Optimisation of energy and brightness transformer stages in a Plasma Wakefield Accelerator

PhD Thesis

Lewis Boulton
University of Strathclyde, Glasgow
Department of Physics

Supervisor: Prof. Bernhard Hidding
Co-supervisor: Prof. Zheng-Ming Sheng

2023

This thesis is the result of the author's original research. It has been composed by the author and has not been previously submitted for examination which has led to the award of a degree.

Signed: Lewis Boulton
Date: 24/01/23

"Any man whose errors take ten years to correct is quite a man."

- J. Robert Oppenheimer, speaking of Albert Einstein.

"I am become Death, the destroyer of cameras."

- the author, having accidentally removed the wrong filter.

Abstract

Capable of sustaining giga-volt-per-metre accelerating gradients, plasma accelerators are a promising technology that offer a path to compact machines for high-energy applications. In the case of a beam-driven plasma wakefield accelerator (PWFA), energy is transferred from a driver particle beam to an existing bunch trailing in its wake; in this way, the plasma accelerator can be seen as an *energy transformer*. Alternatively a beam can also be formed directly inside the wake—a so-called plasma cathode. The quality of this new bunch can furthermore be higher than that of the incoming drive beam, where the plasma stage now acts as a *brightness transformer*.

In the case of an energy transformer, this work focuses on an important quantity in plasma acceleration—the energy-transfer efficiency. Here, a new diagnostic was developed based on the light emitted from a beam-interacted plasma that can—in contrast to conventional dipole spectrometers—non-invasively measure the energy-transfer efficiency with longitudinal resolution on a shot-to-shot basis. After benchmarking this method with the spectrometer-based technique, local energy-transfer efficiencies of up to $(58 \pm 3)\,\%$ were measured. Furthermore, the potential of this method in diagnosing transverse instabilities was investigated, with this diagnostic seen as being key to the monitoring and optimisation of future plasma accelerators.

A PWFA-based plasma cathode stage was then also established and optimised, with the goal of demonstrating the brightness transformation of an input drive beam. Based on optically-generated density downramp injection, bunches were internally injected with high reproducibility and accelerated with $\sim$ GVm^{-1} accelerating gradients. Thorough characterisation of these bunches yielded $\mathcal{O}(10\,\mathrm{pCMeV}^{-1})$ peak spectral densities, percent-level energy spreads and normalised emittances around an order of magnitude less than the drive beam. As a result, the 3D brightness of the injected bunches (i.e. brightness in the horizontal and spectral directions) was shown to be 4.8 times higher than that of the drive beam.

Ultimately, both of the themes explored in this thesis are important if plasma accelerators are to meet the demands of future FEL light sources and linear colliders.

Contents

List of Figures

List of Tables

Acknowledgements

A PhD is a difficult undertaking at the best of times; without support, I might even say its impossible. However, I was lucky enough to find myself in a position with plenty of people who were looking out for me and were generally great to be around.

Thank you first of all to my supervisor Bernhard and the rest of the Strathclyde group for their support and interest, even though my project was not directly connected to most of their work. An enormous amount of thank yous are reserved for all the folk at FLASHForward, past and present, who have either helped me directly or have just made being part of the team a great joy. Thanks to Richard, Stephan and Jens who have all, in their individual ways, provided great examples of how to lead and manage such a technically challenging and ambitious project and in the end gave me the opportunity to contribute to the great work that FLASHForward has achieved. Thanks also to all of those who I whiled away the hours with in the BKR, FLASH2 tunnel or any of the several bars in Altona; the list of these people has grown long by now so I won't mention names, but you know who you are. This being said, I have to reserve special thanks to two people in particular who I feel particularly indebted to.

Thank you first of all to Carl, who I feel extremely lucky to have worked with over the course of my PhD. Despite being busy endlessly churning out great physics results whilst at FLASHForward, he always seemed to find the time and patience to help me with my work. I'm grateful for his trust in having me work on the results that make up Chapter 4 of this thesis which wouldn't have been in the shape it currently is without his guidance. I wish him the best of luck for the future, especially if that means that others like me will get the benefit of his wisdom.

Secondly, there is no way that this thesis would have been what it is without the help of Jon. His stoic leadership of the X-1 experiment, which came about just as I was about to be the only remaining member of the project, transformed my PhD experience for the better. It was at this time that I began to feel a greater sense of ownership of the experiment which, although challenging, was made a great deal easier by Jon's guidance and patience (especially when it came to working in the tunnel). Jon not

only made the experiment successful, but also made it fun whilst doing so—even when it seemed like the project was going to be a lot more simulation-based than we hoped (it worked in the end though!). Ultimately, it's telling that I trust Jon's advice on matters both in- and outside of physics (read korean bbq and history podcasts).

On the homefront I also have a huge amount of people to thank. Thanks to my friends who I don't see enough back home (Grant, Rob, Simon), those who I occasionally see at irish bar karaoke nights (Simms, Ruth, George, Joe Lee , Meirin) and those who I often see as characters in Baldur's Gate 3 and occasionally at the gigs of niche indie bands (Luke and Jack). This list is far from complete but you know who you are and that you're all sound.

Then, there's my family. I don't say it often—I'm certain I've never written it—but I'm immensely proud to be a member of the Boulton-Clifton clan. I'm not sure that many people can say that their family are some of the funniest people they know, but its true in my case; their down-to-earth, razor-sharp wit is something that I sincerely hope I've inherited. In particular, when it comes to my Mum and Dad, one couldn't wish for more supportive parents; they don't know much about physics, all they care about is that I'm happy and that's exactly what I need from them. As for my sister, Lydia: I've realised you're actually quite funny now and this is often a relief when those two are driving me mad, so thank you for that.

To end, I would like to give a special mention to two members of my family who I feel eternally indebted to, but unfortunately won't be able to read this acknowledgment. Grandad Les, who passed away just after I started my undergraduate degree, had a keen interest in history, to the degree that he literally ran out of books to read in Dukinfield library and had to go to Ashton instead. This interest was rivalled only by his keen interest in pints, which I feel was also something to aspire to. Grandma Eileen on the other hand left us during the course of this PhD, but the endless amount of trips she took us on, stories that she told and documentaries that she recorded provided an educational experience that I've only recently come to appreciate. Ultimately, my grandparents were two of the best teachers I've ever had and so, if this thesis is dedicated to anyone, it's to them.

Publications

The author has produced or contributed to the following publications during this study:

1. L. Boulton, C. A. Lindstrøm, J. Beinortaite, J. B. Svensson, J. M. Garland, P. G. Caminal, B. Hidding, G. Loisch, F. Peña, K. Põder, S. Schröder, S. Wesch, J. C. Wood, J. Osterhoff, and R. D'Arcy. *Longitudinally resolved measurement of energy-transfer efficiency in a plasma-wakefield accelerator*. 2022. arXiv: `2209.06690 [physics.acc-ph]`
2. J. Wood, L. Boulton, et al. "Brightness-Boosted Electron Bunches from a Plasma Accelerator". Manuscript Under Review. 2022
3. C. A. Lindstrøm, J. Beinortaite, J. Björklund Svensson, L. Boulton, J. Chappell, J. M. Garland, P. Gonzalez, G. Loisch, F. Peña, L. Schaper, B. Schmidt, S. Schröder, S. Wesch, J. Wood, J. Osterhoff, and R. D'Arcy. "Progress of the FLASHForward X-2 high-beam-quality, high-efficiency plasma-accelerator experiment". In: *PoS* EPS-HEP2021 (2022), p. 880. DOI: `10.22323/1.398.0880`
4. A. F. Habib, T. Heinemann, G. G. Manahan, L. Rutherford, D. Ullmann, P. Scherkl, A. Knetsch, A. Sutherland, A. Beaton, D. Campbell, L. Boulton, A. Nutter, O. S. Karger, M. D. Litos, B. D. O'Shea, G. Andonian, D. L. Bruhwiler, J. R. Cary, M. J. Hogan, V. Yakimenko, J. B. Rosenzweig, and B. Hidding. *Ultrahigh brightness beams from plasma photoguns*. 2021. arXiv: `2111.01502 [physics.acc-ph]`
5. A. Knetsch, B. Sheeran, L. Boulton, P. Niknejadi, K. Põder, L. Schaper, M. Zeng, S. Bohlen, G. Boyle, T. Brümmer, J. Chappell, R. D'Arcy, S. Diederichs, B. Foster, M. J. Garland, P. Gonzalez Caminal, B. Hidding, V. Libov, C. A. Lindstrøm, A. Martinez de la Ossa, M. Meisel, T. Parikh, B. Schmidt, S. Schröder, G. Tauscher, S. Wesch, P. Winkler, J. C. Wood, and J. Osterhoff. "Stable witness-beam formation in a beam-driven plasma cathode". In: *Phys. Rev. Accel. Beams* 24 (10 Oct. 2021), p. 101302. DOI: `10.1103/PhysRevAccelBeams.24.101302`

6. C. A. Lindstrøm, J. M. Garland, S. Schröder, L. Boulton, G. Boyle, J. Chappell, R. D'Arcy, P. Gonzalez, A. Knetsch, V. Libov, G. Loisch, A. Martinez de la Ossa, P. Niknejadi, K. Põder, L. Schaper, B. Schmidt, B. Sheeran, S. Wesch, J. Wood, and J. Osterhoff. "Energy-Spread Preservation and High Efficiency in a Plasma-Wakefield Accelerator". In: *Phys. Rev. Lett.* 126 (1 Jan. 2021), p. 014801. DOI: `10.1103/PhysRevLett.126.014801`
7. B. Hidding, A. Beaton, L. Boulton, S. Corde, A. Doepp, A. F. Habib, T. Heinemann, A. Irman, S. Karsch, G. Kirwan, A. Knetsch, G. Manahan, A. Martinez de la Ossa, A. Nutter, P. Scherkl, U. Schramm, and D. Ullmann. "Fundamentals and Applications of Hybrid LWFA-PWFA". in: *Applied Sciences* 9 (June 2019), p. 2626. DOI: `10.3390/app9132626`

Chapter 1

Introduction

Originally invented as a means of probing fundamental physics [1–4], particle accelerators now perform crucial roles in a wide range of sciences that impact society on a large scale [5]. Particle accelerator technologies can have direct effects on peoples lives, ranging from medical applications such as cancer treatment [6] to industrial applications such as the doping of semiconductors for electronics [7]. However, the role of accelerators as 'engines of discovery' [8] by empowering science at all levels yields benefits that, although not directly tangible for most people, can over time have profound impacts on society. For example, it is worth remembering that CERN [9]—the world's largest particle accelerator lab—is widely credited as the birthplace of the World Wide Web, created in response to the need of scientists to automatically share information with institutes around the globe [10].

Although particle accelerators now find applications in disciplines outside of fundamental science, the appetite of researchers to study the universe at the smallest scales is no less strong. Activities at the cutting edge of many of these fields typically require high energies. For particle physicists this translates to colliding particles together with a large enough centre-of-mass energy to recreate the conditions of the early universe [11]. For photon scientists, resolving the smallest structures relies on high energy particle beams emitting synchrotron radiation at a short enough wavelength [12].

The radio frequency (RF) cavities that act as the work-horse of modern accelerators have an accelerating gradient ($\sim$ 10–100 $\mathrm{MVm^{-1}}$) that is limited by the fields that they can sustain before breakdown of the material at their walls [13]. Increasing the final energy of a particle beam is then simply a matter of propagating the bunches through a series of RF cavities in a so-called *linear accelerator* (linac)[14], but can rapidly lead to long and costly machines [15]. In principle, the energy of the accelerated particles can be increased indefinitely in *circular accelerators*, where the beams will perform multiple orbits of the same machine, gaining energy in each pass [16]. However, the maximum energy is then limited by the strength of the magnets required to steer and focus the beam [17]. Moreover, such machines—known as synchrotrons—are further limited in energy gain due to energy losses via synchrotron radiation, especially for light particles (i.e. electrons and positrons) [18]. Future particle accelerator facilities that need the very highest energies and beam quality are therefore increasingly based on linacs, where the options for reaching the highest energies are either to build larger machines—requiring potentially prohibitive amounts of money and resources—or to address the fundamental limit in accelerating gradient that current RF cavity technology exhibits.

A number of novel accelerator concepts have therefore been proposed in order to overcome these limits. One such scheme involves the use of an acceleration medium that is already in a broken-down state—*plasma*. Often described as the fourth state of matter, plasmas are overall neutral but are able to sustain extremely strong electric fields ($> \mathrm{GVm}^{-1}$) in regions of localised charge separation. Intense particle beams or lasers can be used to generate such structures, known as wakes, which can in turn act as microscopic accelerating cavities. A charged particle bunch placed in the correct part of the wakefields therein can then rapidly gain energy in a fraction of the distance that would be required with conventional machines. Such plasma accelerators [19] are therefore viewed as a promising technology for compact high-energy particle sources that can exist in small-scale laboratories, but also as a solution for future particle colliders and light sources where the footprint and cost of a conventional machine would be unfeasible.

This thesis focuses on plasma accelerators driven with intense charged particle beams, simply known as *plasma wakefield acceleration* (PWFA) [20]. In such accelerators, the mean energy of existing particle beams can be rapidly increased by the transfer of energy from a drive beam. As such, PWFA modules are often termed *energy tranformers*. Amongst other considerations, the energy efficiency of such an acceleration stage is of paramount importance—both from a cost and climate point-of-view.

Whereas the energy of existing particle bunches can be boosted, there is also the opportunity to generate completely new particle bunches within the wake. Such beams can have much higher quality i.e. brightness than the driver; in this case, a PWFA can be viewed as a *brightness transformer*. This is of key importance for future colliders and light sources, both of which require high-quality beams to produce large amounts of collisions or bright pulses of light, respectively.

The work presented here therefore concentrates on the optimisation of the energy and brightness transformer aspects of beam-driven plasma wakefield acceleration. Chapter 2 lays out some of the fundamental concepts required for understanding the processes involved in plasma wakefield acceleration. Chapter 3 introduces the FLASHForward facility at DESY, Hamburg, where the entirety of the experimental studies presented in this thesis were performed. Chapter 4 details a method of measuring wake-to-trailing bunch energy transfer efficiencies that are longitudinally resolved along the length of the accelerator stage, making local diagnosis of the acceleration process now possible in experiment. Chapter 5 presents results of an optically-triggered density downramp (plasma torch) injection experiment, where novel techniques for the optimisation of injection are explored, resulting in a stable generation of witness bunches whose key qualities are then characterised. Extensive simulation studies in support of these results are also included in this chapter. Chapter 6 then concludes the thesis by providing an overarching outlook for all of the work completed in this project.

Chapter 2

Fundamental Concepts

In this chapter, the fundamental concepts required for an understanding of the studies presented in this thesis are given. As a large amount of these involve the interaction of particles with electromagnetic fields, the chapter begins with a brief review of electrodynamics. The particle beam dynamics that underpin particle accelerators of any type are then introduced. A short overview of key plasma physics concepts are then given, including how they can be generated and how they evolve. The formation of plasma wakes is then examined, followed by methods of trailing bunch injection and details of the acceleration process that follows. Finally, some of the computational tools that can be used to model parts of the physics outlined in this chapter are discussed.

2.1 Electrodynamics

Maxwell's equations form the basis of electrodynamics [21], which take the form

$$\boldsymbol{\nabla} \cdot \mathbf{E} = \frac{\rho}{\epsilon_0} \tag{2.1}$$

$$\boldsymbol{\nabla} \cdot \mathbf{B} = 0 \tag{2.2}$$

$$\boldsymbol{\nabla} \times \mathbf{E} + \frac{\partial \mathbf{B}}{\partial t} = 0 \tag{2.3}$$

$$\boldsymbol{\nabla} \times \mathbf{B} = \mu_0 \mathbf{j} + \mu_0 \epsilon_0 \frac{\partial \mathbf{E}}{\partial t}, \tag{2.4}$$

and essentially describe the relation between the electric field $\mathbf{E}$, magnetic field $\mathbf{B}$, charge density ρ and current density $\mathbf{j}$. The constants ϵ_0 and μ_0 are the vacuum permittivity and permeability respectively. By substituting Gauss' law for electricity (Equation 2.1) and Ampére's law (Equation 2.4) into the curl of Faraday's law (Equation 2.3) we obtain the inhomegeneous wave equation for the electric field

$$\frac{1}{c^2}\frac{\partial^2 \mathbf{E}}{\partial t^2} - \nabla^2 \mathbf{E} = -\left(\frac{1}{\epsilon_0}\nabla\rho + \mu_0 \frac{\partial \mathbf{j}}{\partial t}\right), \tag{2.5}$$

where the speed of light in a vacuum is

$$c = \frac{1}{\sqrt{\epsilon_0 \mu_0}} = 2.997\,924\,58 \times 10^8\,\mathrm{m\,s^{-1}}. \tag{2.6}$$

Note that the source terms ρ and $\mathbf{j}$ are related by the continuity equation,

$$\frac{\partial \rho}{\partial t} + \boldsymbol{\nabla} \cdot \mathbf{j} = 0. \tag{2.7}$$

The force acting on a point charge q with mass m and velocity $\mathbf{v}$ due to an electromagnetic field is then determined by the Lorentz equation [22],

$$\mathbf{F} = \frac{d\mathbf{p}}{dt} = \frac{d\gamma m \mathbf{v}}{dt} = q(\mathbf{E} + \mathbf{v} \times \mathbf{B}), \tag{2.8}$$

where $\gamma = \frac{1}{\sqrt{1-v^2/c^2}}$ is the relativistic gamma factor. For most of this thesis, the dynamics of electrons in electromagnetic fields will be considered, with fundamental charge $-e$ and mass m_e.

2.2 Particle beam dynamics

In this section a brief overview of the physics of particle beams is discussed, beginning with the motion of single particles which is then generalised to that of particle-ensembles.

2.2.1 Single-particle dynamics

Before discussing how a particle's motion can be manipulated in a particle accelerator, a coordinate system and corresponding dynamical variables must first be introduced. A curvilinear coordinate system is most commonly adopted in particle accelerator physics (see Figure 2.1), describing the location of a particle in 6D phase space relative to that of a reference particle with longitudinal position s_0 and momentum p_0 travelling along the design trajectory s of the machine. In the curvlinear coordinate system, the particle coordinates are often expressed as the vector $\mathbf{X} = (x, x', y, y', \xi, \delta)$, where for a particle with momenta (p_x, p_y, p_s) in the $(\hat{x}, \hat{y}, \hat{s})$ directions, x and y are the particle's transverse offsets from the design trajectory, $x' = \frac{dx}{ds} \approx \frac{p_x}{p_s}$ and $y' = \frac{dy}{ds} \approx \frac{p_y}{p_s}$ are the angles between the particle's trajectory and the design orbit and $\xi = s - s_0$ and $\delta = \frac{p_s - p_0}{p_0}$ are the particle's longitudinal position and momentum relative to the reference particle. Grouping these variables together, the xx' and yy'-planes constitute the phase spaces formed by the particle positions and angles in the transverse directions—also known as *trace space*—whilst the $\xi\delta$-plane represents the longitudinal phase space.

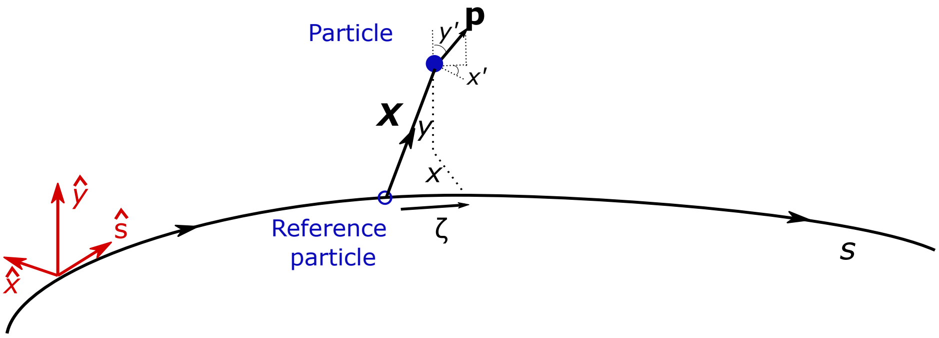

Figure 2.1: Curvilinear Frenet-Serret coordinate system. The position of the particle (blue filled dot) is described relative to the position of the design particle (blue circle) travelling along a design trajectory s. The momentum $\mathbf{p}$ of the particle is also given in terms of the momentum along s and its angles with respect to s [23].

Longitudinal motion

To increase a charged particles momentum in the direction of travel—in other words, to increase its energy—it is well known that only the electric field part of Equation 2.8 can be used; the magnetic field term only ever acts perpendicularly to the direction of motion, and so cannot do work. The increase in energy ΔE of a charged particle by a longitudinal electric field E_z over an acceleration length L is found by integrating along the direction of travel, such that

$$\Delta E = q \int_0^L E_z(s)\, ds \tag{2.9}$$

Note that if $E_z(s)$ is different to what would be experienced by the design particle, this can lead to a non-zero value of δ. Variations in the ξ coordinate can in principle occur if the velocity of the particle is not relativistic; most of the time however, particles in an accelerator travel close to the speed of light in the longitudinal direction. The ξ coordinate in this case can however be manipulated by altering the length of the particle's trajectory, necessarily requiring alteration of the particle's transverse orbit.

Transverse motion

Whereas boosting the energy of charged particles in an accelerator can only be achieved through the use of an electric field, their transverse motion is typically controlled via the application of a magnetic field. Consider the multipole expansion of the vertical component of the magnetic field B_y along the x-axis, about $x = 0$:

$$\begin{aligned} \frac{e}{p}B_y(x) &= \left.\frac{eB_y}{p}\right|_{x=0} + \left.\frac{e}{p}\frac{\partial B_y}{\partial x}\right|_{x=0} \cdot x + \cdots \\ &= \quad \frac{1}{\rho} \quad + \quad k_1 x \quad + \cdots , \end{aligned} \tag{2.10}$$

where B_y has been normalised by the beam rigidity $\frac{p}{e}$—a measure of the amount of deflection experienced by a particle of momentum p and charge e in a magnetic field. Note that only the first two terms of the expansion have been explicitly shown here: the constant dipole term and the quadrupole term that scales linearly with the an offset in x. These constitute the elementary building blocks of an accelerator's magnetic lattice: dipoles that can bend the particle trajectory in one direction, and quadrupoles that can focus particles towards the axis. Normalising by the beam rigidity allows normalised measures of the strengths of these terms to be defined, with the dipole strength quantified by the radius of its bending trajectory ρ and the quadrupole strength given in terms of its normalised focusing gradient k_1. Note that considering the same magnetic field but with $x \to y$, the focusing strength of the quadrupole term will be the same but with the opposite sign—implying that particles cannot be focused in both directions with just one quadrupole element.

The equation of motion for an on-momentum particle ($\delta = 0$) in the x direction is described by the so-called *Hill's equation* [24],

$$x'' = -k(s)x, \tag{2.11}$$

where considering only dipolar and quadrupolar fields, $k(s) = \frac{1}{\rho^2(s)} - k_1(s)$. Equation

2.11 is that of a simple harmonic oscillator with a solution given by

$$x_\beta(s) = \sqrt{\epsilon\beta(s)}\cos\left(\sqrt{k(s)}s + \phi_0\right) \tag{2.12}$$

where ϕ_0 is defined as an initial phase and $\beta(s)$ and ϵ are the varying and constant parts of the amplitude respectively. Known as the *Courant-Snyder invariant* or *single-particle emittance*, ϵ is, for a particle acted on only by linear forces, conserved throughout the accelerator [25]. The beta-function $\beta(s)$ on the other hand varies according to the differential equation

$$\frac{1}{2}\beta(s)\beta''(s) - \frac{1}{4}\beta'(s)^2 + k(s)\beta(s)^2 = 1, \tag{2.13}$$

found by substitution of the solution in Equation 2.12 into 2.11. In combination with the β-function, two extra parameters can be defined as

$$\alpha(s) = -\frac{1}{2}\frac{d\beta(s)}{ds} \tag{2.14}$$

and

$$\gamma(s) = \frac{1+\alpha(s)^2}{\beta(s)} \tag{2.15}$$

which complete the set of functions known as the *Courant-Snyder* or *Twiss parameters*. Together with the single particle emittance, these functions completely specify the motion of a single particle through phase space, following the relation

$$\epsilon(x, x') = \gamma_x x^2 + 2\alpha_x x x' + \beta_x {x'}^2 \tag{2.16}$$

which is the equation of an ellipse of area $\pi\epsilon$ in $x - x'$ trace space.

Off-momentum particles

The motion of off-momentum particles in which $\delta \neq 0$ can also be considered [23]. In this scenario Equation 2.11 is, to first order in x, instead given by

$$x'' + \left(\frac{1-\delta}{\rho^2(1+\delta)} - \frac{k_1(s)}{(1+\delta)}\right)x = \frac{\delta}{\rho(1+\delta)}, \tag{2.17}$$

which is a linearised, inhomegeneous equation with a solution that is given by that of the homogenous equation x_β plus the particular solution $D(s)\delta$, i.e.

$$x(s) = x_\beta(s) + D(s)\delta. \tag{2.18}$$

Simply put, this is a statement that the orbits of particles with different energies will differ by an amount characterised by the dispersion function $D(s)$. In some cases it is necessary to minimise the dispersion in a beamline; on the other hand, dispersion also forms the basis of e.g. spectral measurements. Note also now that x_β satisfies the equation

$$x_\beta'' + (k(s) + \Delta k(s))x_\beta = 0, \tag{2.19}$$

where the term $\Delta k(s)$ includes the δ-dependent terms that now appear. This introduces an additional energy-dependent effect on the dynamics of particles in an accelerator—*chromaticity*. In a similar way to the dispersion, this is simply a consequence of the fact that particles of differing energies will be deflected by different amounts in a magnetic field. In the case of chromaticity, this means that particles will be focused in a way that is dependent on their energy and is described by the Δk term, which can be thought of as a *chromatic perturbation* to the focusing strength of a magnetic element e.g. a quadrupole.

Matrix formalism

For everything discussed so far that has remained linear in the phase space coordinates $(x, x', y, y', \xi, \delta)$, the effects of the elements in an accelerator lattice can be described via linear operators i.e. matrices [26]. In this case, the propagation of a particle represented by an initial phase space vector $\mathbf{X_0}$ through an accelerator lattice consisting of n elements can be written as

$$\begin{aligned} \mathbf{X} &= \mathbf{M_n} \cdots \mathbf{M_2 M_1 X_0} && (2.20) \\ &= \mathbf{R_{ij} X_0} && (2.21) \end{aligned}$$

where $\mathbf{M_i}$ is the *transfer matrix* representing the i^{th} consecutive lattice element and $\mathbf{R_{ij}}$ is a 6×6 matrix describing the linear transformation $\mathbf{X_0} \to \mathbf{X}$.

Often, only transformations involving subsets of the coordinates discussed previously are considered. An illustrative example of a transformation that acts on the transverse phase spaces e.g. the vector (x, x', y, y') is given by the matrix representation of a quadrupole,

$$\mathbf{M_Q} = \begin{pmatrix} \cos\left(\sqrt{k_1}l\right) & \frac{1}{\sqrt{k_1}}\sin\left(\sqrt{k_1}l\right) & 0 & 0 \\ -\sqrt{k_1}\sin\left(\sqrt{k_1}l\right) & \cos\left(\sqrt{k_1}l\right) & 0 & 0 \\ 0 & 0 & \cosh\left(\sqrt{k_1}l\right) & \frac{1}{\sqrt{k_1}}\sinh\left(\sqrt{k_1}l\right) \\ 0 & 0 & \sqrt{k_1}\sinh\left(\sqrt{k_1}l\right) & \cosh\left(\sqrt{k_1}l\right) \end{pmatrix} \tag{2.22}$$

where l is the length of the quadrupole which in this case focuses in the x direction and defocuses in the y direction. In the case that there are no fields present i.e. $k \to 0$, Equation 2.22 becomes

$$\mathbf{M_D} = \begin{pmatrix} 1 & l & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & l \\ 0 & 0 & 0 & 1 \end{pmatrix} \tag{2.23}$$

which is the transfer matrix for a field-free vacuum of length l, commonly known as a *drift*. Naturally, the drift transfer matrices that act on the xx' and yy' trace spaces are equivalent.

In the case of the longitudinal phase space $\xi\delta$, important elements of the $\mathbf{R}$ matrix can be identified by inspecting their effects on the coordinates. The classic example is the element R_{56}, where it can be shown, to first-order, that $\xi \propto R_{56}\delta$. In other words, by inducing a non-zero value of R_{56} in the beamline—also known as the *longitudinal dispersion*—the longitudinal positions of the particles can be manipulated based on their relative momentum δ. This forms the basis of vital accelerator components such as bunch compressors, where the introduction of dispersion (e.g. in a chicane of dipoles) can shorten (or lengthen) a bunch by exploiting correlations in $\xi\delta$.

Note that this treatment only considers transformations that are linear in the dynamical variables. Describing the interaction of particles to higher-order—for example modelling their passage through non-linear fields or considering higher-order dispersive/chromatic effects—requires the inclusion of higher-order matrices. Often, for a full description of particle motion, particle tracking codes are used (see Section 2.5).

2.2.2 Beams of many particles

Rather than the motion of individual particles throughout an accelerator, we are more often than not interested in the orbits of particle ensembles i.e. beams. Consider a smooth, normalised particle distribution in 6D phase space $f(\mathbf{X})$ which we can measure in terms of its moments

$$\langle x_i^n \rangle = \int f(\mathbf{X}) x_i^n dx_i \tag{2.24}$$

and correlations

$$\langle x_i^n x_j^m \rangle = \int f(\mathbf{X}) x_i^n x_j^m dx_i dx_j, \tag{2.25}$$

where x_i and x_j can be any of the six dynamical variables in phase space.

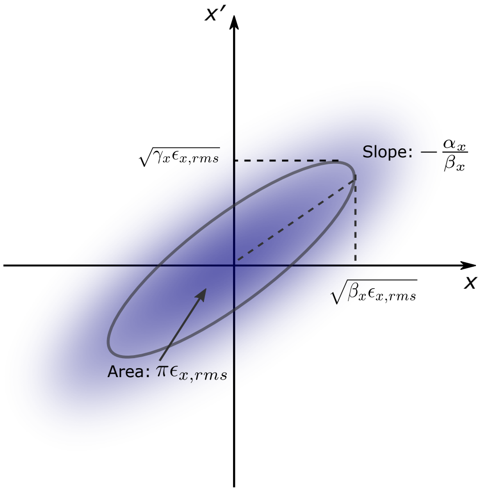

Figure 2.2: Phase-space distribution of a gaussian beam. Surrounding 68 % of the particles in the beam (blue), the phase space ellipse representing the rms trace-space emittance for this beam is shown (black ellipse) and the parameters related to the e.g. beam size and divergence are labelled.

To make the connection between the single- and multi-particle cases, consider a beam with a gaussian distribution in $x - x'$ phase space, depicted in Figure 2.2. The rms *trace-space emittance* of the beam in x is defined as [27]

$$\epsilon_{x,tr} = \sqrt{\langle x^2 \rangle \langle x'^2 \rangle - \langle x x' \rangle^2}. \tag{2.26}$$

This emittance, also known as the *geometric emittance*, can be thought of as that of a single particle who's phase space ellipse—defined by the relation in Equation 2.16—surrounds 68 % of the particles in the distribution (see black ellipse in Figure 2.2). As such, this emittance follows an analagous expression to that in Equation 2.16, only now with $x^2 \rightarrow \langle x^2 \rangle$, $x'^2 \rightarrow \langle x'^2 \rangle$, $xx' \rightarrow \langle xx' \rangle$ and the single particle emittance replaced by the trace-space emittance. The Courant-Snyder parameters can then also be defined in terms of the moments and correlations of the distribution, where

$$\beta_x = \frac{\langle x^2 \rangle}{\epsilon_{x,tr}} \, , \, \gamma_x = \frac{\langle x'^2 \rangle}{\epsilon_{x,tr}} \, , \, \alpha_x = -\frac{\langle xx' \rangle}{\epsilon_{x,tr}}. \tag{2.27}$$

Intuitively, these parameters can be thought of as measures of the rms beam width $\sigma_x = \sqrt{\langle x^2 \rangle}$, rms divergence $\sigma'_x = \sqrt{\langle x'^2 \rangle}$ and the xx' correlation of the distribution, with the emittance factored out.

Note that this only considers the moments and correlations of the distribution in the transverse planes; in general, we can also consider these in all six dimensions of the beam's phase space. For example, two important quantities in longitudinal phase space are the rms bunch length $\sigma_\xi = \sqrt{\langle \xi^2 \rangle}$ and rms energy spread $\sigma_\delta = \sqrt{\langle \delta^2 \rangle}$. Indeed, a finite energy spread can also contribute to the e.g. rms beam size [24]

$$\sigma_x = \sqrt{\epsilon_{x,tr}\beta_x + \eta_x^2 {\sigma_\delta}^2}, \tag{2.28}$$

where $\eta_x = \frac{\langle x\xi \rangle}{\langle \xi^2 \rangle}$ is the *beam dispersion*.

The correlation $\langle \xi\delta \rangle$ is also an important parameter known as the *chirp* which, in the presence of a non-zero R_{56}, can result in bunch compression (or lengthening) in the accelerator. This additionally introduces the concept of *correlated* and *uncorrelated* parameters. In the presence of a chirp, the energy spread σ_δ can be dominated by the correlation term. It is therefore often useful to consider the energy spread of the bunch in the absence of this correlation. In practice, this is equivalent to separating the bunch into longitudinal slices and calculating the energy spreads of each segment individually. The energy spread of these slices is therefore appropriately named the *slice energy spread*, as opposed to the *projected energy spread*, which takes in to full account the $\xi\delta$ correlation.

Similar sliced quantities can also be defined for other beam parameters. Most commonly, the transverse emittance of individual slices a.k.a. the *slice emittance* is considered alongside the *projected emittance*. Both of these parameters are important for understanding sources of emittance growth in particle accelerators.

Beam quality

Depending on the application, there will be certain properties that the accelerated beams are required to have. For demanding applications such as high-energy physics, although the charge and energy of a beam are important parameters, it is also vital that the beam can be transported through the lattice and focused to a small size at the interaction point, requiring necessarily low energy spreads and emittances (see Equation 2.28).

Note that the trace-space emittance defined by Equation 2.26 in fact decreases during acceleration due to *adiabatic damping*, where the angles x' and y' decrease because of an increase in the longitudinal momentum p_z. It therefore becomes useful to separate the emittance change due to acceleration from that due to other sources by defining the *normalised trace space emittance* [27]

$$\epsilon_{x,tr,n} = \gamma \epsilon_{x,tr} = \frac{\langle p_z \rangle}{m_e c} \sqrt{\langle x^2 \rangle \langle x'^2 \rangle - \langle x x' \rangle^2}, \tag{2.29}$$

where the change in momentum of the beam has been factored out. A related and frequently used emittance definition is the *normalised phase space emittance*

$$\epsilon_{x,n} = \frac{1}{m_e c} \sqrt{\langle x^2 \rangle \langle {p_x}^2 \rangle - \langle x p_x \rangle^2}. \tag{2.30}$$

Note that although Equations 2.29 and 2.30 produce the same result when the beam is at a waist i.e $\langle x x' \rangle = 0$, they differ fundamentally in highly divergent beams that also have a significant energy spread (see Ref. [27]). Whereas the normalised trace-space emittance is constant in a drift section, the normalised phase-space emittance can change due to the emergence of correlations between the transverse position and longitudinal momentum. Only when the divergence—or equally energy spread—of the beam is compensated does the normalised trace-space emittance change to once again be equal to the normalised phase-space emittance.

In combination with the high charge and low (normalised) emittances and energy spread, it is often desirable to have short bunches—especially for light source applications where the bunch length needs to be of-order the radiation wavelength. To summarise all of these requirements the *normalised six-dimensional brightness* of a bunch with total charge Q can be defined as [28]

$$B_{6\mathrm{D},n} = \frac{Q}{\epsilon_{x,n} \epsilon_{y,n} \epsilon_{z,n}} \tag{2.31}$$

where the normalised transverse emittances $\epsilon_{x,n}$ and $\epsilon_{y,n}$ represent the focusability of the beams and the normalised longitudinal emittance $\epsilon_{z,n} = \sqrt{\langle z^2 \rangle \langle \delta^2 \rangle - \langle z\delta \rangle^2}$ reflects how short a bunch can be made for a given energy spread. Intuitively, this can be thought of as a measure of the density of particles in 6D phase space, which for the most demanding applications is required to be maximised. It is however common to only consider the brightness for a subset of the dimensions, for instance the *normalised 5D brightness*,

$$B_{5\mathrm{D},n} \propto \frac{I_{pk}}{\epsilon_{x,n}\epsilon_{y,n}}, \tag{2.32}$$

which, via the peak current I_{pk} (equal to $\frac{Q}{\sqrt{2\pi}\sigma_{\xi}}$ in the case of a gaussian bunch), retains the dependence on the longitudinal coordinate without considering the energy spectrum.

A quantity that is related to the brightness but is more relevant for particle colliders is the *luminosity*

$$\mathcal{L} \propto \frac{N^2 f \gamma}{\sigma_x \sigma_y}, \tag{2.33}$$

which summarises the need to collide beams of high-particle number N with high-repetition rate f at an interaction point where the beam sizes $\sigma_{x,y}$ are small. For the most cutting-edge particle physics experiments, the desire is also to maximise the energy of the colliding beams i.e. maximise their mean relativistic factor γ. Equation 2.33 can be recast in a different form where

$$\mathcal{L} \propto \frac{P_{wall}}{\sqrt{\beta_x \beta_y}} \frac{\eta N}{\sqrt{\epsilon_{x,n}\epsilon_{y,n}}}, \tag{2.34}$$

emphasising the need for low normalised emittances and small beta functions. Equation 2.34 also highlights the dependency on the input power, which to be maximised for a given wall-plug-power P_{wall} necessarily requires high energy-efficiency η.

Sources of beam quality degradation

As the beam traverses the lattice, it can experience many different forms of quality degradation. As a simple example, the bunch can experience imperfect *charge coupling* i.e. where charge is lost during transport through the lattice. Most of the time when discussing beam quality decrease we are however referring to an increase in the value of two variables: the energy spread and the transverse emittances.

An increase in energy spread of a particle bunch is in general due to particles within the bunch experiencing different accelerations. This can arise during the acceleration process if the particles are placed at different phases in a non-uniform accelerating field. Energy spread increase can however also arise from e.g. the emission of radiation. Nevertheless, it is sometimes the case that growth of the correlated energy spread can be removed in a process known as *dechirping*.

The normalised emittance on the other hand is a quantity that in most circumstances can at best be preserved. *Liouville's theorem* implies that, in the absence of radiation effects and mutual particle interactions, the 6D phase space volume of a beam is conserved. The shape that this volume occupies however can change, resulting in an altered value of the emittance. Below, examples of emittance growth sources are listed, with emphasis on those that are particularly relevant to plasma accelerators [29]:

1. **Interaction with non-linear fields:** Whereas linear transverse fields are capable of preserving the beam emittance, non-linear focusing fields can result in beam filamentation and thus an emittance growth [30].

2. **Beam mismatch and misalignment:** If a beam is in a uniform focusing channel, the beam envelope does not oscillate and is considered *matched*. In the case that this is not true, betatron oscillations of the beam particles can have different frequencies if the focusing force is not the same for each one, or if they have a spread of energies. This leads to a differential rotation of their ellipses, resulting in a larger area occupied by the beam in phase space; an emittance degrading effect known as *betatron decoherence* [31].

   Closely related to mismatching is misalignment; whilst mismatching is concerned with the second moment of the transverse distribution, misalignment refers to the first moment. Misalignment can also lead to emittance growth via a decoherence effect [32, 33]. However, misalignments can also lead to the growth of instabilities—for example the hosing [34] and beam-breakup [35–37] instabilities—which can increase the emittance of the bunch and essentially destroy the beam.

3. **Chromaticity:** In the case that beams have a large divergence and/or energy spread—often the case in plasma accelerators—a correlation between the transverse coordinates and their momenta can emerge in a drift space, resulting in a correlated increase of ϵ_n [27]. This again can be seen as a decoherence effect; particles of different energies rotate at different rates during the drift, resulting in a growth in the projected phase space emittance that increases with the drift length and energy spread [38]. Both definitions of emittance in Equations 2.29 and 2.30 are therefore relevant; the normalised trace-space emittance grows to the value of the increased normalised phase-space emittance when the beam is refocused, for example when the accelerated beam is captured for transport to applications or subsequent plasma stages. Equally, the phase-space emittance can only be accurately measured after a strong focusing element, where the two emittances are equal.

4. **Scattering:** Collisions of the beam with atoms and ions in an accelerator leads to its particles being deflected at angles. Considering only elastic scattering from a gas of particles, this implies a growth in the divergence at a rate given by [39]

$$\frac{d\langle x'^2\rangle}{ds} = \frac{n_0 e^2}{m_e c^2 \epsilon_0}\frac{r_e}{\gamma^2}\sqrt{Z_i^2 \ln\left(\frac{\lambda_D}{R_a}\right) + 1.78Z(Z+1)\ln\left(\frac{287}{\sqrt{Z}}\right)} \tag{2.35}$$

where n_0 is the atomic density of the gas, Z is the atomic number, Z_i is the ion charge state, λ_D is the Debye length and $R_a \approx 1 \times 10^{-10}\,\mathrm{m}$ is the atomic radius. Whereas the first term in Equation 2.35 represents the contribution from collisions with ions, the second term arises from scattering with neutral atoms. Owing to this increase in divergence, the emittance grows according to

$$\frac{d\epsilon_n}{ds} = \frac{\beta_x}{2}\frac{d\langle x'^2\rangle}{ds}\gamma. \tag{2.36}$$

Due to the dependence on the β-function and the overall dependence on $\frac{1}{\gamma}$, the emittance growth due to scattering is therefore lower for narrow width, high energy beams. Conventional accelerator beamlines are normally kept at high-vacuum and so emittance growth due to gas scattering is often not a concern. However, the very nature of plasma accelerators brings particle beams into contact with gas particles, meaning that this source of emittance growth can be non-negligible.

Beam measurement methods

A variety of methods exist to diagnose the different qualities of the beams produced in an accelerator. For example, the measurement of the energy spread of a beam is often found by dispersing it in a dipole spectrometer. If the dispersion is strong enough—and the beta-function is made sufficiently small—the second term in Equation 2.28 dominates and the measurement of the beam size in the dispersive plane can diagnose the beam's spectrum.

In transverse phase space we are often concerned with the beam's emittance, which can be found by knowing the beam's width and divergence at a given point. A common location where emittance measurements take place is about a beam waist located at a position s_0 in a drift, where the beta function reaches a minimum value β_0 and there are no external focusing forces ($k = 0$). In this case, the β that solves Equation

2.13 is given by

$$\beta(s) = \beta_0 + \frac{(s - s_0)^2}{\beta_0}, \tag{2.37}$$

where the other Twiss parameters are given by $\alpha_0 = 0$ and $\gamma_0 = \frac{1}{\beta_0}$. In the absence of dispersion, substituting Equation 2.37 into Equation 2.28 yields

$$\sigma_x(s) = \sqrt{\epsilon_{x,tr}(\beta_0 + \gamma_0(s - s_0)^2)}. \tag{2.38}$$

By measuring $\sigma_x(s)$ for different s about s_0 and fitting Equation 2.38 to the results, the parameters $\epsilon_{x,tr}$, β_0 and γ_0 can be extracted. These can then be used to calculate other parameters—namely the normalised emittance $\epsilon_{x,n}$, the waist beam size $\sigma_{x,0}$ and the divergence $\sigma_{x'}$ (see Equations in 2.26 and 2.27).

Such measurements can in principle be provided by the insertion of screens into the path of the beam, in order to measure its beam size at multiple points. This however is sometimes not practical; measurements need to take place over a range of locations that span both the waist and divergence regions of the beam's trajectory. A common solution is to instead change the lattice optics downstream of the waist, such that the beam size at a single screen essentially provides measurements of the trace space ellipse for different rotations.

A particular version of this method involves using multiple quadruples to image the beam at an object plane location s_{obj} onto a screen at an imaging plane location s_{im}. The beam size can therefore be measured at different s_{obj} by altering the strengths of the quadrupoles—a so-called *object plane scan*. For a given imaging optic, the s_{obj} is that satisfying the *point-to-point imaging* condition. In xx' space, the transfer matrix for $s_{obj} \to s_{im}$ that fulfils this condition has the form

$$\mathbf{R}_{s_{obj} \to s_{im}} = \begin{pmatrix} R_{11} & 0 \\ R_{21} & R_{22} \end{pmatrix}, \tag{2.39}$$

where the fact that $R_{12} = 0$ implies that the transformation of the beam size at the object plane $\sigma_{x,obj}$ to that at the image plane $\sigma_{x,im}$ is given by $\sigma_{x,im} = R_{11}\sigma_{x,obj}$ i.e. the measured beam size is simply that at the object plane magnified by a value $M_x = R_{11}$ and the effect of any divergence or pointing of the beam does not contribute. The same considerations also apply to the y-direction; in order to independently define the object planes in both directions whilst keeping their magnifications constant, a total of four degrees of freedom i.e. four independently adjustable quadrupoles are required.

Collective effects

Note that throughout this section, beams have been treated as ensembles of non-interacting particles. This however may not always accurately describe the beam evolution throughout an accelerator; sometimes, motion due to the self-fields of the beam needs to be considered.

An example of such a so-called collective effect is that due to space charge [40]. In general, this applies an extra repulsive force to the beam particles. For example, this will produce a contribution to the beam expansion that is not predicted by only considering its emittance. Furthermore, depending on the shape of the beam distribution, this force may be non-linear—potentially leading to emittance growth. To model this effect, computational simulations that solve the electric field due to the beam-charge distribution are needed. However, space-charge forces generally scale with $\frac{1}{\gamma^2}$—meaning that they can often be neglected for high energy beams.

Another collective effect is the interaction of the beam with the electromagnetic radiation that it produces. Such synchrotron radiation is emitted by the beam particles whenever their trajectories are bent e.g. in dipole magnets. If this radiation is emitted incoherently, its power P_I generally scales linearly with the amount of particles in the bunch N. However, in the case that the length of the bunch is shorter than the wavelength of the radiation, the emission becomes coherent; the bunch therefore behaves rather like a single particle emitting radiation with a power P_c given by [41]

$$P_c = NP_I = N^2 \frac{e^2 c}{6\pi\epsilon_0} \frac{\gamma^4}{R^2}, \tag{2.40}$$

where γ is the relativistic factor of the bunch and R is the bending radius of its trajectory. Note that the power of this *coherent sychrotron radiation* (CSR) now scales with N^2; P_c can therefore be much higher than P_I for high-charge, high-energy bunches in a deflecting field. Furthermore, this radiation travels in a straight line after emission; if the bending radius R of the particles is tight enough, CSR from particles in the tail of the bunch can in fact catch up to and interact with those at the head. Such an effect can lead to distortion of the bunch's longitudinal phase space and also increase its emittance, often making CSR a limiting factor in e.g. the compression and transport of high-current bunches. On the other hand, this effect forms the basis for Free Electron Laser (FEL) light sources [42]. By creating a feedback between the longitudinal position of the particles and the radiation they generate in a periodic magnetic field (an undulator), microbunches within the beam can form that produce coherent radiation—a process known as self-amplified spontaneous emission (SASE) [43].

2.3 Plasma physics

Often referred to as the fourth state of matter, plasmas consist of an ionised gas of ions and electrons. Plasmas are often characterised by their response to a perturbation [44]. Consider moving the electrons in a plasma from the stationary ion background by a small amount in one direction; this induces an electric field in the resulting region of charge separation that depends on the ion density n_i and electron density n_e, which in the quasi-neutral case are equal i.e. $n_e = n_i$. As the electrons return to their original positions, they oscillate at the *plasma frequency*

$$\omega_p = \sqrt{\frac{n_e e^2}{\epsilon_0 m_e}}, \tag{2.41}$$

with a corresponding *plasma wavenumber*

$$k_p = \frac{\omega_p}{c} = \sqrt{\frac{n_e e^2}{\epsilon_0 m_e c^2}}, \tag{2.42}$$

the inverse of which k_p^{-1} gives the characteristic scale-length of disturbances in the plasma known as the *skin depth*. This is closely related to the *plasma wavelength*,

$$\lambda_p = \frac{2\pi}{k_p}, \tag{2.43}$$

i.e. the wavelength of the plasma electron oscillations arising from a small disturbance.

2.3.1 Plasma generation

Generation of a plasma requires the input of a large amount of energy to a neutral gas in order for the electrons to be stripped from their corresponding ions. A common method for achieving this is the use of a *high-voltage discharge.* Although different subtypes of electrical discharge exist, they essentially rely on the movement of charge carriers under an electric field applied to a gas, resulting in an avalanche of collision ionisation events [45].

Another plasma generation method, known as *field ionisation*, involves the application of an electric field that directly distorts the potential barrier of the atom—see Figure 2.3 for illustration. This is distinct from single- or multi-photon ionisation [see Figure 2.3 (a)], where the electron directly absorbs the energy from incident photon(s) that exceeds its binding energy. Instead, field ionisation involves the potential being lowered by an applied electric field, increasing the probability that the electron can tunnel through the barrier—specifically referred to as *tunneling ionisation* [see Figure 2.3 (b)]. In the limit that the barrier is lowered so much that the electron can escape classically, this is referred to as *barrier suppression ionisation* [see Figure 2.3 (c)]. One of the distinguishing features between multi-photon and tunneling ionisation is the frequency of the electric field that is applied to the atom, with the classic example being a laser field. In this work, we are mainly concerned with ionisation due to linearly polarised TEM_{00} gaussian laser beams, which have a transverse electric field of amplitude E_0 that can be approximated by the envelope equation

$$E_\perp(r,z,t) = E_0 \frac{w_0}{w(z)} \exp\left(-\frac{r^2}{w(z)^2}\right) \exp\left(-\frac{t^2}{2\tau^2}\right), \tag{2.44}$$

where $w(z) = w_0\sqrt{1+\left(\frac{z}{z_R}\right)^2}$ is the radial beam width at longitudinal position z, w_0 is the waist size at $z = 0$ and $z_R = \frac{\pi w_0^2}{\lambda}$ is the Rayleigh length for a wavelength λ. In the temporal domain t, the pulse duration is denoted by τ. In the case that the period of the laser oscillations is lower than the typical tunneling time of the electron in the Coulomb field of the atom, multi-photon ionisation will dominate. On the other hand, if the oscillation period is longer such that the field can be considered constant over the timescale on which the electron is released, tunneling ionisation prevails. The latter also applies to where the period is so large that the field does not oscillate, meaning that tunneling ionisation theory also applies to e.g. the DC electric fields of particle beams [46]. For a relativistic gaussian electron beam, the transverse electric field is given by

$$E_r(z,r) = \frac{Q}{(2\pi)^{\frac{3}{2}}\sigma_z\epsilon_0 r}\left(1-\exp\left(\frac{-r^2}{2\sigma_r^2}\right)\right)\exp\left(\frac{-z^2}{2\sigma_z^2}\right), \tag{2.45}$$

where Q is the bunch charge and σ_r and σ_z are the bunch widths in the radial r and longitudinal z directions respectively [47]. Radially, this field increases approximately linearly within the bunch, reaching a maximum around $r \approx 1.5852\sigma_r$ before falling as $\frac{1}{r}$ for large distances. Conveniently, this maximum can be written in terms of the peak current $I_p = \frac{Qc}{\sqrt{2\pi}\sigma_z}$ as $E_{r,max} \approx 0.45\frac{I_p}{2\pi\epsilon_0 c\sigma_r}$.

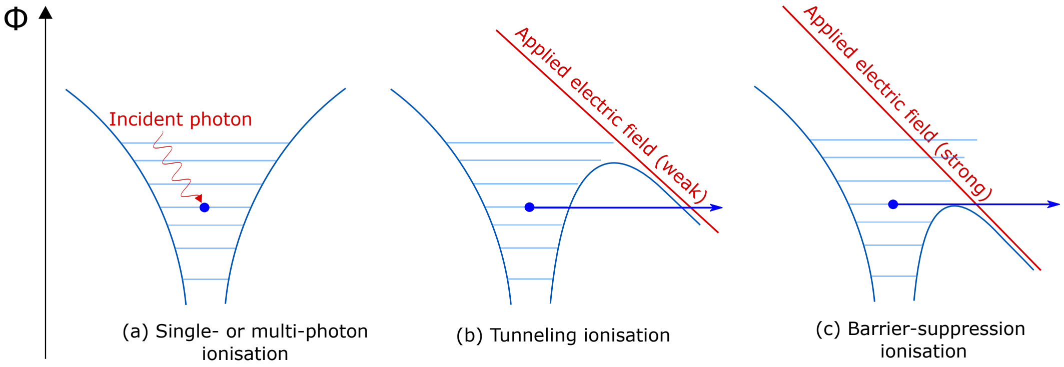

Figure 2.3: Ionisation due to an applied electric field. (a) In single- or multi-photon ionisation, an electron (blue dot) absorbs one or more photons (red curvy arrow) in order to be excited to higher energy levels (light blue lines) and eventually escape the atom. Field ionisation on the other hand involves the distortion of the atomic potential Φ (blue curves) via an applied electric field (red straight lines); in the case of (b) tunneling ionisation, this increases the chance of the electron escaping via tunneling through the potential barrier. (c) Eventually, the potential is reduced so much that the electron can escape classically—an effect known as barrier suppression ionisation.

A model to describe DC tunneling ionisation of complex atoms was formulated by Ammosov, Delone and Krainov [48], with an ionisation rate due to an applied DC electric field E_{dc} to create ions of charge Z given by

$$W_{ion} = \omega_a C^2_{n^*,l^*} \frac{(2l+1)(l+|m|)!}{2^{|m|}(|m|)!(l-|m|)!} \times \left(\frac{U_{ion}}{2U_H}\right) \left[2\frac{E_a}{E_{dc}}\left(\frac{U_{ion}}{U_H}\right)^{\frac{3}{2}}\right]^{2n^*-|m|-1} \times \exp\left[-\frac{2}{3}\frac{E_a}{E_{dc}}\left(\frac{U_{ion}}{U_H}\right)^{\frac{3}{2}}\right] \tag{2.46}$$

where $\omega_a = 4.131 \times 10^{16}\,\mathrm{s}^{-1}$ is the atomic unit frequency, $E_a \simeq 5.1\,\mathrm{GVm}^{-1}$ and $U_H = 13.6\,\mathrm{eV}$ and U_{ion} are the ionisation potentials of ground-state hydrogen and the atom in question. Additionally, l and m are the electron's orbital number and its projection, respectively, $n^* = Z\sqrt{\frac{U_H}{U_{ion}}}$ is the effective principle quantum number, $l^* = n_0^* - 1$ is the effective orbital number (n_0^* is the effective principal quantum number of the ground state) and in the limit that $l^* \ll n^*$,

$$C^2_{n^*,l^*} = \frac{1}{2\pi n^*}\left(\frac{2e}{n^*}\right)^{2n^*}. \tag{2.47}$$

Several variants of this so-called ADK formula exist, for example taking into account AC fields, and their validity depends on the regime being modelled [49]. Figure 2.4 displays Equation 2.46 plotted as a function of applied electric field for different species: the first four ionisation levels of argon (Ar-I–IV), the first two ionisation levels of helium (He-I–II) and the ionisation of molecular hydrogen (H_2-I). Note that the ionisation modelling of even a simple diatomic molecule such as H_2 is complicated; in reality, there are many pathways which can lead to fully dissociated, ionised hydrogen atoms [50, 51]. As a simplification, and as is the case in Figure 2.46, it is therefore common to approximate molecular hydrogen as fully dissociated atoms, with an ionisation energy of 15.4 eV [52–54].

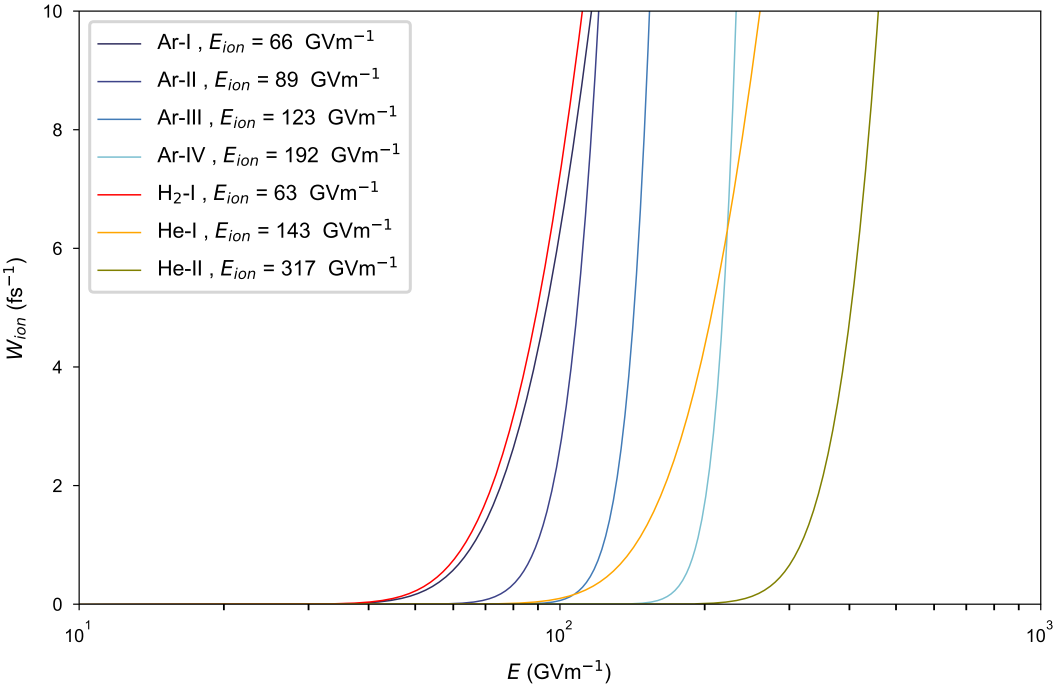

Figure 2.4: ADK ionisation rates for selected elements. The ionisation rates W_{ion} for the lowest ionisation levels argon, molecular hydrogen and helium are shown as a function of the applied field E, as calculated by the ADK formula in Equation 2.46. The electric field threshold for ionisation E_{ion}, defined here as the field at which $W_{ion} = 1\,\mathrm{fs}^{-1}$, are listed in the legend for each species.

One of the key advantages of plasma generation via field ionisation is that the ionisation rate does not explicitly depend on the density of the gas being ionised, in principle allowing for independent control over the shape and maximum density of the resulting plasma [55]. Additionally, the ionisation rate in tunneling ionisation increases over relatively small changes in electric field; the threshold nature of this process affords further opportunities for shaping of the plasma, where strong fields can be applied locally to ionise higher levels of a gas in a specific place [56]. Note that this plasma density profile tailoring can be made more complicated by the presence of non-linear *ionisation defocusing* effects [57], where the laser field is refracted by the non-uniform refractive index of plasma density gradients (see Section 5.1.3 on page 114 for examples).

2.3.2 Plasma evolution

Once generated, the plasma will evolve on the nanosecond timescale—much slower than e.g. the time taken for a particle or laser beam to traverse a typical plasma in the lab. Understanding this evolution though is useful in determining the plasma profile that a beam will interact with if its arrival is delayed significantly relative to plasma generation. Generally, depending on the temperature and density of the plasma immediately after generation, this evolution can also be affected by any additional energy sources after ionisation. The subsequent behaviour of the often non-isotropic, non-thermal plasma therefore involves a multitude of complex processes, including collisional impact ionisation (of electrons and ions) and excitation events [58, 59], bremsstrahlung radiation and hydrodynamic expansion [54, 60]. The latter effect can also include the formation of shock-fronts that expand and further interact with the surrounding material resulting in e.g. more impact ionisation. As a result, magneto-hydrodynamic (MHD) simulations are often required to model the full plasma evolution but are equally dependent on the initial conditions of the problem.

Constantly competing with ionisation processes is the recombination of electrons and ions to produce neutral atoms. Without additional input energy, recombination processes will eventually dominate until the plasma has completely recombined. There are many pathways that electrons and ions can take in order to reach a fully-recombined state; the simplest of these for atomic species is *radiative recombination*

$$A^+ + e^- \rightarrow A + h\nu \tag{2.48}$$

where an ion A^+ captures an electron e^-, forming a neutral atom A whilst emitting a photon with energy $h\nu$ equal to the excess between the kinetic and binding energy of the reactants. The dominant recombination channel is determined by the cross-section of the specific reaction, which in general depends on the atomic configuration of the ion, as well as the temperature of the plasma; the recombination rate is higher for reactants of lower kinetic energy [61]. In addition to the light that is emitted via recombination, radiation at discrete wavelengths (i.e. spectral lines) will also be emitted via the relaxation of electrons from excited states. Combined with the previously mentioned bremsstrahlung radiation, the emission-light due to recombination and de-excitation from the so-called *plasma afterglow* serves to cool the plasma but is also a detectable signal that contains valuable information [62–64]. Note that efforts are being made to understand the light emission from plasma channels used in PWFA—Ref. [64] for instance aims to model and measure the full evolution of a thin laser-generated plasma filament, from creation to recombination, which includes the effects of e.g. surrounding neutral gases that will be present in the experiment.

Overall, the complex dependencies of each stage of the plasma's evolution—from ionisation to recombination—mean that the lifetime of a plasma can vary dramatically. For example, Ref. [62] reports the lifetime of air-plasma emission signals to be on the nanosecond timescale for an 800 nm-wavelength laser of energy 100 µJ and pulse length 100 fs. Ref. [65] on the other hand observes air-plasma lifetimes of up to the millisecond-level for $\mathcal{O}(10\,\mathrm{ns})$-long laser pulses of energy 55 mJ. Discharge plasmas inevitably have a different time evolution; for example, in Ref. [66], the plasma density is measured to decay from $\sim 10^{17}\mathrm{cm}^{-3}$ to $\sim 10^{14}\mathrm{cm}^{-3}$ on the $\mathcal{O}(10\,\mu\mathrm{s})$ timescale—with the time until full recombination likely even longer.

2.4 Plasma wakefield theory

It was realised as early as the 1950s that the strong electric fields that can be generated in a plasma could in theory be used to accelerate charged particles [67, 68]. In 1979, Tajima and Dawson proposed using an intense laser pulse to ponderomotively drive a wake of plasma oscillations; the strong wakefields within could then be used to accelerate electrons to high energies over short distances [69]. As it was originally difficult to produce the short pulses required to drive such a *laser wakefield accelerator* (LWFA), several alternative solutions for generating the required laser intensities were proposed and demonstrated [70–72]. Eventually, the advent of compact, short-pulse laser technology [73] enabled the generation and acceleration of high-quality electron beams—referred to also as *witness* or *trailing bunches*—from an LWFA [74–76].

In the meantime, another method of driving plasma wakefields was proposed by Chen [77] and Ruth [78] et al., using intense charged particle beams rather than a laser pulse. First experimentally demonstrated in 1988 [79], beam-driven plasma wakefield acceleration has since developed rapidly, with notable milestones such as the energy doubling of a 42 GeV electron beam achieved [80]. Most common PWFA schemes use an electron bunch to drive a wakefield that itself accelerates electrons; however, electron acceleration has also been demonstrated using a proton beam driver [81] and in addition positron drivers have been used to accelerate positron bunches [82–85].

In the next sections, the basic theory underlying plasma wakefield is reviewed, starting with the generation of plasma wakes and followed by the injection and acceleration of particle bunches therein. Although large amounts of the fundamental theory applies equally to the cases of both PWFA and LWFA, the focus here is the electron beam-driven case, with comparisons made to LWFA where necessary.

2.4.1 Wakefield generation

Wakes in a plasma accelerator travel with a phase velocity v_p equal to that of their of their driver v_d. In the case of PWFA, the particle beam driver is usually relativistic, implying that, in a plasma of constant density, $v_p = v_d \approx c$. This already marks a contrast with LWFA, whose laser drivers propagate at the laser group velocity in a plasma which in general means that $v_p = v_d = c\sqrt{1 - \frac{\omega_p}{\omega}} < c$, where ω_p and ω are the plasma and the laser frequencies respectively. The consequences of the wake velocity in terms of injection and acceleration into the plasma wakefields are discussed in Sections 2.4.3 and 2.4.2 respectively.

In the mean time, as with any explanation into the generation of wakefields in a plasma accelerator, we begin with a description of linear wakefields, before discussing the theory and properties of those driven in the non-linear regime.

Linear regime

In all treatments of linear PWFA theory it is assumed that the density n_d of the relativistic ($v_d \approx c$) drive beam is much less than the plasma electron density n_e, i.e.

$$n_d \ll n_e. \tag{2.49}$$

Following the derivation in Ref. [86] which is based on the rigorous treatment in Ref. [87], consider a small perturbation δn to the original plasma electron density n_0, such that $n_e = n_0 + \delta n$; the corresponding velocity associated with this perturbation is given by $\mathbf{v} = \delta\mathbf{v}$ for the initially stationary plasma electrons. By substituting these into the continuity equation

$$\frac{\partial n_e}{\partial t} + \nabla \cdot (n_e \mathbf{v}) = 0, \tag{2.50}$$

we obtain a linearised version for the perturbation,

$$-c\frac{\partial \delta n}{\partial \xi} + n_0 \nabla \cdot \delta \mathbf{v} = 0, \tag{2.51}$$

where the transformation of ∂_t into the co-moving coordinate $\xi = s - ct$ has been used. A similar linearisation also applies to the equation of motion (i.e. the Lorentz force in Equation 2.8),

$$-\frac{\partial \delta \mathbf{v}}{\partial \xi} = -\frac{e}{mc}\mathbf{E} \tag{2.52}$$

where we assume the electric field perturbation $\delta \mathbf{E} = \mathbf{E}$ dominates over the magnetic field, implying another constraint to the linear theory—that $|\delta \mathbf{v}| \ll c$ and thus the plasma response is sub-relativistic. Differentiating Equation 2.51 with respect to ξ and substituting in Equation 2.52 we obtain

$$\frac{\partial^2 \delta n}{\partial \xi^2} - \frac{e n_0}{\epsilon_0 m_e}(\nabla \cdot \mathbf{E}) = 0. \tag{2.53}$$

Making use of Gauss' law in Equation 2.1 in terms of the net charge density $\rho = -e(n_b + \delta n)$, we substitute for $(\nabla \cdot E)$ in Equation 2.53 and simplify to get

$$\frac{\partial^2 \delta n}{\partial \xi^2} + k_p^2 \delta n = -k_p^2 n_d. \tag{2.54}$$

Equation 2.54 is that of harmonic density oscillations δn driven by a beam of density n_b, which can be solved via Green's functions to yield

$$\delta n(\mathbf{r}, \xi) = -k_p \int_{\xi}^{\infty} n_d(\mathbf{r}, \xi) \sin\bigl(k_p(\xi - \xi')\bigr) d\xi, \tag{2.55}$$

where $\mathbf{r} = (x, y)$ is the transverse coordinate.

In order to evaluate the longtiudinal and transverse (radial) electric fields that arise due to this perturbation, we substitute Equation 2.55 into the inhomegeneous wave equation for the electric field (see Equation 2.5), yielding

$$\frac{1}{c^2}\frac{\partial^2 \mathbf{E}}{\partial t^2} - \nabla^2 \mathbf{E} = \frac{e}{\epsilon_0}\nabla(\delta n + n_d) + \mu_0 e \frac{\partial}{\partial t}(n_d c \hat{\mathbf{z}} + n_0 \delta \mathbf{v}), \tag{2.56}$$

where the net current due to the beam travelling along the unit vector $\hat{\mathbf{z}}$ and the plasma perturbation is given by $\mathbf{j} = n_d c \hat{\mathbf{z}} + n_0 \mathbf{v}$, keeping only the terms linear in δn and $\delta \mathbf{v}$. A transformation into the co-moving coordinate plus the substitution of $\frac{\partial \delta \mathbf{v}}{\delta t}$ from

Equation 2.52 then produces a characteristic equation for **E**, which simplifies to

$$(\nabla_\perp^2 - k_p^2)\mathbf{E} = -\frac{e}{\epsilon_0}\nabla\delta n - \frac{e}{\epsilon_0}\nabla_\perp n_d. \tag{2.57}$$

Note that the last term on the right-hand-side of Equation 2.57 is that associated with the driver fields; here we are interested in the fields arising from the plasma response and so solve the equation only for the δn term. Assuming cylindrical symmetry, the longitudinal and transverse fields in this case are solved by their Green's functions

$$E_z(r,z) = \frac{e}{\epsilon_0}\int_0^\infty \frac{\partial\delta n(r',z)}{\partial z} K_0(k_p r_>) I_0(k_p r_<) r' dr' \tag{2.58}$$

and

$$E_r(r,z) = -\frac{e}{\epsilon_0}\int_0^\infty \frac{\partial\delta n(r',z)}{\partial r'} K_1(k_p r_>) I_1(k_p r_<) r' dr', \tag{2.59}$$

where I_n and K_n are n^{th}-order Bessel functions of the first and second kind respectively, and $r_<$ is the lesser of r and r' and $r_>$ the greater. In combination with Equation 2.55, Equations 2.58 and 2.59 allow the longitudinal and transverse electric fields due to a beam-driver-induced perturbation to be calculated analytically. Figure 2.5 shows an example of the typical (a) accelerating and (b) focusing fields in a linear wake.

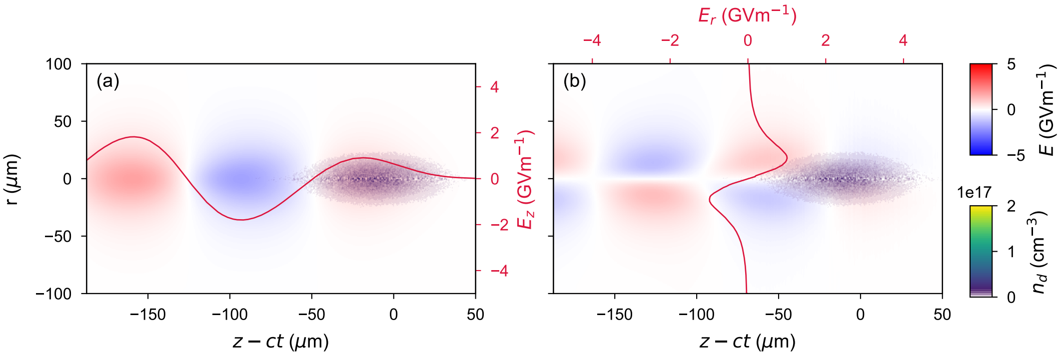

Figure 2.5: Linear regime of PWFA. The (a) accelerating and (b) focusing fields of a wake driven in the linear regime ($\frac{n_d}{n_e} \simeq 0.2$) are shown. Electrons placed in the accelerating phase of the approximately sinusoidal longitudinal field are able to gain energy from the wake. Directly behind the driver, the transverse fields are focusing for electrons.

