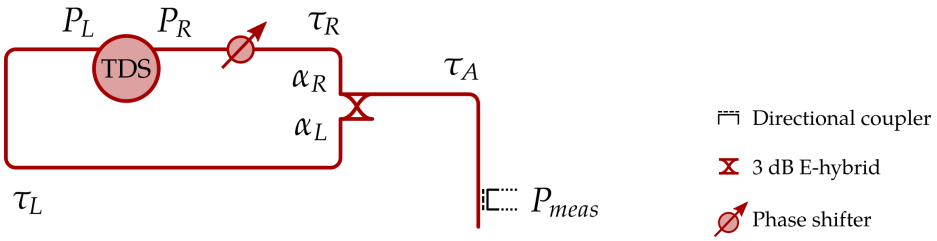

Figure B.1.: Schematic representation the waveguide network from the 3-dB splitter to the E-rotator.

The waveguide attenuation is estimated to be roughly -0.1 dB/m. The attenuation between the directional coupler and the 3-dB E-hybrid is $\tau_A \approx -0.1$ dB, and that of the left and right split arms is $\tau_L \approx -0.2$ dB and $\tau_R \approx -0.1$ dB, respectively. The terms α_L and α_R represent the fraction of power directed towards the left and the right split arms, respectively. The power that reaches the left and right TDS inputs is:

$$P_L = P_{meas} \cdot 10^{\frac{\tau_A}{10}} \cdot \alpha_L \cdot 10^{\frac{\tau_L}{10}} \tag{B.o.1a}$$

$$P_R = P_{meas} \cdot 10^{\frac{\tau_A}{10}} \cdot \alpha_R \cdot 10^{\frac{\tau_R}{10}}. \tag{B.o.1b}$$

On account of the power imbalance measured in Section 3.3, the fractions α_L and α_R can be determined. The obtained values are summarised in Table B, together with the estimated waveguide attenuation of each section. Finally the total power reaching the structure is:

$$P_L + P_R = P_{meas} \cdot 10^{\frac{\tau_A}{10}} \left(\alpha_L \cdot 10^{\frac{\tau_L}{10}} + \alpha_R \cdot 10^{\frac{\tau_R}{10}} \right) = P_{meas} \cdot \alpha_{TDS} \approx P_{meas} \cdot 0.94. \tag{B.o.2}$$

Parameter	Symbol	Units	Value
Attenuation before power splitting	τ_A	dB	-0.1
Power fraction left	α_L	%	56
Power fraction right	α_R	%	44
Attenuation left arm	τ_L	dB	-0.2
Attenuation right arm	τ_R	dB	-0.1
Power fraction at TDS	α_{TDS}	%	94

Table B.1.: Estimated RF-power transmission and attenuation at different locations of the RF-waveguide structure before the PolariX-TDS input ports.

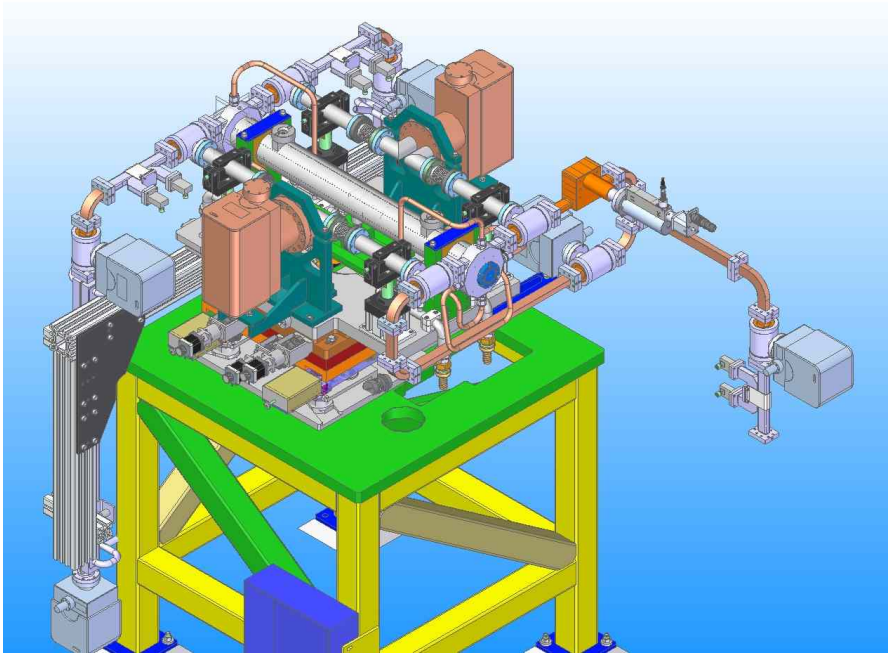

Figure B.2.: CAD model of the PolariX-TDS system around the RF structure.

Bibliography

Aad, G., T. Abajyan, B. Abbott, J. Abdallah, S. Abdel Khalek, A. A. Abdelalim, O. Abdinov, R. Aben, B. Abi, M. Abolins et al. (2012). 'Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC'. In: *Physics Letters, Section B: Nuclear, Elementary Particle and High-Energy Physics* 716.1, pp. 1–29. DOI: 10.1016/j.physletb.2012.08.020. arXiv: 1207.7214 (cit. on p. 1).

Ackermann, W., G. Asova, V. Ayvazyan, A. Azima, Nicoleta Baboi, J. Bähr, Vladimir Balandin, Bolko Beutner, A. Brandt, Andy Bolzmann et al. (2007). 'Operation of a free-electron laser from the extreme ultraviolet to the water window'. In: *Nature Photonics* 1, pp. 336–342. DOI: 10.1038/nphoton.2007.76 (cit. on p. 76).

Adli, Erik, J. P. Delahaye, Spencer J. Gessner, Mark J. Hogan, Tor O. Raubenheimer, W. An, Chandrashekhar Joshi and Warren B. Mori (2013). 'A beam driven plasma-wakefield linear collider from Higgs factory to multi-TeV'. In: *Proceedings of the Snowmass Process CSS2013*. Pp. 3791–3793. DOI: https://doi.org/10.48550/arXiv.1308.1145. arXiv: 1308.1145 (cit. on p. 2).

Agapov, Igor, G. Geloni, Sergey Tomin and Igor Zagorodnov (2014). 'OCELOT: A software framework for synchrotron light source and FEL studies'. In: *Nuclear Instruments and Methods in Physics Research, Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 768, pp. 151–156. DOI: 10.1016/j.nima.2014.09.057 (cit. on pp. 131, 185).

Aicheler, M., P. Burrows, M. Draper, T. Garvey, P. Lebrun, K. Peach, N. Phinney, H. Schmickler, D. Schulte and N. Toge (2012). *A Multi-TeV Linear Collider Based on CLIC Technology*. Tech. rep. CERN-2012-007. Geneva, Switzerland: CERN. DOI: 10.5170/CERN-2012-007 (cit. on pp. 2, 3, 103).

Albert, Félicie, M. E. Couprie, Alexander Debus, Mike C. Downer, Jérôme Faure, Alessandro Flacco, Leonida A. Gizzi, Thomas Grismayer, Axel Huebl, Chandrashekhar Joshi et al. (2021). '2020 Roadmap on Plasma Accelerators'. In: *New Journal of Physics* 23. DOI: 10.1088/1367-2630/abcc62 (cit. on p. 70).

Alesini, David, Giampiero Di Pirro, Luca Ficcadenti, Andrea Mostacci, Luigi Palumbo, James B. Rosenzweig and Cristina Vaccarezza (2006). 'RF deflector design and measurements for the longitudinal and transverse phase space characterization at SPARC'. In: *Nuclear Instruments and Methods in Physics Research, Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 568.2, pp. 488–502. DOI: 10.1016/j.nima.2006.07.050 (cit. on p. 38).

Altarelli, Massimo, Reinhard Brinkmann, Majed Chergui, Winfried Decking, Barry Dobson, Stefan Düsterer, Gerhard Grübel, Walter Graeff, Heinz Graafsma, Janos Hajdu et al. (2007). *The European X-Ray Free-Electron Laser. Technical design report.* Tech. rep. DESY-2006-097. Hamburg, Germany: DESY. DOI: 10.3204/DESY_06-097 (cit. on pp. 1, 76).

Altenmueller, O. H., R. R. Larsen and Gregory A. Loew (1964). 'Investigations of Traveling-Wave Separators for the Stanford Two-mile Linear Accelerator'. In: *Review of Scientific Instruments* 35.4, pp. 438–442. DOI: 10.1063/1.1718840 (cit. on pp. 4, 38, 40, 79).

Andonian, G., A. Cook, M. Dunning, E. Hemsing, G. Marcus, A. Murokh, S. Reiche, D. Schiller, James B. Rosenzweig, M. Babzien et al. (2009). 'Observation of coherent terahertz edge radiation from compressed electron beams'. In: *Physical Review Special Topics - Accelerators and Beams* 12.3. DOI: 10.1103/PhysRevSTAB.12.030701 (cit. on p. 37).

Andrews, H. L., F. Bakkali Taheri, J. Barros, R. Bartolini, V. Bharadwaj, C. Clarke, N. Delerue, G. Doucas, N. Fuster-Martinez, M. Vieille-Grosjean et al. (2014). 'Reconstruction of the time profile of 20 . 35 GeV , subpicosecond long electron bunches by means of coherent Smith-Purcell radiation'. In: *Physical Review Special Topics - Accelerators and Beams* 15.5, pp. 1–13. DOI: 10.1103/PhysRevSTAB.17.052802 (cit. on p. 37).

Arkani-Hamed, Nima, Savas Dimopoulos and Gia Dvali (1998). 'The hierarchy problem and new dimensions at a millimeter'. In: *Physics Letters, Section B: Nuclear, Elementary Particle and High-Energy Physics* 429.3-4, pp. 263–272. DOI: 10.1016/S0370-2693(98)00466-3. arXiv: 9803315 [hep-ph] (cit. on p. 2).

Ayvazyan, V., Nicoleta Baboi, J. Bähr, Vladimir Balandin, Bolko Beutner, A. Brandt, I. Bohnet, A. Bolzmann, Reinhard Brinkmann, O. I. Brovko et al. (2006). 'First operation of a free-electron laser generating GW power radiation at 32 nm wavelength'. In: *European Physical Journal D* 37.2, pp. 297–303. DOI: 10.1140/epjd/e2005-00308-1 (cit. on p. 76).

Baboi, Nicoleta, J. Lund-Nielsen, D. Noelle, W. Riesch, T. Traber, J. Kruse and M. Wendt (2006). 'Resolution studies at beam position monitors at the FLASH facility at DESY'. In: *AIP Conference Proceedings* 868, pp. 227–237. DOI: 10.1063/1.2401409 (cit. on p. 86).

Bane, Karl (1990). *Simulating longitudinal phase space in the SLC, from the damping rings to the final focus*. Tech. rep. SLAC/AP-80. Stanford, CA, USA: SLAC (cit. on p. 46).

Barov, N., James B. Rosenzweig, M. E. Conde, W. Gai and J. G. Power (2000). 'Observation of plasma wakefield acceleration in the underdense regime'. In: *Physical Review Special Topics - Accelerators and Beams* 3.1, p. 011301. DOI: 10.1103/physrevstab.3.011301 (cit. on p. 73).

Behrens, Christopher (2015). *Report on the design of the electron beamline of FLASHForward*. Tech. rep. Hamburg, Germany: DESY (cit. on p. 81).

