



Measurement and optimisation of beam quality from laser wakefield accelerators

Michael Backhouse

December 2022

*Submitted in partial fulfilment of the requirements for the degree of Doctor of Philosophy of
Imperial College London*

Department of Physics
Imperial College London
Prince Consort Road
London
SW7 2BZ

The copyright of this thesis rests with the author. Unless otherwise indicated, its contents are licensed under a Creative Commons Attribution-Non Commercial-No Derivatives 4.0 International Licence (CC BY-NC-ND). Under this licence, you may copy and redistribute the material in any medium or format on the condition that; you credit the author, do not use it for commercial purposes and do not distribute modified versions of the work. When reusing or sharing this work, ensure you make the licence terms clear to others by naming the licence and linking to the licence text. Please seek permission from the copyright holder for uses of this work that are not included in this licence or permitted under UK Copyright Law.

All work presented in this thesis is my own unless otherwise stated and acknowledged.

Michael Backhouse

December 2022

Abstract

This thesis concerns experimental work in the field of laser wakefield acceleration, with a focus on the diagnosis and optimisation of the electron beam quality.

The density length parameter space of a 5 TW, 0.25 J laser driven wakefield accelerator was characterised. Measurements of the electron beams and x-ray pulses were reported, and optimal parameters for various metrics were found. Beam-driven acceleration was identified as the mechanism that produced energies above 211 MeV, and the peak x-ray brilliance was $4.2 \pm 0.8 \times 10^{20} \text{ ph s}^{-1} \text{ mm}^{-2} \text{ mrad}^{-2} 0.1\% \text{BW}^{-1}$. Both the electron energy and the x-ray brilliance are significantly higher than literature values using comparable laser powers. Separately, the parameter scans were used to measure an extended dephasing length of the laser-accelerated beam, attributable to semi-localised depletion of the driving laser pulse, and measure the pulse evolution rate and injection length as a function of plasma density, which was found to be slower than would be expected when only considering the longitudinal evolution.

An emittance diagnostic was developed using a beam mask and electron spectrometer. This was used to measure the spectrally resolved normalised emittance of GeV beams, produced by ionisation injection in a gas jet using a 165 TW, 7.4 J laser. Average emittance values as low as $4 \mu\text{m}$ were measured, which are the lowest emittances recorded using a beam mask technique in the literature, at energies that are close to an order of magnitude higher than other beam mask methods.

The effect of density ramps and plasma mirrors on electron beam divergence was measured in the context of staged wakefield acceleration, using a 242 TW, 11 J laser. Termination of an acceleration stage with a plasma mirror was found to increase total beam divergence from $3.38 \pm 0.07 \text{ mrad}$ to $6.13 \pm 0.13 \text{ mrad}$, and the effect was observed to persist at high energies, up to 2.2 GeV. Using simulations and numerical models, the presence of the density ramp was shown to have a divergence-reducing effect with a magnitude that matched the experiment. The 10^3 tesla magnetic fields generated in plasma mirrors were investigated using simulations, and the effect of these fields on the electron beam was quantified. Compared to normal incidence, a 45 degree angle of the plasma mirror to the beam axis was found to reduce the integrated magnetic fields inside the mirror, with beneficial effects on electron beam emittance.

Contents

Abstract	2
Role of the author	10
Acknowledgements	11
1 Introduction	12
1.1 Conventional particle accelerators	13
1.2 Plasma accelerators	15
1.3 Near-term applications	18
1.4 Towards frontier plasma colliders	19
1.5 Thesis outline	23
2 Theory	24
2.1 Laser pulses	24
2.2 Single particle motion	27
2.3 Ionisation	30
2.4 Plasma waves	32
2.5 Laser pulse evolution	40
2.6 Acceleration	46
2.7 Injection	49
2.8 Transverse dynamics	54
3 Methods	62
3.1 Laser diagnostics	62
3.2 Gas targets	66
3.3 Interferometry	68
3.4 Electron beam diagnostics	75
3.5 Particle-in-cell simulations	82
4 The density-length parameter space of a laser wakefield accelerator	93
4.1 Experimental setup	94
4.2 Density-length scans	104
4.3 Simulations	112

4.4	Discussion	118
4.5	Summary	132
5	Measurement of ionisation injected GeV beam emittance	133
5.1	Experimental setup	134
5.2	Emittance measurement	139
5.3	Experimental results	144
5.4	Simulations	149
5.5	Discussion	152
5.6	Summary	159
6	Beam divergence considerations for staged LWFA	160
6.1	Electron beam measurements	161
6.2	Laser-tape simulations	170
6.3	LWFA simulations	178
6.4	Discussion	184
6.5	Summary	191
7	Conclusions and outlook	192
7.1	Conclusions	192
7.2	Summary and outlook	195
A	Appendix	197
A.1	Fundamental constants	197
A.2	Maxwell's equations	197
A.3	Relativistic single particle motion	198
A.4	Tunnel ionisation rate for specific electrons	199
A.5	Plasma dispersion relation	200
A.6	The quasistatic approximation	200
A.7	Hamiltonian of electron in a 1D plasma wave	201
	Bibliography	226

List of Figures

1.1	Accelerators: then and now	13
1.2	Wakefield simulation snapshot	16
1.3	LWFA electron energies	17
1.4	Plasma collider schematic	22
2.1	Focusing electric field	26
2.2	Zernike polynomials	27
2.3	Single particle trajectories in an EM wave	29
2.4	Ionisation fraction during pulse transit	32
2.5	Linear wake excitation	35
2.6	Non-linear wake excitation	36
2.7	2D particle trajectories	38
2.8	Bubble regime simulation	39
2.9	Relativistic self-focusing	41
2.10	Envelope equation solutions	43
2.11	Plasma focal length	44
2.12	Longitudinal pulse compression	45
2.13	Dephasing length and energy versus plasma density	49
2.14	Phase space trajectories in a plasma wave	50
2.15	Self-injection threshold	52
2.16	Ionisation injection threshold	54
2.17	Emittance phase space plot	58
2.18	Beta function illustration	59
2.19	The angular synchrotron spectrum	61
3.1	Focal spot measurement	64
3.2	Transverse metrics	65
3.3	Gas cell aperture simulations	67
3.4	Parameter fitting fluid simulations	68
3.5	Principle of interferometry	69
3.6	Interferometer diagrams	71
3.7	Fourier filtering of interferograms	72
3.8	Density map retrieval	74

3.9	Stability of interferometry	75
3.10	Electron spectrometer diagram	78
3.11	Electron tracking for spectrometer calibration	79
3.12	Image of electron spectrometer	80
3.13	Beam mask diagram	81
3.14	MC simulation of beam mask	82
3.15	2D numerical dispersion relation	84
3.16	Axial numerical dispersion relation	85
3.17	Azimuthal mode projection of focal spot image	88
3.18	Real versus perfect focal spot waterfalls	88
3.19	Real versus perfect focal spot beam metrics	89
3.20	Longitudinal resolution scan: fields	90
3.21	Longitudinal resolution scan: beam metrics	91
3.22	Azimuthal mode convergence scan	92
4.1	ATA2 experimental layout	94
4.2	Focal spot characterisation	96
4.3	Gas cell interferometry	98
4.5	Electron spectrometer image metrics	99
4.4	Electron spectrometer calibration	99
4.6	Helium gas optimisation	102
4.7	Mixed gas optimisation	104
4.8	Laser energy during data collection	105
4.9	Density length parameter space: charge	105
4.10	Density length parameter space: 95% energy	106
4.11	Density length parameter space: peak energy	107
4.12	Density length parameter space: divergence	108
4.13	Example data from helium peak charge setting	109
4.14	Example data from lowest divergence setting in 1%N ₂	109
4.15	Example data from peak energy setting in 1%N ₂	110
4.16	Comparison of the peak x-ray intensity.	111
4.17	Comparison of the critical energy.	111
4.18	Peak number of photons	112
4.19	SI injection snapshots	113
4.20	PIC spectrum histories for pure helium gas	114
4.21	PIC spectrum histories for doped gas	115
4.22	PIC a_0 evolution	116
4.23	Beam metrics from PIC simulations	117
4.24	Pure helium gas peak energy scaling comparison	119
4.25	Doped gas peak energy scaling comparison	120
4.26	Energy gain versus length in doped gas	121
4.27	PIC charge density maps	122
4.28	Energy versus cell length in doped gas	123

4.29	Cell length calibration	124
4.30	Injection length versus density: I	124
4.31	Comparison of a_0 values at injection	126
4.32	Injection length versus density: II	127
4.33	Estimated x-ray brilliance	128
4.34	PIC beam size for x-rays	129
4.35	Brilliance values using a constant 2 μm source size	130
4.36	Peak electron energies: literature comparison	131
4.37	Peak brilliance values: literature comparison	131
5.1	Diagram of the experimental setup	134
5.2	Photograph and interferogram of a plasma	135
5.3	Plasma density maps	136
5.4	Resolution testing	137
5.5	Images of the emittance mask	138
5.6	Electron spectrometer images	139
5.7	Fitting Gaussian peaks to a signal	140
5.8	Calculation of the mask clear aperture	141
5.9	Simulated emittance measurement versus divergence	142
5.10	Signal-to-noise error analysis	143
5.11	MC simulation-based measurement method	144
5.12	Low-density example emittance measurement	145
5.13	High-density example emittance measurement	146
5.14	Summary of all emittance measurements	147
5.15	Charge weighted median emittance values	148
5.16	Simulated low-density emittance measurement	150
5.17	Simulated high-density emittance measurement	151
5.18	Charge weighted mean simulated emittances	152
5.19	Energy and emittance history plots	153
5.20	Ionisation injection mechanism	154
5.21	Injection into a forming wake	155
5.22	Literature comparison	157
5.23	MC simulated electron spectrometer images	158
6.1	Sketch of the vacuum chamber layout	161
6.2	Diagram of the target	162
6.3	Shadowgrams of the target	162
6.4	Three cell configurations	163
6.5	Focal spot metrics	164
6.6	Electron profile processing	165
6.7	Full beam divergence measurements	166
6.9	Spectrometer particle tracking	168
6.8	Beam centroids	168

6.10 Example electron spectrometer divergence measurements	169
6.11 Divergence measurements	170
6.12 Charge density of laser-tape simulation	171
6.13 Laser temporal profile for the laser-tape simulations	172
6.14 Plasma electron density and magnetic fields	173
6.15 Currents generated in the tape by the laser	175
6.16 Integrated magnetic fields	177
6.17 Electron beam transverse momentum FWHM change	178
6.18 PIC simulated electron spectrum history	179
6.19 PIC simulated density ramp effect on beam	180
6.20 Example PIC electron trajectories	181
6.21 Synthetic spectrographs from PIC simulations	182
6.22 Emittance evolution in PIC simulations	183
6.23 MC simulation of the effect of the second tape	185
6.24 Comparison of results to simulations	186
6.25 Simulated average charge fraction captured	187
6.26 Beam coupling beta function: exponentials	189
6.27 Beam coupling beta function: optimal	191

List of Tables

1.1	Estimated collider parameters	20
2.1	Ionisation thresholds	30
3.1	Scintillator screen properties	76
4.1	Filter pack materials	101
4.2	Simulation parameters	113
4.3	Optimal measurements and settings	132
6.1	Experimental divergence and pointing measurements	167
A.1	Fundamental constants used in this thesis	197

Role of the author

The experiments in this thesis were performed during three experimental campaigns using the Gemini laser at Rutherford Appleton Laboratories, UK. During the first of these experiments, reported in chapter 6, the author worked as part of a large experimental team and was responsible for designing, assembling, and operating the transverse probe diagnostic for both cells, and for live analysis of the plasma density. The second experiment, reported in chapter 4, was set up prior to the author's arrival on the experiment by Matthew Streeter, Robert Shalloo, and Jan-Niclas Gruse, and the automation of the experiment was largely set up by Steven Dann. Their work represents a significant contribution to this thesis. The author assisted on this experiment during the final weeks, performing the parameter scans presented in this thesis. During the experiment reported in chapter 5, the author was the main Target Area Operator for 6 of the 8 weeks of the campaign, performing the daily operation of the facility and assuming responsibility for safety. This experiment was set up by the author and Matthew Streeter. All experimental data analysis was performed by the author, except where otherwise credited. All simulations were performed by the author using the CX1 and CX2 computing clusters at Imperial College London.

Acknowledgements

Looking back, I guess the writing was always on the wall that I would end up doing a PhD, but from my perspective, it has felt somewhat serendipitous. Five years ago, I had just started a job as an ‘Optical Scientist’, a rather grandiose title for a graduate with almost no laboratory experience, and it was then that a burning desire to return to Imperial for years of further study was kindled. Doing so has been one of the best decisions I have ever made. The last four years have been a challenging and exciting time, and I feel extremely lucky to have had this opportunity. To that end, I would first like to thank Zulfikar, both for offering me this opportunity, and for the countless hours reading my work and listening to me present my results.

Some of the most trying times on the PhD were undoubtedly during the experiments, and so a thank you goes to everyone who worked/put up with me throughout these. Particular thanks go to Savio, as the many winter weeks spent isolated at RAL might have been intolerable without someone to laugh about it all with. In the office, thank you, Stuart, for the causal chats we would have that would often lead to interesting avenues of research. Thanks go to Ollie, for countless coffee beaks and the willingness to compete over literally anything, as well as to Nuo, for always providing an interesting perspective, as well as keeping the office clean and tidy. I want to say thank you to Jon, Matt, and Rob, for the guidance you gave me, as well as the examples you set of simultaneously doing good research and having a good time. Finally, thanks to all members of the 2018 kettle cohort; the enjoyable plasma schools and hot-pot nights that accompanied the first years of the PhD will always be remembered.

To my mother and father, Sara and John, and brothers, Nick and Chris, thank you for all the support over the years. There is no chance that I would be where I am today without you all. To my wife, Cristina, thank you for everything. For the encouragement to do the PhD in the first place, the help while doing it, and for almost singlehandedly organising a wedding and honeymoon to what can only be described as military precision. Your conscientiousness and kindness will always be an inspiration.

1 Introduction

THE PARTICLE ACCELERATOR is one of mankind’s primary tools for investigating the structure of the universe. This has been the case since the days of Rutherford¹, who, by simply scattering particles off a gold foil, showed that the true structure of atoms was far from what was expected at the time². This highlighted the use of particle accelerators as microscopes for inspecting the smallest of structures, and advances in fundamental physics have been in lock-step with advances in accelerator technology ever since³. One century after Rutherford’s foil experiment, the Large Hadron Collider (LHC) was used to confirm the existence of the Higgs boson⁴, the final part of the Standard Model⁵. Now the task remains to go beyond this model and explain many pressing phenomena⁶, such as the accelerating expansion of the universe⁷, the mysterious rotational properties of galaxies⁸, and to incorporate gravity into a unified ‘theory of everything’⁹. If these explanations are to be found, it is possible they will be found using particle accelerators.

Beyond pursuing these high physics goals, particle accelerators touch our lives in more direct ways. Proton therapy, a safer method for treating cancerous tumours than gamma ray therapy, is becoming increasingly widespread, owing to the growing accessibility of particle accelerators¹⁰. Further, the fast particles from accelerators are used to create the isotopes needed for Positron-Emission-Tomography¹¹, a technique for monitoring critical bodily functions. Industrial applications are prolific; from shrink-wrapping and food sterilisation¹², to semiconductor production¹³ and material hardening¹⁴. Much of the modern world would not be possible without particle accelerators, and their proliferation is driving growth in many areas.

¹ [1] E. Rutherford (1911) *Lond. Edinb. Dublin philos. mag. j. sci.*

² [2] J. J. Thomson (1904) *Lond. Edinb. Dublin philos. mag. j. sci.*

³ [3] E. F. Hausecker *et al.* (2011) *Physics in Perspective*

⁴ [4] A. Collaboration *et al.* (2015) *Physical Review Letters*

⁵ [5] M. K. Gaillard *et al.* (1999) *Reviews of Modern Physics*

⁶ [6] A. Joyce *et al.* (2015) *Physics Reports*

⁷ [7] A. G. Riess *et al.* (1998) *The Astronomical Journal*

⁸ [8] V. Trimble (1987) *Annual Review of Astronomy and Astrophysics*

⁹ [9] R. B. Laughlin *et al.* (2000) *Proceedings of the National Academy of Sciences*

¹⁰ [10] Y. Matsumoto *et al.* (2021) *Journal of Personalized Medicine*

¹¹ [11] O. Jacobson *et al.* (2015) *Bioconjugate Chemistry*

¹² [12] A Mousavi Khaneghah *et al.* (2020) *Food and Chemical Toxicology*

¹³ [13] J. S. Williams (1998) *Materials Science and Engineering: A*

¹⁴ [14] D. J. Rej *et al.* (1997) *Journal of Vacuum Science Technology A*

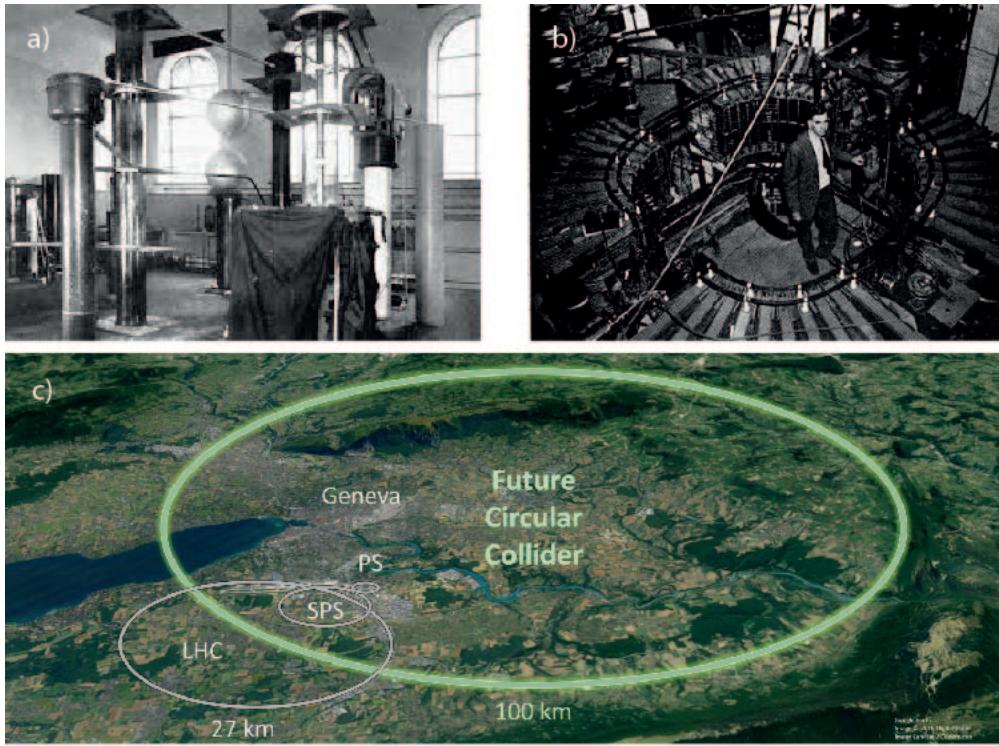


Figure 1.1: Then and now: panel a) shows the original Cockcroft-Walton generator, and panel b) shows the first synchrotron. Panel c) shows the scale of modern-day accelerators.

The LHC, the largest particle accelerator in the world, has a circumference of 27 km, and cost over £3 billion to build. Smaller accelerators that are used for applications are still on the metre scale and can be prohibitively expensive. To reduce the size and cost of these machines, plasma-based accelerators have been proposed. Plasma accelerators can increase the accelerating gradient by over a factor of 1000 when compared to conventional technology¹⁵, which can potentially reduce the size of these machines by the same factor, a promise that has produced a flourishing research community over the last 40 years. This thesis is concerned with one particular class of plasma accelerators, the laser wakefield accelerator (LWFA), which uses ultra-intense laser pulses with near atmospheric density gas to produce relativistic electrons.

1.1 Conventional particle accelerators

The first particle accelerators were the Cockcroft-Walton generators¹⁶, which successfully split the atom for the first time, earning the namesakes the 1951 Nobel Prize in Physics. These are essentially capacitors, where particles are injected using electrical breakdown, and are still used today in electrical appliances that require voltage multiplication from a low AC input¹⁷. An image of the first accelerator is shown in Fig. 1.1, which was capable of producing an accelerating gradient of 440 kV m^{-1} . The main disadvantage of these

¹⁵ [15] E. Esarey *et al.* (2009) *Reviews of Modern Physics*

¹⁶ [16] J. D. Cockcroft *et al.* (1932) *Proc. R. Soc. A: Math. Phys. Eng. Sci.*

¹⁷ [17] Y. Jongen “Review of compact commercial accelerator products and applications” (1997)

accelerators was that the field was only applied to the particle once, such that once the particle reaches the cathode, acceleration terminates. This limited the attainable energy to the product of the acceleration distance and the field strength, $\Delta W = qEl_{\text{acc}}$, where q is the charge of the particle, E is the magnitude of the accelerating electric field, and l_{acc} is the acceleration length. In many modern accelerators, particles travel along a circular path, allowing the same electric field to be applied to the particle repeatedly. This circular path is achieved using dipole magnets, which apply a force to the beam perpendicularly to its motion. This circular geometry increases the energy limit substantially but imposes numerous technical challenges that the field of accelerator physics arose to solve. The problem of keeping the beam focused is addressed by using multi-pole magnets, which can be shown to maintain a small beam size over arbitrary distances¹⁸, while the problem of the phase synchronisation was addressed by gradually ramping up the dipole magnet strengths as the particle accelerates, which turned the cyclotron into the synchrotron¹⁹. The LHC is a synchrotron, which utilises these solutions, among many others, to achieve proton energies of 7 TeV²⁰. The current limit on proton energy is set by the strength of the magnets, such that the only route to increased energy is to increase the accelerator radius or increase the field strength of the magnets.

Electron synchrotrons face a different, more fundamental, limitation. It is a fact of nature that charged particles radiate energy as they accelerate, and this radiation occurs when particles follow curved trajectories. The problem for electrons is that the power of this radiation, $P \propto W^4/m^4R^2$, where W is the particle energy, m is the particle mass, and R is the instantaneous radius of curvature. Since electrons are approximately 2000 times less massive than protons, they radiate energy at 10^{13} times the rate! This hard scaling has limited the maximum attainable electron energy to 100 GeV, which occurred at the Large Electron Positron collider (LEP²¹), built in the same tunnel as the LHC. It is unfortunate that electron-positron energies are limited in this way, as their collisions are much easier to analyse than hadron collisions, owing to the fact that hadrons are composite particles while leptons are fundamental. For this reason, there is great motivation for increasing the maximum electron energy to the multi-TeV level.

Achieving multi-TeV electron energies using a synchrotron would require a planet-sized collider²², and due to the impracticality of constructing such a device, other solutions need to be investigated. A linear accelerating geometry overcomes the radiation problem, but the original limit encountered by the Cockcroft-Walton generators resurfaces. The accelerating field has been increased by using superconducting radio-frequency cavities, which use a moving electromagnetic wave to accelerate the particle. This technology can achieve field strengths of 100 MV m^{-1} , although in practice this is limited to the $20 - 30 \text{ MV m}^{-1}$ range²³. The limitation on the maximum field is caused by electrical

¹⁸ [18] E. D. Courant *et al.* (1952) *Physical Review*

¹⁹ [19] E. M. McMillan (1945) *Physical Review*

²⁰ [20] O. Brüning *et al.* (2012) *Progress in Particle and Nuclear Physics*

²¹ [21] K. Hübner (2004) *Physics Reports*

²² $\sqrt{R_{\text{Earth}}/R_{\text{LHC}}} \sim 40$; an Earth sized LEP could achieve $\sim 4 \text{ TeV}$ using the same technology.

²³ [22] H Padamsee (2017) *Superconductor Science and Technology*

breakdown, as electrons are pulled from the cavity walls by the field, causing irreparable damage to the cavities. Since the field is limited, the desired electron energy sets the accelerator length. The proposed International Linear Collider²⁴ promises to achieve a centre-of-mass energy of 1 TeV using two anti-colinear accelerators, which are each over 10 km in length. The 2021 ILC Report to Snowmass²⁵ states that the cost of the first stage of the project, which will be able to produce electron energies up to 125 GeV, will be \$6.2 – 6.8 billion²⁶, while the upgrade to TeV energies is estimated to cost in the range of \$18 – 30 billion²⁷. To reduce this size and cost, the field strength needs to be increased, and, since the advances in conventional radiofrequency technology are obtaining diminishing returns in this area, there is significant motivation to explore other options.

1.2 Plasma accelerators

If electrical breakdown is limiting field strengths, then why not accelerate particles inside an already broken-down material? This is the idea behind plasma accelerators. A plasma is defined as a collection of positively and negatively charged particles that exhibits collective behaviour and interact electromagnetically. The number of positive and negatively charged particles is approximately equal, such that the plasma is macroscopically neutral. However, over short distances and timescales, large charge separations can be generated, which can support extremely strong electric fields. At a near atmospheric density of $1 \times 10^{24} \text{ m}^{-3}$, a plasma can support electric fields on the 100 GV m^{-1} scale, approximately 3-4 orders of magnitude higher than conventional technology. This holds the promise of shrinking the accelerating sections of kilometre-long facilities to table-top sizes.

The idea of using plasmas to accelerate particles originated in the 1950s, when the potential utility of large electric fields in plasmas was first noted²⁸. However, it was not until 1979 that a seminal paper by Tajima and Dawson²⁹ proposed the use of a laser pulse and a low-density plasma to produce a relativistically moving plasma wave, the electric field of which could accelerate particles to high energies. In order to resonantly excite this plasma wave, the laser pulse length would need to be on the timescale of the plasma frequency, $\omega_p \propto \sqrt{n_e/m}$, which is on the scale of 10s of femtoseconds, in a regime known as Laser Wakefield Acceleration. On this timescale, the ions are approximately static due to their larger mass, so it is the electron motion that dominates. The structure of a large amplitude plasma wave in the wake of an intense laser pulse is illustrated in Fig. 1.2.

²⁴ [23] P. Bambade *et al.* (2019) *arXiv e-prints*

²⁵ [24] A. Aryshev *et al.* (2022)

²⁶ The estimate in the report is explicitly in 2012 dollars; the value quoted here accounts for inflation at the time of writing.

²⁷ [25] M. Chamizo Llatas *et al.* (2022) *arXiv e-prints*

²⁸ [26] V. A. Veksler (1956) *The Soviet Journal of Atomic Energy*

²⁹ [27] T. Tajima *et al.* (1979) *Physical Review Letters*

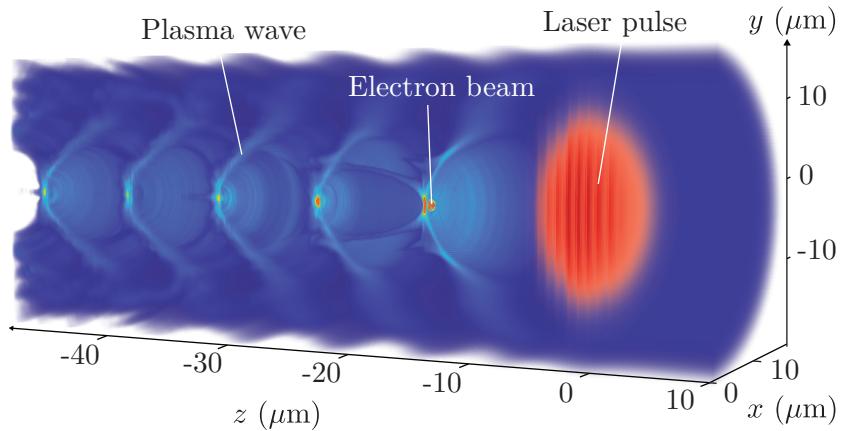


Figure 1.2: An image from a laser wakefield simulation, sliced through the plane $x = 0$ to reveal the structure of the wake. The laser pulse, in red, travels from left to right, driving a wave in the plasma. Electrons from the background plasma have been injected into the wave, forming a beam in the trough behind the laser pulse.

Since single pulses with high intensities were not available in 1979, the Plasma Beat-Wave Accelerator (PBWA) was envisioned. This utilised two laser pulses with slightly different frequencies to produce a train of laser ‘beats’, which could be tuned to the ω_p timescale. While initially successful³⁰, this technique was always destined to encounter problems at large wave amplitudes. The reduced plasma frequency that accompanies the large amplitude waves causes a de-tuning of the interaction between the beats and the plasma, saturating the field strength. More practically, coupling between the long laser pulse and the motion of the ions was observed to limit the attainable field strength to 0.6 GV m^{-1} ³¹. The Self-Modulated Laser Wakefield Accelerator overcame these effects by using the plasma wave itself to modulate the laser pulse to the plasma frequency³². Through various non-linear pulse evolution mechanisms, this pulse would then amplify and inject background electrons into the wake³³, producing electron beams with energies up to 100 MeV ³⁴. However, these electron beams had thermal spectra, making them ill-suited for use in particle accelerators.

The attention of the scientific community was captured in 2004 when it was shown by three groups³⁵ that electron beams with narrow energy spread could be obtained using 100 GeV m^{-1} accelerating gradients. This advance was made possible by the transition to the LWFA regime, facilitated by continued advances in laser technology. The most

³⁰ [28] C. E. Clayton *et al.* (1985) *Physical Review Letters*

[29] C. E. Clayton *et al.* (1993) *Physical Review Letters*

³¹ [30] F. Amiranoff *et al.* (1995) *Physical Review Letters*

³² [31] J. Krall *et al.* (1993) *Physical Review E*

³³ [32] K. Nakajima *et al.* (1995) *Physical Review Letters*

[33] A. Modena *et al.* (1995) *Nature*

[34] C. I. Moore *et al.* (1997) *Physical Review Letters*

[35] W. P. Leemans *et al.* (2001) *Physics of Plasmas*

[36] V. Malka *et al.* (2001) *Physics of Plasmas*

³⁴ [37] D. Gordon *et al.* (1998) *Physical Review Letters*

³⁵ [38] S. P. D. Mangles *et al.* (2004) *Nature*

[39] C. G. R. Geddes *et al.* (2004) *Nature*

[40] J. Faure *et al.* (2004) *Nature*

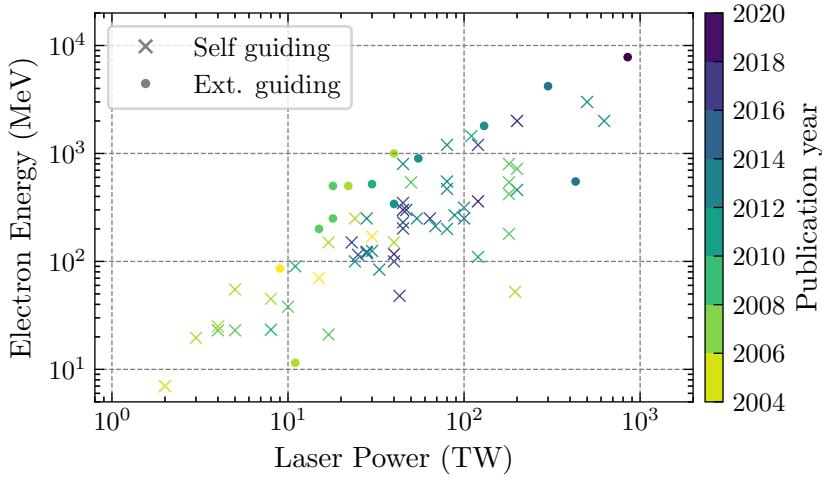


Figure 1.3: Summary of recent LWFA energy measurements as a function of the laser power and publication date.

significant of these advances was the development of Chirped-Pulse Amplification³⁶, a technique which drastically increased the laser power that could be attained before laser amplifiers became damaged³⁷. Since 2004, the LWFA regime has dominated the energy records, which have advanced continuously. Higher energies have been obtained by using increasingly powerful lasers, as shown in Fig. 1.3. This allows large amplitude plasma waves to be driven at lower densities and for longer, which offers favourable scaling for the maximum energy³⁸. Multi-GeV energies require acceleration lengths over centimetre scales, and so methods such as relativistic self-focusing³⁹, electrically ionised capillary waveguides⁴⁰, and laser ionised plasma channels⁴¹, have been used to combat diffraction. Over the last decade, the maximum energies in the literature have been obtained by using capillary waveguides, with the published record standing at 7.8 GeV⁴². Capillaries allow the laser to remain focused at lower densities than self-guiding, and are currently more stable than laser-generated waveguides. These beams are typically limited by a process known as dephasing, which occurs when the electron beam out-runs the slightly slower plasma wave.

³⁶ [41] D. Strickland *et al.* (1985) *Optics Communications*

[42] P. Maine *et al.* (1988) *IEEE Journal of Quantum Electronics*

³⁷ [43] M. D. Perry *et al.* (1994) *Science*

³⁸ [44] W. Lu *et al.* (2007) *Physical Review Special Topics - Accelerators and Beams*

³⁹ [45] A. B. Borisov *et al.* (1992) *Physical Review Letters*

[46] L. M. Chen *et al.* (2007) *Physics of Plasmas*

⁴⁰ [47] A. Zigler *et al.* (1996) *J. Opt. Soc. Am. B*

[48] Y. Ehrlich *et al.* (1998) *J. Opt. Soc. Am. B*

⁴¹ [49] R. Durfee C. G. *et al.* (1993) *Phys Rev Lett*

[50] H. M. Milchberg *et al.* (1996) *Physics of Plasmas*

[51] R. J. Shalloo *et al.* (2018) *Physical Review E*

[52] A. Picksley *et al.* (2020) *Physical Review E*

[53] B. Miao *et al.* (2020) *Physical Review Letters*

[54] B. Miao *et al.* (2022) *Physical Review X*

⁴² [55] A. Gonsalves *et al.* (2019) *Physical Review Letters*

A particle beam can also be used to drive the plasma wave. This can be done with either an electron beam⁴³, such as at SLAC and FLASHForward, or with a proton beam⁴⁴, as is being done at the AWAKE project⁴⁵, using CERN’s SPS to provide the beam. Beam-driven acceleration does not suffer from dephasing, which is one of the main limitations of LWFA. In addition, it benefits from both the impressive stability of the drivers provided by conventional accelerators and the relatively simple interaction between the beam and the plasma, when compared to the laser-plasma interaction. They are also less computationally intensive to simulate because the laser frequency does not need to be resolved. The main downside of these techniques is the requirement for a large conventional accelerator to be placed nearby. Consequently, beam-driven acceleration is well suited for use as an ‘afterburner’⁴⁶, providing a final energy boost to conventionally accelerated beams.

1.3 Near-term applications

The primary near-term application of laser wakefield accelerators is as a source of x-rays, or betatron radiation, which is often produced naturally by LWFA electron beams. The term derives from the oscillation frequency of electrons in a focusing channel, which is known as the betatron frequency, $\omega_\beta = \omega_p / \sqrt{2\gamma}$. First observed at conventional facilities⁴⁷, it was then soon observed in LWFA⁴⁸. The spectrum of the radiation is broad⁴⁹, and similar to that emitted by a synchrotron, with significant frequency components up to a critical energy in the range of 1 - 30 keV, for electrons in the 0.1 - 2 GeV range⁵⁰. The short electron bunch duration⁵¹ means that these beams produce ultra-short pulse x-ray flashes⁵², suitable for time-resolved imaging of processes that proceed on the femtosecond timescale⁵³. The small source size⁵⁴ enables phase contrast imaging⁵⁵, capable of resolving the fine structures inside biological samples⁵⁶, and the broad energy range allows for spectroscopic probing of warm dense matter⁵⁷.

Laser wakefield accelerated beams at the few-GeV level can already be used for frontier physics experiments. The measurement of quantum electrodynamics in strong fields has been performed⁵⁸, experiments in which the electron beam collides with an

⁴³ [56] I. Blumenfeld *et al.* (2007) *Nature*

⁴⁴ [57] R. Assmann *et al.* (2014) *Plasma Physics and Controlled Fusion*

⁴⁵ [58] E. Gschwendtner *et al.* (2016) *Nucl. Instrum. Methods Phys. Res. A: Accel. Spectrom. Detect. Assoc. Equip.*

⁴⁶ [59] S. Lee *et al.* (2002) *Physical Review Special Topics - Accelerators and Beams*

⁴⁷ [60] S. Wang *et al.* (2002) *Physical Review Letters*

⁴⁸ [61] A. Rousse *et al.* (2004) *Physical Review Letters*

⁴⁹ [62] S. Fourmaux *et al.* (2011) *New Journal of Physics*

⁵⁰ [63] M. Schnell *et al.* (2015) *Journal of Plasma Physics*

⁵¹ [64] M. Heigoldt *et al.* (2015) *Physical Review Special Topics - Accelerators and Beams*

⁵² [65] K. T. Phuoc *et al.* (2007) *Physics of Plasmas*

⁵³ [66] J. C. Wood *et al.* (2018) *Scientific Reports*

⁵⁴ [67] K. T. Phuoc *et al.* (2006) *Physical Review Letters*

⁵⁵ [68] F. Albert *et al.* (2008) *Physical Review E*

⁵⁶ [69] S. Fourmaux *et al.* (2011) *Optics Letters*

⁵⁷ [70] J. Cole *et al.* (2018) *Physical Review X*

⁵⁸ [71] M. Z. Mo *et al.* (2013) *Review of Scientific Instruments*

⁵⁹ [72] B. Kettle *et al.* (2019) *Physical Review Letters*

intense laser pulse. Looking slightly further ahead⁵⁹, electron beams with energies in the 10 – 100 GeV range could be used in fixed target experiments to search for dark photons, while high energy electron-proton collisions could be used to probe sub-hadronic structures⁶⁰. These applications benefit from relatively relaxed requirements on the transverse beam quality and hence provide ample near-term motivation to continue to develop plasma acceleration technology.

1.4 Towards frontier plasma colliders

Despite the significant progress that has been made in the field, numerous challenges need to be overcome if a plasma-based collider that pushes the frontier of physics forwards is to be realised. Various ‘roadmaps’ have been laid out by committees around the world with the aim of identifying a path from the current state of the art to a mature machine design⁶¹. There are technical challenges, such as achieving the high repetition rates and stability levels that are required, but also more fundamental physics roadblocks, such as injecting and maintaining high transverse beam quality, reaching TeV energies, and mirroring these achievements for positrons. We will discuss each of these physics issues in turn.

1.4.1 Emittance requirements

The luminosity of a beam, \mathcal{L} , is the ratio between the event rate and the interaction cross section⁶². It is given by

$$\mathcal{L} = \frac{fN^2}{4\pi\sigma_r^2} \propto \frac{fN^2}{\beta\varepsilon}, \quad (1.1)$$

where f is the repetition rate, N is the number of particles in each of the colliding bunches, σ_r is the Gaussian radius of the beam, β is a measure of focusing strength, and ε is the transverse emittance. Here we have assumed the normal incidence collision of two cylindrically symmetric beams with identical numbers of particles. This value is maximised by focusing the beam to a small diameter, and since different electron energy slices focus to different longitudinal positions, the desire for a small energy spread is implicit.

The LHC can produce luminosities in the $10^{38} \text{ m}^{-2} \text{ s}^{-1}$ range⁶³, largely through both the MHz frequencies that can be achieved by colliding bunch trains, and the ability to re-use the beam⁶⁴. Such high repetition rates are unattainable for linear colliders, since

⁵⁸ [70] J. Cole *et al.* (2018) *Physical Review X*

⁵⁹ [73] N. Mounet *European Strategy for Particle Physics - Accelerator RD Roadmap* (2022)

⁶⁰ [74] M. Wing (2019) *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*

⁶¹ [75] P. C. Bhat *et al.* (2019) *Reviews of Accelerator Science and Technology*

[25] M. Chamizo Llatas *et al.* (2022) *arXiv e-prints*

[73] N. Mounet *European Strategy for Particle Physics - Accelerator RD Roadmap* (2022)

[76] E. Adli (2022) *Journal of Instrumentation*

⁶² [77] W. Herr *et al.* (2006) *Cern Document Server*

⁶³ [78] F. Gianotti *et al.* (2005) *The European Physical Journal C - Particles and Fields*

⁶⁴ [79] *Luminosity? Why don't we just say collision rate?*, Quantum Diaries

Table 1.1: Estimated parameter requirements for an electron-positron LWFA collider.

Parameter	Value	Unit
Beam energy	0.5	TeV
Charge	0.64	nC
Normalised emittance	100	nm
Repetition rate	15	kHz
Energy spread	0.35	%
\mathcal{L}	1	$10^{38} \text{ m}^{-2} \text{ s}^{-1}$

each bunch can only be used once, and so frequencies are likely to be restricted to the realm of kHz. The deficiency in repetition rate needs to be accounted for, through both increasing the charge to nanocoulomb levels and reducing the emittance to approximately 100 nm. The projected requirements for plasma-based colliders are shown in Tab. 1.1, taken from a conceptual design by Schroeder *et al.*⁶⁵, and a sketch of this collider is shown in Fig. 1.4.

As will be discussed in Sec. 2.8, the emittance is a quantity that generally only increases, and it is, therefore, necessary to first inject low emittance beams, and then preserve this emittance from injection to the target. Fortunately, plasma accelerators are naturally suited to producing low emittance beams⁶⁶, owing to a coupling of an intrinsically small scale and high accelerating gradient⁶⁷. Emittance values as low as 56 nm have been measured in LWFA⁶⁸, exceeding the requirement given in Tab. 1.1. A review of LWFA emittance measurements is given in Sec. 5.5.2. Achieving such low emittances at the required charge and repetition rate will require further studies into the various injection mechanisms available. In addition, measurement of the ultra-low emittance from plasma-based experiments is challenging and will require the development of the diagnostic methods that are currently available⁶⁹.

1.4.2 Staged acceleration

The highest energies achieved by LWFA use densities of around $1 \times 10^{23} \text{ m}^{-3}$, requiring centimetre acceleration lengths and petawatt class lasers⁷⁰. Higher single-stage energies necessitate driving the plasma wave over longer distances⁷¹, requiring more laser energy, which is antithetical to also attaining high repetition rates due to heat load in the compressor⁷². If driver power was not limited, achieving an energy of 1 TeV in a single stage would require an acceleration length of 240 m, for which the accelerating gradient

⁶⁵ [80] C. B. Schroeder *et al.* (2010) *Physical Review Special Topics - Accelerators and Beams*

⁶⁶ [81] C. M. S. Sears *et al.* (2010) *Physical Review Special Topics - Accelerators and Beams*

⁶⁷ [82] S. Kneip *et al.* (2012) *Physical Review Special Topics - Accelerators and Beams*

⁶⁸ [83] D. H. Dowell (2016) *arXiv e-prints*

⁶⁹ [84] Z. Qin *et al.* (2018) *Physics of Plasmas*

⁷⁰ [85] A. Cianchi *et al.* (2013) *Nucl. Instrum. Methods Phys. Res. A: Accel. Spectrom. Detect. Assoc. Equip.*

⁷¹ [55] A. Gonsalves *et al.* (2019) *Physical Review Letters*

⁷² [44] W. Lu *et al.* (2007) *Physical Review Special Topics - Accelerators and Beams*

⁷² [86] P. Poole *et al.* (2013) *Optics Express*

⁷² [87] C. N. Danson *et al.* (2019) *High Power Laser Science and Engineering*

would have dropped to 3.7 GV m^{-1} . This is similar to the direction taken by AWAKE using beam-driven acceleration and is one possible route for achieving TeV energies.

An alternative is to operate at a higher density, and simply replace the laser pulse when it depletes, or the electron beam dephases. This is known as staging and was demonstrated in 2016 by the group at LBNL⁷³. This proof-of-principle dual-stage experiment captured approximately 3% of the first stage's charge in the second stage. This experiment was performed at the 100 MeV level and used an active plasma lens to focus the beam into the second stage⁷⁴. Staging presents three main challenges⁷⁵:

1. compact laser pulse injection;
2. stability; and,
3. emittance preservation.

The first problem is currently addressed by using a plasma mirror placed near the focus to inject the beam within millimetres of the target⁷⁶. This can make addressing the third problem more difficult, as placing a hot, high-density plasma in the beam will have detrimental effects on the emittance⁷⁷. Curved plasma channels for guiding intense laser pulses could be an alternative to plasma mirrors⁷⁸. Achieving stability is difficult owing to the small scales of the structures involved; the tolerance level has been found to be 0.1 μm in early studies⁷⁹. Effective active stabilisation is likely to be essential to achieve the required level of stability⁸⁰.

While electron beams can be injected with ultra-low emittances, maintaining this emittance over a large number of stages is challenging. Coulomb scattering in the plasma mirror will need to be managed in some way, such as by minimising the thickness of the mirror by using liquid crystal films⁸¹. Further, the beam can scatter off ions in the channel⁸², and emittance will also grow irreversibly due to space charge effects⁸³, effects that are negligible for single stages but may become significant over longer propagation distances. Controlling emittance growth as the beam leaves and enters stages is non-trivial, owing to the large divergence of LWFA beams coupled to finite energy spread⁸⁴.

⁷³ [88] S. Steinke *et al.* (2016) *Nature*

⁷⁴ [89] J. van Tilborg *et al.* (2015) *Physical Review Letters*

⁷⁵ [90] F. Albert *et al.* (2021) *New Journal of Physics*

⁷⁶ [91] C. Thaury *et al.* (2007) *Nature Physics*

⁷⁷ [92] G. Raj *et al.* (2020) *Physical Review Research*

⁷⁸ [93] J. Luo *et al.* (2018) *Physical Review Letters*

⁷⁹ [94] Y. Ma *et al.* (2021) *Physics of Plasmas*

⁸⁰ [95] R. Assmann *et al.* (1998) *Nucl. Instrum. Methods Phys. Res. A: Accel. Spectrom. Detect. Assoc. Equip.*

[96] S. Cheshkov *et al.* (2000) *Physical Review Special Topics - Accelerators and Beams*

[97] C. Chiu *et al.* (2000) *Physical Review Special Topics - Accelerators and Beams*

⁸¹ [98] F. Wu *et al.* (2020) *Optics Laser Technology*

⁸² [99] A. Zingale *et al.* (2021) *Physical Review Accelerators and Beams*

[100] N. Kirby *et al.* “Emittance growth from Multiple Coulomb Scattering in a plasma wakefield accelerator” (2007)

[101] C. B. Schroeder *et al.* (2013) *Physics of Plasmas*

⁸³ [102] C. A. Lindström *et al.* (2022) *Journal of Instrumentation*

⁸⁴ [103] T. Mehrling *et al.* (2012) *Physical Review Special Topics - Accelerators and Beams*

[104] M. Migliorati *et al.* (2013) *Physical Review Special Topics - Accelerators and Beams*

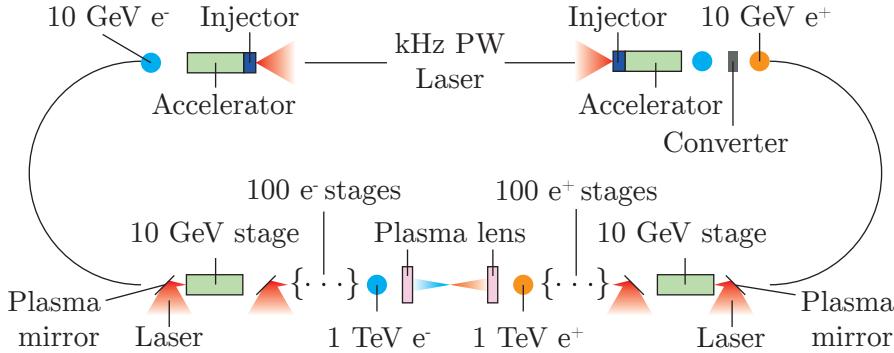


Figure 1.4: Schematic of potential all-plasma electron-positron collider.

If the beam expands between stages, it will need to be re-focused into the next stage, which will require some combination of beam optics and density ramps⁸⁵. The coupling region should not become too large, as this reduces the average accelerating gradient of the accelerator. This is the benefit of active plasma lenses, which are capable of producing focusing fields in a small footprint⁸⁶, but can cause emittance growth through scattering and field non-linearities⁸⁷.

1.4.3 Positron acceleration

It is important to mention positron acceleration, which is an entire half of the plasma collider picture that has been mostly neglected experimentally. Using a high Z material as a converter at the end of a TeV electron accelerator would generate positrons⁸⁸, but a high conversion efficiency is unlikely to be possible if sub-micron emittance is to be preserved⁸⁹. Consequently, low emittance, high charge positron beams will likely require their own accelerator⁹⁰. The highly non-linear regime of plasma wakefield acceleration, which produces emittance preserving linear focusing forces for electrons, is defocusing for positrons. A positron focusing region of the wake exists behind the back of the first bubble, and it has been proposed to use this region to accelerate positrons⁹¹, with recent work suggesting that the focusing forces would be linear in a hot plasma⁹². This scheme would have Coulomb scattering issues, and so acceleration in a hollow core in the centre of the bubble is also being investigated⁹³.

⁸⁵ [105] I. Dornmair *et al.* (2015) *Physical Review Special Topics - Accelerators and Beams*

⁸⁶ [89] J. van Tilborg *et al.* (2015) *Physical Review Letters*

[106] K. Sjobak *et al.* (2021) *Physical Review Accelerators and Beams*

[107] J. van Tilborg *et al.* (2017) *Physical Review Accelerators and Beams*

[108] C. Lindstrøm *et al.* (2018) *Physical Review Letters*

[109] G. Sarri *et al.* (2013) *Plasma Physics and Controlled Fusion*

[110] G. Sarri *et al.* (2017) *Plasma Physics and Controlled Fusion*

[111] M. J. V. Streeter *et al.* (2022) *arXiv e-prints*

[112] S. Corde *et al.* (2015) *Nature*

[113] C. S. Hue *et al.* (2021) *Physical Review Research*

[114] S. Diederichs *et al.* (2019) *Physical Review Accelerators and Beams*

[115] S. Diederichs *et al.* (2020) *Physical Review Accelerators and Beams*

[116] T. Wang *et al.* (2021) *arXiv e-prints*

1.5 Thesis outline

Below are brief descriptions of each of the remaining chapters in this thesis.

Chapter 2

Reviews the theory relevant to the work presented in this thesis. This includes laser pulse descriptions, plasma wave excitation, acceleration dynamics, and transverse beam effects.

Chapter 3

Provides descriptions of the experimental methods used to obtain the data and some of the numerical methods used. Various convergence tests for particle-in-cell simulations are performed in support of the resolution choices made for the simulation results later in the thesis.

Chapter 4

Results from a series of high-resolution density-length parameter scans using a 5 TW laser are presented. These results are used to measure the pulse evolution length, detect beam-driven acceleration, and estimate x-ray brilliance. The optimal parameters for various beam metrics are presented.

Chapter 5

A mask-based emittance measurement technique is described and then used to measure the slice emittance of GeV energy beams. The limitations of ionisation injection are then investigated.

Chapter 6

A novel target design is used to generate electron beams with energies up to 2 GeV. The target configuration is used to deduce the effect of plasma density ramps and plasma mirrors on the electron beam.

Chapter 7

Conclusions from the results chapters are discussed, and suggestions for future work are made.

⁹³ [117] S. Gessner *et al.* (2016) *Nature Communications*
[108] C. Lindström *et al.* (2018) *Physical Review Letters*
[118] T. Silva *et al.* (2021) *Physical Review Letters*

2 Theory

HERE WE REVIEW the theory that will be necessary for the discussions in this thesis. To begin, we will give a brief description of laser pulses, which are central to LWFA. Following this, we will build from the single particle interactions with electromagnetic fields to macroscopic plasma behaviour and the generation of wakefields. Finally, we will transition to viewing these wakefield structures as accelerators, examining particle injection, acceleration, and the transverse dynamics of the electron beam.

2.1 Laser pulses

The behaviour and properties of laser pulses can be deduced from Maxwell's equations¹. In vacuum the source terms are dropped, and, taking the curl, we reproduce the wave equation

$$\nabla^2 \mathbf{E} - \frac{1}{c^2} \frac{\partial^2 \mathbf{E}}{\partial t^2} = 0, \quad (2.1)$$

where \mathbf{E} is the electric field and c is the wave speed. Since there are many solutions to this equation, we are free to choose one where the electric field varies independently in space and time, $\mathbf{E}(\mathbf{r}, t) = \mathbf{R}(\mathbf{r})T(t)$. This allows us to solve Eq. 2.1 by noticing that the derivatives must also be independent, and consequently, the terms must be equal to a constant, k ,

$$\frac{\nabla^2 \mathbf{R}}{\mathbf{R}} = \frac{1}{c^2 T} \frac{\partial^2 T}{\partial t^2} \quad \rightarrow \quad (\nabla^2 + k^2) \mathbf{R} = 0, \quad (2.2)$$

which is the Helmholtz equation. Similarly, we can write down a second-order differential equation for the temporal function, $(\partial_t^2 + \omega^2)T = 0$, where $\omega = ck$, which has oscillatory solutions of the form $T = e^{\pm i\omega t}$.

This motivates choosing a separable trial solution of the form $\mathbf{E} = \mathbf{E}_0(\mathbf{r})e^{ikz}e^{i\omega t}$, allowing us to focus only on the time-independent component. If the length over which the beam envelope varies is much longer than the laser wavelength, then $|\partial_z^2 \mathbf{E}_0| \ll |k \partial_z \mathbf{E}_0|$, which is the paraxial approximation. Substituting $\mathbf{E}_0 e^{ikz}$ into Eq. 2.2 and applying this approximation gives

$$\left(\nabla_{\perp}^2 + 2ik \frac{\partial}{\partial z} \right) \mathbf{E}_0 = 0, \quad (2.3)$$

¹ See appendix A.2.

where $\nabla_{\perp}^2 = (\partial_x^2 + \partial_y^2)$. Transforming to cylindrical coordinates, Eq. 2.3 becomes

$$\left(\frac{\partial^2}{\partial r^2} + \frac{1}{r} \frac{\partial}{\partial r} + 2ik \frac{\partial}{\partial z} \right) \mathbf{E}_0 = 0. \quad (2.4)$$

For the boundary conditions imposed by a laser's spherical mirror cavity, this equation is solved by a series of solutions referred to as Hermite-Gaussians², the lowest order of which is called the Gaussian mode, or TEM_{00} mode, which is usually the dominant mode produced by laser cavities. In linear polarisation, the TEM_{00} mode can be expressed as³

$$E_x(r, z, t) = E_{x,0} \frac{w_0}{w(z)} \exp \left(i \left[(\omega_L t - kz) - \frac{kr^2}{2R(z)} + \tan^{-1} \left(\frac{z}{z_R} \right) \right] - \frac{r^2}{w^2(z)} \right), \quad (2.5)$$

where $E_{x,0}$ is the peak electric field strength, ω_L is the laser frequency, and the wavenumber $k = 2\pi/\lambda$ where λ is the wavelength. Here, the focal plane is at $z = 0$. The value w_0 is the Gaussian waist and refers to the distance from the optical axis at which the field strength drops to $1/e$ of its peak value, at the focal plane. As $|z|$ increases, the $1/e$ intensity width increases as

$$w(z) = w_0 \sqrt{1 + \frac{z^2}{z_R^2}}, \quad (2.6)$$

where z_R is the Rayleigh range, given by

$$z_R = \frac{\pi w_0^2}{\lambda}, \quad (2.7)$$

which is the longitudinal scale length of the laser focus. For $z = z_R$, the peak field strength drops by a factor of $\sqrt{2}$. The wavefronts that are flat at focus become curved as z increases, with a radius of curvature given by

$$R(z) = z \left(1 + \frac{z_R^2}{z^2} \right). \quad (2.8)$$

The term $\tan^{-1}(z/z_R)$ in Eq. 2.5 is known as the Gouy phase shift, and causes a slight increase in the wavelength when the laser passes through focus, an effect that is small enough to be ignored for long Rayleigh lengths⁴. The real component for this solution is plotted in Fig. 2.1, for a 800 nm light focuses to a Gaussian spot size of $w_0 = 1 \mu\text{m}$. Note that these parameters have been chosen for illustrative purposes, allowing the beam envelope changes, wavefront curvature, and wavelength to be observed on the same spatial scale. Focusing a laser this tightly would actually invalidate the paraxial approximation, and would therefore require a more accurate solution⁵.

² [119] F. L. Pedrotti *et al.* *Introduction to Optics* (1987)

³ [120] W. S. C. Chang *Principles of lasers and optics* (2005)

⁴ [121] S. Feng *et al.* (2001) *Optics Letters*

⁵ [122] S. M. Sepke *et al.* (2006) *Optics Letters*

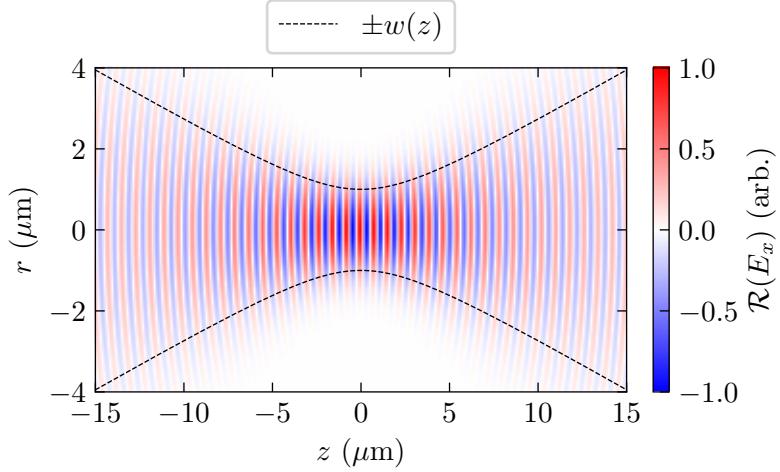


Figure 2.1: Real component of the solution to Eq. 2.5 for 800 nm light focused to a $w_0 = 1 \mu\text{m}$ spot.

2.1.1 Wavefront aberrations

In real laser systems, multiple Hermite-Gaussian modes superpose to produce a transverse field strength that is constant with radius in the near field, as opposed to a Gaussian. This allows for the maximal extraction of energy from an amplifier and also minimises the peak intensity incident on the optics for an equivalent average power. Such a transverse intensity profile in the near field leads to the Airy disk pattern at focus, seen in the first panel of Fig. 2.2. In addition, physical optical systems introduce spatial phase variations, or wavefront aberrations, to the beam, such that $E(r, \theta) = |E(r, \theta)| \exp(i\phi_{abb}(r, \theta))$. The variations add linearly with the phase term imposed by the focusing optic, changing the profile at focus. For circular beam apertures, we can express these aberrations as the sum of a series of functions called Zernike polynomials⁶, $Z(r, \theta)$,

$$\phi_{abb} = \sum_i \alpha_i Z_i(r, \theta), \quad (2.9)$$

where α_i is the coefficient for the i^{th} Zernike polynomial. Each Zernike polynomial has a radial degree, n , and an azimuthal degree, m , with the polynomials being expressed in terms of these indices, Z_m^n . The wavefront aberrations add linearly with the phase variation introduced by a focusing optic, which changes the distribution of energy at focus. The lower-order polynomials are often referred to by names that relate to the focal spot aberrations they produce, such as spherical, astigmatism, or coma. Examples of these aberrated focal spots are plotted in Fig. 2.2. Aberrations from the ideal spot typically reduce the peak intensity and increase the half-maximum beam size. Modifications to the wavefront have been shown to affect the electrons and x-rays produced in laser wakefield

⁶ [123] M. Born *et al. Principles of optics : electromagnetic theory of propagation, interference and diffraction of light* (1999)

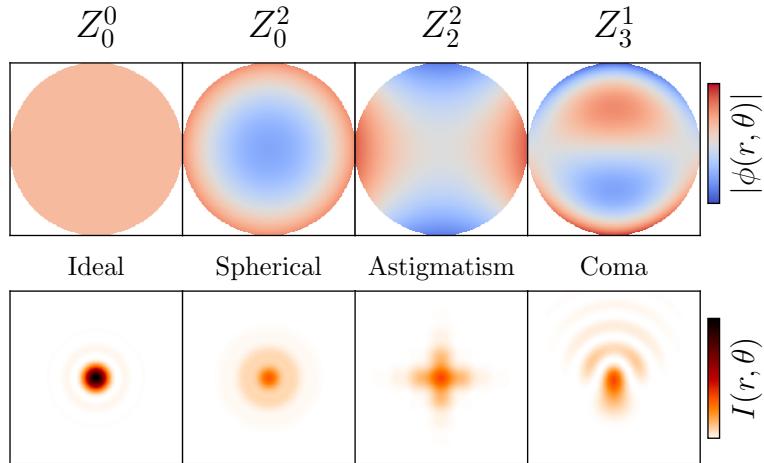


Figure 2.2: Zernike polynomials and the resultant focal spot patterns.

accelerators⁷, while active feedback systems have been used to tailor the wavefront based on the electron beam⁸.

2.2 Single particle motion

The Lorentz equation governs the motion of charged particles in an electromagnetic field,

$$\mathbf{F} = q(\mathbf{E} + \mathbf{v} \times \mathbf{B}), \quad (2.10)$$

where \mathbf{F} is the force vector on the charge. For electromagnetic waves, the magnitude of \mathbf{E} and \mathbf{B} are related by $|\mathbf{B}| = |\mathbf{E}|/c$, so for the second term in Eq. 2.10 to significantly affect the dynamics, the particle speed must approach c .

2.2.1 Non-relativistic motion

First, we will consider the motion at $v \ll c$. In the presence of an oscillating electric field, $\mathbf{E} = E_x \hat{\mathbf{x}} = -E_0 \cos(\omega_L t) \hat{\mathbf{x}}$, an initially motionless electron will undergo harmonic motion in the $\hat{\mathbf{x}}$ direction,

$$v(t) = \frac{eE_0}{m_e \omega_L} \sin \omega_L t. \quad (2.11)$$

Here we see that the magnitude of the electric field determines the peak velocity of the electron. In the non-relativistic limit, the electron velocity is proportional to the vector potential of the laser field, \mathbf{A} , where $\mathbf{E} = -\partial_t \mathbf{A}$. It is helpful to define the parameter a_0 ,

$$a_0 \equiv \frac{eE_0}{m_e \omega_L c}, \quad (2.12)$$

⁷ [124] S. P. D. Mangles *et al.* (2009) *Applied Physics Letters*
 [125] A. Popp *et al.* (2010) *Physical Review Letters*

[126] B. Beaurepaire *et al.* (2015) *Physical Review X*

[127] J. M. Cole *PhD. Thesis* (2016)

⁸ [128] Z. H. He *et al.* (2015) *Nature Communications*
 [129] S. Dann *et al.* (2019) *Physical Review Accelerators and Beams*

which is the normalised peak velocity of the electron quivering in the laser field, or equivalently, the peak amplitude of the normalised laser vector potential, given by $\mathbf{a} = e\mathbf{A}/m_e c$. The parameter allows us to immediately identify the regime of particle motion (non-relativistic for $a_0 \ll 1$)⁹. Numerically, a_0 can be calculated using

$$\begin{aligned} a_0 &\simeq 0.855\lambda[\mu\text{m}]\sqrt{I[10^{22}\text{W m}^{-2}]} \\ &\simeq 0.25E_0[\text{TV m}^{-1}], \end{aligned} \quad (2.13)$$

where, in the second approximation, $\lambda = 800\text{ nm}$. An electron quivering in the field has a cycle-averaged kinetic energy of

$$U_p = \frac{1}{2}m_e\langle v^2 \rangle, \quad (2.14)$$

where U_p is the non-relativistic ponderomotive potential, and the angled brackets denote taking the average value over a cycle. We can see that this can be considered as a potential source of energy if we imagine a scenario where an electron only experiences an odd number of half cycles of the oscillating field, for example, due to a reflection of the field, at which point this potential energy would be suddenly converted to kinetic energy. This potential has an associated force, the ponderomotive force, given by $\mathbf{F}_p = -\nabla U_p$. Physically, we can see this as the electron being ‘kicked’ at a rate of ω_L in alternating directions, but with a consistently weaker kick strength in one direction due to the intensity gradient, causing the electron will gradually accelerate. This is a more gradual mechanism via which the energy in the potential can be converted into kinetic energy. Averaging over the sinusoidal motion of the electron, the non-relativistic ponderomotive force for a linearly polarised laser is

$$\mathbf{F}_p = -\frac{1}{4}m_e c^2 \nabla a^2, \quad (2.15)$$

where $a^2 = |\mathbf{a}|^2$. Note that in circular polarisation the force is a factor of 2 stronger.

2.2.2 Relativistic motion

When the particle speed approaches c , the effect of the magnetic field becomes significant. Using a vector potential $\mathbf{A} = A_x \hat{\mathbf{x}} = a_0 \sin \theta \hat{\mathbf{x}}$, where $\theta = kz - ct$, the electromagnetic fields are computed by $\mathbf{E} = -\partial_t \mathbf{A}$ and $\mathbf{B} = \nabla \times \mathbf{A}$. As shown by the derivation in appendix A.3, the electron momentum in a stationary frame, normalised to $\mathbf{p}/mc = \mathbf{p}$, evolves according to

$$p_x = a_0 \sin \theta, \quad (2.16)$$

$$p_z = \frac{a_0^2}{2} \sin^2 \theta. \quad (2.17)$$

⁹ This is assuming the particle is initially at rest.

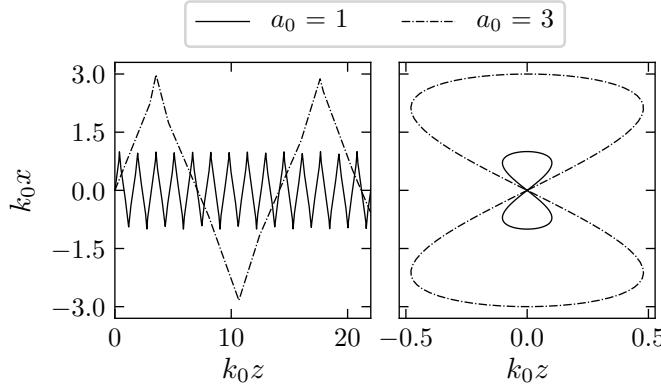


Figure 2.3: Electron trajectories in the presence of an oscillating electromagnetic field with spatial frequency k_0 and field strength a_0 . The left panel shows the motion in the laboratory frame, while the right panel shows the motion in a reference frame that moves at the average drift velocity of the particle.

If the electrons are initially at rest, then their coordinates as a function of θ are

$$x = a_0 \cos \theta \quad (2.18)$$

$$z = \frac{a_0^2}{4} \left(\theta - \frac{1}{2} \sin 2\theta \right), \quad (2.19)$$

where $x = kx$. These solutions are plotted in the left panel of Fig. 2.3, for a_0 values of 1 and 3. When the effects of the magnetic field are included the particle drifts in the laser propagation direction. This is because the particle's transverse motion changes sign at the same time as the magnetic field, so the motion perpendicular to the magnetic field is always in the same direction. The cycle average drift velocity scales with the laser intensity as

$$\frac{v_d}{c} = \frac{\langle p_z \rangle}{\langle \gamma \rangle} = \frac{a_0^2}{4 + a_0^2}. \quad (2.20)$$

In a frame moving with the electron at v_d , the particle's trajectory traces out a 'figure-of-eight', as plotted in the right panel of Fig. 2.3.

Relativistic electrons also experience a weaker force due to intensity gradients than non-relativistic electrons. The relativistic ponderomotive force of the laser was derived by Mora *et al.*¹⁰, which for linear polarisation is given by

$$\mathbf{F}_p = -\frac{m_e c^2}{4\langle \gamma \rangle} \nabla a^2, \quad (2.21)$$

where γ is the Lorentz factor of the electron.

¹⁰ [130] P. Mora *et al.* (1997) *Physics of Plasmas*

Table 2.1: Ionisation intensity thresholds for a selection of ions, calculated from Eq. 2.22. The a_0 values assume a laser wavelength of 800 nm.

Ion	E_{ion} (eV)	I_{BSI} (W m $^{-2}$)	a_0
H $^+$	13.6	1.4×10^{18}	0.0079
He $^+$	24.6	1.5×10^{19}	0.026
He $^{2+}$	54.4	8.8×10^{19}	0.064
C $^+$	11.3	6.4×10^{17}	0.0055
C $^{5+}$	392.1	3.8×10^{22}	1.3
C $^{6+}$	489.0	6.4×10^{22}	1.7
N $^+$	14.5	1.8×10^{18}	0.0091
N $^{6+}$	552.1	1.0×10^{23}	2.2
N $^{7+}$	667.0	1.6×10^{23}	2.7

2.3 Ionisation

The ionisation of initially neutral gas is a necessary precursor to LWFA. Although for weakly bound electrons, such as those in the outer shells of atoms, the laser is typically so intense that these can be assumed to be ionised a long time before the main part of the pulse arrives, large nuclei can produce field strengths that are comparable to those of the laser. We must therefore consider the exact process by which a bound electron ionises.

Consider a simple model of an atom, where the potential from a nucleus of charge Ze is modified by a stationary electric field of strength E_0 . The potential at a position x is $U(x) = -Ze^2/4\pi\epsilon_0 x - eE_0 x$, which has a maximum at $x_{\text{max}} = \sqrt{Ze/4\pi\epsilon_0 E_0}$. This is the position of the potential barrier that separates a bound electron from the continuum. We want to know the field strength where the maximum of the potential dips below the ionisation energy, where $U(x_{\text{max}}) = E_{\text{ion}}$, at which point the electron is no longer bound. This is known as barrier suppression ionisation (BSI), and the intensity at which this occurs, I_{BSI} , serves as a useful boundary above which an unbound electron can be assumed.

$$I_{\text{BSI}} = \frac{\pi^2 c \epsilon_0^3}{2e^6} \frac{E_{\text{ion}}^4}{Z^2} \simeq 4 \times 10^{13} \frac{(E_{\text{ion}}[\text{eV}])^4}{Z^2} \text{W m}^{-2}. \quad (2.22)$$

Ionisation at intensities below this threshold is also possible and can proceed by two mechanisms. A bound electron can be liberated by absorbing a series of high-frequency photons, each with energy $\hbar\omega_L$, which sum to equal or exceed the ionisation energy. This is known as multiphoton ionisation, which allows for ionisation at intensities that are orders of magnitude lower than those listed in Tab. 2.1. Alternatively, an electron can be ionised by tunnelling through a partially suppressed barrier, which can occur at intensities close to I_{BSI} . This occurs because the potential far from the nucleus is lower than the electrons' bound potential, so there is a non-zero probability for the electron to be found outside of the potential well.

Keldysh showed that the ratio of the time taken for an electron to tunnel through the barrier to the period of the laser determines if tunnelling ionisation dominates or if multiphoton ionisation dominates¹¹,

$$\gamma_K = \frac{\tau_t}{T_l} = \sqrt{\frac{E_{\text{ion}}}{2U_p}} \quad (2.23)$$

For $\gamma_K^2 \ll 1$, the frequency is low and the field strength is high, favouring the tunnelling mechanism, while for $\gamma_K^2 \gg 1$ the frequency is high with a relatively unperturbed potential, which favours multiphoton ionisation. For typical field strengths in LWFA, the tunnelling mechanism dominates the ionisation process.

The tunnelling ionisation rate for complex atoms was determined by Ammosov, Delone¹², who built on earlier work by Perelomov *et al.*¹³. An approximation for this rate gives¹⁴

$$W_{\text{ADK}}[\text{s}^{-1}] \approx 1.52 \times 10^{15} \frac{4^{n^*} E_{\text{ion}}[\text{eV}]}{n^* \Gamma(2n^*)} (20.5 E^*)^{2n^*-1} e^{-6.83 E^*}, \quad (2.24)$$

where Γ is the gamma function¹⁵, and

$$E^* = \frac{E_{\text{ion}}^{\frac{3}{2}}[\text{eV}]}{E_0[\text{GV m}^{-1}]}, \quad (2.25)$$

and

$$n^* = 3.69 \frac{Z}{\sqrt{E_{\text{ion}}[\text{eV}]}} \quad (2.26)$$

is the effective principle quantum number for a target ionisation level of Z . See appendix A.4 for a formula that includes the effects of quantum numbers l and m on W_{ADK} .

As a laser pulse passes through a gas or plasma, the fraction of the atoms or ions that are ionised can be calculated using

$$f = 1 - e^{-\left(\int_0^{t'} W_{\text{ADK}}(t') dt'\right)}. \quad (2.27)$$

This integral has been calculated in Fig. 2.4, for the transitions from H^{1+} to H^{2+} , C^{4+} to C^{5+} , and N^{6+} to N^{7+} , for a laser pulse with a Gaussian envelope. The peak a_0 of the laser is 2, and the Gaussian width is 19 fs. The helium ion is completely ionised before the bulk of the pulse arrives, whereas the inner electrons of carbon are liberated later. Note that, despite the peak field strength being below the barrier suppression threshold for the N^{7+} ion, the ultimate ionisation fraction is almost 0.5. Ionisation is often treated

¹¹ [131] L. V. Keldysh (1965) *Sov. Phys. JETP*

¹² [132] M. V. Ammosov *et al.* (1986) *Sov. Phys. JETP*

¹³ [133] A. Perelomov *et al.* (1966) *Sov. Phys. JETP*

[134] A. Perelomov *et al.* (1967) *Sov. Phys. JETP*

¹⁴ [135] D. L. Bruhwiler *et al.* (2003) *Physics of Plasmas*

¹⁵ $\Gamma(z) = \int_0^\infty t^{z-1} e^{-t} dt$; in our case the argument is always real and positive.

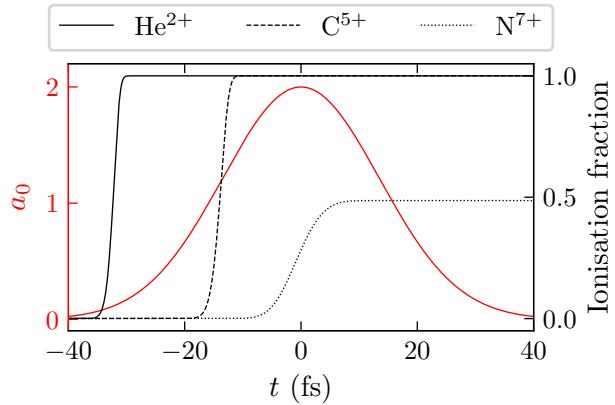


Figure 2.4: Ionisation fraction calculated for different ions during the passage of a 800 nm laser pulse.

as a threshold effect, but for pulse lengths in the 10s of femtoseconds, ionisation fractions at field strengths significantly below threshold can be substantial.

2.4 Plasma waves

A plasma can be defined as a quasi-neutral ensemble of charged particles that interact electromagnetically, exhibiting collective behaviour. They can support many different types of waves because the mode of oscillation depends on the interplay of various conditions, including whether or not an EM wave or magnetic field is present, the species of interest (electrons or ions), and the relative orientation of the magnetic field to the EM wave. In this work, we will restrict ourselves to the domain of electron plasma waves, without an external magnetic field. We will proceed from cold, electrostatic oscillations in 1D, to highly nonlinear electromagnetic waves in 3D.

2.4.1 Plasma oscillations

Consider a slab of cold, neutral plasma with electron and ion densities of n_e and n_i , where the ions are singly charged. Gauss' law states that there will be an electric field associated with any charge imbalance,

$$\nabla \cdot \mathbf{E} = -\frac{e}{\varepsilon_0}(n_e - n_i). \quad (2.28)$$

At equilibrium, $n_e = n_i$, and so no field exists. Perturbing the electron density slightly in the x direction, the field becomes $\varepsilon_0 E = -en_e x$. Using the Lorentz force from Eq. 2.10, for $v \ll c$, the field accelerates the displaced electrons

$$m_e \frac{d^2 x}{dt^2} = -\frac{n_e e^2}{\varepsilon_0} x. \quad (2.29)$$

The electrons, therefore, oscillate around their equilibrium positions at the plasma frequency,

$$\omega_p = \sqrt{\frac{n_e e^2}{m_e \varepsilon_0}}. \quad (2.30)$$

The ions will experience an equal and opposite force to the electrons but will accelerate at a rate reduced by a factor m_i/m_e , where m_i is the ion mass. Consequently, the ions are normally treated as static for timescales on the order of ω_p^{-1} . Since neighbouring regions of plasma are electrostatically shielded from each other, charge oscillations do not propagate in cold plasmas. This can be seen in the lack of wave number dependence in Eq. 2.30, such that the group velocity, $d\omega_p/dk = 0$; consequently, the phase velocity can be anything, and can have any wavelength. The wavelength of a plasma wave moving at $v \simeq c$ is $\lambda_p = 2\pi c/\omega_p$, a quantity often referred to as the plasma wavelength¹⁶, and can be thought of as the distance light would travel in a vacuum during one plasma oscillation period.

2.4.2 Electromagnetic waves

For an electromagnetic wave propagating in a cold plasma, the dispersion relation is¹⁷

$$\omega_L^2 = \omega_p^2 + k^2 c^2. \quad (2.31)$$

The density for which $\omega_p = \omega_L$ and hence $k = 0$ is the critical density, $n_c = m_e \varepsilon_0 \omega_L^2 / e^2$. Above this density, the plasma electrons can respond quickly enough to counteract the incident field oscillations, reflecting the radiation, with only an evanescent field penetrating the plasma up to a length of c/ω_p . Below n_c the plasma is transparent, and the phase and group velocities are

$$v_p = \frac{c}{\sqrt{1 - \frac{n_e}{n_c}}}, \quad v_g = c \sqrt{1 - \frac{n_e}{n_c}} = c \sqrt{1 - \frac{1}{\gamma_g^2}}, \quad (2.32)$$

where γ_g is the relativistic Lorentz factor associated with the group velocity of the radiation. While v_p is always greater than c , the group velocity is always below c , preserving causality. We have also found that the Lorentz factor associated with the group velocity is $\gamma_g = \omega_L/\omega_p = \sqrt{n_c/n_e}$.

2.4.3 Linear wakes

The ponderomotive force of a laser pulse can drive a plasma wave. This can be seen by considering the fluid equation

$$m_e \left(\frac{\partial \mathbf{v}}{\partial t} + (\mathbf{v} \cdot \nabla) \mathbf{v} \right) = -e(\mathbf{E} + \mathbf{v} \times \mathbf{B}). \quad (2.33)$$

¹⁶ Although arguably, there is no such thing as a plasma wavelength.

¹⁷ See appendix A.5.

The transverse velocity is due to the potential of the laser, \mathbf{A} , such that $\mathbf{v}_\perp = e\mathbf{A}/m_e$. Consequently, we see that the difference between the second terms on each side of Eq. 2.33 is the ponderomotive force,

$$m_e(\mathbf{v} \cdot \nabla)\mathbf{v} + e\mathbf{v} \times \mathbf{B} = \frac{e^2}{m_e}((\mathbf{A} \cdot \nabla)\mathbf{A} + \mathbf{A} \times (\nabla \times \mathbf{A})), \quad (2.34)$$

$$= \frac{e^2}{2m_e}\nabla|\mathbf{A}|^2. \quad (2.35)$$

Using this, and the normalisation $\mathbf{A} = m_e c \mathbf{a}/e$, and $|\mathbf{a}|^2 = a^2$, Eq. 2.33 becomes

$$m_e \frac{\partial \mathbf{v}}{\partial t} = -e\mathbf{E} - \frac{m_e c^2}{2}\nabla a^2. \quad (2.36)$$

The current density $\mathbf{J} = -en_e\mathbf{v}$, while the charge density $\rho = -e(n_e - n_i)$, such that charge conservation gives

$$-e \frac{\partial(n_e - n_i)}{\partial t} - e\nabla(n_e\mathbf{v}) = 0. \quad (2.37)$$

For linear wakes, we will assume small perturbations to the electron density, such that $n_e \rightarrow \delta n_e + n_e$, where $\delta n_e \ll n_e$, and small fluid velocities such that $|\mathbf{v}| \ll c$. Differentiating Eq. 2.37 with respect to time, and taking the divergence of Eq. 2.36 (operating with $\nabla \cdot$) gives

$$n_e \nabla \cdot \frac{\partial \mathbf{v}}{\partial t} = -\frac{\partial^2(\delta n_e)}{\partial t^2}, \quad (2.38)$$

$$m_e \nabla \cdot \frac{\partial \mathbf{v}}{\partial t} = -e(\nabla \cdot \mathbf{E}) - \frac{m_e c^2}{2}\nabla^2 a^2. \quad (2.39)$$

Recognising that $\nabla \cdot \mathbf{E} = -e\delta n_e/\epsilon_0$, we reach the wave equation for linear wakes,

$$\left(\frac{\partial^2}{\partial t^2} + \omega_p^2 \right) \frac{\delta n_e}{n_e} = \frac{1}{2}c^2 \nabla^2 a^2. \quad (2.40)$$

A solution to Eq. 2.40 is plotted in Fig. 2.5, where the driving laser has a Gaussian temporal envelope with a width of $c\tau = \lambda_p/4$ and a peak a_0 of 0.01. The oscillation is sinusoidal, with a $\pi/2$ phase difference between the plasma wave and the longitudinal electric field. The electrons receive a ‘kick’ in the forward and backward directions as the laser’s rising and falling edges pass, which excites the oscillation, and consequently, the oscillation will have a larger amplitude if the impulses from the laser add resonantly with the restoring forces from the charge separation, which will occur if the laser pulse length is on the order of $\omega_p^{-1}/2$. The maximum electric field that can be supported by a sinusoidal electron wave can be calculated using Gauss’s law, which gives

$$E_{max} \approx \frac{m_e \omega_p v_g}{e} \equiv \beta_g E_0, \quad (2.41)$$

where E_0 is known as the cold wavebreaking limit. The electric field in Fig. 2.5 is normalised to this limit.

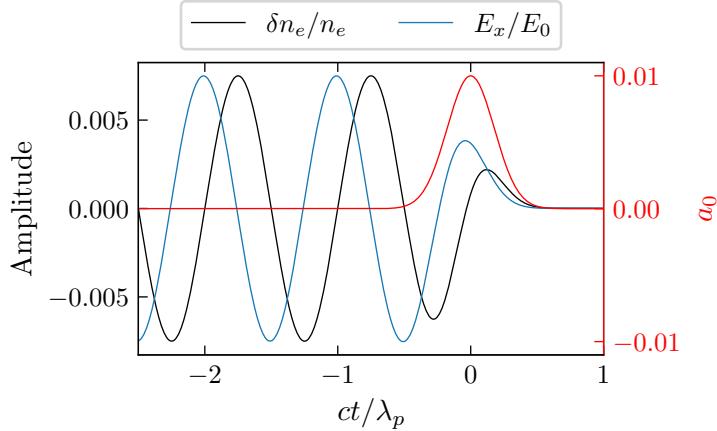


Figure 2.5: Linear electron density wave excited by a laser pulse with a Gaussian temporal envelope with width $c\tau_L = \lambda_p/4$.