Non-linear regime

As n_d increases, the transverse kick that the plasma electrons receive becomes stronger. For dense driver bunches ($n_d \gg n_e$), the plasma electrons are completely radially expelled from the drive beam propagation axis, leaving behind a pure ion channel surrounded by a thin plasma electron sheath—a so-called *blowout* [88]. As the electrons return to axis, their trajectories begin to cross such that their oscillations are no longer coherent—an effect known as *wave-breaking*, the physics of which has been studied as early as the 1950s (see e.g. Ref. [89]). Another characteristic of non-linear wave breaking is that the motion of the electrons can become relativistic. Although this behaviour means that the linear theory no longer applies [90], a characteristic longitudinal field strength can be defined for the non-linear regime in the form of the *wave-breaking field* [69],

$$E_{z,WB} = \sqrt{\frac{m_e c^2 n_0}{\epsilon_0}}, \tag{2.60}$$

which arises from Gauss' law (Equation 2.1) applied to the electric field produced in a charge separation region of characteristic length ${k_p}^{-1}$. Equation 2.60 can be expressed as $E_{z,WB}[\mathrm{V/m}] \approx 96\sqrt{n_0[\mathrm{cm}^{-3}]}$—highlighting the crucial $\sqrt{n_0}$ dependence of the longitudinal field but also suggesting that for a typical plasma density $n_0 = 1 \times 10^{16}\,\mathrm{cm}^{-3}$, accelerating fields as high as $10\,\mathrm{GVm}^{-1}$ can be sustained by the plasma.

Beyond this, the only general approach to finding the fields in the blowout regime revolves around the use of particle-in-cell (PIC) simulations (see Section 2.5.2). There does however exist a phenomenological approach formulated by Lu et al. [91, 92] that is able to analytically calculate these fields in terms of the radial shape $r_b(\xi)$ of the characteristic plasma electron sheath that marks the edge of the blowout and the normalised current-profile of the drive beam, given by

$$\lambda(\xi) = \int_0^{r \gg \sigma_r} n_b(\xi) r' dr' \tag{2.61}$$

where σ_r is the beam width. In the ultra-relativistic blowout regime where the maximum blowout radius r_m satisfies $r_m k_p \gg 1$, it can be shown that the longitudinal and transverse wakefields are given by

$$E_z \propto \frac{k_p r_b}{2} \frac{dr_b}{d\xi} \tag{2.62}$$

and

$$E_r - cB_\theta \propto \frac{k_p r}{2} \tag{2.63}$$

respectively. These two equations demonstrate some of the crucial properties of the accelerating and focusing fields in the blowout regime, typical examples of which are illustrated in Figures 2.6 (a) and (b), respectively. Firstly, Equation 2.62 is independent of the radial coordinate r ($\partial_r E_z = 0$), meaning that particles located at the same ξ but different r will experience the same accelerating field, which is determined purely by the trajectory of the plasma electron sheath r_b. Conversely, the transverse fields in Equation 2.63—which are due to the focusing force of the plasma ions—are ξ-independent ($\partial_\xi(E_r - cB_\theta) = 0$), such that particles experience the same focusing for all longitudinal positions within the blowout. Moreover, Equation 2.63 is linear in r—meaning that the emittance of particle bunches accelerated within the blowout can be preserved (see Section 2.2.2 on page 15).

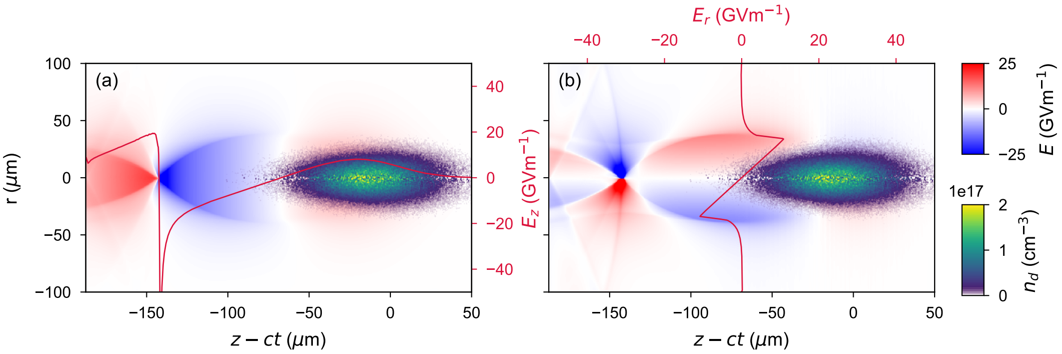

Figure 2.6: Non-linear regime of PWFA. The (a) accelerating and (b) focusing fields of a wake driven in the non-linear regime ($\frac{n_d}{n_e} \simeq 2$) are shown. This regime is characterised by the shape of the longitudinal fields which become very large at the back of the wake and also focusing fields that are linear throughout the blowout, allowing for emittance preservation of accelerated electron bunches.

2.4.2 Injection Methods

Once a wakefield is generated by a driver, injection of a particle bunch into its accelerating phase is the next challenge. Generally, this involves coupling the beam into the small region with characteristic length k_p^{-1} that is moving at a wake phase velocity of v_p. The latter presents a conventional definition of the *trapping condition* for injected bunches—their velocity v_i should be such that they can catch up with the wake i.e. $v_i \geq v_p$. As mentioned in Section 2.4.1, for a uniform plasma density, $v_p \approx c$ in PWFA, making this condition rather stringent when compared to LWFA where generally $v_p < c$. Broadly speaking, methods of injecting a witness beam into a plasma wake can be divided into two categories: *external* and *internal injection*.

External Injection

This injection mode refers to the acceleration of pre-accelerated particles originating from outside of the plasma. In the earliest PWFA experiments, particles in the tail of the drive bunch were accelerated [93]; ideally however, the witness bunch is distinct from the driver, for example originating from a separately generated beam at the linac injector [94] or via separation of the main bunch into a driver-witness pair [95].

One of the main advantages of the external injection of witness beams is that their velocity is usually already relativistic; as such, these beams automatically satisfy the trapping condition. The goal in external injection is therefore to couple the trailing bunch into the wakefield and accelerate it at high-gradient, all whilst crucially preserving its key beam qualities e.g. charge and emittance. PWFA stages relying on this injection mode—for example that studied in Chapter 4—therefore simply act as *energy transformers*, transferring the energy from a high-charge beam to a beam of lower charge in order to increase the latter's average energy.

Internal Injection

In contrast, completely new witness bunches can instead be generated directly inside the wake i.e. they are internally injected. Injected particles in this case generally have initial velocities that are less than that of the wake and so need to accelerate to the wake velocity quickly enough to become trapped. This is generally easier in LWFA, to the extent that unwanted injection (dark current) can occur [96, 97]. In contrast, whereas the strict trapping condition in PWFA makes internal injection challenging, it also means that witness beams can be generated that are largely dark-current free [98].

In order to trigger internal injection in PWFA, electrons can, for example, be released within the blowout via the additional ionisation of the background ions or any extra neutral dopant gas. This tunneling ionisation (see Equation 2.46) can be triggered either by the fields of the drive beam [99–101], by the wakefields themselves [102], or by the introduction of assistive ionisation lasers. The latter of these methods, which include collisions of multiple laser pulses [103] or the laser-ionisation of a dopant neutral gas (the so-called Trojan Horse plasma photocathode technique [104]), have the additional advantage that the injection mechanism is triggered via an external source—a laser—rather than the beam or wakefields which are highly coupled to the acceleration process. On the other hand, they also necessarily rely on the precise alignment and synchronisation of the injection laser(s) to the wake. Particles released on-axis in ionisation injection schemes can have very low transverse momenta, implying also low transverse emittances, which is often termed *cold injection*. By the same token, the particles are also generated with zero longitudinal momentum, making the trapping condition even more difficult to reach. Such injection schemes are therefore mostly limited to high-current drivers ($\gtrsim 5\,\mathrm{kA}$) operating in the strong blowout regime that have the necessary fields to trap at-rest particles.

An internal injection mechanism with relaxed trapping requirements involves the capture of electrons from the wake sheath. As noted in Section 2.4.1 on page 28, in the non-linear regime, these particles—particularly those at the wake vertex—can obtain longitudinal velocities that are comparable to v_p. This can then lead to the injection of particles that begin to extract energy from the wake—specifically known as *longitudinal wave-breaking* [105]. In the laser-driven case, this can occur relatively easily in what is known as *self-injection* [106, 107]. It is however common to trigger this injection by the controlled reduction of the phase velocity in a negative density gradient—a so-called *density downramp*—which has been demonstrated many times in LWFA [108–111]. Derived from the wake phase defined as $\chi = k_p(z)(z - ct)$, the *wake phase velocity* in a plasma of density $n(z)$ is given by [112, 113]

$$\frac{v_p}{c} = -\frac{\left(\frac{\partial \chi}{\partial ct}\right)}{\left(\frac{\partial \chi}{\partial z}\right)} = \frac{1}{\left(1 + \frac{\xi}{k_p}\frac{dk_p}{dz}\right)}, \tag{2.64}$$

where $\frac{dk_p}{dz} = \frac{k_p}{2n}\frac{dn}{dz}$, and $\xi = z - ct < 0$ behind the driver—illustrating the need for a negative density gradient $\frac{dn}{dz} < 0$. Simultaneously, when transitioning the downramp the length of the wake expands due to the increase in the plasma wavelength. As such, electrons originally located at a phase $\chi_i = k_{p,i}\xi$ will move to that located at $\chi_f = k_{p,f}\xi$, such that their *wake phase advance* is given by

$$\Delta\chi = \chi_f - \chi_i = \chi_i\left(1 - \sqrt{\frac{n_f}{n_i}}\right), \tag{2.65}$$

where n_i and n_f are the densities at the top and bottom of the downramp respectively whilst $k_{p,i}$ and $k_{p,f}$ are their corresponding plasma wavenumbers. Rephasing of plasma electrons into the accelerating region of the wakefield enables them to increase their velocity which, in combination with the wake phase velocity reduction, can lead to trapping.

To utilise density downramp injection in PWFA [114], although injection can be achieved with relatively modest driver currents [115, 116], the stringent trapping condition means that sufficiently steep downramps are required. Methods of generating such sharp ramps include the use of hydrodynamic shock fronts [117–119] or, as used in this work, the selective optical-ionisation of higher levels of the base plasma/dopant gas to produce a density spike—also known as a *plasma torch* [56]. Although particles are trapped with some residual momenta during this injection process, it has been shown that, under the right conditions, beams can be generated via density downramp injection that can compete in terms of quality with those produced by ionisation injection [120].

In any case, the use of internal injection in plasma accelerators has the potential to produce particle beams with brightnesses that exceed conventional technologies. Furthermore, if such PWFA-based *plasma cathodes*—for example that studied in Chapter 5—generate bunches that exceed the brightness of their drive bunch, they can be viewed as *brightness transformers*.

2.4.3 Wakefield acceleration

With a particle bunch trapped in the accelerating region of the wakefield, it can begin to rapidly gain energy. The next sections briefly discuss how the quality of acceleration can be defined, as well as some important physical effects that can take place in the process.

Transformer ratio

In their roles as energy transformers, it is important in a PWFA to maximise the energy that is transferred from the driver to the accelerating trailing bunch. In order to quantify this figure-of-merit, the *transformer ratio* is defined as [78, 121]

$$T_R = \frac{E^+_{max}}{E^-_{max}}, \tag{2.66}$$

where E^-_{max} is the maximum decelerating field experienced by the drive beam and E^+_{max} is the maximum accelerating field generated in the wake, which is ideally that sampled by the witness beam. The denominator of T_R is related to an important limitation in PWFA known as *driver energy depletion*, where electrons from the drive beam are decelerated to sub-relativistic velocities such that they fall back into the accelerating region of the wake and disrupt the acceleration process, essentially limiting the energy gain of the trailing bunch to $\sim T_R E_{driver}$, where E_{driver} is the initial average particle energy of the driver. Ideally then, the accelerating fields E^+_{max} need to be maximised whilst the onset of depletion is delayed by minimising E^{-}_{max}—something that is achievable by flattening the decelerating field along the length of the driver such that the drive bunch loses energy uniformly along its extent [122, 123]. Such field-flattening can be produced by tailoring the current-profile of the driver such that it matches the shape of the wakefield—implying that ramped, triangular bunches can maximise the value of T_R [124, 125].

Beam loading

Distortion of the electric fields in an accelerator by the space-charge fields of the accelerating bunch—known as *beam loading*—is an important effect both in conventional and novel accelerators [126]. In the context of wakefield accelerators, this can be seen as the witness bunch driving its own wake that destructively interferes with that generated by the driver. High-current trailing bunches in particular can significantly alter the local accelerating field, implying also a high rate of total energy extraction [112]. Also important is the concept of *optimal beam loading*. With analogy to the maximisation of the transformer ratio, the field experienced by the witness bunch can be flattened along the bunch extent by proper tailoring of the current-profile; again, the optimal profile matches the shape of the wake i.e. a trapezoid [127]. Crucially, this ensures uniform acceleration and thus preservation of the trailing bunch's energy spread.

Acceleration efficiency

A key concern in particle accelerators is their energy efficiency; in general, it is desirable to maximise the output energy of our particle beams for a given amount of input energy. For a wakefield accelerator, the baseline for the overall efficiency is the energy with which the driver is generated, relative to the input energy from the grid. In the case of conventional accelerators, estimates of the wall-plug-to-beam efficiency vary between machines. Ref. [128] quotes a grid power usage of 10 MW for the EuXFEL, which is converted into 40 kW of beam power; an efficiency of only 0.4 %. On the other hand, in the design for CLIC (the Compact Linear Collider), where emphasis is placed on maintaining high overall efficiency, the generation-efficiency of the drive beam which feeds the PETS (Power Extraction and Transfer Structures) that provide the main accelerating RF in CLIC is expected to be up to 55 % [129]; this however does not include overhead for cooling, ventilation and network losses. Such high drive beam acceleration efficiency assumes operation under full beam loading conditions, where almost all of the RF power supplied to the cavities is extracted by the drive beam. In the case of laser drivers, the generation-efficiency is currently typically limited by current technology to less than 0.1 % at low-repetition-rates [130]. The latter may however improve with e.g. multi-pulse LWFA schemes [130, 131] and also as laser technology matures [132, 133]. Note that the initial cost and energy involved in building facilities that house such machines is not considered here, as well as any costs involved that are not directly concerned with powering the machine (e.g control systems, air conditioning etc.)

Ideally, the amount of total energy deposited into the wake by the decelerated driver ΔW_{dec} would be maximised, offering more energy for the witness to potentially absorb. The efficiency of the driver energy deposition can therefore be quantified by the *driver-to-wake energy-transfer efficiency*,

$$\eta_{d\rightarrow wake} = \frac{\Delta W_{dec}}{W_{dec,0}} \tag{2.67}$$

where $W_{dec,0}$ is the initial total energy of the driver.

Beyond this, the efficiency with which the witness is accelerated should then be considered. First of all, the total energy gain of the accelerated bunch ΔW_{acc} can be compared to ΔW_{dec} in what is known as the *wake-to-witness energy-transfer efficiency*,

$$\eta_{wake\rightarrow w} = \frac{\Delta W_{acc}}{\Delta W_{dec}}. \tag{2.68}$$

This can also be seen as an *instantaneous energy-transfer efficiency* as it can be expressed in terms of the energy gain and loss per unit length. However, for a true measure of the energy efficiency of a PWFA, the efficiencies in Equations 2.67 and 2.68 can be combined, yielding

$$\eta_{d\rightarrow w} = \eta_{d\rightarrow wake} \cdot \eta_{wake\rightarrow w} = \frac{\Delta W_{acc}}{W_{dec,0}}, \tag{2.69}$$

which is a measure the witness energy gain relative to the input driver energy. Maximising $\eta_{d\rightarrow w}$ therefore requires simultaneously maximising both $\eta_{d\rightarrow wake}$ and $\eta_{wake\rightarrow w}$. A high T_R means both that a large amount of energy can be extracted from the driver before depletion and also that the witness experiences high accelerating gradients—implying an increase to both $\eta_{d\rightarrow wake}$ and $\eta_{wake\rightarrow w}$. Furthermore, the high rate of energy extraction that arises with strong beam loading consequently implies high $\eta_{wake\rightarrow w}$[134–136].

Note that an effect that is highly consequential to the field of plasma wakefield acceleration as a whole is the *efficiency-instability relation* [137, 138]. Applying to both conventional and novel accelerators, this relation states that the longitudinal wakefields that are associated with strong beam loading and therefore high energy-transfer efficiency also come with an increase in the strength of transverse wakefields. As such, an accelerator operating with high energy-transfer efficiency is more prone to the onset of transverse instabilities from small beam misalignments, resulting in emittance growth as explained in Section 2.2.2 on page 15. The study of this effect is therefore of high importance and is a key application of the work presented in Chapter 4.

Wakefield optimisation parameter

In order to encapsulate all of the qualities of the wakefield discussed so far, a figure-of-merit known as the *wakefield optimisation parameter* was defined as [139]

$$\Omega = \frac{\sigma_\delta}{\eta T_R}, \tag{2.70}$$

where $\eta = \eta_{wake \to witness}$ in this case and σ_δ is the energy spread of the witness normalised to its mean energy gain, such that Ω is purely dimensionless and independent of the acceleration length. This parameter can be seen as a measure of how well optimised the shape of the wake is; a high T_R implies a flattened decelerating field for uniform driver energy loss along with high accelerating fields, whereas a large η and low σ_δ suggests strong energy extraction via beam loading with an accelerating field that is uniform along the witness bunch. A lower value of Ω therefore means that the witness can be accelerated at high-efficiency for longer before the onset of driver depletion, all whilst minimising energy spread growth.

Matching and extraction

In the same way that Equation 2.11 describes the transverse motion of single particles in an accelerator, the *beam envelope equation* [140] describes the evolution of the transverse beam size σ_r, where r can be either of the transverse coordinates or, in the case of a cylindrical beam with symmetric focusing, the radial coordinate. For a beam with relativistic factor γ and normalised emittance ϵ_n, the complete form of the envelope equation is given by [141]

$$\sigma_r'' + \frac{\gamma'}{\gamma}\sigma_r' + k^2\sigma_r = \frac{\epsilon_n^2}{\gamma^2\sigma_r^3} + \frac{k_{sc}}{\gamma^3\sigma_r}, \tag{2.71}$$

where k and k_{sc} are the normalised focusing gradients due to external focusing and space charge forces respectively and the derivatives are with respect to the longitudinal coordinate. Often for low charge, high energy beams, the k_{sc} term can be neglected.

Considering this case whilst also neglecting any changes in the beam energy ($\gamma' = 0$), Equation 2.71 becomes

$$\sigma_r'' + \left(k^2 - \frac{\epsilon_n^2}{\gamma^2 \sigma_r^4}\right) \sigma_r = 0, \tag{2.72}$$

where, in the case of the focusing forces in a PWFA blowout (see Equation 2.59), it can be shown that $k = \frac{\omega_p c}{\sqrt{2\gamma}}$ [142]. Note that Equation 2.72 applies to both witness and drive beams that are focused purely by the ion channel in the blowout regime; furthermore, it can also describe the transverse size evolution of individual beam slices. In order for the beam to be *matched*, the beta function β must by definition not change, implying that α and the derivative term in Equation 2.72 must equal zero. In this case, the *matched beta-function* is given by [31]

$$\beta_m = \frac{\sqrt{2\gamma}}{k_p}, \tag{2.73}$$

where the definitions of the beam size with no dispersion $\sigma_r = \sqrt{\beta \epsilon_{tr}}$ and the normalised emittance $\epsilon_n = \epsilon_{tr}\gamma$ have been used. If these matching conditions are not met, the betatron oscillations of individual particles will lead to oscillations in the beam envelope. Note that for a given k_p, only particles with relativistic factor γ can be simultaneously matched to the plasma; for large energy spread bunches, this means that matching of the full bunch is impossible, leading to emittance growth via betatron decoherence, as described in Section 2.2.2 on page 16.

In general, the plasma density and therefore k_p are not constant throughout the plasma; this is especially true at the plasma entrance and exit, where the density must ramp up and down respectively. Maintaining matching in these regions i.e. $\alpha = 0$ is possible through the use of gradually changing *adiabatic ramps* [143–145]. Another goal in the extraction of accelerated bunches from the exit downramp is to reduce their divergence as much as possible, facilitating easier capture with post-plasma optics [38, 146] and also reduced emittance growth due to correlations between transverse position and longitudinal momentum (see Section 2.2.2 on page 15).

2.5 Computational tools

Although much of the theory discussed here is important for understanding the physical processes in plasma accelerators, accurate modelling of a great deal of the physics—for example collective effects in particle beams or the non-linear blowout regime in PWFA—can only be achieved through numerical methods. The computational tool used then necessarily depends on the nature of the problem.

2.5.1 Accelerator beamline simulations

At the most basic level, the transport of particle beams—which are often simulated as an ensemble of macroparticles—can be modelled with first-order matrices (see Section 2.2.1 on page 10). However, for more accurate results, higher-order and/or collective effects need to be taken into account, often requiring a numerical approaches that are integrated into programs such as OCELOT [147], which is the particle tracking code used in this work.

In OCELOT, the tracking method can be chosen on an element-wise basis; for example, all elements can be modelled with either 1st (linear) or 2nd order matrices. In addition, for the e.g. quadrupoles, a tracking method based on Runge-Kutta (RK) integration can be used that requires the input of a field map for the particular element. Finally, physics effects can be applied over the whole or just defined parts of the beamline. Such an effect is that of space-charge forces, requiring the solution of Poisson's equation in the bunch frame, which can be quite computationally expensive.

2.5.2 Particle-in-cell simulations

A common simulation tool—especially for plasma accelerators—is that based on the Particle-In-Cell (PIC) method [148]. As the name suggests, this involves approximating particle distributions as a large number of macroparticles moving under the influence of electromagnetic fields that are defined on a grid of cells. Whereas the position and momentum of the particles are defined by the electromagnetic forces, likewise the charge and current distributions determine the electromagnetic fields. As such, the PIC method involves the following iterative cycle [149]:

1. **Particle push:** The macro particles are pushed according to discretized versions of the Newton-Lorentz equations of motion in the calculated electromagnetic field.

2. **Current deposition:** From the particles positions and momenta, the charge and current density are calculated.

3. **Field solving:** Maxwell's equations are solved on the simulation grid for the fields arising from the charge and current distributions.

As with any numerical modelling, the resolution of the simulation is key; the cell sizes and also time step of the PIC cycle needs to be small enough to resolve the relevant physical processes. A common way of determining the correct resolution is to decrease the cell size and/or time step until the relevant output parameters do not change significantly; in this case, the simulation is considered *converged*. For the situations requiring full 3D modelling, the PIC code WarpX is used in this project [150]. However, when modelling systems that are approximately cylindrically symmetric—often the case in plasma accelerators—the quasi-3D PIC code FBPIC is utilized [151].

Chapter 3

Experimental Facility

FLASHForward (Future Oriented Wakefield Accelerator Research and Development at FLASH) is an experimental beamline situated at FLASH that aims to exploit the excellent beams that the linac generates for cutting-edge PWFA studies [152, 153].

In this chapter, the FLASHForward experimental facility is introduced, where all of the experimental studies in this thesis were performed. The basis for this facility—the FLASH linac—is introduced first, before then describing in more detail the individual parts of the FLASHForward beamline. Necessary for some of the projects at the facility, the FLASHForward 25 TW laser system is also outlined. Finally, a brief summary of the different experimental activities at FLASHForward is then given.

3.1 FLASH

Initially starting life as the TeV Superconducting Linear Accelerator (TESLA)—a future linear collider test-bed with an integrated FEL—the Free-electron Laser in Hamburg (FLASH) is a facility at DESY that now delivers SASE-FEL radiation in the vacuum ultraviolet (VUV) to soft-xray (XUV) regimes to photon science users from all over the world [154]. As well as empowering photon science for the best part of two decades [155], FLASH has also served as a prototype for its successor, the European XFEL [156]. In the next sections, the key features of FLASH are summarised (see also Figure 3.1 (a)), with particular emphasis on those that make the linac an ideal source of electron bunches for a PWFA experiment.

3.1.1 Photocathode injector

A Cs_2Te photocathode illuminated by UV laser radiation acts as the electron source for the FLASH linac [157]. A choice of three UV lasers are available: two with similar

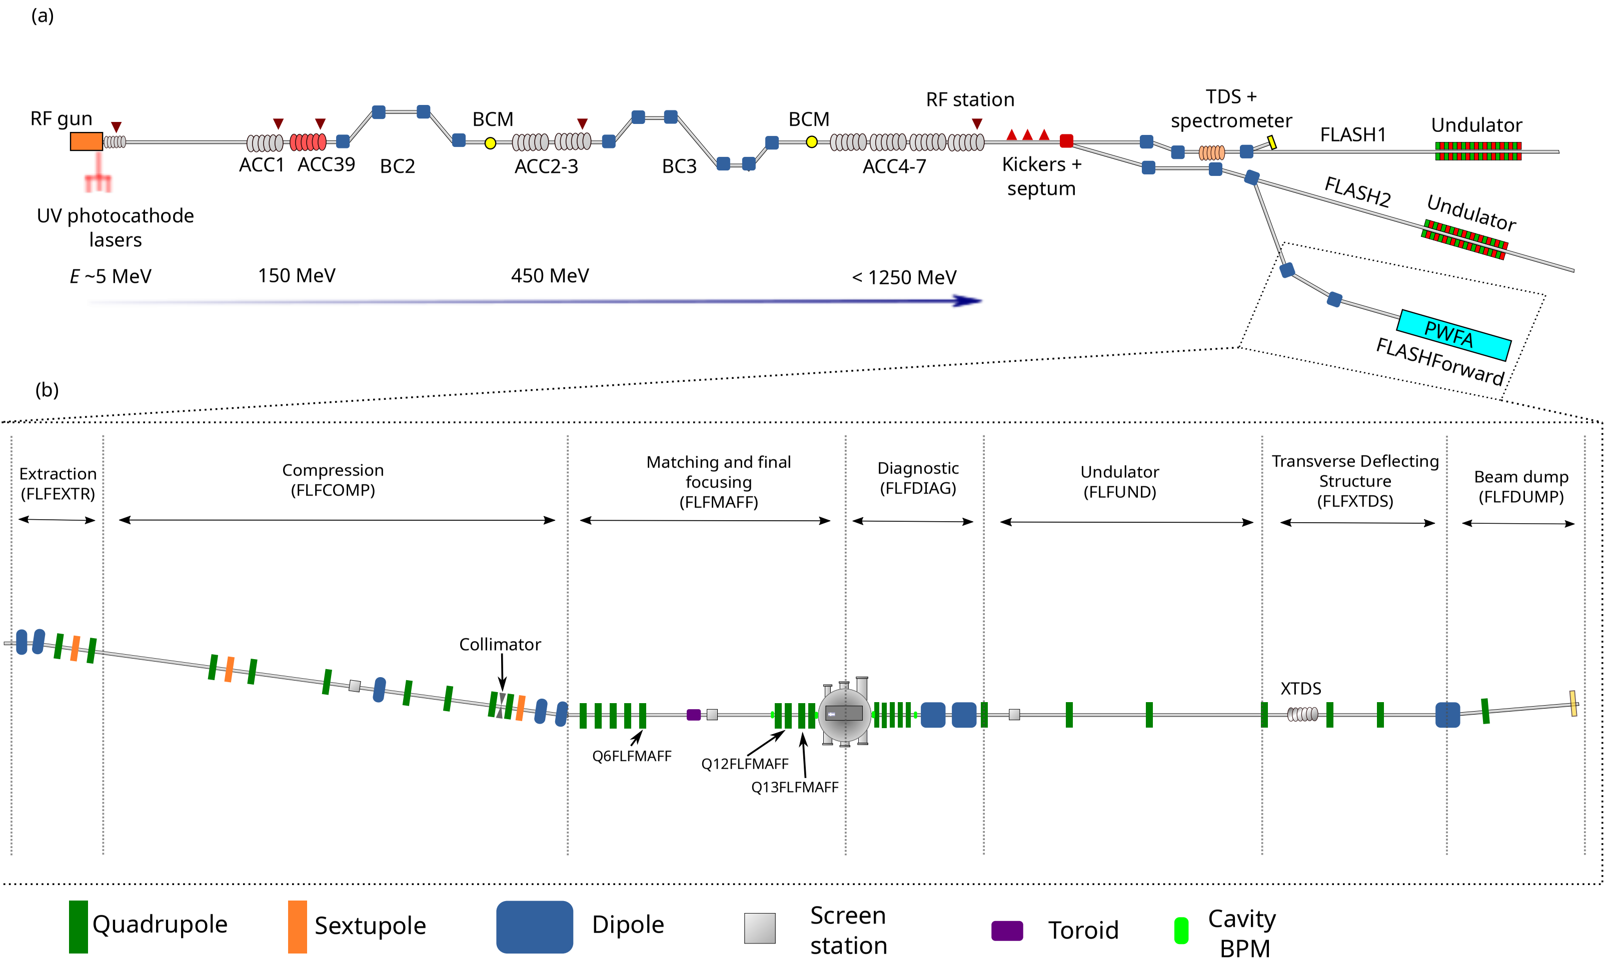

Figure 3.1: FLASH and FLASHForward beamlines. (a) Starting from injection at the photocathode, the FLASH beam is accelerated and compressed before being delivered to a selection of three beamlines: FLASH1, FLASH2 or FLASHForward. (b) A more detailed illustration of the FLASHForward lattice is shown, from extraction from FLASH2 through a radiation protection wall to the dump at the end of the beamline. The names of each beamline section are also shown, along with their short-hand names e.g. FLFCOMP.

pulse lengths of 4.5 ps and 6.5 ps that are able to inject bunches with charges up to 2 nC and another with a shorter pulse length between 0.8 ps and 1.6 ps for low-charge (20–120 pC)), short-pulse operation [158, 159]. Emitted directly into a conventional 1.3 GHz (L-band) RF gun, electrons are rapidly accelerated to energies of 5.6 MeV, which in combination with a focusing solenoid mitigates emittance growth due to space charge effects [154].

3.1.2 Superconducting radio-frequency modules and bunch compressors

A key technology that FLASH leverages are superconducting radio-frequency (SCRF) cavities [160]. Cooled to a temperature of 2 K, these modules generate only small amounts of resistive loss at their walls and so are able to sustain long RF pulses capable of accelerating long trains of electron bunches that can serve more than one beamline (see Section 3.1.3). Designed to achieve accelerating gradients of up to $25\,\mathrm{MV\,m^{-1}}$, a series of up to seven of these modules (see ACC1-7 in Figure 3.1 (a)) operate at the same 1.3 GHz as the RF gun. In addition to these acceleration cavities, a third-harmonic structure operating in a decelerating mode (see ACC39 in Figure 3.1 (a)) is able to eliminate any curvature that is imprinted on the longitudinal phase space during off-crest acceleration in the first accelerating module [161, 162]. While the resulting linearity of the phase space is important in providing compression suitable for lasing, it is also vital for detailed bunch profile tailoring in FLASHForward (see Section 3.2.2).

Compression is achieved in FLASH via two magnetic chicanes (see BC2 and BC3 in Figure 3.1 (a)) operating at nominal energies of 150 MeV and 450 MeV and fixed R_{56} values of 180 mm and 43 mm respectively [154, 163]. Tailoring of the compression provided by each chicane is then possible via tuning of the combined chirp in the preceding acceleration modules; the combination of the ACC1 and ACC39 modules, referred to as ACC139, controls the chirp upon entrance to BC2 whilst the combination ACC2 and ACC3 provides further control before entering BC3.

Such a setup provides electron bunches with energies ranging from 0.3–1.2 GeV, bunch lengths of 50 ps– 5000 ps and peak currents up to 2.5 kA, all whilst maintaining transverse emittances between 0.7–2 mm mrad and a linear, tunable longitudinal phase space with projected energy spreads below 0.1 %. It should be noted that these are the beam parameters at the end of the FLASH linac; properties such as the emittance can naturally change during propagation through the subsequent beamlines. Regardless, these beams possess key features that make FLASH ideal for enabling PWFA experiments in the FLASHForward beamline.

3.1.3 Time structure

Owing to the long RF pulses that can be sustained in the accelerating cavities, FLASH is able to deliver trains—also known as macropulses—of bunches, each with maximum duration 800 µs (see Figure 3.2) (a)) [164]. The photocathode can produce macropulses at repetition rates of 1 Hz, 5 Hz or 10 Hz; the spacing of bunches within a macropulse is then determined by the operating frequency of the injector laser, between values of 0.33 and 25 µs [154]. Furthermore, by adjusting the delay of the injector laser with respect to the RF signal, the time-of-arrival of the accelerated beams can be finely adjusted by 0.77 ns—the separation between acceleration buckets in the 1.3 GHz signal. This feature therefore allows the delay of the electron beam with respect to any other systems synchronised to the RF signal, for instance the FLASHForward laser (see Section 3.3). Moreover, by simultaneously injecting with two of the UV lasers into the same RF flattop but delaying one with respect to the RF signal, bunch separations as low as 0.77 ns can be achieved—a key capability in probing the evolution of the plasma wakefields on such timescales [165].

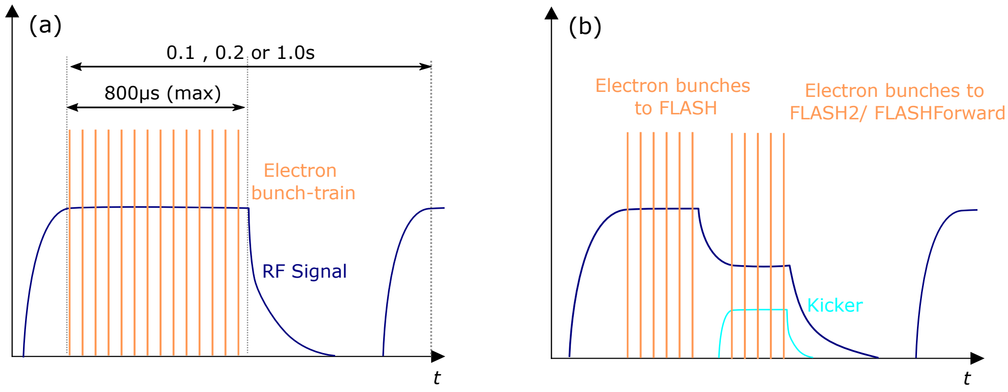

Figure 3.2: Bunch train structure at FLASH. (a) FLASH is able to produce RF pulses (navy curves) with a flattop length of up to 800 µs, at a rate of 1, 5 or 10 Hz. Bunches accelerated within these macropulses (orange lines) can be spaced by between 0.33–25 µs. (b) Activating the kicker magnets for part of the main macropulses duration (light blue curve) allows part of the bunch train to be diverted to one of the parallel beamlines—using a different injector laser affords customised bunch charge and spacing for this section of the macropulse. Modulating the RF signal in this time window provides some control over the energy and compression of these bunches, independent of parameters in FLASH1.

For photon-science applications, the complete train of bunches in a single macropulse is rarely fully used. In order to further exploit the high-average power of the linac, parts of these bunch trains can be kicked into the parallel FLASH2 beamline—and optionally after that the FLASHForward beamline—by way of a combination of kicker magnets and a septum deflecting in the horizontal plane (see Figure 3.1 (a)). This mechanism essentially divides the original macropulse in two (see Figure 3.2 (b)). To control the charge, time-of-arrival and intra-bunch separation of the bunches within the new sub-macropulse, one of the other UV lasers can be used as the injector. The amplitude and phase of each RF station can also be adjusted to modulate the RF signal of the sub-macropulse with respect to the original, providing some degree of longitudinal compression control independent of the bunches that are sent to the FLASH undulators [166].

3.1.4 Beam diagnostics

FLASH is equipped with a variety of diagnostics for measuring the properties of the electron beam. Beam-position monitors (BPM) measure the orbit of the beam at various points in the machine whilst also, along with toroidal current transformers, providing measurements of the bunch charge [167]. Radioactivity induced by charge losses is detected by beam-loss monitors (BLM), while the timing of the beam is monitored by beam-arrival monitors (BAM) [168]. Measurements of the beam's transverse profile are provided by insertable screens located in stations along the beamline [169]. Longitudinal measurements are enabled by bunch-compression monitors (BCMs) [170] and, for full longitudinal phase space tomography, a transverse deflection structure (TDS) combined with a dipole spectrometer [171]. Some elements of this extensive diagnostic suite—for example the BPMs, toroids and screen stations—are then also found in the FLASHForward beamline, integrated with the existing hardware and software infrastructure of FLASH.

3.1.5 Synchronisation

In order to synchronise FLASH to its applications at the femtosecond level, a system is deployed that uses an optical signal derived from a femtosecond mode-locked laser oscillator as a reference [172]. As well as exhibiting excellent short term stability, the advantage of the optical reference is that the temporal diagnostics can be based on optical signal mixing. Example signals that the reference can be mixed with are the electro-optical signals in the BAMs or any of the lasers based at FLASH—including those at the photocathode and those used in pump-probe FEL experiments. The full distribution of the master reference signal—handled by a network of optical fibres whose round trip time is monitored by optical cross correlators—is detailed in Ref. [173]. Ultimately, in combination with an active feedback system, this technique is able to achieve beam-arrival time jitters of 30 fs rms when producing X-ray pulses in the FEL that are 90 fs in length. In principle, the FLASHForward laser can also be included in the list of lasers to benefit from this all-optical synchronisation method. Currently however, synchronisation of this laser to FLASH relies on an RF-based locking mechanism [174] (see Section 5.1.2 for more details), with an upgrade from this precursor setup to the all-optical scheme foreseen in the future.

3.1.6 Feedbacks and beam stabilisation

As well as those concerned with timing, feedbacks are also available for other aspects of the FLASH beam. For example, a low-level RF (LLRF) system is used to precisely regulate the amplitude and phase of the accelerating fields in the SCRF cavities [175], whilst BCMs located after the bunch compressors yield information that can be fed-back on to stabilise the bunch compression. Moreover, and also extending to the parallel beamlines, orbit stabilisation is achieved by an automated control system that adjusts steering magnet strengths in the accelerator lattice in response to orbital changes detected by the BPMs.

3.1.7 Extraction to parallel beamlines

For either dedicated or parallel operation in FLASH2 or FLASHForward, bunches from FLASH are first kicked by the kicker magnet introduced in Section 3.1.3 into the FLASH2 extraction arc (see Figure 3.1) [176]. As well as achieving separation from the main FLASH beamline, the key design goals of this arc are that the longitudinal and vertical dispersions should be closed before its end—the former implying that the original bunch compression should be preserved. Meanwhile, measures must be taken to reduce the effects of CSR in the bending dipoles, such as minimising the beta functions at these locations. Despite this, these effects are often non-negligible—especially if high beam currents are desired in FLASH2 or FLASHForward.

3.2 FLASHForward

Shown in more detail in Figure 3.1 (b), the FLASHForward beamline is installed in the same experimental area as FLASH2. The next sections give a brief overview of the individual beamline segments that make up the FLASHForward beamline and their different purposes during PWFA experiments.

3.2.1 Extraction section

Separation from the FLASH2 beamline is granted by two static horizontal dipoles; as such, FLASHForward and FLASH2 cannot be operated in parallel with each other. Just as with the FLASH2 extraction arc, a balance must be found between providing enough separation and minimising CSR effects. A bending angle of 8 deg was ultimately found to be a suitable compromise between these two factors.

3.2.2 Compression section

Continuing at the same angle as the extraction section, the beamline then leads into the compression section. The name of this region of the lattice reflects one of its primary purposes—to provide further tunability to the longitudinal compression of the beam independent of that in the FLASH linac. Again, a simultaneous requirement is that the dispersion at the end of the section is zero, and that CSR effects are mitigated throughout.

A key element of this part of the beamline is the collimator. Located where the dispersion is large and the beta functions small, the energy coordinate of the beam is transformed into a horizontal position. By inserting a combination of three masks into the beamline at this location, particles of different energies within the beam can be removed (see Figure 3.3) [177].

For a beam with a linear energy chirp, the momenta of the particles are correlated to their longitudinal position, meaning that the energy collimation in this case translates to removal of particles from different longitudinal parts of the bunch. The power of this technique in double-bunch-current-profile tailoring is exemplified in Chapter 4. Additionally, this device enables diagnostic capabilities such as the analysis of individual bunch slices [178] and also the high-resolution sampling of plasma wakefields [179].

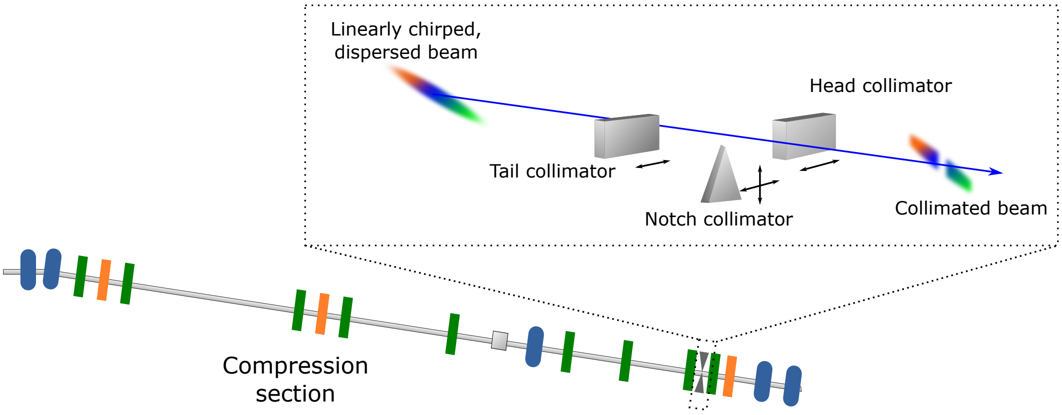

Figure 3.3: Energy collimation in the FLASHForward compression section. For a beam with a negative, linear chirp, the high-energy collimator removes particles from the tail whereas the low-energy collimator removes particles from the head; the amount of charge removed can be adjusted by inserting the masks by different amounts. Additionally, the wedge-shaped notch collimator can be moved both horizontally and vertically to create a notch in the beam with a variable position and width respectively.

3.2.3 Matching and final focusing

At the end of the compression section, the beam is deflected by two dipole magnets into the matching and final focusing region of FLASHForward. A group of five quadrupole magnets allow the measurement of the beam's transverse emittance at a screen station located approximately in the middle of this section. They are also key in establishing large beta functions in the last four final focusing quadrupoles that are designed to demagnify the beam in both directions to a transverse size that is approximately matched to the plasma density, with beta functions usually of order 10 mm. Normally, an optic is established here such that the size and location of the beta function waist can be adjusted in x and y approximately independently (see Q12FLFMAFF and Q13FLFMAFF in Figure 3.1). Additionally, one of the upstream quadrupoles, Q6FLFMAFF, is found to provide good control over the asymmetry of the beam in the transverse plane.

For quick diagnosis of the beta functions at the waist, a novel method was developed that utilises the two high-precision cavity BPMs at the chamber entrance and exit. Relying on the fact that, for small beta functions, the beam phase space and centroid jitter phase spaces are similar, Lindstrøm et al. [178] demonstrated that a measurement of the centroid jitter over multiple shots allows the beta functions in each plane to be approximately reconstructed. This technique therefore proves invaluable in establishing beta functions with approximately the correct waist size and location; further in-situ optimisation of the optic is then performed with the beam interacting with the plasma.

3.2.4 Interaction chamber

The matching and final focusing section of the beamline ends at the centre of the 500 mm-diameter interaction chamber which is itself split into upper and lower sections. Maintained at the Ultra-High vacuum (UHV) environment of FLASH ($\sim 1 \times 10^{-9}$ mbar), the upper part of the chamber contains the plasma targets and diagnostic screens mounted on a baseplate. Manipulation of this baseplate is performed by a parallel mover (hexapod), housed in a $\sim 10^{-3}$ mbar-level chamber below, connected via a feedthrough and bellows system. This allows for precision movement of the components of the experiment that are exposed to high vacuum, without the mover itself requiring UHV compatibility.

Gas cells and pumping

Figure 3.4 shows a key element within the upper chamber—the gas cell. Two sapphire blocks milled with 50 mm and 195 mm-long channels form the capillaries [180] and are held together by a PEEK (Polyether ether ketone) holder. Inlets connected to a buffer volume provide the option to fill the channels with a selection of gas species and mixtures at a range of backing pressures. Both of the experiments detailed in this work use argon as the primary gas species; there is however the possibility of filling the capillaries with other gases such as hydrogen, helium, nitrogen and krypton.

Electrodes at the entrance and exits of the plasma cells allow high-voltage discharge pulses to be passed through the constituent gas to generate a plasma. The initial density profile of this plasma is a function of the pressure inside the capillary and also the discharge voltage used. However, further tailoring of the plasma density that the electron beams actually experience can be achieved by allowing the plasma to evolve on the microsecond time scale [66].

Lasers can also be used to create a plasma in the capillary via field ionisation (see Section 3.3). Whilst a laser co-propagating with the electron beam forms the main accelerating channel, the capillaries also each feature two narrow holes located 20 mm downstream of their entrances that allow a perpendicularly propagating laser to be coupled in, for the purposes of creating a sharp downramp in the X-1 plasma cathode experiment (see Chapter 5).

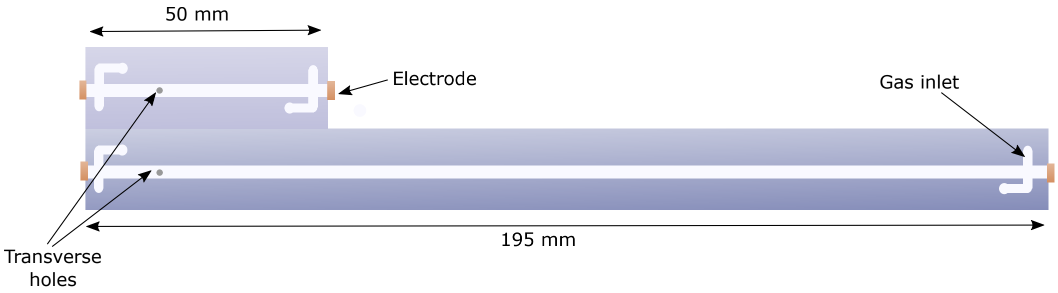

Figure 3.4: Gas cells at FLASHForward. The 50 mm and 195 mm-long gas capillaries are each formed from two milled sapphire blocks, held together by a PEEK holder (not shown).

The absence of windows at either end of the capillaries crucially avoids any scattering-induced degradation of the electron beam, but does also mean that gas is introduced into the beamline vacuum. Two turbo-molecular-pumps connected directly to the chamber remove the majority of this gas load; any that remains is removed by a total of 6 differential pumping stations—three upstream of the chamber and three downstream. Ultimately, the maximum gas load that these pumps can handle sets an upper limit on the pressure allowed in the capillary.

Note that a particular mode of generating plasma with the lasers outside of the capillary can be enabled by flooding the chamber with low density gas. This is possible by closing the shutters to the two chamber turbo-molecular-pumps and introducing gas into the cell at low backing pressures. A low density of ambient gas then fills the chamber whilst the vacuum up- and downstream is still protected by the differential pumping system. Although not providing high enough densities for any acceleration experiments at FLASHForward, this mode allows the detection of laser-generated plasmas via their emission light with no obstruction by the capillary and its holder—useful for gauging the ionisation capabilities of the lasers and also aligning them to the electron beam (see Sections 5.1.2 and 5.1.4).

Upper chamber layout and diagnostics

Figure 3.5 shows the main components—including the capillaries—located in and around the upper part of the interaction chamber. In the nominal position, the electron beam and/or longitudinal laser can pass through one of the capillaries and propagate downstream to the diagnostic section of the beamline (purple arrow); the short or long capillary can be selected by a simple vertical translation of the baseplate. When in use, the transverse laser propagates straight through the holes in the side of the capillary, to be diagnosed by the upstream alignment camera on the other side of the chamber (red arrow). Alternatively, the baseplate and its attachments can be moved completely out of the path of the passing beams.

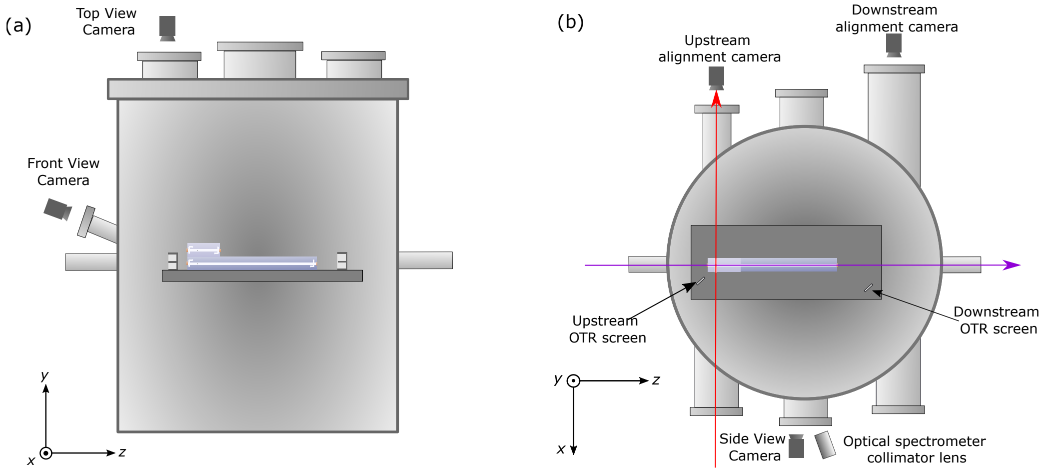

Figure 3.5: Interaction chamber layout at FLASHForward. (a) Side- and (b) top views of the interaction chamber are shown, with the key elements in and around the chamber highlighted. In its current position, the hexapod is located such that the longitudinal laser and/or electron beam can pass downstream through the capillary (purple arrow), whilst the injection laser could pass perpendicularly through the transverse holes (red arrow).

OTR screens are also mounted to the baseplate in such a way that hexapod positions exist that allow for the longitudinally travelling electron beam and/or laser to be intercepted, without collision with any of the other chamber constituents. The signal from the screens can then be imaged via the upstream or downstream alignment cameras—especially useful for the overlap of lasers and electron beam (see Section 5.1.4).

In addition to those for alignment, other cameras are also mounted around the chamber that provide images of the inside of the chamber with a wider field-of-view. Although generally useful as a way of checking the status and position of the chamber components, a number of these cameras prove particularly important in detecting light emitted from the plasmas. Two cameras, mounted at windows on top and at the front of the chamber, capture the light from plasmas generated where the two lasers and the electron beam cross. Another camera with a view from the side of the chamber collects light from plasmas generated either by the longitudinal laser or the discharge inside the capillary. The emission light captured by these cameras, especially after interaction with an electron beam, proves to be a valuable and versatile diagnostic in many parts of this work (see Chapters 4 and 5).

Also located outside of the chamber is a plasma emission spectroscopy diagnostic [66]. Here, light from the plasma is collimated by a lens and transported via optical fibre to the spectrometer, with which broadening of the H-alpha spectral line from ionised H_2 gas can be detected. Using H_2 as a dopant gas—a common gas mixture used in this work is 97:3 Ar:H_2—ultimately enables online plasma density measurements that can be temporally resolved via gating of the spectrometer camera.

3.2.5 Imaging electron spectrometers

Immediately after the interaction chamber—at the start of the diagnostic section (FLF-DIAG) —the electron beam is captured by five XFEL-type quadrupoles [181]. From here, the beam can either be transported onward to the beam dump (see Section 3.2.6) or imaged onto the screen of one of two vertically dispersive dipole electron spectrometers (see Figure 3.6) [182].

Imaging quadrupoles

The strength of the five quadrupoles immediately after the chamber can be set such that the transport matrix between a particular object plane (e.g. the end of the plasma cell) and an image plane (e.g. a spectrometer screen) fulfils the point-to-point imaging condition (see Section 2.2.2 on page 18). The first four quadrupoles in this part of the beamline come in two pairs, with each pair sharing a power supply. As such, the same current is supplied to the quadrupoles within a pair, and so the field values within those magnets are the same and cannot be changed independently. In total, this gives three degrees-of-freedom when performing point-to-point imaging.

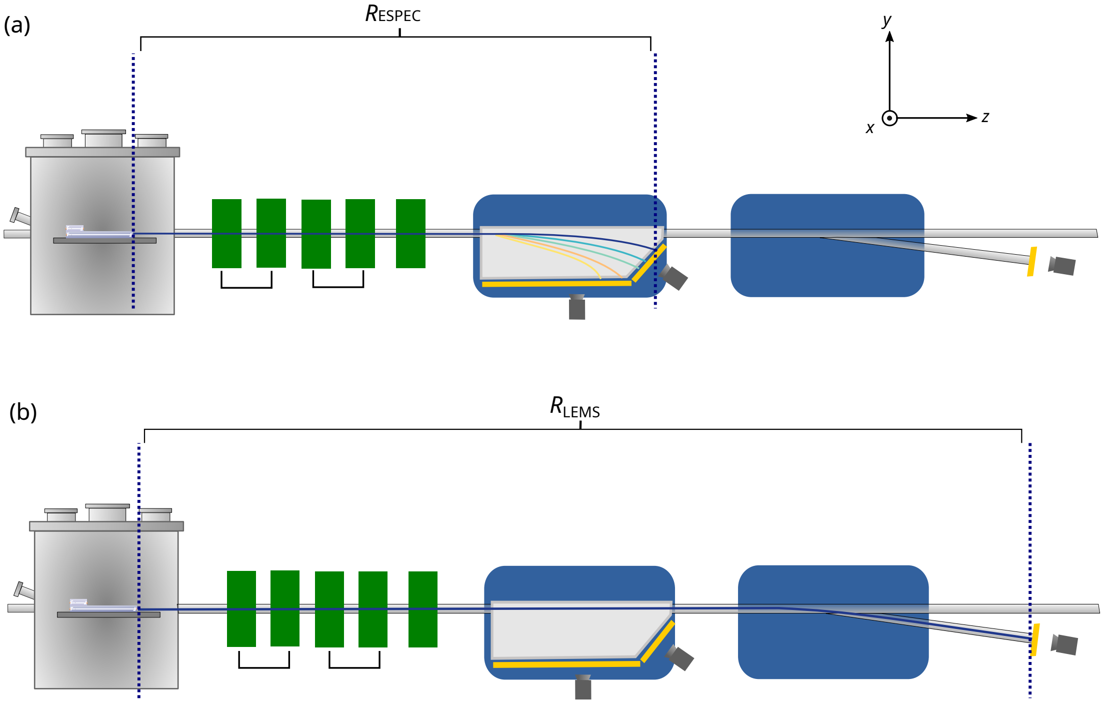

Figure 3.6: Imaging electron spectrometers at FLASHForward. Diagrams of the diagnostic beamline immediately after the plasma are shown with the beams imaged onto (a) the ESPEC and (b) the LEMS. The transport matrices R_{ESPEC} and R_{LEMS} point-to-point imaging at the ESPEC and LEMS respectively will depend on the energy of the transported beam and the location of the measured object plane. R_{ESPEC} additionally depends on the length of the trajectory through the dipole, which for a given field strength and beam energy can be estimated using a rectangular B-field model.

Usually, an optic is found that produces the desired object plane position in both directions, $s_{obj,x,y}$, and the magnification in the non-dispersive direction, M_x. Accounting for the magnification, the horizontal beam profile at a position $s_{obj,x}$ can be retrieved from the horizontal beam profile measured at the image plane. Crucially, this allows the measurement of electron beams 'in focus' i.e. with $s_{obj,x}$ located at the beam waist. Furthermore, scanning the position of $s_{obj,x}$ by systematically adjusting the quadrupole strengths allows the beam size at different positions away from the waist to be measured. Such an object plane scan, as described in Section 2.2.2, yields the beam divergence which, in combination with the waist size, constitutes a measurement of the beam emittance.

It should be noted however that, due to the natural chromaticity of the imaging quadrupoles, $s_{obj,x}$ and M_x are in general functions of the particle momenta. Strictly then, only a single energy within a beam can be properly imaged; particles with energies different to this consequently have different object planes. In practice, this means that only a finite range of energies within the beam can be measured properly at once; away from the 'imaging energy', the beam becomes out-of-focus, essentially decreasing the energy resolution. A solution to this is to measure the energy spectrum for more than one imaging energy—a method that necessarily requires multiple shots. An interesting benefit of this effect however is that for beams with a relatively large energy spread, the beam size for a range of object planes can be measured in a single-shot. As demonstrated in Section 5.2.5, with some assumptions about the properties of energy slices within a bunch, this method can in principle yield shot-to-shot measurements of the beam emittance [183].

Broadband electron spectrometer (ESPEC)

Often referred to simply as the ESPEC, the first spectrometer in the FLASHForward lattice consists of a 1 m-long dipole embedded within a stainless-steel vacuum chamber with walls 1 mm-thick. (see Figure 3.6 (a)).

Electron beams dispersed by this dipole then propagate through the lower walls of this chamber onto scintillating screens (Kodak LANEX fine), the underside of which are imaged by six CCD cameras (model Basler acA2040-35gm; note only two are shown in Figure 3.6). The broadband nature of this diagnostic makes it well-suited for diagnosing the spectra of beams with a large spread of energies, albeit with the constraints of the previously discussed chromatic imaging system.

Measurements in the horizontal beam profile are possible with this spectrometer, but are also limited by the relatively poor spatial resolution of the screens. The high energy screen provides the best resolution, where each pixel corresponds to a region 21 µm-wide at the imaging plane, but this is further degraded by the presence of scattering effects in the chamber wall. These scattering effects will generally depend on the beta function at the surface of the chamber wall and also on the beam energy, becoming more prominent for particles with lower momenta [184]. To some degree this resolution limit can be avoided by increasing the magnification, but this too is limited to what the imaging beamline can accommodate before the beam is apertured by the 10 mm diameter beam-pipe.

Low-Emittance Measurement Station (LEMS)

Located further downstream, another spectrometer, termed the Low-Emittance Measurement Station (LEMS), is designed to overcome the beam size measurement limits present in the ESPEC. Although employing a similar dipole to that in the ESPEC, this spectrometer only accepts particles deflected at an angle of $\sim 10\,\mathrm{deg}$, meaning only a narrow band of energies can be measured at the same time. However, the main feature of this diagnostic is the XFEL-type screen station located $\sim 1\,\mathrm{m}$ downstream of the deflection point [185]. Unimpeded by any scattering effects, an in-vacuum gadolinium aluminium gallium garnet (GAGG) scintillator imaged by a camera equipped with Scheimpflug optics yields a screen resolution measured to be as low as $7\,\mu\mathrm{m}$. This, in addition to the increased magnification that a longer drift length affords, makes this diagnostic ideal for the measurement of small beam sizes and therefore low emittances.

Energy and charge calibrations

In order to diagnose the spectra of the beams in both of the spectrometers, a calibration is needed that can retrieve the energy E of a particle from its screen hit position y and the field strength of the dipole B. Knowing that y is a function of the deflection angle θ and that $\theta \propto \frac{B}{E}$ (from the Larmor radius), it can be written that $y = f(\frac{B}{E})$, where f is a function to be determined. Ideally, f can be found by directly using the FLASH beam with a constant energy known from measurements available in the linac; the B-field is then varied by adjusting the current supplied to the dipole and the values of y measured for each setting. Then, for a measured value of y and B, E can be calculated in the future from $E = \frac{B}{f^{-1}(y)}$. However, it is not always the case that the FLASH beam can be used for this calibration; even for its lowest energy, it is difficult to deflect this beam down to the low energy screens. As an alternative, tracking of particles with varying energy through a simulated dipole can be performed and their y positions recorded on a virtual screen. This therefore allows f to be constructed in the same way, except with E as the parameter being varied.

As well as energy measurements, the spectrometer screens are used to measure the charge density of electron beams [108]. This involves correlating the light detected from the scintillating screen with the charge measured e.g. in an upstream BPM, which can for example be varied by using the collimator (see Section 3.2.2). Complications to this however arise due to the fact that the screen can become saturated, and also that the detected scintillation light can vary spatially across the screen for a fixed amount of charge. To avoid saturation, the simplest solution is to avoid high charge densities by slightly defocusing the beam; this is often only needed for the unscraped, non-interacted FLASH beam. The variation in light signal as a function of screen position however arises because of the non-uniform light collection of the screen-camera optics and also because of the dependence of the scintillation light on the hit angle of the particles incident on the screen. This effect can however be characterised by scanning a beam of known charge (e.g. the FLASH beam), across the screen by varying the dipole field strength, and measuring the light variation. Performing this for each camera yields a correction curve that can account for the combined effect of the camera optics and the hit angle of the electron beams.

3.2.6 Transverse Deflection Structure and beam dump

In the case where the ESPEC or LEMS are not active, the electron beam can be transported further down the beamline—through the section earmarked for the future installation of undulators—to the part of the lattice that hosts the PolariX X-band Transverse Deflecting Cavity (XTDS) [186].

The principle of operation within a TDS is that electrons receive a transverse kick depending on their longitudinal position in an RF field (see Figure 3.7). A bunch propagating on one of the zero crossings of the RF waveform becomes streaked in the polarization direction of the field, where the amount of deflection measured on a screen downstream of the TDS can be used to measure the positions of the particles in the time domain i.e. the bunch current-profile. In general, measurements are performed at the two zero crossings of the RF field and combined (e.g. by averaging); this also necessitates two sets of calibrations, which are found by measuring the centroid kick that the whole bunch receives when placed at different phases in the RF field.

As is the case here, the combination of a TDS with a dipole spectrometer dispersing in the non-streaked direction can then yield shot-to-shot tomography of the bunch longitudinal phase space—its most common application at FLASHForward. Operation and calibration of this spectrometer follows similar principles as for those described in Section 3.2.5. The screen on which the streaking and deflection are measured then constitutes the end of the FLASHForward beamline, where the beam is dumped.

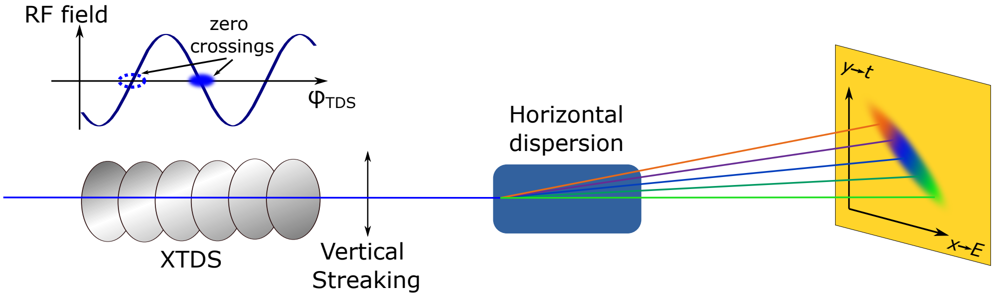

Figure 3.7: Longitudinal phase space measurements with an XTDS and spectrometer. Measurements are usually performed on the two zero crossings of the XTDS RF waveform. A calibration of the vertical deflection y versus the phase of the TDS RF ϕ_{TDS} can be found by shifting the RF phase relative to the bunch and measuring the centroid deflection on the screen. From measurements of the deflection of particles and knowledge of the RF frequency ($f_{TDS} = 11.9888\,\mathrm{GHz}$), the relative longitudinal positions of particles can be calculated. Horizontal dispersion by a dipole then allows the energy spectrum of a bunch to be measured. Additional quadrupoles (shown in Figure 3.1 (b)) allow for transport of beams to this diagnostic and also optimisation of the optics for high time and energy resolution.

Note that suitable optics can be a challenge to find for beams with a hard focus or indeed for those exiting the plasma stage. This is due to the fact that high-strength quadrupoles are required to capture and transport rapidly diverging beams along the relatively long propagation distance between the interaction point and the TDS; such beams can grow to large widths during capture, making propagation through the narrow beam pipe of 10 mm inner diameter a challenge. Furthermore, these issues can be compounded by the limited sensitivity of diagnostics such as BPMs to low charges, which is often the case for accelerated witness bunches. Although diagnosis of such beams is not strictly impossible with the TDS, measurements with this diagnostic most often concern those of the FLASH beam in the absence of plasma interaction, before a hard focus has been established in the interaction chamber.

3.3 FLASHForward laser

Situated in a laboratory above the FLASHForward experimental tunnel, a 25 TW laser system provides ultra-short laser pulses primarily intended for plasma generation in FLASHForward. The following section gives a brief overview of this laser system, before outlining how the pulses it generates are subsequently transported to the experimental hall.

3.3.1 Laser system backend

Figure 3.8 summarises the key sub-systems of the FLASHForward laser. Based on titanium sapphire (Ti:Sa) as the gain medium, the laser chain begins with the oscillator (Venteon — Pulse: One). Pumped by a CW (continous wave) 532 nm solid state laser, the Ti:Sa crystal is located in a cavity consisting of eight mirrors, as shown in the inset of Figure 3.8. Ultra-short, high-intensity pulses are favoured over continuous wave (CW) radiation by the principle of Kerr-lens mode-locking [187, 188]; the repetition-rate of these pulses is defined by the round-trip time of the cavity to be 108 MHz. Small movements of two of the mirrors here allow the cavity length to be adjusted such that the pulses from the oscillator can be synchronised with and delayed with respect to the 1.3 GHz RF signal in FLASH (see Section 3.1.5). In order to compare the signals from the oscillator and FLASH, a small amount of the laser is extracted from the oscillator and converted via a photo-diode detector to an electronic signal that can be mixed with the FLASH RF [174].

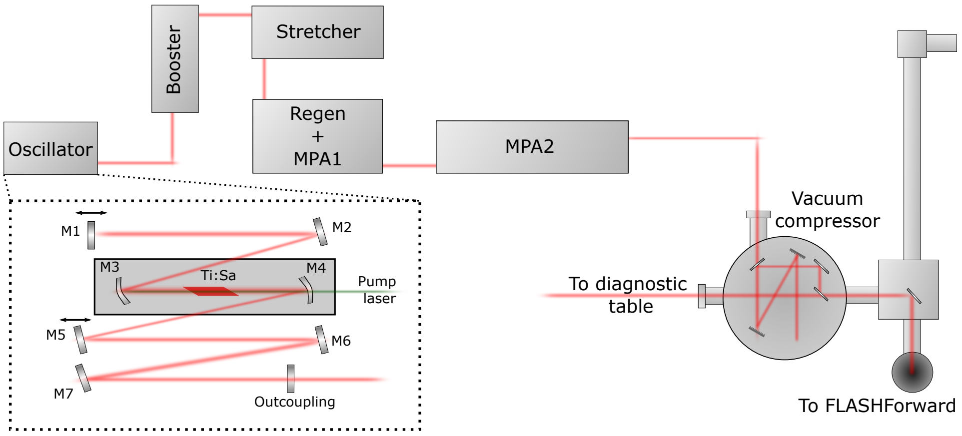

Figure 3.8: Ti:Sa laser system at FLASHForward. The approximate layout of the main stages in the laser chain are shown, with the beam exiting the laboratory down a vacuum pipe bound for the FLASHForward experimental hall. Shown inset are the key components of the oscillator system, the cavity of which consists of the mirrors M1-M7 plus the outcoupling mirror. Two of these mirrors, M1 and M5, can be adjusted via piezo-actuators to modify the length of the cavity and therefore provide synchronisation and controllable delay with respect to the FLASH RF signal.

Upon exiting the oscillator, the spectrally broad pulses are transported to the Booster unit—the first in a chain of multi-pass amplifiers themselves driven by flashlamp-pumped pump lasers. Here, the pulses are further amplified before the repetition rate is reduced to a maximum of 10 Hz by a combination of a Pockels cell and polarizing beam-splitter. Triggering of the Pockels cell is limited to a trigger frequency of 54 MHz, meaning that the arrival time of the down-selected pulse train can only be varied with a granularity of 18.46 ns; in combination with movements of the oscillator cavity mirrors, this pulse-picking completes the set of controls that one has over the laser—electron beam time-of-arrival delay, as summarised in Figure 3.9. After this pulse-picking, the remaining pulses are subsequently focused into a saturable absorber in order to remove pre-pulses and continuous background radiation, thus improving their contrast.

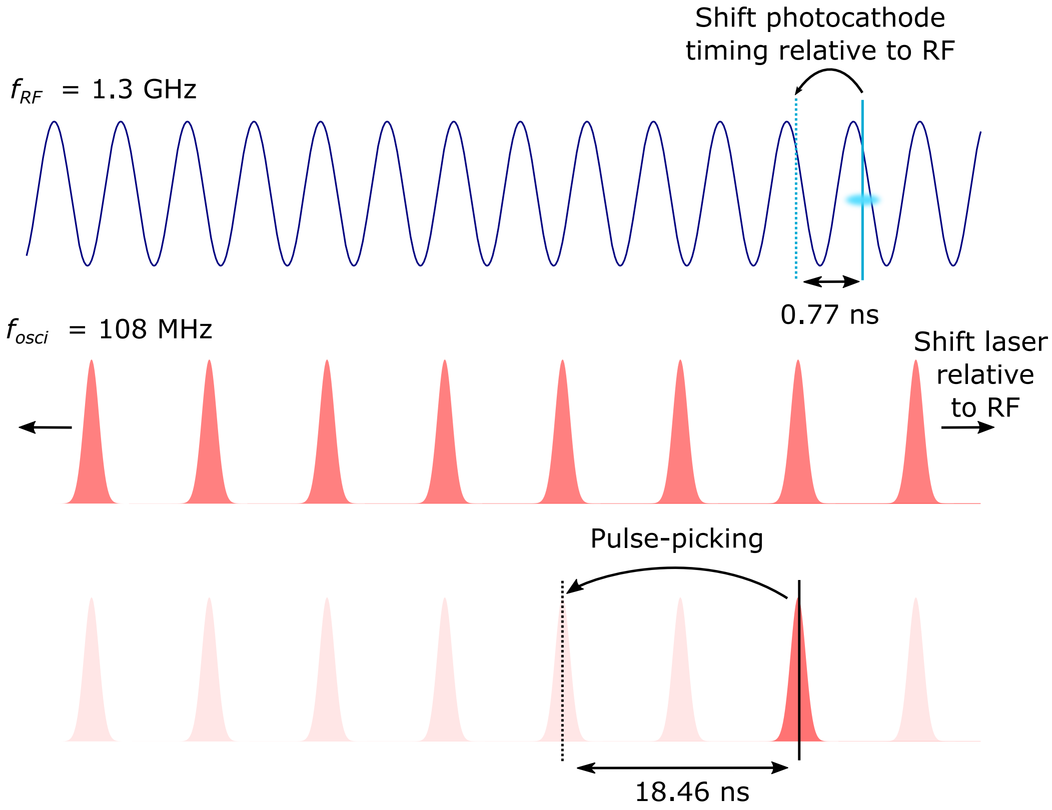

Figure 3.9: Laser—electron beam timing at FLASHForward. The electron beam (blue) time-of-arrival is defined by which acceleration bucket of the RF field (navy sinusoid) it is injected into and can be adjusted by delaying the photocathode laser in steps of 0.77 ns (see Section 3.1.3). Once synchronised to the RF signal, the laser pulse-train timing in the oscillator can be shifted with respect to the RF via movement of the cavity mirrors (see Figure 3.8). Finally, pulses from the laser pulse train can be selected with a granularity of 18.46 ns at the Booster stage.

