Techniques for High Repetition Rate Laser Wakefield Acceleration



Christopher Arran

St John's College, Oxford

Thesis submitted in fulfilment of the requirements for the degree of

Doctor of Philosophy at the University of Oxford

Trinity Term, 2018

Abstract

I present work on three different techniques in order to contribute towards laser wakefield acceleration (LWFA) experiments which could operate at multi-kHz repetition rates and at low plasma densities $n_e \sim 10^{17}\,\mathrm{cm}^{-3}$, enabling the acceleration of higher energy electrons in a single stage. Together, these techniques would help to make the next generation of LWFA facilities substantially more competitive with conventional accelerators.

Firstly, I consider the accuracy required to resonantly excite plasma wakefields with trains of many laser pulses, produced by high average-power but lower peak-power laser systems. Through theory, simulations and experimental results, I show a strong constraint on the systematic error acceptable in the separation of pulses. However, I also explore using chirped trains of laser pulses, where the pulse separation increases over time, to increase the efficacy of resonant excitation when the pulse spacing cannot be perfectly matched to the plasma.

Secondly, I describe a recent technique which can measure plasma waves rapidly and sensitively, using temporally encoded spectral shifting (TESS). I generalise the technique to probe pulses with non-Gaussian profiles and demonstrate that it can be extended to measure quasi-linear wakefields which are not sinusoidal. I detail how a TESS analysis can be conducted in practice and present results from a recent experiment using the technique.

Thirdly, I outline a proposed all-optical method of creating plasma waveguides which would be capable of operating at low densities and high repetition rates and are in principle immune to damage from laser pulses. Start to end simulations demonstrate that these hydrodynamic optical-field-ionized plasma channels could meet all of the requirements for next generation laser wakefield acceleration experiments and are used to propose some optimum experimental parameters.

Acknowledgments

When I started the DPhil, I knew it was going to be a painful marathon of demoralising grunt work, which would pat me on the shoulder with a worthless scrap of paper and a silly hat and convince me to leave academia and never come back. And yet...

I know now, better than I could have imagined, that the whole beautiful but shonky ivory edifice of research is built upon the overburdened backs of underpaid and undervalued grad students. And yet...

I've loved these last four years, and as I write this, sitting in the Rad Cam knowing that thousands of people have written their theses here before me, I can't stop thinking about how lucky I've been, and how much of that is because of the people: you wonderful, passionate, inspiring people. At tea breaks chatting politics, at Elba riding off into the sunset on a pedalo, even at 4 am in the control room at RAL, feeling my heart sink while watching the referendum results come in on one screen, hearing the pop of the flash lamps time after time and staring at hundreds of shots coming in negative on more screens; knowing that we're all working on it together has made it worth it and more; it's been sublime.* So to Rob, Jimmy, Dan, David, Jakob and Gavin, Alex and Alexander, Chris and Laura, and, yes, to Anthony too, thank you, more than I can possibly say.

When I think about the incredible people in the lab, only one explanation can really come to mind. I've seen plenty of my friends have a really shit time throughout the years, without the support they need and deserve, pushing pages into the void without any feedback, or simply being expected to work evenings and weekends without any

*That's not forgetting the original meaning of that word... and yeah, I know, what kind of nerd puts footnotes in the acknowledgements?

kind of thanks. Don't get me wrong, just because I've had some of the best years of my life, doesn't mean I don't think the system is seriously flawed. That I'm proud of this thesis, that I want to stay in academia as long as it'll have me, and that I've loved these years, can only be down to Roman and most of all to Simon. Chances are, if you're reading this, you're probably one of his students, and however long the meetings, and however many drafts you go through before the paper's up to scratch, I'd like you to know that he's one of the good ones.

I'd also like to pass on some advice, dear reader, and that's to do everything else you can, when you can, that isn't physics. There'll definitely be enough physics in your life, and yet not so much that you won't stay late in the evenings just to fit a little more in. When I leave Oxford, though, I know there won't be wine and cheese nights, or cocktail nights to organize, Iffley to climb at, or scholars of mediaeval outlaw ballads to talk to over lunch. A week after I write this, I won't be able to drop in to the Natural History Museum to sit and admire the eagle or nip to Gloucester Green to grab some dumplings. So take the time off that you're entitled to (6–8 weeks a year for grad students, and don't forget it!) and wherever you are, whatever you love, please do as much of it as you can. Thanks to the Simon room, to the climbers, and thanks to the MCR at St John's, and Mikayla, Anthony, Kieran, Laura, Ellie, Sally and Rob, I know I did.

The only reason I can drag myself away from here with a smile is because of what I have to look forward to – I love you, Charlotte, always.

So thanks everyone, it's been a whirl. And thank you reader – so long as you're reading this, those four years aren't dead yet. Good luck with everything!

Contents

List of Figures **vi**

1 Introduction **1**
1.1 Laser Wakefield Acceleration . . . 2
1.1.1 Accelerators . . . 2
1.1.2 Laser Wakefield Accelerators . . . 3
1.1.3 Challenges . . . 4
1.1.4 Chapter Overview . . . 6
1.2 Driving . . . 7
1.2.1 In Theory . . . 7
1.2.2 In Practice . . . 14
1.3 Diagnosing . . . 18
1.3.1 In Practice . . . 18
1.3.2 In Theory . . . 23
1.4 Guiding . . . 27
1.4.1 Problem . . . 27
1.4.2 Solutions . . . 31
1.5 Thesis Overview . . . 35

2 Exciting Laser Wakefields with Resonant Pulses **37**
2.1 Introduction . . . 38
2.2 Background & Linear Theory . . . 39
2.2.1 Systematic Errors . . . 41
2.2.2 Random Errors . . . 42
2.3 Simulations . . . 44
2.3.1 PIC Codes . . . 44
2.3.2 Parameters . . . 48
2.4 Errors in Pulse Spacing . . . 50
2.4.1 Systematic Spacing Error . . . 50
2.4.2 Pulse Spacing Jitter . . . 52
2.5 Auto-Resonance . . . 55
2.5.1 Background . . . 55
2.5.2 Multiple Pulse Auto-Resonance . . . 58
2.5.3 Results . . . 63
2.5.4 Discussion . . . 67
2.6 Conclusions . . . 69

3 Measuring Wakefields with Temporally Encoded Spectral Shifting 72
3.1 Introduction . . . 73
3.2 TESS Background . . . 74
3.2.1 TESS and FDH . . . 75
3.2.2 TESS Signal . . . 78
3.3 Practical Considerations . . . 81
3.3.1 Probe Slippage . . . 81
3.3.2 Retrieving Phase from the Ionisation Front . . . 82
3.3.3 Errors from Finite Bandwidth . . . 85
3.4 Arbitrary Probe and Reference Pulses . . . 86
3.4.1 Non-Gaussian Pulses . . . 86
3.4.2 Non-Equal GDDs . . . 90
3.5 Non-Linear Wakefields . . . 91
3.5.1 Quasi-Linear Wakefields . . . 91
3.5.2 General Wakefields . . . 98
3.6 Conclusions . . . 103

4 A Resonant Excitation Experiment 106
4.1 Introduction . . . 107
4.2 Experimental Set Up . . . 108
4.2.1 Pre-Interaction Point . . . 108
4.2.2 Post-Interaction Point . . . 110
4.3 Modified FDH Analysis . . . 112
4.3.1 Extracting Spectral Phase . . . 112
4.3.2 Extracting Temporal Phase . . . 114
4.4 TESS Analysis . . . 118
4.4.1 Calculating the TESS signal . . . 118
4.4.2 TESS Satellite Peak Locations . . . 119
4.4.3 TESS Satellite Peak Heights . . . 122
4.5 Results . . . 127
4.6 Discussion . . . 131
4.7 Conclusions . . . 133

5 Low Density Plasma Waveguides from Optical Field Ionization 135
5.1 Introduction . . . 136
5.2 Electron Energies from OFI Heating . . . 138
5.2.1 Ionization Models . . . 140
5.2.2 Simpleman Model . . . 142
5.2.3 Simulations with EPOCH . . . 145
5.2.4 Discussion . . . 149
5.3 Shock Expansion in Plasma . . . 150
5.3.1 Fluid Approximation . . . 150
5.3.2 Shock Expansion Theory . . . 153
5.3.3 Simulation Codes . . . 155
5.4 Axicons . . . 157
5.4.1 Optimising Axicon Parameters . . . 160

5.4.2 Constraints on Axicon Angle . . . 166
5.5 Simulated Axicon OFI Channels . . . 170
5.5.1 Channel Properties . . . 174
5.6 Conclusions . . . 175

6 Conclusions **179**
6.1 Original Aims . . . 180
6.2 Driving with Multiple Resonant Pulses . . . 181
6.2.1 Errors in Achieving Resonance . . . 181
6.2.2 Auto-Resonance . . . 183
6.2.3 Future Work . . . 185
6.3 Diagnosing with Generalized TESS . . . 187
6.3.1 Generalizing TESS . . . 187
6.3.2 Future Work . . . 192
6.4 A Resonant Excitation Experiment . . . 193
6.4.1 Discussion . . . 193
6.4.2 Future Work . . . 194
6.5 Guiding with a Hydrodynamic Plasma Waveguide . . . 196
6.5.1 Plasma Column Formation . . . 196
6.5.2 Channel Evolution . . . 197
6.5.3 Future Work . . . 199
6.6 Summary . . . 202

A Derivation of Wakefield Amplitude and Phase in Autoresonance **203**

B Detailed TESS Calculations **208**
B.1 Case $\psi_{\mathrm{a}}^{(2)} = \psi_{\mathrm{b}}^{(2)}$ . . . 209
B.2 Case $\psi_{\mathrm{a}}^{(2)} \approx \psi_{\mathrm{b}}^{(2)}$ . . . 210
B.3 Case $\psi_{\mathrm{a}}^{(2)} \neq \psi_{\mathrm{b}}^{(2)}$ for Gaussian $|\mathcal{E}_{\mathrm{x}}(\omega)|$ . . . 212

References **215**

List of Figures

1.1 Comparison of beatwave and ultra-short pulses for driving LWFA . . . 9
1.2 Driving a wakefield with the ponderomotive force 10
1.3 Electric fields in a linear wakefield 14
1.4 Spectral shifts due to plasma waves 19
1.5 Experimental scheme for FDH and TESS 20
1.6 Spectral shifts for ultra-short pulses 22
1.7 Guiding in step-index and gradient-index waveguides 29

2.1 Driving wakefields with multiple laser pulses 40
2.2 Simulation process for Particle-In-Cell codes 44
2.3 Effect of systematic errors on wakefield amplitude 51
2.4 Effect of random errors on wakefield amplitude 53
2.5 The Euler spiral for chirped pulse trains 61
2.6 Wakefield amplitudes resulting from chirped pulse trains 64
2.7 Wakefield frequencies resulting from chirped pulse trains 65
2.8 Effect of chirp rate on wakefield amplitude 66
2.9 Comparison of chirped pulse train simulations with theory 66
2.10 Optimum chirp rates 68

3.1 Spectral modulation in TESS 77
3.2 Comparison of TESS and FDH processes 78
3.3 Determining the spectral phase of a laser pulse from an ionization front 84

3.4 Effect of the spectral profile of a laser pulse on the TESS overlap factor 89
3.5 Effect of the spectral profile on the inferred wakefield amplitude 90
3.6 Density profiles and TESS signals for quasi-linear plasma waves 92
3.7 Retrieving wakefield amplitudes for quasi-linear plasma waves 96
3.8 TESS signals and reconstructed density for non-linear plasma waves . . 100

4.1 Experimental layout of the Astra 2015 experiment 108
4.2 Measured spectral interferograms of probe and reference pulses 111
4.3 The interference term shown in the TESS signal 114
4.4 The reconstructed temporal phase and electron density in FDH 117
4.5 TESS signals showing signs of a plasma wakefield 120
4.6 The TESS satellite peak resulting from a plasma wakefield 121
4.7 TESS peak locations and inferred plasma periods 122
4.8 Background subtraction in TESS analysis 123
4.9 Calculation of the overlap factor for TESS 125
4.10 Wakefield amplitudes from resonant excitation 129
4.11 Inferring pulse separation from the measured pump pulse spectrum . . . 131

5.1 Atomic potentials in the presence of a strong electric field 139
5.2 Effect of intensity on electron energy spectra 147
5.3 Effect of ellipticity on electron energy spectra 148
5.4 Focussing from an axicon lens 157
5.5 Intensity profiles along an axicon focus 158
5.6 Effect of a mask on the focus of an axicon 161
5.7 Maximising the energy supplied along an axicon focus 162
5.8 Optimal plasma column formed along an axicon focus 163
5.9 Initial plasma columns from simulating an axicon focus 170
5.10 Hydrodynamic plasma column expansion from simulations 172
5.11 Lowest order guided modes inside simulated plasma channels 174

5.12 Guiding properties of plasma channels formed by hydrodynamic expansion 176

6.1 Effect of random and systematic errors on wakefield amplitude 184

6.2 Wakefield amplitudes resulting from linearly chirped pulse trains 186

6.3 Effect of a non-Gaussian spectral profile on a TESS measurement . . . 188

6.4 TESS measurements of quasi-linear plasma waves 190

6.5 TESS measurements of the bubble regime 192

6.6 Measurements of a plasma wave and the plasma period 194

6.7 Measured wakefield amplitudes from resonant excitation 195

6.8 Initial conditions of a plasma column from simulations of OFI 198

6.9 Hydrodynamic expansion of a plasma column formed with OFI 200

Relevant Publications

- J. Cowley, C. Thornton, C. Arran *et al.*

  "Excitation and Control of Plasma Wakefields by Multiple Laser Pulses"

  Phys. Rev. Lett., vol. 119, no. 4, pp. 044802, 2017

- R. J. Shalloo, C. Arran, L. Corner *et al.*

  "Hydrodynamic optical-field-ionized plasma channels"

  Phys. Rev. E, vol. 97, no. 5, pp. 053203, 2018

- C. Arran, N. H. Matlis, R. Walczak *et al.*

  "Reconstructing nonlinear plasma wakefields using a generalized temporally encoded spectral shifting analysis"

  Phys. Rev. Accel. Beams, vol. 21, pp. 103501, 2018

Chapter 1

Introduction

1.1 Laser Wakefield Acceleration

1.1.1 Accelerators

Accelerators have played a crucial role in shaping both modern physics and the world around us. Almost no other tool has been so successful at probing the fundamental states of matter and forces that make up our universe (e.g. [1], [2]) and at laying bare the structures inside our bodies, our pharmaceuticals and our machines, whether an airplane turbine or a virus. In order to probe structures and processes with unprecedented resolution in both time and space, accelerator facilities around the world produce high-brightness tunable x-ray bunches at kHz repetition rates, with femtosecond durations and photon energies on the keV scale [3], [4]. The high energy beams produced by accelerators enable the exploration of physical processes in extreme non-classical regimes.

However, these machines are neither small nor cheap, and while the technology that underpins almost all modern accelerators, radio frequency cavities, continues to improve [5], [6], there are limits to the electric fields, and hence the acceleration gradients, which can be achieved, preventing conventional accelerators from becoming much smaller or cheaper. Electrical breakdown from microscopic imperfections in the cavity becomes a limiting factor, with the highest achievable electric field well below 100 MVm^{-1}, and this implies the devices must be increasingly large to attain the desired energies. Accelerators range from meter scale devices creating MeV electrons for small sample treatment facilities, through the GeV scale for 3rd generation x-ray light sources with beamlines hundreds of meters long (e.g. [7]), to the TeV scale for particle colliders at the forefront of fundamental research, which must be tens of km long. The technology has been effective for almost a century but research is increasingly focussing on alternative acceleration technologies, with a large number of novel accelerator schemes promising high gradient acceleration.

1.1.2 Laser Wakefield Accelerators

Foremost among these is laser wakefield acceleration (LWFA), first suggested by theory and simulations almost four decades ago [8]*. An intense laser pulse passing through a plasma has a strongly varying electric field which drives electrons away from the most intense parts of the beam, while leaving the heavier ions almost stationary. This charge separation drives a powerful electronic wave, which is associated with an electric field on the scale of $E_{\mathrm{wb}} \sim m_e \omega_{\mathrm{p}} c/e$, where $\omega_{\mathrm{p}} = \sqrt{n_e e^2/m_e \epsilon_0}$ is the plasma frequency. For electron densities on the order of $10^{18}\,\mathrm{cm}^{-3}$, this field is approximately $E_{\mathrm{wb}} \approx 100\,\mathrm{GVm}^{-1}$, several orders of magnitude higher than achievable with conventional accelerators. This plasma wave has a phase velocity equal to the group velocity of the driving laser pulse, which is very close to the speed of light. These waves are therefore naturally suited to accelerate highly relativistic particles.

Experiments throughout the last three decades have reported the acceleration of electrons to steadily higher energies with ever greater control [9], [10]. Early work with the low intensity nanosecond duration lasers available at the time used the difference frequency between two longer duration laser beams to generate a plasma wave, in a beatwave process [11], similar to previous stimulated Raman scattering experiments [12]. Another model of the beatwave process is to look at the electric field of the laser as a train of many laser pulses, spaced by the plasma period $\tau_{\mathrm{p}} = 2\pi/\omega_{\mathrm{p}}$, each with a duration shorter than the plasma period. As laser technology improved, research moved towards shorter and more intense laser pulses, which could have durations comparable to the plasma period in a single pulse, without using interference. These laser pulses are both more convenient to build in an experiment, with only a single beamline from a single source laser, and also need not be exactly matched to the plasma period, with a single pulse exciting a wakefield over a wider range of electron densities. Experiments with these short-pulse systems have led to rapid improvements in LWFA, producing

*With the greatest apologies to Immanuel Kant: "[Theorien] ohne [Experimente] sind leer, [Experimente] ohne [Theorien] sind blind."

quasi-monoenergetic electron beams [13], now with energies on the GeV scale [14], [15].

Already laser wakefield accelerators are being employed in a variety of research applications. Transverse oscillations of an electron bunch inside the plasma wave, known as betatron oscillations, give rise to high brightness x-rays, originating from a source size on the μm scale [13]. These x-rays can be used to image delicate samples at high resolution [16]. Furthermore, the x-ray emission is of femtosecond duration and naturally synchronised with the laser pulse, making LWFA systems ideal for pump-probe imaging of ultra-fast phenomena such as shock propagation in solids [17]. Synchronization of a high energy electron beam with the laser pulse also allows study of electrodynamic interactions such as inverse Compton scattering [18]–[21] and radiation reaction [22], [23], which are difficult to achieve with conventional accelerators. Inverse Compton scattering can potentially lead to the excitation of high energy γ radiation on the MeV scale. LWFA has also been used to produce secondary radiation when the electron bunch is incident upon a transformer target made of a heavy element such as lead. In this manner, researchers have generated particles such as positrons [24] and pions [25].

1.1.3 Challenges

However, the powerful ultra-short laser systems used for these experiments, almost ubiquitously built around the amplifier medium of Titanium doped sapphire crystals, are limited by the large amount of heat dumped into the medium, in their efficiency to significantly under 1% from the wall-plug to the optical pulse, and in their repetition rates to around 1 Hz. This is primarily due to the high quantum defect of the amplifier medium, which implies a large amount of the supplied energy is lost to heat, and to the difficulty of efficiently pumping Titanium sapphire systems. Finding alternative laser systems which are more efficient is therefore an active field of research [26], [27]. It has been suggested by my colleagues in Oxford that we might move back towards a resonance scheme, where a train of many lower energy laser pulses, produced by a

more efficient and higher repetition rate laser system, could be spaced by the plasma period and hence constructively drive a wakefield in the linear and quasi-linear regimes [28], [29]. This Multiple Pulse Laser Wakefield Acceleration (MP-LWFA) scheme could enable laser wakefield accelerators to operate at repetition rates on the kHz scale and at efficiencies on the $> 10\%$ scale from the wall-plug to the laser pulse. Because it relies upon resonance with the plasma wave, however, the limitations upon the pulse train driver must be well understood in order to efficiently drive a wakefield.

Laser wakefield accelerators face several other problems in addition to the laser technology. While the laser pulse travels at close to the speed of light, it travels slightly slower than highly relativistic particles. An electron at 1 GeV scale has $\gamma = E/m_e c^2 \approx 2000$ and travels at $v = c\sqrt{1-\gamma^{-2}} \approx (1-10^{-7})c$, whereas a laser pulse with a frequency ω_0 travelling through an underdense plasma travels at a group velocity $v_\mathrm{g} \approx c\sqrt{1-(\omega_\mathrm{p}/\omega_0)^2}$. For an 800 nm laser pulse in a plasma at $n_e = 10^{18}\,\mathrm{cm}^{-3}$ this is $v_\mathrm{g} \approx (1-10^{-4})c$. An electron bunch will therefore move through the plasma faster than a LWFA driver pulse and hence move forward in the plasma wave. This implies that it can only propagate in the accelerating phase of the plasma wave for a certain distance. For an underdense plasma with $n_e \ll n_c$ this is described by the dephasing length $L_\mathrm{dp} = \pi c\omega_0^2/2\omega_\mathrm{p}^3$, which at $n_e = 10^{18}\,\mathrm{cm}^{-3}$ is only 30 millimeters. With the accelerating field on the scale of E_wb, the maximum energy gain of a single LWFA stage is $\Delta W \sim (\omega_0/\omega_\mathrm{p})^2 m_e c^2$. There is therefore a strong motivation to decrease the plasma density of a laser wakefield accelerator, despite the significantly longer acceleration lengths that this requires.

Dephasing is not the only limitation on the length and energy gain of an LWFA accelerator. To reach the high intensities required to drive a plasma wave, lasers are strongly focussed to a small spot size, on the tens of micron scale. The laser driver therefore rapidly diffracts away from focus, with a Rayleigh range described by $Z_R = \pi w_0^2/\lambda_0$ for a laser spot size w_0 and laser wavelength λ_0. This Rayleigh range is typically only a few millimeters. If a laser pulse can be prevented from diffracting by

using a waveguide, the energy gain of an electron bunch is proportional to the length of the waveguide, up to the dephasing limit. A tremendous opportunity exists for waveguides which are long and can operate at low density and high repetition rate.

When operating at low density, however, particularly with wakefields in the linear and quasi-linear regimes, it is difficult to measure a plasma wave or to diagnose issues with the accelerating structure. The phase change associated with an optical probe passing through a plasma is described by $\phi \approx (\omega_p^2 L/2\omega_0 c)(\delta n_e/n_{e0})$ [30] and hence decreases rapidly with plasma density unless the wavelength of the probe pulse ω_0 is also reduced accordingly. A sensitive diagnostic, Temporally Encoded Spectral Shifting (TESS), has been designed [31], using modulation of the spectrum of a probe pulse which co-propagates with the plasma wakefield. Spectral interference between this probe pulse and a reference pulse, which precedes it along the same path without encountering the plasma wakefield, reveals information about the plasma wave. TESS is well suited to operating at high repetition rates, measuring wakefield amplitudes and frequencies rapidly and quantitatively. To date, however, the TESS technique has assumed both that the probe and reference pulses are identical with Gaussian spectral profiles, and that the plasma wave is linear.

1.1.4 Chapter Overview

In the following chapters, I discuss work towards resolving these issues, studying driving, diagnosing and guiding wakefields at lower densities and higher repetition rates. In this chapter, I first write about how a laser pulse drives a plasma wakefield through the ponderomotive force and how the resulting plasma wave depends on the laser driver. I outline the rationale for multiple pulse laser wakefield acceleration and introduce my work on the limitations on pulse spacing in the MP-LWFA scheme. Secondly, I discuss the different methods of diagnosing plasma wakefields, their advantages and disadvantages, and introduce how a plasma affects an optical probe in both TESS and

a related technique, FDH. My own work on generalizing TESS is introduced. Thirdly, I write about waveguides for LWFA, the different methods used to date, and the requirements for waveguides for laser wakefield acceleration. This puts into context my work on hydrodynamic plasma waveguides.

Throughout this thesis I focus on wakefields in the linear and quasi-linear regimes, where electrons in the plasma wave continue to oscillate for many periods after the laser driver has passed, with velocities significantly less than the speed of light, and where the effect of feedback from the plasma wave upon the laser pulse can be neglected. In these regimes wakefields are well described both analytically and numerically, and we can avoid studying effects such as self-injection and self-focussing. There is already a large volume of high-quality work on studying and utilising self-injection and self-focussing in LWFA (e.g. [32], [33]), by operating at higher densities and in the bubble regime, and I shall not attempt to write anything as good. Instead, linear and quasi-linear wakefields provide one of the most promising avenues for achieving controllable, low emittance and low energy spread electron bunches, which can achieve high energy gain in a single stage.

1.2 Driving

1.2.1 In Theory

There are two primary models for how an intense laser pulse drives a plasma wakefield. The first is a wave picture, in which energy is transferred between the travelling wave of the laser's electric field and the electronic Langmuir wave of the plasma. Consider two laser pulses co-propagating with frequencies ω_0 and ω_1. These create an electric field with a beating pattern, with a fast oscillation at the sum frequency $\omega_+ = (\omega_0 + \omega_1)/2$ within an envelope oscillating more slowly at the difference frequency $\omega_- = (\omega_1 - \omega_0)/2$. Each point in the plasma responds to incident electromagnetic waves with a dielectric

response function, and this has a strong peak at the plasma frequency ω_p, corresponding to Langmuir waves. If the difference frequency is tuned to the plasma frequency, with $\omega_1 = \omega_0 + \omega_\mathrm{p}$, the Langmuir wave amplitude begins to grow rapidly, driven by the electric field of the laser. Figure 1.1a shows such a situation, with a strong peak in the power spectrum of the laser pulse at the plasma frequency.

In practice, a broadband laser pulse, such as those used to create ultra-fast laser pulses, contains many frequencies. If the bandwidth limited pulse duration is less than the plasma period $\tau_\mathrm{min} < \tau_\mathrm{p} = 2\pi/\omega_\mathrm{p}$, the spectral width is necessarily larger than the plasma frequency $\sigma_\omega \sim 2\pi/\tau_\mathrm{min} > \omega_\mathrm{p}$. There can therefore be found within the single pulse many pairs of frequencies such that $\omega_1 - \omega_0 = \omega_\mathrm{p}$ and if the laser pulse has a sufficiently flat phase profile that these frequencies occur at the same time, the resulting difference frequency will drive a plasma wave within a single pulse. Figure 1.1b shows an example of a sufficiently short pulse, where the power spectrum overlaps with the plasma frequency.

The second, mechanistic, picture, is based around the ponderomotive force from an intense laser pulse. A laser pulse exerts a force upon charged particles in a plasma described by [34]:

$$\mathbf{F} = -\frac{q^2}{4m\omega_0^2}\nabla\left(E_L^2\right), \tag{1.1}$$

where m and q are the mass and charge of the particle and E_L is the electric field strength of the laser pulse. This can be derived by splitting the motion of the charged particle into a fast oscillation and a slow drift component. Essentially, a particle oscillates faster and further in the more intense parts of a laser pulse, and hence is driven, on average, away from the intensity peak of the beam. It is worth noticing that this force applies to both ions and electrons, but because of the slower oscillations of the heavier ions, they remain essentially stationary, with the acceleration described by $a = F/m \propto m^{-2}$.

Electrons are therefore displaced from around the laser pulse both longitudinally

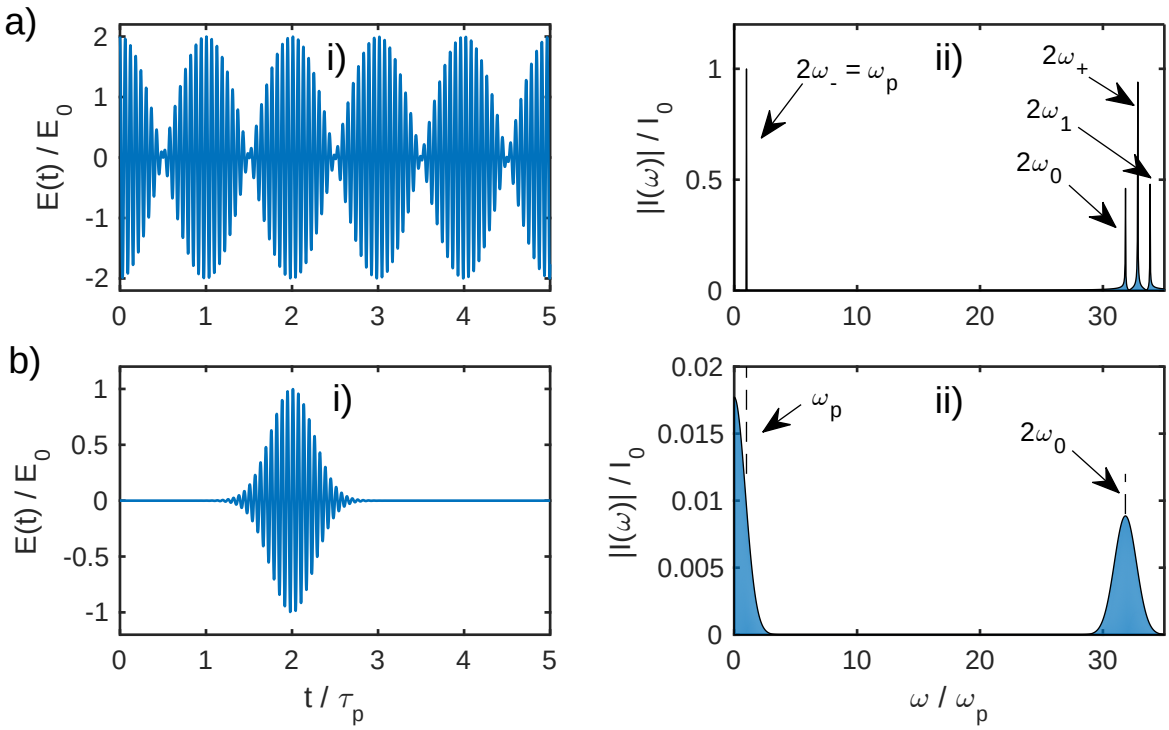


Figure 1.1: i) Electric field of a laser pulse over time and ii) the corresponding power spectrum for a) a beatwave scheme and b) an ultra-short pulse. In the beatwave scheme the difference frequency has been tuned to $\omega_1 - \omega_0 = \omega_\mathrm{p}$, and the ultra-short pulse has a Gaussian profile with a duration of $\sigma_\tau = \tau_\mathrm{p}/4$.

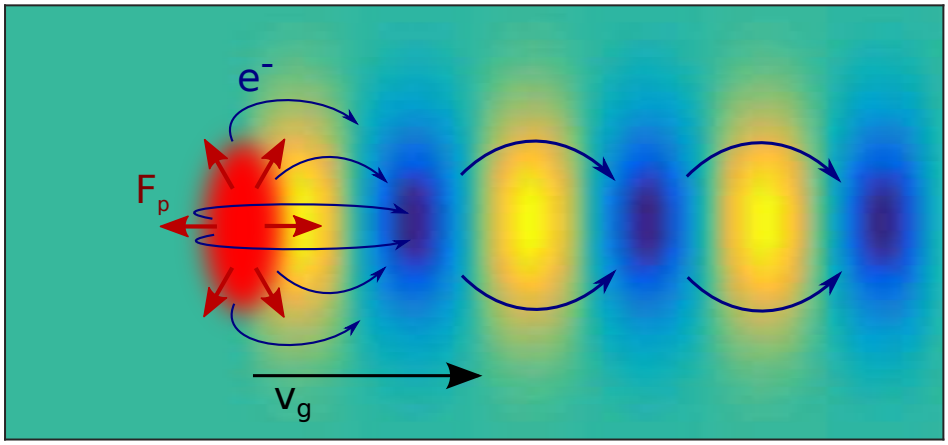


Figure 1.2: A cartoon of a short laser pulse (red) driving a plasma wakefield, shown in the frame of the laser pulse. Electron density is shown with regions of increased electron density in blue and decreased electron density in yellow, with electron trajectories marked with blue arrows.

and transversely, while ions remain approximately stationary. This sets up an extremely strong electric field which attracts electrons back to their original positions and they return to the ions on the timescale of the plasma period. In turn, this drives a strong plasma wave where electrons collectively oscillate at the plasma frequency. Figure 1.2 shows this situation in the frame of the laser pulse; in this picture electrons slip backwards with a velocity v_g while being attracted by regions of positive charge.

Producing a sufficiently short laser pulse requires both that the spectrum is broad-band, with $\sigma_\omega > \omega_p$ as before, and also that the phase profile is sufficiently flat, as for the wave picture. Conversely, a beating wave at a frequency ω_p creates intensity peaks with durations less than half the plasma period, producing a series of ultra-short pulses. The two pictures therefore correspond closely.

For a cold plasma, this plasma wave remains localised, as the Debye length and mean free path of the electrons limit the length scale over which information about the wave can travel. After the laser pulse has passed, the plasma wave therefore continues to oscillate until damped (normally by Landau damping). As the laser pulse moves with a speed $v_g \approx c$, successive regions of the plasma are set to oscillate, with a phase delay

$\omega_{\rm p} z/v_{\rm g}$. This therefore gives the appearance of a wave travelling through the plasma with a wavenumber $k = \omega_{\rm p}/v_{\rm g}$. This gives the wakefield a phase velocity equal to the group velocity of the driving laser pulse $v_{\rm p} = v_{\rm g}$.

We can study the wakefield in the frame of the driving laser pulse, using a co-moving time co-ordinate $\zeta = t - z/v_{\rm g}$, for a laser pulse moving along the positive z direction. The wakefield left behind a laser pulse can be described in the linear regime in one dimension by [35]:

$$\varphi(\zeta) = -\frac{\omega_{\rm p}}{4}\int_{\zeta}^{\infty} |a(\zeta')|^2 \sin\left[\omega_{\rm p}\left(\zeta - \zeta'\right)\right] {\rm d}\zeta', \tag{1.2}$$

where $\varphi(\zeta) = e\Phi(\zeta)/m_e c^2$ is the normalized electric potential in the plasma and $\mathbf{a}(\zeta) = e\mathbf{A_L}(\zeta)/m_e c$ is the normalized vector potential of the laser field. We can relate these to the electric fields of the plasma wave $E_z(\zeta)$ and of the laser $\mathbf{E_L}$ by:

$$\begin{aligned} E_z(\zeta) &= -\frac{m_e c^2}{e}\frac{{\rm d}\varphi}{{\rm d}z} \\ &\approx +\frac{m_e c}{e}\frac{{\rm d}\varphi}{{\rm d}\zeta} \end{aligned} \tag{1.3}$$

$$\begin{aligned} \text{and } \mathbf{E_L}(\zeta) &= -\frac{m_e c}{e}\frac{\partial \mathbf{a}}{\partial t} \\ &= -\frac{m_e c}{e}\frac{{\rm d}\mathbf{a}}{{\rm d}\zeta}, \end{aligned} \tag{1.4}$$

where we have used the standard result $\mathbf{E} = -\nabla\Phi - \partial\mathbf{A}/\partial t$. We can see that if the plasma wave is linear, the wake amplitude becomes $E_z \approx (m_e\omega_{\rm p}c/e)\varphi = \varphi E_{\rm wb}$ and $\varphi = 1$ corresponds to the cold wave-breaking limit. Similarly, if the laser field is sinusoidal, the field strength becomes $E_L \approx (m_e\omega_0 c/e)a$. The velocities of electrons in these electric fields can be estimated by:

$$m_e\dot{\mathbf{v}} = -e\mathbf{E} \tag{1.5}$$

$$\beta \equiv \frac{v}{c} \approx \frac{e}{m_e\omega c}E \tag{1.6}$$

When $a \sim 1$, electrons move at relativistic speeds within the laser pulse, their motions will no longer be sinusoidal, and their effective masses will increase by a factor γ, changing the plasma frequency. Similarly, when $\varphi \sim 1$, electrons will move at relativistic speeds in the plasma wave, implying that the plasma wave is no longer linear or sinusoidal. In most of this thesis we study plasma waves and laser drivers such that $a \ll 1$ and $\varphi \ll 1$.

In this regime, the plasma response to a particular intensity function is a linear combination of the response to its parts. For example, if we describe a laser pulse profile by a sequence of two pulses, as $|a(\zeta)|^2 = |a_1(\zeta - \zeta_1)|^2 + |a_2(\zeta - \zeta_2)|^2$, the plasma wave can be described by:

$$\begin{aligned}
\varphi(\zeta) &= -\frac{\omega_\mathrm{p}}{4}\int_\zeta^\infty \left[|a_1(\zeta' - \zeta_1)|^2 + |a_2(\zeta' - \zeta_2)|^2\right] \sin\left[\omega_\mathrm{p}\left(\zeta - \zeta'\right)\right] \mathrm{d}\zeta' \\
&= -\frac{\omega_\mathrm{p}}{4}\int_{\zeta-\zeta_1}^\infty |a_1(\zeta'')|^2 \sin\left[\omega_\mathrm{p}\left(\zeta - \zeta'' - \zeta_1\right)\right] \mathrm{d}\zeta'' \\
&\qquad -\frac{\omega_\mathrm{p}}{4}\int_{\zeta-\zeta_2}^\infty |a_2(\zeta''')|^2 \sin\left[\omega_\mathrm{p}\left(\zeta - \zeta''' - \zeta_2\right)\right] \mathrm{d}\zeta''' \\
&= \varphi_1(\zeta - \zeta_1) + \varphi_2(\zeta - \zeta_2),
\end{aligned} \tag{1.7}$$

where $\varphi_x(\zeta)$ is the plasma wave produced by a pulse profile $a_x(\zeta)$.

Let us consider a laser pulse with a Gaussian temporal profile described by $a(\zeta) = a_0 \exp\left[-\frac{1}{2}\left(\frac{\zeta}{\sigma_\tau}\right)^2\right]$. A sufficiently long time after the peak of the pulse has passed, $a_0(\zeta)$ becomes arbitrarily small and hence we can expand the lower integration limit to $-\infty$ and simplify the integration. The wakefield is then described by:

$$\varphi(\zeta) = -\frac{\sqrt{\pi}a_0^2}{2}\frac{\omega_\mathrm{p}\sigma_\tau}{2}\exp\left[-\left(\frac{\omega_\mathrm{p}\sigma_\tau}{2}\right)^2\right]\sin(\omega_\mathrm{p}\zeta) \tag{1.8}$$

We can write this in practical terms as:

$$\begin{aligned} E_z &= \frac{\sqrt{\pi}}{2}\frac{m_e\omega_\mathrm{p}c}{e}\left(\frac{eE_0}{m_e\omega_0 c}\right)^2\frac{\omega_\mathrm{p}\sigma_\tau}{2}\exp\left[-\left(\frac{\omega_\mathrm{p}\sigma_\tau}{2}\right)^2\right]\cos(\omega_\mathrm{p}\zeta) \\ &= \sqrt{4\pi}\frac{e}{m_e\epsilon_0 c^2}\frac{I_0}{\omega_0^2\sigma_\tau}\left(\frac{\omega_\mathrm{p}\sigma_\tau}{2}\right)^2\exp\left[-\left(\frac{\omega_\mathrm{p}\sigma_\tau}{2}\right)^2\right]\cos(\omega_\mathrm{p}\zeta), \end{aligned} \tag{1.9}$$

where E_0 and I_0 are the peak electric field strength and laser intensity corresponding to the normalized vector potential a_0 and we have used $I_0 = \epsilon_0 E_0^2 c/2$ for linearly polarized light.

There are therefore two ways of increasing the wakefield amplitude: increasing the intensity of the driving laser pulse, and matching its duration to the plasma period such that $\omega_\mathrm{p}\sigma_\tau \sim 1$. Actually, the intensity of the pulse is proportional to the energy in the pulse as $I_0 \propto E/\sigma_\tau$, and so shorter pulses will always result in a higher wakefield amplitude for the same supplied energy. If the peak intensity is fixed, however, there is an optimal plasma frequency, where $\omega_\mathrm{p}\sigma_\tau = 2$, or $\tau_\mathrm{p} = \pi\sigma_\tau$.

We can compare this to the expression derived for a cylindrically symmetric plasma wave, given by [36]. If the laser pulse is cylindrically symmetric and bi-Gaussian, i.e. it can be described as $I_L(\zeta, r) = I_0\exp(-\zeta^2/2\sigma_\tau^2)\exp(-r^2/2\sigma_r^2)$, the wakefield produced is also cylindrically symmetric and can be described by:

$$E_z(r,\zeta) = E_\mathrm{max}\exp\left[-\left(\frac{r}{\sigma_r}\right)^2\right]\cos\left(\omega_\mathrm{p}\zeta\right) \tag{1.10}$$

$$E_r(r,\zeta) = E_\mathrm{max}\left(\frac{\lambda_\mathrm{p} r}{\pi\sigma_r^2}\right)\exp\left[-\left(\frac{r}{\sigma_r}\right)^2\right]\sin\left(\omega_\mathrm{p}\zeta\right) \tag{1.11}$$

$$E_\mathrm{max} = \frac{2I_0\sqrt{\pi}}{ec^2 n_c\sigma_\tau}\left(\frac{\omega_\mathrm{p}\sigma_\tau}{2}\right)^2\exp\left[-\left(\frac{\omega_\mathrm{p}\sigma_\tau}{2}\right)^2\right], \tag{1.12}$$

where $n_c = \epsilon_0 m_e\omega_0^2/e^2$ is the critical density for light of frequency ω_0.

These fields are shown in Figure 1.3. We can see that in addition to the longitudinal field E_z, the wakefield has a transverse field E_r, which increases in magnitude off axis. Close to the axis, this restoring force is proportional to r. The wakefield therefore has a

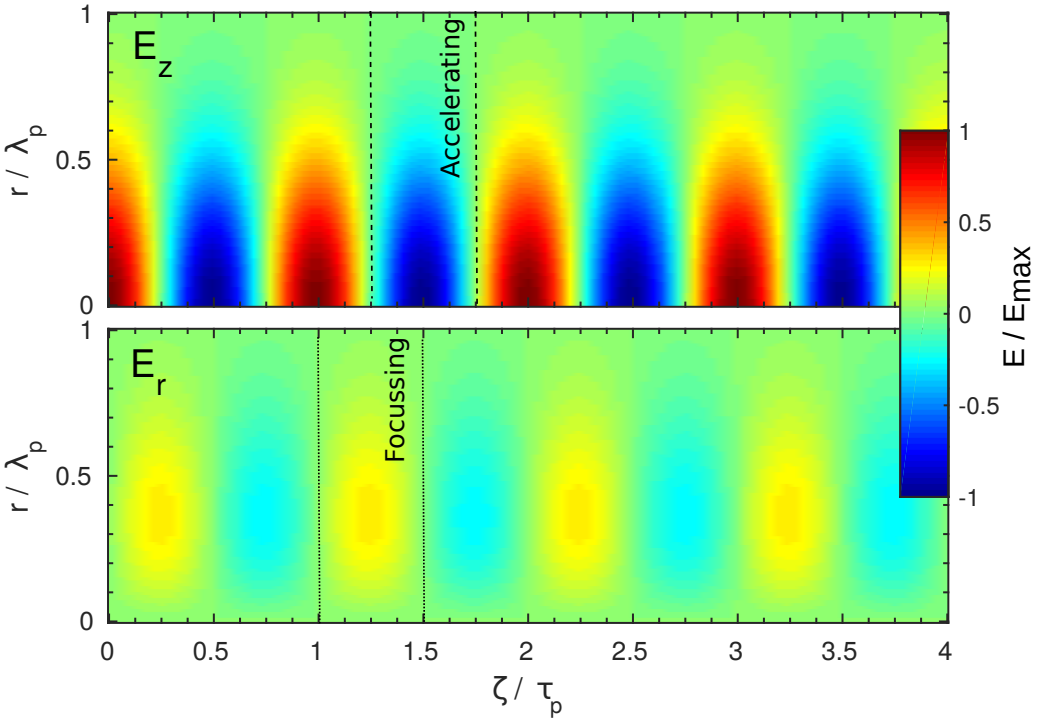


Figure 1.3: Electric fields of a linear wakefield in space and co-moving time, shown for the longitudinal field E_z (top) and the radial field E_r (bottom). Both are shown using the same colour scale, with amplitudes relative to the maximum field strength $E_{\max} = E_z(0,0)$. This plot corresponds to a laser pulse with $\sigma_r = \lambda_p/2$.

phase which accelerates electrons along positive z ($E_z < 0$ for $i + 1/4 < \zeta/\tau_{\rm p} < i + 3/4$ and integer $i \in \mathbb{Z}$) and a phase which focusses electrons onto the axis in the manner of a simple harmonic oscillator ($E_r/r > 0$ for $i < \zeta/\tau_{\rm p} < i + 1/2$). There is therefore a region one quarter of the plasma period in duration where electrons would be both accelerated and retained on axis.

1.2.2 In Practice

There is a strong motivation for using ultra-short laser pulses strongly focussed to produce laser electric fields on the scale of $a_0 \sim 1$ and plasma wakefields on the scale of $E_z \sim E_{\rm wb}$. However, there are several disadvantages to operating in this regime. Firstly, at very high laser intensities, the plasma wave produced is no longer either linear or quasi-linear. Instead, the plasma enters the bubble regime [32], [37], with

the region behind the laser pulse almost entirely evacuated of electrons. While this provides extremely high peak electric fields at the back of the bubble, it also leads to several non-linear effects. For instance, some electrons slightly inside the bubble's rim follow trajectories that bring them into the accelerating phase at the back of the bubble. This self-injection process happens continually as the laser pulse traverses the acceleration length, generally leading to electron bunches with larger energy spread and larger transverse emittance than possible with controlled injection, particularly in a low energy 'tail'.

Even before entering the bubble regime, using highly intense laser pulses can be disadvantageous. As light waves in plasma have both an anomalous group velocity dispersion (with higher frequencies travelling with a higher group velocity) and a positive Kerr non-linearity (with the refractive index increasing with laser pulse intensity due to relativistic effects), the laser pulses are prone to the modulational instability, with deformations to the pulse envelope growing rapidly. For instance, if any portion of the laser pulse becomes red-shifted (which, as we will see later, is a natural consequence of driving a plasma wave), it will travel slower than the rest of the pulse. This disperses the laser pulse, and the portion that falls behind is left with a lower intensity and hence experiences a lower refractive index. In turn, this reduces the group velocity further and further disperses the pulse. This causes erosion of the laser pulse and reduces the length over which a wakefield can be produced. These effects do not rule out LWFA in the non-linear regime and there is a large body of research into either mitigating or better controlling them, but I shall leave that discussion to others, such as ref. [33].

Secondly, and more importantly, wakefield acceleration is limited by the laser systems which drive the plasma wave. Producing ultra-short and high intensity pulses currently relies on Titanium-doped sapphire crystals. While these support both a very large bandwidth and high intensity laser pulses they suffer from an intrinsically low quantum efficiency and are difficult to cool. The Ti:sapphire crystals must be pumped at ≈ 490 nm, and emit at ≈ 800 nm [38]. That means that even in the ideal case only around 60%

of the energy supplied into the crystal is re-emitted and the remainder is lost as heat. Furthermore, none of the lasers at 490 nm which are sufficiently powerful to pump the crystal are yet very efficient, and hence either flash lamps are commonly used, or else frequency doubled light from Neodymium:glass lasers, which are themselves pumped by flash lamps. Of the broadband spectrum of flash lamps only a tiny proportion is useful for pumping the crystal and the rest of the light is wasted. In total, less than 1% of the energy supplied to the laser system reaches the final laser pulse. Worse, all of the heat supplied to the Ti:sapphire crystal must be extracted, and as the crystal is a bulk medium which must be several millimeters long, the cooling rate is slow and the repetition rates for Joule level systems are limited to the scale of 1-10 Hz. Compared to light sources such as synchrotrons and free electron lasers which can run at repetition rates of kHz, Ti:sapphire based laser systems will struggle to be competitive.

One solution is to use other laser systems which are more efficient and have high repetition rates, such as diode pumped fibre and thin disk systems. These systems, with lower peak powers, are not yet capable of driving strong wakefields with a single pulse. Some groups have therefore suggested coherently combining many laser pulses, in order to generate an intense laser pulse with both high average power and high peak power [27], [39]. This is a difficult task, as the laser pulses must retain coherence and be timed to an accuracy significantly better than $2\pi/\omega_0 \sim 1\,\mathrm{fs}$; as more pulses are used this requirement only becomes more difficult.

My group in Oxford has proposed instead using successive pulses separated by multiples of the plasma period to resonantly excite a linear or quasi-linear wakefield [28], [29]. In this way, energy is stored in the plasma, rather than in the lasing medium, and concerns about cooling can avoided. This multiple pulse laser wakefield acceleration scheme could operate at much higher efficiencies and repetition rates, on the scale of 10% wall-plug efficiency and several kHz repetition rate, depending on the laser systems used. The scheme also naturally lends itself to work on controlled injection [40], and on energy recovery, extracting energy from the plasma with laser pulses in anti-phase with

the resonance.

There are, however, several limiting factors on multiple pulse excitation that were extensively studied during work on laser beatwave drivers. Firstly, although this thesis focuses on low amplitude linear wakefields with both $\phi \ll 1$ and $a \ll 1$, electrons will still experience a mass increase due to their relativistic motion, as $m = \gamma m_e$. This will slightly increase the plasma period over time and so uniform pulse trains will eventually fall out of resonance with the plasma wave, with a spacing between pulses less than the true plasma period. Secondly, ions, although experiencing a much weaker ponderomotive force than electrons, experience a strong electric field from the charge displacement in the wakefield. This force almost completely cancels out, with ions experiencing as much of a field in one direction as the other, but as with the ponderomotive force, when the amplitude of the plasma wave is increasing, a slow drift component will eventually be established which will disrupt the plasma wakefield.

Even without these limitations, achieving resonant excitation of a plasma wave by a train of many laser pulses is not trivial. While individual laser pulses need not be coherent, the pulses must be timed to arrive at the correct phase of the plasma wave. This condition is significantly less stringent than timing pulses to within a laser period, as the plasma period is on the scale of $\tau_\mathrm{p} \sim 100\,\mathrm{fs}$, but careful consideration is necessary. Chapter 2 studies the accuracy to which pulse spacing must be maintained in a long train of laser pulses in order to effectively drive a plasma wakefield. I outline a basic theory for the constructive addition of these plasma waves and compare it to particle-in-cell simulations of the MP-LWFA scheme before discussing another possible approach for addressing these limitations, using autoresonance.

1.3 Diagnosing

1.3.1 In Practice

When measuring plasma waves produced by laser drivers, it is common to use optical measurements, which can have pulse durations on the femtosecond scale. As a laser drives a wakefield it loses energy to the plasma, leading to shifts in the spectrum of the pulse [41]. The size of these shifts is determined by the amplitude of the wakefield and hence the accelerating structure can be measured. Similarly, a probe pulse co-propagating with the wakefield will acquire energy from or lose energy to the plasma wave, resulting in blueshifts or redshifts respectively, depending on the phase of the wakefield with which the probe pulse co-propagates; this has been described as 'photon acceleration' [42]. By scanning the delay between the pulse driving the wakefield, the pump pulse, and the probe pulse, the structure of the wakefield can be mapped out. Figure 1.4a shows the spectral shifts for a short duration pulse co-propagating with a plasma wave, experiencing a blue-shift as a result.

The spectral shift can also be thought of as a phase change, and the most sensitive way of measuring this phase change is with interferometry. If a reference laser pulse precedes the pump pulse it measures no phase change, and interfering it with a probe pulse co-propagating with the wakefield reveals the phase change due to the wakefield at a given pump-probe delay. To create this interference pattern, the plasma is imaged onto the entrance slit of a spectrometer, and the interference takes place in the spectral domain, giving rise to the name Frequency Domain Interferometry (FDI) [43]. This allows each measurement to record a one-dimensional spatial density profile of the plasma wave at a given time. Successive measurements build up a picture of the plasma wave in both space and time. Furthermore, if the plasma density varies sufficiently slowly over the duration of the probe pulse, Fourier transforming the signal measured on the spectrometer and isolating the interference sideband yields the temporal phase profile in the vicinity of the probe pulse [44].

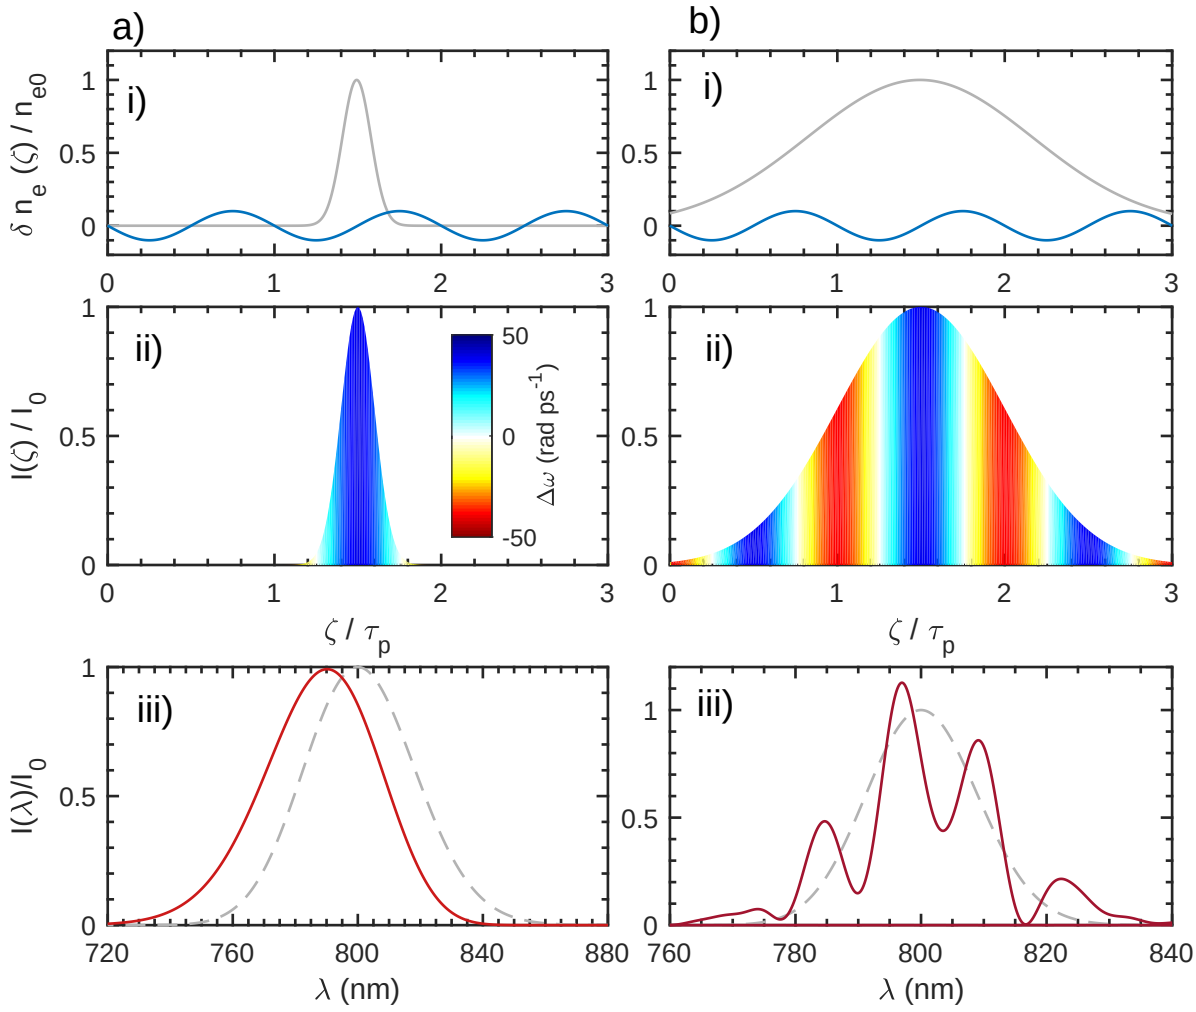


Figure 1.4: Simulated shifts in the spectrum of a laser pulse co-propagating with a plasma wave for a) a short pulse, as in FDI and b) a longer duration chirped pulse, as in FDH and TESS. i) The electron density profiles shown against time lead to the spectral shifts shown in ii), where the colour reflects the direction and scale of the change in frequency. The colour scale is the same in both plots. iii) The resulting spectrum of the laser pulse after co-propagating with the density structure for 3 millimeters, compared to the original spectrum (grey dashed). In both cases the laser and plasma wave are travelling to the left. The plasma density was chosen as $n_e = 10^{18}\,\mathrm{cm}^{-3}$.

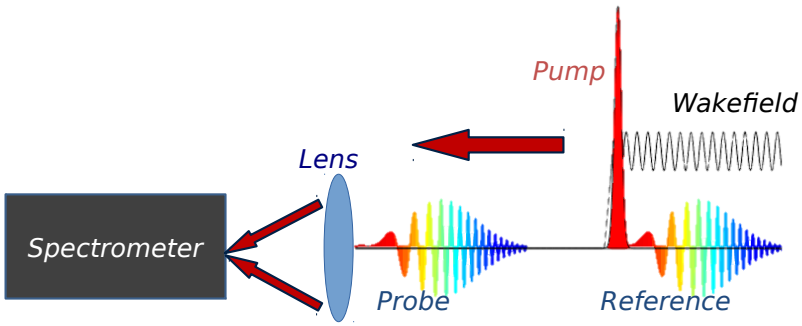


Figure 1.5: A cartoon of the experimental scheme for Frequency Domain Holography and Temporally Encoded Spectral Shifting. A pump pulse, a reference pulse and a probe pulse all co-propagate through a gas before the probe and reference pulses are focussed onto a spectrometer, where they interfere. The probe and reference pulses are stretched to a duration longer than a few plasma periods, and have a different wavelength to the pump pulse, typically from frequency doubling.

It has also been demonstrated that this vicinity can be extended and many plasma periods in the wakefield can be probed in a single shot by using a chirped and stretched probe pulse [45]. For these pulses, the local frequency changes over the duration of the pulse, recording information about a different part of the plasma wave at a different laser wavelength. Figure 1.4b shows the shift in the spectrum of the probe pulse that occurs for a long chirped pulse co-propagating with a plasma wave. When this pulse is interfered with a reference pulse, which is also chirped and stretched, the phase difference allows complete reconstruction of the plasma wave against time and one spatial dimension. Due to the similarity with spatial holograms this has been called Frequency Domain Holography (FDH). This scheme is outlined in Figure 1.5, with pump, probe and reference pulses all co-propagating. Usually the probe and reference pulses are produced from a pick-off of the pump beam, before frequency doubling and stretching them. A method has been recently been proposed to reconstruct both spatial dimensions in addition to the temporal dimension, by using a random mask and a compressed sensing reconstruction algorithm in a so-called 3D spectrometer [46].

The phase change of a stretched laser pulse co-propagating with a plasma wave is comparable to the frequency modulation of radio waves that has been familiar for decades. Satellite copies of the original spectrum are created, separated from the original spectrum by multiples of ω_p. This process, Thomson scattering of light in a plasma, has been extensively studied (e.g. [47]), and the spacing and amplitude of the satellites are used to determine the plasma wave's amplitude and frequency. However, when the original laser pulse has a broadband spectrum and the plasma density is low, the satellite spectra are no longer distinct from the original spectrum. Figure 1.6 shows the resulting spectra for a long laser pulse with a narrow spectral width, and for an ultra-short laser pulse with a broadband spectrum. Whereas the forward and backward Thomson scattering satellites (known as the Stokes and anti-Stokes peaks) are easily visible for a narrowband spectrum, for a broadband spectrum they overlap, causing spectral interference. In this case, recent work has demonstrated that the amplitude and frequency of the plasma wave can instead be extracted by Fourier transforming the signal measured on the spectrometer and looking at the location of the satellites in the temporal domain [31]. This process, called Temporal Encoded Spectral Shifting, allows both sensitive measurement of plasma wakefields and rapid analysis.

All of these optical measurements come down to the same thing –an optical probe pulse is changed by co-propagating with a plasma wave in a well defined way, which depends on the density structure and the gradient of the density structure in the vicinity of the probe pulse (see e.g. [30], [48]). The different terminology and ways of describing this process may seem unnecessarily confusing, but do help to distinguish experimental layouts and analysis techniques which differ in sometimes subtle ways. In order to understand TESS, for instance, one must understand the difference between the phase of the probe in the spectral and temporal domains. I discuss below the phase change of an optical probe pulse resulting from co-propagating with a plasma wave and describe how to map between the spectral and temporal domains.

There are many other techniques of measuring plasma waves and the radiation they

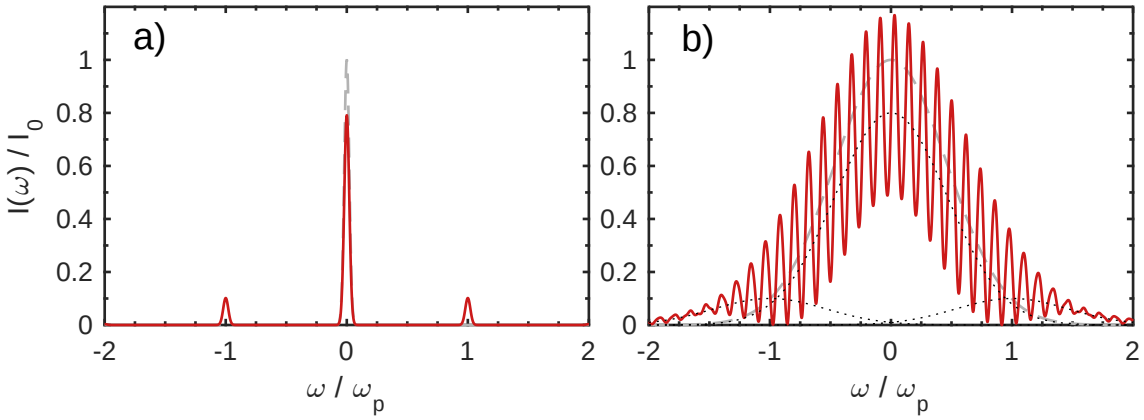


Figure 1.6: Simulated spectra of laser pulses after co-propagating with a plasma wave when the spectral width of the pulse is a) much smaller than the plasma frequency, and b) comparable to the plasma frequency. Both are compared to the original spectrum (grey dashed), and in b) the expected satellite peaks are also shown (black dotted). Both pulses have FWHM durations of one picosecond, but whereas the narrow-band spectrum has a flat phase, the broadband pulse is strongly chirped. The same colour scale is used for both a)i) and b)i).

produce, many of which are described in ref. [49]. This includes geometries where the probe pulse does not travel co-linearly with the pump pulse, such as shadowgraphy [50], Faraday rotation [51], and shallow-angle Thomson scattering [52], but I shall not discuss these in this thesis. This is largely because of the limited sensitivity of these techniques at low densities; the phase change of an optical probe co-propagating with a density structure can be described by $\phi(\zeta) = (\omega_p^2 L/2\omega_0 c)\delta n_e/n_{e0}$, where L is the length over which the structure is probed. As the plasma density drops, this phase change becomes increasingly small. For transverse probes, L is typical small, on the scale of tens of microns, and hence the total phase change becomes too small to measure. Similarly, the scattered power in Thomson experiments is related to the wakefield amplitude by $P_s \propto (\delta n_1/n_{e0})^2 n_{e0}^2)L$.

There is another way of increasing the sensitivity of probe pulses, by increasing the wavelength of the probe laser. If the frequency can be reduced as $\omega_0 \propto \omega_p^2$, the phase change experienced in a given plasma will remain the same. This has the downside of also reducing the spatial resolution possible, as this is proportional to the probe

wavelength. While the scale of features in a wakefield increases as the plasma density falls, either the sensitivity or the resolution will eventually become insufficient, as for a spatial resolution δx at a constant sensitivity, $\delta x/\lambda_{\mathrm{p}} \propto \omega_{\mathrm{p}}/\omega_0 \propto 1/\omega_{\mathrm{p}}$. There is a significant range over which transverse probing will still be possible if laser pulses with sufficiently long wavelength and short duration can be produced, but the longitudinal probe techniques described above, with a longer probing length, manage to retain both sensitivity and spatial resolution.

Of the different longitudinal probing techniques, I work with TESS, as the most recent and probably the most sensitive technique, and the most well suited for high repetition rates. However, the TESS signal expected for a given wakefield frequency and amplitude has until now been derived only for probe and reference pulses which a) are identical but separated in time, and b) have spectral profiles which are perfectly Gaussian. While this gives analytic calculations of the expected TESS signal, it is a poor approximation for many experimental set ups. In this thesis I therefore generalize TESS to probe and reference pulses with arbitrary spectral profiles. Additionally, TESS was derived for purely sinusoidal wakefields in the linear regime. As the wakefield amplitude increases, electrons in the plasma wave become relativistic and the plasma wave is no longer sinusoidal in the quasi-linear regime. In chapter 3, I therefore extend the TESS analysis to quasi-linear plasma waves and look at some of the limitations on using it for non-linear plasma waves. First, however, I outline some of the theory behind the diagnostics I have just described.

1.3.2 In Theory

One of the ways to start describing frequency and phase shifts in plasmas is through the dispersion relation of electro-magnetic waves. For cold and collisionless plasmas in

the absence of an external magnetic field, this can be described by [53]:

$$\omega(k)^2 = \omega_{\mathrm{p}}^2 + k^2 c^2, \tag{1.13}$$

for a wavenumber k. This leads to a phase velocity described by:

$$v_\phi \equiv \frac{\omega(k)}{k} = c\left[1 - \left(\frac{\omega_{\mathrm{p}}}{\omega}\right)^2\right]^{-\frac{1}{2}} \tag{1.14}$$

$$= \frac{c}{n(\omega)}, \tag{1.15}$$

where n is the refractive index, which is less than one and decreases with increasing plasma frequency. This implies that light travelling through a more dense plasma with a higher plasma frequency advances in phase more rapidly than light travelling through a less dense plasma.