2.4.4 Energy gain

Electrons can gain energy by copropagating with the laser pulse. In the boosted frame of the wave moving at the phase velocity, γ_p , an electron will undergo longitudinal simple harmonic oscillations in the field of the wake. While the magnitude of the electric field is invariant under this transformation, the x coordinate is stretched by γ_p , such that the boosted field has the form $E' = E_0 \sin(k_p x'/\gamma_p)$, where the ' denotes boosted values. The electron has maximum velocity when it has propagated from the inflection point of the density wave to its trough¹⁸, over a distance of $\gamma_p \lambda_p/4$,

$$W'_{max} = eE_0 \int_0^{\gamma_p \lambda_p/4} \sin(k_p x'/\gamma_p) dx'' = \gamma_p m_e c^2, \quad (2.42)$$

and the momentum is $cp'_x = \gamma_p \beta_p m_e c^2$. Applying the Lorentz boost to return to the stationary frame, we get

$$\begin{pmatrix} \gamma_p & \gamma_p \beta_p \\ \gamma_p \beta_p & \gamma_p \end{pmatrix} \begin{pmatrix} \gamma_p m_e c^2 \\ \gamma_p \beta_p m_e c^2 \end{pmatrix} = \begin{pmatrix} \gamma_p^2 (1 + \beta_p) m_e c^2 \\ 2\gamma_p^2 \beta_p m_e c^2 \end{pmatrix}, \quad (2.43)$$

$$W_{max} \simeq 2\gamma_p^2 m_e c^2 \quad (2.44)$$

$$= \frac{1.14 \text{ GeV}}{\lambda_L [\mu\text{m}] n_e [10^{24} \text{ m}^{-3}]}, \quad (2.45)$$

where we have taken $\beta_p \simeq 1$, and λ_L is the laser wavelength. This is the energy scaling originally derived in the landmark paper by Tajima and Dawson¹⁹. The physical picture in the lab frame is one where the electron ‘surfs’ on the plasma wave; since both the wave and the particle are moving at relativistic velocities, the energy gain of the electron does not cause it to immediately outrun the wave. Lower densities lead to faster wakes

¹⁸ Regions of high plasma density are defocusing for electrons.

¹⁹ [27] T. Tajima *et al.* (1979) *Physical Review Letters*

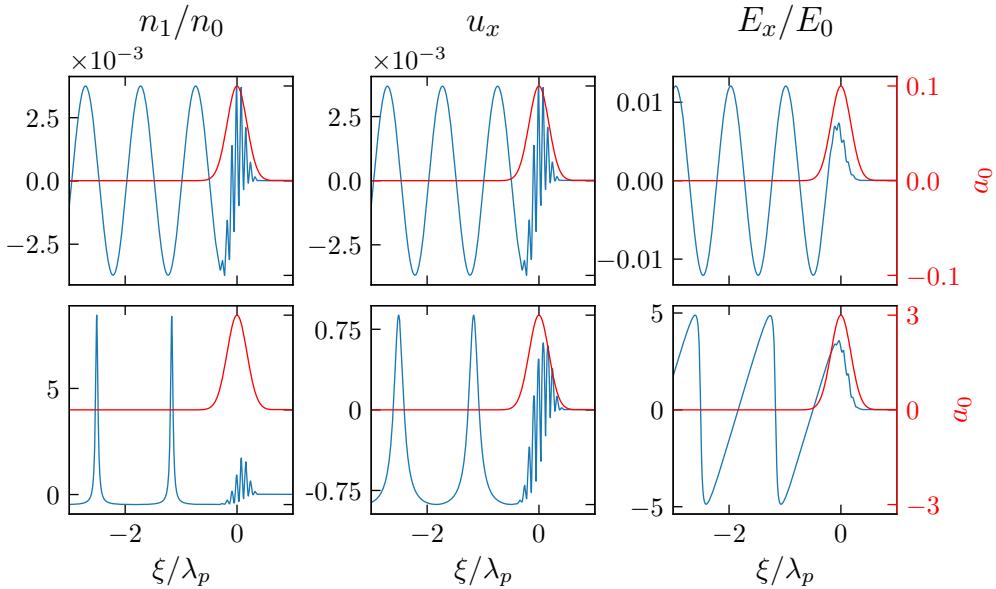


Figure 2.6: Non-linear wake excitation for laser pulses with peak a_0 of 0.1 and 3 for the top and bottom panels respectively. Note the density perturbation $n_1 = n - n_0$.

and hence longer acceleration times, which more than offsets the reduced accelerating field strength.

2.4.5 Non-linear wakes in 1D

For pump strengths where the transverse quiver momentum $p_y \gg m_e c$, a fully relativistic model is required. Using the quasistatic approximation²⁰, we transform the variables (z, t) to (ξ, τ_Q) , where $\xi = x - v_p t$, and $\tau_Q = t$, and assume that the driving envelope varies slowly, such that $\partial_{\tau_Q} = 0$. Under this approximation, the plasma wave is purely a function of ξ . Sprangle *et al.* found the potential for sufficiently underdense plasmas²¹, while for arbitrary wave velocities, the solution for the potential is²²

$$\frac{\partial^2 \phi}{\partial \xi^2} = \gamma_p^2 k_p^2 \left(\beta_p \left[1 - \frac{1 + a^2}{\gamma_p^2 (1 + \phi)^2} \right]^{-\frac{1}{2}} - 1 \right) \quad (2.46)$$

$$= \gamma_p^2 k_p^2 (\beta_p \Phi - 1), \quad (2.47)$$

from which the plasma density can be calculated using $n = \gamma_p^2 \beta_p (\Phi - \beta_p)$, and the fluid velocity $u = (\beta_p - \Phi^{-1}) (1 - \beta_p \Phi^{-1})^{-1}$. The maximum field strength can be predicted using $E_{\max}/E_0 = a_0^2/\sqrt{1 + a_0^2/2}$, where E_0 is the cold wavebreaking limit²³. While analytic solutions for these equations do exist for specific pulse shapes²⁴, we will proceed by solving them numerically.

²⁰ See appendix A.6.

²¹ [136] P. Sprangle *et al.* (1990) *Physical Review Letters*

²² [137] E. Esarey *et al.* (1995) *Physics of Plasmas*

²³ [138] V. I. Berezhiani *et al.* (1990) *Physics Letters A*

²⁴ [139] S. V. Bulanov *et al.* (1989) *JETP Letters*

Plotted in Fig. 2.6 are solutions to the non-linear wake equations for the density, fluid velocity, and electric field. The laser pulse driving these waves is modelled as a linearly polarised plane wave oscillation modulated by a Gaussian envelope, with the form $a(\xi) = a_0 \cos(k_L \xi) \exp(-\xi^2/(\chi \lambda_p)^2)$, where χ is the pulse length in units of λ_p , which takes the value of $\chi = 0.5$ in Fig. 2.6. Within the envelope of the laser pulse, fast oscillations of the density and fluid velocity are driven by the carrier wave, but these subside quickly since the cycle-averaged perturbation is approximately zero. After the pulse, the waveform depends on the peak intensity of the drive pulse. For $a_0 = 0.1$, the wave is sinusoidal with a wavelength of λ_p , while for $a_0 = 3$ the density takes on a spiked waveform with an elongated wavelength. The oscillation wavelength increase is a result of the relativistic mass increase of the electrons, as $m = \gamma m_e$, where $\gamma = \sqrt{1 + a_0^2/2}$, causing $\omega_p \rightarrow \omega_p/\sqrt{\gamma}$.

2.4.6 Non-linear wakes in 3D

The discussion of 1D non-linear waves has been useful for identifying some of the properties of large amplitude plasma waves, such as the formation of the density spike, the elongation of the oscillation wavelength, and the linear longitudinal electric field between the wavefronts. However, when the transverse dimension is included, the dynamics of the interaction change substantially. The total force on the electrons is due to the electromagnetic fields of the wake, and the ponderomotive force,

$$\frac{d\mathbf{p}}{dt} = -[e\mathbf{E} + e\mathbf{v} \times \mathbf{B}]_{\text{plasma}} + \mathbf{F}_p, \quad (2.48)$$

Electrons exactly on the axis will experience purely longitudinal forces, while off-axis electrons will be deflected radially by the gradient of the laser intensity. To solve Eq. 2.48 exactly, the fields need to be computed self-consistently with the electron motion, a task that requires kinetic simulations. In order to gain insight into the electron motion, we make the simplification that the fields from the plasma current are negligible and that consequently, the transverse forces on the electron are due solely to the ion column and the laser. Under the quasistatic approximation²⁵, and assuming cylindrical symmetry, the equation of motion of the electrons in the non-relativistic limit becomes

$$\frac{d^2r}{d\xi^2} = -\frac{r - r_0}{2} - \frac{\nabla|a^2|}{4}, \quad (2.49)$$

where r_0 is the initial radial position of the electron. Here we have used normalised units, $rk_p = r$, $\omega_p t = t$, and $m_e c^2 a = a$.

²⁵ See appendix A.6.

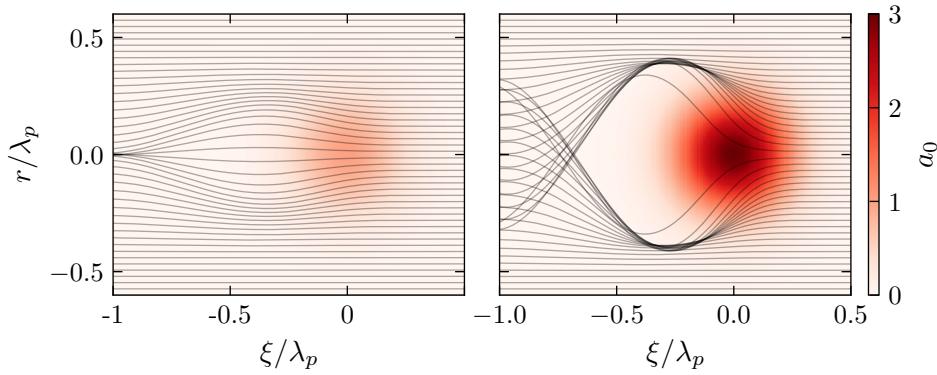


Figure 2.7: Electron trajectories calculated using Eq. 2.49, for a laser pulse with a normalised length and width of $0.25\lambda_p$, the envelope of which is shown in red. The peak a_0 in the left and right panels is 1 and 3 respectively.

Solutions to Eq. 2.49 are plotted in Fig. 2.7 for a laser pulse with a Gaussian pulse length and width of $0.25\lambda_p$, for peak intensities of $a_0 = 1$ and $a_0 = 3$. These plots are similar to those produced by Lu *et al.*²⁶, except that here a laser driver has been used instead of an electron beam²⁷. As the driver intensity increases, electrons are deflected to greater radii, before the electrostatic attraction of the ions pulls them back towards the axis. Once a certain field strength is reached the electron trajectories are seen to cross. More formally, the maximum radius reached by an electron as a function of the initial radius switches from one that steadily increases to one that there is a maximum value. When this occurs, the interaction is said to have entered the ‘blow-out’ or ‘bubble’ regime, with the latter term referring to the circular sheath formed by the electrons as they are deflected around the laser pulse. In simulations reported by Pukhov *et al.* in 2002²⁸, this regime was shown to be able to spontaneously trap some of these crossing trajectories and produce a quasi-monoenergetic energy spectrum, before such observations had been made experimentally.

²⁶ [140] W. Lu *et al.* (2006) *Physics of Plasmas*

²⁷ Interestingly, the laser driver disturbs a much narrower region of plasma than a beam driver, because \mathbf{F}_p decays exponentially, while the electrostatic force decays as $1/r$ for a cylinder of charge.

²⁸ [141] A. Pukhov *et al.* (2002) *Applied Physics B*

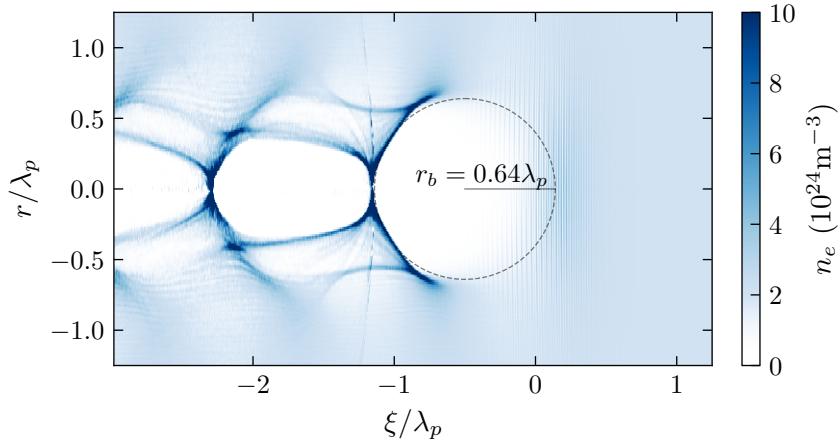


Figure 2.8: Electron density from 3D PIC simulation, illustrating the bubble regime. The laser can be seen to cause fast modulations in the electron density at $\xi \approx 0$, and travels in the $+\xi$ direction. The bulk plasma density is $2 \times 10^{24} \text{ m}^{-3}$, and the laser a_0 is 4.5, producing a bubble with a radius of $0.64\lambda_p$.

The bubble regime is illustrated in Fig. 2.8 using a map of the electron density from a particle-in-cell simulation, where the term ‘bubble’ refers to the approximately spherical sheath of electrons that forms in the laser’s wake. A phenomenological argument can be used to estimate the radius of the bubble. The expulsion of electrons means that the laser cannot self-focus relativistically past a certain point²⁹, and so will reach an equilibrium with a transverse dimension that is approximately equal to the bubble radius r_b . The ponderomotive force at the bubble radius is then

$$\mathbf{F}_p(r_b) = \frac{r_b a_0^2}{2\gamma r_m^2}. \quad (2.50)$$

Equating this with the attractive force from the ions, $r_b/2$, and letting $r_m \approx r_b$, and $\gamma \approx a_0$, the bubble radius is

$$r_b k_p \simeq 2\sqrt{a_0}. \quad (2.51)$$

where the additional factor of 2 was found by Lu *et. al.* using particle-in-cell simulations³⁰. At the moment the snapshot in Fig. 2.8 is taken, the laser peak a_0 is 4.5 and the plasma density is $2 \times 10^{24} \text{ m}^{-3}$, which predicts a bubble radius of $0.64\lambda_p$. As indicated in the figure, this matches the simulation results.

Normalising the electric fields to $E_i = \epsilon_0 E_i / n_e e$ and the magnetic fields to $B_i = B_i / (n_e e \mu_0)$, the fields inside a bubble moving at velocity $v_b \approx c$ are³¹

$$E_z = \frac{\xi}{2}, \quad E_x = \frac{x}{4}, \quad E_y = \frac{y}{4}, \quad (2.52)$$

²⁹ See Sec. 2.5.1.

³⁰ [140] W. Lu *et al.* (2006) *Physics of Plasmas*

[44] W. Lu *et al.* (2007) *Physical Review Special Topics - Accelerators and Beams*

³¹ [142] I. Kostyukov *et al.* (2004) *Physics of Plasmas*

$$B_z = 0, \quad B_x = \frac{y}{4}, \quad B_y = \frac{x}{4}. \quad (2.53)$$

The resulting Lorentz forces on an electron at coordinate (ξ, x, y) , normalised to $m_e \omega_p^2$, are

$$F_z = -\frac{\xi}{2}, \quad F_x = -\frac{x}{2}, \quad F_y = -\frac{y}{2}. \quad (2.54)$$

We see that the bubble provides linear accelerating and focusing forces. Longitudinally, the peak accelerating field scales with the bubble radius, $E_{\max} \propto \sqrt{a_0}$, which is a weaker scaling with intensity than in the 1D case. Transversely, the linear focusing forces can preserve the transverse quality of an electron beam³², making the bubble structure an ideal accelerating cavity.

2.5 Laser pulse evolution

In the preceding section, we have seen how an intense laser pulse can affect an under-dense plasma. Now, we shall look at how this plasma affects the laser pulse. The refractive index in a plasma is given by $\eta = \sqrt{1 - \omega_p^2/\omega_L^2}$, which depends on the relativistic factor of the electrons, the density, and the field frequency.

Variations in any of these quantities will modulate the phase and group velocity, changing the pulse shape. For weak pumps, η can be expressed as³³

$$\eta(n_e, a, \omega) = \left(1 - \frac{1}{2} \frac{\omega_p^2}{\omega_L^2} \left[1 + \frac{\delta n_e}{n_e} - \frac{2\delta\omega}{\omega_L} - \frac{\langle a^2 \rangle}{2} \right] \right), \quad (2.55)$$

where we have used $\omega_p \rightarrow \omega_p / \sqrt{1 + a_\perp^2}$, and performed an expansion to first order using δn_e and $\delta\omega$ as small perturbations to n_e and ω_L . First, we will look at how variations in $\langle a^2 \rangle$ lead to transverse envelope changes, before looking at how the pulse is changed longitudinally by these effects.

2.5.1 Relativistic self-focusing

Refractive index gradients that are transverse to the direction of propagation cause focusing and defocusing of the pulse. Since the normalised potential of the laser varies across the pulse, so does the relativistic plasma frequency $\omega_p/\langle\gamma\rangle^{1/2}$, modifying the refractive index. For an intensity profile that is peaked on-axis, the refractive index gradient will point radially inwards, causing focusing of the pulse. This effect is known as relativistic self-focusing.

³² See Sec. 2.8.

³³ [143] W. B. Mori (1997) *IEEE Journal of Quantum Electronics*

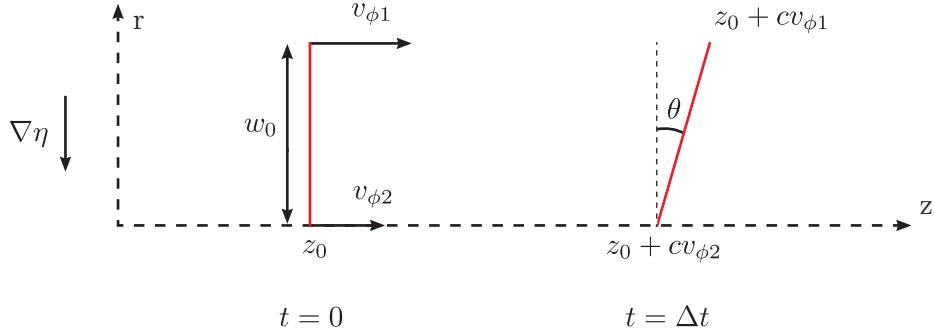


Figure 2.9: The change in angle, θ , of a wavefront propagating in the positive z direction, in the presence of a refractive index gradient.

Following Mori³⁴, we consider an initially flat wavefront propagating through undisturbed plasma, away from z_0 at $t = 0$, as shown in Fig. 2.9. After Δt , the wavefront will acquire an angle to the vertical, given geometrically by $\tan(\theta) \approx \theta = (v_{\phi 1} - v_{\phi 2})\Delta t/w_0$. The difference in phase velocity, $v_{\phi 1} - v_{\phi 2}$, can be written in terms of the gradient, $w_0 \partial_r v_{\phi}$, such that $\partial_t \theta = \partial_r v_{\phi}$. Using $v_{\phi} = c/\eta$, we see that $\partial_t \theta = -c/\eta^2 \partial_r \eta$.

Energy propagates perpendicularly to the wavefronts at the group velocity v_g , which, after acquiring the angle θ , now has a radial component. This radial component has a velocity equal to $v_g \sin \theta \approx v_g \theta$, in the negative r direction, which is equal to the rate of change of the spot size, $\partial_t w = v_g \theta$. Since we are not concerned with variations in the group velocity here, we make the approximation that $v_g \approx c$. Differentiating with respect to time, we find the acceleration of the spot size as a function of the index gradient,

$$\frac{\partial^2 w}{\partial t^2} = \frac{c^2}{\eta^2} \frac{\partial \eta}{\partial r}. \quad (2.56)$$

For small modulations in the refractive index, η can be expressed as

$$\eta \approx 1 - \frac{\omega_L^2}{2\omega_p^2} \left(1 - \frac{\langle a^2 \rangle}{2} \right). \quad (2.57)$$

This can be used to express the acceleration of the spot size in terms of the vector potential

$$\frac{\partial^2 w}{\partial t^2} = \frac{c^2 \omega_p^2}{8\omega_L^2} \frac{\partial a^2}{\partial r}, \quad (2.58)$$

where the assumption that $\eta^2 \approx 1$ has been made. For a Gaussian transverse intensity profile,

$$\frac{\partial a^2}{\partial r} = \frac{1}{w_0} a^2, \quad (2.59)$$

which can be substituted into Eq. 2.58 to give

$$\frac{\partial^2 w}{\partial t^2} = \frac{c^2 \omega_p^2 a^2}{8\omega_L^2 w_0}. \quad (2.60)$$

³⁴ [143] W. B. Mori (1997) *IEEE Journal of Quantum Electronics*

The accelerating contraction of the focal spot is eventually halted and reversed by diffraction. At a certain power level, the diffractive and focusing forces exactly balance, and this level can be found by considering the vacuum diffraction of a pulse as a function of time, given by

$$w(t) = w_0 \sqrt{1 + \left(\frac{ct}{z_R}\right)^2}, \quad (2.61)$$

$$\approx w_0 \left(1 + \frac{1}{2} \left(\frac{ct}{z_R}\right)^2\right), \quad (2.62)$$

where the beam is considered focused at $t = 0$, and the approximation represents a Taylor expansion to first order in $(ct/z_R)^2$. The accelerating increase in the spot size due to diffraction is then,

$$\frac{\partial^2 w}{\partial t^2} = \frac{w_0 c^2}{z_R^2} = \frac{4c^4}{\omega_L^2 w_0^3}, \quad (2.63)$$

where $z_R = \omega w_0^2 / (2c)$ is the Rayleigh range in terms of the laser frequency. Equating this to Eq. 2.60 and rearranging gives the critical condition for constant focal spot size,

$$a^2 w_0^2 = \frac{32c^2}{\omega_p^2}. \quad (2.64)$$

We can relate the term on the left-hand side to the power of the laser pulse by calculating the integral

$$P = \int_A I(r) dA, \quad (2.65)$$

$$= \frac{\epsilon_0 m_e^2 \omega_L^2 c^3 a_0^2}{2 e^2} \int_0^\infty e^{-\frac{2r^2}{w_0^2}} dr, \quad (2.66)$$

$$= \frac{\pi \epsilon_0 m_e^2 \omega_L^2 c^3}{4 e^2} a_0^2 w_0^2. \quad (2.67)$$

$$(2.68)$$

This gives us the critical power for self-focusing to occur,

$$P_c = \frac{8\pi \epsilon_0 m_e^2 c^5}{e^2} \frac{\omega_L^2}{\omega_p^2}, \quad (2.69)$$

$$\approx 17.5 \left(\frac{\omega_L^2}{\omega_p^2}\right) \text{ GW}. \quad (2.70)$$

The threshold is determined by the power, rather than the intensity. Physically, this is because a more tightly focused beam will have a greater intensity and will consequently experience a stronger focusing force from the larger refractive index gradient, but an equally strong diffractive force. The competing forces of self-focusing and diffraction lead to an equation of motion for the spot size. Sprangle *et al.*³⁵ showed that the power

³⁵ [144] P. Sprangle *et al.* (1987) *IEEE Transactions on Plasma Science*

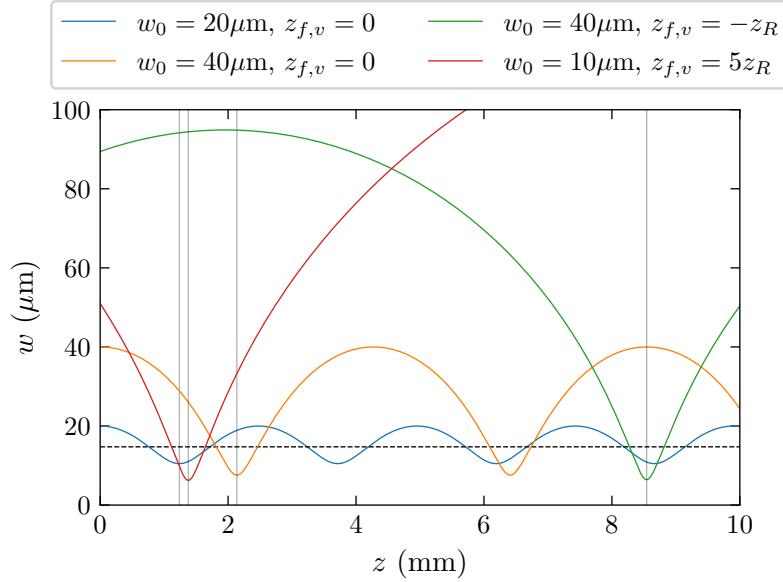


Figure 2.10: Solutions to the envelope equation, Eq. 2.71, for various focusing geometries and focal positions. For all solutions, a_0 is scaled such that $w_0 a_0 = 80 \mu\text{m}$, and the plasma density is $2 \times 10^{24} \text{ m}^{-3}$. The matched spot size at this density $14.7 \mu\text{m}$, indicated in the plot by the dashed horizontal line. The solid vertical lines show the first minima in $w(z)$ for each solution.

normalised waist $X = w/(w_0 a_0)$ evolves according to

$$\frac{d^2X}{dt^2} = V_0 \left(\frac{1}{X^3} - 16\alpha X [h - 1 + 2 \ln 2 + 2 \ln(h + 1)] \right), \quad (2.71)$$

where $h = (1 + X^{-2})^{1/2}$, $V_0 = (2c/(kw_0^2 a_0^2))^2$, and $\alpha = (\omega_p a_0 w_0 / 4c)^2 = P/P_{crit}$. This full expression is required to accurately model the evolution of laser pulses in the moderately relativistic $a_0 \simeq 1$ regime. This equation is stable for $X'' = 0$, which, for $a_0 \gg 1$, occurs for

$$k_p w_m \simeq 2.25 \left(\frac{P}{P_c} \right)^{\frac{1}{6}}, \quad (2.72)$$

where w_m is the matched spot size. The envelope equation has been solved numerically for various initial spot sizes and vacuum focal positions, z_{fv} in Fig. 2.10, where we observe that the laser waist oscillates around the matched spot size for initially mismatched beams. If the initial mismatch is small, such as for the $w_0 = 20 \mu\text{m}$ solution, then the oscillations resemble sinusoidal motion, whereas for larger mismatches the oscillations become asymmetrical, with the pulse being reflected sharply away from focus before gradually refocusing. If the pulse is focused too tightly, then after passing through its plasma focus it can reach an ‘escape velocity’, where it refracts indefinitely, as shown for the $w_0 = 10 \mu\text{m}$, $z_{fv} = 5z_R$.

The vertical lines in Fig. 2.10 indicate the position of the first minimum in $w(z)$ for each solution, which is sensitive to the plasma density, the initial spot size, the laser power, and the distance of the beam from its vacuum focal position. We define the

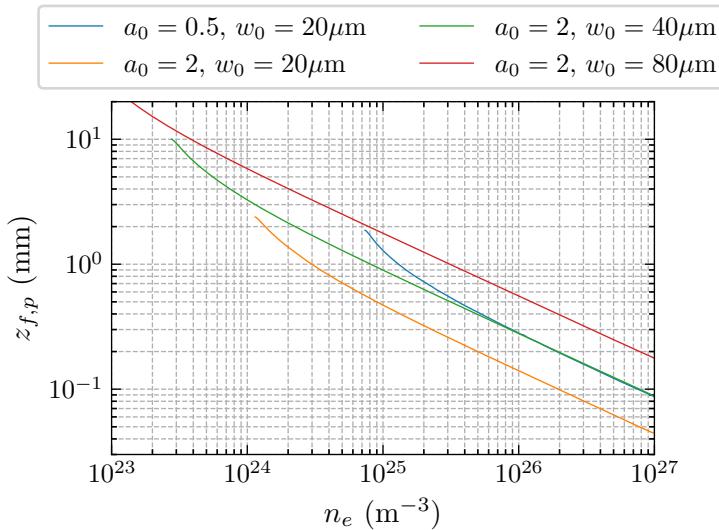


Figure 2.11: Plasma focal length for a wide range of densities, where the laser is at its vacuum focus as it enters the plasma.

plasma focal position, $z_{f,p}$, as the longitudinal position of the first minimum. For a beam with $w_0 = 40 \mu\text{m}$, $z_{f,v} = 0$, we find $z_{f,p} = 2 \text{ mm}$, while a beam with the same w_0 but with a vacuum focus of $z_{f,v} = -z_R$ has $z_{f,p} = 8.5 \text{ mm}$. The vacuum focal position changes both the spot size and the rate of change of spot size when the beam encounters the plasma; an expanding spot will require a longer distance to reach its vacuum focus.

The effect of varying the density on $z_{f,p}$ is shown in Fig. 2.11 for various focal spot sizes and laser intensities, with the lines terminating when the density drops below the threshold density required for self-focusing. As the n_e increases, $z_{f,p}$ reduces due to the increased focusing force from the plasma. At densities significantly above the self-focusing threshold, the approximately linear relationship between $\ln(n_e)$ and $\ln(v_{f,p})$ indicates that a polynomial fit could be used to model this relationship. However, in self-guided laser wakefield experiments, we typically want to operate at low densities, and consequently, predictions of the plasma focal position may require either exact solutions to Eq. 2.71, or a fit to the plot in Fig. 2.11 around the density region of interest.

Guiding up to the point of laser depletion has been observed experimentally³⁶, however the solutions plotted in Fig. 2.10 show envelope behaviour that is observed neither experimentally nor in PIC simulations. Due to the plasma response to the laser, the beam waist does not undergo large amplitude oscillations; instead, the laser focuses to and remains at w_m from Eq. 2.72. This is because, for moderate values of a_0 , the bubble that forms in the wake of the laser helps to guide the pulse. The formation of this bubble was described in Sec. 2.4.6, with Eq. 2.51 predicting $k_p r_b = 2\sqrt{a_0}$. Using this result, in combination with the value of P/P_c , it can be shown that $a_0 \simeq 2(P/P_c)^{1/3}$, or

$$k_p w_m \simeq 2\sqrt{2} \left(\frac{P}{P_c} \right)^{1/6}. \quad (2.73)$$

³⁶ [145] J. E. Ralph *et al.* (2009) *Physical Review Letters*

Interestingly, this is almost identical to Eq. 2.72, except that here we have only used the bubble radius and the critical power for self-focusing.

2.5.2 Pulse compression

Longitudinal refractive index gradients exist within the volume occupied by the laser pulse, which can be due to variations in the plasma density, the carrier wave frequency, or the field strength. Pulse compression can occur when the group velocity at the front of the pulse is lower than at the rear, as shown in Fig. 2.12.

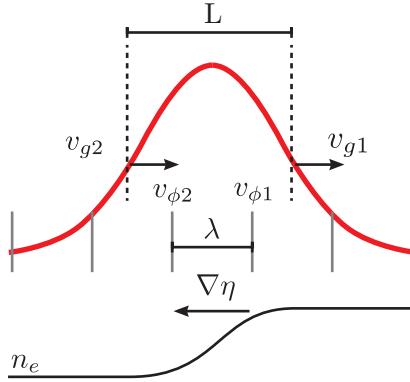


Figure 2.12: Diagram of a laser pulse propagating in a co-moving density ramp, simulating laser-driven plasma wave, with the spatially dependent group velocity denoted by $v_{g1,2}$, and the phase velocities of the wavefronts by $v_{p1,2}$.

A pulse with initial length L changes by $\Delta L = (v_{g1} - v_{g2})\Delta t \approx L\partial_x v_g \Delta t$. The pulse length changes at a rate given by³⁷

$$\frac{\partial L}{\partial t} = L \frac{\partial v_g}{\partial x}. \quad (2.74)$$

On the other hand, wavefronts in regions of higher plasma density will travel faster than those in lower-density regions, leading to a local increase in the wavelength. Considering two wavefronts at x_1 and x_2 initially separated by λ_0 , the local wavelength is $\lambda = x_1 - x_2 = \lambda_0 + (v_{\phi 1} - v_{\phi 2})\Delta t$. The rate of change of the wavelength is then

$$\frac{\partial \lambda}{\partial t} = \lambda \frac{\partial v_{\phi}}{\partial x} = -\frac{c\lambda}{\eta^2} \frac{\partial \eta}{\partial x}. \quad (2.75)$$

For the simplified example in Fig. 2.12, the refractive index gradient is always negative with respect to x , such that only redshifting would occur. However, for intense laser pulses, the gradient in a_0 causes both redshifting and blueshifting independently of the density modulation, a process known as self-phase modulation. In addition, from inspection of the plasma wave density profiles in Fig 2.6, it is clear that both positive and negative density gradients exist, resulting in spectral broadening. This widening of

³⁷ [143] W. B. Mori (1997) *IEEE Journal of Quantum Electronics*

the spectrum is necessary for pulse compression to occur, as the bandwidth and temporal duration of the pulse depend on one another.

In the bubble regime, only the front part of the laser pulse interacts with the plasma, while the rest of the laser energy is guided in a region devoid of electrons. This causes the pulse length to decrease due to the difference in group velocity. Amplification was observed by Streeter *et al.*³⁸ using post-interaction pulse length measurements, where it was found that injection of an electron beam into the wake coincided with the power of the laser reaching a maximum. This was found to occur after the pulse had propagated a distance of

$$L_{\text{evol}} = L \left(\frac{2\omega_L}{3\omega_p} \right) \sqrt{\frac{1}{2} \ln \left(\frac{P}{P_c} \right)}. \quad (2.76)$$

A central frequency shift to $\sim \omega_L/\sqrt{3}$ was observed, in conjunction with a bandwidth increase. Based on these observations, the power amplification was attributed to photons at the front of the pulse transferring energy to the plasma, causing their frequency to drop and recede relative to the front.

2.6 Acceleration

The acceleration of electrons to high energies is the primary purpose of LWFAs. Approximately, the energy that a particle reaches is determined by the product of the accelerating field and the distance over which this field interacts with the electron. In this section we will review the scalings for the maximum accelerating field strength and the acceleration length in various intensity regimes. These will then be combined to give estimates for the electron energies attainable.

2.6.1 Accelerating field

The field strength is determined by the magnitude of the charge separation that can be generated, which will depend on both the plasma density and the driver intensity. Recall that low-intensity drivers ($a_0 \ll 1$) will produce sinusoidal oscillations in the plasma density, which have a maximum amplitude equal to the quiescent plasma density. The maximum field that can be supported by such an oscillation is the cold wavebreaking limit, $E_0 = m_e \omega_p c / e$, which is driver independent because the shape of the wave has been assumed. As was shown in Sec. 2.4, high-intensity drivers produce distorted waves, with spikes in the density profile, which increase the attainable field strength. In the 3D non-linear regime, the cavitation of the bubble of electrons in all directions limits the energy scaling further. To summarise, the field strengths in the different regimes can be

³⁸ [146] M. Streeter *et al.* (2018) *Physical Review Letters*

estimated using³⁹

$$\frac{E}{E_0} = \begin{cases} a_0^2 & (1D, a_0 \ll 1), \\ \frac{a_0^2}{\sqrt{1+a_0^2/2}} & (1D, a_0 \simeq 1), \\ \sqrt{a_0} & (3D, a_0 \gtrsim 2). \end{cases} \quad (2.77)$$

Note that these scalings derive from making various assumptions about the pulse shape, and that the particulars of this will affect the field strength magnitude.

2.6.2 Acceleration limits

As the laser pulse drives the plasma wave, it is continually losing energy, to the plasma electrons on a micro-scale, or to the plasma wave on a macro-scale. Since the energy in the pulse is finite, this process cannot continue indefinitely, and at some point, the laser will no longer support accelerating gradients. If the energy gain of the electron beam is limited by this process, then the acceleration is described as being depletion limited. Shadwick *et al.*⁴⁰ define a depletion length, L_d , as the length over which the laser pulse energy is reduced by a factor of $1/e$. If it is assumed that the plasma wave is being driven by a laser pulse with a Gaussian longitudinal envelope with a pulse length of τ , then the depletion length is⁴¹

$$\frac{L_d}{\lambda_p} = \frac{\omega_L^2}{\omega_p^2} \begin{cases} \frac{1}{2\pi a_0^2} \omega_p \tau & (a_0 \ll 1), \\ \frac{1}{2\pi} \omega_p \tau & (a_0 \gg 1). \end{cases} \quad (2.78)$$

It is interesting to note that the dependence of the depletion length on the driver intensity weakens as the wake amplitude increases, eventually decoupling. Decker *et al.*⁴² identify increasingly localised depletion of the pulse as the process responsible for this decoupling. At high intensities, the front of the laser pulse is interacting with a narrow spike in the plasma density, while the trailing parts of the pulse propagate in a region of low density, or a vacuum in the bubble regime. Consequently, only the pulse front depletes, causing the pulse to etch away from front to back, at a rate that is determined by the width of the density spike and the time required for the energy inside this width to deplete. The pulse front etches away at a speed given by $\beta_{\text{etch}} = n_e/n_c$, such that the pulse of length L will be completely depleted after a time of $L/c\beta_{\text{etch}}$. The field strength no longer affects the depletion length because the width of the density spike and the rate of energy loss in the spike both have the same dependence on a_0 .

Dephasing of the electron beam can also limit energy gain. This is the process of an electron ‘out-running’ the accelerating portion of the wake, due to its velocity being higher than that of the plasma wave. Relative to a wake moving at the linear group

³⁹ [147] L. M. Gorbunov *et al.* (1986) *JETP*

[136] P. Sprangle *et al.* (1990) *Physical Review Letters*

[44] W. Lu *et al.* (2007) *Physical Review Special Topics - Accelerators and Beams*

⁴⁰ [148] B. A. Shadwick *et al.* (2009) *Physics of Plasmas*

⁴¹ [44] W. Lu *et al.* (2007) *Physical Review Special Topics - Accelerators and Beams*

⁴² [149] C. D. Decker *et al.* (1996) *Physics of Plasmas*

velocity, v_g , the electron's speed is $c - v_g \approx 1/2\gamma_p^2$. The accelerating portion of a linear wake is approximately $\lambda_p/2$ long, which means that the electron will move out of this region after propagating a distance of $L_\phi = \lambda_p\gamma_p^2$. This is the linear dephasing length.

For non-linear wakes, the length of the accelerating portion of the wake is modified by the relativistic increase of the oscillation period. In the $a_0 \gg 1$ limit, the non-linear plasma wavelength is⁴³ $\lambda_p^{NL} = 2\lambda_p a_0/\pi$, whereas in the 3D non-linear regime, the accelerating portion of the wake is determined by the bubble radius, $r_b = \lambda_p\sqrt{a_0}/\pi$. In addition to the lengthening of the accelerating portion of the wake, the phase velocity is reduced by the localised etching of the laser pulse⁴⁴, meaning that the electron's velocity relative to the wake is now $3\omega_p^2/2\omega_L^2$. Accounting for these effects, the dephasing lengths in the various regimes are⁴⁵

$$\frac{L_\phi}{\lambda_p} = \frac{\omega_L^2}{\omega_p^2} \begin{cases} 1 & (1D, a_0 \ll 1), \\ \frac{1}{\pi}a_0 & (1D, a_0 \gg 1), \\ \frac{2}{3\pi}\sqrt{a_0} & (3D, a_0 \gtrsim 2). \end{cases} \quad (2.79)$$

The 3D, $a_0 \gtrsim 2$ dephasing length is plotted in the left panel of Fig. 2.13, illustrating the weak effect of a_0 and the strong effect of the plasma density.

2.6.3 Energy gain

Previously, we used the cold wavebreaking field strength and the linear dephasing length to estimate the maximum energy gain from a LWFA as $W_{\max} \approx 2\gamma_p^2 m_e c^2$. We can now see that this energy gain would require a relativistic a_0 if the linear dephasing length is assumed; hence, this is an overly ambitious scaling. Taking into account the appropriate phenomena for each regime, Lu *et al.*⁴⁶ give the following scalings for the maximum energy gain,

$$\frac{W_{\max}}{m_e c^2} = \frac{\omega_L^2}{\omega_p^2} \begin{cases} a_0^2 & (1D, a_0 \ll 1), \\ 4a_0^2 & (1D, a_0 > 1), \\ \frac{2}{3}a_0 & (3D, a_0 \gtrsim 2). \end{cases} \quad (2.80)$$

The 3D ($a_0 \gtrsim 2$) W_{\max} is plotted in the right panel of Fig. 2.13. For the non-linear cases, the average accelerating field experienced throughout the acceleration is half the maximum field, which is generally true for $a_0 \gg 1$. The acceleration has been assumed to be dephasing limited for the 1D linear and 3D non-linear regimes, and depletion limited for the 1D non-linear regime. In the latter case, this is because if the pulse length is set

⁴³ [136] P. Sprangle *et al.* (1990) *Physical Review Letters*

⁴⁴ [149] C. D. Decker *et al.* (1996) *Physics of Plasmas*

⁴⁵ [44] W. Lu *et al.* (2007) *Physical Review Special Topics - Accelerators and Beams*

[136] P. Sprangle *et al.* (1990) *Physical Review Letters*

⁴⁶ [44] W. Lu *et al.* (2007) *Physical Review Special Topics - Accelerators and Beams*

to the bubble radius, $\omega_p \tau = k_p c \tau = 2\sqrt{a_0}$, then

$$L_d = 2 \frac{\omega_L^2}{\omega_p^2} \lambda_p \sqrt{a_0}, \quad (2.81)$$

which has the same scaling with n_e and a_0 as L_ϕ , but is always longer. Hence, in the bubble regime, dephasing will limit the maximum electron energy unless the pulse length is much shorter than the bubble radius.

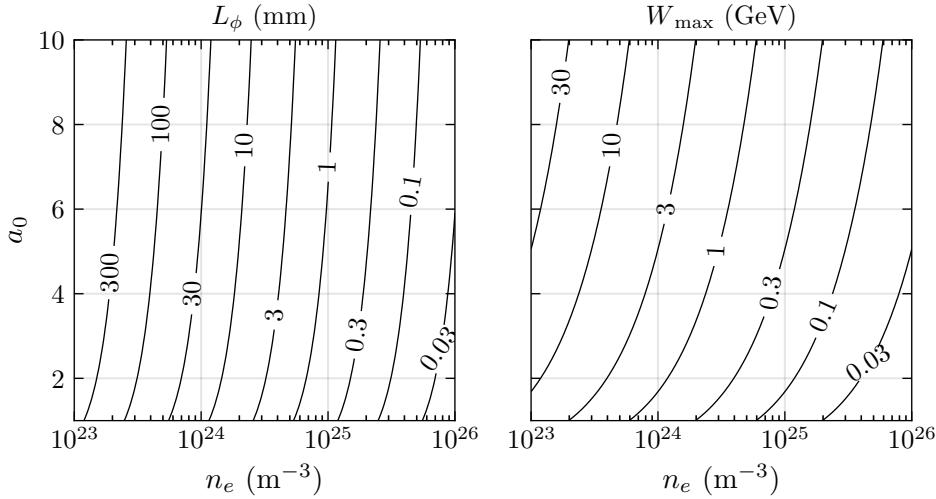


Figure 2.13: Dephasing length and energy gain as functions of plasma density and a_0 .

When attempting to predict energy gain, it is important to use an appropriate value for a_0 . However, this is often difficult to estimate because of the non-linearity of the laser-plasma interaction. For example, in chapter 6, electron beam energies up to 2.2 GeV are reached using a plasma density of $2 \times 10^{24} \text{ m}^{-3}$, with a spot size of $w_0 = 27 \mu\text{m}$. Using the vacuum a_0 value of 2 would predict an energy gain of 1.2 GeV. However, if the matched spot size of 13 μm is used to determine a_0 , then the predicted energy gain is 2.4 GeV, meaning that, in this case, this was the correct spot size to use. Ultimately, these scaling laws need to be interpreted in the context of exactly how and when the electron beam is injected into the wake and the plasma wave evolution that occurs as the beam accelerates.

2.7 Injection

Before a particle can be accelerated it must first be injected into the wave. In this section we review the fundamentals of trapping in a plasma wave, before giving overviews of some of the mechanisms that are used to inject charge in laser wakefield accelerators.

2.7.1 Trapping in a 1D wave

Consider a 1D plasma wave with a potential given by $\phi = \phi_0 \cos(\xi)$, where ϕ_0 is the maximum of the potential, $\xi = k_p(z - v_p t)$ is the phase of the wave, and the wave is

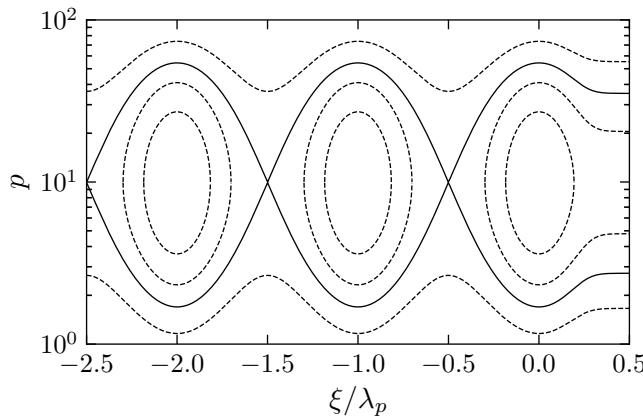


Figure 2.14: Phase space trajectories for electrons in plasma wave with $\gamma_p = 10$. The solid line gives the separatrix between trapped and free trajectories. Electrons follow clockwise trajectories, moving in the $-\xi$ direction below $p = 10$, and in the $+\xi$ direction above $p = 10$.

travelling in the positive direction at v_p . The motion of an electron in this wave is dictated by the Hamiltonian⁴⁷,

$$\mathcal{H}(\gamma, \xi) = \gamma - \beta_p p - \Phi(\xi), \quad (2.82)$$

where γ and p are the electron's Lorentz factor and momentum, $\beta_p = v_g/c$, and $\Phi = e\phi/mc^2$. The derivation of the Hamiltonian can be found in appendix A.7. To determine the conditions for trapping, we can calculate the various phase space orbits available to electrons. Using $\gamma = \sqrt{p^2 + 1}$, we can express Eq. 2.82 as a quadratic in p ,

$$(1 - \beta_p^2)p^2 - 2(\mathcal{H} + \Phi)\beta_p p + 1 - (\mathcal{H} + \Phi)^2 = 0. \quad (2.83)$$

This can be used to calculate trajectories in the (ξ, p) phase space, which has been done in Fig. 2.14 for $\gamma_p = 10$. In this moving frame, particles with low momentum slip backwards at the speed of the wave, while at the other extreme, particles with high momentum overtake the plasma wave.

The critical trajectory exists for $\mathcal{H}_s = \mathcal{H}(\gamma_p, \xi_{\min}) = \gamma_p^{-1} - \Phi_{\min}$, where ξ_{\min} is the point at which $\Phi(\xi_{\min}) = \Phi_{\min}$, and defines a separatrix between phase space orbits that are trapped by the wave and those that are free. This is shown by the solid line in Fig. 2.14. From inspection, we see that trapped orbits must have initial momenta close to γ_p , as the accelerating forces of the wake are too weak to accelerate a particle from rest up to γ_p within a quarter cycle. Amongst the trapped orbits there are two different types; closed and open⁴⁸. The open trapped orbits can be followed by electrons in the bulk plasma that have the necessary momentum, while the closed trapped orbits are inaccessible to these electrons, and instead, particles must be initialised at these points in phase space, such as through external injection. The orbits in Fig. 2.14 are

⁴⁷ [137] E. Esarey *et al.* (1995) *Physics of Plasmas*

⁴⁸ [150] S. V. Bulanov *et al.* (2005) *Physics of Plasmas*

for particles with $\mathcal{H} = 1.5\mathcal{H}_{sep}$ (free), $\mathcal{H} = \mathcal{H}_{sep}/1.5$ (trapped, open), and $\mathcal{H} = \mathcal{H}_{sep}/3$ (trapped, closed). Particles rotate along these trajectories in a clockwise direction, at a speed determined by the particle's distance from γ_p . As stated by Liouville's theorem, the density of particles travelling along these trajectories is constant in time. This means that a group of particles that are injected with a distribution of phases but at the same momentum, will rotate to have a narrower spread of phases but some non-zero spread in momentum. Thus, under the assumptions made here, a small energy spread requires that particles are injected within a narrow region of phase space.

2.7.2 Self-injection

Self-injection is a general term that refers to background plasma electrons somehow becoming trapped in the accelerating and focusing portion of the plasma wave. It was first observed experimentally in 1995 by Coverdale *et al.*⁴⁹, and has since become a common injection method due to its technical simplicity. Self-injection is a highly non-linear process, as it is required for the plasma wave to break in some way, and consequently requires a high intensity drive laser pulse.

There are numerous scaling laws that can be used to predict the onset of self-injection. A model based on observations from PIC simulations, requires the bubble radius, $r_b \gtrsim \sqrt{2}k_p^{-1}\gamma_p$ ⁵⁰. This model is conservative in its predictions of the field strength required for self-injection, as noted in Lu *et al.*⁵¹ and elsewhere, providing an upper bound on the required bubble size. Another model proposed by Thomas *et al.*⁵² finds a less stringent requirement for self-injection, that $r_b \gtrsim 2k_p^{-1}\sqrt{\ln(2\gamma_p^2) - 1}$. Taking the bubble radius $r_b = 2k_p^{-1}\sqrt{a_0}$, the predicted field strength for self-injection as a function of density becomes

$$a_0 \gtrsim \ln(2\gamma_p^2) - 1. \quad (2.84)$$

This threshold is found by calculating the fields inside the bubble in its moving frame, and then integrating the elliptical trajectories of electrons as they travel from the transverse extremes of the bubble to its rear. By injecting transversely, electrons avoid the initially decelerating portion of the bubble, and consequently require less acceleration, or a weaker field, to be trapped. The predicted injection a_0 as a function of density is plotted in Fig. 2.15.

The model by Thomas was found to agree with the experimental results of Mangles *et al.*⁵³, who go on to express the self-injection threshold density in terms of the laser power and the fraction of the laser energy contained within the FWHM of the spot, α ,

$$\alpha \frac{P}{P_c} = \frac{1}{16} \left[\ln \left(\frac{2n_c}{3n_e} \right) - 1 \right]^3, \quad (2.85)$$

⁴⁹ [151] C. A. Coverdale *et al.* (1995) *Physical Review Letters*

⁵⁰ [142] I. Kostyukov *et al.* (2004) *Physics of Plasmas*

⁵¹ [152] I. Kostyukov *et al.* (2009) *Physical Review Letters*

⁵² [44] W. Lu *et al.* (2007) *Physical Review Special Topics - Accelerators and Beams*

⁵³ [153] A. G. R. Thomas (2010) *Physics of Plasmas*

⁵³ [154] S. P. D. Mangles *et al.* (2012) *Physical Review Special Topics - Accelerators and Beams*

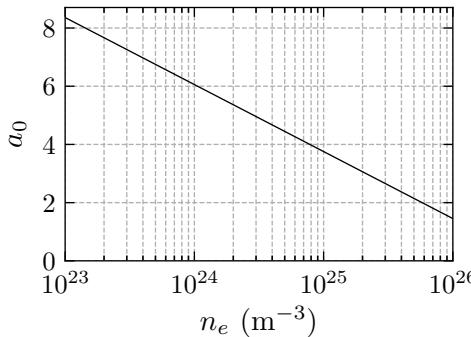


Figure 2.15: The self-injection threshold intensity, as predicted by Thomas *et al.* across the range of densities commonly used for laser wakefield acceleration.

which was shown to have validity over a range of densities from $0.3 - 2 \times 10^{25} \text{ m}^{-3}$. Fundamentally, self-injection is difficult to predict because the laser pulse evolution changes the Hamiltonian of the system. Since pulse evolution is also highly non-linear, cases where we are at the threshold of self-injection often require PIC simulations to support the conclusions of analytic models.

2.7.3 Phase velocity reduction injection

Reductions in the phase velocity of the back of the bubble can relax the intensity requirements for injection because the reduced phase velocity means that the electrons do not need as much longitudinal momentum to become trapped. For example, if the a_0 of the laser is increasing then the bubble radius will expand, reducing the injection threshold⁵⁴. Since ultrashort intense laser pulses tend to amplify as they propagate, some form of injection driven by intensity changes is common, although sometimes undesirable due to its unpredictability.

The bubble radius can also expand due to a decreasing plasma density, which has been observed to trigger injection⁵⁵. Gas targets can be constructed specifically with this aim in mind, as done by Gonsalves *et al.*⁵⁶ and Hansson *et al.*⁵⁷, where multiple gas sources were used to customise the density profile. If the density transition is confined to a small region in space, such as in a hydrodynamic shock, then the energy spread can be reduced⁵⁸, as shown in the work done by Schmid and Buck *et al.*⁵⁹.

⁵⁴ [155] S. Kalmykov *et al.* (2009) *Physical Review Letters*

⁵⁵ [156] H. Suk *et al.* (2001) *Physical Review Letters*

[157] J. Faure *et al.* (2006) *Nature*

[158] A. V. Brantov *et al.* (2008) *Physics of Plasmas*

⁵⁶ [159] A. J. Gonsalves *et al.* (2011) *Nature Physics*

⁵⁷ [160] M. Hansson *et al.* (2015) *Physical Review Special Topics - Accelerators and Beams*

⁵⁸ [150] S. V. Bulanov *et al.* (2005) *Physics of Plasmas*

⁵⁹ [161] K. Schmid *et al.* (2010) *Physical Review Special Topics - Accelerators and Beams*

[162] A. Buck *et al.* (2013) *Physical Review Letters*

2.7.4 Ionisation injection

Ionisation injection offers a technically simple way to gain control over the injection mechanism. By using a low Z gas, such as hydrogen or helium, doped with a certain amount of a higher Z gas, such as nitrogen, control of the injection threshold, as well as over the rate of charge injection, can be obtained. Ionisation injection can reduce the intensity required for injection to occur⁶⁰, as well as increase the beam charge and reduce the divergence⁶¹.

The tightly bound inner shell electrons require a larger field strength to be ionised than the bulk plasma electrons and are therefore ionised at a different phase of the laser pulse and plasma wave. Specifically, cold electrons in the bulk plasma have total energy $\mathcal{H} = 1$, while electrons that are ionised close to the peak of the laser pulse have $\mathcal{H} = 1 - \Phi(\xi_i)$, where $\Phi = e\phi/m_ec^2$ is the normalised potential of the wake at the ionisation phase ξ_i . Electrons will be trapped if $\mathcal{H} \leq \mathcal{H}_{sep}$, where $\mathcal{H}_{sep} = \gamma_p^{-1} - \Phi_{min}$ is the separatrix, and Φ_{min} is the minimum of the potential. The electrons will also have some transverse momentum due to the heating of the laser, $\gamma_{\perp}(\xi_i) = \sqrt{1 + a_{\perp}^2(\xi_i)}$. Including these effects, the trapping condition is⁶²

$$1 - \Phi(\xi_i) \leq \frac{\sqrt{1 + a_{\perp}^2(\xi_i)}}{\gamma_p} - \Phi_{min}. \quad (2.86)$$

The residual transverse momentum a_{\perp} is zero if the electron is ionised at the peak phase of the laser pulse. Generally, this condition cannot be satisfied for $a_0 < 1$ for wake velocities that are appropriate for laser wakefield acceleration, and so the potential of the wake needs to be calculated numerically.

The left panel of Fig. 2.16 shows the potential of a wake, calculated using Eq. 2.46 for a Gaussian laser pulse with an envelope function $a(\xi) = a_0 \exp(-\xi^2/(2k_p\sigma_{\tau})^2)$ for $a_0 = 1.3$ and $k_p\sigma_{\tau} = 1$, the resonant pulse length which produces the largest wave for this a_0 . It is assumed that the electron is ionised at the peak phase, when $a_{\perp} = 0$. For this laser pulse, only electrons that are ionised after the peak of the laser pulse has passed have the potential to be trapped. If an electron is going to be ionised, it is likely to happen before or at the peak of the laser pulse, so, consequently, ionisation injection is unlikely for this laser pulse. A longer laser pulse moves the peak of the field further into the potential, increasing ξ_i , which can increase the chance of ionisation injection occurring, despite the reduction in the wave amplitude that results from moving away from resonance. This is shown in the right panel of Fig. 2.16, which charts the ionisation phase ξ_i as a function of both the laser peak a_0 and the pulse length $k_p\sigma_{\tau} = 1$. The contour of $\xi_i = 0$ can be seen as a soft threshold for ionisation injection.

⁶⁰ [163] T. P. Rowlands-Rees *et al.* (2008) *Physical Review Letters*

[164] C. McGuffey *et al.* (2010) *Physical Review Letters*

⁶¹ [165] A. Pak *et al.* (2010) *Physical Review Letters*

⁶² [166] M. Chen *et al.* (2012) *Physics of Plasmas*

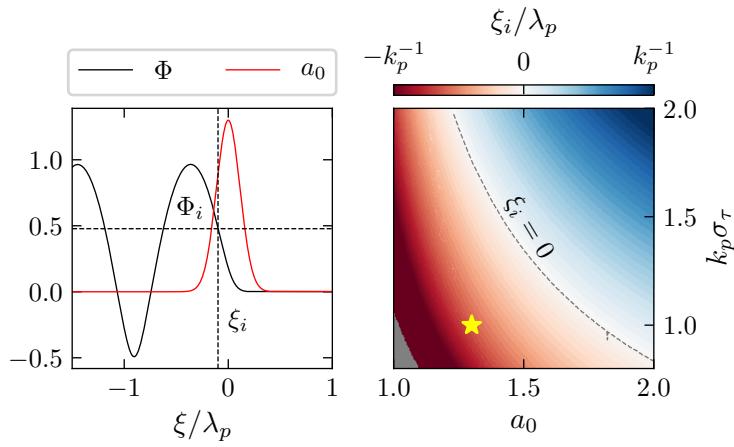


Figure 2.16: Left panel shows the potential of a wake excited by a Gaussian laser pulse with $a_0 = 1.3$, $k_p\sigma_\tau = 1$. The dashed horizontal and vertical lines are the injection condition, $\Phi(\xi_i)$, and injection phase ξ_i . The right panel gives the injection phase as a function of peak a_0 and pulse length, with the yellow star indicating the location of the example in the left panel.

Due to the probabilistic nature of ionisation, electrons will be ionised at different phases of the laser oscillation, when $a_{\perp}^2 > 0$. The electrons are accelerated transversely by the laser in its polarisation direction and retain this momentum once they are captured. In 3D, an additional source of transverse momentum exists for electrons ionised off-axis, where a transverse potential gradient exists. Two-colour ionisation injection has been proposed to alleviate these issues⁶³, with the shorter wavelength, higher field strength pulse being used to ionise the inner shell electrons. The shorter wavelength both reduces the a_{\perp} and allows the laser pulse to be over-focused, reducing the RMS⁶⁴ distance between the ionised electrons and the wake axis, as well as localising injection to a small region in z , reducing energy spread.

2.8 Transverse dynamics

As we saw in Sec. 2.4.6, for large amplitude drivers the strong longitudinal accelerating fields of the wake also exist in the transverse directions. These fields result in LWFA accelerated beams having exceptionally small transverse sizes, on the order of a few microns. For moderate divergences, this means that the beam quality is comparable to conventionally accelerated beams⁶⁵, which is an attractive feature of LWFAs. In addition, the rapid oscillation of high energy electrons in the channel produces high-frequency electromagnetic radiation, typically in the 1 - 50 keV range⁶⁶. These hard x-rays are emitted from a small source, making them suitable for phase-contrast imaging⁶⁷. In this

⁶³

[167] L. L. Yu *et al.* (2014) *Physical Review Letters*

[168] X. Xu *et al.* (2014) *Physical Review Special Topics - Accelerators and Beams*

⁶⁴ Root Mean Squared

⁶⁵ [169] R. Weingartner *et al.* (2012) *Physical Review Special Topics - Accelerators and Beams*

⁶⁶ [170] J. C. Wood *PhD. Thesis* (2016)

⁶⁷ [82] S. Kneip *et al.* (2012) *Physical Review Special Topics - Accelerators and Beams*

section, we will look at the transverse dynamics of electrons in a plasma channel, the definition and behaviour of the beam's emittance, and the properties of the betatron radiation produced by the oscillations.

Consider the motion of an electron in an ion channel, where the electron has some initial displacement from the field axis in the x direction. The focusing forces of the electromagnetic fields inside the bubble are given by Eq. 2.54, such that

$$F_x = -\frac{m_e \omega_p^2}{2} x. \quad (2.87)$$

Using $F_x = \partial_t p_x = \gamma m_e \partial_t^2 x$, the transverse equation of motion is then

$$\frac{d^2 x}{dt^2} = -\omega_\beta^2 x, \quad (2.88)$$

where $\omega_\beta = \omega_p / \sqrt{2\gamma}$ is known as the betatron frequency of the oscillation. At 1 GeV electrons in an ion channel with a plasma density of $1 \times 10^{24} \text{ m}^{-3}$ will oscillate with a wavelength of 2 mm, while at $1 \times 10^{25} \text{ m}^{-3}$ and 100 MeV, this reduces to 0.2 mm. Since electrons are also being accelerated and decelerated, the instantaneous betatron frequency changes with the longitudinal dynamics.

The amplitude of the transverse oscillation depends on the specifics of the electron's injection process, as this determines the initial displacement from the field axis and the initial transverse momentum. For self-injected beams, the injection is mostly from the edges of the bubble, where the initial transverse momentum is low but the displacement is large. For ionisation injection, it will depend on the width of the laser pulse that is above the ionisation threshold, and will be increased by above-threshold heating of the electrons by the laser. The oscillation amplitude will decrease as the electron accelerates, due to the relativistic mass increase of the particle.

2.8.1 Defining emittance

Each electron in a beam has a location in a 6D position-momentum space. As a beam, they occupy a hyper-volume in this phase space, and the magnitude of this volume is defined as the full 6D emittance, which is a value that quantifies the beam quality. It is a useful concept because, according to Liouville's theorem, it is invariant if the beam is only exposed to conservative forces. Examples of non-conservative forces that are relevant to LWFA are elastic scattering and direct laser acceleration, which will increase emittance, and betatron radiation, which reduces emittance⁶⁸. In general, a large emittance is bad for applications, as it reduces the minimum spot size to which the electron beam can be focused, and makes the beam more difficult to transport. Since reducing emittance is challenging, maintaining a low emittance is a priority of accelerator design. As discussed

⁶⁸ This is the purpose of cooling rings for conventional accelerators; electrons 'radiate off' their transverse energy.

in Sec. 2.4.6, the focusing forces in the bubble regime of plasma wakefield acceleration are linear, making this regime well suited to maintaining electron beam emittances.

Floettmann⁶⁹ identifies steps usually taken by physicists to make this 6D volume more tractable. The emittance is projected onto either the longitudinal axis or one of the two transverse axes, and in the transverse directions, the transverse momenta, p_x, p_y are replaced by the quantities $x' = p_x/p_z, y' = p_y/p_z$. In addition, since beams often have long-tailed distributions, the analysis is restricted to a high-density core using some width metric. This is usually the RMS, but for particularly long-tailed distributions a scaling that is more robust might be required, such as the FWHM, or median. It should be noted that only the full 6D distribution strictly obeys Liouville's theorem and that the projections are simply useful shorthands, only being conserved under specific conditions.

If the particle coordinates are known, the emittance can be calculated statistically. The trace space emittance is given by

$$\varepsilon_{tr} = \sqrt{\langle x^2 \rangle \langle x'^2 \rangle - \langle xx' \rangle^2}, \quad (2.89)$$

where the angled brackets denote the RMS. This is the trace space because the x' quantity is being used, as opposed to the transverse momentum. This emittance is considered ‘un-normalised’, because an increase in the longitudinal momentum will decrease the emittance, despite the transverse momentum spread remaining unchanged, a process known as adiabatic damping. The normalised trace emittance is defined as

$$\varepsilon_{tr,n} = \frac{\bar{p}_z}{m_e c} \sqrt{\langle x^2 \rangle \langle x'^2 \rangle - \langle xx' \rangle^2}, \quad (2.90)$$

where \bar{p}_z is the average longitudinal momentum. In contrast, the phase emittance, or geometric emittance, requires no normalisation,

$$\varepsilon_{ph,n} = \frac{1}{m_e c} \sqrt{\langle x^2 \rangle \langle p_x^2 \rangle - \langle xp_x \rangle^2}, \quad (2.91)$$

where the subscript n has been added here for emphasis. Throughout this thesis it will be assumed that it is the normalised emittance being referenced, unless it is either explicitly stated or irrelevant, such as when $\gamma = 1$. The normalised emittance can be more difficult to determine since the transverse momentum is seldom measured directly and so knowledge of the energy is also required so that it can be calculated from the divergence. The projected emittances can be seen as determinants of 2×2 matrices, where the diagonal entries are the area terms, $\langle x^2 \rangle$ and $\langle p_x^2 \rangle$, and the off-diagonal entries are the correlation terms, $\langle xp_x \rangle$. Extending this, the full 6D emittance can be calculated

⁶⁹ [171] K. Floettmann (2003) *Physical Review Special Topics - Accelerators and Beams*

using

$$\varepsilon^2 = \begin{vmatrix} \sigma_{xx} & \sigma_{xp_x} & \sigma_{xy} & \sigma_{xp_y} & \sigma_{xz} & \sigma_{xp_z} \\ \sigma_{p_x x} & \sigma_{p_x p_x} & \sigma_{p_x y} & \sigma_{p_x p_y} & \sigma_{p_x z} & \sigma_{p_x p_z} \\ \sigma_{yx} & \sigma_{yp_x} & \sigma_{yy} & \sigma_{yp_y} & \sigma_{yz} & \sigma_{yp_z} \\ \sigma_{p_y x} & \sigma_{p_y p_x} & \sigma_{p_y y} & \sigma_{p_y p_y} & \sigma_{p_y z} & \sigma_{p_y p_z} \\ \sigma_{zx} & \sigma_{zp_x} & \sigma_{zy} & \sigma_{zp_y} & \sigma_{zz} & \sigma_{zp_z} \\ \sigma_{p_z x} & \sigma_{p_z p_x} & \sigma_{p_z y} & \sigma_{p_z p_y} & \sigma_{p_z z} & \sigma_{p_z p_z} \end{vmatrix} \approx \begin{vmatrix} \sigma_{xx} & \sigma_{xp_x} & 0 & 0 & 0 & 0 \\ \sigma_{p_x x} & \sigma_{p_x p_x} & 0 & 0 & 0 & 0 \\ 0 & 0 & \sigma_{yy} & \sigma_{yp_y} & 0 & 0 \\ 0 & 0 & \sigma_{p_y y} & \sigma_{p_y p_y} & 0 & 0 \\ 0 & 0 & 0 & 0 & \sigma_{zz} & \sigma_{zp_z} \\ 0 & 0 & 0 & 0 & \sigma_{p_z z} & \sigma_{p_z p_z} \end{vmatrix}, \quad (2.92)$$

where $\sigma_{xx} = \langle x^2 \rangle$ and $\sigma_{xp_x} = \langle xp_x \rangle$. The approximation is the assumption that correlations between orthogonal coordinates are small, and the 6D emittance is simply the sum of the individual emittances in each direction.

2.8.2 Drift space emittance growth

In a drift space, there are no forces acting on the beam, and consequently, the full emittance of the beam is conserved. However, if there is any energy spread in the beam, then the phase emittance, ε_{ph} , will grow as the beam spreads out. The magnitude of this effect can be calculated by substituting $x(z) = x_0 + zx'$ into the definitions of $\langle x^2 \rangle$ and $\langle xp_x \rangle$,

$$\langle x^2 \rangle = \langle x_0^2 \rangle + 2z\langle x_0 x' \rangle + z^2\langle x'^2 \rangle, \quad (2.93)$$

$$\langle xp_x \rangle = \langle x_0 p_x \rangle^2 + 2z\langle x_0 p_x \rangle \langle x' p_x \rangle + z^2\langle x' p_x \rangle^2. \quad (2.94)$$

Then, by substituting into 2.91, we derive the phase emittance as a function of drift length,

$$\varepsilon_{ph}(z)^2 = \varepsilon_{ph}(0)^2 + Az + Bz^2, \quad (2.95)$$

where

$$\begin{aligned} A &= \langle x_0 x' \rangle \langle p_x^2 \rangle - \langle x_0 p_x \rangle \langle x' p_x \rangle, \\ B &= \langle x'^2 \rangle \langle p_x^2 \rangle - \langle x' p_x \rangle^2. \end{aligned} \quad (2.96)$$

If the energy spread is small, $1/p_z$ can be factored out of the expressions for A and B , resulting in $A = B = 0$.

Although this growth only affects the phase space emittance, it manifests itself in the trace space emittance if the beam is recollimated by a beam optic, such as a quadrupole. This is due to the γ dependence of a magnetic field's effect on each electron's trajectory. The result is that, upon collimation, a single point in trace space before the optic will map to multiple points after the optic, causing the trace space emittance to increase. When the beam is collimated the trace space and phase space emittances will be equal. The consequence of this is that the trace space emittance measurements of the electron beam, such as the pepper-pot style measurements presented by Sears *et al.*⁷⁰, while valuable

⁷⁰ [81] C. M. S. Sears *et al.* (2010) *Physical Review Special Topics - Accelerators and Beams*

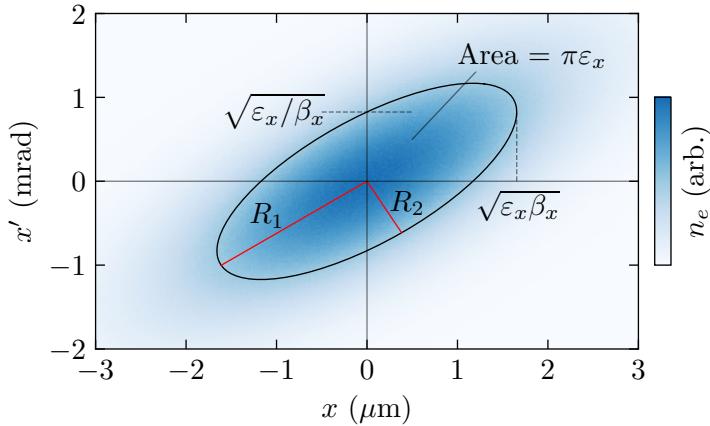


Figure 2.17: Phase space plot of a Gaussian beam. An ellipse with radii R_1 , R_2 has been fit to the RMS contour of the beam. In this plot, the beam has acquired a positive correlation after expanding from focus, where it would have had zero correlation.

for providing information about the emittance of the beam inside the channel, cannot measure the effect of the drift space on the electron beam. To obtain this information, the emittance measurement must be made after a quadrupole, or some other focusing element, where the two spaces are equivalent. Alternatively, the trace emittance of each energy slice of the beam can be measured independently, as is done in chapter 5.

2.8.3 Beam envelope dynamics

The beta function relates the beam size to the emittance at longitudinal position s in the accelerator,

$$\beta_x(s) \equiv \frac{\langle x^2 \rangle}{\varepsilon_x}. \quad (2.97)$$

The physical interpretation of the beta function is shown in Fig. 2.17, with an ellipse fit to the RMS contour of the beam in $x - x'$ phase space. The area of the ellipse is equal to the emittance times a factor of π ; sometimes the units of emittance are given in $\pi \text{ mm mrad}$, and under the action of conservative forces, the area of the ellipse will remain constant, causing the ellipse to rotate, stretch and compress as the beam propagates. For example, when the beam is at focus, the ellipse will be upright, while after some drift space, an initially focused beam will acquire a positive correlation between x and x' , as shown in the Fig. 2.17.

Since the focusing forces depend on the location of the beam in the accelerator, we can recast Eq. 2.88 to depend on s ,

$$\frac{d^2x}{ds^2} + k_\beta^2 x = 0, \quad (2.98)$$

where $k_\beta = \omega_\beta/c$, and we have made the coordinate transformation $s = ct$. This is Hill's equation⁷¹, from which it can be shown that the beta function is given by the solution to

$$\frac{1}{2}\beta'' + k_\beta^2\beta - \frac{1}{\beta} \left[1 + \frac{\beta'^2}{4} \right] = 0. \quad (2.99)$$

This equation often needs to be solved numerically for arbitrary density profiles. For a constant plasma density, the beta function evolves according to

$$\beta = \beta_0 \cos^2(k_\beta s) + \frac{1}{\beta_0 k_\beta^2} \sin^2(k_\beta s), \quad (2.100)$$

where β_0 is the initial condition. This is an oscillating solution unless $\beta_0 = k_\beta^{-1} \equiv \beta_m$, which defines the matched beam size. If an electron's beam size is to remain constant in an accelerating column, it must be injected with the energy-dependent matched beam size. Oscillations in the beta function are undesirable because they produce emittance growth when they couple to energy spread, since different energy slices rotate at different rates in phase space.

Solutions to Eq. 2.100 are plotted in Fig. 2.18 for a tapered density profile using $\gamma = 300$ and $\gamma = 3000$. An initially mismatched solution oscillates around the matched solution, with an oscillation wavelength that scales with the longitudinal momentum. As the density drops, the beta function and oscillation wavelength increase.

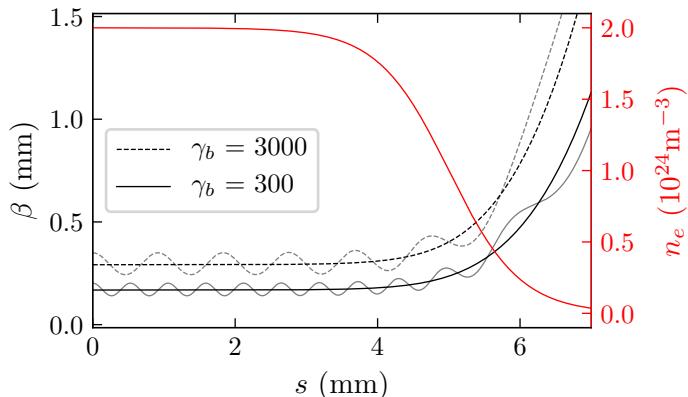


Figure 2.18: Beta function solutions for different Lorentz factors. A matched solution, $\beta_0 = \beta_m$, and a mismatched solution, $\beta_0 = 1.2\beta_m$ has been plotted at each γ . The effect of a decreasing density is shown.

2.8.4 Betatron radiation

As a particle oscillates, it undergoes continual transverse acceleration and emits radiation as a result. The longitudinal acceleration also produces radiation, but the field's strength is a factor of $1/\gamma$ weaker. Consequently, the radiation at any given moment is effectively the same as would be emitted by a particle travelling along the arc segment of a circle

⁷¹ [172] S.-Y. Lee *Accelerator Physics (Fourth Edition)* (2019)

with an equivalent radius of curvature⁷². Due to the particle's relativistic motion, the radiation is collimated by a factor of γ along its direction of motion, such that for highly relativistic particles, the radiation is only 'seen' by an observer when the particle is travelling directly towards them. In the particle's reference frame, the radiation column strafes over the observer in time $\Delta t' \sim \rho/c\gamma$, where ρ is the instantaneous radius of curvature. In the observer's reference frame, this is a factor of γ^2 shorter, so the pulse length is $\Delta t \sim \rho/c\gamma^3$. From Fourier theory, the radiation spectrum must extend up to a critical frequency of $\omega_c \sim \Delta t^{-1} \sim c\gamma^3/\rho$.

The value for the radius of curvature, ρ , depends on the frequency ω_β , and the radius of oscillation, r_β . The radius of oscillation for individual electrons will depend on their specific injection trajectories, so it is more useful to consider the average radius of the beam for a given emittance, ε_n . If the beam is matched to the channel, then the radius will be

$$r_\beta = \sqrt{\frac{\varepsilon_n}{\gamma k_\beta}} = \left(\frac{2\varepsilon_n^2 c^2}{\gamma \omega_p^2} \right)^{\frac{1}{4}}. \quad (2.101)$$

For a typical LWFA emittance of 10 mm mrad, a 1 GeV beam in a $1 \times 10^{24} \text{ m}^{-3}$ plasma will have an oscillation radius of 1.3 μm . For a sinusoidal oscillation, the minimum radius of curvature is

$$\rho = \frac{v^2}{\dot{v}_\perp} = \frac{v^2}{c^2 k_\beta^2 r_\beta} \simeq \frac{\lambda_\beta^2}{4\pi^2 r_\beta}, \quad (2.102)$$

where the approximation comes from taking $v \simeq c$. For our example, $\rho \simeq 0.07 \text{ m}$, which gives a critical frequency of $2.2 \times 10^{19} \text{ s}^{-1}$. This means that significant numbers of photons with energies up to the critical energy of $E_c = 19 \text{ keV}$ should be produced, which are in the hard x-ray part of the spectrum. For the 100 MeV beam in the $1 \times 10^{25} \text{ m}^{-3}$ plasma, the critical energy is an order of magnitude lower, at 1.9 keV.

For N_β electron oscillations in a plasma channel, the intensity spectrum can be expressed in terms of the critical frequency,

$$\frac{d^2 I}{dE d\Omega} \simeq \frac{6N_\beta e^2}{\pi^2 c} \frac{\gamma^2 \zeta^2}{1 + \gamma^2 \theta^2} \left[\frac{\gamma^2 \theta^2}{1 + \gamma^2 \theta^2} K_{1/3}^2(\zeta) + K_{2/3}^2(\zeta) \right], \quad (2.103)$$

where

$$\zeta = \frac{E}{E_c} (1 + \gamma^2 \theta^2)^{\frac{3}{2}}, \quad (2.104)$$

and

$$E_c = \hbar \omega_c = 3\hbar \gamma^3 r_\beta \omega_\beta^2 / c, \quad (2.105)$$

and the functions K_i are modified Bessel functions of the second kind. Eq. 2.103 is plotted in Fig. 2.19 for a 100 MeV electron oscillating in an ion channel with a density of $1 \times 10^{25} \text{ m}^{-3}$. The radiation intensity is strongest on-axis and slightly below the critical energy. The intensity falls at increasing observation angles, as does the critical energy.

⁷² [173] J. D. Jackson *Classical electrodynamics* (1999)

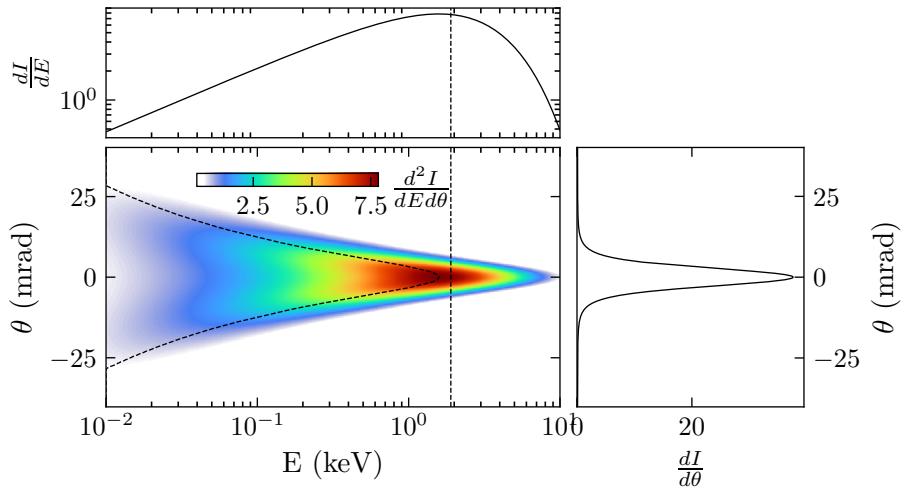


Figure 2.19: X-ray spectrum with a critical energy of 1.9 keV, corresponding to a 100 MeV electron performing a single oscillation in an ion channel with a density of $1 \times 10^{25} \text{ m}^{-3}$. The dotted vertical line shows the critical energy, while the curved dashed line is the maximum intensity contour. The spectrum in the top plot is the on-axis spectrum, while the transverse profile in the right plot is integrated across all energies.

3 Methods

IN THIS CHAPTER, the experimental and numerical methods used in this thesis will be presented. We will follow a similar route through these methods as the route taken through the theory section, starting with the laser diagnostics. Following on from this, we will explain the target configuration and the interferometric methods used for density measurements. The electron beam diagnostics will be discussed, covering the detection of electrons, the measurement of the energy spectrum, and the pepper-pot emittance measurement technique. Finally, due to their extensive use in this thesis, the mechanism behind particle-in-cell simulations will be explained, and various convergence tests will be performed to identify appropriate simulation resolutions.

3.1 Laser diagnostics

The experimental results were obtained using the Gemini laser at the Central Laser Facility (CLF); part of the Rutherford Appleton Laboratories in Oxfordshire. The Gemini laser simultaneously provides 800 nm laser pulses to two target areas, named Astra Target Area 2 and 3 (ATA2/3)¹, with each target area serving a different purpose. Nominally, the laser provided to ATA2 has an energy of 500 mJ, with a pulse duration of 40 fs, at a repetition rate of 5 Hz, while the laser provided to ATA3 is split into two arms, which each deliver 15 J in 30 fs, at a repetition rate of 0.05 Hz. The ATA2 and ATA3 beam diameters are 55 mm and 150 mm respectively. More detailed descriptions of the laser system are given here². The nominal values of the laser system rarely equal those delivered to the target areas during experiments, and consequently, it is necessary to closely monitor the delivered laser energy and pulse duration on a shot-to-shot basis, both for the operation of the experiment and for the analysis of results.

To characterise the laser at focus for each shot, measurements of the energy and temporal duration need to be made, which determines the power. The spatial distribution of the power at the focal plane can then be determined by characterising the focal spot.

¹ The ‘Astra’ refers to the front-end of the Gemini laser, which was previously a laser system in its own right, capable of achieving intensities of 10^{23} W m^{-2} .

² [174] C. J. Hooker *et al.* (2006) *Journal de Physique IV*
[175] K. Poder *PhD. Thesis* (2016)

3.1.1 Power measurements

In ATA2, the shot-to-shot variations in the laser energy are monitored using the residual laser energy that leaks through one of the dielectric mirrors close to the target. The leakage is de-magnified, before being directed onto a camera's CCD, and the integrated number of counts on the camera correlates linearly with the laser energy. The calibration between the number of counts and the laser energy is performed by taking measurements of the laser energy at the target using a Gentec energy meter, which is capable of measuring the absolute energy in the laser. Calibrating this number with the leakage CCD counts accounts for the losses that occur during the beam transport to the target. In ATA3, a similar procedure is performed, except that the measurement is provided by the CLF laser operators, and it is made before compression of the pulse. The compression efficiency is periodically measured and is usually approximately 70 %. This data is then made available to users online³.

