EXPERIMENTAL STUDY OF LONG TIMESCALE PLASMA WAKEFIELD EVOLUTION

INVESTIGATIONS TOWARD HIGH-REPETITION-RATE OPERATION OF PLASMA-BASED ACCELERATORS

JAMES CHAPPELL

Supervised by
PROF. MATTHEW WING

High Energy Physics PhD

Department of Physics and Astronomy
Faculty of Mathematical and Physics Sciences
University College London

July 2021

Submitted to UCL in fulfilment of the requirements for the award of the degree of
Doctor of Philosophy

DECLARATION

I, James Chappell, confirm that the work presented in this thesis is my own. Where information has been derived from other sources, I confirm that this has been indicated in the thesis.

July 2021

James Chappell

ABSTRACT

Over the last century, particle accelerators have significantly enhanced our understanding of the fundamental nature of the universe and, as such, have become ubiquitous tools within society. The continued search for explanations of phenomena beyond our current best models motivates the proposal of ever-larger and more expensive particle accelerators. However, the economic impact of building and operating such machines potentially casts doubt on their delivery. The exploration of alternative acceleration concepts that can potentially provide a reduction in the size and cost of these machines has therefore seen a significant growth in interest over the past few decades.

One such acceleration concept is that of plasma-based wakefield acceleration where high-intensity particle or laser beams strongly perturb a plasma and, in doing so, generate fields in their wake that can be used to accelerate charged particles. This is an attractive concept as plasmas can support accelerating fields multiple orders of magnitude larger than those provided in conventional accelerators, potentially reducing the accelerating length by similar scales. While rapid progress has recently been made with regards to high-quality acceleration in plasma, comparatively little effort has been applied to the study of the frequency at which this is possible.

The measurements presented in this thesis study the evolution of a plasma as the energy imparted into it via the wakefield acceleration process dissipates. Such detailed measurements allow determination of the fundamental mechanisms that will limit high-repetition-rate operation of plasma-based accelerators. As such, these measurements represent a significant first step towards the demonstration of plasma-based acceleration at frequencies comparable to those provided by state-of-the-art conventional accelerators, helping to define the achievable luminosity of future facilities that rely on such technology.

IMPACT STATEMENT

Thanks to recent experimental advances there is little doubt that high-quality acceleration of electron bunches at large gradients will be achieved in plasma-based accelerators in the near future. However, a question that is often asked both within the plasma accelerator community and by the wider community is:

"What is the achievable operating rate of such accelerator devices?".

At present, high-power (joule-level) plasma-based accelerators have only ever operated at tens-of-Hz frequencies at most, and often far below this level. The prospect of wide-scale adoption of plasma-based accelerator technologies will be significantly advanced by the demonstration of operation at repetition-rates on-par with, or beyond, current state-of-the-art conventional accelerator facilities.

The measurements presented within this thesis attempt to directly address the above question by studying the evolution of the plasma following the wakefield acceleration process. This allows the identification of physical mechanisms that can prevent repeatable, high-quality acceleration and study of their dependencies on crucial experimental variables, thus allowing the scaling of these results to parameter ranges of interest for future facilities and applications. Such facilities include plasma-based particle colliders at their most extreme, or compact synchrotron radiation sources and medical and defence applications on smaller scales. The results presented within this thesis therefore have broad reach and applicability and represent an important first step towards the realisation of high-average-power ($>$ kW) plasma-based accelerator facilities.

ACKNOWLEDGEMENTS

The work within this thesis was only achievable thanks to the input, support and collaboration of a large number of people. First and foremost, I would like to thank Matthew for his encouragement, support and supervision of the research I performed throughout my PhD. This work was only possible thanks to his open-mindedness and willingness to let me explore a range of opportunities. A special thanks also to Richard for his supervision and encouragement of my work at FLASHForward. Thanks to my friends at UCL, CERN and DESY (and beyond) for making the PhD experience so enjoyable (for the most part). Finally, thank you to my family for their continuous support.

CONTENTS

LIST OF TABLES

LIST OF SYMBOLS

The following are reference lists of constants and variables that are regularly used within this thesis. The values of the fundamental physics constants are taken from the U.S. National Institute of Standards and Technology (NIST) database:

https://physics.nist.gov/cuu/pdf/wall_2018.pdf

Symbol	Quantity	Value	Unit
e	Elementary charge	1.60×10^{-19}	C
m_e	Electron mass	9.11×10^{-31}	kg
$m_e c^2$	Electron mass energy eq.	0.511	MeV
m_p	Proton mass	1.67×10^{-27}	kg
$m_p c^2$	Proton mass energy eq.	938	MeV
ε_0	Free space permittivity	8.85×10^{-12}	Fm^{-1}
μ_0	Free space permeability	$4\pi \times 10^{-7}$	Hm^{-1}
c	Vacuum speed of light	299 792 458	ms^{-1}
k_B	Boltzmann constant	1.38×10^{-23}	JK^{-1}

Table 1: List of constants and their values.

Symbol	Quantity	Unit
q	Electric charge	C
ρ	Charge density	Cm^{-3}
I	Current	A
$\vec{J}$	Current density	Am^{-2}
$\vec{E}$	Electric field	Vm^{-1}
$\vec{B}$	Magnetic field	T
ϕ	Electric scalar potential	V
$\vec{A}$	Magnetic vector potential	Tm
β	Relativistic normalised velocity	-
γ	Relativistic Lorentz factor	-
λ_D	Debye length	m
n_e	Electron number density	cm^{-3}
n_i	Ion number density	cm^{-3}
n_a	Atomic number density	cm^{-3}
$\omega_{p(e)}$	Plasma (electron) frequency	$rads^{-1}$
k_p	Plasma wavenumber	m^{-1}
λ_p	Plasma wavelength	m
ω_β	Betatron frequency	$rads^{-1}$
k_β	Betatron wavenumber	m^{-1}
T_e	Plasma electron temperature	K
T_i	Plasma ion temperature	K
T	Plasma temperature (1T)	K
κ	Thermal conductivity	$Wm^{-1}K^{-1}$
ν	Collision frequency	s^{-1}
λ_0	Laser pulse central wavelength	m
ω_0	Laser pulse central angular frequency	$rads^{-1}$
$\Delta\tau$	Relative timing/delay	s
ε_N	Normalised transverse emittance	m·rad
σ	Cross-section	m^{-2}
Z	Ionisation fraction	-

Table 2: List of variables regularly used within this thesis.

1

INTRODUCTION

1.1 PARTICLE ACCELERATORS

Particle accelerators are an ubiquitous tool within society with applications ranging from low-energy electron beam radiation sterilisation, particle beam radiotherapy and semiconductor ion implantation, to high-energy ultra-intense light sources and particle colliders. At their most extreme, high-energy particle colliders represent the most powerful microscopes in existence and have enabled study of the fundamental nature of the universe.

The advent of particle accelerators was conceived at the turn of the 20th century with the discovery of the electron by J. J. Thomson [1] via investigation of cathode rays, and measurements of the structure of the atomic nucleus by E. Rutherford [2] using alpha particles generated by nuclear decay. The scientific advances that could be made possible by the use of higher-energy particles in such experiments was immediately obvious to Rutherford, who later remarked in his Address of the President to the Royal Society in 1927 [3] that

> *"It has long been my ambition to have available for study a copious supply of atoms and electrons which have an individual energy far transcending that of the α and β-particles from radioactive bodies."*

Over the last century, Rutherford's ambition was achieved thanks to the development of cyclotrons, synchrotrons and linear accelerators (linacs), amongst other particle accelerator technologies, enabling probing of the fundamental structure of particles and, by extension, the universe. Early discoveries continued with the neutron by J. Chadwick [4] and the positron by C. D. Anderson [5] in 1932, while, most recently, the existence of the Higgs boson [6, 7] was confirmed at the Large Hadron Collider (LHC) [8] at the European Organisation for Nuclear Research (CERN) in 2012. Development of theoretical descriptions accompanied such discoveries, culminating in the derivation of the 'Standard Model' (SM) [9] of particle physics in the latter half of the 20th century. While the SM does not provide a complete description of the universe, it has provided predictions of the existence of fundamental particles and their interactions that are continuously tested at the largest particle accelerator and collider facilities.

Yet it is the knowledge of the incompleteness of the SM in its current form that motivates the pursuit of deviations from its predictions; for example, the existence of 'dark matter' — indirectly confirmed via astrophysical observations of the difference between gravitational and luminous mass in galaxies [10], measurements of the rotation speeds of stars within galaxies [11], and the effects of its gravitational lensing [12], to name but a few examples — is not predicted by the SM. Similarly the non-zero neutrino mass [13], baryon asymmetry [14],

and the observation of the accelerating expansion of the universe [15, 16] are all beyond the Standard Model (BSM). Within the search for possible explanations for these observations, particle accelerators represent a crucial component — at present, a large number of BSM searches are already being undertaken at the LHC [17] and will continue over its entire life-cycle. However, the energy-reach, and hence discovery potential, of the LHC in its current form is limited and plans for future, more powerful facilities must inevitably be made.

1.1.1 *Acceleration of charged particles*

Charged particles are accelerated via their interaction with electromagnetic fields. This was first described at the end of the 19$^{\text{th}}$ century by H. Lorentz [18], via the (now-named) Lorentz force

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

where $\vec{F}$ is the force exerted on the particle, $\vec{E}$ and $\vec{B}$ are the electric and magnetic fields respectively, $\vec{v}$ is the velocity of the particle and q is its electric charge. This equation implies that the energy of a particle cannot be increased by the application of a magnetic field — the resulting force would always be perpendicular to its velocity thanks to the form of the cross-product. Therefore, to accelerate a charged particle along its direction of propagation, and hence increase its energy, an electric field must be used. The total energy gained by a charged particle within an electric field is given by

$$\Delta E = q \int_0^L E_z(s) ds, \tag{1.1.2}$$

where E_z is the electric field along the direction of propagation (the longitudinal field), s is the coordinate along the accelerator and L is the total acceleration length. Therefore, at its simplest level, to increase the energy of particles produced by an accelerator, either the length of the accelerator or the magnitude of the electric field must be increased.

1.1.2 *Limitations of conventional accelerators*

Conventional accelerator facilities exist in two forms; linear and circular. Within a linear accelerator, charged particles propagate through a series of radio-frequency (RF) cavities where the electric field is varied synchronously with the propagation of the particle bunch such that it is continuously accelerated to high energies. Acceleration in this way was first proposed by G. Ising [19] in 1924 and was experimentally demonstrated a few years later by R. Wideröe [20]. The maximum electric field that can be supported within such RF cavities is limited by 'breakdown' [21] of the cavity walls, where electrons are ionised from the surface of the cavity, damaging its structure. This limits the accelerating gradient to a few tens of MVm^{-1}, even in ultra-high vacuum conditions. This effect can be limited by cooling the

cavity [22] or via the use of dielectric structures that are more resistant to breakdown [23], but this still remains a fundamental limitation. However, in principle, a linear accelerator of this form can be made as long as necessary to produce particle bunches of the required energy and are practically limited by the cost of building and running them. The first large-scale linear collider was the Stanford Linear Collider [24], with electron and positron bunches provided by a 3 km-long linear accelerator, achieving a centre-of-mass collision energy of 91 GeV. Future proposals for linear colliders [25, 26] have lengths on the order of tens of km such that higher collision energies (hundreds of GeV) can be achieved.

An alternative to the ever-increasing length of linear accelerators is to recirculate the particle bunch such that the same accelerating cavities can be used many times — this is the principle behind circular accelerators. In this form, the accelerating gradient can be kept modest [$\mathcal{O}(10\,\mathrm{MVm}^{-1})$] — far below the 'breakdown' limit — thus reducing the required power and improving stability, while the energy of particles can simply be increased by increasing the number of passes through the accelerating cavities. The circular trajectory is maintained by dipole magnets that provide a constant magnetic field that bends the trajectory of charged particles, with their radius of curvature, ρ, given by

$$\frac{1}{\rho} = \frac{qB}{p}, \tag{1.1.3}$$

where p is momentum of the particle and B is the magnetic field strength. Therefore, the maximum achievable energy within a circular collider is mainly limited by the strength of available magnets. Superconducting magnets offer the highest magnetic fields, often on the order of 10 T, with those used at the aforementioned LHC consistently operating at 8.3 T [8]. This enables acceleration of protons up to energies of 7 TeV within the 27 km-long circular accelerator.

Circular accelerators therefore seem to provide ideal conditions for the acceleration of charged particles to high energies. However, as the ultra-relativistic charged particles are accelerated transversely (bent) through the dipole magnets, they emit significant amounts of synchrotron radiation — a form of highly-directed, wide-spectrum and often high-power electromagnetic radiation. For each turn around the ring, the amount of energy radiated by charged particles scales as [27]

$$E_{SR} \propto \frac{\gamma^4}{\rho}, \tag{1.1.4}$$

where $\gamma = E/mc^2$ is the relativistic Lorentz factor. Such a strong scaling with the relativistic factor, γ^4, is extremely unfavourable for electrons and positrons as a result of their small mass and hence severely limits their attainable energies in a circular accelerator. For example, the predecessor of the LHC, the Large Electron–Positron Collider [28], shared the same footprint as the current LHC but the particle bunches were limited to energies of 104.5 GeV due to the large synchrotron radiation losses — a factor of 67 lower in energy than that of LHC proton bunches. One proposal for the next generation of circular colliders, the Future Circular Collider (FCC) [29], has a circumference of 100 km and would use magnets of strength $16 - 20$ T

to achieve centre-of-mass collision energies of 100 TeV in proton–proton operation. However, the proposed electron–positron collider in the same footprint (FCC-ee) [30] would be limited to bunch energies of 175 GeV. It is therefore clear that future electron–positron colliders must be linear to reach the TeV-scale.

1.2 PLASMA-BASED PARTICLE ACCELERATION

Using conventional accelerator technology, linear accelerators would need to have accelerating lengths of the order of tens of km to reach the TeV-scale — a potentially unrealistic level of expenditure to both build and operate. As discussed in the previous section, this is due to the occurrence of 'breakdown' within the RF cavities restricting the maximum accelerating gradient that can be sustained. This therefore invites the proposal to change the medium within which particles are accelerated — if 'breakdown' is the limiting factor, move to a medium where this is not possible; plasma.

A plasma represents a fully broken-down state of matter and is a quasi-neutral mixture of ions and electrons that exhibits collective behaviour. This means that at a macroscopic level, ions and electrons within the plasma are distributed such that it remains electrically-neutral. However, on a microscopic level, the thermal motion of the constituents can lead to local regions of charge separation that give rise to electric fields. The approximate scale over which quasi-neutrality is maintained within a plasma can be defined by its Debye length,

$$\lambda_D = \sqrt{\frac{\varepsilon_0 k_B T_e}{n_e e^2}}, \tag{1.2.1}$$

where e is the charge of the electron, ε_0 is the permittivity of free space, k_B is the Boltzmann constant, and T_e and n_e are the temperature and density of electrons within the plasma respectively. The high density of electrically-charged particles within a plasma, typically of order $10^{14} - 10^{18}\,\mathrm{cm}^{-3}$ in plasma accelerators, means electric fields of amplitude far larger than those supported in the near-vacuum of RF cavities can be maintained within its structure provided that an appropriate level of charge separation can be achieved.

1.2.1 *Laser-driven plasma wakefield acceleration*

The first common technique for generating significant charge separation within a plasma was proposed by Tajima and Dawson in 1979 [31], who noted that the ponderomotive force from an intense laser pulse would act (primarily) on electrons within a plasma and excite electron density perturbations in its wake as it propagated. The induced charge separation leads to the generation of large electromagnetic fields within the plasma, with accelerating gradients scaling as

$$E \approx 100\sqrt{n_e\,[10^{18}\,\mathrm{cm}^{-3}]}\;[\mathrm{GVm}^{-1}]. \tag{1.2.2}$$

Clearly, a plasma of density $10^{18}\,\mathrm{cm}^{-3}$ is therefore able to support accelerating gradients of order $100\,\mathrm{GVm}^{-1}$, more than three orders of magnitude greater than typically achieved in traditional RF cavity-based acceleration, thus significantly reducing the required accelerator length. This laser-driven technique would become known as Laser Wakefield Acceleration (LWFA).

LWFA requires the use of a laser pulse that is short with respect to the size of the density perturbation, defined by the plasma-wave wavelength, λ_p, which is of order 30 µm (~ 100 fs) at this density. At the time of publication of the original proposal, sufficiently short, intense laser pulses did not exist to experimentally verify Tajima and Dawson's predictions. Alternative schemes such as the 'beat-wave' [32] were used to demonstrate the laser-driven excitation of wakefields within a plasma via the combination of two co-propagating longer pulses of differing frequency [33, 34]. The development of chirped-pulse amplification [35] in the 1980s enabled first demonstrations of the short-pulse technique in the 1990s [36–38]. Interest in LWFA was intensified in 2004 with the publication of three separate demonstrations of LWFA where electrons bunches were accelerated to 100 MeV-scales in mm-scale plasmas with few-percent energy spreads [39–41] — the so-called 'dream beam' results thanks to their feature on the cover of the scientific journal Nature alongside this headline.

However, a fundamental limitation of LWFA is introduced by the refractive index of the plasma. A laser pulse propagates through plasma at a velocity slightly lower than that in vacuum and hence the wakefield it excites does too. Electrons that are being accelerated within the wakefield quickly reach ultra-relativistic energies and hence propagate through the plasma at velocities close to the speed of light, exceeding the velocity of the wakefield. Over extended propagation distances, electrons outrun the wakefield, transitioning from its accelerating phase to its decelerating phase within the so-called dephasing length [42],

$$L_d \approx \frac{\lambda_p^3}{\lambda_0^2}, \tag{1.2.3}$$

where λ_0 is the wavelength of the laser pulse. This therefore limits the distance over which electrons can be effectively accelerated by the wakefield. The dephasing length can be extended by operating at lower plasma densities (increasing λ_p), and alternative focusing schemes have been proposed to overcome this limit [43] but have yet to be experimentally demonstrated.

1.2.2 *Beam-driven plasma wakefield acceleration*

An alternative technique for inducing plasma wakefields was proposed a few years after the seminal work of Tajima and Dawon and suggested the use of a compact bunch of ultra-relativistic electrons, with their space-charge field providing the perturbing force to generate charge separation in the plasma [44]. This method is known as beam-driven plasma wakefield acceleration (PWFA) [45]. In PWFA, the dephasing limitation is avoided through the use of an ultra-relativistic drive bunch that propagates through the plasma at the same velocity as

electrons being accelerated within the wakefield. Furthermore, the wakefield can be driven within the plasma until all the energy of the drive bunch is extracted, allowing large total energy-transfer efficiencies to be achieved.

First beam-driven PWFA experiments were performed at the Argonne National Laboratory in 1988, with the demonstration of a two-bunch acceleration scheme [46]. In this scheme, a leading "drive" bunch transfers energy to the plasma and excites the wakefield, while a second, lower charge, "witness" bunch follows at the appropriate phase to extract energy from the wakefield and be accelerated. This is the approach often taken within modern PWFA experiments.

In the early 1990s, the non-linear "blowout" regime of PWFA [47, 48] was recognised as having suitable characteristics for high-quality electron acceleration although, again, suitably intense high-charge, highly-compressed bunches that could operate in such a regime were not available until experiments performed at the Final Focus Test Beam (FFTB) facility [49] towards the end of the decade. Initial experiments using a highly-compressed bunch of 30 GeV electrons to drive wakefields in a low density plasma ($n_e \sim 10^{14}\,\mathrm{cm}^{-3}$) demonstrated the acceleration of positrons [50] and multi-GeV acceleration of electrons [51], while the energy doubling of a small fraction of a 42 GeV electron beam in a plasma of length 85 cm was later achieved at higher plasma densities ($n_e \sim 10^{17}\,\mathrm{cm}^{-3}$) [52]. These initial experiments demonstrated that large accelerating gradients could be produced within PWFA schemes, offering the potential for beam-driven plasma accelerators to provide short "afterburner" sections appended to conventional linear accelerators to boost their energies before collision [53] or as the basis for future linear colliders [54]. More recent PWFA experiments have focused on improving the quality and stability of acceleration using upgraded facilities [55, 56]. Of particular relevance to this thesis is the Future-oriented Wakefield Accelerator Research and Development at FLASH (FLASHForward) facility [57] where electron bunches are provided by a conventional accelerator of exceptional stability [58], enabling detailed measurements of the wakefield process and its dependencies to be performed. Such measurements are discussed in detail in Chapters 4 – 6, 8 and 9.

The beam-driven plasma accelerator scheme is not limited to only electron bunches, with proton bunches also providing suitable drive bunches [59]. The use of high-energy proton bunches available at facilities such as CERN could potentially allow the acceleration of witness bunches to TeV-scales in a single extended plasma stage [60]. The Advanced Wakefield Experiment (AWAKE) at CERN [61] was devised to investigate proton-driven plasma wakefield acceleration via the use of seeded self-modulation [62] to transform long proton bunches into effective wakefield drive bunches. Results from this experiment are discussed in detail in Chapters 3 and 7.

1.3 HIGH-REPETITION-RATE OPERATION OF PLASMA-BASED ACCELERATORS

Rapid experimental and theoretical progress has been made within the field of plasma-based accelerators over the past two decades. Advances in the use of non-uniform transverse

guiding structures have enabled the propagation of high-power laser pulses within plasma over extended distances such that injected electrons reach GeV-scale energies in plasmas of lengths on the order of a few centimetres [63–65], with the current state-of-the-art demonstrating an energy gain of 7.8 GeV over 20 cm — an average accelerating gradient of almost $40\,\mathrm{GVm}^{-1}$ [66]. Improvements in control and stability of the drive and witness beams within beam-driven plasma accelerators has similarly enabled significant progress to be made in the quality and efficiency of acceleration [67], with the current state-of-the-art demonstrating simultaneous per-mille-level energy-spread preservation and instantaneous energy-transfer efficiencies in excess of 40% at accelerating gradients of $1.3\,\mathrm{GVm}^{-1}$ [68]. It is widely expected that preservation of the transverse normalised emittance, ε_N, of the witness bunch — a measure of its quality — will be achieved in the near future alongside an increase in total efficiency. Doing so will represent the ultimate demonstration of stable, high-quality acceleration in a plasma-based accelerator.

However, while considerable effort is being applied to improving the quality of plasma-accelerated bunches, this is not the only crucial parameter that contributes towards the realisation of plasma-based accelerators as facilities, be it future particle colliders, next-generation light sources or lower-energy accelerators. Using particle colliders as an example, not only is exquisite beam quality required to enable a large number of collisions to occur per bunch-crossing, but the frequency of bunch-crossings must also be maximised. These effects are summarised by the luminosity of the machine, a measure of the event-rate per cross-section, given by

$$\mathcal{L} = H\frac{N^2 f}{4\pi\sigma_x\sigma_y}, \tag{1.3.1}$$

where H is a small geometrical factor from bunch tilts and pinching effects, N is the number of particles per bunch, f is the frequency of collisions, and σ_x and σ_y are the transverse bunch sizes in the horizontal and vertical planes respectively. Only a small number of particles within each bunch actually undergoes a collision within each crossing and hence, within a circular collider such as the LHC, bunches are circulated many times and undergo many collisions before being dumped. This enables the LHC to operate at very high collision rates with approximately 6×10^8 collisions per second and luminosities in excess of $10^{34}\,\mathrm{cm}^{-2}\mathrm{s}^{-1}$ [8]. Planned upgrades to the injectors [69] to improve the proton bunch quality and corresponding detector upgrades are expected to increase this by a further factor of 10 in the future [70].

In contrast, within a linear collider, the natural scheme for plasma accelerators, bunches cannot be re-circulated and only undergo a single bunch-crossing. Therefore, to maximise the number of collisions and hence luminosity, both the frequency of bunch-crossings and the quality of the colliding bunches must be maximised. Internal injection schemes that allow the generation of ultra-low emittance bunches within the plasma have been proposed [71–73] and their experimental implementation is on-going [74–76] while plasma acceleration schemes exist that will allow preservation of the quality of the bunch [47]. It therefore seems

feasible that high-gradient ($> \mathrm{GVm}^{-1}$) acceleration of ultra-low emittance ($\varepsilon_N \sim 10 - 100\,\mathrm{nm}$) beams will be possible in plasma accelerators over the coming years.

Comparatively little research has been performed to study the maximum achievable rate at which high-quality plasma acceleration can be achieved. This is mostly due to the lack of availability of suitable drive beams at repetition-rates where this could be studied. High-peak-power ($> 100\,\mathrm{TW}$) laser systems that provide laser pulses suitable for GeV-scale acceleration of electrons [77–79] typically operate at few- or sub-single-Hz repetition-rates while beam-driven facilities have previously been operated at 10 Hz repetition-rates and below [55, 61]. Straw-man high-average-power plasma-based facility proposals [54, 80, 81] typically assume operating repetition-rates of $\mathcal{O}(10\,\mathrm{kHz})$ and yet no detailed experimental or theoretical feasibility studies have yet been performed. However, such studies have been made possible at the FLASHForward experiment [57] thanks to the use of a super-conducting radio-frequency (SCRF) linac, FLASH [58], that can provide joule-level electron drive bunches for plasma wakefield experiments at repetition-rates of up to 3 MHz. However, the availability of suitable drive bunches at such a rate does not mean plasma acceleration is also possible at this frequency.

1.3.1 *Extended timescale evolution of plasma wakefields*

Unlike the near-vacuum conditions of a conventional RF cavity, a plasma does not represent a non-evolving accelerating medium. During the acceleration process, the plasma acts as an energy transformer, enabling the transfer of energy from the drive bunch to the trailing witness bunch via the fields driven within it. However, this process is not 100% efficient and hence a significant amount of energy remains within the plasma following the wakefield excitation and particle acceleration process. This energy must dissipate, with a number of evolutionary processes available for this within a plasma occurring over a range of timescales. These are summarised in Figure 1.3.1 and are briefly outlined as follows:

- **Electron plasma response** — the plasma wakefield is effectively a local electron density perturbation driven by, and co-propagating with, an intense drive beam. Plasma electrons respond either to the ponderomotive (laser-driven) or space-charge (beam-driven) force on the shortest timescales (fs – ps) and create the oscillating wakefield structure within the plasma. Plasma electrons are excited by the wakefield, gaining large amounts of energy, and undergo further oscillations following the passage of the beam. These initially coherent oscillations — the wakefield — gradually decohere due to the local plasma density gradients and the wakefield structure breaks down.

- **Ion plasma response** — ions within the plasma also respond to the large-amplitude wakefields and charge density perturbations, although over a longer timescale (ps – ns) due to their increased mass relative to that of electrons. The motion of ions can lead to the development of non-uniform density profiles and dominates mass flow within the plasma.

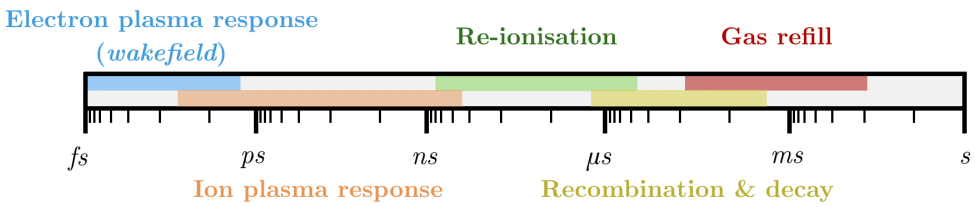

Figure 1.3.1: Representation of an approximate timeline for the mechanisms that dominate the evolution of a plasma wakefield from indications of measurements performed within the X-3 experiment at FLASHForward. Note, these mechanisms and timescales are for beam-driven wakefields at plasma densities of order $10^{16}\,\mathrm{cm}^{-3}$ and will vary for different conditions.

- **Re-ionisation** — plasma electrons and ions, excited by the wakefield, move through the plasma and undergo collisions with each other and recombined neutral species present within the plasma mixture. Such collisions can be of sufficiently high energy that ions and neutrals undergo ionisation, leading to a local increase in plasma density. Appropriate collisional cross-sections [82] indicate this re-ionisation effect is expected over few-ns – μs timescales.

- **Recombination and decay** — plasmas, perturbed by a wakefield or not, only exist for a limited period without additional heating mechanisms to sustain the high temperatures ($T \sim$ eV) required for significant ionisation fractions to be observed. Usually within plasma-based acceleration, the plasma is contained within a solid structure such as a capillary. The solid walls of the structure remain in contact with the plasma and provide one of the dominant sources of heat loss as heavy species (ions and neutrals) undergo collisions with the wall. This reduces the temperature of the plasma bulk and it recombines and decays, typically over μs – ms timescales.

- **Gas refill** — if the plasma is generated and contained within a capillary surrounded by near-vacuum, as is typically the case in plasma-based accelerator environments, the pressure gradient between the high-density plasma and vacuum drives its diffusion. Additionally, the ionisation process heats the plasma and the increased thermal velocity similarly increases diffusive effects. This results in a net loss of gas out of the containing structure which must be replaced before plasma of similar characteristics can be formed again. The refilling process typically occurs over ms timescales.

It is important to note that all of these mechanisms are strongly dependent on the plasma density and plasma species, among other variables, and are given in the context of beam-driven plasma wakefield experiments performed within the X-3 experiment at FLASHForward with argon plasma typically operated at densities of order $10^{16}\,\mathrm{cm}^{-3}$. This also does not represent an exhaustive list of possible dissipation mechanisms, but those that are typically observed to dominate within the FLASHForward experiment.

Clearly, the plasma will evolve significantly following the wakefield acceleration process and the dissipation of this evolution will likely dominate the ultimate repetition-rates achievable with plasma-based accelerators. To date, only two publications [83, 84] have experimentally observed evolution of the plasma wakefield beyond the initial electron plasma response by observing effects from the motion of plasma ions. However, neither represents a detailed study of these effects or their dependencies and hence cannot be used to predict how such mechanisms may limit achievable repetition-rates.

1.4 OUTLINE

This thesis details measurements undertaken by the author at two experimental facilities, the AWAKE experiment at CERN [61] and the FLASHForward experiment at DESY [57], to experimentally study the extended evolution of plasma wakefields.

Chapter 2 introduces a theoretical foundation describing plasma wakefield acceleration, demonstrating the advantageous characteristics of the non-linear wakefield regime.

The second part of the thesis introduces the facilities where measurements key to the results of this thesis were performed.

Chapter 3 introduces the AWAKE experimental facility and the concept of self-modulation within a plasma, crucial for the use of long proton bunches as effective wakefield drive bunches.

Chapter 4 introduces the FLASHForward experimental facility and the experiments being undertaken there. Particular care is taken to highlight experimental and diagnostic techniques that were developed to enable high-precision measurements of the wakefield process and its subsequent evolution.

Part three of the thesis is dedicated to introducing the beam-based plasma characterisation technique developed by the author to enable the study of plasma evolution over a range of timescales.

Chapter 5 motivates the use of short, high-current electron bunches as diagnostics of the plasma state, introduces a novel beam-based plasma characterisation technique, and describes its implementation within the FLASHForward experiment.

Chapter 6 discusses results obtained from application of the aforementioned beam-based characterisation technique to the study of the recombination and decay of various plasmas within sources suitable for beam-driven plasma wakefield acceleration.

The fourth and final part of this thesis is focused on the study of the plasma following its perturbation by the wakefield.

Chapter 7 uses external-injection of witness electrons into the wakefield driven by a self-modulated proton bunch to study the excitation and decay of a resonantly-driven plasma

wakefield on sub-nanosecond timescales. These measurements were performed in the context of the AWAKE experiment.

Chapter 8 applies the beam-based characterisation technique to the study of the plasma ion response to the wakefield excitation process over nanosecond timescales within the context of the FLASHForward experiment. The dependence of the ion response is studied as a function of the wakefield amplitude, plasma density and ion mass.

Chapter 9 again applies the beam-based characterisation technique but to measurements of the plasma evolution over microsecond timescales following wakefield excitation to study re-ionisation and subsequent decay of the plasma. The effects of electron-impact ionisation are measured and its dependencies on the plasma density and species are studied.

The final chapter summarises the findings of the measurements presented within this thesis and their implications for potential future operating modes of plasma-based accelerators.

In summary, the measurements presented within this thesis represent the first detailed experimental study of the entire evolution of the plasma wakefield process and its scaling with relevant experimental parameters, allowing definitive statements about future operation modes of plasma accelerators to be made.

1.4.1 *Contributions of the author*

Part II outlines measurements performed at AWAKE and FLASHForward which are necessary to the comprehension of later chapters within this thesis. The author contributed to these measurements by being part of the team that took the data, analysing the data (or subsets of the data), or discussing the results and subsequent publications at length. However, the author did not directly lead these studies.

Parts III and IV describe measurements where the author either led the analysis of experimental measurements (in the case of AWAKE; Chapter 7) or devised and performed the experiments, developed theoretical understanding, and analysed all results (in the case of FLASHForward; Chapters 5, 6, 8 and 9). All experimental data-taking was performed as part of a small team with other members of the collaborations and the implications of the results of the analysis performed by the author were discussed within the wider collaborations. All simulations described within this thesis were performed by the author unless explicitly stated otherwise.

Part I

THEORY

Part I introduces the fundamental physics that describe a plasma and its interaction with a dense, relativistic, charged particle beam.

2

PLASMA AS AN ACCELERATION MEDIUM

Plasma wakefield theory is typically described using two regimes; the *linear* regime, where relative plasma electron density perturbations are small; and the *non-linear*, 'blowout' regime where large density perturbations are considered. First, a derivation of the linear wakefield theory is outlined following the formalism of Keinigs and Jones [85]. Limitations of the linear regime are then discussed, providing motivation for operating PWFA in the non-linear regime. This regime is described in detail following the formalism of W. Lu *et al.* [86, 87] and experimental demonstrations of its use are briefly discussed. Finally, the use of the particle-in-cell method to describe the non-linear plasma dynamics is outlined while the associated limitations of this method with regards to modelling the extended timescale evolution of plasma wakefields are presented.

2.1 PLASMA WAKEFIELD ACCELERATION IN THE LINEAR REGIME

The linear regime of plasma wakefield is used to describe the response of a plasma to drive beams with densities (n_b) lower than the background plasma electron density (n_e), i.e.

$$n_b = \frac{1}{(2\pi)^{3/2}} \frac{N}{\sigma_r^2 \sigma_z} < n_e, \tag{2.1.1}$$

where N is the total number of charged particles within the beam, and σ_r and σ_z are the RMS width and length of the bi-Gaussian, cylindrically-symmetric bunch respectively. The response of the plasma to the beam can be described using a fluid description, where its motion is described by the continuity equation:

$$\frac{\partial n}{\partial t} + \nabla \cdot (n\vec{v}) = 0, \tag{2.1.2}$$

where n represents the density of the plasma and $\vec{v}$ is its velocity. This statement is an application of charge conservation, such that a local change in the plasma density is equal to the flow of plasma into or out of the region. If electromagnetic fields are present, the plasma response is described by the Lorentz Law:

$$m\frac{\partial (n\vec{v})}{\partial t} = en\left(\vec{E} + \frac{\vec{v} \times \vec{B}}{c}\right) \tag{2.1.3}$$

where the electric ($\vec{E}$) and magnetic ($\vec{B}$) fields themselves evolve via Maxwell's equations:

$$\nabla \cdot \vec{E} = \frac{\rho}{\varepsilon_0} \tag{2.1.4a}$$

$$\nabla \cdot \vec{B} = 0 \tag{2.1.4b}$$

$$\nabla \times \vec{E} = -\frac{\partial \vec{B}}{\partial t} \tag{2.1.4c}$$

$$\nabla \times \vec{B} = \mu_0 \left(\vec{J} + \varepsilon_0 \frac{\partial \vec{E}}{\partial t} \right). \tag{2.1.4d}$$

where ρ represents the charge density, $\vec{J}$ is the current density, and ε_0 and μ_0 are the permittivity and permeability of free space respectively. The combination of the previous six equations can be used to determine the evolution of the plasma and fields in a self-consistent manner, however their solution cannot be obtained analytically for arbitrary sources. Therefore, a number of simplifying assumptions must be made:

- The plasma is initially uniform and neutral, such that $n_e = n_i = n_0$.
- The mass of the ions is much larger than that of the electrons, and hence their motion is not considered on the timescale of interest. They are therefore assumed to be a stationary, positively charged, uniform background throughout.
- The source of the perturbation is an ultra-relativistic charged particle beam with $v = c$. It is therefore considered to be "rigid" and does not evolve in response to the fields generated by the plasma.
- $n_b \ll n_e$. This allows the evolution of the plasma electron density to be solved perturbatively, $n(z, r, t) = n_0 + n_1(z, r, t)$, where $n_1(z, r, t)$ is a perturbation to the initial density.

Substituting this form of the density into Eq. (2.1.2) and (2.1.3) gives

$$\frac{\partial n_1}{\partial t} = -n_0 \nabla \cdot \vec{v}, \tag{2.1.5}$$

$$\frac{\partial \vec{v}}{\partial t} = \frac{e}{m_e} \vec{E}, \tag{2.1.6}$$

where only terms that are first-order in perturbative quantities are kept. The plasma velocity $\vec{v}$, and field quantities $\vec{E}$ and $\vec{B}$ are also considered to be perturbations. Taking the divergence of Eq. (2.1.6) gives

$$\frac{\partial \nabla \cdot \vec{v}}{\partial t} = \frac{e n_0}{m_e} \nabla \cdot \vec{E}. \tag{2.1.7}$$

Substituting the form of $\nabla \cdot \vec{E}$ from Gauss' Law [Eq. (2.1.4a)],

$$\nabla \cdot \vec{E} = \frac{e}{\varepsilon_0} (n_1 + n_b), \tag{2.1.8}$$

and combining with the time derivative of Eq. (2.1.5) gives

$$\frac{\partial^2 n_1}{\partial t^2} = -\frac{n_0 e^2}{\varepsilon_0 m_e}(n_1 + n_b), \tag{2.1.9}$$

giving the definition of the electron plasma frequency,

$$\omega_{pe} = \sqrt{\frac{n_0 e^2}{\varepsilon_0 m_e}}, \tag{2.1.10}$$

which represents the fundamental measure of the timescale over which plasma electrons respond to the perturbing particle bunch. This is also commonly referred to as the plasma frequency, ω_p, dropping the specific mention to electrons as ions are typically assumed to be motionless. In reality ions do respond to the wakefield, typically on a timescale corresponding to their own plasma frequency and this is discussed in detail in Chapter 8. The (electron) plasma frequency not only defines a fundamental timescale, but is also used to define both the typical length scale of the plasma response, via calculation of the plasma wavelength,

$$\lambda_p = \frac{2\pi c}{\omega_p}, \tag{2.1.11}$$

and the amplitude of the maximal field that can be supported within a cold plasma,

$$E_0 = \frac{m_e c \omega_p}{e}. \tag{2.1.12}$$

The majority of current PWFA experiments [56, 57, 61] use plasma densities on the order of $n_e \sim 10^{14} - 10^{17}\,\mathrm{cm}^{-3}$, corresponding to plasma frequencies in the THz range, $\omega_p \sim 5 \times 10^{11} - 2 \times 10^{13}\,\mathrm{s}^{-1}$. The corresponding length scales are given by $\lambda_p \sim 3\,\mathrm{mm} - 100\,\mu\mathrm{m}$, while the maximal field amplitude is $E_0 \sim 100\,\mathrm{MVm}^{-1} - 30\,\mathrm{GVm}^{-1}$. Herein represents a fundamental advantage and inherent challenge associated with plasma-based acceleration — accelerating gradients orders of magnitude larger than those achieved within conventional RF accelerators, $\mathcal{O}(10\,\mathrm{MVm}^{-1})$, are readily available but precise control of the particle bunches on femtosecond- and micron-scales are consequently required.

2.1.1 *Plasma response*

In order to calculate the response of the plasma to a charged-particle drive bunch, we will assume an ultra-relativistic ($\beta = 1$), radially-symmetric drive bunch with density

$$n_b(r, z, t) = R(r)Z(z - ct). \tag{2.1.13}$$

We will solve for the plasma response in the comoving frame of the beam, defined by the coordinate transform $\xi = z - ct$, where ξ represents the position along the drive bunch. Applying this transformation to Eq. (2.1.9) gives

$$\frac{\partial^2 n_1(r,\xi)}{\partial \xi^2} + k_p^2 n_1(r,\xi) = -k_p^2 R(r) Z(\xi), \tag{2.1.14}$$

where $k_p = \omega_p/c$ is defined as the plasma wavenumber. The inverse of this, k_p^{-1}, is known as the plasma skin-depth and represents a measure of the depth in a plasma to which electromagnetic radiation can penetrate. It is often also used as a wakefield length scale for its similarity to the plasma wavelength, $\lambda_p = 2\pi k_p^{-1}$. It is possible to calculate the plasma response by considering a delta-function longitudinal beam profile, $Z(\xi) = \delta(\xi)$. Integrating both sides of Eq. (2.1.14) from $\xi = -\epsilon$ to $\xi = \epsilon$ in the limit $\epsilon \to 0$ gives

$$\left| \frac{\partial n_1}{\partial \xi} \right|_{\xi=0} = -k_p^2 R(r). \tag{2.1.15}$$

For $\xi < 0$, there is no source and hence no plasma response. For $\xi > 0$, the source has passed but the perturbation remains. The solution to Eq. (2.1.14) without a source is $A\sin(k_p\xi) + B\cos(k_p\xi)$ and the delta-function gives the slope of the solution at $\xi = 0$. Therefore, it is possible to infer that $A = -k_p R(r)/(2\pi r)$ and $B = 0$, giving the Green's function solution of [27]

$$n_1^G(r,\xi) = -k_p R(r)\sin(k_p\xi)\Theta(\xi), \tag{2.1.16}$$

where $\Theta(\xi)$ is the Heaviside step function. It is possible to solve the plasma perturbation for an arbitrary source function $n_b = R(r)Z(\xi)$ by performing a convolution integral of the source with the Green's function. For example, for a bi-Gaussian drive beam,

$$n_1(r,\xi) = -\frac{k_p}{(2\pi)^{3/2}} \frac{N_b}{\sigma_r^2 \sigma_z} \exp\left(-\frac{r^2}{2\sigma_r^2}\right) \times \int_\xi^\infty \exp\left(-\frac{\xi'^2}{2\sigma_z^2}\right) \sin k_p\left(\xi - \xi'\right) d\xi', \tag{2.1.17}$$

where N_b is the number of charged particles within the drive bunch, σ_r and σ_z represent the root mean square (RMS) width and length of the bunch respectively, and

$$n_{b0} = \frac{N_b}{(2\pi)^{3/2}\sigma_r^2\sigma_z}, \tag{2.1.18}$$

represents the peak density of the bi-Gaussian particle bunch.

2.1.2 *Fields*

To derive the corresponding fields in the plasma created by the density perturbations, we return to Maxwell's equations. Taking the curl of Faraday's Law [Eq. (2.1.4c)] and the time derivative of the Ampere-Maxwell Law [Eq. (2.1.4d)] and combining gives

$$\nabla^2\vec{E} - \frac{1}{c^2}\frac{\partial^2\vec{E}}{\partial t^2} = \mu_0\frac{\partial\vec{J}}{\partial t} + \frac{1}{\varepsilon_0}\nabla\rho, \tag{2.1.19}$$

via use of the vector identity $\nabla\times\nabla\times\vec{A} = \nabla(\nabla\cdot\vec{A}) - \nabla^2\vec{A}$. There are now two sources, that of the drive beam and the perturbation to the plasma density, and so the source terms are split into components, $\rho = \rho_b + \rho_p$ and $\vec{J} = \vec{J}_b + \vec{J}_p$ where $\vec{J_b} = c\rho_b\hat{z}$ and subscripts are used to represent the external beam contribution, b, and the induced plasma response, p. Using Eq. (2.1.3), under the assumption that B is negligible, the relationship between $\vec{J}_p$ and $\vec{E}$ can be found to be

$$\frac{\partial\vec{J}_p}{\partial t} = \varepsilon_0\omega_p^2\vec{E}. \tag{2.1.20}$$

Substituting this into Eq. (2.1.19) and rearranging gives

$$\left(\nabla^2 - \frac{1}{c^2}\frac{\partial^2}{\partial t^2} - \frac{\omega_p^2}{c^2}\right)\vec{E} = \mu_0 c\frac{\partial\rho_b}{\partial t}\hat{z} + \frac{1}{\varepsilon_0}\nabla\left(\rho_b + \rho_p\right), \tag{2.1.21}$$

which allows the magnetic field to be derived using Faraday's Law [Eq. (2.1.4c)] using the vector identity $\nabla\times\nabla(\phi) = 0$,

$$\left(\nabla^2 - \frac{1}{c^2}\frac{\partial^2}{\partial t^2} - \frac{\omega_p^2}{c^2}\right)\vec{B} = -\mu_0 c\nabla\times\rho_b\hat{z}. \tag{2.1.22}$$

The source term within this equation depends only on the beam current via ρ_b and hence represents a purely electrostatic plasma response such that the beam does not generate electromagnetic radiation. Equations (2.1.21), (2.1.22) and (2.1.9) can then be used to describe the response of the fields and plasma to a non-evolving source.

2.1.2.1 *Longitudinal field*

A useful tool for solving for the fields induced in the plasma by the response to the bunch is the Fourier transform in ξ,

$$E_z(\xi) = \frac{1}{2\pi}\int_{-\infty}^{\infty}\hat{E}_z(k)\exp(ik\xi)dk \tag{2.1.23}$$

as it allows derivatives to be expressed as multiplicative operators, i.e. for $\xi = z - ct$,

$$\frac{\partial}{\partial z} \to ik, \tag{2.1.24}$$

$$\frac{1}{c}\frac{\partial}{\partial t} \rightarrow -ik. \tag{2.1.25}$$

In addition to this, the Laplacian can be separated into its transverse and longitudinal components, $\nabla^2 = \nabla_\perp^2 + \partial_z^2$ and substituted into Eq. (2.1.21) to give

$$\left(\nabla_\perp^2 - k_p^2\right)\hat{E}_z = \mu_0 c\frac{\partial \hat{\rho}_b}{\partial t} + \frac{1}{\varepsilon_0}\frac{\partial}{\partial z}(\hat{\rho}_b + \hat{\rho}_p), \tag{2.1.26}$$

where derivatives with respect to z and t cancel thanks to the assumption that the velocity of the source is ultra-relativistic, i.e. $v = c$. The previously derived relation, $k_p = \omega_p/c$ is also used. The source function is defined to be a delta-function in both r and ξ,

$$\rho_b = e\frac{\delta(r - r_0)}{2\pi r}\delta\xi, \tag{2.1.27}$$

where the delta function in r represents a ring-shaped particle with radius r_0. The Fourier transform of the source is

$$\hat{\rho}_b = e\frac{\delta(r - r_0)}{2\pi r}. \tag{2.1.28}$$

Using the expression for the temporal derivative given in Eq. (2.1.25), the Fourier transform of Eq. (2.1.9) can be written as

$$c^2k^2\hat{\rho}_p = \omega_p^2(\hat{\rho}_p + \hat{\rho}_b), \tag{2.1.29}$$

such that

$$\hat{\rho}_p = \left(\frac{k_p^2}{k^2 - k_p^2}\right)\hat{\rho}_b = e\frac{k_p^2}{k^2 - k_p^2}\frac{\delta(r - r_0)}{2\pi r}. \tag{2.1.30}$$

Inserting the expressions for $\hat{\rho}_b$ and $\hat{\rho}_p$ into Eq. (2.1.26) gives

$$\left(\nabla_\perp^2 - k_p^2\right)\hat{E}_z = e\frac{ikk_p^2}{k^2 - k_p^2}\frac{\delta(r - r_0)}{2\pi r}. \tag{2.1.31}$$

Note that the contributions from ρ_b and J_b cancel here because an ultra-relativistic charge does not source a longitudinal electric field, but rather only a radially-pointing field. The remaining longitudinal electric field is due to the density modulation in the plasma.

In cylindrical coordinates, the $\nabla_\perp^2$ operator is

$$\nabla_\perp^2 = \frac{1}{r}\frac{d}{dr}\left(r\frac{d}{dr}\right) + \frac{1}{r^2}\frac{d^2}{d\phi^2} \tag{2.1.32}$$

but in this case, derivatives with respect to ϕ are zero due to the radial symmetry of the source. Therefore, Eq. (2.1.31) can be rewritten

$$\frac{1}{r}\frac{\partial}{\partial r}\left(r\frac{\partial \hat{E}_z}{\partial r}\right) - k_p^2 \hat{E}_z = e\frac{ikk_p^2}{k^2 - k_p^2}\frac{\delta(r - r_0)}{2\pi r}. \tag{2.1.33}$$

The LHS of this equation is the modified Bessel equation of zeroth order, with solution

$$\hat{E}_z(r) = AI_0(k_p r) + BK_0(k_p r) \tag{2.1.34}$$

if there is no source, where I_m and K_m are the modified Bessel functions of the m^{th} order. The Green's function solution for a radial delta-function can be written as [27]

$$\begin{aligned} g_0(r, r_0) &= [I_0(k_p r)K_0(k_p r_0)\Theta(r_0 - r) + I_0(k_p r_0)K_0(k_p r)\Theta(r - r_0)] \\ &= I_0(k_p r_<)K_0(k_p r_>) \end{aligned} \tag{2.1.35}$$

giving the solution for the Fourier transformed E_z,

$$\hat{E}_z = \frac{1}{2\pi} e\frac{ikk_p^2}{k^2 - k_p^2} g_0(r, r_0). \tag{2.1.36}$$

Performing the reverse transform gives the E_z field due to a delta-function source

$$E_z(r, \xi) = \frac{1}{2\pi} ek_p^2 i g_0(r, r_0) \int_{-\infty}^{\infty} \frac{k}{k^2 - k_p^2} e^{ik\xi} dk. \tag{2.1.37}$$

It is necessary to perform this integral over a contour in the complex plane due to the presence of the two poles at $k = \pm k_p$. The integral is made finite by defining the contour C using two parts, C_1, which is the original integral along the real axis, and C_2, an arc in the imaginary plane with radius that goes to infinity. C_2 is chosen such that its contribution to the integral is zero; by closing in the upper half plane, $ik\xi < 0$ for $\xi > 0$ and hence the exponential goes to zero as $ik\xi \to -\infty$. This will result in only having fields for $\xi > 0$ as required. The contour over which this integral is performed is shown for reference in Figure 2.1.1. The Cauchy Integral Theorem can be used to calculate the value of the integral

$$2\pi i f(z_0) = \oint \frac{f(z)dz}{z - z_0}, \tag{2.1.38}$$

where, using partial fractions, the integrand can be rewritten:

$$\begin{aligned} E_z &= \frac{1}{2\pi} ek_p^2 i g_0(r, r_0) \oint_C \frac{ke^{ik\xi}}{k^2 - k_p^2} dk \\ &= \frac{1}{2\pi} ek_p^2 i g_0(r, r_0) \oint_C \frac{e^{ik\xi}}{2}\left[\frac{1}{k - k_p} + \frac{1}{k + k_p}\right] dk. \end{aligned} \tag{2.1.39}$$

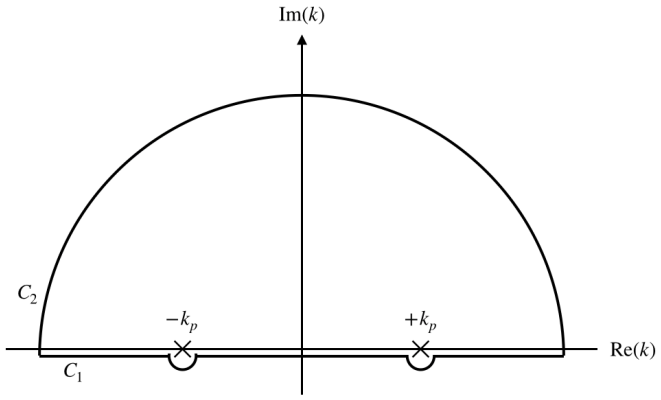

Figure 2.1.1: Contour for the integral given by Eq. (2.1.37).

Through the previous choice of C_2, its contribution to the integral is zero and hence this is equal to the original integral, Eq. (2.1.37). Application of the Integral Theorem [Eq. (2.1.38)] then gives

$$E_z^G(r,\xi) = -ek_p^2 g_0(r,r_0)\cos(k_p\xi)\Theta(\xi), \tag{2.1.40}$$

where the Heaviside function, Θ, appears due to the original choice of the contour. This expression for E_z is often referred to as the single-particle wake function W_z. The longitudinal electric field sourced by any beam distribution is hence given by the convolution of the wake function with the beam density

$$E_z(r,\xi) = -ek_p^2 \int_\xi^\infty \int_0^\infty R(r')g_0(r,r')Z(\xi')\times \cos k_p(\xi-\xi')r'dr'd\xi'. \tag{2.1.41}$$

2.1.2.2 *Transverse Field*

Following the same procedure as for the derivation of the longitudinal field, the transverse components of E and B can be found. Using Eq. (2.1.21) and (2.1.22),

$$\left(\nabla_\perp^2 - k_p^2\right)\hat{E}_r = \frac{1}{\varepsilon_0}\frac{\partial}{\partial r}(\hat{\rho}_b + \hat{\rho}_p), \tag{2.1.42}$$

$$\left(\nabla_\perp^2 - k_p^2\right)\hat{B}_\phi = \mu_0 c\frac{\partial}{\partial r}\hat{\rho}_b, \tag{2.1.43}$$

assuming that the source is radially symmetric and is ultra-relativistic and moves in the $+\hat{z}$ direction. Of most interest is the transverse force experienced by beam particles, referred to as $W_\perp$. Using the definition of the Lorentz Force,

$$\vec{W}_\perp = E_r\hat{r} + \vec{v}\times B_\phi\hat{\phi}. \tag{2.1.44}$$

where, given that we have assumed $\vec{v} = c\hat{z}$, we find $\vec{W}_\perp = (E_r - cB_\phi)\hat{r}$. Combining Eq. (2.1.42) and (2.1.43) gives

$$\left(\nabla_\perp^2 - k_p^2\right) \hat{W}_\perp = \frac{1}{\varepsilon_0} \frac{\partial \hat{\rho}_p}{\partial r}, \tag{2.1.45}$$

where again the beam source cancels as the repulsive space charge force is exactly cancelled by the attractive magnetic force of co-propagating currents. Using the previously defined form of $\hat{\rho}_p$ [Eq. (2.1.30)] gives

$$\left(\nabla_\perp^2 - k_p^2\right) \hat{W}_\perp = e \frac{k_p^2}{k^2 - k_p^2} \frac{\partial \hat{\rho}_b}{\partial r}. \tag{2.1.46}$$

In cylindrical coordinates (r, ϕ, z), the Laplacian components are given by

$$\left(\nabla^2 \vec{A}\right)_r = \nabla^2 A_r - \frac{A_r}{r^2} - \frac{2}{r^2} \frac{\partial A_\phi}{\partial \phi}, \tag{2.1.47a}$$

$$\left(\nabla^2 \vec{A}\right)_\phi = \nabla^2 A_\phi - \frac{A_\phi}{r^2} + \frac{2}{r^2} \frac{\partial A_r}{\partial \phi}, \tag{2.1.47b}$$

$$\left(\nabla^2 \vec{A}\right)_z = \nabla^2 A_z. \tag{2.1.47c}$$

with coupling between the r and ϕ components. This coupling makes solving these equations challenging, but under the assumption of cylindrical symmetry, such that there is no variation in ϕ, the LHS of Eq. (2.1.30) can be expanded

$$\frac{1}{r} \frac{\partial}{\partial r} \left(r \frac{\partial \hat{W}_r}{\partial r} \right) - \left(\frac{1}{r^2} + k_p^2 \right) \hat{W}_r = e \frac{k_p^2}{k^2 - k_p^2} \frac{\partial \hat{\rho}_b}{\partial r}. \tag{2.1.48}$$

The homogeneous solutions to this equation are the modified Bessel functions of the first order, with the corresponding Green's function [27]

$$g_1(r, r_0) = I_1(k_p r_<) K_1(k_p r_>). \tag{2.1.49}$$

The Fourier transformed transverse wake function is hence defined as

$$\hat{W}_\perp = \frac{1}{2\pi} e \frac{k_p^2}{k^2 - k_p^2} g_1(r, r_0), \tag{2.1.50}$$

and inverting the transform gives

$$W_\perp = \frac{1}{2\pi} e k_p^2 g_1(r, r_0) \int_{-\infty}^{\infty} \frac{1}{k^2 - k_p^2} e^{ik\xi} dk, \tag{2.1.51}$$

where the integral is evaluated using the same procedure as to find the longitudinal wake function, this time yielding a sine function

$$W_\perp = -e k_p g_1(r, r_0) \sin(k_p \xi). \tag{2.1.52}$$

Performing the convolution of the transverse wake function with the beam distribution in this case gives

$$F_\perp(r,\xi) = -ek_p \int_\xi^\infty \int_0^\infty \frac{\partial R(r')}{\partial r} g_1(r,r') Z(\xi') \times \sin k_p(\xi-\xi')r'dr'd\xi'. \tag{2.1.53}$$

2.1.3 *Linear regime limitations*

It was shown in Ref. [45] that for a cylindrically-symmetric electron bunch with a parabolic radial distribution described by

$$R(r) \propto 1 - \frac{r}{\sigma_r^2}, \tag{2.1.54}$$

a close approximation to a Gaussian bunch distribution, the integrals in Eqs. (2.1.41) and (2.1.53) could be given by

$$E_z = -k_p R(r) \cos(k_p z - \omega_p t), \tag{2.1.55}$$

$$F_\perp = -\frac{1}{\sigma_r^2}\left[K_2(k_p\sigma_r)I_1(k_p r) - \frac{r}{k_p\sigma_r^2}\right]\sin(k_p z - \omega_p t), \tag{2.1.56}$$

respectively, for $r < \sigma_r$. These indicate that wakefields in the linear regime have cosine-like longitudinal fields with both accelerating and decelerating regions, and sine-like transverse fields with alternating focusing and defocusing regions. The form of the transverse field is closely related to the perturbation to the density profile given in Eq. (2.1.16), while the longitudinal field lags this by a relative phase of $\pi/2$. The net result of this are regions of phase $\pi/2$ that are both accelerating and focusing for electron bunches, allowing the bunch to remain contained within the wakefield and accelerated over extended distances. A sketch indicating the regions of relative phase of the longitudinal and transverse wakefields is shown in Figure 2.1.2 for reference.

At first glance, the linear regime therefore appears to be well-suited for enabling acceleration of electrons at high gradients. However, consider the form of the transverse field given by Eq. (2.1.56). It contains a linearly-focusing component that is inversely proportional to $k_p\sigma_r^2$ and a non-linearly-defocusing component that is proportional to $I_1(k_p r)$. This defocusing component grows with r and hence the transverse bunch size should be small such that the linear component of the field dominates, i.e. $k_p\sigma_r \ll 1$. However, the evolution of the bunch itself due to the transerve field must also be considered. The transverse beam size of an electron bunch within a focusing channel evolves according to the beam envelope equation [88–90]

$$\frac{d^2\sigma_r}{d\xi^2} + k_\beta^2\sigma_r = \frac{\varepsilon^2}{\sigma_r^3}, \tag{2.1.57}$$

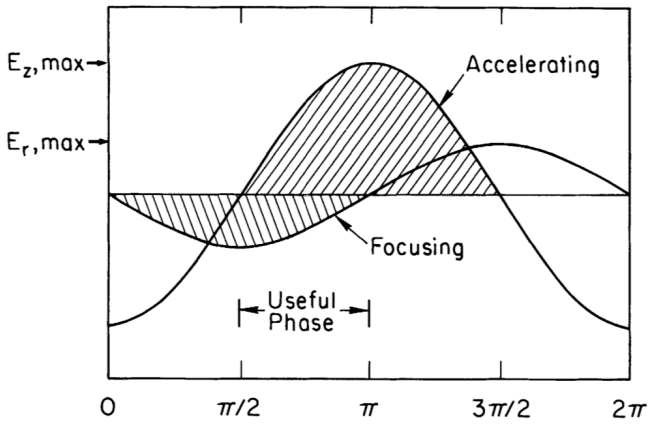

Figure 2.1.2: Demonstration of the longitudinal and transverse fields in a plasma wakefield and their relative phase. *Reproduced from Ref. [45].*

where k_β is the wavenumber of a particle that radially oscillates due to the focusing transverse field and ε is the emittance of the bunch. This implies that there exists a solution where the envelope of the bunch remains constant throughout propagation, corresponding to the balance between the divergence of the bunch due to its emittance and the focusing force supplied by the plasma. This occurs for

$$\sigma_r = \sqrt{\frac{\varepsilon}{k_\beta}} = \sqrt{\beta\varepsilon}, \tag{2.1.58}$$

where $\beta = k_\beta^{-1}$ is the required β-function [91] at focus of the incoming electron bunch. When the bunch satisfies such a constraint, it is said to be 'matched' to the plasma [92, 93]. Away from this value, the transverse size of the bunch oscillates around that of the matched bunch as it propagates through the plasma.

To understand one of the constraints the form of the transverse field in Eq. (2.1.56) imposes on the stable propagation of a bunch, consider a narrow ($k_p\sigma_r \ll 1$) bunch that is not well-matched into the plasma, i.e. its transverse size does not satisfy Eq. (2.1.58). Initially, thanks to its small transverse size, the focusing force on the bunch can be assumed to be linear in r and scales as $F_\perp \propto r/(k_p\sigma_r^4)$. This force strongly depends on the size of the bunch and hence evolves as the transverse bunch size oscillates through the plasma. As the bunch size decreases, the transverse field grows in amplitude, further focusing the bunch in a positive feedback effect. This increases the local bunch density and more strongly perturbs the plasma electrons until none are left within the vicinity of the bunch and only plasma ions remain to provide the focusing force. Such a configuration is referred to as the 'blowout regime' [48, 89, 94, 95]. Hence the linear wakefield regime is unstable and the bunch will naturally evolve such that it produces a non-linear wake. However, the non-linear blowout regime has many advantages for high-quality electron acceleration and is discussed in detail in the following section.

2.2 PLASMA WAKEFIELD ACCELERATION IN THE NON-LINEAR REGIME

The derivation of the linear wakefield regime in the previous section relied on a cold fluid model of the plasma response to the perturbation by a charged-particle bunch. Use of such a model prevents consideration of the relativistic motion of strongly perturbed plasma electrons or complete cavitation of the wakefield, the aforementioned blowout. Hence, a fully relativistic, non-linear kinetic theory was sought as described in Refs. [47, 96]. Such derivations of the non-linear regime of the plasma wakefields relied on a Hamiltonian description of the relativistic plasma electrons given by

$$H = \gamma m_e c^2 + q\phi \tag{2.2.1}$$

where ϕ is the scalar potential of the wakefield, $q = -e$ is the charge of an electron, and the energy of a particle is related to its momentum, p, by

$$\gamma = \sqrt{1 + \frac{p^2}{m_e^2 c^2}}. \tag{2.2.2}$$

The momentum of a particle is related to the canonical momentum $\vec{P}$ via

$$\vec{P} = \vec{p} + \frac{q\vec{A}}{c}, \tag{2.2.3}$$

where $\vec{A}$ is the vector potential of the wakefield. An important approximation with regard to the motion of the particles is now made, that of the quasi-static approximation (QSA) [47]. This states that the motion of plasma electrons that respond to the perturbation from the drive bunch occurs on a timescale far shorter than the perturbation to the drive bunch induced by the plasma response. Use of such an approximation enables the decoupling of the two evolutionary timescales and is commonly used within particle-in-cell (PIC) simulations to enable considerable speedup as the plasma and bunch response can be separately calculated. Under the QSA, the Hamiltonian defined by Eq. (2.2.1) depends on z and t via the combination $\xi = z - ct$. By definition, $dH/dt = \partial H/\partial t$ and hence

$$\frac{\partial H}{\partial t} = -c\frac{\partial H}{\partial \xi} = -c\frac{\partial H}{\partial z} = c\frac{dP_z}{dt}, \tag{2.2.4}$$

via Hamilton's equations, leading to the constant of motion

$$\frac{d}{dt}(H - cP_z) = 0. \tag{2.2.5}$$

Before the drive bunch arrives the plasma electrons are stationary and hence $H - cP_z = m_e c^2$. After the arrival of the bunch,

$$H - cP_z = \gamma m_e c^2 + q\phi - cp_z - qA_z = m_e c^2, \tag{2.2.6}$$

via Eqs. (2.2.3) and (2.2.5). This leads to the definition of the wakefield pseudo-potential $\psi = \phi - A_z$ such that

$$\gamma m_e c^2 - c p_z = m_e c^2 + e\psi. \tag{2.2.7}$$

Application of the Lorenz gauge condition,

$$\nabla \cdot \vec{A} + \frac{1}{c^2}\frac{\partial \phi}{\partial t} = 0, \tag{2.2.8}$$

allows Maxwell's equations to be re-written

$$\left(\frac{1}{c^2}\frac{\partial^2}{\partial t^2} - \nabla^2\right)\begin{bmatrix}\vec{A}\\ \phi\end{bmatrix} = \begin{bmatrix}\mu_0 \vec{J}\\ \rho/\varepsilon_0\end{bmatrix}. \tag{2.2.9}$$

At this point, we adopt the formalism within Ref. [87] and move to normalised 'plasma' units where timescales are normalised to the plasma frequency, ω_p^{-1}, length scales to the plasma skin-depth, k_p^{-1}, velocities to the speed of light, c, and charges to the electron charge, e. In addition, the use of the QSA implies temporal derivatives are much smaller in magnitude than spatial derivatives in the co-moving frame, i.e. $\partial_t \ll \partial_\xi$. This allows Maxwell's equations to be rewritten

$$-\nabla_\perp^2 \begin{bmatrix}\vec{A}\\ \phi\end{bmatrix} = \begin{bmatrix}\vec{J}\\ \rho\end{bmatrix}, \tag{2.2.10}$$

and the wakefield pseudo-potential therefore obeys the Poisson-like equation

$$-\nabla_\perp^2 \psi = \rho - J_z. \tag{2.2.11}$$

The form of the longitudinal field can be extracted from the relation between the electric field and the vector and scalar potentials in the Lorenz gauge

$$\vec{E} = -\nabla\phi - \frac{\partial \vec{A}}{\partial t}, \tag{2.2.12}$$

such that

$$E_z = \frac{\partial \psi}{\partial \xi}. \tag{2.2.13}$$

Hence the form of the wakefield pseudo-potential is of most interest and can be used to describe the fields induced within the plasma. The source of the pseudo-potential depends only on the plasma charge density, ρ, and the axial current, J_z, via Eq. (2.2.11) which themselves depend on the fields via the equation of motion for plasma electrons [96],

$$\frac{d\vec{P}_\perp}{d\xi} = \frac{1}{1 - v_z}\left[-\vec{E}_\perp + (\vec{v} \times \vec{B})_\perp\right]. \tag{2.2.14}$$

The definition for the relativistic factor (in normalised units), $\gamma = (1 + P^2)^{1/2}$, can be combined with the constant of the motion derived earlier,

$$\gamma - P_z = 1 + \psi, \tag{2.2.15}$$

a repeat of Eq. (2.2.7) in normalised units, to link the axial momentum of plasma electrons to their transverse motion via the pseudo-potential,

$$P_z = \frac{1 + P_\perp^2 - (1+\psi)^2}{2(1+\psi)}, \tag{2.2.16}$$

$$\gamma = \frac{1 + P_\perp^2 + (1+\psi)^2}{2(1+\psi)}, \tag{2.2.17}$$

$$1 - v_z = \frac{1}{\gamma}(\gamma - P_z) = \frac{2(1+\psi)^2}{1 + P_\perp^2 + (1+\psi)^2}, \tag{2.2.18}$$

such that once $P_\perp$ is solved via Eq. (2.2.14), the axial momentum, and hence J_z, can be calculated. Note that the drive bunch does not contribute to ψ as $J_{zb} = c\rho_b$, and hence $c\rho_b - J_{zb} = 0$.

2.2.1 *Solving the equation of motion*

The equation of motion for the transverse motion of plasma electrons perturbed by an electron bunch can first be solved via the application of the Dawson sheet model [97] where a plasma is approximated by concentric rings of infinitely long sheets of electrons on a uniform ion background. The force on an electron ring can then be calculated from determination of the electrostatic force from the total charge within that ring. This approach is strictly only valid provided that electron rings do not cross each other, as occurs within the blowout regime, and hence such an approximation enables determination of the transition between the linear and blowout regimes.

By applying the non-relativistic limit ($v_z \ll v_\perp \ll 1$, $v_\perp \approx dr/d\xi$) to Eq. (2.2.14) and assuming that the total force on an individual electron ring includes contributions from the ion charge, electron bunch charge and plasma electron charge within the ring, the equation of motion can be re-written,

$$\frac{d^2 r}{d\xi^2} = -\frac{1}{2}r + \frac{c(r_0, r, \xi)}{r}, \tag{2.2.19}$$

where r_0 is the initial position of the sheet, and r is its radial position at ξ. The first term on the RHS comes from the uniform ion background while the second term represents the force

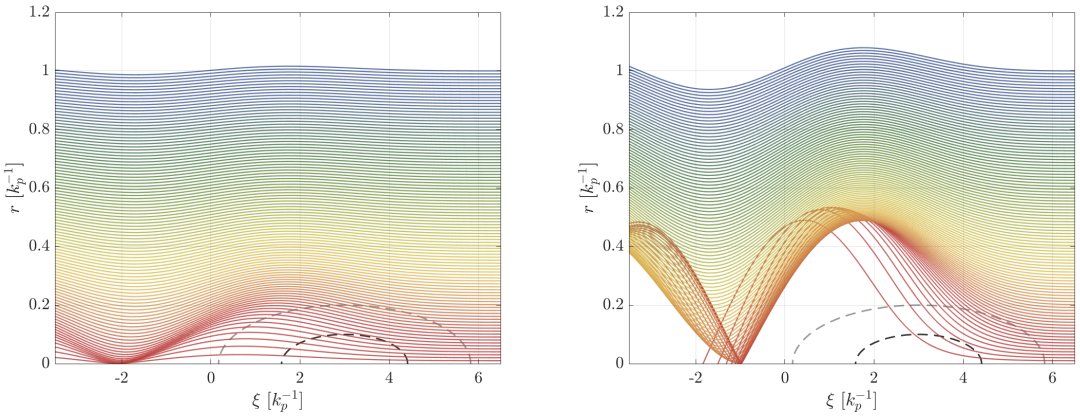

(a) (b)

Figure 2.2.1: Solution of Eq. (2.2.22) over $0 \leqslant r \leqslant 1k_p^{-1}$ for a bi-Gaussian drive beam with $\sigma_r = 0.1$, $\sigma_z = \sqrt{2}$ and (a) $n_{b0} = 1$ and (b) $n_{b0} = 5$. The colour of a given electron trajectory represents the initial radial position of each electron and is the same for both figures. The 1 and 2σ contours of the drive beam are plotted using black and grey dashed lines respectively.

from the electron charge within the ring, consisting of both the drive bunch and the plasma electrons. Under the assumption of no trajectory crossing, this term can be calculated via,

$$c(r_0, r, \xi) = \frac{1}{2}r_0^2 + \int_0^r r' n_b(r', \xi) dr', \tag{2.2.20}$$

where n_b represents the drive bunch density distribution. For a bi-Gaussian drive beam described by

$$n_b(r, \xi) = n_{b0} \exp\left(\frac{-(\xi - \xi_0)^2}{2\sigma_z^2}\right) \exp\left(\frac{-r^2}{2\sigma_r^2}\right) \tag{2.2.21}$$

this integral can be performed analytically and the equation of motion becomes

$$\frac{d^2r}{d\xi^2} = -\frac{1}{2}r + \frac{1}{r}\left\{\frac{1}{2}r_0^2 + n_{b0} \exp\left(\frac{-(\xi - \xi_0)^2}{2\sigma_z^2}\right)\left[\sigma_r^2\left(1 - \exp\left(\frac{-r^2}{2\sigma_r^2}\right)\right)\right]\right\}. \tag{2.2.22}$$

This equation of motion can be solved numerically to give the plasma electron response as shown in Fig. 2.2.1. In this figure, the response of plasma electrons to a bi-Gaussian drive beam with $\sigma_r = 0.1$, $\sigma_z = \sqrt{2}$ and (a) $n_{b0} = 1$ and (b) $n_{b0} = 5$ is solved. In Figure 2.2.1(a) the trajectories of plasma electrons can be seen to converge at the rear of the wakefield but trajectory crossing is not observed. However, when the plasma is more strongly perturbed, as in Figure 2.2.1(b), trajectory crossing can be seen. This clearly indicates a natural transition from no crossing to crossing, and represents the transition from the linear regime to the non-linear regime. It was shown in Refs. [86, 87] that this transition occurs for peak beam densities exceeding $n_b > 1.792 n_{e0}$.

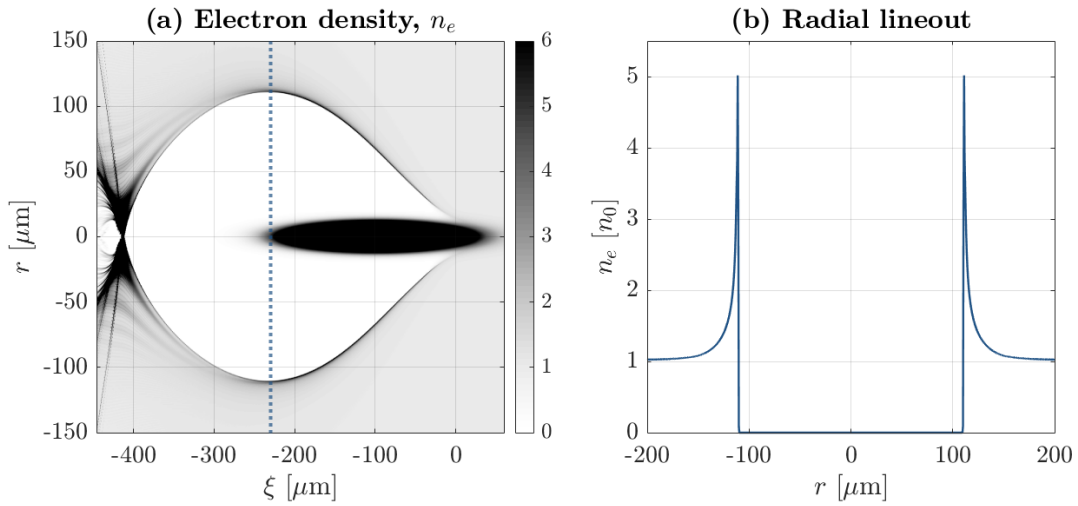

Figure 2.2.2: Example of the non-linear blowout regime excited by a high-density electron bunch. (a) Plasma electron density profile. (b) Radial lineout of the plasma electron density at $\xi = -230\,\mu$m, close to the maximum radius of the wakefield cavity.

The transition to plasma electron trajectory crossing means the simplified Dawson sheet model can no longer be used and a more general model is required to describe wakefield excitation in the non-linear blowout regime. An example of such a blowout is shown in Figure 2.2.2 and demonstrates complete cavitation of the wakefield by the high-density drive bunch such that no plasma electrons remain within the bubble. A detailed derivation of the three-dimensional non-linear wakefield theory follows using the formalism first described in Refs. [86, 87]. This is necessary background as within Chapter 8 the derivation is extended to consider non-uniform ion density profiles to assist interpretation of beam-based measurements of wakefield-induced ion motion. Note, coefficients within this derivation may differ by factors of -1 from those found in Refs. [86, 87] as a different definition of the co-moving frame ($\xi = ct - z$) was used in that case.

The blowout regime can be modelled by using the observation that the plasma density and current profile can be split into three distinct regions that each contribute to the fields experienced by electrons as shown in Figure 2.2.2(b); the first is the wakefield cavity where plasma electrons have been expelled and only the uniform ion background remains. At the edge of this region exists a thin sheath where plasma electrons stream around the cavity and have high density, and beyond this exists a linear response region where the plasma perturbation is small. The shape of the cavity, and hence fields within, can be described by considering the trajectory of the inner-most plasma electron, $r_b(\xi)$.

Within the blown-out cavity, for $r \leqslant r_b$ and for $r \gg \sigma_r$, the forms of the potentials are given by

$$\phi = \phi_0(\xi) - \frac{r^2}{4} + \lambda(\xi)\ln r, \tag{2.2.23}$$

$$A_z = A_{z0}(\xi) + \lambda(\xi)\ln r, \tag{2.2.24}$$

where the second term on the RHS of Eq. (2.2.23) comes from solving Poisson's equation in a cylindrical cavity of uniform ion density, terms with $_0$ subscripts represent on-axis values at $r = 0$, and $\lambda(\xi)$ represents the contribution from the drive bunch density via

$$\lambda(\xi) = \int_0^{r \gg \sigma_r} r' n_b(r', \xi) dr', \tag{2.2.25}$$

while the logarithmic dependence, $\ln r$, appears due to the integration of the $1/r$ dependence in the force due to the bunch. Under the QSA, the gauge condition, Eq. (2.2.8), can be rewritten in normalised units as

$$\frac{1}{r}\frac{\partial(rA_r)}{\partial r} = \frac{\partial}{\partial \xi}(\phi - A_z), \tag{2.2.26}$$

for a cylindrically-symmetric system such that the radial component of the vector potential, and hence the radial plasma current, is defined by

$$A_r = \sigma(\xi) r, \tag{2.2.27}$$

where

$$\sigma(\xi) = \frac{1}{2}\frac{d\psi_0(\xi)}{d\xi}, \tag{2.2.28}$$

and $\psi = \psi_0 - r^2/4$ from the combination of Eqs. (2.2.23) and (2.2.24). Using these relations, the force acting on a plasma electron at $r = r_b(\xi)$ or for an electron within the drive bunch can be written as,

$$\begin{aligned} F_\perp &= -(E_r - v_z B_\phi) = \frac{\partial \phi}{\partial r} - v_z \frac{\partial A_z}{\partial r} + (1 - v_z)\frac{\partial A_r}{\partial \xi} \\ &= -\frac{1}{2} r + (1 - v_z)\frac{\lambda(\xi)}{r} + (1 - v_z)\frac{d\sigma}{d\xi} r, \end{aligned} \tag{2.2.29}$$

where the first term is due to the uniform density ion cavity as before, and the latter two terms represent the electric and magnetic fields from the drive bunch and radial plasma currents respectively. For ultra-relativistic electrons within the drive bunch, $v_z \sim 1$ and hence the force is purely due to the ion cavity whereas plasma electrons, whose velocities can range between $-1 \leqslant v_z \ll 1$, experience the full electromagnetic nature of the wake. To calculate the force on an electron at the boundary of the cavity, $\sigma(\xi)$ must be determined via calculation of $\psi_0(\xi)$.

Application of charge conservation, $\partial \rho/\partial t + \nabla \cdot \vec{J} = 0$, allows the continuity equation to be derived,

$$\frac{\partial}{\partial \xi}(\rho - J_z) + \nabla_\perp \cdot \vec{J}_\perp = 0, \tag{2.2.30}$$

which when integrated gives the requirement

$$\frac{d}{d\xi}\int r(\rho - J_z)dr = 0. \tag{2.2.31}$$

Far in front of the bunch where no perturbation to the plasma has occurred, $\rho - J_z = 0$, and hence the integral must equal zero for all ξ. At each ξ, the source of ψ is given by the combination of the charge due to the ion background, ρ_{ion}, and the plasma electrons, ρ_e, alongside the axial plasma electron current density, J_{ze}, via

$$\rho - J_z = \rho_{ion} + \rho_e - J_{ze}, \tag{2.2.32}$$

where $\rho_{ion} = 1$ for all r as the ions represent a uniform, constant background density. Within the blown-out wakefield cavity, plasma electrons are entirely expelled and hence $\rho_e - J_{ze} = 0$ for $r < r_b$. However, its value sharply rises at the edge of the cavity within the surrounding thin sheath, and then exponentially falls away outside this. It was shown in Refs. [86, 87] that reproduction of the exact form of $\rho - J_z$ in the sheath and surrounding region was not necessary to accurately predict the wakefield response and instead the form of the pseudo-potential could be approximated by

$$\psi = \frac{r_b^2(\xi)}{4}(1 + \beta(\xi)) - \frac{r^2}{4}, \tag{2.2.33}$$

where

$$\beta(\xi) = \frac{(1+\alpha)^2 \ln(1+\alpha)^2}{(1+\alpha)^2 - 1} - 1, \tag{2.2.34}$$

and $\alpha = \Delta/r_b$ is a measure of the combined relative size of the sheath and linear regions outside the wakefield cavity compared to the size of the cavity itself. If the plasma is strongly perturbed by a bunch of very high density, the radius of the cavity far exceeds that of the sheath and $\beta \ll 1$. However, if a weak non-linear blowout is driven, the contributions from the sheath and linear regions outside the cavity must be considered and $\beta \approx 2\ln\alpha$.

The form of the wakefield pseudo-potential is given in Eq. (2.2.33) and can hence be used within the relativistic equation of motion for a plasma electron,

$$\frac{dP_\perp}{d\xi} = \frac{1}{1 - v_z}F_\perp, \tag{2.2.35}$$

the LHS of which can be rewritten,

$$\frac{dP_\perp}{d\xi} = \frac{d(\gamma v_\perp)}{d\xi} = \frac{d}{d\xi}\left[\gamma(1 - v_z)\frac{d}{d\xi}r_\perp\right] = \frac{d}{d\xi}\left[(1+\psi)\frac{d}{d\xi}r_\perp\right], \tag{2.2.36}$$

and Eqs. (2.2.18) and (2.2.29) can be combined to give the equation of motion for a plasma electron at r_b,

$$\frac{d}{d\xi}\left[(1+\psi)\frac{d}{d\xi}r_b\right] = r_b\left(-\frac{1}{4}\left[1+\frac{1}{(1+\psi)^2}+\left(\frac{dr_b}{d\xi}\right)^2\right]+\frac{d\sigma}{d\xi}+\frac{\lambda(\xi)}{r_b^2}\right). \tag{2.2.37}$$

By performing the relevant substitutions, i.e. the form of $\psi(r_b(\xi))$ from Eq. (2.2.33) and σ from Eq. (2.2.28), the equation of motion of a plasma electron at the cavity boundary can be rewritten in the following form,

$$A(r_b)\frac{d^2r_b}{d\xi^2}+B(r_b)r_b\left(\frac{dr_b}{d\xi}\right)^2+C(r_b)r_b=\frac{\lambda(\xi)}{r_b}, \tag{2.2.38}$$

where the coefficients are given by

$$A(r_b)=1+\left[\frac{1}{4}+\frac{\beta}{2}+\frac{1}{8}r_b\frac{d\beta}{dr_b}\right]r_b^2, \tag{2.2.39}$$

$$B(r_b)=\frac{1}{2}+\frac{3}{4}\beta+\frac{3}{4}r_b\frac{d\beta}{dr_b}+\frac{1}{8}r_b^2\frac{d^2\beta}{dr_b^2}, \tag{2.2.40}$$

$$C(r_b)=\frac{1}{4}\left[1+\frac{1}{\left(1+\frac{\beta}{4}r_b^2\right)^2}\right]. \tag{2.2.41}$$

This equation of motion can be solved numerically for a given drive bunch, defined by $\lambda(\xi)$ via Eq. (2.2.25), and the corresponding longitudinal field can be calculated via,

$$E_z(r,\xi)=\frac{d\psi}{d\xi}=\frac{d}{d\xi}\left[\frac{1}{4}r_b^2(1+\beta(\xi))\right]. \tag{2.2.42}$$

An example of the numerical solution to Eqs. (2.2.38) and (2.2.42) for an electron bunch of RMS length $k_p\sigma_z = 1$, width $k_p\sigma_r = 0.1$, and charge 3.8 nC centred at $\xi = -3\,k_p^{-1}$ on a background plasma density of $n_e = 10^{16}\,\mathrm{cm}^{-3}$ is shown in Figure 2.2.3.

The transverse wakefield is provided by the force from the unshielded ions that remain within the cavity and hence is simply given by

$$W_\perp(r,\xi)=-\frac{r}{2}, \tag{2.2.43}$$

in normalised units, assuming a uniform ion density. The above two equations indicate why the blowout regime is ideally suited for the acceleration of electrons. The longitudinal field, Eq (2.2.42), is independent of r and hence all particles within a given longitudinal slice within the wake experience the same longitudinal field amplitude. The transverse field, Eq. (2.2.43), is independent of ξ and linear in r and hence the emittance of a witness bunch can be preserved as it propagates through the plasma [48].

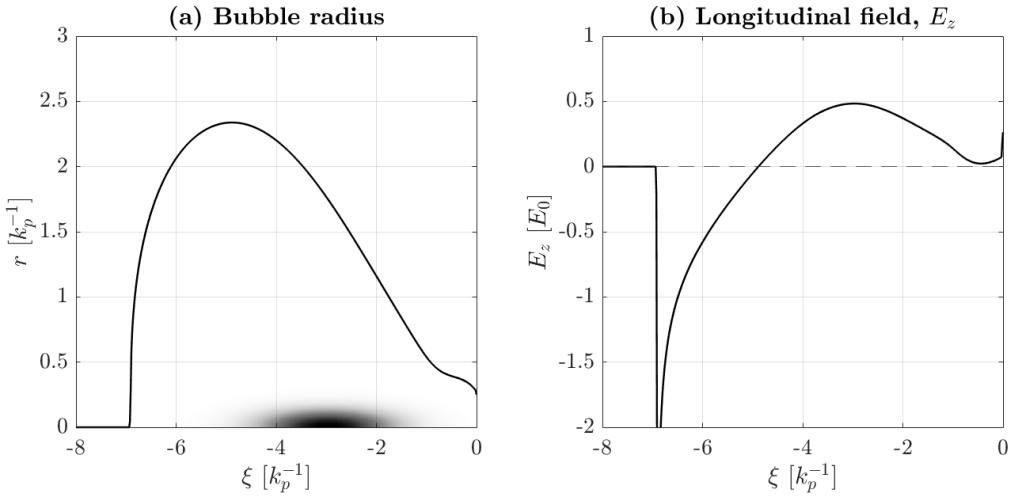

Figure 2.2.3: Solution to (a) Eq. (2.2.38) and (b) Eq. (2.2.42) for an electron bunch of RMS length $k_p\sigma_z = 1$, width $k_p\sigma_r = 0.1$, and charge 3.8 nC centred at $\xi = -3\,k_p^{-1}$ on a background plasma density of $n_e = 10^{16}\,\mathrm{cm}^{-3}$.

2.2.2 *Beam-loading*

The non-linear blowout regime hence represents the optimal regime for electron acceleration thanks to its large amplitude accelerating gradients ($> \mathrm{GVm}^{-1}$) and linear transverse fields. Typical PWFA experiments assume a 'two-bunch' scheme where a short, high-density electron bunch drives a large amplitude wakefield within the plasma and is followed by a second 'witness' electron bunch that samples the wakefield driven by the first bunch and is accelerated to high energies. However, while the longitudinal field within the non-linear blowout has been shown to be independent of r, it has a clear ξ-dependence as can be observed in Figure 2.2.3(b). A witness bunch placed towards the rear of the wakefield cavity will experience a longitudinally-varying accelerating field, causing particles at different ξ-positions within the witness bunch to be accelerated to different energies and will hence induce a large energy spread. However, this effect can be countered by appropriately tailoring the current profile of the witness bunch such that it locally 'flattens' the wakefield in its vicinity, an effect commonly referred to as 'beam-loading' [98, 99].

In the linear wakefield regime, the field experienced by a witness bunch is the sum of the wake from the leading drive bunch and the wake driven by the witness bunch itself [99], which can be written as

$$E_z(\xi) = E_{z0}\cos(k_p\xi) - ek_p^2\int_{\xi_0}^{\xi}\lambda(\xi')\cos[k_p(\xi-\xi')]\mathrm{d}\xi', \tag{2.2.44}$$

using Eq. (2.1.41), where ξ_0 represents the longitudinal position of the leading edge of the witness bunch with respect to the local maximum of the field. It was shown in Ref. [99] that a

current profile of the form $\lambda(\xi) = a\xi + b$ leads to an accelerating field of constant amplitude, E_{acc}, within the vicinity of the witness bunch, where the coefficients are given by

$$a = -\frac{E_{z0}}{e}\cos(k_p\xi_0), \tag{2.2.45}$$

$$b = \frac{E_{z0}}{e}\left[\frac{\sin(k_p\xi_0)}{k_p} - \xi_0\cos(k_p\xi_0)\right]. \tag{2.2.46}$$

Such a current profile is triangular (or trapezoidal if sharply cut at the rear), with a sharp rise in current at its leading edge followed by a linear fall off. This represents the opposite of the 'optimal' drive bunch current profile where the current rises linearly from its leading edge and sharply drops at its rear [100]. Performing the integration within Eq. (2.2.44) using the optimal bunch profile reveals that the value of the accelerating gradient at which the witness bunch loads the wakefield depends on its position via $E_{acc} = E_0\cos(k_p\xi_0)$. Additionally the efficiency with which the witness bunch can extract energy from the wakefield is given by $\eta = 1 - E_{acc}^2/E_{z0}^2 = \sin^2(k_p\xi_0)$ and hence also depends on its relative position [99]. These dependencies indicate that it is not possible to simultaneously maximise both the accelerating gradient and the efficiency in the linear regime — the efficiency only reaches 100% for zero accelerating field and vice versa. Something of a compromise can be found by setting $k_p\xi_0 = \pi/4$ such that approximately 50% of the wake energy can be extracted by the witness bunch while it experiences an accelerating gradient of approximately 70% of the peak.

Building upon the description of the non-linear wakefield developed in Refs. [86, 87], it was shown in Ref. [101] that the same trapezoidal current profile provided local flattening of the wakefield in the non-linear blowout regime. In this case the coefficients were given by,

$$a = -\frac{E_{acc}}{e}, \tag{2.2.47}$$

$$b = \frac{E_{acc}}{e}\left[\xi_0 + \sqrt{\left(\frac{eE_{acc}}{m_e c\omega_p}\right)^4 + \left(\frac{\omega_p R_b}{2c}\right)^4}\right], \tag{2.2.48}$$

where R_b is the maximum radius of the blown-out wakefield cavity and E_{acc} is again a function of ξ_0. In this case, the energy extraction of the witness bunch was shown to be described by

$$Q_{acc}E_{acc} = \frac{m_e c^2}{64 r_e}(k_p R_b)^4, \tag{2.2.49}$$

where r_e is the classical electron radius and Q_{acc} is the total witness bunch charge. The total witness bunch charge that can be accelerated within the wakefield depends on the accelerating gradient. However it was shown that the energy extraction efficiency could be close to 100% for any E_{acc} in the blowout regime [101] — another crucial advantage compared to the linear regime. This was confirmed in PIC simulations that consistently demonstrated

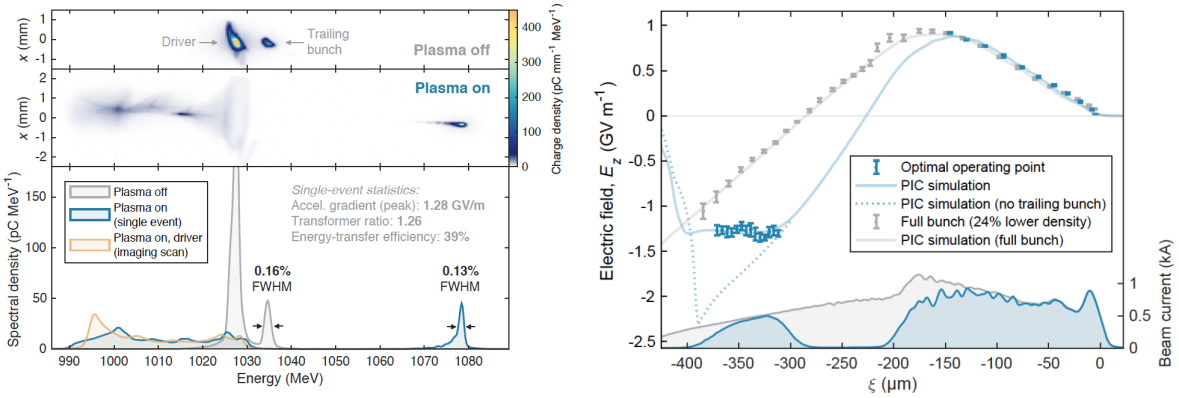

Figure 2.2.4: Experimental demonstration of beam-loading, and hence energy-spread preservation, in the non-linear regime with simultaneous high energy-transfer efficiency within the FLASHForward experiment. *Image reproduced from Ref. [68].*

approximately 90% efficiency for a range of optimally-loaded witness bunches at different longitudinal positions (and hence E_{acc}) within a given wakefield [102].

Experimentally, the FACET experiment [55] first demonstrated the effect of loading a non-linear beam-driven wakefield with witness bunches measured to have energy spreads on the order of 2% following acceleration by approximately 2 GeV, while the average energy-transfer efficiency was measured to be $(29.1 \pm 8.9)\%$ across 92 events [67]. More recently the FLASHForward experiment [57] demonstrated beam-loading and field-flattening in the vicinity of the witness bunch at the few-percent level, allowing energy-spread preservation at the per-mille level [68] as shown in Figure 2.2.4. Here, the energy-transfer efficiency between the drive and witness bunch was measured to be $(42 \pm 4)\%$ across 5000 consecutive events. This result is discussed in more detail in Chapter 4.

2.3 THE PARTICLE-IN-CELL METHOD

As demonstrated by the previous sections, the modelling of plasma-based acceleration is highly complex and involves consideration of the non-linear, relativistic interaction between charged-particle bunches and plasmas. Ever since the seminal work of Tajima and Dawson in 1979 [31], where controlled plasma-based acceleration was first proposed, computational methods have played a crucial role in understanding the physical mechanisms at play. While great advancements have been made in experimental diagnostics that enable measurement of the complex plasma interaction [103], such as shadowgraphy [104] and progressive sampling of the wakefield [105], the interaction is still somewhat treated as a 'black box' with numerical simulations used to inform the analysis and interpretation of experimental results. To this day, it is rare to see publications based on experimental results alone.

The most commonly used numerical method in plasma-based acceleration is the particle-in-cell (PIC) method [106–108]. A plasma is a collection of charged particles, electrons and ions, that interact with each other via self-consistently generated fields. The PIC method

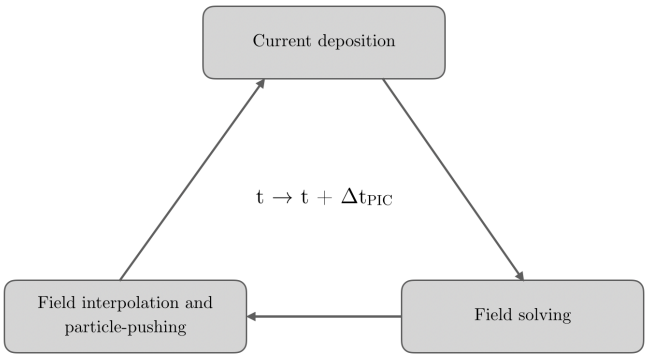

Figure 2.3.1: Representation of a typical time-integration PIC loop. *Modified from Ref. [109].*

attempts to recreate this using a smaller number of 'macroparticles' that each represent a cluster of real particles with the same charge-to-mass ratio as real plasma electrons and ions — in a typical PIC simulation, a single macroparticle represents approximately 1000 beam or plasma particles. The PIC method represents an Eulerian-Lagrangian system where the fields and charge and current densities are defined on a spatial grid (Eulerian) and macroparticles move continuously throughout this grid (Lagrangian). Numerical implementation of the PIC technique proceeds in a loop:

1. Deposit the current and charge at the grid points due to the distribution of the macroparticles.
2. Calculate the associated fields at the grid points via Maxwell's equations.
3. Interpolate the fields onto the positions of the macroparticles.
4. Advance the positions of the macroparticles according to the fields.

