

Laser Wakefield Acceleration in the Hydrogen-Filled Capillary Discharge Waveguide

Thomas Rowlands-Rees
St. John's College, Oxford



Submitted for the degree of Doctor of Philosophy
Hilary Term 2009

Supervised by
Prof. Simon M. Hooker

Clarendon Laboratory
University of Oxford
United Kingdom

Abstract

This thesis presents the results of a detailed study into the generation of relativistic electron beams through laser-wakefield acceleration (LWFA). In particular the benefit of using the hydrogen-filled capillary discharge waveguide, or similar plasma channels, in the LWFA scheme is examined.

A thorough experimental study of LWFA in a waveguide was undertaken using the Astra laser at the Rutherford-Appleton laboratory. The waveguide was demonstrated to guide a $34\ \mu\text{m}$ laser spot over 5 cm. It was found that using 0.6 J laser pulses of duration 45 fs, electrons could be accelerated to 200 MeV if the waveguide was employed. This was attributed to the ability of the waveguide to enable electron acceleration at plasma densities as low as $2 \times 10^{18}\ \text{cm}^{-3}$. By contrast, the density threshold for LWFA for the same laser, but without a plasma channel, was observed to be $\sim 20 \times 10^{18}\ \text{cm}^{-3}$, and the upper limit on electron energy was 110 MeV. An investigation was undertaken to explain how the presence of a plasma channel enabled acceleration at lower plasma densities, and through a combination of theoretical work and careful examination of experimental results it is concluded that the plasma channel allowed more efficient use of the entire laser pulse energy.

A regime was identified during the Astra experiments where the generation of electron beams was relatively stable, with 125 MeV monoenergetic beams detected for 80% of shots, but also extremely sensitive to the state of the plasma channel. Detailed analysis of Raman satellites in the transmitted laser spectra and additional interferometry experiments shows that the stable regime was strongly connected with a drop in the ionization level of the plasma on the waveguide axis. That work is published in Physical Review Letters.

The role of self-focussing was also examined in detail. Simple equations were derived to describe self-focussing along the temporal profile of the laser pulse. From these an expression for a self-focussed matched spot size was derived which was found to agree with the value that had previously been inferred empirically from 3D-simulations. The equations were then used to estimate the form of the Astra laser pulse in a self-focussed ‘steady-state’ inside the waveguide, and the resulting estimates of electron energy were found to be in excellent agreement with the experimental data. In particular, the cross-over between linear and non-linear scaling was successfully predicted.

The work contained in this thesis suggest several ways that the energy gain and stability of LWFA can be improved to fulfil the potential of plasma based accelerators to offer a cheap, compact alternative to existing large-scale facilities.



$$\text{Re}(x^2) \rightarrow I = x^2 \quad \therefore \frac{dT}{dz} = 2x \frac{dx}{dz} \quad \frac{d^2T}{dz^2} = 2 \left[\left(\frac{dx}{dz} \right)^2 + x \frac{d^2x}{dz^2} \right]$$

Re(x²)

I am very sure that this is going to be fucking horrible...

I found the scribbles on page ii in my lab-book. If I didn't know they were by me I would say that they came from a person who was in a pretty dark place — not in the sense of the lights being off and the curtains shut (that's just my limited artistic ability expressing itself), I mean that *psychologically* that person was in a dark place. To some extent its academic — if you knew someone spent their time sat with the lights off, using a biro to scribble pictures of faces and write swears to themselves you'd say they were in a dark place in *both* senses of the words (unless they were blind, in which case it would be inspiring).

Actually, it wasn't so bad. Laser physics is like playing with expensive lego, and I got paid to do it. The last four years have been a lot of fun and I am privileged to have had the opportunity to work in research. There are many people for whose support, whether direct or indirect, I am grateful.

Simon, I still remember our 'potential student meets potential supervisor interview'. I hadn't done any research into what kind of field I wanted to work in and I didn't really understand anything you were saying but what I *did* notice was your ironed shirt and your shiny cufflinks and I thought "Here's a guy who's not gonna let you down when it really counts.". That, in case you ever wondered, was the reason I did a DPhil with you¹, and it turns out you *can* judge a book by its cover. So thank you for all your guidance, patience and time.

The Hooker collective, you all brought something different to the table. In the early days we had Matt's optimism, Claire's (healthy) cynicism and Nic's massive belt. Kevin joined, and whenever we went for lunch I would secretly hope someone would pick a fight with us just so I could see him use his Taekwondo. Charlotte and Tom Robinson came and added culture and stature and fancy phones. When Nicolas, Luis, Andi and Svet arrived they all made me feel clever in different ways: Nicolas by maintaining the charade that I was bilingual; Luis by looking deep in thought however dumb whatever I was saying was; Andi by always coming and asking me questions; Svet simply because having a colleague called Svetoslav Bajlekov made me feel like I was doing real science. A special mention should go to the two guys who had to put up with me the most. Tony Gonsalves, we were a pretty kick-ass team. I learned loads from you not just about physics but about everything to do with making stuff happen. Next time we meet up² you can be *my* wingman. And Tom Ibbotson, to whom I tried to be like Tony was to me. Thank you for all your help and hard work especially with the interferometry and getting to the bottom of the unexpected mysteries of the waveguide. I would have gone mad trying to make it there on my own.

A huge chunk of this thesis is based on an experiment performed at the Rutherford Appleton Laboratory in 2006. Thank you to everyone who worked on all the various aspects of that endeavour: my fellow lab monkeys — Stefan Kneip, Stuart Mangles, Jordan Gallacher, Enrico Brunetti, Chris Murphy and that handsome fellow Christos Kamperidis (see page 38); the bosses who made it all possible — Peter Norreys, Karl Krushelnick, Zulf Najmudin and Dino Jaroszynski; the laser team

¹Seriously I had so many other offers. Loads and loads.

²In Trinidad (hint).

— Chris Hooker, Klaus, Oleg, Jodie, Steve and Edwin; the link guys — Peta Foster and Matthew Streeter; and lastly Darren Neville, Phil Rice, Paul Holligan and Richard Bickerton, without whom I would have ended my days in either a raging hydrogen fireball or a high voltage discharge (or some combination of the two). Thanks also to Raoul Trines for helping me understand the results.

The people I worked with at MPQ, particularly Stefan Karsch, Jens Osterhoff, Antonia Popp, Szuszi Major and Matthias Fuchs: your lab is like something out of James Bond, thanks for the opportunity to come and and experience it too. Did you know I used to secretly (try to) breakdance in the bunker when none of you were looking? What can I say, I was sleep deprived and delusional, I thought I had the moves. Thanks for all the fun, I really enjoyed myself.

Thankyou to all the Clarendon physicists with whom I shared office space, tea and music libraries and without whose friendship the process would have been considerably less fun. There's too many of you to name but I will just say that the thing I miss most about my DPhil is the people. Of the Clarendon staff I would like to give special thanks to Terry Fletcher and George Dancer. Sorry I was always taking the piss George.

Thankyou to Brigitte Cros and Gianluca Gregori for the time spent reading through my thesis. I apologise if for the boring parts.

Thankyou to the friends who have put up with me when I'm up and picked me up when I'm down. Special mention should go to Sarah and Jee for having to tolerate my boyish odours and post-computer-game-victory gloating for three years. Lilian thankyou for always being able to make me laugh and being able to put a cheery spin on life. Shout out to Josh, Sonia, Andrew and Rebecca the remnants of my undergraduate social life. We're still young. Kenneth Gyang thankyou for giving me inspiration to get this thing finished so we can work together. Same goes to Bambo Soyinka. Same to Jairus. Chichi Uzuegbu thankyou for always calling to see how it was going. Kelly Nottage, I haven't forgotten our big plans. Debbie Ostermann thanks for your long emails, I always had fun replying. Thankyou to the Korean Society Football team for taking on a foreign signing. Camille and Imran thankyou for all the enjoyable evenings round your place. Aisha Phoenix thankyou for being such a good listener. Ben Fletcher it used to put things in perspective to hear you saying how badly things were going (and then you went and finished before me, which put things back in their real perspective again). Amita Deb, you are the cleverest person on the planet, and a dude to boot. So thanks for that. Busi, thankyou for your love.

Thankyou to my three companions in the final stages of my write up: Max your positive outlook and lithe physique are an inspiration to everyone who meets you; Coco I don't know what it is you want but I'm pretty sure I don't have it; Daisy we all miss you.

Lastly thankyou to my family. Will for always setting such a high standard, and Mum and Dad for too many things to list here, but I'm grateful for everything.

Contents

Abstract	i
Appreciation	iii
Thesis outline	ix
Role of the author	x
List of publications	xi
List of figures	xiii
List of tables	xiv
List of symbols	xviii
1 Introduction to laser-wakefield acceleration	1
1.1 Introduction	1
1.1.1 Conventional accelerators	1
1.1.2 Plasma accelerators	2
1.2 Basic Theory of LWFA	3
1.2.1 Oscillations in a cold unmagnetised plasma	3
1.2.2 The ponderomotive force	4
1.2.3 The laser wakefield accelerator	6
1.3 Non-linear effects	7
1.3.1 Self-focussing	7
1.3.2 Photon acceleration	8
1.3.3 Pulse compression	11
1.3.4 The Raman instability	12
1.4 Limitations to energy gain	13
1.4.1 Dephasing	13
1.4.2 Laser-depletion	13
1.4.3 Diffraction	15
1.4.4 Wave breaking and beam loading	15
1.5 Energy scaling	16
1.5.1 Linear regime	16
1.5.2 Bubble regime	17
1.6 Key-challenges for the development of LWFA devices	18

2	Waveguides for LWFA	20
2.1	Step refractive index guiding	20
2.2	Graded index guiding	22
2.2.1	Laser-produced plasma channels	22
2.2.2	Two-pulse ponderomotive guiding	24
2.2.3	Discharge ablation capillaries	25
2.2.4	Fast discharge gas cells	25
2.3	The hydrogen-filled capillary discharge waveguide	26
2.3.1	Design and operation	27
2.3.2	Scaling waveguide parameters	28
2.3.3	Recent results using the hydrogen-filled capillary discharge waveguide	29
2.4	Conclusion	31
3	Astra experiment setup and methods of analysis	32
3.1	Experimental setup	32
3.1.1	The Astra laser	32
3.1.2	The waveguide	35
3.1.3	Transmitted laser diagnostics	35
3.1.4	Electron diagnostics	36
3.1.5	Automated control and data collection system	37
3.2	Analysis of diagnostic output	39
3.2.1	Processing of CCD images	39
3.2.2	Estimating fluence	43
3.2.3	Laser spectra	45
3.2.4	Electron spectra	45
3.3	Statistical methods of analysis	50
3.3.1	Averaging for varying t	50
3.3.2	Averaging for varying n_e	50
3.3.3	Beam probability	53
4	Results from Astra experiments	58
4.1	Measurements of the transmitted laser pulse	58
4.1.1	Guiding	58
4.1.2	Transmitted laser spectra	63
4.2	Electron beam measurements	67
4.2.1	Electron energy	68
4.2.2	Electron beam profiles	71
4.3	Correlating measurements of transmitted radiation to LWFA	73
4.3.1	Guiding and electron acceleration	73
4.3.2	Laser red-shifts and LWFA	77
4.4	Capillary durability	77
4.5	Conclusion	84

5	The role of the plasma channel in LWFA	85
5.1	Extending interaction length	85
5.2	Self-focussing regime	86
5.3	Temporal effects	89
5.3.1	Quasistatic equations	89
5.3.2	Deriving a self-focussed matched spot size	90
5.3.3	The quasi-steady state	92
5.3.4	Electron energy scaling	94
5.3.5	Comparison to flat profile	96
5.4	Radial effects	97
5.4.1	Self-focussed propagation code	97
5.4.2	Propagation of an initially Gaussian laser spot	98
5.4.3	Adapting to non-ideal conditions	100
5.5	Conclusion	104
6	Electron trapping at low plasma densities	108
6.1	Diagnosing plasma conditions within the waveguide	110
6.1.1	Interferometry measurements	110
6.1.2	Raman analysis of long pulse spectra	112
6.2	Results of electron density measurements	114
6.2.1	Comparison to existing scaling laws	114
6.2.2	On axis ionization levels	115
6.3	Un-ionized Species	120
6.3.1	Wall hypothesis	120
6.3.2	Recombination hypothesis	121
6.3.3	Gas-slot discharge model	122
6.4	Effect of un-ionized species on laser wakefield acceleration	125
6.4.1	Electrons born within the wakefield	125
6.4.2	Increased wakefield amplitude due to an ionization front	128
6.4.3	Thermal effects	129
6.4.4	Current-switch effects	130
6.5	Conclusion	133
7	Conclusions	135
7.1	Summary	135
7.2	Short term future work	136
7.3	Multi-GeV electron accelerators	136
7.3.1	Scaling to high energy	137
7.3.2	Developing an ideal waveguide	140
A	Imaging the waveguide exit	144

B Chromatic aberration	149
B.1 Quantifying chromatic aberration	149
B.2 Deducing the real transmitted fluence profile	151
B.2.1 Correcting measurements of the laser spectrum	151
B.2.2 Estimating the image plane fluence for an assumed W_o	152
B.2.3 Estimating W_o from the measured fluence profile	152
C Ionization thresholds for elements of interest	155

Thesis outline

The work of this thesis focusses on the use of the hydrogen-filled capillary discharge waveguide plasma for laser wakefield acceleration. The majority of the work relates to experiments performed at the Astra laser facility between June and August 2006.

Chapter 1 The importance of accelerators and the demand for an alternative new technology, are discussed. Key results charting the fast development of laser-wakefield acceleration are discussed. The basic principle behind LWFA is outlined and the theory is then described in detail. The current challenges facing the field are discussed.

Chapter 2 Different techniques and mechanisms for guiding intense lasers through plasmas are compared. The hydrogen-filled capillary discharge waveguide is introduced.

Chapter 3 The experiments at the Astra laser facility are described in detail. The methods for processing the large volume of recorded data from the experimental diagnostics are outlined, and the statistical methods, some of which were developed specifically to analyse the data of this experiment, are explained.

Chapter 4 The results of the Astra experiments are presented in full. Important trends are highlighted, and unexpected results are identified.

Chapter 5 Driven by the experimental observations of Chapter 4, a theoretical investigation is made into the role of the plasma channel in the self-focussing regime. An analytical expression for the self-focussed matched spot size is derived, and this is used to describe the transition between the linear and non-linear regimes. A simple 3d-simulation code was written, and it is used to show that the energy of a driving laser pulse is used more efficiently for LWFA in plasma channels.

Chapter 6 Following from unexpected experimental observations in Chapter 4, the specific conditions that lead to the trapping of electrons in a plasma wake at low density are investigated. Additional experiments and analysis are described. It is found that below a certain density the trapping of electrons is dependent on the ionization levels within the waveguide plasma. Several possible reasons for this are discussed.

Chapter 7 The results of the thesis are summarised, and new directions for future work are outlined.

Role of the author

The majority of experimental measurements presented in this thesis result from a collaborative three month campaign that took place at the Astra laser facility from June to August 2006. The author was assigned the role of target area operator and led the set up and running of the experiments. The analysis of the resulting data was undertaken by the author, and this included developing original statistical methods to process the large volume of data that was collected. All data presented in the thesis were analysed by the author.

The interferometry experiments that measured the waveguide conditions for the Astra experiment were performed with the help of T. Ibbotson and analysed by the author using software written by A. J. Gonsalves.

The theoretical work presented in this thesis is entirely the author's own work. This includes writing the simplified 3D-simulation code that was used to model the self-focussing of a non-ideal laser pulse in a plasma channel.

The author also contributed significantly to two experiments not covered in this thesis. Firstly, the author helped to set up the first transverse interferometry experiments of the hydrogen-filled capillary waveguides, and was responsible for collecting the data from which scaling laws were deduced for the axial density and matched spot of the waveguide. That work is published in *Physical Review Letters* and described fully in the thesis of A. J. Gonsalves. Secondly, the author spent a total of three months at the Max-Planck-Institut für Quantenoptik in Garching, Germany as part of an experimental collaboration with the group of Stefan Karsch. That work led to publications in *New Journal of Physics* and *Physical Review Letters*, and is described fully in the thesis of Jens Osterhoff.

List of publications

In addition to the work listed below, it is expected that the work of this thesis should form the basis of at least one more publication. A further publication based on the work collaborating with the Max-Planck-Institut is currently in preparation.