The pulse duration is more challenging to measure. Since the laser pulse is on the femtosecond timescale, the only readily available reference with which to measure it is the laser itself. This is the principle behind the autocorrelator, which is a common method for measuring the pulse length due to its relative technical simplicity. However, this only gives intensity information, which is ambiguous with respect to the carrier envelope phase and the direction of time. To fully describe the pulse shape, the temporal phase is also required. Two common methods exist for measuring this, the Frequency Resolved Optical Gate (FROG), and Spectral Interferometry for Direct Electric field Reconstruction (SPIDER). The FROG is used by the CLF operators to provide the pulse length information about the ATA3 lasers to the users, while a SPIDER was employed by the users to measure the post-compression pulse length in ATA2. Particular care must be taken when making pulse length measurements with respect to dielectric mirrors, as these can alter the pulse spectrum. One must also take care to account for transmission through dielectric materials since these can alter the pulse length via group velocity dispersion.

3.1.2 Focal spot characterisation

The first step in characterising the focal spot is the measurement of a satisfactory focal spot image. The focal spot should be imaged exactly at its focal plane, and with an appropriate magnification. The focal plane can be determined by monitoring the size of the focal spot as the position of the camera is scanned. For weak focusing geometries, the Rayleigh range of the laser can become comparable to the travel of the focal spot camera, and in such instances, it is best to use quantitative measurements of the beam size as a function of camera position to find the focal plane. The appropriate magnification is a balancing act since the magnification needs to be large enough that the contours of the focal spot can be resolved accurately, but not so large that small amounts of energy in the wings of the pulse are clipped. This will ultimately depend on the dynamic range

³ [176] *eCat2*, Central Laser Facility

and resolution of the camera, as well as the focal spot quality. Of similar importance is to attenuate the laser to a level that uses the full dynamic range of the camera, without causing saturation. Spatial calibration of the camera must also be performed. While it is possible to use the magnification of the imaging system, this can be subject to errors associated with the manufacturing tolerances of the objective. It is often simpler and more reliable to use a physical reference, and a reliable one is the diffraction pattern formed by a grating. For a transmission grating with spacing d placed into the near-field of the laser, a series of focal spots will be produced in the far field with a separation $x \approx f\lambda/d$, where f is the focal length, and λ is the laser wavelength. Finally, it is important to perform a pixel-wise background subtraction of a dark-field image. This image should be created by taking the pixel-wise average of a large number of images with the camera while the laser is blocked.

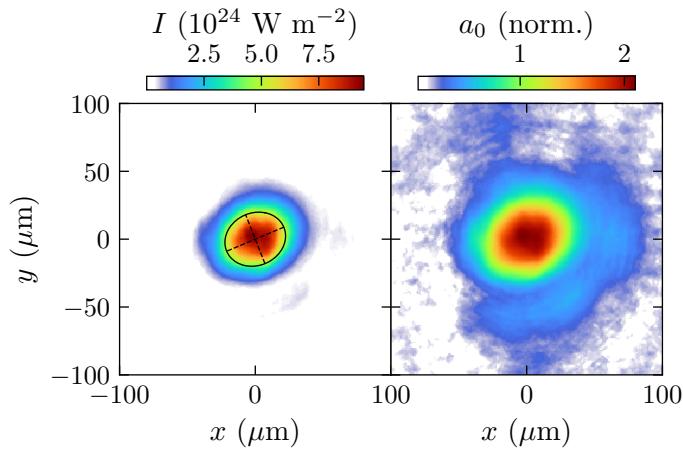


Figure 3.1: Focal spot image (left) and the associated field intensity (right). An ellipse fitted to the FWHM contour has been plotted. The dashed lines indicate the minor and major ellipse axes.

After recording a focal spot image, it is helpful to extract a measure of the beam size and the maximum intensity. The beam size is most often quantified by fitting an ellipse to the full width at half maximum (FWHM) contour, which is a useful measurement for noisy, irregular patterns on account of its insensitivity to long tails in the measurement. The focal spot size can then be quoted in terms of the long (major) and short (minor) axes of this ellipse, which gives a measure of the spot size and its eccentricity. The intensity at each pixel can be determined by stipulating that the power of the laser is shared equally between all the pixel counts in the image. This analysis procedure has been performed for the image in the left panel of Fig. 3.1, using an example focal spot image taken from the experiment described in chapter 6. The conversion from pixel counts to intensity has used a laser energy of 10.9 J and a FWHM temporal envelope width of 45 fs. The right panel is the normalised field strength of the image, proportional to \sqrt{I} , and consequently, the dynamic range of the image is increased.

The intensity and field strength profiles as a function of radius are plotted in Fig. 3.2. This highlights the difference in widths between the two profiles. The laser transverse

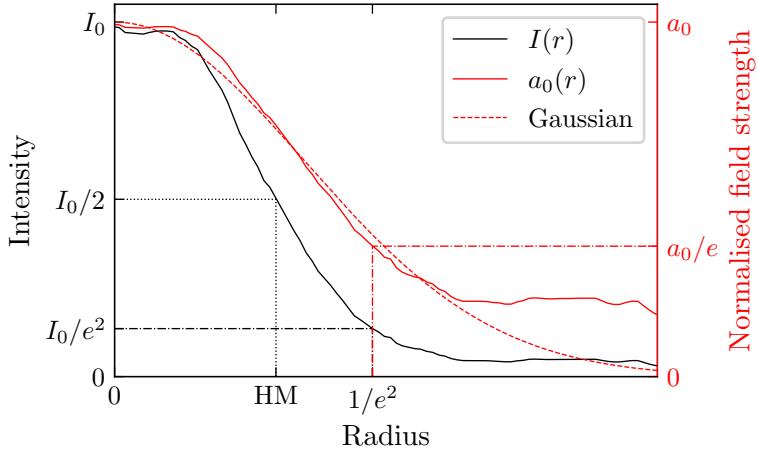


Figure 3.2: Identification of focal spot transverse profile metrics.

profile is often modelled as a Gaussian with a width equal to the $1/e^2$ width of the intensity profile, and Fig. 3.2 shows that using this approximation models the true a_0 profile quite well within the central spot. However, energy in the wings of the spot will not be modelled with this approximation, and some of the consequences of using this approximation will be shown in Sec. 3.5.5.

For reference, it is worth explicitly stating the conversions between the various metrics. A Gaussian field envelope in both space and time has the form

$$E(r, t) = E_0 \exp \left[-\left(\frac{t}{\tau_0} \right)^2 - \left(\frac{r}{w_0} \right)^2 \right], \quad (3.1)$$

where τ_0 and w_0 are defined as the Gaussian widths of the field. The intensity profile, using the same widths, is

$$I(r, t) = I_0 \exp \left[-2 \left(\frac{t}{\tau_0} \right)^2 - 2 \left(\frac{r}{w_0} \right)^2 \right], \quad (3.2)$$

where we are refraining from identifying intensity Gaussian widths to avoid confusion. The two important relationships to identify are those between the intensity focal spot FWHM and the Gaussian width, and the intensity temporal FWHM and the peak intensity. The width relationship can be found by setting $I(r_{\text{HM}}, 0) = 0.5$, giving

$$w_{\text{FWHM}} = \sqrt{2 \ln 2} w_0 \approx 1.18 w_0, \quad (3.3)$$

where $w_{\text{FWHM}} = 2r_{\text{HM}}$, and r_{HM} is the radius at half maximum. The energy in the laser, \mathcal{E}_L , is equal to the 3D integral of the intensity,

$$\mathcal{E}_L = I_0 \int_{-\infty}^{+\infty} I(t') dt' \int_{-\infty}^{+\infty} I(r') r' dr' \int_{-\pi}^{\pi} d\phi' \quad (3.4)$$

$$= \left(\frac{\pi}{2} \right)^{\frac{3}{2}} \tau_0 w_0^2 I_0. \quad (3.5)$$

This expresses the laser energy in terms of the field parameters. In terms of the FWHM measurements, this becomes

$$I_0 = 8 \left(\frac{\ln 2}{\pi} \right)^{\frac{3}{2}} \frac{\mathcal{E}_L}{\tau_{\text{FWHM}} w_{\text{FWHM}}^2}. \quad (3.6)$$

This assumes that the peak of the electric field occurs at the peak of the field envelope, which is approximately true for $c\tau_0 \gg \lambda_0$, where λ_0 is the carrier wave wavelength.

3.2 Gas targets

All the experiments in this thesis use neutral gas as the target. This gas needs to be positioned at the laser focus when the pulse arrives, and control over the density and longitudinal profile is required. In this section, the gas delivery systems are discussed, and fluid simulations are used to model the density variations near apertures.

3.2.1 Target types

Gas jets are a common target in LWFA experiments since they are relatively simple to set up, are easy to diagnose, and can provide the densities and longitudinal profiles that are required. Usually, a longitudinal profile that is approximately constant is desired, and this is achieved by using a converging-diverging (de Laval) nozzle design⁴. The gas flow is initially subsonic in the converging region, accelerating to Mach 1 in the nozzle throat, and then accelerating again to supersonic speeds as it expands in the diverging section⁵. The ratio between the throat diameter and the final nozzle diameter determines the speed of the flow, with faster flows producing flatter density profiles under ideal conditions. The trade-off is that a larger backing pressure will be required to achieve the same density at a higher flow velocity. Gas bottles, which are typically pressurised to 100 bar, can support supersonic flows in the density range of $10^{23} - 10^{25} \text{ m}^{-3}$, which is ideal for LWFA experiments.

In real gas jets, the effects of viscosity and nozzle imperfections need to be accounted for, which can have substantial effects on the density profile. For example, the boundary layer that forms between the flow and the nozzle wall can continue to grow into the flow outside the nozzle⁶. In addition, shocks can form at nucleation points inside the nozzle, producing spikes at the edges of the density profile. Micron-scale density ripples have also been inferred through measurements of local variations in the plasma wave velocity, seeding injection and affecting electron beam quality⁷. Consequently, while gas jets are convenient and robust, they may be poorly suited to applications where fine control of the density profile and the injection process are required.

⁴ [177] S. Semushin *et al.* (2001) *Review of Scientific Instruments*

⁵ [178] D. B. Atkinson *et al.* (1995) *Review of Scientific Instruments*

⁶ [179] N. Lemos *et al.* (2009) *Review of Scientific Instruments*

⁷ [180] S. Kuschel *et al.* (2018) *Physical Review Letters*

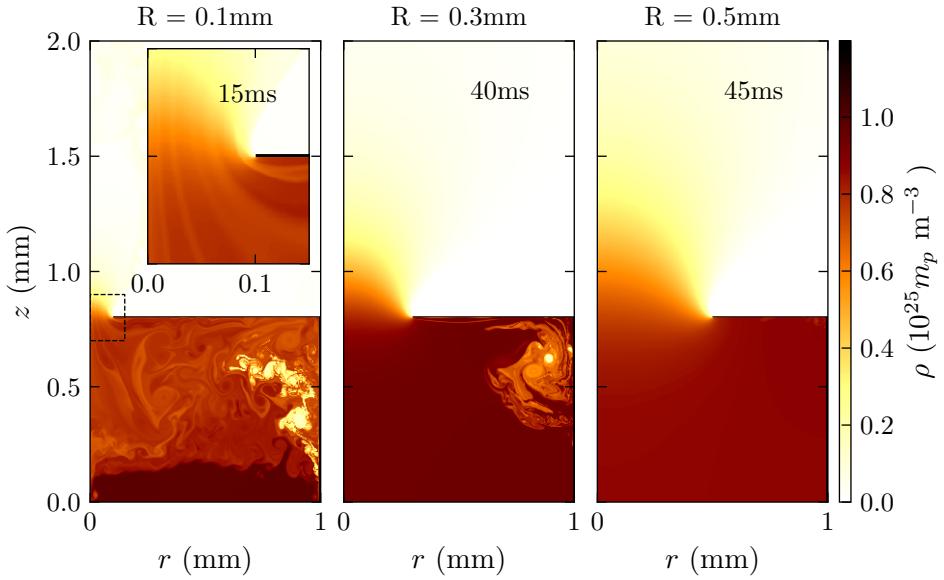


Figure 3.3: Simulations of neutral gas density around gas cell apertures. Each panel shows a simulation with a different aperture radius. These simulations are axisymmetric around $r = 0$, with the gas coming in at the $z = 0$ plane. The snapshots are at different times; 15 ms, 40 ms, and 45 ms moving from left to right, showing the equilibration of the density over time.

Gas cells are another option for LWFA experiments. These are small gas boxes that have entry and exit holes through which the laser propagates. These holes are typically on the 100 μm scale, and they can be subject to laser damage if the laser is misaligned or insufficiently focused. The cell is filled approximately 100 ms prior to the laser's arrival, which is enough time for the density to reach a smooth equilibrium. Cells are naturally suited to producing a constant density profile, and they can also be designed to have a variable length, which provides a useful tool for optimising the interaction, as well as performing parameter scans.

3.2.2 Simulation of cell apertures

One of the disadvantages of gas cells is that the cell casing can obstruct the imaging of the plasma. This is particularly the case at the laser exit and entry apertures, which is where the most variation in the density occurs. To gain insight into the density profile through the apertures, fluid dynamic simulations have been performed using the FLASH code⁸. This code employs adaptive mesh refinement to allow for efficient computation of a large range of length scales. The simulations were performed in 2D cylindrical geometry, where the axis of symmetry was the centre of the aperture. The simulation domain was 2×1 mm in the $z - r$ plane, with the $z = 0$ boundary defined as the inlet with a pressure of 100 mbar, and the $z = 2$ mm and the $r = 1$ mm boundaries defined as outflows. The cell was simulated using a 25 μm thick boundary that extended up the right edge of the domain to $z = 0.8$ mm, and then across to $z = R$, where R is

⁸ [181] B. Fryxell *et al.* (2000) *The Astrophysical Journal Supplement Series*

the aperture radius and the parameter under study. Such a thin boundary was used to simulate the boundary formed by the tape in the target used in chapter 6.

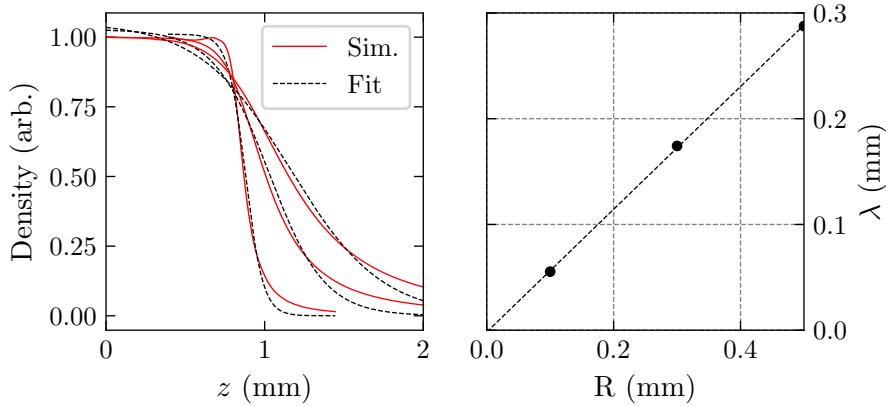


Figure 3.4: Left panel shows exponential fits to on-axis density lineouts. The right panel shows the fitting parameter λ as a function of aperture radius. The dotted black line in the right panel goes through the origin without being forced.

Fig. 3.3 shows outputs from the aperture radii studies, at different times in the simulation. At early times the flow inside the cell is turbulent, as the flow reflects off the cell boundaries and mixes with inflowing gas. Over time, the density stabilises, first at the apertures, and then in the corner of the cell. After approximately 50 ms steady state is achieved throughout the simulation window. The density reached a value of approximately 10^{25} m^{-3} inside the cell, which is similar to the backing pressures used in experiments, and this value was higher for the narrower apertures, as would be expected. The on-axis density profiles are plotted in the left panel of Fig. 3.4. To these profiles, an error function is fit, where the analytic form $\rho = (e^{-\rho/\lambda} + 1)^{-1}$ has been used, where λ is the fitting parameter and characterises the scale length of the density change. Fitting this function to the density profiles gives a linear relationship between λ and R , which can be seen in the right figure of the plot. Approximately, $\lambda \approx R/2$, a scaling that will be used to characterise the density ramps in this thesis.

3.3 Interferometry

Interferometry is a general method that uses the interference between two waves to extract information, and in LWFA experiments, the technique is primarily used for measuring the density profile of the target. The refractive index of the plasma is lower than the vacuum, so light that transmits through the plasma will acquire a relative phase shift. The path length difference of a ray of light propagating through a material relates to the line integral of the refractive index along its path,

$$\lambda_0 \Delta\phi(x, y) = 2\pi \int_A^B [\eta(x, y, z') - 1] dz', \quad (3.7)$$

where the path $A - B$ is the path the light takes through the material with refractive index η . If we assume that the path is straight, that is to say, the deflection of the light rays is negligible, then for a constant density of n_e , the phase shift is

$$\Delta\phi(x, y) = \frac{\lambda_0 e^2}{4\pi c^2 m_e \epsilon_0} n_e L. \quad (3.8)$$

For densities of $1 \times 10^{24} \text{ m}^{-3}$ and plasmas that are $100 \mu\text{m}$ thick, this produces a path length difference on the order of λ_0 , and so the laser wavelength is an appropriate reference to use for the measurement. The interference pattern is generated by overlaying two coherent beams with a slight pointing difference. Perfectly co-linear beams result in no interference pattern in the transverse plane, as the phase difference between the beams is constant everywhere. Only by introducing an angle between the beams, $\Delta\theta$, does an interference pattern emerge, as this is equivalent to having a linearly increasing phase difference across the entire beam. Fig. 3.5 illustrates this effect 1D, and how a phase object can perturb the interference and cause a fringe shift.

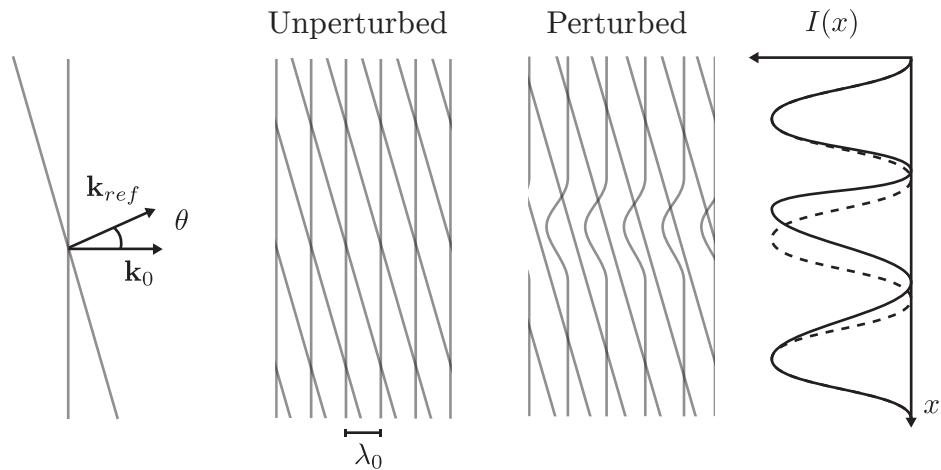


Figure 3.5: Illustration of the principle of interferometry with angled beams.

The distortion of the wavefronts changes the location of the constructive interference between the two waves, which is measured as a phase shift at the detector. In this way, the spatial variation in the phase of the carrier wave produces spatial variations in the intensity of the interference pattern. In our 1D example, the intensity of the pattern is

$$I(x) \propto 1 + \alpha \cos(k_0 x \sin(\theta) + \Delta\phi(x)), \quad (3.9)$$

where α is a factor that relates to the contrast of the fringes. The $k_0 \sin(\theta)$ term is the wavenumber of the interference pattern, arising geometrically. By subtracting the oscillating term, such as by comparing the data to a reference interferogram without a phase object, the phase shift $\Delta\phi$ can be retrieved.

3.3.1 Implementation

Typically, we are interested in the longitudinal density profile of the plasma, and so a probe is set up so that it propagates through the target orthogonally to the drive beam. For our purposes the probe can be the same frequency as the main beam since any $1\omega_L$ light scattered by the plasma is much weaker than the probe beam. The intensity of the probe varies between experiments, but it is normally approximately two orders of magnitude less energetic than the main beam. Once the laser has formed the plasma, it will begin to evolve hydrodynamically; the electrons heat the ions and the whole plasma begins to expand. This expansion reduces the density, and it is therefore important to probe the target immediately after the drive laser completes its interaction with the plasma, so as to accurately measure the plasma density at the moment of interaction. In all experiments in this thesis, the probe beams were obtained from the same laser as the drive beams, synchronising the pulses. To fine-tune the relative delay, a translation stage in the probe beamline was used to provide a variable path length, and modifications were made until the probe beam imaged the plasma immediately after the drive beam had completed its transit.

The design of the interferometer needs to be considered. The plasma is typically long and thin, whereas the probe beam is circular, meaning there is typically a large portion of the beam that propagates through vacuum or neutral gas, undergoing almost no phase shift. Consequently, it is often not necessary to generate the reference beam before the probe interacts with the plasma, as this ‘unused’ portion of the beam can be repurposed as the reference. The advantage of setting up the interferometer in this way is that most of the apparatus can be placed away from the target, facilitating easy access, and liberating space near the target.

The interferometer uses a 50:50 beamsplitter to generate two copies of the probe beam. To produce the desired interferogram, these two copies need to be overlapped on a camera’s sensor with the pointing difference required to produce the fringes, and the positional offset required to overlap the unused portion of the beam with the phase object. In addition, the short pulse length necessitates precise placement of the mirrors; for a 50 fs pulse width, the path lengths of the two arms need to be within $\sim 10 \mu\text{m}$ of each other. These requirements make the geometry of mirrors important, as different geometries offer different degrees of freedom over the interference pattern. For example, a Michelson interferometer, which utilises a central beamsplitting cube to both split and recombine the beams, couples the spatial offset and the pointing angle between the beams, whereas a Mach-Zender (MZ) geometry does not, allowing for independent control of the offset direction and the fringe orientation. The extra control of the MZ may be required when attempting to measure phase shifts near a boundary, such as at the aperture of a gas cell. It is worth noting that the conventionally configured MZ changes the timing between the two arms by changing the pointing in the individual arms, not through the movement of the translation stage, as might be naively assumed⁹.

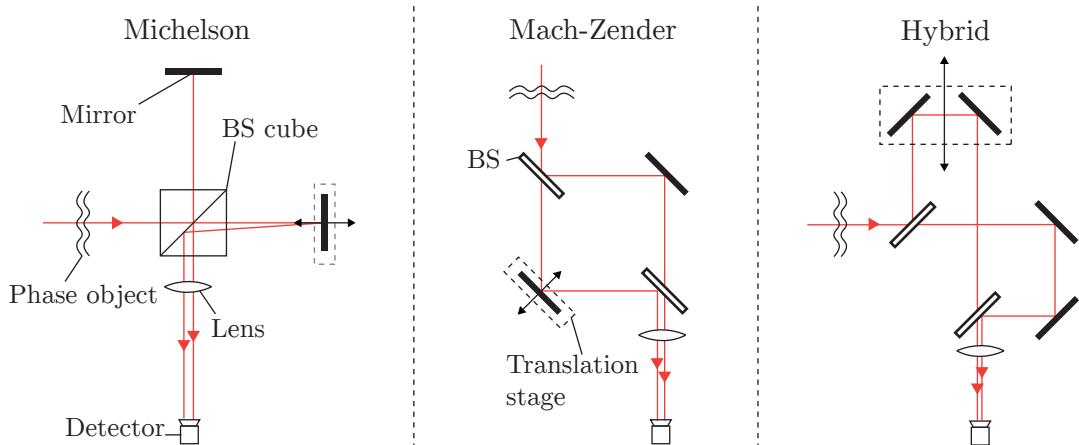


Figure 3.6: Diagrams of Michelson, Mach-Zender, and hybrid interferometers.

This makes achieving temporal overlap more challenging with a Mach-Zender since the alignment needs to be constantly corrected as the pointing offset is scanned in search of temporal overlap. However, there is a modified MZ geometry that achieves the degrees of freedom of the MZ and the ease of timing of the Michelson, at the cost of requiring a slightly more complicated setup. These different arrangements are shown in Fig. 3.6.

To achieve a high-quality measurement, it is important to tune the interferometer to simultaneously maximise spatial resolution and fringe contrast. Fringe contrast determines the phase resolution because small changes in phase are determined not by fringe wavefront movements, but by subtle changes in intensity across the interferogram. Interferometry is therefore surprisingly sensitive; an interferogram with no visible fringe shift can still measure changes in density if the image noise is low enough. Spatial resolution is primarily determined by the imaging system, although fringe spacing can play a role when the image is filtered according to spatial frequency, with thinner fringes allowing for the selection of higher spatial frequencies while retaining a narrow bandwidth¹⁰. Thin fringes result from a larger angle between the two arms of the interferometer, which, for short laser pulses, can result in a spatially varying fringe contrast. This is because the pulses do not overlap temporally across the entire beam, resulting in a region of high contrast bordered by two regions of diminishing contrast. This effect can be mitigated to some extent by using a narrow bandpass filter to reduce the bandwidth of the probe pulse, increasing the pulse length and hence the region of high-contrast fringes.

3.3.2 Density extraction

Mathematically, the interferogram has the form

$$I(x, y) = I_{\text{bg}}(x, y) + \frac{1}{2} I_{\text{plasma}}(x, y) \left(e^{i(k_f x + \Delta\phi(x, y))} + e^{-i(k_f x + \Delta\phi(x, y))} \right), \quad (3.10)$$

⁹ A translation stage is still required in the MZ setup to allow for a large range of pointing angles to be scanned in search of the correct timing.

¹⁰ See Sec. 3.3.2 and Fig. 3.7.

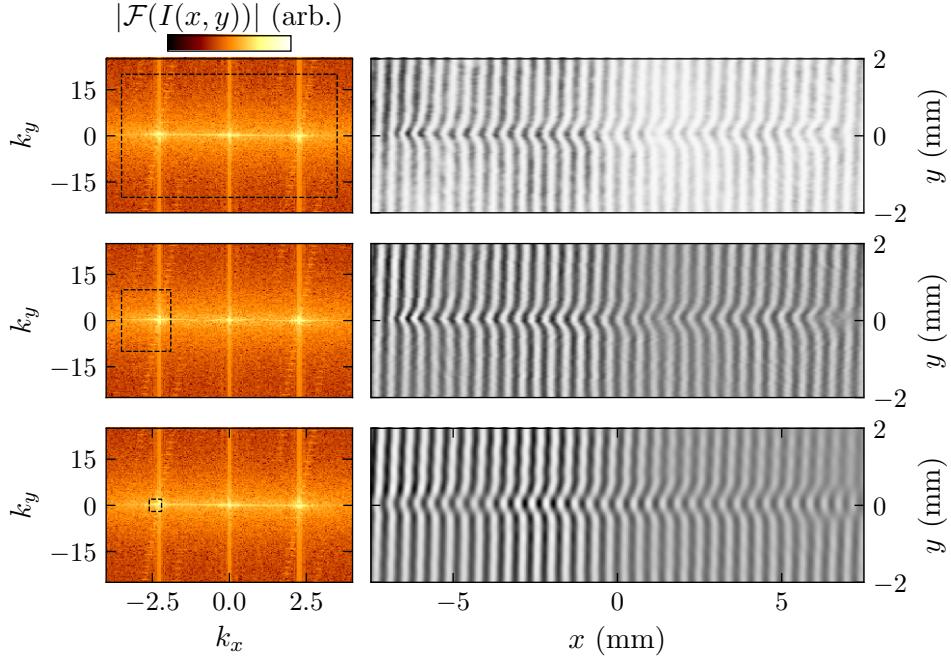


Figure 3.7: The Fourier filtering of interferograms, where the phase object producing the fringe shift is a laser-ionised plasma in a gas jet. The left panels show the absolute value of the Fourier transform of the image, while the right panels show the real components of the respective inverse Fourier transforms, after masking the indicated regions. The frequency space units are mm^{-1} .

where I_{bg} and I_{plasma} are the intensity variations due to the background and plasma respectively, k_f is the fringe spatial frequency, and $\Delta\phi$ is the 2D phase map to be retrieved. We want to remove the noise from this image that is due to the background, and then convert the argument $i\Delta\phi(x, y)$ to an intensity map. This will then be converted to a 2D density profile using an Abel inversion.

The Fourier transform of an example interferogram is shown in the left column of Fig. 3.7. The hot spots at $k_x = \pm 2 \text{ mm}^{-1}$ correspond to the sinusoidal variation in the intensity due to the fringes, which have a wavelength of approximately 0.5 mm. The fringes have the effect of making copies of the image in the Fourier plane at $\pm k_f$,

$$\mathcal{F}(I(x, y)) = \mathcal{F}(I_{\text{bg}}) + \frac{F(k_x + k_f, k_y) + F^*(k_x - k_f, k_y)}{2}, \quad (3.11)$$

where \mathcal{F} denotes the Fourier transform, and

$$F(k_x + k_f, k_y) = \mathcal{F}(I_{\text{plasma}}(x, y)) e^{i\Delta\phi(x, y)}, \quad (3.12)$$

with F^* as the complex conjugate of F . By setting the value of the Fourier transform to 0 everywhere except the region surrounding either $(k_x + k_f, k_y)$ or $(k_x - k_f, k_y)$, most of the background intensity variations are removed, as is the high-frequency part of the noise spectrum. This is shown by the differences between the filtered interferograms in the top

two panels of the right column of Fig. 3.7. Care must be taken so that an appropriately sized region is masked, and there is an inherent trade-off between noise and resolution. If the Fourier plane is cropped too aggressively in the k_y direction, such that the range of spatial frequencies selected is too narrow, then the plasma channel will appear larger due to the loss of high-frequency spatial information, affecting the measurement of the channel radius, and consequently the density. In addition, if the plane is cropped too tightly in the k_x direction, then the sharp change in density at the start and end of the plasma will not be resolved. An example of this loss of resolution can be seen when comparing the second two rows in Fig. 3.7, with the last row being an example of cropping the Fourier plane too tightly and losing important spatial information.

After filtering around the lobe at $k_x = -f_f$, the interferogram now has the form $I(x, y) = I_{\text{plasma}}(x, y) \exp(-i(k_f x + \Delta\phi(x, y)))$. The phase map is then retrieved by dividing by a reference interferogram, to which the exact same filtering has been applied, yielding $I_{\text{ref}}(x, y) \exp(-ik_f x)$. The phase can then be retrieved by dividing by the reference set of fringes and then taking the argument,

$$\Delta\phi(x, y) = \arg\left(\frac{I_{\text{plasma}}(x, y) e^{-i(k_f x + \Delta\phi(x, y))}}{I_{\text{ref}}(x, y) e^{-ik_f x}}\right). \quad (3.13)$$

Using a reference image for this purpose, rather than an artificially generated sinusoidal signal, allows for the warping of the fringes due to imaging system imperfections to be corrected.

The Abel inversion can be used to convert from the phase map to the density profile. In our application, it uses the assumption of cylindrical symmetry and the phase as a function of distance from the axis to determine the density at each point. The density at longitudinal position x is given by

$$n_e(x, r) = \lambda_0 r_e \int_r^\infty \frac{d\Delta\phi(x, y)}{dy} \frac{1}{\sqrt{y^2 - r^2}} dy. \quad (3.14)$$

where r_e is the classical electron radius. In principle, this equation can be used to turn a phase map into a 2D density profile. In practice, brute force calculation of this function is both slow, scaling $\mathcal{O}(N^3)$ with the image size N , and inaccurate, requiring a high sample rate to reach convergence. In addition, it is particularly sensitive to the measurements near the axis, where the denominator in the integrand becomes small, and it is also affected greatly by the presence of noise due to the reliance on a derivative.

Fortunately, numerous algorithms are available to make the transformation more accurate and reliable. The main one used for the density calculations in this thesis is the ‘three-point’ method¹¹, which is well suited to applications where the noise level is not high relative to the change in signal between samples. It cannot resolve single-point values, since the three points smooth over such small fluctuations, but this is not a concern for our application. Methods for computing the Abel inversion rely on an axis

¹¹ [182] C. J. Dasch (1992) *Appl. Opt.*

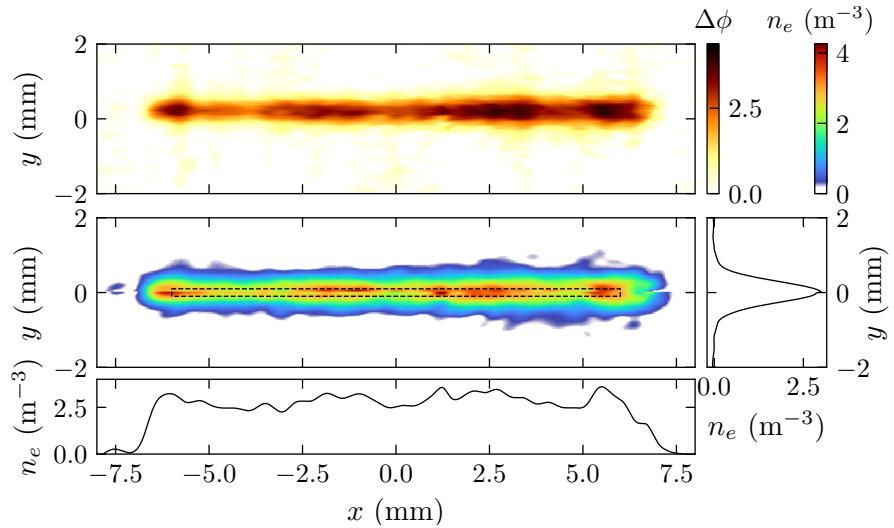


Figure 3.8: The top panel shows the phase extracted from the interferograms in Fig. 3.7 while the lower panel shows the density map retrieved by the three-point method. The side plots show the density profile of the channel, averaging over the region within the black dotted lines.

of symmetry. For real data, this does not exist, therefore a compromise axis must be selected, either manually or computationally. In this thesis, the centre was found by fitting a Gaussian to the transverse phase profile of the channel. An example phase map and the resulting density measurements are plotted in Fig. 3.8.

3.3.3 Accuracy of interferometry

Almost all the errors in interferometry come from the conversion of the phase map to the density map. One large source of error can be incorrect spatial calibration, as this can produce errors due to miscalculation of the channel size, to which the density measurement is sensitive. Spatial calibration is usually performed by placing an object with known dimensions into the focal plane of the diagnostic and then calculating the ratio between the object's size and its pixel width in the image. This can be a problem if the magnification is not equal in each direction, so a first-order improvement to this would be to obtain separate spatial calibrations in both planes individually. However, imperfections in the imaging optics can cause local warping of the image, which could only be properly corrected by using a pixel-wise calibration. A solution could be to image a mesh of known dimensions as a reference and then use this to map each pixel to real space individually.

The accuracy of interferometry could also be improved by increasing the transverse sampling frequency. This would have two beneficial effects: firstly, the central region of the plasma used for measuring the density is typically only a few pixels wide, and so increasing the number of pixels in this region would improve the measurement accuracy; and secondly, it would make identification of the axis of symmetry more reliable, as a single pixel error would have a smaller effect. Typical CCDs are approximately square, while underdense laser plasmas are long and thin, such that a large portion of the

camera's pixels are not utilised. Consequently, the use of different magnifications in each plane would be advantageous.

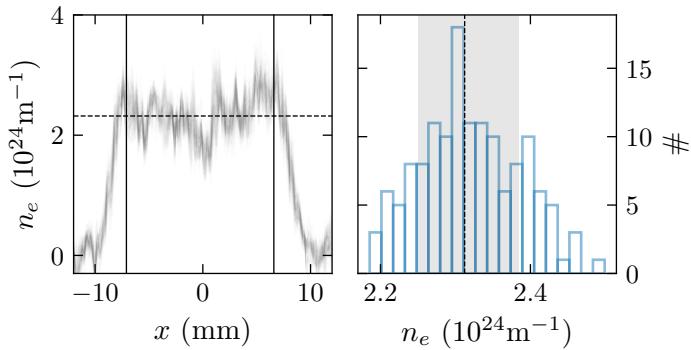


Figure 3.9: Stability of interferometry measurements over 141 shots at the same conditions. The median density of the central region is $2.30 \pm 0.07 \times 10^{24} \text{ m}^{-3}$, where the error is the standard deviation of the measurements.

Fig. 3.9 uses a series of 141 interferometry measurements performed under the same conditions to examine the stability of the diagnostic. This data was taken using the experimental setup in chapter 5. The gas jet target can be assumed to be providing helium molecules with a consistent density profile. While the laser energy variations will affect the channel radius, the plasma density should not be sensitive to these, as the helium at the centre of the channel will be fully ionised for all possible laser energies. The variations in the density measurement are therefore dominated by random errors in the interferometry, which result from a combination of noise, and changes in the identification of the central axis. At this density the standard deviation was $\pm 3\%$; if single shot measurements are to be used, for example, for optimisation procedures, or if it is necessary to monitor fine adjustments to the density profile, less variability would be beneficial.

3.4 Electron beam diagnostics

In this section, we review the electron beam diagnostics employed in this thesis. These come in two forms, beam profile measurements and spectrometer measurements. These both rely on scintillating materials, the properties of which will be discussed first. Finally, a modified pepper-pot emittance measurement technique will be explained.

3.4.1 Detection

A scintillator is a material that emits visible light after the incidence of ionising radiation. The radiation causes electrons within the material to attain an excited state, which then subsequently decay back to the ground state by releasing radiation with a characteristic spectrum¹². Terbium-doped gadolinium oxysulphide ($\text{Gd}_2\text{O}_2\text{STb}$) was used as the scintillating material for all experiments performed, a material that is common in laser

¹² [183] J. Lindström *PhD. Thesis* (2021)

Table 3.1: Scintillator screen relative brightnesses and absolute resolution measurements.

	Lanex Regular	Luminex Ultrafine
Relative brightness	1	0.4
Contrast at 2 lines/mm	0.23	0.65
Contrast at 5 lines/mm	0.05	0.29

plasma accelerator experiments. This is due to its high efficiency¹³ and sufficiently short decay time of 600 μ s, which in principle allows for single shot temporal resolution at kilohertz frequencies. Its emission spectrum has a narrow peak at 545 nm¹⁴, which can be isolated from the continuum with the use of notch band-pass filters, a useful property in a noisy light environment that includes lasers and plasma emissions. The fluorescence intensity per electron is almost completely independent of the incident electron's energy above approximately 1 MeV¹⁵, which is advantageous for charge calibrated measurements, and crucial for charge calibration if the electron beam has unknown energy.

Scintillating screens are plastic sheets coated with a thin layer of the scintillator and have been manufactured for industrial use by numerous companies over the years. Until recently, Kodak manufactured these screens, and their product, Lanex regular¹⁶, was used for the electron detectors in chapters 4 and 6. Note that Lanex is such a common product in the LWFA community that the product name has become genericised. Chapter 5 used the product Luminex Ultrafine from Scintacor¹⁷, due to the higher resolution of this type of screen. The properties of these scintillator screens are compared in Tab. 3.1¹⁸. There is a trade-off between brightness and resolution because resolution improvements are obtained by reducing the thickness of the scintillator, which naturally reduces the emission intensity.

Charge-calibrated measurements are reported in chapter 4, and these use an image plate (IP) to calibrate between the number of counts observed by a camera and the electron flux. It is possible to calibrate the emission from Lanex and Luminex directly¹⁹, however, this involves precise measurements of the capture angle of the imaging system, and the transmission of all of the components in the imaging line, which can be error prone when using multiple optical components. Fujifilm IP, BaFBr_{0.85}I_{0.15}:Eu²⁺, was used²⁰, which is another scintillating material, but with a much longer decay time than Lanex. The decay can be stimulated to occur by irradiating the phosphor with 623 nm light²¹, and scanners have been developed to exploit this property. By fluorescing in

¹³ [184] Y. Glinec *et al.* (2006) *Review of Scientific Instruments*

¹⁴ [185] E.-J. Popovici *et al.* (2004) *Optical Materials*

[186] J. Lindström *et al.* (2020) *Physica Medica*

¹⁵ [184] Y. Glinec *et al.* (2006) *Review of Scientific Instruments*

¹⁶ [187] *LANEX screens*, Carestream

¹⁷ [188] *X-ray Phosphor Screens*, Scintacor

¹⁸ [189] G. C. Tyrrell (2005) *Nucl. Instrum. Methods Phys. Res. A: Accel. Spectrom. Detect. Assoc. Equip.*

[190] *DRZ Screens*, MCI Optonix LLC

¹⁹ [191] A. Buck *et al.* (2010) *Review of Scientific Instruments*

²⁰ [192] S. G. Gales *et al.* (2004) *Review of Scientific Instruments*

a scanner, all the emission light can be measured, so no optical systems need to be calibrated. The conversion from the intensity of the stimulated emission from IP to the deposited energy, and hence the charge, is well understood²², and so by comparing the scans to the scintillating screen images, a conversion between camera counts and charge flux can be obtained. It should also be noted that the emission intensity is angle dependent²³, so for accurate charge calibration across the entire screen, it is necessary for the observation angle to be roughly constant, or for a pixel-wise calibration between the IP and the image to be performed.

3.4.2 Electron profile

The electron profile diagnostic is for simply imaging the beam's entire transverse profile at a certain point in its propagation. Typically, it is set up by placing a scintillating screen on-axis and then imaging the screen onto the CCD of a camera. The beam should be allowed to expand to a size that is larger than the combined resolution limit of the screen and the scintillating system. For an LWFA beam with mrad divergence, and Lanex regular with approximately 200 μm resolution, this distance should be larger than 0.5 m. Normally, the side of the screen facing the target is coated in a protective foil, and it is the rear surface that is imaged, as this reduces the possibility of low energy electrons and residual light skewing the measurement. However, if sufficient care is taken to filter residual light and the distance from the target is large enough, then the front surface can also be used. After recording a satisfactory profile image, the image is corrected for viewing angle distortions, and a background subtraction is performed. The image is also calibrated spatially by using a grid with known dimensions placed on the screen, and the beam axis is determined using an alignment laser.

3.4.3 Electron spectrometer

The electron spectrometer determines the energy spectrum of the beam by converting energy to position with the aid of a deflecting dipole magnet. Electrons follow an energy-dependent trajectory through the magnet, such that their position on a scintillating screen placed after the magnet can be used to determine their energy. Since the dipole deflection is perpendicular to its field line directions, the electron beam transverse properties along the axis of the field lines are relatively unperturbed by the magnet, meaning that the non-dispersed axis of the beam can still be used to measure the 1D divergence and pointing of the beam. The electron spectrometer's ability to simultaneously monitor the spectral and spatial properties of the beam makes it a ubiquitous diagnostic on LWFA experiments.

²¹ [193] I. J. Paterson *et al.* (2008) *Measurement Science and Technology*

²² [194] K. A. Tanaka *et al.* (2004) *Review of Scientific Instruments*

[195] K. Zeil *et al.* (2010) *Review of Scientific Instruments*

[196] S. Kneip *PhD. Thesis* (2010)

[170] J. C. Wood *PhD. Thesis* (2016)

²³ [197] G. E. Giakoumakis *et al.* (1985) *Physics in Medicine Biology*

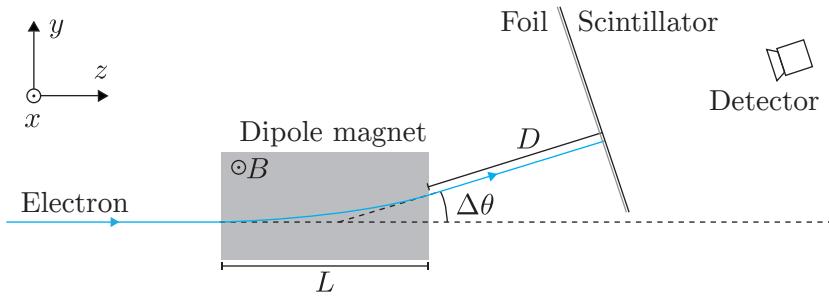


Figure 3.10: Diagram of an electron spectrometer.

Electron spectrometers for GeV experiments are usually set up such that they maximise the resolution at the highest energy. This is because this region of the spectrometer is usually both the most interesting and has the worst resolution, owing to the difficulty of dispersing high-energy electrons. Inside the dipole with field strength B , electrons with $v \approx c$ acquire transverse momentum p_\perp at a rate of ecB . For a magnet of length L , the high energy electrons are deflected by

$$\Delta\theta \approx \frac{\Delta p_\perp}{\gamma m_e c} = \frac{eBL}{\gamma m_e c}. \quad (3.15)$$

The approximation assumes that the electron's longitudinal momentum is much larger than the transverse momentum it acquires due to the magnet. A diagram of this deflection is shown in Fig. 3.10. The spatial dispersion at the scintillating screen is then

$$\frac{d(\Delta\theta D)}{d\gamma} = \frac{e}{m_e} \frac{BLD}{\gamma^2}. \quad (3.16)$$

The γ^{-2} scaling is why the highest energy electrons are the most difficult to resolve. Increasing the magnet strength and size will increase the angular dispersion while moving the screen further away will magnify the spatial dispersion. Considering the beam of electrons with finite divergence along the dispersion axis, if the beam's divergence along the dispersion axis is larger than the angular dispersion, then increasing D will not improve energy resolution. For this reason, if the electron beam divergence is large, it is common to put a slit in front of the magnet to select only those electrons with low divergence. This comes at the cost of the ability to measure all of the beam's charge at each energy.

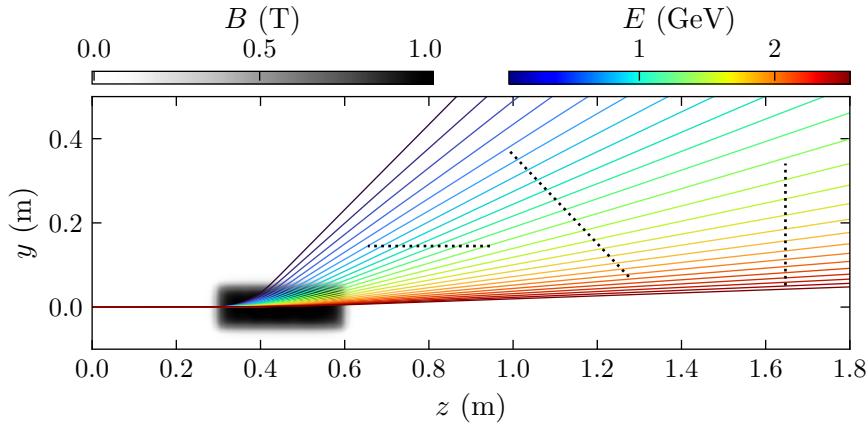


Figure 3.11: Calculation of particle trajectories for different energies through mapped magnetic field. The dotted lines indicate the locations of the scintillating screens.

In order to perform an accurate calibration, the exact magnetic field and spectrometer layout are measured, and these measurements are used to calculate electron trajectories. The magnet map is obtained by scanning the region in and around the dipole with a Hall probe. Measurements are made using two motorised translation stages, such that a grid for the B field can be obtained, which is used to construct the magnetic field map using linear interpolation between the gridpoints. It is important to also scan the region outside the magnet since fringe fields can have an appreciable effect. This thesis used 2D maps, corresponding to the $y - z$ plane in Fig. 3.10, centred in x on the approximate beam axis. An improvement on this would be to use a 3D map, such that off-axis beams are modelled properly. An example of this tracking has been performed for a 300 mm long 1 T peak field magnet in Fig. 3.11.

The arrangement of screens in Fig. 3.11 was the spectrometer setup for the experiment in chapter 5. An image of the setup from inside the target chamber is shown in Fig. 3.12, where the positions of the magnet and the three scintillating screens have been highlighted. The screens all have a vertical extent of approximately 100 mm, and a length between 250 mm and 450 mm. Multiple screens were used so that a large energy range could be covered with a high resolution. In addition, the energy ranges of the screens overlapped significantly, as this allowed for pointing fluctuations along the dispersion axis to be deconvolved.

3.4.4 Emittance measurement

An electron beam mask can be used in conjunction with either the profile or spectrometer to measure the transverse emittance. The mask can be either a series of slits, for 1D measurements²⁴, or pinholes, for 2D measurements²⁵. A wire grid can also be used for low-charge beams, as this will increase the signal. The 2D pinhole mask is known

²⁴ [81] C. M. S. Sears *et al.* (2010) [Physical Review Special Topics - Accelerators and Beams](#)

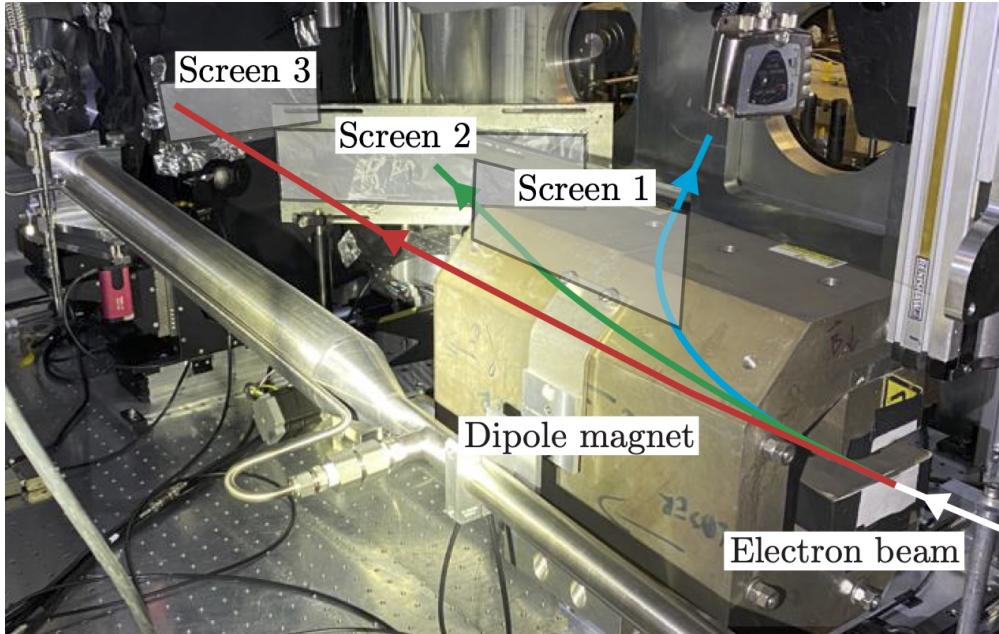


Figure 3.12: Experimental layout of the magnets and electron spectrometer screens for the results in chapter 5, at the ATA3 target area of the Gemini laser.

as a ‘pepper-pot’ plate, and the method of extracting the emittance from the masked beam’s pattern is often referred to as the pepper-pot method, even if performed in 1D. As discussed in Sec. 2.8, the emittance is defined as

$$\varepsilon_x = \sqrt{\langle x^2 \rangle \langle x'^2 \rangle - \langle xx' \rangle^2}, \quad (3.17)$$

where x is the particle transverse pointing and $x' = p_x/p_z$ is the pointing angle of the particle, where z is the longitudinal direction. The angled brackets denote the RMS value of the quantity over the entire distribution of particles.

The mask with p slits produces an equal number of beamlets, each with its own size, pointing, and correlation. It is possible to express each of the terms in Eq. 3.17 in terms of the beamlet measurements²⁶. These are

$$\langle x^2 \rangle \approx \frac{1}{N^2} \sum_{j=1}^p n_j (x_j - \bar{x})^2, \quad (3.18)$$

$$\langle x'^2 \rangle \approx \frac{1}{N^2} \sum_{j=1}^p n_j [\sigma_{x'_j}^2 + (x'_j - \bar{x}')^2], \quad (3.19)$$

and

$$\langle xx' \rangle^2 \approx \frac{1}{N^2} \left[\sum_{j=1}^p n_j x_j \bar{x}'_j - N \bar{x} \bar{x}' \right]^2, \quad (3.20)$$

²⁵ [198] Y. Yamazaki *et al.* (1992) *Nucl. Instrum. Methods Phys. Res. A: Accel. Spectrom. Detect. Assoc. Equip.*

[199] R. P. Shanks *et al.* ()

[200] E. Brunetti *et al.* (2010) *Physical Review Letters*

²⁶ [201] M. Zhang *Emittance Formula for Slits and Pepper-pot Measurement* (1988)

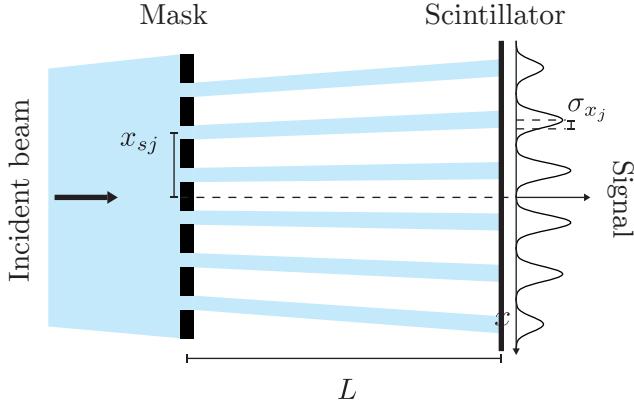


Figure 3.13: A beam is split into a series of beamlets using a mask. This produces a series of peaks on the scintillating screen.

where n_j is the number of electrons that propagate through the j^{th} slit, $N = \sum_{j=1}^p n_j$, x_j is the position of slit j , x'_j is the average pointing of the j^{th} beamlet, \bar{x} is the weighted average position of all the beamlets, and \bar{x}' is the weighted average pointing of the beamlets. The critical quantity is $\sigma_{x'_j}$, which is the RMS divergence of the j^{th} beamlet, which is equal to σ_{x_j}/L , as indicated by the diagram in Fig. 3.13.

The above expressions are derived under the assumption of infinitely small slits, but this is not the case, and so the size of the slit needs to be taken into account. If the beam had zero emittance, then the slit would produce a top-hat profile on the scintillator. For a finite emittance beam, this top-hat transmission function convolves with the divergence that is due to the emittance. The measured divergence of the beamlet is then the quadrature sum of the two RMS divergences, $\sigma_{x'_j,m} = \sqrt{\sigma_{x'_j}^2 + \sigma_s^2}$, where $\sigma_{x'_j}^2$ is the divergence due to the emittance alone in Eq. 3.19. The RMS width of the top-hat function is $Md/\sqrt{12}$, where d is the slit width, and M is the magnification, given by the ratio between the distances of the mask and the screen from the source, $M = L_{m,s}/(L_{m,s} + L)$. Consequently, in terms of the screen measurements, the divergence of the j^{th} beamlet due to the emittance alone is

$$\sigma_{x'_j} = \sqrt{\left(\frac{\sigma_{x_j}}{L}\right)^2 - \left(\frac{L_{m,s}d}{(L_{m,s} + L)\sqrt{12}}\right)^2}. \quad (3.21)$$

To confirm the validity of this scaling, Monte Carlo (MC) simulations of the electron beam drift have been performed. One of these simulations is shown in the two left plots of Fig. 3.14, where an electron beam with a divergence of 1 mrad is masked by 2 μm wide slits, and then propagates 50 mm. The right panel shows the effect of increasing the slit width on the beamlet size for a magnification²⁷ of 34, and the dashed line shows the scaling used to derive Eq. 3.21. This scaling agrees with the MC simulation results well,

²⁷ Corresponding to the geometry used for the emittance measurements presented in chapter 5.

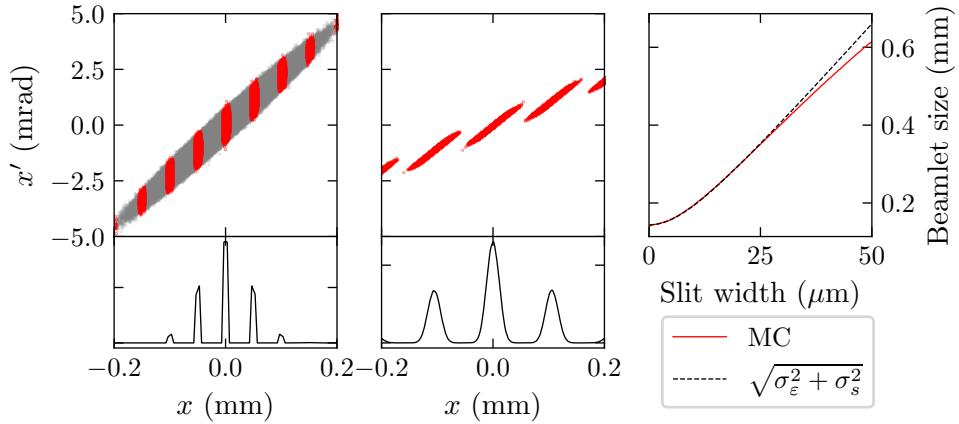


Figure 3.14: Monte Carlo simulation of the electron beam propagation through the mask. The red scatter points indicate the electrons transmitted by the mask.

with a slight discrepancy at larger slit widths that results from the slit size becoming large compared to the beam size.

3.5 Particle-in-cell simulations

Numerical simulations of laser-plasma interactions have been an important tool for obtaining and improving our understanding of physical processes. The first paper on LWFA was supported by numerical simulations²⁸, and they have been used extensively to explore the capabilities and limits of plasma acceleration ever since. Semi-empirical scaling laws have been derived through their use²⁹, and now optimisation methods³⁰ and novel positron acceleration paradigms are being tested³¹. The computational scale of simulations has increased by many orders of magnitude since the 1979 paper, now reaching the exascale³².

The process of LWFA involves laser pulses that are on the order of the plasma period in length, relativistic non-Maxwellian particle beams, and non-linear laser pulse evolution. In order to model this physics as accurately as possible, Maxwell's equations should be solved directly at the positions of each particle of interest, a daunting task when considering the full 6D phase space of the particle distribution at densities of $1 \times 10^{24} \text{ m}^{-3}$. To reduce the computational cost, large ensembles of real particles can be represented by single virtual macroparticles, where each macroparticle represents a sampling of the particle phase space distribution. The electromagnetic fields are typically calculated at

²⁸ [27] T. Tajima *et al.* (1979) *Physical Review Letters*

²⁹ [44] W. Lu *et al.* (2007) *Physical Review Special Topics - Accelerators and Beams*

³⁰ [202] F. Irshad *et al.* (2022) *arXiv e-prints*

³¹ [118] T. Silva *et al.* (2021) *Physical Review Letters*

[115] S. Diederichs *et al.* (2020) *Physical Review Accelerators and Beams*

³² [203] L. Fedeli *et al.* “Pushing the Frontier in the Design of Laser-Based Electron Accelerators with Groundbreaking Mesh-Refined Particle-In-Cell Simulations on Exascale-Class Supercomputers” (2022)

[204] J. L. Vay *et al.* (2018) *Nucl. Instrum. Methods Phys. Res. A: Accel. Spectrom. Detect. Assoc. Equip.*

the vertices of a pre-defined grid, further reducing the computational load. Simulations that perform these two simplifications are known as particle-in-cell (PIC) simulations, which are a type of kinetic simulation.

3.5.1 PIC algorithm

Before the PIC algorithm can be implemented, the simulation must be initialised. This involves defining the starting positions of the particles, the initial values of the fields, and the boundary conditions at the edges of the simulation domain. Once this is done, a recursive algorithm is implemented to evolve the particle positions in time and space, which proceeds for a user-defined number of steps. The structure of the algorithm is as follows:

1. The field values are interpolated at the particle positions.
2. The particle momenta and positions are computed and updated.
3. The current densities are projected onto the grid.
4. The new electromagnetic fields are calculated.

The particle momenta are updated using the Lorentz force, using the newly calculated fields. Care is required to ensure that the simultaneous update of the particle positions and momenta conserves energy, which is typically achieved by using the Boris algorithm³³. This is a leapfrog method, where the particle is first moved, then half of its energy change is applied, then the momentum vector is rotated, and then finally the second half of the energy change is applied. While this algorithm is adequate in Euclidean frames, in Lorentz boosted frames, the Boris algorithm has been shown to perform poorly³⁴, and so other algorithms have gained prevalence³⁵.

Maxwell's equations are typically solved using the finite difference time domain (FDTD) approach. The Yee grid³⁶ is a grid where electric fields are defined at the grid edges and the magnetic fields are defined at the grid faces. This method exhibits long-term stability, and as such is another popular choice for spatio-temporal discretisation. The 3D Cartesian simulations of plasma mirrors performed in this thesis use the Yee implementation of the FDTD method, although another method is used for the laser wakefield simulations, as will be discussed in the next section.

Since the standard PIC implementation only requires local information at each step, the full simulation domain can be split into numerous sub-domains to allow for parallel computation. This is particularly advantageous when a large number of CPUs are available. Each sub-domain is assigned to a separate CPU, which performs the PIC algorithm for that region alone, and the only information that needs to be shared between the CPUs are the field and particle values at the edges of the domains. Not only does

³³ [205] J. P. Boris “Relativistic Plasma Simulation - Optimization of a Hybrid Code” (1970)

³⁴ [206] J. L. Vay (2008) *Physics of Plasmas*

³⁵ [207] B. Ripperda *et al.* (2018) *The Astrophysical Journal Supplement Series*

³⁶ [208] K. Yee (1966) *IEEE Transactions on Antennas and Propagation*

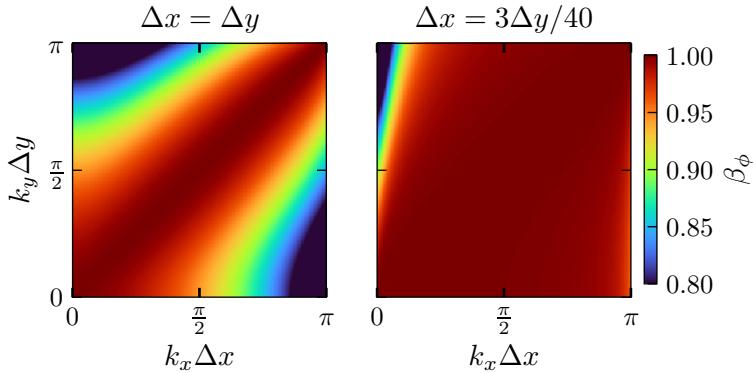


Figure 3.15: Phase velocity of electromagnetic waves as a function of 2D wavevector for $\Delta x = \Delta y$ (left), and for $\Delta x = 3\Delta y/40$ (right).

this speed up the computation, but it also allows for a large amount of memory to be used, permitting simulations with high resolutions to be performed quickly. The full 3D simulations performed for this thesis ran on Imperial College’s cluster, CX1, requiring almost 3 years of clock time, but took less than 72 hours to run due to the number of CPUs employed.

3.5.2 Numerical dispersion

The FDTD algorithm is subject to numerical dispersion³⁷. This is concerning for the simulation of relativistic particles, as a reduced speed of light can lead to particles having super-luminal velocities. By inserting a plane wave solution into the FDTD algorithm, it can be shown that the 2D dispersion relation is given by

$$\frac{\sin^2(\omega\Delta t/2)}{(\Delta t^2)} = \frac{\sin^2(k_x\Delta x/2)}{(\Delta x^2)} + \frac{\sin^2(k_y\Delta y/2)}{(\Delta y^2)}. \quad (3.22)$$

The resulting phase velocity $v_\phi = \omega/k$, is

$$v_\phi = \frac{2}{k\Delta t} \sin^{-1} \left(\Delta t \sqrt{\frac{\sin^2(k_x\Delta x/2)}{(\Delta x^2)} + \frac{\sin^2(k_y\Delta y/2)}{(\Delta y^2)}} \right). \quad (3.23)$$

A feature of the Yee solver is that, for equal grid spacings, waves propagating along the cell diagonal are dispersion free, as can be seen in Fig. 3.15, while waves propagating along the grid axes experience the most dispersion. This is unfortunate for LWFA simulations, as typically the laser pulse runs parallel to one of the grid axes. This is mitigated by choosing an unequal grid spacing, where setting $\Delta x \ll \Delta y$ substantially reduces the dispersion along the x axis. This choice of resolution is also appropriate for LWFA simulations because the scale length of interest along the laser propagation axis is the laser wavelength, which is much shorter than the transverse length scale; the plasma wavelength³⁸. The exact phase velocity for a given 2D resolution can be found by setting

³⁷ [209] R. Nuter *et al.* (2014) *The European Physical Journal D*

³⁸ This may not be the case if the transverse motion of ultra-low emittance beams needs to be resolved.

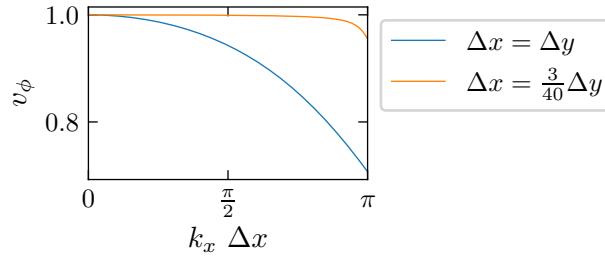


Figure 3.16: Phase velocity along the x axis for different grid spacing ratios.

$k_x = k$, and $k_y = 0$, giving,

$$v_\phi(k, \Delta x, \Delta y) = 2 \frac{\sqrt{\Delta x^{-2} + \Delta y^{-2}}}{k} \sin^{-1} \left(\frac{\sin(k\Delta x/2)}{\sqrt{1 + \frac{\Delta x^2}{\Delta y^2}}} \right). \quad (3.24)$$

As can be seen from Fig. 3.16, for appropriately chosen resolutions, numerical dispersion will not lead to a significant difference in the velocity of the laser pulse and that of the relativistic particles. However, even if the choice of resolution is such that the laser phase velocity and particle velocity are essentially the same, the highest frequencies supported by the grid, where $k_x \Delta x \approx \pi$, will propagate sub-luminally, and high energy particles will propagate faster than these components. This causes a numerical form of Cherenkov radiation to be emitted in the high-frequency band of the simulation; the electric fields of this spurious radiation can act on the particles, producing non-physical emittance growth³⁹. For accurate simulation of transverse beam properties over long propagation lengths, suppression or elimination of this instability is required. Methods for doing so with FDTD solvers include filtering the higher frequencies and increasing the size of the stencil used to calculate the fields⁴⁰. Alternatively, spectral methods can be used to solve Maxwell's equations in Fourier space, which can lead to dispersion-free solutions in vacuum⁴¹.

3.5.3 Specific PIC codes

The code Smilei⁴² was used for the fully 3D laser-solid interactions performed in this thesis. Smilei was chosen due to its ability to output simulation diagnostic information at a different resolution than the main simulation grid, which was useful for performing analysis that required high temporal but low spatial resolution. It has been used to study the Weibel instability, both generally⁴³, and in laser-plasma interactions⁴⁴, and it is for this purpose that Smilei is employed in this thesis. Smilei can be used for performing laser

³⁹ [210] R. Lehe *et al.* (2013) *Physical Review Special Topics - Accelerators and Beams*

⁴⁰ [211] A. D. Greenwood *et al.* (2004) *Journal of Computational Physics*

⁴¹ [212] O. Buneman *et al.* (1980) *Journal of Computational Physics*

⁴² [213] J. Derouillat *et al.* (2018) *Computer Physics Communications*

⁴³ [214] A. Grassi *et al.* (2017) *Physical Review E*

⁴⁴ [215] A. Grassi *et al.* (2017) *Physical Review E*

wakefield simulations, but like all FDTD solvers it can suffer from numerical dispersion, and in its full 3D mode it is overly computationally costly⁴⁵.

The majority of simulations performed in this thesis use the code FBPIC⁴⁶, which is more specialised than Smilei since it was written specifically for simulating wakefield acceleration. FBPIC exploits the near-cylindrical symmetry of LWFA by using Maxwell's equations in cylindrical coordinates $(x, y, z) \rightarrow (z, r, \theta)$, and in addition, solving for the fields in spectral space $(z, r, \theta) \rightarrow (k_z, k_{\perp}, m)$. As a result, the code is dispersion free in all directions in Euclidean frames. Using cylindrical coordinates and spectral space simultaneously requires the use of the Fourier-Hankel transform, which represents the radial variations of the fields as a sum of Bessel functions, J_m ; under this transformation, the fields decouple in spectral space⁴⁷. To exactly model the fields, the transformation should integrate in k_{\perp} from 0 to $+\infty$, in k_z from $-\infty$ to $+\infty$, and sum over all m from $-\infty$ to $+\infty$. However, if the fields exhibit a high degree of cylindrical symmetry then only a small number of modes are required to approximate them reasonably well, decreasing the number of computations required by orders of magnitude. For example, for many LWFA phenomena, only 2 modes are required, and so a 2D cartesian grid with 1000 cells in x and y that would require 10^6 operations to compute would only require 5×10^3 operations, summing over the modes from -2 to $+2$. Once solved, the fields are used to push the particles in 3D cartesian space, such that the macroparticles each have 6 coordinates (x, y, z, p_x, p_y, p_z) . The entire operation amounts to a fully 3D computation of the fields, but with aggressive low-pass filtering of the azimuthal variations around a pre-defined central axis. Finally, FBPIC offers operation in a Lorentz-boosted frame, which can further reduce computational time, particularly for centimetre-length simulations.

3.5.4 Simulating asymmetrical focal spots

It is common practice to model the transverse intensity profile of the laser pulse at focus as a Gaussian. This has the advantage of convenience, as well as being able to initialise the pulse at arbitrary distances from focus. However, if comparison to experimental results is the objective, then the simulation should match the experiment as closely as possible, and to this end, matching the distribution of energy in the laser far field is an improvement over using a perfectly symmetrical profile. It is expected that this is particularly important if the transverse emittance of the injected electron beam is being measured, as is the case in this thesis. Here, we show how experimental focal spot images are used to inject asymmetrical focal spots into FBPIC simulations.

⁴⁵ Smilei has recently introduced azimuthal mode decomposition and spectral solvers, but at the time of writing these were in the beta stages of development. In addition, the spectral solver was not available for all geometries, nor was it omnidirectional.

⁴⁶ [216] R. Lehe *et al.* (2016) *Computer Physics Communications*

⁴⁷ Maxwell's equations in cartesian coordinates only require a Fourier transform to decouple.

For a Gaussian transverse profile, the electric field normalised to its maximum can be written as

$$\frac{E(x, y, z)}{E_0} = \xi^{-1} \exp \left(-ikz - \frac{x^2 + y^2}{w_0^2 \xi} \right), \quad (3.25)$$

where E_0 is the maximum of the field, k is the wavenumber, w_0 is the Gaussian waist. The term ξ is the complex radius of curvature, given by

$$\xi = \left(1 + i \frac{z - z_f}{Z_R} \right), \quad (3.26)$$

where z_f is the focal position in z , and Z_R is the Rayleigh length. This term reduces the field strength as the distance from focus is increased, implements the Gouy phase, and adds curvature to the wavefronts. Separating out the phase components, we can write

$$\frac{E(x, y, z)}{E_0} = \xi^{-1} \exp \left(-ikz - \frac{r^2 - ir^2 \left(\frac{z - z_f}{Z_R} \right)}{w_0^2 \left(1 + \left(\frac{z - z_f}{Z_R} \right)^2 \right)} \right). \quad (3.27)$$

A focal spot image measures the intensity $I(x, y)$, which is proportional to the square of the electric field. This should be used to modify the purely spatial part of Eq. 3.27, while leaving the phase components the same. Therefore, we can write

$$\frac{E(x, y, z)}{E_0} = \xi^{-1} \exp \left(-ikz + \frac{ir^2 \left(\frac{z - z_f}{Z_R} \right)}{w_0^2 \left(1 + \left(\frac{z - z_f}{Z_R} \right)^2 \right)} + \frac{\ln(\sqrt{I(x, y)})}{\left(1 + \left(\frac{z - z_f}{Z_R} \right)^2 \right)} \right), \quad (3.28)$$

where we are making the approximation that the wavefront of the aberrated spot has a radius of curvature that matches a perfect Gaussian. This is only an approximation, and consequently, this formula's validity is limited to $|(z - z_f)/Z_R| < 1$. More accurate simulations of laser pulses could be achieved by measuring the phase of the profile, such as by taking multiple focal spot images at different z and using a phase recovery algorithm, such as the Gerchberg-Saxton algorithm⁴⁸. However, since only the $z - z_f = 0$ images were available, we are limited to initialising focal spots near to focus.

3.5.5 Implementation

The above formula is used to convert focal spot images into a function for calculating the electric field, but the parameters E_0 and w_0 also need to be provided. The maximum field strength is determined from the laser power and the focal spot image pixel values, while the value for w_0 is measured by fitting an ellipse to the FWHM contour of the focal spot image, and then using the mean value of the major and minor axes⁴⁹. The function $I(x, y)$ is a linear interpolator that uses the focal spot image for its data points.

⁴⁸ [217] I. V. Il'ina *et al.* (2009) *Quantum Electronics*

[218] A. Borot *et al.* (2018) *Optics Express*

⁴⁹ Since this beam is not a solution to the paraxial wave equations, w_0 is simply a scaling factor that affects the radius of curvature of the wavefronts, and the rate of change of field strength as focus recedes.

The focal spot image has been artificially modified so that the value of the field drops to 0 at approximately $\sqrt{x^2 + y^2} > 2.5w_0$, ensuring that the electric field is zero at the boundary of the simulation window during initialisation. The code FBPIC performs the azimuthal mode decomposition in the same way that it does for any other function for the fields. Examples of this projection for various numbers of modes are shown in Fig. 3.17, where it can be seen that the projection becomes less symmetrical and more structured as the number of modes increases.

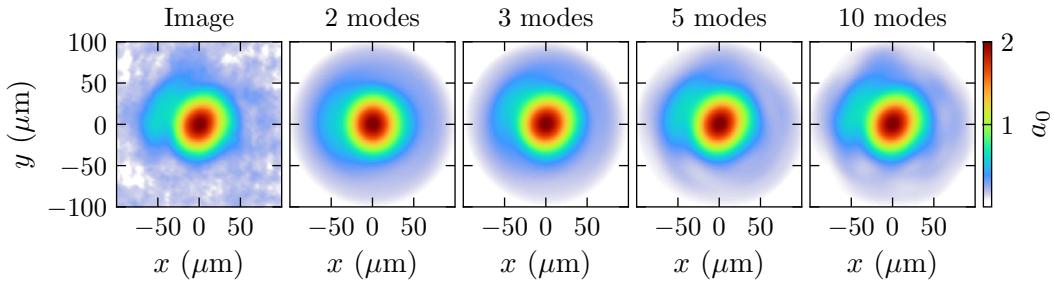


Figure 3.17: Projection of focal spot image with aberrations onto an increasing number of azimuthal modes.

The code has been tested by simulating a centimetre scale LWFA, and comparing the results to an equivalent simulation with a Gaussian focal spot. The resolution for the simulations was 20×4 points per laser wavelength longitudinally and transversely respectively, with 3 azimuthal modes, and 9 particles per cell. The evolution of the electron energy spectrum for each simulation is shown in Fig. 3.18. A similar evolution is observed between the two simulations, where injection occurs at $z \sim 2$ mm, and the maximum electron energy is ~ 2.2 GeV. Interestingly, the charge in the first bunch is slightly higher for the experimental focal spot, and it reaches its highest energy sooner, while the laser pulse depletes slightly later. The extended depletion length is due to the extra energy in the pulse that results from using a realistic profile.

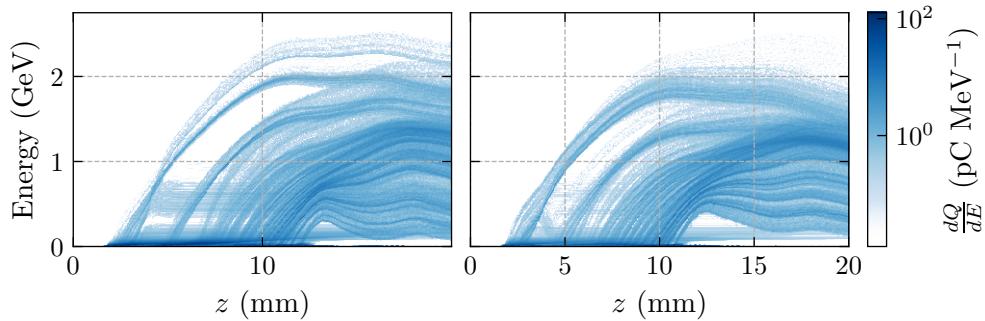


Figure 3.18: Evolution of the electron beam energy spectrum for the aberrated (left) and Gaussian (right) focal spots.

Fig. 3.19 compares the charge and transverse beam measurements for the two simulations, selecting electrons with energies above 5 MeV, and using the median average

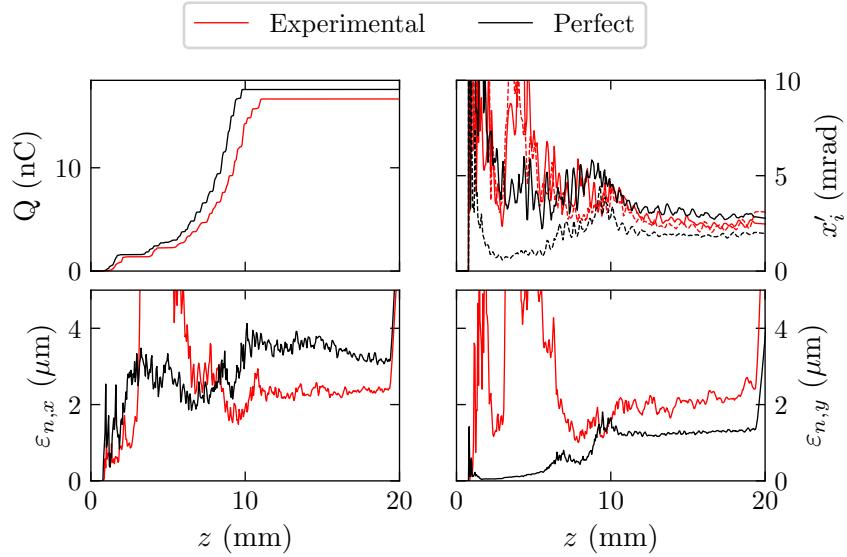


Figure 3.19: Comparing transverse beam properties for aberrated and Gaussian focal spots.

deviation to measure the transverse beam properties. The total amount of charge injected is similar for both simulations, as is the average beam divergence. The only significant difference is seen when the transverse beam measurements are projected onto the x and y directions, where it appears that the introduction of the asymmetrical focal spot has increased the symmetry of the electron beam, a counter-intuitive result. The divergence and emittance in x are slightly larger for the perfect focal spot, while they are slightly smaller in the y direction.

3.5.6 Longitudinal resolution

The minimum acceptable longitudinal resolution of the simulations needs to be determined. Numerical dispersion is not a concern for FBPIC, so we just need to be able to resolve the smallest structures in the simulation. However, choosing this size is non-trivial due to the lack of an analytic description of the plasma wave in the 3D non-linear regime. In particular, the electron bubble that forms in the wake of the laser could be challenging to resolve due to the thin, high-density spike that forms at the bubble's rear.

Simulations have been performed using the longitudinal resolution as the variable of interest. The aim was to determine the resolution required to resolve the bubble thickness and the effect of the resolution on self-injection. A 1 mm plasma was initialised, with 0.1 mm scale length density ramps at either end and a plateau density of $2 \times 10^{25} \text{ m}^{-3}$. A short, high-density simulation was chosen to reduce the computational load, necessary for running abnormally high resolutions in a reasonable amount of time. A Gaussian transverse profile with a width of 4 μm was used for the laser. The temporal profile was also a Gaussian, with a width of 20 fs, while the peak a_0 was 2. These laser and plasma

conditions roughly mirror those used for the pure helium simulations in chapter 4, once the laser pulse has evolved to its plasma-matched parameters.

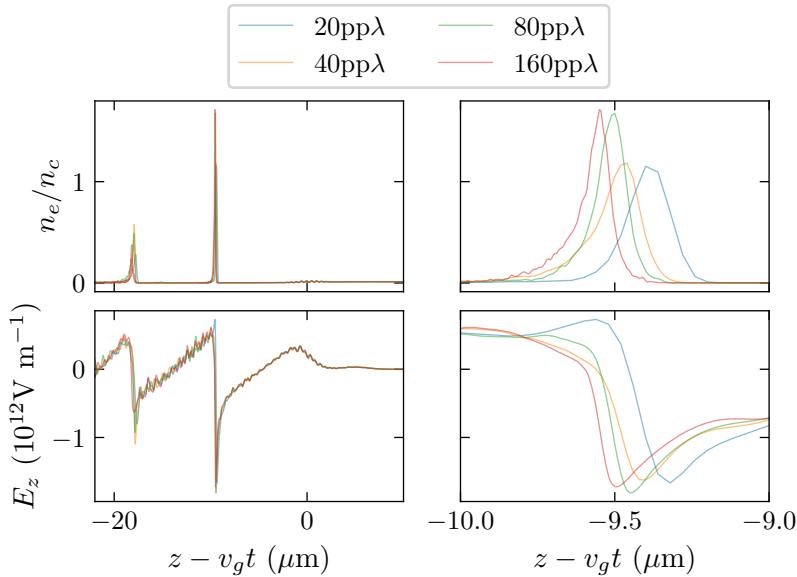


Figure 3.20: Axial plasma density and longitudinal electric field at various longitudinal resolutions. The left panels are on the scale of the plasma wavelength, while the right panels show the $1\text{ }\mu\text{m}$ region near the highest density spike.

Simulations were performed at 20, 40, 80, and 160 grid points per laser wavelength (pp λ). Fig. 3.20 shows the on-axis plasma density and longitudinal electric field at $z = 0.55\text{ mm}$ into the evolution, which is the moment before injection occurs. On the scale of the simulation window, the increasing resolution appears to have no effect on the plasma wavelength or the electric field. Small differences can be observed by zooming in on the back of the bubble, where the effect of increasing the resolution can be seen to both increase the density and decrease the width of the density spike, as well as slightly increase the plasma wavelength. Note that the change in peak density is substantial, approximately 50% when increasing the resolution from 40pp λ to 80pp λ . The change in plasma wavelength slightly shifts the location of the field minimum, while the effect on the magnitude of the accelerating field is small. It seems that the total amount of charge in the spike is relatively constant, which, due to Gauss's law, results in a similar accelerating field strength for each of the simulations, despite the differences in peak density. The change in location of the field minimum is on the order of 100 nm, which is small compared to the millimetre scale dephasing length, and so the effect on the peak energy is also likely to be small. This is confirmed in Fig. 3.21, where the evolution of the charge and the final electron energy spectra are compared. The longitudinal resolution has a negligible effect on these properties, and so a resolution of 20pp λ appears to be sufficient for simulating this set of input parameters.

