Resonant excitation of plasma wakefields in long plasma channels



Aimee Ross

St John's College

University of Oxford

Submitted in partial fulfilment of the requirements for the degree of Doctor of Philosophy at the University of Oxford

March 2023

Acknowledgements

Firstly, thank you to Simon for being such a supportive and kind supervisor. Thanks also to Roman, Laura and Rajeev who have provided many words of advice and encouragement throughout my DPhil research.

Thank you to Jimmy, Jakob, Alex, Alexander, Aaron, James, James, Linus, Warren, Oscar, Emily, Johannes, David, Darren, Sebastian, Harry and Lewis. I have learnt so much from working with you. Thanks to everyone in the Advanced Accelerators field for creating such a welcoming and friendly atmosphere. Thank you to all the friends I have made along the way.

Finally, thanks to Emma for your wisdom, to Olek for putting up with me, and to my parents for everything you have done and continue to do for me. You guys are awesome.

Contents

1 Introduction **1**

1.1 Particle accelerators . . . 1

1.2 Current limitations of laser-plasma accelerators . . . 2

1.2.1 Diffraction of the drive laser pulse . . . 3

1.2.2 Repetition rate of the accelerator . . . 4

1.3 Outline of the thesis . . . 5

1.3.1 Role of the author . . . 6

2 Background **8**

2.1 Laser-plasma accelerators . . . 8

2.1.1 Key concepts . . . 8

2.1.2 Electron injection . . . 11

2.1.3 Limits on energy gain . . . 12

2.2 Resonant excitation of plasma waves . . . 14

2.2.1 Plasma beat-wave acceleration . . . 14

2.2.2 Self-modulated laser wakefield acceleration . . . 16

2.2.3 Multi-pulse laser wakefield acceleration . . . 17

2.2.4 The plasma-modulated plasma accelerator regime . . . 20

2.3 Guiding of lasers pulses . . . 22

2.3.1 Capillary discharge waveguides . . . 24

2.3.2 Hydrodynamic inverse bremsstrahlung plasma channels . . . 24

2.3.3 Hydrodynamic optical-field-ionised plasma channels . . . 25

2.4 Laser-plasma instabilities . . . 26

2.4.1 Self-modulation instability . . . 26

2.4.2 Laser-hosing instability . . . 27

2.4.3 Growth rates . . . 28

3 Theory and Methods **30**

3.1 Pulse train generation and measurement 30

3.1.1 Ultrashort laser pulses 30

3.1.2 Pulse train generation using plasma modulation 32

3.1.3 Pulse train generation from an ultrashort laser pulse 33

3.1.4 Measuring ultrashort pulses 35

3.2 Generating HOFI plasma channels 38

3.2.1 Considerations for the longitudinally-extended focus 40

3.2.2 Physics of optical field ionisation 41

3.2.3 Hydrodynamic expansion of hydrogen plasma 47

3.2.4 Transverse mode profiles of plasma channels 49

3.2.5 Conditioned HOFI plasma channels 51

3.3 Measuring plasma wakefields 53

3.3.1 Phase shifts in a probe beam 53

3.3.2 Spectral shifts in the drive beam 53

3.4 Analysing images of laser transverse intensity profiles 56

4 Novel Optics for the Formation of HOFI Plasma Channels **59**

4.1 Introduction 59

4.2 Theory of longitudinally-extended foci 62

4.2.1 Bessel beams 62

4.2.2 Practical realisation of linear axicons 65

4.3 Modelling optical beam propagation 68

4.3.1 The angular spectrum method 68

4.3.2 Numerical code 69

4.3.3 Code validation 71

4.4 Chromatic effects on axicon-focused beams used to generate HOFI plasma channels 76

4.4.1 Chromatic dispersion 76

4.4.2 B-integral 78

4.4.3 Chromatic effects through a phase plate 80

4.4.4 Discussion 81

4.5 Longitudinally-extended foci with uniform axial intensity 82
4.5.1 Properties of axilenses 83
4.5.2 Practical realisation of axilenses 84
4.5.3 Design of axilens phase plate for HOFI plasma channel generation 86
4.5.4 Results of numerical calculation 87
4.6 Conclusion 91

5 Reducing the Transmission Losses of HOFI Plasma Channels 92
5.1 Introduction 92
5.2 Increasing the spot size of longitudinally-extended foci 96
5.2.1 Constraints on standard single axicon setup 97
5.2.2 Higher order Bessel modes 99
5.2.3 Random phase plates 101
5.3 Optical studies using spatial light modulator 101
5.3.1 Spatial light modulator setup and calibration 102
5.3.2 Investigation of speckle patterns 104
5.3.3 Investigation of higher-order Bessel modes 109
5.3.4 Comparison with design objectives 112
5.4 Hydrodynamic simulations of higher-order Bessel modes 116
5.4.1 The Helios 1D code 116
5.4.2 Ionisation modelling 119
5.4.3 Results of hydrodynamic simulations 122
5.5 Conclusion 130

6 Observation of Resonant Excitation of Plasma Waves Driven in Long Plasma Channels 132
6.1 Introduction 132
6.2 Experimental setup 135
6.2.1 Astra-Gemini TA3 laser 135
6.2.2 Channel-formation 135
6.2.3 Generating a train of laser pulses 136
6.2.4 Drive beam focus properties 137
6.2.5 Gas target 138

6.3 Diagnostics . . . 140
6.3.1 Pulse train measurement . . . 140
6.3.2 Forward diagnostics . . . 142
6.3.3 Input laser beam diagnostics . . . 145
6.4 Guiding Joule-scale pulse trains over 10 centimetres . . . 147
6.4.1 Selecting well-guided shots for analysis . . . 148
6.4.2 Calculating the output energy in the laser beam . . . 151
6.4.3 Guiding measurements . . . 152
6.5 Resonant excitation of plasma waves in a long plasma channel . . . 161
6.5.1 Processing of spectral data . . . 161
6.5.2 Properties of transmitted spectra . . . 162
6.5.3 On-axis density calibration . . . 164
6.5.4 Resonance curve . . . 167
6.6 Extracting the wakefield amplitude . . . 170
6.6.1 Parameters for simulations . . . 171
6.6.2 Calculation results . . . 174
6.7 Discussion . . . 176
6.8 Conclusion . . . 178

7 Conclusion and outlook **180**
7.1 Summary . . . 180
7.2 Outlook . . . 182

A **185**
A.1 Method of Stationary Phase . . . 185
A.2 Phase accumulation through tilted axicon lens . . . 186

Bibliography **188**

Abstract

Laser-plasma accelerators (LPAs) promise a route towards accelerating electrons to GeV-scale energies in centimetre-scale distances. However, the ultra-short laser pulse systems typically used for LPAs are currently limited in repetition rate to $\lesssim 1\,\mathrm{Hz}$. Future laser-plasma accelerators (LPAs) will need to work at the kilohertz level to be useful for the most demanding applications in medicine, industry and research.

This thesis focuses on methods to achieve LPAs operating at the kilohertz level repetition rates. First, hydrodynamic optical-field-ionised (HOFI) plasma channels are studied, which have been shown to be capable of guiding laser pulses over hundreds of millimetres at low plasma densities [1] and are able to operate at kilohertz repetition rates [2]. In this thesis, limitations of the standard optical setups used to form HOFI plasma channels are investigated. Refractive axicon lenses are found to induce significant chromatic and non-linear effects on the ultra-short channel-forming laser pulse. The use of a thin phase plate is shown to overcome this problem. An alternative to the axicon, an axilens, is studied and a parameter regime for an axilens is found that achieves a longitudinally-extended focus over 300 mm with highly uniform axial-intensity and with spot-size variations at the 15% level.

Next, the transmission losses of HOFI plasma channels are studied. It is found that increasing the radius of the initial ionised plasma column is an effective way to deepen the resulting plasma channel, thereby reducing the expected energy leakage from the guided beam as it propagates. Higher-order Bessel modes and random phase plates are investigated as methods to increase the radius of the initial plasma column. Higher-order Bessel modes are shown in calculations to maximise the efficiency of energy conversion from the channel-forming laser pulse to the initial plasma column.

Hydrodynamic simulations show that Bessel modes of order 1 and 2 could generate plasma channel profiles with increased channel depth compared to the standard zeroth-order Bessel mode.

In the second part of the thesis, an experimental campaign relevant to the plasma-modulated plasma accelerator (P-MoPA) scheme is presented [3, 4]. Trains of laser pulses, consisting of ~ 10 pulses spaced at 170 fs with total energy (2.5 ± 0.5) J, are guided over 100 mm-long propagation distances in a HOFI plasma channel. Measurements of the spectrum of the transmitted pulse train show that a plasma wakefield is resonantly excited when the on-axis channel density is tuned to match the plasma period to the separation of pulses in the train. The resonant plasma density is shown to scale with pulse train spacing as expected. Resonant wakefield excitation is not observed when the pulse train is replaced by a temporally-smooth pulse of the same energy and same (1 ps full-width-half-maximum) total duration. Comparison of the experimental results with fluid calculations suggests that wakefield amplitudes up to 10 GV/m were excited by the pulse trains at resonance. These results constitute the first experimental demonstration of the accelerator stage of the P-MoPA scheme [4].

List of author publications

- **Ross, A. J.**, Alejo A., von Boetticher A., Cowley J., Holloway J., Jonnerby J., Picksley A., Walczak R. and Hooker S. M., *Numerical modelling of chromatic effects on axicon-focused beams used to generate HOFI plasma channels.* Journal of Physics: Conference Series. Vol. 1596. No. 1. IOP Publishing, (2020).

- Picksley A., Alejo A., Cowley J., Bourgeois N., Corner L., Feder L., Holloway J., Jones H., Jonnerby J., Milchberg H. M., Reid L. R., **Ross A. J.**, Walczak R., and Hooker S. M., *Guiding of high-intensity laser pulses in 100-mm-long hydrodynamic optical-field-ionized plasma channels.* Physical Review Accelerators and Beams 23.8 (2020): 081303.

- Jonnerby J., Bourgeois N., Corner L., Holloway J., Picksley A., **Ross A. J.**, Walczak R., and Hooker S. M., *Measurement of the decay of laser-driven linear plasma wakefields* [Manuscript in preparation].

- **Ross A. J.**, Chappell J., Archer E., Bourgeois N., Cowley J., Feder L., Jakobsson O., Picksley A., van de Wetering J. J., Wang W., Walczak R., and Hooker S. M., *Resonant wakefield excitation observed in long plasma channels* [Manuscript in preparation].

- Picksley A., Chappell J., Archer E., Bourgeois N., Cowley J., Feder L., Jakobsson O., **Ross A. J.**, Wang W., Walczak R., and Hooker S. M., *GeV electron bunches in low-density plasma channels by all-optical density transition injection* [Manuscript in preparation].

List of Abbreviations

LPA	Laser plasma accelerator
LHC	Large Hadron Collider
XFEL	X-ray free electron laser
HOFI	Hydrodynamic optical-field-ionised (plasma channel)
MP-LWFA	Multi-pulse laser wakefield accelerator
P-MoPA	Plasma-modulated plasma accelerator
PBWA	Plasma beatwave accelerator
SM-LWFA	Self-modulated laser wakefield accelerator
IB	Inverse bremsstrahlung
OFI	Optical field ionisation
CHOFI	Conditioned hydrodynamic optical-field-ionised (plasma channel)
GDD	Group delay dispersion
TOD	Third order delay
FOD	Fourth order delay
AOPDF	Acousto-optic programmable dispersive filter
SSA	Single shot autocorrelator
BSI	Barrier-suppression ionisation
ADK	Ammosov-Delone-Krainov (model)
USAF	United States Air-Force (target)

ROI	Region of interest
CLF	Central Laser Facility
RAL	Rutherford Appleton Laboratory
STFC	Science and Technology Facilities Council
ASM	Angular spectrum method
RMS	Root mean square
FWHM	Full-width half maximum
MSP	Method of stationary phase
HOBM	Higher-order Bessel mode
OAM	Orbital angular momentum
RPP	Random phase plate
SLM	Spatial light modulator
LCOS	Liquid crystal on silicon
LTE	Local thermodynamic equilibrium
OAP	Off-axis paraboloid
DM	Deformable mirror
AO	Adaptic optic
HM	Holed mirror
Ax	Axicon
UVFS	Ultra-violet grade fused silica
IP	Interaction point
LA	Laser area
ND	Neutral density
PIC	Particle-in-cell

List of Symbols

e	Electron charge
m_{e}	Electron mass
ϵ_0	Vacuum permittivity
c	Speed of light in vacuum
k_B	Boltzmann constant
h	Planck's constant
r_{e}	Classical electron radius
λ	Laser beam wavelength
ω	Laser beam angular frequency
k	Laser beam wavenumber
n_{e}	Electron density
η	Plasma refractive index
ω_p	Plasma (angular) frequency

Chapter 1

Introduction

1.1 Particle accelerators

The Large Hadron Collider (LHC) at CERN is a remarkable machine, spanning 27 km in circumference and equipped with technology that can probe nature on scales smaller than the nucleus of an atom [5]. Over the past century, particle accelerators like the LHC have been used to recreate conditions that were present in the seconds after the Big Bang, advancing our understanding of the building blocks of the universe [5].

Apart from applications in particle physics research, accelerator technology has been widely adopted by academia, industry and medicine. For example, in cancer treatment [6] accelerators are used to generate beams of charged particles that are directed towards tumours in the human body, destroying cancerous cells whilst mitigating the impact on surrounding healthy tissues. Additionally, particle accelerators can be adapted to generate x-rays, such as is achieved by synchrotron light sources and x-ray free electron lasers (X-FELs), which are used to probe material properties and chemical processes [7, 8].

Conventional particle accelerators, such as the LHC, consist of a sequence of metal structures known as cavities. Within each cavity, an electromagnetic wave supplies an electric field that transfers energy to the particles. By placing multiple cavities one after another, it is possible to accelerate particles to very high energies, where the energy is proportional to the electric field strength and the travelled distance.

The electric field in a conventional accelerator cavity is limited by the breakdown of the metal walls. During a breakdown event, electrons are pulled out from the metal surface and form a plasma, which disrupts the particle acceleration and can cause damage to the cavity. The highest acceleration gradients achieved by conventional accelerators

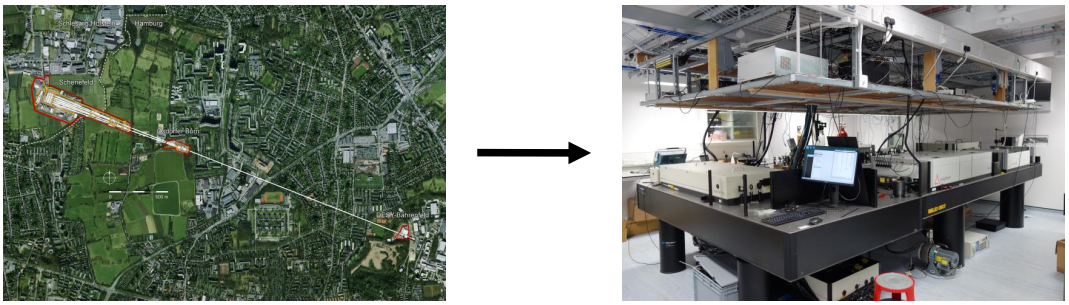


Figure 1.1: Images of the European X-FEL (left) [10], with kilometre-scale accelerator length, and Oxford Laser-Plasma Accelerator Lab (right), with accelerator lengths on the sub-metre scale.

have been demonstrated in the context of the compact linear collider (CLIC) design, where fields of up to 100 MV/m have been achieved [9]. Therefore, to reach multi-GeV particle energies, conventional accelerators need to be tens of metres in length, requiring considerable infrastructure and resources. For example, the linear accelerator which is used for the European X-FEL is 1.7 km long [10] (see Figure 1.1). This state-of-the-art facility is able to accelerate electron bunches to energies up to 17.5 GeV, using superconducting materials for the accelerator cavity structures. Because of the scale and cost of such a project, there are only a few such facilities in the world that are able to accelerate particles to multi-GeV energies.

To increase the availability of multi-GeV particle accelerators, the accelerator length must be reduced. Laser-plasma accelerators (LPAs) offer a way to reach higher electric field strengths than conventional accelerators. In an LPA, a laser beam drives a wave in the plasma (ionised gas) driving the electrons away from their host ions and setting up strong electric fields due to the charge separation (see Figure 1.2). A particle bunch can ride this wave and gain energy as it travels in the wake of the laser beam. Electric fields of the order 100 GV/m are possible using this method [11]. LPAs can therefore accelerate particles to GeV-energies over distances of a few centimetres and thereby provide a route towards compact accelerators that open up opportunities in applications of particle beams and light sources.

1.2 Current limitations of laser-plasma accelerators

The LPA was first proposed in 1979 by Tajima and Dawson [12], who used simulations to show that high-intensity laser pulses were capable of driving wakefields in a plasma with

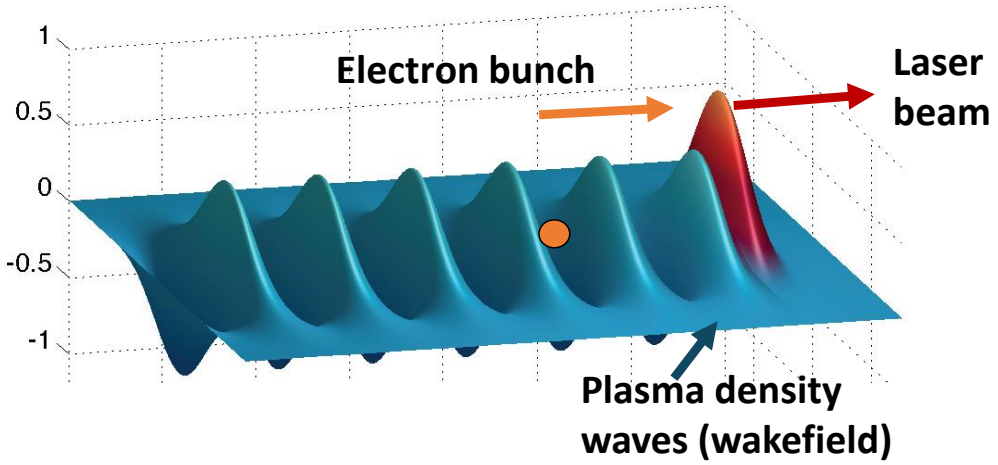


Figure 1.2: Schematic of the laser plasma accelerator, adapted from [3] with permission. The laser intensity envelope is marked in red, the plasma density in blue and the electron bunch in orange. The x-axis shows the laser beam propagation direction, the y-axis shows the normalised amplitudes of the intensity and plasma density.

acceleration gradients of 100 GeV/m and that an electron bunch could be trapped in this wakefield and accelerated. In 2004, three groups experimentally demonstrated that LPAs could be used to accelerate electron beams to hundreds of MeV, whilst simultaneously achieving small energy spreads [13–15]. Since then progress has continued, with GeV electron energies reached in 2006 [16] and the most recent results demonstrating acceleration to 8 GeV in 2019 [17] and 10 GeV in 2022 [18].

However, there are still a number of challenges that need to be overcome for the widespread adoption of LPAs for applications. This thesis focuses on methods to overcome limits owing to the diffraction of the drive laser pulse and the low repetition rates of currently-used laser technology, outlined below.

1.2.1 Diffraction of the drive laser pulse

The laser pulses used in an LPA need to be focused to small spot sizes to efficiently drive the plasma wave. However, tightly focused laser spots will diffract rapidly (on the scale of millimetres for typical laser parameters used), reducing the acceleration length and hence the potential energy gain of the LPA.

Plasma channels have been shown to provide a promising method of extending the length of an LPA by acting as a waveguide for the laser pulse driver (see Figure 1.3). The Oxford Laser Plasma Accelerator Group has developed a plasma channel that can be used to guide high-intensity laser pulses over many centimetres of propagation [1, 2, 19–

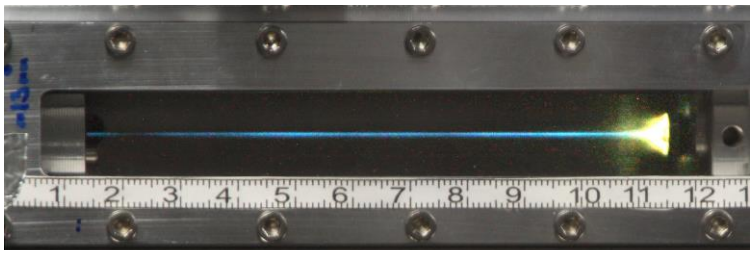


Figure 1.3: Photograph of a plasma channel used to guide a laser pulse. The image shows the time-integrated plasma fluorescence emission from a 100 mm-long plasma channel, measured by the Oxford Laser Plasma Accelerator Group [1]. The guided beam entered the channel from the left of the image.

21]. The remaining limitation on extending these channels to metre-scale propagation distances (and hence high particle energies) has been the rate of energy leakage of the guided beam from the channel [1, 20]. In this thesis, methods to improve the properties of the plasma channels, including methods to reduce energy leakage, will be presented in Chapters 4 and 5.

1.2.2 Repetition rate of the accelerator

Laser plasma accelerators typically utilise ultra-short, high-energy laser systems which are limited to repetition rates of the order $1-10$ Hz due to thermal effects in the laser medium. A future LPA will need to work at much higher repetition rates, ideally at the kilohertz level, to be competitive with conventional accelerators. For example, some cancer treatment modalities rely on high repetition rates to be medically useful [22]. Higher repetition rates would enable a greater number of users to take advantage of the LPA beams and would also help to improve the shot-to-shot stability of the accelerator, as fast feedback mechanisms could be utilised.

A promising method that has been proposed to increase the repetition rate achievable by an LPA is to use a train of laser pulses to excite the plasma wave instead of a single high-intensity laser pulse (see Figure 1.4). This regime, known as the multi-pulse laser wakefield accelerator (MP-LWFA), can achieve strong electric fields through resonant excitation of the plasma wave. MP-LWFA enables laser systems to be used that are more energy efficient and operate at higher repetition rates than the laser systems conventionally used for LPAs [3].

In 2017, the Oxford Laser Plasma Accelerators Group presented the results of the

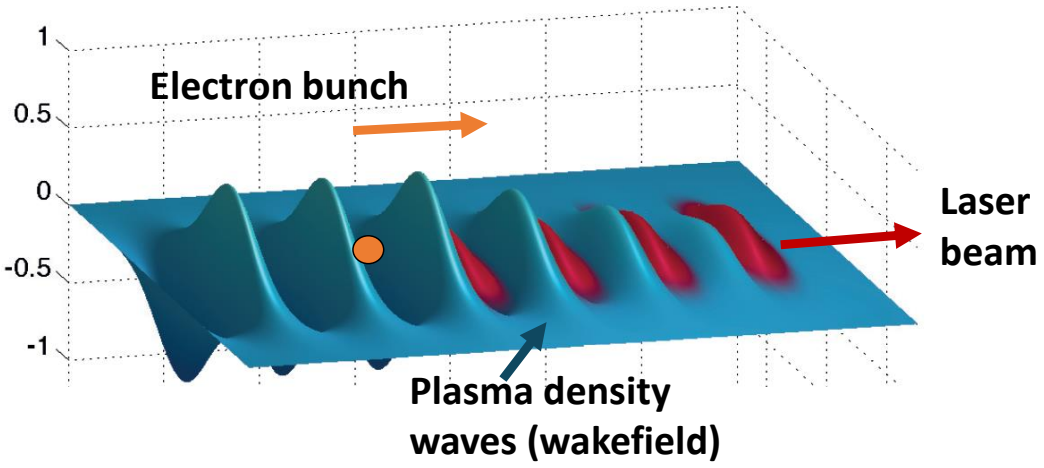


Figure 1.4: Schematic of the multi-pulse laser plasma accelerator, adapted from [3] with permission. The laser intensity envelope is marked in red, the plasma density in blue and the electron bunch in orange.

first experimental investigation of MP-LWFA [23]. The length of the accelerator in this experiment was limited by the diffraction of the drive laser beam to a few millimetres, due to the lack of waveguide structure. Recent simulation results published by the Oxford group suggest that the MP-LWFA technique could be used to drive a GeV-scale accelerator at kilohertz repetition rates, provided the accelerator stage length is on the scale of 100 mm [4]. The first experimental demonstration of the guiding of pulse trains using a plasma channel, and resonant excitation of a plasma wave within the channel, will be described in Chapter 6.

1.3 Outline of the thesis

In this thesis, investigations of relevance to future high-repetition-rate, GeV-scale laser plasma accelerators are presented. The thesis is structured as follows;

Chapter 2 - the operating principles of a laser-plasma accelerator are outlined. A literature review of previous investigations into resonant wakefield excitation and plasma waveguides is presented.

Chapter 3 - experimental methods relevant to the thesis are described, including the generation and measurement of laser pulse trains, the implementation of hydrodynamic optical-field-ionised (HOFI) plasma channels and the measurement of plasma wakefields.

Chapter 4 - a numerical optical beam propagation code developed by the author is used to study the longitudinally-extended foci required for long HOFI plasma channel formation, with suggestions of methods to improve the energy efficiency and uniformity of longitudinally-extended foci.

Chapter 5 - methods to reduce the energy losses of a guided beam as it propagates down a HOFI plasma channel are studied using optical simulations and experimental tests. The chosen parameters are investigated in plasma simulations and shown to generate improved channel properties compared to conventional methods.

Chapter 6 - the first demonstration of guiding a train of laser pulses in a plasma channel is presented, realised using an experiment performed at the Astra-Gemini TA3 facility, Rutherford Appleton Laboratory, UK. Results showing the guiding of the pulse trains as well as the resonant excitation of plasma waves in the plasma channel are shown.

1.3.1 Role of the author

The results presented in Chapter 6 had significant contributions from many scientists and engineers listed here. The experiment was performed at the Central Laser Facility, Rutherford Appleton Laboratory, UK. The experimental design and planning was led by the author, with specific contributions from Nicolas Bourgeois and the team at the Central Laser Facility on the experiment setup, Alexander Picksley on the interaction point and gas target design and Warren Wang and James Cowley on the pulse train generation and measurement. Simulations to define parameter ranges of interest and diagnostic requirements were performed by Oscar Jakobsson. The 6-week experimental campaign was led by the author, who additionally acted as Target Area Officer at the laser facility for the experiment, alongside James Cowley and James Chappell as deputy Target Area Officers. A team of experimentalists consisting of James Cowley, James Chappell, Linus Feder, Warren Wang, Oscar Jakobsson, Harry Jones and the author worked on the day-to-day running of the experiment.

The data analysis for Chapter 6 was led by the author with specific contributions from James Chappell on the pre-processing of the images of the guided mode and the calibration of the on-axis plasma density in the channel. Warren Wang performed the

analysis of the pulse train temporal profile. Simulations used to model the experiment were performed by Johannes van de Wetering, and the results of the calculations were analysed by the author, with many useful discussions with James Chappell.

The experimental data used for validation of the numerical beam propagation code presented in Chapter 4 were taken by Aarón Alejo. Parts of the work presented in Chapter 4 have been published in reference [24].

Chapter 2

Background

2.1 Laser-plasma accelerators

2.1.1 Key concepts

Laser plasma accelerators use the strong electric fields generated in the wake of a high-intensity laser pulse as it propagates through a plasma (ionised gas). The laser pulse propels electrons forward and radially outwards from the laser beam propagation axis via the ponderomotive force (see below). The plasma ions, being much heavier than the plasma electrons, remain approximately stationary on the timescales relevant for the plasma wave ($\sim$ hundreds of femtoseconds).[1] The electrons are attracted back towards the laser axis behind the laser pulse by the electric field of the plasma ions. The electrons move back towards the axis and overshoot, setting up an oscillation at a characteristic frequency known as the plasma frequency, ω_p.

The plasma frequency can be derived from simple Gauss' law considerations and is given by;

$$\omega_p = \sqrt{\frac{n_\mathrm{e} e^2}{m_\mathrm{e} \epsilon_0}} \tag{2.1}$$

where n_e is the plasma electron density and m_e is the electron mass. For plasma density $n_\mathrm{e} = 1 \times 10^{17}\,\mathrm{cm}^{-3}$, the plasma frequency can be calculated as $\omega_p = 0.018\,\mathrm{fs}^{-1}$, corresponding to a plasma period $\tau_p = 2\pi/\omega_p = 350\,\mathrm{fs}$. The plasma period is typically much longer than the period of the laser beam electric field used to drive the wake, $\tau \approx 3\,\mathrm{fs}$.

The refractive index profile of a plasma, η, is dominated by the electron number density of the plasma, n_e. Ions and atoms can also provide contributions to the refractive

[1] Ion motion can become significant when driving wakefields over longer timescales, see discussion in Section 2.2.1.

index, but their polarisability is orders of magnitude lower than that of the electrons and hence shall be neglected here [25, 26]. The refractive index of a plasma is dependent on the angular frequency of the radiation, ω, and is given by [11];

$$\eta = \sqrt{1 - \frac{n_e e^2}{m_e \epsilon_0 \omega^2}} = \sqrt{1 - \frac{\omega_p^2}{\omega^2}} < 1. \tag{2.2}$$

In a plasma, a laser pulse moves with group velocity $v_g = c\eta < c$ (unlike in dielectric media with no free charges, where $\eta > 1$ and $v_g = c/\eta < c$).

The ponderomotive force results from the gradients in the intensity of the laser pulse. The dependence of the force on the laser and plasma parameters can be found by considering a single electron oscillating in the sinusoidal field of the laser pulse, which experiences an acceleration;

$$\ddot{x} = -\frac{eE_0}{m_e} \sin \omega t \tag{2.3}$$

where ω is the angular frequency of the electric field, with amplitude E_0 (neglecting the effect of the magnetic field, which is small for non-relativistic electron motion). The initially stationary electron therefore has a velocity;

$$\dot{x} = \frac{eE_0}{\omega m_e} \cos \omega t \tag{2.4}$$

and time-averaged kinetic energy;

$$U = \langle \frac{1}{2} m_e \dot{x}^2 \rangle = \frac{e^2 E_0^2}{4 m_e \omega^2} \tag{2.5}$$

Therefore, gradients in the local laser electric field strength result in an energy gradient and hence a (conservative) force on the electrons, towards regions of lower intensity. For non-relativistic electron motion, the ponderomotive force per unit volume enacted by the laser pulse with electric field profile E_L is given by [27];

$$F_P = -n_e \nabla U = -\frac{\omega_p^2}{\omega^2} \nabla \frac{\langle \epsilon_0 E_L^2 \rangle}{2}. \tag{2.6}$$

The laser exerts a force in the transverse directions due to the transverse intensity profile, as well as on the leading edge and falling edge of the temporal intensity profile. The plasma wake generated is strongest if the duration of the laser pulse Δt is such that the front and back of the pulse are timed to be in resonance with the plasma frequency, i.e. $\Delta t \sim \tau_p/2$ [28].

The velocity of the plasma electrons repelled by the ponderomotive force can be treated as non-relativistic provided the intensity of the laser is sufficiently low. To quantify this limit, the normalised vector potential a is typically used, where $a = eA/m_{\mathsf{e}}c$ for vector potential A. For a linearly polarised laser pulse, the peak value of the normalised vector potential is given by $a_0 = 0.855\sqrt{I_{18}\lambda^2_{\mu\mathsf{m}}}$ for wavelength $\lambda_{\mu\mathsf{m}}$ in microns and peak intensity I_{18} in units of $10^{18}\,\mathrm{W/cm^2}$ [28].

For $a_0 \ll 1$, the laser pulse drives a linear wakefield which travels in phase with the laser group velocity $v_{\mathsf{wake}} = v_g$. The laser pulse drives a sinusoidal plasma density profile $\delta n_{\mathsf{e}}(\zeta) = \delta n_{\mathsf{e},0} \sin k_p\zeta$ along the laser axis, where $k_p = \omega_p/v_g$, $\delta n_{\mathsf{e},0}$ is the amplitude of the plasma density perturbation above the background plasma density $n_{\mathsf{e},0}$ and $\zeta = z - v_g t$ is the longitudinal coordinate in the frame moving at v_g [28]. The amplitude of the plasma wave can also be given in terms of the associated electric field perturbation δE, which is related to the plasma density perturbation through Poisson's equation;

$$\nabla \cdot \delta E = -\frac{\delta n_{\mathsf{e}} e}{\epsilon_0}. \tag{2.7}$$

In 1D, the solution to Poisson's equation is;

$$\delta E = \frac{\delta n_{\mathsf{e},0} e}{\epsilon_0 k_p} \cos k_p\zeta \tag{2.8}$$

$$= E_{wb} \frac{\delta n_{\mathsf{e},0}}{n_{\mathsf{e},0}} \frac{v_g}{c} \cos k_p\zeta \tag{2.9}$$

where $E_{wb} = m_e c\omega_p/e$ is termed the cold nonrelativistic wave-breaking limit and we have used $\omega_p/k_p = v_g$ and the definition for ω_p^2 from Equation (2.1). For a laser travelling in a plasma at the densities of interest, $v_g \simeq c$ and so the amplitude of the electric field $\delta E_0/E_{wb} \sim \delta n_{\mathsf{e},0}/n_{\mathsf{e},0}$. In the linear regime, the electric field and plasma density amplitudes scale linearly with the peak intensity of the driving laser pulse. The proportionality between wakefield amplitude and laser pulse intensity depends on a number of parameters, but as an example, in the case of a circularly polarised laser pulse with Gaussian envelope, $I = I_0 \exp\{-\zeta^2/L^2\}$ where L is the pulse length and $\zeta = z - v_g t$, in the linear LWFA regime ($a_0^2 \ll 1$), the maximum electric field, E_{max}, of the driven wake behind the laser pulse ($\zeta^2 \gg L^2$) is calculated as [11];

$$\frac{E_{max}}{E_{wb}} = \frac{\sqrt{\pi} a_0^2}{2} k_p L \exp\left\{-k_p^2 L^2/4\right\}. \tag{2.10}$$

In this case, when $L = \lambda_p/(\pi\sqrt{2})$, a maximum electric field of $\frac{E_{max}}{E_{wb}} = a_0^2[\pi/(2e)]^{1/2}$ is achieved.

For short, high-intensity laser drivers with $a_0 \gg 1$, the LPA operates in the non-linear regime. If the laser pulse is intense enough to expel all plasma electrons from the region near the laser axis, the LPA is said to be in the bubble regime [29]. The majority of this thesis will consider cases with $a_0 \lesssim 1$, referred to as the quasi-linear regime.

2.1.2 Electron injection

In order to be accelerated, a particle bunch must be injected into an accelerating region of the wake behind the laser pulse. There are several approaches to electron injection, which can be divided into methods that inject electrons from the background plasma wave and those that inject an external electron bunch into the preformed wakefield. The simplest injection mechanism to achieve experimentally involves using electrons that are naturally injected from the background plasma, as occurs in the bubble regime of LPA. This technique is referred to as self-injection. In self-injection, electrons can be injected into the wake at multiple positions along the accelerator, resulting in electron beams which often have a large momentum spread [30], though narrow spectra with self-injection have been achieved [13]. Self-injected beams also tend to have low shot-to-shot stability [11].

A number of ideas have been developed to achieve controlled injection of electrons from the background plasma at specific locations in the plasma wave [11]. One particularly interesting example is the use of a plasma density transition in the path of the drive laser, which locally reduces the plasma density and hence increases the plasma wavelength [31]. The electrons which would have formed the back of the plasma wave, instead fall into the accelerating region of the plasma wave and get trapped by the wake. When the length scale of the density transition is longer than the plasma wavelength, the density gradient leads to a gradual lengthening and slowing down of the wake [31]. Density transitions with a length scale $\lesssim \lambda_p$ can also trap electrons, resulting in a very short injection region [32]. The electrons in this case are all injected in a similar phase of the wake and hence the resulting beams can have a small energy spread [32].

Injection can also be achieved by using a gas with a high atomic number (Z), either alone or in a mixture with a low Z gas, to constitute the plasma. In this case, the

outer-shell electrons of the high Z gas are easily ionised and form the plasma wave. The inner-shell electrons have a higher ionisation intensity threshold and hence are only ionised when the drive laser beam reaches peak intensity, liberating the electrons into an accelerating part of the wave. This method is known as ionisation injection [33–37]. It is also possible to use additional laser pulses to inject the background electrons, in schemes collectively known as optical injection [11, 38, 39]. For example, in colliding pulse injection, a secondary laser pulse propagates in the opposite direction to the main drive laser and causes a beating in the laser intensity envelope [38]. At the peaks of the beating pattern, the high local laser intensity is able to inject the background electrons.

2.1.3 Limits on energy gain

Three fundamental limits to the maximum energy gain achievable with an LPA are given below. The equations given here are for the quasi-linear regime, $a_0 \lesssim 1$.

Dephasing

The wakefield generated by a laser pulse in a plasma travels at the group velocity of the laser pulse driver which is lower than the speed of light in vacuum, c. The electrons that are accelerated in an LPA rapidly reach relativistic energies $\gtrsim 1\,\mathrm{MeV}$ and so travel at a speed $v_\mathrm{e} \sim c > v_g$. This means that the electrons eventually outrun the plasma wave by one half period and move into a non-accelerating region of the wakefield. This happens at a length termed the dephasing length which depends on plasma density via [11];

$$L_d = \frac{\lambda_p^3}{2\lambda^2} \propto \left(\frac{1}{n_{\mathrm{e},0}}\right)^{3/2}. \tag{2.11}$$

As discussed above, the electric field of an LPA scales approximately with the cold wave-breaking field $E_{wb} \propto \sqrt{n_{\mathrm{e},0}}$. Therefore, the overall energy gain of an LPA that is limited by the dephasing length scales inversely proportional to plasma density as $W_e \sim eE_{wb}L_d \sim 1/n_{\mathrm{e},0}$. That is, lower plasma densities enable greater electron energy gains overall, but require longer propagation distances to reach the dephasing length.

Methods have been suggested to eliminate the problem of dephasing in LPAs, for example by focusing the drive beam to a longitudinally-extended focus and controlling the velocity of the beam focus using spatio-temporal couplings, termed dephasingless laser wakefield acceleration [40]. However, there are serious concerns over the energy

efficiency of such a scheme as very high laser intensities would be required along the extent of the longitudinally-extended focus to drive a wakefield, increasing the drive pulse energy requirement by approximately the ratio of the length of the extended focus L divided by the Rayleigh range of the focus of an ideal lens [40]. In the example given in the paper, to accelerate electrons to 1 GeV, 6 times more energy would be required in the dephasingless regime compared to a traditional LWFA [40]. The laser energy needs and energy efficiency of the accelerator are important properties for any future LPA for applications.

Diffraction

To efficiently drive the wakefield accelerator, the laser pulse is typically focused[2] to a spot size on the scale of the plasma wavelength $\lambda_p = 2\pi c/\omega_p$. Focused laser pulses will diffract on a distance scale which, for a laser pulse with Gaussian transverse intensity profile, is given by the Rayleigh range [11];

$$z_R = \pi w_0^2/\lambda \tag{2.12}$$

which is defined as the distance for a Gaussian beam with focal spot size w_0 to grow in size by a factor of $\sqrt{2}$. Therefore, for typical wavelengths of 1 µm and spot sizes of 40 µm, Rayleigh ranges of just a few millimetres are achievable. Methods to extend the LPA beyond the Rayleigh range are given in Section 2.3.

Depletion

As the laser beam drives the wakefield, it loses energy to the plasma wave. Eventually, assuming the dephasing length is not reached first, all energy in the laser beam will be depleted. The length over which the energy in the laser driver is completely depleted to zero is termed the depletion length, and in the linear regime is given by [11];

$$L_{pd} = \frac{2\lambda_p^3}{a_0^2\lambda^2} \propto \left(\frac{1}{n_{\mathrm{e},0}}\right)^{3/2}. \tag{2.13}$$

In the linear regime, $L_d \ll L_{pd}$ and hence the dephasing length is the limiting factor on the energy gain of the accelerator (provided diffraction is overcome). For an efficient accelerator, where efficiency is defined as the fraction of laser beam energy transferred

[2] Smaller spot sizes enable higher intensity in the focus for given peak power, but too small a spot size can generate strong transverse wakefields which are undesirable.

to the electron beam, it is desirable for $L_d \sim L_{pd}$ which is achieved in the quasi-linear regime with $a_0 \sim 1$.

It is possible to build an LPA which accelerates particles beyond the depletion length, by splitting the LPA into a series of accelerator stages, each driven by a fresh laser pulse driver, a method known as staging. Staging of LPAs is an active area of research [41, 42].

2.2 Resonant excitation of plasma waves

Modern LPAs typically use a single, ultra-short laser pulse as the wakefield driver. The invention of chirped-pulse-amplification in 1985 by Donna Strickland and Gerard Mourou [43] made it possible to generate laser pulses short and intense enough to be used to drive LPAs in the nonlinear regime, with modern Ti:sapphire CPA systems achieving $\Delta t \sim 50\,\text{fs}$ and $I_0 \sim 1 \times 10^{18}\,\text{W cm}^{-2}$. However, high-energy Ti:sapphire laser systems typically suffer from low repetition rates ($\sim 1-10\,\text{Hz}$) and low wall-plug energy efficiencies ($< 0.1\%$).

Resonant excitation of plasma waves enables the constraints on the peak intensity and duration of the pulses from the laser system to be relaxed, providing the opportunity to use alternative laser systems that may be better suited to applications. Since the invention of LPAs, a number of laser-plasma regimes have been investigated that utilise resonant wakefield excitation, including plasma beat-wave acceleration, self-modulated laser wakefield acceleration and multi-pulse laser wakefield acceleration. Each of these regimes is discussed briefly in this section.

2.2.1 Plasma beat-wave acceleration

Plasma beat-wave acceleration (PBWA) was suggested as a mechanism for accelerating electrons in the original LPA paper by Tajima and Dawson in 1979, with relevant theoretical considerations published in 1972 by Rosenbluth and Liu [44] and Cohen *et al.* [45]. In PBWA, the plasma wave is driven by a pair of co-propagating laser pulses with different frequencies, ω_1 and ω_2, which are separated by the plasma frequency, $\omega_2 - \omega_1 = \omega_p$ (see Figure 2.1(b)). The duration of each laser pulse is much longer than the plasma period. The beating between the two pulses results in an electric field modulated at the plasma frequency which resonantly excites the plasma wakefield, resulting in a wake amplitude that grows linearly towards the back of the modulated pulse

[44]. Wakefield excitation and electron acceleration with the PBWA method were experimentally demonstrated in the 1980s and 1990s [46–49]. For example, Dangor *et al*. used two laser beams with wavelengths 1.053 µm and 1.064 µm, peak intensities around $2 \times 10^{14}\,\mathrm{W\,cm^{-2}}$, 50 J energy and 200 ps pulse duration to demonstrate resonant excitation of plasma waves in hydrogen, generating an accelerating gradient of 300 MeV/m over an estimated interaction length of 5 mm [46].

One major drawback of PBWA is that the amplitude of the wakefield saturates once the motion of the electrons forming the plasma wave becomes relativistic, which happens when the amplitude of the wake electric field approaches the cold nonrelativistic wave-breaking limit. The plasma frequency, taking into account relativistic effects, is given by;

$$\omega_{p,\mathrm{rel}} = \sqrt{\frac{n_\mathrm{e} e}{\gamma_v m_\mathrm{e} \epsilon_0}} \tag{2.14}$$

where $\gamma_v = 1/\sqrt{1 - v^2/c^2}$ is the Lorentz factor of the electrons forming the wake. This means that, as the wake amplitude grows with time, the beating pattern of the co-propagating laser beams, which is fixed by the difference in laser wavelengths, will eventually fall out of resonance with the plasma wave and will no longer effectively drive the plasma wakefield. This is known as relativistic detuning. The maximum wake amplitude that can be achieved by a PBWA due to relativistic detuning is termed the Rosenbluth-Liu limit [44].

One method that could be used to overcome the Rosenbluth-Liu limit is to vary the modulation frequency of the beat-wave pattern so that the spacing between peaks grows with time, as suggested by Deutsch *et al*. in 1991 [50]. This could be achieved by using a frequency chirp on each of the two laser pulses in the beat wave, but this technique has not been experimentally verified.

Another difficulty with PBWA is the potential for the ions to move in response to the electric field of the wake and disrupt the wakefield properties (instead of remaining stationary as assumed so far). The characteristic time-scale for the onset of ion motion is given by $\tau_i = 2\pi/\omega_{p,i}$ where $\omega_{p,i} = \sqrt{\frac{n_i Z^2 e^2}{m_i \epsilon_0}}$ is the ion plasma frequency for an ion with atomic number Z. For Hydrogen with density $1 \times 10^{17}\,\mathrm{cm^{-3}}$, $\tau_i \approx 500$ ps. Therefore, ion motion is not significant in LPAs driven by ultra-short pulses, but can become important in the PBWA regime where pulse durations on the order of nanoseconds are used.

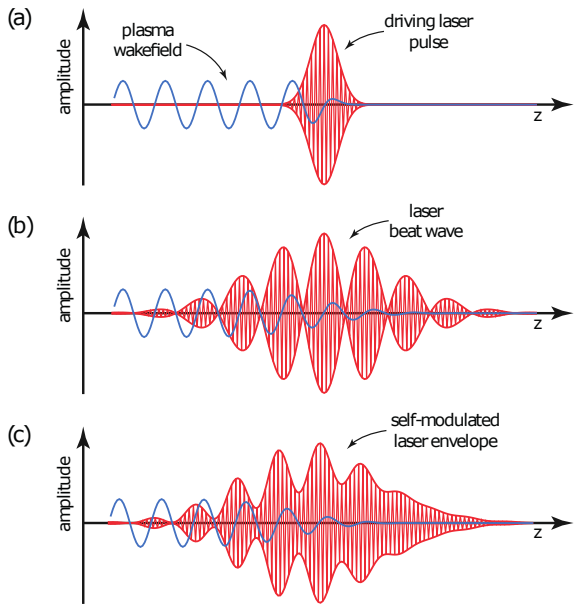


Figure 2.1: Schematics of different laser-plasma accelerator regimes, with the laser driver electric field amplitude marked in red and the plasma wakefield in blue, as a function of propagation distance z. (a) shows an LPA driven by a single ultrashort laser pulse. (b) shows the plasma beat wave regime where the envelope of the laser driver is pre-modulated at the plasma frequency. (c) shows the self-modulated laser wakefield regime, where the envelope modulation grows over time owing to the self-modulation instability. Figure taken from [28] with permission.

A final disadvantage of the PBWA scheme is that, although laser systems with the right wavelength difference exist (e.g. Nd:YAG and Nd:YLF), the repetition-rate and efficiency of these systems is often low [4]. Moreover, the long duration of the pulses generated by these laser systems ($\sim$ ns) results in a very narrow resonance width and hence makes it difficult to tune the plasma to resonance with the beat-wave (see discussion of resonance width in 2.2.3).

2.2.2 Self-modulated laser wakefield acceleration

Another method that utilises long laser pulses to drive the wakefield is self-modulated laser wakefield acceleration (SM-LWFA) [51, 52]. SM-LWFA generates the wakefield via the self-modulation instability described in Section 2.4.1. In this regime, a plasma density wave arises stochastically with frequency ω_p and acts to modulate the envelope of the laser pulse, positively reinforcing the wave in a feedback loop that leads to high-amplitude wakefield generation (see Figure 2.1(c)). Because the laser pulse is modulated at the local value of ω_p, the wake does not saturate when high amplitudes are reached and hence SM-LWFA does not suffer from the Rosenbluth-Liu limit.

SM-LWFA has been investigated in simulations and experiments [51, 53–55]. For

example, in 1994, Nakajima *et al.* [53] used a laser beam with wavelength 1.053 µm, pulse duration 1 ps, power 3 TW and peak intensity approximately 1×10^{17} W cm^{-2} to drive a SM-LWFA, accelerating electrons from 1 MeV to > 18 MeV over a distance of ~ 10 mm in a Helium gas at plasma densities $\sim 1 \times 10^{19}$ cm^{-3}.

To achieve wakefield acceleration using self-modulation, it is desirable that the drive laser pulse have a power above the critical power for self-focusing $P > P_{\text{crit}}$ where $P_{\text{crit}} = 17\omega^2/\omega_p^2$ GW (for discussion of self-focusing see Section 2.3). To avoid the need for unreasonably high laser powers P, the critical power can be reduced by increasing the plasma density. As discussed above, lower densities are desirable to reach long dephasing lengths, hence this requirement is a disadvantage of the scheme over short-pulse LPAs.

Moreover, SM-LWFA is inherently unstable as the process originates from a stochastic perturbation. It has also been found that in the SM-LWFA regime, electrons tend to be injected from the background plasma into the wave, resulting in large momentum spread beams that are less well-suited to applications. For example, simulations modelling the experiment presented by Nakajima *et al.* [53] suggest the electrons would have had a near 100% momentum spread.

2.2.3 Multi-pulse laser wakefield acceleration

Multi-pulse laser wakefield acceleration (MP-LWFA) builds upon the advantages of PBWA but in a different parameter regime, with shorter trains of pulses (~ 10 as opposed to ~ 1000) and higher peak intensity in each pulse [3]. Instead of the drive laser pulse being formed in the beating between two long laser pulses of different wavelengths, MP-LWFA uses a train of short laser pulses with (in principle) control over the spectra, spacing and width of each pulse. Such a pulse train could be generated, for example, by a coherent combination of ultra-short pulses from a fibre laser [3] or from modulation of a picosecond-duration laser pulse from a thin-disk laser [4] (see Section 2.2.4).

For a train of N identical laser pulses spaced by a constant time separation τ, the amplitude of the plasma wave in the linear regime grows with each pulse in the train simply as a sum of the wakefields generated by each delayed pulse, giving;

$$\left(\frac{\delta n_{\text{e},0}}{n_{\text{e},0}}\right)_N = \left(\frac{\delta n_{\text{e},0}}{n_{\text{e},0}}\right)_1 \left|\frac{\sin\left(\frac{1}{2}N\omega_p\tau\right)}{\sin\left(\frac{1}{2}\omega_p\tau\right)}\right| \tag{2.15}$$

which is maximally equal to $N\left(\frac{\delta n_{e,0}}{n_{e,0}}\right)_1$ when $\tau = \tau_p = 2\pi m/\omega_p$ for $m \in \mathbb{Z}$.

The width of the resonance in the ideal case can be approximated by considering a systematic offset of the pulse spacing from the resonant spacing $\tau = \tau_p + \delta\tau$ [56]. The sinusoidal behaviour owing to the numerator in Equation (2.15), $|\sin(\frac{1}{2}N\omega_p\tau)|$, means that when $\delta\tau \neq 0$ there is a maximum number of pulses N_{max} beyond which additional pulses will no longer increase the wake amplitude. To find N_{max}, we calculate the maximum of $|\sin(\frac{1}{2}N\omega_p\tau)|$ with respect to N, which occurs at $N_{\text{max}} = \tau_p/(2\delta\tau)$. When $N = N_{\text{max}}$, Equation (2.15) gives $\left(\frac{\delta n_{e,0}}{n_{e,0}}\right)_N \approx 2N/\pi$ where we have assumed $\delta\tau/\tau_p \ll 1$. That is, the ratio of the imperfect case $\delta\tau \neq 0$, to the perfect case $\delta\tau = 0$, is $2/\pi \approx 64\%$. Therefore, to excite a wakefield effectively (achieving a wake amplitude $> 64\%$ of the perfect case) requires the deviation in pulse spacing to satisfy $\delta\tau/\tau_p < 1/(2N)$. Or equivalently, the deviation in the plasma density from the resonance condition for a given pulse spacing ($n_e \propto \tau_p^2$) needs to be $\delta n_e/n_e < 1/N$.

MP-LWFA was studied theoretically and in simulations in the 1990s [57–64]. The first experimental demonstration of MP-LWFA was published in 2017 by the Oxford Laser Plasma Accelerators Group [23] (see Figure 2.2). In this experiment, a Ti:sapphire laser pulse with central wavelength 800 nm, energy 170 mJ and pulse duration 1.3 ps was modulated into a train of 7 pulses via a Michelson interferometer and used to drive a plasma wakefield over a length of 3 mm in hydrogen gas at a plasma density of $n_e \approx 1 \times 10^{18}\,\text{cm}^{-3}$. The wake amplitude was measured using the phase accumulated by a co-propagating probe beam, calculated as $\delta n_e/n_{e,0} \sim 1\%$. For a discussion of the methods used to generate and measure the pulse train as well as the method used to measure the wake amplitude from the probe beam data, see Chapter 3.

MP-LWFA has a number of benefits. Firstly, MP-LWFA enables LPAs to utilise laser technologies that achieve the required pulse energy and repetition rates, but from which the pulses are too long to drive the wake directly. Moreover, compared to laser wakefield acceleration driven by a single ultra-short laser pulse, the peak intensity required in each pulse in the train is reduced, lessening the risk of laser damage to the optical components in the laser chain.

Compared to PBWA, MP-LWFA allows for control over the spacing between the pulses in the train. This additional flexibility can be used to maintain the resonance between the train and the plasma wave in the relativistic regime, enabling wakefield

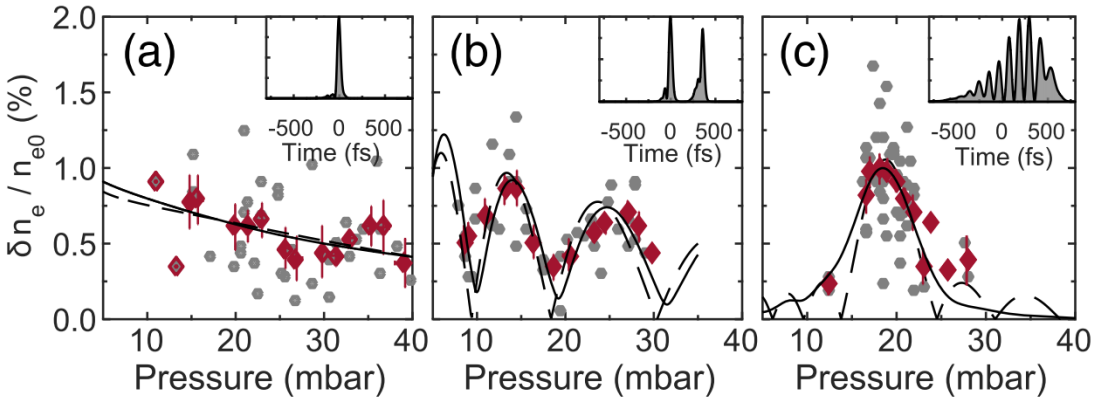


Figure 2.2: Results of the experiment into MP-LWFA conducted by the Oxford Laser Plasma Accelerators Group at the Astra-Gemini TA2 facility, published in 2017 [23]. The measured wakefield amplitude is plotted as a function of pressure for three different laser pulse drivers; (a) for a single laser pulse with energy 270 mJ; (b) for 2 pulses and energy 160 mJ; (c) for 7 pulses and energy 170 mJ. The pulse spacing in (b) was (420 ± 20) fs and in (c) was (112 ± 6) fs. The y-axis is the same for all plots. The insets show the driving pulse trains for each case. Grey dots are individual data points, red diamonds are the data averaged over pressure bins, the dashed line shows fits to Equation (2.15), and the solid line shows the wakefield amplitude calculated for the pulse trains in the inset. Figure taken from reference [23] with permission.

amplitudes to increase beyond the Rosenbluth-Liu limit. Because the trains of laser pulses tend to be shorter than in PBWA, it is easier to tune the pulse train to the resonance condition of the plasma, as the resonance width $\delta n_\mathrm{e}/n_\mathrm{e} \approx 1/N$ is greater.

The shorter pulse trains in MP-LWFA also mean that ion motion is unlikely to have a deleterious effect. Recent simulation and experimental studies conducted by the Oxford Laser Plasma Accelerators Group suggest that the wakefields driven in a hydrogen plasma decay due to ion motion over a timescale of 100 plasma periods [65, 66]. This means that an MP-LWFA can be used with pulse trains up to 100 pulses long, which is greater than the 10s of pulses needed to achieve GeV-scale acceleration [3, 4, 66].

The MP-LWFA scheme can also be adapted to enable later pulses in the train to gain energy from the wake, as a route towards the removal and recycling of residual energy in an accelerator stage after the electron has been accelerated. Currently, the efficiency of energy transfer from the plasma wake to the electron beam is $\sim 20\%$ and hence, for every 1 J of energy given to an electron beam, ~ 1 J is deposited in the plasma [67]. In the few microseconds after the acceleration of the electron beam, the energy in the wake will dissipate as heating of the plasma. If the LPA stage is operated at high repetition rates over prolonged periods of time, as is necessary for many applications, the cumulative heating of the plasma could affect the quality of the stage [68]. Therefore,

there is interest in developing techniques to remove this energy from the wake. The first steps towards energy recovery using MP-LWFA techniques were presented by Cowley *et al.* in 2017 [23], and the results can be seen in Figure 2.2 (b). Here, two pulses with a fixed spacing of (420 ± 20) fs interacted with the plasma, for a range of different gas pressures. The first laser pulse was used to drive a wake. Then, for gas pressures that satisfied $\tau = (m + \frac{1}{2})\tau_p$, the second laser pulse reduced the amplitude of the wake, thereby removing energy from the accelerator (the photons in the second pulse get blue-shifted). If the energy removed from the plasma by the laser pulse could be recycled for other useful applications, the overall energy efficiency of an LPA would also be increased.

Because MP-LWFA enables the excitation of plasma waves at low plasma densities and low a_0, additional techniques are required to maintain the laser pulse spot size over long propagation lengths (see discussion of waveguides in Section 2.3). Also, because electrons are not self-injected from the background plasma, controlled electron injection mechanisms are required [3].

2.2.4 The plasma-modulated plasma accelerator regime

A new regime for resonant wakefield excitation has recently been proposed by the Oxford Laser Plasma Accelerators Group [4]. The regime has been termed plasma-modulated plasma acceleration (P-MoPA). Advances in thin-disk laser technology mean that joule-scale laser pulse energies can now be generated at kilohertz repetition rates, in pulses of picosecond duration. The P-MoPA scheme provides a method of modulating the picosecond pulses into a train of short laser pulses that are resonant with the plasma frequency and using them to accelerate electrons to GeV energies.

P-MoPA consists of two plasma stages, a modulator stage and an accelerator stage (see Figure 2.3). The stages comprise plasma channels with on-axis densities of the order $1 \times 10^{17}\,\mathrm{cm}^{-3}$. The on-axis plasma density of the two stages must be the same to achieve the resonance condition. For information on plasma channels, see the discussion in Section 2.3.

In the modulator stage, a low-amplitude wakefield is introduced into the plasma ahead of the picosecond-long drive pulse by a short-duration, low-energy seed pulse. The drive pulse then co-propagates with the low-amplitude wakefield and is modulated spectrally (for details, see Section 3.1.2). Unlike SM-LWFA, the wakefield in the modulator stage

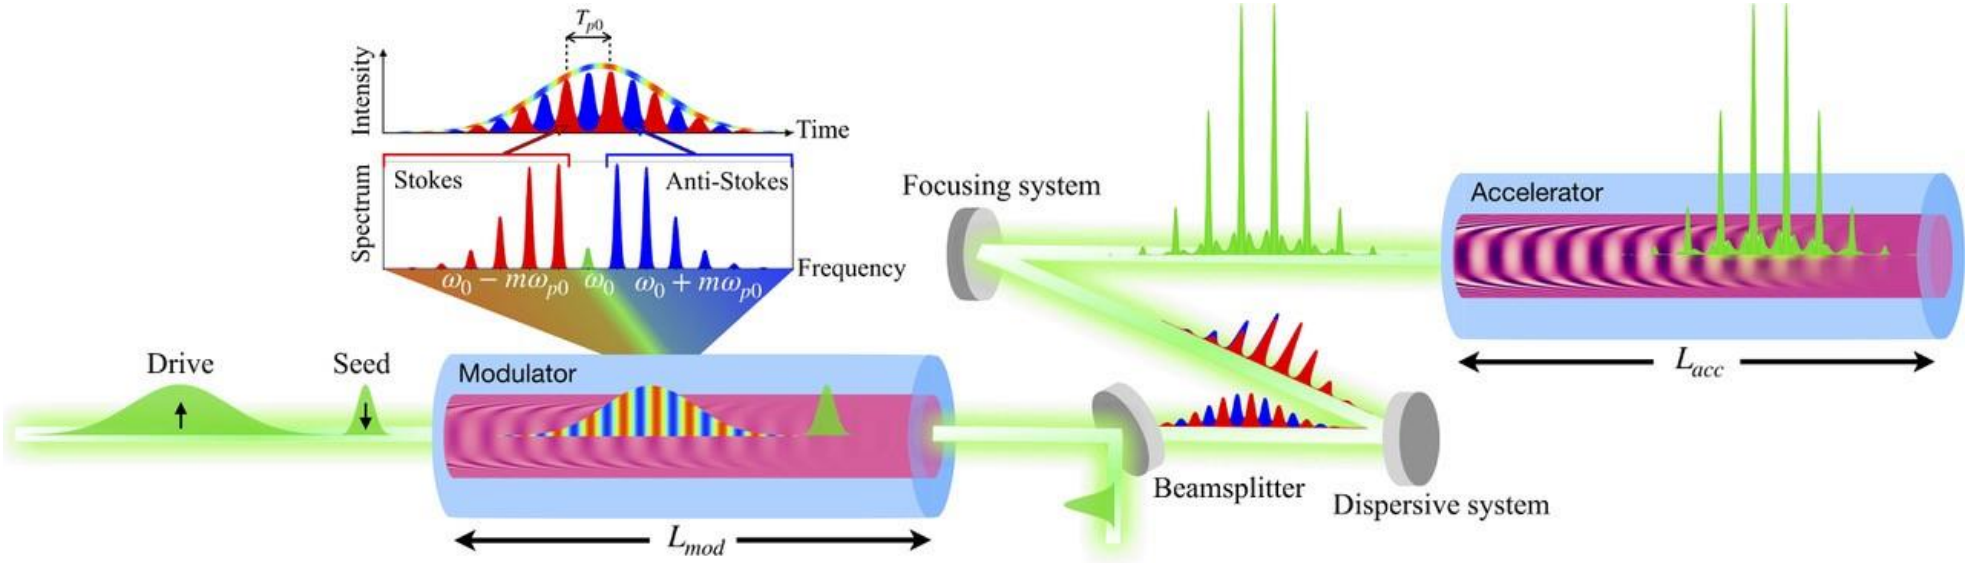


Figure 2.3: Schematic of the plasma-modulated plasma accelerator method, taken from [4] with permission. The low energy, short-pulse seed pulse generates a low amplitude wakefield in the modulator stage. The long drive laser pulse is spectrally modulated by the wakefield, resulting in a series of sidebands in the spectral domain separated by ω_p (referred to as Stoke and anti-Stokes peaks for the red and blue shifted peaks respectively). The seed pulse is removed via a polarising beamsplitter. The spectral modulation in the drive pulse is converted into a temporal modulation via a dispersive system, which modifies the relative phase between the blue-shifted and red-shifted components generated in the modulated long pulse. The drive laser beam is thereby converted into a train of short laser pulses separated by the plasma period which can be used to generate a high amplitude wakefield in the accelerator stage.

is maintained at a low amplitude and does not trigger the self-modulation instability, so the laser envelope of the drive pulse is not modulated temporally but remains temporally smooth [69]. The drive beam is then directed to a chirped optic following the modulator stage which applies spectral phase to the beam, converting the spectral modulations into temporal modulations (for details, see Section 3.1.2). It can be shown that this process generates a pulse train with spacing equal to the plasma period of the modulator stage.

The train of pulses is coupled into the plasma channel of the accelerator stage and excites a high-amplitude wakefield over the length of the stage, through the theoretical considerations outlined for MP-LWFA in Section 2.2.3. Electrons are externally injected into the wakefield and accelerated to GeV-scale energies.

Simulation results presented in the paper by Jakobsson *et al.* showed that a 1.7 J, 1 ps-long driver can be modulated into a train of ~ 10 pulses in a modulator stage of length 120 mm using the P-MoPA technique [4]. The pulse used to seed the modulation in this case was modelled with 140 mJ energy and 40 fs pulse duration. The pulse train was shown in the simulation to drive a wakefield amplitude of $\left(\frac{\delta n_{\mathrm{e},0}}{n_{\mathrm{e},0}}\right) \sim 15\%$ in a hydrogen plasma with on-axis density $2.5 \times 10^{17}\,\mathrm{cm}^{-3}$. A test electron bunch injected into the wake was accelerated to 0.65 GeV over 100 mm in the accelerator stage [4].

P-MoPA therefore provides a potential route for generating GeV-scale electron energies at kilohertz repetition rates, by exploiting state-of-the-art thin-disk laser systems.

2.3 Guiding of lasers pulses

As discussed above, diffraction limits the accelerator length and hence energy gain achievable with an LPA. To overcome diffraction, methods are required to maintain the small spot size of the laser pulse over longer propagation distances.

In a standard dielectric optical waveguide, a beam is confined due to the variation of the refractive index of the guide. For example, in a gradient refractive index fibre, the transverse refractive index profile acts like a lens which exactly counters the tendency of the pulse to diffract. A similar idea can be used in a laser-plasma interaction. As discussed above (Equation (2.2), the refractive index of a plasma is given by;

$$\eta(r) = \sqrt{1 - \frac{n_{\mathrm{e}}(r)e^2}{\gamma_v(r) m_{\mathrm{e}} \epsilon_0 \omega^2}} \tag{2.16}$$

where ω is the angular frequency of the laser radiation. By tailoring the profile of the plasma density so that $n_{\mathrm{e}}(r)$ has a minimum on the laser axis, corresponding to a maximum in refractive index on-axis, it is possible to generate a waveguide from the plasma itself.

One method of utilising the plasma refractive index is 'self-guiding' of the laser pulse, which requires no external guiding structure. Self-guiding occurs due to the combination of relativistic self-focusing and ponderomotive self-channelling [11]. Relativistic self-focusing is the process by which the high-intensity laser pulse causes electrons along the laser axis to gain a higher energy quiver motion and hence develop a greater γ_v-factor than electrons further from the axis. This results in a refractive index profile that is peaked on axis. Ponderomotive self-channelling occurs due to the expulsion of electrons from the vicinity of the laser axis as a result of the ponderomotive force, locally reducing n_{e}, which again causes a peak in η on axis.

Self-guiding requires that the power of the laser pulse is above the critical power for self-focusing, $P_{\mathrm{crit}} = 17(\omega^2/\omega_p^2)$ GW. As with SM-LWFA, satisfying $P > P_{\mathrm{crit}}$ requires either very high peak powers or high plasma densities, the latter resulting in shorter dephasing lengths and the possibility of self-injection of electrons, with high energy spread. For $n_{\mathrm{e}} = 1 \times 10^{17}\,\mathrm{cm}^{-3}$ and $\lambda = 800\,\mathrm{nm}$, the critical power is $\sim 300\,\mathrm{TW}$. For $n_{\mathrm{e}} = 1 \times 10^{18}\,\mathrm{cm}^{-3}$, the critical power is only $\sim 3\,\mathrm{TW}$.

An alternative method to self-guiding is to pre-form the guiding refractive index profile by shaping the transverse plasma density profile ahead of the arrival of the main drive laser beam. Because plasma evolution occurs on the nanosecond timescale, it is possible to form the channel using a separate laser or electron-current timed a few nanoseconds ahead of the drive beam. The plasma density profile of a plasma channel is often approximated as parabolic, given by;

$$n_{\mathrm{e}}(r) = n_{\mathrm{e},0} + \Delta n_{\mathrm{e}} \left(\frac{r}{R}\right)^2 \tag{2.17}$$

where $n_{\mathrm{e},0}$ is the on-axis plasma density of the waveguide and $n_{\mathrm{e},0} + \Delta n_{\mathrm{e}}$ is the plasma density at some radius $r = R$.

There are different types of plasma waveguides, including capillary discharge waveguides, hydrodynamic inverse bremsstrahlung plasma channels and hydrodynamic optical-field-ionised plasma channels. Each type is discussed briefly in this section.

2.3.1 Capillary discharge waveguides

In the gas-filled capillary discharge waveguide, an electric current is generated between two electrodes at either end of a gas-filled capillary tube, ionising and heating the gas on-axis. The resulting hot plasma core expands and evolves into a density profile with an on-axis minimum [16, 17, 70–72]. These waveguides can be used to guide laser pulses over many multiples of the Rayleigh range. Results of the first demonstration of guiding of laser pulses with intensities above $1 \times 10^{17}\,\mathrm{W\,cm^{-2}}$ in a hydrogen-filled capillary discharge waveguide were published in 2002 by the group at the University of Oxford [73], where propagation distances of 50 mm were achieved (equivalent to 42 times the Rayleigh range).

One major difficulty with using a capillary discharge waveguide is the degradation of the capillary itself due to the close proximity of the capillary material to the high-intensity laser-plasma interaction. The drive laser pointing has to be very stable in order to avoid damaging the entrance of the capillary. Moreover, at kilohertz repetition rates, the capillary can heat up over time, resulting in erosion of the capillary walls. Methods suggested to mitigate heating include water-cooling the capillary [74].