Before further amplification can take place, the pulses must first be temporally stretched—the first step in the process of Chirped Pulse Amplification [73]. The Stretcher stage of the laser system in this case deploys an Öffner triplet scheme, providing aberration-free stretching of the laser pulse [189, 190]. Located immediately after the stretcher is also a programmable acousto-optic filter—commonly known as a Dazzler—that allows for precompensation of the group delay dispersion in the subsequent stages of the laser chain [191]. This therefore enables shaping of the spectral phase which, after the compressor, translates into modifications of the pulses temporal profile.

After stretching, the pulses are then amplified in a regenerative amplifier (Regen), followed by two multipass amplifiers (MPA1 and MPA2). This results in a still spectrally broad pulse, but with a nominal energy at the exit of second multipass amplifier of 925 mJ. Note that an adjustable attenuator within this final amplification stage allows the output energy of the laser to be varied continuously (from $\sim$ 0–100 %), whilst another switchable attenuator directly after MPA2 can be used to reduce the laser energy by a factor of 170, putting the laser into a low power alignment mode.

Post-amplification, the pulses then reach the in-vacuum compression stage. Here, the pulses are compressed down to lengths of < 25 fs rms via a grating pair. The main beam propagates in-vacuum to a switch mirror, to be transported either to the BOND (Beam Optimisation and Novel Diagnostics) lab—where mainly LWFA experiments are performed [192, 193]—or to the FLASHForward tunnel.

A small fraction of the beam is however transported to a diagnostic table equipped with, for example, a Wizzler for measuring the laser pulse duration [194, 195]. Based on the principle of self-referenced spectral interferometry (SRSI), the Wizzler works by interfering the main pulse with a reference pulse generated by Cross-Polarized Wave Generation (XPW) [196]. From the interference signal, the pulse's spectral phase and amplitude can be retrieved and its temporal profile therefore fully characterised—all in a single shot. As well as providing pulse duration measurements directly after the compressor, the Wizzler could also be relocated for use closer to the experiment—for example in the FLASHForward tunnel (see Section 5.1.2)—where any dispersion experienced along the beamline will result in a different temporal profile than that measured in the laser lab.

3.3.2 Beam Transport

With the switch mirror inserted, laser light is reflected into the beamline for transport down to FLASHForward (see Figure 3.10). The first three mirrors in this transport line bring the laser down to the level of the FLASHForward experimental hall. Housed in the so-called 'multi-chamber', the next 3 mirrors represent the first that are located in the tunnel. A pipe then connects this chamber to the 'incoupling chamber' where, crucially, a retractable 0.5-inch pick-off mirror diverts the centre of the 34 mm-diameter main beam into a beamline that constitutes the transverse injection laser arm. As described in the next section, the rest of the beam propagates onward for coupling into the FLASHForward electron beamline.

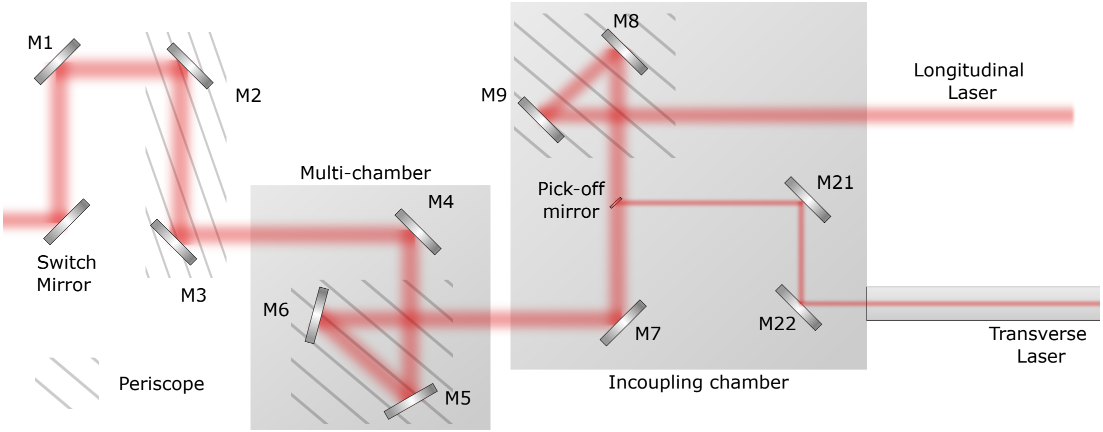

Figure 3.10: Laser transport from the laser lab to the FLASHForward experimental hall. To facilitate the reproducible alignment of the beam through this long beamline, the majority of the mirrors shown are equipped with pointing diagnostic cameras and are also pico-motorised that enable remote steering of the beam back to established ROIs.

3.3.3 Preionisation (longitudinal) laser arm

After two periscope mirrors bring the main beam down to the same vertical level as the electron beam, the laser then enters the focusing chamber (see Figure 3.11). Here, a pair of 750 mm focal length spherical mirrors are arranged to produce an effective focal length of 18 m and are equipped with pico-motors such that pointing adjustments to the longitudinal laser can be made. After reflection from these mirrors, the beam is coupled into the FLASHForward beamline through a window at the beginning of the FLFMAFF section, travelling colinearly with the electron beam path and reaching a wide-diameter focus at the centre of the plasma chamber—see Section 5.1.2 for focal spot measurements of this laser in the context of the plasma cathode experiment. The polarisation of this laser as it travels through the chamber is vertically linear.

After traversing the length of the chamber, the longitudinal laser propagates through the post-interaction section of FLASHForward and is coupled out of the beamline by an insertable silver mirror; an optional mirror with a hole in the centre is also available to allow the electron beam to pass further downstream. Adjacent to this outcoupling station is a breadboard housing near- and far-field diagnostics, with which the transmission of this laser through the beamline can be monitored.

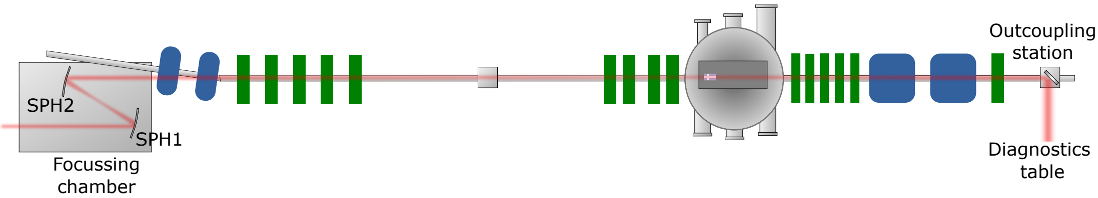

Figure 3.11: Longitudinal laser beamline. After focusing by the two spherical mirrors SPH1 and SPH2, the longitudinal laser enters the FLASH vacuum at the start of the FLFMAFF section. In the configuration shown, the laser propagates through the gas cell and is outcoupled at the start of the FLFUND section. Note that a valve at the entrance of the focusing chamber can be closed to block this laser arm independently of the injection laser.

3.3.4 Injection (transverse) laser arm

After formation at the pick-off mirror, the transverse laser beam reflects off two more mirrors in the in-coupling chamber before propagating along a vacuum pipe and out through an anti-reflection coated window. A delay stage then contains the first two mirrors of the in-air section of the transverse laser beamline, depicted in Figure 3.12, that allow the transverse laser time-of-arrival to be adjusted on the picosecond level relative to the longitudinal laser and electron beam. Following this, a periscope brings the beam down onto a breadboard at approximately the same height as the electron beam in the FLASHForward interaction chamber. After reflection from a high-reflectivity mirror, the beam can optionally be blocked by a beam block mounted on a remotely-controllable flipper. Mirrors then transport the laser beam across the breadboard to a set of motorised optics mounted on a stage that, by virtue of stepper motors, can translate in the $x - z$ plane. The mirrors M212 and M213 of this stage enable precise alignment of the beam onto a gold-coated $f/51$ off-axis parabola (OAP). In addition to actuators that provide angular adjustments to this mirror, a stepper motor was installed on its mount that enabled reproducible vertical translations of the beam's trajectory through the interaction chamber. This is in addition to the pointing control over the beam that is granted by the motorised turning mirror (TM) after the OAP, which turns the beam onto a path that is perpendicular to that of the longitudinal laser and electron beam. Finally, the transverse laser propagates through a

2.4 mm-thick uncoated fused silica window and nominally reaches focus within the interaction chamber at a location perpendicularly intercepting the path of the longitudinal laser and electron beam; the polarisation of this laser as it propagates through the chamber is linear in the vertical direction. The beam then travels out of a window on the other side of the chamber for diagnosis on the upstream alignment camera (see Figure 3.5). Again, characterisation of this laser's focal spot for typical settings in the plasma cathode experiment are presented in Section 5.1.2.

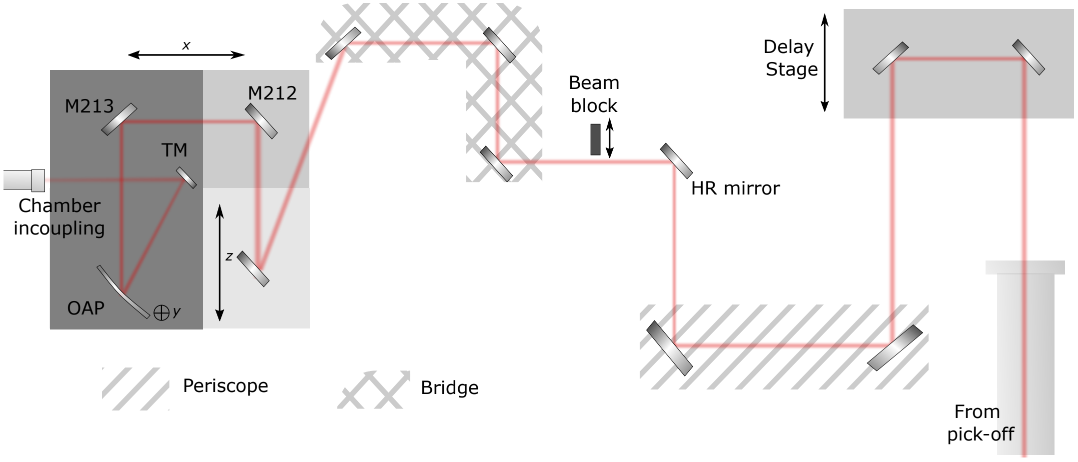

Figure 3.12: Transverse laser beamline. The path of the injection laser in-air across the breadboard located next to the interaction chamber is shown. A large amount of the components can be remotely controlled—most importantly the delay stage, beam block, and the final translation stage with the mirrors that are labelled.

3.4 Experiment Overview

Core scientific activities at FLASHForward can be separated into 3 main experiments:

- X-1: An internal injection experiment aiming to demonstrate a PWFA-based plasma cathode stage that acts as a *brightness booster* for the FLASH drive beam [197].
- X-2: An external injection experiment with the goal of demonstrating a high total efficiency, beam-quality preserving *energy booster* PWFA stage for the existing FLASH beam [139, 182, 198].
- X-3: A high-repetition-rate experiment aiming to exploit the unique bunch-train structure of FLASH to demonstrate *high-average-power* PWFA [153, 199].

In addition to these main research topics, a number of experiments also exist that support the core projects whilst also developing enabling technologies for the future (e.g. plasma-based dechirping [200] , advanced beam diagnostics [186], plasma target characterisation and development [66] etc.) A number of experiments from external collaborators were also proposed, including the Strathclyde-led X-18 experiment focusing on using the emission-light from a plasma afterglow as a diagnostic [201]; this found use as a spatio-temporal alignment diagnostic in the X-1 experiment (see Chapter 5). It is then also linked to the work in Chapter 4 on the X-2 experiment, where the emission-light signal provides a new method of measuring the energy-transfer efficiency of a PWFA stage.

Chapter 4

Longitudinally resolved efficiency measurements in a PWFA stage

In order to become viable energy boosters, plasma accelerator stages must have, amongst other qualities, high overall energy-transfer efficiency. Section 2.4.3 outlined the different factors that are combined to give the overall efficiency of the plasma acceleration process. Experimentally, the most focus has been placed on the energy-transfer efficiency from the driver-generated wake to the accelerated witness bunch. In PWFA, milestone results have demonstrated energy-transfer efficiencies as high as $(42 \pm 4)\,\%$ whilst simultaneously preserving beam qualities such as charge and energy spread [139, 202, 203].

Traditionally, the energy-transfer efficiency is calculated by comparing the respective energy loss and gain of the driver and trailing bunch, measured with a magnetic dipole spectrometer. This however requires the stable transport of the beam to the spectrometer followed by slow, multi-shot measurements that ultimately leave both bunches scattered by the diagnostic screen. Furthermore, the efficiency measured in this way is longitudinally averaged, yielding no information about how the acceleration process evolves within the plasma and leaving measurements subject to systematic errors from imperfect post-plasma transport.

Presented in this chapter is an energy-transfer efficiency diagnostic based on the light emitted from a plasma after beam-interaction, the key results of which are also summarised in Ref. [204]. Previously, it has been shown that an increase in this plasma emission-light—also known as plasma afterglow light—can be related to the driver-to-wake energy deposition [205], and was more recently deployed in the spatio-temporal alignment of particle beams with laser-generated plasma columns [201, 206] (see also Section 5.1.4 for an example of the latter use-case). Although the precise mechanisms that result in this change in emission-light are not yet fully understood, the light itself is a simple and robust observable that can be used for non-invasive measurements.

Using a high-efficiency working point established in the X-2 external injection experiment at FLASHForward as a platform, the change in plasma emission-light was here studied as a function of both energy deposition and extraction to and from the plasma wake. From this, the use of the emission light as a measure of wake energy was demonstrated in calculating energy-transfer efficiency. After comparison with the same efficiency measured on the ESPEC, the novel capabilities of this new diagnostic are then explored—namely that it can provide non-invasive, longitudinally resolved measurements of the acceleration process throughout a plasma accelerator stage.

4.1 Experimental Setup

Figure 4.1 summarises the setup of the energy booster external injection stage in the X-2 project. For this experiment, plasma is generated by firing discharge pulses through neutral gas flowing in one of the capillaries; in this case, the 195 mm-long cell was used. The electron beam, with a current-profile tailored via collimation for optimal double-bunch acceleration, is then focused into the plasma, with the resulting energy loss and gain of the driver-trailing bunch pair then diagnosed on the downstream ESPEC. Simultaneously, and key to these studies, the plasma sideview camera was able to capture the light from the beam-interacted plasma—a signal that, as will be shown, gives valuable insight into the acceleration process. In the following sections, further details are provided regarding the setup of the electron beam, the discharge plasma and the vital sideview plasma-light diagnostic in this experiment.

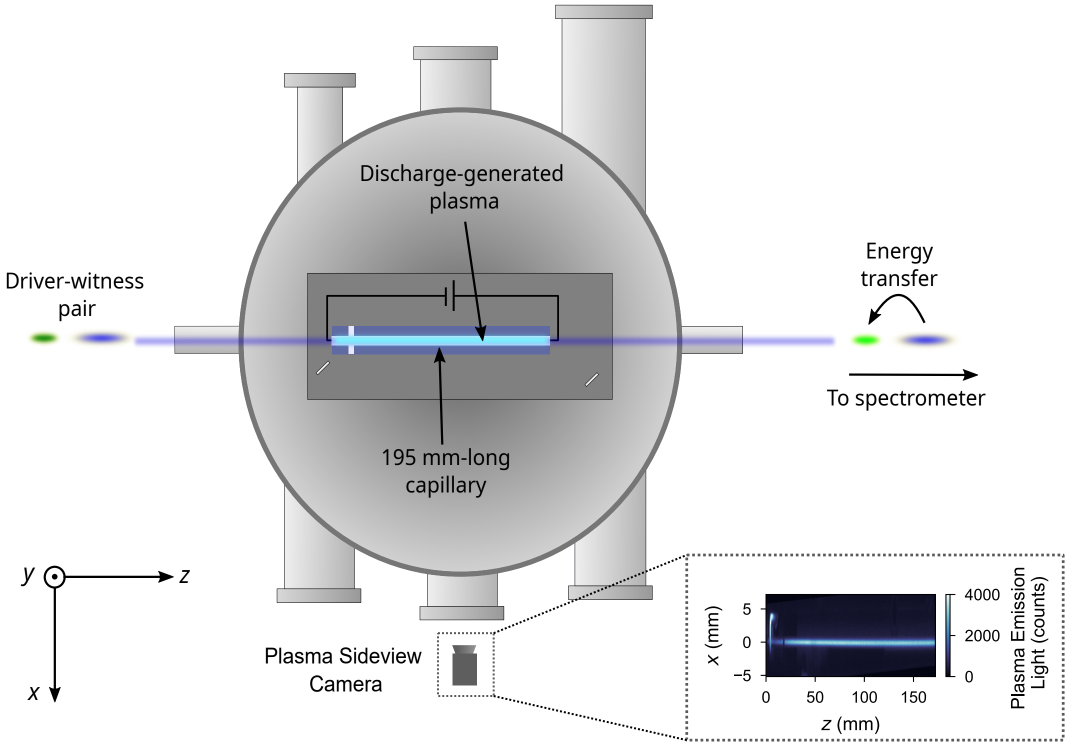

Figure 4.1: External injection PWFA stage setup. An electron beam with a tailored, double-bunch profile is coupled into a discharge-plasma, generated in this case in the 195 mm-long capillary. The energy loss and gain of the driver and witness bunches is then diagnosed downstream on the ESPEC diagnostic. Simultaneous to these measurements, the emission-light from the beam-interacted plasma was detected by a camera observing the side of the plasma cell; an example image from this diagnostic is shown in an inset.

4.1.1 Electron beam setup and characterisation

Dedicated FLASH ARD shifts offered the external injection experiments at FLASHForward the most flexibility in terms of beam parameters. Typical of most external-injection experimental runs, a central beam energy of 1052 MeV with an unscraped bunch charge of (871 ± 3) pC was used for the measurements presented in this chapter.

Ultimately, the optimal electron beam configuration in an external injection experiment was determined by optimising the beam-plasma interaction in-situ. While high acceleration gradients ($\sim$ GVm^{-1}) are important, equally paramount was the stable coupling of trailing-bunch charge into the wake and the subsequent preservation of key bunch parameters such as energy spread—factors that also depend on the plasma density (see Section 4.1.2). Electron beam preparation prior to beam-plasma interaction measurements therefore aimed to find and characterise a suitable base operating point, from which the acceleration process could be iteratively optimised. The following sections illustrate how the longitudinal phase space of the beam was manipulated to generate a tunable double-bunch profile, before discussing how the resulting bunches were focused and straightened for coupling into the plasma stage.

Longitudinal phase space

A key feature of plasma-wakefield experiments relying on external-injection at FLASHForward is the need for a longitudinal phase space that can be manipulated via collimation in the dispersive section (c.f.. the internal injection experiment in Chapter 5). Such a longitudinal phase space is required to have an approximately linear correlation between the energy of the particles and their longitudinal position. While this can be provided via tuning of the phases and amplitudes in the FLASH linac RF cavities (see Section 3.1), such a prerequisite essentially restricts how strongly the bunch can be compressed. External injection experiments using this form of double-bunch generation therefore necessarily involve a trade-off between tunability of the beam's longitudinal phase space and the maximum current that can be achieved.

Figure 4.2 shows example TDS measurements of the linearly chirped longitudinal phase space of the beam in this experiment and illustrates how this can be shaped to produce a double-bunch current-profile. An example of the full, unscraped bunch longitudinal phase space is displayed in Figure 4.2 (a), with the shot-averaged current-profile at the same working point shown in Figure 4.2 (b). By moving collimators into the dispersive section of the FLASHForward beamline, the current-profile of the beam can then be adjusted (see Section 3.2.2).

Figures 4.2 (c–e) demonstrate how movement of the head, wedge and tail collimators independently lead to the removal of charge from the front, middle and back of the beam current-profile, respectively. Of the three, the wedge collimator is particularly important to this experiment as it is responsible for generating the double-bunch structure required for external injection. Whilst movement of this collimator as in Figure 4.2 (d) can be used to adjust the position of the notch in the current-profile, insertion of this wedge by different amounts can also adjust the notch width (not shown), allowing for flexible definition of the charge and spacing of the bunches. In addition to this, the head and tail collimators can then be used to modify the current-profiles of the driver and witness bunches independently.

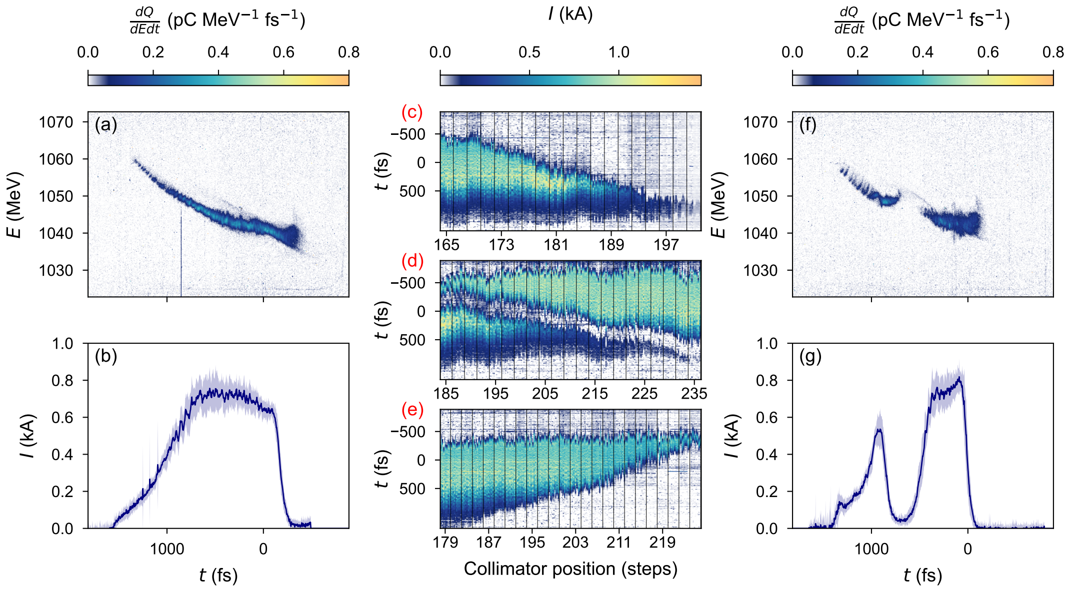

Figure 4.2: Double-bunch current-profile generation in the external-injection experiment. (a) An example of the full, unscraped longitudinal phase space of the beam measured on the TDS and (b) its corresponding shot-averaged current-profile are shown. Via movement of the (c) head , (d) wedge and (e) tail collimators, charge can be removed from different parts of the current-profile. Using a combination of these collimators, a driver-witness pair can be generated; examples of the (f) longitudinal phase space and (g) shot-averaged current-profile measured for a double-bunch working point are shown.

Figure 4.2 (f) shows an example longitudinal phase space measurement of a double-bunch configuration, with a corresponding shot-averaged current-profile depicted in Figure 4.2 (g). Using such a working point as a basis, the bunch profiles can then be further adjusted in order to optimise the acceleration process in the plasma. In general, the tuning of the trailing-bunch current-profile—primarily via the notch collimator width and position—is particularly vital in achieving optimal beam loading and therefore high-efficiency, energy-spread-preserving acceleration (see Section 2.4.3). The optimal configuration is ultimately determined by the shape of the plasma wake, which can be primarily controlled via the plasma density; as well as achieving beam loading, current-profile tailoring was important in minimising the loss of excess charge at the back of the wake. For this, the tail collimator proved extremely useful, whilst also allowing for the sequential introduction/removal of driver and trailing-bunch charge to systematically study the process of energy deposition/extraction to and from the wake.

Tilt correction and focusing

Whilst collimation provides a means of *longitudinally* matching the beam to the length of the plasma wake, the final focusing quadrupoles before the plasma cell enable matching in the *transverse* plane. This involves matching in the traditional sense i.e. coupling the beam into the plasma with the correct width and divergence (see Section 2.4.3), but also cancelling beam tilts. In both of these cases, the aim is to minimise any transverse oscillations that may lead to charge loss and emittance growth—two effects that can both be observed in the plasma-interacted spectra on the ESPEC.

In terms of beam tilt, the optimal conditions are clear; the beam should be as straight as possible i.e. no correlation between the longitudinal and transverse coordinates. This can be controlled via adjustment of the Q23FLFCOMP quadrupole strength in the FLFCOMP section (see Section 3.2.2) which can produce a small amount of 'leaking' lattice dispersion that can be used to cancel beam tilts [207]. Whilst the optimal settings can only truly be found with plasma interaction, preliminary beam straightening can nevertheless be performed beforehand using the TDS and measuring the beam on a screen before the dispersive dipole.

Typically after the completion of measurements on the TDS, a hard focus is then established at the entrance of the plasma cell. Figure 4.3 shows measurements of the approximate sizes and positions of the beta function waists in both the horizontal and vertical directions, made using the two-BPM-tomography method. As explained in Section 3.2.3, an optic was generally found that allowed for the independent modification of the size and positions of the beta function waists through adjustments of certain quadrupole strengths in the FLFMAFF beamline. Beta functions of waist size $\beta_0 = \mathcal{O}(10\,\mathrm{mm})$ were typically used in external injection experiments; the exact width and location of the optimal beam waist depends on the plasma density profile at the entrance of the cell. For example, at a typical plasma density of $1 \times 10^{16}\,\mathrm{cm}^{-3}$, the matched beta function $\beta_m = 3.3\,\mathrm{mm}$. Owing to the presence of the plasma upramp however, the optimal beta function will be somewhat larger than this. Due to their different energies, the effects of chromaticity will also mean that the focal plane of the driver and trailing bunch will be slightly different. In principle, by using the head and tail collimators, the beta functions of the driver and witness can be characterised independently by blocking the other bunch. Although such measurements can be instructive, it is again typically more efficient to optimise the focusing in-situ during acceleration studies; ultimately, it is the final properties of the accelerated bunch that are important.

Finally, more rigorous measurements of the waist at the interaction point can be provided by object plane scans on the spectrometers, yielding both the beta function and also the emittance of the beams, albeit only in the horizontal plane. Typically, the emittance of the incoming trailing-bunch is found to be $\epsilon_{x,n} \sim 1$–$2\,\mathrm{mm} \cdot \mathrm{mrad}$; due to CSR effects, the emittance of the driver is estimated to be larger. Although not vital in these studies, such measurements are crucial when considering issues such as emittance preservation.

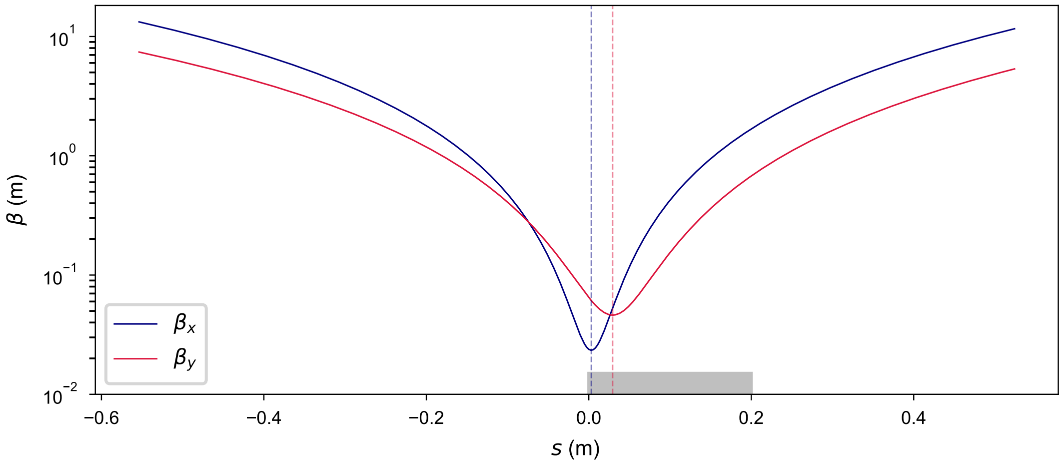

Figure 4.3: Beta function setup at the plasma cell location in the external injection experiment. The two-BPM tomography method was used to establish a tight focus at the plasma entrance. In this case, an asymmetry between the horizontal and vertical waists was measured, with $\beta_{x,0} = 23\,\mathrm{mm}$ and $\beta_{y,0} = 46\,\mathrm{mm}$. In addition, whilst the waist in x is situated at the entrance of the plasma cell (grey rectangle) the waist location in the y direction lies further downstream by $\sim 25\,\mathrm{mm}$.

4.1.2 Discharge plasma generation

Operation of the 195 mm-long capillary is in many ways the same as that of its shorter counterpart. Here the long cell was filled at a backing pressure of 30 mbar with the usual 97:3 Ar:H_2 mixture, corresponding to a pressure in the capillary of 9 mbar. Plasma was then generated by discharging 15 kV pulses through the neutral gas—from measurements of an argon discharge-plasma in a 33 mm-long cell, it can be expected that the first and potentially some of the second level of Ar is ionised [66]. By waiting a certain amount of time for the plasma to decay via recombination and expulsion from the cell, the plasma density experienced by the beam could be finely controlled.

The plasma density in the long cell for this working point was found using H-alpha spectral-line-broadening measurements with the optical spectrometer [66]. In this case, it was determined that a delay in the beam-arrival-time with respect to the initial discharge pulse of 9.64 µs provided the optimal plasma density with which the trailing bunch could be matched and accelerated with reasonable stability. In the central flat-top region of the plasma, this corresponds to a plasma density of approximately $7 \times 10^{15}\,\mathrm{cm}^{-3}$. At the cell entrance and exit, allowing for some plasma evolution will also change the shape of the density up- and downramps, affecting how the bunches are coupled into and extracted from the plasma stage. Note that while use of the long cell naturally affords higher driver-to-witness energy-transfer efficiencies, it also in general increases the sensitivity of the whole process to the input parameters, thus making the stable acceleration of trailing bunches with a high degree of charge coupling more challenging.

4.1.3 Plasma emission-light detection

Light emitted from the plasma in the long cell was detected by the plasma sideview diagnostic: a CMOS camera (model Basler acA2040-25gmNIR) equipped with a 1-inch, f/0.95 lens that provided a spatial resolution of 95 µm/px when imaging the plane of the capillary. In this experiment, we are purely interested in the plasma emission-light integrated over the whole, or part, of the camera field-of-view—a value labelled as I_p and measured simply in units of counts.

The spatial distribution of the measured light signal will not only depend on the plasma distribution at the time of emission; in general, the light collection efficiency varies across the camera lens and will also depend on transmission of light through the capillary walls. To account for these effects, measurements of the emission-light from a plasma immediately after generation with the discharge were made and are presented in Figure 4.4. By measuring the signal at these early times (< 0.1 µs), the plasma will have had no opportunity to evolve (see Section 2.3.2) and so should inherit the shape of the neutral gas profile within the cell, which will be approximately uniform. This was the case in a 33 mm-long cell [see Ref. [66], Figure 8 (a)], where for a measurement at a time 0.1 µs after the start of the discharge, the plasma density is approximately uniform in the centre of the cell, only dropping in the regions in and beyond the gas inlets which are not considered in this analysis (see red region-of-interest in Figure 4.4 (a)). Any spatial variation of the signal across the field-of-view is therefore assumed to be due to the light-detection method, which can then be characterised and accounted for.

Figure 4.4 (a) shows an example image of the emission-light taken immediately after plasma generation. From this image, one can immediately see that the field-of-view does not incorporate the full plasma cell—the end of the capillary beyond a longitudinal position of 175 mm was unfortunately not measured. Furthermore, the upper and lower regions of the image contain no signal. An example improvement to this diagnostic could therefore be optimisation of the field-of-view to observe the whole plasma length, whilst simultaneously better utilising the camera chip in the vertical direction e.g. with cylindrical lenses.

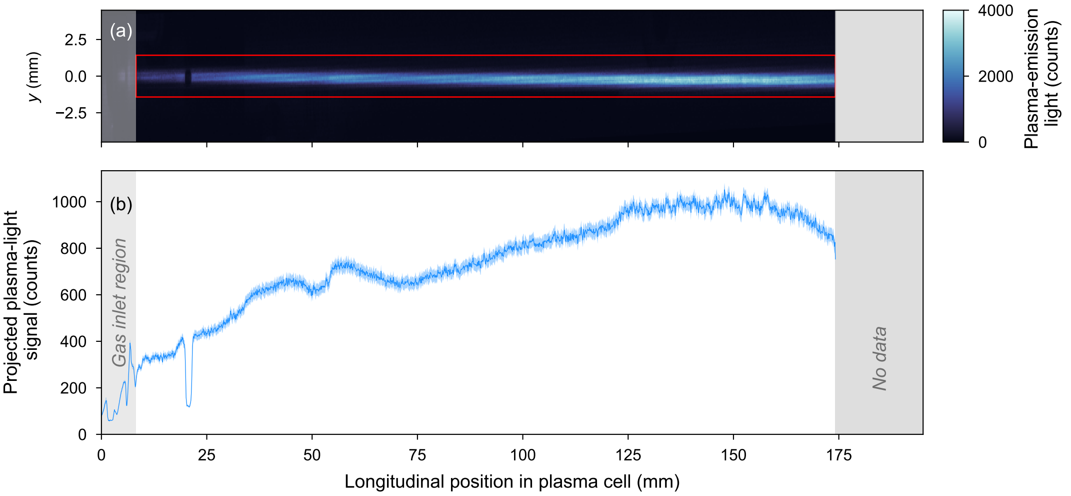

Figure 4.4: Longitudinal variation of plasma emission-light signal from a uniform plasma. (a) An image of the plasma light captured at the beginning of the discharge pulse and its (b) projection over the transverse extent of the capillary is shown. At times shortly after the discharge, the plasma density and therefore the emitted light should be approximately uniform along the capillary. Longitudinal variations in the signal are therefore assumed to be due to the details of transmission through the capillary and also the light-capture optics.

The variation of the signal along the length of the capillary is best quantified by projecting the signal over its transverse extent. Figure 4.4 (b) shows such a projection, averaged over multiple shots; the narrow error band around this mean curve indicates that the light signal emitted by this plasma was highly stable. However, whereas one would expect that the plasma density is uniform for most of the capillary and symmetric about its centre ($\sim$ 100 mm), the light yield in fact decreases towards the entrance of the cell. Furthermore, there are features that are at positions known to have abnormalities in the opacity of the cell walls ($\sim$ 50 mm) and also where the transverse hole is located (20 mm). Assuming that the shape in Figure 4.4 (b) arises purely as a result of the light detection, the curve itself forms an ‘imaging function’ that can be used to divide out any longitudinal dependence in the light collection efficiency;

this correction is therefore applied before any calculation of I_p. Note that, due to the relatively small diameter of the plasma channel, such effects are considered negligible for the transverse direction.

Of course, a measurement of I_p involves not only a spatial sum over part of the image—the signal is also integrated in time and wavelength. Formally, I_p is therefore a measure of the spectral irradiance integrated temporally over an exposure time Δt and spectrally over a bandwidth $\Delta\lambda$, before then summing over an area A of the camera chip i.e. $I_p = \int_{\Delta\lambda} \int_{\Delta t} \int_A \frac{\partial^3 I_p}{\partial\lambda \partial t \partial A}\, d\lambda\, dt\, dA$.

As sketched in Figure 4.5 (a), the exposure window began after the beam-arrival time and extended for a duration of $\Delta t = 24\,\mu\text{s}$. This was to avoid saturation of the camera chip, which could have equally been achieved via filtering; as will be shown in subsequent sections, beam-interaction with the plasma can substantially increase the emission-light yield, and saturation could not be fully avoided through a decrease in the camera gain. Temporally, the early parts of the plasma emission-light were therefore not captured by this diagnostic.

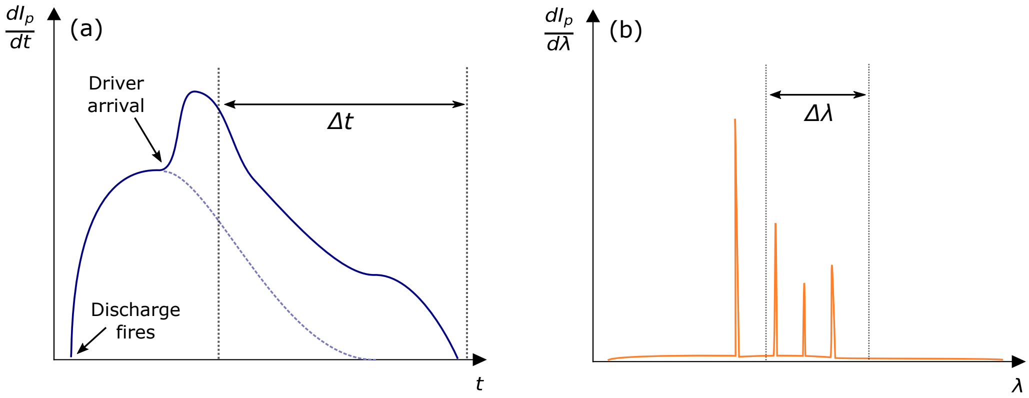

Figure 4.5: Temporal and spectral integration of the emission-light signal. Sketches of the variation in I_p with (a) time and (b) wavelength are shown. Illustrated are the ranges in t and λ that are integrated over. In time, the integration begins after the beam-arrival and lasts until after the plasma has relaxed; the difference between the dishcharge-plus-driver (navy solid line) and discharge-only (navy dashed line) curves integrated over the time window Δt is defined as the excess captured plasma emission light. In wavelength, only a finite region of the plasma-light spectrum is measured.

In a similar way, the spectrum of the emission-light—expected to consist of discrete wavelengths emitted during transitions between atomic states—was also not fully measured by this diagnostic [see Figure 4.5 (b)]. This is naturally due to the finite spectral response of the camera ($\sim$ 400 nm–1000 nm), in addition to any chromaticity of the light capture optics.

Ultimately for this experiment, knowledge of the plasma emission-light properties in the spectral and temporal domains was not crucial. Instead, it was simply required that the exposure window and bandwidth of the diagnostic remained unchanged throughout the experiment; as explained in Section 4.2.3, these parameters become part of a calibration that only holds for the experimental conditions it was measured under. It may be the case that an understanding of the temporal evolution (e.g. with a fast, intensified camera) or the spectral decomposition (e.g. with an optical spectrometer, see Ref. [205]) of the light can provide deeper insights into the long-term behaviour of the plasma. These, in addition to better use of the field-of-view of the sideview diagnostic, therefore make for possible improvements to future iterations of this diagnostic.

4.2 Experimental Results

Following the setup described in Section 4.1, the first step for these experimental studies was to establish a high-efficiency working point. The results of this, in addition to all of the results presented in this section, were taken during the same experimental run.

4.2.1 High-efficiency working point characterisation

A working point was found by tuning of the plasma density via the beam arrival time, as well as the notch scraper width and position. This was similar to the procedure outlined in Ref. [139], with more emphasis on the maximisation of the energy-transfer efficiency, which from now on will refer to the wake-to-witness energy transfer efficiency $\eta_{wake \to w}$.

Figure 4.6 demonstrates how η was calculated with an example shot from this optimised working point. In the absence of plasma (see Figure 4.6 (a)), the initial mean electron energies of the driver and witness are 1044 MeV and 1054 MeV respectively. Upon ignition of the plasma, the driver experiences energy loss whilst the trailing-bunch is accelerated to higher energies. The mean electron energy change of the accelerated and decelerated beams, ΔE_{acc} and ΔE_{dec}, can then be found as the difference between their mean electron energies with and without plasma interaction. Measurements with an electron spectrometer are however complicated by the fact that only a narrow bandwidth of energies can be imaged at a given time, with the energy resolution of the diagnostic decreasing away from the imaging energy. Accurate spectral characterisation of the decelerated and accelerated beams can therefore not occur simultaneously, meaning that separate measurements of the driver and witness bunches are required.

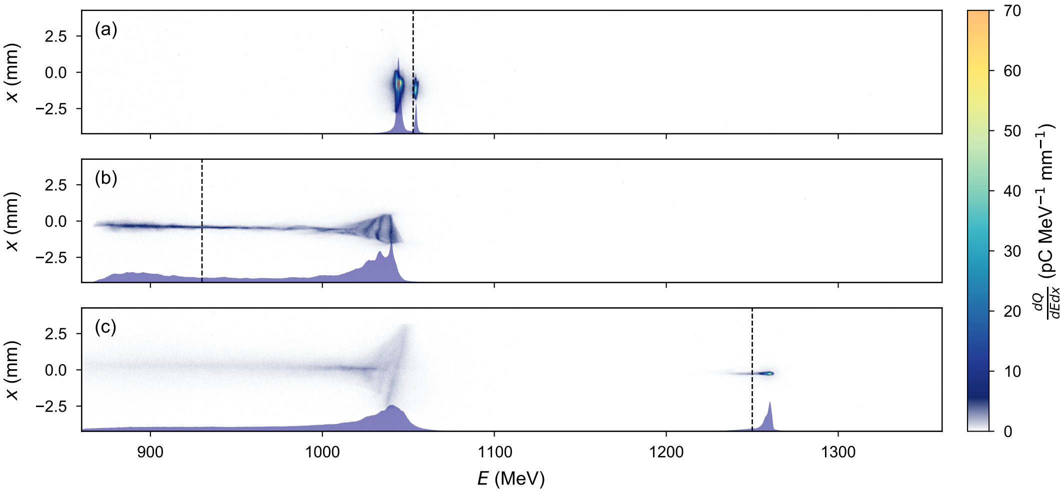

Figure 4.6: High energy-transfer efficiency measurements with the electron spectrometer. Shown are example spectrometer images of (a) the driver and witness in the absence of plasma, (b) the plasma-decelerated drive beam and (c) the accelerated trailing bunch accompanied by the unfocused driver spectrum. The relative spectral densities are also displayed at the bottom of each axis (navy shaded curves), with the black dashed vertical lines representing the imaging energy for each measurement.

Figure 4.6 (b) shows a spectrometer measurement of the decelerated drive beam, with the trailing bunch blocked by the tail collimator. Measured at an imaging energy of 930 MeV, the mean driver electron energy loss was found to be $\Delta E_{dec} = (69 \pm 3)$ MeV for this working point, corresponding to a longitudinally-averaged decelerating field of $(0.35 \pm 0.02)\,\mathrm{GV\,m^{-1}}$. By integrating the signal detected on the screen, the driver charge was calculated as $Q_{dec} = (368 \pm 8)$ pC. This was then combined with the change in driver mean electron energy to give a total driver energy loss of $\Delta W_{dec} = \Delta E_{dec} Q_{dec} = (25.3 \pm 1.4)$ mJ.

A similar procedure was then applied to find the total energy gain of the trailing bunch. As shown in Figure 4.6 (c), the imaging energy was set to 1250 MeV, providing accurate measurements of the witness spectrum whilst leaving the driver out of focus. Another issue with the spectrometer diagnostic is the inevitable loss of charge for bunches that are not perfectly transported through the imaging beamline. This is especially prevalent for the witness beam which, in the case of misalignment, is prone to being defocused or deflected by the strong defocusing fields behind or at the edge of the plasma wake—leading to an underestimate of the average total energy gain of the trailing bunches in the plasma. In order to reduce this systematic error, only shots that display near-perfect charge-coupling—defined here as having above 90% of the maximum measured witness charge—are considered. For the shots that fulfil this criterion, the trailing-bunch charge was measured to be $Q_{acc} = (46.7 \pm 1.4)\,\mathrm{pC}$, with a mean electron energy increase of $\Delta E_{acc} = (194 \pm 4)\,\mathrm{MeV}$ (i.e. an accelerating field of $(0.99 \pm 0.02)\,\mathrm{GV\,m^{-1}}$) and a resultant total energy gain of $\Delta W_{acc} = (9.1 \pm 0.3)\,\mathrm{mJ}$.

With a method of calculating the total energy change of the driver and trailing bunches, the energy-transfer efficiency measured with the spectrometer η_s can be calculated with the formula for $\eta_{wake \to w}$ in Equation 2.68. Assuming that the total driver energy loss remains constant at the value of ΔW_{dec} calculated previously, η_s can be estimated for every measured trailing bunch. Considering only those fulfilling the near-perfect charge-coupling condition, an average energy-transfer efficiency of $\eta_s = (36 \pm 2)\,\%$ was measured. Using the shot with the largest witness energy gain—displayed in Figure 4.6 (c)—the maximum energy transfer efficiency was calculated to be $\eta_{s,max} = (37 \pm 2)\,\%$, where the error remains from the uncertainty in the total driver energy loss.

Although comparable in terms of the average energy-transfer efficiency measured previously in the 50 mm-long capillary (c.f. $(42 \pm 2)\,\%$ in Ref. [139]), the results here are significant because of the factor of 4 higher witness energy gain, owed principally to the 4 times longer plasma cell. This improvement is better quantified by comparing the driver-to-witness energy transfer efficiency $\eta_{d \to w}$ in Equation 2.67, where the denominator is now the total energy of the incoming driver, $W_{dec,0}$. In the case of Ref. [139] with an incoming drive beam of mean electron energy 1050 MeV and total charge $(490 \pm 10)\,\mathrm{pC}$, $W_{dec,0} = 490\,\mathrm{mJ}$, with a total energy extracted by the trailing bunch of $\Delta W_{acc} = (98 \pm 9)\,\mathrm{pC} \times (45.4 \pm 1.4)\,\mathrm{MeV} = 4.4\,\mathrm{mJ}$ and thus a driver—witness energy-transfer efficiency of $\eta_{d->w} = 0.93\,\%$. Although the total incoming driver energy was lower in this experiment—$W_{dec,0} = 387.1\,\mathrm{mJ}$—the accelerated witness charge is also lower by a factor of ~ 2. Despite this, the larger energy gain of the witness in the long cell results in more of the initial energy of the driver being used for acceleration, with $\eta_{d->w} = 2.3\,\%$. This therefore highlights one of the directions for high overall efficiency acceleration—maximisation of the driver energy loss whilst accelerating a large amount of witness charge to greater energies.

4.2.2 Plasma emission-light dependence on driver and witness charge

The high-efficiency working point established in Section 4.2.1 subsequently provided a strong basis for measurements of the change in the plasma emission-light yield as a function of the incoming driver and trailing-bunch charges, which were varied using the tail collimator in the dispersive section of the beamline.

Plasma-light yield versus driver charge

Figure 4.7 presents the results of a tail collimator position scan, measuring the spectra of the driver on the spectrometer whilst simultaneously observing the plasma light signal with the sideview camera.

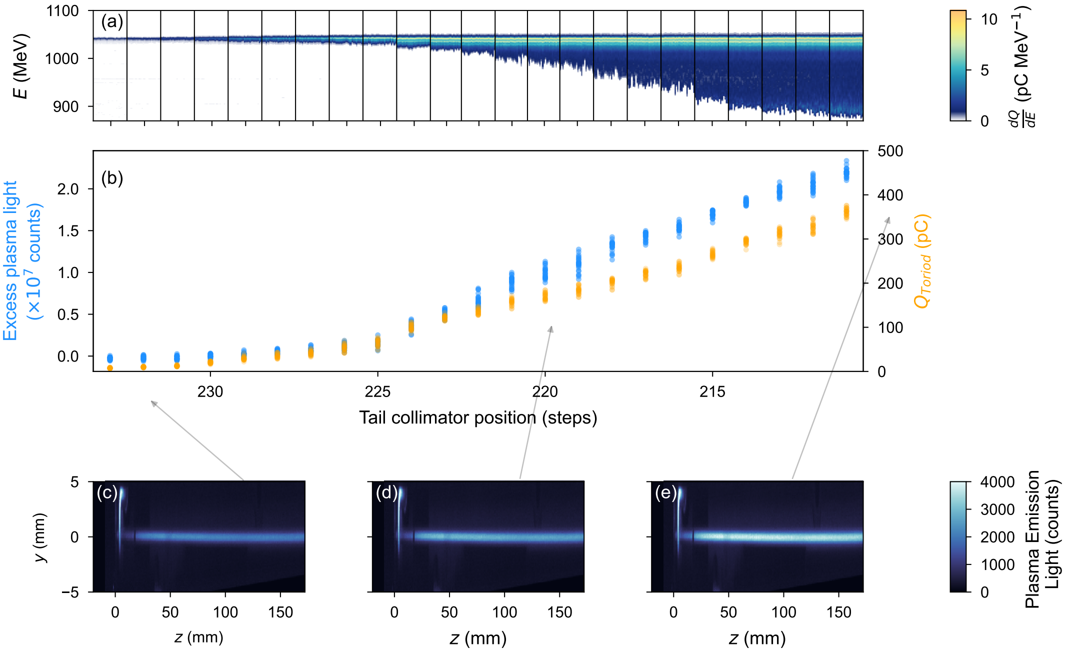

Figure 4.7: Plasma emission-light variation with increasing driver charge. (a) Waterfall plot of decelerated driver spectra, grouped by the position of the tail collimator. (b) Excess plasma light yield (blue points) and the driver charge detected in the upstream toroid Q_{toroid} (orange dots) versus the position of the collimator. Also shown are example images of the plasma light signal observed by the sideview with (c) negligible, (d) partial and (e) full charge in the drive beam, with arrows indicating their corresponding shots in (b).

Measurements began with the collimator fully inserted, such that no witness beam and only negligible amounts of driver charge reached the plasma (see Figure 4.7 (a)). Gradually extracting the tail collimator, longitudinal slices of driver charge were progressively introduced to the plasma, each adding to the energy deposited in the generated wakefield. As shown in Figure 4.7 (b), the excess captured plasma light I_p, defined as the difference between the integrated light signal measured with the side-view camera relative to that with only the discharge firing, increases with the driver charge. Figures 4.7 (c)–(e) show example images of the light signal with (c) negligible, (d) partial and (e) full driver charge reaching the plasma. The higher plasma emission light yield is essentially a manifestation of an increase in the energy contained in the wake; note however that the exact physical mechanism resulting in the extra light is not yet well understood. Once the beam has passed, the energy remaining in the wakefield will eventually be completely converted to the kinetic energy of the plasma electrons [208] that is then dissipated into the bulk plasma [165, 199]. Presently, this extra energy is assumed to be an additional driver of the dynamics described in Ref. [64], leading to more interactions between the plasma constituents and remaining neutrals which ultimately results in more emitted plasma light. One could indeed aim to use the same modelling and measurement techniques as detailed in Ref. [64], in this case including the additional source of energy that is provided by the drive beam interaction. This however was out of the scope of this work—as illustrated by the following sections, the relative change of this plasma light signal can still prove useful, even in the absence of a detailed understanding of the full plasma evolution.

Plasma-light yield versus witness charge

Further extracting the tail collimator, trailing bunch charge was then gradually introduced to the accelerating region of the wake. Unlike in the case of the driver, the slices of witness charge are located at approximately the same phase in the wakefield and so sample approximately the same accelerating gradient, experiencing acceleration up to $\sim 1270\,\mathrm{MeV}$, with a small decrease in energy gain observable as the charge increases [see Figure 4.8 (a)]. This is likely a symptom of beam loading, necessarily implying the efficient extraction of energy from the accelerating field, which again is evident in the plasma emission-light signal.

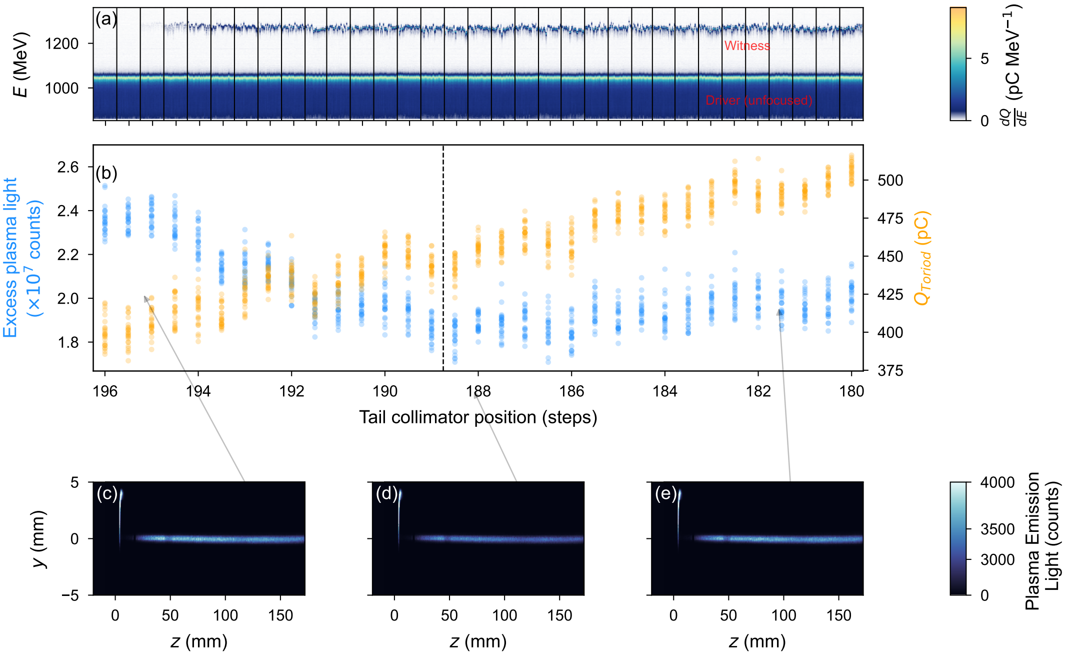

Figure 4.8: Plasma emission-light variation with increasing witness charge. (a) Accelerated trailing-bunch and decelerated (defocused) drive bunch spectra, grouped by the position of the tail collimator. (b) Excess plasma light yield (blue points) and the combined driver and witness charge detected in the upstream toroid Q_{toroid} (orange dots) versus the position of the collimator. Also shown are example images of the plasma light signal observed by the sideview with (c) driver charge only, (d) reduced light due to the acceleration of witness charge and (e) a re-increased light yield as trailing bunch charge is introduced behind the first wake, with arrows indicating their corresponding shots in (b). In order to highlight the change in light signal, a power law colormap normalisation with an exponent of 4 was used in (c–e).

As Figure 4.8 (b) shows, beginning with the driver-only case, the plasma light decreases as more trailing-bunch charge is introduced [see example images in Figures 4.8 (c) and (d)]. This therefore reinforces the hypothesis that the excess emitted plasma-light is connected to the energy remaining in the wake; just as this signal increases with energy deposition (e.g. by a driver), it is also reduced when energy is extracted (e.g. by a witness bunch). Furthermore, this observation suggests that the amount of plasma emission-light can be used to estimate the energy remaining in the wake—an idea that is further pursued in Section 4.2.3.

Interestingly, beyond a certain collimator position (denoted by the black dashed line in Figure 4.8 (b)), the plasma emission-light yield begins to *increase* again [see example image in Figure 4.8 (e)]. Note that although the charge detected by the upstream toroid continues to rise, this charge is in fact not observed on the electron spectrometer screen. A hypothesis explaining this behaviour is that beyond this point, further trailing-bunch charge that is introduced is placed behind the vertex of the first blowout. Whilst some of this charge will naturally be lost due to the defocusing fields behind the bubble, slices of the bunch that are positioned even further back will begin to experience the decelerating fields of the second wake. Such charge will instead lose energy to the plasma—much the same as the driver did in the first wakefield—resulting in a re-increase in the plasma emission-light. In the case that this charge is either decelerated or simply defocused, its detection with the spectrometer would prove difficult; indeed, even if these electrons are successfully transported to the screen, they would likely appear at lower energies that were not well-imaged in this measurement. Therefore, although this presents a novel observation of electron deceleration in the second wake, the exact amount of energy exchanged with the plasma—as is required in further analysis—cannot be accurately calculated for this charge. In practice, for this effect not to be confused with an increase in the energy deposited by the driver when measuring the plasma emission-light, this regime should first be identified by observing that some amount of witness charge is consistently missing on downstream diagnostics (as was done in this case with the spectrometer) and then avoided. From now on, we therefore focus on the data that is comfortably outside of this regime [i.e. $\geq$ 189 steps, to the left of the black dashed line in Figure 4.8 (b)].

4.2.3 Energy-transfer measurements with plasma emission-light

In order to truly relate the energy loss and gain of the drive and trailing bunches to the excess plasma emission-light, a similar procedure as in the analysis of the high-efficiency working point in Section 4.2.1 is applied to the data in the driver and witness tail collimator position scans (see Figure 4.9).

Plasma emission-light enhancement with energy deposition

Figure 4.9 (a) shows the driver total energy loss ΔW_{dec} plotted against the excess plasma emission-light yield for every shot in Figure 4.7. The resulting curve connects the excess captured plasma emission-light yield I_p to the total driver energy loss i.e. $I_p = f(\Delta W_{dec})$. Provided that the total energy lost by the driver is converted purely to wake energy implies that $\Delta W_{dec} = \Delta W_{wake}$ and therefore that $I_p = f(\Delta W_{wake})$. In principle, this suggests that the plasma emission-light yield can be used to estimate the energy remaining in the wake via $\Delta W_{wake} = f^{-1}(I_p)$.

Ultimately, the specific form of $f(\Delta W_{wake})$ should have no impact on its use in measuring the wake energy—it can be viewed simply as a calibration curve. Regardless, the shape of $f(\Delta W_{wake})$ is still worth commenting on. Whereas the plasma emission-light increases rapidly within the first $\sim$ 5 mJ of energy deposition, beyond a total driver energy loss of $\sim$ 10 mJ the curve is well-approximated by a linear function. Such a function is fitted to the data where $\Delta W_{dec} > 10$ mJ (orange dashed line). Speculating, this could for example be a sign of a change in regime (e.g. linear to non-linear wakefield).

It should however be noted that the exact form of the response curve depends on the experimental setup: the amount of detected signal will, for example, be affected by the light collection optics and the sensitivity of the camera chip. Illustrated in Section 4.1.3, the exposure window started somewhat later than beam-arrival, meaning that the signal at earlier times was not measured. In a similar way, the varying spectral response of the camera chip plus the chromaticity of the collection optics meant that only wavelengths within a finite spectral range were detected. Even if the total light signal integrated over all times and wavelengths varies linearly with the energy dumped into the plasma, this is not necessarily true for the I_p measured in the finite detection window used here. In this way, the choice of exposure settings and spectral bandwidth also become imprinted on the shape of $f(\Delta W_{wake})$. While the measurement configuration remained identical throughout the results presented here, any changes to the setup would mean that the response curve would need to be recalibrated.

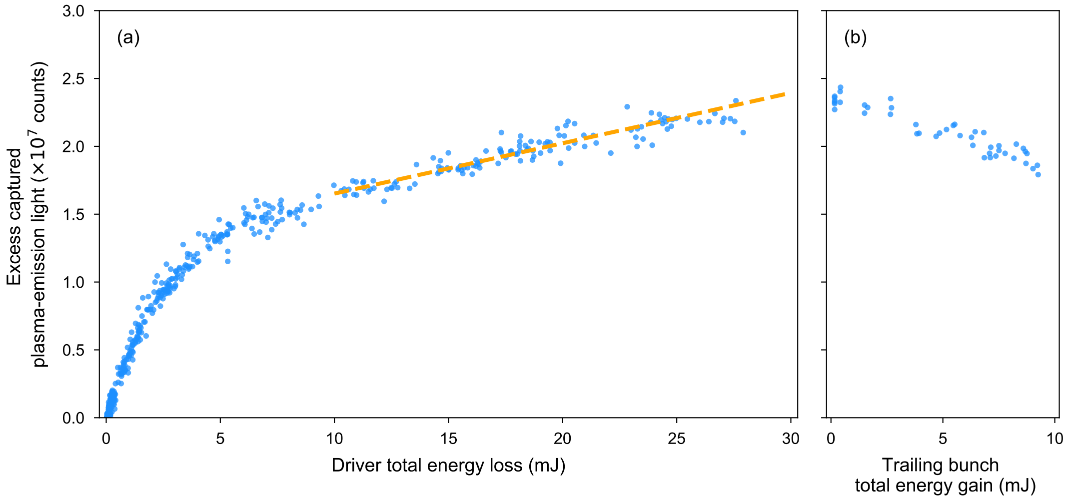

Figure 4.9: Excess captured plasma emission-light I_p is plotted against (a) the total energy loss of the driver and (b) the total energy gain of the witness, varied by adjusting the position of the tail collimator. The response curve f constructed in (a) can be estimated via a linear fit (orange dashed line) for $\Delta W_{dec} = \Delta W_{wake} \gtrsim 10$ mJ.

Although in operation it is likely that the diagnostic settings remain unchanged, the response curve will also in general depend on the plasma and surrounding neutral gas densities, affecting how the wake energy is dissipated and also the amount of material there is to undergo ionisation and excitation processes [201, 205]. Calibrating and applying the response curve therefore requires the initial plasma and gas conditions to be stable—a prerequisite any functioning plasma accelerator stage should already fulfil. If small enough, jitters in the initial state of the plasma, for example due to fluctuations in the ionisation source (in this case a discharge), would therefore simply contribute to the random error of the measurement. It is however reasonable to expect that the practical operation of a plasma accelerator will involve changing the initial conditions of the plasma; for example, the plasma density will likely act as a primary control over the energy gain in a given plasma stage. Provided that the plasma can be reliably reproduced with a given set of experimental settings—in this case the discharge voltage, gas pressure and beam arrival-time—a whole family of response curves can be constructed, one for every point in the parameter space. The $f(\Delta W_{wake})$ measured and utilized here is therefore that corresponding to this specific working point; an array of response curves could, if needed, be calibrated for a whole range of different gas and discharge settings.

On a related note, when constructing or using the response curve, it is important to know with confidence the proportion of the light increase that is due to the excess wake energy, as opposed to e.g. beam-field ionisation or direct ionisation from the wakefields. For the beam intensities at FLASHForward, no beam-field ionisation was anticipated or indeed observed; using the estimate for the maximum radial field of a gaussian bunch derived from Equation 2.45, $E_{r,max} \approx 0.45 \frac{I_p}{2\pi\epsilon_0 c\sigma_r}$ only begins to approach the threshold for the ionisation of Ar-I ($\sim 60\,\mathrm{GV\,m^{-1}}$, see Figure 2.4) for a beam of peak current $I_p = 800\,\mathrm{A}$ when the radial beam size $\sigma_r \lesssim 0.35\,\mu\mathrm{m}$. Given that the beta functions were most likely close to the matched beta function for this optimised working point—$\beta_m \approx 2.9\,\mathrm{mm}$ for $n_e = 7 \times 10^{15}\,\mathrm{cm^{-3}}$ and an energy of 1050 MeV—it can be expected that $\sigma_r \approx 1.2\,\mu\mathrm{m}$ throughout the plasma, if an optimistic normalised emittance of 1 mm mrad is assumed. Similarly, the wakefield strength—calculated in Section 4.2.1 to be approximately $1\,\mathrm{GV\,m^{-1}}$ in the accelerating region of the wake—is far below that needed for field ionisation of argon. As such, it was not possible in this experiment to assess how these extra sources of ionisation would effect the relation between the energy deposition and the measured emission-light signals; if these effects were present, it may be reasonable to expect that the calibration becomes more complicated. However, returning to the case in Ref. [205] where beam-field ionisation is not only present but also necessary for plasma generation, relations between the emission-light and the wake energy can apparently still be drawn provided that the energy loss due to ionisation is comparatively small.

Plasma emission-light reduction with energy extraction

Figure 4.9 (b) plots the excess captured plasma emission-light against the total energy gain of the trailing bunch ΔW_{acc} for the tail collimator position scan shown in Figure 4.8. Only shots where approximately all of the witness charge was transported to the spectrometer were included in this analysis, precluding any for a tail scraper position ≤ 188 steps [i.e. to the right of the black dashed line in Figure 4.8 (b)]. For the remaining data, the systematic error due to imperfect charge-coupling was further mitigated by again discarding shots with less than 90% of the maximum trailing-bunch charge measured on the screen for each collimator setting.

Intuitively, for a driver depositing energy ΔW_{dec} into the wake, the remaining wake energy after a trailing-bunch extracts energy ΔW_{acc} can be written as

$$\Delta W_{wake} = \Delta W_{dec} - \Delta W_{acc} = f^{-1}(I_p), \tag{4.1}$$

implying that the plasma light yield I_p is, for a given response curve f, uniquely determined by the wake energy ΔW_{wake}. In other words, the excess plasma emission-light yield is independent of the details of the energy exchange within the blowout; whether the wake has seen only energy deposition by a driver or also extraction by a witness, it is only the final value of ΔW_{wake} that matters.

4.2.4 Comparison of energy-transfer efficiencies measured with the spectrometer and plasma-light based methods

In order to test the statement about the wake energy in Equation 4.1, it can first be re-cast into a different form. Rearranging Equation 4.1, we find that the total witness energy gain can be inferred from a measurement of I_p in combination with knowledge of the driver total energy loss and the response curve i.e. $\Delta W_{acc} = \Delta W_{dec} - f^{-1}(I_p)$. Substituting this into Equation (2.68), the energy-transfer efficiency can be expressed as

$$\eta_p = 1 - \frac{f^{-1}(I_p)}{\Delta W_{dec}}, \tag{4.2}$$

where the notation η_p is used to distinguish between the efficiency measured with the plasma-light and that measured purely with the spectrometer, η_s. Provided that the spectrometer gives an accurate value of the energy-transfer efficiency, the validity of both Equations (4.1) and (4.2) can therefore be proven by confirming that $\eta_p = \eta_s$.

Both the calculation of η_p and η_s however require knowledge of the total driver energy loss ΔW_{dec}. As in the efficiency calculations in Section 4.2.1, one option is to simply measure the drive beam spectrum separately and assume its value of ΔW_{dec} to be constant throughout measurements of the trailing-bunch energy gain. However, drifts and jitters of the total driver energy loss away from this value will decrease the accuracy of any comparison of η_p and η_s—especially as the rate-of-change in each of these efficiencies have opposite directions for the same error in ΔW_{dec}. Shot-to-shot estimates of ΔW_{dec} simultaneous to measurements of the witness total energy gain are therefore ideally required but, as mentioned in Section 4.2.1, are made difficult by the fact that only part of the spectrum can be accurately measured at a given instance. Naturally, point-to-point imaging of the trailing-bunch is usually a priority when measurements of its spectrum are involved—as was the case here. A way to estimate the total driver energy loss on a shot-to-shot basis using the unfocused drive beam spectrum is therefore explored in the next section.

Driver total energy loss estimates during witness spectral measurements

As a basis for estimates of the driver total energy loss when an accurate spectral characterisation of the drive beam is unavailable, we can in fact turn again to the plasma emission-light. Provided that no witness beam is present and that the experimental conditions are still such that the previously constructed response curve is valid, the driver total energy loss can be estimated from $\Delta W_{dec} = f^{-1}(I_p)$. For the first two positions in the scraper scan of the witness beam the witness remains completely blocked whilst the driver is completely transmitted. Thus, only driver charge is detected on the spectrometer while an average plasma light-yield $\langle I_p \rangle_0 = (2.36 \pm 0.06) \times 10^7$ counts is measured, where the error represents the standard deviation. Using the response curve, the driver total energy loss at the beginning of the witness scraper scan is then calculated to be $\langle \Delta W_{dec} \rangle_0 = f^{-1}(\langle I_p \rangle_0) = (28.38 \pm 1.71)$ mJ. Note that the uncertainty on this value has a contribution from that on $\langle I_p \rangle_0$, but also from the error on the slope and intercept of the linear fit representing f, all of which are combined via Monte-Carlo sampling with 100000 samples to produce the estimate of $\langle \Delta W_{dec} \rangle_0$ quoted here.

To further improve this estimate, we make use of the unfocused drive beam spectra. Although unable to provide *absolute* measurements of the driver total energy loss, these spectra can be utilized to account for drifts and jitters *relative* to the value measured with the plasma light. Figure 4.10 illustrates how this was achieved. Firstly, the charge Q_{dec} and mean electron energy decrease ΔE_{dec} can be extracted from the spectra of the same driver-only shots used in the calculation of $\langle \Delta W_{dec} \rangle_0$ (see Figure 4.10 (a)) and used to calculate $\langle Q_{dec} \Delta E_{dec} \rangle_0$. Although not an accurate estimate of the driver total energy loss, fluctuations relative to this value can be used to correct the estimate of $\langle \Delta W_{dec} \rangle_0$ for drifts and jitters on a shot-to-shot basis. Strictly speaking, the driver charge loss—and therefore the shape of the unfocused driver spectra—in the imaging beamline will depend on the divergence of the driver particles, which will vary both as a function of the slice energy and slice emittance. In this experiment however, for simplicity it is assumed that the spectrum roughly maintains its normalised shape (i.e., charge is lost approximately uniformly across the spectrum). Calculating $Q_{dec} \Delta E_{dec}$ for the driver spectra of the shots considered in the witness collimator scan (see Figure 4.10 (b)), a first-order correction factor for $\langle \Delta W_{dec} \rangle_0$ can be determined as $\frac{Q_{dec} \Delta E_{dec}}{\langle Q_{dec} \Delta E_{dec} \rangle_0}$. Plotted against the tail collimator position in Figure 4.10 (c), this correction factor indicates that the value of the driver total energy loss can drift as low as 86% of the value estimated at the beginning of this scan. Represented by the errorbars, the uncertainty on the correction factor derives from the error on the value of $\langle Q_{dec} \Delta E_{dec} \rangle_0$—estimated here to be on the level of $\approx 3\%$. The corrected value of the driver total energy loss, written as

$$\Delta W_{dec} \approx \langle \Delta W_{dec} \rangle_0 \frac{Q_{dec} \Delta E_{dec}}{\langle Q_{dec} \Delta E_{dec} \rangle_0} = f^{-1}(\langle I_p \rangle_0) \frac{Q_{dec} \Delta E_{dec}}{\langle Q_{dec} \Delta E_{dec} \rangle_0}, \tag{4.3}$$

therefore receives uncertainty contributions from f, $\langle I_p \rangle_0$ and $\langle Q_{dec} \Delta E_{dec} \rangle_0$ that will propagate through to efficiency calculations and can again be estimated via Monte-Carlo sampling.

Note additionally that just as the driver total energy can be corrected using this method, so can e.g. the driver charge. This becomes useful in later analysis (see Section 4.2.5), where the driver-only BPM charge can be combined with a correction factor $\frac{Q_{dec}}{\langle Q_{dec} \rangle_0}$ measured on the spectrometer. Yielding shot-to-shot values of the driver charge, this method therefore enables an estimate of the incoming witness charge via subtraction from an upstream BPM charge measurement.