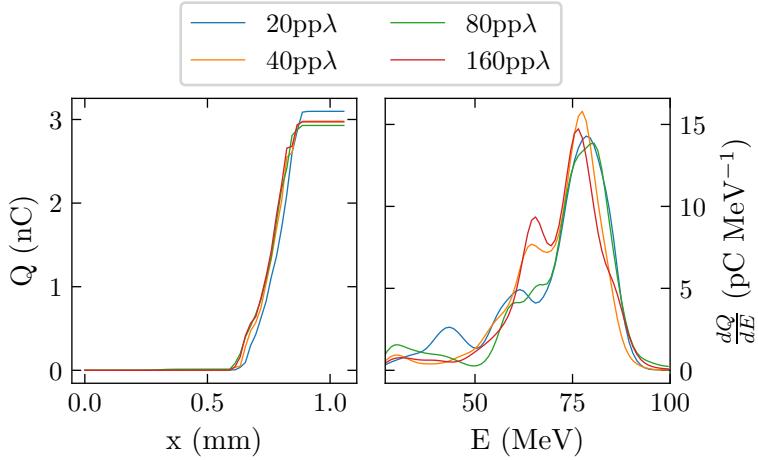


Figure 3.21: Left panel shows the evolution of the charge injected above 1 MeV, while the right panel shows the final energy spectrum of the beam.

3.5.7 Azimuthal modes convergence test

Here we determine the number of azimuthal modes required. It is expected that as the simulation departs further from cylindrical symmetry, more modes will be required to reach convergence. For this thesis, we are particularly concerned with the transverse emittance of electron beams injected via ionisation injection, using an asymmetrical focal spot. Consequently, a scan of the number of azimuthal modes has been performed for a GeV class accelerator with ionisation injection.

These simulations used the same aberrated focal spot used for the simulations in Fig. 3.19, with the a_0 reduced from 2.07 to 1.71. Helium doped with 2% nitrogen by mass was used as the accelerating medium, with the first electron of the helium and the first 5 electrons of nitrogen pre-ionised. When the helium was fully ionised, the plateau plasma density was $1.5 \times 10^{24} \text{ m}^{-3}$, and 0.15 mm scale length ramps were used at both ends of the 12 mm long density profile. The longitudinal and transverse resolution was $20\text{pp}\lambda \times 4\text{pp}\lambda$, and the simulation window spanned $60 \mu\text{m} \times 160 \mu\text{m}$ in z and r respectively. A single macro-particle was used in z and r , while 3 particles were used for each azimuthal mode. This was necessary to avoid under-resolving the higher mode simulations but resulted in an $\mathcal{O}(n^2)$ scaling of the computation time with the number of modes.

The results of these simulations are plotted in Fig. 3.22 for electrons which originate from the inner shells of nitrogen and have $E > 250$ MeV. Note that the $N_m = 5$ did not complete due to its long computational time. A similar amount of charge is produced in each simulation, however, the transverse properties do depend on the number of modes simulated. For the divergence in both x and y to converge, it appears that 3 modes are necessary, while the same is true for ε_y . However, for ε_x , it appears that convergence may require 4 modes. Since the divergence in x does not require 4 modes, this suggests that it is the beam size along the laser polarisation direction that requires the additional resolution. In this thesis, when the transverse properties of the beam were of interest, 3

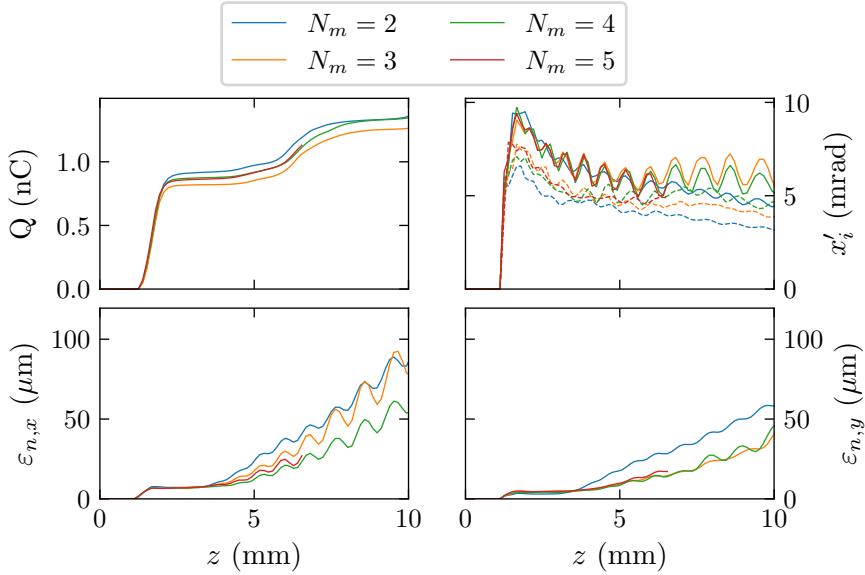


Figure 3.22: The evolution of the charge (top left), the beam divergence (top right), and the normalised emittance (lower panels), for particles that reach the end of the simulation with energy above 250 MeV. In the divergence plot, the solid lines are the divergence in the laser polarised direction, x , while the dashed lines correspond to the polarisation orthogonal direction, y .

modes have been used for simulations. This achieves a balance between computational speed and accuracy. The emittance measurements reported in chapter 5 are in the y direction, and so 3 modes should be sufficient for this purpose.

4 The density-length parameter space of a laser wakefield accelerator

THE PERFORMANCE of a laser wakefield accelerator is, to first order, determined by the laser power¹. However, for most laser powers, there exists a range of densities that will produce desirable electron and x-ray beams, and so usually the density of operation is determined empirically, based on the desired beam metrics. Since the density affects the rate of laser pulse evolution, as well as the electron beam acceleration dynamics, the optimum density is inexorably coupled to the length of the plasma. This means that the plasma length needs to be scanned in conjunction with the density to maximise the performance of the accelerator for any given application. While density scans are frequently reported in the literature, density-length scans are less common², owing to the requirement for a target with a variable length, as well as the time cost of finely scanning a 2D parameter space.

In this chapter, we investigate the density-length parameter space of a 5 TW laser wakefield accelerator. The 250 mJ laser had a central wavelength of 800 nm, a temporal FWHM duration of 45 fs, and was focused to an intensity of approximately $6 \times 10^{21} \text{ W m}^{-2}$. Identical scans were performed for two gas mixtures; pure helium, and a 1% nitrogen-doped helium gas, where the presence of the dopant would permit additional trapping of electrons via the ionisation injection mechanism. The density range was $0.3\text{--}2.8 \times 10^{25} \text{ m}^{-3}$, while a gas cell was used with a length which varied from 0–3 mm. Sec. 4.1 will outline the experimental setup, describing the configuration of the vacuum chamber, interferometer, electron spectrometer, and x-ray diagnostics. An automated Bayesian optimisation procedure was run for each gas mixture prior to the performance of the scans³, with the goal of maximising the signal on the electron spectrometer, and the results of these optimisations are shown here. Following this, in Sec. 4.2 the 2D scan experimental results will be presented; starting with the electron beam measurements,

¹ See Fig. 1.3.

² [64] M. Heigoldt *et al.* (2015) *Physical Review Special Topics - Accelerators and Beams*

[219] S. V. Rozario *PhD. Thesis* (2020)

[175] K. Poder *PhD. Thesis* (2016)

³ [220] R. J. Shalloo *et al.* (2020) *Nature Communications*

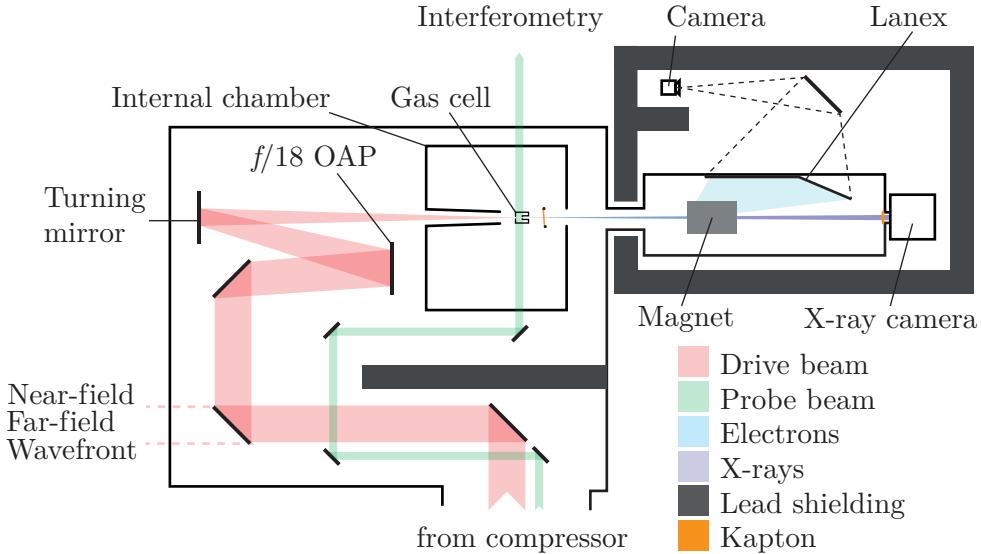


Figure 4.1: Layout of the ATA2 vacuum chamber.

before moving on to the quantification of the x-ray photon flux. Electron beam energies in excess of 200 MeV were measured, which are the highest electron beam energies measured at this laser power to date. Simulations were performed at densities that match the experimental parameters, and the results of these are shown in Sec. 4.3. Finally, Sec. 4.4 will investigate the validity of widely-used scaling laws in this regime, determine the rate of laser pulse evolution and its mechanism, and estimate the x-ray brilliance.

The high electron beam energies observed in the doped gas were attributed to beam acceleration, driven by the ionisation injected electron beam. It is shown that due to the low vacuum a_0 of the laser pulse, the transverse and longitudinal evolution of the laser pulse occur sequentially, delaying the length at which self-injection occurs beyond the typical depletion length. In addition, the low a_0 at low densities was found to lead to a reduction in the etching rate of the beam, which increased the phase velocity at the rear of the bubble. At some parameters, the photon flux exceeded $10.0 \text{ ph mrad}^{-2} 0.1\% \text{ BW}^{-1}$, resulting in a peak brilliance of $4.2_{-0.8}^{+0.8} \times 10^{20} \text{ ph s}^{-1} \text{ mm}^{-2} \text{ mrad}^{-2} 0.1\% \text{ BW}^{-1}$, where PIC simulation measurements have been used to estimate the source size. Compared to other measurements in the literature, this is an exceptionally high brilliance for the laser power.

4.1 Experimental setup

A diagram of the vacuum chamber is shown in Fig. 4.1. The laser energy for this experiment could not exceed $\sim 0.25 \text{ J}$ due to the heat load on the compressor gratings⁴, particularly when operating at repetition rates above 1 Hz. A small hole in the final mirror of the compressor allowed a portion of the pulse to be used for pulse length and phase measurements, which were performed using the SPIDER method⁵. When

⁴ [221] V. Leroux *et al.* (2020) *Optics Express*

optimally compressed, the beam was able to achieve a FWHM duration of 45 fs. However, most of the time, a *Fastlite Dazzler*⁶ was used to modify the 2nd, 3rd, and 4th order spectral phase terms, resulting in a longer pulse than the bandwidth limited length. The SPIDER was used to monitor on-shot changes in the temporal profile; this allowed for the effects of the Dazzler to be measured directly.

The 55 mm diameter beam was focused over a distance of 1 m by an off-axis parabola, constituting a focal geometry of $f/18$. The energy of the laser was measured using a Gentec energy meter placed between the turning mirror and the target. The energy measurements were used to calibrate near and far field measurements of the laser beam leakage, allowing for energy fluctuations to be monitored. The spatial phase of the laser could be controlled using an adaptive optic (AO), while the wavefront was monitored using an *Imagine Optics HASO4*⁷; a commercially available Shack-Hartmann wavefront sensor. During the daily alignment of the system, the wavefront sensor and AO would be used in a closed-loop feedback system to flatten the wavefront, with the goal of producing an intense focal spot. In addition, the AO could also be used actively during shots to optimise certain measurements, such as the number of counts on the x-ray camera or the electron beam divergence, while the HASO could be used to passively monitor these changes, allowing the aberrations introduced by the AO to be measured in terms of Zernike polynomial coefficients.

On the day that the data presented in this chapter was recorded, a series of 100 focal spot images were captured after the closed loop optimisation of the wavefront had been performed. A summary of these metrics is presented in Fig. 4.2, as well as the focal spot with the median peak a_0 . The median peak a_0 , FWHM, and percentage of energy within the FWHM were 0.57 ± 0.02 , $21.7 \pm 0.5 \mu\text{m}$, and $39 \pm 2 \%$, where the error comes from the standard deviation of the data set. For these calculations, it was assumed that the energy in the pulse was 225 mJ, while the temporal FWHM was 49.4 fs, measured using the SPIDER.

⁵ [222] C. Iaconis *et al.* (1998) *Optics Letters*

⁶ [223] *Fastlite Dazzler product information*, Fastlite

⁷ [224] *HASO4 FIRST*, Imagine Optic

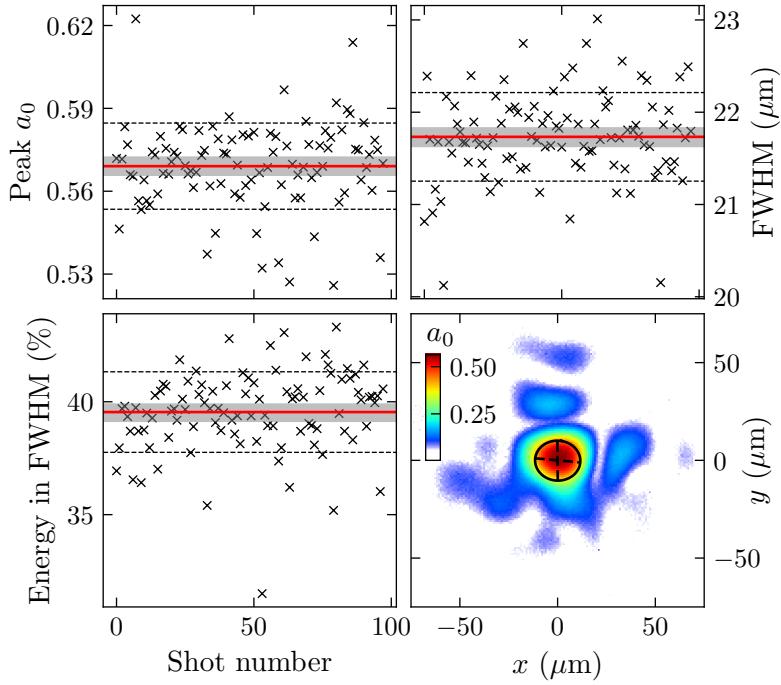


Figure 4.2: Focal spot metrics for the day on which the measurements in this chapter were obtained. The red horizontal lines indicate the median, the dashed grey horizontal lines show the standard deviation of the fluctuations, and the grey shaded region indicates the error on the mean, calculated via bootstrapping. A representative focal spot image, converted to normalised field strength, is shown in the lower right panel. This particular image had the median peak a_0 . An ellipse, fit to the FWHM contour of the intensity profile, has been overlaid on the image, with the dotted lines indicating the major and minor axes.

For this experiment, a relatively high repetition rate of 5 Hz was desired, which put additional strain on the vacuum pumping system. To alleviate this, differential pumping was employed, where the target was enclosed in a second internal chamber inside the main vacuum chamber, which was then pumped directly. Only a small aperture between the main chamber and the small chamber allowed the laser to reach the target, which, combined with the additional pumping, maintained the pressure difference. This caused the internal chamber, where the gas was being injected, to operate at a higher pressure than the rest of the chamber and compressor, which was able to maintain a low pressure of less than 1×10^{-3} mbar, protecting the optics.

A gas cell target was used for this experiment. The cell consisted of a steel frame, with glass windows to allow for imaging and optical probing of the plasma, and a ceramic front plate with a 200 μm diameter pinhole for the laser to enter. During shooting, jitter in the focal spot position and energy in the wings of the spot caused the hole to widen significantly to approximately 1 mm. At the rear of the cell, a piston was used to set the position of the exit pinhole, which again widened from an initial 200 μm to a 1 mm

diameter hole. The piston could move to set the distance between the entrance and exit holes to anywhere from 0 to 10 mm, although, in practice, a cell length of 0 mm meant that there was a buffer of 100 μm between the pinholes, eliminating the risk of the piston damaging the front plate.

After exiting the gas cell, the residual laser energy was dumped using a 25 μm tape-based plasma mirror. The electrons and x-rays propagated through a small aperture into an appendage vacuum chamber, which contained a 0.5 T permanent dipole magnet and a Lanex regular scintillating screen. This was imaged through a window by a 12-bit camera. The x-rays passed through the magnet and through some light-shielding foil before being imaged by a direct detection camera. The electron spectrometer and x-ray camera were housed in a lead enclosure, which shielded the rest of the target chamber from the radiation produced by the experiment.

4.1.1 Density calibration

Transverse interferometry was used to measure the plasma density; an example interferogram is shown in the left panel of Fig. 4.3. This diagnostic made use of a dedicated probe beam that was available in the ATA2 target area, which was taken from the leakage of the main beam after amplification but before compression. This 1 mJ pulse was compressed to approximately 100 fs using a dedicated at-air compressor, before being telescoped down to a diameter of 25 mm and directed into the target chamber. The probe beam backlights the plasma channel, which is imaged by a 750 mm focal length, 100 mm diameter objective lens and a 400 mm imaging lens onto a 12-bit CCD, and a Michelson interferometer was used to generate the fringe pattern⁸. Due to the size of the windows on the target and the obstruction caused by the piston in the gas cell, density measurements could only be made for target lengths greater than 1.7 mm. Consequently, the density values quoted in this chapter are obtained by calibrating the pressure of the gas reservoir to the density in the cell at 3 mm, as shown in the right panel of Fig. 4.3, and then assuming that the fractional change in cell volume that occurs when moving to shorter lengths does not significantly impact the density. A length of 3 mm was chosen for this calibration because this allowed a sufficient number of fringes to be observed, increasing the reliability of the measurement.

⁸ See Sec. 3.3.1.

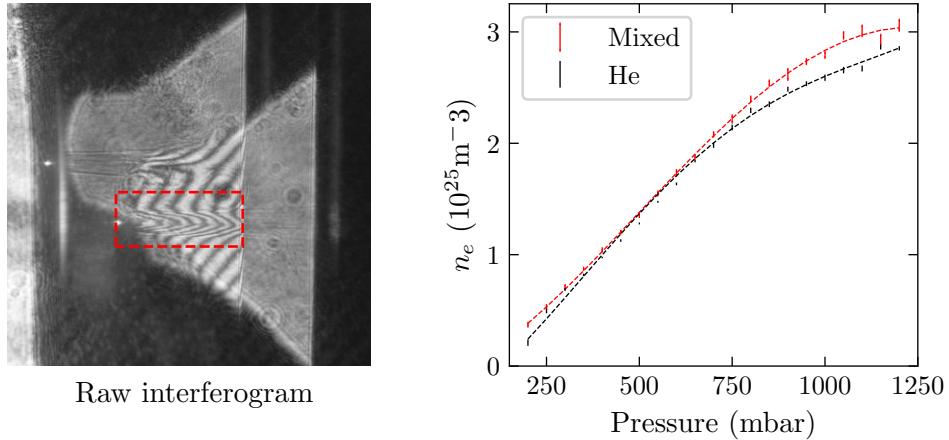


Figure 4.3: The Left panel shows an example interferogram, taken at 1000 mbar, with the red dashed lines indicating the location of the plasma. The right panel shows the measured relationship between pressure and density, where a polynomial has been used to fit the data. Each data point represents the mean of the interquartile range of measurements from 10 shots at that pressure, while the error bars show the standard error, given by σ/\sqrt{N} , where σ is the average error on the measurements and N is the number of samples.

4.1.2 Electron spectrometer

The image of the scintillating screen on the electron spectrometer camera was transformed to a flat, spatially calibrated rectangle using a projective transform. Spectral calibration was performed by tracking the trajectories of energetic electrons through a 2D map of the dipole magnet onto the screen, the results of which are plotted in Fig. 4.4. The charge calibration was done using BAS-TR image plate placed in front of the Lanex screen, and taking 100 shots to allow a large amount of charge to be detected. The absolute charge measured by the image plate⁹ was then compared to the number of counts detected by the camera to provide a conversion coefficient that was used for the data presented here.

⁹ [225] J. Strehlow *et al.* (2019) *Review of Scientific Instruments*

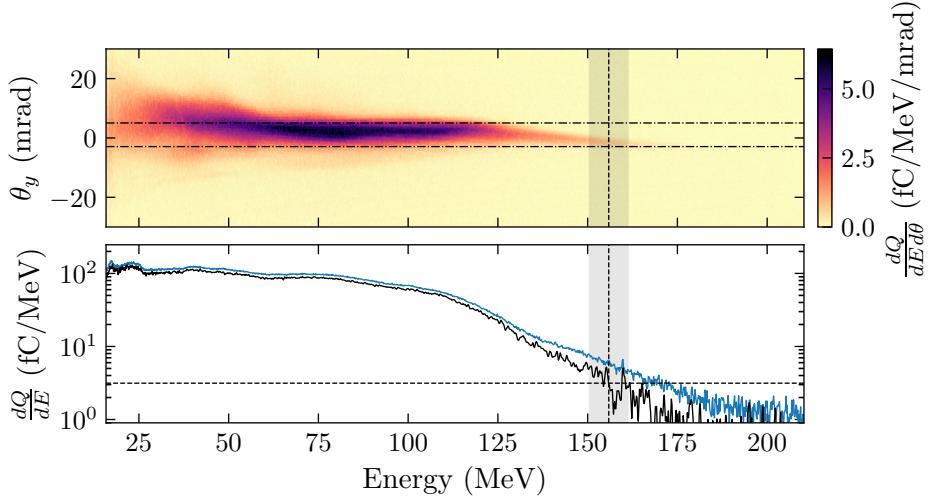


Figure 4.5: Example of electron spectrometer processing. In the upper panel, the dashed horizontal lines indicate the charge weighted FWHM divergence of the beam, while the horizontal line in the lower panel shows the noise floor. The vertical dashed line and grey shaded region are the maximum energy and associated error. The spectrum used for calculating the total charge is shown in blue, while the maximum energy was calculated using the spectrum shown in black, where a gradient background subtraction has been performed.

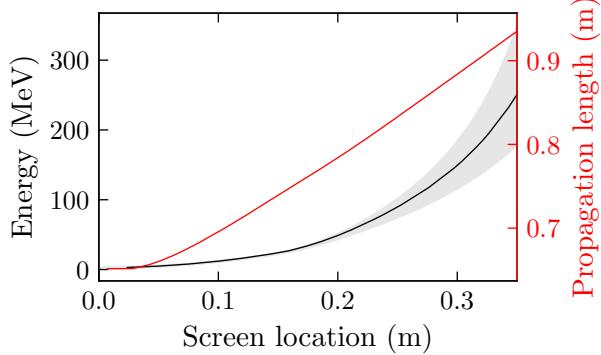


Figure 4.4: Electron energy and propagation length as a function of distance along the Lanex screen. The grey shaded region indicates the change in energy associated with a ± 5 mm translation along the dispersion direction of the electron beam, corresponding to a divergence of 18 mrad.

The upper panel of Fig. 4.5 shows a processed electron spectrometer image, and the lower panel shows two measurements of the spectrum after performing slightly different background subtractions. We define the maximum energy as the highest energy at which the spectrum rises above 3σ , where σ is the standard deviation of the background intensity. Due to light scattering inside the coffin, the value of the background intensity was affected by the brightness of the scintillation light emanating from the Lanex, resulting in a signal-dependent background that was not completely removed when the background subtraction was performed. In the plot of the spectrum, this light contamination takes the form of a pedestal, which complicates the calculation of a maximum energy. To account

for this, a second spectrum was calculated where a gradient background subtraction was performed in addition to the normal background subtraction, which eliminated the pedestal. This second spectrum was only used for calculating the maximum energy, and not for the measurement of the total charge. The value for σ was then determined by taking the standard deviation of the spectrum above 200 MeV and at divergences of $\theta_y > \pm 20$ mrad, a region usually devoid of signal. The error in the maximum energy was calculated by finding the maximum energy using both 2σ and 4σ as thresholds, and then taking the average difference between these values and the maximum energy measured using a 3σ threshold. This method of error analysis captures the main uncertainty in this measurement, namely, the true value of σ , as the maximum energy is more sensitive to this parameter than any other.

The divergence was calculated by splitting the spectra into 10 MeV slices from the bottom of the spectrum up to the maximum energy, determining the FWHM divergence of each slice, and then taking the charge weighted average over each of the slices. The error in this measurement is simply the standard deviation of each of the FWHM measurements. In some instances, beams were produced that completely filled the spectrometer Lanex, such that sometimes the divergence may be underestimated.

4.1.3 X-ray spectrum measurements

For this experiment, the low laser power led to the anticipation of both a low number of photons and a low critical energy, with the latter expected to be in the region of 5 keV. An Andor iKon-M 934 BR-DD camera¹⁰ was used; since this is a direct detection camera, it should be well suited to measuring low intensity, low energy x-rays¹¹.

In conjunction with this camera, a custom filter pack was designed for this experiment. Rectangular strips of each material with widths of 1.96 mm and lengths between 4.12 mm and 7.24 mm were placed in a mount in front of the camera sensor, and the relative signal detected behind each material could then be used to estimate the critical energy. In choosing the filters and their surface area, there existed a trade-off between choosing enough filters so that the spectrum could be measured with a sufficient number of data points, and maximising the number of photons detected per filter, which determines the error on each data point. Since there are expected to be a small number of photons, the filter areas needed to be relatively large to make any meaningful measurement, which resulted in only a small number of filters. Consequently, optimal filter selection was crucial to the spectrum measurement. An optimisation procedure was written to perform this selection, the details of which are given here¹², as are the details of the x-ray analysis procedure, performed by J.-N. Gruse. The final filter material and thicknesses are given in Tab. 4.1. The x-ray spectrum critical energy was retrieved by calculating the expected transmission through each filter for a range of critical energies and then choosing the

¹⁰ [226] *iKon-M 934 BR-DD information*, Andor

¹¹ The K-edge of silicon is 1.89 keV, and direct detection cameras are approximately 100 times more sensitive than comparable indirect detection cameras.

¹² [227] J.-N. Gruse *PhD. Thesis* (2020)

Table 4.1: Selected filter materials for the x-ray camera filter pack, and their thicknesses in μm .

Kapton $\text{C}_{22}\text{H}_{10}\text{N}_2\text{O}_5$	Mylar $\text{C}_{10}\text{H}_8\text{O}_4$	Mg	Al(95%) Mg(5%)	Al(98%) Mg(1%) Si(1%)	W
12.9	21.85	20.2	29.8	33.5	53.6

energy that minimised the χ^2 difference relative to the measured transmission. The error on the critical energy was then estimated by using the standard error on each filter transmission value to calculate low and high transmission bounds, which were then used to calculate lower and upper critical energies. The reported critical energy error is the average difference between these bounds and the centre critical energy.

4.1.4 Initial optimisation

Before performing the density-length parameter scans, the accelerator was first optimised for both the pure helium and mixed gas species to ensure that the scans were performed using parameters that produced a large amount of charge. On most LWFA experiments, optimisation is done entirely manually, and usually takes the form of a series of 1D parameter scans, where the coordinate that produced the best outcome is fixed for subsequent 1D scans. This method suffers from various issues, such as being overly costly and poorly suited to optimising coupled parameters. Instead, after brief manual optimisation to get an initial signal, a 5D Bayesian optimisation was performed, where the dimensions varied were the 2nd, 3rd, and 4th order spectral phase terms of the laser ($\beta^{(2)}$, $\beta^{(3)}$, $\beta^{(4)}$), the vacuum focal position, and the density. The length of the cell was fixed at 1 mm, which was done to reduce the dimensionality of the optimisation, resulting in faster convergence. The $\beta^{(i)}$ refer to the coefficients pulse's spectral phase $\beta(\omega^*)$, where $\omega^* = \omega - \omega_0$,

$$E(\omega) = E_\omega e^{-\left(\frac{\omega^*}{\Delta\omega}\right)^2} e^{-i\beta(\omega^*)}, \quad (4.1)$$

where $\Delta\omega$ is the Gaussian bandwidth and

$$\beta(\omega^*) = \beta^{(0)} + \beta^{(1)}(\omega^*) + \beta^{(2)}\frac{(\omega^*)^2}{2} + \beta^{(3)}\frac{(\omega^*)^3}{6} + \beta^{(4)}\frac{(\omega^*)^4}{24} + \dots. \quad (4.2)$$

The zeroth and first-order spectral phase terms relate to a shift of the phase with respect to the envelope and a shift of the pulse with respect to time, which have negligible effects since the spatial scale of the pulse envelope is much longer than the laser wavelength. However, the addition of $\beta^{(2)}$ such that $\beta(\omega^*) = \beta^{(2)}(\omega^*)^2/2$ results in a time-varying electric field with the form

$$E(t) = \frac{E_\omega}{2\sqrt{\pi\left(\frac{1}{(\Delta\omega)^2} - i\frac{\beta^{(2)}}{2}\right)}} \exp\left[i\omega_0 t - \left(\frac{t^2 + i\frac{\beta^{(2)}(\Delta\omega)^2 t^2}{2}}{(\Delta\omega)^2\left(\frac{1}{(\Delta\omega)^4} + \frac{(\beta^{(2)})^2}{4}\right)}\right)\right] \quad (4.3)$$

We see that there is now a quadratic term in the phase resulting in a chirp, and the pulse length has increased from $(\Delta\omega)^{-1}$ to a length that depends on $\beta^{(2)}$. The field strength has dropped commensurately. Similarly, adding third and fourth-order spectral phase terms results in cubic and quartic terms in the phase, and also modifies the pulse length and peak field. It should be noted that the addition of $\beta^{(3)}$ modifies the pulse shape asymmetrically, in contrast to the $\beta^{(2)}$ and $\beta^{(4)}$ terms.

At each coordinate, a burst of 10 shots was taken at a repetition rate of 1 Hz, after which the average number of counts on the electron spectrometer camera was used as the optimisation goal, with the error being the standard deviation over the burst. Optimisations were halted once a maximum was reached and the algorithm began to search for a new regime to exploit; this point was determined manually. More details of this procedure are given by Shalloo *et al.*¹³; the data for this publication was obtained during the experiment described in this chapter.

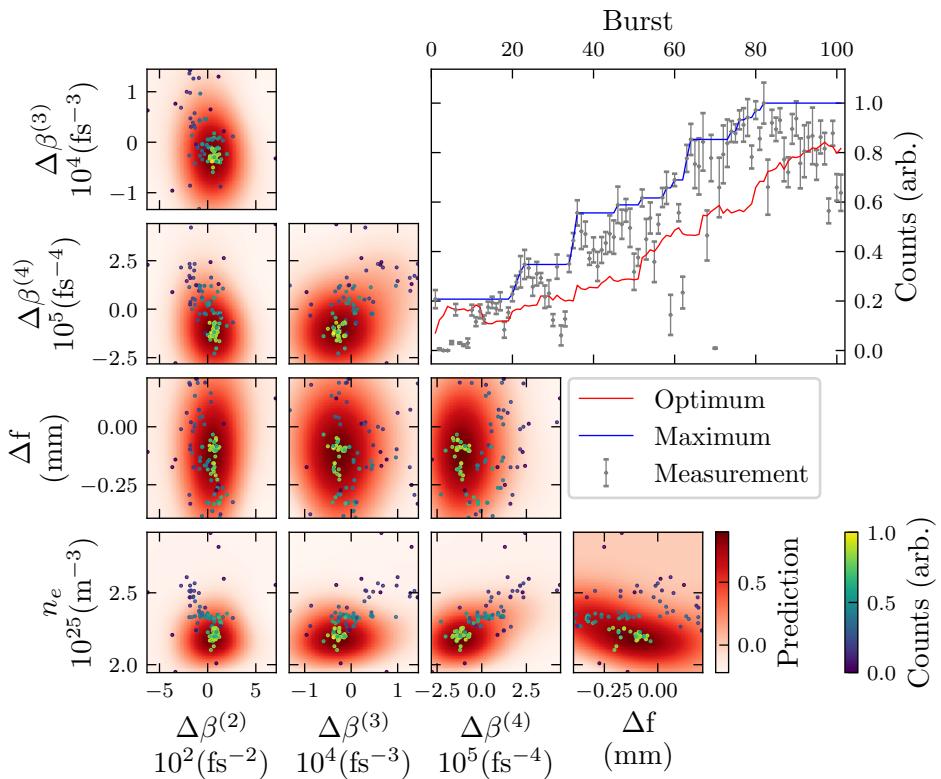


Figure 4.6: Bayesian optimisation of the accelerator using pure helium gas. The upper right plot shows the measurements as a function of the burst number, where each data point is the average of all the data collected during that burst, and the error bar is the standard deviation of the measurements in each burst. The blue line, the ‘Maximum’, refers to the maximum value recorded up to that burst, while the red line, the ‘Optimum’, shows the maximum value of the fitted Gaussian process.

Fig. 4.6 shows a summary of the pure helium gas optimisation. The 10 small scatter plots on the lower left side of the plot show the coordinates at which measurements were

¹³ [220] R. J. Shalloo *et al.* (2020) *Nature Communications*

made, while the marker colour is determined by the number of counts detected at that coordinate. The term $\Delta\beta^{(i)}$ is the change of the i^{th} spectral phase coefficient, relative to the average value of the coordinate, while Δf gives the change in focal position, where positive values correspond to focusing further from the parabola. In each of these plots, as the number of counts increases, the location of the measurements can be seen to gravitate towards a local maximum, indicating the algorithm has found and is attempting to exploit this maximum. During the optimisation, a 5D Gaussian process, $\mathcal{GP}^{(5)}$, was fit to the data and acts as a surrogate model of the accelerator. The 2D colour maps in the background of the scatter plots show the model prediction, where a 2D slice has been taken through the $\mathcal{GP}^{(5)}$ at the optimum position. For example, the upper left plot of $\Delta\beta^{(3)}$ versus $\Delta\beta^{(2)}$ calculates the value of $\mathcal{GP}^{(5)}(\Delta\beta^{(2)}, \Delta\beta^{(3)} | \Delta\beta_{\text{opt}}^{(4)}, \Delta f_{\text{opt}}, n_{e,\text{opt}})$.

The plot in the upper right of the figure shows the measurements and their errors displayed chronologically as a function of burst number. Note that some values consistently produced signal roughly equivalent to the background, resulting in a low measurement with a small error bar. This did not skew the model, because the measured variance at each point was used to construct a second Gaussian process, used to model the true uncertainty at each position, and this second process was used to provide a less noisy estimate of the error. This was done because the measured shot-to-shot fluctuations provide an inaccurate measurement of the variance at that position, given the small number of samples. See the above publication for more details. The predicted optimum, shown in red, is an estimate of the optimum value of $\mathcal{GP}^{(5)}$, calculated iteratively as the result of each burst is supplied to the model. On the other hand, the rolling maximum, shown in blue, is simply the maximum value measured up to a given burst. The predicted optimum of the model does not match the maximum value recorded up to that point because the measurements have errors, and the model has a pre-defined range of scale lengths. Setting up the \mathcal{GP} with a relatively restrictive lower bound on the scale length of the variation is useful in a noisy environment, as it prevents over-fitting and encourages a more gradual but reliable optimisation.

The helium source reached a local optimum after 80 bursts, following which 20 more bursts were taken where the accelerator performance decreased, indicating that the algorithm had converged and had begun to search for a new regime. Since the number of counts increased steadily with the burst number, the path the optimisation took through the parameter space can be seen in the scatter plots. While for most parameters a simple linear shift is seen, such as in the plots of n_e versus $\Delta\beta^{(i)}$, the plots involving Δf follow a slightly more convoluted path from the starting coordinate to the optimum, starting at $+0.2$, then shifting to -0.3 , before arriving at approximately -0.1 .

The optimisation of the doped (mixed) gas is shown in Fig. 4.7. This optimisation appeared to converge quicker than the helium run, as can be seen by the swift rise in the number of counts between bursts 30 and 40, followed by the decline after burst 50, and a levelling of the predicted optimum. As a result, this optimisation was terminated after 55 bursts. In contrast to the helium optimisation, which saw a gradual increase in the

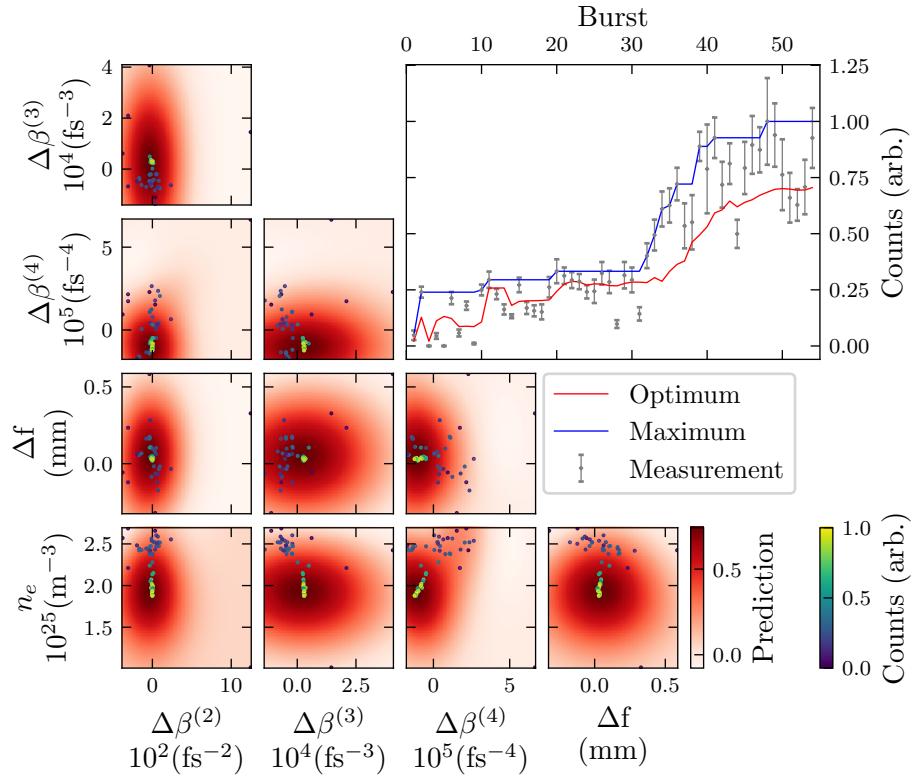


Figure 4.7: Bayesian optimisation of the accelerator for the mixed gas.

number of counts, this optimisation took longer to find an improved regime. After the initial bursts taken at random coordinates (1-5), the maximum value did not improve significantly until burst 30. This is likely due to the different starting points of the optimisations, with the mixed gas starting close to a local optimum, while the helium gas started on a slope, making it easier to find improvements. This can be seen in the n_e vs $\Delta\beta^4$ plot Fig. 4.7, in which two peaks in the \mathcal{GP} can be seen. The smaller number of bursts contributes to the smoother prediction of the \mathcal{GP} , resulting in a smoother kernel being used to construct the surrogate model.

4.2 Density-length scans

The density-length scans were performed immediately after the optimisations. Scans were performed by setting the cell length to 0 mm, and then increasing the backing pressure to the cell from 200 mbar to 1200 mbar in 21 steps. This direction was chosen because increasing the backing pressure required little time and did not require frequent venting, reducing wasted gas. Pressure scans were performed at cell lengths up to 3 mm in increments of 0.2 mm, for a total of 16 cell lengths. At each combination of pressure and length, 10 shots were attempted, for a total of 3,360 shots per scan. The laser failed to produce a beam above 200 mJ on 92 of the 6,720 shots, a failure rate of 1.4%, and consequently, these shots were removed from the analysis. The median laser energy for the helium scan after the failed laser shots were removed was 251.2 ± 7.3 mJ, while 259.0 ± 8.2 mJ was the median energy during the doped gas scan. Laser energy as a

function of shot number is plotted in Fig. 4.8, where it can be seen that no large drifts in the laser energy are observed during the run.

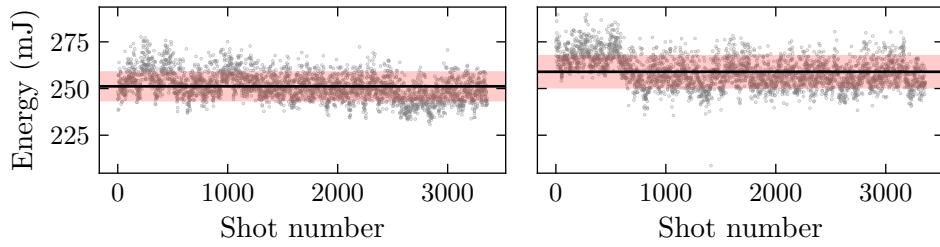


Figure 4.8: Laser energy delivered to target during the helium (left) and doped (right) parameter scans. The horizontal line and red shaded regions show the median and standard deviation respectively.

4.2.1 Electron spectrometer measurements

First, we will review the measurements made by the electron spectrometer. Rather than present individual electron spectrometer images, due to the quantity of data, single-valued metrics of the electron beams will be used to characterise the spectra individually. The average of these metrics over a burst has been calculated, with the results of each burst represented as a single data point on a grid. This allows for a macroscopic, qualitative analysis of the full parameter space to be performed, facilitating the identification of regions with interesting behaviour. These interesting regions will then be investigated quantitatively in Sec. 4.4.

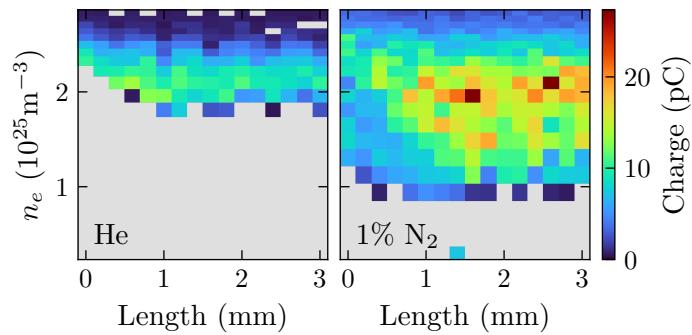


Figure 4.9: Mean electron beam charge for pure helium gas (left) and 1% nitrogen-doped gas (right) in the density length parameter space. The grey colour indicates regions where the charge fell below 0.5 pC. The anomalous low-density measurement at 1.4 mm in the doped scan is due to imperfect venting of the reservoir.

Fig. 4.9 compares the mean electron beam charge for the pure and doped gas density-length scans. Immediately apparent is the larger amount of charge injected by the doped gas, up to a maximum of 27.3 ± 4.0 pC, while the maximum charge in the pure gas scan was 13.5 ± 1.5 pC, where the error is calculated using the standard error, σ / \sqrt{N} , where

N is the number of samples and σ is the standard deviation of those samples. In addition, electron beams are produced over a wider range of parameters for the doped gas, with injection occurring at densities as low as $9.4 \times 10^{24} \text{ m}^{-3}$, and at the minimum cell length at $1.30 \times 10^{25} \text{ m}^{-3}$, while the helium gas required densities above $1.81 \times 10^{25} \text{ m}^{-3}$ for any injection to occur. For both gases, above approximately $2.2 \times 10^{25} \text{ m}^{-3}$, increasing density results in a decrease in charge, independently of the cell length.

One of the features of these plots that will receive the most attention in the following analysis is the coupling between the density and the length at which injection occurs. This can be seen in the slopes on the left-hand side of the plots, where, for lower densities, a longer cell length is required for injection to occur. In pure helium gas, injection is caused by wavebreaking of the plasma wave, while, in the doped gas, inner shell ionisation provides an additional mechanism for injection, and both these processes depend on the dynamics of the driving laser. Consequently, analysis of the injection point provides insight into the pulse evolution rate.

The energy below which 95% of the charge in the spectrum exists is a robust metric for characterising the energy of the beam, as it does not attempt to measure precisely where the signal ends and the background begins, and is therefore well suited to measuring the quasi-thermal electron spectra which are common during these scans. It is calculated by integrating the spectrum from its lowest energy up to the point at which the charge is equal to 95% of the total charge. This metric for the two gas mixtures is shown in Fig. 4.10.

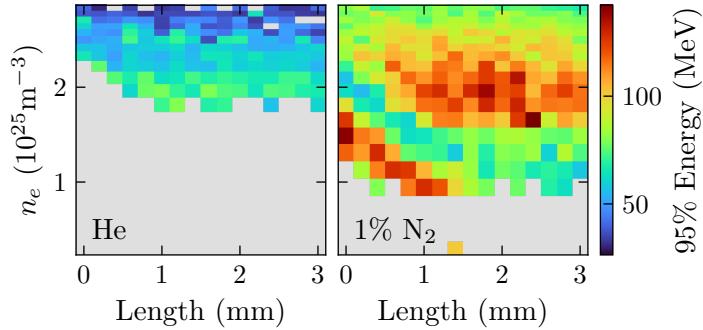


Figure 4.10: Mean 95% cut-off energy of each burst. The regions where the charge did not exceed 0.5 pC have been greyed out, as the spectra in these regions would occasionally produce spurious high energy readings due to noise.

For helium gas, the highest 95% energy value did not exceed 80 MeV, and this metric was optimised at a density of $1.8 \times 10^{24} \text{ m}^{-3}$, with a weak length dependence, provided a beam was injected. When we compare Fig. 4.10 to the charge measurements in Fig. 4.9, we see that the highest energy beams occur at slightly longer lengths and lower densities than the highest charge beams. These two locations in the parameters space are approximately equivalent in terms of the conversion efficiency from laser beam energy to

electron beam energy, indicating the presence of a Pareto front¹⁴, where improvements in one metric come at the cost of another.

The mixed gas plot in Fig. 4.10 has two distinct regions where high energies were recorded; one at low densities and short cell lengths, and another at higher densities and longer cell lengths. In both of these regions, the cut-off energy exceeds 120 MeV, significantly higher than the highest values recorded in the helium gas. In between these two regions, there appears to be a ‘valley’ in the parameter space, where the beam energy decreases before increasing again. Comparing the doped gas results to those from the pure gas, the region of the parameter space occupied by the pure gas beam is similar to the location of the second high-energy region in the doped gas. Taken together, these observations strongly indicate two injection events; an ionisation injection that occurs at shorter plasma lengths and lower densities, and a second injection event that occurs in the same place as self-injection in the pure gas.

While the 95% energy metric is useful for characterising the beam, we would also like to measure the maximum energy that was produced at each coordinate, as this provides quantitative insight into the acceleration dynamics. However, as shown by Fig. 4.4, the error in the energy due to divergence fluctuations becomes significant above approximately 100 MeV, and so care needs to be taken when determining the maximum energy. Simply using the maximum energy detected would not produce a useful result, as the error would be over 50% for the higher energy shots. Another option would be to take the mean maximum energy of the burst, which would mitigate errors due to pointing fluctuations. However, due to shot-to-shot variations in the laser, the peak energy varies significantly across the burst, so this mean value underrepresents the maximum energy. As a compromise, we will use the mean of energies with non-zero values that are above the 25th percentile of the burst, as this should remove the low-quality laser shots while allowing for enough measurements to mitigate the error from pointing fluctuations.

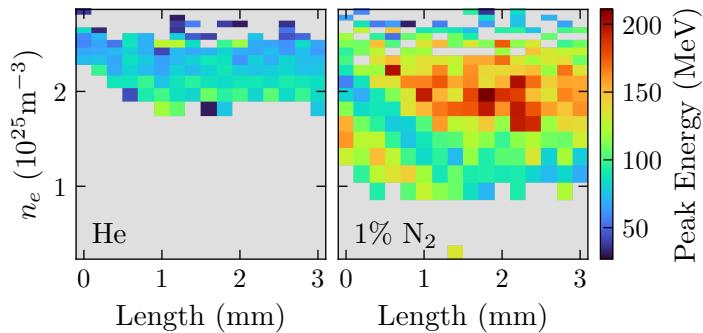


Figure 4.11: Mean peak energy of each burst. The regions where the charge did not exceed 0.5 pC have been greyed out, as the spectra in these regions would occasionally produce spurious high energy readings due to noise. Additionally, the settings that failed to produce a peak energy measurement due to low signal are also shown in grey.

¹⁴ [202] F. Irshad *et al.* (2022) *arXiv e-prints*

These peak energy measurements are shown in Fig. 4.11. The highest energy reached by the pure gas was 125 ± 33 MeV, while the doped gas had an estimated peak energy of 211 ± 7 MeV. The error on the pure gas measurement is higher than the mixed gas because the highest energy burst from the pure gas scan produced only one measurement that met the selection criteria, and, as a result, the error is taken from the error due to pointing fluctuation at that energy, shown as the shaded region in Fig. 4.4. Qualitatively, the doped gas measurements in Fig. 4.11 are similar to those in Fig. 4.10, with two distinct regions of the parameter space producing high-energy beams. The main difference is that the two regions have different peak energies, despite having similar 95% cutoff energies. This occurs due to a difference in the shapes of the spectra; since the 95% metric is insensitive to long tails, we can deduce that the high peak energies are the result of long-tailed spectra.

The average divergences are plotted in Fig. 4.12. For the pure helium gas, the higher electron beam energies coincided with the lowest divergence beams, which were mostly in the range of 15 to 20 mrad. This is to be expected, as electron beams with equivalent transverse momentum spreads will have lower divergences if the energy is higher. The beams produced in the doped gas at short cell lengths and low plasma densities were the most collimated, down to a minimum divergence of 1.4 ± 0.1 mrad, produced at $9.4 \times 10^{24} \text{ m}^{-3}$ at a length of 1 mm. Conversely, the electron beams produced at longer lengths and higher plasma densities in the mixed gas had higher divergences than those produced in the same region of the parameter space in the pure gas. This is despite the higher beam energy, which would be expected to reduce the divergence.

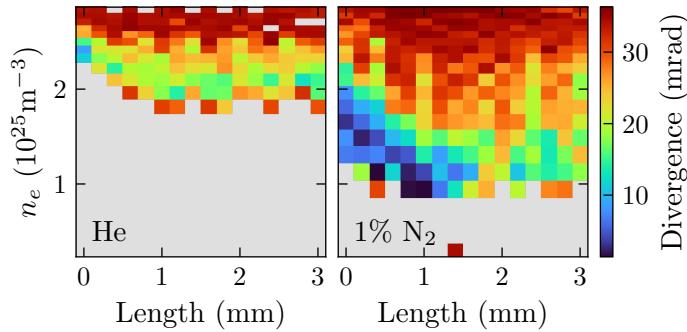


Figure 4.12: Mean divergence of each burst. Regions with charge below 0.5 pC are coloured grey, as these divergence measurements were overly susceptible to noise.

Having examined the parameter space in full, it is worth displaying a selection of the raw data. The shots displayed in Fig. 4.13 constitute the highest charge burst from the pure helium scan, Fig. 4.14 shows the lowest divergence electron beams from the doped gas scan, and Fig. 4.15 shows the high peak energy measurements, also from the doped gas. Note that the three figures all use the same colour and axes scales. Displaying the raw data is important because it reveals information that cannot be deduced from the single metric plots above, such as the interesting structure of the high charge beams

injected in helium gas, and the fact that the highest energy electrons exceed the upper energy threshold of the spectrometer.

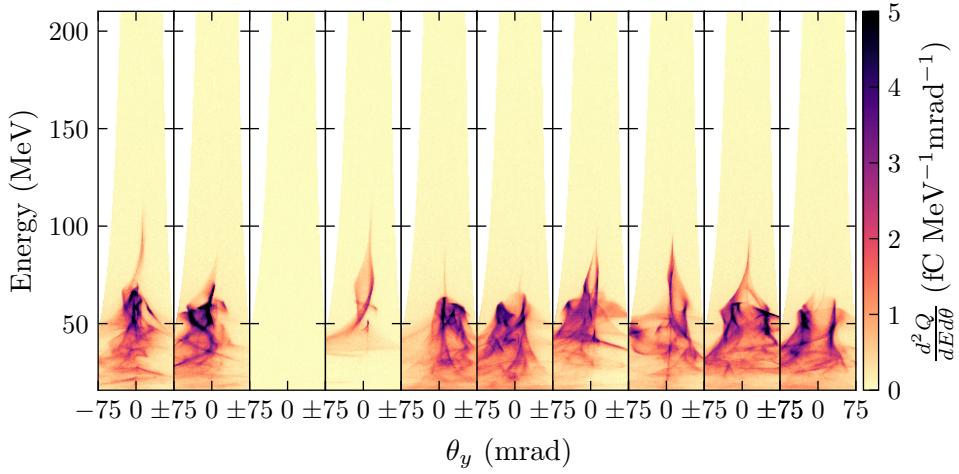


Figure 4.13: Electron spectrometer images from the peak charge burst in pure helium gas, $n_e = 2.1 \times 10^{25} \text{ m}^{-3}$ with a cell length of 1.0 mm. The horizontal axis is the divergence and spans $\pm 75 \text{ mrad}$. A wider divergence range is visible at lower energies due to the shorter propagation length. Empty images show instances where the laser failed to produce an electron beam.

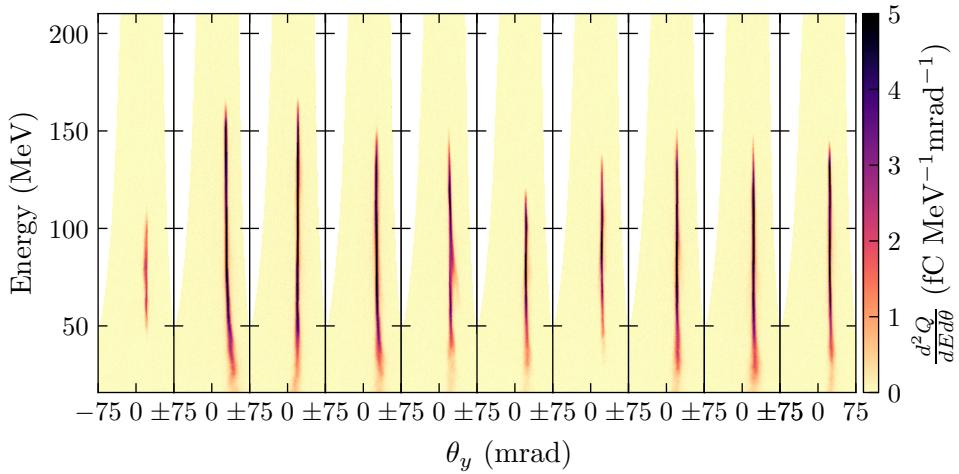


Figure 4.14: Lowest divergence beams produced in the doped gas, $n_e = 0.95 \times 10^{25} \text{ m}^{-3}$ with a cell length of 1.0 mm.

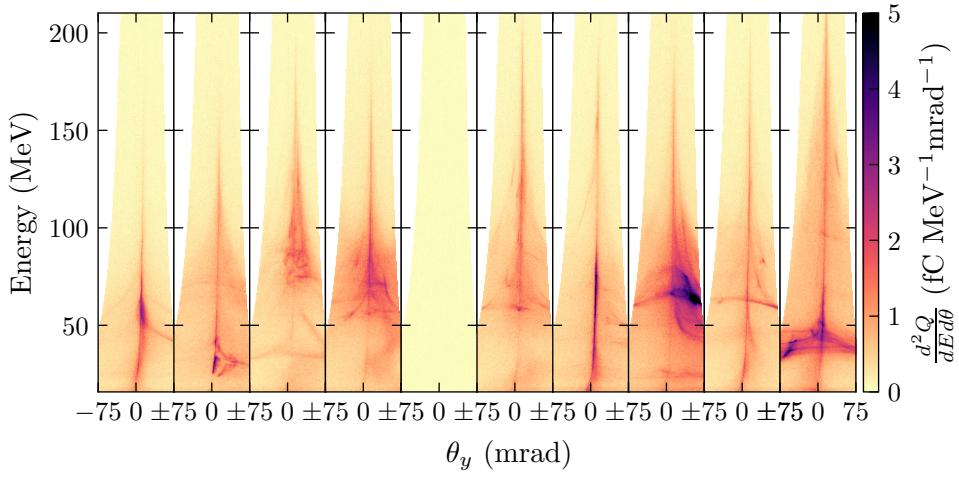


Figure 4.15: Electron spectrometer images from the burst with the highest peak energy, at 1.6 mm, and $n_e = 2.0 \times 10^{25} \text{ m}^{-3}$. The highest energies exceeded the upper energy limit of the spectrometer.

4.2.2 X-ray measurements

A measure of the x-ray beam intensity is shown in Fig. 4.16; the mean peak of the background subtracted counts on the x-ray camera during each burst. Due to shot-to-shot fluctuations, the x-ray intensity varied significantly within each burst, and so, to minimise the error on these measurements, the best 60% of shots from each burst are selected and used for these measurements. Since the error is calculated from the standard deviation across the burst, this reduces the error at the cost of slightly biasing the results. As shall be seen, the error on the x-ray measurements is significant, and so minimising the error was a priority.

Electrons generated during the second injection event in the doped gas appear to be responsible for the majority of the observed x-rays, while the electron beam in the helium gas can only produce signal that rises above the background level over a narrow region. This is partly due to the increased charge and energy of the second beam in the doped gas, which combine to greatly enhance the x-ray flux. Interestingly, electrons from the first injection event in the doped gas fail to produce detectable x-rays, despite having comparable charge and higher energy than the x-ray-producing beam in the helium gas. This is likely to be due to the reduced oscillation amplitude of this beam, as can be seen from the divergence measurements in Fig. 4.12. The 98th percentile mean peak intensity values for the helium gas were 518 ± 2 counts, while for the doped gas these were 4890 ± 10 counts. The uncertainties reflect the standard error on the mean, which is low due to number of pixels, of order 10^3 , and the averaging over multiple shots.

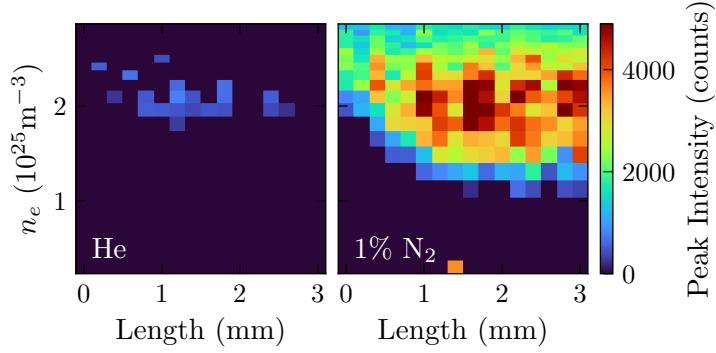


Figure 4.16: Comparison of the peak x-ray intensity.

Fig. 4.17 shows the measurements of the critical energy. The 98th percentile critical energies for the helium gas is 8.0 ± 1.2 keV, and 2.8 ± 0.2 keV for the doped gas. Despite the low signal, the pure helium target was capable of producing higher critical energies than the doped gas at the equivalent parameters. This is surprising because the critical energy scales most strongly with the electron energy, $E_{\text{crit}} \propto \gamma^2$, and the electrons in the doped gas are more energetic in this parameter region. A possible explanation is that the higher charge beam of the doped gas has a large number of low energy electrons that are contributing to the x-ray spectrum, reducing the measured E_{crit} . Alternatively, or perhaps additionally, from inspection of the images in Fig. 4.13, we can see that the beam is clearly undergoing large amplitude oscillations, which would increase the critical energy.

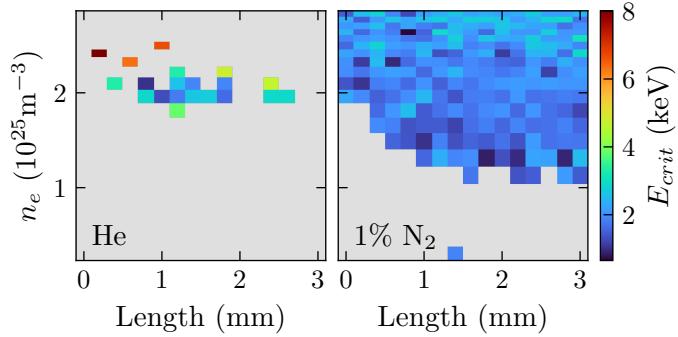


Figure 4.17: Comparison of the critical energy.

The number of photons within a 0.1% bandwidth at the peak of the x-ray spectrum, N_{ph} , is shown in Fig. 4.18, normalised to the unit solid angle. This is calculated by determining the peak number of photons measured per pixel from the peak intensity, dividing by the angular size of that pixel from the source, and then using this value to normalise the integral of the synchrotron spectrum. A 98th percentile value of $2.5^{+0.7}_{-0.1} \text{ ph mrad}^{-2} 0.1\% \text{BW}^{-1}$ was measured for the pure helium gas, while for the doped gas this value was $10.7^{+1.3}_{-1.6} \text{ ph mrad}^{-2} 0.1\% \text{BW}^{-1}$. The unit $0.1\% \text{BW}^{-1}$ means that the

reported N_{ph} is equal to the integration of the intensity spectrum around the energy with the peak intensity, occurring at $E = E_{\text{peak}}$, over a bin width of $E_{\text{peak}} \times 10^{-3}$. The asymmetry of the error results from using the error on the critical energy, ΔE_{crit} , to determine N_{ph} for two different critical energies, $E_{\text{crit}} \pm \Delta E_{\text{crit}}$, with the upper error resulting from using a lower critical energy.

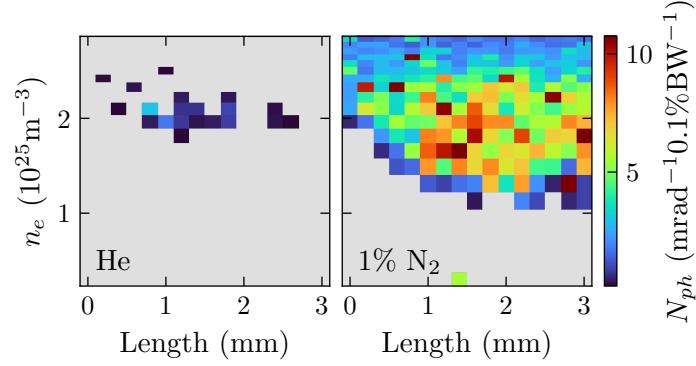


Figure 4.18: Number of photons at the peak of the spectrum, calculated using the peak intensity and the measurements of the spectrum.

4.3 Simulations

In this section, the results of the PIC simulations will be presented. Following an overview of the simulation parameters, the results from the pure helium and doped gas density scans will be presented. Explicit comparison to the experimental data will be delayed until the following discussion section.

4.3.1 Simulation parameters

The simulations were performed using the code FBPIC¹⁵ using the parameters shown in Tab. 4.2, using a single GPU in the infinite order spectral solver mode, making the simulations dispersion free in all directions. Identical parameters were used for both the pure and doped simulations, with the only difference being the presence of the dopant. The helium atoms were initialised with an ionisation state of +1, while the nitrogen atoms were ionised up to +5, which was done to reduce computational time, while the ionisation rates were calculated using the built-in module based on ADK ionisation rates¹⁶. The nitrogen ions were initialised with only 1 particle per cell as opposed to the 24 particles per cell used for the helium and plasma electrons because otherwise tracking the number of particles produced by the dopant exceeded the memory limits of the GPU.

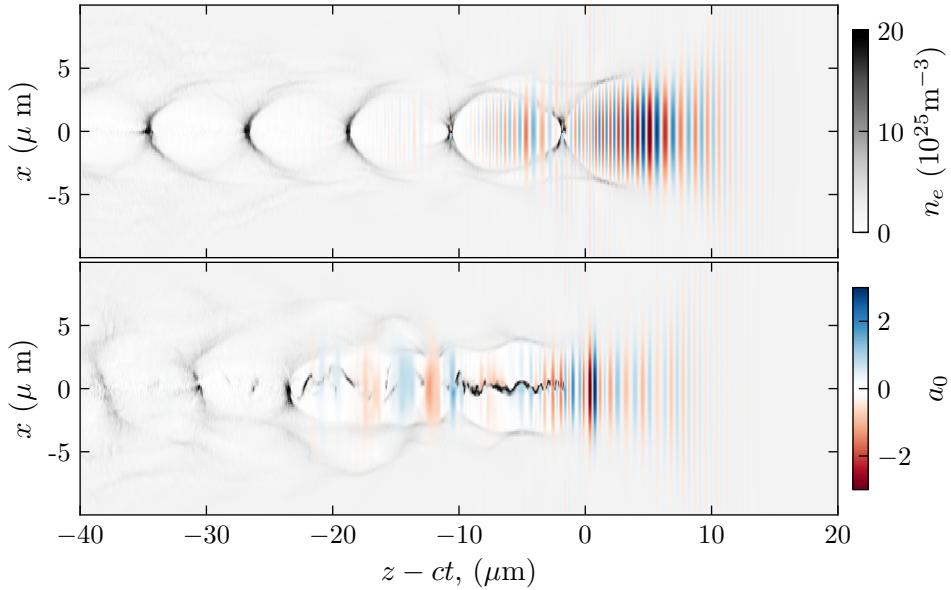
Density ramps with scale lengths of 0.25 mm were used for the transition between the bulk of the plasma and the vacuum. Such long ramps were used to simulate the effects of the large exit and entry holes to the cell, which had been widened to approximately 1 mm by laser ablation. The density profiles were initialised so that the front ramp reached

¹⁵ [216] R. Lehe *et al.* (2016) *Computer Physics Communications*

¹⁶ See Sec. 2.3.

Table 4.2: Simulation parameters

Parameter	Axis (Units)	Number
Window size	$z \times r, (\mu\text{m}^2)$	80×40
Resolution	$z \times r, (\lambda^{-1})$	35×2.5
Azimuthal modes	θ	2
e^-	$z \times r \times \theta$	$2 \times 2 \times 6$
He^{1+}	$z \times r \times \theta$	$2 \times 2 \times 6$
N^{5+}	$z \times r \times \theta$	$1 \times 1 \times 1$

Figure 4.19: Snapshots from the pure helium simulation at $n_e = 2 \times 10^{25} \text{ m}^{-3}$, $t = 4.7 \text{ ps}$, (top) and $t = 5.0 \text{ ps}$ (bottom); before and after the first injection event.

95 % of the plateau plasma density at $z = 2 \text{ mm}$, and the down ramp at the end of the plasma dropped to 95 % at $z = 5 \text{ mm}$.

The laser pulse injected into the simulations had a Gaussian transverse profile with $w_0 = 18.5 \mu\text{m}$ and a peak a_0 of 0.57, chosen to match the median values of the data presented in Fig. 4.2. The vacuum focal position was set to $z = 2 \text{ mm}$, which is over 1 Rayleigh length from the start of the simulation. It is for this reason that a Gaussian transverse profile has been used, as opposed to an experimentally measured one, as the spatial phase information of the focal spot would be required to propagate it accurately over such a long distance. The longitudinal intensity profile was taken from the pulse length measurements made with the SPIDER at the end of the optimisation process, which had $\tau_{\text{FWHM}} 49.4 \text{ fs}$, but required a 266 fs (80 μm) long window to contain the pre and post pulse structure. Two snapshots from one of the simulations are shown in Fig. 4.19. Taken from the moments before and after the first injection event in $n_e = 2 \times 10^{25} \text{ m}^{-3}$ pure helium, these plots illustrate the relative scales of the plasma bubble size, the laser wavelength, and the laser pulse length.

4.3.2 Simulation results

The evolution of the electron energy spectrum for the densities $1.25 - 2.5 \times 10^{25} \text{ m}^{-3}$ are shown in Fig. 4.20 for the pure helium gas. For each density, an initial period of pulse propagation is followed by the injection and rapid acceleration of electrons. These first electrons follow a roughly parabolic trajectory in the energy- z space, characteristic of a beam that is undergoing acceleration followed by dephasing. The curvature radius and maximum energy of the parabolic trajectory appear to be inversely related to the density, with the beam from the lowest density simulation following the smoothest, highest path, while the first injection in the highest density simulation dephases more rapidly and at a lower energy.

Following this first beam, a second beam is injected for $n_e > 1.25 \times 10^{25} \text{ m}^{-3}$. Rather than following a curved trajectory, this beam undergoes constant acceleration, before plateauing abruptly. This is consistent with beam-driven acceleration, where the witness beam does not dephase from the driver. The maximum energy of this second bunch is sensitive to the charge of the driver bunch, the location at which the laser depletes, and the injection location, all of which are the result of highly non-linear processes individually. Consequently, the maximum energy reached by this second beam does not appear to follow a pattern, as the process that produced it is semi-chaotic.

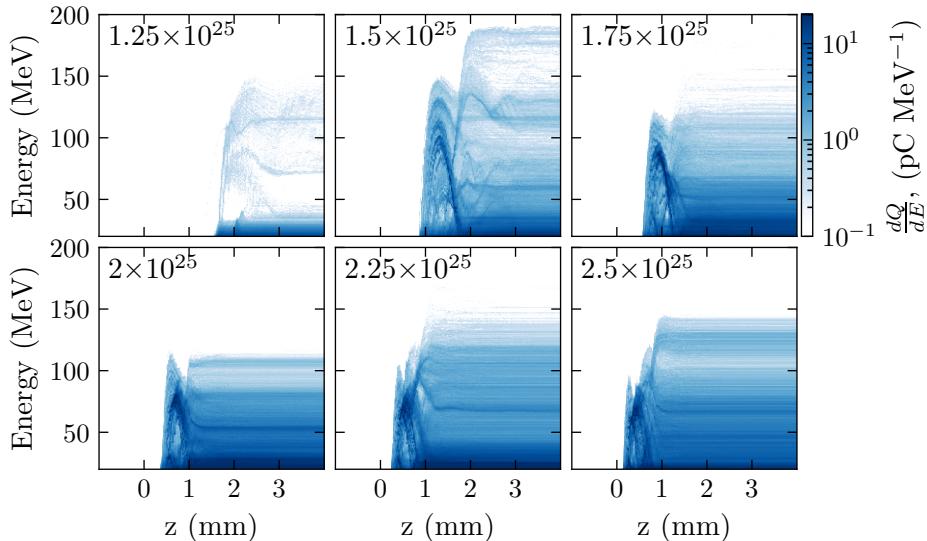


Figure 4.20: Electron beam energy spectrum evolution as a function of longitudinal position z for the pure helium gas at various densities, where a logarithmic colour scale has been used. The simulation for $n_e = 1 \times 10^{25} \text{ m}^{-3}$ did not inject any charge and has been omitted.

The evolution of the energy spectrum for the nitrogen-doped simulations is shown in Fig. 4.21, now for the density range $0.75 - 2.0 \times 10^{25} \text{ m}^{-3}$. In these plots, the electrons which originate from the bulk plasma are shown in the same colour map as in Fig. 4.20, while those which are produced by the ionisation of the N^{5+} ions are shown in red. We see that ionisation injection occurs for densities as low as $0.75 \times 10^{25} \text{ m}^{-3}$, while the bulk

plasma electrons are only injected once the density reaches $1.25 \times 10^{25} \text{ m}^{-3}$, which is the same threshold as that found in the pure helium simulations. The first bunch of electrons injected from the bulk plasma also reach similar energies as those in the first bunches in the pure helium simulations, indicating that the ionisation injected beam has dephased sufficiently so as to not suppress the accelerating field experienced by the self-injected beam. The onset of beam-driven acceleration is also seen to occur at the same density, but appears to reach higher energies for the nitrogen-doped simulations than for the comparable pure simulations. We therefore conclude that, prior to the laser depleting, the ionisation injected beam and self-injected beam do not strongly affect each other, but that once the laser has depleted, the witness bunch energy is sensitive to their combined effects.

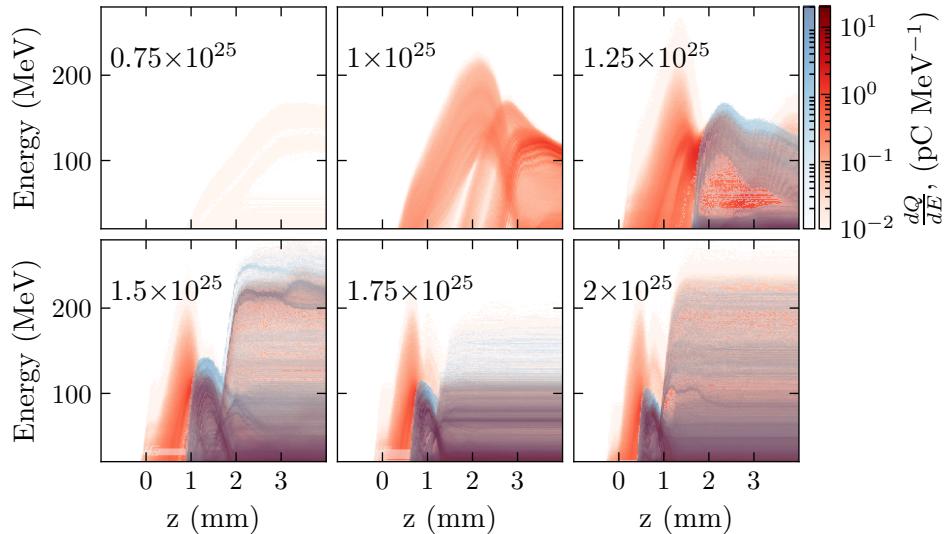


Figure 4.21: Electron beam energy spectrum evolution as a function of longitudinal position z for the nitrogen-doped gas. The simulation for $n_e = 0.5 \times 10^{25} \text{ m}^{-3}$ did not inject any charge, and has been omitted. The red colour denotes the electrons that originate from the inner shells of nitrogen.

The injection, acceleration, and dephasing of the beam is sensitive to the evolution of a_0 , which is plotted in Fig. 4.22. At all densities greater than or equal to $1 \times 10^{25} \text{ m}^{-3}$, a similar evolution is observed; an initial period of increasing a_0 is followed by a plateau, after which a second increase is observed up to the peak value, after which the intensity drops rapidly. The first increase is due to the transverse self-focusing of the pulse as it enters the plasma, while the second increase is due to longitudinal compression, which is soon followed by depletion. The results in the pure and doped gases are plotted side by side, where it can be seen from inspection that the dopant is not having a measurable effect on the dynamics. Some small differences between the two gases are seen for the higher density simulations; the $1.5 \times 10^{25} \text{ m}^{-3}$ simulation has a higher peak a_0 in the doped gas than the pure gas, while the peak a_0 is higher in the pure helium gas simulation at $2 \times 10^{25} \text{ m}^{-3}$. However, these are only seen at the peak a_0 values, which are noisy to measure due to the few-cycle length of the pulse when fully compressed.

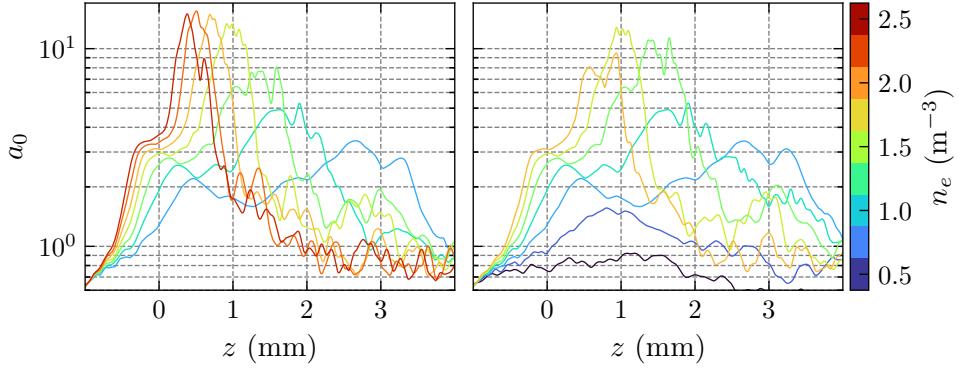


Figure 4.22: Maximum value of the laser a_0 as a function of z , for the pure helium simulations, left, and the doped simulations, right.

Comparing the a_0 evolution plots to the results in Fig. 4.20 and Fig. 4.21, we can see that the first self-injection events occur when $a_0 \approx 3$, while ionisation injection starts¹⁷ when $a_0 \approx 2$. The initial transverse focusing increases a_0 sufficiently for the first ionisation injection event to occur, while the first self-injection event requires the pulse to undergo some longitudinal compression. Further self-injection occurs when the pulse compresses rapidly, after which the onset of beam-driven acceleration aligns with the depletion of the laser pulse.

To provide a quantitative comparison between the experimental 2D scans and the simulations, various electron beam metrics have been measured as a function of longitudinal position, and have been plotted in Fig. 4.23. These metrics were the charge above 27 MeV, the 95% cut-off energy, and the divergence, θ_y . The charge and divergence measurements were made by selecting particles with divergences in x less than 18 mrad and energies above 27 MeV, as these are the only particles which would reach the electron spectrometer screen on the experiment. In contrast, the energy metric is calculated from the 95% energy contours in Fig. 4.20 and Fig. 4.21, which does not set any requirement on the divergence. The fast changes seen in the charge and divergence metrics are a result of this stipulation, as the divergence of the beam is undergoing rapid changes during injection.

From inspection, the simulated data resembles the experimental measurements closely. Firstly, the dependency between density and injection length appears to be similar, for both the self-injected and the ionisation-injected beam, suggesting that these simulations can be used to determine the mechanism behind the pulse evolution rate. In addition, the ionisation injection density threshold occurring somewhere between $0.75 \times 10^{25} \text{ m}^{-3}$ and $1 \times 10^{25} \text{ m}^{-3}$ is in close agreement with the experimental data, and the contour of the peak energy of the first ionisation injected beam is in close agreement with the measurements in Fig. 4.10. The lower divergence of the ionisation injected beam relative

¹⁷ Interestingly, we do see a small amount of ionisation for the $0.75 \times 10^{25} \text{ m}^{-3}$ simulation, despite a_0 not reaching the ionisation threshold, indicating that these electrons were ionised by tunnel ionisation rather than barrier suppression, and the low charge is due to the reduced likelihood of their ionisation.

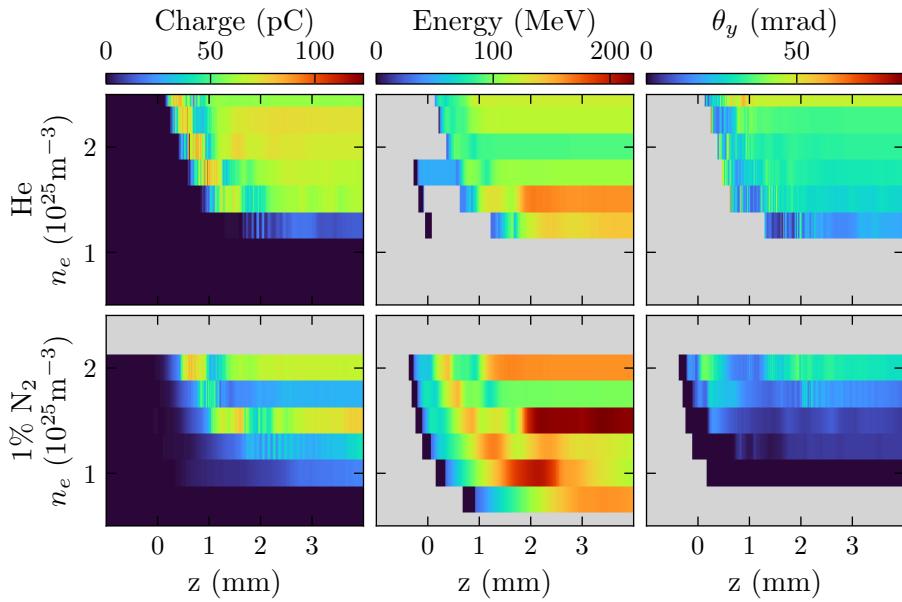


Figure 4.23: Beam metrics as a function of density and longitudinal position, for the pure helium gas, top, and the doped gas, bottom. Grey regions indicate the lack of data or insufficient charge. For comparison, the plots have been extended so that they cover the same density range, even though, for example, the highest density simulation for the doped gas was $2 \times 10^{25} \text{ m}^{-3}$.

to the self-injected beam agrees with the measurements, as well as the average divergence of the doped gas measurements being lower than for the pure gas. These similarities will be used to support the discussion of the data that follows.

Several differences between the experiment and simulation results are also observed. Firstly, although the ionisation injected charges are comparable, at approximately 10 pC, the simulated self-injected charge is roughly 10 times larger than measured experimentally, at 60 pC versus 5-10 pC. It is possible that this is a result of using a perfectly symmetrical transverse profile for the laser pulse, combined with the assumption of cylindrical symmetry, as this could enhance the amount of charge that undergoes self-injection by providing an un-physically stable bubble. Self-injection at the threshold occurs transversely¹⁸, and so it is plausible that instability in the location of the electron sheath that forms the bubble makes trapping electrons transversely more challenging. In contrast, ionisation injection is insensitive to small changes in the electron sheath location, as it is primarily a longitudinal process in which the captured electrons do not get close to the sheath. The self-injection charge discrepancy between simulations and experimental measurements is an area of active investigation. One consequence of this over-injection of charge in simulations may be an overestimate of the energy and charge resulting from any subsequent beam-driven acceleration, and so the high energies at longer lengths in the pure helium simulations are treated as being potentially spurious. Electrons accelerated by beam drivers can be readily identified by the linear

¹⁸ [153] A. G. R. Thomas (2010) *Physics of Plasmas*

relationship between energy and propagation length, compared to the parabolic trajectory of laser-driven acceleration.

Another difference between the experimental and simulated results can be found in the divergence measurements, which are more extreme in the simulations. At the higher densities, these values are too high, while they are too low at the lower densities. The difference at high densities can be explained by realising that the experimental length scan includes the density down ramp at the end of the plasma, whereas here we are only taking snapshots of the particle phase space. In addition, the simulated measurements are more likely to be skewed by long tails in the beam's phase space, which may be below the detection threshold of the spectrometer. The underestimation of the divergence at low densities could also be due to the assumption of a symmetric transverse laser profile in the simulation, as well as the effects of scattering in the tape for the real beam.

4.4 Discussion

In this section, the results presented in the previous sections will be discussed. The resolution of the data allows for many possible avenues of exploration, and consequently, this discussion will not be comprehensive. We will instead focus on the topics of electron beam energy, the rate of laser pulse evolution, and x-ray production. This discussion will be conducted by comparing the data to the appropriate scaling laws, and by employing PIC simulations to investigate the nonlinear phenomena.