Each iteration of this loop represents a single time-step, Δt_{PIC}, of the calculation and is shown in Figure 2.3.1. The size of both the time-step and spatial grid is defined by the smallest temporal and spatial scales that must be fully resolved. For beam-driven plasma accelerators these are typically the inverse plasma frequency, $\tau = \omega_p^{-1}$, and plasma skin-depth, k_p^{-1}, respectively. For laser-driven acceleration this is typically defined by the wavelength of the laser pulse itself, λ_0, and hence much finer spatiotemporal resolution is required in this case. For both cases, a significant number (tens) of spatial and temporal steps is required per appropriate scale. If these requirements are not met, effects of numerical dispersion can lead to unphysical results. This is a result of the finite difference techniques [110] commonly used to solve the field equations, which cause electromagnetic waves of frequencies close to the cutoff frequency, $\omega_c \sim \pi/\Delta t_{PIC}$, to propagate at speeds lower than the speed of light [111]. This can lead to high energy particles emitting 'numerical Cherenkov radiation' and exhibit, for example, unphysical emittance growth [112].

Due to the required resolution ($\lesssim \mu m$) and total number of time-steps, PIC simulations are incredibly computationally demanding. To enable calculations to be performed on reasonable timescales, PIC codes are massively parallelised and simulations are regularly performed on high-performance computing (HPC) facilities capable of performing 10^{16} floating-point operations per second [113]. Despite this, three-dimensional simulations of extended (metre-scale) propagation distances require many millions of time-steps over $\mathcal{O}(1000^3)$ grid points and hence can take days even when using a large number of cores. This fundamental difference between the length-scales within the plasma and the typical acceleration lengths that need to be simulated led to the development of alternative simulation schemes that enable significant speed-up, one of which is the aforementioned quasi-static approximation.

2.3.1 *Quasi-static approximation*

In Section 2.2 it was shown that the quasi-static approximation is applicable when the evolutionary timescales of the response of the plasma and the beam differ significantly [47]. The QSA can be applied to the PIC loop, enabling the treatment of the plasma and particle- or laser-beam using a decoupled time-advance [114, 115], such that time-steps orders of magnitude larger than those in full PIC schemes can be used. The relevant increase in time-step can be calculated via consideration of the typical response rate of a charged-particle bunch to the plasma, given by $\tau_\beta = \sqrt{2\gamma}/\omega_p$. For typical FLASHForward parameters, $\gamma \sim 2500$ and hence the evolution time of the electron bunch is approximately 70 times that of the plasma and a similarly large increase in time-step can be used.

In a quasi-static PIC code, a staggered numerical scheme is introduced where beam and plasma macro-particles are treated separately in a co-moving frame defined by $\xi = z - ct$ [116–118]. After the beam-particles and their current are deposited on the grid they are frozen in place while the plasma-particles are pushed and fields are computed. Plasma-particles are advanced from their initial, unperturbed state ahead of the bunch in the co-moving frame and evolve in the negative ξ-direction along the length of the bunch as demonstrated in Figure 2.3.2. Once the plasma-particles reach the rear of the co-moving frame, the fields and plasma are frozen while the beam-particles are evolved according to the induced fields within the plasma over the increased time-step — in this case, $\Delta t_{QSA} \sim \sqrt{2\gamma}\,\Delta t_{PIC}$. The resulting beam-particle distribution is calculated and its current is deposited on a new grid and frozen, and again a new set of plasma-particles are pushed and evolved in the co-moving frame in a repeating loop. By decoupling the evolution of the plasma and bunch in this way, the full three-dimensional electromagnetic field solve and particle push is effectively reduced to a two-dimensional version, enabling significant speed-up with minimal loss of accuracy [116].

Within this thesis, two quasi-static PIC codes are consistently used for comparison to, and interpretation of, experimental results. Experiments performed at FLASHForward are compared to the three-dimensional quasi-static PIC code HiPACE [117]. However, due to the length of both the plasma ($L_p = 10\,m$) and proton bunch ($\sigma_z \sim 6 - 8\,cm$) in the AWAKE experiment, a three-dimensional simulation of the plasma interaction is too computationally

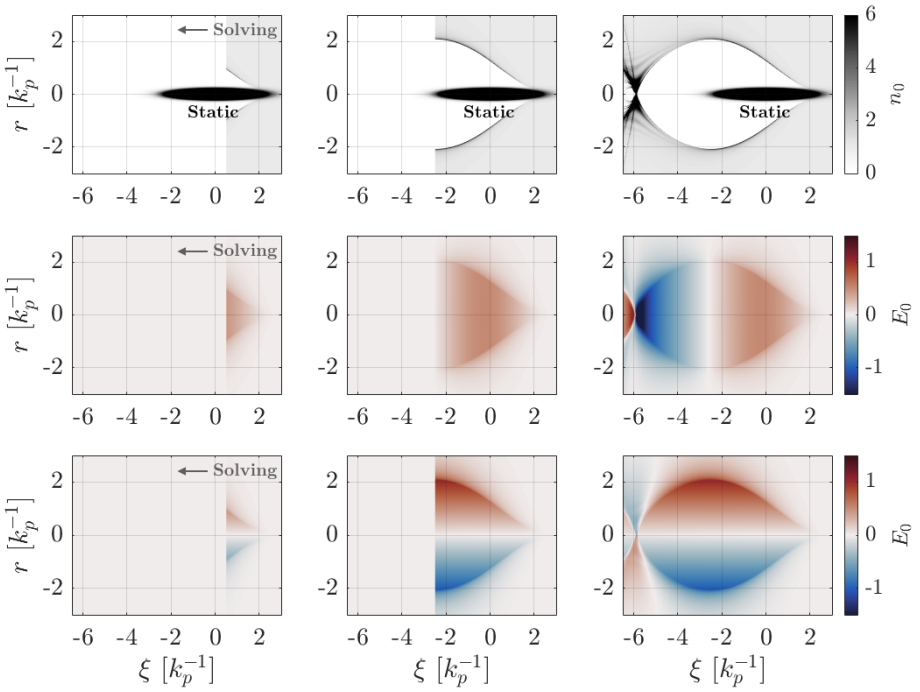

Figure 2.3.2: Demonstration of a single time-step of a quasi-static PIC code. The current from the bunch (black) is deposited on the grid and then frozen. The response of plasma electrons and ions is then calculated from its leading edge to the rear of the simulation window, $\xi = 3 \rightarrow -6.5\,k_p^{-1}$. The upper row represents the normalised plasma electron density, the middle row shows the calculated normalised longitudinal field and the lower row shows the normalised transverse field.

expensive, even when using the quasi-static approach. In this case, the axisymmetric quasi-static PIC code LCODE [118, 119] is used instead. By assuming cylindrical symmetry of the system, the response of only a single radial slice of plasma electrons needs to be calculated along ξ as opposed to the large number of x- and y-slices within a three-dimensional simulation such as those performed in HiPACE. This comes at the cost of the ability to simulate the effects of, for example, asymmetric focusing or beam tilts that can seed the development of transverse instabilities such as hosing [120] over the plasma length. However, within the AWAKE experiment, care is taken to ensure this destructive mode that can lead to the transverse break-up of the beam is not seeded [61, 121].

2.3.2 *PIC simulations on extended timescales*

When considering high-repetition-rate operation of plasma wakefield accelerators, the extended evolution of the plasma over nanosecond to microsecond timescales following the wakefield excitation process must be well understood. Given their success in modelling the initial interaction between the bunch and plasma, it might be expected that PIC simulations could just be performed over an extended timescale to learn how the plasma evolves

after the wakefield. However, there are many issues with this approach. Firstly, a full three-dimensional PIC simulation of only the leading plasma period is already incredibly computationally expensive, so much so that alternative methods such as the quasi-static approximation had to be introduced as discussed in the previous section. Simulating timescales many orders of magnitude larger than this is clearly infeasible.

Use of the quasi-static approximation potentially offers a more suitable method for extended timescale simulations by simply increasing the length of the simulation window. By doing this, the evolution of plasma electrons and ions continues to be self-consistently calculated until the end of the simulation window is reached. However, consider typical FLASHForward operating parameters; the plasma density at interaction is typically $n_e \sim 10^{16}\,\mathrm{cm}^{-3}$ and hence a single plasma wavelength represents a timescale of approximately 1 ps. Therefore, to simulate a single nanosecond the simulation must be extended by a factor of approximately 1000. This is potentially feasible, particularly if using the simplifying assumption of cylindical symmetry or by making use of the speed-up associated with performing such calculations on dedicated graphics processing unit (GPU) nodes that are becoming readily available at HPC facilities. Many state-of-the-art PIC codes are currently being developed or updated such that they are compatible with GPU clusters [122, 123]. However, this does not seem to offer an obvious route towards simulating many nanoseconds or microsecond timescales.

The PIC routine also suffers from numerical scattering and heating effects that, over extended timescales, can lead to unphysical effects being observed within simulations [124]. Numerical scattering is due to the use of discrete macroparticles to represent the plasma rather than a smooth particle distribution function. Small fluctuations in the numbers of macroparticles within a given cell can lead to large changes in its total charge and hence the corresponding calculated field values. During the particle push, these significant changes in the value of the field cause macroparticles to experience an instantaneous force which, when integrated over an extended timescale, leads to heating of the particle ensemble [125]. Numerical heating occurs if the Debye length, λ_D as defined in Eq. (1.2.1), is not sufficiently resolved by the grid. In this case, the dispersion relation is altered and high-frequency modes become aliased to low frequencies, resulting in unphysical heating [126]. This continues until the temperature is such that the Debye length is resolved by the grid [124]. The growth rates of these effects can be reduced by using more macroparticles per cell or applying higher-order particle shapes and field interpolation schemes [124, 127]. However, both solutions significantly increase the computational effort and only slow the growth-rate of the numerical effects rather than removing them.

Additionally, over extended timescales, other important physical processes must be considered within simulations. For example, the standard PIC routine (as in Figure 2.3.1) is collisionless as collisional timescales within the plasma ($\sim$ns [128]) are much longer than those typically associated with wakefield excitation. However, with regards to high-repetition-rate operation, collisions are crucial, especially with regards to how energy from the wakefield is dissipated as is discussed in Chapters 8 and 9. Macro-particle collision routines are already

implemented within some PIC codes, for example EPOCH [129] uses a binary collision operator based on the approach of Sentoku and Kemp [130], however large numbers of macroparticles per cell [$\mathcal{O}(100)$] can be required to achieve convergence [129], vastly increasing computational effort.

It is not clear whether the conventional PIC method can be applied to simulating extended timescales, particularly in its current form. Once the initial wakefield excitation has dissipated, it is possible that further evolution of the system could be handled by hybrid-PIC codes (e.g. see Refs. [131–134]) where, typically, the electron species are treated using a fluid model and ions are represented by macroparticles in a kinetic approach. The use of a fluid model relaxes the restrictions on the sizes of time-steps and grid cells while ionisation and collisional effects can be calculated using rate equations [128]. However, defining how and when to move from a fully-kinetic PIC routine to a hybrid approach and then possibly a full fluid description poses significant challenges.

Part II

EXPERIMENTAL FACILITIES

Part II outlines the two experimental facilities, and recent experimental highlight at each, where measurements discussed in this thesis were performed:

- The Advanced Wakefield Experiment (AWAKE) at CERN.
- The FLASHForward Experiment at DESY.

3

THE AWAKE EXPERIMENT

The Advanced Wakefield Experiment (AWAKE) is a proof-of-principle plasma wakefield acceleration experiment based at CERN [59, 61]. It uses a proton bunch produced in the Super Proton Synchotron (SPS) to drive a wakefield in a laser-ionised 10 m-long Rb plasma. Witness electrons are externally injected into the wakefield and are captured and accelerated.

This chapter first motivates the use of high-energy proton bunches as PWFA drive bunches and then introduces self-modulation, the mechanism through which long bunches are transformed into trains of micro-bunches that can be effective wakefield drivers. A discussion of how the wakefield itself evolves during the self-modulation process follows, with particular attention paid to the effect on the proton micro-bunches and witness electrons captured within the wakefield. Finally, some experimental highlights achieved during Run 1 of the AWAKE experiment ($2016-2018$) are discussed; namely the first demonstrations of self-modulation of a proton bunch by a plasma and the acceleration of electrons in the resulting wakefield.

3.1 PROTON BUNCH DRIVERS

The primary motivation for using a proton bunch as a PWFA drive beam is due to the energy stored within the bunch. The drive bunches used in the electron-driven FLASHForward experiment typically have $\mathcal{O}$(nC) charges with energies around 1 GeV and hence have joule-level energy stored within the bunch [57]. As the bunch drives the wakefield it is continually transferring energy to the plasma and is acted on by the decelerating field driven within its vicinity. Assuming a moderate decelerating field of magnitude $1\,\mathrm{GVm}^{-1}$, an electron in the drive bunch will become depleted within 1 m of propagation, severely limiting the possible accelerating length. To overcome this limitation and allow acceleration of witness particles up to tens of GeV and beyond, such that they are of interest for future particle collider applications, multiple accelerating stages must be coupled together with a new drive bunch used in each accelerating stage [54]. This presents a significant challenge as witness bunches produced by plasma accelerators typically have large [$\mathcal{O}$(mrad)] divergence, due to strong plasma focusing, and percent-level energy spread and, as such, require lengthy, complex beam transport schemes for in- and out-coupling from plasma stages [135, 136]. To date, only one experiment has successfully demonstrated coupling of two independent (laser-driven) plasma acceleration stages [137]. In this case, approximately one-third of the witness charge was captured in the second stage, indicating the inherent complexity of efficiently combining multiple plasma stages. The requirement to stage multiple accelerating sections extends the

overall length of the accelerator and reduces the intrinsic advantage of the large accelerating gradients produced by plasma-based systems.

In contrast, proton beams used in large-scale particle physics experiments have much higher intensities and energies. For example, a proton bunch produced by the SPS used in AWAKE typically has a bunch population of approximately 3×10^{11} protons and an energy of 400 GeV. This corresponds to a stored energy of 19.2 kJ and could drive wakefields over hundreds of metres before becoming depleted. Therefore, using proton bunches in PWFA experiments offers the ability to extend the length of a single acceleration stage such that the coupling of multiple stages is not required. Previous simulations of LHC-type proton bunches ($N_p \sim 1.15 \times 10^{11}$, $E \sim 7$ TeV, $E_s \sim 130$ kJ) have indicated TeV-scale witness electron energies could be achieved in a single km-scale plasma stage using the AWAKE scheme [60].

Proton bunches therefore appear to be suitable PWFA drive bunches for accelerating witness electrons to high energies in a single extended plasma stage. However, an additional complication is introduced due to the length of high energy proton bunches produced by the SPS. In order to avoid longitudinal emittance blow-up and maintain the design transverse emittance ($\varepsilon_N = 3.5$ mm·mrad) and beam size ($\sigma_r = 200 \pm 20$ µm) at the entrance of the plasma cell, the proton bunches produced in the SPS and used for AWAKE experiments had a length of approximately 8 cm RMS [138]. In order to optimally excite the wakefield, drive beams used in PWFA experiments must be short with respect to the plasma wavelength. The beam RMS length, σ_z, should ideally satisfy the relation $\sigma_z \leqslant \lambda_p/2$. The maximum plasma density that satisfies this constraint for a bunch length of 8 cm is $n_e = 4.4 \times 10^{10}$ cm^{-3}, resulting in a cold wave-breaking field of 20 MVm^{-1} and hence offering no increase in accelerating field over traditional RF cavities. However, it is still possible to use the proton bunches to drive wakefields in higher density plasma, and hence access larger accelerating fields, thanks to the process of *self-modulation*.

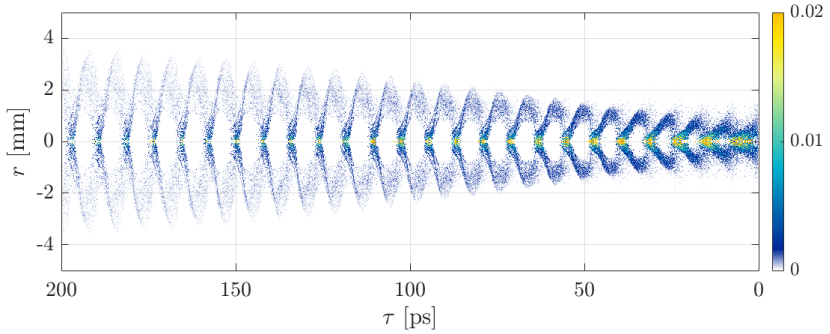

Figure 3.1.1: Simulated self-modulated proton bunch density normalised to the plasma density, n_0, after a propagation distance of 10 m in a plasma of background density $n_0 = 2 \times 10^{14}$ cm^{-3}. Dense micro-bunches have formed on-axis while defocused protons can be observed at larger radii. The initial radial bunch size was $\sigma_r = 200$ µm.

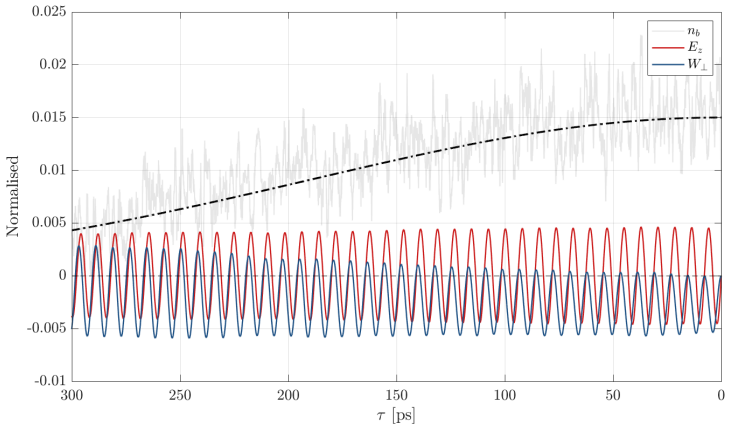

Figure 3.2.1: Seed wakefields for typical AWAKE-SPS proton bunch parameters. Shown are the on-axis longitudinal wakefield (E_z) in red, off-axis transverse wakefield ($W_\perp$) in blue, and the on-axis bunch density (n_b) in grey. The values are normalised to their relevant parameters: $E_z, W_\perp$ are normalised to the cold wave-breaking field, $E_0 = m_e c\omega_p/e = 1.36\,\mathrm{GVm}^{-1}$; n_b is normalised to $n_0 = 2\times10^{14}\,\mathrm{cm}^{-3}$. On-axis values are measured at $r = 0.1\,k_p^{-1}$, off-axis values are measured close to the transverse beam size, $r = 0.5\,k_p^{-1}$. The black dashed line represents a Gaussian fit to the unmodulated bunch density.

3.2 SELF-MODULATION

Self-modulation is a natural plasma response where the fields induced by the interaction between a long ($\sigma_z \gg \lambda_p$) driver and a plasma act to modify the structure of the driver itself such that the fields it drives can be amplified. It was first discussed within the context of laser-driven wakefield acceleration as a method for increasing the achievable accelerating gradient via resonant wakefield excitation driven by modulated laser pulses [139–141]. The density perturbation induced in a plasma by the ponderomotive force of a long, high intensity laser pulse creates regions of enhanced focusing or diffraction of the laser pulse, depending on the local density gradient, and acts to modulate the pulse into a series of beamlets with period equal to the plasma wavelength. These beamlets can then resonantly drive larger amplitude wakefields. A similar response is also observed for charged particle beams and provides the mechanism used within AWAKE to permit the use of long proton bunches as PWFA drive beams. An overview of this process follows and is described in detail in Refs. [62, 120, 142]. Figure 3.1.1 shows an example of a simulated self-modulated proton bunch.

3.2.1 *Seed wakefields*

When a long charged particle beam enters a plasma, its space charge force drives charge separation between electrons and ions and excites low amplitude ($\sim\mathrm{MVm}^{-1}$) wakefields, known as *seed* wakefields. An example is shown in Figure 3.2.1 for typical SPS-AWAKE proton bunch parameters outlined in Section 3.1 on a background plasma density of $n_0 = 2\times10^{14}\,\mathrm{cm}^{-3}$. The extrema of the longitudinal wakefield amplitude is approximately constant along the

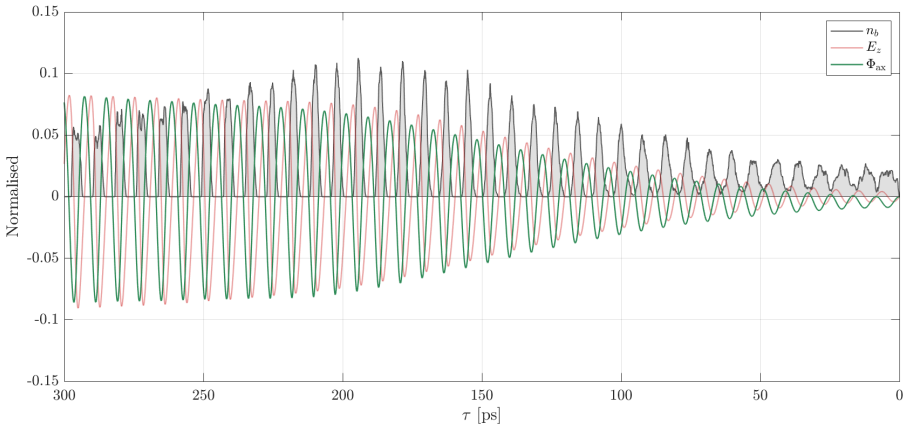

Figure 3.2.2: Resonantly excited longitudinal wakefield (E_z, red) and associated wakefield potential (Φ_{ax}, green) driven by a self-modulated proton bunch (n_b, grey) after a propagation distance of 3 m. The values are normalised to their relevant parameters: E_z, Φ_{ax} are normalised to the cold wave-breaking field, $E_0 = m_e c \omega_p / e = 1.36\,\mathrm{GVm}^{-1}$; n_b is normalised to $n_0 = 2 \times 10^{14}\,\mathrm{cm}^{-3}$. Values are measured at $r = 0.1\,k_p^{-1}$.

entire proton bunch length and equal to approximately $6\,\mathrm{MVm}^{-1}$ while the transverse wakefield, $W_\perp$, is slightly larger in amplitude and initially globally focusing for protons. Electron oscillations naturally occur at the resonant frequency of the plasma in the linear wakefield regime and hence a periodic wakefield structure is excited within the long drive bunch. The low-amplitude transverse wakefield acts on the beam particles and periodically focuses and defocuses alternating regions of the bunch with period equal to the plasma wavelength. This results in modulation of the bunch density profile with micro-bunches of increased density formed on-axis while defocused protons move away from the axis and form a diverging "halo". The micro-bunches are short with respect to the plasma wavelength ($\sigma_z < \lambda_p/2$) and hence are able to more efficiently drive wakefields within the plasma. This in turn increases the transverse wakefield amplitude, further focusing micro-bunches and increasing their density in a positive feedback effect. The micro-bunches are naturally separated by the plasma period and hence their wakefields resonantly combine resulting in a rapid exponential growth of the longitudinal wakefield amplitude driven by the micro-bunches. This can exceed GVm^{-1} at the nominal AWAKE plasma density of $7 \times 10^{14}\,\mathrm{cm}^{-3}$ when using typical SPS proton bunch parameters. An example of the modulated proton bunch density is demonstrated in Figure 3.2.2 after the proton bunch has propagated 3 m through the plasma. The on-axis micro-bunches (grey) are clearly defined with peak densities an order of magnitude larger than the initial unmodulated bunch density and are separated by the plasma wavelength. The longitudinal wakefield (red) grows along τ as it is resonantly excited by the proton micro-bunches, reaching a maximum amplitude of approximately $0.095\,E_0 = 130\,\mathrm{MVm}^{-1}$. The longitudinal wakefield amplitude can also be observed to transition from a sinusoidal oscillation at small τ towards a sawtooth profile at larger τ, indicating transition from the linear regime to the quasi-linear regime along the length of the micro-bunch train.

3.2.2 *Wakefield phase evolution*

To discuss the evolution of micro-bunches and their effect on the wakefield (and vice versa) it is useful to re-introduce the wakefield potential Φ, linked to the longitudinal and transverse wakefield components via:

$$E_z = -\frac{\partial \Phi}{\partial \xi}, \quad E_r - B_\theta = -\frac{\partial \Phi}{\partial r}, \tag{3.2.1}$$

which allows easier visualisation of the regions of the wakefield in which proton micro-bunches will form and remain confined. The force on an axially-moving particle can be calculated from $-\nabla\Phi$ and therefore proton-confining potential wells correspond to regions of negative Φ. This can be observed in Figure 3.2.2 where the proton micro-bunches are clearly aligned with the potential wells. The link to the longitudinal wakefield amplitude is also demonstrated, with the relative phase of the potential being delayed by $\pi/2$ with respect to the longitudinal field.

Consider the effect of the wakefield potential driven by a single micro-bunch on a background wakefield potential of constant amplitude; it has the ability to both amplify the wave and advance its phase, depending on the relative phase of the micro-bunch. Figure 3.2.3 demonstrates this for three different cases:

(a) The micro-bunch is placed at the zero-crossing of the potential (at the maximum decelerating gradient). In this case, the micro-bunch purely amplifies the wave.

(b) The micro-bunch is placed at the potential well maximum. The phase is advanced while the amplitude remains similar.

(c) The micro-bunch is placed in an intermediate position. The wave is amplified while its phase is simultaneously advanced.

These three cases help to illustrate the modification of the phase of the wakefield during the self-modulation process. The micro-bunches start to form from the transverse seed wakefields in the locations of the potential wells, akin to case (b) above. As the bunch density grows the potential associated with the bunch, $\Delta\Phi_{ax}$ in Figure 3.2.3, similarly grows and the phase advance supplied by the micro-bunch modifies the relative phase of the following micro-bunch. The potential slowly shifts backwards with respect to the micro-bunches (its phase velocity is reduced), moving through case (c) and into case (a). At this point, the micro-bunch is in a position to maximally amplify the wave and the longitudinal wakefield amplitude grows exponentially. An example of the evolution of the maximal longitudinal wakefield driven by a self-modulated proton bunch is shown in Figure 3.2.4, with fast amplification of the wakefield amplitude occurring over a distance of $z = 1 \rightarrow 4\,\mathrm{m}$.

However, as can be seen for case (a) in Figure 3.2.3, while the micro-bunch is in a position to maximally amplify the wave its leading edge extends into a region of positive wakefield potential. These protons are no longer confined and the transverse wakefield acts to accelerate

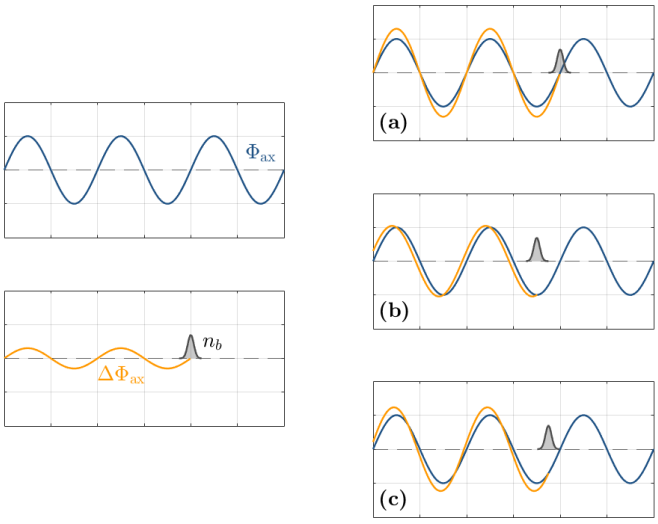

Figure 3.2.3: Simple cartoon representation of the phase-dependent wakefield modification induced by a micro-bunch. Left: background longitudinal wakefield potential, Φ_{ax}, (top) and the additional wakefield potential driven by a bunch of density n_b, $\Delta\Phi_{ax}$, (bottom) drawn separately. Right: linear combination of the additional wakefield potential when the bunch is (a) in phase, (b) exactly $\pi/2$ delayed and (c) exactly $\pi/4$ delayed. Recreated from Ref. [142].

them away from the axis. This reduces the micro-bunch density, shifts the centre of the micro-bunch backwards, and the ability of micro-bunches to amplify the wave is significantly reduced. The maximal longitudinal wakefield amplitude saturates (Figure 3.2.4, $z \sim 6\,m$) before decaying as a significant number of micro-bunches lose the majority of their charge.

As the relative phase of a micro-bunch with respect to the wakefield potential depends on that of all preceding micro-bunches, both the rate and magnitude of the wakefield phase evolution varies along the length of the proton bunch. This is demonstrated in Figures 3.2.5 and 3.2.6 which show the evolution of both the on-axis bunch density and the wakefield potential over the entire plasma length for two temporal regions: Figure 3.2.5 corresponds to the 100 ps-long region directly behind the head of the proton bunch while Figure 3.2.6 corresponds to a 100 ps region around the location of the maximal longitudinal wakefield amplitude ($\tau \sim 250\,ps$). The black contours in the upper panels of these images represent the zero-crossings of the wakefield potential and so define the limits of alternating regions of positive and negative wakefield potential. These therefore indicate the regions where protons are expected to remain confined within the wakefield.

During the initial interaction of the head of the bunch (Figure 3.2.5, $z < 1\,m$) the wakefield potential is mostly negative and micro-bunches begin to be formed. As the micro-bunch density grows and the wakefield amplitude increases, the phase shift induced by the wakefields of the micro-bunches begins to be observed and the potential shifts backwards with respect to the co-moving frame. The positive regions of the wakefield potential grow in amplitude

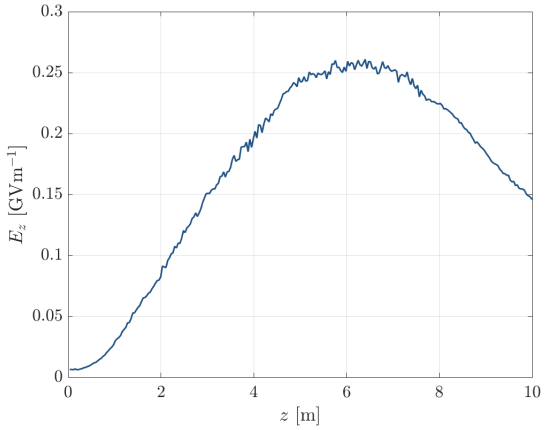

Figure 3.2.4: Evolution of the maximal longitudinal wakefield amplitude driven by a self-modulated proton bunch in a background plasma density of $n_0 = 2 \times 10^{14}$ cm^{-3}.

and extent and cause significant defocusing of bunch particles. Micro-bunches become more clearly defined by this process and the corresponding growth in wakefield amplitude further focuses the bunches, increasing the micro-bunch density to its maximum at around $z \sim 4$ m. The phase velocity of the wakefield potential remains slower than that of the micro-bunches causing the heads of the micro-bunches to move into defocusing regions of the wakefield and charge to be lost.

This effect can be seen even more clearly in Figure 3.2.6 which shows the same evolution of the on-axis bunch density and wakefield potential, but looks at the region around σ_z behind the seed position where maximal wakefields amplitudes are typically observed in simulation [120, 142]. In this case, the micro-bunch peak density is observed slightly earlier during the propagation ($z \sim 3$ m). This is because the wakefield amplitude is resonantly excited and leads to larger amplitude transverse wakefields further along the proton bunch earlier in the propagation length. After the peak micro-bunch density is reached, a large shift in the wakefield phase can be observed ($z \sim 4$ m) causing the leading edge of the micro-bunches to move into regions of positive wakefield potential. This charge is gradually defocused and lost from the micro-bunches by $z \sim 8$ m, corresponding to the decrease in maximal wakefield amplitude observed at this propagation distance in Figure 3.2.4. By the end of the interaction length ($z = 10$ m), the continual change in wakefield phase in this region has caused the micro-bunches to be almost entirely defocused.

The evolution of the wakefield phase both along the proton bunch and the plasma length causes continuous evolution of the micro-bunch density as demonstrated in Figure 3.2.7. The upper panel depicts the on-axis bunch density evolution while the lower panel shows the proton bunch density distribution at $z = 3$ m (upper) and $z = 10$ m (lower). The first micro-bunches are formed at the rear of the proton bunch and the density of all micro-bunches grows as the wakefield amplitude grows ($z = 1 \rightarrow 4$ m) and vice versa in a positive feedback effect. The shift in wakefield phase experienced by a micro-bunch is a cumulative effect from all preceding micro-bunches and micro-bunch charge is lost from the rear of the bunch train

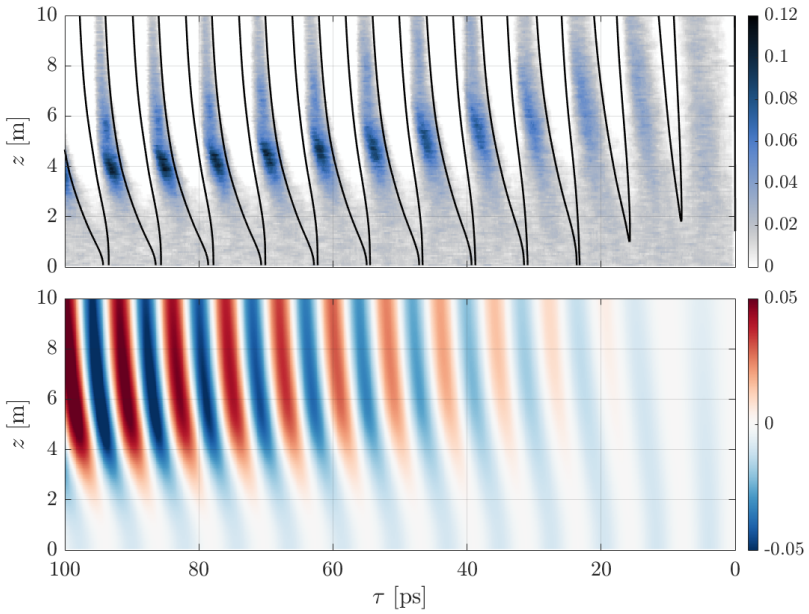

Figure 3.2.5: Upper: evolution of the on-axis proton bunch density at the leading edge of the bunch, $\tau = 0 \rightarrow 100$ ps. Black contours represent zero-crossings in the wakefield potential. Lower: Corresponding evolution of the on-axis wakefield potential. The bunch density is normalised to the plasma density, n_0, while the wakefield potential is normalised to the cold wave-breaking field, E_0.

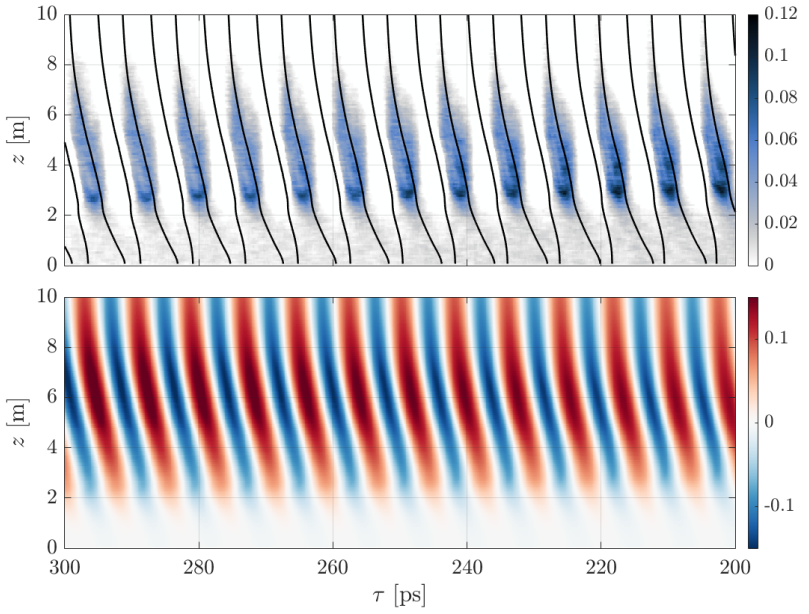

Figure 3.2.6: Upper: evolution of the on-axis proton bunch density around the peak of the wakefield amplitude, $\tau = 200 \rightarrow 300$ ps. Black contours represent zero-crossings in the wakefield potential. Lower: Corresponding evolution of the on-axis wakefield potential. The bunch density is normalised to the plasma density, n_0, while the wakefield potential is normalised to the cold wave-breaking field, E_0.

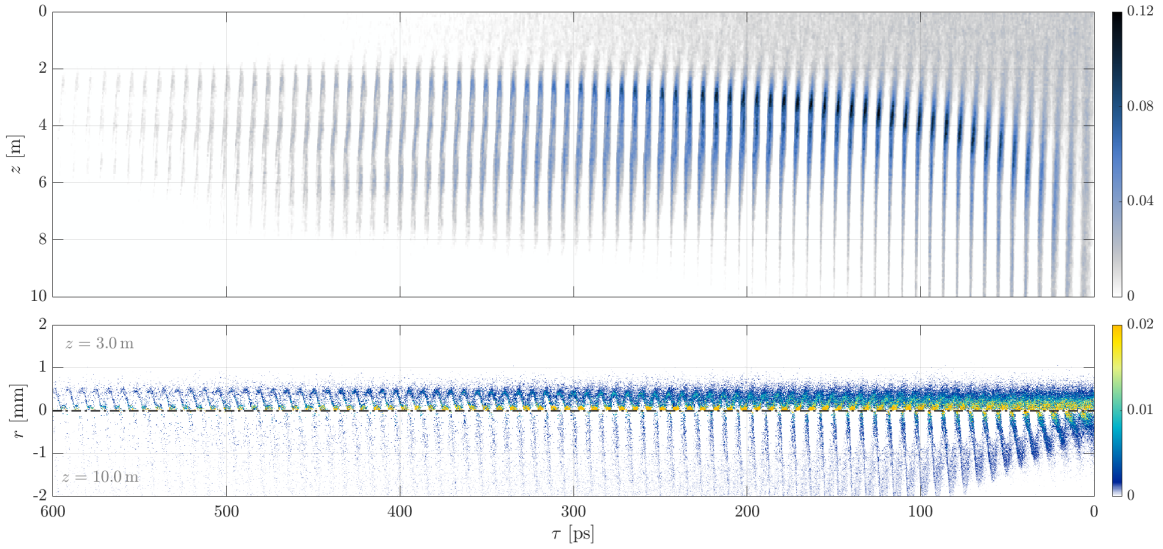

Figure 3.2.7: Upper: evolution of the on-axis proton bunch density measured at $r = 0.1\,k_p^{-1}$ in the co-moving frame over the entire plasma length. Lower: Comparison between the transverse proton bunch profile at $z = 3\,m$ (top) and $z = 10\,m$ (bottom). The colour axis limits are chosen to highlight both the on-axis micro-bunches (yellow) and the defocused protons (blue).

initially. However, as the wakefield amplitude grows, the phase shift is significant along the majority of the bunch length and by the end of the plasma length, high density micro-bunches are only observed on-axis for $\tau \lesssim 200\,ps$. This can be observed in the lower panel of Figure 3.2.7 where at $z = 10\,m$ the on-axis bunch density is low and regions of strongly defocused protons can be observed far off-axis. In comparison, at $z = 3\,m$, a train of micro-bunches can be observed on-axis along the length of the entire proton bunch.

3.3 WITNESS ELECTRONS

The aims of Run 1 of the AWAKE experiment were to first demonstrate self-modulation of the proton bunch as discussed in Section 3.2, and then to accelerate electrons that had been injected into the wakefields driven by the modulated bunch train. The experiment was designed such that a short ($\sigma_z \sim \lambda_p$), $\mathcal{O}(100\,pC)$ electron bunch could be injected at the entrance of the vapour source alongside the proton bunch with variable trajectory and timing with respect to the protons. However, injection at the start of the vapour source meant witness electrons that were captured in the wakefield also experienced the evolution of the phase of the wakefield potential during the self-modulation process. Section 3.2.2 showed the deleterious effect of this on the proton micro-bunches and similar effects are also experienced by captured witness electrons and will be discussed here.

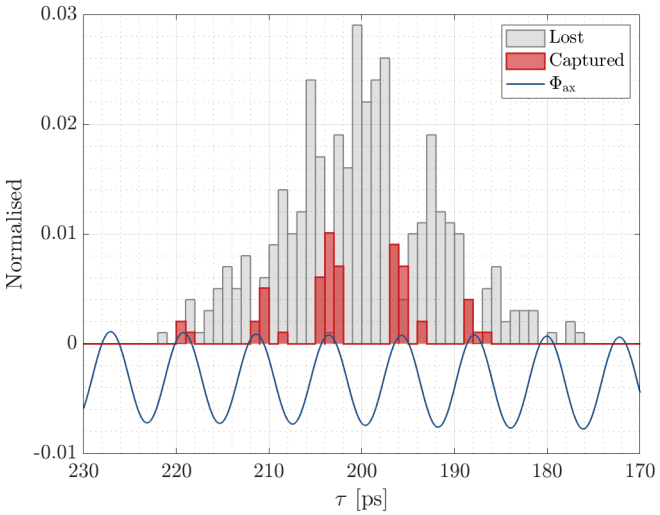

Figure 3.3.1: Histograms of the simulated injection positions of captured (red) and lost (grey) witness electrons in the co-moving frame for an injection delay of $\tau = 200\,\text{ps}$. The corresponding on-axis seed wakefield potential, Φ_{ax}, is plotted in blue. The histogram bins have a width of 1 ps and are normalised to the total number of witness electrons injected in simulation; the wakefield potential is normalised to E_0.

3.3.1 *Capture at injection*

When witness electrons are initially injected they experience the seed wakefields driven by the unmodulated proton bunch. As shown in Figure 3.2.1, the seed transverse wakefields are globally focusing for protons at its head, leading to squeezing of the bunch and the onset of micro-bunch formation. Consequently, the seed wakefields are globally defocusing for electrons and capture within the wakefields is difficult. It isn't until further behind the head of the proton bunch ($\tau \gtrsim 50\,\text{ps}$) that the transverse seed wakefields have regions that are focusing for electrons and the corresponding wakefield potential is positive and allows electrons to be trapped.

Figure 3.3.1 demonstrates this effect for 5000 simulated test-electrons that are injected close to the axis at the entrance of the vapour source ($z = 0\,\text{m}$) at a delay of $\tau = 200\,\text{ps}$ with respect to the leading edge of the proton bunch. The witness electrons are longitudinally gaussian-distributed with a bunch length of $\sigma_z = 8\,\text{ps}$ and initial energy of 19 MeV. Two histograms showing the injection position of the witness electrons are plotted; one represents witness electrons that are captured in the wakefield at the initial injection position and propagate within the wakefield over the entire 10 m plasma length (red); the other represents witness electrons that are defocused and lost from the wakefield (grey), either at the start of the simulation or later during the development of self-modulation. Also plotted is the on-axis seed wakefield potential at $z = 0\,\text{m}$, demonstrating the positions of positive, electron-focusing regions and negative, electron-defocusing regions. Witness capture is achieved in the limited regions of positive wakefield potential as expected, while the vast majority of electrons are defocused by the seed wakefields and lost.

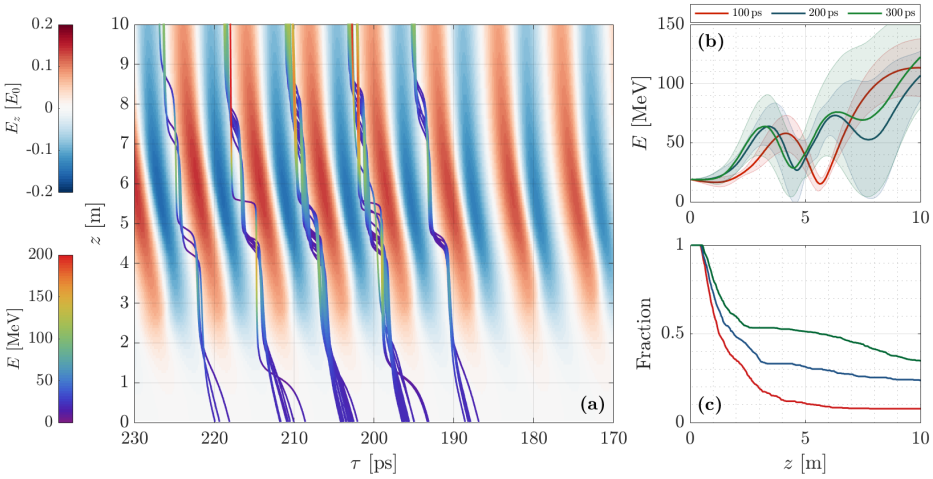

Figure 3.3.2: Simulated captured witness evolution over plasma length. (a) Path of randomly-selected captured witness electrons in the co-moving frame. The color scale represents the energy of a particle at a given z-position. The background shows the evolution of the longitudinal wakefield amplitude, E_z, normalised to E_0. (b) Average captured witness energy and energy spread as a function of propagation distance, z, for three different injection delays; 100 ps (red), 200 ps (blue) and 300 ps (green). (c) Fraction of witness particles that remain within the plasma boundary as a function of propagation distance for the three injection delays.

3.3.2 *Wakefield phase evolution*

However, as can be seen in Figure 3.3.1, witness injection at a position of positive wakefield potential is not a guarantee of successful capture and transport over the entire plasma length. As discussed in Section 3.2, the phase of the wakefield potential continuously evolves over the plasma length as the proton bunch undergoes self-modulation, causing witness electrons that are initially trapped at injection to move into decelerating, defocusing regions where they can be lost from the wakefield.

Figure 3.3.2 demonstrates the simulated evolution of captured witness electrons over the plasma length for an injection delay of $\Delta\tau = 200$ ps. In the co-moving frame (Figure 3.3.2(a)), the electrons are initially accelerated in the low amplitude wakefield until $z \sim 3$ m when a large shift in the phase of the wakefield is observed. This corresponds to the peak micro-bunch density being reached in Figure 3.2.6 and fast growth in the wakefield amplitude in Figure 3.2.4 as noted previously. At this point, witness electrons are observed to undergo dephasing as they move into decelerating regions of the wakefield and a corresponding drop in the mean witness energy is seen in Figure 3.3.2(b). As the electrons are decelerated their velocity drops and they slip further backwards in the co-moving frame into the accelerating phase of the wakefield once again. This re/dephasing continues to happen over the plasma length as the wakefield phase evolves and is not resolved until the point of *self-modulation saturation*. This corresponds to the position at which the wakefield phase stabilises with respect to the micro-bunches and is observed at approximately $z \gtrsim 7.5$ m at this density

($n_0 = 2 \times 10^{14}\,\mathrm{cm}^{-3}$); after this point, witness electrons can be continuously accelerated in the wakefield.

Due to the difference in rate of change of the wakefield phase along the proton bunch, electrons injected at different temporal delays undergo re/dephasing at different z-positions. This is indicated in Figure 3.3.2(b) where the evolution of the mean energy (solid line) and energy spread (shaded region) of all captured witness electrons is plotted as a function of the propagation distance for three different injection delays; 100 ps (red), 200 ps (blue) and 300 ps (green). As shown in Figures 3.2.5 and 3.2.6, both the amplitude of the wakefield and its rate of change of phase is much reduced towards the leading edge of the proton bunch due to the small number of micro-bunches resonantly driving the wakefield. Therefore, for the shortest injection delay ($\Delta\tau = 100\,\mathrm{ps}$) significant shifts in the phase of the wakefield are not observed until larger propagation distances, and the re/dephasing of witness electrons occurs later in the plasma ($z \sim 5\,\mathrm{m}$). In addition to this, the magnitude of the phase shift is reduced at short injection delays and stable acceleration occurs after the first re/dephasing cycle. In contrast, at larger injection delays ($\Delta\tau = 200, 300\,\mathrm{ps}$) two re/dephasing cycles are observed in the witness energy evolution and consistent witness acceleration occurs after $z \gtrsim 8\,\mathrm{m}$.

Figure 3.3.2(c) shows the fraction of witness electrons that remain within the plasma boundary as a function of propagation distance for the three different injection delays discussed previously. In all three cases, the vast majority of witness electrons are lost within the first few metres of propagation. This is due to the nature of the transverse wakefields over the first few metres as the micro-bunches form. As can be seen in Figure 3.2.1 the transverse wakefield is initially globally defocusing for electrons near the head of the proton bunch. However, as the injection delay is increased and the wakefield amplitude grows along the proton bunch, the transverse wakefield develops increasingly large electron-focusing regions and the electron-capture phase-space grows. This means a larger fraction of injected witness electrons follow trajectories that can result in capture within the wakefield as the injection delay is increased. After the first few metres the rate of loss of witness electrons decreases significantly as the wakefield amplitude has grown sufficiently to enable stable propagation. However, a small fraction are lost during the dephasing cycles as captured electrons move into decelerating, defocusing regions of the wakefield.

In general, injecting witness electrons at the entrance of the plasma source is challenging due to the defocusing nature of the seed wakefields and those electrons that are captured experience continuously evolving wakefields due to the development of the proton bunch self-modulation. This leads to low capture efficiencies (Figure 3.3.1) and witness bunches with large energy spreads (Figure 3.3.2). This has motivated a fundamental redesign of the witness injection region for AWAKE Run 2; the vapour source will instead be split into two separate sections with witness injection occurring in vacuum between the two [143]. The first cell is used to modulate the proton bunch such that self-modulation saturates, and the modulated proton bunch is then transported to the second cell to be used for accelerating a high-quality witness bunch.

3.4 EXPERIMENTAL HIGHLIGHTS

The above sections described the self-modulation process of a proton bunch and the intricacies of the subsequent capture and acceleration of electrons within the resulting wakefield. During AWAKE Run 1 these two main experimental goals were successfully demonstrated. This section gives a brief overview of the AWAKE experimental setup during Run 1 followed by a description of these key results.

3.4.1 *Experimental overview*

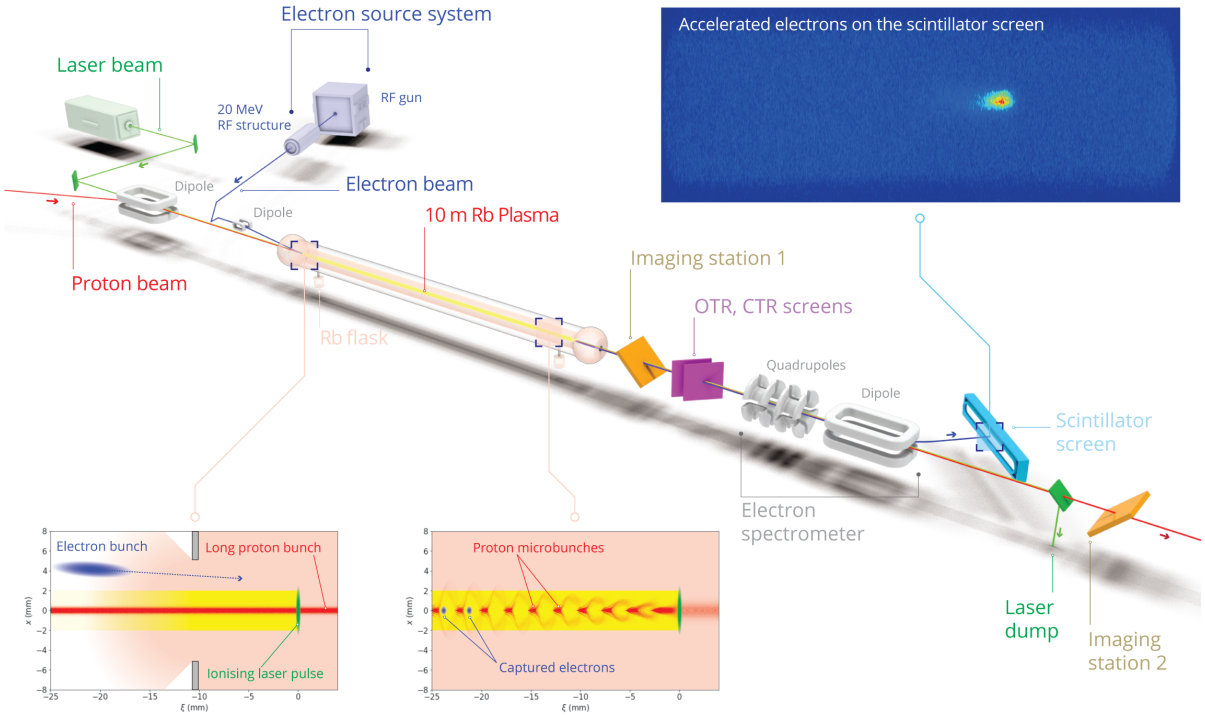

Figure 3.4.1: Overview of the AWAKE beamline setup. *Image reproduced from Ref. [144].*

The proton bunch is extracted from the SPS and transported to the experimental area where it is focused to a transverse size of $\sigma_r \approx 200\,\mu m$ at the entrance of the vapour source and has a transverse emittance of approximately 3.5 mm·mrad. An $\mathcal{O}(100)$ pC witness electron bunch is generated using a frequency-tripled derivative of the main ionising laser pulse to ensure timing stability at the picosecond-level between the electron bunch and seeding laser pulse. The electron bunch is accelerated to 19 MeV using traditional RF cavities and transported along the beam line to the entrance of the vapour source where it is injected into the wakefields driven by the self-modulated proton bunch [145]. The electron bunch trajectory is typically matched to that of the proton bunch with a small vertical offset and injected into the plasma at a tunable angle (~mrad). It is focused to an RMS transverse size of approximately 500 µm at the entrance of the vapour source with a bunch length on the order

of the plasma wavelength, $\sigma_z \approx 8$ ps. Portions of the witness bunch are hence captured within multiple accelerating wakefield buckets and therefore it is not possible to extract detailed longitudinal field structure dependencies below this temporal limit. The relative timing of the ionising seeding laser pulse and witness electron bunch is adjusted using a delay stage in the transport line of the laser pulse to the photocathode. Two quadrupoles are placed 4.48 and 4.98 m downstream of the exit of the plasma to capture and focus the accelerated witness electrons before they are horizontally dispersed by a 1 m-long C-shaped electromagnetic dipole and imaged on a scintillator screen [146]. Light emitted by the scintillator screen is imaged onto an intensified CCD camera. The relationship between the position of an electron in the plane of the scintillator screen and its energy is calculated using the Beam Delivery Simulation (BDSIM) code [147, 148] with measured dipole field maps as input. The energy uncertainty is approximately 2 %, calculated via considerations of the accuracy of the field maps, measurements of the positions of the spectrometer beamline components, and the resolution of the spectrometer imaging system [149].

Calibration of the charge response of the scintillator screen was performed and validated via two independent methods. Firstly, electron beams of variable charge were used to study the scintillator response at the CERN Linear Electron Accelerator for Research (CLEAR) test beam facility [146]. Secondly, beams of mono-energetic electrons were produced by the stripping of high-energy lead ions accelerated in the SPS and transported to the AWAKE experimental area to be imaged by the spectrometer system in-situ [150]. These two complementary measurements permitted calculation of the witness bunch charge from measurements of the integrated light output from the scintillator with an associated uncertainty of 8 %.

The vapour is contained in a 10 m-long cylindrical cell of diameter 40 mm with Rb reservoirs at either end [151]. The cell is heated to provide tunable vapour density. The density of the vapour is monitored by an interferometric measurement at each end of the cell using white light interferometry around the 780 and 795 nm lines of the Rb atom, with an associated uncertainty of 0.5 % [152].

Ionisation of the Rb vapour is achieved using a terawatt-class Ti:sapphire laser. The laser pulse duration is approximately 120 fs with a pulse energy that can be varied from 40 to 450 mJ. The Rayleigh length of the focused laser pulse is 15 m, with a spot size of approximately 2 mm throughout the entire vapour source. It was assumed that the laser singly-ionised the Rb vapour in accordance with previous measurements that demonstrated self-modulation of the proton bunch at a frequency consistent with that of the measured vapour density [153]. The phase of the laser oscillator is locked to the radio frequency of the cavities within the SPS, with synchronisation between the laser pulse and time of arrival of the proton beam measured to be at the picosecond level, thus permitting controlled variations of its arrival with respect to the proton bunch.

3.4.2 *Experimental observation of proton bunch self-modulation*

Proton bunch self-modulation was experimentally observed during AWAKE Run 1 and is detailed in Refs. [153] and [154]. Two complementary diagnostics downstream of the plasma exit were used to confirm self-modulation; a streak camera [155] that observed OTR light emitted by the modulated proton bunch as it passed through an aluminium-coated silicon wafer that allowed temporally-resolved measurements of the fine-scale proton bunch structure to be performed (OTR screen in Figure 3.4.1); and two scintillating Chromox (Al_2O_3:Cr_2O_3) screens with imaging setups that allowed measurement of the time-integrated transverse proton distribution approximately 2 and 10 m downstream of the plasma exit (Imaging station 1 and 2 in Figure 3.4.1).

Streak camera images of the modulated proton bunch captured the bunch structure after it had undergone interaction with the plasma with picosecond resolution [153] as demonstrated in Figure 3.4.2. Discrete Fourier transform (DFT) analysis of the resulting images enabled determination of the modulation frequency of the micro-bunches. Linear wakefield theory predicts that the modulation frequency (f_{mod}) should be equal to the plasma electron frequency (f_{pe}) as it is the action of the wakefield on the long proton bunch that drives the transverse modulation of the bunch:

$$f_{mod} = f_{pe} = \frac{1}{2\pi}\sqrt{\frac{n_{pe}e^2}{m_e\varepsilon_0}}. \tag{3.4.1}$$

Evaluation of the proton bunch modulation frequency from streak camera images over an order-of-magnitude range of Rb vapour densities, $1.3 \times 10^{14}\,\mathrm{cm}^{-3} \leqslant n_{Rb} \leqslant 10.5 \times 10^{14}\,\mathrm{cm}^{-3}$, was consistent with the expected scaling (orange dashed line, Figure 3.4.3). This measurement also demonstrated that the plasma was singly-ionised over the propagation axis of the proton bunch by the ionising laser pulse.

Measurements of the time-integrated transverse proton distribution were used to demonstrate growth in the transverse wakefield amplitude above that of the seed wakefields and thus confirm excitation of the wakefield by the self-modulated proton bunch. The light emitted by each scintillating screen was split and sent to two cameras; one to measure the intense core and the second with a mask in the centre that blocked the light from the core and hence allowed measurement of the defocused "halo" of protons. Example images are shown in Figure 3.4.4. Determination of the maximum radius of defocused protons (green regions in Figure 3.4.4(e)) allows estimation of their divergence angle and hence the average transverse wakefield amplitude driven by the self-modulated proton bunch. This was performed at two different densities, $n_0 = 2.1 \times 10^{14}\,\mathrm{cm}^{-3}$ and $7.7 \times 10^{14}\,\mathrm{cm}^{-3}$, over a range of proton bunch populations, $N_p = (1 - 3.5) \times 10^{11}$, and the transverse wakefield was estimated for each event under two differing assumptions. First assuming that the maximally-defocused protons leave the plasma at the exit of the vapour source ($W_{\perp,av,min}$, black data points in Figure 3.4.5), representing a minimum estimate of the transverse wakefield amplitude. The second assumption is informed by simulations which suggest that strongly defocused pro-

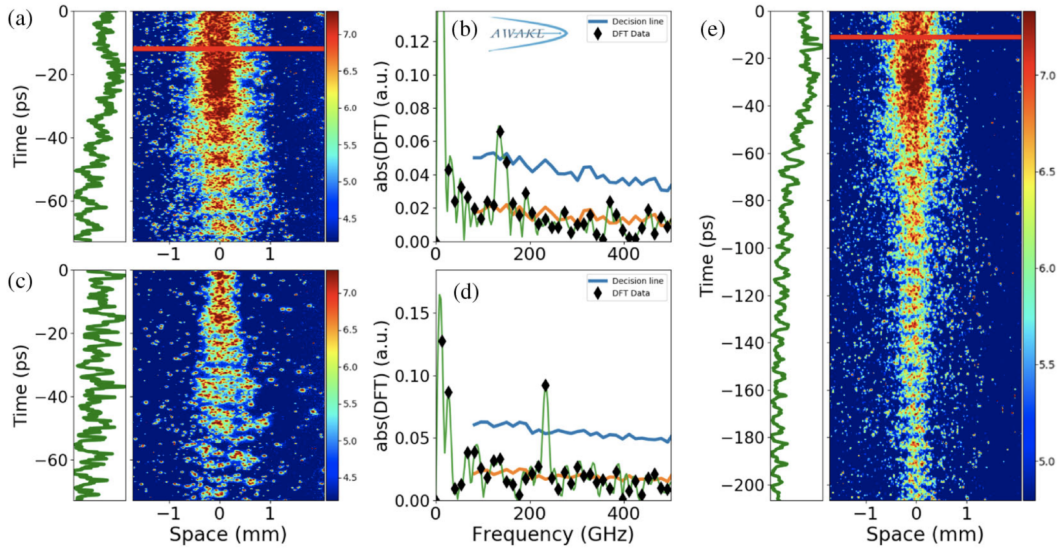

Figure 3.4.2: Streak camera images and associated DFT power spectra of self-modulated proton bunches at various plasma densities. (a) and (b) correspond to a plasma density of $n_{Rb} = 2.457 \times 10^{14}\,cm^{-3}$. (c) and (d) correspond to a plasma density of $n_{Rb} = 6.994 \times 10^{14}\,cm^{-3}$. (e) A full micro-bunch train is observed by combining multiple streak camera images. The plasma density in this case is $n_{Rb} = 2.190 \times 10^{14}\,cm^{-3}$ with an upwards plasma density gradient of 3.4%/10 m. The red lines represent the positions of the ionising laser pulse. *Image reproduced from Ref. [153].*

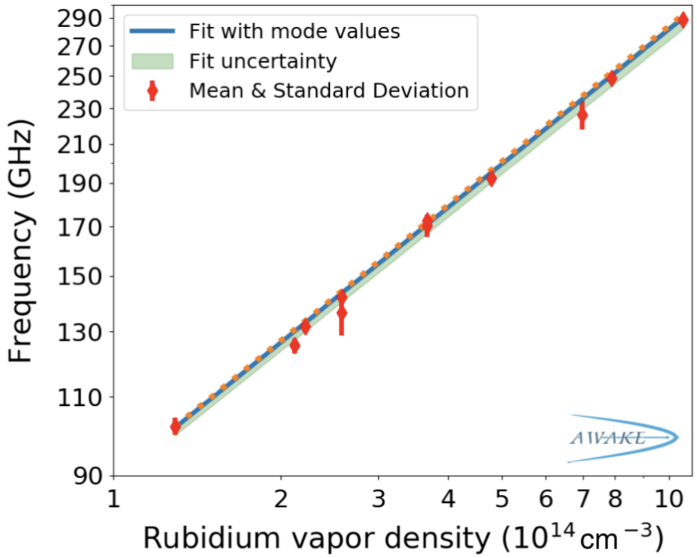

Figure 3.4.3: Measured mean proton bunch modulation frequency as a function of Rb vapour density. The orange line represents the expected scaling under the assumption that the Rb is fully singly-ionised such that $n_{Rb} = n_0$. *Image reproduced from Ref. [153].*

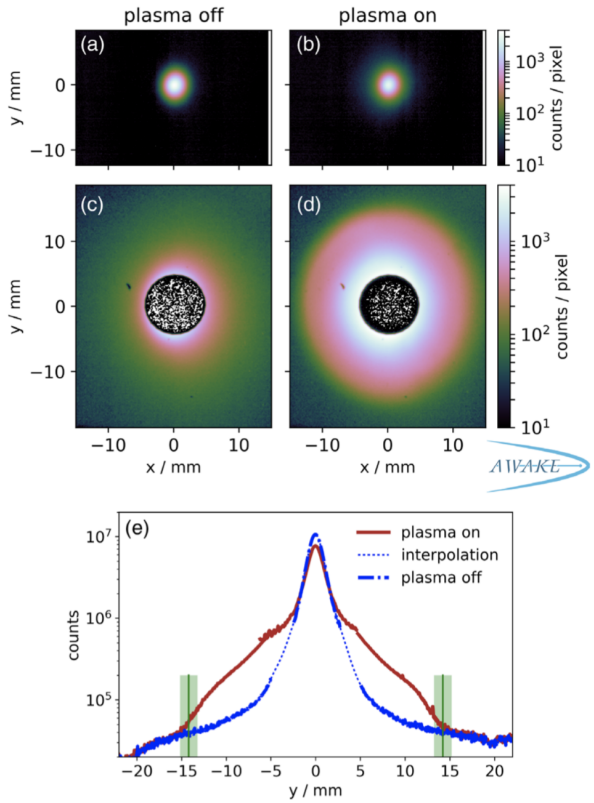

Figure 3.4.4: Time-integrated proton bunch charge distributions demonstrating defocusing of protons in the presence of plasma ($n_0 = 7.7 \times 10^{14}\,\mathrm{cm}^{-3}$). (a) and (b) are measurements of the proton bunch core while (c) and (d) represent "halo" measurements (with central mask). (e) Combined projections of the images from (a) – (d). *Image reproduced from Ref. [154].*

tons gain most of their transverse momentum at a propagation distance of $z \sim 1.5\,\mathrm{m}$ and then radially exit the plasma at $z \sim 4\,\mathrm{m}$ ($W_{\perp,\mathrm{av}}$, blue data points in Figure 3.4.5). This therefore perhaps represents a more realistic estimate of the average transverse wakefield experienced by the maximally-defocused protons. These estimates were then compared to linear wakefield predictions of the seed transverse wakefield driven by the unmodulated proton bunch represented by the red lines in Figure 3.4.5. Both predictions demonstrate average transverse wakefield amplitudes that far exceed the seed wakefield amplitude over a range of plasma densities and proton bunch populations, verification of excitation of the wakefield by the self-modulated proton bunch.

3.4.3 *Acceleration of witness electrons*

Following the demonstration of self-modulation of the proton bunch and wakefield excitation, witness electrons were injected into the plasma to be accelerated in the wakefield; full details of this experiment can be found in Ref. [144].

Electrons were first injected into low density plasma ($n_0 = 1.8 \times 10^{14}\,\mathrm{cm}^{-3}$) in an attempt to aid their capture as the size of the wakefield increases at lower plasma densities

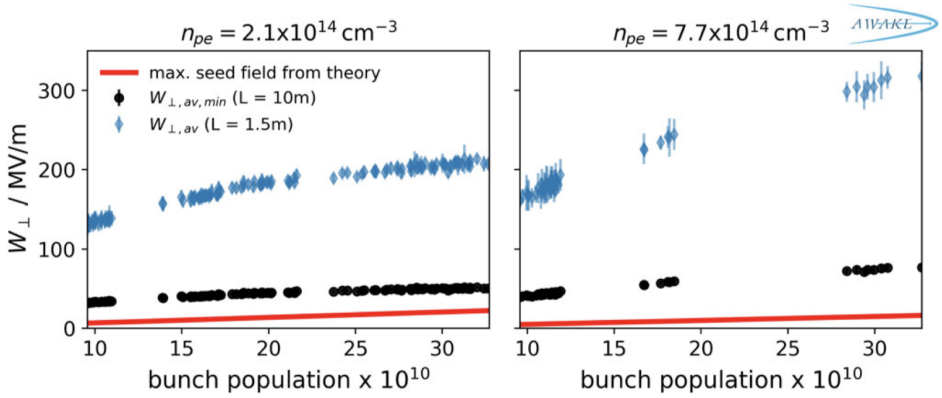

Figure 3.4.5: Comparison between estimated average transverse wakefield experienced by maximally-defocused protons and the seed wakefield amplitude predicted by linear wakefield theory. Blue data points show a best-estimate informed by simulation results, while the black data points represent the minimum value of the average transverse wakefield amplitude. Red lines represent the predicted seed transverse wakefield amplitude. *Image reproduced from Ref. [154].*

($\lambda_p \propto n_e^{-1/2}$). The electrons were injected at an oblique angle of $1.2 - 2\,\text{mrad}$ at an injection delay of 200 ps. Clear signals in the spectrometer scintillating screen corresponding to accelerated electrons were observed as shown in Figure 3.4.6. The accelerated witness signal was stable and reproducible over an extended time period; Figure 3.4.7 shows the spectrometer signal for 194 consecutive electron-injection events with a consistent peak in the spectrum at approximately 600 MeV. A scan of the spectrometer quadrupole current (and hence focal energy) was performed over this time period with only small changes in the resulting energy spectrum observed confirming the finite energy spread ($\sim$ 10%) of the accelerated witness bunch. Over a 10 m plasma length the mean energy of the witness bunch represents an average longitudinal wakefield amplitude of approximately $60\,\text{MVm}^{-1}$, a small fraction ($\sim$ 5%) of the cold wave-breaking field at this density, $E_0 = 1.29\,\text{GVm}^{-1}$. However, as demonstrated in Section 3.3.2 it is likely that the witness electrons underwent multiple dephasing cycles over the plasma length and the effective acceleration length was much shorter, with much larger instantaneous longitudinal wakefield amplitudes.

Following successful capture and acceleration of witness electrons at low plasma densities, modifications to the longitudinal plasma density profile were made in an attempt to increase the energy of the accelerated witness electrons. The simplest modification is to increase the background vapour, and hence plasma, density and allow access to larger wakefield amplitudes. Figure 3.4.8 demonstrates the effect of increasing the plasma density, with larger witness energies achieved with increasing plasma density as expected. The final energy of the witness electrons is a function of two parameters: the average accelerating gradient experienced by the bunch and the length over which electrons are consistently accelerated. While the value of the cold wave-breaking field, E_0, scales predictably with $n_{pe}^{1/2}$ according to linear theory, the acceleration length does not remain constant as function of the plasma density. This is due to the difference in self-modulation evolution for a constant proton beam in dif-

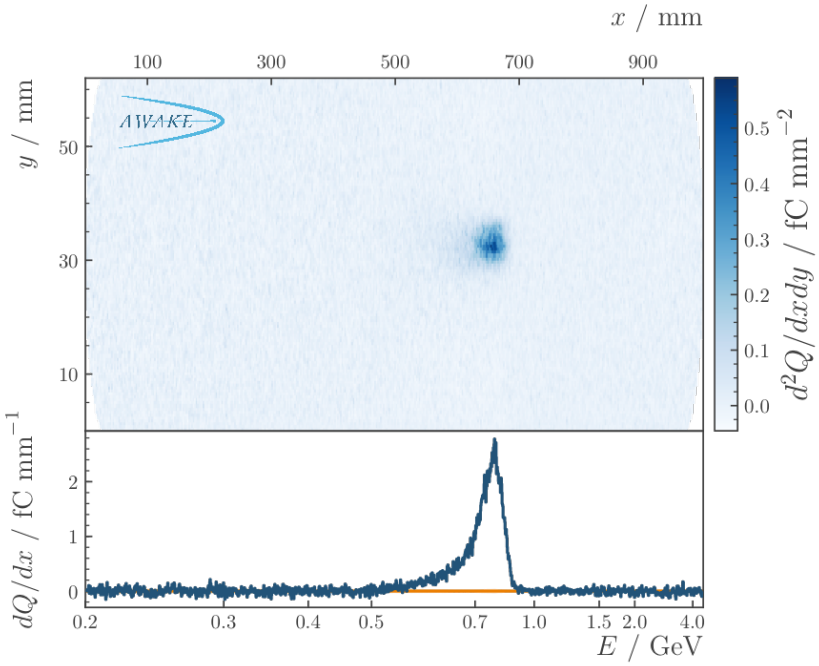

Figure 3.4.6: Signal of electrons accelerated in the wakefield driven by a self-modulated proton bunch. (a) Example scintillator image for a plasma density of $n_0 = 1.8 \times 10^{14}\,\mathrm{cm}^{-3}$ with a density gradient of +5.3%/10 m. (b) Projection in the dispersive plane with associated spectrometer energy axis. *Image reproduced from Ref. [144].*

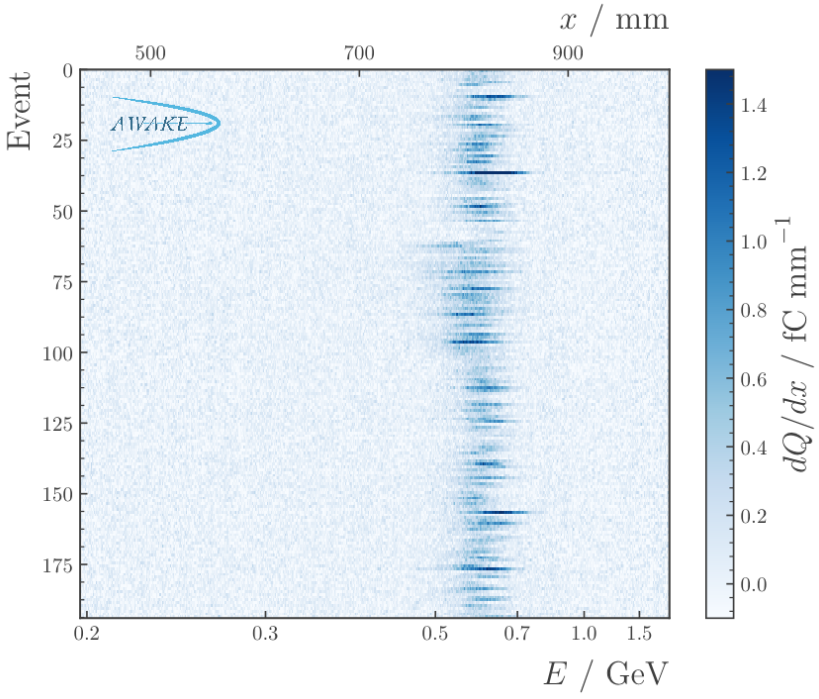

Figure 3.4.7: High statistics dataset demonstrating energy stability of accelerated electrons. The plasma density was kept constant at $n_0 = 1.8 \times 10^{14}\,\mathrm{cm}^{-3}$ over the duration of the dataset with no density gradient applied. *Image reproduced from Ref. [144].*

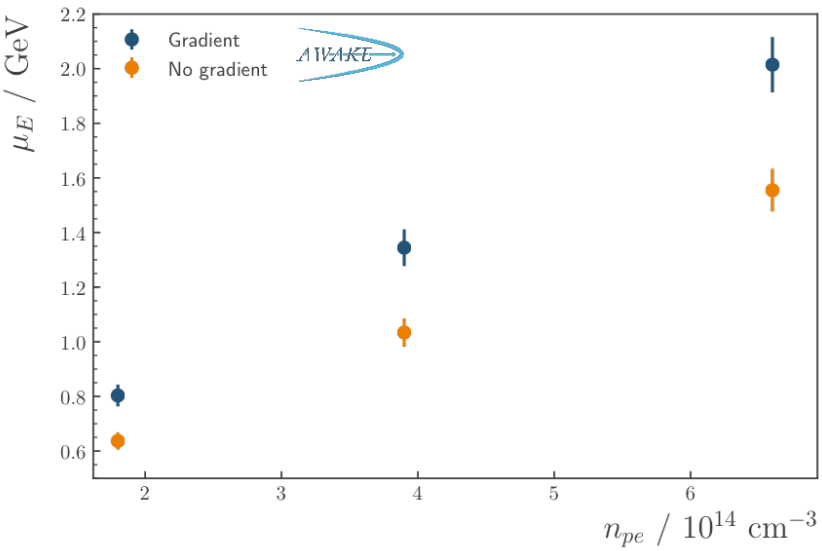

Figure 3.4.8: Measurement of the highest mean witness energy, μ_E, achieved at different plasma densities, n_{pe} with and without the presence of a plasma density gradient. The value of the density gradient was optimised experimentally for each density. *Image reproduced from Ref. [144].*

ferent plasma densities. As discussed in Section 3.3.2, consistent witness acceleration can only occur following saturation of the self-modulation process such that the wakefield phase is constant with respect to the proton micro-bunches and witness electrons. In simulations self-modulation saturation occurs more quickly as the plasma density is increased, leading to extended effective acceleration lengths which, in combination with increased wakefield amplitudes, gives correspondingly higher witness energies.

Longitudinal plasma density gradients were also introduced experimentally by independently adjusting the temperature of the Rb flasks at either end of the vapour source. Small percent-level positive gradients ($n_{pe,exit} > n_{pe,entrance}$) were predicted to counter the wakefield phase evolution discussed in Sections 3.2.2 and 3.3.2 that leads to the loss of proton micro-bunch charge and witness electrons via dephasing. The introduction of a linearly-increasing plasma density continuously shortens the plasma wavelength along its length and increases the phase velocity of the wakefield such that for some "optimum" density gradient the wakefield phase velocity can be matched to that of the micro-bunches, reducing their rate of charge loss and minimising dephasing effects [142, 156]. The exact value of this "optimum" varies with both the plasma density and the injection delay, $\Delta\tau$, and hence experimentally it was varied until the mean witness energy reached a peak. The empirically-discovered values of the density gradient optima were $+5.3\% \pm 0.3\%$, $+2.5\% \pm 0.3\%$ and $+2.2\% \pm 0.1\%$ for plasma densities of $1.8 \times 10^{14}\,\mathrm{cm}^{-3}$, $3.9 \times 10^{14}\,\mathrm{cm}^{-3}$ and $6.6 \times 10^{14}\,\mathrm{cm}^{-3}$ respectively. In all cases, a significant increase in the measured mean witness energy was observed by the introduction of a positive longitudinal plasma density gradient as shown in Figure 3.4.8. The largest witness energies of approximately 2 GeV were observed in a plasma of density $6.6 \times 10^{14}\,\mathrm{cm}^{-3}$ with a longitudinal gradient of $+2.2 \pm 0.1\%$ applied.

3.5 CONCLUSION

Proton bunches represent promising PWFA drive bunches due to their large stored energy negating the requirement for the staging of multiple plasma acceleration sections to reach the high energies necessary for future applications. The process of self-modulation has been shown to transform experimentally-available long proton bunches into effective wakefield drive beams. The evolution of the wakefield phase during the self-modulation process leads to significant micro-bunch charge loss and difficulties in the capture and consistent acceleration of externally-injected witness bunches. However, these effects can be minimised through tuning of the longitudinal plasma density profile and careful choice of the witness injection position.

During Run 1 of the AWAKE experiment, the main experimental goals were achieved: the demonstration of self-modulation of a proton bunch in a plasma [153, 154] and the acceleration of witness electrons injected into the wakefield driven by a self-modulated proton bunch [144]. These experiments were completed by May 2018, leaving three experimental periods before the end of AWAKE Run 1 to perform more detailed parameter scans to better understand both the self-modulation and witness acceleration processes. One priority of the remaining experimental periods was to attempt to map the evolution of the wakefield driven by the self-modulated proton bunch along its length and demonstrate resonant excitation of the longitudinal wakefield while also exploring the evolution of the wakefield on extended timescales. These measurements and their analysis is the subject of Chapter 7.

4

THE FLASHFORWARD EXPERIMENT

The FLASHForward experiment is an electron-beam-driven plasma wakefield acceleration experiment based at the Deutsches Elektronen-Synchotron (DESY) in Hamburg, Germany. It uses high-quality, low-emittance, GeV-level electron beams produced by the FLASH free-electron laser (FEL) facility [58, 157, 158] as drive bunches, with witness bunches generated either via additional cutting and shaping of the drive bunches [159] or via internal injection of electrons within the plasma [76]. This chapter first gives an overview of the experiment, before briefly describing the three main scientific goals that are being explored within the experiment. Particular attention is paid to experimental diagnostics and diagnostic techniques that have been developed at FLASHForward that facilitate detailed reconstruction of the electron bunches, enabling accurate experiment-simulation comparisons in later chapters.

4.1 EXPERIMENTAL OVERVIEW

A detailed overview of the FLASHForward experiment and its primary scientific goals can be found in Ref. [57] and are briefly summarised here. A schematic of the FLASH facility is presented in Figure 4.1.1; the electron bunches are generated at frequencies up to 10 Hz using a photocathode RF gun and accelerated up to energies of 1.25 GeV using seven 1.3 GHz super-conducting radio-frequency (RF) accelerating modules (in three sections). The series of SCRF cavities and bunch compressors within the accelerator permits control of the longitudinal phase space of the bunch [160], with off-crest acceleration in the SCRF cavities typically used in FLASHForward experiments to produce bunches with a correlation between the energy and longitudinal coordinate of particles within the bunch, i.e. an energy *chirp*. Use of a third-harmonic SCRF cavity in the accelerator beamline gives additional control of the first and second derivative of the chirp [161], allowing approximately linear energy chirps to be produced, required for effective drive and bunch pair generation at the collimator [159] as is discussed in more detail later. Typical electron bunches used for FLASHForward experiments have energy spreads of approximately 0.1%, normalised transverse emittances of order 2 µm, bunch charges approaching 1 nC (before collimation), and RMS bunch lengths between $200 - 500$ fs with peak currents (I_{pk}) of order $1 - 2.5$ kA. The exact peak current can be tuned depending on the experiment being performed; external injection experiments require exquisite control of the longitudinal phase space to enable precise bunch shaping and hence operate at lower bunch compressions ($I_{pk} \sim 1$ kA) whereas internal injection experiments relax these requirements on the longitudinal phase space and prioritise peak current ($I_{pk} \gtrsim 2$ kA).

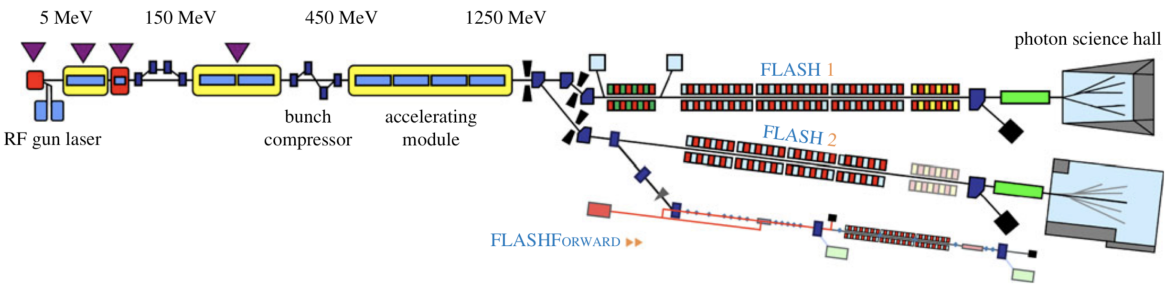

Figure 4.1.1: Schematic of the FLASH SCRF front-end supplying the FLASHFORWARD experimental beamline with high-average-power electron beams. Also shown are the FLASH1 and FLASH2 FEL beamlines. *Image reproduced from Ref. [57].*

The bunch is extracted from the FLASH linac using a kicker magnet and transported to the FLASHFORWARD beamline, an overview of which is presented in Figure 4.1.2. In this dispersive extraction section, three energy collimators can be independently used to remove energy slices and further shape the bunch [159]. These represent two metallic blocks that can be transversely inserted into the beamline to remove charge from the head or tail of the bunch, and a tapered wedge that can be moved into the beamline and translated horizontally and vertically to remove varying amounts of charge from the centre of the bunch, thus allowing generation of separated drive-witness bunch pairs of variable charge. Once in the FLASHFORWARD beamline, the bunch passes through a series of quadrupoles ("Matching quadrupoles" and "Final focusing quadrupoles" in Figure 4.1.2) that allow the bunch to be focused down to micron-scale transverse sizes at the entrance of the plasma.

The plasma is contained within a cylindrical capillary of diameter 1.5 mm, milled into two sapphire slabs. Three different lengths of capillary have been used in experimentation thus far; 33 mm, 50 mm and 195 mm. The capillary is mounted on a hexapod system [162] that permits controlled movement of the capillary with respect to the axis of the electron beam. Two inlets at either end fill the capillary with variable gas species, with gas density controlled by the pressure of a buffer volume connected to gas feedthroughs. The plasma can be generated via two methods; a high-voltage electrical discharge of duration approximately 400 ns, or a longitudinally-propagating 25 TW laser pulse of duration ~ 40 fs which ionises a thin column of plasma contained within the capillary. The discharge voltage can be varied up to 25 kV, with stable operation in argon typically observed for voltages exceeding 8 kV, although this is gas species and pressure dependent. The timing of the generation of the plasma can be controlled independently of the arrival of the electron bunch. The discharge system has been measured to have an $\mathcal{O}$(ns) timing jitter due to the use of a thyratron switch, while the laser system is synchronised to the RF linac to within 100 fs [76]. When the discharge system is operational, in- and out-going current traces are measured on every event to allow monitoring of the stability of the discharge.

After propagating through the plasma, the bunch is captured by a series of quadrupoles and dispersed by a 1 m-long dipole magnet on to a series of scintillating screens attached

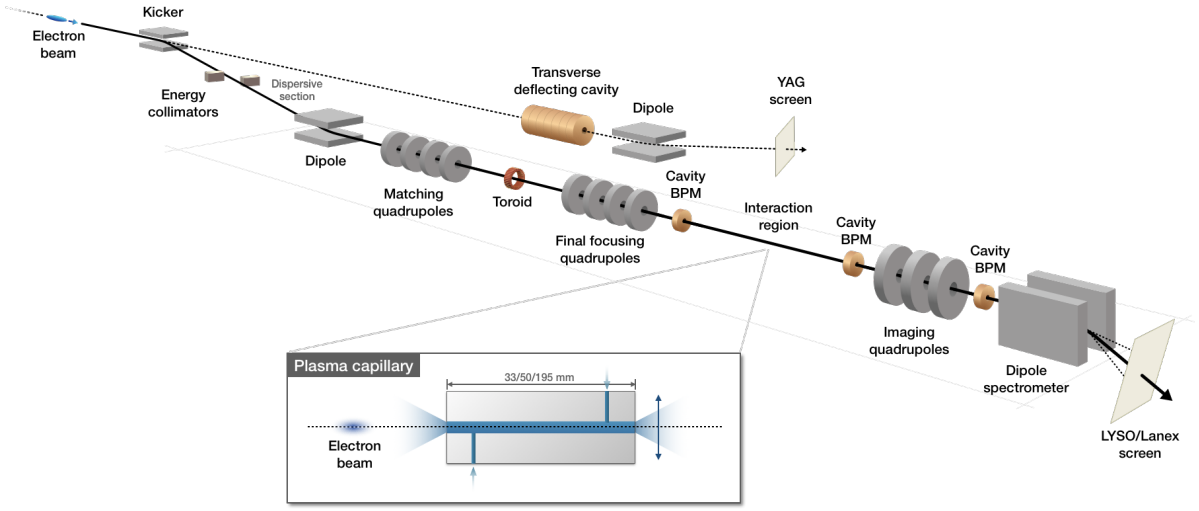

Figure 4.1.2: Overview of the FLASHForward beamline setup. Inset: The electron beam propagates through a 1.5 mm diameter sapphire capillary of variable length. The plasma is generated either via a high-voltage discharge or laser ionisation from the synchronised 25 TW laser system. The capillary sits on an adjustable hexapod mount allowing free movement of the capillary with respect to the axis of propagation of the probe beam.

to the outside of the vacuum chamber to allow measurement of its energy distribution. The dipole accepts electron energies from a few MeV up to 2.4 GeV. The light emitted by the scintillating screens is imaged by a series of six cameras placed around the vacuum chamber; three cameras are placed along the right-side exit of the dipole chamber and represent a high-resolution region, while the other three image the 1 m-long lower exit of the dipole at much reduced resolution. The dipole current can be varied to move the bunch into the high-resolution region of the screen, irrespective of its energy. Two different types of scintillating screens are used depending on the experiment being performed. For typical single bunch experiments, terbium-doped gadolinium oxysulfide (Gd_2O_2S:Tb) 'Lanex' scintillators are used. However, their large decay constant, $\tau \sim 500\,\mu s$, leads to the accumulation of scintillation signal when used in high-repetition-rate, multi-bunch measurements where bunch separations are typically on the order of 1 µs or lower. In this case, a lutetium yttrium orthosilicate (LYSO, Lu_2SiO_5) crystal is preferred as it has a decay constant on the order of 100 ns and thus scintillation signals from multiple bunches can be differentiated for bunch separations exceeding approximately 500 ns. The majority of the scintillating material is 'Lanex' as it is easier to work with and replace in case of damage, but a 4 cm region near the lower end of the high-resolution region of the spectrometer is formed of 200 µm-thick LYSO crystal. This region is imaged by two separate cameras focused on the LYSO crystal that can be triggered independently to measure the scintillation signals of two temporally-separated bunches in the same event.

Within the past 18 months the FLASHForward beamline has been extended beyond the electron spectrometer to include a transverse deflecting structure [163, 164] (X-TDS) that enables the measurement of the longitudinal phase space and transverse slice emittance of

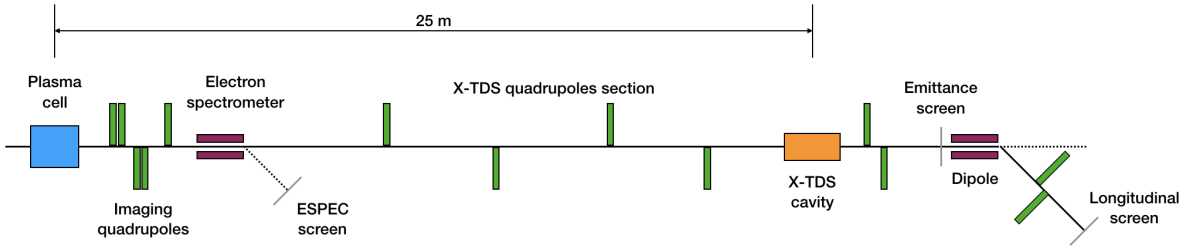

Figure 4.1.3: Schematic of the FLASHForward beamline extension completed in 2019 that introduced the X-band transverse deflecting structure (X-TDS). The beamline was extended by approximately 25 m and included additional quadrupoles, the X-TDS cavity, an additional dipole, and two scintillating screens to enable measurement of the transverse slice emittance and longitudinal phase space of the electron bunch. *Image recreated from Ref. [167].*

the electron bunch, both with and without plasma acceleration, with femtosecond-level resolution [165–167]. A schematic of the beamline extension is illustrated in Figure 4.1.3. The measurement is based on 'streaking' of the electron bunch using high-frequency (X-band) RF fields, such that the field experienced by a particle within the bunch, and hence its deflection, is correlated to its longitudinal position within the bunch (ξ), effectively mapping the longitudinal plane z onto the transverse plane y. The ξ-dependent deflection can be determined by measuring the transverse distribution of the bunch at a screen and, if repeated for multiple phase advances by varying upstream quadrupole strengths, it is possible to reconstruct the transverse emittance of the bunch as a function of ξ, thus measuring its slice emittance [168]. When additionally combined with a dipole bending in the other axis, such that the energy offset δ is mapped to the transverse plane x, the resulting distribution measured at a screen after the dipole can be used to reconstruct the longitudinal phase space of the bunch. The transverse deflecting structure installed within the FLASHForward experiment is unique in that the deflecting field has variable polarisation [169]. This allows the bunch to be streaked at any angle and by combining measurements of the projected profiles at multiple angles it is possible to produce a three-dimensional tomographic reconstruction of the charge density distribution of the bunch [170]. The combination of transverse slice-emittance measurements, longitudinal phase space measurements and bunch tomography techniques enables complete reconstruction of the six-dimensional phase space of the electron bunch, enabling exquisite control over the bunch profile and detailed experiment-simulation comparisons [68, 105] beyond those previously demonstrated in plasma acceleration. Prior to the installation of the X-TDS, another transverse deflecting structure [171, 172] in the neighbouring FLASH1 beamline was used to characterise the phase space of the electron bunches before they entered that FLASHForward beamline and is shown in the upper region of Figure 4.1.2.

Alongside the detailed bunch profile reconstruction measurements enabled by the X-TDS, additional diagnostic techniques have been developed within the FLASHForward experiment to improve control of the electron bunch and hence acceleration process. One such example that is readily used is the method outlined in Ref. [173] where measurements of the centroid jitter of electron bunches between two beam-position monitors (BPMs) can be used

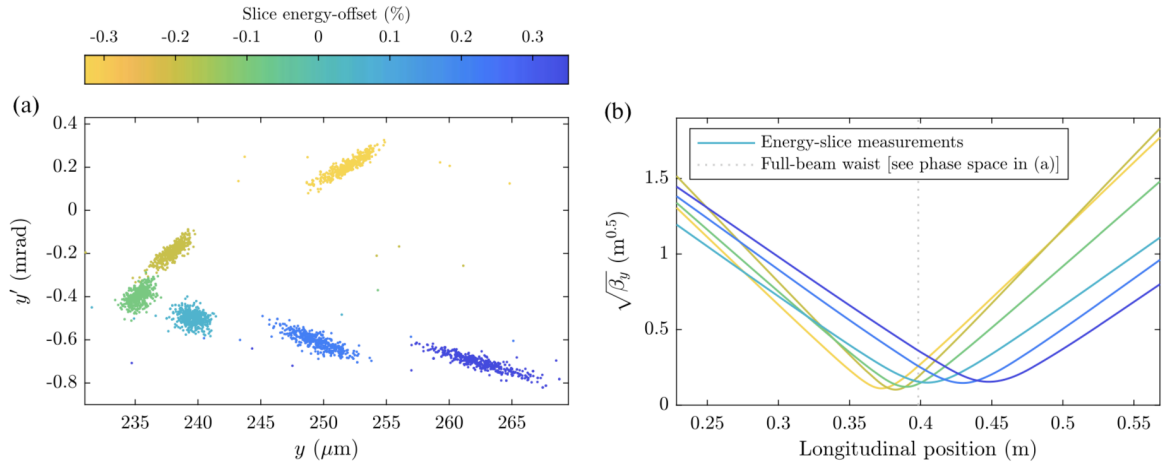

Figure 4.1.4: Example of the two-BPM bunch reconstruction measurement technique. (a) Jitter phase space measurement of a range of energy slices. Significant dispersion (energy-dependent offsets) can be observed. (b) Corresponding evolution of the slice β-function around the focus, indicating a highly chromatic focus. *Image modified from Ref. [173].*

to reconstruct the β-function of the bunch as it is strongly focused into the entrance of the plasma. When combined with the upstream collimators, such that individual energy-slices of the bunch can be selected, this technique enables measurement of the transverse phase space of multiple bunch slices, demonstrating the chromaticity and slice-dependent focusing of the bunch as shown in Figure 4.1.4. As this technique is fast and non-invasive, it is used during the experimental setup to minimise the dispersion of the bunch and perform slice-specific matching into plasma, enabling high-quality acceleration to be performed.