Light travelling through a uniform region of plasma therefore acquires a phase described by [30]:

$$\Delta\phi = \int k \, \mathrm{d}z \tag{1.16}$$

$$= \int \frac{\omega n}{c} \, \mathrm{d}z \tag{1.17}$$

$$= \frac{\omega}{c} \int \left[1 - \left(\frac{\omega_{\mathrm{p}}}{\omega}\right)^2\right]^{\frac{1}{2}} \, \mathrm{d}z \tag{1.18}$$

$$= \frac{\omega L}{c}\left[1 - \left(\frac{\omega_{\mathrm{p}}}{\omega}\right)^2\right]^{\frac{1}{2}} \tag{1.19}$$

If a laser pulse travels through a region of constant density, every part of the pulse experiences the same density along its path and is subject to the same phase change. If a pulse co-propagates with a density structure, however, different parts of the pulse experience different densities, and hence experience different phase changes. For instance, if the front of the pulse co-propagates with a region of high density and the back of the pulse co-propagates with a region of low density, the front experiences a larger phase

change than the back. This can be described by:

$$\Delta\phi(\zeta) = \frac{\omega L}{c}\left[1 - \left(\frac{\omega_\mathrm{p}(\zeta)}{\omega}\right)^2\right]^{\frac{1}{2}} \tag{1.20}$$

$$\frac{\mathrm{d}\Delta\phi(\zeta)}{\mathrm{d}\zeta} = -\frac{1}{2}\frac{\omega L}{c}\frac{2\omega_\mathrm{p}}{\omega^2}\frac{\mathrm{d}\omega_\mathrm{p}}{\mathrm{d}\zeta}\left[1 - \left(\frac{\omega_\mathrm{p}(\zeta)}{\omega}\right)^2\right]^{-\frac{1}{2}} \tag{1.21}$$

$$\Delta\omega(\zeta) = \frac{\omega_\mathrm{p} L}{\omega c n}\frac{\mathrm{d}\omega_\mathrm{p}}{\mathrm{d}\zeta}, \tag{1.22}$$

where we have used $\omega \equiv -\mathrm{d}\phi/\mathrm{d}\zeta$ for the co-moving time ζ.

In this way, the frequency of the pulse is shifted by co-propagating with a plasma wave. If the density increases along the pulse, the laser pulse becomes blue-shifted; this situation commonly arises when the laser pulse is ionizing a plasma. Another way of looking at this situation is that the phase velocity is faster for the back of the pulse, where the density is higher, than the front. The waves therefore become compressed, leading to a blue-shift. Wilks et al. [42] refer to this situation as 'photon acceleration', as the photon energy is slightly increased, but it is the same process of phase changes described by e.g. Esarey et al [30]. A short laser pulse driving a linear wakefield expels electrons from the region behind it and hence sits at a region of decreasing electron density. It therefore acquires a red-shift, losing energy to the plasma wave.

For longer laser pulses co-propagating with a plasma wave, both the phase and the frequency are periodically increased and decreased, resulting in a modulation of the spectrum. If the plasma wave is sinusoidal over the duration of the laser pulse, so too is the change in phase, and we can describe the electric field of the laser pulse as:

$$\begin{aligned} E(\zeta) &= |E(\zeta)| e^{i\Delta\phi\sin(\omega_\mathrm{p}\zeta)} \\ &= |E(\zeta)| \sum_{k=-\infty}^{\infty} J_k(\Delta\phi) e^{ik\omega_\mathrm{p}\zeta}, \end{aligned} \tag{1.23}$$

where we have used the Jacobi-Anger transform to describe the phase change in terms of a sum of carrier waves.

A sinusoidal modulation in the phase corresponds to the creation of new peaks in the frequency spectrum, as for Frequency Modulated (FM) radio. As in FM radio signals, the phase change can be extracted from the carrier waves by a Fourier transform, in a process comparable to TESS. The creation of sidebands in the frequency spectrum, is familiar in plasma physics as Thomson scattering [52]. In this situation a plasma wave scatters light at a frequency ω to frequencies at $\omega \pm \omega_{\mathrm{p}}$, which can be thought of as Doppler shifts from photons colliding with the rapidly oscillating electrons. The plasma frequency is then inferred from the separation of the sidebands, while the wave amplitude is inferred from the amplitude of the sidebands. Again, two different diagnostic techniques are based upon the same physical process described in two different ways.

Up until this point, I've described the phase and frequency in temporal space as $\phi(\zeta)$ and $\omega(\zeta)$. However, in most of the diagnostics I have described, whether photon acceleration, Thomson scattering, or FDH, we measure the spectrum of a laser pulse, and hence the relevant quantity is the *spectral* phase $\psi(\omega)$. This is not the same as the temporal phase, and does not relate directly to it. Instead, they are related through the Fourier transforms of the electric field in the temporal domain $E(\zeta)$ and the spectral domain $\mathcal{E}(\omega)$:

$$\mathcal{E}(\omega) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} E(\zeta) e^{-i\omega\zeta} \, \mathrm{d}\zeta \tag{1.24}$$

$$E(\zeta) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} \mathcal{E}(\omega) e^{+i\omega\zeta} \, \mathrm{d}\omega \tag{1.25}$$

$$\text{where } \mathcal{E}(\omega) = |\mathcal{E}(\omega)| e^{i\psi(\omega)} \tag{1.26}$$

$$\text{and } E(\zeta) = |E(\zeta)| e^{i\phi(\zeta)} \tag{1.27}$$

Without knowing the form of $|E(\zeta)|$ or $|\mathcal{E}(\omega)|$ it is impossible to find $\psi(\omega)$ from $\phi(\omega)$ or vice versa.

It is possible to expand both $\phi(\zeta)$ and $\psi(\omega)$ as Taylor series, as below:

$$\phi(\zeta) \approx \phi^{(0)} + \phi^{(1)}(\zeta - \zeta_0) + \frac{1}{2}\phi^{(2)}(\zeta - \zeta_0)^2 + \dots \tag{1.28}$$

$$\psi(\omega) \approx \psi^{(0)} + \psi^{(1)}(\omega - \omega_0) + \frac{1}{2}\psi^{(2)}(\omega - \omega_0)^2 + \dots, \tag{1.29}$$

where $\phi^{(n)} = \mathrm{d}^n\phi(\zeta)/\mathrm{d}\zeta^n$ and $\psi^{(n)} = \mathrm{d}^n\psi(\zeta)/\mathrm{d}\zeta^n$. We can identify the first order terms as the central frequency $\phi^{(1)} = -\omega_0 = -\omega(\zeta = \zeta_0)$, and as the location of that frequency $\psi^{(1)} = \zeta_0 = \zeta(\omega = \omega_0)$. We can also identify the second order term of the spectral phase $\psi^{(2)} = \mathrm{d}\zeta/\mathrm{d}\omega$ as the Group Delay Dispersion (GDD), which describes the chirp of a laser pulse. A non-zero GDD implies that different frequencies occur at different times, with the pulse duration increasing as a result. GDDs larger in amplitude imply more strongly chirped pulses with longer durations. In normal dispersion, GDD is positive, with the lower (red) frequencies of a pulse occurring earlier in time than the higher (blue) frequencies. Knowledge of this chirp is essential for the FDH and TESS diagnostics.

Over the rest of the thesis, I describe the TESS and FDH diagnostics in terms of the phase changes from plasma in both the temporal and spectral domains. With this theory outlined, I can describe how it affects the signals measured in the TESS diagnostic in chapter 3.

1.4 Guiding

1.4.1 Problem

The third challenge I approach in this thesis is guiding laser pulses over long distances. This is an important step, as lasers naturally diffract at a rate described by:

$$w(z) = w_0\sqrt{1 + \left(\frac{M^2 z}{Z_R}\right)^2} \tag{1.30}$$

$$\text{for } Z_R = \frac{\pi w_0^2}{\lambda}, \tag{1.31}$$

where Z_R is the Rayleigh range, w_0 is the minimum beam waist at focus, defined as the radius for the intensity to decrease by a factor of $e^{-2} \approx 0.14$ from its peak on-axis, and M^2 is the beam quality factor, which is limited as $M^2 \geq 1$.

In order to produce the high intensity beams typically used to drive laser wakefield acceleration, the laser must be focussed to small spot sizes, and the smaller the laser beam is at focus, the faster it diverges. For an 800 nm beam focussed to $w_0 = 30\,\mu$m, the Rayleigh range is just $Z_R \approx 3.5\,$mm. The dephasing length, in contrast, is on the scale of $L_{\rm dp} \sim 30\,$mm at $n_e = 10^{18}\,{\rm cm}^{-3}$ and increases rapidly with decreasing density, as $L_{\rm dp} \sim \lambda_{\rm p}^3/\lambda^2$. If the spot size is comparable to the plasma wavelength $w_0 \sim \lambda_{\rm p}$, the ratio of the dephasing length to the Rayleigh range is $L_{\rm dp}/Z_R \sim \lambda_{\rm p}/\lambda \gg 1$. Increasing the Rayleigh range also requires decreasing the intensity at focus proportionately; for a fixed energy in the beam E the peak intensity is $I_0 \propto E/w_0^2$. In turn, the wakefield amplitude and accelerating field is proportional to intensity, leading to an energy gain restricted by the laser energy as $\Delta W \propto E_z Z_R \propto I_0 w_0^2 \propto E$. No matter how tightly or loosely the beam is focussed, the energy gain of the stage cannot be increased if the beam diffracts. If the diffraction of the beam can be countered, however, it is possible to dramatically increase both the length of the acceleration stage and the energy gain.

As described earlier, the refractive index of a plasma $n = \sqrt{1 - (\omega_{\rm p}/\omega)^2} < 1$, and is smaller for higher density plasmas. It is possible to form waveguides by increasing the refractive index of a structure on axis, such that rays moving off axis are refracted to shallower angles. If there is a step index change in refractive index, such as from a low density plasma encased by a higher density sheath, all incident beams shallower than a critical angle cannot refract and are reflected back towards the axis, as in a simple optical fibre shown in Figure 1.7. Plasmas are particularly interesting as the dispersion relation shown in eqn. (1.13) implies that the plasma becomes opaque to light at a frequency lower than the plasma frequency[†]. As the plasma density approaches critical, the phase velocity $v_\phi = c/n$ approaches infinity while the group velocity $v_{\rm g} = cn$ approaches 0. In

[†]Thank you ionosphere...

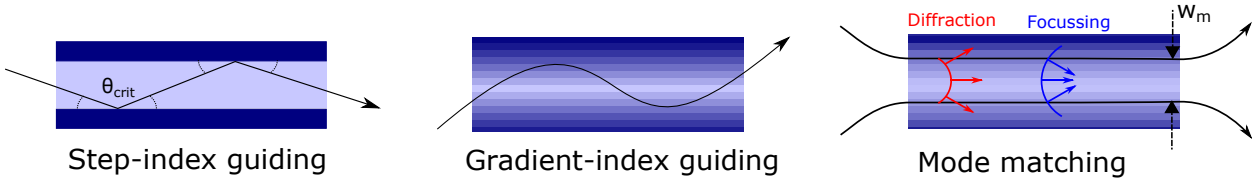


Figure 1.7: Cartoons of rays passing through step-index channels and gradient-index channels, compared to a perfectly matched gradient-index channel, where the beam remains not only confined but also free from divergence.

turn, this implies that if the component of the wavevector perpendicular to the plasma boundary is sufficiently small, the light is reflected off the boundary. This critical density is described by [53]:

$$\begin{aligned} n_{\text{crit}} &= \frac{m_e \epsilon_0 (k_\perp v_\phi)^2}{e^2} \\ &= \frac{m_e \epsilon_0 \omega^2}{e^2} \sin^2 \theta, \end{aligned} \tag{1.32}$$

for an angle of incidence θ. For shallow angles $\theta \ll 1$ it is possible to create total reflection at even very low plasma densities. For instance, at $\theta = 1°$, the critical density for an 800 nm beam is just $n_e \approx 5 \times 10^{17}\,\text{cm}^{-3}$. Conversely, for a given plasma density, the critical angle is described by $\sin\theta_{\text{crit}} = \frac{\omega_{\text{p}}}{\omega}$.

These type of waveguides will transmit light, but the quality of the beam inside the waveguide is by no means assured. If the change in refractive index is gradual, however, the profile can be tuned so as to exactly cancel the divergence of the laser pulse, slowing down the wavefront near the axis just enough such that it travels at the same speed as the wavefront at the edge. Figure 1.7 shows a cartoon of a channel meeting this matching condition compared to step- and gradient-index channels without matching. For a Gaussian beam with a matched spot size of w_m, the ideal plasma density profile is described by [54], [55]:

$$n_e(r) = n_e(0) + \frac{1}{\pi r_e w_m^4} r^2, \tag{1.33}$$

where $r_e \approx 3 \times 10^{-15}\,\mathrm{m}$ is the classical electron radius.

This parabolic plasma density channel will support bound Laguerre-Gaussian modes, with the lowest order mode having a constant spot size w_m. In practice, no plasma will attain a perfect parabolic density profile, particularly as it would have to extend to densities beyond the plasma's critical density. Instead, a beam can be mostly guided, only gradually losing energy to outside the channel, so long as the density is low on-axis and increases with radius in the region of greatest intensity. In practice, this usually means fitting the density profile to an expression of the form described by [54]:

$$n_e(r) = n_e(0) + \frac{1}{\pi r_e w_m^2}\left(\frac{r}{w_m}\right)^{\alpha} \text{ for } r < r_{\max} \tag{1.34}$$

If it is possible to find an exponent α such that $w_m < r_{\max}$, the rate at which the beam loses energy to outside the channel is low, and guiding is likely to be possible. This implies that in order to guide a laser pulse effectively, the density at the edge of the channel must be greater than around $n_e(0)+\Delta n_e$, where $\Delta n_e = 1/\pi r_e w_m^2$. For a matched spot size of $w_m \approx 30\,\mu\mathrm{m}$ this is $\Delta n_e \approx 10^{17}\ \mathrm{cm}^{-3}$, which is more difficult to achieve in low density plasma channels. In low density channels more generally, the change in the refractive index across the density profile becomes smaller, making it significantly more difficult to counter the divergence of the laser pulse. This typically leads to larger matched spot sizes or no guided modes at all, reducing the intensity with which an LWFA driver could drive a wakefield or the length over which this intensity can be maintained. If $w_m \approx \lambda_\mathrm{p}/2$ in an LWFA experiment, however, we require $\Delta n_e \approx 4/\pi r_e \lambda_\mathrm{p}^2 \approx 4n_e(0)/\pi^2$. We see that the important quantity becomes the *relative* height of the plasma channel.

To find out more about the profile of the beam within channels with arbitrary density profiles, it is necessary to solve the wave equation for the propagation of electromagnetic waves in the plasma:

$$\left[v_\phi^2(r,z,t)\nabla^2 - \partial_t^2\right]\mathbf{E}(r,z,t) = \mathbf{0} \tag{1.35}$$

We assume a cylindrically symmetric channel described by $v_\phi(r) = c/n(r)$ and an electric field of the form $\mathbf{E} = \hat{\mathbf{z}} E_0 f(r) \exp\left[i\left(\omega t - \kappa z\right)\right]$, which describes the laser pulse as a continuous wave propagating along positive z with a wavenumber κ. Naively this wavenumber would be $k(r) = \omega/v_\phi(0)$ but inside the waveguide the wavenumber is generally lower. The wave equation reduces to a Helmholtz equation of the form:

$$f''(r) + \frac{f'(r)}{r} + \left[k(r)^2 - \kappa^2\right] f = 0. \tag{1.36}$$

It is then possible to solve the Helmholtz equation numerically over a range of values for κ to obtain the shape of the possible modes $f(r)$. By finding the fundamental mode with maximal amplitude on axis it is then also possible to determine the 1/e attenuation length of the fundamental, $L_{\mathrm{att}} = \left[\Im(-2\kappa)\right]^{-1}$. If a laser pulse is coupled into the waveguide such that a majority of its intensity matches the fundamental mode, this attenuation length describes how long the channel can be before losses become significant.

With this theory outlined in brief, I shall next describe how plasma waveguides are produced in practice, with a focus on creating plasma waveguides which can guide laser pulses over long distances at low plasma densities and are capable of operating at high repetition rate.

1.4.2 Solutions

Before I describe work on producing guiding structures for laser wakefield acceleration, I will take a moment to describe some non-linear processes often used to avoid the need for a plasma waveguide in LWFA experiments. Firstly, as a laser pulse drives a strong plasma wave, it expels electrons from the region immediately behind it through the ponderomotive force, as described earlier. This reduction in density on axis results in an increase in the refractive index, and a slight focussing effect for the second half of the pulse. At high plasma densities and at higher wakefield amplitudes this ponderomotive

self-focussing has a greater effect, increasing the intensity of the laser driver and hence the wakefield amplitude produced, in turn increasing the degree of ponderomotive self-focussing. This effect is desirable in many LWFA experiments, leading to both greater wakefield amplitudes and longer acceleration lengths. However, the front of the laser pulse experiences no such effect, and so the shape of the beam becomes strongly distorted, breaking the quasi-static approximations made in the linear regime. In this thesis I focus on both lower plasma densities and lower wake amplitudes, where this effect can be neglected.

Secondly, at high laser intensities ($a \sim 1$, or $I \sim 10^{18}\,\mathrm{Wcm^{-2}}$ for an 800 nm wavelength beam) electrons in the electric field of the laser begin to oscillate at relativistic velocities. This in turn implies a relativistic increase in the effective mass of the electrons as $m = \gamma m_e$ and a corresponding reduction in the plasma frequency $\omega_\mathrm{p} = \omega_{\mathrm{p}0}/\sqrt{\gamma}$. As the beam has a higher intensity on-axis, the plasma frequency is reduced by more on-axis than off-axis, and this leads to a higher refractive index on-axis. A sufficiently intense beam is therefore 'self-focussed' by the plasma, remaining small or even focussing further. The threshold for self-focussing is described by [55]:

$$P_\mathrm{crit} = \frac{m_e c^5}{e^2}\left(\frac{\omega}{\omega_\mathrm{p}}\right)^2 \approx 17\left(\frac{\omega}{\omega_\mathrm{p}}\right)^2 \mathrm{GW}, \tag{1.37}$$

which for an 800 nm wavelength beam in plasma with an electron density of $n_e = 10^{18}\,\mathrm{cm^{-3}}$ is $P_\mathrm{crit} \approx 30\,\mathrm{TW}$. For a laser pulse focussed to $w_0 = \lambda_\mathrm{p}/2 \approx 15\,\mu\mathrm{m}$, this is a peak intensity of around $I_0 \sim 3 \times 10^{18}\,\mathrm{Wcm^{-2}}$ or $a_0 \approx 1.3$. Clearly it is not possible to both achieve self focussing and work in the linear regime at this density. Instead, the laser pulse is prone to non-linear effects which damage the quality of the driver. Furthermore, as the plasma density falls the power required for self-focussing increases. I shall therefore focus on creating guiding structures which are capable of guiding even low intensity laser pulses, particularly at low densities.

Work on guiding structures for laser wakefield acceleration has included research on

evacuated capillaries with a step in the refractive index (employing total reflection from a critical surface) [56] and gas filled capillaries [36]. In practice, laser pulses capable of driving laser wakefield accelerators will typically ablate material from the capillary walls, creating a gradient in the refractive index. Some work has therefore combined evacuated capillaries with an electrical discharge, allowing the ablation from the capillary walls to be controlled more carefully, producing a gradient index guiding structure before the laser pulse arrives [57]. These waveguides have been constructed with lengths of up to $\sim 100\,\mathrm{mm}$ with densities on the scale of $n_e \sim 10^{18}\,\mathrm{cm}^{-3}$, although achieving densities lower than this has proved challenging.

In order to produce approximately parabolic channels with tunable densities, electrical discharges have also been used in gas filled capillaries [58], [59]. Plasma on the edges of the capillary becomes cooler through thermal conduction to the walls and hence its density is higher than hotter plasma in the centre, which is maintained at the same pressure by collisions. This creates a waveguide structure which persists for nanoseconds, but has been demonstrated to run in bursts at repetition rates of up to 1 kHz [60], although this demonstration took place without high intensity laser pulses. These structures have proved very successful, and have been responsible for acceleration of electrons to the GeV scale in structures up to 10 cm long [14], [15]. Recent measurements of these channels using two-colour spectral interferometry [61] show on-axis densities as low as $n_e(0) \approx 10^{17}\,\mathrm{cm}^{-3}$, although this work is not yet published.

However, all capillary-based waveguides rely on the structural integrity of the capillary. This structure can become easily damaged when high intensity pulses do not perfectly couple into the waveguide - if both the pointing and the transverse profile of the LWFA driver are not carefully controlled, a capillary can become damaged on the scale of tens of thousands of shots [62]. While this performance has proved successful for current 1-10 Hz laser systems, using this approach at kHz repetition rates would be highly impractical, requiring replacing capillaries in-situ every few seconds. Furthermore, the capillaries are not trivial to make, requiring machining sapphire or even diamond, which

can take several hours, and takes even longer for longer waveguides.

It's therefore highly desirable to remove the external structure and instead create a 'free-standing' plasma waveguide. This has been achieved using hydrodynamic plasma channels [63]–[65], where an initial column of hot plasma expands, and leaves a lower density region on-axis. The initial column of plasma can be created by using the line focus of an axicon lens in a gas cell or jet, moving any structures much further away from the beam and reducing the possibility of damage to negligible. This allows in principle an infinite number of shots to be taken without replacing the structure.

However, the methods used to heat the plasma to date rely on collisional heating during picosecond duration laser pulses. The collision rate falls rapidly with plasma density and hence this method becomes ineffective for densities below $n_e \sim 10^{18}\,\mathrm{cm}^{-3}$. The lower end of this range has been achieved using clustered gases from a cryogenic gas jet [66], but no lower has been demonstrated to date. Another heating mechanism is possible using Optical Field Ionization (OFI) from a circularly polarized pulse, which can have a femtosecond scale duration. This leaves ionized electrons with an excess of energy [67], creating a hot column of plasma. However, prior to the work described in this thesis, this had only been demonstrated with a spherical lens [68], [69], where the length of this plasma column was limited by the Rayleigh range of the spherical lens to the millimeter scale. In those experiments the on-axis density was not reduced below $n_e \approx 10^{18}\,\mathrm{cm}^{-3}$.

A potential solution is to combine these two methods and use lower density gas to form long and low density plasma channels. The axicon focus allows columns much longer than produced to date, on the scale of 100s of millimeters or higher, and the heating mechanism places no restrictions on the density of the gas. It should therefore be possible to create plasma channels with densities on the scale of $n_e \sim 10^{17}\,\mathrm{cm}^{-3}$ or below, which could potentially operate at multi-kHz repetition rates without suffering any damage. These plasma waveguides would meet all of the requirements for future LWFA stages.

This thesis studies the potential of these plasma waveguides, using a variety of simulations to study physics from the femtosecond scale ionization of the atoms to the nanosecond scale of the hydrodynamic expansion. I look at the energy spectra of the electrons resulting from OFI and the thermalisation of these electrons. By studying the column of plasma that can be produced using an axicon focus and the limitations on it, I can lay some foundations to describe the subsequent expansion of the plasma column, and to study guiding of beams within the channels formed in this way.

1.5 Thesis Overview

In this chapter I've detailed some of the promise of laser wakefield acceleration, but also some of the challenges facing it before it can compete with conventional accelerators. Three of these can broadly be summarised into driving, guiding, and diagnosing; driving laser wakefield accelerators at high repetition rates and efficiencies, guiding the laser pulses to do this over a long distance, and measuring the plasma wakefields produced by these lasers. Throughout the thesis I focus on linear wakefields at low densities, with the kind of parameters that could be produced by lower energy laser systems in the resonant multiple pulse laser wakefield acceleration scheme. However, the work proposed is not exclusive to these linear wakefields.

First, in Chapter 2, I discuss the MP-LWFA scheme, and the accuracy required to resonantly excite a wakefield using trains of many pulses. For this work I use particle-in-cell simulations to model the wakefield amplitude and compare it to a basic theory of the sum of sinusoids. I can also use this basic theory and simulations to study the effect of a proposed scheme which could potentially increase the robustness to error of the MP-LWFA scheme.

Next, in Chapter 3, I write about my work on the TESS diagnostic for linear plasma waves, describing the scheme and extending it to pulses of arbitrary spectral profile. I also discuss the possible extension of TESS to measure quasi-linear and non-linear

plasma waves which are no longer sinusoidal, and ascertain some of the limits on this process. I have described this work in a journal article, which was recently published [70].

Chapter 4 contains the results of a recent experiment which explored resonant excitation of laser wakefields, using TESS and FDH to measure the plasma waves produced. I explore the practical application of the diagnostics in more detail and compare the results to my work on modelling resonant excitation in Chapter 2. The experimental work was published in a recent paper [71], of which I was a contributing author. My role in this work was primarily data analysis, which I shall describe in detail in this thesis.

To explore guiding, in Chapter 5 I describe the hydrodynamic expansion of plasma channels formed from optical field ionization. I use a variety of simulations and theory to model the process from ionization of the plasma to guiding of the laser pulse through the waveguide. I discuss some of the assumptions made and the optimisation of the technique, as well as the limitations upon it. This allows me to explore a wide variety of physics across many orders of magnitude and suggest optimal experimental parameters. I co-authored an article about this work [72], for which I performed the simulations and wrote about much of the theory.

Chapter 6 concludes this thesis. I hope the reader finds the work throughout interesting; I won't claim to solve all of the challenges facing laser wakefield acceleration, but I believe in the process of studying potential solutions, this thesis explores some interesting phenomena and makes a small improvement to the state of play. Useful and exciting rarely coincide in theses, but I hope there shall be small portions of a little of each in the work ahead.

Chapter 2

Exciting Laser Wakefields with Resonant Pulses

2.1 Introduction

We have seen that laser wakefield acceleration offers an exciting alternative to conventional accelerating structures. Laser drivers are capable of creating much higher accelerating gradient travelling waves in plasma than can be supported in radio-frequency cavities. However, as I described in Chapter 1, the laser drivers necessary, creating Joule level and femtosecond-scale duration pulses, are limited due to the practical difficulties of the lasing medium (today, almost ubiquitously titanium doped sapphire). This material has intrinsically low quantum efficiency and requires strong cooling, which is not rapid enough to achieve high repetition rates.

Recent work [28], [29] has demonstrated how these problems may be overcome using a succession of lower-energy laser pulses, applied one after another in resonance with the plasma wakefield, much like pushing a child on a swing. This Multiple-Pulse Laser Wakefield Acceleration (MP-LWFA) would allow the use of new laser drivers, such as fibre and thin disk laser systems, which do not suffer the same material limitations, and hence could produce high gradient accelerating structures more efficiently and more frequently than in the current LWFA experiments.

If MP-LWFA is to be successful, it is crucial to know both what spacing between successive laser pulses is required in order to achieve resonance with the plasma wakefield, and the accuracy to which this spacing must be controlled; this also implies accurate measurement and control of the plasma source itself. We therefore wish to study the effects of pulse trains which imperfectly excite the plasma resonance, in order to establish the limitations of resonant excitation. This may also allow us to explore ways in which we can ameliorate or overcome some of these limitations by using pulse trains with intentionally non-uniform pulse spacing.

Previous work has used fluid dynamics simulations to study this problem [73], [74] but it is desirable to also simulate MP-LWFA experiments using particle-in-cell (PIC) codes. Unlike fluid simulations, PIC codes do not require assumptions about the velocity

distribution or charge separation, and can deal with complex velocity distributions and non-local phenomena by efficiently sampling the phase space with a large number of pseudoparticles [75].

Here I will study one small part of the problem, using the PIC code EPOCH [76], in one dimension, to look at the effect of variations in the timing of uniform amplitude pulses in a train upon the amplitude of the resulting wakefield. I outline a simple linear theory of the sum of sinusoids corresponding to non-relativistic wakefields in section 2.2, before comparing it to results for systematic errors in pulse spacing, section 2.4.1, and random errors in pulse spacing, section 2.4.2. I also draw on results from earlier beatwave drivers to demonstrate how non-uniformly-spaced pulse trains can increase the tolerance to a mismatched plasma density. In section 2.5, I look at the wakefields produced by these pulse trains and the robustness of the technique to mis-measured and non-uniform plasma densities.

2.2 Background & Linear Theory

We have already seen that in the linear regime a laser drives a sinusoidal wakefield at the plasma frequency, with an electric field on-axis described by:

$$E_z(t) = E_1 \cos\left[\omega_\mathrm{p}(t - t_1)\right], \tag{2.1}$$

where E_1 is the wakefield amplitude caused by the laser pulse and t_1 is the arrival time of the laser pulse.

Furthermore, we showed that for a succession of low power pulses the plasma waves add linearly, each with the same frequency, producing a combined wakefield with an electric field described by:

$$E_z(t) = E_1 \cos\left[\omega_\mathrm{p}(t - t_1)\right] + E_2 \cos\left[\omega_\mathrm{p}(t - t_2)\right] + \ldots, \tag{2.2}$$

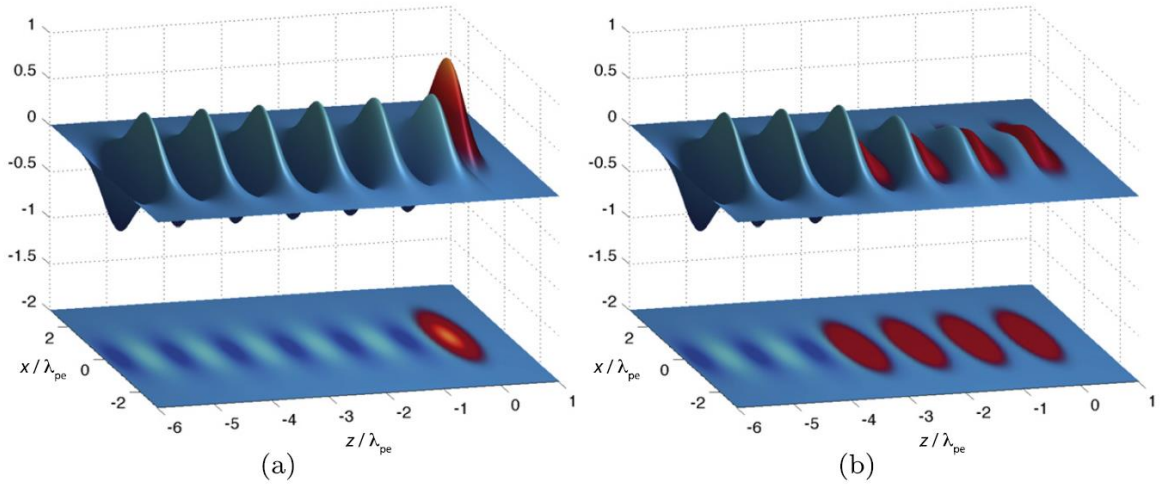


Figure 2.1: A comparison between single and multiple pulse laser wakefield acceleration, showing electron density $\delta n_e/n_{e0}$ in sinusoidal wakefields (blue) behind a laser pulse (red) in the linear regime. The laser pulse is described by the strength of its normalised vector potential a_0^2, where $a_0 = eE_L/(m_e\omega c)$ and propagates from left to right where: a) all of the energy is in one single pulse in conventional LWFA and b) the energy is divided into four identical pulses spaced by the plasma wavelength in MP-LWFA. The same wakefield is achieved with the same laser energy if the pulse separation is perfectly tuned to the plasma. From [29].

where E_n is the wakefield amplitude caused by the n-th laser pulse.

To maximise the strength of this wakefield, the sinusoids should be superposed constructively, with each successive laser pulse one plasma wavelength behind the last, arriving one plasma period later, as in Figure 2.1. The plasma period depends on the electron density, and hence the spacing Δt of pulses in the train must be tuned to the resonance condition $\omega_p \Delta t = 2n\pi$, for integer n, in order to ensure the most constructive superposition.

The result of errors in the pulse train can be studied in one dimension by looking at the sum of multiple misplaced sinusoids. If each individual pulse causes a sinusoidal wakefield with an amplitude $a_n \equiv |\mathbf{E_n}|$, the ideal wakefield, resulting from a train of N pulses perfectly separated by $\Delta t = \tau_p$, is itself sinusoidal with the same frequency and an amplitude of $A \equiv |\mathbf{E}| = \sum_{n=1}^{N} a_n = N\langle a \rangle$. An arbitrary sum of sinusoids with the same frequency will, however, produce a sinusoid with amplitude $0 \leq A \leq N\langle a \rangle$. This can be

modelled by summing arbitrary amplitude sinusoids centred on times $t_n = n\tau_\mathrm{p} + \delta\tau_n$, with $\delta\tau_n$ being the difference from the perfect timing, eqn. (2.4). This is described by the well known result below:

$$\sum_{n=1}^{N} a_n \cos\left[\omega_\mathrm{p}\left(t - n\tau_\mathrm{p} - \delta\tau_n\right)\right] = A\cos(\omega_\mathrm{p} t - \phi) \tag{2.3}$$

$$\text{for} \quad A^2 = \left[\sum_{n=1}^{N} a_n \cos\left(\omega_\mathrm{p}\delta\tau_n\right)\right]^2 + \left[\sum_{n=1}^{N} a_n \sin\left(\omega_\mathrm{p}\delta\tau_n\right)\right]^2 \tag{2.4}$$

$$\tan(\phi) = \left[\sum_{n=1}^{N} a_n \sin\left(\omega_\mathrm{p}\delta\tau_n\right)\right]\left[\sum_{n=1}^{N} a_n \cos\left(\omega_\mathrm{p}\delta\tau_n\right)\right]^{-1}. \tag{2.5}$$

We can also divide the final amplitude into an incoherent part, the sum of the squares of the wakefield amplitudes, and a coherent part, dependent on the spacing between the pulses, as described by:

$$A^2 = \sum_{n=1}^{N} a_n^2 + \sum_{n=1}^{N}\sum_{m\neq n}^{N} a_n a_m \cos\left[\omega_\mathrm{p}(\delta\tau_n - \delta\tau_m)\right]. \tag{2.6}$$

If the positions of the sinusoids are entirely uncorrelated, $\delta\tau_n - \delta\tau_m$ will be entirely random and the second term in eqn (2.6) will average to zero. The resulting wakefield will have an amplitude much lower than the maximum possible value, with $A = \sqrt{N}\, a_\mathrm{rms} \ll N\langle a\rangle$.

Disregarding variations in the strength of individual laser pulses and hence their resulting wakefields, let us consider the sum of independent wakefields with constant unity amplitude $a_n = 1$. The remaining source of error, timing errors in the pulse train, can in general be divided into two categories: random errors around the mean spacing, or 'jitter', and systematic errors which change the mean spacing, or 'tuning'.

2.2.1 Systematic Errors

Systematic errors correspond to using the wrong value for the pulse spacing, $\Delta\tau = \tau_\mathrm{p} + \delta\tau$. This is also the case if the electron density has been measured incorrectly. In this case

there is a constant error in the mean spacing, or a linearly increasing error in the pulse centres $t_n = n\Delta\tau = n\tau_\mathrm{p} + n\delta\tau$, or $\delta\tau_n = n\delta\tau$. Substituting this into eqn. (2.4) results in an amplitude and phase described by:

$$A = \left| \frac{\sin\left(\frac{1}{2}N\omega_\mathrm{p}\delta\tau\right)}{\sin\left(\frac{1}{2}\omega_\mathrm{p}\delta\tau\right)} \right| \tag{2.7}$$

$$\tan\phi = \tan\left(\frac{1}{2}N\omega_\mathrm{p}\delta\tau\right), \tag{2.8}$$

with the resulting amplitude varying periodically with the number of pulses as pulses add first constructively and then destructively.

It is interesting to compare this expression to a diffraction pattern from N slits separated by a spacing d. As with light, we are interfering waves from N different sources, each of which has a different phase. In this analogy, changing the pulse spacing or the plasma density corresponds to moving a short distance off axis, $kd\theta = \omega_\mathrm{p}\delta\tau$, causing the waves to begin to interfere destructively. In the other limit, as $\delta\tau \to 0$, the expression tends, by L'Hôpital's rule, towards $A \to N\cos\left(\frac{1}{2}N\omega_\mathrm{p}\delta\tau\right)/\cos\left(\frac{1}{2}\omega_\mathrm{p}\delta\tau\right) \to N$, as we expect.

2.2.2 Random Errors

Let us now consider random errors, described by some probability distribution $P(\delta\tau_n = \nu) = f_\sigma(\nu)$, which depends on an error parameter σ. In eqn (2.6), the difference in spacing between the m-th and n-th pulses, $\delta\tau_n - \delta\tau_m = \eta$, is therefore described by another probability distribution, $P(\delta\tau_n - \delta\tau_m = \eta) = g_\sigma(\eta)$, which we can write as the autocorrelation of f_σ:

$$\begin{aligned} P(\delta\tau_n - \delta\tau_m = \eta) &= \int \frac{\mathrm{d}\alpha}{\Delta\alpha}\, P(\delta\tau_m = \alpha)P(\delta\tau_n = \alpha + \eta) \\ g_\sigma(\eta) &= (f_\sigma * f_\sigma)(\eta) \end{aligned} \tag{2.9}$$

In the limit of large N we expect the difference in spacings $\delta\tau_n - \delta\tau_m = \eta$ to occur

in the second term of eqn. (2.6) $N(N-1)P(\delta\tau_n - \delta\tau_m = \eta)$ times. The sum then becomes an integral over the possible values of η and we can find the expected wakefield amplitude as:

$$\left\langle A^2 \right\rangle = N + \int \frac{\mathrm{d}\eta}{\Delta\eta}\, N(N-1) g_\sigma(\eta) \cos(\omega_\mathrm{p}\eta) \tag{2.10}$$

This behaves like the real part of a Fourier transform in ω_p, parametrised by σ:

$$\left\langle A^2 \right\rangle = N + N(N-1)\Re\left[\, \widetilde{g}_\sigma\left(\omega_\mathrm{p}\right)\right], \tag{2.11}$$

where $\widetilde{g_\sigma}$ is the Fourier transform of g_σ.

We know, however, that $g_\sigma(\eta)$ is the autocorrelation of $f_\sigma(\eta)$ and, following the Wiener-Khinchin theorem, the Fourier transform of an autocorrelation is the power spectrum: the square modulus of the Fourier transform of the original function $f_\sigma(\eta)$. This gives us an expression for $\widetilde{g_\sigma}$:

$$\widetilde{g}_\sigma\left(\omega_\mathrm{p}\right) = \left|\widetilde{f}_\sigma\left(\omega_\mathrm{p}\right)\right|^2 \tag{2.12}$$

We can substitute this result into eqn. (2.11) to give us the expected amplitude in terms of our probability distribution function $f_\sigma(\eta)$:

$$\begin{aligned} &\left\langle A^2 \right\rangle = N + N(N-1)\left|\widetilde{f}_\sigma\left(\omega_\mathrm{p}\right)\right|^2 \\ &\frac{A_\mathrm{rms}}{N} \to \left|\widetilde{f}_\sigma\left(\omega_\mathrm{p}\right)\right| \quad \text{for} \quad N \gg \left|\widetilde{f}_\sigma\left(\omega_\mathrm{p}\right)\right|^{-1} \end{aligned} \tag{2.13}$$

Once again, this can be compared with the interference of light; in this case the light does not originate from a point source, but from a continuous one dimensional source with varying intensity described by a function f. There is both a term from the incoherent addition of waves, $A^2 \propto N$, and a larger term from the coherent superposition, which depends on the Fourier transform $\tilde{f}$. With these theoretical predictions in mind, we can move on to the simulation of the wakefields driven by imperfect pulse trains.

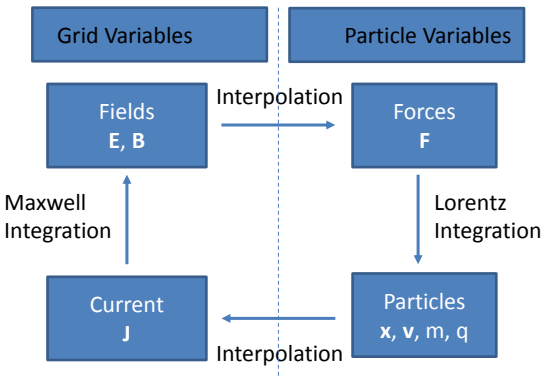


Figure 2.2: Scheme of a Particle-In-Cell code. Electric and magnetic fields are calculated on grid points and are interpolated to calculate the forces upon particles, which move continuously through this grid. These particles and their motions give rise to a current distribution, which is interpolated back to the grid and used to calculate the fields. The whole process is therefore self-consistent without calculating forces between individual pairs of particles. An implicit PIC code

2.3 Simulations

2.3.1 PIC Codes

I simulated the effects of imperfect pulse spacing using the Particle-In-Cell (PIC) code EPOCH [76]. As with other PIC codes, this populates a grid of cells with many pseudo-particles of the different species, each of which represents a cloud of many physical particles. The motion of these pseudo-particles is determined from the electric and magnetic fields as described by the Lorentz force. The electric and magnetic fields are in turn calculated on the grid from Maxwell's equations such that they are self-consistent and consistent with the charge and current distribution of the pseudo-particles. By iteratively calculating the fields and the particle motions, as shown in Figure 2.2, PIC codes describe much of the physics of a real plasma in an intuitive way (see [77] for more details).

To determine particle motion we use the Lorentz equation:

$$\mathbf{F} = \frac{\mathrm{d}\mathbf{p}}{\mathrm{d}t} = q\left(\mathbf{E} + \mathbf{v} \times \mathbf{B}\right) \tag{2.14}$$

$$\frac{\mathrm{d}\mathbf{x}}{\mathrm{d}t} = \frac{\mathbf{p}}{m}, \tag{2.15}$$

where to be fully relativistic $m = \gamma m_0$ and $\mathbf{p} = \gamma m_0 \mathbf{v}$, and where the rest mass m_0 and charge q are actually the sum of the masses and charges of all physical particles within the pseudo-particle.

To determine the evolution of the fields we can rewrite the Maxwell equations in terms of the time derivatives as:

$$0 = \nabla \cdot \mathbf{E} - \frac{\rho}{\epsilon_0} \tag{2.16}$$

$$0 = \nabla \cdot \mathbf{B} \tag{2.17}$$

$$\frac{\partial \mathbf{B}}{\partial t} = -\nabla \times \mathbf{E} \tag{2.18}$$

$$\frac{\partial \mathbf{E}}{\partial t} = \frac{1}{\epsilon_0 \mu_0} \nabla \times \mathbf{B} - \frac{1}{\epsilon_0} \mathbf{J} \tag{2.19}$$

Actually, if our initial conditions for $\mathbf{E}$ and $\mathbf{B}$ satisfy the Poisson equations (2.16) and (2.17), so too do all the subsequent states, so long as charge is conserved, and we only need to explicitly use the Faraday and Ampère laws (2.18) and (2.19). We can

show this as below:

$$\begin{aligned}\frac{\partial}{\partial t}\left(\nabla\cdot\mathbf{E}-\frac{\rho}{\epsilon_0}\right) &= \nabla\cdot\frac{\partial\mathbf{E}}{\partial t}-\frac{1}{\epsilon_0}\frac{\partial\rho}{\partial t}\\ &= \nabla\cdot\left(\frac{1}{\epsilon_0\mu_0}\nabla\times\mathbf{B}-\frac{1}{\epsilon_0}\mathbf{J}\right)-\frac{1}{\epsilon_0}\frac{\partial\rho}{\partial t}\\ &= -\frac{1}{\epsilon_0}\left(\nabla\cdot\mathbf{J}+\frac{\partial\rho}{\partial t}\right)\\ &= 0\end{aligned} \tag{2.20}$$

$$\begin{aligned}\frac{\partial}{\partial t}\left(\nabla\cdot\mathbf{B}\right) &= \nabla\cdot\frac{\partial\mathbf{B}}{\partial t}\\ &= \nabla\cdot\left(-\nabla\times\mathbf{E}\right)\\ &= 0\end{aligned} \tag{2.21}$$

These equations are sufficient to describe the entire behaviour of an ideal plasma, ignoring effects such as collisions and ionization. However, there are several approximations all PIC codes must make which can cause numerical issues:

- **Discrete Grid**. Whereas the particles move on a continuous space, the fields are only calculated exactly at discrete grid points, hence reducing the scale of the problem from $\mathcal{O}(N^2)$ for N particles, to $\mathcal{O}(N\log N_g)$ for $N_g \ll N$ grid points. If the grid points are too far apart, essential physics can be missed, and calculating the force upon a particle at a certain position requires interpolating the fields between several grid points. The higher order the interpolation, generally the more accurate the simulation, but the longer the calculation time required.

- **Pseudo-Particles**. The actual situation to be modelled often involves on the order of $N \sim 10^{16-20}$ particles, which makes reckoning for individual particles computationally impossible*. For that reason the pseudo-particles are weighted

*For comparison, the total amount of digital data created each year is estimated at around ~ 10 ZB or $\sim 10^{22}$ bytes (Forbes, April 23rd 2017), maybe 10^{20} bytes of which is cat memes.

to contain many physical particles, such that the number of pseudo-particles is tractable (although it can still exceed the billions for large simulations, particularly in 3D). However, a large cloud of particles is now being described with just one position and momentum and hence when calculating the fields a shape function describes the distribution of physical particles within the cells. The more accurate the shape function, the more robust the simulation.

- **Numerical Integration**. The forces calculated upon the pseudo-particles and the evolution of the fields give coupled differential equations of motion which must be integrated to find the new fields and the new positions and momenta of these particles. This integration is performed numerically through finite differencing with a certain order and a given timestep, which must adequately resolve all of the physics in the simulation. This timestep must be smaller than the shortest duration process possible, which sets a Courant condition $u\Delta t \leq \Delta x$ for a speed u, a timestep Δt and a grid spacing Δx; the runtime is then inversely proportional to the timestep.

A common consequence of using too poor an interpolation, too simple a shape function, or too coarse an integration step is numerical heating. In this situation, a pseudo-particle moves too far in a single timestep and causes an non-physical increase in the local field strength. This accelerates nearby pseudo-particles further such that the problem is exacerbated and the energy of the system increases over time. There are in general two different ways of removing this issue: either the integration of particle motion during each timestep depends on the fields calculated previously, in an explicit solver, in which case the smallest length scale of the fields, the Debye length, should be properly spatially resolved $\Delta x \ll \lambda_{\mathrm{D}}$ and the Courant condition enforces $\Delta t < \Delta x/c$; or else an implicit solver can be used, which predicts the evolution of the fields from the fields in the previous timestep using a fluid approximation, and therefore allows a particle's motion during a timestep to be calculated from the fields at the end of the

timestep. This implicit method locks in energy conservation and allows much lower spatial and temporal resolution at the cost of a more complex and expensive computation per timestep.

The EPOCH PIC code is developed by the Collaborative Computational Plasma Physics group led by the University of Warwick, and including various UK universities, the UK Atomic Energy Association, the Culham Science Centre, the Rutherford Appleton Laboratory and the Atomic Weapons Establishment[†]. The code is an explicit solver designed to be easily extended by additional features, such as Quantum Electrodynamics interaction models, which will not be described here. The core design is set up in SI units and uses the common Boris pusher scheme to accurately update particle velocities and a split time finite difference scheme to update the electric and magnetic fields.

2.3.2 Parameters

I have made several further approximations in this section in order to simplify the simulations. First, I have conducted all of these simulations in 1D with "simple laser" and "simple outflow" boundary conditions. These allow both radiation and particles to be transmitted while attempting to minimise reflections; this is particularly important under my conditions, when around one hundred laser pulses cross the outward boundary. Throughout, the plasma waves are in the linear regime $\delta n_e/n_{e0} \approx \beta_{\max} \ll 1$. This allows me to treat the plasma wakefields as sums of sinusoids of just one frequency. The situation becomes significantly more complicated once electrons become relativistic.

I also model plasma waves without the dissipative effects of ionization, ion motion or collisions, all of which are explicitly turned off in my simulations, with the plasma starting fully ionized. Landau damping is, however, implicitly included in PIC simulations. In real plasmas these effects act to transfer energy from the plasma wave to the ions and to thermal motion, reducing the wakefield amplitude. We can estimate from the ion plasma frequency, $\omega_{\mathrm{p,i}} = \sqrt{n_i e^2/m_i\epsilon_0} = \omega_{\mathrm{p,e}}\sqrt{Zm_e/m_i}$ (where Z is the average charge

[†]Who use it to study planetary interiors. Nothing to see here...

on each ion, the ion density is n_i, and we have assumed charge neutrality, $Zn_i - n_e = 0$), that ion motion takes place on timescales at least $\sqrt{m_H/m_e} \approx 40$ times slower than the plasma oscillations; I confirmed this behaviour in separate simulations with ion motion turned on. This timescale is longer for heavier ions with lower charge to mass ratios Z/A, where A is the atomic mass number. As the simulation length is around one hundred plasma periods, these simulations are only truly accurate for species where the mass to charge ratio is greater than 6, such as for example N^+ or Ar^{6+}. While achieving these specific ionization states is impractical in most experiments, it allows these simulations to more clearly isolate the effect of pulse spacing errors.

Throughout, I use values outlined in the first proposal of the MP-LWFA scheme [29]. This uses 10 mJ laser pulses focussed to a 40 micron spot size and compressed to a 100 fs duration. This corresponds to a laser intensity of $I_0 = 3.7 \times 10^{15}$ Wcm^{-2}, more than three orders of magnitude lower than the most powerful lasers used in LWFA. The pulses propagate in a fully ionised hydrogen plasma, with electron density $n_\mathrm{e} = 1.75 \times 10^{17}$ cm^{-3}, corresponding to the ideal density for electron acceleration at these laser parameters. This implies a plasma period, and an ideal pulse separation, of $\tau_\mathrm{p} = 266$ fs.

The plasma in these simulations starts at room temperature, and remains cold, with a temperature much less than described by the ponderomotive energy $k_B T_e \ll U_\mathrm{p} = e^2 E_L^2 / 4m\omega_0^2 \sim 100\,\mathrm{eV}$. For these parameters, $\lambda_\mathrm{D} = \sqrt{\epsilon_0 k_B T / n_e e^2} = \lambda_\mathrm{p} \sqrt{k_B T_e / m_e c^2}/2\pi \sim 1\,\mathrm{nm}$, many thousands of times smaller than the plasma wavelength. Properly resolving this is too computationally expensive, requiring tens of thousands of grid cells for each plasma wavelength, and I instead limited numerical heating by limiting the integration step and using a higher order shape function.

2.4 Errors in Pulse Spacing

2.4.1 Systematic Spacing Error

Results

I wish to discover with what accuracy pulses must be separated in the train in order to efficiently drive a wakefield. I do this using a pulse train in which the separation is wrong by a constant offset $\delta\tau$. In this scenario, the n-th pulse is displaced from its ideal position by $n\delta\tau$ as in eqn. (2.7). The resulting wakefield amplitudes from simulations are shown in Fig. 2.3a.

Fig. 2.3a shows that the wakefield amplitude for a pulse train with pulses separated by the plasma period rises almost linearly. It is, however, possible to see the effect of the relativistic change in the plasma period as the trend begins to slowly curve and fall off from linear as the wakefield amplitude increases. The wakefield amplitudes for imperfect pulse trains, on the other hand, oscillate with the number of pulses as $|\sin(\frac{1}{2}N\omega_\mathrm{p}\delta\tau)|$, as predicted by the linear theory, with the maximum wakefields for a non-zero $\delta\tau$ much lower than for a perfect pulse train $\delta\tau = 0$, at $1/N_\mathrm{max} \sin\left(\frac{1}{2}\omega_\mathrm{p}\delta\tau\right) \approx 64\%$ after $N_\mathrm{max} = \tau_\mathrm{p}/2\delta\tau$ pulses. This demonstrates that achieving an accurate pulse separation is necessary for producing a strong wakefield with the multi-pulse method and this strongly limits the length over which a non-ideal pulse train is effective. For instance, at 1% error in separation, adding more pulses is counter-productive after 50 pulses. Conversely, if we know in advance that we wish to use a train N pulses long, we require an accuracy in pulse separation of $\delta\tau/\tau_\mathrm{p} \leq (2N)^{-1}$ to produce a wakefield at least 64% of the resonant case. Alternatively, as $\tau_\mathrm{p} \propto \sqrt{n_e}$, for a given pulse train the density must be correct to within $\delta n_e/n_e \leq N^{-1}$.

To check the analytic predictions against the PIC simulations results more closely, the plasma wave amplitudes, relative to that for a perfect pulse train, are plotted in Fig. 2.3b. This demonstrates a good agreement between eqn. (2.7) and the results

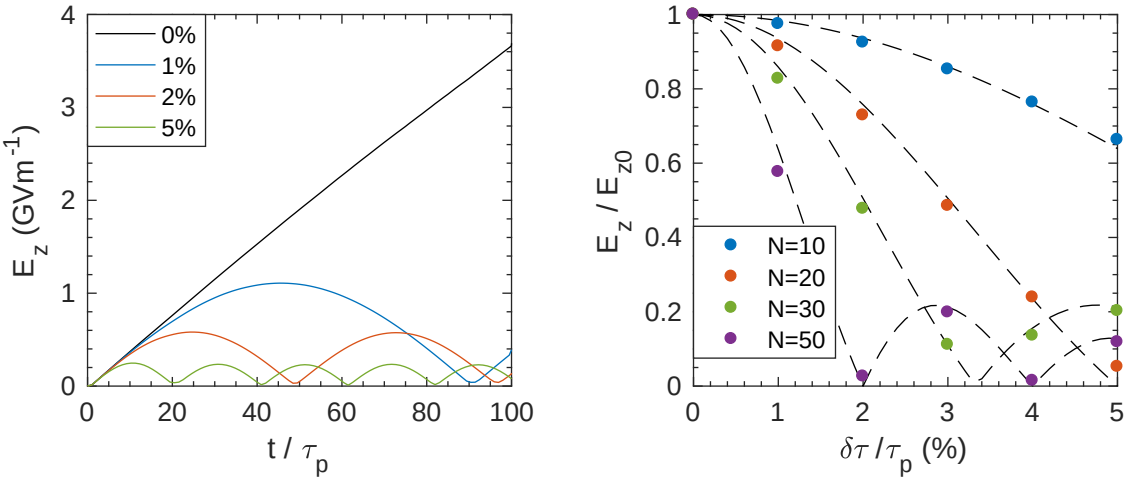


Figure 2.3: One dimensional PIC simulations with uniformly spaced, but mis-tuned, pulse trains. a) Wakefield amplitudes against time, with a perfectly spaced pulse train (black) compared to trains with errors of $\delta\tau/\tau_p = 1\%$ (blue), 2% (red) and 5% (green) in pulse separation. b) Ratio of the wakefield amplitude to that driven by a resonant pulse train as a function of detuning $\delta\tau$ after trains of $N = 10$ (blue), $N = 20$ (red), $N = 30$ (green) and $N = 50$ (purple) laser pulses. The wakefield amplitudes are compared with the analytic prediction (dashed lines) from eqn. (2.7) for pulse separation (or tuning) errors.

of the simulations. However, there are some notable differences, with the simulation results not reaching zero amplitude and the location of the minima some 5 pulses away from the analytic predictions. Firstly, for both a real plasma and the plasma in the PIC simulations, the density profile as a pulse train passes is not purely sinusoidal, but contains many other frequency components, although of much lower amplitude. Energy exchange between these components (such as from Landau damping) ensures that the amplitude of a given component will not be entirely eliminated. The shift in minima location is caused by two competing effects: one is an artefact of the initial lag as pulses enter the simulation window, as shown by an initially quadratic rise in wakefield amplitude in Fig. 2.3a. This moves the entire curves, including the minima, slightly towards later times. The other is a result of the relativistic shift in the plasma frequency. At the highest wake amplitudes the plasma period is slightly longer, and the pulse train is further from resonance. This means the wake amplitude is reduced faster

than if there were no relativistic effect, and the minima occur slightly earlier for higher amplitude wakes, such as at small $\delta\tau/\tau_{\rm p}$, than for lower amplitude wakes and the naive analytic prediction.

Discussion

The basic theory proposed matches the results of simulations well, allowing us to make quick estimates on the precision required for experiments, without running extensive simulations. There will be discrepancies when we extend the one dimensional simulations with stationary ions to a realistic system, but estimates of the scale of the errors should help our experimental design. The error discussed here originates from an incorrect value of the plasma period due either to an error in the density or pressure or from errors in the optical set up of the pulse train generation. Using these results and the linear theory we can set a level of accuracy required in aligning these optics for a given pulse train length and in controlling the pressure of the gas target; for a train N pulses long an accuracy of $1/2N$ is required. For instance, for a train of 20 pulses an accuracy of 2.5% is required in pulse train spacing, or 5% accuracy in the electron density and hence the target gas pressure. It will be important to design experiments such that either the pulse separation or the target pressure are variable in order to tune the experiment until the pulse separation is well matched to the plasma period. This is a similar situation to beam driven plasma wakefield accelerators which rely on the self modulation instability to create trains of particle bunches separated by the plasma period. This has more stringent requirements over many meters of interaction length [78].

2.4.2 Pulse Spacing Jitter

Results

I next studied the effect of random jitter in pulse spacing upon the wakefield. I changed the pulse locations from $t_n = n\tau_{\rm p}$ to $t_n = n\tau_{\rm p} + \delta\tau_n$ where $\delta\tau_n$ is drawn from a Gaussian

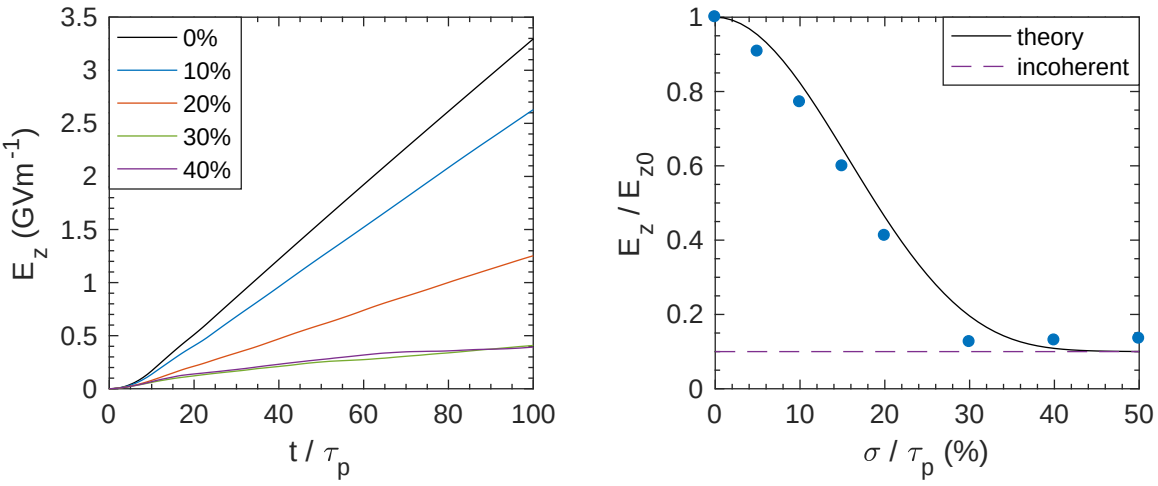


Figure 2.4: a) The mean wakefield amplitude for pulse trains disrupted by random fluctuations in their spacing, described by the standard deviation of the jitter relative to the plasma period. The amplitudes are averaged across 100 simulations. b) Maximum longitudinal electric field of the wake, relative to that produced by an ideal pulse train, versus the standard deviation σ of random fluctuations of the pulse train spacing. The relative amplitude from 100 simulation runs is compared to the analytic prediction from eqn. 2.13 (black line).

distribution with mean 0 and standard error σ, $P(\delta\tau_n = \eta|\sigma) = (\sqrt{2\pi}\sigma)^{-1}\exp\left(-\frac{\eta^2}{2\sigma^2}\right)$. In order to remove random fluctuations, I then ran the simulations 100 times for each value of σ with different sets of random numbers and took the means and calculated the standard deviation for each point in time. The resulting wakefield amplitudes for varying σ are plotted in Fig. 2.4a.

The effect of jitter can clearly be large, with $\sigma/\tau_p = 20\%$ jitter reducing amplitude more than two-fold. However, random fluctuations on the level up to 10% of the plasma period have little effect. Furthermore, random errors do not place any limits on the useful length of the pulse train, as opposed to a systematic pulse spacing error, as in section 2.4.1. This agrees with the linear theory in eqn. (2.13), but to test the theory more accurately I plotted the plasma wave amplitude, relative to that for an ideal pulse train, against σ. I averaged the relative plasma wave amplitude between the 80th and 100th pulses in order to match my theoretical assumption of large N, see Fig. 2.4b.

There is very close agreement between the results from simulations and the analytic

prediction from eqn. (2.13), with the relative amplitude matching a Gaussian in σ, which is the Fourier transform of the Gaussian in η. Increasing the scale of the fluctuations (given as a percentage of the ideal spacing, the plasma period), greatly reduces the plasma wave amplitude up until errors on the scale of 30-40%, where the amplitude is not reduced any further. The amplitude increases over time regardless of the scale of the fluctuations. Within this distribution it seems the data systematically lies beneath the predictions of the theory until $\sigma/\tau_{\mathrm{p}} = 40\%$. Here the relative amplitude has reached a minimum, as the incoherent contribution to the wakefield amplitude described in eqn. (2.13) becomes non-negligible as the assumption $N \gg \left|\widetilde{f}(\sigma)\right|^{-2}$ breaks down, as $\widetilde{f}(40\%) < 0.1$. In this regime, most of the laser pulses passing through the plasma create wakefields which cancel each other out. Enough energy is entering the system, however, that some of the wakefields are not cancelled, regardless of the pulse spacing, in the same manner that even a particle exhibiting a random walk ends up moving gradually further away from its original position.

Discussion

The effect of jitter in the simulations is well predicted by the basic theory. A Gaussian is simply an example of a possible error distribution but is likely representative and provides a scale of the capacity to cope with errors. The simulations show that the wakefield is stable against jitter up to the order of 10%, which agrees with fluid calculations [74] and imposes a limit upon the design of a pulse combination system, where the stability of pulse spacing is important. These constraints are, however, not particularly stringent from an experimental point of view; by using phase locking, jitter can be reduced to the femtosecond, or 1% level[79]. With good stabilisation, the effect of fluctuations around the average pulse spacing is unlikely to have a detrimental effect and would seem to be less important than correctly tuning to the resonance.

2.5 Auto-Resonance

2.5.1 Background

While we have seen the limitation in tuning long pulse trains to the plasma density for a linear wakefield, this effect has been noted before for non-linear wakefields by Rosenbluth and Liu [80] in the context of plasma beatwave acceleration [11], [81]. They describe the effect of relativistic detuning; as the plasma wave amplitude increases, the electrons begin to move at relativistic speeds and hence experience an increase in mass $m_e \to \gamma m_{e0}$, where $\gamma = (1-\beta^2)^{-0.5}$ is the usual relativistic factor for an electron moving at $v = \beta c$. As the mass increases, the effective plasma frequency is reduced, $\omega_\mathrm{p} \to \omega_{\mathrm{p}0}/\sqrt{\bar{\gamma}}$, where $\bar{\gamma}$ is average of the γ factor taken over an oscillation. The ideal laser pulse spacing for resonant amplification therefore changes as the wake amplitude increases. This makes it impossible for a uniformly spaced pulse train to be perfectly tuned to the plasma period at all wakefield amplitudes. The Rosenbluth-Liu limit describes the maximum wakefield amplitude that can be obtained before the pulse train loses resonance; for a laser of normalized amplitude a_0 this limit is:

$$E_\mathrm{RL} = \left(\frac{16a_0^2}{3}\right)^{\frac{1}{3}} E_\mathrm{wb} \tag{2.22}$$

$$E_\mathrm{wb} = \frac{m_e \omega_\mathrm{p} c}{e}, \tag{2.23}$$

where E_wb is the cold non-relativistic wave-breaking limit.

It has since been suggested, e.g. by Deutsch et al. [82], that a *non-uniform* pulse train could overcome this problem. If the plasma period increases over time as the wakefield grows, so too should the separation between laser pulses increase over time. For low amplitude wakefields, the electron velocity varies approximately sinusoidally within a slowly increasing envelope $\beta(t) = \beta_m(t)\cos(\omega_\mathrm{p} t)$. We can approximate the

relativistic mass increase of the electrons, described by:

$$\begin{aligned}\gamma(t) &= \left[1-\beta(t)^2\right]^{-0.5} \\ &\approx \left[1-\beta_m(t)^2\cos^2(\omega_{\mathrm{p0}}t)\right]^{-0.5} \\ &\approx 1+\frac{1}{2}\beta_m(t)^2\cos^2(\omega_{\mathrm{p0}}t) \qquad (2.24)\\ \bar{\gamma} &\approx 1+\frac{1}{4}\beta_m(t)^2 \qquad (2.25)\end{aligned}$$

The plasma period therefore becomes:

$$\begin{aligned}\tau_{\mathrm{p}} &= \sqrt{\bar{\gamma}}\tau_{\mathrm{p0}} \\ &\approx \left(1+\frac{1}{4}\beta_m(t)^2\right)^{0.5}\tau_{\mathrm{p0}} \\ &\approx \left(1+\frac{1}{8}\beta_m(t)^2\right)\tau_{\mathrm{p0}} \qquad (2.26)\end{aligned}$$

If every laser pulse is of equal amplitude, each creating a wake of amplitude $\beta_1 = E_1/E_{\mathrm{wb}}$, the pulse spacing should increase quadratically as $\tau_n/\tau_{\mathrm{p0}} = 1+\beta_1^2 n^2/8$. This implies pulse timings should increase cubically, as $t_n = (n+\beta_1^2 n^3/24)\tau_{\mathrm{p}}$. A pulse train which was designed such that the pulse spacing increased in accordance with this would be able to resonantly excite a wakefield even after the plasma period had changed. However, in practice, measuring E_1 or β_1 would be prohibitively difficult, and in any case once $\beta \sim 1$ most of these approximations break down and the required pulse spacing is no longer analytic.