From the Astra experiments:

T. P. Rowlands-Rees, C. Kamperidis, S. Kneip, A. J. Gonsalves, S. P. D. Mangles, J. G. Gallacher, E. Brunetti, T. Ibbotson, C. D. Murphy, P. S. Foster, M. J. V. Streeter, F. Budde, P. A. Norreys, D. A. Jaroszynski, K. Krushelnick, Z. Najmudin, and S. M. Hooker. **Laser-driven acceleration of electrons in a partially ionized plasma channel.** *Physical Review Letters*, 100(10), 2008.

From interferometry experiments:

A. J. Gonsalves, T. P. Rowlands-Rees, B. H. P. Broks, J. J. A. M. van der Mullen, and S. M. Hooker. **Transverse interferometry of a hydrogen-filled capillary discharge waveguide.** *Physical Review Letters*, 98(2), 2007.

From collaboration with the Max-Planck-Institut:

S. Karsch, J. Osterhoff, A. Popp, T. P. Rowlands-Rees, and Zs. Major, M. Fuchs, B. Marx, R. Hörlein, K. Schmid, L. Veisz, S. Becker, U. Schramm, B. Hidding, G. Pretzler, D. Habs, F. Grüner, F. Krausz and S. M. Hooker. **GeV-scale electron acceleration in a gas-filled capillary discharge waveguide.** *New Journal of Physics* 9(11), 2007.

J. Osterhoff, A. Popp, Zs. Major, B. Marx, T. P. Rowlands-Rees, M. Fuchs, M. Geissler, R. Hörlein, B. Hidding, S. Becker, E. A. Peralta, U. Schramm, F. Grüner, D. Habs, F. Krausz, S. M. Hooker, and S. Karsch. **Generation of stable, low-divergence electron beams by laser-wake field acceleration in a steady-state-flow gas cell.** *Physical Review Letters*, 101(8):085002, 2008

List of Figures

1.1	Examples of accelerators using RF technology.	2
1.2	A schematic representation of laser wakefield acceleration.	6
1.3	Different examples of photon acceleration in a plasma	10
1.4	The etching of a laser pulse front.	14
1.5	The bubble regime.	17
1.6	Recent results showing improved electron beam stability.	19
2.1	Step index guiding in hollow glass capillaries	21
2.2	The igniter-heater scheme	23
2.3	Two-pulse ponderomotive guiding.	24
2.4	Fast discharge gas-cell	26
2.5	The hydrogen-filled capillary discharge waveguide	27
2.6	Scaling laws and interferometry measurements of waveguide parameters	29
2.7	Demonstration of guiding in the hydrogen filled capillary-discharge waveguide	30
2.8	Acceleration of electrons to 1 GeV in the hydrogen-filled capillary discharge waveguide	30
3.1	Schematic layout of the Astra experiments.	33
3.2	Focal scan of the Astra laser.	33
3.3	Electron spectrometer calibration image with electron beam data superimposed. . .	38
3.4	CCD image at different stages of processing.	40
3.5	Background subtraction from a CCD array.	42
3.6	Comparing vacuum fluence profiles for different NDs.	44
3.7	Methods of deconvolving the electron spectrum.	46
3.8	Applying transmission biased averaging to two samples of real data.	51
3.9	Applying transmission biased averaging across a range of densities.	54
3.10	Comparison of methods for calculating $P(t)$	55
3.11	The form used for making Bayesian estimates of $P(t)$	57
4.1	Short pulse guiding data for [$L = 50$ mm, $D = 250$ μ m, $p = 160$ mbar, $V = 25$ kV]. . .	59
4.2	Guiding data for varying L and D	60
4.3	Transmission rise and fall delays for varying L and D	62
4.4	Spectra of transmitted laser pulses	64
4.5	Spectral broadening versus plasma density.	66

4.6	Overview of electron beam generation for the Astra experiments	67
4.7	A selection of electron spectrometer lanex screen images.	69
4.8	Electron energy of monoenergetic beams versus density	70
4.9	Electron beams recorded by the electron pointing monitor.	72
4.10	Transmission T and beam probability $P(t)$ versus delay t for $n_e \gtrsim 6.5 \times 10^{18} \text{ cm}^{-3}$	74
4.11	Transmission T and beam probability $P(t)$ versus delay t for $n_e \lesssim 6.5 \times 10^{18} \text{ cm}^{-3}$	75
4.12	Sensitivity of transmission and electron generation to delay against density	76
4.13	Electron signal and red-limit Λ^R versus density n_e	78
4.14	$\times 1$ magnification scans for two used capillaries.	80
4.15	$\times 5$ magnification scans of the used $[L = 15 \text{ mm}, D = 200 \mu\text{m}]$ capillary.	81
4.16	$\times 5$ magnification scans of the used $[L = 33 \text{ mm}, D = 200 \mu\text{m}]$ capillary.	82
4.17	$\times 10$ magnification scans near the $[L = 33 \text{ mm}, D = 200 \mu\text{m}]$ capillary exit.	83
5.1	Scaling L_d and P/P_c against density	87
5.2	Laser pulses and wakefields in the quasi-steady state.	93
5.3	Scaling for the quasi-steady state	95
5.4	Fraction of laser energy trapped within a plasma wavelength radius.	99
5.5	Propagation of an intense Gaussian laser in flat plasmas and channels	101
5.6	Measured input and output transverse fluence profiles of the long pulse.	102
5.7	Simulated initial laser pulse evolution using the real laser profile and quartic channel.	105
5.8	Propagation of an intense laser with the Astra profile in flat plasmas and channels.	106
6.1	Electron spectra for $[L = 15 \text{ mm}, D = 200 \mu\text{m}, p = 250 \text{ mbar}]$	109
6.2	Schematic diagram of transverse interferometry experiments.	111
6.3	Long pulse spectra and fluence profiles for $[L = 15 \text{ mm}, D = 200 \mu\text{m}, p = 250 \text{ mbar}]$	113
6.4	Comparison of n_e^I and n_e^R to predictions.	114
6.5	Measured data as a function of t for $[L = 15 \text{ mm}, D = 200 \mu\text{m}, p = 600 \text{ mbar}]$	116
6.6	Measured data as a function of t for $[L = 15 \text{ mm}, D = 200 \mu\text{m}, p = 400 \text{ mbar}]$	117
6.7	Measured data as a function of t for $[L = 15 \text{ mm}, D = 200 \mu\text{m}, p = 250 \text{ mbar}]$	118
6.8	Measured data as a function of t for $[L = 33 \text{ mm}, D = 200 \mu\text{m}, p = 80 \text{ mbar}]$	119
6.9	Schematic diagrams of ‘clean’ and ‘unclean’ discharges.	122
6.10	Diagram of the Butler waveguide. After [1].	124
6.11	Trajectories of axial electrons ionized for varying ξ	127
6.12	Probability a plasma electron does not collide whilst being accelerated.	132
7.1	Optimum acceleration parameters for future experiments.	138
7.2	Matched spots of different waveguide schemes against plasma density.	141
A.1	Lens system for imaging the waveguide exit.	145
B.1	The design of a doublet lens.	150
B.2	Estimates of the real exit fluence profile and corrected laser spectrum.	154
C.1	Above Barrier Ionization Thresholds for selected elements	157

List of Tables

1.1	Conventional interpretations of spectral measurements after laser-plasma interaction	9
3.1	Astra laser parameters	34
3.2	Comparison of methods for calculating P_0 and S	56
4.1	Electron beam pointing and divergence.	71
7.1	Optimum electron accelerations for a selection of existing laser systems.	139
A.1	Magnitude of errors for the target method and quick method.	147
C.1	Above Barrier Ionization Thresholds for selected elements	156

List of Symbols

- β_z $\beta_z c$ is the longitudinal velocity of electrons as part of plasma wave motion, page 90
- $\Delta\phi$ Difference in phase, page 110
- ϵ_0 Permittivity of free space, page 3
- η Refractive index, page 4
- γ Relativistic factor, $(1 - v^2/c^2)^{-1/2}$, page 5
- γ_g Relativistic γ factor for $v = v_g$, page 13
- λ Laser wavelength, page 8
- Λ^R, Λ^B The red, blue wavelength limit, page 45
- λ_p Plasma wavelength, page 6
- μ_0 Permeability of free space, page 3
- $\nu_{ei}, \nu_{ee}, \nu_{eh}$ Collision frequency of electrons off ions, electrons and hydrogen atoms, page 129
- ω Angular frequency of electromagnetic radiation, page 4
- ω_p Plasma frequency, page 4
- Φ Potential, page 89
- τ Laser pulse length, full-width half maximum (unless referring to data averaging), page 3
- τ When referring to data averaging: gaussian width of averaging window, page 50
- \underline{A}_\perp Transverse component of the vector potential, page 5
- ξ Convenient coordinate for quasistatic calculations given by $\xi = z - ct$, page 89
- A Capillary cross-sectional area, page 62
- a_0 Amplitude of \underline{a} , page 5
- $\underline{\mathcal{B}}$ Magnetic field, page 3
- c The speed of light, page 4

D	Capillary diameter, page 20
dn	Perturbation to the plasma density, page 90
$\underline{\mathcal{E}}$	Electric field, page 3
E	Energy, page 5
e	Electron charge, page 4
ΔE	Electron energy gain over one dephasing length, page 16
\mathcal{E}_{lwb}	The longitudinal wave-breaking limit, page 16
$\underline{\mathcal{E}}_0$	Electric field amplitude, page 5
\mathcal{E}_{wb}	The linear wave-breaking limit, page 16
$E(x)$	Relationship between electron energy and displacement (x) for the electron spectrometer, page 37
\mathcal{E}_z	Longitudinal component of the wakefield, page 16
\underline{F}	Force, page 5
h	Convenient expression, $h = 1 - \eta^2$, page 90
I	Laser intensity, page 5
$I(x)$	For deconvolving electron spectra: the measured signal on the lanex screen, page 47
\underline{J}	Current density, page 3
\underline{k}	Wave vector of plane wave solutions to Maxwell's equations, page 4
L	Capillary length, page 27
L_c	Length of the capillary between the gas slots, page 62
$L_d, L_d^{\text{lin}}, L_d^{\text{bub}}$	Dephasing length, the superscripts indicate the linear ($a_0 \ll 1$) and bubble ($a_0 > 2$) regimes, page 13
$L_{\text{pd}}, L_{\text{pd}}^{\text{mnl}}, L_{\text{pd}}^{\text{bub}}$	Pulse-depletion length, the superscripts indicate the moderately non-linear ($a_0 \sim 1$) and bubble ($a_0 > 2$) regimes, page 14
m_e	Electron mass, page 4
m_p	Proton mass, page 5
n_c	The critical plasma density for which $\omega_p = \omega$. For the majority of work of this thesis $\lambda = 812 \text{ nm}$, corresponding to $n_c = 1.75 \times 10^{21} \text{ cm}^{-3}$, page 4

n_e	Unperturbed plasma density (in the case of a plasma channel, the unperturbed <i>axial</i> density), page 4
n_e^I	Axial electron density measured by interferometry, page 112
n_e^R	Axial electron density measured by analysis of Raman satellites in the long pulse spectrum, page 112
n_{H_2}	Initial density of hydrogen molecules inside the capillary, page 28
$n(r)$	Transverse plasma channel density profile, page 91
$O(x)$	For deconvolving electron spectra: the 'underlying' electron spread, page 47
p	Electron momentum, page 5
P	Laser power, page 8
p	Pre-discharge hydrogen pressure inside the capillary, page 27
P_0	Peak value of $P(t)$, page 53
P_c	The critical laser power for relativistic self-focussing, page 8
$P(t)$	The probability of generating an electron signal or monoenergetic beams (denoted by superscripts) as a function of delay t , page 53
P_x	For deconvolving electron spectra: the instrument function, page 47
R_b	Bubble radius in the bubble regime, page 18
R_c	Resistance of the capillary between the gas slots as a function of delay t , page 123
R_g^*	Resistance of the 'gas slot path' after ionization, page 125
R_c^*	Resistance of the capillary between the gas slots once the hydrogen has been ionized, page 123
r_e	Classical electron radius, page 28
R_g	Resistance of the 'gas slot path' as a function of delay t , page 123
S	The sensivity of $P(t)$ to delay, page 53
s	CCD signal, page 41
t	Time (except for when referring to the capillary discharge), page 3
t	When referring to the capillary discharge: the time delay between the onset of the discharge current flow and the arrival of the laser pulse at the capillary entrance, page 28
T	Fraction of laser input energy transmitted through the capillary, page 36
t_r, t_f	The delays at which $T(t)$ rises and falls above and below 60% of its peak value, page 61

\underline{v}	Velocity of plasma electrons, page 4
V	Voltage across discharge capacitor, page 27
v_d	The velocity at which the leading edge of a self-guiding laser pulse diffracts away, page 96
v_{etch}	Speed at which the laser pulse front is etched back by interacting with the plasma in the bubble regime, page 14
v_g	Group velocity, page 4
v_ϕ	Phase velocity, page 4
W	Transverse $1/e^2$ spot radius, page 6
W_0	W at focus, page 15
W_c	The matched spot size of a plasma channel, page 22
W_m	Matched spot size (general expression including effects of both channel, relativistic self-focussing and plasma response), page 90
W_s	The matched spot size due to self-focussing effects, page 18
z	Space coordinate in the direction of laser propagation, page 6
Z_R	The Rayleigh range, page 15

Chapter 1

Introduction to laser-wakefield acceleration

1.1 Introduction

For over 100 years the acceleration of charged particles has been at the heart of science and technology. The earliest accelerators are cathode ray tubes, whose development at the end of the 19th century led directly to the discovery of the X-ray in 1895 and the electron in 1896. Accelerator physics has existed, by name, since the late 1920s, and in its first decade led to the development of the Van der Graff generator, the Cyclotron and the Betatron all of which gave rise to significant advances in physics. Additionally these techniques for accelerating particles all have their own scientific, medical, industrial and commercial applications, for which they are still in use today. The importance of accelerator technology cannot be overstated.

1.1.1 Conventional accelerators

State-of-the-art accelerating technology applies the principle of phase stability, which was developed in the mid 1940s. In these radio-frequency (RF) accelerators, particles pass through tubes called cavities, to which an alternating voltage is applied. The frequency and phase of the alternating voltage is controlled so that the particle bunches only experience an accelerating field. High energies can be reached as particles pass through many cavities, being accelerated repeatedly by the same voltage. There are two types of RF accelerator: linear accelerators (linacs), where particles are accelerated in a straight line and synchrotrons, where magnets are used to make the particles follow a circular path¹.

An example of a linac is the Stanford Linear Collider (SLAC), which is shown in figure 1.1(a). The maximum energy of particles from a linac is the length of the accelerator multiplied by the accelerating gradient. An example of a synchrotron is the Diamond Light Source, shown in figure 1.1(b). For these accelerators the limit to the possible energy of the particles is not the available length for acceleration (which can be infinite in the ring structure) but the strength of the magnets

¹In fact the path is that of a many sided polygon, with alternating straight sections and bends, forming an effective circle.



Figure 1.1: Examples of accelerators using RF technology.

SLAC image courtesy of the U.S. Geological Survey. Diamond Light Source image courtesy of Diamond Light Source Ltd.

for bending the electrons, and the power loss due to ‘synchrotron radiation’ as particles follow the circular path. The effect of radiation loss is far greater for lighter particles meaning that in general linacs are advantageous for accelerating electrons and positrons.

Although RF technology goes some way to solving the problem of electrical breakdown (which is more serious in static field-devices) it is still a problem above $\sim 50 \text{ MV m}^{-1}$. This fundamentally limits the accelerating gradient that can be achieved using RF technology and the only way to accelerate electrons to higher energy is to increase the acceleration length. The length of the SLAC linac is 3.2 km (giving $\sim 50 \text{ GeV}$ electrons) and the circumference of the LEP synchrotron at CERN was 27 km (giving $\sim 100 \text{ GeV}$ electrons). The construction of larger accelerators is prohibitively expensive. Even for applications requiring lower energy electrons the space and cost is a problem — for example the Diamond Light Source, which uses undulators to generate X-rays from electrons accelerated to approximately 3 GeV in a synchrotron, occupies the space of approximately 5 football pitches and cost $\sim \$600 \text{ million}^2$ to construct.