The use of the FLASH FEL linac that provides short, high-current, high-quality, stable, PWFA-suitable electron bunches, combined with the array of advanced diagnostic techniques available within the FLASHForward experiment enables it to perform detailed, high-statistics studies of electron-beam-driven plasma wakefield acceleration and make large strides towards optimisation of the acceleration process. Initial measurements in this direction have already been performed, demonstrating the operation of internal injection mechanisms with the potential to produce even higher-quality bunches, quality-preserving acceleration of externally-injected bunches and detailed studies of the evolution of plasma after the wakefield process that are relevant for high-repetition-rate operation. These measurements, in the context of the scientific goals of the FLASHForward experiment, are outlined in the following section.

4.2 SCIENTIFIC GOALS

The primary scientific goals of the FLASHForward experiment can be summarised by consideration of the three (current) 'core' experiments being undertaken:

- X-1: Controlled internal injection of a high-quality witness bunch.
- X-2: Quality-preservation of an externally-injected witness bunch.
- X-3: High-repetition-rate plasma wakefield acceleration.

An outline of each experiment and recent experimental highlights follows.

4.2.1 *X-1: Internal injection*

The X-1 experiment aims to demonstrate the controlled injection of high-brightness witness bunches generated by density-downramps (DDRs) within beam-driven plasma wakefields. DDR injection uses sharp, imposed gradients in the longitudinal plasma density profile to locally reduce the phase velocity of the wakefield, extending its length and facilitating the injection of plasma electrons forming the wake into its accelerating phase, trapping them and forming a witness bunch [174, 175] as demonstrated in Figure 4.2.1. Injecting into the wakefield in this way has been shown to result in ultra-high quality witness bunches with sub-µm normalised emittances in 3D PIC simulations [72, 73], and is highly tunable via modification of the profile of the density-downramp [176, 177]. This technique is commonly used in laser-driven wakefield accelerators [178–180] but stable operation was not previously demonstrated experimentally in beam-driven wakefield accelerators. Previous beam-driven DDR studies were typically limited to few successful events as large shot-to-shot jitters prevented consistent injection [75, 181–183].

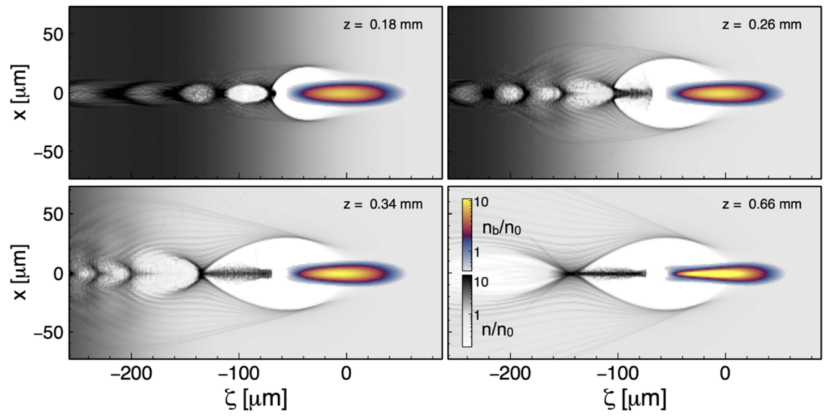

Figure 4.2.1: 3D PIC simulation demonstrating the evolution of a beam-driven wakefield as it propagates through a density-downramp. As the beam propagates, the wakefield cavity extends as the density decreases, locally injecting and trapping a witness bunch into the rear of the cavity. *Image reproduced from Ref. [176].*

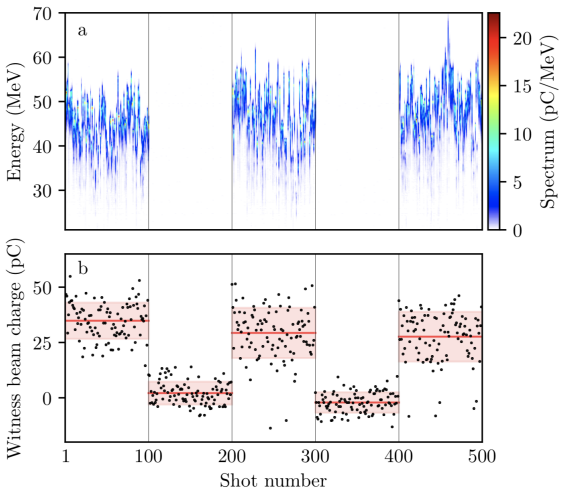

Figure 4.2.2: (a) Energy spectra and (b) charge of the density-downramp injected witness beam. The injection is switched off by blocking the transverse laser pulse in shot numbers $101-200$ and $301-400$. *Image reproduced from Ref. [76].*

Within the FLASHForward experiment, generation of a DDR was achieved using two independently-focused arms of the 25 TW laser pulse [76]. One propagated longitudinally through the gas-filled capillary while the second entered the capillary transversely through a small hole 20 mm along the capillary. The capillary was filled with a mixture of argon and helium gas at a partial-pressure ratio of $1:2$ to enable selective local ionisation; the intensity of the longitudinally-propagating laser pulse was kept sufficiently low to ionise only the first level of argon, while the transversely-propagating laser intensity was increased such that it could ionise both the second level of argon and the first level of helium. This therefore allowed generation of a region of increased plasma density around the focus of the transversely-propagating laser pulse, representing the DDR [184, 185]. The drive electron bunch was maximally compressed without concern for linearising its longitudinal phase space, achieving a peak current of 2.1 kA [76].

A brief discussion of the stable demonstration of DDR injection in a beam-driven wakefield accelerator at FLASHForward follows; all results mentioned can be found in Ref. [76]. Laser-controlled injection was first demonstrated, with witness bunches of charge 32.1 ± 9.6 pC generated when the transverse laser was present and disappearing when the transverse laser was blocked, as demonstrated in Figure 4.2.2. Following this, a dataset of 1885 consecutive injection attempt events was recorded with witness bunches observed in 95.4% of the events with a mean charge of 33 ± 10 pC, mean energy of 45 ± 5 MeV and an RMS relative energy spread of 4.4%. These witness bunches were accelerated over the remaining capillary length after the density-downramp, indicating an average accelerating gradient of approximately $2\,\mathrm{GVm}^{-1}$. The transverse RMS normalised emittance of the witness bunch projected over all events was measured to be $\varepsilon_x^n = 9.3 \pm 0.3\,\mu\mathrm{m}$.

Two parameter scans were then performed to investigate dependencies of the witness bunch on the experimental conditions; the energy of the laser pulses (in both arms), and the

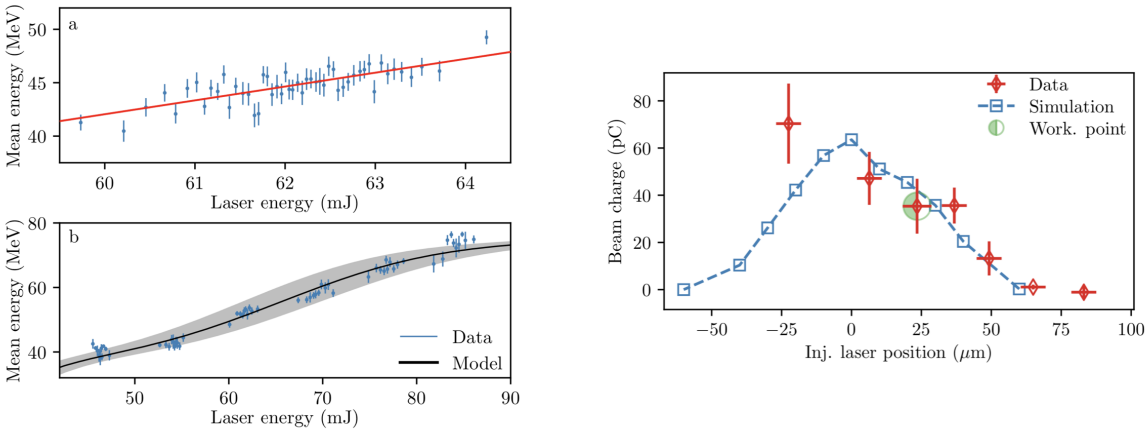

Figure 4.2.3: Left: (a) Binned variation of mean witness bunch energy with the laser energy over 1885 consecutive events. (b) The laser pulse energy was scanned over a wider range. Measurements (blue) indicate good agreement with a simple model (black) predicting the variation of plasma density as a function of the laser energy, based on the assumption that the longitudinal field scales as $n_e^{1/2}$. Right: Effect of the positioning of the laser-spot centroid on the measured charge of the witness bunch. *Images reproduced from Ref. [76].*

vertical position of the transversely-propagating laser pulse within the longitudinal plasma column. The mean energy of the witness bunch was observed to increase with increasing laser pulse energy as demonstrated in the left side of Figure 4.2.3. The measured increase in mean witness energy could be reconstructed using a model (black line, Figure 4.2.3(b)) that calculated the expected plasma density as a function of the laser energy via the ADK ionisation model [186], and the resultant effect this has on the longitudinal wakefield amplitude under the assumption that it scales as $n_e^{1/2}$. The vertical position of the transverse laser arm was modified by moving its entire focusing assembly in the y-direction to ensure its focus was not modified. This changed the overlap between the DDR and the propagation axis of the drive electron bunch, reducing the amount of charge injected into the wakefield. The density ramp produced by the transverse laser pulse has a finite length and spatial structure influenced by the focal spot of the laser pulse. As the overlap between the electron bunch and DDR changes, the region over which injection occurs is modified, resulting in changes in the witness bunch charge. Experimental measurements of this effect are presented in the right side of Figure 4.2.3 and compared to 3D PIC simulations of the injection region that use the measured transverse laser pulse spot as input to predict the (asymmetric) downramp profile, showing reasonable agreement.

The X-1 experiment has successfully demonstrated generation of a DDR-injected witness bunch in a beam-driven plasma wakefield [76]. Future studies will focus on improving control of the injection process, particularly the brightness of the resultant witness bunch by reducing its emittance and energy spread. This will be achieved by optimising both the drive bunch properties and the properties of the density-downramp.

4.2.2 *X-2: External injection*

The X-2 experiment aims to maintain the quality of an externally-injected witness bunch, particularly demonstrating simultaneous preservation of its emittance and energy spread, while operating at high total acceleration efficiency. Theoretically, this has been shown to be possible in the non-linear blowout wakefield regime [48, 86, 87] via (i) beam-loading of the wakefield [98–101, 187], where the longitudinal field is locally flattened by the presence of the witness bunch, and (ii) correct matching of the bunch into the plasma [188–190], where the focusing provided by the ion column balances the divergence of the beam such that the beam envelope experiences no oscillations as it propagates.

To be able to achieve preservation of the quality of an externally-injected witness bunch, precise control of the drive and witness bunches and the plasma wakefield itself must be achieved; this is only possible by using similarly precise diagnostic techniques. Within the X-2 experiment, a technique for measuring the longitudinal wakefield amplitude experienced by the drive and witness bunches was developed, allowing sampling of the wakefield with femtosecond-level resolution [105]. This measurement involved progressively removing the rear of a linearly-chirped electron bunch using the energy collimators in the dispersive section of the beamline [159]. Removing charge from the rear of the bunch does not alter the wakefield experienced by the remaining charge as demonstrated in Figure 4.2.4. Therefore, by comparing measurements of the energy spectra as charge is progressively removed from the tail of the bunch, it is possible to determine the energy and thus longitudinally-averaged wakefield amplitude experienced by each slice of the bunch. This measurement technique was shown to be consistent with 3D PIC simulations of the experiment over a range of plasma densities and bunch profiles as demonstrated in Figure 4.2.5.

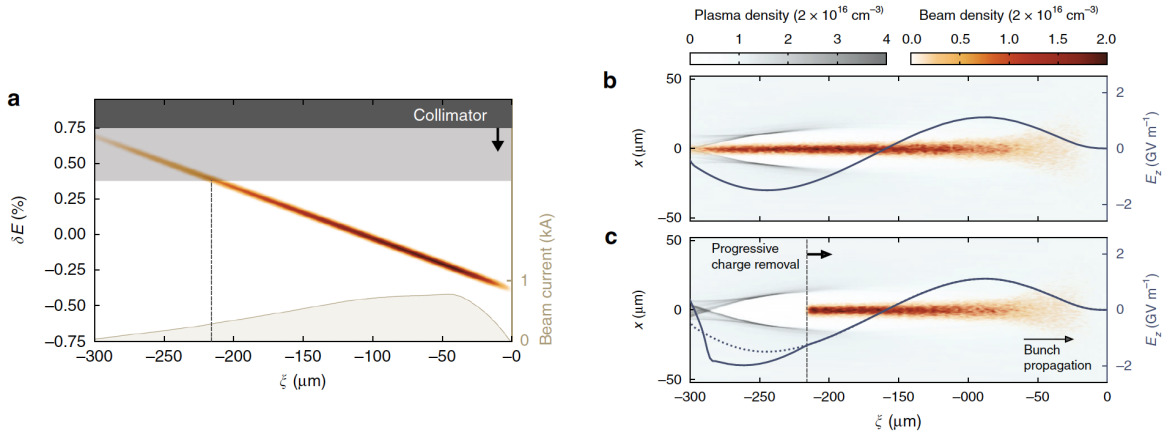

Figure 4.2.4: Example of the collimation-based wakefield sampling technique. (a) A strongly correlated longitudinal phase space allows progressive removal of energy slices. (b) The electron bunch interacts with plasma and excites a density wake with large amplitude longitudinal fields. (c) Removing charge from the bunch tail alters the wakefield behind the bunch, but not that experienced by the remaining charge. *Image modified from Ref. [105].*

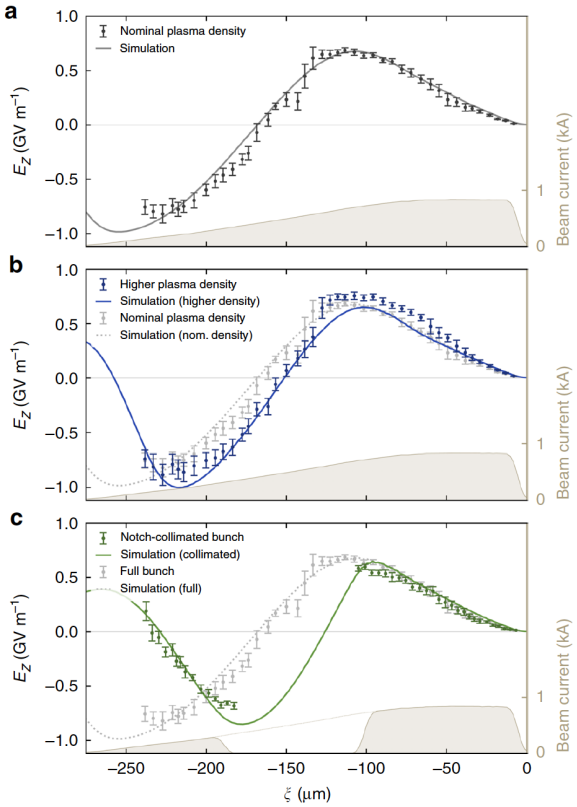

Figure 4.2.5: Measured longitudinally averaged plasma wakefields for (a) a nominal bunch profile and plasma density. (b) The measurement is repeated at an 80% higher plasma density, demonstrating shortening of the wakefield. (c) The centre of the bunch is removed to produce a drive-witness pair at the nominal plasma density. The wakefield is identical at the head, but significantly altered at the rear after charge has been removed. All measurements are compared to 3D PIC simulations using the experimentally measured bunch profile and plasma density. *Image reproduced from Ref. [105].*

Armed with the ability to accurately tune the longitudinal phase space [161] and shape the current profile [159] of an electron bunch, measure its slice emittance and reconstruct its Twiss parameters at focus [173], and finely sample the longitudinal wakefield it drives [105], the X-2 experiment is in an ideal position to demonstrate the preservation of the quality of an externally-injected witness bunch. First steps towards achieving this were taken in the measurements presented in Ref. [68] where energy-spread preservation and high efficiency acceleration of an externally-injected witness bunch in a plasma wakefield were simultaneously demonstrated for the first time.

In these measurements, the electron bunch produced in the FLASH linac was divided into a drive-witness bunch pair using the wedge-shaped notch collimator [159] in the dispersive section of the beamline. A parameter scan of the position and width of the notch collimator (to vary the drive-witness bunch charge ratio) and the plasma density was performed to enable identification of the point of optimal acceleration of the witness bunch. This three-dimensional parameter scan represented 12675 events in total. Once the optimal setup had

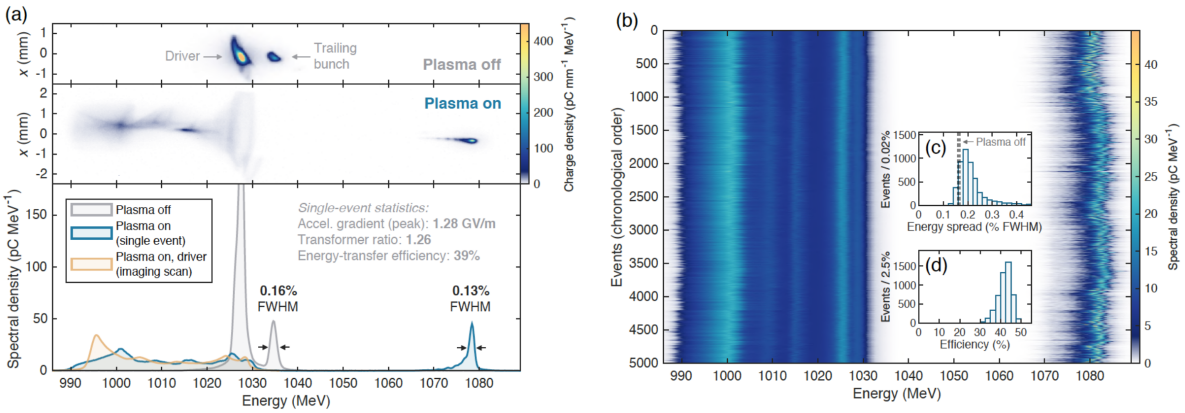

Figure 4.2.6: (a) Spectrometer images at the optimal operating point as well as the corresponding energy spectra for events with and without plasma interaction. (b) High stability is observed across 5000 events. (c) The energy spread is preserved in 6.4% of these events and nearly preseved in the remaining events. (d) Simulataneously, high energy-transfer efficiency is observed, distributed between 30% and 50%. *Image reproduced from Ref. [68].*

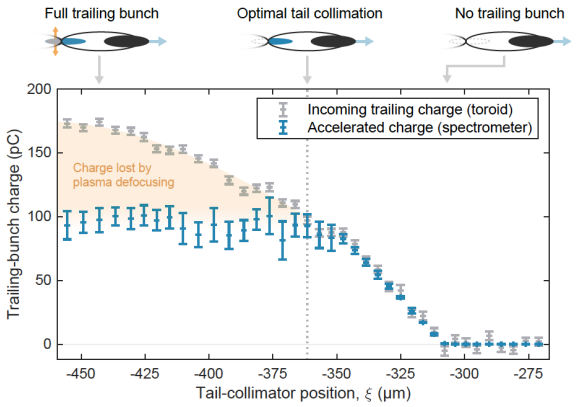

Figure 4.2.7: Demonstration of full charge coupling through comparison between the incoming witness bunch charge (grey) and the accelerated charge measured on the spectrometer screen (blue) via a scan of the tail-collimator position. Full charge-coupling is observed until $\xi \lesssim 360\,\mu\text{m}$ at which point the additional charge extends outside the wakefield cavity and is strongly defocused by returning sheath electrons. *Image reproduced from Ref. [68].*

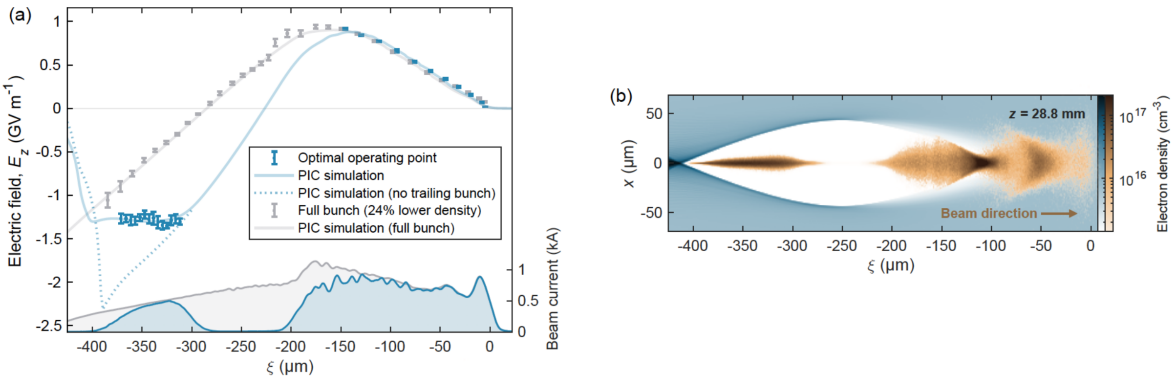

Figure 4.2.8: (a) Longitudinally-averaged wakefield measurement, both for optimal beam-loading (blue points) and the full bunch (grey points) in excellent agreement with PIC simulations (lines), demonstrating flattening of the longitudinal wakefield. (b) PIC simulation snapshot demonstrating the optimal beam-loading working point. *Image reproduced from Ref. [68].*

been identified (shown in Figure 4.2.6(a)), 5000 consecutive acceleration events were recorded and demonstrated 45.4 ± 1.4 MeV energy gain in a plasma of length 34.2 mm, representing an average accelerating gradient in excess of 1.3 GVm^{-1}. Figure 4.2.6(b) shows a waterfall image of the energy spectra of all events. Of these 5000 events, the initial FWHM relative energy spread of the witness bunch of 0.16% was preserved in 6.4%, while the rest showed a 28% median energy-spread growth to 0.2% FWHM as shown in Figure 4.2.6(c). The acceleration efficiency is calculated by the ratio of the energy gained by the charge in the witness bunch to the energy lost by the charge in the drive bunch. At the optimal working point, the energy-transfer efficiency was measured to be 42 ± 4%, demonstrated in Figure 4.2.6(d).

One-hundred percent charge coupling of the witness bunch into the wakefield was additionally demonstrated at the optimal acceleration configuration via a tail-collimator position scan as shown in Figure 4.2.7. The charge of the accelerated witness bunch was measured on the charge-calibrated scintillating spectrometer screen and compared to measurement of the bunch charge upstream of the plasma interaction. The measurements agree until the tail of the witness bunch extends outside the wakefield cavity, causing it to be defocused by returning sheath electrons. Confirmation of beam-loading of the wakefield by the witness bunch at the few-percent level was demonstrated via the wakefield sampling technique introduced previously [105] and showed significant flattening of the longitudinal wakefield as can be seen in Figure 4.2.8.

Given the excellent recent experimental progress that has been enabled by improved control and measurement of the wakefield process, the X-2 experiment is in position to demonstrate emittance preservation of the witness bunch in the near future. Following this, demonstration of GeV-level energy gain while simultaneously preserving witness bunch quality and maximising the total efficiency of the acceleration process will be the focus of future experimental studies.

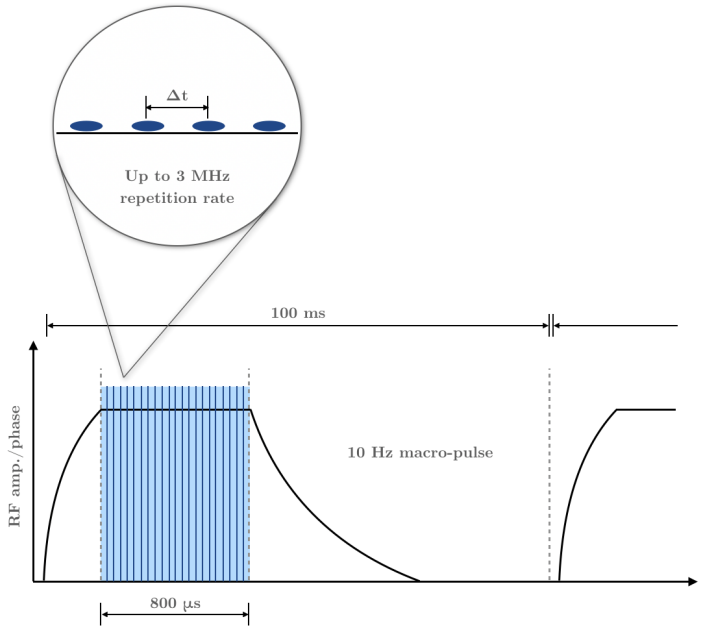

Figure 4.2.9: Schematic of the 10 Hz macro-pulse and MHz bunch train structure available to FLASHForward for high-power experimentation. *Image adapted from Ref. [57].*

4.2.3 *X-3: High-repetition-rate plasma wakefield acceleration*

Building on the progress made by the X-2 experiment, the X-3 experiment aims to demonstrate high-quality acceleration of an externally-injected witness bunch at the highest possible repetition-rate, eventually demonstrating the operation of a > 10 kW average power plasma-based accelerator.

FLASHForward is uniquely positioned within the plasma-based accelerator community to study high-repetition-rate operation thanks to the capabilities of the FLASH linac that provides electron beams for experimentation. The FLASH photocathode gun and linac modules operate with a 10 Hz macro-pulse structure, where each macro-pulse contains an approximately 800 μs-long RF flat-top in which electron bunches can be accelerated with exquisite stability at up to 3 MHz micro-pulse frequencies [57, 58]. A schematic of the macro-pulse structure is shown in Figure 4.2.9. At full capacity, this therefore represents 2400 bunches per macro-pulse and thus the potential to deliver 24,000 bunches per second to the FLASHForward experimental area. Additional flexibility is introduced thanks to the presence of three separate photo-injector lasers that can be used in parallel to produce electron beams of similar characteristics [58]. By using more than one photo-injector, the FLASH linac is able to provide two electron bunches at the limit of the frequency of the SCRF modules, 1.3 GHz, and hence bunch separations of 769 ps are achievable. The X-3 experiment has developed the necessary expertise to readily make use of this functionality, enabling the generation of two similar (but not identical) PWFA-suitable drive bunches at separations ranging from 769 ps to the full extent of the RF flat-top, in steps of 1.3 GHz, representing a six order-of-magnitude temporal

range. This has enabled experimental study of the evolution of plasma wakefields over such a range for the first time and is the focus of later chapters of this thesis.

Plasma wakefield acceleration has never previously been performed at MHz repetition-rates; beam-driven wakefield accelerators typically operate at 10 Hz and below [56, 61] and high-power (> 100 TW) laser-driven wakefield accelerators are typically operating at single-Hz rates and below [77–79]. Lower-power LWFAs (few-TW) have been operated at kHz repetition-rates but typically use short gas jet targets and have not yet demonstrated acceleration of electrons beyond a few MeV [191–193]. It is also not yet possible to simulate plasma wakefield evolution over extended timescales using the PIC technique due to the accumulation of numerical effects that cause such simulations to lose accuracy. Hence, the initial purview of the X-3 experiment is to develop diagnostic techniques that can be used to probe the evolution of the plasma at such repetition-rates — one such example is discussed in Chapter 5 — to develop an understanding of the physics that will limit the ultimate repetition-rates available to plasma-based accelerators.

Alongside study of the plasma evolution at high-repetition-rates, a number of technological advances must also be made. At present, the plasma is often generated at FLASHForward using a high-voltage discharge that, in its current setup, cannot yet be operated at MHz repetition-rates. Modifications must be made to enable this, then stable operation at such rates must be demonstrated. Additionally, the effect of this load on the capillary itself must be studied, as has previously been performed for similar capillary discharge waveguides that aim to be operated at kHz repetition-rates [194]. Furthermore, it has previously been shown that a significant amount of plasma is expelled from capillaries following the discharge, leading to significant evolution of the longitudinal plasma density profile [195]. Either this expulsion needs to be minimised or the capillary refilled (or likely both) to enable reproducible density profiles to be generated at high-repetition-rates, motivating redesign of the capillary and associated systems to enable this. An additional consideration is that, during its lifetime, the plasma transfers a significant amount of heat to the capillary walls [196, 197], an effect that will likely be amplified by the transfer of energy from the bunches driving the wakefield. This must be efficiently dissipated, possibly via external cooling systems, to ensure the capillary does not become damaged during high-average-power operation.

4.3 CONCLUSION

The FLASHForward experiment is uniquely positioned to move beam-driven plasma accelerators from research facilities towards user facilities. This will require stable operation of high-quality acceleration at high-repetition-rates and, as demonstrated in the previous sections, rapid progress is already being made on both fronts via the three 'core' experiments. Generation of internally-injected witness bunches via density-downramp injection in a beam-driven wakefield accelerator has already been demonstrated within the X-1 experiment and, with further optimisation, promises the production of ultra-high quality witness bunches with sub-μm normalised emittances. The X-2 experiment has demonstrated precise measure-

ment and control of the wakefield, preserving per-mille energy spreads while simultaneously operating at high instantaneous efficiencies. It is expected that the recent progress of the X-2 experiment will enable the demonstration of emittance preservation of the witness bunch during the acceleration process in the near future, with large strides made in improving the total efficiency of the accelerator in unison. The X-3 experiment has made significant progress towards understanding the limitations that will be placed on high-repetition-rate operation of beam-driven plasma wakefield accelerators thanks to the development of a beam-based diagnostic technique, discussed in detail in Chapter 5, that has enabled study of the plasma over extended timescales. First experiments studying the evolution of plasma and ways to adjust its lifetime are detailed in Chapter 6, while detailed studies of the response of the plasma to the wakefield process over a range of timescales are discussed in Chapters 8 and 9. This has enabled the FLASHForward experiment to make significant progress towards the demonstration of PWFA at high-repetition-rates, although further technological advances must be made in tandem to make the operation of a $> 10\,\mathrm{kW}$ average-power plasma accelerator a reality.

Part III

UNPERTURBED PLASMA EVOLUTION

Part III details the development of a beam-based measurement technique at FLASH-Forward that enabled studies of the evolution of a plasma over extended timescales. It is first bench-marked against a standard plasma density diagnostic and then used to study a range of unperturbed plasma states with a view to adjusting the lifetime of a plasma such that it can be optimised for high-repetition-rate operation.

5

BEAM-BASED PLASMA CHARACTERISATION TECHNIQUE

As described in the previous chapter, the X-3 experiment aims to demonstrate stable plasma wakefield accelerator operation at the highest possible repetition-rate. In order to do this, the bunches used as drive beams within the PWFA process must be generated, and the plasma itself must be refreshed, at the appropriate frequency. State of the art super-conducting radio-frequency facilities [157, 198, 199] have demonstrated the ability of conventional linear accelerators to generate trains of electron bunches suitable for use as PWFA drive-bunches at MHz frequencies and beyond, with exceptional stability. However, comparatively little research has been performed to study the extended evolution of the plasma following the wakefield acceleration process. For typical FLASHForward experimental parameters, energy is deposited into the plasma by the drive bunch via the wakefield at a rate of multiple Joules per metre, driving subsequent evolution of the plasma and heating it significantly. This must dissipate such that the plasma can recover to its initial state before the arrival of the following bunch for a stable, repeatable acceleration process to be possible.

In order to study the rate at which this occurs, diagnostics that are sensitive enough to detect local, small-scale changes to the plasma state over a wide range of timescales are essential. Via the X-2 experiment, in particular Ref. [68], the FLASHForward experiment has demonstrated the sensitivity of the plasma acceleration process to the exact bunch and plasma profile, with exquisite control of both necessary to achieve high-quality, repeatable acceleration. This inadvertently highlights the suitability of the drive bunch itself as an effective probe of the plasma state. By improving our understanding of the coupling between the state of the plasma and the drive bunch, it would be possible to extract information about the plasma from measurements of the bunch. This approach can be used to study the evolution of the plasma by adjusting the arrival time of the bunch with respect to its generation, with potentially significantly improved sensitivity compared to other diagnostics.

This chapter outlines the development of exactly such a beam-based diagnostic technique. First, existing plasma density diagnostics are introduced before the beam-based method is motivated and compared. Next, a description of the experimental implementation of this technique and all associated analysis is outlined. This technique underpins further measurements of the extended timescale evolution of plasma sources and plasma wakefields that are discussed in later chapters.

5.1 EXISTING PLASMA DENSITY DIAGNOSTICS

In order to understand and place limitations on the repetition rate of future plasma-based accelerator facilities, studies of the evolution of plasma on extended timescales must take

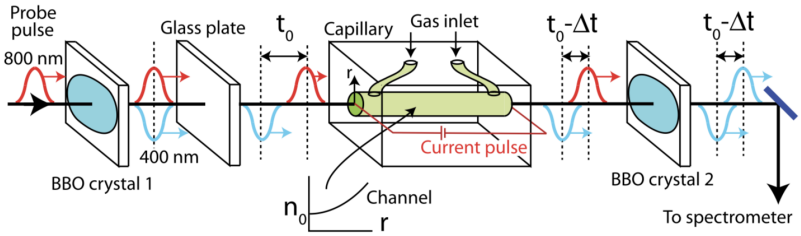

Figure 5.1.1: Sketch of the two-colour laser interferometry measurement principle. A detailed description can be found in the text. *Reproduced from Ref. [200].*

place such that their lifetime, and the lifetime of any perturbations induced by the wakefield process, can be identified. Many diagnostic techniques exist for probing plasma densities, a selection of which are outlined in the following section that are readily used in the context of PWFA experiments and hence provide useful comparison points for the beam-based method introduced here. The diagnostic techniques discussed are two-colour laser interferometry (TCI) and spectroscopic measurements of the broadening of characteristic emission lines, both of which are available within a dedicated plasma source development laboratory at DESY, while previous beam-based measurements are also presented.

5.1.1 *Two-colour laser interferometry*

The principle behind the common-path two-colour laser interferometry technique relies on the difference between the shift in phase and envelope accumulated by two laser pulses of different frequencies as they propagate through a plasma [200, 201]. A sketch of the measurement principle is demonstrated in Figure 5.1.1. Within a TCI measurement, two short pulses of differing wavelength are generated via the propagation of a single laser pulse through a frequency-doubling beta-barium borate (BBO) crystal. This creates two pulses that are intrinsically locked in phase and timing, one with a wavelength equal to that of the original pulse ("red"), and one with half this wavelength ("blue"). The two pulses are directed through a glass block to provide a constant temporal offset, t_0 in Figure 5.1.1, as a result of their differing frequency, "blue" trailing "red". The pulses then propagate through the plasma accumulating a density-dependent phase shift and envelope slippage, according to their frequencies, and are then passed through a second BBO crystal. The "red" pulse is frequency-doubled in the second crystal but its envelope slippage and phase advance are conserved [201]. The two 400 nm pulses are then imaged on to a slit into a spectrometer such that an interference pattern can be observed as demonstrated in Figure 5.1.2.

By extracting the phase shift observed in the interferogram as a function of time, it is possible to extract the evolution of the plasma density via the relation:

$$n_e = \frac{4\varepsilon_0 m_e c}{3e^2}\frac{\omega_0}{L_p}\Delta\phi, \tag{5.1.1}$$

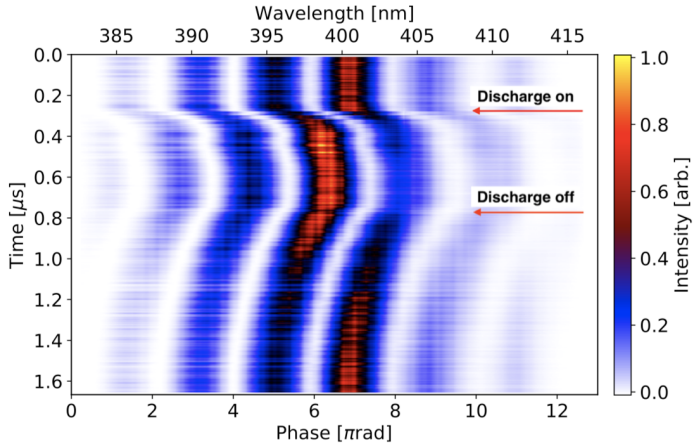

Figure 5.1.2: Example measurement of the evolution of the interference pattern created via the TCI technique as a function of time during a plasma discharge. *Reproduced from Ref. [195].*

where ω_0 is the fundamental frequency of the laser ($\lambda_0 = 800\,\text{nm}$), L_p is the plasma length, and $\Delta\phi$ is the total integrated phase shift accumulated during propagation through the plasma [200, 201]. As indicated by Eq. (5.1.1), the minimum measurable plasma density is a function of the resolution of the spectrometer, i.e. the measurement of $\Delta\phi$, and the plasma length L_p. At FLASHForward, the length of the capillaries used are typically tens of mm, limiting the minimum plasma density that can be resolved using this technique to densities significantly exceeding $10^{15}\,\text{cm}^{-3}$ [195]. This technique is therefore well-suited to accurately determining operating plasma densities ($n_e \sim 10^{16}\,\text{cm}^{-3}$), but its suitability for measuring the extended evolution of plasma sources, crucial for X-3 experiments studying high-repetition-rate operation, is limited.

5.1.2 *Spectral line broadening*

A second technique that is commonly used to measure the density of a plasma is the broadening of characteristic emission lines via the Stark effect [202], where spectral lines can be shifted or split by the presence of an external electric field. Within a plasma, even if the macroscopic electric field is zero, ions and electrons within the plasma move generating local electric fields that lead to the broadening of spectral emission lines [203]. By measuring the temporally- and spectrally-resolved emission characteristics of a plasma, it is possible to study its density evolution.

Typically, measurements of spectral line broadening are performed around the H_α line of the Balmer series as this exhibits a temperature-dependent power law relationship between the broadening observed and the plasma electron density over the range of interest for PWFA experiments ($n_e \sim 10^{15} - 10^{18}\,\text{cm}^{-3}$) [195]. To do this, gases must be doped with few-percent levels of hydrogen and care taken to determine emission lines characteristic of the different gas species. Figure 5.1.3 demonstrates such a measurement; (a) shows the emission spectra of a pure argon plasma, used to identify peaks within the spectra that could contribute to the

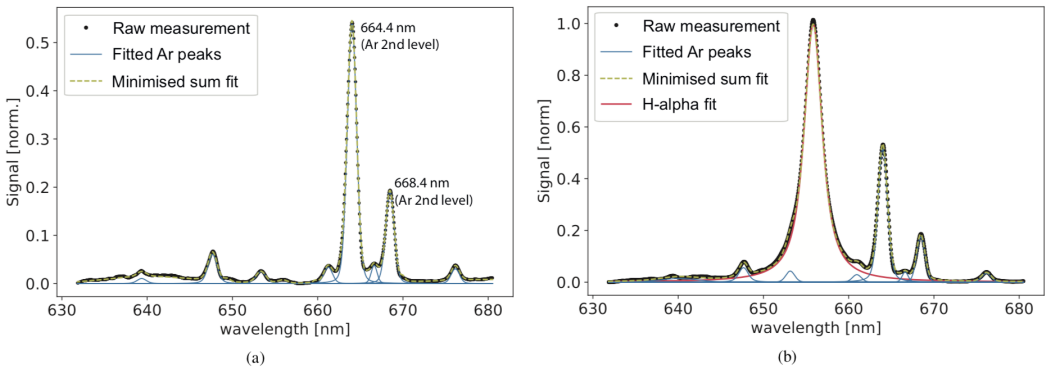

Figure 5.1.3: Example measurement of the emission spectra from (a) pure argon and (b) argon doped with 5% hydrogen. *Reproduced from Ref. [195].*

appearance of spectral broadening of the H_α line shown in (b) for an argon plasma doped with 5% hydrogen. By measuring the full width at half maximum (FWHM) of the H_α spectral line, $\Delta\lambda$, the plasma density can be extracted using methodology developed by, for example, Gigosos and Cardenoso [204, 205], using the simple relation:

$$n_e = A \cdot \Delta\lambda^B, \tag{5.1.2}$$

where A and B are constants derived from fits to simulated data that is bench-marked against experimental measurements [206]. Emission light from the plasma can be collected from any region of the plasma, allowing this technique to be used to spatially-resolve the plasma density profile. For example, the evolution of the longitudinal density profile of discharge-ionised plasmas was measured in this way in Ref. [195]. The light emitted by the plasma must be collected over a limited temporal period such that the plasma density does not vary significantly during the measurement [$\mathcal{O}(10)$ ns]. This limits the sensitivity of this technique as the amount of light emitted by the plasma decreases as the plasma density decreases. Within the dedicated plasma source diagnostic laboratory at FLASHForward, the signal-to-noise ratio becomes too low to accurately determine the plasma density for densities below $10^{16}\,\mathrm{cm}^{-3}$ despite the use of an intensified-CCD camera [195]. This can be improved by extending the exposure time of the camera, but at the cost of temporal resolution. Therefore, while spectral line broadening represents an accurate plasma density diagnostic at typical operating densities that additionally permits extraction of spatial information, in its current setup it does not have sufficient temporal (and hence density) resolution to allow measurement of the extended evolution of plasmas and the wakefield process necessary to advance the X-3 experimental goals.

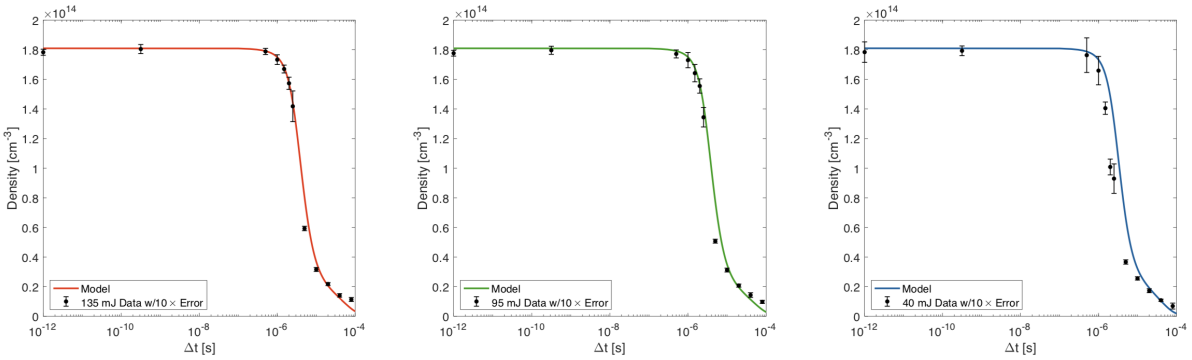

Figure 5.1.4: Evolution of the on-axis density of a laser-ionised plasma column measured using the modulation frequency of a self-modulated proton bunch. Three different ionising laser pulse energies were measured, 135 mJ (left), 95 mJ (centre) and 40 mJ (right). *Reproduced from Ref. [207].*

5.1.3 *Previous beam-based plasma density measurement techniques*

Two experiments have previously used PWFA drive bunches to probe the density of plasmas [207, 208]. Both measurements relied on the modulation of a long ($\sigma_z \gg \lambda_p$) bunch induced by plasma wakefields driven in the linear regime.

It was previously demonstrated in the AWAKE experiment that a long proton bunch can become transversely modulated at the plasma frequency by the wakefield it drives when propagating through a plasma [153]. Therefore, the plasma density can be extracted by measuring the modulation frequency of the proton bunch. This is the approach taken in Ref. [207]. During this experiment, the arrival time of the proton bunch was varied with respect to the generation of the plasma and its modulation frequency was measured using a streak camera. The rubidium plasma was generated using an ionising laser pulse and the density evolution of the resulting plasma column was studied as a function of the laser pulse energy as shown in Figure 5.1.4. The measurements demonstrated reasonable agreement with a model describing the expansion and recombination of the plasma column. This measurement technique is highly accurate but is limited in its density range that can be probed. Within the current AWAKE experimental setup, the streak camera used to measure the modulation of the proton bunch has picosecond-level resolution [153] and therefore its temporal resolution is not sufficient to allow plasma densities above approximately $10^{15}\,\mathrm{cm}^{-3}$ to be measured. As discussed extensively in Chapter 3, it takes a significant propagation distance within the plasma for the proton bunch to become self-modulated. As the plasma density lowers, the distance over which this occurs increases and hence multiple-metre length plasmas must be used (as a reminder, the plasma length at AWAKE is 10 m). The minimum plasma density measured within Ref. [207] was approximately $10^{13}\,\mathrm{cm}^{-3}$, multiple orders of magnitude lower than TCI or spectral broadening techniques. However, if TCI were used on a plasma of similar length, the sensitivity achieved in Ref. [195] would suggest plasma densities approaching $10^{12}\,\mathrm{cm}^{-3}$ could be accurately measured.

Ref. [208] similarly used modulation of a long bunch in plasma to estimate the plasma density. In this case, a low energy electron bunch was used and its longitudinal phase space was measured at the exit of the plasma using a transverse deflecting structure combined with a dipole. The argon plasma was contained within a cell of length 100 mm and was ionised using a 2.4 kV discharge. The longitudinal phase space of the bunch became modulated at the plasma frequency and hence Fourier analysis allowed determination of the plasma density as a function of the relative arrival time of the bunch with respect to the discharge initiation. This technique enabled plasma density measurement over an approximate range of $10^{13} - 10^{16}\,\mathrm{cm}^{-3}$, although determination of the modulation frequency at high densities was complicated by the appearance of multiple peaks within the Fourier spectra due to the evolution of the phase velocity of the wakefield during modulation. This effect is minimised at AWAKE due to the high rigidity of the highly-relativistic proton bunch. The beam-based measurements in Ref. [208] were compared to spectroscopic measurements of H_α line-broadening, showing reasonable agreement.

Both beam-based measurements discussed here rely on the measurement of the modulation of long bunches and hence are not necessarily compatible with high-gradient electron-driven PWFA facilities where the bunch is compressed to produce peak currents in excess of 1 kA. They are also limited at high plasma densities by the inability to accurately determine the modulation frequency. The beam-based characterisation technique developed at FLASHForward and discussed in the rest of this chapter represents a far simpler measurement that is compatible with high-gradient PWFA facilities and gives access to a larger range of plasma densities. Additionally, thanks to its simplicity and sensitivity, it enables online plasma density monitoring that can be used while experiments are being performed.

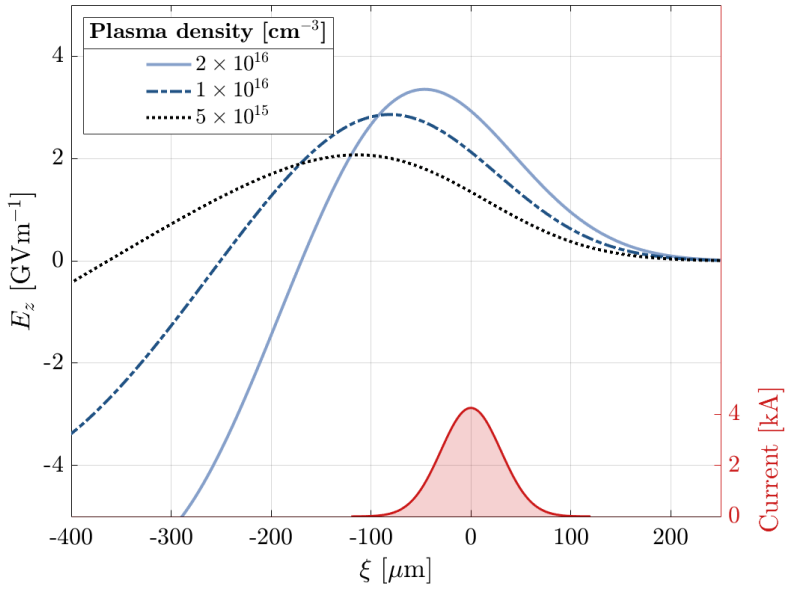

Figure 5.2.1: Comparison between the longitudinal field driven by a bi-Gaussian distributed electron bunch of RMS length 100 fs and charge 1 nC focused to a transverse size of 10 µm in plasma of density equal to $2 \times 10^{16}\,\mathrm{cm}^{-3}$ (light blue, solid), $1 \times 10^{16}\,\mathrm{cm}^{-3}$ (dark blue, dashed), and $5 \times 10^{15}\,\mathrm{cm}^{-3}$ (black, dotted) according to Eq. (5.2.1). The corresponding bunch current, $I(\xi)$, is shown in red.

5.2 BEAM-BASED PLASMA CHARACTERISATION

5.2.1 *Beam response to plasma of varying density*

The amplitude of the longitudinal wakefield driven by a highly-relativistic charged-particle bunch as it propagates through a plasma is dependent on the plasma density and beam current. For a thin bunch ($k_p \sigma_r \lesssim 0.1$), the on-axis longitudinal wakefield can be approximated through the relation

$$\frac{E_z(\xi)}{E_0} \simeq [1 - k_p \sigma_r K_1(k_p \sigma_r)] \int_\infty^\xi d\xi' \cos\left[k_p\left(\xi - \xi'\right)\right] \rho_{\parallel}(\xi'), \tag{5.2.1}$$

where σ_r is the RMS transverse bunch size, K_1 is the first-order modified Bessel function and $\rho_{\parallel}(\xi)$ describes the longitudinal bunch profile. Equation (5.2.1) represents the Green's function solution for the plasma response to a relativistic charged-particle bunch [85]. An example of the wakefield driven by a bi-Gaussian electron bunch of RMS length 100 fs and 1 nC charge focused to an RMS transverse size of 10 µm in plasma of varying density is shown in Figure 5.2.1. The relation in Eq. (5.2.1) implies that for an unchanging bunch profile, i.e. constant $I_b(\xi)$, the density of the plasma can be estimated via measurement of the amplitude of the wakefield it drives. While it is not possible to directly measure the exact amplitude of the wakefield at a given position within the plasma, it is possible to infer the integrated wakefield amplitude by measuring the change in the energy spectrum of a bunch after it has

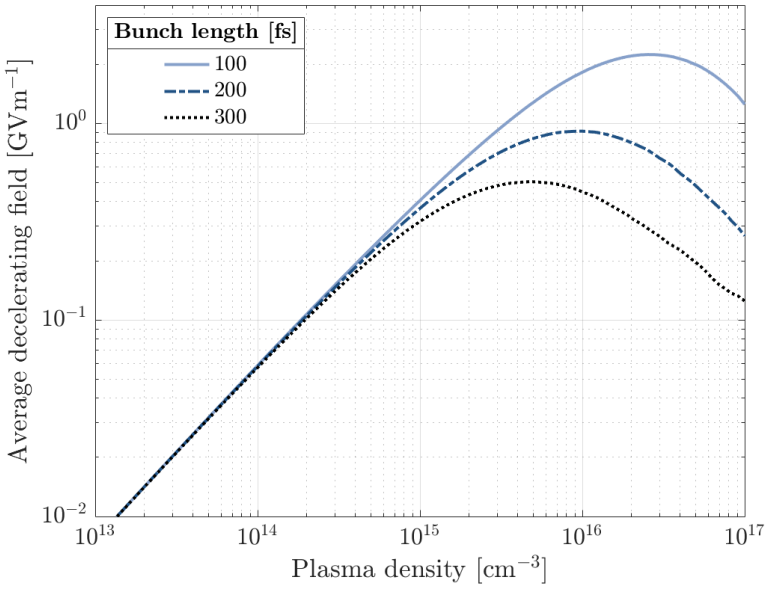

Figure 5.2.2: Average decelerating field experienced by a Gaussian electron bunch of charge 1 nC and RMS length of 100 fs (light blue, solid), 200 fs (dark blue, dashed) and 300 fs (black, dotted) as a function of the background plasma density.

propagated through a plasma of known length. This therefore allows determination of the longitudinally-integrated *effective* plasma density experienced by the bunch.

If the bunch is short with respect to the plasma wavelength ($\sigma_z < \lambda_p/2$), its entirety will fill the decelerating phase of the wakefield and the integrated wakefield amplitude can be calculated by measuring the energy loss of the bunch. However, at high densities the plasma wavelength is shorter ($\lambda_p \propto n_e^{-1/2}$) and the rear of the bunch can extend into the accelerating phase of the wakefield. At this point, the average decelerating field experienced by the bunch begins to decrease. The evolution of the average decelerating field experienced by an electron bunch of varying bunch lengths as a function of the background plasma density according to Eq. (5.2.1) is demonstrated in Figure 5.2.2. The average decelerating field is calculated using

$$\langle E_{dec} \rangle = \frac{\Sigma_\xi E_{z,\xi} dQ_\xi}{\Sigma_\xi dQ_\xi} \tag{5.2.2}$$

where $E_{z,\xi}$ is the amplitude of the wakefield at a position ξ in the co-moving frame, dQ_ξ is the bunch charge in slice ξ and the sum is performed over the electron-decelerating region ($E_z > 0$) of the wakefield. As expected, as the bunch length is reduced the bunch is able to drive significantly larger amplitude average decelerating wakefields, with the peak deceleration occurring at higher plasma densities. At the lowest plasma densities, the bunch size is significantly shorter than the plasma skin-depth, $k_p^{-1} = c/\omega_p$, and the plasma electron response (and hence wakefield) becomes less sensitive to the exact profile of the bunch, resulting in similar average decelerating fields for the three considered bunch lengths. Furthermore, the approximation in Eq. (5.2.1) becomes less accurate as the plasma density is reduced as the bunch density will begin to significantly exceed the plasma density. Once

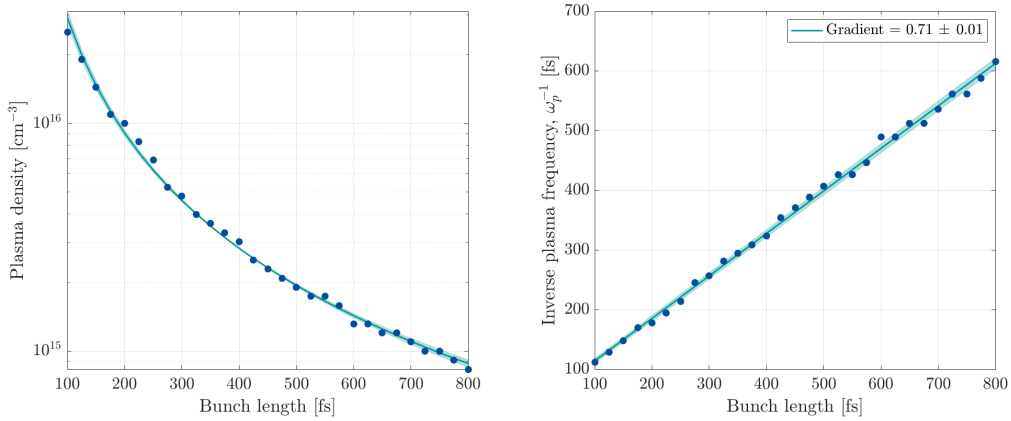

Figure 5.2.3: Left: Plasma density corresponding to the maximum decelerating field as a function of the Gaussian electron bunch RMS length. Right: Inverse plasma frequency at the maximum decelerating field as a function of RMS bunch length. A linear fit reveals the expected dependence, $\sigma_\tau \propto \sqrt{2}\,\omega_p^{-1}$.

$n_b \gtrsim 10 n_e$, linear wakefield theory has previously been shown to over-estimate the wakefield amplitude [209].

Figure 5.2.2 demonstrates that there is a plasma density at which the bunch is maximally decelerated for a given bunch length. The expected dependence of this density on the RMS length of the electron bunch can be calculated from Eq. (5.2.1) when assuming a Gaussian longitudinal profile of the form $\rho_{\parallel}(\xi) = -e n_b \exp(-\xi^2/2\sigma_z^2)$ [209]:

$$E_z(0,\xi) = -\left(\sqrt{2\pi}\,\frac{m_e c \omega_p}{e}\,\frac{n_b}{n_e}\left[k_p \sigma_z \exp(-k_p^2\sigma_z^2/2)\right] \times \left[1 - k_p\sigma_r K_1(k_p\sigma_r)\right]\right)\cos(k_p\xi). \tag{5.2.3}$$

Hence the expression for the longitudinal field is maximised for

$$\frac{dE_z(0,\xi)}{d(k_p\sigma_z)} \propto \left(1 - \frac{1}{2}k_p^2\sigma_z^2\right)\exp(-k_p^2\sigma_z^2/2) = 0, \tag{5.2.4}$$

$$\Rightarrow \sigma_z = \sqrt{2}\,k_p^{-1}, \tag{5.2.5}$$

a well-known result of linear wakefield theory [85]. This is demonstrated in Figure 5.2.3 which shows the plasma density at which the maximum decelerating field is observed as a function of the RMS bunch length according to Eq. (5.2.1). As expected, the inverse plasma frequency, ω_p^{-1}, scales linearly with the RMS bunch length, σ_τ, with gradient equal to $1/\sqrt{2}$. Hence, the magnitude of the average decelerating field experienced by a bi-Gaussian electron bunch in plasma can be approximated by [209, 210]

$$\langle E_z \rangle \approx \frac{m_e c \omega_p}{e}\,\frac{n_b}{n_e}\,k_p\sigma_z \exp(-k_p^2\sigma_z^2/2)\;\left[1 - k_p\sigma_r K_1(k_p\sigma_r)\right]. \tag{5.2.6}$$

5.2.2 *Using a short beam to measure plasma evolution*

Figure 5.2.2 demonstrates that for fixed bunch parameters the average decelerating field that acts on an electron bunch varies monotonically with the plasma density up until the RMS bunch length exceeds approximately

$$\sigma_z \gtrsim \sqrt{2}\, k_p^{-1} = \sqrt{2}\, c \left(\frac{n_e e^2}{m_e \varepsilon_0} \right)^{-1/2} . \tag{5.2.7}$$

This therefore means that short ($\sigma_z \sim 100\,\mathrm{fs} < k_p^{-1}$), high current ($\sim$kA) electron bunches can be used as a probe for a large range of plasma densities. By changing the relative time of arrival of an electron bunch with respect to the generation of a plasma and measuring the energy spectrum of the bunch after it has interacted with the plasma, the *effective* density of the plasma can be calculated as a function of time. This method is very similar to two-colour laser interferometry in that it measures the longitudinally-integrated effect of a plasma but has the distinct advantage that an electron bunch will still drive a wakefield of significant amplitude at low plasma densities ($n_e < 10^{16}\,\mathrm{cm}^{-3}$) allowing density measurement over much longer timescales provided the energy spectrum of the bunch can be measured with sufficient resolution following its interaction with plasma. In addition to this, it only probes the plasma density local to the propagation axis of the bunch — the region relevant to acceleration experiments — the relaxation of which is crucial to understanding high-repetition-rate operation limitations.

5.3 EXPERIMENTAL IMPLEMENTATION

Experiments to test and develop this beam-based diagnostic technique were undertaken at the FLASHForward facility. The electron bunch typically used within PWFA experiments was used as a probe bunch and allowed measurement of the evolution of a plasma over timescales up to 80 µs and densities as low as $10^{13}\,\mathrm{cm}^{-3}$, approximately two orders of magnitude lower than achievable using two-colour laser interferometry within the same setup. This section outlines the typical experimental setup used to execute these measurements and the analysis performed to determine the plasma density evolution. As an example, measurements performed to study the evolution of a krypton-filled discharge capillary are used, but the same technique is applied to all datasets discussed thereafter.

5.3.1 *Experimental setup*

The experiments at FLASHForward used a PWFA drive beam interacting with a plasma generated via high-voltage discharge. The electron bunch had a charge of $982.4 \pm 1.9\,\mathrm{pC}$ and was accelerated to an energy of $\mu_0 = 1030.1 \pm 0.1\,\mathrm{MeV}$ with an energy spread of $\sigma_0 = 1.5 \pm 0.1\,\mathrm{MeV}$. The bunch was compressed to an RMS length of $\sigma_\tau = 373 \pm 5\,\mathrm{fs}$ and the longitudinal phase space of the bunch was measured using an X-band transverse deflecting structure within the FLASHForward beamline and is shown in Figure 5.3.1. The bunch was strongly focused at the entrance of the plasma cell achieving β-functions on the order of 10 mm, estimated using the jitter of beam position measurements on either side of the interaction chamber [173]. After interacting with the plasma, the beam was captured and focused by a quadrupole triplet and dispersed vertically by a 1 m-long dipole magnet on to a 200 µm-thick scintillating crystal (LYSO). The dipole current was set to 390.0 ± 0.1 A allowing energies between $866.44 \pm 0.13\,\mathrm{MeV}$ and $1130.16 \pm 0.21\,\mathrm{MeV}$ to be measured as shown in Figure 5.3.2. The crystal was imaged by a digital camera with a sensor made up of 2048×1024 pixels. The corresponding physical size of the pixels in the image plane was measured using a resolution target and determined to be 25.3 ± 0.1 µm. Scattering of the electron bunch as it passed through the vacuum chamber separating the accelerator environment and the scintillating screen limited the spatial resolution of the spectrometer to 57 µm. The theoretical energy resolution of the spectrometer across the LYSO crystal is presented in Figure 5.3.2 for reference.

The discharge delivered a pulse of voltage 25 kV and duration of approximately 400 ns. The discharge is switched using thyratron with an intrinsic jitter measured to be on the order of a nanosecond. The discharge trigger is locked to the RF of the FLASH linac and operates at a frequency of 108.3 MHz, allowing temporal steps of 9.23 ns. The gas pressure within the capillary is not measured but the buffer volume pressure was kept constant at 15.1 ± 0.1mbar throughout the measurement.

To perform the measurement, the timing of the initiation of the discharge was changed with respect to the arrival of the electron bunch. Two scans were performed; the first at

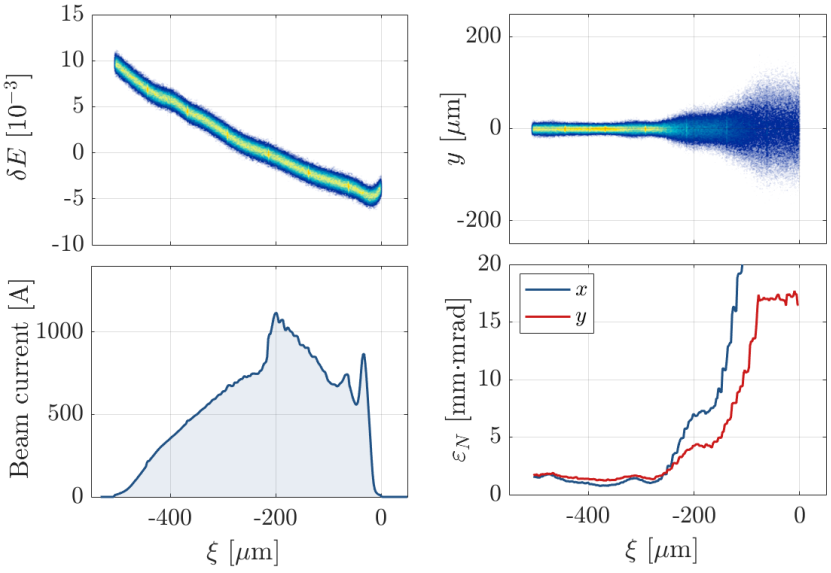

Figure 5.3.1: Reconstruction of the probe electron beam according to longitudinal and beam jitter phase space measurements. Upper left: Longitudinal phase space. Upper right: Transverse beam profile in the $y - \xi$ plane. Lower left: Beam current distribution. Lower right: Slice emittance in x and y planes.

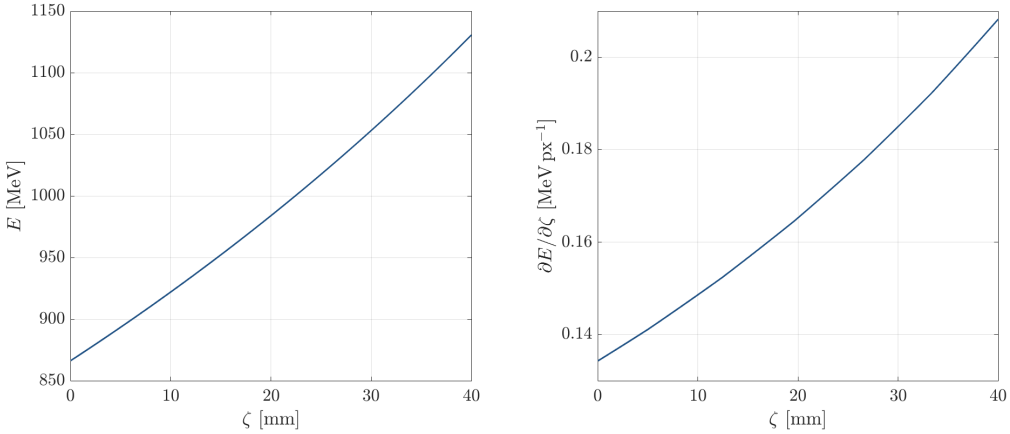

Figure 5.3.2: Left: energy-position relation in the plane of the scintillating crystal for the dipole spectrometer settings used in this measurement calculated via single particle tracking through measured dipole field maps. Right: Theoretical energy resolution of the spectrometer system, not including effects from scattering of the beam within the vacuum chamber.

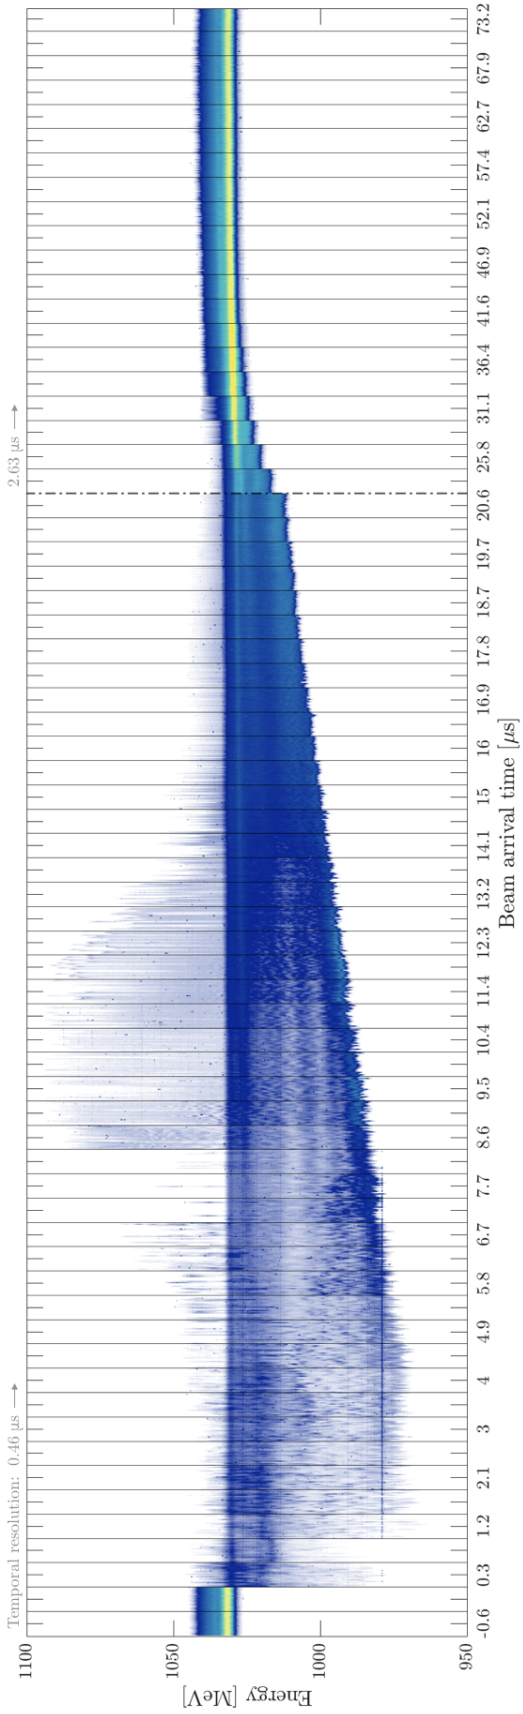

Figure 5.3.3: Waterfall plot demonstrating the measured probe beam energy spectra as a function of the beam arrival time with respect to the firing of the discharge in a krypton plasma with a buffer volume pressure of 15.1 mbar. This data was taken in two separate scans: the first at early times with a temporal resolution of 0.46 µs and the second over longer timescales with resolution of 2.63 µs. 20 events were recorded per step.

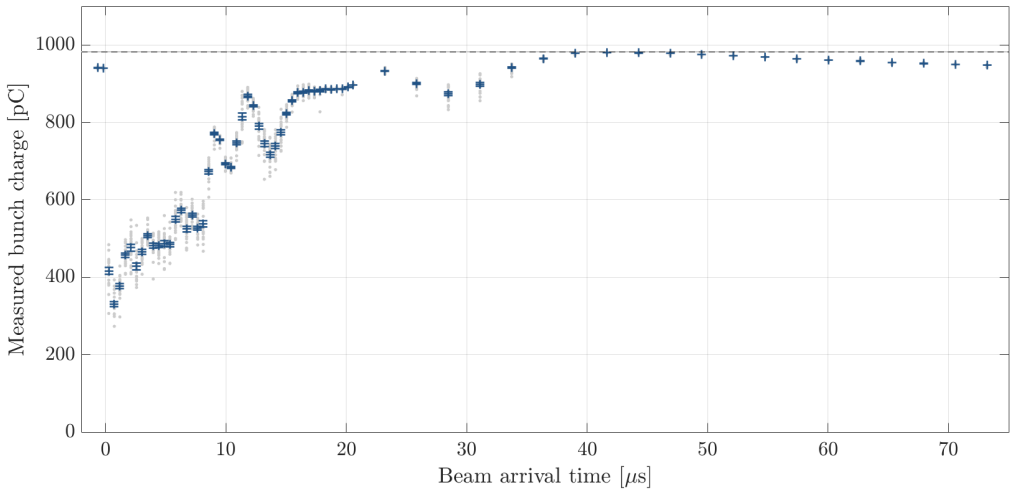

Figure 5.3.4: Electron bunch charge measured as a function of the beam arrival time during the discharge timing scan shown in Fig 5.3.3. The mean bunch charge is represented by blue crosses with error bars equal to the standard error on the mean. The individual measurements are shown by light grey dots. The dashed grey line represents the mean bunch charge measured at a toroid upstream of the plasma.

arrival times between $-0.65 - 20.58\,\mu s$ with a resolution of $0.46\,\mu s$, and the second at larger separations with arrival times between $20.58 - 73.19\,\mu s$ with a resolution of $2.63\,\mu s$. 20 events were recorded at each step of the scan, with the entire dataset comprising of 940 events. The measured energy spectra of the bunches following interaction with the plasma as a function of the relative arrival time of the bunch are displayed in Figure 5.3.3.

For negative arrival times, where the bunch arrives before the initiation of the discharge, the bunch is unaffected and passes straight through the capillary and is imaged on the spectrometer. These events represent a reference measurement of the energy distribution of the bunch without any plasma interaction and are used to calculate μ_0 and σ_0, the mean energy and energy spread of the non-interacted bunch respectively. Figure 5.3.4 shows the electron bunch charge measured at a BPM approximately 2 m downstream of the exit of the capillary for the discharge timing scan presented in Figure 5.3.3. Without any plasma present, the bunch charge is slightly lower than that measured by a toroid upstream of the plasma cell, represented by the dashed grey line. This is a result of the small β-functions required to achieve matching of the bunch into the plasma. To be able to produce β-functions on the order of 10 mm at the entrance of the plasma, the bunch must be strongly focused and hence also has a large divergence. When there is no plasma present in the capillary, the beam has no external focusing force acting on it and freely diverges. Therefore, a small fraction of particles hit the beam pipe during transport to the spectrometer and are lost, causing a reduction in the charge measured after the interaction chamber by approximately 4%.

At short times following the culmination of the discharge ($\lesssim 4\,\mu s$), the plasma density is highest and its wavelength, and correspondingly the size of the wakefield cavity, is shortest. As can be seen in the lower left panel of Figure 5.3.1, the beam extends over approximately

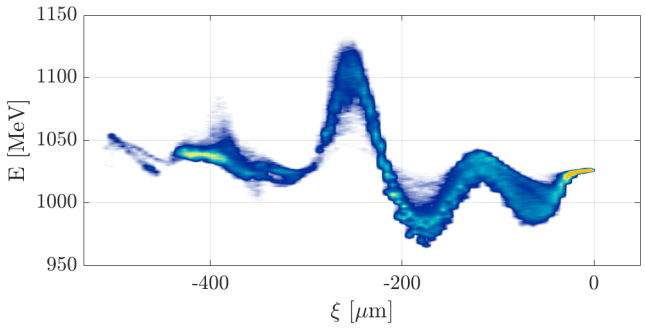

Figure 5.3.5: Simulated longitudinal phase space of the electron bunch following 50 mm propagation in plasma of density $7 \times 10^{16}\,\mathrm{cm}^{-3}$. The phase space becomes modulated as the bunch extends over multiple accelerating and decelerating phases.

500 μm. This distance corresponds to the wavelength of a plasma of density approximately $4.5 \times 10^{15}\,\mathrm{cm}^{-3}$. Hence, for densities larger than this value, the bunch will extend over more than one plasma period. The bunch extends into multiple decelerating and accelerating regions and the longitudinal phase space becomes modulated. This effect was used in Ref. [208] to infer the plasma density as discussed previously. An example of this is demonstrated in Figure 5.3.5 which shows the simulated longitudinal phase space of the electron bunch after propagating for 50 mm in a plasma of density equal to $7 \times 10^{16}\,\mathrm{cm}^{-3}$. When the bunch is longer than a single plasma period, significant parts of the bunch experience regions of transversely defocusing fields induced by the return of sheath plasma electrons at the rear of the wakefield cavities, causing a reduction in the amount of charge that is transported through the plasma and captured by the post-plasma quadrupoles as shown at short delays in Figure 5.3.4.

As the beam arrival time moves further from the initiation of the discharge, the plasma recombines and the density at the point of interaction decreases. At an arrival time of around 9 μs following the initiation of the discharge, the plasma has recombined such that its density is low enough that the wakefield cavity is of a similar length to the bunch. At this point, the majority of the bunch is in the transversely focusing phase of the wakefield and an increasing bunch charge is successfully transported through the plasma and imaged by the spectrometer. A significant amount of accelerated "*witness*" charge with energies exceeding those of the bunch without any plasma interaction ($E > 1035$ MeV) is also observed in the spectrometer images in Figure 5.3.3. At these densities, the wakefield is in the blow-out regime and the rear of the electron bunch is approximately matched into the resulting ion column [188–190]. This means the envelope of the bunch does not oscillate in this region and the witness particles are simultaneously focused and accelerated by the wakefield and are successfully transported to the electron spectrometer.

A modulation in the measured bunch charge is observed in Figure 5.3.4, with local minima at approximately 10.5 μs and 13.5 μs and maxima at approximately 9 μs and 11.8 μs after the initiation of the discharge. In the blow-out regime the focusing force in the rear of the wakefield is dependent on the density of ions in the channel. Electrons in the bunch undergo

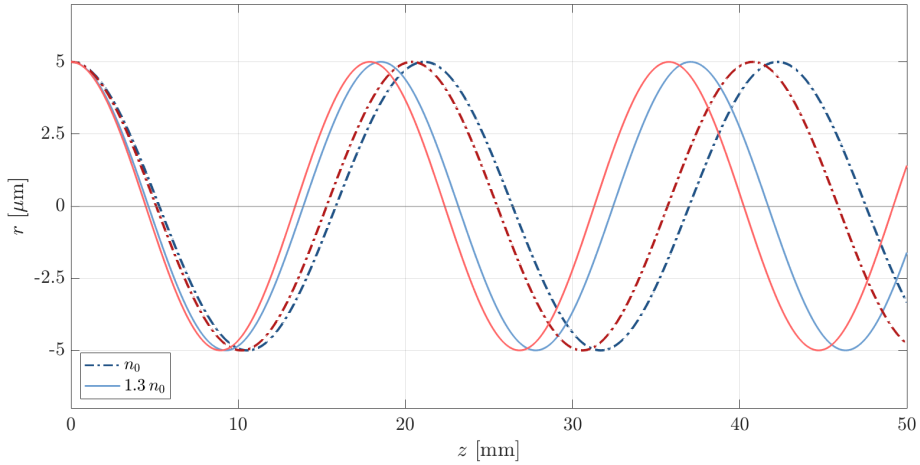

Figure 5.3.6: Example single electron betatron oscillation trajectories for different plasma densities and electron energies. As the density is changed (dashed or continuous lines) the phase of the oscillation at the exit of the plasma changes. The difference in phase advance induced by a small change in the density varies for different electron energies (blue or red). These trajectories assume a background density of $1.0 \times 10^{16}\,\mathrm{cm}^{-3}$ (dashed line) and $1.3 \times 10^{16}\,\mathrm{cm}^{-3}$ (continuous line) and electron energies of 1020 MeV (blue) and 950 MeV (red) and are calculated using Eq. (5.3.1).

transverse betatron oscillations in the plasma with a frequency dependent on their energy and the plasma density:

$$\omega_\beta = \frac{\omega_p}{\sqrt{2\gamma}}. \tag{5.3.1}$$

Therefore, different slices of the bunch (with different energies) oscillate at different frequencies and exit the plasma having acquired a different total betatron phase. An electron is most likely to be captured by the post-plasma quadrupoles if its angle to the longitudinal axis of propagation as it leaves the plasma is small, i.e. it is near the maximum amplitude of its betatron oscillation, and vice versa if the electron is near a zero-crossing of its trajectory. As the bunch arrival time, and hence plasma density, changes, the number of oscillations that individual energy slices of the beam undergo in the plasma correspondingly changes as shown in Figure 5.3.6. This results in different portions of the bunch having small enough divergence at the plasma exit to be successfully captured by the post-plasma quadrupoles and transported to the spectrometer as a function of the plasma density, resulting in oscillations in the measured post-plasma bunch charge as a function of beam arrival time as observed in Figure 5.3.4. The frequency of the measured bunch-charge oscillation in Figure 5.3.4 lowers as the plasma density decreases and the betatron frequency correspondingly lowers.

At even lower plasma densities, approximately 40 μs after the initiation of the discharge in Figure 5.3.4, the entire bunch is successfully transported through the plasma and captured by the post-plasma quadrupoles, leading to agreement between measurements of the bunch charge before the plasma (grey dashed line) and after (blue data points). In this case,

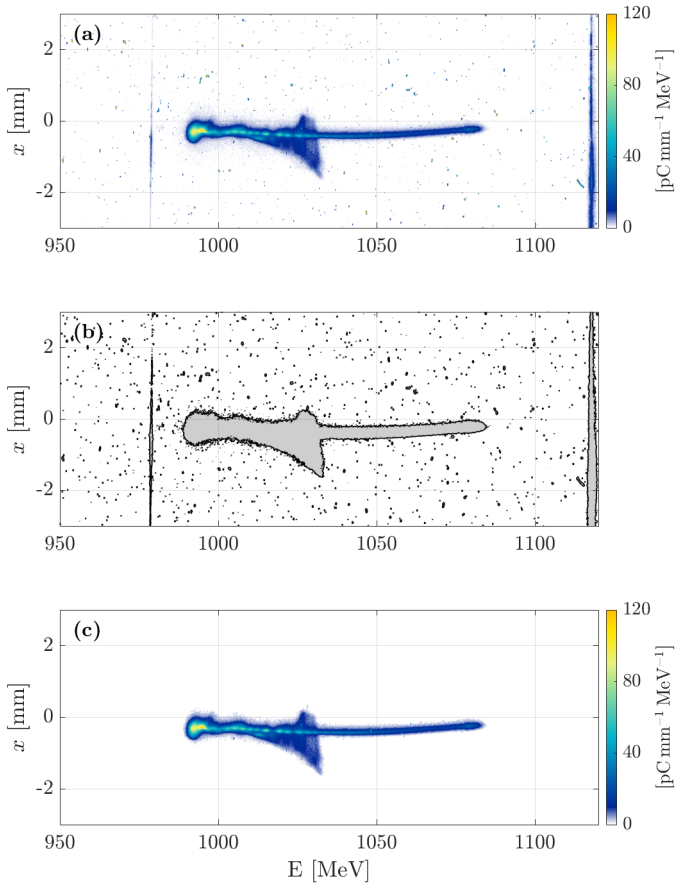

Figure 5.3.7: Example of spectrometer image analysis and signal identification techniques. (a) Raw spectrometer image from the dataset shown in Figure 5.3.3 corresponding to a beam arrival time of 11.35 µs. (b) All contours identified within the image with a continuous raw pixel count above the cutoff value. (c) Spectrometer image following the signal identification procedure outlined in the text.

the interaction between the bunch and the low density plasma means the plasma acts as a passive plasma lens [94, 211] such that the low-amplitude, radially-focusing transverse wakefield response balances the natural divergence of the bunch induced by its small β-function and results in zero loss of charge. After this, as the plasma density continues to decrease, the wakefield interaction strength further decays and the charge measured after the plasma decreases and returns to the value with no plasma present ($\Delta\tau < 0$, Figure 5.3.4).

5.3.2 *Analysis of spectrometer images*

Use of a LYSO scintillating crystal is advantageous due to its short decay time, however its increased sensitivity is consequently particularly susceptible to the effect of background particles interacting with the scintillating material. This is demonstrated in Figure 5.3.7(a) which shows a typical raw spectrometer image taken during the dataset presented in Fig-

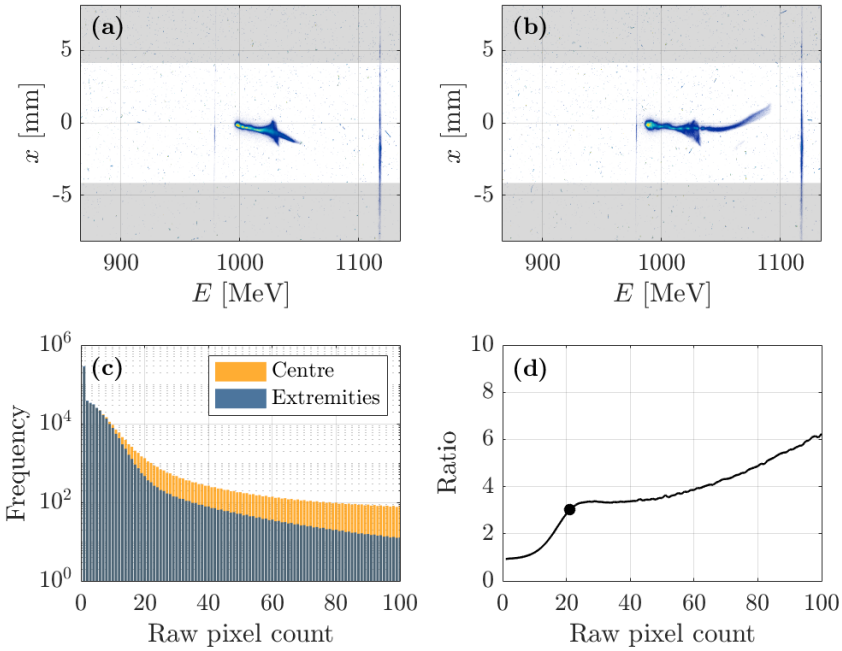

Figure 5.3.8: Example of the technique used to identify the pixel cutoff value when identifying the signal region in a spectrometer image for the dataset shown in Figure 5.3.3. (a) and (b) Raw spectrometer images from within the dataset, with the upper and lower extremities indicated in grey. (c) Histogram of the average pixel count across all images in the dataset. The histogram of the central region is plotted in orange and the histogram of the upper and lower extremities of the image (grey regions in (a) and (b)) is plotted in blue. (d) The ratio of the frequency of raw pixel counts in the central and upper and lower extremeties of the spectrometer images. The raw pixel count for which this ratio exceeds 3 is used as the pixel cutoff value, in this case 21, indicated by the black circle.

ure 5.3.3. A large number of small regions of pixels with low counts are observed away from the main signal region and can significantly alter measurements of the energy distribution of the bunch if not removed. The position of these regions varies on a shot-to-shot basis, indicating its source is the interaction of background particles with the crystal rather than an issue with the camera sensor itself. In addition to this, light bleeding out of the edge of the scintillating crystal can be observed at the right hand edge of the image and at a pixel position corresponding to an energy of approximately 980 MeV.

For these reasons, signal identification based on extracting contours of the image was used. A minimum cutoff pixel count value was defined, and all connected areas in the image with pixel counts above this value are identified. The cutoff pixel value was chosen by comparing the occurrence frequency of raw pixel counts in different regions of the spectrometer images as shown in Figure 5.3.8. For each image in the dataset, for example those shown in Figure 5.3.8(a) and (b), the upper and lower quarters of the image (shown in grey) are separated and considered to be regions that do not consist of the signal from the bunch which is purposefully centred within the image during the experiment. A histogram of the aver-

age frequency of raw pixel counts in these upper and lower regions across all images in the dataset is then calculated and is shown in blue in Figure 5.3.8(c). Identical analysis is then performed on the remaining central region of all images, where the scintillating signal from the bunch is observed, and plotted in orange in Figure 5.3.8(c). By comparing the two resulting histograms, it is clear that for the smallest raw pixel counts ($0-10$) the distribution in both regions is very similar, but beyond this, higher raw pixel counts occur more frequently in the central region of the images than in the extremities. The ratio of these two distributions can therefore be used to define a raw pixel count value at which the corresponding pixel likely belongs to a region of signal from the bunch rather than representing background. The cutoff pixel count value was chosen to correspond to the raw pixel count at which the ratio of the frequency exceeds 3. For this particular dataset, this corresponds to a raw pixel count of 21 and a calibrated spectral density of approximately $1\,\mathrm{pC\,mm^{-1}\,MeV^{-1}}$, although this is calculated separately for every dataset to account for possible changes in the camera setup. Note, the cameras used within the spectrometer at FLASHForward have 12-bit sensors, and hence it is possible for a pixel to have any value in the range $0-4095$. However the background effects discussed here typically have low pixel values and hence only raw pixel counts in the range $0-100$ are plotted in Figure 5.3.8(c) and (d).

Once the minimum pixel cutoff value is defined, regions of continuous areas with pixel counts above this value within a spectrometer image are identified. An example of the result of this identification procedure is shown in Figure 5.3.7(b). A large number of contours are found (shown in grey with black outline), and subsequently sorted according to their area. Background hits can be seen to be restricted to small regions of few pixels and hence the contour with the largest area is automatically selected as the signal of the bunch, with the pixel counts of the rest of the image set to zero. The result of this signal identification produce is shown in Figure 5.3.7(c) and demonstrates successful isolation of the signal region from the background noise.

Following identification of the signal, the energy distribution of the bunch can be calculated. To do this, the spectrometer image in Figure 5.3.7(c) is summed over the transverse plane to give its projection in the dispersive plane as demonstrated in Figure 5.3.9. The charge-weighted mean energy of the bunch is calculated according to:

$$\mu_E = \frac{\sum_i E_i \cdot dQ_i}{\sum_i dQ_i}, \tag{5.3.2}$$

while the RMS energy spread is similarly calculated via:

$$\sigma_E = \left(\frac{\sum_i E_i^2 \cdot dQ_i}{\sum_i dQ_i} - \mu_E^2 \right)^{1/2}, \tag{5.3.3}$$

where i represents the columns of the signal region of the image, E_i is the energy associated with column i, and dQ_i is the total charge measured within column i. As discussed previously, to estimate the density of the plasma, the average decelerating field experienced by bunch must be calculated. Therefore, the summation is performed over the decelerated

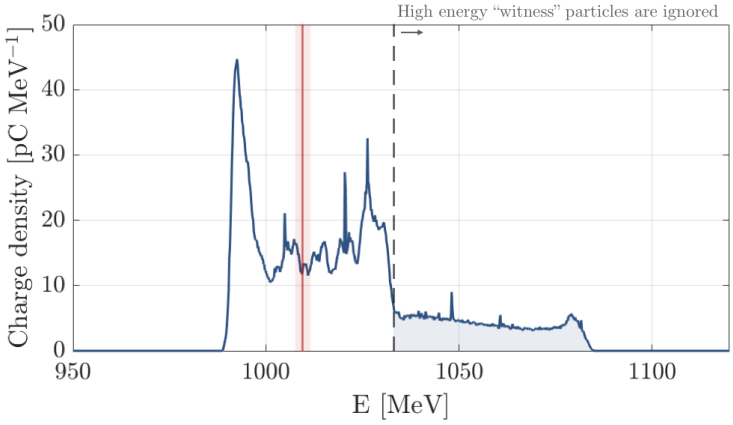

Figure 5.3.9: Energy projection of a spectrometer image with the mean energy, $\mu_E = 1009.5 \pm 1.3\,\text{MeV}$ marked in red. Any "witness" particles with energies exceeding $\mu_0 + 2\sigma_0 = 1033.1\,\text{MeV}$ (represented by the grey dashed line) are removed from the calculation.

region of the bunch with a cutoff energy defined by $\mu_0 + 2\sigma_0$ used to identify accelerated particles as shown by the shaded blue region in Figure 5.3.9. The uncertainty associated with these calculations includes contributions from the energy resolution of the spectrometer (see Figure 5.3.2) and the charge calibration of the spectrometer screen which is typically of order 10%. The calculated mean energy of $\mu_E = 1009.5 \pm 1.3\,\text{MeV}$ for the event in Figure 5.3.7 is represented by the red vertical line and shaded region, while its RMS energy spread is calculated to be $\sigma_E = 13.4 \pm 0.3\,\text{MeV}$. Assuming a plasma length of 50 mm, equal to the length of the capillary, this therefore represents an average decelerating gradient of $411 \pm 27\,\text{MVm}^{-1}$ for this particular event.

For this dataset, twenty events were measured per relative bunch arrival time. Identical analysis is performed on the spectrometer image from each event, and these results are combined to give a mean bunch energy of $1009.7 \pm 0.3\,\text{MeV}$ where the total uncertainty, $\sigma = 0.3\,\text{MeV}$, is calculated from a combination of the uncertainty on each individual measurement, σ_i, via:

$$\frac{1}{\sigma^2} = \sum_{i=1}^{20} \frac{1}{\sigma_i^2}. \tag{5.3.4}$$

This therefore results in the determination of an average decelerating gradient of $407 \pm 8\,\text{MVm}^{-1}$ for this value of the relative bunch arrival time. This procedure is repeated for each value of the relative bunch arrival time to give the average decelerating gradient as a function of the time after the initiation of the discharge.

5.3.3 *Plasma density reconstruction*

The longitudinally-averaged effective plasma density experienced by the bunch can be reconstructed by comparing the experimentally-determined average decelerating field experienced

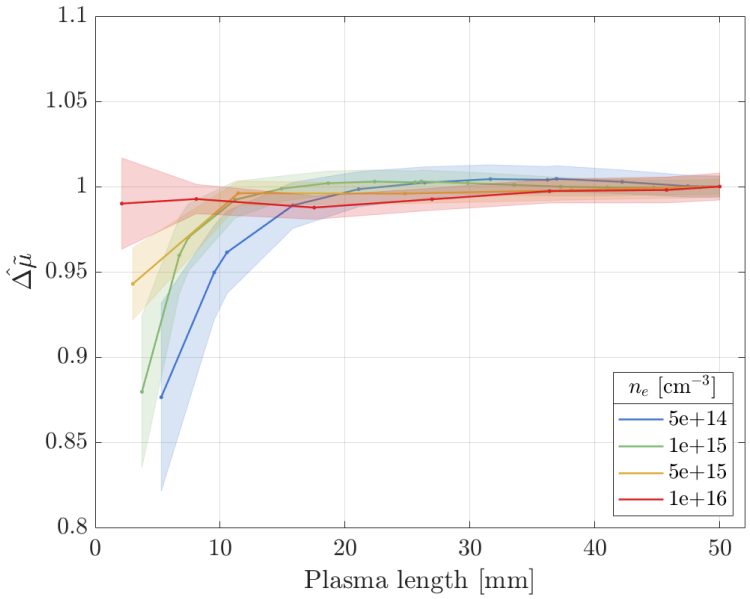

Figure 5.3.10: Simulated evolution of the value of $\Delta\tilde{\mu}$ as a function of the propagation distance of a bunch through plasma for a range of plasma densities. $\hat{\Delta\tilde{\mu}} = \Delta\tilde{\mu}(z)/\Delta\tilde{\mu}|_{L_p}$ represents the relative value of $\Delta\tilde{\mu}$ at a particular propagation distance, z, compared to the value at the plasma exit, $\Delta\tilde{\mu}|_{L_p}$. The shaded regions represent the uncertainty associated with the mean energy of the bunch calculated in simulation.

by the bunch to that predicted by three-dimensional PIC simulations. To be able to do this, the phase space of the bunch must be accurately measured using the diagnostic techniques introduced in Chapter 4. This allows accurate reconstruction of not only the six-dimensional phase space of the bunch in simulation but also its slice-dependent focusing at the start of the plasma. The detailed reconstruction of the bunch used in the measurements discussed in this chapter are shown in Figure 5.3.1.

The metric introduced to compare the experimental measurements to the output of the PIC simulations is the *fractional energy perturbation per unit length*,

$$\Delta\tilde{\mu} = \frac{\mu_0 - \mu_E}{\mu_0 \cdot L_p}, \tag{5.3.5}$$

where μ_0 is the mean energy of the bunch without plasma interaction, μ_E is the mean energy of the bunch after plasma interaction, and L_p is the plasma length, assumed to be equal to the length of the capillary, 50 mm in this case. This metric represents a measure of the coupling strength between the electron bunch and the plasma and can be seen in simulation to saturate at a constant value as the bunch propagates through the plasma. This is illustrated in Figure 5.3.10 which shows the simulated evolution of the value of $\Delta\tilde{\mu}$ as a function of the propagation distance for a range of plasma densities. The value of $\Delta\tilde{\mu}$ initially varies as the bunch enters the plasma and undergoes transverse beam size oscillations as it is not optimally matched to the plasma density [190]. This beam size oscillation results in similar oscillations in the bunch density and hence wakefield amplitude as the bunch focuses and approaches

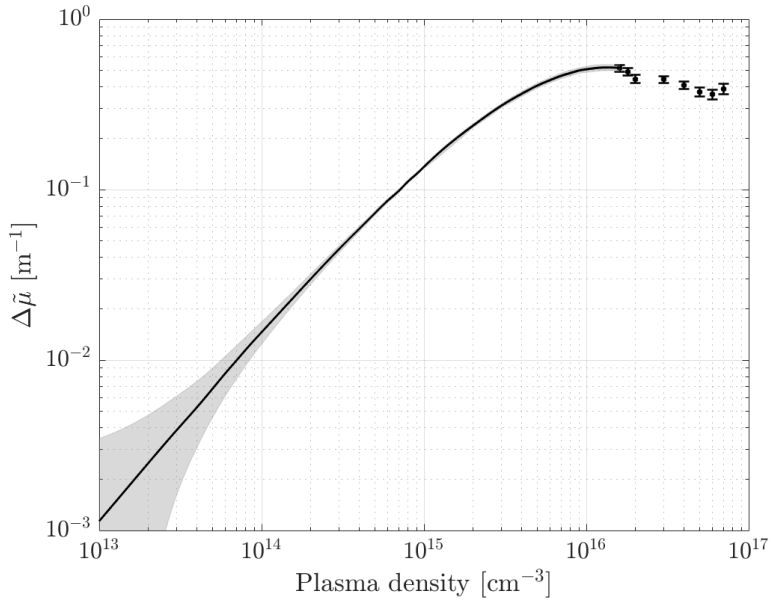

Figure 5.3.11: Simulated relative energy perturbation per unit length, $\Delta\tilde{\mu}$, as a function of the plasma density for the bunch profile reconstructed in Figure 5.3.1. The grey shaded regions represent the uncertainty on the calculation. Data points at high densities ($n_e \gtrsim 2 \times 10^{16}\,\mathrm{cm}^{-3}$) where a large fraction of bunch charge is lost experimentally are plotted for reference.

a stable propagation mode. The rate at which these transverse bunch size oscillations occur is defined by the betatron oscillation frequency, $\omega_\beta = \omega_p/\sqrt{2\gamma}$, and hence scales with the square root of the plasma density. This means the beam reaches a stable propagation mode, allowing it to drive a wakefield of consistent amplitude, more quickly at higher densities, as demonstrated in Figure 5.3.10. For the bunch used in these measurements, simulations indicate the saturated $\Delta\tilde{\mu}$ value is reached within 30 mm of propagation within the plasma for plasma densities exceeding $10^{14}\,\mathrm{cm}^{-3}$. At densities lower than this, $\Delta\tilde{\mu}$ saturation is not reached until further propagation distances in simulation and hence larger uncertainties are associated with the inferred beam-based plasma density measurement.