Behrens, Christopher, F.-J. Decker, Yuantao Ding, Valery A. Dolgashev, Josef Frisch, Zhirong Huang, Patrik Krejcik, Henrik Loos, A. Lutman, T. J. Maxwell et al. (2014). 'Few-femtosecond time-resolved measurements of X-ray free-electron lasers'. In: *Nature Communications* 5. DOI: 10.1038/ncomms4762 (cit. on p. 38).

Behrens, Christopher, N Gerasimova, Christopher Gerth, Bernhard Schmidt, E A Schneidmiller, S Serkez, Stephan Wesch and M V Yurkov (2012). 'Constraints on photon pulse duration from longitudinal electron beam diagnostics at a soft x-ray free-electron laser'. In: *Physical Review Special Topics - Accelerators and Beams* 15.3, p. 030707. DOI: 10.1103/PhysRevSTAB.15.030707 (cit. on p. 79).

Behrens, Christopher and Christopher Gerth (2009). 'On the Limitations of Longitudinal Phase Space Measurements usign a Transverse Deflecting Structure'. In: *Proceedings of the 9th European Workshop on Beam Diagnostics and Instrumentation for Particle Accelerators - DIPAC09*. Basel, Switzerland: JACoW, p. 3 (cit. on p. 51).

Benedikt, Michael, Paul Collier, V. Mertens, John Poole and Karlheinz Schindl (2004). *LHC Design Report*. Tech. rep. CERN-2004-003. Geneva, Switzerland: CERN. DOI: 10.5170/CERN-2004-003-V-3 (cit. on p. 68).

Benedikt, Michael, Daniel Schulte and Frank Zimmermann (2015). 'Optimizing integrated luminosity of future hadron colliders'. In: *Physical Review Special Topics - Accelerators and Beams* 18.10, p. 101002. DOI: 10.1103/PhysRevSTAB.18.101002 (cit. on p. 2).

Bertone, Gianfranco, Dan Hooper and Joseph Silk (2005). 'Particle dark matter: Evidence, candidates and constraints'. In: *Physics Reports* 14.5, pp. 279–390. DOI: 10.1016/j.physrep.2004.08.031. arXiv: 0404175 [hep-ph] (cit. on p. 2).

Blumenfeld, Ian, Christopher E. Clayton, Franz Josef Decker, Mark J. Hogan, Chengkun Huang, Rasmus Ischebeck, Richard Iverson, Chandrashekhar Joshi, Thomas Katsouleas, Neil Kirby et al. (2007). 'Energy doubling of 42 GeV electrons in a metre-scale plasma wakefield accelerator'. In: *Nature* 445, pp. 741–744. DOI: 10.1038/nature05538 (cit. on pp. 3, 71, 73).

Brau, James, Yasuhiro Okada and Nicholas Walker (2007). *ILC Reference Design Report Volume 1: Executive Summary*. Tech. rep. DESY-07-046, ILC GDE. Hamburg, Germany: DESY. arXiv: 0712.1950 (cit. on p. 76).

Brown, Karl L. (1982). *A First- and Second-Order Matrix Theory for the Design of Beam Transport Systems and Charged Particle Spectrometers*. Tech. rep. SLAC-75, Rev.4 UC-28 (M). Stanford, CA, USA: SLAC, pp. 72–134 (cit. on pp. 9, 13–15, 25, 35, 44).

Buon, Jean (1993). 'Multi-dimensional beam emittance and β-functions'. In: *Proceedings of the 15th Particle Accelerator Conference - PAC1993*. Vol. 1. Washington DC, USA: JACoW Publishing, pp. 469–471. DOI: 10.1109/pac.1993.308879 (cit. on p. 23).

Canetti, Laurent, Marco Drewes and Mikhail Shaposhnikov (2012). 'Matter and antimatter in the universe'. In: *New Journal of Physics* 14.9, p. 095012. DOI: 10.1088/1367-2630/14/9/095012. arXiv: 1204.4186 (cit. on p. 2).

Chatrchyan, S., V. Khachatryan, A. M. Sirunyan, A. Tumasyan, W. Adam, E. Aguilo, T. Bergauer, M. Dragicevic, J. Erö, C. Fabjan et al. (2012). 'Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC'. In: *Physics Letters, Section B:*

Nuclear, Elementary Particle and High-Energy Physics 716.1, pp. 30–61. DOI: 10.1016/j.physletb.2012.08.021. arXiv: 1207.7235 (cit. on p. 1).

Chen, Francis F. (1984). *Introduction to Plasma Physics and Controlled Fusion. Volume 1: Plasma Physics*. 2nd. New York, NY, USA: Plenum Publishing Corporation. DOI: https://doi.org/10.1007/978-1-4757-5595-4 (cit. on pp. 53, 55).

Chen, Pisin, John M. Dawson, R. W. Huff and Thomas Katsouleas (1985). 'Acceleration of Electrons by the Interaction of a Bunched Electron Beam with a Plasma'. In: *Physical Review Letters* 54.7, pp. 693–696. DOI: https://doi.org/10.1103/PhysRevLett.54.693 (cit. on p. 2).

Christie, Florian, Juliane Roensch-Schulenburg and Mathias Vogt (2019). 'A PolariX-TDS for the FLASH2 Beamline Layout of the TDS Diagnostic'. In: *Proceedings of the 39th International Free-Electron Laser Conference - FEL2019*. Hamburg, Germany: JACoW Publishing, pp. 328–331. DOI: 10.18429/JACoW-FEL2019-WEP006 (cit. on p. 103).

Clayton, Christopher E., Erik Adli, J. Allen, W. An, C. I. Clarke, Sebastian Corde, J. Frederico, Spencer J. Gessner, S. Z. Green, Mark J. Hogan et al. (2016). 'Self-mapping the longitudinal field structure of a nonlinear plasma accelerator cavity'. In: *Nature Communications* 7.12483. DOI: 10.1038/ncomms12483 (cit. on pp. 3, 71, 73, 199).

Colby, Eric R. and L. K. Len (2017). 'Roadmap to the future'. In: *Reviews of Accelerator Science and Technology* 9, pp. 1–18. DOI: 10.1142/S1793626816300012 (cit. on p. 3).

Courant, E.D and H.S Snyder (1958). 'Theory of the alternating-gradient synchrotron'. In: *Annals of Physics* 3.1, pp. 1–48. DOI: 10.1016/0003-4916(58)90012-5 (cit. on pp. 9, 16, 17).

Craievich, Paolo, Markus Bopp, H. H. Braun, Alessandro Citterio, Renato Fortunati, Romain Ganter, T. Kleeb, F. Marcellini, Marco Pedrozzi, Eduard Prat et al. (2020). 'Novel X-band transverse deflection structure with variable polarization'. In: *Physical Review Accelerators and Beams* 23.11, p. 112001. DOI: 10.1103/PhysRevAccelBeams.23.112001 (cit. on pp. 4, 104, 106, 107, 109).

Craievich, Paolo, R Ischebeck, Florian Loehl, G L Orlandi and Eduard Prat (2013). 'Transverse Deflecting Structures for Bunch Length and Slice Emittance Measurements on SwissFEL'. In: *Proceedings of the 35th Free-Electron Laser Conference - FEL2013*. New York, NY, USA: JACoW Publishing, pp. 236–241 (cit. on p. 38).

Crick, Francis and James Watson (1953). 'Molecular Structure of Nucleic Acids'. In: *Nature* 171, pp. 737–738. DOI: 10.1038/171737a0 (cit. on p. 1).

Cros, Brigitte, Patric Muggli, Carl B. Schroeder, S. M. Hooker, Philippe Piot, R. J. England, Spencer J. Gessner, Jorge Vieira, Edda Gschwendtner, Jean-Luc Vay et al. (2019). 'Towards an Advanced Linear International Collider'. In: *ALEGRO collaboration*. arXiv: 1901.10370 (cit. on p. 3).

Curcio, A., M. Bergamaschi, R. Corsini, W. Farabolini, D. Gamba, L. Garolfi, R. Kieffer, T. Lefevre, S. Mazzoni, K. Fedorov et al. (2020). 'Noninvasive bunch length measurements exploiting Cherenkov diffraction radiation'. In: *Physical Review Accelerators and Beams* 23.2, p. 022802. DOI: 10.1103/PhysRevAccelBeams.23.022802 (cit. on p. 37).

D'Arcy, Richard, Alexander Aschikhin, Simon Bohlen, Gregor J. Boyle, T. Brümmer, James Chappell, Severin Diederichs, Brian Foster, Mathew James Garland, L. Goldberg et al. (2019). 'FLASHForward: Plasma wakefield accelerator science for high-average-power applications'. In: *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 377. DOI: 10.1098/rsta.2018.0392 (cit. on pp. 4, 75).

D'Arcy, Richard, James Chappell, Judita Beinortaite, Severin Diederichs, Gregor J. Boyle, Brian Foster, Mathew James Garland, Pau Gonzalez Caminal, Carl Andreas Lindstrøm, Gregor Loisch et al. (2022). 'Recovery time of a plasma-wakefield accelerator'. In: *Nature* 603.7899, pp. 58–62. DOI: 10.1038/s41586-021-04348-8 (cit. on pp. 72, 73, 75).

D'Arcy, Richard, Stephan Wesch, Alexander Aschikhin, Simon Bohlen, Christopher Behrens, Mathew James Garland, L. Goldberg, Pau Gonzalez Caminal, Alex Knetsch, Vladyslav Libov et al. (2019). 'Tunable Plasma-Based Energy Dechirper'. In: *Physical Review Letters* 122.3, p. 034801. DOI: 10.1103/physrevlett.122.034801 (cit. on pp. 72, 73).

Dale, John and Lucas Schaper (2017). *Tender for FLASHForward Target Chamber and Differential Pumping System Vacuum pumps*. Tech. rep. Hamburg, Germany: DESY, p. 3 (cit. on p. 100).

Dawson, John M. (1959). 'Nonlinear Electron Oscillation in a Cold Plasma'. In: *Physical Review* 113.2, pp. 383–387. DOI: https://doi.org/10.1103/RevModPhys.55.403 (cit. on p. 56).

Dawson, John M. (1983). 'Particle simulation of plasmas'. In: *Reviews of Modern Physics* 55.2, pp. 403–447. DOI: 10.1103/RevModPhys.55.403 (cit. on p. 61).

Di Mitri, Simone and M. Cornacchia (2014). 'Electron beam brightness in linac drivers for free-electron-lasers'. In: *Physics Reports* 539.1, pp. 1–48. DOI: 10.1016/j.physrep.2014.01.005 (cit. on p. 2).

Diederichs, Severin, Carlo Benedetti, Axel Huebl, Rémi Lehe, Andrew Myers, Alexander Sinn, Jean-Luc Vay, Weiqun Zhang and Maxence Thévenet (2021). 'HiPACE++: a portable, 3D quasi-static Particle-in-Cell code'. In: *arXiv*, pp. 1–10. arXiv: 2109.10277 (cit. on p. 199).

Ding, Yuantao, Christopher Behrens, Paul Emma, Josef Frisch, Zhirong Huang, Henrik Loos, Patrik Krejcik and M. H. Wang (2011). 'Femtosecond x-ray pulse temporal characterization in free-electron lasers using a transverse deflector'. In: *Physical Review Special Topics - Accelerators and Beams* 14.12, p. 120701. DOI: 10.1103/PhysRevSTAB.14.120701 (cit. on p. 38).