Deutsch et al. suggest instead approximating the change in plasma period by increasing the spacing linearly over time, which is much easier to achieve experimentally with the beatwave set up, by using a chirped laser pulse. As described in Chapter 1, in the beatwave scheme two laser pulses co-propagate with the same field strength $E_1 = E_2$ and with similar frequencies ω_0 and ω_1, giving rise to a fast oscillation at a frequency $\omega_+ = (\omega_1+\omega_0)/2$ inside a slowly oscillating envelope at $\omega_- = (\omega_1-\omega_0)/2$. This envelope

has maxima at $\omega_- t_n = n\pi$, and so if we set $\omega_1 - \omega_0 = \omega_{\mathrm{p0}}$, the maxima are separated by the plasma period, $t_n = 2n\pi/\omega_{\mathrm{p0}} = n\tau_{\mathrm{p0}}$, as we require for resonance. If one of the laser pulses is now chirped with a frequency described by $\omega_1 = \omega_0 + \omega_{\mathrm{i}} - \alpha\omega_{\mathrm{p0}}^2 t$ (where ω_{i} is the initial frequency of the *beating* in the laser intensity), the frequency of the slow oscillation decreases over time and the separation of the maxima increases. The maxima in the intensity of the driving laser are now located at:

$$2\omega_- t_n = 2\pi n = \left(\omega_{\mathrm{i}} - \alpha\omega_{\mathrm{p0}}^2 t_n\right) t_n \tag{2.27}$$

$$\Rightarrow 0 = \alpha\omega_{\mathrm{p0}}^2 t_n^2 - \omega_i t_n + 2\pi n$$

$$\Rightarrow t_n = \frac{1}{2\alpha\omega_{\mathrm{p0}}^2}\left(\omega_{\mathrm{i}} \pm \sqrt{\omega_{\mathrm{i}}^2 - 8\pi\alpha\omega_{\mathrm{p}}^2 n}\right)$$

$$= \frac{\tau_{\mathrm{p0}}}{4\pi\alpha}\left[\frac{\omega_{\mathrm{i}}}{\omega_{\mathrm{p0}}} - \sqrt{\left(\frac{\omega_{\mathrm{i}}}{\omega_{\mathrm{p0}}}\right)^2 - 8\pi\alpha n}\right], \tag{2.28}$$

where we have taken the negative root such that the $n = 0$ pulse arrives at $t_0 = 0$ and higher n pulses arrive later in time.

Note that these times are those at which pulses arrive. The spacing between pulses can therefore be approximated as the derivative:

$$\tau_n = t_n - t_{n-1} \approx \frac{\mathrm{d}t_n}{\mathrm{d}n}$$

$$= \frac{\tau_{\mathrm{p0}}}{4\pi\alpha}\frac{1}{2}8\pi\alpha\left[\left(\frac{\omega_{\mathrm{i}}}{\omega_{\mathrm{p0}}}\right)^2 - 8\pi\alpha n\right]^{-0.5}$$

$$= \tau_{\mathrm{p0}}\left[\left(\frac{\omega_{\mathrm{i}}}{\omega_{\mathrm{p0}}}\right)^2 - 8\pi\alpha n\right]^{-0.5} \tag{2.29}$$

If we choose $\omega_{\mathrm{i}} = \omega_{\mathrm{p0}}$, for small chirp rates $8\pi\alpha n \ll 1$:

$$\tau_n \approx \tau_{\mathrm{p0}}\left(1 + 4\pi\alpha n + \ldots\right) \tag{2.30}$$

This means that the spacing between pulses increases over time and by choosing α we

can approximate the true plasma period, depending on the growth rate of the wakefield.

Lindberg et al. [83], [84] go further, and suggest exploiting the phenomenon of auto-resonance. Under certain conditions, a non-linear oscillating system driven by an oscillatory force through a resonant state begins to follow the frequency of the driver rather than oscillating at its natural frequency, in the same manner as a driven harmonic oscillator exhibits a response to frequencies away from its natural resonance. The driver must start above resonance with the initial beat frequency $\omega_{\mathrm{i}} > \omega_{\mathrm{p0}}$ and the plasma must be initially quiescent, with its response to the driver linear in the low amplitude limit. The driver then excites a plasma wave which increases in amplitude while the driver changes frequency. This change in frequency must be sufficiently slow, with the chirp rate α less than the critical chirp rate α_c described below:

$$\alpha \leq \alpha_c \approx 0.169\epsilon^{\frac{4}{3}} \tag{2.31}$$

Here the strength of the driving pulses is described by $\epsilon = a_1 a_2$ for the normalized vector potentials of the two drive lasers $a_i = eE_i/m_e\omega_0 c$. The more strongly the plasma wave is driven, the faster the non-linearity becomes evident and the faster the chirp can be.

2.5.2 Multiple Pulse Auto-Resonance

Let us now again consider multiple pulse drivers. If the current state of the plasma wave, its frequency and phase, is known exactly at any given time, the laser pulse can be timed to arrive at exactly the right moment to reinforce the existing wakefield. In principle this allows a multiple pulse system to drive a laser wakefield more effectively than any simple beatwave system, including a system which exploits auto-resonance[‡].

[‡]It's worth pointing out that it is possible to replicate any multiple pulse driver with a beatwave scheme between two equal amplitude laser pulses, so long as one can be given a completely arbitrary phase profile. First calculating and then producing this phase profile is, however, very difficult. In order to use the results from earlier in this chapter we will continue to consider a succession of identical pulses.

However, in practice such a measurement and calculation is extremely difficult if the wakefield amplitude is not negligible and we have seen that using an incorrect pulse spacing dramatically reduces the wakefield amplitude produced after a number of pulses.

We will consider applying the results from auto-resonance to this situation, using a train of pulses chosen such that the spacing between pulses increases over time in the same way as for an auto-resonant beatwave driver. We therefore place the n-th pulses at a time t_n when the equivalent beatwave driver is at a maximum, as described by equation (2.28). Here, the chirp rate α is the same as for a beatwave driver but the critical chirp rate has now changed. To reach the same laser intensity with one laser as with the beatwave scheme requires $a_0 = 2a_1$ and hence $\epsilon = a_0^2/4$, giving $\alpha_c = 0.0266 a_0^{\frac{8}{3}}$.

We can use our previous work on the linear sum of sinusoids to estimate the wake amplitude resulting from these chirped drive pulses. Note that this sum ignores the order in which pulses arrive, therefore disregarding any dynamic effects from the chirp. First, in the limit of a low chirp rate, $8\pi\alpha n \ll (\omega_\mathrm{i}/\omega_\mathrm{p0})^2$, we can expand the pulse timing to second order as:

$$\begin{aligned} t_n &= \frac{\tau_\mathrm{p0}}{4\pi\alpha}\frac{\omega_\mathrm{i}}{\omega_\mathrm{p0}}\left[1-\sqrt{1-8\pi\alpha n\left(\frac{\omega_\mathrm{i}}{\omega_\mathrm{p0}}\right)^{-2}}\right] \\ &\approx \frac{\tau_\mathrm{p0}}{4\pi\alpha}\frac{\omega_\mathrm{i}}{\omega_\mathrm{p0}}\left[\frac{1}{2}8\pi\alpha n\left(\frac{\omega_\mathrm{i}}{\omega_\mathrm{p0}}\right)^{-2}+\frac{1}{2\cdot 2}\left(8\pi\alpha n\right)^2\left(\frac{\omega_\mathrm{i}}{\omega_\mathrm{p0}}\right)^{-4}+\ldots\right] \\ &= \tau_\mathrm{p0}\frac{\omega_\mathrm{p0}}{\omega_\mathrm{i}}n+\frac{4\pi^2\alpha n^2}{\omega_\mathrm{i}}\left(\frac{\omega_\mathrm{i}}{\omega_\mathrm{p0}}\right)^{-2}+\ldots \end{aligned} \tag{2.32}$$

$$\delta t_n \equiv t_n - n\tau_\mathrm{p0} = 2\pi n\left(\frac{1}{\omega_\mathrm{i}}-\frac{1}{\omega_\mathrm{p0}}\right)+\frac{4\pi^2\alpha n^2}{\omega_\mathrm{i}}\left(\frac{\omega_\mathrm{i}}{\omega_\mathrm{p0}}\right)^{-2}+\ldots \tag{2.33}$$

$$\omega_\mathrm{p0}\delta t_n = 2\pi n\left(\frac{\omega_\mathrm{p0}}{\omega_\mathrm{i}}-1\right)+4\pi^2\alpha n^2\left(\frac{\omega_\mathrm{i}}{\omega_\mathrm{p0}}\right)^{-3}+\ldots \tag{2.34}$$

This is the phase of the wakefield created by the n-th laser pulse alone, expanded to second order. It is slowly varying when $\omega_\mathrm{p0} \approx \omega_\mathrm{i}$ and αn is small. Assuming as before that all the pulses are of equal strength, creating wakefields of amplitude $a_n = 1$, we can use this to find the wakefield amplitude from our expression in eqn. (2.4). The details

of this calculation are in appendix A, with the result that the wakefield amplitude and phase can be described to second order by:

$$A = \frac{1}{c}\left|\mathrm{Eu}\left[v(N)\right] - \mathrm{Eu}\left[v(0)\right]\right| \tag{2.35}$$

$$\phi = \arg\left\{\mathrm{Eu}\left[v(N)\right] - \mathrm{Eu}\left[v(N)\right]\right\} - \chi \tag{2.36}$$

Here we have defined the additional parameters:

$$\mathrm{Eu}(v) = C(v) + iS(v) \tag{2.37}$$

$$C(v) = \int_{x=0}^{v} \mathrm{d}v' \ \cos\left(\frac{\pi v'^2}{2}\right) \tag{2.38}$$

$$S(v) = \int_{x=0}^{v} \mathrm{d}v' \ \sin\left(\frac{\pi v'^2}{2}\right) \tag{2.39}$$

$$v(x) = cx + v(0) \tag{2.40}$$

$$c = \sqrt{8\pi\alpha}\left(\frac{\omega_\mathrm{i}}{\omega_\mathrm{p0}}\right)^{-\frac{3}{2}} \tag{2.41}$$

$$v(0) = -\frac{1}{\sqrt{2\pi\alpha}}\left(\frac{\omega_\mathrm{i}}{\omega_\mathrm{p0}}\right)^{+\frac{1}{2}}\left(\frac{\omega_\mathrm{i}}{\omega_\mathrm{p0}} - 1\right) \tag{2.42}$$

$$\chi = \frac{1}{4\alpha}\left(\frac{\omega_\mathrm{i}}{\omega_\mathrm{p0}}\right)\left(\frac{\omega_\mathrm{i}}{\omega_\mathrm{p0}} - 1\right)^2 \tag{2.43}$$

$\mathrm{Eu}(v)$ is a complex function of a parameter v. Its shape is the Euler spiral, otherwise known as a Cornu spiral, and might be familiar for its role in Fresnel diffraction [85], [86], the transition curves of railway tracks and highways [87], and elastic beams [88][§]. In this situation it allows us to determine the wakefield amplitude by measuring the distance between two points on the curve defined by $v(N)$ and $v(0)$. It's a neat result; whereas the wakefield amplitude of a uniformly spaced pulse train can be described by linear diffraction theory, taking a Fourier transform as in the Fraunhofer regime, the amplitude caused by a linearly chirped train requires the next order correction, moving to the Fresnel regime.

[§]While there's not strictly a need for me to cite Euler, I'll hardly pass up the chance...

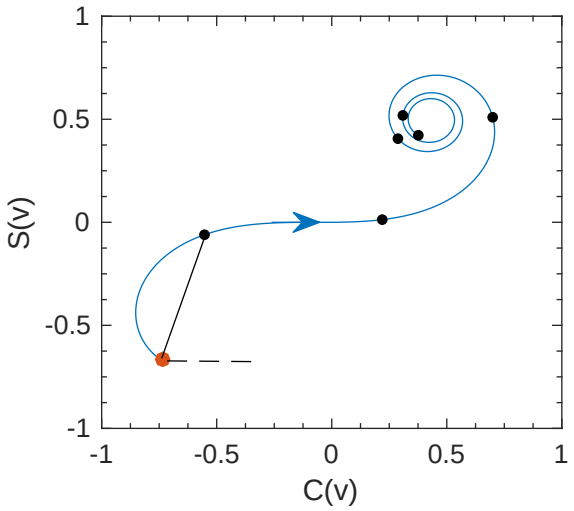


Figure 2.5: The Euler Spiral Eu(v) created by the sum of sinusoids from a chirped pulse train, parameterised by v. Eu[$v(0)$] is marked by a red dot and Eu[$v(50m)$] are marked by black dots for $m = \{1, .., 6\}$, with both N and $v(N)$ increasing along the direction of the arrow. The amplitude of the wake after N pulses is proportional to the distance between Eu[$v(0)$] and Eu[$v(N)$], as demonstrated by the solid black line for $A = 50$. The phase of the plasma wave is described by the angle between this line and the horizontal, the dashed black line, and hence the frequency is the derivative of this angle with increasing N.

Similarly, the change in the phase of the wakefield is the argument of the difference between the two points on the Euler spiral (that is, the angle between the end point at $v(N)$, the start point at $v(0)$, and the positive horizontal). We have assumed that the wakefield is at a constant frequency ω_{p0} while summing sinusoids, and therefore the interesting information about the wakefield frequency is hidden in the phase ϕ. After n pulses, the wakefield is described by:

$$\mathbf{E}(t) \propto \cos\left[\omega_{\mathrm{p0}} t + \phi(n)\right] \tag{2.44}$$

After the nth pulse the frequency is therefore:

$$\begin{aligned}
\omega(t) &= \frac{\mathrm{d}}{\mathrm{d}t}\left[\omega_{\mathrm{p0}} t + \phi(n)\right] \\
&= \omega_{\mathrm{p0}} + \frac{\mathrm{d}\phi}{\mathrm{d}n}\frac{\mathrm{d}n}{\mathrm{d}t} \\
&\approx \omega_{\mathrm{p0}} + \frac{\mathrm{d}\phi}{\mathrm{d}n}\frac{1}{\tau_{\mathrm{p0}}} \\
&= \omega_{\mathrm{p0}}\left[1 + \frac{1}{2\pi}\frac{\mathrm{d}\phi}{\mathrm{d}n}\right] \\
&= \omega_{\mathrm{p0}}\left[1 + \frac{1}{2\pi}\frac{\mathrm{d}}{\mathrm{d}n}\arg\Delta\mathrm{Eu}\right]
\end{aligned} \tag{2.45}$$

By calculating the Euler spiral we can not only predict the wakefield's amplitude, but also its frequency. In this way a simple theoretical argument has given us a prediction for the wakefields produced by a train of multiple pulses, now with non-uniform spacing. Note that the predictions of this model are independent of the order in which laser pulses arrive in the medium, and hence it ignores any dynamic effects, such as, ironically, auto-resonance itself, the original motivation for this work. By comparing this prediction to results from simulations we can isolate and study the effect of dynamic auto-resonance, as well as predicting the most effective multiple pulse drivers.

2.5.3 Results

I simulated a chirped pulse train in EPOCH with 200 pulses, using a variety of different chirp rates on the scale of $\alpha \sim 10^{-5}$, with the initial driver frequency slightly higher then the plasma frequency, $\omega_{\mathrm{i}}/\omega_{\mathrm{p}} = 1.01$. The other parameters were the same as described earlier, with $a_0 = 0.052$ resulting in a critical chirp rate of $\alpha_c = 1.00 \times 10^{-5}$. The resulting plasma wakefields were fitted to sinusoids of variable amplitude, frequency and phase, of the form $E_w(t) = A(t)\cos\left[\omega(t)t + \phi(t)\right]$ with these fitting coefficients shown in Figures 2.6 and 2.7. The wakefield remains sinusoidal and non-relativistic, with $\beta < 0.13$, and remains below the Rosenbluth-Liu limit, which for these parameters is $E_{\mathrm{RL}} \approx 10$ GVm^{-1}.

None of the mistuned pulse trains therefore perform as well as a perfectly tuned uniform pulse train ($\omega(t) = \omega_{\mathrm{p}}$), with the most effective train, chirped at $\alpha = 10^{-5}$ (red line), driving a wakefield which after 160 pulses reaches 87% of the amplitude of the wakefield driven by perfectly tuned train. However, whereas the wakefield from a mis-tuned and uniform ($\omega(t) = \omega_{\mathrm{i}}$) pulse train never exceeds 2 GVm^{-1}, chirped pulse trains can reach wakefield amplitudes of 5 GVm^{-1}. This exceeds the expectation using the basic theory, which predicts that after 160 pulses a train chirped at $\alpha = 10^{-5}$ will drive a wakefield with an amplitude 76% of the perfectly tuned pulse train. Both plots show the wakefield amplitude normalized to the electric field strength generated by a single laser pulse, E_1; for our PIC simulation parameters this was $E_1 = (38.7 \pm 0.3)$ MVm^{-1}. This allows us to see that the wakefield amplitude measured in simulations is in all cases very slightly larger than that predicted from analytic theory.

In Fig. 2.7, the wakefield frequency changes in much the way predicted by eqn. (2.45), starting halfway between the plasma frequency and the driver frequency at $\approx 1.005\,\omega_{\mathrm{p0}}$ and subsequently following the frequency of the driver rather than the plasma frequency until losing resonance around 0.2% below the plasma frequency. For both theory and simulations, the maximum wake amplitude is achieved just before the wake slips from

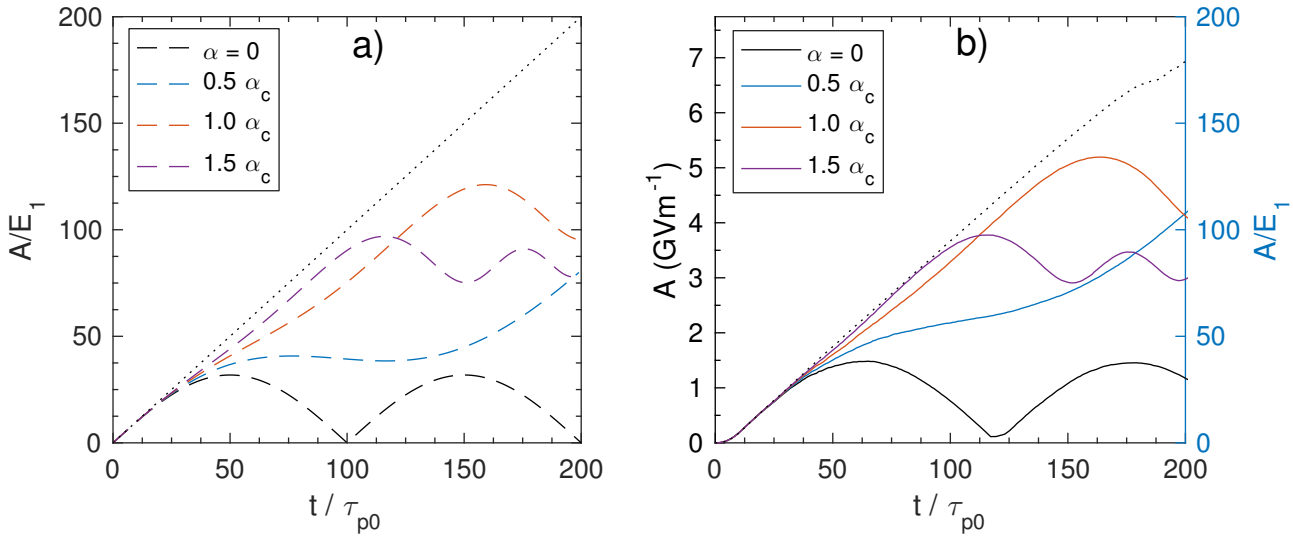


Figure 2.6: The wakefield amplitude over time expected for plasmas driven through resonance by trains of linearly chirped pulse trains using a) the sum of the individual sinusoids (dashed) and b) the results from simulations (solid).

the driving frequency. Thereafter, the minima in frequency correspond to the maxima in amplitude and vice versa, and the higher the chirp rate, the more oscillations the frequency undergoes. This corresponds to the oscillatory region at the edge of the Euler spiral –at higher chirp rates v increases faster and each cycle is traversed sooner. However, in simulations the size of these oscillations is smaller than predicted by the analytic theory, particularly for the sub-critically chirped pulse train driver ($\alpha = 0.5\,\alpha_c$) where no oscillation is seen in simulations.

Some of the differences between the analytic theory and the PIC simulations are due to numerical limitations. The simulation box contains ten plasma periods in order to avoid edge effects and better estimate the wakefield amplitude and frequency, and hence both the frequency and amplitude measured in the simulations are effectively averaged over ten plasma periods, reducing the amplitude of oscillations. The amplitude and frequency also suffer a lag of five plasma periods while the laser pulses pass through the simulation box.

We can run these simulations for a wide range of chirps, between $\alpha = 0$ and $\alpha = 4 \times 10^{-5}$, producing surfaces of the wakefield amplitude as shown in Figure 2.8. We

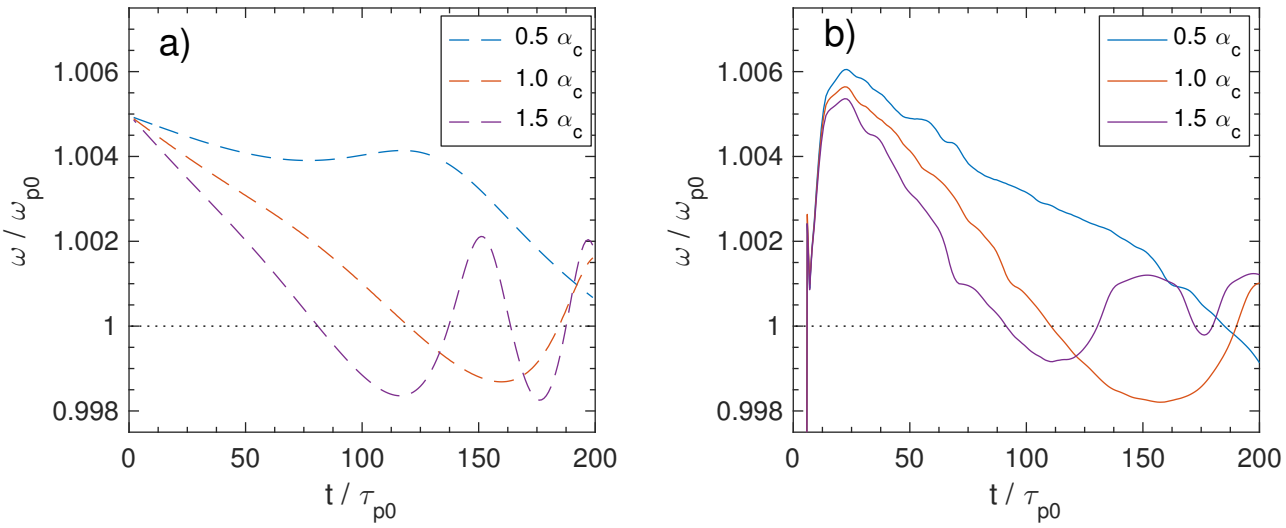


Figure 2.7: The wakefield frequency over time expected for plasmas driven through resonance by linearly chirped pulse trains using a) the analytic theory described in equation (2.35) (dashed) and b) the results from simulations (solid).

see that the analytic theory gives results that look qualitatively very similar across the whole range of chirps and times; the Fresnel interference picture accurately describes the wakefield amplitude from a chirped multiple pulse driver. For all chirps, the wakefield amplitude initially rises with time at the same rate. For very low chirps the amplitude then falls as the driving laser pulses fall out of phase with the plasma wakefield, before rising again later, as for the uniform pulse trains shown in Figure 2.3. Immediately below the critical chirp $\alpha_c = 1.00 \times 10^{-5}$, the amplitude instead only rises with time, reaching the maximal wakefield amplitude after 200 pulses. However, for higher chirp rates the driver frequency changes too quickly and loses resonance with the plasma wave sooner. This results in an earlier maximum amplitude and a ringing effect, where the wakefield amplitude oscillates around a final value lower than the maximum, of the same form as seen in Fresnel diffraction.

Only when we subtract the theoretical predictions from the results from simulations can we see the effect of the dynamic processes, including auto-resonance. Here the results from simulations have been transposed by five plasma periods and divided by $E_1 = 38.7$ MVm^{-1} so as to normalize the amplitudes to the same units as the

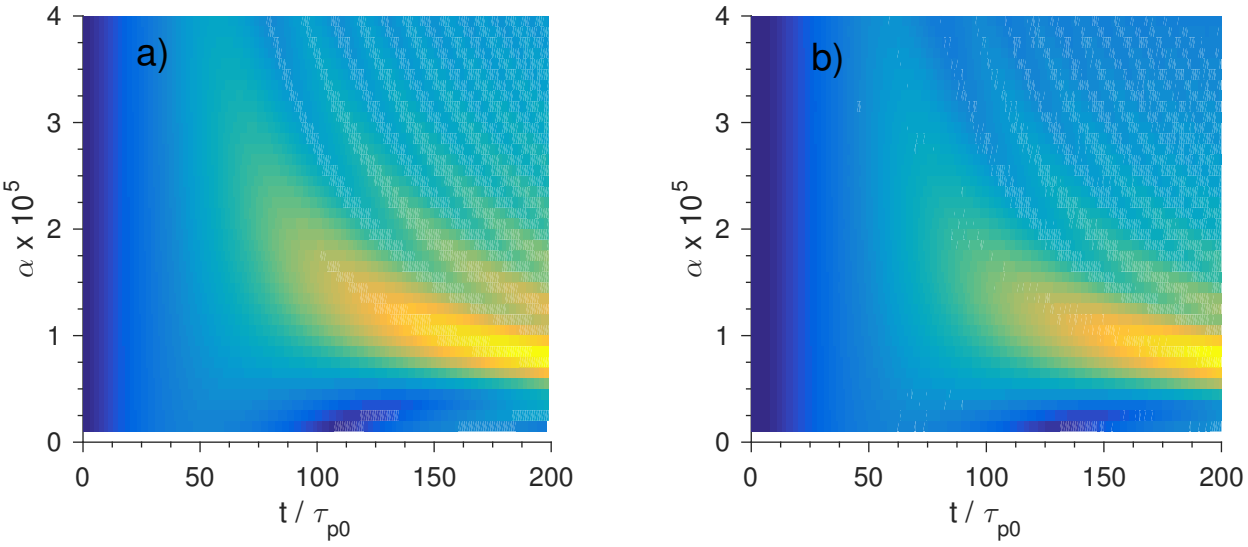


Figure 2.8: The wakefield amplitudes over time against the chirp of the multiple pulse driver, expected from a) the sum of the individual sinusoids and b) simulations.

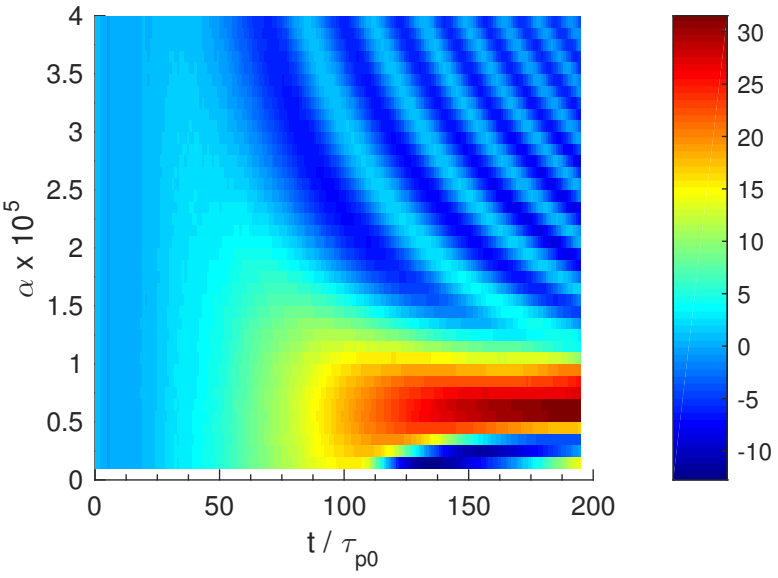


Figure 2.9: Difference between the wakefield amplitudes calculated in PIC simulations and the analytic theory. To compare like-for-like, the results from simulations have been transposed to 5 plasma periods earlier in time, $t \to t - 5\tau_{p0}$, and have been normalized by the wakefield amplitude from a single pulse, $A \to A/E_1$.

theoretical prediction. The ringing feature is strongly evident at high chirp rates and later times and we see that the simulations give wakefield amplitudes the same or lower than those predicted by theory everywhere but around $\alpha \approx \alpha_{\text{crit}} = 10^{-5}$, where the wakefield amplitude is substantially higher by up to $30E_1$ or 15% (≈ 1.2 GVm^{-1} for these simulation parameters). This demonstrates that auto-resonance is active here and is a dynamic process that increases the wakefield beyond the sum of its sinusoidal parts.

2.5.4 Discussion

We have seen earlier in this chapter that uniform pulse trains mis-tuned to the plasma frequency by as little as 1% become ineffective after some 50 laser pulses, or more generally for a pulse train of N pulses the pulse spacing must be accurate to within $1/(2N)$. For the chirped pulse trains we have introduced here, with a chirp α and pulse train length N, the pulse spacing changes across the train by a total $\Delta\tau/\tau_{\text{p0}} \approx 4\pi\alpha N$, and hence the driver can explore a range of frequencies and excite a wakefield even when the plasma period is not known exactly. We have formulated a basic prediction for the wakefield amplitude resulting from this kind of non-uniformly spaced train of laser pulses based on the second order expansion of a sum of sinusoids and predict that the chirp and initial driving frequency must be chosen in such a way as to maximise the distance across an Euler spiral.

This was observed in simulations, with the critically chirped pulse train significantly outperforming a uniformly spaced but mis-tuned pulse train with the same initial frequency, producing wakefield amplitudes up to 3 times greater as the wakefield frequency decreased from around $\omega = 1.006\,\omega_{\text{p0}}$ to around $\omega = 0.998\,\omega_{\text{p0}}$ before losing resonance with the pulse train driver. However, a critical chirp rate limits the range of frequencies that the driver can explore, and hence the maximum density error in the plasma which can be tolerated. Above the critical chirp rate, in this case $\alpha_c = 10^{-5}$, pulse trains fall out of resonance with the wakefield after a finite time and the wakefield

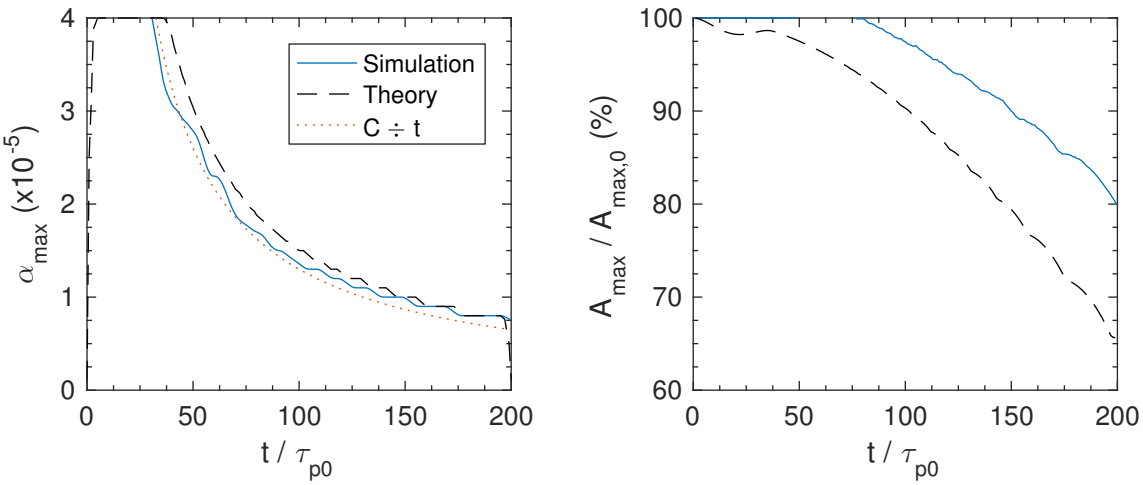


Figure 2.10: a) The chirp rate which maximises the wakefield amplitude after a given time, drawn from Figure 2.8, comparing theory with simulations and with the trend $\alpha_{\max} = C/t$ for $C = 130 \times 10^{-5}\tau_{p0}$. b) The relative wakefield amplitude at this optimal chirp against time, compared to that produced by a uniform perfectly tuned pulse train driver.

amplitude is then substantially reduced. This matches the second order approximation well over a wide range of chirps, demonstrating that there is a significant regime where we can accurately describe the plasma without invoking complex dynamic behaviour. However, near the critical chirp and after more laser pulses the dynamic nature of auto-resonance is not described with this approximation and the wakefield amplitude can be substantially higher than predicted with the simple model.

We can also determine the optimal chirp rate if trying to excite a wake with a train of a given number of pulses, see Figure 2.10. For a fixed amount of energy per pulse, shorter pulse trains can support a higher chirp rate, with the wakefields adding constructively and the wakefield not yet losing resonance with the pulse train driver. To use longer pulse trains, however, the chirp must decrease towards $\alpha \approx \alpha_c = 10^{-5}$. This agrees well with the analytic theory from the sum of sinusoids, which predicts that the number of pulses can be higher without tipping around the end of the Euler spiral. However, as the pulse train becomes longer the efficiency of the driver is reduced; whereas after 100 pulses the simulations found that the critically chirped train produced a wakefield with an amplitude 97% of that caused by the ideal perfectly spaced driver,

after 200 pulses, this efficiency had dropped to 80%. Because of auto-resonance, this is significantly higher than the efficiency predicted by the basic theory, which is just 66% of the theoretical maximum.

As the optimal chirp rate is approximately inversely proportional to the number of pulses in the train that have passed, the total frequency explored by the driver is close to constant, corresponding to $\Delta\tau/\tau_{\mathrm{p0}} \approx 1.6\%$ for these parameters, whereas the initial error in the driver frequency was set as $(\omega_{\mathrm{i}} - \omega_{\mathrm{p0}})/\omega_{\mathrm{p0}} = 1\%$. If the error on the plasma frequency is higher than this, auto-resonance will cease to be effective as it can no longer be guaranteed that the drive frequency will pass through resonance. Instead, the critical chirp rate must instead be increased by using more intense laser pulses, with $\Delta\tau/\tau_{\mathrm{p0}} \propto \alpha_c \propto a_0^{\frac{8}{3}}$.

2.6 Conclusions

In summary, I considered two different ways in which errors in the spacings of pulses in a train will change the wakefield generated through multi-pulse wakefield acceleration: a systematic error in mean pulse separation and random fluctuations around the mean pulse separation. The findings match both theoretical expectations from a basic linear theory of the sum of sinusoids and previous findings from hydrodynamic simulations in both cases.

- To produce a train N pulses long the mean pulse separation must be controlled to within $(2N)^{-1}$ of the plasma period.
- The process is stable to random fluctuations in pulse separation on the scale of up to 10%, which would reduce the wakefield amplitude to 80% of that from the ideal pulse train.

Many methods of combining pulses can suffer from random fluctuations in pulse separation. Good stability is required to achieve a strong wakefield but current methods

can achieve this. For instance, previous work in [79] has reduced random variations to the scale of femtoseconds, or around 1% of the plasma period. This would reduce the wakefield amplitude to down to 99.8% of that from the ideal pulse train, a negligible effect. The main limit upon the pulse train is therefore set by the accuracy to which we can tune the pulse separation and/or the density and pressure of the plasma. A 5% error when measuring the density of the plasma implies the maximum effective pulse train is only 20 pulses long and the resulting wakefield is only 64% of the ideal amplitude. This is a strong limitation on the experimental accuracy required for MP-LWFA to succeed. If we wish to use a train 50 pulses long, we must achieve the same density shot to shot and tune the pulse separation to within 1%.

If measuring and controlling the plasma density to better than this accuracy is difficult we have an alternative: exploiting auto-resonance. A chirped pulse train drives a wakefield at a frequency which becomes locked to that of the driver rather than the natural plasma frequency, and hence excites a wakefield more effectively than a uniformly spaced but mis-tuned train. This phenomenon can be modelled accurately to second order using a linear sum of sinusoids and the resulting Fresnel integrals, or graphically using the Euler spiral.

The dynamic process not described by this approximation increases the wakefield amplitude for near-critical chirps, such that a pulse train chirped at the critical rate can, despite a mis-matched plasma density, produce a wakefield with an amplitude up to 91% of that caused by a perfectly tuned driver after over 150 pulses, despite an error on the initial pulse spacing of 1%. The optimal chirp rate should be chosen to be inversely proportional to the number of pulses in the train, but higher errors on the plasma density require more strongly chirped pulse trains using more intense laser pulses.

This discussion of resonant wakefield excitation comes with several provisos, foremost among which is the limitation from ion motion. As the length of a train of laser pulses increases it is inevitable that the ions, however much slower their motion than the electrons, will begin to move. This will disrupt the plasma wakefield and reduce its

amplitude. Future work will focus on how the onset of ion motion can be forestalled by using heavy ions with large charge to mass ratios. Interestingly, this depends strongly on the dimensionality of the system, with ion motion a greater problem in physical situations where the ions also experience a strong transverse force.

Chapter 3

Measuring Wakefields with Temporally Encoded Spectral Shifting

3.1 Introduction

As described in Chapter 1, it is important to measure linear and quasi-linear wakefields directly, rather than through the radiation they produce and Frequency Domain Holography (FDH) [45] and Temporally Encoded Spectral Shifting (TESS) [31] provide a sensitive measurement of such plasma density structures. In these techniques, chirped and stretched probe and reference pulses travel along the same path as a laser pulse driving a wakefield, such that the reference pulse crosses the interaction region before the driver pulse but the probe pulse co-propagates with the wakefield. Information about the density profile of the wakefield is encoded into the spectral phase of the probe, with different frequencies containing information about different points in the wakefield. Interfering the probe and reference pulses onto a spectrometer allows us to extract this spectral phase and hence to recover the density profile of the wakefield in a single shot.

The TESS and FDH techniques both use the spectral interferogram measured with the same experimental setup, but approach the analysis in different ways, with FDH attempting to reconstruct first the spectral phase of the probe pulse, and then the temporal phase, whereas the TESS analysis aims by assuming a sinusoidal wakefield to extract the frequency and amplitude of the wakefield directly from the temporal domain without reconstructing the phase. TESS is therefore well suited to rapid quantitative measurements which could be taken at high repetition rate. Given the small signal strength of weak wakefields acquired over a short interaction length, FDH and TESS are very effective, but both the measurement and the subsequent analysis can be challenging. Although FDH has been described for many years, little has been written about TESS, a much more recent technique. I therefore hope in Section 3.2 to explain clearly how the TESS signal is derived and how to use it to calculate the wakefield's amplitude and phase. In the process I shall look at some of the assumptions and reference measurements in Section 3.3.

As it stands, TESS has been derived assuming identical probe and reference pulses

with Gaussian profiles. This is clearly not a valid approximation in many situations, and so I shall attempt in Section 3.4 to determine its limitations and extend the technique to general probe and reference pulses, which need not be either identical or Gaussian. This is demonstrated with both simulated and experimental data. Similarly, TESS was only originally described for linear, purely sinusoidal wakefields but, in practice, even a moderately non-linear wakefield is no longer sinusoidal. In Section 3.5 I have described an extension to TESS which can calculate the wakefield frequency and amplitude for non-linear wakefields; this is demonstrated with simulated TESS measurements. This chapter largely derives from a paper of which I am the lead author, which is now published [70].

3.2 TESS Background

As we saw earlier, there are several different ways of describing the same physical process. Wilks et al. [42] describe how a density gradient can cause frequency blue-shifting in a probe pulse and hence refer to the process as photon acceleration. Other sources [41] describe the modulation of a probe pulse by a plasma wave in terms of Thomson scattering. We shall describe the same effect purely in terms of phase changes to a probe pulse which co-propagates with a plasma wakefield, before demonstrating how this leads to the measurement of a TESS spectrum, from which a wakefield's amplitude and phase can be calculated.

As a probe laser pulse co-propagates with an underdense plasma density structure such as a wakefield, different parts of the probe pulse experience a different refractive index and therefore acquire phase at a different rate. We have seen in the introduction

that the phase acquired can be described by:

$$\phi(\zeta) = \frac{\omega L}{c}\left[1 - \left(\frac{\omega_{\mathrm{p}}(\zeta)}{\omega}\right)^2\right]^{\frac{1}{2}} \tag{3.1}$$

$$\begin{aligned}\Delta\phi(\zeta) &\approx \frac{\omega L}{2c}\left(\frac{\omega_{\mathrm{p}}(\zeta)}{\omega}\right)^2 \\ &= \frac{\omega_{\mathrm{p0}}^2 L}{2\omega c}\frac{n_e(\zeta)}{n_{e0}}\end{aligned} \tag{3.2}$$

FDH considers a general density variation $n_e(\zeta)$, but for the purposes of TESS we will initially assume that the plasma wave is in the low amplitude linear regime, and hence may be written as a sinusoid. The final phase change is therefore split into constant and oscillating terms. For a density variation described by $n_e(\zeta) = n_{e0} + \delta n_e \sin(\omega_{\mathrm{p0}}\zeta + \theta_1)$, the phase change can be written as:

$$\begin{aligned}\Delta\phi(\zeta) &= -\frac{\omega_{\mathrm{p0}}^2 L}{2\omega_0 c}\left[1 + \frac{\delta n_e}{n_{e0}}\sin(\omega_{\mathrm{p0}}\zeta + \theta_1)\right] \\ &= \phi_0 + \phi_1 \sin(\omega_{\mathrm{p0}}\zeta + \theta_1)\end{aligned} \tag{3.3}$$

$$\text{for} \quad \phi_0 = -\frac{\omega_{\mathrm{p0}}^2 L}{2\omega_0 c} \tag{3.4}$$

$$\text{and} \quad \phi_1 = -\frac{\omega_{\mathrm{p0}}^2 L}{2\omega_0 c}\frac{\delta n_e}{n_{e0}}, \tag{3.5}$$

3.2.1 TESS and FDH

As described above and by Matlis et al. [31], a probe pulse is transformed by co-propagating with a linear plasma wave in a well defined way. Information about a linear wakefield with a plasma frequency ω_{p0} is encoded in the modulation of the spectrum of the probe by an additional phase $\phi_{\mathrm{wake}}(\zeta) = \phi_0 + \phi_1 \sin(\omega_{\mathrm{p0}}\zeta + \theta_1)$. This leads to the transformations between the electric field of the probe before entering the plasma, in temporal space $E_{\mathrm{pr}}(\zeta)$ and in spectral space $\mathcal{E}_{\mathrm{pr}}(\omega)$, and after leaving the plasma, $E'_{\mathrm{pr}}(\zeta)$ and $\mathcal{E}'_{\mathrm{pr}}(\omega)$, shown below:

$$
\begin{aligned}
E'_{\mathrm{pr}}(\zeta) &= E_{\mathrm{pr}}(\zeta) e^{i\phi_0} e^{i\phi_1 \sin(\omega_{\mathrm{p0}}\zeta + \theta_1)} \\
&= E_{\mathrm{pr}}(\zeta) \sum_{k=-\infty}^{\infty} J_k(\phi_1) e^{ik\omega_{\mathrm{p}}\zeta} e^{i(\phi_0 + k\theta_1)}. \qquad (3.6) \\
\mathcal{E}'_{\mathrm{pr}}(\omega) &= \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} E_{\mathrm{pr}}(\zeta) e^{i\phi_{\mathrm{wake}}(\zeta)} e^{-i\omega\zeta} \, \mathrm{d}\zeta \\
&= \sum_{k=-\infty}^{\infty} J_k(\phi_1) \mathcal{E}_{\mathrm{pr}}(\omega - k\omega_{\mathrm{p0}}) e^{i(\phi_0 + k\theta_1)}. \qquad (3.7)
\end{aligned}
$$

After it interacts with the plasma wave, the spectrum of the probe can be expanded as a sum of carrier waves at frequencies shifted by $k\omega_{\mathrm{p0}}$, as shown in the cartoon in Fig. 3.1a. The reference, meanwhile, remains unchanged, co-propagating with the probe a time Δt earlier:

$$
\begin{aligned}
E'_{\mathrm{r}}(\zeta) &= E_{\mathrm{r}}(\zeta + \Delta t). \qquad (3.8) \\
\mathcal{E}'_{\mathrm{r}}(\omega) &= \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} E_{\mathrm{r}}(\zeta + \Delta t) e^{-i\omega\zeta} \, \mathrm{d}\zeta \\
&= \mathcal{E}_{\mathrm{r}}(\omega) e^{+i\omega\Delta t}. \qquad (3.9)
\end{aligned}
$$

The two pulses interfere within the spectrometer to create a spectral interferogram $\mathcal{S}(\omega) = \left|\mathcal{E}'_{\mathrm{pr}}(\omega)\right|^2 + |\mathcal{E}'_{\mathrm{r}}(\omega)|^2 + \mathcal{E}'_{\mathrm{pr}}(\omega)\mathcal{E}'^{*}_{\mathrm{r}}(\omega) + c.c.$, where $c.c.$ denotes the complex conjugate of the previous term.

In FDH the interferogram is analyzed as follows, shown in the diagram in Fig. 3.2:

1. First, one uses a Fourier transform to the temporal domain to isolate the interference term, $\mathcal{E}'_{\mathrm{pr}}(\omega)\mathcal{E}'^{*}_{\mathrm{r}}(\omega)$, with a bandpass filter, as described in Takeda et al. [89], before Fourier transforming back to the spectral domain.

2. Next, one removes the reference pulse, measured through its spectral amplitude and phase, to recover the electric field of the probe after it has passed through the plasma, $\mathcal{E}'_{\mathrm{pr}}(\omega)$.

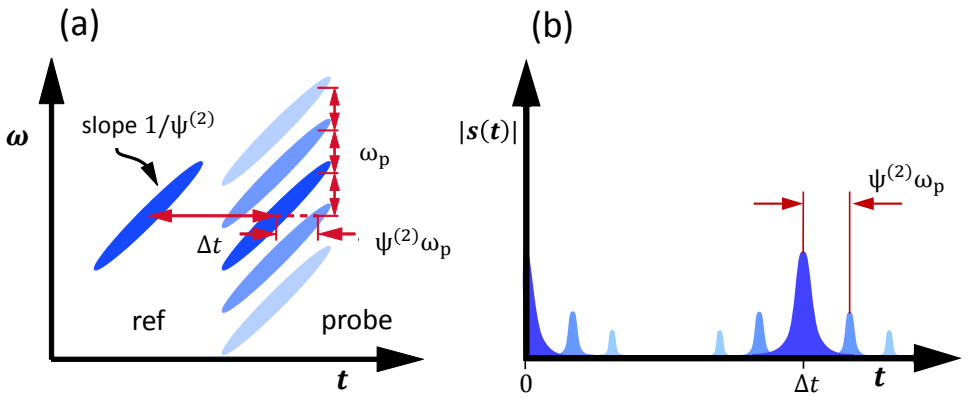


Figure 3.1: a) Frequency-time domain plots of the reference pulse, and the probe pulse after it has interacted with a sinusoidal plasma wave of frequency $\omega_{\rm p}$. Modulation of the probe pulse generates copies of the incident probe pulse, spectrally shifted by multiples of $\omega_{\rm p}$. b) The TESS signal, obtained by a Fourier transform of the recorded spectrum of the transmitted probe and reference pulses. The temporal separation of the probe and reference pulses yields a DC term at $t = 0$ and a sideband at $t = \Delta t$, and modulation of the probe causes a series of satellites (lighter blue) separated from the sideband by multiples of $\omega_{\rm p}\psi^{(2)}$

3. One can then Fourier transform this to $E'_{\rm pr}(\zeta)$, obtaining the correctly chirped probe pulse in the temporal domain.

4. Taking the argument of this and removing the temporal phase of the probe extracts the phase due to the wakefield $\phi_{\rm wake}(\zeta)$.

5. Finally, one can use this to determine the density structure, using $n_e(\zeta)/n_{e0} = (2\omega c/\omega_{\rm p0}^2 L)\phi_{\rm wake}(\zeta)$.

For an FDH reconstruction we therefore need three Fourier transforms and a complete description of the reference pulse in the spectral domain, both its amplitude and phase, and a measurement of probe pulse's temporal phase. Whereas an amplitude measurement in the spectral domain is straightforward with a spectrometer, taking measurements of the phase is not straightforward in either spectral or temporal domains, and is described in Section 3.3.2.

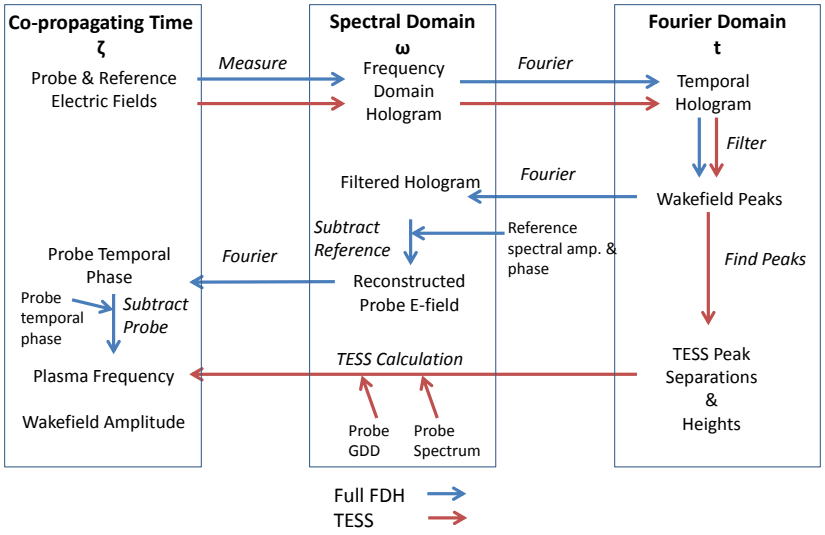


Figure 3.2: A comparison of the FDH (blue) and TESS (red) processes. Both take place across three different domains, the co-moving time ζ, the measured probe frequency ω, and the Fourier co-domain of this, the relative time t. Both require Fourier transforming the measured spectrum in order to filter away information in the relative time t, but whereas TESS becomes a peak finding problem, FDH then requires subtracting information about the reference pulse in the spectral domain, transforming to the co-moving time and then subtracting information about the probe pulse.

3.2.2 TESS Signal

In TESS however, it is only necessary to make the first Fourier transform to the temporal domain to obtain the TESS signal. As the probe electric field was modulated in frequency and both the probe and the reference pulses are chirped, the probe's electric field and the resulting spectrum are also modulated in time, as shown in the cartoon in Fig. 3.1. The peaks in the TESS signal, their heights and locations, can therefore characterize

the wakefield. Taking the Fourier transform of the spectrum $\mathcal{S}(\omega)$ yields:

$$\begin{aligned} s(t) &= \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} \mathcal{S}(\omega)\ \mathrm{d}\omega \\ &= H_{\mathrm{r,r}}(t,0) + \sum_{m=-\infty}^{\infty} g_m(\phi_1, t, \omega_\mathrm{p}) H_{\mathrm{pr,pr}}(t, m\omega_\mathrm{p}) \\ &\quad + \sum_{k=-\infty}^{\infty} J_k(\phi_1) \Big[H_{\mathrm{pr,r}}(t-\Delta t, k\omega_\mathrm{p}) e^{i(\phi_0 + k\theta_1)} \\ &\quad + H^*_{\mathrm{pr,r}}(-t-\Delta t, k\omega_\mathrm{p}) e^{-i(\phi_0+k\theta_1)} \Big] \end{aligned} \tag{3.10}$$

$$\text{where} \quad H_{\mathrm{a,b}}(t,\Omega) \equiv \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} \mathcal{E}_\mathrm{a}(\omega-\Omega)\mathcal{E}^*_\mathrm{b}(\omega) e^{i\omega t}\ \mathrm{d}\omega \tag{3.11}$$

$$\text{and} \quad g_m(\phi_1, \theta_1, t, \Omega) \equiv \sum_{n=-\infty}^{\infty} J_n(\phi_1) J_{m+n}(\phi_1) e^{in\Omega t} e^{in\theta_1} \tag{3.12}$$

Matlis et al. demonstrate that for identical Gaussian probe and reference pulses described by:

$$\mathcal{E}_\mathrm{pr}(\omega) = \mathcal{E}_\mathrm{r}(\omega) = \mathcal{E}_0(\omega) = A \exp\left[-\frac{1}{2} \left(\frac{\omega-\omega_0}{\delta\omega} \right)^2 - \frac{i}{2}\psi^{(2)}(\omega-\omega_0)^2 \right], \tag{3.13}$$

where $\psi^{(2)}$ is the group delay dispersion (GDD), we can find analytic expressions for the functions $H_{\mathrm{a,b}}(t,\Omega)$:

$$\begin{aligned} H_{\mathrm{pr,r}}(t,0) &= \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} |\mathcal{E}_0(\omega)|^2 e^{i\omega t}\ \mathrm{d}\omega \\ &= \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} A^2 \exp\left[-\frac{1}{\delta\omega^2} (\omega-\omega_0)^2 \right] e^{i\omega t}\ \mathrm{d}\omega \\ &= \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} A^2 \exp\left[-\frac{1}{\delta\omega^2} \left(\omega - \omega_0 - \frac{1}{2}\delta\omega^2 it \right)^2 + \frac{1}{\delta\omega^2} \left(\frac{\delta\omega^2 it}{2} \right)^2 \right] e^{i\omega_0 t}\ \mathrm{d}\omega \\ &= \frac{1}{\sqrt{2\pi}} \sqrt{\pi\delta\omega^2} \exp\left(-\frac{1}{4}\delta\omega^2 t^2 \right) e^{i\omega_0 t} \\ &= 2^{-0.5}\delta\omega A^2 \exp\left(-\frac{1}{4}\delta\omega^2 t^2 \right) e^{i\omega_0 t}, \end{aligned} \tag{3.14}$$

and:

$$\begin{aligned}
H_{\mathrm{pr,r}}(t,\Omega) &= \frac{1}{\sqrt{2\pi}}\int_{-\infty}^{\infty} A^2 \exp\left\{-\frac{1}{2\delta\omega^2}\left[(\omega-\Omega-\omega_0)^2+(\omega-\omega_0)^2\right]\right\} \\
&\qquad \cdot \exp\left\{-\frac{i\psi^{(2)}}{2}\left[(\omega-\Omega-\omega_0)^2-(\omega-\omega_0)^2\right]\right\} e^{i\omega t}\,\mathrm{d}\omega \\
&= \frac{1}{\sqrt{2\pi}}\int_{-\infty}^{\infty} A^2 \exp\left\{-\frac{1}{2\delta\omega^2}\left[2\left(\omega-\omega_0-\frac{1}{2}\Omega\right)^2+\frac{1}{2}\Omega^2\right]\right\} \\
&\qquad \cdot \exp\left\{-\frac{i\psi^{(2)}}{2}\left[-2\Omega(\omega-\omega_0)+\Omega^2\right]+i\omega t\right\}\,\mathrm{d}\omega \\
&= \frac{1}{\sqrt{2\pi}}\int_{-\infty}^{\infty} A^2\left[-\frac{1}{\delta\omega^2}\left(\omega-\omega_0-\frac{1}{2}\Omega\right)^2\right]\exp\left[+i\omega\left(t+\psi^{(2)}\Omega\right)\right]\,\mathrm{d}\omega \\
&\qquad \cdot \exp\left(-\frac{1}{4}\frac{\Omega^2}{\delta\omega^2}\right)\exp\left[-i\psi^{(2)}\Omega\left(\omega_0+\frac{1}{2}\Omega\right)\right] \\
&= H_{\mathrm{pr,r}}\left(t+\psi^{(2)}\Omega,0\right)\exp\left(\frac{i\Omega t}{2}\right)\exp\left(-\frac{1}{4}\frac{\Omega^2}{\delta\omega^2}\right) \qquad (3.15)
\end{aligned}$$

For Gaussian pulses $H_{\mathrm{pr,r}}(t,0)$ is peaked at $t=0$ and hence $H_{\mathrm{pr,r}}(t,\Omega)$ is peaked at $t=-\psi^{(2)}\Omega$. The phase change due to the wakefield therefore causes a series of equally spaced satellite peaks in the Fourier domain, which occur at times $t_k=\Delta t+\psi^{(2)}k\omega_{p0}$ and with heights that are in ratios to the $k=0$ satellite peak described by:

$$r_k \equiv \left|\frac{s(t_k)}{s(\Delta t)}\right| = \frac{J_k(\Delta\phi_{\mathrm{p}})}{J_0(\Delta\phi_{\mathrm{p}})}\mathcal{F}_{\mathrm{Gauss}}(k\omega_{\mathrm{p0}}) \qquad (3.16)$$

for an overlap factor $\mathcal{F}_{\mathrm{Gauss}}=\exp\left[-\frac{1}{4}\left(\frac{k\omega_{\mathrm{p0}}}{\delta\omega}\right)^2\right]$. Measuring the peak locations and heights allows us to extract the wakefield frequency ω_{p0} and amplitude $\delta n_e/n_{e0}\propto\Delta\phi_{\mathrm{p}}$.

The TESS analysis procedure is therefore as follows:

1. First, one must Fourier transform the spectrum to the time domain to obtain the TESS signal $|s(t)|$.

2. Next, one finds the first order satellite at $k=1$ and measures its separation from the sideband $t_1-\Delta t$; the plasma frequency can then be calculated using the measured GDD as $\omega_{\mathrm{p0}}=(t_1-\Delta t)/\psi^{(2)}$.

3. To find the wakefield amplitude, one must measure the ratio r_1 of the heights of the first order satellite and the sideband and calculate the overlap factor $\mathcal{F}(\omega_{\mathrm{p}0})$ from the measured spectra

4. The phase ϕ_1 can then be extracted by solving $J_1(\phi_1)/J_0(\phi_1) = r_1/\mathcal{F}(\omega_{\mathrm{p}0})$

5. Finally, one can calculate the wakefield amplitude from $\delta n_e/n_{e0} = \left(2\omega_0 c/\omega_{\mathrm{p}0}^2 L\right)\phi_1$

In TESS the phase reconstruction problem has therefore become a peak finding problem. Good filtering and background subtraction is essential as the wakefield can be very weak, with even first order $k = 1$ satellite peaks many times smaller than the main sideband peak. Once the peaks have been found, however, the calculation is fairly simple, and requires only the GDD, $\psi^{(2)}$, to find the plasma frequency. Again, measuring this is described in Section 3.3.2. Finding the wakefield amplitude requires an additional measurement of the probe's spectrum and of the interaction length.

3.3 Practical Considerations

3.3.1 Probe Slippage

As the group velocity of a laser pulse at a frequency ω_{L} is $v_{\mathrm{gL}} = c\left(1 - \frac{\omega_{\mathrm{p}}^2}{\omega_{\mathrm{L}}^2}\right)^{\frac{1}{2}}$, the difference between a pump pulse at ω_1 and a probe pulse at ω_0 travelling through an underdense plasma at $\omega_{\mathrm{p}} \ll \omega_1 < \omega_0$ is approximately $\Delta v_{\mathrm{g}} = \frac{\omega_{\mathrm{p}}^2 c}{2}\left({\omega_1}^{-2} - {\omega_0}^{-2}\right)$. The two pulses therefore slip by a full plasma wavelength over a length approximately $L_{\mathrm{slip}} = 2\lambda_{\mathrm{p}}^3(\lambda_1^2 - \lambda_0^2)^{-1}$. If the probe is frequency doubled from the same wavelength as the pump pulse, $\omega_0 = 2\omega_1$, this becomes $\Delta v_{\mathrm{g}} = \frac{3}{8}\frac{\omega_{\mathrm{p}}^2 c}{\omega_1^2}$ and $L_{\mathrm{slip}} = \frac{8}{3}\frac{\lambda_{\mathrm{p}}^3}{\lambda_1^2}$. For an 800 nm pump pulse and 400 nm probe pulse in an $n_e = 10^{18}$ cm^{-3} plasma, the difference in velocities is just $\Delta v_{\mathrm{g}} = 0.0002c$, and the pulses can travel a distance of up to $L_{\mathrm{slip}}/10 = 1.6$ cm without slipping by more than a tenth of the plasma wavelength.

This is highly relevant to measuring laser wakefields as the phase velocity of the laser wakefield is set by the group velocity of the driving laser pulse. As we saw earlier, we

wish a probe pulse to co-propagate with this laser wakefield over a substantial interaction length, and we want the phase of the wakefield at a particular point ζ in the probe to remain constant. This means we wish to minimise the slippage between the group velocity of the pump and probe laser pulses over a substantial co-propagation length.

However, the dephasing length of a laser wakefield accelerator driven by a laser at a frequency ω_{L} is $L_{\mathrm{dp}} = \frac{1}{2}\frac{\lambda_{\mathrm{p}}^3}{\omega_{\mathrm{L}}^2}$, and currently there is no additional advantage in making a laser wakefield acceleration stage longer than this limit. Over the full acceleration length of an optimum laser plasma accelerator driven by a pump pulse at ω_1 the slippage between the probe pulse and the plasma wave will therefore be $\frac{3}{16}\lambda_{\mathrm{p}}$. This would blur the accumulated phase and reduce the measured wakefield amplitude by just 5% and shows that pump-probe slippage can safely be neglected for practical laser wakefield acceleration experiments.

3.3.2 Retrieving Phase from the Ionisation Front

For both FDH and TESS we need to know how the comoving time ζ is mapped onto the frequency $\omega(\zeta)$ which encodes the density variation at that time. In FDH this means calculating the full function for both the temporal phase $\phi(\zeta)$ and the spectral phase $\psi(\omega)$ whereas in TESS only the 2nd order coefficient of the spectral phase need be calculated, the GDD $\psi^{(2)}$, and higher order contributions are assumed to be negligible.

$$\phi(\zeta) \approx \phi^{(0)} + \omega_0\zeta + \frac{1}{2}\phi^{(2)}\zeta^2 + ... \tag{3.17}$$

$$\psi(\omega) \approx \psi^{(0)} + \psi^{(1)}(\omega - \omega_0) + \frac{1}{2}\psi^{(2)}(\omega - \omega_0)^2 + ... \tag{3.18}$$

$$\omega(\zeta) = \frac{d\phi}{d\zeta}, \quad \zeta(\omega) = \frac{d\psi}{d\omega}$$

We can calculate the full temporal phase $\phi(\zeta)$ up to an additive constant by overlapping the probe with the pump and measuring the change in the spectrum as the delay between the pump and probe is changed, hence changing the frequency at which

the ionisation front occurs in the probe spectrum. In this method, we choose a feature which occurs at a specific point in co-propagating time ζ, for instance the point where the probe is most strongly blue shifted near the ionisation front, and measure it at a certain frequency ω on the spectrum, see Fig. 3.3. The accuracy of this measurement is however limited by the difficulty in identifying the precise frequency of a feature in the spectrum, particularly if the spectrum is highly structured and the frequency depends on the spatial location. This results in a significant error on the GDD measurement.

By changing a stage position we change the time at which the feature occurs and hence measure $\omega(\zeta)$. This is the first derivative of the temporal phase, so integrating it yields $\phi(\zeta)$ up to a constant. However, in order to subtract information about the reference we also require the spectral phase, which we can obtain in one of several ways. If the function $\omega(\zeta)$ is bijective, we can invert it to obtain $\zeta(\omega)$, which is the first derivative of the spectral phase. Integrating this gives $\psi(\omega)$ up to a constant.

The constraint that $\omega(\zeta)$ is bijective over the domain is equivalent to saying that $\omega(\zeta)$ is single valued and has no turning points in the domain. This is achieved if the quadratic term in the temporal phase dominates at small displacements in ζ: $|\phi^{(2)}\zeta^2| \gg \mathcal{O}(\zeta^3)$.* If there are no turning points where $\frac{\mathrm{d}\omega}{\mathrm{d}\zeta} = 0$, we can then identify the GDD as the reciprocal of the gradient of our measured $\omega(\zeta)$:

$$\psi^{(2)} = \frac{d^2\psi}{d\omega^2}\bigg|_{\omega_0} = \frac{d\zeta}{d\omega}\bigg|_{\omega_0} = \left(\frac{d\omega}{d\zeta}\bigg|_{\zeta_0}\right)^{-1} = \left(\frac{d^2\phi}{d\zeta^2}\bigg|_{\zeta_0}\right)^{-1} = \frac{1}{\phi^{(2)}} \tag{3.19}$$

In all the cases we measured, the linear term in the measured $\lambda(\zeta)$ (proportional to the second order phase $\phi^{(2)}$) strongly dominated and hence the wavelength was monotonically increasing (see e.g. Fig. 3.3). The mapping between frequency and time is therefore bijective and invertible as required, allowing us to easily calculate temporal

*To see this, first recognise that we also want $\omega(\zeta)$ to be smooth and continuous over the domain. It must therefore be monotonic within the domain, which means its turning points $\frac{d\omega}{d\zeta} = 0$ lie outside the domain. Including only third order terms and lower, its turning points lie at $\zeta_{\mathrm{tp}} - \zeta_0 = -\frac{\phi^{(2)}}{\phi^{(3)}}$ so we require $|\zeta - \zeta_0| \ll \left|\frac{\phi^{(2)}}{\phi^{(3)}}\right|$.