4.4.1 Self-injected beam energy

We will first consider the self-injected beam in the pure helium gas. The energy gain for an electron being accelerated in the 3D, $a_0 > 2$ regime¹⁹ is given by,

$$\Delta W = mc^2 \frac{2}{3} \frac{\omega_0}{\omega_p} a_0. \quad (4.4)$$

Recall²⁰ that this scaling law is derived from a model where an electron is propagating from the peak of the accelerating field to the equilibrium point at the centre of the bubble, over the distance of the bubble radius. For the self-injected beam in pure helium, we can use the model by Thomas to estimate the laser a_0 at injection²¹,

$$\frac{2c\sqrt{\ln(2\gamma_p^2) - 1}}{\omega_p} \gtrsim r_b, \quad (4.5)$$

where r_b is the bubble radius. In the bubble regime, $r_b = 2\sqrt{a_0}c/\omega_p$, and $\gamma_p^2 = n_c/3n_e$, where γ_p is the phase velocity of the plasma wave, or the velocity of the back of the bubble, when the laser pulse is etching. Therefore, for the self-injected beam, the modified

¹⁹ [44] W. Lu *et al.* (2007) *Physical Review Special Topics - Accelerators and Beams*

²⁰ See Sec. 2.6.3.

²¹ [153] A. G. R. Thomas (2010) *Physics of Plasmas*

energy scaling law becomes

$$\Delta W_{SI} = mc^2 \frac{2}{3} \frac{\omega_0}{\omega_p} \left(\ln \left[\frac{2}{3} \left(\frac{\omega_0}{\omega_p} \right)^2 \right] - 1 \right). \quad (4.6)$$

This expression for the energy gain scales with the density slightly faster than Eq. 4.4, as at lower densities a larger bubble is driven, resulting in a longer acceleration length. For a density of $2 \times 10^{25} \text{ m}^{-3}$, this model predicts injection at $a_0 = 3.06$, giving an energy gain of 90.8 MeV. This model also assumes that the value of a_0 , and hence the electric field strength at the back of the bubble, is constant throughout the acceleration, which is likely to be an oversimplification. However, an increasing a_0 will result in an increasing bubble radius, accelerating dephasing, while a decreasing a_0 will slightly rephase the electrons, enhancing the acceleration, so the effect of a variable a_0 may be partially mitigated.

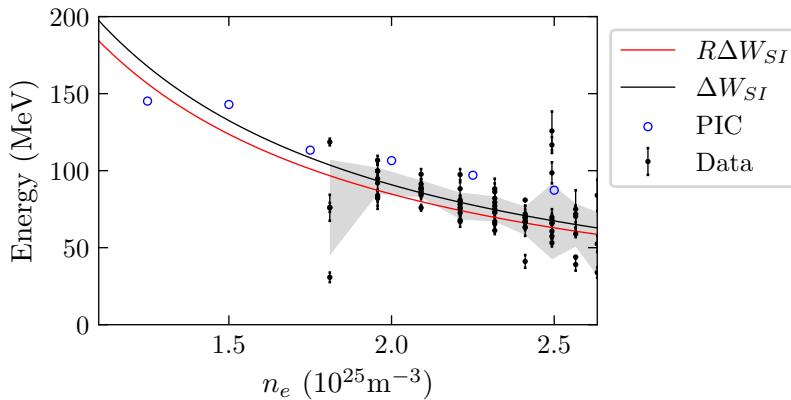


Figure 4.24: Comparing the pure helium gas peak energy data to theoretical scaling. The fitted value of R in this plot is 93%.

In Fig. 4.24 the experimental data and the results of the pure helium PIC simulations are compared to ΔW_{SI} . Each experimental data point is the peak energy and its associated error from the left panel of Fig. 4.11, excluding the densities at which no beams were injected. The grey shaded region indicates the mean error at each density, calculated using the errors on the data points, and centred on the mean energy at that density. The results of the simulations are also included, where the energy plotted here is the maximum energy reached by the first electron beam. Attempting to fit Eq. 4.6 to the experimental data, a prefactor of $R = 0.93$ minimises the square of the error-weighted distance of the line from the data points. Both ΔW_{SI} and $R\Delta W_{SI}$ are plotted in Fig. 4.24, and since both lines lie within the region of error, we conclude that the model behind Eq. 4.6 is a good descriptor of the acceleration and injection process.

4.4.2 Doped gas beam energy

We now turn our attention to the peak electron energies from the doped gas. Again, we will assume we are in the 3D, $a_0 > 2$ regime²² and will use Eq. 4.4 to predict the energy

gain. For ionisation injection of nitrogen, we simply need to determine the electric field intensity for the barrier suppression ionisation of the first inner shell electron. These will be the first electrons to be injected, and consequently should reach the highest energies due to beam loading effects. This value is found in Tab. 2.1, and gives an $a_0 = 2.2$. This value is used to predict the ionisation injected energy gain ΔW_{II} , and is plotted alongside a selection of the experimental data in Fig. 4.25.

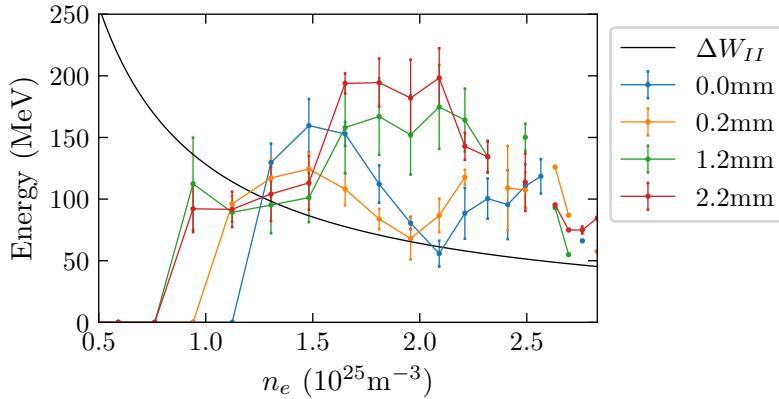


Figure 4.25: Electron beam peak energy as a function of density for a selection of cell lengths. The model behind ΔW_{II} assumes $a_0 = 2.2$. Experimental data is discontinuous at some densities due to lack of data.

When tracing the relationship between density and peak energy in Fig. 4.25 for various lengths, a pattern emerges. As density increases, the energy rises until it hits a peak, then either plateaus or falls, before rising again to produce a second peak, then falling again at high energies. This does not follow the prediction of ΔW_{II} , which is not surprising given the more complicated injection behaviour. The simulations suggest that the first peak corresponds to electrons injected via ionisation injection, while the second peak is due to injection that occurs when the laser pulse depletes, and that further, the height of the second peak is a result of beam-driven acceleration. It is not expected for the second peak to follow the predictions of ΔW_{II} , since the acceleration mechanism is completely different. However, the first peak is produced by laser-driven acceleration when the a_0 of the laser is known, so it is surprising that even this acceleration does not follow the theoretical prediction. For example, the density scan at $L = 0$ mm produces a prominent peak of 160 MeV at $1.1 \times 10^{25} \text{ m}^{-3}$, much higher than predicted by the scaling law. This means either that the accelerating gradient is stronger than predicted, or the acceleration length is longer than predicted.

To further investigate the source of this extra energy gain, we can recast the data into a length scan. This has been done for the data taken at $1.1 \times 10^{25} \text{ m}^{-3}$ in Fig. 4.26. In the fully blown-out model, the longitudinal electric field E_z has a magnitude of

$$E_z(z) = \left(1 - \frac{z}{L_{acc}}\right) E_{z,max}, \quad (4.7)$$

²² [44] W. Lu *et al.* (2007) *Physical Review Special Topics - Accelerators and Beams*

where $E_{z,max}$ is the maximum electric field found at the back of the bubble, z is the distance from the back of the bubble, and L_{acc} is the length of acceleration. The energy gain as a function of z is

$$\frac{\Delta W}{m_e c^2} = E_{z,max} \int_0^z \left(1 - \frac{z'}{L_{acc}}\right) dz'. \quad (4.8)$$

This length is determined by the bubble radius and the phase velocity of the plasma wave,

$$L_{acc} = \left(\frac{1}{1 - \beta_\phi}\right) r_b. \quad (4.9)$$

If, for now, we assume that the bubble radius is constant, and determined by the value of a_0 at injection, and further take the phase velocity to be the etching velocity, $\beta_\phi = 1 - (3n_e/2n_c)$, where n_c is the critical density, the energy gain from Eq. 4.8 follows the dashed line in Fig. 4.26. We see that the form of the prediction is similar to the measurement, which is approximately parabolic until $z = 1.2$ mm. After this point, the measured energy gain falls slowly, which is a consequence of the laser depleting before the beam fully dephases.

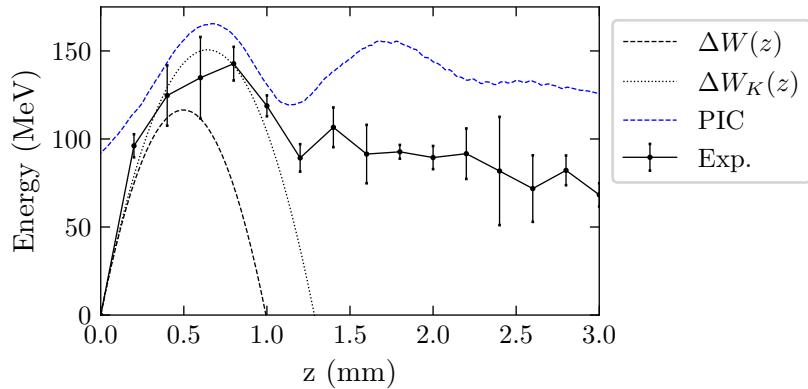


Figure 4.26: Energy gain as a function of length, for experimental data recorded at $1.1 \times 10^{25} \text{ m}^{-3}$, and simulation results at $1.25 \times 10^{25} \text{ m}^{-3}$. The error bars on the data points reflect the shot-to-shot fluctuations at each position. The energy gain model with a bubble velocity reduced by etching, $\Delta W(z)$, has been plotted, in addition to the modified energy gain model with a higher bubble velocity, $\Delta W_K(z)$.

The peak of the predicted and measured energy gains occur at different points in z , with a lower and earlier predicted peak energy than measured. A different phase velocity could account for this, which we can find by adding a fitting parameter, K , to the expression for the etching velocity,

$$\beta_\phi(K) \simeq 1 + \left(K - \frac{3}{2}\right) \left(\frac{\omega_p}{\omega_L}\right)^2. \quad (4.10)$$

A good fit is obtained for $K = 2/3$, giving $\beta_\phi \simeq 1 - 5n_e/6n_c$. This velocity is slower than the linear group velocity, for which $K = 1$, but it is closer to the group velocity than the etching velocity $K = 0$.

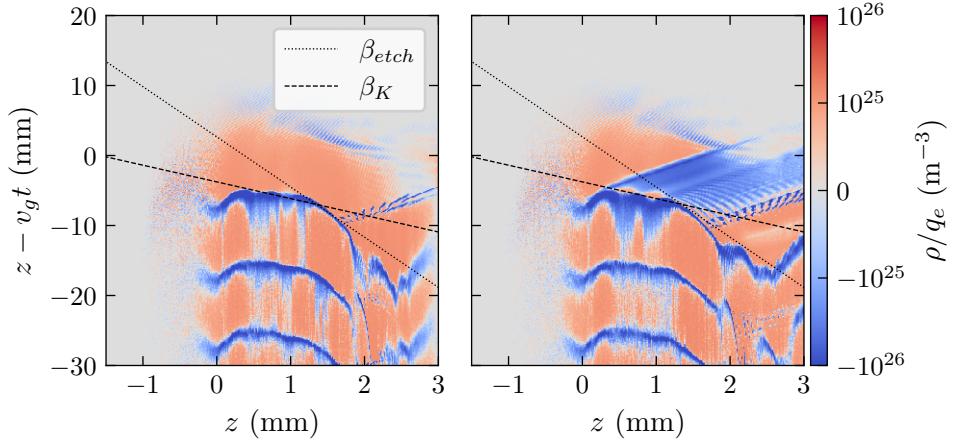


Figure 4.27: On-axis charge density map for the undoped simulation, left, and doped gas simulation, right, for $n_e = 1.25 \times 10^{25} \text{ m}^{-3}$. The laser pulse moves from left to right across the plot.

By examining the PIC simulations, we can directly measure the plasma wave velocity at the moment of injection. Fig. 4.27 gives the on-axis charge density as a function of longitudinal position for the $1.25 \times 10^{25} \text{ m}^{-3}$ simulations, where the central 3 cells have been averaged. Moving from left to right across the plots, we can see that the plasma wave structure forms at $z = 0.2 \text{ mm}$, and that the first ionisation injection event happens at $z \approx 0.2 \text{ mm}$. The injected electrons move at c through the simulation, while the plasma wave recedes at approximately $\beta_\phi (K = 2/3)$. Note that the front of the wave is also moving at this velocity, and that the phase velocity of the plasma wave is slightly higher in the pure helium gas than in the doped gas due to the effects of beam loading. In the pure helium simulation, the phase velocity drops to β_{etch} at $z = 1.2 \text{ mm}$, while in the doped gas simulation the phase velocity decelerates below the etching velocity, where the difference in velocity between the two simulations is attributed to further beam loading of the wake due to ionisation injection.

We conclude that we are in a regime where pulse depletion is semi-localised. The pulse does not etch at the rate given by Decker *et al.*²³ because the laser intensity at the front of the pulse is too low, nor does it move at the linear group velocity. For $a_0 \gg 1$, a narrow density spike forms at the front of the pulse, which causes rapid local depletion of the slice of the pulse propagating in that region of the plasma. However, we can see that for our conditions the density increase at the front of the plasma wave is not narrow, and consequently, a wide but still localised region of the pulse is being depleted. This wide region moves backwards through the pulse at a slower rate than would a narrow spike, increasing the plasma wave velocity relative to β_{etch} . An advantage of this regime is that particles injected into the faster wake can attain higher energies due to the extended dephasing length.

An increased dephasing length is also the source of the 200 MeV energies seen at $n_e = 2 \times 10^{25} \text{ m}^{-3}$. This can be seen from the plot in Fig. 4.28, where the data for

²³ [149] C. D. Decker *et al.* (1996) *Physics of Plasmas*

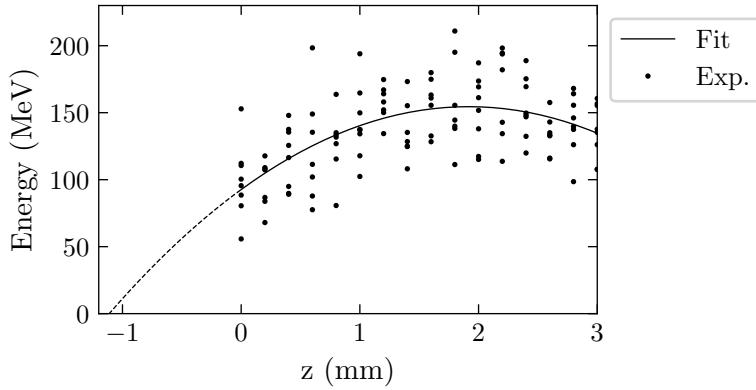


Figure 4.28: Energy as a function of cell length for all densities above $1.5 \times 10^{25} \text{ m}^{-3}$.

densities above $1.5 \times 10^{25} \text{ m}^{-3}$ has been plotted as a function of cell length. Fitting a parabola to this data, we find that the acceleration length is 3.0 mm, which is significantly longer than the group velocity dephasing length, $\lambda_p n_c / n_e = 0.6 \text{ mm}$ for $2 \times 10^{25} \text{ m}^{-3}$. Such a long acceleration length is indicative of beam-driven acceleration, where the acceleration length is determined by the energy in the driving beam, rather than the velocity difference between the driver and the witness. Beam-driven acceleration was also observed in the simulations, verifying the experimental observations.

4.4.3 Pulse evolution rate

We will now investigate the pulse evolution rate. To do this, the cell length coordinates are converted to target lengths, which is complicated by the presence of the density ramps at the front of the cell. For the purposes of this discussion, we will define the start of the plasma as the point at which the density is high enough for the laser pulse to self-focus, which occurs at $6 \times 10^{24} \text{ m}^{-3}$ for the 5 TW laser used. The form and scale length of the ramp was obtained from the fluid simulations reported in Sec. 3.2.2. The effect of this definition is that the cell length now depends on the density; for a plateau density of $1.0 \times 10^{25} \text{ m}^{-3}$, the laser pulse will begin to focus 0.6 mm before the cell, while for a density of $2.5 \times 10^{25} \text{ m}^{-3}$, the laser pulse will begin to focus 1 mm before the cell. This non-linear relationship has been plotted in Fig. 4.29.

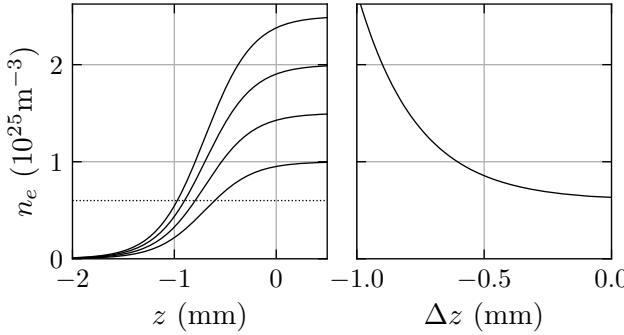


Figure 4.29: Left panel shows some of the exponential density ramps used in the simulations, with the self-focusing threshold for the 5 TW laser at $0.6 \times 10^{25} \text{ m}^{-3}$. The right panel shows the location at which the density rises above the focusing threshold as a function of plateau density, relative to $z = 0$.

The injection length as a function of density is plotted in Fig. 4.30 for both the pure helium gas and the nitrogen-doped gas. The experimentally measured injection length was determined by arranging the data into a series of density scans, see Fig. 4.25 for an example, and then determining the first point at which the charge reached 10% of the maximum charge for that density scan, linearly interpolating between data points. The error in n_e is calculated by performing the same procedure using thresholds of 5% and 15% and then using the mean of these two values. The injection point for the simulations is determined from the data in the first column of Fig. 4.23. For all injection length measurements, the correction for the effect of the density ramp from Fig. 4.29 is applied.

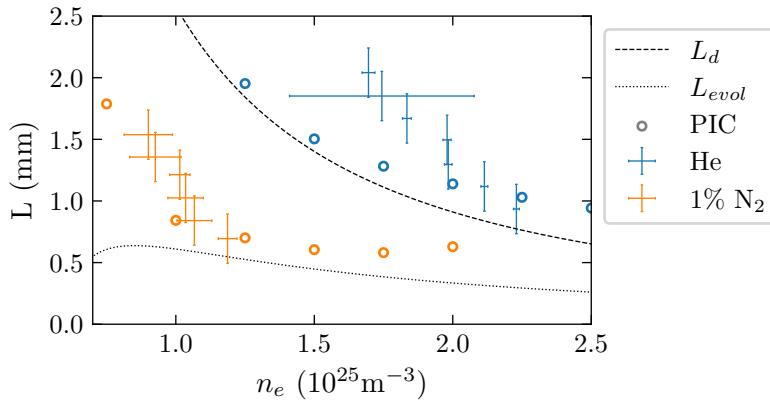


Figure 4.30: Injection length versus density, for the experimental data and from the simulations.

The depletion length²⁴, $L_d = 1.4\lambda_p n_c / n_e$ is also plotted in Fig. 4.30, as is the pulse evolution length from Streeter *et al.*²⁵,

$$L_{\text{evol}} = \lambda_p \frac{2}{3} \frac{n_c}{n_e} \sqrt{\frac{1}{2} \ln \left(\frac{P_0}{P_c} \right)}, \quad (4.11)$$

²⁴ [148] B. A. Shadwick *et al.* (2009) *Physics of Plasmas*

²⁵ [146] M. Streeter *et al.* (2018) *Physical Review Letters*

where P_0 and P_c are the peak and critical power respectively. This model requires moderate values of P_0/P_c , and will not be valid for $P_0 \lesssim P_c$, as seen by the fall in L_{evol} at lower densities. The evolution length gives the point at which maximum power amplification should occur, which is after approximately half of the laser pulse has been etched under the conditions studied by the authors, while the depletion length is the length at which all of the energy in the pulse has been lost ($L_d \approx 2L_{evol}$). Comparing L_{evol} and L_d to our measurements in Fig. 4.30, we see that they predict pulse depletion before the injection length. In the most extreme case at $n_e = 1.7 \times 10^{25} \text{ m}^{-3}$, the depletion length is 1.16 mm, while the injection length in the pure helium gas is 2.04 mm. It is unphysical for the first electrons to be injected at $z = 1.76L_d$, and so the longitudinal evolution mechanism should be re-evaluated.

First, let us consider the transverse evolution of the pulse. The matched spot size $w_m \approx 2\sqrt{2}k_p^{-1}$, and the a_0 at this size will be given by $a_{0,m} = a_{0,v}(w_0/w_m)^2$, where $a_{0,v}$ and w_0 are the vacuum a_0 and Gaussian waist. The median FWHM vacuum spot size is $21.7 \pm 0.5 \mu\text{m}$, corresponding to $w_0 = 18.4 \pm 0.4 \mu\text{m}$, with $a_{0,v}$ equal to 0.57. The matched spot size as a function of density is plotted in Fig. 4.31, as well as the corresponding a_0 value reached by the laser at the plasma focus. This is compared to the self-injection threshold predicted by Thomas *et al.*²⁶,

$$a_{0,i} = \ln(2\gamma_p^2) - 1. \quad (4.12)$$

The a_0 at the point of injection from the pure helium simulations is also plotted in Fig. 4.31. Compared to the theoretical self-injection threshold, the simulations are in reasonable agreement, although it appears that slightly higher values than the threshold are required, particularly at higher densities. This difference could be due to the higher bubble velocity that results from being in the semi-localised pulse etching regime discussed previously. The simulations indicate that, in the density range studied, transverse evolution alone does not amplify the field strength enough to cause self-injection; it will require longitudinal evolution to further increase a_0 .

For ionisation injection, the field strength needs to be strong enough to cause a sizeable fraction of the N^{5+} ions to ionise, and to drive a sufficiently strong plasma wave to capture these electrons²⁷. The a_0 value required to ionise N^{5+} through the barrier suppression mechanism is 2.2, however, over the duration of the pulse, some ionisation can be expected to occur through tunnelling at intensities significantly below this value, an effect shown in Fig. 2.4. This is why ionisation injection can be seen to occur at a_0 values far below 2.2, as shown by the red crosses in Fig. 4.31. The a_0 at which ionisation injection occurs in simulations roughly follows the $a_{0,m}$ line, since, although the intensity is high enough to cause ionisation for $a_0 \gtrsim 1.5$, the plasma wave is still forming before

²⁶ [153] A. G. R. Thomas (2010) *Physics of Plasmas*
See Sec. 2.7.2.

²⁷ [164] C. McGuffey *et al.* (2010) *Physical Review Letters*
See Sec. 2.7.4.

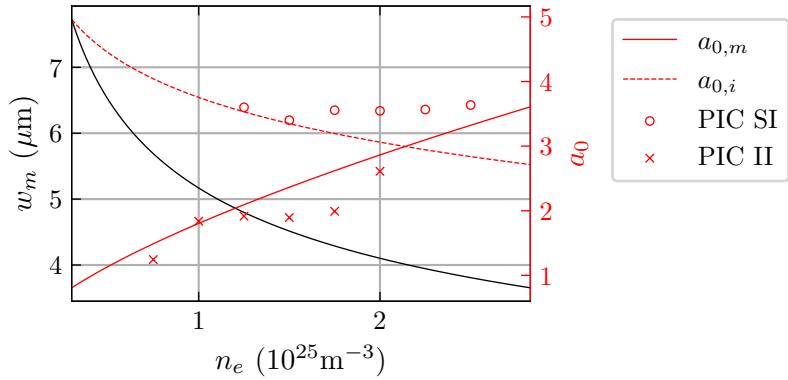


Figure 4.31: The black line shows the matched spot size, w_m , as a function of n_e . The solid red line, $a_{0,m}$, is the a_0 of the laser at this matched spots size, while the dashed red line is the theoretical a_0 required for self-injection. The ‘PIC SI’ and ‘PIC II’ markers give the a_0 at which injection occurs in the simulations.

the laser pulse reaches $a_{0,m}$, and so the field strength needs to reach an intensity that allows the ionised electrons to be captured.

Since a_0 crosses from the upper end of the linear regime ($a_0 \lesssim 1$) into the lower end of the nonlinear regime ($a_0 \gtrsim 1$), the pulse first self-focuses and then longitudinally compresses. For densities where longitudinal compression is required for injection, the injection length is determined by the sum of the plasma focal length, the distance taken for the laser pulse to reach the matched spot size, and the longitudinal scaling, either L_{evol} or L_d . The self-focusing rate is determined by Eq. 2.71, and the distance to the plasma focal length, L_f , can be determined by finding the first minimum in the spot size, as was done in Fig. 2.11 for various spot sizes and densities. By performing a fit to the solutions of Eq. 2.71, we find that, over the density range studied here, the plasma focal position scales with $(n_c/n_e)^{3/2}$, which is the same scaling as the depletion length,

$$L_f[\text{mm}] \simeq 3.1 \times 10^{-4} \left(\frac{n_c}{n_e} \right)^{\frac{3}{2}} + 0.34. \quad (4.13)$$

This fitting uses only the plasma density and the initial focal spot parameters, not the experimental data. A numerical solution was required because the transverse envelope equations of motion cannot be solved analytically. The injection length is then the sum of this transverse evolution length and the longitudinal evolution length

$$L_{inj} \simeq L_f + L_d. \quad (4.14)$$

In Fig. 4.32 the simulated and measured injection lengths are compared to this new scaling. The injection point for the pure helium gas now occurs at some point between the evolution length and the depletion length, and the scaling with the density is approximately correct. In simulations, we observe that the largest a_0 values occur at

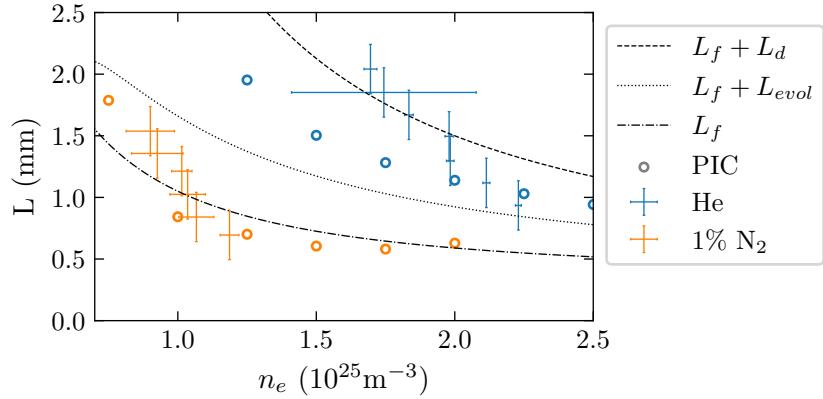


Figure 4.32: Injection length versus density, for the experimental data and from the simulations. Now compared to new injection length predictions.

this point in its evolution; when the pulse is highly compressed but yet to fully deplete. Only the lowest density injected beams lie above the depletion length, and this can be accounted for by realising that the evolution fit we derived assumed a constant density, whereas the average density during the focusing is slightly lower than this due to the ramp. The ionisation injected beam injects at or slightly before the plasma focal length, in agreement with the injection a_0 values in Fig. 4.31.

4.4.4 X-ray source

We will now discuss the x-rays produced by the accelerator. The peak x-ray intensity produced by the doped gas was a factor of 9.4 times higher than that produced by the pure helium gas, while the peak photon flux was a factor of 4.4 times higher in the doped gas. The difference in performance can be explained by considering the electron beam properties; the charge and the electron beam energy were both approximately a factor of 2 larger in the doped gas than in the helium. The spectral intensity of the x-rays scale as²⁸

$$\frac{d^2I}{dEd\Omega} \propto N_\beta \gamma^2, \quad (4.15)$$

where N_β is the total number of betatron oscillations performed by all of the radiating electrons. This scaling explains the difference in the measured intensity between the gases. The high critical energies measured in the pure helium gas reduced the peak photon flux, and are likely to be due to the larger amplitude betatron oscillations, as indicated by the electron spectrometer images in Fig. 4.13, and as shown in the injection mechanism in Fig. 4.19.

The x-ray measurements show that the doped gas was capable of producing approximately $11.0^{+0.8}_{-4.2}$ ph mrad $^{-2}$ 0.1%BW $^{-1}$, but we would like to compare this value to the x-ray production of other sources. A broadly used figure of merit for x-rays is the peak brilliance of the beam, and so, to compare our results, we will estimate the brilliance experimentally, and then again separately using PIC simulations.

²⁸ See Sec. 2.8.4.

First, we will attempt to make this measurement using only experimental measurements. Since the brilliance requires accounting for the spatiotemporal source size, and this was not directly measured, we must make some assumptions to estimate this size. The first assumption will be that the x-ray pulse length is equivalent to a plasma period, $\tau_X \sim 2\pi/\omega_p$. Although electron bunch lengths have been measured to be as short as a few femtoseconds²⁹, large x-ray fluxes are caused by large amounts of charge, which, as seen in simulations, results from continuous injection over a large distance. As a result, the plasma period is probably an appropriate approximation for the x-ray pulse length.

The transverse dimension of the bunch can be estimated using the relation,

$$s_\beta = 2r_\beta = \frac{2E_{crit}\lambda_p}{3\pi^2\gamma^2c\hbar}, \quad (4.16)$$

where γ is an estimate of the electron energy. The peak brilliance is then

$$B = \frac{N_{ph}}{\tau_X\pi r_\beta^2}, \quad (4.17)$$

where the photon flux N_{ph} , was calculated previously and plotted in Fig. 4.18. Using the measurement of the 95% cut-off energy in Fig. 4.10 and the density measurements, the brilliance can be determined. These are shown in Fig. 4.33.

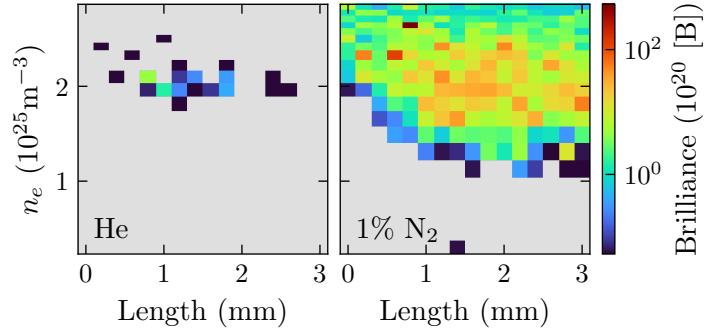


Figure 4.33: Estimated x-ray brilliance. The logarithmic colour scale spans the entire range of values, while the grey regions indicate the regions with insufficient signal for the brilliance to be calculated. The units of brilliance are $[B] = \text{ph s}^{-1} \text{mm}^{-2} \text{mrad}^{-2} 0.1\% \text{BW}^{-1}$.

The 98th percentile brilliance for the helium gas was $3.3_{-2.8}^{+3.0} \times 10^{20} [\text{B}]$, while for the doped gas this value was over an order of magnitude larger, at $45_{-19}^{+18} \times 10^{20} [\text{B}]$, where the units of brilliance are $[B] = \text{ph s}^{-1} \text{mm}^{-2} \text{mrad}^{-2} 0.1\% \text{BW}^{-1}$. Such high brilliance values are mainly due to the small source sizes; $0.35 \pm 0.14 \mu\text{m}$ and $0.18 \pm 0.07 \mu\text{m}$ for the pure and doped gases respectively, derived using Eq. 4.16. Unsurprisingly, the measured brilliance has large errors, as both the source size and the photon flux are sensitive to the critical energy, which itself has significant uncertainty. In addition, the average error due

²⁹ [64] M. Heigoldt *et al.* (2015) *Physical Review Special Topics - Accelerators and Beams*

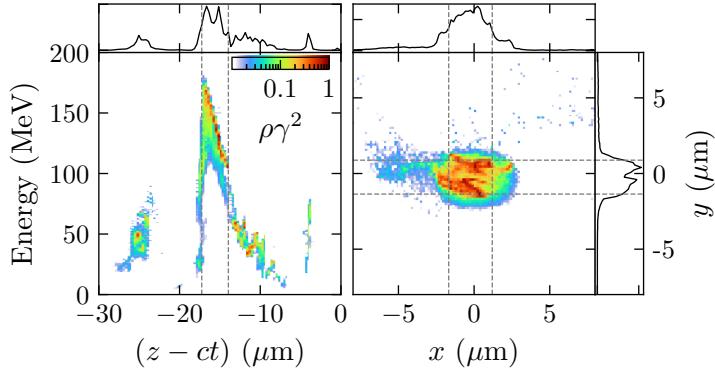


Figure 4.34: Longitudinal (left) and transverse (right) beam size from doped PIC simulation at $2 \times 10^{25} \text{ m}^{-3}$. The lineouts to the sides of the histograms are the averaged values along their respective axes, in arbitrary units. The dashed lines show the extent of the FWHM.

to uncertainty on $\gamma \sim 10\%$, which further increases the uncertainty on the source size. The error on the critical energy could be reduced by acquiring more signal, while the error on the source size could be reduced by making a knife-edge measurement, although this would require a large magnification if the source truly is approximately $0.1 \mu\text{m}$.

Instead of using the measurements of the electron beam energy to estimate the source size, we can use the PIC simulations. In Fig. 4.34, weighted histograms of the particle distributions from the $2 \times 10^{25} \text{ m}^{-3}$ are shown, which give an indication of the longitudinal and transverse beam size. In these histograms, the weight of each particle has been multiplied by γ^2 , as this is the scaling of the intensity of radiation emitted by an accelerating particle. We will use this radiation density as a metric to measure the beam size.

In the longitudinal phase space, we see a multi-bunch structure, corresponding to the multiple injection events. The FWHM length of the bunch with the highest radiation density is $5.4 \mu\text{m}$, which is approximately equal to the plasma wavelength, $\lambda_p = 7.5 \mu\text{m}$. This suggests that our initial estimate that the x-ray pulse length is equal to the plasma wavelength is valid. However, this is only a static picture of the phase space, since, from observing the waterfall plots in Fig. 4.21, we see that electrons have a range of energies over the course of the interaction. It is possible that each bunch contributes to the total x-ray flux, which would mean that the x-ray pulse is composed of multiple flashes, separated by 10s of femtoseconds, significantly reducing the average brightness.

Transversely, an elliptical radiation density profile is observed, with the major axis of the ellipse aligned to the laser polarisation axis. The FWHM beam size is $2.24 \times 2.28 \mu\text{m}$ in the x and y directions respectively. Eq. 4.16 for a $\gamma = 200$ beam with $E_c = 2 \text{ keV}$ predicts a beam size of $1.2 \mu\text{m}$, significantly lower than for the simulated beam. Likewise, for an electron beam with $\gamma = 200$, and a source size that matches the beam from the PIC simulation, then we would predict $E_c = 4.2 \text{ keV}$. This discrepancy can be addressed by realising that the value of γ that has been used to calculate r_β may not be appropriate.

While the maximum intensity per electron may be radiated by those with the highest energy, the x-ray spectrum that is measured is a composite of many different synchrotron spectra, resulting from the significant energy spread in the electron bunch. In addition, since electrons radiate while they are being accelerated and dephased, a scaled metric of the energy, such as the maximum of $\gamma^2(dQ/d\gamma)$, would also be inaccurate for predicting r_β . Instead, either experimental measurements of the source size, such as using a knife edge, or simulations using slightly aberrated focal spots should be used to determine this quantity, and hence the brightness.

For completeness, the values for the brilliance in Fig. 4.33 have been re-calculated using a constant transverse source size of $s_\beta = 2\text{ }\mu\text{m}$ and plotted in Fig. 4.35. This decreases the brilliance by approximately two orders of magnitude for a maximum value in the doped gas of $4.2_{-0.8}^{+0.8} \times 10^{20}[\text{B}]$.

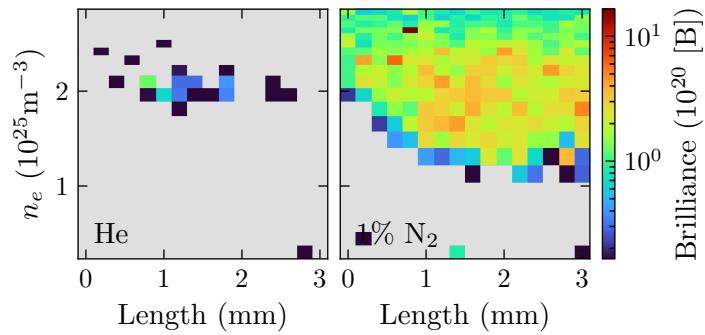


Figure 4.35: Brilliance values using a constant $2\text{ }\mu\text{m}$ source size.

4.4.5 Literature comparisons

In Fig. 4.36 the highest electron beam energy measured in this chapter is compared to the published literature, as well as results from recent theses³⁰. The theses used the same laser system as was used in this chapter but with more than twice the energy. Across all the results, the relationship between laser power and peak electron energy is clear and was highlighted in the introduction³¹. Higher powers produce higher energies because they allow for non-linear plasma wave excitation to occur at lower densities, extending the dephasing length. The results presented here were able to achieve high energies through the use of ionisation injection seeded beam-driven acceleration, extending the dephasing length through a different mechanism. While the high energy beams produced in this chapter had thermal spectra, if the injection could be localised, such as with the use of structured density profiles³², then it might be possible to produce self-seeded narrow energy spread beams that exceed 200 MeV with 5 TW lasers. Further, the increased

³⁰ [175] K. Poder *PhD. Thesis* (2016)

[170] J. C. Wood *PhD. Thesis* (2016)

[219] S. V. Rozario *PhD. Thesis* (2020)

³¹ See Sec. 1.2.

³² See Sec. 2.7.3.

phase velocity of the bubble due to the semi-localised depletion of the pulse could be used to further enhance the laser-driven energy gain of the electrons, through the use of pulse shapes optimised to maximise the phase velocity.

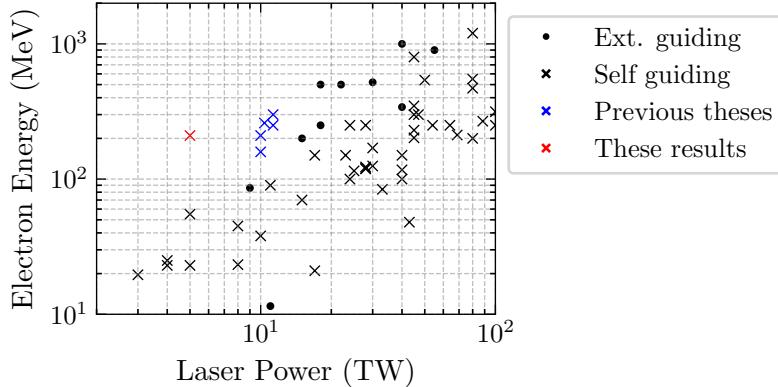


Figure 4.36: Comparison of the peak electron beam energy measured in this chapter to published literature results (black markers) and previous theses from the group using the same laser system (blue crosses).

The peak brilliance measurements reported in this chapter are compared to the literature in Fig. 4.37. These measurements are almost an order of magnitude higher than measurements performed at more than double the laser power³³ and are the lowest power at which betatron beams capable of imaging have been detected. Since lower-power lasers can operate at higher repetition rates, the achievement of high repetition rate imaging, as well as for the use of laser-driven sources as imaging devices in small-scale facilities, such as hospitals and universities. Since the high photon fluxes observed here result from the use of the ionisation injected beam, optimisation of the dopant species and percentage would be expected to yield even higher performance. Consequently, for experiments in this low- a_0 regime, a greater degree of control over the dopant would be useful.

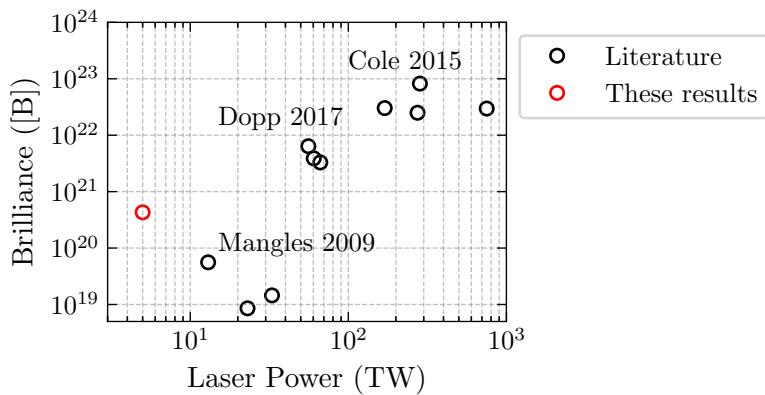


Figure 4.37: Comparison of the peak x-ray brilliance measured in this chapter to published literature results (black markers).

³³ [124] S. P. D. Mangles *et al.* (2009) *Applied Physics Letters*

4.5 Summary

The density-length parameter space of a laser wakefield accelerator has been characterised, both in terms of the produced electron beams and betatron radiation. Measurements were performed using pure helium gas and 1% nitrogen-doped helium gas as the target, with automated Bayesian optimisation procedures used to tune the experimental parameters before recording the data. The data was used in conjunction with simulation results to identify various phenomena, such as semi-localised pulse depletion and beam-driven acceleration, and to construct a model for the pulse evolution rate that combines the transverse and longitudinal effects. The findings are most applicable to laser pulses with a pulse length on the order of the plasma wavelength and have a vacuum $a_0 < 1$, which rely on relativistic self-focusing to drive a non-linear plasma wave and achieve ionisation and self-injection. The electron beam peak energies and x-ray brilliance measurements are the highest in the literature at the laser power used. A summary of the optimal beam metrics and the parameters that produced them are displayed in Tab. 4.3.

Table 4.3: Peak values obtained for various electron and x-ray beam metrics. The left column gives the measurements, while the right column gives the coordinates, in the form ($n_e(10^{25} \text{ m}^{-3})$, Length (mm)).

	Measurement		10^{25} m^{-3} , mm	
	He	1% N ₂	He	1% N ₂
Charge (pC)	13.0 ± 1.5	27 ± 4	2.1, 0.6	2.1, 2.6
Max. Energy (MeV)	125 ± 33	211 ± 7	2.5, 1.0	1.9, 1.8
Min. Divergence (mrad)	8.6 ± 1.1	1.4 ± 0.1	2.4, 0.0	0.9, 1.0
X-ray signal (counts)	518 ± 2	4890 ± 10	2.0, 0.8	2.1, 1.6
Critical energy (keV)	8.0 ± 1.2	2.8 ± 0.2	2.4, 0.2	2.5, 0.4
Flux (ph mrad ⁻² 0.1%BW ⁻¹)	$2.5_{-0.1}^{+0.7}$	$10.7_{-1.6}^{+1.3}$	2.1, 0.8	1.3, 2.8
Brilliance $10^{20}[\text{B}]$	$0.2_{-0.04}^{+0.03}$	$4.22_{-0.8}^{+0.8}$	2.2, 1.2	2.1, 1.2

5 Measurement of ionisation injected GeV beam emittance

THE TRANSVERSE EMITTANCE is a useful figure of merit for the quality of a particle beam. Laser wakefield accelerators naturally produce beams with normalised emittances that are typically below $10\text{ }\mu\text{m}$, which is comparable to conventionally accelerated beams. However, measuring beams from wakefield accelerators is comparatively challenging, owing to the small source size combined with the large divergence of the beam, as well as the large energy spread. Several methods have been implemented¹, but all of the measurements reported in the literature are at sub-GeV energies and often calculated using x-ray measurements, and are hence indirect. In addition, the vast majority of the measurements concern electrons injected from the background plasma via wavebreaking, as opposed to those injected by the ionisation injection mechanism.

This chapter presents direct emittance measurements of GeV laser wakefield accelerated beams. All the data presented here was obtained using the Gemini laser in ATA3 during the summer of 2021, where the average on-shot energy 7.4 J , reaching a power of 165 TW and a vacuum intensity of approximately $6 \times 10^{22}\text{ W m}^{-2}$. The primary goal of the experiment was to provide a high-charge, high-energy electron beam to a commercial user group, and the measurements reported here were obtained in parallel to that experiment. A helium gas doped with 2% nitrogen was used as the target, so beams were primarily injected via ionisation injection. A tungsten mask placed close to the target was used to scatter the beam, and the scattered pattern on the electron spectrometer was used to determine the spatially resolved beam size, from which the emittance was calculated. A Monte Carlo-based technique was used to confirm these measurements. The first part of this chapter is dedicated to describing this method and characterising its uncertainty. Possible improvements to this technique are discussed at the end of the chapter.

The second part of this chapter reports the emittance measurements and develops a simple model for predicting the emittance of ionisation-injected beams. The normalised emittance measurements were as low as $4\text{ }\mu\text{m}$ in some regimes but could also be as high as $30\text{ }\mu\text{m}$ in others, and the emittance was found to depend on both the plasma

¹ See Sec. 5.5.2.

density and the gas jet backing pressure independently. Particle-in-cell simulations were used to investigate the emittance evolution at a range of densities and examine electron trajectories as they are injected into the wake. A scaling was derived that predicts that the emittance should be particularly sensitive to the beam size at injection, which is effectively the ionisation radius of the laser pulse.

5.1 Experimental setup

5.1.1 Vacuum chamber layout

The vacuum chamber layout is shown in Fig. 5.1. Only the south beam was in use for this experiment, which enters the vacuum chamber from the roof. Upon reflection, a beam with a polarisation direction in the plane of the page is produced, which is retained to the target. After entering the chamber it is reflected by an Adaptive Optic (AO), which is used to correct for wavefront aberrations by using a feedback loop between it and a HASO wavefront sensor². The AO is also used to perform parameter scans of various aberrations using the electron spectrometer for feedback. The beam was focused using an $f/40$ off-axis parabolic mirror, which focused the 150 mm diameter beam over 6 m to a spot size of $41.5 \pm 1.0 \mu\text{m}$ average FWHM. To fit this focusing geometry into the vacuum chamber, a pair of mirrors was used to reflect the beam mid-focusing, both of which required periodic changing due to burn damage.

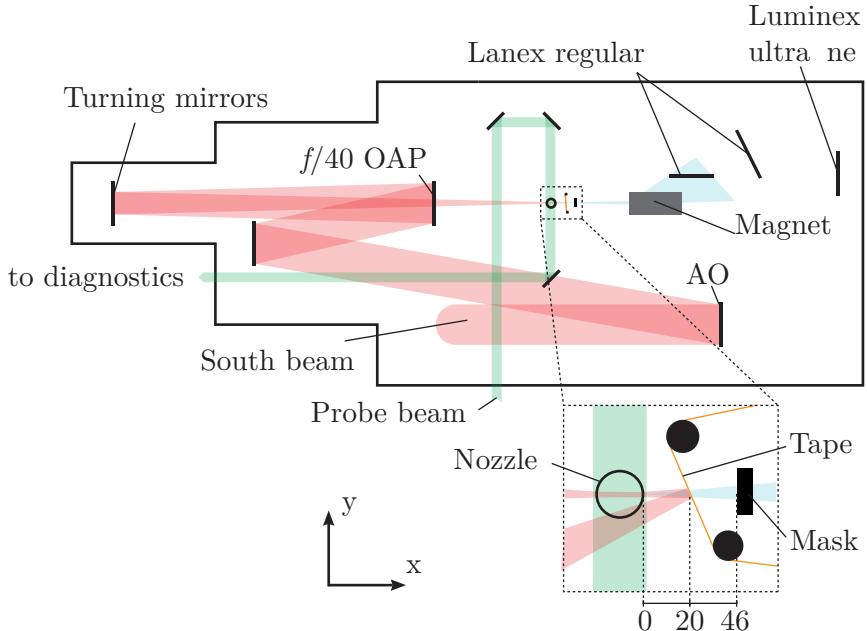


Figure 5.1: Experimental setup for the data reported in this chapter. The inset shows the arrangement of the tape and emittance mask near the nozzle, with measurements in millimetres.

² [224] HASO4 FIRST, Imagine Optic

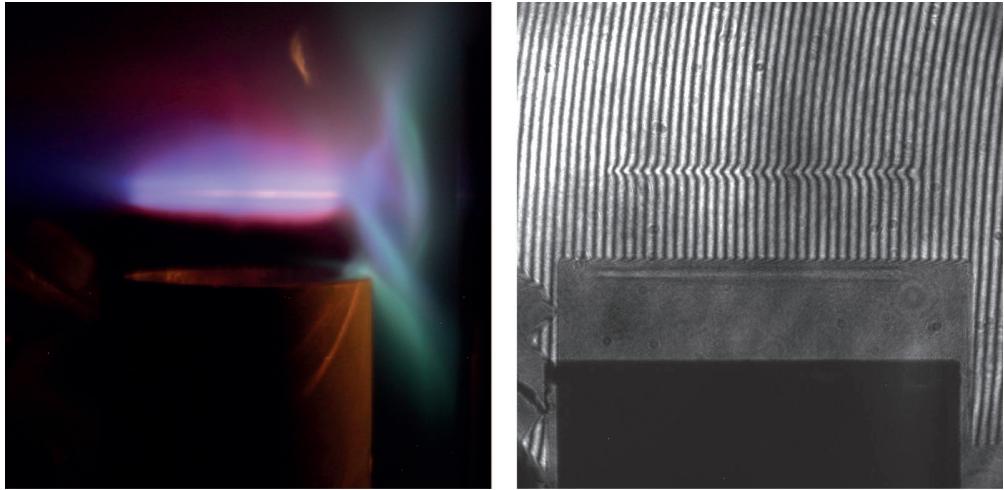


Figure 5.2: Left image is a true-colour photograph of the target during a shot. The laser is propagating from left to right, producing the purple plasma in the centre of the image. An interferogram of the plasma is shown on the right.

A tape was positioned 40 mm after the end of the gas jet, which was used to dump the residual laser energy onto a ceramic plate, protecting the emittance mask. The tape was composed of Kapton, was 25 μm thick, and angled at approximately 10° to the optical axis to avoid reflecting light back up the optical path, where it may cause damage in the upstream amplification chain. This tape was necessary to protect the emittance mask, as well as the Kapton window at the end of the vacuum chamber. This window formed part of the vacuum seal; catastrophic damage would result if it was punctured by the laser. The tape was spooled between each shot so as to provide an undamaged surface for reflecting the laser, and cameras monitored the position of the tape to ensure that it did move between each shot as intended.

5.1.2 Gas jet characterisation

A time-integrated true colour image of the target during a shot is shown in the left panel of Fig. 5.2. The laser is propagating from left to right in the image, first encountering the helium-nitrogen mixed gas above the nozzle before being reflected off the plasma mirror to the right of the image (out of shot). The different colours are due to the various species that are undergoing recombination, with the blue-pink light being due to the helium-nitrogen mix, while the green light is due to scattered laser light ablating and ionising some steel in the nozzle. Note that the exact colours are sensitive to the white balance setting of the camera, therefore the true colours may be different to those depicted.

The density profile was measured using transverse interferometry; an example interferogram is shown in the right panel of Fig. 5.2. This made use of a dedicated probe beam, formed from the leakage of the main laser. The pulse length of this beam was reduced to approximately 100 fs, which was the minimum attainable pulse length given the compressor geometry, before propagating into the target chamber, where it

was first expanded to a 30 mm diameter before traversing the gas jet, orthogonally to the main beam. The 10 mJ beam had an intensity of approximately $1 \times 10^{14} \text{ W m}^{-2}$, bright enough to overpower the plasma self-emission and forgo frequency doubling, and the polarisation of the probe at the target was orthogonal to that of the main beam, suppressing any pump-probe scattering³. An image relay system composed of three lenses transported the probe and image rays out of the chamber, ensuring that the image rays were perpendicular to the vacuum window to avoid distortion. A final objective lens focused the beam onto 16-bit CCDs, one for shadowgraphy and one for interferometry. A Mach-Zender geometry was chosen for the interferometer to allow independent control of the fringe spacing and angle, and a narrow band filter was used to expand the region of high fringe contrast.

Measurements of the longitudinal plasma density profile at various heights and backing pressures are shown in Fig. 5.3. This was taken from a scan where 3 shots were taken at a range of backing pressures from 40 bar to 100 bar, and from a range of heights from 5 mm to 17 mm in 7 steps. Each horizontal line of the density maps in Fig. 5.3 is the average longitudinal density at each location. The longitudinal resolution is pixel limited to $19 \mu\text{m}$. These plots show how the gas expands and decreases in density as the distance from the nozzle increases. Longitudinally, the density profile is approximately trapezoidal, but as the backing pressure increases, the profile takes on a distorted form, with high-density sides and a lower-density core. This effect has been observed in neutral gas density measurements at high pressures⁴ and is attributed to shocks propagating in the nozzle⁵.

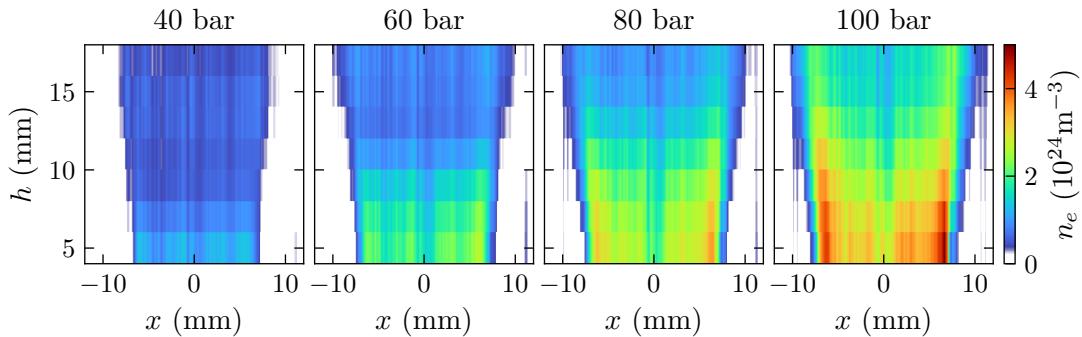


Figure 5.3: Plasma density maps of the gas jet for a range of backing pressures.

5.1.3 Electron spectrometer

The electron spectrometer for this experiment was composed of a 300 mm \times 100 mm 1 T magnet, and 3 separate scintillating screens. The first screen covered an energy range of 116 - 309 MeV, the second covered from 101 - 1039 MeV, while the third measured

³ [127] J. M. Cole *PhD. Thesis* (2016)

⁴ [36] V. Malka *et al.* (2001) *Physics of Plasmas*

⁵ See Sec. 3.2.1.

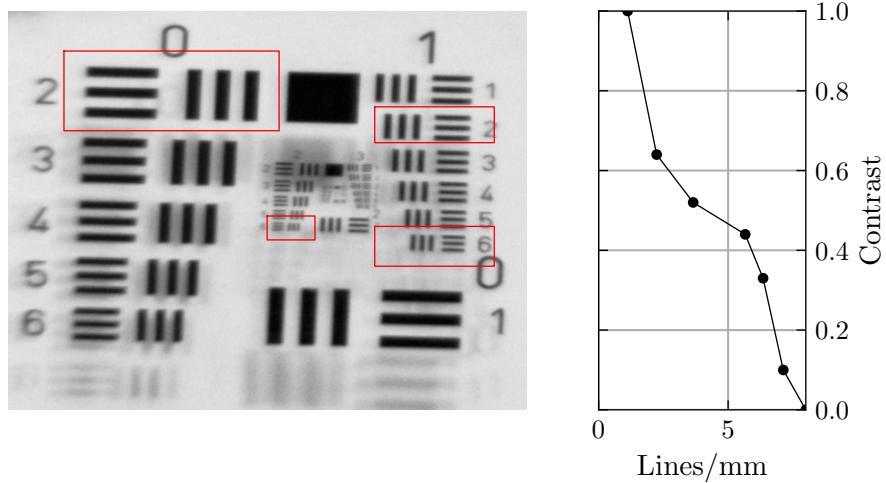


Figure 5.4: Image of resolution target taken by the high magnification spectrometer camera, with the red boxes indicating some of the line sets used for the MTF measurement. The right plot shows the measured MTF of the imaging system.

electrons with energies from 317 - 2130 MeV, where the energy ranges assume the electron beam is initially aligned with the optical axis on entering the magnetic field. Note that the final 1000 MeV of the high energy screen was detected by a length of only 30 mm, so this energy range was particularly susceptible to errors due to electron beam pointing variations. The two lower energy screens used Lanex Regular, while the highest energy screen used Luninex Ultrafine⁶, which reduced the sensitivity of this screen but increased the resolution, as required for imaging the shadow of the mask. A 16-bit CCD camera imaged each scintillating screen, and projective transforms were used to correct viewing angle distortion. The purpose of the second Lanex was to provide pointing information along the dispersion axis of the spectrometer, but this could not be used for emittance measurements as the additional scattering it caused destroyed the structure imprinted on the beam by the mask. Hence, divergence fluctuations introduce uncertainty into these measurements.

To improve the spatial resolution of the spectrometer, a second 16-bit camera with higher magnification was used to image the upper energy range of the high-energy screen. This camera could detect electrons with energies from 702 - 1393 MeV, which was the typical energy range of the high energy features of the electron beam. This improved the resolution by a factor of approximately 3, which was necessary to resolve the emittance. Fig. 5.4 shows an image of a resolution target placed on the Lanex screen, as well as a measurement of the modulation transfer function (MTF), which measures the contrast as a function of spatial frequency for the purposes of determining the resolution. Based on this measurement, the optical system should be able to resolve $\sim 80 \mu\text{m}$ features, so long as the contrast is sufficiently high.

⁶ See Sec. 3.4.1.

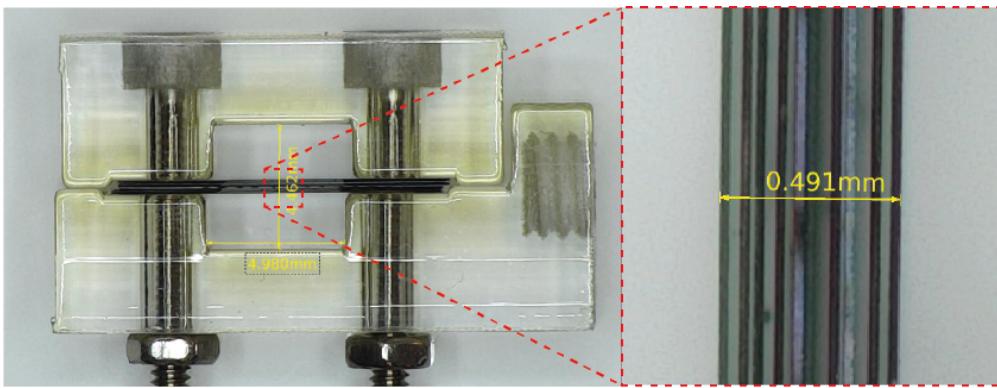


Figure 5.5: Images of the emittance mask. Each foil is $25\text{ }\mu\text{m}$ thick, and extends for 5 mm.

5.1.4 Emittance mask

The emittance mask was constructed from a stack of tungsten foil by the target fabrication team at the Central Laser Facility. Each foil had a thickness of $25\text{ }\mu\text{m}$, and was cut to a length of 5 mm. The separation between foils was provided by using the same foils as spacers, and the stack was clamped in place by two screws in a 3D-printed mount. The total number of slits was 8, limited by the time taken to stack the foils, as well as the requirement for some amount of clear space between the foils when viewed at an angle.

The design and placement of the emittance mask was determined using the resolution limits of the imaging system, the size of the vacuum chamber, and the expected electron beam properties. Electron beams with divergences of approximately 2 mrad were expected, with energies in the 1 - 2 GeV range, while the source size was expected to be on the order of $1\text{ }\mu\text{m}$. This meant that, for the $\sim 40\text{ }\mu\text{m}$ resolution of the imaging system, a magnification of $M = 40$ would be required to resolve the source. During the experiment, the mask was positioned 46 mm after the edge of the gas jet, while the high energy (closest) end of the electron spectrometer screen was 1.6 m from the same point, for a magnification of 35.

Placing the mask closer to the edge of the gas jet would have increased the magnification, however, this would have been at the cost of either sampling frequency or the scattering material thickness. For example, at 46 mm, a 2 mrad beam would be $92\text{ }\mu\text{m}$ wide, and so a mask spacing of $25\text{ }\mu\text{m}$ or smaller would be needed to provide sufficient samples to perform a reliable emittance measurement. Masks with smaller spacing than this can be constructed, such as with an array of thin wires or a laser-drilled sheet of metal, but these would not have the thickness required to scatter a 2 GeV beam. Using a thinner foil thickness would have been possible, but this would make the mask overly susceptible to warping of the foils, and would have made alignment difficult due to the reduced inter-foil spacing.

The emittance mask was placed 2 mm after the tape, protecting it from laser damage. The mask was attached to a vertical motorised translation stage, which allowed for it to

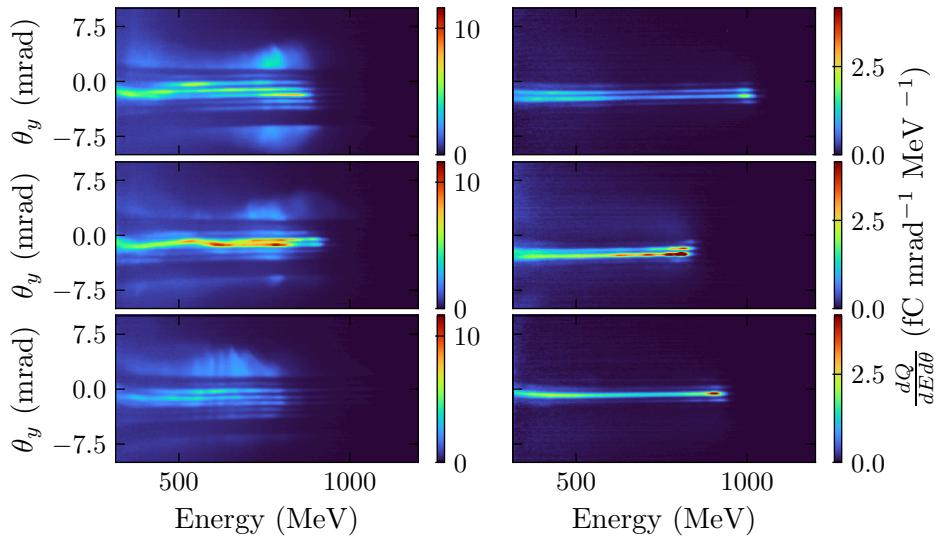


Figure 5.6: Electron spectrometer images from the highest energy scintillating screen, with the emittance mask inserted into the beam. The three beams on the left were produced with a plasma density of $2.5 \times 10^{24} \text{ m}^{-3}$, while the beams on the right used a plasma density of $1.5 \times 10^{24} \text{ m}^{-3}$.

be inserted into and removed from the beam under vacuum. Due to the small beam size at this point, accurate positioning of the mask was required to ensure that the electron beam hit the centre of the mask as often as possible. This was done by first setting the vertical position using the laser as a reference and then using the mask's shadow on the electron spectrometer to perform fine adjustments. Due to the length of the foils and the small foil spacing, control of the mask's pitch angle was also required. A picomotor attached to the mask's mount was used to finely adjust the pitch under vacuum, where the signal intensity on the electron spectrometer was used to provide feedback.

The emittance mask slits were aligned with the dispersion direction of the dipole magnet. This allowed for the emittance for each energy slice of the beam to be measured independently, permitting a full measurement of the geometric beam emittance, since the transverse momentum spread of each slice can be deduced from the divergence and energy. Example images from the electron spectrometer with the mask in the beam are shown in Fig. 5.6, where two different plasma densities are being compared. The beams produced in the higher-density plasma are more divergent, overfilling the mask so that its entire shadow is visible, whereas, for the lower divergence beams, the beam is fully masked.

5.2 Emittance measurement

After recording a satisfactory spectrometer image of the masked beam, the task remains to extract the emittance. This will first be done using a modified pepper-pot method for an example image, and a detailed error analysis will be performed, in which Monte Carlo (MC) simulations of the scattering process are used to determine the effect of

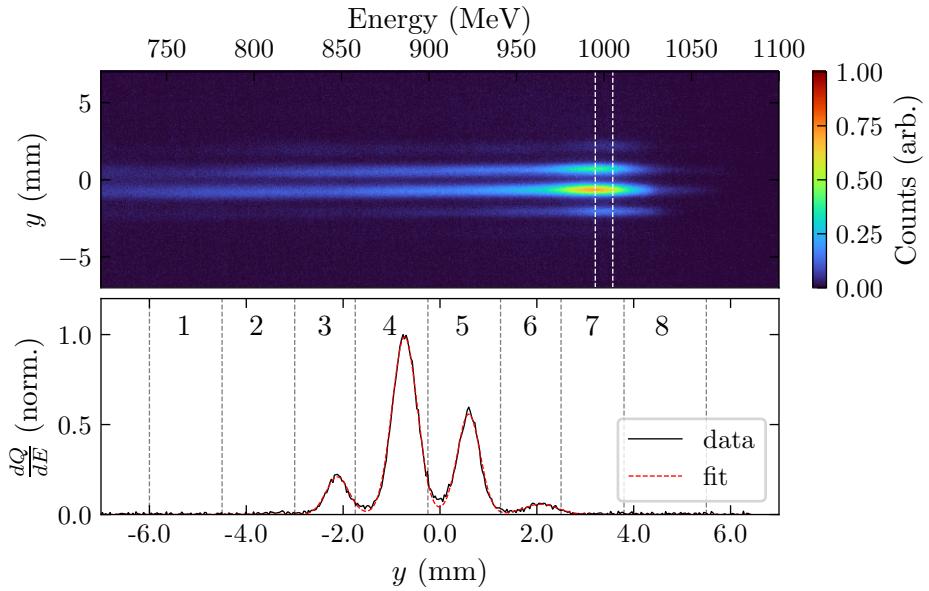


Figure 5.7: The upper panel shows a processed electron spectrometer image from the high-resolution camera. The lower panel shows the multi-Gaussian fit to the data at 1 GeV and the identification of the slits, numbered from 1 – 8.

input parameter variation on the measurement. Finally, it will be shown how these MC simulations could be used to measure the emittance directly through an iterative process.

5.2.1 Gaussian fit method

The main method for determining the emittance from the images will be to fit a series of Gaussian peaks to the transverse lineouts of the image at each energy. The tacit assumption here is that the transverse coordinates of the electrons in the beam at the accelerator exit are normally distributed. This process is performed for the beam from the top right panel in Fig. 5.6. The high magnification image of the beam is shown in the top panel of Fig. 5.7, and the bottom panel of the figure shows a Gaussian peak fit to the lineout of the image at 1000 MeV, where a ± 5 MeV window has been averaged to produce the data lineout.

To perform the Gaussian fit, the number of peaks is first identified. Peaks are detected by finding where the signal rises above 3σ of the background level, where σ is the standard deviation of the background noise, calculated using the edge of the lineout. Rather than fitting a Gaussian to each peak individually, a multi-Gaussian function is defined,

$$G(\boldsymbol{\mu}, \boldsymbol{\sigma}) = \sum_i^N a_i \mathcal{G}(\mu_i, \sigma_i), \quad (5.1)$$

where $\mathcal{G}(\mu, \sigma)$ is a Gaussian with mean μ and standard deviation σ , and the vectors $\boldsymbol{\mu}$, $\boldsymbol{\sigma}$, are vectors populated by the individual μ_i and σ_i , and a_i are fitting coefficients. This function is then fit to the lineout for N peaks. Fitting to the lineout in this way was necessary due to the overlap between the beamlets, which can be substantial for beams

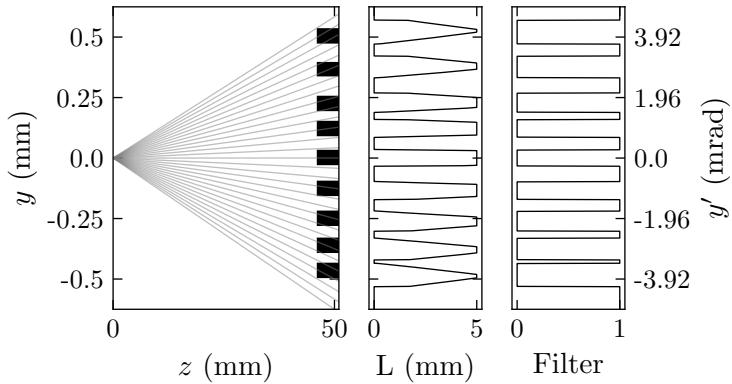


Figure 5.8: Calculation of the mask clear aperture. The left panel shows simulated test electron trajectories emanating from a point source. The central panel plots the length of scattering material encountered as a function of divergence. The right plot shows the binary filter used to determine the effective slit widths.

with a large source size. As can be seen in the lower panel of Fig. 5.7, this function fits the data well, indicating that modelling the beamlets as having Gaussian divergences is appropriate.

The divergence of a beamlet is related to its intrinsic emittance, but other factors also contribute and must be accounted for in order to determine the amount that is due solely to the emittance. These factors are the finite width of the slit, scattering in the tape, and incomplete scattering in the mask. The latter is not dealt with in this method, but shall later be shown to only have a small effect. The other two factors add in quadrature to increase the measured divergence,

$$\sigma_m^2 = \sigma_\varepsilon^2 + \sigma_W^2 + \sigma_S^2, \quad (5.2)$$

where σ_m is the measured divergence, σ_ε is the divergence due to the emittance, σ_W is due to the finite slit width, and σ_S is due to the tape scattering. The effect of finite slit width can be realised by considering a perfect beam, one with zero emittance, incident on a slit. This beam will produce a magnified image of the slit on the spectrometer screen, which is a top-hat function. This has an RMS value of $\sigma_W = \sqrt{12}Md$, where M is the magnification due to the distance between the screen and the mask, and d is the slit width.

Due to the length of the foils and the variation in their spacing, d must be determined individually for each slit. In Fig. 5.8, the method for doing this is shown. The emittance mask is modelled using the measurements provided by the mask manufacturer, which detailed the foil thicknesses and their spacing, and then the thickness of tungsten encountered along straight trajectories emanating from the source is calculated as a function of divergence angle. The slit width is the range of divergence angles which encounter no tungsten when viewed from the accelerator exit. It should be noted that, if the origin of the beam shifts, then these values will change slightly, becoming less reliable. This is a disadvantage of using a mask with this geometry, an issue that will be addressed

in the discussion. The peaks are then assigned a slit number and the slit width is used to correct for the effect of the finite slit width. The variation in the position of the mask shadow on the electron spectrometer screen is small enough that the locations of the various peaks can be determined from the high divergence beams, and then hardcoded into the analysis procedure. The value for σ_S is calculated using the length of Kapton traversed, $25 \mu\text{m} / \cos 10^\circ$, and the energy of the electrons being analysed. Knowing the energy allows scattering effects to be accounted for, which is one of the advantages of energetically resolved measurements. The emittance can now be calculated using the formula given in Sec. 3.4.4, using the value σ_ε for the beamlet divergence.

5.2.2 Sampling error

The mask has 8 slits in total, setting the largest number of source size measurements that can be made for a given beam. Further, the lowest divergences measured during the experiment were $\sim 0.5 \text{ mrad}$, which only illuminate 2 - 3 slits. Such a small number of samples can make the measured emittance sensitive to exactly which parts of the beam are transmitted by the mask. In addition, if the number of samples drops to 1, then the formula will place the average position of the beam exactly where the only measurement is being made, giving a measured emittance of 0.

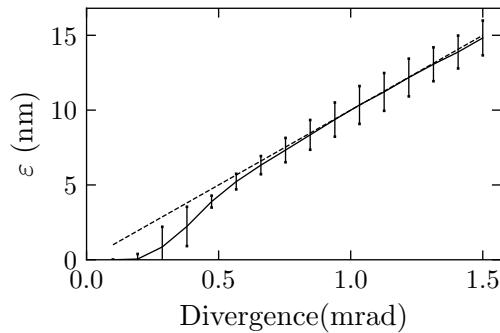


Figure 5.9: The effect of divergence on the systematic error of the emittance measurement. The measured emittance is shown in the black line, while the dashed line shows the simulated beam's input emittance.

This effect is shown in Fig. 5.9. Here, beams with a $10 \mu\text{m}$ source size and a random initial pointing are initialised and scattered off the tungsten mask model shown in Fig. 5.8. The lineouts they produce are measured using the method described above, returning a value for the emittance. Over the range of random pointings at each divergence, the average emittance and standard deviation are plotted in Fig. 5.9, with the error bars denoting the standard deviation of the measurements. For large enough divergences, there are sufficient samples, and although there is significant variation in the measurement due to the pointing fluctuations, the average of these measurements is equivalent to the true emittance. Once the divergence of the beam drops to the point where sometimes only a single beamlet is measurable, at approximately 0.5 mrad in Fig. 5.9, the average

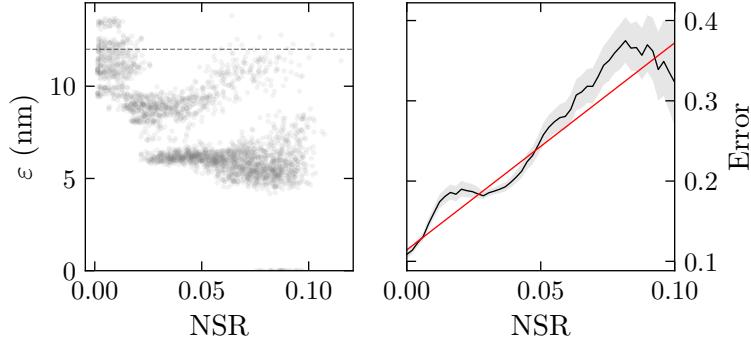


Figure 5.10: Left plot shows individual emittance measurements of a simulated beam for random input pointing angles, as a function of noise-to-signal ratio. The right panel performs a linear fit to the rolling standard deviation of these measurements. The grey shaded region is the error on the standard deviation, calculated using bootstrapping.

measured emittance begins to drop relative to the true emittance. This can be thought of as a resolution limit of the mask and should be considered when designing emittance masks.

The effect of noise must also be considered. Noise makes fitting the series of Gaussian peaks to the lineout less reliable, but also has an effect on the sampling number. Since peaks are detected relative to a threshold that is calculated using the noise floor, a noisier lineout will lead to fewer peaks being detected. To quantify these effects, simulated beams with emittances⁷ of 12 nm and random pointing angles have been scattered off the model mask. The lineouts produced have artificial noise added to them, and then an emittance measurement is attempted. The ratio between the standard deviation of the noise and the highest signal in the lineout, NSR, is calculated. The emittance at each NSR value is plotted in the left panel of Fig 5.10, where it can be observed that, as the NSR increases, the error on the emittance measurement increases, and the average value of the emittance measured also starts to decrease systematically. The increase in the range of emittance values is due to the increased difficulty of accurately fitting Gaussian peaks to the lineout, while the systematic shift is a manifestation of the effect shown in Fig. 5.9.

5.2.3 Monte Carlo fit

The scattering model that was used to measure the clear aperture of the emittance mask can also be used to measure the beam directly. If we assume that the electron beam particle coordinates have a Gaussian distribution in $y - py$ phase space, then we can scatter this beam off the virtual model of the tape and mask to obtain a simulated scatter pattern. By comparing the simulated pattern to the measured one, we can define a cost in terms of the difference between the two, and seek to minimise this cost using any number of algorithms. Since the number of free parameters is small⁸, yet individual simulations

⁷ Unnormalised at 1 GeV

are somewhat costly, Bayesian optimisation was chosen to tune the parameters. An example of this optimisation is performed in Fig. 5.11 for a single energy slice. After 300 iterations, each of which using 1×10^6 particles, the optimisation converges and a good fit is obtained. The small differences between the optimum and the data can be attributed to imperfect mask modelling.

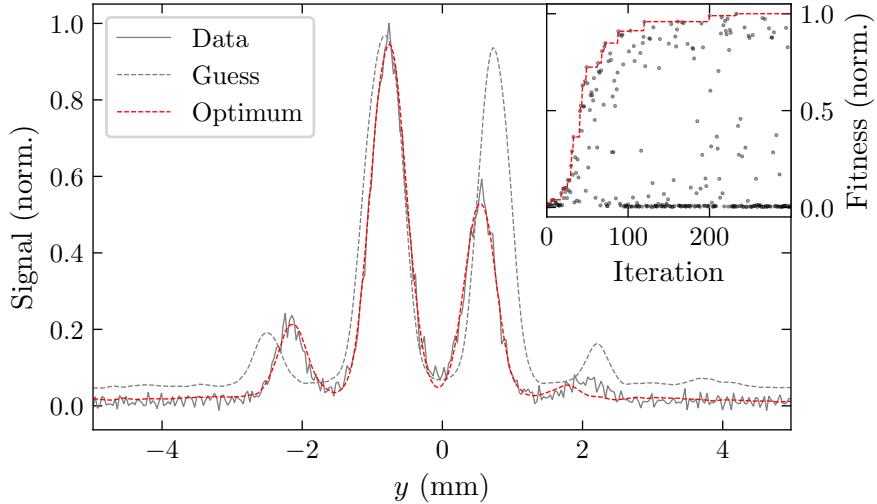


Figure 5.11: Alternate method for measuring the emittance. The main plot shows the experimental data, the initial guess, and the result with the highest fitness value. The inset shows the fitness as a function of iteration number, with the red line indicating the rolling maximum.

This analysis procedure produced results that were in agreement with the pepper-pot method, verifying the modifications made to account for the scattering in the tape. It also shows that the effects of partial scattering in the edges of the foil are not significant; the foils are functioning as almost perfect beam masks. Since the optimisation method takes substantially longer, it was not used for the bulk of the analysis. However, we will return to this method in the discussion, as it has the potential to improve on the pepper-pot method if the computational limitations can be overcome.

5.3 Experimental results

Now that the method for extracting the emittance measurements from the electron spectrometer images has been explained, the experimental measurements can be presented. In this section, we will first present two example measurements before proceeding to show the results of a parameter scan, where the gas jet backing pressure and the height above the nozzle were varied.

⁸ There were five parameters, the two Gaussian phase space widths, the initial beam position and pointing, and a shift associated with the background subtraction performed on the data.

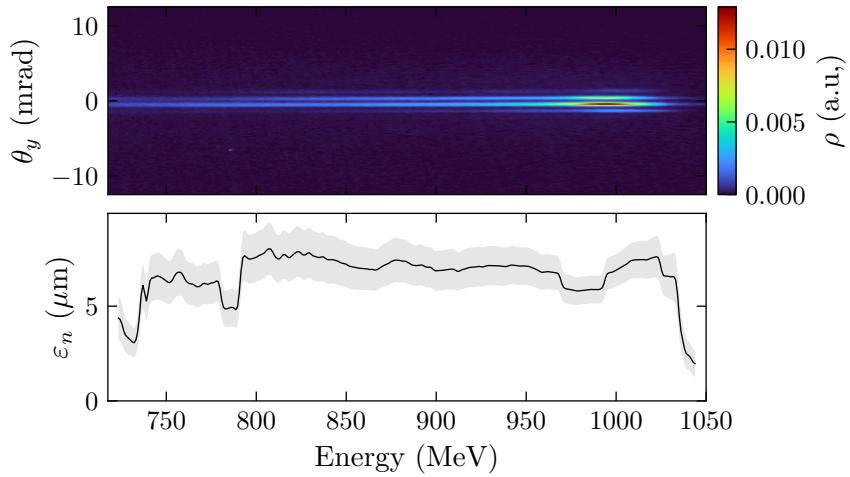


Figure 5.12: Normalised emittance measurement, for a gas jet backing pressure of 40 bar, producing a plasma density of $1.53 \pm 0.06 \times 10^{24} \text{ m}^{-3}$.

5.3.1 Example measurements

The normalised emittance measurement for the example image is shown in Fig. 5.12. This measurement was taken 5 mm above the nozzle, with a backing pressure of 40 bar, which produced an average plasma density of $1.53 \pm 0.06 \times 10^{24} \text{ m}^{-3}$. The measurement was performed using an energy bin width of 10 MeV, with measurements taken at every energy. This was performed up to the point at which the signal of the transversely integrated spectrum fell to within 3σ of the background. The normalised emittance is then calculated by multiplying the phase emittance by the γ of the energy slice. The weighted median emittance for this measurement was $\varepsilon_n = 6.9 \pm 0.6 \mu\text{m}$, where the error was calculated using the charge weighted median average deviation from the emittance measurement.

The grey shaded region shows the error on the individual energy slice measurements, which uses the signal-to-noise ratio energy scaling derived in Fig. 5.10. The effects of sampling error can be clearly observed in this measurement, since, as the charge of the beam changes, different numbers of peaks are detected. This causes jumps in the emittance measurement, which do not appear to correspond to substantial changes in the beam profile. These jumps are larger in magnitude and more frequent at places where the signal is low. As can be seen by the analysis in Fig. 5.9, when the divergence drops below approximately 0.5 mrad, the measured emittance drops spuriously, and this effect is also observed at energies below ~ 750 MeV and above 1025 MeV. However, since the charge is low in these regions, the charge-weighted median of the emittance measurement is insensitive to these errors for this shot.

The emittance measurement of a beam that over-fills the mask is shown in Fig. 5.13. This measurement was taken with a backing pressure of 80 bar, shooting at 5 mm above the nozzle, producing an average plasma density of $2.86 \pm 0.08 \times 10^{24} \text{ m}^{-3}$. This generated a beam with a thermal spectrum, where electron energies up to approximately 1 GeV

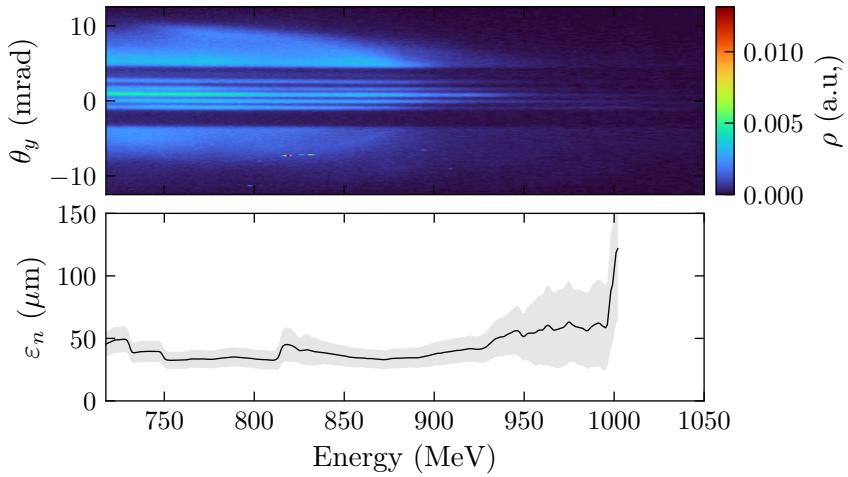


Figure 5.13: Normalised emittance measurement, for a gas jet backing pressure of 80 bar, producing a plasma density of $2.86 \pm 0.08 \times 10^{24} \text{ m}^{-3}$.

were detected, and a median emittance of $35 \pm 6 \mu\text{m}$. For beams that overfill the mask, the emittance was calculated by extrapolating the measurements of the masked regions to the unmasked regions. This was done by first ignoring the regions of the beam that are outside of the mask, and calculating the emittance for the masked region only. This measurement is then scaled by the divergence of the entire beam. This is equivalent to using the mask to measure the source size, and then using the divergence and source size to calculate the emittance. Since extrapolation is involved, the reliability of this measurement is reduced. However, due to the increased number of samples of the beam, the emittance of the central region is characterised more accurately, which partially compensates for the loss of reliability due to extrapolation. The error on the measurement grows substantially above $\sim 950 \text{ MeV}$. This is due to the low signal in this region, but an emittance measurement is still attempted since the integrated signal is above the background threshold.