Dohlus, Martin and Torsten Limberg (1997). 'Emittance growth due to wake fields on curved bunch trajectories'. In: *Nuclear Instruments and Methods in Physics Research, Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 393.1-3, pp. 494–499. DOI: 10.1016/S0168-9002(97)00552-4 (cit. on p. 30).

Dohlus, Martin, Torsten Limberg and Paul Emma (2005). 'Bunch Compression for Linac-based FELs'. In: *International Committee for Future Accelerators - Beam Dynamics Newsletter* 38, pp. 15–36 (cit. on pp. 1, 27, 28, 30).

Dolgashev, Valery A., Gordon Bowden, Yuantao Ding, Paul Emma, Patrick Krejcik, James Lewandowski, Cecile Limborg, Michael Litos, Juwen Wang and Dao Xiang (2014). 'Design and application of multimegawatt X-band deflectors for femtosecond electron beam diagnostics'. In: *Physical Review Special Topics - Accelerators and Beams* 17.10, p. 102801. DOI: 10.1103/PhysRevSTAB.17.102801 (cit. on p. 38).

Dorda, Ulrich, Ralph Assmann, Reinhard Brinkmann, Klaus Floettmann, I. Hartl, Markus Hüning, Franz Kärtner, A. Fallahi, Barbara Marchetti, Y. Nie et al. (2016). 'SINBAD—The accelerator R&D facility under construction at DESY'. In: *Nuclear Instruments and Methods in Physics Research, Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 829, pp. 233–236. DOI: 10.1016/j.nima.2016.01.067 (cit. on p. 103).

Dowell, D. H., P. R. Bolton, J. E. Clendenin, Paul Emma, S. M. Gierman, W. S. Graves, C. G. Limborg, B. F. Murphy and J. F. Schmerge (2003). 'Slice emittance measurements at the SLAC gun test facility'. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 507.1, pp. 327–330. DOI: 10.1016/S0168-9002(03)00939-2 (cit. on pp. 3, 71).

Downer, Mike C., R. Zgadzaj, A. Debus, U. Schramm and M. C. Kaluza (2018). 'Diagnostics for plasma-based electron accelerators'. In: *Reviews of Modern Physics* 90.3, p. 035002. DOI: 10.1103/RevModPhys.90.035002 (cit. on p. 70).

Edwards, H., Christopher Behrens and E. Harms (2010). '3.9 GHz Cavity Module for Linear Bunch Compression at FLASH'. In: *Proceedings of the 25th Linear Accelerator Conference - LINAC2010*. Tsukuba, Japan: JACoW Publishing, pp. 41–45 (cit. on p. 77).

Elder, F. R., A. M. Gurewitsch, R. V. Langmuir and H. C. Pollock (1947). 'Radiation from Electrons in a Synchrotron'. In: *Physical Review* 71.11, pp. 829–830. DOI: 10.1103/PhysRev.71.829.5 (cit. on p. 1).

Ellenberger, Urs, Ludwig Paly, Heinrich Blumer, Charles Zumbach, Florian Loehl, Markus Bopp and Hansrudolf Fitze (2013). 'Status of the manufacturing process for the SwissFEL c-band accelerating structures'. In: *Proceedings of the 35th Free-Electron Laser Conference - FEL2013*. New York, NY, USA: JACoW Publishing, pp. 245–249 (cit. on p. 103).

Emma, Paul, Josef Frisch and Patrik Krejcik (2000). *A transverse RF deflecting structure for bunch length and phase space diagnostics*. Tech. rep. LCLS-TN-00-12. Stanford, CA, USA: SLAC (cit. on pp. 38, 41, 46).

Esarey, E., P Catravas and Wim P. Leemans (2001). 'Betatron radiation from electron beams in plasma focusing channels'. In: *AIP Conference Proceedings* 569, pp. 473–486. DOI: 10.1063/1.1384377 (cit. on p. 57).

Esarey, E., Carl B. Schroeder and Wim P. Leemans (2009). 'Physics of laser-driven plasma-based electron accelerators'. In: *Reviews of Modern Physics* 81.3, pp. 1229–1285. DOI: 10.1103/RevModPhys.81.1229 (cit. on pp. 3, 61).

Evans, Lyndon and Philip Bryant (2008). 'LHC Machine'. In: *Journal of Instrumentation* 3.08, S08001. DOI: 10.1088/1748-0221/3/08/S08001 (cit. on p. 1).

Ferran Pousa, Ángel (2019). 'Novel concepts and theoretical studies for high-quality plasma-based accelerators'. PhD thesis. Universität Hamburg. DOI: 10.3204/PUBDB-2020-04106 (cit. on p. 60).

Ferran Pousa, Ángel, Ralph Assmann and Alberto Martinez de la Ossa (2019). 'Wake-T: A fast particle tracking code for plasma-based accelerators'. In: *Journal of Physics: Conference Series* 1350.1. DOI: 10.1088/1742-6596/1350/1/012056 (cit. on pp. 62, 63, 66).

Ferrario, M., M. Migliorati and L. Palumbo (2014). 'Space charge effects'. In: *Proceedings of the CAS-CERN Accelerator School: Advanced Accelerator Physics*. Ed. by W. Herr. Geneva, Switzerland: CERN, pp. 331–356. DOI: 10.5170/CERN-2014-009.331. arXiv: 1410.7991 (cit. on pp. 29, 30).

Floettmann, Klaus (2003). 'Some basic features of the beam emittance'. In: *Physical Review Special Topics - Accelerators and Beams* 6.3, pp. 80–86. DOI: 10.1103/PhysRevSTAB.6.034202 (cit. on pp. 10, 20).

Floettmann, Klaus (2014). 'Adiabatic matching section for plasma accelerated beams'. In: *Physical Review Special Topics - Accelerators and Beams* 17.5, p. 054402. DOI: 10.1103/PhysRevSTAB.17.054402 (cit. on p. 68).

Franklin, Rosalind K. and R. G. Gossling (1953). 'Molecular Configuration in Sodium Thymonucleate'. In: *Nature* 171, pp. 740–741. DOI: 10.1038/171740a0 (cit. on p. 1).

Frederico, Joel, Mark J. Hogan and Tor Raubenheimer (2017). 'Beam match and emittance saturation within plasma wakefield acceleration'. In: *AIP Conference Proceedings* 1812. DOI: 10.1063/1.4975881 (cit. on p. 68).

Ganter, Romain, Rafael Abela, Arturo Alarcon, Jürgen Alex, Christopher Arrell, Vladimir Arsov, Simona Bettoni, Markus Bopp, Christoph Bostedt, Hans Heinrich Braun et al. (2019). 'The SwissFEL soft X-ray free-electron laser beamline: Athos'. In: *Journal of Synchrotron Radiation* 26, pp. 1073–1084. DOI: 10.1107/S1600577519003928 (cit. on pp. 37, 103).

Garland, Mathew James, Gabriele Tauscher, Simon Bohlen, Gregor J. Boyle, Richard D'Arcy, L. Goldberg, Krjstian Põder, Lucas Schaper, Bernhard Schmidt and Jens Osterhoff (2021). 'Combining laser interferometry and plasma spectroscopy for spatially resolved high-sensitivity plasma density measurements in discharge capillaries'. In: *Review of Scientific Instruments* 92.1, p. 013505. DOI: 10.1063/5.0021117 (cit. on pp. 75, 84, 147, 199, 200, 204).

Gerigk, Frank (2012). 'Cavity types'. In: *Proceedings of the CAS-CERN Accelerator School: RF for accelerators*. Ed. by Roger Bailey. Ebeltoft, Denmark: CERN, pp. 277–298. DOI: 10.5170/CERN-2011-007. arXiv: 1111.4897 (cit. on p. 40).

Goldstein, Herbert, Charles P. Jr. Poole and John L. Safko (2001). *Classical Mechanics*. 3rd ed. Addison Wesley (cit. on p. 18).

Gross, Matthias, J. Engel, J. Good, H. Huck, I. Isaev, G. Koss, M. Krasilnikov, O. Lishilin, Gregor Loisch, Y. Renier et al. (2018). 'Observation of the Self-Modulation Instability via Time-Resolved Measurements'. In: *Physical Review Letters* 120.14. DOI: 10.1103/PhysRevLett.120.144802 (cit. on p. 73).

Grudiev, Alexej (2016). *Design of compact high power RF components at X-band*. Tech. rep. CLIC-Note-1067. Geneva, Switzerland: CERN (cit. on pp. 4, 5, 103–105).

Grudiev, Alexej (2017). *RF Design of X-band TDS with variable polarization*. Tech. rep. CERN (cit. on p. 103).

Guetg, Marc Walter, Bolko Beutner, Eduard Prat and Sven Reiche (2015). 'Optimization of free electron laser performance by dispersion-based beam-tilt correction'. In: *Physical Review Special Topics - Accelerators and Beams* 18.3, p. 030701. DOI: 10.1103/PhysRevSTAB.18.030701 (cit. on pp. 29, 133).

Guetg, Marc Walter, F. J. Decker, Y. Ding, Z. Huang and T. J. Maxwell (2016). 'Measurement of advanced dispersion-based beam-tilt correction'. In: *Proceedings of the 7th International Particle Accelerator Conference - IPAC2016*, pp. 813–816. DOI: 10.18429/JACoW-IPAC2016-MOPOW045 (cit. on p. 133).

Hahn, H. (1963). 'Deflecting Mode in Circular Iris-Loaded Waveguides'. In: *Review of Scientific Instruments* 34.10, pp. 1094–1100. DOI: https://doi.org/10.1063/1.1718141 (cit. on p. 40).

Higgs, Peter W. (1964). 'Broken symmetries and the masses of gauge bosons'. In: *Physical Review Letters* 13.16, pp. 508–509. DOI: 10.1103/PhysRevLett.13.508 (cit. on p. 1).

Hilbert, Christian (2017). *VPD Dimensioning Report*. Tech. rep. Agilent Technologies GmbH (cit. on p. 100).

Hogan, Mark J., C. D. Barnes, Christopher E. Clayton, F. J. Decker, S. Deng, Paul Emma, Chengkun Huang, R. H. Iverson, D. K. Johnson, Chandrashekhar Joshi et al. (2005). 'Multi-GeV energy gain in a plasma-wakefield accelerator'. In: *Physical Review Letters* 95.5, p. 054802. DOI: 10.1103/PhysRevLett.95.054802 (cit. on p. 73).

Hogan, Mark J., Tor O. Raubenheimer, A. Seryi, Patric Muggli, Thomas Katsouleas, C. Huang, Wei Lu, W. An, Kenneth A. Marsh, Warren B. Mori et al. (2010). 'Plasma wakefield acceleration experiments at FACET'. In: *New Journal of Physics* 12.5, p. 055030. DOI: 10.1088/1367-2630/12/5/055030 (cit. on pp. 3, 71).

Hooker, S. M., R. Bartolini, S. P.D. Mangles, A. Tünnermann, L. Corner, J. Limpert, A. Seryi and R. Walczak (2014). 'Multi-pulse laser wakefield acceleration: A new route to efficient,

high-repetition-rate plasma accelerators and high flux radiation sources'. In: *Journal of Physics B: Atomic, Molecular and Optical Physics* 47.23, p. 234003. DOI: 10.1088/0953-4075/47/23/234003. arXiv: 1401.7874 (cit. on p. 3).

Huang, Zhirong, Karl Bane, Yuantao Ding and Paul Emma (2010). 'Single-shot method for measuring femtosecond bunch length in linac-based free-electron lasers'. In: *Physical Review Special Topics - Accelerators and Beams* 13.9, p. 092801. DOI: 10.1103/PhysRevSTAB.13.092801 (cit. on p. 46).