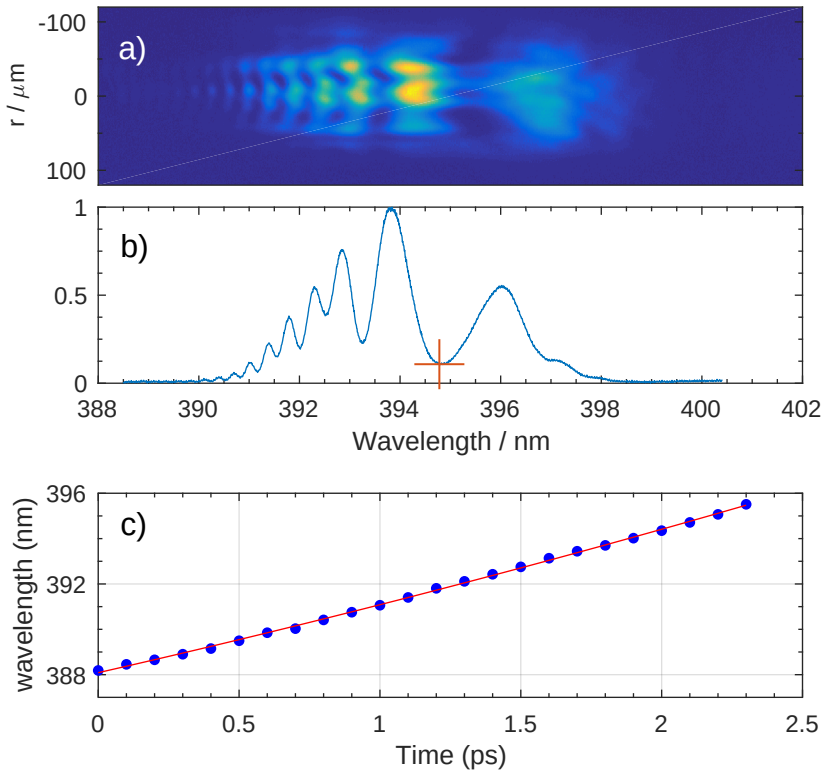


Figure 3.3: An illustration of the process needed to find the temporal and spectral phase of the probe pulse. a) The intensity measured on the spectrometer when the pump and probe pulses are overlapped, showing the signal from an ionization front. As we change the pump-probe delay we can track the wavelength at which there is an intensity minimum caused by strong blue shifting, shown in the lineout in b). This gives us a mapping $\lambda(\zeta)$, which is shown in c). Data for c) from group member Rob Shalloo.

and spectral phases and the GDD $\psi^{(2)}$ for both the probe and reference pulses.

3.3.3 Errors from Finite Bandwidth

An issue common to both FDH and TESS is interpreting the spectrometer signal. This is measured as a function of λ on a uniform grid because of the diffraction condition $kd\sin\theta = 2n\pi$. The Fourier transform, however, assumes that the signal is measured on a grid uniform in ω in order to convert to the Fourier co-domains t and ζ. As the bandwidth of the probe pulse is substantial for short pulse lasers these two measures begin to diverge, and correcting for this requires interpolating from the spectrometer grid to a new grid regular in ω. Without this interpolation there is an error on the stated frequency, and hence on the measured plasma frequency, on the order of $\delta\lambda/\lambda_0$, which can be on the level of a few percent.

We have made another approximation when calculating the phase accumulated by the probe pulse. As this pulse is chirped, the phase accumulated is actually:

$$\begin{aligned}\Delta\phi(\zeta) &= -\frac{\omega_{\mathrm{p}0}^2 L}{2\omega(\zeta)c}\left[1+\frac{\delta n_e}{n_{e0}}\sin(\omega_{\mathrm{p}}\zeta+\psi)\right] \\ &= -\frac{\omega_{\mathrm{p}0}^2 L}{2\omega_0 c}\left(1+\frac{\delta\omega(\zeta)}{\omega_0}\right)^{-1}\left[1+\frac{\delta n_e}{n_{e0}}\sin(\omega_{\mathrm{p}}\zeta+\psi)\right] \\ &\approx \phi_0\left[1+\frac{\delta n_e}{n_{e0}}\sin(\omega_{\mathrm{p}}\zeta+\psi)-\frac{\delta\omega(\zeta)}{\omega_0}\right], \end{aligned} \tag{3.20}$$

where $\delta\omega(\zeta) = \omega(\zeta) - \omega_0$ depends on the phase profile of the probe pulse.

Again, this error is on the order of $\delta\lambda/\lambda_0$ and hence can be substantial for short pulse laser systems with large bandwidths of 10s of nanometres. If the wakefield amplitude is on the few percent level it can be washed out by this change, making interpreting an FDH phase reconstruction difficult. However, while this error is substantial, it is not sinusoidal. If we neglect phase higher than second order we can approximate $\delta\omega(\zeta) \approx (\zeta - \zeta_0)/\psi^{(2)}$. This affects the transformation of the electric field of the probe

with an additional term which looks like a frequency shift $e^{ia(\zeta-\zeta_0)}$ for $a = \phi_0/\psi^{(2)}\omega_0$:

$$\begin{aligned} E'_{\rm pr}(\zeta) &\approx E_{\rm pr}(\zeta)e^{i\phi_0}e^{i\phi_1\sin(\omega_{\rm p}\zeta+\theta_1)}e^{ia(\zeta-\zeta_0)} \\ &= E_{\rm pr}(\zeta)\sum_{k=-\infty}^{\infty} J_k(\phi_1)e^{ik\omega_{\rm p}\zeta}e^{i(\phi_0+k\theta_1)}e^{ia(\zeta-\zeta_0)} \qquad (3.21) \\ \mathcal{E}'_{\rm pr}(\omega) &= \frac{1}{\sqrt{2\pi}}\int_{-\infty}^{\infty} E'_{\rm pr}(\zeta)e^{-i\omega\zeta}\,{\rm d}\zeta \\ &= \sum_{k=-\infty}^{\infty} J_k(\phi_1)\mathcal{E}_{\rm pr}(\omega+a-k\omega_{\rm p})e^{i(\phi_0+k\theta_1)}e^{-ia\zeta_0}. \qquad (3.22) \end{aligned}$$

While this changes the position of the TESS peaks as $k\omega_{\rm p0} \to k\omega_{\rm p0} - a$ and hence $t_k \to t_k - \psi^{(2)}a = t_k - \omega_{\rm p0}^2 L/2\omega_0^2 c$, it neither changes the separation between the peaks nor their relative amplitudes, and hence has no effect on TESS calculations. This reflects the slightly different group velocities for frequencies at the front and back of the probe pulse. Indeed, for a 400 nm probe pulse co-propagating with an $n_e \sim 10^{18}$ cm^{-3} plasma for $L = 1$ mm, the change in the position of the peaks is just 2 fs. That TESS measures the separation of peaks and their relative heights, rather than their positions or absolute heights, makes the process robust to some of the other phase variations that might arise in a probe pulse.

3.4 Arbitrary Probe and Reference Pulses

3.4.1 Non-Gaussian Pulses

So far we have followed the previous work by Matlis et al. [31] and assumed that the incident probe and reference pulses are identical Gaussian pulses with no 3rd or higher order spectral phase: $\mathcal{E}_{\rm pr}(\omega) = \mathcal{E}_{\rm r}(\omega) = \mathcal{E}_0(\omega) \equiv Ae^{-\frac{1}{2}(1+i\sigma)\left(\frac{\omega-\omega_0}{\delta\omega}\right)^2}$. Here we remove the restriction on the spectral shape of the probe and reference pulses, whilst retaining the approximation that third and higher order phases can be neglected. Hence we may

write:

$$\mathcal{E}_{\mathrm{a}}(\omega) = |\mathcal{E}_{\mathrm{a}}(\omega)| \exp\left\{ i \left[\psi_{\mathrm{a}}^{(0)} + \psi_{\mathrm{a}}^{(1)} (\omega - \omega_0) + \frac{1}{2}\psi_{\mathrm{a}}^{(2)} (\omega - \omega_0)^2 + \ldots \right] \right\}, \tag{3.23}$$

where the first order spectral phase $\psi_{\mathrm{a}}^{(1)}$ is the group delay, and the second order spectral phase $\psi_{\mathrm{a}}^{(2)}$ is the GDD. The group delay describes the arrival time of the central frequency ω_0 before the delay Δt has been introduced; for our incident probe and reference pulses $E_{\mathrm{pr}}(\zeta)$ and $E_{\mathrm{r}}(\zeta)$ we have defined $\psi_{\mathrm{pr}}^{(1)} = \psi_{\mathrm{r}}^{(1)}$ without loss of generality.

Equation 3.10 is still valid for these generalized probe and reference pulses, and one can show (see Appendix B) that, if the difference in the 2nd order spectral phase of the probe and reference pulses is sufficiently small, $\psi_{\mathrm{r}}^{(2)} \approx \psi_{\mathrm{pr}}^{(2)} \approx \psi^{(2)}$, the function $H_{\mathrm{pr,r}}(t, \Omega)$ peaks at $t_\Omega = \psi^{(2)}\Omega + \psi_{\mathrm{r}}^{(1)} - \psi_{\mathrm{pr}}^{(1)} = \psi^{(2)}\Omega$, at which time the peak amplitude is described by a cross-correlation:

$$|H_{\mathrm{pr,r}}(t_\Omega, \Omega)| = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} |\mathcal{E}_{\mathrm{pr}}(\omega - \Omega)| \, |\mathcal{E}_{\mathrm{r}}(\omega)| \, \mathrm{d}\omega \tag{3.24}$$

From equations (3.10) and (3.24) we see that in general the TESS signal is again composed of the DC peak at $t = 0$, the two familiar delay sidebands at $\pm\Delta t$, and a series of equally spaced satellites either side of the delay sidebands, with the k-th order satellite of the $t_0 = \Delta t$ peak located at $t_k = \Delta t + k\omega_{\mathrm{p0}}\psi^{(2)}$, see Fig. 3.1. Again, this gives us a measurement of the wakefield frequency, using $\omega_{\mathrm{p0}} = (t_k - \Delta t)/k\psi^{(2)}$, and hence an estimate of the electron density of the plasma, as $n_{e0} = \omega_{\mathrm{p0}} m_e \epsilon_0 / e^2$.

Additionally, the height of the satellite relative to that of the sideband is described by:

$$r_k = \frac{J_k(\phi_1)}{J_0(\phi_1)} \mathcal{F}(k\omega_{\mathrm{p}}), \tag{3.25}$$

where the spectral overlap factor $\mathcal{F}$ is now given by:

$$\mathcal{F}(\Omega) \equiv \frac{\int_{-\infty}^{\infty} |\mathcal{E}_{\mathrm{pr}}(\omega-\Omega)|\,|\mathcal{E}_{\mathrm{r}}(\omega)|\,\mathrm{d}\omega}{\int_{-\infty}^{\infty} |\mathcal{E}_{\mathrm{pr}}(\omega)|\,|\mathcal{E}_{\mathrm{r}}(\omega)|\,\mathrm{d}\omega} = \frac{\int_{-\infty}^{\infty} \sqrt{I_{\mathrm{pr}}(\omega-\Omega)I_{\mathrm{r}}(\omega)}\mathrm{d}\omega}{\int_{-\infty}^{\infty} \sqrt{I_{\mathrm{pr}}(\omega)I_{\mathrm{r}}(\omega)}\mathrm{d}\omega}, \tag{3.26}$$

where I_{pr} and I_{r} are the spectral intensities of the *input* probe and reference pulses. For probe and reference pulses with a bandwidth $\delta\omega$, we would expect this spectral overlap factor to become important when $\omega_{\mathrm{p0}} \gtrsim \delta\omega$.

For the case of identical and idealised Gaussian probe and reference pulses described in the spectral domain as $\sqrt{I(\omega)} = |\mathcal{E}_0(\omega)| = Ae^{-\frac{1}{2}\left(\frac{\omega-\omega_0}{\delta\omega}\right)^2}$ with a bandwidth $\delta\omega$, we find:

$$\mathcal{F}_{\mathrm{Gauss}}(k\Omega) = \exp\left[-\left(\frac{1}{4}\frac{k\Omega}{\delta\omega}\right)^2\right], \tag{3.27}$$

which agrees with the result found by Matlis et al. described earlier.

Using the expression for Gaussian pulses can therefore result in a substantial miscalculation of the wakefield amplitude. In recent experiments [71] we measured low amplitude wakefields with $\delta n_e/n_{e0} \sim 1\%$ using frequency doubled probe and reference pulses, the spectra of which were far from Gaussian, as shown in Fig. 3.5a. The spectrum had a bandwidth of approximately $\delta\omega \approx 40$ rad ps^{-1} and therefore an electron density of only $n_e \approx 0.5 \times 10^{18}$ cm^{-3} produces a plasma frequency $\omega_{\mathrm{p}} \approx \delta\omega$. As a consequence we found that, even at low pressures, $\mathcal{F}_{\mathrm{Gauss}}(\omega_{\mathrm{p}})$ differed from the true overlap factor calculated from the measured spectrum, making it unsuitable for accurately measuring wakefield amplitudes in real experiments. For example, for an electron density of $n_e \approx 2.5 \times 10^{18}\,\mathrm{cm}^{-3}$, a Gaussian approximation reduced the calculated wakefield amplitude by a factor of approximately two. This is shown in Fig. 3.5b-c.

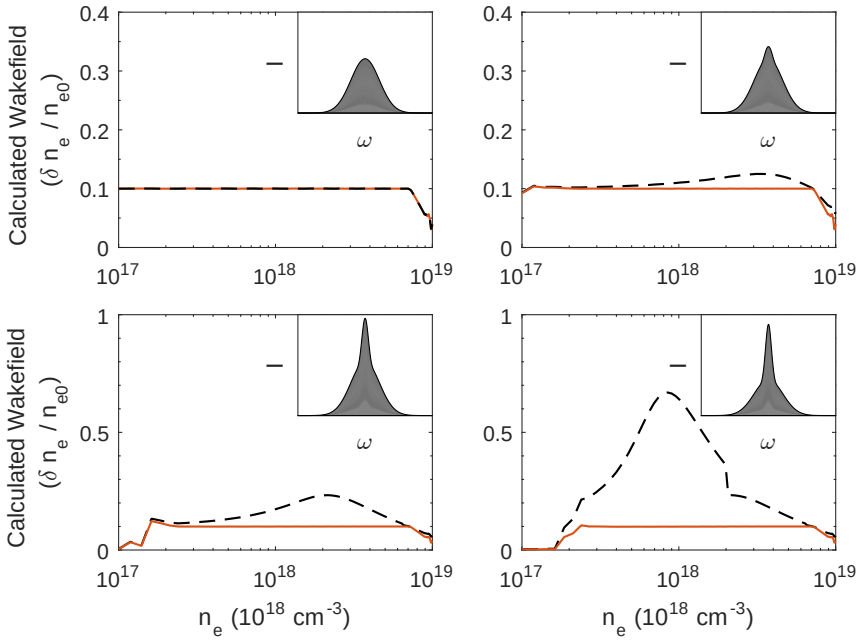


Figure 3.4: The recovered wakefield amplitude as a function of electron density, calculated from a simulated spectral hologram using TESS. The hologram was calculated for a $\delta n_e/n_{e0} = 10\%$ plasma wakefield and for a variety of example probe spectra (shown in insets). The solid black line describes the wakefield amplitude recovered using the true overlap factor calculated from eqn. (3.26), whereas the dashed red line is calculated using the Gaussian approximation in eqn. (3.27).

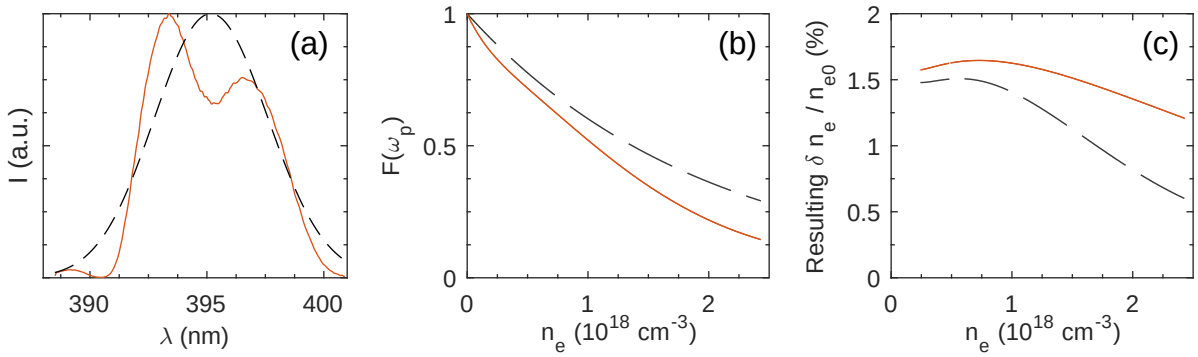


Figure 3.5: a) The measured spectrum of the probe pulse (solid red line) and a Gaussian fit to it (dashed black line). b) Comparison of the spectral overlap factors $\mathcal{F}(\omega_{\rm p})$ (solid, red) and $\mathcal{F}_{\rm Gauss}(\omega_{\rm p})$ (dashed, black) evaluated at the first order TESS peak as a function of electron density, assuming that the hydrogen gas was fully ionized by the driving laser. c) Deduced wakefield amplitude as a function of plasma density assuming spectral overlap factors $\mathcal{F}(\omega_{\rm p})$ (solid, red) and $\mathcal{F}_{\rm Gauss}(\omega_{\rm p})$ (dashed, black).

3.4.2 Non-Equal GDDs

In general, the probe and reference pulses will have different GDDs; this difference, which is usually small, arises from differences in the materials present in their optical paths. Whereas the previous work by Matlis et al. [31] has ignored these effects, they may cause changes in the locations of the satellites, and hence on the deduced plasma frequency. Using the Fourier shift theorem it can be shown that $H_{\rm pr,r}(t, k\omega_{\rm p0}) = H^*_{\rm r,pr}(-t, -k\omega_{\rm p0})e^{{\rm i}k\omega_{\rm p0}t}$ and the satellite peak location must depend equally on the probe and reference pulses. I show in Appendix B that for the case of non-identical Gaussian pulses the satellite peak is located at $t = \Delta t + k\omega_{\rm p0}\psi^{(2)}_{\rm eff}$, where $\psi^{(2)}_{\rm eff}$ is the mean of the probe and reference GDDs, weighted by the square of their bandwidths:

$$\psi^{(2)}_{\rm eff} = \frac{\delta\omega^2_{\rm pr}\psi^{(2)}_{\rm pr} + \delta\omega^2_{\rm r}\psi^{(2)}_{\rm r}}{\delta\omega^2_{\rm pr} + \delta\omega^2_{\rm r}} \tag{3.28}$$

In the case where probe and reference pulses have the same bandwidth $\delta\omega_{\rm pr} \approx \delta\omega_{\rm r}$,

this reduces to the arithmetic mean GDD: $\psi^{(2)}_{\text{eff}} = \frac{1}{2}\left[\psi^{(2)}_{\text{pr}} + \psi^{(2)}_{\text{r}}\right]$. In our experiments [71] we measured a difference in GDDs of around $1,000$ fs^2 for GDDs on the scale of $20,000\,\text{fs}^2$. In this case, using the probe GDD alone and assuming $\psi^{(2)}_{\text{eff}} \approx \psi^{(2)}_{\text{pr}}$ would result in an error in the calculated plasma frequency of around 2%.

3.5 Non-Linear Wakefields

It is now worth considering the extension of TESS to the characterization of non-linear plasma wakefields, where electrons in the plasma wave travel at relativistic velocities. The mass of the electrons increases as $m_e \rightarrow \gamma m_{e0}$, using the usual relativistic factor $\gamma = (1-\beta^2)^{-0.5}$. The plasma frequency is therefore reduced $\omega_{\text{p}} \rightarrow \omega_{\text{p0}}/\sqrt{\bar{\gamma}}$, where $\bar{\gamma}$ is the effective γ factor, averaged over one electron oscillation. If we know the non-relativistic plasma frequency ω_{p0} and can measure the wakefield frequency ω_{p}, we can deduce $\bar{\gamma}$ and hence estimate the wakefield amplitude, with larger $\bar{\gamma}$ corresponding to higher amplitude wakes. However, if we wish to characterize the wakefield more accurately, we must calculate the TESS spectrum in more detail, for first quasi-linear and then strongly non-linear wakes.

3.5.1 Quasi-Linear Wakefields

In the quasi-linear regime, we can calculate plasma wave solutions in one dimension in a cold and underdense plasma, where the phase velocity of the plasma wave $v_{\text{p}} \approx c$. Following Akhiezer and Noble [90], [91], the plasma wave amplitude can be described by the maximum electron velocity through $\beta_{\text{max}} = v_{\text{e,max}}/c$, where $\delta n_e/n_{e0} = \beta/(1-\beta)$ and $|\beta| \leq \beta_{\text{max}}$. The plasma waves are linear in the regime $\beta_{\text{max}} \ll 1$. Fig. 3.6a shows the calculated relative wave amplitude $\delta n_e/n_{e0}$ as a function of the co-moving coordinate ζ for two values of β_{m} at a plasma density of $n_{e0} = 10^{18}$ cm^{-3}. It can be seen that as β_{m} increases, the wave becomes more sharply peaked and the period of the plasma wave is increased as $\tau_{\text{p}} = \sqrt{\bar{\gamma}}\tau_{\text{p0}}$.

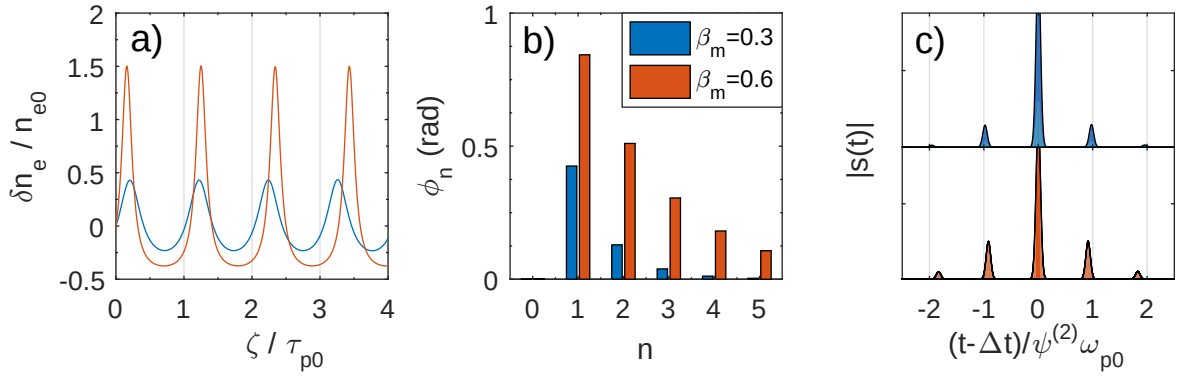


Figure 3.6: a) Calculated density profiles of quasi-linear plasma waves against time for wakefield amplitudes of $\beta_{\max} = 0.3$ (blue) and $\beta_{\max} = 0.6$ (red). b) The phase change resulting from these plasma waves expanded as a sum of harmonics $\phi(\zeta) = \sum_n \phi_n \sin(n\omega_p + \theta_n)$. c) The TESS spectra calculated for the plasma waves in a). The upper panel shows the case $\beta_{\max} = 0.3$ (blue) and the lower panel shows $\beta_{\max} = 0.6$ (red).

As the wakefield remains both periodic, with a period τ_p, and continuous, the phase change acquired by the probe can be decomposed into a linear combination of a finite number of harmonics: $\phi_{\text{wake}}(\zeta) \approx \sum_{n=1}^{N} \phi_n \sin(n\omega_p\zeta + \theta_n)$, where $\omega_p = \omega_{p0}/\sqrt{\bar{\gamma}}$ and N can be arbitrarily large. Fig. 3.6b shows the amplitudes ϕ_n of the first five harmonics resulting from a 400 nm probe co-propagating over a distance of $L = 1$ mm with the plasma waves shown in Fig. 3.6a at a plasma density of $n_{e0} = 10^{18}$ cm^{-3}. Whereas at $\beta_m = 0.3$ the phase change is dominated by the fundamental $n = 1$, as β_m increases the wave becomes more non-linear and the higher harmonics at $n > 1$ become relatively more significant. The Jacobi-Anger expansion can be applied to each of these N harmonics as in the case of a linear wakefield, and so the electric field of the transmitted probe can

be expanded as:

$$\begin{aligned} E'_{\rm pr}(\zeta) &= E_{\rm pr}(\zeta) \prod_{n=1}^{N} e^{i\phi_n \sin(n\omega_{\rm p}\zeta+\theta_n)} \\ &= E_{\rm pr}(\zeta) \prod_{n=1}^{N} \left[\sum_{k_n=-\infty}^{\infty} J_{k_n}(\phi_n)\ e^{ik_n(n\omega_{\rm p}\zeta+\theta_n)} \right] \\ &= E_{\rm pr}(\zeta) \sum_{k_1=-\infty}^{\infty} \dots \sum_{k_N=-\infty}^{\infty} J_{k_1}(\phi_1) \dots J_{k_N}(\phi_N) e^{i\left(\sum_{n=1}^{N} k_n n\omega_{\rm p}\right)\zeta} e^{i\sum_{n=1}^{N} k_n\theta_n} \end{aligned} \tag{3.29}$$

$$\mathcal{E}'_{\rm pr}(\omega) = \sum_{k_1=-\infty}^{\infty} \dots \sum_{k_N=-\infty}^{\infty} J_{k_1}(\phi_1) \dots J_{k_N}(\phi_N) e^{i\sum_{n=1}^{N} k_n\theta_n} \mathcal{E}_{\rm pr}\left(\omega - \sum_{n=1}^{N} k_n n\omega_{\rm p}\right) \tag{3.30}$$

The transmitted probe spectrum therefore contains contributions from the N (in principle, infinite) harmonics, each of which in turn produces an infinite set of carrier waves, labelled by the integers k_n, which can be both positive or negative. By truncating the phase expansion at finite N, each term is uniquely labelled with a set of N different indices, from k_1 to k_N. However, as $|k_n|$ increases, the amplitude of the carrier wave decreases since $|J_{k_n}(\phi_n)| \to 0$ as $|k_n| \to \infty$. For any finite ϕ_n the carrier wave expansion can therefore be terminated at a finite value of $|k_n| \leq K$. The larger the phase ϕ_n, the larger K must be for accurate reconstruction of the wakefield. Every term in this expansion can therefore be described with a point on a discrete N-dimensional grid with $2K+1$ points on each side, although both K and N can be arbitrarily large; each point is described by $\mathbf{k} = (k_1, k_2, k_3, k_4, \dots, k_N)$ and there are $(2K+1)^N$ such points in the grid. The reference pulse is the same as before, and the interference between the two gives rise to a TESS signal. We shall consider only the term corresponding to the peaks

near $t = \Delta t$, which arises from the interference term $\mathcal{E}'_{\rm pr}(\omega)\mathcal{E}'^{*}_{\rm r}(\omega)$, i.e.:

$$s(t) = \ldots + \sum_{k_1=-\infty}^{\infty} \ldots \sum_{k_N=-\infty}^{\infty} J_{k_1}(\phi_1)\ldots J_{k_N}(\phi_N) e^{i\sum_{n=1}^{N} k_n\theta_n} H_{\rm pr,r}\left(t - \Delta t, \sum_{n=1}^{N} k_n n\omega_{\rm p}\right) + \ldots$$

$$= \ldots + \sum_{\kappa=-\infty}^{\infty} Z_\kappa\left(\{\phi_n\},\{\theta_n\}\right) H_{\rm pr,r}\left(t - \Delta t, \kappa\omega_{\rm p}\right) + \ldots \tag{3.31}$$

$$\text{where} \quad Z_\kappa\left(\{\phi_n\},\{\theta_n\}\right) = \sum_{\mathbf{k}\in S_\kappa} J_{k_1}(\phi_1)\ldots J_{k_N}(\phi_N) e^{i\sum_{n=1}^{N} k_n\theta_n} \tag{3.32}$$

$$\text{for a set} \quad S_\kappa = \left\{(k_1,\ldots,k_N) : \sum_{n=1}^{N} n k_n = \kappa\right\} \tag{3.33}$$

In eqn. (3.31) we have divided the set of points $\mathbf{k}$ in the grid into subsets, each labelled with the new integer κ, which is the effective harmonic order. Each of these subsets S_κ comprises those points from the original N dimensional grid which also lie upon a particular $N-1$ dimensional plane, which is perpendicular to the vector $\mathbf{v} = (1, 2, 3, \ldots, N)$. The new label κ is then the plane number, $\kappa = \mathbf{v}\cdot\mathbf{k}$, which can be both positive or negative. For instance, the $\kappa = 2$ subset contains points $\mathbf{k} = (2,0,0,0,\ldots)$, $(0,1,0,0,\ldots)$, $(-1,0,1,0,\ldots)$ and $(2,-2,0,1,\ldots)$, amongst many others. To completely span the original grid we must consider all planes up to $|\kappa| \leq \kappa_{\rm max} = K + 2K + \ldots + NK = KN(N+1)/2$.

The TESS signal $s(t)$ now looks very similar to that for a linear wakefield, with the substitution of a new complex number Z_κ, which includes contributions from all of the wakefield's frequency components, instead of the Bessel function J_k, which only accounts for the fundamental frequency. The crucial part of this signal again consists of a sideband peak at $t_0 = \Delta t$ and a series of satellite peaks spaced around it at times $t_\kappa = \Delta t + \kappa\psi^{(2)}\omega_{\rm p}$. For a quasi-linear wakefield, however, each peak will contain contributions from many different frequency components; the κ-th order peak will have contributions from all terms in S_κ, including all harmonics such that $\sum_{n=1}^{N} n k_n = \kappa$.

Fig. 3.6c shows a simulation of the TESS satellites resulting from the quasi-linear plasma waves with $\beta_{\rm m} = 0.3$ and $\beta_{\rm m} = 0.6$ shown in Fig. 3.6a. While the relative heights

of the satellite peaks are slightly greater for $\beta_{\mathrm{m}} = 0.6$, the form of the signal is largely unchanged from that of a linear wakefield. The satellites are located slightly closer to the origin because the plasma frequency is reduced as $\omega_{\mathrm{p}} = \omega_{\mathrm{p0}}/\sqrt{\bar{\gamma}}$. The satellite peak heights, relative to the sideband, are now:

$$r_\kappa = \left| \frac{Z_\kappa\left(\{\phi_n\},\{\theta_n\}\right)}{Z_0\left(\{\phi_n\},\{\theta_n\}\right)} \right| \mathcal{F}\left(\kappa\omega_{\mathrm{p}}\right). \tag{3.34}$$

In general, the wakefield is difficult to recover from measured r_κ, as every combination of $(k_1, k_2, \ldots, k_N) \in S_\kappa$ must be accounted for. This set, which we have limited to the points of an N dimensional grid of side length $2K+1$ which also lie on the $(N-1)$ dimensional plane $\kappa = \mathbf{v}\cdot\mathbf{k}^\dagger$, contains on the order of $|S_\kappa| \sim (2K+1)^{(N-1)}$ members and the contribution of each must be calculated. For example, limiting the set to $N = 5, K = 5$ gives $|S_2| = 2583$ solutions to $\kappa = 2$.

Once the sets S_0 and S_κ have been found for each of N satellite peaks, obtaining the wakefield profile from the measured peak heights would involve measuring the amplitude and phase of N satellite peaks, yielding $2N$ non-linear simultaneous equations, and solving for the $2N$ unknowns $\{\phi_n\}$ and $\{\theta_n\}$. In practice, however, it is difficult to measure more than a few satellite peaks since the overlap factor $\mathcal{F}(\kappa\omega_{\mathrm{p}})$ becomes small as κ increases. Given a set of experimental parameters and a particular model of the wake, such as Akhiezer and Noble quasi-linear plasma waves, it is nonetheless possible to tackle this problem numerically by calculating the values $\{\phi_n\}$ and $\{\theta_n\}$ for a range of wakefield amplitudes. Instead of $2N$ unknowns the wakefield is described by a single amplitude β_{m}, and using $\{\phi_n\}$ and $\{\theta_n\}$, a lookup table of ratios $|Z_\kappa/Z_0|$ against wakefield amplitudes β_{m} can be constructed. Comparing measured ratios to these values then yields the wakefield amplitudes. In the following section I demonstrate this approach for simulated data.

[†]If we let N and K increase without limit, this is also a Diophantine equation, studied for the past 1800 years, and the subject of a Hilbert problem only solved in 1970. I would love to cite Diophantus of Alexandria but I'm not at all sure how I would...

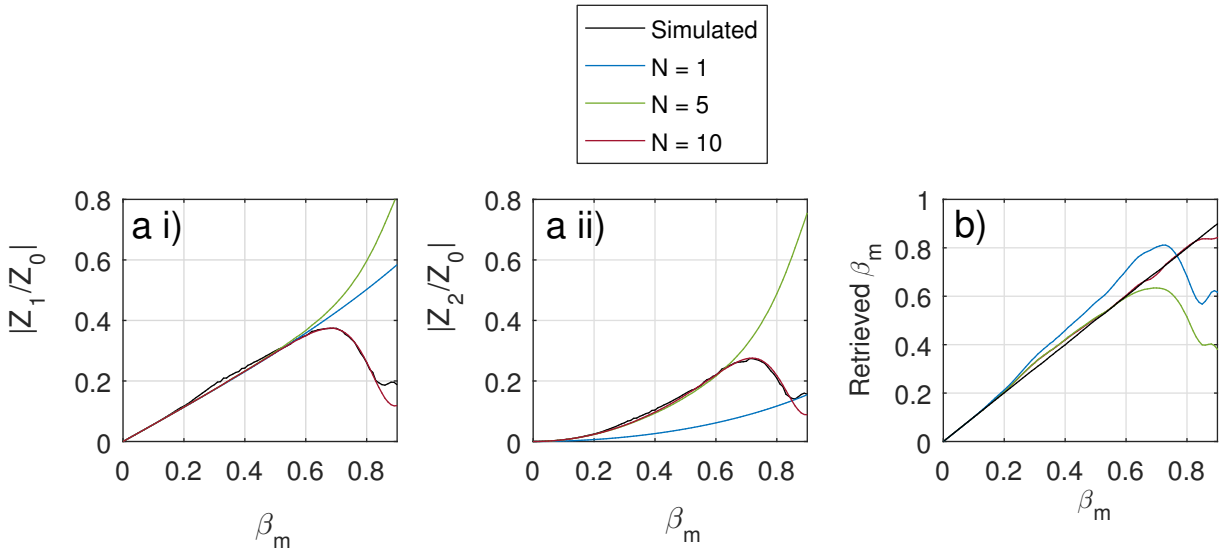


Figure 3.7: a) The calculated ratios $|Z_\kappa/Z_0|$ against plasma wave amplitude β_m for i) the $\kappa = 1$ and ii) $\kappa = 2$ satellite peaks. The ratios are approximated using $N = 1$ (blue), $N = 5$ (green) and $N = 10$ (red) and are compared to simulations of the ideal ratios (black). b) The wakefield amplitude retrieved using the peak ratios $|Z_\kappa/Z_0|$ against the simulated plasma wave amplitude.

Simulated TESS Analysis

In order to demonstrate this procedure, the TESS signals, similar to those shown in Fig. 3.6c, were calculated for simulated plasma waves with maximum electron velocity β_m of up to 0.9, using the same probe frequency, plasma density and interaction length as above. From these TESS signals the peak ratios r_1 and r_2 were extracted and divided by the overlap factors to calculate $|Z_1/Z_0|$ and $|Z_2/Z_0|$. These simulated ratios were compared to expectations from eqn. (3.32), where Z_0, Z_1 and Z_2 were calculated for cold 1D plasma waves as above. In all cases I limited the Bessel expansion to third order, $|k_n| \leq 3$ as $J_4(x)/J_3(x) < 0.13$ for $x < 1$. I considered three possible analyses, with truncation of the harmonic expansion at: N = 1, corresponding to assumption of a linear wakefield; N = 5, for which there are around $|S_\kappa| \approx 400$ combinations; and N = 10, for which $|S_\kappa| \approx 3$ million. The results, shown in Fig. 3.7a, demonstrate that as wakefield amplitude increases more harmonics must be accounted for in eqn. (3.32) in order to correctly calculate the ratios $|Z_1/Z_0|$ and $|Z_2/Z_0|$. Whereas the linear approximation

$N = 1$ fails for $\beta_{\mathrm{m}} > 0.1$, using $N = 10$ the simulations match the calculations very closely up to a wakefield amplitude of around $\beta_{\mathrm{m}} \approx 0.8$.

Next, these ratios were used to calculate lookup tables of $|Z_1/Z_0|$ and $|Z_2/Z_0|$ over different wakefield amplitudes; each value of N gave a different lookup table. The wakefield amplitude of the simulated TESS signals were then estimated by using these lookup tables and the known interaction length and wakefield frequency. For each simulated TESS spectrum I chose the value of β_{m} which minimised the distance (or 2-norm) between the measured ratios and the ratios on the lookup table. For $N = 1$ I ignored Z_2 and followed the same procedure as for linear plasma waves. As shown in Fig. 3.7b, the linear assumption used in Section 3.4 works surprisingly well for non-linear plasma waves, and the accuracy can be increased by further expanding N. The linear TESS procedure, $N = 1$, begins to diverge from the true value at wakefield amplitudes of $\beta_{\mathrm{m}} \approx 0.2$, or $\delta n_e/n_{e0} \approx 25\%$, and subsequently overestimates the wakefield amplitude by around 10% of the true value. Using the $N = 5$ expansion, however, accurately measures wakefield amplitudes up to $\beta_{\mathrm{m}} \approx 0.6$, and the expansion to $N = 10$ only fails at $\beta_{\mathrm{m}} > 0.8$.

There is therefore a trade off between accuracy and computation time. We have demonstrated that an expansion to the 10th harmonic of the plasma frequency allows us to accurately calculate the amplitude of a quasi-linear wakefield up to $\beta_{\mathrm{m}} \approx 0.8$. However, it requires calculating the sum of around 3 million components for each satellite peak at each wakefield amplitude, which takes several minutes on a desktop computer. When truncating the expansion to $N = 5$ harmonics, however, the process requires only around 400 components and takes less than 0.1 seconds. These considerations are important as each lookup table is only valid for a certain experimental set up, with a given plasma density, probe frequency and interaction length. Once the lookup table is calculated, however, retrieval of the wakefield amplitude and frequency is extremely fast.

The analysis can be simplified in two cases. If the wakefield is not too non-linear, $\beta_{\mathrm{m}} \ll 1$, the original frequency component at ω_{p} dominates, such that $\forall n > 1, \phi_n \ll \phi_1$.

In this case, it is possible to truncate the expansion at $N = 1$ and use the same procedure as for linear wakefields, with $r_1 \approx [J_1(\phi_1)/J_0(\phi_1)]\,\mathcal{F}(\omega_\mathrm{p})$. This sacrifices only a small amount of accuracy, on the few percent level. On the other hand, if the phase change is sufficiently small, $\phi(\zeta) \ll 1$, such as at low densities or short interaction lengths, it is feasible to limit the order of the Bessel peaks $|k_n| \leq 1$, as $J_2(x)/J_1(x) < 0.13$ for $x < 0.5$. If we further allow only one k_n to be non-zero, the ratio of heights of the κ-th satellite to the sideband is much easier to calculate:

$$r_\kappa = \frac{J_1(\phi_\kappa)}{J_0(\phi_\kappa)}\mathcal{F}(\kappa\omega_\mathrm{p}) \tag{3.35}$$

This situation can be achieved for a given experiment by reducing the interaction length until $\phi(\zeta) \ll 1$. As with conventional interferometry, phase changes which are too large make reconstruction difficult, but small phase changes are difficult to measure.

3.5.2 General Wakefields

In general, wakefields need not be periodic, and hence contain many frequencies which aren't multiples of the plasma frequency. If we approximate this with a finite set of N frequencies, which are not uniformly spaced, the phase change due to the wakefield can be written as $\phi_\mathrm{wake}(\zeta) = \sum_{n=1}^{N} \phi_n \sin(\omega_n\zeta + \theta_n)$. By comparison with the results for a quasi-linear wakefield with $n\omega_\mathrm{p} \to \omega_n$ we can find the resulting TESS signal:

$$s(t) = \ldots + \sum_{k_1=-\infty}^{\infty} \ldots \sum_{k_N=-\infty}^{\infty} J_{k_1}(\phi_1)\ldots J_{k_N}(\phi_N) e^{\mathrm{i}\sum_{n=1}^{N} k_n\theta_n} H_\mathrm{pr,r}\left(t - \Delta t, \sum_{n=1}^{N} k_n\omega_n\right) + \ldots \tag{3.36}$$

While this looks very similar to the result for quasi-linear wakefields, the peaks for general wakefields will lie either side of the original TESS satellites at locations $t = \Delta t + \psi^{(2)}\left(\sum_{n=1}^{N} k_n\omega_n\right)$. These positions are not equally spaced and therefore each of the TESS satellites will be split into new peaks. However, as one of the

sinusoidal components is at the plasma frequency, $\omega_1 = \omega_\mathrm{p}$, there will still be peaks at $t_k = \Delta t + k\psi^{(2)}\omega_\mathrm{p}$, from the cases $\{k_n\} = \{k, 0, 0, \ldots\}$.

This is demonstrated for example simulated wakefields, in i) a weakly non-linear regime and ii) the bubble regime, shown in Fig. 3.8. In the bubble regime the ponderomotive force from an intense laser pulse is sufficient to completely evacuate a region of electrons, leaving only ions within a bubble with a diameter equal to the plasma period. The density maps were simulated with the EPOCH Particle-In-Cell code [76], with a bi-Gaussian drive laser with a peak beam intensity of i) 2×10^{18} Wcm^{-2} and ii) 4×10^{19} Wcm^{-2}, passing through a high density plasma at $n_e = 10^{19}$ cm^{-3}. The pulse was matched to the plasma wavelength and period with a $w_0 = 3$ μm spot and a duration of $t_\mathrm{FWHM} = 13$ fs. The resulting density profiles are neither sinusoidal nor periodic and hence will contain many frequency components. The resulting TESS signals in both space and time were simulated for identical and Gaussian probe and reference pulses with 400 nm wavelength and 10 fs bandwidth limited duration, each stretched to 1 ps FWHM duration.

Fig. 3.8a shows the density maps in the two cases: in the weakly non-linear case several plasma periods are captured by the simulation, whereas in the bubble regime only two bubbles are shown, separated by τ_p. Because the wakefield is strongly non-linear the electrons are relativistic and the plasma period is longer than expected for the density, $\tau_\mathrm{p} = \sqrt{\bar{\gamma}}\tau_\mathrm{p0}$, giving $\omega_\mathrm{p} = \omega_\mathrm{p0}/\sqrt{\bar{\gamma}}$. Fig. 3.8b shows the TESS signal for each spatial position of the simulation, with many peaks present at different spatial positions. In the bubble regime, the 1st and 2nd order satellites can be seen at locations $t_1 \approx \Delta t + 0.8\,\omega_\mathrm{p0}\psi^{(2)}$ and $t_2 \approx \Delta t + 1.6\,\omega_\mathrm{p0}\psi^{(2)}$, which implies that $\bar{\gamma} \approx 1.6$. From this it is possible to estimate the wakefield amplitude as $\bar{\beta} \approx 0.8$, or $\delta n_e/n_{e0} \approx 4$, which is a substantial underestimate.

Recovering the wakefield amplitude more accurately is difficult as in general many frequencies are required to completely reconstruct the wakefield; the amplitude of the fundamental at a frequency ω_p is much smaller than the total height of the density

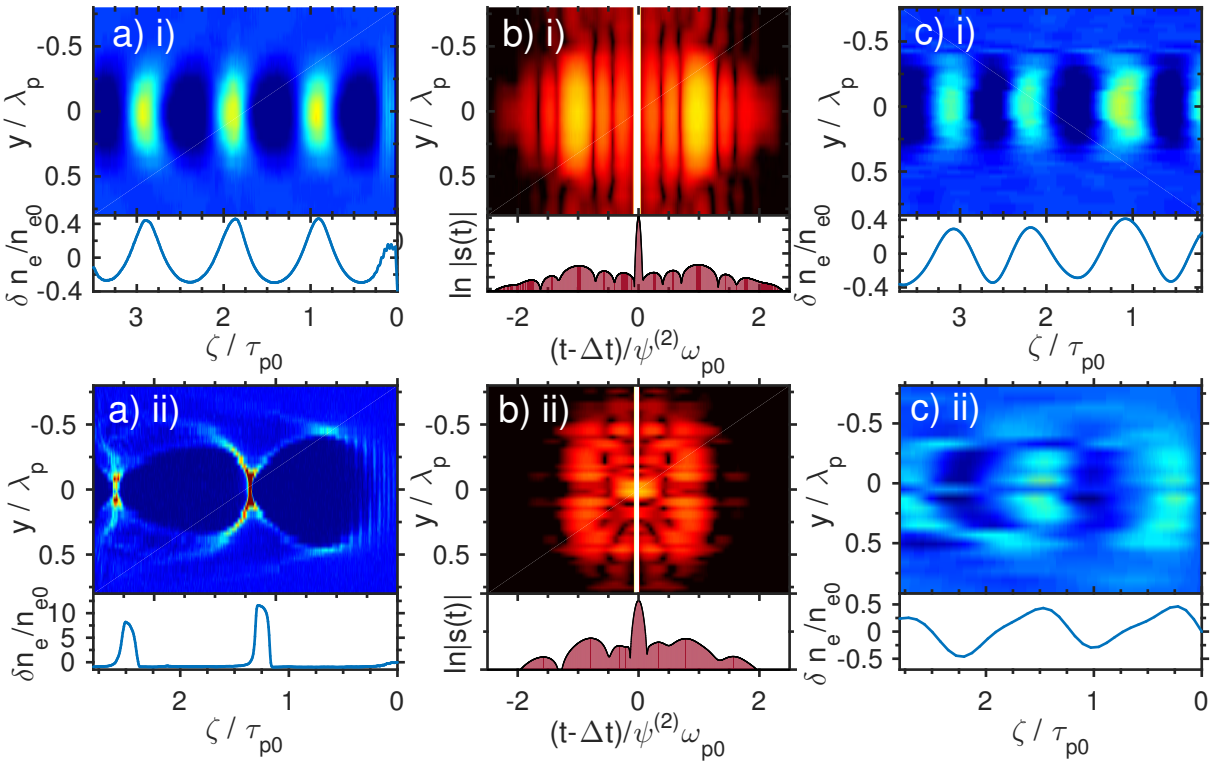


Figure 3.8: a) Simulated electron density maps of a wakefield in i) a weakly non-linear regime and ii) the strongly non-linear bubble regime. b) Maps of the calculated TESS signal resulting from these density profiles, plotted in both space and in time against the expected peak positions. The magnitude of the TESS signal is plotted on a logarithmic scale. c) A reconstruction of the electron density map of the wakefield using the TESS analysis, using $N = 4$ peaks from the TESS signal. In all plots the profile on-axis is shown below.

fluctuation. Unlike periodic quasi-linear wakefields, there is no way to know from the satellite peak heights alone whether the amplitude is reduced because of the shape of the wakefield or because there are only a few plasma periods present. With only two plasma periods present in the strongly non-linear regime, applying the linear analysis of Section 3.4 to the TESS signal in Fig. 3.8b measures the density fluctuation as $\delta n_e/n_{e0} \approx 0.9$, which is a factor of 10 lower than the true value. In the weakly non-linear regime, however, the wake amplitude is estimated as $\delta n_e/n_{e0} \approx 0.4$, close to the true value. In this situation the first TESS peak is sufficient for a first order approximation.

Fig. 3.8c shows attempted reconstructions by instead measuring the location, amplitude and phase of all peaks up to $t = \Delta t + 3\psi^{(2)}\omega_{\mathrm{p0}}$ at each spatial position of the

TESS spectrum. The frequency of each component was deduced as $\omega_n = (t_n - \Delta t)/\psi^{(2)}$, where t_n was the location of the n-th peak; the amplitude was deduced from $r_n = [J_1(\phi_n)/J_0(\phi_n)]\mathcal{F}(\omega_n)$, where $r_n = |s(t_n)/s(\Delta t)|$ is the measured ratio of peak heights; and the phase was deduced from $\theta_n \approx \arg[s(t_n)/s(\Delta t)]$. This allowed us to reconstruct the phase change due to the wakefield as $\phi_{\text{wake}} \approx \sum_{n=1}^{N} \phi_n \sin(\omega_n \zeta + \theta_n)$, for each spatial position, and hence to reconstruct the density map of the wakefield.

While some of the general features of the wakefield are reproduced, only a limited number of peaks ($N = 4$) were captured in the TESS signal and so this reconstruction inevitably fails to reproduce the true profile present in the simulation, particularly small spatial features. Whereas for the weakly non-linear wakefield it is still possible to reconstruct the density with reasonable accuracy, for the simulation in the bubble regime the reconstruction fails and the calculated wakefield amplitude is more than an order of magnitude too small. The accuracy is limited by the number of peaks that can be captured in the TESS signal, which in turn is effectively limited by the bandwidth of the probe pulse and the overlap factor $\mathcal{F}(\omega_n)$. Sharp and non-periodic features with durations less than the bandwidth limited duration of the probe pulse, for instance features such as an ionization front, are poorly reproduced by TESS, and are a source of noise when trying to extract the wakefield amplitude. Capturing shorter duration features requires using a shorter duration probe pulse with a broader bandwidth.

When attempting to reconstruct the density profile of a non-linear wakefield with a continuous range of frequency components, FDH is therefore likely to be more effective than TESS. We can see this by considering another approach to this problem, where we can approximate the phase with N uniformly spaced frequencies separated by an arbitrarily small frequency, δ_ω. The phase can then be written as $\phi_{\text{wake}}(\zeta) = \sum_{n=1}^{N} \phi_n \sin(n\delta_\omega \zeta + \theta_n)$, where many of the amplitudes ϕ_n are small or zero. This is a discrete Fourier transform (DFT), and approximates the true phase profile increasingly well at larger N and smaller δ_ω. All of the information about the wakefield is contained in the region of the TESS signal surrounding the sideband at $t = \Delta t$, which can be

written as:

$$s(t) = \ldots + \sum_{\kappa=-\infty}^{\infty} Z_\kappa\left(\{\phi_n\},\{\theta_n\}\right) H_{\mathrm{pr,r}}\left(t-\Delta t, \kappa\delta_\omega\right) + \ldots, \tag{3.37}$$

where as before $\kappa = \sum_{n=1}^{N} k_n n$ and Z is defined in eqn. (3.32). The potential TESS peaks are then separated by $\psi^{(2)}\delta_\omega$, although only some of these peaks will be non-zero. Fourier transforming this region of the TESS signal to the spectral domain yields:

$$\begin{aligned}\mathcal{S}(\omega) &= \ldots + \sum_{\kappa=-\infty}^{\infty} Z_\kappa\left(\{\phi_n\},\{\theta_n\}\right) \mathcal{E}_{\mathrm{pr}}\left(\omega-\kappa\delta_\omega\right) \mathcal{E}_{\mathrm{r}}^{*}\left(\omega\right) e^{-\mathrm{i}\omega\Delta t} + \ldots \\ &= \ldots + \mathcal{E}'_{\mathrm{pr}}(\omega)\mathcal{E}'^{*}_{\mathrm{r}}(\omega) + \ldots\end{aligned} \tag{3.38}$$

The reference pulse in the spectral domain is unchanged, but the presence of the wakefield has modulated the electric field of the probe, creating many copies of the probe pulse, spectrally shifted by multiples of δ_ω and with amplitudes and phases described by Z_κ. δ_ω can therefore be considered as the resolution of the DFT in the spectral domain. However, as the DFT becomes continuous with $\delta_\omega \to 0$, these copies, and the resulting satellite peaks in the TESS signal, are no longer distinct. In the continuum limit we cannot measure relative peak heights and must instead continue full FDH reconstruction of the phase due to the wakefield using our now isolated spectral component $\mathcal{E}'_{\mathrm{pr}}(\omega)\mathcal{E}'^{*}_{\mathrm{r}}(\omega)$.

Longitudinal Variation

We have just discussed a situation where the wakefield is not periodic in the co-moving frame. The same phase change can arise, however, if the background plasma density varies longitudinally and the plasma frequency changes along the path of the probe pulse. By the nature of the experimental set up, both FDH and TESS average the density profile along this path, with each point ζ in the co-moving frame corresponding to a line of points in space described by $z = v_g(\tau - \zeta)$. Whereas an FDH reconstruction risks

obscuring the signal from one region by overlaying it with the signal from another region, TESS separates the density profile into its different frequency components. If the probe pulse encounters distinct regions of varying electron density along its path, such as in two stage injection-acceleration set-ups described in refs. [92]–[94], the spectral interferogram will again contain components from several frequencies described by $\omega_n = \omega_{\mathrm{p}}(z_n)$.

As described above, this will have the effect of creating new peaks in the TESS spectrum, where each corresponds to the plasma frequency at a particular region along the path of the probe pulse. However, if we are aware of the longitudinal variation of the plasma density and have measured the length of each region, we can again look only at uniformly spaced peaks from the sets $\{(k_1, 0, 0\ldots)\}$, $\{(0, k_2, 0, 0\ldots)\}$, $\{(0, 0, k_3, 0\ldots)\}$ etc. These relate to distinct regions along the path of the probe pulse, each with a distinct plasma density. So long as these do not overlap, if the density ramps between different regions are sufficiently short, it is possible to reconstruct the wakefield frequency and amplitude within each region as before, using:

$$t_{x,k_x} = \Delta t + k_x \psi^{(2)} \omega_{\mathrm{p},x}, \tag{3.39}$$

$$r_{x,k_x} = \frac{J_{k_x}(\phi_x)}{J_0(\phi_x)} \mathcal{F}(k_x \omega_{\mathrm{p}0}), \tag{3.40}$$

$$\text{and } \phi_x = \frac{\omega_{\mathrm{p},x} L_x}{2\omega_0 c} \frac{\delta n_e}{n_{e,x}}, \tag{3.41}$$

where L_x is the length of a region with plasma density $n_{e,x}$ and plasma frequency $\omega_{\mathrm{p}x}$. This causes a phase change of amplitude ϕ_x and peaks in the TESS spectrum at locations t_{x,k_x}.

3.6 Conclusions

In this chapter I have described the TESS technique and shared some (hopefully) useful information for applying it in a real experiment. While some other techniques for measuring wakefields look for qualitative changes, I have shown the highly quantitative

nature of TESS, deriving the TESS signal analytically for identical Gaussian probe and reference pulses co-propagating with a linear sinusoidal wakefield, and shown how to extract the wakefield amplitude and phase from the measured TESS signal. This involves making measurements of several properties of the experimental set-up, such as the phase profiles of the probe and reference pulses and the interaction length. While the previous literature on measuring these is somewhat scanty, I have hopefully described how we approached the problem.

I have described the extension of the TESS analysis technique to probe and reference pulses of arbitrary temporal and spectral profile, which need not be identical. This allows more accurate measurement of the frequency and amplitude of the wakefield in real situations by using the measured spectra of the probe and reference pulses instead of a Gaussian approximation. In turn this allows us to calculate the electron density of the plasma and to infer information about the laser pulse driving the wakefield. In calculating wakefield amplitudes using TESS, the generalized spectral overlap factor given in eqn. (3.26) can be calculated straightforwardly from the measured spectra of the probe and reference pulses. Using recent experimental results, we have seen that the assumption of Gaussian probe and reference pulse spectra can lead to errors in the deduced amplitude of the plasma wakefield by a factor of around two. We have also demonstrated that when the probe and reference have different GDDs the TESS peak separation is described by the effective GDD, which is the mean of the probe and reference GDDs weighted by the square of their bandwidths.

We have also explored applying TESS to measurements of non-linear plasma waves by decomposing the wakefield into harmonics of the plasma frequency. Simulations showed that the extension of TESS to quasi-linear plasma waves allowed accurate reconstruction of cold plasma waves with amplitudes as high as $\beta_{\mathrm{m}} \approx 0.8$. For high wakefield amplitudes finding the peak height ratios involves solving a Diophantine equation and summing over millions of contributions, but at lower amplitudes or smaller phase shifts the wakefield amplitude can be extracted rapidly. TESS therefore retains its advantages over FDH

for wakefields in the quasi-linear regime. It is also possible to follow these principles to reconstruct a full spatio-temporal map of the density for weakly non-linear wakefields. For strongly non-linear wakefields, however, this reconstruction was poor and it was only possible to measure the wakefield frequency and not the amplitude. To reconstruct non-linear wakefields an FDH phase reconstruction is therefore required. On the other hand, for wakefields in density profiles with distinct regions of different densities, TESS has the capability to extract the wakefield amplitude in each region separately, but only if the length of these regions are known.

The extension of TESS to quasi-linear plasma waves is particularly relevant for measuring strong wakefields generated at low plasma densities. As laser wakefield experiments attempt to increase the electron energy gain through increasing the interaction length, TESS provides a means of rapidly diagnosing problems with the wakefield on-shot, without requiring electron injection. We have previously demonstrated [71] that TESS can work effectively at densities of $n_e \sim 10^{18}$ cm^{-3} and below, accurately measuring wakefields with a relative amplitude as small as 1%. By demonstrating that TESS can also be effective for plasma waves with density fluctuations on the scale of $\delta n_e \sim n_{e0}$, it is possible to envisage applying it to accurately measure large amplitude wakefields at electron densities below $n_e \sim 10^{17}$ cm^{-3}.

Chapter 4

A Resonant Excitation Experiment

4.1 Introduction

We have seen in the previous chapters results from simulations addressing two of the concerns of laser wakefield acceleration: driving plasma waves resonantly using trains of multiple pulses, and measuring these plasma waves sensitively using the Temporally Encoded Spectral Shifting (TESS) technique. This chapter goes on to describe experimental results where we successfully demonstrated both the production and measurement of plasma waves at the Astra Gemini (TA2) laser facility at the Rutherford Appleton Laboratory in 2015; this work is published in ref. [71]. I will not, however, go into very great detail; full details can be found in Christopher Thornton's DPhil thesis [62].

The rationale for experimental research is fairly self-evident – we intend both the resonant excitation technique and TESS measurements to be useful and practical tools for laser wakefield acceleration studies and facilities and hence must be willing to demonstrate both of these tools in practice. In the course of the experiment we encountered various challenges, and understanding how we overcame these will hopefully be useful for any future researchers attempting to use the same techniques. While this was by no means the first experiment to resonantly excite a wake, with work over many years on laser beatwave acceleration (e.g. [81]), during this experiment we exhibited a greater degree of control over the pulse train and measured the wakefield amplitude accurately, and hence showed clearly the effect of resonance.

My primary contribution was analysing the data using both the TESS technique and Frequency Domain Holography (FDH), and so I shall focus on how I recovered the wakefield amplitudes in practice, before describing the results and comparing them to expectations from the theory described in Chapter 2. First, however, I will describe the experimental set up.

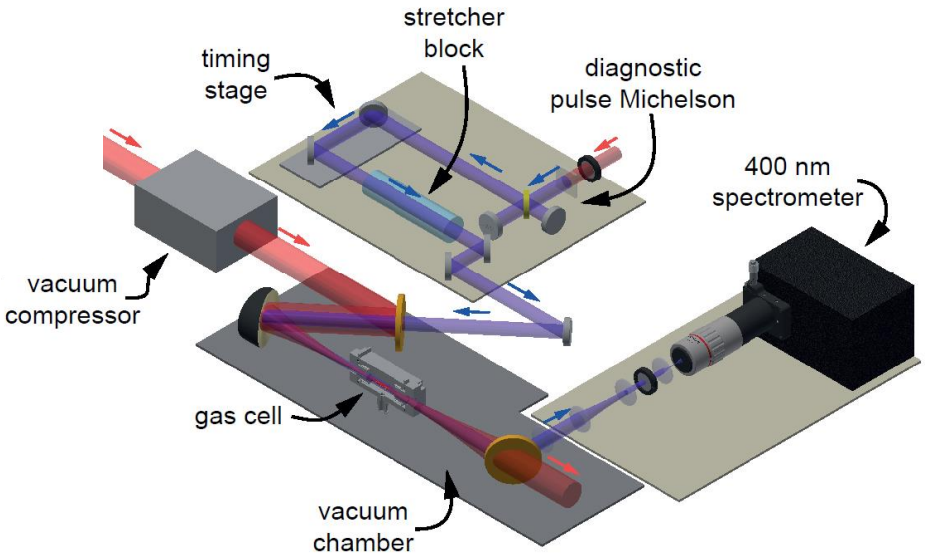


Figure 4.1: Experimental layout of the Astra 2015 experiment, from [71] with thanks to Rob Shalloo. The pump beam shown in red enters on the left, whereas the probe beam, shown in blue after the second harmonic generation crystal, enters from the top right.

4.2 Experimental Set Up

4.2.1 Pre-Interaction Point

The experiment used two co-propagating beams to create and measure the wakefield. A 55 mm diameter 600 mJ 'red' pump beam at 800 nm passed through a Michelson interferometer located before the compressor, which split the beam into two approximately equal energy pulses; by controlling the delay introduced between these two pulses and by changing the compressor settings we could combine two pulses of a controlled duration and chirp, as described in ref. [95]. An acousto-optic programmable dispersive filter (or Dazzler) [96] earlier in the laser chain also allowed us to control third and higher order phase. In this way we could create a pump composed of a) a single compressed pulse, approximately 50 fs in duration b) two compressed pulses separated by a delay $\delta t \approx 400$ fs, and c) a pulse train formed by interference between two stretched and chirped pulses. In case c) the delay introduced in the Michelson acted as a spectral

filter, such that the pulses interfered constructively where $\omega(t)\delta t = 2n\pi$. By controlling the spectral profile $\omega(t)$ with the compressor and Dazzler, the temporal profile of the pulse train could be tuned to create a roughly uniform pulse train approximately 1 ps in duration, with pulses separated by $\delta t' \approx 110$ fs. The temporal profiles of the pulse trains were inferred from a measurement on a single shot autocorrelator (SSA) in combination with a model of the compressor [95].

A separate probe beam of approximately 50 fs duration was passed through a soft aperture in order to reduce the beam diameter to 3.1 mm. This beam was then frequency doubled to 400 nm by passing the beam through a 600 μm thick beta barium borate (BBO) crystal, in order to aid combining and separating the pump and probe beams. This 'blue' beam was then passed through a Michelson interferometer, splitting the beam into two pulses: a probe pulse and a reference pulse, separated by a time $\Delta t \approx 5$ ps. Timing between the pump and probe beams was controlled by reflecting the split probe beam off a retro-reflector, which was mounted on a stage driven by a stepper motor. Finally, both pulses were passed through a 160 mm long block of BK7, which stretched and chirped both pulses to around 1.5 ps duration before they were directed to the vacuum chamber. The group delay dispersion (GDD) of the probe and reference pulses was measured in the manner described in Chapter 3 to be around $\psi^{(2)} \approx 22000$ fs^2, although this changed over the course of the experiment due to differences in pointing.

In the chamber, the probe beam was passed through the back of a dichroic mirror with high reflectance in red and high transmission in blue, and was in this way combined with the pump beam, which was reflected off the front surface. The two beams co-propagated to an $f = 1$ m focal length $f/18$ off-axis paraboloid, which focussed the beams into the centre of a 3 mm long hydrogen gas cell. Focal scans demonstrated that at focus the pump beam had a spot size of $w_0 = (35 \pm 5)$ μm and a Rayleigh range of approximately 5 mm, while the probe beam had a spot size on the scale of 100 μm. The beams entered and exited through 250 μm diameter pinholes in the gas cell.

4.2.2 Post-Interaction Point

After propagating through the interaction region, the probe and reference pulses were separated from the pump pulse by reflectance off a dichroic mirror with high reflectance in the blue and high transmission in the red and passed out of the vacuum chamber through a 25 mm thick fused silica window. Light leaving the gas cell was collimated by an $f = 1$ m focal length lens of 75 mm diameter and reduced in diameter by a factor of three by a Gallilean telescope. Reflection from a third dichroic mirror removed any residual light from the pump beam and a prism polariser removed any blue light created by the pump beam within the plasma. Using an $f = 300$ mm, 25 mm diameter lens, the exit pinhole of the gas cell was imaged in front of a 10× microscope objective (Mitutoyo M plano APO NUV10×) and tube lens and this image was brought onto the entrance slit of a 750 mm focal length Czerny-Turner spectrograph (Acton SpectraPro 2750).

This spectrograph employed a 2400 lines per mm grating, which dispersed the spectra onto a CCD with 2048 × 512 pixels, each of size 13.5 μm × 13.5 μm (Andor Newton DU940N-BU). The spectral dimension of the CCD was calibrated with a Mercury lamp, demonstrating that it captured 12 nm bandwidth from 388.5 nm to 400.9 nm, or around 0.006 nm per pixel, which was sufficient to record the entire spectral range of the frequency doubled beam. The spatial dimension was calibrated with a USAF test chart, which showed that one pixel corresponded to 1.3 μm in the object plane. This was sufficient to resolve features on the scale of the focal spot of the pump beam while capturing the entire spatial extent of the probe beam at focus, which was on the scale of 150 μm.