A series of three-dimensional PIC simulations of the interaction of the electron bunch with plasmas of density varying from $10^{13} - 10^{17}\,\mathrm{cm}^{-3}$ were performed. A constant resolution of $\Delta x = \Delta y = \Delta\xi = 0.02 k_p^{-1}$ was used, except at low densities where this resolution would be insufficient to fully resolve the profile of the electron bunch — in this case, a constant number of grid cells was used, with $N_x = N_y = N_\xi = 100$. The simulations used longitudinal and transverse plasma profiles of constant density in a plasma of length 50 mm. The bunch distribution at the exit of the plasma was used to recreate a simulated energy spectrum measurement with resolution equal to that available experimentally ($\delta E \sim 0.2$ MeV). The simulated $\Delta\tilde{\mu}$ was calculated from this as a function of the plasma density and is shown in Figure 5.3.11.

As expected, the simulated $\Delta\tilde{\mu}$ varies monotonically over a wide range of plasma densities until the plasma density is sufficiently high ($n_e \gtrsim 2 \times 10^{16}\,\mathrm{cm}^{-3}$) that the bunch signifi-

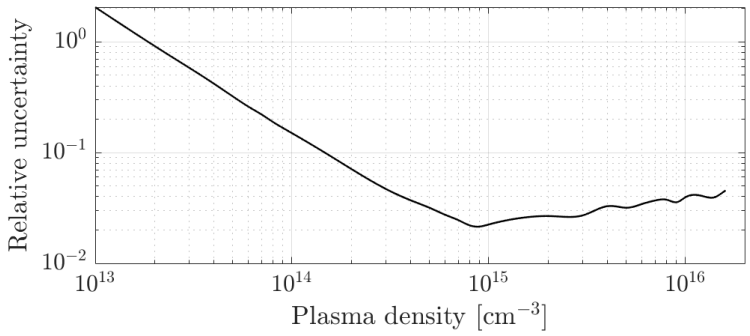

Figure 5.3.12: Relative uncertainty on the simulated measurement of $\Delta\tilde{\mu}$ presented in Figure 5.3.11. At higher plasma densities ($n_e \gtrsim 3 \times 10^{14}\,\mathrm{cm}^{-3}$) the relative uncertainty is below 5%. At plasma densities below this, the relative uncertainty grows significantly as $\Delta\tilde{\mu}$ has not saturated within the plasma length.

cantly extends beyond the first wakefield cavity and a large amount of charge is strongly defocused by returning sheath plasma electrons at the rear of the cavity. In the experiment this defocused charge has large divergence and hence is not transported through the plasma and captured by the imaging quadrupoles. In simulation, it remains within the simulation window and can still be included in the analysis of the energy spectrum and is represented by the separated data points in Figure 5.3.11. For densities below this, the simulated relation between $\Delta\tilde{\mu}$ and the plasma density bears a very similar resemblance to that predicted by linear theory via Eq. (5.2.6) shown in Figure 5.2.2, although the linear relation is suppressed at high density by its $\exp(-k_p^2\sigma_z^2/2)$ dependence. For consistency with experimental measurements, the energy spectrum of the simulated bunch at the plasma exit is binned according to the finite energy resolution of the electron spectrometer ($\delta E \sim 0.2\,\mathrm{MeV}$). The uncertainty associated with a measurement of the mean energy of the simulated bunch, and hence $\Delta\tilde{\mu}$, is calculated using the same technique as in the experiment, although does not include the uncertainty on the spectrometer charge calibration. At higher densities ($n_e \gtrsim 3 \times 10^{14}\,\mathrm{cm}^{-3}$) the relative uncertainty is on the order of a few percent as shown in Figure 5.3.12, but increases significantly at low densities where the value of $\Delta\tilde{\mu}$ has not yet saturated within the simulated plasma length. In this case, as we do not know the exact plasma length experimentally, the fractional energy perturbation is calculated for multiple timesteps towards the end of each simulation ($z = 30 - 50\,\mathrm{mm}$). These values are used to find its variation over this period and provide reasonable bounds for the value of $\Delta\tilde{\mu}$ and a corresponding increase of the associated uncertainty. This leads to a significant increase in the relative uncertainty of the simulated $\Delta\tilde{\mu}$ for plasma densities between $10^{13} - 10^{14}\,\mathrm{cm}^{-3}$.

The simulated values of the fractional energy perturbation can then be used to estimate the longitudinally-integrated effective plasma density experienced by the bunch. Using the analysis technique discussed in Section 5.3.2, the experimentally measured $\Delta\tilde{\mu}$ can be extracted from the spectrometer images presented in Figure 5.3.3 with results shown in Figure 5.3.13. As discussed previously, when the separation between the discharge initiation and the arrival of the bunch is short ($\Delta\tau \lesssim 7\,\mu\mathrm{s}$) the plasma density is sufficiently high that

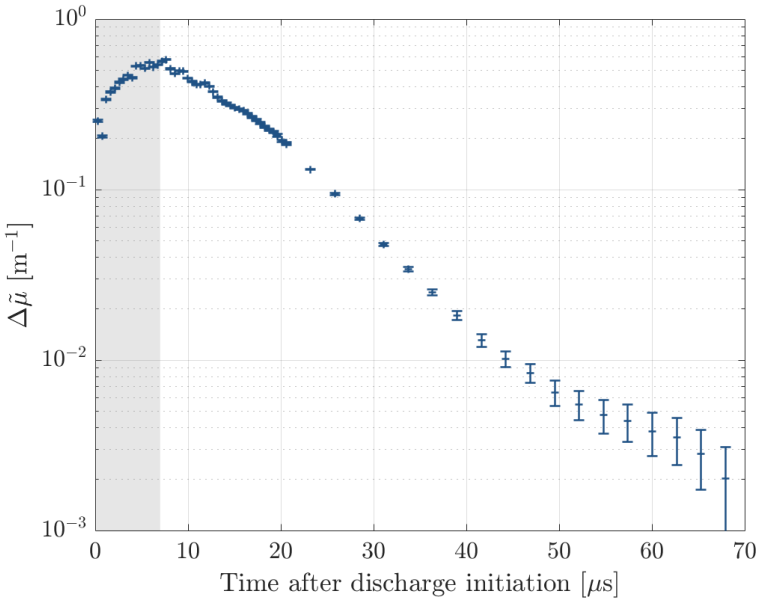

Figure 5.3.13: Measurement of the fractional energy perturbation to the probe bunch as a function of time after initiation of the discharge in a krypton plasma with a buffer volume pressure of 15.1 mbar. The data corresponds to that presented in the waterfall image in Figure 5.3.3. Error bars represent the standard error on the mean. The resolution limit for this measurement is approximately $3 \times 10^{-3}\,\mathrm{m}^{-1}$.

the wakefield cavity is significantly shorter than the bunch length and a large fraction of the bunch charge is defocused in the wakefield and lost. This period is represented by the grey shaded region at short timescales. The peak coupling between the bunch and plasma occurs approximately 8 µs after the initiation of the discharge, with good agreement demonstrated between the experimentally measured maximum $\Delta\tilde{\mu} = 0.528 \pm 0.012\,\mathrm{m}^{-1}$ and that predicted by simulation in Figure 5.3.11 to occur at a plasma density of $(1.33 \pm 0.06) \times 10^{16}\,\mathrm{cm}^{-3}$. The uncertainty on the inferred plasma density is calculated by combining the fractional uncertainties of both the experimentally-measured and simulated $\Delta\tilde{\mu}$ in quadrature. For each experimentally-measured value of the fractional energy perturbation, the plasma density that gives this value of $\Delta\tilde{\mu}$ in simulation is determined, allowing reconstruction of the evolution of the plasma density experienced by the bunch.

The resulting beam-based reconstruction of the plasma density evolution for a discharge-ionised krypton plasma of buffer volume pressure 15.1 mbar is presented in Figure 5.3.14(a). Thanks to the sensitivity of the beam-based method, plasma densities approaching $10^{13}\,\mathrm{cm}^{-3}$ can be measured, two orders of magnitude larger than would be possible with two-colour laser interferometry techniques in the same setup [195]. The physics and implications of the extended timescale evolution of such plasma sources is investigated and discussed further in Chapter 6.

Figure 5.3.14(b) shows the relative uncertainty associated with each measurement of the plasma density. At the highest densities, the relative uncertainty is consistently below 5%, indicating the accuracy of the method at typical PWFA operating densities ($n_e \gtrsim 10^{15}\,\mathrm{cm}^{-3}$).

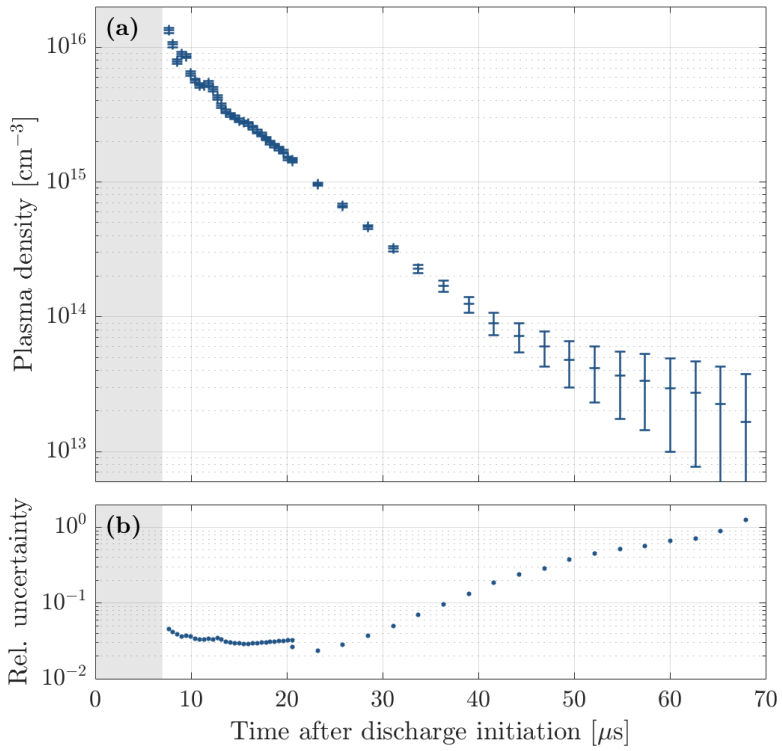

Figure 5.3.14: Evolution of the longitudinally-integrated effective plasma density for discharge-ionised krypton with a buffer volume pressure of 15.1 mbar. (a) Reconstructed plasma density as a function of time after the initiation of the discharge, measured using the beam-based technique. This allows access to extended timescales with density ranges approaching 10^{13} cm^{-3}. (b) Relative uncertainty of the reconstructed plasma density using the beam-based technique.

For plasma densities below 10^{14} cm^{-3}, the relative uncertainty increases significantly due to the associated increase in the relative uncertainty of the simulated relative energy perturbation indicated in Figure 5.3.12. As discussed before, this is because the wakefield amplitude continuously evolves over the plasma length as the bunch undergoes matching, meaning the value of $\Delta\tilde{\mu}$ does not saturate. The sensitivity of this method can of course be improved by extending the plasma such that $\Delta\tilde{\mu}$ can saturate within its length at lower plasma densities and the accumulated perturbation to the beam is more significant.

The beam-based method has shown to be an effective technique for inferring the plasma density over a wide range of densities. However, as it relies on a simple measurement of the difference in energy spectrum of the bunch before and after plasma interaction, it represents a longitudinally-integrated density measurement and it is not possible to extract information about the form of the longitudinal plasma density profile. It has previously been shown that the longitudinal density profile of such capillary discharges evolves significantly over the lifetime of the plasma as material is expelled from the ends of the capillary [195]. This motivates the use of a length-normalised metric, $\Delta\tilde{\mu}$, to interpret the plasma density. As shown in Figure 5.3.10, the value of $\Delta\tilde{\mu}$ saturates during the wakefield process after the

beam stabilises within the plasma, and thus represents a metric that is independent of the plasma length provided saturation has been reached. Therefore, as long as the plasma length is sufficient that saturation of $\Delta\tilde{\mu}$ occurs at that density, it is still possible to accurately infer the longitudinally-averaged plasma density, irrespective of the true plasma length.

The upper and lower limits of the density range over which this technique can be used are defined by the bunch and plasma lengths respectively, as well as the spectrometer energy resolution. Using shorter bunches means the constraint $\sigma_z \gtrsim \sqrt{2}\,k_p^{-1}$ is reached at higher densities (see Figure 5.2.3) while an increased plasma length leads to detectable changes in μ_E at lower densities.

5.4 CONCLUSION

A beam-based plasma characterisation technique has been developed within the context of the X-3 experiment at FLASHForward. The fast, simple technique relies on comparing measurements of the energy spectrum of an electron bunch with and without plasma interaction and is compatible with short, high peak current ($>$ kA) electron bunches typically used in high-gradient wakefield accelerators unlike previous beam-based techniques [207, 208]. When combined with detailed measurement of the six-dimensional phase space of the electron bunch, PIC simulations can be used to infer the longitudinally-averaged plasma density experienced by the bunch to an accuracy of a few percent.

This technique was developed as a method for studying the extended timescale evolution of plasmas, made possible thanks to its increased sensitivity compared to other traditional techniques such as two-colour laser interferometry. Use of this method allows easy comparison between plasmas of varying species and ionisation mechanism, enabling identification of mechanisms that can be used to adjust their lifetime as is explored further in Chapter 6. In addition to this, thanks to the capabilities of the FLASH accelerator described in Section 4.2.3, this beam-based technique can be used to perform detailed studies of the evolution of plasma that has been perturbed by a previous beam, mapping its relaxation and creating the opportunity to investigate PWFA-operation at high-repetition-rates. This is the subject of measurements presented in Chapters 8 and 9.

6

CHARACTERISATION OF UNPERTURBED PLASMA

This chapter applies the beam-based plasma characterisation method introduced in Chapter 5 to measurements of the extended timescale evolution of capillary-based discharge-ionised plasmas. It begins with a short introduction to capillary discharges, motivating their use in plasma-based accelerators. The beam-based technique is then used to measure the evolution of a discharge-ionised argon plasma and is benchmarked against a standard plasma density diagnostic, two-colour laser interferometry, showing good agreement over the common measurement period while extending the plasma density sensitivity by two orders of magnitude. Next, the beam-based technique is applied to the study of the extended timescale evolution of discharge capillary plasmas of four different species, demonstrating the dependence of their evolution on the ion mass. Finally, additional measurements performed using the beam-based technique are discussed alongside proposals for future studies to better understand methods for optimising the plasma lifetime to achieve maximal repetition-rates within plasma-based accelerators.

6.1 CAPILLARY-BASED DISCHARGE-IONISED PLASMAS

Gas-filled capillary discharges [63–65, 212, 213] are common plasma sources used within plasma accelerators due to their relative simplicity and flexibility. Thanks to the radial [214] and longitudinal [195] plasma density profiles typically produced within capillary discharges, they can be used to guide high-power laser pulses over extended distances [63–65], enable schemes for preserving the quality of charged-particle bunches during acceleration using plasma density ramps [215, 216] and symmetrically focus charged-particle bunches in plasma lenses [137, 217–219].

A schematic of a typical capillary used at FLASHForward is shown in Figure 6.1.1. A central cylindrical region is milled into sapphire slabs and then filled with gas via inlets maintained at a constant pressure to ensure gas uniformity within the region. Electrodes are placed near to the extrema of the central region and a high-voltage [$\mathcal{O}(10\,\mathrm{kV})$] discharge is initiated between the two. When the discharge is initially fired, the uniform gas density within the capillary results in a uniform current density being applied across its radius. This heats the gas via ohmic heating as electrons, accelerated by the applied potential, undergo collisions with neutral atoms and ions, causing the gas to become ionised and a uniform radial density profile to be established. After around 50 ns for typical discharge parameters [196, 197], the gas approaches full ionisation (of its first level) and the plasma temperature increases significantly as the thermal energy is no longer readily absorbed via ionisation. The solid capillary wall remains cooler, inducing a large temperature gradient across the capil-

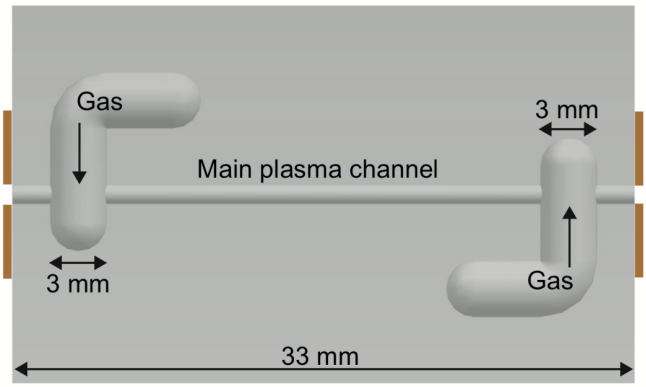

Figure 6.1.1: Schematic of a typical capillary discharge source used at FLASHForward for beam-driven plasma wakefield acceleration studies. The central cylindrical channel has a diameter of 1.5 mm and is milled from sapphire slabs. Gas is fed into the central region via the two inlets which are connected to buffer volumes filled with the relevant gas at variable pressure. Electrodes are placed at either end of the capillary (brown) and have an opening with the same diameter as the central capillary. *Image reproduced from Ref. [195].*

lary radius that results in the build-up of a pressure gradient. In order to return to uniform pressure conditions, the density is redistributed with a minimum close to the centre of the capillary where the temperature is highest and a density maximum close to the wall where the temperature is lowest. The resulting approximately parabolic radial density profile is a characteristic of capillary-based discharge-ionised plasmas [196, 197, 214]. The timescale over which the development of the non-uniform radial profile occurs during the discharge has previously been demonstrated to critically depend on the mass of the plasma ion, with heavier species exhibiting a uniform radial density profile over extended timescales, allowing the production of uniformly radially-focusing plasma lenses [219]. The presence of non-uniform radial density and temperature profiles following the culmination of the discharge has important implications for the extended timescale evolution of the plasma and is discussed in more detail in the subsequent sections.

Previous studies of capillary-based discharge-ionised plasmas have typically been focused on the discharge dynamics [194, 196, 197, 220, 221] and the generation (or suppression) of the non-uniform radial density profile [214, 219]. However, recently there has been renewed interest in their evolution over extended timescales [195] as plasma accelerators begin to be operated at high-repetition-rates. Proposals for future high-average-power plasma-based facilities [54, 57, 80, 81] typically assume operating repetition-rates of $\mathcal{O}(10\,\mathrm{kHz})$ with plasma generation, use, and decay expected to occur within a timescale compatible with these frequencies. However, detailed theoretical descriptions and experimental measurements of relevant plasma sources on this timescale have not yet been performed. The beam-based measurement technique introduced here naturally lends itself to these studies thanks to the extended plasma density range which can be accessed.

6.2 BEAM-BASED MEASUREMENT TECHNIQUE BENCH-MARKING

To validate the beam-based approach to measuring plasma evolution, direct comparison to a well-benchmarked plasma density diagnostic must be made. Such a comparison was made within the FLASHForward [57] experiment at DESY where plasma density retrieval via a two-colour laser interferometry (TCI) technique is additionally available [195]. Both measurement techniques rely on a longitudinally-integrated density-dependent response, either by the accumulated phase shift associated with two propagating laser pulses of differing frequency or by the deceleration of a drive beam by its own wakefield, and as such provide a direct comparison.

6.2.1 *Experimental setup*

Acceleration and compression of the electron beam was performed by the FLASH superconducting linear accelerator [158]. This provided energy-chirped electron beams of mean energy $\mu_0 = 1122.1 \pm 1.1$ MeV with an RMS energy spread of $\sigma_0 = 2.0 \pm 0.1$ MeV at a repetition rate of 10 Hz. The bunch charge was initially 599 ± 1 pC, compressed to 212 ± 2 fs. The longitudinal phase space of the bunch was characterised in the neighbouring FLASH1 beamline using a transverse deflecting cavity [164, 171], allowing accurate recreation of the bunch properties with approximately 10 fs temporal resolution [172]. The bunch was extracted from the FLASH beamline and passed through a dispersive section where energy collimators [159] were used to remove the tail of the bunch and shorten it such that higher densities could be accurately measured (see Figure 5.2.2), reducing the electron bunch charge to 328 ± 18 pC. Once in the FLASHForward beamline, a series of quadrupoles were used to focus the bunch on the entrance of the plasma cell to a transverse size of approximately 10 µm. The argon plasma was produced within a sapphire capillary of diameter 1.5 mm and length 33 mm by a 25 kV electrical discharge of approximately 400 ns duration, with the gas pressure of the buffer volume kept constant at 40.0 ± 0.1 mbar. The relative arrival time of the bunch with respect to the initiation of the discharge was increased until no significant perturbation to the energy spectrum of the bunch could be measured, approximately 83 µs after the discharge was fired. Approximately 50 events were recorded for each separation.

The TCI measurement and analysis was performed by a colleague in a separate dedicated plasma source measurement facility [195]. The Ti:sapphire laser pulse had a central wavelength of 800 nm with a spot size of approximately 0.5 mm in the plasma, with the central 100 µm of the spatial projection used for density retrieval. Before the plasma the laser pulse first propagated through a BBO crystal, converting approximately 10% of the pulse into the second harmonic at 400 nm, and then propagated through a 1 mm-thick glass plate to generate a temporal offset between the pulses of approximately 150 fs. After the plasma, a second BBO crystal was used to double the frequency of a fraction of the 800 nm pulse. The two perturbed 400 nm pulses were imaged on to a 10 µm slit into a spectrometer, such that the spectral interference pattern could be observed on a CCD camera. The plasma density was

then retrieved from measurements of the phase shift as outlined in Section 5.1.1. Measurements of the plasma density were performed until the resolution limit of the technique was reached, approximately 18 μs after the initiation of the discharge.

6.2.2 *Density reconstruction*

For each event, the mean energy of the decelerated region of the bunch was calculated according to Eq. (5.3.2), with a cutoff energy representing an upper bound of the energy of the bunch without plasma interaction ($E_{cut} = \mu_0 + 2\sigma_0$) applied to remove any accelerated witness particles from the analysis. The uncertainty associated with an individual measurement of μ_E includes contributions from the energy resolution of column i and the scintillator charge calibration which is typically of the order of a few percent. As discussed previously in Section 5.3.3, the metric introduced to quantify the coupling strength between the bunch and plasma was the *fractional energy perturbation per unit length*, given by Eq. (5.3.5). The length of the plasma continuously evolves as material is expelled out of the capillary into the surrounding vacuum [195], however the longitudinal plasma profile could not be measured in-situ and a fixed plasma length, equal to the length of the capillary, is used in the analysis of both the beam-based and TCI measurements for consistency. For this reason, the density inferred by these methods represents a longitudinally-integrated *effective* density measurement.

The measured $\Delta\tilde{\mu}$ for the discharge-ionised argon plasma with a buffer volume pressure of 40 mbar is shown in Figure 6.2.1. As the delay between the discharge initiation and the arrival of the bunch increases, the plasma undergoes recombination, is expelled from the capillary, and the strength of the interaction decays. Measurements began approximately 4.5 μs after the initiation of the discharge (at the peak coupling strength) and continued until the resolution limit of the measurement was reached at approximately 83 μs. The same density reconstruction technique as discussed in detail in Section 5.3.3 was used to extract the longitudinally-integrated effective plasma density. For the measurements presented in Figure 6.2.2, a series of simulations were performed in HiPACE [117] using the measured longitudinal phase space as the input beam distribution. Three-dimensional simulations were performed in background densities ranging from 10^{13}–10^{17} cm^{-3} with a fixed grid size of $\Delta\xi = \Delta x = \Delta y = 0.02\, k_p^{-1}$. At the lowest densities where this resolution was insufficient to properly resolve the bunch distribution a fixed number of grid cells was used instead. In simulation, the beam propagates through a uniform plasma of length 33 mm with identical analysis to that of the experimental measurements performed on the resulting beam distribution. The relation between the simulated $\Delta\tilde{\mu}$ and n_e (as in Figure 5.3.11) then allows determination of the effective plasma density measured in the experiment. The uncertainty on the density includes contributions from $\Delta\tilde{\mu}$ measured experimentally and found in simulation, combined in quadrature. The largest contribution to the experimental uncertainty is from the positional jitter of the beam in the dispersive section resulting in an increased variability of the beam charge following collimation (~ 5%). Simulations with bunches rep-

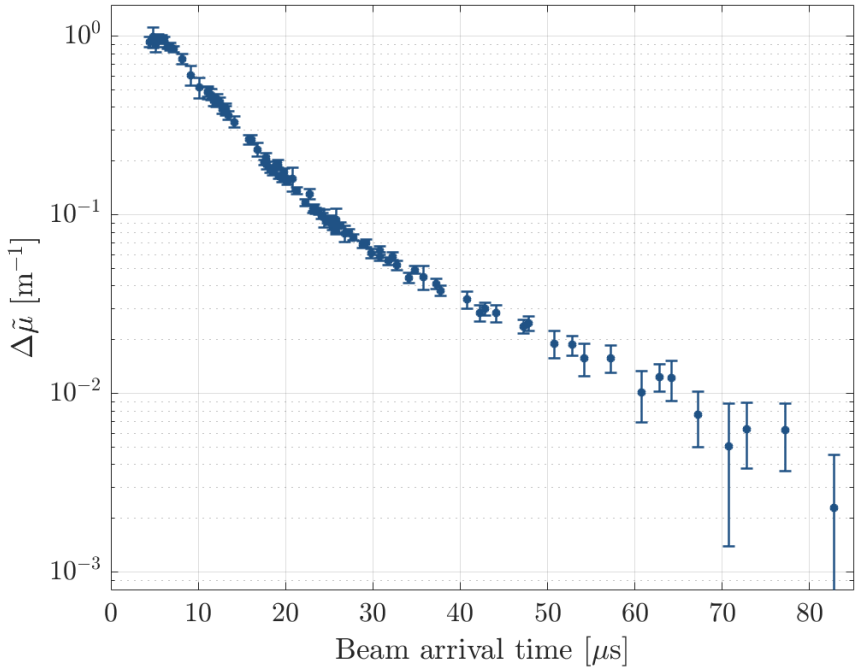

Figure 6.2.1: Measured $\Delta\tilde{\mu}$ as a function of the beam arrival time with respect to the initiation of the discharge for an argon plasma with a buffer volume pressure of 40 mbar. Error bars represent the standard error on the mean.

resenting the upper and lower bounds of the measured bunch charge were performed and additionally included within the simulation uncertainty analysis.

A comparison between the inferred plasma density measured using the TCI and beam-based techniques over the first 20 µs after the initiation of the discharge is shown in Figure 6.2.2, with the entire measured density evolution presented in Figure 6.2.3. Within the first microsecond, the plasma density grows as the discharge fires, heating the gas via Ohmic heating and ionising. Once the discharge switches off, the plasma quickly cools and recombines as ions and neutral species within the plasma undergo collisions with the capillary walls. At the highest plasma densities ($n_e \gtrsim 10^{17}\,\mathrm{cm}^{-3}$), the plasma wavelength is shorter than the probe bunch length and hence it cannot be used to accurately determine the plasma density as a large fraction of the bunch charge will extend outside the first wakefield cavity, be defocused by the transverse wakefield, and lost in transmission through the plasma. In contrast, at high densities the interferometry technique is at its most sensitive as the laser pulses acquire the largest phase shift as they propagate through the plasma. A comparison between the relative uncertainties on the inferred plasma density for the two measurement techniques is demonstrated in Figure 6.2.4, demonstrating sub-percent fractional uncertainties at densities of approximately $10^{17}\,\mathrm{cm}^{-3}$ when using the laser interferometry technique. The uncertainty on the laser interferometry measurements includes two main contributions; the shot-to-shot fluctuations induced by, for example, the nanosecond-level jitter on the thyratron that is used to trigger the discharge, and the so-called *instrument function* that defines the

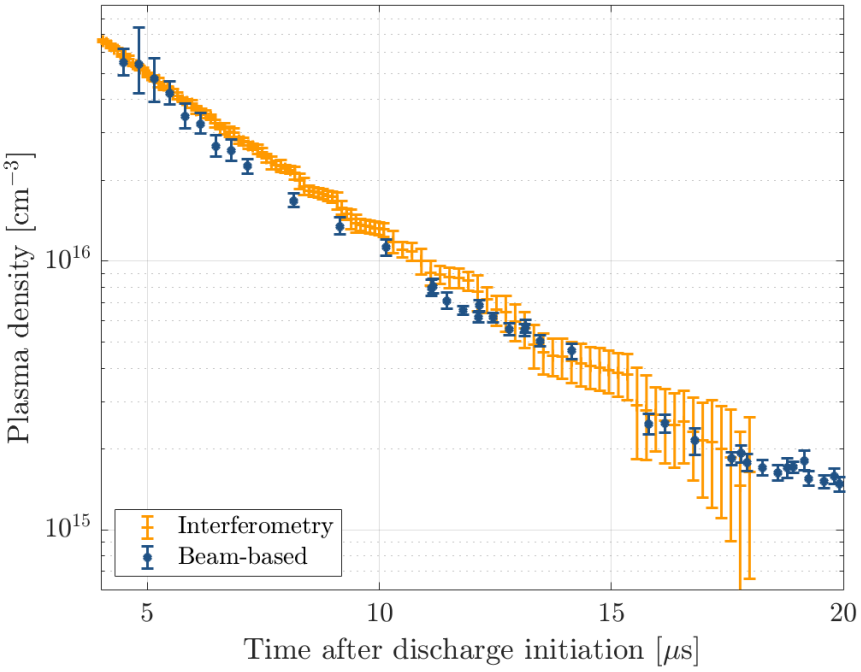

Figure 6.2.2: Comparison between TCI [195] (orange) and beam-based plasma density (blue) measurement techniques for a pure argon plasma with a buffer volume pressure of 40 mbar over a time period of 20 µs.

minimum resolvable plasma density. The value of this is calculated by consideration of the effects of the spectrometer and camera imaging setups and was calculated to be approximately $2.0 \times 10^{14}\,\mathrm{cm}^{-3}$ as outlined in Ref. [195].

Approximately 4 µs after the initiation of the discharge, the plasma has undergone sufficient recombination that its density has reduced to the point at which the beam-based technique can be used to determine the longitudinally-integrated effective plasma density. Over the next 15 µs measurements of the plasma density were able to be performed with both diagnostic techniques as the plasma continued to undergo recombination, with reasonable agreement observed between the two. Towards the end of the measurement range presented in Figure 6.2.2, the TCI measurement begins to reach its resolution limit as the accumulated phase shift of the laser pulses is too small to be accurately measured. This, in combination with the increased relative magnitude of the instrument function, leads to a significant increase in the uncertainty on an individual measurement. Figure 6.2.4 demonstrates that below $10^{16}\,\mathrm{cm}^{-3}$ ($\Delta\tau \gtrsim 10\,\mu\mathrm{s}$, Figure 6.2.3), the relative uncertainty of the TCI measurement already exceeds 10% and grows quickly as the density decreases further. In contrast, the beam-based technique reaches a relative uncertainty of 10% at densities approaching $10^{14}\,\mathrm{cm}^{-3}$ ($\Delta\tau \gtrsim 50\,\mu\mathrm{s}$, Figure 6.2.3), almost two orders of magnitude lower than the TCI method. Thanks to this increased sensitivity, the beam-based method continues to provide accurate measurements of the perturbation to the bunch energy spectrum (and hence the plasma density) over a far larger time period, considerably extending the measurement of

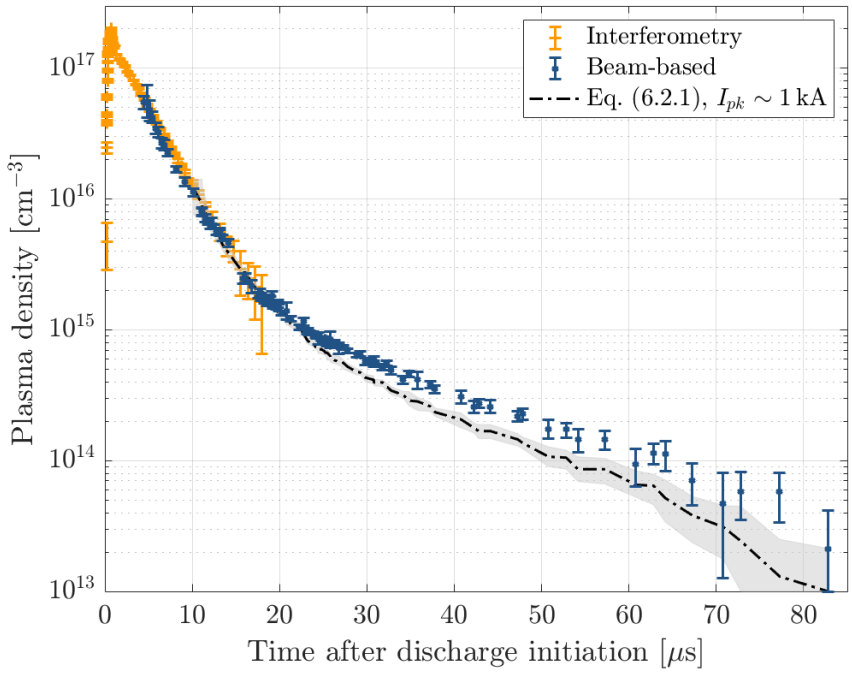

Figure 6.2.3: Comparison between TCI [195] (orange) and beam-based plasma density (blue) measurement techniques for a pure argon plasma with a buffer volume pressure of 40 mbar over a time period of 83 µs, the limit of the beam-based measurement technique. The black dashed line represents the inferred plasma density when assuming a bi-Gaussian bunch according to Eq. (6.2.1) with associated uncertainty represented by the grey shaded region. The first 20 µs are as in Figure 6.2.2.

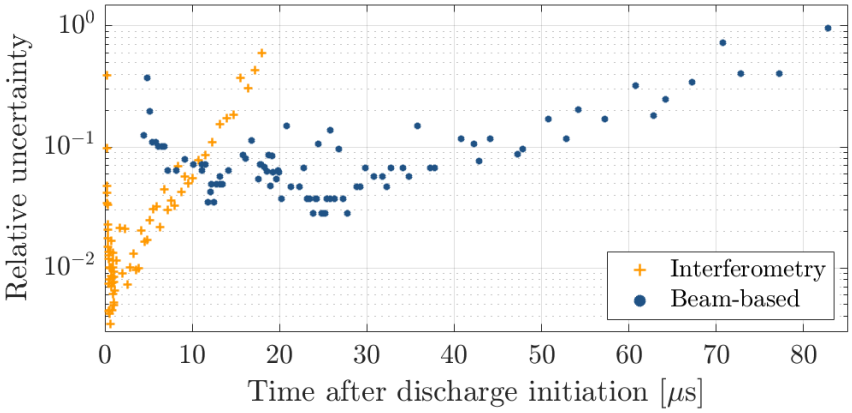

Figure 6.2.4: Comparison between the relative uncertainty of TCI [195] (orange) and beam-based plasma density (blue) measurement techniques for a pure argon plasma with a buffer volume pressure of 40 mbar over a time period of 83 µs, the limit of the beam-based measurement technique.

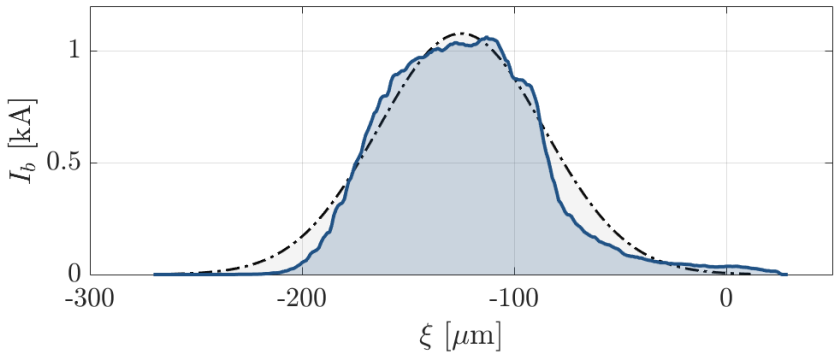

Figure 6.2.5: Comparison between reconstructed probe bunch longitudinal current profile (blue) and a bi-Gaussian approximation (black, dashed) with a bunch length of 130 fs. The bi-Gaussian approximation is used in conjunction with Eq. (6.2.1) to predict the plasma density shown in Figure 6.2.3 via measurements of $\Delta\tilde{\mu}$ presented in Figure 6.2.1.

the plasma density as shown in Figure 6.2.3. Over this extended timescale, the recombination of the plasma is observed to slow considerably and moves away from the approximately exponential decay initially observed. This is in line with similar beam-based measurements of plasma decay on extended timescales which demonstrated transition towards a power-law temporal dependency [207, 208]. This considerably extends the lifetime of the plasma when compared to assuming an exponential decay and has significant implications for the operation of high-repetition-rate systems which involve the generation of a new plasma for each acceleration event, potentially limiting them to few-kHz rates for these particular discharge and gas parameters.

It was previously shown in Chapter 5 that for a thin beam ($k_p\sigma_r \lesssim 0.1$), the average on-axis decelerating field in the vicinity of the beam, i.e. a measure of its ability to effectively couple to the plasma, can be approximated by [210]:

$$\langle E_z \rangle \approx \frac{m_e c \omega_p}{e} \frac{n_b}{n_e} k_p \sigma_z \exp(-k_p^2 \sigma_z^2/2) \ [1 - k_p \sigma_r K_1(k_p \sigma_r)] , \tag{6.2.1}$$

where $n_b = (Q/e)[(2\pi)^{3/2}\sigma_r^2\sigma_z]^{-1}$ is the beam density, Q is the bunch charge, σ_r and σ_z are the RMS transverse and longitudinal beam sizes, and K_1 is the first-order modified Bessel function. Use of this approximation allows estimation of the longitudinally-integrated plasma density without detailed measurement of the phase-space of the electron bunch and subsequent simulation of its interaction with the plasma. An example of this is also shown in Figure 6.2.3, where the plasma density is inferred from the measurements of $\Delta\tilde{\mu}$ in combination with Eq. (6.2.1) under the assumption of a bi-Gaussian bunch with longitudinal current profile shown by the black, dashed line in Figure 6.2.5. Identical bunch parameters as those measured in the experiment are used, with an RMS bunch length of 130 fs assumed to emulate the peak current of the measured non-Gaussian electron bunch ($\sim$ 1 kA). The maximum density that can be inferred using this approximation is lower than from performing simulations due to the $\exp(-k_p^2\sigma_z^2/2)$ dependence in Eq. (6.2.1) and is approximately

$10^{16}\,\mathrm{cm}^{-3}$ (i.e. measurements from a time of approximately 10 µs after the discharge onwards). Good agreement is observed at the highest densities (Figure 6.2.3, $10 \lesssim \Delta\tau \lesssim 25\,\mu\mathrm{s}$) but for $n_e < 10^{15}\,\mathrm{cm}^{-3}$, $n_b/n_e > 10$ and the linear approximation is no longer valid [209]. At this point, the wakefield amplitude is overestimated by linear theory, leading to an underestimate of the effective plasma density. However, this example demonstrates that complete characterisation of the phase space of the probe bunch is not necessary to reasonably infer plasma densities provided that $n_b/n_e < 10$.

6.2.3 *Further application of the beam-based technique*

The beam-based measurement technique has been shown to compare well to an existing plasma density diagnostic technique in TCI, while extending the applicable density range by two orders of magnitude allowing detailed study of the longer timescale evolution of the plasma. Unlike previous beam-based plasma density measurements [207, 208] that rely on measuring the plasma-density-dependent modulation of long bunches, this method uses short [$\mathcal{O}(100\,\mathrm{fs})$], high peak current ($\sim$kA) electron bunches as a probe and, as such, is compatible with existing high-gradient PWFA drive bunches, requiring only an energy spectrometer as a diagnostic. With detailed characterisation of the probe bunch this metric can be additionally transformed into an effective plasma density measurement as demonstrated in the preceding sections.

The identification of the *fractional energy perturbation per unit length*, $\Delta\tilde{\mu}$, as a measure of the coupling efficiency between a given bunch and the plasma through which it propagates allows fast, comparative measurements of different plasmas using the same probe bunch via interpretation of this metric alone. This forms the basis for the measurements that follow in this chapter; a constant bunch is used to measure the evolution of different plasmas on extended timescales and identify key experimental variables that can be used to minimise the lifetime of a plasma. Minimisation of this is important as the operation of plasma-based accelerators at repetition-rates that rival state-of-the-art conventional particle accelerators [157, 158, 199, 222] is considered, i.e. tens of kHz to MHz, as the plasma must have relaxed such that no perturbation from a previous acceleration event can affect any subsequent event.

6.3 VARYING PLASMA SPECIES

The plasma species was varied in an attempt to understand the dependence of the lifetime of the plasma on the mass of the ion. Four different plasma species were investigated across two datasets (listed in order of increasing mass):

- Helium; $m_{He} = 4.00$,
- Neon; $m_{Ne} = 20.18$,
- Argon; $m_{Ar} = 39.95$,
- Krypton; $m_{Kr} = 83.80$,

where the masses are given in atomic mass units; $1\,\text{amu} = 1.66 \times 10^{-27}\,\text{kg}$. The first dataset compared the extended timescale evolution of discharge-ionised neon and krypton, while the second compared helium and argon. The same probe bunch was used across each dataset where comparisons are made. The bunch did however differ between the two datasets as the measurements were performed in different experimental periods due to time constraints and the availability of the different gas species. A brief description of the probe bunch used in each measurement is given at the start of each dataset for reference.

6.3.1 *Neon vs. Krypton*

For this dataset, an electron beam of charge $982 \pm 2\,\text{pC}$ and energy distribution $\mu_0 = 1030.1 \pm 1.1\,\text{MeV}$ and $\sigma_0 = 1.5 \pm 0.1\,\text{MeV}$ was compressed to an RMS bunch length of $\sigma_\tau = 373 \pm 4\,\text{fs}$ and used to probe discharge-ionised plasmas of differing species and pressure: neon at a buffer volume pressure of 30.1 mbar and krypton at 15.1 mbar in capillaries of length 50 mm. The longitudinal current profile of the probe bunch is shown in the inset of Figure 6.3.1. Both measurements are performed using the same probe beam and so comparison between the measured $\Delta\tilde{\mu}$ is sufficient to understand the relative evolution of the plasma states.

Immediately after the culmination of the discharge the plasma density of krypton exceeds that of neon, indicated by the lower values of $\Delta\tilde{\mu}$ in the first μs in the lead up to the peak interaction. This is expected due to the difference in their ionisation energies: the first two ionisation energies for krypton are $[14.0, 24.4]\,\text{eV}$ compared to $[21.6, 41.0]\,\text{eV}$ for neon [223]. The lower ionisation energies of krypton means ions readily reach higher ionisation states under similar discharge conditions. In addition to this, the lower initial gas pressure leads to increased plasma electron temperatures (T_e), further increasing the plasma density relative to neon. This can be quantified using a simple model [224] that solves for the evolution of the average plasma density during the discharge according to the Saha equation [225] using representative experimental parameters as input. This model predicts an average ionisation fraction of $Z = 2.67$ for krypton ($T_{e,max} = 3.4\,\text{eV}$) compared to $Z = 1.04$ for neon

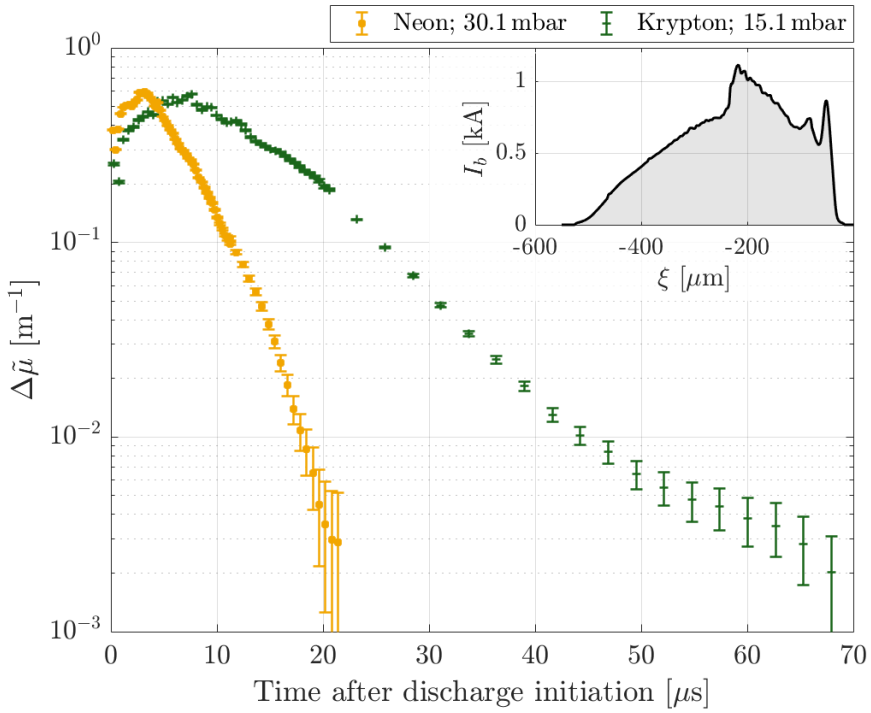

Figure 6.3.1: Comparison between beam-based measurements of the evolution of neon (orange squares) and krypton (green crosses) discharge-ionised plasma. Error bars represent the standard error on the mean. Inset: measured longitudinal current profile of the probe bunch.

($T_{e,max} = 2.7\,\text{eV}$) immediately after the discharge when using experimentally-measured discharge current traces as input. This large discrepancy in the average ionisation fraction more than accounts for the difference in gas pressure within the buffer volume.

At these early times, when the plasma density is highest, the electron bunch extends over more than one plasma period. Regions of the bunch experience defocusing forces and up to 60 % of the bunch charge is lost during transmission through the plasma, measured downstream of the plasma exit. The peak interaction between the beam and the plasma is the same for both datasets within experimental uncertainties and corresponds to a plasma density of approximately $1.4 \times 10^{16}\,\text{cm}^{-3}$ according to simulation. The minimum densities measured using the beam-based method approach $10^{13}\,\text{cm}^{-3}$ in both neon and krypton for beam arrival times of 21.4 µs and 67.9 µs respectively, again demonstrating the large density range accessible to this technique.

The long timescale behaviour following peak interaction demonstrates a significantly faster decay of neon compared to krypton. The main contributing reason for this is the difference in ion mass (m_i) between the two species. Following culmination of the discharge, the density evolution is expected to be dominated by longitudinal outflow [195] and radial diffusion [196, 197] of the plasma towards the wall, driven by the pressure gradients that are created during the discharge as discussed in Section 6.1. Mass flow is controlled by the ion diffusion (D_i) and viscosity coefficients which both scale as $m_i^{-1/2}$ [226, 227]. In the plasma bulk, quasi-

neutrality is maintained as the space charge fields induced by electrons with large thermal velocities act to increase (decrease) the diffusion rate of ions (electrons). Both species diffuse at the ambipolar rate [228]:

$$D_a \approx D_i \left(1 + \frac{T_e}{T_i}\right), \tag{6.3.1}$$

tending towards $2D_i$ over microsecond timescales as the ion and electron temperatures equalise at a temperature T via inter-species collisions. The ion diffusion rate is given by [128]

$$D_i = \frac{T}{m_i(\nu_{ii} + \nu_{in})}, \tag{6.3.2}$$

where ν_{ab} represent collision rates between species a and b. The dominant heat loss mechanism of the plasma is expected to be via collisions with the capillary wall, described by Fourier's law:

$$q = -\kappa \frac{\partial T}{\partial r} \tag{6.3.3}$$

where q represents the heat flux to the wall, κ is the thermal conductivity and $\partial T/\partial r$ represents the value of the radial temperature gradient at the wall. The thermal conductivity can be described by a sum of contributions from each species [196]:

$$\begin{aligned} \kappa &= \kappa_e + \kappa_i + \kappa_n, \\ &= \frac{n_e T}{m_e(\nu_{ei} + \nu_{en})} + \frac{n_i T}{m_i(\nu_{ii} + \nu_{in})} + \frac{n_n T}{m_i(\nu_{ni} + \nu_{nn})}. \end{aligned} \tag{6.3.4}$$

In the vicinity of the capillary wall, the temperature is sufficiently low that the plasma is weakly-ionised and the dominant heat exchange comes via collisions between heavy plasma species (ions and neutrals) and the wall. The relevant collision rates are

$$\nu_{ii} = \frac{4}{3}\left(\frac{\pi}{m_i}\right)^{1/2} \frac{e^4 n_i \ln \lambda_{ii}}{(4\pi\varepsilon_0)^2 (k_B T)^{3/2}}, \tag{6.3.5}$$

$$\begin{aligned} \nu_{in} &= n_n \left(\frac{k_B T}{m_i}\right)^{1/2} 4\pi a^2, \\ &= \nu_{nn}, \end{aligned} \tag{6.3.6}$$

$$\nu_{ni} = n_i \left(\frac{k_B T}{m_i}\right)^{1/2} 4\pi a^2, \tag{6.3.7}$$

where a is the electron-neutral scattering length and $\ln \lambda_{ii}$ is the ion-ion Coulomb logarithm, an approximation of a diverging collision integral that is typically of order 10 and is defined as

$$\ln \lambda_{ii} = \ln \left[\frac{3}{4\sqrt{2\pi}} \frac{(4\pi\varepsilon_0)^{3/2} (k_B T)^{3/2}}{e^3 n_e^{1/2}}\right]. \tag{6.3.8}$$

These transport coefficients all indicate an inverse square root dependence on the ion mass, i.e. $m_i^{-1/2}$. The ratio of the decay rate of the interaction following the peak for the two species is found to be 2.37 ± 0.14, calculated from exponential fits to the measurements presented in Figure 6.3.1 over timescales of $5 - 15\,\mu s$ in neon and $10 - 45\,\mu s$ in krypton. This is close to the square root of their mass ratio, $\sqrt{m_{Kr}/m_{Ne}} = 2.04$, consistent with mass flow and the resultant heat transfer to the capillary wall being one of the dominant mechanisms that defines the long timescale evolution of discharge-ionised plasma. Further differences between the two species are likely due to recombination mechanisms that are strongly dependent on the electron temperature (e.g. the three-body recombination rate $\propto T_e^{-9/2}$ [128]). Hence, recombination will be accelerated for neon as a result of its lower expected T_e immediately after the discharge and its higher ionisation energies.

6.3.2 *Helium vs. Argon*

For the second dataset, an electron beam of charge $455 \pm 3\,pC$ and energy distribution $\mu_0 = 1115.5 \pm 1.2\,MeV$ and $\sigma_0 = 1.7 \pm 0.1\,MeV$ was compressed to an RMS bunch length of $\sigma_\tau = 200 \pm 2\,fs$, giving a peak current of approximately 1 kA. This bunch was used to probe discharge-ionised helium at a buffer volume pressure of 25.3 mbar and argon at 8.6 mbar in capillaries of length 50 mm. These buffer volume pressures were chosen such that an equal plasma density was reached approximately 1 μs after the initiation of the discharge, identified by comparing values of the fractional energy perturbation per unit length, with the peak $\Delta\tilde{\mu} = 0.63 \pm 0.01$ occurring at a plasma density of approximately $n_e = 2.4 \times 10^{16}\,cm^{-3}$ according to simulation. Note, the use of a shorter probe bunch for the measurements presented in this dataset result in the peak coupling strength occurring at a higher plasma density than those presented in Section 6.3.1 as expected. The large difference in required pressure to reach the same plasma density in helium and argon is due to the difference in breakdown voltage required to initiate the plasma generation for the two species as described by Paschen's Law [213, 229]. A 25 kV discharge pulse of duration approximately 400 ns was used to generate the plasma in both cases.

The resulting measurement of the fractional energy perturbation per unit length as function of time after the initiation of the discharge is shown in Figure 6.3.2. These measurements present the familiar initially exponential dependence of $\Delta\tilde{\mu}(t)$ as also measured in Figures 6.2.1 and 6.3.1. As in Section 6.3.1, an exponential is fitted to this initial temporal region to compare the decay rate of the two plasmas; for helium this represents the measurements of $\Delta\tilde{\mu}$ between $2 - 5\,\mu s$, while for argon this represents the time period between $2 - 20\,\mu s$. The ratio of the decay rates was calculated to be 3.23 ± 0.13, consistent with the mass ratio of the two species: $\sqrt{m_{Ar}/m_{He}} = 3.16$. This represents a second independent measurement of two different plasma species using a different probe bunch and demonstrates a similar scaling of the extended timescale evolution of the discharge-ionised plasma, thus providing confidence in both the simple $m_i^{-1/2}$ scaling and also the beam-based measurement technique itself.

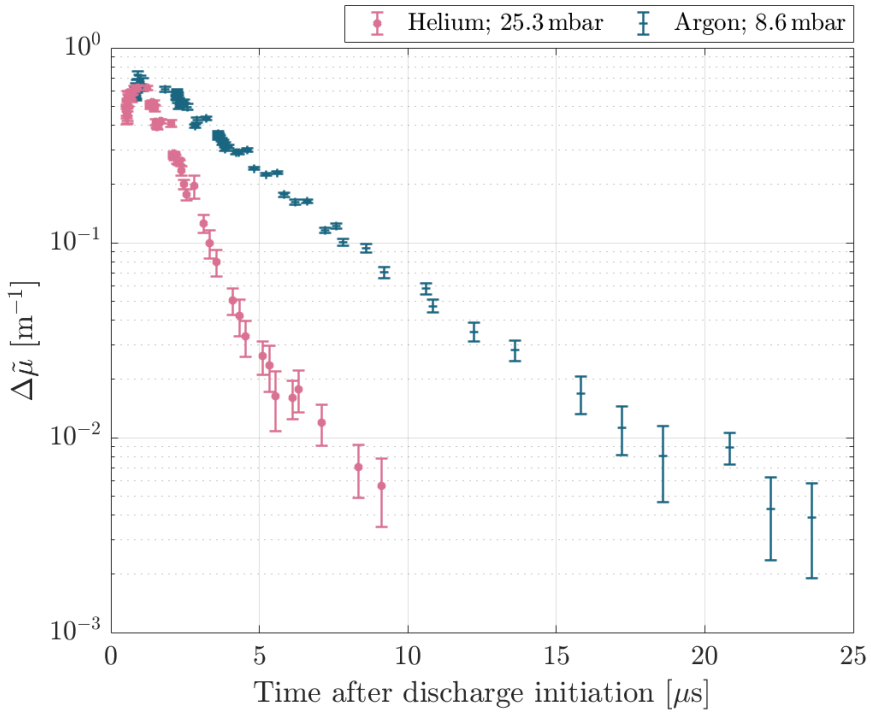

Figure 6.3.2: Comparison between beam-based measurements of the evolution of helium (pink circles) and argon (blue crosses) discharge-ionised plasma. Error bars represent the standard error on the mean. Inset: measured longitudinal current profile of the probe bunch.

The identification of a fractional energy perturbation metric, $\Delta\tilde{\mu}$, that is easy to interpret and requires only a single calculation that can be performed during experimentation has hence revealed important underlying physics, indicating its usefulness as a technique. These measurements have demonstrated that through careful choice of the plasma species, the lifetime of discharge-ionised capillaries can be optimised for the operation of plasma-based accelerators at high-repetition-rates.

6.3.3 *A comment on the plasma "lifetime"*

While the beam-based measurement technique has been shown to offer an extension in plasma density sensitivity by two orders of magnitude over the TCI technique, it still reaches its limit at plasma densities of approximately $10^{13}\,\mathrm{cm}^{-3}$ within the described experimental setup. We therefore cannot define the "lifetime" of the plasma using this technique as it will continue to exist and decay beyond this point. However, consider typical values measured in these datasets; the minimum measurable $\Delta\tilde{\mu}$ approached $2 \times 10^{-3}\,\mathrm{m}^{-1}$, which for a 1 GeV bunch in a plasma of length 50 mm represents a mean energy change of 0.1 MeV – a 0.01% change, highlighting the sensitivity of this technique.

Herein lies an intrinsic advantage of this beam-based technique and the motivation for the use of an energy perturbation metric. When considering high-repetition-rate operation of plasma-based accelerators, we are not solely interested in the exact plasma density at some

specific time after the discharge, but rather whether the plasma density remains sufficiently high over a sufficiently long time period to significantly affect a subsequent acceleration event. By using the PWFA drive bunch itself to probe the plasma state and identify the time at which sub-per-mille perturbations to its energy spectrum are no longer observed, the time period over which this is the case, and hence the maximum possible repetition-rate, is automatically determined.

This of course assumes that the chosen high-repetition-rate operation scheme relies on the generation of a new plasma for every acceleration event and that it is possible to do this at such a frequency. As discussed in Section 1.3.1, additional measurements performed within the X-3 experiment at FLASHForward (not discussed within this thesis) indicate that gas refill of the capillary following firing of the discharge occurs over millisecond timescales. This would significantly reduce achievable repetition-rates in schemes that rely on regeneration of the plasma in this way although further optimisation of the capillary design and plasma conditions can relax this restriction. Alternative high-repetition-rate operation modes based on using the same plasma more than once during its lifetime also must be considered and provide the basis for further measurements presented in Chapters 8 and 9.

6.4 EXTENSIONS OF THE BEAM-BASED MEASUREMENT TECHNIQUE

6.4.1 *Additional dependencies*

It is clear from the measurements discussed in the previous sections that the extended timescale evolution of a discharge-ionised plasma is a function of not only the plasma species but a myriad of other variables; for instance, the pressure of the gas within the capillary. Consider the measurement of the fractional energy perturbation to the probe bunch presented in Figures 6.2.1 and 6.3.2 where discharge-ionised argon plasmas of different pressures, 40 mbar and 8.6 mbar respectively, are measured. Note, these measurements were performed with different probe bunches and hence direct comparison of their $\Delta\tilde{\mu}$ is not valid, although the peak currents of the bunches was similar in both cases.

Even without direct comparison of the measured $\Delta\tilde{\mu}$ values, it is clear that the timescale over which the plasma decays significantly differs in the two cases. The resolution limit of the measurement is reached at approximately 83 µs for a buffer volume pressure of 40 mbar compared to approximately 25 µs at 8.6 mbar. This can be explained by consideration of the ion diffusion rate given in Eq. (6.3.2). The ion-ion collision rate within this term is proportional to the ion density and hence the ion diffusion rate is inversely proportional, $D_i \propto n_i^{-1} m_i^{-1/2}$. As the gas pressure is reduced, the ion density after the culmination of the discharge correspondingly decreases and the diffusion rate increases. This results in an increased flow of the plasma bulk both radially towards the capillary walls, where the plasma quickly cools and recombines, and longitudinally out the ends of the capillary. This increases the rate of loss of the plasma via both recombination and expulsion. In addition to this, the discharge parameters used in these measurements were the same; 25 kV and 400 ns duration. A reduction in the gas pressure would therefore be expected to result in an increase in the temperature of the plasma at the culmination of the discharge. This has two effects; (i) further increasing the ion diffusion rate via Eq. (6.3.2), and (ii) increasing the ionisation fraction of atoms within the plasma bulk. Higher ionisation states typically recombine more quickly and at high densities ($n_e > 10^{14}\,\mathrm{cm}^{-3}$), three-body recombination is expected to be the dominant recombination mechanism [230]. Its rate has a strong electron temperature dependence, $\alpha_3 \propto T_e^{-9/2}$ [128], suggesting that an increased recombination rate would be expected for lower gas pressures thanks to the increased rate of temperature loss following the discharge.

It is therefore expected that reducing the gas pressure within the capillary offers an avenue for reducing the plasma lifetime, provided that the required on-axis plasma density can still be reached following the discharge. To confirm this dependency, additional measurements of the extended evolution of discharge-ionised plasmas under varying gas pressures using the same probe beam should be performed. When doing so, the experimental setup should be modified to allow measurement of the gas pressure within the capillary itself rather than within the buffer volume. This is necessary as the buffer volume is connected to the capillary via long (metre-scale) gas pipes. Directly measuring the pressure within the capillary removes the possibility that a non-linear relation between the pressure of the buffer volume

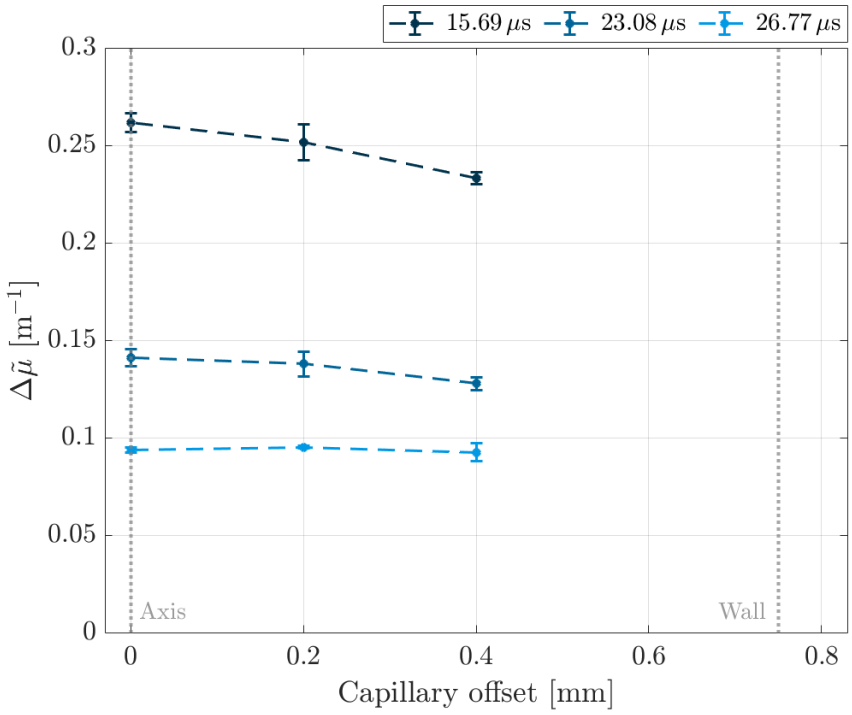

Figure 6.4.1: Beam-based measurement of the evolution of the radial plasma density profile of an argon plasma at a buffer volume pressure of 40 mbar. Different colour data points represent three times after discharge initiation: 15.69 µs (dark blue), 23.08 µs (blue) and 26.77 µs (light blue). Experimental conditions were identical to those presented in Figure 6.2.1. Error bars represent the standard error on the mean. The two grey, dashed vertical lines represent the central axis of the capillary (left) and the capillary wall (right).

and the pressure within the capillary itself could obfuscate interpretation of the resulting measurements.

6.4.2 *Radial evolution*

A feature of the beam-based measurement technique is that it samples only the plasma density local to the propagation axis of the probe bunch — even at the lowest plasma densities, the radial extent of the wakefield is expected to be $\mathcal{O}(100\,\mu\text{m})$. This therefore means the radial evolution of the plasma density within the capillary can be probed. Experimentally, this is achieved by changing the position of the capillary with respect to the propagation axis of the probe bunch, here defined as the *capillary offset*. At FLASHForward this is possible as the base-plate on which the capillary is mounted is connected to a hexapod [162], allowing consistent translation of the base-plate in all three planes. This measurement concept was tested during the measurement of the extended timescale evolution of argon at a buffer volume pressure of 40 mbar presented in Figure 6.2.1, with results shown in Figure 6.4.1. Three different capillary offsets were used in 0.2 mm steps at three different timings with respect to the initiation of the discharge. Measurements were also attempted at a capillary offset of

0.6 mm but significant bunch charge was lost due to the proximity to the wall (at 0.75 mm). These initial measurements indicate a reduction in the coupling strength between the probe bunch and plasma as the bunch is moved closer to the capillary wall, indicating a decrease in plasma density with increasing radius at extended timescales after the culmination of the discharge. Note, this radial density dependence directly contrasts to that produced *during* the discharge where a density minimum appears on-axis with an increasing parabolic radial density dependence towards the wall [196, 197, 224]. As discussed previously, during the discharge the presence of the capillary wall cools plasma in its vicinity and induces a characteristic radial temperature (and, hence, density) profile while the plasma near the axis continues to be heated by the discharge current. However, after the discharge terminates such radial profiles are no longer supported and the plasma near the wall recombines, inverting the radial density profile.

Measurements such as these could be repeated with the temporal granularity typically used in on-axis measurements, such as those presented in Figure 6.2.1, with finer capillary offset steps to reconstruct the complete radial evolution of the plasma density profile and compliment the on-axis measurements. Recent advances with control of the probe bunch focusing and trajectory optimisation [173] would also allow measurements to be performed closer to the capillary wall. Reconstruction of the radial evolution of the plasma density in this way would provide useful comparison to 1D MHD simulations [194, 196] that are frequently used to predict the behaviour of discharge-ionised plasmas and optimise experimental parameters.

6.4.3 *Alternative ionisation mechanisms*

Another approach for minimising the lifetime of plasmas could be to use a different ionisation mechanism, thereby modifying the initial plasma state and potentially its extended evolution. One such alternative is ionisation by a short, high-intensity laser pulse, which can produce spatially-confined plasmas of tunable temperature (e.g. via variation of the laser pulse polarisation ellipticity [231]), allowing additional control over those typically produced via discharge-ionisation. While lasers of suitable performance to match recent progress in high-repetition-rate discharge systems are not yet readily available, there has been recent progress in their development with high peak power ($>$ 100 TW), high average power ($>$ 1 kW) laser systems expected within the next five to ten years [81]. Laser-driven wakefield acceleration experiments have already demonstrated operation at kHz repetition-rates using few-mJ laser pulses [191–193], indicating the potential to develop laser systems that could be suitable for the ionisation of plasma sources at tens of kHz repetition-rates in the future.

Within the FLASHForward experiment, a 25 TW Ti:sapphire laser system is available for studies of controlled internal injection within plasma [76], and hence provides an alternative ionisation method that can be compared to discharge-ionisation. Exploratory measurements of the evolution of a laser-ionised plasma were performed in an argon-filled capillary of length 195 mm at a buffer volume pressure of 10.4 mbar using an electron bunch with a

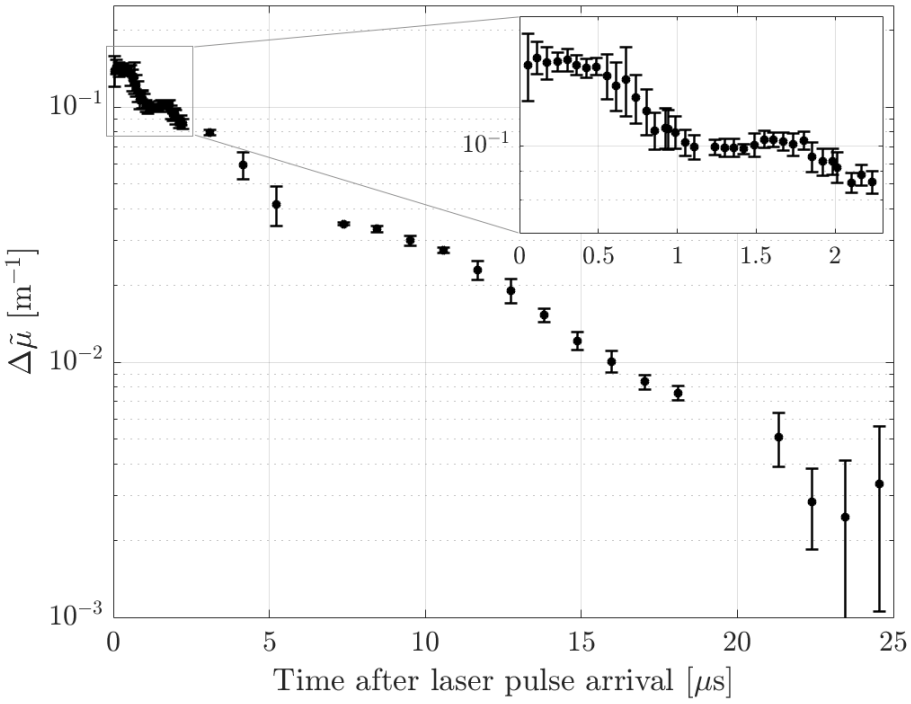

Figure 6.4.2: Beam-based measurement of the evolution of laser-ionised argon plasma at a buffer volume pressure of 10.4 mbar. Inset: short timescale evolution shows the appearance of two density plateaus. Error bars represent the standard error on the mean.

lower peak current ($I_{pk} \sim 500$ A) and are presented in Figure 6.4.2. The measurements were performed in two separate scans; the first focused on short timescales (up to 2.5 µs after the arrival of the laser pulse) to study the initial evolution of the plasma column and the second over an extended timescale (25 µs) until no significant perturbation to the probe bunch energy spectrum was observed.

The reduced stability of the laser-ionised plasma as compared to the discharge-ionised can be immediately noticed in these measurements. The error bars associated with the measurements presented in Figure 6.4.2 fluctuate significantly in size between data points, a result of the increased shot-to-shot variation of the energy spectrum of the drive bunch following its interaction with the plasma. This is likely due to fluctuations in the overlap between the ionised plasma column and the drive bunch caused by pointing jitters of the laser pulse. The focal length of the ionising laser pulse is of order 18 m due to spatial constraints from the extended accelerator beam-line and hence small pointing jitters upstream caused by vibrations can lead to significant shifts in the position of the focal spot. The overlap between the thin plasma column and the electron bunch that is focused to transverse sizes of $\mathcal{O}(10\,\mu\text{m})$ can therefore vary significantly shot-to-shot. This effect is likely amplified by ionisation defocusing effects that are discussed in more detail later.

The peak coupling strength between the probe bunch and the plasma is lower than was typically observed in measurements of discharge-ionised plasmas for two reasons: (i) the

use of a probe bunch with a lower peak current and (ii) the effect of ionisation defocusing reducing the plasma length. Ionisation defocusing occurs due to the non-uniform transverse density profile produced by the leading edge of the laser pulse as it ionises and focuses in the resultant plasma [232, 233]. The refractive index, n, of a plasma is given by

$$n(r,z,t) = \left[1 - \left(\frac{\omega_p^2}{\omega_0^2}\right)\right]^{1/2}, \tag{6.4.1}$$

where $\omega_0 = 2\pi c/\lambda_0$ is the angular frequency of the laser pulse. Hence regions of high plasma density correspond to lower refractive indices. The intensity of the laser pulse varies across its leading edge, producing a non-uniform plasma electron density profile with a density maximum along its propagation axis where the intensity is highest. A density maximum on-axis corresponds to a refractive index minimum, radially increasing towards the wings of the pulse. The resultant plasma density profile therefore acts as a lens to the trailing part of the laser pulse and causes it to refract, reducing its intensity and modifying its structure. This effect is particularly significant as the pulse comes into focus as only the central region has sufficient intensity to ionise the gas, creating large local plasma density gradients within the pulse. This is a particular problem at FLASHForward as the focal length of the ionising laser pulse is of order 18 m, resulting in metre-scale Rayleigh lengths with a focal spot size of approximately 350 µm [76]. Ionisation defocusing effects can be dominant if the defocusing length, l_d, is shorter than the Rayleigh length [233],

$$l_d = \frac{\lambda_0}{2}\frac{n_c}{n_e}, \tag{6.4.2}$$

where λ_0 is the central wavelength of the laser pulse and $n_c = (\varepsilon_0 m_e/e^2)\omega_0^2$ is the critical plasma density for a wave of angular frequency ω_0. Typical FLASHForward operating parameters assume a plasma density of order $10^{16}\,\mathrm{cm}^{-3}$ while the central wavelength of the ionising laser pulse is approximately 800 nm, giving a defocusing length of $l_d \approx 7$ cm, considerably shorter than the Rayleigh length of the pulse and indicating that ionisation defocusing is expected to be significant for this parameter range. Therefore, we do not expect the plasma to be ionised along the entire length of the 195 mm capillary and the plasma density is likely to vary radially across the width of the plasma column.

Despite the expected non-uniformity of the initial plasma density profile complicating interpretation of the measurement, effects of the plasma evolution can be extracted from the measurements presented in Figure 6.4.2 and comparisons to similar discharge-ionised measurements can be made. Over the first 0.5 µs, the coupling between the probe bunch and the plasma, and hence the plasma density, appears to remain constant. This is likely due to the initial spatial distribution of the plasma. When ionising using a laser pulse, only a thin filament of plasma is ionised around the laser focus. Conversely, during a discharge the plasma is generated throughout the entirety of the capillary. Therefore, the laser-ionised plasma can initially rapidly expand into the surrounding volume unlike the discharge-ionised plasma. The expansion of the plasma column begins from its outer edge, etching away towards the

centre of the plasma column, and a finite amount of time passes before the on-axis density (that the probe bunch experiences) reduces, here approximately 0.5 µs. A similar effect is observed in Ref. [207] where the plasma density of the central laser-ionised column remains approximately constant over the first microsecond.

A second, small density plateau is observed approximately 1 µs after the laser propagates through the gas. This possibly corresponds to the time at which the expanding plasma column reaches the capillary walls, marking the transition from the initial rapid expansion of the column to a diffusion- and recombination-dominated decay as is observed in discharge-ionised plasmas. It is possible to estimate the electron temperature of the laser-ionised plasma from determination of this expansion time. The front of the plasma filament expands radially into the capillary and the rate at which this occurs has previously been shown to be described well by the Sedov-Taylor solution for the expansion of ideal gases [231, 234, 235]. This describes the idealised expansion of a cylindrical region of gas of temperature T and following the derivation in Ref. [231], the evolution of the position of the shock front can be described by:

$$r_s(t) = (\gamma+1)^{1/2}\left(\frac{ZE_k}{m_i}\right)^{1/4}[r_0(t+\tau_0)]^{1/2} \tag{6.4.3}$$

where $\gamma = 5/3$ is the adiabatic index, Z is the average ionisation fraction (here assumed to be unity), $E_k = (3/2)k_B T$ is the mean electron energy, r_0 is the initial plasma radius and τ_0 is a constant that can be calculated from the initial conditions:

$$\tau_0 = \frac{1}{\gamma+1}\left(\frac{m_i}{ZE_k}\right)^{1/2} r_0. \tag{6.4.4}$$

Using an initial plasma radius of $r_{plas} \approx 150\,\mu m$ and an expansion time of $\Delta t \approx (1 \pm 0.2)\,\mu s$, the mean temperature can be estimated to be $0.35 < T < 0.80\,eV$. While this estimate encapsulates a large range of temperatures, it is important to note the almost order of magnitude difference with respect to discharge-ionised plasmas where multi-eV electron temperatures have been inferred from comparisons between laser interferometry and spectroscopic measurements of the plasma density immediately following the culmination of the discharge [195]. Despite this difference in initial plasma temperature and the fast initial drop in plasma density as the column expands as demonstrated in Figure 6.4.2, on extended timescales it appears as though the plasma exists over a similar timescale to discharge-ionised plasmas. This suggests that the plasma lifetime is not strongly dependent on the ionisation mechanism for these parameters as collisional and diffusive effects equilibriate the plasma states over microsecond timescales and heat loss to the capillary walls and subsequent recombination becomes the dominant evolution mechanism. It could be expected that this would not be the case for a laser-ionised plasma that is not confined within a capillary, as the initial expansion of the plasma column would continue uninhibited. It is unfortunately not possible to measure this within the FLASHForward experiment due to stringent vacuum

requirements within the accelerator limiting the maximum achievable gas pressure of the chamber outside of the capillary.

To be able to make definitive comparisons between discharge- and laser-ionised plasmas, these measurements should be repeated in conditions where ionisation defocusing effects can be minimised. For example, reducing the capillary length to 50 mm as was used in the discharge-ionised measurement presented previously. Additionally, an alternative gas species with a reduced ionisation threshold, for example hydrogen, could be used to minimise the laser intensity required to ionise the gas and hence achieve ionisation further in the wings of the laser spot, reducing the ionisation defocusing effect. Measurements of discharge- and laser-ionised plasmas over a range of gas pressures should be performed using the same probe bunch to allow easy direct comparison.

The effect of the plasma electron temperature on the extended evolution of the plasma could additionally be studied by controlling the ellipticity of the ionising laser pulse by passing the pulse through a quarter-wave plate. It has previously been shown that the average electron temperature can be varied by changing the ellipticity of the pulse as electrons released during ionisation acquire kinetic energy in the plane transverse to the axis of propagation of the laser pulse [231, 236]. Given that the mass flow rate of the plasma depends on its temperature, it would be expected that its density evolution can be altered in this way, particularly during the initial expansion of the plasma column within the first microsecond.

6.5 CONCLUSION

The beam-based measurement technique introduced in Chapter 5 has been applied to the study of the extended evolution of plasma. The use of a simple metric that corresponds to the strength of the interaction between a bunch and the plasma [Eq. (5.3.5)] has been shown to provide an effective probe of the plasma state and allow comparison between different plasmas. When combined with a measurement of the six-dimensional phase space of the bunch, possible at FLASHFORWARD thanks to the recent advances in available diagnostic techniques, this metric can be transformed into an integrated plasma density measurement that compares well with TCI techniques, extending the plasma density sensitivity by two orders of magnitude within this setup. Even without exact reconstruction of the probe bunch, the plasma density can be inferred from linear theory when assuming a bi-Gaussian distribution provided that $n_b/n_e < 10$.

Use of the beam-based measurement technique indicated that the extended timescale evolution of the discharge-ionised plasma is dominated by its mass flow, longitudinally and transversely, exhibiting the expected inverse square root dependency on the mass of the ion. Lowering the capillary pressure was similarly observed to reduce the plasma lifetime, likely due to the induction of larger relative pressure gradients and temperatures driving increased mass flow rates. A scheme for measuring the radial dependency of the plasma density within the capillary was defined with initial measurements presented, demonstrating a decreasing plasma density as the bunch approached the capillary walls as the reduced temperature in this region leads to a reduction in the ionisation fraction. Application of this measurement scheme would enable reconstruction of the radial evolution of a discharge-ionised plasma over an extended timescale, providing an invaluable comparison to one dimensional MHD simulations that are readily used. Initial measurements of the extended evolution of a laser-ionised plasma are presented, demonstrating initial expansion of the ionised plasma column and a transition to the capillary-wall-driven decay observed in discharge-ionised plasma following this. Combining this measurement technique with ionising laser pulses of varying ellipticity would enable the study of the dependency of the evolution of the plasma column on the electron temperature.

In summary, the beam-based technique has been shown to provide a fast, easy-to-interpret measurement of the evolution of a plasma that can be applied to a wide range of systems. The measurements presented here indicate that the lifetime of the discharge-ionised plasma can be minimised by using light gas species, for instance hydrogen, at low pressures. Given that measurements of a discharge-ionised helium plasma demonstrated no significant effect on the bunch energy distribution within 10 µs of its generation (Figure 6.3.2), it is reasonable to expect a similarly short lifetime in hydrogen and hence a maximum possible plasma regeneration rate approaching 100 kHz, on-par with current and proposed conventional accelerator facilities that operate at rates of tens of thousands of bunches per second [58, 198, 199]. However, it is important to note that these facilities achieve this by operating in a pulsed mode where a large number of bunches ($\sim$ 1000) are accelerated at much higher frequencies

(~MHz) but contained within RF macro-pulses that operate at lower frequencies (~ 10 Hz), as discussed in Chapter 4, rather than operating continuously at $\mathcal{O}(10\,\text{kHz})$ frequencies.

This therefore raises the question of whether a similar "burst" operation mode would be possible in plasma-based accelerators such that they are compatible with state-of-the-art conventional accelerators that are likely to provide the drive and witness beams, particularly in the case of beam-driven plasma wakefield accelerators. In this operation mode a single plasma would be created, used multiple times to accelerate particles within its lifetime and then decay, after which a subsequent plasma would be created and the process repeated *ad infinitum*. The crucial difference though is the presence of the plasma as an acceleration medium as opposed to the ultra-high vacuum environments of conventional accelerators. A plasma does not represent a fixed, non-evolving medium as demonstrated by the measurements within this chapter, and hence its evolution during and after the acceleration process must be considered and possibly utilised. A large amount of energy is transferred to the plasma from the bunches via the wakefield and this will drive subsequent evolution of the plasma, akin to the heating provided by the discharge or laser during ionisation. To date, only two publications [83, 84] have experimentally observed further evolution of a plasma over sub-nanosecond timescales driven by the energy transferred into the plasma from the bunches. These do not represent detailed studies of this evolution or its dependencies, and measurements over longer microsecond timescales that could be crucial to high-repetition-rate operation of plasma accelerators have yet to be performed.

Coupled with the unique capabilities of the FLASH accelerator systems described in Section 4.2.3, the beam-based measurement technique presented here lends itself to such detailed measurements thanks to its inherent sensitivity to the plasma conditions. Detailed measurements of the evolution of a plasma following the wakefield acceleration process over timescales from few-nanoseconds to tens of microseconds using the beam-based characterisation technique are presented in Chapters 8 and 9.

Part IV

PERTURBED PLASMA EVOLUTION

Part IV describes detailed measurements of plasma wakefield evolution both during and after its perturbation by a drive beam and spans the entire temporal range described in the introduction in Chapter 1.

In Chapter 7, measurements of the growth and decay of the wakefield amplitude along a self-modulated proton bunch are discussed. These measurements were performed at AWAKE and demonstrate the growth in amplitude of resonantly combined wakefields and the onset and effects of plasma electron trajectory-crossing and phase mixing that lead to decay of the wakefield amplitude on sub-nanosecond timescales.

Chapter 8 presents detailed measurements of the motion of ions within the plasma driven by the wakefield excitation process on nanosecond timescales. Chapter 9 extends these measurements up to tens of microseconds, investigating the re-ionisation of the plasma induced by the energy transfer from the drive bunch. These represent the first detailed measurements of their kind, were performed at FLASHForward, and are made possible using the beam-based technique discussed in Part III.

7

SUB-NANOSECOND RESONANT WAKEFIELD EXCITATION

This chapter details experiments performed at AWAKE to attempt to measure the evolution of the longitudinal wakefield driven by a self-modulated proton bunch along both the length of the plasma and the length of the bunch itself. This was achieved by varying the temporal delay between the arrival of the ionising laser pulse and the witness electron bunch, the *laser–electron delay*, $\Delta\tau$. The experimental results are split into two sections according to the timescales over which they were performed:

- The first corresponds to a dataset taken in September 2018 and focuses on "short" timescales where the delay between the ionising pulse and the witness bunch was limited to $\Delta\tau \leqslant 365\,\text{ps}$. This measurement aimed to demonstrate resonant excitation of the longitudinal wakefield amplitude.
- The second dataset was taken in November 2018 and focused on "extended" timescale evolution of the wakefield, $\Delta\tau \leqslant 800\,\text{ps}$. An additional experimental variable was included within these measurements: modifying the position of the ionising laser pulse within the proton bunch to change the number of protons that underwent interaction with the plasma and hence modify the maximal wakefield amplitude. This is referred to as the *seeding position*.

The measurements and analysis presented in this chapter were recently published,

> [237] J. Chappell *et al.* (AWAKE Coll.), Phys. Rev. Accel. Beams **24**, 011301 (2021).

7.1 MEASUREMENT OVERVIEW

An overview of the entire AWAKE Run 1 setup can be found in Section 3.4.1, however specifics of the experimental setup for the measurements discussed in this chapter follow and are summarised in Table 3 for reference. The witness electron bunch had a charge of approximately 300 pC and its trajectory was matched to that of the proton bunch with a small vertical offset. This allowed the electron bunch to be injected into the wakefield at a small vertical angle of approximately 0.5 mrad with its focus in the plane of the entrance of the vapour source. The plasma density was kept constant at $n_0 = 2 \times 10^{14}\,\text{cm}^{-3}$ with a constant longitudinal density profile maintained throughout the measurements; while larger witness energies could be observed at higher plasma densities (see Section 3.4.3 or Ref. [144]), achieving consistent witness capture was more challenging at higher densities and hence lower operating densities were preferred for systematic studies.

Plasma		
Density [cm^{-3}]	2×10^{14}	
Radius [mm]	2	
Ionising laser pulse		
Energy [mJ]	56.1 ± 1.5	
Pulse length [fs]	120	
Electron bunch		
Energy [MeV]	19	
Norm. emittance [mm·mrad]	10	
Transverse size at focus [μm]	500	
Charge [pC]	300 ± 8	
RMS bunch length [ps]	8	
Focal position [m]	0	
Trajectory at focus	On-axis; −0.5 mrad (vert.)	
Proton bunch	**Section 7.3**	**Section 7.4**
Energy [GeV]	400	
Norm. emittance [mm·mrad]	3.5	
Transverse size at focus [μm]	200	
Bunch population [10^{11}]	3.13 ± 0.16	2.83 ± 0.14
RMS bunch length [ps]	257 ± 3	255 ± 3
Seeding position [ps]	+100	$[-200, 0, +200]$

Table 3: Summary of experimental parameters. These also represent the values used in the simulations outlined in Section 7.2. The seeding positions are defined relative to the centre of the proton bunch with positive values representing seeding ahead of the centre.

The electron spectrometer dipole current was set to 100 A allowing electron energies ranging from 84 MeV to 2 GeV to be measured as shown in Figure 7.1.1. The spectrometer optical system was designed to achieve a resolution of 1 mm in the dispersive plane, corresponding to a theoretical energy resolution of $< 2\%$ except at the high energy region of the scintillating screen ($\zeta < 100$ mm). During the experiment, a protective fire window was added to the optical line of the spectrometer for safety reasons. Unfortunately, warping of this window caused a reduction of the resolution of the spectrometer system in the dispersive axis to approximately 2 mm, measured in situ using a resolution target in combination with a fixed light source that mimicked the emission profile of the scintillating screen [146, 149].

Typically, the timing of the arrival of the ionising laser pulse was chosen to be close to the centre of the proton bunch. This meant the plasma was generated in the region around the

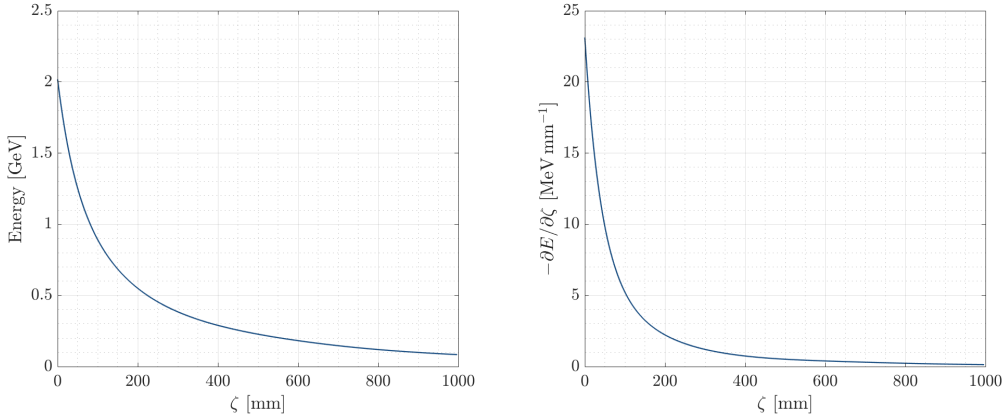

Figure 7.1.1: AWAKE electron spectrometer position-energy relation (left) and theoretical resolution (right) for a dipole current of 100 A. These relations are calculated via beam tracking simulations performed using BDSIM using measured dipole field maps as input. $\zeta = 0$ mm corresponds to the end of the scintillating screen closest to the proton beam trajectory.

maximal unmodulated proton bunch density and ensured that the wakefields driven by the unmodulated bunch were sufficiently large in amplitude to seed self-modulation rather than allowing non-axisymmetric modes such as hosing to develop from small-scale plasma non-uniformities [61, 120]. Three different seeding positions were used in the study of extended timescale wakefield evolution, corresponding to ionising laser pulse arrival times of 200 ps ahead of, behind, or in the centre of the proton bunch as shown in Figure 7.1.2. The ionising laser pulse energies for these seeding positions were measured using an energy meter to be 55.1 ± 0.8 mJ (200 ps ahead), 58.7 ± 1.2 mJ (centre) and 54.5 ± 0.9 mJ (200 ps behind), hence the plasma radius is expected to be consistent across all three measurements [207].

Moving the seeding position within the proton bunch modifies the amplitude of the seed wakefields driven by the unmodulated proton bunch due to the change in bunch density that initially interacts with the plasma. When seeding 200 ps ahead of, or behind, the centre of the proton bunch, the seed wakefield amplitude is reduced compared to the central seeding position as shown in Figure 7.1.3. However, in all three cases the seed wakefield is sufficiently large in amplitude to seed self-modulation and results in reproducible wakefield phase, essential for consistent acceleration of high-quality witness bunches [61]. In addition to this, changes in the seeding position modifies the number of protons within the bunch that can undergo interaction with the plasma and hence the wakefield amplitude driven by the modulated bunch. By moving the seeding position ahead of the centre of the proton bunch more protons undergo interaction with the plasma and the maximal wakefield amplitude is increased (and vice versa) as demonstrated in the right panel of Figure 7.1.2.

Analysis of spectrometer images

For an event to be included in this analysis, a minimum witness charge of 50 fC needed to be detected on the spectrometer screen in order to eliminate events where sufficient witness

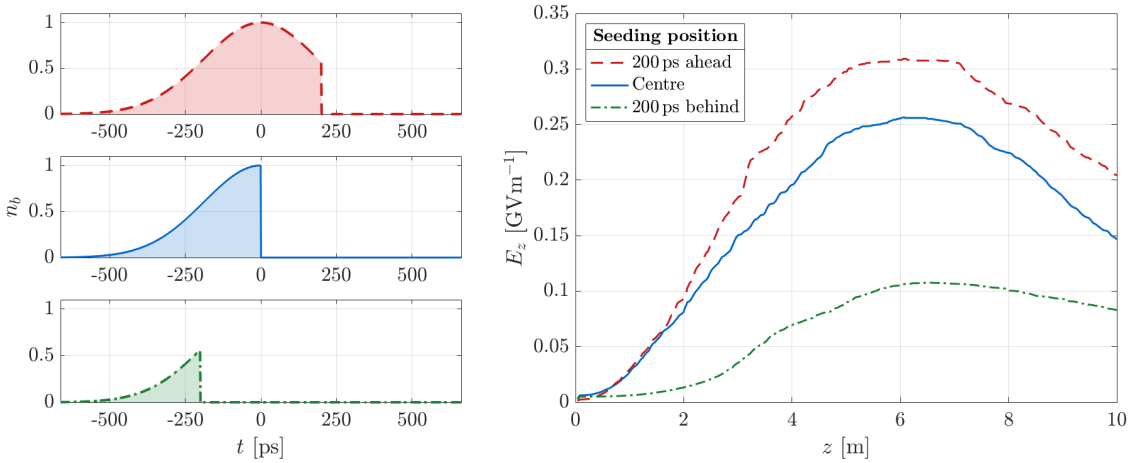

Figure 7.1.2: Left: Illustration of the normalised beam density, n_b, corresponding to the seeding positions used in this study. Shaded regions indicate the region of the beam that propagates through, and hence interacts with, the plasma. Left, upper: Seeding 200 ps ahead of the centre. Left, middle: Seeding in the centre of the proton bunch. Left, lower: Seeding 200 ps behind the centre. Right: Simulated evolution of the maximum longitudinal field over the entire plasma length for the three different seeding positions for a plasma density of $n_e = 2 \times 10^{14}$ cm^{-3}. The proton bunch parameters are the same as those used for studying extended timescale wakefield evolution described in Section 7.4.

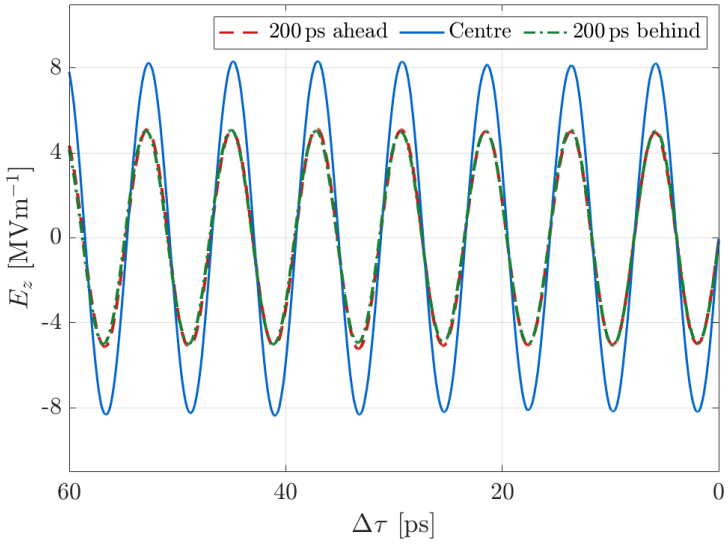

Figure 7.1.3: Seed longitudinal wakefield amplitude for the three seeding positions used in this study.