Capillary discharge waveguides can also be used with an additional laser pulse which comes before the main drive beam to cause further heating of the plasma. This laser-assisted method enables lower on-axis plasma densities to be achieved [75]. An LPA using this technique was able to guide laser pulses with peak intensities of $\sim 8 \times 10^{18}\,\mathrm{W\,cm^{-2}}$ over 20 cm (15 Rayleigh ranges) with on-axis plasma densities $\sim 3 \times 10^{17}\,\mathrm{cm^{-3}}$ [17]. In the same experiment, the LPA was shown to accelerate electron beams to energies of 7.8 GeV [17].

2.3.2 Hydrodynamic inverse bremsstrahlung plasma channels

Hydrodynamic inverse bremsstrahlung plasma channels are a type of freestanding plasma channel which do not require a capillary structure [76–81]. Instead of an electric current, a laser pulse is used to ionise and heat the gas along a longitudinally-extended focus (for details of longitudinally-extended foci, see Chapter 4). Within a few nanoseconds, the hot plasma expands radially outwards from the axis and evolves into an approximately parabolic plasma density profile, in which the drive laser pulse can be guided.

Hydrodynamic inverse bremsstrahlung (IB) plasma channels developed by researchers at the University of Maryland were the first pre-formed plasma waveguides to be demonstrated [76]. A focused laser beam ionises the plasma along the laser axis via multi-photon ionisation and then heats the liberated electrons through inverse bremsstrahlung absorption [79]. In the IB process, the coherent electron motion in the laser electric field is converted to thermal energy via collisions with other electrons and ions. The pressure differential between the hot plasma and the cold gas surrounding it results in a blast wave forming. The blast wave propagates away from the laser axis, leaving a minimum plasma density on-axis.

In one experiment [81], the plasma channel was formed using a laser pulse with energy 500 mJ, pulse duration 100 ps and peak intensity at focus $5 \times 10^{13}\,\mathrm{W\,cm^{-2}}$. The channel was formed in a gas mixture of helium and nitrous oxide over a distance of 1 cm and guided a laser pulse with energy 10 mJ, pulse duration 90 fs and peak focus intensity $5 \times 10^{15}\,\mathrm{W\,cm^{-2}}$. On-axis plasma densities in such channels were measured by the same group to be $\sim 5 \times 10^{18}\,\mathrm{cm^{-3}}$ [79].

The IB heating mechanism is strongly dependent on plasma density and becomes ineffective at producing plasma channels with on-axis plasma densities below $\sim 1 \times 10^{18}\,\mathrm{cm^{-3}}$ [1]. This restriction severely limits the dephasing lengths and hence energy gains that would be achievable with an LPA using a hydrodynamic IB plasma channel.

2.3.3 Hydrodynamic optical-field-ionised plasma channels

Hydrodynamic optical-field-ionised (HOFI) plasma channels utilise a mechanism for ionisation and heating of the plasma column that is independent of plasma density. Optical field ionisation both ionises and heats the electrons instantaneously (see Section 3.2.2). The hot plasma then expands radially outwards as a blast wave in the same way as the hydrodynamic IB method. Since optical field ionisation is independent of density, it allows HOFI plasma channels to be generated with much lower on-axis plasma densities than are achievable with hydrodynamic IB plasma channels [19, 82, 83].

A recent experiment [1] demonstrated that HOFI plasma channels could be generated with on-axis plasma densities as low as $7 \times 10^{16}\,\mathrm{cm^{-3}}$ and used to guide laser pulses with peak input intensity $6 \times 10^{17}\,\mathrm{W\,cm^{-2}}$ over 100 mm. The channel was formed by a

beam with pulse duration 40 fs, pulse energy 160 mJ focused to a peak axial intensity of $7 \times 10^{15}\,\mathrm{W\,cm^{-2}}$ using an axicon lens.

Hydrodynamic IB and HOFI plasma channels are typically leaky waveguides, resulting in the guided laser pulse energy radiating away over time as the pulse propagates down the channel. One way to quantify the energy leakage is via the power attenuation length of the channel, which is the distance over which the power in the guided beam falls by a factor of e. The attenuation length in the experiment described above was ~ 100 mm [1]. To generate LPAs on the metre scale, attenuation lengths much greater than 1 m are required.

Methods have recently been proposed to increase the attenuation length of HOFI plasma channels using a process known as conditioning [21, 84] (for details see Section 3.2.5). In this method, a secondary laser pulse deepens the plasma channel ahead of the arrival of the drive laser pulse, and this deepening significantly reduces energy leakage from the channel. Calculations suggest that attenuation lengths on the order of 20 m can be generated using the conditioning technique [21].

2.4 Laser-plasma instabilities

Laser-plasma instabilities can arise in an LPA via a number of mechanisms. An instability is caused by a positive feedback loop which amplifies noise in the laser-plasma interaction to a level where it can affect the wakefield or laser beam driver.

One objective of this thesis is to experimentally demonstrate the guiding of laser pulses in a plasma channel and the resonant excitation of plasma waves within the channel. There were concerns that laser-plasma instabilities might arise during the interaction and have a detrimental effect on the plasma wakefield or the propagation of the laser, or both. Laser-plasma instabilities that have been found to be significant in LPAs are outlined here.

2.4.1 Self-modulation instability

The (envelope) self-modulation instability occurs when a laser pulse with a duration longer than the plasma period undergoes envelope modulations at the plasma frequency. To see how this occurs, we use the equation for the refractive index of a plasma in the

presence of a small density perturbation $\delta n_{\mathrm{e},0}/n_{\mathrm{e},0}$, which for a laser pulse with angular frequency ω is given by;

$$\eta = 1 - \frac{1}{2}\frac{\omega_p^2}{\omega^2}\left(1 + \frac{\delta n_{\mathrm{e},0}}{n_{\mathrm{e},0}}\right) \tag{2.18}$$

where we have assumed $\omega_p \ll \omega$, $a_0 \ll 1$ and we have ignored spectral shifts of the laser pulse [85]. The group velocity of an electromagnetic wave in a plasma is given by $v_g = c\eta$ and the phase velocity is given by $v_\phi = c/\eta$.

As discussed in the section above on SM-LWFA, the self-modulation instability starts with a small density perturbation at the plasma frequency that arises stochastically, $\delta n_{\mathrm{e}} = \delta n_{\mathrm{e},0} \sin(k_p \zeta)$ where ζ is the longitudinal coordinate in the frame of the laser. This density modulation results in a refractive index modulation through Equation (2.18). On-axis, the density modulation is greater than away from the axis, which results in a transverse variation in the refractive index. The difference in phase velocity of the laser beam on-axis versus away from the axis leads to focusing of the laser beam towards the axis. These focusing forces are stronger when $|\delta n_{\mathrm{e}}|$ is greater, i.e. when $k_p\zeta$ is a multiple of π. Hence, the laser beam focuses strongest at the peaks of the plasma wave, leading to a longitudinal modulation of the laser beam envelope with a wave-vector k_p. The modulated laser beam then drives a wakefield via the ponderomotive force which is in phase with the original density perturbation, enhancing the wave and further modulating the laser envelope.

2.4.2 Laser-hosing instability

The laser-hosing instability describes the process by which small oscillations in the centroid position of the laser beam with respect to the propagation axis tend to grow over time (see Figure 2.4) [11, 85]. Consider a Gaussian laser pulse with pulse duration much longer than the plasma period, and Gaussian transverse profile described by a centre position x_{c} and spot size x_s. The normalised vector potential of the pulse can be written $a \sim a_0 \exp\{-(x - x_{\mathrm{c}})^2/x_s^2\}$.

The laser-hosing instability starts with a small centroid oscillation at the plasma frequency $x_{\mathrm{c}} = x_{\mathrm{c},0} \sin k_p\zeta$ for x the transverse coordinate with respect to the laser propagation axis and ζ the longitudinal coordinate in the frame of the laser. From the perspective of a plasma element at a fixed position on the x-axis, x_0, the laser profile

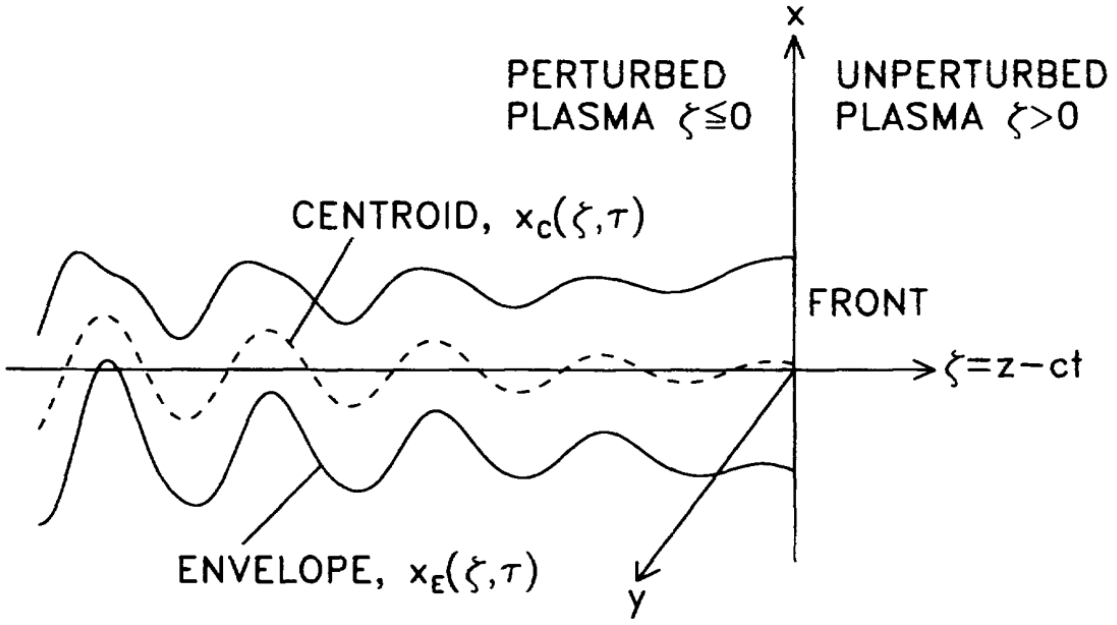


Figure 2.4: Motion of the laser centroid due to the laser-hosing instability and the corresponding variation in the laser envelope in the frame of the wakefield (which moves from left to right). Taken from Sprangle *et al*. 1994 [86] with permission.

appears to oscillate in time as;

$$a^2 \sim a_0^2 \left(1 + \frac{4x_0 x_{\mathrm{c},0}}{x_s^2} \sin k_p \zeta\right) \exp\left(\frac{-2x_0^2}{x_s^2}\right) \tag{2.19}$$

where we have expanded the Gaussian exponential for small centroid oscillations [87]. The oscillation in a drives a plasma wave at the plasma frequency along the line $x = x_0$, with amplitude $\delta n_{\mathrm{e}}(x) \sim -x_0/x_s \cos k_p \zeta$. Similarly, plasma waves are driven at positions above and below the laser axis with amplitude proportional to x, such that the waves driven in $x > 0$ and $x < 0$ are π out of phase with each other.

The perturbations in the plasma density result in perturbations in the refractive index via Equation (2.18). The laser pulse tends to focus into the regions of lower plasma density, i.e. further from the laser axis. This results in a reinforcement of the centroid oscillation and hence a positive feedback loop [87].

2.4.3 Growth rates

The equations governing the growth rate of an instability depend on the laser-plasma regime [11]. Both the self-modulation and the laser-hosing instability grow at a rate which increases with P/P_{crit}. As discussed in the context of SM-LWFAs, $P_{\mathrm{crit}} \propto 1/n_{\mathrm{e}}$ and hence the instabilities should be less significant at lower plasma densities.

Most studies into laser-plasma instabilities have been in a regime involving high plasma densities ($n_{\mathrm{e}} \sim 1 \times 10^{19}\,\mathrm{cm}^{-3}$), uniform plasma density profile (i.e. no plasma

channel) and short propagation distances. For the experiment in Chapter 6, much lower plasma densities were used ($n_e \sim 1 \times 10^{17}\,\mathrm{cm}^{-3}$) and the pulse train propagated over much longer distances, and within a plasma channel, in keeping with the P-MoPA scheme. Simulations performed by Oscar Jakobsson in [4] show that the P-MoPA accelerator stage can successfully accelerate electrons over 100 mm, suggesting that instabilities do not become significant in the parameter regime studied in the publication, at least in 2D. It is important to demonstrate that this is the case experimentally. Research into the region of parameter space for which instabilities are not a problem for P-MoPA is ongoing [69].

Chapter 3

Theory and Methods

3.1 Pulse train generation and measurement

The plasma-modulated plasma accelerator (P-MoPA) provides a scheme by which picosecond-long laser pulses from a thin-disk laser can be used to efficiently drive an LPA. In the experiment described in Chapter 6, we aimed to demonstrate the accelerator stage of the P-MoPA scheme by guiding a train of laser pulses in a plasma channel and resonantly exciting a plasma wave using the pulse train.[1] The laser available for the experiment in Chapter 6 was a Ti:sapphire system, with a bandwidth-limited pulse duration of 40 fs. Therefore, to make a train of pulses with parameters similar to that of the P-MoPA regime ($N \sim 10$, $\tau \sim 200$ fs), a different method of generating the train of pulses was required.

This section introduces the properties of ultrashort laser pulses and outlines the mechanism used in P-MoPA to generate a pulse train. The process used to mock-up a similar pulse train for the experiment in Chapter 6, and measure its properties, is then described.

3.1.1 Ultrashort laser pulses

We consider the properties of a laser pulse with a Gaussian spectrum and central frequency ω_0. The electric field of the pulse in the angular spectrum (ω) domain can be written;

$$\tilde{E}(\omega) = \tilde{E}_0(\omega_0) \exp\left\{ \frac{-(\omega - \omega_0)^2}{2\sigma_\omega^2} \right\} \exp\{\mathrm{i}\phi(\omega)\} \tag{3.1}$$

for $\Delta\omega = 2\sqrt{\ln 2}\sigma_\omega$ the full-width-half-maximum (FWHM) bandwidth of the power spectrum (spectral intensity). The spectral phase $\phi(\omega)$ can be expanded about the

[1] Investigation of the modulator stage will be carried out in a separate, future experiment.

central frequency;

$$\phi(\omega) = \phi(\omega_0) + \phi^{(1)}(\omega_0)(\omega - \omega_0) + \frac{1}{2}\phi^{(2)}(\omega_0)(\omega - \omega_0)^2 + \dots \tag{3.2}$$

where $\phi^{(1)}(\omega_0)$ is known as the group delay, $\phi^{(2)}(\omega_0)$ is the group delay dispersion (GDD) or quadratic phase, $\phi^{(3)}(\omega_0)$ would be the third order delay (TOD), $\phi^{(4)}(\omega_0)$ the fourth order delay (FOD) and so on [88].

If we assume that terms higher than second order can be neglected, we can find a simple correspondence between the properties of the pulse in the spectral domain and the properties in the temporal domain. Let the pulse in the temporal domain be written;

$$E(t) = E_0 \exp\{-\Gamma t^2\} \exp\{-\mathrm{i}\omega_0 t\} \tag{3.3}$$

where $\Gamma = a + \mathrm{i}b$ is the complex Gaussian parameter [89]. The FWHM time duration of the intensity of the pulse $\Delta t = 4\ln 2\sqrt{1/a}$. The temporal phase of the pulse $\psi(t) = -(\omega_0 t + bt^2)$.

It can be helpful to define the instantaneous frequency in the temporal domain as the time derivative of the temporal phase, $\omega(t) \equiv -d\psi(t)/dt = \omega_0 + 2bt$. This can be understood as the local effective frequency in the pulse. For a positive chirp ($b > 0$), the back of the pulse (later in time) has a higher frequency (more blue) than the front.

To understand the pulse train generation, it is useful to analyse the properties of an ultrashort Gaussian pulse in the spectral domain. The spectrum of a pulse can be calculated using a Fourier transform;

$$\tilde{E}(\omega) = E_0 \frac{1}{\sqrt{2\Gamma}} \exp\left\{-\mathrm{Re}\left\{\frac{1}{4\Gamma}\right\}(\omega - \omega_0)^2\right\} \exp\left\{-\mathrm{Im}\left\{\frac{1}{4\Gamma}\right\}(\omega - \omega_0)^2\right\}. \tag{3.4}$$

Identifying the second-order term in the spectral phase, we find that the chirp on a laser pulse is equivalent to;

$$\phi^{(2)} = \frac{1}{2b}\frac{1}{1 + a^2/b^2}. \tag{3.5}$$

A useful limit to consider is $a^2/b^2 \ll 1$, which gives $\phi^{(2)} = 1/(2b)$. This limit is equivalent to the limit where the FWHM time duration of the pulse is much longer than the bandwidth-limited time duration t_{bw} (i.e. the pulse duration for a pulse with no spectral chirp, $b = 0$). The time duration of a Gaussian pulse obeys an invariant known as the time-bandwidth product $\Delta t \Delta\omega = 4\ln 2\sqrt{1 + b^2/a^2}$. The FWHM spectral bandwidth $\Delta\omega$ is unchanged by the spectral phase and hence $\Delta t = t_{bw}\sqrt{1 + b^2/a^2}$. Therefore $a^2/b^2 \ll 1$ is equivalent to $\Delta t \gg t_{bw}$.

3.1.2 Pulse train generation using plasma modulation

To understand how the P-MoPA scheme discussed in Chapter 2 forms a pulse train, we briefly cover the physics involved in the modulation of a picosecond-long laser pulse by a plasma density wave, following the arguments in the supplemental material of reference [4]. The theoretical considerations are also relevant to methods of diagnosing plasma wakefields (discussed below).

The phase accumulated by a laser pulse with central frequency ω co-propagating with a refractive index structure $\eta(r,z)$ over a distance L is given by;

$$\phi(r,L) = \frac{\omega}{c}\int_0^L \eta(r,z)\,dz. \tag{3.6}$$

We consider the case where the phase velocity of the wakefield is equal to the group velocity of the laser driver, v_g. We define the coordinate $\zeta = z - v_g t$ as the longitudinal position in the frame moving at the group velocity of the laser pulse. Assuming the laser envelope is slowly varying, it can be shown that the phase accumulation becomes [90];

$$\phi(r,\zeta) = -\frac{c^2}{2\omega}\int_0^{t_L}[k_p(r,\zeta)^2 - k_{p,0}^2]\,dt' \approx -\frac{c^2}{2\omega}[k_p(r,\zeta)^2 - k_{p,0}^2]t_L \tag{3.7}$$

where $k_{p,0} = \omega_{p,0}/c$ and $\omega_{p,0}$ is the plasma frequency of the ambient plasma. The time $t_L = L/v_g$ is the time over which the laser and wakefield are co-propagating. For a plasma with $k_p^2(r,\zeta) - k_{p,0}^2 = \delta n_\mathsf{e}(r,\zeta)\frac{k_{p,0}^2}{n_{\mathsf{e},0}}$ where δn_e is the perturbation of the plasma density from the ambient density $n_{\mathsf{e},0}$, Equation (3.7) becomes;

$$\phi(r,\zeta) \approx -\frac{\omega}{2}\frac{\delta n_\mathsf{e}(r,\zeta)}{n_\mathsf{crit}}\frac{L}{v_g} \tag{3.8}$$

where we have used $\omega_p^2/\omega^2 = n_{\mathsf{e},0}/n_\mathsf{crit}$. For the case of a sinusoidally varying wakefield with no radial dependence, the plasma density perturbation has profile $\delta n_\mathsf{e}(r,\zeta) = \delta n_{\mathsf{e},0}\sin(k_p z - \omega_p t)$. Hence we can write;

$$\phi(r,\zeta) \approx -\beta\sin(k_p z - \omega_p t) \tag{3.9}$$

where $\beta = \frac{1}{2}\frac{\omega_p^2}{\omega}\frac{\delta n_{\mathsf{e},0}}{n_{\mathsf{e},0}}\frac{L}{v_g}$ and we have used $n_\mathsf{crit} = n_{\mathsf{e},0}\omega^2/\omega_p^2$.

To understand the effect of this phase on the spectrum of the laser pulse, we consider the full form of the laser pulse electric field, $E = E_0\exp\{\mathsf{i}(kz - \omega t + \phi)\}$. We can use the Jacobi-Anger expansion on the new phase term;

$$\exp(\mathsf{i}\phi) = \exp(-\mathsf{i}\beta\sin(k_p z - \omega_p t)) = \sum_{m=-\infty}^{\infty} J_m(-\beta)\exp\{\mathsf{i}m(k_p z - \omega_p t)\} \tag{3.10}$$

where J_m are the m^{th} order Bessel functions of the first kind. Hence, the phase modulation provided by the plasma wave results in the pulse developing side-bands in the spectrum at $\pm m\omega_p$ [4].

If the redshifted side-bands ($m < 0$) could be isolated from the rest of the spectrum, they would correspond to a train of laser pulses in the temporal domain. The same is true for the blueshifted side-bands ($m > 0$) when isolated. However, these two trains are π out of phase with each other so the overall pulse envelope remains smooth. The spectral phase of the side-bands can be calculated as $\psi_m = \theta_m + mk_p z - |m|\pi/2$ where θ_m is the argument of $J_{|m|}(\beta) \approx 0$ for $\beta < x_{00}$, where $x_{00} = 3.83171$ is the first non-zero solution of $J_1(x_{00}) = 0$ [4]. The spectral phase has a term $mk_p z$ that is linear in m and hence linear in angular frequency, and a term $-|m|\pi/2$ which is not linear in m. As discussed below, a linear frequency chirp enables peaks in the spectral domain to result in peaks in the temporal domain. The P-MoPA scheme proceeds by removing the non-linear spectral phase. This can be achieved using a dispersive system such as a specialised chirped optic with a prescribed frequency-dependent phase supplying $+|m|\pi/2$ [4].

3.1.3 Pulse train generation from an ultrashort laser pulse

The method of pulse train generation used in the experiment in Chapter 6 was developed by Warren Wang and James Cowley, based on methods employed in MP-LWFA experiments performed by the Oxford Laser Plasma Accelerator Group previously [23, 91].

Consider a laser pulse with a Gaussian power spectrum and full-width-half-maximum (FWHM) bandwidth given by $\Delta\omega$ (see Figure 3.1). This pulse is propagated through a Michelson interferometer where the two arms have a difference in optical path length given by ΔX_{mich}. In the spectral domain, wavelengths for which ΔX_{mich} is a multiple of the wavelength will coherently interfere on exit from the Michelson. That is, the output from the Michelson will have peaks in the spectrum at $m\lambda = \Delta X_{\text{mich}}$ for integers m and wavelengths λ, and troughs at $(m + 1/2)\lambda = \Delta X_{\text{mich}}$. Equivalently, the peak positions can be written in terms of the optical frequency $\omega = 2\pi cm/\Delta X_{\text{mich}}$. Therefore, in the spectral domain, the peaks are spaced by;

$$\Omega = \frac{2\pi c}{\Delta X_{\text{mich}}}. \tag{3.11}$$

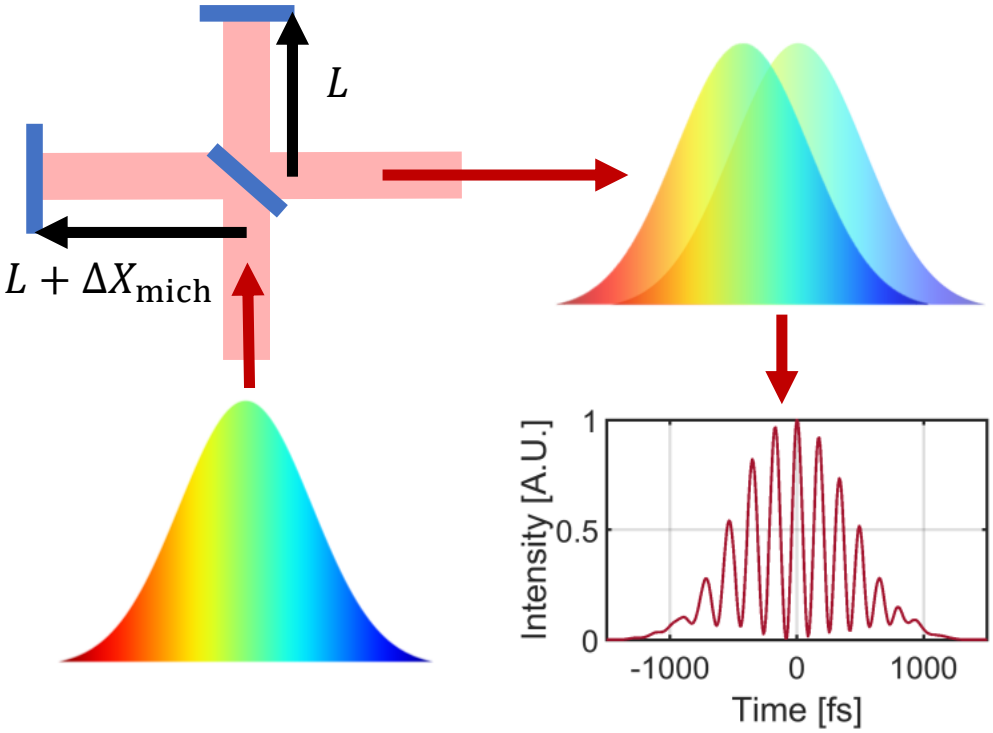


Figure 3.1: Schematic of the method used to generate a mock-up pulse train for the experiment in Chapter 6. The chirped laser pulse bottom left passes through a Michelson interferometer, with the two arms separated in optical path length by ΔX_{mich}. The interference pattern following the Michelson can be used to form a pulse train, with an example intensity temporal profile shown bottom right.

To convert the peaks from the spectral domain into peaks in the time domain, the Gaussian pulse has to have a chirp. Assuming $a^2/b^2 \ll 1$ for discussion following Equation (3.5), we find that the spacing of the pulses in the temporal domain;

$$\tau = \frac{\Omega}{2b} = \phi^{(2)}\Omega = \phi^{(2)}\frac{2\pi c}{\Delta X_{\text{mich}}}. \tag{3.12}$$

The linear relation makes it easy to identify the number of pulses in the train from the spectral information, as the number of peaks observed in the spectrum is the same as the number of peaks in the time domain. The spacing of the pulses in the spectral domain is determined by the Michelson spacing only. The spacing of the pulses in the temporal domain can therefore be controlled by setting the Michelson spacing and the GDD ($\phi^{(2)}$) of the pulse.

During the experiment described in Chapter 6, we set the GDD of the pulse using the grating compressor of the chirped pulse amplification system in the laser chain. Full compression would give zero GDD and the shortest possible pulse duration, which for the laser system used is typically ~ 40 fs. To generate the pulse train, we set the compressor for a partial compression to 1 ps FWHM pulse duration and then achieved different pulse

train spacings by varying the separation of the Michelson arms ΔX_{mich}. The overall duration of the pulse was unchanged throughout the experiment, but the number of pulses N and their spacing τ were varied, such that $N \times \tau$ was approximately constant.

Unfortunately, using the compressor away from full compression settings not only introduces GDD, but can also introduce unwanted higher-order spectral terms into the pulse, particularly third- and fourth-order delay (TOD and FOD). These terms act to change the temporal spacing of the pulses along the pulse train, so that for example the first two pulses may be closer together in time than the final two pulses. To achieve resonant excitation of plasma waves, it is important to have uniform pulse spacing throughout the pulse train. In the experiment, we compensated for the added TOD and FOD terms using an acousto-optic programmable dispersive filter (AOPDF) which can be used to tune the spectral phase terms of a pulse by small amounts. To understand how much to tune the TOD and FOD terms in practice, information about the temporal profiles of the pulse train had to be extracted. This was achieved using an autocorrelation signal, as discussed below.

3.1.4 Measuring ultrashort pulses

Measuring the temporal profile of ultrashort laser pulses and laser pulse trains with variations on the femtosecond timescale is very challenging. Standard temporal diagnostics such as photodiodes do not have the temporal resolution required. Various techniques have been developed to overcome this problem [89]. In the experiment described in Chapter 6, a single shot autocorrelator (SSA) was used in the measurement of the pulse trains. Here, the key principles of autocorrelators are described, following the explanation in [89].

Autocorrelators

An autocorrelator involves taking the pulse to be measured and splitting it into two copies which are interfered with each other. In a scanning autocorrelator, one copy of the pulse is delayed relative to the other by a time δt. The two pulses are passed through a nonlinear medium, such as a second harmonic generating crystal, which outputs light with intensity proportional to the square of the intensity of the interference pattern, $s(t, \delta t) \propto (|E(t) + E(t + \delta t)|^2)^2$. The time-integrated amplitude of the signal $s(t, \delta t)$ is

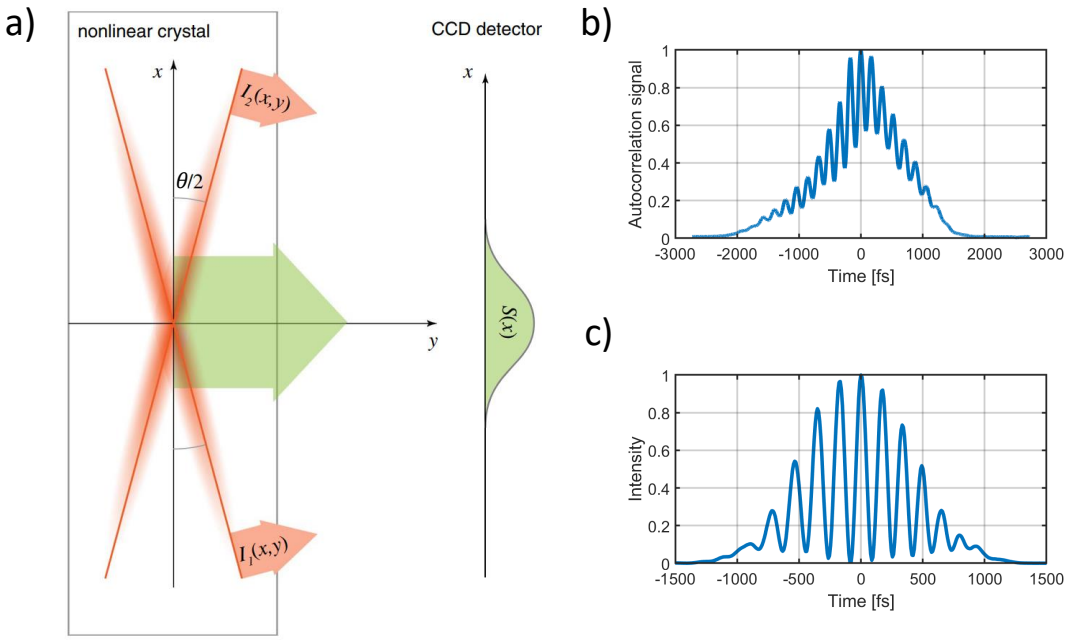


Figure 3.2: (a) Schematic of setup for single-shot autocorrelator, taken from [92] with permission. (b) Example autocorrelation signal measured during the experiment discussed in Chapter 6. (c) shows the corresponding temporal intensity profile of the pulse train generated using the retrieval algorithm developed by Warren Wang.

recorded as the autocorrelation signal $S(\delta t) = \int s(t, \delta t)dt$. The amplitude $S(\delta t)$ can be measured by placing a photo-diode after the second harmonic generating crystal. Taking measurements at different values of δt produces the autocorrelation signal profile $S(\delta t)$.

A single-shot autocorrelator uses a similar idea to the scanning autocorrelator but instead of taking multiple data points at different values of δt, a single-shot autocorrelator captures the full autocorrelation profile in one measurement (see diagram in Figure 3.2(a)). In the single-shot method, the pulse and the copy are tilted with respect to each other as they pass through the second harmonic crystal. The tilt angle introduces a correspondence between the transverse position on the crystal and the relative delay between the two pulses $\delta t = 2x \sin(\theta/2)$ for θ the angle between the two pulses and x the transverse position along the crystal. The crystal plane can then be imaged onto a camera. The signal on the camera as a function of x is proportional to the autocorrelation signal profile, $S(\delta t)$.

An example measured autocorrelation signal for a pulse train is shown in Figure 3.2(b).

Pulse train retrieval

It is not possible to extract the full complex field $E(t)$ directly from $S(\delta t)$. However, by implementing advanced retrieval methods it is possible to extract the key parameters of the field by using the SSA signal in combination with other measured information about the pulse. For the experiment in Chapter 6, Warren Wang directed efforts to develop retrieval methods to extract the parameters of the pulse trains. The process used on the experiment to calculate the temporal profile of the pulse trains is explained here.

First, the time axis of the autocorrelation signal was calibrated. This was achieved using two fully-compressed laser pulses, one from each arm of the Michelson. One Michelson arm was scanned using a motorised stage to vary the delay time ($\Delta X_{\text{mich}}/c$) between the two pulses by a known amount. SSA readings were taken at each value of ΔX_{mich} and the variation of spacing of the features on the autocorrelation signal was calibrated against the known variation in delay between the two pulses.

The spectrum of the fully-compressed pulse was measured with a spectrometer, to be used as a constraint in the retrieval algorithm. The compressor grating separation and Michelson spacing were then set to generate a pulse train, and the spectrum of the pulse train was measured. As noted earlier, the number of pulses N could be read off from the pulse train spectrum directly. The next step in the experimental setup was to reduce the TOD term in the spectral phase of the train as much as possible so that the pulse train spacing between each pair of pulses was uniform. First, a model of the grating compressor was used to estimate the TOD term expected for the grating separation.[2] This gave an estimate of the TOD value that needed to be applied with the AOPDF. The uniformity was then optimised by scanning the TOD value about the estimate from the compressor model and analysing the fringe visibility of each pair of peaks in the autocorrelation signal, $V = (I_{\text{max}} - I_{\text{min}})/(I_{\text{max}} + I_{\text{min}})$ where I_{max} and I_{min} were the maximum and minimum intensities of each peak. For a pulse train with equal pulse spacing, the autocorrelation peaks from each pair of pulses would appear at precisely the same position on the x-axis of the crystal. If the pulse spacing varies

[2] A grating compressor consists of two (reflective) diffraction gratings which are set parallel to each other and spaced by a distance called the grating separation. The input broadband stretched (chirped) beam is reflected from the gratings in series. There is one grating separation that will result in full compression of the laser pulse, leaving zero second or higher order spectral phase (in the ideal case). For other grating separations, the geometry of the optical beam paths for each wavelength in the pulse spectrum can be used to calculate the resulting GDD, TOD and FOD expected on the output pulse.

between pulses, the peaks in the autocorrelation function would be slightly offset from each other in x, resulting in a smearing of the peaks. Therefore, maximising the average value of V results in the greatest uniformity of the pulse spacing (and hence the smallest TOD and FOD terms in the spectral phase).

Assuming the TOD and FOD terms were zero after applying this method, the only remaining parameter to be determined was the average spacing of the pulses in the train $\tau = \phi^{(2)} 2\pi c / \Delta X_{\text{mich}}$. The value of ΔX_{mich} was first estimated from the spectrum of the pulse train, by fitting the modulated spectrum to the spectrum of the fully-compressed pulse multiplied by a modulation function with a period of $\Omega = 2\pi c / \Delta X_{\text{mich}}$ (from Equation (3.11)). We could also estimate $\phi^{(2)}$ from models of the grating compressor. These estimates for ΔX_{mich}, $\phi^{(2)}$ and TOD were used to give the starting points and bounds on the retrieved pulse train properties (e.g. the retrieved value for ΔX_{mich} was bounded to be within 10% of the estimated value, and the value for $\phi^{(2)}$ was bounded to be within 3000 fs^2 of the estimated value). The retrieval method progressed via a numerical optimisation algorithm (developed by Simon Hooker and Warren Wang), similar to one described in reference [23], to find the values for ΔX_{mich}, $\phi^{(2)}$ and TOD that would generate the measured auto-correlation profile $S(\delta t)$. Because the retrieval procedure is bounded to within 10% of the expected value for the experimental setup, further analysis is required to determine whether the local minimum found for the fit of $\phi(2)$ and ΔX_{mich} to the auto-correlation is also the global minimum.

The calculated values for the spectral phase terms along with the measured power spectrum amplitude give a full description of the pulse train in the spectral domain, and hence the temporal profile can be calculated via a Fourier transform. Figure 3.2(c) shows the temporal intensity profile retrieved from the autocorrelation in Figure 3.2(b) using the algorithms developed by Warren Wang.

3.2 Generating HOFI plasma channels

In this section, the formation of HOFI plasma channels is described, along with an overview of their guiding properties.

For an LPA experiment involving HOFI plasma channels, two laser beams are required, one to form the plasma channel (the channel-forming beam) and one to drive

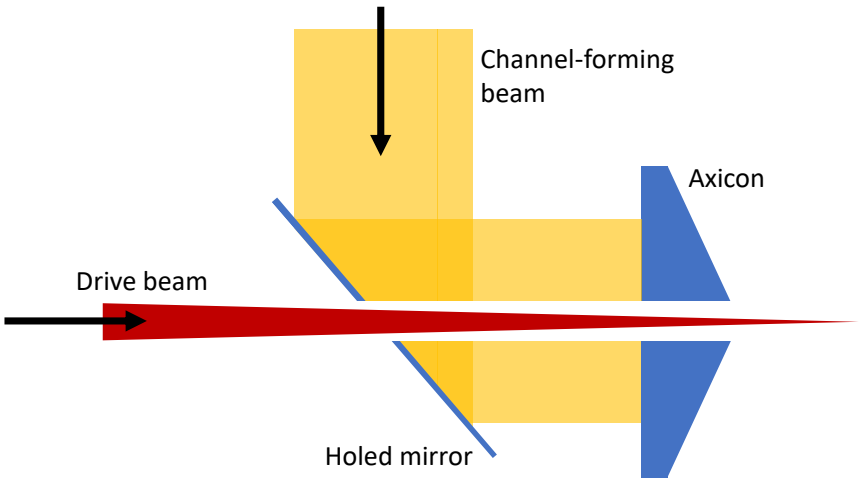


Figure 3.3: Schematic of experimental layout used to couple drive beam and channel-forming beam with a holed mirror. The axicon can be placed before or after the holed mirror depending on the experimental parameters, but if after the holed mirror, the axicon must also have a hole in the centre.

the plasma wakefield within the channel (the drive beam). The two laser beams typically originate from a common laser pulse from the oscillator that is split into two beams somewhere along the laser chain. This ensures that the two pulses are temporally synchronised. The temporal delay between the two beams is typically set as $2-3$ ns, with the channel-forming beam coming earlier in time. The delay allows the hot plasma column to expand into a transverse density profile that is useful for guiding. In practice, the time delay is chosen during the experiment by scanning the timing between the two pulses using a delay stage (with all other parameters of the experiment fixed) and finding the delay which maximises the energy in the guided beam at the end of the plasma channel.

The channel-forming beam is focused to a longitudinally-extended focus using an optic such as an axicon (more details on axicons can be found in Chapter 4). The drive beam is focused to a spot at the start of the channel using a parabolic mirror. To combine the two beams collinearly at focus, it is common to use a mirror with a hole in the centre. The channel-forming beam is reflected from the front surface of the mirror, and the drive beam is focused through the centre hole (see Figure 3.3.

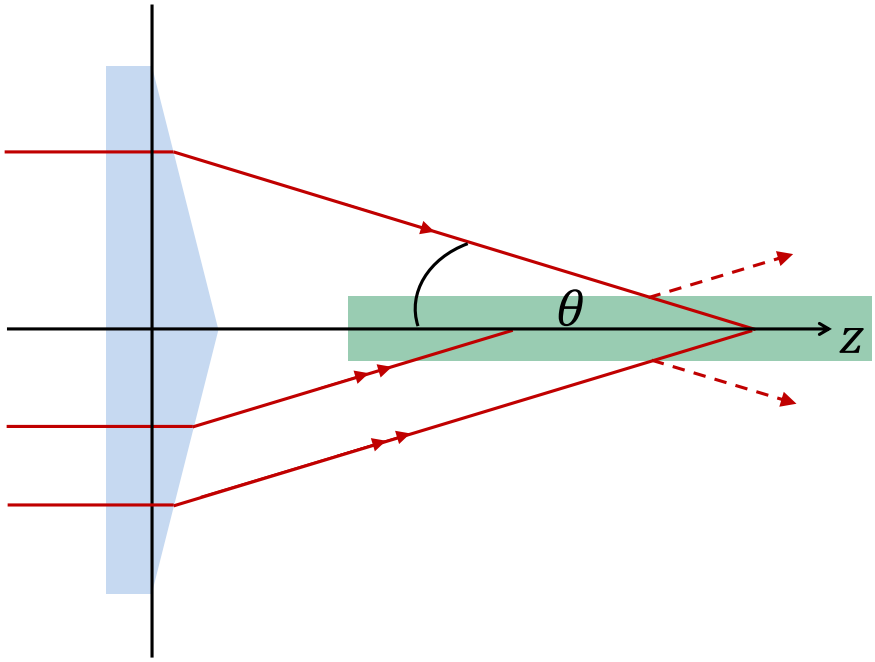


Figure 3.4: Schematic of the experimental setup for the generation of a plasma column. The focusing optic, here a refractive axicon (shown in blue), focuses rays (red) to the axis at a fixed angle θ to the laser axis. The green region indicates the plasma formed as the beam comes to focus. If $\theta < \theta_{\text{crit}}$, rays can reflect off the plasma interface and hence do not reach the laser axis, instead following the path indicated by the dashed red lines.

3.2.1 Considerations for the longitudinally-extended focus

The properties of the longitudinally-extended focus are typically set by the angle of approach of the rays to the laser axis, θ. The value that θ is allowed to take is constrained by the critical density condition of the plasma, described here. We shall consider the case where a linear axicon is used to form the longitudinally-extended focus, a conically-shaped optic which diverts all rays to the axis at the same angle θ (for more information on axicons see Chapter 4). The setup is shown schematically in Figure 3.4.

When an electromagnetic (EM) wave propagates towards a vacuum-plasma interface, it can either be reflected, transmitted or absorbed by the plasma. What happens to the wave depends on the refractive index of the plasma, which for an EM wave with frequency ω is given by;

$$\eta = \sqrt{1 - \frac{n_{\text{e}} e^2}{m_{\text{e}} \epsilon_0 \omega^2}} \tag{3.13}$$

where n_{e} is the electron density. Hence for $n_{\text{e}} > \frac{m_{\text{e}} \epsilon_0 \omega^2}{e^2}$, the refractive index of the plasma is imaginary and the reflectivity at the vacuum-plasma interface is unity. Therefore, we

define the critical density $n_{\text{crit}} = \frac{m_e \epsilon_0 \omega^2}{e^2}$.

The critical density condition for a wave incident at an oblique angle to the plasma surface is modified to account for the smaller component of the wavevector of the radiation propagating perpendicular to the plasma surface. Because the plasma refractive index is less than 1, total internal reflection can occur at the vacuum plasma interface. In this case, for the electromagnetic wave to propagate into the plasma, the following relation needs to be satisfied; $n_e < n_{\text{crit}} \sin(\theta)^2 \sim n_{\text{crit}} \theta^2$ for small angles θ [93]. This constraint can be written in terms of a critical angle of the wave to the plasma surface, below which the wave will be reflected at the interface, $\theta_{\text{crit}} = 0.54\sqrt{n_e}\lambda$ degrees for n_e in units of 10^{17} cm^{-3} and λ in μm. As θ approaches θ_{crit} from above, the rays are more and more bent away from the normal to the plasma surface due to refraction and when $\theta = \theta_{\text{crit}}$ the rays are totally reflected from the boundary.

In the case of the channel-forming beam, rays propagating towards the axis begin to ionise the plasma as they reach the axial line and come to focus (see Section 3.2.2). This generates the interface between the unionised gas ($\eta \sim 1$) and plasma, which acts back on the rays via the plasma refractive index (depicted in Figure 3.4). Hence, if the approach angle of the axicon rays goes below θ_{crit}, the rays will be reflected from the plasma interface, reducing the ionisation of the plasma on-axis. The critical density condition, therefore, puts a lower bound on the allowed values θ can take to ensure the longitudinally-extended focus effectively ionises the plasma.[3]

3.2.2 Physics of optical field ionisation

The physics involved in the ionisation and heating of plasma channels via OFI is discussed in this section.

Ionisation of plasma with high-intensity pulses

The laser beam used to form a HOFI plasma channel (referred to as the channel-forming beam) is typically taken from the same laser system as the drive laser and hence has a wavelength ~ 1 μm. At this wavelength, the energy per photon $hc/\lambda = 1.6$ eV is less than the ionisation energy of a hydrogen atom, $e_I = 13.6$ eV, or molecule, $e_I = 15.4$ eV,

[3] At slightly larger angles, the rays are expected to undergo refraction and be focused slightly further from the axicon than in vacuum. However, in previous experiments using HOFI plasma channels, axicon approach angles close to θ_{crit} have been used to successfully generate the plasma channel [1] (see Section 5.2).

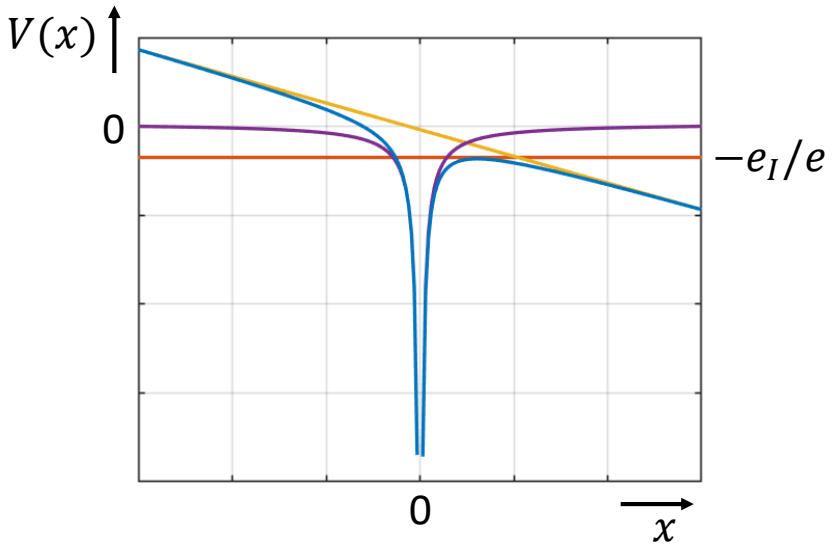


Figure 3.5: Schematic diagram of barrier suppression ionisation for an electric field polarised in the x direction. The electric potential of the electron in the atom is shown in blue and the binding potential is shown in red. The original binding potential of the atom is shown in purple and the potential due to the externally applied electric field in yellow.

and hence the photons of the beam cannot ionise the atoms directly. The ionisation instead occurs due to the strength of the electric fields present at focus.

There are a number of pathways by which OFI can occur, dependent on the intensity, optical frequency and pulse duration of the laser pulse. One way to quantify the different regimes of ionisation is using the Keldysh parameter $\gamma_k = \sqrt{\frac{e_I}{2U_p}}$ [94, 95] where U_p is the ponderomotive energy of the laser beam. For a linearly polarised laser beam with intensity I_{14} in units of $10^{14}\,\mathrm{W/cm^2}$ and wavelength $\lambda_{\mu\mathrm{m}}$ in microns, the ponderomotive energy can be calculated as $U_p = 9.3\, I_{14}\lambda^2_{\mu\mathrm{m}}$ in units of eV.

Quantum mechanical expressions of the ionisation rates are given by the Ammosov-Delone-Krainov (ADK) model [96]. For regimes with $\gamma_k \ll 1$, the system is modelled in the multi-photon regime, where the atom absorbs many photons simultaneously. For interactions with $\gamma_k \gg 1$, the system is modelled in the tunnelling regime where the electron escapes via quantum tunnelling through the potential barrier that has been partially suppressed by the electric field of the laser.

It is also possible to model the ionisation process using a classical approach, provided the electric field of the laser is greater than the binding electric field of the atom. This model is known as barrier suppression ionisation (BSI), so-called because in this regime, the electric potential of the atom is completely suppressed by the electric field, allowing

the electron to escape classically (see Figure 3.5) [97]. The threshold electric field required can be approximated in the case of a hydrogen-like atom as follows. Consider the 1D electric potential of an electron at a distance x from the nucleus of an atom, sitting in an electric field with strength E, given by;

$$V(x) = -\frac{Ze}{4\pi\epsilon_0|x|} - Ex \tag{3.14}$$

where Z is the atomic number of the atom. The potential has a local maximum value of $V_{\text{max}} = -2\sqrt{\frac{ZeE}{4\pi\epsilon_0}}$. The electron is bound to the atom with potential e_I/e. Therefore, if $V_{\text{max}} \geq -e_I/e$, the electron can classically escape the potential well [97]. The threshold electric field is therefore given by the equality condition;

$$E_{\text{th}} = \frac{\pi\epsilon_0}{Ze^3}e_I^2. \tag{3.15}$$

For hydrogen atoms with $e_I = 13.6\,\text{eV}$, the threshold electric field is calculated to be $32\,\text{GV/m}$, which corresponds to a laser intensity $I_{\text{th}} = 1.4 \times 10^{14}\,\text{W cm}^{-2}$.

For HOFI plasma channels, $I \sim 10^{14}\,\text{W/cm}^2$ and hence the BSI regime is dominant. However, tunnelling ionisation can also play a role, for example at the edges of the focus where $I < I_{\text{th}}$. Additionally, if the laser pulse is linearly polarised, the electric field strength of the laser pulse oscillates in time and at each moment, there is a certain probability of ionisation given by the tunnelling or BSI rate equations (see section on OFI rate equations below). For a circularly polarised laser pulse, the electric field strength is constant.

Hydrogen molecules

It should be noted that the theory described above has been developed for the case of hydrogen-like atoms as opposed to hydrogen molecules. For short-pulse lasers, the distinction becomes important. The energy required to dissociate hydrogen molecules into atoms is small, of the order of eV [98], however the timescales for the molecule to dissociate are longer than the timescales for OFI. Research into the ionisation of molecular gases in extreme laser fields is ongoing [99–102].

Experimental and numerical results have suggested the following dominant ionisation process of molecular hydrogen for laser conditions relevant to this thesis chapter, namely laser pulse of duration $\sim$50 fs, wavelength 800 nm and intensity $5 \times 10^{14}\,\text{W cm}^{-2}$ [102]

The first step in this process is for one electron to be liberated from the molecule via OFI processes. The ionised molecule then undergoes 'bond-softening' where the molecular bond between the two hydrogen nuclei is weakened as a result of the first electron being lost. The bond length increases up to a critical length, at which point the second electron is liberated through an effect known as 'enhanced ionisation'. The resulting doubly ionised molecular ion then undergoes Coulomb explosion. For intensities lower than $1 \times 10^{14}\,\mathrm{W\,cm^{-2}}$, it is possible for the first electron to be liberated and the bond softened, without resulting in the liberation of the second electron [102].

For tractability, a simplified model is used in this thesis, following that employed by Shalloo *et al*. [19, 56, 83]. The hydrogen molecule is treated as a hydrogen-like atom, with the ionisation energy of the system equal to the first ionisation energy of the hydrogen molecule, $e_I = 15.4\,\mathrm{eV}$, as opposed to that of the hydrogen atom $e_I = 13.6\,\mathrm{eV}$ [98]. We also assume that the ionisation rates are identical for both electrons on the molecule at the intensities of interest, so either both electrons are liberated or neither are. This simple model has been found to produce reasonable results that agree with experiment for the case of HOFI plasma channels in hydrogen gas [19].

OFI rate equations

The equations used to calculate the rates of ionisation of a hydrogen-like gas with nuclear charge Z via OFI are given in this section. These equations are used in Chapter 5 to calculate electron number density profiles resulting from different laser beam foci. The ionisation model used here is based on that of the simulation code EPOCH which includes both tunnelling and barrier-suppression ionisation, as described below [103, 104].

The ionisation rate for BSI for electric field strength E can be written [104, 105];

$$W_{BSI} = \frac{\left(\frac{e_I}{e_H}\right)^{3/2}}{4\pi Z}\left(1 - \frac{\left(\frac{e_I}{2e_H}\right)^2}{2Z\frac{E}{E_H}}\right)\chi^{-1} \tag{3.16}$$

where $e_H = 13.6\,\mathrm{eV}$ is the ionisation energy of a hydrogen atom, $E_H = 5.1 \times 10^9\,\mathrm{Vcm^{-1}}$ is the electric field in a hydrogen atom at the Bohr radius from the nucleus and $\chi = 2.42 \times 10^{-17}\,\mathrm{s}$ is the factor for converting between the Hartree and SI unit of time.

The ionisation rate in the tunnelling regime in the presence of a static electric field

can be written using the Ammosov-Delone-Krainov (ADK) model as [104, 106];

$$W_{st}(E) = 4\omega_0 \left(\frac{e_I}{e_H}\right)^{5/2} \frac{E_H}{E} \exp\left[-\frac{2}{3}\left(\frac{e_I}{e_H}\right)^{3/2} \frac{E_H}{E}\right]. \tag{3.17}$$

If the electric field is oscillatory, Equation (3.17) can be integrated over a period to give;

$$W_{ac}(E) = \sqrt{\frac{3}{\pi}} \left[\left(\frac{e_I}{e_H}\right)^{3/2} \frac{E_H}{E}\right]^{1/2} W_{st}(E) \tag{3.18}$$

For the case of a circularly polarised laser pulse, the electric field strength is a constant with time, in which case the static electric field rate is used, $W_{ADK} = W_{st}$ [96]. For a linearly polarised laser pulse, the oscillatory solution is used, $W_{ADK} = W_{ac}$.

The overall ionisation rate is then calculated in the different regimes so as to have a continuous transition from the tunnelling regime to barrier suppression. This is achieved in practice by using;

$$W(E) = \begin{cases} W_{ADK}(E) & E < E_{\text{th}} \\ \min(W_{ADK}(E), W_{BSI}(E) + W_{ADK}(E_{\text{th}})) & E_{\text{th}} \leq E < E_B \\ W_{BSI}(E) + W_{ADK}(E_{\text{th}}) & E \geq E_B \end{cases} \tag{3.19}$$

where E_B is the turning point in the ADK rate $dW_{ADK}/dE = 0$.

The rate of ionisation of hydrogen with $e_I = 15.4\,\text{eV}$ using a circularly-polarised laser pulse for different laser intensities is shown in Figure 3.6, calculated with a code written in MATLAB. The rate of ionisation is found to increase rapidly with intensity up to an intensity $I \approx I_{\text{th}} = 2.3 \times 10^{14}\,\text{Wcm}^{-2}$, above which the rate has a much weaker dependence on intensity.

Electron energy

The kinetic energy of an electron following optical-field-ionisation is initially zero. After the passage of the laser, the electron is left with kinetic energy equal to the ponderomotive energy in the laser field at the moment of ionisation, $U_p(t_0)$, which for a circularly polarised pulse is $E_K = U_p(t_0) = 9.3\, I_{14}(t_0)\lambda_{\mu\text{m}}^2$ [56, 97]. The laser intensity used in the equation is the intensity of the laser pulse at the moment of liberation t_0. As shown in Figure 3.6, the ionisation rate is greatest for intensities $I \gtrsim I_{\text{th}}$ and occurs on a timescale of femtoseconds. Therefore, most electrons are liberated at the moment when $I(t_0) \approx I_{\text{th}}$.

For a laser pulse with a peak intensity greater than the threshold, the leading edge of the pulse in time will rapidly ionise the plasma for times t such that $I(t) \gtrsim I_{\text{th}}$, and

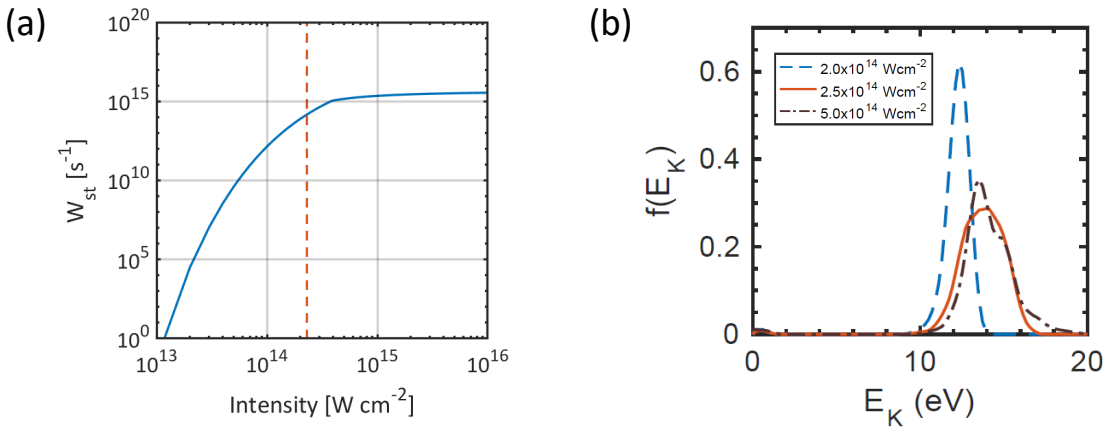


Figure 3.6: Ionisation and heating of electrons generated from optical field ionisation (OFI). (a) Calculated ionisation rates W as a function of laser intensity I for the case of a circularly polarised laser pulse incident on hydrogen gas with $e_I = 15.4\,\text{eV}$. The red dotted line indicates the BSI threshold, $I_{\text{th}} = 2.3 \times 10^{14}\,\text{Wcm}^{-2}$. (b) Electron energy distributions following ionisation by OFI, taken from [19]. The calculation was performed using the EPOCH code by Chris Arran for molecular hydrogen with density $1 \times 10^{18}\,\text{cm}^{-3}$ following propagation of a circularly polarised laser pulse, $\lambda = 0.8\,\mu\text{m}$, FWHM pulse duration 30 fs [56].

hence there will be no un-ionised atoms remaining by the time $I = I_{\text{peak}}$. This means that the electron energy dependence is dominated by the threshold intensity of the atom as opposed to the peak intensity of the laser pulse in this model, provided $I_{\text{peak}} \gtrsim I_{\text{th}}$. Therefore, the electron energy spectrum is expected to have a peak at $E_K = 9.3\, I_{\text{th}} \lambda_{\mu\text{m}}^2$ for I_{th} in units $10^{14}\,\text{W/cm}^2$. For hydrogen molecules with $e_I = 15.4\,\text{eV}$, ionised by a laser with $\lambda = 0.8\,\mu\text{m}$, the electron kinetic energy is expected to be peaked at $E_K \approx 13.5\,\text{eV}$.

The electron spectrum calculated for the case of molecular hydrogen is shown in Figure 3.6 (b) for different laser intensities, taken from [19]. These calculations were performed by Chris Arran using the simulation code EPOCH [56, 103]. The maximum energy of the liberated electrons is found to vary minimally with laser intensity, provided $I \gtrsim I_{\text{th}}$, as expected. The peak in the energy spectrum for $I = 2.5 \times 10^{14}\,\text{W cm}^{-2}$ is at $E_K \approx 13.7\,\text{eV}$, closely matching the value predicted by the simple model above (13.5 eV [56].

Following ionisation, the electron energy spectrum subsequently evolves via collisions of the plasma electrons with each other. This distributes the momentum of the electrons, which was initially in the plane of the laser polarisation, into 3 dimensions. The collisions also thermalise the electrons to form a Maxwellian distribution. The time scale for this

isotropisation and thermalisation process is the electron-electron collision time, which can be calculated using Spitzer rates [107];

$$\tau_{\mathrm{e,e}} \propto \frac{T_{\mathrm{e}}^{3/2}}{n_{\mathrm{e}} \ln \Lambda} \tag{3.20}$$

where $\ln \Lambda$ is the Coulomb logarithm with $\Lambda = (4/3)\pi n_{\mathrm{e}}(\epsilon_0 k_B T_{\mathrm{e}}/n_{\mathrm{e}}e^2)^{3/2}$. For parameters of relevance to HOFI plasma channels, we have that $T_{\mathrm{e}} = 10\,\mathrm{eV}$ and $n_{\mathrm{e}} = 3 \times 10^{18}\,\mathrm{cm}^{-3}$, giving $\tau_{\mathrm{e,e}} \approx 0.3\,\mathrm{ps}$.

The resulting Maxwellian distribution in electron energy has an effective thermodynamic temperature T_{e} roughly equal to $2/3$ of the initial peak in the electron kinetic energy distribution, where the effective temperature is the parameter defining the 3D Maxwellian distribution which in terms of electron velocity is given by $f(v)d^3v = [m_{\mathrm{e}}/2\pi k_B T_{\mathrm{e}}]^{3/2} \exp\{-mv^2/2k_B T_{\mathrm{e}}\} d^3v$. For the distribution shown in Figure 3.6(b), this gives $T_{\mathrm{e}} \simeq 9\,\mathrm{eV}$ [56]. The ion-electron collision timescale can be calculated similarly, and varies as [107];

$$\tau_{\mathrm{e,i}} \approx 5 \times 10^6 \frac{m_{\mathrm{p}}}{m_{\mathrm{e}}} \frac{A T_{\mathrm{e}}^{\frac{3}{2}}}{n_{\mathrm{i}} Z^2 \ln \Gamma_{e,i}} \tag{3.21}$$

where m_{p} is the proton mass, $n_{\mathrm{i}} = n_{\mathrm{e}}$ is the ion density, A is the mass number and Z is the atomic number of the ion and $\Gamma_{\mathrm{e,i}}$ is the electron-ion Coulomb logarithm. Here, $A = Z = 1$ and $\Gamma_{\mathrm{e,i}} \approx 6$, giving $\tau_{\mathrm{e,i}} \sim 0.5\,\mathrm{ns}$ for the parameters of interest for HOFI, as above. Hence the ions remain cold over the period of electron thermalisation.

These temperature calculations, although approximate, have been supported by simulations and experimental measurements [19].

3.2.3 Hydrodynamic expansion of hydrogen plasma

Following ionisation and thermalisation of the plasma electrons, the pressure differential between the hot plasma core and cold neutral gas surrounding it leads to the development of a cylindrical blast wave (see the example plasma density transverse profiles as a function of time in Figure 3.7 from an experiment by Shalloo et al. [19]). The dynamics of the blast wave can be approximated by the Sedov-Taylor theory. The theory assumes that the heat deposited is contained within a small volume, such that the mass of the material initially heated is much smaller than the eventual mass of material swept up by the blast wave. Additionally, the initial energy density deposited is assumed to be

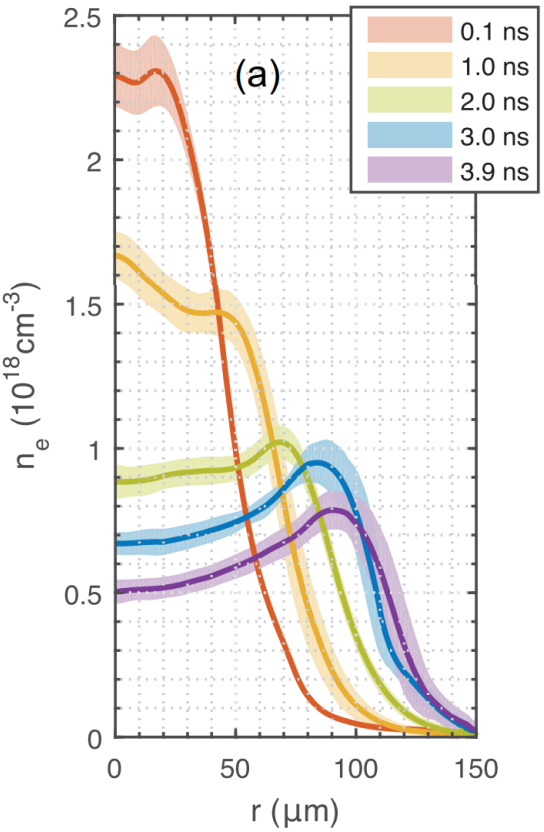


Figure 3.7: Measured electron density profiles following ionisation of hydrogen plasma by OFI at different time delays, showing the blast wave expansion. Figure taken from [19] with permission.

much larger than the total energy density of the ambient medium [108]. In the case of HOFI plasma channel generation, the initial volume of ionised plasma has a radius of $\sim 5 - 10$ µm and the resulting blast wave is expected to propagate to radial distances of the order 50 µm, hence the first condition is expected to be satisfied. The initial energy density is given by the temperature of the ionised electrons, 9.7 eV, which is considerably larger than the ambient temperature, 0.025 eV, and hence the second condition is also satisfied.

For a cylindrical blast wave, Sedov-Taylor theory predicts that the position of the shock front r_s as a function of time t is given by [108, 109];

$$r_s(t) \propto \sqrt[4]{ZE_K/m_\mathrm{i}}\sqrt{t+\tau_0}\sqrt{r_{e,0}} \tag{3.22}$$

where Z is the atomic number of the plasma species, E_K is the kinetic energy of the electrons after ionisation and m_i is the mass of the ions of the plasma species, $r_{e,0}$ is the initial radius of the column of ionised electrons and the parameter τ_0 is added to ensure $r_s(t=0) = r_{e,0}$. The proportionality constant in Equation (3.22) is of the order of 1 (dimensionless) [108].

The blast wave will stagnate when the energy in the initial column of plasma has been depleted by the work done against the background pressure of the gas. Using conservation of energy, the final radius can be calculated as;

$$r_f = (1 + \sqrt{ZE_K/k_B T_0})\, r_{e,0} \tag{3.23}$$

where T_0 is the temperature of the ambient gas, usually assumed to be room temperature [56, 109].

The amplitude of the blast wave in electron density can be calculated using the Rankine-Hugoniot relations for a strong shock [110]. These relations are the conservation equations for energy, momentum and mass at the shock front. It is found that the amplitude of a blast wave as a function of time is dependent on the velocity of the shock, the mass density of the unshocked region and the adiabatic index of the medium (where the adiabatic index is defined as the ratio of the specific heat capacity at constant pressure to the specific heat capacity at constant volume). Hence, once the initial conditions are specified, the Sedov-Taylor expression for the radius of the blast wave allows for the velocity and thus the amplitude of the shock to be estimated. In this thesis, the hydrodynamic code Helios 1D [111] is used to calculate the amplitudes of the blast waves as a function of time, which is built upon conservation equations for the fluid (see Section 5.4.1).

3.2.4 Transverse mode profiles of plasma channels

A plasma channel is found to support certain transverse electric field profiles or 'modes'. To investigate the properties of these modes, it is helpful to consider the case of an infinite parabolic plasma density profile given by;

$$n_{\mathrm{e}}(r) = n_{\mathrm{e},0} + \frac{1}{\pi r_e w_m^4} r^2 \tag{3.24}$$

where r is the radial coordinate, r_e is the classical electron radius, $n_{\mathrm{e},0}$ is the on-axis density and w_m is a characteristic length scale known as the matched spot size. The modes of the infinite parabolic plasma channel can be calculated in cylindrical coordinates and are found to take the form $E(\vec{r}) = u_{p,m}(r, \varphi) \exp(\mathrm{i}k_z z)$ where z is the laser propagation axis, φ is the azimuthal coordinate and k_z is the component of the laser

beam wavevector along the z-axis [80]. The transverse field profile is a Laguerre Gauss mode;

$$u_{p,m}(r,\varphi) = u_0 \left(\frac{2r^2}{w_m^2}\right)^{\frac{m}{2}} \exp\left(\frac{-r^2}{w_m^2}\right) L_p^m \left(\frac{2r^2}{w_m^2}\right) \exp(\mathrm{i}m\varphi) \tag{3.25}$$

where u_0 is the mode amplitude and L_p^m is the Laguerre polynomial (m and p label the radial and azimuthal modes respectively). In the infinite parabolic channel, all modes propagate without loss down the extent of the waveguide.

It is not possible to generate an infinite parabolic waveguide in practice, but the plasma channels made by hydrodynamic IB and HOFI methods are well approximated by parabolic channels with finite radial extent;

$$n_{\mathsf{e}}(r) = n_{\mathsf{e},0} + \Delta n_{\mathsf{e}} \left(\frac{r}{R}\right)^2, \qquad r < R \tag{3.26}$$

where $\Delta n_{\mathsf{e}} = n_{\mathsf{e}}(R) - n_{\mathsf{e},0}$ is the height of the plasma density wall at $r = R$ [78]. For a finite parabolic channel, energy in the guided beam can leak out during its propagation down the channel. Some modes will leak less rapidly than others and these are referred to as quasi-bound modes. The quasi-bound modes of a plasma channel can be found by considering the paraxial Helmholtz equation of the laser with angular frequency ω in a plasma with refractive index $\eta(r)$ [80];

$$(\nabla_\perp^2 + k_r^2) u(r,\varphi) = 0 \tag{3.27}$$

where $\nabla^2 = \nabla_\perp^2 + \partial^2/\partial z^2$, $\nabla_\perp^2 = \partial^2/\partial r^2 + (1/r)\partial/\partial r + (1/r^2)\partial^2/\partial\varphi^2$ and $k^2 = k_r^2 + k_z^2$ where k is the amplitude of the laser wavevector in the plasma $ck = \omega\eta$. Using $\eta(r) = \sqrt{1 - n_{\mathsf{e}}(r)/n_{\mathsf{crit}}}$ from Equation (2.2) with $n_{\mathsf{crit}} = \omega^2 m_{\mathsf{e}} \epsilon_0 / e^2$, we can write;

$$k_r^2(r,\varphi) = \frac{\omega^2}{c^2}\left(1 - \frac{k_z^2}{k^2} - \frac{n_{\mathsf{e}}(r)}{n_{\mathsf{crit}}}\right) \tag{3.28}$$

where we have ignored contributions to the refractive index from the ions as before. If for some mode, $k_r^2 > 0$ for all r, the mode can propagate through all space and hence is termed a radiated mode. If $k_r^2 < 0$ for some range of r, the mode will be quasi-bound. As the channel is finite in radial extent, the quasi-bound mode will have $k_r^2 > 0$ for some $r \gtrsim R$ and hence some beam energy can leak out and freely propagate outside of the channel, leading to a reduction in energy in the mode over time [80].