The first steps in setting the diagnostic up were maximising the visibility of the interference fringes between the probe and reference pulses, and choosing a probe-reference delay Δt such that these fringes were well resolved by the spectrometer. A delay of $\Delta t = 5$ ps gave a fringe period of $\Delta\omega\Delta t = 2\pi$ or $\Delta\lambda \approx \lambda_0^2/(c\Delta t) \approx 0.4$ nm, well resolved by the spectrometer. Next, the delay between the pump and probe beams was

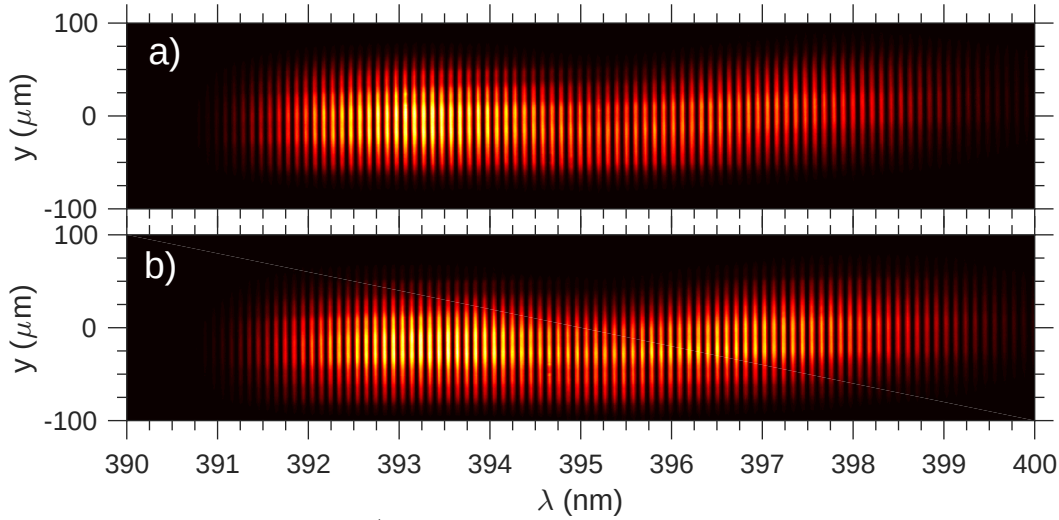


Figure 4.2: Typical spectra from a) a shot with a strong wakefield signal, compared to b) a shot taken without gas.

adjusted such that the pump beam arrived on target between the probe and reference pulses. This was achieved by taking laser shots with a filled gas cell and increasing the probe-pump delay until an ionization front signature was seen on the recorded spectrum of the combined probe and reference pulses. At this stage position the probe and pump pulses were overlapped at focus in both time and space. Increasing the delay further allowed a second ionization front signature to be found, at which stage position there was temporal and spatial overlap between the pump and the reference pulses. Between these two stage positions the pump arrived on target between the probe and reference pulses. Scans taken of the ionization front signatures changing with pump-probe beam delay allowed the spectral phases of the probe and reference to be determined up to a constant, as described in Chapter 3.

4.3 Modified FDH Analysis

4.3.1 Extracting Spectral Phase

Several modifications to the standard FDH analysis described in [45] were made in analysing the data. Rather than taking background measurements with the probe and reference pulses separately, we took background shots with both probe and reference pulses, but with no gas in the target. Fig. 4.2 compares the spectra resulting from a signal shot with a wakefield present and a background shot taken with no gas present. Note that the difference between the two raw spectra is almost imperceptible and the wakefield signal can only be revealed after further analysis.

The first step of the analysis was regridding the measured spectra to a grid which was uniform in frequency, rather than in wavelength, as described in Chapter 3. During this step the measured intensity was linearly interpolated onto the new grid points. However, the spectrometer CCD had a non-zero background; the level of this background was estimated using the mean of 10 pixels in the bottom left of the CCD, far from the spectrum, and this value was subtracted from the entire CCD before performing the interpolation. In constructing the new grid, I chose to increase the number of pixels across from 2048×512 to 8192×512, with the surrounding area filled with zeros. This increased the frequency range from 0.150 rad/fs to 0.602 rad/fs and hence increased the resolution in the temporal domain from 41.9 fs to 10.4 fs.

As described in Chapter 3 and by Matlis et al. [31], [45], the interference term in which we are interested is distinct from the other terms in the temporal domain, where it occurs in the vicinity of $t = \Delta t$. Using the procedure described by Takeda et al. [89], the interferogram was Fourier transformed to the temporal domain and the interference term was isolated with a top-hat bandpass filter between times of 3.5 ps $< t <$ 8.5 ps. The interference term was then inverse Fourier transformed back to the spectral domain. Fig. 4.3 shows the isolation of this interference term in a shot with a wakefield present.

The spectral interferogram can be described by:

$$\mathcal{S}(\omega) = \left|\mathcal{E}'_{\mathrm{pr}}(\omega)\right|^2 + \left|\mathcal{E}'_{\mathrm{r}}(\omega)\right|^2 + \mathcal{E}'_{\mathrm{pr}}(\omega)\mathcal{E}'^{*}_{\mathrm{r}}(\omega) + c.c., \tag{4.1}$$

where c.c. denotes the complex conjugate of the preceding term. The sideband, once isolated, is therefore written as:

$$\mathcal{S}_{\mathrm{sb}}(\omega) = \mathcal{E}'_{\mathrm{pr}}(\omega)\mathcal{E}'^{*}_{\mathrm{r}}(\omega), \tag{4.2}$$

with an argument given by:

$$\arg(\mathcal{S}_{\mathrm{sb}})(\omega) = \psi_{\mathrm{pr}}(\omega) + \psi_{\mathrm{wake}}(\omega) - \psi_{\mathrm{r}}(\omega) + \omega\Delta t, \tag{4.3}$$

A background shot taken without gas, however, gives a sideband described by:

$$\mathcal{S}_{\mathrm{sb,bg}}(\omega) = \mathcal{E}_{\mathrm{pr}}(\omega)\mathcal{E}'^{*}_{\mathrm{r}}(\omega), \tag{4.4}$$

with an argument given by:

$$\arg(\mathcal{S}_{\mathrm{sb,bg}})(\omega) = \psi_{\mathrm{pr}}(\omega) - \psi_{\mathrm{r}}(\omega) + \omega\Delta t \tag{4.5}$$

Putting these together:

$$\therefore \psi_{\mathrm{wake}}(\omega) = \arg(\mathcal{S}_{\mathrm{sb}})(\omega) - \arg(\mathcal{S}_{\mathrm{sb,bg}})(\omega) \tag{4.6}$$

Equations (4.3) and (4.5) describe the interference terms resulting from shots with a wakefield signal and shots taken without gas, respectively. In theory, the interference terms are identical apart from the additional phase due to the wakefield. This allows the spectral phase of the wakefield to be obtained easily as the difference in their arguments. In practice, however, there are normally small variations in $\psi_{\mathrm{pr}}(\omega)$ from shot to shot,

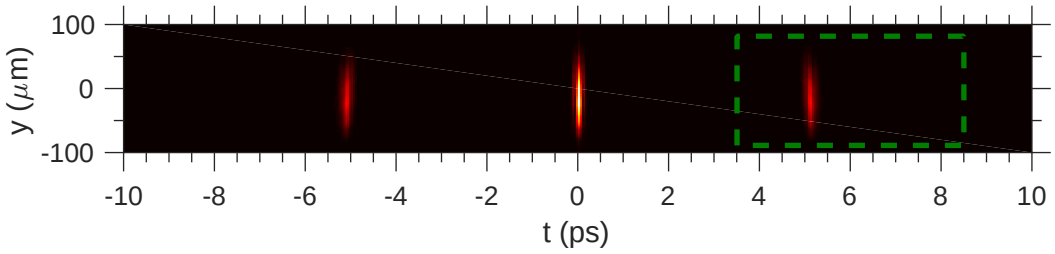


Figure 4.3: The Fourier transform of the measured spectrum against time, showing the isolation of the interference term around $t \approx \Delta t$ (green dashed box).

for instance due to changes in pointing of the probe and reference pulses. I accounted for this by finding the spatial locations of the peak in the measured spectrum for both a signal shot with a wakefield and a reference shot without gas. The reference spectrum was then translated vertically to lie in the same spatial position as the signal before the subtraction was carried out.

4.3.2 Extracting Temporal Phase

Here phase reconstruction for spatial interferometry would stop, but because one cannot typically directly relate the spectral phase of the wakefield to the temporal phase, one must first reconstruct the electric field of the probe in the spectral domain, $\mathcal{E}'_{\mathrm{pr}}(\omega)$. Because the probe and reference pulses were split from the same beam, the measured probe and reference spectra were the same before passing through the gas. I therefore estimated the probe amplitude by dividing the signal spectrum by the square root of the background spectrum. Because we measured the 3rd and higher order terms of the spectral phase of the probe to be negligible (as described in Chapter 3), I estimated the spectral phase $\psi_{\mathrm{pr}}(\omega)$ truncated to 2nd order using the measured GDD $\psi^{(2)}_{\mathrm{pr}}$. This gives an estimate for the probe pulse after passing through the plasma which can be written

as below:

$$\begin{aligned}\mathcal{E}'_{\mathrm{pr}}(\omega) &= \left|\varepsilon'_{\mathrm{pr}}(\omega)\right| e^{\mathrm{i}\psi_{\mathrm{pr}}(\omega)} e^{\mathrm{i}\psi_{\mathrm{wake}}(\omega)} \\ &\approx \left|\frac{\mathcal{S}_{\mathrm{sb}}(\omega)}{\sqrt{\mathcal{S}_{\mathrm{sb,bg}}(\omega)}}\right| e^{\frac{1}{2}\mathrm{i}\psi^{(2)}_{\mathrm{pr}}(\omega-\omega_0)^2} e^{\mathrm{i}\psi_{\mathrm{wake}}(\omega)}. \end{aligned} \tag{4.7}$$

At this stage I increased the number of pixels in the spectral dimension with zero padding again, to 32768×512, and hence increased the temporal resolution to 2.6 fs. The electric field of the probe in the temporal domain is then given by:

$$E'_{\mathrm{pr}}(\zeta) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} \mathcal{E}'_{\mathrm{pr}}(\omega) e^{+\mathrm{i}\omega\zeta} \mathrm{d}\omega. \tag{4.8}$$

The argument of $E'_{\mathrm{pr}}(\zeta)$ contains the temporal phase of the wakefield, written as:

$$\arg(E'_{\mathrm{pr}})(\zeta) = \phi_{\mathrm{pr}}(\zeta) + \phi_{\mathrm{wake}}(\zeta). \tag{4.9}$$

To isolate the temporal phase due to the wakefield, however, one must subtract the intrinsic temporal phase of the probe. This can be calculated from the electric field of the probe in the spectral domain, as above. To estimate $\mathcal{E}_{\mathrm{pr}}(\omega)$, I again took the probe's amplitude as approximately equal to the square root of the measured background spectrum, and the spectral phase of the probe truncated to 2nd order using the measured GDD, giving:

$$\mathcal{E}_{\mathrm{pr}}(\omega) \approx \left|\sqrt{\mathcal{S}_{\mathrm{sb,bg}}(\omega)}\right| e^{\frac{1}{2}\mathrm{i}\psi^{(2)}_{\mathrm{pr}}(\omega-\omega_0)^2} \tag{4.10}$$

The inverse Fourier transform of this yields the electric field of the probe before the interaction:

$$E_{\mathrm{pr}}(\zeta) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} \mathcal{E}_{\mathrm{pr}}(\omega) e^{+\mathrm{i}\omega\zeta} \mathrm{d}\omega, \tag{4.11}$$

with an argument given by:

$$\arg(E_{\mathrm{pr}})(\zeta) = \phi_{\mathrm{pr}}(\zeta) \tag{4.12}$$

This can be subtracted from the argument of $E'_{\mathrm{pr}}(\zeta)$ to leave $\phi_{\mathrm{wake}}(\zeta)$, which is directly proportional to the density structure of the wakefield.

$$\therefore \phi_{\mathrm{wake}}(\zeta) = \arg(E'_{\mathrm{pr}})(\zeta) - \arg(E_{\mathrm{pr}})(\zeta) \tag{4.13}$$

Equations (4.7) and (4.10) describe estimates of the probe pulses in the spectral domain resulting from a shot with a wakefield signal and from a reference shot without gas, respectively. These estimates allow us to reconstruct the probe pulses in the temporal domain in these two cases, and hence to extract the temporal phase due to the wake as the difference in their arguments.

However, when dealing with such small wakefields one must remember to account for the chirped probe pulse by using the local frequency of the probe: $\omega(\zeta) \approx \omega_0 - \frac{\zeta}{\psi^{(2)}}$. The error from using $\omega(\zeta) = \omega_0$ is only a few percent, but can be of comparable scale to the wakefield amplitude $\delta n_e/n_{e0}$. As shown in Chapter 3, the phase due to the wakefield is described by:

$$\begin{aligned}\Delta\phi(\zeta) &= -\frac{\omega_{\mathrm{p0}}^2 L}{2\ \omega_0 c}\left[1 + \frac{\delta n_e}{n_{e0}}\sin(\omega_{\mathrm{p}}\zeta + \psi) - \frac{\delta\omega(\zeta)}{\omega_0}\right] \\ &\approx -\frac{\omega_{\mathrm{p0}}^2 L}{2\ \omega_0 c}\left(1 + \frac{\delta n_e(\zeta)}{n_{e0}} + \frac{\zeta}{\omega_0\psi^{(2)}}\right).\end{aligned} \tag{4.14}$$

We can therefore reconstruct the density profile as:

$$\frac{\delta n_e(\zeta)}{n_{e0}} \approx \frac{2\ \omega_0 c}{\omega_{\mathrm{p}}^2 L}\phi_{\mathrm{wake}}(\zeta) - \frac{\zeta}{\omega_0\psi^{(2)}} - 1 \tag{4.15}$$

Fig. 4.4 compares the temporal phase of the wakefield with the relative electron density deduced from eqn. (4.15). It can be seen that the chirp in the probe pulse caused

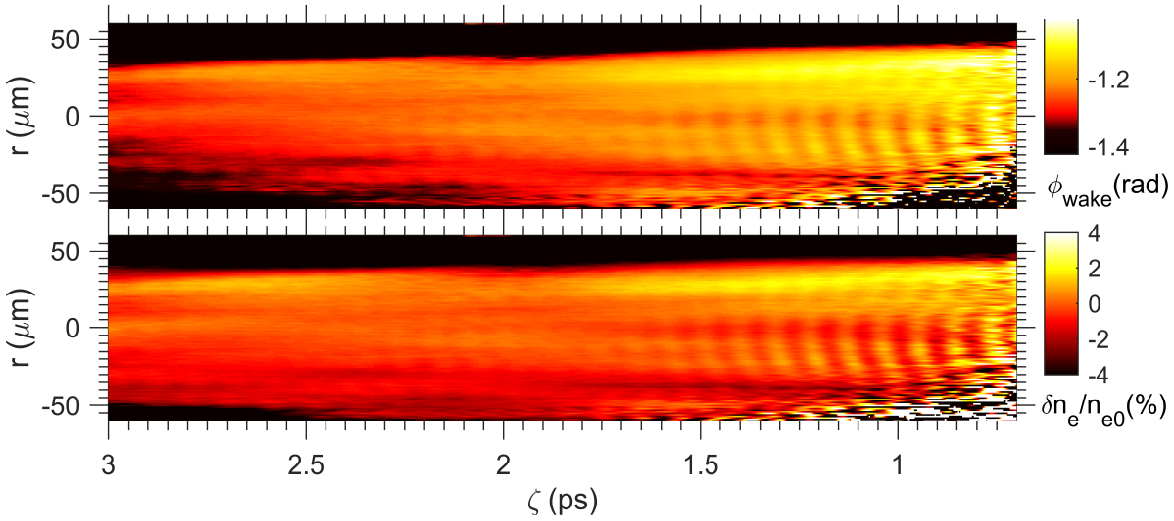


Figure 4.4: a) The temporal phase of the wakefield signal $\phi_{\text{wake}}(\zeta)$ compared to b) the density variation in the wakefield $\delta n_e(\zeta)/n_{e0}$ once the chirp of the probe pulse has been accounted for.

a slowly changing phase gradient in addition to the wakefield; removing it leads to a clearer view of the density oscillation.

Despite cleaning up the signal in this process, this density map still contains errors from a number of sources. Firstly, the background removal in eqn. (4.6) is not perfect, as the phase of the probe and reference pulses change shot-to-shot. This is mainly due to small differences in pointing, which change the material the pulses pass through over their long optical paths, and results in a slowly varying background remaining. The background subtraction will also not remove effects from the neutral gas and the plasma surrounding the wakefield, which aren't present when the gas cell is empty. Secondly, the probe pulse has been reconstructed from an estimate of the probe spectrum and GDD in eqn. (4.7). An incorrect estimate of the probe spectrum again results in a slowly varying change to the final phase map, which is hard to estimate quantitatively; some previous work attempted to recover the wakefield more accurately by subtracting a polynomial fit to the phase map, leaving just the wakefield oscillations. An incorrect measurement of the GDD, on the other hand, results in a mis-scaling of the final phase

map in time; in practice both the GDD and the final time axis had an error on the scale of 5%.

4.4 TESS Analysis

4.4.1 Calculating the TESS signal

Once the preliminary set up was complete, spectral holograms could be taken with the gas cell filled to a range of pressures and wakefield signatures could be extracted using the TESS technique. As we saw in Chapter 3, TESS analysis requires first Fourier transforming the measured spectrum before finding and measuring peaks in the TESS signal. While this process may seem trivial, it is complicated by background noise, particularly for weak wakefields. In analysing the results of this experiment I employed several tactics to identify weak signals amongst background and to remove this background.

The first step was regridding the measured spectrum to a grid which was uniform in frequency rather than in wavelength, as described in Chapter 3 and in more detail above. The signal due to a wakefield was again too small to be seen in the raw spectral hologram, as previously shown in Fig. 4.2. The spectra both show the fringes from interference between the probe and reference pulses. In the TESS signals resulting from these spectra, the TESS satellite peaks can't be seen when plotting using a linear colour scale (as shown in Fig. 4.3), and so Fig. 4.5 instead shows the TESS signals using a logarithmic colour scale. The five strongest peaks are the DC peak, the delay sidebands at $t = \pm\Delta t \approx \pm 5$ ps, and some additional peaks at higher times $t \approx \pm 10$ ps due to back reflections from mirrors. Only after close inspection are the $k = 1$ TESS satellite peaks due to the wakefield evident at around $t \approx \pm 6$ ps. These peaks are present only in the TESS signal with a wakefield present, and not in the shot taken without gas. They are much smaller spatially, around 25 μm high, indicating that they are a feature caused by

the pump pulse, which has a tighter focal spot. The sharp eyed reader might also be able to spot the $k = -1$ satellites at around $t \approx \pm 4$ ps.

These TESS satellites are more obvious after the first stage of background subtraction; Fig. 4.5c shows such a TESS signal after the background has been subtracted. A TESS signal taken from a shot without gas was shifted in the spatial dimension and scaled in amplitude to match a signal from a shot with a wakefield, before being subtracted from the wakefield positive signal. This transformation was made by first finding the spatial lineouts of both the background shot and the wakefield positive shot at the $t = \Delta t$ delay peak and fitting these to Gaussians. The TESS signal from the background shot was then shifted and scaled until the peak of the Gaussian was the same amplitude and location. This accounted for changes in the pointing and intensity of the probe and reference pulses between the two shots. The success of the background subtraction can be seen in the complete removal of the back reflection signals at $t \approx 10$ ps and the removal of most of the noise near the DC peak, allowing the two satellites at $t \approx \pm 1$ ps to be seen.

4.4.2 TESS Satellite Peak Locations

The next step of the TESS analysis was to find the locations of the satellite peaks. The signal was filtered to isolate the $k = 1$ satellite, looking only in the region above the delay peak given by $\Delta t + 0.5$ ps $< t < 8.5$ ps; the largest peak in this range was then found. This filter removed the $t = \Delta t$ delay peak and any remaining back reflections, while avoiding some of the noise present near the $k = -1$ satellite. Brownian noise on the original measured spectrum, for instance, gives rise to a background noise level in the TESS signal proportional to $1/t^2$, and hence the noise level is lower above the $t = \Delta t$ delay peak than below it. Given the measured GDD of $\psi^{(2)} \approx 22000$ fs^2 this range allowed the measurement of plasma frequencies between 24 rad/ps and 160 rad/ps, or electron densities between 1.8×10^{17} cm^{-3} and 8.3×10^{18} cm^{-3}. This range could be

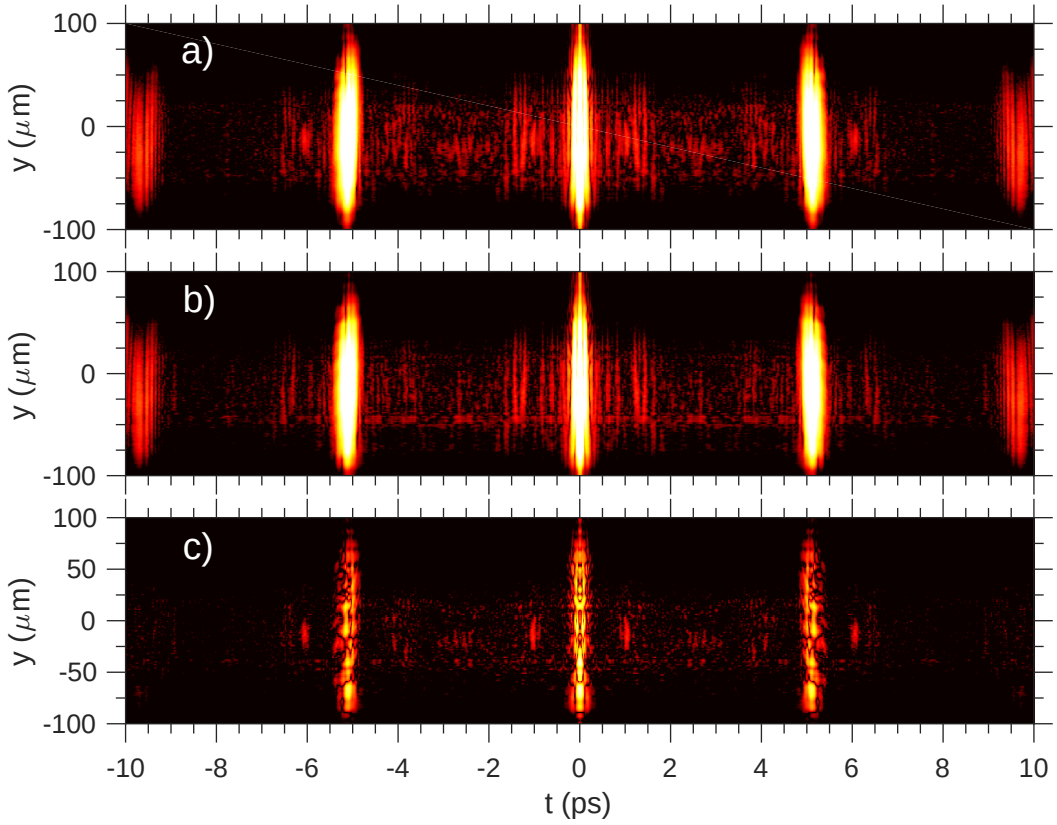


Figure 4.5: A typical TESS signal from a) a shot with a strong wakefield signal, compared to b) a shot taken without gas. c) The TESS signal from (a) after subtracting the background signal shown in (b). The colour scale is logarithmic, with the same scale applied to all three plots.

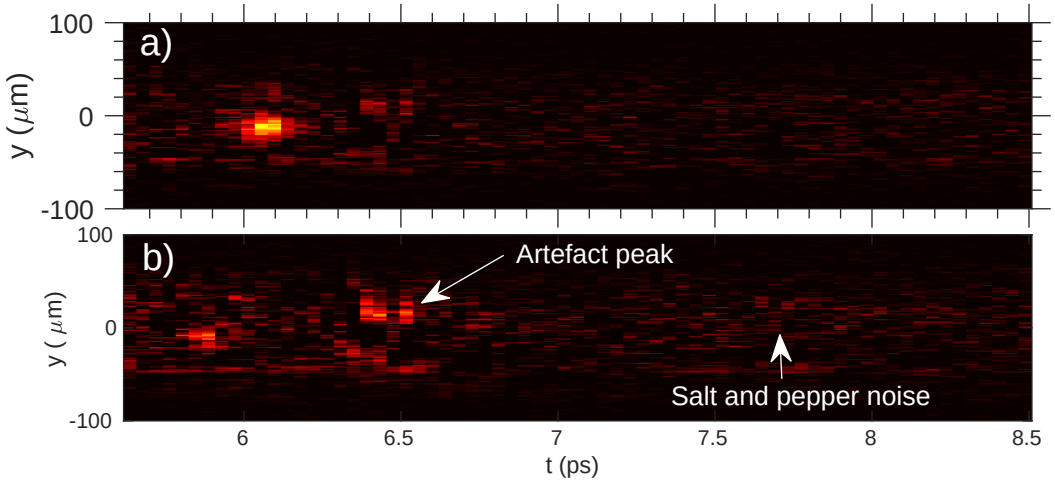


Figure 4.6: TESS satellite peaks from a) a shot with a strong and isolated wakefield signal and b) a shot with a weaker wakefield signal below $t = 6$ ps and a strong artefact peak around $t \approx 6.5$ ps. The colour scale is linear and is the same in both plots.

extended on both the lower and upper limits at the cost of finding more noise peaks.

Not all shots were as clear as that shown in Fig. 4.5a. Fig. 4.6 compares this shot with another taken a short while earlier in the same run, with a slightly lower pressure. In this shot, the background subtraction has not been as successful and there remain two peaks: a peak in the expected location at $t \approx 5.8$ ps and a larger artefact peak at around $t \approx 6.5$ ps. In an attempt to remove peaks from salt and pepper noise (randomly distributed single pixels with high values, such as can be seen across the TESS signal for $|y| \lesssim 50$), only peaks with a spatial width of at least three pixels were accepted, but this filter would not be sufficient to remove the artefact peak shown in Fig. 4.6b.

This results in a number of peaks being falsely identified as TESS satellites. Fig. 4.7a shows the locations of peaks taken over the course of a single run, with these false peaks distinctly away from the trend of the data. This allows artefact peaks to be identified and removed, with peaks more than 0.2 ps away from the trend line neglected in the following analysis. Once the TESS satellite peaks are identified, the plasma period can be calculated from their locations, using $t_1 = \Delta t + \psi^{(2)}_{\text{eff}}\omega_{\text{p}}$ as described in Chapter 3.

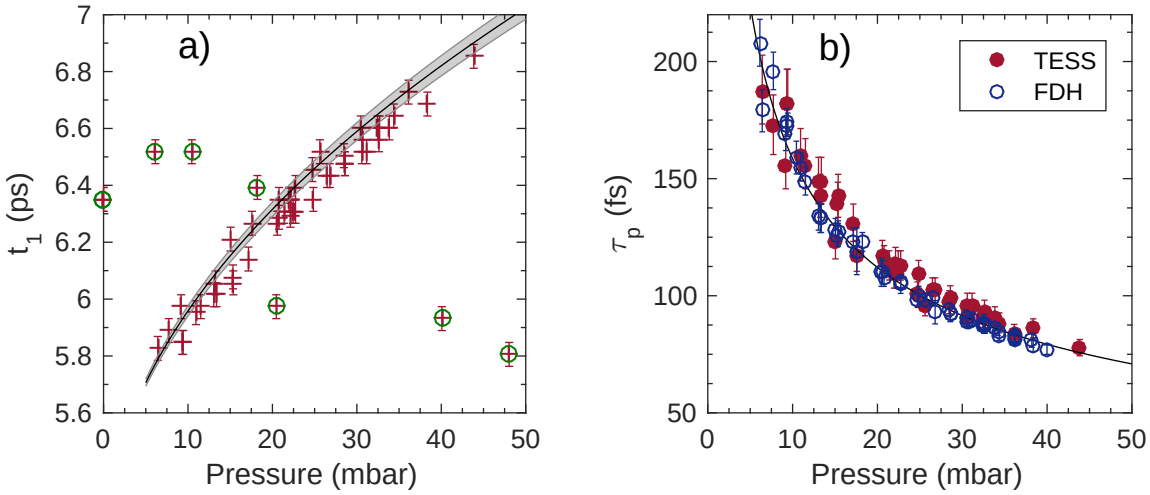


Figure 4.7: a) Measured peak locations in the TESS signal against backing pressure of the gas cell with the trend expected from the measured GDD and probe-reference delay (dashed black line) and suspected artefact peaks marked (green circles). b) Plasma period calculated from the TESS satellite peak locations (solid red circles) and measured from the FDH phase reconstruction (open blue) compared to the plasma period expected (solid black line).

The GDDs of the probe and reference pulses were measured as $\psi^{(2)}_{\mathrm{pr}} = (20900 \pm 300)$ fs^2 and $\psi^{(2)}_{\mathrm{r}} = (22600 \pm 600)$ fs^2 respectively, giving $\psi^{(2)}_{\mathrm{eff}} = (21750 \pm 900)$ fs^2. Fig. 4.7b then shows the plasma periods inferred from the TESS satellite positions compared to those measured from the FDH phase reconstruction and to the expected plasma period in hydrogen. With no free parameters on either measurement, the fit to expectations is very good.

4.4.3 TESS Satellite Peak Heights

Now that these peaks are confirmed as TESS satellites, a measurement of their heights can yield the wakefield amplitude. Before making this measurement, any remaining background must be subtracted with care, as in our measurements the TESS satellite peak was smaller than one hundredth the amplitude of the $t = \Delta t$ delay peak. Three regions of the TESS signal were selected, each 8 pixels by 8 pixels in size: one centred on the satellite; and two regions separated in the spatial dimension by 20 pixels (i.e. above

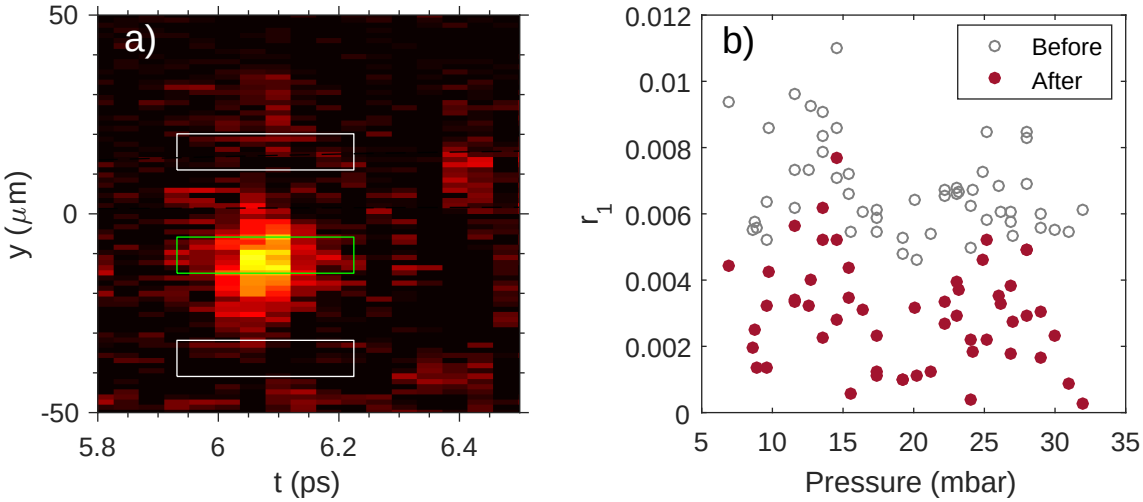


Figure 4.8: a) Regions selected for background subtraction (white boxes) surrounding a signal region (green box). b) Measured TESS satellite peak heights against gas cell backing pressure before and after background subtraction.

and below the satellite). The two regions above and below the satellite were averaged and the resulting average smoothed with a locally weighted linear regression fit. The region surrounding the satellite was smoothed in the same way. The 8×8 grid in the background region was then subtracted from the grid in the signal region and the height of the TESS satellite peak was measured.

Fig. 4.8a shows the three regions selected for a shot with a strong wakefield signal; whereas the central signal region captures the peak, the two surrounding background regions must be sufficiently far away to be unaffected by the wakefield. Dividing the measured peak height by the height of the $t = \Delta t$ sideband at the corresponding spatial location then gives the ratio of peak heights r_1. This is shown in Fig. 4.8b both before and after background subtraction. As we can see, before applying the procedure described above, values of r_1 were measured with heights down to around $r_1 \approx 0.005$ but no lower, implying that this is the level of the noise floor. After applying the full background subtraction the noise floor has effectively been removed, with the smallest measured peak heights now on the scale of $r \approx 0.0005$. These measurements of $r_1 < 0.01$ demonstrate the sensitivity of the TESS measurement.

As described in Chapter 3, in order to calculate the wakefield amplitude from the measured TESS satellites, the overlap factor $\mathcal{F}(\omega_\mathrm{p})$ must be measured, using a cross-correlation of the spectra of the incident probe and reference pulses. As the probe and reference pulses are formed from the same beam split in a Michelson, they have nearly identical spectra before passing through the plasma, albeit of different intensities and pointing, $I_\mathrm{pr}(\omega) \propto I_\mathrm{r}(\omega)$. The cross correlation can therefore be estimated by an auto-correlation. For convenience, we used the background measurement of the spectrum with no gas present in order to estimate the incident spectra, after taking the mean along the spatial direction of the central 40 pixels and passing this lineout through a low-pass filter to remove the fringes from the probe-reference interference. Whereas the TESS analysis primarily uses the interference term $S_\mathrm{sb}(\omega) = \mathcal{E}'_\mathrm{pr}(\omega)\mathcal{E}'^*_\mathrm{r}(\omega)$, to estimate the incident spectra we use the DC terms $S_\mathrm{0,bg}(\omega) = |\mathcal{E}_\mathrm{pr}(\omega)|^2 + |\mathcal{E}_\mathrm{r}(\omega)|^2 \propto I_\mathrm{r}(\omega)$.

This gives an estimate for the spectral overlap factor given by:

$$\begin{aligned}\mathcal{F}(\omega_\mathrm{p}) &= \frac{\int_{-\infty}^{\infty}\sqrt{I_\mathrm{pr}(\omega-\omega_\mathrm{p})I_\mathrm{r}(\omega)}\mathrm{d}\omega}{\int_{-\infty}^{\infty}\sqrt{I_\mathrm{pr}(\omega)I_\mathrm{r}(\omega)}\mathrm{d}\omega} \\ &\approx \frac{\int_{-\infty}^{\infty}\sqrt{S_\mathrm{0,bg}(\omega-\omega_\mathrm{p})S_\mathrm{0,bg}(\omega)}\mathrm{d}\omega}{\int_{-\infty}^{\infty}\sqrt{S_\mathrm{0,bg}(\omega)S_\mathrm{0,bg}(\omega)}\mathrm{d}\omega}\end{aligned} \tag{4.16}$$

Fig. 4.9 shows the effect of this approximation on the overlap factor. The spectra of the probe and reference pulses, shown in Fig. 4.9a, are similar but not identical, with the combined spectrum $S_\mathrm{0,bg}(\omega)$ containing features of both. When the overlap factor $\mathcal{F}(\omega_\mathrm{p})$ is estimated by an auto-correlation of $S_\mathrm{0,bg}(\omega)$ this leads to an error on the order of 10%. This level of error is not ideal, but as we did not possess measurements of the probe and reference pulse spectra taken individually for all of our data runs, it is a level of error we must accept.

Once we have calculated the spectral overlap factor, we can recover the phase and wakefield amplitude as described in Chapter 3. In the limit of small phase shifts $\phi \ll 1$ we can approximate the zeroth and first order Bessel functions as $J_0(\phi_1) \approx 1$ and

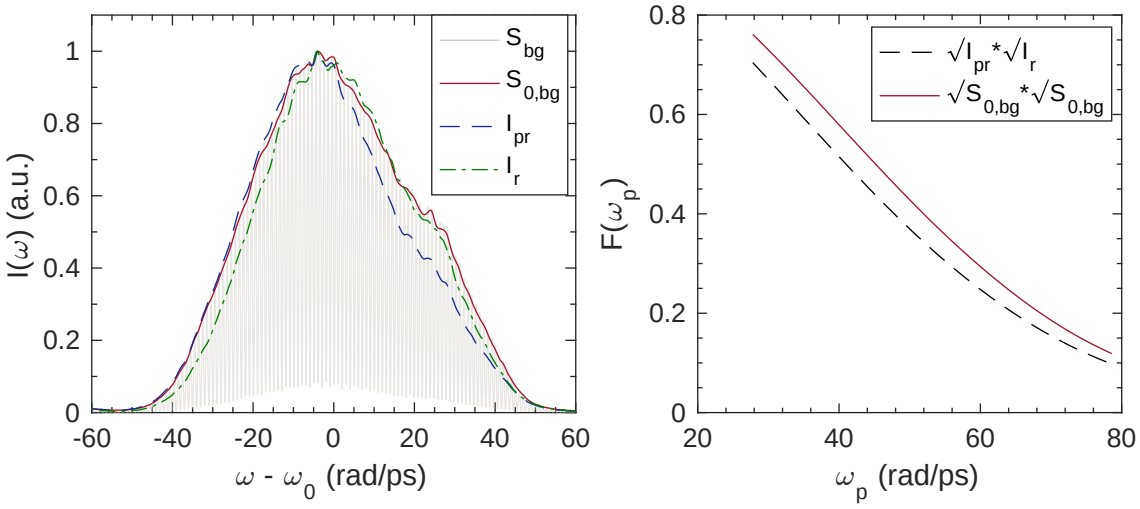


Figure 4.9: a) The measured spectrum in the absence of gas from both probe and reference pulses before applying a low pass filter (solid grey) and afterwards (solid red), compared to the measured spectra of the probe pulse alone (dashed blue) and the reference pulse alone (dash dot green). b) The calculated spectral overlap factor against plasma density using the true cross correlation (dashed black) and an approximation using the auto-correlation of the combined spectra (solid red).

$J_1(\phi_1) \approx \phi_1/2$ respectively. The relative errors in these approximations are on the scale of $\phi_1^2/4$ and $\phi_1^2/8$ respectively, which for $\phi_1 \sim 0.01$ are on the scale of 10^{-4}. Eqn. (4.18) then describes how to recover the wakefield amplitude from the measured relative height of the TESS satellite peak.

The ratio of Bessel functions can be estimated as:

$$\begin{aligned} \frac{1}{\mathcal{F}(\omega_\mathrm{p})} r_1 &= \frac{J_1(\phi_1)}{J_0(\phi_1)} \\ &\approx \frac{\phi_1}{2}. \end{aligned} \tag{4.17}$$

This allows us to extract the wakefield amplitude as:

$$\begin{aligned} \frac{\delta n_e}{n_{e0}} &= -\frac{2\ \omega_0 c}{\omega_\mathrm{p}^2 L}\phi_1 \\ &\approx -\frac{4\ \omega_0 c}{\omega_\mathrm{p}^2 L}\frac{1}{\mathcal{F}(\omega_\mathrm{p})} r_1. \end{aligned} \tag{4.18}$$

This calculation results in errors from a number of sources. First, the plasma density ω_p was calculated from the measured backing pressure of the gas cell. As this measurement was taken some tens of centimeters back from the gas cell itself, this had an accuracy limited to around 1 mbar, resulting in a random error on the scale of 10% at the lowest target densities, or 2% at the highest target densities. Secondly, the interaction length was estimated by measuring the distance between the pinholes directly, giving $L \approx 3.0 \pm 0.1$ mm, with a error of around 3%. However, the gas extended outside the pin-holes, and so the systematic error on this measurement was actually on the scale of 10%. Thirdly, the overlap factor $\mathcal{F}(\omega_p)$ was estimated, as described above, from a low-pass filter of the spectral interferogram, resulting in a systematic error around 15% and a random error from an in-accurate measurement of pressure around 5%. Fourthly, the ratio of peak heights can be mis-measured due to noise. Although we have attempted to control for this using background subtraction, the random error is hard to estimate and due to the low signal remains high, as large as 50% judging by Fig. 4.8. In total the systematic error from the contributions described will be around 20% and the random error around 50%. However, the systematic error is either constant (from gas cell length) or slowly varying (from the overlap factor) and will not affect the measurements of the resonance behaviour of the wakefield. The extremely large random errors, on the other hand, will make it more difficult to determine the resonance behaviour, and necessitate binning the data and fitting it to broad trends.

Finally, there is a physical variation in the wakefield amplitude from shot-to-shot changes in the energy and pointing of the pump pulse. We measured the pump pulse energy on-shot with a molectron and estimated the change in pointing by looking at pointing changes on the 'red' spectrometer centered around 800 nm. Neither of these correlated strongly with the measured fluctuations in peak height in Fig. 4.8 and the scale of the variation in pointing and energy were much lower than 50% and therefore insufficient to explain the scale of the variations in the measured peak height; I therefore concluded that their effect was negligible compared to the other 'hidden' errors and

made no attempt to compensate for variation in the pump laser.

We have therefore described the complete process of recovering a wakefield amplitude from a measured spectral hologram, using a background measurement of a spectral hologram taken without gas present. Now we can see how these wakefield amplitudes change with the pressure of the gas and study the resonances in wakefield amplitude with changing plasma period.

4.5 Results

We measured the wakefield amplitudes against pressure for three different cases: a single pump pulse, a pair of pump pulses separated by a time δt, and a pulse train containing around seven roughly equally separated pulses. We expect the plasma to respond very differently to these temporal profiles depending on the plasma period. There are two ways of calculating the resonances expected: using the work in Chapter 2 to predict the resonance from multiple identical pulses; and calculating the amplitude of the plasma wave driven by the (non-ideal) pulse train measured by the single shot auto-correlator (SSA).

From Chapter 2 we expect the wakefield amplitude following a train of N identical pulses to be described by:

$$\left(\frac{\delta n_e}{n_{e0}}\right)_N = \left(\frac{\delta n_e}{n_{e0}}\right)_{N=1} \left| \frac{\sin\left(\frac{1}{2}N\omega_\mathrm{p}\delta t\right)}{\sin\left(\frac{1}{2}\omega_\mathrm{p}\delta t\right)} \right|, \tag{4.19}$$

where we fix the number of pulses N but leave the separation δt as a fitting factor.

However, a single laser pulse is not expected to produce the same wakefield amplitude at all pressures and has a resonance curve of its own. For a bi-Gaussian pulse with a full width at half maximum duration of τ_0 and a spot size w_0, the wakefield amplitude

in the linear regime can be described by [36], [97]:

$$\left(\frac{\delta n_e}{n_{e0}}\right)_{N=1} = A\omega_{\mathrm{p}}\tau_0 \left[1 + 8\left(\frac{c}{\omega_{\mathrm{p}} w_0}\right)^2\right] \exp\left[-\frac{(\omega_{\mathrm{p}}\tau_0)^2}{16\ln 2}\right], \tag{4.20}$$

where A is proportional to the laser intensity. We have measured $w_0 = (35 \pm 5)$ μm but leave both the duration and the pre-factor A as fitting parameters.

We can also calculate the resonance curve from the inferred temporal profile. In one dimension, the linear response to a laser pulse with a profile $a(\zeta)$ after the laser pulse has passed is described by [35]:

$$\frac{\delta n_e(\zeta)}{n_{e0}} \propto \int_{-\infty}^{+\infty} |a(\zeta')|^2 \cos\left[\omega_{\mathrm{p}}\left(\zeta - \zeta'\right)\right] \, \mathrm{d}\zeta' \tag{4.21}$$

To calculate the wakefield amplitude we can perform a convolution between the inferred temporal profile and a sinusoid for a range of values of ζ. This results in another sinusoid, the amplitude of which is proportional to the wakefield amplitude. However, the SSA itself has an error on the scaling of the temporal axis, leading to an error $\zeta = \alpha\zeta_{\mathrm{meas}}$ for $\alpha \approx 1$. To account for this we allowed the plasma frequency to scale to counteract this effect as $\omega_{\mathrm{p}} \to \omega_{\mathrm{p}}/\alpha$, or $P \to P/\alpha^2$. The final resonance curve then has two fitting parameters: the scaling factors of pressure and of wakefield amplitude.

Fig. 4.10 shows these two resonance curves compared to the measured wakefield amplitudes in the three cases. The scatter in the data arises from both remaining background in the TESS measurement and from physical variation in the wakefield amplitude. However, when binned by pressure the trends become clearer, revealing dramatic differences in the resonance curves. The change in wakefield amplitude with pressure is small for a single pulse, where the primary effect is from matching the finite pulse duration to the plasma period. This resonance in eqn. (4.20) leads to an estimate of the pulse duration of $\tau_0 = (49 \pm 8)$ fs, in excellent agreement with the measurement from the SSA of $\tau_0 = (46 \pm 7)$ fs.

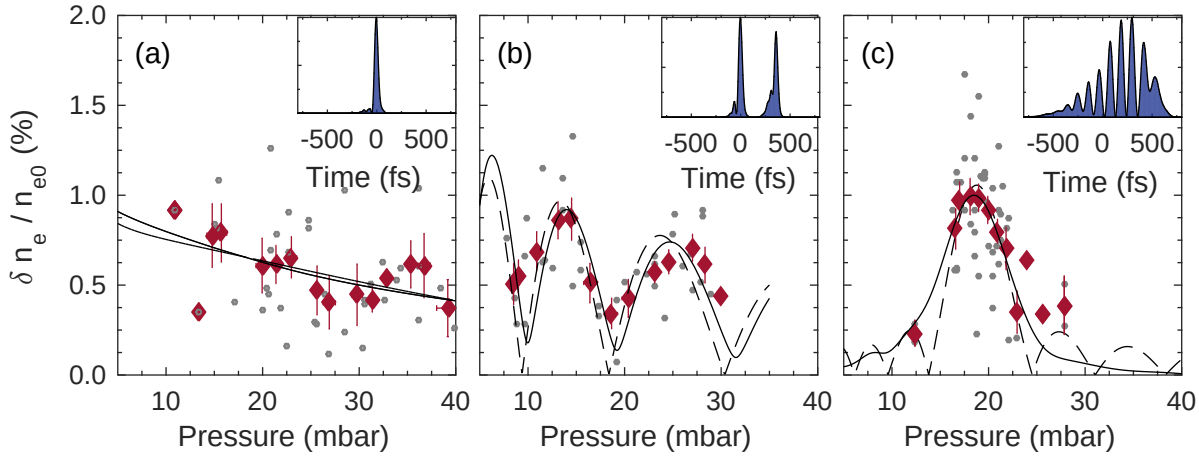


Figure 4.10: Relative wakefield amplitudes against gas cell pressure for a driving pulse train comprising a) a single pulse, b) two pulses separated by $\delta t = (420 \pm 20)$ fs and c) a pulse train with around seven pulses separated by $\delta t = (112 \pm 6)$ fs. Grey circles show the wakefield amplitudes for individual shots and red diamonds show these amplitudes binned by pressure. The insets show the inferred profile of the pump beam. The dashed lines show fits of eqn. 4.19, and the solid lines show the wake amplitudes calculated from the pulse trains shown in the figure insets. The plots are all shown on the same x and y-axis scales.

For a pulse train, however, Fig. 4.10c shows a clear resonance, with wakefield amplitude diminishing rapidly away from the matched pressure at around 20 mbar, at a plasma period of $\tau_{\mathrm{p}} \approx 110$ fs. Indeed, using the fit described in eqn. (4.19), the pulse separation is estimated as $\delta t = (116 \pm 2)$ fs, in good agreement with the SSA measurement of $\delta t = (112 \pm 6)$ fs. In comparison, when using the inferred temporal profile, a fit to the data gave a scaling for the plasma frequency of $\alpha = 1.04 \pm 0.02$, implying that our measurements of the pulse train driver and of the resulting wakefield amplitude are very consistent. The measured non-ideal pulse train gives rise to a smoother resonance curve without nodes, as the non-uniform peak spacing and unequal peak amplitudes do not allow a situation where subsequent laser pulses exactly cancel the wakefield created by earlier pulses. While this should give a better agreement to the measured wakefield amplitudes, there is not enough data here to prove this either way, as the wakefield amplitudes would fall beneath the noise floor.

In the intermediate case of two pulses, the resonance from multiple pulses can be reduced using the double angle formula as:

$$\begin{aligned}\left(\frac{\delta n_e}{n_{e0}}\right)_N &= \left(\frac{\delta n_e}{n_{e0}}\right)_{N=1} \left|\frac{\sin\left(\omega_{\mathrm{p}}\delta t\right)}{\sin\left(\frac{1}{2}\omega_{\mathrm{p}}\delta t\right)}\right| \\ &= 2\left(\frac{\delta n_e}{n_{e0}}\right)_{N=1} \left|\cos\left(\frac{1}{2}\omega_{\mathrm{p}}\delta t\right)\right| \end{aligned} \tag{4.22}$$

The sequence of nodes and anti-nodes resulting from this can be seen in the measured wakefield amplitudes in Fig. 4.10b. Fitting to eqn. (4.19) gives an estimate of $\delta t = (407 \pm 6)$ fs. The SSA, however, measured approximately $\delta t = (365 \pm 40)$ fs. It is possible to gain a third estimate by looking at the measured spectrum of the pump pulse, shown in Fig. 4.11a. As the two pulses are coherent with each other they produce spectral interference, where the separation between fringes is determined by the separation in time. As with the TESS analysis, a Fourier transform of the spectrum, shown in Fig. 4.11b, has a peak at the delay between the pulses. This gives an estimate of

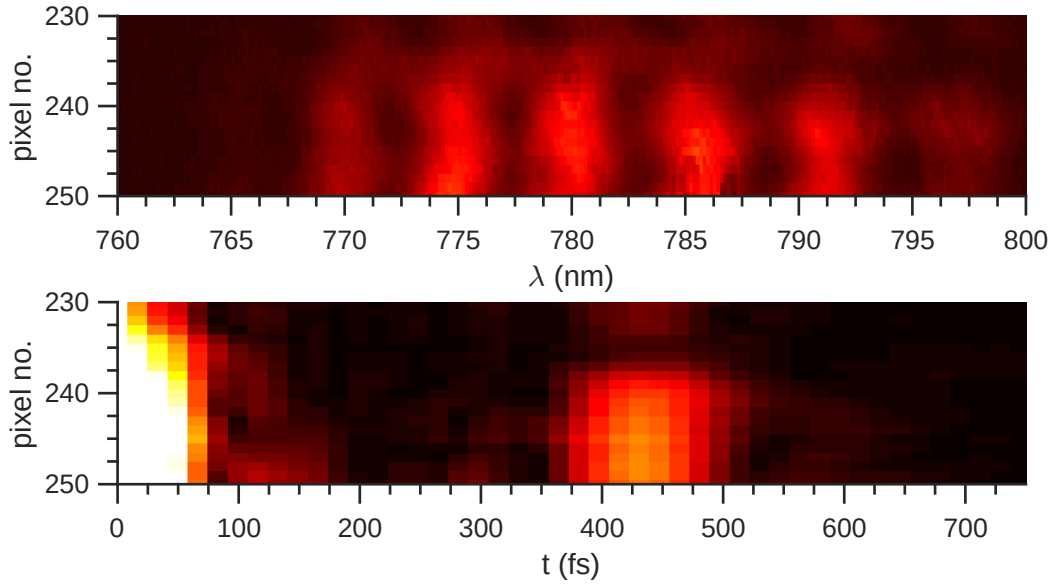


Figure 4.11: a) The measured spectrum of the combination of two pulses, shown against wavelength. b) The Fourier transform of this spectrum shown against time.

the separation as $\delta t = (420 \pm 20)$ fs, and so all of these estimates agree within errors. Similarly, fitting the calculated plasma response from the measured pulse train profile results in scaling the time of the SSA by $\alpha = 1.11 \pm 0.02$, giving a good fit to the data and indicating that the SSA axis was slightly mis-calibrated for this data. For all three cases, of a single pulse, a pair of pulses, and a pulse train, the large variation in the wakefield amplitudes belies clear trends in the resonance behaviour which agree with our theoretical expectations and allow us to accurately extract pulse parameters.

4.6 Discussion

Our two theoretical models both agree well with the data but make different assumptions and have different advantages. Eqn. (4.19) assumes identical pulses in the pump beam, and eqn. (4.20) further assumes that these pulses are bi-Gaussian. This has the advantage of a simple analytic model which is intuitive, following the same behaviour as an idealised diffraction grating, and requires no measurement of the exact temporal profile of the pulse train. There are, however, three fitting parameters: for the pulse separation, the pulse

duration, and the wakefield amplitude created by each pulse. Making the analogy to a diffraction grating, the pulse separation is comparable to the period of the grating and determines the distance between maxima and minima of wakefield amplitude, whereas the pulse duration is comparable to the slit width and hence determines the envelope of the resonance curve. The fourth parameter, the beam waist w_0, was found to affect only the low pressure behaviour of the resonance curve, and so was under-constrained by the data. We therefore fixed it to our measurement of 35 μm.

The second model from eqn. (4.21) is essentially a Fourier transform of the measured intensity profile and requires only two fitting parameters, one again for the wakefield amplitude, the second allowing for a mis-calibration of the SSA. This has the advantage of capturing more information about the resonance, if a measurement can be made of the pump beam's temporal profile. On the other hand, this model is 1D and neglects the spatial size of the pump beam. Our measurement of the temporal profile was indirect, using a model of the compressor as described by Shalloo et al. [95], but nonetheless demonstrated that the pulses in the train were not identical. The assumptions made by this model are therefore probably more realistic than assuming a train of uniformly spaced and identical laser pulses.

In particular, a train of identical pulses results in perfect wakefield extinction at certain plasma frequencies. The plasma responses to the measured temporal profiles lead to a resonance that looks very similar, but does not reach zero wakefield amplitude. At perfect extinction, or perfect visibility of the nodes, successive laser pulses exactly destroy the wakefield created by the preceding pulses, extracting all the energy contained in the plasma wave and becoming blue shifted in response. This 'energy recovery' would be an extremely useful property when moving towards high repetition rate LWFA systems with very high average power inputs: failure to remove energy from the plasma could lead to catastrophic heating and damage. Ideally, however, an electron bunch or bunches may be placed between the constructive and destructive portions of a train of laser pulses, such that immediately after accelerating the electrons any remaining

plasma energy may be efficiently extracted.

We can calculate the extinction ratio from the resonance curve for a pair of pulses, using the maximum and minimum wakefield amplitudes. Where a_1 and a_2 are the strongest and weakest wakefield amplitudes respectively, the maximum amplitude is the sum of the amplitudes resulting from both pulses $A_{\max} = a_1 + a_2$, whereas the minimum amplitude is the difference $A_{\min} = a_1 - a_2$. The energy extinction ratio is the proportion of the energy removed from the wake, or $\epsilon = [a_1^2 - (a_1 - a_2)^2]/a_1^2$. Using $A_{\max} = (0.87 \pm 0.12)\%$ and $A_{\min} = (0.34 \pm 0.08)\%$ we can calculate $\epsilon = 0.68^{+0.15}_{-0.20}$. A significant majority of the energy has therefore been extracted by the second laser pulse.

That a later portion of a pulse train might act to reduce or eliminate the wakefield created by the start of a pulse train was a concern for proponents of the beatwave technique. As described in Chapter 2, the relativistic change in the plasma frequency at high wakefield amplitudes leads to the Rosenbluth-Liu limit on wakefield amplitude, making long and uniformly spaced trains of laser pulses counter-productive for acceleration. In this work, we used the interference between two ps duration and strongly chirped laser pulses to create a pulse train in a manner similar to the beatwave technique. We specifically tuned the spectral profile of the pulses $\omega(t)$ in order to obtain a train of pulses with close to uniform spacing. In future, it would be possible to use 3rd and higher order phase in the Dazzler to chirp the pulse train spacing, exploiting auto-resonance and overcoming some of the limitations of a laser beatwave accelerator.

4.7 Conclusions

Our results are a clear indication that plasma waves can be excited by multiple pulses acting in resonance with the plasma wave. The measured wakefield amplitudes for a) a single pulse, b) a pair of pulses, and c) a train of multiple pulses, agreed well with expected resonance curves using i) a simple model of identical pulses and ii) the plasma response to the inferred temporal profile. We have therefore made a physical

demonstration of the techniques and results discussed in Chapter 2, opening the door to further studies of multiple pulse laser wakefield acceleration. In future, instead of producing a train of pulses 'destructively' from a single laser pulse (that is, by spectrally filtering the pulse and selectively dumping energy) using the powerful Astra laser, we may be able to constructively produce a pulse train by combining a succession of laser pulses from lower peak power laser systems, which could potentially operate with higher efficiency and repetition rates.

We have also experimentally demonstrated the TESS measurement technique described in Chapter 2. This was used to measure extremely weak wakefields, on the scale of just 1% in density variation, over a wide range of pressures. The locations of satellite peaks in the TESS signal changed as expected and in agreement with FDH measurements, with our results the first replication and confirmation of the TESS technique. However, both the more accurate calculation of the spectral overlap factor using the measured non-Gaussian spectrum, as described in Chapter 3, and the effective background subtraction described in this chapter were essential to an accurate measurement of the wakefield amplitudes. This allowed us to accurately extract parameters of the driving laser pulses. This sensitive measurement would allow us in future to tune wakefield accelerators to the desired specifications, without requiring production of electrons or other radiation.

Chapter 5

Low Density Plasma Waveguides from Optical Field Ionization

5.1 Introduction

For acceleration of electrons to high energies in the linear and quasi-linear regimes of laser-plasma experiments the limiting factor is generally the dephasing length, over which a relativistic electron bunch slips relative to the laser driver. As described in Chapter 1 this length is given by $L_{\mathrm{dp}} = \lambda_p^3/2\lambda_0^2$, which is significantly longer than typical Rayleigh ranges, which are on the scale of $Z_R = \pi w_0^2/\lambda_0 \sim \lambda_p^2/\lambda_0$. To employ laser wakefield acceleration effectively it is therefore essential to guide the laser driver over a distance much greater than the Rayleigh range. As achieving higher energy gains requires both longer acceleration lengths and therefore lower plasma densities, the importance of guiding only increases for the next generation of laser wakefield acceleration experiments, such as BELLA 10 GeV [98] and EuPRAXIA [99].

These waveguides should fulfill several desirable properties: they should be long (on the scale of 10s of centimetres), with low density on-axis (around 10^{17} cm^{-3}), and capable of operating at high repetition rates (kHz and above) without suffering damage. For this purpose, many laser-plasma accelerators employ external guiding structures, such as capillary waveguides (e.g. [36], [57]), particularly gas-filled electrical discharge capillaries [58], [59]. These structures offer great advantages, consistently achieving the largest energy gains demonstrated by LWFA systems (e.g. [14], [15]), guiding intense laser pulses over acceleration lengths up to 9 centimetres and producing electrons up to 4 GeV, with on-axis densities in the channel as low as $n_e \approx 2 \times 10^{17}$ cm^{-3} [60]*. Furthermore, some capillary discharges have achieved repetition rates of kHz, although without guiding high intensity laser pulses [100]. However, achieving all of these properties at once remains challenging, and under working conditions with intense laser pulses their physical structures can suffer from damage if the transverse profile or pointing of the laser driver are not ideal, despite commonly using structures

*Recent work by the same group, yet to be published, has demonstrated acceleration to 8 GeV over 27 cm with an on-axis density as low as $n_e \approx 1 \times 10^{17}$ cm^{-3}.

made out of sapphire or even diamond[†].

There is an alternative scheme, creating 'freestanding' plasma waveguides, without an external capillary structure to damage. In this process, an initial laser pulse ionizes a column of plasma along its focus, and a cylindrical shock forms from the hot plasma expanding radially into a cold and un-ionized gas, which leaves a cavity on-axis that can act as a guiding structure. This all-optical system alleviates the requirements on any external physical structure, greatly reducing the risk of damage. Moreover, as plasma recombines on the timescale of microseconds, these channels would be capable of operating at multi-kHz repetition rates. The repetition rate and length of the waveguide are now limited by the repetition rate of the laser and the length of the focus. An axicon, or conical lens, produces a line focus, which is not limited by the Rayleigh range and can be much longer; at each point along the focus, the intensity peak on axis is caused by constructive interference between wavefronts arriving from a different annulus of the axicon. For wavefronts from an axicon of radius R approaching the focus at an angle α to the axis, the length of the axicon focus is $\zeta = R \cot \alpha$. In this way, plasma columns and channels on the scale of tens of centimetres and above can be produced by using shallow approach angles and large diameter axicons.

This work has been developed for many years, as described in refs. [54], [63]–[66], [101]. Until recently, however, these experiments relied upon using picosecond duration laser pulses to heat the plasma collisionally. The collision rate is proportional to the plasma density and diminishes rapidly at low densities, making it impossible in practice to use this technique to create plasma channels with on-axis densities around 10^{18} cm^{-3} or below, where the time between collisions becomes longer than the duration of the laser pulse. Previous work has used clustered gases to maintain high local densities and collision rates while decreasing the average on-axis density to $n_e \approx 10^{18}$ cm^{-3} [66], but lower densities prove difficult. Some other experiments have been conducted using instead a femtosecond laser pulse focussed by a spherical lens, which ionizes the gas

[†]The DPhil thesis of Christopher Thornton [62] discusses using diamond structures in more detail.

and creates a hot plasma using Optical Field Ionization (OFI) [68], [69]. This work was conducted at densities around 10^{18} cm^{-3} and produced channels limited by the Rayleigh range to lengths of a few millimetres. However, the OFI process is independent of the electron density and could potentially be used to create much lower density channels by hydrodynamic expansion, while an axicon lens could extend the length of the plasma waeguides by orders of magnitude.

In this chapter, I will describe combining Optical Field Ionization with an axicon focus to create freestanding plasma channels which are both long and low density. As an all optical system this would be capable of operating at high repetition rate without damage. In order to outline this scheme I will describe the physics of OFI and of the subsequent hydrodynamic expansion, and the advantages and limitations an axicon focus imposes upon the resulting plasma channels. The work is described elsewhere, in a recent paper [72] and in Rob Shalloo's thesis [72], and this chapter contains much of the same work. Whereas I discuss the theory and simulations towards this goal, Rob Shalloo focusses primarily on the experimental design and results demonstrating this scheme.

5.2 Electron Energies from OFI Heating

To understand the expansion of plasma channels formed by OFI driven shocks, we must first find the electron energy distribution that results from OFI heating. This has two contributions: the rate at which electrons are liberated from the atoms, and the energy which the now free electrons accumulate in the electric field of the laser. The first contribution is a complex area of study, with several different regimes of ionization depending on the laser intensity and carrier frequency [102]. The regime of ionization is described by the Keldysh parameter $\gamma = \sqrt{\frac{\varepsilon_I}{2U_p}}$, where ε_I is the ionization energy of the electron, and U_p is the ponderomotive energy of the laser. If $\gamma \ll 1$ the electric field of the laser is comparable to the electric field binding an electron to the atom, and

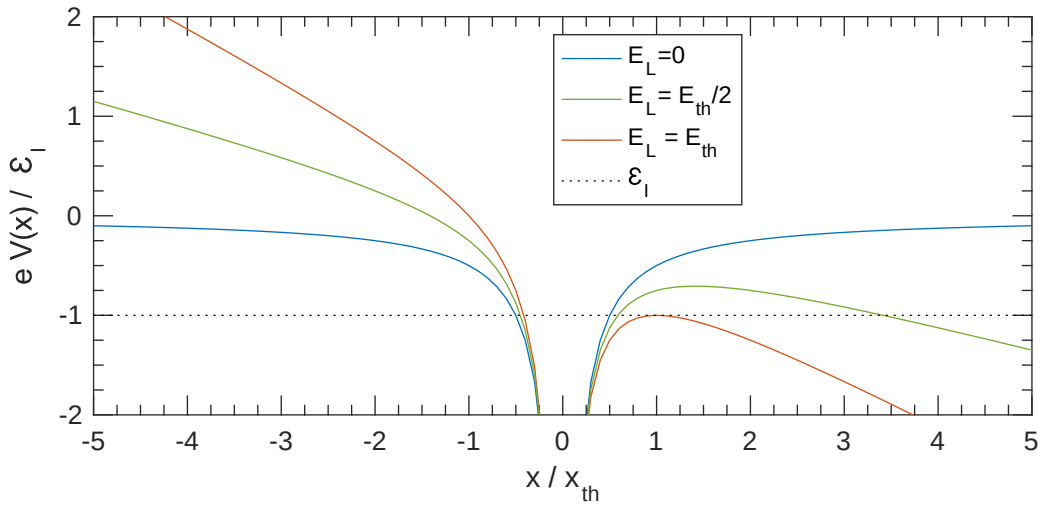


Figure 5.1: Perturbed atomic potentials, $V(x) = (Ze/4\pi\epsilon_0|x|) - E_L x$, along the axis of the laser polarisation for: no applied electric field (blue), an applied electric field half the amplitude of the BSI threshold field (green), and a field at the BSI threshold (red).

the potential barrier can be greatly reduced or suppressed, as shown in Fig. 5.1. The time taken to tunnel through the barrier becomes small compared to the oscillation period of the laser. When $\gamma \gg 1$, however, there is insufficient time for an electron to tunnel through the potential barrier, and instead it becomes ionized by the absorption of incident photons, gaining enough energy to raise it from the ionization state to the continuum. As an individual optical wavelength photon has insufficient energy to overcome the ionization barrier (an 800 nm photon has an energy of only 2 eV), multiple photons must be absorbed at once, in Multi-Photon Ionization (MPI).

As an example consider an ionization energy $\varepsilon_I = 13.6$ eV for ionizing atomic hydrogen to H^+ and a ponderomotive potential $U_p = \frac{e^2}{2m_e c\epsilon_0}\frac{I_L}{\omega^2} \sim 100$ eV for a laser intensity of $I_L \sim 10^{15}$ Wcm^{-2} and wavelength of $\lambda = 800$ nm. This results in $\gamma \sim 0.1$ and allows us to neglect MPI, and focus instead on the barrier suppression ionization (BSI) [103] and Ammosov, Delaine and Krainov (ADK) tunnelling models [104]. This is generally true when ionizing the target gas at optical wavelengths and longer.

5.2.1 Ionization Models

Barrier Suppression Ionization (BSI)

The simplest model we can use for strong field ionization is barrier suppression ionization [105]. Here we describe when the electric field of the laser pulse is sufficiently strong to overcome the Coulomb potential of the atom, at which all electrons in a given band would be liberated at once. The combined potential a free electron experiences at a distance x from the atom is:

$$V(x) = -\frac{Ze}{4\pi\epsilon_0|x|} - E_L x, \tag{5.1}$$

where E_L is the laser electric field strength. The potential has a local maximum with a value of

$$V_{\max} = -2\sqrt{\frac{Ze}{4\pi\epsilon_0}E_L}. \tag{5.2}$$

If this peak is less than the bound potential of the electron, which is ε_I/e below the continuum state, there is no barrier to the electron escaping, giving a threshold electric field and laser intensity:

$$eV_{\max} < -\varepsilon_I \tag{5.3}$$

$$E_L > E_{\text{th,BSI}} = \frac{2\pi\epsilon_0}{Ze^2}\varepsilon_I^2, \tag{5.4}$$

For linear polarization this occurs at an intensity of:

$$I_L > I_{\text{th,BSI}} = \frac{\pi^2\epsilon_0^3 c}{2Z^2e^6}\varepsilon_I^4. \tag{5.5}$$

For a hydrogen atom with an ionization energy of $\varepsilon_I = 13.6\,\text{eV}$, the threshold field is approximately $E_L = 32\,\text{GVm}^{-1}$, requiring a minimum laser intensity of $I_L \approx 1.4 \times 10^{14}\,\text{Wcm}^{-2}$

for linear polarization.