5.3.2 Density scan

A series of scans were performed where the height above the nozzle and the gas jet backing pressure was varied. Backing pressures of 40, 60 and 80 bar were used, while 4 heights in the range from 5 mm to 11 mm were used. At each setting, 10 shots were taken, which was necessary since the electron beam was intermittent and had substantial pointing fluctuations, meaning that it only occasionally produced a measurable pattern on the electron spectrometer screen. The spectrally resolved emittance measurements obtained from this scan are plotted in Fig. 5.14. From inspection, it is clear that many of the measurements are too noisy and discontinuous to warrant single-shot analysis, and so the shots at each setting have been averaged to produce a single spectrally resolved emittance measurement for that setting, shown by the red lines. The charge was used

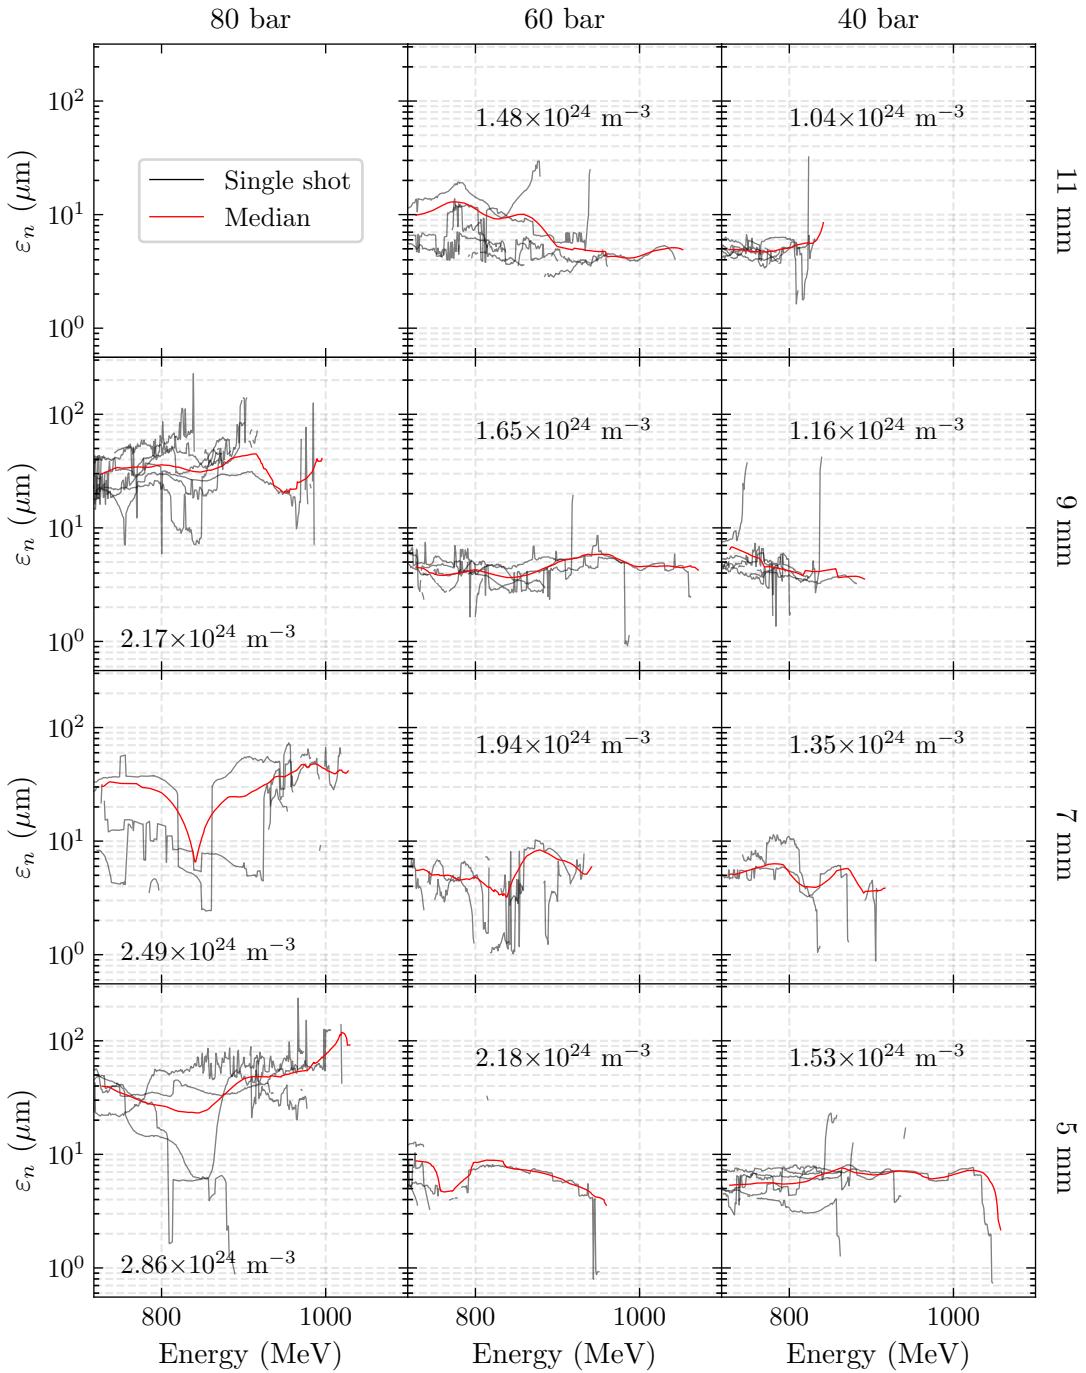


Figure 5.14: Summary of all emittance measurements. Each plot corresponds to a particular target configuration, with height above the nozzle increasing, moving from low to high, and gas pressure decreasing from left to right; the coordinates are labelled on the top and right of each column and row, respectively. The grey lines are the single-shot measurements, while the red lines are the charge-weighted average measurements. No satisfactory emittance measurements were obtained at (80 bar, 11 mm), due to lack of signal.

to weight the average, reducing the effect of slice emittance measurements that are spuriously low due to low charge⁹.

Examining the scan as a whole, the measurements bifurcate; low densities tend to produce average emittances lower than 10 μm , while the higher density measurements produce beams with average emittance above 20 μm . In addition, for the latter group, the emittance gradually increases as a function of energy, while for the lower-density group, emittance falls slightly as a function of energy. Within the low-density group, further reducing the density appears to reduce the emittance while also reducing the peak electron energy, which is indicated by the energy at which the measurements stop. In the high-density group, there is no clear trend in the emittance as the energy increases.

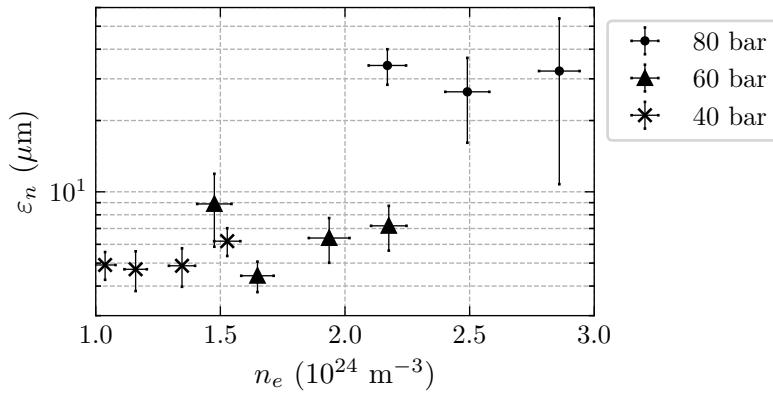


Figure 5.15: Charge weighted median emittance values for each density and gas jet backing pressure. The error bars represent the charge-weighted standard deviation of the slice emittance measurements.

The charge-weighted median emittance for each setting is plotted in Fig. 5.15. The lowest median emittance was $4.42 \pm 0.60 \mu\text{m}$, with the error value representing the standard deviation of the slice emittance measurements, in a plasma with a density of $1.65 \pm 0.06 \times 10^{24} \text{ m}^{-3}$, while the highest emittance was $34 \pm 6 \mu\text{m}$, which used a plasma density of $2.17 \pm 0.07 \times 10^{24} \text{ m}^{-3}$. This would indicate that the emittance is particularly sensitive to density changes in this region, which may be true. However, the variable of the gas jet backing pressure also changed between these two measurements; from 60 bar to 80 bar, and the height above the nozzle was the same, at 9 mm. Further, a measurement at $2.18 \pm 0.07 \times 10^{24} \text{ m}^{-3}$ with a pressure of 60 bar was made, which produced a median emittance of $7.2 \pm 1.3 \mu\text{m}$. Taken together, these results indicate that the backing pressure has an independent effect on the emittance.

⁹ See Sec. 5.2.2.

5.4 Simulations

5.4.1 Simulation setup

Simulations of the interaction have been performed using FBPIC¹⁰. This has been done to provide a comparison to our results, as well as to determine the independent effect of density on emittance for our experimental parameters. The simulations used a $100\text{ }\mu\text{m} \times 100\text{ }\mu\text{m} z \times r$ moving window with a velocity equal to the group velocity of the laser in the bulk plasma. The grid resolution was 20×4 points per laser wavelength in the longitudinal and transverse dimensions respectively, and three azimuthal modes were used. The gas jet target is initialised as a 12.8 mm long region of constant density gas, with exponential density ramps at both ends. The ramps had scale lengths of $150\text{ }\mu\text{m}$, a value that was determined from interferometry. The helium atoms were initially singly ionised, while the nitrogen atoms were ionised up to $+5$. The density of the helium and nitrogen atoms was determined by stipulating that the plasma density was achieved when the helium was fully ionised and the first 5 electrons of the nitrogen were ionised.

A total of 48 macroparticles per cell was used for the helium, while only 1 macroparticle per cell was used for the nitrogen. This was done to reduce the discrepancy in each species' resolution. Helium has 0.29 times the weight of nitrogen yet composes 98% of the gas mixture by weight, meaning that the ratio of helium atoms to nitrogen atoms is 344. This means that even with 48 times fewer particles per cell, the nitrogen atoms are still represented by approximately seven times more macroparticles. For example, in the highest density simulation performed, a total of 234 pC of charge from the bulk plasma was injected and accelerated to energies above 10 MeV, while 591 pC of charge was injected through the ionisation of N^{5+} ions. The electrons injected from the plasma were represented by 2.7×10^4 macroparticles, while the ionisation injected charge was represented by 5.8×10^6 macroparticles. This means that each simulation particle represents 1.4×10^5 electrons in the self-injected beam, but 1.8×10^3 electrons in the ionisation injected beam.

The laser pulse is initialised using a real focal spot image, using the method outlined in Sec. 3.5.4. Although focal spot data was not available for the day on which the data in this chapter was recorded, many focal spot images had been taken on other days, and the day-to-day variation was not significant. The focal spot image for these simulations was taken from this collection, which had a FWHM of $41\text{ }\mu\text{m}$ and a peak a_0 of 1.724. The temporal envelope was Gaussian and had a width of 38.1 fs. The focal position of the laser was set to the point at which the density reached 90% of its plateau value.

Simulations were performed at 5 density intervals between $1 \times 10^{24}\text{ m}^{-3}$ and $3 \times 10^{24}\text{ m}^{-3}$. All electrons in the simulation were tracked so that the injection mechanism could be studied in detail. The emittance was calculated using the median average deviation of the particle coordinates.

¹⁰ [216] R. Lehe *et al.* (2016) *Computer Physics Communications*

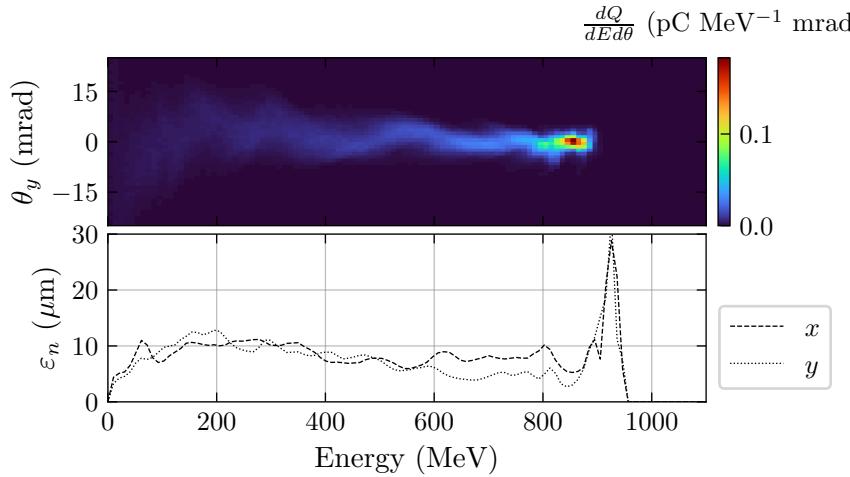


Figure 5.16: The top panel shows a synthetic electron spectrometer image from the low-density $1 \times 10^{24} \text{ m}^{-3}$ simulation. The lower panel shows the spectrally resolved emittance in x and y . The laser is polarised in the x direction.

5.4.2 Simulation results

In the top panel of Fig. 5.16, a synthetic electron spectrometer image from the $1 \times 10^{24} \text{ m}^{-3}$ simulation is shown, below which the spectrally resolved emittances are plotted. A quasi-monoenergetic beam is produced, where the energy peak occurs at 900 MeV, and the divergence is 1.7 mrad for the 55 pC of charge above 725 MeV¹¹. The similarity between this and the results presented in the right panel of Fig. 5.6 implies that the simulation is modelling the experiment reasonably well and that comparisons between the simulated beam's emittance and the experimental measurements are appropriate. The spectrally resolved emittance was calculated using a bin width of 20 MeV, and the number of macroparticles was sufficient that the standard deviation and median average deviation metrics gave the same measurements. The betatron oscillations in the spectrum, which were not present when using a Gaussian transverse profile for the laser, did not affect the spectrally resolved emittance significantly with this bin width. For the energies above 725 MeV, the emittance of the simulated beam in the y direction is approximately 5 μm , while in the x direction the emittance is larger, at approximately 9 μm . The spike in the emittance that occurs at the highest energy is due to a small number of high-energy electrons that have interacted with the laser pulse. Note that at this density, all electrons are injected via ionisation injection.

The synthetic electron spectrometer image from the $3 \times 10^{24} \text{ m}^{-3}$ simulation is shown in the upper panel of Fig. 5.17, where different colours have been used to denote the electrons from nitrogen and those from the background plasma. The opacity of the overlaid beam in blue has been reduced to make both beams visible in the plot. The total charge above 725 MeV was 294 pC, with 1.8×10^6 macroparticles modelling the 214 pC coming from the nitrogen, compared to the 80 pC coming from the plasma electrons, which had only 1×10^4 macroparticles. This discrepancy in the resolution of the species

¹¹ Threshold chosen to match the low energy end of the experimental emittance measurements.

is despite using many more macroparticles per cell for the background plasma than the dopant. The particles injected from the nitrogen have a large divergence of 5.8 mrad, reducing the charge density on the spectrometer, while the plasma electrons have a low energy slice divergence of 2.0 mrad, but with spectrally resolved pointing fluctuations that exceed ± 15 mrad. Again, this is qualitatively similar to the experimental measurements at high density, shown in the left panel of Fig. 5.6. From comparison with the simulations, we can now identify that the high divergence part of the beam is from ionisation injection, while the low divergence core of the beam is from background plasma electron injection.

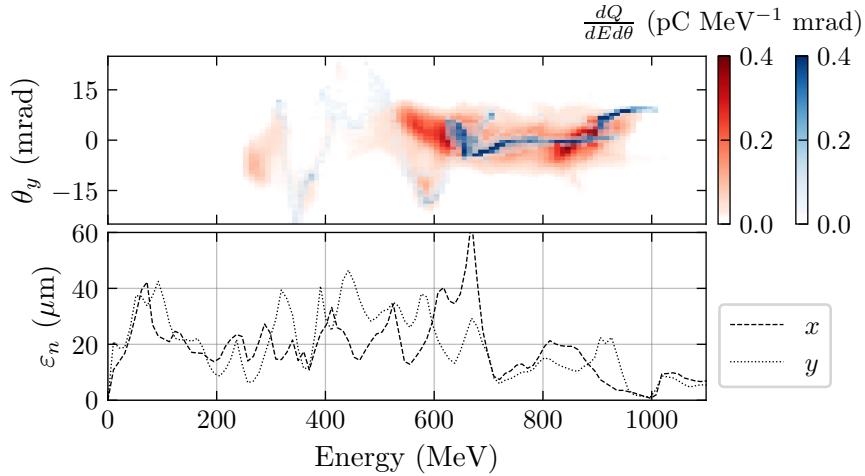


Figure 5.17: Synthetic electron spectrometer image (top) and emittance measurements (bottom) from the $3 \times 10^{24} \text{ m}^{-3}$ simulation. Different colour maps have been used for the different electron origins; red for electrons from nitrogen K-shells and blue for electrons from the bulk plasma.

The spectrally resolved emittance of the beam produced in the $3 \times 10^{24} \text{ m}^{-3}$ plasma is shown in the lower panel of Fig. 5.17. This is the combined emittance of both the ionisation-injected beam and the self-injected beam, although the former dominates the measurement due to its higher charge. The energy slice emittance of the plasma electron injected beam is lower than that of the ionisation injected beam, and consequently, its inclusion slightly reduces the emittance measurement. The large emittance in the low energy part of the spectrum is due to the observed pointing fluctuations, which are present in both beams. Although the emittance is lower above 725 MeV, it is still higher than the comparative part of the other beam; $(\varepsilon_x, \varepsilon_y) = (7.5 \pm 2.5, 4.8 \pm 1.5) \text{ } \mu\text{m}$ from the low-density plasma, versus $(11.4 \pm 8.0, 11.5 \pm 3.5) \text{ } \mu\text{m}$ in the high-density plasma.

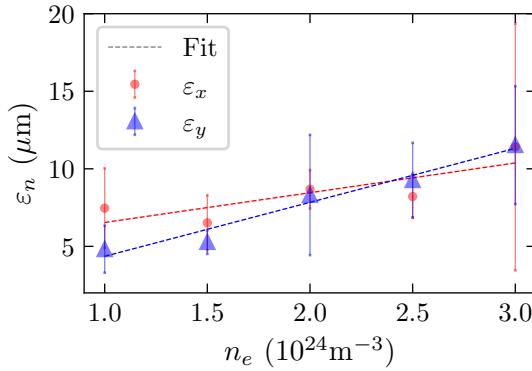


Figure 5.18: Charge weighted mean emittances for energies above 725 MeV in the x and y planes for a range of simulations. The error bars denote the charge-weighted standard deviation.

The mean emittances for the 5 densities simulated are plotted in Fig. 5.18. A linear fit to the data yields gradients of 1.9 for ε_x and 3.5 for ε_y , with correlation coefficients¹² of 0.82 and 0.98 respectively. The low density $< 2 \times 10^{24} \text{ m}^{-3}$ measurements in the y plane agree closely with the experimental measurements, indicating that the transverse beam dynamics are being simulated accurately over this energy range. In contrast, for densities $> 2 \times 10^{24} \text{ m}^{-3}$, these simulated emittances are much lower than those measured experimentally. This is also the density at which significant numbers of electrons begin to be injected from the background plasma, as opposed to ionisation-injected electrons being the most common.

5.5 Discussion

In this section, we will use the low-density simulations to determine the source of the emittance. The measurements will then be compared to the literature, and the advantages and disadvantages of the various emittance measurement methods will be discussed. Finally, possible developments of this diagnostic method are discussed.

5.5.1 Emittance source

The laser power is above the critical power for self-focusing, which is given by $P_c \approx 17.5\gamma_p^2 \text{ GW}$, such that for our 150 W laser, $P/P_c \approx 5$ at a plasma density of $1 \times 10^{24} \text{ m}^{-3}$. The matched spot size is $w_m \simeq 2.94k_p^{-1}$, equal to $15.6 \mu\text{m}$, which is a factor of approximately $\sqrt{2}$ lower than the vacuum spot size. The a_0 is boosted from 1.7 to 2.4 via focusing alone, and which then undergoes further amplification via longitudinal compression¹³. This means that the laser is usually intense enough to cause near complete cavitation of electrons in its wake, providing an emittance preserving linear focusing structure for the electron beam.

¹² Pearson's r value.

¹³ [146] M. Streeter *et al.* (2018) *Physical Review Letters*

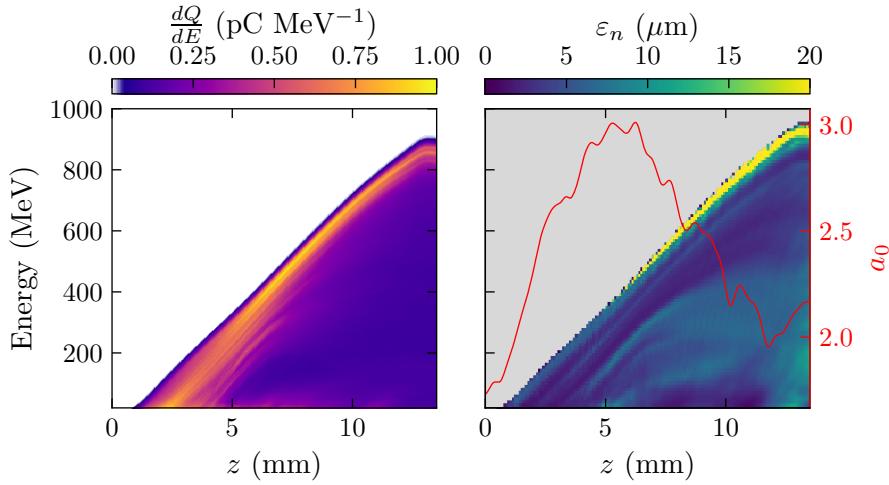


Figure 5.19: History of the electron beam energy spectrum (left) and spectrally resolved emittance (right) in the y direction. The right plot also shows the evolution of the laser a_0 .

The evolution of the energy spectrum and the emittance of the beam as it accelerates is plotted in Fig. 5.19. The high energy electron bunch appears to be accelerated at a constant rate, an effect that is due to a combination of the long dephasing length at this density (approximately 15 mm) and the rephasing of the electron beam as a_0 drops and the bubble shrinks. This coincides with some beam loss, which begins at $z = 5$ mm, as electrons close to the back of the bubble leave the accelerating and focusing portion of the wake. The emittance of electrons with the highest energies starts to increase rapidly at approximately 6 mm. However, this growth is limited to electrons that lie outside the energy peak, with a total charge of 0.01 pC above 900 MeV, compared to the 55 pC above 725 MeV. This emittance growth is due to the interaction of electrons with a large displacement from the axis interacting with the electron sheath of the wave.

For the high charge, high energy peak of the beam, almost all of the beam's emittance is generated at the moment of injection. We will now use a simple phenomenological argument to predict this value. In ionisation injection, an electron's transverse energy is determined by the potential into which it is ionised. In the direction transverse to the laser polarisation, only the plasma wave potential contributes to this, in contrast to the direction aligned with the laser polarisation, where electrons are heated by the instantaneous electric field of the laser, which provides an additional source of transverse energy. In the bubble regime, the potential is quadratic with distance from the axis of the wake, $\phi = en_e y^2/3\epsilon_0$, where y is the transverse distance from the axis of the wake. From the solution to Eq. 2.88, $y(t) = y_0 \cos(\omega_\beta t)$, where we recall that $\omega_\beta = \omega_p/\sqrt{2\gamma}$. Using the Lorentz force, $\partial_t p_y = e\nabla\phi$, the transverse momentum RMS is given by

$$\sigma_{p_y} = \frac{2\sqrt{2\gamma}}{3} \frac{n_e e^2}{\epsilon_0 \omega_p} \sqrt{\langle y_0^2 \rangle}, \quad (5.3)$$

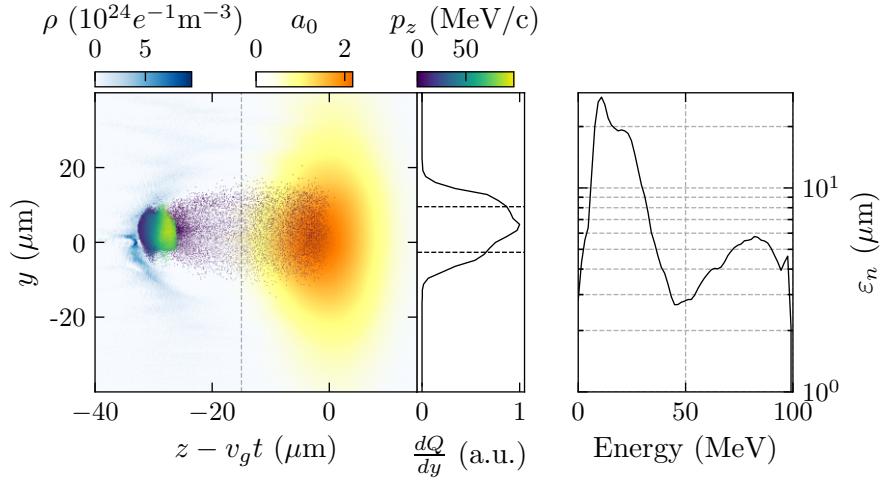


Figure 5.20: Injection process in the $1 \times 10^{24} \text{ m}^{-3}$ plasma. This snapshot is taken at the moment the highest energy electrons are injected, at $t = 6.0 \text{ ps}$. The histogram on the right is of the y coordinates of the electrons to the right of the dashed line, which has an RMS width of $6.1 \mu\text{m}$. The plot on the right shows the emittance of the beam.

The geometric emittance at injection ($\gamma = 1$) is then

$$\varepsilon_{ph} = \frac{2\sqrt{2}}{3} k_p \langle y_0^2 \rangle. \quad (5.4)$$

Since electrons are initially stationary in the laboratory frame, this is also the normalised emittance. At a plasma density of $1 \times 10^{24} \text{ m}^{-3}$, a normalised emittance of $5 \mu\text{m}$ would be given by a beam size of $y_\sigma \approx 5.3 \mu\text{m}$ using this scaling.

We can use the simulations to estimate y_σ . A snapshot of the injection process is shown in Fig. 5.20, where this snapshot is taken at $t = 6.0 \text{ ps}$; the moment the electrons that produce the high energy peak in Fig. 5.16 are ionised. At this moment, the peak a_0 of the laser is 2.18, which drives a plasma wave that is close to a fully cavitated bubble. The electrons produced by the ionisation of N^{5+} nitrogen ions are plotted as scatter points, with a colour that depends on their longitudinal momentum in the laboratory frame. Note that, since the laser polarisation points in the x direction, the electrons are initially cold in the $y - z$ plane. Since the peak a_0 of the laser is slightly below the intensity required for barrier suppression ionisation, these electrons are ionised via the tunnelling mechanism, the rate of which depends strongly on the laser intensity. Hence, most electrons are ionised close to the peak of the laser intensity, and this value falls off radially. The RMS beam size at the moment of ionisation is $6.1 \mu\text{m}$, which, using the simple model we have derived, would predict a normalised emittance of $2.7 \mu\text{m}$. This is equivalent to the lowest value of the injected beam's emittance, seen in the right panel of Fig. 5.20, indicating that the model is accurately predicting the lower bound of the emittance. An additional source of emittance comes from early injection, into a wake structure that is undergoing formation. Fig. 5.21 shows this early injection, where the wake is capable of capturing electrons, but the focusing forces are non-linear. To avoid

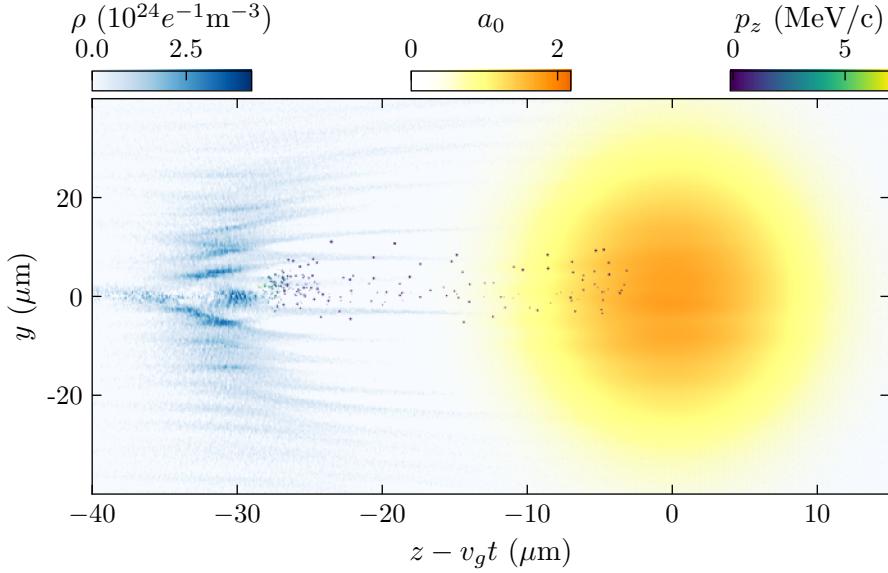


Figure 5.21: Injection into a wake as it forms, resulting in higher emittance. This snapshot is at $t = 1.5$ ps.

this early injection, and to ensure a low emittance at all energies, it would be beneficial to have a region of the plasma without the dopant present, to allow a stable wake to form before injection occurs.

We can now explain the low emittance electron beams produced in the experiment and simulation, but this does not explain why the emittance of the beams produced using 80 bar of backing pressure have such large emittances. One possible explanation is the development of ripples in the gas flow, which disrupt the stable structure of the wake, increasing emittance. For example, Kuschel *et al.*¹⁴ have used high-resolution shadowgraphy to compare the structure of the plasma wake in gas jets to gas cells and found that the plasma wave can have an irregular wavelength in gas jets. They attribute this variation to ripples in the gas flow that are too small to be resolved by density measurements. This spatially varying wavelength causes self-injection to occur at lower densities than predicted¹⁵, a form of downramp injection, and also increases the beam's divergence, as measured by a beam profile monitor. It is possible that a similar process is occurring here and that the backing pressure threshold for the onset of the ripples occurs somewhere between 60 bar and 80 bar.

Although the lower backing pressure measurements had lower emittances, it is possible that these values could also be reduced further by using a gas cell instead of a jet. Aside from this, to further reduce the emittance in the polarization-perpendicular plane, a reduced ionisation radius would be necessary. To this end, two-colour ionisation injection schemes have been suggested¹⁶, where a second, higher frequency laser pulse is focused

¹⁴ [180] S. Kuschel *et al.* (2018) *Physical Review Letters*

¹⁵ [153] A. G. R. Thomas (2010) *Physics of Plasmas*

¹⁶ [167] L. L. Yu *et al.* (2014) *Physical Review Letters*

into the bubble behind the laser pulse driving the plasma wave. Only this second laser pulse is capable of ionising electrons near its focus, and as such, the emittance could be substantially reduced by the small focal spot of the second laser. However, such a scheme comes at the cost of reducing the charge. Maintaining the charge while reducing the emittance would necessitate a large ionisation radius and a shallower potential well, such as those produced in lower-density plasmas, and this would need to be balanced against the increased trapping difficulty. Carefully structured gas targets, with different gas densities for the injection and acceleration regions, as well as localised dopants, would be necessary to achieve this.

5.5.2 Literature comparison

In Fig. 5.22 the experimental measurements presented here are compared to a non-exhaustive list of measurements from the literature¹⁷. The beam mask methods encompass measurements that rely on some type of pepper-pot technique, either in 1D or 2D, while the quadrupole methods use a combination of magnetic lenses to re-image the beam at the exit of the accelerator onto a high-resolution detector. The x-ray source size measurement used the shadow of a knife edge in the x-ray profile to estimate the size of the beam in the plasma channel, while the x-ray spectroscopy measurements use a combination of the x-ray spectrum and the electron spectrum to infer an oscillation radius, and hence a source size. In general, the beam mask measurements have been used to measure emittance values of more than 5 μm , at energies of 100 MeV. The quadrupole methods have recorded lower emittance values of $\sim 1 \mu\text{m}$, while the x-ray methods have recorded emittances as high as 10 μm , by Kneip *et al.* and as low as 0.056 μm , by Qin *et al.* Except for the Weingartner *et al.* study, which used a gas cell, and the work by Plateau *et al.*, which used a capillary discharge, all the measurements used gas jets to provide the plasma for the target. This is presumably due to their ease of operation, which was discussed in Section 5.1. As implied by Fig. 1.3, the energy of the electron beams roughly correlates to the laser power used in each experiment; the Sears 2010 used a 10 TW laser, the Kneip 2012 paper used a 70 TW system, and the Qin 2018 result used powers of 120 TW. Relative to the laser polarization, the Weingartner 2012 measurement was perpendicular to the laser polarisation, while the Barber 2017 measurements were parallel, and the other measurements were either unstated (Sears, Qin) or relied on some measurements from one plane and other measurements from another (Kneip).

¹⁷ [228] S. Fritzler *et al.* (2004) *Physical Review Letters*
 [81] C. M. S. Sears *et al.* (2010) *Physical Review Special Topics - Accelerators and Beams*
 [200] E. Brunetti *et al.* (2010) *Physical Review Letters*
 [82] S. Kneip *et al.* (2012) *Physical Review Special Topics - Accelerators and Beams*
 [169] R. Weingartner *et al.* (2012) *Physical Review Special Topics - Accelerators and Beams*
 [229] G. R. Plateau *et al.* (2012) *Physical Review Letters*
 [230] S. Barber *et al.* (2017) *Physical Review Letters*
 [231] A. Curcio *et al.* (2017) *Physical Review Accelerators and Beams*
 [84] Z. Qin *et al.* (2018) *Physics of Plasmas*

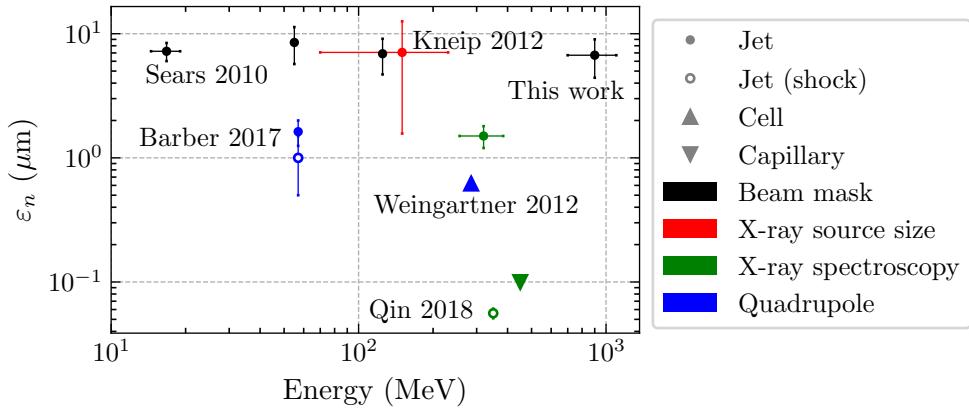


Figure 5.22: Laser wakefield acceleration beam emittance measurements from the literature. Except for this work, and the gas jet work performed by Barber *et al.*, all experiments used either pure helium or pure hydrogen targets. The error bars represent the range of emittance values recorded, and the energy range studied.

To our knowledge, the energy at which the measurements in this chapter were recorded is higher than any other in the literature. They are also the lowest emittance measurements among the other beam mask method measurements and are the only spectrally resolved measurements amongst this subset. In terms of injection method, the only other study that investigated ionisation injection was the work by Barber *et al.*, which was performed at an energy that is over an order of magnitude lower.

The different measurement methods used in the literature each have strengths and weaknesses. Arguably, the quadrupole method is the gold standard, since this directly measures the minimum beam size that can be obtained; one of the primary reasons for caring about the emittance in the first place. However, the requirement for a beam optic can make this method unattractive, due to both the space that needs to be reserved for the optic in the vacuum chamber, and the increased complexity involved with its installation. It is also normally limited to some extent by the field strengths it can produce, and the proximity to the target that can be achieved given mechanical constraints¹⁸. The x-ray measurements have a naturally high resolution owing to either the relative ease of stopping x-rays compared to electrons (in the case of the source size method), or the accuracy with which the x-ray and electron energy spectra can be determined. The downside of these measurements is that they often rely on a downstream electron beam measurement to determine the divergence of the source, while the source size is estimated while the beam is inside the channel, meaning that since the transition to vacuum is not properly accounted for, these methods underestimate the emittance.

¹⁸ [85] A. Cianchi *et al.* (2013) *Nucl. Instrum. Methods Phys. Res. A: Accel. Spectrom. Detect. Assoc. Equip.*

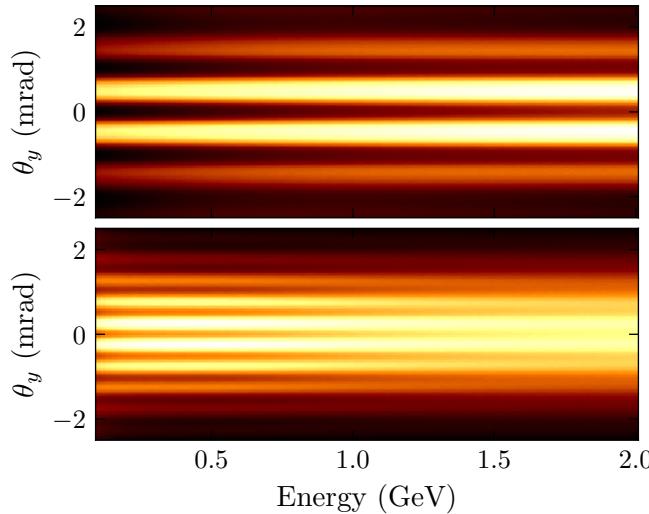


Figure 5.23: Upper panel shows MC simulated electron spectrometer image for current mask geometry. The bottom panel shows the scattering pattern for an alternative mask geometry.

5.5.3 Diagnostic development

The diagnostic used in this chapter was able to measure the emittance of beams that were aligned with the laser axis and had a low divergence. However, if the divergence caused the beam to overfill the mask, the measurements had to be extrapolated, which reduces the accuracy significantly. Further, since the size of the gaps depended on the initial position of the electron beam, pointing fluctuations of the laser introduce a source of error. Finally, the low sampling rate can lead to large errors, and if the sampling rate dropped too low then the measurement could not be made at all.

One of the main weaknesses of the pepper-pot technique is that it relies on the complete scattering of the beam. This places restrictions on the thickness of the mask, as a thicker mask will be required as the electron energy increases. At GeV energies, millimetre scale thicknesses are required, while, to achieve the required magnification, the mask needs to be placed where the beam is only a few 100s of microns in diameter. Therefore, to use the pepper-pot method with a sampling rate of 10, the ratio between the mask thickness and the spacing needs to be approximately 100, which means the foils will be liable to warping. There are two solutions to this problem for a fixed vacuum chamber size; either place the mask further away, and image the spectrometer pattern with a higher resolution, or change the analysis method. Imaging the spectrometer pattern with higher resolution would require switching to a higher resolution scintillating material, which, while suitable for small energy ranges¹⁹, becomes prohibitively expensive if a large energy range is desired.

An alternative is to use a different analysis method. In this chapter, we have used simulations of the electron beam scattering off the mask for various purposes, such as

¹⁹ [230] S. Barber *et al.* (2017) *Physical Review Letters*

determining the clear aperture of the mask and verifying the measurements made by the pepper-pot technique. However, with sufficient computational resources, the electron beam phase space could be determined with any scattering pattern, complete or not, provided the structure of the mask is known. This lifts the requirement for complete scattering, making a high sample rate possible. For example, a thin tungsten sheet with laser-machined slits could be used, allowing for narrower, more closely spaced slits. In Fig. 5.23, the scattering caused by a $10\text{ }\mu\text{m}$ foil with $10\text{ }\mu\text{m}$ slits is compared to the mask used in this experiment. Even though the scattering is almost imperceptible at high energies, the emittance can still be extracted because some modulation has occurred.

5.6 Summary

In this chapter, an emittance measurement diagnostic based on the pepper-pot method has been developed and used to measure the emittance of GeV beams. At low gas jet backing pressures of either 40 or 60 bar, the measured normalised emittances were in the range of $4 - 9\text{ }\mu\text{m}$, while at 80 bar, the emittance measurements exceeded $30\text{ }\mu\text{m}$. The lower pressure emittance values closely agreed with simulations, with the emittance being primarily due to the radial distance from the axis at which the electrons were ionised. The large emittance values were attributed to ripples in the gas flow that become significant at higher backing pressures. The low emittance values measured are the lowest reported in the literature using a mask method, at energies over an order of magnitude higher. Finally, improving the resolution of the diagnostic using MC simulation-based analysis was discussed.

6 Beam divergence considerations for staged LWFA

THE STAGING of two independent 250 TW laser wakefield accelerators was investigated during a 2019 experiment at the Astra-Gemini laser facility. While multi-stage acceleration has already been demonstrated¹, it has not been done at the GeV level or with a large capture efficiency. Although energy gain due to the second accelerator was not achieved, numerous experiments were performed on the dual-stage target to find optimal operating conditions.

The results of one of these experiments are reported in this chapter, where it was found that the target configuration at the end of the first cell had a significant effect on the divergence of the electron beam. This divergence change was immediately noticeable during the experiment, but the primary mechanism behind it was not known. A large divergence is undesirable for most applications, and particularly so for multi-stage accelerators because the divergence will accelerate emittance growth that occurs between stages, lowering the beam quality. In addition, if there is no focusing beam optic between the stages, as was the case for our experiment, a large divergence will cause the beam to overfill the wake in the second stage, reducing the capture efficiency. Understanding the dominant cause for this change in divergence is important for making design modifications for future multi-stage acceleration experiments. More generally, understanding all of the factors that affect the beam quality between stages becomes increasingly important as the number of stages increases, as even small effects could become significant if they are allowed to compound.

In Section 6.1, the experimental setup and target configurations are described, and the experimental measurements are reported. Following this, two potential causes of divergence increase are investigated using simulations. The first of these is the effect of magnetic fields generated in the tape by the laser, reported on in Section 6.2, while the second is the effect of longitudinal density ramps on the electron beam envelope, reported on in Section 6.3. Finally, in Section 6.4, the preceding analyses are combined to determine the relative contributions of the different mechanisms to the observed divergence changes. The experimental measurements are then used to estimate the effect

¹ [88] S. Steinke *et al.* (2016) *Nature*

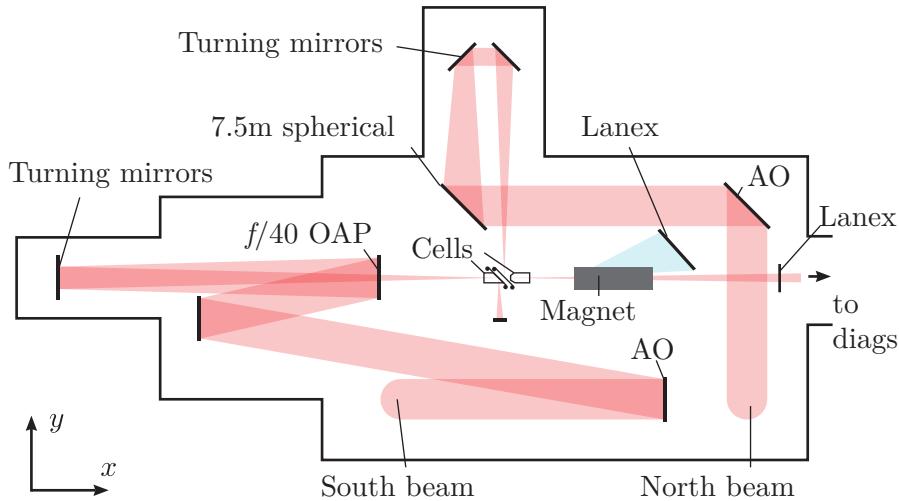


Figure 6.1: Sketch of the vacuum chamber layout. The laser beams enter the vacuum chamber from above (propagating into the page). The two Lanex screens indicate the positions of the electron beam diagnostics. Note that the physical beam paths were not exactly as sketched; the angle of incidence onto the 7.5 m spherical was not 45° .

the divergence increase would have on the capture efficiency of a second stage. Control of the electron beam envelope between stages is also discussed, where density ramps with a range of scale lengths and different functional forms are compared.

6.1 Electron beam measurements

6.1.1 Experimental setup

The layout of the laser beams and the positioning of the electron beam diagnostics is shown in Fig. 6.1. The South beam was focused over a distance of 6 m by an off-axis parabola (OAP), while the North beam was focused over 7.5 m by a spherical mirror. This constituted a focal geometry of $f/40$ and $f/50$ for the South and North beams, respectively. Turning mirrors were used between the focusing optics and the target to allow this geometry to fit inside the vacuum chamber. Since the intensity of the beam was approximately four times higher on the turning mirrors than the other optics due to the halved beam diameter, these optics required regular replacement due to damage caused by the laser.

Each laser was focused into separate gas cells using a configuration shown in Fig. 6.2. The South beam was responsible for the injection and initial acceleration of the electron beam, while the North beam was to provide the secondary energy boost. Two independent Kapton tapes were used between the two cells to facilitate the coupling between the lasers and the electron beam. Upon encountering the tapes, the focused lasers ionise the surface, producing an above-critical density plasma that functions as a plasma mirror. The first tape (extraction tape) was used to dump residual laser energy at the end of the first cell, while the second tape (injection tape) injected the North beam into the second cell. The extraction tape also served the secondary role of sealing the end of the first

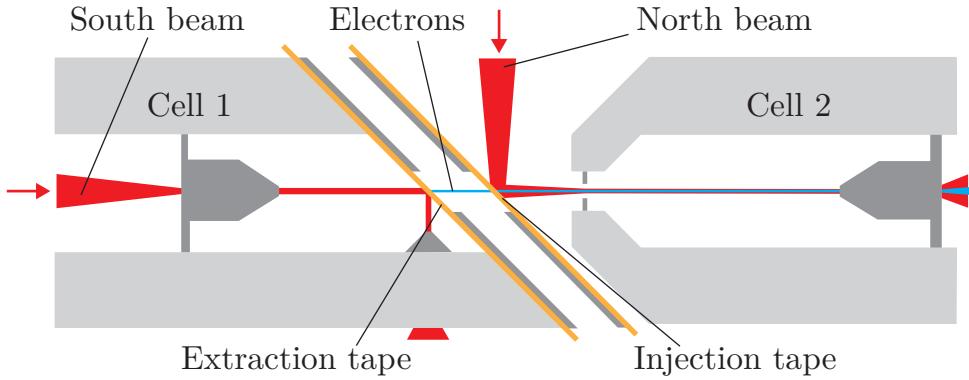


Figure 6.2: Diagram of the target. Electrons injected and accelerated in cell 1 pass through the tapes into cell 2. The red arrows indicate the initial laser propagation direction.



Figure 6.3: Shadowgrams of the target during a shot. The lasers propagate from left to right across the images.

cell, which allowed the two cells to be positioned as close to each other as possible. This tape had a thickness of $25\text{ }\mu\text{m}$, while the injection tape had a thickness of $125\text{ }\mu\text{m}$, and both were laid at 45° to the optical axis. Since the reflection of a single laser shot caused the near or complete puncturing of the tapes, during normal operation, the tapes were translated between each shot. A composite shadowgraphy image of the two cells during a shot is shown in Fig. 6.3.

When the laser pulses enter the vacuum chamber, they are both polarised in the y direction, as defined in Fig. 6.1, and they initially propagate in the z direction, entering the chamber vertically. After reflection from the first optic, the South beam's polarisation remains in the plane of the page, while the North beam's polarisation points in the z direction, and this orientation is maintained until the laser pulses reach the target. Consequently, the South beam is p -polarised with respect to the tape, while the North beam is s -polarised. This is advantageous for maintaining energy in the North beam, as energy is not lost via Brunel absorption² that is enhanced in p -polarisation and degrades the reflection quality³.

² [91] C. Thaury *et al.* (2007) *Nature Physics*

[232] F. Brunel (1987) *Physical Review Letters*

³ [233] B. H. Shaw *et al.* (2016) *Physics of Plasmas*

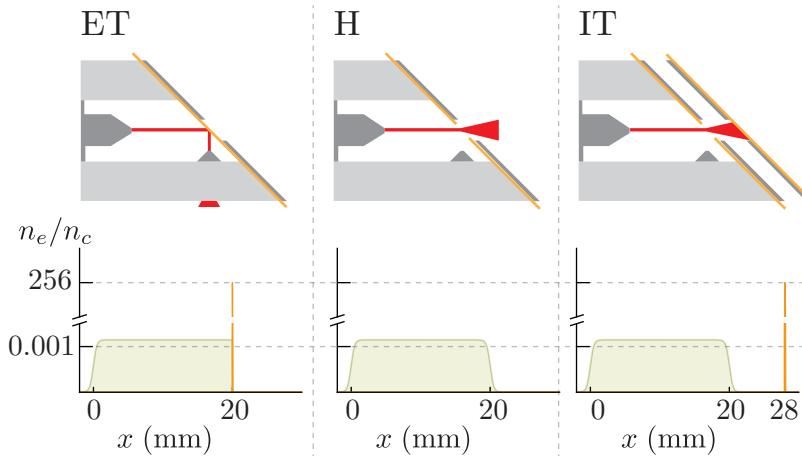


Figure 6.4: The three cell 1 configurations used in this experiment, extraction tape (ET), hole (H), and injection tape (IT), are displayed across the top row. The corresponding longitudinal density profiles are shown below. The underdense plasma from the gas is shaded green, while the overdense plasma from the plasma mirror is orange. The axis is broken to accommodate the range of densities.

To investigate the effect of the tape on the electron beam, measurements of the electron beam profile and spectrum were made under various target configurations. The three main configurations are depicted in Fig. 6.4, where, for all data reported here, cell 2 was removed, and the North beam was not used⁴. For the extraction tape configuration (ET), the injection tape was removed, and 50 shots were taken while the extraction tape was translated between shots. To create the hole configuration (H), 10 shots were taken without translating the extraction tape, allowing a hole of stable size to form in the tape. A total of 40 shots were then taken without moving the extraction tape. Finally, for the injection tape configuration (IT), the injection tape is placed in its normal position, 8 mm after the extraction tape, and 19 shots were taken while translating this tape, whilst the extraction tape remained static.

The leakage of the South beam through a dielectric mirror was used to probe the interaction. The probe beam diameter was reduced to ~ 75 mm with a Galilean telescope and was then sent through a delay line ahead of reaching the target. The beam was split using a 50:50 beamsplitter before propagating through the cells orthogonally to the direction of the main beam, as shown in Fig. 6.3. This allowed transverse interferometry and shadowgraphy to be performed independently for each cell. This, alongside pressure sensors in the cells, was used to monitor the plasma density. A 300 mm motorised translation stage was used to adjust the delay between the probe beam and drive beam, ensuring that the plasma was imaged close to its time of formation. For all measurements reported in this chapter, the plasma density was $2.0 \pm 0.5 \times 10^{24} \text{ m}^{-3}$.

Before taking data, the focal spot of each beam was optimised by correcting for wavefront aberrations using the adaptive optic (AO in Fig. 6.1). The focal spot quality was

⁴ The properties of the North beam were described because they pertain to the discussion later in the chapter.

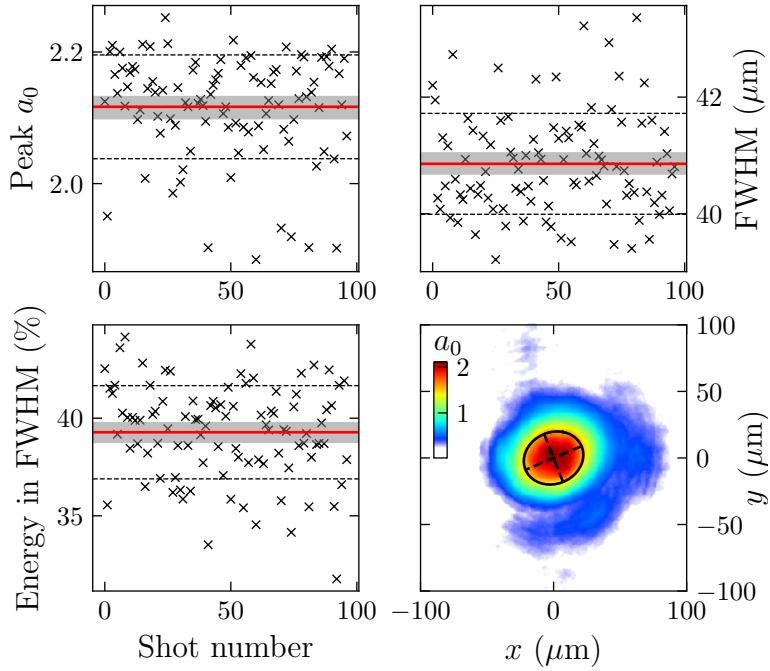


Figure 6.5: Focal spot metrics of all the focal spot images taken on that day, and the median focal spot image a_0 . The red horizontal lines indicate the median value, the black dashed lines show the shot-to-shot standard deviation, and the grey shaded region is the estimated error on the mean, calculated via bootstrapping.

measured by taking 100 images of the focal spot using a spatially calibrated microscope. The relevant measurements for this chapter are shown in Fig. 6.5. Accounting for the 70% compressor efficiency, 10.9 ± 0.2 J of energy was delivered to the target by the $f/40$ OAP, with a central wavelength of 808 nm. When optimally compressed, the pulse length was 45 ± 2 fs (FWHM), however, to maximise the energy and charge of the electron beam a Fastlite dazzler⁵ was used to add modest second, third, and fourth-order spectral phase terms, stretching the pulse slightly. The median FWHM spot size was 41.2 ± 0.9 μm , with 39 ± 2 % of the energy contained within the FWHM. Assuming an optimally compressed beam with 10.9 J of energy, the laser intensity was $I = 9.58 \pm 0.70 \times 10^{22} \text{ W m}^{-2}$, corresponding to a vacuum $a_0 = 2.12 \pm 0.08$.

6.1.2 Electron profile

The electron beam profile monitor was composed of a 75 mm circular Lanex screen placed 2.15 m from the end of cell 1, imaged off-axis by a 12-bit CCD. Although placing the Lanex screen before the magnet would allow for simultaneous measurement of the beam profile and the spectrum, placing it further from the source magnifies the beam, improving the measurement resolution. The advantage of the electron profile diagnostic over the spectrometer is that it can measure the divergence and pointing of the beam in both transverse dimensions. This was particularly desirable for our experiment since

⁵ [223] *Fastlite Dazzler product information*, Fastlite

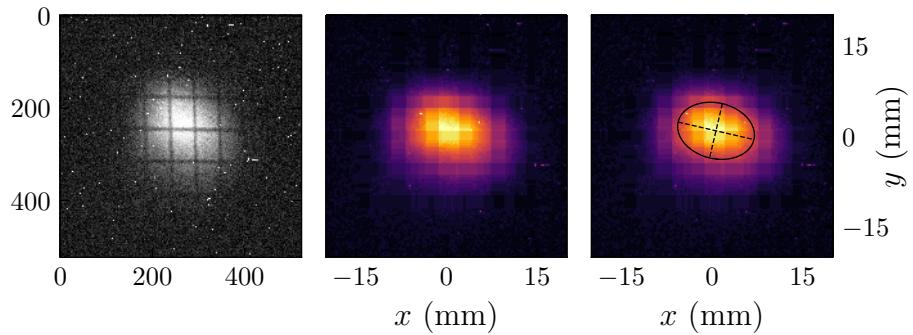


Figure 6.6: Electron profile processing and measurement example. The cropped raw image is displayed in the left panel (axes are in pixels), while the transformed and median-filtered profile is shown in the central panel. Note the wire grid shadow seen in the raw image has been filled in. The black ellipse in the right panel is fitted to the FWHM intensity contour.

both the asymmetrical target and the linearly polarised laser pulse had the potential to produce asymmetrical electron beam behaviour in the transverse plane.

At the sub-MeV energy level, the response of the scintillating screen becomes non-linear⁶, and as such, there is the possibility that low energy particles bias the measurements. For the ET and IT configurations, 1 MeV electrons would have a scattering angle of over 160 mrad as they propagate through 25 μm of Kapton⁷, which is much larger than the 35 mrad wide beam profile screen. However, in the H configuration, there is nothing in the beam path between the cell and the diagnostic, so it is possible that a beam of ultra-low transverse momentum electrons could reach the Lanex. If this were the case, we would expect the integrated counts from the H configuration to be significantly higher than for the other configurations. Since this is not the case, we assume that the low-energy electrons are not biasing the measurements in any of the configurations. This assumption is verified by comparing beam profile divergence measurements with those from the electron spectrometer, which are in agreement.

A beam profile image at different stages of processing is shown in Fig. 6.6. Using the known dimensions of the screen, a projective transform was used to correct for distortions resulting from the viewing angle. Two background subtractions were performed; a pixel-wise one, using a reference image taken without the laser pulse firing, and a global one, using pixels at the edges of the Lanex. This second one was required to remove the effects of shot-to-shot variations in the background level, caused by imperfect timing of the camera shutter. Hard hits were removed using a median filter with a kernel width of 3 pixels. As can be observed from the left panel in Fig. 6.6, the shadow of a grid is imprinted on the beam, which was from a tungsten wire grid that was deliberately placed in the beam path. In principle, the edge spread of the shadow could be used to

⁶ [184] Y. Glinec *et al.* (2006) *Review of Scientific Instruments*

⁷ [234] M. B. Reid (1991) *Journal of Applied Physics*

[235] V. L. Highland (1975) *Nuclear Instruments and Methods*

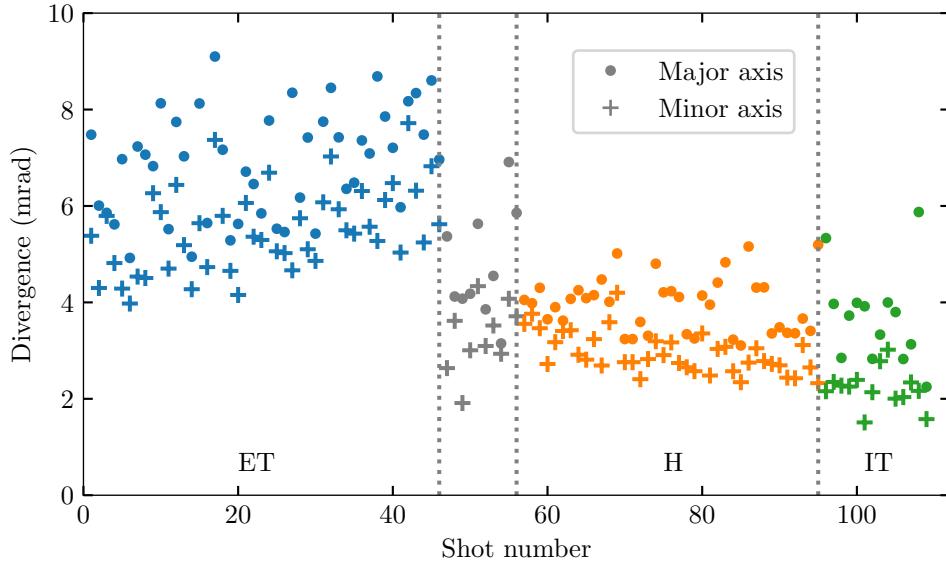


Figure 6.7: The major and minor axis divergence for every measured shot, arranged chronologically. The grey markers indicate those shots taken during the hole formation process.

determine the spatially resolved divergence, allowing for a measurement of the emittance to be obtained⁸. In practice, the electron energies, the energy spread, and the resolution of the profile monitor made this measurement impossible. More details of why this measurement failed, and a method for performing it successfully, are given in chapter 5.

The beam size was measured by fitting an ellipse to the FWHM contour of the beam's profile. This procedure was complicated by the shadow of the wire grid, as the beam was effectively sliced into numerous small beamlets, each with its own FWHM contour. Simply applying a Gaussian blur to merge the beamlets is a crude method of addressing this issue, because the standard deviation of the kernel used for the blurring spuriously affects the measurement. Instead, a synthetic profile was generated from each beam profile image, where the regions masked by the wire array are detected and filled using information in the neighbouring parts of the image. This synthetic image was then used to calculate the FWHM, requiring no Gaussian filtering. An example of this synthetic image is shown in the central panel of Fig. 6.6, while the right panel shows the fitting of an ellipse to the synthetic image.

6.1.3 Beam profile measurements

The divergence measurements from every measured shot are shown chronologically in Fig. 6.7. The total number of shots for which a measurable electron beam profile was obtained was 46, 39, and 14, for the ET, H, and IT configurations respectively. Moving from the ET to the H configuration results in a substantial reduction in divergence, while moving from the H to the IT configuration results in a further slight drop in divergence.

⁸ [236] D. Marx *et al.* (2018) *Physical Review Accelerators and Beams*

The charge-weighted mean and standard deviation for each configuration are given in Tab. 6.1. Taking the average of the major and minor axes, the divergence of the ET data set was 80% higher than the hole data set, and the shot-to-shot fluctuations were three times greater (12% vs. 4%). The beams were slightly elliptical, with the major axis predominantly aligned with the polarisation axis of the laser, and ellipticity increased slightly with shot number.

Table 6.1: Experimental divergence and pointing measurements. For each measurement, the mean, μ , and standard deviation, σ , are given, and the errors are calculated using bootstrapping. The key results are highlighted in yellow.

Divergence (mrad)			Pointing (mrad)		
	Major	Minor	Mean		
	μ	μ	μ	x	
ET	μ	6.85 ± 0.16	5.41 ± 0.13	6.13 ± 0.13	1.39 ± 0.27
	σ	1.16 ± 0.08	0.72 ± 0.13	0.76 ± 0.08	1.85 ± 0.36
H	μ	3.86 ± 0.09	2.90 ± 0.07	3.38 ± 0.07	0.19 ± 0.21
	σ	0.28 ± 0.05	0.16 ± 0.05	0.16 ± 0.04	1.31 ± 0.18
IT	μ	3.38 ± 0.25	2.14 ± 0.10	2.76 ± 0.14	0.51 ± 0.34
	σ	0.67 ± 0.2	0.11 ± 0.08	0.23 ± 0.09	1.25 ± 0.29

The pointing angle of each electron beam measurement is shown in Fig. 6.8, where pointing is defined relative to the optical axis. An ellipse has been fit to the 1σ confidence level of each data set and centred on the average pointing of each configuration, reported in Tab. 6.1. This ellipse is slightly larger for the ET configuration than for the H and IT configurations, and there is also a slight shift in the average pointing. However, the difference in ellipse size is small, and the pointing shift is within the 1σ confidence contour. We therefore conclude that the differences in electron beam pointing are not statistically significant. Repeating this experiment and either taking more data or using a laser with less jitter would help determine if the small changes we observe are due to the target configuration.

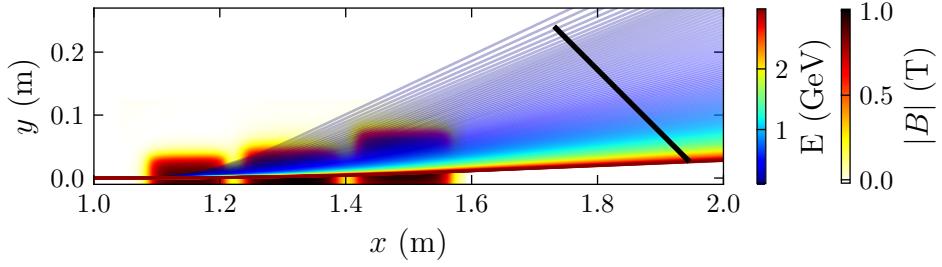


Figure 6.9: Particle tracking for calibration of the electron spectrometer used in this experiment. The black line shows the position of the Lanex screen.

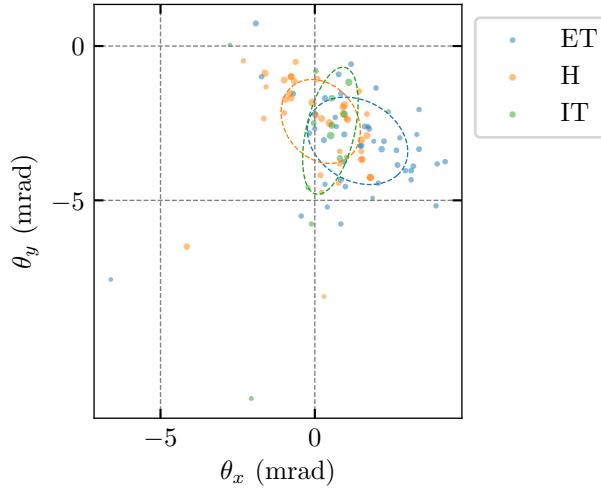


Figure 6.8: Beam profile centroids relative to the laser axis. The marker size scales with the charge of the beam, and the ellipses give the 1σ confidence contour of the measurements, weighted by the charge.

6.1.4 Electron spectrometer

The setup for the electron spectrometer is shown in Fig. 6.9. Three 1 T dipole magnets extending over 500 mm were used to disperse the beam, and two Lanex scintillating screens were used to extend the energy range up to 5 GeV. The first Lanex screen covered the 0.1–2.6 GeV energy range, with an error of ± 0.135 GeV for a 1 mrad \hat{y} deviation at the highest energy. The second Lanex screen was not used in the measurements presented here, because, firstly, the electron energies did not exceed ~ 2.2 GeV, and secondly, the divergence of the beam on the second screen at energies below 2.2 GeV would be modified by multiple scattering as the electrons propagate through the first screen. Correcting for this would needlessly complicate the analysis and introduce errors. To process the images from the spectrometer camera, a projective transform was performed to correct for viewing angle distortion, and a median filter was applied to remove hard hits. Since the background varied significantly across the image, particularly for shots where there was a hole in the extraction tape, a gradient background subtraction was used, where a strip of pixels along each edge of the Lanex was used to calculate the local background.

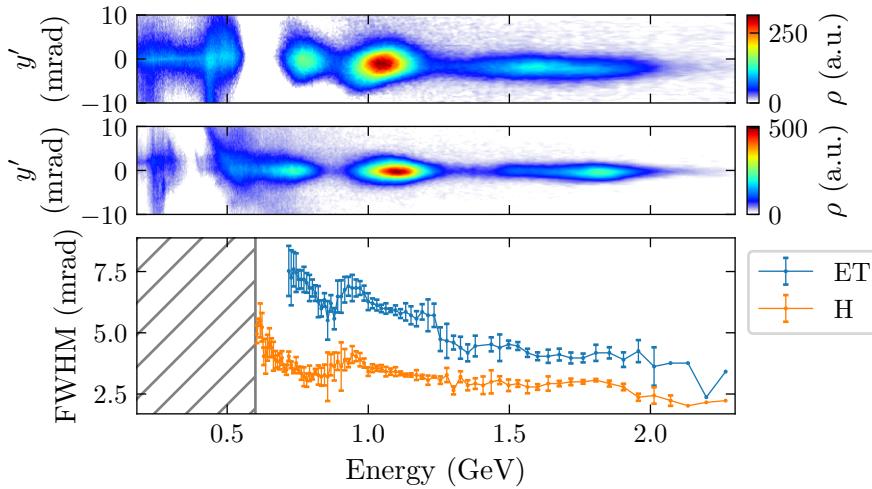


Figure 6.10: The top two panels show example electron spectrometer measurements, where the top panel is from the ET configuration and the lower panel is from the H configuration. The lower panel shows the spectrally resolved divergence measurements for each beam.

Example electron spectrometer images are shown in the upper two panels of Fig. 6.10, and the lower panel shows the energetically resolved FWHM divergence of the beams. The data points are calculated by taking the mean transverse beam profile over a sample width of 10 pixels, which has the result of reducing the sampling rate at higher energies. An estimate of the error for this measurement is calculated by first measuring the divergence for a sample width of 1 pixel at all possible locations, then calculating the rolling standard deviation of these values with a window length of 10 pixels. The error bars in the figure are sampled from this rolling measurement. The region below 0.6 GeV was polluted by stray light, which varied transversely across the scintillator, and hence could not be completely removed via background subtraction. Consequently, these beam divergence measurements will only use energies above 0.6 GeV.

6.1.5 Electron spectrometer measurements

The averaged electron spectrometer data is shown in Fig. 6.11. In the ET configuration, 14 shots were recorded, while for the H configuration, 13 shots were recorded after 10 shots were taken beforehand to form the hole in the tape. The upper panel shows the mean spectrum for each configuration and the associated uncertainty. There was considerable shot-to-shot variation in the energy spectrum, but the consistent features were a peak at approximately 1 GeV, and a tail extending past 2 GeV, similar to the example images in Fig. 6.10. The total charge above 0.6 GeV for the H configuration was $\sim 25\%$ lower than for the ET shots. Still, the similarity of the averaged spectra is reasonable, making comparisons between shots taken in the different configurations appropriate. The spectrally resolved FWHM divergence is measured using the procedure depicted in Fig. 6.10, producing an array of divergence measurements at each energy with an associated error. The lines and shaded regions in the lower panel of Fig. 6.11 represent the mean and standard deviation of these arrays, where the mean is weighted by charge.

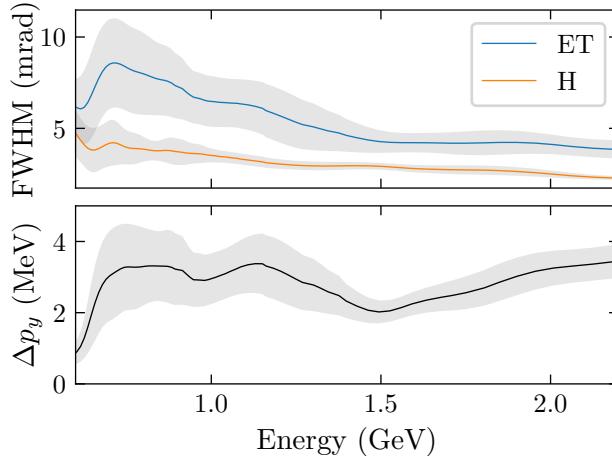


Figure 6.11: The top panel shows the mean FWHM divergence measurements for the ET and H configurations. The bottom panel shows the inferred change in divergence. The grey shaded regions indicate the $\pm 1\sigma$ error.

The divergence reduction due to the hole's presence is observed to persist at energies above 0.6 GeV. The average divergence across the spectrum drops from 6.20 ± 0.18 mrad for the ET data to 3.90 ± 0.06 mrad for the H data, with errors calculated using bootstrapping, in close agreement with the beam profile measurements reported in Tab. 6.1. The lower panel of Fig. 6.11 shows the average difference in transverse momentum between the two data sets as a function of energy, which is consistently in the 2 - 3 MeV range.

To summarise the main experimental findings, the electron profile diagnostic observed a 2.8 ± 0.5 mrad decrease in the total beam divergence when changing from the ET configuration to the H configuration. The electron spectrometer corroborated this effect, which observed a 2.3 ± 0.6 mrad average difference between the configurations. Critically, this latter observation means that the effects of low-energy electrons are not biasing the beam profile data; the presence of the tape is affecting the high-energy electrons that are of most interest for staging applications.

6.2 Laser-tape simulations

Particle-in-cell simulations have been performed to investigate the mechanism causing the divergence change observed between the ET and H configurations. The range of relevant densities extends from $10^{-3}n_{crit}$ in the wakefield accelerator to over $250n_{crit}$ in the plasma mirror, and the spatial scales of interest range from the 2 cm length of the under-dense plasma down to the 10 nm skin depth of the laser in the overdense plasma on the surface of the tape. To accurately resolve this range of length scales and to do so in a computationally efficient way, separate simulations using different simulation codes were performed. This section presents the results of the overdense laser-plasma simulations.

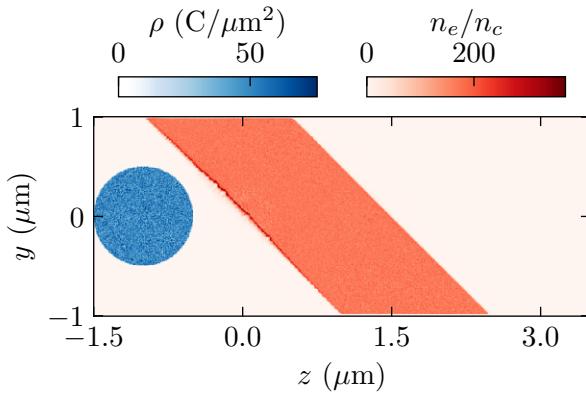


Figure 6.12: Snapshot of the laser-tape simulation charge density in the 45° geometry, after interaction with the laser pulse and before the electron beam enters the tape. The simulation time at this point is 0.17 ps. The laser pulse and electron initially travel in the positive z direction.

6.2.1 Simulation setup

The code Smilei⁹ was used to model the interaction between the laser, tape, and electron beam. Separate simulations were performed for the tape angled at 0° and 45° to the direction of laser propagation, and these simulations were performed in a full 3D Cartesian geometry. The 25 μm thick Kapton tape used in the experiment has the chemical formula $\text{C}_{22}\text{H}_{10}\text{N}_2\text{O}_5$, and a density of $1.42 \times 10^3 \text{ kg m}^{-3}$. In these simulations, the tape is modelled as a carbon strip with a number density of $7.3 \times 10^{28} \text{ m}^{-3}$, such that the plasma density reaches $256n_{\text{crit}}$ when fully ionised, equivalent to the fully ionised plasma density of Kapton. This simplification of the target's chemical composition will have the effect of placing a higher proportion of the electrons in the inner shell, making the target slightly more difficult to ionise. However, this is offset by the lower nuclear charge of carbon, lowering their ionisation energy. As such, the simplification to an all-carbon target is not expected to have a significant effect on the physics of the interaction. Although pre-ionising the target would reduce the computational load, this has been shown to reduce the magnitude of the fields of interest¹⁰, and also introduces uncertainty as to the appropriate initialisation temperature of the plasma electrons, which is important for the Weibel instability. As such, the target is initialised as a block of initially neutral carbon atoms. Ionisation is determined by the ADK rates¹¹, which is handled by Smilei's ionisation module. The tape thickness was set to 1.5 μm , allowing it to fit inside the 5 μm long simulation window, with room on each side for fields outside the tape to be observed. The transverse extent of the simulation was 2 μm in x and y , where, for the angled tape simulations, the tape was angled when viewed in the $y - z$ plane. A snapshot of the full simulation window is shown in Fig. 6.12. Similar simulations have been performed by Raj *et al.*¹², but with differences in the plasma density, laser properties, and geometry.

⁹ [213] J. Derouillat *et al.* (2018) *Computer Physics Communications*

¹⁰ [237] Y. Lang *et al.* (2018) *Plasma Physics and Controlled Fusion*

¹¹ [132] M. V. Ammosov *et al.* (1986) *Sov. Phys. JETP*

¹² [92] G. Raj *et al.* (2020) *Physical Review Research*

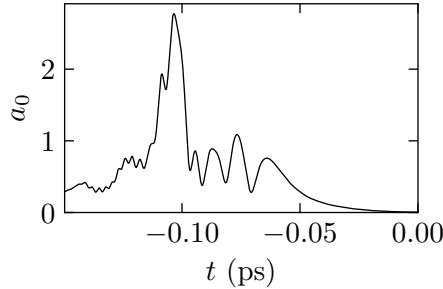


Figure 6.13: Laser temporal profile for the laser-tape simulations, taken from the end of the simulations reported in Sec. 6.3.

The smallest length scale that needs to be resolved is the plasma skin depth, $\sim c/\omega_p$, which, for the fully ionised Kapton, is $8\text{ nm} \approx \lambda_L/100$. Consequently, the resolution of the simulations was $\Delta = \lambda/103$ in all directions. Due to the currents generated, the plasma density can reach considerably higher values than the fully ionised plasma density, up to and above $400n_c$ at the edges of the tape, so some regions of the simulation may be improperly resolved. This can be seen on the front surface of the tape in Fig. 6.12. The magnetic fields of primary interest here have a spatial scale of approximately $\lambda/10$ and should therefore be resolved sufficiently. For the carbon, 25 macroparticles per cell were used, while 4 particles per cell were used for the electrons, and open boundary conditions were used for the simulation edges. Open boundary conditions were required due to the asymmetry in the 45° geometry, but would not be required if the laser pulse were rotated instead of the tape.

The laser pulse undergoes depletion¹³ and amplification¹⁴ during its interaction with the 20 mm of underdense plasma that precedes the tape. Since the interaction is highly nonlinear, it is difficult to predict the energy that remains in the pulse. Further, the pulse amplification process will produce a highly distorted temporal profile that cannot be calculated analytically. To account for these effects, the longitudinal laser pulse envelope for the tape simulations was taken from the wakefield simulations discussed later in this chapter. This profile is plotted in Fig. 6.13, containing 2.1 J of energy with a peak $a_0 = 2.76$, higher than the vacuum focus a_0 of 2.07. This pulse was focused onto the tape's front surface with a Gaussian spot size of $20\text{ }\mu\text{m}$.

An electron beam was also simulated, propagating across the tape after the laser interaction. The beam was modelled as a sphere with a radius of $0.5\text{ }\mu\text{m}$, 50 pC of charge, and an energy of 2 GeV. The electron beam enters the simulation at $t = 0.17\text{ ps}$, approximately $20\text{ }\mu\text{m}$ after the peak of the laser pulse encounters the tape, mimicking the separation between the laser pulse and the electron bunch observed in wakefield simulations, and in experiments¹⁵.

¹³ [148] B. A. Shadwick *et al.* (2009) *Physics of Plasmas*

¹⁴ [146] M. Streeter *et al.* (2018) *Physical Review Letters*

¹⁵ [64] M. Heigoldt *et al.* (2015) *Physical Review Special Topics - Accelerators and Beams*

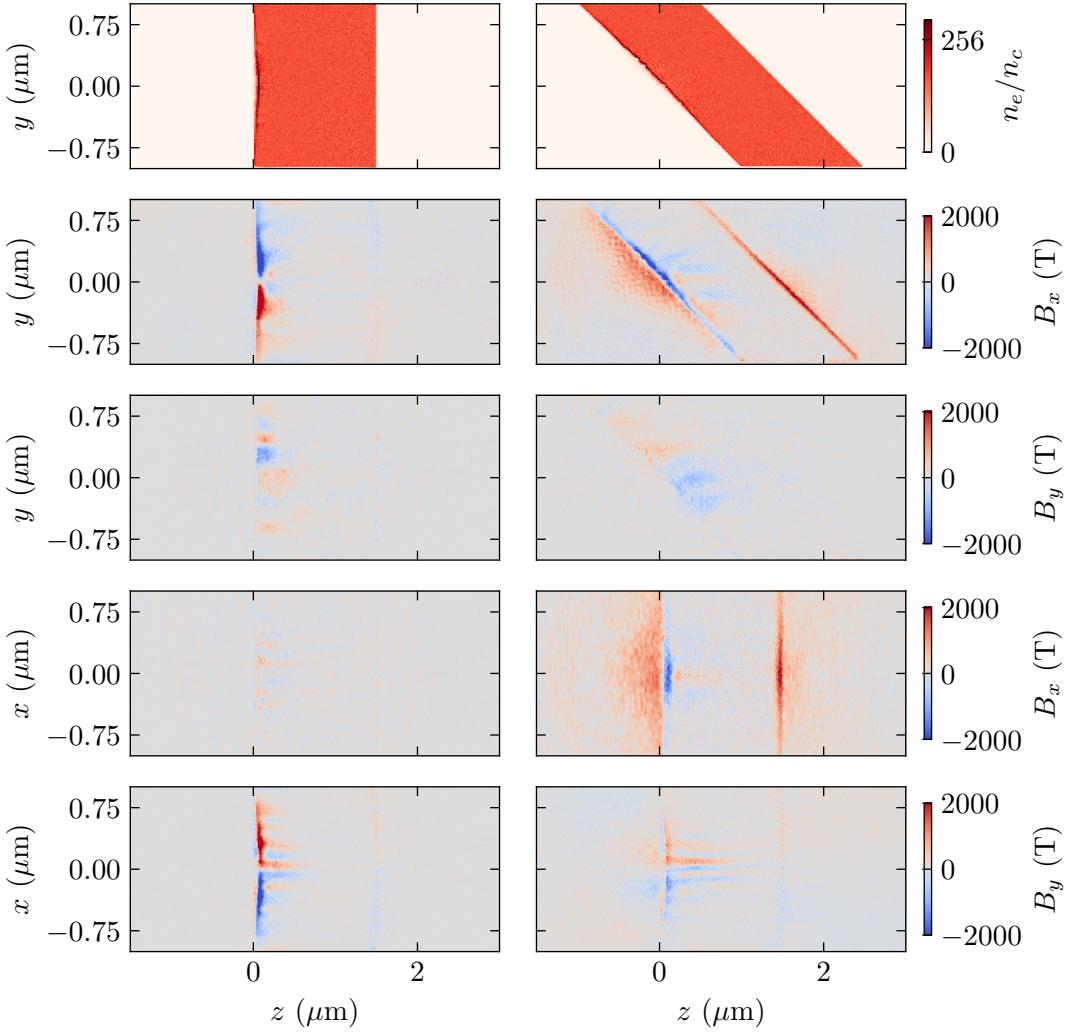


Figure 6.14: Plasma electron density and magnetic fields at the front surface of the tape, at the moment the peak of the laser envelope encounters the tape. The left column of plots is for the 0° geometry, and the right column of plots is for the 45° geometry.

6.2.2 Field structure in the tape

A snapshot of the plasma density and magnetic fields at $t = 0.17$ ps is shown in Fig. 6.14. The plasma density from the carbon electrons is shown in the $y - z$ plane, while the B_x and B_y fields are shown in both the $y - z$ and the $x - z$ planes, for $x = 0$ and $y = 0$ respectively. The B_z fields have not been plotted since these do not deflect electrons travelling in the z direction. It should also be noted that the electric fields have been neglected; since the magnetic fields are on the order of 500 T, $cB \simeq 10^{12} \text{ V m}^{-1}$, and no electric fields approaching this magnitude were observed ($|E|_{\text{max}} \ll |cB|_{\text{max}}$). A thin region on the front of the tape is observed to be fully ionised, with the bulk of the tape having a median ionisation state of +3.9, which corresponds to the $2 \times 2p$ and $2 \times 2s$ outer shell electrons of carbon.