1.1.2 Plasma accelerators

The electric field gradients within plasma waves are approximately 4 orders of magnitude greater than those that can be achieved with RF technology. Electrical breakdown is not an issue since the plasma is, in effect, already broken down. A method for harnessing these fields was proposed by Tajima and Dawson in 1979 [2]. They proposed that electrons could be accelerated by the wave that trailed an intense laser pulse as it propagated through a plasma. At that time lasers did not exist to provide sufficiently intense pulses for such ‘laser wakefield acceleration’ (LWFA). However rapid developments in laser technology, most notably the development of chirped pulse amplification in 1985 [3], have allowed rapid developments in LWFA, particularly in the last five years. A notable breakthrough came in 2004 when three groups used LWFA to generate mono-energetic electron beams with energy $\sim 100 \text{ MeV}$ [4–6]. Two years later electrons were accelerated to 1 GeV, over a length of just 3 cm, using a plasma waveguide [7, 8]. Another experiment in 2006 demonstrated a technique for controlling LWFA electron energy [9], and in 2008 synchrotron radiation was generated by passing LWFA accelerated electrons through an undulator [10]. Simulations suggest that the generation of mono-energetic multi-GeV femtosecond electron beams is a realistic short term

²Obviously the accelerator itself was only a part of the expenditure for this project.

possibility [11].

These results all suggest that in the near term LWFA may allow compact devices, equivalent to facilities such as the Diamond Light Source, at a size and cost affordable to individual universities, laboratories and hospitals. The demand for such ‘lower energy’ facilities is extremely high since they are used for condensed matter physics, biophysics, biochemistry, chemistry and medicine. Aside from cost, LWFA also offers the additional advantage of significantly shorter electron bunches than are possible using RF technology. This may lead to a new level of temporal resolution for X-ray diffraction experiments, allowing time-resolved study of chemical reactions on a molecular scale.

At present the development of LWFA devices that could match or exceed the energies of high energy particle physics accelerators is a long term goal. Using particle-driven wakefield acceleration (PWFA), where an electron beam is used to drive the plasma wave, rather than a laser pulse [12], it has been shown that plasma waves can, in principle, be used to accelerate particles to energies at the particle physics frontier. In 2007 it was reported that passing the 42 GeV beam at SLAC through a meter long plasma led to the acceleration of some electrons within the beam to 80 GeV [13] — effectively the front of the electron bunch excited a plasma wave, and the electrons at the back of the bunch were accelerated as a result [13]. The electrons that were accelerated gained the same amount of energy in a meter of plasma as in the previous 3.2 km of RF accelerator. This demonstrates the potential of plasma accelerators, whether driven by lasers or particles, for high energy physics.

1.2 Basic Theory of LWFA

In this section the key physical concepts are outlined to describe the excitation of the wakefield that follows an intense pulse as it passes through a plasma. Only the LWFA scheme, where the laser pulse length τ is of order the plasma frequency ω_p^{-1} (defined below) is discussed. Self-modulated laser wakefield acceleration (SMLWFA) [14] and plasma beat wave acceleration (PBWA) [15] are related schemes for driving plasma wakes with intense laser pulses that were useful before the laser technology had been developed to generate such short laser pulses. This is no longer the case and so, for the sake of simplicity and brevity, these historic schemes are not described.

1.2.1 Oscillations in a cold unmagnetised plasma

By defining a plasma as cold it is assumed that on the timescales under examination collisions can be ignored. Since ions are much heavier than electrons their motion is also ignored. According to the Maxwell-Faraday equation and Ampère’s law, an electric field $\underline{\mathcal{E}}$ and magnetic field $\underline{\mathcal{B}}$ can be described by

$$\underline{\nabla} \times \underline{\mathcal{E}} = -\frac{\partial \underline{\mathcal{B}}}{\partial t} \quad (1.1)$$

$$\underline{\nabla} \times \underline{\mathcal{B}} = \mu_0 \underline{J} + \mu_0 \epsilon_0 \frac{\partial \underline{\mathcal{E}}}{\partial t} \quad (1.2)$$

where \underline{J} is the current density. If the plasma is assumed to be (approximately) homogenous then the general solution can be written as a sum of plane wave solutions of form $\underline{\mathcal{E}}_{\underline{k}} e^{i(\underline{k} \cdot \underline{r} - \omega t)}$, $\underline{\mathcal{B}}_{\underline{k}} e^{i(\underline{k} \cdot \underline{r} - \omega t)}$

and $\underline{J}_k e^{i(\underline{k}\cdot\underline{r}-\omega t)}$.

In the non-relativistic case the velocity \underline{v} of plasma electrons can also be written as a sum of solutions of the form $\underline{v}_k e^{i(\underline{k}\cdot\underline{r}-\omega t)}$ with $\underline{v}_k = -i \frac{e \underline{\mathcal{E}}_k}{\omega m_e}$ where e and m_e are the mass and charge of an electron. This means that $\underline{J} = -en_e \underline{v}$ can be written in the same form as $\underline{\mathcal{E}}$ and $\underline{\mathcal{B}}$ with $\underline{J}_k = in_e \frac{e^2 \underline{\mathcal{E}}_k}{\omega m_e}$. Here n_e represents the unperturbed plasma electron density. Substituting these expressions into equations (1.1) and (1.2) gives

$$\underline{k} \times \underline{\mathcal{E}}_k = \omega \underline{\mathcal{B}}_k \quad (1.3)$$

$$\underline{k} \times \underline{\mathcal{B}}_k = \mu_0 n_e \frac{e^2 \underline{\mathcal{E}}_k}{\omega m_e} - \omega \mu_0 \epsilon_0 \underline{\mathcal{E}}_k \quad (1.4)$$

Inspection of these equations shows that a cold unmagnetised plasma can sustain two types of waves, those for which \underline{k} and $\underline{\mathcal{E}}_k$ are parallel, and those for which they are perpendicular.

For the first case, the left hand side of equation (1.3) is equal to zero, meaning that $\underline{\mathcal{B}}_k = 0$. Equation (1.4) can then be rearranged to give

$$\omega^2 = \omega_p^2 = \frac{e^2 n_e}{m_e \epsilon_0} \quad (1.5)$$

These are longitudinal electrostatic waves oscillating at the plasma frequency ω_p , which depends only on n_e . They are known as plasma waves or Langmuir waves and it is their longitudinal electric fields that accelerate electrons in LWFA. Since ω_p is independent of \underline{k} the phase velocity $v_\phi = \omega/k$ is undefined.

Other modes, for which \underline{k} and $\underline{\mathcal{E}}_k$ are perpendicular, equation (1.3) can be substituted into (1.4) to give

$$k^2 = \frac{\omega^2 - \omega_p^2}{c^2} \quad (1.6)$$

This then, is the dispersion relation for electromagnetic waves propagating in a plasma. For $\omega < \omega_p$, k^2 becomes negative, so waves below that cut off frequency will not propagate through the plasma. Similarly equation (1.6) can be inverted to give an expression for a critical plasma density $n_c(\omega) = \frac{\omega^2 m_e \epsilon_0}{e^2}$ above which the plasma cannot sustain waves of frequency ω . For the work in this thesis the laser wavelength was 812 nm, for which $n_c = 1.75 \times 10^{21} \text{ cm}^{-3}$. For $\omega > \omega_p$ the phase velocity of these electromagnetic waves is $v_\phi = \omega/k = c(1 - \omega_p^2/\omega^2)^{-1/2}$ meaning that the refractive index of a plasma is $\eta = (1 - \omega_p^2/\omega^2)^{1/2}$. The group velocity is $v_g = \partial\omega/\partial k = c(1 - \omega_p^2/\omega^2)^{1/2}$, so laser pulses propagate more slowly through denser plasmas. LWFA operates in the ‘underdense’ regime where $\omega_p^2/\omega^2 \ll 1$ and $v_g \approx c$.

1.2.2 The ponderomotive force

Electrons exposed to laser radiation oscillate in the electric field. The equation for this ‘quiver motion’ in a laser field of frequency ω and amplitude $\underline{\mathcal{E}}_0$ is

$$\frac{d\underline{p}}{dt} = -e \underline{\mathcal{E}}_0 \cos[\omega t] \quad (1.7)$$

where \underline{p} is the electron momentum. Assuming that the field is linearly polarised³ along the x -axis and is sufficiently strong that the quiver motion is dominant it is straightforward to show that the time-averaged expression for p^2 is

$$\langle p^2 \rangle = \frac{1}{2} \left(\frac{e\mathcal{E}_0}{\omega} \right)^2 \quad (1.8)$$

This means that an electron in a laser field has a time averaged energy given by

$$\langle E^2 \rangle = m_e^2 c^4 + \frac{1}{2} \left(\frac{e\mathcal{E}_0}{\omega} \right)^2 c^2 \quad (1.9)$$

The relativistic form of these equations has been maintained because the electron motion in an intense laser field is relativistic. Where there is a gradient in laser intensity there is a gradient in electron energy, and if there is a gradient in energy there is a force given by $\underline{F} = -\underline{\nabla}E$, so that

$$\underline{F} = - \frac{e^2 c^2}{4\omega^2 \sqrt{m_e^2 c^4 + \frac{1}{2} \left(\frac{e\mathcal{E}_0}{\omega} \right)^2 c^2}} \underline{\nabla} \mathcal{E}^2 \quad (1.10)$$

The force describes motion of the electrons averaged over oscillations at the laser frequency. The normalized vector potential \underline{a} is given by

$$\underline{a} = \frac{|e|\underline{A}_\perp}{m_e c^2} = \frac{|e|\mathcal{E}_0}{\omega m_e c} \cos[\omega t] \quad (1.11)$$

where \underline{A}_\perp is the transverse component of the vector potential. Conventionally the amplitude of \underline{a} is written as a_0 and is related to the laser intensity I by the expression

$$a_0 = \frac{|e|\mathcal{E}_0}{\omega m_e c} = \frac{|e|}{\omega m_e c} \sqrt{\frac{2I}{\epsilon_0 c}} \quad (1.12)$$

Note that writing equation (1.8) in terms of a_0 yields the result $\langle p^2 \rangle = \frac{1}{2} a_0^2 m_e^2 c^2$ which can be rearranged to give $\langle \gamma^2 \rangle = 1 + \frac{a_0^2}{2}$, where γ is the relativistic mass factor. This shows that as $a_0 \gtrsim 1$ the motion of the electrons in the laser field is relativistic. Substituting into (1.10) gives

$$\underline{F} = - \frac{m_e c^2}{4\sqrt{1 + a_0^2/2}} \underline{\nabla} a_0^2 \quad (1.13)$$

A full electromagnetic derivation for the ponderomotive force can be found in [16] but the above arguments are sufficient to demonstrate that an intense laser pulse exerts a force on electrons as it propagates through a plasma. Protons and positive ions are considerably heavier and so will be moved less by the ponderomotive force — not only that, but equation (1.10) shows that the force itself is also less on heavier particles. This means that the displacement of protons by the ponderomotive force is approximately $(m_e/m_p)^2 \approx 3 \times 10^{-7}$ less than that for electrons. Therefore ions are treated as stationary when considering the ponderomotive force.

³To date all LWFA experiments have used a linearly polarised field. For simplicity expressions are derived using linear polarisation. It can be shown that the expressions for the ponderomotive force are the same for other polarisations when written in terms of the intensity (not the field amplitude).




Figure 1.2: A schematic representation of laser wakefield acceleration.

A laser pulse (solid black lines) propagates from left to right at velocity v_g . The plasma density (grey fill) oscillates due to the ponderomotive force of the laser pulse, causing a wake to trail the pulse. A co-propagating bunch of electrons (green) gains energy from the accelerating phase of the wake's longitudinal electric field (dotted line).

1.2.3 The laser wakefield accelerator

From the above it is clear that as a laser pulse propagates through a plasma it will excite the longitudinal electrostatic oscillation modes of the plasma as well as the transverse electromagnetic modes. Since the initial perturbation to the plasma is caused by the ponderomotive force at the front edge of the laser pulse, the phase velocity of this plasma wave must equal the laser's group velocity — thus this wave, the ‘wakefield’ appears to trail the laser pulse very close to the speed of light. Its wavelength is defined as the plasma wavelength $\lambda_p = \frac{2\pi}{k_p} = \frac{2\pi c}{\omega_p}$.

The principle of a laser wakefield accelerator is demonstrated schematically in figure 1.2. Because the wake is travelling close to the speed of light co-propagating relativistic particles remain within the accelerating phase⁴. Rather than experiencing the oscillating fields seen in the stationary frame these particles only experience the accelerating fields. Therefore electrons co-propagating with the wakefield are accelerated and gain energy.

For simplicity a square pulse is considered in the 1D linear limit ($a_0^2 \ll 1$ and the transverse pulse radius $W \gg \lambda_p$). Equation 1.13 can be used to estimate the amplitude of the electrostatic waves excited by the leading edge of this pulse. If the plasma is perturbed by the leading edge of a laser pulse for which the intensity rises from 0 to a_0^2 in a time $\Delta t \approx c\Delta z$ then the momentum gain of plasma electrons is

$$F\Delta t \approx -\frac{m_e c^2}{4} \frac{a_0^2}{\Delta z} \Delta t \approx -\frac{m_e c}{4} a_0^2 \quad (1.14)$$

The initial perturbation to the momentum of plasma electrons leads to longitudinal plasma oscillations of the form $p_z = F\Delta t \cos[\omega_p t] = -\frac{m_e c}{4} a_0^2 \cos[\omega_p t]$. The force driving these oscillations is the electric field of the plasma wave, which is given by

$$\mathcal{E}_z = \frac{1}{e} \frac{dp_z}{dt} = \frac{\omega_p m_e c}{e} \frac{a_0^2}{4} \quad (1.15)$$

⁴This is a simplification. Dephasing between the electron and the wakefield and the laser are discussed in section 1.4.

If the ponderomotive force of the trailing edge of the plasma is also taken into account, and the pulse length is given by $c\tau = \lambda_p/2$ then this amplitude is doubled. For $a_0 = 0.5$, $\tau = \lambda_p/2c = 45$ fs the field is 15 GV m^{-1} . As shall be shown, the electric field in LWFA can easily exceed this, yet this is still far greater than the maximum possible accelerating gradient that can be achieved in RF accelerators.

1.3 Non-linear effects

Section 1.2 described the reaction of a plasma to a laser pulse. The laser pulse is also affected by the plasma, in the form of self-focussing, temporal compression of the laser pulse and frequency shifts to the laser spectrum. Such interactions affect the intensity distribution of the laser pulse (in spatial and temporal dimensions) which in turn affects the plasma wake generated by the laser pulse. These changes to the plasma further influence the evolution of the laser pulse, so that these processes can feed back on themselves becoming highly non-linear.

1.3.1 Self-focussing

Generally self-focussing (and self-guiding) refers to a combination of relativistic and ponderomotive focussing effects that occur when the power of the laser pulse is sufficiently high. Most LWFA experiments to date have relied on self-focussing to counteract diffraction of the laser pulse and extend the acceleration length [4, 5, 17]. The mechanisms of self-focussing are outlined briefly below. The role of self-focussing is considered in detail in chapter 5.

Relativistic self-focussing

The derivation of the refractive index in section 1.2 ignored relativistic effects when relating \underline{J} to $\underline{\mathcal{E}}$. If these are included, the expression for the refractive index becomes dependent on the laser intensity. By substituting (1.1) into (1.2) and considering the electromagnetic modes gives

$$-k^2 \underline{\mathcal{E}}_{\underline{k}} = i\mu_0 \omega \underline{J}_{\underline{k}} - \omega^2 \mu_0 \epsilon_0 \underline{\mathcal{E}}_{\underline{k}} \quad (1.16)$$

Substituting a non-relativistic expression for \underline{J} leads to equation (1.6). In the relativistic case one can write

$$\omega^2 - c^2 k^2 = \frac{i\omega \underline{J}_{\underline{k}}}{\epsilon_0 \underline{\mathcal{E}}_{\underline{k}}} \quad (1.17)$$

In the relativistic equations of motion $-e\underline{\mathcal{E}} = \frac{d\underline{p}}{dt}$ meaning that $-e\underline{\mathcal{E}}_{\underline{k}} = i\omega \underline{p}_{\underline{k}}$ where \underline{p} is the electron momentum, given by $\underline{p} = m_e \underline{v} \gamma$. Since $\underline{J}_{\underline{k}} = -en_e \underline{v}$, this leads to the expression $\underline{J}_{\underline{k}} = \frac{-ie^2 n_e}{\omega m_e \gamma} \underline{\mathcal{E}}_{\underline{k}} = \frac{i\epsilon_0 \omega_p^2}{\omega \gamma} \underline{\mathcal{E}}_{\underline{k}}$. Substituting this expression into (1.17) gives

$$\omega^2 - c^2 k^2 = \frac{\omega_p^2}{\gamma} \quad (1.18)$$

The dispersion relation has been modified by the relativistic mass of the electrons. Making the substitution $\gamma \rightarrow \langle \gamma \rangle = (1 + a_0^2/2)^{1/2}$ gives a refractive index [18]

$$\eta = \frac{kc}{\omega} = \left(1 - \frac{1}{(1 + a_0^2/2)^{1/2}} \frac{\omega_p^2}{\omega^2} \right)^{1/2} \quad (1.19)$$

Relativistic effects mean that at high intensity the refractive index depends on intensity. An intense beam experiences a higher refractive index on axis where the intensity is peaked than towards the edges where the intensity is lower. The refractive index gradient can guide or focus the laser, depending on its strength, for reasons discussed in detail in chapter 2. This is called relativistic self-guiding, or relativistic self-focussing. The critical power P_c at which the relativistic self-focussing effect counteracts diffraction is $P_c = \frac{8\pi\epsilon_0 m_e^2 c^5}{e^2} \frac{n_c}{n_e} \approx 17.4 \frac{n_c}{n_e}$ [GW] [18]. Note that the critical power is reduced as the density of the plasma increases.