In this simple one dimensional BSI model all the electrons are liberated at once, which does not account for the spatial distribution of electrons around the nucleus when in the field of the laser. In practice, the probability of a electron being in the right place to escape the atomic potential is substantially less than one, and only a subset of electrons escape. This can be modelled as described in ref. [106] by looking at the disruption to the shape of the atomic potential and finding the proportion of the volume where electrons can escape. We shall, however, use the naive BSI estimate as a guide to the order of magnitude of the threshold intensity for ionization and for the resulting energy of the electrons liberated.

Tunnelling Ionization (ADK model)

Clearly the BSI model is only a first estimate and electrons can begin to tunnel through the barrier before it is completely suppressed, implying a finite rate of ionization. To predict this more accurately we can use the Ammosov, Delone and Krainov (ADK) model [104], which describes tunnelling by the transition between the bound electron states and the free electron Volkov states in the quasi-static limit. This predicts an ionization rate with an exponential dependence on the field described by:

$$w(t) \propto \exp\left[-\frac{2\left(2\epsilon_I\right)^{1.5}}{3F(t)}\right], \tag{5.6}$$

where $F(t) = \frac{ea_0}{E_\mathrm{h}} E_L(t)$ is the electric field strength and $\epsilon_I = \frac{\varepsilon_I}{E_\mathrm{h}}$ is the ionization potential, both measured in Hartree atomic units where $E_\mathrm{h} = \alpha^2 m_\mathrm{e} c^2$, the Bohr radius $a_0 = \hbar/(\alpha m_\mathrm{e} c)$, and α is here the fine structure constant[‡].

The rate of ionization therefore increases rapidly around a threshold electric field

[‡]Note that everywhere else in the thesis, F is the F-number of an optic, a_0 is the normalized laser potential and α is the axicon approach angle, but I'm not going to presume to re-label fundamental constants...

strength, where the exponent is equal to -1, described by:

$$F_{\text{th,tnl}} = \frac{2}{3}(2\epsilon_I)^{1.5} \tag{5.7}$$

$$E_{\text{th,ADK}} = \frac{4}{3}\sqrt{\frac{2m_\text{e}e}{\hbar^2}}\varepsilon_I^{1.5}, \tag{5.8}$$

where again, the ionization energy is measured in electron volts. For the hydrogen atom, the threshold field is $E_L \approx 340\,\text{GVm}^{-1}$, requiring $I_L \approx 1.6 \times 10^{16}\,\text{Wcm}^{-2}$. Tunnelling ionization occurs at intensities well below this threshold at a much slower rate.

For low ionization energies, this intensity is significantly higher than the threshold for barrier suppression ionization and the rate of tunnelling is small. As the ionization energy increases, however, the BSI threshold increases faster than the tunnelling threshold, and the situation can no longer be described classically. In this regime, electrons tunnel through the potential barrier significantly before the barrier suppression field is reached. The cross-over between these two regimes occurs when the thresholds are roughly equal, giving an ionization energy described by:

$$\begin{aligned}\frac{\pi\epsilon_0}{Ze}\varepsilon_I^2 &= \frac{4}{3}\sqrt{\frac{2m_\text{e}}{e^2\hbar^2}}\varepsilon_I^{1.5} \\ \varepsilon_I &= \left(\frac{4}{3}\sqrt{\frac{2m_\text{e}e}{\hbar^2}}\frac{Ze}{\pi\epsilon_0}\right)^2 \\ &\approx 1500Z^2\,\text{eV}\end{aligned} \tag{5.9}$$

In practice, however, tunnelling ionization plays a significant role at a much lower ionization energy than this.

5.2.2 Simpleman Model

To understand the energy distribution of electrons produced from Optical Field Ionization we shall use the Simpleman model as described by [67]: a classical approximation of the energy accumulated by an electron as it moves in the electric field of the laser. If an

electron is liberated from the atom at rest at a time t_0, integrating the Lorentz equation we obtain:

$$\mathbf{p}(t) = -\int_{t_0}^{t} e\mathbf{E_L}(t')\mathrm{d}t' \tag{5.10}$$

We can break the electric field into a sinusoidal oscillation within a slowly changing envelope $\mathbf{E_L}(t) = \Re\left[\mathbf{E_0} f(t) e^{-i\omega t}\right]$, where $\mathbf{E_0} \in \mathbb{C}^3$ is a time independent vector describing the polarisation of the laser. We can then integrate by parts, ignoring the derivative $f'(t) \ll \omega f(t)$:

$$\begin{aligned}
\mathbf{p}(t) &= \Re\left(\frac{e\mathbf{E_0}}{i\omega}\left\{\left[f(t)e^{-i\omega t'}\right]_{t_0}^{t} - \int_{t_0}^{t} f'(t')e^{-i\omega t'}\mathrm{d}t'\right\}\right) \\
&\approx \frac{e}{\omega^2}\left[-i\omega f(t)e^{-i\omega t'}\right]_{t_0}^{t} && (5.11) \\
&= \frac{e}{\omega^2}\left[\mathbf{E'_L}(t) - \mathbf{E'_L}(t_0)\right] && (5.12)
\end{aligned}$$

If we look at the final energy of the electron E_K after the laser pulse has passed, $E_L(t) \to 0$, the energy is therefore related to the derivative of the electric field at the moment it is ionized:

$$\begin{aligned}
E_K &= \frac{p(t)^2}{2m_\mathrm{e}} && (5.13) \\
&\approx \frac{e^2}{2m_\mathrm{e}\omega^4}\left|\mathbf{E'_L}(t_0)\right|^2 && (5.14)
\end{aligned}$$

That is, the oscillations in the electric field almost entirely cancel each other out, such that the final energy of the electron is dependent on the derivative of the electric field at the moment of ionization, and not on the laser pulse profile thereafter.[§] This derivative depends on the polarisation state of the laser pulse. An elliptically polarised

[§]It is also possible to reach this result neatly by considering the conservation of the canonical momentum $\mathbf{p} + e\mathbf{A}$.

laser pulse is described by:

$$\mathbf{E_L}(t) = f(t)\left(\cos\{\omega t\}, \epsilon \sin\{\omega t\}, 0\right), \tag{5.15}$$

where $f(t)$ is a slowly varying envelope and $0 \leq \epsilon \leq 1$ is a measure of the ellipticity: $\epsilon = 0$ corresponds to linear polarization and $\epsilon = 1$ to circular. If $f(t)$ is sufficiently slowly varying, the derivative can be written as:

$$\mathbf{E'_L}(t_0) \approx \omega f(t_0)\left(-\sin\{\omega t_0\}, \epsilon \cos\{\omega t_0\}, 0\right) \tag{5.16}$$

In the BSI case an electron will be ionized where the electric field first reaches the threshold E_{th}. If the envelope is sufficiently slowly varying compared to the laser frequency ω this will be where $\cos\{\omega t_0\} = 1$ and $f(t_0) = E_{\mathrm{th}}$, giving:

$$\mathbf{E'_L}(t_0) = \omega E_{\mathrm{th}}\left(0, \epsilon, 0\right) \tag{5.17}$$

$$E_K \approx \frac{e^2}{2m_{\mathrm{e}}\omega^2}\epsilon^2 E_{\mathrm{th}}^2 \tag{5.18}$$

$$\approx \frac{\pi^2\epsilon_0^2}{2m_{\mathrm{e}}\omega^2}\frac{\epsilon^2\varepsilon_I^4}{Z^2} \tag{5.19}$$

The energy distribution will therefore have a sharp peak at an energy which varies with the laser ellipticity and the charge and ionization energy of the species in question, but is largely insensitive to the temporal profile of the laser pulse.

For the tunnelling case, however, electrons are ionized continuously. t_0 therefore takes a range of values and the energy will be distributed in a broader peak. The electron energy distribution will therefore depend more strongly on the temporal profile of the laser pulse, rather than just the threshold intensity. As the ionization energy approaches the point described by eqn. (5.9), tunnelling ionization begins to play a greater role and this will result in broader peaks in the energy distribution of the electrons, the shape of which will depend more strongly on the temporal profile of the laser pulse.

5.2.3 Simulations with EPOCH

We can combine these models with simulations using the Particle-In-Cell (PIC) code EPOCH [76]. This models ionization with a rate equation [107], with the ionization rate calculated using both the ADK model [104] and the Posthumus et al. description of the BSI model [106], as well as Multi-Photon Ionization, depending on the field strength. The subsequent motion of particles during the passage of the laser pulse is calculated using the Lorentz force equation, updating the electro-magnetic fields self-consistently with Maxwell's equations. See Chapter 2 for more details on EPOCH and the PIC scheme in general.

We therefore have three different ways of estimating the electron energy spectra: using naive BSI alone gives a single estimate for the peak energy; numerically integrating the ionization rates from the ADK model would give an energy distribution based on the Simpleman model; and an EPOCH simulation uses both ADK and the improved BSI model from Posthumus et al to simulate the energies of a population of electrons, which can be binned to form a distribution. As I am considering ionization energies $\varepsilon_I \ll 1000\,\mathrm{eV}$ I shall focus on BSI calculations and EPOCH simulations. Simulations in EPOCH used a circularly polarised laser pulse with a duration of 30 fs FWHM passing through hydrogen, helium and nitrogen plasmas with initial densities such that after full ionization $n_e = 10^{18}\,\mathrm{cm}^{-3}$. For simplicity I treated all of these species as being atomic, but used the first ionization energy of *molecular* hydrogen, which is diatomic[¶]. I therefore used the ionization energies and ion charges of hydrogen as $\varepsilon_I = 15.4\,\mathrm{eV}$, for $Z = 1$, helium as $\varepsilon_I = \{24.6\,\mathrm{eV}, 54.4\,\mathrm{eV}\}$, for $Z = \{1, 2\}$, and nitrogen

[¶]The situation for ionizing diatomic molecules is substantially more complex than for atoms. In practice there are several different ionization pathways depending on when the molecular bond breaks and the atoms dissociate (see for instance [108]). While the molecular bond is associated with a much lower energy change than ionization, and hence has a much lower threshold for the electric field strength, it involves the motions of ions, which takes place on a much longer timescale than electron motion. This means when using very short pulses we may in practice have to consider the ionization of molecules, such as in the process $H_2^+ \rightarrow H_2^{2+}$. This has a much higher ionization energy than $H \rightarrow H^+$ and will drastically increase the laser intensity required to fully ionize hydrogen. However, in this chapter I assume that the laser pulses are sufficiently long that we can neglect this effect.

as $\varepsilon_I = \{14.5\,\text{eV}, 29.6\,\text{eV}, 47.4\,\text{eV}, 77.5\,\text{eV}, 97.9\,\text{eV}\}$, for $Z = \{1, 2, 3, 4, 5\}$.

First, I ran EPOCH simulations to determine the threshold for full ionization and compared those thresholds to those expected from the naive BSI model described by eqn. 5.5. I defined the threshold as the minimum intensity required to ionize more than 99% of the electrons in the species during the passage of the laser pulse. The results are shown in table 5.1.

			$I_{\text{th}}(\text{Wcm}^{-2})$	
Species	Z_{max}	$\varepsilon_{I,\text{max}}$	BSI	EPOCH
Hydrogen	1	15.4 eV	4.5×10^{14}	2.5×10^{14}
Helium	2	54.4 eV	1.8×10^{16}	2.0×10^{16}
Nitrogen	5	97.9 eV	2.9×10^{16}	4.5×10^{16}

Table 5.1: Threshold intensities for ionizing different species to a full valence shell from the BSI model and from EPOCH simulations.

This demonstrates that while both estimates agree on the order of magnitude of intensity required, they can differ substantially on the exact thresholds. The intensity required to ionize increases dramatically with the ionization energy and hence hydrogen becomes fully ionized at a much lower threshold intensity than helium or nitrogen. However, the higher charge of nitrogen ionized to N^{5+} means that the intensity required is on the same order of magnitude as for He^{2+}, despite the significantly higher ionization energy.

We can also look at the energy distributions of electrons when ionized with fields above and below the threshold intensity required to ionize the species to noble-gas-like electron configurations. I ran EPOCH simulations with laser intensities below, at, and twice the threshold intensity determined in EPOCH. The resulting energy distributions, shown in Fig. 5.2, contain several peaks, one for each ionization level of the species in question. Each peak is associated with electrons ionized at a different intensity, with the energy of the peak close to the ponderomotive potential at that point. Higher ionization potentials require higher threshold intensities and result in higher energy electrons.

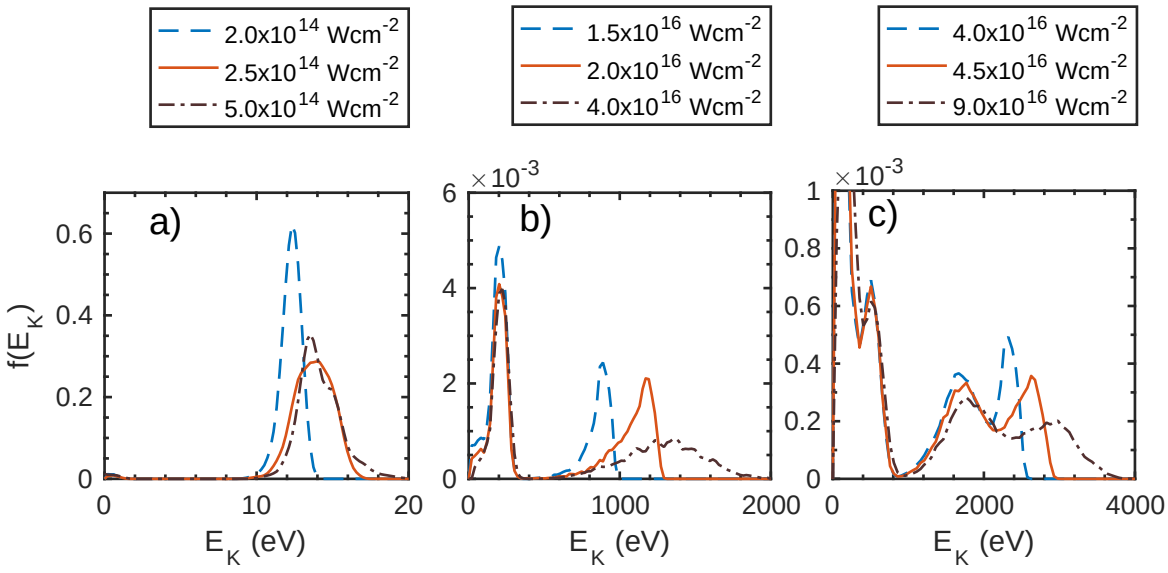


Figure 5.2: Simulated electron energy distributions resulting from OFI for a) hydrogen, b) helium, and c) nitrogen, with laser intensities at $< I_{\text{th}}$ (blue dashed), $\approx I_{\text{th}}$ (red) and $\approx 2I_{\text{th}}$ (brown dash-dotted). Simulations were conducted in EPOCH using a circularly polarised laser pulse, with a final electron density of $10^{18}\,\text{cm}^{-3}$.

We can also see the effect of a transition between BSI and tunnelling ionization. The lower energy peaks in the energy distribution, corresponding to the first electrons, are sharper and vary much less with laser intensity once the intensity is above threshold. As electrons are liberated very rapidly around this low threshold intensity, the subsequent profile of the laser pulse does not affect the energy distribution. However, for the higher energy peaks, and intensities below the BSI threshold, tunnelling ionization plays a much greater role, and the peaks in the distribution are much broader, becoming broader still as the laser intensity increases.

Next, we can study the effect of the ellipticity of the laser pulse. I ran simulations in EPOCH at the threshold laser intensity for circular polarisation, changing between $\epsilon = 0$ (linear polarisation), 0.5 (elliptical), and 1 (circular). The threshold laser intensity was taken from the EPOCH simulations shown above, and set to be $2.5 \times 10^{14}\,\text{Wcm}^{-2}$

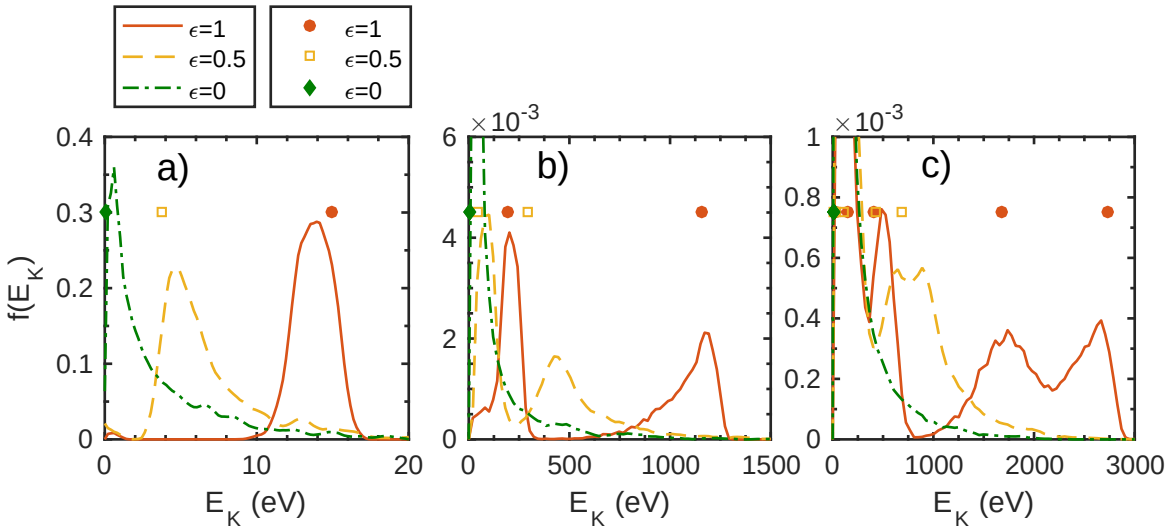


Figure 5.3: Simulated electron energy distributions resulting from OFI for a) hydrogen ionized to H^+ using a $2.5 \times 10^{14}\,\mathrm{Wcm^{-2}}$ intensity laser pulse, b) helium ionized to He^{2+} at an intensity of $2.0 \times 10^{16}\,\mathrm{Wcm^{-2}}$, and c) nitrogen ionized to N^{5+} using an intensity of $4.5 \times 10^{16}\,\mathrm{Wcm^{-2}}$. The polarisation state of the laser was varied between linear ($\epsilon = 0$ in green dash-dotted), elliptical ($\epsilon = 0.5$ in yellow dashed), and circular ($\epsilon = 1$ in red). Predictions from the Simpleman model using the simulated laser intensity are marked for circular polarisation (filled circles), elliptical (empty squares) and linear (filled diamonds).

for hydrogen, $2.0 \times 10^{16}\,\mathrm{Wcm^{-2}}$ for helium, and $4.5 \times 10^{16}\,\mathrm{Wcm^{-2}}$ for nitrogen, while the duration of the pulse was set as 30 fs FWHM, as before. The Simpleman model for BSI can also predict the electron energies using eqn. (5.18), where the value of E_{th} is taken from the threshold intensities used in the EPOCH simulations.

The expected energies are plotted along with the energy distributions resulting from simulations in Fig. 5.3. For all species, reducing the ellipticity results in a decrease in the energy of the electrons, in agreement with the predictions of the Simpleman model. For linearly polarised laser pulses the most likely energy for the electrons is zero. However, electrons are not ionized instantaneously and therefore carry a range of energies. For a linearly polarised laser pulse with an intensity of $4.5 \times 10^{16}\,\mathrm{Wcm^{-2}}$ passing through nitrogen, the electrons are left with a mean energy of $\langle E_K \rangle \approx 190\,\mathrm{eV}$, which is still substantial. The mean energy of electrons ionized by a circularly polarised laser pulse of the same intensity is, however, much higher, at $\langle E_K \rangle \approx 1000\,\mathrm{eV}$. We can

summarise the change in electron energies in table 5.2.

	$\langle E_K \rangle$ (eV)		
Species	$\epsilon = 0$ (linear)	$\epsilon = 0.5$ (elliptical)	$\epsilon = 1$ (circular)
Hydrogen	3.5	6.7	13.7
Helium	140	320	630
Nitrogen	190	480	1000

Table 5.2: Average electron energies resulting from ionizing species using different laser polarizations, calculated from EPOCH simulations.

5.2.4 Discussion

From these simulations, we expect that it is possible to generate extremely hot electrons on the femtosecond timescale through optical field ionization using elliptically polarised laser pulses. The threshold intensity for OFI can be estimated using the barrier suppression ionization model, although tunnelling ionization also plays an important role, particularly at higher ionization energies. The energy of the resulting electrons scales strongly with the threshold intensity for ionization rather than with the peak intensity of the laser pulse. This has a strong dependence on the ionization energy, giving electrons with energies $E_K \propto \varepsilon_I^4$. Different species give rise to very strongly varying electron energy distributions, particularly for circularly polarized laser pulses, with peaks in the energy distribution corresponding to the different ionization levels of the species. Furthermore, using the same species and the same laser intensity it is possible to control the energy of electrons produced by optical field ionization simply by controlling the ellipticity of the polarisation state, with the mean electron energy varying by a factor of four or more between linear and circular polarization.

We therefore have a means of rapidly creating hot plasma in order to drive hydrodynamic shocks, independently of the density of the target plasma. The laser intensities required are much lower than commonly used for wakefield acceleration, and hence are not too demanding for applications. In the rest of this chapter I shall focus only

on hydrogen, as the simplest and easiest to ionize of these species, and on circularly polarized light, as the most effective way of heating a plasma.

5.3 Shock Expansion in Plasma

5.3.1 Fluid Approximation

Most theoretical models for studying shock expansion use a fluid model of the plasma and the majority of simulation codes used for this purpose, such as FLASH [109] and HELIOS [110], as well as a code developed by our collaborator at the University of Maryland, Howard Milchberg [63], [64], [111], are hydrodynamic, studying the movement of fluid elements which are assumed to have Maxwellian energy distributions with well defined temperatures. However, the non-isotropic and non-thermal electron momentum distributions we have predicted from OFI heating are not well described by these approximations and this complicates the calculation of the subsequent expansion of the plasma columns. For the general case where species are multiply ionized to form ions with atomic charge $Z > 1$, ionization of each valence electron gives rise to a 'class' of electrons with a mean energy determined by its ionization energy, and the resulting energy distributions contain several distinct peaks.

However, particles in a class collide with each other and hence for each class the momentum distribution will over time become isotropic, and the energy distribution will thermalize. The timescale for both processes is of the same order as the Spitzer electron self collision time [112], which is described by:

$$\begin{aligned}\tau_{\mathrm{c}}^{\mathrm{e}} &\approx 1.4\frac{v_{\mathrm{e}}^3}{8\pi r_{\mathrm{e}}^2 c^4 n_{\mathrm{e}} \ln\Lambda} \\ &\approx 3.3\times 10^5 \frac{T_{\mathrm{e}}[\mathrm{eV}]^{\frac{3}{2}}}{n_{\mathrm{e}}[\mathrm{cm}^{-3}]\ln\Lambda}\ \mathrm{s},\end{aligned} \tag{5.20}$$

where $\ln\Lambda$ is the Coulomb logarithm, $v_{\mathrm{e}} = \sqrt{2E_K/m_{\mathrm{e}}}$ is the r.m.s. electron velocity,

$T_\mathrm{e} = 2E_K/3$ is the effective temperature of the class in electron volts, and the electron density is measured in cm^{-3}.

For $T_\mathrm{e} = 10\,\mathrm{eV}$ and $n_\mathrm{e} = 10^{18}\,\mathrm{cm}^{-3}$, we find $\ln\Lambda \approx 12$ and hence $\tau_\mathrm{c}^\mathrm{e} \approx 0.9\,\mathrm{ps}$. Meanwhile, electrons from the different energy classes also collide and will exchange energy, eventually forming a single thermal distribution at a rate of order the electron equilibriation frequency, described by [113]:

$$\nu_\epsilon^{\alpha/\beta} \approx 1.8\times10^{-19}\frac{(m_\alpha m_\beta)^{\frac{1}{2}} Z_\alpha^2 Z_\beta^2 n_\beta \ln\Lambda_{\alpha\beta}}{(m_\alpha T_\beta + m_\beta T_\alpha)^{\frac{3}{2}}}\mathrm{s}^{-1}, \tag{5.21}$$

where all units are in c.g.s. and the temperature is again measured in eV. For the case of electron-electron collisions this reduces to an equilibriation time described by:

$$\tau_\epsilon^{\mathrm{e/e}} \approx 1.7\times10^{5}\frac{(T_\alpha[\mathrm{eV}] + T_\beta[\mathrm{eV}])^{\frac{3}{2}}}{n_\beta[\mathrm{cm}^{-3}]\ln\Lambda_\mathrm{ee}}\mathrm{s}, \tag{5.22}$$

where T_α and T_β are the temperatures in electron volts of the two different electron classes after both have thermalised, and the electron density n_β is again measured in cm^{-3}.

This timescale is on the same order as the self-collision time. For example, for two classes with temperatures of $T_\alpha = 10\,\mathrm{eV}$ and $T_\beta = 20\,\mathrm{eV}$ at a density of $n_\beta = 10^{18}\,\mathrm{cm}^{-3}$, the equilibriation time is $\tau_\epsilon^{\mathrm{e/e}} \approx 2\,\mathrm{ps}$. In this situation, after a few picoseconds electrons will comprise a single thermal and isotropic population, which is well described by the fluid approximation[‖]. However, if we were to generate electrons with energies on the order of 100 eV, the collision times increase to $\tau_\mathrm{c} \sim \tau_\epsilon \sim 10$ ps. The thermal velocity $v_\mathrm{e} = \sqrt{2E_K/m_\mathrm{e}} \sim 10^7\,\mathrm{ms}^{-1}$ also increases, and in the time taken to equilibriate electrons can travel around $100\,\mu$m. A non-thermal population of electrons could therefore distinctly change the physics of the shock evolution. In fact, the mean-free path scales

[‖]Recent, as yet unpublished, experiments exploring this process through Thomson scattering measurements have measured that this process is actually faster than expected due to kinetic instabilities [114].

as $\lambda_{\mathrm{mfp}} \propto {T_{\mathrm{e}}}^2$ and hence in the higher energy regime the hydrodynamic simulations may be less accurate; this would be an interesting topic for future research.

Electrons and ions will also exchange energy over time; the timescale for equilibration of energy between electrons and ions is,

$$\nu_{\epsilon}^{e/i} \approx 1.8 \times 10^{-19} \frac{(m_{\mathrm{e}} m_{\mathrm{i}})^{\frac{1}{2}} Z_{\mathrm{i}}^2 n_{\mathrm{i}} \ln \Lambda_{\mathrm{ei}}}{(m_{\mathrm{e}} T_{\mathrm{i}} + m_{\mathrm{i}} T_{\mathrm{e}})^{\frac{3}{2}}} \mathrm{s}^{-1}. \tag{5.23}$$

As $m_{\mathrm{i}} \gg m_{\mathrm{e}}$ and $T_{\mathrm{e}} > T_{\mathrm{i}}$ we can reduce this to:

$$\tau_{\epsilon}^{\mathrm{e/i}} \approx 3.1 \times 10^{8} \frac{A T_e^{\frac{3}{2}}}{Z_i^2 n_i [\mathrm{cm}^{-3}] \ln \Lambda_{\mathrm{ei}}} \mathrm{s}, \tag{5.24}$$

For a density $n_{\mathrm{e}} = 10^{18}\,\mathrm{cm}^{-3}$ and a singly charged hydrogen ion $A = Z_i = 1$, $\ln \Lambda_{\mathrm{ei}} \approx 6$, giving $\tau_{\epsilon}^{\mathrm{e/i}} \approx 2\,\mathrm{ns}$, significantly larger than the electron equilibriation time. This implies that as the shock expansion begins the ions remain cold, and only thermalise with the electrons during the course of the expansion.

A further difficulty occurs with charge separation, which is not accounted for in our hydrodynamic simulations. The hot electrons produced by OFI will start to stream out of the initial plasma column but will be held back by the cold, positive ions. However, as the Debye length $\lambda_D = \sqrt{\frac{\epsilon_0 k_B T}{n_e e^2}} \sim 10$ nm, the charge separation is minimal and the ions follow the electrons closely. Some of the electrons will have sufficiently high energies to escape from the plasma column and stream freely outwards, but if a fraction α of high energy electrons have already escaped a plasma column of radius r_{c}, only electrons with a kinetic energy of $E_K > \alpha m_{\mathrm{e}} \omega_{\mathrm{p}}^2 r_c^2$ can escape. For a plasma column on the scale of $r_{\mathrm{c}} \approx 10\,\mu\mathrm{m}$, with a density $n_e \sim 10^{18}\,\mathrm{cm}^{-3}$, this condition yields $\alpha \lesssim 10^{-6} E_K[\mathrm{eV}]$, and hence we expect only a very small proportion of electrons, on the scale of a few parts per million, to stream out of the expanding plasma, and for the plasma to remain locally neutral.

The considerations above show that after a time on the order of a few ps the electron

distribution will be isotropic and thermal, with a temperature in electron volts given by $eT_\mathrm{e} = (2/3)\langle E_K \rangle$, where $\langle E_K \rangle$ is the mean electron kinetic energy immediately after OFI. The ions will remain cold and the electrons will remain confined by the ions, with the plasma approximately charge neutral. The subsequent expansion of the plasma can therefore be modelled by fluid codes.

5.3.2 Shock Expansion Theory

Now we have described the formation of a thermal distribution resulting from OFI heating we can use the results of some standard shock theory. First, if the plasma column has an initial radius r_0, contains a total energy per unit length $E_\sigma = Zn_0E_K\pi r_0^2$, and is doing work against the background gas at a pressure $p = n_0k_BT_0$, we can find the maximum possible radius to which the column can expand, r_f.

The total work done against the background pressure is:

$$\Delta W = \iiint_{r_0}^{r_f} p\mathrm{d}V. \tag{5.25}$$

This implies an initial energy per unit length of:

$$E_\sigma = n_0k_BT_0 \iint_{r_0}^{r_f} \mathrm{d}A \tag{5.26}$$

$$(Zn_0E_K)\,\pi r_0^2 = n_0k_BT_0\pi\left(r_f^2 - r_0^2\right) \tag{5.27}$$

Rearranging this gives an estimate for r_f described by:

$$\frac{r_f}{r_0} = 1 + \sqrt{\frac{ZE_K}{k_BT_0}} \tag{5.28}$$

For molecular hydrogen we have found $E_K \approx 10\,\mathrm{eV}$, which at room temperature implies an initial column can expand to $r_f \approx 21r_0$ before it expends all of its energy. This is actually a significant overestimate as it assumes that no energy crosses the shock

boundary, where it would be dissipated to the background without driving the expansion. In practice any heat carried out of the plasma column will substantially reduce the size a shock can reach.

It is also possible to describe the rate at which a shock expands. A useful analytic expression for the temporal evolution of the radial position of the shock front is provided by the Sedov-Taylor solution for expansion of ideal gases, assuming $T_\mathrm{e} = T_\mathrm{i}$ [115]. For an idealized system comprising an initial cylindrical region of infinitesimal radius and energy per unit length E_σ, expanding into an unshocked region of density ρ_0 we expect that the radius of the shock front r_s and time τ can be related by a dimensionless constant θ:

$$\theta = \frac{E_\sigma \tau^2}{\rho_0 r_s^4} \tag{5.29}$$

where $\tau = 0$ is defined to be the hypothetical point when the shock is infinitely small. We can also set our energy to be the kinetic energy of the electrons E_K, in an initial plasma column of radius r_0:

$$E_\sigma = (Z n_0 E_K)\, \pi r_0^2 \tag{5.30}$$

Using $n_0 \approx \rho/m_\mathrm{i}$, where $m_\mathrm{i} \gg m_\mathrm{e}$ is the mass of the ions, allows us to describe the shock radius as:

$$r_s(t) \propto \sqrt[4]{\frac{Z E_K}{m_\mathrm{i}}} \sqrt{r_0(t + \tau_0)}, \tag{5.31}$$

where $\tau_0 \equiv \tau - t > 0$ is an offset in time to account for a plasma column starting with a finite size at a measured time $t = 0$. If we could track the evolution of the shock front backwards, the shock would be infinitesimally small at $t = -\tau_0$.

We can find the pre-factor by conserving mass, energy and momentum flow over the shock boundary, assuming a sufficiently thin and adiabatic shock front. This gives the

Rankine-Hugoniot shock conditions [116], from which we can find [117]:

$$r_s(t) = \sqrt[4]{\frac{(\gamma+1)^2}{\pi}\frac{ZE_K}{m_\mathrm{i}}}\sqrt{r_0(t+\tau_0)}, \tag{5.32}$$

where γ is the adiabatic constant for the species in question, normally taken as $\gamma = 5/3$ for a plasma with no additional degrees of freedom.

We can additionally use the initial boundary condition $r_s(t=0) = r_0$ to find:

$$\tau_0 = r_0\left[\frac{(\gamma+1)^2}{\pi}\frac{ZE_K}{m_\mathrm{i}}\right]^{-\frac{1}{2}} \tag{5.33}$$

If we know both the initial radius and energy of a plasma column we can predict its future evolution with no free parameters. However, in practice, a plasma column lacks a hard edge and fitting this expression to real data with r_0 as a coefficient gives an estimate of the effective hard-edge radius of the plasma column, if all of the energy were concentrated into a uniform cylinder.

5.3.3 Simulation Codes

We have modelled the hydrodynamic expansion of OFI plasma columns using two different Lagrangian single fluid codes: HELIOS [110], and a code developed by our collaborator at the University of Maryland, Howard Milchberg, called Bessel-ATI [63], [64], [111]. Both codes assume Maxwellian energy distributions for electrons and ions, but do not assume $T_\mathrm{e} = T_\mathrm{i}$, which we have justified in Section 5.3.1.

Bessel-ATI includes a model of OFI and solves the Helmholtz wave equation for the propagation of the channel-forming pulse in the evolving plasma [111]. The channel-forming pulse is assumed to be focused by an axicon lens, and the plasma expansion is subsequently calculated using the single-fluid equations for mass, momentum and energy conservation. In addition to modelling OFI, the code includes inverse-bremsstrahlung heating, collisional ionization and recombination, and thermal conduction. However,

in practice we found it was difficult to accurately model thermal conduction across the shock boundary using this code, with accurate simulations freezing after 1 ns of evolution. We therefore used this code to simulate the early part of the evolution and switched to HELIOS for later times.

The HELIOS code does not include OFI and provides only an opacity model for energy deposition. Therefore for the HELIOS simulations the initial electron density profile simulated by Bessel-ATI was used to describe the initial fraction of ionized atoms and the initial electron temperature was calculated from the electron energy spectrum shown in Figure 5.3. The ion temperature, and the temperature of the neutral gas, were assumed to be $T_\mathrm{i} = 298\,\mathrm{K}$ everywhere.

In the course of this work we also attempted to use two other simulation codes, the PIC code EPOCH and the hydrodynamic code FLASH. EPOCH does not treat collisions with neutral particles and so cannot accurately convey the pressure behaviour of the system. Instead, we modelled the neutral gas outside the plasma column as being ionized. Unfortunately, we discovered a bug in the EPOCH collisions routine that led to extreme numerical heating under this regime. We therefore attempted to use FLASH, which has the advantages over HELIOS and Bessel-ATI of being able to model two and three dimensional plasmas. However, while it can handle both ionization, using rate equations between the different populations, and distinct electron and ion temperatures, using a three temperature model, it cannot handle both problems simultaneously, and therefore failed to accurately model our system.

Before we look at simulations made using Bessel-ATI and HELIOS, I will first describe the axicon lens and the focus that results. This sets the initial conditions of our simulations, which are described in Section 5.5.

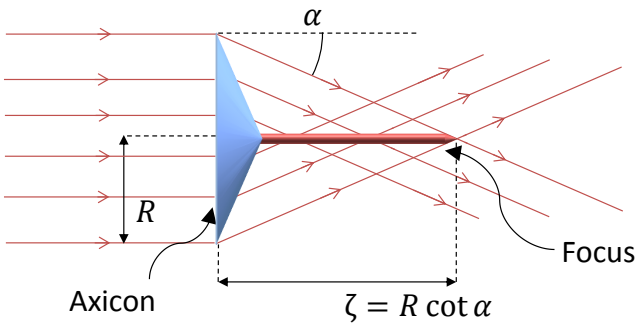


Figure 5.4: Schematic of an axicon focus, with rays converging onto the axis at an angle α.

5.4 Axicons

An axicon lens is a conical prism which causes all parts of a beam with an initially flat wavefront to approach the axis at a constant angle α, forming a concave conical wavefront. A ray arriving at the axicon a radius r' from the center therefore reaches the axis at a point $z = r' \cot \alpha$. All the rays from a radius r', contained within an annulus of area $2\pi r' \delta r'$, reach the axis and constructively interfere at this point. If the beam has a radius R, there is a region of constructive interference extending along the axis from the tip of the axicon at $z = 0$ to $z = \zeta \equiv R \cot \alpha$. For an incident beam with a wavenumber k_1, a total power P, and a uniform intensity I_0 over a top-hat transverse profile with a radius R, the intensity profile at focus is described by:

$$\begin{aligned} I(r,z) &= I_0 \cdot 2\pi z \tan\alpha \cdot k_1 \sin\alpha \cdot J_0^2(k_1 r \sin\alpha) \\ &= \frac{2Pk_1}{R^2} z \sin\alpha \tan\alpha J_0^2(k_1 r \sin\alpha) \\ &= \frac{2Pk_1 \sin\alpha}{R} \frac{z}{\zeta} J_0^2(k_1 r \sin\alpha) \\ &= I_{\max} \frac{z}{\zeta} J_0^2(k_1 r \sin\alpha), \end{aligned} \tag{5.34}$$

where J_0 is the zero-th order Bessel function of the first kind and we have used $P = \pi R^2 I_0$ for a top hat incident beam, and defined $\zeta \equiv R \cot\alpha$, and $I_{\max} \equiv 2Pk_1 \sin\alpha / R$.

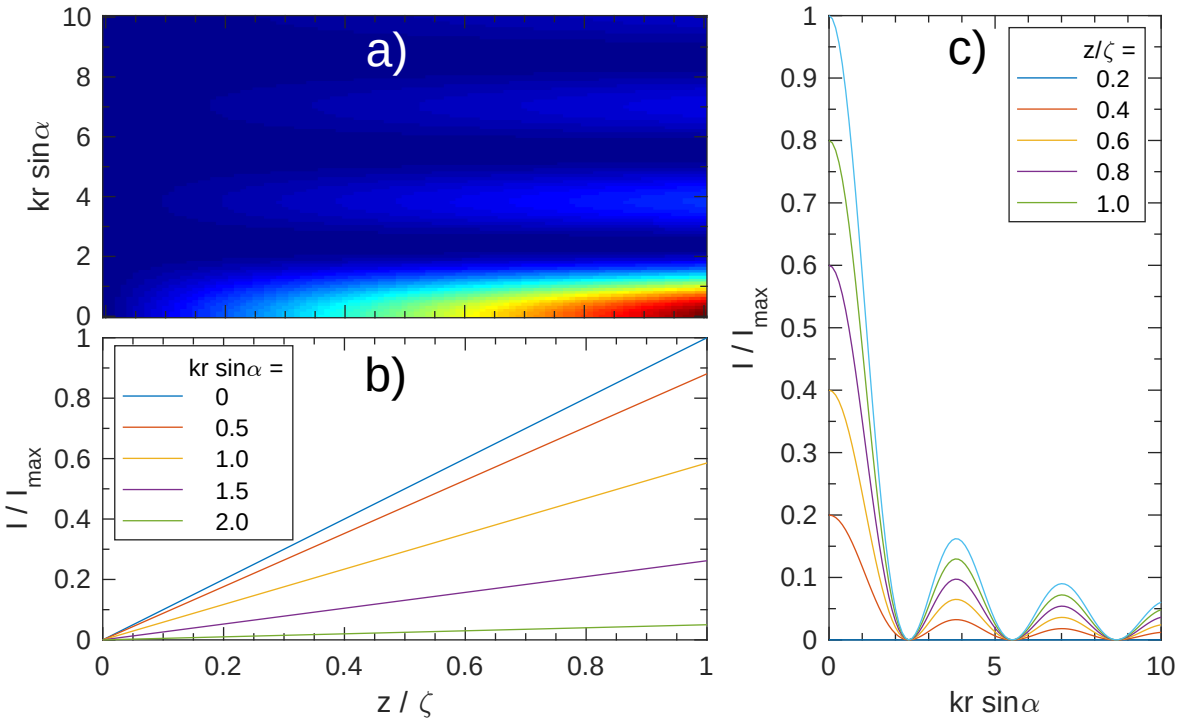


Figure 5.5: Intensity of an ideal axicon focus, showing a) plan view, with higher intensities in red, b) longitudinal profiles and c) transverse profiles. Normalized co-ordinates and intensities are used throughout.

The intensity increases linearly along the axis until reaching a value $I_{\max}$ at the end of the focal region at $z = \zeta$. Moving off-axis, however, the intensity diminishes rapidly, reaching a zero at a radius ρ described by $k_1 \rho \sin\alpha = b_0$, where $b_0 \approx 2.4$ is the first zero of the zero-th order Bessel function J_0. This typically leads to very long but narrow focii, which have radii $\rho = b_0 \lambda_1 / 2\pi \sin\alpha$ and lengths $\zeta = R \cot\alpha$. For an 800 nm wavelength incident pulse and a two-inch diameter axicon with an approach angle $\alpha = 2°$, this gives a spot of radius $\rho \approx 9\,\mu$m along a focus $\zeta = 73\,$cm long. Similar to decreasing the numerical aperture of a spherical lens, decreasing the approach axicon of the axicon both increases the size of the focal spot and extends the length of the focal region, but reduces the intensity at the focus.

However, the choice of axicon angle limits the density of plasma column that can be formed. As the laser pulse ionizes the gas close to focus, a boundary forms between the neutral gas (with a refractive index close to 1) and the plasma (with a refractive index less than 1). This has the potential to be a critical density surface, with total reflection off the plasma column. The component of the wavevector perpendicular to this barrier is $k_\perp = k_1 \sin\alpha$, and hence the effective critical density is substantially reduced, to $n_{\rm crit}(\alpha) = n_{\rm crit}(90^\circ)\sin^2\alpha$, where $n_{\rm crit}(90^\circ) = \pi/r_e\lambda_1^2$ is the usual critical plasma density. The initial density can always be reduced below this, hence reducing the final on-axis density, but reaching higher densities requires increasing the approach angle. For instance, for an approach angle of $\alpha = 2^\circ$, the critical density is $2 \times 10^{18}\,{\rm cm}^{-3}$.

It is also worth mentioning that the axicon focus is a travelling object. As light from the axicon must travel a distance $\zeta \sec\alpha$ in order to reach the end of the focus ζ along the axis, there is a time difference of approximately $\zeta\alpha^2/2c \approx R\alpha/2c$ between light from the axicon and light travelling on axis which has not encountered an axicon. For a two-inch diameter axicon with an approach angle of $\alpha = 2^\circ$, this is a time difference of just 1.4 ps, which is negligible is most contexts. If the maximum allowable window for guiding has a duration Δt, dephasing between the channel forming beam and a beam travelling inside the channel will not be an issue until $\zeta = 2c\Delta t/\alpha^2$.

The finite time it takes an axicon beam to traverse the focus also implies that the plasma column at the start of the focus is older than the plasma column at the end of the focus, in the example above by approximately 2.4 ns, comparable to the time scale over which the shock evolves. If a guided beam were to travel along a waveguide formed from this plasma column, the waveguide might not be formed until just before the moment the guided beam reaches it. This makes it difficult to use axicon formed waveguides in a counter-propagating geometry, where the dephasing becomes an issue if $\zeta \gtrsim c\Delta t$.

That digression aside, in this way, we can form a long and thin plasma column. This will expand, forming a shock, and leave a cavity on axis as described on refs. [63], [64].

By choosing the size of the axicon and the resulting approach angle, we can tune this column to produce the length and radius of cavity that we require to guide a given drive laser. The section below describes how we might choose an axicon in order to produce a given plasma column. This in turn implies a particular set of requirements on the laser pulse which forms the initial plasma column.

5.4.1 Optimising Axicon Parameters

When designing a plasma experiment involving an axicon, there are two main possibilities for the plasma source: a gas cell or a gas jet. When using a gas cell the length of the plasma is generally set by the distance between two pinholes at the entrance and exit of the cell and between the two pinholes the plasma density can be very uniform. Coupling both an LWFA driver beam and the beam from the axicon into the cell can be challenging, however, and we have explored several different options, from using a large front pinhole, to replacing the front pinhole with a thin glass plate; these are described in more detail in refs. [72], [118]. When using a gas jet, on the other hand, it is convenient to couple both beams into the plasma, but there is no clear start or end to the plasma column produced and the gas density might vary strongly with position. When using an axicon though, the plasma column can be controlled more accurately by using a spatial mask. As each annulus of the axicon focusses to a different point on axis, a circular mask will cause the laser intensity on axis to increase suddenly from zero at a certain point and hence create a plasma column, and guiding channel, which starts at a well defined point, for instance at the start of a region of uniform density, as shown in Fig. 5.6.

This intensity increases along the axis and decreases with radius, and so to form a plasma column of length L and radius r_0, this intensity must be greater than a

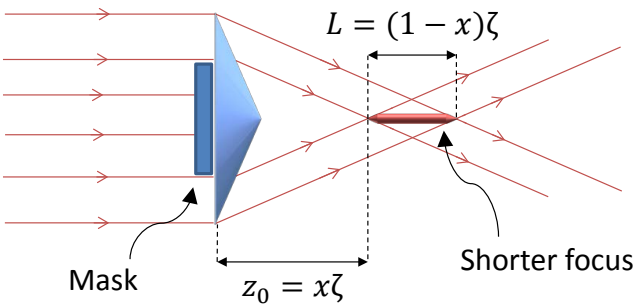


Figure 5.6: Schematic of an axicon focus with a mask, with a length of the focus obscured.

threshold I_{th} at a point (r_0, z_0), where $L = \zeta - z_0$. We can describe this point with the dimensionless parameters $x = z_0/\zeta$ and $y = r_0/\rho$, writing it as:

$$\begin{aligned} I_{\text{th}} = I(r_0, z_0) &= \frac{2Pk_1 \sin\alpha}{\zeta \tan\alpha} x J_0^2(b_0 y) \\ &\approx 2Pk_1 \frac{1-x}{L} x J_0^2(b_0 y). \end{aligned} \tag{5.35}$$

Rearranging this allows us to find P in terms of x and y as:

$$P = \frac{\lambda_1 I_{\text{th}} L}{4\pi} J_0^{-2}(b_0 y) \frac{1}{x(1-x)}. \tag{5.36}$$

This is minimised by $x = x_0 = 0.5$, which demonstrates that the mask we should use should obscure half the radius of the beam and therefore create a plasma column half the length of the axicon focus. This allows us to find the minimum power required to create such a plasma column with a top hat beam **.

$$P_{\text{min}} = \frac{\lambda_1 I_{\text{th}} L}{\pi} J_0^{-2}(b_0 y). \tag{5.37}$$

This power is also minimised by $y = 0$, but this would imply a plasma column of zero radius with no energy. In practice we want to transfer as much energy as possible

**Actually, with appropriate manipulation it would be possible to create an annular beam profile, with no mask required. This would use $\frac{3}{4}$ of the energy of a top hat beam, with no energy dumped onto the mask.

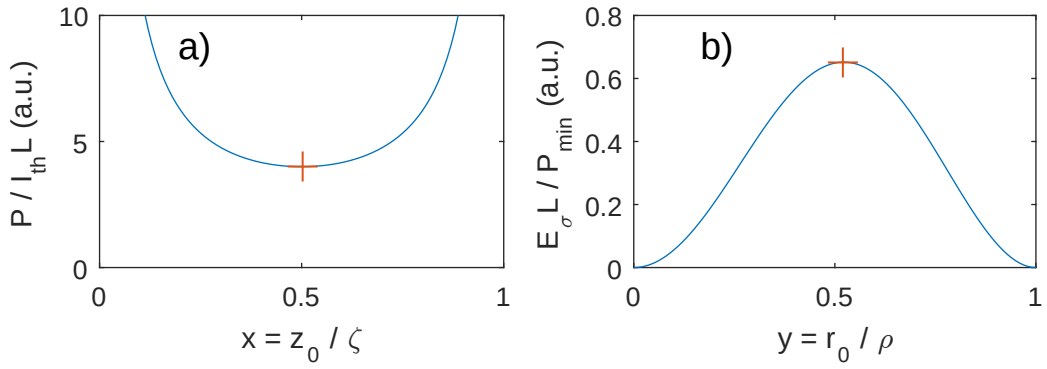


Figure 5.7: a) The required power against x to ionize a channel of a given length at a threshold intensity of I_{th} with an axicon focus. b) The energy supplied to a plasma column by OFI along an axicon focus against y for a given laser power. The minimum laser power at $x = x_0$ and the maximum supplied energy at $y = y_0$ are both marked (red crosses).

from the channel forming beam into the plasma and hence we calculate the energy per unit length gained by the plasma from OFI ionization at critical density, where the energy deposited per unit length is described by:

$$\begin{aligned} E_\sigma &= E_K n_{\text{crit}}(\alpha)\pi r_0^2 \\ &= \frac{r_e\lambda_1^2}{2\pi c} I_{\text{th}} \frac{\pi^2}{r_e\lambda_1^2} \sin^2\alpha \left(\frac{b_0 y \lambda_1}{2\pi \sin\alpha}\right)^2 \\ &= \frac{1}{8\pi c} I_{\text{th}} \lambda_1^2 \left(b_0 y\right)^2 \end{aligned} \tag{5.38}$$

This is maximised by $y \to \infty$, but this would require infinite input power. Instead we wish to maximise the energy per unit length supplied at a given power, for which we can write:

$$\frac{E_\sigma}{P_{\text{min}}} = \frac{1}{8c}\frac{\lambda_1}{L}\left(b_0 y\right)^2 J_0^2(b_0 y). \tag{5.39}$$

To maximise this for a given input power we must maximise $b_0 y J_0(b_0 y)$, the solution to which can be found numerically, giving $y = y_0 \approx 0.522$. This is equivalent to saying

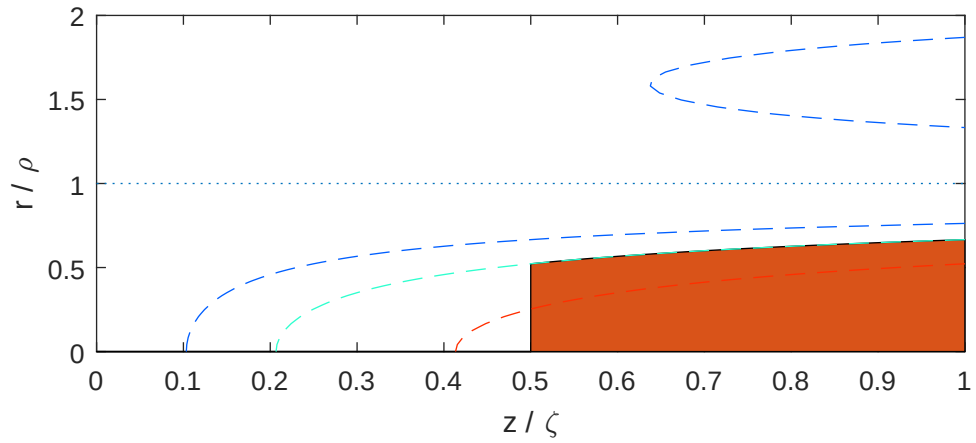


Figure 5.8: The region of the initial plasma column where $I(r, z) > I_{\text{th}}$ (shaded area), compared to contours (dashed lines) without a mask at $0.5I_{\text{th}}$ (blue), I_{th} (green) and $2I_{\text{th}}$ (red).

that we should aim to make the spot size of the axicon focus ρ roughly twice the radius of the desired initial plasma column. From y_0 and x_0 we can determine the optimum initial plasma column, described by the contour (r_0, z_0). This occurs at the threshold intensity, where we can now write:

$$\begin{aligned} I_{\text{th}} = I(r_0, z_0) &= I_{\text{max}} x_0 J_0^2(b_0 y_0) \\ &\approx 0.207\, I_{\text{max}} \end{aligned} \tag{5.40}$$

This ensures that we create a sufficiently wide initial column over a sufficiently long initial length. This is shown in Fig. 5.8, with a mask in front of the axicon preventing ionization before the start of the channel. In practice this could also be created by only supplying target gas above $z \geq 0.5\,\zeta$.

If we require a given plasma column radius size r_0 over a length L, we must therefore

choose an axicon with parameters described by:

$$\begin{aligned} \sin\alpha &= \frac{b_0 y_0 \lambda_1}{2\pi r_0} \\ &\approx 0.2\frac{\lambda_1}{r_0} \end{aligned} \tag{5.41}$$

$$\begin{aligned} R &= \frac{L}{x_{\min}}\tan\alpha \\ &= 2L\tan\alpha \\ &\approx 0.4L\frac{\lambda_1}{r_0} \end{aligned} \tag{5.42}$$

We must then use a laser with a power and total energy in the incident beam described by:

$$P_{\min} \approx 0.770\lambda_1 I_{\text{th}} L \tag{5.43}$$

$$E_{\text{in}} \approx 0.820\lambda_1 t_{\text{FWHM}} I_{\text{th}} L, \tag{5.44}$$

where $E_{\text{in}} \approx 1.06\, P_{\min} t_{\text{FWHM}}$ is the total energy in the beam incident upon the axicon.

For an 800 nm laser with 40 fs FWHM duration, heating hydrogen with to the BSI threshold intensity of $4.5 \times 10^{14}\,\text{Wcm}^{-2}$, this is a total of just 1.18 mJ per centimetre of channel, which is not at all demanding.

The analysis above also allows us to determine the total energy supplied per unit length, described by:

$$E_{\sigma,\max} \approx \frac{0.815}{c}\frac{\lambda_1 P_{\min}}{L}. \tag{5.45}$$

This implies the total efficiency of using OFI over an axicon focus to heat a plasma column can be described by:

$$\eta \equiv \frac{E_{\sigma,\max} L}{E_{\text{in}}} \approx 0.765\frac{\lambda_1}{c t_{\text{FWHM}}}. \tag{5.46}$$

Using an 800 nm wavelength laser pulse with a FWHM duration of 40 fs this efficiency is around 5%.

It is worth noting from this that there are several ways of increasing the energy deposited in the channel. The energy per electron is proportional to $E_K \propto I_{\mathrm{th}}\lambda_1^2$, which can be increased either by changing the gas species to increase the intensity required to ionize, or by increasing the wavelength of the channel forming beam. Whereas changing the gas species requires proportionately more energy in the channel forming beam without changing the efficiency, changing the wavelength leads to an increase in the efficiency. In fact, from eqn. (5.46) the relevant quantity is the number of cycles in the laser pulse, $\eta \propto \lambda_1/ct_{\mathrm{FWHM}} \propto 1/N$. Moving to few cycle pulses with either shorter duration or longer wavelength would increase the fraction of the energy transferred to the plasma column. However, a pulse must still have a duration long enough to fully ionize the plasma. Increasing the wavelength of the channel forming beam, on the other hand, requires more energy per unit length of channel, the axicon angle must increase to maintain the same intensity on axis and critical density, and therefore the axicon radius must also increase to maintain the same plasma column length; there is substantial room for further work on this subject.

In the course of this discussion I have not mentioned another limitation to axicon design: so-called 'B integral'. As an intense laser pulse propagates through a medium the refractive index of the material can change, in the optical Kerr effect. This causes self-phase modulation during the course of the propagation and damages the quality of the wavefront. In order to avoid these effects the B integral must be kept low, below around 1, where $B = k \int n_2 I(z)\ \mathrm{d}z$ and n_2 is the second order non-linear refractive index, which quantifies the Kerr effect for that material [119]. In practice, then, we can

estimate the B integral as:

$$\begin{aligned} B &= k_1 \int n_2 I(z)\, \mathrm{d}z \\ &\sim k_1 n_2 I_0 R \tan\alpha \\ &= k_1 n_2 \frac{P_{\min}}{\pi R} \tan\alpha \\ &= k_1 n_2 \cdot 2.42 \frac{\lambda_1 I_{\mathrm{th}} L}{\pi^2 R} \tan\alpha \\ &= 1.54 n_2 I_{\mathrm{th}} \frac{L \tan\alpha}{R} \\ &= 0.77 n_2 I_{\mathrm{th}}, \end{aligned} \tag{5.47}$$

where we have used $P = I_0 \pi R^2$, estimated the length of the widest part of the axicon as $\sim R \tan\alpha$, and used $L = x_0 R \cot\alpha$.

For fused silica $n_2 \approx 3 \times 10^{-16}$ cm^{2}W^{-1}, and hence B integral should not be an issue so long as $I_{\mathrm{th}} \lesssim 10^{15}$ Wcm^{-2}. While this is true for ionizing hydrogen and the first electron of many species, it could prove a challenge for more strongly bound electrons. In this situation we must consider using a non-transmissive optic, such as a reflective axicon. There are several possible designs for this optic, which is at its simplest a concave conical mirror.

5.4.2 Constraints on Axicon Angle

We have described how to choose an axicon approach angle in order to create a plasma column of a given size. However, this angle is limited in several different ways. First, to guide a laser beam, the focus must be formed at the start of the channel. If the beam is co-propagating with a channel forming beam from an axicon, the numerical aperture of the guided beam must ordinarily be smaller than the approach angle of the channel forming beam[††], giving the constraint $\alpha \geq F_1/2F$ for a desired guided beam F-number

[††]I can think of a geometry which would avoid this, using a pellicle beam splitter, but that would dump a substantial portion of the energy; this constraint holds true for the current experimental set-ups.

$F = f/D$ and a safety factor F_1 which is greater than 1. The spot size of this guided beam is then described by $w_0 = F_2(2\lambda_0/\pi)F$, where λ_0 is the wavelength of the guided beam (which need not equal that of the channel forming beam) and the second safety factor $F_2 = M^2$ is the beam quality factor, which is also greater than 1. Putting these constraints together gives us a limit on the axicon angle required to couple in a beam of spot size w_0.

If w_0 is described by:

$$\begin{aligned} w_0 &= F_2\frac{2\lambda_0}{\pi}F \\ &\geq F_2\frac{2\lambda_0}{\pi}\frac{F_1}{2\alpha}, \end{aligned} \tag{5.48}$$

then the axicon angle is described by:

$$\alpha \geq F_1F_2\frac{\lambda_0}{\pi w_0} \tag{5.49}$$

The spot size ρ of the channel forming beam, however, is limited by the axicon angle as $2\pi\rho\sin\alpha = b_0\lambda_1$. The initial plasma column must be smaller than this spot size, $r_0 < \rho$, and we have described the optimum plasma column radius as $r_0 = y_0\rho$, where $y_0 \approx 0.522$. This plasma column can only expand a certain distance, limiting the ratio r_f/r_0 as described in Section 5.3.2. As the matched spot size is necessarily smaller than the radius of the shock front, $w_m < r_f$, we have a constraint on the matched spot size which we will describe as $w_m \leq \mathcal{C}\rho$ for a parameter $\mathcal{C}$. For molecular hydrogen as discussed in Section 5.2.3, $r_f/r_0 \approx 20$, and hence $r_f/\rho \approx 10$, giving $\mathcal{C} \approx 10$. In practice, a matched spot size smaller than the shock front and dissipation of energy through the shock front will both make $\mathcal{C}$ substantially lower, so I shall introduce a safety factor $w_m/\rho \leq \mathcal{C}/F_3$, where $F_3 \geq 1$ and $\mathcal{C} \approx 0.522\left(1 + \sqrt{ZE_K/k_BT_0}\right)$ depends on the species being ionised and the wavelength of the laser. To reach a certain matched spot size the initial plasma column radius must therefore be sufficiently large, and the approach

angle of the channel forming beam sufficiently small. With the two previous conditions this angle is now constrained from both sides. The matched spot size and the guided spot size can be written as:

$$w_m \leq \frac{\mathcal{C}}{F_3}\rho_0 \tag{5.50}$$

$$w_0 \leq \frac{\mathcal{C}}{F_3}\frac{b_0\lambda_1}{2\pi\sin\alpha}, \tag{5.51}$$

giving a constraint on the axicon angle described by:

$$\sin\alpha \leq \frac{\mathcal{C}}{F_3}\frac{b_0\lambda_1}{2\pi w_0}. \tag{5.52}$$

Combined with the previous condition this gives:

$$F_1F_2\frac{\lambda_0}{\pi w_0} \lesssim \sin\alpha \leq \frac{\mathcal{C}}{F_3}\frac{b_0\lambda_1}{2\pi w_0} \tag{5.53}$$

For some example parameters $F_1 = F_2 = F_3 = 2$, aiming to guide a $w_0 = 30\,\mu$m spot using two 800 nm beams gives a condition $1.9^\circ \lesssim \alpha \leq 2.9^\circ$. Choosing, for instance, a 2-inch diameter axicon with $\alpha = 2.5^\circ$, the plasma column produced is then $L \approx 30\,$cm long, whereas the corresponding Rayleigh range for the guided beam is just $Z_R \approx 3\,$mm.

We can also write these constraints in terms of the F number of the focussing optic as:

$$F_1\frac{1}{2F} \lesssim \sin\alpha \leq \frac{\mathcal{C}}{F_2F_3}\frac{b_0\lambda_1}{4\lambda_0F} \tag{5.54}$$

For instance, moving from an F/20 focussing optic to an F/40 would increase the spot size of the guided beam and hence force α to decrease.

These constraints also limit our safety factors F_1, F_2 and F_3. This can be described

by:

$$F_1 \frac{1}{2F} \leq \frac{\mathcal{C}}{F_2 F_3} \frac{b_0 \lambda_1}{4\lambda_0 F} \tag{5.55}$$

$$F_1 F_2 F_3 \leq \mathcal{C} \frac{b_0 \lambda_1}{2\lambda_0} \tag{5.56}$$

If we use the same wavelength for the guided beam and the channel forming beam, $\lambda_0 = \lambda_1$, we can state that:

$$F_1 F_2 F_3 \lesssim 1.2\mathcal{C} \approx 12 \tag{5.57}$$

For example, if the final matched spot can be no more than $\mathcal{C}/F_3 = 5$ times larger than the radius of the axicon spot ρ and the most tightly we can couple in the guided beam is with an angle $F_1 = 2$ times smaller than the axicon approach angle, the spot quality of the guided beam M^2 must be less than $F_2 \leq 3$, which is ordinarily easily achievable. So long as the safety factors satisfy eqn. (5.56), an axicon can be found that will create a plasma column which will go on to form the required channel. Increasing the wavelength of the channel forming beam, decreasing the wavelength of the guided beam, or changing the gas species to produce more energetic electrons will make this condition easier to satisfy.

On the other hand, changes to the wavelength of the channel forming beam will have repercussions on the critical density, and hence on the on-axis density of the channel that may be formed with a particular axicon. $n_{\text{crit}} \propto (\sin\alpha/\lambda_1)^2 \propto \rho^{-2}$, and so if a certain plasma density is required, the initial plasma column cannot be arbitrarily large without changing the target gas species, with the resulting increase in the laser energy required .

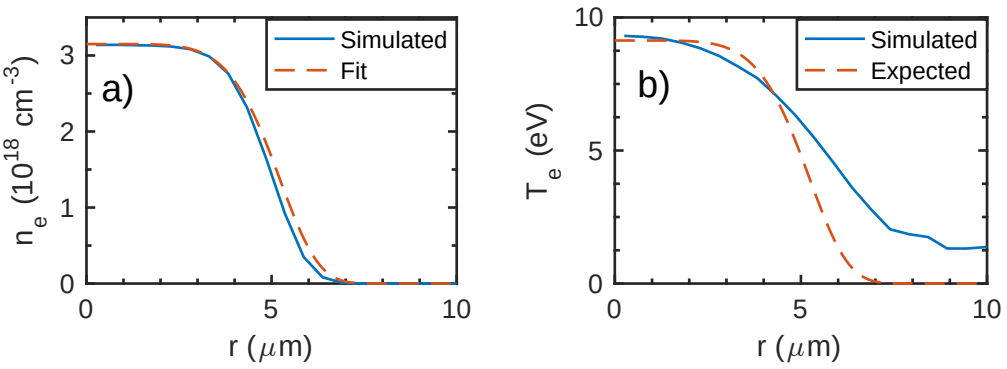


Figure 5.9: Simulated radial profiles for the initial plasma column of: (a) the electron density; and (b) the electron temperature. The results from a simulation using Bessel-ATI, with a laser pulse focussed by an $\alpha = 2.5°$ approach angle axicon into hydrogen gas at 60 mbar. The beam was modelled with a wavelength of 800 nm, a peak intensity of $4.5 \times 10^{14}\,\mathrm{Wcm}^{-2}$ and a FWHM duration of 40 fs, corresponding to 15 mJ in a two-inch diameter beam.

5.5 Simulated Axicon OFI Channels

From the work in Section 5.4.1 above, we can calculate the ideal axicon parameters for a given channel. Imagine we wish to create a channel in hydrogen to guide a matched spot of $w_m = 30\,\mu$m over a distance of $L = 20\,$cm and require an initial plasma column $r_0 \approx 4\,\mu$m in radius. With $x = x_0$ and $y = y_0$ we should then use an axicon with an approach angle of $\alpha \approx 2.5°$ and a channel forming beam with an energy of just 24 mJ in a diameter of $2R \approx 34\,$mm. This gives us a safety factor on the expansion $F_3 \approx 2.3$.