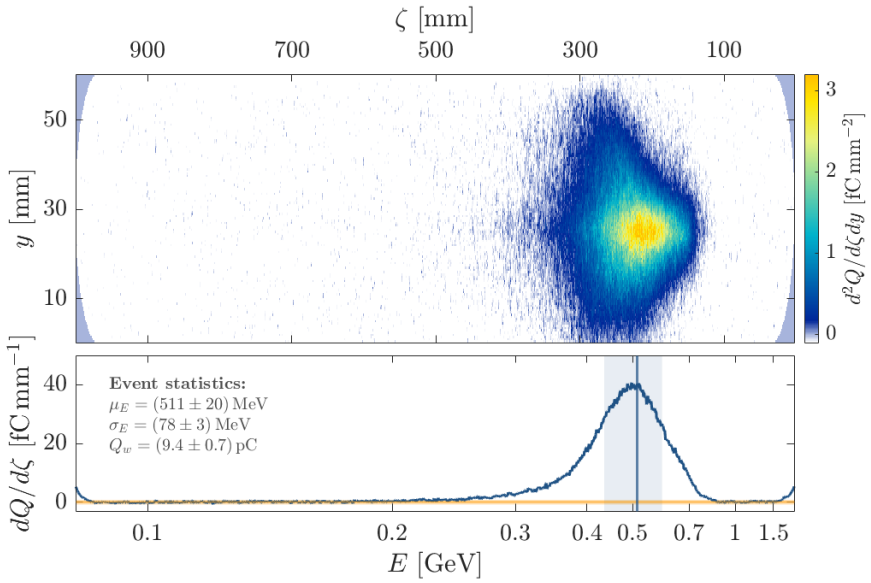

Figure 7.1.4: Upper: Example electron spectrometer image. The light blue regions near the edges represent the mask used to isolate the scintillating screen in spectrometer images. Lower: projection in the dispersive plane, ζ, with the corresponding non-linear energy axis. The blue vertical line corresponds to the measured mean energy of the witness bunch, μ_E, and the blue shaded region represents the energy spread, σ_E. The orange shaded region represents the $3\sigma_{bkg}$ signal cutoff level.

capture was not achieved. An example is shown in Figure 7.1.4. For each event, background subtraction and geometric corrections [149] were applied to the image of the scintillator screen and the pixel count was integrated over the vertical (non-dispersive) plane. The region with signal that exceeded the expected background by $3\sigma_{bkg}$ was identified. The mean energy of the captured witness electrons was calculated according to

$$\mu_E = \frac{(\Sigma_i E_i \cdot dQ_i)}{(\Sigma_i dQ_i)}, \tag{7.1.1}$$

where i corresponds to the index of the column of the image in the signal region, E_i is the energy associated with column i in the plane of the scintillator screen and dQ_i is the integrated charge measured in column i. The energy spread was similarly calculated according to

$$\sigma_E = \left(\frac{(\Sigma_i E_i^2 \cdot dQ_i)}{(\Sigma_i dQ_i)} - \mu_E^2 \right)^{1/2}. \tag{7.1.2}$$

The integrated charge of a column was calculated by summing the total CCD pixel counts of the background-subtracted image and applying the calibrated scintillator charge response value, $(4.22 \pm 0.33) \times 10^5\,\mathrm{pC}^{-1}$.

7.2 SIMULATION OVERVIEW

The experimental measurements are compared to 2D cylindrical, quasi-static simulations performed using LCODE [118]. These simulations solve for the plasma response in the co-moving frame, defined by $\xi = z - ct$, and use the experimental proton and electron bunch parameters as input. The simulation domain spans $0 \leqslant \xi \leqslant 800\, k_p^{-1}$, $0 \leqslant r \leqslant 25\, k_p^{-1}$ with a resolution of $\Delta\xi = \Delta r = 0.02\, k_p^{-1}$ where $k_p = \omega_p/c$. A fixed timestep of $\Delta t = 100\, \omega_p^{-1}$ is used to update both the plasma state and proton beam; for witness electrons, an energy-dependent reduced timestep is used to fully resolve their betatron oscillations. Approximately 30 radius-weighted macro-particles per cell per species are used to model the response of plasma electrons and ions with 2.7×10^6 equal-weighted beam macro-particles used to model the proton bunch. The proton bunch is initialised with Gaussian longitudinal and radial distributions. Modelling of the varying seeding positions used in this study is achieved by using a step function in the proton bunch density at the relevant seeding position within the bunch. This negates the need to model the interaction between the ionising laser pulse and the Rb vapour and hence saves computational resources.

Simulation domain size

Use of a wide simulation domain that extends far beyond the boundary of the plasma ($r_p = 5.3\, k_p^{-1} = 2\,\text{mm}$) is necessitated by the requirement to continue to track plasma electrons that are expelled from the plasma [238]. In LCODE, when an electron reaches the radial boundary at $r = r_{max}$ it is returned to a location nearby with much reduced momentum [119]. The use of a radial boundary that is too small will lead to expelled electrons hitting the simulation boundary and will remove any effect from electrons that can later return to the plasma which, as discussed in more detail later, can be vital to understanding the extended timescale evolution of the wakefield. An example of this effect is shown in Figure 7.2.1 where the on-axis longitudinal wakefield at a propagation distance of $z = 3\,\text{m}$ is plotted for three different values of r_{max}. The simulations are identical in all other aspects. While no difference between the simulations is observed on short timescales ($\Delta\tau \leqslant 200\,\text{ps}$), the longitudinal wakefield amplitude differs significantly on extended timescales with considerable damping of the wakefield amplitude in the case of wider simulation windows. The mechanism by which this occurs is discussed in more detail in the following sections, but this example illustrates the importance of using a sufficiently wide simulation window when comparing to experimental results where the boundary of the vapour source itself (20 mm) extends far beyond the expected plasma boundary (2 mm). Tests indicated that the radial domain needed to exceed $r_{max} = 20\, k_p^{-1}$ for the extended timescale evolution of the wakefield to converge for the experimental bunch and plasma parameters. For this reason, all simulations were performed with $r_{max} = 25\, k_p^{-1} = 9.4\,\text{mm}$.

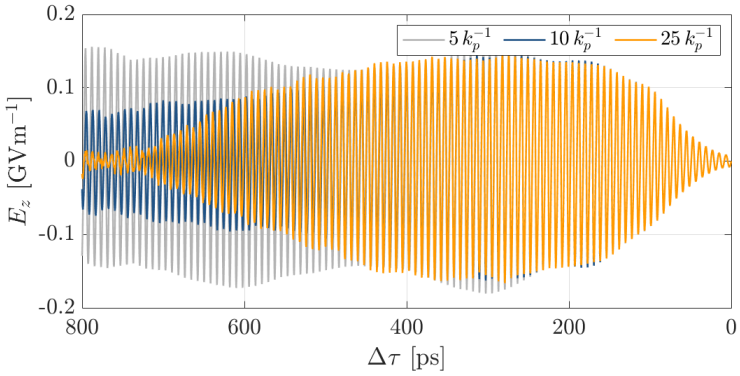

Figure 7.2.1: Comparison between simulated longitudinal wakefield amplitude in the co-moving frame for different radial domain sizes; $r_{max} = 5\,k_p^{-1}$ (grey), $r_{max} = 10\,k_p^{-1}$ (blue) and $r_{max} = 25\,k_p^{-1}$ (orange) for a propagation distance of $z = 3\,m$.

2D vs. 3D simulations

The process of external electron injection is inherently a complicated 3D problem [239] and therefore perfect agreement between 2D simulations and experiment is not expected. In 2D cylindrical simulations, the witness electrons are represented by a ring of charge being injected into the wakefield toward the axis as opposed to an electron bunch with pointing jitters as in the experiment. As such, the total captured charge in 2D is expected to exceed that observed experimentally. $\mathcal{O}$(pC) captured witness charges were typically measured during the experiment and hence no significant modification of the wakefield amplitude experienced by the witness electrons via beam loading is expected. For these two reasons, test witness electrons that can experience the wakefield but do not alter its amplitude or drive their own wake are used in simulation in an attempt to recreate measured experimental trends.

It is additionally expected that 2D axisymmetric simulations underestimate turbulent effects due to the imposed symmetry and hence likely provide an overestimate of the wakefield amplitude at large laser–electron delays following phase mixing. However, at present, three-dimensional quasi-static simulations of the extended plasma region required for comparison to experimental observations are prohibitively computationally expensive. In spite of this, axisymmetric simulations still provide useful comparison and insight into the underlying physical mechanisms.

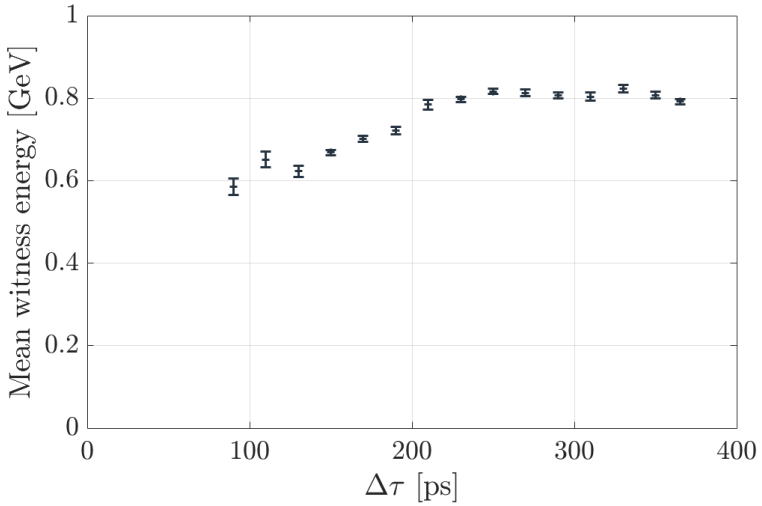

Figure 7.3.1: Experimental results showing mean witness energy for $\Delta\tau$ up to 365 ps when seeding 100 ps ahead of the centre of the proton bunch. The growth in mean witness energy with increasing $\Delta\tau$ is consistent with resonant excitation of the longitudinal wakefield. Error bars represent the standard error on the mean.

7.3 RESONANT WAKEFIELD EXCITATION

As discussed in Chapter 3, as the long proton bunch becomes modulated it is expected that the wakefield is resonantly excited by the combined action of micro-bunches driving wakefields in the plasma. It was demonstrated in Ref. [154] by measurement of the proton bunch "halo" that the transverse wakefield amplitude grows far above the predicted seed wakefield amplitude driven by a long unmodulated proton bunch, indicating further excitation of the wakefield. However, this measurement was temporally-integrated and hence did not measure the dependence of the wakefield amplitude along the length of the proton bunch and demonstrate *resonant* excitation.

For the first measurement scheme, the delay between the arrival time of the ionising laser pulse and the witness bunch, $\Delta\tau$, was varied between $50 - 365$ ps in steps of 15 ps, probing the integrated longitudinal wakefield after every two micro-bunches ($\lambda_p/c \approx 8$ ps) by measuring the energy spectrum of the resulting witness bunch. The proton bunch population was $(3.13 \pm 0.16) \times 10^{11}$ with an RMS length of (7.70 ± 0.08) cm $= (257 \pm 3)$ ps and the seeding position corresponded to 100 ps ahead of the centre of the proton bunch. The results are shown in Figure 7.3.1.

For the smallest delays, $\Delta\tau < 150$ ps, achieving consistent witness capture was challenging, resulting in low detected charge with larger mean energy variability and energy spread. This was likely due to the globally defocusing nature of the seed wakefields at short delays with respect to the leading edge of the proton bunch as discussed in Section 3.3.1. Witness capture was not achieved for $\Delta\tau < 90$ ps, while approximately 30 % of injection attempts achieved sufficient charge capture to contribute to the measurements for each value of the laser–electron delay for $90 \leqslant \Delta\tau < 150$ ps. However, for $\Delta\tau \geqslant 150$ ps, successful witness

capture was observed in more than 75 % of events per delay. Nevertheless, the trend for increasing witness mean energy with increasing laser–electron delay was observed on short timescales ($\Delta\tau < 250\,\mathrm{ps}$) and directly demonstrated resonant excitation of the longitudinal wakefield. At around $\Delta\tau = 250\,\mathrm{ps} \approx \sigma_z$, the witness energy was observed to saturate as predicted by previous simulations [120] and is discussed in more detail in the following sections.

7.4 EXTENDED TIMESCALE WAKEFIELD EVOLUTION

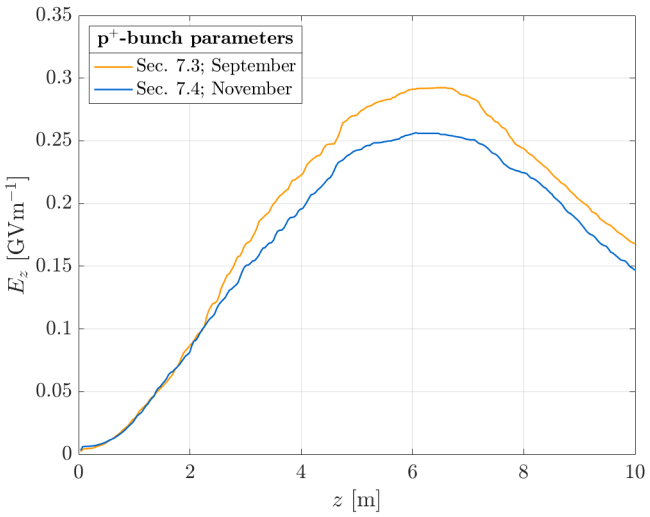

Figure 7.4.1: Comparison between the maximal longitudinal wakefield amplitude evolution for the proton bunch parameters in Sections 7.3 (orange) and 7.4 (blue). The proton bunch parameters from Section 7.4 correspond to a central seeding position.

The second measurement campaign was dedicated to studying extended timescale evolution of the wakefield beyond that presented in Section 7.3. The measured proton bunch parameters differed between the two measurements and three different seeding positions were used as discussed previously. For the following measurements the proton bunch population was measured to be $(2.83 \pm 0.14) \times 10^{11}$ while its RMS bunch length was $(7.65 \pm 0.08)\,\mathrm{cm}$. In addition to the change in seeding position between the measurements, the reduction in bunch population while maintaining approximately the same bunch length reduces the peak current of the beam by approximately 10%. This results in a reduction in the wakefield amplitude over the plasma length as shown in Figure 7.4.1 and, as a consequence, a difference in the self-modulation saturation length. Therefore, witness electrons captured in the experimental conditions discussed in Section 7.3 will experience a larger amplitude wakefield and undergo consistent acceleration over a longer distance than those for subsequent measurements, resulting in larger witness energies. For this reason, direct comparison between the mean captured witness energy of the two measurements cannot be made. However, for the measurements studying the effect of various seeding positions on extended timescales (up to

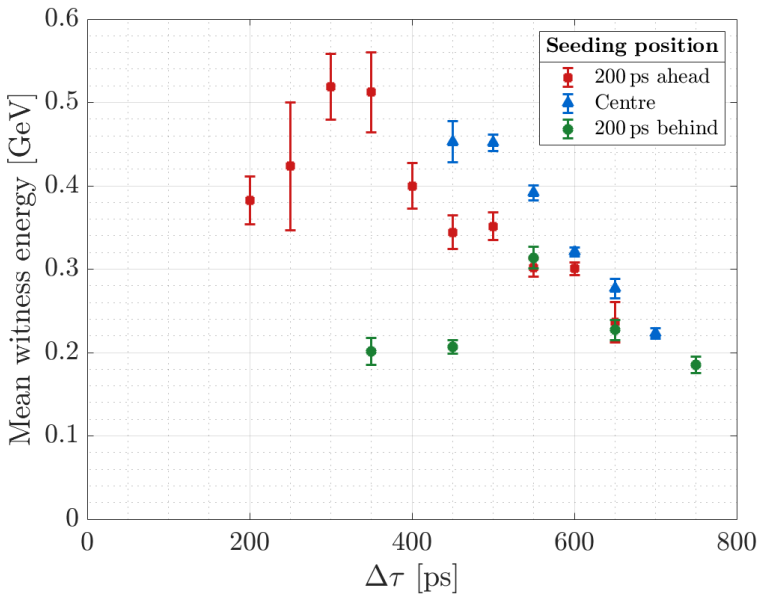

Figure 7.4.2: Experimental results showing the mean witness energy with increasing laser–electron delays for the three measured seeding positions. Error bars represent the standard error on the mean.

800 ps) detailed in this section, the experimental parameters were kept as consistent as possible to allow direct comparison. In this case, for the smallest delays ($\Delta\tau < 200$ ps) capture of electrons was not realised while fewer than 30 % of attempted events per delay achieved sufficient witness capture to contribute to the measurements of laser–electron delays between $200 - 400$ ps. For larger delays, witness electron bunches of sufficient charge were measured in approximately 70 % of events per laser–electron delay.

7.4.1 *Longitudinal wakefield amplitude*

For laser–electron delays larger than 500 ps, the effects of further plasma wakefield evolution were observed experimentally and are demonstrated in Figure 7.4.2. On these extended timescales a decreasing witness energy was measured with increasing laser–electron delay, indicative of a decreasing longitudinal wakefield amplitude along the self-modulated proton bunch.

Plasma electron trajectory crossing and expulsion

Simulations indicate the decrease in longitudinal wakefield amplitude is as a result of the onset of plasma electron trajectory crossing within the wakefield. An example of this is shown in Figure 7.4.3 which exhibits the simulated plasma electron density (upper) and a selection of corresponding plasma electron trajectories in the co-moving frame (lower) within 200 ps of the leading edge of the proton bunch after a propagation distance of 3 m. As the wakefield is

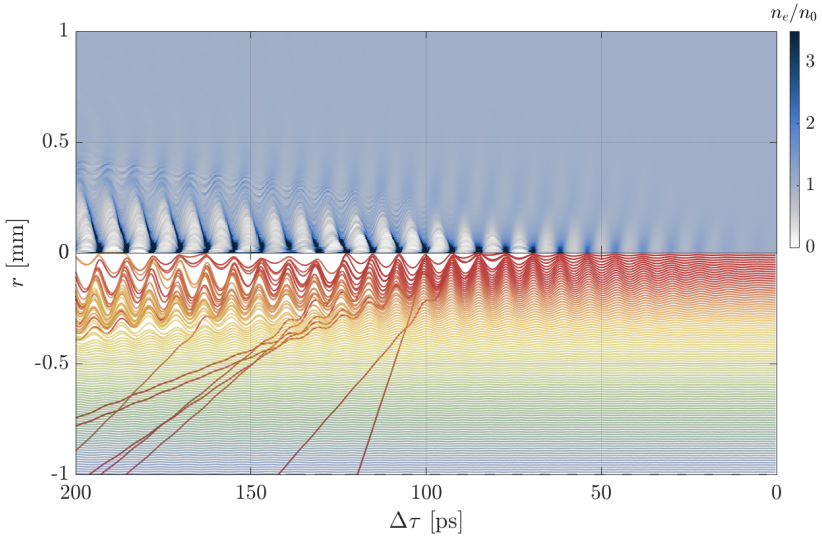

Figure 7.4.3: Simulated normalised plasma electron density map (upper) with corresponding plasma electron trajectories in the co-moving frame (lower) at a propagation distance of $z = 3\,\mathrm{m}$. The plasma electron trajectories are coloured according to their initial radial position and are separated by $\Delta r = 0.03\,k_p^{-1}$.

resonantly excited, the plasma electron density perturbation grows close to the axis leading to the development of an increasing radial density gradient. This is plotted in Figure 7.4.4(a) which depicts the simulated transverse plasma electron density profile for increasing delays behind the leading edge of the proton bunch. Due to the presence of the radial density gradient created by the wakefield, plasma electrons experience a spatially-dependent radial force [Figure 7.4.4(b)] and their initially coherent oscillations begin to mix [97]. This induces curvature in the transverse wakefield [Figure 7.4.4(c)], modifiying its resonant frequency across the radial plane, while the longitudinal field continues to be excited at the initial plasma frequency [240]. The wavefront of the density perturbation becomes deformed and electron trajectories start to cross as shown in Figure 7.4.3. For the plasma channel parameters simulated here approximately 5 % of plasma electrons are radially expelled from the plasma as innermost electrons cross trajectories of outer electrons and experience an increased negative charge density. Example trajectories of radially expelled electrons are shown in red in the lower panel of Figure 7.4.3. As plasma electron phase mixing is a cumulative effect that occurs over multiple plasma periods, it first begins at the rear of the proton bunch and quickly propagates forward as the wakefield potential, and hence density perturbation, associated with each micro-bunch grows as self-modulation develops along the length of the plasma [154]. After these inner electrons are expelled as demonstrated in the upper panel of Figure 7.4.5, outer plasma electrons move inwards to replace them and the plasma becomes charged at its boundary [241]. This, in combination with the presence of electrons outside the boundary of the plasma, induces a positive potential in the surrounding volume as shown in the lower panel of Figure 7.4.5. The potential is attracting for electrons and acts to accelerate the previously ejected electrons back towards the plasma.

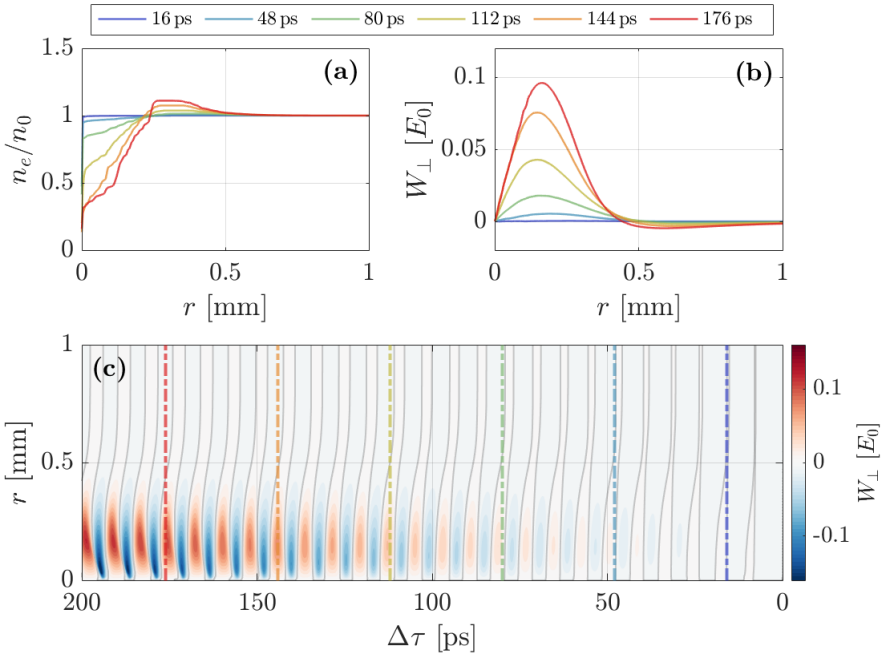

Figure 7.4.4: (a) Evolution of the transverse plasma electron density normalised to the plasma density, n_0, at increasing delays behind the leading edge of the proton bunch. (b) Corresponding transverse wakefield, normalised to the cold wave-breaking field, E_0. (c) Map of the transverse wakefield at a propagation distance of $z = 3\,\mathrm{m}$. The dashed, coloured vertical lines correspond to the timings shown in (a) and (b). Grey lines represent the zero-crossings of the transverse wakefield and highlight the development of its curvature.

As these electrons re-enter the plasma, they return back on-axis and interfere with the resonantly driven wakefield. The combination of this interference effect and the phase mixing induced by the radial density gradient damps the wakefield amplitude, causing the decay of the longitudinal wakefield observed in the lower panel of Figure 7.4.5 and the experimentally measured decrease in witness energy observed for large laser–electron delays shown in Figure 7.4.2. As the attracting force is low in amplitude when compared to the radial wakefield near the axis, expelled electrons re-enter the plasma far behind the position at which they are ejected from the plasma in the co-moving frame. For this reason, while trajectory crossing is observed to occur earlier than σ_z behind the seeding position in this simulation, significant effects from the return of plasma electrons on the wakefield amplitude do not become apparent until larger laser–electron delays, $\Delta\tau \geqslant 500\,\mathrm{ps}$.

This mechanism therefore splits the evolution of the wakefield amplitude into three distinct regions along the proton bunch:

(i) Initially for small laser–electron delays, $\Delta\tau \leqslant 250\,\mathrm{ps}$, the wakefield amplitude grows as an increasing number of proton micro-bunches contribute to driving the wakefield and plasma electron oscillations remain coherent. This is the region that is typically studied experimentally and theoretically as it offers the largest stable wakefield amplitudes (e.g. [120, 144, 242]). As the proton beam undergoes self-modulation and the wakefield

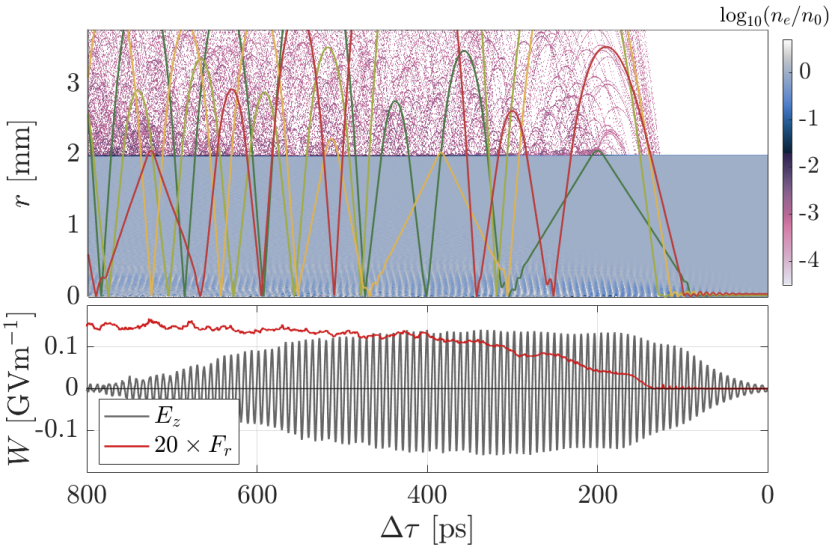

Figure 7.4.5: Upper: Simulated plasma electron density map at $z = 2.9\,\mathrm{m}$. Electrons are ejected from the plasma as the wakefield amplitude is resonantly excited and trajectory crossing occurs near the axis. Example plasma electron trajectories are plotted for reference (solid lines). The electron density perturbation driven by the self-modulated proton bunch can be seen near the axis. Lower: Corresponding on-axis longitudinal field (grey) shows decay for large $\Delta\tau$ alongside the development of a radial field outside the plasma boundary (red), measured at $r = 7.5\,k_p^{-1}$.

amplitude and density perturbation grow, trajectory crossing is observed in simulation towards the rear of this region and electrons are ejected from the plasma.

(ii) At around 250 ps behind the seeding position, the wakefield amplitude saturates. The wakefield amplitude is maintained in this region ($250 < \Delta\tau < 500\,\mathrm{ps}$) as outer plasma electrons replace those ejected via trajectory crossing. Phase mixing effects induced by the radial density gradient cause electron oscillations to decohere and limit further growth of the wakefield amplitude along the proton bunch despite an increasing number of micro-bunches driving the wakefield.

(iii) Finally, for $\Delta\tau \geqslant 500\,\mathrm{ps}$, a significant number of ejected electrons return to the plasma and their interference with the resonantly driven wakefield on-axis, in combination with accumulated phase mixing effects, cause further decay in the wakefield amplitude along the bunch.

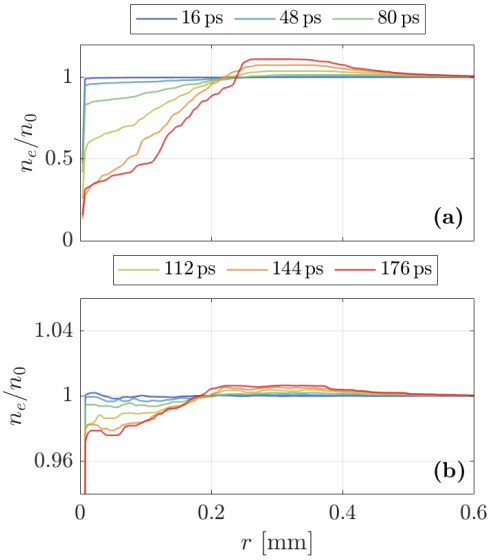

Figure 7.4.6: Comparison between the development of the radial electron density gradient at densities of (a) $n_0 = 2 \times 10^{14}\,\mathrm{cm}^{-3}$ and (b) $n_0 = 7 \times 10^{14}\,\mathrm{cm}^{-3}$. The transverse density lineout is plotted at specific delays with respect to the leading edge of the proton bunch, taken at equal normalised propagation distances: $t = 8000\,\omega_p^{-1}$.

Plasma density dependence

The appearance of the onset of plasma electron trajectory crossing so early within the resonantly excited wakefield ($\Delta\tau \approx 70\,\mathrm{ps}$) is a consequence of the low plasma density used in these studies. At a density of $2 \times 10^{14}\,\mathrm{cm}^{-3}$ the proton beam is narrow relative to the skin depth of the plasma, $\sigma_r \approx 200\,\mu\mathrm{m} = 0.53\,k_p^{-1}$, quickly exciting the large transverse gradients necessary to induce plasma electron trajectory crossing. At the AWAKE nominal density of $7 \times 10^{14}\,\mathrm{cm}^{-3}$ where the transverse bunch size is optimised for the plasma density ($\sigma_r = k_p^{-1}$), the relative bunch density is reduced meaning the growth of the wakefield amplitude, and hence the transverse density gradient, is far slower. This is highlighted in Figure 7.4.6 which shows the simulated evolution of the transverse electron density profile at specific delays behind the leading edge of the proton bunch propagating in background plasma densities of (a) $2 \times 10^{14}\,\mathrm{cm}^{-3}$ and (b) $7 \times 10^{14}\,\mathrm{cm}^{-3}$ at equal *normalised* propagation distances of $t = 8000\,\omega_p^{-1}$. At the operating density used for these measurements ($2 \times 10^{14}\,\mathrm{cm}^{-3}$), the transverse density gradient grows quickly and the wakefield quickly approaches full blowout near the axis. In contrast, at the nominal density ($7 \times 10^{14}\,\mathrm{cm}^{-3}$) electron density perturbations are typically only of the order of a few percent. The reduction in transverse density gradient means that plasma electron oscillations remain coherent over longer timescales and the onset of plasma electron trajectory crossing is not observed in simulation until further behind the leading edge of the proton bunch ($\Delta\tau \approx 400\,\mathrm{ps}$). At the nominal density, it is rather expected that motion of plasma ions near the axis causes decoherence and decay of the wakefield on extended timescales [243–245], an effect not observed in the simulations performed for this study at the operating density.

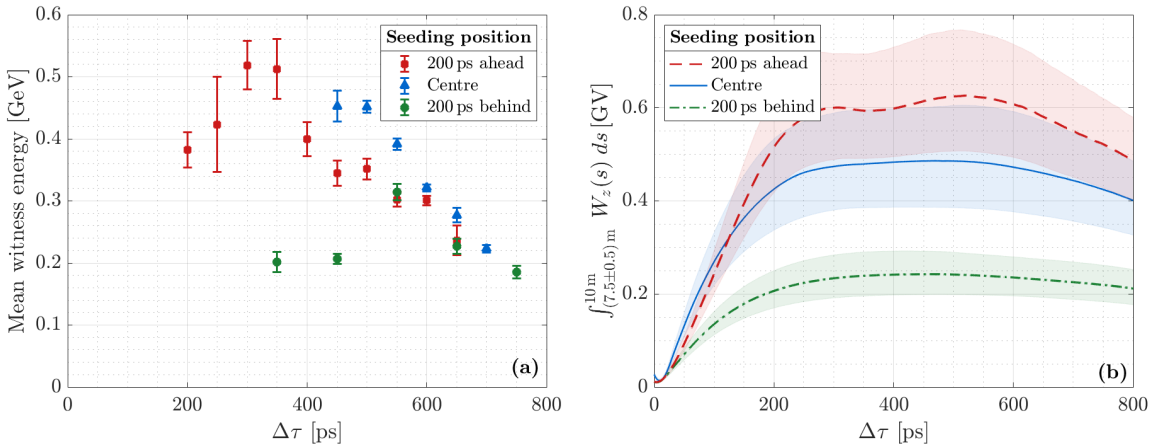

Figure 7.4.7: (a) Experimental results showing the mean witness energy with increasing laser–electron delays for the three measured seeding positions (repeat of Figure 7.4.2 for ease of comparison). (b) Corresponding simulation results showing the envelope of the integrated longitudinal wakefield following self-modulation saturation. Error bands represent the integrated longitudinal wakefield calculated for upper and lower bounds of the estimated saturation length $z_s = 7.5 \pm 0.5\,\mathrm{m}$.

Comparison to simulation

Figure 7.4.7 shows a comparison between the experimental measurements of the mean witness energy in Figure 7.4.2 and the envelope of the integrated longitudinal wakefield for different laser–electron delays predicted by LCODE simulations using the proton bunch parameters of the experiment. The wakefield amplitude is integrated over the plasma length following self-modulation saturation, identified in simulation by the position after which the wakefield phase remains approximately constant with respect to the proton micro-bunches. The procedure for identifying the saturation position, z_s, is demonstrated in Figure 7.4.8. This figure shows the evolution of the positions of the zero-crossings of the on-axis longitudinal wakefield, a useful proxy for mapping the phase evolution of the wakefield. The colour of a line at a given position represents the instantaneous rate of change of the phase of the wakefield at that location, with darker colours representing larger values. As discussed in Section 3.2.2, as the proton bunch undergoes self-modulation large changes in the wakefield phase are observed. However, the wakefield phase is observed to stabilise at approximately $z_s = 8.0 \pm 0.5\,\mathrm{m}$, with the exact value of z_s depending on the delay with respect to the leading edge of the proton bunch, τ. In general, z_s increases slightly with increasing τ as the phase shift at a given value of τ is accumulated from all preceding micro-bunches.

The self-modulation saturation position corresponds to the position in the plasma at which consistent acceleration of witness electrons is expected to be possible as discussed in Sec-

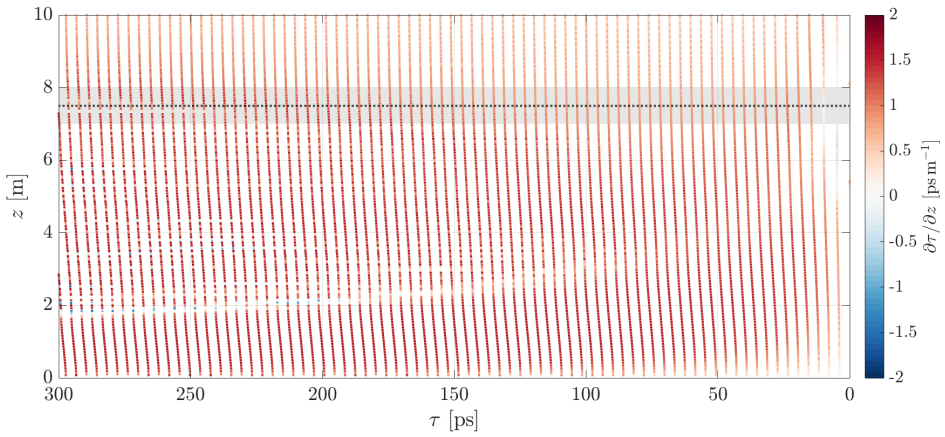

Figure 7.4.8: Simulated evolution of the positions of the zero-crossings of the on-axis longitudinal wakefield amplitude for $0 \leqslant \tau \leqslant 300\,\mathrm{ps}$ over the entire plasma length in the co-moving frame. The colour of a line corresponds to the instantaneous rate of change of phase of the wakefield, $\partial\tau/\partial z$, with darker colours representing larger values. The wakefield stabilises between approximately $7.0 < z_s < 8.0\,\mathrm{m}$ (grey region), at which point the self-modulation process is said to have saturated.

tion 3.3.2. For this reason, it is chosen as the lower bound for the longitudinal wakefield integration for comparison to experimental results:

$$\int_{z_s}^{10\,\mathrm{m}} W_z(s)\,\mathrm{d}s; \quad z_s = 7.5 \pm 0.5\,\mathrm{m}. \tag{7.4.1}$$

The large uncertainty associated with the exact value of z_s is chosen to represent reasonable bounds on the expected witness energy. Even so, the integrated wakefield amplitude shown in Figure 7.4.7(b) represents an upper bound on the expected witness energy as it assumes witness electrons remain on-axis in the position of maximal electric field amplitude over the relevant integrated plasma length. It is rather expected that witness electrons oscillate within the accelerating region of the wakefield and therefore have lower energy at the exit of the plasma than predicted by integration of the simulated longitudinal wakefield, consistent with the data presented in Figure 7.4.7(a). The trends observed experimentally (Figure 7.4.7(a)) are reproduced in simulation (Figure 7.4.7(b)) with three distinct regions: initial growth of the wakefield as it is resonantly excited ($\Delta\tau < 250\,\mathrm{ps}$), a region of field saturation ($250 \leqslant \Delta\tau < 550\,\mathrm{ps}$), and decay on longer timescales. The experimental results presented in Figure 7.3.1 for a seeding position 100 ps ahead of the centre of the proton bunch are also consistent with these findings, but show only the initial growth and saturation regions of the integrated wakefield amplitude due to the limited range of laser–electron delays measured ($\Delta\tau \leqslant 365\,\mathrm{ps}$).

The magnitude of the decay of the integrated wakefield amplitude on long timescales ($\Delta\tau > 500\,\mathrm{ps}$) shown in Figure 7.4.7(b) is dependent on the seeding position. This is expected as trajectory crossing and the subsequent ejection of plasma electrons is a direct result of the phase mixing induced by the increasing radial density gradient from the cumulative ex-

citation of the wakefield. When seeding 200 ps behind the centre of the proton bunch, the initial proton bunch density is lower and decreasing along the bunch when compared to the other two seeding positions. Consequently the wakefield amplitude driven by the modulated proton bunch is lower as demonstrated in Figure 7.1.2 and simulations indicate that trajectory crossing does not occur until later in the plasma ($z > 3$ m) and further behind the seed position ($\Delta\tau > 300$ ps). This is exemplified in Figure 7.4.9 which shows the simulated development of the radial field outside the plasma boundary ($r = 7.5\, k_p^{-1}$) over the entire plasma length for the three different seeding positions tested experimentally. The peak amplitude of the potential induced outside the boundary of the plasma when seeding 200 ps behind the centre of the proton bunch (Figure 7.4.9(a)) is an order of magnitude lower than for the other seeding positions, indicating far fewer electrons being ejected from the plasma. Therefore, the long timescale wakefield decay induced by the return of ejected plasma electrons is reduced for this seeding position as shown in Figure 7.4.7(b). This was similarly observed in the experimental measurements shown in Figure 7.4.7(a) where the mean witness energy is approximately constant with increasing delay when seeding 200 ps behind the centre.

In comparison, when seeding 200 ps ahead of the centre of the bunch the number of protons contributing to driving the wakefield is far higher. Larger amplitude wakefields are driven earlier in the plasma (Figure 7.1.2) and the transverse potential outside the plasma boundary begins to develop within 2 m (Figure 7.4.9(c)). The amplitude of the potential also exceeds that of the central seeding position shown in Figure 7.4.9(b). Larger on-axis wakefield amplitudes are driven when seeding 200 ps ahead of the centre of the bunch and correspondingly more electrons undergo trajectory crossing and can later return to the plasma and damp the wakefield on-axis. This results in faster decay of the integrated wakefield amplitude for large laser–electron delays than for the central seeding position as demonstrated in Figure 7.4.7(b).

In general, the decay of the integrated longitudinal wakefield amplitude on extended timescales is slower than that of the experimentally-measured mean witness energy. There are two likely reasons for this:

(i) The definition of the saturation length, z_s. For simplicity, a constant value of z_s is assumed across the entire range of $\Delta\tau$ within the calculation of the integrated longitudinal wakefield. As discussed previously the saturation length will vary along the proton bunch, slightly increasing with increasing $\Delta\tau$. A later saturation length reduces the effective acceleration length of witness electrons that are injected at large laser–electron delays, further reducing their expected energy at the plasma exit. This motivates the large range in z_s used in the calculations: for smaller values of $\Delta\tau$, the witness energy would be expected to be closer to the upper bound of the integration calculation ($z_s = 7.0$ m) but tends towards the lower bound ($z_s = 8.0$ m) as $\Delta\tau$ increases.

(ii) Use of axisymmetric simulations means that turbulent effects from plasma electron phase mixing and the return of plasma electrons are likely underestimated leading to an overestimation of the longitudinal wakefield amplitude, particularly on extended timescales.

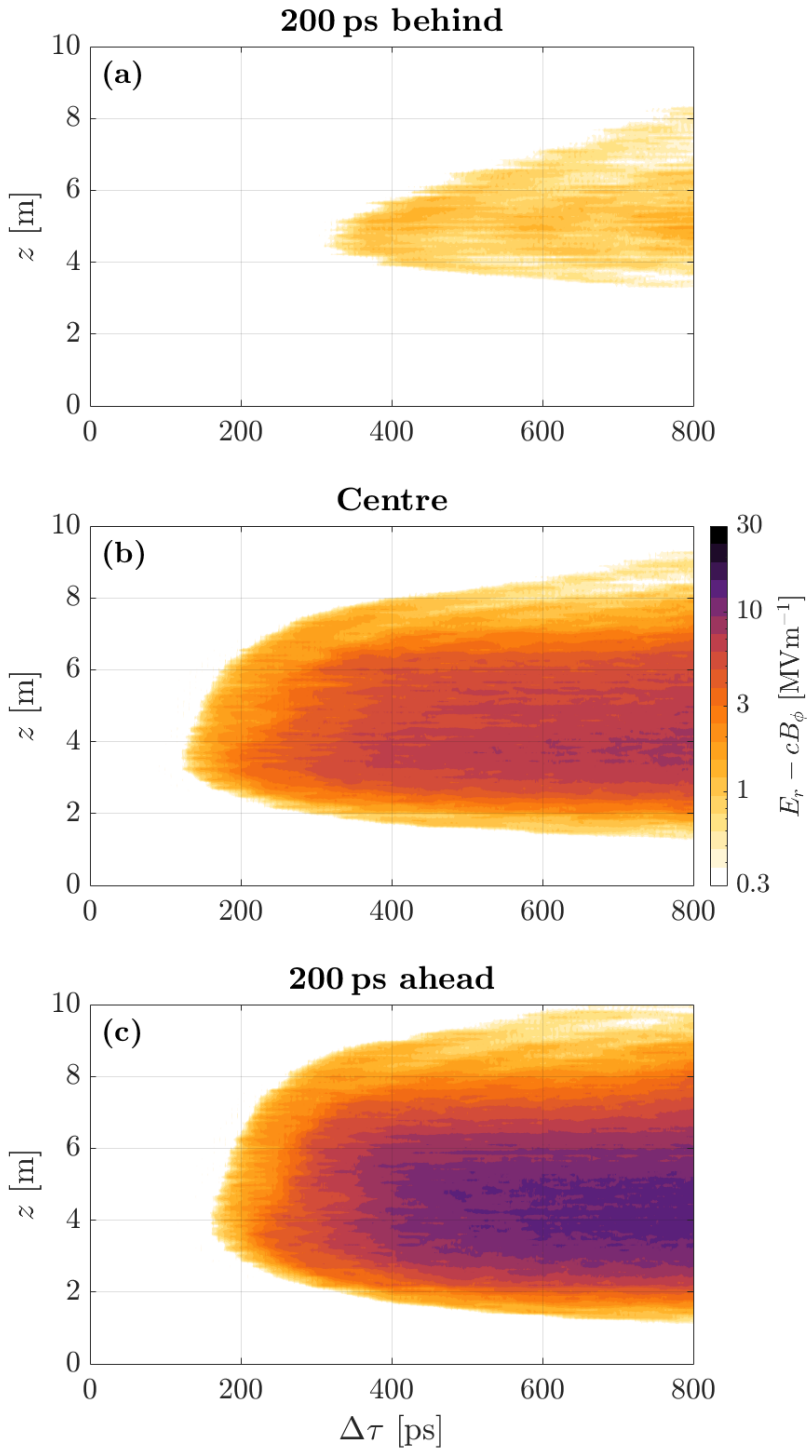

Figure 7.4.9: Simulated development of the transverse field, $W_\perp = E_r - cB_\phi$, outside the boundary of the plasma measured at $r = 7.5\,k_p^{-1}$ for three seeding positions: (a) 200 ps behind the centre of the proton bunch, (b) in the centre of the proton bunch and (c) 200 ps ahead of the centre of the proton bunch. All images use the same colour scale, logarithmically-spaced to highlight the difference in magnitude of the field for the three cases.

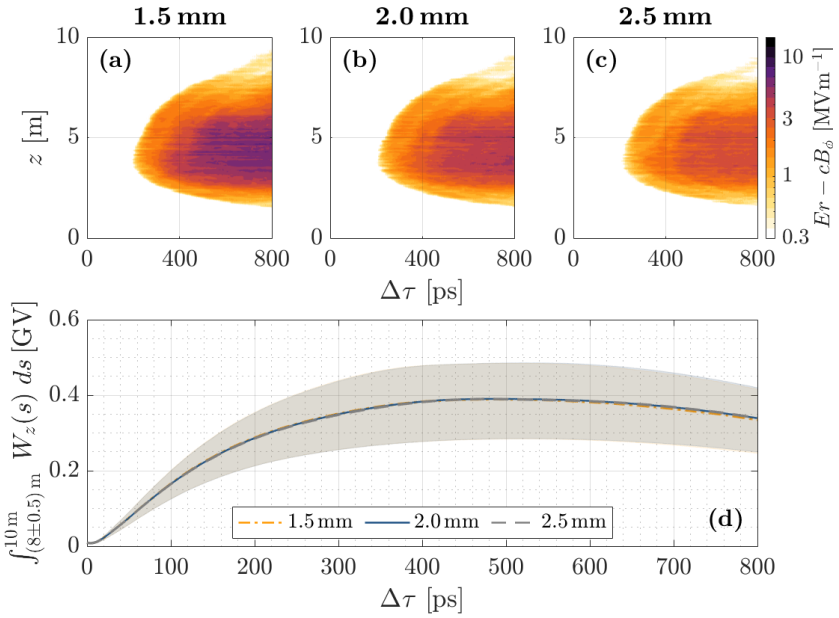

Figure 7.4.10: Effect of moving the position of the plasma boundary (r_p) in simulation. Upper: Simulated development of the transverse field, $W_\perp = E_r - cB_\phi$, at a position 0.8 mm ($2.2\,k_p^{-1}$) outside the plasma boundary for $r_p = 1.5$ mm (a), 2.0 mm (b) and 2.5 mm (c). (d) Corresponding integrated on-axis longitudinal wakefield following self-modulation saturation.

Sensitivity to plasma boundary position

The position of the plasma boundary, assumed to be $r_p = 2$ mm in this study, is not measured experimentally and hence has a large uncertainty. Simulations varying the position of the boundary between reasonable experimental limits, $r_p = 1.5 \rightarrow 2.5$ mm, were performed to investigate the effect of this for a central seeding position and are shown in Figure 7.4.10. When reducing the boundary position to $r_p = 1.5$ mm, a 35% increase in the average amplitude of the field generated outside the plasma boundary was observed as an increased number of plasma electrons gained sufficient transverse momentum following trajectory crossing to be ejected from the plasma. In contrast, a 20% drop in the average field amplitude was observed when increasing the boundary position to 2.5 mm. However, a correspondingly smaller change in the integrated longitudinal wakefield amplitude (%-level) was observed as shown in Figure 7.4.10(d). Therefore, over the range of positions tested in this study, the amplitude of the integrated longitudinal wakefield was found to be relatively insensitive to the exact position of the plasma boundary. Despite this, the position of the plasma boundary may be expected to have a significant effect on the average witness charge and is discussed in more detail in the following section.

7.4.2 *Witness capture*

The development of the potential outside the plasma boundary also affects the amount of witness charge captured within the wakefield. As discussed previously in Section 3.4.3, simulations of witness injection have indicated that a significant fraction of witness electrons injected at the entrance of the plasma can be defocused near the plasma boundary [239] or by the seed wakefields driven by the unmodulated proton bunch [246]. These electrons are then observed to oscillate around or near the boundary of the plasma and continue to propagate along its length. In addition to this, some witness electrons that are captured in the wakefield during the initial injection can be lost from the wakefield later in the plasma due to dephasing caused by the continuous evolution of the wakefield phase during self-modulation. In simulation these electrons leave the plasma quickly and are then also observed to continue to propagate along the plasma length near its boundary. In this study it was found that simulations indicate that this population of witness electrons can be re-injected following the development of the potential induced outside the boundary of the plasma, increasing the captured charge at the exit of the plasma.

This is demonstrated in Figure 7.4.11 which shows the simulated maps of witness capture positions along the plasma for the three different seeding positions measured experimentally for variable laser–electron delays. Example witness electron trajectories are additionally shown in Figure 7.4.12(d); some electrons remain close to the axis and trapped within the wakefield over the entire plasma length, while others are observed to propagate near the plasma boundary and are injected following the development of the potential outside the plasma boundary at $z > 5\,\mathrm{m}$.

When seeding in the centre of the proton bunch for small laser–electron delays (Figure 7.4.11 (b), $\Delta\tau \leqslant 200\,\mathrm{ps}$), capture only occurs at the position of the initial injection of the witness bunch at the start of the plasma ($z < 1\,\mathrm{m}$). This is the region of the proton bunch ahead of the development of the potential outside the plasma boundary demonstrated in Figure 7.4.9 (b). However, for $\Delta\tau \geqslant 300\,\mathrm{ps}$ a growing amount of charge is re-injected and captured in the wakefield later in the plasma ($6 \leqslant z \leqslant 9\,\mathrm{m}$) with increasing laser–electron delay. This directly corresponds to the evolution of the electron-attracting potential outside the boundary of the plasma illustrated in Figure 7.4.9. Similarly, when seeding 200 ps ahead of the centre of the proton bunch (Figure 7.4.11 (c)), significant re-injection and capture is observed following the development of the potential outside the plasma boundary ($z > 5\,\mathrm{m}$). In this case, an increased fraction is re-injected and captured when compared with seeding in the centre of the proton bunch, consistent with the larger amplitude potential induced outside the plasma boundary demonstrated in Figures 7.4.9 (b) and (c). In simulation, for both seeding positions, there is not a significant difference in the final energy of witness electrons that are captured either at the start of the plasma and remain within the wakefield throughout the entire development of the self-modulation ($z < 1\,\mathrm{m}$), or those that are re-injected and captured in the wakefield following the development of the potential outside the plasma boundary ($z > 5\,\mathrm{m}$) as shown in Figure 7.4.12. Witness electrons that are captured in the

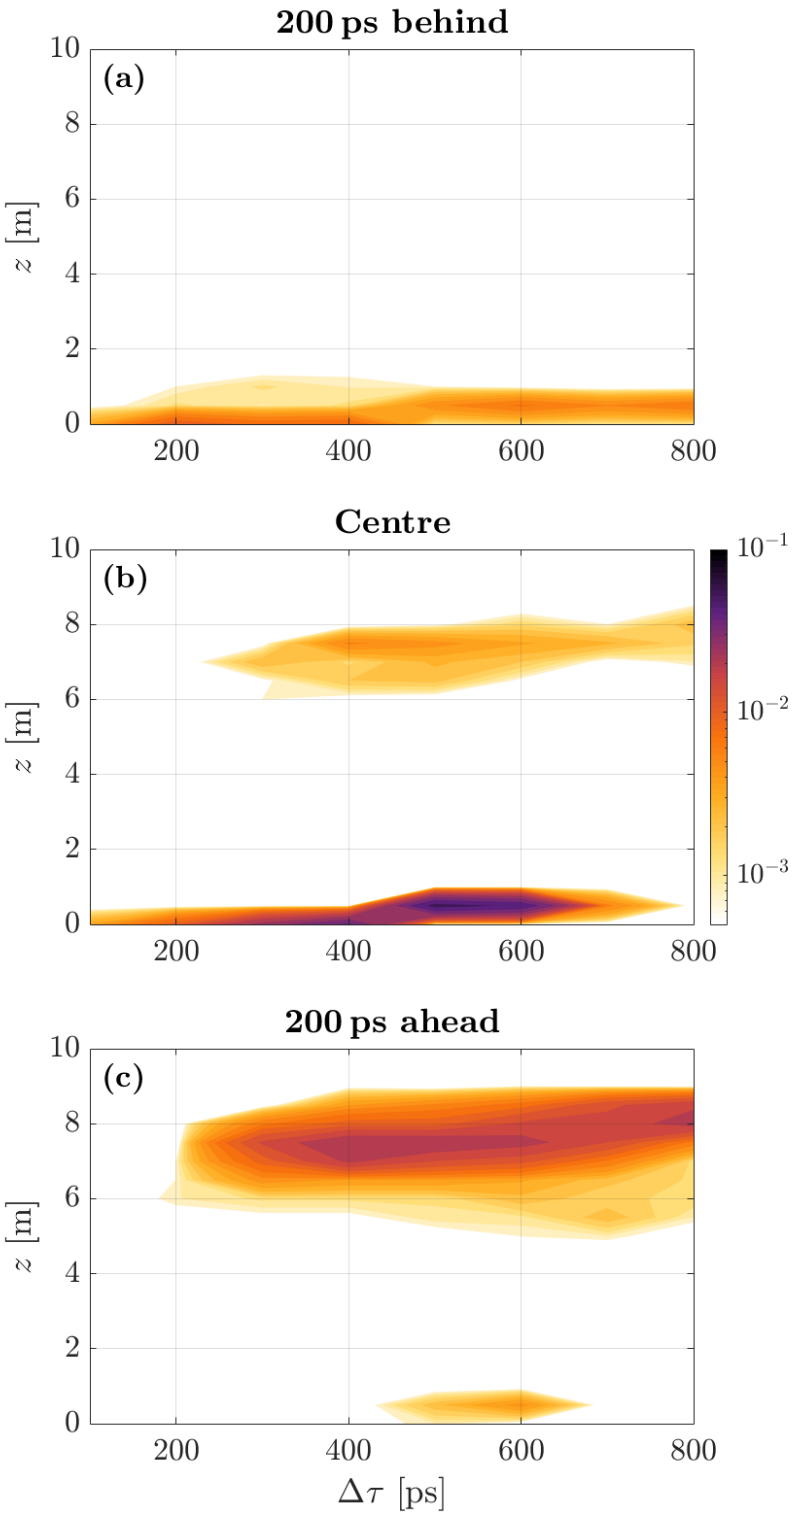

Figure 7.4.11: Maps of final injection position for witness electrons from simulations of electron injection. A separate simulation is performed for each laser–electron delay in 100 ps steps. The colourscale represents the fraction of electrons that are captured in the wakefield at a particular z-position, normalised to the total number of witness electrons that are injected at the start of the simulation. A cut on the witness electron energy of $E \geqslant 84$ MeV is used to mirror the experimental conditions for the spectrometer dipole.

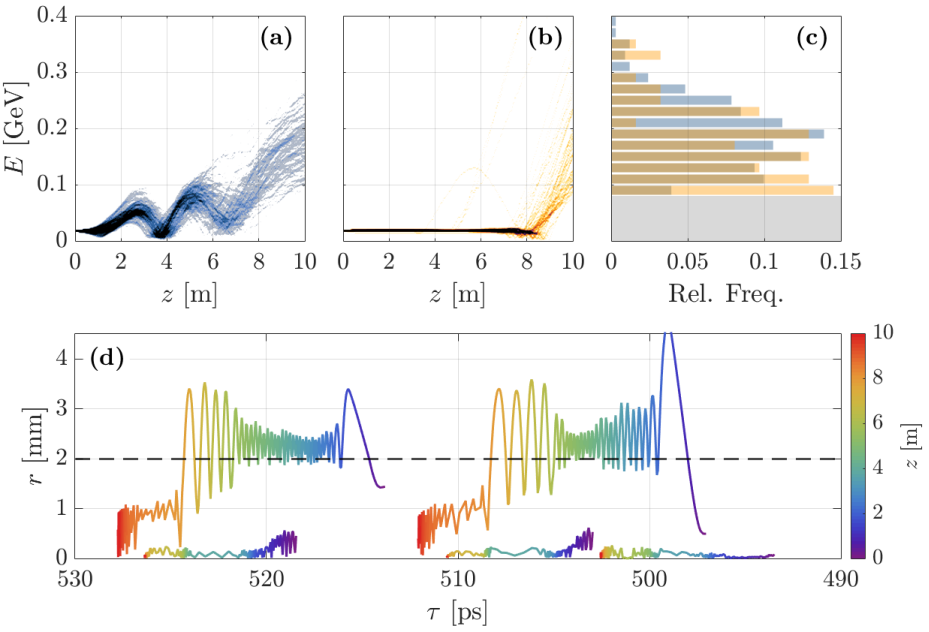

Figure 7.4.12: Simulated evolution of the witness energy for (a) electrons that remain trapped in the wakefield throughout self-modulation, and (b) electrons that are injected following development of the electron-focusing potential outside the plasma boundary. (c) Normalised histogram of witness energy at the plasma exit. Values are normalised to the total number of electrons in each injection population to allow comparison of their distributions. (d) Example witness electron trajectories in the co-moving frame; the colour scale represents the z-position of an electron. The black horizontal dashed line represents the position of the plasma boundary. This simulation corresponds to a central seeding position with $\Delta\tau = 500\,\text{ps}$.

wakefield at the start of the plasma undergo dephasing as the wakefield phase continuously evolves and are decelerated multiple times before self-modulation saturates at $z \approx 7\,\text{m}$, at which point consistent acceleration is observed. Therefore, the effective acceleration length for witness electrons captured either at the start of the plasma (Figure 7.4.12(a)) or following re-injection by the potential outside the plasma boundary (Figure 7.4.12(b)) is similar. This further motivates using the saturation position ($z_s = 7.5 \pm 0.5\,\text{m}$) as the lower limit within the calculation of the integrated wakefield in Figure 7.4.7(b) when comparing to experimental results. The energy distribution of the two injection populations at the plasma exit are shown in Figure 7.4.12(c) normalised to the number of electrons in each population and demonstrates a similar energy distribution for both. When seeding 200 ps behind the centre of the proton bunch, successful capture of re-injected witness electrons late in the plasma is not observed in simulation (Figure 7.4.11 (a), $z > 1\,\text{m}$). This is due to the reduced amplitude of the focusing potential outside the boundary of the plasma in this case, as demonstrated in Figure 7.4.9 (a).

For comparison, identical analysis to that presented in Figure 7.4.12 is performed for a simulation using a central seeding position and an injection delay of $\Delta\tau = 200\,\text{ps}$ and is shown in Figure 7.4.13. In this case, the initial injection position is ahead of the effects from the

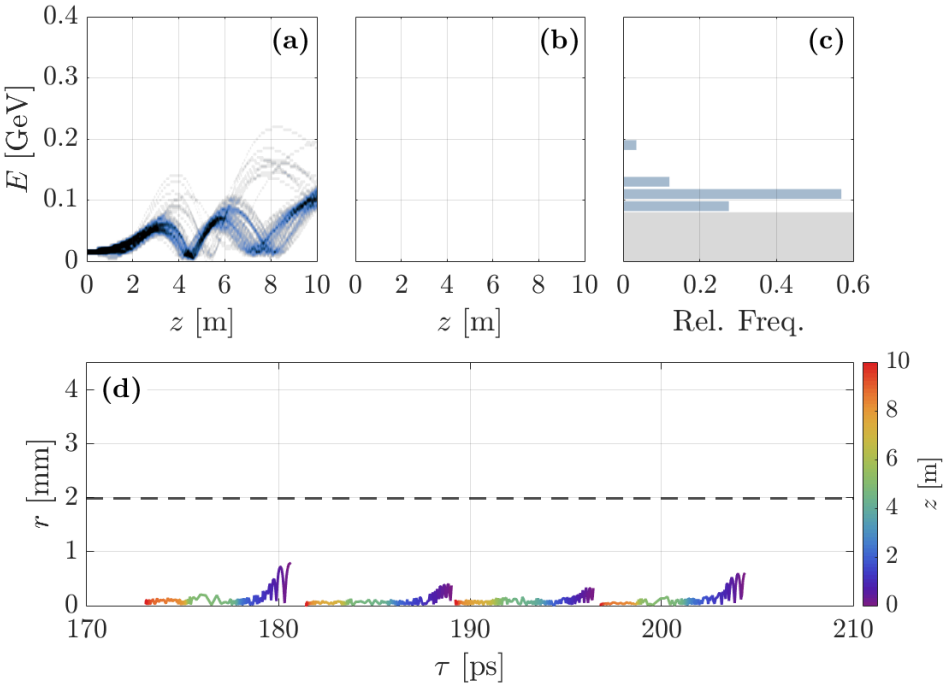

Figure 7.4.13: Identical analysis of simulated witness electrons as in Figure 7.4.12, but for $\Delta\tau = 200\,\text{ps}$ for a central seeding position. In this case, no witness electrons are re-injected later in the plasma as the initial injection position is ahead of the effects due to the onset of plasma electron ejection.

onset of plasma electron trajectory crossing and ejection as can be seen in Figure 7.4.9(b) and hence no electrons are re-injected at later propagation distances. The trajectories of randomly selected witness electrons in the co-moving frame in panel (d) all correspond to electrons that are captured at the beginning of the plasma and remain within the wakefield, propagating close to the axis over the entire plasma length. Note also the effect of the difference in saturation length for the two injection delays in these figures. In panel (a) in Figure 7.4.13, consistent acceleration of witness electrons is observed after the second rephasing cycle at $z \sim 7.5 - 8.5\,\text{m}$, whereas it occurred at approximately $z \sim 7\,\text{m}$ in Figure 7.4.12(a). The reason for this has been discussed at length previously, but these results further motivate the use of a wide uncertainty range on the definition of the saturation length for calculations of the integrated longitudinal wakefield amplitude in Figure 7.4.7.

Experimentally, corresponding signatures consistent with the seed-dependent effect of witness re-injection following the development of a radial potential outside the plasma boundary were observed and are demonstrated in Figure 7.4.14. This figure shows (a) the average witness charge measured experimentally and (b) observed in simulation when varying the laser–electron delay. A minimum witness energy of 84 MeV at the exit of the plasma was applied in simulation analysis in order to mimic the spectrometer dipole settings used in the experiment. When seeding ahead of, or in, the centre of the proton bunch, an increasing witness charge was measured with increasing laser–electron delay. This is consistent with the

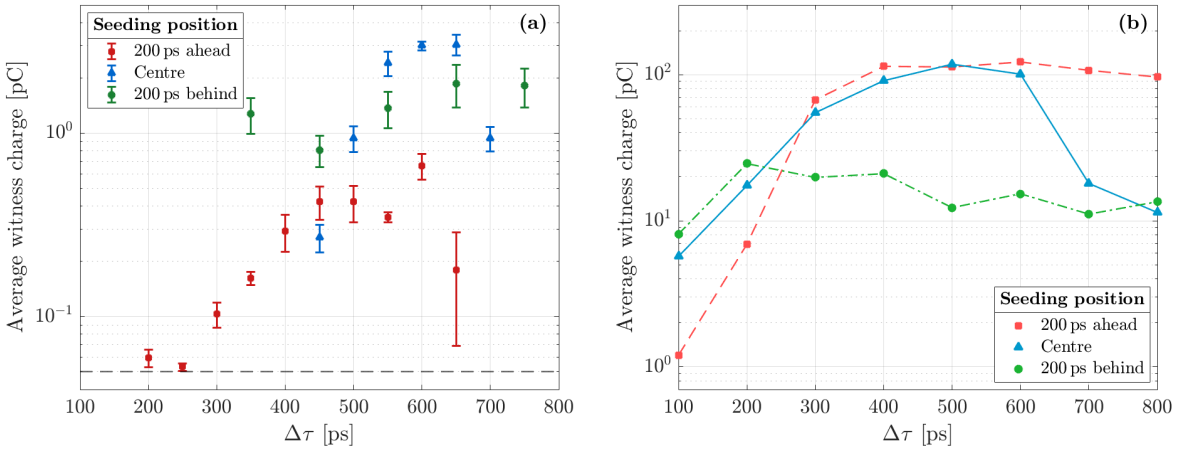

Figure 7.4.14: (a) Experimental and (b) simulation results demonstrating the mean witness charge as a function of the laser–electron delay. Trends observed experimentally are consistent with those produced in simulation. Experimental error bars represent the standard error on the mean. The horizontal dashed line at 50 fC represents the minimum required charge measured on the spectrometer screen for an event to be included in the analysis.

evolution of the electron focusing field outside the boundary of the plasma created following plasma electron ejection (Figure 7.4.9) acting to re-inject witness electrons where they can be captured within the wakefield. However, when seeding behind the centre of the proton bunch approximately constant witness capture with increasing delay was observed both experimentally and in simulations. These experimental observations directly correspond to the trends predicted by the witness injection simulations for the three seeding positions demonstrated in Figure 7.4.14(b). For $\Delta\tau = 700\,\text{ps}$ with a central seeding position, a drop in the average witness charge is observed both experimentally and in simulation. This corresponds to a large reduction in the amount of charge captured close to the entrance of the plasma ($z < 1\,\text{m}$), as shown in Figure 7.4.11(b). In contrast, when seeding 200 ps ahead of the centre of the proton bunch a smaller fraction of witness electrons are captured for $z < 1\,\text{m}$ at all $\Delta\tau$, resulting in a more consistent average witness charge at large $\Delta\tau$ both experimentally and in simulation as the re-injection mechanism dominates.

As discussed in Section 7.4.1, changing the position of the plasma boundary in simulation results in $\mathcal{O}(10\%)$ changes in the average amplitude of the transverse field generated outside the boundary by the expulsion of plasma electrons. While this was not observed to have a significant effect on the integrated on-axis longitudinal wakefield amplitude in simulation, it may be expected to have a larger effect on the captured witness charge. Reducing the transverse extent of the plasma and the corresponding increase in the transverse field amplitude could lead to a greater number of witness electrons gaining enough transverse momentum outside the plasma boundary to be successfully captured in the wakefield upon their re-injection into the plasma and vice versa. This could be explored further in future experiments by varying the ionising laser pulse energy to modify the position of the plasma boundary [207].

As expected, the simulations predict average captured witness charges that are more than an order of magnitude larger than those measured experimentally for geometric reasons discussed previously. While the development of this re-injection mechanism appears to offer a method for increasing charge capture experimentally, it is important to note that this secondary injection is not controlled. It will not result in high quality witness bunches and as such is not desirable.

Wider implications for AWAKE Run 2

The injection scheme for AWAKE Run 2 differs from that used in Run 1 in order to improve capture, and aid in the preservation of the emittance and minimisisation of the energy spread of the injected witness bunch [143]. This will be achieved by splitting the plasma interaction in two; the proton bunch will undergo self-modulation and reach saturation in a first vapour source before being transported in vacuum to a second vapour source. Witness electrons will be injected on-axis in vacuum in the drift section between the two vapour sources, and accelerated in the wakefield driven by the proton micro-bunch train in the second plasma section. The electron beam will be injected into large amplitude wakefields driven by the saturated micro-bunch train (rather than the seed wakefields in Run 1) and hence will experience large electron-focusing regions of the wakefield at the entrance of the plasma meaning capture efficiency should grow significantly. This will reduce the population of witness electrons that are defocused and end up outside the plasma boundary and hence minimise sources of the secondary injection mechanism identified in this study that can degrade the witness bunch quality.

The largest acceleration gradients were observed ahead of the onset of effects created by trajectory crossing and as such the long timescale wakefield effects explored here are not expected to limit the accelerating gradient, energy gain or charge capture for AWAKE Run 2. The use of longitudinal plasma density gradients/variations will be limited to the first vapour source where self-modulation occurs while the second vapour source will have a constant longitudinal density and hence no further modification of the secondary injection mechanism would be expected by their use.

7.5 CONCLUSION

The evolution of the amplitude of the wakefield driven along a self-modulated proton bunch in plasma is measured experimentally by varying the relative timing of the wakefield seeding position and the injection position of witness electrons. Measurements of the mean witness energy at the plasma exit demonstrate resonant excitation of the longitudinal wakefield on timescales below the RMS bunch length. The use of a low plasma density ($n_e = 2 \times 10^{14}\,\mathrm{cm}^{-3}$) and a correspondingly narrow drive bunch ($\sigma_r = 0.53\,k_p^{-1}$) results in the fast development of strong radial transverse plasma density gradients. The presence of radial gradients causes phase mixing and plasma electron trajectory crossing within the wakefield, leading to the expulsion of plasma electrons. An electron-focusing potential is induced in the volume outside the plasma boundary and acts to accelerate electrons back into the plasma. The return of electrons damps the wakefield amplitude while increasing witness capture on extended timescales.

The seeding position is changed to investigate the effect of the longitudinal wakefield amplitude on the development of the electron-focusing potential outside the plasma boundary showing amplified effects for higher amplitude wakefields and vice versa. Agreement between the witness energy and charge capture trends observed both experimentally and in simulation provides evidence that the mechanisms identified in this study are the cause of the decrease in energy gain and increase in charge capture measured on extended timescales. This study therefore contributes useful information for optimisation of the acceleration process for AWAKE Run 2 and beyond.

It is expected that the occurrence of trajectory crossing within the resonantly excited wakefield is a result of using a narrow drive beam and hence should be suppressed at higher plasma densities where the relative transverse beam size is larger. Therefore, to further study the development of this mechanism, similar measurements could be repeated at a range of plasma densities. This would allow determination of the relative transverse beam size at which the experimental signatures associated with trajectory crossing can be observed, thus providing a useful cross-check with simulation predictions and helping to optimise beam and plasma parameters for future applications.

8

FEW-NANOSECOND PLASMA WAKEFIELD EVOLUTION

Generally, in PWFA operation, plasma ions are assumed to form a static background due to their large mass and hence their motion is not usually considered. The timescale on which their motion should be considered can be approximated in the same way as for plasma electrons, according to their plasma frequency:

$$\omega_{pi} = \sqrt{\frac{n_i Z^2 e^2}{m_i \varepsilon_0}}, \tag{8.0.1}$$

identical to that of plasma electrons, but with their charge represented by Ze where Z is the charge of the ion, and m_i is the mass of the ion. For a singly-ionised plasma where $Z = 1$ and $n_e = n_i$, the plasma ion frequency is therefore calculated by scaling the plasma electron frequency by the square root of the electron-ion mass ratio:

$$\omega_{pi} = \sqrt{\frac{m_e}{m_i}}\omega_{pe}. \tag{8.0.2}$$

In the case of hydrogen plasma ($m_i = m_p = 1836 m_e$) the timescales therefore differ by a factor of approximately 40, increasing for plasma species with higher masses. At typical FLASHForward operating plasma densities of $n_e \sim 10^{15} - 10^{16}\,\mathrm{cm}^{-3}$ in argon ($m_i = 40 m_p$) plasma, the characteristic ion timescale is $\omega_{pi}^{-1} \sim 50 - 150\,\mathrm{ps}$, indicating the motion of ions driven by the plasma wakefield process will occur over nanosecond timescales. This therefore potentially represents an impediment to the ability of plasma wakefield accelerators to demonstrate repeatable acceleration at high-repetition-rates if their motion does not dissipate on compatible timescales.

This chapter details measurements performed at FLASHForward within the context of the X-3 experiment using the beam-based plasma characterisation technique introduced in Chapter 5 that aimed to investigate the development of ion motion following the wakefield excitation process. Signatures within the probe bunch energy and transverse spectra are used to identify ion motion evolution, allowing its dependence on the wakefield amplitude, plasma density and plasma species to be investigated.

8.1 ION MOTION IN PWFA

Previous numerical studies of the effects of ion motion within PWFA have typically been limited to short timescale effects that occur within the first plasma electron oscillation period as this can be more easily simulated for reasons discussed in Section 2.3.2. Such motion occurs for the case of ultra-intense drive bunches where the bunch density is many orders of mag-

nitude larger than the background plasma density [247]. In this case, the charge density of the bunch is sufficiently high that ions are attracted towards the axis of propagation within the excitation period of the wakefield, causing non-linear focusing fields to be generated at the rear of the cavity and severely degrading the emittance of the trailing witness beam. This realisation led to the development of alternative acceleration schemes including the tailoring of the transverse bunch phase space [248], or using multiple plasma sections of reducing ion mass to adiabatically match the witness bunch into the accelerating section [249], in spite of the presence of background ion motion to allow the witness bunch quality to still be maintained throughout the acceleration process. In contrast, an unexpected benefit of ion motion was detailed in Ref. [250] where the head-to-tail variation in focusing force created by the development of ion motion within the wake caused the decoherence of beam slice betatron oscillations and suppressed growth of the hosing instability. The above examples generally assume collider-relevant bunch parameters with nano-coulomb charges compressed to tens of kA peak currents and (sub-)micron-scale normalised emittances and hence represent normalised bunch densities on the order of 10^4, far beyond those available at FLASHForward.

For more modest bunch parameters, ion motion will not develop within the first wakefield period, but rather over extended timescales. This is of particular importance for wakefield acceleration concepts based on resonant excitation of the wakefield over multiple periods such as in the AWAKE experiment [61] or in multi-pulse laser wakefield acceleration schemes [251, 252]. For the AWAKE experiment, it was shown in simulation that even for wakefields excited in the linear regime, small perturbations in the ion background ($\delta n_i/n_0 \sim 10^{-2}$) induced by transverse wakefields following the development of self-modulation caused turbulent mixing of plasma electrons and suppressed the wakefield amplitude towards the rear of the bunch [243, 244]. This effect could be overcome by using a heavier plasma species, motivating the use of rubidium within the final experiment. However, ion motion effects are still expected to play a part on extended timescales ($\sim$ns) in Rb at the nominal AWAKE operating density ($n_0 = 7 \times 10^{14}\,\mathrm{cm}^{-3}$), possibly enhancing the wakefield amplitude as wake energy is redistributed on-axis via the motion of plasma ions [245]. In the measurements presented in Chapter 7, a lower operating density of $n_0 = 2 \times 10^{14}\,\mathrm{cm}^{-3}$ meant such ion motion effects were not observed in simulation over the nanosecond timescale studied experimentally.

Ion motion on extended timescales has been directly observed in two experiments to date: a cone-like feature was observed to trail electron bunch-driven wakefields in shadowgraphy measurements in Ref. [83], while the presence of a transversely propagating plasma boundary over nanosecond timescales was measured using a probe laser pulse propagating at a shallow angle in Ref. [84]. These two measurements confirm the presence of ion dynamics on extended timescales following the PWFA process, but do not make any attempt to study its dependencies on relevant experimental parameters such as the plasma density, drive bunch current, or plasma species. Nor are they able to comment on the lifetime of the perturbation created by the motion of ions. The measurements outlined in this chapter attempted to perform such studies using an extension to the beam-based plasma characterisation technique introduced in Chapter 5.

8.2 GENERATION OF NON-UNIFORM TRANSVERSE DENSITY PROFILES

We start by considering the effects of the wakefield on the motion of ions within the plasma. The wakefield excitation process leads to the development of non-uniform background plasma densities, in particular the presence of peaks within the transverse density profile. Regions of high electron density appear at two positions in the wakefield; (i) close to the axis of propagation due to a combination of the intense electron drive and witness beams and the high density of plasma electrons as they return to the axis at the rear of the wakefield cavities; and (ii) at the maximum radius of the wakefield cavity as streaming plasma electrons collect at its boundary. These structures form over the first cavity but remain for multiple wakefield periods before phase-mixing [97] of plasma electrons causes dissipation of the wakefield structure as measured in Chapter 7. These regions of high electron density attract plasma ions and, with their effects, and those of the corresponding transverse fields, averaged over multiple plasma periods as the wakefield structure is maintained, the ions redistribute with density peaks appearing close to the axis and at the maximum radius of the plasma wakefield cavity. In order to maintain quasi-neutrality of the plasma, plasma electrons move with the ions and non-uniform transverse density structures are produced. The dissipation of these peaks is driven by the pressure gradients induced in the plasma as a result of the combination of the local density gradients and the energy imparted to the plasma by the beams. This energy is mostly transferred to plasma electrons involved in the generation of the wakefield structure, but is gradually distributed over the plasma bulk via their resultant collisions with other plasma constituents. Therefore, over extended timescales, the initial sharp ion density spikes diffuse and the evolving transverse density profile can be approximated by a series of local parabolic channels, akin to those generated on short timescales within discharge-ionised plasmas contained in capillaries [63, 65] although over much smaller spatial scales in this case.

The effect of the wakefield on the transverse plasma density profile can be seen by considering an analogous process where a region of increased temperature near the axis is imposed on a background plasma of constant density. The evolution of the plasma can be considered via the fluid approximation, by solving the continuity and heat equations in one dimension:

$$\frac{\partial n}{\partial t} + \frac{1}{r}\frac{\partial}{\partial r}(r\Gamma_i) = 0, \tag{8.2.1}$$

$$\frac{3}{2}k_B n\frac{\partial T}{\partial t} = \frac{1}{r}\frac{\partial}{\partial r}\left(r\kappa\frac{\partial T}{\partial r}\right) \tag{8.2.2}$$

where Γ_i represents the ion flux density, formed of the sum of ambipolar diffusion and fluxes due to the presence of density and temperature gradients [228]:

$$\Gamma_i = -\mathcal{M}_i n\left(\frac{D_i - D_e}{\mathcal{M}_e + \mathcal{M}_i}\frac{1}{n}\frac{\partial n}{\partial r}\right) + D_i\frac{\partial n}{\partial r} + \frac{nD_i}{T}\frac{\partial T}{\partial r}, \tag{8.2.3}$$

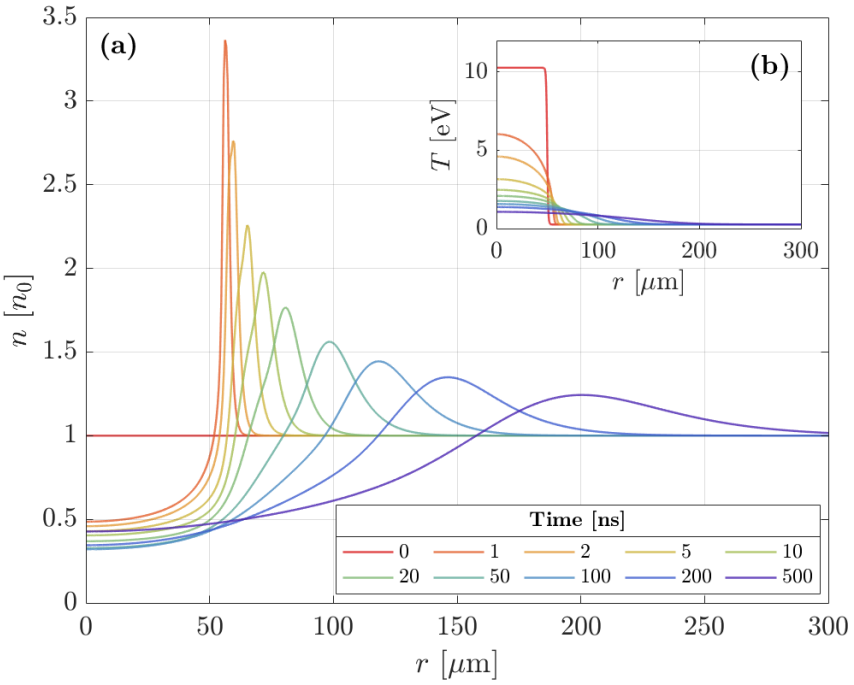

Figure 8.2.1: Evolution of the transverse (a) plasma density profile and (b) temperature when a 50 µm region of plasma on-axis is heated to 10 eV.

where $\mathcal{M}$ and D are the mobility and diffusion coefficients of ions and electrons within the plasma respectively. An example of such evolution, solved using a fourth-order explicit Runge-Kutta routine within the MATLAB pdepe package [253], is displayed in Figure 8.2.1. In this example, an initial temperature of 10 eV is imposed on a region of width 50 µm near the axis, similar to the local effect of the wakefield excitation that leads to heating of the plasma within its vicinity. The consequent evolution of the transverse plasma density profile and temperature are shown in Figure 8.2.1(a) and (b) respectively. The increased temperature near the axis generates a local pressure gradient and, in an attempt to return to uniform pressure conditions, the plasma density on-axis rapidly drops while a density peak develops at the edge of the region of increased temperature. As the plasma continues to evolve, the off-axis peak dissipates and propagates radially outwards while the on-axis density reaches a minimum before beginning to return to the initial condition as the heat continues to dissipate through the plasma bulk.

The timescales over which this evolution occurs in this case are approximately two orders of magnitude larger than is expected in the case of perturbation by a plasma wakefield. This is because in reality the wakefield has a much larger perturbative effect on the plasma and this considerably-simplified model does not consider the effect of the time-averaged transverse wakefield on the initial motion of plasma ions. However, it demonstrates a prediction for the types of transverse density profiles that could be induced on extended timescales by wakefield excitation. After the initial evolution ($t \geqslant 10$ ns in this case), the transverse density profile within the vicinity of the axis can be reasonably well-approximated by a series of

evolving parabolic profiles. As such, the effect of such parabolic density profiles on the PWFA process is considered within the following section.

8.3 PWFA ON PARABOLIC TRANSVERSE DENSITY PROFILES

The presence of parabolic transverse plasma density profiles significantly modifies the PWFA process. Its effect can be derived by considering the blowout regime of plasma accelerators where the perturbation driven by the bunch is sufficiently large to cause complete cavitation of the wake with all background plasma electrons expelled, leaving only an ion background. This section builds on the non-linear wakefield theory introduced in Chapter 2, comparing the motion of electrons both within the bunch and within the plasma for a uniform ion density profile with that derived here for a parabolic ion channel. The equation of motion for the wakefield cavity boundary in a parabolic ion channel is derived, enabling calculation of the longitudinal and transverse fields within the wakefield following the formalism presented in Refs. [86, 87].

8.3.1 *Electron bunch response*

8.3.1.1 *Particle motion in a blown-out uniform ion channel*

The effect on electrons within the drive and witness bunches can be calculated from the Lagrangian describing the motion of a single particle in a potential, ϕ:

$$\mathcal{L} = -m_e c^2 \sqrt{1-\beta^2} + e\phi. \tag{8.3.1}$$

The form of the potential can be found using Poisson's equation, which for a cylindrically symmetric system is given by

$$\frac{1}{r}\frac{\partial}{\partial r}\left(r\frac{\partial \phi}{\partial r}\right) = \frac{-e\,(n_i - n_e)}{\varepsilon_0}. \tag{8.3.2}$$

In the case of complete blowout, the electron density $n_e = 0$, giving

$$\phi = -\frac{e n_i r^2}{4\varepsilon_0}. \tag{8.3.3}$$

The radial equation of motion can be found from the Euler-Lagrange equation,

$$\frac{d}{dt}\left(\gamma m_e \dot{r}\right) = -\frac{e^2 n_i}{2\varepsilon_0} r \tag{8.3.4}$$

and assuming no accelerating field ($\dot{\gamma} = 0$) for a particle with a small radial velocity when compared to its longitudinal velocity (i.e. $v_r \ll v_z \sim c$), giving

$$\ddot{r} = -\omega_\beta^2 r \tag{8.3.5}$$

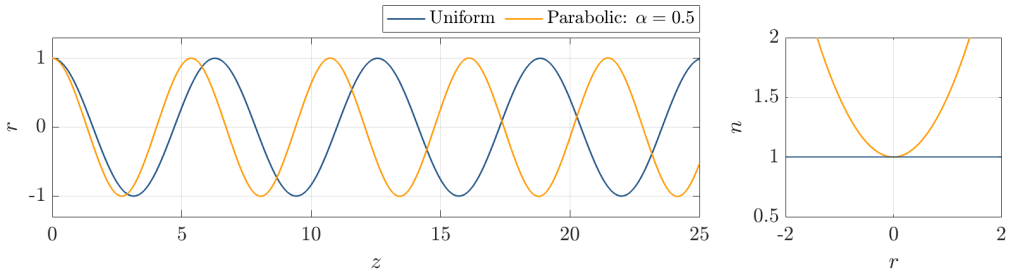

Figure 8.3.1: Example betatron trajectories in uniform (blue) and parabolic ($\alpha = 0.5$, orange) transverse density profiles for identical initial conditions: $r_0 = 0.5$, $\dot{r}_0 = 0$. The corresponding transverse density profiles are shown in the right panel.

where $\omega_\beta = \omega_p/\sqrt{2\gamma_{z0}}$ is known as *betatron* frequency and represents the frequency at which electrons undergo transverse oscillations within the large amplitude transverse wakefields. This has solution

$$r = r_\beta \cos(\omega_\beta t) \tag{8.3.6}$$

for initial conditions $r(t = 0) = r_\beta$, $\dot{r}(t = 0) = 0$. The radial oscillations of the particle modify the longitudinal equation of motion. The previous condition $\dot{\gamma} = 0$ implies that the square of the electron velocity is constant, i.e.

$$\beta^2 = \beta_r^2 + \beta_z^2 \tag{8.3.7}$$

thus allowing calculation of the effect of radial oscillations on the longitudinal velocity of the particle. Using the expansion $\beta_z = \sqrt{\beta^2 - \beta_r^2} \approx \beta\left(1 - \frac{\beta_r^2}{2}\right)$, the z-component of the velocity can be found to be

$$\beta_z \approx \beta\left(1 - \frac{r_\beta^2\omega_\beta^2}{4c^2}\right) + \beta\frac{r_\beta^2\omega_\beta^2}{4c^2}\cos(2\omega_\beta t). \tag{8.3.8}$$

Integrating gives

$$\frac{z}{c} \approx \frac{z_0}{c} + \beta\left(1 - \frac{r_\beta^2\omega_\beta}{4c^2}\right)t + \beta\frac{r_\beta^2\omega_\beta^2}{8c^2}\sin(2\omega_\beta t). \tag{8.3.9}$$

8.3.1.2 *Particle motion in a blown-out parabolic ion channel*

If instead of a uniformly-distributed ion background, the channel has a parabolic shape due to, for example, ion motion created by the wakefield from a preceding drive bunch, the form of the potential is modified:

$$\frac{1}{r}\frac{\partial}{\partial r}\left(r\frac{\partial\phi}{\partial r}\right) = \frac{-e\,n_{i0}(1 + \alpha r^2)}{\varepsilon_0}. \tag{8.3.10}$$

where α is a scaling factor that describes the shape of the parabolic channel, $n_i(r) = n_{i0}(1 + \alpha r^2)$. In this case, the potential becomes

$$\phi = -\frac{e n_{i0} r^2}{4\varepsilon_0}\left(1 + \frac{\alpha r^2}{4}\right). \tag{8.3.11}$$

Following the procedure for a uniform channel, the equation of motion of an electron in a parabolic ion channel is given by

$$\ddot{r} = -\frac{e^2 n_{i0}}{2 m_e \varepsilon_0 \gamma} r\left(1 + \frac{\alpha r^2}{2}\right) \tag{8.3.12}$$

which can be solved numerically. Using the initial conditions $r_0 = 0.5$, $\dot{r}_0 = 0$ for a parabolic channel with $\alpha = 0.5$, the difference in trajectory between a uniform and parabolic channel is shown in Fig. 8.3.1. The increasing ion density off-axis modifies the local focusing force and causes an increase in the resulting betatron frequency ω_β.

Constant longitudinal field

It is possible to also include the effect from a constant longitudinal field by modifying the Lagrangian to include an extra term in ϕ such that

$$\frac{\partial \phi}{\partial z} = -E_z. \tag{8.3.13}$$

The equations of motion are again found using the Euler-Lagrange equations and are

$$\frac{d}{dt}(\gamma m_e \dot{r}) = -\frac{e^2 n_{i0}}{2\varepsilon_0} r\left(1 + \frac{\alpha r^2}{2}\right) \tag{8.3.14}$$

$$\frac{d}{dt}(\gamma m_e \dot{z}) = -eE_z. \tag{8.3.15}$$

With the introduction of the longitudinal field, the previous assumption $\dot{\gamma} = 0$ is no longer valid. Therefore Eq. (8.3.14) becomes

$$\dot{\gamma}\beta_r + \gamma\dot{\beta_r} = -\frac{e^2 n_{i0}}{2\varepsilon_0 m_e c} r\left(1 + \frac{\alpha r^2}{2}\right). \tag{8.3.16}$$

Equation (8.3.15) can be solved trivially to give

$$\gamma\beta_z = -\frac{e}{m_e c} E_z t + \gamma_0 \beta_0. \tag{8.3.17}$$

The electron is considered to be highly-relativistic ($\beta_z \approx 1$) and hence

$$\dot{\gamma} = -\frac{eE_z}{m_e c} \tag{8.3.18}$$

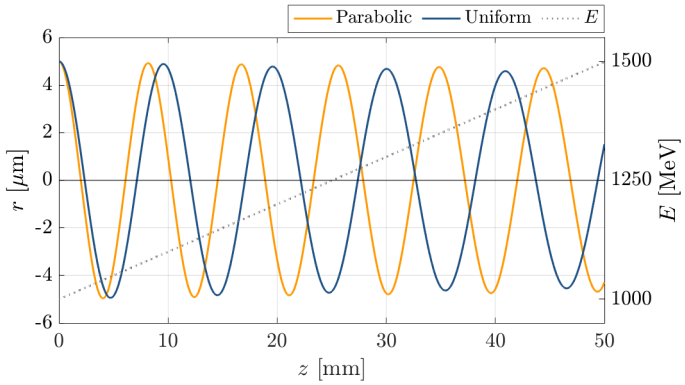

Figure 8.3.2: Example betatron trajectories in uniform (blue) and parabolic ($\alpha = 0.5$, orange) transverse density profiles for identical initial conditions: $r_0 = 5\,\mu m$, $\dot{r}_0 = 0$ in a background density of $n_{i0} = 5 \times 10^{16}\,cm^{-3}$ with a constant longitudinal field of $10\,GVm^{-1}$ ($\sim 0.46E_0$) applied. The corresponding energy increase is plotted in grey.

which along with the relativistic assumption can be substituted into Eq. (8.3.16) to give

$$-\frac{eE_z}{m_e c}\beta_r + \left(\gamma_0\beta_0 - \frac{eE_z}{m_e c}t\right)\dot{\beta_r} = -\frac{e^2 n_{i0}}{2\varepsilon_0 m_e c} r\left(1 + \frac{\alpha r^2}{2}\right). \tag{8.3.19}$$

Consequently, the radial equation of motion of an electron in a blown out parabolic ion column described by the form $n_i(r) = n_{i0}(1 + \alpha r^2)$ under a constant longitudinal field (E_z) is

$$(\gamma_0\beta_0 m_e c - eE_z t)\,\ddot{r} - eE_z\dot{r} + \frac{e^2 c n_{i0}}{2\varepsilon_0} r\left(1 + \frac{\alpha r^2}{2}\right) = 0. \tag{8.3.20}$$

This equation can be solved numerically and example trajectories for a 1 GeV electron in an accelerating field of $10\,GVm^{-1}$ in a plasma of density $n_0 = 5 \times 10^{16}\,cm^{-3}$ and length 50 mm are shown in Figure 8.3.2. For this calculation $r_0 = 5\,\mu m$ and the parabolic channel is parameterised by $\alpha = 0.5$. As the energy of the electron grows due to the longitudinal field, the frequency of the oscillation decreases while its radius also decreases.

8.3.2 *Plasma electron response*

The presence of a parabolic transverse density profile has an effect on not only the betatron oscillations of electrons within the drive and witness beams but also on the formation of the wakefield itself. This can be seen by considering the modification to the equation of motion for plasma electrons, derived in Ref. [86, 87], and previously given in Eq. (2.2.19). As a reminder, this equation was derived using the Dawson sheet model [97] where the plasma can be treated as being made up of a set of concentric, infinitely-long cylindrical sheets of electrons and ions. It calculates plasma electron trajectories during perturbation by the drive beam under the assumption that the total force on any individual electron ring is due solely to the electrostatic force from the total charge within the ring; i.e. the ion charge,

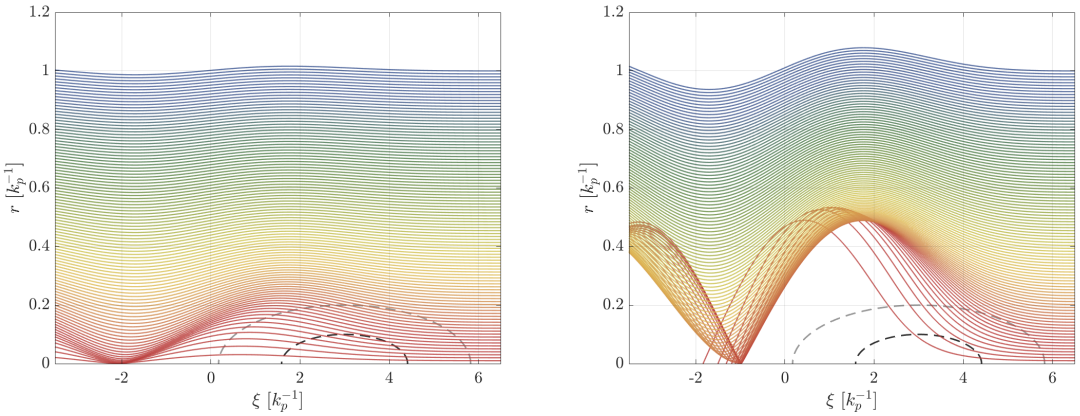

Figure 8.3.3: Solution of Eq. (2.2.22) over $0 \leqslant r \leqslant 1k_p^{-1}$ for a bi-Gaussian drive beam with $\sigma_r = 0.1$, $\sigma_z = \sqrt{2}$ and (a) $n_{b0} = 1$ and (b) $n_{b0} = 5$. The colour of a given electron trajectory represents the initial radial position of each electron and is the same for both figures. The 1 and 2σ contours of the drive beam are plotted using black and grey dashed lines respectively.

plasma electron charge and drive beam charge. It is only valid under the condition that the trajectory of a given ring never crosses that of a ring outside of it (i.e. that trajectory crossing is not observed), and is limited to the non-relativistic limit where $v_z \ll v_\perp \ll c$, and $v_\perp \approx dr/d\xi$. The leading term in Eq. (2.2.19), $-r/2$, comes from the background uniform ion density and as such can be modified for the case of a parabolic density profile. The second term, $c(r_0, r, \xi)/r$, represents the two-dimensional cylindrical electrostatic force from the total electron charge within the ring at r and is calculated via Eq. (2.2.20) under the assumption of no trajectory crossing.

In Chapter 2 it was shown that the equation of motion describing the plasma electron response to a bi-Gaussian drive bunch is given by Eq. (2.2.22). This equation of motion can be solved numerically to give the plasma electron response as shown in Fig. 8.3.3. In this figure, the response of plasma electrons to a bi-Gaussian drive beam with $\sigma_r = 0.1$, $\sigma_z = \sqrt{2}$ and (a) $n_{b0} = 1$ and (b) $n_{b0} = 5$ is solved assuming a uniform background ion density. In (a), formation of the wakefield structure due to the perturbation by the drive beam can be observed, with trajectories converging towards the rear of the cavity. In (b), the increased drive beam density perturbs the plasma more strongly and evacuation of the wakefield structure is observed towards the rear of the drive bunch. In this case, significant trajectory crossing is observed at the extremities of the blown-out cavity thanks to the strong perturbative effect from the high-density drive bunch and hence the solution is not strictly valid. However, it is still somewhat representative of the wakefield cavity produced by the bunch and can be used to qualitatively compare to the case of a parabolic ion density profile.