It is found that the quasi-bound modes of the finite parabolic channel are approximately the Laguerre-Gauss modes given in Equation (3.25), but the number of modes

that will propagate is limited to a subset of p, m values, given by;

$$(2p + m + 1)^2 < \Delta n_e R^2 \pi r_e. \tag{3.29}$$

In general, higher order modes have larger values of k_r^2 and hence have faster rates of energy leakage from the channel. The energy leakage rate is typically quantified using an attenuation length L_{att}, which is defined as the distance over which the power in the mode falls by a factor of e.

In practice, the experimental setup for guiding of laser pulses in a plasma channel is designed to maximise the coupling of the guided beam into the lowest-order channel mode. This is achieved by optimising the focal spot of the guided beam to be as close to a Gaussian profile as possible and aligning the focus position and pointing of the guided beam to be matched to the plasma channel axis. The delay time, t_d, between the channel forming and guided beams is scanned to maximise the energy throughput of the guided beam. This ensures that the plasma channel transverse profile is best matched to the experimental laser focus profile.

In this setup, higher-order quasi-bound modes can still be excited in the channel due to misalignment of the guided beam into the channel or imperfect matching between the laser transverse profile and the channel fundamental mode. The modes interfere with each other, generating a varying transverse beam profile as a function of longitudinal position (due to their different values of k_z).[4] The highest-order modes are radiated away most rapidly, and hence the longitudinal interference reduces with propagation distance. For sufficiently long and leaky plasma channels, the mode measured at the output of the channel is the fundamental mode of the channel.

In the ideal case for LPAs, the guided beam would be perfectly matched to the fundamental mode of the channel and the attenuation length of the fundamental mode would be $\gg 1$ m.

3.2.5 Conditioned HOFI plasma channels

Recent work has shown that the formation of the HOFI plasma channels results in a peak in neutral gas density just outside the electron plasma density blast wave. In

[4] In the context of LPAs, mode beating within the channel would likely be highly disruptive to the generation of a wakefield. Even if electrons could become trapped in the wake at one point in the propagation, they would likely be lost from the accelerating structure as the transverse shape of the wakefield varied over time.

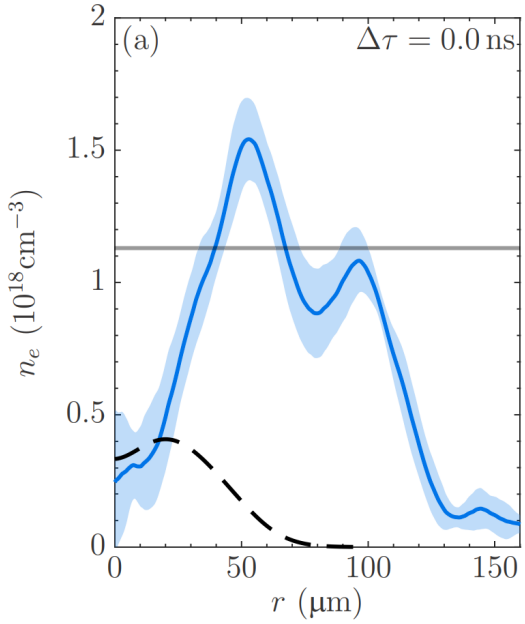


Figure 3.8: Measured electron density profile before (black dashed) and after (blue) conditioning of the HOFI plasma channel. The horizontal grey line is the density of the ambient gas before ionisation by the channel-forming beam. Figure taken from [21] with permission.

work published in 2020 by the Oxford Laser Plasma Accelerator Group, the peak in neutral gas density arises due to a 'snow-plough' effect of the blast wave as it expands radially outwards [21]. The effect of this neutral gas density collar was also observed by Feder *et al*. from the University of Maryland in 2020, where they measured the electron and neutral density profiles, showing that the amplitude of the neutral density peak is significantly greater than that of the electron density peak [112]. It is desirable to ionise the neutral density peak to create a deeper plasma channel (larger Δn_e) which more closely approximates the infinite parabolic waveguide. The ionisation of the neutral collar has been termed 'conditioning' of the plasma channel (see Figure 3.8).

Conditioned HOFI plasma channels have been experimentally demonstrated to reduce the rate of energy leakage from the plasma channel. In one experiment by the Oxford Laser Plasma Accelerator Group, the ionisation of the neutral gas collar was achieved using the guided laser pulse itself, which had sufficient intensity in the wings of the focus to ionise the gas beyond the plasma channel wall [21]. In another experiment by researchers at the University of Maryland, the conditioning pulse was formed from a part of the channel-forming beam that was shaped into a ring focus profile with radius $r \approx R$ [84]. Both experiments demonstrated that a conditioning pulse enabled the attenuation lengths of HOFI plasma channels to be extended from the centimetre-scale

to the metre-scale.

3.3 Measuring plasma wakefields

For the experiment described in Chapter 6, it was necessary to measure the wakefield amplitude generated by a train of laser pulses propagating inside a plasma channel. This presented a new challenge as wake amplitudes inside channels of this length had not previously been diagnosed.

3.3.1 Phase shifts in a probe beam

In the Oxford group's previous experiment into multi-pulse laser wakefield acceleration [23], over propagation distances of a few millimetres, the wakefield amplitude was measured using a probe-reference laser pulse pair that travelled collinearly with the wakefield-driving laser pulse. The reference pulse travelled ahead of the wake driver and the probe behind. The probe accumulated spectral phase via Equation (3.8). Interference between the probe and reference then revealed information about the phase accumulation of the probe in the presence of the wake.

The probe and reference pulses were stretched and chirped in time, allowing the probe pulse to sample many oscillations of the plasma wave. Each time slice in the probe corresponded to a different local frequency and hence temporal information about the wake could be unwrapped in the spectral domain. The wake period and amplitude were interpreted from the spectral interferograms using Fourier analysis techniques [23].

3.3.2 Spectral shifts in the drive beam

In the experiment described in Chapter 6, the probe-reference method was considered for the wakefield diagnostic. However, the incorporation of the plasma channel into the experiment in Chapter 6, meant that the probe-reference technique would have been substantially more complicated than used in the previous MP-LWFA experiment [23]. In the ideal case, the channel-forming beam, the drive beam and the probe would be collinear and the probe beam and drive beam would travel at the same group velocity. Any misalignment between the three beam-lines would result in the probe moving with respect to the wakefield as it propagates down the channel, resulting in a washing-out of the phase imprint over the centimetres of propagation.

The probe-reference pair used by Cowley *et al.* [23] was designed to have a wavelength of 400 nm so that it could be separated from the drive laser pulse (with wavelength 800 nm) using dichroic mirrors. One potential problem that could be encountered when taking this technique to longer propagation distances is that the probe-reference beams could slip in phase relative to the drive laser pulse due to the difference in the group velocities of the two different wavelengths in plasma, $v_g(\omega) = c\eta(\omega)$. In an ideal setup, the propagation length over which the pulses slip by an optical path length of $\lambda_p/2$ is given by $L_{slip} = 2\lambda_p^3/|\lambda_1^2 - \lambda_2^2|$. For $\lambda_1 = 800\,\text{nm}$ and $\lambda_2 = 400\,\text{nm}$, the slip length can be calculated for $n_\text{e} = 1 \times 10^{17}\,\text{cm}^{-3}$ as $L_{slip} = 4.9\,\text{m}$ and for $n_\text{e} = 1 \times 10^{18}\,\text{cm}^{-3}$ as $L_{slip} = 0.15\,\text{m}$. Therefore, in the experiment in Chapter 6 where densities between $1-12 \times 10^{17}\,\text{cm}^{-3}$ and propagation lengths of 100 mm were used, the phase-slippage between probe and drive beams could start to become a problem at the highest plasma densities.

The additional interaction length provided by the experiment in Chapter 6 compared to Cowley *et al.* [23] meant that signatures of the wakefield on the spectrum of the drive beam itself were enhanced. Such signatures have been used previously to diagnose a plasma wave, for example by Wojda *et al.* [113] who excited plasma waves in a 80 mm-long hydrogen-filled capillary tube and analysed the spectrum of the drive laser beam to interpret the amplitude of the wake. Therefore, in Chapter 6, the drive beam spectrum was used to diagnose the plasma wakefield instead of the probe-reference technique, thereby removing the complication of a third beam-line.

To estimate the significance of spectral shifts on the drive beam, we can again consider the phase accumulated by a beam co-propagating with a plasma wakefield which, following on from the discussion in Section 3.1.2, is given by;

$$\phi(r,\zeta) = -\frac{\omega}{2}\frac{\delta n_\text{e}(r,\zeta)}{n_\text{crit}}\frac{L}{v_g} \tag{3.30}$$

where $\zeta = z - v_g t$. We can define the (change in) instantaneous frequency in the laser pulse as the time-derivative of the (change in) temporal phase [90];

$$\Delta\omega(r,\zeta) = -\frac{d\phi(r,\zeta)}{dt} \approx \frac{\omega}{2n_\text{crit}}\frac{L}{v_g}\left(-v_g\frac{d\delta n_\text{e}(r,\zeta)}{d\zeta}\right). \tag{3.31}$$

Therefore, we can see that the gradients in the density profile result in shifting of the instantaneous frequency. This process is often referred to as photon acceleration or

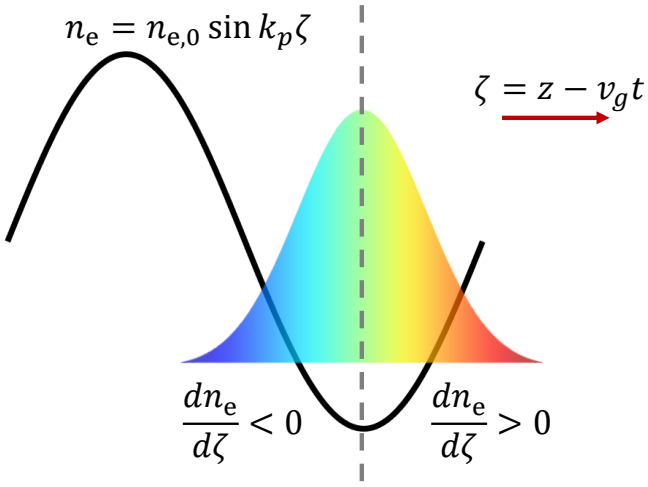


Figure 3.9: Depiction of the spectral shifts that occur within a laser pulse, that is initially spectrally flat, as it propagates collinearly with a sinusoidal plasma wakefield. The black sinusoidal curve represents the plasma density and the rainbow Gaussian profile represents the instantaneous frequency within the laser pulse envelope. The laser pulse travels from left to right and propagates at the same velocity as the plasma wave.

deceleration as the photons of the laser pulse are gaining or losing energy ($\Delta E_{\text{photon}} = \hbar\Delta\omega$) to the wake, depending on the sign of the refractive index gradient sampled by the photon.

For the case of a sinusoidally varying wakefield with no radial dependence, the plasma density has profile $\delta n_e(r,\zeta) = \delta n_{e,0} \sin(k_p\zeta)$. Equation (3.31) becomes;

$$\Delta\omega(\zeta) \approx -\frac{\omega L}{2n_{\text{crit}}} k_p \delta n_{e,0} \cos(k_p\zeta) \tag{3.32}$$

and using $n_{e,0}/n_{\text{crit}} = \omega_p^2/\omega^2$ from Equation (2.1) and the definition of the critical density, $n_{\text{crit}} = \omega^2 m_e \epsilon_0/e^2$;

$$\Delta\omega(\zeta) \approx -\frac{\omega_p^2}{2\omega}\frac{\delta n_{e,0}}{n_{e,0}} k_p L \cos(k_p\zeta). \tag{3.33}$$

The effect of spectral shifting is depicted in Figure 3.9. A Gaussian laser pulse which is initially spectrally flat propagates in phase with a wakefield. The front of the laser pulse experiences $dn_e/d\zeta > 0$ and hence is redshifted $\Delta\omega < 0$. The back of the laser pulse experiences $dn_e/d\zeta < 0$ and is blueshifted. This results in a broadening of the overall spectrum.

The equations given above hold in the linear regime, where $a_0 \ll 1$ and relativistic effects can be neglected. Additionally, Equation (3.33) only holds for small spectral shifts $\Delta\omega(\zeta)/\omega \ll v_g/c$ [90]. Beyond this limit, the laser envelope variation also becomes important and the full wave equation needs to be considered.

We can use Equation (3.33) as a first approximation to the spectral shifts that might be expected in the experiment in Chapter 6, to check that they would not be too small to measure. Based on the parameters in the P-MoPA paper [4], for wake amplitudes $\delta n_{\mathrm{e},0}/n_{\mathrm{e},0} = 10\%$, after 100 mm of propagation in a plasma density wave with $n_{\mathrm{e},0} = 5 \times 10^{17}\mathrm{cm}^{-3}$, we find that Equation (3.33) predicts spectral shifts of $\Delta\omega/\omega \sim 20\%$. Therefore, it was expected that spectral shifts would be large enough to detect in the spectrum of the drive pulse in the experiment in Chapter 6.

3.4 Analysing images of laser transverse intensity profiles

In the experiments described in Chapter 5 and Chapter 6, part of the data-set consisted of images of the far-field laser transverse intensity profile. Methods used to quantify the properties of the image and the laser transverse profile are described here.

Spatial resolution

The spatial resolution of an imaging system is defined as the smallest distance between two spots in the object plane that can be distinguished in the image. Commonly, the resolution is quantified by placing in the object plane a USAF target [114], which consists of multiple sets of triplet lines, each set having a different, known line spacing. The USAF target is illuminated and imaged onto the camera. The user finds the set of lines with the smallest spacing that can be distinguished on the image (before the lines blur together). The line spacing of this set corresponds to the spatial resolution of the image.

The spatial resolution of an optical system can be estimated in the design of an experiment using the Rayleigh criterion. To resolve two spots in the image, the spots must be separated by a distance $\geq \delta x$, for δx given by;

$$\delta x \approx \frac{1.22 f \lambda}{D} \tag{3.34}$$

where λ is the wavelength of the laser beam, f is the focal length of the collecting optic (the lens or parabolic/spherical mirror used to re-collimate the light from the object plane) and D is the diameter of the limiting aperture in the system (typically the diameter of the collecting optic). Therefore, to resolve features in the object plane with

a resolution of 10 µm using a laser with wavelength 0.8 µm, the collecting optic needs to have an f-number $f/D \lesssim 10$.

Spatial calibration

The spatial calibration of an image is given in terms of microns per pixel in each of the horizontal and vertical directions. This can be achieved by measuring the number of pixels corresponding to features in the object plane of known size. In Chapter 6, a thin wire was placed in the object plane of the imaging system and translated using a motorised stage with micron-scale step sizes. By moving the motorised stage by known amounts, and taking images after each step, the correspondence between wire translation at the object plane and the shift in pixels on the image could be calculated.

Pre-processing and spot identification

The pre-processing and spot identification algorithms described in this section were developed and implemented by James Chappell for the analysis of the images in Chapter 6. The images consisted of one or more Gaussian-like spots.

The exit mode images were background corrected using a 50×50 square in the corner of the image to define the background level. The spot position was identified on the image by finding the pixel with the 99^{th} percentile counts (using the pixel with the maximum counts would sometimes pick up single-pixel noise). About this spot centre, an ellipse was fitted to a contour of $1/e$ of the 99^{th} percentile counts value using the python openCV package's ellipse fitting tool. The ellipse minor and major axes were scaled up by a factor of 3 to define a region of interest (ROI), within which the spot size could be extracted as discussed below.

In addition, the analysis identified shots for which there was more than one 'spot' in the output mode from the channel. To do this, contours were drawn within the ROI of the image using the openCV find contours tool. Each contour was used to define new ROI masks for calculating the spot size in the same way as for the single-spot mode case.

Spot size

To calculate the spot size, a region the size of $1.5\times$ the ellipse ROI was considered and pixels with counts below some threshold were ignored (here, the threshold was set to be 20% of the signal at the spot centre). Within this region, the $D4\sigma$ method was used to determine the spot size [115]. The $D4\sigma$ measure is equivalent to finding the $1/e^2$ beam diameter of the intensity distribution in the case of a Gaussian transverse beam profile and has advantages over other methods of measuring the spot size as it is well-defined for aberrated and asymmetric beam foci.

To take into account the ellipticity of the spot, the $D4\sigma$ moments in the horizontal and vertical directions were found separately. The components are defined;

$$\sigma_{xx}^2 = \frac{\int_{-\infty}^{\infty} I(x,y)(x-\bar{x})^2 dxdy}{\int_{-\infty}^{\infty} I(x,y)dxdy} \tag{3.35}$$

where x, y are the transverse coordinates in the original image (before coordinate rotation), $I(x,y)$ is the transverse intensity distribution, and $\bar{x}$ is the beam centroid defined;

$$\bar{x} = \frac{\int_{-\infty}^{\infty} I(x,y)\, x\, dxdy}{\int_{-\infty}^{\infty} I(x,y)dxdy}. \tag{3.36}$$

Similarly, σ_{yy} can be calculated by swapping x in the above equations for y. The component σ_{xy} is given by;

$$\sigma_{xy}^2 = \frac{\int_{-\infty}^{\infty} I(x,y)(x-\bar{x})(y-\bar{y})dxdy}{\int_{-\infty}^{\infty} I(x,y)dxdy}. \tag{3.37}$$

The major and minor axes are then calculated as;

$$\sigma_{1,2}^2 = 8(\sigma_{xx}^2 + \sigma_{yy}^2 \pm \sqrt{(\sigma_{xx}^2 - \sigma_{yy}^2)^2 + 4\,\sigma_{xy}^4}). \tag{3.38}$$

The $D4\sigma$ spot size is defined $w_0 = (\sigma_1 + \sigma_2)/2$.

Chapter 4

Novel Optics for the Formation of HOFI Plasma Channels

4.1 Introduction

HOFI plasma channels provide all-optical guiding structures with metre-scale lengths and low ($\sim 1 \times 10^{17}\,\mathrm{cm}^{-3}$) on-axis plasma densities, suitable for future laser-plasma accelerator stages [1, 19, 20] (see Section 2.3). The first proof-of-principle experiments of HOFI plasma channels performed by Lemos *et al.* in 2013 [116], and the Oxford Laser-Plasma Accelerators Group in 2018 [19], used a lens to focus the light from a short-pulse laser beam into a plasma to initiate the plasma channel formation. These experiments generated short plasma channels over lengths of 8 mm [19] and 4 mm [117] respectively.

In a subsequent experiment performed at the Astra-Gemini TA2 facility, RAL, in 2018 [20] (henceforth referred to as 'the Astra 2018 experiment'), the Oxford group replaced the lens with a refractive axicon lens, an optic which forms a longitudinally-extended focus [118]. The use of axicons enabled plasma channels to be formed over much longer distances, achieving guiding of laser pulses (with energy on-target 167 mJ, pulse duration 44 fs and peak axial intensity $4 \times 10^{17}\,\mathrm{W\,cm}^{-2}$) over 16 mm in the Astra 2018 experiment [20]. Guiding over 100s of millimetres was achieved for laser pulses with energy 1.1 J, pulse duration 40 fs and peak axial intensity $6 \times 10^{17}\,\mathrm{W\,cm}^{-2}$ in a follow-up experiment at the Astra-Gemini TA3 facility, RAL, conducted in 2019 [1].

The Astra 2018 experiment was the first time an axicon lens had been used to form a HOFI plasma channel. The high peak intensities required for HOFI plasma channel formation, of order $1 \times 10^{14} - 1 \times 10^{16}\,\mathrm{W\,cm}^{-2}$, mean that the pulse duration of the beams used must be short (typically ~ 40 fs) for the energy of the channel-

forming pulse to be kept to a reasonable value. This is in contrast to the laser pulses previously used for plasma channels formed by inverse-bremsstrahlung (IB) heating of the plasma column, which had comparatively long pulse durations ($\sim 100\,\mathrm{ps}$) and low on-axis intensity requirements ($\sim 1 \times 10^{13}\,\mathrm{W\,cm^{-2}}$)[79]. As discussed in Chapter 1, IB-heating involves collisional processes that only work at high plasma density, hence the resulting plasma channels are limited to on-axis densities above $1 \times 10^{18}\,\mathrm{cm^{-3}}$ [79, 119]. The HOFI method, utilising short high-intensity laser pulses, is capable of generating channels with on-axis densities below $1 \times 10^{17}\,\mathrm{cm^{-3}}$ [1] and hence is better suited as a method of generating low-density plasma channels required for future LPA stages [120]. The difference in laser pulse parameters required for HOFI plasma channel formation compared to the IB method has implications for the optics used in the channel forming beam-line. In particular, chromatic and nonlinear effects may become important at the shorter pulse durations and higher beam intensities.

In the Astra 2018 experiment, the HOFI plasma channel was formed by a 45 fs titanium:sapphire laser pulse that was passed through a refractive axicon lens with base angle 5.6° and focused into a hydrogen gas target. The fill-pressure in the target was varied from 40 mbar up to a maximum of 250 mbar, corresponding to a background plasma density of between $1.6 \times 10^{18}\,\mathrm{cm^{-3}}$ and $1 \times 10^{19}\,\mathrm{cm^{-3}}$ (and on-axis plasma densities in the channel between $\sim 1.5 \times 10^{17}\,\mathrm{cm^{-3}}$ and $1.5 \times 10^{18}\,\mathrm{cm^{-3}}$). During the experiment, it was found that to generate a plasma waveguide, the intensity in the channel-forming beam had to be a factor of 10 higher than the intensity predicted for the axicon focused beam to ionise the hydrogen molecules [83] (see Section 3.2). It has been suggested that the chromatic aberrations introduced by passing the ultra-short pulse through the refractive axicon lens could help account for this observation [83].

Following the Astra 2018 experiment, it was of interest to investigate alternative optics that could be used to form the longitudinally-extended foci required for HOFI plasma channel formation. One drawback of the linear axicon is that, although the transverse spatial profile of the focus is independent of longitudinal position, the axial intensity of the focus tends to increase with longitudinal distance from the axicon tip (see Section 4.2.2). In previous experiments, the intensity variation has been minimised by choosing to use only part of the axicon longitudinal focus starting at a distance $z_{\mathrm{start}} \sim z_{\mathrm{end}}/2$ from the axicon where z_{end} is the furthest point of the longitudinal focus

from the axicon.[1] In this case, the difference in intensity between the start and end of the longitudinal focus can be kept to a factor of ~ 2 [56]. The OFI mechanism is expected to be independent of intensity to first order, as discussed in Section 3.2. Hence the axial intensity variation of the axicon has been considered to have negligible impact on HOFI plasma waveguide generation. However, non-uniform axial intensity is undesirable for plasma channels, firstly from an energy efficiency perspective. Positions along the longitudinally-extended focus with axial intensity above the ionisation intensity threshold correspond to wasted energy in the channel-forming beam as this energy is not usefully conveyed to the plasma. Secondly, the tunnelling ionisation, which dominates for intensities $I < I_{\text{thresh}}$, does have an intensity dependence in ionisation rate and may result in cold electrons being born outside of the central plasma column. The colder electrons may reduce the pressure differential between the hot material and the neutral gas, which is expected to weaken the resulting blast wave, generating lower amplitude, slower moving shocks (see Section 3.2). For future laser plasma accelerator (LPA) stages, which may require channel lengths on the metre scale, the increasing axial intensity with longitudinal position may result in different guiding properties at the start and end of the plasma waveguide. LPAs driven by pulse trains are particularly sensitive to the transverse and longitudinal profiles of the waveguide, so it is important that the plasma channels formed are as uniform as possible over the full length of the LPA stage.

One way to improve the uniformity of the longitudinally-extended focus could be to use optics related to the axicon, such as axilenses and axiparabolae, which have been shown in recent experiments to generate uniform axial intensity profiles [121]. At the time of the investigation outlined in this chapter, these optics had not been applied to the generation of long HOFI plasma channels.

In this chapter, methods to improve the energy efficiency of plasma channel formation and uniformity of the longitudinally-extended foci are presented. The theoretical properties of an extended focus are outlined and a numerical optical propagation code developed by the author is described. The severity of chromatic effects on beams focused by a refractive axicon lens are studied. An investigation into the suitability of axilenses and axiparabolae for HOFI plasma channel generation is presented.

[1] This can be achieved for example by using an annular beam incident on the axicon, with the inner ring of the annulus being at half the radius of the outer ring.

The experimental data used for validation of the numerical code presented in Section 4.3 were taken by Aarón Alejo. Parts of this work have been published in reference [24].

4.2 Theory of longitudinally-extended foci

To generate a longitudinally-extended focus instead of the Gaussian focus typically formed by a standard lens, a different solution to the scalar wave equation is required. In this section, starting from simple assumptions about the form of the focus, it is shown that the transverse intensity profile of an azimuthally-symmetric extended focus takes the form of a Bessel function of the first kind. Bessel foci are sometimes referred to as 'non-diffracting' as they preserve their shape over extended propagation distances.

4.2.1 Bessel beams

The scalar wave equation can be written;

$$\nabla^2 u - \frac{n^2}{c^2}\frac{d^2 u}{dt^2} = 0 \tag{4.1}$$

where $u(\vec{r}, t)$ is the real scalar field component as a function of position $\vec{r}$ and time t, and n is the refractive index of the medium [122]. We restrict to monochromatic wave solutions to Equation (4.1), of the form $u(\vec{r}, t) = \mathrm{Re}\{U(\vec{r})\exp\{-\mathrm{i}\omega t\}\}$ where $U(\vec{r})$ is a complex function and ω is the optical frequency of the wave.

Figure 4.1(a) shows the setup of the optical system being considered. A plane wave is incident on an optical component which applies a phase $\Psi(\xi, \eta)$, where (ξ, η) are the transverse coordinates in the optic plane. We wish to find the form of the field in the focal plane, $U(x, y)$, which is separated from the optic plane by some longitudinal distance z. For this setup, the Fresnel diffraction integral solution to the scalar wave equation can be used [122];

$$U(x, y) = \frac{1}{\mathrm{i}\lambda z}\exp(\mathrm{i}kz)\iint U_{\mathsf{op}}(\xi, \eta)\exp\left\{\frac{\mathrm{i}k}{2z}\left[(x-\xi)^2 + (y-\eta)^2\right]\right\}d\eta d\xi \tag{4.2}$$

where $U_{\mathsf{op}}(\xi, \eta)$ is the complex field at the optic plane, λ is the wavelength and $k = 2\pi/\lambda$ is the angular wavenumber. This form of the diffraction integral assumes that the rays are paraxial (i.e. that $[(x-\xi)^2 + (y-\eta)^2] \ll z^2$), which is true for the optical setups discussed in this section. Transforming Equation (4.2) into cylindrical coordinates

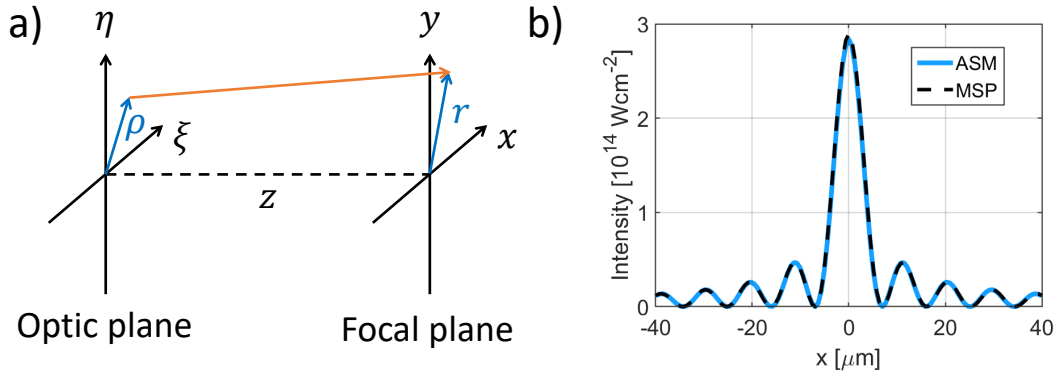


Figure 4.1: The stationary phase method. (a) The coordinate system used for the diffraction calculation. (b) The transverse intensity profile at the focus of a laser beam with transverse phase $\Psi = -k\beta\rho + C$, calculated using the numerical optical beam propagation code (ASM, see Section 4.3) and using the method of stationary phase (MSP, Equation (4.10)). Here, $k = 2\pi/\lambda$ for $\lambda = 800\,\text{nm}$, $\beta = 2.5 \times \pi/180$, $z = 15\,\text{mm}$ and the incident transverse intensity profile is given by a Gaussian with peak intensity $6 \times 10^{11}\,\text{W cm}^{-2}$ and beam waist 2.54 mm.

relevant for a HOFI plasma channel, we denote;

$$x = r\cos(\varphi) \qquad y = r\sin(\varphi)$$

$$\xi = \rho\cos(\Phi) \qquad \eta = \rho\sin(\Phi)$$

and obtain;

$$U(r,\varphi) = \frac{1}{\mathrm{i}\lambda z}\exp(\mathrm{i}kz)\exp\left(\frac{\mathrm{i}kr^2}{2z}\right) \times \iint U_{\text{op}}(\rho,\Phi)\exp\left(\frac{\mathrm{i}k\rho^2}{2z}\right)\exp\left\{-\frac{\mathrm{i}k}{z}\cos(\varphi-\Phi)\right\}\rho d\rho d\Phi. \quad (4.3)$$

Considering the case where the input field has no azimuthal dependence, this equation can be simplified using the definition of a zeroth order Bessel function of the first kind;

$$J_0(v) = \frac{1}{2\pi}\int_0^{2\pi} d\theta\exp\{\mathrm{i}v\cos(\theta)\} \quad (4.4)$$

to give;

$$U(r) = \frac{2\pi}{\mathrm{i}\lambda z}\exp(\mathrm{i}kz)\exp\left(\frac{\mathrm{i}kr^2}{2z}\right)\int U_{\text{op}}(\rho)\exp\left(\frac{\mathrm{i}k\rho^2}{2z}\right)J_0\left(\frac{kr\rho}{z}\right)\rho d\rho. \quad (4.5)$$

If we write the complex input field in terms of a real amplitude and phase, $U_{\text{op}}(\rho) = A(\rho)\exp\{\mathrm{i}\Psi(\rho)\}$, we find that the integral takes the form;

$$f = \int \exp\{\mathrm{i}G(\rho)\}F(\rho)d\rho$$

with $F(\rho) = A(\rho)J_0(kr\rho/z)\rho$ and $G(\rho) = k\rho^2/2z + \Psi(\rho)$. We can then proceed to solve the integral by the method of stationary phase (MSP) [123, 124] in the following

way. The function $F(\rho)$ is slowly varying whereas the exponential component is rapidly oscillating. Hence the integral averages to zero over much of ρ-space. The main contributions to f therefore arise at points where $G(\rho)$ is stationary. By finding the stationary points of $G(\rho)$ and evaluating the integral at these points, a first approximation to the value of f can be found. The contribution to the integral at a stationary point ρ_s can be shown to take the form (see Appendix A);

$$f_s = F(\rho_s)\exp\{\mathrm{i}G(\rho_s)\}\sqrt{\frac{2\pi}{|G_s''|}}\exp(\mathrm{i}\,\mathrm{sign}[G_s'']\pi/4) \tag{4.6}$$

where $G_s'' \equiv \frac{d^2G}{d\rho^2}(\rho_s) \neq 0$ and $\mathrm{sign}[X] = +1$ for $X > 0$ and -1 for $X < 0$ (and 0 if $X = 0$). The full evaluation of f is then given by $f = \sum_{\rho_s\in\Omega} f_s$ where Ω is the set of stationary points of $G(\rho)$.

To use Equation (4.6) to approximate the solution to $U(r,\phi)$ in Equation (4.5), the stationary points of $k\rho^2/2z + \Psi(\rho)$ must be calculated. The stationary points occur at;

$$\Psi'(\rho_s) = -k\rho_s/z. \tag{4.7}$$

The solution under the MSP is then given by;

$$U(r,\phi) = \frac{2\pi}{\mathrm{i}\lambda z}\exp(\mathrm{i}kz)\exp\left(\frac{\mathrm{i}kr^2}{2z}\right)\sum_{\rho_s\in\Omega} U_{\mathrm{op}}(\rho_s)\exp\left(\frac{\mathrm{i}k\rho_s^2}{2z}\right)J_0\left(\frac{kr\rho_s}{z}\right)\rho_s \times\sqrt{\frac{2\pi}{|\Psi_s''+k/z|}}\exp(\mathrm{i}\,\mathrm{sign}[\Psi_s''+k/z]\pi/4). \tag{4.8}$$

The simplest non-trivial result for $U(r,\phi)$ occurs when Equation (4.7) has only one solution. The intensity profile at focus can then be written;

$$I(r,\phi) = \frac{8\pi^3\rho_s^2}{\lambda^2z^2}\frac{1}{|\Psi_s''+k/z|}I_{\mathrm{op}}(\rho_s)J_0^2\left(\frac{kr\rho_s}{z}\right). \tag{4.9}$$

We can see that, in this case, the intensity profile at the focus takes the form of the square of a zeroth order Bessel function.

Returning to the motivation of generating a longitudinally-extended focus, we enforce that the transverse size of the intensity profile is constant over a region of longitudinal extent $z_{\mathrm{start}} < z < z_{\mathrm{end}}$. If we force the argument of the Bessel function in Equation (4.9) to be independent of z, we find that $\rho_s = \beta z$ for some proportionality constant β. Substituting this back into Equation (4.7) and integrating once with respect to ρ gives $\Psi = -k\beta\rho + C$. Hence, assuming the input beam has a flat spatial phase, the optic

must apply a phase that is linearly proportional to ρ. This is precisely the phase applied by a linear axicon. The parameter β is interpreted as the angle of approach of the rays to the axis (see Section 4.2.2). Equation (4.9) then becomes;

$$I(r) = 2\pi\beta^2 kz I_{\text{op}}(\beta z) J_0^2 (kr\beta)\,. \tag{4.10}$$

The intensity profile in Equation (4.10) is shown in Figure 4.1(b) for the case of $\lambda = 800\,\text{nm}$, $\beta = 2.5 \times \pi/180$, $z = 15\,\text{mm}$ and the incident intensity profile $I(\rho)$ takes the form of a Gaussian with peak intensity $6 \times 10^{11}\,\text{W cm}^{-2}$ and beam waist 2.54 mm. The position of the first minimum of the transverse profile is given by $w_0 = 2.4048/k\beta$. Here, $w_0 = 7.0\,\mu\text{m}$. The result from the MSP is compared to the calculation by the numerical optical beam propagation code described in Section 4.3 and found to have excellent agreement. Hence, the approximations made in deriving the form of the transverse intensity profile for a longitudinally-extended focus (Equation (4.10)) provide a reasonable result.

4.2.2 Practical realisation of linear axicons

Optics which apply a phase that increases linearly with radius are known as (linear) axicons [118], which include refractive axicon lenses and reflective axicon mirrors (shown in Figure 4.2). Linear axicons have been found to be very effective in the generation of quasi-Bessel focus patterns over extended longitudinal regions [125].[2] Different methods of incorporating a linear axicon into a laser beam-line are discussed below. Their relative advantages and disadvantages are briefly discussed in 4.4.3. Other methods of generating longitudinally-extended foci apart from linear axicons are discussed in 4.5.

Refractive axicon lens

An axicon lens is characterised by a base angle α, substrate thickness T and refractive index $n(\lambda)$ (see Figure 4.2(a)). The phase applied to a normally-incident beam by an axicon lens is given by;

$$\Psi(\rho) = -(n-1)k\rho\alpha + nkT \simeq -k\rho\theta + nkT. \tag{4.11}$$

[2]Note that a perfect Bessel function profile which has rings extending to infinite radius would require infinite energy in the incident beam, which is non-physical. Hence the transverse focus profiles considered in this section are 'quasi-Bessel' functions, being well approximated by a Bessel function at small radii.

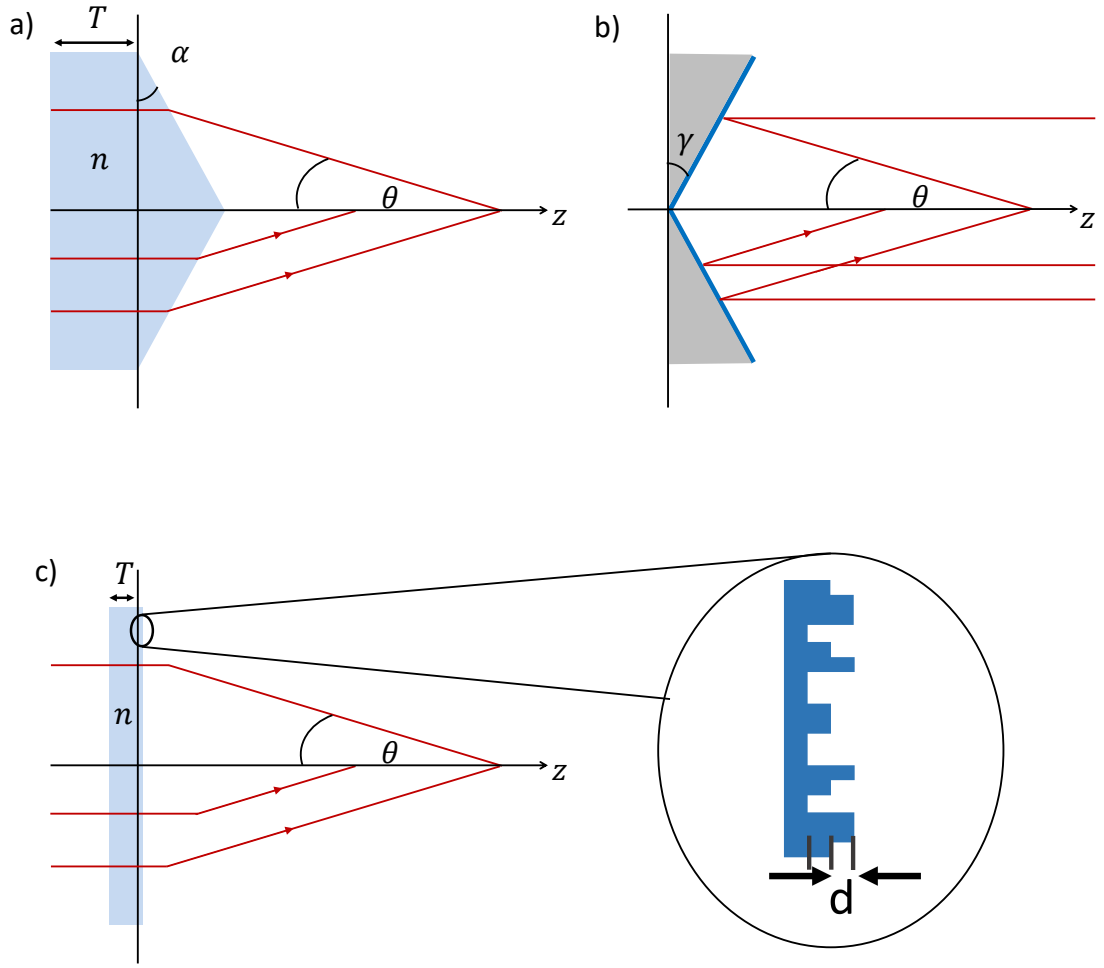


Figure 4.2: Experimental setups for linear axicons, showing geometric ray propagation of a plane wave incident on; (a) a refractive axicon lens; (b) a reflective axicon mirror (with angle of incidence =0°); (c) a transmissive axicon phase plate with layer thickness d. In each case, all rays after the optic propagate towards the axis at a fixed angle θ.

where $\theta \simeq (n-1)\alpha$ (for small angles). In this case, the stationary point in the method of stationary phase, $\rho_s = z/\theta$, described above can be understood from a geometric optics description. Each ray on the incident plane wave leaves the axicon and propagates towards the laser axis at the same angle, θ. This provides a one-to-one mapping between the input radial position of the ray on the optic, ρ_s, and the position it crosses the focal line z (as depicted in Figure 4.2(a)). The angle θ can be calculated from the properties of the axicon using Snell's law at the conical surface to give the relation $\theta \simeq (n-1)\alpha$.

For a Gaussian beam of beam waist ρ_{out} incident on an axicon lens, the output field at position z takes the form [126];

$$U(r,z) = -\mathrm{i}CU_0(2\pi kz)^{1/2}(n-1)\alpha \exp\{(\mathrm{i}\pi/4)\}\exp\{(\mathrm{i}kz)\}\exp\{[-\mathrm{i}kz(n-1)^2\alpha^2/2]\}$$
$$\times \exp\left(\mathrm{i}kr^2/2z\right)\exp\{[-z^2(1-n)^2\alpha^2/\rho_{\text{out}}^2]\}J_0[(n-1)kr\alpha] \tag{4.12}$$

where $C = \Pi_1\Pi_2\exp\{(\mathrm{i}knT)\}$ and Π_1 and Π_2 are the amplitude transmission coeffi-

cients for the air-glass and glass-air boundaries, which are approximated as $\Pi_1 \simeq \Pi_2 \simeq 1$.

Reflective axicon mirror

It is also possible to generate a linear phase from a reflective axicon mirror (depicted in Figure 4.2(b)). The reflective axicon is a conical reflective surface with a base angle γ. The radial dependence of the phase is given by;

$$\Psi(\rho) = -2k\gamma\rho = -k\theta\rho$$

where now the approach angle $\theta = 2\gamma$. Reflective axicons can be manufactured to give the longitudinally-extended focus either on-axis (0° angle of incidence) or off-axis [127].

Axicon phase plates

Recently, axicon phase plates have been used instead of the refractive axicon lens or reflective axicon mirror [128]. Phase plates are typically thin transmissive optics which can be used to apply an approximation to $\Psi(\xi, \eta)$ onto the incident beam (see Figure 4.2(c)). The phase plate consists of a thin substrate onto which successive layers of a refractive material are applied, each layer having a different binary pattern. The layer has thickness d and refractive index n_l (which may be different to n of the substrate) such that the phase applied by a layer of material to an incident beam of wavelength λ is given by;

$$\delta\psi(\xi, \eta) = \frac{2\pi}{\lambda}(n_l - 1)d \times \mathbb{1}(\xi, \eta). \tag{4.13}$$

where $\mathbb{1}(\xi, \eta)$ is the function that determines the pattern on the layer, and equals either 1 or 0. The layers are typically designed such that $\delta\psi(\xi, \eta)$ is a fraction of π (i.e. d is on the scale of the wavelength). Each layer can be considered as a 2D array of pixels in the (ξ, η) plane, with each pixel value being 0 or 1, where a value of 1 corresponds to a position where material is applied. Phase plates are often made with multiple layers of material (or 'levels'), with a greater number of levels resulting in a better approximation of the desired phase function $\Psi(\xi, \eta)$ [129]. Typically, the number of levels varies between 1 and 8.

In a recent paper, Miao *et al.* used an axicon phase plate with 4 levels and a substrate thickness of 0.5 mm to form line foci over 650 mm [128]. The beam used for channel-formation in this case had incident energy 0.9 J and pulse duration 45 fs.

4.3 Modelling optical beam propagation

In order to investigate the chromatic effects through a refractive axicon, it was necessary to simulate the optical propagation of short laser pulses through arbitrary phase optics. The most popular commercial software for modelling optical beam propagation is Zemax OpticStudio [130], which enables the user to input custom or commercially available optical components and optimise the output of the optical system for their requirements. Zemax OpticStudio predominantly uses ray-based algorithms but it can also be used to calculate effects of diffraction and chromatic aberration. However, the software is not well suited to ultra-short pulse propagation and does not include nonlinear, high-peak-intensity effects. Other codes do handle effects owing to ultra-short pulse duration, such as VirtualLab Fusion by LightTrans International [131], but these codes are not freely available.

Therefore, for the purposes of this study, a numerical optical beam propagation code was written by the author. The angular spectrum method was selected as a particularly convenient, fast and robust method of optical propagation that takes advantage of fast Fourier transform methods [132]. The method and the code are described in this section. The code was validated using experimental data of axicon foci under angular misalignments of the optic.

4.3.1 The angular spectrum method

The angular spectrum method (ASM) of free-space propagation, assumes that the scalar wave equation holds (Equation (4.1)). ASM involves taking the 2D Fourier transform of the transverse spatial profile at the input plane (the angular spectrum) and propagating each wavenumber component individually to the output plane via a phase evolution. A second inverse Fourier transform can then be used to generate the new spatial field distribution at the output plane.

This method utilises the simplicity of the Helmholtz equation in wave-number space [122]. The Helmholtz equation comes directly from the scalar wave equation and is written;

$$(\nabla^2 + k^2)U = 0 \tag{4.14}$$

for wave-vector $\vec{k}$, with $k = |\vec{k}| = 2\pi/\lambda$. In transverse wave-number, or 'angular', space (k_x, k_y), this becomes;

$$(-k_x^2 - k_y^2 + \frac{d^2}{dz^2} + k^2)\tilde{U} = 0 \tag{4.15}$$

where $\tilde{U}$ is the field as a function of (k_x, k_y) such that;

$$U = \frac{1}{4\pi^2} \iint \tilde{U}(k_x, k_y) \exp\{\mathrm{i}(k_x x + k_y y)\} dk_x dk_y. \tag{4.16}$$

Equation (4.15) for $\tilde{U}$ has a simple solution of the form;

$$\tilde{U}(k_x, k_y; z) = \tilde{U}(k_x, k_y; 0)\mathcal{H}(k_x, k_y, z) \tag{4.17}$$

where $\mathcal{H}(k_x, k_y, z) = \exp\Big\{(\mathrm{i}\sqrt{k^2 - (k_x^2 + k_y^2)}\ z)\Big\}$ is known as the propagator function [122]. That is, each point in (k_x, k_y) space accumulates a phase $\sqrt{k^2 - (k_x^2 + k_y^2)}\ z = k_z z$ as it propagates.

4.3.2 Numerical code

The numerical code developed in this study models the evolution of the field of a laser pulse as it propagates through a sequence of optical elements to an output plane. The initial transverse and longitudinal shape of the pulse is defined in the input script, as well as the optic type and material properties. The incident laser pulse is taken to have a flat spatial phase front unless otherwise stated. The transverse space is discretised into a 2D Cartesian grid.[3]

Free space propagation is modelled using the angular spectrum method described above. A discrete inverse 2D Fourier transform recovers the field at the output plane. Optical elements are modelled by multiplying the incident amplitude by a complex amplitude transmission function. Equivalently, the optic can be introduced in the spectral domain via a convolution of the spectral field amplitude with the discrete Fourier transform of the transmission function of the optic. Optics are treated as 'thin', i.e. the total phase accumulated is applied at one z-position and the laser does not evolve within the optic itself.[4]

[3] In future, the code could be adapted to 2D polar coordinates to take advantage of the azimuthal symmetry of most use cases of interest. However, it was of interest with this code to be able to model misalignments of an optical system which break azimuthal symmetry, therefore a 2D Cartesian grid was used.

[4] Phase accumulation within the optic is addressed for the inclusion of non-linear effects in Section 4.4.2.

The process described is valid for monochromatic light, i.e. a continuous wave laser. To understand the behaviour of a broadband, short-pulse laser, it is necessary to model the full spectrum of the pulse. In our code, the frequency spectrum is generated via a discrete Fourier transform of the incident temporal profile and a number of discrete angular frequencies ω are sampled. For each frequency, the initial field amplitude is scaled by the spectral shape, for example for a Gaussian spectrum;

$$U_\omega(x, y, z = 0, \omega) = U(x, y, z = 0) \exp\left[-\frac{(\omega - \omega_0)^2}{2\sigma^2}\right] \tag{4.18}$$

where $\omega = 2\pi c/\lambda$, ω_0 is the central frequency and σ is the $1/e^2$ bandwidth of the spectral intensity. U_ω then evolves via the steps described above. The field at a particular transverse position in the output plane (x, y) may be extracted for each frequency. The temporal profile may then be retrieved by an inverse Fourier transform;

$$U_t(x, y, z, t) = \mathcal{F}_t^{-1}[U_\omega(x, y, z, \omega)]. \tag{4.19}$$

The simulation box size and the number of grid points must be chosen so as to be consistent with the desired resolution in the transverse domain. Because the code is based on Fourier transforms, the number of pixels is conserved from the input plane to the output plane. This means that, to model beam sizes of order 10s of millimetres and foci with structure on the micron scale, spatial grids with $\sim 10,000$ pixels are required. Moreover, in ASM, the grid in the angular spectrum domain has the same number of pixels as the spatial domain. The resolution in the angular spectrum domain must be sufficient to resolve the propagator function. This results in the Nyquist sampling conditions, which put an additional constraint on the number of pixels [132]. For sufficient Nyquist sampling, the box size should be less than or equal to the Nyquist box size $L \leq L_N = \sqrt{\lambda z_{prop} M}$ for z_{prop} the propagation distance and M the number of pixels [133].

In this work, a typical simulation employed a transverse spatial grid of 20 000 $\times$ 20 000 pixels, covering a transverse spatial extent of 20 $\times$ 20 mm^2. Propagation distances were of order 15 mm and a spectral grid of 1000 wavelengths was used. The code was written in MATLAB and run on a personal computer with 256 GB random-access memory, 20 cores and 40 logical processors. To reduce running time, the propagation of each wavelength was calculated in parallel, utilising the MATLAB parallel computing capabilities.

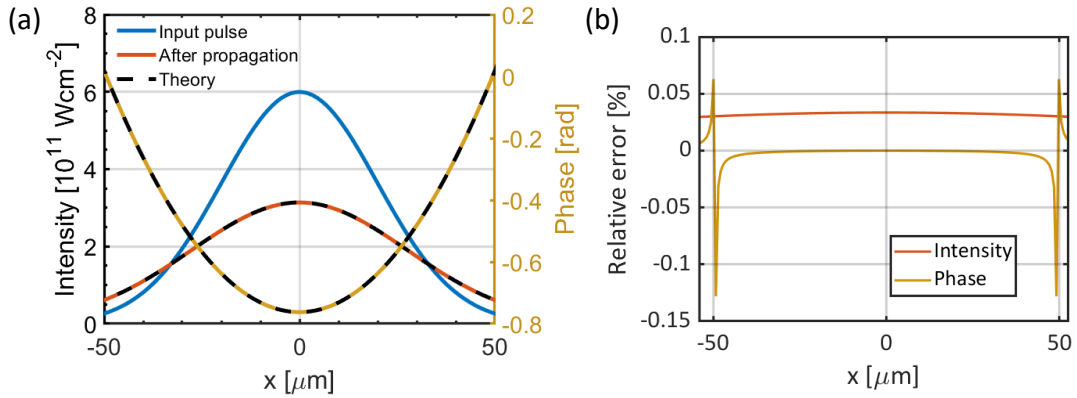


Figure 4.3: Validation of the ASM code for the case of a monochromatic Gaussian beam propagating in free space, with wavelength 805 nm and initial spot size $w_0 = 40\,\mu$m. (a) The intensity profile and transverse spatial phase profile are plotted after a propagation distance $z = 6\,$mm. The outputs of the ASM code are overlaid with known analytic results for the free-space propagation of a Gaussian beam (see text). (b) The difference between the ASM and analytic results are plotted as a percentage of the ASM value as a function of transverse position.

With these settings, a simulation with a single propagation step of a broadband pulse took approximately 30 minutes to run.

4.3.3 Code validation

The ASM code was tested against a number of simple optical setups with known analytic behaviour. Figure 4.3 shows an example validation check, here for the case of a Gaussian monochromatic beam with initially flat spatial phase propagating in free-space, which was used to check that the free-space phase evolution and energy conservation were achieved by the code implementation. A more sophisticated validation test is performed in the following subsection.

For this text, the Gaussian beam had wavelength 805 nm and initial spot size $w_0 = 40\,\mu$m and hence a Rayleigh length of $z_R = 6.3\,$mm. The simulation box size had side lengths 6.9 mm, with $10,000 \times 10,000$ pixels used, giving a resolution of 0.7 μm in both transverse dimensions. The beam was propagated in free-space for a distance of 6 mm. The transverse intensity profile of the beam after propagation is plotted in the figure, showing that the Gaussian shape is retained. The width of the Gaussian is increased to $w_0 = (55.3 \pm 0.7)\,\mu$m by the ASM code, within the uncertainty given by the spatial resolution, compared to the analytic spot size $w(z) = w_0\sqrt{1 + \left(\frac{z}{z_R}\right)^2} = 55.31\,\mu$m to 4 significant figures [89]. The peak intensity of the Gaussian profile is reduced, as expected from energy conservation arguments, resulting in a profile that matches the expected

profile to within a relative error of $< 0.1\%$, as can be seen in Figure 4.3 (b). The transverse spatial phase profile across the electric field at $z = 6\,\text{mm}$ is also plotted. The expected dependence of the phase on radial position for a Gaussian beam propagating in free space is given by $\phi(r,z) = \frac{kr^2 z}{2(z^2+z_R^2)} + kz - \arctan\left(\frac{z}{z_R}\right)$ [89]. The phase profile output by the ASM code is found to match the known analytic result.

Experimental verification of code output

The ASM code was also compared to more complex optical setups and compared to experimental results. It is found experimentally that small angular misalignments of an axicon degrade the focal spot quality substantially. This behaviour can be investigated using the numerical code. For the experimental demonstration, the beam from a helium-neon laser, operating at $\lambda = 633\,\text{nm}$, was expanded and propagated through a refractive axicon with $\alpha = 5.52°$ and $t = 3.0\,\text{mm}$. The axicon focus was imaged by a microscope objective onto a CCD camera at $z = (140 \pm 2)\,\text{mm}$ downstream of the axicon.

To set the angle δ (see Figure 4.4(a)), the axicon alignment was first optimised to give a focus best resembling the ideal Bessel shape. This fixed $\delta = 0.0\,\text{mrad}$. The screw on the axicon mount was then used to introduce small angular offsets and the camera was moved transversely to keep the focal spot on the screen. The angles between the axicon and incident beam were determined by calibrating the screw turns, which was done by replacing the axicon with a mirror and making 10 turns of the screw, marking the change is angular position of the beam a long distance away from the mirror mount. The angular change was divided by two (as the change in mirror angle gave double the change in beam angle) and divided by the number of screw turns (10) to give an angle per turn. This procedure gave $(5.5 \pm 0.4)\,\text{mrad}$ per turn, where the error represents the uncertainty on the angle measurement. The spots displayed in Figures 4.4(c-e) were produced at $\delta = 22.0$, 33.0, $44.0 \pm 0.4\,\text{mrad}$ respectively.

In the numerical code, the angular misalignments are implemented by applying a tilt to the incident beam before it reaches the axicon. This is equivalent to varying the angle of the axicon with respect to the incident beam as in the experiment. For a tilt about the y-axis, a phase was applied to the incident beam as a function of transverse position x of the form $\phi_\delta = kx\tan(\delta)$. Additional phase introduced by refraction inside the axicon can be estimated geometrically and was found to be negligible for the angles

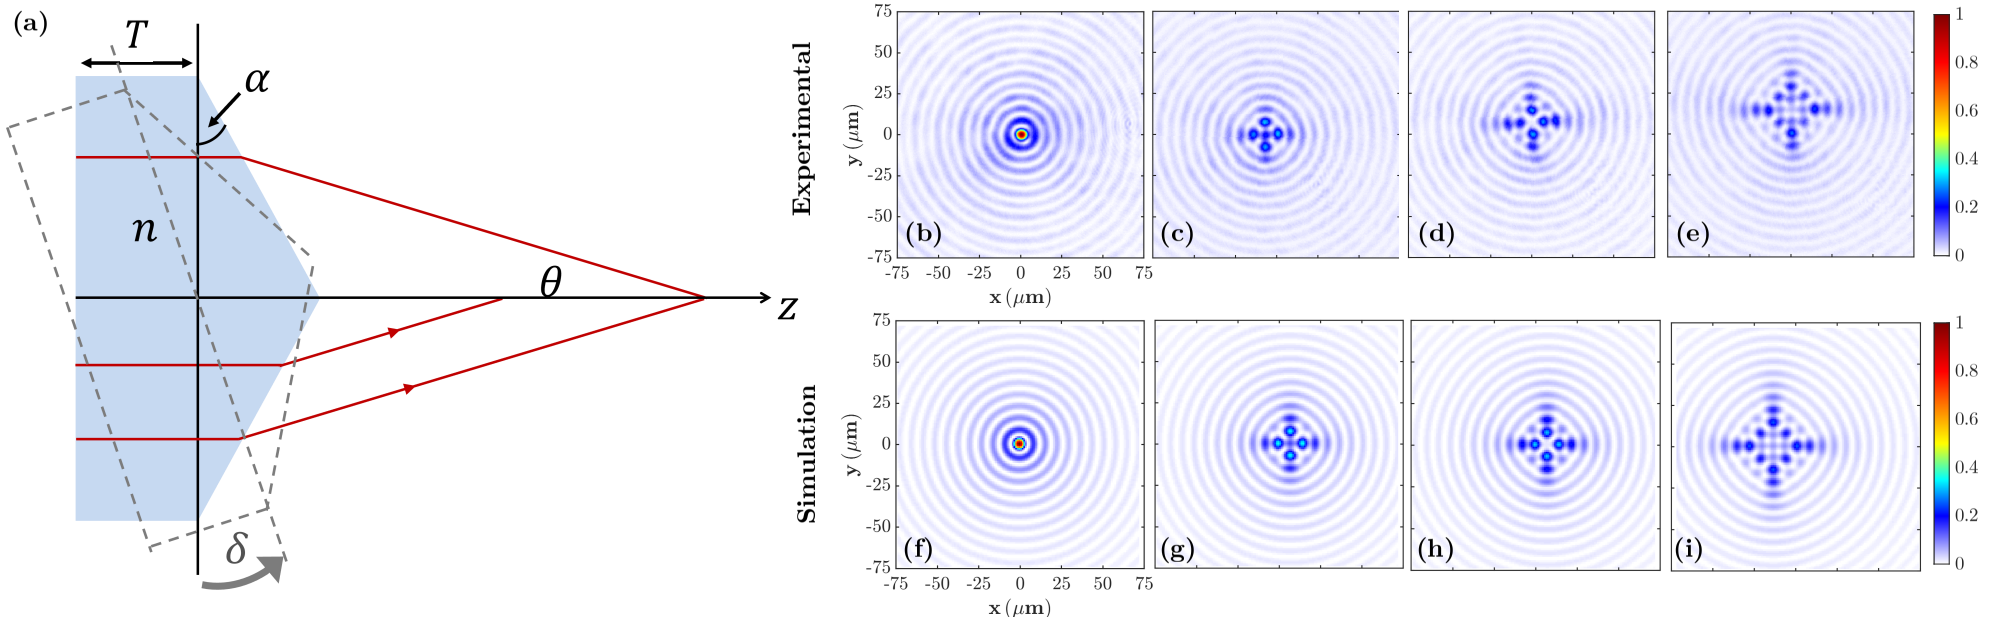


Figure 4.4: Validation of the code for the case of a refractive axicon as shown in (a), where α is the axicon base angle, T is the substrate thickness, $n(\lambda)$ is the refractive index and θ is the approach angle. (b-e) are the focal spots produced for different values of misalignment angle, δ, with (b) being the $\delta = 0.0$ mrad case. These images were found experimentally for $\lambda = 633$ nm at $z = (140 \pm 2)$ mm downstream of the axicon, for misalignment angles $\delta = 22.0$, 33.0, 44.0 $\pm$ 0.4 mrad respectively (c-e). These can be compared to the focal spots generated by the numerical code shown in (f-i) recorded at 140 mm from the axicon, at angles $\delta = 0.0$, 49.5, 55.0, 66.0 mrad. Possible reasons for the difference between the δ angles measured in experiment compared to those found in the simulation for similar focal spot quality are discussed in the text. The colour scales in (b-i) are normalised to the peak of (f).

considered here (see Appendix A). The laser and axicon parameters in the simulation were set to match those in the experiment. The refractive index was calculated from the Sellmeier equation for fused silica [134]. The resulting focal spots were compared to the experimental results in Figures 4.4(c-e). It was found that using the same angle δ as in the experiment did not produce the same focal spot profiles when modelled in the code. Larger angles were needed in the code to generate the same focal spot degradation.

To quantify the difference between the code and the experiment, the misalignment angle was treated as a free parameter in the code, varied in steps of 5.5 mrad between 44.0 mrad and 71.5 mrad, and the simulated image which most closely corresponded to each of the experimental images was determined. This best-match image was found by taking the average root-mean-square (RMS) deviation between a line-out along the x-axis of the experimental image and the line-outs of the simulated images. The simulated image which gave the minimum average RMS deviation was selected. The RMS deviation could have instead been calculated from the entire image of the experimental and simulated foci, however this would have given greater weight to the behaviour of the focus at large radii, which contains less information about the aberration and is more susceptible to noise. The best-match simulated images are presented in Figures 4.4(g-i) and were generated using $\delta = 49.5$, 55.0, and 66.0 mrad respectively, however improved matching may be found for values of δ between the steps used in the simulation.

A possible reason for the increased values for δ required in the simulation to match the images found in experiment is that the beam in the experimental suffered from additional phase aberrations not modelled in the code. Small aberrations in the beam can result in much more rapid spot quality degradation with δ than without any aberration, whilst only having a small effect on the focal spot at $\delta = 0$.

The phase error on the beam was not measured on this experiment, so the numerical code cannot exactly replicate the experimental conditions. However, the code can be used to demonstrate the effect that additional aberrations can have on the focal spot quality degradation. For this demonstration, the code was run with two example aberrations, namely a tip error δ_x (angular rotation about the x-axis) and an astigmatism component $\phi_{astig} = A\frac{kr^2}{R}[\sin(2\varphi) + \cos(2\varphi)]$ where φ is the azimuthal angle at the input plane and A is an astigmatism parameter. The aberration severity was small such that the focal spot at $\delta = 0$ mrad was still reasonably azimuthally symmetric. Figures

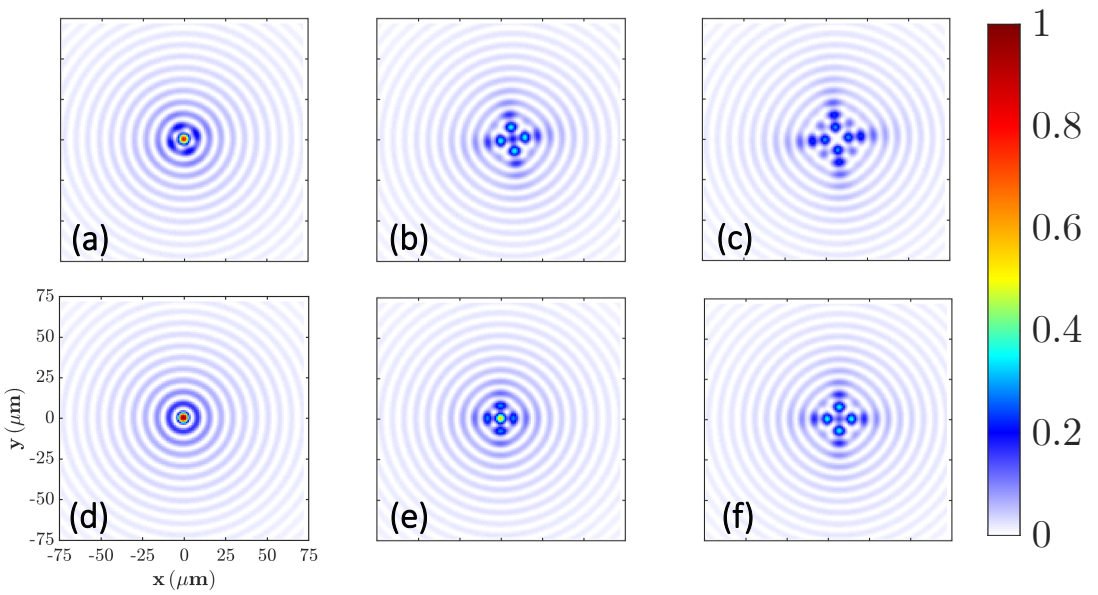


Figure 4.5: (a-c) show the axicon focal spots generated by the numerical code for different values of tilt error δ as described in the text, but with additional phase aberrations added to the beam, namely a tip error $\delta_x = 1$ mrad and astigmatism with astigmatism parameter $A = 10^{-5}$. (a-c) correspond to $\delta = 0.0$ mrad, 38.5 mrad, and 49.5 mrad respectively. (d-f) show the focal spots for the same values of δ but with the astigmatism and tip removed. The colour scales are normalised to the peak of (d).

4.5 (a-c) show the simulated focal spots when additional aberrations with $\delta_x = 1$ mrad and $A = 10^{-5}$ were included, compared to Figures 4.5 (d-e), where the same values of δ are used in the code but without the additional aberrations. The focal spot in Figure 4.5 (a) is qualitatively similar to the spot in Figure 4.4 (b), with a clear maximum in the centre and no significant azimuthal asymmetry. However, the focal spot in Figure 4.5 (b) with $\delta = 38.5$ mrad and tip and astigmatism errors included is significantly more degraded than the spot in Figure 4.5 (e) with the same δ but no additional aberrations. Indeed, the spot in Figure 4.5 (b) is more qualitatively similar to that in Figure 4.5 (f) with $\delta = 49.5$ mrad and no other aberrations. In other words, the inclusion of additional aberrations on the beam results in a reduction in the tilt error required to generate the same focal spot degradation. This demonstration shows that the additional aberrations increase the rate of spot degradation with increasing δ.

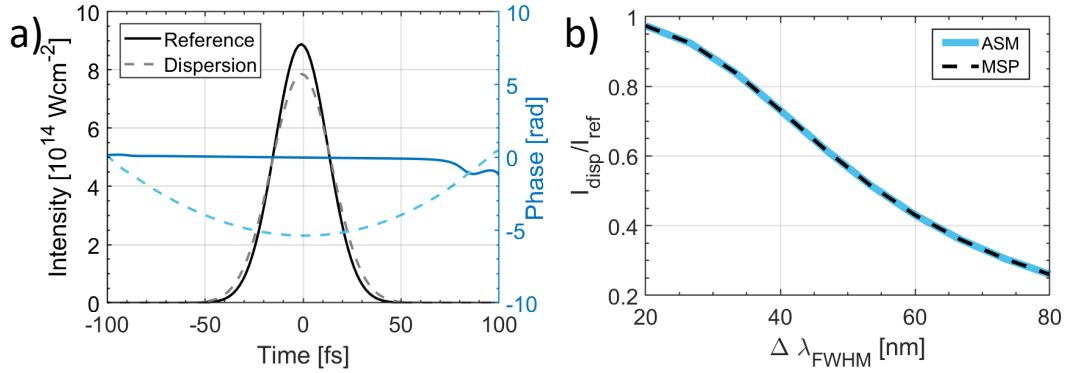


Figure 4.6: The effect of dispersion on the temporal intensity profile of the channel forming pulse compared to a zero-dispersion reference case at $z = 20\,\text{mm}$. (a) shows the temporal intensity profile and corresponding temporal phase for $\Delta\lambda_{\text{FWHM}} = 30\,\text{nm}$. The reference is calculated assuming $n(\lambda) = n(\lambda_0)$. In (b), the calculated ratio of the peak on-axis intensities obtained with dispersive (I_{disp}) and non-dispersive (I_{ref}) axicons is plotted as a function of the bandwidth of the incident pulse, calculated using the numerical code (ASM) and method of stationary phase expression (MSP).