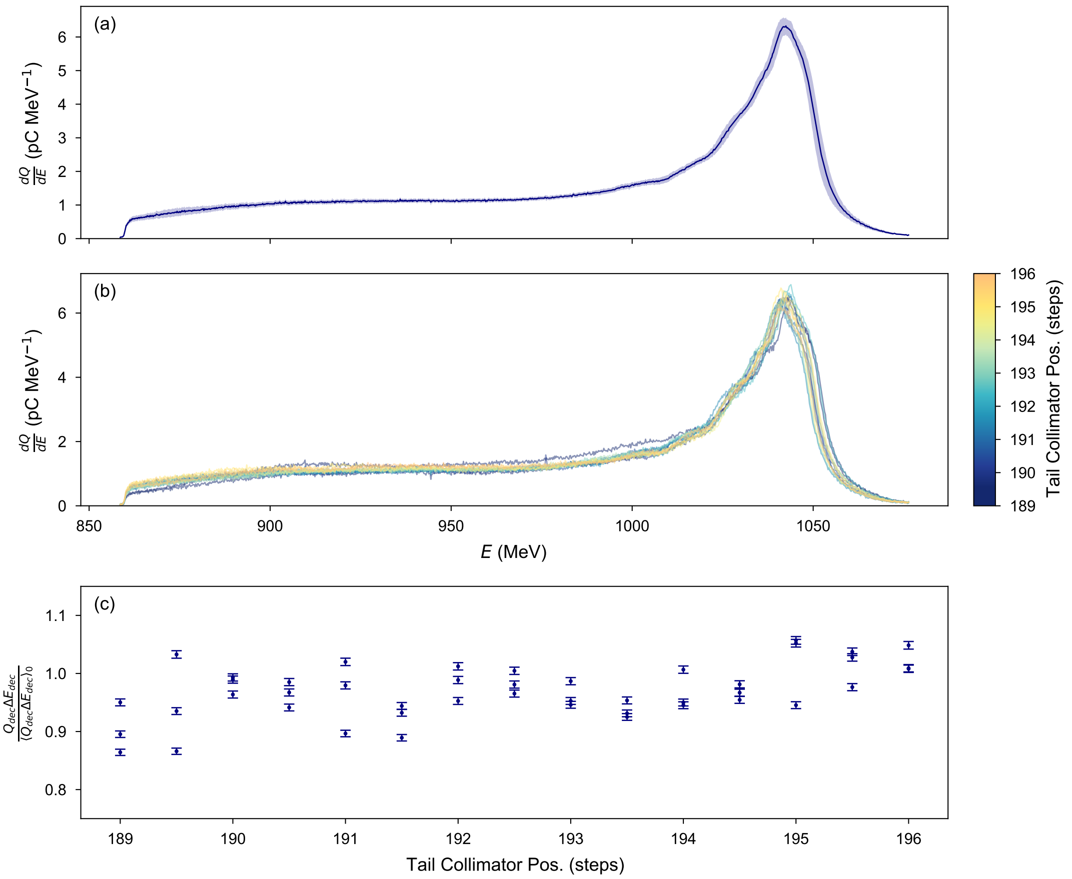

Figure 4.10: Driver total energy loss correction using unfocused drive beam spectra. (a) Mean spectrum of the unfocused driver on the spectrometer screen for shots in the witness collimator position scan containing only drive beam charge, where the error band represents the standard deviation of the spectra. (b) Unfocused driver spectra and (c) drive beam total energy loss correction factors for the points in Figure 4.9 (b) are also shown, with the corresponding tail collimator positions indicated.

Spectrometer versus plasma-light based efficiency measurements

For consistency, when calculating the values of η_s and η_p for each of the shots in Figure 4.9 (b), the same estimate of the driver total energy loss is used in both Equations (2.68) and (4.2).

Figure 4.11 compares the energy-transfer efficiencies calculated from the electron spectrometer with those derived from the plasma emission-light. Figure 4.11 (a) first of all shows the efficiency values for the case where the driver total energy loss is estimated from the plasma light yield when only the driver is present, but assumes that this remains constant throughout the scan i.e. no shot-to-shot correction is applied.

Although the results here demonstrate a direct correlation between η_p and η_s, there is also a disagreement in their absolute values, particularly at higher efficiencies where the points begin to stray from the line $\eta_p = \eta_s$ that represents 100% agreement. Indeed, rather than a gradient equal to 1, fitting a straight line passing through (0,0) to this data via Orthogonal Distance Regression (orange line) yields a slope of (1.33 ± 0.03), indicating that η_s is being underestimated and/or η_p is being overestimated. As explained previously, this can be attributed mainly to an underestimation of ΔW_{dec}.

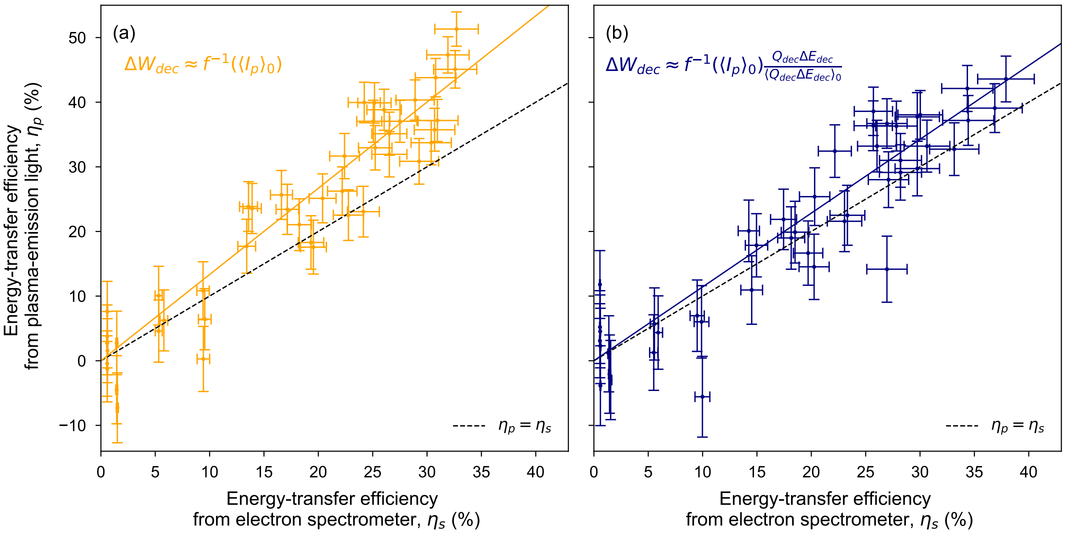

Figure 4.11: Agreement between spectrometer and plasma-light based energy-transfer efficiency methods. Energy-transfer efficiencies from the electron spectrometer η_s and plasma emission-light η_p are plotted against each other, with the driver total energy loss estimated (a) as constant and (b) with an additional correction factor. Agreement between the two methods is signified by linear fits to the points (solid lines) following $\eta_p = \eta_s$ (black dashed line), which is improved by the inclusion of the correction factor in (b).

Figure 4.11 (b) demonstrates how accounting for a changing ΔW_{dec} improves the agreement between the two methods. By calculating the driver total energy loss with the correction factor (see Equation 4.3), what appears to be an overall decrease in ΔW_{dec} (see Figure 4.10 (c)) is accounted for. The improvement in the agreement that this affords is clear; most of the points in Figure 4.11 lie within experimental error of the agreement line $\eta_p = \eta_s$, including those at higher efficiencies. Due to the additional error contributions in the calculation of the correction factor, the errorbars are slightly larger here than in Figure 4.11 (a); the uncertainty on η_p is on average larger than that on η_s, with a mean error of 4.8 percentage points. Nevertheless, a line fitted to this data (blue line) in a similar way to that in Figure 4.11 (a) has a slope of

(1.14 ± 0.04)—an improvement of 19% compared to the previous estimate. A potential reason for the 14% discrepancy that remains may be a consequence of imperfect witness charge coupling. Although this effect is already reduced by considering only those shots with a high percentage of witness charge transmission to the spectrometer, this does not guarantee 100% of the trailing-bunch charge is measured in every shot, necessarily leading to a systematic underestimate of ΔW_{acc} and therefore η_s.

Note that in both of the cases in Figure 4.11 (a) and (b), there exists some data points with $\eta_p < 0$. Physically, this would correspond to the situation where a witness beam actually deposits more energy than it extracts. However, the negative efficiencies measured here are simply a result of the calculation of $\langle \Delta W_{dec} \rangle_0 = f^{-1}(\langle I_p \rangle_0)$ and $\langle Q_{dec} \Delta E_{dec} \rangle_0$. By examining Equation (4.2) with Equation (4.3) substituted in, we find that η_p can be negative if the term $\frac{f^{-1}(I_p)}{\Delta W_{dec}} = \frac{f^{-1}(I_p)}{f^{-1}(\langle I_p \rangle_0)} \frac{\langle Q_{dec} \Delta E_{dec} \rangle_0}{Q_{dec} \Delta E_{dec}}$ is greater than unity. This can happen if $Q_{dec} \Delta E_{dec} < \langle Q_{dec} \Delta E_{dec} \rangle_0$ or, given that f^{-1} is a monotonic increasing function, if $I_p > \langle I_p \rangle_0$. For the shots used to calculate $\langle I_p \rangle_0$ and $\langle Q_{dec} \Delta E_{dec} \rangle_0$, there are by definition some shots with $Q_{dec} \Delta E_{dec} < \langle Q_{dec} \Delta E_{dec} \rangle_0$ and/or $I_p > \langle I_p \rangle_0$, which can therefore yield $\eta_p < 0$.

Capabilities of spectrometer and plasma-light based efficiency-transfer diagnostics

At this stage, it is useful to compare and contrast the different capabilities of efficiency measurements performed with a conventional electron spectrometer to those made with the plasma-light based method.

Electron spectrometers are, and will remain, vital components of any plasma accelerator, especially where accurate diagnosis of bunch spectra is needed. Furthermore, they are still required for the construction of the wake-energy-to-plasma-light response curve $I_p = f(\Delta W_{wake})$ that forms the foundation of the plasma-light based method. However, as a consequence of the natural chromaticity of any quadrupoles that are used to transport the beam to the spectrometer screen, only a finite bandwidth of the spectrum can be properly imaged at once—problematic for shot-to-shot measurements of bunches with a wide spread of energies. Moreover, any measurement performed with the spectrometer must contend with potential beam losses between the location-of-interest—e.g. the plasma stage—and the spectrometer, at which point the beam is necessarily destroyed after being dispersed by the dipole and scattered by the diagnostic screen. The latter is particularly relevant for multi-stage plasma accelerators and applications, where the trailing-bunch must be transported between stages and/or to the point-of-use unimpeded.

The plasma-light based method on the other hand essentially exchanges a measurement of the witness total energy gain for that of the wake energy. This makes the method potentially non-invasive to the trailing-bunch; although an estimate of the driver total energy loss is still needed, the drive beam could in principle still be diagnosed with a spectrometer, with an imaging energy better suited for measuring its spectrum. Another benefit of the plasma-light method is that measurements require only a small amount of additional equipment—an extra camera and some collection optics per plasma stage. Furthermore, measurements are performed directly at the location of acceleration with this method, in contrast to trailing bunches that are measured on a downstream spectrometer that are prone to imperfect charge-coupling.

One of the key features of the plasma-light method stems from the fact that each measurement is an image that contains spatial information. Whereas the spectrometer only measures the efficiency downstream of the acceleration process, the plasma-light is detected along the full length of the accelerator stage. Whilst the spectrometer can therefore only give a *longitudinally averaged* energy-transfer efficiency, the plasma emission-light measurement can in principle yield local information about the acceleration process that is *longitudinally resolved*.

4.2.5 Longitudinally resolved efficiency measurements using the plasma-light based method

By separating the plasma into longitudinal sections, the energy-transfer efficiency in the i^{th} slice can be generally defined as

$$\eta_i = 1 - \frac{\Delta W_{wake,i}}{\Delta W_{dec,i}}, \tag{4.4}$$

where $\Delta W_{wake,i}$ and $\Delta W_{dec,i}$ are the energy transferred to the wake and the driver total energy loss within the i^{th} segment. With analogy to the longitudinally averaged case, in order to measure $\Delta W_{wake,i}$ with the plasma-light emitted from the i^{th} segment $I_{p,i}$, a relationship between the two must be found, relying on a driver-only measurement of $\Delta W_{dec,i} = \Delta W_{wake,i}$. Here however, a calculation of $\Delta W_{dec,i}$ was impossible; only the total driver energy loss after the plasma could be measured. In principle, one could foresee an experimental configuration in which segments of plasma are isolated and $\Delta W_{dec,i}$ measured for each e.g. variable length gas cells, but no such capability was available here.

An assumption must therefore be made about the relationship between ΔW_{dec} and $\Delta W_{dec,i}$ that at least holds during the construction of the response curve. It can be expected that the drive beam energy deposited in each segment simply scales linearly with the driver total energy loss at the end of the plasma,

$$\Delta W_{dec,i} = k_i \Delta W_{dec}, \tag{4.5}$$

where k_i is a constant of proportionality for the i^{th} segment. Next, response curves f_i for each longitudinal slice that are functions of the longitudinally integrated driver total energy loss i.e. $I_{p,i} = f_i(\Delta W_{dec})$ can be constructed, such that $\Delta W_{dec} = f_i^{-1}(I_{p,i})$. Substituting this into Equation (4.5), the driver energy deposited in the i^{th} segment can be written as $\Delta W_{dec,i} = k_i f_i^{-1}(I_{p,i})$. This is true in the driver-only case, but applies more generally written in terms of the wake energy,

$$\Delta W_{wake,i} = k_i f_i^{-1}(I_{p,i}). \tag{4.6}$$

Substituting Equations (4.5) and (4.6) into Equation (4.4), the longitudinally resolved energy-transfer efficiency can be written as

$$\eta_{p,i} = 1 - \frac{f_i^{-1}(I_{p,i})}{\Delta W_{dec}}, \tag{4.7}$$

where $\eta_{p,i}$ is the energy-transfer efficiency throughout the i^{th} segment, calculated from the segment emission-light yield and the longitudinally integrated driver energy deposition—both of which can be measured experimentally. Note that this is not a statement about the specific shape of $\Delta W_{dec,i}$; it is purely a way of estimating the total driver energy deposited in a given segment from the total driver energy loss at the end of the plasma. It is not, for example, required that the energy deposition is uniform along the direction of propagation, as it indeed won't be for e.g. the plasma up- and downramps. Additionally, although relying on Equation (4.5) holding true, any error in this assumption will simply imprint a systematic trend on $\eta_{p,i}$, meaning that comparisons of the local energy-transfer efficiency between shots will therefore still be valid.

Response curve construction for longitudinally resolved efficiency measurements

Figure 4.12 illustrates how the plasma-light signal was separated into longitudinal segments and how using the driver-only collimator scan, response curves can be constructed for each of the resulting slices.

Figure 4.12 (a) displays an example plasma-light image with only the drive bunch present. This was separated into 55 segments, each 3 mm long, providing a reasonable balance between resolution and statistical noise. The extent of the capillary from its entrance to the edge of the camera field-of-view is shown; unfortunately, no data was available for the last $\sim 20\,\mathrm{mm}$ of the cell.

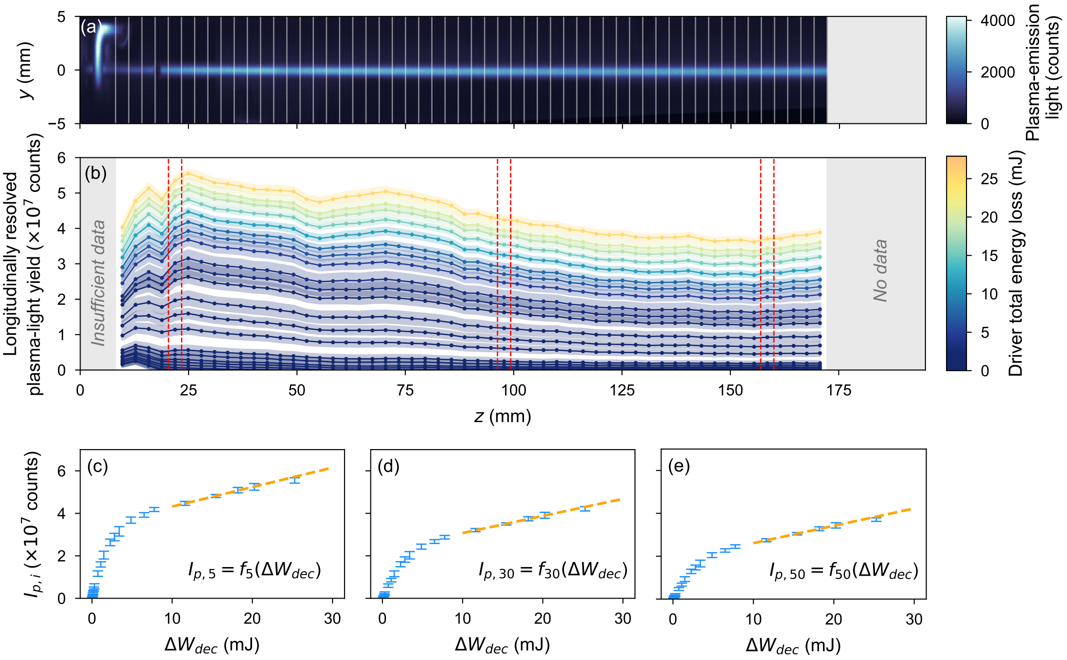

Figure 4.12: Longitudinal segment response curve construction. (a) Plasma-light image with only driver interaction, separated into 55 longitudinal segments, each 3 mm-long. (b) Plasma-light yield as a function of longitudinal position z for different amounts of longitudinally integrated driver energy deposition ΔW_{dec}. Each curve was found by averaging over multiple shots for a given collimator setting, with the errorband representing the standard deviation of $I_{p,i}$; the value of ΔW_{dec} corresponds to the mean for that collimator position. Example response curves of $I_{p,i}$ versus ΔW_{dec} (light blue errorbars) for the (c) 5th, (d) 30th and (e) 50th segments. Straight lines (orange dashed lines) are fitted to the linear region of each curve.

For each segment, the captured plasma emission-light yield $I_{p,i}$ can be measured shot-to-shot, where i is the segment number, increasing from upstream ($i = 0$) to downstream ($i = 54$). Figure 4.12 (b) plots the values of $I_{p,i}$ for every segment against their longitudinal positions z, where each curve constitutes an average over shots taken at the same collimator position, with the shaded bands signifying standard deviations. These curves are then colour-coded according to the mean value of the driver total energy loss ΔW_{dec} measured on the spectrometer for the corresponding shots. By plotting $I_{p,i}$ against ΔW_{dec} for each individual i, a response curve $I_{p,i} = f_i(\Delta W_{wake})$ for every segment can be constructed, remembering that in the driver-only case, $\Delta W_{dec} =$

ΔW_{wake}. Figures 4.12 (c–e) show example response curves for the 5th, 30th and 50th segments.

Longitudinally resolved efficiency measurements

Figure 4.13 shows examples of the longitudinally resolved energy-transfer efficiency calculated with Equation 4.7 and the response curves f_i generated in Figure 4.12.

An example plasma-light image with both driver and witness beam interaction is shown in Figure 4.13 (a), separated once again into 55, 3 mm-long segments—the resulting $\eta_{p,i}$ of which are displayed in Figure 4.13 (b) for four different shots, labelled (I)–(IV). Each curve in Figure 4.13 represents the mean value of $\eta_{p,i}$ for every segment, with the errorbands depicting the systematic uncertainty resulting from errors in the models f_i and the driver total energy loss estimates ΔW_{dec}—similar to that explained in Section 4.2.4. Dashed horizontal lines in Figure 4.13 (b) indicate the longitudinal average of $\eta_{p,i}$ for each shot, which correspond almost exactly to the value of η_p calculated with the whole image. In addition, the longitudinally averaged efficiency values measured with the ESPEC η_s are also displayed in Figure 4.13 (c) for comparison.

Note that for this analysis, it is useful to have an estimate of the incoming witness charge. In a similar way as in Section 4.2.4, the drive beam charge can be estimated using the last BPM before the plasma chamber for driver-only collimator settings, and corrected shot-to-shot by a factor $\frac{Q_{dec}}{\langle Q_{dec} \rangle_0}$ measured using the unfocused driver spectra on the spectrometer. With an estimate of the driver charge for every shot, this can be subtracted from the upstream BPM charge, yielding the charge of the incoming trailing-bunch.

All of the shots here are taken from the same collimator position [189 steps in Figure 4.8 (b)]; the requirement that they have high charge-coupling is however relaxed. As a result, although the shots have similar input conditions—the incoming trailing-bunch charge $Q_{acc,in}$ for each shot lies in the range 48–56 pC—the evolution of the local plasma efficiency and also the efficiency measured at the spectrometer can vary dramatically.

Before looking at their behaviour in detail, it is first useful to highlight features that are common to all shots. For example, in all cases there is a rapid initial increase in the local efficiency (for up to $s \sim 20$ mm)—likely due to the plasma density upramp at the start of the cell caused by the expulsion of gas into vacuum. After acceleration, the efficiency measured on the spectrometer for a given shot is also consistently lower than its longitudinally averaged value from the plasma. As discussed previously, an obvious reason for this is imperfect charge-coupling, the degree of which is dependent on transport both in and after the plasma. However, another factor affecting the value of η_p is that the cell exit was not captured by the field-of-view of the camera. This low-density downramp region, as with the plasma entrance, would be expected to have a lower local energy-transfer efficiency than the bulk plasma; the absence of this part of the measurement could therefore be responsible for a small overestimate of η_p, which also applies to the results in Figure 4.11.

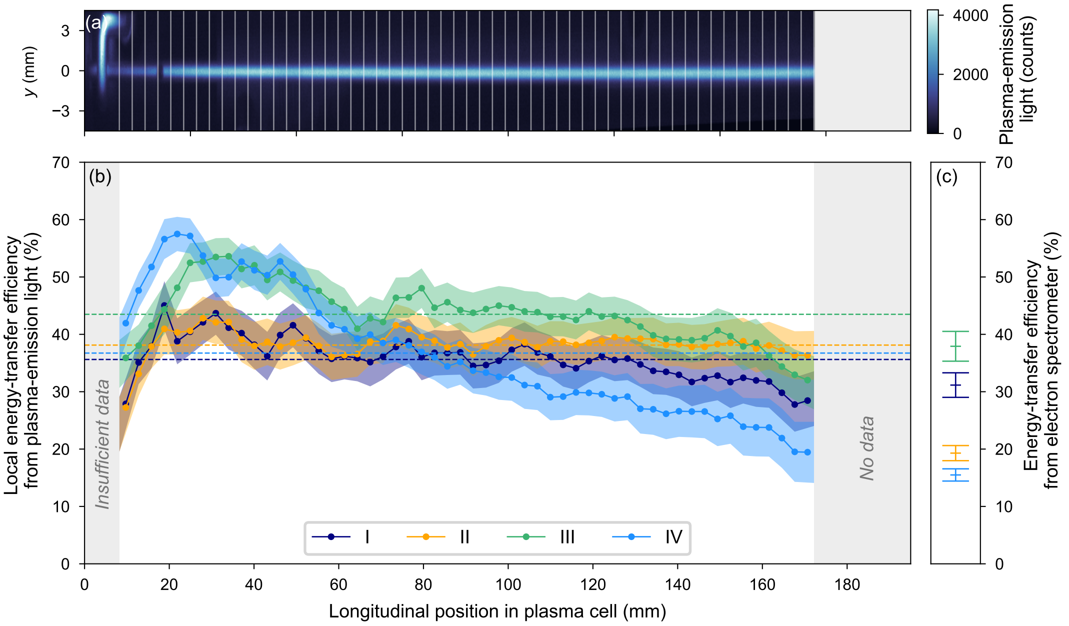

Figure 4.13: Longitudinally resolved energy-transfer efficiency measurements from the plasma emission-light. (a) Plasma-light image with driver and witness interaction, separated into 55, 3 mm-long longitudinal segments. (b) Local energy-transfer efficiency $\eta_{p,i}$ as a function of longitudinal position in the plasma cell, measured for four separate shots. The errorband on each curve represents the systematic uncertainty on the driver total energy loss and the response curves. Longitudinally averaged values of the efficiency $\langle \eta_{p,i} \rangle = \eta_p$ are also shown for each shot (horizontal dashed lines). (c) Longitudinally averaged efficiencies measured on the spectrometer corresponding to each shot.

Beginning with shots (I) and (II), the local and averaged efficiencies measured with the plasma-light and spectrometer respectively can be compared between cases. With incoming witness charges $Q_{acc,in}$ equal to (I) (48 ± 2) pC and (II) (56 ± 2) pC, the local energy-transfer efficiencies of these shots show a similar evolution, remaining relatively uniform, with the exception of a slight decrease towards the end of the cell in the case of (I). The similarity of $\eta_{p,i}$ for these shots is also reflected in the longitudinally averaged efficiencies from the plasma light, where η_p equals (36 ± 4) % and

$(38 \pm 4)\,\%$ for shots (I) and (II) respectively. However, a significant difference between these cases is evident in the longitudinally averaged efficiency measurements on the spectrometer. Whereas for shot (I) the efficiency measured on the spectrometer $\eta_s = (31 \pm 2)\,\%$ lies within experimental error of its corresponding η_p value, for shot (II) $\eta_s = (19 \pm 1)\,\%$—much less than measured with the plasma-light. As the behaviour of $\eta_{p,i}$ is relatively similar in each case, this would suggest that the reason for the discrepancy arises in the transport from the plasma to the spectrometer. Indeed, by calculating the witness charge detected on the spectrometer screen as a fraction of that detected upstream of the plasma chamber, we find charge-couplings of $(85 \pm 4)\,\%$ for (I) and $(46 \pm 2)\,\%$ for (II)—implying that increased trailing-bunch charge losses, likely in the spectrometer imaging beamline, are responsible for the low value of η_s measured in the latter case.

Similar comparisons can also be extended to shots (III) and (IV). Whereas both of these shots reach higher local efficiencies than in (I) and (II) initially—as much as $(58 \pm 3)\,\%$ in the case of (IV)—the value of $\eta_{p,i}$ decreases as a function of longitudinal position in the plasma cell, possibly as a result of an evolution of the wakefield generated by the driver. However, despite having the same incoming trailing-bunch charge of $(50 \pm 2)\,\mathrm{pC}$, the rate of this decline, found by fitting straight lines to the data points from $\sim 20\,\mathrm{mm}$ onwards, was found to be $0.23\,\%\mathrm{mm}^{-1}$ for shot (IV)—around twice as high as that in shot (III). Naturally, this also results in a lower η_p for (IV) where $\eta_p = (37 \pm 4)\,\%$, than for (III) in which $\eta_p = (44 \pm 4)\,\%$. Looking again at the efficiencies measured on the ESPEC, for shot (III) $\eta_s = (38 \pm 3)\,\%$ is in agreement with η_p, whereas the η_s in the case of (IV) is much lower at a value of $(16 \pm 1)\,\%$. This again is likely a symptom of low charge-coupling; whereas $(98 \pm 4)\,\%$ of the incoming witness charge is detected on the spectrometer in (III), this was only $(38 \pm 2)\,\%$ for shot (IV). Whilst the fact that $\eta_p > \eta_s$ for the latter again indicates imperfect transport between the plasma and spectrometer, the rapidly decreasing $\eta_{p,i}$ from around 50–70 mm suggests that continuous charge loss already begins *within* the plasma. As the input parameters of (III) and (IV) are approximately the same, the behaviour of the low charge-coupling shot is therefore consistent with the growth of a transverse instability. Moreover, as the local energy-transverse efficiency is initially high, this could be connected to the *efficiency-instability relation* [137, 138], in which trailing bunches accelerated with high-efficiency are inherently prone to transverse instabilities.

Independent of whether the behaviour exhibited by shot (IV) is truly a result of a transverse instability, the fact that this may be discussed at all highlights an important and novel capability of the plasma-light based diagnostic; with this method, the evolution of the acceleration process can, for the first time, be studied experimentally with longitudinal resolution. This includes, but is not limited to, the investigation of transverse instabilities, which have been historically difficult to diagnose.

4.3 Outlook

Looking ahead, the study of instabilities is just one potential application of the plasma-light-based diagnostic. As a measure of the energy remaining in the plasma wake, the plasma emission-light signal may also be relevant in high repetition rate PWFA, where any excess heat dissipated into the plasma will prolong its relaxation time and even damage components such as the gas cell [165, 199, 208]. Mentioned previously in Section 4.1.3, greater insight into the long-term evolution of the plasma may also be afforded by decomposing the emission-light temporally and/or spectrally [201, 205]. For instance, the change in the relative intensity of certain spectral lines may indicate the levels of the plasma that are reionised by the energy deposited by the driver, and if there is any difference when part of that energy is extracted by a trailing bunch. Already then, there are potentially multiple ways in which this diagnostic can be extended to study interesting physics in a single PWFA stage.

However, as alluded to in the comparison to traditional dipole spectrometers (see Section 4.2.4), the nature of the plasma-light-based diagnostic that makes it non-invasive to the trailing bunch particularly lends to applications and also multi-stage accelerators [209–212]. Figure 4.14 shows a schematic of how the individual stages in a plasma-based linac could be diagnosed with this new method on a shot-to-shot basis. In each stage, the witness bunch is accelerated in a plasma wake driven by a fresh drive beam. At the end of a stage, the decelerated driver is separated from the trailing bunch, after which the driver energy deposited in that stage can be estimated from the imaged drive beam spectrum. The emission-light collected by cameras observing each plasma can then provide single-shot measurements of the longitudinally resolved energy-transfer efficiency measurements along the accelerator. Regardless of the details of the acceleration scheme, it is foreseen that the plasma-light-based diagnostic will be key in the operation of future plasma-based accelerators, enabling the online monitoring and optimisation of the acceleration process in a way that integrated, invasive measurements of the trailing-bunch do not allow.

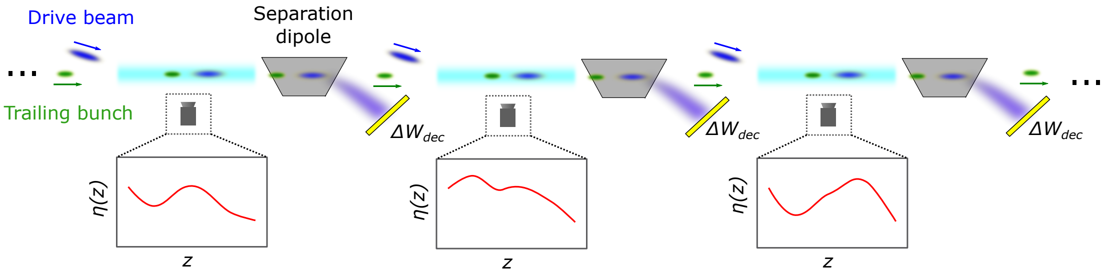

Figure 4.14: Plasma-light based efficiency diagnosis of a multi-stage plasma accelerator. Shown is a conceptual drawing of how the emission-light from consecutive PWFA stages could be used to provide longitudinally resolved energy-transfer efficiency measurements (red lines) along a plasma-based linac. In each stage, drive beams (blue) are used to generate plasma wakes before being dumped and diagnosed; the trailing-bunch (green) on the other hand is free to propagate to subsequent stages and eventually applications.

As a final outlook, one may ask whether the plasma-light based diagnostic could also be applied to Laser Wakefield Accelerators, with either a single or multiple stages [213, 214]. Fundamentally this would seem possible; just as with a particle driver, a laser driver deposits energy in the wake that can either be extracted by a witness bunch or dissipated into the surrounding plasma, which should manifest itself as excess emission-light. A relation between the emitted light and the wake energy could then be constructed, but relies on the experimental conditions being stable enough such that the calibration holds. The applicability of this technique to LWFA is therefore likely dependent on practical issues, such as whether reliable response curves can be generated. For instance, the driver in LWFA is often used to also produce the plasma channel via field ionisation, meaning the initial plasma density itself can also depend on the properties of the driver. This could then complicate finding the connection between wake energy and plasma emission-light— the initial ionisation due to the laser driver would have to be stable enough such that the plasma could be considered constant shot-to-shot, which may be the case when using low-Z gases such as hydrogen and helium.

4.4 Summary

Beginning from a high energy-transfer efficiency working point, it was shown in this chapter that the excess light emitted from a plasma after beam-interaction depends on how much energy was deposited in the driven wake. Moreover, a reduction in this light was shown to be a signature of loading of the wake by an energy-extracting trailing-bunch. By constructing a light-versus-wake-energy response curve, the emission-light signal could therefore be used to non-intrusively estimate the wake energy remaining in the plasma and in turn the energy-transfer efficiency of the acceleration process. After confirming that efficiency measurements based on the plasma light agree with those of a traditional electron spectrometer, the capabilities of the two methods were compared.

In addition to the non-invasiveness of the plasma-light-based method to the witness bunch, this technique was found to be uniquely capable of measuring the longitudinally resolved energy-transfer efficiency. Using this, the acceleration process for a number of shots was studied with longitudinal resolution, revealing local efficiencies of up to 58 % for one shot and also potential evidence of transverse instabilities in another. It is therefore foreseen that this technique will be useful in the short-term in studying the acceleration process in more detail at the single-stage level. In the long-term however, it is believed that this diagnostic will become a vital part of the online operation and optimisation of multi-stage plasma accelerators, ultimately enabling them to serve practical applications.

Chapter 5

Optimisation of a plasma cathode PWFA stage

As well as boosting the energy of existing particle bunches, plasma wakefield accelerators can themselves also act as particle sources, generating witness beams directly in the blowout structure in what is often termed a plasma cathode. Different methods of achieving internal injection in a plasma wakefield are discussed in Section 2.4.2, as well as the merits of using a PWFA-based plasma cathode rather than one based on LWFA. In general, most of these schemes promise injection of beams with low emittance and high current, resulting in correspondingly high brightnesses. One advantage of PWFA-based plasma cathodes is that the properties of the injected beams are in principle largely independent of the quality of the drive beam; all that is needed is a driver that can generate a strong enough wake to trap particles that are not yet travelling close to the speed of light. Furthermore, if the resulting witness beams have brightnesses that surpass that of the drive beam, such injection stages may be seen as *brightness transformers*.

The X-1 plasma cathode experiment at FLASHForward aims to demonstrate such an internal injection stage [152]—with the eventual goal to act as a brightness transformer for the FLASH drive beam. Here, internal injection of electron beams is achieved on the downramp of a sharp plasma density spike, generated via optical ionisation of higher levels of a base gas species with a transversely propagating injection laser—often referred to as a 'plasma torch' [56]. First demonstrated at FACET [215, 216], this method was then deployed in the X-1 experiment, producing internally injected bunches in a PWFA with unprecedented stability for a plasma accelerator [197]. Some of the main properties of these bunches were measured, but never fully optimised.

In this chapter, further studies from the X-1 project are presented. This includes the establishment of the plasma cathode stage that, through improved optimisation techniques, produced beams with enhanced properties that were more thoroughly characterised. Data presented for this experiment was obtained over a campaign consisting mainly of runs taking place in April, June, September, October and November of 2021. These runs took place largely in parallel to FLASH user experiments; although this restricted some of the beam parameters (e.g. the beam energy), the flexibility offered by the FLASH linac (see Section 3.1.3) allowed largely independent control over e.g. the driver compression. Detailed tailoring of the longitudinal phase space was also not a priority in the plasma cathode experiment, making these 'grey' shifts well-suited for internal injection studies. In particular, June 2021 yielded results that culminated in the most complete characterisation of internally injected beams; the presented measurements are therefore primarily from this run, with measurements from other shifts included where illustrative or necessary. Supplementary to experimental work, simulation studies are shown that aim to replicate the real-life conditions of the experiment and thus recreate and understand the injected bunches produced by the plasma cathode. Ultimately, the final product of the work in this chapter is that the brightness transformation of an input drive beam was successfully demonstrated, which forms the core result of Ref. [217].

5.1 Experimental Setup

Figure 5.1 summarises the setup of the plasma cathode stage in the X-1 project. The hexapod was carefully positioned for maximum transmission of all beams—namely the driver, along with the preionisation and injection lasers—through the capillary. In almost all cases, the 50 mm-long channel was used in this experiment, filled with the standard 97:3 Ar:H_2 mixture. At the buffer pressure of 46 mbar commonly used in all internal injection experimental runs ANSYS fluid flow simulations corroborate this, indicate a corresponding neutral gas density in the centre of the channel of $3.5 \times 10^{17}\,\mathrm{cm}^{-3}$ (see Section 5.1.3). Production of the accelerating plasma channel in the context of this experiment was performed by the longitudinal laser—in principle the threshold nature of field ionisation should allow better control of the ionisation levels of the argon than the discharge (see Section 2.3.1). The transverse laser, with its significantly more intense focus, was used to ionise higher levels of the neutral gas, forming a sharp density spike at a position 20 mm along the channel generated by its longitudinal counterpart. Simulations of the ionisation degree achieved by each of these lasers are presented in Section 5.1.3.

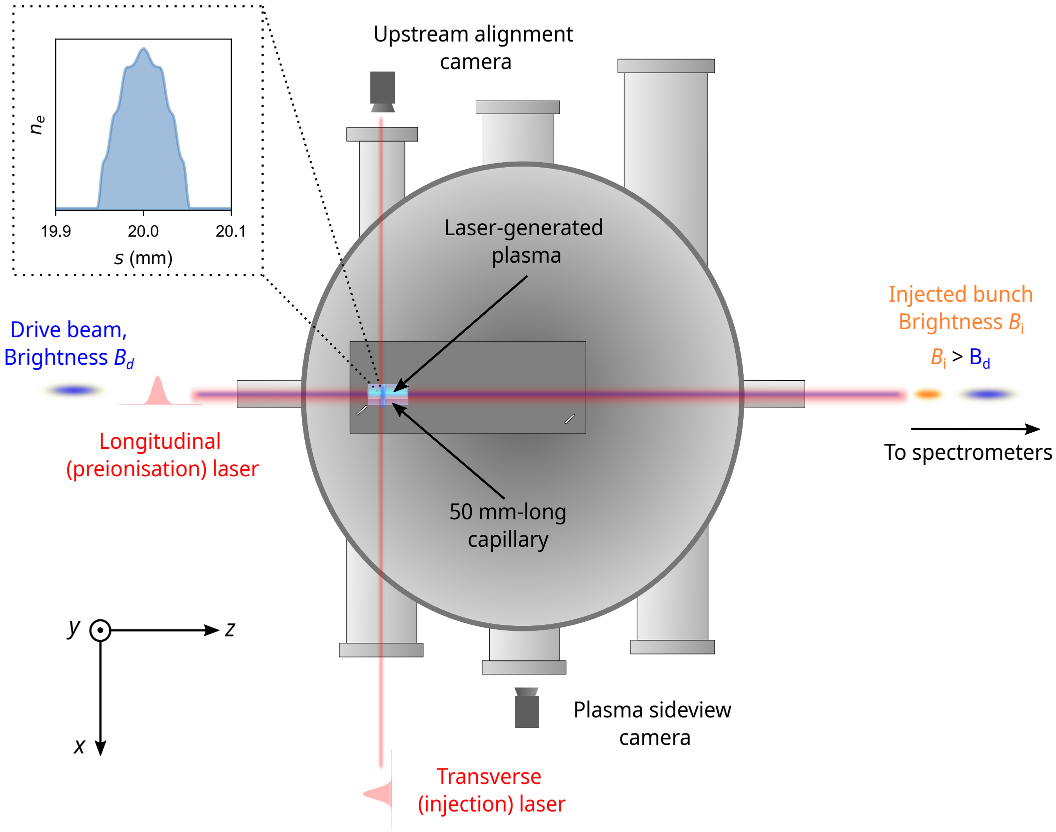

Figure 5.1: Plasma cathode stage setup. The 50 mm-long capillary is positioned such that the longitudinal laser and drive beam propagate freely through the channel (red and navy horizontal lines). Travelling perpendicular to these beams through the hole 20 mm downstream of the cell entrance, the transverse laser propagates directly to the upstream alignment camera for continual diagnosis (red vertical line). The plasma sideview camera also provides observations of the plasma generated by the longitudinal laser within the capillary.

A wake generated by the drive beam in this plasma could then experience internal injection of a witness beam when traversing the density downramp produced by the transverse laser (blue filled curve in inset of Figure 5.1). After acceleration in the remaining 30 mm of the plasma cell, this witness beam, as well as the decelerated driver, were then transported to and diagnosed by the downstream electron spectrometers. In addition to these, the upstream alignment diagnostic provided online monitoring of the transverse laser spot, whilst the plasma sideview camera proved further useful for observing the longitudinal plasma and its interaction with the drive beam. Ultimately, this setup formed the basis for our PWFA-based plasma cathode stage; if the brightness of the injected bunch B_i exceeds that of the driver B_d, the stage can be considered a *brightness booster*. The preparation of the individual components of this setup is described in more detail in the following sections.

5.1.1 Driver setup and characterisation

Experimental runs generally began with the compression of the drive beam to high peak currents, followed by the focusing of the driver to a spot size that was small enough for coupling into the plasma. The following section illustrates how the driver was compressed and focused for the purposes of the plasma cathode experiment, along side characterisation of both the longitudinal and transverse properties of the resulting bunches.

Driver longitudinal phase space

Although important for generating wakefields that can accelerate witness beams with high gradients, high-current drive bunches are also key for the trapping of sub-relativistic electrons that are internally injected into the wake. Drive bunches for this experiment were in general therefore highly compressed, primarily via tuning of the applied chirp in the ACC139 SRF module in order to produce strong compression in the linac bunch compressors, in which the R_{56} is kept constant (see Section 3.1.2). This is illustrated by a TDS measurement made in the June 2021 experimental run, presented in Figure 5.2. Figure 5.2 (a) shows an example of the (304 ± 2) pC drive beam's longitudinal phase space at a highly compressed working point, with a shot-averaged current-profile made at the same compression setting displayed in Figure 5.2 (b). For this particular working point, a peak current of (1.9 ± 0.2) kA and a (96 ± 6) fs rms duration were measured. Doubling also as a measurement of the spectrum, the mean electron energy of its driver was found to be 689 MeV, with am rms energy spread of 5 MeV and a peak spectral density 47 pCMeV^{-1}.

The relatively large driver currents measured here did however come at a price. Firstly, note the complicated longitudinal phase space; in general, such strong compression was essentially mutually exclusive from a linear longitudinal phase space, such as that used in the experiment described in Chapter 4. This therefore precluded detailed tailoring of the bunch profile via collimation (see Section 3.2.2), but also meant that the focusing of the beam was generally more complicated in the presence of chromaticity. Furthermore, such strongly compressed beams are susceptible to effects such as CSR, which will inevitably increase the transverse emittance and can potentially result in a beam that can no longer be focused tightly enough to drive a strong wake.

Tuning of the compression is therefore often a trade-off between maximising the peak current of the driver whilst keeping undesirable side effects to tolerable levels.

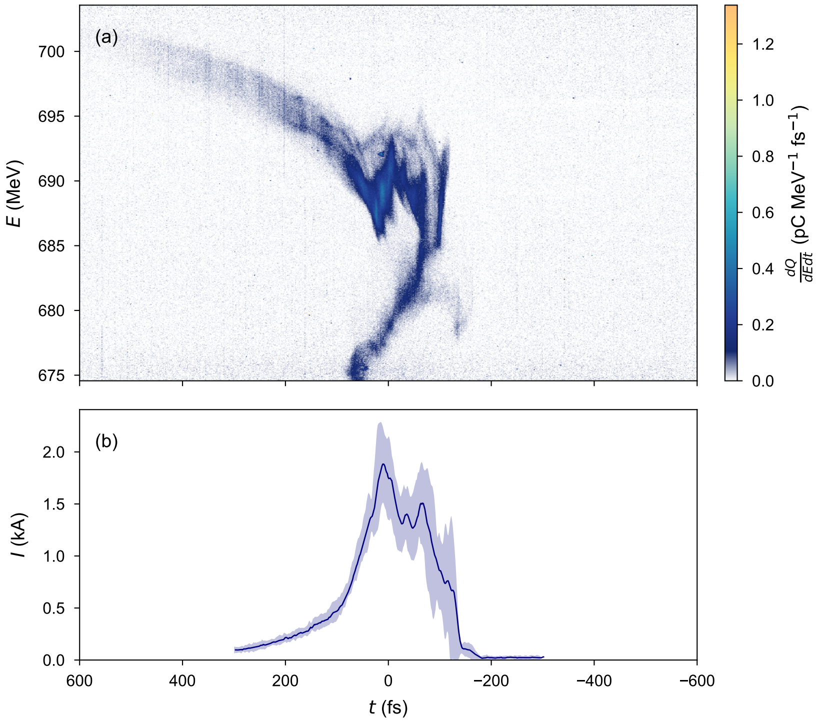

Figure 5.2: Driver longitudinal phase space characterisation at high compression. (a) Example of the longitudinal phase space of the drive beam, measured on the TDS at a ACC139 chirp setting of $-4.38\,\mathrm{m}^{-1}$. (b) Shot-averaged current-profile of the driver at the same setting as in (a) (navy line), with the shaded band representing the shot-to-shot rms variation of the profile. Data displayed here was acquired during the June 2021 experimental run.

An example of a compression scan performed in the June 2021 experimental run by varying the ACC139 chirp is presented in Figure 5.3. Figure 5.3 (a) shows the current-profiles measured by the TDS for different ACC139 chirp values. As shown in Figure 5.3 (b), the average rms duration measured by the TDS (navy points) decreases for stronger compressions (more negative chirps), reaching a minimum value of 77 fs at a chirp of $-4.44\,\mathrm{m}^{-1}$ before increasing again. This minimum point is likely that at which the bunch is fully compressed, with either side of this value representing under- and over- compressed regimes [218]. The minimum bunch length when fully-compressed is essentially a fundamental limit, ultimately governed by the longitudinal emittance and uncorrelated energy spread of the drive bunch. Before reaching this point however, the effects of CSR are likely to be non-negligible, increasing the emittance of the beam and therefore its spot size at focus. This therefore means a practical limit on the compression may be more restrictive than that of over-compression; in the context of this experiment, this limit is found as the point where increasing the compression begins to negatively impact the injection process (see Section 5.2.3 later).

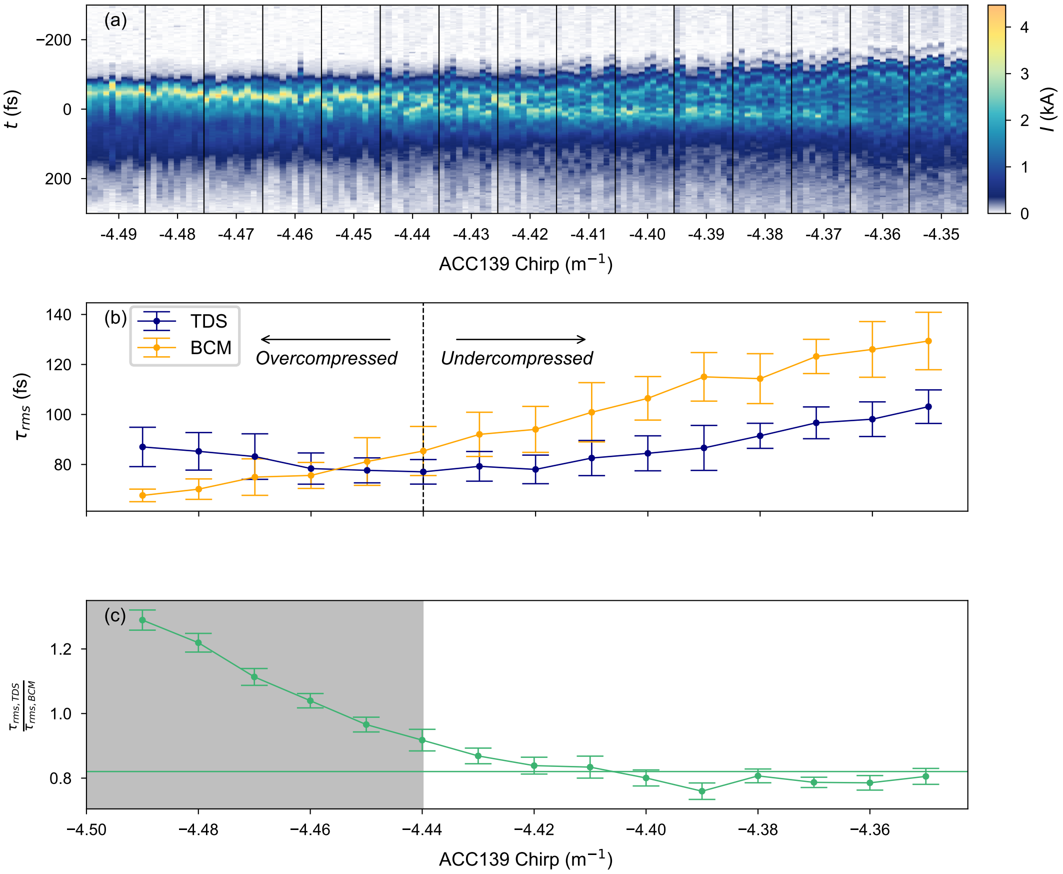

Figure 5.3: Driver bunch length measurements with the TDS and BCM, during a compression scan. (a) Waterfall plot of the driver current-profile, measured by the TDS and grouped by the chirp in the ACC139 accelerator module. (b) Mean driver rms bunch length measured by the TDS and BCM, as a function of the compression, where the errorbars represent the standard deviation. The TDS bunch length reaches a minimum at a chirp setting of -4.44, delineating between over- and undercompressed regimes (black dashed line). (c) Mean ratio between the two bunch length measurements for each compression setting, with each errorbar representing the standard error on the mean. In the undercompressed region, the ratio is approximately constant at an average value of $R_{\tau_{rms},0} \sim 0.86$ (green horizontal line). In the overcompressed regime, the ratio increases (grey shaded region).

In addition to the TDS measurements, rms bunch lengths were simultaneously obtained with a BCM located shortly after BC3 (orange points in Figure 5.3 (b)). Note that the discrepancy between the two diagnostics is likely a result of a non-zero R_{56} in the FLFCOMP section; the shorter bunch length at the TDS when undercompressing in the linac suggests that the bunch is further compressed in the FLASHForward beamline. If however the bunch is beginning to be overcompressed (i.e positive chirp) after BC3, the effect of the same R_{56} would then be to decompress the bunch, such that the length measured by the TDS increases again. Although the absolute values of the rms bunch lengths from the TDS and BCM are therefore not in full agreement, important comparisons can still be drawn between the trends they exhibit. Figure 5.3 (c) shows the ratio of the two bunch length measurements for each compression setting. From this data, one can see that the ratio between the two bunch lengths $R_{\tau_{rms}} = \frac{\tau_{rms,TDS}}{\tau_{rms,BCM}}$ is fairly constant in the undercompressed regime at a mean value of $R_{\tau_{rms},0} \sim 0.86$. In the overcompressed regime this ratio increases, suggesting that the BCM measurement is not as reliable in this region. Nevertheless, whereas TDS measurements of a driver are impossible after hard focusing into a plasma, the BCMs can always provide online bunch length monitoring. Readings from the BCM can therefore be multiplied by $R_{\tau_{rms},0}$ to estimate the true bunch length when TDS measurements are unavailable, with the caveat that this becomes unreliable in the overcompressed regime.

Driver focusing and transverse emittance

With a suitable high compression working point found and characterised on the TDS, the next step in any experimental run was to establish a hard focus at the location of the plasma cell. Starting from the optical arrangement used to transport the beam to the TDS, a setting was typically found that produced a waist beta function in both planes of size $\beta_{0,d} = \mathcal{O}(10\,\mathrm{mm})$ (see example from November 2021 experimental run in Figure 5.4.) As described in Section 3.2.3, an optic was often chosen that also allowed for the position and size of the waist to be tuned approximately independently in each transverse plane, using only the final focusing quadrupoles—an important capability when optimising the internal injection process (see Section 5.2.3 later).

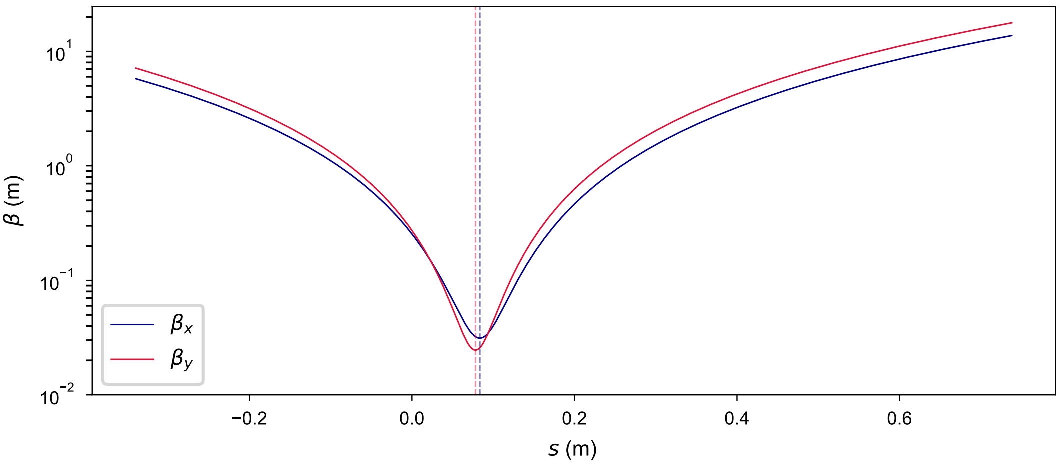

Figure 5.4: Driver beta function setup at the plasma cell location in the internal injection experiment. Measured using the two-BPM tomography method, these beta functions are purely estimates used as a quick way of placing the driver waist at the correct position with a small focus.

With the driver now strongly focused at the position of the plasma cell, its transverse properties could be more accurately characterised via an object plane scan measurement (see Section 2.2.2 on page 18); to ensure the highest resolution, this was performed with the LEMS diagnostic (see Section 3.2.5). Making such measurements before attempting to couple the beam in to the plasma was important for understanding whether the beam density at focus would be enough to drive a strong blowout. They were also vital in quantifying the emittance and therefore brightness of the incoming drive beams. During optimisation of internal injection, these properties were subject to change. Consequently, these measurements were also performed immediately after any characterisation of the witness beams transverse properties, such that the brightness of the driver and witness could be compared.

An example of such a measurement performed at the end of the April 2021 experimental run is presented in Figure 5.5. Figure 5.5 (a) shows the drive beam x-projection on the LEMS diagnostic for each shot in the scan, sorted in order of increasing object plane position $s_{obj,x}$. The horizontal drive beam size $\sigma_{x,d}$ at the location of the object plane was found by dividing the rms widths of the profiles in Figure 5.5 (a) by the calculated magnification M_x. Plotting $\sigma_{x,d}$ against $s_{obj,x}$ (navy points) and fitting Equation 2.38 characterises the driver waist at the location of the plasma (see Figure 5.5 (b)). Here, a tight horizontal waist size of $\sigma_{x,0,d} = (35 \pm 1)$ µm was measured, owing to a small waist beta function $\beta_{x,0,d} = 60$ mm—the measurement of which also corroborates the values obtained via the two-BPM method. The driver however has a large normalised emittance of $\epsilon_{n,x,d} = (21.7 \pm 0.4)\,\mathrm{mm \cdot mrad}$—likely a result of deterioration due to CSR effects when strongly compressing in the linac. Such driver emittances were common for any of the experimental runs presented here. As will be demonstrated, provided that it can still generate a strong enough wakefield, the quality of the witness beams that can be produced via internal injection is largely independent from—and can substantially exceed—that of the drive beam.

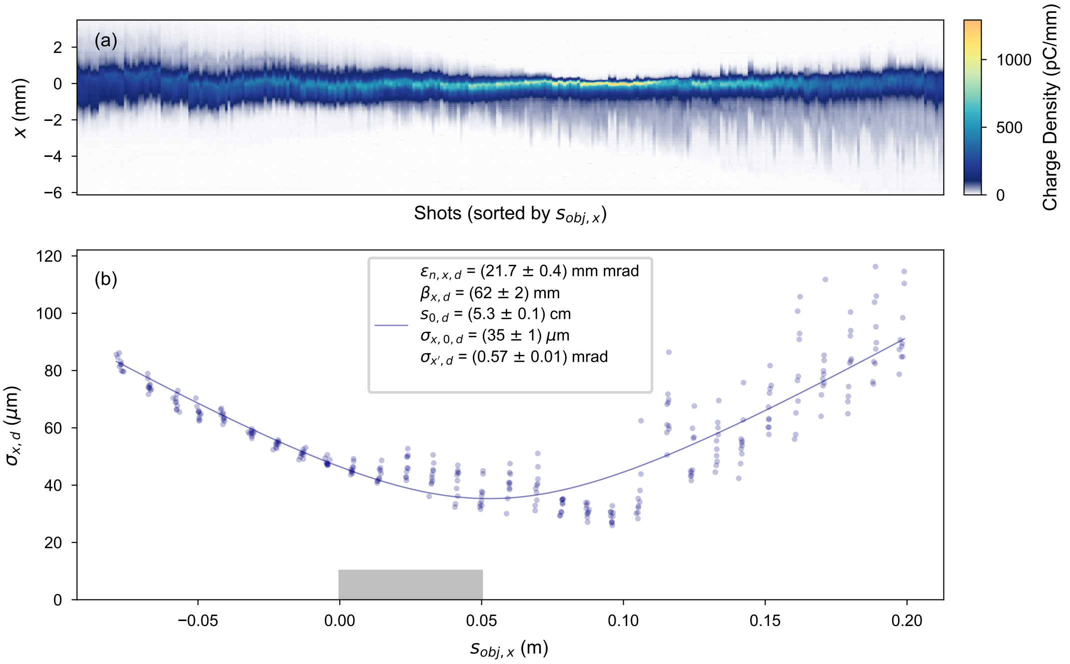

Figure 5.5: Driver horizontal emittance measurement via an object plane scan on the LEMS diagnostic, for an internal injection working point. (a) Waterfall plot of the driver x-projections, sorted in order of ascending object plane position $s_{obj,x}$. (b) Drive beam horizontal widths $\sigma_{x,d}$ versus $s_{obj,x}$ (navy points), to which Equation 2.38 was fit (navy curve); the extracted fit parameters are displayed in the legend. The extent of the plasma cell (grey box) is also shown, with its entrance at $s = 0\,\mathrm{mm}$ and its exit at $s = 50\,\mathrm{mm}$.

5.1.2 Laser setup and characterisation

A key part of preparations for an experimental run was also the setup and characterisation of the longitudinal and transverse laser arms, most of which could be performed before or in parallel to the establishment of the drive beam. The initial stage of this preparation was the activation of the laser system, followed by the alignment of the laser beamline, as described in Section 3.3.2. From here, the lasers were optimised for the purposes of the plasma cathode experiment. A useful tool for gauging the ability of the lasers to ionise the neutral gas mixture again involved the plasma afterglow light: the stronger the ionisation, the greater the light yield [201]. Figure 5.6 illustrates how the plasma afterglow light signal was observed in this context. With the baseplate and its attachments moved clear of the paths of all beams, the chamber was flooded with a low density gas—in this case, the Ar:H_2 mixture. Upon the release of the lasers and ionisation of a region of the neutral gas, the plasma afterglow light signal could be detected with the various cameras located around the chamber. For this purpose, the 'top' (Basler model acA640-120gm) and 'front' (Basler model acA2040-35gm) view cameras were best suited, each with a field-of-view that could in principle capture the light from both laser-generated plasmas (orange box in Figure 5.6). In particular, the top view camera provided a good view of the upstream part of the chamber, and could prove particularly useful in overlapping the two plasmas in the x-direction, forming a cross where they met perpendicularly. Importantly, this camera was able to view a large section of the length of the longitudinal plasma (rightmost inset in Figure 5.6). On the other hand, the front view camera was best suited to observe the transverse plasma (leftmost inset in Figure 5.6) as it could image the whole of the $x - y$ plane that the column could potentially move within. Note that for all cameras observing the light emission from laser-generated plasmas in this chapter, the camera exposure window was set to begin $\mathcal{O}(10\,\mathrm{ns})$ after the arrival of the laser, such that any stray laser light—or indeed other effects that occur on the timescale of the laser propagation e.g. Thomson scattering—would be gated-out temporally. The exposure duration used was then dependent on the camera and experimental run, but in all cases lasted much longer than the lifetime of the plasma. Additionally, as with the camera used in Chapter 4, no spectral filters were used in front of these cameras; the spectral response was purely determined by the camera chip, spanning $\sim 400\,\mathrm{nm}$–$1000\,\mathrm{nm}$ with a maximum between $\sim 500\,\mathrm{nm}$–$600\,\mathrm{nm}$ for all cameras. Ultimately, the appearance and relative changes in this signal were important in this context, rather than its absolute values. The optimisation that this signal enabled, along with measurements of the temporal and spatial profiles for each laser arm, forms the basis of the next sections.

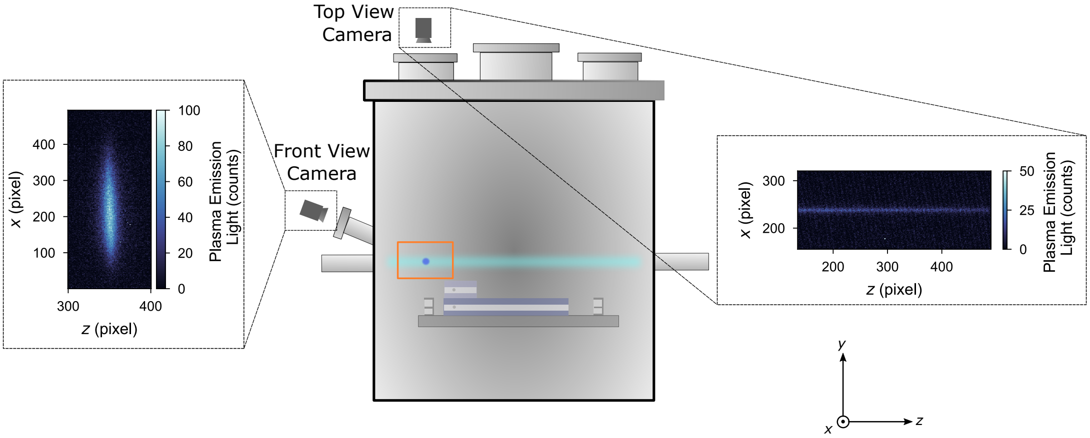

Figure 5.6: Setup for measuring plasma afterglow light signals for driver-laser spatio-temporal alignment. Center: A side view of the plasma chamber is shown, with a sketch of the longitudinal (light blue column, axis in z direction) and transverse (dark blue column, axis in x direction) plasmas generated by their respective lasers in a low density gas; the baseplate and its attachments are moved clear of all beams in this configuration. The front and top view cameras each have a field-of-view (orange box) that allows observations of the emission light signals from the longitudinal and transverse plasmas. Left: an example image from the front view camera of the light emitted by the transverse laser-generated plasma. Right: an example image from the top view camera of the light emitted by the longitudinal laser-generated plasma.

Laser pulse compression and temporal profile

Tuning of the pulse length of the lasers was typically achieved through adjustments of the second order spectral phase using the Dazzler (see Section 3.3.1 on page 59). Optimisation of the pulse length could in principle be performed by observing the plasma light generated by either of the laser arms since they are coupled; owing to higher light yield, the longitudinal signal was most commonly used. Starting at the lowest energy available in the high power mode, the longitudinal laser was released into the flooded plasma chamber. The laser energy was then gradually increased until a plasma was ignited and its light visible on the diagnostic cameras. From there, the second order spectral phase was adjusted until the plasma light yield was maximised—indicating the optimal pulse length that would be used in internal injection experiments.

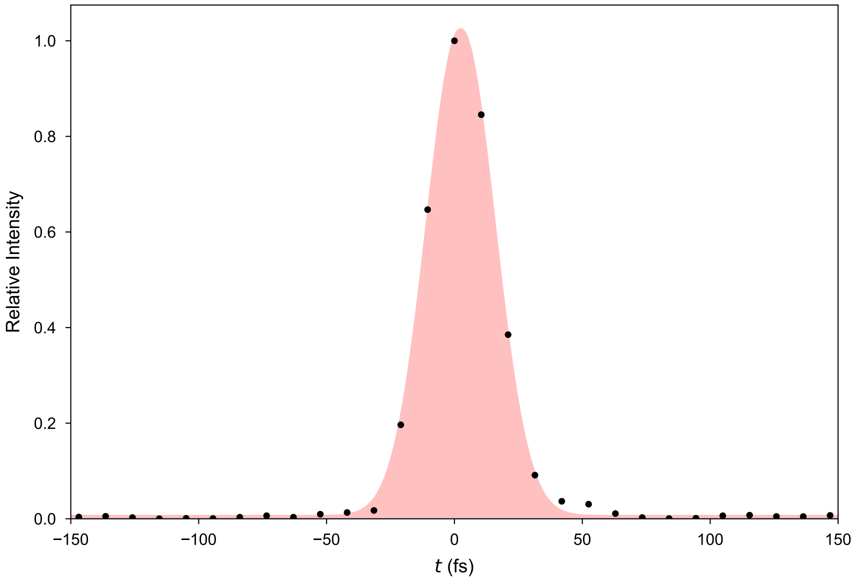

Figure 5.7: Laser pulse length measurement at optimal compression settings. An example Wizzler measurement for a single shot is shown (black points) with a gaussian fitted to the main part of the pulse (red shaded curve).

Assuming that B-integral effects e.g. in the windows were not too severe, the pulse length could be considered to be similar in both the transverse and longitudinal beamlines. A Wizzler pulse length measurement performed immediately after the first periscope in the transverse beamline (see Figure 3.12) therefore served as an estimate of the pulse duration for both lasers (see Section 3.3.1 on page 59 for details of the Wizzler diagnostic). Configuring the Dazzler with the same settings used to produce the maximum ionisation signal, a total of 101 shots were measured with this diagnostic. Gaussian profiles were then fit to the main pulse of each shot, an example of which is presented in Figure 5.7. From each fit the intensity FWHM duration was extracted, with a shot-averaged value of (32.6 ± 0.2) fs, where the error represents the standard deviation. As expected in the absence of any major changes to the compressor, the same Dazzler settings were found to be optimal for all the experimental runs presented here. This measurement is therefore assumed to provide a reasonable estimate of the laser pulse length that was used for all the experimental data presented in this chapter.

Laser focusing and spatial profiles

As well as information about the duration of the pulses, characterisation of the intensity distribution also required measurements of the spatial profile and energy content of the lasers. To estimate the latter, the mean pulse energy and its jitter at the exit of MPA2 were measured regularly. If such a measurement was unavailable at the start of an experimental run, the values recorded most recently were used e.g. after routine maintenance. In any case, combined with an estimate of the transmission to each of the laser arms, these measurements provided a means to calculate the energy for any laser power setting in a given experimental run. An estimate of the energy throughput began with a measurement in the transverse laser beamline. At a laser power of 100% in low power mode, a pulse energy of 102.2 µJ was measured by focusing the beam down onto an energy meter immediately before the M212 mirror (see Figure 3.12 on page 62). Accounting for the 4 mirrors that follow (3 silver and 1 gold with approximate reflectances of 96%), the energy in the chamber was estimated as $102.2\,\mu\mathrm{J} \times 0.96^4 = 86.8\,\mu\mathrm{J}$; the 93% reflectances of the two surfaces of the fused silica chamber window were assumed to be effectively mimicked by those of the lens that was used in the measurement. From this, the transverse laser energy in the plasma chamber at a 100% high power setting was found via multiplication by 170—the attenuation factor between low and high power mode. The maximum energy available at the interaction point of the transverse laser was in this instance therefore $86.8\,\mu\mathrm{J} \times 170 = 14.8\,\mathrm{mJ}$. To calculate this energy at any other occasion, this value was scaled by the ratio of the MPA2 energy, measured here as $(775.0 \pm 9.3)\,\mathrm{mJ}$, to that measured at the other given time. Adjustments could then be made for the amount of attenuation used for a given setting.

With a means of estimating its pulse energy for given settings, the intensity distribution of the transverse laser could be measured. This was achieved by direct measurement of the transverse laser on the upstream alignment camera after propagation straight through the chamber. Figure 5.8 shows how this was done, along with an example spot from the June 2021 experimental run. The translation stage position of the upstream alignment camera was generally set such that the waist of the beam was imaged i.e. where the spot size was at a minimum and the intensity highest. By dividing the summed pixel count for a given image by the estimated energy in the transverse laser pulse at that time, the amount of energy measured for each pixel was calculated—with this calibration of course only valid for the same camera settings, including any camera filtering. From this, the flux was calculated by dividing by the area that one pixel represents in the object plane. Assuming a gaussian temporal profile, the intensity of the laser could then be found from the previously detailed pulse length measurements.

Generally, transverse laser peak intensities of approximately $5 \times 10^{15}\,\mathrm{Wcm}^{-2}$ were measured for the power settings typically used in plasma cathode experiments , corresponding to a pulse energy of 7.4 mJ. The core widths of each spot, estimated by taking the FWHM of gaussians fitted to the central lineout along each axis, were found to be consistently between $40-50\,\mu\mathrm{m}$, with some asymmetry in the y and z directions. Limited improvements to these values could be achieved through further optimisation of the OAP, with efforts made to produce a tight focus with as little astigmatism as possible. The amount of light emitted from the transverse plasma column could also be used as an indicator of the ionising power of this laser, and thus the quality of the focal spot. This being said, in later shifts the quality of this spot did appear to deteriorate—possibly as a result of optical damage in the transverse laser beamline.

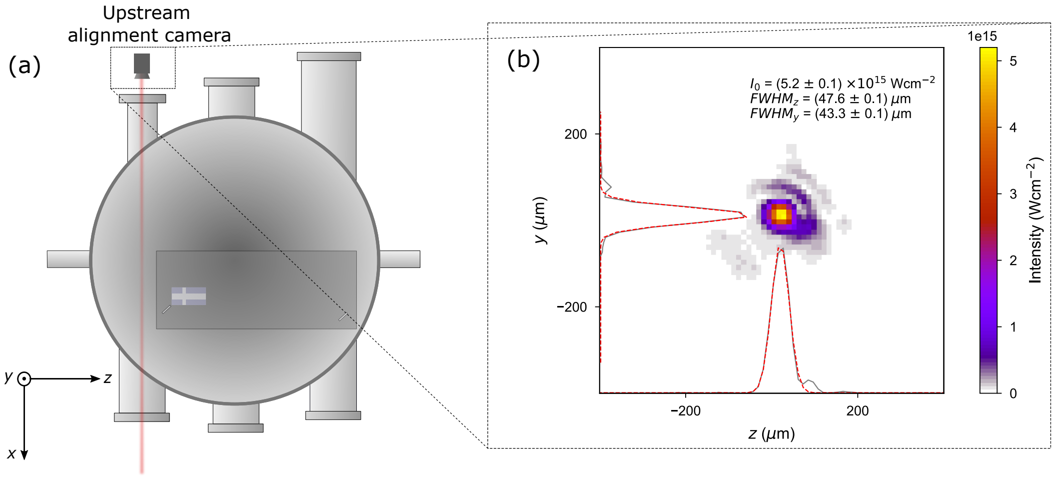

Figure 5.8: Transverse laser intensity characterisation. (a) Plasma chamber configuration for measurements of the transverse laser spot. (b) Example of the transverse laser intensity distribution at focus from the June 2021 experimental run. Plotted along each axis are also the relative intensities along a central lineout in each direction (black curves), with a gaussian fitted to each (red curves). The FWHM of these fits, plus the maximum intensity I_0 of the spot, are included within the image, where the values and their errors represent the mean and standard deviation calculated over multiple shots.

Reflecting the longitudinal laser off the upstream OTR screen into the upstream alignment camera also allowed the intensity of this beam to be characterised. Figure 5.9 shows the chamber configuration for this, as well as an example longitudinal laser spot for the June 2021 experimental run. Provided that the same camera settings were used as for the measurements of the transverse laser on that occasion, the same calibration of energy per pixel count could be applied and the intensity calculated in the same way. For simplicity, the OTR screen is assumed to have been 100% reflective. At the same attenuation setting that produces the transverse laser intensity shown in Figure 5.8 (b), the pulse energy in the longitudinal laser was $\sim 29\,\mathrm{mJ}$. However, with larger core spot widths of $\sim 300 - 450\,\mu\mathrm{m}$ in each direction (see Figure 5.9 for exact values), the intensity of the longitudinal laser ($\sim 3 \times 10^{14}\,\mathrm{Wcm}^{-2}$) was around and order of magnitude lower than that of the transverse laser. The reason for the apparently poor transport of energy in the longitudinal laser arm remains unknown; options for optimising the quality of the longitudinal laser spot were limited, with no obvious adjustments that could be made without intrusively opening the longitudinal laser beamline. On the other hand, further optimisation of this spot was not found to be a priority, with its performance in generating the longitudinal plasma channel found to be quite consistent, albeit with a low intensity.