Extending the equation of motion, Eq. (2.2.22), to account for parabolic density profiles can be simply implemented by considering the modification of the focusing force on a plasma

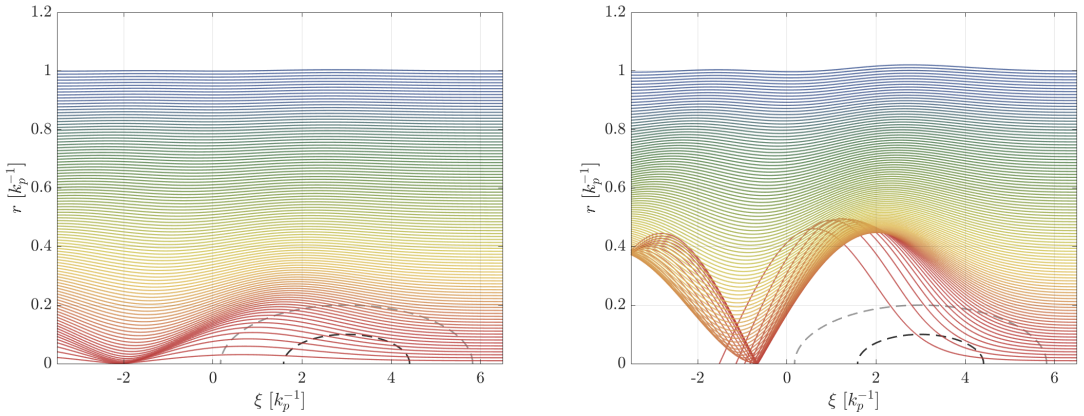

(a) (b)

Figure 8.3.4: Solution of Eq. (2.2.22) with modifications for a parabolic ion density profile outlined in Eqs. (8.3.21) & (8.3.22) for $\alpha = 1.0$.

electron due to the parabolic ion channel. From the modification to the potential, Eq. (8.3.10), this becomes

$$\vec{F}_\perp \propto -\frac{1}{2} r \left(1 + \frac{\alpha}{2} r^2\right) \tag{8.3.21}$$

and a similar modification must be made to $c(r_0, r, \xi)$ to account for the change in transverse density profile

$$c(r_0, r, \xi) = \frac{1}{2} r_0^2 \left(1 + \frac{\alpha}{2} r_0^2\right) + \int_0^r r' \, n_b(r', \xi) \, dr'. \tag{8.3.22}$$

The equation of motion, Eq. (2.2.22), is otherwise identical and again can easily be solved numerically as shown in Fig. 8.3.4 for identical conditions as in Fig. 8.3.3 but with a parabolic density profile parameterised by $\alpha = 1.0$. In this case, the increased focusing at larger radii reduces the transverse extent of trajectories of electrons strongly perturbed by the drive beam. This additionally reduces the longitudinal extent of the wakefield cavity and as such significantly modifies the wakefield amplitude. Note also that the introduction of a non-linear term into the form of the transverse wakefield in Eq. (8.3.21) will prohibit emittance preservation of a witness bunch accelerated within the wakefield, reducing the bunch quality and highlighting the need for uniform transverse density profiles to achieve high-quality acceleration. This therefore indicates that any ion motion perturbation must dissipate before subsequent acceleration events can occur to demonstrate repeatable acceleration at high-repetition-rates.

8.3.3 *Wakefield cavity shape*

The previous derivation can be extended to non-linear theory and applied to the perturbation from high-density drive bunches through use of the quasi-static approximation, as was

discussed in Chapter 2. There, the solution for the boundary of the plasma bubble (r_b) on a uniform ion background was found from solving the relativistic equation of motion for a plasma electron, Eq. (2.2.37), for the wakefield pseudo-potential described by Eq. (2.2.33). This reduced to solving Eq. (2.2.38) with coefficients given by Eqs. (2.2.39) – (2.2.41). A similar derivation can again be made but for the case of the parabolic transverse density profile given by $n_i(r) = n_{i0}(1 + \alpha r^2)$. In this case, the pseudo-potential within the cavity has a different form:

$$\begin{aligned}\psi(0,\xi) &= \int_0^{r_b} \frac{dr}{r} \int_0^{r} \rho_{ion} r' dr' \\ &= \int_0^{r_b} \frac{dr}{r} \int_0^{r} (1 + \alpha r'^2) r' dr' = \frac{r_b^2}{4}\left(1 + \frac{\alpha r_b^2}{4}\right),\end{aligned} \tag{8.3.23}$$

such that at the boundary of the wakefield cavity, the pseudo-potential is now given by

$$\psi = \frac{r_b^2}{4}\beta\left(1 + \frac{\alpha r_b^2}{4}\right), \tag{8.3.24}$$

where β is again given by Eq. (2.2.34). The above equation can then be inserted into Eq. (2.2.37) to calculate the equation of motion for a plasma electron at the boundary of the cavity. In this case, the LHS of the equation is given by

$$\begin{aligned}\frac{d}{d\xi}\left[(1+\psi)\frac{d}{d\xi}r_b\right] &= \frac{d\psi}{d\xi}\frac{dr_b}{d\xi} + (1+\psi)\frac{d^2 r_b}{d\xi^2} \\ &= \left[1 + \frac{r_b^2}{4}\beta\left(1 + \frac{\alpha r_b^2}{4}\right)\right]\frac{d^2 r_b}{d\xi^2} \\ &+ \left[\frac{r_b}{2}\beta\left(1 + \frac{\alpha r_b^2}{4}\right) + \frac{r_b^2}{4}\left(1 + \frac{\alpha r_b^2}{4}\right)\frac{d\beta}{dr_b} + \frac{r_b^3}{8}\alpha\beta\right]\left(\frac{dr_b}{d\xi}\right)^2,\end{aligned} \tag{8.3.25}$$

while the major change to the RHS is introduced by the updated form of $\sigma = -(1/2)d\psi_0/d\xi$, such that,

$$\begin{aligned}\frac{d\sigma}{d\xi} = -\frac{1}{2}\Bigg(&\left[\frac{r_b}{2}\left(1 + \beta\left(1 + \frac{\alpha r_b^2}{4}\right)\right) + \frac{r_b^2}{4}\left(1 + \frac{\alpha r_b^2}{4}\right)\frac{d\beta}{dr_b} + \frac{r_b^3}{8}\alpha\beta\right]\frac{d^2 r_b}{d\xi^2} \\ &+ \left[\frac{1}{2}\left(1 + \beta\left(1 + \frac{\alpha r_b^2}{4}\right)\right) + r_b\left(1 + \frac{\alpha r_b^2}{4}\right)\frac{d\beta}{dr_b} + \frac{5r_b^2}{8}\alpha\beta + \frac{r_b^3}{8}\alpha\frac{d\beta}{dr_b}\right]\left(\frac{dr_b}{d\xi}\right)^2 \\ &+ \frac{r_b^3}{8}\alpha\frac{dr_b}{d\xi}\Bigg).\end{aligned} \tag{8.3.26}$$

Inserting these terms into Eq. (2.2.37) gives a similar equation to that found for the uniform case [Eq. (2.2.38)] but now includes an additional term that arises due to the form of the parabolic profile:

$$A_p(r_b)\frac{d^2 r_b}{d\xi^2} + B_p(r_b) r_b\left(\frac{dr_b}{d\xi}\right)^2 + \frac{r_b^4}{16}\alpha\frac{dr_b}{d\xi} + C_p(r_b) r_b = \frac{\lambda(\xi)}{r_b}, \tag{8.3.27}$$

while the coefficients are also modified:

$$A_p(r_b) = 1 + \left[\frac{1}{4} + \frac{\beta}{2}\left(1 + \frac{\alpha r_b^2}{4}\right) + \frac{1}{8}\left(1 + \frac{\alpha r_b^2}{4}\right) r_b \frac{d\beta}{dr_b} + \frac{r_b^2}{16}\alpha\beta\right] r_b^2, \tag{8.3.28}$$

$$B_p(r_b) = \frac{1}{2} + \frac{3}{4}\beta + \frac{3}{4} r_b \frac{d\beta}{dr_b} + \alpha\left(\frac{9 r_b^2}{16}\beta + \frac{r_b^3}{16}\beta + \frac{r_b^3}{4}\frac{d\beta}{dr_b}\right) + \mathcal{O}\left(\frac{d^2\beta}{dr_b^2}\right), \tag{8.3.29}$$

$$C_p(r_b) = \frac{1}{4}\left[1 + \frac{1}{\left(1 + \frac{r_b^2}{4}\beta\left(1 + \frac{\alpha r_b^2}{4}\right)\right)^2}\right]. \tag{8.3.30}$$

Note that Eq. (8.3.27), and all coefficients within, return to the uniform case given in Eq. (2.2.38) when $\alpha = 0$ as expected. The longitudinal field within the wakefield cavity can again be found via the relation

$$E_z(r, \xi) = \frac{d\psi(r = 0, \xi)}{d\xi}, \tag{8.3.31}$$

and hence is given by

$$E_z = \frac{1}{2} r_b \frac{dr_b}{d\xi} \tag{8.3.32}$$

for a uniform ion background and

$$E_z = \frac{1}{2} r_b \frac{dr_b}{d\xi}\left(1 + \frac{\alpha r_b^2}{4}\right) \tag{8.3.33}$$

in the case of a parabolic density profile. Therefore, once Eq. (8.3.27) is solved, the solution for r_b can be inserted into Eq. (8.3.33) to determine the form of the longitudinal field for a wakefield driven within a parabolic radial ion density profile.

In the case of an ultra-relativistic blowout, $n_b \gg n_e$, the maximum radius of the wakefield cavity exceeds the plasma skin-depth, i.e. $r_m \gg k_p^{-1}$, and the width of the cavity far exceeds that of its surrounding sheath. In this case, $\beta \ll 1$ and the coefficients in the equations of motion can be vastly simplified. For both a uniform and parabolic transverse ion density profile they reduce to

$$A(r_b) = A_p(r_b) = \frac{1}{4} r_b^2, \tag{8.3.34}$$

$$B(r_b) = B_p(r_b) = \frac{1}{2}, \tag{8.3.35}$$

$$C(r_b) = C_p(r_b) = \frac{1}{4}, \tag{8.3.36}$$

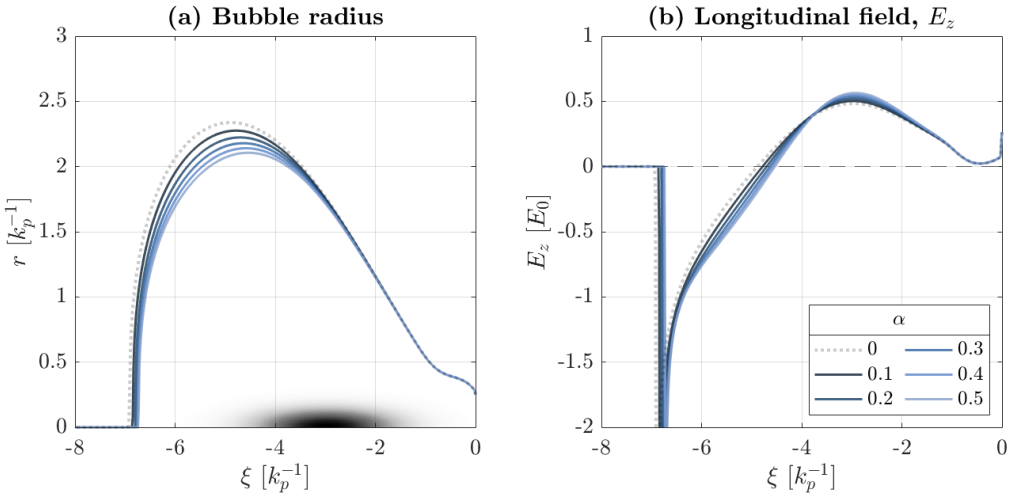

Figure 8.3.5: (a) Numerical solution to Eq. (8.3.37) (grey, dashed) and Eq. (8.3.38) (coloured) for increasing values of α demonstrating narrowing and shortening of the wakefield cavity. (b) The solutions from (a) are inserted into Eq. (8.3.31) to calculate the longitudinal field within the cavity.

leading to the familiar equation of motion [86, 87]

$$r_b \frac{d^2 r_b}{d\xi^2} + 2\left[\frac{dr_b}{d\xi}\right]^2 + 1 = \frac{4\lambda(\xi)}{r_b^2}, \tag{8.3.37}$$

for a uniform background ion density and

$$r_b \frac{d^2 r_b}{d\xi^2} + 2\left[\frac{dr_b}{d\xi}\right]^2 + \frac{r_b^3}{4}\alpha\frac{dr_b}{d\xi} + 1 = \frac{4\lambda(\xi)}{r_b^2}, \tag{8.3.38}$$

for a parabolic ion density profile. Equations (8.3.37) and (8.3.38) are solved in Figure 8.3.5(a) for increasing values of $\alpha = 0 - 0.5$, demonstrating narrowing and shortening of the wakefield cavity with increasing curvature as previously observed with regards to the discussion of plasma electron trajectories in Section 8.3.2. The calculation assumes a bi-Gaussian electron bunch of RMS length $k_p\sigma_z = 1$, width $k_p\sigma_r = 0.1$, and charge 3.8 nC centred at $\xi = -3\,k_p^{-1}$ on a background plasma density of $n_0 = 10^{16}\,\mathrm{cm}^{-3}$.

The solutions to the equations of motion are then inserted into Eq. (8.3.31) to calculate the longitudinal wakefield within the cavity and are shown in Figure 8.3.5(b). The addition of the parabolic profile shifts the position of the zero-crossing of the longitudinal field, defined by the position of maximum extent of the wakefield cavity i.e. $dr_b/d\xi = 0$, forwards within the wakefield and correspondingly increases its amplitude in the shortened decelerating region while increasing its gradient, $dE_z/d\xi$, beyond this and introducing additional curvature to the form of the longitudinal field towards the rear of the cavity. As discussed before, the transverse wakefield within the parabolic ion cavity also includes a non-linear term and is given by Eq. (8.3.21). These two modifications to the form of the wakefield within the cavity have the potential to alter both the ability of the witness bunch to optimally load

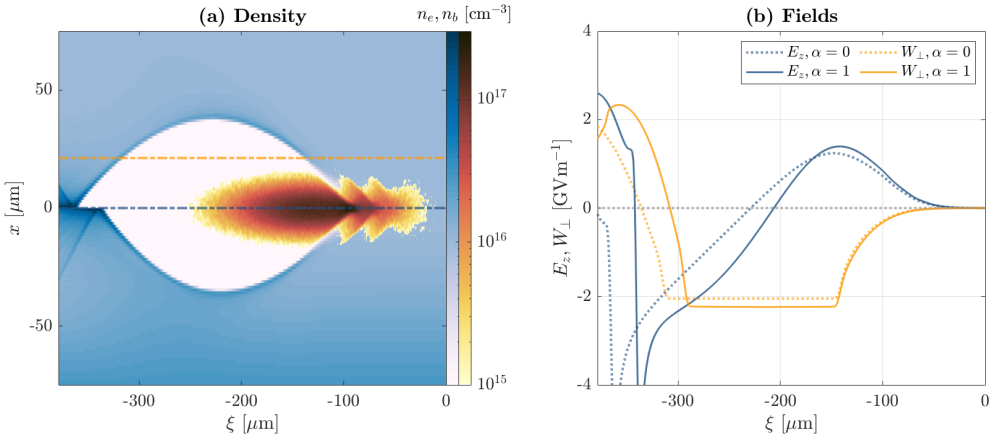

Figure 8.3.6: Effect of the introduction of a parabolic transverse density profile parameterised by $\alpha = 1$ on the wakefield driven by a bi-Gaussian bunch of charge 310 pC and RMS length 95 fs ($I_{pk} \sim 1$ kA) in a plasma of density $n_0 = 10^{16}$ cm^{-3} according to PIC simulations performed using HiPACE. (a) Plasma electron (blue) and bunch (orange) density maps. (b) On-axis longitudinal field (blue) and transverse field (orange) measured at $0.4\,k_p^{-1}$ for $\alpha = 0$ (dashed) and $\alpha = 1$ (solid). These positions are represented by the dotted-dashed lines in (a) for reference.

the wakefield and the evolution of its transverse emittance during propagation through the plasma.

The theoretically-predicted effects of introducing parabolic transverse density profiles are confirmed via PIC simulations of the wakefield driven by a bi-Gaussian electron bunch with parameters similar to those expected in the FLASHForward experiment as shown in Figure 8.3.6. The electron bunch has a charge of 310 pC with an RMS bunch length of 95 fs, giving a peak current of approximately 1 kA. The bunch is focused at the plasma entrance ($n_0 = 10^{16}$ cm^{-3}) with a β-function of 30 mm in x and y assuming a transverse normalised emittance of $\varepsilon_N = 2$ mm·mrad. Figure 8.3.6(a) shows the plasma electron (blue) and bunch (orange) density maps, demonstrating that the introduction a parabolic density profile parameterised by $\alpha = 1$ (lower half) has the expected effect on the wakefield cavity shape, significantly reducing both its transverse and longitudinal extent. Figure 8.3.6(b) shows the effect on the longitudinal on-axis field (blue) and the transverse field calculated at $r = 0.4\,k_p^{-1} \approx 20\,\mu$m (orange) with the case of the uniform density profile plotted using dashed lines and that of the parabolic density profile plotted using solid lines. Similarly to Figure 8.3.5(b), the peak decelerating field grows in amplitude for $\alpha = 1$ while the position of the zero-crossing is moved forwards within the cavity and the form of the longitudinal field develops a curvature towards the rear of the cavity that is not seen in the case of the uniform density profile. The amplitude of the transverse field within the cavity also increases by approximately 8% from 2.05 GVm^{-1} ($\alpha = 0$) to 2.22 GVm^{-1} ($\alpha = 1$), consistent with the expected scaling of $1 + (\alpha/2)r^2 = 1.08$ at $r = 0.4\,k_p^{-1}$ predicted in Eq. (8.3.21).

This section has demonstrated that significant modifications to the wakefield excitation process are expected when a parabolic radial density profile is present within the plasma.

The beam-based measurement technique introduced in Chapter 5 therefore seems an ideal diagnostic, as modifications to the longitudinal and transverse fields within the cavity will present themselves within measurements of the energy and transverse spectra of the probe bunch after its interaction with the plasma. Such measurements are the subject of the following sections within this chapter, enabling the study of wakefield-induced ion motion within plasma over nanosecond timescales and beyond.

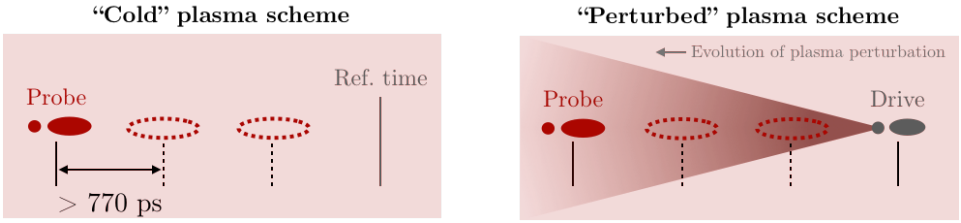

Figure 8.4.1: Cartoon representing the beam-based perturbed plasma measurement scheme. In the "cold" plasma scheme, the probe bunch pair samples the evolution of an unperturbed plasma with respect to a reference time. In the "perturbed" plasma scheme the probe bunch pair samples the plasma state following the passage of a high peak current drive bunch pair at the reference time. Differences in both the energy and transverse spectrum of the probe bunch between the "cold" and "perturbed" plasma states allow determination of the extended timescale effect of the PWFA process on the plasma.

8.4 EXTENSION OF THE BEAM-BASED PLASMA MEASUREMENT TECHNIQUE

The beam-based *perturbed* plasma measurement technique is a simple extension to the technique outlined in Chapter 5 and uses two drive-witness bunch pairs with a small, variable separation to probe plasma evolution on extended timescales. As discussed in Chapter 4, multiple photocathode lasers are available within the FLASH facility, allowing two electron bunches with similar parameters to be produced at the cathode and accelerated and transported to the FLASHForward experimental area with temporal separations at the frequency of the super-conducting RF cavities within the FLASH beamline: multiples of 1.3 GHz (769 ps) up to the length of the RF flat top, approximately 800 µs.

The measurement principle, illustrated in Figure 8.4.1, represents a simple drive-probe system, where the first ("drive" or "leading") bunch pair drives a large amplitude wakefield (GVm^{-1}) and the second ("probe") bunch pair samples the plasma state at variable times behind the first. By measuring both the energy and transverse spectra of the probe bunch pair after interaction with the plasma both with ("perturbed"), and without ("cold" or "unperturbed"), the presence of the leading bunch pair, the effect of the leading bunch pair on the plasma can be ascertained.

8.4.1 *Spectrometer signal subtraction*

For small temporal separations between the leading and probe bunch pairs, the spectrometer signals overlap due to the decay rate of the scintillating material, approximately 100 ns for the LYSO crystal used in these experiments. In order to retrieve the signal of the probe bunch pair for bunch separations where the signals overlap, a signal prediction and subtraction technique was developed. For each measurement, three different datasets were taken:

1. "Subtraction" dataset: a high-statistics (typically $>$ 100 shots) dataset, recording *only* the leading bunch pair signal on the spectrometer following its interaction with the plasma.

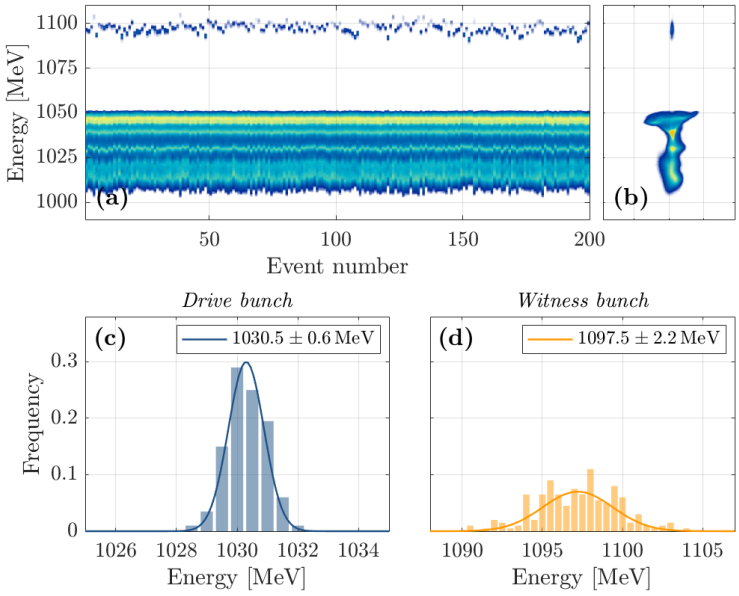

Figure 8.4.2: Example subtraction dataset for the interaction between a leading bunch pair and argon plasma. (a) Waterfall image of the energy spectrum of all 200 events in the dataset. (b) Average spectrometer image. (c) Histogram of the drive bunch mean energy following plasma interaction. (d) Histogram of the witness bunch mean energy following plasma interaction.

2. "Perturbed" dataset: a scan of the temporal separation between the two bunch pairs over the desired range, recording *both* the leading and probe bunch pair signals on the spectrometer. Typically 50 shots per bunch separation.

3. "Cold" dataset: a repeat of the "perturbed" dataset scan but with *only* the probe bunch pair present, recording its signal on the spectrometer.

The subtraction dataset is then used in post-analysis to predict the signal from the leading bunch pair in the perturbed dataset. This allows subtraction of its predicted signal to leave only the probe bunch signal in the case of a perturbed plasma. An example subtraction dataset is shown in Figure 8.4.2 for the interaction between argon plasma at a backing pressure of 10.6 mbar and the leading bunch pair used in ion motion studies presented in Sections 8.6.1 and 8.6.2.

A subtraction method such as this can only be used if the interaction between the bunch and plasma is stable on a shot-to-shot basis, typically an issue for plasma-based accelerators due to the complexity of the non-linear plasma interaction and its sensitivity to the bunch profile. However, Figure 8.4.2 is an example of the stability of the plasma acceleration process achieved as a result of the stability of both the electron bunch profile provided by the FLASH accelerator, and the plasma profile produced by the high-voltage discharge. The mean energy of the drive bunch following interaction with the plasma was measured to be $\mu_E = 1030.5 \pm 0.6$ MeV, an energy-jitter-to-loss ratio of approximately 3% across 200 events. The witness bunch mean energy following plasma interaction was measured to be $\mu_E = 1097.5 \pm 2.2$ MeV,

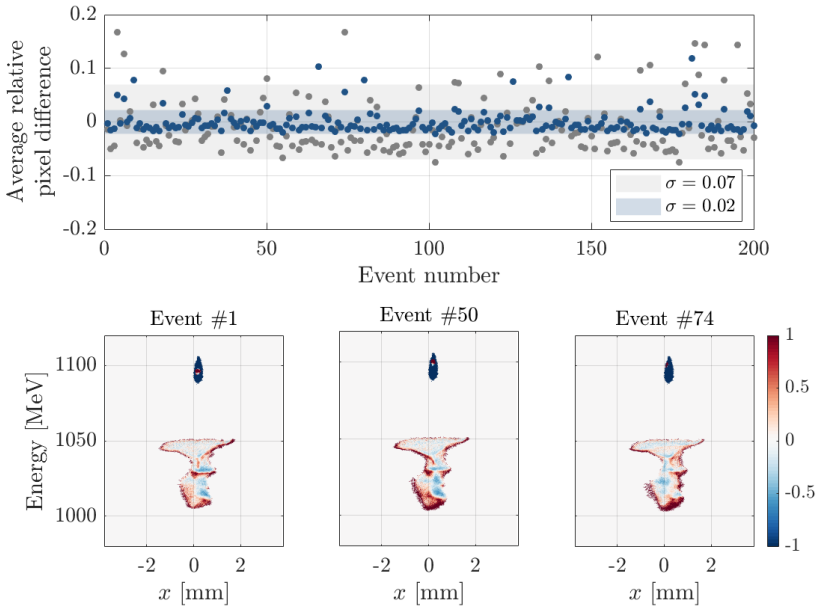

Figure 8.4.3: Similarity of the predicted bunch profile and that of spectrometer images from individual events. Upper: average relative pixel difference between the predicted and true bunch profile for the dataset shown in Figure 8.4.2. Grey data points represent the average deviation over the entire bunch profile, whereas blue data points correspond to the decelerated drive bunch only. The grey and blue shaded regions represent the standard deviation across the dataset for the entire bunch and drive regions respectively. Lower: three example bunch profiles showing the difference between the predicted and true spectrometer signal calculated according to Eq. (8.4.1). The average relative pixel difference for the three events are: −0.014 (left), 0.081 (centre) and 0.167 (right).

giving a slightly larger energy-jitter-to-gain ratio of approximately 5%. The reason for the increased jitter of the witness bunch is discussed in detail later.

Given the few-percent repeatability of the plasma interaction, a bunch signal subtraction technique was developed that predicted the signal that would be associated with the leading bunch pair. This was used in events where both the leading and probe bunch pair signals overlap on the spectrometer image to allow reconstruction of only the probe bunch signal. The predicted signal represented an average bunch profile calculated via a charge-weighted combination of all events in the subtraction dataset, an example of which is shown in Figure 8.4.2(b). In order to estimate the uncertainty introduced by the use of this "predicted bunch" subtraction technique, the predicted spectrometer signal profile was compared to the true bunch profiles from the subtraction dataset. For each event, the average bunch profile is rescaled according to its measured bunch charge and compared to the true spectrometer signal on a pixel-by-pixel basis:

$$\langle \delta n \rangle = \sum_{i,j} \left(\frac{n_{ij} - \tilde{n}_{ij}}{\tilde{n}_{ij}} \right) \Big/ N \,, \tag{8.4.1}$$

where i and j represent the rows and columns of the spectrometer image, n_{ij} is the *true* pixel count, $\tilde{n}_{ij}$ is the *predicted* pixel count from the rescaled average bunch image and N is the number of pixels in the signal region identified using the background subtraction routine described in Section 5.3.2. An example of the result of this calculation is shown in Figure 8.4.3. The top panel shows the calculated $\langle \delta n \rangle$ for each event in the subtraction dataset shown in Figure 8.4.2. Despite the previously discussed stability of the energy projection of the bunch spectrum, there is significant variation in the average relative pixel difference across the dataset. Closer analysis of the differences in the predicted and true bunch profiles (Figure 8.4.3, lower panel) reveals that the majority of the difference in signal corresponds to the accelerated witness bunch. This is expected as the exact form of the wakefield experienced by a bunch slice is an integral over all preceding bunch slices. Therefore, subtle changes in the position or current of each bunch slice due to, for example, small fluctuations in the laser spot profile on the photocathode or RF jitters in the conventional accelerating cavities, are amplified along the bunch as it interacts with the plasma. It is the bunch tail that experiences the accelerating phase of the wakefield and is most sensitive to these jitters, resulting in larger shot-to-shot variability in this region of the spectrometer signal. The averaging procedure has the effect of blurring the shot-to-shot witness energy jitter observed in Figure 8.4.2(a) and leads to larger witness energy spreads in the averaged image that contribute significantly to the difference between true and predicted images. The average relative pixel difference is shown to reduce considerably when considering only the drive bunch, represented by the blue data points in the upper panel of Figure 8.4.3. In this case, the standard deviation of $\langle \delta n \rangle$ across the dataset is found to be $\sigma = 0.02$ compared to $\sigma = 0.07$ when including the accelerated witness bunch in the analysis. Three examples of the pixel-by-pixel difference between the true and predicted bunch profiles are shown in the lower panel of Figure 8.4.3, corresponding to $\langle \delta n \rangle = -0.014$ (left), 0.081 (centre) and 0.167 (right). These represent cases where the average relative pixel difference is close to its smallest (left) and largest (right) for comparison. Clearly, the variation within the drive bunch is much reduced compared to that in the witness bunch where the energy spread is overestimated due to the averaging procedure. The measured standard deviation of the average relative pixel difference is then used as an estimate of the fractional uncertainty introduced within the subtraction method; 7% for this particular dataset. This value is calculated for each separate subtraction dataset and is typically in the range of $5 - 10\,\%$.

8.5 OBSERVED ION MOTION SIGNATURES

Throughout experimentation, four signatures were consistently observed in the electron spectrometer measurements of the probe bunch pair following their reconstruction.

1. Oscillations in the witness beam transverse size.
2. Modifications to the energy of slices that are maximally focused by the quadrupoles.
3. Perturbations to the witness bunch energy spectrum.
4. Perturbations to the drive bunch energy spectrum.

Possible causes for these signatures and their implications for learning about the evolution of the ion distribution within the plasma are now discussed.

8.5.1 *Witness beam size oscillations*

Figure 8.5.1 shows typical projections of the witness bunch spectra in the transverse (non-dispersive) plane of the electron spectrometer as a function of the separation between the drive and probe bunch pairs. In this dataset the spectrometer quadrupoles image at an energy of 1100 MeV, close to the mean energy of the witness bunch. The upper panel shows this projection in the cold measurement scheme where only the probe bunch pair undergoes interaction with the plasma. The witness bunch transverse size is approximately constant across the entire 47 ns measurement range. In comparison, the projections of the perturbed plasma scheme show significant oscillations in the transverse size of the witness bunch. In this case, measurements were performed at two different imaging energies as the witness bunch mean energy was observed to significantly change as a function of the separation between the drive and probe bunch pairs as is discussed in more detail later. The middle panel shows the transverse projection of the witness bunch at an imaging energy of 1100 MeV, the same as in the top panel, while the lower panel shows it for an imaging energy of 1040 MeV. While the absolute value of the beam size in the plane of scintillating screen differs significantly for the two imaging setups, consistent oscillations can be observed when compared to the cold plasma scheme as shown in Figure 8.5.2. This figure shows the relative perturbation to the witness bunch size, $\Delta\tilde{\sigma}_x$, calculated according to:

$$\Delta\tilde{\sigma}_x = \frac{\sigma_{x,p} - \sigma_{x,c}}{\sigma_{x,c}} \tag{8.5.1}$$

as a function of the separation between the drive and probe bunch pairs. $\sigma_{x,p}$ represents the RMS transverse bunch size for the "*perturbed*" scheme and $\sigma_{x,c}$ represents the RMS transverse bunch size for the "*cold*" scheme.

The measured bunch size at the scintillating screen will be a function of both the imaging setup and the divergence of the witness bunch at the plasma exit. As discussed in Section 8.3.1, electrons within the bunches undergo betatron oscillations as they propagate

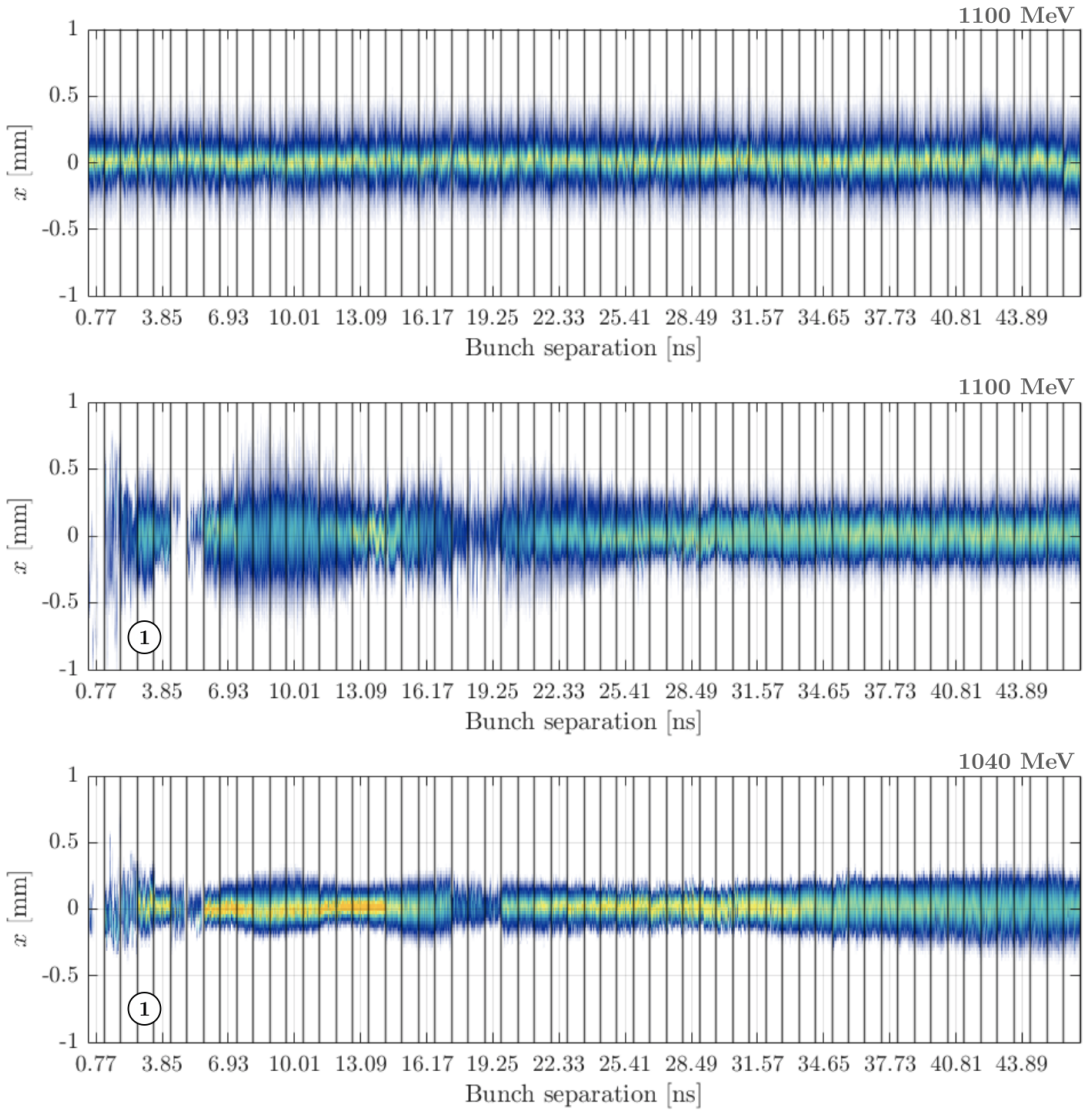

Figure 8.5.1: Comparison between typical cold (upper) and perturbed (middle and lower) scheme transverse spectra indicating signatures of ion motion. (1) Oscillations in the witness beam size. The perturbed scheme measurements are performed at two different imaging energies; 1100 MeV (middle) and 1040 MeV (lower). The cold scheme measurement is performed at an imaging energy of 1100 MeV.

within the plasma as the large amplitude transverse wakefield acts on them. The frequency of these oscillations, ω_β, depends on the energy of an individual particle via:

$$\omega_\beta = \frac{\omega_p}{\sqrt{2\gamma_z}}. \tag{8.5.2}$$

Therefore, slices of the bunch with different energies undergo oscillations at different frequencies and exit the plasma at different phase advances of their betatron oscillation. Consider the evolution of a single energy-slice of a bunch at position ξ_1 within the co-moving frame;

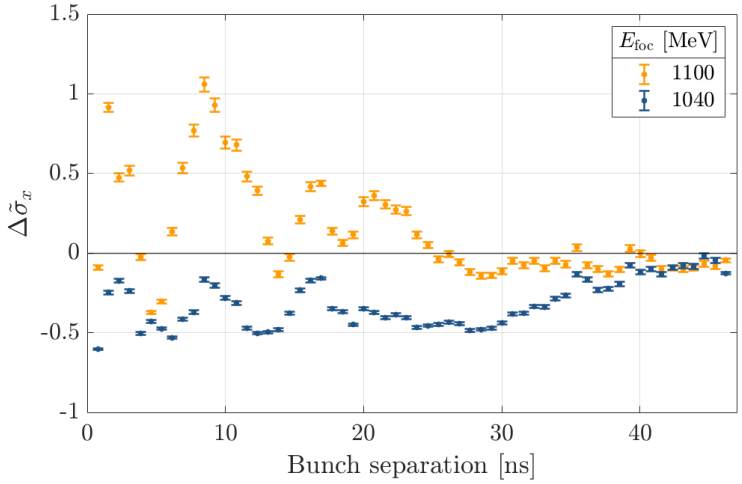

Figure 8.5.2: Perturbations to the transverse witness bunch size as a function of the separation between the drive and probe bunch pairs for two different imaging energies: 1100 MeV (orange) and 1040 MeV (blue).

its transverse size continually oscillates as it propagates through the plasma. Its divergence is given by $d\sigma_x(\xi_1)/dz$ and depends on its betatron phase advance [254]:

$$\phi(\xi) = \int_0^{L_p} k_\beta(\xi, s)ds = \frac{1}{c\sqrt{2}} \int_0^{L_p} \omega_p(s) \left[\gamma_0 + \frac{eE_z(\xi)}{mc^2}s\right]^{-1/2} ds, \tag{8.5.3}$$

where $k_\beta = \omega_\beta/c$, L_p is the plasma length, γ_0 is the initial slice relativistic factor and $E_z(\xi)$ is the magnitude of the longitudinal field at slice ξ. If at the plasma exit the slice has undergone an integer number of oscillations, that is its phase advance $\phi(\xi_1) \approx n\pi$ where $n \in \mathbb{Z}$, it will leave the plasma with a large beam size and hence minimal divergence. Particles from this slice will be captured and focused by the spectrometer quadrupoles to a small size on the scintillating screen and correspond to regions of maximal intensity. In contrast, a different slice (ξ_2) will experience a different longitudinal field and exit the plasma with a betatron phase advance of $\phi(\xi_2) = (n - 1/2)\pi$ such that it is at a spot-size minimum at the plasma exit. In this case, its divergence at the plasma exit is large and it is imaged to a large spot size on the scintillating screen with reduced intensity.

In this experiment, the operating parameters were chosen such that the witness bunch is close to optimally-loading the wakefield in the unperturbed plasma and hence maintains a small energy spread over the acceleration length. Therefore, the amplitude of the longitudinal field is approximately constant over the length of the witness bunch ($E_z(\xi) \approx E_z$) and hence there should not be significant shifts in the betatron phase advance of different slices of the witness bunch. This implies that oscillations in the transverse size of the witness bunch as a function of the bunch separation are instead mainly a result of changes in the on-axis plasma density that the probe bunch pair experiences, i.e. via the $\omega_p(s)$ term in Eq. (8.5.3). As discussed previously, after the leading drive-witness pair propagate through the plasma, the plasma ions are redistributed and a series of approximately parabolic channels are expected

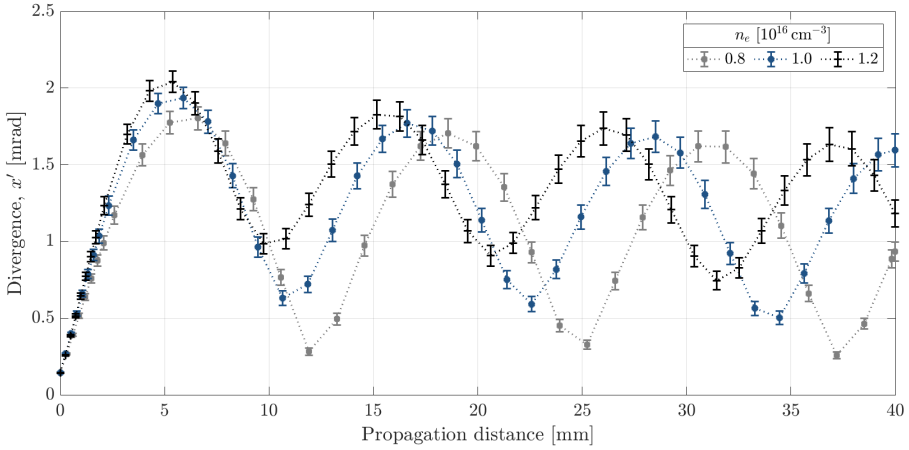

Figure 8.5.3: Simulated development of the witness bunch divergence as a function of propagation distance for three different background plasma densities: 8×10^{15} cm^{-3} (grey, squares), 1.0×10^{16} cm^{-3} (blue, circles) and 1.2×10^{16} cm^{-3} (black, crosses).

to evolve. Over this nanosecond timescale, there is not expected to be significant re-ionisation or decay of the plasma and hence the total number of ions in the region close to the propagation axis should remain approximately constant. Therefore, to develop large peaks off-axis, ions from close to the axis must propagate outwards, leading to a correspondingly lower on-axis density. The probe bunch pair will propagate through this lower on-axis density and hence experience a different plasma frequency, ω_p, and acquire a different betatron phase advance at the plasma exit. As the parabolic profile dissipates due to its intrinsic pressure gradient, the on-axis density gradually returns to its original value assuming no further ionisation or recombination occurs.

The effect of this is shown in Figure 8.5.3 which plots the simulated evolution of the witness bunch divergence for an identical bunch propagating in three different background plasma densities. The divergence of the bunch oscillates as a function of the propagation distance, with faster oscillation frequencies observed for higher plasma densities as $\omega_\beta \propto \omega_p \propto \sqrt{n_e}$. The effect of the plasma density on the betatron phase advance can be seen by considering a fixed plasma length; for example, let $L_p \approx 32$ mm. At this point in the propagation length, at a density of 0.8×10^{16} cm^{-3} the divergence is close to a local maximum and hence a maximal transverse beam size and minimum intensity would be measured on the spectrometer screen. In contrast, a divergence minimum is observed for a plasma density of 1.2×10^{16} cm^{-3} and would correspond to a transverse beam size minimum and intensity maximum at the scintillating screen. Therefore, oscillations in the transverse size of the witness bunch as measured in Fig. 8.5.2 allow the reconstruction of the evolution of the on-axis density experienced by the probe bunch, provided that the bunch profile is sufficiently well known to be simulated accurately. This method additionally relies on the assumption that the plasma length remains constant for different separations between the leading and probe bunch pairs, a reasonable assumption over the few-nanosecond timescales considered.

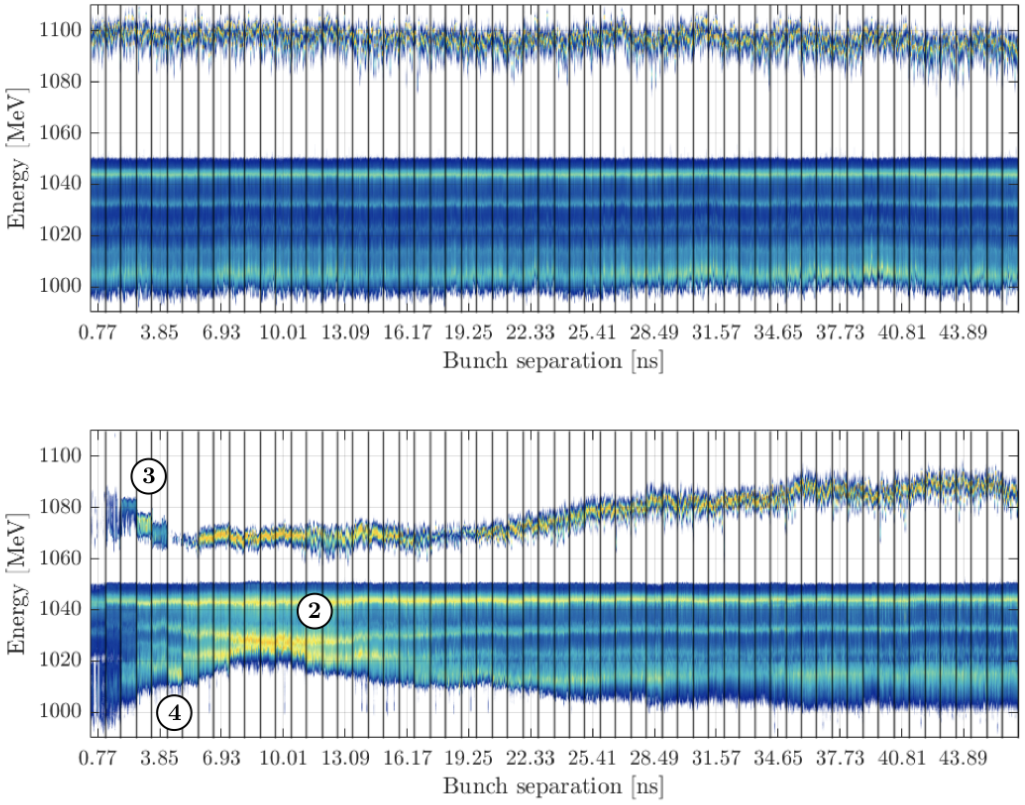

Figure 8.5.4: Comparison between typical cold (upper) and perturbed (lower) scheme energy spectra indicating signatures of ion motion. (2) Perturbation to the spectral "focal" lines. (3) Perturbation to the witness bunch energy. (4) Perturbation to the drive bunch energy. The imaging energy is fixed at 1100 MeV for both measurements.

8.5.2 *Drive bunch focal line perturbations*

The second signature corresponds to perturbations in the energies of slices that are maximally focused in the drive bunch energy spectrum, represented by the (2) marker in Figure 8.5.4. While the witness bunch is short and loads the wakefield such that the variation in longitudinal field along its length is minimised, the same is not true of the drive bunch which experiences a large range of wakefield phase and hence longitudinal field amplitude. This means the betatron phase advance [Eq. (8.5.3)] varies significantly across the drive bunch at the plasma exit and the divergence of individual energy slices similarly varies. In the cold measurement scheme (upper panel) this leads to the observation of constant horizontal lines of increased intensity in the decelerated drive bunch spectrum, corresponding to energy slices that are minimally-divergent at the plasma exit and hence are focused to maximal intensity on the scintillating screen. Experimentally, these lines have been observed to appear at approximately constant energies over order of magnitude changes to the plasma density (not shown here), indicating that the amplitude of the transverse wakefield that leads to their

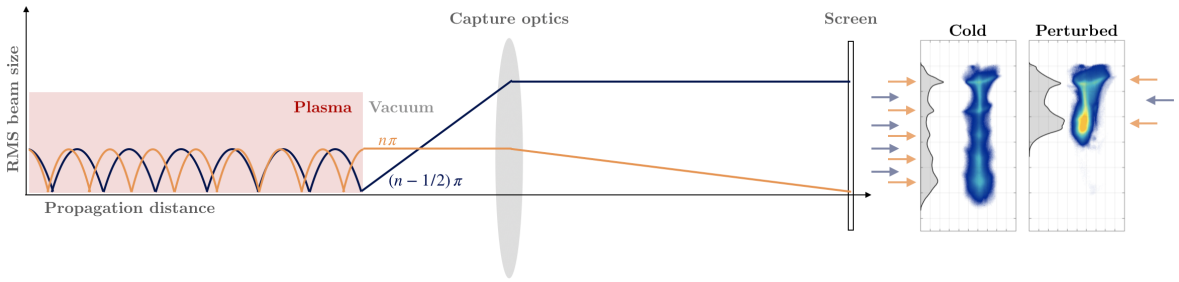

Figure 8.5.5: Simplified drawing demonstrating the effect of the betatron phase of individual energy slices on the beam size and intensity on the scintillating screen. Two different energy slices (blue and orange) undergo different numbers of betatron oscillations within the plasma, exiting with maximal (minimal) divergence. The capture optics then focus the slices to minimal (maximal) beam sizes on the screen. Example spectrometer images are shown demonstrating regions of maximal and minimal intensity represented by orange and blue arrows respectively. Shifts in the position of the regions of maximal intensity can be observed between the cold and perturbed measurement schemes.

appearance is independent of the background plasma density and is rather a function of only the bunch density. This suggests a (quasi-)linear wakefield regime and is attributed to the head of the bunch which typically has a larger emittance (and hence lower density) due to coherent synchrotron radiation effects induced during the transport of the bunch to the plasma, and hence drives a weaker wakefield. The position of lines of increased intensity within the energy spectrum are additionally only observed to remain constant for slices that lose only a small amount of energy, again indicating that low amplitude wakefields cause their appearance.

In contrast, the positions and intensities of these focal lines are observed to move as a function of the bunch separation in the perturbed measurement scheme. The largest changes in focal energy observed in Figure 8.5.4 occur close to the maximum perturbation to the drive bunch mean energy at a bunch separation of around 10 ns (when compared to the cold measurement scheme) before slowly returning to their original energy as the perturbation to the drive bunch mean energy similarly decays. As discussed in Sec. 8.3.1, the introduction of parabolic density profiles will modify the betatron oscillation frequency of energy slices of the bunch due to the increased off-axis focusing from the radially-increasing ion density. This modifies the betatron phase advance [Eq. (8.5.3)] of all energy slices such that, for a fixed plasma length, a different energy slice will have undergone an integer number of betatron oscillations and exit the plasma at a divergence minimum to be focused to a minimum size and maximal intensity on the spectrometer screen. This is exemplified in Fig. 8.5.5 which illustrates the divergent-dependent focusing on the spectrometer screen for two different energy slices. In the example bunch image in the cold measurement scheme shown in the right of the figure, there is a clear modulation in the intensity of the spectrometer image, indicated by orange and blue arrows, corresponding to energy slices of the bunch that are maximally and minimally divergent at the plasma exit. A similar intensity modulation is observed in the perturbed measurement scheme, however the position of the intensity maximum has

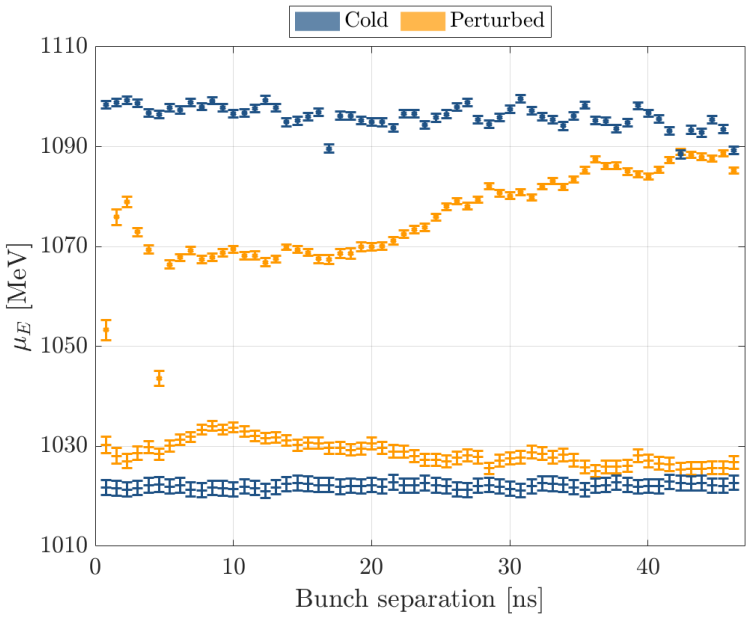

Figure 8.5.6: Mean energy evolution for the cold (blue) and perturbed (orange) measurement schemes of the drive (crosses) and witness (circles) bunches shown in Fig. 8.5.4. The error bars represent the standard error on the mean.

shifted to a lower energy due to the change in betatron oscillation frequency induced by the non-uniform transverse density profile.

The energies at which maximal intensities are measured in the perturbed energy spectrum can be determined and compared to those in the unperturbed plasma spectrum. This shift in focal energy, in combination with Eq. (8.3.20), can be used to reconstruct the curvature of the parabolic density profile at each bunch separation, $\alpha(\Delta\tau)$.

8.5.3 *Energy perturbations*

The final signatures observed experimentally are the perturbations to both the witness and drive bunch mean energy, represented by the (3) and (4) markers in Figure 8.5.4, indicative of the modification of the longitudinal wakefield structure by the continuously evolving transverse density profile as discussed in Section 8.3.3. The extracted mean energy of the drive and witness bunches as a function of the separation between the drive and probe bunch pairs are plotted in Fig. 8.5.6. In the cold measurement scheme, the drive bunch energy is observed to slowly trend upwards over the entire measurement range while the witness bunch mean energy slightly decreases, a result of the recombination of the plasma leading to slightly lower interaction densities. In contrast, the perturbed witness bunch mean energy initially reduces when compared to the cold plasma measurement scheme, reaching a minimum at around 10 ns, before tending back towards the unperturbed value. A corresponding increase in the mean energy of the drive bunch is similarly observed over this timescale, clearly demonstrating the evolution of the ion motion perturbation as it builds over the first few nanoseconds,

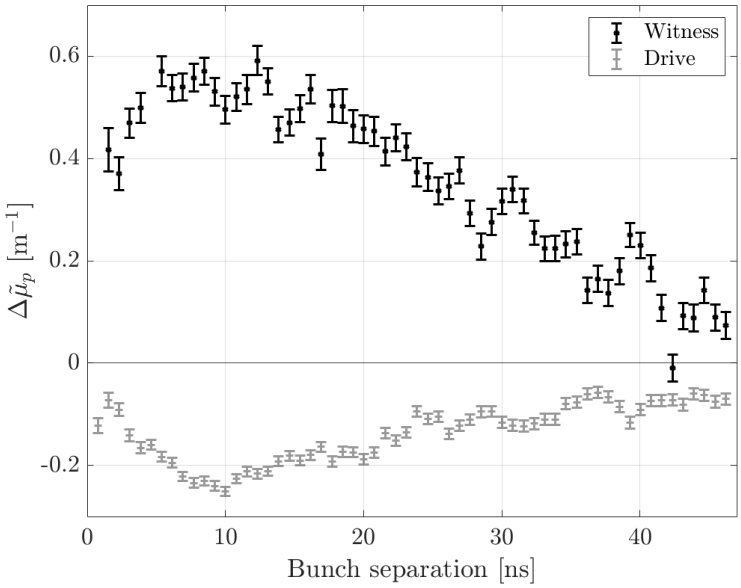

Figure 8.5.7: Evolution of the relative perturbative effect, $\Delta\tilde{\mu}_p$, calculated via Eq. (8.5.4), for the measurements presented in Figure 8.5.6. The perturbation to the witness bunch is marked with black circles while the perturbation to the drive bunch is marked with grey crosses.

peaks after approximately 10 ns in this case, and then dissipates over the following tens of nanoseconds. At the shortest bunch separations ($\Delta\tau = 0.77, 4.62$ ns), the majority of the charge of the witness bunch is not successfully transported through to the spectrometer (as can be seen in Figure 8.5.4) and hence the witness bunch is not correctly identified by the automated image analysis algorithm.

The relative perturbative effect induced by ion motion can be calculated from measurements of the drive and witness bunch mean energies (Figure 8.5.6) via the equation

$$\Delta\tilde{\mu}_p = \frac{\mu_{E,c} - \mu_{E,p}}{\mu_{E,c} \cdot L_p}, \tag{8.5.4}$$

where $\mu_{E,c}$ and $\mu_{E,p}$ are the mean energies of the (drive or witness) bunch in the cold and perturbed measurement schemes, and L_p is the plasma length, here assumed to be equal to the length of capillary, 50 mm. This is shown in Figure 8.5.7 and again demonstrates growth of the perturbative effect over the first few nanoseconds, a peak at around 10 ns, and dissipation over tens of nanoseconds. The perturbative effect is significantly greater in the case of the witness bunch as it is short with respect to the wakefield cavity and experiences only a small region of the evolving longitudinal wakefield, whereas the longer drive bunch is subject to the majority of the decelerating phase and hence is less sensitive to small changes. The use of a relative perturbation metric facilitates easy comparison between measurements performed to investigate the dependency of the ion motion signatures on different experimental variables and is explored in Section 8.6.

The evolution of the plasma ion distribution induced by the leading bunch pair creates a combination of non-uniform transverse density profiles and a changing on-axis plasma den-

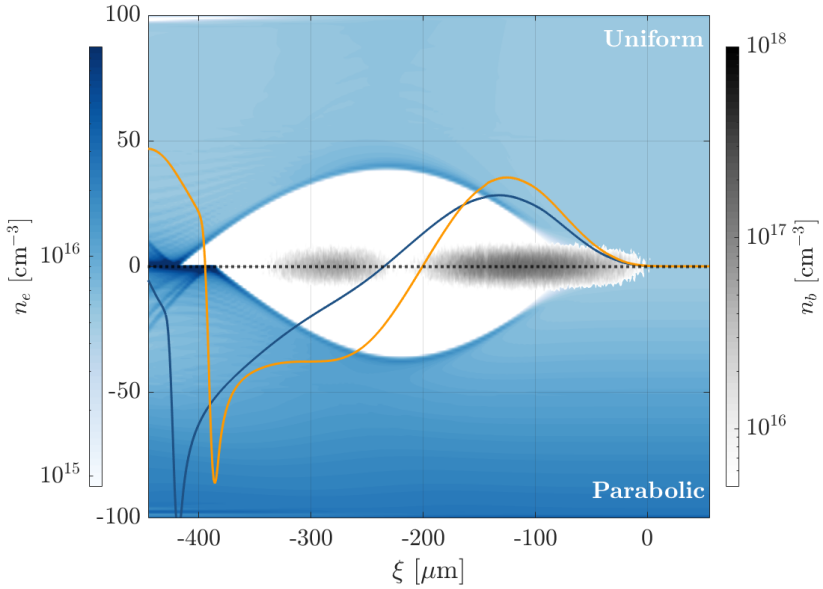

Figure 8.5.8: Simulated comparison between the wakefield driven in a uniform transverse density profile ($n_0 = 5.6 \times 10^{15}$ cm^{-3}) and that with a parabolic profile parameterised by $\alpha = 1.3$. The upper half of the image represents the plasma electron density (n_e) in the uniform density case, whereas the lower half represents the parabolic density case. The bunch density (n_b) is plotted in grey and is identical in both simulations. The on-axis longitudinal wakefield is overlaid in blue (uniform) and orange (parabolic).

sity that significantly alters the wakefield structure driven, and experienced, by the probe bunch drive-witness pair as shown in Section 8.3.3. As the ion channel deepens over the first few nanoseconds, the on-axis density reduces leading to a reduction in the wakefield amplitude and lengthening of the wakefield cavity. This additionally modifies the relative phase of the witness bunch within the wakefield, further reducing the accelerating gradient it experiences. In contrast, as discussed in Sec. 8.3.2, the presence of parabolic transverse density profiles act to truncate the wakefield cavity due to the increased focusing off-axis. The effect of this is shown in Figure 8.5.8 which compares the simulated wakefield response to an identical drive-witness pair for two cases: (i) a uniform transverse density profile (upper), $n_0 = 5.6 \times 10^{15}$ cm^{-3}, and (ii) a parabolic density profile (lower) of the same on-axis density but with curvature parameterised by $\alpha = 1.3$. In the case of the parabolic profile the wakefield cavity is truncated as predicted in Section 8.3.2 and this significantly modifies the form of the on-axis longitudinal wakefield profile, plotted as solid lines. The peak decelerating field amplitude is increased by approximately 25% in the parabolic case and the position of the zero-crossing shifts forwards within the wakefield cavity, as previously discussed in Section 8.3.3. This alters the relative phase of the witness bunch within the field and allows it to significantly load the wakefield, simultaneously increasing its energy and reducing its energy spread at the plasma exit relative to the uniform case. This effect is observed experimentally (Figure 8.5.4) as the energy spread of the witness bunch remains small as the perturbation evolves despite the large change in on-axis density experienced by the probe bunches.

8.5.4 *Transverse plasma density profile reconstruction*

The four signatures described above can be used to reconstruct the evolution of the transverse plasma density. This is because the two parameters required to reconstruct the parabolic shape, n_{e0} and α, can be independently determined. The evolution of the on-axis density, n_{e0}, is reconstructed from measurements of the oscillations in the size of the witness bunch on the spectrometer screen, while the evolution of the channel curvature, α, is reconstructed from the change in focal lines within the drive bunch spectrum. These two effects combined then result in the perturbations to the witness and drive bunch energy spectra. A general description of the reconstruction process follows:

1. Reconstruct the bunch profile from experimental measurements of the longitudinal and transverse phase space to be used in simulation. These measurements and their reconstruction are described in Chapters 4 and 5 and are shown for the probe bunch bunch used in these measurements in Figures 8.5.9 and 8.5.10.

2. Perform a detailed simulation scan where the reconstructed bunch propagates through a series of transversely-uniform plasma density profiles over a large range of on-axis densities (at least one order of magnitude). This scan is used to find the plasma density which best recreates the experimentally measured energy spectrum in the cold plasma measurement scheme.

3. From the plasma density scan, determine the divergence of the witness bunch as a function of the plasma density (e.g. Figure 8.5.3). Match the extrema of the experimental measurements of the witness beam transverse size (e.g. Figure 8.5.2) to approximate on-axis plasma densities in simulation.

4. Use the measured deviations to the position of the focal energies within the spectrum combined with Eq. (8.3.20) to predict the curvature of the density profile as a function of the bunch separation.

This procedure will return a series of on-axis density values along with their expected curvature that can be used to recreate the evolution of the transverse density profile. These are then used in three-dimensional PIC simulations in an attempt to reconstruct the measured probe bunch transverse and energy spectra. An example for the measurements of the evolution of the probe bunch presented in Figures 8.5.1 and 8.5.4 is discussed below.

Firstly, the probe bunch must be accurately reconstructed for use in the simulations. Measurements of the focusing of different slices of the probe bunch are shown in Figure 8.5.9, calculated using the positional jitter of the slice between two BPMs either side of the interaction chamber (with no plasma present). The Twiss parameters of each slice are determined using the methodology presented in Ref. [173], enabling calculation of the focal plane of the bunch in both transverse planes, represented by the black dashed lines in the right panels in Figure 8.5.9. Note, the entrance to the capillary corresponds to a position of approximately $s = -0.2\,\mathrm{m}$ within this frame. These measurements enable determination of the

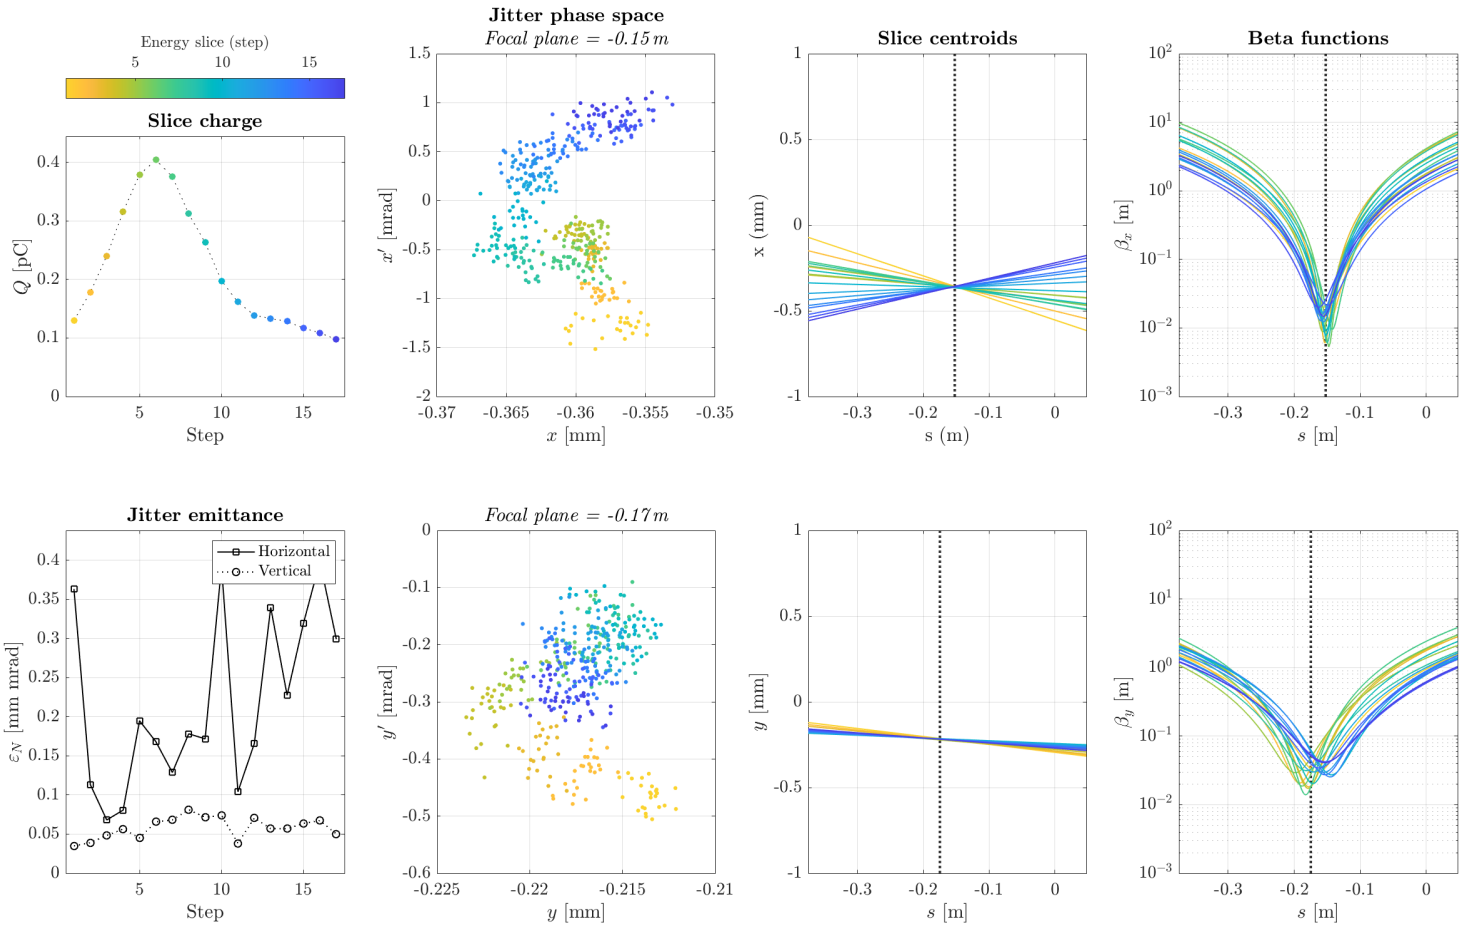

Figure 8.5.9: Reconstruction of the slice-dependent focusing of the probe bunch pair used in measurements described in Sections 8.6.1 and 8.6.2 using the two-BPM tomography technique developed at FLASHForward by C. A. Lindstrøm. Full details of this technique can be found in Ref. [173].

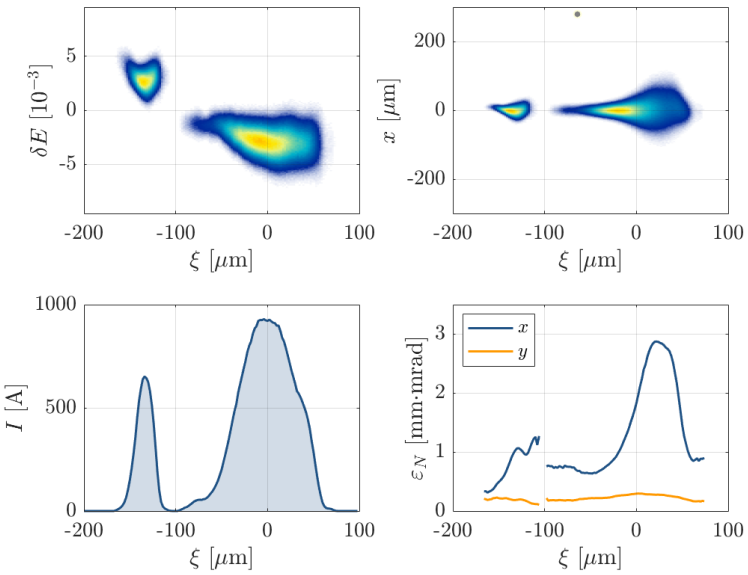

Figure 8.5.10: Reconstruction of the probe bunch pair at the entrance of the plasma used in measurements described in Sections 8.6.1 and 8.6.2. Upper left: longitudinal phase space. Upper right: bunch distribution in the transverse $x - \xi$ plane. Lower left: current profile. Lower right: slice-emittance in both the x and y transverse planes.

slice-dependent focusing of the bunch which can then be applied to the reconstructed bunch profile that is used in the three-dimensional PIC simulations.

Measurements of the phase space of the bunch performed using the X-TDS are recreated in Figure 8.5.10 and are similarly applied to the simulated bunch profile. The reconstructed bunch profile is then propagated through longitudinally- and transversely-uniform plasma of varying density in simulation in an attempt to determine the operating plasma density in the unperturbed plasma case. The result of such simulations is demonstrated in Figure 8.5.11.

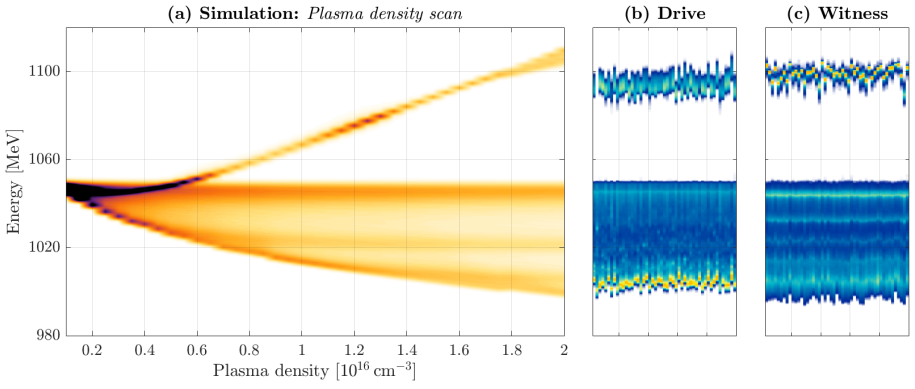

Figure 8.5.11: Comparison between (a) simulated plasma density scan and (b), (c) experimental measurements of the interaction between the probe bunch pair and the unperturbed plasma for imaging energies of 1040 MeV (drive) and 1100 MeV (witness) respectively. This scan is used to identify the operating plasma density.

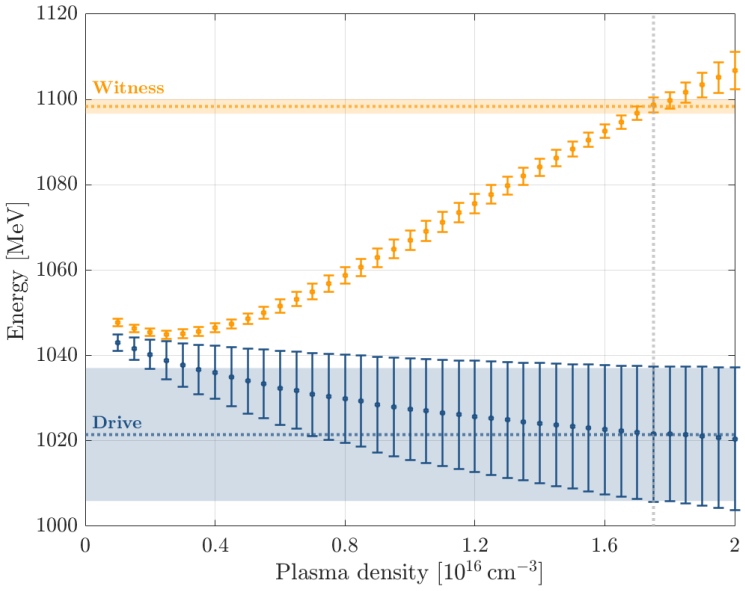

Figure 8.5.12: Mean energy (data points) and RMS energy spread (error bars) of the simulated probe drive (blue) and witness (orange) bunches as a function of the plasma density. The horizontal dashed lines represent the mean energy of the experimentally-measured probe bunch for the drive (blue) and witness (orange) bunches respectively, while their RMS energy spreads are represented by the shaded regions. The vertical dashed line at $n_e = 1.75 \times 10^{16}\,\mathrm{cm}^{-3}$ corresponds to the simulation that shows best agreement and hence indicates the operating density.

In panel (a), Figure 8.5.11 shows the simulated energy spectrum at the plasma exit after the probe bunch propagates through 32 mm of plasma of varying density, $n_e = 1.0 \times 10^{15} - 2.0 \times 10^{16}\,\mathrm{cm}^{-3}$, and that measured experimentally in panel (b) and (c) for imaging energies of 1040 MeV (drive) and 1100 MeV (witness) respectively. Comparison between the two enables determination of the longitudinally-averaged plasma density experienced by the bunch pair in the unperturbed plasma state as shown in Figure 8.5.12. In this case, an operating density of approximately $n_0 = 1.75 \times 10^{16}\,\mathrm{cm}^{-3}$ is found to provide the best comparison to the experimentally measured energy spectrum. This can be compared to measurements of the plasma density performed *in situ* using the spectral line broadening technique discussed in Section 5.1.2. This diagnostic was added to the experimental setup and used a single fibre optic and collimating lens to transport light emitted by the plasma through a view-port to a spectrometer located in the neighbouring technical corridor. Measurements of the light emitted by the hydrogen-doped (2%) argon plasma around the Hα line in the Balmer series ($\lambda = 656.28$ nm) indicated a plasma density of $(1.77 \pm 0.25) \times 10^{16}\,\mathrm{cm}^{-3}$ for these operating conditions, in agreement with the longitudinally-integrated density found in PIC simulations.

Once the operating plasma density has been determined, the simulated relative divergence of the witness bunch at the plasma exit as a function of the plasma density can be calculated.

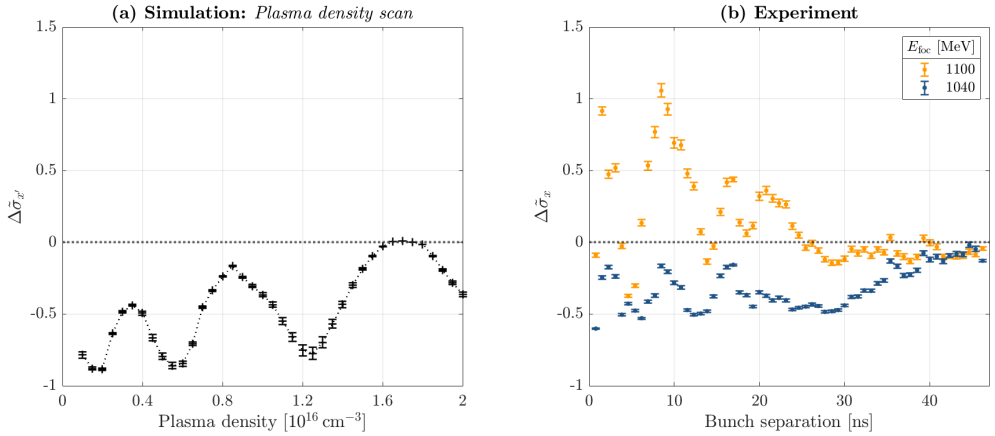

Figure 8.5.13: Comparison between (a) the relative divergence of the witness bunch during a simulated plasma density scan and (b) experimental measurements of the witness bunch size in the plane of the scintillating screen as a function of the bunch separation.

Figure 8.5.13(a) shows this for the plasma density scan presented in Figure 8.5.11. The divergence of the bunch is calculated relative to that found at the operating density, n_0, via

$$\Delta\tilde{\sigma}_{x'} = \frac{\sigma_{x',n_e} - \sigma_{x',n_0}}{\sigma_{x',n_0}}, \tag{8.5.5}$$

where σ_{x',n_0} is the witness bunch divergence at the operating density and σ_{x',n_e} is the same at a plasma density of n_e. Such a scan can be used to estimate the on-axis plasma density as a function of the bunch separation for the experimental measurements presented in Figure 8.5.13(b) as discussed in Section 8.5.1. Measurements of the energy spectra of the drive and witness bunch in the perturbed plasma state shown in Figure 8.5.4 indicate that the peak perturbation, and hence on-axis density minimum experienced by the probe bunch pair, occurs at approximately 9 ns. This corresponds to a transverse bunch size maximum according to Figure 8.5.13(b). At bunch separations larger than this, one further oscillation in the transverse bunch size is observed before the perturbation tends back to that measured in the unperturbed plasma state. This behaviour matches well with the simulated witness bunch divergence over the range of plasma densities from $n_e \sim 3.5 \times 10^{15} - 1.75 \times 10^{16}\,\mathrm{cm}^{-3}$. The extrema of the oscillations can therefore be matched to these values, with the peak at $\Delta\tau \sim 9$ ns occurring at an on-axis plasma density of $n_e \sim 3.5 \times 10^{15}\,\mathrm{cm}^{-3}$, and following extrema at $\Delta\tau \sim 13, 18, 29$ ns corresponding to on-axis densities of $n_e \sim 5.5, 8.5, 12.5 \times 10^{15}\,\mathrm{cm}^{-3}$ respectively. The intervening densities can be interpolated to give the evolution of the on-axis density for each bunch separation.

The evolution of the curvature of the density profile is extracted from comparisons of the position of the focal energies within the drive bunch energy spectrum in the perturbed and cold cases. This is shown in Figure 8.5.14, where waterfalls of the drive bunch energy spectra for bunch separations $\Delta\tau = 0.77 - 27.72$ ns are plotted in (a) for the unperturbed plasma and (b) in the perturbed measurement scheme. The focal energies in the perturbed case are calculated by performing peak fitting to the projection of the spectrometer images in the

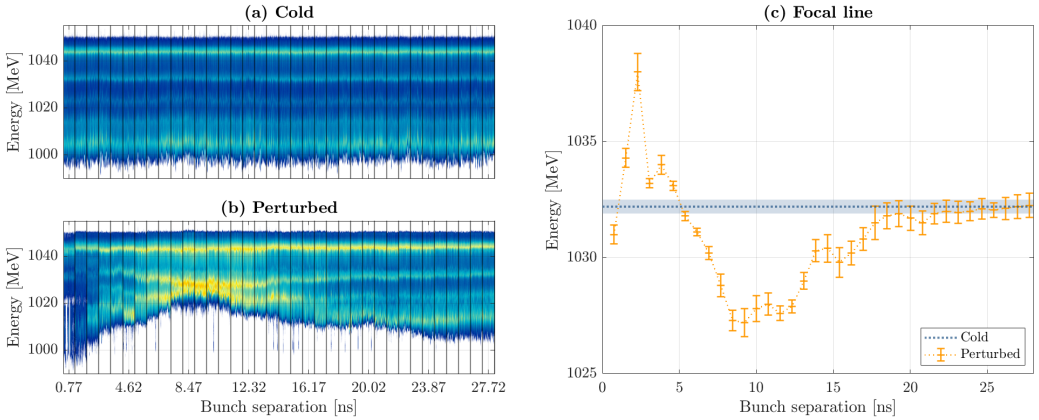

Figure 8.5.14: Extraction of the focal energies within the measured drive bunch spectra. (a) Drive bunch energy spectra within the cold measurement scheme. (b) Drive bunch energy spectra within the perturbed measurement scheme. (c) Extracted energy of the focal line as a function of the bunch separation for the cold (blue) and perturbed (orange) schemes.

dispersive plane at each bunch separation and are shown in Figure 8.5.14(c) for this case, with error bars representing the average FWHM of the peak within each bunch separation dataset. At bunch separations around 9 ns, multiple focal lines appear in a small energy range within the spectra, leading to an increase in the FWHM as the automatic peak finding routine cannot always correctly identify individual peaks. At the shortest timescales a large fraction of the drive bunch charge is lost and the hence identification of the peaks in the spectra is more challenging. Determination of the change in focal energies within the spectra, as in Figure 8.5.14(c), can then be combined with Eq. (8.3.20) to extract the value of the curvature of the transverse density profile as a function of the bunch separation.

The resulting evolution of the transverse ion density profile extracted from the measurements of the spectra of the probe bunch in Figures 8.5.1 and 8.5.4 is shown in Figure 8.5.15. At short timescales, $\Delta\tau < 4\,\text{ns}$, the on-axis density is assumed to have increased as ions are initially attracted towards the axis by the regions of high electron density within the wakefield. The on-axis density then decreases to a minimum at approximately $\Delta\tau \sim 9\,\text{ns}$ before slowly returning towards the background operating density of $n_0 = 1.75 \times 10^{16}\,\text{cm}^{-3}$ found via the plasma density scan in Figure 8.5.11. Initially the curvature of the density profile is assumed to be small and grows to a maximum of $\alpha \approx 1.85$, at approximately $\Delta\tau \sim 9\,\text{ns}$, calculated via the perturbation to the focal energies shown in Figure 8.5.14. Over extended timescales the curvature approaches zero as the off-axis peak, and hence non-uniformity of the transverse density profile, dissipates. A map of the radial density profile, normalised to n_0, is shown in Figure 8.5.15(b) for reference.

These parabolic density profiles are then used as input to three-dimensional PIC simulations of the interaction of the reconstructed probe bunch. A separate simulation is performed for each bunch separation using the corresponding predicted parabolic channel in Figure 8.5.15(b) as the background transverse density profile. The output of these simula-

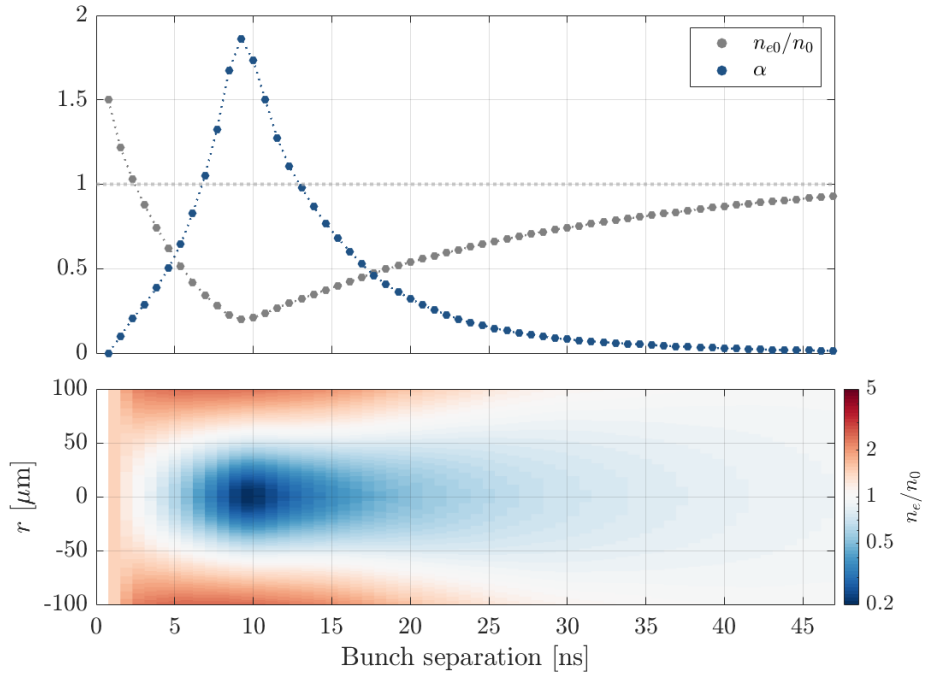

Figure 8.5.15: Reconstructed evolution of the transverse ion density profile from the beam-based measurements presented in Figures 8.5.1 and 8.5.4. (a) Extracted values of the normalised on-axis density, n_{e0} (grey), and the curvature of the parabolic density profile, α (grey). (b) Map of the radial ion density profile as a function of the bunch separation. Plasma densities are normalised to the background operating density, $n_0 = 1.75 \times 10^{16}\,\mathrm{cm}^{-3}$.

tions is then compared to the experimentally measured probe bunch energy and transverse spectra. Such a comparison is shown in Figures 8.5.16 and 8.5.17.

Figure 8.5.16 shows the mean energy of the probe (drive and witness) bunch as a function of separation between the leading and probe bunch pairs. The experimental measurements (as in Figure 8.5.6) are represented by the data points while the mean energy of the simulated drive and witness bunches are represented by the dotted lines. As discussed previously, at the shortest separations the on-axis density is initially expected to increase as ions are attracted to regions of increased electron density on-axis as the wakefield structure dissipates. This leads to highly non-uniform transverse density profiles over small ($\sim$ µm) scales and a significant amount of charge is lost from the drive and witness bunches of the probe bunch pair as can be seen in Figure 8.5.4. This loss of charge makes it difficult to reconstruct the density profile at short separations and significant disagreement between the experimental measurements and simulation predictions (using an *assumed* profile) is observed in Figure 8.5.16. It is possible that this short timescale evolution ($\sim$ns) could be within reach of quasi-static PIC simulations in the near future and hence should be investigated further to see if improved agreement can be found. Following this, the evolution of the mean energy of both the drive and witness probe bunches agrees well with simulation predictions indicating that the transverse profile reconstruction technique provides a good estimate of the evolution of the plasma.

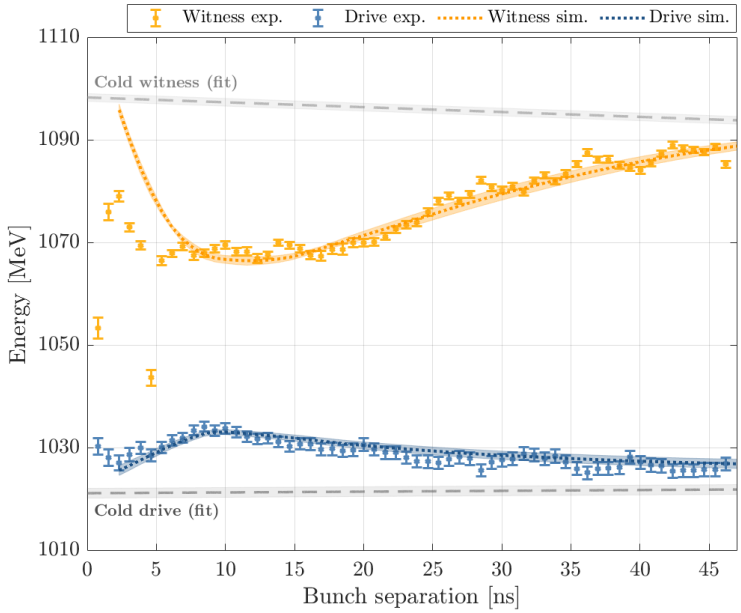

Figure 8.5.16: Comparison between the experimentally measured probe bunch mean energies (data points) and that predicted by three-dimensional PIC simulations (dotted lines) using the reconstructed transverse density profile extracted via signatures within the measured spectra as shown in Figure 8.5.15. The error bars (exp.) and bands (sim.) represent the standard error on the mean. The grey dashed lines represent linear fits to the measurements of the probe drive (darker) and witness (lighter) bunches in the unperturbed measurement scheme with uncertainties represented by the shaded regions.

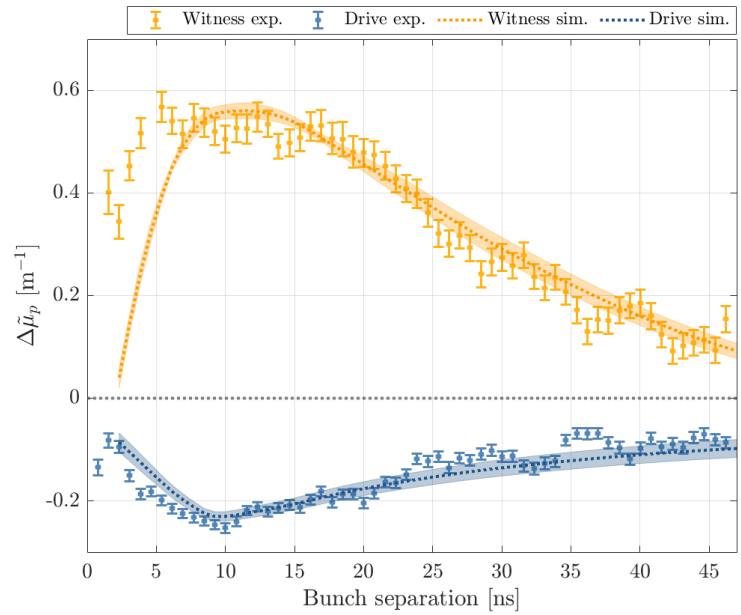

Figure 8.5.17: Comparison between the experimentally measured probe bunch energy perturbation ($\Delta\tilde{\mu}_p$) and that predicted by three-dimensional PIC simulations. Both the experimental and simulated datasets are compared to the linear fits to the mean energies in the unperturbed measurement scheme as shown by the grey dashed lines in Figure 8.5.16.

Figure 8.5.17 shows the evolution of the energy perturbation metric defined by Eq. (8.5.4) for both the experimental measurements and the predictions of simulations using the transverse density profile evolution given in Figure 8.5.15. In both cases, the value of the perturbation metric is calculated via measurements of the mean energy of the perturbed bunches with respect to linear fits of the evolution of the mean energy of the probe bunches in the case of the unperturbed plasma. These fits are represented by the grey dashed lines in Figure 8.5.16. Such fits are used in an attempt to minimise the effect of periodic fluctuations observed in the measurements of both the cold and perturbed measurement schemes as can be seen in Figure 8.5.6. The origin of these oscillations is not confirmed, but one possible cause is the disparity between the RF frequency of the triggers for the bunch and the discharge systems. While performing these measurements, the arrival time of the leading bunch and the discharge trigger are moved in tandem to maintain a constant interaction plasma density for the leading bunch. However, the bunch arrival time is defined to the nearest $\Delta\tau_{min} = 0.77\,\text{ns}$ (1.3 GHz) while the discharge trigger can only be moved in steps of 0.923 ns (1.083 GHz) as this is derived from the primary diagnostics trigger at 108.3 MHz. This can lead to temporal discrepancies between the two of up to $\sim 0.5\,\text{ns}$ which can have a small effect on the plasma density at interaction. Peaks in the unperturbed witness bunch energy are typically separated by approximately $6 \times \Delta\tau_{min} = 4.61\,\text{ns}$ which corresponds to 5 steps at the discharge trigger frequency. The combined "beating" effect of these two differing frequencies is therefore consistent with the periodic fluctuations observed in the experimental measurements. However, the initiation of the discharge itself would be expected to have $\mathcal{O}$(ns) temporal jitter thanks to the thyratron that is used as a switch, and hence it is likely there are other additional contributors to this effect. Despite the presence of this oscillation, good agreement between the experimentally measured values of $\Delta\tilde{\mu}_p$ and those found via simulations are observed following the initial evolution.

Comparisons between the transverse and energy spectra measured experimentally and recreated using three-dimensional PIC simulations are shown in Figure 8.5.18(a) and (b) respectively. The experimental spectra correspond to measurements of the probe bunch pair energy distribution for an imaging energy of 1100 MeV. The simulated distributions are calculated by propagating all PIC-simulated bunch macroparticles from the plasma exit through a simple matrix model of the post-plasma quadrupoles and spectrometer dipole and reconstructing the measurement of the distribution in the plane of the scintillating screen. This model recreates the imaging setup used in the experiment and assumed an energy of 1100 MeV with a horizontal magnification factor of -5 in the non-dispersive plane of the spectrometer and a vertical magnification of -0.97. Reasonable agreement between the simulated energy spectra and those measured experimentally is again observed except in the witness bunch on short timescales ($\Delta\tau < 5\,\text{ns}$) for reasons discussed previously. The position of the focal line of increased intensity is observed to shift as observed in the experiment, particularly around the peak perturbation ($\Delta\tau \sim 10\,\text{ns}$), as the steep parabolic channel modifies the betatron frequency of beam particles.

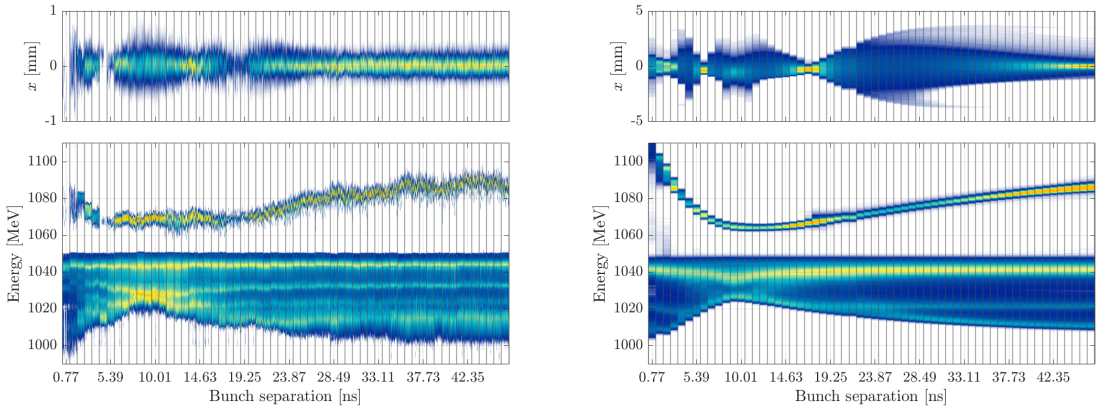

(a) (b)

Figure 8.5.18: Comparison between the perturbed transverse (upper) and energy (lower) spectra (a) measured experimentally and (b) recreated in three-dimensional PIC simulations using the evolution of the transverse density profile shown in Figure 8.5.15.

Use of a simple spectrometer dipole model leads to a difference in the minimum energy of the decelerated drive bunch between the two spectra. This apparent change in maximal energy loss is caused by the vertical divergence of particles with energies significantly different to the imaging energy. Its effect can also be seen in Figure 8.5.11 by comparing the energy spectra of the decelerated drive bunch for the two different imaging energies shown in panels (b) and (c) of that figure. For this reason, measurements were performed at two imaging energies, close to the mean energy of the drive bunch (1040 MeV) and the witness bunch (1100 MeV), to improve the energy resolution of each bunch individually. There is also a significant difference between the size of the witness bunch in the plane of the scintillating screen in the experiment and in simulation. This is likely due to the choice of a longitudinally-uniform plasma density profile within simulations. It is rather expected that there will be extended, lower density ramps at either end of the plasma as measured in Ref. [195] that act to reduce the divergence of the bunch as it exits the plasma [190], leading to less extreme changes in bunch size as it propagates through the imaging optics. Despite this difference in absolute size, the oscillation in transverse witness bunch size caused by the evolving on-axis density can still be observed.

The beam-based plasma characterisation method introduced in Chapter 5 has been successfully applied to measurements of the evolution of the plasma state over nanosecond timescales following the wakefield excitation process. Identification of specific features within the transverse and energy spectra of the probe bunch has enabled reconstruction of the evolution of the transverse plasma density profile with good agreement observed between three-dimensional PIC simulations and experimental measurements. Such a beam-based technique is therefore suitably sensitive to be used to study how the evolution of the transverse density profile depends on different experimental parameters and such measurements are the subject of subsequent sections within this chapter.