Ponderomotive guiding

The refractive index of a plasma depends on its density, so the displacement of plasma by the ponderomotive force of a passing laser pulse leads to variations in the refractive index. Unfortunately, at the front of the laser pulse there is an intensity-dependent increase in the plasma density that exactly cancels the intensity-dependent relativistic contribution to the refractive index [19, 20]. This means that even when the laser power $P > P_c$ the front of the laser pulse diffracts as if it was of low intensity, exposing a fresh slice of the temporal profile of the laser to the unperturbed plasma. At later phases in the plasma wave created by the laser pulse the plasma profile is inverted, meaning that the plasma wave makes a focussing contribution to the refractive index, enhancing self-focussing.

In LWFA experiments without external guiding structures the spot size is typically chosen so that the Rayleigh range of the laser is roughly equal to the dephasing length for electron acceleration (both of these terms are defined below) [4]. This means that the front of the laser pulse does not diffract significantly, and the back of the laser pulse is focussed tightly to a radius $\sim \lambda_p$. In a sense the back of the pulse can self-guide because of the presence of the relatively wide front of the pulse.

1.3.2 Photon acceleration

The refractive index of a plasma is a function of both the radiation intensity and the electron density. This means that as a laser pulse propagates through a plasma the resulting changes to the refractive index co-propagate. Gradients in refractive index change the wavelength of radiation. This can be demonstrated by considering two phase fronts separated by 2π , separated by the laser wavelength λ . If there is a difference in the phase velocity at each of these phases the distance between them changes at a rate Δv_ϕ . The change of distance between two phase fronts corresponds to a change in wavelength. Assuming that $\frac{d^2 v_\phi}{dz^2} \ll \frac{1}{\lambda} \frac{dv_\phi}{dz}$, it can be written that

$$\Delta v_\phi = \lambda \frac{dv_\phi}{dz} = \frac{d\lambda}{dt} \quad (1.20)$$

Observation	Interpretation
No spectral shifts.	The pulse is not interacting with the plasma significantly.
Red shift only.	The pulse is driving a wakefield, $c\tau \lesssim \lambda_p/2$ as in figure 1.3(a).
Red and blue shifts.	The pulse is driving a wakefield, $c\tau \sim \lambda_p$, as in figure 1.3(b).
Blue shift only only.	The pulse is ionising species as it propagates, as in figure 1.3(c).
New peaks in the spectrum, all at an even spacing.	The peaks are the result of the Raman instability, described in section 1.3.4. The separation between the peaks corresponds to a frequency shift of ω_p . This process requires that $c\tau \gg \lambda_p$, as in figure 1.3(d).

Table 1.1: Conventional interpretations of spectral measurements after laser-plasma interaction

Using the substitution $v_\phi = c/\eta$ and considering the changing wavelength as light travels along that gradient it is found that

$$\frac{d\lambda}{dt} = -\frac{\lambda c}{\eta} \frac{d\eta}{dz} \quad (1.21)$$

In the case of a time-stationary gradient the change in wavelength does not correspond to a change in frequency. Equation (1.21) simply describes the change in wavelength as light crosses the boundary between media of different refractive indices.

On the other hand a time-varying refractive index alters the frequency of light. For $\frac{\partial\eta}{\partial z} = 0$ and $\frac{\partial\eta}{\partial t} \neq 0$ the phase velocity is constant across space meaning that the wavelength is unaffected. So with a change in phase velocity there must be a corresponding change in laser frequency. Since $\omega = \frac{2\pi c}{\eta\lambda}$ then it is straightforward to show that

$$\frac{d\omega}{dt} = -\frac{\omega}{\eta} \frac{d\eta}{dt} \quad (1.22)$$

Spatial variations in the refractive index lead to changes in wavelength and temporal variations lead to changes in frequency. In the lab frame the refractive index gradients that are associated with a wakefield vary in both space and time. The changes in frequency that result correspond to changes in photon energy as the laser pulse does work to drive the wakefield. This process is generally described as ‘photon acceleration’ (in the case of photons gaining energy) and ‘photon deceleration’ (in the case of photons losing energy) [21, 22]. Photon deceleration is the mechanism by which the laser pulse loses energy to the wakefield, and in the photon-kinetic model the ponderomotive force results from momentum change of red-shifted photons.

As well as being a necessary element to the theory of LWFA in describing energy conservation, the spectral shifts in the driving laser pulse are a useful diagnostic [23]. The processes that lead to changes in the laser spectrum are summarised in table 1.1 and figure 1.3. Firstly, if the driving pulse length is shorter than half a plasma wavelength then the laser radiation propagates entirely in the initial red-shifting phase of the wakefield. The result is that some of the pump laser spectrum is transferred to longer wavelengths. As a general rule, if there is no laser energy transferred to longer wavelengths in the measured spectrum, it is unlikely that a wakefield of significant amplitude was being driven by the laser pulse.

When the pulse extends into the accelerating phase of the wakefield the back of the pulse is blueshifted. The energy gained at the back of the pulse originates from the front of the pulse,

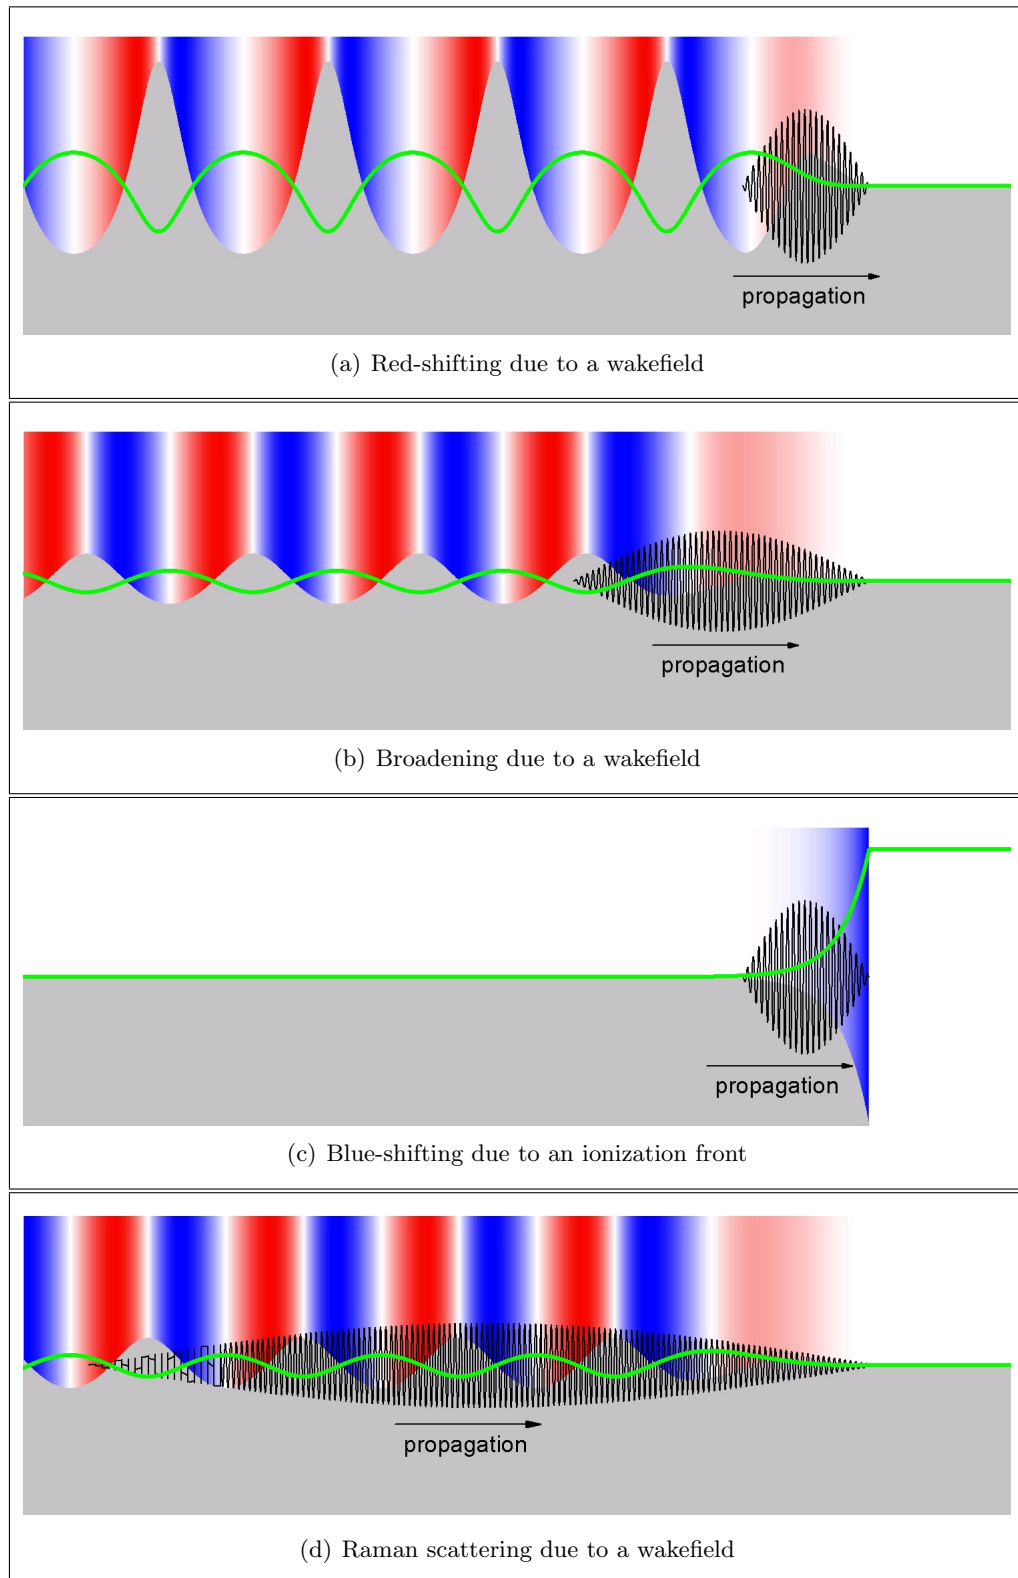


Figure 1.3: Different examples of photon acceleration in a plasma. The figures depict a laser pulse (black line) travelling from left to right, causing co-propagating changes to the plasma density (grey fill) and refractive index (green line). The refractive index gradient is represented on a colour scale ranging from red (photon decelerating gradient) to blue (photon accelerating gradient) through white (zero gradient).

which drives the wakefield, and is red-shifted. As a result the laser spectrum broadens. This case is depicted in figure 1.3(b).

It is also possible for the pulse to be exclusively blue-shifted by photon acceleration if it is also ionising the plasma through which it propagates. The ionization front, which co-propagates with the laser pulse, has an accelerating refractive index gradient. Despite the overall frequency increase of the light energy is conserved because some photons are absorbed in the ionization process. Generally an overall shift of the spectrum towards blue is interpreted as evidence that the laser is creating electrons through ionization as it propagates. This is often referred to as an ‘ionization blueshift’. Predicting the expected blue-shift is difficult due to instabilities at the ionization front [24].

These processes affecting the spectrum of the laser pulse as it propagates through a plasma are not mutually exclusive, for example a laser pulse can ionise a gas and drive a wakefield in the resulting plasma. The fact that more than one of these processes occurs simultaneously makes quantitative interpretation of the spectral observations difficult, particularly if a spectrum is analysed in isolation without other data with which it can be compared. Despite this, changes to the laser spectrum as the pulse propagates through plasma are useful as a qualitative indication that certain processes are occurring, or not. This is especially the case when comparing the spectra of many different laser shots for varying conditions.

1.3.3 Pulse compression

The refractive index gradients that cause distortion to the laser spectrum also distort the temporal profile of the laser pulse. In fact it would be impossible to affect the spectrum without affecting the pulse profile since both are related to one another through fourier transforms.

Physically pulse compression is the result of temporal variations of the group velocity of the laser pulse. For a plasma $v_g = \eta c$, meaning that

$$\frac{dv_g}{dt} = c \frac{d\eta}{dt} \quad (1.23)$$

A pulse travelling from left to right compresses where $\frac{dv_g}{dt}$ is positive. This corresponds to the phases of the wake where the pulse is red-shifted. Since the front of a laser pulse is always in the red-shifting phase of the wakefield the leading edge steepens, increasing the ponderomotive force. Although the blue-shifting phase might disperse the pulse initially, once those photons are at a higher frequency the refractive index they experience is higher meaning that they can catch the main part of the pulse, leading to compression. As light passes over an ionization front, it is accelerated to a higher frequency and reaches a lower density plasma so that its group velocity becomes faster than the light that is following behind. An ionization front must therefore stretch the front of a laser pulse. In the case where $\lambda_p \gg c\tau$ the pulse overlaps with many cycles of the plasma wave, as illustrated in figure 1.3(d). Pulse compression means that, in the intense case, the beam is broken longitudinally into a series of pulses, each compressed in its respective plasma bucket. The self-modulated laser wakefield (SM-LWFA) scheme relies on this process to achieve short pulses of sufficiently high intensity [2]. When this was initially proposed the laser technology

did not exist to produce pulses as short as a plasma wavelength. This is an extreme case of the Raman instability, which is described in section 1.3.4.

Since the plasma density, laser frequency and laser intensity all affect the refractive index pulse compression effects are very difficult to quantify analytically. Generally simulations are used to evaluate their significance.

1.3.4 The Raman instability

The Raman instability (also known as stimulated Raman scattering) is a resonant interaction between the laser field and the plasma wave. In terms of particle scattering an incident photon from the driving laser beam decays to become a scattered photon plus a plasmon (a plasma wave quasi-particle):

$$\omega_0 \rightarrow \omega_1 \pm \omega_p \quad (1.24)$$

$$\underline{k}_0 \rightarrow \underline{k}_1 + \underline{k}_p \quad (1.25)$$

It should be clear that, from the point of view of energy conservation, the frequency of the scattered photon ω_1 should differ from the pump frequency ω_0 by ω_p . Forward stimulated Raman scattering (where \underline{k}_0 , \underline{k}_1 and \underline{k}_p are co-directional) manifests itself experimentally by the presence of peaks (known as Raman satellites or Stokes/anti-Stokes peaks) ω_p either side of the central frequency in the spectrum of a laser pulse that has propagated through a plasma.

In terms of a wave picture the scattering results from the overlap of the driving laser radiation with the oscillating refractive index modulations due to the plasma wave. An important aspect of this process is its inherent instability: as radiation is scattered to frequencies ω_p either side of the central frequency the electromagnetic waves beat against one another resulting in modulations to the intensity profile that overlap with the plasma wave. This increases the amplitude of the plasma wave and thus the mechanism feeds back on itself. An analysis of the instability shows that as a laser pulse propagates through a plasma the amplitude of the Raman satellites grows exponentially with a form $e^{\Gamma t}$ where [16]

$$\Gamma \approx \frac{\omega_p^2 a_0}{2\sqrt{2}\omega_0} \quad (1.26)$$

In the strongly non-linear case pulselets with length λ_p/c form as the Raman instability grows. In the earlier self-modulated laser wakefield accelerator scheme this effect was key for generating short laser pulses within the plasma.