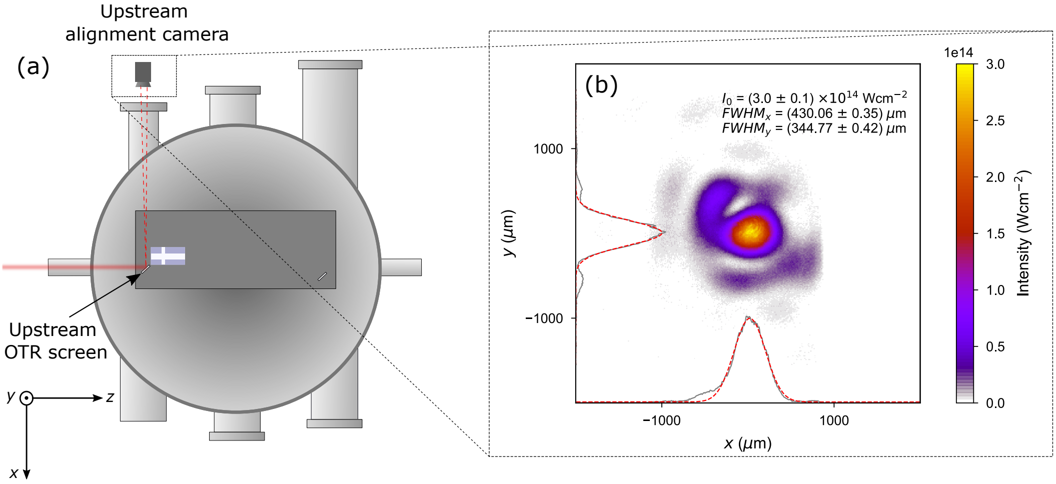

Figure 5.9: Longitudinal laser intensity characterisation. (a) Plasma chamber configuration for measurements of the longitudinal laser spot. (b) Example of the longitudinal laser intensity distribution at focus from the June 2021 experimental run. Plotted along each axis are also the relative intensities along a central lineout in each direction (black curves), with a gaussian fitted to each (red curves). The FWHM of these fits, plus the maximum intensity I_0 of the spots, are included within the image, where the values and their errors represent the mean and standard deviation calculated over multiple shots.

5.1.3 Plasma generation via field ionisation

Sufficiently accurate density measurements of the laser-generated plasma were unfortunately unavailable in this experiment. Simulations of the neutral gas profile for a typical working point were instead performed that then served as an input for ionisation simulations involving the longitudinal and transverse lasers. The results of these simulations are then combined to model the complete plasma profile from experiments.

Neutral gas profile

One dimensional simulations of the neutral gas profile along the axis of the 50 mm-long capillary were performed in the Computational Fluid Dynamics code ANSYS. The input parameters for these reflected those typical of the experiment; a gas of pure Argon was modelled, with a maximum pressure in the capillary of 46 mbar. Note that in experiment, a mixture of 97% argon with 3% molecular hydrogen was used in order to enable density measurements via Optical Emission Spectroscopy [66]. However, the small amount of hydrogen is not expected to change the obtained neutral gas profile by a noticeable amount.

Figure 5.10 (a) shows the density profile of the neutral gas along the axis of the capillary. As shown, the density has a maximum of $3.57 \times 10^{17}\,\mathrm{cm}^{-3}$, with a slight dip down to a value of $3.52 \times 10^{17}\,\mathrm{cm}^{-3}$ at a position of $z = 20\,\mathrm{mm}$, where the 1.5 mm-diameter transverse hole is located. For the purposes of transverse laser ionisation simulations, the density profile across the width of the main capillary at this point is modelled as a trapezoid, with flattop density equal to that at $z = 20\,\mathrm{mm}$ (see inset Figure 5.10(b)). The flattop length of this profile is assumed to be equal to the diameter of the capillary, with linear up- and downramps of length twice that of the transverse hole diameter as a rule-of-thumb estimate. Towards the entrance and exits of the capillary, the density begins to taper off as the gas flows out into vacuum, forming ramps that extend beyond the length of the capillary. After $\sim 10\,\mathrm{mm}$ beyond the capillary extent, the density already drops by approximately 2 orders of magnitude (see inset Figure 5.10(c)). Although only a low density, this gas can still be ionised by the long focal length longitudinal laser and produce a plasma that the beams interact with. This is particularly important in the region beyond the exit of the capillary, where small amounts of plasma can still have a noticeable focusing effect on injected witness beams. In order to simulate this low density gas, the simulation domain is extended as far as possible beyond the ends of capillary, until the flow becomes turbulent and the results no longer reliable.

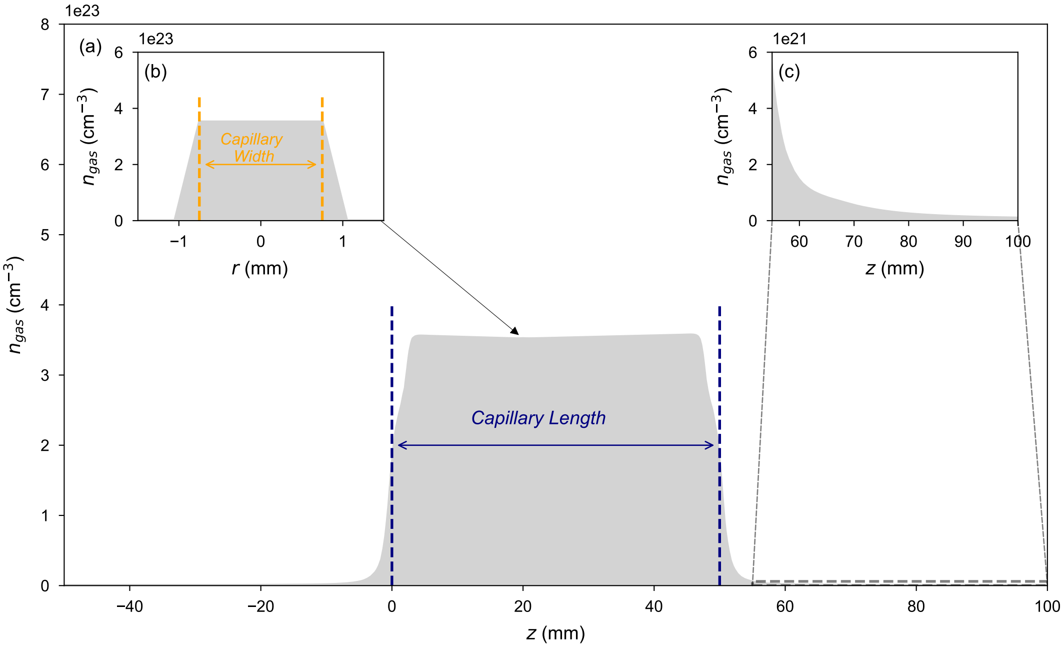

Figure 5.10: Simulations of the neutral gas profile in the 50 mm capillary for internal injection experiments. (a) Neutral gas density profile along axis of the 50 mm capillary. (b) Estimated density profile as a function of the transverse position *r* within the capillary, at the location of the transverse hole. (c) Low density gas ($\sim 1 \times 10^{15}$ cm^{-3}) beyond the exit of the capillary. ANSYS fluid flow simulations were performed by T. Parikh.

Longitudinal laser ionisation

In order to model the shape of the main accelerating plasma channel, ionisation simulations involving the longitudinal laser were performed in the PIC code FBPIC [151]. The gas initialised in the simulation follows the same profile modelled in Section 5.1.3, with only a longitudinal dependence. To accurately model the experiment, a gas of 97% Ar and 3% H_2 by volume is used. As discussed in Section 2.3.1, the hydrogen is approximated in these simulations as fully dissociated atoms, with an ionisation energy of 15.4 eV. Accounting for the dissociated hydrogen also slightly changes the gas mixture. In terms of atoms, the ratio is 97:6 Ar:H, and the *atomic* number density is $1.03 \times n_{neutral}$, where $n_{neutral}$ is the gas density obtained in Section 5.1.3.

The simulated longitudinal laser pulse was reflective of the optimal working point found in the June 2021 experimental run (see Section 5.2.3). A linearly polarised, 800 nm-wavelength gaussian beam with a 35 fs FWHM pulse length, peak intensity of 3×10^{14} Wcm^{-2} and a radially symmetric spatial distribution of FWHM width 400 µm was initialised 15 mm upstream of the capillary entrance, at which point the low density gas at the capillary entrance was made to decrease linearly to zero such that the laser began in vacuum. Owing to the long rayleigh length of this beam, the laser is assumed to be initially at focus and in the absence of plasma does not diverge significantly over the propagation distances modelled here. A co-propagating window of radius 600 µm and length 45 µm with open boundaries was used—enough to contain the pulse in both dimensions for the duration of the simulation. The simulation domain was split into (120×1200) cells in (r, z), with 2 azimuthal modes used. For the neutral gas, the number of macroparticles per cell in (r, z, θ) was $(4 \times 1 \times 4)$.

Figure 5.11 demonstrates the evolution of the laser pulse and its ability to ionise the neutral gas in the simulation. Initially, the laser pulse enters the gas and has enough intensity to partially ionise the Ar and H_2 species (see Figure 5.11 (a)). The resulting plasma is non-uniform, with density gradients in the radial direction that cause parts of the pulse to begin to diffract—an effect referred to as ionisation defocusing (see Figure 5.11 (b)) [57]. An obvious consequence of this is that the intensity at the centre of the pulse begins to drop, such that the fraction of the neutral gas that is ionised decreases along the length of the capillary. In order to create a map of the resulting density, the plasma electron density field for each species is averaged longitudinally over a 20 µm window at the back of the simulation box for every dump. This yields a series of radial density distributions corresponding to different distances along the resulting plasma. Combining these distributions produces the density profile in (r, z) shown in Figure 5.11 (c). Whilst the plasma density is initially higher on axis, as the energy in the laser pulse begins to move outwards, the density gradually becomes lower in the centre, with an almost parabolic radial profile towards the end of the capillary. A density lineout along $r = 0$ (orange dashed line, Figure 5.11) shows the ionisation fractions in more detail, which is displayed in Figure 5.11 (d). After reaching a maximum of 0.44 initially, the ionisation fraction of argon gradually decreases with propagation, with only a fraction of 0.05 being ionised by the end of the capillary. With its slightly lower ionisation energy, the fraction of H_2 that is ionised is marginally larger but has only a small contribution to the overall plasma profile due to the hydrogen only constituting 6% of the available atoms. As the laser leaves the capillary, it nevertheless still has enough intensity to ionise some of the residual gas that exists in the transition to vacuum—albeit only in small amounts ($< 1 \times 10^{14}\,\mathrm{cm}^{-3}$ for $z > 60\,\mathrm{mm}$).

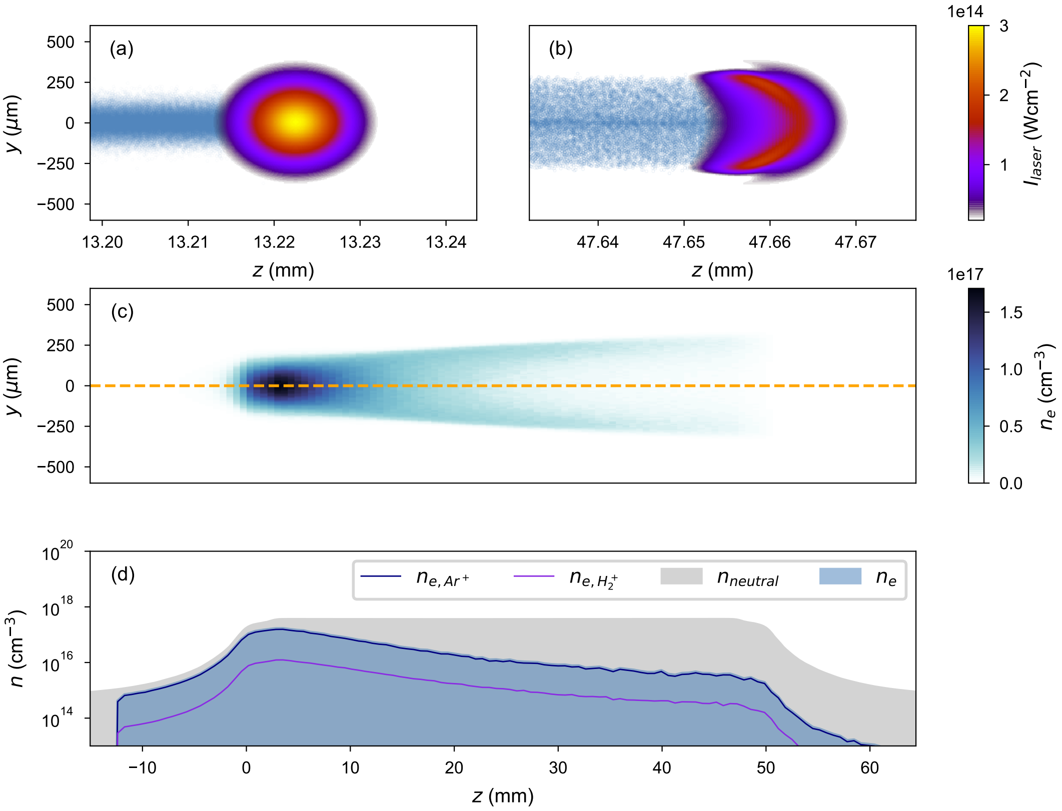

Figure 5.11: Ionisation simulations with the longitudinal laser. (a) Intensity distribution of the laser pulse at a distance of 13.22 mm into the capillary. The distribution is initially gaussian, with electron macroparticles released by ionisation shown as light blue dots. (b) Intensity distribution of the laser pulse at a distance of 47.66 mm into the capillary. As the pulse interacts with the generated plasma, the density gradients cause energy to be defocused away from the centre of the intensity distribution. (c) Density distribution of the released plasma electrons as a function of (r, z). As the laser undergoes ionisation defocusing, plasma is generated at larger radii, with less plasma generated on-axis (orange dashed line). (d) On-axis ionisation fractions of argon and hydrogen atoms. The laser only partially ionises hydrogen and the first level of argon, with the ionisation fraction decreasing along the direction of propagation.

Due to this ionisation defocusing, the resulting plasma profile is a non-linear function of the laser intensity and neutral gas density. For example, whereas a higher neutral gas density exacerbates the defocusing of the laser pulse, the reduced ionisation fraction is to some degree countered by more atoms being available to be ionised. An obvious consequence of ionisation defocusing is that it limits the amount by which a longer capillary will be useful; at some point, the laser will have defocused so much that it is no longer able to ionise any more neutral gas. Ultimately, to fully understand the ramifications of this effect on the experiment, full PWFA simulations are needed, with the resulting profile in this section used as the main plasma channel.

Transverse laser ionisation

Similar simulations as those outlined for the longitudinal laser were also performed for the transverse laser. In this case, the target neutral gas followed the estimated profile shown in Figure 5.10 (b), i.e. the neutral gas profile that the transverse laser propagates through as it crosses the capillary diameter. For simplicity, the gas sampled by the laser is assumed to be completely neutral. Identical gas mixture considerations as in Section 5.1.3 also apply here. The parameters of the transverse laser are again derived from the June 2021 experimental run; a 800 nm-wavelength, linearly polarised gaussian pulse with a FWHM length of 37.2 fs, peak intensity of $5.2 \times 10^{15}\,\mathrm{Wcm}^{-2}$ and a radially symmetric spatial distribution of FWHM width 45 µm was simulated. Starting in vacuum, the pulse is initially converging, with a focal plane located at the centre of the capillary diameter. The co-propagating simulation box had a radius of 150 µm and a length of 60 µm, with open boundary conditions. (272×2400) cells in (r, z) were used, with 2 azimuthal modes. Again, for the neutral gas, the number of macroparticles per cell in (r, z, θ) was $(4 \times 1 \times 4)$.

Figure 5.12 illustrates the ionising behaviour of the transverse laser as it propagates across the diameter of the capillary. Similar to the longitudinal case, the intensity distribution is initially gaussian but starts to experience the effects of ionisation defocusing as it generates and interacts with the plasma (see Figure 5.12 (a-b)). However, due to the relatively short amount of gas that this pulse traverses, the consequences of this effect are not as severe as in Figure 5.11. Figure 5.12 (c) shows the resulting plasma electron density from the transverse laser, constructed in the same way as that in Figure 5.11 (c). By the end of the neutral gas, the width of this plasma column begins to increase, with the central density decreasing slightly—both symptoms of ionisation defocusing. Lineouts of the relative contributions of each ionisation level to the electron density profile at $x = 0$ (i.e. the centre of the capillary diameter) are displayed at the bottom of Figure 5.12. Owing to the relatively high intensity of this pulse, the argon is fully ionised up to a level of Ar^{3+} across a large part of the transverse laser's radial extent. This is also true for hydrogen, although as discussed before, the relatively low volume of this species means that it only contributes a small amount of electrons to the overall density. At the very centre of the pulse, the intensity is also high enough to partially ionise the 4th level of Argon. However, due to ionisation defocusing, the degree to which this happens also relies on the exact position along the propagation axis x. Regardless, the ability of this laser to locally ionise multiple levels of argon in principle makes it ideal for generating a sharp downramp for the purpose of internal injection into a PWFA, but this also depends on the height of this ramp relative to the base density of the main accelerating channel.

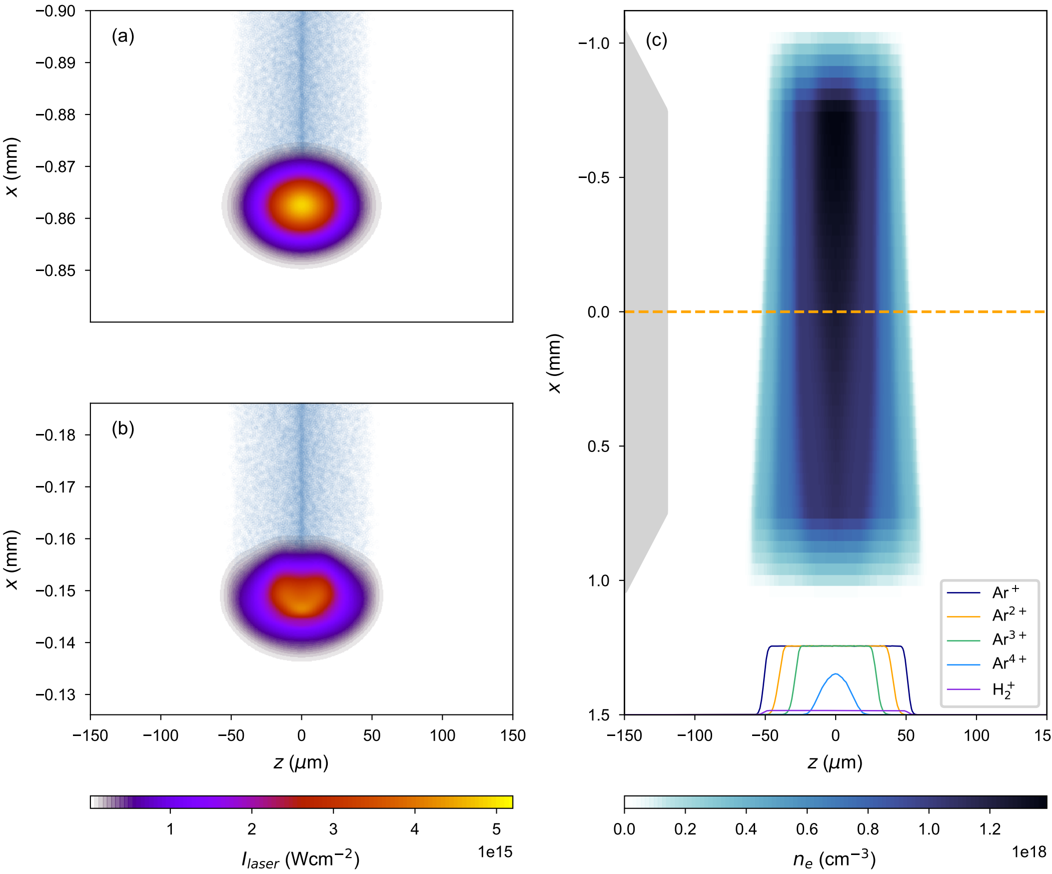

Figure 5.12: Ionisation simulations with the transverse laser. (a) Intensity distribution of the transverse laser pulse at a position $x = -0.86$ mm relative to the centre of the capillary. The distribution is initially gaussian, with electron macroparticles released by ionisation shown as light blue dots. (b) Intensity distribution of the laser pulse at a position $x = 0.90$ mm relative to the centre of the capillary. Ionisation defocusing begins to move intensity away from the centre of the laser spot, out to larger radii. (c) Density distribution of the released plasma electrons as a function of (x, r), alongside the relative neutral gas density (grey filled curve, left). Lineouts of the different relative ionisation fractions of each species at $x = 0$ (orange dashed line) are shown at the bottom of the panel.

Combined plasma profile

In order to model the plasma that the drive beam experiences in the experiment, the plasma profiles created by the longitudinal and transverse lasers (see Figures 5.11 and

5.12 respectively) are combined. The plasma column formed by the transverse laser is placed at the position of the transverse hole in the capillary i.e. $z = 20\,\text{mm}$. For simplicity, the column is made uniform along the x-axis (i.e. the direction of transverse laser propagation), with ionisation fractions equal to those at $x = 0$ in Figure 5.12. The overall plasma shape is found by taking the maximum plasma density of the two profiles at every point in space . This treatment essentially assumes that the generation of one plasma is independent of the other. In the case that the main channel is formed *after* the transverse plasma, the presence of the small density column is unlikely to influence the ionisation from the longitudinal laser. If the longitudinal laser arrives *before* the transverse, the relatively low plasma density that exists at $z = 20\,\text{mm}$ is not expected to dramatically change the formation of the transverse plasma column.

Figure 5.13 shows the results of this combination. In the case that the transverse plasma is vertically aligned with the longitudinal axis of the main plasma channel, the centre of the transverse column sits with its centre at (0 mm , 20 mm) in the (y, z) plane (see Figure 5.13 (a)). The column extends to a radius of $\sim 50\,\mu\text{m}$ about its centre in this plane, meaning the downramp that is generated has finite extent in the y-direction, with some curvature (see inset in Figure 5.13 (a)). In the x-direction (i.e. into the page), the column has infinite extent; the rayleigh length of the laser is much longer than any expected driver misalignment in the x-direction, and the evolution caused by ionisation defocusing is considered small.

By taking a lineout along the longitudinal axis of this combined profile, the relative densities around the transverse plasma position can be compared (see Figure 5.13 (b)). In this region, whereas the longitudinal laser only partially ionises the first level of argon producing a density of $0.06 n_{neutral} = 2.0 \times 10^{16}\,\text{cm}^{-3}$, the transverse laser reaches into the 4th ionisation level of argon, generating a peak density of $\sim 3.5 n_{neutral} = 1.3 \times 10^{18}\,\text{cm}^{-3}$. The peak of the density spike is therefore higher than the surrounding base density by a factor of ~ 65. Combined with the short ramp lengths of this density spike ($\sim 50\,\mu\text{m}$ peak to base), this makes the downramps generated here rather extreme in the context of internal injection into PWFA, where previous studies have been limited to ramp-base height ratios ~ 10. Ultimately, PWFA simulations with the plasma profile shown here are required to understand the consequences of these conditions for downramp injection.

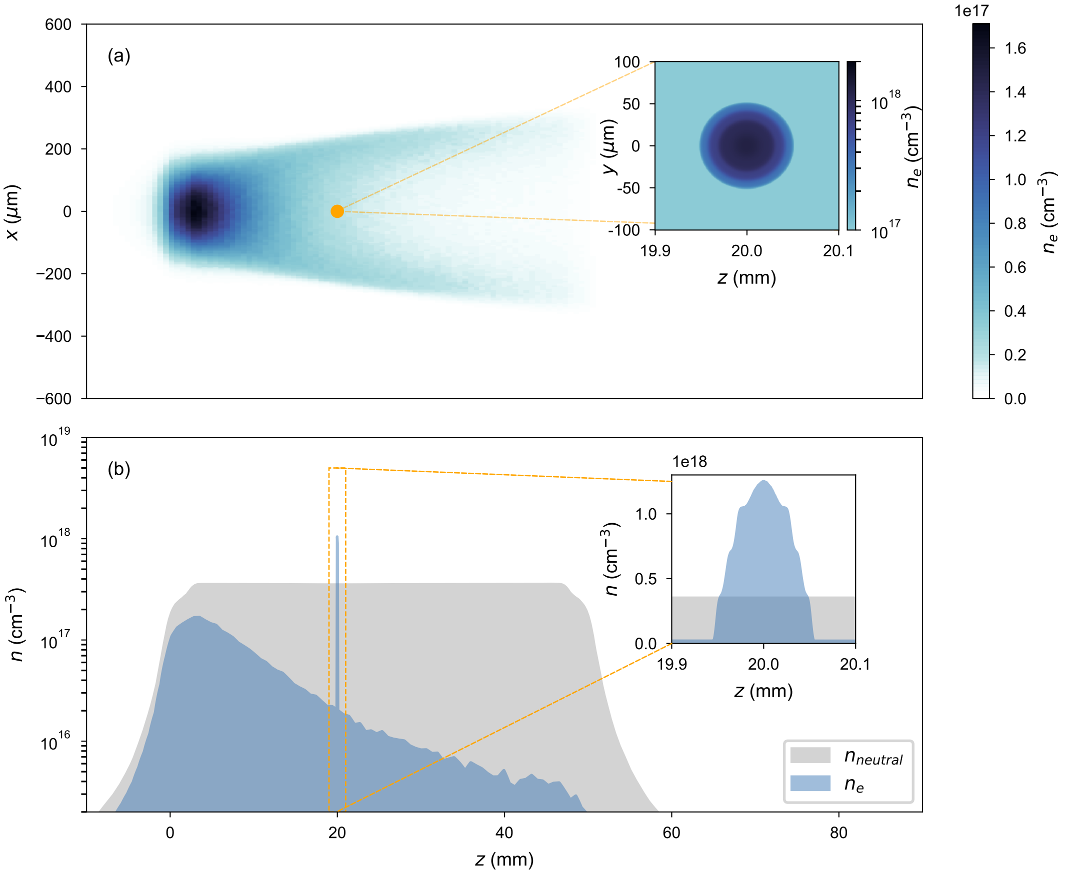

Figure 5.13: Combined plasma profile generated by the longitudinal and transverse lasers. (a) Plasma density in (y, z), from a slice across $x = 0$. The axis of the transverse column intersects perpendicularly to that of the longitudinal channel at the position $(y, z) = (0\,\text{mm}, 20\,\text{mm})$. Inset is the density profile in a 200 μm × 200 μm window around the centre of the transverse plasma (orange dot). (b) Lineout of the combined density profile along the z-axis at $(x, y) = (0, 0)$ (blue shaded profile). For comparison, the neutral gas density is also displayed (grey shaded profile). Inset is a 200 μm-wide section of the density lineout centred on the position of the transverse plasma column.

Other sources of field ionisation

Note that—if strong enough—the electric fields due to the driver and the wake also have the potential to ionise the neutral gas. It is therefore worth briefly considering these additional sources of ionisation which, if present, could in principle lead to changes in the plasma profile and even the unwanted injection of electrons into the plasma wake.

In the case of the driver fields, rough estimates can be made for the maximum transverse fields of a gaussian bunch. From Equation 2.45 and the estimate for the maximum radial field $E_{r,max}$ that follows, $E_{r,max} \approx 0.45\frac{I_p}{2\pi\epsilon_0 c\sigma_r} = \mathcal{O}(5\,\mathrm{GV\,m^{-1}})$ for a peak current $I_p = 2\,\mathrm{kA}$ and radial size of $\sigma_r = \mathcal{O}(10\,\mu\mathrm{m})$. Consulting Figure 2.4, this is far below the required field strength to ionise even the first level of argon ($\sim 60\,\mathrm{GV\,m^{-1}}$); only if the radial beam size were to reduce to $< 1\,\mu\mathrm{m}$ would this field begin to approach the required threshold. In principle, an unmatched drive beam can potentially exhibit 'pinching' of the beam envelope within the plasma, where the transverse size of a beam slice decreases sharply producing an increased radial field strength [219]. Given the large transverse emittance of the drive beam in question however, it is unlikely that pinching extreme enough as to result in ionisation would occur in this case.

As for ionisation due to wakefields [102], simple estimates can be made based on the wave-breaking field (see Equation 2.60). As the simulation of the longitudinal laser ionisation shows (see Figure 5.11), the maximum expected plasma density in the main plasma channel is $\mathcal{O}(10^{17}\mathrm{cm}^{-3})$, resulting in a maximum wakefield strength estimate of $\mathcal{O}(30\,\mathrm{GV\,m^{-1}})$. Like the driver field then, it is unlikely that any ionisation due to wakefields will occur in the context of this experiment; the plasma profile is therefore assumed to arise purely from laser ionisation, resembling one like that presented in Figure 5.13.

5.1.4 Electron and laser beam overlap

Before coupling the focused drive beam and lasers into the plasma cell, the three separate beams first required overlapping in time and space. As explained in the following sections, this was typically achieved in two stages: coarse overlap, followed by fine synchronisation and alignment.

Coarse Overlap

Coarse overlap of the electron and laser beams was primarily achieved with the OTR screens, located at up- and downstream positions on the baseplate in the chamber (see Figure 5.14).

At the injection position, the upstream OTR screen could be positioned to intercept both the driver and longitudinal laser, whilst the transverse laser passed through a gap between the screen and its mounting (see Fig 5.14 (a)). Imaging the plane of this screen, the upstream alignment camera was simultaneously able to observe the focal spots of the drive beam and both lasers. With the waist of the driver set approximately at this longitudinal position, only small adjustments could be made to its focal spot

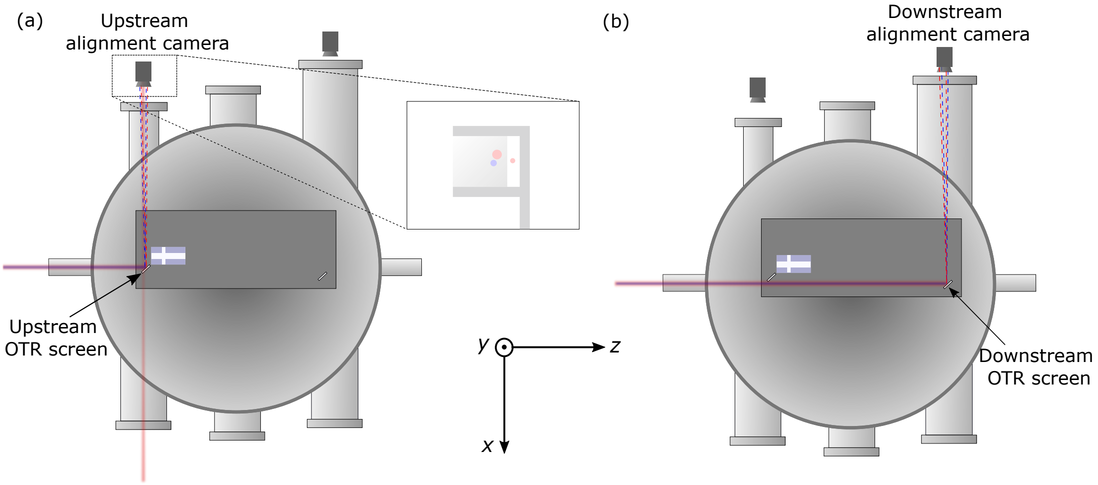

Figure 5.14: Up- and downstream alignment of the driver and laser beams on the OTR screens. (a) Chamber configuration for the upstream alignment of the longitudinal and transverse lasers (red) with the drive beam (blue). (b) Chamber configuration for the downstream alignment of the longitudinal laser with the driver.

position in the $x - y$ plane; therefore, the positions of the laser spots were adjusted in order to achieve spatial overlap with the drive beam. For the longitudinal laser, angular movements of the final spherical mirror were used to directly overlap its spot with that of the drive beam in x and y. The transverse laser was then brought to the same height as the longitudinal laser and driver via vertical translation of the OAP mirror and/or adjusting the angle of the turning mirror.

Whilst the procedure on the upstream alignment camera ensured that the paths of the lasers and drive beam intersected at the injection position, the orbits of the longitudinal laser and the driver would ideally be overlapped for at least the length of the capillary. This was ensured by also viewing the drive and longitudinal laser beams on the downstream OTR, with the downstream alignment camera (see Figure 5.14 (b)). Misalignment at this position could here be corrected by moving the driver, with little effect on its transverse position at the injection point.

Temporally, the laser pulses and drive bunch could be approximately brought together by simply gating the signals on the upstream alignment camera. At this stage, the time-of-arrival of the driver was considered fixed. By first shifting the delay of the the camera until the signal from the drive beam on the OTR disappeared, its timing was found in units of camera delay. The laser pulse timing could then be adjusted in steps of 18.46 ns by pulse-picking in the Booster stage of the laser (see Section 3.3.1 on page 57) until it reached the edge of the upramp of the camera exposure. At this point, the lasers and drive bunch were timed within ~ 100 ns of each other.

Fine alignment and synchronisation using plasma afterglow light enhancement

Although the previous procedure provided a good basis for the overlap of the laser and drive beams, more was required. This particularly applied to the temporal synchronisation; it was necessary—at least initially—to find the relative time-of-arrival $t_{driver} - t_{laser}$ of the beams at around the $\sim$ps level, where t_{driver} and t_{laser} refer to the absolute time-of-arrival of the driver and laser respectively. Control with this granularity was possible, again through more shifting of the pulses that leave the laser oscillator. However, the alignment cameras were not 'fast' enough to be able to provide this kind of time-of-arrival information, and thus a different method was needed. We therefore turn again to the plasma afterglow light technique, relying on the increase in emitted plasma light when an electron beam transfers energy to the plasma [201].

The flooded chamber configuration introduced in Figure 5.6 was again adopted, with the difference now that the drive beam would also be present. Beginning with only the longitudinal plasma being generated, the driver was released. If the drive bunch arrived before the laser ($t_{driver} - t_{laser} < 0$ in Figure 5.15), the driver would not interact with the plasma and the emission light signal thus remained the same as that with only the laser present. However, by bringing the arrival of the laser before that of the driver ($t_{driver} - t_{laser} > 0$), the drive bunch would begin to interact with the plasma and produce a substantial increase in the light emitted once the plasma had decayed ($\sim$ 2 orders of magnitude higher in the region-of-interest shown in Figure 5.6). This dramatic transition occurred at what is defined here as $t_{driver} - t_{laser} = 0$ i.e. the delay at which the longitudinal laser and drive beams had the same time-of-arrival. In practice, this transition was first roughly found by more scanning of the laser timing via pulse-picking (i.e. to a resolution of 18.46 ns), but was then located on the picosecond level by adjusting the delay via movements of the oscillator cavity mirrors. Finely scanning the delay over this region yielded a curve like that in Figure 5.15 (b)—a step-like transition from which the exact definition of $t_{driver} - t_{laser} = 0$ was determined. In this case, this was achieved by fitting an error function to the data around where the increase in light occured (solid light blue line); the transition timing was then defined as that where the integrated light signal was half way between the upper and lower plateaus (black dashed line).

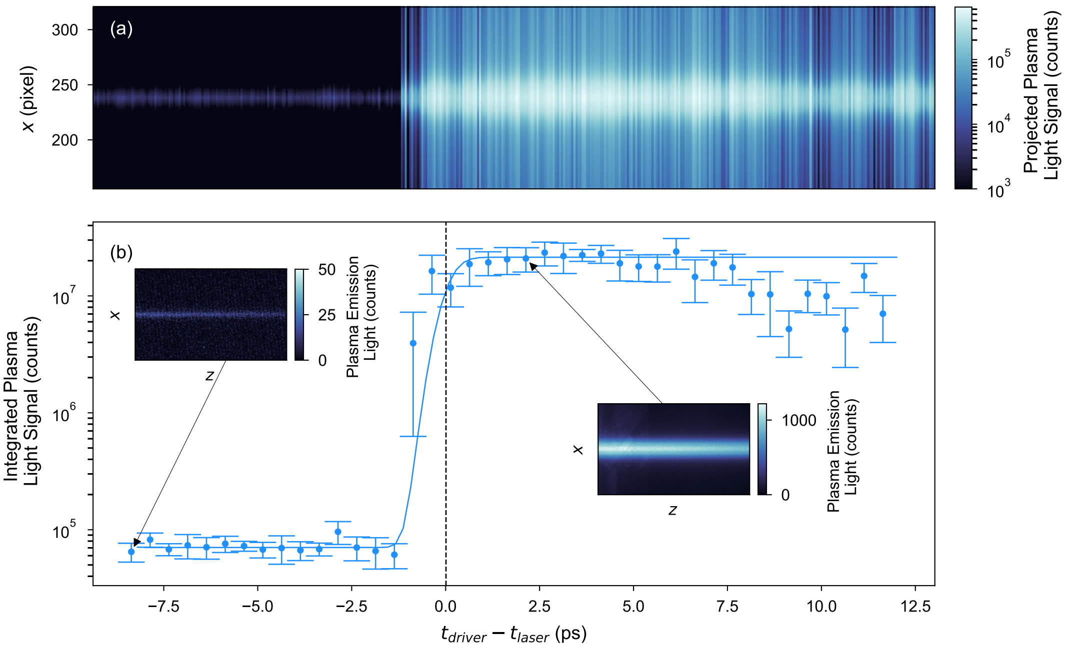

Figure 5.15: Beam-interacted afterglow light signal from the longitudinal plasma as a function of driver—laser delay. (a) Waterfall plot of the longitudinal plasma afterglow light signal measured by the top view camera, projected along the plasma column axis and sorted by the driver—laser delay. (b) Log-lin plot of the average integrated plasma light yield versus the delay between the drive beam and the longitudinal laser, with errorbars that represent the standard deviation. An error function (solid blue line) is fit to data points around the transition time (dashed black line). Inset are example images of the longitudinal plasma light with and without drive beam interaction (right and left respectively).

A perhaps interesting observation in the scan shown in Figure 5.15 was that for $t_{driver} - t_{laser} > 7.5\,\text{ps}$, the light emitted by the plasma post-interaction began to decrease slightly. The reasons for this are unclear; it may be that this was simply a symptom of the spatial overlap found in the previous section deteriorating, and thus changing the interaction.

A similar procedure could then also be applied to the transverse plasma. Figure 5.16 shows the results of a picosecond-level timing scan, this time involving the transverse laser and driver. $t_{driver} - t_{laser} = 0$ is still defined as the delay at which the longitudinal laser and drive bunch coincided in time. Scanning the delay revealed again a sigmoid-like transition between a region of interaction and no-interaction (see Figure 5.16 (b)).

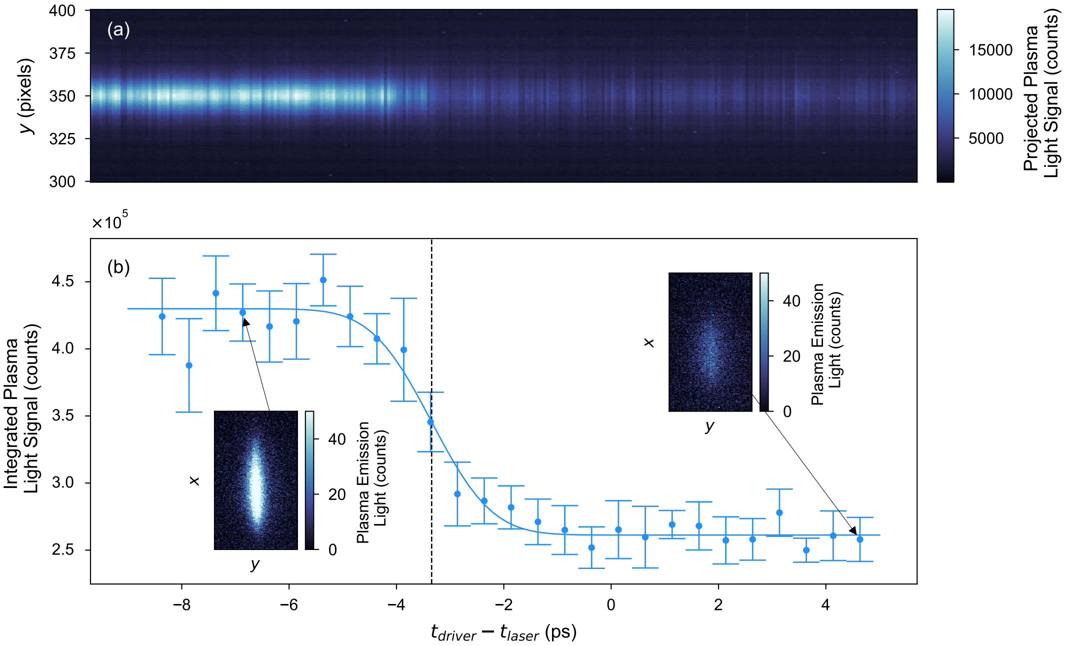

Figure 5.16: Beam-interacted afterglow light signal from the transverse plasma as a function of driver—laser delay. (a) Waterfall plot of the transverse plasma emission light signal measured by the front view camera, projected along the plasma column axis and sorted by the driver-laser delay. (b) The average integrated transverse plasma light yield versus the delay between the drive beam and the transverse laser, with errorbars that represent the standard deviation. An error function fit to these data points gives a transition time (black dashed line) of $t_{driver} - t_{laser} = -3.34\,\text{ps}$, where $t_{driver} - t_{laser} = 0$ remains the timing at which the longitudinal laser coincides with the driver. Inset are example images of the transverse plasma light with and without drive beam interaction (left and right respectively).

Peculiarly however, the case where the transverse plasma and drive beam interact was in this instance characterised by a *lower* plasma light yield than when the driver did not interact. This is in contrast not just to the case of the longitudinal plasma, but also to observations of the same interaction in previous iterations of this experiment (see Appendix B of Ref. [197]). Such behaviour suggests that the plasma was somehow receiving energy from the drive beam, but not transferring this to the surrounding neutral gas in a way that enhanced the plasma light yield. Speculating, the interaction of the strong drive beam with a narrow-width, low-density plasma column is likely to expel the plasma electrons with high-velocities, leaving an exposed column of ions that then expands due to coulomb forces. If the kinetic energies of the plasma column consituents are such that they have lower collisionality compared to the laser-only case, this could explain the reduction in emitted light upon relaxation. Nevertheless, this interaction, albeit unusual, was still one that could be used to find the time-of-arrival between the transverse laser and drive bunch at the picosecond level. Fitting an error function to the data yields a transition time that in this case was $t_{driver} - t_{laser} = -3.34\,\mathrm{ps}$ after that of the longitudinal laser. With this measurement, the time-of-arrival difference between the longitudinal and transverse lasers could then be adjusted with the delay stage described in Section 3.3.4, and if desired measured again with the same method.

With the transverse laser arriving before the drive beam, the spatial overlap between the two could also be further optimised using the plasma afterglow light enhancement method. Figure 5.17 shows the variation in the emission light from the transverse plasma column as the vertical alignment between the transverse laser and electron beam was adjusted. In a similar way to the timing scan in Figure 5.16, when the beam and plasma were better overlapped, the light yield decreased. As shown in Figure 5.17 (b), the amount of emission light is reduced for offsets $|\Delta y_{trans}| \lesssim 400\,\mu\mathrm{m}$. Compared to the widths of the transverse laser and drive beam at focus, the extent of this region where the signal drops suggests some long-range interaction between the driver and plasma column that occurred without the two truly overlapping in space. However, around $|\Delta y_{trans}| \lesssim 20\,\mu\mathrm{m}$ the signal experienced a small increase which was found to be repeatable; a finer scan of this region is shown in the inset of Figure 5.17 (b). This spike in the signal was assumed to be an indication that the actual spatial extent of the driver and plasma column were overlapped. The centre of this spike was then taken as the point where the driver and transverse plasma were optimally aligned in y i.e. $\Delta y_{trans} = 0$. Again, this behaviour was unusual in the context of aligning the transverse laser and drive beam, and has not been observed in previous iterations of this experiment. Regardless, the interaction was still such that it provided a useful figure-of-merit for initial alignment.

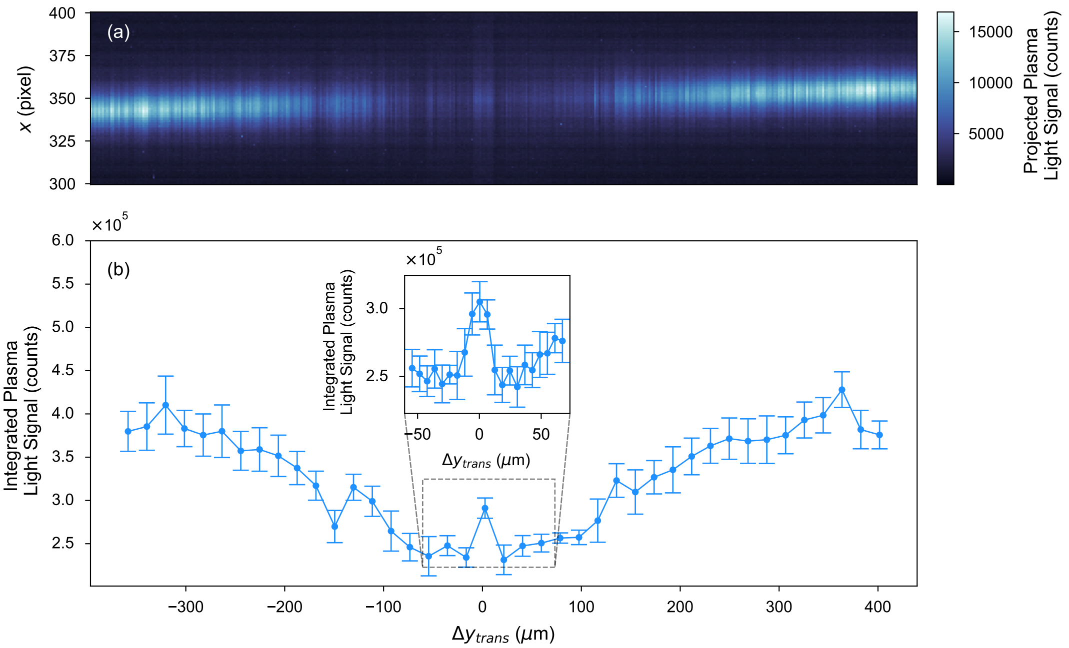

Figure 5.17: Beam-interacted afterglow light signal from the transverse plasma as a function of transverse laser vertical offset. (a) Waterfall plot of the transverse plasma afterglow light signal measured by the front view camera, projected along the plasma column axis and sorted by the transverse laser vertical position. (b) The average integrated transverse plasma light yield versus the vertical offset of the transverse laser, with errorbars that represent the standard deviation. As the driver—plasma overlap grows, the plasma emission light drops, with a small spike in the signal at the centre. Inset is data from a finer scan of this region.

It should be noted that all of the alignment and synchronisation discussed in this section was subject to change due to drifts, particularly in the spatial alignment. Nevertheless, careful alignment was still crucial to have any chance of initially achieving internal injection. Once this injection was achieved, various experimental parameters—including the spatio-temporal alignment—could be adjusted with the goal of optimising the witness beams that were produced.

5.2 Experimental Results

With the different components of the experiment in place, studies of internal injection could begin. In a similar way to Section 5.1, this section introduces these studies in the order that they typically took place in an experimental run.

5.2.1 Driver interaction

Figure 5.18 presents a measurement of the driver interaction with the laser-generated plasma, performed during the June 2021 experimental run. Shown is an example drive beam measured on the high-energy ESPEC screen with an imaging energy of 685 MeV (red dashed line). Note that some clipping of the beam in the imaging quadrupoles, particularly at lower energies, was likely present in this measurement. Nevertheless, the FWHM width of the spectrum (blue filled curve) gives a useful estimate of the maximum deceleration that the driver electrons experienced. Roughly 100 MeV for a typical shot, this energy loss corresponds to a decelerating field of $\sim 2\,\mathrm{GVm}^{-1}$, longitudinally averaged over the 50 mm-long cell. Observe also the *acceleration* of driver charge to energies of $\gtrsim 725$ MeV—suggesting that for this shot and others like it, part of the drive beam was located in the accelerating phase of the plasma wake.

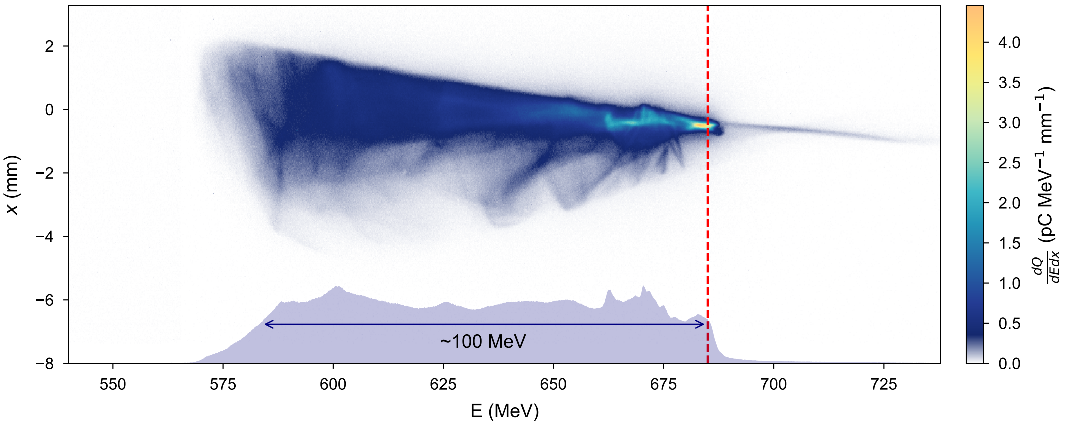

Figure 5.18: Driver deceleration measurement in the internal injection experiment. The spectrum (navy filled curve) was imaged at an energy of 685 MeV (red dashed line). On this occasion, the charge calibration was found by dividing the total number of counts on the screen by the charge measured a BPM just before the ESPEC.

5.2.2 Observation of internally injected charge

In the first instance, releasing the transverse laser often resulted in low-charge witness beams that were only intermittently injected. Figure 5.19 presents witness beams that were measured just after injection was established in the June 2021 experimental run. Here, the longitudinal and transverse lasers were initially set to promptly arrive on the order of picoseconds before the drive beam and the laser system was set to a power setting of 45 %. Any evolution of the injection behaviour—as in Figure 5.19 (a)—was likely a result in a drift in the experimental conditions e.g. the laser—driver spatial alignment. Histograms of the peak spectral densities $(\frac{dQ_i}{dE})_{peak}$, as well as the energy at which these peaks occur E_{peak}, are plotted in Figures 5.19 (b) and (c). Considering only shots with a non-negligible injected charge ($Q_i > 0.1\,\mathrm{pC}$), an average charge $\langle Q_i \rangle = (0.6 \pm 0.6)\,\mathrm{pC}$ was measured, with an average peak spectral density $\langle (dQ_i/dE)_{peak} \rangle = (0.9 \pm 0.9)\,\mathrm{pCMeV^{-1}}$ and average peak energy $\langle E_{peak} \rangle = (28 \pm 2)\,\mathrm{MeV}$. Here, the errors quoted represent the standard deviation of each quantity, indicating large jitters of order ±100 % on the injected charge and peak spectral density. Initially, the focus of optimisation was therefore the stability of the injection process. In addition, the goal would be to maximise the spectral density, with a maximum of only $4.6\,\mathrm{pCMeV^{-1}}$ measured for the best shot in this dataset (see Figure 5.19 (d)).

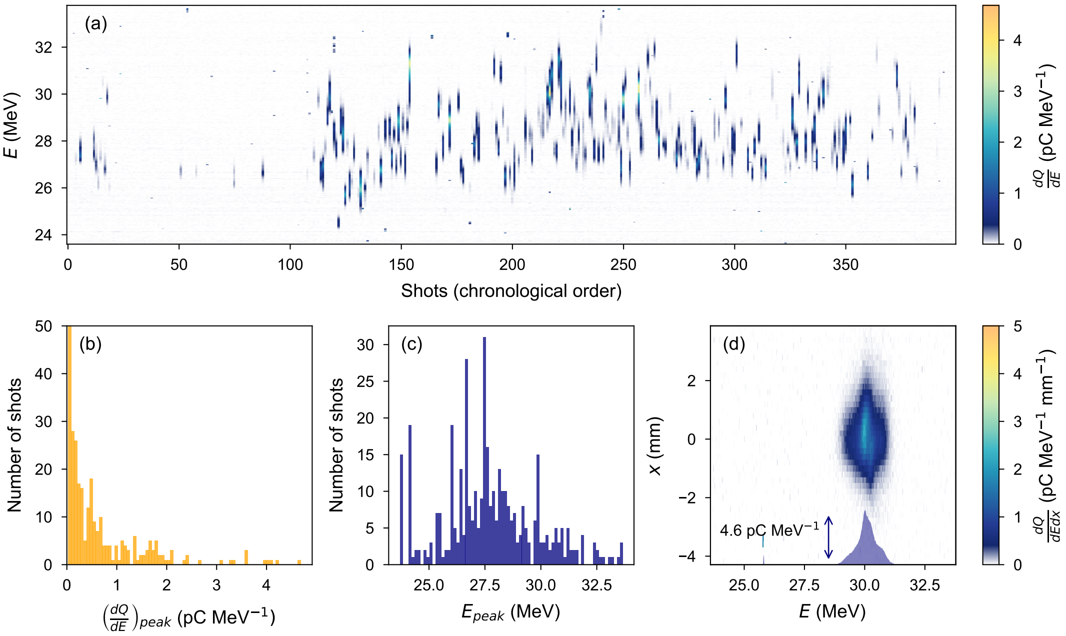

Figure 5.19: Initial injection statistics example. (a) Waterfall plot of the witness beam energy spectra over 400 consecutive shots, measured on the ESPEC and displayed in chronological order. Histograms of the (b) peak spectral densities and (c) peak energies are shown, each using 75 bins. The shot with the highest $(\frac{dQ_i}{dE})_{peak}$ measured for this dataset—$(\frac{dQ_i}{dE})_{peak} = 4.6\,\mathrm{pCMeV^{-1}}$—is plotted in (d), along with its spectrum (navy shaded curve).

5.2.3 Injection optimisation

Using the initially injected beams as a basis for optimisation, vast improvements could be made, specifically to the stability of injection and the charge of the injected bunches; it was generally found that the latter correlated well with the peak spectral density of the bunches and so initially optimising for charge was the first step in producing high brightness beams. The following section therefore explores the impact of the different experimental variables on the measured injected charge, before presenting witness beam data from a fully optimised working point.

Laser Energy

Figure 5.20 investigates the effect of increasing laser energy on the injection by adjusting the attenuation in the main laser beamline with data acquired in the June 2021 experimental run. As the two are coupled, this affects the intensities of both the longitudinal and transverse lasers simultaneously. As evident in the spectra in Figure 5.20 (a), the injection of charge was detected only for peak intensities $\gtrsim 1.8 \times 10^{14}\,\mathrm{W/cm^2}$ and $\gtrsim 3 \times 10^{15}\,\mathrm{W/cm^2}$ in the longitudinal and transverse lasers respectively. Any less than this may have resulted in a downramp that was not steep enough to inject and/or a longitudinal channel with insufficiently high density for acceleration.

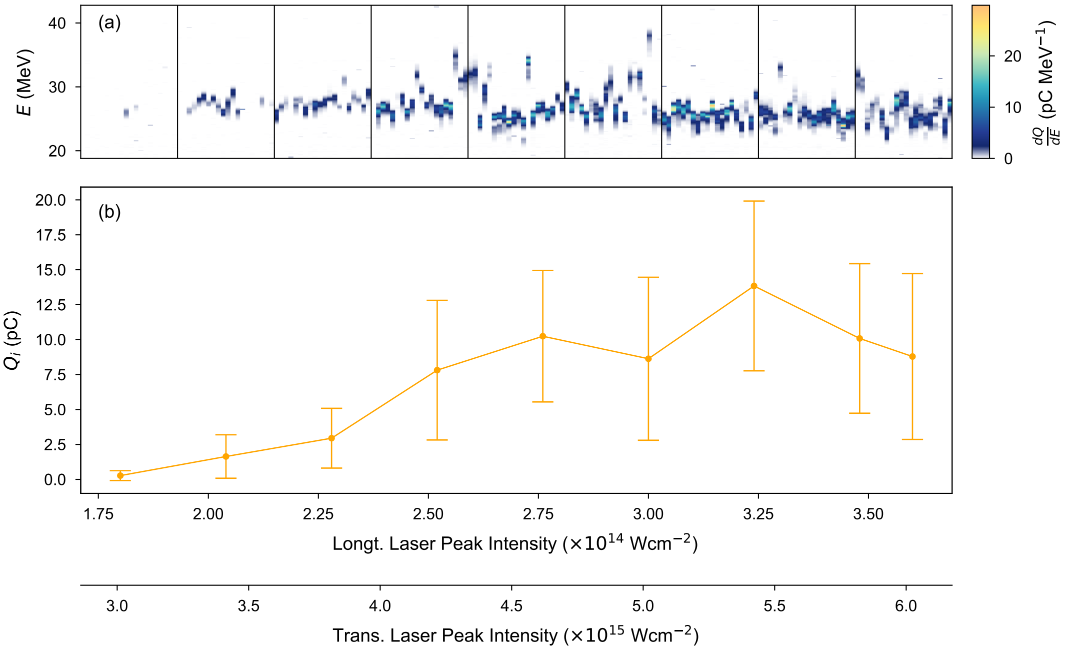

Figure 5.20: Injection dependence on laser intensities. (a) Waterfall plot of the injected witness beam spectra, measured on the ESPEC and grouped by the laser attenuation setting. (b) Mean injected charge detected on the ESPEC plotted against both the longitudinal and transverse laser peak intensities, with errorbars indicating the rms variation of the charge at a given attenuation setting.

The injected charge Q_i roughly increased for peak intensities up to $\sim 2.8 \times 10^{14}\,\mathrm{W/cm^2}$

in the longitudinal laser and $\sim 4.6 \times 10^{15}\,\mathrm{W/cm^2}$ in the transverse laser, beyond which there was no clear improvement. Interestingly, although one may naively expect an increase in plasma density and therefore accelerating field with increasing longitudinal laser intensities, the energy of the injected bunches remained saturated at a level of $\sim 28\,\mathrm{MeV}$. This then is likely a result of the ionisation defocusing effect of the longitudinal laser pulse discussed in Section 5.1.3. Nevertheless, this still corresponds to average accelerating fields of order $\sim 1\,\mathrm{GVm^{-1}}$ in the latter 30 mm of the short plasma cell, and was found to be common for almost all datasets from this iteration of the plasma cathode. Wary to avoid high laser intensities that may deteriorate the laser optics, a peak intensities of $3 \times 10^{14}\,\mathrm{W/cm^2}$ and $5 \times 10^{15}\,\mathrm{W/cm^2}$ in the longitudinal and transverse lasers were generally used, producing a reasonable amount of charge with relatively high stability.

Transverse laser alignment

Although previously aligned via the OTR and plasma afterglow light-based methods, the alignment between the transverse laser and drive beam could typically be further improved. Figure 5.21 shows the variation in the injected beams whilst vertically translating the transverse laser during the June 2021 experimental run.

Intuitively, the injection was optimal both in terms of spectral density and stability around the centre of a $\sim 50\mu$m-wide region—approximately the FWHM spot size of the transverse laser. Either side of this region however, the quality and stability of the injection diminished as the overlap between the wake and the downramp generated by the transverse laser decreased [197, 216]. Typical operation involved performing such a scan and identifying the centre of this region that was characterised by large amounts of charge injected with high stability. Placing the transverse laser here meant that the injection was improved at that time but was also more resistant to inevitable drifts in the transverse laser and drive beam alignment; a drift of $\sim 25\mu$m in either direction was necessary to dramatically effect the injection. Nevertheless, such drifts could occur and so this procedure was often revisited when an otherwise unexplained drop in injected charge was detected.

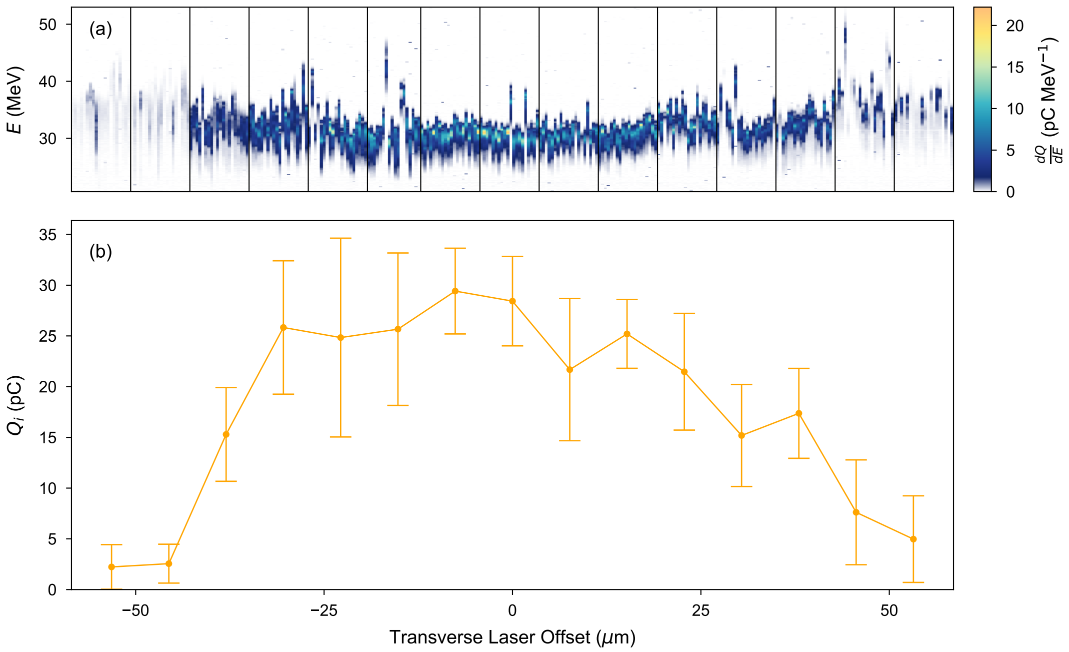

Figure 5.21: Injection dependence on transverse laser offset. (a) Waterfall plot of the injected witness beam spectra, measured on the ESPEC grouped by the transverse laser offset position. (b) The mean injected charge detected on the ESPEC for each transverse laser offset position, with errorbars indicating the rms variation of the charge at a given offset position. Here, the transverse laser position equal to zero was set as the point in the middle of the region with high injection stability.

Longitudinal laser alignment

In a similar way to the transverse laser, the pointing of the longitudinal laser could also be further optimised for injected charge despite already being aligned to the drive beam during setup. By tilting the last spherical mirror in the longitudinal laser beamline, the overlap between the longitudinal plasma and the driver could be varied and optimised. In the figures that follow, the spherical mirror was tilted in the horizontal direction; similar scans along the other axis also exist and show similar trends. As this movement was performed with a picomotor, the absolute positions that correspond to the given steps are not completely reproducible. Nevertheless, the scans still illustrate how changes in the alignment of the drive beam and longitudinal laser impacted the interaction.

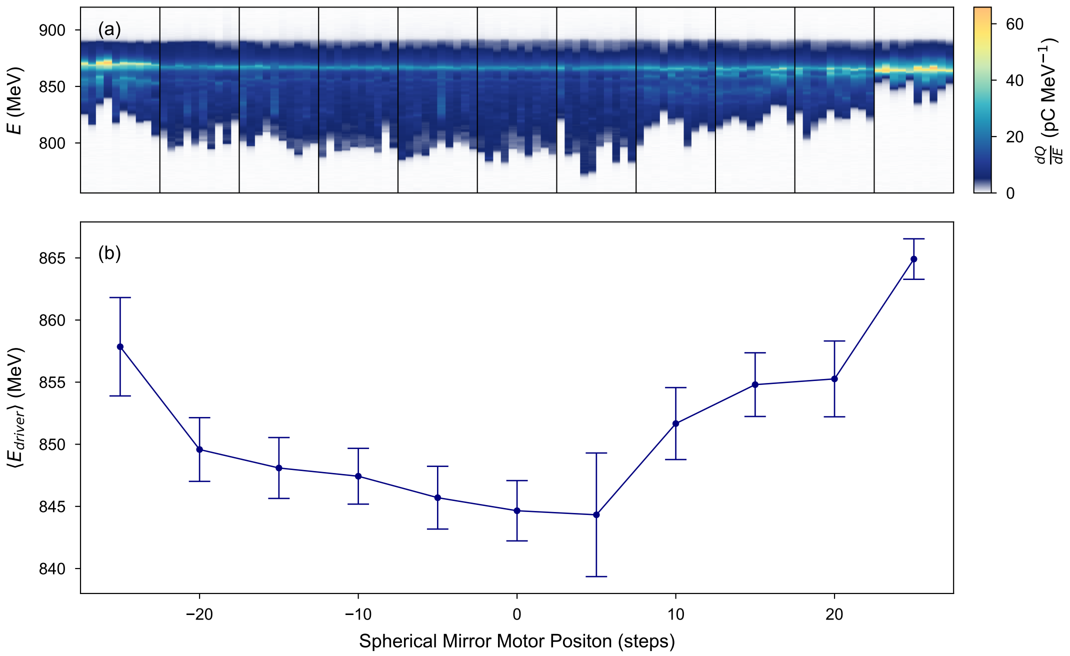

Figure 5.22: Driver energy loss variation with longitudinal laser pointing. (a) Waterfall plot of the drive beam energy spectra, measured on the ESPEC and grouped by the spherical mirror horizontal motor position. (b) Mean energy of the drive beam electrons, averaged for each spherical mirror motor position, with errorbars indicating the rms variation of the mean driver electron energy each position. The spherical mirror motor position corresponding to zero was set as the centre of the scan with the most strong and stable driver interaction. Data presented in this figure was acquired during the October 2021 experimental run.

Misalignment of the longitudinal laser with respect to the drive beam resulted in a variation in the wake that is being driven. Figure 5.22 illustrates this with a systematic study of the effect of the longitudinal laser pointing on the driver energy spectrum, carried out in the October 2021 experiment run.

The increased energy loss at the centre of this scan range indicates that the beam was sampling a higher plasma density, averaged along the length of the longitudinal plasma. The approximate symmetry about this point suggests that this was the position where the drive beam was most central with respect to the plasma channel. Any interaction of the wake with the edges of the plasma would likely effect the injection and transport of the witness beams, and so driving a wake as close to the middle of the channel as possible was crucial.

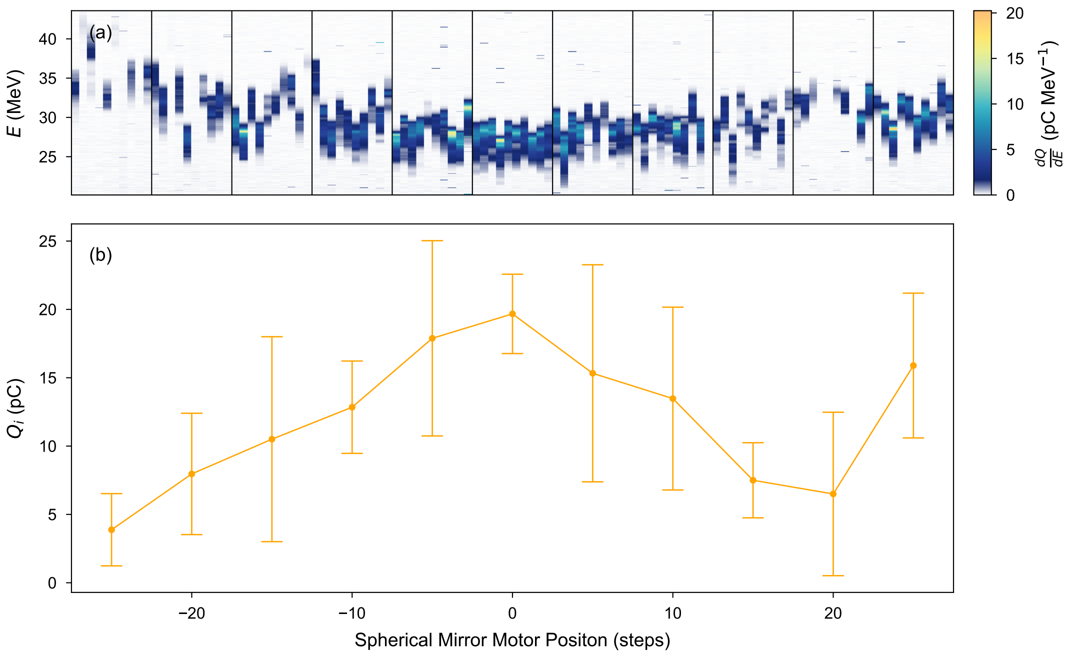

Figure 5.23: Injection dependence on longitudinal laser pointing. (a) Waterfall plot of the witness beam energy spectra, measured on the ESPEC and grouped by the spherical mirror horizontal motor position. (b) Mean injected charge detected on the ESPEC for each longitudinal laser offset position, with errorbars indicating the rms variation of the charge at a given offset position. The spherical mirror motor position corresponding to zero was again the optimal for injection, and should roughly correspond to the same zero position in Figure 5.22.

This is supported by an analogous scan presented in Figure 5.23, this time studying the witness beam injection as a function of longitudinal laser pointing. Both in terms of the amount of charge and the stability with which it was injected, it is clear that maintaining alignment between the driver and longitudinal laser was crucial. Again, an optimal position existed at the centre of the scan range, where deviations away from this resulted in poorer injection. An interesting observation here is that as the alignment became sub-optimal, the average energy of the injected beams also increased. An obvious explanation for this is beam loading; as the charge decreased, beam loading was reduced and the charge that was injected saw a higher accelerating field. However, both the acceleration of the witness beams and the energy loss of the drive beam will be dependent on the details of the laser-generated longitudinal plasma—something explored in more detail in simulations in Section 5.1.3. Figure 5.13 shows that the longitudinal plasma channel initially has a maximum density on-axis; transverse misalignment of this plasma would therefore translate to lower deceleration of the drive beam. However, due to ionisation defocussing, the radial profile that the injected bunches experience is in fact *lower* at the centre. Transverse misalignment of the accelerating plasma channel with respect to the injector laser would therefore result in injected bunches experiencing a higher plasma density during acceleration, offering another potential explanation as to why the accelerated bunches gain more energy when the longitudinal laser has a radial offset.

From the previous results, it is clear that the measured witness beams were sensitive to any misalignment between the driver and the longitudinal laser. As this laser propagated through a relatively long beamline, it was also particularly susceptible to drifts and jitters. However, the degree of overlap could also be monitored via another observation—the amount of plasma light measured after interaction (see Figure 5.24) [201, 206].