The magnetic fields on the inside of the front surface of the tape have the characteristic structure of the current filamentation instability¹⁶. These fields extend up to $\sim 1\text{ }\mu\text{m}$ into the tape, and at this time have a transverse wavelength of $\sim 0.4\text{ }\mu\text{m}$. At earlier times, the transverse wavelength of these fields is shorter, in agreement with the Weibel instability growth rate, which predicts a higher growth rate for shorter wavelength perturbations. Since this transverse wavelength is sub-micron, these fields will impart spatially dependent deflections of the electron beam, increasing divergence. The magnetic field strength magnitude of the Weibel instability at saturation given by Fiuza *et al.*¹⁷ is

$$B_{\text{sat}} = \sqrt{2a_0\alpha n_c \mu_0 m_e c^2}, \quad (6.1)$$

where αn_c is the density of hot electrons driving the instability, with the value of α determining the number of high-energy electrons driving the instability. This scaling was derived by matching the energy density in the hot electron flow to the energy density in the magnetic fields. In the simulations presented here, field strengths of approximately 1 kT are reached, with a peak a_0 of 2.73, indicating $\alpha \approx 10^{-3}$ when the electron beam enters the tape. In contrast, at the moment the peak of the laser is incident on the tape, the magnetic fields are an order of magnitude stronger, indicating $\alpha \approx 0.1$, which is in reasonable agreement with the results of Lang *et al.*¹⁸.

The currents generated in the tape by the laser are shown in Fig. 6.15. This snapshot is taken at the moment the peak of the laser pulse is incident on the tape, as this time most clearly shows the various currents that are produced. The laser's electric field is polarised in the y direction and can be seen to produce currents in this direction on the front surface of the tape. In the 0° geometry, this is uniform across the surface, while for the 45° geometry, the J_y field on the front surface of the tape varies as the phase of the laser changes across the surface. Large amplitude plane waves in the current are observed inside the bulk plasma, which have a wavelength of approximately 60 nm, equal to the plasma wavelength of the partially ionised carbon. In the 0° geometry, these currents are entirely in the z direction, while in the 45° geometry they are in both the z and y direction, indicating that these plane waves are emanating from the front surface of the tape. These are simply Langmuir waves, and carry pulses of current across the tape.

Currents extending out of the front surface of the tape are observed for the 45° geometry. This is due to electrons being pulled out of the plasma by the electric field of the laser, before being accelerated into the tape, in a process known as vacuum heating, or Brunel heating¹⁹. This enhances laser energy absorption²⁰, and should result in extra heating of the target compared to the 0° geometry since this heating mechanism cannot

¹⁶ [238] E. S. Weibel (1959) *Physical Review Letters*
[239] B. D. Fried (1959) *The Physics of Fluids*

[92] G. Raj *et al.* (2020) *Physical Review Research*
[240] F. Fiuza *et al.* (2012) *Physical Review Letters*

[237] Y. Lang *et al.* (2018) *Plasma Physics and Controlled Fusion*

[232] F. Brunel (1987) *Physical Review Letters*

[241] P. Gibbon *Short pulse laser interactions with matter* (2005)
[242] N. Dover *PhD. Thesis* (2013)

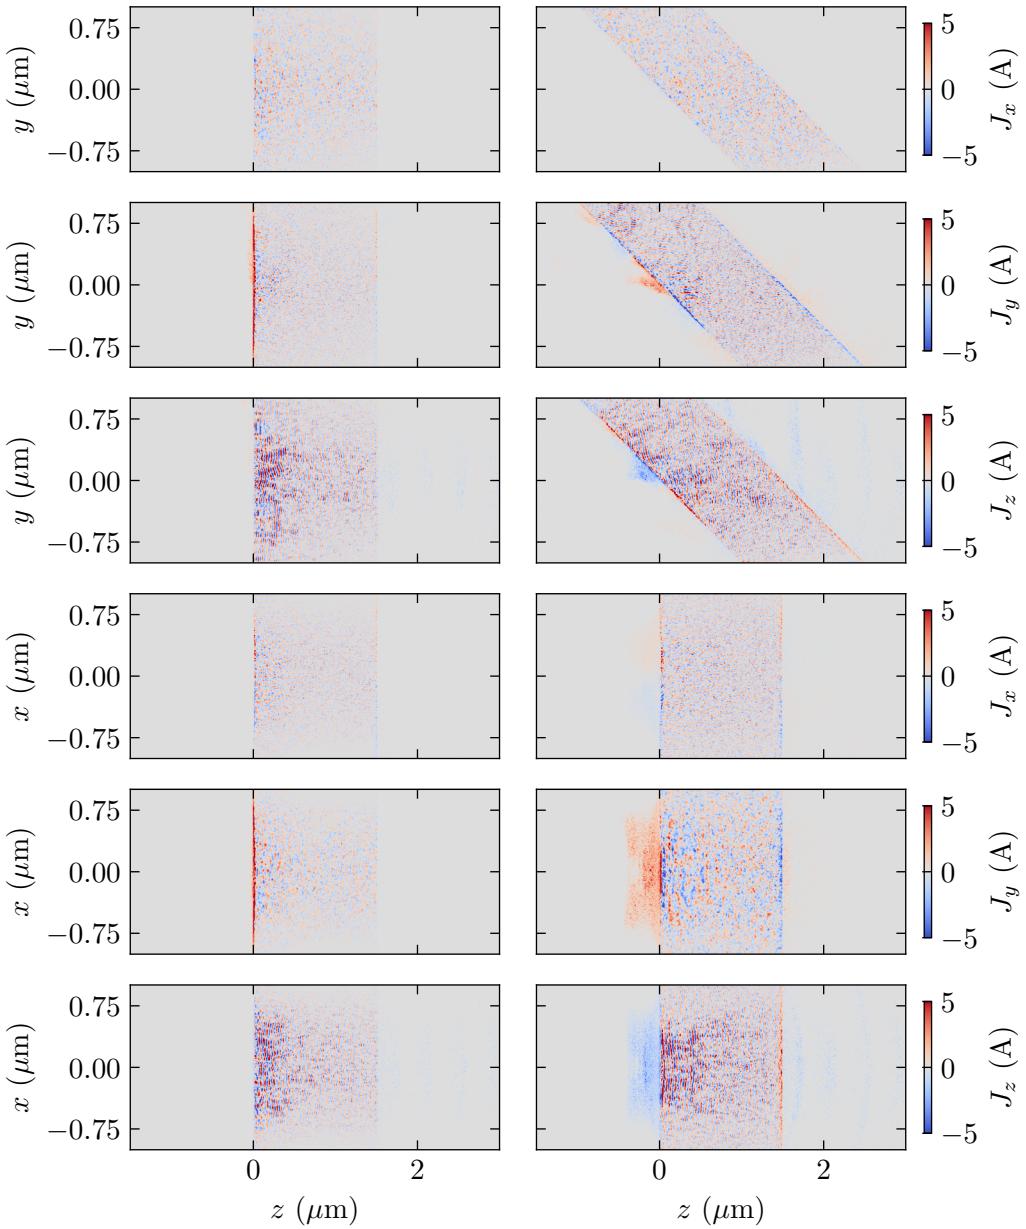


Figure 6.15: Currents generated in the tape by the laser. The left column plots are from the 0° geometry, while the right column plots are from the 45° geometry. If the current coordinate is orthogonal to both axes, positive current is directed out of the page.

occur at normal incidence. However, since the laser is relativistic, $J \times B$ heating becomes significant, and this is most efficient at normal incidence, so this may compensate for the lack of vacuum heating. The J_z currents emanating from the rear surface of the tapes at a frequency of 2ω are characteristic of this heating mechanism²¹.

Once the laser-tape interaction has finished, it is the residual electron currents that drive the long-term magnetic fields with which the electron beam interacts. There are two sources of these currents; a temperature anisotropy, normal to the tape surface, and a drift of electrons out of the simulation, parallel to the tape surface. The former causes the Weibel fields that are observed extending from the front surface of the tape. The latter produces the magnetic fields that can be observed on the front and back surfaces inside the tape, and the fields outside the tape in the 45° geometry. The drift of electrons out of the simulations occurs in all directions symmetrically for the 0° geometry, while it occurs predominantly in the negative y , positive z direction for the 45° geometry. This net drift of electrons is why there are significant magnetic fields outside the tape in the 45° simulations but not in the 0° simulations. Although it is produced by the removal of electrons at a boundary, these currents and associated fields could be physical, since, in an infinite tape irradiated in a small region by a laser at an angle, the net flow of electrons will still be influenced by the initial laser pulse direction.

6.2.3 Electron beam deflection

An electron beam passing through the tape will experience the fields at different points in space, due to the beam's transverse size, and time, due to the time taken for the beam to traverse the tape. Individual electrons moving with $v \approx c$ will be deflected by the Lorentz force

$$\frac{dp_y}{dt} = -ecB_x, \quad (6.2)$$

and because the transverse scale of the fields is smaller than that of the beam, the deflection will have a spatial dependence. If we assume that the fields are frozen in time, the increase in transverse momentum spread in y as the electron beam crosses the tape is given by

$$\begin{aligned} \Delta p_y^2 &= e^2 c^2 \langle B_{x,int}^2 \rangle, \\ &= \frac{e^2 c^2}{Q} \iint q(x, y) dx dy \int (B_x(x, y, z) - \bar{B}_x(z))^2 dz, \end{aligned}$$

where

$$\bar{B}_x(z) = \frac{1}{Q} \iint B_x(x, y, z) q(x, y) dx dy, \quad (6.3)$$

is the charge-weighted mean magnetic field, and $q(x, y)$ is the transverse electron beam profile, such that $Q = \int_{beam} q(A) dA$.

²¹ [243] S. C. Wilks *et al.* (1992) *Physical Review Letters*

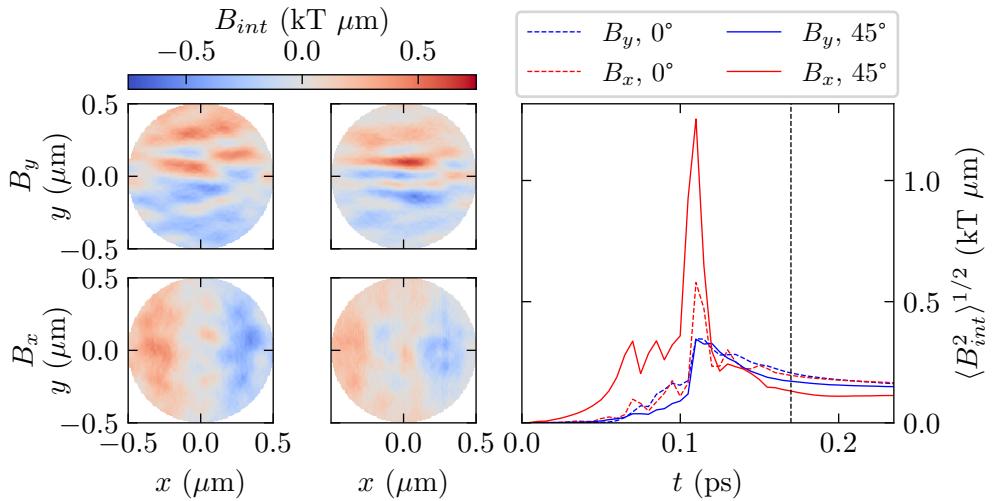


Figure 6.16: The left four panels show the z integrated magnetic field over the electron beam radius of $0.5 \mu\text{m}$, with the left column showing the output from the 0° simulation, and the right column showing the output from the 45° simulation. The right panel shows the evolution of the variation in B_{int} during the simulation. The dashed vertical line indicates the time at which the electron beam enters the tape, which is the moment at which the four plots on the left are extracted.

Fig. 6.16 shows the integrated B fields at the moment the electron beam enters the tape and the temporal evolution of their variation. The fields in y have a higher spatial frequency than those in x , and the fields for the 45° geometry are slightly weaker, with a larger difference for the B_x fields. This larger difference in x is due to the integration of the tilted fields along z , causing these fields to ‘wash out’ slightly. Both of these differences can be seen numerically in the plot of the temporal evolution. The fields in the 45° geometry are larger at earlier times due to the earlier interaction with the tape, and the effect of $J \times B$ heating, which is most prominent at 0.11 ps . However, once the laser has passed, the fields have a similar magnitude, stabilising at approximately $\sqrt{\langle B_{int}^2 \rangle} = 0.1 \text{ kT } \mu\text{m}$.

Finally, the electron beam that propagates across the tape can be used to estimate the effects of these fields. The evolution of the transverse momentum FWHM of the beam is plotted in Fig. 6.17, which shows that the deflection of the beam occurs differently in each plane and each geometry. In the 0° simulations, the majority of the divergence increase occurs when the beam enters the tape at $\bar{z} = 1.5 \text{ mm}$, and the behaviour in x and y is similar. In contrast, in the 45° geometry, the divergence increases more gradually, with markedly different behaviour in the x and y directions. The p_x FWHM increases as it enters the tape²², and then continues to increase while inside the tape, while the p_y FWHM increases early, at 0.5 mm , remains constant while entering the tape, and then increases gradually while transiting the tape. The p_y increase at 0.5 mm is due to the B_x fields that extend beyond the tape, while the gradual increases inside the tape are due to the fields extending further into the 45° tape than the 0° tape. Considering only

²² This region is wider in the 45° geometry because the top of the beam enters the tape before the bottom.

the effects of the fields inside the tape, responsible for divergence changes from 1 mm to 3 mm, we see that the reduced integrated B_x fields observed in Fig. 6.16 result in a reduced divergence increase in the y direction.

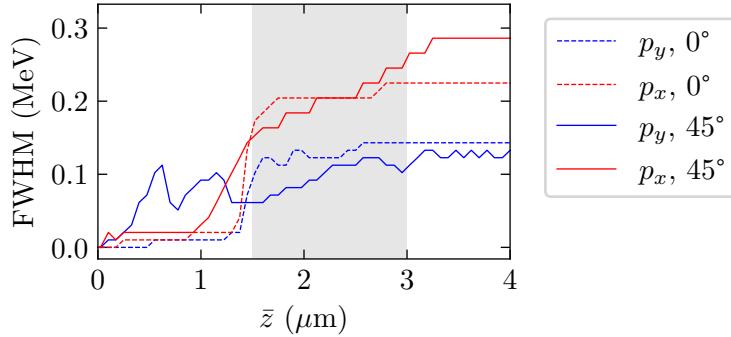


Figure 6.17: FWHM electron beam transverse momentum change as it traverses the simulation. The grey shaded region indicates the location of the tape.

In summary, these simulations estimate that the fields in and around the tape can increase the electron beam transverse momentum by a FWHM value of approximately 0.2 MeV, or a 0.1 mrad increase in the divergence of the 2 GeV electrons in the beam. Since this is too small to explain the > 2 mrad divergence change, other factors must be considered.

6.3 LWFA simulations

Simulations of the wakefield accelerator were performed in FBPIC²³. The main purpose of these simulations was to quantify the effects of a density ramp at the end of the cell on the beam divergence. To do this, simulations were run using density ramps of different scale lengths, where the ramp was modelled as an exponential. The analytic form of the longitudinal density profile at the end of the cell was $n(s) = n_0/(e^{(s/\lambda_R)} + 1)$, where λ_R is the scale length of the ramp. This function was found to fit fluid simulations of gas leaking through circular apertures²⁴, where the radius of the aperture $R \sim 2\lambda_R$. Thus, a 1 mm diameter hole would be expected to have a ramp with a 250 μm scale length. This modelling does not account for the effect of a ragged or moving aperture, which may have been the case for a hole formed in thin tape by a distorted laser pulse. The ramp at the front of the cell had $\lambda_R = 250 \mu\text{m}$, and the plateau density was $n_0 = 2 \times 10^{24} \text{ m}^{-3}$.

These simulations use the method described in Sec. 3.5.4 to inject a realistic focal spot, where the focal spot image with the median peak intensity was chosen, shown in Fig. 6.5. A Gaussian temporal profile was used, with $\tau_0 = 38 \text{ fs}$, and the peak a_0 was 2.07. The vacuum focal position was set to coincide with the start of the density plateau, defined as the point at which the density reaches 90% of n_0 , at approximately $2\lambda_R$ from the ramp midpoint. This required that the laser was initialised $\sim 0.2Z_R$ from focus,

²³ [216] R. Lehe *et al.* (2016) *Computer Physics Communications*

²⁴ See Sec. 3.2.2.

which is within the method's domain of applicability. The full 20 mm interaction was simulated for ramps with scale lengths which varied from 0.0–0.5 mm. The simulation window was $z = 100 \mu\text{m}$ long and had a radial extent of $r = 100 \mu\text{m}$, which contained 2000×400 cells, for a resolution of $20\lambda^{-1} \times 4\lambda^{-1}$. In z , r , and θ , each cell contained $2 \times 2 \times 6$ macroparticles, and 3 azimuthal modes were used.

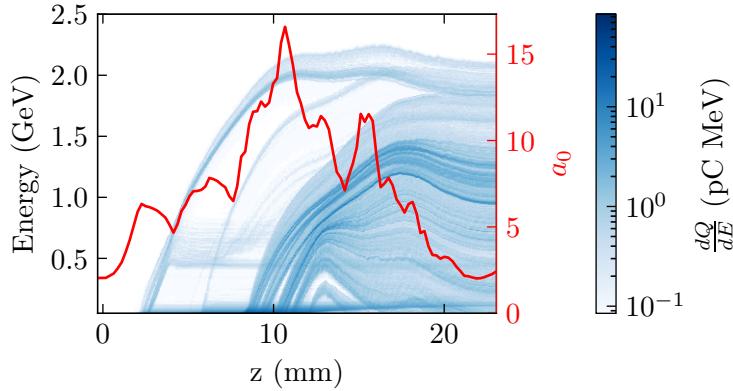


Figure 6.18: Evolution of the electron energy spectrum and laser a_0 during the simulation. The plasma density profile starts at 0 mm and ends at 20 mm

A waterfall plot of one of the simulations is shown in Fig. 6.18, which reveals the mechanism by which the high-energy electron beam is generated. Two main injection events are observed; one at 2 mm, which produces the highest energy electrons, and another that starts at 8 mm and ends at 12 mm, which produces a high charge beam with a large energy spread, roughly centred at 1 GeV. The first electrons are injected when the laser pulse self-focuses to its plasma focal spot size, which results in an $a_0 \sim 6$. Note that this is equivalent to the self-injection threshold predicted by Thomas *et al.*²⁵, $\ln(2\gamma_p^2) - 1$. This beam is then accelerated for the next 8 mm, which is similar to the dephasing length for a beam undergoing etching²⁶. The second injection event is caused by longitudinal compression and amplification of the laser, which causes bubble expansion, reducing the velocity of the back of the bubble, making it easier for injection to occur²⁷. Despite the high a_0 at this point, this part of the beam is injected into a pre-loaded wake²⁸, with a laser pulse that is approaching depletion, resulting in less energy gain than experienced by the first beam. Interestingly, the dephasing length of this second beam is similar to the first, indicating that the laser pulse is still driving a non-linear plasma wave. The depletion length at this density predicted by Shadwick *et al.*²⁹ is 28.7 mm, and so we should expect the laser pulse to still contain appreciable energy after 20 mm of propagation. This is important when we consider the ability of the laser pulse to drive a wake in the ramp at the end of the cell.

²⁵ [153] A. G. R. Thomas (2010) *Physics of Plasmas*

²⁶ [149] C. D. Decker *et al.* (1996) *Physics of Plasmas*

²⁷ [155] S. Kalmykov *et al.* (2009) *Physical Review Letters*

[146] M. Streeter *et al.* (2018) *Physical Review Letters*

²⁸ [244] M. Tzoufras *et al.* (2008) *Physical Review Letters*

²⁹ [148] B. A. Shadwick *et al.* (2009) *Physics of Plasmas*

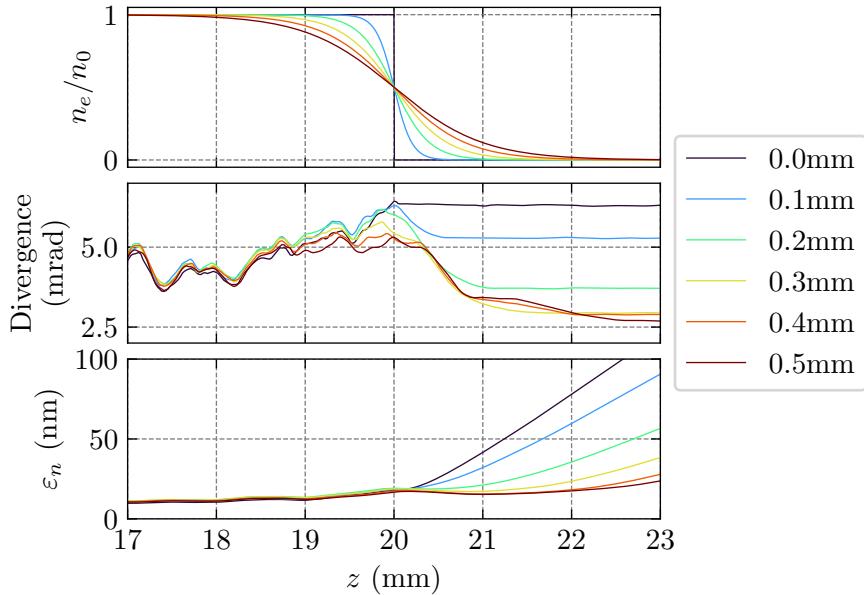


Figure 6.19: Plasma density ramps at the end of the simulations (top), electron beam divergence across all energies (middle), and geometric (phase) normalised emittance (bottom). The divergence and emittance values plotted here are the average of the values projected into the \hat{x} and \hat{y} directions.

Every macroparticle that reached an energy of 1 MeV, and remained within the moving window at the end of the simulation was selected for tracking. For the approximately 3 mm drift space between the end of the plasma and the end of the simulation, this selection criterion was equivalent to a requirement that particles had a divergence of less than 33 mrad, which is larger than the ~ 20 mrad measurement range of the experimental beam profile diagnostic. This resulted in an average of 5.6×10^4 tracked particles in each simulation. The divergence was calculated by measuring the median average deviation, σ_{MAD} , and then scaling to the RMS or FWHM as required³⁰. The trace emittance projected into the \hat{y} direction was calculated by computing $\varepsilon_{tr,y} = \sqrt{\langle y^2 \rangle \langle y'^2 \rangle - \langle yy' \rangle^2}$, where the angled brackets denote averaging over the tracked electrons, and $y' = p_y/p_z$. Similarly, the normalised phase emittance is calculated using $\varepsilon_{ph,y} = (m_e c \bar{p}_z)^{-1} \sqrt{\langle y^2 \rangle \langle p_y^2 \rangle - \langle yp_y \rangle^2}$, where \bar{p}_z is the mean longitudinal momentum.

6.3.1 Beam divergence

The density profiles corresponding to exit ramp scale lengths from 0 mm to 0.5 mm are plotted in the upper panel of Fig. 6.19, while the lower panels of this figure show the evolution of the divergence and emittance. Between $z = 17$ mm and $z = 20$ mm in the central panel of Fig. 6.19, the divergence is observed to gradually increase, as well as undergo small scale oscillations. For such a long propagation length, it might be expected that the range of energies present in the beam would cause the divergence to stabilise, as the different energy slices of the electron beam oscillate incoherently. However, although the energy spread is large, there are still peaks in the energy spectrum that will oscillate

³⁰ $\sigma_{\text{RMS}} \approx 1.4826\sigma_{\text{MAD}}$, and $\sigma_{\text{FWHM}} = 2\sqrt{2 \ln 2} \sigma_{\text{rms}}$ for normally distributed data.

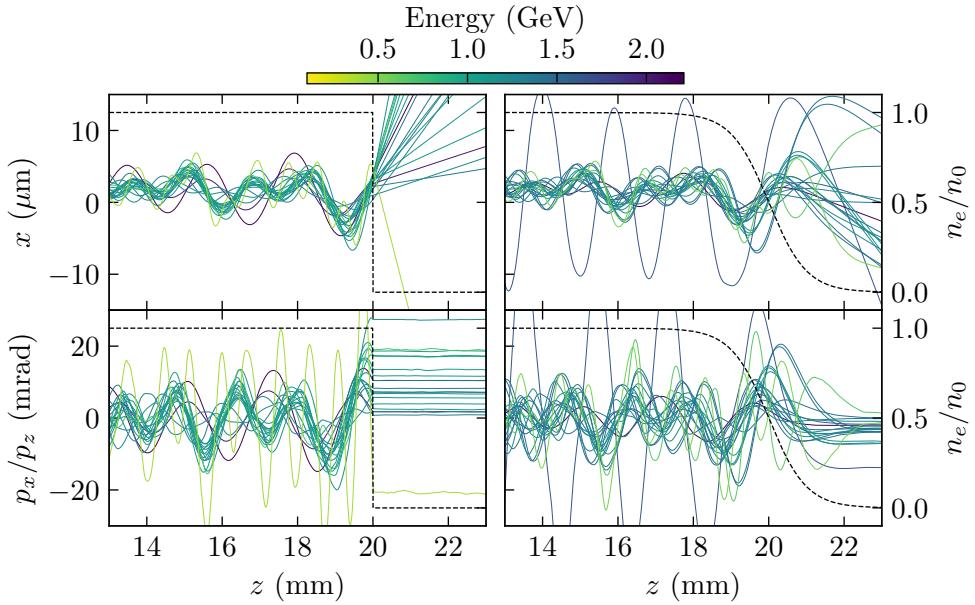


Figure 6.20: Sample of particle trajectories at the accelerator exit. The trajectory colour depends on the instantaneous electron longitudinal momentum. The density profiles are depicted by the black dotted lines.

coherently for long propagation lengths, and it is these high charge features that are responsible for the observed oscillations in the divergence. The gradual increase in the divergence that occurs before the beam enters the vacuum is due to the onset of the hosing instability for parts of the electron beam³¹.

For the step density function ($\lambda_R = 0$), the divergence of the beam in the vacuum is determined by its instantaneous divergence at which it encounters the density transition. For the short ramps ($\lambda_R = 0.1, 0.2\text{mm}$), the divergence is similar to the step density simulation up until $z = 20\text{ mm}$, but then undergoes a smooth decrease as it enters the vacuum, stabilising rapidly. Finally, for the longer ramps ($\lambda_R \geq 0.3$), the divergence can be seen to undergo multiple oscillations as it is gradually reduced, requiring a longer distance to reach a stable value. The bottom panel of the figure shows that the emittance is unaffected by the differences in the divergence before $z = 20\text{ mm}$, but is sensitive to the behaviour of the beam in the vacuum. For all ramps, the emittance is growing in the vacuum, but this growth occurs at different rates and starts at different positions, both of which depend strongly on the ramp scale length.

Fig. 6.20 has been produced to show the behaviour of particles with different energies as they leave the plasma; for the step density transition and the 0.5 mm ramp. A random sub-sample of 20 particles from each simulation has been plotted, and the selection index is kept constant for the two simulations. Consequently, the selected particles from each simulation should have equivalent energies, although, since the particles are from two different simulations, the individual trajectories are slightly different. Within the plasma, the particles are oscillating with an energy-dependent spatial frequency, and the beam is

³¹ [170] J. C. Wood *PhD. Thesis* (2016)

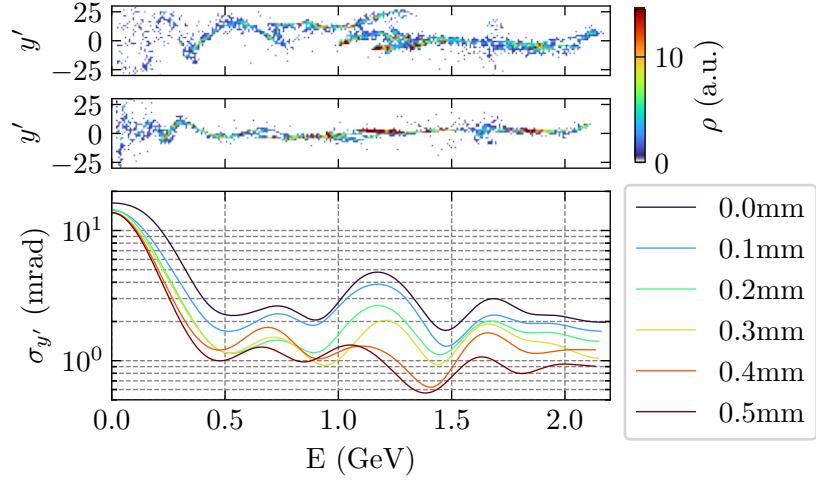


Figure 6.21: Energy-divergence spectrographs for the $\lambda_R = 0$ mm and $\lambda_R = 0.5$ mm simulations are shown in the top and middle panels respectively. The lower panel compares the simulated FWHM divergence for the 0.0 mm, 0.2 mm, and 0.4 mm ramps to the experimental data.

also undergoing collective oscillations. For the case of the step density transition, the transverse momentum of each electron is frozen at the value it has when it encounters the transition. Since there is a large range of transverse momentum values that can be supported in the channel, the spread in transverse momentum after the plasma is high. The density ramp is able to damp the transverse momentum oscillations, resulting in a more focused beam. This is because, as the density drops, the focusing forces in the wake weaken, reducing the range of transverse momentum values that can be supported. The electron beam that then emerges from the plasma has a reduced transverse momentum spread. Note that for the low-energy electrons, which undergo multiple oscillations in the ramp, the amplitude of the oscillations gradually reduces, while for the higher-energy electrons, which only undergo one or two oscillations in the ramp, it is the final oscillation that appears to be affected the most strongly.

Synthetic electron spectrometer images have been generated from the particle beams at the end of the 0.0 mm and 0.5 mm simulations, and are plotted in the upper two panels of Fig. 6.21. These mimic the experimental data plotted in Fig. 6.10. The beam divergence as a function of energy is plotted in the lower panel of Fig. 6.21, calculated statistically using the median average deviation, and then smoothed with a Gaussian weighting function, which had a width of 100 MeV. This shows that the reduction in the divergence is consistent at all energies in the beam, in agreement with the experimental measurements.

6.3.2 Emittance

The evolution of the transverse emittance is plotted in Fig. 6.22, projected onto the x and y axes. For $z < 20$ mm, both the trace and phase space emittances are stable, staying in the 20-30 μm range, while in the drift space for $z > 20$ mm, the phase space

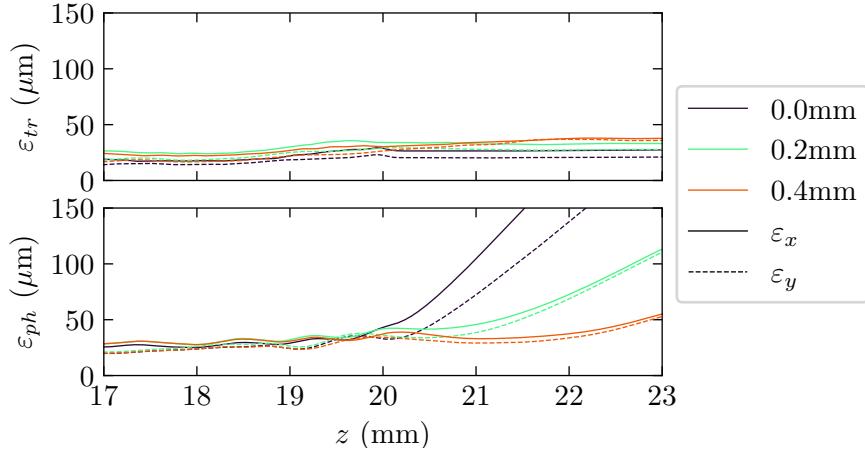


Figure 6.22: Comparison of the behaviour of the trace (top) and phase (bottom) emittances in the density ramp. The solid lines show the projection onto the \hat{x} direction, while the dashed lines show the projection into the \hat{y} direction.

emittance grows linearly in both x and y . The trace space emittance is constant in this region, roughly equivalent to its value at $z = 20$ mm. The x projected emittance is slightly larger due to the interaction of some electrons with the tail of the x polarised laser pulse. This increases the emittance because part of the electron beam is accelerated transversely by the spatially varying field, increasing the transverse momentum spread.

The fact that the trace space and phase space emittances behave differently in the drift space is simply due to their different definitions³². The phase space emittance measures the beam under the assumption that there is a correlation between transverse position and transverse momentum, while the trace space emittance assumes a correlation between transverse position and divergence. The phase space emittance grows in the drift because two particles with the same transverse momentum but different longitudinal momentum will move different amounts transversely, leading to a spreading out in phase space, while two particles with the same divergence will, by definition, move the same amount transversely, preserving the phase space area.

If it were possible to exactly reverse the sign of the divergence of every particle with some focusing optic, then it would be possible to preserve the trace space emittance, causing the phase space emittance to fall to equal the trace space. Since this is not the case, and instead, focusing strength depends on the particle energy, any focusing of the beam will instead cause the trace space emittance to rise to meet the phase space value. Consequently, for multi-stage acceleration applications, the phase space emittance is the more appropriate measure. Migliorati *et al.*³³ derive an equation that parameterises the phase space emittance growth of a beam in a drift space in terms of its energy spread, σ_E , and its divergence, $\sigma_{y'}$, as

$$\varepsilon_n^2 = \langle \gamma \rangle^2 (z^2 \sigma_E^2 \sigma_{y'}^4 + \epsilon_0^2), \quad (6.4)$$

³² [171] K. Floettmann (2003) *Physical Review Special Topics - Accelerators and Beams*

³³ [104] M. Migliorati *et al.* (2013) *Physical Review Special Topics - Accelerators and Beams*

where ϵ_0 is the unnormalised trace space emittance when the beam is at focus inside the plasma channel. For small energy spreads, $\varepsilon_n \simeq \langle \gamma \rangle \epsilon_0$, which is simply the normalised trace space emittance. However, in the presence of energy spread, the emittance growth rate scales with the square of the divergence, $\varepsilon'_n \approx \langle \gamma \rangle \sigma_E \sigma_y^2$. This scaling agrees with the results of our PIC simulations; the emittance growth rate of the 6.5 mrad beam is approximately $100 \mu\text{m mm}^{-1}$, while for the 3.5 mrad beam the growth rate is approximately $25 \mu\text{m mm}^{-1}$. The sensitivity of the emittance growth rate to the divergence makes controlling the divergence a priority for maintaining low emittance.

6.4 Discussion

We will now discuss the results of the preceding analysis. First, we will quantify and compare the effects of the various mechanisms that contribute to the divergence increase, following which we will solve the beta function equation for density ramps of various scale lengths. The effect of divergence change on charge capture efficiency will be estimated, and future multi-stage experiment design considerations will be discussed. Finally, the benefits of expanding the beam optimal ramp shapes will be explained.

6.4.1 Divergence change mechanism

There are three main contributors to the divergence difference between the ET and H configurations: Coulomb scattering, magnetic field deflection, and the effect of the density ramp. These effects add in quadrature, such that the divergence change can be written as $\theta_{ET}^2 - \theta_H^2 = \Delta\theta^2 = \Delta\theta_S^2 + \Delta\theta_B^2 + \Delta\theta_R^2$, where S , B , and R depict the divergence change due to scattering, magnetic fields, and the ramp, respectively. Coulomb scattering should increase electron beam divergence by³⁴

$$\Delta\theta_S(z) = \frac{17 \text{ MeV}}{E[\text{MeV}]} \sqrt{\frac{z}{L_R}} \left[1 + \frac{9}{8} \log_{10} \left(\frac{z}{L_R} \right) \right], \quad (6.5)$$

where z is the length of material traversed, and L_R is the radiation length of the material. For a $25 \mu\text{m}$ thick Kapton tape³⁵, angled at 45° , $\Delta\theta_S = 0.16 \text{ mrad GeV}^{36}$. From the plasma mirror simulations in Sec. 6.2, the magnetic fields in the tape are estimated to increase the divergence by 0.1 mrad. This means that the ramp of length $\sim 0.25 \mu\text{m}$ should be able to account for a divergence change of at least 2.98 mrad of the 3 mrad divergence change observed at 1 GeV. The simulations in Fig. 6.21 indicate that this is possible.

While on the subject of multiple scattering, it is worth commenting on the effect of the injection tape. Counter-intuitively, its placement in the electron beam path reduced the divergence slightly; from $3.38 \pm 0.22 \text{ mrad}$ in the H configuration to $2.76 \pm 0.39 \text{ mrad}$ in

³⁴ [234] M. B. Reid (1991) *Journal of Applied Physics*

[235] V. L. Highland (1975) *Nuclear Instruments and Methods*

³⁵ Kapton radiation length $L_R = 28.57 \text{ cm}$

[245] *Kapton Properties*, Particle Data Group, LBNL

³⁶ The units denote that a 1 GeV beam would experience a divergence increase of 0.16 mrad.

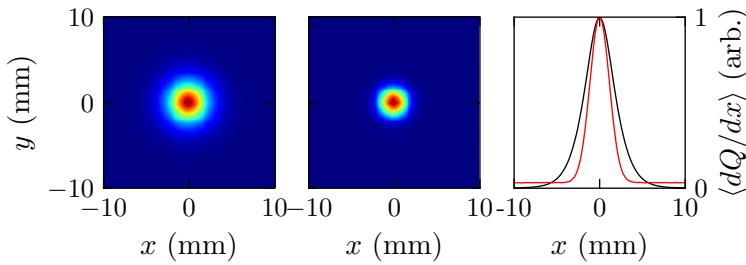


Figure 6.23: Monte Carlo simulation results showing a possible effect of electron beam scattering. The left panel is without scattering while the centre panel includes scattering. Mean transverse intensity profiles are shown in the right panel, with the red line indicating the scattered beam.

the IT configuration. This observation is unlikely to be due only to random fluctuations, given the 39 and 14 shots recorded in the respective setups. Instead, scattering could cause this observation by preferentially removing low energy, high divergence electrons from the beam.

This effect is demonstrated using a Monte Carlo simulation in Fig. 6.23. An electron beam is initialised with two mono-energetic peaks in its spectrum; one at $\gamma = 30$, and the other peak at $\gamma = 3000$. The lower energy component of the beam has a divergence of 2 mrad, and contains 66% of the total charge, while the high energy component's divergence is 1 mrad. The combined beam starts from a point source, propagates 8 mm, scatters by 1 mrad GeV, and then propagated a further 2 m. For this particular configuration, the effect of the scattering is to remove the high divergence, low energy electrons, leaving the unperturbed high energy core. It is possible that this is the mechanism behind the observed effect of the IT configuration.

6.4.2 Ramp scale length

The effect of the density ramp scale length on a beam with a given emittance can be determined numerically. Recall³⁷ the beta function of the electron beam evolves according to

$$\beta'' = -2k_\beta(z)\beta + \frac{2}{\beta} \left[1 + \left(\frac{\beta'}{2} \right)^2 \right], \quad (6.6)$$

where $\beta = \langle x^2 \rangle / \varepsilon$ represents the size of the beam normalised by the emittance, and $k_\beta^2(z) = \omega_p^2(z) / 2\gamma_e c^2$ is the betatron wavenumber for an electron with relativistic factor γ_e . In Fig. 6.24, this equation is solved numerically for emittances between 10 μm and 30 μm , assuming a mean $\gamma = 2000$. For comparison, the results of the PIC simulations and the experimental results are also plotted. The PIC simulation results are in close agreement with the 20 μm solution to Eq. 6.6, as are the experimental median values, while the error bars fit within the solutions for an emittance variation of $\pm 10 \mu\text{m}$.

³⁷ See Sec. 2.8.

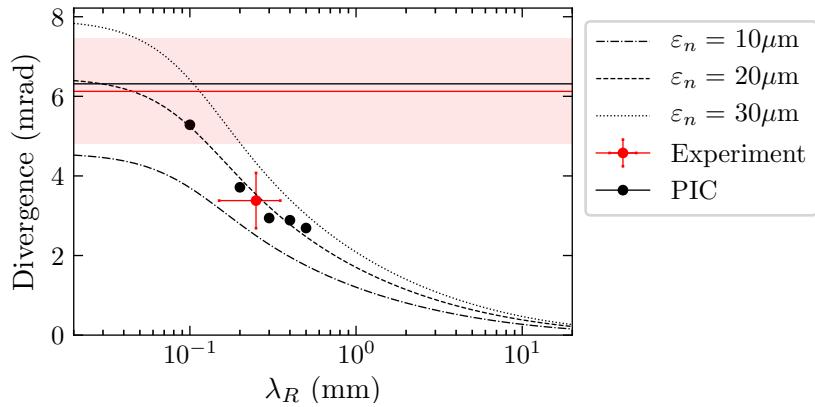


Figure 6.24: Comparison of the experimental results to the PIC simulation results and the numerical model, for various values of the emittance.

The close agreement between the experimental results, the PIC results, and the numerical model suggests a method for measuring emittance experimentally. If the functional form of the density ramp at the end of the accelerator is known accurately, then so is the beta function, and consequently, a single measurement of the beam size should be sufficient to obtain an estimation of the emittance. Further, if the scale length of the density ramp can be controlled, then a parameter scan of this scale length would produce a series of divergence measurements that could be used to measure the emittance via a fitting line, similar to those in Fig. 6.24. In our example, this method is being used to measure the total beam emittance, but could also be used to measure the energetically resolved emittance if it were paired with the electron spectrometer. It should be noted that the solutions in Fig. 6.24 assume a matched beam, which is not necessarily the case. Beam mismatch introduces an additional free parameter, which has the potential to either add noise to the measurement or to shift it systematically.

6.4.3 Charge capture

We will now consider the effect that the increase in divergence would have on the amount of charge that could be captured in a second stage. A minimum requirement for an electron to be captured is that its transverse coordinates are close enough to the second laser pulse's beam axis that the plasma wake produced by this laser pulse is able to focus the electron. This 'close enough' distance will be called the acceptance of the second stage and will be roughly on the scale of the plasma wavelength. Using this minimum requirement, we can compute the maximum fraction of charge that can be captured for a given acceptance. The relevant factors are the divergence of the electron beam, the distance between the stages, the pointing jitter of the electron beam, the spatial jitter of the second laser at the injection plane, and the average distance between the electron beam and the second laser at the injection plane. If we model the transverse fluctuations as being normally distributed, then we can produce an expression for the average electron beam size, with respect to the centre of the acceptance aperture of the

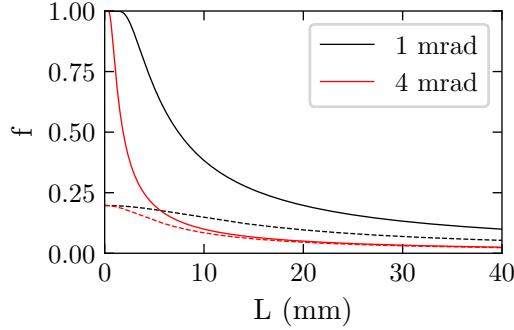


Figure 6.25: Average charge capture fraction as a function of inter-stage separation L , without jitter (solid lines), and with jitter (dashed lines). The 1 mrad beam (RMS) corresponds to the lowest divergence beams, the H configuration, while the 4 mrad beam corresponds to the highest divergence measurements, the ET configuration.

second wake. The expected fraction of charge captured is then given by

$$f = \int_{-R_A}^{R_A} \frac{1}{\sigma_T \sqrt{2\pi}} e^{-\frac{|\mathbf{r}-\mathbf{r}_0|^2}{2\sigma_T^2}} d\mathbf{r}, \quad (6.7)$$

where \mathbf{r}_0 is the vector connecting the centres of the shot patterns of the lasers, and σ_T represents the average electron beam size, averaging over the range of possible shots, given by

$$\sigma_T^2 = \sigma_L^2 + (L\phi_B)^2 + (L\theta_B)^2, \quad (6.8)$$

where σ_L is the standard deviation of the laser focal spot position, ϕ_B and θ_B are the electron beam pointing and divergence standard deviations, and L is the inter-stage distance.

To obtain a value for the acceptance radius, we will assume that the second wake is in the bubble regime, driven by a laser with $a_0 = 3$, at a plasma density of $1 \times 10^{24} \text{ m}^{-3}$. This gives a bubble radius of³⁸

$$r_b \approx \frac{c\sqrt{a_0}}{\omega_p} \sim 10 \mu\text{m}. \quad (6.9)$$

For the calculations presented here, we will approximate that the acceptance of the accelerator is half of this value. This accounts for the need to inject into the accelerating portion of the accelerator, which is narrower than r_b .

In Fig. 6.25, equation Eq. 6.7 is plotted for increasing values of stage separation, L , using an acceptance radius of $5 \mu\text{m}$. The value for ϕ_B is taken from the measured values of the electron beam pointing fluctuations in Tab. 6.1, while the value for σ_L is taken from spatially calibrated focal spot measurements of the North beam fluctuations. These have values of 1.5 mrad and $20 \mu\text{m}$ respectively. For the solid lines in Fig. 6.25, these fluctuations are set to zero, to see the maximum fraction that could be captured

³⁸ [153] A. G. R. Thomas (2010) *Physics of Plasmas*

for a perfectly aligned shot. For $L < 1$ mm, we see that 100% charge capture is possible, but will only happen a small fraction of the time. At large values of L , the difference between the best possible shot and the average vanishes, as at this point the electron beam is much larger than the target wake, overlapping it uniformly but weakly. For the inter-stage separation during our experiment, we would have expected an average charge capture of about 10% for 1 mrad beams, with the best shots capturing about 25% of the charge. Here, the beam divergence refers to the standard deviation of a Gaussian, as opposed to the FWHM used to measure the data. This simple model demonstrates that if the stages are close enough, then 100% charge capture should be possible on occasion, but if a reliable coupling is desired, then the combined pointing jitter of the laser and the electron beam must be reduced.

6.4.4 Future staging design considerations

The target design chosen for this experiment attempted to use the proximity of the stages to maximise charge capture. Even for the low divergence of approximately 1 mrad, this would have limited charge capture to 25%, while the normal operation set-up, with the extraction tape moving, would have reduced the peak capture even more. However, since thousands of shots were taken, we would still have expected to observe energy gain in the second stage, given the signal-to-noise ratio of the spectrometers (of order 1×10^3).

The main limiting factor for the experiment was the effect of aberrations in the near field of the North beam. These were strong enough to rapidly damage the dielectric mirrors if shots were taken above ~ 4 J. This greatly reduced the intensity that could be obtained at the focus of the North beam. Further, these aberrations reduced the focal spot quality such that only 20% of this energy was contained within the FWHM, even before reflection off the tape. Even with the injection tape's high reflectivity of 70%³⁹, the vacuum $a_0 \approx 0.6$, resulting fields of strength 30 MeV m^{-1} . The reduced bubble radius that results from the low a_0 would have had a negative effect on capture efficiency, and the highly aberrated focal spot would have driven a distorted wakefield, further reducing the captured beam quality. All of these factors contributed to the failure to detect staged acceleration.

As in conventional accelerators, other multi-stage LPA concepts utilise an active focusing element in between the stages, such as a plasma lens⁴⁰. The advantage of using such an optic is that it gives a degree of control over the electron beam size as it enters the next stage, potentially allowing for high charge capture and, in principle, injection with a beam size matched to the focusing strength of the second wake. Adding such an optic significantly increases the technical complexity of the experiment; if each plasma stage now requires its own plasma optic, then the number of gas targets has doubled, complicating alignment and increasing the gas load on the pumping system. Further, in the presence of energy spread and energy variation, active focusing does not stand a

³⁹ [227] J.-N. Gruse *PhD. Thesis* (2020)

⁴⁰ [88] S. Steinke *et al.* (2016) *Nature*

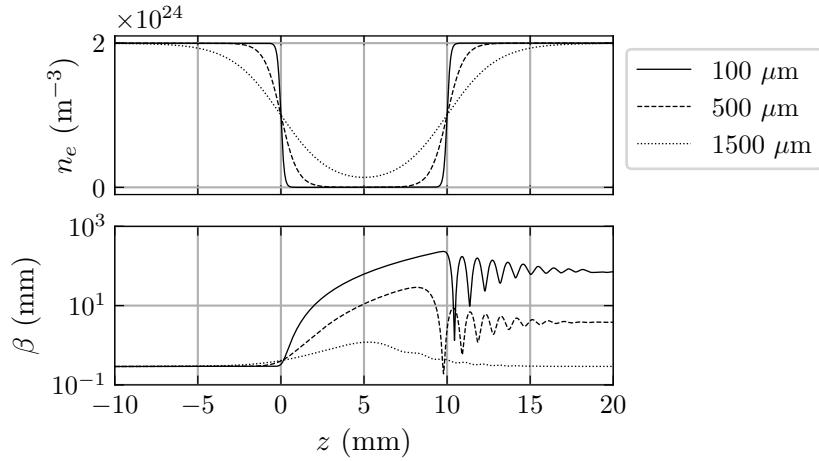


Figure 6.26: Beta function solutions for various ramp scale lengths connecting two plasmas. Beams are initially matched to the plasma with a mean $\gamma = 3000$ and 10% energy spread (RMS).

realistic chance of being able to preserve emittance, as a slight change in energy will cause the injection of a mismatched beam, and any energy spread in the beam will produce uncorrelated emittance growth when the beam is collimated. A method for passively and compactly transferring the electron beam from one stage to another is required if multi-stage acceleration is to be robust.

Connected plasma density ramps offer a method for achieving this. Fig. 6.26 shows two regions of high-density plasma with a 10 mm drift space between them, representing the end and start of two plasma accelerating stages, where the edges of the stages have exponential density ramps of various scale lengths. For each density profile, we consider a beam with a Gaussian energy spectrum with $\langle \gamma \rangle = 3000$ and $\sigma_\gamma = 300$, initially matched to the density profile. The plotted value of the beta function is the charge-weighted mean of the beta function for each energy slice. Between the stages, the beta function expands at a rate that depends inversely on the scale length of the ramp. As the beam enters the second stage its expansion is halted and it instead undergoes damped betatron oscillations, eventually reaching a constant value. This occurs because each energy slice is entering the second stage at a mismatched beam size, consequently undergoing betatron oscillations at an energy-dependent frequency, resulting in a smearing out of the beam in phase space and an increase in emittance. For the longest ramp, this does not occur; instead, the beam remains matched throughout the transfer, such that $\beta(z = -10) = \beta(z = 20)$. Note that this perfect matching only occurs because there is no drift space between the stages. If there is any drift space then the emittance will necessarily increase in the presence of energy spread.

6.4.5 The benefit of beam expansion

If allowing the beam to expand too quickly leads to emittance growth, then the question that naturally arises is: why have a coupling section at all? It might seem more

straightforward to have the plasma mirror bisect a region of constant plasma density, as this would circumvent the need to manage the beta function. Further, removing the coupling section would have the additional benefit of maximising the average accelerating gradient over the two stages.

Unfortunately, such a system would be overly sensitive to alignment, both in terms of charge capture and emittance preservation. To maximise the per-stage energy gain, the electrons should be injected into the bubble as close to the back as possible, but this region becomes increasingly narrow. At acceleration densities of $1 \times 10^{24} \text{ m}^{-3}$, pointing jitter would need to be limited to a few microns if the high accelerating gradients at the back of the bubble are to be exploited. Injecting into an expanded bubble relaxes this requirement to some extent, as once the beam is within the expanded bubble, the focusing forces can steer the beam onto the axis. Even if the laser pulse-electron beam alignment fluctuations could be limited to a few microns, expanding the beam in between the stages would still be advantageous. Off-axis injection will cause betatron oscillations, which will lead to emittance growth. This is because the electrons are not injected at the bottom of the transverse potential well, and this will increase the total transverse momentum in the beam. For the same spatial offset, this increase in transverse momentum will be smaller if the plasma density is lower at the point of injection. Consequently, expanding the beam between the stages reduces emittance growth that occurs due to misalignment.

6.4.6 Optimal density ramps

Finally, we can find the optimal functional form of these ramps. So far, we have considered only exponential profiles, which are convenient from an experimental perspective but are inefficient in terms of space. The betatron wavelength in a ramp of the form $\eta(s) = n_0 \exp(-s/\lambda_R)$ is

$$\lambda_\beta(s) = 2\pi \left(\frac{2\gamma_e c^2}{\omega_{p,0}^2} \right)^{\frac{1}{2}} \frac{1}{\sqrt{\eta(s)}} = \frac{\lambda_{\beta,0}}{\sqrt{\eta(s)}}. \quad (6.10)$$

In contrast, the exponential ramp varies on the scale of $\eta/\eta' = \lambda_R$, such that the beam will eventually cease to see the ramp as varying slowly. To negate emittance growth in an exponential ramp, the scale length needs to be longer than strictly necessary since the ramp will change too slowly when the density is high. An optimal ramp will maintain a constant ratio between the betatron wavelength and variation scale. Ariniello *et al.*⁴¹ define an adiabaticity parameter related to the beta function as being equal to this ratio,

$$\alpha_m = \lambda_\beta(s) \left| \frac{\eta'}{\eta} \right| \ll 1 \quad (6.11)$$

where $\alpha = -\beta'/\varepsilon$ is the normalised correlation in the beam, and α_m is defined as the correlation that keeps the emittance growth within acceptable limits. An optimal ramp

⁴¹ [246] R. Ariniello *et al.* (2019) *Physical Review Accelerators and Beams*

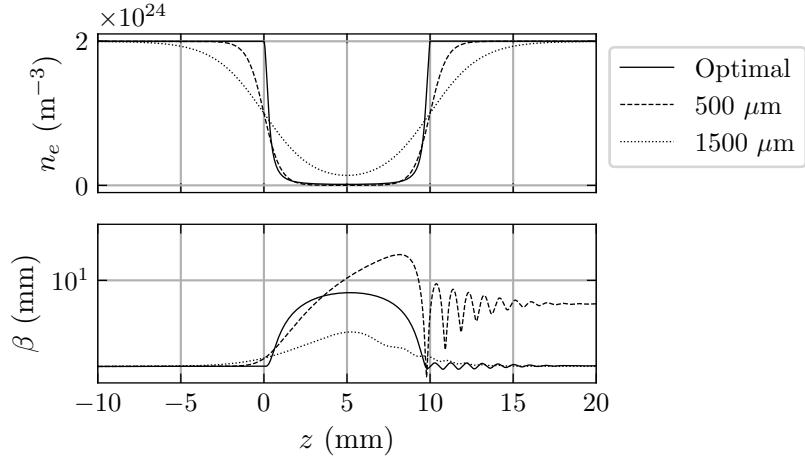


Figure 6.27: Optimal density ramp compared to the exponential ramps. The parameter $a = 10^{-3}k_\beta$.

will have a constant α_m , which was shown to be achieved for

$$\eta(s) = \frac{1}{1 + (as)^2}, \quad (6.12)$$

where the parameter a can tune the acceptable level of emittance growth.

By using a ramp that is optimally tailored to the electron energy, a large increase in the beam size while maintaining the emittance can be achieved with a much smaller ramp. This is shown in Fig. 6.27 for a ramp with $a = 10^{-3}k_\beta$. This ramp is shorter than the 500 μm exponential ramp and achieves a larger beta function between the stages while also performing similarly to the 1500 μm in the second stage. Such a ramp would allow for a larger beam size to be achieved between the stages, or a reduced inter-stage distance for the same beam size.

6.5 Summary

Beam divergence measurements have been made using various cell configurations that are of interest for staged wakefield acceleration. Termination of the accelerator with a plasma mirror was found to increase total beam divergence from 3.38 ± 0.22 mrad to 6.13 ± 0.94 mrad, an effect that was found to still be significant at the high energies that are of interest for staging. The presence of the density ramp was shown to account for the observed divergence change, while the effects of magnetic fields generated in the plasma mirror on the beam were shown to be small for the studied experimental configuration. The angle of the tape was found to reduce the impact of the mirror's internal magnetic fields on the electron beam divergence in the y direction. The results were discussed in the context of beam transfer between laser wakefield stages, and the use of density ramps was suggested to facilitate coupling.

7 Conclusions and outlook

IN THIS THESIS, the quality of electron beams produced by laser wakefield accelerators have been measured and optimised. To this end, the density length parameter space of a laser wakefield accelerator was characterised, an emittance diagnostic was developed and implemented, and target designed for the purpose of multi-stage acceleration were investigated. In this chapter, we present a summary of the findings and discuss future directions.

7.1 Conclusions

7.1.1 The density-length parameter space

The density-length parameter space of a 5 TW laser wakefield accelerator has been explored in chapter 4, using plasma densities from $0.3 - 2.8 \times 10^{24} \text{ m}^{-3}$, and cell lengths from $0 - 3 \text{ mm}$. Two different gas species were used; pure helium and helium doped with 1% nitrogen, and each 2D scan took 5 laser shots at 16×21 locations, for a total of 3,360 shots over the two scans. Before performing the scans, an automated Bayesian optimisation routine was used to ensure that the laser pulse shape and focal position maximised the signal on the electron spectrometer camera.

In general, it was found that the doped gas produced both higher charge and higher energy electron beams than the pure helium gas, although the beams from the pure gas were more collimated. Notably, electron beams with peak energies above 200 MeV were measured, which is much larger than would be predicted by typically used scaling laws. The combination of higher charge and energy resulted in a higher x-ray signal from the doped gas, while the pure helium gas produced x-ray spectra with higher critical energy. See Tab. 6.1 for details. The electron beam results were found to be in reasonable agreement with PIC simulations, with the exception of the self-injected beam charge, which was much higher in simulations.

The dephasing length of the ionisation injected beam was found to be longer than expected, resulting in higher energies. This was attributed to an increase in the phase velocity of the plasma wave, from $\beta_\phi = \beta_{\text{etch}} \simeq 1 - 3n_e/2n_c$ to $\beta_\phi \simeq 1 - 5n_e/6n_c$, which was found to accurately predict the speed of the back of the bubble in PIC simulations. The enhanced phase velocity was attributed to the laser pulse depletion process being

only semi-localised, such that the phase velocity fell between the linear group velocity and β_{etch} . The exceptionally high (> 200 MeV) energies measured using the doped gas were attributed to beam-driven acceleration, identified by observing increasing energy with plasma length up to approximately 2 mm at 2×10^{25} , much longer than the depletion length at this density.

The pulse evolution rate was measured by monitoring the injection point as a function of density over the length scans. The depletion length scaling of $L_d = \lambda_p n_c / n_e$ was found to be too short to explain the pulse evolution length. However, if the pulse first needed to self-focus over a length L_f before depletion could begin, then it was found that the combined length $L_f + L_d$ fitted the experimental data well. It was concluded that, in LWFA regimes where the a_0 of the laser pulse crosses from $a_0 \lesssim 1$ to $a_0 \gtrsim 1$, then the evolution rate is reduced by the need for the pulse to self-focus almost fully before it can begin to compress longitudinally.

The x-ray data was analysed with the goal of determining a value for the brilliance. From PIC simulations, the source size was determined to be $\sim 1 \mu\text{m}$, which led to brilliance measurements in the region of $10^{20} [\text{B}]$, where $[\text{B}] = \text{ph s}^{-1} \text{mm}^{-2} \text{mrad}^{-2} 0.1\% \text{BW}^{-1}$. In addition, a commonly used method for determining the source size using the electron beam energy was used, which produced brilliance measurements that were approximately two orders of magnitude too high. It was concluded that such a measurement method is too inaccurate for beams with high energy spread since, although the highest energy electrons are responsible for the majority of the photon flux, large amounts of charge at intermediate energies could offset this, biasing the critical energy measurement.

7.1.2 Ionisation injection emittance measurements

In chapter 5, we developed an emittance diagnostic and used it to measure the emittance of an ionisation injected GeV electron beam. At moderate backing pressures, the normalised emittance was measured to be in the range of $4 - 9 \mu\text{m}$, over a density range of $1 - 2 \times 10^{24} \text{ m}^{-3}$. At a higher backing pressure, this emittance jumped to approximately $30 \mu\text{m}$, over a density range of $2 - 3 \times 10^{24} \text{ m}^{-3}$. Simulations performed using the experimental conditions produced results that agreed with the low-density measurements, but did not agree with high-density measurements. When compared to the literature, these are the lowest emittance measurements made using a beam mask method and are made at the highest energies. The sources of error for the diagnostic method were explored. It was found that, for low divergence beams, the small number of samples contributed significantly to random error, and if the sample number dropped too low, then the measured emittance was systematically low. To overcome the sampling errors, for reasonable vacuum chamber sizes, an analysis method that only requires partial scattering of the beam can be used.

A scaling was derived to estimate the emittance of ionisation injected beams for the direction orthogonal to the laser polarisation, which predicted $\varepsilon \propto \omega_p^2 y_\sigma^3$. This scaling

produced results that agree with the experimental results in the $1 - 2 \times 10^{24} \text{ m}^{-3}$, provided that the r.m.s. beam size at the moment of injection was approximately $7 \mu\text{m}$. This beam size was verified using the simulations. Larger emittances than predicted by this model could occur at some energies due to injection into a partially formed wake, suggesting that a localised dopant would be best placed at a point where a stable wake has formed.

7.1.3 Beam divergence considerations for staged LWFA

Finally, in chapter 6 we investigated a novel gas cell design for use in staged acceleration experiments. Experimental measurements of the electron beam profile and spectrum have indicated that terminating a LWFA accelerator with a $25 \mu\text{m}$ thick tape instead of a 1 mm aperture increases electron beam divergence by approximately 75%, from $3.38 \pm 0.22 \text{ mrad}$ to $6.13 \pm 0.94 \text{ mrad}$. An additional, thicker tape, placed 8 mm after the aperture further reduced total beam divergence, from $3.38 \pm 0.22 \text{ mrad}$ to $2.76 \pm 0.39 \text{ mrad}$.

Through PIC simulations of the wakefield acceleration and a numerical model, it was shown that the density ramp that forms by gas leaking through the hole accounted for the experimentally observed divergence change. Fluid simulations indicated that an exponential density ramp is formed at the plasma-vacuum boundary, with a scale length determined by the hole radius. For a scale length of approximately 0.25 mm , corresponding to a 1 mm diameter hole, the divergence reduction observed in simulations and the numerical model agreed well with the experimental measurements. The mechanism for the divergence reduction is the adiabatic reduction of the plasma wave's focusing forces, which allows the beam to expand to a large beam size with a lower divergence before entering the vacuum. This was shown to effectively reduce emittance growth that occurs in the drift space after the plasma.

The magnetic fields generated in the tape by the laser were investigated using high-resolution 3D PIC simulations, for both 0° and 45° geometries. The input temporal profile of the laser was obtained from LWFA simulations, accurately accounting for the compression and depletion that occurred. Magnetic fields were generated in the tape via the Weibel instability, which had grown to micron-scale structures by the time the electron beam reached the tape. When probed with a test electron beam, the fields caused a spatially dependent deflection that increased divergence by approximately 0.1 mrad for a 2 GeV beam. The fields were geometry dependent; stronger fields were observed inside the tape in the 0° orientation, while the 45° degree orientation resulted in weaker internal fields, but additional fields external to the tape. The RMS variations of the internal integrated field strengths reached values as high as $1 \text{ kT} \mu\text{m}$ when the peak of the laser was incident, but these decayed rapidly before the arrival of the electron beam. The reduced integrated B_x field variation in the 45° tape resulted in a reduced divergence increase in the y direction, indicating that tape angle has an effect on emittance preservation.

7.2 Summary and outlook

7.2.1 Outlook

The quantity of data presented in chapter 4 permits many possible future directions of study. One of these is determining the exact mechanism that produced the enhanced x-ray flux measured in the doped gas, as understanding this mechanism in detail could help design experiments that directly exploit it. X-ray production is one of the primary use cases of wakefield accelerators at this scale, and so developing this facet of the performance would be of great interest to imaging applications.

The charge injected in FBPIC simulations by ionisation injection matched the amount measured experimentally, while the self-injected charge was significantly higher in simulations. This issue is common in simulations of self-injection, with all convergence tests performed so far giving consistently high results. The effect of this over-injection of charge is to cast doubt on the simulation of electrons injected after the first electrons, as beam loading and beam-driven acceleration are sensitive to charge. The explanation given in chapter 4 was that the overly stable bubble caused by the assumption of cylindrical symmetry is responsible. Investigating this using a fully-3D code and simulating the spatiotemporal phase of a realistic laser pulse would be of wide interest.

The emittance diagnostic presented in chapter 5 obtained proof of principle measurements. As discussed, using more sophisticated analysis methods would enable this diagnostic to operate more precisely and achieve a higher resolution while making the effect of the mask less invasive, such that it could be used as with the electron spectrometer to perform charge-calibrated measurements. In addition, using this diagnostic with a larger magnification and an optimised imaging system would increase its resolution to the sub-micron level. This could then be used to investigate various injection mechanisms, or implemented in real-time so that reducing emittance could be an optimisation goal. This latter application would likely require advanced numerical methods to quickly extract the electron beam phase space distribution from the data.

The natural progression from the work presented in chapter 6 would be to perform the same experiment again, except using long density ramps between the cells and a single tape to facilitate both the extraction and injection of the laser pulses. Further, additional considerations should be made for the fluence of the North laser on the dielectric optics to maximise the laser power in the second cell. While the achievement of this aim is an important milestone on the road to high-energy multi-stage LWFA, it would not address the emittance preservation issue. Consequently, it would be extremely valuable to demonstrate coupling that not only has a high charge capture efficiency but also maintains emittance, even in the presence of energy spread and pointing jitter. The development of all-plasma beam coupling solutions might be able to achieve this, provided the driver intensity could be maintained during the coupling.

The effect of magnetic fields generated in a plasma mirror on electron beams warrants further study. While one literature study that pertains to these results exists¹, the effects of tape angle, density, pre-plasma, and laser variations remain unexplored. In chapter 6, simulations were performed for the specific parameters investigated experimentally, but this is only a small look into the range of parameters that could be of interest for staging applications.

7.2.2 Summary

In the introduction of this thesis, the near, medium, and long-term use cases for LWFAs were outlined. Chapter 6 focused on the most ambitious of these aims, that of developing multi-stage accelerators such that a useful TeV plasma-based collider could, one day, be constructed. Due to the relatively low repetition rate of linear colliders, a small emittance is paramount, and although LWFAs are naturally suited to generating low emittance electron beams, extracting these beams from the wakefield and injecting them into subsequent micron-scale accelerating stages while maintaining the emittance is challenging. The experimental results show the beneficial effect that density ramps have on the divergence, and consequently, the emittance growth rate, motivating their inclusion in future beam transport designs. In addition, the effects of angled plasma mirrors, used for injecting and extracting laser pulses between accelerating stages, were characterised, and these effects will become more important as the number of accelerating stages increases.

To improve the emittance, it is helpful to be able to measure it, and current methods are either technically involved and require sophisticated apparatus (using quadrupoles), or rely on indirect methods that are likely to underrepresent the emittance (inference from x-ray measurements). Beam mask methods suffer from neither of these issues, but are difficult to implement at high energies owing to the penetration depth of high-energy electrons. In chapter 5 the beam mask method was developed so that it could measure the energetically resolve emittance of GeV energy beams. This method offers a means by which smaller-scale LWFA facilities can monitor and develop the transverse beam quality.

In contrast, chapter 4 focused on developing a LWFA source for the most near-term applications; that of using a compact electron beam and x-ray source. This chapter showed that lower power, small footprint laser systems could produce remarkably bright x-ray beams when optimised, such that they could find widespread use in industrial and medical imaging. Near-term use cases such as these are necessary, not only to solve problems in the present but also for providing motivation and funding to continue to develop the technology for the decades that will be required if the most ambitious aims of the field are to be reached.

¹ [92] G. Raj *et al.* (2020) *Physical Review Research*

A Appendix

A.1 Fundamental constants

Tab. A.1 lists the constants used in this thesis.

Table A.1: Fundamental constants used in this thesis.

Symbol	Name	Value	Unit
c	Speed of light	2.998×10^8	m s^{-1}
m_e	Electron rest mass	9.109×10^{-31}	kg
$m_e c^2$	Electron rest energy	0.5110	MeV
e	Elementary charge	1.602×10^{-19}	C
ε_0	Vacuum permittivity	8.854×10^{-12}	F m^{-1}
μ_0	Vacuum permeability	1.256×10^{-6}	N A^2
h	Planck's constant	6.626×10^{-34}	$\text{m}^2 \text{kg s}^{-1}$
\hbar	hbar	1.054×10^{-34}	$\text{m}^2 \text{kg s}^{-1}$
r_e	Classical electron radius	2.818×10^{-15}	m

A.2 Maxwell's equations

Maxwell's equations are a set of four equations that are foundational to electromagnetism:

$$\nabla \cdot \mathbf{E} = \frac{\rho}{\varepsilon_0}, \quad (\text{A.1})$$

$$\nabla \cdot \mathbf{B} = 0, \quad (\text{A.2})$$

$$\nabla \times \mathbf{E} = -\frac{\partial \mathbf{B}}{\partial t}, \quad (\text{A.3})$$

$$\nabla \times \mathbf{B} = \mu_0 \mathbf{J} + \mu_0 \varepsilon_0 \frac{\partial \mathbf{E}}{\partial t}. \quad (\text{A.4})$$

The first of these is Gauss's law, which describes the electric fields, \mathbf{E} , from distributions of charge, ρ . The second, known as Gauss's law for magnetic fields, \mathbf{B} , is equivalent to the statement that there are no magnetic monopoles. Faraday's law relates a changing magnetic field to the curl of the electric field, and is the principle by which electricity is generated. Finally, the fourth equation is Ampère's law with Maxwell's addition, which describes how magnetic fields are produced by current, \mathbf{J} , and time-varying electric fields

The fields \mathbf{E} and \mathbf{B} can be written in terms of a scalar electrical potential, ϕ , and a vector potential, \mathbf{A} , using

$$\mathbf{E} = -\nabla\phi - \frac{\partial\mathbf{A}}{\partial t}, \quad (\text{A.5})$$

$$\mathbf{B} = \nabla \times \mathbf{A}. \quad (\text{A.6})$$

The wave equation in a vacuum can be produced by taking the curl of Eq. A.3 with $\mathbf{J} = 0$,

$$\nabla \times (\nabla \times \mathbf{E}) = \frac{\partial(\nabla \times \mathbf{B})}{\partial t}, \quad (\text{A.7})$$

$$\nabla(\nabla \cdot \mathbf{E}) - \nabla^2 \mathbf{E} = \mu_0 \varepsilon_0 \frac{\partial^2 \mathbf{E}}{\partial t^2}. \quad (\text{A.8})$$

In vacuum, $\nabla \cdot \mathbf{E} = 0$, and since $\mu_0 \varepsilon_0 = c^{-2}$, we retrieve the wave equation, Eq. 2.1.

In addition, charge conservation is equivalent to the statement that

$$\frac{\partial \rho}{\partial t} + \nabla \cdot \mathbf{J} = 0. \quad (\text{A.9})$$

A.3 Relativistic single particle motion

We consider the motion of an electron in a relativistic field, where $a_0 > 1$. Following Gibbon¹, the electromagnetic fields will be expressed in terms of the vector potential² $\mathbf{E} = -\partial_t \mathbf{A}$, $\mathbf{B} = \nabla \times \mathbf{A}$. We choose $\mathbf{A} = A_x \hat{\mathbf{x}} = a_0 \sin \theta \hat{\mathbf{x}}$, where we have used, and will be using, the normalised units of $e\mathbf{A}/m_e c = \mathbf{A}$, $\mathbf{v}/c = \mathbf{v}$, and $\mathbf{p}/mc = \mathbf{p}$. The phase of the wave is given by $\theta = \omega t - kz = ct - z$. Using these definitions, the Lorentz force can be written as,

$$\frac{d\mathbf{p}}{dt} = -\frac{\partial \mathbf{A}}{\partial t} + [\mathbf{v} \times (\nabla \times \mathbf{A})], \quad (\text{A.10})$$

$$= -\frac{\partial A_x}{\partial t} \hat{\mathbf{x}} + \left(-v_z \frac{\partial A_x}{\partial z} \hat{\mathbf{x}} + v_x \frac{\partial A_x}{\partial z} \hat{\mathbf{z}} \right), \quad (\text{A.11})$$

and the rate of energy change is given by the power, $\mathbf{v} \cdot \mathbf{F}$,

$$\frac{d\gamma}{dt} = -v_x \frac{\partial A_x}{\partial t}. \quad (\text{A.12})$$

If we first look at the $\hat{\mathbf{x}}$ component, by using the definition of the convective derivative³, we can see that this is equal to the total derivative,

$$\frac{dp_x}{dt} = \frac{\partial A_x}{\partial t} + v_z \frac{\partial A_x}{\partial z} = \frac{dA_x}{dt}. \quad (\text{A.13})$$

¹ [241] P. Gibbon *Short pulse laser interactions with matter* (2005)

² See appendix A.2.

³ $\frac{d}{dt} = \frac{\partial}{\partial t} + \mathbf{v} \cdot \nabla$.

Integrating, the transverse momentum is simply the sum of the potential and the initial momentum,

$$p_x = A_x + p_{x,0}. \quad (\text{A.14})$$

To determine the longitudinal momentum, we compare the $\hat{\mathbf{z}}$ component of Eq. A.10 to the rate of change of energy,

$$\frac{dp_z}{dt} - \frac{d\gamma}{dt} = v_x \left(\frac{\partial A_x}{\partial t} + \frac{\partial A_x}{\partial z} \right) = 0. \quad (\text{A.15})$$

Integrating with respect to time gives the relation

$$\gamma - p_z = \alpha, \quad (\text{A.16})$$

where α is a constant of the integration. Using the definition of $\gamma^2 = 1 + |\mathbf{p}|^2$, Eq. A.16 can be rearranged to show that,

$$p_z = \frac{1 + p_x^2 - \alpha^2}{2\alpha}. \quad (\text{A.17})$$

Taking $p_{x,0} = 0$ in Eq. A.14, and $\alpha = 1$, the equations of motion are

$$p_x = a_0 \sin \theta, \quad (\text{A.18})$$

$$p_z = \frac{a_0^2}{2} \sin^2 \theta. \quad (\text{A.19})$$

To find the coordinates, these equations are integrated with respect to time. Recognising $p_z = \gamma \dot{z}$, the total derivative of the phase is

$$\frac{d\theta}{dt} = \frac{\partial \theta}{\partial z} + \frac{dz}{dt} \frac{\partial \theta}{\partial z} = \frac{\alpha}{\gamma}, \quad (\text{A.20})$$

where in the second equality, Eq. A.16 has been used. From here, we can express the momentum as

$$p_x = \gamma \frac{d\theta}{dt} \frac{dx}{d\theta} = \alpha \frac{dx}{d\theta}. \quad (\text{A.21})$$

For $\alpha = 1$, the coordinates are given by

$$x = a_0 \cos \theta, \quad (\text{A.22})$$

$$z = \frac{a_0^2}{4} \left(\theta - \frac{1}{2} \sin 2\theta \right). \quad (\text{A.23})$$

A.4 Tunnel ionisation rate for specific electrons

The rate given by Eq. 2.24 can be used to estimate the tunnel ionisation rate for electrons at an effective principle quantum number n^* . To calculate the rate for an electron with

quantum numbers m and l , the ADK rate is

$$W_{\text{ADK}}(E) = \omega_0 C_{xl}^2 X_1 \left(\frac{2E_0}{E} \right)^{X_2} e^{-\frac{2}{3} \frac{E_0}{E}} \quad (\text{A.24})$$

$$\frac{E_0}{E_H} = \left(\frac{\omega}{\omega_H} \right)^{\frac{3}{2}}, \quad X_1 = \frac{(2l+1)(l+|m|)!}{2^{|m|}|m|!(l-|m|)!}, \quad X_2 = 2Z\sqrt{\frac{\omega_H}{\omega_0}} - |m| - 1, \quad (\text{A.25})$$

where Z is the post-ionisation charge of the ion, and the dimensionless constant C_{xl} is difficult to calculate for complex atoms, but taking the hydrogen value of $C_{xl} = C_0 = 2$ will be sufficient for most applications. The binding energy of the electron is equal to $\hbar\omega$, and E_0 is the scaled Coulomb field, where the subscript H indicates energy and field of the ionisation of hydrogen, $\hbar\omega_H = 13.6\text{eV}$.

A.5 Plasma dispersion relation

Consider an electromagnetic plane wave which has an electric component of the form $\mathbf{E} = \mathbf{E}_0 e^{i(\mathbf{k} \cdot \mathbf{r} - \omega_L t)}$. The operations ∂_t and $\nabla \times$ become $-i\omega_L \mathbf{\hat{r}}$ and $i\mathbf{k} \times$. Using these operations, Faraday's law, Eq. A.3, becomes

$$i\mathbf{k} \times \mathbf{E} = i\omega_L \mathbf{B}. \quad (\text{A.26})$$

Taking the curl, and using the orthonormality of \mathbf{E} and \mathbf{k} , ($\mathbf{k} \cdot \mathbf{E} = 0$), this becomes

$$-k^2 \mathbf{E} = -i\omega_L \mathbf{k} \times \mathbf{B}. \quad (\text{A.27})$$

Using this equation, Ampère's law becomes

$$-\frac{k^2}{\omega_L} \mathbf{E} = \mu_0 \frac{\partial \mathbf{J}}{\partial t} - i\omega_L \mu_0 \varepsilon_0 \mathbf{E}, \quad (\text{A.28})$$

Finally, recognising that the current $\mathbf{J} = -en_e \mathbf{v}$, and from the Lorentz force, $m_e \partial_t \mathbf{v} = -e \mathbf{E}$, the current is given by the electric field.

$$\mathbf{J} = \frac{n_e e^2}{m_e \omega_L} \mathbf{E}, \quad (\text{A.29})$$

Substituting into Eq. A.28 gives

$$-c^2 k^2 \mathbf{E} = \omega_p^2 \mathbf{E} - \omega_L^2 \mathbf{E}, \quad (\text{A.30})$$

which reduces to the dispersion relation.