4.4 Chromatic effects on axicon-focused beams used to generate HOFI plasma channels

4.4.1 Chromatic dispersion

As discussed in the introduction to this chapter, in 2018 an experiment was performed into the generation of HOFI plasma channels at the Astra-Gemini TA2 facility, RAL. In this experiment, the channel-forming beam had energy 114 mJ, pulse duration 45 fs and central wavelength 805 nm [20, 83]. The channel-forming beam was passed through a refractive UV fused-silica axicon with approach angle 2.5° to form a longitudinally-extended Bessel focus that was designed to ionise hydrogen gas via OFI to form the plasma column. Prior to the experiment, the intensity required in the incident channel-forming beam on the axicon to generate an intensity on-axis matching the threshold ionisation intensity for hydrogen molecules, $I \gtrsim I_{\text{thresh}} \approx 2 \times 10^{14}\,\text{Wcm}^{-2}$, was calculated using Equation 4.12. During the experiment, it was found that the incident intensity in the channel-forming beam had to be increased by a factor of 10 to effectively ionise the hydrogen gas [83]. One possible explanation for this observation is that the refractive axicon resulted in significant chromatic aberration that reduced the on-axis intensity in the longitudinal focus below that predicted by Equation 4.12. Therefore, in this section, chromatic effects on a beam passing through a refractive axicon are explored. Alternative reasons for the experimental observation are discussed in Section 6.7.

To investigate the chromatic effects of a refractive axicon, the numerical ASM code

was run with inputs chosen to model the Astra 2018 experiment [20]. The channel-forming laser beam was modelled as a bi-Gaussian pulse, with central wavelength $\lambda_0 =$ 805 nm, bandwidth $\Delta\lambda_{\mathrm{FWHM}} = 30$ nm and incident peak intensity $I_0 = 6 \times 10^{11}\,\mathrm{W\,cm^{-2}}$. In the experiment, the laser beam waist was 25.4 mm and the focus started 200 mm downstream of the axicon (due to a hole in the centre of the axicon of radius 9 mm) [83]. In the numerical calculation, the beam waist and propagation distance were scaled down by a factor of 10 to $\sigma = 2.54$ mm and $z = 20$ mm to enable adequate resolution without overloading the computer memory. This is equivalent to analysing the focus closer to the axicon tip than in the experimental setup. The axicon had the same properties as for the code validation above, consisting of fused silica with a substrate thickness of 3.0 mm and base angle of 5.52°. Dispersion was calculated through both the substrate and the conical part of the axicon, using the Sellmeier equation for $n(\lambda)$ of fused silica [134].

The results of the simulation for the on-axis intensity are shown in Figure 4.6. The reference case shown is for zero-dispersion through the refractive axicon, which is generated by setting $n(\lambda) = n(\lambda_0)$. As seen in Figure 4.6(a), it is found that, for the parameters used in this simulation, the effect of dispersion is to reduce the peak value to 89% of the reference value. The temporal phase is plotted on the same figure. The phase for the reference case is almost perfectly flat over the region of the pulse, suggesting that when dispersion is ignored the pulse is near diffraction-limited. In the dispersive case, the pulse develops a parabolic phase dependence, as is expected for group delay dispersion [89].

In order to gain further insight into the role of dispersion, we calculate the ratio of the peak axial intensity with and without dispersion, r, as a function of bandwidth of the incident pulse. Figure 4.6(b) plots this ratio for $z = 20$ mm. As expected, increased bandwidth leads to greater reduction in peak intensity. For an incident beam with a Gaussian transverse profile, the analytic expression from Equation (4.12) can also be used to investigate the effects of dispersion. Calculating the on-axis field at a fixed z-position for a number of wavelengths in the spectrum, and computing the inverse Fourier transform to the time domain (as in Equation 4.19) generates the temporal intensity profile at that focus position. Again, dispersion was introduced via $n(\lambda)$ using the Sellmeier equation for fused silica, and the reference case was fixed with $n(\lambda) =$

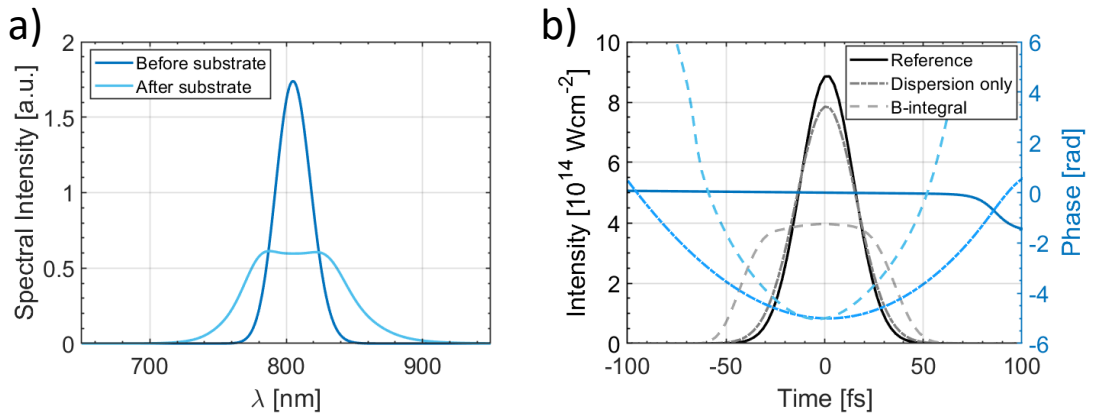


Figure 4.7: Nonlinear effects on the focus of a refractive axicon lens, with same parameters as Figure 4.6. The B-integral was calculated within the 3 mm thick substrate of the axicon. (a) Spectrum of the short-pulse before and after propagation through the axicon substrate. (b) Temporal intensity profile and corresponding temporal phase of the pulse when B-integral is included in the calculation, compared to the dispersive and reference cases which assume $n_2 = 0$.

$n(\lambda_0)$. Figure 4.6(b) shows the excellent agreement between the numerical and analytic calculations, providing further validation of the numerical code.

4.4.2 B-integral

The calculations provided so far do not include nonlinear effects. In propagating through a medium of thickness T, a pulse accumulates nonlinear phase (also known as the B-integral) $\phi_{NL} \approx kn_2IT$ where n_2 is the nonlinear refractive index and I is the peak intensity of the pulse [89]. The nonlinear phase leads to self focusing and self phase modulation, the latter of which further increases the pulse bandwidth. This can be seen by considering the instantaneous frequency of the pulse $\omega(t) = -d\phi(t)/dt \approx kn_2T\,dI/dt$. For a short pulse, dI/dt can be significant, leading to modifications in the frequency content of the pulse.

To avoid these problems it is therefore usual to seek to keep $\phi_{NL} < 1$ as a rule of thumb [89]. However, for the parameters considered here, assuming an intensity of $6 \times 10^{11}\,\mathrm{W\,cm^{-2}}$, with a nonlinear refractive index of fused silica $n_2 = 2.6 \times 10^{-16}\,\mathrm{cm^2\,W^{-1}}$, the thickest part of the axicon introduces $\phi_{NL} \simeq 5$. Therefore it is important to also consider nonlinear phase effects on the short pulse used for HOFI plasma channel formation.

To investigate the combination of B-integral and chromatic effects through an axicon, the numerical code was modified to include a calculation of the nonlinear phase

accumulation. For a lower bound calculation of the effects, only the B-integral accumulated via propagation through the flat substrate was considered, and nonlinear effects in the conical section were ignored (the thickest part of the conical section the beam passed through had a thickness of 1.6 mm compared to the substrate thickness $T = 3$ mm). The substrate of thickness T was split into a series of slices of thickness δT. After each slice, the nonlinear phase was calculated in the time-domain via $\phi_{NL}(t) \approx k_0 n_2 I(t) \delta T$ where k_0 is the central wavenumber of the pulse. A Fourier transform was then performed from time to frequency space and each frequency component was propagated through the slice, developing a phase $\tilde{\phi}(\omega) = -kn(\omega)\,\delta T$. An inverse Fourier transform back from frequency space to time space then generated the new temporal intensity profile for the next nonlinear phase calculation step. Once the phase accumulation through all slices of the substrate had been calculated, the phase through the conical section of the axicon was applied (Equation (4.11)). The beam then propagated to the output plane via the angular spectrum method.

The parameters of the setup were the same as for the chromatic effects study. The substrate of thickness $T = 3$ mm was split into 50 slices. The results of the calculation are shown in Figure 4.7, for $z = 20$ mm. Figure 4.7(a) shows the spectrum of the laser pulse before and after the pulse has accumulated B-integral in the axicon substrate. The nonlinear effects cause a spectral broadening of the pulse to a full-width-half-maximum (FWHM) of 85 nm, compared to the 30 nm at the input. In Figure 4.7(b) is shown the temporal intensity profile of the laser pulse at the focus of the refractive axicon, with nonlinear and dispersive effects applied, compared to the cases of dispersion only and the reference non-dispersive case. It is clear that including B-integral in the calculation has a significant effect on the resulting intensity at focus, with chromatic effects and dispersion reducing the peak intensity to 50% of the case with dispersion only. The overall reduction in focus intensity compared to the non-dispersive, no B-integral reference case is then 45%. [5]

It is worth noting that increasing the energy of the incident beam alone in an attempt to compensate for the reduced intensity in the focus also increases the nonlinear

[5] Note that the reduction in peak intensity due to dispersion and non-linear effects in this case ($r \sim 45\%$) is slightly less severe than would be predicted by the results in Figure 4.6(b) for a beam with 80 nm FWHM bandwidth due to dispersion only ($r \sim 25\%$). This is because the spectral shape of the beam due to non-linear effects, shown in Figure 4.7(a) is not a Gaussian as was assumed for the calculations in Figure 4.6.

phase accumulation in the refractive axicon ($\phi_{NL} \propto I$). This leads to further spectral broadening of the pulse and an even smaller value for r, as per Figure 4.6(b). Therefore, the resulting intensity at the focus of a refractive axicon, which scales as $r(I_{\text{op}}) \times I_{\text{op}}$, depends on the competing effects of increased energy in the beam and increased pulse duration owing to the nonlinear and chromatic effects. The full form of the dependence of r on the incident intensity of the laser pulse is beyond the scope of this investigation, but the considerations outlined here would suggest that a simple doubling of the incident peak intensity would not be sufficient to make up for a 50% reduction caused by chromatic effects through the refractive axicon.

4.4.3 Chromatic effects through a phase plate

The reduction in on-axis intensity resulting from chromatic effects through an axicon lens motivates the use of alternative optics. One option would be to use a reflective axicon mirror, which does not introduce B-integral or chromatic aberration. However, reflective axicons are more difficult to employ in an optical setup than refractive axicons [128]. This was made evident during an experiment performed by the Oxford Laser Plasma Accelerators Group at Astra-Gemini TA3 in 2019 where we attempted to use a reflective axicon for HOFI plasma channel generation, but found that aligning both the off-axis angle of the axicon mirror and achieving an ideal Bessel focus during day to day alignment was time consuming and hence the reflective axicon was eventually swapped with the refractive axicon on this experiment. We later discovered that the reflective axicons we were supplied did not have the requested angle of incidence, suggesting potential manufacturing difficulties. A simpler solution would be to replace the refractive axicon lens with a transmissive axicon phase plate. It is therefore important to calculate the severity of chromatic effects on a beam that is passed through a transmissive axicon phase plate.

Phase plates with substrate thickness 0.5 mm have recently been used in HOFI plasma channel experiments for linear axicons [128]. The B-integral accumulated through a fused silica substrate of this thickness would be $\phi_{NL} \sim 0.6$ for a beam with the same parameters as used in the refractive axicon study above, which is much less than through the refractive axicon substrate and also lower than 1, making it acceptable for use in this beam-line (provided B-integral through the other optics in the beam-line is < 0.4).

However, diffractive optical elements such as phase plates are very sensitive to wavelength. From Equation (4.13), a phase-plate designed for a specific wavelength λ_1, will result in a different phase accumulation for a different wavelength $\lambda_2 = \lambda_1 + \delta\lambda$. In the limit of small $\delta\lambda$, the phase applied to λ_2 will be different by an amount $\delta\psi$ given by;

$$\left(\frac{\delta\psi}{\psi}\right)^2 = \left(\frac{\delta n_l}{n_l(\lambda_1) - 1}\right)^2 + \left(\frac{\delta\lambda}{\lambda_1}\right)^2 \tag{4.20}$$

where $\delta n_l = n_l(\lambda_2) - n_l(\lambda_1)$.

From this equation, we can estimate the effects on a broadband pulse as it propagates through a fused silica phase plate. For a laser pulse from a titanium-doped sapphire laser, with centre wavelength 800 nm and FWHM bandwidth 30 nm, 99.7% of the energy of the pulse is contained in a wavelength range of 760 – 840 nm, i.e. a variation of 5% either side of the centre wavelength. An example material used for the layers of the phase plate (chromium [135]) has refractive index variation over this range of wavelengths of 2% [134]. Hence the maximum relative 'phase error' introduced for different wavelengths within the pulse is $\sim 5\%$. Manufacturing errors in the phase jumps printed on phase plates are typically of the order 10%, and it has been demonstrated that these manufacturing errors do not degrade the focus properties by a significant amount [136]. Therefore we can expect that chromatic effects through a phase plate for HOFI plasma channel generation would be minimal. Indeed, phase plates have now been used successfully with titanium-doped sapphire laser beams [128, 137]. It would be possible to model such phase plate axicons in the ASM code in a future investigation.

4.4.4 Discussion

In the context of the Astra 2018 experiment, we have found from numerical calculations that the B-integral accumulated in the axicon substrate, coupled with dispersion, was significant enough to result in a reduction in axial intensity to half of that expected from Equation (4.12). Future experiments incorporating refractive axicons will need to consider the nonlinear effects when designing the fluence requirements for the plasma-channel-forming laser beam. The loss in intensity caused by chromatic and nonlinear effects means that more energy is needed to achieve the threshold intensity for ionisation, reducing the energy-efficiency of the setup compared to the case without chromatic

effects. To achieve efficient HOFI plasma channels for future LPA stages, it may be preferable to move to a reflective axicon or thin phase plate setup instead.

The intensity reduction resulting from B-integral and chromatic effects does not completely account for the experimental observation found on the Astra 2018 experiment, where 10 times more energy was needed to ionise the plasma than theoretically predicted for an axicon focus in vacuum [83]. It is possible that, as the beam focused through the hydrogen gas and formed a plasma, the refractive index of the plasma itself may have acted back on the laser pulse and refracted light away from the axis. Particularly at the higher plasma densities, the approach angle used was close to the critical angle of the plasma θ_{crit} (see Chapter 3). For $\theta < \theta_{\text{crit}}$, the focusing beam would be reflected by the plasma as it formed and would not reach the laser axis to efficiently ionise and heat the plasma column. To fully understand the discrepancy observed, these and other factors would need to be considered.

Thin phase plate optics are found to introduce significantly reduced nonlinear effects compared to the refractive axicon. Replacing the refractive axicon with a phase plate axicon would be a simple and effective change to the channel-forming beam-line and could be considered for future experiments.

4.5 Longitudinally-extended foci with uniform axial intensity

The linear axicon discussed so far can be used to generate a longitudinally-extended focus with a uniform transverse spatial scale. However, the axial intensity of the extended focus of a linear axicon (Equation (4.10)) maintains a z-dependence. In the case where the incident beam has a top-hat transverse profile, the intensity along the axis increases linearly with z. For HOFI plasma channel generation, we would ideally like to achieve a constant intensity along the longitudinally-extended focus to consistently ionise the plasma and form a uniform channel over long LPA stage lengths.

One way to achieve a uniform axial intensity profile would be to use an incident transverse beam profile that falls with radius, $I(\rho_s) \propto 1/\rho_s$, to exactly cancel the z dependence in Equation (4.10). However, forming such a transverse beam profile from a Gaussian or top-hat beam that is typically output by a laser system would require an

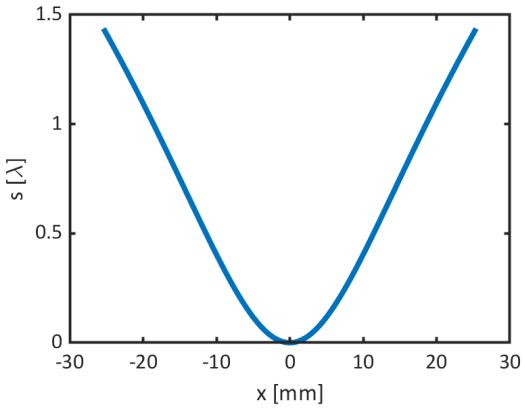


Figure 4.8: Illustrative variation of an axilens surface thickness s as a function of transverse position x.

amplitude mask with a transmission dependent on radial position, reducing the total energy available for channel formation [138, 139].

As discussed below, by relaxing the constraint on the spot size uniformity of the line focus, it is possible to achieve a longitudinally-extended focus with a uniform axial intensity profile. The optics used to generate such foci are very closely related to the axicons described so far and are referred to as axilenses (transmissive version) and axiparabolae (reflective version). For convenience, 'axilens' is used to refer to both axilenses and axiparabolae for the remainder of the chapter.

4.5.1 Properties of axilenses

To motivate the form of the uniform axial intensity solution, we follow the method of Sochacki *et al.* [140] and consider the axicon setup from an energy conservation viewpoint. Let the incident beam have a top-hat profile with energy density P_σ, such that the energy in an infinitesimal ring on the axicon at radius ρ is given by $u_\sigma(\rho) = P_\sigma 2\pi\rho\delta\rho$. Let the energy density on-axis in the axicon-like focus be given by P_z such that the energy along an element of the axis $u_z(z) = P_z(z)\delta z$. We make the approximation that all of the energy in a ring at radius ρ contributes to the on-axis energy at position $z(\rho)$, such that we can write;

$$P_\sigma 2\pi\rho\delta\rho = P_z(z)\delta z. \tag{4.21}$$

If we now force the energy density along the line focus to be independent of z and take the limit where $\delta\rho \to 0$ and $\delta z \to 0$, we find a relationship between z and ρ required for a uniform longitudinal intensity, namely $z \propto \rho^2$. This can be contrasted with $z \propto \rho$

for the linear axicon case. This new relationship between z and ρ can now be substituted as the stationary points in Equation (4.7). For $z = \zeta\rho_s^2 + z_c$, we get;

$$\Psi(\rho_s) = -\frac{k}{2\zeta}\log\left(\zeta\rho_s^2 + z_c\right). \tag{4.22}$$

This phase corresponds to a variation of surface thickness for a refractive or reflective optic as illustrated in Figure 4.8. Substituting the result for Ψ back into Equation (4.9) gives;

$$I(r,\phi) = 2\pi k^2 I(\rho_s) J_0^2\left(\frac{kr\rho_s}{z_c + \zeta\rho_s^2}\right). \tag{4.23}$$

Hence, for the case of a beam with a top-hat transverse intensity profile incident on an optic with $\Psi(\rho_s)$ given by Equation (4.22), the axial intensity of the extended focus is independent of z.

However, by using the phase for the axilens derived above (Equation (4.22)), the uniformity in the spot size w_0 is lost. Instead, we find that the first minimum in the Bessel-function varies as;

$$w_0(z) = \frac{2.4048\sqrt{\zeta}}{k}\frac{z}{\sqrt{z - z_c}}. \tag{4.24}$$

Going along the extended focus, w_0 initially decreases with z for $z_c < z < 2z_c$, reaching a minimum at $z = 2z_c$ and then increases with z for $z > 2z_c$. Operating around the stationary point in w_0 at $z = 2z_c$ would give the least variation in spot-size. The behaviour of the spot size of the axilens is demonstrated using numerical calculations in Section 4.5.3.

4.5.2 Practical realisation of axilenses

Studies into the use of axilenses for plasma channel generation are limited. One research group has attempted to manufacture a reflective axilens, which they term the axiparabola [121, 143, 144]. The axiparabola applies a phase to the beam $\Psi(\rho) = 2ks(\rho)$ where $s(\rho)$ is the relative offset of the mirror surface from the flat substrate as a function of radial position (for a linear axicon $s(\rho) \propto \rho$). The researchers designed a surface profile using a series of Zernike polynomials to approximate $s(\rho) = \Psi(\rho)/2k = -\frac{1}{4\zeta}\log(\zeta\rho^2 + z_c)$. They also added a tilt phase to $s(\rho)$ so as to generate the focus line at a non-zero angle of incidence. The surface was then manufactured using diamond turning of aluminium.

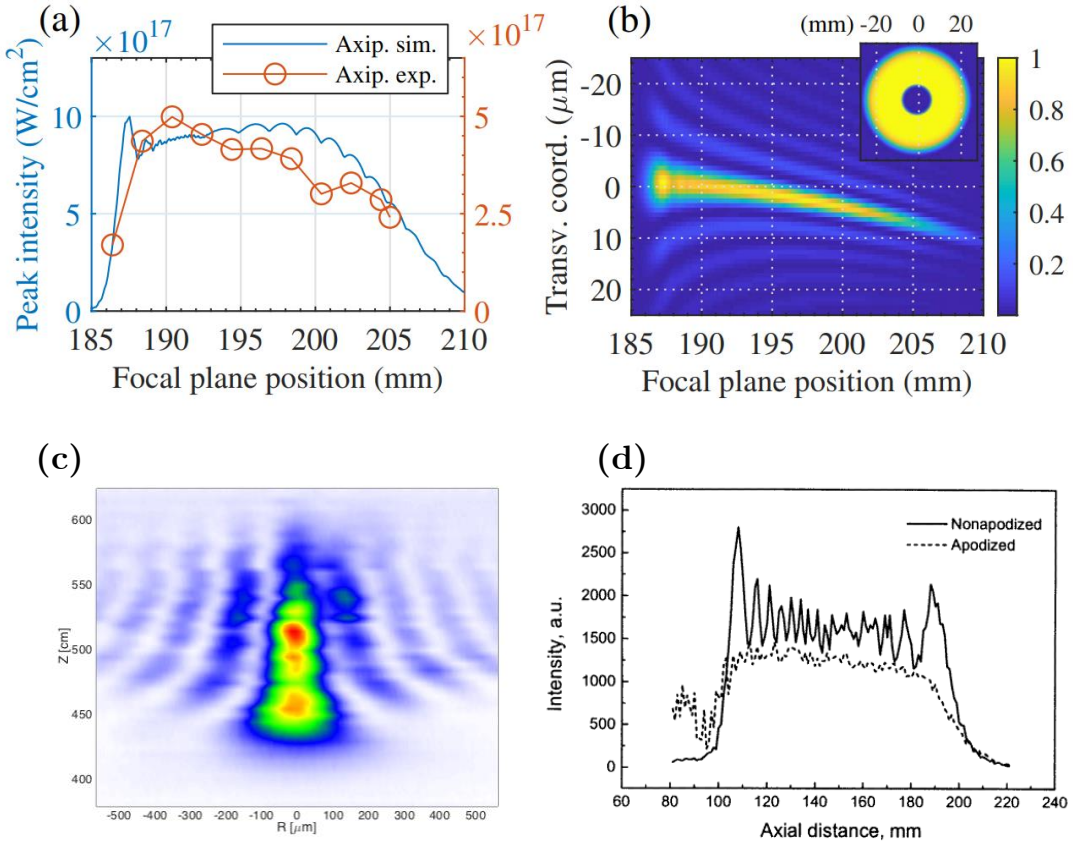


Figure 4.9: Experimental implementations of axilenses from other work; (a) the peak intensity at the focus of the axiparabola as a function of longitudinal position, figure taken from [121]; (b) the simulated intensity profile of the axiparabola focus in the (x, z) plane, with the assumed incident laser profile shown in the inset, figure taken from [121]; (c) the intensity cross section of the focus of the axilens phase plate developed by Gessner *et al.*, figure taken from [141]; (d) the axial intensity profile of the axilens phase plate developed by Sochacki *et al.*, figure taken from [142].

With this optic, the group have demonstrated longitudinally-extended foci of length $\sim$ 20 mm [121] (see Figure 4.9 (a) and (b) reproduced from publication [121]).

As acknowledged by the authors of the publication [121], the axial intensity from the axiparabola decreased with longitudinal position, rather than providing the uniform intensity designed (as can be seen in Figure 4.9 (a)). The focus also did not form a straight line longitudinally as anticipated, but displayed a curvature away from the laser axis (as can be seen in Figure 4.9 (b)). The reasons for these discrepancies do not seem to have been fully understood. Until these issues can be rectified, using axiparabolae for HOFI plasma channel generation over longitudinal extents on the order of 10s of centimetres would appear challenging.

An alternative method would be to use a thin phase-plate axilens. In 2016, Gessner *et al.* developed an axilens phase plate to be used to form hollow plasma channels

over metre length scales [141]. The phase plate was manufactured with 8 levels on a 1 mm thick fused-silica substrate. In this case, the length of the axilens focus was found to be substantially shorter than the designed length (1.3 m as opposed to 2 m), an observation that was not fully understood (see Figure 4.9 (c) reproduced from thesis [141]). Moreover, the axial intensity distribution was not uniform, but showed a ringing pattern as a function of longitudinal position.

The ringing in axial intensity was also observed in initial studies of axilenses published in the 1990s. In 1992, Sochacki et al [142] were able to generate a focal region of length 140 mm using a binary computer-generated hologram with phase approximating the full logarithmic form in Equation (4.22). They also observed the ringing in axial intensity found by Gessner *et al.*. To suppress the ringing, Sochacki *et al.* applied a mask on the aperture to soften the edges of the beam and reduce the diffraction and interference effects from a hard-edge on the beam [145–147]. The axial intensity profiles with and without the aperture mask (referred to as apodisation) are shown in Figure 4.9 (d) taken from the publication [145].

4.5.3 Design of axilens phase plate for HOFI plasma channel generation

In this section, a design for an axilens phase plate for HOFI plasma channel generation is described. The axilens phase is written in the form [145];

$$\Psi(\rho_s) = -\frac{k}{2}\frac{\rho_{\text{out}}^2 - \rho_{\text{in}}^2}{z_L}\log\left(\frac{\rho_s^2 - \rho_{\text{in}}^2}{\rho_{\text{out}}^2 - \rho_{\text{in}}^2}z_L + z_0\right). \tag{4.25}$$

where ρ_{out} is the outer radius of the annular aperture of the incident beam, ρ_{in} is the inner radius of the annular aperture, z_0 is the start of the focal line and z_L is the length of the focal line. In terms of the parameters used in Equation (4.22), $z_c = z_0 - z_L\rho_{\text{in}}^2/(\rho_{\text{out}}^2 - \rho_{\text{in}}^2)$ and $\zeta = z_L/(\rho_{\text{out}}^2 - \rho_{\text{in}}^2)$. From the discussion in Section 4.5.1, the spot size of the Bessel focus is smallest at $z = 2z_c$ where it takes the value;

$$w_0(z = 2z_c) = \frac{2.4048\lambda}{\pi}\sqrt{\zeta z_c}. \tag{4.26}$$

As discussed in Chapter 3.2, it is important to consider the critical angle of the plasma in the design parameters of the optic. In the case of axicon-like optics (such as an axilens), the requirement can be written $\theta_{\text{eff}} \equiv \rho_s/z > \theta_{\text{crit}}$, where

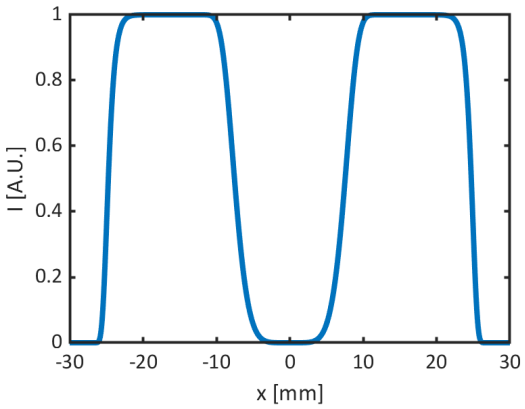


Figure 4.10: Normalised intensity variation as a function of transverse position in the beam used for the axilens numerical calculation.

$\theta_{\text{crit}} = 0.54\sqrt{n_{\text{e,17}}}\lambda_{\mu\text{m}}\,\pi/180$ is the critical angle of rays to the laser axis in radians for a given plasma density $\sqrt{n_{\text{e,17}}}$ in units of $1 \times 10^{17}\,\text{cm}^{-3}$ and wavelength $\lambda_{\mu\text{m}}$ in microns. For a typical plasma density used in HOFI plasma channel formation of $n_{\text{e}} = 3 \times 10^{18}\,\text{cm}^{-3}$, and with $\lambda = 800\,\text{nm}$, the critical angle is $\theta_{\text{crit}} = 2.37°$. For the axilens, a good choice for z_0 therefore would be $z_0 = \rho_{\text{in}}/\theta_{\text{crit}}$ (or z_0 such that $z_0 + z_L = \rho_{\text{out}}/\theta_{\text{crit}}$, whichever is the smallest). This ensures that $\theta_{\text{eff}} > \theta_{\text{crit}}$ for the length of the focus.

An axilens phase plate was designed for experimental parameters similar to those used in the Astra 2018 experiment. The incident beam was modelled with peak intensity $6 \times 10^{11}\,\text{W cm}^{-2}$ and wavelength 800 nm. The incident beam had an annular transverse intensity profile (see Figure 4.10) defined by a difference of super-Gaussians, $I(\rho) = \exp\{-2\,(\rho/\rho_{\text{out}})^{2g}\} - \exp\{-2\,(\rho/\rho_{\text{in}})^{2d}\}$ with $g = 20$ and $d = 3$, inner radius $\rho_{\text{in}} = 9\,\text{mm}$ and outer radius $\rho_{\text{out}} = 25.4\,\text{mm}$.

In the Astra 2018 experiment, a linear axicon was used with approach angle 2.5°. For the beam described above, the axicon focus would have length 375 mm and $w_0 = 7\,\mu\text{m}$. The parameters of the axilens focus were designed to roughly match the focus parameters of the linear axicon, with a minimum spot size $w_0(z = 2z_c) = 6\,\mu\text{m}$ and a focal length $z_L = 400\,\text{mm}$. The value for z_0 of the axilens was set using the constraint on θ_{eff} outlined above, giving $z_0 = 200\,\text{mm}$.

4.5.4 Results of numerical calculation

The ASM code (see Section 4.3) was run for the chosen parameters of the axilens phase plate as described above. The calculation had a transverse resolution of 3.33 µm and a longitudinal step size of 11 mm. The axilens was modelled as having a continuous phase

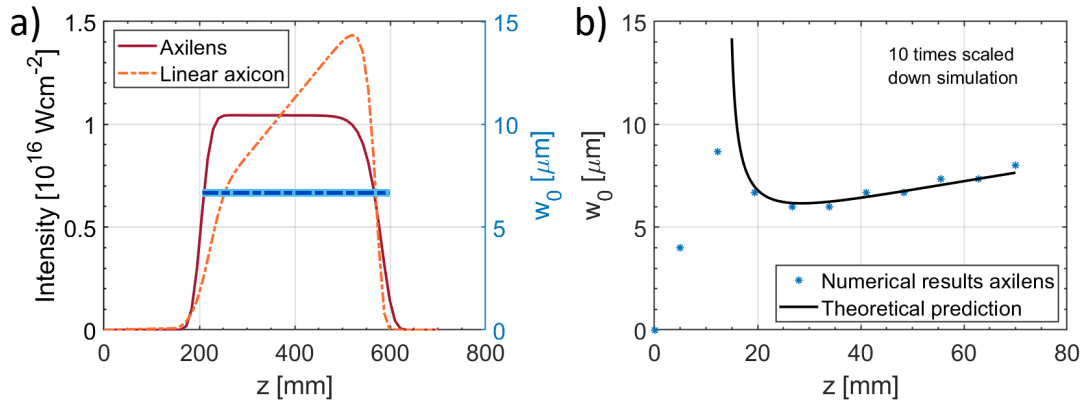


Figure 4.11: Focus properties for an axilens compared to a linear axicon, as generated by the numerical ASM simulation. (a) shows the axial intensity profile and the spot radius w_0 variation with longitudinal position. The incident beam for this calculation was monochromatic, with wavelength 800 nm, and had peak intensity 6×10^{11} W cm^{-2}. The beam had an annular transverse intensity profile with inner radius 9 mm and an outer radius of 25.4 mm, defined by a super-Gaussian-smoothed top-hat profile with $g = 20$ and $d = 3$ (see text). The axilens was defined with $z_0 = 200$ mm and $z_L = 400$ mm and $\rho_{\text{in}} = 9$ mm and $\rho_{\text{out}} = 25.4$ mm. The axicon was defined to have approach angle 2.5°. The simulation had a transverse resolution of 3.33 µm and longitudinal resolution of 11 mm. (b) shows the spot size variation of the axilens with longitudinal position for a 10 times scaled down simulation with higher resolution of 0.67 µm, overlaid with the theoretical curve (Equation (4.24)).

profile that looked like the profile in Figure 4.8 (i.e. consisting of infinitely many layers on the phase plate). The code was run for one wavelength (800 nm) and so chromatic and B-integral effects were not included.

Figure 4.11 shows the result of the numerical calculation. (a) shows the axial intensity profile as a function of longitudinal position in the longitudinally-extended focus of the designed axilens phase plate. For the following discussion we shall take the focus region to be the volume for which $I > 1 \times 10^{15}$ W cm^{-2} (which is $\sim 5\%$ of the peak axial intensity). The focus of the axilens extends between $z = 180$ mm and $z = 610$ mm, close to the designed parameters of $z_0 = 200$ mm and $z_0 + z_L = 600$ mm. The axial intensity in the focus is found to be highly uniform over the full longitudinal extent of the focus. This can be compared to the axial intensity of the focus of a linear axicon with approach angle 2.5° for the same incident beam properties, which increases linearly over the focal region for the same incident beam parameters.

In the context of HOFI plasma channels, the ionisation of the gas to form the plasma column is expected to be roughly independent of intensity provided the ionisation intensity threshold is reached (see Chapter 3). Therefore, both the linear axicon and

axilens case presented here could be used with the incident beam described to generate a HOFI plasma channel of ~ 300 mm length.

However, the flat axial intensity profile of the axilens design enables a more efficient use of the energy in the input beam compared to the linear axicon. To see this,[6] consider the case of a gas with threshold intensity for ionisation $I_{\text{thresh}} = 1 \times 10^{16}\,\text{W cm}^{-2}$. With the axilens focusing optic, the intensity in the incident beam used in this calculation would be sufficient to ionise the gas from $z = 230$ mm to 520 mm, generating a plasma column of length 290 mm. For the axicon to achieve the same length of plasma column, the intensity of the focus would have to be increased by a factor of $\sim 4/3$ (according to the results in Figure 4.11(a)). Hence, the intensity requirements for the incident beam in the axicon case are increased by $4/3$ to achieve a plasma column of an equivalent longitudinal extent. As the spot size of both the axilens and linear axicon foci are equal and constant with z to within 1 μm (as discussed in the next paragraphs), the increased intensity requirement for the linear axicon results in an increased energy requirement, by the same factor of $4/3$. For future LPA designs, where channels may need to be generated over longer length scales than have been used up to now, the energy efficiency gained by using the axilens phase instead of the linear axicon phase becomes even more significant.

A final point to note about the axial intensity profiles is the gradient in axial intensity with longitudinal position as the beam reaches focus (around $z = 200$ mm). In Figure 4.11 (a), it is clear that the axial intensity gradient is greater in the case of the axilens than the linear axicon. This rapid increase in intensity at the start of the focus may be of interest for the truncated channel injection method that has recently been proposed [148]. In the truncated channel injection method, the blast wave generated in the longitudinal direction at the front of the plasma channel is used to form a density transition for electron injection. The additional precision with which the start of the longitudinally-extended focus can be set with the axilens may enable greater control over the properties of this density transition and hence enhanced electron beam quality.

Figure 4.11 (a) also shows w_0 of the longitudinally-extended focus of the axilens compared to the linear axicon. With the resolution of this simulation (3.33 μm), both the axilens and the linear axicon are found to have $w_0 = (6.67 \pm 1.67)$ μm over the full

[6] As the axial intensity is linearly proportional to the incident beam intensity, the threshold value of $2 \times 10^{16}\,\text{W cm}^{-2}$ chosen here is arbitrary and the whole calculation can be scaled to achieve the threshold intensity of the gas species of interest.

length of the focus, as designed. To investigate how the spot size of the longitudinally-extended foci vary on the micron scale, a higher resolution simulation was run by scaling the transverse beam parameters by a factor of $1/10$ ($\rho_{\text{in}} = 0.9\,\text{mm}$, $\rho_{\text{out}} = 2.54\,\text{mm}$) and equivalently scaling the axilens parameters to preserve the angles of the rays to the axis ($z_0 = 20\,\text{mm}$ and $z_L = 40\,\text{mm}$). A resolution of 0.67 µm could then be achieved. The equivalence of this setup to the full-scale simulation can be checked in Equation (4.26), where the factor of $1/10$ in ζ and z_c cancels out.

The spot size variation for the higher resolution simulation is shown in Figure 4.11 (b). The spot size is found to vary from 6 µm to 8 µm over the length of the focal region. This is a relative variation of $\sim 15\%$. The prediction for the variation in w_0 from Equation (4.26) is overlaid with the results of the numerical calculation, demonstrating very good agreement. Therefore, it can be concluded that the spot size variation over the length of the longitudinally-extended focus in the full-scale simulation would have had the same relative variation of 15% over the length of the focus from $z = 200\,\text{mm}$ to 600 mm. The impact of the variation in spot size on the properties of the resulting plasma channel would need to be investigated experimentally and/or in plasma simulations.

In the design here, the angle of the rays to the axis varies between $\theta_{\text{eff}} = 2.4^\circ$ and 2.85°, with the smallest angle occurring at the end of the plasma channel. To extend to longer plasma channels in the same density plasma ($\theta_{\text{crit}} = 2.37^\circ$) would require the input beam to have a larger outer radius. This is equivalent to the requirements of the standard axicon, where the maximum length of the plasma channel is limited by the critical angle and the radial dimensions of the input beam. Therefore, the axilens does not provide a method of extending the length of the longitudinally-extended focus for the same input beam parameters.

In the numerical calculations presented, the inner and outer edges of the annular transverse intensity profile were softened using a super-Gaussian function with $g = 20$ and $d = 3$. The super-Gaussian softening reduced the effects of hard-edge diffraction at the inner and outer edges of the beam, which have been found to contribute to a ringing in the axial intensity profile in previous experimental implementations [141, 142]. To realise soft edges in experiment, serrated apertures could be used to define the inner and outer radii of the annulus, as have been used in [19]. The effects of hard-edges in the transverse intensity profile cannot be well modelled by the numerical ASM code

because of the resulting high spatial frequencies required in the angular domain. As the number of grid points in the angular domain is the same as the number in the spatial domain due to the Fourier transform techniques employed, the simulations would be too computationally expensive. Therefore, different techniques would be needed to understand the severity of ringing in the axial intensity for the axilens case compared to the axicon case for realisable incident beam profiles.

4.6 Conclusion

In this chapter, the chromatic and B-integral effects through an axicon lens have been studied using a numerical optical propagation code written by the author and shown to have a significant impact on the on-axis intensity achievable in a standard HOFI plasma channel setup. Reflective and phase plate axicons could provide an alternative solution to enable more energy efficient generation of HOFI plasma channels in the future.

Axilens phase has been studied as an alternative phase profile to the linear axicon phase for future HOFI plasma channel experiments. A design procedure for generating an axilens phase plate has been presented, alongside simulations of an example parameter set. The axilens is shown to have minimal variation in spot size over the extent of the focal line and highly uniform axial intensity. An axilens phase plate following the design described could enable more energy efficient and longitudinally-uniform HOFI plasma channels for future accelerator stages. Utilising the thin phase plate technique, only a small change to the linear axicon phase plates used for example in Miao *et al*. [128] would be required to generate the axilens phase instead.

Chapter 5

Reducing the Transmission Losses of HOFI Plasma Channels

5.1 Introduction

In the previous chapter, optical methods to form HOFI plasma channels were investigated, including ways to couple more energy from the axicon beam into the extended focus, by reducing losses due to chromatic aberration and increasing the longitudinal uniformity of the axial intensity. In the present chapter, methods to reduce the transmission losses of the HOFI plasma channels themselves are considered. As discussed in Chapter 3, HOFI plasma channels are 'leaky' waveguides and hence the guided pulse loses energy as it propagates. It is important that waveguides used for laser-plasma accelerators (LPAs) have minimal transmission losses so that the guided beam can drive a high amplitude wakefield along the extent of the LPA stage. The rate of energy leakage from a waveguide is typically quantified by the power attenuation length, L_{att}, which is the distance over which the pulse energy in a waveguide mode falls by a factor of e.

Designs for future laser wakefield accelerators for applications, for example the EuPRAXIA design, require accelerator lengths of $\gtrsim$ 100 mm [120]. For the channels to work effectively for such a laser wakefield accelerator, the attenuation lengths of the waveguide should be considerably greater than 100 mm. In the most recent experiments performed using a standard single axicon setup for HOFI plasma channel generation, the attenuation length was measured to be $L_{att} = 84$ mm [1].

A number of methods have recently been proposed to increase the attenuation length of the plasma channels (for details see Section 3.2). Experiments described in [21, 84] demonstrate that an additional laser pulse can be used to deepen the plasma channel, prior to the arrival of the main pulse to drive the wakefield. In one case [21], the additional

laser pulse can be provided by a pre-pulse on the guided beam or the leading-edge of the guided beam itself, as demonstrated by the Oxford Laser Plasma Accelerators Group. This is known as conditioned HOFI plasma channel generation or CHOFI. In another design demonstrated by the group at the University of Maryland [84], the channel-forming beam is split into two parts separated by a time delay, and the second beam is shaped into a ring profile to maximise the energy efficiency of the channel deepening. Both methods provide promising routes to achieving attenuation lengths on the order of tens of metres.

In this investigation, an alternative direction is taken towards improving the waveguide transmission efficiency of HOFI plasma channels, by beam-shaping the channel-forming laser pulse. It is found that increasing the diameter of the focus of the channel-forming laser pulse results in a deeper plasma channel. The HOFI channels formed in this way are expected to have considerably higher energy confinement, enabling the attenuation lengths to be extended beyond the $\sim 100\,\text{mm}$ so far achieved without conditioning.

In this chapter, the design objectives for the extended focus are first determined, using the theory of HOFI plasma channel evolution. Constraints on the current optical setups used for HOFI plasma channel formation are described, leading to a discussion of optical techniques that can overcome these constraints. Methods of achieving the desired extended focus are then tested using the ASM numerical optical propagation code (see Section 4.3) and in an optical experiment. The techniques that best match the design objectives are modelled in a plasma simulation code to determine the improvement in the resulting HOFI plasma channels and decide the final parameters to be tested in a future laser-plasma experiment.

Design objectives

As discussed in Chapter 3, HOFI plasma channels are gradient index waveguides for the drive laser [1, 19, 20]. The refractive index profile of the plasma waveguide, $\eta(r)$, is dominated by the electron number density of the plasma, $n_\text{e}(r)$. HOFI plasma channels can be approximated as parabolic in shape, with an on-axis minimum electron density and a rising channel wall, with Equation 3.26 repeated here for convenience;

$$n_\text{e}(r) = n_{\text{e},0} + \Delta n_\text{e} \left(\frac{r}{R}\right)^2, \qquad r < R \tag{5.1}$$

which corresponds to a refractive index variation via;

$$\eta(r) = \sqrt{1 - \frac{n_{\mathrm{e}}(r)e^2}{m_{\mathrm{e}}\epsilon_0\omega^2}} \tag{5.2}$$

where R is the radial position marking the maximum in the plasma density profile, Δn_{e} is the height of the channel wall at $r = R$ relative to the on-axis density and ω is the angular frequency of the incident laser.

As discussed in Chapter 3, because HOFI plasma channels are not infinite parabolic channels, the modes of the channel are not bound but are known as 'leaky' modes. Each mode has a different attenuation length L_{att}, with higher order modes having shorter L_{att} and hence are radiated away more rapidly [1, 80, 149]. As the channel is increased in depth and width, the channel profile more closely approximates the infinite parabolic channel with true bound modes [77]. Hence, it is expected that increasing the depth of the channels will increase the attenuation length of the modes.

Enhancing the channel in this way will also have the effect of confining more of the higher-order channel modes (c.f. Equation (3.29)). If multiple modes propagate within the channel, the transverse profile of the guided beam will change as a function of longitudinal position, owing to beating between the modes. Therefore, for an LPA, it is desirable to only excite the fundamental mode of the plasma channel. The excitation of higher-order modes can be minimised in practice by tuning the parameters of the guided beam, so that the spatial profile of the focus best matches the shape of the lowest-order channel mode. For a parabolic channel, the beam will be maximally coupled into the lowest order mode if it is a Gaussian beam with spot size equal to the matched spot size of the channel $w_0 = w_m = 1/\sqrt{\pi r_{\mathrm{e}} \Delta n_{\mathrm{e}}}$ where r_{e} is the classical electron radius. It is worth noting that to maintain the matched spot size of the parabolic channel whilst increasing the attenuation length, the degree of parabolic curvature in $n_{\mathrm{e}}(r)$ must be maintained, hence both Δn_{e} and R must be increased simultaneously [77].

To investigate the factors that determine the depth and width of a HOFI plasma channel, the hydrodynamic behaviour of the plasma must be studied. After the initial ionisation of the gas by an extended focus, a column of electrons is formed on-axis with a high temperature (for more details, see Section 3.2). These hot electrons cause a shock front to travel radially outwards, setting up a blast wave which generates the plasma

channel profile (see Figure 5.1). As discussed in Chapter 3, the progression of the blast wave has been found to be well approximated by Sedov-Taylor theory [108, 109];

$$r_s(t) = \frac{(\gamma_C + 1)^2}{\pi^{\frac{1}{4}}} \left(\frac{Z E_K}{m_\mathrm{i}} \right)^{\frac{1}{4}} \sqrt{t + \tau_0} \sqrt{r_{e,0}} \tag{5.3}$$

where $r_s(t)$ is the position of the shock front at time t, Z is the atomic number of the plasma species (assuming full ionisation), E_K is the kinetic energy of the electrons after ionisation, m_i is the mass of the ions of the plasma species, $r_{e,0}$ is the initial radius of the column of ionised electrons and τ_0 is a factor to ensure $r_s(t = 0) = r_{e,0}$. The proportionality constant is approximated here as $(\gamma_C + 1)^2/\pi^{1/4}$ with $\gamma_C = 1.66$ being the specific heat ratio (c_P/c_V where c_P is the heat capacity at constant pressure and c_V is the heat capacity at constant volume) of a monatomic plasma [108]. The radial extent of the plasma channel is then $R \sim r_s(t_d)$ for the chosen delay t_d between the ionisation of the gas and the arrival of the guided beam to the channel entrance.

Equation (5.3) demonstrates that to increase $R \sim r_s(t_d)$ of the plasma channel using the same gas, either E_K, t_d or $r_{e,0}$ must be increased. For optical field ionisation, E_K is fixed for a given species (see Section 3.2). Going to later times, t_d, gives a wider but less deep channel, and is fundamentally limited by the stagnation radius r_f given by Equation 3.23. Hence, to generate channels that are simultaneously wider and deeper, a promising route would appear to be to increase $r_{e,0}$.

The benefit of a larger initial column radius is demonstrated in Figure 5.1(b) using data generated from a hydrodynamic code (for details, see Section 5.4.3). The blast wave for the larger initial column radius in Figure 5.1(b) travels more quickly and has a greater amplitude than the case in Figure 5.1(a). Note that in reality, the amplitudes of the blast waves are expected to be lower than shown here and the shock profile softened as a result of a number of non-local heat loss mechanisms, such as via radiative emission [150]. However, the trend is expected to be correct.

The considerations outlined here suggest that increasing the initial radius of the plasma would be an effective way to generate HOFI plasma channels with lower transmission losses. The aim of this chapter is to investigate different methods for increasing $r_{e,0}$ that would be practical for a laser wakefield accelerator. To make the requirements more concrete, the design objectives for the channel-forming beam focus were defined;

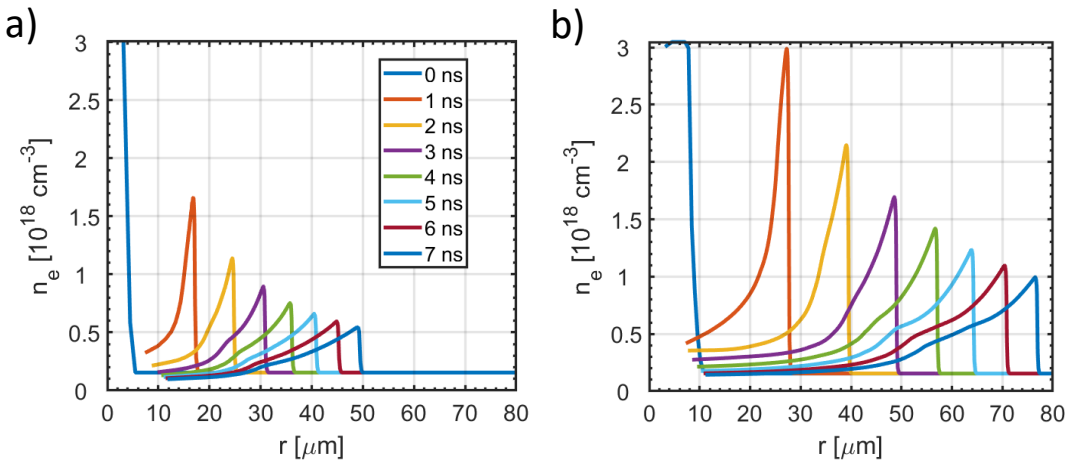


Figure 5.1: Calculated electron density profiles $n_{\mathrm{e}}(r)$, at various delays t, after the formation of an initial plasma column of radius (a) $r_{e,0} = 3\,\mu\mathrm{m}$, (b) $r_{e,0} = 6\,\mu\mathrm{m}$. The simulations were performed using the Helios code for hydrogen at initial pressure 62 mbar (corresponding to $n_{\mathrm{e}} = 3 \times 10^{18}\,\mathrm{cm}^{-3}$) and electron temperature $T_{\mathrm{e}} = 9.7\,\mathrm{eV}$.

- The focus must be able to ionise the plasma to form a column at least 100 mm in length.
- The initial radius of ionisation should be significantly larger than the standard $\sim$ 3 µm radius used previously. [1, 20]. A radius of 10 µm was decided upon, to roughly double the expected shock position at late times ($r_s(t) \propto \sqrt{r_{\mathrm{e},0}}$).
- The region ionised by the channel-forming beam should have a high degree of uniformity, both azimuthally and longitudinally.
- The focus must be practical to realise in a future LPA.
- The laser energy per unit length required to form the channel should not be too large.

5.2 Increasing the spot size of longitudinally-extended foci

In this section, the constraints on the width of an extended focus from the axicon and plasma parameters are described. Two optical setups are presented which overcome these constraints, namely higher-order Bessel modes and random phase plates.

5.2.1 Constraints on standard single axicon setup

As discussed in Chapter 4, experiments for HOFI plasma channels have typically used an axicon optic to focus the channel-forming beam. An axicon has a conical surface which bends all incident rays so that they approach the optical axis with the same angle, θ (see Figure 4.2). Hence the rays which start at different radii focus on the axis at different longitudinal positions $z = \rho/\tan\theta \simeq \rho/\theta$. For this discussion, θ is small such that the paraxial approximation applies. The field profile of a tophat beam focused by an axicon can be calculated using the Fresnel approximation and method of stationary phase (see Chapter 4), and is repeated here for convenience;

$$\begin{aligned} U(r,z) &= -\mathrm{i}CU_0(2\pi kz)^{1/2}\theta\exp\left(\frac{\mathrm{i}\pi}{4}\right)\exp(\mathrm{i}kz)\exp\left(-\frac{\mathrm{i}k\theta^2}{2}\right) \\ &\quad\times\exp\left(\frac{\mathrm{i}kr^2}{2z}\right)J_0(k\theta r) \\ &= A_z(z)\exp\left(\frac{\mathrm{i}kr^2}{2z}\right)J_0(k\theta r) \end{aligned} \tag{5.4}$$

where J_0 is the zeroth order Bessel function of the first kind, and C is a constant dependent on the transmission properties of the axicon [151]. Equation (5.4) shows that the intensity of the axicon focus increases linearly with longitudinal position for an incident tophat beam. The radius of the first zero of the Bessel function is given by $w_0 = 1.2\lambda/\pi\theta$. In a plasma, w_0 and $A_z(z)$ set the radius of ionisation $r_{e,0}$ (see Section 3.2).

To increase the spot size of the focus w_0, one must either increase the wavelength of the radiation used for the channel formation beam, λ, or choose an axicon with a reduced approach angle θ. However, the critical density of the plasma puts a constraint on the minimum value that θ can take for a given wavelength λ (as explained in Section 3.2). When $\theta = \theta_{\mathrm{crit}} \propto \sqrt{n_{\mathrm{e}}}\lambda$, the λ dependence of the spot size falls out to give $w_0 \propto 1/\sqrt{n_{\mathrm{e}}}$. That is, the maximum spot size achievable by an axicon focus in a plasma is determined by the plasma density alone. In the setups typically used for HOFI plasma channel generation, $n_{\mathrm{e}} \simeq 3\times10^{18}\,\mathrm{cm}^{-3}$, $w_0 \simeq 7\,\mu\mathrm{m}$, $\lambda = 0.8\,\mu\mathrm{m}$ and $\theta = 2.5^\circ$ ($\theta_{\mathrm{crit}} = 2.36^\circ$ for $n_{\mathrm{e}} = 3\times10^{18}\,\mathrm{cm}^{-3}$).

The electron number density for the background plasma is fixed by requirements of the laser wakefield accelerator, for example through the dephasing length as described in Chapter 2. Hence, it is not possible to increase the radius of the focus using the

standard setup beyond the limit set by the critical angle. It is already the case that in recent experiments, the approach angle has been set only slightly larger than the critical angle [1, 20]. Therefore, further improvements to the focus cannot be achieved using the parameters already available, namely θ, n_e and λ.

Other optics are also used to form longitudinally-extended foci, such as axilenses and axiparabolae, as discussed in Chapter 4. Where the axicon phase is linear with radius, axilenses and axiparabolae have phases $\Psi \propto \log(\zeta\rho^2 + C)$ [121, 140]. The foci of these optics must satisfy the same minimum angle constraints outlined above. Therefore, the transverse intensity profile of the axilens focus $\propto J_0^2(kr\rho/z)$ also has a maximum spot-size $w_0 \propto \lambda z/\rho \propto 1/\sqrt{n_\mathrm{e}}$.

Considerable research has been conducted into methods to control the transverse and longitudinal profiles of extended foci [152–154]. By using a combination of zeroth order Bessel beams it is possible in principle to simultaneously generate arbitrary transverse and longitudinal profile shapes [155]. However, for the case of HOFI plasma channel formation, the minimum critical angle described above corresponds to a fundamental maximum transverse scale size of variations in the focus (minimum transverse spatial frequency present in the intensity profile). We can see this by returning to the general form of the focus of an optic with azimuthal symmetry which, for monochromatic waves in the paraxial regime, is given by Equation (4.5) from Chapter 4;

$$\begin{aligned} U(r,\phi) &= \frac{2\pi}{\mathrm{i}\lambda z}\exp(\mathrm{i}kz)\exp\left(\frac{\mathrm{i}kr^2}{2z}\right)\int_0^\infty U(\rho)\exp\left(\frac{\mathrm{i}k\rho^2}{2z}\right)J_0\left(\frac{kr\rho}{z}\right)\rho d\rho \\ &= \frac{\lambda z}{2\pi\mathrm{i}}\exp(\mathrm{i}kz)\exp\left(\frac{\mathrm{i}kr^2}{2z}\right)\int_0^\infty U(k_\rho z/k)\exp\left(\frac{\mathrm{i}k_\rho^2 z}{2k}\right)J_0\left(k_\rho r\right)k_\rho dk_\rho. \end{aligned}$$

where we have made the substitution for $k_\rho \equiv k\rho/z \equiv k\theta$. When written this way, the integral takes the form of a zeroth-order Hankel transform;

$$f(r) = \int_0^\infty F(k_\rho)J_0\left(k_\rho r\right)k_\rho dk_\rho.$$

Just as in the case of a Fourier transform, the largest scale features in the output of a Hankel transform are given by the lowest frequency components in the function $F(k_\rho)$. To achieve large scale features in our focus therefore, we need contributions in the field that have a small value of k_ρ, which in turn corresponds to a small value of θ. Hence, in general, the transverse scale of features in the focus is fixed by the smallest angles of rays allowed in the plasma, given by the critical angle θ_crit above. Combinations of zeroth

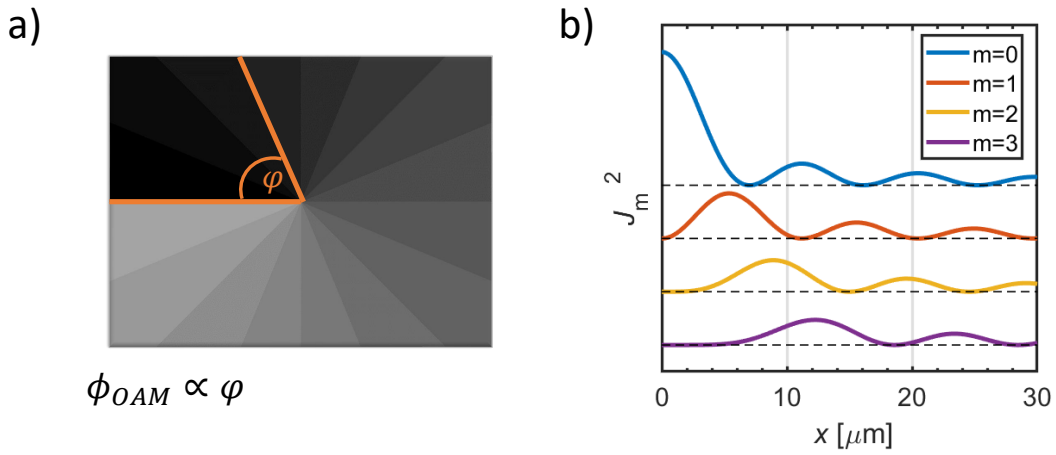


Figure 5.2: Higher order Bessel mode focus patterns. (a) The discretised spiral phase applied to the wavefront of the laser to generate orbital angular momentum (OAM). (b) Normalised transverse intensity profiles formed at the focus of an axicon for different orders, m, which take the form of Bessel modes, $J_m(k\theta r)^2$. The cases drawn are for a wavelength of 0.8 µm and approach angle $\theta = 2.5°$.

order Bessel beams with larger angles, and hence smaller transverse scales, cannot be used to generate a larger focal spot than that of the axicon at the critical angle.

In the following sections, methods are considered which overcome this constraint by (a) using Bessel foci beyond the zeroth order, and (b) by using random phase plates to generate speckle pattern foci.

5.2.2 Higher order Bessel modes

Higher order Bessel modes (HOBMs) are a set of diffractionless beams which can be formed by an axicon [156]. Unlike the standard zeroth-order Bessel function described above, the HOBMs have a minimum on-axis surrounded by a bright ring profile. Such profiles can be generated at the focus of a linear axicon or axilens by applying a spiral phase to the laser pulse, which increases azimuthally from 0 to $2\pi m$, i.e. $\phi_{\text{OAM}}(\rho, \varphi) = m\varphi$ where φ is the azimuthal coordinate (see Figure 5.2(a)). The wavefront of the pulse is then said to carry orbital angular momentum (OAM), with topological charge denoted by integer m.[1]

A laser beam with orbital angular momentum $m \in \mathbb{Z}^+$ that is passed through an axicon will form a longitudinally-extended focus with electric field profile;

$$U(r, \phi, z) = A_z(z) \exp\left(\frac{\mathrm{i}kr^2}{2z}\right) \exp(\mathrm{i}m\varphi) J_m(k\theta r) \tag{5.5}$$

[1]OAM beams with non-integer m are also achievable, but result in transverse intensity profiles that are not azimuthally symmetric [157].

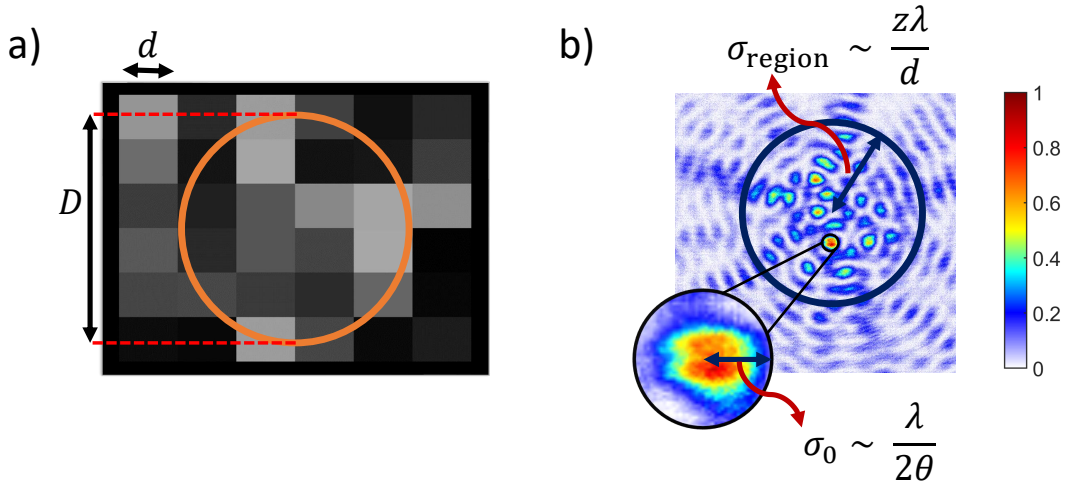


Figure 5.3: Generation of speckle patterns. (a) Diagram of laser beam of aperture diameter D incident on a random phase plate with element size d. (b) Measured transverse intensity profile of an example speckle pattern generated in experiments with a spatial light modulator and an axicon for approach angle $\theta = 2.5°$. x and y axes go from $-80\,\mu m$ to $+80\,\mu m$ (see Section 5.3.2).

where J_m is the m^{th} order Bessel function of the first kind. The first three Bessel functions are depicted in Figure 5.2(b). It can be observed that as m is increased, the ring maxima are of lower amplitude and have a larger radius. The width of the ring also increases with m, such that the overall area of the ring goes roughly as m^2. Hence, HOBM beams have the possibility of generating a larger ionisation region than the standard zeroth order beam, without varying the angle of approach of the rays to the axis, θ.

It would be possible to combine HOBMs of different m-values with different amplitudes to generate foci of differing transverse profiles. However, due to the azimuthal phase variation on the field at focus, coherent combinations of Bessel foci of different m will interfere as a function of the azimuthal coordinate φ and hence the resulting foci will no longer be azimuthally symmetric. This interference can be avoided by combining the beams incoherently, for example via a time separation. The foci generated are then of the form;

$$I(r,z) = \sum_m |A_m A_z(z)|^2 J_m(k\theta r)^2 \tag{5.6}$$

where A_m denote the relative amplitudes of the different orbital angular momentum modes.

5.2.3 Random phase plates

A method conventionally used to increase the spot size of a Gaussian beam is to generate a speckle pattern at the focus [158], as for example in Figure 5.3(b). This can be achieved by combining a focusing lens with a random phase plate (RPP), which offsets different parts of the incident beam profile by a different random phase between 0 and π (depicted in Figure 5.3(a)). Each speckle of the pattern has the original diffraction-limited spot size of the focusing system, $\sigma_0 \sim f\lambda/D$ for D the diameter of the optic [159]. The speckles are spread out over a region defined by the properties of the random phase plate. In particular, if the RPP is segmented into square elements of size d, the resulting distribution of speckles in the focus will occupy a region $\sigma_{\text{region}} \sim f\lambda/d$ at the focus of a lens with focal length f. Therefore, the RPP causes the region size to increase by a factor $\sim D/d$. It is desirable for the number of speckles over the bright region to be as large as possible to smooth out the speckling effect, hence D/d tends to be $\gg 1$.

A similar idea can be used for the case of an RPP combined with an axicon [160, 161]. However, in this case, the equations for the diffraction-limited and region size are adjusted to the parameters of the axicons focus. In particular, we know from geometric optics that for a lens $f/D = 1/(2\theta_{\text{lens}})$ where θ_{lens} is the largest incident angle of rays to the axis. The spot size of the focus generated by a lens can therefore be written $w_{0,\,\text{lens}} \sim f\lambda/D \sim \lambda/(2\theta_{\text{lens}})$. For an axicon $\rho/z = \theta$ is a constant and the focus has spot size $w_0 \sim \lambda/(2\theta)$. Each speckle in the focus of an axicon therefore has scale size $\sigma_0 \sim \lambda/(2\theta) \sim w_0$, i.e. the individual speckles in the pattern have the transverse length scale of the original Bessel focus pattern. This dependence has been shown by Belyi *et al.*, [161]. The study by Belyi *et al.* did not give an expression for the region size of the speckle pattern in the axicon case. However, continuing the analogy to the lens case, we might expect the transformation to be $f \equiv z$, giving $\sigma_{\text{region}} \sim z\lambda/d$. This relation predicts that the speckle region size would increase linearly with z for fixed element size on the random phase plate, d.

5.3 Optical studies using spatial light modulator

In order to investigate HOBMs and speckle patterns for applications to HOFI plasma channels, a number of studies were made. The numerical beam propagation code devel-

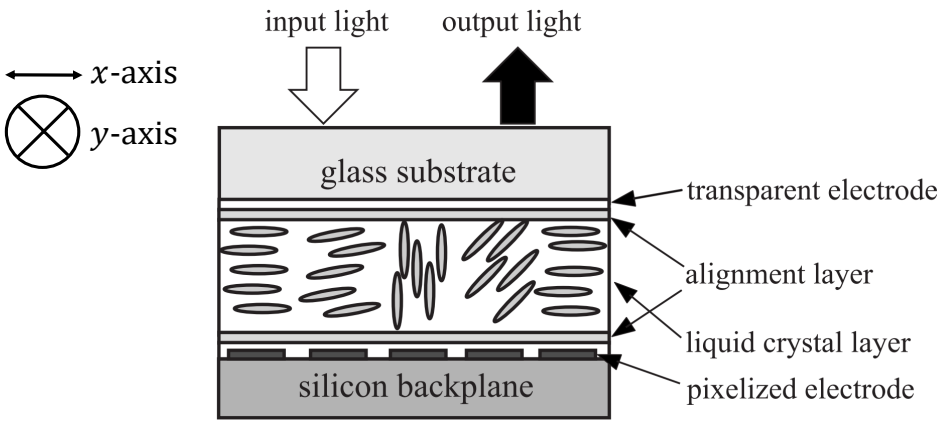


Figure 5.4: Diagram illustrating important aspects of liquid crystal on silicon SLM chip, reproduced from publication with permission [162].

oped by the author (described in Section 4.3) was used to investigate the optical setups required to generate the desired focal spot qualities [24]. The focal spots were then generated experimentally using a spatial light modulator and helium-neon laser setup.[2] The focus parameters that best fit the design objectives for the longitudinally-extended focus were then tested in hydrodynamic codes to model the ionisation and subsequent evolution of the HOFI plasma channels.

In this section, the optical studies using the ASM code and spatial light modulator are described. The hydrodynamic simulations are presented in Section 5.4.3.

5.3.1 Spatial light modulator setup and calibration

A spatial light modulator (SLM) is a device which can be used to apply user-defined phase profiles to a wavefront. The SLM chip typically consists of a medium with a voltage-dependent refractive index, such as a liquid crystal. The working principle of a liquid crystal on silicon (LCOS) SLM is displayed in Figure 5.4. Pixelised electrodes on the SLM chip apply a position-dependent voltage with resolution limited by the transverse dimension of the electrode pixels, typically on the order of 20 microns. The laser beam is incident normally onto the SLM chip, passes through the liquid crystal layer and develops a phase dependent on the voltage applied at each position. The beam then reflects off the back surface and traverses the liquid crystal again before leaving from the front surface of the SLM chip.

[2]Differences in the properties of the focus generated with the continuous wave helium-neon laser compared to a short-pulse titanium-sapphire laser would result both from the difference in central wavelength of the two systems (633 nm compared to 800 nm respectively) and the broad spectrum of the titanium-sapphire laser, leading to blurring of the transverse features by a factor of $\Delta\lambda/\lambda_0 \sim 30/800 \lesssim 5\%$.

The LCOS SLM displays birefringence due to the anisotropy of the liquid crystal molecules, which have a rod-like shape. The applied voltage changes the extraordinary refractive index but not the ordinary refractive index of the medium, i.e. the voltage causes a rotation of the liquid crystal director about the axis marked y in Figure 5.4. Therefore, only light polarised parallel to the x-axis will be modulated by the spatial light modulator. The maximum amount of phase gained by the wavefront through the liquid crystal layer can be written, $\phi_{\text{SLM, max}} = 2k\eta_{\text{ex}}T$ where k is the wavenumber of the light, η_{ex} is the extraordinary refractive index of the liquid crystal and T is the thickness of the crystal layer (the minimum amount is $\phi_{\text{SLM, min}} = 2k\eta_{\text{o}}T$ where η_{o} is the ordinary refractive index of the liquid crystal). The actual phase applied as a function of position on the chip, $\phi_{\text{SLM}}(x, y)$, depends on the rotation of the LCD in each pixel. The perpendicular polarisation (along the y-axis) will gain a trivial phase independent of position as it traverses the liquid crystal layer, $\phi_{\perp} = 2k\eta_{\text{o}}T = \text{const.}$

The spatial light modulator used for this study was a Hamamatsu X10468-02. The chip size was 15.8×12 mm, consisting of 792×600 pixels. A helium-neon laser, $\lambda =$ 633 nm, was used for the experiments. We take the electric field of the laser pulse to have polarisation purely in the plane of the SLM chip (the (x, y) plane).