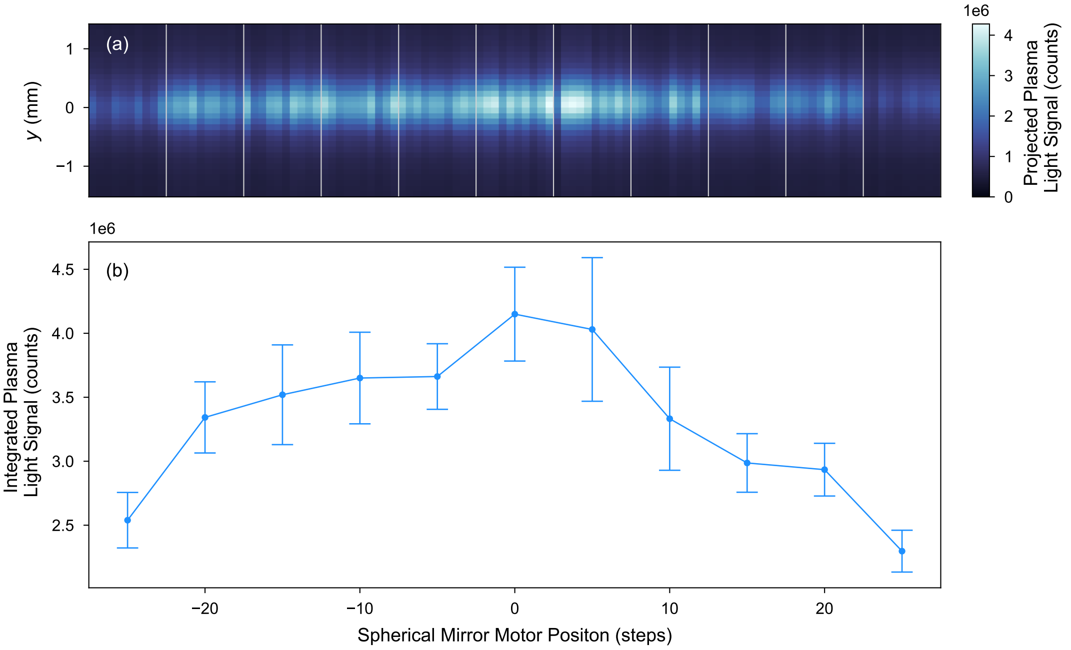

Figure 5.24: Integrated plasma light yield variation with longitudinal laser pointing. (a) Waterfall plot of the plasma light signal measured with the side view camera, projected along the capillary and grouped by the spherical mirror horizontal motor position. (b) Mean integrated plasma light yield for each longitudinal laser offset position, with errorbars indicating the rms variation of the signal at a given offset position. The data displayed in this figure was acquired simultaneously to that in Figure 5.22.

As the driver energy loss increased with improved overlap, so did the excess plasma light detected through the side of the capillary—as in Chapter 4. The plasma afterglow light therefore again served as a useful, non-intrusive diagnostic—this time as a proxy for the overlap between driver and longitudinal plasma. Typical operation often involved frequently monitoring the plasma light and, after a drift occured, re-aligning the longitudinal laser until this signal was maximised. In most cases, such a correction substantially improved injection. Combatting shot-to-shot jitters during online operation was however more difficult; although a well aligned laser and drive beam was less susceptible to random jitter, it could not be mitigated fully. Nevertheless, the excess plasma light could then be effective in identifying shots with poor overlap in subsequent analysis.

Driver compression

A key factor affecting the amount of injected charge is the driver current-profile; higher currents generally mean a stronger wake with a deeper trapping potential. In this experiment, the driver current-profile was modified by adjusting the compression settings in the FLASH linac, as described in Section 5.1.1.

Figure 5.25 shows the dependence of the injected charge on the driver compression in the June 2021 experimental run, where the driver bunch length was obtained from a BCM reading, multiplied by the factor $R_{BL,0}$ that was determined from the data shown in Figure 5.3 (c). For bunch lengths $\gtrsim 160\,\mathrm{fs}$, only very small amounts of charge were injected. Compressing the bunch length down to $\sim 100\,\mathrm{fs}$, the amount of injected charge increased approximately linearly to a value of $\sim 35\,\mathrm{pC}$. For bunch lengths shorter than this, the injected charge began to decrease again before completely ceasing at a value of $\sim 70\,\mathrm{fs}$, approximately coinciding with the corrected BCM bunch length measured at the point where overcompression begins in Figure 5.3 (b) (black dashed line at 73 fs). This was likely a manifestation of the negative side effects of strong compression, discussed previously in Section 5.1.1. Ultimately, this data illustrates the importance of driver compression as a factor affecting injected charge, as well as the limit beyond which stronger compression proved counter-productive to the injection process.

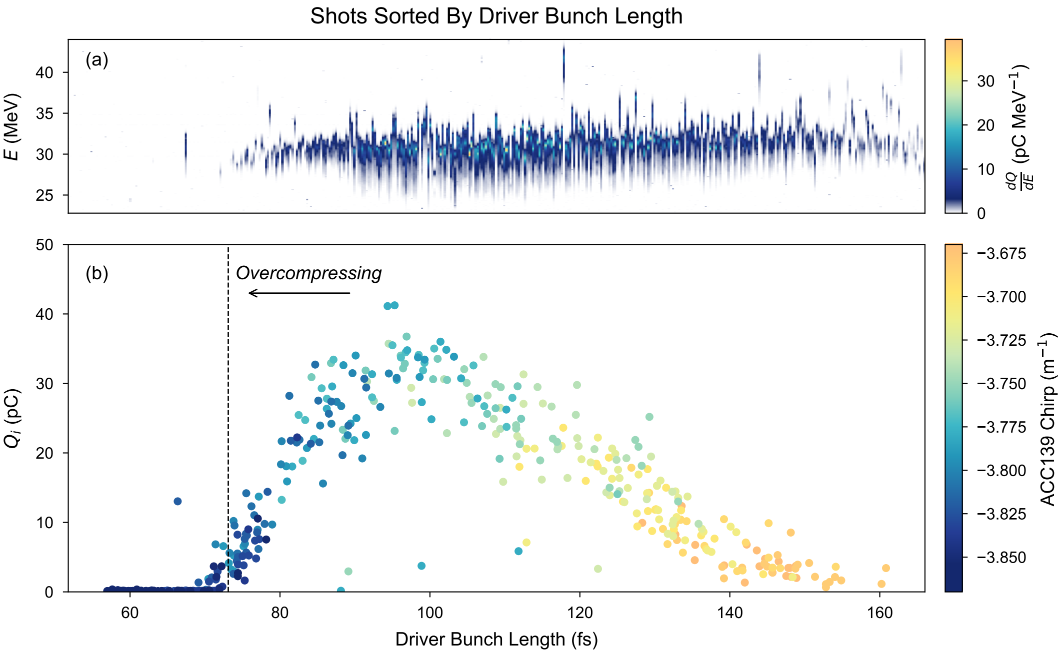

Figure 5.25: Injection dependence on the drive beam compression. (a) Waterfall plot of the witness beam energy spectra, measured on the ESPEC and sorted by driver bunch length. (b) Injected charge detected on the ESPEC versus the driver bunch length measured by the BCM with the correction factor calculated from Figure 5.3. A clear optimum in the injected charge was observed for measured bunch lengths $\sim$ 100 fs; approaching the overcompression regime (black dashed line), the injected charge decreased again. Note that drifts in the machine meant that the chirp settings required to compress the driver were different to the data in Figure 5.3, which was acquired some time earlier in the experimental run.

Driver focusing

Although the current-profile of the driver plays a vital part in how the wake is generated, it more generally depends on how much of this current is located transversely within the wake ($\sigma_{r,d} \lesssim k_p^{-1}$)—meaning that the transverse phase space of the beam is also equally important. In experiment, this translated to proper focusing of the driver into the plasma, in order to produce a strong enough wake.

Figure 5.26 shows a 2D scan of the two final focusing quadrupoles, performed in the June 2021 experimental run. As explained in Section 5.1.1, the optical configuration was such that these magnets were able to change the longitudinal position of the driver waist in the x and y planes approximately independently, with little effect on the driver waist size. Accordingly, the normalised strengths of these quadrupoles are denoted as Δk_x and Δk_y, controlling the position of the waist in the horizontal and vertical planes respectively. Whilst calculations of the absolute waist position from the quadrupole currents would not be accurate enough at this scale, an important observation can be made about the sensitivity of the injection process to the driver focusing. As shown in Figure 5.26, changes to the quadrupole strengths on the order of 1% were enough to move the system from an optimised working point to one where injection had completely ceased. This therefore highlights the need for precise control over the focusing of the driver—something that is currently only achievable with the low energy spread, highly stable beams produced by conventional machines like FLASH.

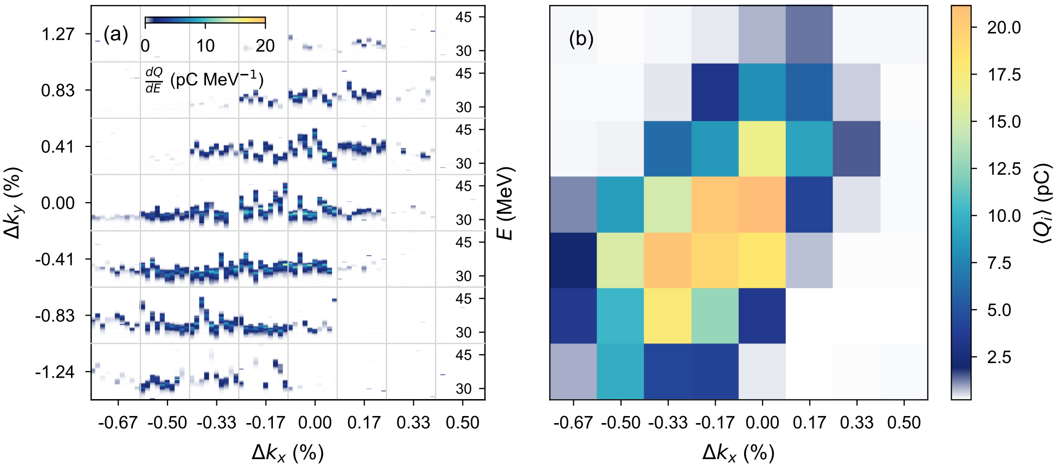

Figure 5.26: Injection dependence on the drive beam focusing. (a) 2D grid of waterfall plots of the witness beam energy spectra, measured on the ESPEC and grouped by the relative focusing strengths of the final focusing quadrupoles Δk_x and Δk_y that control the position of the driver waist in the x and y planes respectively. (b) 2D histogram of the average injected charge detected on the ESPEC for each combination of Δk_x and Δk_y.

Laser—electron beam timing

During experimentation, a further powerful method of increasing the injected charge was discovered. In previous iterations of this experiment, both lasers arrived promptly before the drive beam ($\sim$ 10 ps earlier), interacting with the plasma just after its creation. However, it was found that by letting the plasma evolve over timescales on the order of tens of nanoseconds, the quality of injected bunches was dramatically improved.

Figure 5.27 shows a demonstration of this from the June 2021 experimental run, where the lasers were made to arrive earlier than the driver by multiple nanoseconds. This was easily achieved by selecting earlier laser pulses from the oscillator to be amplified, but meant that the delay could only be varied in steps of 18.46 ns (see Figure 3.9 on page 58 for the different options in adjusting the laser—electron beam timing). Despite this, a significant improvement in the injected charge was observed as the lasers arrived earlier in time to the driver, with an optimum found at a setting of 18.46 ns. For arrival times earlier than this, the amount of injected charge decreased again.

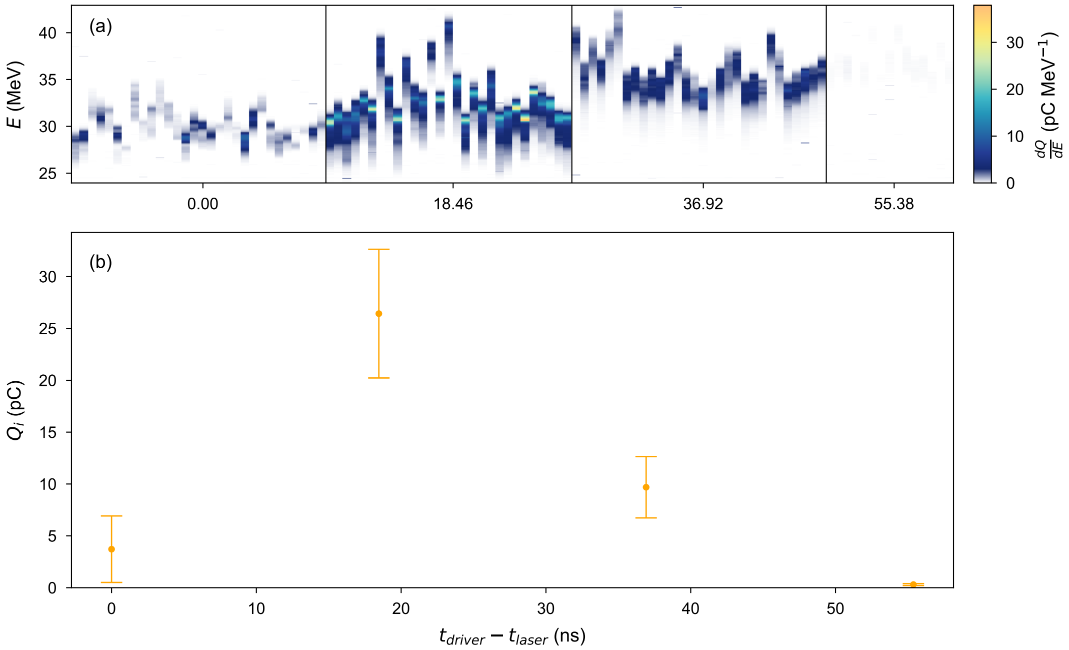

Figure 5.27: Injection dependence on driver-laser arrival time difference, varied by picking earlier laser pulses. (a) Waterfall plot of the witness beam energy spectra, measured on the ESPEC and grouped by the difference in laser and driver arrival time. (b) Mean injected charge detected on the ESPEC for each arrival time difference, with errorbars indicating the rms variation of the charge at a given timing.

The reason for the dramatic change in injected charge is not immediately obvious; over these timescales, both the longitudinal and transverse plasma columns can be expected to evolve in a way that is dependent on their initial density and temperature distributions after generation—parameters that themselves depend on the neutral gas mixture, the laser intensities and the laser polarisation [54]. Recall that the lasers are both linearly polarised in the vertical directions in this case; as investigated in Ref. [54], any ellipticity of the polarisation would indeed be expected to increase the temperature of the resulting plasma. Note that internal injection of witness bunches using laser-seeded, hydrodynamically generated downramps has been experimentally demonstrated before in both LWFA [110, 220, 221] and PWFA [118]. In this case however, the injection mechanism does not rely purely on hydrodynamic expansion; the transverse plasma column is still able to inject $\sim$ ps after it is generated by ionisation, before any evolution has taken place. Rather, hydrodynamic expansion modifies the existing 'plasma torch' structure, leading to in this case an increase in the injected charge. Section 5.3.2 specifically explores the effect of a changing downramp on the injection process in this context and attempts to explain how the evolution of these downramps led to a change in the injected charge.

Another option to vary the arrival time between driver and lasers on the tens of nanoseconds timescale was to make the drive beam arrive later. Adopted in later experimental runs, this method involved selecting different RF buckets to accelerate the drive beam in, with the advantage that this could be used to delay the drive beam in steps of 0.77 ns—a much higher resolution than could be achieved with the laser pulse selection technique (see again Figure 3.9 on page 58).

Figure 5.28 displays data from the experimental run in November 2021, where this method was deployed. Again, by allowing the driver to arrive later than the lasers by multiple nanoseconds, the injected charge increased, reaching an optimal value before declining again. The higher resolution of this scan provides more confidence that the timing with the highest amount of injected injected charge was a true optimum in the parameter space. However, one immediately notices that the timescale of this behaviour was shorter than in the previous case; the injection was optimal with 5.39 ns drive beam delay, and almost completely ceased after $\sim$ 18.5 ns. The causes for these disparate timescales are not clear, but it should be stressed that these measurements were taken from separate experimental runs, with slightly different experimental conditions. For instance, the drive beam peak current in November 2021 was measured on the TDS to be $\sim$ 1.25 kA, as opposed to the 1.9 kA measured in June 2021 (see Figure 5.2). Moreover, changes in the focal spots of the lasers between shifts—specifically deterioration of the transverse laser focus quality mentioned in Section 5.1.2—would also be imprinted on the resulting plasma channels, inevitably impacting their subsequent evolution.

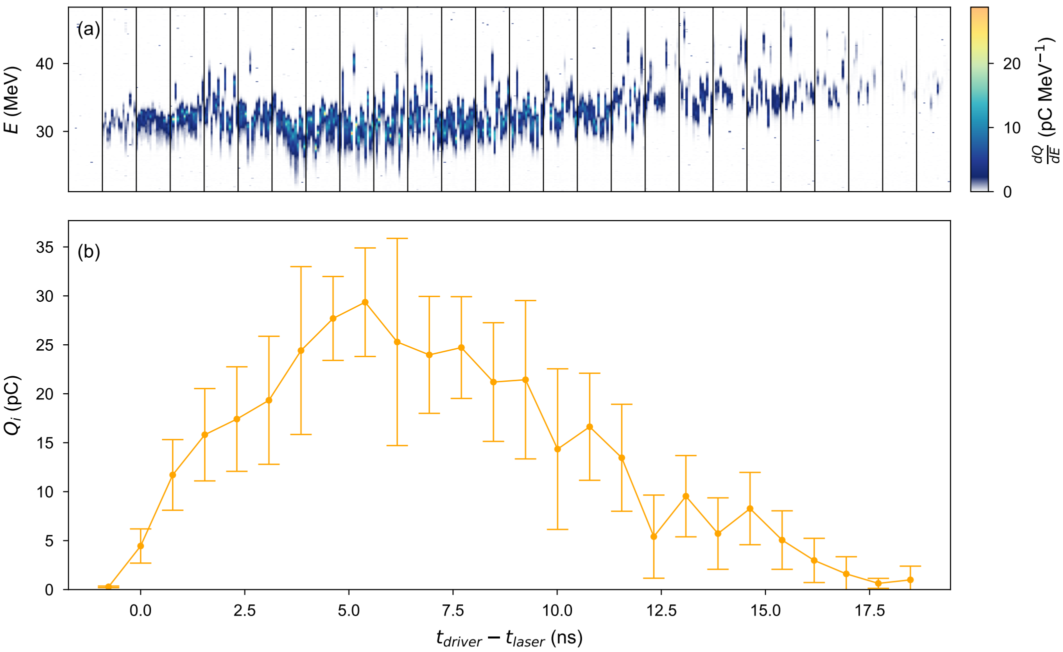

Figure 5.28: Injection dependence on driver-laser arrival time difference, varied by accelerating the driver beam in later RF buckets. (a) Waterfall plot of the witness beam energy spectra, measured on the ESPEC and grouped by the difference in laser and driver arrival time. (b) Mean injected charge detected on the ESPEC for each arrival time difference, with errorbars indicating the rms variation of the charge at a given timing.

5.2.4 Optimised working point

Using the methods detailed in the previous section, the injection process was optimised for high charge and consequently high spectral density. Such an optimised working point with 1000 consecutive shots from the June experimental run is displayed in Figure 5.29. Injection here was highly reproducible, generating witness beams with low mean energy spreads of 1.3% FWHM and peak beam energies narrowly distributed around a mean of $(29.4 \pm 1.5)\,\mathrm{MeV}$. A high average charge of $(19 \pm 7)\,\mathrm{pC}$ translated also to a high average peak spectral density $\langle (dQ_i/dE)_{peak} \rangle = (14 \pm 6)\,\mathrm{pC\,MeV^{-1}}$, with a maximum measured value of $36\,\mathrm{pC\,MeV^{-1}}$ (see Figure 5.29 (d)).

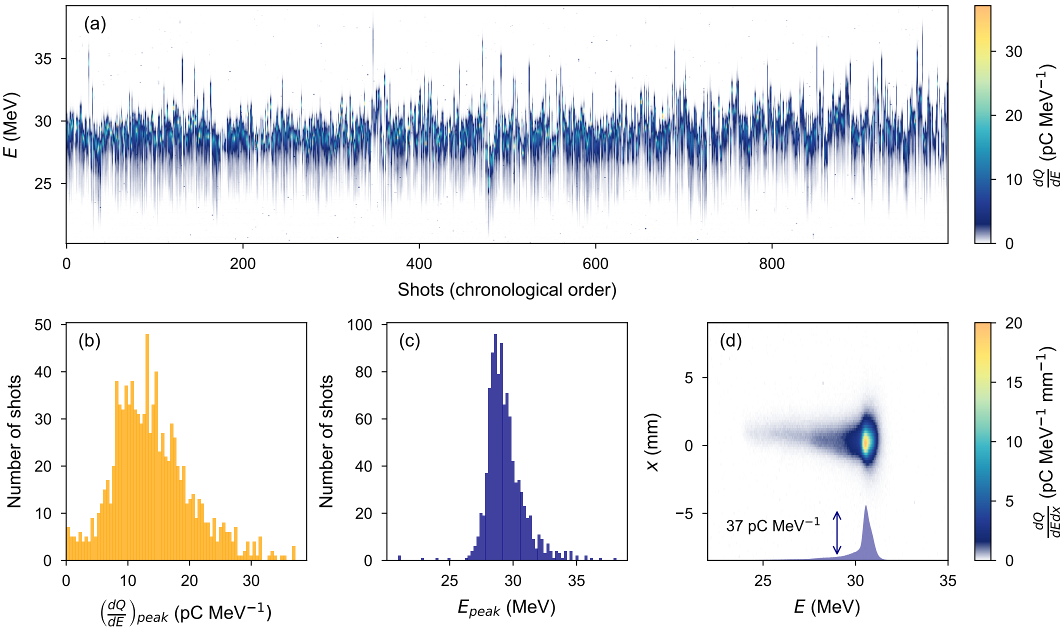

Figure 5.29: Optimised injection working point. (a) Waterfall plot of the witness beam energy spectra over 1000 consecutive shots, measured on the ESPEC and displayed in chronological order. Histograms of the (b) peak spectral densities and (c) peak energies are shown, each using 75 bins. (d) Example ESPEC image of a beam and its spectral density profile (navy shaded curve) that reaches a peak value of 37 pC MeV^{-1}—the highest of all the shots in this dataset.

As alluded to in Section 5.2.3, the influences of the longitudinal laser pointing could also be observed in the integrated plasma light yield. The plasma light signal could therefore be used to identify shots where the alignment of the drive beam and the main accelerating channel were non-optimal for injection. Figure 5.30 demonstrates this by sorting the shots by increasing integrated plasma light yield. For the shots with the lowest integrated plasma light yields, little to no charge was injected. Beyond this, one can see from the waterfall plot in Figure 5.30 that the injection became visibly more stable for higher plasma light yields; for example, considering only the 20 % of shots with the highest plasma light yields, the mean peak energy is (28.6 ± 0.8) MeV—a lower energy than overall but with a relative rms jitter of 2.8 % rather than 5 %. This then suggests that the longitudinal laser—driver alignment was at least partly responsible for the overall jitter of the injection process and so decreasing the jitter of this parameter should increase the stability of the injected witness beams.

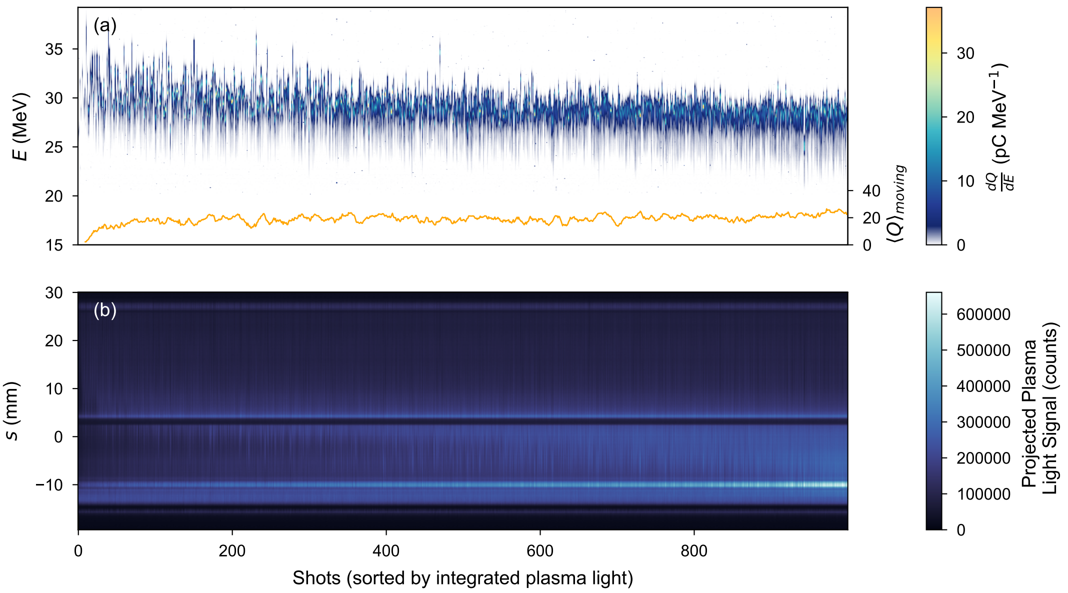

Figure 5.30: Optimised injection working point, with shots sorted by increasing plasma light yield. (a) Waterfall plot of the witness beam energy spectra over 1000 consecutive shots, measured on the ESPEC and sorted in order of increasing plasma light yield. Overlayed is a moving average of the injected charge (orange line), with a window size of 10 shots. (b) Waterfall plot of the plasma light signal, projected over the transverse axis and so illustrating the longitudinal variation of the plasma light.

5.2.5 Emittance measurements

A key part of characterising high-brightness beams is a measurement of their emittance, typically involving the retrieval of the divergence and waist size of the beams as they expand in free space. In this section, the previously optimised working point is used as a basis for multi-shot object plane scan measurements (see Section 2.2.2 on page 18) of the witness beam divergence on the ESPEC. The dependence of this divergence on a key experimental parameter—the laser-to-driver timing—is also investigated. Ultimately, the emittance measurements are limited on the ESPEC due to an inability to resolve the waist size of the witness beams at the plasma exit. Measurements on the LEMS for a particular high-brightness candidate working point are therefore used, where a higher resolution makes true estimates of the witness waist size possible. For the same data, the single-shot butterfly emittance measurement is also attempted [183]—the results of which are then compared with its multi-shot counterpart.

ESPEC object plane scans

In order to ensure the highest resolution possible when characterising the transverse properties of the injected witness beams, measurements at the ESPEC are performed on the high energy screens (see Section 3.2.5 on page 53)

For the low-energy electron beams produced in this experiment, this involved operating the ESPEC dipole with only small currents in order to deflect them to the higher resolution screen. A drawback of this method was that hysteresis effects therefore became non-negligible, meaning previous calibrations obtained e.g. with a dipole current scan of the high-energy drive beam would develop a systematic error. Although the effect of hysteresis could in principle be mitigated by 'cycling' the magnet—ramping the field up and down to maximum and minimum values several times before returning to the desired setting—this would need to be performed many times before the field could be reliably known for such small values. Additionally, the field of this magnet often needed to be changed when switching between different screens, requiring a cycling procedure each time which would ultimately become unfeasible. However, this error could instead be accounted for using knowledge of the witness beam mean energy that was already reliably measured on the low-energy screen. For datasets taken on the high-energy screen, the mean energy was measured and compared with the mean energy of witness beams obtained on the low energy screens, with experimental conditions that were as close as possible to the working point in question. The ratio of these two mean energies then served as a correction factor to be multiplied with the energy axis of beams measured on the high energy screen, essentially enforcing that the same mean energy was measured as that on the low-energy screen.

Figure 5.31 shows an object plane scan measurement of the injected beams on the ESPEC, performed in the June 2021 experimental run. The working point is that displayed in Figure 5.29. Figure 5.31 (a) displays the injected beams on the ESPEC screen, projected onto the x-axis. In order to quantify the sizes of the beams on the screen $\sigma_{x,i,screen}$, a gaussian is fitted to each projection whose width is taken to represent that of the beam in the non-dispersive direction. The real size of each beam $\sigma_{x,i} = \frac{\sigma_{x,i,screen}}{|M_x|}$ is plotted as a function of $s_{obj,x}$ in Figure 5.31 (b), where the magnification M_x and object plane position $s_{obj,x}$ were calculated using the mean witness electron energy and quadrupole strengths for each shot, following the procedure in Section 3.2.5.

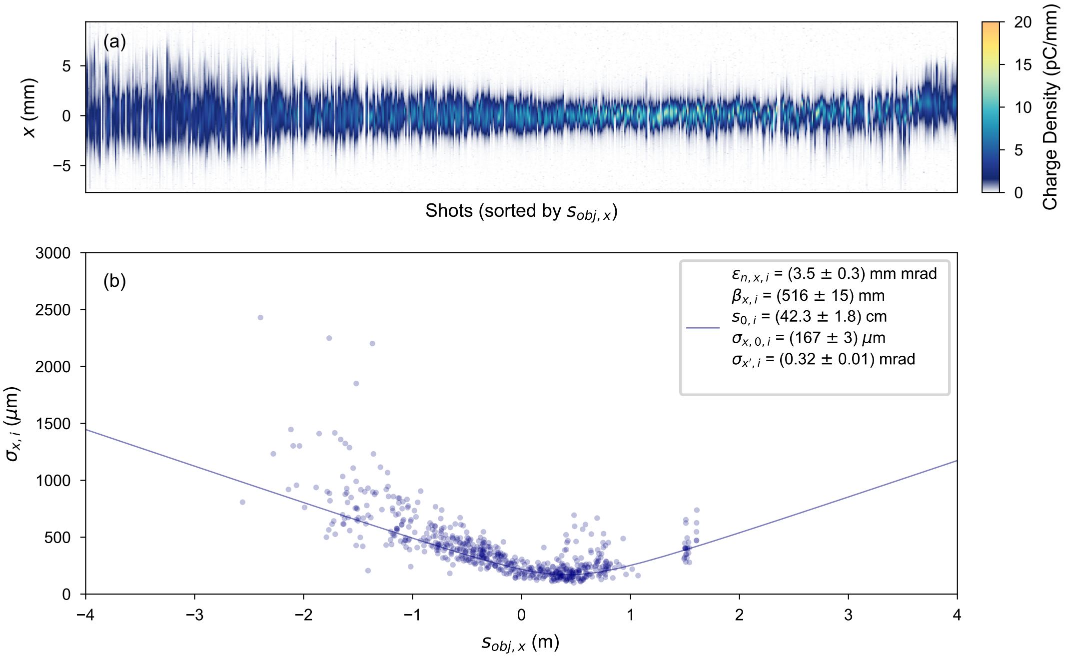

Figure 5.31: Example object plane scan measurement of internally injected beams on the ESPEC. The working point is the same as that displayed in Figure 5.29. (a) Waterfall plot of the screen x-projections, sorted in order of ascending object plane position $s_{obj,x}$. (b) Beam widths in the x-direction (navy dots), found from gaussian fits to the projections in (a) and plotted versus $s_{obj,x}$, with the magnification of each beam M_x accounted for. Equation 2.38 is fitted to the data (navy curve), with the errors on each width (from the gaussian fits) used as variance-defined weights. From this fit, values relating to the waist and divergence of the beam are extracted and displayed in the legend, along with their fit errors.

By fitting Equation 2.38 to the points in Figure 5.31 (b), the expansion of the beam into free-space was characterised, yielding a divergence $\sigma_{x'} = (0.32 \pm 0.02)$ mrad and waist size $\sigma_{x,0} = (167 \pm 3)\,\mu\text{m}$. Due primarily to the large waist size, the emittance estimated here is also relatively large—$\epsilon_{x,n,i} = (3.5 \pm 0.3)\,\text{mm} \cdot \text{mrad}$. Note that, the waist position returned by this method also appears downstream of the end of the plasma cell—by approximately (370 ± 20) mm in this case. This effect was reproduced in many of the injected beam object plane scans performed in this experiment, including in Ref. [197], where the presence of long exit ramps formed by the laser-ionisation of residual gas beyond the capillary extent was suggested as a mechanism with which the beam could be refocused downstream of the plasma cell. The feasibility of this explanation is further discussed with the aid of simulations in Section 5.3.1.

In addition to the optimised working point, object plane scans were also performed for different values of the laser—driver arrival time delay in order to investigate whether further improvements to the injected beam's divergence could be made. Figure 5.32 shows the divergences and waist sizes from object plane scans performed at 3 different laser delay values—including the working point measured at the delay $t_{driver} - t_{laser} = 18.46\,\text{ns}$ in Figure 5.31. As well as having a significant effect on the injected charge (see Figures 5.27 and 5.28), allowing the driver to arrive on the order of tens of nanoseconds after plasma generation also influenced the transverse properties of the internally injected beams. Figure 5.32 (a) shows that the divergence of the injected beams decreased from a value of $(0.71 \pm 0.01)\,\text{mrad}$, where the lasers arrived promptly before the driver, down to a value of $(0.30 \pm 0.01)\,\text{mrad}$, where the driver arrived 36.92 ns after the lasers. Just as with the variation in injected charge, the reason for the change in the divergence of the injected beams is not immediately clear. Over these timescales, both the longitudinal and transverse plasma channels are expected to evolve, in principle affecting both the initial injection event but also how the beams are transported in and released from the plasma. In any case, the decreased divergence—in combination with the increase in charge—suggests that allowing the driver to arrive after some evolution of the plasma resulted in an enhanced brightness of the injected beams. As a compromise between low divergence and high charge, the working point at a delay of $t_{driver} - t_{laser} = 18.46\,\text{ns}$ still proved optimal. The real brightness however could only be determined with knowledge of the true emittance, necessarily requiring a reliable measure of the witness beam waist size. Figure 5.32 (b) shows that although also decreasing with laser delay, the injected beam waist size measured on the ESPEC during object plane scans remained relatively large ($> 150\,\mu\text{m}$).

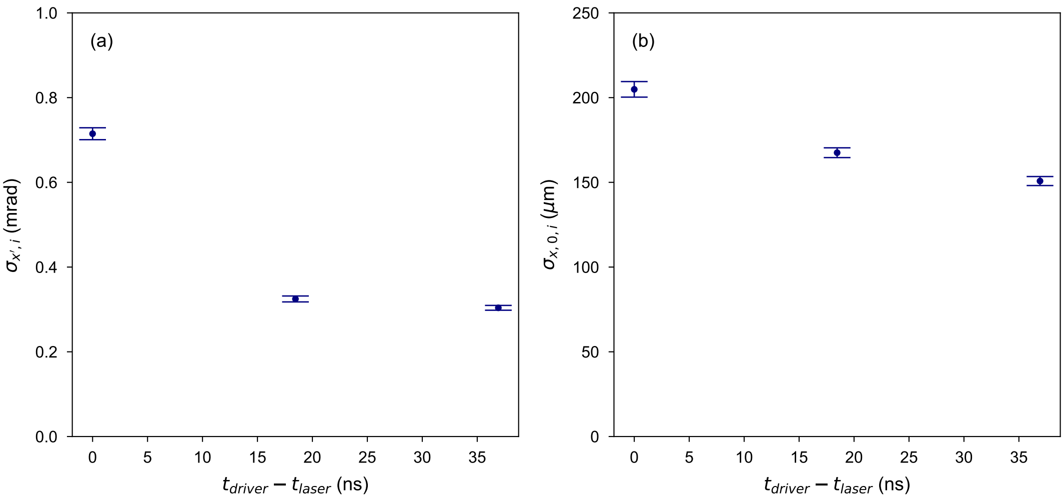

Figure 5.32: Variation of horizontal divergence and waist size with driver—laser arrival delay, measured on the ESPEC. As the laser arrives progressively earlier than the driver, both the (a) divergence and (b) waist size of the injected witness beams decrease, as estimated with the same method demonstrated in Figure 5.31. In the case of both the divergence and waist size, the values (navy dots) are those extracted from the fits of Equation 2.38, with the error bars representing the fit errors on these parameters.

Owing to the expansion of the beam at the plasma—vacuum transition, the waist size would indeed be expected to be larger than the transverse beam size inside the blowout. Nevertheless, the values measured here are likely overestimated due to limits in the measurement technique. One such limit was the spatial resolution of the diagnostic, which was further degraded due to scattering of the relatively low-energy electron beams in the chamber wall (see Section 3.2.5). Another limit however came from the natural chromaticity of the imaging system, coupled with the relatively high energy spread of the injected bunches. Due to this chromaticity, only a small bandwidth of energies within the electron beam could be properely imaged at a given time. Any beam slices at energies away from the imaging energy would therefore appear larger, resulting in a projected beam size that was an overestimate of the real beam size. This is especially true around the waist, where the signal would form a 'butterfly' shape (see imaging beamline tracking simulations in Section 5.3.3). Ultimately, whilst object plane scans on the ESPEC could provide a robust measurement of the beam divergence, the combination of factors resulting in low spectral and spatial resolution made this diagnostic unsuitable for diagnosing the waist size of the injected beams produced in this experiment.

LEMS measurements

In contrast to the ESPEC, the LEMS is an in-vacuum diagnostic with a large amount of pixels in the field-of-view, removing the measurement limits due to scattering and poor screen resolution (see Section 3.2.5 on page 53). This results in both far higher spatial and also spectral resolution, meaning that the widths of individual beam energy slices could be measured accurately. The beam waist size could therefore be better characterised with this diagnostic, even in the presence of the natural chromaticity of the imaging system. As is often the case however, this increased resolution comes at a price; the bandwidth of the LEMS is much narrower than that of the ESPEC, meaning that beams with an energy jitter often partially or even fully miss the screen. Furthermore, the LEMS is positioned further downstream of the ESPEC in a beamline designed for high-energy, 100 pC-charge bunches, making the stable transport of these injected beams even more challenging. Such issues essentially precluded the application of the multi-shot object plane scan method on the LEMS to the beams generated in this experiment. Nevertheless, by collecting as many shots as possible whilst attempting to image the waist of the injected beams, a more accurate value of the beam waist size was estimated. Using Equation 2.29 in the absence of the correlation term,

$$\epsilon_{x,n} \approx \gamma \cdot \sigma_{x,0} \cdot \sigma_{x'}, \tag{5.1}$$

this waist size $\sigma_{x,0}$ could be combined with the mean electron relativistic factor γ and the divergence $\sigma_{x'}$ measured in Figure 5.31 to obtain upper limits on the average emittance of the injected beams at this optimised working point.

Due to low signal-to-noise ratio, identification of shots that contain a reasonable amount of charge was difficult. Even after subtraction of a background image acquired with all beams blocked, the ambient noise around the signal when an injected beam was present could fluctuate significantly. Consequently, a simple filter based on the pixel count of the images does not provide a clear enough threshold. A technique that instead worked well was to simply fit a gaussian to the projection each of the beams in the x-direction. Shots were then discarded by filtering based on the fit peak value and fit width. For example, if the fit has a high peak value but very narrow width, the likelihood is that no beam signal is present and so the fit instead converges on hot pixels. Conversely, if the fit peak is a similar level to that of the background and the width of the fit very large, there is again likely no signal on the screen and the fit is instead dominated by the background noise. Ultimately, this procedure identifies 154 images out of a set of 620 that contain a reasonable amount of injected beam signal;

a manual inspection also confirms that little to no viable shots are rejected by this method.

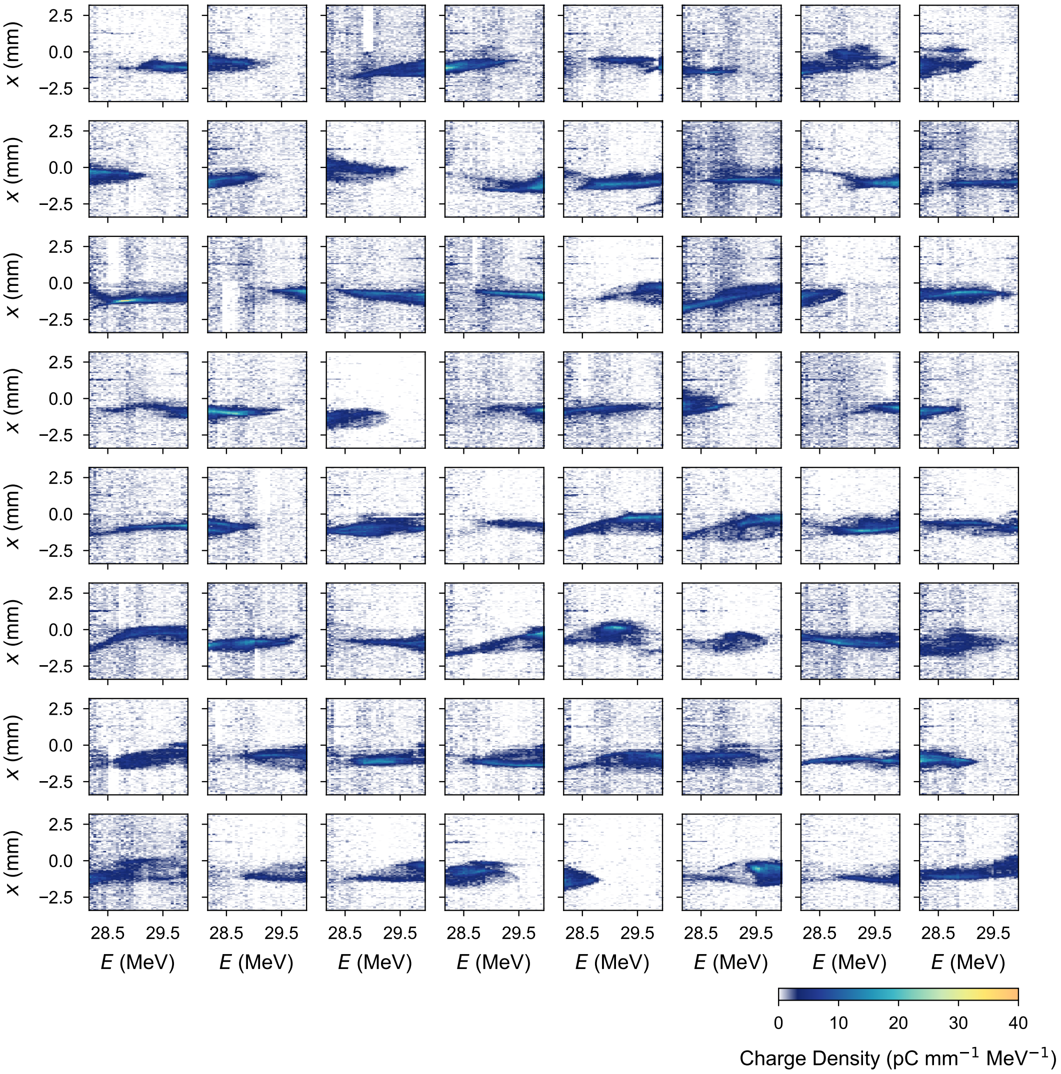

Figure 5.33: Example injected beams measured on the LEMS diagnostic. A subset of 64 images are shown, with the median filter and binning procedure described in the text applied.

Figure 5.33 shows a subset of the images that were identified with the aforementioned routine. Just as in the case of the ESPEC, the horizontal beam size for each shot can be estimated from the widths of the gaussian fits to the projections in the previous step. However, the width of the beam in a given shot is not constant across energy slices, primiarly as a consequence of the chromaticity of the imaging system. Furthermore, not all beams are perfectly straight, often possessing a slight tilt. Both of these effects will result in an inflated projected width that contributes to an overestimate of the actual beam waist. A more accurate measurement of the waist size would instead be provided by the width of the beam slice that is properly imaged on the screen, minimising the effects of chromaticity and tilt.

Identification of energy slices that are well-imaged is however non-trivial. Ideally, each beam would make a clear 'butterfly' shape, from which a waist could be identified and measured. Although some shots show hints of this behaviour, a large amount do not. This is likely because the divergence of the beam has to be large enough such that the 'butterfly' structure appears over a relatively narrow bandwidth of energies, which may only occur for a small number of shots. Furthermore, precisely imaging a particular slice was challenging; the transport of the beams was very sensitive to changes to the quadrupole settings, and optimisation of the imaging was made harder still by the relatively unstable nature of the injected beams.

To obtain a fair measurement of the waist size from the images, a statistical approach was therefore adopted. In order to combat the relatively low signal-to-noise ratio of this measurement, an image processing routine was first executed for each shot. This involved the application of a 3×3 median filter, followed by the binning of the image along the energy axis into slices 11 pixels wide. The background noise level C_n of this image was then estimated by taking the maximum pixel value of a region far from the signal. Only slices that have a maximum value that is greater than $N_{min}C_n$ are considered to contain enough charge to be viable, where N_{min} is a free parameter. Moreover, any images that contain no viable slices are completely rejected. With this in mind, a value of $N_{min} = 2.5$ was found to minimise the amount of low charge slices that were selected, whilst also resulting in relatively few images being completely discarded. Gaussians were then fit to the remaining energy slices and the width of the narrowest slice in the x-direction was taken as the estimate of the true waist size.

The resulting widths found using this method are displayed in Figure 5.34 (a), along with the projected beam sizes for comparison. Figure 5.34 (b) shows these widths binned into histograms. Extracting the mean and standard deviation from these distributions yields a shot-averaged projected beam width $\langle \sigma_{x,proj} \rangle = (143 \pm 8)$ µm and minimum energy slice width $\langle \sigma_{x,slice} \rangle = (69 \pm 4)$ µm, indicating that the projected beam width does indeed result in an overestimate of the beam waist size. As a test of the robustness of the minimum slice width method, the free parameters of the algorithm—namely the median filter kernel size, bin width and the factor N_{min}—were varied. It was found that by independently adjusting these parameters by a factor of 2 resulted in a change of around only 10% in the value of $\langle \sigma_{x,slice} \rangle$, suggesting that the procedure is reliable. Whilst the performance of this algorithm will have some impact on the widths that are obtained, the range of values measured will also be due to the fact that each beam will differ in quality, and also that the imaging is quite often imperfect—the latter of which is likely to increase the estimated beam size.

Using $\langle \sigma_{x,i,slice} \rangle$ therefore as an upper limit on the waist size $\sigma_{x,0,i}$ of the beams, this can be combined in Equation 5.1 with the divergence measured for this working point on the ESPEC to yield an upper bound on the emittance of $\epsilon_{x,n,i} = (1.2 \pm 0.1)\,\text{mm} \cdot \text{mrad}$. As a more optimistic estimate, taking the modal value $(\sigma_{x,slice,i})_{mode} = 43\,\mu\text{m}$ from the histogram in Figure 5.34 as the value of $\sigma_{x,0,i}$ yields $\epsilon_{x,n,i} = 0.8\,\text{mm} \cdot \text{mrad}$.

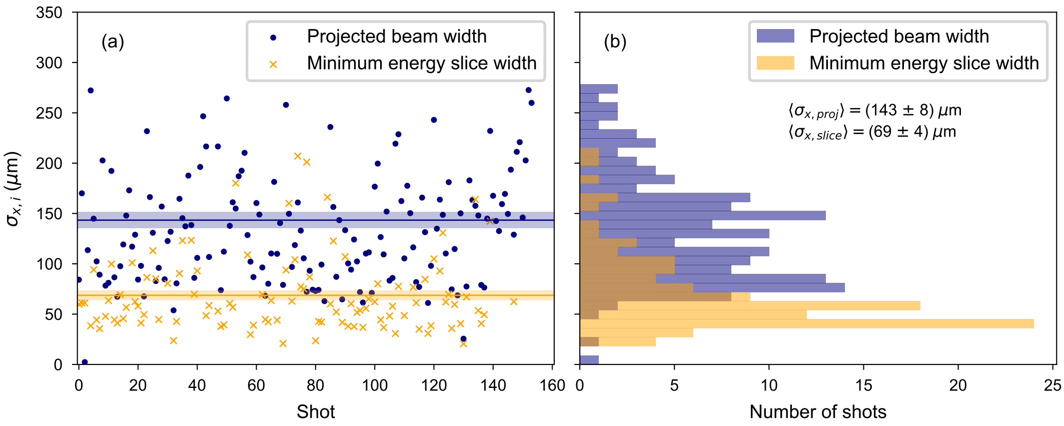

Figure 5.34: Width measurements of injected beams on the LEMS. (a) Projected beam widths (navy dots) and, for viable shots, minimum energy slice widths (orange crosses). The mean values from the two width definitions are displayed as two horizontal lines with colours that correspond to their respective data points; the shaded band around each represents the standard error on the mean. (b) Histograms of the projected (navy) and minimum energy slice (orange) widths, each made using the same 40 bins.

In addition to these multi-shot measurements of the emittance, another method was also attempted that could be applied on a shot-by-shot basis [183]. In this case, the chromaticity of the beam was used to our advantage, as beam slices with different energies could be attributed their own object planes (and magnifications) for a given set of quadrupole strengths. This essentially formed an object plane scan for a single shot, from which the waist size and divergence could be extracted and the slice emittance then calculated. The caveat of this method is that it necessarily relies on the assumption that the emittance of each slice was equal. Furthermore, the appearance of the 'butterfly' structure necessary to perform this analysis is not a given; only beams that exhibited a clear waist on the screen can reasonably be used. This required beams to be well-imaged, but also favoured those with a larger divergence, such that a variation in the slice size due to chromaticity could be detected over a relatively narrow energy bandwidth.

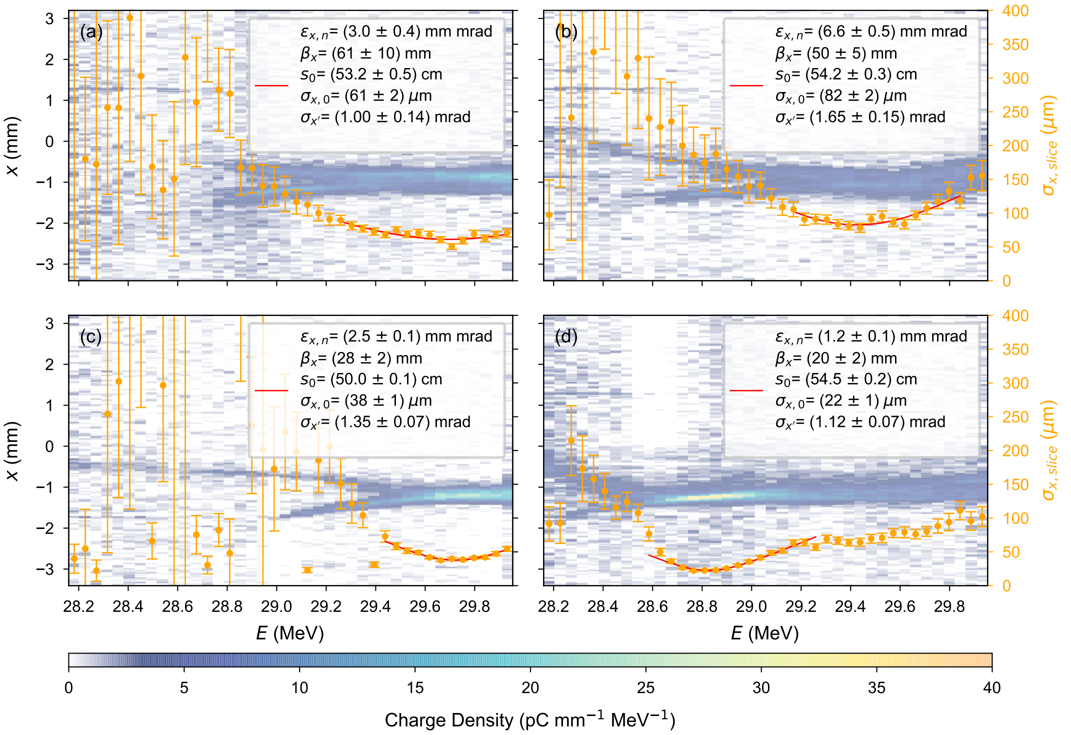

Figure 5.35: Single-shot 'butterfly' emittance analysis of internally-injected beams on the LEMS diagnostic. (a-d) Four examples of beams detected on the LEMS diagnostic to which the 'butterfly' emittance measurement was applied. For each of the 40 energy slices in a given image, a gaussian was fit to the x-profile; the width of each slice, divided by the slice magnification, is plotted as a function of the energy E (orange points), with the errorbars representing the error on the fitted gaussian widths. A fit of Equation 2.38 to the widths of the high-charge slices around the apparent waist of each beam is also plotted as a function of E (red curves), with the corresponding fit parameters and their errors displayed in the legend.

Figure 5.35 presents four beams to which this treatment was applied, as well as the results that the analysis yields. Again, for reasons of low signal-to-noise ratio, a median filter with a 3 × 3 kernel size was applied to each image, before dividing the image into 40 energy bins—each of width 0.05 MeV. Gaussians were then fitted to every energy slice, the widths of which divided by the slice magnification resulted in the slice beam sizes $\sigma_{x,i,slice}$. In a plot of $\sigma_{x,i,slice}$ versus the slice object plane $s_{obj,x,slice}$, a fit of Equation 2.38 to the high-charge, well-imaged energy slices yielded the waist size, divergence and slice emittance.

A common feature of the four beams presented in Figure 5.35 was the relatively high divergence, with $\sigma_{x\prime} \gtrsim 1\,\text{mrad}$ in all cases. Again, by exhibiting the butterfly structure required for this analysis, beams necessarily have a relatively high divergence, implying also large emittances. This is especially true if the measured waist size is also large; $\sigma_{x,0,slice} = (82 \pm 2)\,\mu\text{m}$ for the beam in Figure 5.35 (c), which combined with a large divergence $\sigma_{x\prime} = (1.7 \pm 0.2)\,\text{mrad}$ results in an emittance of $\epsilon_{n,x,slice} = (6.6 \pm 0.5)\,\text{mm} \cdot \text{mrad}$. On the other hand, the shot shown in Figure 5.35 (d) provides a more optimistic example. First of all, this beam acts as an exemplary candidate for this type of analysis, possessing a clearly imaged waist with high charge density, away from which the beam shows clear evidence of divergence. Moreover, the waist size of this beam is one of the smallest measured with this technique—$\sigma_{x,0,slice} = (22 \pm 1)\,\mu\text{m}$, with a peak spectral density of $(\frac{dQ_i}{dE})_{peak} = 8\,\text{pCMeV}^{-1}$ detected on the screen. Combining the waist size of this beam with its modest divergence $\sigma_{x\prime} = (1.1 \pm 0.1)\,\text{mrad}$ results in a slice emittance of $\epsilon_{n,x,slice} = (1.2 \pm 0.1)\,\text{mm} \cdot \text{mrad}$.

5.2.6 Brightness comparison

As a final step in the characterisation of the plasma cathode, the brightnesses of the injected bunches can be compared to that of the incoming drive beam. Due to a lack of information about the emittance of the beams in the y direction and also the length of the injected bunches, a full 6D brightness estimate $B_{6\text{D}}$ was impossible for either the driver or witness. However, using the properties that could be commonly measured for both beams—namely the charge Q, energy spread σ_E and horizontal emittance $\epsilon_{x,n}$—the *3D brightness* could instead be calculated as

$$B_{3\text{D}} = \frac{Q}{\epsilon_{x,n}\sigma_E} \approx \frac{(\frac{dQ}{dE})_{peak}}{\epsilon_{x,n}}, \tag{5.2}$$

where $\frac{Q}{\sigma_E}$ is approximated by the peak spectral density $(\frac{dQ}{dE})_{peak}$ for a given beam. Since $\epsilon_{x,n}$ was almost exclusively determined via multi-shot measurements and so roughly constitutes a shot-averaged measurement of the horizontal emittance, mean values of the peak spectral density are also used when calculating the 3D brightness.

For the June 2021 experimental run, the driver peak spectral density before plasma interaction was found on average to be $\langle(\frac{dQ_d}{dE})_{peak}\rangle = 47\,\text{pC}\,\text{MeV}^{-1}$. Although a direct driver emittance measurement was not available from June 2021, the results from April 2021 shown in Figure 5.5, where $\epsilon_{x,n,d} = (21.7 \pm 0.4)\,\text{mm} \cdot \text{mrad}$, are believed to be representative of the typical driver emittances used in the plasma cathode experiment. Combining these emittance and spectral density estimates using Equation 5.2 therefore yields a 3D driver brightness of $B_{3\mathrm{D},d} = 2.4\,\mathrm{pC/MeV/mm \cdot mrad}$.

In the case of the injected bunches, the average peak spectral density was measured to be $\left(\frac{dQ_i}{dE}\right)_{peak} = (14 \pm 6)\,\mathrm{pCMeV}^{-1}$ for the optimised working point presented in Figure 5.29. Characterisation of these beams on the ESPEC and LEMS then gave a number of estimates of their emittances (see Section 5.2.5), each resulting in different values of the witness 3D brightness $B_{3\mathrm{D},i}$ that are summarised in Table 5.1. Conservatively estimating the injected beam emittance via the shot-averaged minimum energy slice width detected on the LEMS yields $B_{3\mathrm{D},i} = 11.6\,\mathrm{pC/MeV/mm \cdot mrad}$—implying a 3D brightness boost compared to the driver by a factor $\frac{B_{3D,i}}{B_{3D,d}} = 4.8$. Moreover, using the modal minimum energy slice width from the same data results in a more optimistic estimate $B_{3\mathrm{D},i} = 17.5\,\mathrm{pC/MeV/mm \cdot mrad}$, amplifying the driver brightness by an factor of 7.3 on average. Additionally, the brightness estimated using the slice emittance $\epsilon_{x,n,slice}$ calculated for the shot in Figure 5.35 (d) is also included; its emittance, combined with its peak spectral density of $8\,\mathrm{pCMeV}^{-1}$ yields a brightness of $B_{3\mathrm{D},i} = 10.0\,\mathrm{pC/MeV/mm \cdot mrad}$, corresponding to a brightness boost by a factor of 4.2.

Emittance estimate	$\left(\frac{dQ_i}{dE}\right)_{peak}$	$\epsilon_{x,n,i}$ (mm · mrad)	$B_{3\mathrm{D},i}$ (pC/MeV/mm · mrad)	$\frac{B_{3\mathrm{D},i}}{B_{3\mathrm{D},d}}$
$\gamma \cdot \langle\sigma_{x,0,slice}\rangle \cdot \sigma_{x'}$	(14 ± 6)	(1.2 ± 0.1)	11.6	4.8
$\gamma \cdot (\sigma_{x,0,slice})_{mode} \cdot \sigma_{x'}$	(14 ± 6)	0.8	17.5	7.3
$\epsilon_{x,n,slice}$	8	(1.2 ± 0.1)	10.0	4.2

Table 5.1: Driver—witness brightness comparison. Each row represents a different method of estimating the emittance and therefore 3D brightness of the injected beams.

Even in the most conservative case, these results therefore act as a demonstration that the brightness of a drive beam in the 3D phase space defined by the energy dimension and a single transverse plane can be boosted via generation of a new, high-quality witness beam directly inside the plasma wake. It can be expected that an amplification in $B_{3\mathrm{D}}$ likely also implies an increase in $B_{6\mathrm{D}}$. For a driver that longitudinally fills the decelerating region of the blowout and a witness that is injected over all accelerating phases of the wake, the lengths of the two bunches can be considered approximately equal. Assuming that the driver emittance was similarly deteriorated in the vertical plane and that the emittances in each direction were approximately symmetric for the injected beam, a similar brightness increase can be expected from the reduced emittance in the vertical dimension. Although the witness emittance is expected to be slightly higher in y due to the asymmetry of the injection mechanism (see simulations in Section 5.3.1), it is not predicted to differ by orders of magnitude. Owing to a similar emittance reduction in both transverse planes, these estimates predict a 6D driver—witness brightness amplification by a factor of ~ 85. Naturally, to truly validate these predictions, experimental measurements are necessary. Possible methods for measuring the currently inaccessible dimensions of the beams involved here are therefore discussed as an outlook in Section 5.4.

It is also demonstrative to compare the quality of the injected bunches produced in this experiment with that from other internal injection experiments. Although producing beams of slightly higher mean energy ($\sim 50\,\mathrm{MeV}$) and higher charge ($\sim 30\,\mathrm{pC}$), the previous iteration of this experiment presented in Ref. [197] only managed to measure an injected bunch emittance of $\epsilon_{x,n,i} = (9.3 \pm 0.3)\,\mathrm{mm\,mrad}$. Whilst demonstrating an injection stability that was similar or even better than in Ref. [197], the major step of the current work was to produce and measure injected beams with around a factor of 10 lower emittance—owed primarily to improved optimisation techniques and better characterisation e.g. with the high resolution LEMS.

One can also draw a comparison with the internal injection experiment at FACET reported in Ref. [215]. Here, a best-case estimate of the emittance of $\epsilon_{n,x,i} \approx 1\,\mathrm{mm\,mrad}$ was found—a similar level to the conservative estimates resulting from the multi-shot characterisation in this work. On the other hand, the absolute spectral density of these bunches is lower than in this work; a maximum of around $0.45\,\mathrm{pC\,MeV^{-1}}$ is shown in Figure 4 (a) of Ref. [215]. A major difference between the work here and in Ref. [215] is the much lower energy gain achieved in the former; it is possible that the relatively high spectral density achieved in the current experiment would decrease if accelerated to higher energies in a non-uniform accelerating field. As a fairer comparison, we can therefore compare the spectral densities in units of charge divided by *percentage energy spread*. In the case of Ref. [215] with energy gains of $\sim 500\,\mathrm{MeV}$, the peak spectral density is estimated as $\left(\frac{dQ_i}{dE}\right)_{peak} = 0.45\,\mathrm{pC\,MeV^{-1}} \times 500\,\mathrm{MeV}/100 = 2.25\,\mathrm{pC\,\%^{-1}}$. For the current work with energy gains on the level of $\sim 30\,\mathrm{MeV}$, a similar calculation yields $\left(\frac{dQ_i}{dE}\right)_{peak} = 4.2\,\mathrm{pC\,\%^{-1}}$—still larger than Ref. [215], although by a lower amount than when considering absolute energy spreads.

Ultimately, a future challenge will be to maintain the relatively high absolute spectral density of the injected bunches in this experiment whilst accelerating them to higher energies. The low witness energies in this experiment were attributed to the method of generating the longitudinal plasma, with a more thorough understanding obtained through simulations presented in Section 5.1.3. Additional results in the next section however demonstrate a promising alternative for forming the main channel in a plasma cathode experiment that can overcome the limits of the laser-only method discussed so far.

5.2.7 Discharge preionised case

During the June 2021 experimental run, it was found that by firing the discharge in addition to both lasers resulted in a significant increase in energy of the injected beams. Figure 5.36 shows witness beam spectra for a discharge timing scan, measured on the low energy ESPEC screen. Peculiarly, in this case no injected charge was detected with the longitudinal laser blocked. With the discharge arriving after the drive beam and lasers, the plasma density that the driver initially experienced was identical to the laser-only case. Firing the discharge such that its arrival preceded that of the driver then began to increase the energy gain of the injected bunches up to a measured maximum of $\sim$ 180 MeV. This was presumably a result of an increase in the degree of ionisation; by scanning the rising edge of the discharge pulse over the time-of-arrival of the driver, the average plasma density in which the wake was driven increased, along with the average accelerating gradient. This came at the expense of charge, with only $\sim$ 1 pC-level beams being detected for the highest energies. Whilst a non-optimal imaging energy for these beams could be partly responsible for this, it should be emphasised that the driver was still optimised for the laser-only case, and not for the completely different plasma that the additional discharge inevitably generated. One may expect that the density profile generated by the discharge—which increases with the rising current of the discharge pulse—is slightly modified by the longitudinal laser, which will also be able to ionise neutral gas beyond the discharge electrodes at the exits of the capillary. From the perspective of the injector, it is also unclear how much of the upper levels of argon remained unionised after the discharge for the transverse laser to generate the density downramp. For an improved understanding of this density profile, detailed simulations of both the discharge plasma which the lasers then propagate through are ultimately required, but are beyond the scope of this work. Nevertheless, this data demonstrates that a discharge is compatible with optically-generated downramp injection, and with it, energy gains that are more expected of a PWFA can be achieved.

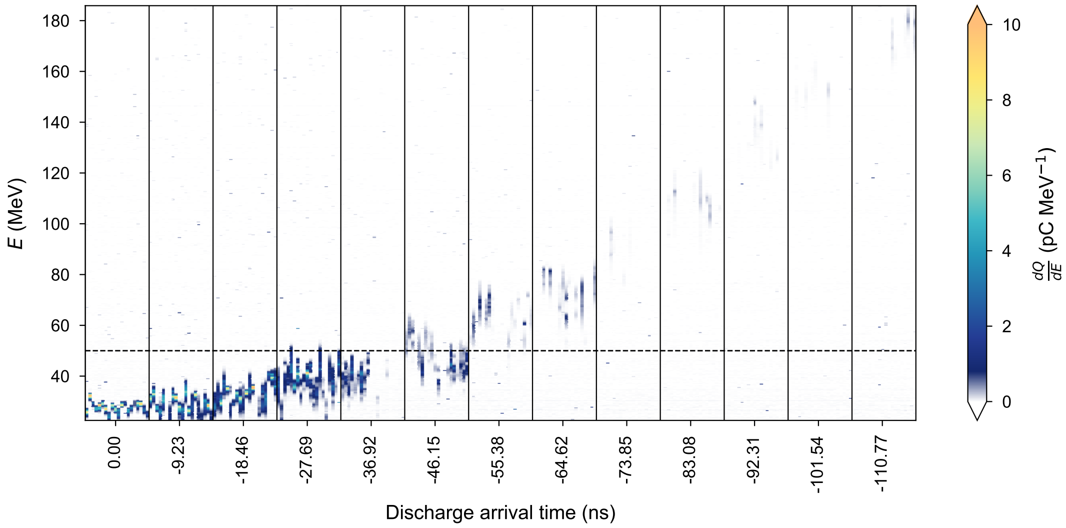

Figure 5.36: Injection dependence on discharge timing in the laser-plus-discharge preionised plasma cathode. Witness beam energy spectra measured on the ESPEC diagnostics show an increasing energy gain as the discharge is brought earlier in time. Throughout the scan, the imaging energy remains at 50 MeV (horizontal black dashed line).

In the September 2021 experimental run, another attempt at a discharge-assisted plasma cathode was made—this time demonstrating that internal injection was possible *without* the need for the longitudinal laser. Figure 5.37 displays one of a handful of beams that were detected on the low energy ESPEC screen, with only the discharge and the transverse laser firing. Like those in Figure 5.36, the energy gain of the beams here already more than doubled those achieved in the laser only case. As in the laser-plus-discharge case, the spectral density of the bunches was also not very high, with a peak of only 0.3 pCMeV^{-1} for the 2 pC-charge beam in Figure 5.37. Similar reasons for the low amounts of injected charge apply here as for the data in Figure 5.36, but conditions of the injection on this occasion were also made less clear by the fact that the discharge triggered the desynchronisation of the laser to FLASH—likely the result of an electromagnetic disturbance from the discharge pulser unit that interfered with the laser synchronisation electronics in the room above. Ultimately, this meant that the relative time-of-arrival of the injector laser to that of the driver and discharge was completely unknown. In spite of these drawbacks, this data demonstrates that a plasma cathode can in principle be operated with discharge preionisation and a low energy injector laser. The advantages of such an operation mode, and how it could be achieved in the future, are discussed in Section 5.4.

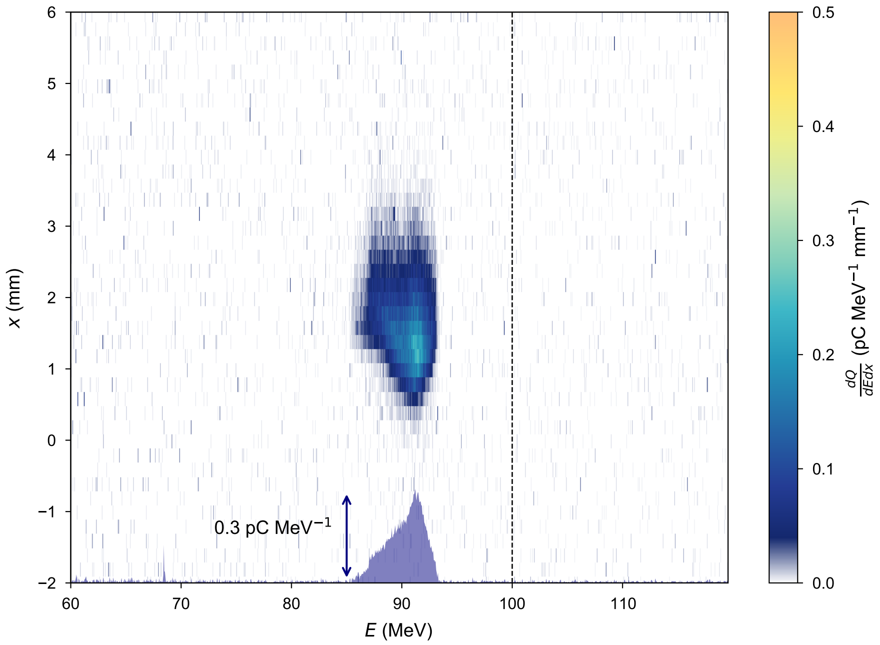

Figure 5.37: Example injected beam from a purely discharge-preionised plasma cathode stage. The spectrum of the bunch is also shown (navy shaded curve) , with a peak spectral density of 0.3 pCMeV^{-1}. The imaging energy for this measurement was set to 100 MeV (black dashed line).

5.3 Simulations

In the following section, a variety of simulations are presented that are intended to give a deeper understanding of the internally injected witness beams that were measured in experiment and also provides information about properties of the injected bunches that were impossible to measure in the experiment.

5.3.1 Plasma cathode stage modelling

Simulations of the PWFA-based plasma cathode stage were performed in the full 3D PIC code WarpX [150]. A 200 µm × 200 µm × 400 µm co-moving simulation domain was used, with cell sizes $dx = dy = 0.39$ µm and $dz = 0.14$ µm shown to produce convergence in key parameters e.g. the injected bunch charge and emittances.

Based on that characterised for the June 2021 experimental run in Section 5.1.1, the 304 pC, 689 MeV drive beam in this simulation was initialised in vacuum with

3×10^5 equal-weight macroparticles. The mean profile in Figure 5.2 (b) was used as the current-profile of the driver. For simplicity, the driver was modelled with a gaussian energy distribution of rms width of 5 MeV. Similarly, the transverse phase space of the beam was gaussian and symmetric in x and y, with emittance $\epsilon_{n,x,d} = \epsilon_{n,y,d} = 20\,\mathrm{mm\,mrad}$; the driver was initialised such that it reaches a symmetric focus that would have an rms width in vacuum of $\sigma_{x,0,d} = \sigma_{y,0,d} = 16\,\mu\mathrm{m}$ at the location $z = 0$ i.e. the start of the plasma cell.