For the work of this thesis a short pulse laser was available, meaning that the Raman instability was not necessary for the acceleration of electrons. However the Raman instability is useful as a plasma diagnostic. Measurement of the Raman satellites can be used to estimate the plasma frequency, and from this, the plasma density. It has also been suggested that the temperature of the plasma can be measured by comparing the frequency of the backscattered satellites and the forward scattered satellites [25]. That is because the forward scattered satellites are separated from the central frequency by ω_p , which is a function of plasma density only, but the back scattered satellite is separated by $\sqrt{\omega_p^2 + 3k^2 v_{th}^2}$ where v_{th} is the thermal electron velocity.

1.4 Limitations to energy gain

Unfortunately the maximum possible energy gain of a laser wakefield accelerator is limited by more than just the length of plasma and the electric field. The factors that limit the energy gain are outlined below.

1.4.1 Dephasing

Relativistic electrons remain in the accelerating phase of a co-propagating wake-field because the wake phase velocity and electron velocity are both approximately equal to the speed of light. The relativistic factor relating to v_g , the group velocity of the driving laser pulse is given by $\gamma_g = \omega/\omega_p$. To give an example, for $n_e = 5 \times 10^{18} \text{ cm}^{-3}$ and $\lambda = 800 \text{ nm}$, $\gamma_g = 19$. For those conditions electrons with energy greater than 10 MeV propagate faster than the laser pulse. Accelerated electrons will eventually outrun the accelerating field and enter the decelerating phase of the wake, at which point they lose energy.

If the velocity of the electrons being accelerated is taken to be c and the phase velocity of the wakefield is taken to be the group velocity of the laser pulse, then the difference in velocity between the two is

$$\Delta v = c - c\sqrt{1 - \left(\frac{\omega_p}{\omega}\right)^2} \approx \frac{c}{2} \left(\frac{\omega_p}{\omega}\right)^2 \quad (1.27)$$

In a linear wakefield, the length of the accelerating phase is $\lambda_p/2$, so the time that an electron can remain in the accelerating phase is given by $\lambda_p/2\Delta v$. This gives an expression for the linear ‘dephasing length’⁵

$$L_d^{\text{lin}} = c \frac{\lambda_p}{2\Delta v} = \frac{\lambda_p^3}{\lambda^2} \quad (1.28)$$

If electrons continue to interact with the plasma wave beyond the dephasing length they lose energy (or are defocussed).

An expression also exists for L_d^{bub} , which is the dephasing length in the non-linear ‘bubble’ regime (for $a_0 > 2$). This applies to much of the work of this thesis. Reference [11] estimates that

$$L_d^{\text{bub}} = \frac{2\sqrt{a_0} \lambda_p^3}{3\pi \lambda^2} \quad (1.29)$$

The difference between this and the linear expression is due to a change in the effective plasma wavelength, and the fact that the phase velocity of the plasma wave is not equal to the group velocity of the laser pulse [11, 26, 27], due to the etching of the pulse front.

1.4.2 Laser-depletion

Unsurprisingly a laser pulse driving a plasma wave loses energy. If this was not the case energy would not be conserved by LWFA. The laser pulse loses energy through the photon deceleration

⁵There is also a radial electric field within the linear wake that can focus or defocus the accelerating electron bunch to and from the axis. The expression for L_d^{lin} is sometimes quoted as being half of that given in equation (1.28) due to the fact that the wake is only accelerating *and* focussing for a quarter of a plasma wavelength. The importance of the radial forces as an issue depends on the radius of the laser pulse, so this factor was not included in the above derivation.




Figure 1.4: The etching of a laser pulse front.

Results from a 1D simulation with $a_0 = 10$, $\omega/\omega_p \approx 5$ and $\tau = 125/\omega$. The plots show the transverse electric field \mathcal{E} in units of $e/m_e c$ versus $\xi = z - ct$ for $ct =$ (a) 0.08 mm, (b) 0.16 mm (c) 0.24 mm, (d) 0.32 mm. After [27]

mechanism outlined in section 1.3.2. Since the number of photons is conserved [28] a decrease in the mean photon frequency corresponds to the transfer of energy from the pulse to the plasma wave.

For $a_0 \sim 1$ photons exchange energy with the plasma wave throughout the length of the laser pulse [29]. As the laser pulse loses energy the mean value of a_0^2 within the pulse increases [29, 30]. This is because $a_0^2 \propto I/\omega^2$, and $I \propto \hbar|\omega|$ meaning that $a_0^2 \propto 1/\omega$. From equations (1.13) and (1.15) it is clear that the wakefield amplitude depends on a_0^2 rather than I , so as the pulse loses energy the wakefield amplitude grows, thus accelerating the process of pulse depletion. In this ‘moderately non-linear’ regime the laser intensity ‘explodes’, but the process cannot continue indefinitely since there comes a point where no energy is left. The length scale over which the explosion of intensity and depletion of energy occurs is given by [29]

$$L_{\text{pd}}^{\text{mnl}} \approx \frac{L_{\text{d}}^{\text{lin}}}{2a_0^2} \quad (1.30)$$

Although there are different expressions quoted elsewhere for the depletion length in this regime [11], all have the same qualitative result, which is that in the limit where $a_0 \ll 1$, $L_{\text{pd}}^{\text{mnl}} > L_{\text{d}}^{\text{lin}}$. Although at first glance it might appear beneficial that the pulse depletion length is longer than the dephasing length, it in fact shows that the laser energy is not being used efficiently. It also means that the length of the plasma must be chosen carefully to terminate the interaction after a dephasing length so that electrons are not decelerated once they have reached their maximum energy.

In the bubble regime ($a_0 > 2$) the description of pulse depletion is simpler. As shall be discussed in section 1.5.2, there is a cavity behind the laser pulse front from which almost all electrons are evacuated [31, 32]. As a result the only interaction between the laser and the plasma occurs at the front of the laser pulse, and the depletion is restricted to the leading edge. This manifests itself as an ‘etching’ of the pulse front [27], from where photons are strongly decelerated. This is demonstrated in figure 1.4, which is the result of 1D simulations $a_0^2 \gg 1$, taken from [27]. It has been found for both 1D and 3D simulations that for $a_0 > 2$ the front of the pulse etches back at a velocity of $v_{\text{etch}} \approx c\lambda^2/\lambda_p^2$ [11, 27]. The time it takes for a pulse of length τ to be etched away is

$c\tau/v_{\text{etch}}$, meaning that the bubble regime depletion length is [11]

$$L_{\text{pd}}^{\text{bub}} \approx \frac{\lambda_{\text{p}}^2}{\lambda^2} c\tau = \frac{3\pi}{2\sqrt{a_0}} \frac{c\tau}{\lambda_{\text{p}}} L_{\text{d}}^{\text{bub}} \quad (1.31)$$

In the non-linear regime it is generally the case that $L_{\text{pd}}^{\text{bub}} \sim L_{\text{d}}^{\text{bub}}$ meaning that the energy of the laser pulse is used efficiently, and the constraint that the length of plasma should be chosen to match L_{d} is relaxed.

1.4.3 Diffraction

For efficient LWFA the driving laser pulse is usually focussed to a small spot, so that the driving intensity is increased. It has also been shown that charge of the accelerated bunch that can be supported by a wakefield is maximised when the laser spot radius W is similar to λ_{p} [33]. The propagation, and diffraction of a laser beam with a Gaussian radial intensity profile of the form $I(r, z) = I(0, z)e^{-2r^2/W(z)^2}$, propagating in the z -direction and focussed to a minimum radius W_0 at $z = 0$ is described simply by

$$W(z) = W_0 \sqrt{1 + (z/Z_{\text{R}})^2} \quad (1.32)$$

where $Z_{\text{R}} = \pi W(0)^2/\lambda$ is the Rayleigh range. The beam maintains its Gaussian radial profile, but as it spreads out due to diffraction its peak intensity drops, according to the equation

$$I(0, z) = \frac{I(0, 0)}{1 + (z/Z_{\text{R}})^2} \quad (1.33)$$

The result of equation (1.33) is that an unguided beam can only drive a wakefield for lengths $\sim Z_{\text{R}}$. Using $W \approx \lambda_{\text{p}}$ it can be shown that the diffraction-limited acceleration length is

$$L_{\text{dif}} = Z_{\text{R}} = \frac{\pi\lambda_{\text{p}}^2}{\lambda} = \frac{\pi\lambda}{\lambda_{\text{p}}} L_{\text{d}}^{\text{lin}} \quad (1.34)$$

This means that in general $L_{\text{dif}} \ll L_{\text{d}}$, so diffraction of the driving laser pulse strongly limits the energy gain from LWFA. To some extent self-guiding can offset this effect. However it is notable that monoenergetic beams from LWFA were observed for the first time when W was chosen so that $Z_{\text{R}} \approx L_{\text{d}}$ [4, 5], suggesting that self-guiding alone was not enough to maintain the axial intensity over an entire dephasing length. This was despite the fact that $P \gg P_{\text{c}}$ for these experiments. The use of external guiding channels, which are discussed in Chapter 2, has been shown to increase the energy that can be achieved through laser wakefield acceleration [6–8] for a given pump power. They also allow the possibility of extended acceleration at densities for which $P \lesssim P_{\text{c}}$, and at these lower densities a higher energy gain is possible.

1.4.4 Wave breaking and beam loading

There is an upper limit to the electric field gradient of a laser wakefield. If, during wave motion, the velocity of some plasma electrons exceeds the phase velocity of the wake then ‘wavebreaking’ occurs [16] — this is due to the fact that plasma electrons moving at the wave velocity remain in the

accelerating phase of the wakefield, meaning that their motion is no longer part of the plasma wave motion, but that of ‘trapped electrons’. This process is extremely important to LWFA because it means that electrons from the background plasma can become trapped and accelerated, eliminating the need for an external source of electrons (there are a range of technical problems associated with external injection schemes).

However this also limits the maximum achievable accelerating gradient within the plasma wave. That is because the electric field of the trapped electrons suppresses the electric field of the wakefield, which is known as ‘beam-loading’ [33]. If the amplitude of the wakefield is above the wavebreaking threshold it can continue to trap electrons until the field of the trapped bunch is sufficient to suppress the field of the wakefield so that it is (just) below the wavebreaking threshold.

The linear wave-breaking limit to the wakefield amplitude \mathcal{E}_0 , is given by [34]

$$\mathcal{E}_{\text{wb}} = \frac{\omega_{\text{p}}c}{e} \quad (1.35)$$

When non-linear effects are taken into consideration this becomes the longitudinal wavebreaking limit

$$\mathcal{E}_{\text{lwb}} = \mathcal{E}_{\text{wb}} \sqrt{2 \left(\frac{\omega}{\omega_{\text{p}}} - 1 \right)} \quad (1.36)$$

In fact the above description is simplified. For a fully 3D wakefield transverse wavebreaking occurs at lower amplitudes than longitudinal wavebreaking. Electron trapping and beam loading are complex problems which are still the subject of considerable theoretical study [35, 36], meaning that the above description can only be treated as qualitative. However equations (1.12) and (1.36), together with (1.15) which is the expression for the wakefield amplitude, do show that electrons are more easily trapped at higher plasma densities.

1.5 Energy scaling

Having discussed the limitations to the energy gain of a wakefield accelerator, it is possible to predict the maximum energy gain (corresponding to one dephasing length) of a laser wakefield accelerator. This is given by

$$\Delta E = -e \int_0^{L_{\text{d}}} \mathcal{E}_z(z) dz \quad (1.37)$$

where \mathcal{E}_z is the field experienced by the electrons after propagating a distance z from the point at which they were trapped (or injected). Note that in (1.37) it has been implicitly assumed that any problems relating to depletion and diffraction have been overcome, and the energy gain is limited by the more fundamental problem of dephasing.

1.5.1 Linear regime

For $a_0^2 \ll 1$ the plasma wave is sinusoidal, so the value of $|\mathcal{E}_z(z)|^{\text{lin}}$ averaged over one dephasing length is $2\mathcal{E}_z/\pi$, where \mathcal{E}_z is given by equation (1.15). The resulting expression for ΔE_{lin} is

$$\Delta E^{\text{lin}} = -e |\mathcal{E}_z(z)|^{\text{lin}} L_{\text{d}}^{\text{lin}} = m_e c^2 \frac{n_c}{n_e} a_0^2 \quad (1.38)$$




Figure 1.5: The bubble regime.

On-axis electron density in the xy plane from a 3D-PIC simulation of LWFA in the bubble regime. The panels represent propagation after (a) $25c/\omega_p$ and (b) $442c/\omega_p$. The coordinates are given in c/ω_p . Taken from reference [37].

Despite the fact that the electric field of the plasma wave increases with $n_e^{1/2}$ the maximum energy gain from a single stage is greater for low density because the dephasing length scales as $n_e^{-3/2}$, meaning that for fixed laser parameters $\Delta E^{\text{lin}} \propto n_e^{-1}$. Unsurprisingly ΔE^{lin} is greater for higher laser intensities.

1.5.2 Bubble regime

The scaling changes as the laser intensity becomes relativistic. For $a_0 > 2$ the key parameters scale according to the ‘bubble’ regime⁶ description of LWFA [11, 31, 32, 37]. When the ponderomotive force is sufficiently strong electrons are expelled radially from the path of the laser. As a result a cavity forms behind the driving pulse that is virtually free of electrons. This is illustrated by the simulation results of figure 1.5, taken from [37]. In 1.5(b), which is after the laser has propagated $221 \lambda_p$ inside the plasma, a dense electron bunch has been trapped within the evacuated cavity. The cavity is larger than in 1.5(a) due to the electrostatic repulsion of the electron bunch. A principal reason that the bubble regime has attracted interest is that the evacuated cavity is a particularly stable structure that can sustain large numbers of accelerating electrons without losing its form [36].

Deriving scaling laws in the non-linear limit is complicated by the fact that the action of the plasma on the laser pulse needs to be considered as well as the action of the laser pulse on the plasma. The amplitude of the wakefield, and hence the accelerating gradient, depends on the axial intensity of the laser pulse. However this axial intensity is subject to change due to the evolution of the laser pulse, particularly through the process of self-focussing.

⁶Also known as the ‘blowout’ regime.

This problem is confronted in reference [11], where scaling laws in the non-linear regime were derived. The approach of that work is adapted in Chapter 5, so their derivation shall be outlined briefly. It was deduced empirically from simulation that for $a_0 > 2$ a stable self-focussed spot of radius W_s , which is approximately equal to the blowout radius R_b , is given by

$$W_s \approx R_b \approx \frac{2\sqrt{a_0}}{k_p} \quad (1.39)$$

For $W = W_s$ it was also found that the mean accelerating field is $|\mathcal{E}_z(z)|^{\text{bub}} = -(\sqrt{a_0}/2)\mathcal{E}_0$. Their value of L_d^{bub} , which is that given in equation (1.29) is the length over which electrons with velocity c propagate R_b further than the wake structure propagating at a velocity $v_g - v_{\text{etch}} \approx c \left(1 - \frac{3\omega_p^2}{2\omega^2}\right)$. The resulting expression for the energy gain is

$$\Delta E^{\text{bub}} = -e|\mathcal{E}_z(z)|^{\text{bub}} L_d^{\text{bub}} = m_e c^2 \frac{2}{3} \frac{n_c}{n_e} a_0 \quad (1.40)$$

If it is assumed that even where initially $W \neq W_s$ the laser pulse evolves, through self-focussing, towards this state where $W = W_s$ then a more general rule can be written in the form

$$\Delta E^{\text{bub}} = m_e c^2 \frac{4}{3} \frac{n_c}{n_e} \left(\frac{P}{P_c}\right)^{1/3} \quad (1.41)$$

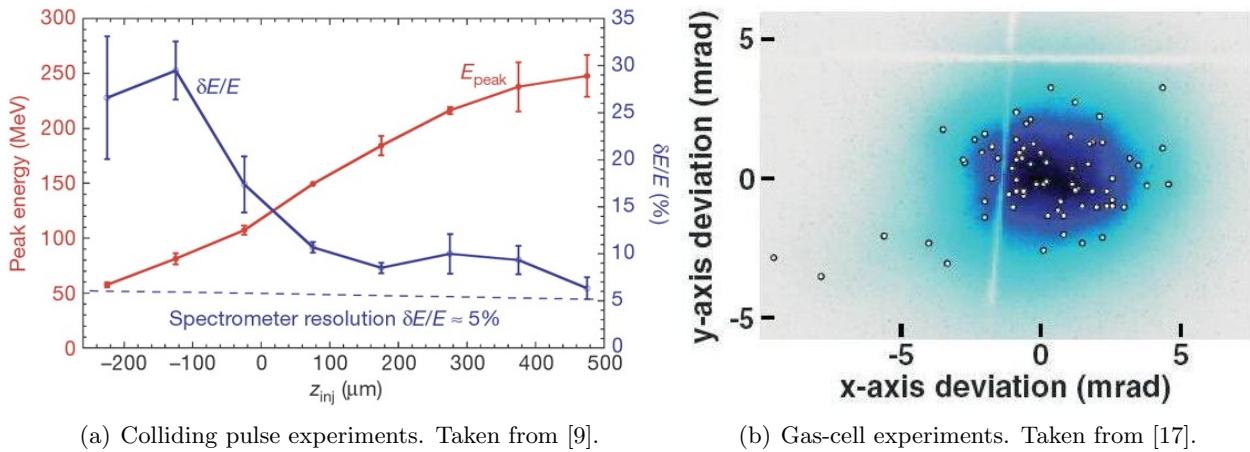
Like in the linear case, ΔE^{bub} increases as n_e decreases. Taking into consideration both P_c and ω_p , ΔE^{bub} scales as $n_e^{-2/3}$.