To use the SLM effectively, the phase applied must be calibrated for the wavelength of the laser. The calibration is achieved by sending a laser pulse with polarisation 45° to the x-axis (i.e. split equally along x and y) and supplying a homogeneous phase mask to the SLM. The phase on the SLM is only applied to the electric field component with polarisation parallel to the $y-$axis and hence the two components of polarisation change their phase relationship. This results in a change in the polarisation of the field leaving the SLM, which can be converted to an intensity readout using a simple optical setup.

The experimental layout for the phase calibration of the SLM is shown in Figure 5.5(a). The beam was expanded using a two-lens telescope to a diameter of 21 mm prior to the diagram in Figure 5.5(a). The expanded beam was first sent through a linear polariser (LP1), set at 45° to the x-axis. The beam passed through a beamsplitter onto the SLM at normal incidence and was reflected. The reflected beam was then sent back to the beamsplitter to reflect towards the second linear polariser (LP2), set parallel to LP1, before being focused onto the camera (*Allied Vision Mako G-503B*).

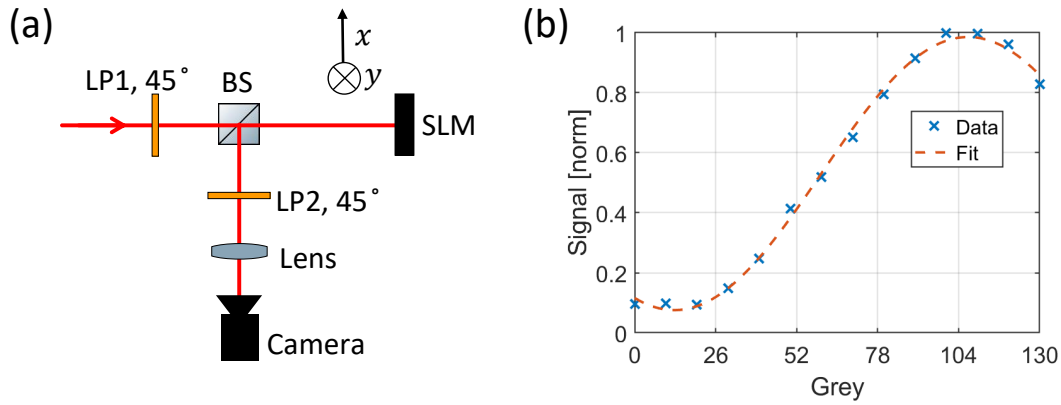


Figure 5.5: Phase calibration of the spatial light modulator (SLM). (a) Experimental layout for performing calibration, where LP1 and LP2 are linear polarisers set at 45° to the liquid crystal director of the SLM and BS is a beamsplitter so that the SLM can be used at normal incidence. (b) Signal measured on the camera as a function of Grey value of the bitmap supplied to the SLM. The data is fitted with a cosine function calibration curve.

Phase was applied to the SLM using a bitmap array of values, denoted here as Grey values, which could be varied between 0 and 255. Depending on the Grey value supplied, the phase accumulated by the light polarised parallel to the x-axis, ϕ_{SLM}, was varied with respect to that of the y-axis, $\phi_{\perp}$. When passed through LP2, the change in polarisation resulted in a change of intensity and this was detected by the camera.

The results of the phase calibration scan can be seen in Figure 5.5(b). Fitting the result to a cosine curve gave the calibration factor ξ in $\phi_{\text{SLM}} = \xi \times Grey$, as $\xi = \frac{\pi}{94 \pm 5}$, where the error is determined by 95% confidence intervals on the fit.

5.3.2 Investigation of speckle patterns

Numerical calculations were performed using the ASM code described in Section 4.3 to model a random phase plate in combination with an axicon. First, the qualitative features of the focal intensity distributions generated with an RPP and axicon were compared to the foci achieved by an RPP and lens. The ASM code was run to model the effect of an RPP with element size $d = 200\,\mu\text{m}$. The incident beam had wavelength 800 nm and a Gaussian transverse intensity profile with radius 20 mm. The lens was modelled with focal length $f = 175\,\text{mm}$ and the axicon with approach angle $\theta = 2.5°$. The transverse intensity profile of the focus was output at $z = 175\,\text{mm}$ in each case. The resolution of the simulation was $2.5\,\mu\text{m}$.

Figure 5.6 shows the resulting focus profiles. In (a), the speckle pattern at the focus of a lens is shown, demonstrating the tightly clustered arrangement of speckles, extend-

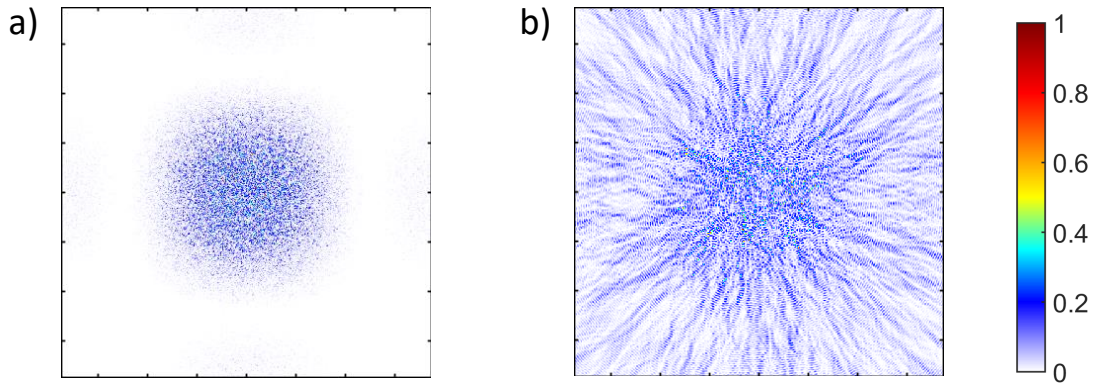


Figure 5.6: Speckle pattern foci calculated using the ASM code for a random phase plate with element size $d = 200\,\mu m$ at the focus of (a) a lens of focal length 175 mm; (b) a linear axicon with approach angle 2.5°. The incident beam had wavelength 800 nm and Gaussian transverse intensity profile with beam radius 20 mm. The propagation distance was 175 mm in each case. The ASM simulation had resolution 2.5 µm. The x and y axes of both images have lengths of 1.5 mm with ticks spaced by 0.2 mm and the colour axis is normalised to the peak counts of the image in each case.

ing over a region of $\sigma_{\text{region}} \sim 500\,\mu m$, qualitatively matching the speckle patterns found in experimental demonstrations of RPPs in conjunction with a focusing lens [158]. In comparison, at the focus of an axicon with the same RPP, shown in (b), the pattern consists of chains of speckles. The contrast between the central region and the surrounding area is lower than in the lens case, with speckles with normalised count values > 0.2 extending far outside of the central region. From these images, it is clear that the speckle patterns at the focus of an axicon are less like a spot than at the focus of a lens. Nevertheless, due to the threshold nature of the OFI ionisation mechanism, it was of interest to investigate whether there was a parameter regime for RPPs which would generate the design objectives in Section 5.1.

The region size achieved in Figure 5.6(b) with $d = 200\,\mu m$ and axicon approach angle $\theta = 2.5°$ was too large for the design objectives. Therefore, the element size and the axicon approach angle were both increased in an attempt to scale down the speckle region size. The increase in axicon approach angle could not be too large, as sharper angled axicons result in shorter lengths of longitudinally-extended foci for the same input beam (the axicon focal line length $z_L \sim \rho_{\text{out}}/\theta$ in geometric optics for ρ_{out} the outer radius of the incident transverse intensity profile). A compromise was made between the number of speckles and the length of axicon focus, resulting in an axicon approach angle of 9° to achieve a speckle size of $\sigma_0 = 1.5\,\mu m$. This was expected to generate a reasonable number of speckles (~ 10) within the 10 µm region, whilst shortening the

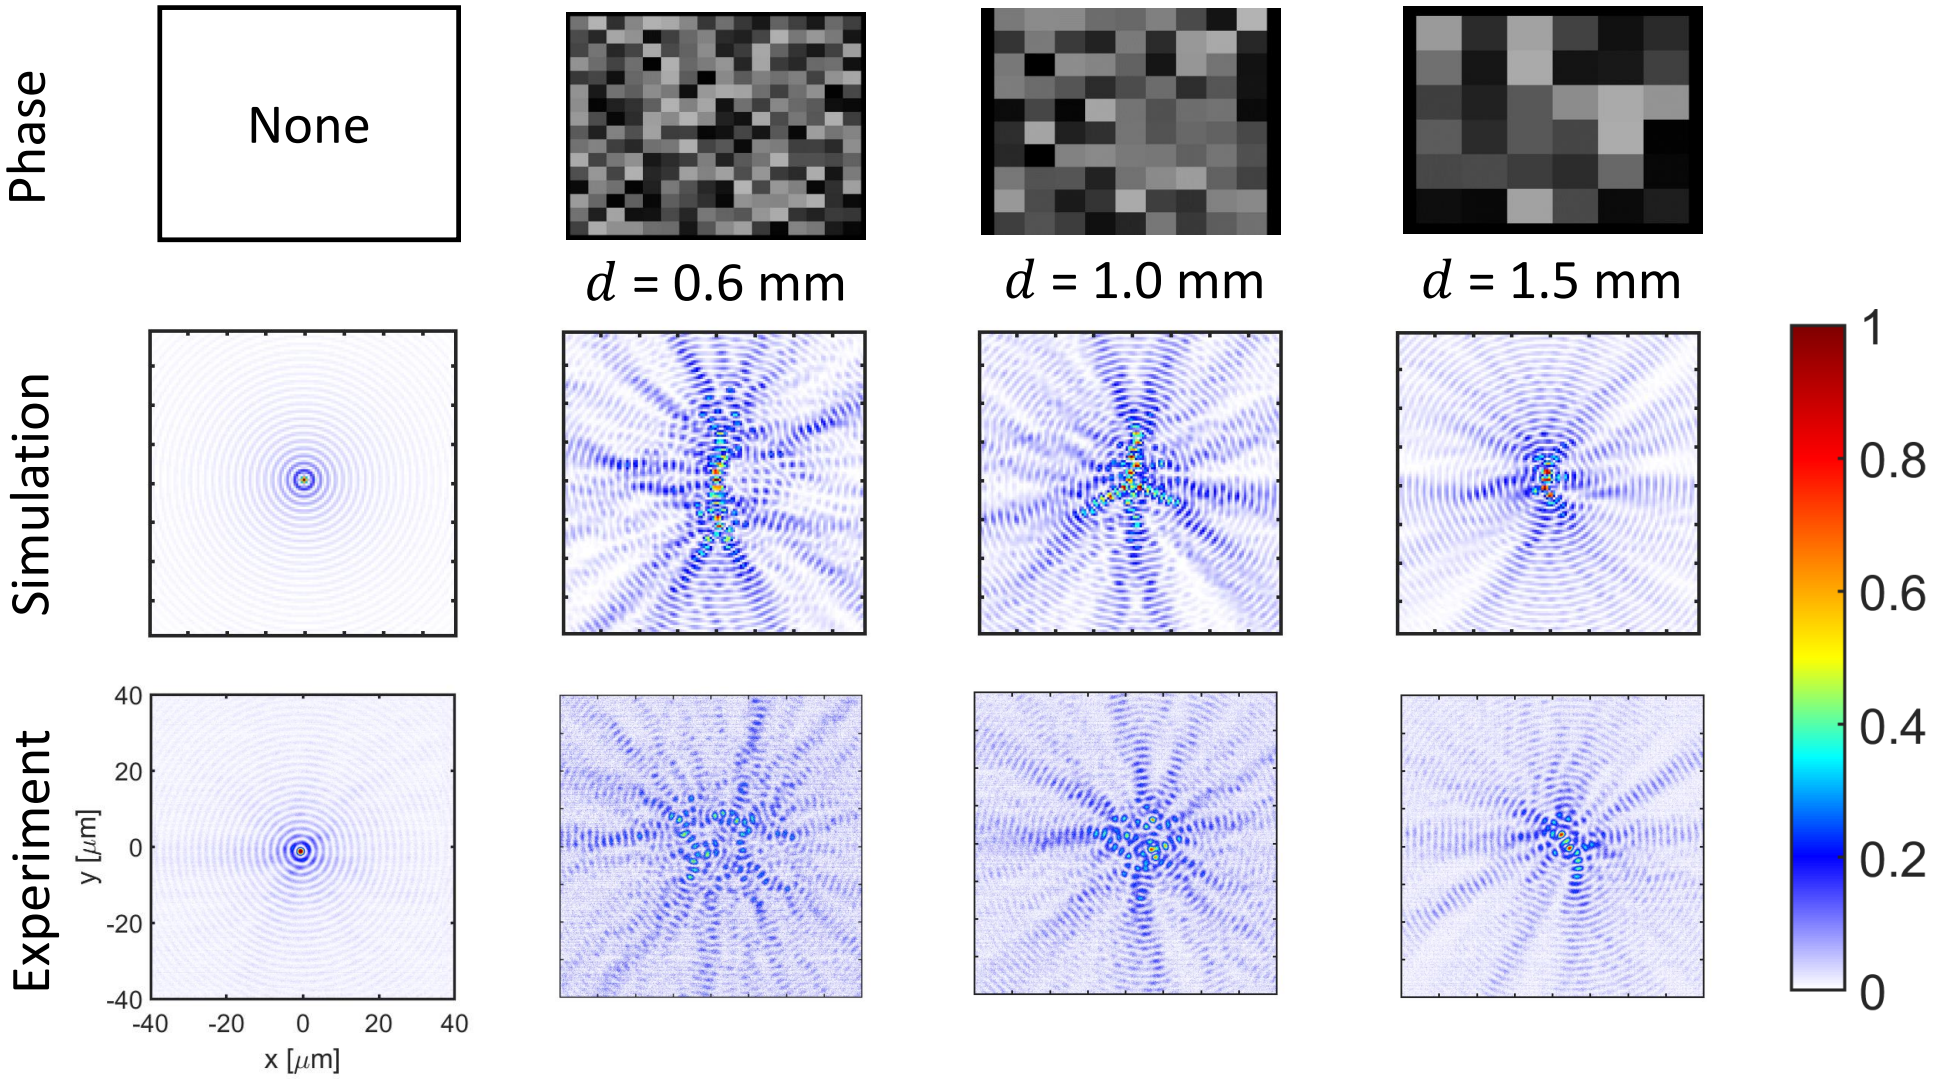


Figure 5.7: Focal spots generated at the focus of a linear axicon, base angle 20°, after application of a random phase plate with element size d to the wavefront. Foci are imaged at $z = (40 \pm 2)$ mm with a beam of wavelength $\lambda = 0.633\,\mu$m. Top row: bitmaps of the random phase distributions used. Middle row: results of numerical calculations using ASM code. Bottom row: focal spots imaged in SLM experiment, with parameters matching those of the ASM simulation. Images in the middle row are self-normalised, images in the bottom row are normalised to the peak of the experimental J_0^2 image. The x and y axes on the middle and bottom row images go from $-40\,\mu$m to $+40\,\mu$m.

achievable length of focus by a factor of ~ 4.

For the following ASM simulations, the incident laser beam was modelled[3] to approximate the SLM experiment described above, namely as having a Gaussian transverse intensity profile with $1/e^2$ diameter of 21 mm and wavelength 633 nm. Scans of element size d were performed to find values that would generate spot region sizes of approximately 10 µm at a longitudinal position $z = 40$ mm. Element sizes of 1 mm were found to give the speckle region size desired. The cases of $d = 0.6$ mm, 1.0 mm and 1.5 mm were chosen to be tested experimentally.

The experimental setup for the study of speckle pattern foci was the same as for the phase calibration of the SLM (see Figure 5.5(a)), except that LP1 was set to 0° such that the laser polarisation was purely along the x-axis, LP2 was removed and the lens was replaced with an ultraviolet-grade fused silica (UVFS) refractive axicon lens. The axicon base angle was $\alpha = 20.0°$, to achieve the chosen approach angle $\theta = 9°$. A pinhole was placed within the beam expander to remove high-frequency noise in the near field of the beam. The camera (*Allied Vision Mako G-503B*) had a $40\times$ microscope objective attached. A US Air Force target was used to perform a spatial calibration of the camera, giving (0.090 ± 0.005) µm per pixel in the horizontal direction and (0.086 ± 0.005) µm per pixel in the vertical direction. In the following, a value of 0.09 µm per pixel is used in both horizontal and vertical.

The results for random phase plates with element sizes $d = 0.6$ mm, 1.0 mm and 1.5 mm are shown in Figure 5.7. The distance between the axicon and the imaged focus here was set to $z = (40 \pm 2)$ mm. Figure 5.7 also displays the results of the numerical calculations alongside the experimental focal spots. For the numerical calculations, the ASM code was used with a spatial resolution of 0.6 µm.

The focal spots for the numerical and experimental cases do not match perfectly for the random phases applied, despite using identical phase masks in both. This is likely due to the sensitivity of the focus to phase non-uniformities across the beam, which were not measured in the experimental setup and hence could not be modelled in the ASM code. Nevertheless, the speckle size and the trends in spot region size appear qualitatively similar for the numerical and experimental foci. Both show significant azimuthal

[3]Note that the transverse profile of the incident beam mostly affects the longitudinal intensity profile of the longitudinally-extended focus of an axicon and is not expected to have any significant effect on the transverse properties of the focus.

asymmetry, as seen in previous experimental implementations of speckle patterns over longitudinally-extended foci [160].

The spot region size σ_{region} of the speckle patterns in the case of an axicon is not well defined, as noted above. To get an idea of how the spot region size varies with element size d, σ_{region} was quantified for the experimental data shown in Figure 5.7 using the following method. The centre of the image was set using the reference image, taking the position of the peak of the J_0^2 beam. The imaging system was not moved between the successive images so that the coordinate of the pixel labelled as the centre position was the same for each image. A background was subtracted from the image data and a median filter applied to reduce the electrical noise (median value taken over a 10 by 10 square of pixels, about each pixel position). The parameter σ_{region} was then defined as the outermost radial position containing a pixel with counts greater than half the maximal counts in each image. This gave $\sigma_{\text{region}} \sim 23\,\mu\text{m}$, $13\,\mu\text{m}$ and $8\,\mu\text{m}$ for the cases of $d = 0.6\,\text{mm}$, $1.0\,\text{mm}$ and $1.5\,\text{mm}$ respectively.

As discussed in Section 5.2, there were concerns that RPPs would generate different speckle patterns as a function of longitudinal position from the axicon tip. Therefore, images were taken at a few different longitudinal positions to visualise the evolution of the speckle region. Figure 5.8 shows the experimentally measured focal regions at $z = (25 \pm 2)\,\text{mm}$, $(36 \pm 2)\,\text{mm}$ and $(45 \pm 2)\,\text{mm}$ for the case of $d = 1.0\,\text{mm}$. Going from (a) to (c), the region size of the speckles is found to increase as a function of longitudinal position as anticipated.[4] The focal region in (a) at $z = 25\,\text{mm}$ consists of only a few speckles with counts greater than half the maximum counts in the image. The image in (b) at $z = 36\,\text{mm}$ shows a clustering of spots within a roughly $10\,\mu\text{m}$ spot region, most closely matching the desired specification. By $z = 45\,\text{mm}$ in (c), the speckles are spread further from each other and appear to form longer chain-like structures as opposed to a clustered spot region. Note that the images in (a) to (c) are self-normalised, which removes the effect of the axial intensity variation with z given by an axicon. In practice, both the variation in axial intensity and the variation in speckle pattern region size with longitudinal position would contribute to the non-uniformity of the longitudinally-extended focus. Methods that could be used to reduce the variation

[4] Note that in the images in Figure 5.8, the centre of the images are not meaningful because the camera was physically moved between each z-position.

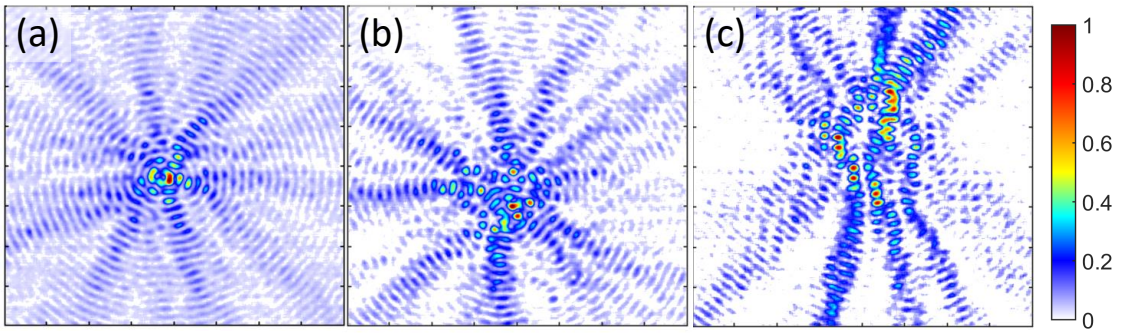


Figure 5.8: Measured focal spots generated at the focus of a linear axicon, base angle 20°, after application of a random phase plate on the SLM with element size $d = 1.0\,\text{mm}$ to the wavefront, imaged at varying longitudinal positions from the axicon; (a) $z = 25\,\text{mm}$; (b) $z = 36\,\text{mm}$; (c) $z = 45\,\text{mm}$ (where z values are given to $\pm\,2\,\text{mm}$). The pixel counts are self-normalised by dividing by the peak pixel count in each image. The x and y axes have a range of 72 µm.

of speckle region with longitudinal position are discussed in Section 5.3.4.

The speckle foci consist of significant areas of darkness in between the speckles. The gas in these areas would not be ionised by the laser beam. The effect of this initial ionisation distribution is difficult to predict due to the complexity of the ionisation and heating mechanisms involved. To reduce the fraction of the focal area that is dark, smoothing techniques can be used. One way to do this is to incoherently combine two or more beams with different random phase patterns to generate a higher density of speckles, a method known as spatial incoherence [158]. The resulting foci are expected to be more uniform azimuthally. It was not possible to use spatial incoherence smoothing in the experiments because it would require more than one SLM. Smoothing methods increase the number of beams that need to be used to form the focus, increasing the complexity and energy requirements of the channel-forming beam setup (see discussion in Section 5.3.4).

5.3.3 Investigation of higher-order Bessel modes

The experimental setup for the study of higher-order Bessel modes was the same as that used for the speckle patterns, except that the axicon of base angle 20.0° was replaced with an axicon of base angle 5.0°, comparable to that typically used for HOFI plasma channel experiments [1, 20]. For the standard zeroth order Bessel mode at wavelength 0.633 µm, this would give $\theta = 2.3°$ and $w_0 = 6\,\mu\text{m}$.

The spiral phase bitmaps for the SLM were designed to have 16 discrete levels of

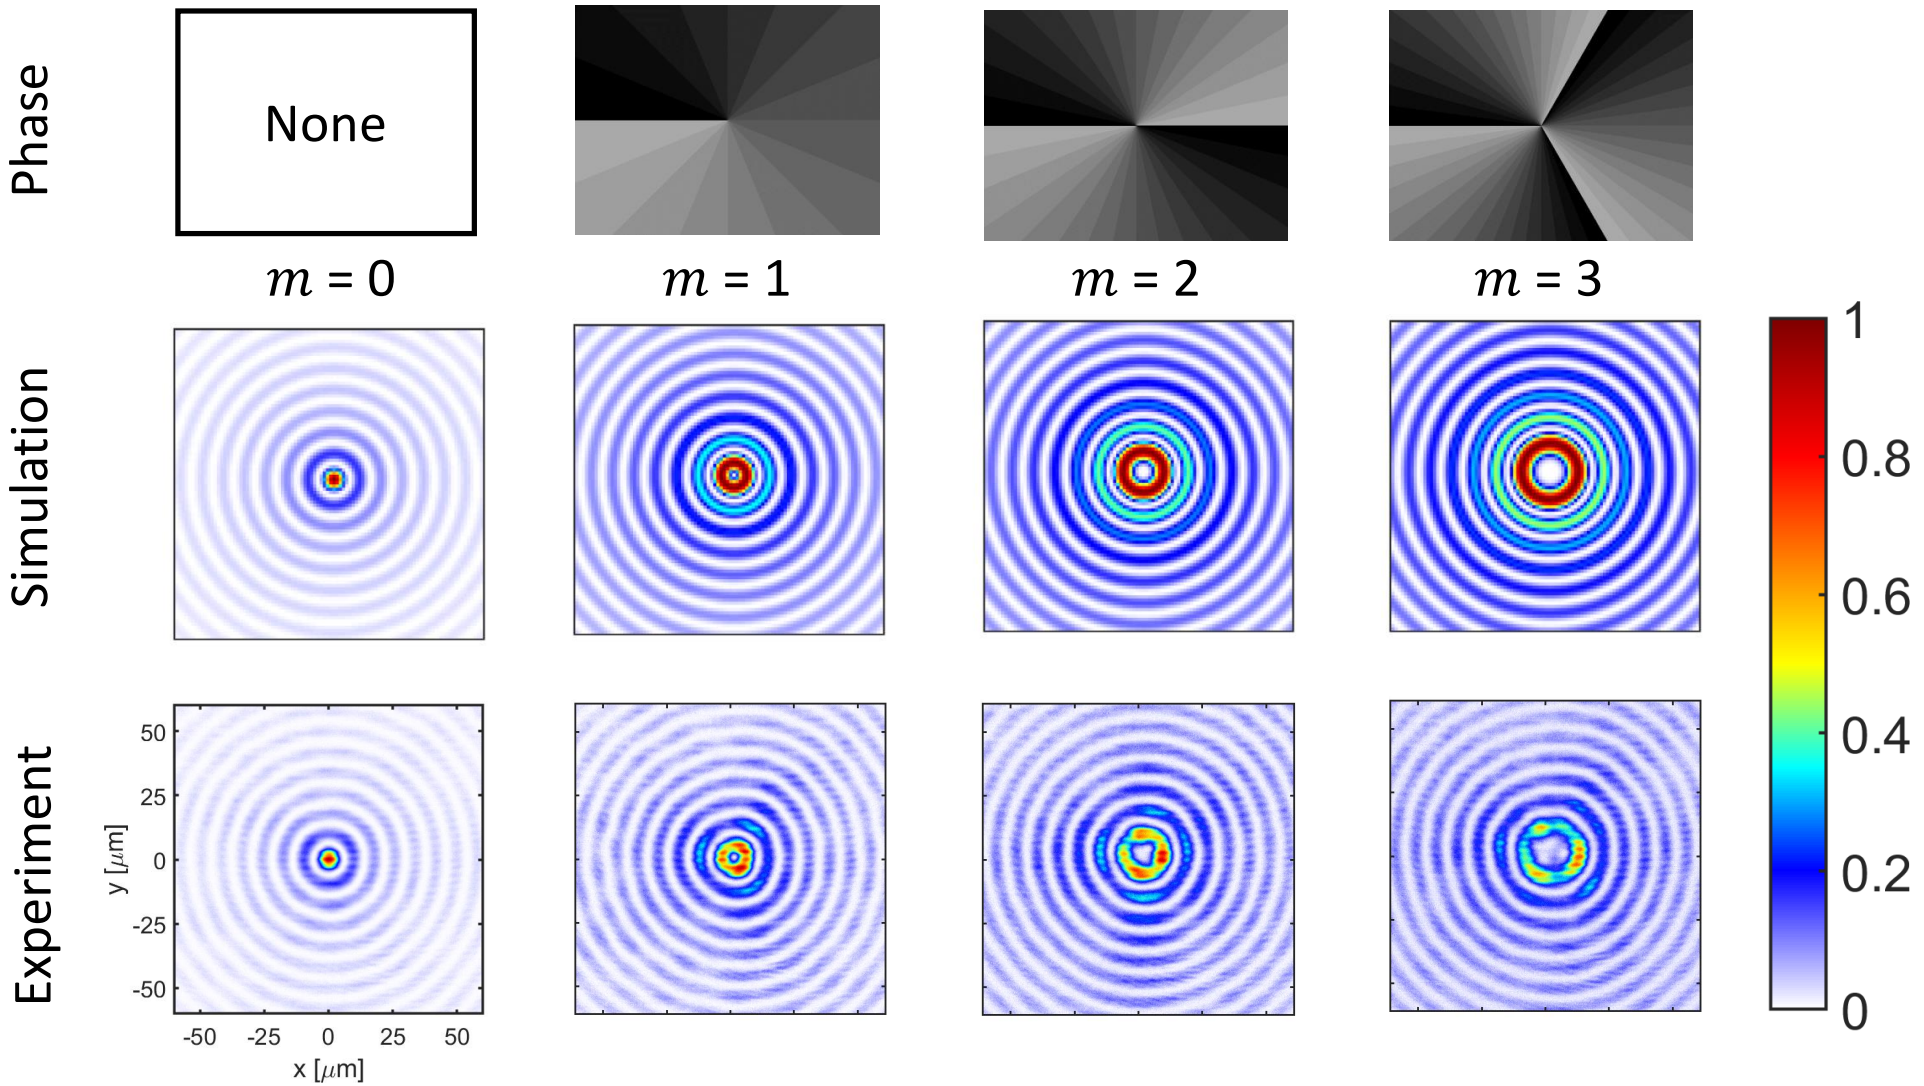


Figure 5.9: Focal spots generated at the focus of a linear axicon, base angle 5.0° after application of a spiral phase plate of mode m to the wavefront. Foci are imaged at $z = (40 \pm 2)$ mm with a beam of wavelength $\lambda = 0.633\,\mu$m. Top row: bitmaps of the spiral phase plates used. Middle row: results of numerical calculations using ASM code. Bottom row: focal spots imaged in SLM experiment. Images in the middle row are self-normalised, images in the bottom row are normalised to the peak of the experimental J_0^2 image. The x and y axes on the middle and bottom row images go from $-60\,\mu$m to $+60\,\mu$m.

phase instead of a continuous phase between 0 and 2π, so as to best match the spiral phase plates available commercially [137]. Topological charge values $m = 1$, 2 and 3 were chosen to generate higher order Bessel modes (HOBM) of the first kind, J_1, J_2 and J_3. The resulting focal spots are shown in Figure 5.9. The numerical calculations were run to model the same experimental setup. The resolution of the numerical simulations was 1.2 µm.

It can be seen that the HOBM focal spots produced experimentally display a small amount of azimuthal asymmetry. The azimuthal symmetry was optimised for each m-value by adjusting the position of the centre of the bitmap on the SLM with respect to the beam, as well as the transverse position and pointing of the axicon. The asymmetry observed in Figure 5.9 is due to the remaining imperfections in the beam that could not be removed with these methods. In a more sophisticated experiment, a wavefront sensor and adaptive optic could be used before the spiral phase optic to remove aberrations from the laser beam and further improve the HOBM focal spot quality.

To compare the results of the numerical calculation and the SLM experiment, the ring radius was determined for each of the topological charge values $m = 1$, 2 and 3, where ring radius here means half of the distance between the maxima on a line through the centre of the ring. For the experimental images, the ring radius was calculated from two perpendicular line-outs and averaged. The results are shown in Table 5.1. The values for the experiment and numerical calculation agree within the uncertainty.

m-values	Calculation [µm $\pm$ 0.6 µm]	Experiment [µm $\pm$ 0.5 µm]
1	4.5	4.5
2	7.1	7.3
3	10.8	10.8

Table 5.1: Ring radii of focal spots generated using the ASM numerical calculation and SLM experiment for 3 values of spiral charge m. The uncertainty on the numerical calculation is due to the resolution of the simulation.

Figure 5.10 shows images of the $m = 1$ HOBM at three different longitudinal positions, $z = 88$ mm, 108 mm and 128 mm respectively. For these images, a median filter has been applied to remove electrical noise and make the key features more distinct. The HOBM focus is found to be highly uniform over the 40 mm longitudinal position scan shown, as expected for a Bessel mode.

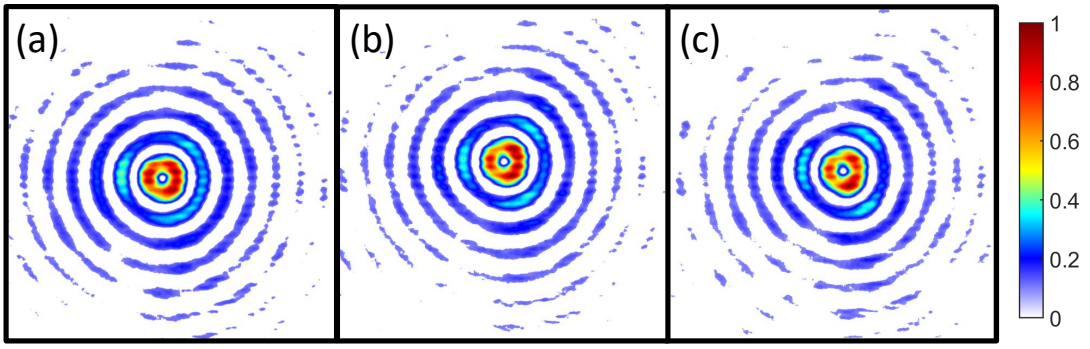


Figure 5.10: Measured focal spots generated at the focus of a linear axicon, base angle 5.0°, after application of a spiral phase plate on the SLM with topological charge $m = 1$, imaged at varying longitudinal positions from the axicon; (a) $z = 88$ mm; (b) $z = 108$ mm; (c) $z = 128$ mm (where z values are given to $\pm$ 2 mm). The pixel counts are self-normalised by dividing by the peak pixel count in each image. The x and y axes have a range of 110 µm.

5.3.4 Comparison with design objectives

The properties of the speckle beams and higher order Bessel modes can be compared to the design objectives described in Section 5.1.

Length of longitudinally-extended focus

The speckle patterns and HOBMs are both capable of producing extended foci of 100 mm and longer using an axicon. The HOBMs, being formed with the same axicon as the standard J_0^2 focus, can form foci over the same longitudinal extent as has been used in previous HOFI channel experiments, without changing the incident beam size.

The angle required for the speckle patterns is larger than for the standard J_0^2 focus. The larger approach angle is required in order to achieve a sufficient number of speckles across the focal region. Therefore, to form the same length of extended focus with the parameter regime found above, a larger radius beam would be required, by a factor roughly equal to the ratio of the angles. In the case described here, the axicon base angle for the speckle pattern was 20° compared to the 5.0° base angle used in the HOBM case, hence the radius of the incident beam required to form the same length of extended focus with the speckle pattern would be ~ 4 times greater than with the HOBM. The energy in the incident beam therefore has to be increased by a factor of $\sim 4^2$ to maintain the fluence over the larger required area.

Transverse size and uniformity of longitudinally-extended focus

The ring radius for each HOBM focus was constant with longitudinal position in the focus. Scaling the results of the SLM experiment to wavelengths of 800 nm, as opposed to the 633 nm used here, gives a ring radius of 5.7 µm for the case of J_1^2, 9.2 µm for the case of J_2^2 and 13.6 µm for J_3^2. If we consider the region with $I > I_{\text{peak}}/2$ as being ionised, the outer ionisation radius for a perfect HOBM would be 7.9 µm for the case of J_1^2 and 11.9 µm for the J_2^2, compared to the 3.2 µm for the J_0^2. Hence HOBMs can achieve the ionisation radius of $\sim$ 10 µm desired. However, HOBMs leave a region of unionised gas in the centre of the ring. The effects of the unionised gas are investigated using hydrodynamic simulations in Section 5.4.

The longitudinally increasing intensity from the linear axicon can cause the secondary wings of the Bessel focus to reach the intensity for ionisation. In Figure 5.2, it can be seen that higher m-values have larger secondary peaks relative to the dominant peak, hence secondary peaks will ionise the gas more easily at higher m. The tunnelling ionisation of electrons in the secondary peaks may give rise to free electrons outside the central core, as discussed in Chapter 3. These free electrons may conduct energy away from the hot core and reduce the pressure differential with the neutral gas. This might be expected to weaken the resulting blast wave, generating lower-amplitude, slower shocks. To offset the problems related to the increasing intensity with longitudinal position, the linear axicon could be exchanged for an axilens or the incident channel-forming beam could be tailored so that the intensity decreases with radius (see Chapter 4).

For the speckle patterns, an element size $d = 1$ mm was found to give a region size $\sigma_{\text{region}} \simeq 13$ µm at $z = 40$ mm. Hence in this configuration, the required ionisation radius can indeed be achieved. However, it was found that σ_{region} varies significantly with longitudinal position owing to two effects; firstly, the intensity in the elongated focus of an axicon increases linearly with z, leading to speckles further from the axis becoming intense enough to ionise the gas; secondly, the speckle region size has an approximately linear dependence on z due to considerations outlined in Section 5.2. The former effect could be mitigated using an axilens or tailored incident intensity profile. To address the latter effect, a more sophisticated arrangement of phase elements on the random phase plate would be needed. For example, instead of a square grid of phase elements, a polar

grid could be used with element size that increases with radius, so that the number of elements within each thin annulus on the transverse profile of the incident beam is constant.

The effective area ionised by the foci in each case is discussed in the section on energy efficiency below.

Practical realisation

Both the HOBM and speckle focus would be simple to implement, requiring a single phase plate to be added to the beam-line for channel formation. The focus of a HOBM was more sensitive to beam alignment than the speckle pattern due to the requirement that the beam be centred both on the spiral phase pattern and on the axicon. The same was not true for the RPP due to the self-similar nature of the phase distribution on the RPP as a function of x and y. In the future, a single phase plate which incorporates both axicon (or axilens) and spiral phase could be designed which would remove this difficulty.

Energy and energy efficiency

Finally, the energy requirements and efficiency may be considered. The energy per unit length coupled into the initial plasma column is given by;

$$\beta_{\mathsf{coupled}} = E_K n_{\mathsf{e}} A \tag{5.7}$$

where E_K is the mean kinetic energy of the plasma electrons after ionisation and A is a measure of the area of the region ionised by the incident channel-forming laser pulse. The values for A were calculated using experimental data generated with the SLM (see Section 5.3). Any regions where the normalised intensity was greater than $1/2$ were included in A. It should be noted that, due to the z-evolution of the intensity at focus with a linear axicon, the effective area ionised would differ with longitudinal position. For these calculations, $z = 40\,\mathsf{mm}$ and $\lambda = 633\,\mathsf{nm}$.

The other important figure of merit is the energy required in the channel-forming beam per unit length of the channel, which takes into account that not all of the focus is used to ionise the plasma. The energy per unit length is given by (up to a constant factor);

$$\beta_{\mathsf{incident}} = I_0\,\delta t\,\pi\rho_{\mathsf{out}}^2/L = I_0\,\delta t\,\pi\theta^2 L. \tag{5.8}$$

where ρ_{out} is the outer radius of the incident beam and δt is the pulse duration. To allow for a fair comparison, the incident intensity profile is taken to be a uniform tophat and the length of channel is taken to be $L = 200\,\text{mm}$. Here I_0 is an approximate measure of the intensity required to achieve $I_{th} = 4.5 \times 10^{14}\,\text{W cm}^{-2}$ at $z = 40\,\text{mm}$, taken from previous experimental parameters [1, 20] and is calculated using the results of ASM simulations for each case.

The ratio $\zeta = \beta_{\text{coupled}}/\beta_{\text{incident}}$ can be taken as a measure of the coupling efficiency of the channel-formation methods, i.e. the fraction of the energy in the incident laser beam that is coupled into electron energy in the ionised plasma core. The values for each case are shown in Table 5.2. Here ζ_n is the normalised efficiency compared to the standard J_0^2 method.

	β_{coupled}[mJ / m]	β_{incident}[J / m]	ζ_n
$m = 0$	0.09	0.5	1
$m = 1$	0.40	1.4	1.6
$m = 2$	0.50	2.2	1.3
speckle	0.01	9.3	0.002

Table 5.2: Calculated energy parameters for the HOBMs and speckle patterns compared to the standard $m = 0$ axicon focus without speckle. See text for details.

These calculations demonstrate that the higher order modes couple significantly more energy into the plasma compared to the speckle patterns. The low efficiency of the speckle patterns is due to the significant amount of light that is present at large radii outside of the region of interest, with intensity below threshold for ionisation (as demonstrated in Figure 5.6(b)). This energy is not useful for channel formation and hence leads to inefficiency. The value of β_{coupled} for the speckle patterns could be increased with spatial incoherence methods, with a number of beams with random phase being incoherently combined at focus such that overall the random distribution of speckles from one focus 'fill in the gaps' between the speckle distribution of the other beams. However each beam in the spatial incoherence setup would require the same initial intensity in order to ionise the gas at focus, leading to a corresponding increase in β_{incident}. Hence, the efficiency of the speckle patterns cannot be greatly improved.

The higher order modes are found to be more efficient than the single J_0^2 beam, with slightly lower efficiency for the J_2^2 compared to the J_1^2 case. The energy requirements

for the HOBMs are higher than for the J_0^2, but are modest compared to the multi-joule laser systems typically used for laser wakefield accelerators. Effects owing to B-integral would be more severe for beams with higher energy (see Chapter 4), therefore thin phase plates or reflective optics would likely be required for the axicon and spiral phase implementation in the HOBM channel-forming beam-line.

Conclusion

Considering the different design objectives, the HOBMs with $m = 1$ or $m = 2$ were decided upon as the most promising candidates to try in a HOFI plasma channel generation experiment. This is due to the much higher energy efficiencies of the HOBMs compared to the speckle patterns and the significantly greater uniformity of the focal properties of the HOBMs with azimuthal and longitudinal position.

5.4 Hydrodynamic simulations of higher-order Bessel modes

Having decided on a set of parameters to use for HOFI plasma channel generation, hydrodynamic simulations were performed to understand the evolution of the HOFI plasma channels following ionisation by a J_1^2 or J_2^2 focus compared to the standard J_0^2 focus. One concern with using a single higher-order Bessel mode is that the on-axis minimum of the initial ionisation profile may have a detrimental effect on the resulting plasma channel profiles. Any neutral gas remaining on-axis after the channel has formed will be ionised by the guided laser pulse and lead to a high on-axis plasma density. This will tend to refract the laser pulse out of the channel and hence disrupt the waveguide properties. Therefore, foci generated from an incoherent combination of J_1^2 or J_2^2 beams with a zeroth order Bessel mode were also considered, henceforth referred to as $(J_1^2+J_0^2)$ and $(J_2^2 + J_0^2)$. This was to provide information on any differences in channel evolution that arise when the on-axis gas is ionised (when the J_0^2 is present) compared to when it is not (for HOBMs alone).

5.4.1 The Helios 1D code

The hydrodynamic code used for the simulations was Helios 1D, a commercial code that is only accessible via the Science and Technology Facilities Council (STFC) servers [111].

Due to the sensitive nature of the calculations that Helios can be used to perform, such as for military applications, access to the code is restricted. Users submit their calculation requests as input decks to the STFC server and receive the results.

Helios calculates the time evolution of the plasma properties using conservation of mass, energy and momentum, along with an equation of state for the plasma [111]. Helios employs the Lagrangian specification of the flow, which means that the calculation follows individual fluid 'parcels' as they evolve over time. This is as opposed to the Eulerian specification, where the coordinate system is fixed and the fluid moves through the grid. For the Lagrangian specification, the grid is initialised into a discrete mesh of zones representing the fluid parcels, and these zones move and change size as the fluid moves and changes density. Due to the relatively long temporal and spatial scales involved with hydrodynamic motion (unlike for wakefield generation), the electrons and ions are assumed to move together as a single fluid in these simulation and hence the plasma remains neutral. This single-fluid approximation is valid when the Debye length $\lambda_D = \sqrt{\frac{\epsilon_0 k_B T}{ne^2}}$ is small compared to the scale length of the density variations, where $1/T = 1/T_\mathrm{e} + 1/T_i$ for T_e the electron temperature and T_i the ion temperature. Here $\lambda_D \approx 1\,\mathrm{nm}$ and the plasma variations are on the order of microns, so the single-fluid approximation holds.

Although the electron and ion species move together, Helios does allow for the temperature of the electrons and the ions to be set independently. However, Helios does not include OFI physics. Therefore, for the simulations in this chapter, the distribution of electron temperature following ionisation by the different Bessel foci was found using a numerical code written by the author (see section below about ionisation modelling). The electron-ion collision timescales are long ($\sim 0.5\,\mathrm{ns}$) compared to the timescales for the isotropisation and thermalisation of the electrons ($< 1\,\mathrm{ps}$) and hence the ions were initialised at room temperature everywhere (see Section 3.2.2 for details of the plasma timescales).

For the simulations presented in this chapter, the Helios code uses a local thermodynamic equilibrium (LTE) model, where the electron number density at each position is given as a simple function of the local temperature and background atomic density, with no consideration of the time-dependence of ionisation and recombination processes.

This function is based on the Saha equation [111], given by;

$$\frac{n_{\mathrm{e}} n_{\mathrm{i}}}{n_{\mathrm{H2}} - n_{\mathrm{e}}} = 2\,\frac{g_1}{g_0}\left(\frac{2\pi m_{\mathrm{e}} k_B T}{h^2}\right)^{\frac{3}{2}} \exp\left(-\frac{e_I}{k_B T}\right). \tag{5.9}$$

where $n_{\mathrm{i}} = n_{\mathrm{e}}$ is the plasma density (neutral plasma) and n_{H2} is the combined density of hydrogen atoms and ions, e_I is the ionisation energy of the atom and g_1 and g_0 are the degeneracy of states for the ionised atom and the ground state respectively. Here T is the temperature of the plasma following the thermalisation of the electrons and ions with each other (which occurs on the timescale of 0.5 ns).

The LTE model assumes that the ionisation and recombination timescales are short compared to the timescales of the plasma evolution so that the ionisation fraction of the plasma, $f_n = n_{\mathrm{e}}/n_{\mathrm{H2}}$, responds instantaneously to changes in temperature and background density. This approximation holds well in the high-density plasma regimes that Helios was designed for. In the case of HOFI plasma channels, where low plasma densities are used, this assumption needs to be investigated. At electron temperatures of $\sim$ 10 eV and plasma densities $\sim 3 \times 10^{18}\,\mathrm{cm}^{-3}$, the radiative recombination and collisional ionisation timescales are given by the electron-ion and electron-neutral atom collision timescale, which are both $\sim$ 0.5 ns (as calculated in Section 3.2).[5] The blast wave evolves over a few nanoseconds and hence, in this low-density regime, the LTE assumption is at the edge of its applicability. [6]

Previous studies have shown that the results of Helios simulations for the position of the HOFI channel blast wave, in the same conditions as considered here, match well with experimental results and Sedov-Taylor theory [19]. However, it is well known that, in regions of sharp density transitions, the LTE model is not valid because the timescales of the plasma variation at the density transition edge are short [163]. Indeed, the amplitude and sharpness of the shock front in HOFI plasma channels given by the code have been found to be overestimated compared with experimental results [19]. More realistic models of the plasma dynamics give blast waves which have reduced amplitude and smoother, wider profiles [150].

[5] It is also possible for recombination to occur through collisional recombination, requiring two electrons and one ion to collide, but this rate is slower at low densities than the radiative recombination rate as it is a three-body process.

[6] For analysis of applicability of the LTE assumption for times $t < 1$ ns, see the discussion regarding Figure 5.15.

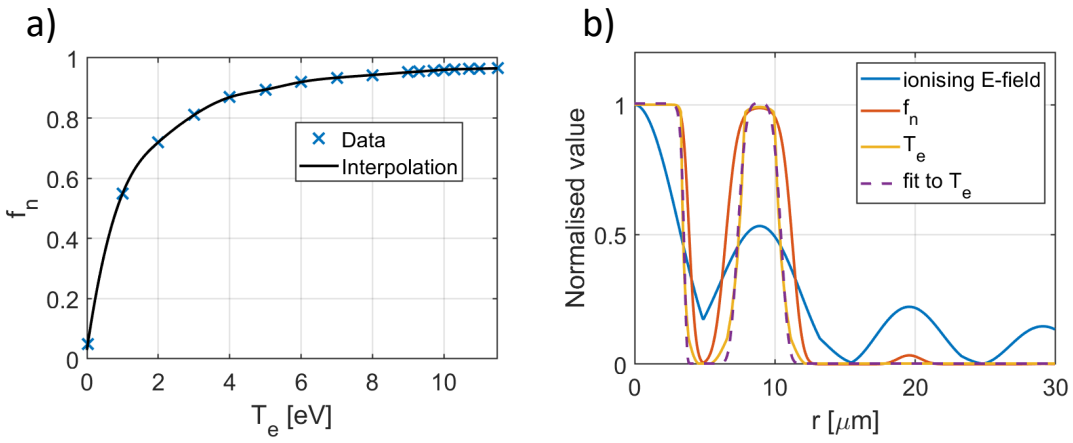


Figure 5.11: Modelling of ionisation of the gas by Bessel foci. (a) shows the calibration curve between electron temperature T_e and ionisation fraction f_n in Helios 1D, extracted from a series of simulations of hydrogen with constant density $n_{H2} = 3 \times 10^{18}\,\mathrm{cm}^{-3}$ and ionisation energy of 15.4 eV. The line through the data points is a 1D interpolation. (b) shows the profiles used in the ionisation modelling for an example case of a combination focus $J_0^2 + J_2^2$. The plotted profiles show the input electric field of the Bessel focus at the focus plane (blue), the ionisation fraction f_n calculated via ADK ionisation rates (red) and the corresponding temperature profiles T_e (yellow) required in Helios to generate the ionisation fraction profile. The fit to T_e (purple dashed) takes the form of a sum of super-Gaussian functions (see text).

More sophisticated non-LTE simulations of the HOFI plasma channels were attempted for the study presented in this chapter, using the Helios 1D non-LTE models [111]. However, it was found that non-LTE simulations of the HOFI channels were infeasible due to the significant computation time required. A non-LTE Helios simulation modelling 0.1 ns of blast wave evolution took 1 week to run on the STFC servers, hence simulations of HOFI channels to $t \gtrsim 5$ ns would be expected to have run-times on the order of a year. Research is ongoing into the development of more advanced simulation tools to better model the processes involved in HOFI plasma channel formation [164]. The Helios simulations presented in this chapter allow for the visualisation of the physics given by the Rankine-Hugoniot conservation relations (see Section 3.2) and are expected to provide a first approximation to the position and amplitude of the blast waves. Experiments and more advanced simulations will be required to fully understand the properties of the blast waves generated by higher-order Bessel modes.

5.4.2 Ionisation modelling

The hydrodynamic simulations were initialised to generate a pre-ionised plasma distribution. To determine the shape of the plasma distribution for each HOBM setup, a numerical code was written by the author, based on the ADK ionisation rates described

in Section 3.2. The code takes as an input the transverse and temporal electric field profile of the laser pulse at a focus plane and outputs the ionisation fraction f_n of a hydrogen plasma in this plane, after the passage of the laser pulse. The calculation proceeds via a series of time steps, where the plane of the plasma is fixed and the laser pulse moves through the plane. The ionisation rate associated with the transverse intensity profile at each time step is calculated using the ADK rate. The number of hydrogen atoms is updated at each time step by multiplying the number of available atoms by $1 - f_n$ to ensure that ions cannot be ionised a second time. The hydrogen gas was modelled as monatomic but with the ionisation energy of molecular hydrogen, i.e. 15.4 eV. In effect, this model assumes that if one electron is liberated from the hydrogen molecule then both are (see Section 3.2.2 for discussion of molecular hydrogen ionisation).

In an ideal model of OFI, the simulation would allow the electron temperature and ionisation fraction of the plasma to be set independently, to account for the fact that the electrons are not expected to recombine with their ions on the picosecond timescales of the isotropisation and thermalisation (as discussed in Section 3.2). However, under the LTE approximation used in the Helios simulations, f_n and T_e cannot be set independently. The pre-ionised plasma density profile must be generated in Helios via the inputs of the background plasma density and the electron and ion temperature profiles at the initial time step.

Therefore, the initial condition for the hydrodynamic simulations was set as follows.[7] A series of calibration Helios simulations were run with the same initial density of hydrogen as used for the main simulation runs below ($n_\mathrm{e} = 3 \times 10^{18}\,\mathrm{cm}^{-3}$), but with a range of different electron temperatures. The ionisation fraction of the plasma was extracted for each simulation run and plotted against electron temperature, to give the calibration curve $f_n(T_\mathrm{e})$ shown in Figure 5.11 (a) (closely related to the Saha equation, see Equation (5.9)). Interpolation of $f_n(T_\mathrm{e})$ enabled the desired ionisation fraction profile to be converted into an electron temperature profile. The maximum of the electron temperature profile corresponding to $f_n = 1$ could be set to any temperature $T_\mathrm{e} > 9\,\mathrm{eV}$ due to the degeneracy in the calibration curve in Figure 5.11 (a). Therefore, the

[7] The choice made in this section for the initial conditions of the hydrodynamic simulation ensures that the number density of electrons matches that calculated by OFI and that the kinetic energy per electron closely matches the model discussed in Section 3.2.2. Hence the total energy contained in the plasma core ($\sim n_\mathrm{e} T_\mathrm{e}$) at the start of the Helios simulation is in good alignment with OFI theory.

maximum temperature was set as $T_{\mathrm{e,max}} = 9.7\,\mathrm{eV}$, to match the temperature expected from theoretical considerations of OFI of hydrogen molecules (see Section 3.2). Note that the ion temperature throughout this discussion is initialised as room temperature, $T_i = 0.025\,\mathrm{eV}$.

The resulting transverse temperature profiles were fitted with either a single super-Gaussian or a sum of super-Gaussians depending on whether the single Bessel focus or a combination focus was used. For a combination focus, the fit function took the form $T_{\mathrm{e}}(r) = T_{\mathrm{e,max}}(\exp\{(-r/\sigma_1)^{g_1}\} + \exp\{(-(r-x_0)/\sigma_2)^{g_2}\})$, with the offset ($x_0$), widths ($\sigma_i$) and order of the super-Gaussians (g_i) being parameters of the fit. In the case of a single HOBM, the first of the two terms was dropped. The parameterisation $T_{\mathrm{e}}(r)$ was then used as the input for Helios. Figure 5.11 (b) demonstrates the use of the code for an example case with the $J_0^2 + J_2^2$ focus, showing the transverse electric field of the focus, the calculated ionisation fraction profile, the associated electron temperature profile and the double super-Gaussian fit $T_{\mathrm{e}}(r)$.

As demonstrated in Figure 5.2, HOBMs generate a lower peak intensity at focus than a J_0^2 beam with the same incident beam intensity and axicon. Therefore, the intensity in the HOBM incident beam was set in the simulations to be higher than that of the J_0^2. The intensity ratio between HOBM and J_0^2 was chosen to be the same for both the J_1^2 and J_2^2 cases, for ease of implementation in a future experiment. An intensity ratio of 2.25 (electric field ratio of 1.5) ensured that the central rings of the J_1^2 and J_2^2 foci were fully ionised, whilst the secondary ring of the J_1^2 was still below the ionisation threshold.

For the simulations below, the following parameters were used in the numerical ionisation code, chosen to match closely the values used in previous HOFI experiments [19]; peak intensity of the zeroth order Bessel foci $I_0 = 1 \times 10^{15}\,\mathrm{W\,cm^{-2}}$; peak intensity of the HOBM foci $I_{\mathrm{HOBM}} = 2.25 \times 10^{15}\,\mathrm{W\,cm^{-2}}$; laser wavelength 800 nm; axicon approach angle 2.5°; Gaussian laser temporal envelope with pulse duration 32 fs; time-step in ionisation calculation 0.05 fs; spatial resolution in ionisation calculation 50 nm. For the combination foci, the transverse electric field profile was chosen to model the case where the two foci were combined incoherently, i.e. at each radial position, the electric field amplitude used in the ionisation calculation was taken as either that of the J_0^2 or that of the HOBM, whichever was the higher. In practice, the two foci would be separated either in time or by polarisation and the resulting ionisation dynamics between the two

cases might be slightly different than assumed here. However, as the combination foci were being used mostly as a comparison for the on-axis density behaviour of the HOBM alone, the approximation outlined here was deemed sufficient for the present study.

5.4.3 Results of hydrodynamic simulations

The following simulations were performed using the Helios 1D fluid code, with parameters closely matching the regime used by the Oxford group to model HOFI plasma channels previously [19]. The grid was discretised into a mesh of 1000 fluid zones covering a 100 µm radial extent. The radial position of a zone in the subsequent results is defined as the centre of the zone at each time step and hence the plots of transverse profiles do not have a data point plotted at $r = 0$. The ionisation energy was set to that of molecular hydrogen, i.e. 15.4 eV. The initial electron ionisation distribution was calculated for different focus intensity profiles using the numerical ionisation code. The ion species was initialised at room temperature everywhere, $T_i = 0.025$ eV. The gas pressure was set at 62 mbar, corresponding to $n_\mathrm{e} = 3 \times 10^{18}\,\mathrm{cm}^{-3}$.

Trends in blast wave evolution

The electron density profiles, n_e, calculated at a series of time steps are plotted in Figure 5.12. The combined atomic and ion density profiles, denoted n_{H2}, at the same time steps are presented in Figure 5.13. In every case where a higher-order Bessel mode is used, the amplitude and speed of the blast wave is greater than in the reference case, resulting in a deeper plasma channel. When the HOBM alone is used, the blast wave propagates both towards and away from the laser axis ($r = 0$). This is particularly evident for the J_2^2 case in Figure 5.12(c), which shows two peaks in the electron density profile at early times. The collapsing blast wave rebounds from the axis and propagates outwards, eventually catching up with the initially outward-propagating blast wave at $t \approx 2.0$ ns. Again, this can be seen most clearly in Figure 5.12(c), where the plasma density profile at $t = 1.2$ ns has a bump just behind the main shock front, which grows in amplitude by $t = 1.6$ ns and then combines with the main shock front at $t = 2.0$ ns.

For ease of comparison, the properties of the shock front in each case are extracted and plotted in Figure 5.14. Figures 5.14(a) and (b) show the position r_s and amplitude n_s respectively of the peak in the transverse plasma density profiles as a function of

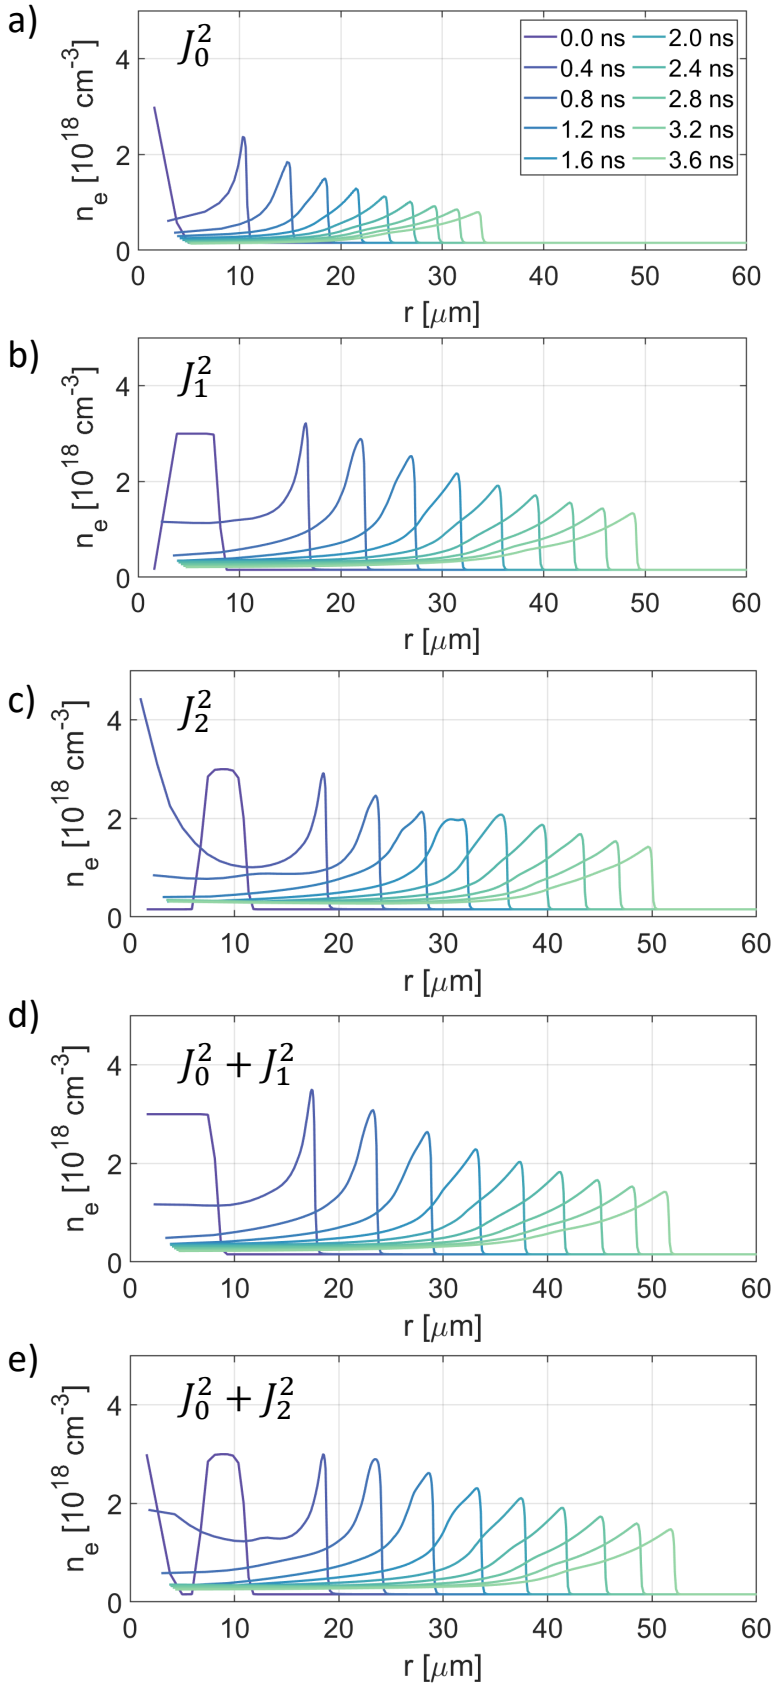


Figure 5.12: Calculated electron density profiles $n_e(r)$, at various delays t, after the formation of an initial plasma column generated using Bessel beam foci consisting of; (a) J_0^2 only, (b) J_1^2 only, (c) J_2^2 only, (d) $(J_0^2 + J_1^2)$, (e) $(J_0^2 + J_2^2)$. The simulations were performed using the Helios code for hydrogen at initial pressure 62 mbar and electron temperature $T_e = 9.7$ eV.

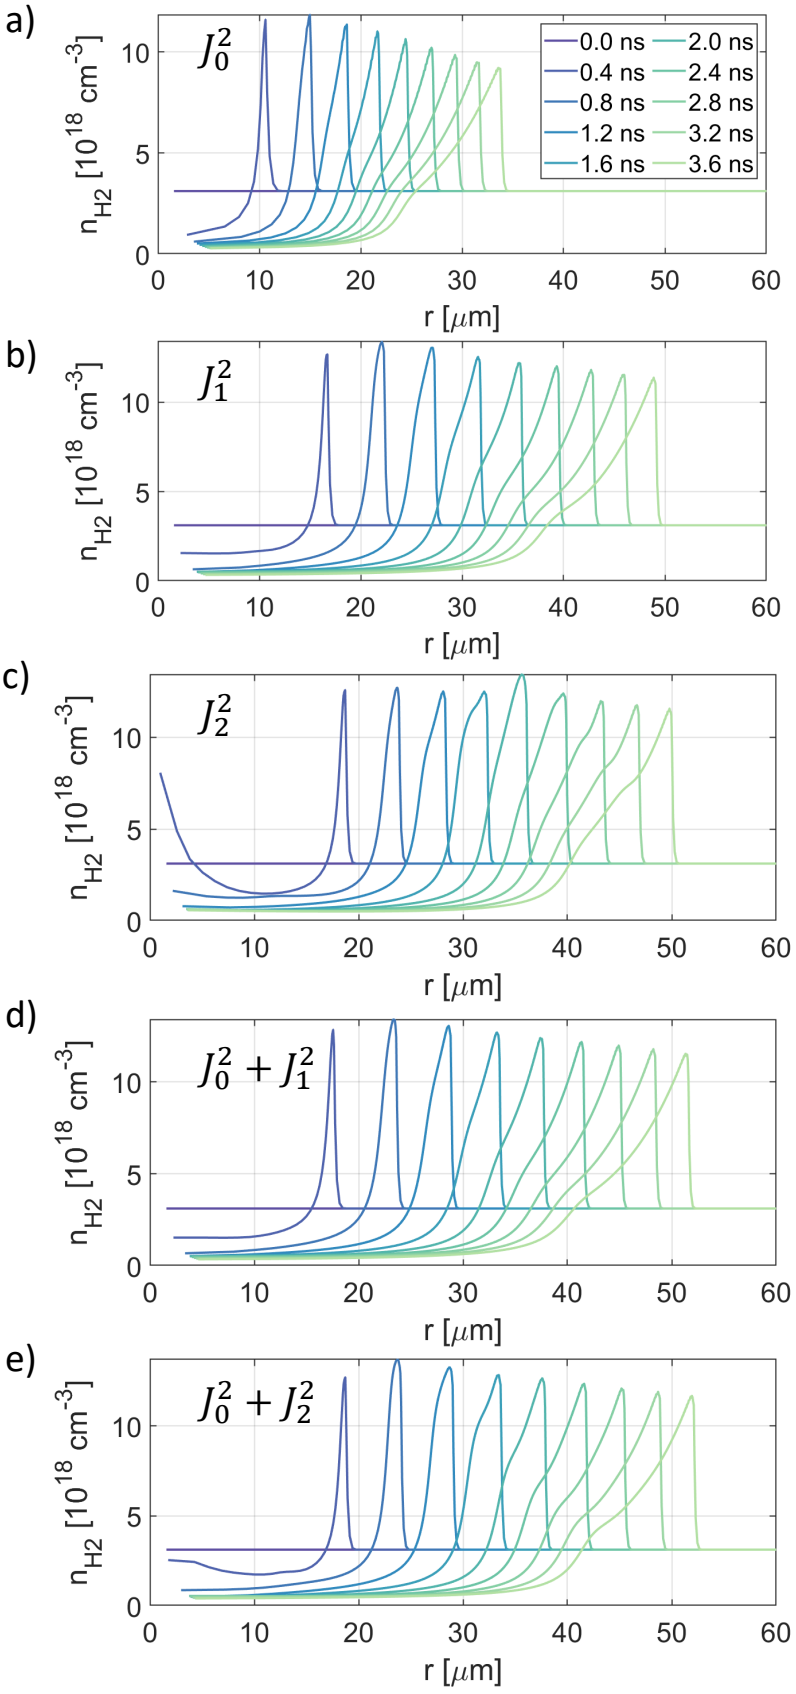


Figure 5.13: Calculated atomic and ion density profiles $n_{\mathrm{H2}}(r)$, at various delays t, after the formation of an initial plasma column generated using Bessel beam foci consisting of; (a) J_0^2 only, (b) J_1^2 only, (c) J_2^2 only, (d) $(J_0^2 + J_1^2)$, (e) $(J_0^2 + J_2^2)$. The simulations were performed using the Helios code for hydrogen at initial pressure 62 mbar and electron temperature $T_{\mathrm{e}} = 9.7\,\mathrm{eV}$.