8.6 ION MOTION DEPENDENCIES

After developing the beam-based technique that allows mapping of the motion of ions within the plasma following a wakefield event as discussed in Section 8.5, the technique was used to study the dependency of ion motion on three key experimental variables:

1. The amplitude of the wakefield driven by the leading bunch.
2. The background plasma density at the interaction time of the leading bunch.
3. The mass of the ions of the plasma in which the wakefield is driven.

Such measurements were performed to investigate scaling laws that could descibe the evolution of the ion motion perturbation in parameter regimes that are currently inaccessible to the FLASHForward experiment. One such example is those assumed in plasma-based linear collider proposals [54], where operating plasma densities of order $10^{16} - 10^{17}\,\mathrm{cm}^{-3}$ are predicted using drive bunches with peak currents of order 10 kA.

8.6.1 *Wakefield strength*

The strength of the interaction between the leading bunch and the plasma was modified by changing the timing of the photocathode laser that generated the electron bunch with respect to the RF phase in the conventional accelerator section of the FLASH beamline. Modifying the time-of-arrival (TOA) of the bunch changes its relative RF phase resulting in modification of the compression of the bunch. It was found that picosecond-level shifts in the TOA led to significant changes in the peak current of the bunch and hence wakefield interaction strength. Generation of the drive and witness bunches was performed using the wedge, head and tail scrapers within the dispersive section of the FLASHForward beamline [159]. For ease of experimentation, the positions of these were not modified as the TOA was changed. This therefore meant that the amount of charge in the drive and witness bunches varied as the longitudinal phase space of the bunch was changed by modifying its TOA. Three stable operating points were found, judged by the stability of the plasma interaction, corresponding to TOAs of 225.87 ps, 226.27 ps and 227.27 ps as shown in Figure 8.6.1. This figure shows the average spectrometer image of the drive and witness bunches for the three different TOA operating values without plasma interaction.

The lowest TOA (225.87 ps) corresponded to the lowest peak current of the drive bunch and hence the weakest plasma interaction. In contrast, for a TOA of 227.27 ps the drive bunch charge and peak current is much increased, with only a small fraction of the bunch charge forming a witness. In this case, the witness is not successfully captured after the plasma interaction as shown in Figure 8.6.2(c). For the other two TOAs, the accelerated witness bunch can be observed at approximately 1095 MeV and 1085 MeV for TOAs of 226.27 ps and 225.87 ps respectively, demonstrating average accelerating gradients of approximately $0.8\,\mathrm{GVm}^{-1}$ and $0.6\,\mathrm{GVm}^{-1}$ over an *assumed* 50 mm plasma length. It is rather expected that the plasma length

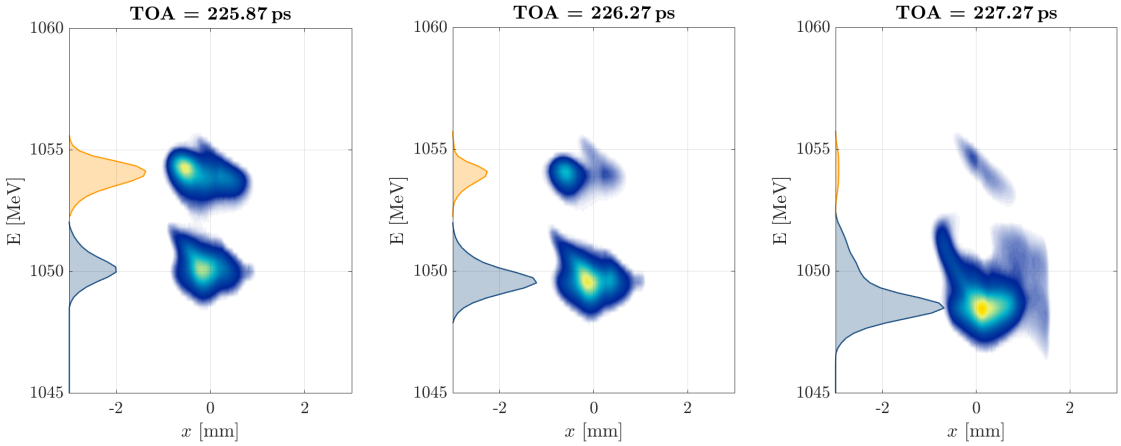

(a) (b) (c)

Figure 8.6.1: Average spectrometer images for Laser 1 time-of-arrivals (TOAs) of (a) 225.87 ps, (b) 226.27 ps and (c) 227.27 ps *without* plasma interaction. All images use the same colour scale. The projections in the dispersive plane are represented by blue and orange regions for the drive and witness bunches respectively.

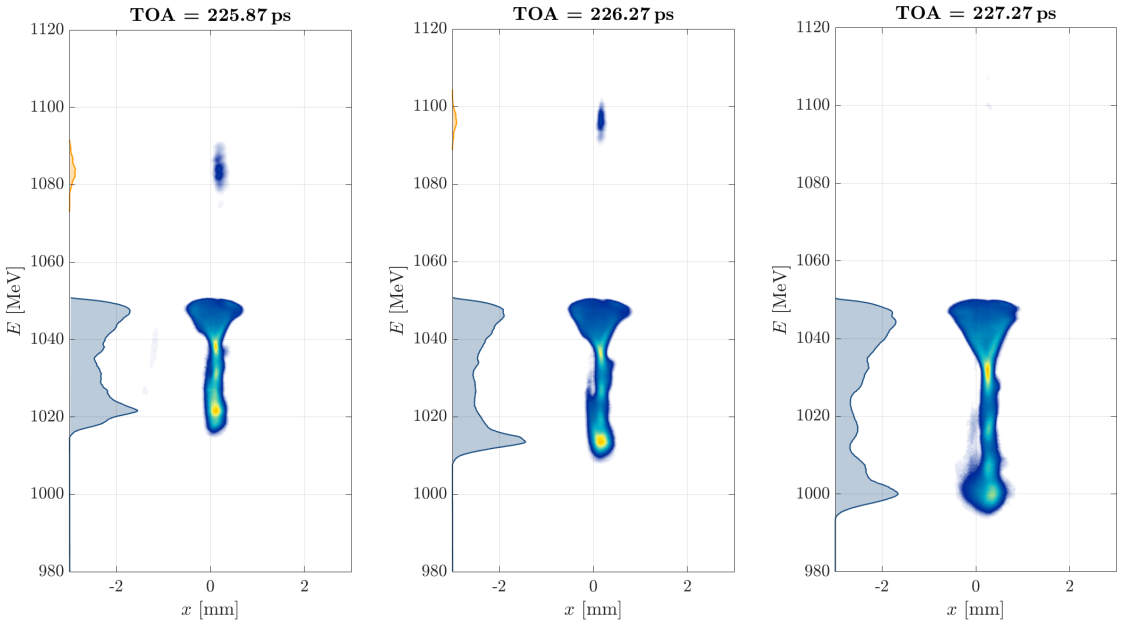

(a) (b) (c)

Figure 8.6.2: Average spectrometer images for Laser 1 time-of-arrivals (TOAs) of (a) 225.87 ps, (b) 226.27 ps and (c) 227.27 ps *with* plasma interaction. For TOA = 227.27 ps the low charge witness bunch is lost during plasma interaction. All images use the same colour scale; a factor of 10 lower than that used in Figure 8.6.1. The projections in the dispersive plane are represented by blue and orange regions for the drive and witness bunches respectively.

is shorter than this at the interaction point due to longitudinal expulsion of the plasma out of the open ends of the capillary. This was studied in detail in Ref. [195] and it was found that the longitudinal plasma profile could be well approximated in simulation by a flat-top density profile with length close to the distance between the gas inlets, 32 mm [68], as discussed in the previous section. In this case the average accelerating gradients are approximately $1\,\mathrm{GVm^{-1}}$ and $1.3\,\mathrm{GVm^{-1}}$ for TOAs of 225.87 ps and 226.27 ps respectively. To allow comparison between all three TOAs, the longitudinally-averaged maximal decelerating gradient (E_{dec}) was calculated for each setup, with values of 1.11, 1.32, and $1.74\,\mathrm{GVm^{-1}}$ calculated for TOAs of 225.87, 226.27, and 227.27 ps respectively. The plasma density at the interaction point was controlled by varying the arrival time of the bunch with respect to the trigger that initiated the discharge and was kept constant throughout these measurements.

The second probe bunch pair followed the leading bunch pair at variable separation ($\Delta\tau$) and drove a wakefield within the perturbed plasma state. Its current profile is not affected by changes to the TOA of Laser 1 (L1) as a separate photocathode laser (Laser 2) is used to produce it. The resulting probe bunch energy spectra are shown in Figures 8.6.3 and 8.6.4 for two different quadrupole focusing energies. Figure 8.6.3 uses a quadrupole focusing energy of 1040 MeV to better image the energy spectrum of the drive bunch following its interaction with the plasma whereas Figure 8.6.4 uses a focal energy of 1100 MeV to better image the witness bunch. In all subsequent analysis, calculations relating to the drive bunch spectrum use data from the 1040 MeV imaging energy measurements and calculations relating to the witness bunch spectrum use data from the 1100 MeV imaging energy measurements as these should be more accurate for the relevant bunch.

The first panels of these figures, (a), correspond to the measurement of the energy spectrum of the probe bunch pair after propagation through the plasma without any preceding bunch, referred to as a cold or unperturbed plasma. The witness bunch mean energy can be observed to slowly decrease from $\mu_{E,w} = 1098.4 \pm 0.8\,\mathrm{MeV}$ to $\mu_{E,w} = 1093.5 \pm 0.8\,\mathrm{MeV}$ over this time period. This is caused by the recombination of the plasma; as the relative arrival time of the bunch is moved later, the plasma density with which it interacts slightly decreases, resulting in a reduced wakefield strength and hence lower accelerating gradient. A corresponding effect is seen in the drive bunch energy spectrum with its mean energy slightly increasing from $\mu_{E,d} = 1021.5 \pm 1.0\,\mathrm{MeV}$ to $\mu_{E,d} = 1022.2 \pm 0.9\,\mathrm{MeV}$. As discussed previously, the witness bunch energy spectrum is more sensitive to the change in plasma density as it is short with respect to the wakefield cavity length and only samples a small region of the longitudinal wakefield, hence is highly sensitive to shifts in its relative phase incurred by the increasing cavity length as the plasma density decreases. In contrast, the drive bunch samples the entirety of the decelerating phase of the wakefield so is less sensitive to small density changes, although differences in the energy spectrum are still observed over the 47 ns measurement period.

Panels (b), (c) and (d) in Figures 8.6.3 and 8.6.4 show the evolution of the energy spectrum of the probe bunch for increasing wakefield interaction strength of the leading drive-witness pair. In all three cases, initially ($\Delta\tau < 2\,\mathrm{ns}$) the drive bunch energy loss is observed to increase

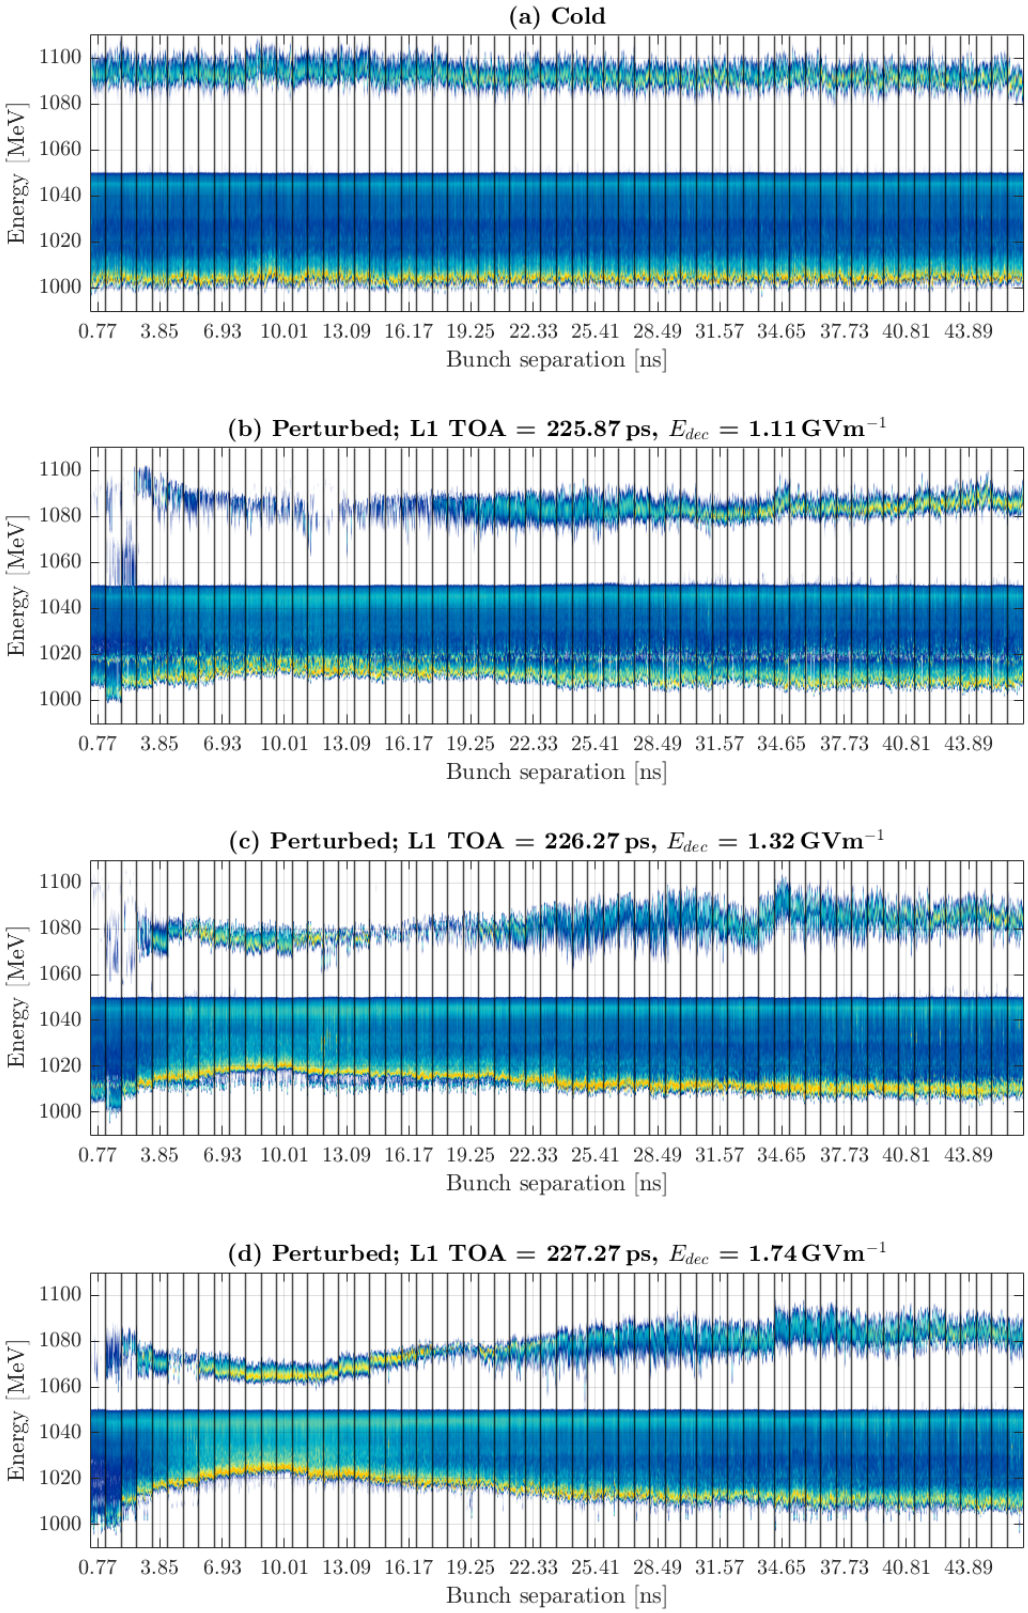

Figure 8.6.3: Waterfall spectrometer images showing the evolution of the probe bunch energy profile as a function of separation with respect to the leading bunch for increasing drive wakefield strengths. The post-plasma quadrupoles are set to focus at 1040 MeV, imaging the decelerated probe drive bunch. (a) Probe bunch only. (b) L1 TOA = 225.87 ps, $E_{dec} = 1.11\,\mathrm{GVm}^{-1}$. (c) L1 TOA = 226.27 ps, $E_{dec} = 1.32\,\mathrm{GVm}^{-1}$. (d) L1 TOA = 227.27 ps, $E_{dec} = 1.74\,\mathrm{GVm}^{-1}$.

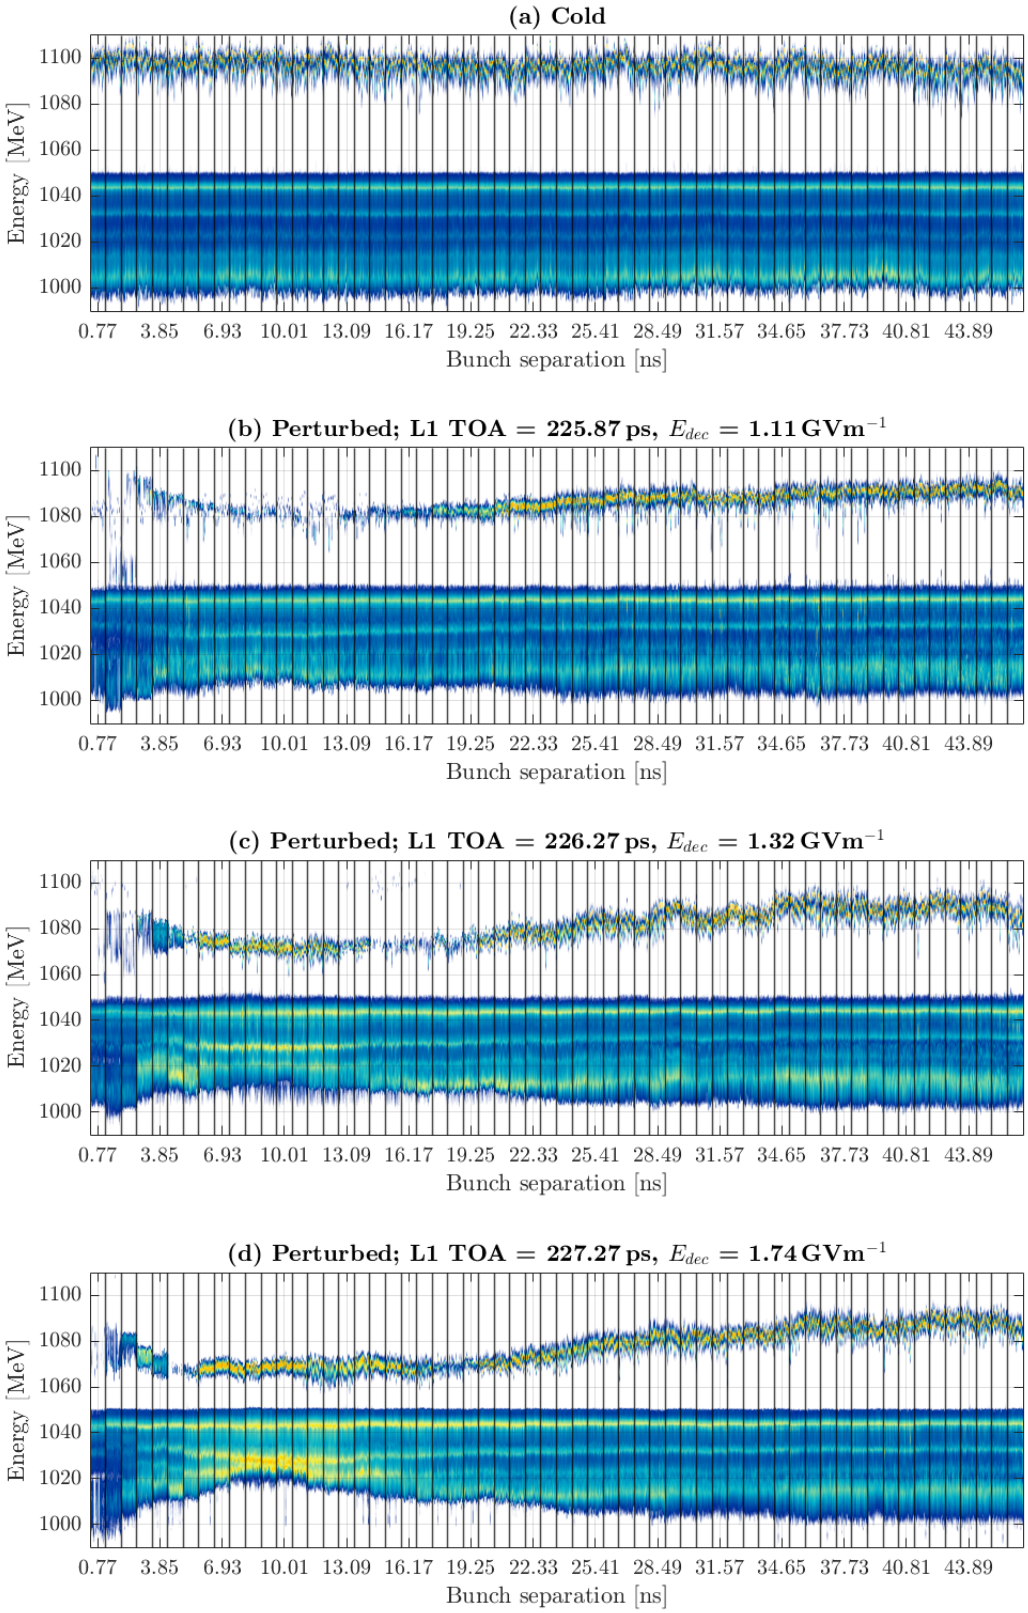

Figure 8.6.4: Waterfall spectrometer images showing the evolution of the probe bunch energy profile as a function of separation with respect to the leading bunch for increasing drive wakefield strengths. The post-plasma quadrupoles are set to focus at 1100 MeV, imaging the accelerated probe witness bunch. (a) Probe bunch only. (b) L1 TOA = 225.87 ps, $E_{dec} = 1.11\,\mathrm{GVm}^{-1}$. (c) L1 TOA = 226.27 ps, $E_{dec} = 1.32\,\mathrm{GVm}^{-1}$. (d) L1 TOA = 227.27 ps, $E_{dec} = 1.74\,\mathrm{GVm}^{-1}$.

while a large fraction of the witness bunch is lost from the spectrometer signal. This is indicative of an initial increase in on-axis plasma density experienced by the probe bunch pair. This is expected as ions within the wakefield cavity that are initially close to the propagation axis will be attracted towards the axis where large electron densities are induced by the return of electrons at the rear of the wakefield cavities. The increase in on-axis density acts to shorten the wakefield cavity driven by the probe bunch pair such that significant portions of its witness bunch extend beyond the cavity length and experience transversely-defocusing fields due to the return of plasma electrons at the rear of the cavity. This witness charge is strongly defocused, is not successfully transported through the length of the plasma, and hence cannot be captured by the spectrometer quadrupoles.

Over the following few nanoseconds, the on-axis ion density begins to decrease as ions continue to propagate transversely. They are driven by the thermal pressure gradient present in the plasma induced by the energy that is locally deposited near the axis via the wakefield driven by the leading drive-bunch pair. This reduction in on-axis density reduces the interaction strength between the probe bunch pair and the plasma, resulting in a reduction in the energy loss (gain) of the drive (witness) bunch. As can be seen in Figure 8.6.3, the perturbative effect to the drive bunch increases with increasing wakefield strength (i.e. from L1 TOAs of (b) 225.87 ps, $E_{dec} = 1.11\,\mathrm{GVm}^{-1}$ → (d) 227.27 ps, $E_{dec} = 1.74\,\mathrm{GVm}^{-1}$). The same effect is observed in the witness bunch spectra shown in Figure 8.6.4, with the largest perturbations to the witness bunch energy spectra, when compared to the unperturbed plasma measurements, observed for the largest wakefield strength. This is expected as the motion of ions is driven by two main effects:

(i) Initially, over the first few wakefield periods, ions are attracted to regions of high plasma electron density in the wakefield and are driven by the time-averaged, large amplitude, transverse fields. If the plasma is more strongly perturbed by the leading bunch, larger electron densities are observed at the sheath and rear of the cavity, larger amplitude wakefields are driven, and ions are hence accelerated more quickly.

(ii) Following phase-mixing of plasma electrons, the continued motion of ions is driven by the thermal pressure gradient present in the plasma. As the leading bunch pair propagate through the plasma, they transfer a significant amount of energy into the plasma via the fields excited in it, the majority of which is transferred to plasma electrons. These high-energy electrons undergo collisions with both ions and electrons within the plasma, thermalising and cooling, while heating the bulk of the plasma (both electrons and ions) within the vicinity of the wakefield. The local increase in temperature, combined with the density gradient excited by the wakefield itself, leads to the generation of a significant pressure gradient in the plasma that drives its redistribution as discussed in Section 8.2. A stronger wakefield interaction, i.e. increased energy transfer between the leading bunch pair and plasma, drives larger amplitude pressure gradients and more significant perturbations to the density profile.

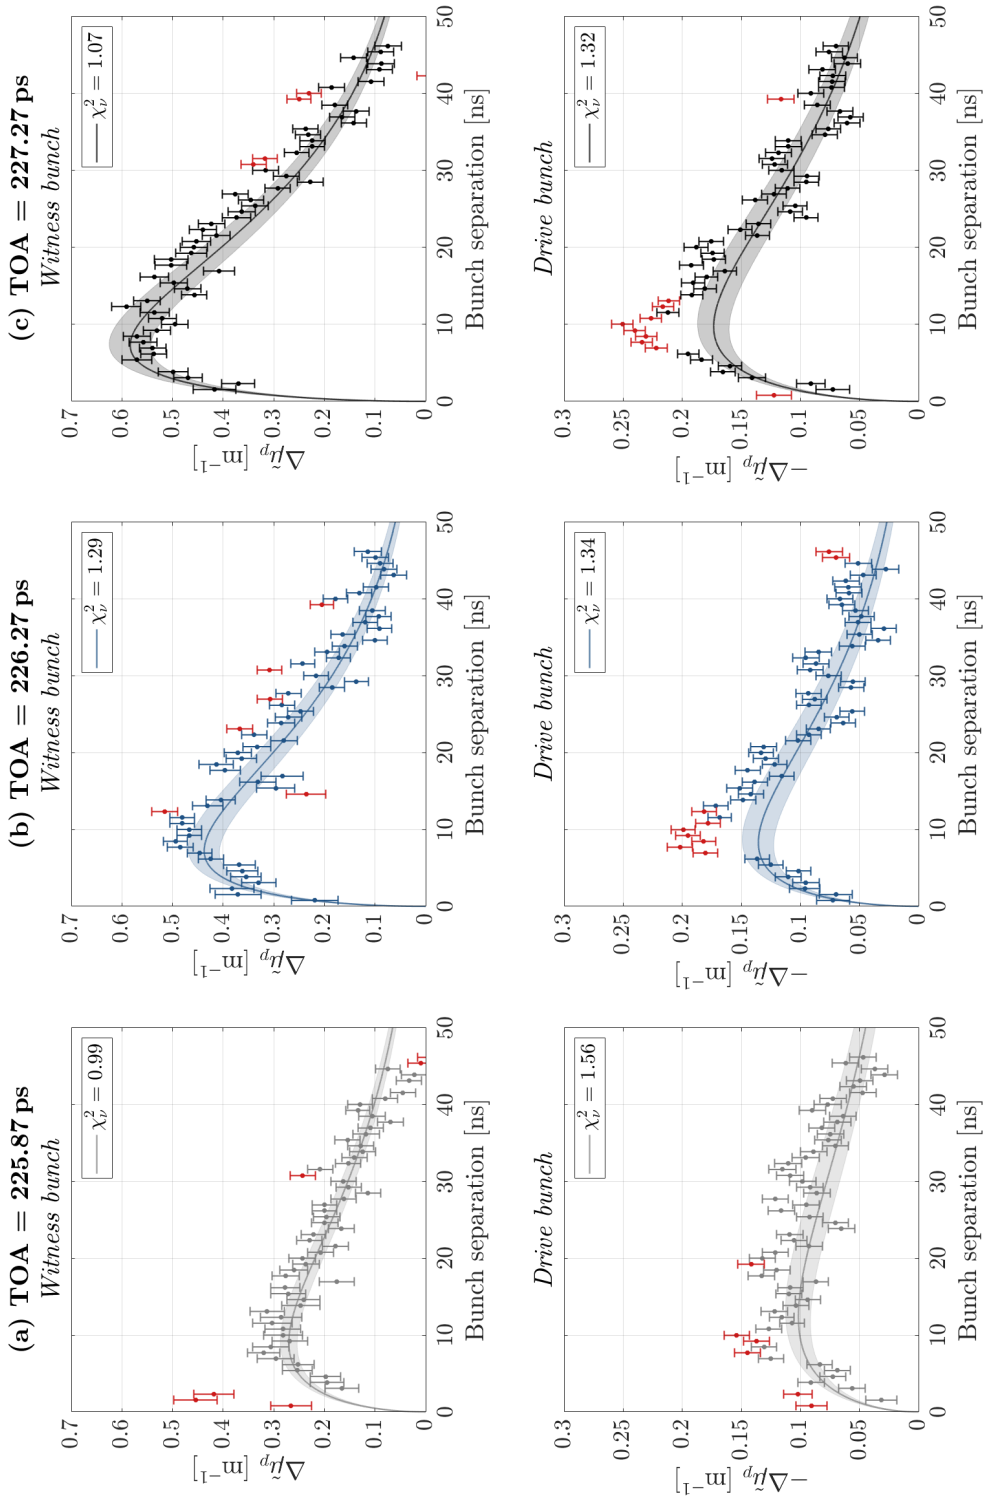

Figure 8.6.5: Evolution of the perturbation to the mean energy of the probe witness (upper) and drive (lower) bunches for varying wakefield strengths. L1 TOAs of (a) 225.87 ps, $E_{dec} = 1.11\,GVm^{-1}$, (b) 226.27 ps, $E_{dec} = 1.32\,GVm^{-1}$, and (c) 227.27 ps, $E_{dec} = 1.74\,GVm^{-1}$. Each dataset is fit to the function $f(t) = At^{1/2}\exp(-Bt)$ shown by the solid line with uncertainties on the fit represented by the shaded bands. The calculated reduced chi-squared statistic, χ^2_ν, for each fit is shown in the upper right of each panel. Outliers are plotted in red.

L1 TOA [ps]	E_{dec} [GVm^{-1}]	Drive		Witness	
		A [10^{-3}]	B [10^{-3}]	A [10^{-3}]	B [10^{-3}]
225.87	1.11	47.6 ± 3.4	40.3 ± 2.6	148.5 ± 7.5	55.5 ± 1.2
226.27	1.32	77.6 ± 5.8	60.4 ± 2.8	267.8 ± 14.5	69.0 ± 2.2
227.27	1.74	91.6 ± 5.2	51.4 ± 2.0	356.4 ± 21.1	68.6 ± 1.5

Table 4: Calculated parameters A and B for fits to the relative perturbation, $\Delta\tilde{\mu}_p$, to the function $f(t) = At^{1/2}\exp(-Bt)$ for the three measured L1 TOAs.

The effect of this can be quantified by comparing the relative perturbation to the drive and witness bunches for the three different values of the L1 TOA as outlined in Section 8.5.3. The relative perturbation is calculated by comparing the drive and witness mean energy in the cold and perturbed measurement schemes via Eq. (8.5.4). The results are shown in Figure 8.6.5 for all three values of the L1 TOA. Measurements of the drive bunch perturbation use values of the mean energy of the drive bunch calculated for an imaging energy of 1040 MeV while the witness bunch perturbation uses an imaging energy of 1100 MeV to improve the accuracy of the measurements in each case.

Clearly, the strength of the perturbation grows with increasing wakefield interaction of the leading bunch pair and a significantly larger perturbation is consistently observed for the witness bunch when compared to the drive bunch for reasons discussed in more detail later. The measurements of the relative perturbation are fitted to the function

$$f(t) = At^{1/2}\exp(-Bt), \tag{8.6.1}$$

with each individual measurement, i, weighted by the square of its relative uncertainty, i.e. $(\sigma_i/\Delta\tilde{\mu}_{p,i})^2$. The uncertainty on each individual measurement, σ_i, combines the uncertainty on the measurements of the mean energy of the relevant bunch in both the perturbed and unperturbed plasma states. The function, $f(t)$, is chosen as it empirically matches the evolution of the position of the zero-crossing within the longitudinal wakefield found in the simulations discussed in Section 8.5.4. The two fitting parameters, A and B, can be considered to correspond to the magnitude of the on-axis density reduction as the ion channel forms (A), and the rate of formation and dissipation of the non-uniform structure (B). The extracted parameter values are presented in Table 4 with the reduced chi-squared statistic, χ^2_ν, for each fit given in the legends in Figure 8.6.5. Data points plotted in red in Figure 8.6.5 represent those whose value exceeds three times the combined uncertainty of the individual measurement and the uncertainty of the fitting function at the relevant bunch separation and hence are considered to be outliers and are not included in the fit.

Increases in the wakefield interaction strength via adjustments of the L1 TOA lead to significant increases in the A parameter, suggesting larger on-axis ion density perturbations as expected. However, less significant changes in the timescale of the evolution of the per-

turbation are observed over the measured parameter range. The peak perturbation occurs at

$$f'(t) = \frac{1}{2}At^{-1/2}\exp(-Bt) - ABt^{1/2}\exp(-Bt) = 0, \tag{8.6.2}$$

$$\Rightarrow t = \frac{1}{2B}. \tag{8.6.3}$$

As can be seen in Figure 8.6.5, the fit function $f(t)$ does not recreate the strongly-peaked perturbation to the drive bunch energy spectra at $\Delta\tau \sim 10\,\text{ns}$, leading to an underestimation of the peak perturbation to the drive bunch (A), however the timescale over which it is measured to dissipate (B) appears to be well represented by the fit. In general, the fit of $f(t)$ to the measured energy perturbation to the witness bunch is good ($\chi^2_\nu \sim 1$) and approximately recreates both the peak perturbation and rate of dissipation. The witness bunch is most sensitive to the perturbative effect as it is short and hence samples only a limited phase of the wakefield, while the wakefield it experiences is a sum over any perturbative effect experienced by each slice of the drive bunch. In contrast, the drive bunch experiences a large fraction of the entire wakefield and hence a measurement of only its mean energy has the effect of averaging over any perturbative effect. For this reason, return of the witness bunch to its non-perturbed state is a more effective measure of the perturbative lifetime.

It is expected that the perturbation should be longer-lasting for lower interaction strengths. This is because the motion of plasma ions over extended timescales should be driven by the pressure gradient present within the plasma after the leading interaction. As discussed previously, the magnitude of this pressure gradient should be reduced in the case of lower interaction strengths as less energy is transferred from the leading bunch pair to the plasma. This appears to be the case in these measurements when comparing the lowest interaction strength (L1 TOA = 225.87 ps, $E_{dec} = 1.11\,\text{GVm}^{-1}$) to the others measured as it has a significantly lower value of the fitting parameter B, particularly for the drive bunch measurements, indicating slower generation and dissipation of the perturbed density profile. However, the values of B measured for the witness bunches at L1 TOAs of 226.27 ps ($E_{dec} = 1.32\,\text{GVm}^{-1}$) and 227.27 ps ($E_{dec} = 1.74\,\text{GVm}^{-1}$) agree with each other within their uncertainties, suggesting similar perturbative lifetimes despite the large difference in the magnitude of their perturbation. In future, measurements must be performed over increased timescales to confirm return of the witness bunch energy spectrum to that of the unperturbed plasma state and demonstrate relaxation of the plasma. This may enable more conclusive statements to be made about the dependency of the perturbative lifetime on the wakefield strength.

The evolution of the relative perturbation to the transverse size of the witness beam in the plane of the spectrometer for the three L1 TOAs is shown in the upper panels of Figure 8.6.6 for both measured imaging energies, calculated according to Eq. (8.5.1). In both cases, the transverse bunch size is compared to the transverse size measured in the unperturbed plasma at an imaging energy of 1100 MeV. As discussed in Section 8.5, the change in witness beam size has been shown in simulation to be mostly a function of the evolution

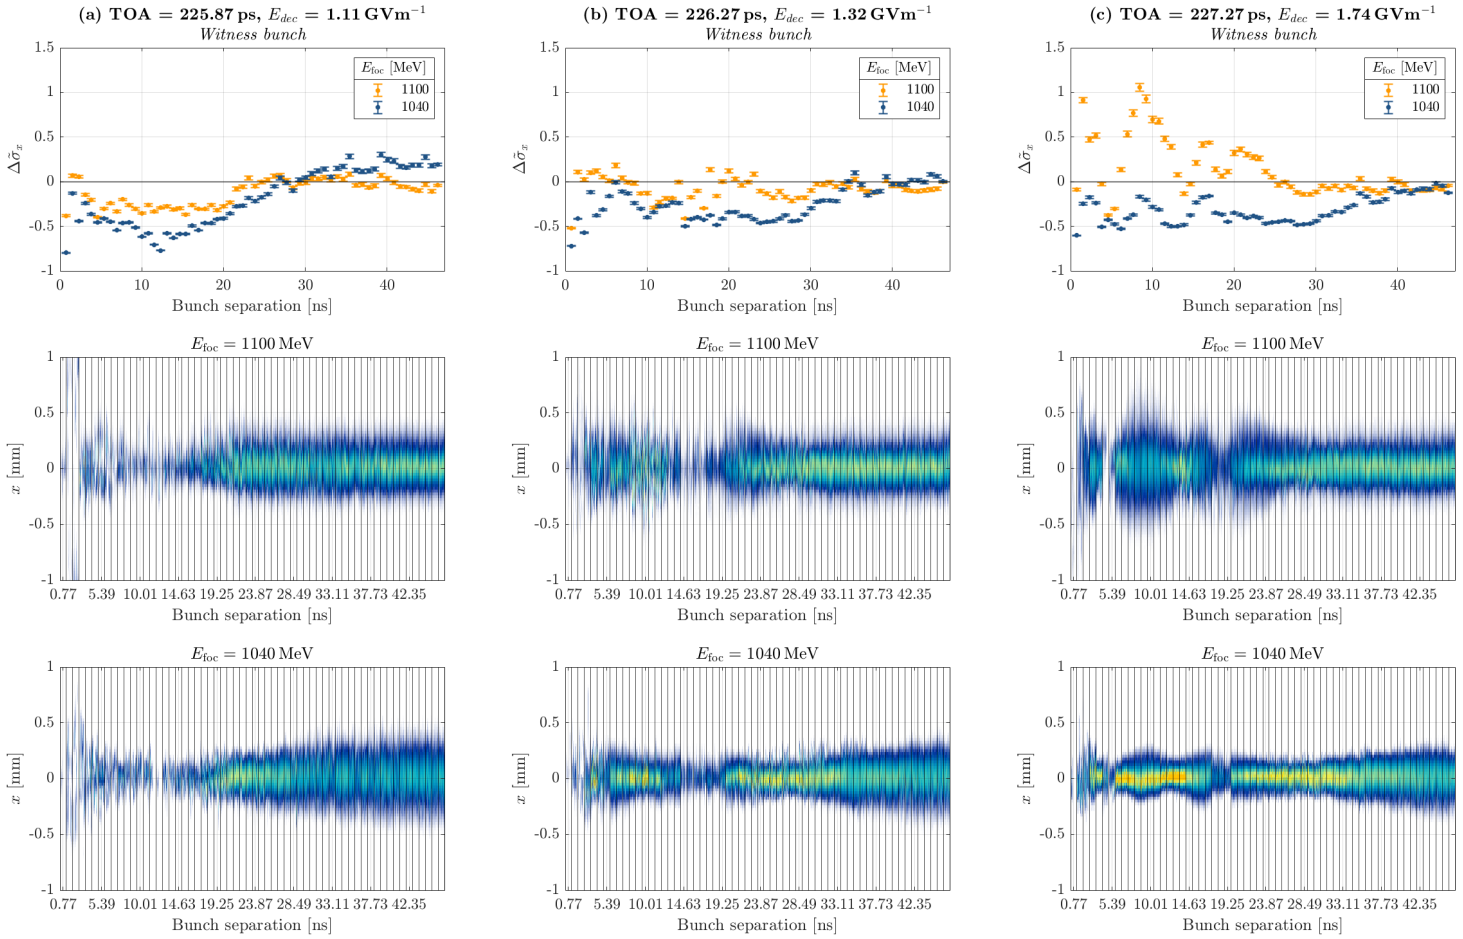

Figure 8.6.6: Evolution of the relative perturbation to the transverse witness beam size for varying wakefield strengths, calculated via Eq. (8.5.1). L1 TOAs of (a) 225.87 ps, $E_{dec} = 1.11\,\mathrm{GVm}^{-1}$, (b) 226.27 ps, $E_{dec} = 1.32\,\mathrm{GVm}^{-1}$, and (c) 227.27 ps, $E_{dec} = 1.74\,\mathrm{GVm}^{-1}$. The lower panels show waterfall projections of the witness beam size for imaging energies of 1100 MeV and 1040 MeV respectively.

of the on-axis plasma density. For the strongest perturbation ($E_{dec} = 1.74\,GVm^{-1}$), Figure 8.6.6(c), three bunch size oscillations are observed over the measurement period as the probe bunch pair experiences a significant range of on-axis densities. The peak perturbation to the energy spectrum, and hence likely the lowest on-axis density experienced by the bunch, occurs at approximately 10 ns, with the on-axis density increasing from there onwards. At timescales shorter than this, the ion channel is expected to be rapidly deepening and, particularly over the first few ns, large variations in the amount of witness charge transported to the spectrometer lead to large fluctuations in the measurements of the relative bunch size.

As the leading wakefield interaction strength is reduced by reducing the L1 TOA, the perturbative ion motion effect similar reduces for reasons discussed previously. This is mirrored in the measurements of the relative transverse bunch size shown in Figure 8.6.6(a) and (b) as the number of oscillations in bunch size is reduced. For a L1 TOA of 226.27 ps ($E_{dec} = 1.32\,GVm^{-1}$), Figure 8.6.6(b), one complete oscillation of the witness bunch size is measured over the time period $\Delta\tau \sim 8 - 40$ ns, whereas this is further reduced for a L1 TOA of 225.87 ps ($E_{dec} = 1.11\,GVm^{-1}$) over the measured period [Figure 8.6.6(a)]. This reduction in the number of witness bunch size oscillations is indicative of a decrease in the range of on-axis plasma densities experienced by the probe bunch, consistent with a reduced perturbative effect. In all three cases, the relative perturbation to the transverse size of the witness beam is close to zero at the end of the measurement period indicating the on-axis density is close to that measured in the cold measurement scheme after approximately 50 ns, consistent with the analysis of the energy perturbation to the probe bunches.

Figure 8.6.7 shows the fits to the relative energy perturbation shown in Figure 8.6.5 extrapolated over 100 ns and illustrates the similarity of the perturbative lifetime for all three L1 TOAs measured in these datasets. In all cases, the perturbative lifetime appears to be very similar and equal to approximately 100 ns. Therefore, while the magnitude of the perturbation shows a clear dependency on the amplitude of the wakefield, the dependence of its lifetime is less clear within these measurements. It is possible that the range of wakefield amplitudes investigated in this measurement was not sufficiently wide to see significant effects on the perturbative lifetime and this should be investigated further in future measurements in an attempt to demonstrate the expected dependency.

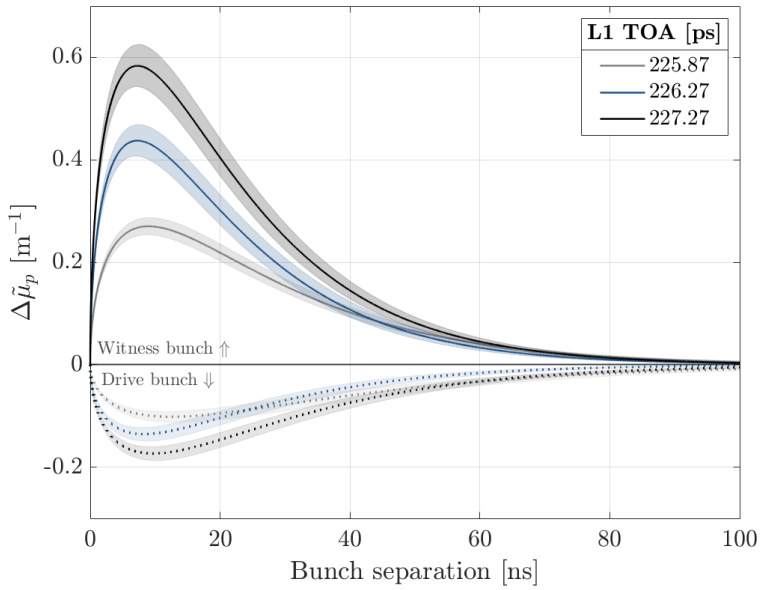

Figure 8.6.7: Comparison of the fits to the measured relative perturbation in Figure 8.6.5 extrapolated over 100 ns. Solid lines represent the witness bunch perturbation and dotted lines represent the drive bunch perturbation.

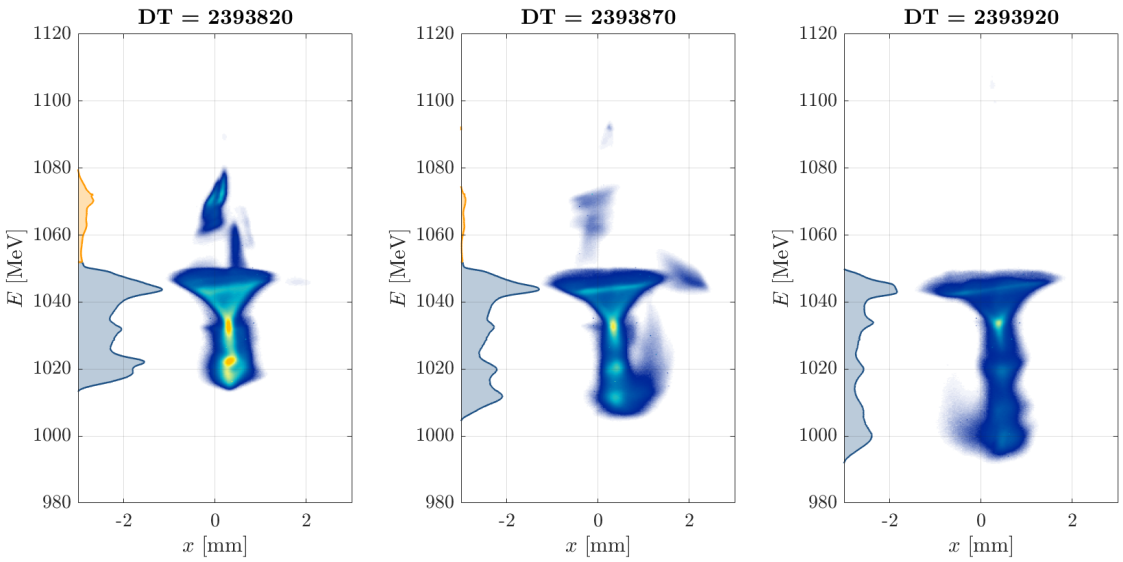

(a) (b) (c)

Figure 8.6.8: Average spectrometer images of Laser 1 for discharge timings (DT) of (a) 2393820 ($n_{b0}/n_0 \sim 10$), (b) 2393870 ($n_{b0}/n_0 \sim 5$) and (c) 2393920 ($n_{b0}/n_0 \sim 3$). An L1 TOA of 227.27 ps was used in all datasets, giving a peak bunch density of approximately $n_{b0} \sim 5 \times 10^{16}\,\mathrm{cm}^{-3}$. All images use the same colour scale. The bunch projections in the dispersive plane are represented by blue and orange regions for the drive and witness bunches respectively.

8.6.2 *Plasma density*

The second dependency investigated was that of the ion motion signal on the interaction density. This was controlled experimentally by varying the timing of the discharge trigger with respect to the arrival of the leading bunch pair. By increasing this delay, the plasma undergoes additional recombination such that the interaction density has decreased. The L1 TOA was kept constant throughout these measurements at 227.27 ps resulting in a leading bunch of charge approximately 590 pC, RMS length $\sigma_z = 195\,\mu\mathrm{m}$, and RMS transverse size of approximately 5 µm — a peak density of approximately $n_{b0} \sim 5 \times 10^{16}\,\mathrm{cm}^{-3}$ assuming Gaussian longitudinal and transverse distributions. Three discharge delay timings were used; (a) 2393820, (b) 2393870 and (c) 2393920, where this number represents a digital delay applied to the discharge trigger. The trigger has an associated frequency of 108.3 MHz and hence each individual step represents 9.23 ns, i.e. 50 step changes as used here represent additional delays of 461.5 ns between the firing of the discharge and the arrival of the leading bunch pair. Lower discharge delay numbers represent firing the discharge earlier and so correspond to lower plasma densities at interaction. The corresponding plasma densities were measured *in situ* using the spectral line broadening technique discussed in Section 5.1.2, giving approximate plasma densities of (a) $(4.93 \pm 0.69) \times 10^{15}\,\mathrm{cm}^{-3}$, (b) $(9.34 \pm 1.30) \times 10^{15}\,\mathrm{cm}^{-3}$ and (c) $(1.77 \pm 0.25) \times 10^{16}\,\mathrm{cm}^{-3}$ respectively. The large relative uncertainties in these measure-

ments are due to difficulties collecting a significant amount of light from the plasma meaning the measured density range was close to the resolution limit of the diagnostic. These three working points therefore represented relative operating plasma densities of approximately $n_{b0}/n_0 \sim$ (a) 10, (b) 5 and (c) 3.

These discharge timing (DT) values were chosen as they produced stable interactions between the probe bunch pair and the plasma. The corresponding average spectrometer images of the interaction between the leading bunch and the plasma are shown in Figure 8.6.8. Reducing the interaction density reduces the amplitude of the wakefield driven by the bunch, leading to reduced energy loss for the drive bunch and reduced energy gain for the witness bunch as expected. Again, comparing the maximal energy loss of the drive bunch in each case to its energy without plasma interaction gives longitudinally-integrated maximal decelerating gradients of 1.16, 1.39, and 1.74 GVm^{-1} for discharge delay timings of 2393820 ($n_{b0}/n_0 \sim 10$), 2393870 ($n_{b0}/n_0 \sim 5$), and 2393920 ($n_{b0}/n_0 \sim 3$) respectively, very close to those measured in the wakefield strength scan discussed in Section 8.6.1. This similarity enables comparison of the relative effects of changes to the wakefield strength and plasma density on the ion motion perturbation and is discussed in more detail in the following section.

Measurements of the probe bunch energy spectra are shown in Figures 8.6.9, 8.6.10 and 8.6.11 for discharge delay timings of 2393820 ($n_{b0}/n_0 \sim 10$), 2393870 ($n_{b0}/n_0 \sim 5$) and 2393920 ($n_{b0}/n_0 \sim 3$) respectively. Unlike the scan of the wakefield interaction strength presented in Section 8.6.1 where the interaction strength of only the leading bunch pair was changed, by changing the plasma density at the point of interaction the cold probe bunch measurement scheme is also affected in each case. For this reason, the corresponding cold scheme energy spectra are shown in panel (a) of each measurement for reference. These measurements were all performed with the spectrometer quadrupoles set to image close to the mean energy of the probe witness bunch in order to better measure its energy spectrum; 1070 MeV, 1085 MeV and 1100 MeV for discharge delay timings of 2393820 ($n_{b0}/n_0 \sim 10$), 2393870 ($n_{b0}/n_0 \sim 5$) and 2393920 ($n_{b0}/n_0 \sim 3$) respectively. While the chosen discharge delay timings resulted in stable interaction between the probe bunch pair and the plasma, the interaction between the leading bunch pair was far less stable and made automatic identification of the probe bunch signal much more challenging, particularly in the case of the probe witness bunch. This is particularly apparent in Figure 8.6.9(b) at the shortest bunch separations where occasionally portions of the leading witness bunch were not completely removed during the signal subtraction procedure outlined in Section 8.4.1 and these remaining regions of signal (around 1060 MeV) were incorrectly identified as the probe witness bunch. For future measurements, a more sophisticated subtraction routine based on the identification of "optimal" individual shots that can be used in the subtraction process is under development, and additionally more care can be taken during the experiment to ensure that stable spectrometer signals are observed for both the leading and probe bunch pairs.

Despite the imperfect subtraction procedure, clear dependencies of the ion motion signatures on the interaction density can be extracted from the measurements. The most obvious

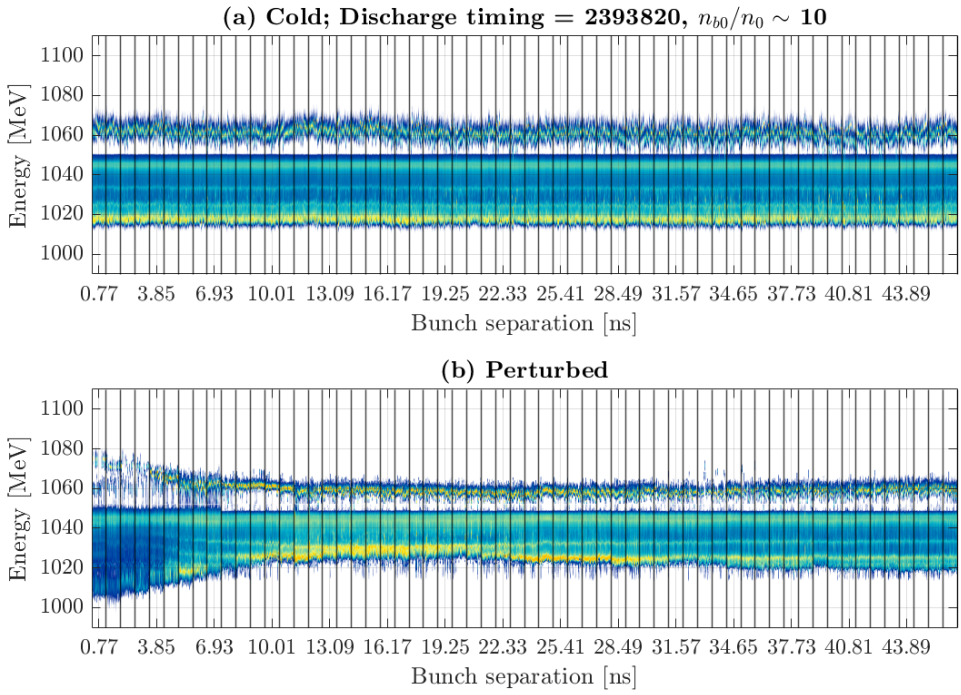

Figure 8.6.9: Waterfall spectrometer images showing the evolution of the probe bunch energy profile as a function of separation with respect to the leading bunch for a discharge timing delay of 2393820 ($n_{b0}/n_0 \sim 10$). The post-plasma quadrupoles are set to focus at 1070 MeV, imaging the accelerated probe witness bunch. (a) Probe bunch only. (b) Probe bunch following perturbation from leading bunch.

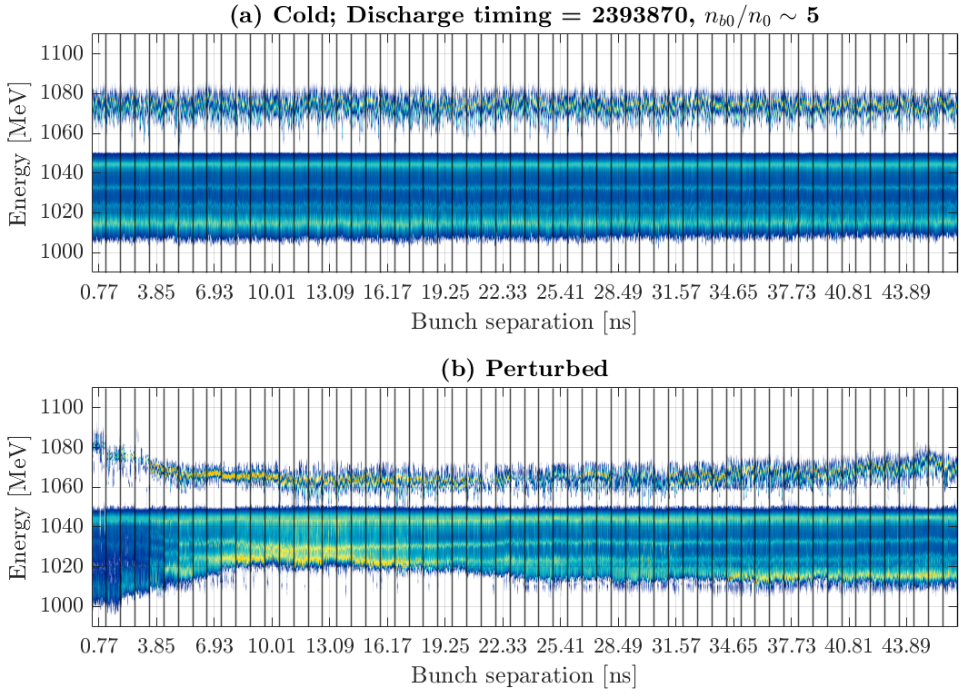

Figure 8.6.10: As in Figure 8.6.9 for a discharge timing delay of 2393870 ($n_{b0}/n_0 \sim 5$). The post-plasma quadrupoles are set to focus at 1085 MeV.

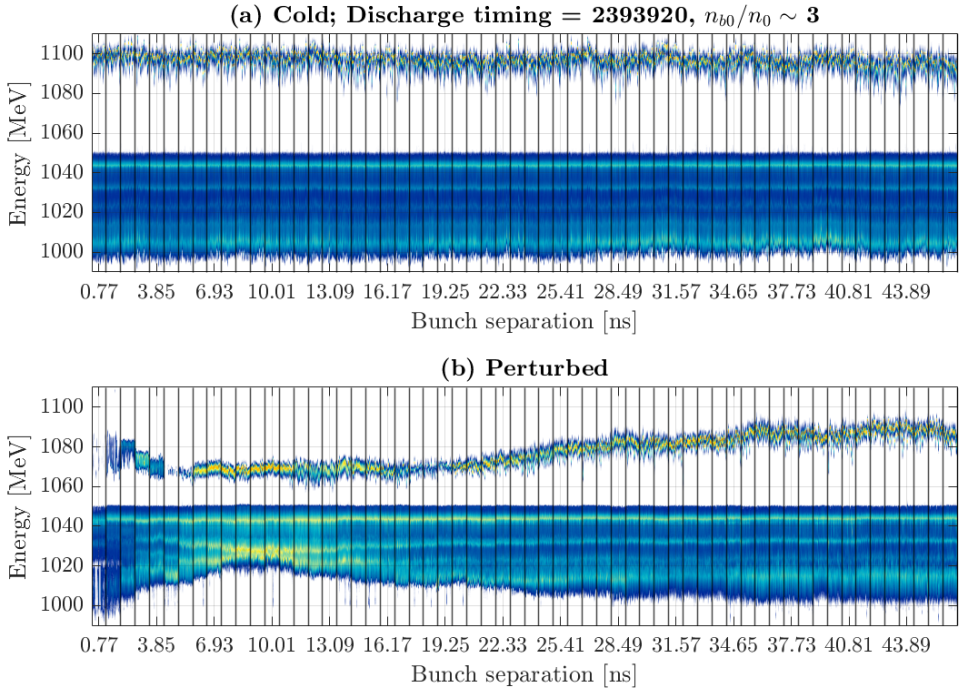

Figure 8.6.11: As in Figure 8.6.9 for a discharge timing delay of 2393920 ($n_{b0}/n_0 \sim 3$). The post-plasma quadrupoles are set to focus at 1100 MeV.

of these is the timescale over which the perturbations are observed, with the perturbative lifetime appearing to extend significantly for lower plasma densities. During the discussion of the wakefield strength dependence in Section 8.6.1, it was noted that at the very shortest bunch separations the probe bunch appeared to experience larger on-axis plasma densities from ions that are initially attracted towards the axis by the large electron densities that are induced at the extrema of the wakefield cavities. This effect is exacerbated by moving to lower interaction densities. For a discharge delay timing of 2393870 ($n_{b0}/n_0 \sim 5$, Figure 8.6.10), for at least the first three nanoseconds the minimum energy of the drive bunch is lower in the perturbed measurement scheme than in the cold scheme, indicating an increase in the on-axis plasma density. This is mirrored in the probe witness bunch energy where the mean energy of the fraction of the witness bunch that is successfully transported to the spectrometer in the perturbed measurement scheme exceeds that of the cold scheme over the same time period. As the plasma density is lowered further still at a discharge delay timing of 2393820 ($n_{b0}/n_0 \sim 10$, Figure 8.6.9), the perturbed on-axis density appears to be above that of the cold measurement scheme for approximately 5 ns before the ion channel begins to form and the on-axis density reduces. The deepening of the ion channel can be observed via the position of the focal lines within the probe drive bunch energy spectrum. Between bunch separations of 4 ns and 13 ns the focal lines can be seen to gradually shift lower in energy, indicating a growth in the curvature of the ion channel as the on-axis density reduces. A similar effect is also observed for the other interaction densities on reduced timescales. At the lowest interaction density ($n_{b0}/n_0 \sim 10$, Figure 8.6.9), the maximum perturbation to the probe bunch

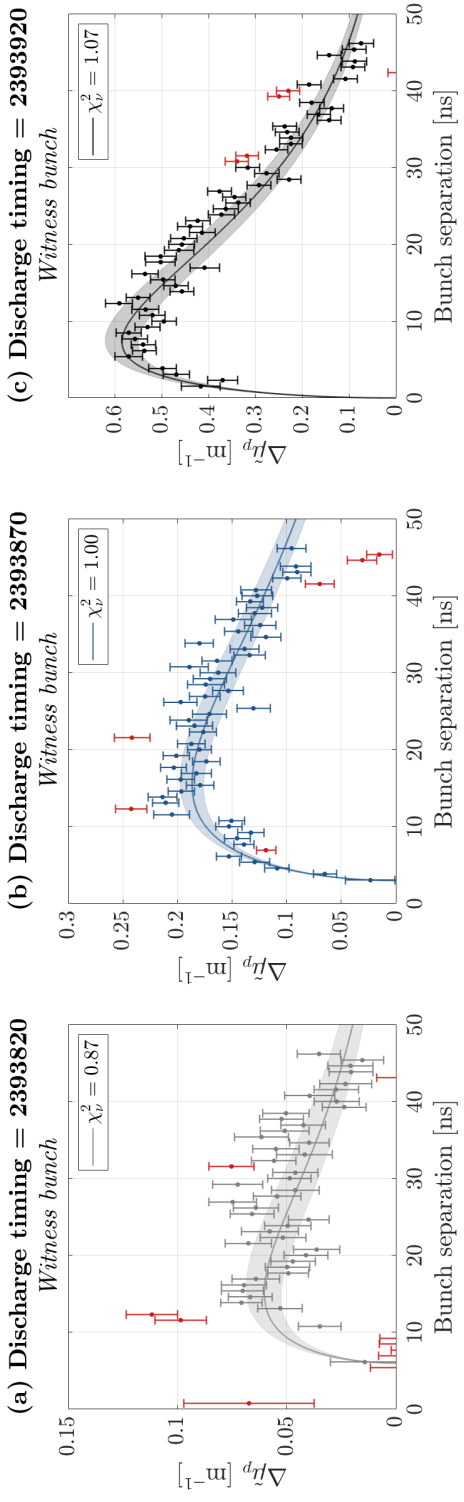

Figure 8.6.12: Evolution of the perturbation to the mean energy of the probe witness bunches for varying plasma densities at interaction. Discharge delay timings of (a) 2393820 ($n_{b0}/n_0 \sim 10$), (b) 2393870 ($n_{b0}/n_0 \sim 5$) and (c) 2393920 ($n_{b0}/n_0 \sim 3$). Each dataset is fit to the function $f(t) = At^{1/2}\exp(-Bt)$ shown by the solid line with uncertainties on the fit represented by the shaded bands. The calculated reduced chi-squared statistic, χ^2_ν, for each fit is shown in the upper right of each panel. Outliers are plotted in red. Note the change in y-axis scale between each panel as the magnitude of the perturbation varies significantly.

energy spectrum occurs at approximately 18 ns. As the interaction density is increased the time of maximum perturbation decreases, occurring at approximately 14 ns for a discharge delay of 2393870 ($n_{b0}/n_0 \sim 5$) and approximately 9 ns for a delay of 2393920 ($n_{b0}/n_0 \sim 3$).

The evolution of the relative perturbation to the probe witness bunch for each interaction density is plotted in Figure 8.6.12. Similar measurements of the perturbation to the drive bunch do not exist as measurements with the spectrometer imaging set to properly resolve the drive bunch energy spectra were not performed due to time constraints. Here, the effects of the imperfect subtraction procedure can be clearly observed with a larger witness energy variability within datasets leading to a larger relative uncertainty on measurements of the witness mean energy in the perturbed plasma scheme. This propagates to a larger uncertainty on $\Delta\tilde{\mu}_p$ at the lower plasma densities shown in Figure 8.6.12(a) and (b) when compared to similarly small values of $\Delta\tilde{\mu}_p$ in Figure 8.6.5 during the wakefield amplitude scan in Section 8.6.1 and highlights the need for a more sophisticated image subtraction procedure.

The probe witness bunch measurements are again fit to Eq. (8.6.1), however the fitting begins at the time of the first positive value of $\Delta\tilde{\mu}_p$ to remove the earliest data points where the on-axis density is higher than in the cold measurement scheme. This analysis illuminates two effects on the perturbation induced by changes in the interaction density: (i) the lifetime of the perturbation is reduced for higher interaction densities as discussed before, and (ii) the magnitude of the perturbation is significantly larger at higher interaction densities. Similar to the discussion of the dependency of the ion motion signature on the wakefield interaction strength in Section 8.6.1, this can be explained by consideration of the pressure gradient that is induced in the plasma by the leading bunch pair. At the highest density considered here, the ability of the leading bunch pair to effectively couple to the plasma is increased, resulting in an increased energy transfer between the bunches and the plasma. This induces larger pressure gradients within the plasma and leads to more significant perturbations to the density profile which propagate more quickly due to the increased ion thermal velocity. As the plasma density is lowered, not only is the coupling between the plasma and bunches reduced and hence less energy transferred between the two, but the magnitude of the pressure within the plasma is also reduced as $P \propto nT$. The reduction in both the pressure and the gradient induced within the plasma by the leading bunch pair leads to less extreme modifications to the density profile alongside slower evolution of the perturbation.

8.6.3 *Interaction strength vs. interaction density*

The dependency of the ion motion perturbation on the plasma density as measured in the previous section is a convolution of both the coupling efficiency of the leading bunches to the plasma (via the energy transfer) and the magnitude of the pressure gradient induced within the plasma (via the plasma density). Hints to the relative contribution of each effect can be gleaned by comparing the relative perturbations observed when modifying the interaction strength (Section 8.6.1) and those observed when modifying the interaction density (Section 8.6.2). This is because during the variation of the interaction strength, the background

density was kept constant and so the resulting dependence of the perturbation should be a function of *only* the energy transferred to the plasma.

Average spectrometer images of leading drive bunches following their interaction with the plasma, shown in Figure 8.6.2 for the interaction strength and Figure 8.6.8 for the interaction density, are very similar despite the difference in spectrometer imaging setups and indicate a similar energy loss by the drive bunch (and hence energy transferred to the plasma) for the variations considered. This is highlighted by the similarity of the values of the longitudinally-integrated maximal decelerating wakefield amplitude discussed in each section. This suggests the coupling efficiency of the leading bunches to the plasma was similar in the mid (b) and lower (a) perturbation variants in the measurements presented in Sections 8.6.1 and 8.6.2. Note, the discharge timing (2393920) and L1 TOA (227.27 ps) for the largest perturbation in both cases represent the same datasets, i.e. panel (c) in all relevant figures.

Consistently larger magnitude perturbations to the transverse density profile are observed when modifying only the interaction strength and keeping the interaction density constant (Figure 8.6.5). The relative perturbation to the witness bunch was measured to peak at $\Delta\tilde{\mu}_p = 0.283 \pm 0.020\,m^{-1}$ and $\Delta\tilde{\mu}_p = 0.454 \pm 0.031\,m^{-1}$ for the low and mid variants respectively when modifying only the interaction strength compared to $\Delta\tilde{\mu}_p = 0.062 \pm 0.009\,m^{-1}$ and $\Delta\tilde{\mu}_p = 0.192 \pm 0.013\,m^{-1}$ when modifying the interaction density (Figure 8.6.12). For reference, the peak relative perturbation at the highest density and interaction strength was measured to be $\Delta\tilde{\mu}_p = 0.566 \pm 0.022\,m^{-1}$. These measurements indicate that an *additional* damping of the relative perturbative effect by approximately 80% and 45% is observed when modifying the plasma density beyond the initial approximately 50% and 20% reduction observed by purely changing the wakefield interaction strength. This suggests that the dependence of the ion motion perturbation on the plasma density is stronger than that of the wakefield interaction strength, however further measurements must be made before definitive statements regarding its dependencies on experimental variables can be made.

These measurements do however indicate that the lifetime of the perturbation induced by ion motion can be minimised by operating plasma accelerators at higher densities. This is advantageous as higher densities also offer higher accelerating gradients. These measurements also demonstrate that even at these relatively modest accelerating gradients ($> 1\,GVm^{-1}$) and plasma densities ($5 \times 10^{15} - 2 \times 10^{16}\,cm^{-3}$), the ion motion perturbation exists on the timescale of tens of nanoseconds and therefore does not appear to prohibit the operation of plasma accelerators at MHz repetition-rates.

8.6.4 *Ion mass*

The final dependency that was measured was that of the lifetime of the ion motion perturbation on the mass of the ion. Similar to the discussion of the extended evolution of discharge-ionised plasmas of different species in Chapter 6, changes in the mass of the ion are expected to modify the perturbation due to ion motion as the diffusion rate of ions within plasma is expected to scale with the inverse square root of the ion mass, $D_i \propto m_i^{-1/2}$ [see

Eq. (6.3.2)]. The use of lighter ions is expected to lead to two main changes in the perturbative signatures measured within this experiment: (i) the magnitude of the perturbation is increased while, (ii) the lifetime of the perturbation is reduced. The magnitude of the perturbation is expected to increase as ions are able to respond to the wakefield on shorter timescales [see Eq. (8.0.1)]. After the initial wakefield is driven, the plasma electron motion it drives becomes decoherent as phase mixing occurs, leading to a continuous decay in the wakefield amplitude as demonstrated in Chapter 7. Ions within the plasma initially respond to the temporally-averaged transverse wakefield, hence a reduction in the ion mass effectively shortens the period over which the transverse wakefield is averaged. Ions therefore respond to larger amplitude averaged-transverse-wakefields and more extreme perturbations to their density profile are induced. This amplifies the local pressure gradient excited within the plasma which drives further diffusion and dissipation of the density profile, leading to shorter perturbative lifetimes.

Exploratory measurements were performed to investigate this effect. Three different gas species were used; helium, neon, and argon. These measurements were performed in experimental periods prior to those presented in Sections 8.6.1 and 8.6.2 and hence their full scope was not yet realised; for example, during these measurements, the witness bunch was removed and only a drive bunch was used to probe the plasma state, while the minimum temporal resolution of the measurements was not utilised over the entire measurement period. This meant two key ion motion signatures associated with the evolution of the energy spectrum of the witness bunch could not be measured, however the evolution of the perturbation could still be inferred from perturbations to the drive bunch mean energy and shifts in the energy of focal lines within its spectra.

Each gas species has a different discharge profile due to their different breakdown characteristics and hence the discharge parameters varied for different species. While operating with helium, the stable discharge operating point was found at a buffer volume pressure of 50.6 mbar using a discharge voltage of 20 kV. In neon, the buffer volume pressure was reduced to 31.2 mbar using a discharge voltage of 25 kV, and in argon the buffer volume pressure was 20.0 mbar, again using a discharge voltage of 25 kV. During these measurements, the aim was to reach the same interaction density for the leading bunch, irrespective of the exact nature of the discharge. To do this, discharge delay timing scans were performed for each gas species (as in Figure 5.3.3) and the relative energy perturbation of the bunch as a function of the discharge timing was extracted (as in Figure 5.3.13). This was then used to identify the relative discharge delay timing that gave a particular value of the relative energy perturbation and ensured that approximately constant plasma density was used across all subsequent measurements. The evolution of the ion motion perturbation was measured over 21 ns in helium, and 30 ns in both neon and argon, with a minimum temporal resolution of 769 ps.

The resulting energy spectra for both the cold and perturbed plasma schemes are shown in Figures 8.6.13, 8.6.14 and 8.6.15 for helium, neon and argon respectively. Over all three cases, the measurement of the energy spectrum of the probe bunch in the unperturbed plasma

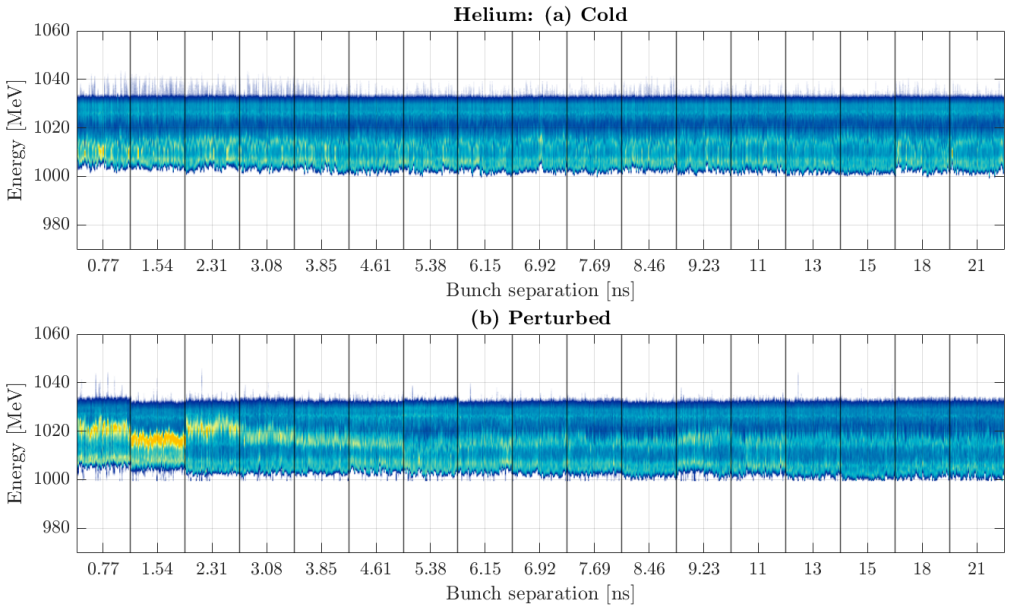

Figure 8.6.13: Waterfall spectrometer images showing the evolution of the probe bunch energy profile as a function of separation with respect to the leading bunch in helium. The post-plasma quadrupoles are set to focus at 1015 MeV, imaging the central energy of the decelerated probe bunch. (a) Probe bunch only. (b) Probe bunch following perturbation from leading bunch.

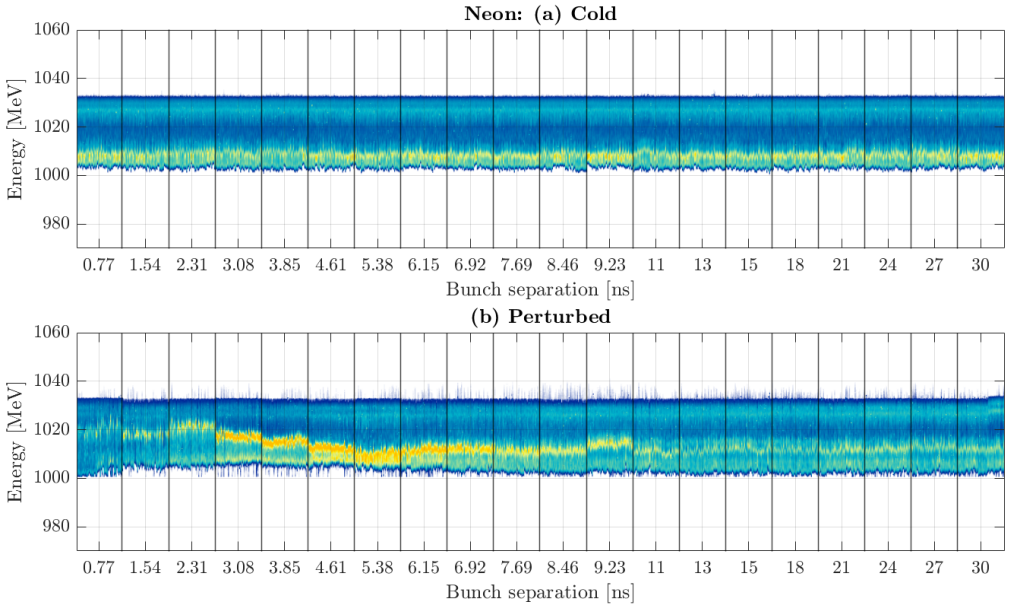

Figure 8.6.14: Waterfall spectrometer images showing the evolution of the probe bunch energy profile as a function of separation with respect to the leading bunch in neon. The post-plasma quadrupoles are set to focus at 1015 MeV, imaging the central energy of the decelerated probe bunch. (a) Probe bunch only. (b) Probe bunch following perturbation from leading bunch.

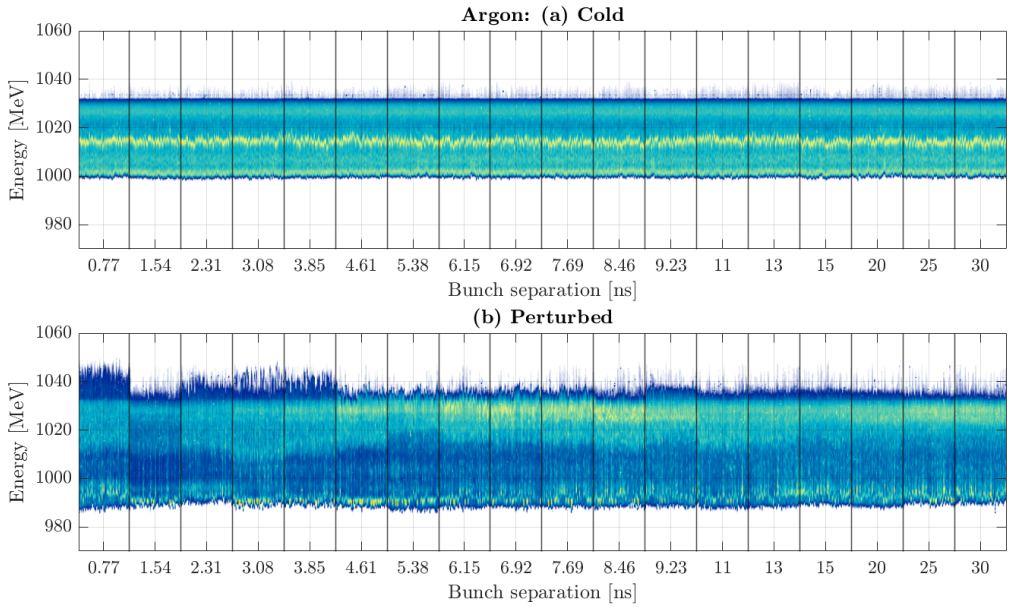

Figure 8.6.15: Waterfall spectrometer images showing the evolution of the probe bunch energy profile as a function of separation with respect to the leading bunch in argon. The post-plasma quadrupoles are set to focus at 1015 MeV, imaging the central energy of the decelerated probe bunch. (a) Probe bunch only. (b) Probe bunch following perturbation from leading bunch.

demonstrates a similar energy loss, indicating that similar interaction plasma densities were found in all three species despite the difference in gas and discharge parameters.

The lifetime of the ion motion perturbation was significantly shortened by reducing the ion mass. In helium, Figure 8.6.13(b), significant perturbations to the drive bunch energy spectra are observed for bunch separations smaller than approximately 4 ns. During this period, the energy loss of the bunch is observed to only increase, indicating an increasing on-axis density experienced by the probe bunch, while the focal line within its energy spectrum tends towards an energy of approximately 1015 MeV as seen in the unperturbed plasma. These observations indicate that the ion motion perturbation is already in the process of dissipating within 769 ps and the temporal resolution is insufficient to measure the development of the peak perturbation as was measured in argon in Sections 8.6.1 and 8.6.2.

In contrast, the measurements of ion motion in neon shown in Figure 8.6.14(b) show both the development and dissipation of the ion motion channel. In this measurement, the channel appears to deepen over approximately the first 5 ns, before dissipating over the following 10 ns such that the probe bunch spectrum returns to those measured in the cold measurement scheme without the presence of the leading bunch. This therefore demonstrates the increase in perturbative lifetime with an increase in ion mass as predicted.

The energy loss of the probe bunch in argon is significantly increased in the perturbed plasma when compared to the unperturbed case for all bunch separations as shown in Figure 8.6.15. This indicates an increase in plasma density and leads to the appearance of an

accelerated tail in the energy spectra as the wakefield cavity is shortened at higher density. The mechanism by which the plasma density increases is discussed in Chapter 9 where more detailed measurements of this effect are performed and hence is not commented on here. The increase in on-axis density masks the perturbation to the drive bunch energy spectrum induced by the presence of an ion channel, however significant changes in the location of increased charge density in the energy spectra over the entire measurement period (30 ns), particularly when compared to measurements of the unperturbed plasma in Figure 8.6.15(a), potentially indicate the continued evolution of an ion channel.

While the measurement conditions were significantly different in Sections 8.6.1 and 8.6.2 and hence direct comparison to the measurements in this section is not possible, those measurements were also performed in argon but in conditions that minimised the re-ionisation effect observed in Figure 8.6.15. Those measurements consistently demonstrated perturbative lifetimes in excess of 50 ns, considerably longer than those in helium and neon demonstrated here, again consistent with an increase in perturbative lifetime with increasing ion mass.

The measurements presented in Figures 8.6.13 and 8.6.14 demonstrate a reduction in the lifetime of the ion motion perturbation induced by wakefield excitation by reducing the ion mass. In helium, the perturbation to the drive bunch was measured to dissipate within approximately 5 ns, while in neon dissipation occured within approximately 15 ns. The lifetime in helium was too short to allow measurement of the peak perturbative effect and hence the expected increase in the magnitude of the ion motion perturbation with reducing ion mass could not be confirmed within the temporal resolution available within the experiment. Given the experimental progress made between the measurements presented in this section and those in the preceding two sections — particularly with regards to the inclusion of witness bunches — the beam-based method appears capable of confirming the expected $m_i^{-1/2}$ scaling of the perturbative lifetime. Future measurements should be performed at lower plasma densities to extend the perturbative lifetime (as demonstrated in Section 8.6.2) to allow measurement of the formation of the ion channel in helium. This would also enable comparison of the amplitude of the peak perturbation and confirm its expected dependence on the ion mass. Measurements in argon should be performed at lower gas pressures to minimise the re-ionisation effect observed in Figure 8.6.15 that masked the ion motion effect. This is discussed further in Chapter 9.

8.7 CONCLUSION

The motion of plasma ions driven by the wakefield excitation process leads to the development of highly non-uniform transverse plasma density profiles that can significantly alter and diminish the wakefield acceleration process. For this reason, ion motion must dissipate and a uniform transverse density profile must be re-established for high-repetition-rate operation of plasma accelerators to be realised. As such, ion motion appears to define the *fundamental* timescale over which identical acceleration can be observed. This chapter has demonstrated the derivation of the effect of parabolic density profiles, produced by an evolving ion channel, on the wakefield process, highlighting the suitability of the beam-based plasma characterisation technique introduced in Chapter 5 to probing the evolution of ion motion.

Experiments were performed at FLASHForward using a drive-probe setup where the separation between the leading bunch pair and a following probe bunch pair was varied from 769 ps to tens of nanoseconds. Four signatures were consistently observed in the reconstructed probe bunch transverse and energy spectra following its interaction with the perturbed plasma. Oscillations in the size of the witness bunch were attributed to the evolving on-axis density profile as ions were driven by the presence of a local pressure gradient induced by the wakefield process. The position of regions of increased intensity in the probe bunch energy spectra were observed to shift with the development of the radial density profile, caused by changes to the betatron oscillation frequency induced by the increasing off-axis ion density and focusing effect. The combination of these two effects significantly modified the shape and structure of the wakefield cavity, leading to significant perturbations to the energy of the probe drive and witness bunches. Identification of these signatures enabled the reconstruction of the evolution of the transverse density profile and three-dimensional PIC simulations using the reconstructed profile demonstrated good agreement with experimental measurements.

The beam-based plasma characterisation method was then applied to measurements of the dependence of the evolution of the plasma state following the wakefield excitation process on three key experimental variables: the wakefield interaction strength of the leading bunch, the background plasma density, and the ion mass. Increasing the wakefield interaction strength increased the relative perturbation experienced by the probe bunch as more energy was transferred to the plasma by the leading bunch. Increasing the plasma density similarly led to an increase in the perturbative effect for the same reason thanks to the increased coupling strength. The lifetime of the perturbation was significantly extended by lowering the plasma density as the magnitude of the induced pressure gradient was reduced. Comparisons between these two effects suggested that the dependence of the ion motion perturbation on the plasma density exceeded that of the wakefield strength. Reductions in the ion mass were shown to reduce the perturbative lifetime as expected, though more detailed measurements should be performed to confirm the expected $m_i^{-1/2}$ dependency.

Future measurements should extend the timescale over which the perturbation is measured to ensure return of the perturbed plasma to the state of the unperturbed plasma in all

cases. Furthermore, the divergence and emittance of the witness bunch should be directly measured experimentally using scans of the transverse bunch size as a function of the object plane of the spectrometer quadrupole imaging system. This will provide a direct comparison of the transverse dynamics with PIC simulations compared to the transverse bunch size measurement which is convolved with the resolution and magnification of the imaging system and assumes that the object plane is consistent across all measurements. Measurement of the emittance in particular could demonstrate its increase due to the non-linear focusing force introduced by the approximately parabolic transverse density profile and confirm theoretical predictions.

The measurements within this section represent the first detailed experimental study of ion motion within beam-driven plasma accelerators and have significant implications for high-repetition-rate operation. The dissipation of the perturbative effect within approximately 100 ns for all measurements presented here indicates that ion motion will not limit operation at MHz repetition-rates. Furthermore, the dependencies measured within Section 8.6 indicate that for collider-relevant parameters, where operating densities and wakefield strengths are increased from those studied here, the perturbative effect to the transverse density profile induced by ion motion will dissipate on even shorter timescales, with sub-nanosecond lifetimes potentially possible when operating in hydrogen, indicating the possibility of GHz repetition-rates if suitable bunches can be provided at such a rate. However, the measurements presented here indicate that ion motion would likely prohibit the use of proton bunch trains within the AWAKE experiment. Such trains are readily available at CERN but have bunch separations of 25 ns. The combination of the low operating plasma density ($n_0 = 7 \times 10^{14}\,\mathrm{cm}^{-3}$) and heavy ion mass ($m_{Rb} = 85.47\,m_p$) necessary for stable self-modulation of the proton bunch would lead to ion motion perturbations over timescales significantly in excess of available bunch separations. It is likely that the perturbation due to ion motion would dissipate within a microsecond for AWAKE parameters and hence MHz repetition-rates may still be possible if suitable proton bunches could be provided at a corresponding rate. However, initial rapid expansion of the laser-ionised column at AWAKE [207], similar to those measured in the evolution of laser-ionised plasma presented in Chapter 6, may prevent re-use of the plasma within its lifetime in this case.

The dissipation of ion motion does not, however, represent the end of the evolution of the plasma state following the wakefield excitation process. The energy imparted into the plasma by the drive bunch, responsible for inducing ion motion, must dissipate. Other mechanisms by which dissipation occurs, possibly over longer timescales than that of ion motion as discussed in Section 1.3.1, must also be considered. Such mechanisms can also lead to the development on non-uniform plasma states that can degrade the acceleration process and prohibit relaxation of the plasma before the following acceleration event. Measurements of an example of just such a mechanism, electron-impact ionisation, are discussed in Chapter 9.

9

MICROSECOND PERTURBED PLASMA EVOLUTION

Chapter 8 demonstrated the evolution of the plasma ion response to the wakefield excitation process and its dissipation within approximately 100 ns for typical FLASHForward experimental parameters. However, as detailed in the introduction to this thesis (Chapter 1), the plasma continues to evolve beyond this point. The energy imparted to the plasma by the drive bunch not only drives ion motion, but also has the potential to lead to re-ionisation of the plasma on extended timescales. Measurements of this effect have previously been reported in Ref. [84] but have not been studied in detail. This chapter first introduces the expected re-ionisation mechanism and then presents measurements of this effect using the beam-based characterisation technique described in Chapter 5. The re-ionisation effect is studied as a function of plasma density and species, enabling a demonstration of its minimisation and indicating the potential for control of the re-ionisation process.

9.1 PERTURBED PLASMA EVOLUTION ON EXTENDED TIMESCALES

It is expected that on extended timescales, the plasma can become re-ionised following wakefield excitation due to the energy transfer from the drive bunch to the plasma. The dominant re-ionisation mechanism is expected to be electron-impact ionisation [84] where high-energy plasma electrons, excited by the wakefield, collide with either neutral recombined atoms or ions within the plasma mixture. The timescale over which this is observed can be approximated by considering the electron-impact ionisation rate coefficient, $\langle \sigma v \rangle$, which has previously been shown to be well-described by the function [82],

$$\langle \sigma v \rangle = A \frac{(1 + P \cdot U^{1/2})}{(X + U)} U^K \exp(-U), \tag{9.1.1}$$

where $U = E_T/T_e$ is the dimensionless relative temperature, E_T is the threshold ionisation energy, T_e is the mean electron temperature and A, K, and X are parameters found from a fit

Ion	E_T [eV]	A [$cm^3 s^{-1}$]	X	K
Ar	15.8	0.599×10^{-7}	0.136	0.26
Ar^+	27.6	0.607×10^{-7}	0.544	0.21
Ar^{2+}	40.9	0.343×10^{-7}	0.834	0.17

Table 5: Fit parameters for the ionisation rates of argon atoms and ions by electron-impact. Values extracted from Ref. [82].

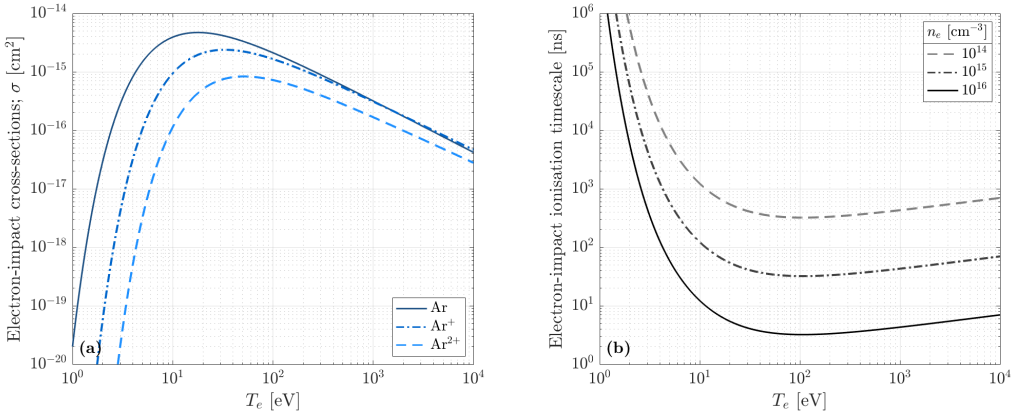

Figure 9.1.1: (a) Dependence of the electron-impact cross-sections on the mean electron temperature (T_e) for the first three ionisation levels of argon. (b) Timescale over which electron-impact ionisation would be expected, calculated via Eq. (9.1.4), for a range of electron densities.

to experimental data [255, 256] and given in Table 5. The parameter P takes a value of either 0 or 1 and is included to improve the fit to the cross-section behaviour for certain ions near the ionisation threshold. For the first three ionisation levels of argon considered here, $P = 1$.

The electron-impact ionisation cross-sections, σ, are highly sensitive to the mean electron temperature and are plotted in Figure 9.1.1(a) for the first three ionisation levels of argon. The rate coefficient is calculated via the overlap integral

$$\langle \sigma v \rangle \sim \int_0^\infty \epsilon \sigma(\epsilon) f(\epsilon) d\epsilon, \tag{9.1.2}$$

where $\epsilon = (3/2) k_B T_e$ is the mean electron energy, $\sigma(\epsilon)$ is the cross-section and $f(\epsilon)$ is the electron energy distribution function. For a highly-ionised plasma in thermodynamic equilibrium, $f(\epsilon)$ is assumed to follow a Maxwell-Boltzmann distribution,

$$f(\epsilon) = 2\sqrt{\frac{\epsilon}{\pi}} \left(\frac{1}{k_B T_e} \right)^{3/2} \exp\left(-\frac{\epsilon}{k_B T_e} \right), \tag{9.1.3}$$

which has a tail at high-energies as demonstrated in Figure 9.1.2. The presence of this high-energy tail leads to a small fraction of electrons with energies above the ionisation threshold, and hence non-zero cross-sections, even at mean electron energies far below the ionisation threshold. As the mean electron temperature increases, the fraction of electrons with energies above the threshold energy significantly increases resulting in a rapid increase in the rate coefficient overlap integral, with the peak of the cross-section occurring at electron temperatures just above the threshold energy. As T_e continues to increase, $f(\epsilon)$ shifts towards higher energies and the cross-section plotted in Figure 9.1.1(a) starts to decrease. The timescale over which the effects from electron-impact ionisation would be expected to be observed can be calculated from the rate coefficient via,

$$\tau \sim \frac{1}{n_e \cdot \langle \sigma v \rangle}. \tag{9.1.4}$$

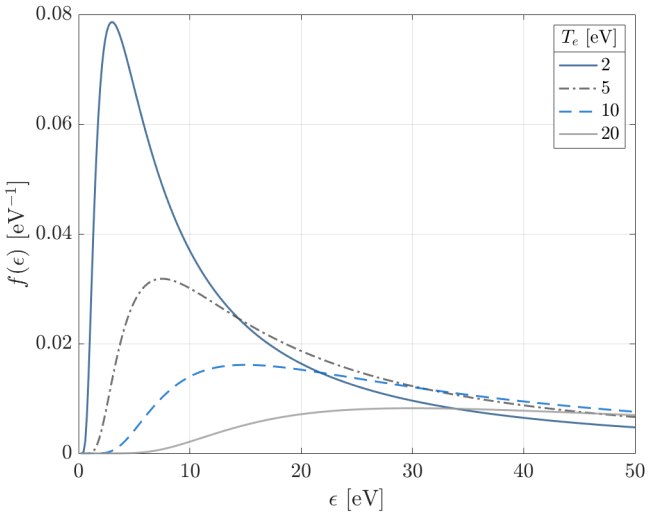

Figure 9.1.2: Example Maxwell-Boltzmann distributions for a range of mean electron temperatures, T_e, calculated via Eq. (9.1.3).