1.6 Key-challenges for the development of LWFA devices

Improving stability and control

A key requirement for practical applications of LWFA technology is that the output of any LWFA should be reproducible between laser shots. Even when operating outside of the strongly non-linear bubble regime there are non-linear aspects to the laser-plasma interactions meaning that small shot-to-shot fluctuations in the input laser pulse and target plasma (or gas) lead to large fluctuations in the properties of the accelerated electron bunches. Improving the stability of LWFA means improving the reproducibility of the parameters of the accelerated electron beam. Ideally electron beams should be generated on every shot, fluctuations in their pointing should be minimised, and their energy and energy spread should be constant.

The stability of the electron beam has been the focus of recent work. Experiments [9] showed that a counter-propagating laser pulse can stabilise the electron trapping process, as predicted by [38, 39]. Additionally it was seen that by varying the timing between the pump-pulse and the counter-propagating injection pulse the acceleration length, and therefore electron energy, could be controlled, as shown in figure 1.6(a). Control of the injection of electrons into the wakefield also offers the prospect of accelerating electrons at densities lower than the self-trapping threshold which, according to the scaling considerations of section 1.5, would increase the energy gain of the LWFA.



(a) Colliding pulse experiments. Taken from [9].

(b) Gas-cell experiments. Taken from [17].

Figure 1.6: Recent results showing improved electron beam stability.

Figure (a): electron trapping was controlled through use of a counter-propagating pulse [9]. The graph shows the accelerated electron energy versus the time delay between the pump and injection pulse. Figure (b): stable electrons were generated using a gas-cell target [17]. The graph is a false colour image of the summed signal of 74 electron beams and the peak positions of the individual shots (dots).

Another recent experiment, the results of which are shown in figure 1.6(b), demonstrated that the stability of the electron beam could be greatly improved by minimising fluctuations in the target gas by using a gas cell rather than a gas jet [17]. The advantage of this scheme is that it is relatively simple to implement.

Improving single stage energy gain

Equations (1.38) and (1.41) show that the LWFA can accelerate electrons to higher energy at lower plasma density. This is because the dephasing length increases as density decreases. However it is only possible to take advantage of the increase to L_d if the laser can drive a wakefield for that entire length. This is many Rayleigh ranges, and although self-focussing can extend the interaction length to some extent, its effect diminishes as n_e is decreased. If the laser pulse is guided by an external guiding structure, then electron densities can be chosen for which $L_d \gg Z_R$ and $P \ll P_c$. The result would be that for a given laser power, electrons could be accelerated to higher energy over the extended propagation length.

Initial work using laser plasma channels produced by lasers [6] and discharges [7, 8] has shown that the energy gain of a single LWFA stage can be increased significantly if the laser is guided by an external structure. Methods for preventing the diffraction of the laser beam are discussed in chapter 2, and the subject of this thesis is application of one of these, the hydrogen-filled capillary discharge waveguide, to laser wakefield acceleration.

Chapter 2

Waveguides for LWFA

In chapter 1 the need for maintaining the high intensity of a focussed laser beyond the Rayleigh range was established. In this chapter the various ways by which an intense laser pulse can be guided for LWFA are reviewed. One particular device for creating a plasma channel is the hydrogen-filled capillary discharge waveguide which, being the subject of this thesis, is described in detail.

2.1 Step refractive index guiding

A cylindrical waveguide with a core material of refractive index η_{core} surrounded by a cladding material of refractive index η_{clad} can extend the propagation of a laser pulse considerably through reflection at the boundary between the materials. In the case of optical fibres, with a glass core for which $\eta_{\text{core}} > \eta_{\text{clad}}$, total internal reflection can occur at the walls, meaning that laser pulses can be guided over considerable distance without loss. The impact of such waveguides on communications is well known.

For a plasma $\eta_{\text{plasma}} < 1$. Presuming that the cladding is a dielectric material (such as glass) it is only possible to create a step index waveguide for LWFA in which $\eta_{\text{core}} < \eta_{\text{clad}}$. A schematic diagram of such a device, a hollow glass capillary into which gas is flowed through small holes, is shown in figure 2.1(a). Although total internal reflection is not possible the reflectivity is high for grazing incidence and a laser can be guided a considerable distance in such a waveguide if conditions are chosen to minimise the angle of reflection at the capillary wall.

The coupling of different transverse laser modes into the capillary, and their transmission losses at the capillary wall, are well understood theoretically [40, 42]. Experimental results, shown in figure 2.1(b), are in excellent agreement with predictions [41, 43]. Theory predicts that 98% of a transverse Gaussian beam of radius W is coupled into the 0th order mode of the capillary when

$$W = 0.64D/2 \tag{2.1}$$

where D is the capillary diameter. Transmission losses for the n^{th} mode can be characterised as an exponential decay in the beam energy along the length of the capillary with a coefficient k_{Im} (the

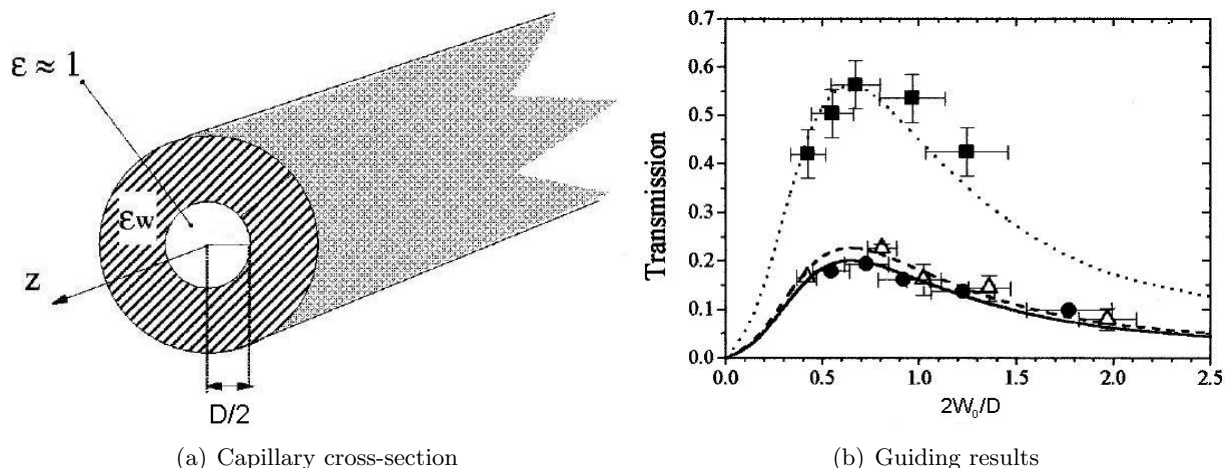


Figure 2.1: Step index guiding in hollow glass capillaries

Figure (a): cross-section of the hollow glass capillary. Taken from [40]. Figure (b): theoretical predictions (lines) against measured values (points) of the transmission through a hollow glass capillary against $2W/D$, for capillary parameters [$L = 8.3$ cm, $D/2 = 35$ μm] (squares and dotted line); [$L = 5.4$ cm, $D/2 = 22.5$ μm] (triangles and dashed line); [$L = 8.3$ cm, $D/2 = 25$ μm] (circles and plain line). Taken from [41].

imaginary part of the wavenumber) given by

$$k_{\text{Im}} \approx -\Im \left[\frac{\alpha_{0,n}^2}{2k^2(D/2)^3} \frac{1 + \epsilon_w}{\sqrt{1 - \epsilon_w}} \right] \quad (2.2)$$

where ϵ_w is the relative permittivity of the cladding material, $\alpha_{0,n}$ is the n^{th} order solution to the equation $J_0(\alpha_{0,n}) = 0$, and J_0 is a Bessel function. The transmission decreases as the capillary becomes narrower. Since $\alpha_{0,n}$ increases with n , transmission losses are minimised for the 0^{th} order mode, which is very close to a Gaussian.

If such a waveguide can be used effectively for LWFA it should be that $1/k_{\text{Im}} \gtrsim L_{\text{d}}^{\text{lin}}$, so that the beam can maintain its intensity over at least a dephasing length. If D is expressed in the form $D/2 = \beta\lambda_p$ and $1/k_{\text{Im}} = L_{\text{d}}^{\text{lin}}$ then it is found that for the 0^{th} mode that

$$\beta = \left(\frac{\alpha_{0,0}^2}{8\pi^2} \frac{1 + \epsilon_w}{\sqrt{\epsilon_w - 1}} \right)^{1/3} \approx 0.45 \quad (2.3)$$

where a value of $\epsilon_w = 2.25$ (corresponding to glass) was used. Combining this with equation (2.1) shows that transmission losses are not an issue (for this particular application) so long as $W \gtrsim 0.3\lambda_p$. The ideal spot size for LWFA applications is $W \sim \lambda_p/2$ [11, 32, 33, 44], so hollow glass capillaries can guide an ideal spot size for LWFA over the required distance.

The great advantage of the hollow glass step index waveguide is that it can guide small spots. However there are several practical issues that need to be overcome for such a waveguide if it is to be applied successfully to LWFA. Experiments have shown that the transmission of the capillary is less for high laser energies [41]. This was attributed to the pre-pulse ablating material at the

capillary entrance which then obstructed the main pulse. With an improved laser focal spot quality and contrast ratio such an effect should not be an issue.

2.2 Graded index guiding

The diverging effects of diffraction can be counteracted if the beam propagates through a medium whose refractive index decreases smoothly away from the axis. This is the case for a ‘plasma channel’ for which the plasma density is at a minimum on axis. Ideally the plasma density is parabolic, with form $n(r) = n_e + \frac{1}{2}n''r^2$. In such a case the transverse profile of the laser pulse remains Gaussian, with the radius $W(z)$ varying as [18]

$$\frac{d^2W}{dz^2} = \frac{\lambda^2}{\pi^2W^3} \left[1 - \frac{W^4}{W_c^4} \right] \quad (2.4)$$

where W_c is the matched spot of the channel, given by

$$W_c = \left[\frac{8n_e}{k_p^2 n''} \right]^{1/4} \quad (2.5)$$

When $W_0 = W_c$ the transverse profile of the pulse does not vary as it propagates through the channel. For $W_0 \neq W_c$ a general solution to equation (2.4) is [45]

$$[W(z)]^2 = \frac{W_0^2}{2} + \frac{W_c^4}{2W_0^2} + \left(\frac{W_0^2}{2} - \frac{W_c^4}{2W_0^2} \right) \cos \left(\frac{2\pi z}{Z_{\text{osc}}} \right) \quad (2.6)$$

where

$$Z_{\text{osc}} = \frac{\pi^2 W_c^2}{\lambda} \quad (2.7)$$

This means that even where the plasma channel and initial spot size are not exactly matched, the plasma channel can still guide the laser pulse, but the transverse size of the spot oscillates during propagation.

The waveguide that is the focus of this thesis utilises such a plasma channel, and is described in section 2.3. There are numerous alternative techniques for creating plasma channels, and several of those that have attracted recent interest are described below.

2.2.1 Laser-produced plasma channels

Early techniques for creating guiding channels used laser radiation to ionise a gas and then heat the resulting plasma through inverse bremsstrahlung. The plasma then expands in a shockwave during which the plasma density is minimised on axis, resulting in a profile suitable for guiding a laser pulse.

Initial methods for creating such a plasma channel used a single beam line focus that both heated and ionised the plasma [46, 47]. This was used to channel laser pulses of peak intensity $\sim 10^{14} \text{ W cm}^{-2}$ for distances of approximately $24Z_R$. Although successful at demonstrating the potential of plasma channels to guide laser pulses there were limits to this method stemming

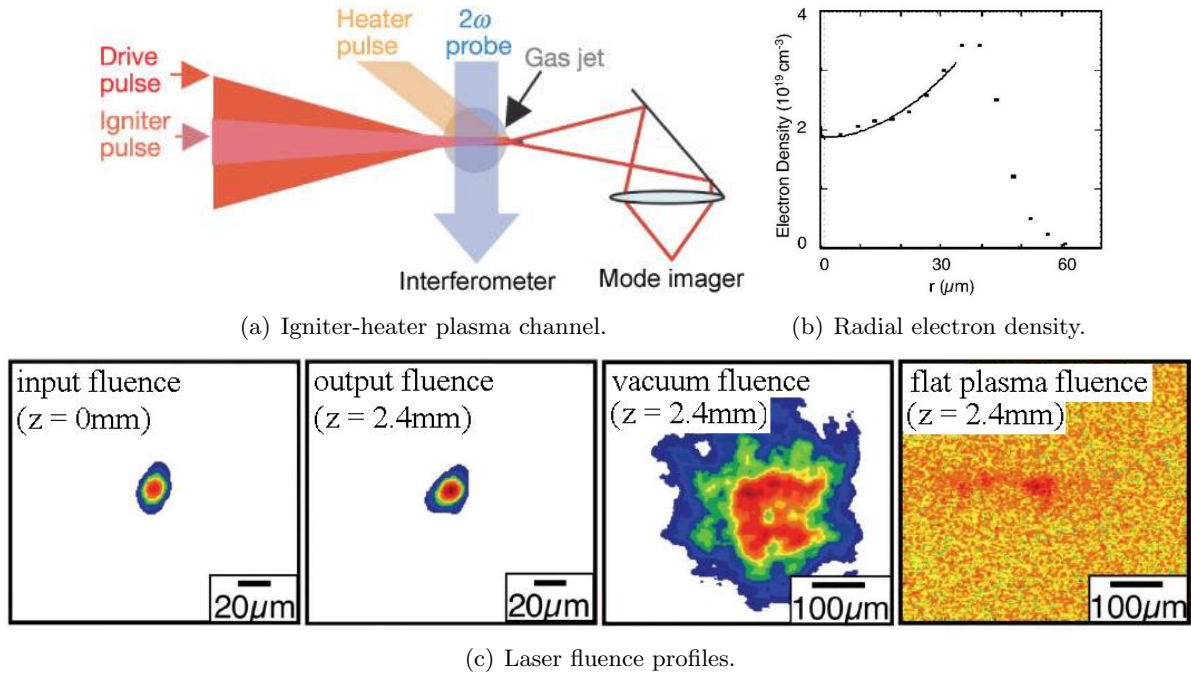


Figure 2.2: The igniter-heater scheme

Figure (a): a schematic diagram of the igniter heater scheme. After [6]. Figure (b): interferometric measurements of the radial density profile using the probe laser pulse. After [48]. Figure (c): measured fluence profiles, comparing input and output modes in scenarios, as labeled. After [49].

from the fact that a high- Z gas was required for efficient inverse bremsstrahlung heating. In the experiment Argon was ionised by the initial pulse to Ar^{8+} , meaning that if the guided pulse was of intensity $\gtrsim 10^{18} \text{ W cm}^{-3}$ further ionisation would occur. Since the rate of ionisation would be dependent on the local intensity of the laser more electrons would be ionised on axis, inverting the electron density profile for some region near the axis and causing refractive defocussing.