Hudson, Stephen, Jeffrey Larson, Stefan M Wild, David Bindel and John-luke Navarro (2021). *libEnsemble User Manual*. Tech. rep. Revision 0.8.0+dev. Argonne National Laboratory (cit. on p. 203).

Iselin, F. Christoph (1994). *The MAD Program (Methodical Accelerator Design) Physical Methods Manual*. Tech. rep. CERN/SL/92-?? (AP). Geneva, Switzerland: CERN, p. 79 (cit. on p. 15).

Joshi, Chandrashekhar, B. Blue, Christopher E. Clayton, E. Dodd, Chengkun Huang, Kenneth A. Marsh, Warren B. Mori, S. Wang, Mark J. Hogan, C. O'Connell et al. (2002). 'High energy density plasma science with an ultrarelativistic electron beam'. In: *Physics of Plasmas* 9.5, pp. 1845–1855. DOI: https://doi.org/10.1063/1.1455003 (cit. on pp. 3, 71).

Kallos, Efthymios, Tom Katsouleas, Wayne D. Kimura, Karl Kusche, Patric Muggli, Igor Pavlishin, Igor Pogorelsky, Daniil Stolyarov and Vitaly Yakimenko (2008). 'High-gradient plasma-wakefield acceleration with two subpicosecond electron bunches'. In: *Physical Review Letters* 100.7, p. 074802. DOI: 10.1103/PhysRevLett.100.074802 (cit. on p. 73).

Kandasamy, Kirthevasan, Karun Raju Vysyaraju, Willie Neiswanger, Biswajit Paria, Christopher R. Collins, Jeff Schneider, Barnabás Póczos and Eric P. Xing (2020). 'Tuning hyperparameters without grad students: Scalable and robust bayesian optimisation with dragonfly'. In: *Journal of Machine Learning Research* 21, pp. 1–25. arXiv: 1903.06694 (cit. on p. 203).

Kersevan, Roberto (2007). 'Analytical and numerical tools for vacuum systems.' In: *Proceedings of the CAS-CERN Accelerator School and ALBA Synchrotron Light Facility : Course on Vacuum in Accelerators*. Ed. by D. Brandt. Platja d'Aro, Spain: CERN, pp. 285–312. DOI: 10.5170/CERN-2007-003 (cit. on p. 100).

Kersevan, Roberto and Marton Ady (2019). 'Recent Developments of Monte-Carlo Codes Molflow+ and Synrad+'. In: *Proceedings of the 10th International Particle Conference - IPAC2019*. Melbourne, Australia: JACoW Publishing, pp. 1327–1330. DOI: 10.18429/JACoW-IPAC2019-TUPMP037 (cit. on p. 100).

Knetsch, Alex, Bridget Sheeran, Lewis Boulton, Pardis Niknejadi, Krjstian Põder, Lucas Schaper, Ming Zeng, Simon Bohlen, Gregor J. Boyle, Theresa Brümmer et al. (2021). 'Stable witness-beam formation in a beam-driven plasma cathode'. In: *Physical Review*

Accelerators and Beams 24.10, p. 101302. DOI: 10.1103/physrevaccelbeams.24.101302 (cit. on p. 75).

Krausz, Ferenc and Misha Ivanov (2009). 'Attosecond physics'. In: *Reviews of Modern Physics* 81.1, pp. 163–234. DOI: 10.1103/RevModPhys.81.163. arXiv: 1102.1291 (cit. on p. 57).

Krejcik, Patrik, F.-J. Decker, Yuantao Ding, Josef Frisch and Zhirong Huang (2013). 'Commissioning the new LCLS X-Band Transverse Deflecting Cavity with Femtosecond Resolution'. In: *Proceedings of the 2nd International Beam Instrumentation Conference - IBIC2013*. Oxford, UK: JACoW Publishing, pp. 308–311 (cit. on p. 38).

Kube, Gero, Sasa Bajt, A. P. Potylitsyn, L. G. Sukhikh, A. V. Vukolov, I. A. Artyukov and W. Lauth (2015). 'Transverse Beam Profile Imaging of Few-micrometer Beam Sizes Based on a Scintillator Screen'. In: *Proceedings of the 4th International Beam Instrumentation Conference - IBIC2015*. Melbourne, Australia: JACoW Publishing, pp. 330–334. DOI: DOI: 10.18429/JACoW-IBIC2015-TUPB012 (cit. on p. 86).

Lai, R., U. Happek and A. J. Sievers (1994). 'Measurement of the Longitudinal Asymmetry of a Charged Particle Bunch from the Coherent Synchrotron or Transition Radiation Spectrum'. In: *Physical Review E* 50.6, R4294–R4297. DOI: https://doi.org/10.1103/PhysRevE.50.R4294 (cit. on p. 37).

Libov, Vladyslav, Alexander Aschikhin, John Dale, Richard D'Arcy, Kai Ludwig, Alberto Martinez de la Ossa, Timon Johannes Mehrling, Jan-Hendrik Röckemann, Lucas Schaper, Bernhard Schmidt et al. (2018). 'FLASHForward X-2 : Towards beam quality preservation in a plasma booster'. In: *Nuclear Inst. and Methods in Physics Research, A* 909, pp. 80–83. DOI: 10.1016/j.nima.2018.02.063 (cit. on pp. 75, 137).

Limberg, Torsten, H. Weise, A. Molodozhentsev and V. Petrov (1996). 'The bunch compression system at the TESLA test facility FEL'. In: *Nuclear Instruments and Methods in Physics Research, Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 375.1-3, pp. 322–324. DOI: 10.1016/0168-9002(95)01470-5 (cit. on p. 77).

Lindstrøm, Carl Andreas (2019a). 'Emittance Growth and Preservation in a Plasma-based Linear Collider'. PhD thesis. arXiv: 1802.02750 (cit. on pp. 69, 73).

Lindstrøm, Carl Andreas (2019b). *Low-emittance measurement station at FLASHForward*. Tech. rep. Hamburg, Germany: DESY (cit. on p. 87).

Lindstrøm, Carl Andreas and Erik Adli (2016). 'Design of general apochromatic drift-quadrupole beam lines'. In: *Physical Review Accelerators and Beams* 19.7, p. 071002. DOI: 10.1103/PhysRevAccelBeams.19.071002 (cit. on p. 25).

Lindstrøm, Carl Andreas, Erik Adli, Gregor J. Boyle, R. Corsini, A. E. Dyson, W. Farabolini, S. M. Hooker, Martin Meisel, Jens Osterhoff, Jan-Hendrik Röckemann et al. (2018). 'Emittance Preservation in an Aberration-Free Active Plasma Lens'. In: *Physical Review Letters* 121.19, p. 194801. DOI: 10.1103/PhysRevLett.121.194801. arXiv: 1808.03691 (cit. on pp. 71, 75).

Lindstrøm, Carl Andreas, Erik Adli, J. Pfingstner, E. Marin and D. Schulte (2016). 'Transverse tolerances of a multi-stage plasma wakefield accelerator'. In: *Proceedings of the 7th International Particle Accelerator Conference - IPAC2016*. Busan, Korea: JACoW Publishing, pp. 2561–2564. DOI: DOI:10.18429/JACoW-IPAC2016-WEPMY009 (cit. on pp. 3, 69).

Lindstrøm, Carl Andreas, Judita Beinortaite, Jonas Björklund Svensson, Lewis Boulton, James Chappell, Mathew James Garland, Pau Gonzalez Caminal, Gregor Loisch, Felipe Peña, Lucas Schaper et al. (2021). 'Progress of the FLASHForward X-2 high-beam-quality, high-efficiency plasma-accelerator experiment'. In: *The European Physical Society Conference on High Energy Physics*. DOI: 10.22323/1.398.0880. arXiv: 2111.08384 (cit. on pp. 70, 75, 88, 137).

Lindstrøm, Carl Andreas, Richard D'Arcy, Mathew James Garland, Pau Gonzalez Caminal, Bernhard Schmidt, Sarah Schröder, Stephan Wesch and Jens Osterhoff (2020). 'Matching small beta functions using centroid jitter and two beam position monitors'. In: *Physical Review Accelerators and Beams* 23.5, p. 52802. DOI: 10.1103/PhysRevAccelBeams.23.052802. arXiv: 2002.06022 (cit. on pp. 5, 36, 86, 94).

Lindstrøm, Carl Andreas, Mathew James Garland, Sarah Schröder, Lewis Boulton, Gregor J. Boyle, James Chappell, Richard D'Arcy, Pau Gonzalez Caminal, Alex Knetsch, Vladyslav Libov et al. (2021). 'Energy-Spread Preservation and High Efficiency in a Plasma-Wakefield Accelerator'. In: *Physical Review Letters* 126.1, p. 14801. DOI: 10.1103/PhysRevLett.126.014801 (cit. on pp. 4, 72, 73, 75, 137, 154, 156, 204, 215, 218).

Lipka, D. (2009). 'Cavity BPM Designs, Related Electronics and Measured Performances'. In: *Proceedings of the 9th European Workshop on Beam Diagnostics and Instrumentation for Particle Accelerators - DIPAC2009*. Basel, Switzerland: JACoW Publishing, pp. 280–284 (cit. on p. 86).

Lipka, D., Nicoleta Baboi, D. Noelle, G. Petrosyan, S. Vilcins, R. Baldinger, R. Ditter, B. Keil, W. Koprek, R. Kramert et al. (2014). 'Flash undulator BPM commissioning and beam characterization results'. In: *Proceeding of the 3rd International Beam Instrumentation Conference - IBIC2014*. Monterey, CA, USA: JACoW Publishing, pp. 315–319 (cit. on p. 86).

Litos, M. D., R. Ariniello, C. E. Doss, K. Hunt-Stone and J. R. Cary (2019). 'Beam emittance preservation using Gaussian density ramps in a beam-driven plasma wakefield accelerator'. In: *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 377.2151. DOI: 10.1098/rsta.2018.0181 (cit. on p. 68).

Litos, Michael, Erik Adli, J. M. Allen, W. An, C. I. Clarke, Sebastian Corde, Christopher E. Clayton, Joel Frederico, Spencer J. Gessner, S. Z. Green et al. (2016). '9 GeV energy gain in a beam-driven plasma wakefield accelerator'. In: *Plasma Physics and Controlled Fusion* 58.3. DOI: 10.1088/0741-3335/58/3/034017 (cit. on p. 73).

Litos, Michael, Erik Adli, W. An, C. I. Clarke, Christopher E. Clayton, Sebastian Corde, J. P. Delahaye, R. J. England, A. S. Fisher, J. Frederico et al. (2014). 'High-efficiency acceleration of an electron beam in a plasma wakefield accelerator'. In: *Nature* 515.7525, pp. 92–95. DOI: 10.1038/nature13882 (cit. on pp. 3, 71, 73).

Lockmann, Nils Maris, Christopher Gerth, Bernhard Schmidt and Stephan Wesch (2020). 'Noninvasive THz spectroscopy for bunch current profile reconstructions at MHz repetition rates'. In: *Physical Review Accelerators and Beams* 23.11, p. 112801. DOI: 10.1103/PhysRevAccelBeams.23.112801 (cit. on p. 37).

Loehl, Florian (2005). 'Measurements of Transverse Emittance at the VUV-FEL'. Diplom. Universität Hamburg. DOI: 10.3204/DESY-THESIS-2005-014 (cit. on pp. 31, 32, 79).