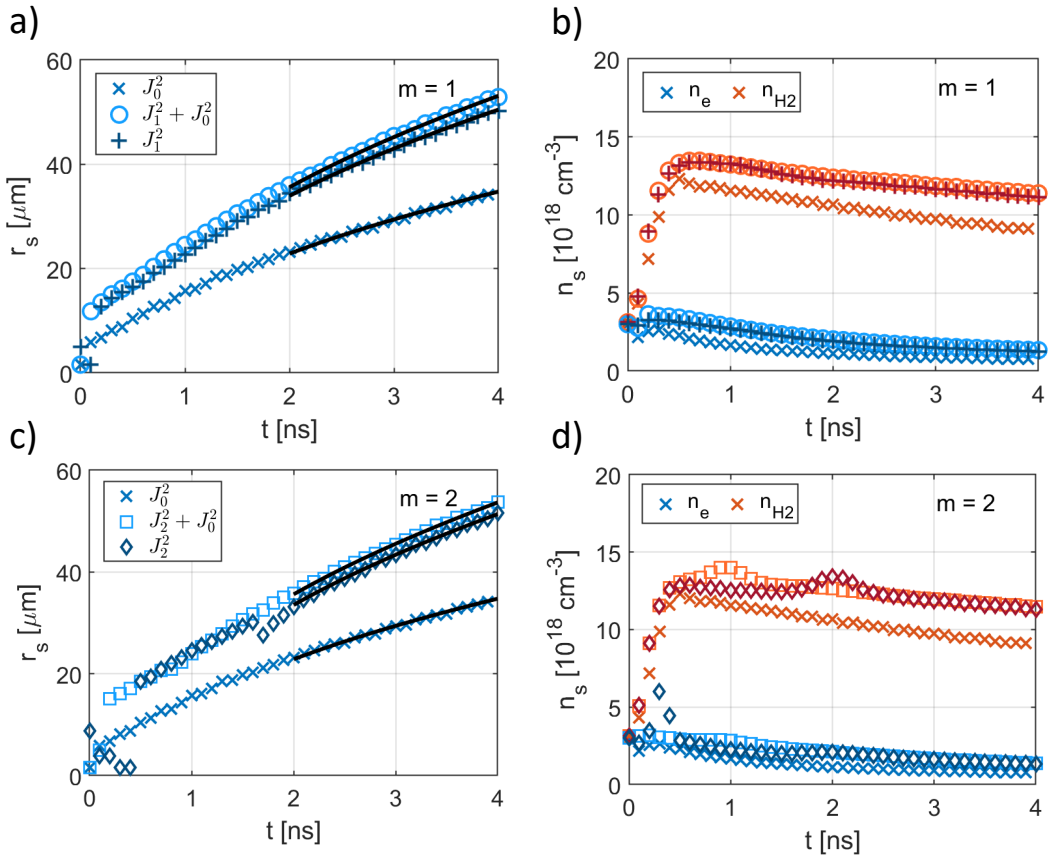


Figure 5.14: Calculated parameters of the blast waves shown in Figures 5.12 and 5.13 as a function of delay t, after the formation of an initial plasma column generated using Bessel beams with; (a) & (b) J_0^2 only, $(J_1^2 + J_0^2)$, and J_1^2 only; (c) & (d) J_0^2 only, $(J_2^2 + J_0^2)$, and J_2^2 only. (a) and (c) show the radial position r_s of the peak of the outward-propagating blast wave. (b) and (d) show the amplitude of the blast waves n_s in the electron density profile (n_e, blue) and ion density profile (n_{H2}, red). Note that r_s was the same for both the electron and ion density profiles. In (a) and (c), the radial position of the shock wave for $t > 2.0$ ns has been fit to Equation (3.22) for each case (see text), with the fit plotted as a black line.

delay time t, for ionisation by J_1^2 and $(J_1^2 + J_0^2)$ foci compared to J_0^2. Figures 5.14(c) and (d) show the blast wave properties for the J_2^2 and $(J_2^2 + J_0^2)$ Bessel beams. The difference in r_s given by the electron density profile compared to the hydrogen atom and ion density profile was negligible, hence only the peak position of the electron density profile is shown in (a) and (c). For the cases where a blast wave propagates both towards and away from the axis, the peak plotted here is that of the outward-propagating wave.

Figure 5.14(a) and (c) demonstrate that using a HOBM focus significantly increases the rate of expansion of the blast wave compared to the reference case with J_0^2. The radial position r_s evolves with an approximately square-root dependence on delay time, as expected from the Sedov-Taylor theory (see Equation (3.22)). The dip in the value of r_s at $t = 1.7$ ns in the J_2^2 case is a result of the initially inward-propagating wave

catching up with the outward-propagating wave and distorting the peak position.

The data for each $r_s(t)$ curve has been fit to Equation (3.22) to give an idea of the effective $r_{e,0}$ of the ionised region, i.e. the initial spot size that would have been required to generate the same blast wave expansion rate following ionisation by a standard J_0^2. In the single HOBM cases, the rate of expansion of the blast wave increases when the collapsing blast wave catches up with the initially outward-propagating wave (at $t \sim 2$ ns). Therefore, the fit was only performed over the domain $t > 2.0$ ns for each of the five data sets. The initial kinetic energy of the plasma electrons was set as 10 eV (see Section 3.2) and the mass and charge of the ions were set to model hydrogen plasma ($Z = 1$, $m_i = 1.66 \times 10^{-27}$ kg). The effective initial radius of the column $r_{e,0}$ and the time offset from the start of the evolution τ_0 were parameters of the fit. The fitted values are tabulated below.

	$r_{e,0}$ [μm ± 0.1 μm]	τ_0 [ns ± 0.02 ns]
J_0^2	7.4	-0.48
J_1^2	15.0	-0.35
$(J_1^2 + J_0^2)$	16.7	-0.37
J_2^2	16.3	-0.52
$(J_2^2 + J_0^2)$	17.3	-0.43

Table 5.3: Blast wave parameters calculated from fits to the curves $r_s(t)$ presented in Figure 5.14. See text for details.

The fitted values for $r_{e,0}$ in the HOBM cases are at least double that of the reference case. The J_2^2 beam produces a slightly stronger blast wave than the J_1^2, but the benefit is small compared to the improvement from using either HOBM over the J_0^2. The combined HOBM and J_0^2 foci give a slightly greater $r_{e,0}$ over using the HOBM alone in each case, as expected from the greater energy density coupled into the initially ionised region in the combined foci. The values for τ_0 are given for completeness and show that the blast wave evolution follows the Sedov-Taylor theory after approximately 0.3 – 0.5 ns.

The amplitudes of the blast waves are also enhanced by the HOBM foci, as shown in Figures 5.14(b) and (d). The maximum amplitude reached by the HOBMs in both n_e and n_{H2} is greater in each HOBM case compared to the reference. In the case of the J_2^2, the bump in blast wave amplitude at $t = 2.0$ ns corresponds to the moment the collapsing blast wave catches up with the initially outward-propagating wave. The temporary increase in blast wave amplitude does not have a significant effect on the

evolution after this time step, with the trends in both the radial position and amplitude of the shock wave continuing monotonically to later times.

To quantify the enhancement in amplitude when using a HOBM vs the reference case, the amplitudes of the blast waves were extracted at a delay time typically used between channel-forming and guided beams in HOFI experiments, $t = 3.0$ ns. The blast wave amplitudes in n_{e} and n_{H2} are tabulated below.

	$n_{s,\mathrm{e}}$ [10^{17}cm^{-3}]	$n_{s,\mathrm{H2}}$ [10^{18}cm^{-3}]
J_0^2	8.9	9.75
J_1^2	14.9	11.6
$(J_1^2 + J_0^2)$	15.9	11.9
J_2^2	16.1	11.9
$(J_2^2 + J_0^2)$	16.6	12.0

Table 5.4: Blast wave amplitude extracted from n_{e} and n_{H2} at $t = 3.0$ ns presented in Figure 5.14. See text for details.

In n_{e}, the amplitudes of the HOBM blast waves were almost double that of the J_0^2 blast wave. Adding the J_0^2 to the HOBM in each case only increased the amplitude of the blast wave marginally (by $< 1 \times 10^{17}$ cm^{-3}). Similarly, in n_{H2}, the difference in using the J_0^2 in combination with the HOBM compared to using the HOBM alone was negligible ($< 3 \times 10^{17}$ cm^{-3}). It is worth re-emphasising that the absolute values of the shock amplitudes output by the Helios code are not expected to be realised in an experiment, but the general trends are expected to be correct.

On-axis density

In the simulations described here, the initial annular ionisation region generated by the J_1^2 and J_2^2 Bessel beams leads to a blast wave that travels towards and away from the laser axis, with the inwards-propagating blast wave reaching $r = 0$ in a fraction of a nanosecond. The hot electrons from the blast wave cause the gas on-axis to be collisionally ionised and subsequently move away from the axis. The extent to which Helios in the LTE regime is able to model this behaviour is investigated further in this section.

The evolution of the ion and neutral atom density can be extracted from Helios to check for the presence of unionised matter on-axis. Figure 5.15 (a) and (b) show the electron number density and hydrogen ion and atom number density respectively for the simulations of the J_2^2 Bessel beam alone. The collapsing blast wave is found to generate

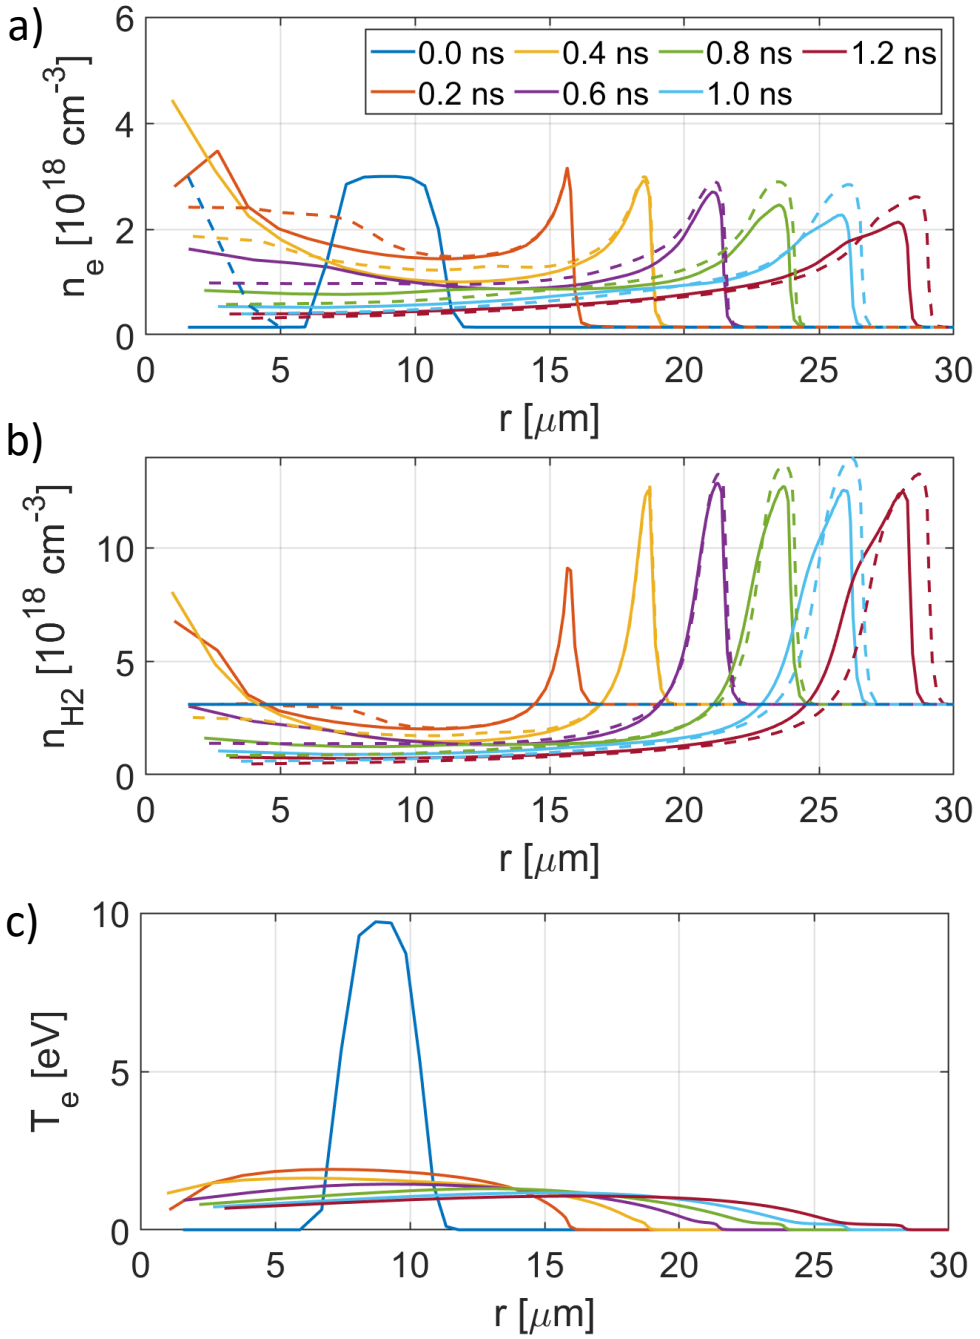


Figure 5.15: Evolution of blast wave at early time steps. (a) shows the calculated electron number density profile and (b) the calculated hydrogen ion and atom number density profile for the same Helios simulation parameters as in Figure 5.14. The solid lines show the case J_2^2 Bessel beam only and the dashed line shows the case $(J_2^2 + J_0^2)$ for comparison. (c) shows the electron temperature profiles for the J_2^2 Bessel beam only at the same time steps as (a) and (b).

a peak in number density at $r = 0\,\mu$m for time-step 0.4 ns. By $t = 1.2$ ns, the blast wave has rebounded from the axis, leaving a low on-axis density. The dashed lines show the density profiles at the same time step for the combined $(J_2^2 + J_0^2)$ Bessel focus, showing that by $t = 1.2$ ns, the profiles in both electron and hydrogen atom and ion density are nearly identical to those generated by the J_2^2 Bessel beam alone. Hence the addition of the J_0^2 Bessel beam to the HOBM is not necessary to achieve a plasma channel with

a low on-axis density. This is beneficial in terms of the experimental simplicity of the scheme.

To check the validity of the Helios LTE model for modelling the on-axis physics of the collapsing blast wave, we consider the collisional ionisation timescale for the plasma conditions at $t = 0.2\,\text{ns}$. Figure 5.15 (c) displays the electron temperature as a function of time, showing that the initial hot core of electrons cools to a temperature $T_\text{e} < 2\,\text{eV}$ by $t = 0.2\,\text{ns}$. The on-axis ion and atomic density at $t = 0.2\,\text{ns}$ is $n_{\text{H2}} \sim 6 \times 10^{18}\,\text{cm}^{-3}$. Therefore, using Equation (3.21), we find that $\tau_{\text{e,i}} \approx 0.02\,\text{ns}$. Therefore, the collisional ionisation process is expected to be reasonably well captured by the Helios code in this regime.

HOBMs and conditioned-HOFI

We now consider the results of the hydrodynamic simulations of HOBM foci in the context of conditioned HOFI plasma channels [21]. In the unconditioned HOFI plasma channel regime, the electron number density profile determines the guiding structure experienced by the guided beam. For a sufficiently high-intensity guided beam ($I \gtrsim 6 \times 10^{16}\,\text{W cm}^{-2}$ [21]) or with an additional conditioning pulse, the neutral hydrogen atoms also become ionised, liberating electrons with a density $n_\text{e}(r) = n_{\text{H2}}(r)$, and hence it is the $n_{\text{H2}}(r)$ profile (combined neutral atoms and ions) that determines the guiding structure (conditioned HOFI plasma channel or CHOFI regime). In the J_0^2 case, conditioning of HOFI plasma channels enables the depth of the plasma channel to be increased significantly compared to the unconditioned case [21]. However, the radius of the plasma channel is only increased marginally upon conditioning because the radius of the neutral gas collar is close to that of the HOFI blast wave. In comparison, by using a HOBM, the radius R of the plasma channel itself can be increased. Hence, given a desired channel width, the HOBM allows for a shorter delay time between channel-forming and guided beams to be used, thereby utilising the blast wave at an earlier point in its evolution, where the amplitude of the blast wave is greater and hence the attenuation length is expected to be longer.

Moreover, by using a HOBM focus in conjunction with the conditioning effect, HOBM-CHOFI plasma channels can be formed that are both significantly wider and deeper than the CHOFI channel generated by a J_0^2 beam alone. For example, using the

results in Figure 5.14(c), to achieve $r_s = 30\,\mu m$, the channel formed by the J_2^2 beam could be used at delay $t_d = 1.4$ ns giving a conditioned plasma channel depth of $n_{s,\mathrm{H2}} = 13 \times 10^{18}\,\mathrm{cm}^{-3}$, whereas for the J_0^2 beam, the delay required is 3.4 ns, giving $n_{s,H2} = 9 \times 10^{18}\,\mathrm{cm}^{-3}$. Therefore, in the simulations shown here, the J_2^2 HOBM-CHOFI plasma channel is predicted to give a 45% enhancement on the channel depth of the standard CHOFI plasma channel.

5.5 Conclusion

The theory of HOFI plasma channels has been studied with the aim of increasing the power attenuation length of the lowest-order mode guided by the plasma channels. It was found that increasing the initial spot size of the focus used to form the plasma channels would lead to wider waveguides that would confine the guided beams more effectively. The design parameters for the channel-forming beam focus were proposed. Two methods of generating extended foci with large diameter focal spots were discussed, namely random phase plates and higher-order Bessel modes.

A series of studies were then conducted to investigate the utility of these two methods. Experiments and simulations showed that both speckle patterns and HOBMs could indeed generate wider longitudinally-extended foci. A comparison of the focal spots with the design objectives revealed that the most promising candidates were the higher-order Bessel modes with $m = 1$ and $m = 2$. To investigate the expected differences between the possible higher-order Bessel mode setups, hydrodynamic simulations were performed to calculate the evolution of the HOFI plasma channel. The simulations demonstrated that using the HOBM compared to a standard J_0^2 Bessel focus significantly increased the final width and depth of the plasma channel. Combined J_0^2 and HOBM foci were also investigated and found to give negligible improvement in the channel properties compared to the HOBM alone. Implementing a single spiral phase plate is much simpler than combining two different Bessel orders spatially and temporally at focus. Hence, the J_1^2 or J_2^2 beam alone was proposed as the simplest, most energy-efficient method of achieving higher transmission plasma channels.

This work has informed the design of a future experimental investigation in which the temporal evolution and guiding properties of HOFI plasma channels produced by HOBMs

will be explored. The experiment will be used to generate HOFI plasma channels with and without the spiral phase applied, in order to determine whether the use of the HOBM does indeed increase the waveguide transmission efficiency as expected. If successful, the HOBM-enhanced HOFI or CHOFI plasma channel may be integrated into future plasma accelerator experiments.

Chapter 6

Observation of Resonant Excitation of Plasma Waves Driven in Long Plasma Channels

6.1 Introduction

As discussed in Chapter 1, multi-pulse laser wakefield accelerators (MP-LWFA) could enable GeV-scale electron energies to be reached using laser wakefield drivers from kilohertz-repetition-rate laser systems. Recent simulation work [4] performed by Oscar Jakobsson *et al.* demonstrated the potential of a new scheme for MP-LWFA, termed the plasma-modulated plasma accelerator (P-MoPA). In the P-MoPA scheme, a joule-level, picosecond-long laser pulse, such as that available from state-of-the-art thin disk lasers, co-propagates with a low-amplitude wakefield driven by a low-energy seed pulse in a plasma and develops spectral modulation (see Chapter 3.1.2). Following compression by a dispersive optic, a pulse train is formed which can then be used to resonantly drive a wakefield. In simulations, the wakefield driven by the pulse train accelerated a test bunch of electrons to 0.65 GeV in 100 mm [4].

Prior to the experiment described in this chapter, demonstrations of resonant wakefield excitation by trains of laser pulses had been limited to stage lengths of a few millimetres [23] and guiding of pulse trains in plasma waveguides had not been attempted. Short single laser pulses had been successfully guided and used to drive a wakefield over 100 mm-long propagation distances [1, 84], but laser hosing and self-modulation only become significant for pulses with duration greater than the plasma period as discussed in Chapter 2 [11] and hence these instabilities had not been probed by the short-pulse guiding experiment. Hence, it was not known whether laser-plasma instabilities would arise in the P-MoPA regime in practice and modify the laser pulse and/or the wakefield

properties. In this experiment, we aimed to demonstrate that trains of laser pulses could be guided in a plasma waveguide and that plasma wakefields could be resonantly excited by the pulse trains over many Rayleigh ranges. These are key steps in achieving the P-MoPA scheme, allowing the challenges relating to guiding of pulse trains to be understood and parameters of the accelerator to be optimised, ahead of future experiments to inject and accelerate electrons in the wakefield.

The experiment was performed at the Central Laser Facility (CLF) in the UK, using the 300 TW Astra-Gemini TA3 laser. The experimental design and planning were led by the author, with specific contributions from Nicolas Bourgeois and the team at the CLF on the experiment setup, Alexander Picksley on the interaction point and gas target design and Warren Wang and James Cowley on the pulse train generation and measurement. Simulations to define parameter ranges of interest and diagnostic requirements were performed by Oscar Jakobsson. The 6-week experimental campaign was led by the author, who additionally acted as Target Area Officer at the laser facility for the experiment, alongside James Cowley and James Chappell as deputy Target Area Officers. A team of experimentalists consisting of James Cowley, James Chappell, Linus Feder, Warren Wang, Oscar Jakobsson, Harry Jones and the author worked on the day-to-day running of the experiment.

Due to the short time periods awarded for experiments at the CLF (typically 6 weeks for the entire process of building the optical setups and taking data), experiment designs must necessarily be robust, reliable and quick to assemble. The layouts for the experiment described in this chapter were evolved from designs used for the Oxford LPA group's previous pulse train experiments [23] and plasma channel guiding experiments [1]. The experiment layout was made to be as compatible as possible with the preceding experiment on the Astra-Gemini TA3 laser, also conducted by the Oxford laser-plasma accelerator group and led by Alexander Picksley, on the acceleration of electrons in a plasma channel [165].

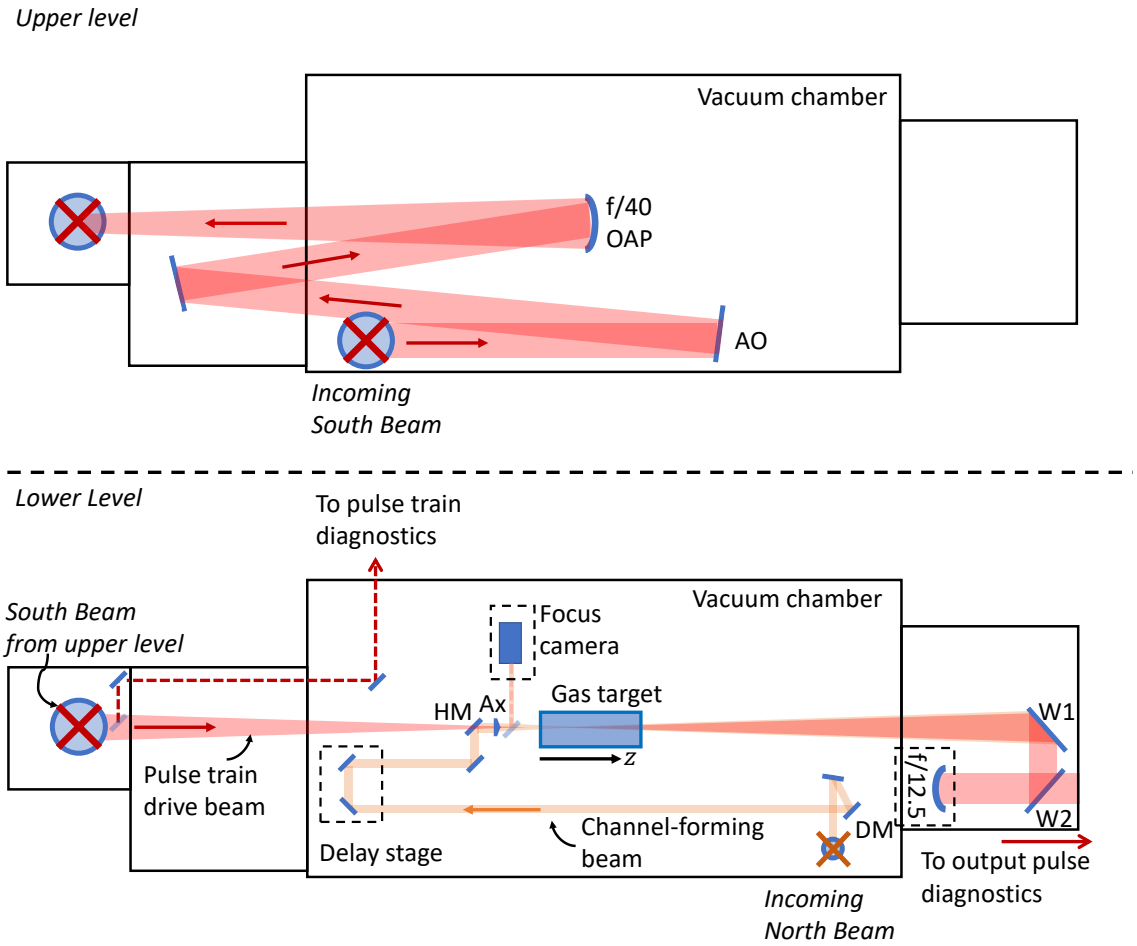


Figure 6.1: Experiment layout, looking from above. Note that the pulse train generation setup and grating compressor were located in the separate laser room, such that the South beam entered the experimental setup already modulated into a pulse train. The optical setup inside the vacuum chamber had two levels at different heights, the upper and lower level. On the upper level, the South beam entered from above and was directed towards the adaptive optic (AO) and then to the off-axis paraboloid (OAP) with $f = 6\,\mathrm{m}$ used at $f/40$ focusing. The focusing South beam was then sent to the lower level of the chamber via a downwards-going periscope. The North beam entered the vacuum chamber on the lower level and was directed via a deformable mirror (DM) which was used to increase the wavefront flatness of the beam to achieve optimised focal spots for channel formation. The beam was then sent to a fine-tunable delay stage which was used to set the time delay between the channel-forming and drive beams. The North beam was reflected from the front surface of the holed mirror (HM) and was focused by the UVFS axicon lens with base angle 3.6° (Ax) to a longitudinally-extended focus inside the gas target, where it formed the plasma channel. The focusing South beam propagated through the holes in the HM and Ax and came to focus at the start of the gas target. A pick-off mirror could be driven into the South beam line when the laser was in a lower power setting to direct the beam to the pulse train diagnostics. A mirror on a flip-mount could be inserted before the gas target to image the North and South beams on the focal camera in a low-power alignment mode. Laser pulses leaving the plasma channel were partially dumped by the two uncoated wedge substrates (W1 and W2) and the remaining light was collected by a spherical mirror used at $f/12.5$. The beam was then directed out of the vacuum chamber towards the forward diagnostics.

6.2 Experimental setup

6.2.1 Astra-Gemini TA3 laser

The experimental layout is shown in Figure 6.1. The Astra-Gemini TA3 laser is able to provide two synchronised beams (known as the 'North' and 'South' beams) which are generated via chirped pulse amplification of pulses derived from a common laser oscillator pulse. The laser gain medium is titanium-doped sapphire, and the amplified pulses have central wavelength $\lambda_0 = 800\,\text{nm}$ and FWHM bandwidth of 30 nm. Each beam is specified to reach pulse energies up to 15 J pre-compression (giving $\approx 9\,\text{J}$ post-compression) and pulse durations down to 40 fs post-compression. The laser beams have a diameter of 150 mm and can be operated at a maximum repetition-rate of one pulse every 20 s. In this chapter, the word 'shot' is used to refer to a single event following the firing of the laser system.

6.2.2 Channel-formation

The HOFI plasma channel was generated by the North beam (henceforth referred to as the channel-forming beam). The pulse duration was 80 fs, used in single short-pulse mode[1]. The 150 mm-diameter beam was apertured down to a beam size of 2 inch-diameter by a circular transmission mask placed immediately before the beam was directed into the target chamber. The channel-forming beam was reflected by a holed mirror (hole radius 15 mm) to enable coupling of the two beams into the gas target co-axially. The beam energy remaining at the interaction point was measured using an energy metre to be $\sim$100 mJ. The channel-forming beam was focused to a longitudinally-extended focus by an UVFS axicon lens (*Eksma optics*) with 3.6° base angle, resulting in 1.6° approach angle of the rays to the axis. The axicon had a hole in the centre of radius 13 mm, and was placed (560 ± 2) mm from the front plate of the gas target. The beam from the axicon focused through a 6 mm diameter hole in the front plate of the target. The axicon focus had a Bessel-function transverse profile (see Chapter 4) with a FWHM spot size of $(9.8 \pm 0.1)\,\mu\text{m}$.

[1] The pulse duration of the channel-forming beam was longer than the nominal pulse duration of the system (40 fs) because of additional spectral phase on the beam. This was introduced intentionally to modify the pulse trains using an acousto-optic programmable dispersive filter, but had the undesired side-effect of also impacting the channel-forming beam (see Section 6.2.3).

6.2.3 Generating a train of laser pulses

The South beam was used as the wakefield driver. In a future P-MoPA accelerator [4], the drive laser beam would be generated from a joule-scale, pico-second duration, kilohertz repetition-rate laser pulse. This pulse would be modulated using the P-MoPA technique to be resonant with the plasma density in the accelerator stage. However for this experiment, we did not have access to a joule-scale, pico-second pulse. Therefore instead, we formed the pulse train from the titanium-doped sapphire laser pulse as a mock-up of the pulse trains needed for the accelerator stage of P-MoPA.

The method of pulse train generation in the experiment in this chapter follows the same method as Cowley *et al.*, as described in Section 3.1.3 [23, 91]. In the South beamline, a Michelson interferometer was installed after the final amplifier in the laser chain, before the grating compressor. A flat (unwedged) compensator plate was used in the 'reflected' arm of the Michelson to account for the one extra pass of the beamsplitter glass made by the 'transmitted' arm, to ensure the beams in each arm experienced the same amount of group delay dispersion. The mirror on the transmitted arm was placed on a motorised stage for remote control of the Michelson arm optical path difference, δX_{mich}. This mirror also had piezoelectric motorised actuators for remote tip/tilt alignment. It was also possible to bypass the Michelson in the laser chain, resulting in a smooth, unmodulated pulse of the same FWHM pulse duration of 1 ps.

The compressor was set for partial compression, to leave a spectral chirp on the laser beam. Using the compressor in this way, with settings away from optimal compression, resulted in additional unwanted spectral phase terms being present in the partially-compressed pulse (in addition to the desired GDD). These third and fourth order delay (TOD and FOD) spectral terms would have affected the uniformity of the pulse spacing in the pulse trains, as discussed in Section 3.1.3. An acousto-optic programmable dispersive filter (AOPDF or 'Dazzler') was used to tune the TOD and FOD terms, using the method

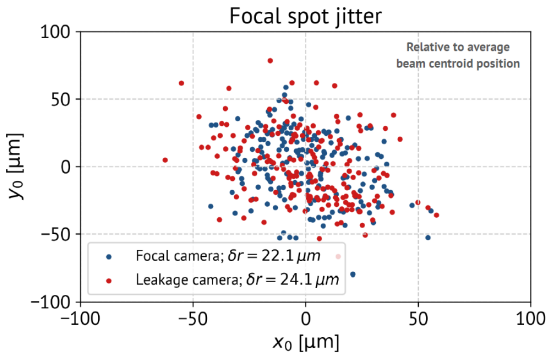


Figure 6.2: Spatial jitter of the drive beam focus measured at 0.05 Hz, plotted by James Chappell. The data taken using the focal camera (in blue) are used to calculate the spatial jitter in the text. The leakage camera data is not used in this chapter.

outlined in Section 3.1.4.[2]

6.2.4 Drive beam focus properties

The drive beam had (2.5 ± 0.5) J on-target energy and was focused into the pre-formed plasma channel via an off-axis parabola with 6 m focal length, used at $f/40$. The time delay between the arrivals of the channel-forming beam and the drive beam at the gas target position was set as $t_d = 3.5$ ns. The spot size of the drive beam at focus (equivalent to the $1/e^2$ radius of the beam intensity profile, for a Gaussian beam) was measured using the $D4\sigma$ method to be $(36.8 \pm 3.4) \times (42.1 \pm 4.5)$ μm^2, along the minor and major axes of the ellipse respectively (see Section 3.4 for the spot size definition).

The vacuum Rayleigh range was found by scanning the longitudinal position, z, of the focus camera and taking images of the drive beam on either side of the focus. The $D4\sigma$ spot size $w(z)$ was measured at each z position as the average of the beam diameters along the major and minor axes respectively, and plotted against z. A fit to the spot size evolution of a Gaussian beam, given by $w(z) = w_0\sqrt{1 + z/z_R^2}$, allowed the Rayleigh range to be measured as $z_R = (6.4 \pm 0.8)$ mm (analysis performed by James Chappell).

The RMS spatial jitter of the drive beam focus was measured to be 25.1 μm × 37.2 μm when the laser was operated at a repetition rate of 0.05 Hz (analysis performed

[2]It is worth noting that the AOPDF also affected the channel-forming beam, due to the location of the AOPDF early in the laser chain. The TOD and FOD phase components acted to increase the pulse duration of the channel-forming beam beyond the shortest possible duration. The temporal requirements on the channel-forming beam were simply that the threshold intensity for ionisation ($\simeq 4 \times 10^{14}$ W cm^{-2} [83]) was achieved over the length of the line focus. Hence, the spectral phase applied had to be kept at a low enough level so as to not increase the pulse duration of the channel-forming beam beyond $\simeq$ 200 fs. The pulse duration of the channel-forming beam was checked by the laser staff at the CLF after TOD and FOD settings were applied to ensure that this was the case. The TOD and FOD phase components themselves were not expected to have a significant effect on the HOFI plasma channel generation or evolution dynamics.

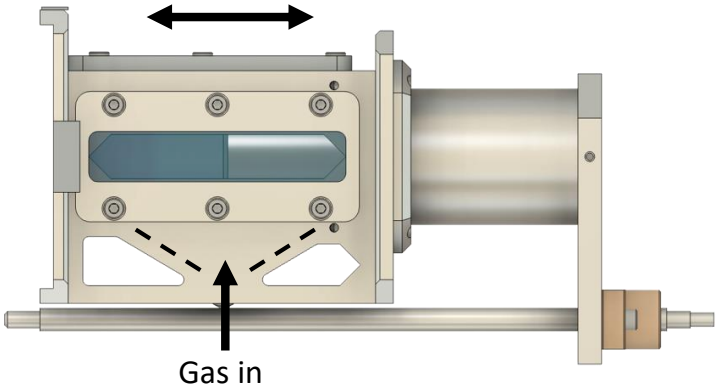


Figure 6.3: 3D rendering of the gas target used on the experiment, designed by Alexander Picksley (Oxford group) and Barnaby Matthews (CLF). The plunger on the right of the image could be moved to vary the length of the interaction. Gas entered the target volume from an inlet at the bottom of the chamber, via a conical volume marked with dotted lines.

by James Chappell). The data used for the measurement of the jitter is shown in Figure 6.2, generated by James Chappell. The jitter in the channel-forming beam focus position was much smaller than that of the drive beam, measured to be 5.2 µm × 2.4 µm, which can be explained by the small effective f-number of the axicon focus ($f_{\#} = 1/(2\theta) = 11.5$) compared to the drive beam focus ($f_{\#} = 40$). As shown in Figure 6.2, there was a significant fraction of shots for which the drive beam jitter was comparable to the drive beam focal spot size ($\approx$ 40 µm) and hence these shots would not be efficiently coupled into the plasma channel. Despite best efforts, jitter correction on the drive beam provided by the CLF was not effective at mitigating this issue and so was not used. Long-term ($\simeq$ hour) drifts of focal position were also observed in the experiment which needed to be compensated by regular checks and realignment of the beams during data taking.

6.2.5 Gas target

The gas target used in this experiment is shown in Figure 6.3. The target consisted of a cylindrical chamber of length 110 mm, closed on each end by a plate with a 6 mm diameter circular hole in the centre, to enable the laser pulses to enter and leave. The length of the gas target could be varied using a motorised plunger attached to the back plate of the target which could be driven into the cylindrical chamber. The back plate was mounted on the plunger. The gas inlet was designed in a cone-shape to disperse the gas uniformly in the target. For more details about the gas target design, please see

Alexander Picksley's thesis [165].

Hydrogen gas was let into the target via a solenoid valve (*Parker Series 9*) which was opened 150 ms before the arrival of the laser pulses, and remained open for a total duration of 200 ms to ensure that the laser-plasma interaction occurred in a steady-state gas condition. The in-target gas pressure was recorded using two pressure transducers (*RS PRO*) attached via gas tubing to the top of the gas target, one placed at roughly the centre of the target ($z = 59$ mm from the front plate) and one near the front of the target ($z = 30$ mm from the front plate). The gas tubing had 6 mm inner diameter and was a few centimetres in length. The transducers gave on-shot voltage readings that could be converted to pressure values via a calibration curve, measured in the Oxford lab prior to the experiment. Measured pressures in the target were in the range 10 – 150 mbar and had an uncertainty of 3 mbar. The longitudinal gas profile inside the target is discussed in Section 6.6.1.

For this experiment, it was important to know the on-axis plasma density of the HOFI plasma channel seen by the drive laser pulse. Previous experiments have demonstrated that the on-axis density of a HOFI plasma channel is roughly linearly proportional to the initial background gas density, with the ratio of background gas density to on-axis plasma density being somewhere between 5 and 10 [1, 112]. In previous experiments, the on-axis density has been retrieved for each laser shot using interferometry of a transverse-propagating probe beam [1]. The phase accumulated by the probe beam is dependent on the refractive index profile of the channel and so, by making assumptions about the likely shape of the plasma channel, the on-axis plasma density can be extracted. The experiment in this chapter was designed with a transverse probe included. The probe was to be developed by the laser facility for the experiment but, unfortunately, the commissioning of the probe beam was delayed. When the probe was provided, two weeks were spent attempting to identify the interferometric signal of the plasma channel on the probe, but these attempts were unsuccessful. Due to the limited time available for the experiment, the reasons for the lack of signal could not be identified and solved during the experimental campaign and hence it was not possible to collect the interferometric data. Therefore the on-axis density of the plasma channel had to be calculated from other measurements. In the results shown below, features in the

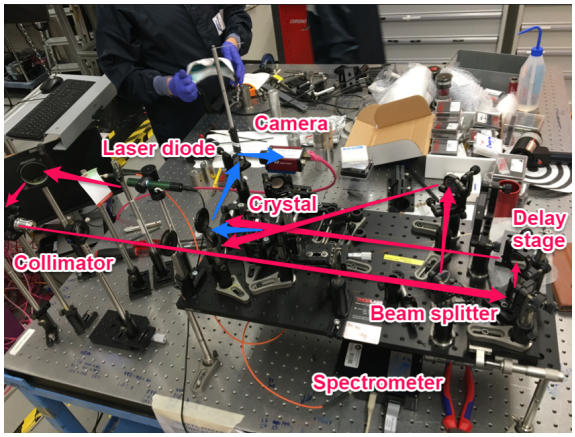


Figure 6.4: Layout of the SSA breadboard used on the experiment. The laser diode and collimator shown were used for alignment purposes. When used to measure pulse trains during the experiment, the drive beam was directed to enter the raised breadboard via two mirrors (bypassing the collimator). Photo and labels courtesy of Warren Wang.

spectrum of the transmitted pulse trains were used to extract the proportionality factor between pressure and on-axis plasma density of the plasma channels (see Section 6.5.3).

6.3 Diagnostics

6.3.1 Pulse train measurement

Warren Wang led the design of the diagnostic for pulse train measurement and the development of novel retrieval algorithms to extract the pulse train temporal intensity profile and phase. The pulse train spacing was measured using a single shot autocorrelator (SSA) which, in conjunction with information about the settings of the grating compressor and measurements of the input spectrum, could be used to retrieve the temporal profile of the pulse. See Chapter 3.1.3 for details of the retrieval method.

The measurements of the pulse train were conducted as follows. With the laser in a pulsed alignment mode, a motorised stage was driven into the path of the drive beam, after the final optic before the IP. A lens with focal length $f = -1000\,\mathrm{mm}$ was used to re-collimate the beam after the parabola. The beam was then sent out of the vacuum chamber via a thin optical window, and onto the SSA breadboard, which is shown in Figure 6.4.

The SSA breadboard consisted of a 50/50 beamsplitter (*KnightOptical*, 4 mm thickness), where one arm was directed via a delay stage and the other via a mirror such that the two beams crossed at an angle 4° to each other. At the crossing point, a nonlinear crystal (BBO Type 1, 38.8° cut angle from *InradOptics*) was placed for second

harmonic generation, with the generated 400 nm light leaving the crystal perpendicular to the plane of the crystal. The unconverted 800 nm light was dumped. The blue light was focused onto a camera (*Allied Vision Mako*) to image the plane of the crystal. To increase the efficiency of blue light generation, a pair of cylindrical lenses (*Thorlabs* UVFS, f =30 mm) were placed in the path of the two crossing beams, one before and one after the crystal, to focus the 800 nm beams in the plane perpendicular to their crossing angle.

The temporal profiles of the pulse trains retrieved from the autocorrelation traces are presented in Figure 6.5. The retrievals of the pulse train temporal profiles were performed by Warren Wang. In Figure 6.5(a), the pulse train consists of $N = 10$ pulses, where N is defined as the number of peaks within the $1/e^2$ temporal envelope of the intensity. The pulses are spaced by an average spacing of 170 fs. The second pulse train, shown in (b), has $N = 8$ pulses, spaced by an average of 203 fs. The pulse train spacing was switched from 170 fs to 200 fs by changing the spacing of the Michelson arms, as discussed in Section 3.1.3.

We can extract a measure of uniformity for these pulse trains, U defined as $U \equiv 1-(\tau_{\text{max}}-\tau_{\text{min}})/2\tau_{avg}$ where τ_{max} is the maximum spacing, τ_{min} is the minimum spacing and τ_{avg} is the average spacing between the pulses. A perfectly uniform train would have $U = 1$. The profiles shown in Figure 6.5 are examples of the temporal profiles that would fit the autocorrelation data, however other profiles with a different degree of uniformity would also fit the data. Therefore, unfortunately, the retrieval algorithm was not sensitive to the third-order (TOD) spectral phase on the pulse train spectrum, and hence the uniformity could not be calculated directly. However, by scanning the TOD with the Dazzler and maximising the average visibility of the peaks as discussed in Section 3.1.3, an upper bound on the TOD was set as $\pm$ 1 $\times$ 10^4 fs^3. A TOD value of 1 $\times$ 10^4 fs^3 corresponds to a uniformity measure of $U \sim 90\%$ on each train.

The uncertainty on the uniformity of the pulse train is expected from theoretical considerations to impact the degree to which the pulse train can resonantly drive the plasma wave. However, as the number of pulses in the train used on this experiment was relatively low, on the order of 10, the width of the resonance curve with plasma density was expected to be wide enough to account for variations in uniformity on the scale of 10% (see Section 2.2.3). Hence the uniformity of the pulse trains achieved by

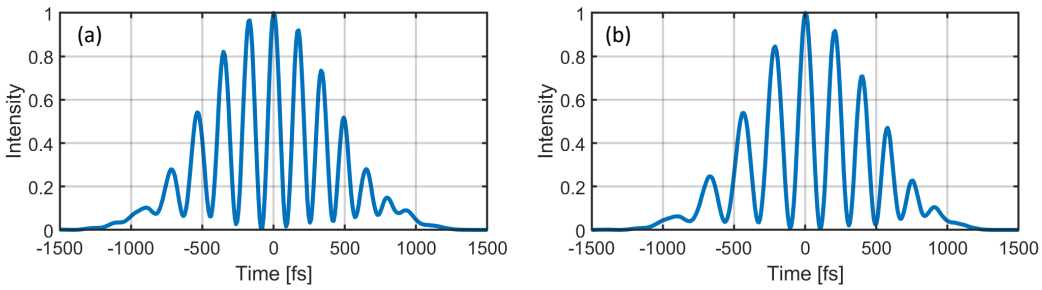


Figure 6.5: Example temporal intensity profiles of the pulse trains retrieved from the experimental autocorrelation data; (a) for the case of 170 fs pulse spacing; (b) for 205 fs pulse spacing. Retrieval performed by Warren Wang.

the above method was deemed to be sufficient for the experimental demonstration of resonant excitation of plasma wakefields.

As discussed in 2.2, resonance is achieved when the spacing of the pulses in the train (τ) is approximately equal to the time period of the plasma wave, that is;

$$\tau = \frac{2\pi}{\omega_p} = 2\pi \times \sqrt{\frac{m_e \epsilon_0}{n_{e,0} e^2}} \tag{6.1}$$

where $n_{e,0}$ is the on-axis plasma density. The resonance condition for the two pulse trains used in the experiment, with pulse spacings $\tau = 170\,\text{fs}$ and $205\,\text{fs}$, is met for on-axis plasma densities calculated as $n_{e,0} = 4.3 \times 10^{17}\,\text{cm}^{-3}$ and $2.9 \times 10^{17}\,\text{cm}^{-3}$ respectively.

6.3.2 Forward diagnostics

Following the interaction, the drive beam was directed out of the vacuum chamber and to the forward diagnostics, which were set up to image the exit plane of the plasma channel. The layout of the optics to the forward diagnostics is shown in Figure 6.6. The imaging system consisted of a spherical mirror with $f = 3.175\,\text{m}$ and an on-axis parabola with $f = 2\,\text{m}$. The object plane of the imaging system was set using a delay stage on the spherical mirror. The laser beams from the on-axis parabola were focused onto a camera (*Andor NEO 5.5 sCMOS*) and a fibre spectrometer (*Ocean Optics Maya 2000Pro*). The camera (henceforth referred to as the exit mode diagnostic) was used with a lens assembly (*OLYMPUS Plan N 4X*). The power incident on the camera chip was reduced using ND filters, which could be changed via a motorised filter wheel placed directly in front of the camera. Intensity attenuation of the fibre spectrometer was achieved by placing pieces of lens tissue in front of the collecting fibre.

It was not possible to prevent light from the channel-forming beam reaching the forward diagnostics, as the channel-forming and drive beams propagated co-axially and had

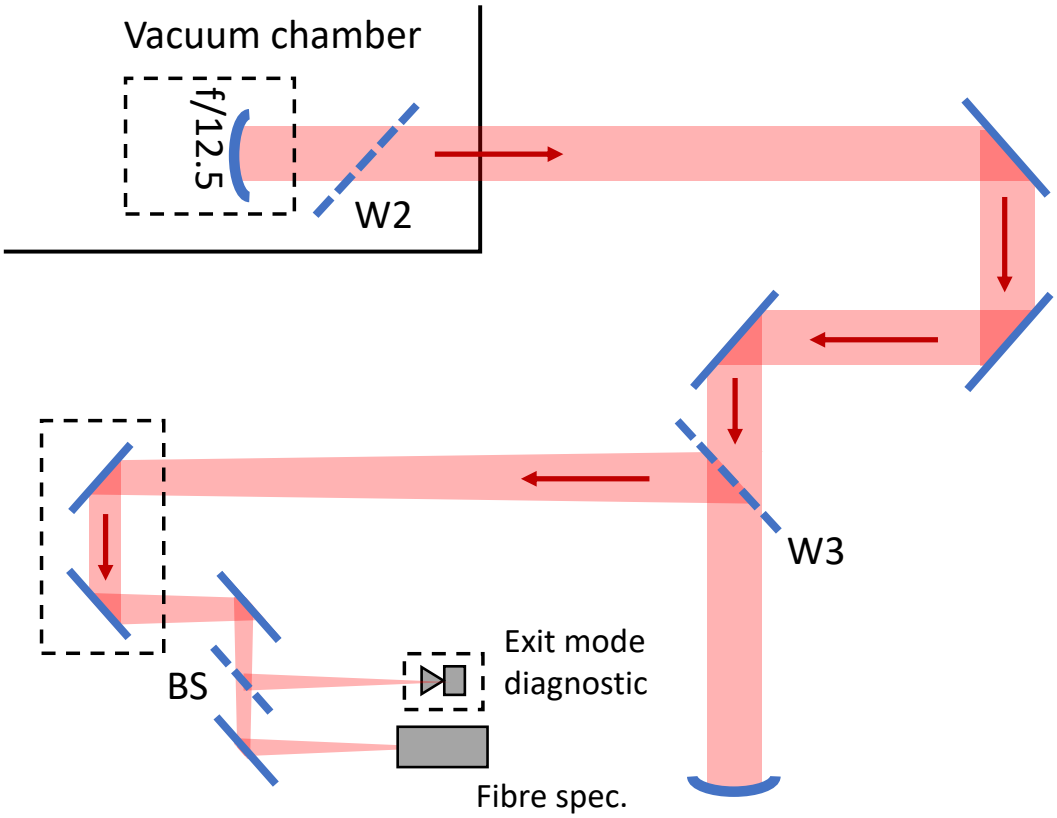


Figure 6.6: Simplified diagram of the layout to the forward diagnostics, starting from the $f/12.5$ gold-coated spherical mirror inside the chamber as shown in Figure 6.1. The beam passed through the wedge W2 and out of the vacuum chamber. The beam path after the chamber consisted of reflection from three silver-coated mirrors, transmission through a third uncoated wedge (W3), and focusing by a silver-coated on-axis parabola with $f = 2\,\mathrm{m}$. The focusing beam was reflected from the third uncoated wedge and directed to a 50:50 beamsplitter BS (*Thorlabs BSW30*) which was used to split the beam, with the reflection directed towards a camera used as the exit mode diagnostic (*Andor NEO 5.5 sCMOS*) and the transmission directed towards a fibre spectrometer (*Ocean Optics Maya 2000Pro*).

the same wavelength. We attempted to polarise the beams perpendicularly and prevent the light from channel-forming beam reaching the forward diagnostics using a polarising beam-splitter. However, these attempts were ultimately unsuccessful, perhaps due to small polarisation rotations of the beams during propagation that left their polarisations mixed. The energy in the channel-forming beam at the end of the channel was much less than that of the drive beam and hence gave a low-level background noise to the forward diagnostic data.

The exit mode diagnostic was spatially calibrated using a USAF target placed at focus. The resolution of the imaging system was measured to be $12.5 \times 20\,\mu\mathrm{m}^2$. The difference in resolution in the vertical and horizontal directions can be explained by the

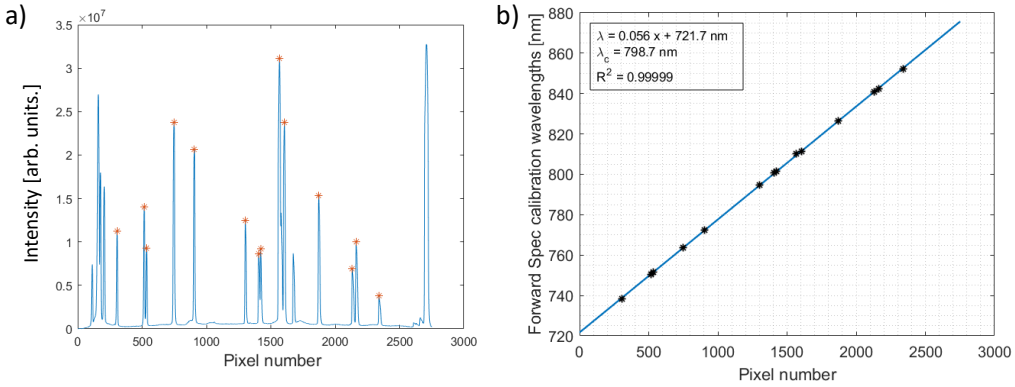


Figure 6.7: Example wavelength calibration procedure for a spectrometer using a Mercury (Argon) lamp.

fact that the limiting aperture in the optical system (W1) is placed at 45° to the beam direction and hence the horizontal axis of the beam sees a smaller aperture diameter than the vertical axis, leading to a larger 'f-number' (f/D) and hence worse resolution (see Equation (3.34)). The image consisted of 2560 × 2160 pixels, with each pixel covering $4.52 \times 4.52\,\mu m^2$ in the object plane.

The fibre spectrometer wavelength axis was calibrated using a light source with a known emission spectrum consisting of well-defined emission lines, here a Mercury (Argon) lamp. The peaks in the recorded spectrum were identified with the peaks in the known spectrum, thereby giving a calibration curve of wavelength against pixel number. Example calibration data for a forward spectrometer is shown in Figure 6.7. In Figure 6.7 (a), the intensity reading on the spectrometer is plotted as a function of pixel value showing clearly defined peaks in the spectrum. A peak finding algorithm in Matlab is used to identify the peaks. The pixel positions of the peaks on the spectrometer are compared to the known Mercury (Argon) peaks in Figure 6.7 (b). The fibre spectrometer (*Ocean Optics Maya 2000Pro*) had a manufacturer resolution of 0.7 nm, which is a combination of the dispersion per pixel and the width of the slit that is imaged onto the detector.

The spectral transmission of the optics in the forward diagnostic beam path was not flat across the wavelength range of interest and hence had to be accounted for. In particular, the beamsplitter (BS) used between the exit mode diagnostic and the beam path to the fibre spectrometer did not have a spectrally flat response over the wavelength range of interest (700 nm to 900 nm). This was achieved by placing a white light source (*Thorlabs SLS201L/M*) with known spectral emission in the beamline and recording the response of the optical system. Let the white light source spectral emission

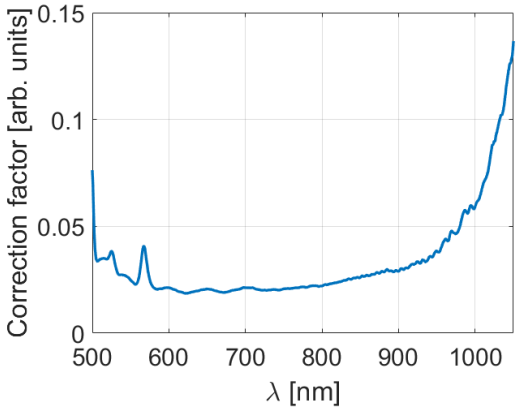


Figure 6.8: The fibre spectrometer raw data was multiplied by the calibration curve shown in the figure, extracted from measurements of the spectral response of the forward diagnostics beam path when illuminated by a white light source, as described in the text.

curve be denoted $A_W(\lambda)$. The spectrum of the white light source recorded by the fibre spectrometer was processed (wavelength calibration and background correction, as described below, in Section 6.5.1) to give the measured white light spectrum $S_W(\lambda)$. The spectrum could then be corrected by multiplying the measured spectrum on-shot by a calibration curve given by $A_W(\lambda)/S_W(\lambda)$. The curve is shown in Figure 6.8.

Ideally, for the calibration of the spectral response, the white light source would be placed at the object plane of the forward diagnostics so that the entire spectral transmission from the object plane to the spectrometer could be accounted for. However, the three uncoated wedges in the beam path to the forward diagnostics had very low reflectivity which meant that if the white light source was placed upstream of the wedges, the light from the source would not have produced a strong enough signal on the spectrometer. Therefore, the white light source had to be placed after the final wedge (W3). The optics not included in the calibration (the uncoated glass substrates, the spherical mirror and the on-axis parabola) were expected to have very little variation in spectral response over the 700 – 900 nm range of the drive beam, according to manufacturer data.

6.3.3 Input laser beam diagnostics

The energy of each of the South and North input laser beams was measured on-shot in the laser area (henceforth referred to as the LA energy readings) and provided by the laser facility (for schematic of the layout in the laser area see Figure 6.9. The LA energy reading had to be calibrated to the energy in the beams measured at the position of the start of the gas target (henceforth referred to as the interaction point, IP). This

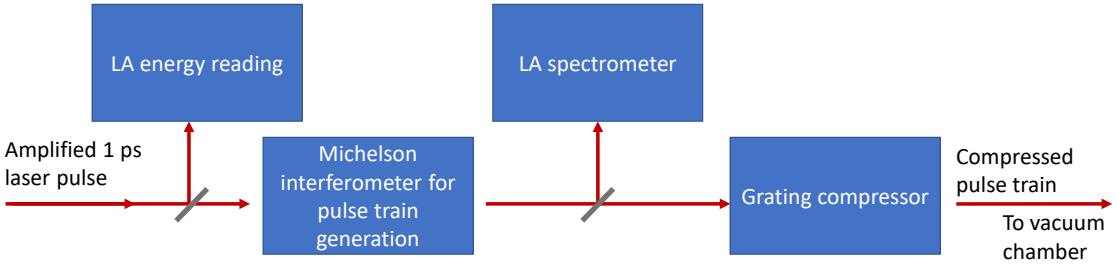


Figure 6.9: Schematic of the layout in the laser area, for the South beamline, before being directed into the vacuum chamber.

was to account for any energy lost by the beam in propagating from the LA energy reading position to the IP. The calibration was performed in a pulsed alignment mode of the laser beam using an energy meter (*Gentec QE12SP-S-MB-D0*) placed at the IP. The reading on the energy meter at the IP was averaged over 100 shots for a given LA energy reading to give the correction factor. The transmission factor from the LA energy reading pre-compressor to the IP, including the 50% energy lost in passing through the pulse-train-generating Michelson interferometer, was 19.8%.

In addition, we installed an on-shot spectrometer on the drive beam, placed directly after the pulse-train-generating Michelson. We found in the experiment that the spectral peaks in the drive beam shifted within the spectral envelope of the beam from shot to shot (see Figure 6.10). This spectral jitter was believed to arise from small ($\lesssim 1\,\mu\text{m}$) fluctuations in the optical path lengths in the Michelson arms. The six shots randomly selected for Figure 6.10 appear to show a bimodal distribution, with shots 1, 2 and 3 having one set of spectral peaks and shots 4, 5 and 6 another. Investigation of the behaviour of the spectral shifts over short and long timescales is beyond the scope of this chapter.

In the analysis of the spectral data described below, the shot-to-shot spectral shifts on the input beam were accounted for with the following procedure. First, a reference spectrum for the data set was selected (the first data shot in the set) and the peaks in the spectrum identified. Here, a data set means a series of data taken within a few hour time period, with fixed settings for the pulse-generating Michelson and compressor. Then, on each shot, the peak positions in the input spectrum of that shot were compared to the positions of the reference spectrum to give a shift value for each peak, $\delta\lambda$. The mean of the peaks shift values was calculated, $\overline{\delta\lambda}$. The output spectrum, $f_{\text{raw}}(\lambda)$ for that shot was then translated as $f(\lambda) = f_{\text{raw}}(\lambda + \overline{\delta\lambda})$. Accounting for the shifts in the

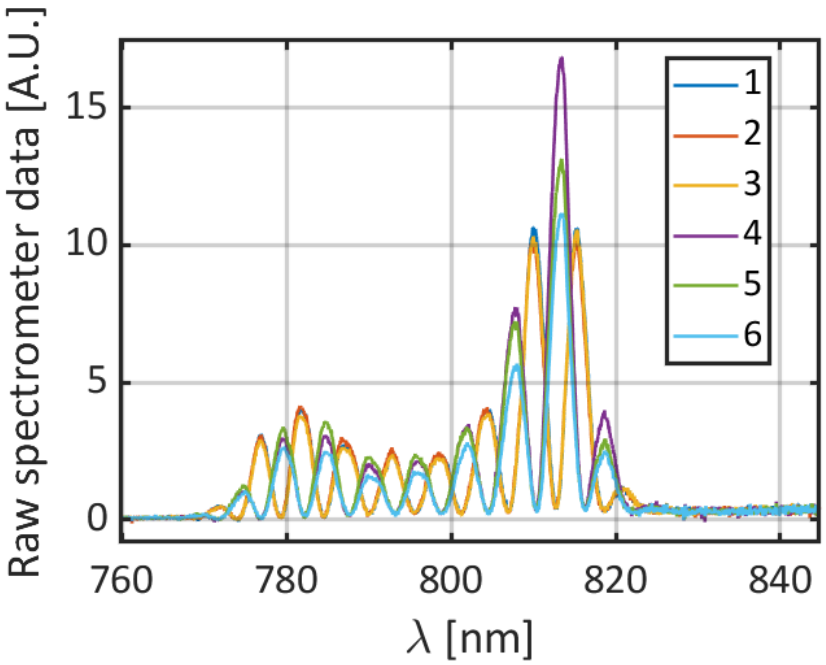


Figure 6.10: Raw data of the pulse-train input spectrum measured by a spectrometer in the laser area for a series of 6 successive shots, demonstrating the shot-to-shot variation in the position of the spectral peaks. Note that the envelope is not the true spectral intensity envelope of the pulse train as the measurements here have not been corrected for the non-uniform white-light response of the measurement system.

input spectrum in this way enabled the features in the transmitted spectra to be more directly compared.[3]

6.4 Guiding Joule-scale pulse trains over 10 centimetres

In this section, results of the first demonstration of guiding of pulse trains in long plasma channels are presented. Unless stated otherwise, the pulse trains used for this work had a pulse spacing of 170 fs. The images presented here were processed by James Chappell using the methods discussed in Chapter 3. The calibration factor between gas-target pressure reading and on-axis plasma density was calculated by James Chappell and discussed in Section 6.5.3.

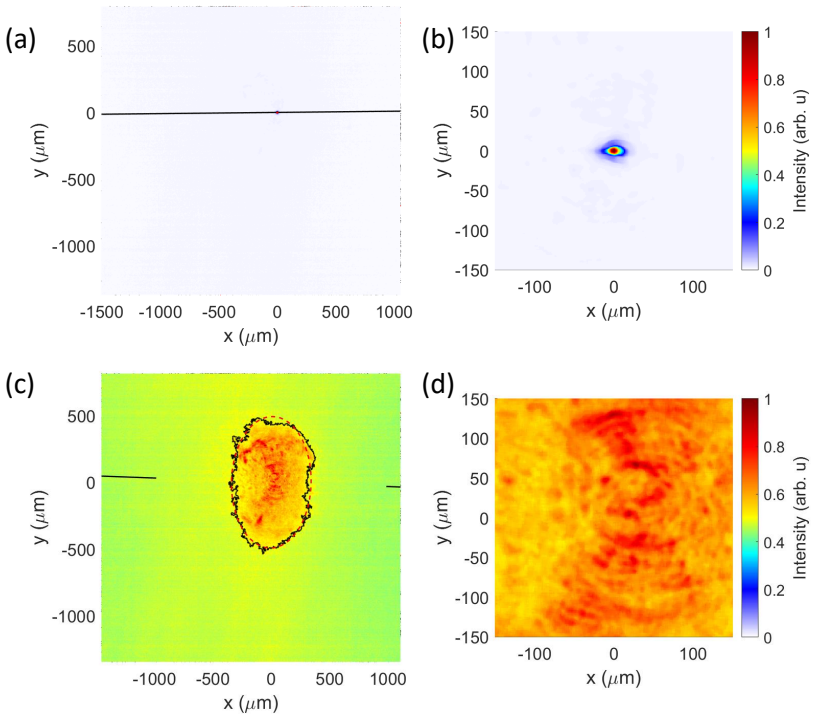


Figure 6.11: Example exit mode images of shots were the drive laser was guided, (a) and (b), compared to when the laser was not guided, (c) and (d), by the plasma channel due to drive beam pointing fluctuations at the input of the channel. (a) and (c) are the raw images and (b) and (d) are the same data zoomed in to a region of interest centred on the highest counts pixel. The image colour-bar is normalised to the highest counts pixel in each image. Plots generated by Linus Feder.

6.4.1 Selecting well-guided shots for analysis

Due to the significant transverse jitter of the drive beam focus position ($\sim 30\,\mu m$), there were a number of shots for which the drive beam was not coupled into the plasma channel. Ideally, a diagnostic would be used to measure the shot-to-shot jitter of the drive beam to enable the badly-aligned shots to be independently identified and removed in the analysis. However, unfortunately this data was not collected on the experiment. Therefore, to distinguish the signal in the data corresponding to well-guided shots, a selection procedure was employed to filter the data before the analysis. The images from the exit mode diagnostic were used to define the selection criteria. When the drive beam was well-guided, the image showed a well-defined spot or collection of spots, as shown in Figure 6.11 (a) and (b). The image from each shot was analysed to give

[3]Note that the spacing of the spectral peaks in the input beam is linearly related to the spacing of the peaks of the pulse train in the time domain from Equation (3.12). The spectral jitter measured here appeared to change the position of the peaks in the spectrum relative to the envelope, but the spacing of the spectral peaks remained approximately constant. Therefore, we interpret that the spacing of the pulses in the train was not significantly affected by this spectral jitter (within the uncertainty on the pulse train temporal spacing measurement).

the number of spots, the spot sizes and the average counts per pixel within each spot using the methods described in Section 3.4. When the beam was not guided, the image consisted of a low-amplitude, speckled light distribution, as shown in Figure 6.11 (c) and (d). The analysis of the images for the non-guided shots identified many low-signal spot regions due to the non-uniform light distribution in the image. For the shot to be considered well-guided, the image was required to satisfy the following criteria;

1. The image must contain fewer than 4 spots;

2. All spots identified in the image must have a spot size smaller than some threshold value for that data set;

3. The average counts per pixel of each identified spot must be above some threshold value for that data set.

The threshold values for spot size and average counts per pixel were determined for each data set separately. One data set refers to a series of shots taken with the same input beam properties over a few-hour time period. This accounted for any differences in the alignment of the imaging system to the forward diagnostics between data sets. Histograms of the spot size[4] and average fluence distributions for one data set are shown in Figure 6.12. A Gaussian mixture model (GMM) [166] was fit to the data in each histogram to identify peaks in the distribution. The GMM output a combination of n Gaussian peaks for each distribution, with the i^{th} peak having a mean $\overline{x}_i$ and standard deviation σ_i. The number of peaks in the fitted GMM was allowed to vary between 1 and 6. The best fit value for n was selected using the Akaike information criterion (AIC), which is a metric used for GMMs that includes both a goodness-of-fit measure and a preference for smaller models (fewer components) over larger models [166]. A full description of GMMs and their associated criteria is beyond the scope of this chapter, and further details can be found in [166].

The Gaussian peaks in the histogram going from those with smaller mean values to larger are labelled as $i = 1$ to n. In the histogram for average counts, for example that shown in Figure 6.12 (b), the Gaussian peak in the GMM with the lowest mean position was taken as the set of failed shots. The threshold value was set at $T_C = \overline{x}_{C,1} + \sigma_{C,1}$

[4]When multiple spots were identified in the image analysis, the spot with the largest spot size was included in the histogram.

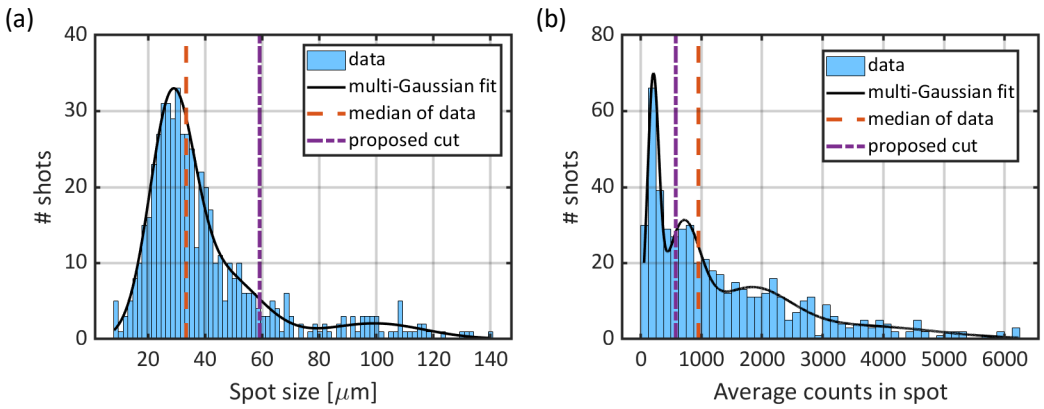


Figure 6.12: Histograms showing; (a) the spot size; and (b) the average counts per pixel in the spot, extracted from the analysis of the exit mode diagnostic images for the data set with pulse spacing 170 fs and gas target length 110 mm. When multiple spots were detected in the image, the spot with the largest spot size is plotted. The black curve indicates the fitted Gaussian mixture model. The red dashed line shows the median value of the data in each case. The purple dashed line shows the threshold values calculated by the method described in the text. All data with spot size greater than the spot size threshold value and average counts lower than the average counts threshold value were discarded.

such that all shots with average counts per pixel $< T_C$ were removed from the data set. In the histogram of spot sizes, for example that shown in Figure 6.12 (a), the Gaussian peak in the GMM with the largest mean position was taken as the set of failed shots. In this case, the n^{th} peak typically had a very large width which made cutting based on the properties of the n^{th} peak unrepresentative with different effects on each data set. Therefore, the threshold spot size was calculated as $T_S = \overline{x}_{S,n-1} + 3 \times \sigma_{S,n-1}$ such that all shots with spot size $> T_S$ were removed from the data set. By using $3 \times \sigma_{S,n-1}$, as opposed to for example $1 \times \sigma_{S,n-1}$ increased the number of shots retained in the data set so that fewer shots were falsely removed.

This procedure removed between 25% and 50% of the total shots from each data set. For example, in the data set associated to the scan of backing pressure for a pulse train with 170 fs pulse spacing in a channel of length 110 mm, there was a total of 957 data shots, of which 260 were 'null' shots (where 'null' here means data for which the gas inlet valve was not opened), 411 were 'well-guided' and 286 were 'unguided'. For data shots where the channel-forming beam or gas was switched off (i.e. without a plasma channel), the selection procedure marked 100% of these shots as unguided.

For the following results, only the data which satisfy the criteria above are included in the analysis.