The plasma profile used in the simulation is based on that derived from laser-plasma generation simulations in Figure 5.13. For simplicity, the case where the driver propagates on-axis with the main plasma column is modelled; the main plasma channel is assumed to vary only in the longitudinal direction [i.e. as in Figure 5.13 (b)], whereas the transverse plasma's shape in (y,z) is still fully captured [see inset of Figure 5.13 (b)]. The profile continues until $z = 60\,\mathrm{mm}$, after which the exiting driver and injected bunch propagate for around an extra 10 mm in vacuum. For the majority of the plasma, $1 \times 1 \times 1$ macroparticles per cell in (x, y, z) are used. This is with the exception of a 400 µm-long region centred on the middle of the plasma density spike, which has a higher macroparticle-per-cell number of $1 \times 1 \times 4$—i.e a 4 times higher macroparticle density along the z axis—in order to better sample the injected bunch. Although a non-zero plasma temperature is known to effect the fine structure of the wake [222] and the fields therein [223–226], the initial plasma temperature in these simulations was assumed to be 0 K for simplicity—a common practice for simulations of the plasma acceleration of electrons. Whether this would effect the injection and acceleration of electron bunches in this particular case was out-of-scope for this project, but could be investigated in the future.

Figure 5.38 summarises the key events in the plasma cathode stage simulation. As depicted in Figure 5.38 (a), the driver is focused into the plasma and drives a wake in the relatively high density region at the start of the plasma cell. At a longitudinal position of $z = 20\,\mathrm{mm}$, the driver reaches the density spike formed by the transverse laser, where the plasma wake rapidly contracts before again extending on the down-ramp [see Figure 5.38 (b)]. This promotes the injection of plasma electrons into the wake, some of which become trapped and accelerated. For a short distance after injection, the injected witness bunch is located at the back of the plasma wake with some evidence of beam loading apparent in the flattened shape of the longitudinal electric field (red lineout), as shown in Figure 5.38 (c). This however is short-lived; as the plasma density decreases with propagation, the wake expands such that the witness is no longer located near the wake vertex and instead experiences a gradually reducing accelerating field [see Figure 5.38 (d)]. Eventually, the driver and injected bunches reach the end of the plasma and are released into vacuum.

The next sections cover the key aspects of the plasma cathode stage simulation in more detail, namely the generation of the plasma wake by the driver, followed by the injection, acceleration and extraction of the injected witness bunches.

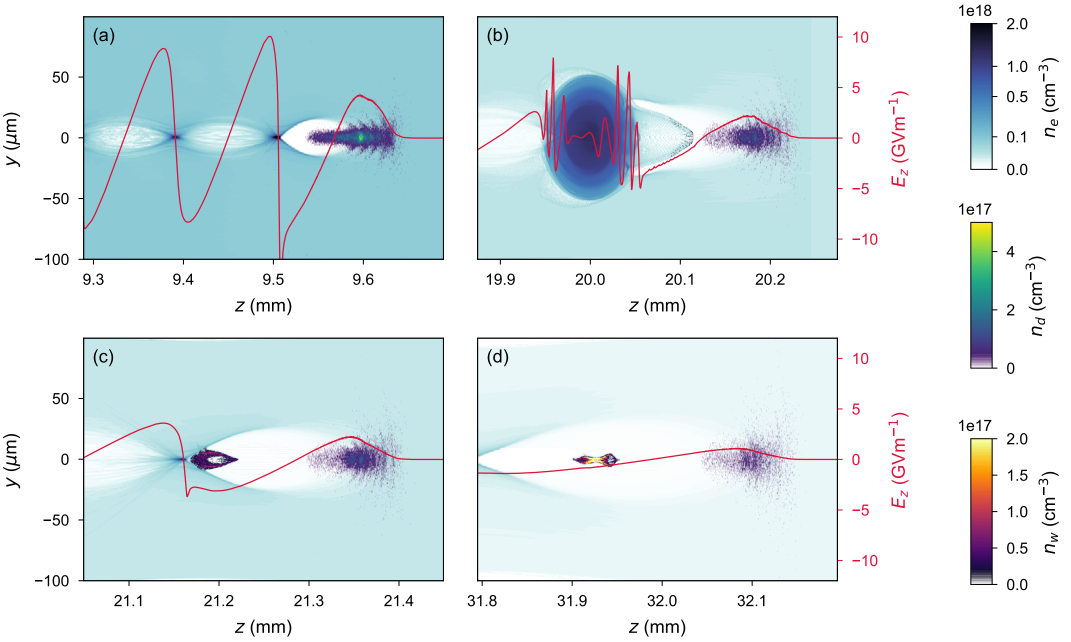

Figure 5.38: Plasma cathode stage simulation snapshots. Four steps in the simulation are shown: (a) initial wake generation by the driver, (b) downramp injection of witness beam electrons, (c) acceleration of an injected bunch at the back of the wake with beam loading and (d) acceleration of the witness electrons in a low density plasma. The simulation is visualised by plotting a slice at $x = 0$ of the driver n_d , plasma n_e and injected n_i electron number densities. A lineout of the longitudinal electric field E_z along $(x, y) = (0, 0)$ is also plotted (red lines).

Wake generation

Figure 5.39 shows how the drive beam is coupled into the plasma in the simulation and its transverse dynamics within the wakefield that it generates. The evolution of the driver beam width in the absence of plasma is represented by the navy dashed line in Figure 5.39 (a). However, with the plasma present (filled blue curve), the driver experiences an extra focusing force to due plasma lensing (navy curve), reaching a minimum rms width of 9.5 μm at $z = 2$ mm. After this however, the driver continually expands as the focusing force of the wakefield driven in the decreasing plasma density profile is smaller than the effective defocusing force due to the beam's large emittance (see Equation 2.71). At all points along the accelerator, the drive beam width is less than the plasma skin depth k_p^{-1} [orange line in Figure 5.39 (a)]—one of the preconditions for driving a non-linear wakefield [77, 227].

The evolution of the longitudinal wakefield E_z along the propagation axis is depicted by the waterfall plot in Figure 5.39 (b). The co-moving coordinate $z - ct$ is defined such that ct is the longitudinal centroid position of the driver. In the first $\sim$ 10 mm after the start of the plasma cell, wakefields of amplitude $\mathcal{O}(10\,\mathrm{GVm}^{-1})$ are driven in the first blowout; initially, multiple periods of the plasma oscillation are present in the simulation box. However, as the plasma density decreases, the wake expands, leading also to a decrease in the wakefield strength. For example, at $z = 30\,\mathrm{mm}$ the wakefield amplitude has decreased to $\mathcal{O}(1\,\mathrm{GVm}^{-1})$, and the first wakefield period only just fits into the simulation box. This behaviour is interrupted only momentarily at the injection point due to the high density plasma spike, but this has no discernible effect on the driver evolution downstream.

Figure 5.39 (c) shows the consequences of the wakefield evolution on the slice transverse size of the drive beam $\sigma_{x,slice,driver}$. Slices towards the front of the beam are focused into the plasma, before expanding due to the driver's emittance. This expansion is indeed faster for charge at the very head of the beam that resides outside of the wake cavity and so does not experience the focusing force of the exposed plasma ions. At the tail, the full longitudinal extent of the driver does not initially fit inside the first wake cavity; some slices are, for a period of time, located behind the first wake vertex (solid red line) and as such experience a defocusing force, causing their $\sigma_{x,slice,driver}$ to increase before eventually being lost from the simulation. This defocusing is also evidenced by a temporary rapid increase in the projected beam size $\sigma_{x,driver}$ [see Figure 5.39 (a)], before decreasing again as the charge is lost. Slices that are contained within the first wake but are located behind the first zero crossing (dashed red line) will at that point reside in the accelerating region of the wakefield, actually gaining energy. This energy gain, as well as the energy loss that most of the driver electrons experience, is summarised in Figure 5.40.

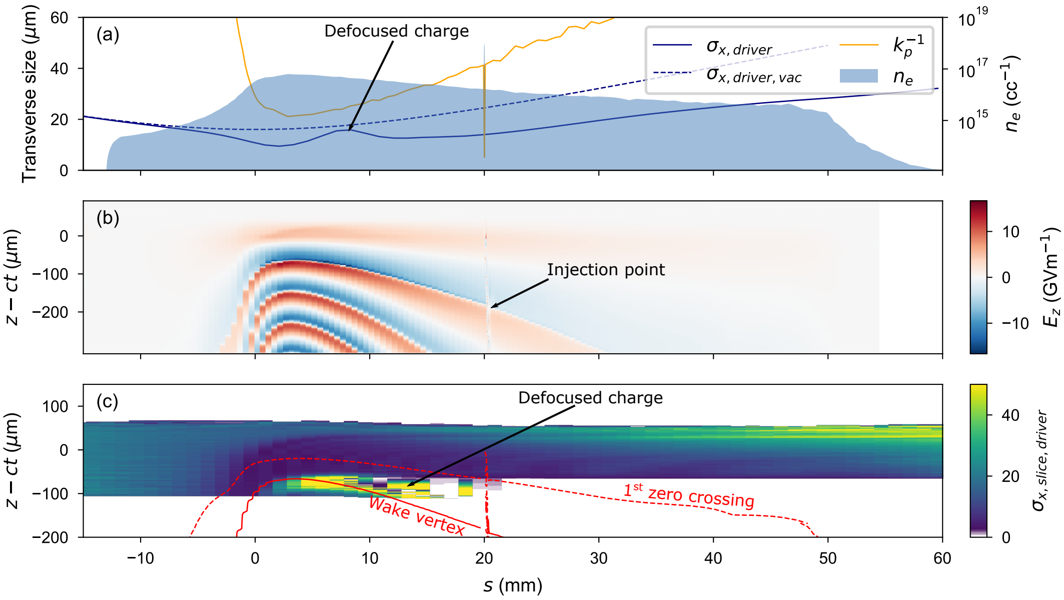

Figure 5.39: Wake generation in the plasma cathode stage simulation. (a) The evolution of the driver rms width in vacuum $\sigma_{x,driver,vac}$ (dashed navy line) and in the plasma $\sigma_{x,driver}$ (solid navy line) are shown. For comparison, the skin depth k_p^{-1} (orange line) calculated from the plasma density n_e (blue filled curve) is also plotted. (b) The longitudinal electric field E_z is plotted as a function of propagation distance, decreasing in amplitude as n_e reduces. (c) A waterfall of the slice driver width $\sigma_{x,slice,driver}$ variation with propagation is shown. Extracted from the plot in (b), the positions of the first zero crossing of E_z (dashed red line) and the first wake vertex (solid red line) are also plotted.

Figure 5.40 (a) shows the evolution of the driver energy spectrum as a function of propagation distance. As the driver enters the plasma, a large amount of the charge decelerates; a small amount resides for a time in the accelerating region of the wake, reaching energies as high as $\sim$ 800 MeV. Due to the decreasing plasma density, the rate of deceleration decreases as the driver propagates. Charge that was originally accelerated also eventually becomes placed in the decelerating region of the expanded wake and so is decelerated back down to lower energies.

The longitudinal phase space of the drive beam at the end of the simulation is depicted in Figure 5.40 (b). As highlighted, the final, maximum energy loss of the driver electrons was around 130 MeV—somewhat larger than the experimental measurement shown in Figure 5.18. This discrepancy can likely be attributed to the idealised, simplified input beam used in this simulation, as well as any imperfect modelling of the plasma profile. On the other hand, the experimental measurement is also imperfect, where only a finite range of energies can be imaged simultaneously. Overall however, these results indicate that the main features of the drive beam interaction with the plasma are approximately captured.

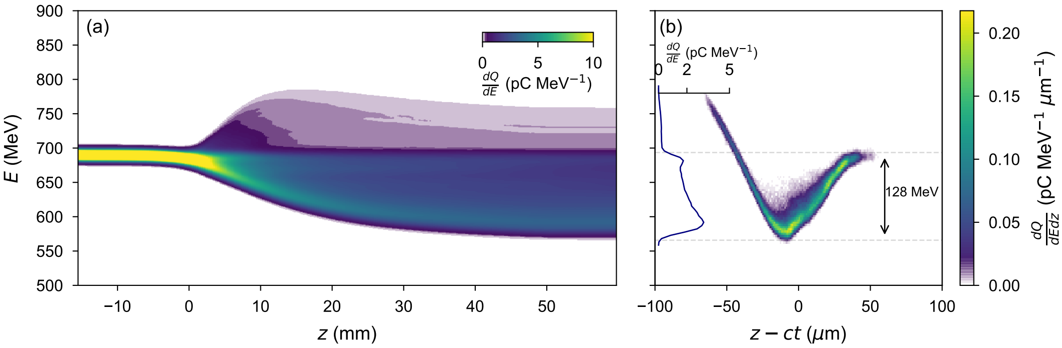

Figure 5.40: Driver deceleration in the plasma cathode stage simulation. As shown by the evolution of the drive beam spectrum in (a), the overall driver loses energy to the wakefield, with a small amount of electrons located in the accelerating region of the wake gaining energy—resulting in the longitudinal phase space at the end of the simulation depicted in (b).

Longitudinal phase space evolution of injected bunches

After injection into the wake, the trapped witness bunch is accelerated by the wakefield. This acceleration process is summarised in Figure 5.41. As shown by the witness energy spectrum evolution in Figure 5.41 (a), the mean energy and rms energy spread of the injected bunch increases to values of 27.3 MeV and 2.0 MeV respectively. The rate of growth of both of these quantities depends on the plasma density n_e [filled blue curve in Figure 5.41 (a)] at the given position, determining both the wakefield amplitude and also the accelerating phase at which the injected bunch is located within the blowout.

Figure 5.41 (b) shows a similar plot of the longitudinal electric field as that in Figure 5.39 (b) but with more emphasis on the the accelerating field that the injected bunch experiences. The position and extent of the witness bunch in the co-moving coordinate is represented by the orange band; after acceleration to relativistic velocities, the bunch stays at approximately the same comoving position relative to the driver. As the plasma density reduces however, the decreasing wakefield amplitude, as well as the fact that the witness is located at phases further forward in the expanding wake, decreases the acceleration gradient that the bunch experiences until the evolution of its energy spectrum essentially plateaus.

As well as the overall, correlated energy spread of the injected bunch, also of interest is its uncorrelated, slice energy spread $\delta E_{rms,slice}$—the evolution of which is depicted in Figure 5.41 (c). Just after the formation of the witness, the slice energy spread across the bunch is on the level of $\sim$ 0.1–0.2 %. The energy spread of slices at the front of the bunch however increases with propagation distance, reaching values of $\sim$ 0.4 % by the end of the plasma. An apparent reason for this is phase slippage; note how in Figure 5.41 (c), the bunch slices appear to move backwards in the co-moving frame. This is more pronounced for slices at the head of the bunch, implying differential phase slippage across the bunch's length. A consequence of this is that the bunch becomes slightly shorter as electrons originally located further forward begin to slip into slices further back. These particles will also have experienced a lower energy gain to those at later accelerating phases, such that the slippage of these electrons into later slices leads to an increase in the slice energy spread. This is due to the bunch having a spread of velocities due to its energy spread, with electrons at the head travelling slower. Specifically, once the bunch has reached $\beta_z \sim 1$ and is near-phase-locked at around $z = 30\,\mathrm{mm}$, the rms spread of the longitudinal velocities due to the bunch energy spread is given by $\sigma_{\beta_z} \approx 2.7 \times 10^{-4}$, implying a decrease in the rms bunch length of around 0.5 µm between $z = 30\,\mathrm{mm}$ and $z = 50\,\mathrm{mm}$. This is in agreement with the bunch length decrease observed in simulation; whereas the rms bunch length is around $\sigma_z = 10.7\,\mu\mathrm{m}$ at $z = 30\,\mathrm{mm}$, the bunch length decreases to 10.2 µm by $z = 50\,\mathrm{mm}$, suggesting that the phase slippage is a result of the bunch's velocity spread rather than e.g. betatron motion [228].

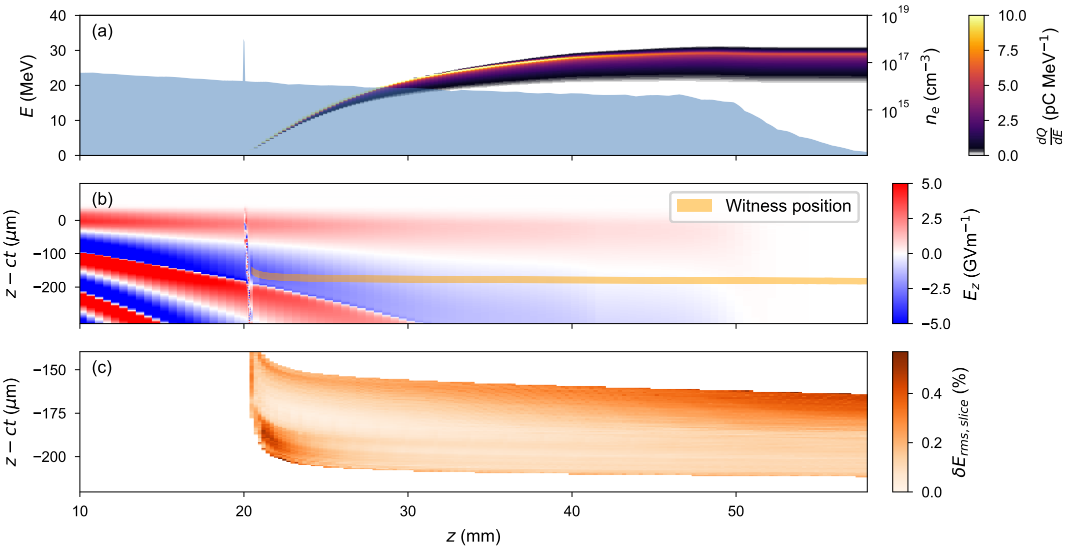

Figure 5.41: Acceleration of injected bunches in the plasma cathode simulation. (a) The evolution of the injected bunch's spectrum is shown, overlaid with the plasma density profile (filled blue curve). The decreasing rate of acceleration can be explained by the evolution of the wakefield shown in (b), with the location of the witness bunch therein represented by the orange band. (c) Due to phase slippage of the witness electrons, the slice energy spread also evolves during acceleration.

Transverse phase space evolution of injected bunches

As the injected bunch accelerates, its transverse phase space also evolves. The evolution of key transverse properties of the injected bunch—namely its width, divergence and emittance—are summarised in Figure 5.42.

After injection, the beam width and divergence oscillate, as shown by the solid blue and red lines in Figure 5.42 (a) and (b). Overall, the size of the beam envelope and also its divergence is larger in y than in x, presumably due to the curvature of the downramp in the y-direction. As the injected bunch propagates into regions of lower plasma density, the restoring focusing force of the plasma ions decreases, allowing the beam to slightly expand whilst simultaneously reducing its divergence. Beyond the end of the plasma cell ($z > 50$ mm), the plasma density decreases gradually to zero; in this region, the divergence of the beam in x and y is reduced from around 1.9 mrad and 4.1 mrad to 0.8 mrad and 1.4 mrad respectively. Note that beyond $z = 60$ mm—where the divergence has already began to plateau—the plasma density is assumed to be zero. After this the divergence no longer decreases and the beam width grows as the bunch diverges into free-space.

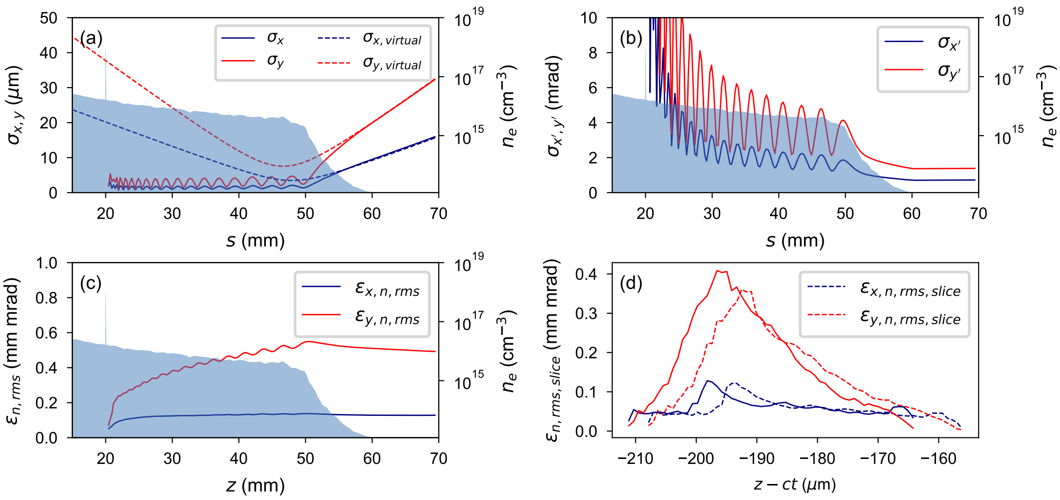

Figure 5.42: Transverse phase space evolution of injected bunches in the plasma cathode stage simulation. The evolution of the injected bunch's rms (a) width, (b) divergence and (c) normalised emittance is plotted, with the blue and red solid lines representing the respective quantities in each direction. In addition, the dashed lines in (a) depict the virtual waist of the outgoing injected bunch. Slice emittances across the bunch are also plotted in (d) for $z = 30$ mm (dashed lines) and $z = 50$ mm (solid lines).

By back-propagating the injected bunch's Twiss functions, the virtual waist of the beam can be plotted in both dimensions (red and blue dashed lines in Figure 5.42 (a)). Object plane scan measurements of the beam should in principle reconstruct this virtual waist, which shares the same divergence as the bunch has in vacuum but has a minimum beam size that is necessarily larger than the bunch width inside the wake. In this case, the virtual waist sizes in x and y are $\sigma_{x,virtual,0} = 3.5\,\mu\text{m}$ and $\sigma_{y,virtual,0} = 7.5\,\mu\text{m}$.

Figure 5.42 (c) shows the evolution of the normalised emittance in x and y throughout the plasma stage. Again, there is an asymmetry between the two transverse directions; the emittance in y is greater than that in x. In both x and y the beam also experiences some projected emittance growth, up to values of 0.14 mm mrad and 0.55 mm mrad respectively at $z = 50$ mm. In order to determine the source of this increase, the slice emittance can also be studied.

The lines in Figure 5.42 (d) depict the slice emittance in x (blue) and y (red), at propagation distances of $z = 30$ mm (dashed) and $z = 50$ mm (solid). Although there is some increase in the slice emittance—the average slice emittance in y for example increases from 0.14 mm mrad to 0.17 mm mrad—this growth likely arises from the phase slippage of electrons between slices, in a similar way to the increase in slice energy spread. Moreover, the slice emittance change due to phase slippage would not impact the projected emittance and even for the maximum emittance slice in y, its value only grows to a value of 0.4 mm mrad and so cannot account for the full increase of the projected emittance growth in that direction.

Rather than an increase in the intrinsic emittance of the bunch, this growth in projected emittance appears more consistent with that due to *betatron decoherence*, which also explains the oscillations that are visible in $\epsilon_{x,n}$ and $\epsilon_{y,n}$ [31]. Figure 5.43 explores this by extracting the Twiss parameters of the individual slices of the injected bunch and using these to plot their corresponding trace space ellipses. As shown in Figure 5.43 (a) and (b), the slice ellipses at $z = 22\,\text{mm}$ in both x and y trace spaces are relatively well aligned, with the exception of some low-charge slices.

However, this situation changes as the bunch propagates. Figure 5.42 (c) and (d) plot the evolution of the slice Twiss α function for each direction $\alpha_{x,slice}$ and $\alpha_{y,slice}$—a parameter quantifying the position-angle correlation and thus the orientation of the slice trace space ellipses. As with the oscillations in the beam width and divergence in Figure 5.42 (a) and (b), the oscillatory behaviour of the α function shows that the slices were at no point matched to the plasma; if this were the case, this parameter would equal zero. In addition, the particles in different slices undergo betatron oscillations at different frequencies across the length of the bunch, highlighted by the curvature of the vertical bands in the plots of $\alpha_{x,slice}$ and $\alpha_{y,slice}$. Due to the fact that the wake is driven in the blowout regime, the focusing field that each bunch slice experiences is the same. The reason for the variation in betatron frequency is therefore due to the energy spread of the bunch, rather than the bunch length. In any case, the result is that as the bunch propagates, the slice trace space ellipses will rotate at different rates, resulting in an overall larger projected emittance.

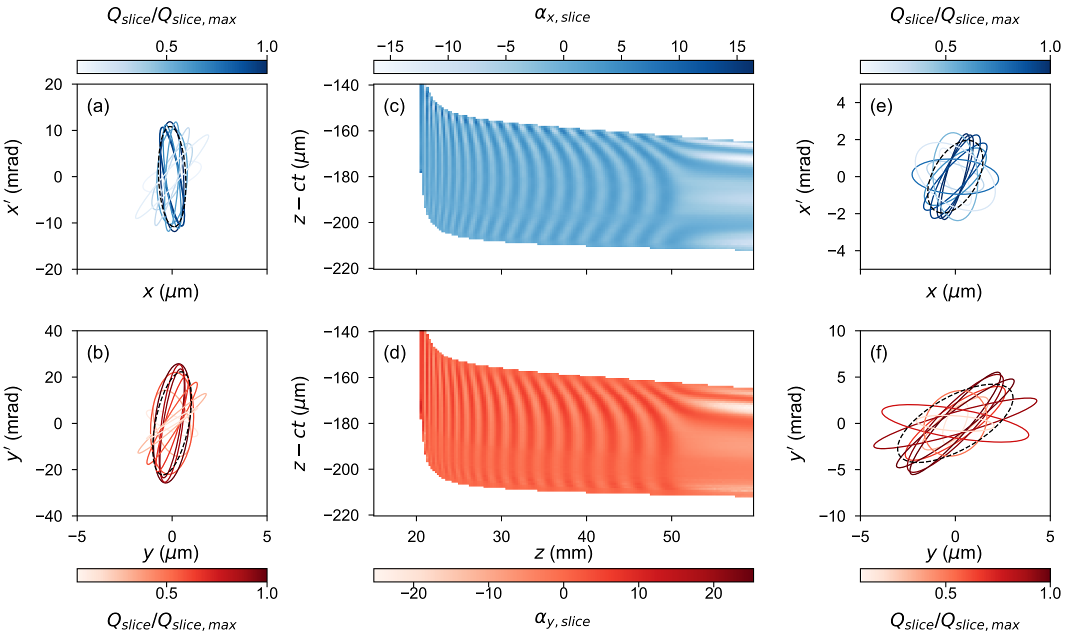

Figure 5.43: Betatron decoherence of injected bunches in the plasma cathode stage simulation. At a position in the plasma of $z = 22\,\text{mm}$, the slice trace space ellipses in the (a) x-x' and (b) y-y' planes have approximately the same orientation, with the projected trace space ellipse in each plane plotted with a black dashed line. As the bunch propagates, different slices oscillate with different frequencies, as shown by the evolution of (c) $\alpha_{x,slice}$ and (d) $\alpha_{y,slice}$. At $z = 47\,\text{mm}$, the slice phase trace ellipses were therefore misaligned in both (e) x-x' and (f) y-y', leading to projected ellipses (black dashed lines) with larger areas (i.e. larger emittances). Note that all of the ellipses shown are colour-coded based on the charge in their respective slice Q_{slice}, normalised by that of the maximum-charge slice $Q_{slice,max}$.

This is illustrated by the slice trace space ellipses in Figure 5.43 (e) and (f), taken at $z = 47\,\text{mm}$. Observe that the ellipses are in general less well aligned than those in Figure 5.43 (a) and (b), which includes also those slices with significant charge. Although the ellipses have not fully decohered (which would lead to a saturation in the emittance growth), this suggests that betatron decoherence is indeed the primary source of the in-plasma projected emittance increase observed in these simulations.

Note that a related source of emittance growth is that due to misalignment, as explained in Section 2.2.2 on page 16. Any misalignment of the injected bunch—for example due to any asymmetries in the wake as it transitions the density downramp—could therefore lead to an even higher growth in the projected emittance than that shown here.

Another potential source of in-plasma emittance growth that is not explicitly modelled in the PIC simulation is that due to scattering of the injected beam by on-axis atoms and ions [see Section 2.2.2 on page 15] [39, 229, 230]. As the plasma generated consists mainly of relatively heavy, singly ionised argon ions, the ion scattering term in Equation 2.35 is considered negligible. The neutral atom scattering term may however contribute to emittance growth; due to the fact that only a small fraction of the gas is ionised over the later parts of the plasma cell, a significant amount of neutral atoms are present in the path of the injected bunch.

Equation 2.35 is therefore applied—considering only neutral atom collisions—to the beam parameters extracted from the injection simulation in order to estimate the emittance growth in each direction due to scattering, the results of which are presented in Figure 5.44. The neutral density n_{gas} (grey filled area in Figure 5.44) is inferred by subtracting the longitudinal plasma profile from Figure 5.11 from the initial neutral gas profile in Figure 5.10. Although initially high as the bunch is injected, the emittance growth per unit distance in both planes, $\frac{d\epsilon_{x,n}}{ds}$ and $\frac{d\epsilon_{y,n}}{ds}$ [red and navy dashed lines in Figure 5.44], on average decreases as the bunch accelerates. Due to its dependence on the beta functions in each plane $\beta_{x,y}$ [navy and red solid lines in Figure 5.44], the emittance growth-rate oscillates with the bunch width. Although the beta functions naturally expand at the plasma exit, the low amount of neutral gas at this point means that the emittance growth-rate does not increase dramatically.

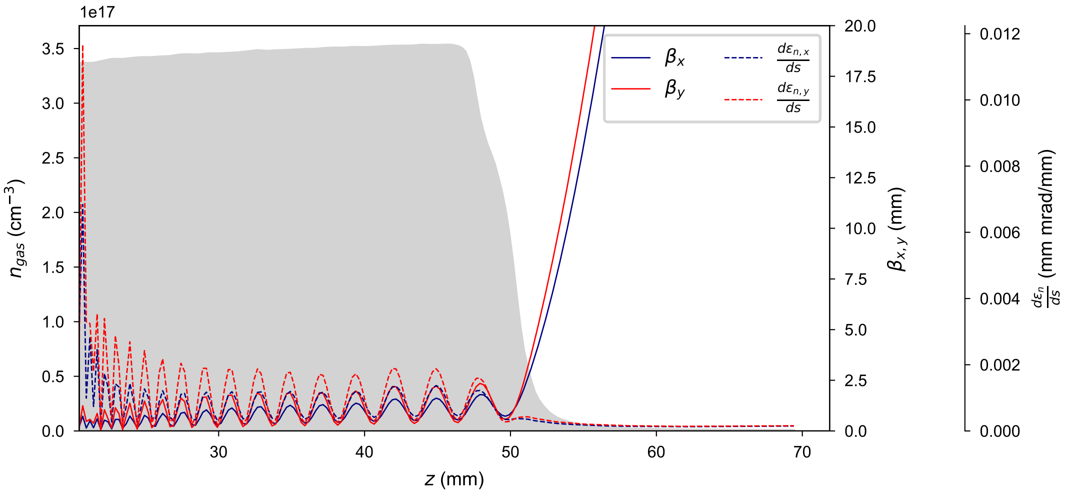

Figure 5.44: Neutral gas scattering-induced emittance growth estimates in the plasma cathode stage simulation. The emittance growth rate in x and y (blue and red dashed lines) is calculated from the beta functions in the respective directions (blue and red solid lines) and the remaining on-axis neutral gas density after laser ionisation (grey filled curve).

By integrating $\frac{d\epsilon_n}{ds}$ over all s, the total amount of emittance growth due to scattering is found to be $\Delta\epsilon_{x,n} = 0.03\,\mathrm{mm\,mrad}$ and $\Delta\epsilon_{y,n} = 0.04\,\mathrm{mm\,mrad}$ in x and y respectively—on the level of 10 % of the final injected beam emittance in the PIC simulations. This in principle means that neutral gas scattering could contribute a non-negligible amount to the in-plasma emittance growth in experiment. To explore this further, the effect of scattering could be computed during the PIC cycle—a capability that is indeed offered by WarpX [231, 232]—such that any evolution of e.g. the injected beam's beta functions due to the scattering-induced emittance degradation are taken into account in the calculations.

Finally, note that increasing the ionisation degree of the plasma would have a two-fold positive effect in mitigating emittance growth from neutral gas collisions. Firstly, a higher ionisation degree naturally means a lower amount of neutral gas for the beam to collide with. Secondly, the bunch would be accelerated by a greater average accelerating field, leading to higher energies and thus further reducing the effect of scattering. Any neutral gas scattering-induced emittance degradation present in this experiment is therefore yet another consequence of the sub-optimal plasma generation method.

Injected bunch overview

At this stage, it is instructive to review the key properties of the injected bunches at the end of the simulated plasma stage and also compare them with experimental results. Figure 5.45 depicts the longitudinal phase space of the bunch at a position $z = 50.4\,\mathrm{mm}$, with key properties such as the spectral density (filled purple curve) and slice quantities also plotted. Remembering that the working point being modelled is one in which the plasma has not been allowed to evolve after it is created, the simulation results here should be compared to those in experiment where the lasers arrive just before the drive beam.

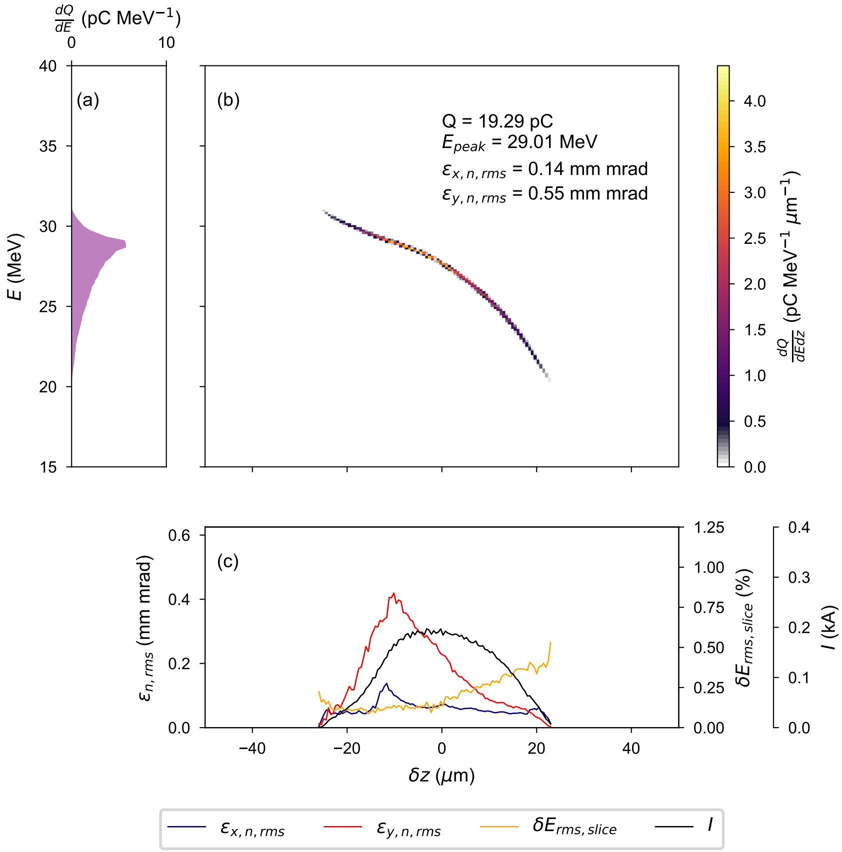

Figure 5.45: Injected bunch spectrum, longitudinal phase space and slice properties at the exit of the simulated plasma cathode stage. The injected bunch (a) spectrum and its corresponding (b) longitudinal phase space are plotted, where E is the particle energy and δz is the longitudinal position relative to the bunch centroid. Included also in (c) are lines representing the slice emittance in x and y (blue and red), the slice energy spread (orange) and the current (black).

In the case of the injected charge, the charge $Q = 19.3\,\text{pC}$ simulated here should be compared to the first scan point corresponding to a delay $t_{driver} - t_{laser} = 0$ in Figure 5.27 on page 139, where the average charge is measured to be $Q_i = (3.7 \pm 3.2)\,\text{pC}$ and the maximum charge (not shown) is 15.6 pC. As the simulation represents a rather idealised case (e.g. with optimal driver—downramp alignment and a transversely symmetric driver), it is reasonable to make the comparison to the maximum measured injected charge , which is closer than the average but still around 20 % lower than in the simulated case.

In terms of the spectrum of the injected bunch, the peak spectral density $\left(\frac{dQ}{dE}\right)_{peak}$ in the simulation is 5.7 pCMeV^{-1}. In spite of the higher charge in simulation, this is in fact lower than that of the maximum charge bunch from experiment discussed previously, where $\left(\frac{dQ_i}{dE}\right)_{peak} = 10.3\,\mathrm{pCMeV}^{-1}$. This is due to a discrepancy in the energy spread of the two spectra, which is lower in the experimental case (1.5 MeV FWHM) than in simulation (2.4 MeV FWHM). On the other hand, the energy gain of the simulated injected bunch is remarkably consistent with experiment. As found in nearly all of the experimental results that use only the longitudinal laser to preionise the plasma cathode stage, the energy gain was limited to around 30 MeV. In the simulation, the energy E_{peak} at which the peak spectral density occurs is 29.0 MeV, giving further confidence that the modelling of the plasma profile and the wake generated therein is representative of the experiment.

A noteable discrepancy between the injected beams measured in experiments and those modelled in simulations lies in their transverse properties. Considering only the x direction (i.e. the only direction in which transverse beam measurements could be made with the spectrometers), the 0.80 mrad rms divergence at the end of the simulated plasma cathode stage is at a similar level to the (0.7 ± 0.1) mrad measured for this working point in experiment [see first data point in Figure 5.32 (a)]. The size of the virtual waist at the plasma exit is however smaller in simulations [$\sigma_{x,virtual,0} = 3.5\,\mu\mathrm{m}$ in Figure 5.42] than at any working point measured in experiment; one of the smallest waist size estimates was provided by the minimum energy slice beam size of the shot displayed in Figure 5.35 (d), at a value of $\sigma_{x,slice,0} = (22 \pm 1)\,\mu\mathrm{m}$. Consequently, the emittance of the simulated injected bunch $\epsilon_{x,n} = 0.14\,\mathrm{mm\,mrad}$ is much smaller than any measured in experiment. This may indicate further limits to the resolution of the experimental measurements that have not been considered, but is likely also a result of emittance growth sources that are not captured by the current model.

Some sources of in-plasma emittance growth have already been investigated in simulation—namely those due to betatron decoherence and neutral gas scattering. Albeit studied in a rather idealised model, these effects alone do not appear to explain the larger emittances measured in the experiment. Further sources of emittance growth, whether real or apparent, may however also arise in the transport and measurement of the injected beams on the spectrometer screen—something studied in Section 5.3.3.

For now, brightness calculations can be performed for the simulated beam as it exits the plasma. As defined by Equation 5.2, the 3D brightness of the injected beams in these simulations is $41\,\mathrm{pC/MeV/mm \cdot mrad}$; an amplification of ~ 34 compared to the $1.2\,\mathrm{pC/MeV/mm \cdot mrad}$ 3D brightness of the simulated driver. However, with full access to the full 6D phase space in simulations, calculations of the 6D brightness can also be performed. Using the 0.19 kA peak current of the simulated witness bunch and including also its emittance in y, the 6D brightness of the injected bunch from Equation 2.31 is $1.2\,\mathrm{kA\,MeV^{-1}\,mm^{-2}\,mrad^{-2}}$. Despite the reduction in current, this represents a brightness boost by a factor of ~ 1300 compared to the simulated driver and its initial 6D brightness of $9.5 \times 10^{-4}\,\mathrm{kA\,MeV^{-1}\,mm^{-2}\,mrad^{-2}}$.

5.3.2 Injection downramp evolution

As noted, the working point modelled up to this point is one in which the driver samples a plasma immediately after the latter is generated by the lasers i.e. $t_{driver} - t_{laser} = 0$. However, it was found in experiment that allowing the plasma to evolve on the tens of nanoseconds scale yielded both an increase (initially) in the measured injected bunch charge (see Figure 5.27) and a decrease in its divergence [see Figure 5.32 (a)]—implying therefore an increase in the resulting brightness.

To understand how the plasmas might evolve on this time scale, consider that a plasma column generated by a gaussian laser pulse will be hottest along its axis; from longituidnal and transverse laser ionisation simulations in Figures 5.11 and 5.12, the mean axial temperatures were found to be $\sim 2\,\mathrm{eV}$ and $\sim 10\,\mathrm{eV}$ respectively. This results in a pressure gradient that leads to radial expansion. At least initially, this will lead to a drop in the peak density of the column. Eventually, the supersonic flow of the plasma will cause a Sedov-Taylor-like shock front to form [54], but will be unsupported and its amplitude will decay as it propagates.

Exact modelling of this effect would require magneto-hydrodynamics simulations of both the longitudinal and transverse plasma channels with precise knowledge of the initial conditions of the system—something currently out of reach (and out of scope) in this work. Nevertheless, progress can be made by constructing a simple, idealised model of the injection process and considering how hydrodynamic evolution of the plasma would change the simulation parameters and therefore the injected bunches.

Such a model was created using FBPIC, concentrating in this case purely on the evolution of the injection downramp and its ramifications for the injected bunches. Here, the simulation box had dimensions in (r, z) of $(106\,\mu\mathrm{m} \times 354\,\mu\mathrm{m})$, with corresponding cell sizes in each direction of $(0.4\,\mu\mathrm{m} \times 0.25\,\mu\mathrm{m})$ and 2 azimuthal modes. The gaussian driver initialised in this simulation is the same as in the WarpX case (see Section 5.3.1 on page 158 for parameters) in all respects apart from the current; although a peak current of 1.9 kA was again used, the driver was modelled here with a gaussian temporal profile. A total of 3×10^5 fixed-weight macroparticles were used to model the driver electrons.

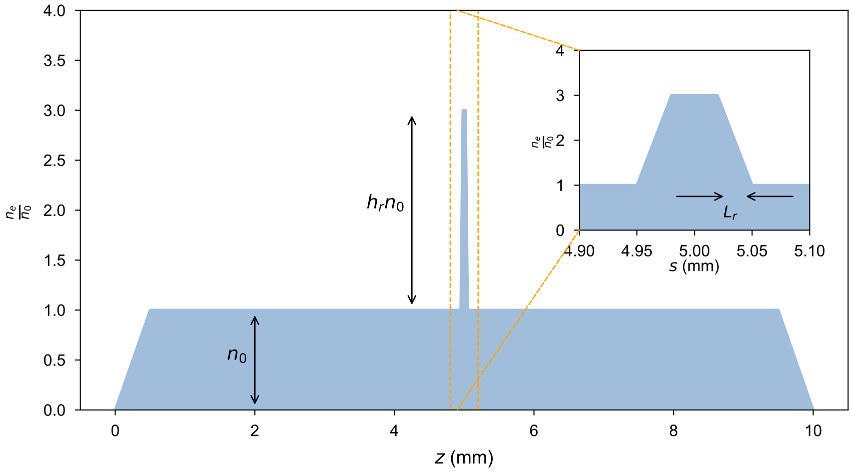

Figure 5.46: Idealised injection downramp simulation plasma profile. The main plasma channel is modelled as a trapezoid varying only in the longitudinal direction with a flattop density of n_0. An additional trapezoid in the centre of this profile with a peak density $h_r n_0$ above the base provides an idealized injection downramp (shown in more detail in inset), with ramp height h_r and length L_r as parameters that can be varied.

This driver comes to focus at the end of the entrance upramp of the plasma profile presented in Figure 5.46. In this case, the accelerating channel was modelled by a trapezoid function with a 9 mm-long flattop and 0.5 mm-long linear entrance and exit ramps. The maximum density of this channel n_0 is $2 \times 10^{16}\,\mathrm{cm}^{-3}$—approximately the density of the main plasma channel in the injection region (see Figure 5.13). At the centre of this profile is a density spike modelled by a trapezoid of additional peak density $h_r n_0$ and ramp length L_r, where h_r and L_r are parameters to be varied. The flattop length of this density spike is fixed at 40 µm; changing this parameter to lengths of 25 µm , 10 µm and 2.5 µm was however not shown to have a large effect over the parameter space considered here. For simplicity, and for cylindrical symmetry in FBPIC, the ramp was considered to only vary longitudinally; rather than the plasma column with finite y-extent that was previously modelled, the transverse plasma is instead here a density slab, extending infinitely in both x and y. For the majority of the plasma $(2 \times 2 \times 8)$ macroparticles per cell were used in (r, z, θ), with the exception of a 1 mm-wide region around the injection spike where there were $(4 \times 4 \times 8)$. Again, as in the Warp-X case, the plasma was initialised with zero temperature.

Using this model as a basis, a 2D scan of L_r and h_r was performed, with the results presented in Figure 5.47. Three key parameters were extracted from the witness bunch straight after injection: the (a) charge Q (and peak spectral density $\left(\frac{dQ}{dE}\right)_{peak}$), (b) horizontal emittance $\epsilon_{x,n}$ and (c) the corresponding 3D brightness defined by Equation 5.2.

A downramp with $h_r = 60$ and $L_r = 30\,\mu\text{m}$ best approximates that modelled for the $t_{driver} - t_{laser} = 0$ case in Figure 5.13—a working point labelled here as (I). A charge of $Q_I = 22\,\text{pC}$ is injected for these parameters [see Figure 5.47 (a)], whereas the optimum in the parameter space is in fact provided by downramps of lower peak density ($h_r = 8$–12) and longer ramp length ($L_r = 45\,\mu\text{m}$). As shown by Figure 5.47 (b), $\epsilon_{x,n}$ varies only weakly between ~ 0.4–$0.6\,\text{mm mrad}$. The resulting optimum in $B_{3\text{D}}$ [see Figure 5.47 (c)] therefore occurs at a similar place to the optimal charge, labelled (II) with $h_r = 12$ and $L_r = 45\,\mu\text{m}$ where $Q_{II} = 45\,\text{pC}$.

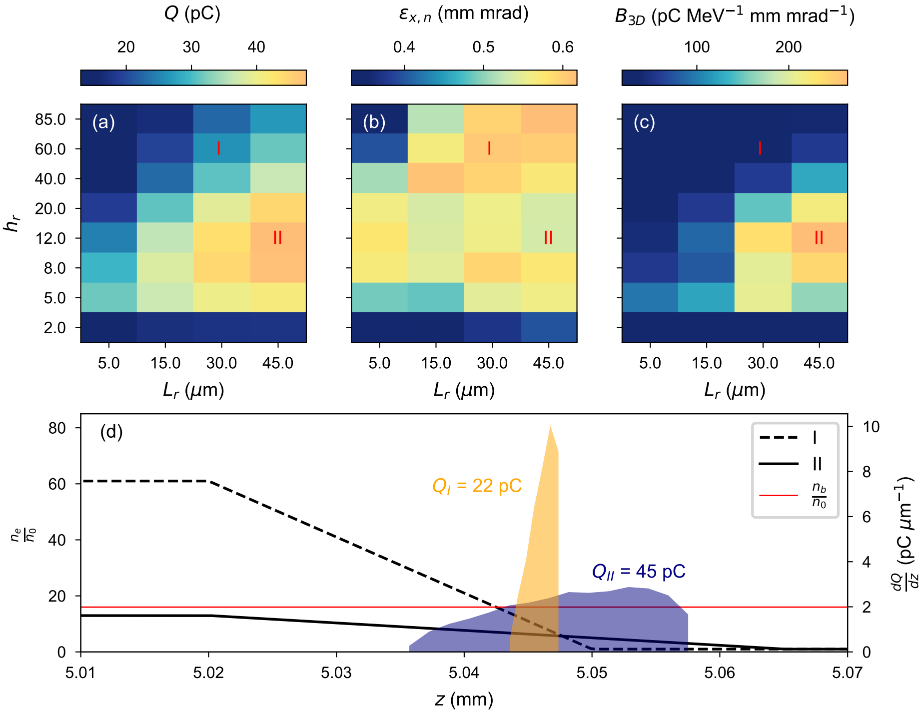

Figure 5.47: Injected bunch dependence on downramp height and length in idealised injection downramp simulations. The injected bunch (a) charge Q, (b) horizontal emittance $\epsilon_{x,n}$ and (c) 3D brightness B_{3D} are plotted as a function of injection downramp height h_r and length L_r. Two cases are highlighted: (I) with parameters most resembling the downramp modelled in Figure 5.13 and (II) with parameters that provide the highest value of B_{3D} within the parameter space. The downramp profiles for (I) (black dashed line) and (II) (black solid line) are plotted in (d), along with the initial positions of the injected electrons before the arrival of the driver given as line charge densities $\frac{dQ}{dz}$ (orange and blue filled curves).

At first glance, this observation runs counter to the downramp dynamics predicted by Equation 2.64, where a steeper downramp leads to a greater reduction in the wake phase velocity, allowing more charge to be trapped. This can however be explained by looking in detail at the initial positions of the injected electrons and the relative densities of the driver and plasma in the injection region [see Figure 5.47 (d)]. In the case of (I), the majority of the injection spike profile (black dashed line) has a higher density than that of the drive beam (red line) at the injection point ($n_d = 3.2 \times 10^{17}\,\mathrm{cm}^{-3} < h_r n_0$); where this is true, the driver is not able to generate a strong blowout. Only at the end of the downramp is the drive beam density comparable to the plasma density, where once again a blowout forms and injection can occur. Electrons are injected in this case at a high rate, but only over a short distance confined to the end of the downramp, as shown by the initial longitudinal locations of the witness electrons before the arrival of the driver [filled orange curve in Figure 5.47 (d)]. For case (II) on the other hand, the density in the downramp (solid black line) is similar to that of the driver, meaning a wake is driven throughout and injection occurs over a longer region. Although the rate of injection is less overall, this—along with the fact that the downramp length is slightly longer—leads to more injected charge. Case (I) then represents a rather unusual parameter regime; previous studies have confined themselves to cases where $h_r = 1$–2 [56, 216] and where $n_d \gg n_e$ throughout. This also applies to the short ramp lengths in this case; in other studies where $L_r > k_p^{-1}$ (k_p^{-1} being the skin depth of the base plasma density), the injected charge is shown to decrease with ramp length [116].

To explain how allowing the transverse plasma to expand moves the system from point (I) to different regions in the parameter space, consider the previous discussion of how a hot column of laser-generated plasma would evolve. At the beginning of the expansion, the column height would decrease, increasing the amount of injected charge [down in Figure 5.47 (a)]. However, once a shock forms L_r will decrease, reducing the injected charge along with it [left in Figure 5.47 (a)]. Eventually, this shockwave will dissipate and injection will consequently cease.

Note that, although this serves as an explanation for the evolution in charge in experiment when allowing the plasma to evolve, it does not explain the decrease in divergence. This then could perhaps be explained by evolution of the longitudinal plasma, which has not yet been considered. Although a lower temperature than the transverse plasma column and thus expected to evolve less quickly, expansion of this channel will inevitably have an effect on the shape of the exit ramps of the longitudinal plasma and subsequently the divergence of the injected bunches leaving the plasma cathode stage. Along with more accurate modelling of the transverse plasma (e.g. with full MHD simulations), studies of the longitudinal plasma evolution would therefore be required to fully understand the experimental results.

5.3.3 Plasma-to-spectrometer tracking simulations

In order to understand how the injected bunch parameters at the end of the plasma stage may translate into measured parameters on the spectrometer diagnostics, beam tracking simulations were performed using OCELOT [147]. For simplicity, a transversely symmetric beam is assumed that is gaussian in all directions. Based on an average injected beam at the optimised experimental working point in Figure 5.29, the simulated bunch is given 20 pC of charge, a central energy of 30 MeV and a 1.5 MeV rms energy spread. Additionally, the divergence of this beam is set as 0.3 mrad from the measurement in Figure 5.31. Modelled as diverging from a point in vacuum, a waist size of 20 µm is used—similar to the smallest waist sizes measured in Figure 5.34 (a) and that of the single beam in Figure 5.35 (d)—implying an emittance of $\epsilon_{x,n} = \epsilon_{y,n} = 0.36\,\text{mm}\,\text{mrad}$. Finally, the rms bunch length—of which no experimental measurement is available—is estimated from simulations as 10 µm, giving a peak current of 0.24 kA.

Figure 5.48 summarises the simulated optics of the ESPEC and LEMS beamlines during an object plane scan about the plasma exit. Based on the beamlines described in Figure 3.6, the transport from the plasma exit to the ESPEC and LEMS [see Figure 5.48 (a) and (b) respectively] is modelled.

The dipoles are in both cases dispersive only in the y-direction [see red lines in Figure 5.48 (c) and (d)], increasing the vertical beam size of the beam on the screen according to the energy spread. An object plane scan is then performed by altering the imaging quadrupole strengths in the same way as in experiment [see Section 2.2.2]. Example beta functions in x (blue lines) and y (red lines) for different object planes are displayed in Figure 5.48 (e) and (f) for the ESPEC and LEMS respectively. Accounting for the magnification for each quadrupole setting, the beam size in x at the object plane $s_{obj,x}$ can be calculated from the rms beam sizes on the virtual screen.

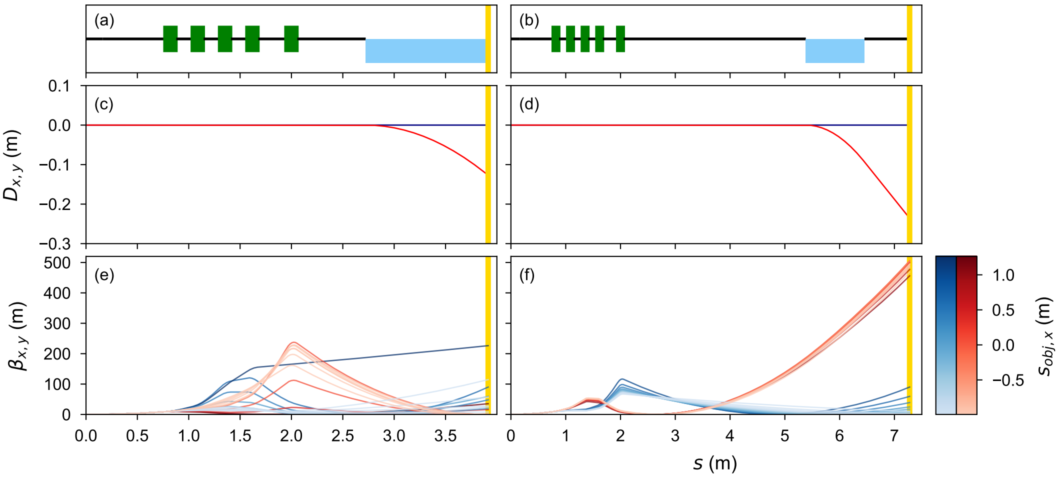

Figure 5.48: Beta and dispersion function evolution during object plane scans in plasma-to-spectrometer tracking simulations. The ESPEC (a) and LEMS (b) beamlines are modelled by a series of drifts (black lines), quadrupole magnets (green boxes) and a bending dipole magnet (blue boxes) before reaching the virtual diagnostic screen (yellow line). In both cases the dispersion (blue and red lines for x and y respectively) of the (c) ESPEC and (d) LEMS dipoles is in y-only and remains constant throughout measurements. Changing the quadrupole strengths alters the beta functions in x and y (blue and red lines respectively), changing the object plane $s_{obj,x}$ imaged on the (e) ESPEC and (f) LEMS diagnostics.

Note that whereas in the ESPEC case the beta function in y at the screen location is consistently small throughout the scan ($\sim$ 10 m maximum), the same is not true on the LEMS. Here, the β_y at the screen was as large as $\sim$ 500 m. Whilst allowing larger beta functions in y at the diagnostic screen can in some cases make finding an imaging solution for the desired object plane and magnification in x easier, this in principle will provide worse energy resolution in the dispersive plane. Nevertheless, it can be shown that the dispersive contribution to the vertical beam size $D_y\sigma_\delta$, where σ_δ is the relative energy spread, is much larger than that from the optics, $\frac{\beta_y\epsilon_{y,n}}{\gamma}$. With $D_y = 0.2$ m and $\beta_y = 500$ m at the LEMS screen, typical parameters $\gamma \sim 60$ and $\sigma_\delta \sim 1\,\%$ lead to a vertical beam size $\sigma_y = \sqrt{(D_y\sigma_\delta)^2 + (\frac{\beta_y\epsilon_{y,n}}{\gamma})^2} = 2.0017$ mm, where a pessimistic value of the vertical normalised emittance $\epsilon_{y,n} = 10$ mm mrad was used. This is then only around 0.085 % larger than the vertical beam size with only the dispersive contribution present—a negligible contribution.

An object plane scan of the previously described, idealised beam about its waist is first simulated for the ESPEC case; the results of this are presented in Figure 5.49. Three different simulations are performed considering increasingly higher order effects, represented by the three different sets of data (different colours) in Figure 5.49

(a). For each case, two measures of the beam size are considered. First of all, the rms projected beam size in x is calculated for each imaging setting; the resulting data is plotted and Equation 2.38 fitted (points and solid lines), allowing the waist size $\sigma_{x,0}$ and divergence $\sigma_{x'}$ to be extracted for each case (see legend at bottom of Figure 5.49). Secondly, the size of the central energy slice for each imaging setting is also extracted [i.e. the width of a slice taken along the black dashed lines in the example virtual spectrometer images in Figure 5.49 (b–g)]. The width of this slice at the object plane scan waist $\sigma_{x,slice,0}$ is plotted as horizontal dashed lines in Figure 5.49 (a), and is also displayed in the legend.

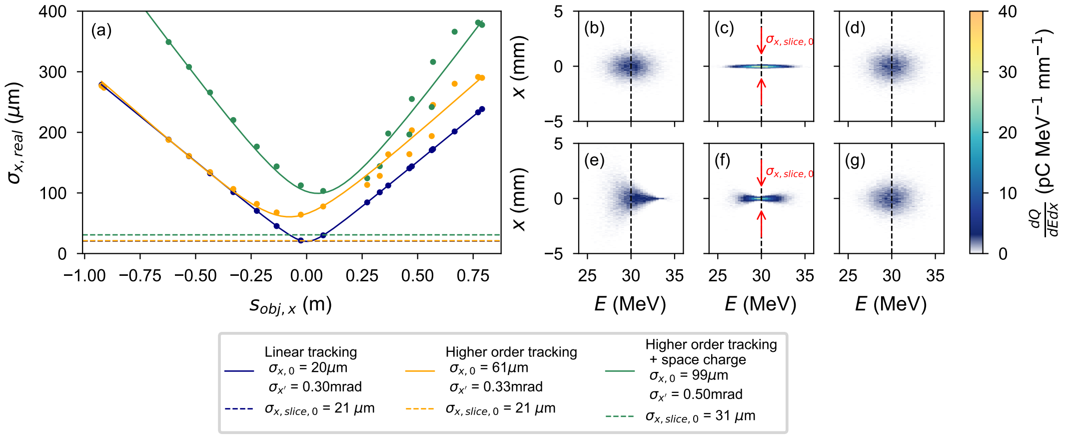

Figure 5.49: Simulated object plane scan of injected bunches on the ESPEC. (a) Virtual measurements of the horizontal beam size σ_x as a function of object plane position $s_{obj,x}$ are displayed for three cases: first order tracking (navy), higher order tracking (orange) and higher order tracking with space charge (green); the relevant fit parameters extracted in each case are displayed in the legend below. In addition, virtual spectrometer images are displayed showing the beam in and out of focus for (b–d) the first order tracking case and (e–g) the higher order tracking case. Black vertical dashes lines here highlight the central energy from which $\sigma_{x,slice}$ is extracted, which is equal to $\sigma_{x,slice,0}$ at a waist.

In the first simulated object plane scan, the tracking is performed entirely with first order transfer maps, with the projected beam sizes perfectly reconstructing the waist of the beam at the plasma exit (navy points and fit). This is due to the fact that first order transfer maps in OCELOT consider all of the particles to be at the design energy when focusing in the quadrupoles, such that no effect of chromaticity is visible on the spectrometer screen and the energy slices all reach a waist at the same imaging setting [see Figure 5.49 (b–d)], such that the energy slice waist size is approximately equal to the projected waist size ($\sigma_{x,0} \approx \sigma_{x,slice,0}$).

This situation changes however in the second case, where high order tracking—specifically treating the drifts with second order matrices and the quadrupoles with a Runge-Kutta integrator—means that particles of different energy will be imaged differently on the spectrometer screen. The chromatic effect that this induces is demonstrated by Figure 5.49 (e–g), where around the beam waist a 'butterfly' shape appears because not all of the energy slices can be focused simultaneously. Naturally, as noted in Section 5.2.5, this chromaticity leads to an overestimation of the beam waist when using the projected beam size, as shown by the orange data points and fit in Figure 5.49 (a) where $\sigma_{x,0} = 61\,\mu\text{m}$. Considering only a single energy slice however removes this chromatic effect, with $\sigma_{x,slice,0}$ the same for both this case and the first order tracking case. Note that the divergence measured by looking at the projected beam width slightly increases (by around 10 %) compared to the true divergence at the waist—indicating that the divergence measured in e.g. Figure 5.31 is at worst slightly overestimated due to this effect.

In addition to higher order tracking, preliminary studies have taken place into the effect of space charge on the beam transport, which can have a non-negligible effect on the apparent emittance of low energy beams measured via the object plane scan method [233]. As evidenced by the green data points and fit line in Figure 5.49, space charge can increase the apparent divergence ($\sigma_{x'} = 0.50\,\text{mrad}$) and waist size ($\sigma_{x,0} = 99\,\mu\text{m}$, $\sigma_{x,slice,0} = 31\,\mu\text{m}$) of the beam, with the waist appearing further downstream than the true waist location. This effect however is dependent on a number of different beam properties, for example: the beam energy, the initial beam size and divergence, the current (which is only approximated here from simulations) and the amount of dispersion the beam experiences in the dipole (i.e. an energy spread dependence). More accurate knowledge of the input parameters would be required to conclude for sure whether this effect is present in experiment; in the case that it does, it provides another source of overestimation of the apparent measured emittance.

Albeit subject to some potential over-estimations, the divergence $\sigma_{x'}$ can therefore be measured from the fit to the projected beam size whereas a more accurate value of the waist size is provided by the minimum slice energy size $\sigma_{x,slice,0}$. Note however that none of the resolution issues on the ESPEC—namely due to the low resolution of the camera chip paired with scattering in the vacuum chamber wall—are modelled in this simulation, the presence of which results in a blurred, round beam [see example in Figure 5.29 (d)] on the screen from which $\sigma_{x,slice,0}$ cannot be identified. This therefore motivates the measurements of $\sigma_{x,slice,0}$ with the high-resolution LEMS diagnostic, as was presented in Section 5.2.5.

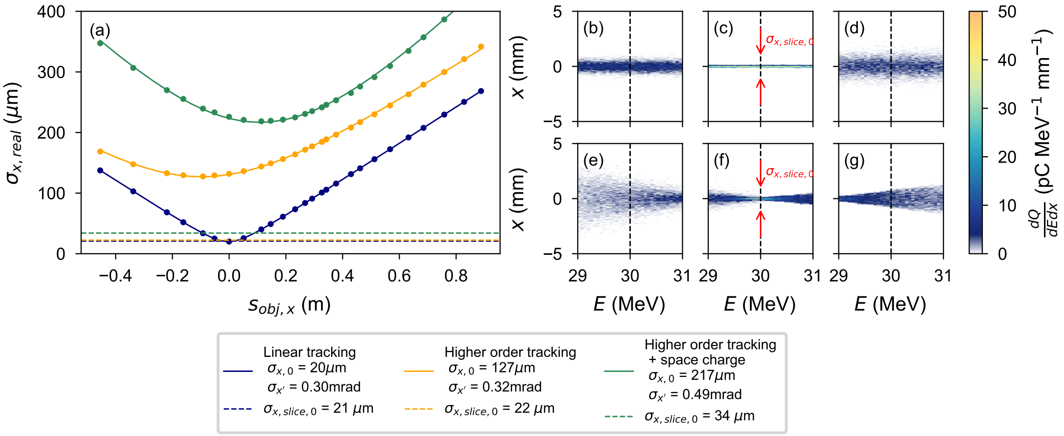

Figure 5.50: Simulated object plane scan of injected bunches on the LEMS. Following the same format as for the ESPEC case (see Figure 5.49 for more details), the object plane scans for the three different tracking cases are shown in (a). Similarly, virtual spectrometer images are shown for the first order tracking (b–d) and higher order tracking (e–g) cases; note that a smaller energy range is plotted to reflect the narrow bandwidth ($\sim 2\,\text{MeV}$) of the LEMS diagnostic.

Modelling of the LEMS object plane scans was performed in the same way as described for the ESPEC, except now transporting the beams further to the LEMS diagnostic; the results are presented in a similar format in Figure 5.50. With the results displayed in Figure 5.50 (a), the object plane scan simulations are performed for the same three cases: first order tracking (navy data), higher order tracking (orange data) and higher order tracking with space charge (green data). Again, in the first order case, there is no effect of chromaticity [see virtual spectrometer images in Figure 5.50 (b–d)] and so using the projected beam sizes perfectly reconstructs the true waist of the beam. In a similar fashion to the ESPEC case, the modelling of higher order chromatic effects [see Figure 5.50 (e–g)] overestimates the waist size when using the projected profile, where $\sigma_{x,0} = 127\,\mu\text{m}$ [c.f. the mean projected beam size measured on the LEMS in Figure 5.34, where $\langle\sigma_{x,proj}\rangle = (143 \pm 8)\,\mu\text{m}$]. This effect can be negated by again measuring the minimum energy slice width $\sigma_{x,slice,0}$—which the high resolution LEMS diagnostic is able to achieve in reality. As the divergence can be provided more easily by the ESPEC measurements, the estimation of the waist size from $\sigma_{x,slice,0}$ is the main goal here. Finally, space charge can potentially lead to an overestimation of this value; $\sigma_{x,slice,0} = 34\,\mu\text{m}$ in this case, which is comparable to some of the measurements displayed in Figure 5.34.

Just as in Section 5.2.5, we can combine the divergence $\sigma_{x'}$ obtained from the ESPEC with the minimum slice energy width $\sigma_{x,slice,0}$ from the LEMS in Equation 5.1. In both cases, the effects of space charge would appear to lead to overestimates of the values, with $\sigma_{x,slice,0} = 34\,\mu\text{m}$ from the LEMS and $\sigma_{x'} = 0.50\,\text{mrad}$ from the ESPEC—consequently leading to an overestimation of the emittance as $\epsilon_{x,n} \approx 1\,\text{mm}\,\text{mrad}$.

5.4 Outlook

The results reported here constitute another important step in beam-driven plasma cathode development. As well as acting as a proof-of-principle that the brightness of a drive beam can be boosted (at least in three dimensions) by internal injection of a low emittance bunch, there are many lessons that can be drawn from this experiment, with many different avenues of research to explore in the future.

One of the main drawbacks in this iteration of a plasma cathode stage was the rather low energy ($\sim 30\,\text{MeV}$) of the injected bunches—understood now to be a consequence of the plasma generation method. In previous experiments, the preionisation laser was found to produce an accelerating channel that was too narrow, resulting in broken wakefields and therefore unstable transport of injected beams [215]. Here, the use of a long-rayleigh-length, large-focal-spot preionisation laser afforded the generation of a wide accelerating channel, but unfortunately aggravated the issue of ionisation defocusing, resulting in a diminished average accelerating field and an inhomogeneous plasma that may have also contributed to the unstable transport of injected bunches—an effect now better understood via simulations. One possible solution to this issue is an improvement of the longitudinal laser focusing method (e.g. some variant of the axicon optics in Ref. [215] but with the avoidance of plasma channels that are again too narrow). An attractive alternative has however been briefly studied here—using a discharge to form the accelerating channel, whilst an independent laser triggers the injection. As the example beam in Figure 5.37 proves, a discharge alone can be used to produce the longitudinal plasma whilst still allowing a density spike to be generated by the transverse laser, injecting bunches that are then accelerated to higher energies ($\sim 90\,\text{MeV}$) than the laser-only case. This setup was far from optimised and, as noted, the exact experimental conditions were unclear; it therefore remains an open question as to whether this method can produce high-charge bunches with enough stability to properly measure their qualities. If this proves to be the case however, not only would this method overcome the energy gain limitations of the laser-generated plasma here, but it could provide a pathway to internal injection at the high-repetition rates compatible with discharge systems and low energy lasers—an area of significant interest at FLASHForward and indeed the entire plasma wakefield acceleration community.

Regardless of exactly how it is achieved, an increase in the energy of the injected bunches in this experiment would greatly benefit their diagnosis; for example, any uncertainties about the effects of e.g. space charge would quickly be eliminated. This would then provide a more suitable platform to achieve the next experimental milestones in this research: the diagnosis of the longitudinal and ultimately six-dimensional phase spaces of the injected beams. Such an endeavour would inevitably require transport of the bunches to the TDS further downstream and would therefore naturally benefit from increased injected charge (such that the beams can be detected on e.g. the BPMs) and of course increased stability. Section 5.2.3 illustrated several methods for optimising both of these aspects of the injected beams. Further improvements could be made through the use of e.g. an astigmatic focus to remove the dependence on driver—downramp alignment [118] and also the reintroduction of multiple gas species, affording more tunability of the downramp height and length [56].

In parallel to the improved diagnosis of injected bunches in the experiment, more thought can also be given to increasing their 6D brightness. Whilst more charge should, if the injected bunch length is constant (bunch length $\lesssim$ the extent of the wake's accelerating phase at injection), yield higher currents, improvements in the transverse and spectral directions should also be considered.

For the transverse planes, this would involve limiting the amount of in- and out-of-plasma emittance growth, but would also no doubt benefit from an improvement of the internal injection method. Still considering downramp injection, it has been suggested that using gentler downramps can in fact generate high-current bunches that are injected with low transverse momentum, resulting in a corresponding decrease in the emittance [117]; this however may require higher driver currents to ensure that reasonable amounts of charge are still trapped. Although ionisation injection schemes could also be considered, the peak currents that are required to trap from rest may remain prohibitively high for a FLASHForward-like machine (although schemes have been proposed that may relax this restriction e.g. Ref. [234]). Furthermore, these methods are yet to match the stability demonstrated with the downramp injection method demonstrated here, which is necessary for the multi-shot characterisation of the injected beams. For those methods relying on the precise synchronisation of the injector to the wake (e.g. a Trojan Horse plasma photocathode), upgrades to e.g. the laser—driver synchronisation would also be required [176].

Large energy spreads remain one of the inherent drawbacks of plasma accelerators and make it difficult to compete with the relatively monoenergetic beams produced by conventional machines like FLASH. Removal of the correlated energy spread of the injected bunches is therefore a key step in truly boosting the 6D brightness of the FLASH linac (i.e. before the FLASH bunches are subjected to CSR effects in strong compression). Ideally, this could be accomplished via proper shaping of the bunches to achieve optimal beam loading [127] via control over their injection positions [216, 235]. Alternatively however, plasma-based dechirping methods may also provide a solution to the large energy spread of the injected beams; for schemes that currently exist, these would require either multiple plasma stages [200, 236] or the injection of multiple bunches [237]. Without the energy spread reduction afforded by these methods, using the injected bunches to achieve FEL gain—seen as a key application of this technology—would be typically seen as impossible. A notable exception to this however is a novel scheme that in fact utilises the large time-energy correlation and small slice energy spread of internally injected bunches from a PWFA (see e.g. those in Figure 5.45) to its advantage. By compressing the strongly chirped bunch with a chicane consisting of only small bend angles, the Plasma-driven Attosecond X-ray (PAX) scheme allows the generation of ultra-short bunches with a peak current that is in principle limited only by the bunch's uncorrelated energy spread [238, 239]. Driving an x-ray FEL light source with such pre-bunched electron beams also relaxes the restrictive tolerances on e.g. the emittance, energy spread and pointing stability which have previously made plasma accelerator-driven XFELs difficult to realise, making PAX a yet more attractive scheme.