Loisch, Gregor, Galina Asova, Prach Boonpornprasert, Reinhard Brinkmann, Ye Chen, Johannes Engel, James Good, Matthias Gross, Florian Grüner, Holger Huck et al. (2018). 'Observation of High Transformer Ratio Plasma Wakefield Acceleration'. In: *Physical Review Letters* 121.6, p. 064801. DOI: 10.1103/PhysRevLett.121.064801 (cit. on pp. 67, 73).

Loos, Henrik (n.d.). *Reconstruction of a Filamentary Phase Space from two Projections*. Tech. rep. Stanford, CA, USA: SLAC (cit. on p. 46).

Lorbeer, B. (2006). *Stability of the master oscillator for FLASH at DESY*. Tech. rep. TESLA-FEL 2006-12. Hamburg, Germany: DESY (cit. on p. 79).

Lotov, Konstantin V. (2005). 'Efficient operating mode of the plasma wakefield accelerator'. In: *Physics of Plasmas* 12.5, p. 053105. DOI: 10.1063/1.1889444 (cit. on p. 67).

Lu, Wei, Chengkun Huang, Miaomiao Zhou, Warren B. Mori and Thomas Katsouleas (2006). 'Nonlinear theory for relativistic plasma wakefields in the blowout regime'. In: *Physical Review Letters* 96.16, p. 165002. DOI: 10.1103/PhysRevLett.96.165002 (cit. on pp. 61, 62, 65).

Madey, John M.J. (1971). 'Stimulated emission of bremsstrahlung in a periodic magnetic field'. In: *Journal of Applied Physics* 42.5, pp. 1906–1913. DOI: 10.1063/1.1660466 (cit. on p. 1).

Marchetti, Barbara, Ralph Assmann, Bolko Beutner, J. Branlard, Florian Christie, Richard D'Arcy, W. Decking, U. Dorda, J. Herrmann, Matthias Hoffmann et al. (2017). 'X-Band TDS Project'. In: *Proceedings of the 8th International Particle Accelerator Conference - IPAC2017*. Copenhagen, Denmark: JACoW Publishing, pp. 184–187. DOI: https://doi.org/10.18429/JACoW-IPAC2017-MOPAB044 (cit. on p. 103).

Marchetti, Barbara, Alexej Grudiev, Paolo Craievich, Ralph Assmann, H. H. Braun, Nuria Catalan Lasheras, Florian Christie, Richard D'Arcy, Renato Fortunati, Romain Ganter et al. (2021). 'Experimental demonstration of novel beam characterization using a polarizable X-band transverse deflection structure'. In: *Scientific Reports* 11.1, pp. 1–14. DOI: 10.1038/s41598-021-82687-2 (cit. on p. 4).

Marsh, Kenneth A., Christopher E. Clayton, D. K. Johnson, Chengkun Huang, Chandrashekhar Joshi, Wei Lu, Warren B. Mori, M. Zhou, C. D. Barnes, Franz Josef Decker et al. (2005). 'Beam matching to a plasma wake field accelerator using a ramped density profile at the plasma boundary'. In: *Proceedings of the IEEE Particle Accelerator Conference - PAC2005*, pp. 2702–2704. DOI: 10.1109/PAC.2005.1591234 (cit. on p. 68).

Martinez de la Ossa, A., Timon Johannes Mehrling and Jens Osterhoff (2018). 'Intrinsic Stabilization of the Drive Beam in Plasma-Wakefield Accelerators'. In: *Physical Review Letters* 121.6, p. 64803. DOI: 10.1103/PhysRevLett.121.064803. arXiv: 1806.11352 (cit. on p. 70).

Mehrling, Timon Johannes (2014). 'Theoretical and numerical studies on the transport of transverse beam quality in plasma-based accelerators'. PhD thesis. Universität Hamburg. DOI: dx.doi.org/10.3204/DESY-THESIS-2014-040 (cit. on pp. 56–59).

Mehrling, Timon Johannes, R. A. Fonseca, Alberto Martinez de la Ossa and Jorge Vieira (2017). 'Mitigation of the Hose Instability in Plasma-Wakefield Accelerators'. In: *Physical Review Letters* 118.17, pp. 1–5. DOI: 10.1103/PhysRevLett.118.174801. arXiv: 1704.02188 (cit. on p. 70).

Mehrling, Timon Johannes, J. Grebenyuk, F. S. Tsung, Klaus Floettmann and Jens Osterhoff (2012). 'Transverse emittance growth in staged laser-wakefield acceleration'. In: *Physical Review Special Topics - Accelerators and Beams* 15.11, p. 111303. DOI: 10.1103/PhysRevSTAB.15.111303 (cit. on pp. 3, 68).

Meisel, Martin (2018). 'Emittance Measurement of Electron Beams from Laser Wakeeld Acceleration using an Active Plasma Lens'. MSc thesis. Universität Hamburg. DOI: 10.3204/PUBDB-2018-03854 (cit. on p. 75).

Meyer, Kenneth R., Glen R. Hall and Dan Offin (2009). *Introduction to Hamiltonian Dynamical Systems and the N-Body Problem*. 2nd Ed. New York, NY, USA: Springer. DOI: 10.1007/978-0-387-09724-4 (cit. on p. 14).

Minty, Michiko G. and Frank Zimmermann (2003). *Measurement and Control of Charged Particle Beams*. Berlin, Heidelberg, New York: Springer. DOI: https://doi.org/10.1007/978-3-662-08581-3 (cit. on pp. 26, 31, 32).

Mockus, Jonas (1982). 'The Bayesian approach to global optimization'. In: *System Modeling and Optimization, Vol 38*. Ed. by R.F. Drenick and F. Kozin. Berlin, Heidelberg: Springer, pp. 473–481. DOI: https://doi.org/10.1007/BFb0006170 (cit. on p. 202).

Montague, Brian William (1979). *Linear Optics for Improved Chromaticity Correction*. Tech. rep. CERN-LEP-Note-165. Geneva, Switzerland: CERN, p. 16 (cit. on p. 25).

Muggli, Patric, B. E. Blue, Christopher E. Clayton, S. Deng, F. J. Decker, Mark J. Hogan, C. Huang, R. Iverson, Chandrashekhar Joshi, T. C. Katsouleas et al. (2004). 'Meter-Scale Plasma-Wakefield Accelerator Driven by a Matched Electron Beam'. In: *Physical Review Letters* 93.1, p. 014802. DOI: 10.1103/PhysRevLett.93.014802 (cit. on pp. 3, 73).

Nakajima, K., A. Enomoto, H. Kobayashi, H. Nakanishi, Y. Nishida, A. Ogata, S. Ohsawa, T. Oogoe, T. Shoji and T. Urano (1990). 'Plasma wake-field accelerator experiments at KEK'. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 292.1, pp. 12–20. DOI: https://doi.org/10.1016/0168-9002(90)91729-U (cit. on p. 73).

O'Connell, C. L., C. D. Barnes, F. J. Decker, Mark J. Hogan, R. Iverson, P. Krejcik, R. Siemann, D. R. Walz, Christopher E. Clayton, C. Huang et al. (2006). 'Plasma production via field ionization'. In: *Physical Review Special Topics - Accelerators and Beams* 9.10, p. 101301. DOI: 10.1103/PhysRevSTAB.9.101301 (cit. on p. 57).

Odom, B., D. Hanneke, B. D'urso and G. Gabrielse (2006). 'New measurement of the electron magnetic moment using a one-electron quantum cyclotron'. In: *Physical Review Letters* 97.3, p. 030801. DOI: 10.1103/PhysRevLett.97.030801 (cit. on p. 2).

Panofsky, W. K.H. and W. A. Wenzel (1956). 'Some Considerations Concerning the Transverse Deflection of Charged Particles in Radio-Frequency Fields'. In: *Review of Scientific Instruments* 27, p. 967. DOI: 10.1063/1.1715427 (cit. on p. 38).

Paramonov, Valentin V. (2013). *Field distribution analysis in deflecting structures.* Tech. rep. DESY13-018. Hamburg, Germany: DESY. arXiv: arXiv:1302.5306v1 (cit. on p. 40).

Pellegrini, C., A. Marinelli and S. Reiche (2016). 'The physics of x-ray free-electron lasers'. In: *Reviews of Modern Physics* 88.1, p. 015006. DOI: 10.1103/RevModPhys.88.015006 (cit. on p. 1).

Plath, T., Juliane Rönsch-Schulenburg, Holger Schlarb, Siegfried Schreiber, Bernd Richard Steffen and Jörg Rossbach (2013). 'Commissioning and Diagnostics Development for the New Short-Pulse Injector Laser at FLASH'. In: *Proceedings of the 2nd International Beam Instrumentation Conference - IBIC2013*. Oxford, UK: JACoW Publishing, pp. 353–356 (cit. on p. 77).

Prat, Eduard (2009). 'Spurious Dispersion Effects at FLASH'. PhD thesis. Universität Hamburg. DOI: 10.3204/DESY-THESIS-2009-026 (cit. on pp. 19, 138).

Prat, Eduard (2014). 'Symmetric single-quadrupole-magnet scan method to measure the 2D transverse beam parameters'. In: *Nuclear Instruments and Methods in Physics Research, Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 743, pp. 103–108. DOI: 10.1016/j.nima.2014.01.021 (cit. on p. 34).

Prat, Eduard, Rafael Abela, Masamitsu Aiba, Arturo Alarcon, Jürgen Alex, Yunieski Arbelo, Christopher Arrell, Vladimir Arsov, Camila Bacellar, Carl Beard et al. (2020). 'A compact and cost-effective hard X-ray free-electron laser driven by a high-brightness and low-energy electron beam'. In: *Nature Photonics* 14.12, pp. 748–754. DOI: 10.1038/s41566-020-00712-8 (cit. on p. 103).

Ramakrishnan, V. (2010). 'Unraveling the Structure of the Ribosome (Nobel Lecture)'. In: *Angewandte Chemie - International Edition* 49, pp. 4355–4380. DOI: 10.1002/anie.201001436 (cit. on p. 1).

Randall, Lisa (1999). 'Large Mass Hierarchy from a Small Extra Dimension'. In: *Physical Review Letters* 83.17, pp. 3370–3373. DOI: https://doi.org/10.1103/PhysRevLett.83.3370 (cit. on p. 2).

Rappoccio, Salvatore (2019). 'The experimental status of direct searches for exotic physics beyond the standard model at the Large Hadron Collider'. In: *Reviews in Physics* 4, p. 100027. DOI: 10.1016/j.revip.2018.100027. arXiv: 1810.10579 (cit. on p. 2).

Rasmussen, C.E. and C.K.I. Williams (2006). *Gaussian Process for Machine Learning*. Ed. by Thomas Dietterich. Cambridge, MA (USA): MIT Press (cit. on p. 202).

Reiser, Martin (2008). *Theory and Design of Charged Particle Beams*. 2nd Ed. Weinheim, Germany: WILEY-VCH Verlag GmbH & Co. KGaA. DOI: 10.1002/9783527622047 (cit. on p. 10).

Reukauff, Matthias (2022). 'An X-band LLRF Front-end for PolariX TDS'. PhD thesis (to be published). Technische Universität Hamburg, p. 103 (cit. on pp. 94, 109).

Richard, F., Jochen R. Schneider, D. Trines and A. Wagner (2001). *TESLA Technical Design Report: Executive Summary*. Tech. rep. DESY-01-011. Hamburg, Germany: DESY (cit. on p. 76).