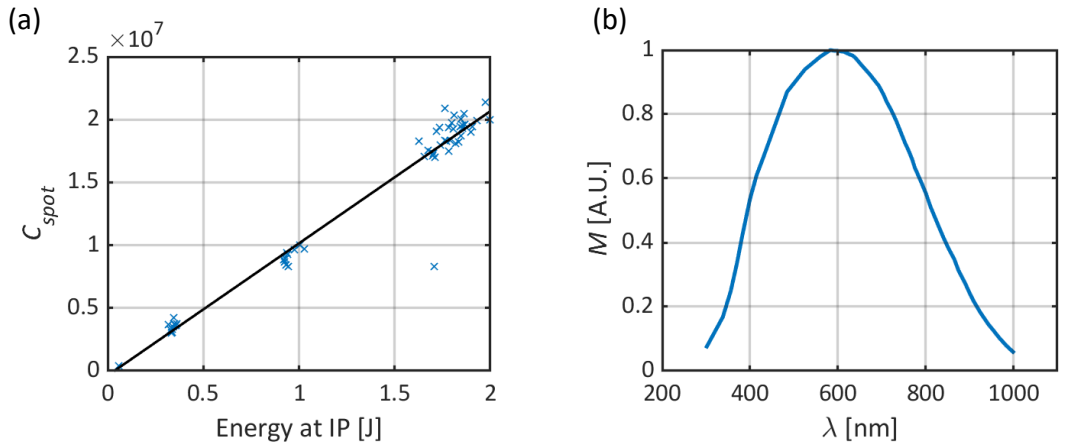


Figure 6.13: Calibrating the energy in the transmitted laser beam using the exit mode diagnostic. (a) shows the calibration between the total counts in the image of a laser spot at the drive beam focus (IP) without the gas target (C_{spot}) against the energy in the beam at the IP. The black line shows the linear fit $U = (C_{spot} - p_0)/p_1$ with $p_1 = 1.0 \pm 0.1 \times 10^7$ and p_0 consistent with zero. (b) shows the spectral response $M(\lambda)$ of the exit mode diagnostic camera chip, given by the manufacturer.

6.4.2 Calculating the output energy in the laser beam

The energy contained in the spot at the object plane is found by summing all the counts per pixel within the ellipse ROI defined via the procedure in Section 3.4 (C_{spot}) and calibrating the values of C_{spot} against energy. For the energy calibration, the gas target was moved out of the laser beam path and the imaging system was set up to image the focus of the drive beam. Images of the drive beam focus were taken whilst simultaneously recording the laser area (LA) energy reading on-shot (which had already been calibrated for the corresponding energy at the IP, as discussed in Section 6.3.3). Figure 6.13 (a) shows the calculated values for C_{spot} plotted against the energy at the IP in each shot. These data are fitted with a linear dependence. The fitted line was calculated as $U = (C_{spot} - p_0)/p_1$ with $p_1 = 1.0 \pm 0.1 \times 10^7$ and p_0 consistent with zero ($p_0 = -0.04 \pm 0.1 \times 10^7$). Note that the outlier in Figure 6.13 (a) with energy at IP of $\approx$ 1.7 J was due to a strange output of the laser on this shot giving a aberrated focal spot on this shot (the analysis for the calibration data assumed the exit mode image consisted of a single spot). The outlier does not have a significant effect on the calibration.

The spectral response of the camera chip also had to be taken into account. The spectrum of the drive beam transmitted by the plasma channel varied shot-to-shot (for example spectra, look ahead to Figure 6.18). Because the camera did not have a flat spectral response across the wavelength range of interest (700 nm to 900 nm), the

reading of counts per pixel would change shot-to-shot depending on the spectrum of the transmitted beam,. The white light response of the camera could not be directly measured during the experiment. Therefore, the spectral response curve of the camera chip (*Andor NEO 5.5 sCMOS*) from the manufacturer, $M(\lambda)$, was used, as shown in Figure 6.13 (b).[5] The spectral correction factor, F, for each shot was then calculated as;

$$F = \frac{\sum_\lambda S(\lambda)}{\sum_\lambda S(\lambda) \times M(\lambda)} \tag{6.2}$$

where $S(\lambda)$ is the background corrected spectral intensity, measured by the fibre spectrometer in the forward diagnostics. The sum over λ here is the discretised integral underneath the spectrum curve. The calculation was also performed for the input spectra to give F_{in}, using the same data shots as for the absolute energy calibration in Figure 6.13 (a). The factor for the transmitted beam was normalised to the average of F_{in}, to give $F_{norm} = F/\langle F_{in} \rangle$ for each transmitted beam. By calculating F_{norm} for each shot, the difference in multiplicative factor between the data used for the absolute energy calibration and the data for the transmitted beam was accounted for. The transmitted energy on each shot was then calculated from C_{spot} as $(C_{spot}/p_1) \times F_{norm}$.

The procedure described above provided an energy calibration for one particular ND filter setting on the camera. The ND filters were changed throughout the experiment, in order to get the highest contrast image data in the different laser-plasma regimes. Calibration images were taken with each ND filter setting at a known IP energy. Ratios were calculated between the values of C_{spot} for each ND filter setting to the setting used in the absolute energy calibration in Figure 6.13 (a).

6.4.3 Guiding measurements

Figure 6.14 shows the transverse intensity profiles for a pulse train with pulse spacing 170 fs and on-target energy 2.5 J, at the input (a) and output (b) of a plasma channel of length 110 mm. For this shot, the transverse intensity profile of the pulse train at the end of the plasma channel is a clearly defined spot, indicating that the plasma

[5] Note that for the true correction factor, the spectral response of the optics from the IP to the camera should also be taken into account in $M(\lambda)$. The spectral response of the optics can be calculated by taking the white light calibration factor of the fibre spectrometer and optics combined, $S_W(\lambda)/A_W(\lambda)$, and dividing it by the spectral response curve of the Maya spectrometer from manufacturer data. However, in practice, the spectral responses of the forward diagnostics in the range of wavelengths of interest, 700 – 900 nm, were found to be dominated by that of the fibre spectrometer chip and camera chip respectively, and hence the optics response was neglected here.

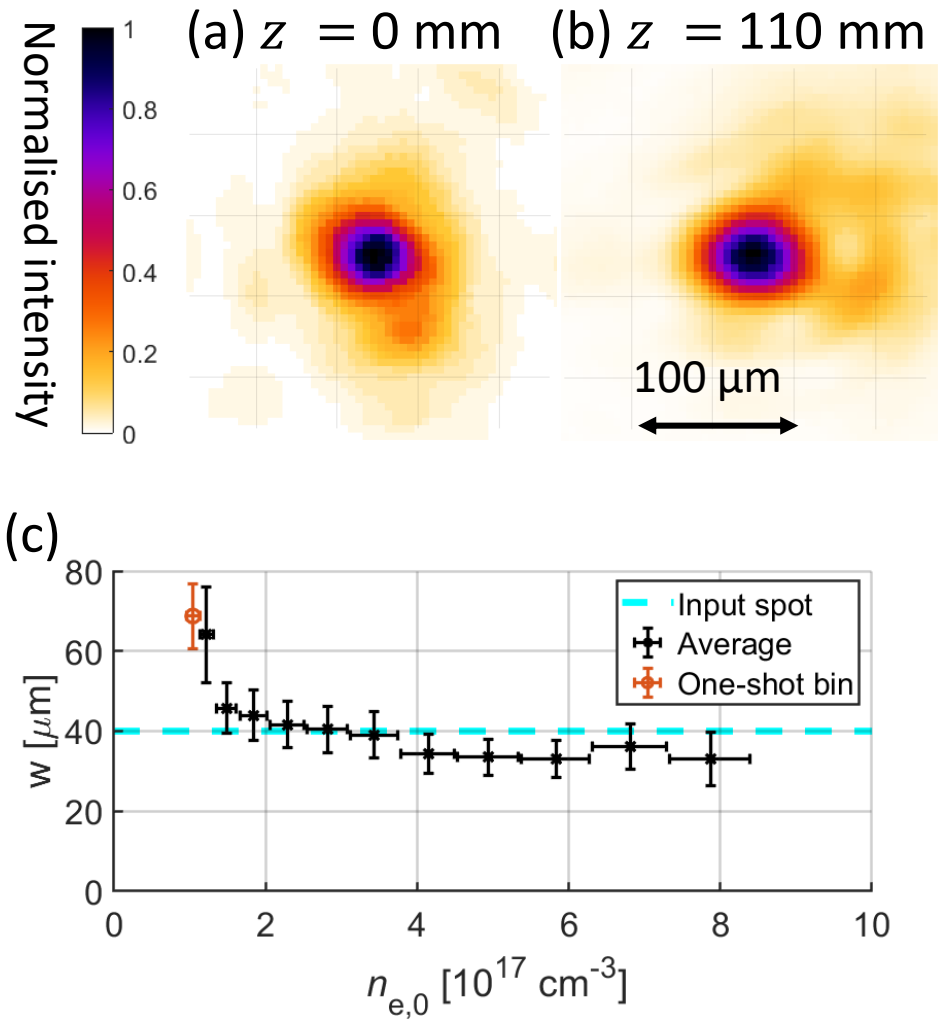


Figure 6.14: Guiding of 170 fs-spaced pulse train with 2.5 J on-target energy over 110 mm. (a) Transverse intensity profile of the input laser spot of the pulse train, measured by the exit mode diagnostic. (b) Transverse intensity profile of the output laser spot after propagation in 110 mm gas target, demonstrating that guiding has been achieved. (c) Measured variation of spot size of the transmitted pulse train as a function of on-axis plasma density inside the channel. The dashed line shows the input spot size. Red markers indicate pressure bins with only one shot contributing.

channel successfully guided the pulse train. The propagation length here corresponds to 17 Rayleigh ranges of the input drive beam. In absence of the plasma channel, the spot size of the drive beam in the exit plane of the plasma channel would have been approximately $w(z = 17z_R) \simeq 680\,\mu$m.

In Figure 6.14 (c), the spot size of the transmitted pulse train at the output of the plasma channel is plotted against the on-axis density in the plasma channel,[6] for a range of densities relevant to this experiment. The spot size is calculated using the $D4\sigma$ method described in Section 3.4. The horizontal error bars represent the uncertainty on the measured pressure, combined with the uncertainty of the on-axis density calibration.

[6]See Section 6.5.3 for method used to calibrate the on-axis density.

The vertical error bars represent the standard error of the spot size calculated for different shots.

The spot size of the pulse train at the input of the channel is approximately recreated at the output over a range of densities between $1.5 \times 10^{17}\,\mathrm{cm}^{-3}$ and $8 \times 10^{17}\,\mathrm{cm}^{-3}$. There is a tendency for the output spot size to decrease with increasing on-axis density. This behaviour could be explained by the dependence of the matched spot size of the plasma channel on the gas pressure in the target. In the model discussed in Section 3.2, as the drive beam propagates, the higher order modes are radiated away such that the output beam predominantly consists of the fundamental channel mode. In 2020, Picksley *et al*. found that the matched spot size of the plasma channel decreases as fill pressure is increased [1]. Here, $n_{\mathrm{e},0}$ is taken to be proportional to fill pressure (see Section 6.5.3) and hence, the decrease in spot size with density observed in this experiment is consistent with that in reference [1].

Matching of input beam to lowest-order channel mode

It is desirable to match the input beam profile with the lowest-order mode of the plasma channel, as the lowest-order mode has the longest attenuation length of all the channel modes (see Section 3.2). To estimate what fraction of the input beam energy is coupled to the lowest-order mode of the plasma channel in a well-aligned case, an overlap integral can be calculated between the input beam amplitude profile and an estimate of the lowest-order channel mode profile. This estimate can be used as a sanity check for the coupling efficiency extracted from energy transmission measurements below (see text regarding Figure 6.16. Due to the lack of a transverse probe on this experiment, the plasma channel transverse profile was not known and so the modes of the channel could not be calculated from measurement. Therefore, to approximate the lowest-order mode of the channel for this calculation, the transverse amplitude profile of an output beam was used. This method assumes that the higher-order modes radiated away completely by the end of the channel, leaving only the lowest-order channel mode remaining. Simulations run for previous experiments into a similar HOFI plasma channel, used for the guiding of single short laser pulses, suggested that higher order modes were attenuated over a length of $\approx 30\,\mathrm{mm}$ and the fundamental mode over a length-scale of $\approx 100\,\mathrm{mm}$ [1]. Such simulations could not be run for the experiment in this chapter, because the

transverse profile of the plasma channel was not measured, however the information from the simulations in [1] can be used to motivate the reasonableness of this approach in giving a rough approximation to the lowest order mode for the case where the drive beam is well aligned into the channel.

The output beam from the plasma channel varied in shape shot-to-shot, due to fluctuations in the drive beam pointing leading to differences in beam propagation in the channel. A systematic analysis of which data shots had output profiles that best approximated a single lowest order mode would have been a difficult process. Therefore, to achieve an approximation of the lowest-order mode of the channel, a shot was visually selected from the data set for which the transverse intensity profile of the guided pulse train appeared to be well described by a single bright spot with smooth transverse intensity profile. For this calculation, the spatial phases of the input beam and the transmitted beam were assumed to be flat. The images of the input and transmitted beams were first pre-processed as described in Section 3.4 and centred on the centre of the fitted ellipse. The overlap integral was then calculated as $\eta = \frac{(\int \sqrt{I_1 \times I_2} dA)^2}{\int I_1 dA \int I_2 dA}$ where I_1 and I_2 are the transverse intensity profiles of the input and channel modes respectively and A is the area of the integral, here taken to be a circle. The radius of the circle was increased until the value of the overlap integral converged.

The on-axis density for the calculation was $n_{e,0} = 2.4 \times 10^{17}\,\mathrm{cm}^{-3}$, as for this density the spot size of the transmitted beam most closely matched that of the input beam, shown in Figure 6.14 (c). This analysis suggests that the overlap of the input transverse profile and the fundamental channel mode for a channel with $n_{e,0} = 2.4 \times 10^{17}\,\mathrm{cm}^{-3}$ was $85\% \pm 20\%$.

This is a highly oversimplified analysis of the coupling of the input beam into the plasma channel, but provides an upper bound estimate. In reality, a number of effects including density ramps at the channel entrance, the effects of conditioning of the channel by the drive beam, and the significant spatial jitter of the drive beam would all be expected to impact the coupling of the beam into the channel, as discussed below.

Energy transmission

The energy transmission as a function of pressure for a gas target length of 110 mm is shown in Figure 6.15 for the case of a pulse train with 170 fs spacing. The energy

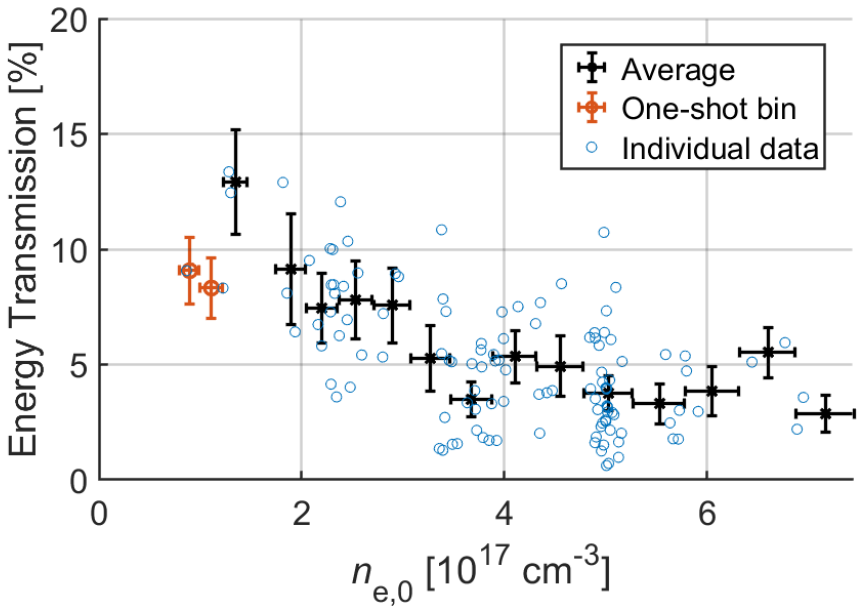


Figure 6.15: Average energy transmission as a function of on-axis plasma density for a target-length of 110 mm and pulse trains of energy 2.5 J and spacing 170 fs.

transmission values are defined by the total counts in the ROI on the exit mode camera as described above, calibrated to energy and divided by the energy on target measured shot-to-shot. Only the well-guided shots are included, as defined in the shot selection criteria in Section 6.4.1. The vertical error bars represent the standard error of the energy transmission calculated for different shots (where standard error is defined as the the standard deviation divided by the square root of the number of data points contributing), combined with the uncertainty of the energy calibration. The relative uncertainty due to the energy calibration was 32%, which was multiplied by the average energy transmission for each density bin, resulting in higher transmission values having larger uncertainty.

The energy transmission in Figure 6.15 appears to fall with density between $1 \times 10^{17}\,\mathrm{cm}^{-3}$ and $3 \times 10^{17}\,\mathrm{cm}^{-3}$. This may be explained in part by the decrease in channel-matched spot size with density as discussed above, which would reduce the overlap integral between the input beam and the channel mode. Figure 6.15 shows minimal variation in energy transmission with density between $3 \times 10^{17}\,\mathrm{cm}^{-3}$ and $8 \times 10^{17}\,\mathrm{cm}^{-3}$, with an average transmission efficiency of $\approx 5\%$ of the input laser beam energy. Note that the reduction in channel-matched spot size with increased density over this density range is expected to still be present, but the resulting difference in transmission may be smaller than the standard deviation of the data for each density.

To understand where the energy in the beam was lost, the energy transmission as

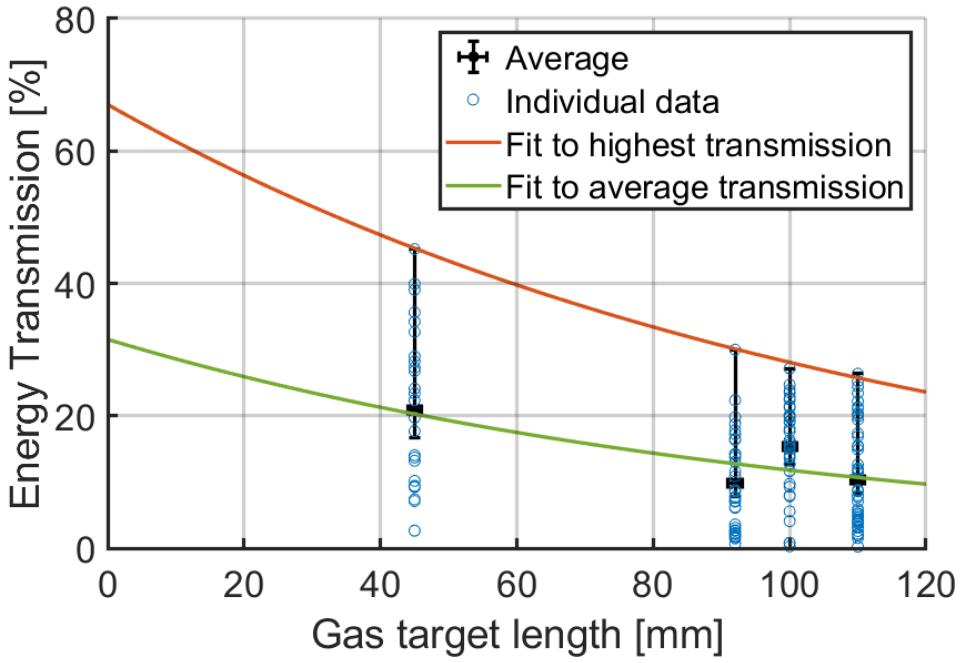


Figure 6.16: Measured energy transmission of the pulse train with spacing 200 fs as a function of plasma channel length, for on-axis plasma density of $2.9 \times 10^{17}\,\mathrm{cm}^{-3}$. The upper error-bar corresponds to the maximum transmission for each position, the lower error-bar is the standard-error on the mean.

a function of length of the plasma channel can be investigated. The plasma channel length was varied by moving the position of the motorised plunger on the gas target. The energy transmission as a function of gas target length is plotted in Figure 6.16, for the case of a 200 fs-spaced pulse train at an on-axis density of $2.9 \times 10^{17}\,\mathrm{cm}^{-3}$. The lower part of the vertical error bar is defined as for Figure 6.15. The upper part of the vertical error bar is the maximum transmission value for each on-axis density. Note that the distribution of spots about the mean is not necessarily a normal distribution as it is defined by the jitter and drift of the drive beam pointing during the time that the data set was taken, as well as the acceptance angle of the plasma channel, which is not known. The mean plotted here is used later in the Chapter to understand the spectral shifts of the transmitted beam, in Section 6.6.2. Also note that the data in Figure 6.16 was taken on a different day to that in Figure 6.15 and hence the experimental conditions, in particular the shot-to-shot pointing jitter of the drive beam, may have been different, leading to a different spread of throughput values in each density bin.

It can be seen that within each gas target length setting, there is a significant variation in energy transmission. This is expected to be partly a result of the significant pointing jitter of the drive beam, which causes the transverse position of the input pulse to vary with respect to the axis of the plasma channel. Although the shots where the laser beam

was not guided were mostly removed by the shot-selection process outlined in Section 6.4.1, some poorly guided shots were still present in the data set.

The highest transmission for each channel length is expected to correspond to conditions where the drive beam was most closely aligned to the position of the channel entrance. For the shots with the highest energy transmission at each target length, the data in Figure 6.16 were fitted to an exponential decay curve of the form $T(z) = T_0 \exp(-z/L_{att})$ with T the energy transmission, T_0 the coupling efficiency, z the target length and L_{att} the power attenuation length. Note that T_0 is not equal to 1 due to coupling losses that arise as a result of a number of experimental factors, including the input beam transverse profile not being perfectly matched to the lowest-order mode of the plasma channel as well as the possibility of the drive beam being affected in the gas that flows out beyond the pinholes of the gas target. A constant offset term is not included in this fit because it would be nonphysical, as the transmission of a plasma channel of infinite length is zero. The fit yields $T_0 = 60 \pm 10\%$ and $L_{att} = (114 \pm 20)$ mm. The coupling efficiency found in this data scan at $n_{e,0} = 2.9 \times 10^{17}\,\mathrm{cm}^{-3}$ can be compared to the overlap integral result above of $85\% \pm 20\%$ which was calculated for a comparable plasma density, $n_{e,0} = 2.4 \times 10^{17}\,\mathrm{cm}^{-3}$. The values are consistent within the errors, suggesting that for the best-aligned shots, the coupling efficiency of the beam into the channel was determined by the fraction of the input beam energy that was in the lowest-order channel mode. A fit to the average transmission values in Figure 6.16 to the equation $T(z) = T_0 \exp(-z/L_{att})$ gives a coupling efficiency of $T_0 = 32 \pm 30\%$, which takes into account the coupling losses due to spatial jitter of the drive beam with respect to the channel axis.

Comparison to guiding of a single short pulse

The energy transmission results of the experiment in this chapter can be compared to previous results obtained in the guiding of single, short laser pulses in HOFI channels. Picksley *et al.* [1] present the attenuation length and energy transmission for a single ultrashort laser pulse, in a regime similar to the one used in the experiment in this chapter. The experiment in [1] was also performed on the Astra-Gemini TA3 laser facility with very similar plasma channel parameters and with $f/40$ focusing of the drive beam, but

with a short 40 fs FWHM pulse duration pulse. The experimental parameters for the two experiments are presented in the table below.

	I [10^{17} W/cm^2]	E [J]	w [μm]	$n_{e,0}$ [10^{17} cm^{-3}]	Delay [ns]
Ref [1]	12	1.1	38	1.1	2.7
This experiment	1.6	2.5	40	2.9	3.5

The peak laser intensity I given in the table is calculated as follows. For a single laser pulse, the peak intensity is given by;

$$I = \frac{2P}{\pi w^2} \tag{6.3}$$

where w is the radius at which the transverse intensity profile falls to $1/e^2$ of the peak intensity (as defined in Chapter 3) and P is the laser beam power $P = E/\Delta t$ for E the total pulse energy and Δt the FWHM pulse duration of the intensity. For the pulse train, the temporal intensity profile retrieved from the SSA measurements was scaled such that the integral of the temporal intensity profile over time was equal to the beam fluence, $2E/(\pi w^2)$. Figure 6.17 shows the scaled intensity profile of the pulse train with $\tau = 170$ fs, $E = 2.5$ J and $w = 40$ μm. The unmodulated laser pulse with $\Delta t = 1$ ps, $E = 2.5$ J and $w = 40$ μm is shown in the same figure. The modulations in the intensity profile of the pulse train result in the peak intensity being $\sim 70\%$ greater than the unmodulated laser pulse with the same GDD and total energy. The peak intensity given in the table for this experiment is that of the pulse train with 170 fs pulse spacing.

The parameters of the two experiments are found to be very similar in all aspects other than the peak intensity of the laser beam, which is a factor ≈ 8 smaller in the experiment in this chapter due to the much longer pulse duration, and the on-axis density which is a factor of ≈ 3 larger in the experiment in this chapter. In the next table, the attenuation length and energy transmission at the end of the channel z_{end} are presented, where $z_{end} = 100$ mm in reference [1] and $z_{end} = 110$ mm in this chapter. The transmission for this chapter is the average transmission of the case shown in Figure 6.16, with $n_{e,0} = 2.9 \times 10^{17}$ cm^{-3}.

	L_{att} [mm]	$T(z_{end})$ [%]
Ref [1]	102 ± 38	$15 \pm 4\%$
This experiment	114 ± 20	$8 \pm 2\%$

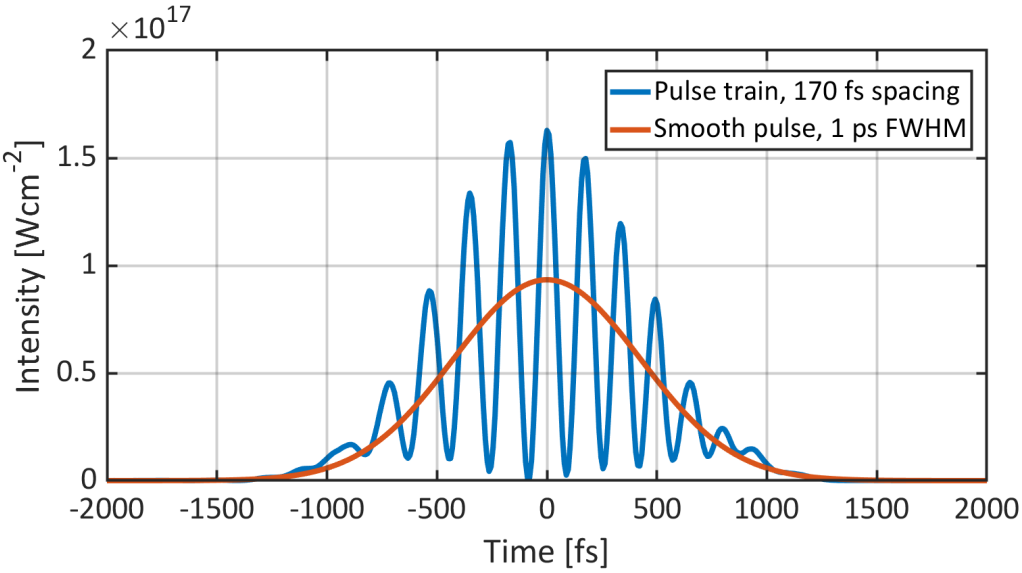


Figure 6.17: Temporal intensity profiles of; blue: the laser pulse train used in this experiment with 170 fs pulse spacing, 2.5 J energy and 40 µm spot size; red: a smooth Gaussian laser pulse with FWHM pulse duration 1 ps, energy 2.5 J and spot size 40 µm.

Both experiments found an attenuation length of around 100 mm. The energy transmission to the end of the channels were both of order $\sim 10\%$. Note that the transmission in this experiment for the lower plasma density of $\approx 1 \times 10^{17}\,\mathrm{cm}^{-3}$ (as used in Ref [1]) was higher than that of $\approx 2.9 \times 10^{17}\,\mathrm{cm}^{-3}$ by about a factor ≈ 2, as seen in Figure 6.15, which would bring the values for $T(z_{end})$ into closer agreement. The results in the table suggest that the guiding efficiency of the pulse trains in the experiment in this chapter was not significantly different to the guiding of single short laser pulses in similar plasma channels.

In other experiments, it has been found that the power attenuation length of a plasma channel is significantly extended if the guided laser pulse is intense enough to condition the channel as it propagates [21, 112] (see Section 3.2 for details on conditioned HOFI plasma channels). Simulations have been performed by the Oxford group [21] to investigate the significance of conditioning in the case of a single short laser pulse drive beam, for plasma channel conditions similar to the experiment in this chapter ($n_{\mathrm{e},0} = 1 \times 10^{17}\,\mathrm{cm}^{-3}$, $t_d = 1.5\,\mathrm{ns}$). The simulations suggest that for a single short (40 fs) pulse to effectively condition the channel, the peak intensity of the pulse needs to be $\gtrsim 3.5 \times 10^{17}\,\mathrm{W\,cm}^{-2}$. The study also found that a single pulse with peak intensity above $0.6 \times 10^{17}\,\mathrm{W\,cm}^{-2}$ was able to partially condition the channel as the very peak of the laser beam was intense enough to ionise the neutral gas collar. In the experiment in this chapter, the peak intensity of the pulse train ($\sim 1.6 \times 10^{17}\,\mathrm{W\,cm}^{-2}$) was not above

the threshold for effective conditioning found in [21], but was above that for partial conditioning. Additionally, the simulation study was not set up to exactly model the conditions of the experiment in this chapter, and hence it is unclear the extent to which the plasma channel was conditioned by the pulse train. Simulation studies to investigate conditioning effects in the case of the pulse train drive beam are ongoing.

Note that in the perfect experimental setup, the drive beam would achieve 100% coupling efficiency into the plasma channel and close to zero attenuation due to channel losses over the propagation distance. However, the drive beam would still lose energy during its propagation due to plasma wakefield generation. An estimate of energy lost to plasma wave generation is postponed to Section 6.7, as the calculation requires an estimate of the wake amplitude driven by the pulse train which is extracted from the spectrum of the transmitted drive beam presented in the next section.

6.5 Resonant excitation of plasma waves in a long plasma channel

In this section results are presented of the first demonstration of resonant excitation of plasma wakefields by a train of laser pulses guided in a plasma channel. The data have been filtered to include the well-guided shots only, as per the criteria in Section 6.4.1.

6.5.1 Processing of spectral data

The diagnostic for wakefield excitation was the spectral intensity of the transmitted drive beam. The raw spectral intensity was recorded by the spectrometer, consisting of a 1D strip of pixels. A background was measured using a series of data taken without any laser light on the spectrometer and subtracted. The spectral intensity profile was normalised in the following way; the background-corrected spectral intensity $S(\lambda)$ was divided by the energy per photon ($\propto 1/\lambda$) to convert spectral energy density to spectral photon density. This is equivalent to multiplying $S(\lambda)$ by wavelength. The area under the curve $\lambda S(\lambda)$ was calculated and used as the normalisation factor. This generated the photon-normalised spectrum defined as $f(\lambda) = \frac{\lambda S(\lambda)}{\sum_\lambda \lambda S(\lambda)}$ which can be understood as the fractional number of photons per wavelength interval. The spectra presented in this chapter are the photon-normalised spectra, $f(\lambda)$, unless otherwise stated. The choice to use the photon-normalised spectrum in this chapter highlights that the laser-plasma

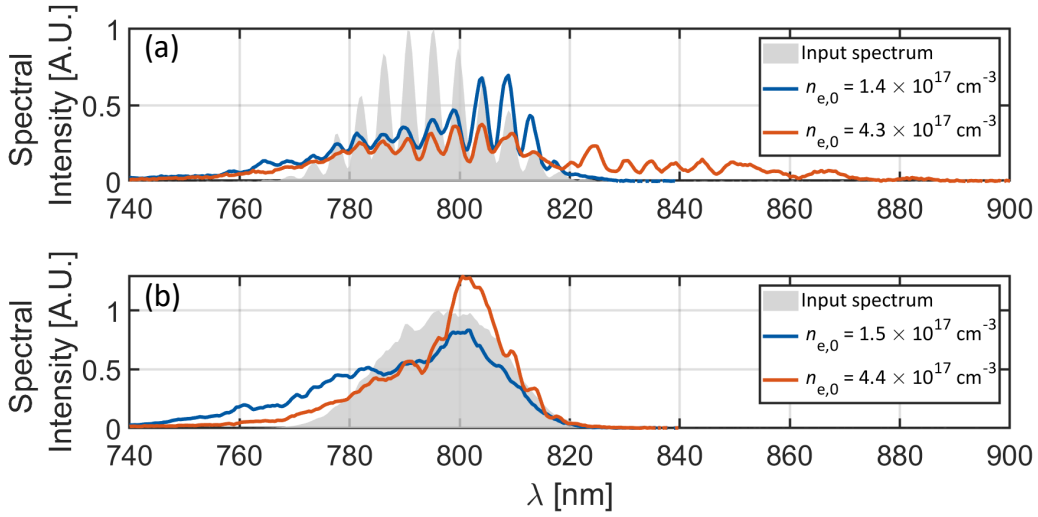


Figure 6.18: Comparison of the input spectrum with that of a beam transmitted through a 110 mm-long HOFI plasma channel for; (a) a pulse train with 170 fs pulse spacing and energy-on-target 2.5 J; (b) the unmodulated pulse with FWHM duration of 1 ps and energy-on-target of 2.7 J. The spectra of the transmitted beams are shown for an on-axis density near the resonance condition of the pulse train ($n_{\mathrm{e,\ res}} = 4.3 \times 10^{17}\,\mathrm{cm}^{-3}$) and an off-resonant density. Measured on-axis density values are given to $\pm\ 0.3 \times 10^{17}\,\mathrm{cm}^{-3}$. The spectrum in each case has been photon-normalised and scaled so that the maximum value of the input spectrum in each of (a) and (b) is equal to 1.

interaction acts to change the spectral content in the distribution of photons but does not change the number of photons (unlike energy leakage from the plasma waveguide which does result in a reduction in the number of photons).

6.5.2 Properties of transmitted spectra

Figure 6.18 (a) shows typical output spectra of the pulse train following propagation in a plasma channel. In this case, the pulse spacing was 170 fs and the channel length was 110 mm. The resonant plasma density for a pulse train with this spacing is expected to be $n_{\mathrm{e,res}} = 4.3 \times 10^{17}\,\mathrm{cm}^{-3}$. Two cases are shown, one for an on-axis plasma density far from the resonance condition and one with an on-axis density near to resonance. The spectra can be compared to the spectrum of the pulse train on the input to the plasma channel.

The input spectrum shown is characterised by clear peaks within the bandwidth of the laser beam. These peaks are a result of the spectral modulation used to generate the pulse trains in this experiment, as explained in Section 3.1.3. The output spectra retain this peaked structure but have additional features. For a plasma density far from resonance, most of the light in the output spectrum is concentrated in the same peaks

as the original input spectrum. The peaks at lower wavelengths are more strongly suppressed then at higher wavelengths and there is a tail of light into the shorter wavelength region. This may be a result of ionisation-induced blueshifting [167] in the residual neutral gas beyond the pinholes of the gas target. In ionisation-induced blueshifting, as the front of the laser pulse ionises the gas, it creates an electron density gradient which in turn induces a refractive index gradient which acts to shift the spectrum of the beam to shorter wavelengths, following the same principles as described in Section 3.3. The presence of neutral gas beyond the pinholes of the gas target is expected from the gas flow dynamics from the inlet and outlet holes (see Section 6.6.1). The preferential suppression of the blue wavelengths in the train may have been due to the fact that the shorter wavelengths arrived earlier in time compared to the long wavelengths (due to the chirp direction on the pulse train) and hence suffered greater losses due to ionisation. If the pulse train induced some degree of conditioning of the HOFI plasma channel, the later, longer wavelength peaks would see a less-lossy plasma channel and hence have greater transmission than the shorter wavelengths.

In contrast, when the electron density is close to $n_{\mathrm{e,res}}$, the spectrum looks distinctly different to the input spectrum. In the case $n_{\mathrm{e},0} = (4.3 \pm 0.3) \times 10^{17}\,\mathrm{cm}^{-3}$ in Figure 6.18 (a), very substantial redshifting of the spectrum is observed, with light extending beyond 860 nm, far outside the bandwidth of the original pulse, which falls to an amplitude close to zero by 820 nm. Significant plasma density gradients are needed to generate such a large spectral shift in the laser pulse, via the arguments in Section 3.3. Therefore, this strong redshift observed at near-resonant densities is a clear signature that a wakefield was driven in the plasma channel.

To confirm that the laser pulse train was responsible for driving the wakefield, experiments were also undertaken with a temporally smooth laser pulse. The unmodulated laser pulse was generated by placing a mirror in the South beamline to bypass the Michelson interferometer. The energy of the beam in the laser area was halved to ensure that the same energy at the IP was achieved as for the pulse train. The measured energy of the unmodulated laser pulse at the IP was (2.7 ± 0.5) J and the FWHM pulse duration was 1 ps. The transmitted spectra of the unmodulated pulse for the same axial densities are presented in Figure 6.18 (b). As expected, the input spectrum of the unmodulated pulse is smooth, and the bandwidth matches that of the pulse train. At both densities,

the transmitted spectra have the tail of blueshifted light. However, there is very little redshifted light at either density, with the maximum wavelength reaching $\lesssim$ 830 nm. Hence we can conclude that the unmodulated pulse does not drive a strong wakefield for either density, as expected. This gives further supporting evidence to the hypothesis that it is the pulse train structure that is resonantly exciting a wakefield when the density is tuned correctly.

6.5.3 On-axis density calibration

The Fourier transform of the transmitted spectra near the resonant density, as shown in Figure 6.18 (a), contained information about the frequency of the plasma wave. The following method, developed and implemented by James Chappell, was used to extract the plasma frequency from the spectra and calibrate the on-axis plasma density in the channel.

As discussed in Section 6.2.5, it was known that the on-axis density of a HOFI plasma channel varies approximately linearly with the background gas density in the gas target [1]. The proportionality factor, Ξ, between the plasma density at the moment of ionisation and the resulting on-axis plasma density of the channel for the conditions used in the experiment is expected to be in the range of 5 to 10 from previous experiments [1].

The spectral intensity of the input pulse train (e.g. shown in Figure 6.19 (a)) is modulated at $\Omega = 2\pi c/\Delta X_{\text{mich}}$, as discussed in Section 3.1.3. Therefore, a Fourier transform of the spectral intensity,[7] $FT[\tilde{I}(\omega)]$, has a dominant peak at $2\pi/\Omega = \Delta X_{\text{mich}}/c$, as shown in Figure 6.19 (b) for an example case with pulse spacing 170 fs. Near resonance, the high-amplitude wakefield acts to spectrally modulate the laser pulse train, generating side-bands in the spectrum at $\pm m\omega_p$, in the same way as discussed for P-MoPA in Section 3.1.2. According to the theory, these side-bands have the same spectral shape as the input laser spectrum but have a reduced amplitude and are offset by multiples of ω_p. In the case where only the highest amplitude side-bands are visible (with $m = 1$), we would expect the spectral intensity of the transmitted pulse to consist of the input pulse

[7] Note that without spectral phase information, the Fourier transform of the spectral intensity does not give the temporal intensity profile (the pulse train), but instead gives a function proportional to the temporal intensity profile of a fully-compressed pulse. The chirp (second order spectral phase) on the laser pulse generates the pulse train in the time domain.

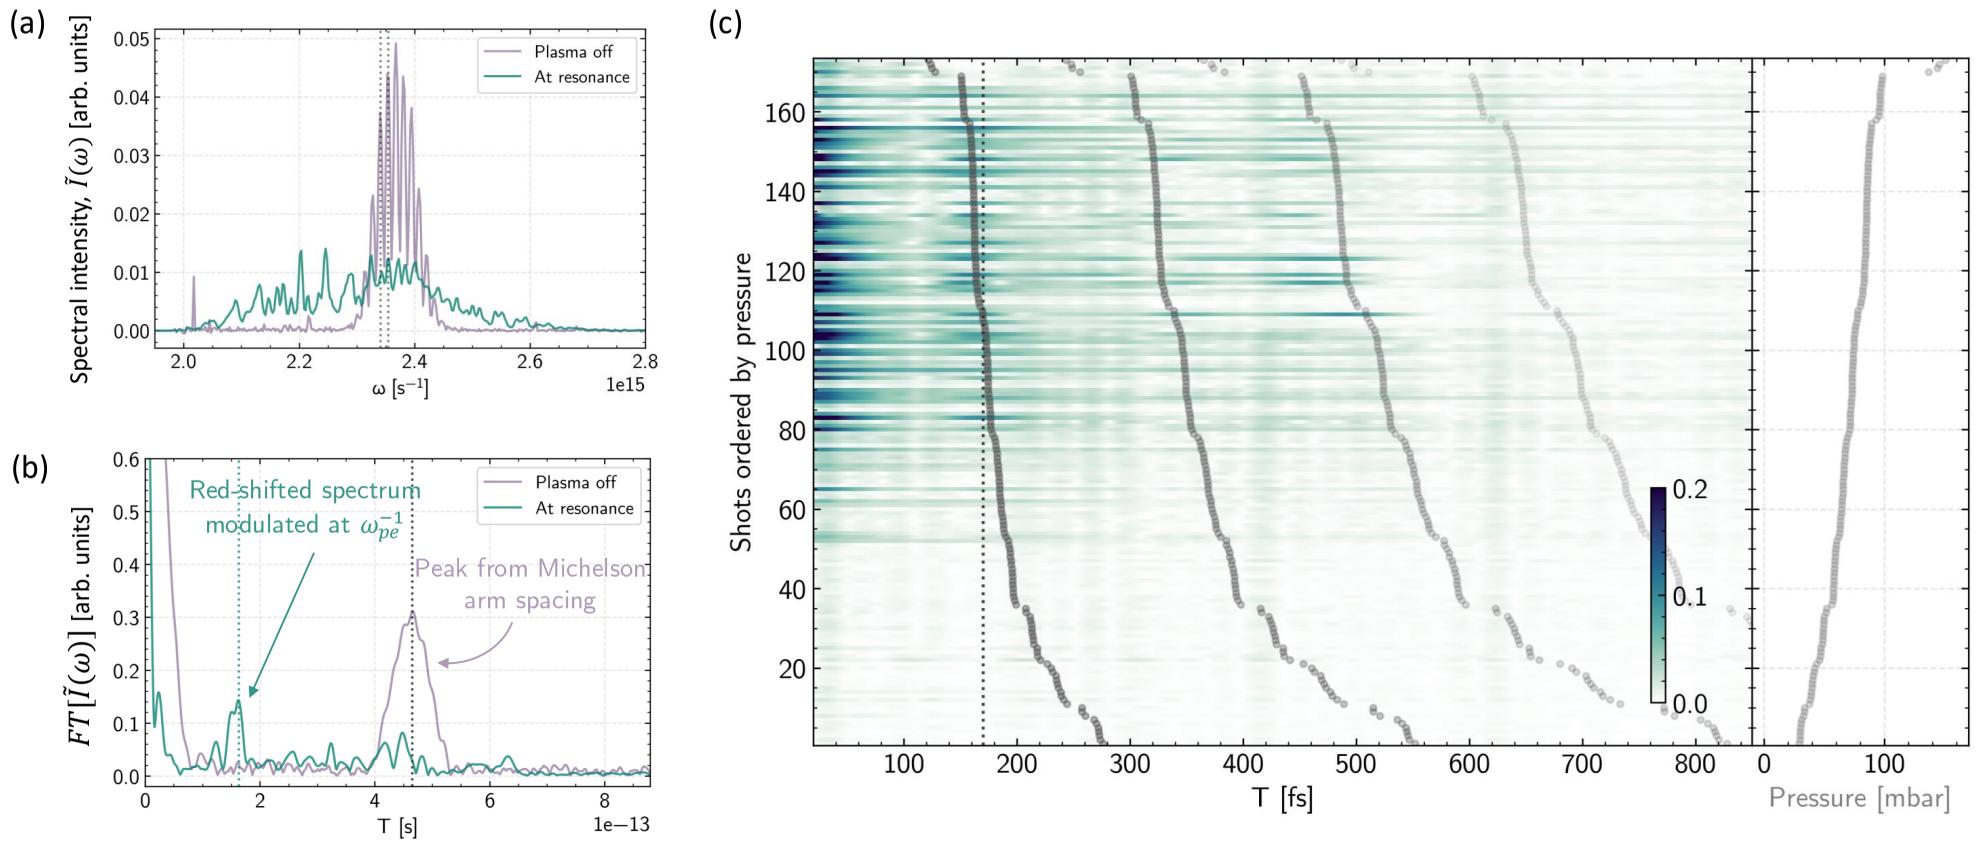


Figure 6.19: Procedure for the on-axis density calibration using signatures in the spectrum of the transmitted pulse train, for an example data set with laser pulse spacing 170 fs and gas target length 110 mm. (a) shows the spectral intensity $\tilde{I}(\omega)$ of the input pulse train (plasma off, purple) and a pulse train at resonance (green) plotted as a function of ω. (b) shows the Fourier transform of the two curves in (a), with a strong peak in the plasma off case at $\Delta X_{\text{mich}}/c$ and a new peak in the resonant case at $2\pi/\omega_p$. (c) shows $FT[\tilde{I}_{red}(\omega)]$ for a series of shots ordered by gas pressure, where $\tilde{I}_{red}(\omega)$ is the spectral intensity for the redshifted light only. The gas pressure for each shot is plotted in the panel on the right of the figure. The grey circles in the left panel indicate the position of τ_p calculated using the on-axis density calibration for each shot, as well as the positions of the higher order peaks that would be expected at $m\omega_p$ for $m = 2, 3$ and 4. The black dotted line indicates the temporal spacing of the pulses in the train $\tau = 170$ fs. The analysis and plots presented here were generated by James Chappell.

train spectrum, and two copies of the input pulse train spectrum, shifted by ω_p on either side of the central frequency of the pulse train spectrum ω_0. Therefore, the Fourier transform of the spectral intensity at resonance is expected to consist of the peak at $\Delta X_{\text{mich}}/c$, as well as a peak at $2\pi/\omega_p = \tau_p$, as indicated in Figure 6.19 (b).[8] Tracking the position of the peak at τ_p as a function of gas pressure around the resonance enables the on-axis plasma density to be extracted.

James Chappell performed the following analysis, using all data with measurable redshift (i.e. with gas pressure close to the resonance). The spectral intensity below $\approx$ 820 nm (corresponding to $\omega \gtrsim 2.30 \times 10^{15}\,\text{s}^{-1}$) was set to zero, to focus on the new redshifted light only, $\tilde{I}_{red}(\omega)$. In Figure 6.19 (c), the function $FT[\tilde{I}_{red}(\omega)]$ is plotted for a series of data taken with pulse spacing 170 fs and gas target length 100 mm, showing shots for a range of pressures, not only those that were close to resonance (figure by James Chappell). The corresponding gas pressure, P, for each shot is plotted in the panel on the right of the figure.

The conversion factor from gas pressure to density can be calculated from the ideal gas law, $n_{H2} = P/(k_B T)$ where n_{H2} is the background hydrogen molecular number density and T is the temperature of the gas before ionisation, assumed to be room temperature, $T = 300\,\text{K}$. In convenient units, this gives $n_{17} = P_{\text{mbar}}/4.14$ for n_{17} the hydrogen molecule density in units of $1 \times 10^{17}\,\text{cm}^{-3}$ and P_{mbar} the pressure in units of mbar. The electron density immediately following ionisation is therefore $n_{\text{e}} = 2 \times n_{H2}$. The relation between on-axis plasma density and $\tau_p(P)$ is given by the plasma frequency equation, $\tau_p(P) = 2\pi \times \sqrt{m_{\text{e}}\epsilon_0/(n_{\text{e},0}e^2)}$. Therefore, using $P = 2n_{\text{e}}k_B T = 2\Xi n_{\text{e},0}k_B T$, the fit equation is given by;

$$\tau_p(P) = 2\pi\sqrt{\Xi}\sqrt{\frac{2k_B T m_{\text{e}}\epsilon_0}{Pe^2}} \tag{6.4}$$

with one free parameter, Ξ. Here, only the lowest order peak position in the Fourier transform is considered (higher order peaks are also expected to be present but were not used for the analysis as they were less evident in the data).

The fitting parameter was calculated from the data shots with measurable redshift. The value extracted was $\Xi = n_{\text{e}}/n_{\text{e},0} = (8.1 \pm 0.6)$. This value is within the range

[8] These expressions can be compared to the spacing of peaks in the temporal domain of the input beam due to the chirp $\phi^{(2)}$, $\tau = \phi^{(2)}\Omega = \phi^{(2)}\frac{2\pi c}{\Delta X_{\text{mich}}}$ from Equation 3.12. At the resonant plasma density for the pulse train spacing, $\tau = \tau_p$.

expected for HOFI plasma channels [1]. Therefore, to convert the measured pressure values to on-axis plasma density values in units of $1 \times 10^{17}\,\mathrm{cm}^{-3}$, the pressure in mbar must be divided by $(8.1 \pm 0.6) \times 4.14/2$. The peak position corresponding to τ_p with this calibration factor is plotted for the series of data shots in Figure 6.19 (c). The expected peak positions from the model roughly align with the lowest order peaks in the spectrum. A more sophisticated model would allow for the relationship between $n_{\mathrm{e},0}$ and n_{e} to be a high order polynomial as opposed to a linear dependence as assumed here.

When the gas target was set shorter than full length, the gas reading on the pressure transducers was reduced for the same backing pressures. This may have been due to a partial covering of the transducer outlet by the gas target plunger and/or a change in the flow dynamics. The on-axis density calibration factor was re-calculated for the shorter gas target length used, L =70 mm, to account for any differences in the pressure transducer reading. The calibration factor in this case came to $\Xi = (6.2 \pm 0.5)$.

6.5.4 Resonance curve

To investigate the resonant response of the plasma further, a detailed pressure scan was performed, the results of which are shown in Figure 6.20 (a). These data were recorded for the same conditions as Section 6.5.2, with τ = 170 fs and channel length 110 mm. In the left part of the figure is a waterfall plot, where for each density bin, the averaged output spectrum is shown. The bins are defined in terms of a backing pressure range of 4 mbar which is then converted to on-axis density using the calibration factor Ξ discussed in the previous section. The spectral axis is cut at 818 nm, to show the details of the new redshifted light. It is evident that the spectra show a maximum redshift for densities in the range $4-5 \times 10^{17}\mathrm{cm}^{-3}$.

To quantify this redshift, a redshift metric R was defined as $R = \sum_{\lambda,\mathrm{min}}^{\infty} f(\lambda)$ where $f(\lambda)$ is the photon-normalised spectral intensity and λ, min is the threshold wavelength above which the light is said to be redshifted. Here λ, min was defined as the highest wavelength present in the input spectrum above the noise. For the pulse train with 170 fs pulse spacing, λ, min was calculated to be 825 nm.

The variation of R with on-axis density is presented as an error-plot in the right panel of Figure 6.20 (a). The R values for each shot were weighted by the energy transmission of the shot to give the weighted mean value. This weighting was chosen to give more

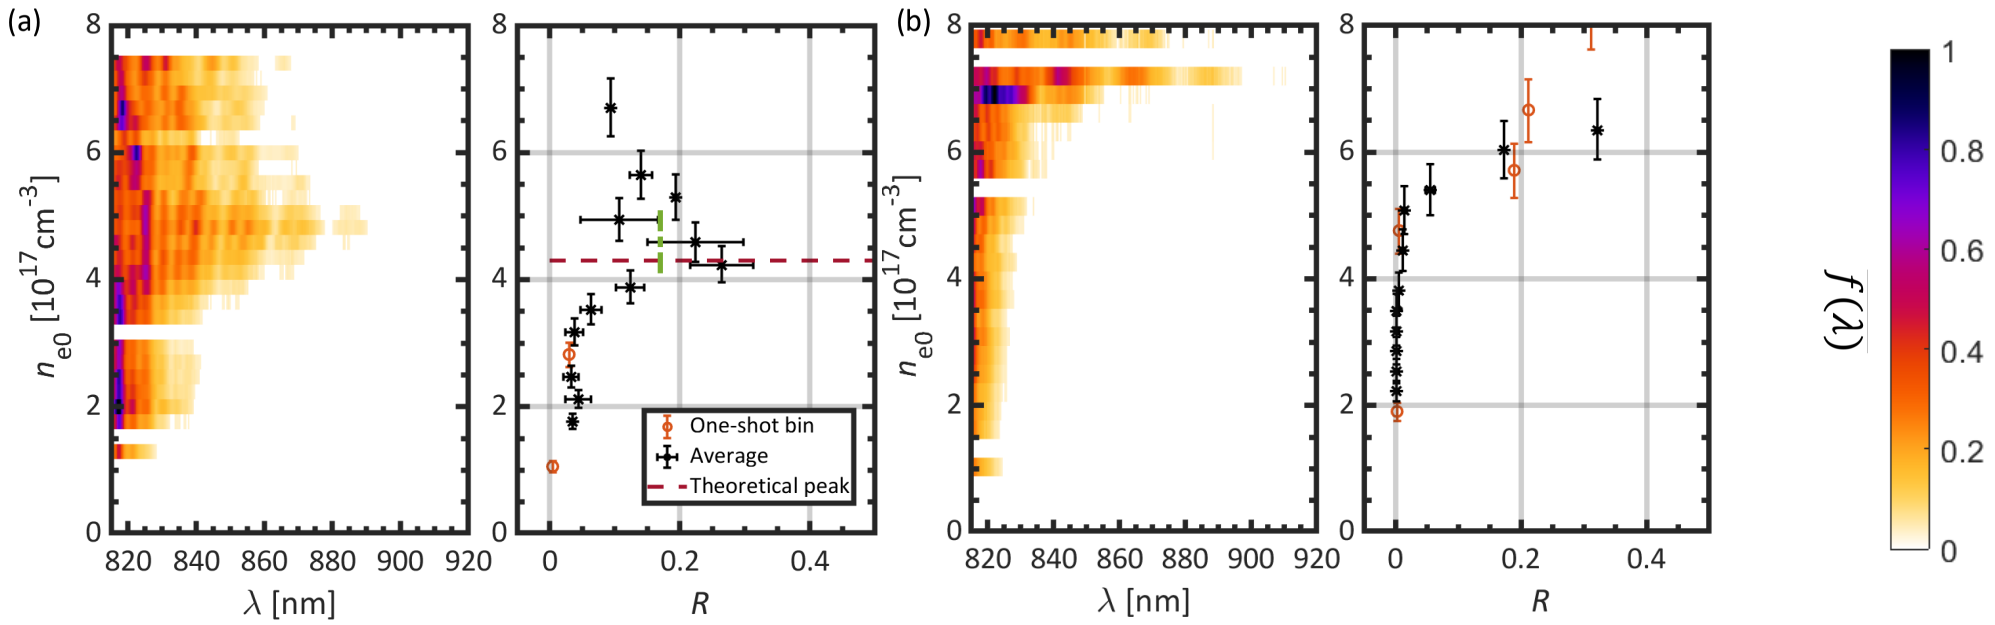


Figure 6.20: Transmitted spectra of the laser beam following 110 mm propagation in the plasma channel for; (a) the 170 fs pulse train; and (b) the unmodulated pulse. The spectra are averaged in plasma density bins of width 4 mbar. The calculated values for the redshift metric R (defined in the text) are plotted against plasma density to the right of the waterfall plot in both (a) and (b). The black crosses give the weighted mean of the data for each on-axis density, where the weighting is given by the energy transmission for each shot. The maroon dashed line in (a) is the expected resonant density for a 170 fs pulse train. The green dashed line indicates the full-width at 64% amplitude of the resonance peak. The spectra shown have been photon-normalised. A scaling is applied so that the maximum value within each data set is equal to one.

importance to the well-guided laser shots as these should drive the wakefield within the channel more effectively. Again, the peak in redshift is clearly apparent. Note that the peak position in the resonance curve in R is not expected to line up with the peak visually identified in the waterfall plot because R is calculated from the sum under the curve of the redshifted spectrum, as opposed to the maximum wavelength present in the spectrum, and the average values of R are calculated using a weighting function, whereas the data shown in the waterfall is a simple average without weighting. Also included on plot of R is the expected position of the resonance peak for the pulse-spacing of 170 fs. There is a very good agreement found between the expected resonance peak position and the peak in the R value. This provides strong evidence that the pulse train was resonantly exciting the wakefield at the expected plasma density.

In an ideal setup, the half-width of the resonance peak in wakefield amplitude would be given by $\delta n_{\mathrm{e},0}/n_{\mathrm{e,res}} \approx 1/N$ as discussed in Section 2.2.3. In this case, with $N \sim 10$ and at the resonant density $n_{\mathrm{e,res}} = 4.3 \times 10^{17}\,\mathrm{cm}^{-3}$, the half-width of the resonance would be expected to be $\delta n_{\mathrm{e},0}/2 \approx 0.43 \times 10^{17}\,\mathrm{cm}^{-3}$, giving a full-width $\delta n_{\mathrm{e},0} \approx 0.86 \times 10^{17}\,\mathrm{cm}^{-3}$. In the data shown in Figure 6.20 (a), the width of the peak, defined as the full-width at 64% amplitude in R, is $\delta n_{\mathrm{e},0} \sim 1.0 \times 10^{17}\,\mathrm{cm}^{-3}$. The correspondence between the width in the wakefield amplitude and the width in the resonance peak in R is not expected to be one-to-one because of the non-linear processes involved in generating the redshifted light. Numerical calculations are used to better understand the resonance curve in the experiment (see Section 6.6 below). Experimental factors, such as the non-uniformity of the longitudinal gas density profile (also discussed in Section 6.6), would be expected to increase the width of the peak.

We can compare the results obtained with a pulse train to those obtained with the unmodulated 1 ps pulse. This comparison is shown in Figure 6.20 (b). In the unmodulated case, there is very little redshift apparent at densities below $6 \times 10^{17}\,\mathrm{cm}^{-3}$, with R values consistently close to zero. The peak at $4-5 \times 10^{17}\mathrm{cm}^{-3}$ is only present when the pulse train is used, indicating resonant wakefield excitation.

The increasing redshift at higher plasma densities, $n_{\mathrm{e},0} > 6 \times 10^{17}\,\mathrm{cm}^{-3}$, in Figure 6.20 (b) is likely to be caused by self-modulation of the unmodulated laser pulse (see Section 2.4). As $n_{\mathrm{e},0}$ is increased, the variation in the plasma refractive index increases

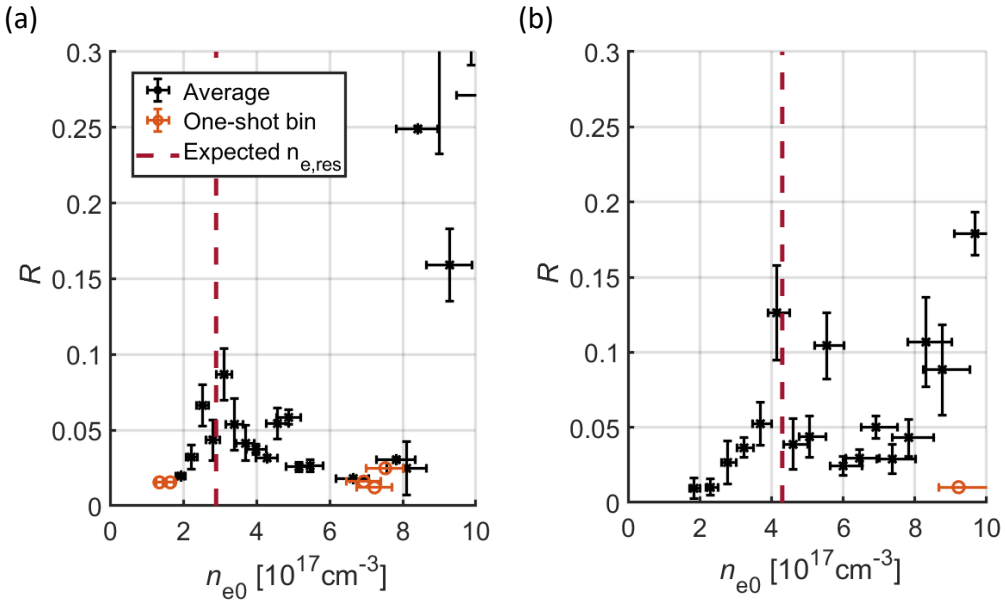


Figure 6.21: Redshift metric R resonance curves against on-axis density for the spectra after 70 mm propagation, for the case of (a) the 170 fs pulse train and (b) the 205 fs pulse train. The black crosses are the weighted mean of the data, where the weighting is given by the energy transmission for each shot. The dashed lines are the expected resonant densities for each case.

and hence the self-modulation instability becomes stronger, leading to the rapid rise in R seen here at higher plasma densities.

To provide further evidence for the resonant excitation of plasma wakefields, the pulse spacing in the train was varied to compare the positions of the resonance peaks. Figure 6.21 shows the plots of R against axial density for the case of 170 fs and 205 fs pulse spacings, here for a plasma channel length of 70 mm. In both cases, the onset of self-modulation appears at $n_{e,0} \approx 7 \times 10^{17}\,\mathrm{cm}^{-3}$, showing the redshift increasing to larger values at higher densities. This can be compared to the onset of self-modulation in the smooth pulse case seen in Figure 6.20 (b) at $n_{e,0} \approx 6 \times 10^{17}\,\mathrm{cm}^{-3}$. At lower densities, a clear resonance in R is evident for both data sets in Figure 6.21. Furthermore, the density at which the peak value of R is obtained is seen to be in good agreement with the expected resonant densities of $2.9 \times 10^{17}\,\mathrm{cm}^{-3}$ and $4.3 \times 10^{17}\,\mathrm{cm}^{-3}$ for pulse train spacings of 205 fs and 170 fs respectively.

6.6 Extracting the wakefield amplitude

In this section, we attempt to use the measured spectra of the transmitted pulse trains to deduce the amplitude of the wakefield driven. To do this, the experimental setup

is modelled in simulations and the spectral results compared to the experimental data. Simulations were performed by Johannes van de Wetering, and the results of the calculations were analysed by the author, with many useful discussions with James Chappell. The code used was a 2D cylindrical code written by Johannes van de Wetering based on qausi-static fluid theory, benchmarked against an open-source particle-in-cell (PIC) code WarpX [168]. The workings of a quasi-static fluid code are beyond the scope of this thesis, for further information please see the Supplemental Materials of ref. [4].

6.6.1 Parameters for simulations

The drive beam was modelled with a Gaussian transverse intensity profile with a spot size of 40 µm and energy of 2.5 J. The spectrum matched that of the experimental pulse, with the spectral phase terms extracted from the SSA data. The plasma channel was modelled as a pre-formed density structure (see below). The exact shape of the plasma density structure was not known. In the following subsections, the choices for longitudinal and transverse density profiles used in the simulation are discussed.

The code assumes that the laser envelope of the drive beam does not vary during the propagation. This means that self-modulation was not included and so the effects of self-modulation on the spectra of the transmitted pulses were not observed in the simulations. The assumption of fixed laser envelope is expected to hold for plasma densities below the onset of self-modulation, which in the experimental data was found to occur at $n_{e0} \approx 6 \times 10^{17}\,\mathrm{cm}^{-3}$ (for example in Figure 6.20).

Longitudinal density profile of the plasma channel

The longitudinal plasma density profile depended predominantly on the longitudinal profile of the gas density inside the gas target at the moment of ionisation. In the experiment, the gas pressure was measured at two locations by pressure transducers connected to the top edge of the gas target, one near the entrance ($z = 29\,\mathrm{mm}$ from the front pinhole) and one near the centre ($z = 59\,\mathrm{mm}$). The ratio of these two readings as a function of pressure is shown in Figure 6.22 (a). The mean ratio of these data is 0.79, suggesting that the centre pressure was consistently higher than that at the front by $\sim 20\%$.

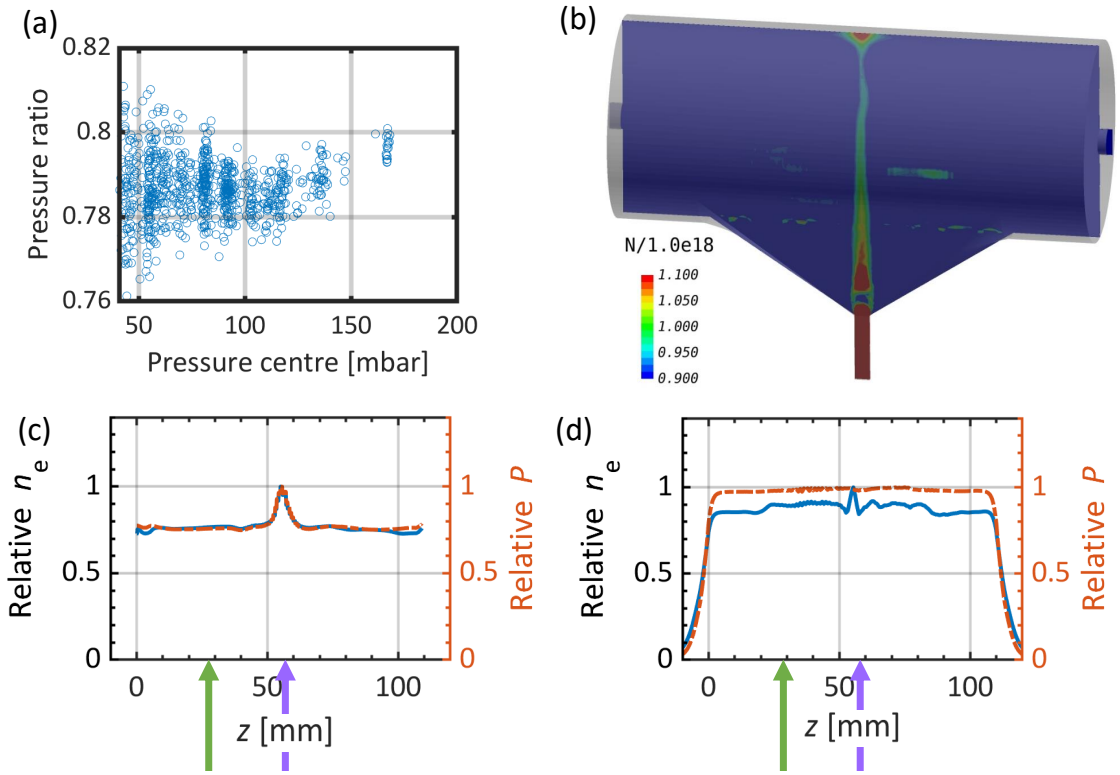


Figure 6.22: Gas pressure and density profiles inside the gas target. (a) shows the experimentally measured ratio of the pressure measured by the front transducer to that measured by the centre transducer. (b) shows the 3D gas density distribution inside the gas target for 1 bar backing pressure in steady-state conditions. (c) and (d) show the gas pressure and density profile lineouts extracted from the 3D fluid simulations in (b). The lineout in (c) is taken at the top edge of the target and the lineout in (d) is taken at the centre line of the target (laser axis). Each lineout is normalised to its peak value. The green arrow indicates the position of the front transducer and the purple arrow the centre transducer. Data shown in (b-d) was calculated by Xiao-Jun Gu at the Scientific Computing Department, STFC Daresbury Laboratory, using Code_Saturne [169], included here with permission.

To gain a better understanding of the full gas profile inside the gas target, calculations were performed by Xiao-Jun Gu at the Scientific Computing Department, STFC Daresbury Laboratory, using the computational fluid dynamic code Code_Saturne [169]. Figure 6.22 (b) shows the calculated gas density distribution for the case of 1 bar backing pressure, in steady-state conditions. The simulations show that the gas density directly above the inlet position was greater than in the surrounding gas. In simulation, the ratio of the central pressure to the pressure at the position of the front transducer measured on the *top edge* of the target shown in Figure 6.22 (c) was found to be 0.77, close to that found in the experiment. However, the pressure ratio between the centre and front transducer positions on the laser axis, shown in Figure 6.22 (d), was found to be 0.98. Therefore, in the following calculations, the on-axis longitudinal density profile was taken to be uniform. The density ramps at the ends of the target that can be seen at $z <$ 0 mm and $z >$ 110 mm in Figure 6.22 (d), with a scale length of $\sim$ 10 cm, were not included in these calculations but could be incorporated in future simulation studies.[9]

Transverse channel profile

The radial electron density profile of the plasma channel depends on the evolution dynamics of the blast wave. These are difficult to measure and simulate (as discussed in Chapter 5). Previous work suggests that the transverse plasma density profiles of HOFI channels are approximately parabolic, and better approximate the infinite parabolic case when the laser drive beam intensity is high enough to generate a conditioned HOFI channel [112]. As discussed in Section 6.4.3, the temporal intensity profile of the drive beam meant that the highest intensity parts of the pulse train were expected to partially condition the plasma channel. Therefore, for the calculations outlined in this section, the plasma channel was modelled as an infinite parabolic channel with matched spot size taken to be equal to that of the drive laser spot size of 40 µm. Ionisation effects on the drive laser beam were not included in the code. Therefore, the effects of ionisation blue-shift and energy attenuation of the pulses in the train that arrive earliest in time is not included in the code. If the plasma channel was instead modelled as a finite parabolic channel, while still neglecting envelope variation and ionisation effects, would not have a significant effect on the simulation results as the plasma density experienced by the

[9] Note that the difference between the profiles in electron density and gas pressure at the centre of the target in Figure 6.22 (d) can be explained by the upwards velocity of the gas above the inlet.

laser pulse as a function of transverse position would not be changed within the main high intensity region of the pulse transverse profile.