For any of the above endeavours, simulations will continue to compliment experiments, either by providing understanding of a working plasma stage (as was done here) or by testing the feasibility of any of the suggested schemes before they are implemented in reality. Ultimately, start-to-end simulations will be necessary. Specifically, this means starting with a driver that is known to be obtainable in the machine, propagating this through a simulated plasma profile motivated by experimental measurements that has the injection method of choice implemented, before ending with simulations of the diagnostics to ensure that the beams that are generated can be properly measured. Naturally, this can also be extended to applications of the injected bunches (e.g. FEL simulations).

Looking briefly outside of the context of this project, it is clear that PWFA-based plasma cathode research will also continue to rapidly develop at other experimental facilities. For instance, FACET II aims to utilize its extremely high driver currents to further test some of the most demanding internal injection schemes [240, 241]. On the other hand, the use of short, high charge LWFA-generated beams [242, 243] as drivers for PWFA stages in so-called hybrid plasma accelerators unlocks the possibility of investigating beam-driven plasma acceleration with high drive currents at small scale facilities [244–246]—something that is already empowering internal injection studies [118, 119]. In addition to this, such schemes also profit from the inherent synchronisation of the PWFA drive beam to any lasers (e.g. those used as probes or injectors) that are derived from the LWFA driver. Whether the LWFA-produced beams have the qualities and tunability needed to drive a truly stable plasma cathode stage remains to be seen however. For instance, whilst the high energy spread of drivers from an LWFA stage can make them resistant to instabilities such as hosing [247], the focusing of these beams into the PWFA stage—shown in Section 5.2.3 on page 137 to be an important control over the injected charge—may be a challenge [248]. Moreover, hybrid plasma accelerators still inherit the currently low wall-plug-efficiency and repetition-rate of their LWFA stages which, until addressed, may limit their future applicability. Facilities based on conventional machines therefore remain a key part of PWFA-based internal injection research.

5.5 Summary

Building upon a previous iteration of the X-1 plasma cathode experiment at FLASHForward [197], the work presented in this chapter has studied in greater detail the experimental dependencies of downramp-injected bunches. Armed with this knowledge, a stable working point injecting on-average high spectral density beams was established. From this, robust multi-shot measurements of the injected beam's transverse properties were performed. Ultimately, this led to the demonstration that the 3D brightness of a drive beam—measurable here only in the spectral and horizontal dimensions—can be boosted via the injection of a low-emittance witness bunch.

Throughout these studies, simulations of the real-life experiment—from the generation of the plasma, to the injection and acceleration of witness bunches, to their detection on the spectrometers—have substantially increased the understanding of the plasma cathode stage. Equipped with the tools and experience that experiments and simulations have afforded, a multitude of different improvements to the plasma cathode stage have been suggested for the future. For instance, a change of the preionisation method from laser to discharge—a technique that has been shown here to work in principle—would increase the energy of the injected bunches whilst also providing a potential pathway to a high-repetition-rate plasma cathode stage. Combined with further optimisations of the stability and charge, this would then allow for even more rigorous characterisation of the generated beams, with the eventual goal being the measurement of their full 6D phase space. Improvements to the transverse properties of the injected beams may also be possible through the deployment of more advanced injection techniques. Ultimately, this could then enable the overarching goals of the plasma cathode experiment to be realised—the demonstration of a full 6D brightness boost of the electron beams in the FLASH linac that could then be used to drive e.g. an FEL light source.

Chapter 6

Conclusion

In summary, two features of a beam-driven plasma accelerator have been studied throughout the course of this project. In Chapter 4, a new method for measuring the energy transfer efficiency of a PWFA-based energy transformer stage was developed, based on the light emitted from a plasma after interaction with the beam. In Chapter 5, a PWFA-based plasma cathode stage was established and optimised to produce new electron bunches which, after characterisation, demonstrated a brightness transformation of the incoming drive beam. Outlooks specific to each each of these sub-topics have been given at the end of their respective chapters in Sections 4.3 and 5.4.

To conclude, we can briefly reflect on the overlaps between the two main concepts of this work and also hypothesise as to how they may become more connected in the future. An obvious link is the use of the emitted light from a plasma afterglow as a diagnostic for plasma accelerators. Whereas this forms the basis of the longitudinally-resolved efficiency measurements in Chapter 4, this signal had been previously exploited in the alignment and synchronisation of laser-generated plasmas and electron beams, and indeed proved useful again in this way in Chapter 5—both as a tool for setup and online monitoring. In general, this therefore illustrates the potential of this diagnostic as a simple, versatile tool for the real-life operation of plasma accelerators.

Looking even further ahead, it is widely perceived that future plasma accelerators—especially those intended for high-energy physics applications—will require multiple stages. Such multi-stage accelerators pose a supreme technical challenge to the plasma wakefield accelerator community. As noted in Section 4.3, the non-destructive, longitudinally-resolved plasma light-based diagnostic may go some way to understanding and monitoring the instabilities in such a machine. Even if such research progresses using primarily externally injected trailing bunches, the majority of the lessons learnt will also apply to the case where the accelerated beams are provided via internal injection. Indeed, by using a robust plasma cathode stage as the first step in a multi-stage plasma accelerator, internal injection could provide bunches with qualities that surpass those provided by conventional technologies in a much more compact configuration than, say, a damping ring. The accelerator providing the driver bunches—whether it be a conventional linac or an LWFA in a hybrid configuration—can then focus on optimising the drive beam such that the maximum energy is transferred to the trailing bunch. However, if such a plasma cathode stage is to become the de-facto injector for such a machine, it must be proven to reliably produce the required electron bunches with also the necessary tunability. Research therefore needs to continue into increasing the stability of both energy and brightness transformer PWFA stages; with the eventual goal that they can be brought together to meet the demands of future FEL light sources and eventually linear colliders.

Bibliography

[1] G. Ising. "Prinzip einer Methode zur Herstellung von Kanalstrahlen hoher Voltzahl". In: *Ark. Mat. Astron. Fys.* 18.30 (1924), pp. 1–4. URL: `https://cds.cern.ch/record/433984`.

[2] R. Wideröe. "Über ein neues Prinzip zur Herstellung hoher Spannungen". In: *Archiv für Elektrotechnik* 21.4 (July 1928), pp. 387–406. ISSN: 1432-0487. DOI: `10.1007/BF01656341`.

[3] J. D. Cockcroft, E. T. S. Walton, and E. Rutherford. "Experiments with high velocity positive ions.—(I) Further developments in the method of obtaining high velocity positive ions". In: *Proceedings of the Royal Society of London. Series A, Containing Papers of a Mathematical and Physical Character* 136.830 (1932), pp. 619–630. DOI: `10.1098/rspa.1932.0107`. eprint: `https://royalsocietypublishing.org/doi/pdf/10.1098/rspa.1932.0107`.

[4] J. D. Cockcroft, E. T. S. Walton, and E. Rutherford. "Experiments with high velocity positive ions. II. -The disintegration of elements by high velocity protons". In: *Proceedings of the Royal Society of London. Series A, Containing Papers of a Mathematical and Physical Character* 137.831 (1932), pp. 229–242. DOI: `10.1098/rspa.1932.0133`. eprint: `https://royalsocietypublishing.org/doi/pdf/10.1098/rspa.1932.0133`.

[5] O. Barbalat. *Applications of particle accelerators; rev. version*. Tech. rep. 1994. DOI: `10.5170/CERN-1994-001.841`.

[6] G. Cuttone. *Applications of Particle Accelerators in Medical Physics*. Tech. rep. Geneva: CERN, 2008. URL: `https://cds.cern.ch/record/1514553`.

[7] R. W. Hamm and M. E. Hamm. *Industrial Accelerators and Their Applications*. World Scientific, 2012. DOI: `10.1142/7745`. eprint: `https://www.worldscientific.com/doi/pdf/10.1142/7745`.

[8] A. Sessler and E. Wilson. *Engines of Discovery*. Revised and Expanded. World Scientific, 2014. DOI: `10.1142/8552`. eprint: `https://www.worldscientific.com/doi/pdf/10.1142/8552`.

[9] P. Petitjean. "Pierre Auger and the Founding of CERN". In: *Sixty Years of Sciences at Unesco, 1945-2005*. Ed. by Petitjean et al. Unesco, 2006, pp. 57–60. URL: `https://shs.hal.science/halshs-00166533`.

[10] T. J. Berners-Lee and R. Cailliau. *WorldWideWeb: proposal for a HyperText Project*. Tech. rep. Geneva: CERN, 1990. URL: `http://cds.cern.ch/record/2639699`.

[11] K. Hubner. "Accelerators for high-energy physics". In: *2004 European School of High-Energy Physics*. Jan. 2005, pp. 361–394.

[12] A. W. Chao and W. Chou. *Reviews of Accelerator Science and Technology*. World Scientific, 2011. DOI: `10.1142/8078`. eprint: `https://www.worldscientific.com/doi/pdf/10.1142/8078`.

[13] W. Wuensch. "High Gradient Breakdown in Normal-Conducting RF Cavities". In: *Proc. EPAC'02* (Paris, France). European Particle Accelerator Conference 8. JACoW Publishing, Geneva, Switzerland, Aug. 2007, pp. 134–138. URL: `https://jacow.org/e02/papers/MOYGB003.pdf`.

[14] N. Pichoff. "Introduction to RF linear accelerators; 2000 ed." In: *CAS - CERN Accelerator School : Basic Course on General Accelerator Physics*. CERN. Geneva: CERN, 2005. DOI: `10.5170/CERN-2005-004.53`.

[15] C. Adolphsen et al. *The International Linear Collider Technical Design Report*. Tech. rep. Geneva, 2013. arXiv: `1306.6328`. URL: `https://cds.cern.ch/record/1601969`.

[16] G. Trubnikov. "Introduction to Circular Accelerators—Basic Science and Applied Research". In: *American Institute of Physics Conference Series*. Ed. by C. Granja and C. Leroy. Vol. 1204. American Institute of Physics Conference Series. Jan. 2010, pp. 91–97. DOI: `10.1063/1.3295682`.

[17] D. Tommasini et al. "Status of the 16 T Dipole Development Program for a Future Hadron Collider". In: *IEEE Transactions on Applied Superconductivity* 28.3 (2018), pp. 1–5. DOI: `10.1109/TASC.2017.2780045`.

[18] R. P. Walker. "Synchrotron radiation". In: *CERN Yellow Rep. School Proc.* 1994. DOI: `10.5170/CERN-1994-001.437`.

[19] E. Gschwendtner and P. Muggli. "Plasma wakefield accelerators". In: *Nature Reviews Physics* 1.4 (Apr. 2019), pp. 246–248. ISSN: 2522-5820. DOI: `10.1038/s42254-019-0049-z`.

[20] M. J. Hogan. "Electron and Positron Beam–Driven Plasma Acceleration". In: *Reviews of Accelerator Science and Technology*, pp. 63–83. DOI: `10.1142/9789813209589_0003`. eprint: `https://www.worldscientific.com/doi/pdf/10.1142/9789813209589_0003`.

[21] J. D. Jackson. *Classical electrodynamics; Third edition*. New York: John Wiley & Sons, Inc., 1999, xxi, 808 pages : illustrations, diagrams. ISBN: 9780471309321. URL: `https://bib-pubdb1.desy.de/record/386317`.

[22] H. A. Lorentz. *Versuch einer Theorie der electrischen und optischen Erscheinungen in bewegten Körpern*. Cambridge Library Collection - Physical Sciences. Cambridge University Press, 2013. DOI: `10.1017/CBO9781139381406`.

[23] S. Y. Lee. *Accelerator Physics*. 3rd. World Scientific, 2011. DOI: `10.1142/8335`. eprint: `https://www.worldscientific.com/doi/pdf/10.1142/8335`.

[24] H. Wiedemann. *Particle Accelerator Physics*. SpringerLink: Springer e-Books. Springer Berlin Heidelberg, 2007. ISBN: 9783540490456. URL: `https://books.google.de/books?id=S8CfmLe87RAC`.

[25] E. Courant and H. Snyder. "Theory of the alternating-gradient synchrotron". In: *Annals of Physics* 3.1 (1958), pp. 1–48. ISSN: 0003-4916. DOI: `https://doi.org/10.1016/0003-4916(58)90012-5`.

[26] K. L. Brown. "A First and Second Order Matrix Theory for the Design of Beam Transport Systems and Charged Particle Spectrometers". In: *Adv. Part. Phys.* 1 (1968), pp. 71–134.

[27] K. Floettmann. "Some basic features of the beam emittance". In: *Phys. Rev. ST Accel. Beams* 6 (3 Mar. 2003), p. 034202. DOI: `10.1103/PhysRevSTAB.6.034202`.

[28] S. Di Mitri. "On the Importance of Electron Beam Brightness in High Gain Free Electron Lasers". In: *Photonics* 2 (June 2015), pp. 317–341. DOI: `10.3390/photonics2020317`.

[29] C. Lindstrøm and M. Thévenet. "Emittance preservation in advanced accelerators". In: *Journal of Instrumentation* 17.05 (May 2022), P05016. DOI: `10.1088/1748-0221/17/05/P05016`.

[30] M. Reiser. *Theory and design of charged particle beams; Second, updated and expanded edition*. Weinheim: Wiley-VCH Verlag, 2008, XXVII, 647 pages : illustrations. ISBN: 9783527407415. DOI: `10.1002/9783527622047`.

[31] T. Mehrling et al. "Transverse emittance growth in staged laser-wakefield acceleration". In: *Phys. Rev. ST Accel. Beams* 15 (11 Nov. 2012), p. 111303. DOI: `10.1103/PhysRevSTAB.15.111303`.

[32] M. Thévenet et al. "Emittance growth due to misalignment in multistage laser-plasma accelerators". In: *Phys. Rev. Accel. Beams* 22 (5 May 2019), p. 051302. DOI: `10.1103/PhysRevAccelBeams.22.051302`.

[33] C. Lindstrøm et al. "Transverse Tolerances of a Multi-Stage Plasma Wakefield Accelerator". In: *Proc. of International Particle Accelerator Conference (IPAC'16), Busan, Korea, May 8-13, 2016* (Busan, Korea). International Particle Accelerator Conference 7. Geneva, Switzerland: JACoW, June 2016, pp. 2561–2564. ISBN: 978-3-95450-147-2. DOI: `doi:10.18429/JACoW-IPAC2016-WEPMY009`.

[34] D. H. Whittum et al. "Electron-hose instability in the ion-focused regime". In: *Phys. Rev. Lett.* 67 (8 Aug. 1991), pp. 991–994. DOI: `10.1103/PhysRevLett.67.991`.

[35] A. Chao. *Physics of Collective Beam Instabilities in High Energy Accelerators*. Wiley Series in Beam Physics and Accelerator Technology. Wiley, 1993. ISBN: 9780471551843. URL: `https://books.google.de/books?id=MjRoQgAACAAJ`.

[36] W. K. H. Panofsky and M. Bander. "Asymptotic Theory of Beam Break-Up in Linear Accelerators". In: *Review of Scientific Instruments* 39.2 (1968), pp. 206–212. DOI: `10.1063/1.1683315`. eprint: `https://doi.org/10.1063/1.1683315`.

[37] K. L. F. Bane and M. Sands. "The Short-Range Resistive Wall Wakefields". In: *AIP Conference Proceedings* 367.1 (1996), pp. 131–149. DOI: `10.1063/1.50300`. eprint: `https://aip.scitation.org/doi/pdf/10.1063/1.50300`.

[38] M. Migliorati et al. "Intrinsic normalized emittance growth in laser-driven electron accelerators". In: *Phys. Rev. ST Accel. Beams* 16 (1 Jan. 2013), p. 011302. DOI: `10.1103/PhysRevSTAB.16.011302`.

[39] B. W. Montague. "Emittance growth from multiple scattering in the plasma beat-wave accelerator". In: *CAS-ECFA-INFN Workshop : Generation of High Fields for Particle Acceleration to Very-high Energies*. 1985. DOI: `10.5170/CERN-1985-007.208`.

[40] M. Ferrario, M. Migliorati, and L. Palumbo. “Space Charge Effects”. In: *CAS - CERN Accelerator School: Advanced Accelerator Physics*. 2014, pp. 331–356. DOI: `10.5170/CERN-2014-009.331`. eprint: `1601.05214`.

[41] S. Di Mitri. “Coherent Synchrotron Radiation and Microbunching Instability”. In: *CERN Yellow Rep. School Proc.* 1 (2018). Ed. by R. Bayley, p. 381. DOI: `10.23730/CYRSP-2018-001.381`.

[42] E. A. Seddon et al. “Short-wavelength free-electron laser sources and science: a review”. In: *Reports on Progress in Physics* 80.11 (Oct. 2017), p. 115901. DOI: `10.1088/1361-6633/aa7cca`.

[43] R. Bonifacio, C. Pellegrini, and L. Narducci. “Collective instabilities and high-gain regime in a free electron laser”. In: *Optics Communications* 50.6 (1984), pp. 373–378. ISSN: 0030-4018. DOI: `https://doi.org/10.1016/0030-4018(84)90105-6`.

[44] L. Tonks and I. Langmuir. “Oscillations in Ionized Gases”. In: *Phys. Rev.* 33 (2 Feb. 1929), pp. 195–210. DOI: `10.1103/PhysRev.33.195`.

[45] S. J. Thomson. “XLII. Ionization by moving electrified particles”. In: *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science* 23.136 (1912), pp. 449–457. DOI: `10.1080/14786440408637241`. eprint: `https://doi.org/10.1080/14786440408637241`.

[46] C. L. O’Connell et al. “Plasma production via field ionization”. In: *Phys. Rev. ST Accel. Beams* 9 (10 Oct. 2006), p. 101301. DOI: `10.1103/PhysRevSTAB.9.101301`.

[47] G. Wittig. “Start-to-end simulations of hybrid laser- and beam-driven plasma wakefield acceleration and free electron lasers”. PhD thesis. Universität Hamburg, 2017. URL: `https://ediss.sub.uni-hamburg.de/handle/ediss/7487`.

[48] M. V. Ammosov, N. B. Delone, and V. P. Krainov. “Tunnel Ionization Of Complex Atoms And Atomic Ions In Electromagnetic Field”. In: *High Intensity Laser Processes*. Ed. by J. A. Alcock. Vol. 0664. International Society for Optics and Photonics. SPIE, 1986, pp. 138–141. DOI: `10.1117/12.938695`.

[49] M. Chen et al. “Numerical modeling of laser tunneling ionization in explicit particle-in-cell codes”. In: *Journal of Computational Physics* 236 (2013), pp. 220–228. ISSN: 0021-9991. DOI: `https://doi.org/10.1016/j.jcp.2012.11.029`.

[50] G. Tauscher et al. "Theoretical and experimental studies of plasma generation. Tailoring plasmas for wakefield accelerators". In: 3rd European Advanced Accelerator Concepts Workshop, La Biodola, Isola d'Elba (Italy), 24 Sep 2017 - 30 Sep 2017. Sept. 24, 2017. DOI: `10.3204/PUBDB-2017-11827`.

[51] A. Zavriyev et al. "Ionization and dissociation of H_2 in intense laser fields at 1.064 μm, 532 nm, and 355 nm". In: *Phys. Rev. A* 42 (9 Nov. 1990), pp. 5500–5513. DOI: `10.1103/PhysRevA.42.5500`.

[52] D. Shiner et al. "H_2, D_2, and HD ionization potentials by accurate calibration of several iodine lines". In: *Phys. Rev. A* 47 (5 May 1993), pp. 4042–4045. DOI: `10.1103/PhysRevA.47.4042`.

[53] C. Arran. "Techniques for high repetition rate laser wakefield acceleration". PhD thesis. University of Oxford, 2018.

[54] R. J. Shalloo et al. "Hydrodynamic optical-field-ionized plasma channels". In: *Phys. Rev. E* 97 (5 May 2018), p. 053203. DOI: `10.1103/PhysRevE.97.053203`.

[55] S. Z. Green et al. "Laser ionized preformed plasma at FACET". In: *Plasma Physics and Controlled Fusion* 56.8 (July 2014), p. 084011. DOI: `10.1088/0741-3335/56/8/084011`.

[56] G. Wittig et al. "Optical plasma torch electron bunch generation in plasma wakefield accelerators". In: *Phys. Rev. ST Accel. Beams* 18 (8 Aug. 2015), p. 081304. DOI: `10.1103/PhysRevSTAB.18.081304`.

[57] S. Rae. "Ionization-induced defocusing of intense laser pulses in high-pressure gases". In: *Optics Communications* 97.1 (1993), pp. 25–28. ISSN: 0030-4018. DOI: `https://doi.org/10.1016/0030-4018(93)90611-8`.

[58] W. Lotz. "Electron-impact ionization cross-sections and ionization rate coefficients for atoms and ions from hydrogen to calcium". In: *Zeitschrift für Physik* 216.3 (June 1968), pp. 241–247. ISSN: 0044-3328. DOI: `10.1007/BF01392963`.

[59] J. D. Gillaspy. "Highly charged ions". In: *Journal of Physics B: Atomic, Molecular and Optical Physics* 34.19 (Sept. 2001), R93. DOI: `10.1088/0953-4075/34/19/201`.

[60] C. G. Durfee and H. M. Milchberg. "Light pipe for high intensity laser pulses". In: *Phys. Rev. Lett.* 71 (15 Oct. 1993), pp. 2409–2412. DOI: `10.1103/PhysRevLett.71.2409`.

[61] H. A. Kramers. "XCIII. On the theory of X-ray absorption and of the continuous X-ray spectrum". In: *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science* 46.275 (1923), pp. 836–871. DOI: `10.1080/14786442308565244`. eprint: `https://doi.org/10.1080/14786442308565244`.

[62] A. E. Martirosyan et al. "Time evolution of plasma afterglow produced by femtosecond laser pulses". In: *Journal of Applied Physics* 96.10 (2004), pp. 5450–5455. DOI: `10.1063/1.1803920`. eprint: `https://doi.org/10.1063/1.1803920`.

[63] P. Scherkl. "High-brightness plasma-based Compton backscattering source for high energy physics". PhD thesis. University of Strathclyde, 2020. DOI: `10.48730/e7nv-ek71`.

[64] V. Lee et al. *Temporal Evolution of the Light Emitted by a Thin, Laser-ionized Plasma Source*. Sept. 2023. arXiv: `2309.10723 [physics.plasm-ph]`.

[65] S. S. Harilal, B. E. Brumfield, and M. C. Phillips. "Lifecycle of laser-produced air sparks". In: *Physics of Plasmas* 22.6 (2015), p. 063301. DOI: `10.1063/1.4922076`. eprint: `https://doi.org/10.1063/1.4922076`.

[66] J. M. Garland et al. "Combining laser interferometry and plasma spectroscopy for spatially resolved high-sensitivity plasma density measurements in discharge capillaries". In: *Review of Scientific Instruments* 92.1 (2021), p. 013505. DOI: `10.1063/5.0021117`. eprint: `https://doi.org/10.1063/5.0021117`.

[67] V. I. Veksler. "Coherent principle of acceleration of charged particles". In: *CERN Symposium on High Energy Accelerators and Pion Physics*. 1956. DOI: `10.5170/CERN-1956-025.80`.

[68] I. B. Fainberg. "The use of plasma waveguides as accelerating structures in linear accelerators". In: *CERN Symposium on High Energy Accelerators and Pion Physics*. 1956. DOI: `10.5170/CERN-1956-025.84`.

[69] T. Tajima and J. M. Dawson. "Laser Electron Accelerator". In: *Phys. Rev. Lett.* 43 (4 July 1979), pp. 267–270. DOI: `10.1103/PhysRevLett.43.267`.

[70] C. E. Clayton et al. "Ultrahigh-gradient acceleration of injected electrons by laser-excited relativistic electron plasma waves". In: *Phys. Rev. Lett.* 70 (1 Jan. 1993), pp. 37–40. DOI: `10.1103/PhysRevLett.70.37`.

[71] A. Modena et al. "Electron acceleration from the breaking of relativistic plasma waves". In: *Nature* 377.6550 (Oct. 1995), pp. 606–608. ISSN: 1476-4687. DOI: `10.1038/377606a0`.

[72] V. Malka et al. "Electron Acceleration by a Wake Field Forced by an Intense Ultrashort Laser Pulse". In: *Science* 298.5598 (2002), pp. 1596–1600. DOI: `10.1126/science.1076782`. eprint: `https://www.science.org/doi/pdf/10.1126/science.1076782`.

[73] D. Strickland and G. Mourou. "Compression of amplified chirped optical pulses". In: *Optics Communications* 55.6 (1985), pp. 447–449. ISSN: 0030-4018. DOI: `https://doi.org/10.1016/0030-4018(85)90151-8`.

[74] S. P. D. Mangles et al. "Monoenergetic beams of relativistic electrons from intense laser–plasma interactions". In: *Nature* 431.7008 (Sept. 2004), pp. 535–538. ISSN: 1476-4687. DOI: `10.1038/nature02939`.

[75] C. G. R. Geddes et al. "High-quality electron beams from a laser wakefield accelerator using plasma-channel guiding". In: *Nature* 431.7008 (Sept. 2004), pp. 538–541. ISSN: 1476-4687. DOI: `10.1038/nature02900`.

[76] J. Faure et al. "A laser–plasma accelerator producing monoenergetic electron beams". In: *Nature* 431.7008 (Sept. 2004), pp. 541–544. ISSN: 1476-4687. DOI: `10.1038/nature02963`.

[77] P. Chen et al. "Acceleration of Electrons by the Interaction of a Bunched Electron Beam with a Plasma". In: *Phys. Rev. Lett.* 54 (7 Feb. 1985), pp. 693–696. DOI: `10.1103/PhysRevLett.54.693`.

[78] R. D. Ruth et al. "A Plasma Wake Field Accelerator". In: *Part. Accel.* 17 (1985), p. 171.

[79] J. B. Rosenzweig et al. "Experimental Observation of Plasma Wake-Field Acceleration". In: *Phys. Rev. Lett.* 61 (1 July 1988), pp. 98–101. DOI: `10.1103/PhysRevLett.61.98`.

[80] I. Blumenfeld et al. "Energy doubling of 42 GeV electrons in a metre-scale plasma wakefield accelerator". In: *Nature* 445.7129 (Feb. 2007), pp. 741–744. ISSN: 1476-4687. DOI: `10.1038/nature05538`.

[81] E. Adli et al. "Acceleration of electrons in the plasma wakefield of a proton bunch". In: *Nature* 561.7723 (Sept. 2018), pp. 363–367. ISSN: 1476-4687. DOI: `10.1038/s41586-018-0485-4`.

[82] M. J. Hogan et al. "Ultrarelativistic-Positron-Beam Transport through Meter-Scale Plasmas". In: *Phys. Rev. Lett.* 90 (20 May 2003), p. 205002. DOI: `10.1103/PhysRevLett.90.205002`.

[83] B. E. Blue et al. "Plasma-Wakefield Acceleration of an Intense Positron Beam". In: *Phys. Rev. Lett.* 90 (21 May 2003), p. 214801. DOI: `10.1103/PhysRevLett.90.214801`.

[84] S. Corde et al. "Multi-gigaelectronvolt acceleration of positrons in a self-loaded plasma wakefield". In: *Nature* 524.7566 (Aug. 2015), pp. 442–445. ISSN: 1476-4687. DOI: `10.1038/nature14890`.

[85] S. Gessner et al. "Demonstration of a positron beam-driven hollow channel plasma wakefield accelerator". In: *Nature Communications* 7.1 (June 2016), p. 11785. ISSN: 2041-1723. DOI: `10.1038/ncomms11785`.

[86] I. Blumenfeld. "Scaling of the longitudinal electric fields and transformer ratio in a non-linear Plasma Wakefield Accelerator". English. PhD thesis. 2009, p. 148. ISBN: 978-1-109-45010-1. URL: `https://www.proquest.com/dissertations-theses/scaling-longitudinal-electric-fields-transformer/docview/305009807/se-2`.

[87] R. Keinigs and M. E. Jones. "Two-dimensional dynamics of the plasma wakefield accelerator". In: *The Physics of Fluids* 30.1 (1987), pp. 252–263. DOI: `10.1063/1.866183`. eprint: `https://aip.scitation.org/doi/pdf/10.1063/1.866183`.

[88] J. B. Rosenzweig et al. "Acceleration and focusing of electrons in two-dimensional nonlinear plasma wake fields". In: *Phys. Rev. A* 44 (10 Nov. 1991), R6189–R6192. DOI: `10.1103/PhysRevA.44.R6189`.

[89] J. M. Dawson. "Nonlinear Electron Oscillations in a Cold Plasma". In: *Phys. Rev.* 113 (2 Jan. 1959), pp. 383–387. DOI: `10.1103/PhysRev.113.383`.

[90] W. Lu et al. "Limits of linear plasma wakefield theory for electron or positron beams". In: *Physics of Plasmas* 12.6 (2005), p. 063101. DOI: `10.1063/1.1905587`. eprint: `https://doi.org/10.1063/1.1905587`.

[91] W. Lu et al. "Nonlinear Theory for Relativistic Plasma Wakefields in the Blowout Regime". In: *Phys. Rev. Lett.* 96 (16 Apr. 2006), p. 165002. DOI: `10.1103/PhysRevLett.96.165002`.

[92] W. Lu et al. "A nonlinear theory for multidimensional relativistic plasma wave wakefields". In: *Physics of Plasmas* 13.5 (2006), p. 056709. DOI: `10.1063/1.2203364`. eprint: `https://doi.org/10.1063/1.2203364`.

[93] N. Barov et al. “Observation of plasma wakefield acceleration in the underdense regime”. In: *Phys. Rev. ST Accel. Beams* 3 (1 Jan. 2000), p. 011301. DOI: `10.1103/PhysRevSTAB.3.011301`.

[94] K. Nakajima et al. “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 (1990), pp. 12–20. ISSN: 0168-9002. DOI: `https://doi.org/10.1016/0168-9002(90)91729-U`.

[95] E. Kallos et al. “High-Gradient Plasma-Wakefield Acceleration with Two Subpicosecond Electron Bunches”. In: *Phys. Rev. Lett.* 100 (7 Feb. 2008), p. 074802. DOI: `10.1103/PhysRevLett.100.074802`.

[96] C. B. Schroeder et al. “Trapping, dark current, and wave breaking in nonlinear plasma waves”. In: *Physics of Plasmas* 13.3 (2006), p. 033103. DOI: `10.1063/1.2173960`. eprint: `https://doi.org/10.1063/1.2173960`.

[97] S. Y. Kalmykov et al. “Laser plasma acceleration with a negatively chirped pulse: all-optical control over dark current in the blowout regime”. In: *New Journal of Physics* 14.3 (Mar. 2012), p. 033025. DOI: `10.1088/1367-2630/14/3/033025`.

[98] G. G. Manahan et al. “Hot spots and dark current in advanced plasma wakefield accelerators”. In: *Phys. Rev. Accel. Beams* 19 (1 Jan. 2016), p. 011303. DOI: `10.1103/PhysRevAccelBeams.19.011303`.

[99] E. Oz et al. “Ionization-Induced Electron Trapping in Ultrarelativistic Plasma Wakes”. In: *Phys. Rev. Lett.* 98 (8 Feb. 2007), p. 084801. DOI: `10.1103/PhysRevLett.98.084801`.

[100] N. Vafaei-Najafabadi et al. “Evidence for high-energy and low-emittance electron beams using ionization injection of charge in a plasma wakefield accelerator”. In: *Plasma Physics and Controlled Fusion* 58.3 (Feb. 2016), p. 034009. DOI: `10.1088/0741-3335/58/3/034009`.

[101] N. Kirby et al. “Transverse emittance and current of multi-GeV trapped electrons in a plasma wakefield accelerator”. In: *Phys. Rev. ST Accel. Beams* 12 (5 May 2009), p. 051302. DOI: `10.1103/PhysRevSTAB.12.051302`.

[102] A. Martinez de la Ossa et al. “High-Quality Electron Beams from Beam-Driven Plasma Accelerators by Wakefield-Induced Ionization Injection”. In: *Phys. Rev. Lett.* 111 (24 Dec. 2013), p. 245003. DOI: `10.1103/PhysRevLett.111.245003`.

[103] F. Li et al. "Generating High-Brightness Electron Beams via Ionization Injection by Transverse Colliding Lasers in a Plasma-Wakefield Accelerator". In: *Phys. Rev. Lett.* 111 (1 July 2013), p. 015003. DOI: `10.1103/PhysRevLett.111.015003`.

[104] B. Hidding et al. "Ultracold Electron Bunch Generation via Plasma Photocathode Emission and Acceleration in a Beam-Driven Plasma Blowout". In: *Phys. Rev. Lett.* 108 (3 Jan. 2012), p. 035001. DOI: `10.1103/PhysRevLett.108.035001`.

[105] S. Bulanov et al. "Particle injection into the wave acceleration phase due to nonlinear wake wave breaking". In: *Phys. Rev. E* 58 (5 Nov. 1998), R5257–R5260. DOI: `10.1103/PhysRevE.58.R5257`.

[106] F. S. Tsung et al. "Near-GeV-Energy Laser-Wakefield Acceleration of Self-Injected Electrons in a Centimeter-Scale Plasma Channel". In: *Phys. Rev. Lett.* 93 (18 Oct. 2004), p. 185002. DOI: `10.1103/PhysRevLett.93.185002`.

[107] S. Kalmykov et al. "Electron Self-Injection and Trapping into an Evolving Plasma Bubble". In: *Phys. Rev. Lett.* 103 (13 Sept. 2009), p. 135004. DOI: `10.1103/PhysRevLett.103.135004`.

[108] A. Buck et al. "Absolute charge calibration of scintillating screens for relativistic electron detection". In: *Review of Scientific Instruments* 81.3 (2010), p. 033301. DOI: `10.1063/1.3310275`. eprint: `https://doi.org/10.1063/1.3310275`.

[109] M. Hansson et al. "Down-ramp injection and independently controlled acceleration of electrons in a tailored laser wakefield accelerator". In: *Phys. Rev. ST Accel. Beams* 18 (7 July 2015), p. 071303. DOI: `10.1103/PhysRevSTAB.18.071303`.

[110] J. Faure et al. "Injection and acceleration of quasimonoenergetic relativistic electron beams using density gradients at the edges of a plasma channel". In: *Physics of Plasmas* 17.8 (2010), p. 083107. DOI: `10.1063/1.3469581`. eprint: `https://doi.org/10.1063/1.3469581`.

[111] K. Schmid et al. "Density-transition based electron injector for laser driven wakefield accelerators". In: *Phys. Rev. ST Accel. Beams* 13 (9 Sept. 2010), p. 091301. DOI: `10.1103/PhysRevSTAB.13.091301`.

[112] T. Katsouleas. "Physical mechanisms in the plasma wake-field accelerator". In: *Phys. Rev. A* 33 (3 Mar. 1986), pp. 2056–2064. DOI: `10.1103/PhysRevA.33.2056`.

[113] G. Fubiani et al. "Improvement of electron beam quality in optical injection schemes using negative plasma density gradients". In: *Phys. Rev. E* 73 (2 Feb. 2006), p. 026402. DOI: `10.1103/PhysRevE.73.026402`.

[114] H. Suk et al. "Plasma Electron Trapping and Acceleration in a Plasma Wake Field Using a Density Transition". In: *Phys. Rev. Lett.* 86 (6 Feb. 2001), pp. 1011–1014. DOI: `10.1103/PhysRevLett.86.1011`.

[115] J. Grebenyuk et al. "Beam-driven plasma-based acceleration of electrons with density down-ramp injection at FLASHForward". In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 740 (2014), pp. 246–249. ISSN: 0168-9002. DOI: `https://doi.org/10.1016/j.nima.2013.10.054`.

[116] A. Martinez de la Ossa et al. "Optimizing density down-ramp injection for beam-driven plasma wakefield accelerators". In: *Phys. Rev. Accel. Beams* 20 (9 Sept. 2017), p. 091301. DOI: `10.1103/PhysRevAccelBeams.20.091301`.

[117] C. Zhang et al. "Effect of fluctuations in the down ramp plasma source profile on the emittance and current profile of the self-injected beam in a plasma wakefield accelerator". In: *Phys. Rev. Accel. Beams* 22 (11 Nov. 2019), p. 111301. DOI: `10.1103/PhysRevAccelBeams.22.111301`.

[118] F. M. Foerster et al. "Stable and High-Quality Electron Beams from Staged Laser and Plasma Wakefield Accelerators". In: *Phys. Rev. X* 12 (4 Nov. 2022), p. 041016. DOI: `10.1103/PhysRevX.12.041016`.

[119] J. P. Couperus Cabadağ et al. "Gas-dynamic density downramp injection in a beam-driven plasma wakefield accelerator". In: *Phys. Rev. Research* 3 (4 Oct. 2021), p. L042005. DOI: `10.1103/PhysRevResearch.3.L042005`.

[120] X. L. Xu et al. "High quality electron bunch generation using a longitudinal density-tailored plasma-based accelerator in the three-dimensional blowout regime". In: *Phys. Rev. Accel. Beams* 20 (11 Nov. 2017), p. 111303. DOI: `10.1103/PhysRevAccelBeams.20.111303`.

[121] K. L. F. Bane, P. B. Wilson, and T. Weiland. "Wake fields and wake field acceleration". In: *AIP Conference Proceedings* 127.1 (1985), pp. 875–928. DOI: `10.1063/1.35182`. eprint: `https://aip.scitation.org/doi/pdf/10.1063/1.35182`.

[122] P. Chen et al. "Energy Transfer in the Plasma Wake-Field Accelerator". In: *Phys. Rev. Lett.* 56 (12 Mar. 1986), pp. 1252–1255. DOI: `10.1103/PhysRevLett.56.1252`.

[123] K. L. F. Bane, P. Chen, and P. B. Wilson. "On Collinear Wake Field Acceleration". In: *IEEE Transactions on Nuclear Science* 32.5 (1985), pp. 3524–3526. DOI: `10.1109/TNS.1985.4334416`.

[124] G. Loisch et al. "Observation of High Transformer Ratio Plasma Wakefield Acceleration". In: *Phys. Rev. Lett.* 121 (6 Aug. 2018), p. 064801. DOI: `10.1103/PhysRevLett.121.064801`.

[125] R. Roussel et al. "Single Shot Characterization of High Transformer Ratio Wakefields in Nonlinear Plasma Acceleration". In: *Phys. Rev. Lett.* 124 (4 Jan. 2020), p. 044802. DOI: `10.1103/PhysRevLett.124.044802`.

[126] P. B. Wilson and J. E. Griffin. "High energy electron linacs; application to storage ring RF systems and linear colliders". In: *AIP Conference Proceedings* 87.1 (1982), pp. 450–555. DOI: `10.1063/1.33620`. eprint: `https://aip.scitation.org/doi/pdf/10.1063/1.33620`.

[127] M. Tzoufras et al. "Beam Loading in the Nonlinear Regime of Plasma-Based Acceleration". In: *Phys. Rev. Lett.* 101 (14 Sept. 2008), p. 145002. DOI: `10.1103/PhysRevLett.101.145002`.

[128] M. Seidel. "Towards Efficient Particle Accelerators - A Review". In: *JACoW* IPAC2022 (2022), FRPLYGD1. DOI: `10.18429/JACoW-IPAC2022-FRPLYGD1`.

[129] M. Aicheler et al. *A Multi-TeV Linear Collider Based on CLIC Technology: CLIC Conceptual Design Report*. CERN Yellow Reports: Monographs. Geneva: CERN, 2012. DOI: `10.5170/CERN-2012-007`.

[130] S. M. Hooker et al. "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 (Nov. 2014), p. 234003. DOI: `10.1088/0953-4075/47/23/234003`.

[131] J. Cowley et al. "Excitation and Control of Plasma Wakefields by Multiple Laser Pulses". In: *Phys. Rev. Lett.* 119 (4 July 2017), p. 044802. DOI: `10.1103/PhysRevLett.119.044802`.

[132] T. Nubbemeyer et al. "1 kW, 200 mJ picosecond thin-disk laser system". In: *Opt. Lett.* 42.7 (Apr. 2017), pp. 1381–1384. DOI: `10.1364/OL.42.001381`.

[133] C. Jauregui, J. Limpert, and A. Tünnermann. "High-power fibre lasers". In: *Nature Photonics* 7.11 (Nov. 2013), pp. 861–867. ISSN: 1749-4893. DOI: `10.1038/nphoton.2013.273`.

[134] S. van der Meer. *Improving the power efficiency of the plasma wakefield accelerator*. Tech. rep. Geneva: CERN, 1985. URL: `https://cds.cern.ch/record/163918`.

[135] T. C. Katsouleas et al. "Beam Loading in Plasma Accelerators". In: *Part. Accel.* 22 (1987), pp. 81–99.

[136] K. V. Lotov. “Efficient operating mode of the plasma wakefield accelerator”. In: *Physics of Plasmas* 12.5 (2005), p. 053105. DOI: `10.1063/1.1889444`. eprint: `https://doi.org/10.1063/1.1889444`.

[137] V. Lebedev, A. Burov, and S. Nagaitsev. “Efficiency versus instability in plasma accelerators”. In: *Phys. Rev. Accel. Beams* 20 (12 Dec. 2017), p. 121301. DOI: `10.1103/PhysRevAccelBeams.20.121301`.

[138] S. S. Baturin. *On "Efficiency versus Instability in plasma accelerators"*. 2020. arXiv: `2010.01046 [physics.acc-ph]`.

[139] C. A. Lindstrøm et al. “Energy-Spread Preservation and High Efficiency in a Plasma-Wakefield Accelerator”. In: *Phys. Rev. Lett.* 126 (1 Jan. 2021), p. 014801. DOI: `10.1103/PhysRevLett.126.014801`.

[140] K. Steffen. *High Energy Beam Optics*. Interscience monographs and texts in physics and astronomy. Interscience Publishers, 1965. URL: `https://books.google.co.uk/books?id=xB9RAAAAMAAJ`.

[141] M. Ferrario. “Injection, Extraction and Matching”. In: *CAS - CERN Accelerator School 2019: High Gradient Wakefield Accelerators*. July 2020. arXiv: `2007.04102 [physics.acc-ph]`.

[142] C. E. Clayton et al. “Transverse Envelope Dynamics of a 28.5-GeV Electron Beam in a Long Plasma”. In: *Phys. Rev. Lett.* 88 (15 Apr. 2002), p. 154801. DOI: `10.1103/PhysRevLett.88.154801`.

[143] K. Floettmann. “Adiabatic matching section for plasma accelerated beams”. In: *Phys. Rev. ST Accel. Beams* 17 (5 May 2014), p. 054402. DOI: `10.1103/PhysRevSTAB.17.054402`.

[144] I. Dornmair, K. Floettmann, and A. R. Maier. “Emittance conservation by tailored focusing profiles in a plasma accelerator”. In: *Phys. Rev. ST Accel. Beams* 18 (4 Apr. 2015), p. 041302. DOI: `10.1103/PhysRevSTAB.18.041302`.

[145] K. Marsh et al. “Beam Matching to a Plasma Wake Field Accelerator using a Ramped Density Profile at the Plasma Boundary”. In: *Proceedings of the 2005 Particle Accelerator Conference*. 2005, pp. 2702–2704. DOI: `10.1109/PAC.2005.1591234`.

[146] P. Antici et al. “Laser-driven electron beamlines generated by coupling laser-plasma sources with conventional transport systems”. In: *Journal of Applied Physics* 112.4 (2012), p. 044902. DOI: `10.1063/1.4740456`. eprint: `https://doi.org/10.1063/1.4740456`.

[147] I. Agapov et al. "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 (2014), pp. 151–156. ISSN: 0168-9002. DOI: `https://doi.org/10.1016/j.nima.2014.09.057`.

[148] C. K. Birdsall and A. B. Langdon. *Plasma Physics Via Computer Simulation*. Adam Hilger series on plasma physics. McGraw-Hill, 1985. ISBN: 9780070053717. URL: `https://books.google.de/books?id=7TMbAQAAIAAJ`.

[149] J.-L. Vay and R. Lehe. "Simulations for Plasma and Laser Acceleration". In: *Reviews of Accelerator Science and Technology* 09 (2016), pp. 165–186. DOI: `10.1142/S1793626816300085`. eprint: `https://doi.org/10.1142/S1793626816300085`.

[150] A. Myers et al. "Porting WarpX to GPU-accelerated platforms". In: *Parallel Computing* 108 (2021), p. 102833. ISSN: 0167-8191. DOI: `https://doi.org/10.1016/j.parco.2021.102833`.

[151] R. Lehe et al. "A spectral, quasi-cylindrical and dispersion-free Particle-In-Cell algorithm". In: *Computer Physics Communications* 203 (2016), pp. 66–82. ISSN: 0010-4655. DOI: `https://doi.org/10.1016/j.cpc.2016.02.007`.

[152] A. Aschikhin et al. "The FLASHForward facility at DESY". In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 806 (2016), pp. 175–183. ISSN: 0168-9002. DOI: `https://doi.org/10.1016/j.nima.2015.10.005`.

[153] R. D'Arcy et al. "FLASHForward: plasma wakefield accelerator science for high-average-power applications". In: *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 377.2151 (2019), p. 20180392. DOI: `10.1098/rsta.2018.0392`. eprint: `https://royalsocietypublishing.org/doi/pdf/10.1098/rsta.2018.0392`.

[154] S. Schreiber and B. Faatz. "The free-electron laser FLASH". In: *High Power Laser Science and Engineering* 3 (2015), e20. DOI: `10.1017/hpl.2015.16`.

[155] J. Rossbach, J. R. Schneider, and W. Wurth. "10 years of pioneering X-ray science at the Free-Electron Laser FLASH at DESY". In: *Physics Reports* 808 (2019), pp. 1–74. ISSN: 0370-1573. DOI: `https://doi.org/10.1016/j.physrep.2019.02.002`.

[156] R. Abela et al. *XFEL: The European X-Ray Free-Electron Laser - Technical Design Report*. Hamburg: DESY, 2006, pp. 1–646. ISBN: 978-3-935702-17-1. DOI: `10.3204/DESY_06-097`.

[157] S. Schreiber et al. "Update on the Lifetime of Cs2Te Photocathodes Operated at FLASH". In: *Proc. of International Free Electron Laser Conference (FEL'17), Santa Fe, NM, USA, August 20-25, 2017* (Santa Fe, NM, USA). International Free Electron Laser Conference 38. Geneva, Switzerland: JACoW, Feb. 2018, pp. 415–418. ISBN: 978-3-95450-179-3. DOI: `https://doi.org/10.18429/JACoW-FEL2017-WEP003`.

[158] J. Rönsch-Schulenburg et al. "Short SASE-FEL Pulses at FLASH". In: *FEL 2013: Proceedings of the 35th International Free-Electron Laser Conference*. Aug. 2013, pp. 379–382.

[159] T. Plath et al. "Commissioning and diagnostics development for the new short-pulse injector laser at FLASH". In: *Proceedings, 2nd International Beam Instrumentation Conference*. Sept. 2013, pp. 353–356.

[160] J. Rossbach. "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 E. J. Jaeschke et al. Cham: Springer International Publishing, 2020, pp. 323–348. ISBN: 978-3-030-23201-6. DOI: `10.1007/978-3-030-23201-6_10`.

[161] H. Edwards et al. "3.9 GHz Cavity Module for Linear Bunch Compression at FLASH". In: 25th International Linear Accelerator Conference, Tsukuba (Japan), 12 Sep 2010 - 17 Sep 2010. KEK, Sept. 12, 2010. URL: `https://bib-pubdb1.desy.de/record/91880`.

[162] E. Vogel et al. "Test and Commissioning of the Third Harmonic RF System for FLASH". In: *Proc. IPAC'10* (Kyoto, Japan, May 2010). JACoW Publishing, Geneva, Switzerland, pp. 4281–4283. URL: `http://accelconf.web.cern.ch/IPAC10/papers/THPD003.pdf`.

[163] T. Limberg et al. "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 (1996), pp. 322–324. ISSN: 0168-9002. DOI: `https://doi.org/10.1016/0168-9002(95)01470-5`.

[164] B. Faatz et al. "Simultaneous operation of two soft x-ray free-electron lasers driven by one linear accelerator". In: *New Journal of Physics* 18.6 (June 2016), p. 062002. DOI: `10.1088/1367-2630/18/6/062002`.

[165] R. Zgadzaj et al. "Dissipation of electron-beam-driven plasma wakes". In: *Nature Communications* 11.1 (Sept. 2020), p. 4753.

[166] V. Ayvazyan et al. "Low level RF control implementation and simultaneous operation of two FEL undulator beamlines at FLASH". In: *Proc. of ICALEPCS2015 (Melbourne, Australia)* (2015), pp. 42–5.

[167] N. Baboi et al. "Resolution Studies at Beam Position Monitors at the FLASH Facility at DESY". In: *AIP Conference Proceedings* 868.1 (2006), pp. 227–237. DOI: `10.1063/1.2401409`. eprint: `https://aip.scitation.org/doi/pdf/10.1063/1.2401409`.

[168] M. Viti et al. "The Bunch Arrival Time Monitor at FLASH and European XFEL". In: *Proc. of International Conference on Accelerator and Large Experimental Control Systems (ICALEPCS'17), Barcelona, Spain, 8-13 October 2017* (Barcelona, Spain). International Conference on Accelerator and Large Experimental Control Systems 16. Geneva, Switzerland: JACoW, Jan. 2018, pp. 701–705. ISBN: 978-3-95450-193-9. DOI: `https://doi.org/10.18429/JACoW-ICALEPCS2017-TUPHA125`.

[169] F. Loehl. "Measurements of the transverse emittance at the VUV-FEL". Dipl. Hamburg: Universität Hamburg, 2005, p. 56. DOI: `10.3204/DESY-THESIS-2005-014`.

[170] S. Wesch et al. "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 (2011), pp. 40–47. ISSN: 0168-9002. DOI: `https://doi.org/10.1016/j.nima.2011.11.037`.

[171] O. H. Altenmueller, R. R. Larsen, and G. A. Loew. "Investigations of Traveling-Wave Separators for the Stanford Two-Mile Linear Accelerator". In: *Review of Scientific Instruments* 35.4 (1964), pp. 438–442. DOI: `10.1063/1.1718840`. eprint: `https://doi.org/10.1063/1.1718840`.

[172] S. Schulz et al. "Femtosecond all-optical synchronization of an X-ray free-electron laser". In: *Nature Communications* 6.1 (Jan. 2015), p. 5938. ISSN: 2041-1723. DOI: `10.1038/ncomms6938`.

[173] S. Schulz et al. “Past, Present and Future Aspects of Laser-Based Synchronization at FLASH”. In: 2nd International Beam Instrumentation Conference, Oxford (United Kingdom), 16 Sep 2013 - 19 Sep 2013. Geneva: JaCow, Sept. 16, 2013. ISBN: 978-3-95450-127-4. URL: `https://bib-pubdb1.desy.de/record/165641`.

[174] M. Felber et al. “RF-based Synchronization of the Seed and Pump-Probe Lasers to the Optical Synchronization System at FLASH”. In: *Proc. FEL'10* (Malmö, Sweden, Aug. 2010). JACoW Publishing, Geneva, Switzerland, pp. 544–547. URL: `https://jacow.org/FEL2010/papers/THOA3.pdf`.

[175] C. Schmidt et al. “Performance of the MTCA.4 Based LLRF System at FLASH”. In: *Proc. 5th International Particle Accelerator Conference (IPAC'14), Dresden, Germany, June 15-20, 2014* (Dresden, Germany). International Particle Accelerator Conference 5. Geneva, Switzerland: JACoW, July 2014, pp. 2433–2435. ISBN: 978-3-95450-132-8. DOI: `https://doi.org/10.18429/JACoW-IPAC2014-WEPME067`.

[176] M. Scholz et al. “Extraction Arc for FLASH II”. In: *Proc. FEL'12* (Nara, Japan, Aug. 2012). JACoW Publishing, Geneva, Switzerland, pp. 305–307. URL: `https://jacow.org/FEL2012/papers/TUPD33.pdf`.

[177] S. Schröder et al. “Tunable and precise two-bunch generation at FLASHForward”. In: *Journal of Physics: Conference Series* 1596.1 (July 2020), p. 012002. DOI: `10.1088/1742-6596/1596/1/012002`.

[178] C. A. Lindstrøm et al. “Matching small β functions using centroid jitter and two beam position monitors”. In: *Phys. Rev. Accel. Beams* 23 (5 May 2020), p. 052802. DOI: `10.1103/PhysRevAccelBeams.23.052802`.

[179] S. Schröder et al. “High-resolution sampling of beam-driven plasma wakefields”. In: *Nature Communications* 11.1 (Nov. 2020), p. 5984. ISSN: 2041-1723. DOI: `10.1038/s41467-020-19811-9`.

[180] D. J. Spence and S. M. Hooker. “Investigation of a hydrogen plasma waveguide”. In: *Phys. Rev. E* 63 (1 Dec. 2000), p. 015401. DOI: `10.1103/PhysRevE.63.015401`.

[181] I. Okunev, I. Morozov, and N. Nefedov. “X-FEL Quadrupole with Gradient of 100 T/m”. In: *Physics Procedia* 84 (Dec. 2016), pp. 101–107. DOI: `10.1016/j.phpro.2016.11.018`.

[182] C. A. Lindstrøm et al. “Progress of the FLASHForward X-2 high-beam-quality, high-efficiency plasma-accelerator experiment”. In: *PoS* EPS-HEP2021 (2022), p. 880. DOI: `10.22323/1.398.0880`.

[183] R. Weingartner et al. “Ultralow emittance electron beams from a laser-wakefield accelerator”. In: *Phys. Rev. ST Accel. Beams* 15 (11 Nov. 2012), p. 111302. DOI: `10.1103/PhysRevSTAB.15.111302`.

[184] D. Möhl. “Sources of emittance growth”. In: *CAS - CERN Accelerator School: Intermediate Course on Accelerator Physics*. CERN, 2006.

[185] G. Kube et al. “Transverse Beam Profile Imaging of Few-Micrometer Beam Sizes Based on a Scintillator Screen”. In: *4th International Beam Instrumentation Conference*. 2016, TUPB012. DOI: `10.18429/JACoW-IBIC2015-TUPB012`.

[186] B. Marchetti et al. “X-Band TDS Project”. In: (2017), MOPAB044. 4 p. DOI: `10.18429/JACoW-IPAC2017-MOPAB044`.

[187] D. E. Spence, P. N. Kean, and W. Sibbett. “60-fsec pulse generation from a self-mode-locked Ti:sapphire laser”. In: *Opt. Lett.* 16.1 (Jan. 1991), pp. 42–44. DOI: `10.1364/OL.16.000042`.

[188] T. Brabec et al. “Kerr lens mode locking”. In: *Opt. Lett.* 17.18 (Sept. 1992), pp. 1292–1294. DOI: `10.1364/OL.17.001292`.

[189] A. Offner. *Unit Power Imaging Catoptric Anastigmat*. U.S. patent 3,748,015 (June 21, 1971).

[190] G. Cheriaux et al. “Aberration-free stretcher design for ultrashort-pulse amplification”. In: *Opt. Lett.* 21.6 (Mar. 1996), pp. 414–416. DOI: `10.1364/OL.21.000414`.

[191] P. Tournois. “Acousto-optic programmable dispersive filter for adaptive compensation of group delay time dispersion in laser systems”. In: *Optics Communications* 140.4 (1997), pp. 245–249. ISSN: 0030-4018. DOI: `https://doi.org/10.1016/S0030-4018(97)00153-3`.

[192] T. Staufer et al. “Development of a laser-wakefield Thomson x-ray source for x-ray fluorescence imaging”. In: *Laser Acceleration of Electrons, Protons, and Ions V*. Ed. by E. Esarey, C. B. Schroeder, and J. Schreiber. Vol. 11037. International Society for Optics and Photonics. SPIE, 2019, p. 110370M. DOI: `10.1117/12.2520685`.

[193] S. Bohlen et al. "Stability of ionization-injection-based laser-plasma accelerators". In: *Phys. Rev. Accel. Beams* 25 (3 Mar. 2022), p. 031301. DOI: `10.1103/PhysRevAccelBeams.25.031301`.

[194] T. Oksenhendler et al. "Self-referenced spectral interferometry". In: *Applied Physics B* 99.1 (Apr. 2010), pp. 7–12.

[195] X. Shen et al. "Self-Referenced Spectral Interferometry for Femtosecond Pulse Characterization". In: *Applied Sciences* 7.4 (2017). ISSN: 2076-3417. DOI: `10.3390/app7040407`.

[196] N. Minkovski et al. "Nonlinear polarization rotation and orthogonal polarization generation experienced in a single-beam configuration". In: *J. Opt. Soc. Am. B* 21.9 (Sept. 2004), pp. 1659–1664. DOI: `10.1364/JOSAB.21.001659`.

[197] A. Knetsch et al. "Stable witness-beam formation in a beam-driven plasma cathode". In: *Phys. Rev. Accel. Beams* 24 (10 Oct. 2021), p. 101302. DOI: `10.1103/PhysRevAccelBeams.24.101302`.

[198] V. Libov et al. "FLASHForward X-2: Towards beam quality preservation in a plasma booster". In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 909 (2018), pp. 80–83. ISSN: 0168-9002. DOI: `https://doi.org/10.1016/j.nima.2018.02.063`.

[199] R. D'Arcy et al. "Recovery time of a plasma-wakefield accelerator". In: *Nature* 603.7899 (Mar. 2022), pp. 58–62.

[200] R. D'Arcy et al. "Tunable Plasma-Based Energy Dechirper". In: *Phys. Rev. Lett.* 122 (3 Jan. 2019), p. 034801. DOI: `10.1103/PhysRevLett.122.034801`.

[201] P. Scherkl et al. "Plasma photonic spatiotemporal synchronization of relativistic electron and laser beams". In: *Phys. Rev. Accel. Beams* 25 (5 May 2022), p. 052803. DOI: `10.1103/PhysRevAccelBeams.25.052803`.

[202] M. Litos et al. "High-efficiency acceleration of an electron beam in a plasma wakefield accelerator". In: *Nature* 515.7525 (Nov. 2014), pp. 92–95. ISSN: 1476-4687. DOI: `10.1038/nature13882`.

[203] M. Litos et al. "9 GeV energy gain in a beam-driven plasma wakefield accelerator". In: *Plasma Physics and Controlled Fusion* 58.3, 034017 (Mar. 2016), p. 034017. DOI: `10.1088/0741-3335/58/3/034017`. arXiv: `1511.06743 [physics.plasm-ph]`.

[204] L. Boulton et al. *Longitudinally resolved measurement of energy-transfer efficiency in a plasma-wakefield accelerator*. 2022. arXiv: `2209.06690 [physics.acc-ph]`.

[205] E. Oz et al. "Optical Diagnostics for Plasma Wakefield Accelerators". In: *AIP Conference Proceedings* 737.1 (2004), pp. 708–714. DOI: `10.1063/1.1842612`. eprint: `https://aip.scitation.org/doi/pdf/10.1063/1.1842612`.

[206] E. Adli et al. "Long-range attraction of an ultrarelativistic electron beam by a column of neutral plasma". In: *New Journal of Physics* 18.10 (Oct. 2016), p. 103013. DOI: `10.1088/1367-2630/18/10/103013`.

[207] M. W. Guetg et al. "Optimization of free electron laser performance by dispersion-based beam-tilt correction". In: *Phys. Rev. ST Accel. Beams* 18 (3 Mar. 2015), p. 030701. DOI: `10.1103/PhysRevSTAB.18.030701`.

[208] A. A. Sahai et al. "Long Term Evolution of Plasma Wakefields". In: *Proc. NA-PAC'13* (Pasadena, CA, USA, Sep.-Oct. 2013). JACoW Publishing, Geneva, Switzerland, pp. 90–92. URL: `https://jacow.org/PAC2013/papers/MOPAC10.pdf`.

[209] J. Rosenzweig et al. "Towards a plasma wake-field acceleration-based linear collider1Work supported by U.S. Dept. of Energy grants DE-FG03-93ER0796, and the Alfred P. Sloan Foundation grant BR-3225.1". In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 410.3 (1998), pp. 532–543. ISSN: 0168-9002. DOI: `https://doi.org/10.1016/S0168-9002(98)00186-7`.

[210] S. Pei et al. "Conceptual Design of the Drive Beam for a PWFA-LC". In: *Proc. PAC'09* (Vancouver, Canada, May 2009). Particle Accelerator Conference 23. JACoW Publishing, Geneva, Switzerland, pp. 2682–2684. URL: `https://jacow.org/PAC2009/papers/WE6PFP079.pdf`.

[211] C. A. Lindstrøm. "Staging of plasma-wakefield accelerators". In: *Phys. Rev. Accel. Beams* 24 (1 Jan. 2021), p. 014801. DOI: `10.1103/PhysRevAccelBeams.24.014801`.

[212] J.-P. Delahaye et al. "A Beam Driven Plasma-wakefield Linear Collider from Higgs Factory to Multi-TeV". In: *Proc. IPAC'14* (Dresden, Germany, Jun. 2014). JACoW Publishing, Geneva, Switzerland, pp. 3791–3793. DOI: `10.18429/JACoW-IPAC2014-THPRI013`.

[213] C. B. Schroeder et al. "Physics considerations for laser-plasma linear colliders". In: *Phys. Rev. ST Accel. Beams* 13 (10 Oct. 2010), p. 101301. DOI: `10.1103/PhysRevSTAB.13.101301`.

[214] S. Steinke et al. “Multistage coupling of independent laser-plasma accelerators”. In: *Nature* 530.7589 (Feb. 2016), pp. 190–193. ISSN: 1476-4687. DOI: `10.1038/nature16525`.

[215] A. Deng et al. “Generation and acceleration of electron bunches from a plasma photocathode”. In: *Nature Physics* 15.11 (Nov. 2019), pp. 1156–1160. ISSN: 1745-2481. DOI: `10.1038/s41567-019-0610-9`.

[216] D. Ullmann et al. “All-optical density downramp injection in electron-driven plasma wakefield accelerators”. In: *Phys. Rev. Research* 3 (4 Dec. 2021), p. 043163. DOI: `10.1103/PhysRevResearch.3.043163`.

[217] J. Wood, L. Boulton, et al. “Brightness-Boosted Electron Bunches from a Plasma Accelerator”. 2022.

[218] S. Di Mitri. “Bunch Length Compressors”. In: *CERN Yellow Rep. School Proc.* 1 (2018). Ed. by R. Bayley, p. 363. DOI: `10.23730/CYRSP-2018-001.363`.

[219] N. Vafaei-Najafabadi et al. “Beam Loading by Distributed Injection of Electrons in a Plasma Wakefield Accelerator”. In: *Phys. Rev. Lett.* 112 (2 Jan. 2014), p. 025001. DOI: `10.1103/PhysRevLett.112.025001`.

[220] T.-Y. Chien et al. “Spatially Localized Self-Injection of Electrons in a Self-Modulated Laser-Wakefield Accelerator by Using a Laser-Induced Transient Density Ramp”. In: *Phys. Rev. Lett.* 94 (11 Mar. 2005), p. 115003. DOI: `10.1103/PhysRevLett.94.115003`.

[221] S. Fourmaux et al. “Quasi-monoenergetic electron beams production in a sharp density transition”. In: *Applied Physics Letters* 101.11 (2012), p. 111106. DOI: `10.1063/1.4752114`. eprint: `https://doi.org/10.1063/1.4752114`.

[222] K. V. Lotov. “Fine wakefield structure in the blowout regime of plasma wakefield accelerators”. In: *Phys. Rev. ST Accel. Beams* 6 (6 June 2003), p. 061301. DOI: `10.1103/PhysRevSTAB.6.061301`.

[223] T. Katsouleas and W. B. Mori. “Wave-Breaking Amplitude of Relativistic Oscillations in a Thermal Plasma”. In: *Phys. Rev. Lett.* 61 (1 July 1988), pp. 90–93. DOI: `10.1103/PhysRevLett.61.90`.

[224] J. B. Rosenzweig. “Trapping, thermal effects, and wave breaking in the nonlinear plasma wake-field accelerator”. In: *Phys. Rev. A* 38 (7 Oct. 1988), pp. 3634–3642. DOI: `10.1103/PhysRevA.38.3634`.

[225] C. B. Schroeder, E. Esarey, and B. A. Shadwick. "Warm wave breaking of nonlinear plasma waves with arbitrary phase velocities". In: *Phys. Rev. E* 72 (5 Nov. 2005), p. 055401. DOI: `10.1103/PhysRevE.72.055401`.

[226] C. B. Schroeder and E. Esarey. "Relativistic warm plasma theory of nonlinear laser-driven electron plasma waves". In: *Phys. Rev. E* 81 (5 May 2010), p. 056403. DOI: `10.1103/PhysRevE.81.056403`.

[227] N. Barov et al. "Energy loss of a high-charge bunched electron beam in plasma: Analysis". In: *Physical Review Accelerators and Beams* 7.6, 061301 (June 2004), p. 061301. DOI: `10.1103/PhysRevSTAB.7.061301`. arXiv: `physics/0205007 [physics.acc-ph]`.

[228] A. Ferran Pousa, A. Martinez de la Ossa, and R. W. Assmann. "Intrinsic energy spread and bunch length growth in plasma-based accelerators due to betatron motion". In: *Scientific Reports* 9.1 (Nov. 2019), p. 17690. ISSN: 2045-2322. DOI: `10.1038/s41598-019-53887-8`.

[229] N. Kirby et al. "Emittance growth from Multiple Coulomb Scattering in a plasma wakefield accelerator". In: *2007 IEEE Particle Accelerator Conference (PAC)*. 2007, pp. 3097–3099. DOI: `10.1109/PAC.2007.4440680`.

[230] C. B. Schroeder et al. "Control of focusing forces and emittances in plasma-based accelerators using near-hollow plasma channels". In: *Physics of Plasmas* 20.8 (2013), p. 080701. DOI: `10.1063/1.4817799`. eprint: `https://doi.org/10.1063/1.4817799`.

[231] C. Birdsall. "Particle-in-cell charged-particle simulations, plus Monte Carlo collisions with neutral atoms, PIC-MCC". In: *IEEE Transactions on Plasma Science* 19.2 (1991), pp. 65–85. DOI: `10.1109/27.106800`.

[232] Y. Zhao et al. "Modeling of emittance growth due to Coulomb collisions in plasma-based accelerators". In: *Physics of Plasmas* 27.11 (2020), p. 113105. DOI: `10.1063/5.0023776`. eprint: `https://doi.org/10.1063/5.0023776`.

[233] S. G. Anderson et al. "Space-charge effects in high brightness electron beam emittance measurements". In: *Phys. Rev. ST Accel. Beams* 5 (1 Jan. 2002), p. 014201. DOI: `10.1103/PhysRevSTAB.5.014201`.

[234] A. Knetsch et al. *Downramp-assisted underdense photocathode electron bunch generation in plasma wakefield accelerators*. Dec. 2014. arXiv: `1412.4844 [physics.acc-ph]`.

[235] N. Vafaei-Najafabadi et al. "Beam Loading by Distributed Injection of Electrons in a Plasma Wakefield Accelerator". In: *Phys. Rev. Lett.* 112 (2 Jan. 2014), p. 025001. DOI: `10.1103/PhysRevLett.112.025001`.

[236] A. Ferran Pousa et al. "Compact Multistage Plasma-Based Accelerator Design for Correlated Energy Spread Compensation". In: *Phys. Rev. Lett.* 123 (5 July 2019), p. 054801. DOI: `10.1103/PhysRevLett.123.054801`.

[237] G. G. Manahan et al. "Single-stage plasma-based correlated energy spread compensation for ultrahigh 6D brightness electron beams". In: *Nature Communications* 8.1 (June 2017), p. 15705. ISSN: 2041-1723. DOI: `10.1038/ncomms15705`.

[238] C. Emma et al. "Terawatt attosecond x-ray source driven by a plasma accelerator". In: *APL Photonics* 6.7 (2021), p. 076107. DOI: `10.1063/5.0050693`. eprint: `https://doi.org/10.1063/5.0050693`.

[239] C. Emma et al. "PAX: A Plasma-Driven Attosecond X-Ray Source". In: *Proc. IPAC'21* (Campinas, SP, Brazil). International Particle Accelerator Conference 12. JACoW Publishing, Geneva, Switzerland, Aug. 2021, WEPAB072, pp. 2755–2758. ISBN: 978-3-95450-214-1. DOI: `10.18429/JACoW-IPAC2021-WEPAB072`.

[240] C. Joshi et al. "Plasma wakefield acceleration experiments at FACET II". In: *Plasma Physics and Controlled Fusion* 60.3 (Jan. 2018), p. 034001. DOI: `10.1088/1361-6587/aaa2e3`.

[241] A. F. Habib et al. *Ultrahigh brightness beams from plasma photoguns*. 2021. arXiv: `2111.01502 [physics.acc-ph]`.

[242] O. Lundh et al. "Few femtosecond, few kiloampere electron bunch produced by a laser–plasma accelerator". In: *Nature Physics* 7.3 (Mar. 2011), pp. 219–222. ISSN: 1745-2481. DOI: `10.1038/nphys1872`.

[243] J. P. Couperus et al. "Demonstration of a beam loaded nanocoulomb-class laser wakefield accelerator". In: *Nature Communications* 8.1 (Sept. 2017), p. 487. ISSN: 2041-1723. DOI: `10.1038/s41467-017-00592-7`.

[244] B. Hidding et al. "Monoenergetic Energy Doubling in a Hybrid Laser-Plasma Wakefield Accelerator". In: *Phys. Rev. Lett.* 104 (19 May 2010), p. 195002. DOI: `10.1103/PhysRevLett.104.195002`.

[245] B. Hidding et al. "Fundamentals and Applications of Hybrid LWFA-PWFA". In: *Applied Sciences* 9 (June 2019), p. 2626. DOI: `10.3390/app9132626`.

[246] T. Kurz et al. "Demonstration of a compact plasma accelerator powered by laser-accelerated electron beams". In: *Nature Communications* 12.1 (May 2021), p. 2895. ISSN: 2041-1723. DOI: `10.1038/s41467-021-23000-7`.

[247] T. J. Mehrling et al. "Mitigation of the Hose Instability in Plasma-Wakefield Accelerators". In: *Phys. Rev. Lett.* 118 (17 Apr. 2017), p. 174801. DOI: `10.1103/PhysRevLett.118.174801`.

[248] T. Heinemann et al. "Investigating the Key Parameters of a Staged Laser- and Particle Driven Plasma Wakefield Accelerator Experiment". In: *Proc. of International Particle Accelerator Conference (IPAC'17), Copenhagen, Denmark, 14â19 May, 2017* (Copenhagen, Denmark). International Particle Accelerator Conference 8. Geneva, Switzerland: JACoW, May 2017, pp. 1703–1706. ISBN: 978-3-95450-182-3. DOI: `https://doi.org/10.18429/JACoW-IPAC2017-TUPIK010`.