A.6 The quasistatic approximation

A convenient frame of reference to use for many underdense laser-plasma interactions is one that moves at approximately the speed of the laser. This is usually taken as the group velocity of the laser pulse, or the phase velocity of the plasma wave. For low-density

plasmas, $v_g \simeq c$. In addition to moving to this frame of reference, it is also common to neglect the evolution of the driver, such that the behaviour of the plasma is a function of co-moving position only, and not time. Taken together, the coordinate shift and slowly-varying approximation form the quasistatic approximation. The coordinates (x, t) become (ξ, τ_Q) , where $\xi = x - ct$ and $\tau_Q = t$. The spatial derivative is left unchanged,

$$\frac{d}{dx} = \frac{d}{d\xi}; \quad \frac{d^2}{dx^2} = \frac{d^2}{d\xi^2}, \quad (\text{A.31})$$

while the first-order temporal derivative becomes

$$\frac{d}{dt} = \frac{d}{d\tau_Q} - c \frac{d}{d\xi} \simeq -c \frac{d}{d\xi}, \quad (\text{A.32})$$

where the approximation in the second line is the slowly changing approximation. The second-order temporal derivative is

$$\frac{d^2}{dt^2} = c^2 \frac{d^2}{d\xi^2}. \quad (\text{A.33})$$

A.7 Hamiltonian of electron in a 1D plasma wave

Following Esarey and Pilhoff⁴, consider a test electron moving in a plasma wave. The electron gains energy $\delta\gamma$ by moving δx in the longitudinal field given by the gradient of the potential,

$$m_e c^2 \delta\gamma = -e E_x \delta x, \quad (\text{A.34})$$

$$= e \frac{d\phi}{dx} \delta x, \quad (\text{A.35})$$

$$= -e \frac{d\phi}{d\xi} \delta x, \quad (\text{A.36})$$

where in the last line, we have switched to the moving frame. We now construct the Hamiltonian $\mathcal{H} = \mathcal{H}(p(t), q(t))$, where we use γ and ξ as proxies for the momentum and position, p and q , while x replaces the parametric variable t . Using Hamilton's equations,

$$\frac{\partial \mathcal{H}}{\partial \xi} = -\frac{d\gamma}{dx}, \quad \frac{\partial \mathcal{H}}{\partial \gamma} = \frac{d\xi}{dx}, \quad (\text{A.37})$$

the Hamiltonian is given by

$$\mathcal{H} = \int \frac{d\xi}{dx} d\gamma. \quad (\text{A.38})$$

⁴ [137] E. Esarey *et al.* (1995) *Physics of Plasmas*

From the definition of the moving frame,

$$\frac{d\xi}{dx} = \frac{d}{dx}(x - v_p t), \quad (\text{A.39})$$

$$= 1 - v_p \frac{dt}{dx}, \quad (\text{A.40})$$

$$= 1 - \frac{\beta_p}{\beta}. \quad (\text{A.41})$$

Using $\beta = \sqrt{\gamma^2 - 1}/\gamma$,

$$\mathcal{H} = \int (1 - \beta_p \frac{\gamma}{\sqrt{\gamma^2 - 1}}) d\gamma, \quad (\text{A.42})$$

$$= \gamma - \beta_p \sqrt{\gamma^2 - 1} + f(\xi), \quad (\text{A.43})$$

$$= \gamma - \beta_p \sqrt{\gamma^2 - 1} - \Phi, \quad (\text{A.44})$$

where in the last line we have used Eq. A.36, with $\phi = mc^2\Phi/e$. From here, Eq. 2.82 can be derived using $\gamma = \sqrt{1 + p^2}$,

$$\mathcal{H}(\gamma, \xi) = \gamma - \beta_p p - \Phi(\xi). \quad (\text{A.45})$$

Bibliography

- [1] E. Rutherford. “LXXIX. The scattering of α and β particles by matter and the structure of the atom”. In: *Lond. Edinb. Dublin philos. mag. j. sci.* 21, pp. 669–688, 1911, pp. 669–688. (See p. 12)
- [2] J. J. Thomson. “XXIV. On the structure of the atom: an investigation of the stability and periods of oscillation of a number of corpuscles arranged at equal intervals around the circumference of a circle; with application of the results to the theory of atomic structure”. In: *Lond. Edinb. Dublin philos. mag. j. sci.* 7, pp. 237–265, 1904, pp. 237–265. (See p. 12)
- [3] E. F. Haussecker and A. W. Chao. “The Influence of Accelerator Science on Physics Research”. In: *Physics in Perspective* 13, p. 146, 2011, p. 146. (See p. 12)
- [4] A. Collaboration, C. M. S. Collaboration, G. Aad, B. Abbott, J. Abdallah, O. Abdinov, R. Aben, M. Abolins, O. AbouZeid, H. Abramowicz, H. Abreu, R. Abreu, Y. Abulaiti, B. Acharya, L. Adamczyk, D. Adams, J. Adelman, S. Adomeit, T. Adye, A. Affolder, T. Agatonovic-Jovin, J. Aguilar-Saavedra, S. Ahlen, F. Ahmadov, G. Aielli, H. Akerstedt, T. Åkesson, G. Akimoto, A. Akimov, *et al.* “Combined Measurement of the Higgs Boson Mass in pp Collisions at $\sqrt{s} = 7$ and 8 TeV with the ATLAS and CMS Experiments”. In: *Physical Review Letters* 114, p. 191803, 2015, p. 191803. (See p. 12)
- [5] M. K. Gaillard, P. D. Grannis, and F. J. Sciulli. “The standard model of particle physics”. In: *Reviews of Modern Physics* 71, S96–S111, 1999, S96–S111. (See p. 12)
- [6] A. Joyce, B. Jain, J. Khoury, and M. Trodden. “Beyond the cosmological standard model”. In: *Physics Reports* 568, pp. 1–98, 2015, pp. 1–98. (See p. 12)
- [7] A. G. Riess, A. V. Filippenko, P. Challis, A. Clocchiatti, A. Diercks, P. M. Garnavich, R. L. Gilliland, C. J. Hogan, S. Jha, R. P. Kirshner, B. Leibundgut, M. M. Phillips, D. Reiss, B. P. Schmidt, R. A. Schommer, R. C. Smith, J. Spyromilio, C. Stubbs, N. B. Suntzeff, and J. Tonry. “Observational Evidence from Supernovae for an Accelerating Universe and a Cosmological Constant”. In: *The Astronomical Journal* 116, p. 1009, 1998, p. 1009. (See p. 12)
- [8] V. Trimble. “Existence and Nature of Dark Matter in the Universe”. In: *Annual Review of Astronomy and Astrophysics* 25, pp. 425–472, 1987, pp. 425–472. (See p. 12)

- [9] R. B. Laughlin and D. Pines. "The Theory of Everything". In: *Proceedings of the National Academy of Sciences* 97, pp. 28–31, 2000, pp. 28–31. (See p. 12)
- [10] Y. Matsumoto, N. Fukumitsu, H. Ishikawa, K. Nakai, and H. Sakurai. "A Critical Review of Radiation Therapy: From Particle Beam Therapy (Proton, Carbon, and BNCT) to Beyond". In: *Journal of Personalized Medicine* 11, 2021. (See p. 12)
- [11] O. Jacobson, D. O. Kiesewetter, and X. Chen. "Fluorine-18 Radiochemistry, Labeling Strategies and Synthetic Routes". In: *Bioconjugate Chemistry* 26, pp. 1–18, 2015, pp. 1–18. (See p. 12)
- [12] A Mousavi Khaneghah, M Hashemi Moosavi, C. A. F. Oliveira, F Vanin, and A. S. Sant'Ana. "Electron beam irradiation to reduce the mycotoxin and microbial contaminations of cereal-based products: An overview". In: *Food and Chemical Toxicology* 143, p. 111557, 2020, p. 111557. (See p. 12)
- [13] J. S. Williams. "Ion implantation of semiconductors". In: *Materials Science and Engineering: A* 253, pp. 8–15, 1998, pp. 8–15. (See p. 12)
- [14] D. J. Rej, H. A. Davis, J. C. Olson, G. E. Remnev, A. N. Zakoutaev, V. A. Ryzhkov, V. K. Struts, I. F. Isakov, V. A. Shulov, N. A. Nochevnaya, R. W. Stinnett, E. L. Neau, K. Yatsui, and W. Jiang. "Materials processing with intense pulsed ion beams". In: *Journal of Vacuum Science Technology A* 15, pp. 1089–1097, 1997, pp. 1089–1097. (See p. 12)
- [15] E. Esarey, C. B. Schroeder, and W. P. Leemans. "Physics of laser-driven plasma-based electron accelerators". In: *Reviews of Modern Physics* 81, pp. 1229–1285, 2009, pp. 1229–1285. (See p. 13)
- [16] J. D. Cockcroft, E. T. S. Walton, and E. Rutherford. "Experiments with high velocity positive ions.—(I) Further developments in the method of obtaining high velocity positive ions". In: *Proc. R. Soc. A: Math. Phys. Eng. Sci.* 136, pp. 619–630, 1932, pp. 619–630. (See p. 13)
- [17] Y. Jongen. "Review of compact commercial accelerator products and applications". In: *Proceedings of the 1997 Particle Accelerator Conference*. Vol. 3. 1997. 3770–3774 vol.3. (See p. 13)
- [18] E. D. Courant, M. S. Livingston, and H. S. Snyder. "The Strong-Focusing Synchrotron—A New High Energy Accelerator". In: *Physical Review* 88, pp. 1190–1196, 1952, pp. 1190–1196. (See p. 14)
- [19] E. M. McMillan. "The Synchrotron—A Proposed High Energy Particle Accelerator". In: *Physical Review* 68, pp. 143–144, 1945, pp. 143–144. (See p. 14)
- [20] O. Brüning, H. Burkhardt, and S. Myers. "The large hadron collider". In: *Progress in Particle and Nuclear Physics* 67, pp. 705–734, 2012, pp. 705–734. (See p. 14)
- [21] K. Hübner. "Designing and building LEP". In: *Physics Reports* 403-404, pp. 177–188, 2004, pp. 177–188. (See p. 14)
- [22] H Padamsee. "50 years of success for SRF accelerators—a review". In: *Superconductor Science and Technology* 30, p. 053003, 2017, p. 053003. (See p. 14)

[23] P. Bambade, T. Barklow, T. Behnke, M. Berggren, J. Brau, P. Burrows, D. Denisov, A. Faus-Golfe, B. Foster, K. Fujii, J. Fuster, F. Gaede, P. Grannis, C. Grojean, A. Hutton, B. List, J. List, S. Michizono, A. Miyamoto, O. Napoly, M. Peskin, R. Poeschl, F. Simon, J. Strube, J. Tian, M. Titov, M. Vos, A. White, G. Wilson, *et al.* “The International Linear Collider: A Global Project”. In: *arXiv e-prints*, arXiv:1903.01629, 2019, arXiv:1903.01629. (See p. 15)

[24] A. Aryshev, T. Behnke, M. Berggren, J. Brau, N. Craig, A. Freitas, F. Gaede, S. Gessner, S. Gori, C. Grojean, S. Heinemeyer, D. Jeans, K. Kruger, B. List, J. List, Z. Liu, S. Michizono, D. W. Miller, I. Moult, H. Murayama, T. Nakada, E. Nanni, M. Nojiri, H. Padamsee, M. Perelstein, M. E. Peskin, R. Poeschl, S. Posen, A. Robson, *et al.* “The International Linear Collider: Report to Snowmass 2021”. In: arXiv:2203.07622, 2022, arXiv:2203.07622. (See p. 15)

[25] M. Chamizo Llatas, S. Dasu, U. Heintz, E. A. Nanni, J. Power, and S. Wagner. “Report of the Snowmass 2021 e^+e^- -Collider Forum”. In: *arXiv e-prints*, arXiv:2209.03472, 2022, arXiv:2209.03472. (See pp. 15, 19)

[26] V. A. Veksler. “Principles of acceleration of charged particles”. In: *The Soviet Journal of Atomic Energy* 1, pp. 77–83, 1956, pp. 77–83. (See p. 15)

[27] T. Tajima and J. M. Dawson. “Laser Electron Accelerator”. In: *Physical Review Letters* 43, pp. 267–270, 1979, pp. 267–270. (See pp. 15, 35, 82)

[28] C. E. Clayton, C. Joshi, C. Darrow, and D. Umstadter. “Relativistic Plasma-Wave Excitation by Collinear Optical Mixing”. In: *Physical Review Letters* 54, pp. 2343–2346, 1985, pp. 2343–2346. (See p. 16)

[29] C. E. Clayton, K. A. Marsh, A. Dyson, M. Everett, A. Lal, W. P. Leemans, R. Williams, and C. Joshi. “Ultrahigh-gradient acceleration of injected electrons by laser-excited relativistic electron plasma waves”. In: *Physical Review Letters* 70, pp. 37–40, 1993, pp. 37–40. (See p. 16)

[30] F. Amiranoff, D. Bernard, B. Cros, F. Jacquet, G. Matthieussent, P. Miné, P. Mora, J. Morillo, F. Moulin, A. E. Specka, and C. Stenz. “Electron Acceleration in Nd-Laser Plasma Beat-Wave Experiments”. In: *Physical Review Letters* 74, pp. 5220–5223, 1995, pp. 5220–5223. (See p. 16)

[31] J. Krall, A. Ting, E. Esarey, and P. Sprangle. “Enhanced acceleration in a self-modulated-laser wake-field accelerator”. In: *Physical Review E* 48, pp. 2157–2161, 1993, pp. 2157–2161. (See p. 16)

[32] K. Nakajima, D. Fisher, T. Kawakubo, H. Nakanishi, A. Ogata, Y. Kato, Y. Kitagawa, R. Kodama, K. Mima, H. Shiraga, K. Suzuki, K. Yamakawa, T. Zhang, Y. Sakawa, T. Shoji, Y. Nishida, N. Yugami, M. Downter, and T. Tajima. “Observation of Ultrahigh Gradient Electron Acceleration by a Self-Modulated Intense Short Laser Pulse”. In: *Physical Review Letters* 75, pp. 984–984, 1995, pp. 984–984. (See p. 16)

[33] A. Modena, Z. Najmudin, A. E. Dangor, C. E. Clayton, K. A. Marsh, C. Joshi, V. Malka, C. B. Darrow, C. Danson, D. Neely, and F. N. Walsh. “Electron acceleration from the breaking of relativistic plasma waves”. In: *Nature* 377, pp. 606–608, 1995, pp. 606–608. (See p. 16)

[34] C. I. Moore, A. Ting, K. Krushelnick, E. Esarey, R. F. Hubbard, B. Hafizi, H. R. Burris, C. Manka, and P. Sprangle. “Electron Trapping in Self-Modulated Laser Wakefields by Raman Backscatter”. In: *Physical Review Letters* 79, pp. 3909–3912, 1997, pp. 3909–3912. (See p. 16)

[35] W. P. Leemans, D. Rodgers, P. E. Catravas, C. G. R. Geddes, G. Fubiani, E. Esarey, B. A. Shadwick, R. Donahue, and A. Smith. “Gamma-neutron activation experiments using laser wakefield accelerators”. In: *Physics of Plasmas* 8, pp. 2510–2516, 2001, pp. 2510–2516. (See p. 16)

[36] V. Malka, J. Faure, J. R. Marquès, F. Amiranoff, J. P. Rousseau, S. Ranc, J. P. Chambaret, Z. Najmudin, B. Walton, P. Mora, and A. Solodov. “Characterization of electron beams produced by ultrashort (30 fs) laser pulses”. In: *Physics of Plasmas* 8, pp. 2605–2608, 2001, pp. 2605–2608. (See pp. 16, 136)

[37] D. Gordon, K. C. Tzeng, C. E. Clayton, A. E. Dangor, V. Malka, K. A. Marsh, A. Modena, W. B. Mori, P. Muggli, Z. Najmudin, D. Neely, C. Danson, and C. Joshi. “Observation of Electron Energies Beyond the Linear Dephasing Limit from a Laser-Excited Relativistic Plasma Wave”. In: *Physical Review Letters* 80, pp. 2133–2136, 1998, pp. 2133–2136. (See p. 16)

[38] S. P. D. Mangles, C. D. Murphy, Z. Najmudin, A. G. R. Thomas, J. L. Collier, A. E. Dangor, E. J. Divall, P. S. Foster, J. G. Gallacher, C. J. Hooker, D. A. Jaroszynski, A. J. Langley, W. B. Mori, P. A. Norreys, F. S. Tsung, R. Viskup, B. R. Walton, and K. Krushelnick. “Monoenergetic beams of relativistic electrons from intense laser–plasma interactions”. In: *Nature* 431, pp. 535–538, 2004, pp. 535–538. (See p. 16)

[39] C. G. R. Geddes, C. Toth, J. van Tilborg, E. Esarey, C. B. Schroeder, D. Bruhwiler, C. Nieter, J. Cary, and W. P. Leemans. “High-quality electron beams from a laser wakefield accelerator using plasma-channel guiding”. In: *Nature* 431, pp. 538–541, 2004, pp. 538–541. (See p. 16)

[40] J. Faure, Y. Glinec, A. Pukhov, S. Kiselev, S. Gordienko, E. Lefebvre, J. P. Rousseau, F. Burgy, and V. Malka. “A laser–plasma accelerator producing monoenergetic electron beams”. In: *Nature* 431, pp. 541–544, 2004, pp. 541–544. (See p. 16)

[41] D. Strickland and G. Mourou. “Compression of amplified chirped optical pulses”. In: *Optics Communications* 56, pp. 219–221, 1985, pp. 219–221. (See p. 17)

[42] P. Maine, D. Strickland, P. Bado, M. Pessot, and G. Mourou. “Generation of ultrahigh peak power pulses by chirped pulse amplification”. In: *IEEE Journal of Quantum Electronics* 24, pp. 398–403, 1988, pp. 398–403. (See p. 17)

[43] M. D. Perry and G. Mourou. “Terawatt to Petawatt Subpicosecond Lasers”. In: *Science* 264, pp. 917–924, 1994, pp. 917–924. (See p. 17)

[44] W. Lu, M. Tzoufras, C. Joshi, F. S. Tsung, W. B. Mori, J. Vieira, R. A. Fonseca, and L. O. Silva. “Generating multi-GeV electron bunches using single stage laser wakefield acceleration in a 3D nonlinear regime”. In: *Physical Review Special Topics - Accelerators and Beams* 10, p. 061301, 2007, p. 061301. (See pp. 17, 20, 39, 47, 48, 51, 82, 118, 120)

[45] A. B. Borisov, A. V. Borovskiy, V. V. Korobkin, A. M. Prokhorov, O. B. Shiryaev, X. M. Shi, T. S. Luk, A. McPherson, J. C. Solem, K. Boyer, and C. K. Rhodes. “Observation of relativistic and charge-displacement self-channeling of intense subpicosecond ultraviolet (248 nm) radiation in plasmas”. In: *Physical Review Letters* 68, pp. 2309–2312, 1992, pp. 2309–2312. (See p. 17)

[46] L. M. Chen, H. Kotaki, K. Nakajima, J. Koga, S. V. Bulanov, T. Tajima, Y. Q. Gu, H. S. Peng, X. X. Wang, T. S. Wen, H. J. Liu, C. Y. Jiao, C. G. Zhang, X. J. Huang, Y. Guo, K. N. Zhou, J. F. Hua, W. M. An, C. X. Tang, and Y. Z. Lin. “Self-guiding of 100TW femtosecond laser pulses in centimeter-scale underdense plasma”. In: *Physics of Plasmas* 14, p. 040703, 2007, p. 040703. (See p. 17)

[47] A. Zigler, Y. Ehrlich, C. Cohen, J. Krall, and P. Sprangle. “Optical guiding of high-intensity laser pulses in a long plasma channel formed by a slow capillary discharge”. In: *J. Opt. Soc. Am. B* 13, pp. 68–71, 1996, pp. 68–71. (See p. 17)

[48] Y. Ehrlich, C. Cohen, D. Kaganovich, A. Zigler, R. F. Hubbard, P. Sprangle, and E. Esarey. “Guiding and damping of high-intensity laser pulses in long plasma channels”. In: *J. Opt. Soc. Am. B* 15, pp. 2416–2423, 1998, pp. 2416–2423. (See p. 17)

[49] r. Durfee C. G. and H. M. Milchberg. “Light pipe for high intensity laser pulses”. In: *Phys Rev Lett* 71, pp. 2409–2412, 1993, pp. 2409–2412. (See p. 17)

[50] H. M. Milchberg, T. R. Clark, C. G. Durfee, T. M. Antonsen, and P. Mora. “Development and applications of a plasma waveguide for intense laser pulses”. In: *Physics of Plasmas* 3, pp. 2149–2155, 1996, pp. 2149–2155. (See p. 17)

[51] R. J. Shalloo, C. Arran, L. Corner, J. Holloway, J. Jonnerby, R. Walczak, H. M. Milchberg, and S. M. Hooker. “Hydrodynamic optical-field-ionized plasma channels”. In: *Physical Review E* 97, p. 053203, 2018, p. 053203. (See p. 17)

[52] A. Picksley, A. Alejo, R. J. Shalloo, C. Arran, A. von Boetticher, L. Corner, J. A. Holloway, J. Jonnerby, O. Jakobsson, C. Thornton, R. Walczak, and S. M. Hooker. “Meter-scale conditioned hydrodynamic optical-field-ionized plasma channels”. In: *Physical Review E* 102, p. 053201, 2020, p. 053201. (See p. 17)

[53] B. Miao, L. Feder, J. Shrock, A. Goffin, and H. Milchberg. “Optical Guiding in Meter-Scale Plasma Waveguides”. In: *Physical Review Letters* 125, p. 074801, 2020, p. 074801. (See p. 17)

[54] B. Miao, J. Shrock, L. Feder, R. Hollinger, J. Morrison, R. Nedbailo, A. Picksley, H. Song, S. Wang, J. Rocca, and H. Milchberg. “Multi-GeV Electron Bunches from an All-Optical Laser Wakefield Accelerator”. In: *Physical Review X* 12, p. 031038, 2022, p. 031038. (See p. 17)

[55] A. Gonsalves, K. Nakamura, J. Daniels, C. Benedetti, C. Pieronek, T. de Raadt, S. Steinke, J. Bin, S. Bulanov, J. van Tilborg, C. Geddes, C. Schroeder, C. Tóth, E. Esarey, K. Swanson, L. Fan-Chiang, G. Bagdasarov, N. Bobrova, V. Gasilov, G. Korn, P. Sasorov, and W. Leemans. “Petawatt Laser Guiding and Electron Beam Acceleration to 8 GeV in a Laser-Heated Capillary Discharge Waveguide”. In: *Physical Review Letters* 122, p. 084801, 2019, p. 084801. (See pp. 17, 20)

[56] I. Blumenfeld, C. E. Clayton, F.-J. Decker, M. J. Hogan, C. Huang, R. Ischebeck, R. Iverson, C. Joshi, T. Katsouleas, N. Kirby, W. Lu, K. A. Marsh, W. B. Mori, P. Muggli, E. Oz, R. H. Siemann, D. Walz, and M. Zhou. “Energy doubling of 42GeV electrons in a metre-scale plasma wakefield accelerator”. In: *Nature* 445, pp. 741–744, 2007, pp. 741–744. (See p. 18)

[57] R. Assmann, R. Bingham, T. Bohl, C. Bracco, B. Buttenschön, A. Butterworth, A. Caldwell, S. Chattopadhyay, S. Cipiccia, E. Feldbaumer, R. A. Fonseca, B. Goddard, M. Gross, O. Grulke, E. Gschwendtner, J. Holloway, C. Huang, D. Jaroszynski, S. Jolly, P. Kempkes, N. Lopes, K. Lotov, J. Machacek, S. R. Mandry, J. W. McKenzie, M. Meddahi, B. L. Militsyn, N. Moschuering, P. Muggli, *et al.* “Proton-driven plasma wakefield acceleration: a path to the future of high-energy particle physics”. In: *Plasma Physics and Controlled Fusion* 56, p. 084013, 2014, p. 084013. (See p. 18)

[58] E. Gschwendtner, E. Adli, L. Amorim, R. Apsimon, R. Assmann, A. M. Bachmann, F. Batsch, J. Bauche, V. K. Berglyd Olsen, M. Bernardini, R. Bingham, B. Biskup, T. Bohl, C. Bracco, P. N. Burrows, G. Burt, B. Buttenschön, A. Butterworth, A. Caldwell, M. Cascella, E. Chevallay, S. Cipiccia, H. Damerau, L. Deacon, P. Dirksen, S. Doeberl, U. Dorda, J. Farmer, V. Fedosseev, *et al.* “AWAKE, The Advanced Proton Driven Plasma Wakefield Acceleration Experiment at CERN”. In: *Nucl. Instrum. Methods Phys. Res. A: Accel. Spectrom. Detect. Assoc. Equip.* 829, pp. 76–82, 2016, pp. 76–82. (See p. 18)

[59] S. Lee, T. Katsouleas, P. Muggli, W. B. Mori, C. Joshi, R. Hemker, E. S. Dodd, C. E. Clayton, K. A. Marsh, B. Blue, S. Wang, R. Assmann, F. J. Decker, M. Hogan, R. Iverson, and D. Walz. “Energy doubler for a linear collider”. In: *Physical Review Special Topics - Accelerators and Beams* 5, p. 011001, 2002, p. 011001. (See p. 18)

[60] S. Wang, C. E. Clayton, B. E. Blue, E. S. Dodd, K. A. Marsh, W. B. Mori, C. Joshi, S. Lee, P. Muggli, T. Katsouleas, F. J. Decker, M. J. Hogan, R. H. Iverson, P. Raimondi, D. Walz, R. Siemann, and R. Assmann. “X-Ray Emission from Betatron Motion in a Plasma Wiggler”. In: *Physical Review Letters* 88, p. 135004, 2002, p. 135004. (See p. 18)

[61] A. Rousse, K. T. Phuoc, R. Shah, A. Pukhov, E. Lefebvre, V. Malka, S. Kiselev, F. Burgy, J.-P. Rousseau, D. Umstadter, and D. Hulin. “Production of a keV X-Ray Beam from Synchrotron Radiation in Relativistic Laser-Plasma Interaction”. In: *Physical Review Letters* 93, p. 135005, 2004, p. 135005. (See p. 18)

[62] S. Fourmaux, S. Corde, K. Ta Phuoc, P. M. Leguay, S. Payeur, P. Lassonde, S. Gnedyuk, G. Lebrun, C. Fourment, V. Malka, S. Sebbar, A. Rousse, and J. C. Kieffer. “Demonstration of the synchrotron-type spectrum of laser-produced Betatron radiation”. In: *New Journal of Physics* 13, p. 033017, 2011, p. 033017. (See p. 18)

[63] M. Schnell, A. Sävert, I. Uschmann, O. Jansen, M. C. Kaluza, and C. Spielmann. “Characterization and application of hard x-ray betatron radiation generated by relativistic electrons from a laser-wakefield accelerator”. In: *Journal of Plasma Physics* 81, p. 475810401, 2015, p. 475810401. (See p. 18)

[64] M. Heigoldt, A. Popp, K. Khrennikov, J. Wenz, S. Chou, S. Karsch, S. Bajlekov, S. Hooker, and B. Schmidt. “Temporal evolution of longitudinal bunch profile in a laser wakefield accelerator”. In: *Physical Review Special Topics - Accelerators and Beams* 18, p. 121302, 2015, p. 121302. (See pp. 18, 93, 128, 172)

[65] K. T. Phuoc, R. Fitour, A. Tafzi, T. Garl, N. Artemiev, R. Shah, F. Albert, D. Boschetto, A. Rousse, D. E. Kim, A. Pukhov, V. Seredov, and I. Kostyukov. “Demonstration of the ultrafast nature of laser produced betatron radiation”. In: *Physics of Plasmas* 14, p. 080701, 2007, p. 080701. (See p. 18)

[66] J. C. Wood, D. J. Chapman, K. Poder, N. C. Lopes, M. E. Rutherford, T. G. White, F. Albert, K. T. Behm, N. Booth, J. S. J. Bryant, P. S. Foster, S. Glenzer, E. Hill, K. Krushelnick, Z. Najmudin, B. B. Pollock, S. Rose, W. Schumaker, R. H. H. Scott, M. Sherlock, A. G. R. Thomas, Z. Zhao, D. E. Eakins, and S. P. D. Mangles. “Ultrafast Imaging of Laser Driven Shock Waves using Betatron X-rays from a Laser Wakefield Accelerator”. In: *Scientific Reports* 8, p. 11010, 2018, p. 11010. (See p. 18)

[67] K. T. Phuoc, S. Corde, R. Shah, F. Albert, R. Fitour, J.-P. Rousseau, F. Burgy, B. Mercier, and A. Rousse. “Imaging Electron Trajectories in a Laser-Wakefield Cavity Using Betatron X-Ray Radiation”. In: *Physical Review Letters* 97, p. 225002, 2006, p. 225002. (See p. 18)

[68] F. Albert, R. Shah, K. T. Phuoc, R. Fitour, F. Burgy, J.-P. Rousseau, A. Tafzi, D. Douillet, T. Lefrou, and A. Rousse. “Betatron oscillations of electrons accelerated in laser wakefields characterized by spectral x-ray analysis”. In: *Physical Review E* 77, p. 056402, 2008, p. 056402. (See p. 18)

[69] S. Fourmaux, S. Corde, K. T. Phuoc, P. Lassonde, G. Lebrun, S. Payeur, F. Martin, S. Sebbar, V. Malka, A. Rousse, and J. C. Kieffer. “Single shot phase contrast imaging using laser-produced Betatron x-ray beams”. In: *Optics Letters* 36, pp. 2426–2428, 2011, pp. 2426–2428. (See p. 18)

[70] J. Cole, K. Behm, E. Gerstmayr, T. Blackburn, J. Wood, C. Baird, M. Duff, C. Harvey, A. Ilderton, A. Joglekar, K. Krushelnick, S. Kuschel, M. Marklund, P. McKenna, C. Murphy, K. Poder, C. Ridgers, G. Samarin, G. Sarri, D. Symes, A. Thomas, J. Warwick, M. Zepf, Z. Najmudin, and S. Mangles. “Experimental Evidence of Radiation Reaction in the Collision of a High-Intensity Laser Pulse with a Laser-Wakefield Accelerated Electron Beam”. In: *Physical Review X* 8, p. 011020, 2018, p. 011020. (See pp. 18, 19)

[71] M. Z. Mo, Z. Chen, S. Fourmaux, A. Saraf, K. Otani, J. C. Kieffer, Y. Y. Tsui, A. Ng, and R. Fedosejevs. “Laser wakefield generated X-ray probe for femtosecond time-resolved measurements of ionization states of warm dense aluminum”. In: *Review of Scientific Instruments* 84, p. 123106, 2013, p. 123106. (See p. 18)

[72] B. Kettle, E. Gerstmayr, M. Streeter, F. Albert, R. Baggott, N. Bourgeois, J. Cole, S. Dann, K. Falk, I. Gallardo González, A. Hussein, N. Lemos, N. Lopes, O. Lundh, Y. Ma, S. Rose, C. Spindloe, D. Symes, M. Šmíd, A. Thomas, R. Watt, and S. Mangles. “Single-Shot Multi-keV X-Ray Absorption Spectroscopy Using an Ultrashort Laser-Wakefield Accelerator Source”. In: *Physical Review Letters* 123, p. 254801, 2019, p. 254801. (See p. 18)

[73] N. Mounet. *European Strategy for Particle Physics - Accelerator RD Roadmap*. Vol. 1. 2022. (See p. 19)

[74] M. Wing. “Particle physics experiments based on the AWAKE acceleration scheme”. In: *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 377, p. 20180185, 2019, p. 20180185. (See p. 19)

[75] P. C. Bhat and R. Rubinstein. “The International Committee for Future Accelerators (ICFA): History and the Future”. In: *Reviews of Accelerator Science and Technology* 10, pp. 303–309, 2019, pp. 303–309. (See p. 19)

[76] E. Adli. “Towards a PWFA linear collider — opportunities and challenges”. In: *Journal of Instrumentation* 17, T05006, 2022, T05006. (See p. 19)

[77] W. Herr and B. Muratori. “Concept of luminosity”. In: *Cern Document Server*, 2006. (See p. 19)

[78] F. Gianotti, M. L. Mangano, T. Virdee, S. Abdullin, G. Azuelos, A. Ball, D. Barberis, A. Belyaev, P. Bloch, M. Bosman, L. Casagrande, D. Cavalli, P. Chumney, S. Cittolin, S. Dasu, A. De Roeck, N. Ellis, P. Farthouat, D. Fournier, J. B. Hansen, I. Hinchliffe, M. Hohlfeld, M. Huhtinen, K. Jakobs, C. Joram, F. Mazzucato, G. Mikenberg, A. Miagkov, M. Moretti, *et al.* “Physics potential and experimental challenges of the LHC luminosity upgrade”. In: *The European Physical Journal C - Particles and Fields* 39, pp. 293–333, 2005, pp. 293–333. (See p. 19)

[79] J. Gillies. *Luminosity? Why don't we just say collision rate?* Quantum Diaries, 2011. 2011 URL: <http://cds.cern.ch/record/1997001> (visited on Nov. 15, 2022) (see p. 19)

[80] C. B. Schroeder, E. Esarey, C. G. R. Geddes, C. Benedetti, and W. P. Leemans. “Physics considerations for laser-plasma linear colliders”. In: *Physical Review Special Topics - Accelerators and Beams* 13, p. 101301, 2010, p. 101301. (See p. 20)

[81] C. M. S. Sears, A. Buck, K. Schmid, J. Mikhailova, F. Krausz, and L. Veisz. “Emittance and divergence of laser wakefield accelerated electrons”. In: *Physical Review Special Topics - Accelerators and Beams* 13, p. 092803, 2010, p. 092803. (See pp. 20, 57, 79, 156)

[82] S. Kneip, C. McGuffey, J. L. Martins, M. S. Bloom, V. Chvykov, F. Dollar, R. Fonseca, S. Jolly, G. Kalintchenko, K. Krushelnick, A. Maksimchuk, S. P. D. Mangles, Z. Najmudin, C. A. J. Palmer, K. T. Phuoc, W. Schumaker, L. O. Silva, J. Vieira, V. Yanovsky, and A. G. R. Thomas. “Characterization of transverse beam emittance of electrons from a laser-plasma wakefield accelerator in the bubble regime using betatron x-ray radiation”. In: *Physical Review Special Topics - Accelerators and Beams* 15, p. 021302, 2012, p. 021302. (See pp. 20, 54, 156)

[83] D. H. Dowell. “Sources of Emittance in RF Photocathode Injectors: Intrinsic emittance, space charge forces due to non-uniformities, RF and solenoid effects”. In: *arXiv e-prints*, arXiv:1610.01242, 2016, arXiv:1610.01242. (See p. 20)

[84] Z. Qin, C. Yu, W. Wang, J. Liu, W. Li, R. Qi, Z. Zhang, J. Liu, M. Fang, K. Feng, Y. Wu, L. Ke, Y. Chen, Y. Xu, Y. Leng, C. Wang, R. Li, and Z. Xu. “Ultralow-emittance measurement of high-quality electron beams from a laser wakefield accelerator”. In: *Physics of Plasmas* 25, p. 023106, 2018, p. 023106. (See pp. 20, 156)

[85] A. Cianchi, M. P. Anania, M. Bellaveglia, M. Castellano, E. Chiadroni, M. Ferrario, G. Gatti, B. Marchetti, A. Mostacci, R. Pompili, C. Ronsivalle, A. R. Rossi, and L. Serafini. “Challenges in plasma and laser wakefield accelerated beams diagnostic”. In: *Nucl. Instrum. Methods Phys. Res. A: Accel. Spectrom. Detect. Assoc. Equip.* 720, pp. 153–156, 2013, pp. 153–156. (See pp. 20, 157)

[86] P. Poole, S. Trendafilov, G. Shvets, D. Smith, and E. Chowdhury. “Femtosecond laser damage threshold of pulse compression gratings for petawatt scale laser systems”. In: *Optics Express* 21, pp. 26341–26351, 2013, pp. 26341–26351. (See p. 20)

[87] C. N. Danson, C. Haefner, J. Bromage, T. Butcher, J.-C. F. Chanteloup, E. A. Chowdhury, A. Galvanauskas, L. A. Gizzi, J. Hein, D. I. Hillier, N. W. Hopps, Y. Kato, E. A. Khazanov, R. Kodama, G. Korn, R. Li, Y. Li, J. Limpert, J. Ma, C. H. Nam, D. Neely, D. Papadopoulos, R. R. Penman, L. Qian, J. J. Rocca, A. A. Shaykin, C. W. Siders, C. Spindloe, S. Szatmári, *et al.* “Petawatt and exawatt class lasers worldwide”. In: *High Power Laser Science and Engineering* 7, e54, 2019, e54. (See p. 20)

[88] S. Steinke, J. van Tilborg, C. Benedetti, C. G. Geddes, C. B. Schroeder, J. Daniels, K. K. Swanson, A. J. Gonsalves, K. Nakamura, N. H. Matlis, B. H. Shaw, E. Esarey, and W. P. Leemans. “Multistage coupling of independent laser-plasma accelerators”. In: *Nature* 530, pp. 190–3, 2016, pp. 190–3. (See pp. 21, 160, 188)

[89] J. van Tilborg, S. Steinke, C. Geddes, N. Matlis, B. Shaw, A. Gonsalves, J. Huijts, K. Nakamura, J. Daniels, C. Schroeder, C. Benedetti, E. Esarey, S. Bulanov, N. Bobrova, P. Sasorov, and W. Leemans. “Active Plasma Lensing for Relativistic Laser-Plasma-Accelerated Electron Beams”. In: *Physical Review Letters* 115, p. 184802, 2015, p. 184802. (See pp. 21, 22)

[90] F. Albert, M. E. Couplie, A. Debus, M. C. Downer, J. Faure, A. Flacco, L. A. Gizzi, T. Grismayer, A. Huebl, C. Joshi, M. Labat, W. P. Leemans, A. R. Maier, S. P. D. Mangles, P. Mason, F. Mathieu, P. Muggli, M. Nishiuchi, J. Osterhoff, P. P. Rajeev, U. Schramm, J. Schreiber, A. G. R. Thomas, J.-L. Vay, M. Vranic, and K. Zeil. “2020 roadmap on plasma accelerators”. In: *New Journal of Physics* 23, p. 031101, 2021, p. 031101. (See p. 21)

[91] C. Thaury, F. Quéré, J. P. Geindre, A. Levy, T. Ceccotti, P. Monot, M. Bougeard, F. Réau, P. d’Oliveira, P. Audebert, R. Marjoribanks, and P. Martin. “Plasma mirrors for ultrahigh-intensity optics”. In: *Nature Physics* 3, pp. 424–429, 2007, pp. 424–429. (See pp. 21, 162)

[92] G. Raj, O. Kononenko, M. F. Gilljohann, A. Doche, X. Davoine, C. Caizergues, Y. Y. Chang, J. P. Couperus Cabadăg, A. Debus, H. Ding, M. Förster, J. P. Goddet, T. Heinemann, T. Kluge, T. Kurz, R. Pausch, P. Rousseau, P. San Miguel Claveria, S. Schöbel, A. Siciak, K. Steiniger, A. Tafzi, S. Yu, B. Hidding, A. Martinez de la Ossa, A. Irman, S. Karsch, A. Döpp, U. Schramm, *et al.* “Probing ultrafast magnetic-field generation by current filamentation instability in femtosecond relativistic laser-matter interactions”. In: *Physical Review Research* 2, p. 023123, 2020, p. 023123. (See pp. 21, 171, 174, 196)

[93] J. Luo, M. Chen, W. Wu, S. Weng, Z. Sheng, C. Schroeder, D. Jaroszynski, E. Esarey, W. Leemans, W. Mori, and J. Zhang. “Multistage Coupling of Laser-Wakefield Accelerators with Curved Plasma Channels”. In: *Physical Review Letters* 120, p. 154801, 2018, p. 154801. (See p. 21)

[94] Y. Ma, D. Seipt, K. Krushelnick, and A. G. R. Thomas. “Generation of straight and curved hollow plasma channels by laser-generated nonlinear wakefields and studies of ultra-intense laser pulse guiding”. In: *Physics of Plasmas* 28, p. 063104, 2021, p. 063104. (See p. 21)

[95] R. Assmann and K. Yokoya. “Transverse beam dynamics in plasma-based linacs”. In: *Nucl. Instrum. Methods Phys. Res. A: Accel. Spectrom. Detect. Assoc. Equip.* 410, pp. 544–548, 1998, pp. 544–548. (See p. 21)

[96] S. Cheshkov, T. Tajima, W. Horton, and K. Yokoya. “Particle dynamics in multistage wakefield collider”. In: *Physical Review Special Topics - Accelerators and Beams* 3, p. 071301, 2000, p. 071301. (See p. 21)

[97] C. Chiu, S. Cheshkov, and T. Tajima. “High energy laser-wakefield collider with synchronous acceleration”. In: *Physical Review Special Topics - Accelerators and Beams* 3, p. 101301, 2000, p. 101301. (See p. 21)

[98] F. Wu, Z. Zhang, X. Yang, J. Hu, P. Ji, J. Gui, C. Wang, J. Chen, Y. Peng, X. Liu, Y. Liu, X. Lu, Y. Xu, Y. Leng, R. Li, and Z. Xu. “Performance improvement of a 200TW/1Hz Ti:sapphire laser for laser wakefield electron accelerator”. In: *Optics Laser Technology* 131, p. 106453, 2020, p. 106453. (See p. 21)

[99] A. Zingale, N. Czapla, D. Nasir, S. Barber, J. Bin, A. Gonsalves, F. Isono, J. van Tilborg, S. Steinke, K. Nakamura, G. Cochran, J. Purcell, W. Leemans, C. Geddes, C. Schroeder, E. Esarey, and D. Schumacher. “Emittance preserving thin film plasma mirrors for GeV scale laser plasma accelerators”. In: *Physical Review Accelerators and Beams* 24, p. 121301, 2021, p. 121301. (See p. 21)

[100] N. Kirby, M. Berry, I. Blumenfeld, M. J. Hogan, R. Ischebeck, and R. Siemann. “Emittance growth from Multiple Coulomb Scattering in a plasma wakefield accelerator”. In: *2007 IEEE Particle Accelerator Conference (PAC)*. 2007. Pp. 3097–3099. (See p. 21)

[101] C. B. Schroeder, E. Esarey, C. Benedetti, and W. P. Leemans. “Control of focusing forces and emittances in plasma-based accelerators using near-hollow plasma channels”. In: *Physics of Plasmas* 20, p. 080701, 2013, p. 080701. (See p. 21)

[102] C. A. Lindstrøm and M. Thévenet. “Emittance preservation in advanced accelerators”. In: *Journal of Instrumentation* 17, P05016, 2022, P05016. (See p. 21)

[103] T. Mehrling, J. Grebenyuk, F. S. Tsung, K. Floettmann, and J. Osterhoff. “Transverse emittance growth in staged laser-wakefield acceleration”. In: *Physical Review Special Topics - Accelerators and Beams* 15, 2012. (See p. 21)

[104] M. Migliorati, A. Bacci, C. Benedetti, E. Chiadroni, M. Ferrario, A. Mostacci, L. Palumbo, A. R. Rossi, L. Serafini, and P. Antici. “Intrinsic normalized emittance growth in laser-driven electron accelerators”. In: *Physical Review Special Topics - Accelerators and Beams* 16, p. 011302, 2013, p. 011302. (See pp. 21, 183)

[105] I. Dornmair, K. Floettmann, and A. Maier. “Emittance conservation by tailored focusing profiles in a plasma accelerator”. In: *Physical Review Special Topics - Accelerators and Beams* 18, p. 041302, 2015, p. 041302. (See p. 22)

[106] K. Sjobak, E. Adli, R. Corsini, W. Farabolini, G. Boyle, C. Lindstrøm, M. Meisel, J. Osterhoff, J. H. Röckemann, L. Schaper, and A. Dyson. “Strong focusing gradient in a linear active plasma lens”. In: *Physical Review Accelerators and Beams* 24, p. 121306, 2021, p. 121306. (See p. 22)

[107] J. van Tilborg, S. Barber, H. E. Tsai, K. Swanson, S. Steinke, C. Geddes, A. Gonsalves, C. Schroeder, E. Esarey, S. Bulanov, N. Bobrova, P. Sasorov, and W. Leemans. “Nonuniform discharge currents in active plasma lenses”. In: *Physical Review Accelerators and Beams* 20, p. 032803, 2017, p. 032803. (See p. 22)

[108] C. Lindstrøm, E. Adli, J. Allen, W. An, C. Beekman, C. Clarke, C. Clayton, S. Corde, A. Doche, J. Frederico, S. Gessner, S. Green, M. Hogan, C. Joshi, M. Litos, W. Lu, K. Marsh, W. Mori, B. O’Shea, N. Vafaei-Najafabadi, and V. Yakimenko. “Measurement of Transverse Wakefields Induced by a Misaligned Positron Bunch in a Hollow Channel Plasma Accelerator”. In: *Physical Review Letters* 120, p. 124802, 2018, p. 124802. (See pp. 22, 23)

[109] G. Sarri, W. Schumaker, A. Di Piazza, K. Poder, J. M. Cole, M. Vargas, D. Doria, S. Kuschel, B. Dromey, G. Grittani, L. Gizzi, M. E. Dieckmann, A. Green, V. Chvykov, A. Maksimchuk, V. Yanovsky, Z. H. He, B. X. Hou, J. A. Nees, S. Kar, Z. Najmudin, A. G. R. Thomas, C. H. Keitel, K. Krushelnick, and M. Zepf. “Laser-driven generation of collimated ultra-relativistic positron beams”. In: *Plasma Physics and Controlled Fusion* 55, p. 124017, 2013, p. 124017. (See p. 22)

[110] G. Sarri, J. Warwick, W. Schumaker, K. Poder, J. Cole, D. Doria, T. Dzelzainis, K. Krushelnick, S. Kuschel, S. P D Mangles, Z. Najmudin, L. Romagnani, G. M Samarin, D. Symes, A. G R Thomas, M. Yeung, and M. Zepf. “Spectral and spatial characterisation of laser-driven positron beams”. In: *Plasma Physics and Controlled Fusion* 59, p. 014015, 2017, p. 014015. (See p. 22)

[111] M. J. V. Streeter, C. Colgan, N. Cavanagh, E. Los, A. F. Antoine, T. Audet, M. D. Balcazar, L. Calvin, J. Carderelli, H. Ahmed, B. Kettle, Y. Ma, S. P. D. Mangles, Z. Najmudin, P. P. Rajeev, D. R. Symes, A. G. R. Thomas, and G. Sarri. “Laser Generation of Near-GeV Low Emittance Positron Beams”. In: *arXiv e-prints*, arXiv:2205.13850, 2022, arXiv:2205.13850. (See p. 22)

[112] S. Corde, E. Adli, J. M. Allen, W. An, C. I. Clarke, C. E. Clayton, J. P. Delahaye, J. Frederico, S. Gessner, S. Z. Green, M. J. Hogan, C. Joshi, N. Lipkowitz, M. Litos, W. Lu, K. A. Marsh, W. B. Mori, M. Schmeltz, N. Vafaei-Najafabadi, D. Walz, V. Yakimenko, and G. Yocky. “Multi-gigaelectronvolt acceleration of positrons in a self-loaded plasma wakefield”. In: *Nature* 524, pp. 442–445, 2015, pp. 442–445. (See p. 22)

[113] C. S. Hue, G. J. Cao, I. A. Andriyash, A. Knetsch, M. J. Hogan, E. Adli, S. Gessner, and S. Corde. “Efficiency and beam quality for positron acceleration in loaded plasma wakefields”. In: *Physical Review Research* 3, p. 043063, 2021, p. 043063. (See p. 22)

[114] S. Diederichs, T. Mehrling, C. Benedetti, C. Schroeder, A. Knetsch, E. Esarey, and J. Osterhoff. “Positron transport and acceleration in beam-driven plasma wakefield accelerators using plasma columns”. In: *Physical Review Accelerators and Beams* 22, p. 081301, 2019, p. 081301. (See p. 22)

[115] S. Diederichs, C. Benedetti, E. Esarey, J. Osterhoff, and C. Schroeder. “High-quality positron acceleration in beam-driven plasma accelerators”. In: *Physical Review Accelerators and Beams* 23, p. 121301, 2020, p. 121301. (See pp. 22, 82)

[116] T. Wang, V. Khudik, and G. Shvets. “Positron Acceleration in an Elongated Bubble Regime”. In: *arXiv e-prints*, arXiv:2110.10290, 2021, arXiv:2110.10290. (See p. 22)

[117] S. Gessner, E. Adli, J. M. Allen, W. An, C. I. Clarke, C. E. Clayton, S. Corde, J. P. Delahaye, J. Frederico, S. Z. Green, C. Hast, M. J. Hogan, C. Joshi, C. A. Lindstrøm, N. Lipkowitz, M. Litos, W. Lu, K. A. Marsh, W. B. Mori, B. O’Shea, N. Vafaei-Najafabadi, D. Walz, V. Yakimenko, and G. Yocky. “Demonstration of a positron beam-driven hollow channel plasma wakefield accelerator”. In: *Nature Communications* 7, p. 11785, 2016, p. 11785. (See p. 23)

[118] T. Silva, L. Amorim, M. Downer, M. Hogan, V. Yakimenko, R. Zgadzaj, and J. Vieira. “Stable Positron Acceleration in Thin, Warm, Hollow Plasma Channels”. In: *Physical Review Letters* 127, p. 104801, 2021, p. 104801. (See pp. 23, 82)

[119] F. L. Pedrotti, L. M. Pedrotti, and L. S. Pedrotti. *Introduction to Optics*. 3rd ed. 1987. (See p. 25)

[120] W. S. C. Chang. *Principles of lasers and optics*. Cambridge: Cambridge University Press, 2005. (See p. 25)

[121] S. Feng and H. G. Winful. “Physical origin of the Gouy phase shift”. In: *Optics Letters* 26, pp. 485–487, 2001, pp. 485–487. (See p. 25)

[122] S. M. Sepke and D. P. Umstadter. “Exact analytical solution for the vector electromagnetic field of Gaussian, flattened Gaussian, and annular Gaussian laser modes”. In: *Optics Letters* 31, pp. 1447–1449, 2006, pp. 1447–1449. (See p. 25)

[123] M. Born and E. Wolf. *Principles of optics : electromagnetic theory of propagation, interference and diffraction of light*. 7th (expanded) ed. Cambridge ; Cambridge University Press, 1999. (See p. 26)

[124] S. P. D. Mangles, G. Genoud, S. Kneip, M. Burza, K. Cassou, B. Cros, N. P. Dover, C. Kamperidis, Z. Najmudin, A. Persson, J. Schreiber, F. Wojda, and C. G. Wahlström. “Controlling the spectrum of x-rays generated in a laser-plasma accelerator by tailoring the laser wavefront”. In: *Applied Physics Letters* 95, p. 181106, 2009, p. 181106. (See pp. 27, 131)

[125] A. Popp, J. Vieira, J. Osterhoff, Z. Major, R. Hörlein, M. Fuchs, R. Weingartner, T. P. Rowlands-Rees, M. Marti, R. A. Fonseca, S. F. Martins, L. O. Silva, S. M. Hooker, F. Krausz, F. Grüner, and S. Karsch. “All-Optical Steering of Laser-Wakefield-Accelerated Electron Beams”. In: *Physical Review Letters* 105, p. 215001, 2010, p. 215001. (See p. 27)

[126] B. Beaurepaire, A. Vernier, M. Bocoum, F. Böhle, A. Jullien, J. P. Rousseau, T. Lefrou, D. Douillet, G. Iaquaniello, R. Lopez-Martens, A. Lifschitz, and J. Faure. “Effect of the Laser Wave Front in a Laser-Plasma Accelerator”. In: *Physical Review X* 5, p. 031012, 2015, p. 031012. (See p. 27)

[127] J. M. Cole. “Diagnosis and Application of Laser Wakefield Accelerators”. Thesis. 2016. (See pp. 27, 136)

[128] Z. H. He, B. Hou, V. Lebailly, J. A. Nees, K. Krushelnick, and A. G. R. Thomas. “Coherent control of plasma dynamics”. In: *Nature Communications* 6, p. 7156, 2015, p. 7156. (See p. 27)

[129] S. Dann, C. Baird, N. Bourgeois, O. Chekhlov, S. Eardley, C. Gregory, J. N. Gruse, J. Hah, D. Hazra, S. Hawkes, C. Hooker, K. Krushelnick, S. Mangles, V. Marshall, C. Murphy, Z. Najmudin, J. Nees, J. Osterhoff, B. Parry, P. Pourmoussavi, S. Rahul, P. Rajeev, S. Rozario, J. Scott, R. Smith, E. Springate, Y. Tang, S. Tata, A. Thomas, *et al.* “Laser wakefield acceleration with active feedback at 5 Hz”. In: *Physical Review Accelerators and Beams* 22, p. 041303, 2019, p. 041303. (See p. 27)

[130] P. Mora and J. Thomas M. Antonsen. “Kinetic modeling of intense, short laser pulses propagating in tenuous plasmas”. In: *Physics of Plasmas* 4, pp. 217–229, 1997, pp. 217–229. (See p. 29)

[131] L. V. Keldysh. “Ionization in the field of a strong electromagnetic wave”. In: *Sov. Phys. JETP* 20, pp. 1307–1314, 1965, pp. 1307–1314. (See p. 31)

[132] M. V. Ammosov, N. B. Delone, and V. P. Krainov. “Tunnel ionization of complex atoms and of atomic ions in an alternating electromagnetic field”. In: *Sov. Phys. JETP* 64, p. 4, 1986, p. 4. (See pp. 31, 171)

[133] A. Perelomov, V. Popov, and M. Terent’ev. “Ionization of Atoms in an Alternating Electric Field”. In: *Sov. Phys. JETP* 23, 1966. (See p. 31)

[134] A. Perelomov, V. Popov, and M. Terent’ev. “Ionization of Atoms in an Alternating Electric Field: II”. In: *Sov. Phys. JETP* 24, 1967. (See p. 31)

[135] D. L. Bruhwiler, D. A. Dimitrov, J. R. Cary, E. Esarey, W. Leemans, and R. E. Giaccone. “Particle-in-cell simulations of tunneling ionization effects in plasma-based accelerators”. In: *Physics of Plasmas* 10, pp. 2022–2030, 2003, pp. 2022–2030. (See p. 31)

[136] P. Sprangle, E. Esarey, and A. Ting. “Nonlinear theory of intense laser-plasma interactions”. In: *Physical Review Letters* 64, pp. 2011–2014, 1990, pp. 2011–2014. (See pp. 36, 47, 48)

[137] E. Esarey and M. Piloff. “Trapping and acceleration in nonlinear plasma waves”. In: *Physics of Plasmas* 2, pp. 1432–1436, 1995, pp. 1432–1436. (See pp. 36, 50, 201)

[138] V. I. Berezhiani and I. G. Murusidze. “Relativistic wake-field generation by an intense laser pulse in a plasma”. In: *Physics Letters A* 148, pp. 338–340, 1990, pp. 338–340. (See p. 36)

[139] S. V. Bulanov, V. I. Kirsanov, and A. S. Sakharov. “Excitation of ultrarelativistic plasma waves by pulse of electromagnetic radiation”. In: *JETP Letters* 50, pp. 176–178, 1989, pp. 176–178. (See p. 36)

- [140] W. Lu, C. Huang, M. Zhou, M. Tzoufras, F. S. Tsung, W. B. Mori, and T. Katsouleas. “A nonlinear theory for multidimensional relativistic plasma wave wakefields”. In: *Physics of Plasmas* 13, p. 056709, 2006, p. 056709. (See pp. 38, 39)
- [141] A. Pukhov and J. Meyer-ter Vehn. “Laser wake field acceleration: the highly non-linear broken-wave regime”. In: *Applied Physics B* 74, pp. 355–361, 2002, pp. 355–361. (See p. 38)
- [142] I. Kostyukov, A. Pukhov, and S. Kiselev. “Phenomenological theory of laser-plasma interaction in “bubble” regime”. In: *Physics of Plasmas* 11, pp. 5256–5264, 2004, pp. 5256–5264. (See pp. 39, 51)
- [143] W. B. Mori. “The physics of the nonlinear optics of plasmas at relativistic intensities for short-pulse lasers”. In: *IEEE Journal of Quantum Electronics* 33, pp. 1942–1953, 1997, pp. 1942–1953. (See pp. 40, 41, 45)
- [144] P. Sprangle, C. M. Tang, and E. Esarey. “Relativistic Self-Focusing of Short-Pulse Radiation Beams in Plasmas”. In: *IEEE Transactions on Plasma Science* 15, pp. 145–153, 1987, pp. 145–153. (See p. 42)
- [145] J. E. Ralph, K. A. Marsh, A. E. Pak, W. Lu, C. E. Clayton, F. Fang, W. B. Mori, and C. Joshi. “Self-Guiding of Ultrashort, Relativistically Intense Laser Pulses through Underdense Plasmas in the Blowout Regime”. In: *Physical Review Letters* 102, p. 175003, 2009, p. 175003. (See p. 44)
- [146] M. Streeter, S. Kneip, M. Bloom, R. Bendoyro, O. Chekhlov, A. Dangor, A. Döpp, C. Hooker, J. Holloway, J. Jiang, N. Lopes, H. Nakamura, P. Norreys, C. A. Palmer, P. Rajeev, J. Schreiber, D. Symes, M. Wing, S. P. D. Mangles, and Z. Najmudin. “Observation of Laser Power Amplification in a Self-Injecting Laser Wakefield Accelerator”. In: *Physical Review Letters* 120, p. 254801, 2018, p. 254801. (See pp. 46, 124, 152, 172, 179)
- [147] L. M. Gorbunov and V. I. Kirsanov. “Excitation of plasma waves by an electromagnetic wave packet”. In: *JETP* 66, p. 290, 1986, p. 290. (See p. 47)
- [148] B. A. Shadwick, C. B. Schroeder, and E. Esarey. “Nonlinear laser energy depletion in laser-plasma accelerators”. In: *Physics of Plasmas* 16, p. 056704, 2009, p. 056704. (See pp. 47, 124, 172, 179)
- [149] C. D. Decker, W. B. Mori, K. Tzeng, and T. Katsouleas. “The evolution of ultra-intense, short-pulse lasers in underdense plasmas”. In: *Physics of Plasmas* 3, pp. 2047–2056, 1996, pp. 2047–2056. (See pp. 47, 48, 122, 179)
- [150] S. V. Bulanov, M. Yamagiwa, T. Z. Esirkepov, J. K. Koga, M. Kando, Y. Ueshima, K. Saito, and D. Wakabayashi. “Spectral and dynamical features of the electron bunch accelerated by a short-pulse high intensity laser in an underdense plasma”. In: *Physics of Plasmas* 12, p. 073103, 2005, p. 073103. (See pp. 50, 52)

[151] C. A. Coverdale, C. B. Darrow, C. D. Decker, W. B. Mori, K. C. Tzeng, K. A. Marsh, C. E. Clayton, and C. Joshi. “Propagation of Intense Subpicosecond Laser Pulses through Underdense Plasmas”. In: *Physical Review Letters* 74, pp. 4659–4662, 1995, pp. 4659–4662. (See p. 51)

[152] I. Kostyukov, E. Nerush, A. Pukhov, and V. Seredov. “Electron Self-Injection in Multidimensional Relativistic-Plasma Wake Fields”. In: *Physical Review Letters* 103, p. 175003, 2009, p. 175003. (See p. 51)

[153] A. G. R. Thomas. “Scalings for radiation from plasma bubbles”. In: *Physics of Plasmas* 17, p. 056708, 2010, p. 056708. (See pp. 51, 117, 118, 125, 155, 179, 187)

[154] S. P. D. Mangles, G. Genoud, M. S. Bloom, M. Burza, Z. Najmudin, A. Persson, K. Svensson, A. G. R. Thomas, and C. G. Wahlström. “Self-injection threshold in self-guided laser wakefield accelerators”. In: *Physical Review Special Topics - Accelerators and Beams* 15, p. 011302, 2012, p. 011302. (See p. 51)

[155] S. Kalmykov, S. A. Yi, V. Khudik, and G. Shvets. “Electron Self-Injection and Trapping into an Evolving Plasma Bubble”. In: *Physical Review Letters* 103, p. 135004, 2009, p. 135004. (See pp. 52, 179)

[156] H. Suk, N. Barov, J. B. Rosenzweig, and E. Esarey. “Plasma Electron Trapping and Acceleration in a Plasma Wake Field Using a Density Transition”. In: *Physical Review Letters* 86, pp. 1011–1014, 2001, pp. 1011–1014. (See p. 52)

[157] J. Faure, C. Rechatin, A. Norlin, A. Lifschitz, Y. Glinec, and V. Malka. “Controlled injection and acceleration of electrons in plasma wakefields by colliding laser pulses”. In: *Nature* 444, pp. 737–9, 2006, pp. 737–9. (See p. 52)

[158] A. V. Brantov, T. Z. Esirkepov, M. Kando, H. Kotaki, V. Y. Bychenkov, and S. V. Bulanov. “Controlled electron injection into the wake wave using plasma density inhomogeneity”. In: *Physics of Plasmas* 15, p. 073111, 2008, p. 073111. (See p. 52)

[159] A. J. Gonsalves, K. Nakamura, C. Lin, D. Panasenko, S. Shiraishi, T. Sokollik, C. Benedetti, C. B. Schroeder, C. G. R. Geddes, J. van Tilborg, J. Osterhoff, E. Esarey, C. Toth, and W. P. Leemans. “Tunable laser plasma accelerator based on longitudinal density tailoring”. In: *Nature Physics* 7, pp. 862–866, 2011, pp. 862–866. (See p. 52)

[160] M. Hansson, B. Aurand, X. Davoine, H. Ekerfelt, K. Svensson, A. Persson, C. G. Wahlström, and O. Lundh. “Down-ramp injection and independently controlled acceleration of electrons in a tailored laser wakefield accelerator”. In: *Physical Review Special Topics - Accelerators and Beams* 18, p. 071303, 2015, p. 071303. (See p. 52)

[161] K. Schmid, A. Buck, C. M. S. Sears, J. M. Mikhailova, R. Tautz, D. Herrmann, M. Geissler, F. Krausz, and L. Veisz. “Density-transition based electron injector for laser driven wakefield accelerators”. In: *Physical Review Special Topics - Accelerators and Beams* 13, p. 091301, 2010, p. 091301. (See p. 52)

[162] A. Buck, J. Wenz, J. Xu, K. Khrennikov, K. Schmid, M. Heigoldt, J. M. Mikhailova, M. Geissler, B. Shen, F. Krausz, S. Karsch, and L. Veisz. “Shock-Front Injector for High-Quality Laser-Plasma Acceleration”. In: *Physical Review Letters* 110, p. 185006, 2013, p. 185006. (See p. 52)

[163] T. P. Rowlands-Rees, C. Kamperidis, S. Kneip, A. J. Gonsalves, S. P. D. Mangles, J. G. Gallacher, E. Brunetti, T. Ibotson, C. D. Murphy, P. S. Foster, M. J. V. Streeter, F. Budde, P. A. Norreys, D. A. Jaroszynski, K. Krushelnick, Z. Najmudin, and S. M. Hooker. “Laser-Driven Acceleration of Electrons in a Partially Ionized Plasma Channel”. In: *Physical Review Letters* 100, p. 105005, 2008, p. 105005. (See p. 53)

[164] C. McGuffey, A. G. R. Thomas, W. Schumaker, T. Matsuoka, V. Chvykov, F. J. Dollar, G. Kalintchenko, V. Yanovsky, A. Maksimchuk, K. Krushelnick, V. Y. Bychenkov, I. V. Glazyrin, and A. V. Karpeev. “Ionization Induced Trapping in a Laser Wakefield Accelerator”. In: *Physical Review Letters* 104, p. 025004, 2010, p. 025004. (See pp. 53, 125)

[165] A. Pak, K. A. Marsh, S. F. Martins, W. Lu, W. B. Mori, and C. Joshi. “Injection and Trapping of Tunnel-Ionized Electrons into Laser-Produced Wakes”. In: *Physical Review Letters* 104, p. 025003, 2010, p. 025003. (See p. 53)

[166] M. Chen, E. Esarey, C. B. Schroeder, C. G. R. Geddes, and W. P. Leemans. “Theory of ionization-induced trapping in laser-plasma accelerators”. In: *Physics of Plasmas* 19, p. 033101, 2012, p. 033101. (See p. 53)

[167] L. L. Yu, E. Esarey, C. Schroeder, J. L. Vay, C. Benedetti, C. Geddes, M. Chen, and W. Leemans. “Two-Color Laser-Ionization Injection”. In: *Physical Review Letters* 112, p. 125001, 2014, p. 125001. (See pp. 54, 155)

[168] X. Xu, Y. Wu, C. Zhang, F. Li, Y. Wan, J. Hua, C. H. Pai, W. Lu, P. Yu, C. Joshi, and W. Mori. “Low emittance electron beam generation from a laser wakefield accelerator using two laser pulses with different wavelengths”. In: *Physical Review Special Topics - Accelerators and Beams* 17, p. 061301, 2014, p. 061301. (See p. 54)

[169] R. Weingartner, S. Raith, A. Popp, S. Chou, J. Wenz, K. Khrennikov, M. Heigoldt, A. R. Maier, N. Kajumba, M. Fuchs, B. Zeitler, F. Krausz, S. Karsch, and F. Grüner. “Ultralow emittance electron beams from a laser-wakefield accelerator”. In: *Physical Review Special Topics - Accelerators and Beams* 15, p. 111302, 2012, p. 111302. (See pp. 54, 156)

[170] J. C. Wood. “Betatron Radiation from Laser Wakefield Accelerators and its Applications”. Thesis. 2016. (See pp. 54, 77, 130, 181)

[171] K. Floettmann. “Some basic features of the beam emittance”. In: *Physical Review Special Topics - Accelerators and Beams* 6, p. 034202, 2003, p. 034202. (See pp. 56, 183)

- [172] S.-Y. Lee. *Accelerator Physics (Fourth Edition)*. Accelerator Physics. World Scientific, 2019. (See p. 59)
- [173] J. D. Jackson. *Classical electrodynamics*. 3rd ed. New York ; Wiley, 1999. (See p. 60)
- [174] C. J. Hooker, J. L. Collier, O. Chekhlov, R. Clarke, E. Divall, K. Ertel, B. Fell, P. Foster, S. Hancock, A. Langley, D. Neely, J. Smith, and B. Wyborn. “The Astra Gemini project A dual-beam petawatt Ti:Sapphire laser system”. In: *Journal de Physique IV* 133, pp. 673–677, 2006, pp. 673–677. (See p. 62)
- [175] K. Poder. “Characterisation of self-guided laser wakefield accelerators to multi-GeV energies”. Thesis. 2016. (See pp. 62, 93, 130)
- [176] *eCat2*. Central Laser Facility URL: <https://gemini.clf.stfc.ac.uk/ecat2/> (visited on Nov. 1, 2022) (see p. 63)
- [177] S. Semushin and V. Malka. “High density gas jet nozzle design for laser target production”. In: *Review of Scientific Instruments* 72, pp. 2961–2965, 2001, pp. 2961–2965. (See p. 66)
- [178] D. B. Atkinson and M. A. Smith. “Design and characterization of pulsed uniform supersonic expansions for chemical applications”. In: *Review of Scientific Instruments* 66, pp. 4434–4446, 1995, pp. 4434–4446. (See p. 66)
- [179] N. Lemos, N. Lopes, J. M. Dias, and F. Viola. “Design and characterization of supersonic nozzles for wide focus laser-plasma interactions”. In: *Review of Scientific Instruments* 80, p. 103301, 2009, p. 103301. (See p. 66)
- [180] S. Kuschel, M. Schwab, M. Yeung, D. Hollatz, A. Seidel, W. Ziegler, A. Sävert, M. Kaluza, and M. Zepf. “Controlling the Self-Injection Threshold in Laser Wakefield Accelerators”. In: *Physical Review Letters* 121, p. 154801, 2018, p. 154801. (See pp. 66, 155)
- [181] B. Fryxell, K. Olson, P. Ricker, F. X. Timmes, M. Zingale, D. Q. Lamb, P. MacNeice, R. Rosner, J. W. Truran, and H. Tufo. “FLASH: An Adaptive Mesh Hydrodynamics Code for Modeling Astrophysical Thermonuclear Flashes”. In: *The Astrophysical Journal Supplement Series* 131, p. 273, 2000, p. 273. (See p. 67)
- [182] C. J. Dasch. “One-dimensional tomography: a comparison of Abel, onion-peeling, and filtered backprojection methods”. In: *Appl. Opt.* 31, pp. 1146–1152, 1992, pp. 1146–1152. (See p. 73)
- [183] J. Lindström. “Radioluminescence: A simple model for fluorescent layers - analysis and applications”. Thesis. 2021. (See p. 75)
- [184] Y. Glinec, J. Faure, A. Guemnie-Tafo, V. Malka, H. Monard, J. P. Larbre, V. De Waele, J. L. Marignier, and M. Mostafavi. “Absolute calibration for a broad range single shot electron spectrometer”. In: *Review of Scientific Instruments* 77, p. 103301, 2006, p. 103301. (See pp. 76, 165)

[185] E.-J. Popovici, L. Muresan, A. Hristea-Simoc, E. Indrea, M. Vasilescu, M. Nazarov, and D. Y. Jeon. “Synthesis and characterisation of rare earth oxysulphide phosphors. I. Studies on the preparation of Gd₂O₂S:Tb phosphor by the flux method”. In: *Optical Materials* 27, pp. 559–565, 2004, pp. 559–565. (See p. 76)

[186] J. Lindström, G. Alm Carlsson, E. Wåhlin, Carlsson Tedgren, and G. Poludniowski. “Experimental assessment of a phosphor model for estimating the relative extrinsic efficiency in radioluminescent detectors”. In: *Physica Medica* 76, pp. 117–124, 2020, pp. 117–124. (See p. 76)

[187] *LANEX screens*. Carestream URL: <https://www.carestream.com/en/us/medical/products/radiography/film-systems/film-solutions-for-radiology/lanex-screens> (visited on Dec. 5, 2022) (see p. 76)

[188] *X-ray Phosphor Screens*. Scintacor URL: <https://scintacor.com/products/x-ray-scintillator-screens/> (visited on Dec. 5, 2022) (see p. 76)

[189] G. C. Tyrrell. “Phosphors and scintillators in radiation imaging detectors”. In: *Nucl. Instrum. Methods Phys. Res. A: Accel. Spectrom. Detect. Assoc. Equip.* 546, pp. 180–187, 2005, pp. 180–187. (See p. 76)

[190] *DRZ Screens*. MCI Optonix LLC URL: <http://www.mcio.com/Products/drz-screens.aspx> (visited on Dec. 5, 2022) (see p. 76)

[191] A. Buck, K. Zeil, A. Popp, K. Schmid, A. Jochmann, S. D. Kraft, B. Hidding, T. Kudyakov, C. M. S. Sears, L. Veisz, S. Karsch, J. Pawelke, R. Sauerbrey, T. Cowan, F. Krausz, and U. Schramm. “Absolute charge calibration of scintillating screens for relativistic electron detection”. In: *Review of Scientific Instruments* 81, p. 033301, 2010, p. 033301. (See p. 76)

[192] S. G. Gales and C. D. Bentley. “Image plates as x-ray detectors in plasma physics experiments”. In: *Review of Scientific Instruments* 75, pp. 4001–4003, 2004, pp. 4001–4003. (See p. 76)

[193] I. J. Paterson, R. J. Clarke, N. C. Woolsey, and G. Gregori. “Image plate response for conditions relevant to laser–plasma interaction experiments”. In: *Measurement Science and Technology* 19, p. 095301, 2008, p. 095301. (See p. 77)

[194] K. A. Tanaka, T. Yabuuchi, T. Sato, R. Kodama, Y. Kitagawa, T. Takahashi, T. Ikeda, Y. Honda, and S. Okuda. “Calibration of imaging plate for high energy electron spectrometer”. In: *Review of Scientific Instruments* 76, p. 013507, 2004, p. 013507. (See p. 77)

[195] K. Zeil, S. D. Kraft, A. Jochmann, F. Kroll, W. Jahr, U. Schramm, L. Karsch, J. Pawelke, B. Hidding, and G. Pretzler. “Absolute response of Fuji imaging plate detectors to picosecond-electron bunches”. In: *Review of Scientific Instruments* 81, p. 013307, 2010, p. 013307. (See p. 77)

[196] S. Kneip. “Laser Plasma Accelerator and Wiggler”. Thesis. 2010. (See p. 77)

[197] G. E. Giakoumakis and D. M. Miliotis. “Light angular distribution of fluorescent screens excited by x-rays”. In: *Physics in Medicine Biology* 30, p. 21, 1985, p. 21. (See p. 77)

[198] Y. Yamazaki, T. Kurihara, H. Kobayashi, I. Sato, and A. Asami. “High-precision pepper-pot technique for a low-emittance electron beam”. In: *Nucl. Instrum. Methods Phys. Res. A: Accel. Spectrom. Detect. Assoc. Equip.* 322, pp. 139–145, 1992, pp. 139–145. (See p. 80)

[199] R. P. Shanks, M. P. Anania, E. Brunetti, S. Cipiccia, B. Ersfeld, J. G. Gallacher, R. C. Issac, M. R. Islam, G. Vieux, G. H. Welsh, S. M. Wiggins, and D. A. Jaroszynski. “Pepper-pot emittance measurement of laser-plasma wakefield accelerated electrons”. In: *Proc. SPIE*. Vol. 7359. P. 735907. (See p. 80)

[200] E. Brunetti, R. P. Shanks, G. G. Manahan, M. R. Islam, B. Ersfeld, M. P. Anania, S. Cipiccia, R. C. Issac, G. Raj, G. Vieux, G. H. Welsh, S. M. Wiggins, and D. A. Jaroszynski. “Low Emittance, High Brilliance Relativistic Electron Beams from a Laser-Plasma Accelerator”. In: *Physical Review Letters* 105, p. 215007, 2010, p. 215007. (See pp. 80, 156)

[201] M. Zhang. *Emittance Formula for Slits and Pepper-pot Measurement*. Technical Report. FERMILAB-TM, 1988 (see p. 80)

[202] F. Irshad, S. Karsch, and A. Döpp. “Multi-objective and multi-fidelity Bayesian optimization of laser-plasma acceleration”. In: *arXiv e-prints*, arXiv:2210.03484, 2022, arXiv:2210.03484. (See pp. 82, 107)

[203] L. Fedeli, A. Huebl, F. Boillod-Cerneux, T. Clark, K. Gott, C. Hillairet, S. Jaure, A. Leblanc, R. Lehe, A. Myers, C. Piechurski, M. Sato, N. Zaim, W. Zhang, J. Vay, and H. Vincenti. “Pushing the Frontier in the Design of Laser-Based Electron Accelerators with Groundbreaking Mesh-Refined Particle-In-Cell Simulations on Exascale-Class Supercomputers”. In. 2022. (See p. 82)

[204] J. L. Vay, A. Almgren, J. Bell, L. Ge, D. P. Grote, M. Hogan, O. Kononenko, R. Lehe, A. Myers, C. Ng, J. Park, R. Ryne, O. Shapoval, M. Thévenet, and W. Zhang. “Warp-X: A new exascale computing platform for beam–plasma simulations”. In: *Nucl. Instrum. Methods Phys. Res. A: Accel. Spectrom. Detect. Assoc. Equip.* 909, pp. 476–479, 2018, pp. 476–479. (See p. 82)

[205] J. P. Boris. “Relativistic Plasma Simulation - Optimization of a Hybrid Code”. In: *Numerical Simulations of Plasmas*. Ed. by J. P. Boris. Vol. 3. 1970. (See p. 83)

[206] J. L. Vay. “Simulation of beams or plasmas crossing at relativistic velocity”. In: *Physics of Plasmas* 15, p. 056701, 2008, p. 056701. (See p. 83)

[207] B. Ripperda, F. Bacchini, J. Teunissen, C. Xia, O. Porth, L. Sironi, G. Lapenta, and R. Keppens. “A Comprehensive Comparison of Relativistic Particle Integrators”. In: *The Astrophysical Journal Supplement Series* 235, p. 21, 2018, p. 21. (See p. 83)

[208] K. Yee. “Numerical solution of initial boundary value problems involving maxwell’s equations in isotropic media”. In: *IEEE Transactions on Antennas and Propagation* 14, pp. 302–307, 1966, pp. 302–307. (See p. 83)

[209] R. Nuter, M. Grech, P. Gonzalez de Alaiza Martinez, G. Bonnaud, and E. d’Humières. “Maxwell solvers for the simulations of the laser-matter interaction”. In: *The European Physical Journal D* 68, p. 177, 2014, p. 177. (See p. 84)

[210] R. Lehe, A. Lifschitz, C. Thaury, V. Malka, and X. Davoine. “Numerical growth of emittance in simulations of laser-wakefield acceleration”. In: *Physical Review Special Topics - Accelerators and Beams* 16, p. 021301, 2013, p. 021301. (See p. 85)

[211] A. D. Greenwood, K. L. Cartwright, J. W. Luginsland, and E. A. Baca. “On the elimination of numerical Cerenkov radiation in PIC simulations”. In: *Journal of Computational Physics* 201, pp. 665–684, 2004, pp. 665–684. (See p. 85)

[212] O. Buneman, C. W. Barnes, J. C. Green, and D. E. Nielsen. “Principles and capabilities of 3-D, E-M particle simulations”. In: *Journal of Computational Physics* 38, pp. 1–44, 1980, pp. 1–44. (See p. 85)

[213] J. Derouillat, A. Beck, F. Pérez, T. Vinci, M. Chiaramello, A. Grassi, M. Flé, G. Bouchard, I. Plotnikov, N. Aunai, J. Dargent, C. Riconda, and M. Grech. “Smilei : A collaborative, open-source, multi-purpose particle-in-cell code for plasma simulation”. In: *Computer Physics Communications* 222, pp. 351–373, 2018, pp. 351–373. (See pp. 85, 171)

[214] A. Grassi, M. Grech, F. Amiranoff, F. Pegoraro, A. Macchi, and C. Riconda. “Electron Weibel instability in relativistic counterstreaming plasmas with flow-aligned external magnetic fields”. In: *Physical Review E* 95, p. 023203, 2017, p. 023203. (See p. 85)

[215] A. Grassi, M. Grech, F. Amiranoff, A. Macchi, and C. Riconda. “Radiation-pressure-driven ion Weibel instability and collisionless shocks”. In: *Physical Review E* 96, p. 033204, 2017, p. 033204. (See p. 85)

[216] R. Lehe, M. Kirchen, I. A. Andriyash, B. B. Godfrey, and J.-L. Vay. “A spectral, quasi-cylindrical and dispersion-free Particle-In-Cell algorithm”. In: *Computer Physics Communications* 203, pp. 66–82, 2016, pp. 66–82. (See pp. 86, 112, 149, 178)

[217] I. V. Il’ina, T. Y. Cherezova, and A. V. Kudryashov. “Gerchberg—Saxton algorithm: experimental realisation and modification for the problem of formation of multimode laser beams”. In: *Quantum Electronics* 39, p. 521, 2009, p. 521. (See p. 87)

[218] A. Borot and F. Quéré. “Spatio-spectral metrology at focus of ultrashort lasers: a phase-retrieval approach”. In: *Optics Express* 26, pp. 26444–26461, 2018, pp. 26444–26461. (See p. 87)

[219] S. V. Rozario. “Novel injection and targetry in laser wakefield acceleration”. Thesis. 2020. (See pp. 93, 130)

[220] R. J. Shalloo, S. J. D. Dann, J. N. Gruse, C. I. D. Underwood, A. F. Antoine, C. Arran, M. Backhouse, C. D. Baird, M. D. Balcazar, N. Bourgeois, J. A. Cardarelli, P. Hatfield, J. Kang, K. Krushelnick, S. P. D. Mangles, C. D. Murphy, N. Lu, J. Osterhoff, K. Pöder, P. P. Rajeev, C. P. Ridgers, S. Rozario, M. P. Selwood, A. J. Shahani, D. R. Symes, A. G. R. Thomas, C. Thornton, Z. Najmudin, and M. J. V. Streeter. “Automation and control of laser wakefield accelerators using Bayesian optimization”. In: *Nature Communications* 11, p. 6355, 2020, p. 6355. (See pp. 93, 102)

[221] V. Leroux, T. Eichner, and A. R. Maier. “Description of spatio-temporal couplings from heat-induced compressor grating deformation”. In: *Optics Express* 28, pp. 8257–8265, 2020, pp. 8257–8265. (See p. 94)

[222] C. Iaconis and I. A. Walmsley. “Spectral phase interferometry for direct electric-field reconstruction of ultrashort optical pulses”. In: *Optics Letters* 23, pp. 792–794, 1998, pp. 792–794. (See p. 95)

[223] *Fastlite Dazzler product information*. Fastlite URL: <https://fastlite.com/produits/dazzler-ultrafast-pulse-shaper/> (visited on Oct. 12, 2022) (see pp. 95, 164)

[224] *HASO4 FIRST*. Imagine Optic URL: <https://www.imagine-optic.com/products/haso4-first-wavefront-sensor/> (visited on Nov. 17, 2022) (see pp. 95, 134)

[225] J. Strehlow, P. Forestier-Colleoni, C. McGuffey, M. Bailly-Grandvaux, T. S. Daykin, E. McCary, J. Peebles, G. Revet, S. Zhang, T. Ditmire, M. Donovan, G. Dyer, J. Fuchs, E. W. Gaul, D. P. Higginson, G. E. Kemp, M. Martinez, H. S. McLean, M. Spinks, H. Sawada, and F. N. Beg. “The response function of Fujifilm BAS-TR imaging plates to laser-accelerated titanium ions”. In: *Review of Scientific Instruments* 90, p. 083302, 2019, p. 083302. (See p. 98)

[226] *iKon-M 934 BR-DD information*. Andor URL: <https://www.photonicsonline.com/doc/deep-depletion-ccd-camera-ikon-m-934-br-dd-0002> (visited on Dec. 12, 2022) (see p. 100)

[227] J.-N. Gruse. “Development of Laser Wakefield Accelerators”. Thesis. 2020. (See pp. 100, 188)

[228] S. Fritzler, E. Lefebvre, V. Malka, F. Burgy, A. E. Dangor, K. Krushelnick, S. P. D. Mangles, Z. Najmudin, J. P. Rousseau, and B. Walton. “Emittance Measurements of a Laser-Wakefield-Accelerated Electron Beam”. In: *Physical Review Letters* 92, p. 165006, 2004, p. 165006. (See p. 156)

[229] G. R. Plateau, C. G. R. Geddes, D. B. Thorn, M. Chen, C. Benedetti, E. Esarey, A. J. Gonsalves, N. H. Matlis, K. Nakamura, C. B. Schroeder, S. Shiraishi, T. Sokollik, J. van Tilborg, C. Toth, S. Trotsenko, T. S. Kim, M. Battaglia, T. Stöhlker, and W. P. Leemans. “Low-Emittance Electron Bunches from a Laser-Plasma Accelerator Measured using Single-Shot X-Ray Spectroscopy”. In: *Physical Review Letters* 109, p. 064802, 2012, p. 064802. (See p. 156)

[230] S. Barber, J. van Tilborg, C. Schroeder, R. Lehe, H. E. Tsai, K. Swanson, S. Steinke, K. Nakamura, C. Geddes, C. Benedetti, E. Esarey, and W. Leemans. “Measured Emittance Dependence on the Injection Method in Laser Plasma Accelerators”. In: *Physical Review Letters* 119, p. 104801, 2017, p. 104801. (See pp. 156, 158)

[231] A. Curcio, M. Anania, F. Bisesto, E. Chiadroni, A. Cianchi, M. Ferrario, F. Filippi, D. Giulietti, A. Marocchino, M. Petrarca, V. Shpakov, and A. Zigler. “Trace-space reconstruction of low-emittance electron beams through betatron radiation in laser-plasma accelerators”. In: *Physical Review Accelerators and Beams* 20, p. 012801, 2017, p. 012801. (See p. 156)

[232] F. Brunel. “Not-so-resonant, resonant absorption”. In: *Physical Review Letters* 59, pp. 52–55, 1987, pp. 52–55. (See pp. 162, 174)

[233] B. H. Shaw, S. Steinke, J. van Tilborg, and W. P. Leemans. “Reflectance characterization of tape-based plasma mirrors”. In: *Physics of Plasmas* 23, p. 063118, 2016, p. 063118. (See p. 162)

[234] M. B. Reid. “Electron beam emittance growth in thin foils: A betatron function analysis”. In: *Journal of Applied Physics* 70, pp. 7185–7187, 1991, pp. 7185–7187. (See pp. 165, 184)

[235] V. L. Highland. “Some practical remarks on multiple scattering”. In: *Nuclear Instruments and Methods* 129, pp. 497–499, 1975, pp. 497–499. (See pp. 165, 184)

[236] D. Marx, J. Giner Navarro, D. Cesar, J. Maxson, B. Marchetti, R. Assmann, and P. Musumeci. “Single-shot reconstruction of core 4D phase space of high-brightness electron beams using metal grids”. In: *Physical Review Accelerators and Beams* 21, p. 102802, 2018, p. 102802. (See p. 166)

[237] Y. Lang, X. H. Yang, H. Xu, Z. Jin, and H. B. Zhuo. “Influence of field ionization effect on the divergence of laser-driven fast electrons”. In: *Plasma Physics and Controlled Fusion* 60, p. 075002, 2018, p. 075002. (See pp. 171, 174)

[238] E. S. Weibel. “Spontaneously Growing Transverse Waves in a Plasma Due to an Anisotropic Velocity Distribution”. In: *Physical Review Letters* 2, pp. 83–84, 1959, pp. 83–84. (See p. 174)

[239] B. D. Fried. “Mechanism for Instability of Transverse Plasma Waves”. In: *The Physics of Fluids* 2, pp. 337–337, 1959, pp. 337–337. (See p. 174)

- [240] F. Fiúza, R. A. Fonseca, J. Tonge, W. B. Mori, and L. O. Silva. “Weibel-Instability-Mediated Collisionless Shocks in the Laboratory with Ultraintense Lasers”. In: *Physical Review Letters* 108, p. 235004, 2012, p. 235004. (See p. 174)
- [241] P. Gibbon. *Short pulse laser interactions with matter*. London: Imperial College Press, 2005. (See pp. 174, 198)
- [242] N. Dover. “Exploring novel regimes for ion acceleration driven by intense laser radiation”. Thesis. 2013. (See p. 174)
- [243] S. C. Wilks, W. L. Kruer, M. Tabak, and A. B. Langdon. “Absorption of ultra-intense laser pulses”. In: *Physical Review Letters* 69, pp. 1383–1386, 1992, pp. 1383–1386. (See p. 176)
- [244] M. Tzoufras, W. Lu, F. S. Tsung, C. Huang, W. B. Mori, T. Katsouleas, J. Vieira, R. A. Fonseca, and L. O. Silva. “Beam Loading in the Nonlinear Regime of Plasma-Based Acceleration”. In: *Physical Review Letters* 101, p. 145002, 2008, p. 145002. (See p. 179)
- [245] *Kapton Properties*. Particle Data Group, LBNL URL: <https://pdg.lbl.gov/2005/AtomicNuclearProperties/substances/kapton.html> (visited on Dec. 10, 2022) (see p. 184)
- [246] R. Ariniello, C. Doss, K. Hunt-Stone, J. Cary, and M. Litos. “Transverse beam dynamics in a plasma density ramp”. In: *Physical Review Accelerators and Beams* 22, p. 041304, 2019, p. 041304. (See p. 190)