Resulting calculations for argon are shown in Figure 9.1.1(b) for three different electron densities where the total cross-section is the sum of the cross-sections for the first three ionisation levels presented in Figure 9.1.1(a), i.e.

$$\sigma(T_e) = \sigma_{Ar}(T_e) + \sigma_{Ar^+}(T_e) + \sigma_{Ar^{2+}}(T_e). \tag{9.1.5}$$

The electron-impact ionisation cross-sections presented in Figure 9.1.1(a) indicate that the dominant contribution to the growth of the plasma density will likely be collisions between electrons and neutral argon atoms, with a growing contribution from second-level ionisation of argon ions at high electron energies. Plasma electrons near the axis that are excited by the wakefield are expected to gain large energies (> keV), however these quickly move from the vicinity of the wakefield due to their large thermal velocity and undergo many collisions with lower-energy plasma electrons and ions, reducing their energy. This process acts to heat the bulk of the plasma over a timescale given by the inverse of the electron-ion collision frequency [128, 196]:

$$\nu_{ei} = \frac{4}{3}\sqrt{\frac{2\pi}{m_e}}\frac{e^2 n_e \ln\lambda_{ei}}{(4\pi\varepsilon_0)^2(k_b T_e)^{3/2}}, \tag{9.1.6}$$

where $\ln\lambda_{ei}$ is the electron-ion Coulomb logarithm given by

$$\ln\lambda_{ei} = \ln\left[\frac{3}{2\sqrt{2\pi}}\frac{(4\pi\varepsilon_0)^{3/2}(k_b T_e)^{3/2}}{e^3 n_e^{1/2}}\right], \tag{9.1.7}$$

and is typically of order 10. Example timescales for a range of plasma densities are shown in Figure 9.1.3 assuming the ion temperature is initially sub-eV. This figure indicates that high-energy plasma electrons ($T_e \sim$ keV) excited by the wakefield would be expected to cool to eV-scales (where their impact-ionisation cross-section is highest) on timescales of the order

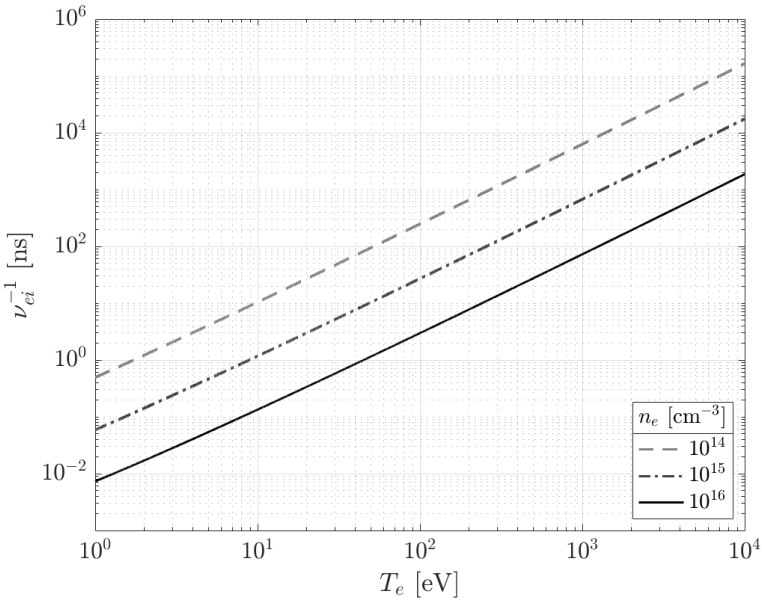

Figure 9.1.3: Electron-ion collision timescale, ν_{ei}^{-1}, assuming cold ions ($T_i < 1$ eV) for a range of plasma densities, calculated according to Eq. (9.1.6).

of tens-to-hundreds of nanoseconds for typical FLASHForward operating plasma densities. Given that re-ionisation of the plasma begins to be observed on timescales shorter than this, as will be shown in Section 9.2, this suggests that it is not the highest-energy electrons that contribute to the onset of impact-ionisation but rather bulk plasma electrons that are heated to $\mathcal{O}(10)$ eV scales via collisions with wakefield-excited electrons and ions.

For typical FLASHForward operating densities, $n_e \sim 10^{15} - 10^{16}\,\mathrm{cm}^{-3}$, it would be expected that re-ionisation of the plasma via electron-impact ionisation would be observed on a timescale of order few-to-tens of nanoseconds following the excitation of the wakefield. Initial measurements of this effect were performed using the beam-based plasma characterisation technique and are presented in the following sections.

9.2 PLASMA DENSITY DEPENDENCE

Beam-based measurements of the re-ionisation of the plasma following wakefield excitation were performed in conjunction with the measurements of the extended evolution of unperturbed plasma discussed in Chapter 6. For these measurements, the "perturbed" measurement scheme introduced in Chapter 8 was used where a leading bunch drives a wakefield within the plasma and is then followed by a second bunch at variable delay that probes the evolution of the plasma. In all cases presented here, both the leading and probe bunch comprised of a drive bunch only and no witness bunches were present. The spectrometer signal subtraction technique outlined in Section 8.4.1 was used in the analysis of spectrometer images as the bunch separations were sufficiently short that the scintillating signals of the bunches overlapped. In all datasets, the "cold" measurement scheme, where the leading

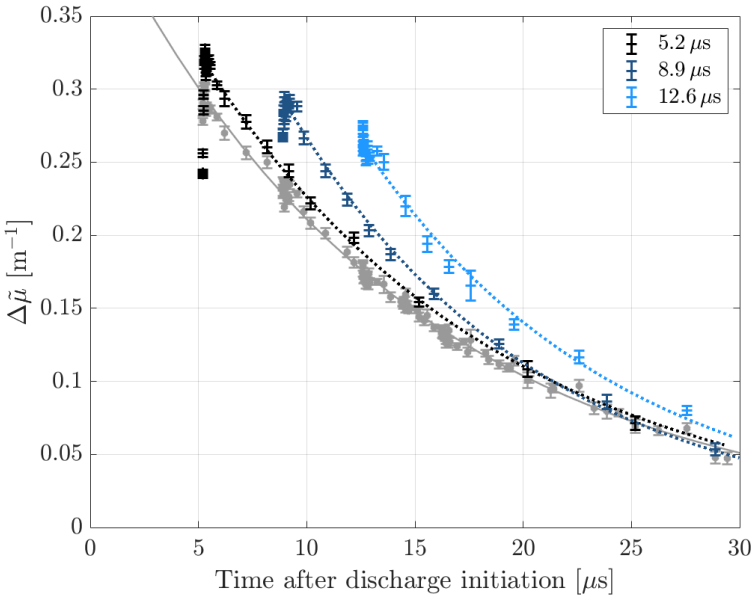

Figure 9.2.1: Beam-based measurement of the re-ionisation of plasma following wakefield excitation by a leading bunch. Three different background plasma densities are measured by changing the timing of the arrival of the leading bunch with respect to the initiation of the discharge; 5.2 µs (black), 8.9 µs (dark blue) and 12.6 µs (light blue). The decay of the unperturbed plasma is plotted in grey for reference. Exponentials are fitted to all datasets and plotted.

bunch is removed, is compared to the "perturbed" measurement scheme to determine the effect of the wakefield on the evolution of the plasma.

Figure 9.1.1(b) indicates that the onset of electron-impact ionisation should show a clear dependency on the density of high-energy electrons (and hence plasma density), with ionisation occurring sooner at higher densities. This was explored experimentally by varying the arrival time of the leading bunch with respect to the initiation of the discharge, as an increased delay allows the plasma more time to recombine resulting in a lower interaction density. First measurements were performed in discharge-ionised argon at a buffer volume pressure of 40 mbar, similar to results discussed in Section 6.2 although measured using a different probe bunch. Three different relative arrival times were used; 5.2 µs, 8.9 µs and 12.6 µs, with relative timing between the leading and probe bunches, $\Delta\tau$, varied between 9.23 ns and 15 µs with a minimum temporal resolution of 9.23 ns.

The evolution of the fractional energy perturbation per unit length, $\Delta\tilde{\mu}$, for the "cold" (grey) and "perturbed" measurements schemes are plotted in Figure 9.2.1, and repeated in Figure 9.2.2 with a logarithmic relative timing axis to highlight the early timescale evolution. The increased value of $\Delta\tilde{\mu}$ in all three "perturbed" cases when compared to the "cold" scheme indicate an increase in plasma density following excitation of the wakefield.

At the highest interaction density, Figure 9.2.2(a), at the shortest timescales ($\Delta\tau \lesssim 40$ ns), $\Delta\tilde{\mu}$ is initially lower in the "perturbed" case than the "cold" case. This is due to ion motion reducing the on-axis density experienced by the probe bunch as discussed in detail within

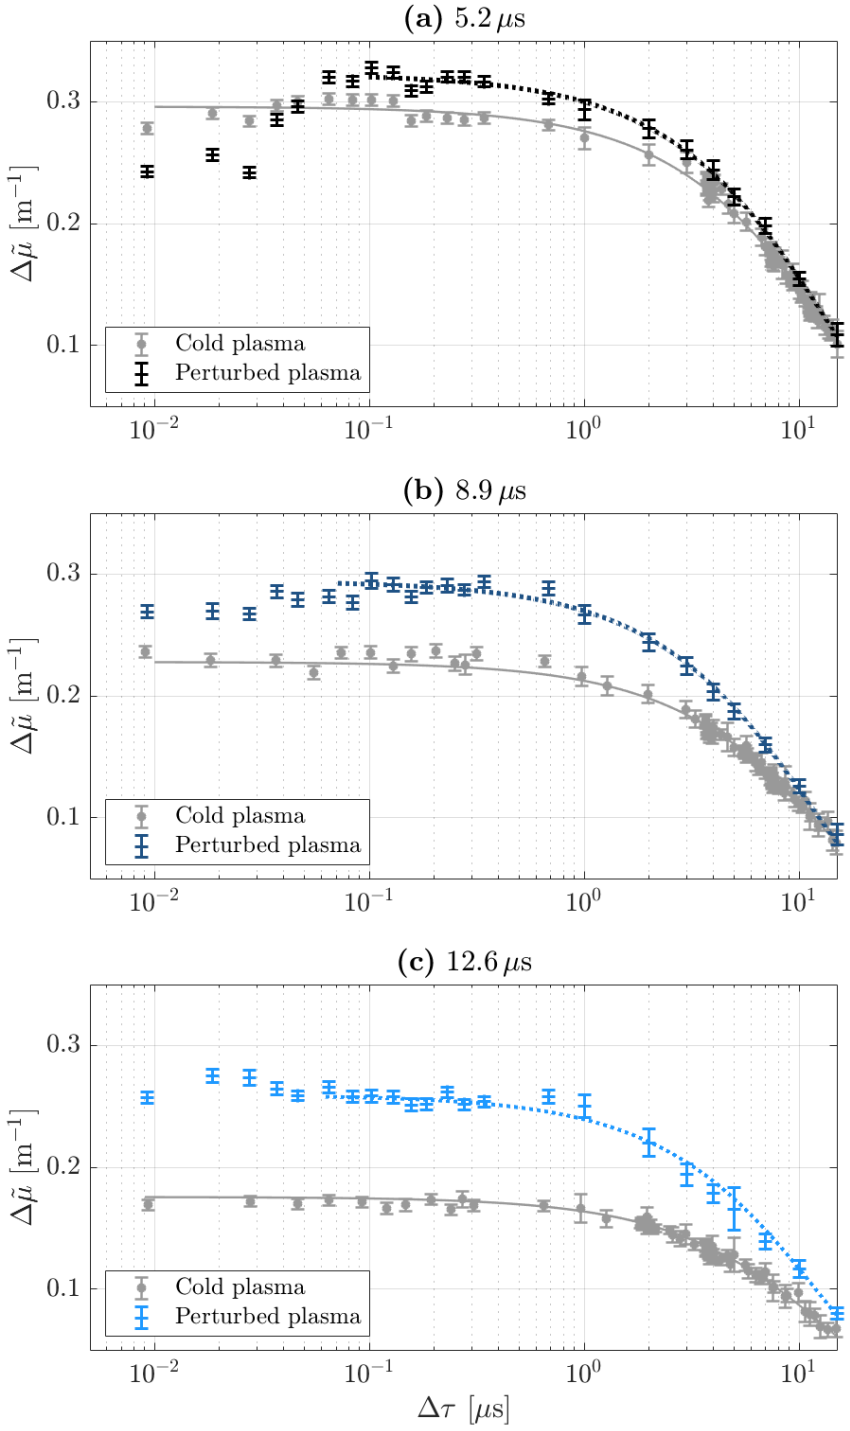

Figure 9.2.2: Datasets from Figure 9.2.1 re-plotted separately and using a logarithmic temporal axis to show short timescale detail. Leading bunch relative arrival time of (a) 5.2 μs, (b) 8.9 μs and (c) 12.6 μs. The decay of the unperturbed plasma is plotted in grey for reference. Exponentials are fitted to all datasets and plotted. $\Delta\tau$ represents the relative timing between the leading and probe electron bunches.

Measurement	A	B $[\mu s^{-1}]$	χ^2_ν
"Cold"	0.427 ± 0.003	$(7.04 \pm 0.07) \times 10^{-2}$	1.09
(a)	0.471 ± 0.007	$(7.31 \pm 0.22) \times 10^{-2}$	1.10
(b)	0.639 ± 0.016	$(8.70 \pm 0.24) \times 10^{-2}$	0.92
(c)	0.709 ± 0.028	$(7.98 \pm 0.29) \times 10^{-2}$	1.36

Table 6: Results of fits of the measurements of $\Delta\tilde{\mu}$ presented in Figure 9.2.2 to the exponential function $f(t) = A\exp(-Bt)$.

Chapter 8. As the ion motion perturbation dissipates, the plasma density experienced by the probe bunch increases, indicated by the increase in $\Delta\tilde{\mu}$, demonstrating the apparent onset of re-ionisation of the plasma. The timescale over which this is observed, approximately 50 ns, is consistent with both the dissipation of ion motion measured in Chapter 8 and the expected electron-impact ionisation timescale presented in Figure 9.1.1(b) for an electron density of order 10^{15} cm^{-3}. The increased plasma density, relative to the "cold" measurement scheme, is maintained over the remainder of the measurement period and exhibits a similar exponential relation between $\Delta\tilde{\mu}$ and $\Delta\tau$ as the unperturbed plasma. The measurements of the perturbed plasma $\Delta\tilde{\mu}$ from $\Delta\tau = 100$ ns -15 µs were fitted to an exponential function of form $f(t) = A\exp(-Bt)$, with each data point weighted by the inverse of its relative error, and compared to the same fit to the unperturbed plasma. The resulting fit parameter values are listed in Table 6 with fits represented by the lines in Figures 9.2.1 and 9.2.2 with dotted lines showing the fit to the perturbed measurement and the solid line to the unperturbed plasma.

At the highest density the exponential decay rate, given by B in Table 6, is consistent between the two measurements, indicating a similar extended timescale decay of both the "cold" and "perturbed" plasma and suggesting that the effect of the wakefield would not significantly extend the lifetime of the plasma in this case.

As the interaction plasma density is reduced, the onset of re-ionisation of the plasma appears earlier, so much so that within the first measurement at a relative bunch separation of 9.23 ns the plasma is already significantly re-ionised in both cases in Figure 9.2.2(b) and (c). As discussed in Chapter 8, it would be expected that the ion motion perturbation would exist for a longer time period at lower densities but have a lower perturbative effect. In the measurements at lower densities presented here, the re-ionisation effect appears significantly larger than the perturbative ion motion effect as no effects that have previously been attributed to ion motion are observed. The apparent demonstration of a shorter re-ionisation onset at lower plasma densities is in direct contradiction with calculations in Figure 9.1.1(b) which suggested that the electron-impact re-ionisation timescale increases at lower plasma densities as there are fewer high-energy electrons that can undergo such collisions. This contradictory result therefore suggests that the onset of re-ionisation is not restricted by the population density of high-energy electrons but rather the density of ions and neutral

atoms within the vicinity of the axis of propagation of the probe bunch that can undergo electron-impact ionisation. Two mechanisms would be expected to significantly alter the density distribution of ions and neutral atoms within the capillary: (i) ion motion and (ii) the background plasma (and neutral) density itself.

At higher plasma densities where, in Section 8.6.2, ion motion perturbations were shown to be more significant, the density of ions that remain within the vicinity of the propagation axis of the probe bunch on nanosecond timescales after the wakefield has been driven is much reduced. Neutrals will not respond to the average transverse wakefield that initially drives the motion of plasma ions as they have no charge, but they do respond to the pressure gradient within the plasma that is induced by the motion of the ions and electrons. Therefore, neutrals will also be transversely perturbed by the wakefield excitation. It is therefore likely that any re-ionisation that occurs on few-nanosecond timescales takes place away from the propagation axis of the probe bunch, and hence only affects the measurements of the on-axis plasma density once the ion motion has dissipated. At lower plasma densities where the on-axis density perturbation due to ion motion is less extreme, there still exists a high density of ions and neutrals near the propagation axis of the bunches that can be ionised via collisions with plasma electrons of sufficient energy.

Furthermore, the measurements presented here are performed in discharge-ionised plasma which continuously recombines and evolves following the culmination of the discharge as discussed with regards to the measurements presented in Chapter 6. During, and immediately after, the discharge a significant fraction of the plasma is expelled out of the ends of the capillary due to the thermal pressure gradient [195]. Following this expulsion, the remaining plasma recombines and the density of neutral argon atoms within the capillary grows. Within these measurements, as the delay between the discharge initiation and the arrival of the leading bunch is increased, the plasma density decreases via recombination and the neutral density correspondingly increases. Therefore, when the wakefield is driven in a lower density plasma by the leading bunch, there is a larger density of neutral argon atoms distributed within the capillary that can undergo electron-impact ionisation collisions and be ionised, leading to an earlier apparent onset of re-ionisation.

It is therefore likely that the apparent decrease in the timescale of the onset of electron-impact ionisation as a function of the interaction plasma density shown in Figure 9.2.2 is as a result of using a beam-based plasma characterisation technique where only the plasma density local to the propagation axis of the beam is probed. Measurements of regions away from the leading and probe bunch axis would likely indicate the ionisation of neutrals and ions on few-nanosecond timescales at the highest plasma density measured here, similar to results found in Ref. [84].

Exponential fits to the extended timescale evolution of the perturbed plasma at these lower densities indicate a faster decay when compared to the unperturbed plasma, as can be seen from the values of the fit parameters in Table 6. This is likely due to the increase in the bulk plasma temperature driven by the energy transferred by the leading bunch to the plasma via the wakefield. A local increase in temperature in the region near the propagation axis drives

a pressure gradient, increasing the diffusion rate of electron and ions within the plasma. Their redistribution towards the cooler capillary walls is accelerated, resulting in an increased cooling and recombination effect relative to the unperturbed plasma. The measurements presented here only demonstrate the perturbative effect on timescales up to 15 μs and hence it is not known whether this perturbation continues to drive redistribution and recombination of the plasma at a higher rate and would hence reduce the overall plasma lifetime. Similar measurements should be repeated until the resolution limit of the beam-based technique is reached, $\Delta\tilde{\mu} \sim 10^{-3}\,\mathrm{m}^{-1}$ for this setup, to determine whether this is in fact the case or whether the extended timescale evolution mirrors that of the unperturbed plasma.

9.3 PLASMA SPECIES DEPENDENCE

The plasma re-ionisation effect demonstrated in the previous section should be dependent on the re-ionisation mechanism cross-section. For example, in the case of electron-impact re-ionisation, the cross-section can be significantly reduced by operating in a different plasma species with a higher ionisation threshold. Measurements to investigate this effect were performed in discharge-ionised argon and helium plasmas at buffer volume pressures of 8.6 mbar and 25.3 mbar respectively, as in Section 6.3.2. The capillary had length 50 mm and a discharge voltage of 25 kV was used. The operation of the discharge at such low pressures resulted in a less stable discharge, with large amplitude current reflections from impedance mismatching within the discharge electronics observed in the measured current traces. This leads to slightly larger variability in the energy spectra of the bunches following plasma interaction, however the evolution of the plasma can still be clearly observed. The probe electron beam had charge 455 ± 3 pC with energy distribution $\mu_0 = 1115.5 \pm 1.2$ MeV and $\sigma_0 = 1.7 \pm 0.1$ MeV, and was compressed to an RMS bunch length of $\sigma_\tau = 200 \pm 2$ fs, giving a peak current of approximately 1 kA. The measurements of the extended evolution of perturbed argon and helium plasmas were each performed at two different interaction densities using the same drive and probe bunches such that direct comparison between the datasets via measurements of $\Delta\tilde{\mu}$ is possible. The higher and lower interaction density working points are characterised by similar values of $\Delta\tilde{\mu}$ for the unperturbed plasma in both species, indicating similar (but not identical) interaction densities in both cases. The relative arrival time of the leading and probe bunches, $\Delta\tau$, was varied between 769 ps -5 μs with a minimum temporal resolution of 769 ps.

The electron-impact ionisation cross-sections of the first and second ionisation levels of helium are shown in Figure 9.3.1 alongside those for argon that were previously shown in Figure 9.1.1(a). The corresponding fit parameters to Eq. (9.1.1) for helium are given in Table 7 for reference. The electron-impact ionisation cross-section for neutral helium is approximately an order of magnitude lower than for argon over the calculated range of electron temperatures, indicating a significantly lower probability of electron-impact ionisation occurring following the excitation of the wakefield.

Ion	E_T [eV]	A [$cm^3 s^{-1}$]	X	K
He	24.6	0.175×10^{-7}	0.180	0.35
He^+	54.4	0.205×10^{-7}	0.265	0.25

Table 7: Fit parameters for the ionisation rates of helium atoms and ions by electron-impact. Values extracted from Ref. [82].

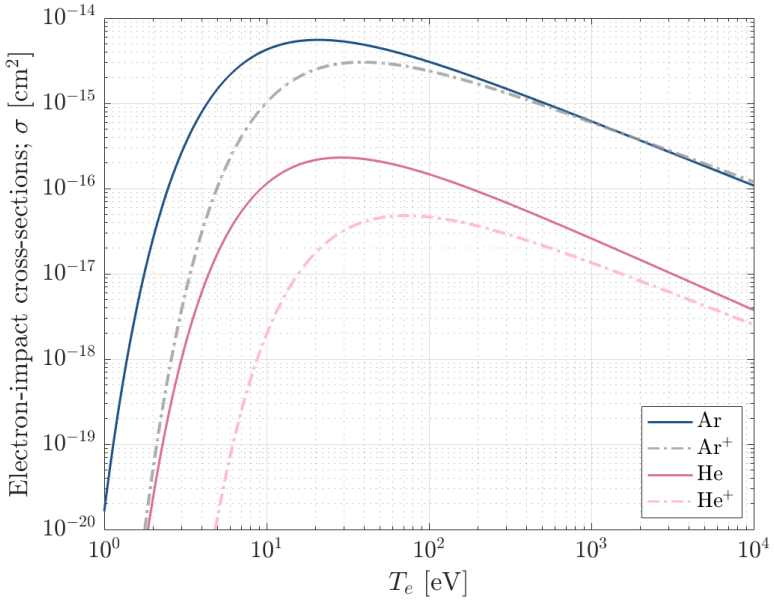

Figure 9.3.1: Dependence of the electron-impact cross-sections on the mean electron temperature (T_e) for the first two ionisation levels of argon and helium.

The measured fractional energy perturbation per unit length to the probe bunch, $\Delta\tilde{\mu}$, for the four datasets are shown in Figure 9.3.2. In all cases, the measurement of the unperturbed plasma is shown in grey for reference. For all datasets, the effects of ion motion are observed at the shortest timescales. Within argon, ion motion signatures are typically observed over the first approximately 25 ns in these datasets, whereas in helium they appear to dissipate within the first few nanoseconds.

At the higher interaction density in argon, minimal re-ionisation of the plasma is observed over the entire measurement period. This contrasts to the measurements in the previous section where significant re-ionisation was observed at all measured plasma densities. This is likely due to the difference in buffer volume pressure between the two measurements, with a pressure of 8.1 mbar used here compared to 40 mbar in Section 9.2. The reduction in buffer volume pressure correspondingly reduces the gas pressure within the capillary. Under identical discharge parameters, it would be expected that the reduction in gas pressure would result in higher temperatures following culmination of the discharge and hence an increased expulsion of plasma and gas out of the ends of the capillary. The density of neutral argon

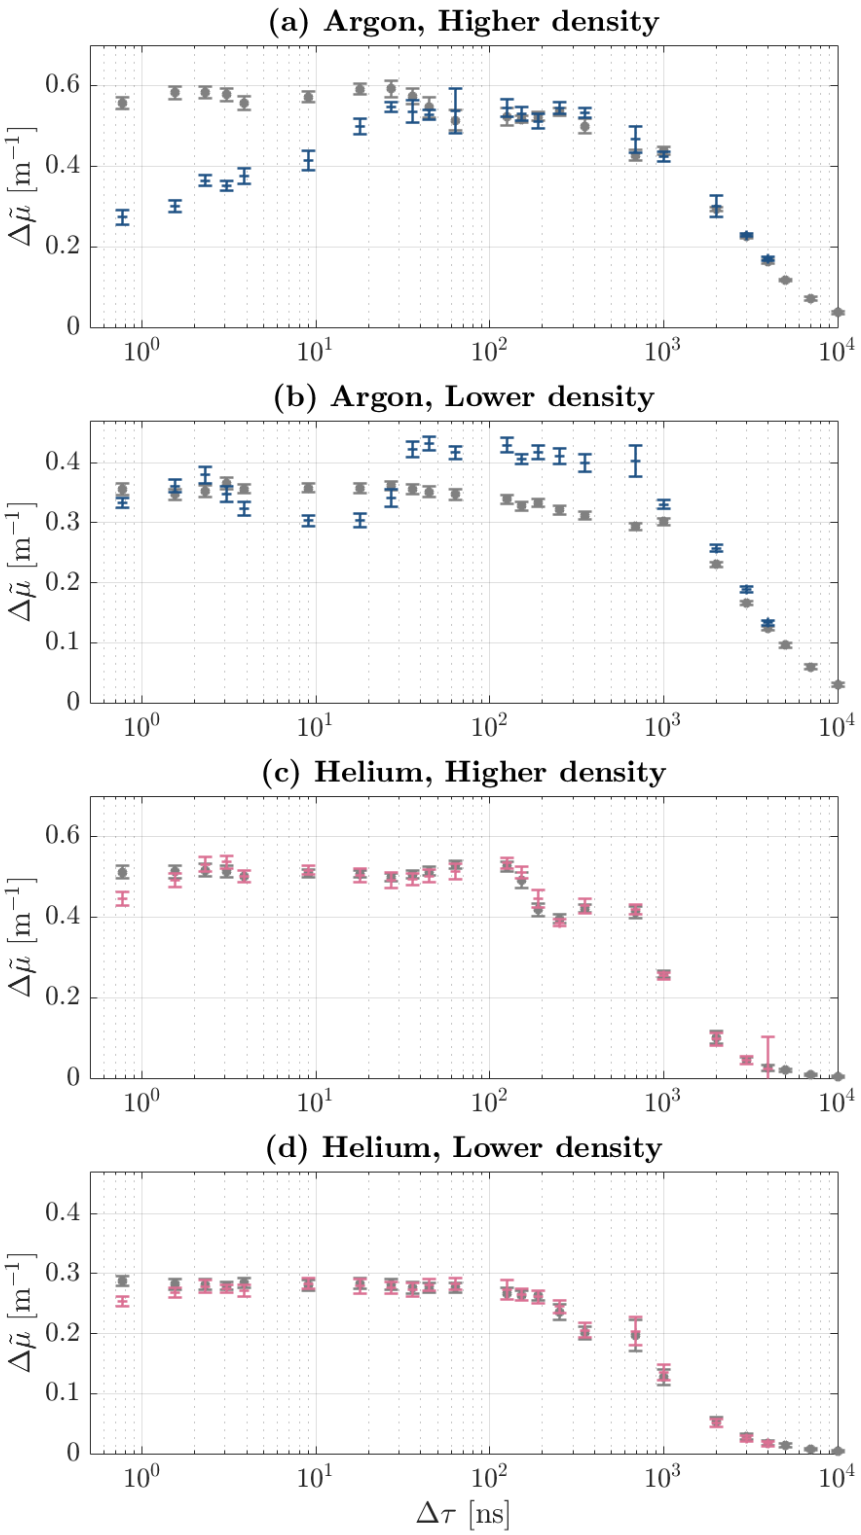

Figure 9.3.2: Comparison between the re-ionisation measured after wakefield excitation in differing plasma species. Panels (a) and (b) represent higher and lower plasma density working points in argon respectively. Panels (c) and (d) represent higher and lower working points in helium respectively. The unperturbed plasma measurements are plotted in grey for all datasets for reference.

species within the capillary at the arrival of the leading bunch would therefore be reduced when compared to the 40 mbar buffer volume pressure, and hence the re-ionisation effect is reduced. It can still be observed however by increasing the delay between the initiation of the discharge and the arrival of the leading bunch, hence reducing the interaction plasma density but increasing the neutral density. This is the case in Figure 9.3.2(b), where significant re-ionisation of the argon plasma is observed following wakefield excitation. The gas pressure dependence of the re-ionisation mechanism was utilised in the measurements of ion motion presented in Chapter 8 where a buffer volume pressure of 10.6 mbar was selected to minimise the re-ionisation effect of the argon plasma and enable the study of only the ion motion effect.

These measurements were repeated at similar interaction densities in helium and are presented in Figure 9.3.2(c) and (d). At the higher interaction density, similarly to argon, ion motion effects resulting in a decrease in the measured $\Delta\tilde{\mu}$ are observed on the shortest timescales. However, following their dissipation the measurement of the evolution of the perturbed plasma agrees with the unperturbed plasma over the entire measurement range. Unlike the measurements in argon, increasing the delay of the arrival of the leading bunch to lower the interaction density, such that a greater number of recombined helium atoms would be present within the capillary, did not demonstrate significant re-ionisation of the plasma within the resolution afforded by the beam-based plasma characterisation technique. This is in-line with the order of magnitude reduction in the electron-impact ionisation cross-section for helium compared to argon presented in Figure 9.3.1. On extended timescales, the decay of the perturbed helium plasma closely matched that of the cold plasma indicating no significant difference in the lifetime of the plasma would be expected. Rapid changes in the measured $\Delta\tilde{\mu}$, for instance at $\Delta\tau \sim 350$ ns in Figure 9.3.2(c), are due to reflections of the discharge current within the plasma caused by an impedance mismatch. This short secondary current pulse heats and re-ionises the plasma, resulting in an increased $\Delta\tilde{\mu}$ before the plasma decay is continued. These are present in these measurements due to the reduced gas pressure within the capillary, but are consistent across both the "cold" and "perturbed" measurement schemes indicating the repeatability of the current reflection and highlighting the sensitivity of the beam-based method to changes in the plasma density.

9.4 CONCLUSION

The plasma wakefield process generates highly-energetic plasma electrons that undergo inter- and intra-species collisions with other constituents of the plasma, heating the plasma bulk. It is expected that collisions between plasma electrons and neutral atomic species within the capillary result in their ionisation and an increase in the plasma density. Initial exploratory measurements of the extended evolution of plasma following the wakefield excitation process have demonstrated such re-ionisation of the plasma over timescales that are consistent with those expected for electron-impact ionisation for the plasma densities considered.

It is expected from theory [82] that the timescale over which electron-impact ionisation is observed should decrease significantly at increased plasma densities. Beam-based measure-

ments of the re-ionisation of a discharge-ionised argon plasma over a range of plasma densities indicated the opposite relation, with the effects of plasma re-ionisation being observed on shorter timescales at lower densities. It is believed this was due to the interplay between ion motion and re-ionisation effects, with stronger ion motion perturbations observed at higher plasma densities reducing the local population of ions and neutral atoms in the vicinity of the wakefield driven by the probe bunch, masking the re-ionisation effect. The decay rate of the perturbed plasma was measured to increase slightly with decreasing plasma density, but further measurements over a longer timescale should be performed to determine whether this continues over its entire lifetime.

The plasma species was varied experimentally to explore the effect of reducing the electron-impact ionisation cross-section. Comparisons between discharge-ionised argon and helium plasmas demonstrated significant re-ionisation in the case of argon while no significant re-ionisation effect was observed in helium, even at reduced plasma densities where the largest effects are measured in argon. This was consistent with theoretical calculations of the electron-impact ionisation cross-section of helium which indicated an order of magnitude decrease in the probability of electron-impact ionisation occurring compared to argon. Reducing the gas pressure within the capillary via a reduction of the buffer volume pressure was shown to further minimise the re-ionisation effect within the plasma thanks to a reduction in the density of neutral gas atoms present within the capillary.

The measurements presented within this chapter demonstrate not only the re-ionisation of plasma following the wakefield process, but also its tunability. For example, it can be minimised by either reducing the gas pressure within the capillary or using a gas species with a lower electron-impact ionisation cross-section. Conversely, the re-ionisation effect can be increased by increasing the density of neutral atoms within the plasma source, either by increasing the gas pressure or allowing more recombination to occur before driving a wakefield within the plasma. This suggests that via careful tuning of the plasma source and density, the re-ionisation effect can be controlled and potentially even utilised to enable high-repetition-rate PWFA operation in the "burst" mode often used at state-of-the-art conventional accelerator facilities. Measurements of this type could be additionally combined with temporally-resolved spectroscopic measurements of the light emitted by the plasma following the wakefield excitation process. Such measurements could allow the determination of the evolution of the plasma electron temperature [257] which, combined with the electron-impact ionisation cross-sections, would permit modelling of the re-ionisation process and allow the study of its dependence on additional experimental parameters. This could be used to optimise experimental parameters (e.g. the plasma species and density) to enable sufficient re-ionisation to occur between the arrival of consecutive bunches and hence maintain the operating plasma density over extended periods of time such that long, $\mathcal{O}(100\,\mu s)$, MHz-repetition-rate bunch trains could be used.

Part V

CONCLUSIONS

See the individual chapters of this thesis for specific conclusions and outlooks with regards to each set of measurements. Here, the wider implications of the measurements presented within this thesis are discussed, particularly with regards to potential operating modes of future plasma-based accelerators.

IMPLICATIONS FOR HIGH-REPETITION-RATE PWFA OPERATION

The measurements presented within this thesis seem to indicate three different plasma wakefield accelerator high-repetition-rate operation modes may be possible in the future. These are represented in Figure C1.

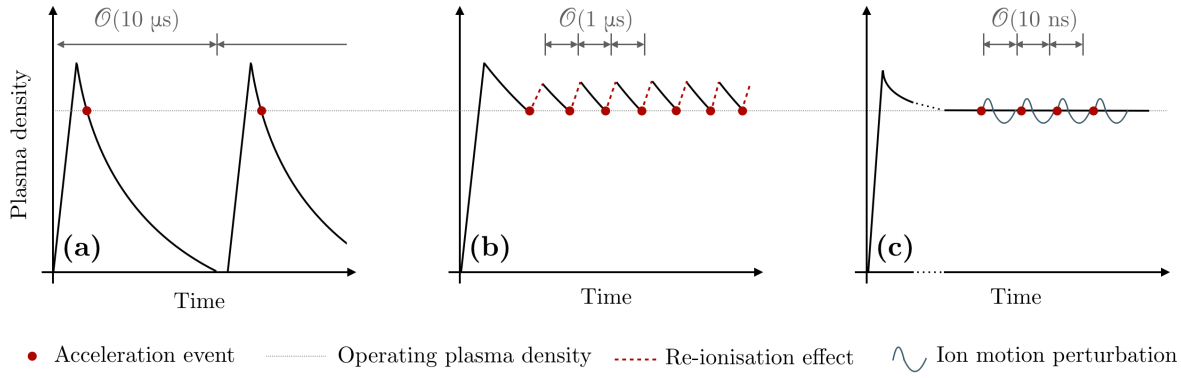

Figure C1: Potential operating schemes for future plasma-based accelerators. (a) A fresh plasma is generated for every acceleration event. (b) The re-ionisation effect of the wakefield is used to maintain the plasma density. (c) Bunch trains are used, with ion motion dissipating between the arrival of the following bunch. Note, different temporal scales are represented in each panel.

(a) A PWFA-based accelerator could be operated under the principle that the plasma is freshly generated for every acceleration event. In this case, the measurements in Chapters 6, 8 and 9 indicate that helium provides a suitable accelerating medium. The lifetime of discharge-ionised helium with a peak density in excess of $10^{16}\,\mathrm{cm}^{-3}$ has been shown to be of order 10 µs, its low mass means ion motion effects dissipate within a few nanoseconds and its low electron-impact ionisation cross-section means re-ionisation of the plasma is avoided on extended timescales. These results indicate that a helium-based plasma wakefield accelerator, where a fresh plasma is generated for each event, could be operated at repetition-rates approaching 100 kHz, of similar order to state-of-the-art conventional accelerator facilities. However, continuous operation at 100 kHz is not necessarily compatible with current facilities that could provide electron bunches as they typically operate in a short-pulse [$\mathcal{O}$(100 µs)] "burst" mode. This scheme is therefore possibly more suited to laser-driven facilities provided that suitably powerful, high-repetition-rate laser systems can be developed.

(b) In contrast, argon appears to be a suitable medium for a "burst" operation mode where the re-ionisation of the plasma induced by the preceding wakefield event is used to increase the plasma density between events such that the same plasma density is reached when a subsequent bunch arrives within $\mathcal{O}(1\,\mu s)$. Measurements in Chapter 9 indicated that the level of re-ionisation induced in the plasma by the wakefield process could be controlled via changes in the density of neutral atoms within the capillary, either by adjusting the gas pressure or the time over which the plasma recombines before the interaction occurs. This therefore suggests that suitable operating conditions could be found where short MHz-frequency bunch trains can be used thanks to their re-ionisation of the plasma between consecutive acceleration events. Small discharge currents (in comparison to the initial discharge) could additionally be applied to further boost the plasma density if necessary. Hence, plasma wakefield accelerators could be operated at MHz-frequencies via use of a "burst" operation mode, akin to the operation modes of state-of-the-art super-conducting RF accelerators. As such, high-gradient plasma wakefield sections could be considered as standalone extensions to existing facilities, acting as a "plasma booster" stage to significantly increase their energy with minimal extension to the footprint of the accelerator.

(c) The effects of ion motion within the plasma lead to the generation of highly non-uniform transverse density profiles and therefore seem to define the *fundamental* limit at which repeatable acceleration events could be observed. Measurements of this effect in argon within Chapter 8 indicated a perturbation lifetime of order 70 ns and demonstrated a reduction in this lifetime similar to the expected scaling of $m_i^{1/2}$. This therefore implies that for hydrogen plasma, a perturbative lifetime significantly below 10 ns could potentially be achieved. It is therefore possible that through careful choice of operating conditions (high density, low ion mass), plasma-based accelerators could be compatible with ultra-high frequency bunch trains similar to those produced at CERN where the bunch separation is typically 25 ns. They would need to be operated in a "burst" mode similar to the MHz-frequencies discussed previously, but theoretically it appears frequencies approaching GHz repetition-rates may be possible.

A word of caution

While the measurements presented in this thesis represent a significant first step towards identifying and mitigating physical mechanisms that could limit high-repetition-rate operation of plasma-based accelerators, the three potential operation modes outlined previously have a number of associated caveats.

Principle among the "burst" operation modes [(b) and (c)] is that all measurements presented here were performed with, at most, two bunches (or bunch pairs) in the plasma accelerator and hence the cumulative effect of multiple acceleration processes on a plasma could not be studied. It is improbable that a third acceleration event would necessarily give

an identical signal to the second, let alone the n^{th} event, without further optimisation of the acceleration process.

However, the FLASHFORWARD experiment is in an ideal position to investigate exactly this effect thanks to the flexibility provided by the FLASH linac. The use of multiple photocathode lasers permits the generation of a stable bunch train of length up to 800 µs at $\mathcal{O}$(MHz) repetition-rates with a first photocathode laser, and generation of a further similar probe bunch with a second photocathode laser that can be placed anywhere within that bunch train with a temporal resolution of 769 ps. This enables the beam-based perturbation method developed within the X-3 experiment to be used to study the perturbation to the plasma after the first, second, fifth, hundredth... bunch in the main train and demonstrate not only consistent acceleration signatures but also consistent plasma relaxation (or otherwise).

Other than the lack of compatibility with the operating modes of current accelerator facilities, the principle caveat to the first operation mode discussed above that relies on the fresh-generation of plasma for each event is the refill rate of the capillary. It has been shown that a significant fraction of the gas within the capillary is expelled during the discharge. This must either be refilled at a rate matching the operating rate or significant redesign of the capillary must be made to minimise this expulsion.

Furthermore, the measurements presented in this thesis are specific to the relaxation rate of the plasma that acts as the acceleration medium. As outlined in Chapter 4, a number of technical challenges remain before high-repetition-rate operation can be demonstrated such as cooling of the capillary and ionisation of the plasma at an appropriate rate. Such studies will be performed within the context of the X-3 experiment over the coming years as it aims to achieve its ultimate goal of the demonstration of an $\mathcal{O}$(10 kW) average-power plasma stage.

BIBLIOGRAPHY

[1] J. J. Thomson. "XL. Cathode Rays." *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science* 44.269 (1897), pp. 293–316. DOI: 10.1080/14786449708621070.

[2] E. Rutherford. "LXXIX. The scattering of α and β particles by matter and the structure of the atom." *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science* 21.125 (1911), pp. 669–688. DOI: 10.1080/14786440508637080.

[3] E. Rutherford. "Address of the President, Sir Ernest Rutherford, O. M., at the Anniversary Meeting, November 30, 1927." *Proceedings of the Royal Society of London* 117.777 (1928), pp. 300–316. DOI: 10.1098/rspa.1928.0001.

[4] J. Chadwick. "The existence of a neutron." *Proceedings of the Royal Society of London* 136.830 (1932), pp. 692–708. DOI: 10.1098/rspa.1932.0112.

[5] C. D. Anderson. "The Positive Electron." *Phys. Rev.* 43 (1933), pp. 491–494. DOI: 10.1103/PhysRev.43.491.

[6] G. Aad et al., (ATLAS Collaboration), "Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC." *Phys. Lett. B* 716.1 (2012), pp. 1 –29. DOI: 10.1016/j.physletb.2012.08.020.

[7] S. Chatrchyan et al., (CMS Collaboration), "Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC." *Phys. Lett. B* 716.1 (2012), pp. 30 –61. DOI: 10.1016/j.physletb.2012.08.021.

[8] L. Evans and P. Bryant. "LHC Machine." *JINST* 3 (2008), S08001. DOI: 10.1088/1748-0221/3/08/S08001.

[9] M. D. Schwartz. *Quantum Field Theory and the Standard Model*. Cambridge University Press, 2014. ISBN: 978-1-107-03473-0.

[10] F. Zwicky. "On the Masses of Nebulae and of Clusters of Nebulae." *Astrophys. J.* 86 (1937), p. 217. DOI: 10.1086/143864.

[11] V. C. Rubin and W. K. Ford Jr. "Rotation of the Andromeda Nebula from a Spectroscopic Survey of Emission Regions." *Astrophys. J.* 159 (1970), p. 379. DOI: 10.1086/150317.

[12] A. S. Bolton et al. "The Sloan Lens ACS Survey. V. The Full ACS Strong-Lens Sample." *Astrophys. J.* 682.2 (2008), pp. 964–984. DOI: 10.1086/589327.

[13] Y. V Kozlov, V. P Martem'yanov, and K. N Mukhin. "Neutrino mass problem: the state of the art." *Phys.-Usp.* 40.8 (1997), pp. 807–842. DOI: 10.1070/pu1997v040n08abeh000273.

[14] L. Canetti, M. Drewes, and M. Shaposhnikov. "Matter and antimatter in the universe." *New J. Phys.* 14.9 (2012), p. 095012. DOI: 10.1088/1367-2630/14/9/095012.

[15] A. G. Riess et al. "Observational Evidence from Supernovae for an Accelerating Universe and a Cosmological Constant." *Astron. J.* 116.3 (1998), pp. 1009–1038. DOI: 10.1086/300499.

[16] S. Perlmutter et al., (The Supernova Cosmology Project), "Measurements of Ω and Λ from 42 High-Redshift Supernovae." *Astrophys. J.* 517.2 (1999), pp. 565–586. DOI: 10.1086/307221.

[17] K. Whalen, (ATLAS Collaboration), "BSM searches at the LHC" (2019). URL: http://cds.cern.ch/record/2654369.

[18] H. A. Lorentz. *Versuch einer Theorie der electrischen und optischen Erscheinungen in bewegten Körpern*. Cambridge University Press, 2013. DOI: 10.1017/CBO9781139381406.

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

[20] R. Wideröe. "Über ein neues Prinzip zur Herstellung hoher Spannungen." *Arch. Elektrotech.* 21.4 (1928), pp. 387–406. DOI: 10.1007/BF01656341.

[21] J. W. Wang and G. A. Loew. "Field emission and RF breakdown in high gradient room temperature linac structures." 1997, pp. 768–794. DOI: 10.2172/663321.

[22] V. A. Dolgashev et al. "Study of RF Breakdown in Normal Conducting Cryogenic Structure." *Conf. Proc. C* 1205201 (2012), pp. 3368–3370. URL: https://accelconf.web.cern.ch/IPAC2012/papers/THPPC039.PDF.

[23] M. C. Thompson et al. "Breakdown Limits on Gigavolt-per-Meter Electron-Beam-Driven Wakefields in Dielectric Structures." *Phys. Rev. Lett.* 100 (2008), p. 214801. DOI: 10.1103/PhysRevLett.100.214801.

[24] J. T. Seeman. "The Stanford Linear Collider." *Ann. Rev. Nucl. Part. Sci.* 41 (1991), pp. 389–428. DOI: 10.1146/annurev.ns.41.120191.002133.

[25] T. Behnke et al. *The International Linear Collider Technical Design Report*. Tech. rep. ILC-REPORT-2013-040. 2013. URL: https://cds.cern.ch/record/1601970.

[26] P. N. Burrows et al., (CLIC Collaboration), *The Compact Linear Collider (CLIC) - 2018 Summary Report*. CERN Yellow Reports: Monographs. 2018. DOI: 10.23731/CYRM-2018-002.

[27] J. D. Jackson. *Classical Electrodynamics*. 3rd ed. New York: John Wiley & Sons, 1999. ISBN: 978-0-471-3093201.

[28] R. Aßmann, M. Lamont, and S. Myers. "A brief history of the LEP collider." *Nucl. Phys. B - Proc. Sup.* 109.2 (2002), pp. 17–31. DOI: 10.1016/S0920-5632(02)90005-8.

[29] A. Abada et al., (FCC Collaboration), "FCC-hh: The Hadron Collider: Future Circular Collider Conceptual Design Report Volume 3." *Eur. Phys. J. ST* 228.4 (2019), pp. 755–1107. DOI: 10.1140/epjst/e2019-900087-0.

[30] A. Abada et al., (FCC Collaboration), "FCC-ee: The Lepton Collider: Future Circular Collider Conceptual Design Report Volume 2." *Eur. Phys. J. ST* 228.2 (2019), pp. 261–623. DOI: 10.1140/epjst/e2019-900045-4.

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

[32] C. M. Tang, P. Sprangle, and R. N. Sudan. "Excitation of the plasma waves in the laser beat wave accelerator." *Appl. Phys. Lett.* 45.4 (1984), pp. 375–377.

[33] C. E. Clayton et al. "Relativistic Plasma-Wave Excitation by Collinear Optical Mixing." *Phys. Rev. Lett.* 54 (1985), pp. 2343–2346. DOI: 10.1103/PhysRevLett.54.2343.

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

[35] D. Strickland and G. Mourou. "Compression of amplified chirped optical pulses." *Opt. Commun.* 55.6 (1985), pp. 447–449. DOI: 10.1016/0030-4018(85)90151-8.

[36] K. Nakajima et al. "A proof-of-principle experiment of laser wakefield acceleration." *Phys. Scr.* T52 (1994), pp. 61–64. DOI: 10.1088/0031-8949/1994/t52/009.

[37] A. Ting et al. "Plasma wakefield generation and electron acceleration in a self-modulated laser wakefield accelerator experiment." *Phys. Plasmas* 4.5 (1997), pp. 1889–1899. DOI: 10.1063/1.872332.

[38] F. Amiranoff et al. "Observation of Laser Wakefield Acceleration of Electrons." *Phys. Rev. Lett.* 81 (1998), pp. 995–998. DOI: 10.1103/PhysRevLett.81.995.

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

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

[41] J. Faure et al. "A laser–plasma accelerator producing monoenergetic electron beams." *Nature* 431.7008 (2004), pp. 541–544. DOI: 10.1038/nature02963.

[42] S. M. Hooker. "Developments in laser-driven plasma accelerators." *Nat. Photonics* 7.10 (2013), pp. 775–782. DOI: 10.1038/nphoton.2013.234.

[43] J. P. Palastro et al. "Dephasingless Laser Wakefield Acceleration." *Phys. Rev. Lett.* 124 (2020), p. 134802. DOI: 10.1103/PhysRevLett.124.134802.

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

[45] R. D. Ruth et al. "A Plasma Wake Field Accelerator." *Part. Accel.* 17 (1985), p. 171. URL: https://cds.cern.ch/record/157249/files/p171.pdf.

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

[47] P. Sprangle, E. Esarey, and A. Ting. "Nonlinear interaction of intense laser pulses in plasmas." *Phys. Rev. A* 41 (1990), pp. 4463–4469. DOI: 10.1103/PhysRevA.41.4463.

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

[49] M. J. Hogan et al. "E-157: A 1.4-m-long plasma wake field acceleration experiment using a 30 GeV electron beam from the Stanford Linear Accelerator Center Linac." *Phys. Plasmas* 7.5 (2000), pp. 2241–2248. DOI: 10.1063/1.874059.

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

[51] M. J. Hogan et al. "Multi-GeV Energy Gain in a Plasma-Wakefield Accelerator." *Phys. Rev. Lett.* 95 (2005), p. 054802. DOI: 10.1103/PhysRevLett.95.054802.

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

[53] S. Lee et al. "Energy doubler for a linear collider." *Phys. Rev. ST Accel. Beams* 5 (2002), p. 011001. DOI: 10.1103/PhysRevSTAB.5.011001.

[54] E. Adli et al. *A Beam Driven Plasma-Wakefield Linear Collider: From Higgs Factory to Multi-TeV*. 2013. arXiv: 1308.1145 [physics.acc-ph].

[55] M. J. Hogan et al. "Plasma wakefield acceleration experiments at FACET." *New J. Phys.* 12.5 (2010), p. 055030. DOI: 10.1088/1367-2630/12/5/055030.

[56] C. Joshi et al. "Plasma wakefield acceleration experiments at FACET II." *Plasma Phys. Control. Fusion* 60 (2018), p. 034001. DOI: 10.1088/1361-6587/aaa2e3.

[57] R. D'Arcy et al. "FLASHForward: Plasma-wakefield accelerator science for high-average-power applications." *Philos. T. R. Soc. A* 377 (2019), p. 2018.0392. DOI: 10.1098/rsta.2018.0392.

[58] B. Faatz et al. "Simultaneous operation of two soft x-ray free-electron lasers driven by one linear accelerator." *New J. Phys.* 18.6 (2016), p. 062002. DOI: 10.1088/1367-2630/18/6/062002.

[59] A. Caldwell et al. "Proton-driven plasma-wakefield acceleration." *Nat. Phys.* 5 (2009), pp. 363–367. DOI: 10.1038/nphys1248.

[60] A. Caldwell and K. V. Lotov. "Plasma wakefield acceleration with a modulated proton bunch." *Phys. Plasmas* 18 (2011), p. 103101. DOI: 10.1063/1.3641973.

[61] P. Muggli et al., (AWAKE Collaboration), "AWAKE readiness for the study of the seeded self-modulation of a 400 GeV proton bunch." *Plasma Phys. Control. Fusion* **60** (2017), p. 014046. DOI: 10.1088/1361-6587/aa941c.

[62] C. B. Schroeder et al. "Growth and Phase Velocity of Self-Modulated Beam-Driven Plasma Waves." *Phys. Rev. Lett.* 107 (2011), p. 145002. DOI: 10.1103/PhysRevLett.107.145002.

[63] A. Butler, D. J. Spence, and S. M. Hooker. "Guiding of High-Intensity Laser Pulses with a Hydrogen-Filled Capillary Discharge Waveguide." *Phys. Rev. Lett.* 89 (2002), p. 185003. DOI: 10.1103/PhysRevLett.89.185003.

[64] S. Karsch et al. "GeV-scale electron acceleration in a gas-filled capillary discharge waveguide." *New J. Phys.* 9.11 (2007), pp. 415–415. DOI: 10.1088/1367-2630/9/11/415.

[65] W. P. Leemans et al. "Multi-GeV Electron Beams from Capillary-Discharge-Guided Subpetawatt Laser Pulses in the Self-Trapping Regime." *Phys. Rev. Lett.* 113 (2014), p. 245002. DOI: 10.1103/PhysRevLett.113.245002.

[66] A. J. Gonsalves et al. "Petawatt Laser Guiding and Electron Beam Acceleration to 8 GeV in a Laser-Heated Capillary Discharge Waveguide." *Phys. Rev. Lett.* 122 (2019), p. 084801. DOI: 10.1103/PhysRevLett.122.084801.

[67] M. Litos et al. "High-efficiency acceleration of an electron beam in a plasma wakefield accelerator." *Nature* 515.7525 (2014), pp. 92–95. DOI: 10.1038/nature13882.

[68] C. A. Lindstrøm et al. "Energy-Spread Preservation and High Efficiency in a Plasma-Wakefield Accelerator." *Phys. Rev. Lett.* 126 (2021), p. 014801. DOI: 10.1103/PhysRevLett.126.014801.

[69] K. Hanke et al. "The LHC Injectors Upgrade (LIU) Project at CERN: Proton Injector Chain." *8th International Particle Accelerator Conference*. 2017, WEPVA036. DOI: 10.18429/JACoW-IPAC2017-WEPVA036.

[70] I. Béjar Alonso et al. "High-Luminosity Large Hadron Collider (HL-LHC): Technical design report" (2020). DOI: 10.23731/CYRM-2020-0010.

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

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

[73] J. Grebenyuk et al. "Beam-driven plasma-based acceleration of electrons with density down-ramp injection at FLASHForward." *Nucl. Instrum. Meth. A* 740 (2014), pp. 246–249. DOI: 10.1016/j.nima.2013.10.054.

[74] B. Hidding et al. "First Measurements of Trojan Horse Injection in a Plasma Wakefield Accelerator." *8th International Particle Accelerator Conference*. 2017, TUYB1. DOI: 10.18429/JACoW-IPAC2017-TUYB1.

[75] A. Deng et al. "Generation and acceleration of electron bunches from a plasma photocathode." *Nat. Phys.* 15.11 (2019), pp. 1156–1160. DOI: 10.1038/s41567-019-0610-9.

[76] A. Knetsch et al. *Controlled density-downramp injection in a beam-driven plasma wakefield accelerator*. 2020. arXiv: 2007.12639 [physics.acc-ph].

[77] K. Nakamura et al. "Diagnostics, Control and Performance Parameters for the BELLA High Repetition Rate Petawatt Class Laser." *IEEE J. Quantum Elect.* 53.4 (2017), pp. 1–21. DOI: 10.1109/JQE.2017.2708601.

[78] N. Delbos et al. "LUX – A laser–plasma driven undulator beamline." *Nucl. Instrum. Meth. A* 909 (2018), pp. 318 –322. DOI: 10.1016/j.nima.2018.01.082.

[79] C. J. Hooker et al. "The Astra Gemini project - A dual-beam petawatt Ti:Sapphire laser system." *J. Phys. IV France* 133 (2006), pp. 673–677. DOI: 10.1051/jp4:2006133135.

[80] W. Leemans et al. "Report of workshop on Laser Technology for k-BELLA and Beyond, Lawrence Berkeley National Laboratory" (2017). URL: https://www2.lbl.gov/LBL-Programs/atap/Report_Workshop_k-BELLA_laser_tech_final.pdf.

[81] W. Leemans et al. "KALDERA – High average power laser plasma accelerator project at DESY." 4th European Advanced Accelerator Concepts Workshop. 2019. URL: https://agenda.infn.it/event/17304/contributions/97644/.

[82] G. S. Voronov. "A practical fit formula for ionization rate coefficients of atoms and ions by electron impact: Z= 1–28." *Atom. Data Nucl. Data* 65.1 (1997), pp. 1 –35. DOI: https://doi.org/10.1006/adnd.1997.0732.

[83] M. F. Gilljohann et al. "Direct Observation of Plasma Waves and Dynamics Induced by Laser-Accelerated Electron Beams." *Phys. Rev. X* 9 (2019), p. 011046. DOI: 10.1103/PhysRevX.9.011046.

[84] R. Zgadzaj et al. "Dissipation of electron-beam-driven plasma wakes." *Nat. Commun.* 11.1 (2020), p. 4753. DOI: 10.1038/s41467-020-18490-w.

[85] R. Keinigs and M. E. Jones. "Two-dimensional dynamics of the plasma wakefield accelerator." *Phys. Fluids* 30 (1987), pp. 252–263. DOI: 10.1063/1.866183.

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

[87] W. Lu et al. "A nonlinear theory for multidimensional relativistic plasma wave wakefields." *Phys. Plasmas* 13.5 (2006), p. 056709. DOI: 10.1063/1.2203364.

[88] E. P. Lee and R. K. Cooper. "General Envelope Equation For Cylindrically Symmetric Charged Particle Beams." *Part. Accel.* 7 (1976), pp. 83–95. URL: https://core.ac.uk/download/pdf/44196704.pdf.

[89] J. Krall and G. Joyce. "Transverse equilibrium and stability of the primary beam in the plasma wake-field accelerator." *Phys. Plasmas* 2.4 (1995), pp. 1326–1331. DOI: 10.1063/1.871344.

[90] E. Esarey et al. "Synchrotron radiation from electron beams in plasma-focusing channels." *Phys. Rev. E* 65 (2002), p. 056505. DOI: 10.1103/PhysRevE.65.056505.

[91] E. D. Courant and H. S. Snyder. "Theory of the alternating-gradient synchrotron." *Ann. Phys.* 3.1 (1958), pp. 1–48. DOI: 0.1016/0003-4916(58)90012-5.

[92] C. E. Clayton et al. "Transverse Envelope Dynamics of a 28.5-GeV Electron Beam in a Long Plasma." *Phys. Rev. Lett.* 88 (2002), p. 154801. DOI: 10.1103/PhysRevLett.88.154801.

[93] P. Muggli et al. "Meter-Scale Plasma-Wakefield Accelerator Driven by a Matched Electron Beam." *Phys. Rev. Lett.* 93 (2004), p. 014802. DOI: 10.1103/PhysRevLett.93.014802.

[94] N. Barov and J. B. Rosenzweig. "Propagation of short electron pulses in underdense plasmas." *Phys. Rev. E* 49 (1994), pp. 4407–4416. DOI: 10.1103/PhysRevE.49.4407.

[95] K. V. Lotov. "Blowout regimes of plasma wakefield acceleration." *Phys. Rev. E* 69 (2004), p. 046405. DOI: 10.1103/PhysRevE.69.046405.

[96] P. Mora and T. M. Antonsen Jr. "Kinetic modeling of intense, short laser pulses propagating in tenuous plasmas." *Phys. Plasmas* 4.1 (1997), pp. 217–229. DOI: 10.1063/1.872134.

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

[98] S. van der Meer. *Improving the power efficiency of the plasma wakefield accelerator*. Tech. rep. CERN-PS-85-65-AA. 1985. URL: https://cds.cern.ch/record/163918.

[99] T. Katsouleas et al. "Beam loading in plasma accelerators." *Part. Acc.* 22 (1987), pp. 81–99. URL: https://s3.cern.ch/inspire-prod-files-3/30e3b8e467eb7298ce19d6bbc95b457f.

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

[101] M. Tzoufras et al. "Beam loading by electrons in nonlinear plasma wakes." *Phys. Plasmas* 16.5 (2009), p. 056705. DOI: 10.1063/1.3118628.

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

[103] M. C. Downer et al. "Diagnostics for plasma-based electron accelerators." *Rev. Mod. Phys.* 90 (2018), p. 035002. DOI: 10.1103/RevModPhys.90.035002.

[104] M. B. Schwab et al. "Few-cycle optical probe-pulse for investigation of relativistic laser-plasma interactions." *Appl. Phys. Lett.* 103.19 (2013), p. 191118. DOI: 10.1063/1.4829489.

[105] S. Schröder et al. "High-resolution sampling of beam-driven plasma wakefields." *Nat. Commun.* 11.1 (2020), p. 5984. DOI: 10.1038/s41467-020-19811-9.

[106] J. M. Dawson. "Particle simulation of plasmas." *Rev. Mod. Phys.* 55 (1983), pp. 403–447. DOI: 10.1103/RevModPhys.55.403.

[107] R. W. Hockney and J. W. Eastwood. *Computer Simulation Using Particles*. CRC Press, 1988. ISBN: 9781439822050.

[108] C. K. Birdsall and A. B. Langdon. *Plasma Physics via Computer Simulation*. CRC Press, 2018. ISBN: 9781482263060.

[109] T. J. Mehrling. "Theoretical and numerical studies on the transport of transverse beam quality in plasma-based accelerators." PhD thesis. Hamburg U., 2014. DOI: 10.3204/DESY-THESIS-2014-040.

[110] K. Yee. "Numerical solution of initial boundary value problems involving maxwell's equations in isotropic media." *IEEE T. Antenn. Propag.* 14.3 (1966), pp. 302–307. DOI: 10.1109/TAP.1966.1138693.

[111] B. B. Godfrey. "Numerical Cherenkov instabilities in electromagnetic particle codes." *J. Comput. Phys.* 15.4 (1974), pp. 504 –521. DOI: https://doi.org/10.1016/0021-9991(74)90076-X.

[112] R. Lehe et al. "Numerical growth of emittance in simulations of laser-wakefield acceleration." *Phys. Rev. ST Accel. Beams* 16 (2013), p. 021301. DOI: 10.1103/PhysRevSTAB.16.021301.

[113] H. Meuer et al. *Top 500 supercomputers*. 2020. URL: https://top500.org/lists/top500/.

[114] P. Mora and T. M. Antonsen. "Electron cavitation and acceleration in the wake of an ultraintense, self-focused laser pulse." *Phys. Rev. E* 53 (1996), R2068–R2071. DOI: 10.1103/PhysRevE.53.R2068.

[115] D. H. Whittum. "Transverse two-stream instability of a beam with a Bennett profile." *Phys. Plasmas* 4.4 (1997), pp. 1154–1159. DOI: 10.1063/1.872202.

[116] C. Huang et al. "QUICKPIC: A highly efficient particle-in-cell code for modeling wakefield acceleration in plasmas." *J. Comput. Phys.* 217.2 (2006), pp. 658 –679. DOI: 10.1016/j.jcp.2006.01.039.

[117] T. J. Mehrling et al. "HiPACE: a quasi-static particle-in-cell code." *Plasma Phys. Control. Fusion* 56 (2014), p. 084012. DOI: 10.1088/0741-3335/56/8/084012.

[118] K. V. Lotov. "Fine wakefield structure in the blowout regime of plasma wakefield accelerators." *Phys. Rev. ST Accel. Beams* 6 (2003), p. 061301. DOI: 10.1103/PhysRevSTAB.6.061301.

[119] K. V. Lotov and A. P. Sosedkin. *LCODE user manual*. 2020. URL: https://lcode.info/site-files/manual.pdf.

[120] N. Kumar, A. Pukhov, and K. Lotov. "Self-Modulation Instability of a Long Proton Bunch in Plasmas." *Phys. Rev. Lett.* 104 (2010), p. 255003. DOI: 10.1103/PhysRevLett.104.255003.

[121] J. Vieira, W. B. Mori, and P. Muggli. "Hosing Instability Suppression in Self-Modulated Plasma Wakefields." *Phys. Rev. Lett.* 112 (2014), p. 205001. DOI: 10.1103/PhysRevLett.112.205001.

[122] J.-L. Vay et al. "Warp-X: A new exascale computing platform for beam–plasma simulations." *Nucl. Instrum. Meth. A* 909 (2018), pp. 476 –479. DOI: 10.1016/j.nima.2018.01.035.

[123] H. Burau et al. "PIConGPU: A Fully Relativistic Particle-in-Cell Code for a GPU Cluster." *IEEE T. Plasma Sci.* 38.10 (2010), pp. 2831–2839. DOI: 10.1109/TPS.2010.2064310.

[124] E. Cormier-Michel et al. "Unphysical kinetic effects in particle-in-cell modeling of laser wakefield accelerators." *Phys. Rev. E* 78 (2008), p. 016404. DOI: 10.1103/PhysRevE.78.016404.

[125] R. W. Hockney. "Measurements of collision and heating times in a two-dimensional thermal computer plasma." *J. Comput. Phys.* 8.1 (1971), pp. 19 –44. DOI: 10.1016/0021-9991(71)90032-5.

[126] A. B. Langdon. "Effects of the spatial grid in simulation plasmas." *J. Comput. Phys.* 6.2 (1970), pp. 247 –267. DOI: 10.1016/0021-9991(70)90024-0.

[127] C. K. Birdsall and N. Maron. "Plasma self-heating and saturation due to numerical instabilities." *J. Comput. Phys.* 36.1 (1980), pp. 1 –19. DOI: 10.1016/0021-9991(80)90171-0.

[128] J. D. Huba. *NRL Plasma Formulary*. Naval Research Laboratory, 2019, pp. 1–71. URL: http://wwwppd.nrl.navy.mil/nrlformulary/.

[129] T. D. Arber et al. "Contemporary particle-in-cell approach to laser-plasma modelling." *Plasma Phys. Control. Fusion* 57.11 (2015), p. 113001. DOI: 10.1088/0741-3335/57/11/113001.

[130] Y. Sentoku and A. J. Kemp. "Numerical methods for particle simulations at extreme densities and temperatures: Weighted particles, relativistic collisions and reduced currents." *J. Comput. Phys.* 227.14 (2008), pp. 6846 –6861. DOI: 10.1016/j.jcp.2008.03.043.

[131] D. W. Hewett and C. W. Nielson. "A multidimensional quasineutral plasma simulation model." *J. Comput. Phys.* 29.2 (1978), pp. 219 –236. DOI: 10.1016/0021-9991(78)90153-5.

[132] D. W. Hewett. "A global method of solving the electron-field equations in a zero-inertia-electron-hybrid plasma simulation code." *J. Comput. Phys.* 38.3 (1980), pp. 378 –395. DOI: 10.1016/0021-9991(80)90155-2.

[133] J. J. Honrubia and J. Meyer ter Vehn. "Fast ignition of fusion targets by laser-driven electrons." *Plasma Phys. Control. Fusion* 51.1 (2008), p. 014008. DOI: 10.1088/0741-3335/51/1/014008.

[134] B. I. Cohen, A. J. Kemp, and L. Divol. "Simulation of laser–plasma interactions and fast-electron transport in inhomogeneous plasma." *J. Comput. Phys.* 229.12 (2010), pp. 4591 –4612. DOI: 10.1016/j.jcp.2010.03.001.

[135] C. A. Lindstrøm et al. "Staging optics considerations for a plasma wakefield acceleration linear collider." *Nucl. Instrum. Meth. A* 829 (2016), pp. 224 –228. DOI: https://doi.org/10.1016/j.nima.2015.12.065.

[136] C. A. Lindstrøm. "Staging of plasma-wakefield accelerators." *Phys. Rev. Accel. Beams* 24 (2021), p. 014801. DOI: 10.1103/PhysRevAccelBeams.24.014801.

[137] S. Steinke et al. "Multistage coupling of independent laser-plasma accelerators." *Nature* 530 (2016), pp. 190–193. DOI: 10.1038/nature16525.

[138] C. Bracco et al. "Beam studies and experimental facility for the AWAKE experiment at CERN." *Nucl. Instrum. Meth. A* 740 (2014), pp. 48 –53. DOI: 10.1016/j.nima.2013.10.060.

[139] P. Sprangle, E. Esarey, and A. Ting. "Nonlinear theory of intense laser-plasma interactions." *Phys. Rev. Lett.* 64 (1990), pp. 2011–2014. DOI: 10.1103/PhysRevLett.64.2011.

[140] E. Esarey, J. Krall, and P. Sprangle. "Envelope analysis of intense laser pulse self-modulation in plasmas." *Phys. Rev. Lett.* 72 (1994), pp. 2887–2890. DOI: 10.1103/PhysRevLett.72.2887.

[141] N. E. Andreev, V. I. Kirsanov, and L. M. Gorbunov. "Stimulated processes and self-modulation of a short intense laser pulse in the laser wake-field accelerator." *Phys. Plasmas* 2.6 (1995), pp. 2573–2582. DOI: 10.1063/1.871219.

[142] K. V. Lotov. "Physics of beam self-modulation in plasma wakefield accelerators." *Phys. Plasmas* 22 (2015), p. 103110. DOI: 10.1063/1.4933129.

[143] P. Muggli et al., (AWAKE Collaboration), "Physics to plan AWAKE Run 2." *J. Phys.: Conference Series* 1596 (2020), p. 012008. DOI: 10.1088/1742-6596/1596/1/012008.

[144] E. Adli et al., (AWAKE Collaboration), "Acceleration of electrons in the plasma wakefield of a proton bunch." *Nature* 561 (2018), pp. 363–367. DOI: 10.1038/s41586-018-0485-4.

[145] K. Pepitone et al. "The electron accelerators for the AWAKE experiment at CERN – Baseline and Future Developments." *Nucl. Instrum. Meth. A* 909 (2018), pp. 102 –106. DOI: 10.1016/j.nima.2018.02.044.

[146] J. Bauche et al. "A magnetic spectrometer to measure electron bunches accelerated at AWAKE." *Nucl. Instrum. Meth. A* 940 (2019), pp. 103 –108. DOI: 10.1016/j.nima.2019.05.067.

[147] L. J. Nevay et al. "BDSIM: An accelerator tracking code with particle–matter interactions." *Comput. Phys. Commun.* 252 (2020), p. 107200. DOI: 10.1016/j.cpc.2020.107200.

[148] S. Agostinelli et al. "Geant4 – a simulation toolkit." *Nucl. Instrum. Meth. A* 506 (2003), pp. 250 –303. DOI: 10.1016/S0168-9002(03)01368-8.

[149] F. Keeble. "Measurement of the electron energy distribution at AWAKE." PhD thesis. University College London, 2019. URL: https://discovery.ucl.ac.uk/id/eprint/10082378.

[150] D. Cooke et al. "Calibration of the AWAKE electron spectrometer with electrons derived from a partially stripped ion beam." *10th International Particle Accelerator Conference (IPAC2019)*. 2019. DOI: 10.18429/JACoW-IPAC2019-WEPGW089.

[151] E. Öz and P. Muggli. "A novel Rb vapor plasma source for plasma wakefield accelerators." *Nucl. Instrum. Meth. A* 740 (2014), pp. 197 –202. DOI: 10.1016/j.nima.2013.10.093.

[152] F. Batsch et al. "Interferometer-based high-accuracy white light measurement of neutral rubidium density and gradient at AWAKE." *Nucl. Instrum. Meth. A* 909 (2018), pp. 359 –363. DOI: 10.1016/j.nima.2018.02.067.

[153] E. Adli et al., (AWAKE Collaboration), "Experimental Observation of Proton Bunch Modulation in a Plasma at Varying Plasma Densities." *Phys. Rev. Lett.* 122 (2019), p. 054802. DOI: 10.1103/PhysRevLett.122.054802.

[154] M. Turner et al., (AWAKE Collaboration), "Experimental Observation of Plasma Wakefield Growth Driven by the Seeded Self-Modulation of a Proton Bunch." *Phys. Rev. Lett.* 122 (2019), p. 054801. DOI: 10.1103/PhysRevLett.122.054801.

[155] A. Campillo and S. Shapiro. "Picosecond streak camera fluorometry - A review." *IEEE J. Quantum Elect.* 19.4 (1983), pp. 585–603. DOI: 10.1109/JQE.1983.1071909.

[156] A. Petrenko, K. Lotov, and A. Sosedkin. "Numerical studies of electron acceleration behind self-modulating proton beam in plasma with a density gradient." *Nucl. Instrum. Meth. A* 829 (2016), pp. 63 –66. DOI: 10.1016/j.nima.2016.01.063.

[157] W. Ackermann et al. "Operation of a free-electron laser from the extreme ultraviolet to the water window." *Nat. Photonics* 1.6 (2007), pp. 336–342. DOI: 10.1038/nphoton.2007.76.

[158] K. Tiedtke et al. "The soft x-ray free-electron laser FLASH at DESY: beamlines, diagnostics and end-stations." *New J. Phys.* 11.2 (2009), p. 023029. DOI: 10.1088/1367-2630/11/2/023029.

[159] S. Schröder et al. "Tunable and precise two-bunch generation at FLASHForward." *J. Phys.: Conference Series* 1596 (2020), p. 012002. DOI: 10.1088/1742-6596/1596/1/012002.

[160] P. Piot et al. "Generation and Characterization of Electron Bunches with Ramped Current Profiles in a Dual-Frequency Superconducting Linear Accelerator." *Phys. Rev. Lett.* 108 (2012), p. 034801. DOI: 10.1103/PhysRevLett.108.034801.

[161] K. Floettmann, T. Limberg, and P. Piot. *Generation of Ultrashort Electron Bunches by Cancellation of Nonlinear Distortions in the Longitudinal Phase Space*. Tech. rep. TESLA FEL 2001-06. 2001, p. 15. URL: https://bib-pubdb1.desy.de/record/409209.

[162] J. Dale et al. "Vorrichtung Mit Beweglicher Aufnahme Für Vakuumkammern." Pat. DE 102014116476 A1 (Germany). May 12, 2016. URL: https://lens.org/179-396-520-666-03X.

[163] G. A. Loew and O. H. Altenmueller. "Design And Applications Of R.F. Deflecting Structures At SLAC." *5th International Conference on High-Energy Accelerators*. 1965, pp. 551–555. URL: https://inspirehep.net/files/c1be43e925db8226d80778005361b3a7.

[164] P. J. Emma, J. Frisch, and P. Krejcik. *A Transverse RF Deflecting Structure for Bunch Length and Phase Space Diagnostics*. Tech. rep. LCLS-TN-00-12. 2000. URL: https://www-ssrl.slac.stanford.edu/lcls/technotes/lcls-tn-00-12.pdf.

[165] R. D'Arcy, V. Libov, and J. Osterhoff. "A Transverse Deflecting Structure for the Plasma Wakefield Accelerator Experiment, FLASHForward." *5th International Beam Instrumentation Conference*. 2017. DOI: 10.18429/JACoW-IBIC2016-WEPG51.

[166] B. Marchetti et al. "X-Band TDS Project." *8th International Particle Accelerator Conference*. 2017. DOI: 10.18429/JACoW-IPAC2017-MOPAB044.

[167] R. D'Arcy et al. "Longitudinal Phase Space Reconstruction at FLASHForward Using a Novel Transverse Deflection Cavity, PolariX-TDS." *9th International Particle Accelerator Conference*. 2018. DOI: 10.18429/JACoW-IPAC2018-TUPML017.

[168] M. Röhrs et al. "Time-resolved electron beam phase space tomography at a soft x-ray free-electron laser." *Phys. Rev. ST Accel. Beams* 12 (2009), p. 050704. DOI: 10.1103/PhysRevSTAB.12.050704.

[169] A. Grudiev. *Design of compact high power RF components at X-band*. Tech. rep. CERN-ACC-NOTE-2016-0044. 2016. URL: https://cds.cern.ch/record/2158484.

[170] D. Marx et al. "Reconstruction of the 3D charge distribution of an electron bunch using a novel variable-polarization transverse deflecting structure (TDS)." *J. Phys.: Conference Series* 874 (2017), p. 012077. DOI: 10.1088/1742-6596/874/1/012077.

[171] A. Bolzmann. "Investigation of the Longitudinal Charge Distribution of Electron Bunches at the VUV-FEL using the Transverse Deflecting Cavity LOLA." MA thesis. Wurzburg U., 2005. DOI: 10.3204/DESY-THESIS-2005-046.

[172] C. Behrens et al. "Constraints on photon pulse duration from longitudinal electron beam diagnostics at a soft x-ray free-electron laser." *Phys. Rev. ST Accel. Beams* 15 (2012), p. 030707. DOI: 10.1103/PhysRevSTAB.15.030707.

[173] C. A. Lindstrøm et al. "Matching small β functions using centroid jitter and two beam position monitors." *Phys. Rev. Accel. Beams* 23 (2020), p. 052802. DOI: 10.1103/PhysRevAccelBeams.23.052802.

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

[175] R. J. England, J. B. Rosenzweig, and N. Barov. "Plasma electron fluid motion and wave breaking near a density transition." *Phys. Rev. E* 66 (2002), p. 016501. DOI: 10.1103/PhysRevE.66.016501.

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

[177] 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." *Phys. Rev. Accel. Beams* 22 (2019), p. 111301. DOI: 10.1103/PhysRevAccelBeams.22.111301.

[178] A. J. Gonsalves et al. "Tunable laser plasma accelerator based on longitudinal density tailoring." *Nat. Phys.* 7.11 (2011), pp. 862–866. DOI: 10.1038/nphys2071.

[179] J. Faure et al. "Injection and acceleration of quasimonoenergetic relativistic electron beams using density gradients at the edges of a plasma channel." *Phys. Plasmas* 17.8 (2010), p. 083107. DOI: 10.1063/1.3469581.

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

[181] E. Oz et al. "Ionization-induced electron trapping in ultrarelativistic plasma wakes." *Phys. Rev. Lett.* 98 (2007), p. 084801. DOI: 10.1103/PhysRevLett.98.084801.

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

[183] N. Vafaei-Najafabadi et al. "Evidence for high-energy and low-emittance electron beams using ionization injection of charge in a plasma wakefield accelerator." *Plasma Phys. Control. Fusion* 58.3 (2016), p. 034009. DOI: 10.1088/0741-3335/58/3/034009.

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

[185] G. Wittig et al. "Electron beam manipulation, injection and acceleration in plasma wakefield accelerators by optically generated plasma density spikes." *Nucl. Instrum. Meth. A* 829 (2016), pp. 83 –87. DOI: 10.1016/j.nima.2016.02.027.

[186] M. V. Ammosov, N. B. Delone, and V. P. Krainov. "Tunnelling ionization of complex atoms and of atomic ions in an alternating electromagnetic field." *Sov. Phys. JETP* 64 (1986), p. 1191. DOI: 10.1117/12.938695.

[187] K. V. Lotov. "Efficient operating mode of the plasma wakefield accelerator." *Phys. Plasmas* 12.5 (2005), p. 053105. DOI: 10.1063/1.1889444.

[188] X. L. Xu et al. "Physics of Phase Space Matching for Staging Plasma and Traditional Accelerator Components Using Longitudinally Tailored Plasma Profiles." *Phys. Rev. Lett.* 116 (2016), p. 124801. DOI: 10.1103/PhysRevLett.116.124801.

[189] J. Frederico, M. Hogan, and T. Raubenheimer. "Beam match and emittance saturation within plasma wakefield acceleration." *AIP Conf. Proc.* 1812.1 (2017), p. 070001. DOI: 10.1063/1.4975881.

[190] M. D. Litos et al. "Beam emittance preservation using Gaussian density ramps in a beam-driven plasma wakefield accelerator." *Philos. T. R. Soc. A* 377.2151 (2019), p. 20180181. DOI: 10.1098/rsta.2018.0181.

[191] F. Salehi et al. "MeV electron acceleration at 1 kHz with < 10 mJ laser pulses." *Opt. Lett.* 42.2 (2017), pp. 215–218. DOI: 10.1364/OL.42.000215.

[192] D. Gustas et al. "High-charge relativistic electron bunches from a kHz laser-plasma accelerator." *Phys. Rev. Accel. Beams* 21 (2018), p. 013401. DOI: 10.1103/PhysRevAccelBeams.21.013401.

[193] L. Rovige et al. "Demonstration of stable long-term operation of a kilohertz laser-plasma accelerator." *Phys. Rev. Accel. Beams* 23 (2020), p. 093401. DOI: 10.1103/PhysRevAccelBeams.23.093401.

[194] A. J. Gonsalves et al. "Demonstration of a high repetition rate capillary discharge waveguide." *J. Appl. Phys.* 119.3 (2016), p. 033302. DOI: 10.1063/1.4940121.

[195] J. M. Garland et al. "Combining laser interferometry and plasma spectroscopy for spatially resolved high-sensitivity plasma density measurements in discharge capillaries." *Rev. Sci. Instrum.* 92.1 (2021), p. 013505. DOI: 10.1063/5.0021117.

[196] N. A. Bobrova et al. "Simulations of a hydrogen-filled capillary discharge waveguide." *Phys. Rev. E* 65 (2001), p. 016407. DOI: 10.1103/PhysRevE.65.016407.

[197] B. Broks, K. Garloff, and J. van der Mullen. "Nonlocal-thermal-equilibrium model of a pulsed capillary discharge waveguide." *Phys. Rev. E* 71 (2005), p. 016401. DOI: 10.1103/PhysRevE.71.016401.

[198] W. Decking et al. "A MHz-repetition-rate hard X-ray free-electron laser driven by a superconducting linear accelerator." *Nat. Photonics* 14.6 (2020), pp. 391–397. DOI: 10.1038/s41566-020-0607-z.

[199] J. Stohr et al. *Linac Coherent Light Source II (LCLS-II) Conceptual Design Report*. Tech. rep. SLAC-R-978. 2011. DOI: 10.2172/1029479.

[200] J. van Tilborg et al. "Density characterization of discharged gas-filled capillaries through common-path two-color spectral-domain interferometry." *Opt. Lett.* 43.12 (2018), pp. 2776–2779. DOI: 10.1364/OL.43.002776.

[201] J. van Tilborg et al. "High-sensitivity plasma density retrieval in a common-path second-harmonic interferometer through simultaneous group and phase velocity measurement." *Phys. Plasmas* 26.2 (2019), p. 023106. DOI: 10.1063/1.5080269.

[202] J. Stark and G. Wendt. "Beobachtungen über den Effekt des elektrischen Feldes auf Spektrallinien. II. Längseffekt [Observations of the effect of the electric field on spectrial lines II. Transverse effect]." *Ann. Phys.* 348.7 (1914), pp. 983–990. DOI: 10.1002/andp.19143480703.

[203] H. Griem. *Spectral Line Broadening by Plasmas*. Elsevier Science, 2012. ISBN: 9780323150941.

[204] M. A. Gigosos and V. Cardeñoso. "New plasma diagnosis tables of hydrogen Stark broadening including ion dynamics." *J. Phys. B: At. Mol. Opt.* 29.20 (1996), pp. 4795–4838. DOI: 10.1088/0953-4075/29/20/029.

[205] M. A. Gigosos, M. Á. González, and V. Cardeñoso. "Computer simulated Balmer-alpha, -beta and -gamma Stark line profiles for non-equilibrium plasmas diagnostics." *Spectrochim. Acta B* 58.8 (2003), pp. 1489–1504. DOI: 10.1016/S0584-8547(03)00097-1.

[206] G. Boyle et al. *Plasma Sources and Diagnostics*. 2020. arXiv: 2007.08184 [physics.acc-ph].

[207] S. Gessner et al., (AWAKE Collaboration), *Evolution of a plasma column measured through modulation of a high-energy proton beam*. 2020. arXiv: 2006.09991 [physics.acc-ph].

[208] G. Loisch et al. "Plasma density measurement by means of self-modulation of long electron bunches." *Plasma Phys. Control. Fusion* 61.4 (2019), p. 045012. DOI: 10.1088/1361-6587/ab04b9.

[209] W. Lu et al. "Limits of linear plasma wakefield theory for electron or positron beams." *Phys. Plasmas* 12.6 (2005), p. 063101. DOI: 10.1063/1.1905587.

[210] J. B. Rosenzweig et al. "Energy loss of a high charge bunched electron beam in plasma: Simulations, scaling, and accelerating wakefields." *Phys. Rev. ST Accel. Beams* 7 (2004), p. 061302. DOI: 10.1103/PhysRevSTAB.7.061302.

[211] C. E. Doss et al. "Laser-ionized, beam-driven, underdense, passive thin plasma lens." *Phys. Rev. Accel. Beams* 22 (2019), p. 111001. DOI: 10.1103/PhysRevAccelBeams.22.111001.

[212] D. J. Spence, A. Butler, and S. M. Hooker. "Gas-filled capillary discharge waveguides." *J. Opt. Soc. Am. B* 20.1 (2003), pp. 138–151. DOI: 10.1364/JOSAB.20.000138.

[213] M. A. Lieberman and A. J. Lichtenberg. *Principles of Plasma Discharges and Materials Processing*. John Wiley & Sons, Ltd., 2005. DOI: 10.1002/0471724254.

[214] A. J. Gonsalves et al. "Transverse Interferometry of a Hydrogen-Filled Capillary Discharge Waveguide." *Phys. Rev. Lett.* 98 (2007), p. 025002. DOI: 10.1103/PhysRevLett.98.025002.

[215] I. Dornmair, K. Floettmann, and A. R. Maier. "Emittance conservation by tailored focusing profiles in a plasma accelerator." *Phys. Rev. ST Accel. Beams* 18 (2015), p. 041302. DOI: 10.1103/PhysRevSTAB.18.041302.

[216] R. Ariniello et al. "Transverse beam dynamics in a plasma density ramp." *Phys. Rev. Accel. Beams* 22 (2019), p. 041304. DOI: 10.1103/PhysRevAccelBeams.22.041304.

[217] W. K. H. Panofsky and W. R. Baker. "A Focusing Device for the External 350-Mev Proton Beam of the 184-Inch Cyclotron at Berkeley." *Rev. Sci. Instrum.* 21.5 (1950), pp. 445–447. DOI: 10.1063/1.1745611.

[218] J. van Tilborg et al. "Active Plasma Lensing for Relativistic Laser-Plasma-Accelerated Electron Beams." *Phys. Rev. Lett.* 115 (2015), p. 184802. DOI: 10.1103/PhysRevLett.115.184802.

[219] C. A. Lindstrøm et al. "Emittance Preservation in an Aberration-Free Active Plasma Lens." *Phys. Rev. Lett.* 121 (2018), p. 194801. DOI: 10.1103/PhysRevLett.121.194801.

[220] M. P. Anania et al. "Design of a plasma discharge circuit for particle wakefield acceleration." *Nucl. Instrum. Meth. A* 740 (2014), pp. 193 –196. DOI: 10.1016/j.nima.2013.10.053.

[221] A. Esaulov et al. "MHD simulation of a fast hollow cathode capillary discharge." *Plasma Phys. Control. Fusion* 43.4 (2001), pp. 571–588. DOI: 10.1088/0741-3335/43/4/313.

[222] M. Altarelli et al. *XFEL: The European X-Ray Free-Electron Laser. Technical design report*. Tech. rep. DESY-06-097. 2006. DOI: 10.3204/DESY_06-097.

[223] D. R. Lide et al. "CRC Handbook of Chemistry and Physics, 84th edition." *J. Am. Chem. Soc.* 126 (2004), p. 1586. DOI: 10.1021/ja0336372.

[224] G. J. Boyle et al. "Reduced model of plasma evolution in hydrogen discharge capillary plasmas." *Phys. Rev. E* 104 (2021), p. 015211. DOI: 10.1103/PhysRevE.104.015211.

[225] M. C. M. van de Sanden et al. "Thermodynamic generalization of the Saha equation for a two-temperature plasma." *Phys. Rev. A* 40 (1989), pp. 5273–5276. DOI: 10.1103/PhysRevA.40.5273.

[226] A. Fick. "Ueber Diffusion." *Ann. Phys.* 170.1 (1855), pp. 59–86. DOI: 10.1002/andp.18551700105.

[227] R. Byron Bird, W. E. Stewart, and E. N. Lightfoot. *Transport Phenomena*. 3rd ed. New York: John Wiley & Sons, 2006. ISBN: 0470115394.

[228] A. V. Phelps. "The Diffusion of Charged Particles in Collisional Plasmas: Free and Ambipolar Diffusion at Low and Moderate Pressures." *J. Res. Natl. Inst. Stan.* 95.4 (1990), pp. 407–431. DOI: 10.6028/jres.095.035.

[229] F. Paschen. "Ueber die zum Funkenübergang in Luft, Wasserstoff und Kohlensäure bei verschiedenen Drucken erforderliche Potentialdifferenz." *Ann. Phys.* 273 (1889), pp. 69–96. DOI: 10.1002/andp.18892730505.

[230] L. Frommhold and M. A. Biondi. "A mathematical study of the electron decay in diffusion and recombination controlled afterglows." *Ann. Phys.* 48.3 (1968), pp. 407–440. DOI: 10.1016/0003-4916(68)90140-1.

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

[232] T. Auguste et al. "Defocusing effects of a picosecond terawatt laser pulse in an underdense plasma." *Opt. Commun.* 89.2 (1992), pp. 145 –148. DOI: 10.1016/0030-4018(92)90148-K.

[233] S. C. Rae. "Ionization-induced defocusing of intense laser pulses in high-pressure gases." *Opt. Commun.* 97.1 (1993), pp. 25 –28. DOI: 10.1016/0030-4018(93)90611-8.

[234] G. I. Taylor. "The formation of a blast wave by a very intense explosion I. Theoretical discussion." *Proc. R. Soc. Lond. A.* 201 (1950), pp. 159 –174. DOI: 10.1098/rspa.1950.0049.

[235] G. J. Hutchens. "Approximate near-field blast theory: A generalized approach." *J. Appl. Phys.* 88.6 (2000), pp. 3654–3658. DOI: 10.1063/1.1288785.

[236] P. B. Corkum, N. H. Burnett, and F. Brunel. "Above-threshold ionization in the long-wavelength limit." *Phys. Rev. Lett.* 62 (1989), pp. 1259–1262. DOI: 10.1103/PhysRevLett.62.1259.

[237] J. Chappell et al., (AWAKE Collaboration), "Experimental study of extended timescale dynamics of a plasma wakefield driven by a self-modulated proton bunch." *Phys. Rev. Accel. Beams* 24 (2021), p. 011301. DOI: 10.1103/PhysRevAccelBeams.24.011301.

[238] A. A. Gorn et al., (AWAKE Collaboration), "Proton beam defocusing in AWAKE: comparison of simulations and measurements." *Plasma Phys. Control. Fusion* 62 (2020), p. 125023. DOI: 10.1088/1361-6587/abc298.

[239] N. Moschuering et al. "First fully kinetic three-dimensional simulation of the AWAKE baseline scenario." *Plasma Phys. Control. Fusion* 61 (2019), p. 104004. DOI: 10.1088/1361-6587/ab411e.

[240] A. A. Gorn et al. "Response of narrow cylindrical plasmas to dense charged particle beams." *Phys. Plasmas* 25 (2018), p. 063108. DOI: 10.1063/1.5039803.

[241] K. V. Lotov, A. P. Sosedkin, and A. V. Petrenko. "Long-Term Evolution of Broken Wakefields in Finite-Radius Plasmas." *Phys. Rev. Lett.* 112 (2014), p. 194801. DOI: 10.1103/PhysRevLett.112.194801.

[242] A. Pukhov et al. "Phase Velocity and Particle Injection in a Self-Modulated Proton-Driven Plasma Wakefield Accelerator." *Phys. Rev. Lett.* 107 (2011), p. 145003. DOI: 10.1103/PhysRevLett.107.145003.

[243] J. Vieira et al. "Ion Motion in Self-Modulated Plasma Wakefield Accelerators." *Phys. Rev. Lett.* 109 (2012), p. 145005. DOI: 10.1103/PhysRevLett.109.145005.

[244] J. Vieira et al. "Ion motion in the wake driven by long particle bunches in plasmas." *Phys. Plasmas* 21.5 (2014), p. 056705. DOI: 10.1063/1.4876620.

[245] V. A. Minakov, A. P. Sosedkin, and K. V. Lotov. "Accelerating field enhancement due to ion motion in plasma wakefield accelerators." *Plasma Phys. Control. Fusion* 61.11 (2019), p. 114003. DOI: 10.1088/1361-6587/ab41a7.

[246] K. V. Lotov et al. "Electron trapping and acceleration by the plasma wakefield of a self-modulating proton beam." *Phys. Plasmas* 21.12 (2014), p. 123116. DOI: 10.1063/1.4904365.

[247] J. B. Rosenzweig et al. "Effects of Ion Motion in Intense Beam-Driven Plasma Wakefield Accelerators." *Phys. Rev. Lett.* 95 (2005), p. 195002. DOI: 10.1103/PhysRevLett.95.195002.

[248] C. Benedetti et al. "Emittance preservation in plasma-based accelerators with ion motion." *Phys. Rev. Accel. Beams* 20 (2017), p. 111301. DOI: 10.1103/PhysRevAccelBeams.20.111301.

[249] R. Gholizadeh et al. "Preservation of Beam Emittance in the Presence of Ion Motion in Future High-Energy Plasma-Wakefield-Based Colliders." *Phys. Rev. Lett.* 104 (2010), p. 155001. DOI: 10.1103/PhysRevLett.104.155001.

[250] T. J. Mehrling et al. "Suppression of Beam Hosing in Plasma Accelerators with Ion Motion." *Phys. Rev. Lett.* 121 (2018), p. 264802. DOI: 10.1103/PhysRevLett.121.264802.

[251] 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." *J. Phys. B: At. Mol. Opt.* 47.23 (2014), p. 234003. DOI: 10.1088/0953-4075/47/23/234003.

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

[253] R. D. Skeel and M. Berzins. "A Method for the Spatial Discretization of Parabolic Equations in One Space Variable." *SIAM J. Sci. Stat. Comp.* 11.1 (1990), pp. 1–32. DOI: 10.1137/0911001.

[254] C. E. Clayton et al. "Self-mapping the longitudinal field structure of a nonlinear plasma accelerator cavity." *Nat. Commun.* 7.1 (2016), p. 12483. DOI: 10.1038/ncomms12483.

[255] K. L. Bell et al. "Recommended Data on the Electron Impact Ionization of Light Atoms and Ions." *J. Phys. Chem. Ref. Data* 12.4 (1983), pp. 891–916. DOI: 10.1063/1.555700.

[256] M. A. Lennon et al. "Recommended Data on the Electron Impact Ionization of Atoms and Ions: Fluorine to Nickel." *J. Phys. Chem. Ref. Data* 17.3 (1988), pp. 1285–1363. DOI: 10.1063/1.555809.

[257] X.-M. Zhu and Y.-K. Pu. "Optical emission spectroscopy in low-temperature plasmas containing argon and nitrogen: determination of the electron temperature and density by the line-ratio method." *J. Phys. D: Appl. Phys.* 43.40 (2010), p. 403001. DOI: 10.1088/0022-3727/43/40/403001.