6.6.2 Calculation results

Figure 6.23 (a) shows the results of the 2D cylindrical fluid calculation for the case of a 170 fs-spaced pulse train over 110 mm propagation. The value for the redshift metric R is calculated in an equivalent way to that discussed in Section 6.5.4. The spectra displayed in the waterfall plot on the left are qualitatively similar to those seen in the experimental data, reprinted in Figure 6.23 (b).

The R curve from Figure 6.23 (a) is overlaid with the experimental values for R in Figure 6.23 (b). The position of the resonance peak in the simulation data matches up well with that of the experimental data. However, the amplitude of the resonance peak in the simulation is much greater than that of the experiment. The lower amplitude realised on experiment could be explained by the low coupling efficiency of the laser beam into the plasma channel. The guiding results presented in this chapter (Figure 6.16) suggest that the average coupling efficiency achieved during the experiment was of order 30%. Therefore, a reduced effective energy of the drive beam, $30\%\times$ 2.5 J $\sim$ 800 mJ, was input in the code.

The results of the same calculation with 800 mJ input energy are shown in Figure 6.23 (b). The simulation curve with the lower effective energy shows a much better agreement with the experimental data, particularly for the R-value at the peak of the resonance curve. The amplitude of the wake excited at different positions along the pulse train can be calculated using the fluid code for the conditions shown in the figure, i.e. a drive beam pulse train with 800 mJ, pulse spacing 170 fs at the resonant plasma density $4.3 \times 10^{17}\,\mathrm{cm}^{-3}$. The maximum wakefield amplitude achieved by the pulse train was calculated to be 3 GV/m. The individual data points plotted in Figure 6.23 (b) demonstrate that, for certain shots, redshift values closer to the 2.5 J simulation were achieved. The large variation in redshift measured for the individual data can be attributed to the significant pointing jitter of the drive beam at the entrance to the plasma channel. The wakefield amplitude driven by the pulse train with the full 2.5 J energy is calculated to be 10 GV/m. Therefore, we can interpret that wakefield amplitudes up to 10 GV/m were driven by the pulse trains on this experiment.

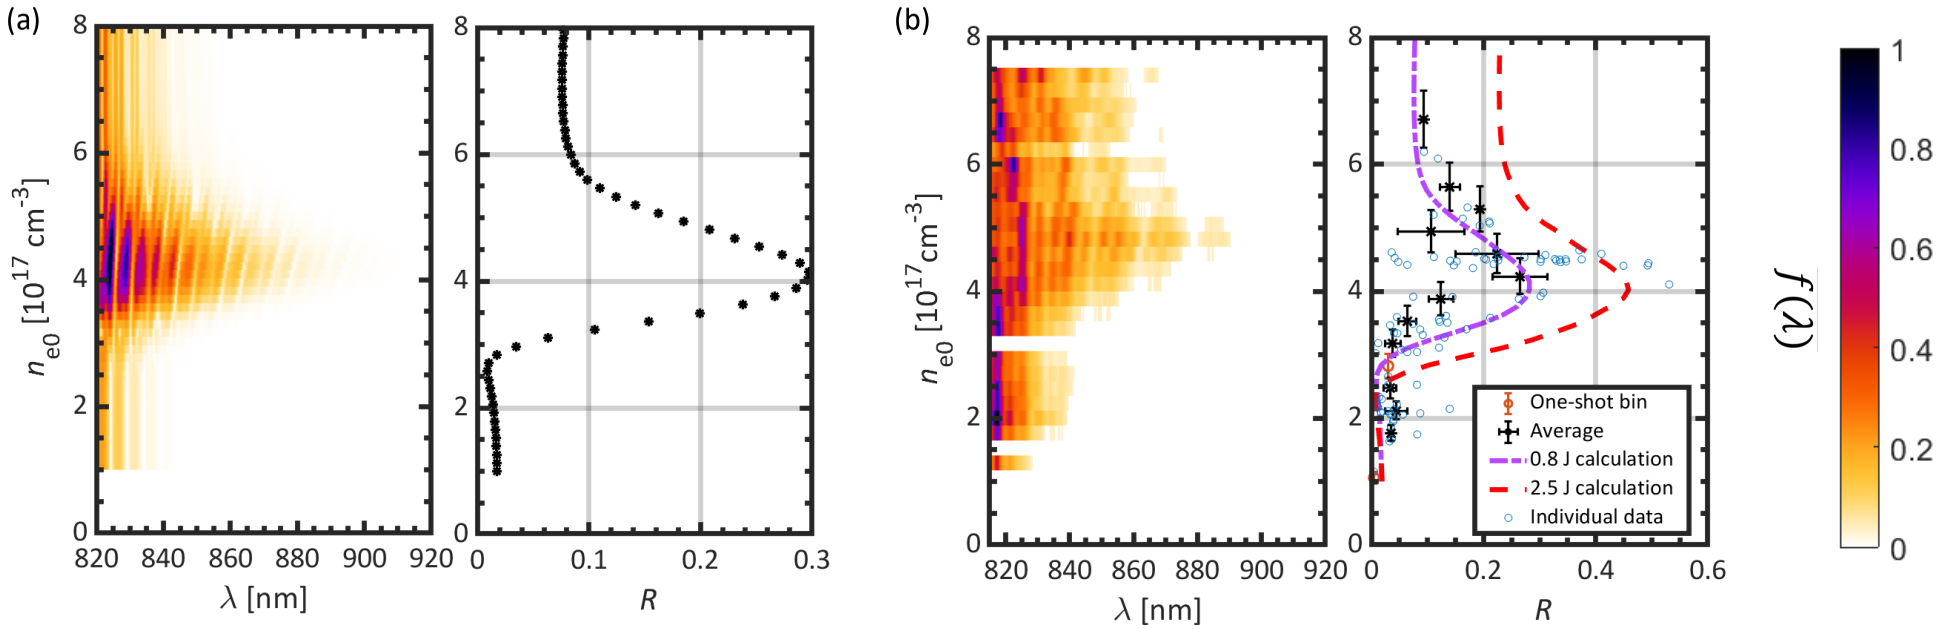


Figure 6.23: Transmitted spectra after 110 mm propagation binned against on-axis plasma density for the 170 fs pulse train from (a) a 2D-cylindrical fluid calculation and (b) the experimental data, as shown in Figure 6.20. The R-values for the individual data are included in (b), for all shots that were selected by the criteria in Section 6.4.1. The fluid calculation was performed for two different drive beam energies, 2.5 J and 800 mJ (see text for details). The calculations were performed by Johannes van de Wetering.

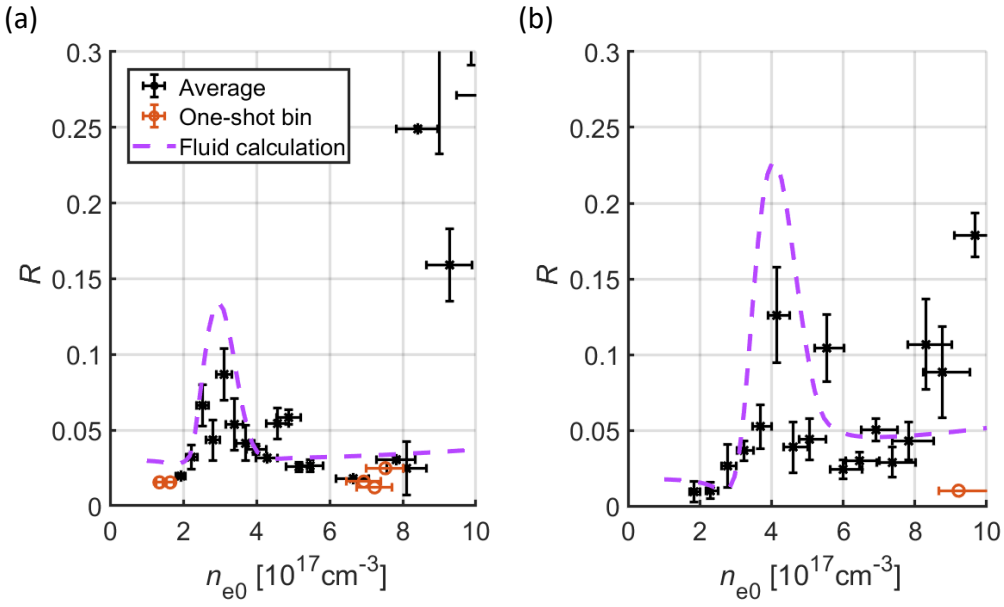


Figure 6.24: Redshift metric R resonance curves against on-axis density for the spectra after 70 mm propagation, for the case of (a) the 170 fs pulse train and (b) the 200 fs pulse train. RZ calculation results with 800 mJ effective energy overlaid. The calculations were performed by Johannes van de Wetering.

Similarly, the calculation can be run for the two other resonance curves, with channel length of 70 mm and pulse spacings of 170 fs and 200 fs as shown in Figure 6.24. The energy of the drive beam was again set to 800 mJ. The calculated curves shown in the figure are consistent with the experimental data in both cases. As noted before, the simulation cannot capture self-modulation effects because it assumes the laser envelope is fixed, and hence the high redshift values at high plasma density observed in experiment are not seen in the calculation.

6.7 Discussion

The experimental data has been found to be explained well using the 2D cylindrical fluid code. However, the physical processes involved in the propagation of the laser pulse trains within the plasma channel in the conditions of the experiment were likely more complicated than the analysis presented here.

The spatial jitter of the beam at the entrance to the plasma channel may have caused the pulse train to oscillate in centroid position and spot size as it propagated down the channel. As discussed in Chapter 2, there is the potential for the laser-hosing instability to arise in experiments which involve long propagation distances. Laser centroid oscillations can also occur as a result of the density gradients in the transverse profile of the plasma

channel, which acts to correct any laser centroid offset from the axis of the plasma channel [170]. Research into the interaction of these two effects on the pulse train propagation in a plasma channel is ongoing. The experimental results presented in this chapter show that, even with non-negligible spatial jitter of the pulse train beam, there are at least some shots for which the laser-hosing instability does not prohibit the guiding of pulse trains and generation of wakefields in this regime.

The energy lost by the drive beam during the propagation in the plasma channel occurs due to a combination of energy leakage from the lossy channel and energy transferred from the beam to the wakefield. The energy lost to the wakefield can be estimated by considering the energy density in the wake electric field, $(1/2)\epsilon_0 E_z^2$ for E_z the maximum electric field of the wakefield driven, which is generated over an area approximately equal to the drive beam transverse profile area πw_0^2 for w_0 the spot size of the laser beam driver. Hence, the energy lost by the driver per unit length $\approx (1/2)\epsilon_0 E_z^2 \pi w_0^2$ [165]. For $E_z = 10\,\mathrm{GV/m}$, and $w_0 = 40\,\mu\mathrm{m}$, the energy loss comes to 2.2 mJ per millimetre of propagation. Therefore $\approx 10\%$ of the initial energy in the drive beam is converted to energy in the wakefield over the 110 mm stage.

The energy leakage from the channel could be reduced by pre-conditioning the plasma channel, before the arrival of the drive beam pulse train. This could be achieved using a separate conditioning laser pulse that arrives after the channel-forming beam and before the drive beam, to deepen the channel before the drive beam arrives. The studies in Picksley *et al*. [21] suggest that to effectively precondition a HOFI plasma channel, an additional laser beam with $7\,\mathrm{mJ\,cm^{-1}}$ would be needed. Therefore, for the 110 mm plasma channel length used in this experiment, the extra laser beam would need a total energy of 77 mJ. This is about 3% of the 2.5 J energy used in the drive beam. Alternatively, the pulse train temporal intensity envelope could be designed such that the first pulse in the train has the highest intensity. Other methods of increasing the attenuation length of plasma channels, such as using a higher order Bessel mode to initialise the plasma column as discussed in Chapter 5, could also be considered.

One of the objectives of the experiment was to demonstrate parameters relevant to the accelerator stage of the P-MoPA scheme, as presented in [4]. In the table below, the parameters of the experiment in this chapter are compared to the parameters of the P-MoPA accelerator stage [4]. The parameters achieved on the experiment are indeed

	E [J]	w [μm]	N	τ [fs]	$n_{e,0}$ [10^{17} cm^{-3}]	L [mm]
Ref [4]	1.7	50	~ 10	220	2.5	100
This experiment	2.5	40	~ 10	170	4.3	110

similar to those proposed in [4]. One difference between the two studies not captured in the table is that the method of generating the pulse trains in the experiment in this chapter resulted in pulses with a larger temporal width than would be achieved using the P-MoPA technique (for the same pulse spacing). The narrower temporal width of the pulses in P-MoPA is expected to make the pulse trains less sensitive to instabilities as the pulses sample a smaller region of the plasma wave profile. Research into the effect of the temporal width of the pulses on the LPA properties is ongoing.

6.8 Conclusion

In this chapter, experiments to investigate the guiding of pulse trains and resonant excitation of plasma waves within the plasma channels were presented. These results constitute the first demonstration of guiding of pulse trains inside a plasma channel. It was not clear before these experiments were performed whether laser instabilities such as laser-hosing would result in deleterious effects on the laser and the wakefield over the propagation of the pulse train. These results show that, in the regime of this experiment, laser-plasma instabilities did not prohibit the guiding of pulse trains. Pulse trains with 2.5 J total energy, 1 ps pulse duration and $N \sim 10$ pulses were successfully guided over propagation distances of 110 mm, equivalent to 18 Rayleigh ranges. The HOFI plasma channels used to form the waveguide had on-axis densities of order 5×10^{17} cm^{-3}, and the pulse trains were guided with attenuation lengths on the order 100 mm.

The spectra of the transmitted beams demonstrate that wakefields were resonantly excited within the plasma channels. The redshift metric was found to be peaked at a plasma density close to the resonant density for the measured plasma train spacing as predicted by plasma theory. This peak disappeared when a picosecond-long, unmodulated laser pulse with the same total energy was guided in the same plasma channels. Moreover, the plasma density at which the peak occurred scaled with pulse train spacing in agreement with theoretical expectation. Simulations suggest that the wakefield driven reached amplitudes up to 10 GV/m.

The parameters of the pulse trains and plasma channels used in this experiment are directly relevant to the P-MoPA accelerator stage and hence this demonstration can be seen as an important step towards the realisation of kilohertz-repetition-rate, GeV-scale electron acceleration using LPAs.

Chapter 7

Conclusion and outlook

7.1 Summary

In this thesis, techniques useful for future high-repetition-rate laser-plasma accelerators were presented.

In Chapter 2, the historical progress in laser-plasma accelerators was summarised, with a particular focus on the development of resonant wakefield excitation techniques and plasma waveguides. The multi-pulse laser wakefield accelerator (MP-LWFA), was described as a potential route towards achieving kilohertz repetition-rate laser-plasma accelerators using currently available laser technology [3, 4]. In 2017, the Oxford Laser Plasma Accelerators group published the first experimental demonstration of an MP-LWFA [23], showing that wakefields of amplitude $\sim 1\%$ could be resonantly excited in a short gas cell by a train of laser pulses ($N \sim 7$) tuned to the plasma frequency. This experiment was limited to an accelerator length of a few millimetres due to the diffraction of the drive beam. In 2020, the Oxford group published theoretical and simulation results proposing a regime for a kilohertz, GeV-scale accelerator termed the plasma-modulated plasma accelerator (P-MoPA) [4]. The accelerator stage of P-MoPA requires a laser pulse train to resonantly excite a wakefield within a plasma channel over lengths of order 100 mm, something which had yet to be demonstrated. Potential laser-plasma instabilities that could arise over such long laser-plasma interaction lengths were briefly discussed in this chapter, motivating experimental verification of the P-MoPA scheme.

Chapter 2 also highlighted the need for plasma channels to guide high-intensity laser pulses for LPAs at low plasma density. A few different methods of achieving laser waveguides were outlined. The hydrodynamic optical-field-ionised (HOFI) plasma channel developed by the Oxford group [19, 82, 83] was shown to be a promising technology for

this application, with experiments demonstrating that HOFI plasma channels could guide single short laser pulses over propagation lengths of 100 mm [1]. Current constraints on the usable length of HOFI plasma channels due to energy losses of the guided beam were described, motivating investigations into alternative optical setups that could be used to enhance the attenuation length of a HOFI channel.

In Chapter 4, the optics used to form longitudinally-extended foci were presented, with a particular focus on the refractive axicon lens used in the experimental demonstration of HOFI plasma channels performed by the Oxford group in 2018 [20]. The chromatic and non-linear effects through the axicon lens on a beam with properties similar to that used on the experiment (central wavelength $\lambda_0 = 805\,\text{nm}$, bandwidth $\Delta\lambda_{\text{FWHM}} = 30\,\text{nm}$ and incident peak intensity $I_0 = 6 \times 10^{11}\,\text{W cm}^{-2}$) were studied using a numerical optical beam propagation code written by the author. The combined chromatic and non-linear effects were shown to have a significant impact on the peak intensity of the axicon focus, reducing it by almost 50% of the focus intensity for a non-dispersive system. Alternative focusing optics, such as thin phase plates, were discussed as a way to reduce chromatic effects.

The chapter also discussed the longitudinal non-uniformity of the axial intensity of the axicon focus and the potential implications for generating high-quality HOFI plasma channels on metre-length scales. The axilens was studied as an alternative optic for generating longitudinally-extended foci. Axilenses are capable of generating highly uniform axial intensity profiles at the cost of a small variation in the spot-size of the longitudinally-extended focus with longitudinal position. A parameter set of interest for a future implementation of an axilens for HOFI plasma channel formation was suggested. The designed axilens was simulated using the numerical optical beam propagation code and shown to achieve a highly uniform axial intensity over $\sim 300\,\text{mm}$ with spot-size $\sim 7\,\mu\text{m}$, and spot-size variations on the scale of 15%. The axilens focus profile also demonstrated a much sharper turn-on of intensity at the leading edge of the longitudinally-extended focus, which may have applications in the design of sharp density-transition injection of electrons into plasma channels [148].

In Chapter 5, methods to increase the power attenuation length of a HOFI plasma channel were studied. It was found that increasing the radius of the initially-ionised plasma column provided a promising route to increasing the strength of the blast wave

and hence the depth of the resulting HOFI plasma channel. Therefore, ways to increase the initial spot size of the longitudinally-extended focus were studied. It was shown that simply changing the available parameters of the standard focus (approach angle of the rays from the axicon, laser beam wavelength) is not a viable option due to the critical angle constraint of the plasma, which is fixed for a given plasma density. Therefore, alternative methods to increase the spot size of the focus are required. Random phase plates and higher-order Bessel mode foci (HOBMs) were both shown to achieve larger focal spot regions than the standard zeroth-order Bessel focus. The two methods were investigated using numerical optical beam propagation simulations and optical tests using a spatial light modulator to investigate the achievable focus properties. HOBMs with orders 1 and 2 were identified as the most promising foci for HOFI plasma channel generation as they achieved high optical energy efficiency and longitudinal uniformity. In hydrodynamic simulations with the Helios code, the HOBMs generated significantly deeper channels than the standard axicon focus, while maintaining low on-axis densities.

Finally, in Chapter 6, the results of an experimental campaign at the Astra-Gemini TA3 facility, Rutherford Appleton Laboratory, were presented, showing, for the first time, resonant excitation of a plasma wave by a channel-guided pulse train. The parameters of the experimental demonstration were relevant to the P-MoPA accelerator stage, with pulse trains of total duration $\sim$1 ps and energy 2.5 J, consisting of ~ 10 pulses. The plasma channels were generated with on-axis densities in the range $1-12\times10^{17}\,\mathrm{cm}^{-3}$ and used to guide the pulse trains over 110 mm. The pulse trains were successfully guided, with attenuation lengths on the order 100 mm.

Clear resonances in the wakefield amplitude as a function of gas pressure were observed, using the spectral shifts in the drive beam as the diagnostic of the wakefield excitation. The resonances appeared at the plasma densities expected for the two different pulse train spacings tested. Simulations suggest that the pulse train with spacing 170 fs drove a wakefield with amplitude up to 10 GV/m.

7.2 Outlook

In Chapter 5, the modelling of the hydrodynamic behaviour of the HOFI plasma channels after ionisation was limited due to the local-thermal-equilibrium (LTE) assumption em-

ployed in the Helios code. It was found that the inclusion of non-LTE physics in Helios would have resulted in one simulation taking on the scale of a year to complete. However, non-LTE processes such as radiative emission are expected to be important to model the processes occurring on the leading edge of the blast wave, and their inclusion has been shown to result in smoother and lower amplitude shock fronts compared to LTE models [150]. Further development of hydrodynamic simulation codes for the plasma conditions relevant to HOFI plasma channels will enable greater confidence in the modelling of the blast wave. Research into codes that can better model HOFI plasma channel generation is ongoing [164]. Ideally, a future model would include the optical field ionisation of the hydrogen molecules, the isotropisation and electron-electron thermalisation physics that occurs in the first few hundred femtoseconds after ionisation and the subsequent evolution of the blast wave in the electron, ion and neutral atom/molecule species, which occurs on the nanosecond timescale.

It is planned that the work on novel optical schemes for HOFI plasma channel formation will be investigated by members of the Oxford group in experiments in the new Oxford Plasma Accelerator Laboratory. Phase plate axilenses, following the description in Chapter 4, will be investigated to study whether they can successfully generate HOFI plasma channels for guiding laser pulses and whether any difficulties arise in practice, for example in the manufacturing or alignment of the axilenses. The higher-order Bessel modes with orders 1 and 2 will be generated experimentally using a spiral phase plate and the quality of the HOFI plasma channels they generate will be compared to the standard zeroth-order Bessel focus. It is hoped that the HOBMs will result in significantly improved energy throughput of the channels, providing a simple way to enhance the waveguide in future experiments.

It is also anticipated that the methods developed for this thesis, in particular the numerical optical beam propagation code, could be applied to other novel optical schemes of interest for LPAs. For example, there is interest in generating curved HOFI plasma channels, which could be used to couple-out the spent drive laser pulse at the end of a plasma accelerator stage, and to couple-in a fresh drive pulse into a subsequent stage, while allowing the accelerated bunch to propagate in a straight line. The phase required on the incident beam to generate a curved longitudinally-extended focus could be investigated using the ASM code.

One limitation of the experiment described in Chapter 6 was the severity of the drive beam pointing fluctuations and the impact of these fluctuations on the coupling efficiency of the drive laser into the plasma channel. Future efforts to improve laser-pointing stability will be needed to increase the average energy efficiency of the scheme. Moreover, the delays to the commissioning of the transverse probe beam for the experiment meant that the plasma channel properties could not be well characterised. Future experiments incorporating a plasma channel probe could be used to investigate the degree to which the pulse train induces conditioning of the HOFI plasma channel in the regime of interest.

This thesis has presented experimental investigations into the parameters relevant to the accelerator stage of P-MoPA. Methods will now need to be developed to achieve controlled injection and acceleration of electrons in the plasma wave of a P-MoPA. Further experimental demonstrations will be needed for the modulator stage of P-MoPA, to investigate the properties of the pulse trains that can be generated by this method. Finally, the operation of the P-MoPA at kilohertz repetition rates will need to be studied, using the thin-disk laser systems it was designed for. These studies will be attempted by the Oxford Laser Plasma Accelerators group in the future. It is hoped that the work in this thesis contributes towards the achievement of the P-MoPA scheme and the ambition to generate GeV-scale electron beams at kilohertz repetition rates for applications.

Appendix A

A.1 Method of Stationary Phase

The method of stationary phase can be used to approximate the solution to integrals of the form;

$$f = \int \exp(\mathrm{i}G(\omega))F(\omega)d\omega \tag{A.1}$$

where $G(\omega)$ is a large valued function and $F(\omega)$ is slowly varying [123, 124]. In this case, the exponential oscillates rapidly over most of the domain and hence the integral in these regions averages to zero. The dominant contributions to f arise from the stationary points of $G(\omega)$, hence by evaluating the integral about these stationary points, we can get a first approximation to f over the whole domain.

First we expand $G(\omega)$ as a Taylor series about the stationary points of $G(\omega)$, which are labelled ω_s;

$$G(\omega) = G(\omega_s) + \frac{1}{2}G_s'' \times (\omega - \omega_s)^2 + ... \tag{A.2}$$

where $G_s'' = \frac{d^2G}{d\omega^2}(\omega_s)$. For the case of one stationary point only (can be generalised to multiple stationary points later);

$$f_s = F(\omega_s)\exp(\mathrm{i}G(\omega_s))\int_{-\infty}^{\infty}\exp\big(\mathrm{i}G_s''(\omega-\omega_s)^2/2\big)d\omega. \tag{A.3}$$

To solve this integral, consider the change of variable $\pi t^2/2\mathsf{sign}(G_s'') = G_s''(\omega-\omega_s)^2/2$, resulting in;

$$f_s = 2F(\omega_s)\exp(\mathrm{i}G(\omega_s))\sqrt{\frac{\pi}{|G_s''|}}\int_0^{\infty}\exp\big(\mathrm{i}\mathsf{sign}(G_s'')\pi t^2/2\big)dt \tag{A.4}$$

$$= F(\omega_s)\exp(\mathrm{i}G(\omega_s))\sqrt{\frac{\pi}{|G_s''|}}(1+\mathsf{sign}(G_s'')\mathrm{i}) \tag{A.5}$$

$$= F(\omega_s)\exp(\mathrm{i}G(\omega_s))\sqrt{\frac{2\pi}{|G_s''|}}\exp(\mathrm{i}\mathsf{sign}(G_s'')\pi/4). \tag{A.6}$$

Summing over contributions f_s for all stationary points ω_s gives the solution for f. This solution assumes that all stationary points are non-degenerate and real valued.

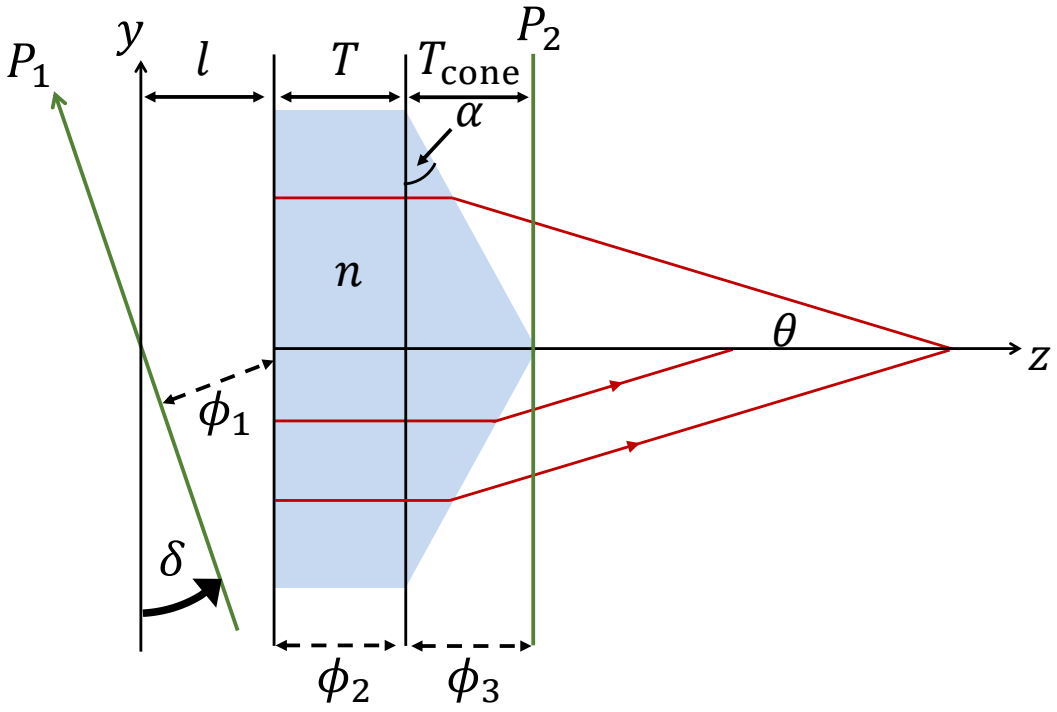


Figure A.1: Diagram of beam propagating at a tilt angle δ to the back plane of an axicon lens, with thickness T, refractive index n, and base angle α.

A.2 Phase accumulation through tilted axicon lens

In Chapter 4, it was claimed that for small tilt angles δ about the x-axis, the phase accumulated by a wavefront with angular wavenumber k through a tilted axicon is equivalent to the sum of a tilt phase, $\phi_{\text{tilt}} = ky\tan\delta$, and the standard axicon phase, $\phi_{\text{ax}}(\rho) = knT + knT_{\text{cone}} - k\alpha\rho(n-1)$ for axicon with substrate thickness T, base angle α and refractive index n. The parameter T_{cone} introduced here is the maximum thickness of the conical section of the axicon from the back of the cone to the tip, as drawn in the diagram in Figure A.1. In this section, the claim above is studied using a geometric optics analysis. We first consider the 2D setup in the (y, z) plane defined by $x = 0$ through the centre of the axicon, and generalise to the full 3D setup at the end.

Consider a tilted wavefront incident on the axicon, with phase contour along a plane at angle δ to the plane of the back of the axicon. This plane is labelled P_1 in Figure A.1. We want to calculate the phase along the beam at the point of the axicon, along plane P_2, which is parallel to the back face of the axicon. We split up the phase accumulation into three components, ϕ_1 accumulated between P_1 and the back of the axicon, ϕ_2 through the substrate of the axicon, and ϕ_3 through the conical section of the axicon to P_2.

The phase is calculated as $\phi_1(y) = k(l + y\tan\delta)$ where l is an arbitrary length from P_1 to the back of the axicon substrate. As the phase kl is a constant with respect to y, we choose to set $l = 0$ for the rest of the derivation, such that $\phi_1(y) = ky\tan\delta$.

At the interface between the vacuum and the axicon substrate, the wavefront is refracted so that it now propagates at an angle δ/n in the glass. This can be calculated using Snell's law. Therefore, $\phi_2 = nkT/\cos(\delta/n)$.

Similarly, the phase accumulation within the cone is modified by the angle of the wavefront. The wavefront in the glass propagates at an angle δ/n, but in the vacuum on exit of the axicon, the wavefront is again refracted to angle δ. In the standard non-tilted case, the small deviation of the rays at the exit of the axicon owing to the axicon cone is ignored when calculating the phase accumulation from the conical surface to the plane P_2. Therefore, we make this same approximation here. We then find;

$$\phi_3 = \frac{kn(T_{\text{cone}} - \alpha y')}{\cos(\delta/n)} + \frac{k\alpha y''}{\cos(\delta)} \tag{A.7}$$

where $y' = y + T\tan(\delta/n)$ is the new transverse position that a ray originating at position y on the back face of the axicon would be after propagating at angle δ/n through the substrate. Similarly, $y'' = y' + (T_{\text{cone}} - \alpha y')\tan(\delta/n)$ is the position on the conical surface for the ray originating at position y.

The phase accumulated from P_1 to P_2 is then calculated as the sum of the three phase contributions;

$$\phi_{\text{tot}}(y) = ky\tan\delta + nkT/\cos(\delta/n) + \frac{kn(T_{\text{cone}} - \alpha y')}{\cos(\delta/n)} + \frac{k\alpha y''}{\cos(\delta)}. \tag{A.8}$$

In the limit of small angles δ, $\cos\delta = 1 - \delta^2/2 + O(\delta^4)$ and $\tan(\delta) = \delta + O(\delta^3)$. Ignoring terms in δ^2, α^2 and $\alpha\delta$ and above (both angles small), gives;

$$\phi_{\text{tot}}(y) = ky\delta + nkT + kn(T_{\text{cone}} - \alpha(y + T\frac{\delta}{n}) + k\alpha(y + \frac{\delta}{n}(T + T_{\text{cone}} - \alpha y)) \tag{A.9}$$

$$= ky\delta + nkT + nkT_{\text{cone}} + k(1-n)\alpha y \tag{A.10}$$

$$= \phi_{\text{tilt}}(y) + \phi_{\text{ax}}(y) \tag{A.11}$$

in the 2D plane at $x = 0$. This is equivalent to the sum of the tilt phase and the axicon phase in the 2D plane. Extending to 3D, the tilt phase is unchanged and the axicon phase becomes the full 3D axicon phase by setting $y \to \rho$.

Bibliography

[1] Picksley A, Alejo A, Cowley J, Bourgeois N, Corner L, Feder L, Holloway J, Jones H, Jonnerby J, Milchberg H M, Reid L R, Ross A J, Walczak R and Hooker S M 2020 *Phys. Rev. Accel. Beams* **23**(8) 081303

[2] Alejo A, Cowley J, Picksley A, Walczak R and Hooker S M 2022 *Phys. Rev. Accel. Beams* **25**(1) 011301 URL `https://link.aps.org/doi/10.1103/PhysRevAccelBeams.25.011301`

[3] Hooker S M, Bartolini R, Mangles S P D, Tünnermann A, Corner L, Limpert J, Seryi A and Walczak R 2014 *Journal of Physics B: Atomic, Molecular and Optical Physics* **47** 234003

[4] Jakobsson O, Hooker S M and Walczak R 2021 *Phys. Rev. Lett.* **127**(18) 184801 URL `https://link.aps.org/doi/10.1103/PhysRevLett.127.184801`

[5] Lopes A and Perrey M L 2022 FAQ-LHC The guide URL `https://cds.cern.ch/record/2809109`

[6] Bailey R (ed) 2017 *CAS - CERN Accelerator School: Accelerators for Medical Applications: Vösendorf, Austria 26 May - 5 Jun 2015. Proceedings, CAS-CERN Accelerator School: Accelerators for Medical Applications. Vösendorf, Austria, May 26 - June 5, 2015. CAS - CERN Accelerator School: Accelerators for Medical Applications* CERN (Geneva: CERN) 18 lectures, 346 pages, published as CERN Yellow Report a this url: https://e-publishing.cern.ch/index.php/CYRSP/issue/view/33 URL `https://cds.cern.ch/record/2271793`

[7] CERN 1998 *CAS - CERN Accelerator School : Synchrotron Radiation and Free Electron Lasers: Grenoble, France 22 - 27 Apr 1996. CAS - CERN Accelerator*

School : Synchrotron Radiation and Free Electron Lasers (Geneva: CERN) URL https://cds.cern.ch/record/362891

[8] CERN 2018 *CAS - CERN Accelerator School: Free Electron Lasers and Energy Recovery Linacs: CERN, Geneva, Switzerland 31 May - 10 Jun 2016. CAS - CERN Accelerator School: Free Electron Lasers and Energy Recovery Linacs* (Geneva: CERN) URL https://cds.cern.ch/record/2312429

[9] Zha H and Grudiev A 2016 New clic-g structure design Tech. rep. CERN Geneva URL http://cds.cern.ch/record/2157079

[10] EU X-FEL Facility URL https://www.xfel.eu/facility/accelerator/index_eng.html#:~:text=The%20European%20XFEL%20accelerator%20currently,over%20a%201.7%20km%20length.

[11] Esarey E, Schroeder C B and Leemans W P 2009 *Rev. Mod. Phys.* **81**(3) 1229–1285

[12] Tajima T and Dawson J M 1979 *Phys. Rev. Lett.* **43**(4) 267–270 URL https://link.aps.org/doi/10.1103/PhysRevLett.43.267

[13] Mangles S P D, Murphy C D, Najmudin Z, Thomas A G R, Collier J L, Dangor A E, Divall E J, Foster P S, Gallacher J G, Hooker C J, Jaroszynski D A, Langley A J, Mori W B, Norreys P A, Tsung F S, Viskup R, Walton B R and Krushelnick K 2004 *Nature* **431** 535–538 ISSN 1476-4687

[14] Geddes C G R, Toth C, van Tilborg J, Esarey E, Schroeder C B, Bruhwiler D, Nieter C, Cary J and Leemans W P 2004 *Nature* **431** 538–541 ISSN 1476-4687

[15] Faure J, Glinec Y, Pukhov A, Kiselev S, Gordienko S, Lefebvre E, Rousseau J P, Burgy F and Malka V 2004 *Nature* **431** 541–544 ISSN 1476-4687

[16] Leemans W P, Nagler B, Gonsalves A J, Tóth C, Nakamura K, Geddes C G R, Esarey E, Schroeder C B and Hooker S M 2006 *Nature Physics* **2** 696–699 ISSN 1745-2481

[17] Gonsalves A J, Nakamura K, Daniels J, Benedetti C, Pieronek C, de Raadt T C H, Steinke S, Bin J H, Bulanov S S, van Tilborg J, Geddes C G R, Schroeder C B,

Tóth C, Esarey E, Swanson K, Fan-Chiang L, Bagdasarov G, Bobrova N, Gasilov V, Korn G, Sasorov P and Leemans W P 2019 *Phys. Rev. Lett.* **122**(8) 084801

[18] Aniculaesei C, Ha T, Yoffe S, McCary E, Spinks M M, Quevedo H J, Labun L, Labun O Z, Sain R, Hannasch A, Zgadzaj R, Pagano I, Franco-Altamirano J A, Ringuette M L, Gaul E, Luedtke S V, Tiwari G, Ersfeld B, Brunetti E, Ruhl H, Ditmire T, Bruce S, Donovan M E, Jaroszynski D A, Downer M C and Hegelich B M 2023 High-charge 10 gev electron acceleration in a 10 cm nanoparticle-assisted hybrid wakefield accelerator (*Preprint* 2207.11492)

[19] Shalloo R J, Arran C, Corner L, Holloway J, Jonnerby J, Walczak R, Milchberg H M and Hooker S M 2018 *Phys. Rev. E* **97**(5) 053203

[20] Shalloo R J, Arran C, Picksley A, von Boetticher A, Corner L, Holloway J, Hine G, Jonnerby J, Milchberg H M, Thornton C, Walczak R and Hooker S M 2019 *Phys. Rev. Accel. Beams* **22**(4) 041302

[21] Picksley A, Alejo A, Shalloo R J, Arran C, von Boetticher A, Corner L, Holloway J A, Jonnerby J, Jakobsson O, Thornton C, Walczak R and Hooker S M 2020 *Phys. Rev. E* **102**(5) 053201

[22] Nakajima K, Yuan J, Chen L and Sheng Z 2015 *Applied Sciences* **5** 1–20 ISSN 2076-3417 URL https://www.mdpi.com/2076-3417/5/1/1

[23] Cowley J, Thornton C, Arran C, Shalloo R J, Corner L, Cheung G, Gregory C D, Mangles S P D, Matlis N H, Symes D R, Walczak R and Hooker S M 2017 *Phys. Rev. Lett.* **119**(4) 044802

[24] Ross A J, Alejo A, von Boetticher A, Cowley J, Holloway J, Jonnerby J, Picksley A, Walczak R and Hooker S M 2020 Numerical modelling of chromatic effects on axicon-focused beams used to generate HOFI plasma channels *Journal of Physics: Conference Series* vol 1596 (IOP Publishing) p 012049

[25] Sanginés de Castro R, Sobral H, Sánchez-Aké C and Villagrán-Muniz M 2006 *Physics Letters A* **357** 351–354 ISSN 0375-9601

[26] Schwerdtfeger P and Nagle J K 2019 *Molecular Physics* **117** 1200–1225 (*Preprint* https://doi.org/10.1080/00268976.2018.1535143)

[27] Chen F F 1984 *Introduction to Plasma Physics and Controlled Fusion* 2nd ed (Plenum Press, New York) ISBN 978-3-319-22308-7 rev. ed. of: Introduction to plasma physics. 1974

[28] Hooker S M 2013 *Nature Photonics* **7** 775–782 ISSN 1749-4893 URL https://doi.org/10.1038/nphoton.2013.234

[29] Pukhov A and Meyer-ter Vehn J 2002 *Applied Physics B* **74** 355–361 ISSN 1432-0649 URL https://doi.org/10.1007/s003400200795

[30] Modena A, Najmudin Z, Dangor A E, Clayton C E, Marsh K A, Joshi C, Malka V, Darrow C B, Danson C, Neely D and Walsh F N 1995 *Nature* **377** 606–608 ISSN 1476-4687 URL https://doi.org/10.1038/377606a0

[31] Bulanov S, Naumova N, Pegoraro F and Sakai J 1998 *Phys. Rev. E* **58**(5) R5257–R5260 URL https://link.aps.org/doi/10.1103/PhysRevE.58.R5257

[32] Chien T Y, Chang C L, Lee C H, Lin J Y, Wang J and Chen S Y 2005 *Phys. Rev. Lett.* **94**(11) 115003 URL https://link.aps.org/doi/10.1103/PhysRevLett.94.115003

[33] Pollock B B, Clayton C E, Ralph J E, Albert F, Davidson A, Divol L, Filip C, Glenzer S H, Herpoldt K, Lu W, Marsh K A, Meinecke J, Mori W B, Pak A, Rensink T C, Ross J S, Shaw J, Tynan G R, Joshi C and Froula D H 2011 *Phys. Rev. Lett.* **107**(4) 045001 URL https://link.aps.org/doi/10.1103/PhysRevLett.107.045001

[34] Rowlands-Rees T P, Kamperidis C, Kneip S, Gonsalves A J, Mangles S P D, Gallacher J G, Brunetti E, Ibbotson T, Murphy C D, Foster P S, Streeter M J V, Budde F, Norreys P A, Jaroszynski D A, Krushelnick K, Najmudin Z and Hooker S M 2008 *Phys. Rev. Lett.* **100**(10) 105005 URL https://link.aps.org/doi/10.1103/PhysRevLett.100.105005

[35] Pak A, Marsh K A, Martins S F, Lu W, Mori W B and Joshi C 2010 *Phys. Rev. Lett.* **104**(2) 025003 URL https://link.aps.org/doi/10.1103/PhysRevLett.104.025003

[36] McGuffey C, Thomas A G R, Schumaker W, Matsuoka T, Chvykov V, Dollar F J, Kalintchenko G, Yanovsky V, Maksimchuk A, Krushelnick K, Bychenkov V Y, Glazyrin I V and Karpeev A V 2010 *Phys. Rev. Lett.* **104**(2) 025004 URL https://link.aps.org/doi/10.1103/PhysRevLett.104.025004

[37] Liu J S, Xia C Q, Wang W T, Lu H Y, Wang C, Deng A H, Li W T, Zhang H, Liang X Y, Leng Y X, Lu X M, Wang C, Wang J Z, Nakajima K, Li R X and Xu Z Z 2011 *Phys. Rev. Lett.* **107**(3) 035001 URL https://link.aps.org/doi/10.1103/PhysRevLett.107.035001

[38] Esarey E, Hubbard R F, Leemans W P, Ting A and Sprangle P 1997 *Phys. Rev. Lett.* **79**(14) 2682–2685 URL https://link.aps.org/doi/10.1103/PhysRevLett.79.2682

[39] Umstadter D, Kim J K and Dodd E 1996 *Phys. Rev. Lett.* **76**(12) 2073–2076 URL https://link.aps.org/doi/10.1103/PhysRevLett.76.2073

[40] Palastro J P, Shaw J L, Franke P, Ramsey D, Simpson T T and Froula D H 2020 *Phys. Rev. Lett.* **124**(13) 134802 URL https://link.aps.org/doi/10.1103/PhysRevLett.124.134802

[41] Steinke S, van Tilborg J, Benedetti C, Geddes C G R, Schroeder C B, Daniels J, Swanson K K, Gonsalves A J, Nakamura K, Matlis N H, Shaw B H, Esarey E and Leemans W P 2016 *Nature* **530** 190–193 ISSN 1476-4687 URL https://doi.org/10.1038/nature16525

[42] Foerster F M, Döpp A, Haberstroh F, Grafenstein K v, Campbell D, Chang Y Y, Corde S, Couperus Cabadağ J P, Debus A, Gilljohann M F, Habib A F, Heinemann T, Hidding B, Irman A, Irshad F, Knetsch A, Kononenko O, Martinez de la Ossa A, Nutter A, Pausch R, Schilling G, Schletter A, Schöbel S, Schramm U, Travac E, Ufer P and Karsch S 2022 *Phys. Rev. X* **12**(4) 041016 URL https://link.aps.org/doi/10.1103/PhysRevX.12.041016

[43] Strickland D and Mourou G 1985 *Optics Communications* **56** 219–221 ISSN 0030-4018 URL https://www.sciencedirect.com/science/article/pii/0030401885901208

[44] Rosenbluth M N and Liu C S 1972 *Phys. Rev. Lett.* **29**(11) 701–705 URL https://link.aps.org/doi/10.1103/PhysRevLett.29.701

[45] Cohen B I, Kaufman A N and Watson K M 1972 *Phys. Rev. Lett.* **29**(9) 581–584 URL https://link.aps.org/doi/10.1103/PhysRevLett.29.581

[46] Dangor A E, Dymoke-Bradshaw A K L and Dyson A E 1990 *Physica Scripta* **1990** 107 URL https://dx.doi.org/10.1088/0031-8949/1990/T30/015

[47] Clayton C E, Joshi C, Darrow C and Umstadter D 1985 *Phys. Rev. Lett.* **54**(21) 2343–2346 URL https://link.aps.org/doi/10.1103/PhysRevLett.54.2343

[48] Amiranoff F, Laberge M, Marquès J R, Moulin F, Fabre E, Cros B, Matthieussent G, Benkheiri P, Jacquet F, Meyer J, Miné P, Stenz C and Mora P 1992 *Phys. Rev. Lett.* **68**(25) 3710–3713 URL https://link.aps.org/doi/10.1103/PhysRevLett.68.3710

[49] Clayton C E, Marsh K A, Dyson A, Everett M, Lal A, Leemans W P, Williams R and Joshi C 1993 *Phys. Rev. Lett.* **70**(1) 37–40 URL https://link.aps.org/doi/10.1103/PhysRevLett.70.37

[50] Deutsch M, Meerson B and Golub J E 1991 *Physics of Fluids B: Plasma Physics* **3** 1773–1780 (*Preprint* https://doi.org/10.1063/1.859992) URL https://doi.org/10.1063/1.859992

[51] Andreev N, Gorbunov L, Kirsanov V, Pogosova A, Ramazashvili R and Parsons D 1992 *JETP letters* **55** 571–576

[52] Esarey E, Krall J and Sprangle P 1994 *Phys. Rev. Lett.* **72**(18) 2887–2890 URL https://link.aps.org/doi/10.1103/PhysRevLett.72.2887

[53] Nakajima K, Fisher D, Kawakubo T, Nakanishi H, Ogata A, Kato Y, Kitagawa Y, Kodama R, Mima K, Shiraga H, Suzuki K, Yamakawa K, Zhang T, Sakawa Y,

Shoji T, Nishida Y, Yugami N, Downer M and Tajima T 1995 *Phys. Rev. Lett.* **74**(22) 4428–4431 URL https://link.aps.org/doi/10.1103/PhysRevLett.74.4428

[54] Coverdale C, Darrow C, Decker C, Mori W, Tzeng K, Marsh K, Clayton C and Joshi C 1995 *Physical Review Letters* **74** 4659

[55] King P M, Miller K, Lemos N, Shaw J L, Kraus B F, Thibodeau M, Hegelich B M, Hinojosa J, Michel P, Joshi C, Marsh K A, Mori W, Pak A, Thomas A G R and Albert F 2021 *Phys. Rev. Accel. Beams* **24**(1) 011302 URL https://link.aps.org/doi/10.1103/PhysRevAccelBeams.24.011302

[56] Arran C 2018 *Techniques for high repetition rate laser wakefield acceleration* Ph.D. thesis University of Oxford

[57] Nakajima K 1992 *Phys. Rev. A* **45**(2) 1149–1156 URL https://link.aps.org/doi/10.1103/PhysRevA.45.1149

[58] Berezhiani V I and Murusidze I G 1992 *Physica Scripta* **45** 87 URL https://dx.doi.org/10.1088/0031-8949/45/2/007

[59] Umstadter D, Esarey E and Kim J 1994 *Phys. Rev. Lett.* **72**(8) 1224–1227 URL https://link.aps.org/doi/10.1103/PhysRevLett.72.1224

[60] Johnson D A, Cairns R A, Bingham R and de Angelis U 1994 *Physica Scripta* **1994** 77 URL https://dx.doi.org/10.1088/0031-8949/1994/T52/013

[61] Dalla S and Lontano M 1994 *Plasma Physics and Controlled Fusion* **36** 1987 URL https://dx.doi.org/10.1088/0741-3335/36/12/003

[62] Bonnaud G, Teychenné D and Bobin J L 1994 *Phys. Rev. E* **50**(1) R36–R39 URL https://link.aps.org/doi/10.1103/PhysRevE.50.R36

[63] Cairns R, Johnson D and Bingham R 1995 *Laser and Particle Beams* **13** 451–458

[64] Umstadter D, Kim J, Esarey E, Dodd E and Neubert T 1995 *Phys. Rev. E* **51**(4) 3484–3497 URL https://link.aps.org/doi/10.1103/PhysRevE.51.3484

[65] Jonnerby J 2021 *Multi-pulse laser wakefield acceleration* Ph.D. thesis University of Oxford

[66] von Boetticher A, Walczak R and Hooker S M 2023 *Phys. Rev. E* **107**(2) L023201 URL https://link.aps.org/doi/10.1103/PhysRevE.107.L023201

[67] Streeter M J V, Ma Y, Kettle B, Dann S J D, Gerstmayr E, Albert F, Bourgeois N, Cipiccia S, Cole J M, Gallardo González I, Hussein A E, Jaroszynski D A, Falk K, Krushelnick K, Lemos N, Lopes N C, Lumsdon C, Lundh O, Mangles S P D, Najmudin Z, Rajeev P P, Sandberg R, Shahzad M, Smid M, Spesyvtsev R, Symes D R, Vieux G and Thomas A G R 2022 *Phys. Rev. Accel. Beams* **25**(10) 101302 URL https://link.aps.org/doi/10.1103/PhysRevAccelBeams.25.101302

[68] Chappell J 2021 *Experimental study of long timescale plasma wakefield evolution* Ph.D. thesis University College London

[69] van de Wetering J J, Hooker S M and Walczak R 2023 Stability of the modulator in a plasma-modulated plasma accelerator (*Preprint* 2303.14032)

[70] Spence D J and Hooker S M 2000 *Phys. Rev. E* **63**(1) 015401

[71] Ibbotson T P A, Bourgeois N, Rowlands-Rees T P, Caballero L S, Bajlekov S I, Walker P A, Kneip S, Mangles S P D, Nagel S R, Palmer C A J, Delerue N, Doucas G, Urner D, Chekhlov O, Clarke R J, Divall E, Ertel K, Foster P S, Hawkes S J, Hooker C J, Parry B, Rajeev P P, Streeter M J V and Hooker S M 2010 *Phys. Rev. ST Accel. Beams* **13**(3) 031301

[72] Leemans W P, Gonsalves A J, Mao H S, Nakamura K, Benedetti C, Schroeder C B, Tóth C, Daniels J, Mittelberger D E, Bulanov S S, Vay J L, Geddes C G R and Esarey E 2014 *Phys. Rev. Lett.* **113**(24) 245002

[73] Butler A, Spence D J and Hooker S M 2002 *Phys. Rev. Lett.* **89**(18) 185003 URL https://link.aps.org/doi/10.1103/PhysRevLett.89.185003

[74] Gonsalves A J, Liu F, Bobrova N A, Sasorov P V, Pieronek C, Daniels J, Antipov S, Butler J E, Bulanov S S, Waldron W L, Mittelberger D E and Leemans W P 2016 *Journal of Applied Physics* **119** 033302 (*Preprint* https://doi.org/10.1063/1.4940121) URL https://doi.org/10.1063/1.4940121

[75] Pieronek C V, Gonsalves A J, Benedetti C, Bulanov S S, van Tilborg J, Bin J H, Swanson K K, Daniels J, Bagdasarov G A, Bobrova N A, Gasilov V A, Korn G, Sasorov P V, Geddes C G R, Schroeder C B, Leemans W P and Esarey E 2020 *Physics of Plasmas* **27** 093101 (*Preprint* https://doi.org/10.1063/5.0014961) URL https://doi.org/10.1063/5.0014961

[76] Durfee C G and Milchberg H M 1993 *Phys. Rev. Lett.* **71**(15) 2409–2412

[77] Durfee C, Lynch J and Milchberg H 1994 *Optics letters* **19** 1937–1939

[78] Durfee C G, Clark T R and Milchberg H M 1996 *J. Opt. Soc. Am. B* **13** 59–67

[79] Clark T R and Milchberg H M 1997 *Phys. Rev. Lett.* **78**(12) 2373–2376

[80] Clark T and Milchberg H 2000 *Physical Review E* **61** 1954

[81] Nikitin S P, Antonsen T M, Clark T R, Li Y and Milchberg H M 1997 *Opt. Lett.* **22** 1787–1789 URL https://opg.optica.org/ol/abstract.cfm?URI=ol-22-23-1787

[82] Hooker S 2016 Simulations of hydrodynamic optical-field-ionised plasma channels Advanced Accelerator Concepts Workshop URL https://aip.scitation.org/doi/abs/10.1063/1.4975840

[83] Shalloo R 2018 *Hydrodynamic optical-field-ionized plasma waveguides for laser plasma accelerators* Ph.D. thesis University of Oxford

[84] Miao B, Feder L, Shrock J E, Goffin A and Milchberg H M 2020 *Phys. Rev. Lett.* **125**(7) 074801

[85] Mori W 1997 *IEEE Journal of Quantum Electronics* **33** 1942–1953

[86] Sprangle P, Krall J and Esarey E 1994 *Phys. Rev. Lett.* **73**(26) 3544–3547 URL https://link.aps.org/doi/10.1103/PhysRevLett.73.3544

[87] Ren C and Mori W B 2001 *Physics of Plasmas* **8** 3118–3119 (*Preprint* https://doi.org/10.1063/1.1373419) URL https://doi.org/10.1063/1.1373419

[88] Hooker S M and Webb C 2010 *Laser Physics* Volume 9 of Oxford Master Series in Physics (Oxford University Press) ISBN 0198506910

[89] Hooker S M and Webb C 2010 *Laser Physics* (Oxford University Press)

[90] Esarey E, Ting A and Sprangle P 1990 *Phys. Rev. A* **42**(6) 3526–3531 URL https://link.aps.org/doi/10.1103/PhysRevA.42.3526

[91] Shalloo R, Corner L, Arran C, Cowley J, Cheung G, Thornton C, Walczak R and Hooker S 2016 *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* **829** 383–385 ISSN 0168-9002 2nd European Advanced Accelerator Concepts Workshop - EAAC 2015 URL https://www.sciencedirect.com/science/article/pii/S0168900216002114

[92] Figueira G, Braga L, Ahmed S, Boyle A, Galimberti M, Galletti M and Oliveira P 2019 *J. Opt. Soc. Am. B* **36** 366–373 URL https://opg.optica.org/josab/abstract.cfm?URI=josab-36-2-366

[93] Kruer W 2019 *The physics of laser plasma interactions* (crc Press)

[94] Keldysh L *et al.* 1965 *Sov. Phys. JETP* **20** 1307–1314

[95] Popruzhenko S 2014 *Journal of Physics B: Atomic, Molecular and Optical Physics* **47** 204001

[96] Ammosov M V, Delone N B and Krainov V P 1986 Tunnel ionization of complex atoms and atomic ions in electromagnetic field *High intensity laser processes* vol 664 (International Society for Optics and Photonics) pp 138–141

[97] Corkum P, Burnett N and Brunel F 1989 *Physical review letters* **62** 1259

[98] Herzberg G 1969 *Physical review letters* **23** 1081

[99] Emmanouilidou A and Staudte A 2009 *Physical Review A* **80** 053415

[100] Emmanouilidou A, Lazarou C, Staudte A and Eichmann U 2012 *Physical Review A* **85** 011402

[101] Spanner M, Gräfe S, Chelkowski S, Pavičić D, Meckel M, Zeidler D, Bardon A B, Ulrich B, Bandrauk A D, Villeneuve D M, Dörner R, Corkum P B and Staudte

A 2012 *Journal of Physics B: Atomic, Molecular and Optical Physics* **45** 194011 URL https://dx.doi.org/10.1088/0953-4075/45/19/194011

[102] Brabec T and Kapteyn H 2008 *Strong field laser physics* vol 1 (Springer)

[103] Arber T, Bennett K, Brady C, Lawrence-Douglas A, Ramsay M, Sircombe N, Gillies P, Evans R, Schmitz H, Bell A *et al.* 2015 *Plasma Physics and Controlled Fusion* **57** 113001

[104] Lawrence-Douglas A 2013 *Ionisation effects for laser-plasma interactions by particle-in-cell code* Ph.D. thesis University of Warwick

[105] Miladinović T B and Petrović V M 2015 *Journal of Russian Laser Research* **36** 312–319

[106] Burnett N and Corkum P B 1989 *JOSA B* **6** 1195–1199

[107] Spitzer L 2006 *Physics of fully ionized gases* 2nd ed Dover books on physics. (Mineola, NY: Dover Publications) ISBN 0486449823 (pbk) rev. ed. of: New York : Interscience Publishers, 1956.

[108] Hutchens G J 1995 *Journal of Applied Physics* **77** 2912–2915 URL https://doi.org/10.1063/1.358706

[109] Taylor G I 1950 *Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences* **201** 159–174

[110] Zel'dovich Y and Raizer Y 2012 *Physics of Shock Waves and High-Temperature Hydrodynamic Phenomena* Dover Books on Physics (Dover Publications) ISBN 9780486135083

[111] MacFarlane J, Golovkin I and Woodruff P 2006 *Journal of Quantitative Spectroscopy and Radiative Transfer* **99** 381–397

[112] Feder L, Miao B, Shrock J E, Goffin A and Milchberg H M 2020 *Phys. Rev. Res.* **2**(4) 043173 URL https://link.aps.org/doi/10.1103/PhysRevResearch.2.043173

[113] Wojda F, Cassou K, Genoud G, Burza M, Glinec Y, Lundh O, Persson A, Vieux G, Brunetti E, Shanks R P, Jaroszynski D, Andreev N E, Wahlström C G and Cros B 2009 *Phys. Rev. E* **80**(6) 066403 URL https://link.aps.org/doi/10.1103/PhysRevE.80.066403

[114] USAF targets - Edmund Optics Knowledge Center URL https://www.edmundoptics.co.uk/knowledge-center/application-notes/imaging/testing-and-targets/

[115] ISO 11146 Beam Size Definitions URL https://laserbeamsize.readthedocs.io/en/latest/01-Definitions.html

[116] Lemos N, Grismayer T, Cardoso L, Figueira G, Issac R, Jaroszynski D A and Dias J M 2013 *Physics of Plasmas* **20** 063102 (*Preprint* https://doi.org/10.1063/1.4810797) URL https://doi.org/10.1063/1.4810797

[117] Lemos N, Cardoso L, Geada J, Figueira G, Albert F and Dias J M 2018 *Scientific Reports* **8** 3165 ISSN 2045-2322 URL https://doi.org/10.1038/s41598-018-21392-z

[118] McLeod J H 1954 *J. Opt. Soc. Am.* **44** 592–597

[119] Lemos N, Grismayer T, Cardoso L, Geada J, Figueira G and Dias J M 2013 *Physics of Plasmas* **20** 103109 (*Preprint* https://doi.org/10.1063/1.4825228) URL https://doi.org/10.1063/1.4825228

[120] Assmann R W *et al.* 2020 *The European Physical Journal Special Topics* **229** 3675–4284 ISSN 1951-6401

[121] Smartsev S, Caizergues C, Oubrerie K, Gautier J, Goddet J P, Tafzi A, Phuoc K T, Malka V and Thaury C 2019 *Optics letters* **44** 3414–3417

[122] Goodman J W 2005 *Introduction to Fourier Optics* (Roberts & Company Publishers)

[123] Born M and Wolf E 1964 *Principles of optics: electromagnetic theory of propagation, interference and diffraction of light; 2nd ed.* (Oxford: Pergamon)

[124] Fitzpatrick R Classical Electromagnetism URL https://farside.ph.utexas.edu/teaching/jk1/Electromagnetism/

[125] Durnin J 1987 *J. Opt. Soc. Am. A* **4** 651–654 URL https://opg.optica.org/josaa/abstract.cfm?URI=josaa-4-4-651

[126] Wang Y, Yan S, Friberg A T, Kuebel D and Visser T D 2017 *J. Opt. Soc. Am. A* **34** 1201–1211

[127] Edmund Optics Reflective Axicon URL https://www.edmundoptics.com/p/20-axicon-angle-1000---1100nm-off-axis-reflective-axicon/42646/

[128] Miao B, Shrock J E, Feder L, Hollinger R C, Morrison J, Nedbailo R, Picksley A, Song H, Wang S, Rocca J J and Milchberg H M 2022 *Phys. Rev. X* **12**(3) 031038 URL https://link.aps.org/doi/10.1103/PhysRevX.12.031038

[129] Swanson G J and Veldkamp W B 1989 *Optical Engineering* **28** 286605 URL https://doi.org/10.1117/12.7977008

[130] Zemax Optic Studio https://www.zemax.com/pages/opticstudio [Accessed 24-January-2022]

[131] LightTrans VirtualLab Fusion https://www.lighttrans.com/ [Accessed 24-January-2022]

[132] Voelz D 2011 *Computational Fourier Optics* (Society of Photo-Optical Instrumentation Engineers)

[133] Wang C, Ma J, Cai C and Su P 2023 *Sensors* **23** ISSN 1424-8220 URL https://www.mdpi.com/1424-8220/23/3/1662

[134] Refractive Index Information https://refractiveindex.info [Accessed 24-January-2020]

[135] Miao B 2020 *Laser wakefield accelerator experiments: coherent injection radiation and optical field ionization-based plasma waveguides* Ph.D. thesis University of Maryland

[136] Longman A 2020 *Under-dense laser-plasma interactions in relativistic optical vortices* Ph.D. thesis University of Alberta

[137] Longman A, Salgado C, Zeraouli G, Apiñaniz J I, Pérez-Hernández J A, Eltahlawy M K, Volpe L and Fedosejevs R 2020 *Optics letters* **45** 2187–2190

[138] Davidson N, Friesem A and Hasman E 1992 *Optics Communications* **88** 326–330 ISSN 0030-4018 URL https://www.sciencedirect.com/science/article/pii/0030401892900502

[139] Zhang N, Ye J S, Feng S F, Wang X K, Han P, Sun W F, Zhang Y and Zhang X C 2021 *OSA Continuum* **4** 1223–1233 URL https://opg.optica.org/osac/abstract.cfm?URI=osac-4-4-1223

[140] Sochacki J, Bara S, Jaroszewicz Z and Kołodziejczyk A 1992 *Optics letters* **17** 7–9

[141] Gessner S 2016 *Demonstration of the hollow channel plasma wakefield accelerator* Ph.D. thesis Stanford University

[142] Sochacki J, Kołodziejczyk A, Jaroszewicz Z and Bará S 1992 *Appl. Opt.* **31** 5326–5330 URL https://opg.optica.org/ao/abstract.cfm?URI=ao-31-25-5326

[143] Oubrerie K, Andriyash I A, Lahaye R, Smartsev S, Malka V and Thaury C 2022 *Journal of Optics* **24** 045503 URL https://dx.doi.org/10.1088/2040-8986/ac57d2

[144] Caizergues C, Smartsev S, Malka V and Thaury C 2020 *Nature Photonics* **14** 475–479

[145] Sochacki J, Jaroszewicz Z, Staroński L R and Kołodziejczyk A 1993 *J. Opt. Soc. Am. A* **10** 1765–1768 URL https://opg.optica.org/josaa/abstract.cfm?URI=josaa-10-8-1765

[146] Herman R M and Wiggins T A 1992 *Appl. Opt.* **31** 5913–5915 URL https://opg.optica.org/ao/abstract.cfm?URI=ao-31-28-5913

[147] Cox A J and D'Anna J 1992 *Opt. Lett.* **17** 232–234 URL https://opg.optica.org/ol/abstract.cfm?URI=ol-17-4-232

[148] Picksley A, Chappell J, Archer E, Bourgeois N, Cowley J, Feder L, Jakobsson O, Ross A J, Wang W, Walczak R and Hooker S GeV electron bunches in low-density plasma channels by all-optical density transition injection [In preparation]

[149] Fan J, Parra E and Milchberg H 2000 *Physical Review Letters* **84** 3085

[150] Scott R, Booth N, Hawkes S, Symes D, Hooker C, Doyle H, Olsson-Robbie S, Lowe H, Price C, Bigourd D *et al.* 2020 *Physics of Plasmas* **27** 033301

[151] Wang Y, Yan S, Friberg A T, Kuebel D and Visser T D 2017 *JOSA A* **34** 1201–1211

[152] Rico-Botero V M and Negrete-Regagnon P 2020 *Applied Optics* **59** 5553–5559

[153] Grunwald R and Bock M 2020 *Advances in Physics: X* **5** 1736950

[154] Shen C, Hong Q, Zhu Q, Zu C and Wei S 2019 *Optics & Laser Technology* **120** 105682

[155] Zamboni-Rached M 2021 *Optics Letters* **46** 1205–1208

[156] Paterson C and Smith R 1996 *Optics Communications* **124** 121–130

[157] Götte J B, O'Holleran K, Preece D, Flossmann F, Franke-Arnold S, Barnett S M and Padgett M J 2008 *Opt. Express* **16** 993–1006 URL https://opg.optica.org/oe/abstract.cfm?URI=oe-16-2-993

[158] Dixit S, Thomas I M, Woods B W, Morgan A, Henesian M, Wegner P and Powell H 1993 *Applied optics* **32** 2543–2554

[159] Hu X B, Dong M X, Zhu Z H, Gao W and Rosales-Guzmán C 2020 *Scientific reports* **10** 1–11

[160] Dudley A, Vasilyeu R, Belyi V, Khilo N, Ropot P and Forbes A 2012 *Optics Communications* **285** 5–12

[161] Belyi V, Kazak N, Khilo N and Ropot P 2007 Vortex structure of elongated speckles *ICONO 2007: Coherent and Nonlinear Optical Phenomena* vol 6729 (SPIE) pp 41–50

[162] Matsumoto N, Ando T, Inoue T, Ohtake Y, Fukuchi N and Hara T 2008 *JOSA A* **25** 1642–1651

[163] Larsen J 2017 *Foundations of High-Energy-Density Physics: Physical Processes of Matter at Extreme Conditions* (Cambridge University Press)

[164] Shalloo R J, Mewes M, Boyle G, Arran C, Corner L, Osterhoff J, Walczak R, Hooker S and Thévenet M 2022 Simulations of hydrodynamic optical-field-ionised plasma channels Advanced Accelerator Concepts Workshop URL https://indico.classe.cornell.edu/event/2108/contributions/1971/contribution.pdf

[165] Picksley A 2021 *Low Density Plasma Waveguides for Multi-GeV Laser Wakefield Accelerators* Ph.D. thesis University of Oxford

[166] McLachlan G J and Peel D 2000 *Finite Mixture Models* Wiley Series in Probability and Statistics (John Wiley & Sons, Inc) ISBN 978-0-471-00626-8

[167] Mittelberger D E, Nakamura K, Lehe R, Gonsalves A J, Benedetti C, Mao H S, Daniels J, Dale N, Venkatakrishnan S V, Swanson K K, Esarey E and Leemans W P 2016 *J. Opt. Soc. Am. B* **33** 1978–1985 URL https://opg.optica.org/josab/abstract.cfm?URI=josab-33-9-1978

[168] Fedeli L, Huebl A, Boillod-Cerneux F, Clark T, Gott K, Hillairet C, Jaure S, Leblanc A, Lehe R, Myers A, Piechurski C, Sato M, Zaim N, Zhang W, Vay J L and Vincenti H 2022 Pushing the frontier in the design of laser-based electron accelerators with groundbreaking mesh-refined particle-in-cell simulations on exascale-class supercomputers. sc22: International conference for high performance computing, networking, storage and analysis (sc). *2022 SC22: International Conference for High Performance Computing, Networking, Storage and Analysis (SC) (SC)* pp 25–36 ISSN 2167-4337 URL https://www.computer.org/csdl/proceedings-article/sc/2022/544400a025/1I0bSKaoECc

[169] Archambeau F, Méchitoua N and Sakiz M 2004 *International Journal on Finite Volumes* **1** URL https://hal.archives-ouvertes.fr/hal-01115371/document

[170] Gonsalves A J, Nakamura K, Lin C, Osterhoff J, Shiraishi S, Schroeder C B, Geddes C G R, Tóth C, Esarey E and Leemans W P 2010 *Physics of Plasmas* **17** 056706 (*Preprint* https://doi.org/10.1063/1.3357175) URL https://doi.org/10.1063/1.3357175