Röckemann, Jan-Hendrik, Lucas Schaper, S. K. Barber, N. A. Bobrova, Gregor J. Boyle, S. Bulanov, N. Delbos, Klaus Floettmann, Gero Kube, W. Lauth et al. (2018). 'Direct measurement of focusing fields in active plasma lenses'. In: *Physical Review Accelerators and Beams* 21.12, p. 122801. DOI: 10.1103/physrevaccelbeams.21.122801 (cit. on p. 75).

Röhrs, Michael (2008). 'Investigation of the Phase Space Distribution of Electron Bunches at the FLASH-Linac Using a Transverse Deflecting Structure'. PhD thesis. Universität Hamburg. DOI: 10.3204/DESY-THESIS-2008-012 (cit. on pp. 39, 130).

Röhrs, Michael, Christopher Gerth, Holger Schlarb, Bernhard Schmidt and Peter Schmüser (2009). 'Time-resolved electron beam phase space tomography at a soft x-ray free-electron laser'. In: *Physical Review Special Topics: Accelerators and Beams* 12.5, p. 050704. DOI: 10.1103/PhysRevSTAB.12.050704 (cit. on p. 38).

Romano del Pozo, V., H. Bursali, Nuria Catalan Lasheras, Alexej Grudiev, S. L. Pitman, I. Syratchev and M. Volpi (2019). 'High Power Conditioning of X-band Variable Power Splitter and Phase Shifter'. In: *Proceedings of the 10th International Particle Conference - IPAC2019*. Melbourne, Australia: JACoW Publishing, pp. 2964–2967. DOI: 10.18429/JACoW-IPAC2019-WEPRB064 (cit. on pp. 109, 110, 114).

Rönsch-Schulenburg, Juliane, Bart Faatz, Katja Honkavaara, Marion Kuhlmann, Siegfried Schreiber, Rolf Treusch and Mathias Vogt (July 2017). 'Experience with Multi-Beam and Multi-Beamline FEL-Operation'. In: *Journal of Physics: Conference Series* 874, p. 12023. DOI: 10.1088/1742-6596/874/1/012023 (cit. on p. 76).

Rönsch-Schulenburg, Juliane, E. Hass, A. Kuhl, T. Plath, M. Rehders, Jörg Rossbach, G. Brenner, Christopher Gerth, U. Mavric, Holger Schlarb et al. (2013). 'Short SASE-FEL Pulses at FLASH'. In: *Proceedings of the 35th Free-Electron Laser Conference - FEL2013*. New York, NY, USA, pp. 379–382 (cit. on p. 77).

Röntgen, W. C. (1898). 'Ueber eine neue Art von Strahlen'. In: *Annalen der Physik* 300, pp. 12–17. DOI: 10.1002/andp.18983000103 (cit. on p. 1).

Rosenzweig, James B., B. Breizman, T. Katsouleas and J. J. Su (1991). 'Acceleration and focusing of electrons in two-dimensional nonlinear plasma wake fields'. In: *Physical Review A* 44.10, R6189–R6192. DOI: 10.1103/PhysRevA.44.R6189 (cit. on pp. 2, 3).

Rosenzweig, James B., D. B. Cline, B Cole, H. Figueroa, W. Gai, R. Konecny, J. Norem, P. Schoessow and J. Simpson (1988). 'Experimental Observation of Plasma Wake-Field Acceleration'. In: *Physical Review Letters* 61.1, pp. 98–101. DOI: https://doi.org/10.1103/PhysRevLett.61.98 (cit. on pp. 71, 73).

Rosenzweig, James B., A. M. Cook, A. Scott, M. C. Thompson and R. B. Yoder (2005). 'Effects of ion motion in intense beam-driven plasma wakefield accelerators'. In: *Physical Review Letters* 95.19, p. 195002. DOI: 10.1103/PhysRevLett.95.195002 (cit. on p. 192).

Rossbach, Jörg (2006). 'Results from FLASH'. In: *Proceedings of the 10th European Particle Accelerator Conference - EPAC2006*. Edinburgh, Scotland: JACoW Publishing, pp. 34–38 (cit. on p. 76).

Rossbach, Jörg (2016). 'FLASH: The First Superconducting X-Ray Free-Electron Laser'. In: *Synchrotron Light Sources and Free-Electron Lasers: Accelerator Physics, Instrumentation and Science Applications*. Ed. by Eberhard J. Jaeschke, Shaukat Khan, Jochen R. Schneider and Jerome B. Hastings. Springer, Cham, pp. 303–328. DOI: 10.1007/978-3-319-14394-1_10 (cit. on p. 76).

Rossbach, Jörg and Peter Schmüser (1994). 'Basic Course on Accelerator Physics'. In: *Proceedings of the CAS-CERN Accelerator School: 5th General Accelerator Physics Course*. Jyvaskyla, Finland: CERN, pp. 17–79. DOI: 10.5170/CERN-1994-001 (cit. on pp. 8, 17).

Rossbach, Jörg, Jochen R. Schneider and Wilfried Wurth (2019). '10 years of pioneering X-ray science at the Free-Electron Laser FLASH at DESY'. In: *Physics Reports* 808, pp. 1–74. DOI: 10.1016/j.physrep.2019.02.002 (cit. on p. 76).

Roussel, R., G. Andonian, W. Lynn, K. Sanwalka, R. Robles, C. Hansel, A. Deng, G. Lawler, James B. Rosenzweig, G. Ha et al. (2020). 'Single Shot Characterization of High Transformer Ratio Wakefields in Nonlinear Plasma Acceleration'. In: *Physical Review Letters* 124.4, p. 044802. DOI: 10.1103/PhysRevLett.124.044802. arXiv: 1910.08601 (cit. on p. 67).

Ruth, R. D., Alex Chao, P. L. Morton and P. B. Wilson (1985). 'A Plasma Wake Field Accelerator'. In: *Particle Accelerators* 17, pp. 171–189 (cit. on pp. 2, 67).

Rutherford, Ernest (1911). 'The scattering of α and β particles by matter and the structure of the atom'. In: *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science* 21.125, pp. 669–688. DOI: 10.1080/14786440508637080 (cit. on p. 1).

Saldin, E. L., E. A. Schneidmiller and M. V. Yurkov (2004). 'Longitudinal space charge-driven microbunching instability in the TESLA Test Facility linac'. In: *Nuclear Instruments and Methods in Physics Research, Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 528.1-2, pp. 355–359. DOI: 10.1016/j.nima.2004.04.067 (cit. on p. 30).

Sands, Mattew (1991). *A Beta Mismatch Parameter*. Tech. rep. SLAC-AP-85. Stanford, CA, USA: SLAC (cit. on p. 26).

Schmidt, Bernhard, Nils Maris Lockmann, Peter Schmüser and Stephan Wesch (2020). 'Benchmarking coherent radiation spectroscopy as a tool for high-resolution bunch shape reconstruction at free-electron lasers'. In: *Physical Review Accelerators and Beams* 23.6, p. 62801. DOI: 10.1103/PhysRevAccelBeams.23.062801 (cit. on p. 37).

Schmidt, Christian, V. Ayvazyan, J. Branlard, Butkowski, M. Grecki, M. Hoffmann, F. Ludwig, U. Mavrič, S. Pfeiffer, K. Przygoda et al. (2014). 'Performance of the MicroTCA.4 Based LLRF System at FLASH'. In: *Proceedings of the 5th International Particle Accelerator Conference - IPAC2014*. Dresden, Germany: JACoW Publishing, pp. 2433–2435. DOI: 10.18429/JACoW-IPAC2014-WEPME067 (cit. on p. 79).

Schneider, Jochen R. (2010). 'Photon Science at Accelerator-Based Light Sources'. In: *Reviews of Accelerator Science and Technology* 03.01, pp. 13–37. DOI: 10.1142/s1793626810000348 (cit. on p. 1).

Scholz, Matthias (2013). 'Design of the Extraction Arc for the 2nd Beam Line of the Free-Electron Laser FLASH'. PhD thesis. Universtität Hamburg. DOI: 10.3204/DESY-THESIS-2014-002 (cit. on p. 80).

Schöps, A., O. Bilani, T. Ramm, S. Tripathi, P. Vagin and M. Tischer (2014). 'Design and commissioning of the FLASH2 undulators'. In: *Proceedings of the 5th International Particle Accelerator Conference - IPAC2014*. Dresden, Germany: JACoW Publishing, pp. 2007–2009. DOI: 10.18429/JACoW-IPAC2014-WEPRO031 (cit. on p. 76).

Schreiber, Siegfried and Bart Faatz (2015). 'The free-electron laser FLASH'. In: *High Power Laser Science and Engineering* 3. DOI: 10.1017/hpl.2015.16 (cit. on pp. 4, 71, 77–79, 94).

Schreiber, Siegfried, Bart Faatz and Katja Honkavaara (2008). 'Operation of flash at 6.5nm wavelength'. In: *Proceedings of the 11th European Particle Accelerator Conference - EPAC2008*. Genoa, Italy: JACoW Publishing, pp. 133–135 (cit. on p. 76).

Schreiber, Siegfried, Sven Lederer, L. Monaco, D. Sertore and P. Michelato (2017). 'Update on the Lifetime of Cs2Te Cathodes Operated at the FLASH Facility'. In: *Proceedings of the 38th International Free Electron Laser Conference - FEL2017*. Santa Fe, NM, USA: JACoW Publishing, pp. 407–410. DOI: 10.18429/JACoW-FEL2017-WEP003 (cit. on p. 77).

Schröder, Sarah (2021). 'External injection of electron beams into plasma-wakefield accelerators'. PhD thesis. Universität Hamburg. DOI: 10.3204/PUBDB-2021-05396 (cit. on p. 84).

Schröder, Sarah, Carl Andreas Lindstrøm, Simon Bohlen, Gregor J. Boyle, Richard D'Arcy, Severin Diederichs, Mathew James Garland, Pau Gonzalez Caminal, Alex Knetsch, Vladyslav Libov et al. (2020). 'High-resolution sampling of beam-driven plasma wakefields'. In: *Nature Communications* 11.1. DOI: 10.1038/s41467-020-19811-9 (cit. on pp. 72, 82).

Schröder, Sarah, Kai Ludwig, Alexander Aschikhin, Richard D'Arcy, M. Dinter, Pau Gonzalez Caminal, Sven Karstensen, Alex Knetsch, Vladyslav Libov, Carl Andreas Lindstrøm et al. (2020). 'Tunable and precise two-bunch generation at FLASHForward'. In: *Journal of Physics: Conference Series* 1596.1, p. 7. DOI: 10.1088/1742-6596/1596/1/012002 (cit. on p. 82).

Schroeder, Carl B., E. Esarey, C. G. R. Geddes, Carlo Benedetti and Wim P. Leemans (2010). 'Physics considerations for laser-plasma linear colliders'. In: *Physical Review Special Topics - Accelerators and Beams* 13.10, pp. 1–11. DOI: 10.1103/PhysRevSTAB.13.101301 (cit. on p. 2).

Schroer, Christian, Ralf Röhlsberger, Ilya Agapov, Rainer Wanzenberg, Reinhard Brinkmann, Wim Leemans and Edgar Weckert (2019). *PETRA IV Conceptual Design Report*. Tech. rep. PUBDB-2019-03613. Hamburg, Germany: DESY. DOI: 10.3204/PUBDB-2019-03613 (cit. on p. 1).

Schulz, Sebastian (2011). 'Implementation of the Laser-Based Femtosecond Precision Synchronization System at FLASH'. PhD thesis. Universität Hamburg. DOI: 10.3204/DESY-THESIS-2014-012 (cit. on p. 79).