The Bessel-ATI axicon propagation code [63] was used to determine the initial electron density profile created from an axicon focus with an approach angle of 2.5^o. This was incident upon room temperature hydrogen at 60 mbar, chosen so as to produce critical electron density after ionization ($n_e = n_0 = 3 \times 10^{18}\,\mathrm{cm}^{-3}$). The 800 nm wavelength beam was chosen to have an intensity of $4.5 \times 10^{14}\,\mathrm{Wcm}^{-2}$, chosen to match the BSI threshold intensity for molecular hydrogen. As the simulated beam had a Gaussian profile in both time and space, this implies a beam energy of 15 mJ for a FWHM diameter of 2 inches and FWHM duration of 40 fs.

The density and temperature profiles after the laser pulse had passed, 1 ps after the peak of the laser pulse, are shown in Fig. 5.9. The density profile was approximated by a 6th order super-Gaussian with a radius of $r_0 \approx 5\,\mu$m and this was used to set the initial conditions of the HELIOS simulation. This is not a hard-edge column and is slightly broader than predicted from the basic theory described in Section 5.4.1, with partial ionization further off-axis. The temperature profile measured in the simulation was $T_e(r = 0) = 9.3\,$eV, which corresponds to a mean electron energy of $\langle E_K \rangle = 14\,$eV and is very close to expectations from the simulated density profile and Section 5.2.3. The broadening in the electron temperature profile compared to the density profile is due to heat transfer between the ionized plasma and the small number of electrons surrounding the plasma column. This takes place on the picosecond timescale.

The subsequent expansion of the plasma column was modelled using HELIOS, using as initial conditions the fitted electron density and temperature profiles shown in Fig. 5.9. The evolution of the plasma in the simulation is shown in Fig. 5.10, with lineouts showing the density profile changing over time. A shock rapidly forms within the first 0.5 ns after the laser pulse has passed, leaving a cavity on axis behind a sharp shock front. This shock front position moves out over time at a diminishing velocity as the initial energy does work against the neutral gas. The position of this shock front agrees well with the Taylor-Sedov solution, where the effective hard edge plasma column radius was fitted as $r_0 = 4.5\,\mu$m. In the absence of any heat dissipation, eqn. 5.28 predicts that the shock front could therefore reach a maximum radius of around $r_f \approx 100\,\mu$m before the temperature falls to room temperature and the shock stops.

As described in Chapter 1, the matched spot size w_m can also be estimated for the channel by fitting the density profile to the expression below:

$$n_e = n_{e0} + \frac{1}{\pi r_e w_m^2} \left| \frac{r}{w_m} \right|^n , \tag{5.58}$$

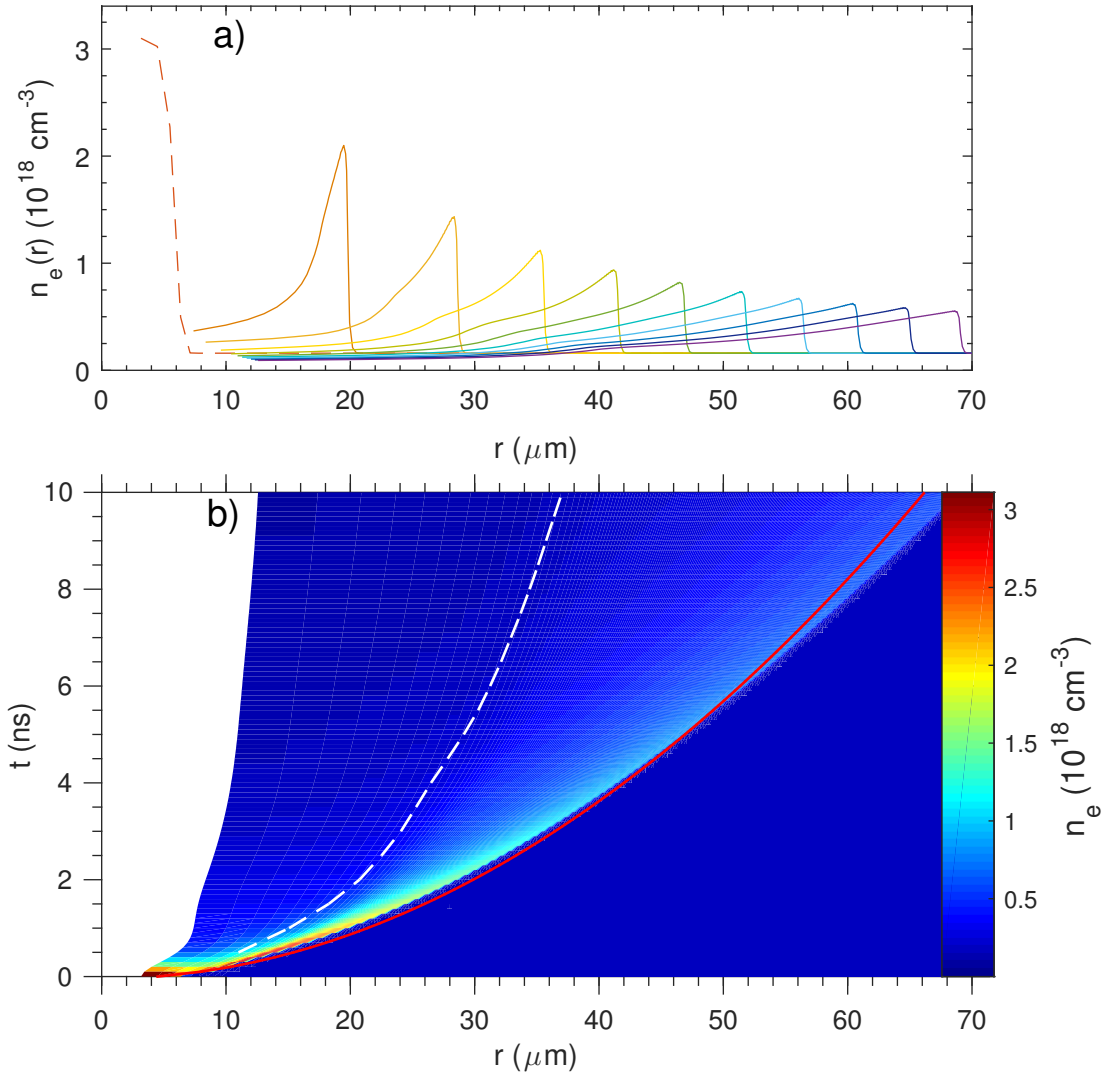


Figure 5.10: Simulated temporal evolution of the radial electron density profiles for an OFI channel produced along an axicon focus. Electron density is shown through a) lineouts at 0.5 ns intervals starting at 0 ns (red dashed), and b) an electron density surface, overlaid with the shock front position $r_s(t)$ expected from the Taylor-Sedov solution (red) and the estimated matched spot size of the plasma channel (dashed white).

where r_e is the classical electron radius. The density profile at each timestep was fitted to eqn. (5.58), using both n_{e0} and w_m as fitting parameters, and terminating the fit at the peak of the shock front. The resulting spot size is overlaid on Fig. 5.10 and predicts that channels can be produced with matched spot sizes of $w_m \approx (20 - 30)\,\mu$m with on-axis densities around $n_{\rm e}(0) \approx 10^{17}\,{\rm cm}^{-3}$. This spot size is approximately half as large as the position of the shock front and the shock front, which along with the good agreement to the Taylor-Sedov solution implies that the shock expansion safety factor is actually around $F_3 \approx 2$, demonstrating that this waveguide would effectively guide laser pulses with spot sizes of $w_0 = 35\,\mu$m.

However, we do not know from the matched spot size alone how effective these channels will be at guiding a laser pulse since the channels are finite, and hence are lossy. For this purpose, a mode analysis code called Leaky, developed by our collaborators at the University of Maryland [120], was used to find solutions to the Helmholtz equation given the different channel profiles, assuming an exponential decay in the power of a mode with length, with the electric field of the form $E(r,z) = f(r)e^{i\kappa z}$. As described in Chapter 1, this allowed us to determine the shape of the modes without assuming eqn. (5.58), and also to determine the length over which a mode leaks energy to outside the channel; the length over which a mode decreases in power by a factor of $1/e$ is given by $L_{\rm att} = [2\Im(\kappa)]^{-1}$. These modes, and the attenuation length for each mode, are plotted in Fig. 5.11. The modes found were close to Gaussian, with attenuation lengths reaching up to the scale of one metre for the larger channels. This demonstrates that the channels formed are suitable for guiding intense laser pulses at low densities over distances much longer than existing waveguides.

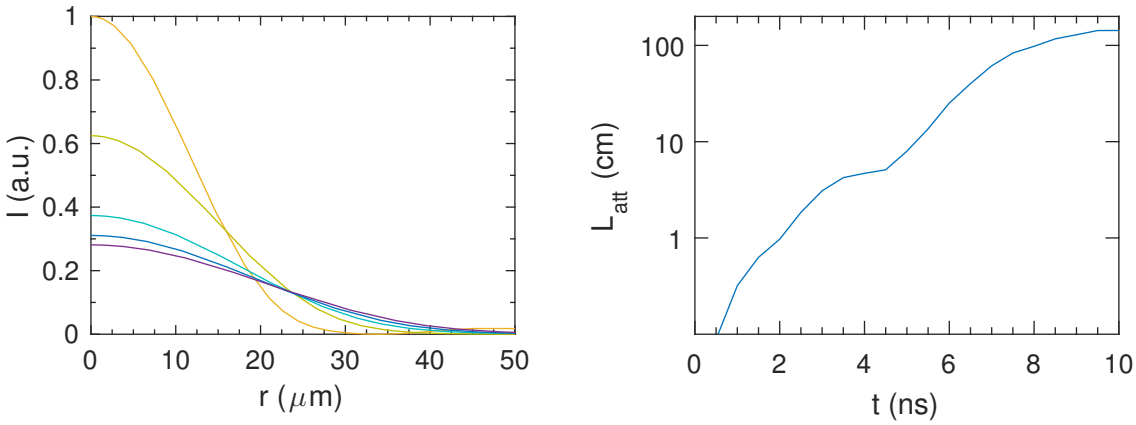


Figure 5.11: Lowest order guided modes of the plasma channel over time, taken at 2 ns intervals from 2 ns until 10 ns (red to blue), with the intensity normalized such that all modes have the same power. Inset: The 1/e power attenuation length of these modes over time, calculated by solving the Helmholtz equation assuming a uniform channel.

5.5.1 Channel Properties

These low densities and long guiding lengths are highly desirable for laser wakefield acceleration, but as yet we have not determined what combination of channel parameters is possible. In Fig. 5.12 we plot curves of the simulated channel parameters – on-axis density, matched spot size, and attenuation length – for a range of axicon angles, calculated at 0.5 ns intervals. The density of the hydrogen target was chosen in each case to create critical electron density after ionization. For axicon approach angles of 2.5°, 5° and 7.5°, these critical densities are $3.3 \times 10^{18}\,\mathrm{cm}^{-3}$, $1.3 \times 10^{19}\,\mathrm{cm}^{-3}$ and $3.0 \times 10^{19}\,\mathrm{cm}^{-3}$, respectively. It will always be possible to reduce the electron density below this point by dropping the target density, but not to increase it. The required energy in the laser pulse remains constant for a given channel length and threshold intensity, and was always 15 mJ in a bi-Gaussian beam with a FWHM diameter of 2 inches and a FWHM duration of 40 fs.

The different axicon angles produce different characteristic curves, with lower axicon angles giving lower density channels with slightly higher matched spot sizes. For matched spots smaller than half the plasma period, it is always possible to choose an axicon angle which will produce a given combination of matched spot size and on-axis density.

These channels also have long attenuation lengths, on the scale of tens of centimetres, and are increasingly low loss as the channels expand. Close to the limit $w_0 \approx \lambda_p/2$, the attenuation lengths reach the metre scale, and are equal to the laser wakefield dephasing length at $n_e = 10^{17}\,\mathrm{cm}^{-3}$, or significantly greater than the dephasing length at the higher densities. These channels would therefore be well suited to driving laser pulses for LWFA experiments even as they move to lower densities and longer stage lengths than previously accomplished; the design points for BELLA 10 GeV and the EuPRAXIA project are for on-axis densities of 1×10^{17} cm^{-3} and 1.5×10^{17} cm^{-3} respectively and acceleration stage lengths of 97 cm and 30 cm respectively [98], [99].

For a given set of channel parameters, Fig. 5.12 can be used to estimate the axicon angle which will best provide that channel. For instance, to guide a $w_0 \approx 30\,\mu$m spot at an on-axis density of $n_e(0) \approx 2 \times 10^{17}\,\mathrm{cm}^{-3}$, it is most appropriate to use an axicon with an approach angle of $\alpha = 5^\circ$ and an initial density of $n_0 \approx 1 \times 10^{19}\,\mathrm{cm}^{-3}$. At this point the attenuation length is approximately one metre and hence losses will be negligible. To create a matched spot size of $w_m = 10\,\mu$m at an on-axis density of $10^{18}\,\mathrm{cm}^{-3}$, however, an axicon angle of greater than 7.5° would be required.

5.6 Conclusions

In this chapter we have described a new technique to achieve long and low density plasma waveguides, which could operate at high repetition rates while requiring little laser energy. Optical field ionization along the focus of an axicon lens can create hot plasma columns with lengths on the scale of 10s of centimetres, and the subsequent evolution of these columns leads to channels with on-axis densities as low as $n_e(0) = 10^{17}\,\mathrm{cm}^{-3}$. In order to study these plasma structures we have studied the electron energies resulting from OFI and compared these to those expected from the simple barrier suppression ionization model. The BSI and Simpleman models describe much of the behaviour of

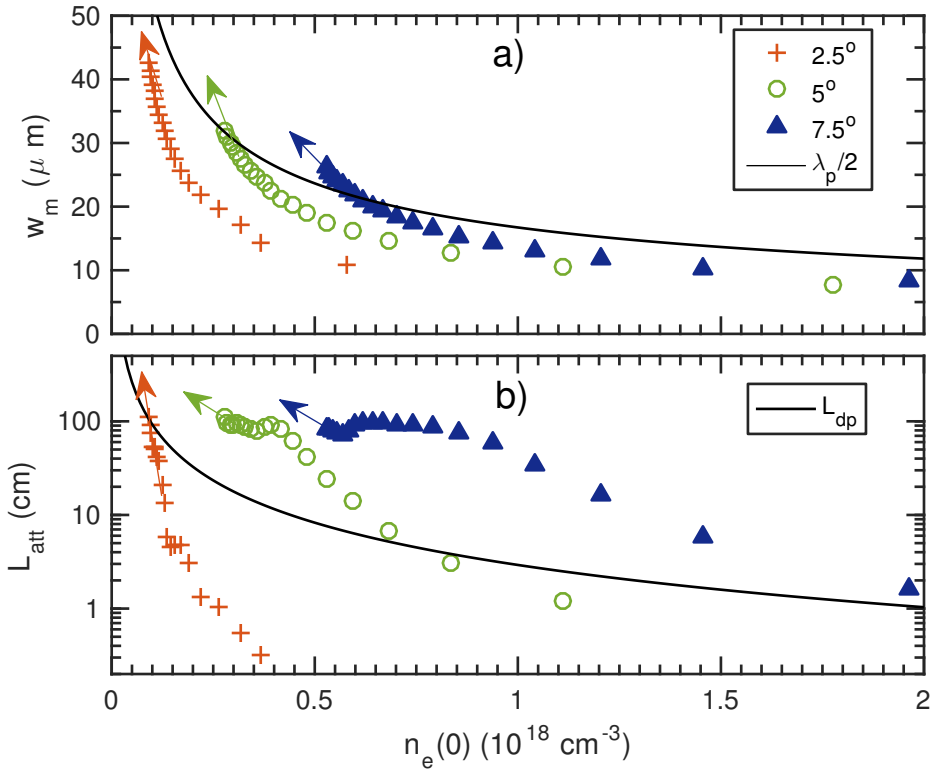


Figure 5.12: a) Calculated matched spot size of the plasma channels against the on-axis density for axicon approach angles of $\alpha = 2.5\,^\circ$ (red crosses), $5\,^\circ$ (green open circles) and $7.5\,^\circ$ (blue closed triangles), plotted at 0.5 ns intervals. As time passes the on-axis density drops and the matched spot size increases. These are compared to half the plasma wavelength (black line). b) The attenuation length of the guided modes against the on-axis density, for the same three axicon approach angles. Each marker represents the same timestep as the equivalent marker at the same density in the plot above. The laser wakefield dephasing length (black line) is overlaid, assuming an 800 nm drive pulse.

the electron heating, but simulations in EPOCH reveal that when the ionization energy is high, tunnelling ionization plays a substantial role and results in much broader peaks in the electron energy distribution than predicted by BSI. For low ionization energies, such as in hydrogen, the effect is largely insensitive to the laser pulse profile, depending only on the ionization energy. The polarization state of the ionizing laser pulse also plays a strong role, with circularly polarised laser pulses creating much higher energy electrons than pulses with linear polarisation.

The initial plasmas produced by OFI are far from thermal equilibrium, but at energies on the scale of $E_K \sim 10\,\mathrm{eV}$, for hydrogen ionized by OFI, the collision frequency is sufficiently fast to thermalise electrons to a single thermal distribution on the ps timescale, while leaving the ions cold. Despite using this to make a fluid approximation, finding a simulation code that could capture the physics of this situation was very challenging. Ultimately we used one code, Bessel-ATI, to model the formation of the plasma column and another code, HELIOS, to capture its subsequent expansion. This resulted in a model for the expansion of the plasma column which agreed well with the Sedov-Taylor theory of shock expansion.

In order to set the initial conditions more carefully, I described the structure of an axicon focus in more detail and derived some expressions to establish the axicon parameters which should be used to create a plasma column of a given length and radius, while using minimum input energy to transfer as much energy as possible to the plasma. For top-hat beam profiles, optimum axicons should be chosen such that the plasma column begins halfway along the focus and has a radius approximately half that of the first node in the Bessel beam profile. To create a column of hydrogen plasma with this optimised axicon focus requires only around 1 mJ per cm length of column. This increases proportionately with the threshold ionization intensity of the species in question.

Several other constraints on the axicon approach angle were derived for an experimental context. Firstly, counter-propagating geometries are restricted by the time

difference encountered by the axicon beam compared to a guided beam. This in turn normally implies that the guided beam must be able to fit inside the numerical aperture of the channel forming beam. The guided beam must also have a sufficiently small spot and the plasma channel must be able to expand large enough to match and guide this spot. This set of constraints interlock, and yield three safety factors which determine what range of axicon angles is necessary for a given matched spot size, or indeed whether the experiment will be possible at all.

Finally, I have used all of this work to run simulations in a range of codes that describe OFI electron heating along an axicon focus, channel formation, and hydrodynamic expansion, and I have found the lowest order guided modes for these channels. Channel formation and evolution is described well by the theoretical fluid models, with plasma channels evolving on the timescale of 10 ns, and forming waveguides with on-axis densities of $n_e(0) = (10^{17}\text{–}10^{18})\,\mathrm{cm}^{-3}$ and matched spot sizes of $w_m = (15\text{–}40)\,\mu\mathrm{m}$. Changing the axicon angle and the time at which the guided beam arrives allows control of the on-axis density and matched spot size, whereas changing the axicon radius controls the length of the channel. Together, this demonstrates the potential of these OFI axicon waveguides to guide LWFA drivers at low densities over long distances with a great degree of control. These waveguides could operate at multi-kHz repetition rates without suffering damage.

Chapter 6

Conclusions

6.1 Original Aims

The primary aim of this thesis has been to consider techniques to allow laser wakefield accelerators to operate at higher repetition rates and at lower densities, which would allow a higher energy gain in a single acceleration stage. These two developments would substantially improve the utility of LWFA. Higher repetition rates support imaging applications in particular; currently, competing facilities such as synchrotron light sources can operate at kHz repetition rates, and accessing this with LWFA would allow significant improvement in signal to noise ratios, as well as making more practical techniques such as tomography, which require many images to be taken. High energy gain, on the other hand, is important not only for building particle colliders, but also to access new interaction regimes, to increase the production of secondary radiation from solid targets, and to explore non-classical processes such as inverse Compton scattering.

The work in this thesis has tried to make progress towards these two objectives by working on three fronts: resonantly driving plasma waves with multiple laser pulses in order to make use of laser systems with higher average power, but lower peak power; guiding these pulses over long distances, at low densities, and at high repetition rates; and diagnosing the resulting wakefields rapidly and quantitatively. In order to drive laser wakefield accelerators with trains of multiple pulses, I studied the accuracy required to tune the pulse spacing to the resonance in the plasma. I compared the limitations from random and systematic errors in pulse train spacing to a basic linear theory and used this linear theory to explore a way in which those limitations might be overcome: auto-resonance. I then discussed a diagnostic which is both fast and extremely sensitive to the small phase shifts produced at low densities: temporally encoded spectral shifting (TESS). I have generalised this new diagnostic to pulses of arbitrary spectral profile, and suggested using it to measure quasi-linear and non-linear wakefields which are not sinusoidal. A recent experiment at the Astra Gemini TA2 laser facility allowed my work on both of these to be tested, with the TESS diagnostic being used to measure

plasma waves which were resonantly excited by a train of laser pulses, though imperfectly matched to the plasma wave. I described the experiment in brief, and my analysis techniques in detail, and have demonstrated the efficacy of my work on driving and diagnosing linear plasma waves.

Finally, I worked on guiding laser pulses at high repetition rates and low densities. I have here described a new technique which builds on previous work on the formation of hydrodynamic plasma waveguides. My simulations show that optical field ionization along the focus of an axicon lens can lead to plasma channels which meet all of the requirements for future laser wakefield accelerators. These structures achieve on-axis densities lower than achieved to date, over lengths longer than possible with current means. Furthermore, while some of the existing techniques are at high risk of damage by the driving laser pulses, these waveguides would be largely immune to this threat even when operating at multi-kHz repetition rates.

Through this chapter I will summarise the work described in the rest of the thesis, bringing my work on these three fronts to a close and suggesting future directions research in these areas can take.

6.2 Driving with Multiple Resonant Pulses

6.2.1 Errors in Achieving Resonance

To study the accuracy required in a train of laser pulses in order to resonantly excite a plasma wave, I took two approaches: first, I described a simple linear model, treating the total plasma wave as the sum of the plasma waves from each individual pulse, and then I compared this to simulations with the Particle-In-Cell (PIC) code EPOCH. To begin with, I will summarise the results of the linear model.

For random errors, where the pulse position 'jitters' around the correct spacing, the wakefield has contributions from both coherent and incoherent addition. If the

position is described by a probability distribution $P(\delta\tau_n = \nu) = f_\sigma(\nu)$, where σ describes the standard deviation away from ideal pulse positions, we can estimate the wakefield amplitude in these two cases as:

$$A_{\mathrm{rms}} \approx \begin{cases} N\left|\widetilde{f}_\sigma\left(\omega_{\mathrm{p}}\right)\right| & \text{for } N \gg \left|\widetilde{f}_\sigma\left(\omega_{\mathrm{p}}\right)\right|^{-1} \\ \sqrt{N} & \text{for } \left|\widetilde{f}_\sigma\left(\omega_{\mathrm{p}}\right)\right|^{-1} \gg N \end{cases}, \tag{6.1}$$

where $\widetilde{f}_\sigma\left(\omega\right)$ is the Fourier transform of the original probability distribution $f_\sigma(\nu)$.

For sufficiently large N and small σ, the coherent contribution dominates and the amplitude rises linearly with time, regardless of the scale of the errors. Pulse trains are not limited in length and can still drive a wakefield effectively so long as $\omega_{\mathrm{p}}\sigma \lesssim 1$. However, if the scale of the random errors is too large, on the scale of $\sigma \sim \tau_{\mathrm{p}}$, such that $f_\sigma(\nu)$ is broad and $\widetilde{f}_\sigma\left(\omega\right)$ narrow, the coherent part vanishes, and the wakefield is driven extremely inefficiently by the incoherent contribution.

For systematic errors, where the pulse train is imperfectly 'tuned' to the plasma resonance, the error in pulse timing is described by $\delta\tau_n = n\delta\tau$, and the amplitude and phase of the wakefield can be described by:

$$A = \left|\frac{\sin\left(\frac{1}{2}N\omega_{\mathrm{p}}\delta\tau\right)}{\sin\left(\frac{1}{2}\omega_{\mathrm{p}}\delta\tau\right)}\right| \tag{6.2}$$

$$\tan\phi = \tan\left(\frac{1}{2}N\omega_{\mathrm{p}}\delta\tau\right), \tag{6.3}$$

which implies that the wakefield amplitude is periodic over time, with the period inversely proportional to the error in pulse train spacing. Pulse trains longer than $N_{\mathrm{max}} = \tau_{\mathrm{p}}/2\delta\tau$ are redundant or even counter-productive, with the maximum wakefield amplitude limited to $A(N_{\mathrm{max}}) = 1/\left|\sin\left(\frac{1}{2}\omega_{\mathrm{p}}\delta\tau\right)\right|$.

Both of these results reflect that the wakefield amplitude in the linear regime can be described similarly to a diffraction pattern, replacing the interference of electro-magnetic waves with longer wavelength electro-static plasma waves. The effect of systematic

errors is analogous to an angular shift away from a diffraction peak, with the number of pulses equivalent to the number of grooves in the grating. The expression for stability to random errors, on the other hand, corresponds to the intensity of light diffracting from a slit of finite size, with the size of the slit standing in for the width of the distribution of errors in pulse spacing.

These linear models were tested with PIC simulations, and found to accurately describe the situation when the plasma wave remains close to linear. Figure 6.1 shows a summary of the results, with simulations agreeing closely with the linear theory. Random errors do not place any limit on the length of a pulse train, but do reduce its effectiveness, with errors larger than $\sigma/\tau_{\mathrm{p}} \approx 10\%$ dramatically reducing the wakefield amplitude. High systematic errors, however, reduce the length of a pulse train that can be used and reduce the maximum wakefield amplitude that can be generated to $A(N_{\mathrm{max}}) \approx 0.64 N_{\mathrm{max}}$.

In general, then, errors in tuning are more serious, but are also easier to account for, by simply changing the plasma density to match the pulse train spacing. For a train comprising of N pulses, the plasma density must be controlled to an accuracy of $1/N$. Only if the plasma density cannot be measured or controlled to this accuracy, or if the scale of random errors is comparable to the plasma period, are errors in pulse train spacing serious.

6.2.2 Auto-Resonance

I have explored a technique which may be able to overcome some of the limitations on pulse train length and maximum wakefield amplitude for cases where the tuning error cannot be eliminated. By increasing the spacing between laser pulses in the train over time, the laser can explore a range of different resonances and drive the wakefield to higher amplitudes than possible with a uniform but mis-tuned train. Using my linear model to describe a pulse train where the spacing between pulses increases over time, I

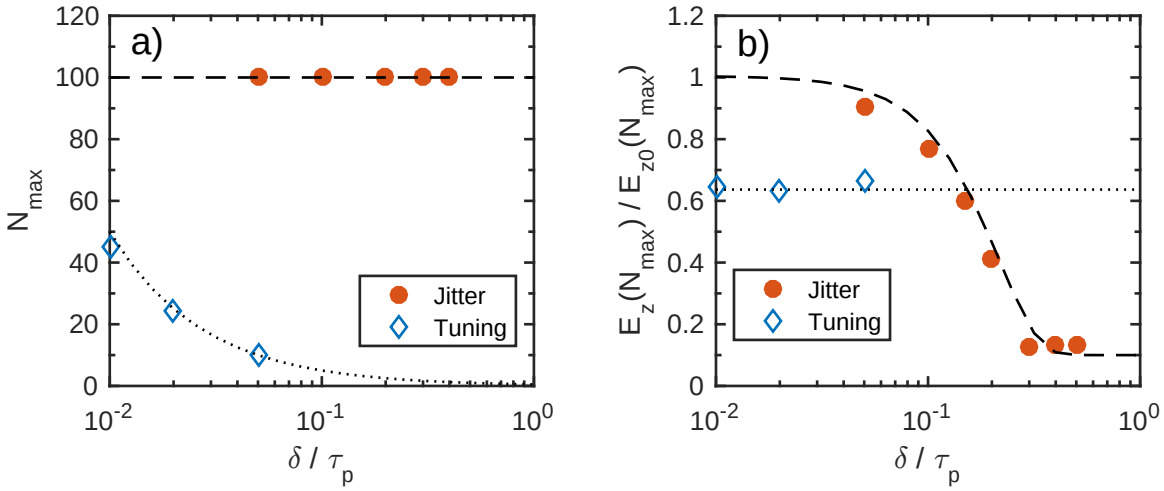


Figure 6.1: a) The number of pulses $N_{\max}$ which produced the maximum wakefield amplitude in simulations, plotted against the error in pulse train positions on a log scale. Jitter (filled red circles) corresponds to random errors in pulse timing from a normal distribution with a standard deviation $\sigma = \delta$. Tuning (open blue diamonds) corresponds to systematic errors with successive pulses separated by $\Delta\tau = \tau_p + \delta$. Theoretical predictions are also marked for jitter (dashed) and tuning (dotted). b) The wakefield amplitude in simulations after $N_{\max}$ pulses, compared to the wakefield amplitude after $N_{\max}$ pulses produced by an ideal pulse train with $\delta = 0$. Again, the results from PIC simulations are compared with the theoretical predictions for random errors (filled red circles, dashed line) and systematic errors (open blue diamonds, dotted line).

determined an estimate for the resulting wakefield amplitude of:

$$A \approx \frac{1}{c}|\mathrm{Eu}\,[v(N)] - \mathrm{Eu}\,[v(0)]|, \tag{6.4}$$

where $\mathrm{Eu}(v) = C(v) + iS(v)$ describes the Euler spiral, and v is a parameter which increases linearly with the number of pulses as $v(N) = cN + v(0)^*$. This, while a complicated expression, corresponds to the Fresnel expression for the amplitude of a diffracted beam in the near-field regime, extending the analogy with electro-magnetic waves. The amplitude of the wakefield can be estimated by measuring the length of a line across the Euler spiral.

Again, I can compare this with the results from PIC simulations, shown in Fig. 6.2, demonstrating excellent agreement. The maximum wakefield amplitude is obtained at a non-zero chirp rate of $\alpha \approx 10^{-5}$, demonstrating that in this case, when the initial pulse spacing in the train was just 1% away from the plasma period, a chirped pulse train performs better than a pulse train which is uniform but mis-tuned. This optimum chirp rate is close to the critical chirp rate of auto-resonance, but curiously, the linear model does not depend on the order in which pulses arrive. This implies that the phenomenon measured here is not strictly auto-resonance; exploring true auto-resonance, where the frequency of the plasma wave becomes locked to the frequency of the driver, requires further work.

6.2.3 Future Work

Future work could focus on moving towards the higher amplitude, quasi-linear plasma waves, where auto-resonance becomes more relevant. As electrons move relativistically, the shape of the plasma wave becomes more saw-toothed, and timing each pulse such that the wakefields constructively superimpose becomes increasingly difficult. Moving towards higher amplitude wakefields would require either higher intensity pulses, which

*$C(v)$, $S(v)$, c and $v(0)$ are described in detail in Chapter 2.

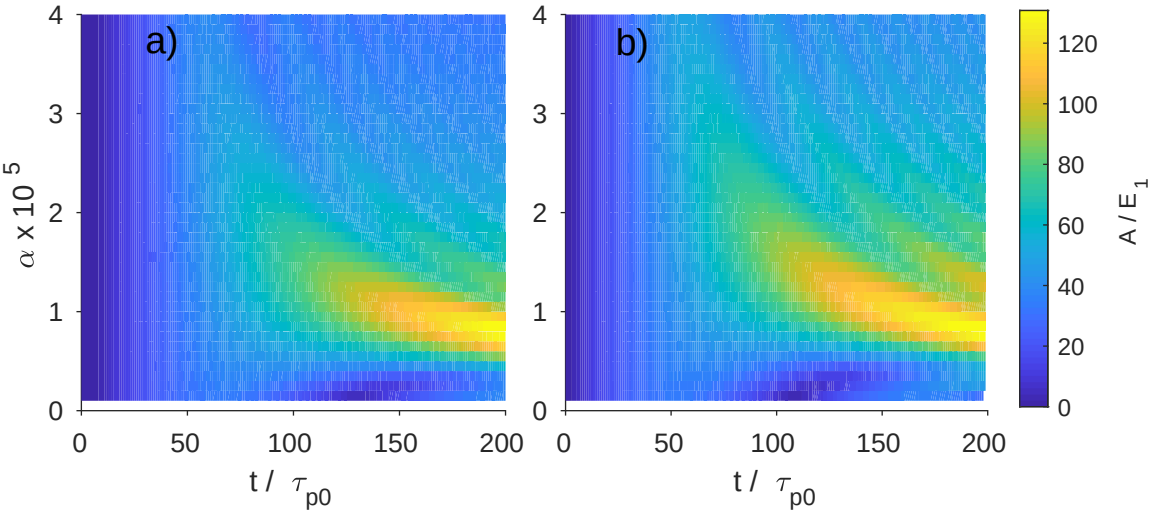


Figure 6.2: Wakefield amplitude over time against chirp rate predicted by a) the linear theory and b) PIC simulations. The colour describes the wakefield amplitude relative to the amplitude resulting from a single pulse.

would cause non-linear effects in the plasma if $a_0 \sim 1$, or longer pulse trains. In both these cases, chirp and auto-resonance become more important, with higher intensity pulses increasing the critical chirp rate, while longer pulse trains require more accurate tuning with the plasma period.

Longer pulse trains are also prone to other kinds of instability, with ion motion becoming relevant on long timescales. Throughout this work I ignored this effect by unphysically fixing ions to be stationary in PIC simulations, but in practice this could place strong limits on the length of pulse trains that could be used. Normally ion motion is relevant for beam driven acceleration using long particle bunches, and any future work should compare the results for laser driven plasma accelerators. The creation of a density structure in the wake of the pulse train also leads to problems for pulses at the back of the train. These pulses can experience strong spectral shifts, such as from Raman scattering, and can experience a hosing instability, where laser pulses later in a pulse train are deflected off-axis by the changes in refractive index. Density profiles which vary longitudinally, where there is no single resonant frequency, would prove a significant challenge to multiple pulse LWFA, but with further study auto-resonance

may be able to overcome this.

6.3 Diagnosing with Generalized TESS

6.3.1 Generalizing TESS

The next subject studied by this thesis was generalizing the Temporally Encoded Spectral Shifting (TESS) analysis technique to non-Gaussian pulses and to non-linear plasma waves. This was mainly achieved with analytic work, building some basic numerical models of TESS measurements to demonstrate the results. Although TESS is closely related to Frequency Domain Holography (FDH), which has been established for several years, TESS is itself a relatively new technique and this work comprises the first development upon it. I therefore first spent some time describing TESS, the differences between TESS and Frequency Domain Holography (FDH), and how TESS can be used in practice. This included the phase acquired by a chirped laser pulse over time, the effect of different group velocities of probe and pump laser pulses of different frequencies, and how to measure the group delay dispersion (GDD) of probe pulses. I then moved on to extending the TESS analysis, first by looking at probe and reference laser pulses which have arbitrary spectral profiles which need not be identical.

Arbitrary Spectral Profiles

Whereas the initial description of TESS assumed probe and reference laser pulses that were both identical and Gaussian, I showed that the peak height of TESS satellites depends strongly on the overlap factor between the original spectrum and the copies of it created by modulation by the plasma wave. This overlap factor is calculated by the auto-correlation of the two spectra, described by:

$$\mathcal{F}(\Omega) \equiv \frac{\int_{-\infty}^{\infty} |\mathcal{E}_{\mathrm{pr}}(\omega-\Omega)|\,|\mathcal{E}_{\mathrm{r}}(\omega)|\,\mathrm{d}\omega}{\int_{-\infty}^{\infty} |\mathcal{E}_{\mathrm{pr}}(\omega)|\,|\mathcal{E}_{\mathrm{r}}(\omega)|\,\mathrm{d}\omega} = \frac{\int_{-\infty}^{\infty} \sqrt{I_{\mathrm{pr}}(\omega-\Omega)I_{\mathrm{r}}(\omega)}\mathrm{d}\omega}{\int_{-\infty}^{\infty} \sqrt{I_{\mathrm{pr}}(\omega)I_{\mathrm{r}}(\omega)}\mathrm{d}\omega}, \tag{6.5}$$

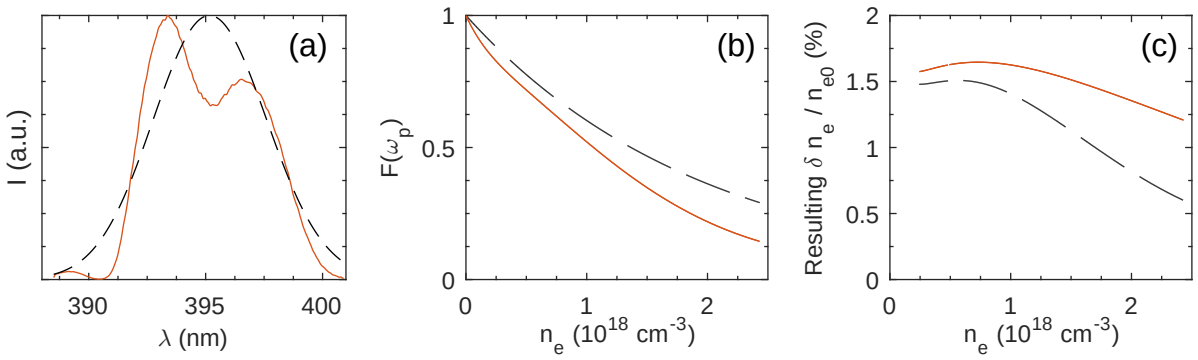


Figure 6.3: a) Comparison of a measured spectrum (red) with the Gaussian approximation to it (black dashed). b) The resulting overlap factor $\mathcal{F}(\omega_\mathrm{p})$ plotted against plasma density for the true spectrum and for the Gaussian approximation. c) The inferred wakefield amplitude using the true overlap factor and using the Gaussian approximation.

where $I_x(\omega)$ and $\mathcal{E}_x$ are the intensity and electric field respectively of a laser pulse in the spectral domain.

This overlap factor is significantly smaller for asymmetric and highly structured spectra, and this leads to significant under-estimates of the amplitude of the plasma wave when using the Gaussian approximation. Fig. 6.3 shows the effect on the overlap factor and the inferred plasma wave amplitude of assuming a Gaussian pulse, using a spectrum and plasma wave measured in the recent experiment described in Chapter 4. In order to accurately measure plasma waves with TESS it is vital to use an accurate calculation of the overlap factor, with the wakefield amplitude related to the overlap factor as:

$$\frac{\delta n_e}{n_{e0}} \approx \frac{2\omega_0 c}{\omega_\mathrm{p}^2 L} J_k^{-1}\left[\frac{r_k}{\mathcal{F}(k\omega_\mathrm{p})}\right], \tag{6.6}$$

where $r_k = |s(t_k)/s(t_0)|$ is the ratio of peak amplitudes between the k-th satellite peak and the main interference peak at $t_0 = \Delta t$.

I also showed that when probe and reference pulses have different GDDs, the peaks in the TESS spectrum will now lie at points $t_k = \Delta t + k\psi^{(2)}_\mathrm{eff}$. It is important to use this

effective GDD, rather than the measured GDD of just the probe pulse. The effective GDD is described by:

$$\psi^{(2)}_{\text{eff}} \approx \frac{\delta\omega^2_{\text{pr}}\psi^{(2)}_{\text{pr}} + \delta\omega^2_{\text{r}}\psi^{(2)}_{\text{r}}}{\delta\omega^2_{\text{pr}} + \delta\omega^2_{\text{r}}}, \tag{6.7}$$

where $\delta\omega_{\text{pr}}$ is the spectral bandwidth of the probe pulse and $\delta\omega_{\text{r}}$ that of the reference pulse.

Non-linear Wakefields

Next, I worked on extending TESS to quasi-linear and non-linear wakefields. TESS was originally developed for use on plasma waves which were purely sinusoidal; these modulate the probe pulse with a single frequency and result in a TESS signal which can be described fairly simply, with multiple equally spaced satellite peaks occurring either side of the original interference peak. When a plasma wave is no longer sinusoidal, but remains periodic, the probe pulse is instead modulated by multiple frequencies, each an integer multiple of the plasma frequency. Fig. 6.4a shows some example wakefields, with the wakefield amplitude described by the maximum electron velocity through $\beta_m = v_{e,\max}/c$, and Fig. 6.4b shows the frequency components of the phase change decomposed into harmonics such that $\phi(\zeta) = \sum_n \phi_n \sin(n\omega_{\text{p}}\zeta + \theta_n)$. As the wake amplitude increases and electrons become relativistic, the plasma wave changes from being approximately linear, dominated by the ϕ_1 component, to being sharply peaked with many harmonics playing a significant role.

The TESS signal remains comprised of equally spaced satellite peaks, but now each peak contains contributions from many frequency components. I showed that if we can approximate the plasma wave by limiting the number of frequency components to N, it is possible to describe these contributions by finding the points which lie both on an N dimensional grid of integer values of k_n, and on a $N-1$ dimensional plane described by $\sum_n nk_n = \kappa$. Each of these points contribute to the peak lying at $t_\kappa = \Delta t + \kappa\psi^{(2)}\omega_{\text{p}}$,

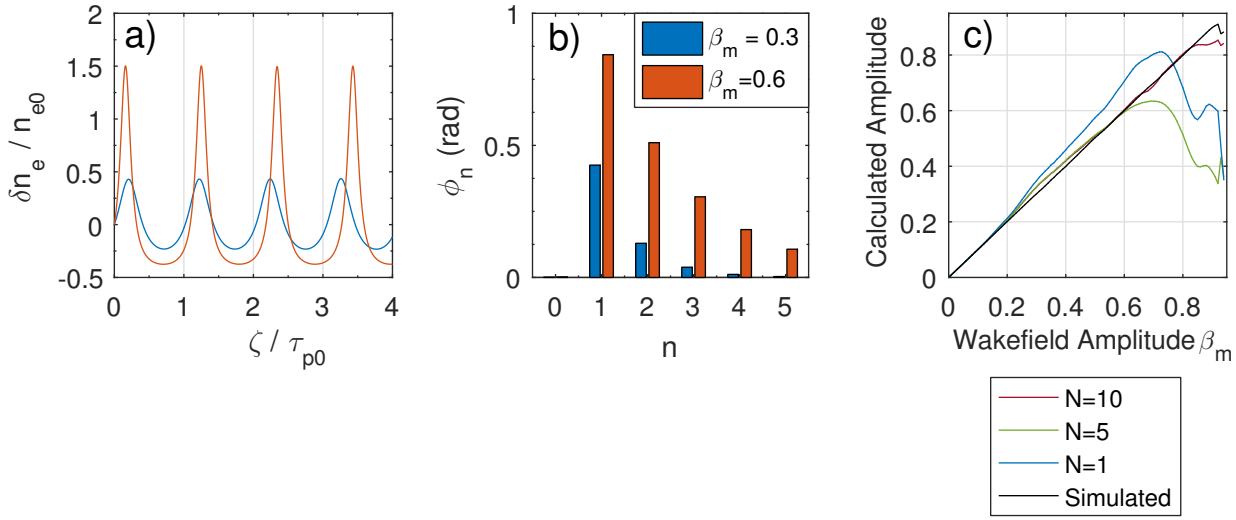


Figure 6.4: a) Simulated relative density of plasma waves in the quasi-linear regime, with electrons moving at $v_{e,\max} = 0.3\,c$ (blue) and $0.6\,c$ (orange). b) Harmonic amplitudes ϕ_n of the phase change experienced by a 400 nm probe co-propagating with these plasma waves for 1 mm at $n_e = 10^{18}\,\mathrm{cm}^{-3}$. c) The wakefield amplitude retrieved using a lookup table calculated for $N = 10$ (red), $N = 5$ (green) and $N = 1$ (blue) harmonics, compared to the value taken from simulations (black).

with its amplitude described by:

$$r_\kappa \equiv \left|\frac{s(t_\kappa)}{s(t_0)}\right| = \left|\frac{Z_\kappa\left(\{\phi_n\},\{\theta_n\}\right)}{Z_0\left(\{\phi_n\},\{\theta_n\}\right)}\right| \mathcal{F}\left(\kappa\omega_\mathrm{p}\right) \tag{6.8}$$

$$\text{where} \quad Z_\kappa\left(\{\phi_n\},\{\theta_n\}\right) = \sum_{\mathbf{k}\in S_\kappa} J_{k_1}(\phi_1)\ldots J_{k_N}(\phi_N) e^{i\sum_{n=1}^{N} k_n\theta_n} \tag{6.9}$$

$$\text{for a set} \quad S_\kappa = \left\{(k_1,\ldots,k_N) : \sum_{n=1}^{N} n k_n = \kappa\right\} \tag{6.10}$$

This makes reconstructing a quasi-linear plasma wave from the TESS spectrum difficult, but progress can be made by calculating a lookup table using a given number of harmonics N, each limited up to a maximum expansion order $|k_n| \leq K$. This technique was tested for wakefields with amplitudes ranging from linear to highly relativistic. This demonstrated that increasing the number of harmonics increases the accuracy of the reconstruction at the cost of a longer computation time required when making the lookup table. For $N = 10$ harmonics, the lookup table had to include the contributions of around 3 million components for each satellite peak, and took several minutes on a

desktop computer. This created a lookup table that was accurate for those particular experimental parameters (a 400 nm probe co-propagating with a plasma wave at a density of $n_{e0} = 10^{18}\,\mathrm{cm}^{-3}$ for 1 mm). Once the lookup table is calculated, however, the analysis is rapid. Wakefield amplitudes were successfully retrieved up to $\beta_m = 0.8$, demonstrating the promise of this extension to the TESS technique. If a TESS analysis can rapidly retrieve the amplitude and frequency of these relativistic wakefields, it can be applied to the next generation of high gradient laser wakefield accelerators at low densities and at high repetition rates.

I also attempted applying the TESS analysis to a wakefield in the bubble regime, simulated in EPOCH, as shown in Fig. 6.5. The TESS signal looks somewhat similar and the frequency of the wakefield can still be extracted along with estimates of the wakefield amplitude. By looking at the phase of the TESS peaks at different points in space, it is also possible to reconstruct the curvature of the wavefront and obtain a spatio-temporal map of the wakefield, as in FDH. On the other hand, the TESS analysis naturally averages the frequency components over both the interaction length and the duration of the probe pulse. When the wakefield is not periodic this makes it substantially more difficult to estimate the amplitude of the plasma wave from the peak of the TESS satellite peaks, and many satellite peaks need to be measured to accurately reconstruct the wakefield. Measuring satellite peaks from high frequency components is, however, limited by the spectral width of the probe pulse through the overlap factor $\mathcal{F}(\omega)$. It is therefore difficult to reconstruct any temporal features smaller than the bandwidth limited duration of the probe pulse. With non-periodic wakefields frequencies are no longer harmonics of the plasma frequency and TESS satellite peaks increasingly overlap, until it is necessary to perform the full FDH reconstruction of the wakefield.

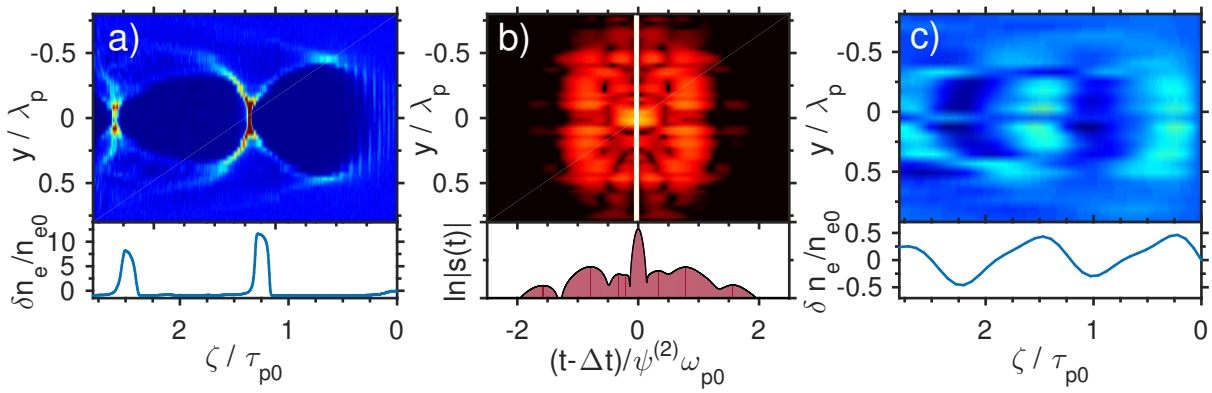


Figure 6.5: a) Simulated electron density maps of a wakefield in the strongly non-linear bubble regime. b) Maps of the calculated TESS signal resulting from this density profiles, plotted in both space and in time against the expected peak positions. The magnitude of the TESS signal is plotted on a logarithmic scale. c) A reconstruction of the electron density map of the wakefield using the TESS analysis, using $N = 4$ peaks from the TESS signal. In all plots the profile on-axis is shown below.

6.3.2 Future Work

With non-periodic wakefields currently inaccessible to TESS analysis, it may be worth developing the work on non-linear but periodic wakefields. The lookup tables were calculated for one dimensional cold plasma waves, described by Akhiezer and Polovin [90], but in practice plasma waves take different forms, especially when the driving laser pulse has a spot size smaller than the plasma wavelength, or the plasma waves are driven inside a density structure such as a waveguide. It would be helpful to look at the effectiveness of the TESS reconstruction in these cases, where the spatial dependence of the wakefield is important.

As more information becomes encoded in the TESS signal, it is possible to compare it to other diagnostics such as Thomson scattering. In this diagnostic, as with TESS, the location of a peak encodes information about the frequency of a plasma wave, and from the amplitude of the peak the amplitude of the plasma wave can be estimated. In addition, the width of a Thomson scattering peak contains information about the temperature of the plasma, and in a similar way the shape of TESS satellites may be altered by the form of the plasma wave. If this could be studied in more detail it may

be possible to infer more information about the plasma.

6.4 A Resonant Excitation Experiment

6.4.1 Discussion

An experiment performed by the group in 2015 allowed us to put together some of the work from the previous two chapters. Plasma wakefields were resonantly excited using a train of laser pulses and measured using Temporally Encoded Spectral Shifting (TESS). In the first part of this chapter I discussed the experimental set up in brief, and then described in detail how I conducted the data analysis. I primarily used the TESS analysis after generalising it to calculate the overlap factor for the measured, non-Gaussian, probe and reference spectra. I explained each step of my analysis, which successfully removed much of the background to extract extremely weak signals. The $k = 1$ satellite peaks we measured were smaller than one hundredth the amplitude of the main $k = 0$ interference peak, corresponding to wakefield amplitudes on the scale of $\delta n_e/n_{e0} \approx 1\%$, demonstrating that TESS is extremely sensitive.

In addition, I described the Frequency Domain Holography (FDH) technique and the alterations I made to it in order to subtract our measured background data. I applied it to an example shot, extracting a plasma wakefield with an amplitude on the scale of $\delta n_e/n_{e0} \approx 1\%$, shown in Fig. 6.6 and comparing the plasma periods extracted from this FDH reconstruction to those calculated in the TESS analysis, showing good agreement between the methods and with the expected plasma frequency.

In the second part of the chapter, I showed the results from the resonant excitation of plasma waves, comparing the resonance curves with plasma density for a single pulse, a pair of pulses, and a train of approximately seven pulses, shown in Fig. 6.7. The variation of the wake amplitude with pressure shows clear differences between these three cases, and agrees well with the theoretical predictions from Chapter 2. This

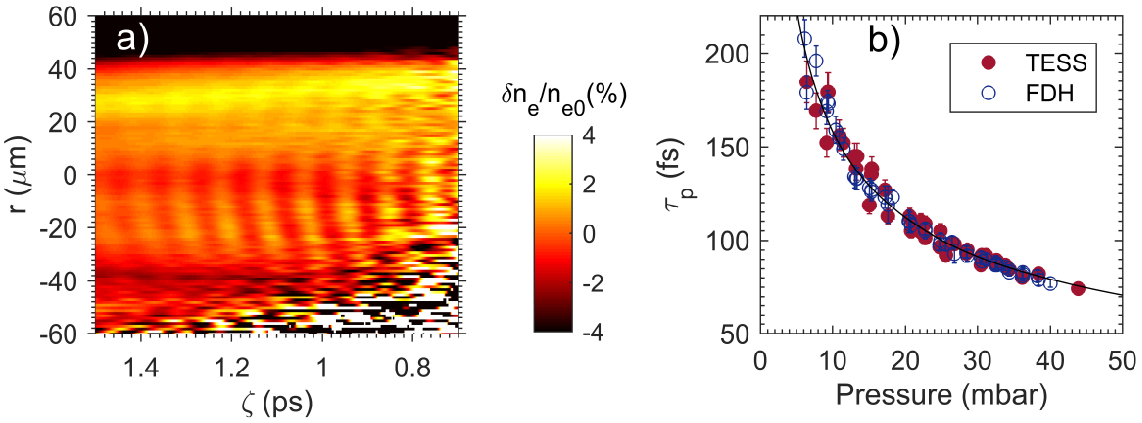


Figure 6.6: a) An FDH reconstruction of the relative electron density $\delta n_e/n_{e0}$ of a plasma wave against space and co-moving time ζ. b) The plasma periods against backing pressure extracted from FDH (open blue circles) and TESS analyses (closed red) compared with the expected values (black line).

confirms not only that multiple laser pulses can constructively drive plasma waves, but also that the TESS measurement of wakefield amplitudes and the theoretical predictions for these amplitudes are both accurate. In the process the experiment has demonstrated a step towards energy recovery, with a second laser pulse reducing the amplitude of the wakefield created by the first. This experiment has therefore verified some of the results described earlier in the thesis.

6.4.2 Future Work

This work, however, has been conducted with relatively few pulses and at extremely low wakefield amplitudes. Future resonant excitation experiments will attempt to work with much longer trains of pulses to excite a much higher amplitude wakefield, demonstrating the capacity of the technique to compete with existing single laser driver experiments.[†] This is challenging, requiring not only more laser energy but also, as we have described in Chapter 2, much better control of the plasma density, in order to maintain resonance over long pulse trains. If the plasma density cannot be matched to the pulse train spacing

[†]As I write this, the rest of my team is working hard on such an experiment at Astra Gemini, so I hope that by the time you read this it is already out of date!

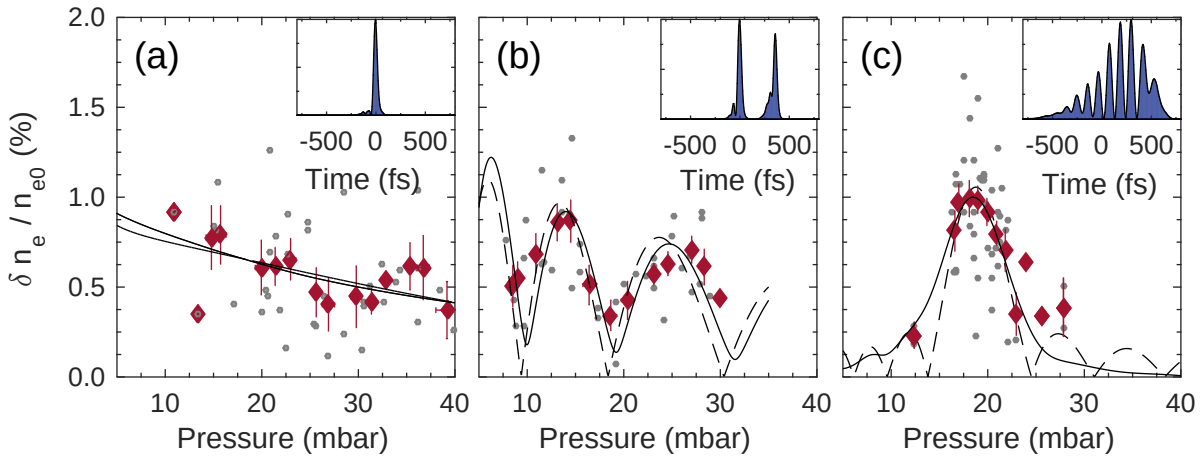


Figure 6.7: Relative wakefield amplitudes against gas cell pressure for a driving pulse train comprising a) a single pulse, b) two pulses separated by $\delta t = (420 \pm 20)$ fs and c) a pulse train with around seven pulses separated by $\delta t = (112 \pm 6)$ fs. Grey circles show the wakefield amplitudes for individual shots and red diamonds show these amplitudes binned by pressure. The insets show the inferred profile of the pump beam. The dashed lines show fits of eqn. (6.2), and the solid lines show the wake amplitudes calculated from the pulse trains shown in the figure insets. The plots are all shown on the same x and y-axis scales.

with sufficient accuracy, or if the plasma density varies over the length of the acceleration stage, the wake amplitude will be strongly limited. In this case, auto-resonance might be exploited to overcome this limitation, testing further the results from earlier in this thesis.

Measuring these higher amplitude wakefields will also be challenging if the plasma waves become non-linear. Chapter 3 has described a way in which even relativistic wakefields may be measured, however, and it would be interesting to test this work experimentally. The difficulties of measuring a plasma wave with a frequency which varies over time, as is produced in auto-resonance, may require further developments to the TESS technique. There are a number of other physical effects which will also play a role at this point, with ions moving on the timescale of tens of plasma periods. Measuring this effect will help in discovering more about the limitations upon multiple pulse plasma wave excitation. On the other hand, it may also be possible to directly

measure energy recovery in future experiments, by looking for frequency shifts in the spectrum of the second laser pulse. This would better demonstrate one of the potential advantages of the resonant excitation scheme.

6.5 Guiding with a Hydrodynamic Plasma Waveguide

The final area I studied was the creation of long and low density plasma waveguides which could operate at high repetition rates without damage. I described a proposed scheme to use optical field ionization (OFI) from circularly polarized light along the focus of an axicon to generate a column of hot plasma. The hydrodynamic expansion of this column forms a shock which propagates outwards leaving a cavity on-axis, suitable for guiding a subsequent laser pulse. My simulations suggest that with just mJ scale laser pulses this technique could generate channels which are 10s of centimeters long with on-axis densities on the order of $n_e(0) \sim 10^{17}\,\mathrm{cm}^{-3}$. These plasma channels do not rely on any external structure and the plasma itself recombines on the μs timescale without any degradation.

6.5.1 Plasma Column Formation

An axicon creates a long line focus described by the Bessel function J_0, as shown in the cartoon in Fig. 6.8a; in the ideal case this beam is diffraction free, and so can produce a much longer region of high intensity than is possible with focussing optics such as spherical lenses. In turn, this allows a very long column of plasma to be produced using little energy in the laser beam. I studied the optimum axicon parameters for creating these columns, and determined that to use the minimum possible laser power to create a plasma column of a given length L, the axicon radius should be chosen as $R = 2L\tan\alpha$, for an approach angle α. This approach angle limits both the radius of

the plasma column and the highest density at which the column can be formed, the critical density, and in order to maximise the energy transferred from the laser beam to a plasma column of a given radius r_0, this angle should be chosen as $\sin\alpha \approx 0.2\lambda_1/r_0$, where λ_1 is the wavelength of the ionizing beam. The laser power required is then proportional to the threshold intensity for ionization and the length of the channel, as $P_{\min} \approx 0.77\lambda_1 I_{\text{th}} L$. For an 800 nm laser with 40 fs FWHM duration, ionizing hydrogen at $I_{\text{th}} = 4.5 \times 10^{14}\,\text{Wcm}^{-2}$, this is just 1.18 mJ per centimeter length of the plasma column.

I then calculated the energy distributions of the electrons in this plasma column which result from optical field ionization, using simulations in the PIC code EPOCH. These demonstrate that electrons can be generated in hydrogen with energies around 10 eV, as shown in Fig. 6.8b; the mean energy created by circularly polarized light is $\langle E_K \rangle = 13.7\,\text{eV}$. These energy distributions depend strongly on both the ellipticity of the laser pulse and on the ionization energy of the electrons; elliptically and linearly polarized laser pulses produce lower energy electrons than circularly polarized pulses, whereas species with higher ionization energies such as nitrogen and helium lead to much higher electron energies, in the 1000s of eV. Ionization in these regimes is largely a threshold effect which is fairly well described by barrier suppression ionization, and the resulting electron energy distributions agree reasonably well with the expectations from the so-called 'Simpleman' model for OFI.

6.5.2 Channel Evolution

Once this plasma is created, collisions act to thermalize the electrons on the timescale of picoseconds, while the ions remain cold. The thermal distribution allows us to simulate these conditions using fluid codes, using as the initial conditions the output of the previous simulations. First, the Bessel-ATI fluid code developed by one of our collaborators, Howard Milchberg, was used to calculate the electron density profile

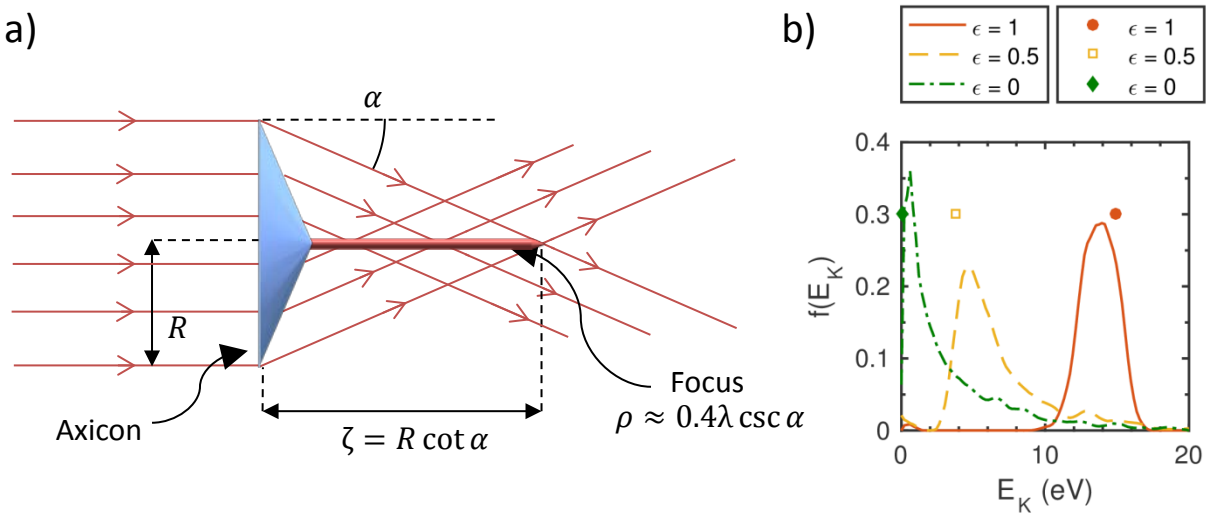


Figure 6.8: a) A cartoon of the line focus of an axicon, with dimensions marked. b) Simulated electron energy spectra resulting from OFI for hydrogen ionized to H^+ using a $2.5 \times 10^{14}\,\mathrm{Wcm}^{-2}$ intensity laser pulse. The polarisation state of the laser was varied between linear ($\epsilon = 0$ in green dash-dotted), elliptical ($\epsilon = 0.5$ in yellow dashed), and circular ($\epsilon = 1$ in red). Predictions from the Simpleman model using the simulated laser intensity are marked for circular polarisation (filled circles), elliptical (empty squares) and linear (filled diamonds).

created by the Bessel focus from an axicon with an approach angle of $\alpha = 2.5°$. A super-Gaussian fit to this profile was then fed into a HELIOS simulation of the column expansion, using $\langle E_K \rangle = 13.7\,\mathrm{eV}$ from the earlier simulations. This simulation, the results of which are shown in Fig. 6.9a, demonstrated the rapid formation of a shock, which propagates outwards to leave a channel on-axis on the timescale of nanoseconds. These channels expand to 10s of microns in radius, with on-axis densities falling to $n_e(0) \approx 10^{17}\,\mathrm{cm}^{-3}$.

A mode solver, also written by Howard Milchberg, could then be used to find the fundamental guided modes for these plasma channels by solving the Helmholtz equation in the simulated density profiles. As the plasma channel expands and the on-axis density decreases, the matched spot size of the guided mode increases, such that each set of initial conditions – axicon angle and initial plasma density – gives rise to a different contour of channel properties, including on-axis density, matched spot size, and attenuation length:

the length over which the intensity of a guided beam will decrease by a factor $1/e$.

The properties of these modes for several different axicon approach angles are shown in Fig. 6.9b, with each simulation starting at the critical density corresponding to the axicon approach angle. Over time the matched spot size increases as the on-axis density falls, with the spot size consistently close to the typical scale for driving a laser wakefield, $\lambda_p/2$. Similarly, the attenuation length of these channels increases, and is on the scale of 10s of centimeters, well above the LWFA dephasing length for all but the lowest on-axis densities. This combination of channel properties is ideally suited to meet all of the requirements for guiding laser pulses to drive laser wakefield accelerators.

I have therefore described and simulated from start to finish a scheme for guiding high intensity laser pulses, using hydrodynamic expansion of a hot plasma column formed by optical field ionization along the focus of an axicon. Using a combination of different simulation codes, I have demonstrated the potential of this technique for producing long and low density plasma channels, which would be resistant to damage and able to operate at kilohertz repetition rates, and are therefore ideally suited to the next generation of laser wakefield accelerators. Furthermore, I described how the axicon angle limits both the plasma density and the radius of the channel and used this to determine some of the optimum axicon angles to use for given requirements, as well as the energy required to create channels of a given length in a given species. However, this scheme is in the early stages and much work lies ahead.