For a low- Z gas like hydrogen or helium the radiation for creating the channel must be sufficiently intense to ionise the gas ($\gtrsim 10^{14} \text{ W cm}^{-2}$), but low enough for efficient IB heating¹ ($\lesssim 10^{13} \text{ W cm}^{-2}$). Clearly, this is impossible with a single pulse if a high- Z buffer gas cannot be used and this led to the ‘igniter-heater’ scheme [49, 50], which is shown in figure 2.2(a). An initial fast ‘igniter’ pulse ionises helium from a gas jet. A longer, transverse low intensity pulse, arriving shortly after the initial ionisation then heats the plasma, driving the expansion and formation of the plasma channel. A probe pulse was used to measure the radial plasma profile through interferometry and the results of this measurement are shown in figure 2.2(b) — a guiding structure is clearly formed. The results of using this channel to guide a pulse of intensity $\sim 5 \times 10^{18} \text{ W cm}^{-2}$ over $10Z_R$ are shown in figure 2.2(c). The output fluence profile is almost identical to that of the input. The effect of the plasma channel is clear when comparing the fluence profile of the laser measured in the same plane as the output but after propagating through vacuum. Similarly the laser pulse was not guided when propagating through a flat plasma profile even though for this experiment $P/P_c \approx 2$. Even in the self-focussing regime a plasma channel was required to maintain

¹The absorption of photons by colliding charges.




Figure 2.3: Two-pulse ponderomotive guiding.

Schematic diagram showing the guiding pulse (narrow red cone) and the co-propagating drive pulse (wide red cone) focussed into a gas jet (green). The resulting electron beam is also shown (blue cone). Taken from [51].

the intensity of the laser pulse (for the particular set of conditions in question). The ‘igniter-heater’ scheme was used to create plasma channels for LWFA in 2004 [6, 48] and it was shown that quality and energy of the electron bunches were significantly improved through the presence of a plasma channel.

The igniter-heater technique is effective for creating plasma channels with axial densities $\sim 10^{19} \text{ cm}^{-3}$. Without any external structure for the plasma channel (such as a capillary) there is no issue of damage, meaning that at relatively high plasma density this is an ideal method for creating a plasma channel for LWFA. Unfortunately at lower density ($\sim 10^{18} \text{ cm}^{-3}$), for which a higher energy gain is possible, the heater pulse does not transfer energy to the plasma efficiently [7], meaning that igniter-heater technique is not effective in that regime.

2.2.2 Two-pulse ponderomotive guiding

In a recent experiment it was demonstrated that ponderomotive guiding could be adapted for a two pulse scheme [51]. A leading pulse, with a spot radius W_{lead} chosen so that $Z_R \gtrsim L_d$, was focussed into a gas jet. Behind the lead pulse its wake has a channel like profile whose matched spot size is relatively narrow. A co-propagating trailing pulse of radius $W_{\text{trail}} \approx \lambda_p/2$ was also focussed into the gas jet, following in the path of the lead pulse. It was demonstrated experimentally and theoretically that when the delay between the two pulses was set correctly, the trailing pulse was guided in the wake of the lead pulse, enhancing the generation of electron beams.

The principal advantage of this guiding scheme is that no external structure is required. The disadvantage is that a certain fraction of the total laser energy is required for the leading pulse, meaning that there is less energy for the trailing pulse that drives the accelerating wake. However, this scheme is relatively new and further work may determine that the conditions can be optimised so that only a small fraction of the total laser energy is required for the lead pulse.

2.2.3 Discharge ablation capillaries

A plasma channel can be formed by flowing a current pulse of peak current ~ 200 A and duration ~ 500 ns through a capillary. The source of the plasma is material ablated from the walls of the capillary, which has diameter ~ 400 μm and is made from a soft material, for example polypropene. This plasma fills the capillary and is heated by the current. Heat is conducted to the capillary walls, meaning that the temperature is peaked on axis. Since the pressure within the capillary does not vary transversely this means that the density is minimised on axis and a plasma channel, suitable for guiding laser pulses, is formed.

This technique was first demonstrated to guide a 30 μm spot of intensity $\sim 4 \times 10^{17}$ W cm^{-2} over 2 cm with a transmission $\sim 75\%$ [52, 53]. Simulations show that at high intensities ionization-induced refractive defocussing could be a significant problem for this type of waveguide [54]. However electrons were recently accelerated to 560 MeV in such a device with a 4 J, 27 fs laser pulse [55], demonstrating their usefulness for LWFA.

The discharge ablated waveguide is similar to that which is the subject of this thesis, introduced in section 2.3. The advantage of the discharge-ablated scheme is the relative simplicity of not requiring an external supply of gas, and a reduction in the gas load to the vacuum system. The principal disadvantage is that by its very nature the waveguide must have a limited lifetime, because material is ablated from the wall. The lifetime of the waveguide is typically found to be less than 1000 shots.

2.2.4 Fast discharge gas cells

In section 2.2.1 a technique was described in which a plasma channel was created during the expansion of a plasma column, heated by a laser pulse through the inverse bremsstrahlung mechanism. Similarly, plasma channels can be created by the expansion of plasma columns heated by a fast² current pulse.

The principal technical issue associated with the fast expanding plasma column is the controlled initiation of the current with a jitter that is small relative to the timescale of plasma expansion, so that the laser pulse can be timed to arrive at the correct moment during the formation of the plasma channel. In a prototype device this problem was overcome by field ionising the gas between two high voltage electrodes with an initial laser pulse, so that the discharge breakdown was seeded [56]. Interferometry measurements showed that a plasma channel, suitable for guiding, was formed. The disadvantage of this method was that the initial laser pulse was required to arrive between the electrodes ~ 100 ns before the main pulse. That corresponds to an optical path difference of 30 m, which can be impractical to implement.

Recently this technique was improved by introducing a fast rise time voltage pulse, which ionised the initial hydrogen gas with minimal jitter. The straightness of the plasma column was defined by a series of 300 μm apertures in dielectric plates placed evenly between the electrodes that restricted the plasma diameter and position. A schematic diagram of the device is given in figure 2.4(a).

Initial results using this device are very promising, showing that a low intensity laser spot of radius ~ 25 μm could be guided over 16.6 mm with almost 100% transmission. The output fluence

²Compared to the plasma expansion time.




Figure 2.4: Fast discharge gas-cell

Figure (a): a schematic diagram of the gas cell, showing the charged conical electrodes (red) and the thin dielectric plates within which the $300\ \mu\text{m}$ apertures are drilled (blue). Figure (b): output fluence profiles without and with the plasma channel. The circular markings represents the $300\ \mu\text{m}$ diameter of the final aperture. Both figures are taken from [57]

mode of figure 2.4(b) show that the spot remained tightly focussed. Most impressively, the axial plasma density was measured interferometrically to be $3 \times 10^{17}\ \text{cm}^{-3}$ [57], so that $W_c \sim \lambda_p/2$. Most competing techniques for creating plasma channels are unable to produce steep density gradients at such a low density. Additional advantages to this technique are transverse optical access (for diagnostics), long device lifetime (most probably) and the possibility of creating longitudinal variations in plasma density (by filling different cells within the device to different pressure). Currently efforts are being made to create plasma channels longer than 16.6 mm by this technique, and to the knowledge of the author these fast discharge gas cells have not yet been tested for guiding high intensity laser pulses. However, devices of this type may be the answer for guiding tight laser spots at low plasma densities to extend the acceleration length of multi-GeV laser wakefield accelerators.

2.3 The hydrogen-filled capillary discharge waveguide

The hydrogen-filled capillary discharge waveguide, whose application to LWFA is the main subject of this thesis, was developed by Spence and Hooker [58, 59] in order to overcome the problems of refractive defocussing and limited lifetime inherent to the discharge-ablated waveguide. As in the discharge ablated capillary, a plasma channel is formed in the hydrogen-filled capillary discharge waveguide as a slow current pulse is flowed through the capillary. This is due to the fact that heat is conducted away from the plasma by the capillary walls, meaning that the temperature is peaked on axis. Unlike the discharge ablated capillary the wall material for the hydrogen-filled discharge capillary waveguide is sapphire (or in earlier designs alumina), and the source of plasma




Figure 2.5: The hydrogen-filled capillary discharge waveguide
Schematic diagram showing the sapphire plates (dark blue), the electrodes (yellow), the guided laser radiation (red) and the waveguide housing (light blue). Taken from [45].

is not material ablated from the wall but an initial gas fill of hydrogen gas that is ionised by the discharge. Since the hydrogen gas is easily ionised the problem of refractive defocussing is eliminated. Sapphire is considerably more resistant to ablation than polypropylene walls, meaning that the lifetime of the device is greatly increased.

2.3.1 Design and operation

There have been several designs of the hydrogen-filled capillary discharge waveguide. The version used for the work of this thesis is that designed by A. J. Gonsalves, who also developed the automated optical system used in the group for laser machining the capillaries. The main features of the waveguide design are outlined below, and a more detailed description of the design and the machining of capillaries can be found in the thesis of A. J. Gonsalves [45].

Figure 2.5 depicts the design and operation of the waveguide schematically. A capillary is formed by machining semi-circular channels of diameter D into two flat sapphire blocks, which are then sandwiched together to form a narrow capillary of length L . Hydrogen gas is flowed into the capillary through gas slots machined into the sapphire plates near the capillary ends. The fill pressure, p , is measured in the pipe leading to the waveguide unit. It has been verified, using a method outlined in [45], that the pressure in the pipe is approximately equal to the pressure inside the capillary.

A 1.7 nF capacitor, charged to a voltage $V \sim 25$ kV, is connected across the electrodes situated at either end of the capillary using a thyatron switch. This results in the breakdown of the gas within the waveguide to form a plasma. As the current flows the density profile evolves into a stable plasma channel. The laser is guided if the time delay t between the onset of the discharge and the arrival of the laser pulse is set so that a plasma channel has formed by the time the laser

pulse arrives at the capillary entrance.

2.3.2 Scaling waveguide parameters

There are two key parameters to the plasma channel that forms inside the waveguide. These are the matched spot of the channel, W_c , and the electron density on the channel axis, n_e . Initial magnetohydrodynamic simulations by Bobrova et al showed that as current flowed through the capillary the plasma would evolve towards a steady state in which the heat conduction away from the plasma balanced heat deposited into the plasma from the current [60]. An estimate of the steady state profile could be made without the need for simulations by considering the heat flow equation

$$-\nabla [\chi \nabla] = \sigma \mathcal{E}_{\text{app}}^2 \quad (2.8)$$

where χ is the thermal conductivity of the plasma, σ is the electrical conductivity and \mathcal{E}_{app} is the applied electric field. The plasma was treated as being fully ionised and the temperature at the capillary wall was set equal to 0 K. Using the expressions for χ and σ , defined by Braginskii [61] it was found that

$$n_e = 1.47 n_{\text{H}_2} \quad (2.9)$$

$$W_c = 0.643 \left(\frac{D^2}{n_{\text{H}_2} r_e} \right)^{1/4} \quad (2.10)$$

where n_{H_2} is the initial density of hydrogen molecules and r_e is the classical electron radius.

An improved magnetohydrodynamic model was developed by Broks which could track the evolution of the plasma from the initial hydrogen gas state with a time varying current [62]. This model took into consideration effects such as the absorption of hydrogen gas by the capillary walls and the fact that the plasma was not necessarily fully ionised. By repeating the simulations for differing conditions the following expressions for W_c and n_e were inferred [63]

$$n_e = 0.736 n_{\text{H}_2} + 0.28 \times 10^{24} \text{ m}^{-3} \quad (2.11)$$

$$W_c [\mu\text{m}] = 3.08 \times 10^6 \left(\frac{D [\mu\text{m}]}{2} \right)^{0.5625} (n_{\text{H}_2} [\text{m}^{-3}])^{-0.25} \quad (2.12)$$

The values of W_c and n_e predicted by the Bobrova quasistatic model and the Broks simulation are plotted against n_{H_2} in figure 2.6. The two models predict approximately the same value of W_c , but the values of n_e differ considerably. Interferometric experiments were conducted in order to determine which of these theories was correct [45, 64], and the resulting data are plotted in 2.6. The values of W_c measured by interferometry are in agreement with the predictions of both models. However the measured values of n_e are closer to the value predicted by Broks. Fits to the measured data yielded the following formulae for n_e and W_c

$$n_e [\text{m}^{-3}] = 0.87 \times n_{\text{H}_2} [\text{m}^{-3}] + 0.11 \times 10^{24} \quad (2.13)$$

$$W_c [\mu\text{m}] = 6.6 \times 10^4 \left(\frac{D [\mu\text{m}]}{2} \right)^{0.651} (n_{\text{H}_2} [\text{m}^{-3}])^{-0.1875} \quad (2.14)$$

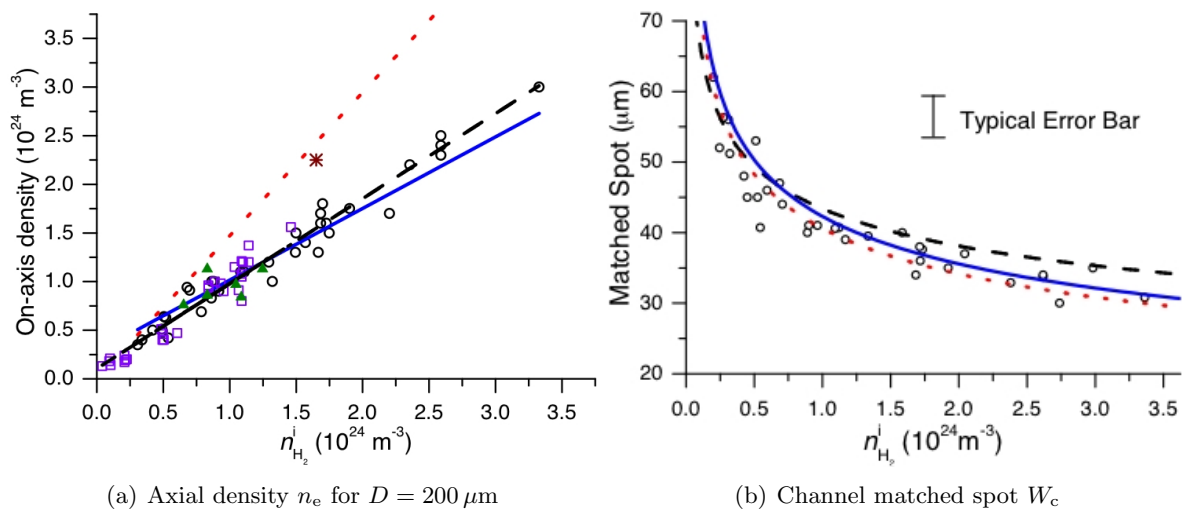


Figure 2.6: Scaling laws and interferometry measurements of waveguide parameters. The axial electron density, n_e , and matched spot, W_c , against the initial fill density of hydrogen molecules, n_{H_2} . Plotted are: measured values for $D = 125 \mu\text{m}$ (triangles), $D = 200 \mu\text{m}$ (circles) and $D = 465 \mu\text{m}$ (squares); the predictions of equations (2.9) and (2.10), which is the Bobrova static model [60] (red dotted line); the predictions of equations (2.11) and (2.12), which are the fits to Broks simulation results [63] (blue solid line); the predictions of equation (2.13) and (2.14), which are fits to interferometry data [45, 64] (dashed line). After [45].

For the remaining sections of this thesis the symbols n_e and W_c are, unless otherwise stated, those calculated using equations (2.13) and (2.14). The value of n_{H_2} is calculated using the measured value of p and the ideal gas law $p = n_{\text{H}_2} k_{\text{B}} T$ where $T = 290 \text{ K}$. The validity of using these formulae to convert from p and D to n_e and W_c is examined in Chapter 6.

The scaling laws can be used to predict the relationship between W_c and n_e , and hence estimate the ratio between W_c and λ_{p} that can be achieved with the waveguide. The formulae for W_c are all a function of D and theoretically W_c can be set to any size by varying D . However in 2.1 it was shown that guiding by grazing incidence reflections at the capillary wall occur for $D \approx 3W$. For capillaries of this size or smaller, the interaction between the capillary wall and the laser pulse must dominate. This sets a lower limit on useful values of D for the hydrogen-filled capillary discharge waveguide. In the experiments of this thesis $D \approx 5W$ was used and, as shall be shown, the guiding effect of the plasma channel dominated. Substituting this value of D into all of the equations for n_e and W_c gives $W_c \approx 3\lambda_{\text{p}}/2$. Ideally the matched spot size of the channel would be smaller, but as discussed in the following section, the waveguide greatly improves LWFA performance [7, 8].