Schulz, Sebastian, I. Grguraš, Christopher Behrens, H. Bromberger, J. T. Costello, Marie Kristin Czwalinna, M. Felber, M. C. Hoffmann, M. Ilchen, H. Y. Liu et al. (2015). 'Femtosecond all-optical synchronization of an X-ray free-electron laser'. In: *Nature Communications* 6. DOI: 10.1038/ncomms6938 (cit. on p. 79).

Shahriari, Bobak, Kevin Swersky, Ziyu Wang, Ryan P. Adams and Nando De Freitas (2016). 'Taking the human out of the loop: A review of Bayesian optimization'. In: *Proceedings of the IEEE* 104.1, pp. 148–175. DOI: 10.1109/JPROC.2015.2494218 (cit. on p. 202).

Spence, D. J. and S. M. Hooker (2000). 'Investigation of a hydrogen plasma waveguide'. In: *Physical Review E - Statistical, Nonlinear, and Soft Matter Physics* 63.1, p. 015401. DOI: 10.1103/PhysRevE.63.015401 (cit. on p. 57).

Sprangle, P., E. Esarey and A. Ting (1990). 'Nonlinear interaction of intense laser pulses in plasmas'. In: *Physical Review A* 41.8, pp. 4463–4469. DOI: https://doi.org/10.1103/PhysRevA.41.4463 (cit. on p. 58).

Steffen, Bernd Richard, V. Arsov, G. Berden, W. A. Gillespie, S. P. Jamison, A. M. MacLeod, A. F.G. Van Der Meer, P. J. Phillips, Holger Schlarb, Bernhard Schmidt et al. (2009). 'Electro-optic time profile monitors for femtosecond electron bunches at the soft x-ray free-electron laser FLASH'. In: *Physical Review Special Topics - Accelerators and Beams* 12.3, p. 032802. DOI: 10.1103/PhysRevSTAB.12.032802 (cit. on p. 37).

Stephan, Frank, C. H. Boulware, M. Krasilnikov, J. Bähr, G. Asova, A. Donat, U. Gensch, H. J. Grabosch, M. Hänel, L. Hakobyan et al. (2010). 'Detailed characterization of electron sources yielding first demonstration of European x-ray free-electron laser beam quality'. In: *Physical Review Special Topics - Accelerators and Beams* 13.2, p. 020704. DOI: 10.1103/PhysRevSTAB.13.020704 (cit. on pp. 3, 71).

Strickland, Donna and Gerard Mourou (1985). 'Compression of amplified chirped optical pulses'. In: *Optics Communications* 55.6, pp. 447–449. DOI: 10.1016/0030-4018(85)90151-8 (cit. on p. 3).

Stulle, Frank (2004). 'A bunch compressor for small emittances and high peak currents at the VUV Free-Electron Laser'. PhD thesis. Universität Hamburg. DOI: 10.3204/DESY-THESIS-2004-041 (cit. on p. 78).

Tajima, T. and John M. Dawson (1979). 'Laser Electron Accelerator'. In: *Physical Review Letters* 43.4, pp. 267–270. DOI: https://doi.org/10.1103/PhysRevLett.43.267 (cit. on p. 2).

Thomson, Joseph John (1897). 'XL. Cathode Rays'. In: *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science* 44.269, pp. 293–316. DOI: 10.1080/14786449708621070 (cit. on p. 1).

Treyer, D. M., R. Baldinger, R. Ditter, B. Keil, W. Koprek, G. Marinkovic, M. Roggli, D. Lipka, D. Noelle and S. Vilcins (2013). 'Design and Beam Test Results of Button BPMs for the European XFEL'. In: *Proceedings of the 2nd International Beam Instrumentation Conference - IBIC2013*. Oxford, UK: JACoW Publishing, pp. 723–726 (cit. on p. 86).

Tzoufras, M., Wei Lu, F. S. Tsung, C. Huang, Warren B. Mori, Thomas Katsouleas, Jorge Vieira, R. A. Fonseca and L. O. Silva (2008). 'Beam loading in the nonlinear regime of plasma-based acceleration'. In: *Physical Review Letters* 101.14, p. 145002. DOI: 10.1103/PhysRevLett.101.145002. arXiv: 0809.0227 (cit. on pp. 3, 65–67, 174, 210).

Viti, M., Marie Kristin Czwalinna, H. Dinter, Christopher Gerth, K. Przygoda, R. Rybaniec and Holger Schlarb (2017). 'The Bunch Arrival Time Monitor At FLASH and European XFEL'. In: *Proceedings of the 16th International Conference on Accelerator and Large Experimental Control Systems - ICALEPS2017*. Barcelona, Spain: JACoW Publishing, pp. 701–705. DOI: 10.18429/JACoW-ICALEPCS2017-TUPHA125 (cit. on p. 79).

Vogel, E., C. Albrecht, Nicoleta Baboi, Christopher Behrens, T. Delfs, J. Eschke, Christopher Gerth, M. G. Hoffmann, M. Hoffmann, Markus Hüning et al. (2010). 'Test and Commissioning of the Third Harmonic RF System for FLASH'. In: *Proceedings of the 1st International Particle Accelerator Conference - IPAC2010*. Kyoto, Japan: JACoW Publishing, pp. 4281–4283 (cit. on p. 77).

Vogt, Mathias, Bart Faatz, Katja Honkavaara, Marion Kuhlmann, Juliane Rönsch-Schulenburg, Rolf Treusch and Siegfried Schreiber (2015). 'Status of the soft X-ray free electron laser FLASH'. In: *Proceedings of the 8th International Particle Accelerator Conference - IPAC2015*. Richmond, VA, USA: JACoW Publishing, pp. 2628–2630. DOI: 10.18429/JACoW-IPAC2015-TUPWA033 (cit. on p. 76).

Vogt, Mathias and Johann Zemella (2021). 'A New 2nd Bunch Compression Chicane for the FLASH2020+ Project'. In: *Proceedings of the 12th International Parcticle Accelerator Conference - IPAC2021*. Campinas, SP, Brazil: JACoW Publishing, pp. 1618–1621. DOI: 10.18429/JACoW-IPAC2021-TUPAB102 (cit. on p. 133).

Wangler, Thomas P. (2008). *RF Linear Accelerators*. 2nd ed. Weinheim, Germany: WILEY-VCH Verlag GmbH & Co. KGaA (cit. on pp. 7, 26, 38–40, 55).

Weinberg, Steven (1967). 'A Model of Leptons'. In: *Physical Review Letters* 19.21, pp. 1264–1266. DOI: https://doi.org/10.1103/PhysRevLett.19.1264 (cit. on p. 1).

Wesch, Stephan (2012). 'Echtzeitbestimmung longitudinaler Elektronenstrahlparameter mittels absoluter Intensitäts- und Spektralmessung einzelner kohärenter THz Strahlungspulse'. PhD thesis. Universität Hamburg (cit. on pp. 28, 79).

Wesch, Stephan, Bernhard Schmidt, Christopher Behrens, Hossein Delsim-Hashemi and Peter Schmüser (2011). 'A multi-channel THz and infrared spectrometer for femtosecond electron bunch diagnostics by single-shot spectroscopy of coherent radiation'. In: *Nuclear Instruments and Methods in Physics Research, Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 665, pp. 40–47. DOI: 10.1016/j.nima.2011.11.037. arXiv: 1109.0458 (cit. on p. 37).

Weyl, Hermann (1946). *The Classical Groups: Their Invariants and Representations*. 2nd ed. Princeton, NJ, EUA: Princeton University Press (cit. on p. 14).

Whittum, David H., William M. Sharp, Simon S. Yu, Martin Lampe and Glenn Joyce (1991). 'Electron-hose instability in the ion-focused regime'. In: *Physical Review Letters* 67.8, pp. 991–994. DOI: 10.1103/PhysRevLett.67.991 (cit. on pp. 3, 69).

Wiebers, Ch., M. Holz, Gero Kube, G. Priebe and H.-Ch. Schröder (2013). 'Scintillating Screen Monitors for Transverse Electron Beam Profile Diagnostics At the European Xfel'. In: *Proceedings of the 2nd International Beam Instrumentation Conference - IBIC2013*. Oxford, UK: JACoW Publishing, pp. 807–810 (cit. on pp. 86, 87).

Wiedemann, Helmut (2007). *Particle Accelerator Physics*. 3rd ed. Berlin, Heidelberg: Springer, Cham. DOI: 10.1007/978-3-319-18317-6 (cit. on pp. 9–12, 18).

Wille, Klaus (2000). *The Physics of Particle Accelerators: an Introduction*. 1st ed. Oxford University Press (cit. on pp. 15, 17, 18).

Wolski, Andzej (2014). *Beam dynamics in High Energy Particle Accelerators*. 1st ed. London, UK: Imperial College Press (cit. on pp. 16, 24).

Woodley, M. D. and Paul Emma (2000). 'Measurement and Correction of Cross-Plane Coupling in Transport Lines'. In: *Proceedings of the 20th Linear Accelerator Conference - LINAC2000*. Monterey, CA, USA: JACoW Publishing, pp. 196–198. arXiv: 0008194 [arXiv:physics] (cit. on p. 23).

Xu, X. L., J. F. Hua, Y. P. Wu, C. J. Zhang, F. Li, Y. Wan, C. H. Pai, W. Lu, W. An, P. Yu et al. (2016). 'Physics of Phase Space Matching for Staging Plasma and Traditional Accelerator Components Using Longitudinally Tailored Plasma Profiles'. In: *Physical Review Letters* 116.12, p. 124801. DOI: 10.1103/PhysRevLett.116.124801. arXiv: 1411.4386 (cit. on pp. 68, 70).

Zennaro, Riccardo, J. Alex, A. Citterio and J.-Y. Raguin (2014). 'Measurements and High-power Test of the First C-Band Accelerating Structure for SwissFEL'. In: *Proceedings of the 27th Linear Accelerator Conference - LINAC2014*. Geneva, Switzerland: JACoW Publishing, pp. 333–335 (cit. on p. 103).