6.5.3 Future Work

Firstly, I have performed simulations which assume radial symmetry, and ignore the longitudinal variation of the plasma channels. In contrast, I have based my simulations upon a top hat beam incident upon the simplest form of axicon, which creates plasma columns with a significant longitudinal dependence. I have also conducted the majority of the simulations assuming a hydrogen gas source for the plasma, as this requires the lowest

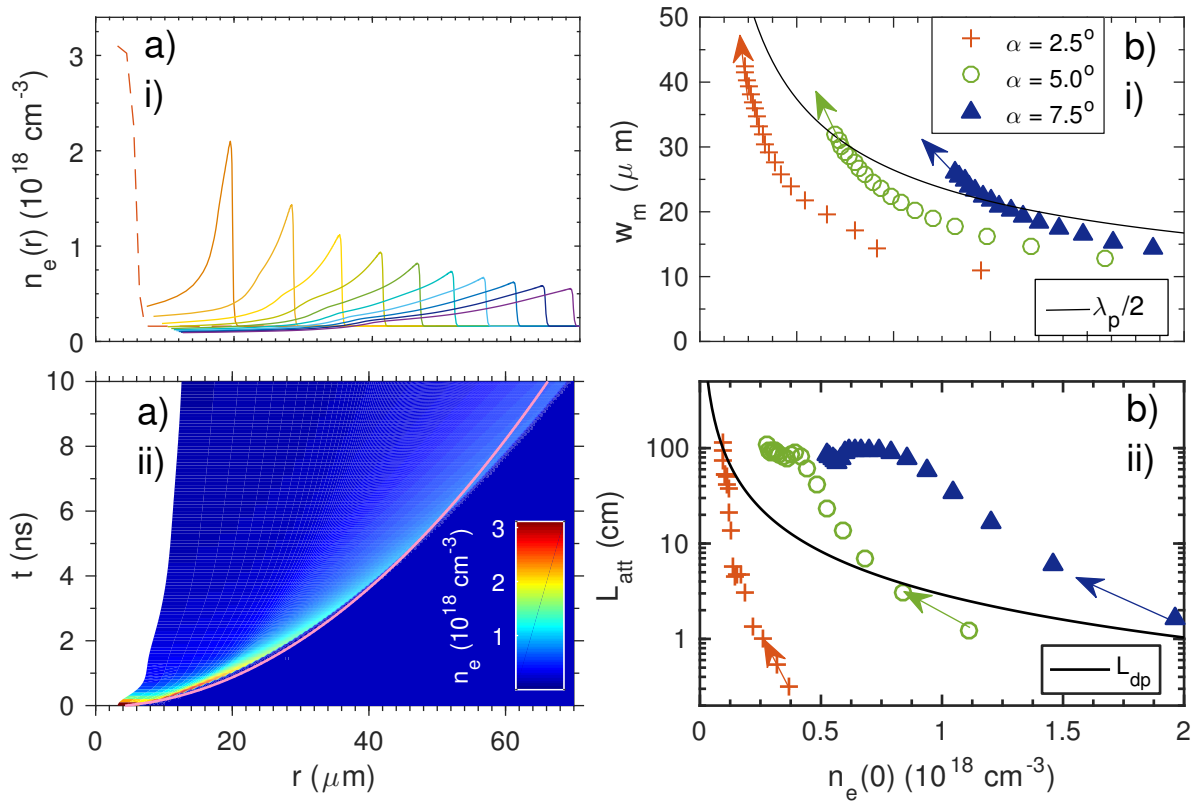


Figure 6.9: a) Simulated evolution of a hydrogen plasma column formed by an axicon with an approach angle of $\alpha = 2.5°$ shown by i) density profiles over time and ii) a color map of the density. The initial conditions (dashed density profile) are drawn from a previous fluid simulation in Bessel-ATI, with electrons at $\langle E_K \rangle = 13.7\,\text{eV}$. The predicted shock front position from a Sedov-Taylor model is overlaid on the color map (red line). b)i) Matched spot size and ii) 1/e attenuation length of the fundamental guided mode of the plasma channels, both plotted against the on-axis density of the channel. The direction of time is shown by arrows with starting conditions calculated for axicons with approach angles of $\alpha = 2.5°$ (red crosses), $\alpha = 5.0°$ (green open circles), and $\alpha = 2.5°$ (blue filled triangles). The matched spot size is compared to the typical spot size for LWFA, $\lambda_p/2$, whereas the 1/e attenuation length is compared to the dephasing length $L_{dp} = \lambda_p^3/2\lambda_0^2$ (black lines).

laser intensity to reach full ionization, but I have neglected the complicated ionization pathways of molecular hydrogen which can have a significantly higher ionization threshold. It would be helpful to consider this in more detail, and to study alternative gas species, with higher ionization energies. When stronger laser pulses are required, the energy required per unit length of channels will increase dramatically, and so too will the energy of the electrons which are produced. These electrons will take longer to reach a thermal distribution and the shock expansion rate will be faster; in these situations fluid codes may not adequately describe the physical situation. Instead, there may be a succession of shocks from the different 'classes' of electrons, with different energies. If the gas is not fully ionized by the initial channel forming pulse the subsequent guided beam may ionize fresh plasma, which would again cause multiple shocks.

These are interesting conditions, and the physics involved may allow us to exploit new dynamic processes. For instance, similar conditions are proposed for ionization injection into laser wakefields, in schemes such as Two Pulse Ionization Injection (2PII) [121] and Trojan Horse [122]. In disregarding longitudinal effects I have also neglected the start of the plasma column, where a hemi-spherical shock will propagate outwards. This too, may create interesting conditions for processes such as density downramp injection. To explore these situations, full Particle-In-Cell simulations could study the coupling of the guided pulse into the plasma waveguide and the propagation of the guided pulse once it is inside the channel. These would be substantially more informative than the numerical mode-solving conducted in this thesis.

Finally, this scheme must crucially be tested experimentally. This has begun with work using a spherical lens [72], and the first experiments with an axicon lens were recently completed and are being analyzed as I write this thesis.[‡] Clearly this is only the first step on the beginning of a long journey.

[‡]See Rob Shalloo's thesis for a description of these experiments and the results [118] – they're extremely exciting.

6.6 Summary

In the course of this thesis I have described three techniques which could help laser wakefield acceleration experiments to succeed at higher repetition rates:

- Firstly, I demonstrated that MP-LWFA enables the excitation of plasma waves with trains of high repetition rate but low energy laser pulses, but that it is essential to correctly tune the spacing between pulses. I showed that a chirped pulse train can, however, help to excite a wakefield even when the tuning is imperfect.

- Secondly, I showed that the TESS technique is effective for fast, sensitive and quantitative measurements of laser wakefields, even at low densities. I also described how to generalise the TESS technique, both for realistic laser pulse profiles and for measuring quasi-linear plasma waves.

- Thirdly, I outlined a proposed scheme to produce long and low density plasma waveguides which would be capable of high repetition rate operation without becoming damaged by the laser wakefield driver. I simulated the process from start to finish and constrained the optimum experimental parameters.

All of these components together contribute to building low density laser wakefield acceleration experiments which could operate at high repetition rates. This would allow the production of higher energy electrons at faster rates, both allowing current experiments to perform better and enabling new experiments which are currently impossible. Clearly this thesis is but a stone in the path which leads towards these aims, but I'm glad to have been able to contribute in this small way and I'm excited by the work which lies ahead. Thank you for reading this far.

Appendix A

Derivation of Wakefield Amplitude and Phase in Autoresonance

We wish to use the basic sum of sinusoids described in chapter 2 to calculate the wakefield amplitude and phase resulting from a chirped pulse train. To do this, we have expanded the pulse timing to second order as:

$$t_n = \tau_{\mathrm{p0}} \frac{\omega_{\mathrm{p0}}}{\omega_{\mathrm{i}}} n + \frac{4\pi^2 \alpha n^2}{\omega_{\mathrm{i}}} \left(\frac{\omega_{\mathrm{i}}}{\omega_{\mathrm{p0}}} \right)^{-2} + \dots \tag{A.1}$$

$$\delta t_n \equiv t_n - n\tau_{\mathrm{p0}} = 2\pi n \left(\frac{1}{\omega_{\mathrm{i}}} - \frac{1}{\omega_{\mathrm{p0}}} \right) + \frac{4\pi^2 \alpha n^2}{\omega_{\mathrm{i}}} \left(\frac{\omega_{\mathrm{i}}}{\omega_{\mathrm{p0}}} \right)^{-2} + \dots \tag{A.2}$$

$$\omega_{\mathrm{p0}} \delta t_n = 2\pi n \left(\frac{\omega_{\mathrm{p0}}}{\omega_{\mathrm{i}}} - 1 \right) + 4\pi^2 \alpha n^2 \left(\frac{\omega_{\mathrm{i}}}{\omega_{\mathrm{p0}}} \right)^{-3} + \dots \tag{A.3}$$

We can then use the expression described earlier to calculate the expected wakefield amplitude:

$$A^2 = \left[\sum_{n=1}^{N} a_n \cos\left(\omega_{\mathrm{p}} \delta \tau_n \right) \right]^2 + \left[\sum_{n=1}^{N} a_n \sin\left(\omega_{\mathrm{p}} \delta \tau_n \right) \right]^2 \tag{A.4}$$

First we can calculate the sum:

$$\begin{aligned} \sum_{n=1}^{N} \sin\left(\omega_{\mathrm{p0}} \delta t_n \right) &\approx \int_{x=0}^{N} \mathrm{d}x \ \sin\left(\omega_{\mathrm{p0}} \delta t_x \right) \\ &= \int_{x=0}^{N} \mathrm{d}x \ \sin\left[2\pi x \left(\frac{\omega_{\mathrm{p0}}}{\omega_{\mathrm{i}}} - 1 \right) + 4\pi^2 \alpha x^2 \left(\frac{\omega_{\mathrm{i}}}{\omega_{\mathrm{p0}}} \right)^{-3} \right] \end{aligned} \tag{A.5}$$

This looks like a Fresnel integral and hence we make the substitution below:

$$\begin{aligned}
\frac{\pi v(x)^2}{2} - \chi &\equiv 2\pi x\left(\frac{\omega_{\mathrm{p0}}}{\omega_{\mathrm{i}}} - 1\right) + 4\pi^2\alpha x^2\left(\frac{\omega_{\mathrm{i}}}{\omega_{\mathrm{p0}}}\right)^{-3} \\
&= \frac{\pi}{2}\left[8\pi\alpha\left(\frac{\omega_{\mathrm{i}}}{\omega_{\mathrm{p0}}}\right)^{-3} x^2 + 4x\left(\frac{\omega_{\mathrm{p0}}}{\omega_{\mathrm{i}}} - 1\right)\right] \\
&= \frac{\pi}{2}\left[\sqrt{8\pi\alpha}\left(\frac{\omega_{\mathrm{i}}}{\omega_{\mathrm{p0}}}\right)^{-\frac{3}{2}} x + \frac{1}{\sqrt{8\pi\alpha}}\left(\frac{\omega_{\mathrm{i}}}{\omega_{\mathrm{p0}}}\right)^{+\frac{3}{2}} 2\left(\frac{\omega_{\mathrm{p0}}}{\omega_{\mathrm{i}}} - 1\right)\right]^2 \\
&\quad - \frac{\pi}{2}\frac{1}{8\pi\alpha}\left(\frac{\omega_{\mathrm{i}}}{\omega_{\mathrm{p0}}}\right)^{3} 4\left(\frac{\omega_{\mathrm{p0}}}{\omega_{\mathrm{i}}} - 1\right)^2 \\
&= \frac{\pi}{2}\left[\sqrt{8\pi\alpha}\left(\frac{\omega_{\mathrm{i}}}{\omega_{\mathrm{p0}}}\right)^{-\frac{3}{2}} x - \frac{1}{\sqrt{2\pi\alpha}}\left(\frac{\omega_{\mathrm{i}}}{\omega_{\mathrm{p0}}}\right)^{+\frac{1}{2}}\left(\frac{\omega_{\mathrm{i}}}{\omega_{\mathrm{p0}}} - 1\right)\right]^2 \\
&\quad - \frac{1}{4\alpha}\left(\frac{\omega_{\mathrm{i}}}{\omega_{\mathrm{p0}}}\right)\left(\frac{\omega_{\mathrm{i}}}{\omega_{\mathrm{p0}}} - 1\right)^2
\end{aligned} \tag{A.6}$$

Here we have:

$$v(x) = cx + v(0) \tag{A.7}$$

$$c = \sqrt{8\pi\alpha}\left(\frac{\omega_{\mathrm{i}}}{\omega_{\mathrm{p0}}}\right)^{-\frac{3}{2}} \tag{A.8}$$

$$v(0) = -\frac{1}{\sqrt{2\pi\alpha}}\left(\frac{\omega_{\mathrm{i}}}{\omega_{\mathrm{p0}}}\right)^{+\frac{1}{2}}\left(\frac{\omega_{\mathrm{i}}}{\omega_{\mathrm{p0}}} - 1\right) \tag{A.9}$$

$$\chi = \frac{1}{4\alpha}\left(\frac{\omega_{\mathrm{i}}}{\omega_{\mathrm{p0}}}\right)\left(\frac{\omega_{\mathrm{i}}}{\omega_{\mathrm{p0}}} - 1\right)^2 \tag{A.10}$$

We can then approximate the sum in terms of the Fresnel integrals $S(v) = \int \sin(\pi v^2/2)\mathrm{d}v$

and $C(v) = \int \cos(\pi v^2/2)\mathrm{d}v$:

$$\begin{aligned}
\sum_{n=1}^{N} \sin\left(\omega_{\mathrm{p}0}\delta t_n\right) &\approx \int_{v=v(0)}^{v(N)} \frac{\mathrm{d}x}{\mathrm{d}v}\mathrm{d}v \ \sin\left(\frac{\pi v^2}{2} - \chi\right) \\
&= \frac{1}{c}\int_{v=v(0)}^{v(N)} \mathrm{d}v \ \left[\sin\left(\frac{\pi v^2}{2}\right)\cos\chi - \cos\left(\frac{\pi v^2}{2}\right)\sin\chi\right] \\
&= \frac{1}{c}\cos\chi\int_{v=v(0)}^{v(N)} \sin\left(\frac{\pi v^2}{2}\right)\mathrm{d}v - \frac{1}{c}\sin\chi\int_{v=v(0)}^{v(N)} \cos\left(\frac{\pi v^2}{2}\right)\mathrm{d}v \\
&= \frac{1}{c}\cos\chi\left\{S\left[v(N)\right] - S\left[v(0)\right]\right\} - \frac{1}{c}\sin\chi\left\{C\left[v(N)\right] - C\left[v(0)\right]\right\} \\
&= \frac{1}{c}\Delta S\cos\chi - \frac{1}{c}\Delta C\sin\chi
\end{aligned} \tag{A.11}$$

Similarly, we can approximate the cosine sum:

$$\begin{aligned}
\sum_{n=1}^{N} \cos\left(\omega_{\mathrm{p}0}\delta t_n\right) &\approx \int_{v=v(0)}^{v(N)} \frac{\mathrm{d}x}{\mathrm{d}v}\mathrm{d}v \ \cos\left(\frac{\pi v^2}{2} - \chi\right) \\
&= \frac{1}{c}\cos\chi\left\{C\left[v(N)\right] - C\left[v(0)\right]\right\} + \frac{1}{c}\sin\chi\left\{S\left[v(N)\right] - S\left[v(0)\right]\right\} \\
&= \frac{1}{c}\Delta C\cos\chi + \frac{1}{c}\Delta S\sin\chi
\end{aligned} \tag{A.12}$$

Putting these together we estimate the wakefield amplitude:

$$\begin{aligned}
A^2 &= \left[\sum_{n=1}^{N} \sin\left(\omega_{\mathrm{p}0}\delta t_n\right)\right]^2 + \left[\sum_{n=1}^{N} \cos\left(\omega_{\mathrm{p}0}\delta t_n\right)\right]^2 \\
&= \frac{1}{c^2}\Delta S^2\cos^2\chi + \frac{1}{c^2}\Delta C^2\sin^2\chi - \frac{1}{c^2}2\Delta S\Delta C\sin\chi\cos\chi \\
&\quad + \frac{1}{c^2}\Delta C^2\cos^2\chi + \frac{1}{c^2}\Delta S^2\sin^2\chi + \frac{1}{c^2}2\Delta S\Delta C\sin\chi\cos\chi \\
&= \frac{1}{c^2}\left(\Delta S^2 + \Delta C^2\right)
\end{aligned} \tag{A.13}$$

We can therefore write:

$$\begin{aligned} A &= \frac{1}{c}\left|\Delta C + i\Delta S\right| \\ &= \frac{1}{c}\left|\mathrm{Eu}\left[v(N)\right] - \mathrm{Eu}\left[v(0)\right]\right| \end{aligned} \tag{A.14}$$

$$\text{where } \mathrm{Eu}(v) = C(v) + iS(v) \tag{A.15}$$

We can go further and calculate the phase of the wakefield excited by a chirped pulse train in the same manner:

$$\begin{aligned} \tan\phi &= \left[\sum_{n=1}^{N}\sin\left(\omega_{\mathrm{p0}}\delta t_n\right)\right]\left[\sum_{n=1}^{N}\cos\left(\omega_{\mathrm{p0}}\delta t_n\right)\right]^{-1} \\ &= \left[\Delta S\cos\chi - \Delta C\sin\chi\right]\left[\Delta C\cos\chi + \Delta S\sin\chi\right]^{-1} \\ &= \left[\frac{\Delta S}{\Delta C} - \frac{\sin\chi}{\cos\chi}\right]\left[1 + \frac{\Delta S}{\Delta C}\frac{\sin\chi}{\cos\chi}\right]^{-1} \\ &= \left\{\tan\left[\arg\left(\Delta\mathrm{Eu}\right)\right] - \tan\chi\right\}\left\{1 + \tan\left[\arg\left(\Delta\mathrm{Eu}\right)\right]\tan\chi\right\}^{-1} \\ &= \tan\left[\arg\left(\Delta\mathrm{Eu}\right) - \chi\right] \end{aligned} \tag{A.16}$$

We have therefore calculated to second order the expected wakefield amplitude and phase resulting from a linearly chirped pulse train.

Appendix B

Detailed TESS Calculations

We're interested in extending the TESS calculations to general probe and reference pulses, which are not identical. We start by looking at the change in phase accumulated by the probe compared to the reference. In particular, we're interested in the difference in probe and reference spectral phases, $\Delta\psi_{\mathrm{a,b}}(\omega, \Omega)$, where ψ_{x} is the spectral phase of a pulse x: $\mathcal{E}_{\mathrm{x}}(\omega) = |\mathcal{E}_{\mathrm{x}}(\omega)| \exp(i\psi_{\mathrm{x}})$. We can expand $\Delta\psi_{\mathrm{a,b}}(\omega, \Omega)$ to a second order polynomial in ω, assuming that the higher order contributions are small.

$$\begin{aligned}
\Delta\psi_{\mathrm{a,b}}(\omega,\Omega) &\equiv \psi_{\mathrm{a}}(\omega-\Omega) - \psi_{\mathrm{b}}(\omega) \\
&= \left[\psi_{\mathrm{a}}^{(0)} + \psi_{\mathrm{a}}^{(1)}(\omega-\omega_0-\Omega) + \frac{1}{2}\psi_{\mathrm{a}}^{(2)}(\omega-\omega_0-\Omega)^2 + \ldots\right] \\
&\quad - \left[\psi_{\mathrm{b}}^{(0)} + \psi_{\mathrm{b}}^{(1)}(\omega-\omega_0) + \frac{1}{2}\psi_{\mathrm{a}}^{(2)}(\omega-\omega_0)^2 + \ldots\right] \\
&= \frac{1}{2}\left(\psi_{\mathrm{a}}^{(2)} - \psi_{\mathrm{b}}^{(2)}\right)(\omega-\omega_0)^2 + \left(\psi_{\mathrm{a}}^{(1)} - \psi_{\mathrm{b}}^{(1)} - \psi_{\mathrm{a}}^{(2)}\Omega\right)(\omega-\omega_0) \\
&\quad + \left(\psi_{\mathrm{a}}^{(0)} - \psi_{\mathrm{b}}^{(0)} - \psi_{\mathrm{a}}^{(1)}\Omega + \frac{1}{2}\psi_{\mathrm{a}}^{(2)}\Omega^2\right) + \ldots \\
&= A(\omega-\omega_0)^2 + B(\Omega)(\omega-\omega_0) + C(\Omega) + \ldots \qquad \text{(B.1)}
\end{aligned}$$

$$\text{for} \quad A = \frac{1}{2}\left(\psi_{\mathrm{a}}^{(2)} - \psi_{\mathrm{b}}^{(2)}\right) \qquad \text{(B.2)}$$

$$B = \psi_{\mathrm{a}}^{(1)} - \psi_{\mathrm{b}}^{(1)} - \psi_{\mathrm{a}}^{(2)}\Omega \qquad \text{(B.3)}$$

$$C = \psi_{\mathrm{a}}^{(0)} - \psi_{\mathrm{b}}^{(0)} - \psi_{\mathrm{a}}^{(1)}\Omega + \frac{1}{2}\psi_{\mathrm{a}}^{(2)}\Omega^2 \qquad \text{(B.4)}$$

First we will consider the simplest case, where the two pulses a and b have equal GDD, before exploring the effect of a finite difference in GDD:

B.1 Case $\psi_{\mathrm{a}}^{(2)} = \psi_{\mathrm{b}}^{(2)}$

When the two pulses have equal GDDs the coefficient $A = 0$ and hence the function $H_{\mathrm{a,b}}(t, \Omega)$, which describes the TESS satellite peaks, appears like a Fourier transform with a translated time:

$$\begin{aligned} H_{\mathrm{a,b}}(t,\Omega) &\equiv \frac{1}{\sqrt{2\pi}}\int_{-\infty}^{\infty} \mathcal{E}_{\mathrm{a}}(\omega-\Omega)\mathcal{E}_{\mathrm{b}}^{*}(\omega)\exp\left(i\omega t\right)\mathrm{d}\omega \\ &= \frac{1}{\sqrt{2\pi}}\int_{-\infty}^{\infty} \left|\mathcal{E}_{\mathrm{a}}(\omega-\Omega)\right|\left|\mathcal{E}_{\mathrm{b}}(\omega)\right|\exp\left[i\psi_{\mathrm{a}}(\omega-\Omega)-i\psi_{\mathrm{b}}(\omega)+i\omega t\right]\mathrm{d}\omega \\ &= \frac{1}{\sqrt{2\pi}}\int_{-\infty}^{\infty} \left|\mathcal{E}_{\mathrm{a}}(\omega-\Omega)\right|\left|\mathcal{E}_{\mathrm{b}}(\omega)\right|\exp\left[iB(\Omega)(\omega-\omega_0)+iC(\Omega)+i\omega t\right]\mathrm{d}\omega \\ &= \frac{1}{\sqrt{2\pi}}e^{iC(\Omega)}e^{-i\omega_0 B(\Omega)}\int_{-\infty}^{\infty} \left|\mathcal{E}_{\mathrm{a}}(\omega-\Omega)\right|\left|\mathcal{E}_{\mathrm{b}}(\omega)\right|\exp\left\{i\omega\left[B(\Omega)+t\right]\right\}\mathrm{d}\omega \end{aligned} \quad \text{(B.5)}$$

Since $\left|\mathcal{E}_{\mathrm{a}}(\omega-\Omega)\right|\left|\mathcal{E}_{\mathrm{b}}(\omega)\right|$ is real, this integral is maximised when the phase $\omega\left[B(\Omega)+t\right]=0$, and hence there is a peak at $t_\Omega = -B(\Omega) = \psi_{\mathrm{a}}^{(2)}\Omega + \psi_{\mathrm{b}}^{(1)} - \psi_{\mathrm{a}}^{(1)}$. At this peak, the integral has an amplitude:

$$\left|H_{\mathrm{a,b}}(t_\Omega,\Omega)\right| = \frac{1}{\sqrt{2\pi}}\int_{-\infty}^{\infty}\left|\mathcal{E}_{\mathrm{a}}(\omega-\Omega)\right|\left|\mathcal{E}_{\mathrm{b}}(\omega)\right|\mathrm{d}\omega \quad \text{(B.6)}$$

B.2 Case $\psi_{\mathrm{a}}^{(2)} \approx \psi_{\mathrm{b}}^{(2)}$

The situation when the pulses have different GDDs is more complex, but we can approach it by treating $H_{\mathrm{a,b}}(t,\Omega)$ as a Fourier transform of a product, which is equal to the convolution of the individual Fourier transforms:

$$\begin{aligned} H_{\mathrm{a,b}}(t,\Omega) &= \frac{1}{\sqrt{2\pi}}\int_{-\infty}^{\infty} \left|\mathcal{E}_{\mathrm{a}}(\omega-\Omega)\right|\left|\mathcal{E}_{\mathrm{b}}(\omega)\right|\exp\left[i\psi_{\mathrm{a}}(\omega-\Omega)-i\psi_{\mathrm{b}}(\omega)+i\omega t\right]\mathrm{d}\omega \\ &= \frac{1}{\sqrt{2\pi}}\int_{-\infty}^{\infty}\left[\frac{1}{\sqrt{2\pi}}\int_{-\infty}^{\infty}\left|\mathcal{E}_{\mathrm{a}}(\omega'-\Omega)\right|\left|\mathcal{E}_{\mathrm{b}}(\omega')\right|e^{i\omega' t'}\mathrm{d}\omega'\right] \\ &\qquad \cdot\left\{\frac{1}{\sqrt{2\pi}}\int_{-\infty}^{\infty}\exp\left[i\Delta\psi_{\mathrm{a,b}}(\omega'',\Omega)+i\omega''(t-t')\right]\mathrm{d}\omega''\right\}\mathrm{d}t' \\ &= \frac{1}{\sqrt{2\pi}}\int_{-\infty}^{\infty} K_{\mathrm{a,b}}(t',\Omega)\cdot L_{\mathrm{a,b}}(t-t',\Omega)\ \mathrm{d}t' \end{aligned} \quad \text{(B.7)}$$

The phase term $\Delta\psi_{\mathrm{a,b}}(\omega'',\Omega)+i\omega''(t-t')$ is a real quadratic and so the integral

$L_{\mathrm{a,b}}(t, t', \Omega)$ in Eq (B.7) is a Gaussian integral with a purely imaginary argument, or a Fresnel integral, which admits an analytic solution:

$$\begin{aligned} L_{\mathrm{a,b}}(t-t', \Omega) &= \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} \exp\left\{ iA(\omega'' - \omega_0)^2 + i\left[B(\Omega) + t - t'\right](\omega'' - \omega_0) \right\} \mathrm{d}\omega'' \\ &\qquad \cdot \exp\left[+iC(\Omega) + i\omega_0(t-t')\right] \\ &= \frac{1}{\sqrt{2\pi}} \sqrt{\frac{\pi \mathrm{i}}{A}} \exp\left\{ -i\frac{\left[B(\Omega) + t - t'\right]^2}{4A} \right\} e^{iC(\Omega)} e^{i\omega_0(t-t')} \\ &= \sqrt{\frac{\mathrm{i}}{2A}} e^{iC(\Omega)} e^{i\omega_0(t-t')} e^{\lambda g(t-t', \Omega)} \end{aligned} \tag{B.8}$$

Where $g(\tau, \Omega) = -\frac{\mathrm{i}}{4}\left(B(\Omega) + \tau\right)^2$ and $\lambda = A^{-1} = 2\left(\psi_{\mathrm{a}}^{(2)} - \psi_{\mathrm{b}}^{(2)}\right)^{-1}$. This means that for a given t and Ω the integral $H_{\mathrm{a,b}}(t, \Omega)$ is of the form $\int_\Gamma f(x) e^{\lambda g(x)} \mathrm{d}x$ and if the difference between the probe and reference pulse GDDs is sufficiently small, $\lambda \gg 1$. $g(t-t', \Omega)$ is an exact quadratic in t' and so if $K_{\mathrm{a,b}}(t', \Omega)$ is sufficiently well behaved we can extend t' to the complex plane while ensuring that $f(t', \Omega)$ and $g(t-t', \Omega)$ are holomorphic. This allows us to deform the contour of integration Γ and use the method of steepest descent, with $g(t-t', \Omega)$ having a single non-degenerate saddle point at $t - t'_0 = -B(\Omega)$ at which point $g(t-t'_0, \Omega) = 0$ and $g''(t-t'_0, \Omega) = -\frac{\mathrm{i}}{2}$:

$$\begin{aligned} H_{\mathrm{a,b}}(t, \Omega) &= \frac{1}{\sqrt{2\pi}} \sqrt{\frac{\mathrm{i}}{2A}} e^{iC(\Omega)} \int_{-\infty}^{\infty} K_{\mathrm{a,b}}(t', \Omega) e^{i\omega(t-t')} e^{\lambda g(t-t', \Omega)} \mathrm{d}t' \\ &\approx \sqrt{\frac{\mathrm{i}}{4\pi A}} e^{iC(\Omega)} K_{\mathrm{a,b}}(t'_0, \Omega) e^{i\omega_0(t-t'_0)} \sqrt{\frac{2\pi}{\lambda}} e^{\lambda g(t-t'_0, \Omega)} \left[-g''(t-t'_0, \Omega)\right]^{-\frac{1}{2}} \\ &= e^{iC(\Omega)} e^{i\omega_0(t-t'_0)} K_{\mathrm{a,b}}(t'_0, \Omega) \\ &= \frac{1}{\sqrt{2\pi}} e^{iC(\Omega)} e^{-i\omega_0 B(\Omega)} \int_{-\infty}^{\infty} \left|\mathcal{E}_{\mathrm{a}}(\omega' - \Omega)\right| \left|\mathcal{E}_{\mathrm{b}}^*(\omega')\right| e^{i\omega'(t+B(\Omega))} \mathrm{d}\omega' \end{aligned} \tag{B.9}$$

This expression is identical to Eqn (B.5) for the case $\psi_{\mathrm{a}}^{(2)} = \psi_{\mathrm{b}}^{(2)}$ and hence when the difference between GDDs is small but finite, the integral will again be maximised

at approximately $t_\Omega = -B(\Omega) = \psi_{\rm a}^{(2)}\Omega + \psi_{\rm b}^{(1)} - \psi_{\rm a}^{(1)}$ with an amplitude described in Eqn (B.6).

B.3 Case $\psi_{\rm a}^{(2)} \neq \psi_{\rm b}^{(2)}$ for Gaussian $|\mathcal{E}_{\rm x}(\omega)|$

The final case I will look at is when pulses a and b can both be assumed to be Gaussian but are not identical. This allows me to treat the integrals exactly without using the approximation in the method of steepest descent. First define the pulse envelopes as below, following the conventions of Matlis et al. but for the purposes of generality also allowing the pulse central frequency to vary with $\kappa_{\rm x}$:

$$|\mathcal{E}_{\rm x}(\omega)| = \mathcal{E}_{\rm x0} \exp\left[-\frac{1}{2}\left(\frac{\omega-\omega_0}{\delta\omega_{\rm x}}\right)^2 - \kappa_{\rm x}(\omega-\omega_0)\right] \tag{B.10}$$

$$\begin{aligned}
|\mathcal{E}_{\rm a}(\omega-\Omega)|\,|\mathcal{E}_{\rm b}(\omega)| &= \mathcal{E}_{\rm a0}\mathcal{E}_{\rm b0} \exp\left[-\frac{1}{2}\left(\frac{\omega-\omega_0-\Omega}{\delta\omega_{\rm a}}\right)^2 - \kappa_{\rm a}(\omega-\omega_0) - \frac{1}{2}\left(\frac{\omega-\omega_0}{\delta\omega_{\rm b}}\right)^2 - \kappa_{\rm b}(\omega-\omega_0)\right] \\
&= \mathcal{E}_{\rm a0}\mathcal{E}_{\rm b0} \exp\left[-\frac{1}{2}\left(\frac{1}{\delta\omega_{\rm a}^2}+\frac{1}{\delta\omega_{\rm b}^2}\right)(\omega-\omega_0)^2 - \left(\kappa_{\rm a}+\kappa_{\rm b}-\frac{\Omega}{\delta\omega_{\rm a}^2}\right)(\omega-\omega_0) - \frac{\Omega^2}{2\delta\omega_{\rm a}^2}\right] \\
&= \mathcal{E}_{\rm a0}\mathcal{E}_{\rm b0} \exp\left[-D\,(\omega-\omega_0)^2 - E(\Omega)\,(\omega-\omega_0) - F(\Omega)\right]
\end{aligned} \tag{B.11}$$

This allows us to construct the integrand of $H_{\rm a,b}(t,\Omega)$ as a Gaussian, using $\alpha \equiv D - iA$, $\beta(\Omega) \equiv E(\Omega) - iB(\Omega)$ and $\gamma \equiv F(\Omega) - iC(\Omega)$, and hence to integrate it exactly, with the positive real part of α ensuring convergence:

$$
\begin{aligned}
H_{\rm a,b}(t,\Omega) &= \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} |\mathcal{E}_{\rm a}(\omega-\Omega)|\,|\mathcal{E}_{\rm b}^{*}(\omega)| \exp\left[i\psi_{\rm a}(\omega-\Omega) - i\psi_{\rm b}(\omega) + i\omega t\right] {\rm d}\omega \\
&= \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} \mathcal{E}_{\rm a0}\mathcal{E}_{\rm b0} \exp\left[-\alpha\left(\omega-\omega_0\right)^2 - \beta(\Omega)\left(\omega-\omega_0\right) - \gamma(\Omega) + i\omega t\right] {\rm d}\omega \\
&= \frac{1}{\sqrt{2\pi}} \mathcal{E}_{\rm a0}\mathcal{E}_{\rm b0}\; e^{-\gamma(\Omega)} e^{i\omega_0 t} \int_{-\infty}^{\infty} \exp\left\{-\alpha\left(\omega-\omega_0\right)^2 + \left[it - \beta(\Omega)\right]\left(\omega-\omega_0\right)\right\} {\rm d}\omega \\
&= \frac{1}{\sqrt{2\pi}} \mathcal{E}_{\rm a0}\mathcal{E}_{\rm b0}\; e^{-\gamma(\Omega)} e^{i\omega_0 t} \sqrt{\frac{\pi}{\alpha}} \exp\left\{\frac{\left[it - \beta(\Omega)\right]^2}{4\alpha}\right\}
\end{aligned}
\tag{B.12}
$$

This expression demonstrates what we expect; $H_{\rm a,b}(t,\Omega)$ is a Gaussian peak, which we might naively expect to be centred at $t_\Omega = \Im\left[\beta(\Omega)\right] = -B(\Omega)$. However, the argument of the exponential is more complicated and needs unpacking:

$$
\begin{aligned}
\Re\left\{\frac{\left[it - \beta(\Omega)\right]^2}{4\alpha}\right\} &= \Re\left\{\frac{(D+iA)}{4\left(A^2+D^2\right)}\left[it - E(\Omega) + iB(\Omega)\right]^2\right\} \\
&= \Re\left(\frac{(D+iA)}{4\left(A^2+D^2\right)}\left\{-\left[t+B(\Omega)\right]^2 + \left[E(\Omega)\right]^2 - 2iE(\Omega)\left[t+B(\Omega)\right]\right\}\right) \\
&= \frac{1}{4\left(A^2+D^2\right)}\left\{-D\left[t+B(\Omega)\right]^2 + D\left[E(\Omega)\right]^2 + 2AE(\Omega)\left[t+B(\Omega)\right]\right\} \\
&= \frac{1}{4\left(A^2+D^2\right)}\left\{-D\left[t+B(\Omega) - \frac{AE(\Omega)}{D})\right]^2 + \frac{\left[AE(\Omega)\right]^2}{D} + D\left[E(\Omega)\right]^2\right\}
\end{aligned}
\tag{B.13}
$$

This means the peak occurs at a position t_Ω slightly shifted from $-B(\Omega)$:

$$\begin{aligned}
t_\Omega &= -B(\Omega) + \frac{AE(\Omega)}{D} \\
&= \left(\psi_{\mathrm{a}}^{(2)}\Omega + \psi_{\mathrm{b}}^{(1)} - \psi_{\mathrm{a}}^{(1)}\right) + \frac{1}{2}\left(\psi_{\mathrm{a}}^{(2)} - \psi_{\mathrm{b}}^{(2)}\right)\left(\kappa_{\mathrm{a}} + \kappa_{\mathrm{b}} - \frac{\Omega}{\delta\omega_{\mathrm{a}}^2}\right)\left[\frac{1}{2}\left(\frac{1}{\delta\omega_{\mathrm{a}}^2} + \frac{1}{\delta\omega_{\mathrm{b}}^2}\right)\right]^{-1} \\
&= \left[\psi_{\mathrm{b}}^{(1)} - \psi_{\mathrm{a}}^{(1)} + \left(\psi_{\mathrm{a}}^{(2)} - \psi_{\mathrm{b}}^{(2)}\right)(\kappa_{\mathrm{a}} + \kappa_{\mathrm{b}})\left(\frac{\delta\omega_{\mathrm{a}}^2\delta\omega_{\mathrm{b}}^2}{\delta\omega_{\mathrm{a}}^2 + \delta\omega_{\mathrm{b}}^2}\right)\right] \\
&\quad + \Omega\left[\psi_{\mathrm{a}}^{(2)} - \left(\psi_{\mathrm{a}}^{(2)} - \psi_{\mathrm{b}}^{(2)}\right)\frac{1}{\delta\omega_{\mathrm{a}}^2}\left(\frac{\delta\omega_{\mathrm{a}}^2\delta\omega_{\mathrm{b}}^2}{\delta\omega_{\mathrm{a}}^2 + \delta\omega_{\mathrm{b}}^2}\right)\right] \\
&= t_0 + \Omega\left[\left(\delta\omega_{\mathrm{a}}^2 + \delta\omega_{\mathrm{b}}^2\right)\psi_{\mathrm{a}}^{(2)} - \delta\omega_{\mathrm{b}}^2\left(\psi_{\mathrm{a}}^{(2)} - \psi_{\mathrm{b}}^{(2)}\right)\right]\left(\frac{1}{\delta\omega_{\mathrm{a}}^2 + \delta\omega_{\mathrm{b}}^2}\right) \\
&= t_0 + \Omega\frac{\delta\omega_{\mathrm{a}}^2\psi_{\mathrm{a}}^{(2)} + \delta\omega_{\mathrm{b}}^2\psi_{\mathrm{b}}^{(2)}}{\delta\omega_{\mathrm{a}}^2 + \delta\omega_{\mathrm{b}}^2} \qquad \text{(B.14)} \\
&= t_0 + \Omega\psi_{\mathrm{eff}}^{(2)} \qquad \text{(B.15)}
\end{aligned}$$

This implies that instead of the GDD of pulse a, $\psi_{\mathrm{a}}^{(2)}$, we should use the mean GDD of pulses a and b, weighted by the square of their bandwidths. If the probe and reference pulses are similar but not identical, $\delta\omega_a \approx \delta\omega_b$, this tends towards the mean GDD, $\frac{1}{2}\left(\psi_{\mathrm{a}}^{(2)} + \psi_{\mathrm{b}}^{(2)}\right)$. While the position of the main sideband t_0 is dependent on the central frequencies of the two pulses through κ_a and κ_b, the separation of the peak from the main sideband, $\Omega\psi_{\mathrm{eff}}^{(2)}$, is unaffected by changes in the central frequency of the two pulses.

References

[1] J. J. Thomson, "Cathode Rays," *London, Edinburgh, Dublin Philos. Mag. J. Sci.*, vol. 44, no. 269, p. 293, 1897.

[2] CMS Collaboration, "Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC," *Phys. Lett. Sect. B Nucl. Elem. Part. High-Energy Phys.*, vol. 716, no. 1, pp. 30–61, 2012. arXiv: 1207.7235.

[3] (). European X-Ray Free Electron Laser, [Online]. Available: https://www.xfel.eu/.

[4] (). LINAC Coherent Light Source, [Online]. Available: https://lcls.slac.stanford.edu/.

[5] P. Kneisel, "High gradient superconducting niobium cavities. A review of the present status," *IEEE Trans. Appl. Supercond.*, vol. 9, no. 2, pp. 1023–1029, 1999.

[6] A. Grassellino, A. Romanenko, Y. Trenikhina, *et al.*, "Unprecedented quality factors at accelerating gradients up to 45 MV/m in niobium superconducting resonators via low temperature nitrogen infusion," *Supercond. Sci. Technol.*, vol. 30, no. 9, p. 94 004, 2017. arXiv: 1701.06077.

[7] (). Diamond Light Source, [Online]. Available: https://www.diamond.ac.uk/Home.html.

[8] T Tajima and J. Dawson, "Laser electron accelerator," *Phys. Rev. Lett.*, vol. 43, no. 4, pp. 267–270, 1979.

[9] E. Esarey, C. Schroeder, and W. Leemans, "Physics of laser-driven plasma-based electron accelerators," *Rev. Mod. Phys.*, vol. 81, no. 3, pp. 1229–1285, 2009.

[10] S. M. Hooker, "Developments in laser-driven plasma accelerators," *Nat. Photonics*, vol. 7, no. 775, pp. 1–11, 2013.

[11] C. E. Clayton, C. Joshi, C. Darrow, and D. Umstadter, "Relativistic plasma-wave excitation by collinear optical mixing," *Phys. Rev. Lett.*, vol. 54, no. 21, pp. 2343–2346, 1985.

[12] C. J. Walsh, D. Villeneuve, and H. Baldis, "Electron plasma-wave production by stimulated Raman scattering: Competition with stimulated Brillouin scattering," *Phys. Rev. Lett.*, vol. 53, no. 15, pp. 1445–1448, 1984.

[13] S. P. Mangles, C. D. Murphy, and Z. Najmudin, "Monoenergetic beams of relativistic electrons from intense laser-plasma interactions," *Nature*, vol. 431, pp. 535–538, 2004.

[14] W. Leemans, B. Nagler, A. Gonsalves, *et al.*, "GeV electron beams from a centimetre-scale accelerator," *Nat. Phys.*, vol. 2, no. 10, pp. 696–699, 2006.

[15] W. Leemans, A. Gonsalves, H.-S. Mao, *et al.*, "Multi-GeV Electron Beams from Capillary-Discharge-Guided Subpetawatt Laser Pulses in the Self-Trapping Regime," *Phys. Rev. Lett.*, vol. 113, no. 24, p. 245 002, 2014.

[16] J. M. Cole, J. C. Wood, N. C. Lopes, *et al.*, "Laser-wakefield accelerators as hard x-ray sources for 3D medical imaging of human bone.," *Sci. Rep.*, vol. 5, p. 13 244, 2015.

[17] J. C. Wood, D. J. Chapman, K. Poder, *et al.*, "Ultrafast Imaging of Laser Driven Shock Waves using Betatron X-rays from a Laser Wakefield Accelerator," *Sci. Rep.*, vol. 8, no. 1, pp. 1–10, 2018. arXiv: 1802.02119.

[18] F. V. Hartemann, D. J. Gibson, W. J. Brown, *et al.*, "Compton scattering x-ray sources driven by laser wakefield acceleration," *Phys. Rev. Spec. Top. - Accel. Beams*, vol. 10, no. 1, pp. 1–8, 2007.

[19] N. D. Powers, I. Ghebregziabher, G. Golovin, *et al.*, "Quasi-monoenergetic and tunable X-rays from a laser-driven Compton light source," *Nat. Photonics*, vol. 8, no. 1, pp. 28–31, 2013.

[20] A. Döpp, E. Guillaume, C. Thaury, *et al.*, "Single Shot Radiography Using an All-optical Compton Backscattering Source," *Phys. Procedia*, vol. 77, pp. 9–14, 2015.

[21] H. E. Tsai, X. Wang, J. Shaw, *et al.*, "Compact tunable Compton x-ray source from laser wakefield accelerator and plasma mirror," *Phys. Plasmas*, vol. 22, no. 2, p. 023 106, 2015. arXiv: 1411.2134.

[22] J. M. Cole, K. T. Behm, E. Gerstmayr, *et al.*, "Experimental Evidence of Radiation Reaction in the Collision of a High-Intensity Laser Pulse with a Laser-Wakefield Accelerated Electron Beam," *Phys. Rev. X*, vol. 8, no. 1, pp. 8–13, 2018. arXiv: 1707.06821.

[23] K. Poder, J. M. Cole, J. C. Wood, *et al.*, "Measurements of self-guiding of ultrashort laser pulses over long distances," *Plasma Phys. Control. Fusion*, vol. 60, no. 1, p. 14 022, 2018.

[24] G. Sarri, W. Schumaker, A. Di Piazza, *et al.*, "Table-top laser-based source of femtosecond, collimated, ultrarelativistic positron beams," *Phys. Rev. Lett.*, vol. 110, no. 25, pp. 1–5, 2013. arXiv: 1304.5379.

[25] W. Schumaker, T Liang, R Clarke, *et al.*, "Making pions with laser light," *New J. Phys.*, vol. 20, no. 7, p. 73 008, 2018.

[26] W. Leemans, "Laser Technology for k-BELLA and Beyond," LBNL, Tech. Rep., 2017.

[27] J. Le Dortz, A Heilmann, M Antier, *et al.*, "Highly scalable femtosecond coherent beam combining demonstrated with 19 fibers," *Opt. Lett.*, vol. 42, no. 10, p. 1887, 2017.

[28] L. Corner, R. Walczak, L. J. Nevay, *et al.*, "Multiple pulse resonantly enhanced laser plasma wakefield acceleration," *AIP Conf. Proc.*, vol. 1507, no. 2012, pp. 872–873, 2012.

[29] S. M. Hooker, R Bartolini, S. P. Mangles, *et al.*, "Multi-pulse laser wakefield acceleration: a new route to efficient, high-repetition-rate plasma accelerators and high flux radiation sources," *J. Phys. B At. Mol. Opt. Phys.*, vol. 47, no. 23, p. 234 003, 2014.

[30] E. Esarey, A. Ting, and P. Sprangle, "Frequency Shifts Induced in Laser Pulses by Plasma Waves," *Phys. Rev. A*, vol. 42, no. 6, pp. 3526–3531, 1990.

[31] N. H. Matlis, A. Maksimchuk, V. Y. Yanovsky, W. Leemans, and M. Downer, "Analysis of sinusoidally-modulated chirped laser pulses by temporally-encoded spectral shifting," *Opt. Lett.*, vol. 41, no. 23, p. 5503, 2016.

[32] A. M. Pukhov, S. Gordienko, S. Kiselev, and I. Kostyukov, "The bubble regime of laser-plasma acceleration: Monoenergetic electrons and the scalability," *Plasma Phys. Control. Fusion*, vol. 46, no. 12 B, 2004. arXiv: `0409089 [physics]`.

[33] W. Lu, M. Tzoufras, C. Joshi, *et al.*, "Generating multi-GeV electron bunches using single stage laser wakefield acceleration in a 3D nonlinear regime," *Phys. Rev. Spec. Top. - Accel. Beams*, vol. 10, no. 6, pp. 1–12, 2007. arXiv: `0612227 [physics]`.

[34] W. Kruer, *The Physics of Laser Plasma Interactions*, 1st ed. Westview Press, 2003, p. 60.

[35] P. Gibbon, *Short pulse laser interactions with matter*. Imperial College Press, 2005, p. 68.

[36] A Modena, F Dorchies, F Amiranoff, *et al.*, "Acceleration of Injected Electrons in a Laser Wakefield Experiment," *Phys. Plasmas*, vol. 6, no. 7, 1999.

[37] A. M. Pukhov and J. Meyer-ter Vehn, "Laser wake field acceleration: The highly non-linear broken-wave regime," *Appl. Phys. B Lasers Opt.*, vol. 74, no. 4-5, pp. 355–361, 2002.

[38] P. Albers, E. Stark, and G. Huber, "Continuous-wave laser operation and quantum efficiency of titanium-doped sapphire," *J. Opt. Soc. Am. B*, vol. 3, no. 1, p. 134, 1986.

[39] G. Mourou, B. Brocklesby, T. Tajima, and J. Limpert, "The future is fibre accelerators," *Nat. Photonics*, vol. 7, no. 4, pp. 258–261, 2013.

[40] P. Tomassini, S. De Nicola, L. Labate, *et al.*, "The resonant multi-pulse ionization injection," *Phys. Plasmas*, vol. 24, no. 10, 2017. arXiv: 1708.04957.

[41] C. D. Murphy, R. Trines, J. Vieira, *et al.*, "Evidence of photon acceleration by laser wake fields," *Phys. Plasmas*, vol. 13, no. 3, pp. 1–8, 2006.

[42] S. Wilks, J. Dawson, W. B. Mori, T Katsouleas, and M. E. Jones, "Photon Accelerator," *Phys. Rev. Lett.*, vol. 62, no. 22, pp. 2600–2603, 1989.

[43] J. P. Geindre, P Audebert, A. Rousse, *et al.*, "Frequency-domain interferometer for measuring the phase and amplitude of a femtosecond pulse probing a laser-produced plasma.," *Opt. Lett.*, vol. 19, no. 23, pp. 1997–1999, 1994.

[44] C. W. Siders, S. Paul, L. Blanc, *et al.*, "Plasma-Based Accelerator Diagnostics Based upon Longitudinal Interferometry with Ultrashort Optical Pulses," vol. 24, no. 2, 1996.

[45] N. H. Matlis, S. Reed, S. Bulanov, *et al.*, "Snapshots of laser wakefields," *Nat. Phys.*, vol. 2, no. 11, pp. 749–753, 2006.

[46] M. F. Kasim and P. Norreys, "Three dimensional spectrometer," pp. 1–6, 2018. arXiv: 1802.00504.

[47] S. Le Blanc, M. Downer, R. Wagner, *et al.*, "Temporal characterization of a self-modulated laser wakefield accelerator," *Phys. Rev. Lett.*, vol. 77, no. 27, pp. 5381–5384, 1996.

[48] J. M. Dias, L. Oliveira e Silva, and J. Mendonca, "Photon acceleration versus frequency-domain interferometry for laser wakefield diagnostics," *Phys. Rev. Spec. Top. - Accel. Beams*, vol. 1, no. 3, p. 031 301, 1998.

[49] M. Downer, R. Zgadzaj, A. Debus, U. Schramm, and M. Kaluza, "Diagnostics for plasma-based electron accelerators," *Rev. Mod. Phys.*, vol. 90, no. 3, p. 035 002, 2018.

[50] A. Sävert, S. P. Mangles, M. Schnell, *et al.*, "Direct Observation of the Injection Dynamics of a Laser Wakefield Accelerator Using Few-Femtosecond Shadowgraphy," *Phys. Rev. Lett.*, vol. 115, no. 5, p. 055 002, 2015.

[51] M. C. Kaluza, H. P. Schlenvoigt, S. P. Mangles, *et al.*, "Measurement of magnetic-field structures in a laser-wakefield accelerator," *Phys. Rev. Lett.*, vol. 105, no. 11, pp. 3–6, 2010. arXiv: `1007.3241`.

[52] R. E. Slusher and C. M. Surko, "Study of density fluctuations in plasmas by small-angle CO2 laser scattering," *Phys. Fluids*, vol. 23, no. 3, pp. 472–490, 1980.

[53] W. Kruer, *The Physics of Laser Plasma Interactions*. Westview Press, 2003, p. 29.

[54] C. Durfee, J. Lynch, and H. M. Milchberg, "Mode Properties of a Plasma Waveguide for Intense Laser Pulses," *Opt. Lett.*, vol. 19, no. 23, pp. 1937–1939, 1994.

[55] E. Esarey, P. Sprangle, J. Krall, and A. Ting, "Self-focusing and guiding of short laser pulses in ionizing gases and plasmas," *IEEE J. Quantum Electron.*, vol. 33, no. 11, pp. 1879–1914, 1997.

[56] S. Jackel, R. Burris, J. Grun, *et al.*, "Channeling of terawatt laser pulses by use of hollow waveguides," *Opt. Lett.*, vol. 20, no. 10, p. 1086, 1995.

[57] A Zigler, Y. Ehrlich, C. Cohen, J. Krall, and P. Sprangle, "Optical guiding of high-intensity laser pulses in a long plasma channel formed by a slow capillary discharge," *J. Opt. Soc. Am. B*, vol. 13, no. 1, p. 68, 1996.

[58] D. J. Spence and S. M. Hooker, "Investigation of a hydrogen plasma waveguide.," *Phys. Rev. E*, vol. 63, no. 1 Pt 2, p. 015 401, 2001.

[59] A Butler, D. J. Spence, and S. M. Hooker, "Guiding of high-intensity laser pulses with a hydrogen-filled capillary discharge waveguide.," *Phys. Rev. Lett.*, vol. 89, no. 18, p. 185 003, 2002.

[60] J Daniels, A. Gonsalves, C. Pieronek, *et al.*, "Laser-assisted capillary discharge for enhanced guiding of tightly-focused laser pulses at low densities," *SPIE Proc.*, vol. 10240, 2017.

[61] J. van Tilborg, A. Gonsalves, E. Esarey, C. B. Schroeder, and W. P. Leemans, "Density characterization of discharged gas-filled capillaries through common-path two-color spectral-domain interferometry," *Opt. Lett.*, vol. In Press, 2018.

[62] C. Thornton, "Experimental aspects of plasma wakefields driven in the linear regime," PhD thesis, 2016, p. 64.

[63] C. Durfee and H. M. Milchberg, "Light Pipe for High Intensity Laser Pulses," *Phys. Rev. Lett.*, vol. 71, no. 15, 1993.

[64] C. Durfee, J. Lynch, and H. M. Milchberg, "Development of a plasma waveguide for high intensity laser pulses," *Phys. Rev. E*, vol. 51, no. 3, pp. 2368–2387, 1995.

[65] H. M. Milchberg, C. Durfee, T. M. Antonsen, and P. Mora, "Channel guided lasers for plasma accelerators," *Proc. Part. Accel. Conf.*, pp. 621–625, 1995.

[66] H. Sheng, K. Y. Kim, V. Kumarappan, B. D. Layer, and H. M. Milchberg, "Plasma waveguides efficiently generated by Bessel beams in elongated cluster gas jets," *Phys. Rev. E*, vol. 72, no. 3, p. 036 411, 2005.

[67] P. B. Corkum, N. H. Burnett, and F Brunel, "Above-Threshold Ionization in the Long-Wavelength Limit," *Phys. Rev. Lett.*, vol. 62, no. 11, pp. 1259–1262, 1989.

[68] N. Lemos and T. Grismayer, "Plasma expansion into a waveguide created by a linearly polarized femtosecond laser pulse," *Phys. Plasmas*, vol. 20, no. 6, p. 063 102, 2013.

[69] N. Lemos, T. Grismayer, L. Cardoso, *et al.*, "Effects of laser polarization in the expansion of plasma waveguides," *Phys. Plasmas*, vol. 20, no. 10, p. 103 109, 2013.

[70] C. Arran, N. H. Matlis, R. Walczak, and S. M. Hooker, "Reconstructing nonlinear plasma wakefields using a generalized temporally encoded spectral shifting analysis," *Phys. Rev. Accel. Beams*, vol. 21, p. 103 501, 2018.

[71] J Cowley, C Thornton, C Arran, *et al.*, "Excitation and Control of Plasma Wakefields by Multiple Laser Pulses," *Phys. Rev. Lett.*, vol. 119, no. 4, p. 044 802, 2017.

[72] R. J. Shalloo, C Arran, L Corner, *et al.*, "Hydrodynamic optical-field-ionized plasma channels," *Phys. Rev. E*, vol. 97, no. 5, p. 053 203, 2018.

[73] R. Walczak, N. Ratan, I. Buchanan, *et al.*, "Fluid simulations of multi-pulse laser wakefield acceleration : excitation and design tolerances," Laser Plasma Accelerators Workshop, 2013.

[74] R. Walczak and S. M. Hooker, "Ion motion , detuning and hosing in Multi-Pulse Laser Wakefield Accelerators .," Advanced Accelerators Concepts, 2014.

[75] H. C. Kim, F Iza, S. S. Yang, M Radmilović-Radjenović, and J. K. Lee, "Particle and fluid simulations of low-temperature plasma discharges: benchmarks and kinetic effects," *J. Phys. D. Appl. Phys.*, vol. 38, no. 19, R283–R301, 2005.

[76] T. D. Arber, K Bennett, C. S. Brady, *et al.*, "Contemporary particle-in-cell approach to laser-plasma modelling," *Plasma Phys. Control. Fusion*, vol. 57, no. 11, p. 113 001, 2015.

[77] J. M. Dawson, "Particle simulation of plasmas," *Rev. Mod. Phys.*, vol. 55, no. 2, p. 403, 1983.

[78] A. Caldwell, E. Gschwendtner, K. Lotov, P. Muggli, and M. Wing, "AWAKE Design Report A Proton-Driven PlasmaWakefield Acceleration Experiment at CERN," pp. 1–38, 2013.

[79] A. Leitenstorfer, C. Fürst, and A. Laubereau, "Widely tunable two-color mode-locked Ti:sapphire laser with pulse jitter of less than 2 fs.," *Opt. Lett.*, vol. 20, no. 8, pp. 916–8, 1995.

[80] M. N. Rosenbluth and C. Liu, "Exitation of Plasma Waves by Two Laser Beams," *Phys. Rev. Lett.*, vol. 29, no. 11, pp. 701–705, 1972.

[81] Y. Kitagawa, T. Matsumoto, T. Minamihata, *et al.*, "Beat-wave excitation of plasma wave and observation of accelerated electrons," *Phys. Rev. Lett.*, vol. 68, no. 1, pp. 48–51, 1992.

[82] M. Deutsch, B. Meerson, and J. E. Golub, "Strong plasma wave excitation by a "chirped" laser beat wave," *Phys. Fluids B Plasma Phys.*, vol. 3, no. 7, p. 1773, 1991.

[83] R. R. Lindberg, A. E. Charman, J. S. Wurtele, and L. Friedland, "Robust Autoresonant Excitation in the Plasma Beat-Wave Accelerator," *Phys. Rev. Lett.*, vol. 93, no. 5, p. 055 001, 2004.

[84] R. R. Lindberg, A. E. Charman, J. S. Wurtele, L. Friedland, and B. A. Shadwick, "Autoresonant beat-wave generation," *Phys. Plasmas*, vol. 13, no. 12, p. 123 103, 2006.

[85] M. Born and E. Wolf, *Principles of Optics*, 7th ed. London: Cambridge University Press, 1999, p. 480.

[86] M. Cornu, "Méthode nouvelle pour la discussion des problémes de diffraction dans le cas d'une onde cylindrique," *J. Phys. Théorique Appliquée*, vol. 3, no. 1, p. 5, 1874.

[87] A. Talbot, *The Railway Transition Spiral*, 5th ed. New York: Engineering News Publishing, 1901.

[88] L. Euler, "De Curvis Elasticis," in *Methodus Inven. lineas curvas maximi minimive Propr. gaudentes, sive Solut. Probl. isoperimetrici latissimo sensu Accept.* 1744, ch. Additament, p. 245.

[89] M. Takeda, H. Ina, and S. Kobayashi, "Fourier-transform method of fringe-pattern analysis for computer-based topography and interferometry," *J. Opt. Soc. Am.*, vol. 72, no. 1, pp. 156–160, 1982.

[90] A. Akhiezer and R. Polovin, "Theory of Wave Motion of an Electron Plasma," *J. Exp. Theor. Phys.*, vol. 3, no. 5, p. 696, 1956.

[91] R. Noble, "Plasma-wave generation in the beat-wave accelerator," *Phys. Rev. A*, vol. 32, no. 1, p. 460, 1985.

[92] V. Malka, "Manipulating relativistic electrons with lasers," *Europhys. Lett.*, vol. 115, no. 5, p. 54 001, 2016.

[93] G. Golovin, S. Chen, N. D. Powers, *et al.*, "Tunable monoenergetic electron beams from independently controllable laser-wakefield acceleration and injection," *Phys. Rev. Spec. Top. - Accel. Beams*, vol. 18, no. 1, p. 011 301, 2015.

[94] K. K. Swanson, H.-E. Tsai, S. K. Barber, *et al.*, "Control of tunable , monoenergetic laser-plasma-accelerated electron beams using a shock-induced density downramp injector," *Phys. Rev. Accel. Beams*, vol. 20, no. 5, p. 051 301, 2017.

[95] R. J. Shalloo, L. Corner, C. Arran, *et al.*, "Generation of laser pulse trains for tests of multi-pulse laser wakefield acceleration," *Nucl. Instruments Methods Phys. Res. Sect. A Accel. Spectrometers, Detect. Assoc. Equip.*, vol. 829, pp. 383–385, 2016.

[96] P. Tournois, "Acousto-optic programmable dispersive filter for compensation of group delay time dispersion in laser systems," *Opt. Commun.*, vol. 140, no. 4-6, pp. 245–249, 1997.

[97] L. M. Gorbunov and V. I. Kirsanov, "Excitation of Plasma Waves by an Electromagnetic Wave Packet," *Sov. Phys. JETP*, vol. 66, no. 2, pp. 290–294, 1987.

[98] W. Leemans, R. Duarte, E. Esarey, *et al.*, "The Berkeley Lab Laser Accelerator (BELLA): A 10 GeV laser plasma accelerator," *AIP Conf. Proc.*, vol. 1299, pp. 3–11, 2010.

[99] X. Li, A. Mosnier, P. Anh, and P. Nghiem, "Design of a 5 GeV Laser Plasma Accelerating Module in the Quasi-linear Regime," *Nucl. Instruments Methods Phys. Res. Sect. A Accel. Spectrometers, Detect. Assoc. Equip.*, vol. In Press, 2018. arXiv: `arXiv:1802.08158v1`.

[100] A. Gonsalves, F. Liu, N. A. Bobrova, *et al.*, "Demonstration of a high repetition rate capillary discharge waveguide," *J. Appl. Phys.*, vol. 119, no. 3, p. 033 302, 2016.

[101] C. Durfee, T. R. Clark, and H. M. Milchberg, "Mode control in a two-pulse-excited plasma waveguide," *J. Opt. Soc. Am. B*, vol. 13, no. 1, pp. 59–67, 1996.

[102] L. Keldysh, "Ionization in the field of a strong electromagnetic wave," *J. Exp. Theor. Phys.*, vol. 20, no. 5, pp. 1307–1314, 1965.

[103] N. Delone and V. P. Krainov, "Tunneling and barrier-suppression ionization of atoms and ions in a laser radiation field," *Phys. - Uspekhi*, vol. 41, no. 5, 1998.

[104] M. V. Ammosov and V. P. Krainov, "Tunnel Ionization of Complex Atoms in an Alternating Electromagnetic Field," *Sov. Phys. JETP*, vol. 64, no. December 1986, pp. 4–7, 1987.

[105] S Augst, D. D. Meyerhofer, D Strickland, and S. L. Chint, "Laser ionization of noble gases by Coulomb-barrier suppression," *J. Opt. Soc Am. B*, vol. 8, no. 4, pp. 858–867, 1991.

[106] J. Posthumus, M. Thompson, I. Frasinki, and K. Codling, "Molecular dissociative ionisation using a classical over-the-barrier approach," in *Inst. Phys. Conf. Ser. No. 154*, IOP Publishing Ltd, 1997, p. 298.

[107] A. Lawrence-Douglas, "Ionisation Effects for Laser-Plasma Interactions by Particle-in-Cell Code," PhD thesis, University of Warwick, 2013, p. 161.

[108] J. Posthumus, A. Giles, M. Thompson, and K. Codling, "Field-ionization, Coulomb explosion of diatomic molecules in intense laser fields," *J. Phys. B At. Mol. Opt. Phys.*, vol. 29, p. 5811, 1996.

[109] A. Dubey, K. Antypas, M. K. Ganapathy, *et al.*, "Extensible component-based architecture for FLASH, a massively parallel, multiphysics simulation code," *Parallel Comput.*, vol. 35, no. 10-11, pp. 512–522, 2009. arXiv: `0903.4875`.

[110] J. J. MacFarlane, I. E. Golovkin, and P. R. Woodruff, "HELIOS-CR – A 1-D radiation-magnetohydrodynamics code with inline atomic kinetics modeling," *J. Quant. Spectrosc. Radiat. Transf.*, vol. 99, pp. 381–397, 2006.

[111] J. Fan, E. Parra, and H. M. Milchberg, "Resonant Self-Trapping and Absorption of Intense Bessel Beams," *Phys. Rev. Lett.*, vol. 84, pp. 3085–3088, 2000.

[112] L. Spitzer, *Physics of Fully Ionized Gases*. Interscience Publishers, 1967, p. 133.

[113] J. D. Huba, "NRL Plasma Formulary," 2013.

[114] C Huang, C Zhang, K. A. Marsh, C. E. Clayton, and C Joshi, "Dynamics of Optical-field Ionized Helium Plasmas," Advanced Accelerators Concepts, 2018.

[115] G. Taylor, "The formation of a blast wave by a very intense explosion," *Proc. R. Soc. A*, vol. 201, pp. 159–174, 1950.

[116] W. J. Macquorne Rankine, "On the Thermodynamic Theory of Waves of Finite Longitudinal Disturbance," *Philos. Trans. R. Soc.*, vol. 160, no. January, pp. 277–288, 1870.

[117] G. J. Hutchens, "Approximate cylindrical blast theory: Near-field solutions," *J. Appl. Phys.*, vol. 77, no. 7, pp. 2912–2915, 1995.

[118] R. J. Shalloo, "Hydrodynamic Optical-Field-Ionized Waveguides for Laser Plasma Accelerators," PhD thesis, University of Oxford, 2018.

[119] S. M. Hooker and C. Webb, *Laser Physics*. Oxford: Oxford University Press, 2010, p. 472.

[120] T. R. Clark and H. M. Milchberg, "Time- and Space-Resolved Density Evolution of the Plasma Waveguide," *Phys. Rev. Lett.*, vol. 78, no. 12, pp. 2373–2376, 1997.

[121] N. Bourgeois, J. Cowley, and S. M. Hooker, "Two-pulse ionization injection into quasilinear laser wakefields," *Phys. Rev. Lett.*, vol. 111, no. 15, p. 155 004, 2013.

[122] B. Hidding, G. Pretzler, J. B. Rosenzweig, *et al.*, "Ultracold electron bunch generation via plasma photocathode emission and acceleration in a beam-driven plasma blowout," *Phys. Rev. Lett.*, vol. 108, no. 3, p. 035 001, 2012.