List of Figures

List of Tables

List of publications

- Craievich, Paolo, M. Bopp, H. Braun, Romain Ganter, T. Kleeb, M. Pedrozzi, Eduard Prat, Sven Reiche, Riccardo Zennaro, Alexej Grudiev, Núria Catalan Lasheras, G. Mcmonagle, W. Wuensch, Barbara Marchetti, Ralph Assmann, Florian Christie, Richard D'Arcy, Ulrich Dorda, M. Foese, **Pau González Caminal**, Matthias Hoffmann, Mark Huening, Rolf Jonas, O. Krebs, Sven Lederer, Vladyslav Libov, Daniel Marx, Jens Osterhoff, F Poblotzki, Matthias Reukauff, Holger Schlarb, Siegfried Schreiber, Gerd Tews, Mathias Vogt and Antonio Wagner de Zauberrieta [2018]. 'Status of the PolariX-TDS Project'. In: *Proceedings of the 9th International Particle Accelerator Conference - IPAC2018*. Vancouver, BC, Canada: JACoW Publishing, pp. 3808–3811. doi: 10.18429/JACoW-IPAC2018-THPAL068

- Craievich, Paolo, M. Bopp, H. Braun, A. Citterio, Romain Ganter, F. Marcellini, T. Kleeb, Marco Pedrozzi, Eduard Prat, Sven Reiche, Alexej Grudiev, W. Lee Millar, G. Mcmonagle, S. Pitman, V. Romano del Pozo, K. T. Szypula, W. Wuensch, Barbara Marchetti, Ralph Assmann, Florian Christie, Richard D'Arcy, Ulrich Dorda, M. Foese, **Pau González Caminal**, Matthias Hoffmann, Markus Hüning, Rolf Jonas, O. Krebs, Sven Lederer, Vladyslav Libov, Daniel Marx, Jens Osterhoff, Matthias Reukauff, Holger Schlarb, Siegfried Schreiber, Gerd Tews, Mathias Vogt and Antonio Wagner de Zauberrieta [2019]. 'The PolariX-TDS Project: Bead-pull Measurements and High-power Test on the Prototype'. In: *Proceedings of the 39th Free-Electron Laser Conference - FEL2019*. Hamburg, Germany: JACoW Publishing, pp. 396–399. doi: 10.18429/JACoW-FEL2019-WEP036

- D'Arcy, Richard, Stephan Wesch, Alexander Aschikhin, Simon Bohlen, Christopher Behrens, Mathew James Garland, L. Goldberg, **Pau González Caminal**, Alex Knetsch, Vladyslav Libov, Alberto Martínez de la Ossa, Martin Meisel, Timon Johannes Mehrling, Pardis Niknejadi, Kristjan Põder, Jan-Hendrik Röckemann, Lucas Schaper, Bernhard Schmidt, Sarah Schröder, C. Palmer, J. P. Schwinkendorf, Bridget Sheeran, M. J. V. Streeter, Gabriele Tauscher, V. Wacker and Jens Osterhoff [2019]. 'Tunable Plasma-Based Energy Dechirper'. In: *Physical Review Letters* **122**.3, p. 034801. doi: 10.1103/physrevlett.122.034801

- D'Arcy, Richard, Alexander Aschikhin, Simon Bohlen, Gregor J. Boyle, T. Brümmer, James Chappell, Severin Diederichs, Brian Foster, Mathew James Garland, L. Goldberg, **Pau González Caminal**, Sven Karstensen, Alex Knetsch, P. Kuang, Vladyslav Libov, Kai Ludwig, Alberto Martínez de la Ossa, Frank Marutzky, Martin Meisel, Timon Johannes Mehrling, Pardis Niknejadi, Krjstian Põder, P. Pourmoussavi, M. Quast, Jan-Hendrik Röckemann, Lucas Schaper, Bernhard Schmidt, Sarah Schröder, J. P. Schwinkendorf,

Bridget Sheeran, Gabriele Tauscher, Stephan Wesch, M. Wing, Paul Viktor Winkler, Ming Zeng and Jens Osterhoff [2019]. 'FLASHForward: Plasma wakefield accelerator science for high-average-power applications'. In: *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* **377**. doi: 10.1098/rsta.2018.0392

- Lindstrøm, Carl Andreas, Richard D'Arcy, Mathew James Garland, **Pau González Caminal**, Bernhard Schmidt, Sarah Schröder, Stephan Wesch and Jens Osterhoff [2020]. 'Matching small beta functions using centroid jitter and two beam position monitors'. In: *Physical Review Accelerators and Beams* **23**.5, p. 52802. doi: 10.1103/PhysRevAccelBeams.23.052802

- Schröder, Sarah, Carl Andreas Lindstrøm, Simon Bohlen, Gregor J. Boyle, Richard D'Arcy, Severin Diederichs, Mathew James Garland, **Pau González Caminal**, Alex Knetsch, Vladyslav Libov, Pardis Niknejadi, Kristian Põder, Lucas Schaper, Bernhard Schmidt, Bridget Sheeran, Gabriele Tauscher, Stephan Wesch, Johann Zemella, Ming Zeng and Jens Osterhoff [2020]. 'High-resolution sampling of beam-driven plasma wakefields'. In: Nature Communications **11**.1. doi: 10.1038/s41467-020-19811-9

- Craievich, Paolo, Markus Bopp, H. H. Braun, Alessandro Citterio, Renato Fortunati, Romain Ganter, T. Kleeb, F. Marcellini, Marco Pedrozzi, Eduard Prat, Sven Reiche, K. Rolli, R. Sieber, Alexej Grudiev, W. L. Millar, Núria Catalan-Lasheras, G. McMonagle, S. Pitman, V. Romano del Pozo, K. T. Szypula, W. Wuensch, Barbara Marchetti, Ralph Assmann, Florian Christie, Barbara Conrad, Richard D'Arcy, M. Foese, **Pau González Caminal**, Matthias Hoffmann, Markus Huening, Rolf Jonas, O. Krebs, Sven Lederer, Daniel Marx, Jens Osterhoff, Matthias Reukauff, Holger Schlarb, Siegfried Schreiber, Gerd Tews, Mathias Vogt, Antonio Wagner de Zubiaurre and Stephan Wesch [2020]. 'Novel X-band transverse deflection structure with variable polarization'. In: *Physical Review Accelerators and Beams* **23**.11, p. 112001. doi: 10.1103/PhysRevAccelBeams.23.112001

- Lindstrøm, Carl Andreas, Mathew James Garland, Sarah Schröder, Lewis Boulton, Gregor J. Boyle, James Chappell, Richard D'Arcy, **Pau González Caminal**, Alex Knetsch, Vladyslav Libov, Gregor Loisch, Alberto Martínez de la Ossa, Pardis Niknejadi, Krjstian Põder, Lucas Schaper, Bernhard Schmidt, Bridget Sheeran, Stephan Wesch, Jonathan C. Wood and Jens Osterhoff [2021]. 'Energy-Spread Preservation and High Efficiency in a Plasma-Wakefield Accelerator'. In: *Physical Review Letters* 126.1, p. 14801. doi: 10.1103/PhysRevLett.126.014801

- Marchetti, Barbara, Alexej Grudiev, Paolo Craievich, Ralph Assmann, H. H. Braun, Núria Catalan Lasheras, Florian Christie, Richard D'Arcy, Renato Fortunati, Romain Ganter, **Pau González Caminal**, Matthias Hoffmann, Markus Huening, Sonja M. Jaster-Merz, Rolf Jonas, F. Marcellini, Daniel Marx, G. McMonagle, Jens Osterhoff, Marco Pedrozzi, Eduard Prat, Sven Reiche, Matthias Reukauff, Siegfried Schreiber, Gerd Tews, Mathias Vogt, Stephan Wesch and W. Wuensch [2021]. 'Experimental demonstration of novel beam characterization using a polarizable X-band transverse deflection structure'. In: *Scientific Reports* **11**.1, pp. 1–14. doi: 10.1038/s41598-021-82687-2

- Knetsch, Alex, Bridget Sheeran, Lewis Boulton, Pardis Niknejadi, Krjstian Põder, Lucas Schaper, Ming Zeng, Simon Bohlen, Gregor J. Boyle, Theresa Brümmer, James Chappell, Richard D'Arcy, Severin Diederichs, Brian Foster, Mathew James Garland, **Pau González Caminal**, B. Hidding, Vladyslav Libov, Carl Andreas Lindstrøm, Alberto Martínez de la Ossa, Martin Meisel, T. Parikh, Bernhard Schmidt, Sarah Schröder, Gabriele Tauscher, Stephan Wesch, Paul Viktor Winkler, Jonathan C. Wood and Jens Osterhoff [2021]. 'Stable witness-beam formation in a beam-driven plasma cathode'. In: *Physical Review Accelerators and Beams* **24**.10, p. 101302. doi: 10.1103/physrevaccelbeams.24.101302
- Lindstrøm, Carl Andreas, Judita Beinortaite, Jonas Björklund Svensson, Lewis Boulton, James Chappell, Mathew James Garland, **Pau González Caminal**, Gregor Loisch, Felipe Peña, Lucas Schaper, Bernhard Schmidt, Sarah Schröder, Stephan Wesch, Jonathan Wood, Jens Osterhoff and Richard D'Arcy [2021]. 'Progress of the FLASHForward X-2 high-beam-quality, high-efficiency plasma-accelerator experiment'. In: *The European Physical Society Conference on High Energy Physics*. doi: 10.22323/1.398.0880
- D'Arcy, Richard, James Chappell, Judita Beinortaite, Severin Diederichs, Gregor J. Boyle, Brian Foster, Mathew James Garland, **Pau González Caminal**, Carl Andreas Lindstrøm, Gregor Loisch, Siegfried Schreiber, Sarah Schröder, Rob Shalloo, Maxance Thévenet, Stephan Wesch, M. Wing and Jens Osterhoff [2022]. 'Recovery time of a plasma-wakefield accelerator'. In: *Nature* **603**.7899, pp. 58–62. doi: 10.1038/s41586-021-04348-8

Aknowledgements

This thesis is the result of a permanent exchange of ideas and information with many people, all of whom I have enjoyed working with during the past few years and to whom I would like to express my most sincere gratitute for their support. First and foremost to **Dr. Stephan Wesch**, **Dr. Carl A. Lindstrøm** and **Dr. Richard D'Arcy**, who have been most heavily involved in the work presented in this dissertation and played a fundamental role in its realisation. Furthermore, I would like to thank:

- Dr. Jens Osterhoff and Dr. Bernhard Schmidt for giving me the opportunity to do my doctoral research in their group.
- the whole FLASHForward group for these years of exciting and thriving scientific research.
- the X-2 team for their support, many fruitful discussions, successful beamtime shifts and amusing time spent together: Carl A. Lindstrøm, Sarah Schröder, Jonas Björklund Svensson, Felipe Peña, Judita Beinortaite, and Lewis Boulton.
- the people that provided the most valuable hands-on support during the experimental campaigns dedicated to the PolariX-TDS: Florian Christie, Sonja M. Jaster-Merz, Mathias Vogt, Stephan Wesch, Carl A. Lindstrøm, Johann Zemella, and Rolf Jonas.
- Severin Diederichs, Alberto Martínez de la Ossa, and Sergey Tomin for their priceless help with simulation codes.
- Bolko Beutner for introducing me to the nuances of the TDS operation and for many fruitful discussions about its data analysis, to which Lars Frölich and Josef Wilgen made helpful contributions.
- Alexej Grudiev, Paolo Craievich, Markus Hüning, and Rolf Jonas for their support with the interpretation of the PolariX-TDS beam-based-commissioning data.
- the entire PolariX-TDS collaboration involved in the development of the PolariX-TDS prototype at CERN, PSI and DESY, with which it has been a pleasure to collaborate.
- the FLA technical team for their dedication and permanent support: Kai Ludwig, Frank Marutzky, Maik Dinter, Sven Karstensen, Sandra Kottler, Vladimir Rybnikov, and Andrej Schleiermacher.
- Sven Lederer, Antonio de Zubiaurre Wagner, Nils Plambeck and Benno Zeitler for their help with the vacuum design of the new beamline.

- the DESY technical staff involved in the installation and commissioining of the PolariX-TDS and its beamline (MPC1, MIN, MVS, MCS, MSK, ...).
- the FLASH coordination and operator teams for their support with the countless measurements shifts.
- Prof. Dr. Wolfgang Hillert for accepting the request to act as my doctoral supervisor and reviewer of the thesis, and for his support and the useful discussions we had at different stages of this work.
- Prof. Dr. D. Pfannkuche, Prof. Dr. G. Moortgat-Pick and Dr. J. Osterhoff for their commitment as part of the examination board.

Finally, I would like to express my gratitude to my dear family and my bunch of weird, beloved friends, who were a massive emotional support for me during the past years and helped me to overcome all the different kinds of frustration that one encounters as a doctoral student.