2.3.3 Recent results using the hydrogen-filled capillary discharge waveguide

Experiments show that the hydrogen-filled capillary discharge waveguide can guide intense spots over many Rayleigh ranges with negligible loss of energy [1, 45, 58, 59, 65]. For example, figure 2.7 shows the results of guiding a spot of radius $W = 31 \mu\text{m}$, peak intensity $1.2 \times 10^{17} \text{ W cm}^{-2}$ and measured Rayleigh range 1.2 mm through capillaries of diameter $D = 400 \mu\text{m}$ and lengths $L = 30 - 50 \text{ mm}$. The transmission through the waveguide was close to 100%. For negative t ,

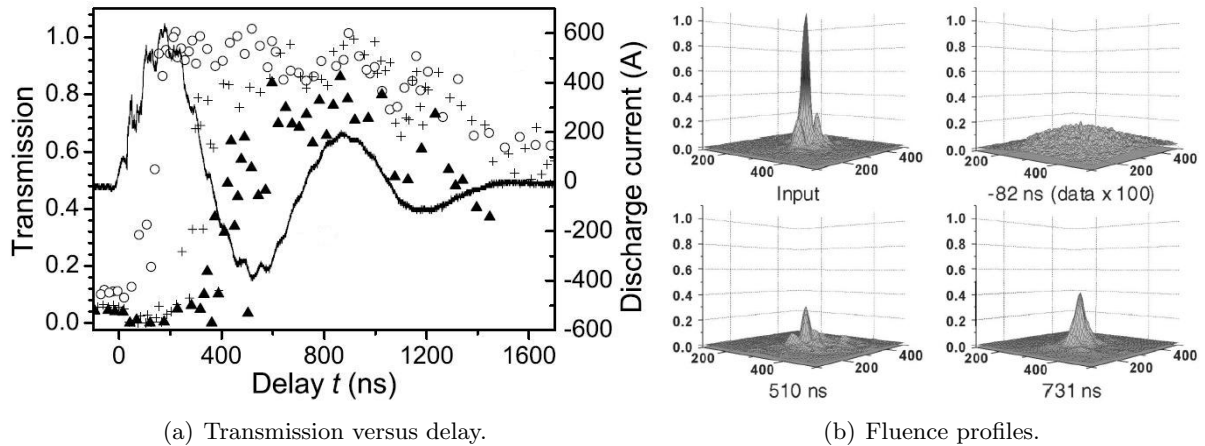


Figure 2.7: Demonstration of guiding in the hydrogen filled capillary-discharge waveguide. The results of guiding a laser spot of $W = 31 \mu\text{m}$, peak intensity $1.2 \times 10^{17} \text{ W cm}^{-2}$ and Rayleigh range $Z_R = 1.2 \text{ mm}$. Figure (a): the fraction of energy transmitted through a waveguide against delay t for $[D = 400 \mu\text{m}, L = 30 \text{ mm}, p = 110 \text{ mbar}]$ (circles), $[D = 400 \mu\text{m}, L = 30 \text{ mm}, p = 330 \text{ mbar}]$ (crosses) and $[D = 400 \mu\text{m}, L = 50 \text{ mm}, p = 330 \text{ mbar}]$ (triangles). Figure (b): transmitted fluence profiles for selected t for $[D = 400 \mu\text{m}, L = 30 \text{ mm}, p = 330 \text{ mbar}]$. After [1].

where no plasma channel was formed, virtually no energy was transmitted, showing that a plasma channel must have been responsible for the high transmission at later delays.

The output modes of the plasma channel were approximately Gaussian, but with a lower peak fluence than the input. Equation (2.14) predicts that for those conditions $W_c = 44 \mu\text{m}$. Equation (2.6), which describes the oscillations in the size of a Gaussian spot during propagation through a parabolic channel, can be used to show that the peak intensity should vary between $1.2 \times 10^{17} \text{ W cm}^{-2}$ and $0.3 \times 10^{17} \text{ W cm}^{-2}$ (assuming an ideal Gaussian beam). The predicted peak intensity at the capillary exit is $0.5 \times 10^{17} \text{ W cm}^{-2}$ when the actual value was $\sim 0.4 \times 10^{17} \text{ W cm}^{-2}$.

The hydrogen-filled capillary discharge waveguide has been used to enhance x-ray lasing in a hydrogen plasma doped with xenon by maintaining the intensity of the driving 800 nm laser pulse




Figure 2.8: Acceleration of electrons to 1 GeV in the hydrogen-filled capillary discharge waveguide. The electron signal from an electron spectrometer measuring electrons accelerated by LWFA with a 1 J, 33 fs laser pulse guided in a $[D = 310 \mu\text{m}, L = 30 \text{ mm}]$ hydrogen-filled capillary discharge waveguide. Taken from [7].

over many Z_R [66]. In 2006 the hydrogen-filled capillary discharge waveguide was used in a LWFA experiment for the first time. It was found that monoenergetic electron beams could be accelerated to 1 GeV, as seen in figure 2.8 using a 1.5 J, 38 fs (40 TW) laser beam, and to 500 MeV using a 0.9 J, 73 fs (12 TW) guided by a 33 mm waveguide [7, 8]. Prior to that the highest energy of monoenergetic electron beams generated by LWFA was 170 MeV, using a 1.5 J, 38 fs (30 TW), produced in a gas jet with no external guiding structure [5]. The improvement due to the hydrogen-filled capillary discharge waveguide was very significant, and the energies achieved were only a factor of 3 less than that at a facility such as the Diamond Light Source. Additionally the stability of the beams generated in the waveguide appeared to be better than the gas jet experiments of 2004, although it was not established if this improved stability could be attributed to laser quality or the presence of the plasma channel.

2.4 Conclusion

Different methods for guiding intense laser pulses for LWFA were reviewed, and one of these, the hydrogen-filled capillary discharge waveguide, was described in detail. That waveguide has been attracting attention, due in part to the recent acceleration of monoenergetic beams to 1 GeV, as a possible way of producing compact devices that can replace large scale synchrotron facilities. The work in the following chapters focusses on a comprehensive experimental study, conducted at the 12 TW Astra Laser Facility in 2006, in which electrons were accelerated in the hydrogen-filled capillary discharge waveguide for a range of different conditions. This was motivated by the need to better understand the effect of that particular waveguide, and of plasma channels in general, facilitating the design of compact devices able to generate stable, monoenergetic electron beams with energy above 3 GeV.

Chapter 3

Astra experiment setup and methods of analysis

From the start of June 2006 to the end of August 2006 experiments were conducted using the Astra laser facility at the Rutherford Appleton Laboratory. The aim of these experiments was to undertake a thorough investigation into the effect of the hydrogen-filled capillary discharge waveguide (and more generally plasma channels) on laser wakefield acceleration across a range of conditions. The experimental setup used is described in this chapter. A very large volume of data was collected over the course of the three months, requiring reliable automated processing of the diagnostic outputs to extract the essential information, and the development of statistical methods to interpret that information meaningfully. These methods of analysis are also the subject of this chapter.

3.1 Experimental setup

The experimental layout after the laser compressor is outlined in figure 3.1. The Astra laser was focussed into the entrance of the hydrogen-filled capillary discharge waveguide. The exit of the waveguide was imaged onto diagnostics to measure the fluence profile, spectrum and energy of radiation emerging from the waveguide. Co-propagating electron beams entered into an electron spectrometer which allowed measurement of their energy spectrum. In this way simultaneous observation of the guiding properties of the waveguide and the resultant electron bunches was made for a range of conditions, allowing detailed comparison of the two. The description of the experimental setup that follows is broken down into five different elements: the Astra laser; the waveguide; the optical diagnostics; the electron diagnostics; the control software.

3.1.1 The Astra laser

The Astra laser delivered pulses of diameter 6 cm that were focused using an $f/27$ off axis paraboloid. The exact position of the laser focus was determined by attenuating the beam and imaging its profile with a 12-bit camera, which was mounted on a longitudinal translation stage and placed directly in the line of the beam. Images of the laser profile were recorded as the z position of the camera was

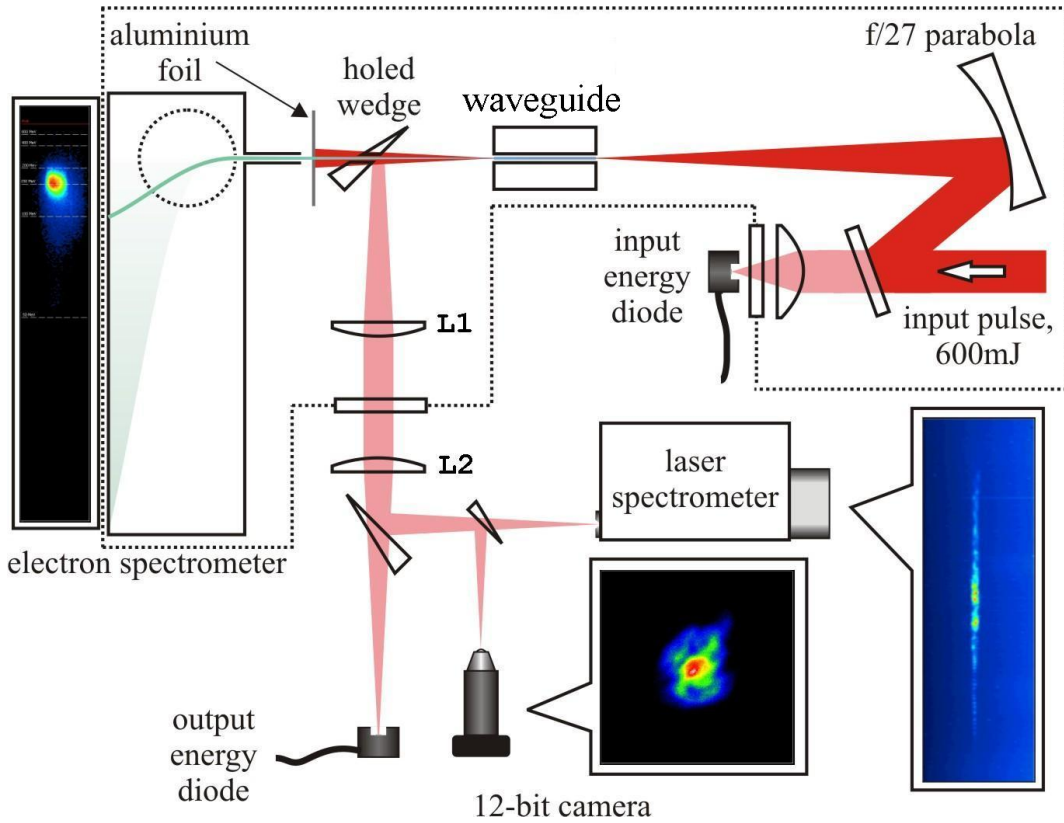


Figure 3.1: Schematic layout of the Astra experiments.

The laser beam (red) arrives at the right hand side of the image having emerged from the compressor. It is focussed into the waveguide. The electron beam (green line) generated inside the waveguide passes through the axial hole in the wedge and aluminium foil to the electron spectrometer where it is deflected onto a phosphorescent screen. A fraction of the laser radiation emerging from the waveguide is reflected by the holed wedge and imaged onto the optical diagnostics using achromat lenses L1 and L2. Components within the dotted lines are under vacuum. Some turning mirrors and wedge reflections (used to attenuate the transmitted beam) are not included so as to maintain the simplicity of the diagram.




Figure 3.2: Focal scan of the Astra laser.

The above focal scan was conducted with a 12-bit camera at focus at low power setting. Each square in this plot represents a square in the focal plane of side $125 \mu\text{m}$. The scan was used to estimate the peak fluence, spot size, rayleigh range, pointing stability and M^2 of the Astra laser.

Shot energy		600 mJ
Central wavelength (λ_0)		818 nm
Spot radius (W)		$34 \mu\text{m}$
Rayleigh range (Z_R)		4.5 mm
M^2		1.5
Pointing r.m.s.		$4.5 \mu\text{m}$
		$2.8 \mu\text{rad}$
Peak fluence		$4 \times 10^4 \text{ J cm}^{-2}$
FWHM pulse length (τ)	Long pulse	150 fs
	Short pulse	45 fs
Peak Power (P)	Long pulse	3.5 TW
	Short pulse	12 TW
Peak intensity (I_0)	Long pulse	$2.4 \times 10^{17} \text{ W cm}^{-2}$
	Short pulse	$7.8 \times 10^{17} \text{ W cm}^{-2}$
a_0	Long pulse	0.33
	Short pulse	0.60

Table 3.1: Astra laser parameters

changed. The results of this focal scan can be seen in figure 3.2. The position of focus is the position for which the peak fluence of the measured laser radiation was highest. Since the distance of the cameras object plane relative to itself was known, the z -position of the focus inside the vacuum chamber could be determined. The z -position of the focus was found not to drift noticeably. The $1/e^2$ spot radius (W), Rayleigh range (Z_R), pointing stability and M^2 were estimated from the focal scan.

Experiments were conducted for two different laser pulse lengths: a *short pulse*, for which the laser was fully compressed; a *long pulse*, which was achieved by removing a glass block from the beam line before the compressor, resulting in a longer, chirped pulse. The purpose of changing the pulse length was to change the intensity of the radiation without varying the energy or alignment of the beam. Maintaining the beam energy meant that there was no need to change the attenuation or calibration of any of the diagnostics, allowing for direct comparison between short pulse and long pulse data. The long pulse was used as a probe to measure the guiding properties of the waveguide at intensities for which processes such as self-focussing and pulse depletion were insignificant. A comparison of the electrons accelerated by the short pulse and the guiding of the long pulse for the same conditions could be used to determine the degree to which the LWFA process was affected by the guiding channel.

The full-width at half maximum of the temporal pulse envelope was measured by auto-correlation. The peak power P , peak intensity I_0 and normalised vector potential a_0 were estimated by assuming a sech^2 form for the pulse envelope. The laser parameters are summarised in table 3.1.

3.1.2 The waveguide

The principal of the hydrogen-filled capillary discharge waveguide is described in chapter 2. In these experiments sapphire capillaries of lengths $L = 7$ mm, 15 mm, 33 mm and 50 mm were employed, with inner diameters, D , ranging from $200 \mu\text{m}$ to $300 \mu\text{m}$. The gas slots had diameter $650 \mu\text{m}$ and length 10 mm located 2 or 4 mm from each end¹. The gas was pulsed with each laser shot, to reduce the load on the vacuum pumps. The maximum value of p that could be used in pulsed mode without overloading the pumps was 600 mbar. The 1.7 nF capacitor was charged to a voltage (V) between 15 and 30 kV and the current that resulted from connecting this capacitor across the capillary was measured by a Rogowski coil connected to what shall be referred to as the ‘timing oscilloscope’.

A fast diode, connected to the timing oscilloscope, detected the leakage through the back of a turning mirror located before the capillary. During setup a second fast diode was placed at the laser focus to record Δt , the delay between the measured signal from the turning mirror diode and the arrival of the laser pulse at focus. During the experiment t , the delay between the onset of the discharge current and the arrival of the laser pulse at focus, could be found by adding Δt to the measured delay between the discharge current signal and the turning diode signal. The timing oscilloscope was triggered by the discharge current, and $t = 0$ was defined as the point at which the current exceeded 28 A (the peak was ~ 300 A).

The waveguide unit was positioned so that the laser focus was located at the center of the capillary entrance and angled so that the capillary and direction of laser propagation were co-linear. This was performed when the laser beam was switched to low power mode and the exit plane of the capillary was imaged on the fluence monitor. Both guiding [1, 58, 59] and electron acceleration [8] are sensitive to the relative alignment and position of the laser and waveguide. Drift in the laser pointing due to temperature variations meant that adjustments were required approximately once per hour.

3.1.3 Transmitted laser diagnostics

A portion of the laser radiation transmitted by the capillary was reflected by an optically-flat ‘pick-off’ wedge placed 0.6 m behind the capillary exit. After a second reflection from the surface of a wedge, the beam was collimated by an f/15 achromatic lens L1 of focal length 1.2 m. The collimated beam passed normally through a window and was refocused by a second lens L2 which was identical to lens L1. After this point the beam was split using wedges and steered onto different optical diagnostics, as detailed in figure 3.1. Details of the setup of the imaging system can be found in Appendix A.

The branch of the beam that was only reflected from wedges (and never transmitted) was free from astigmatism. The resulting image of the waveguide exit was magnified onto a 12-bit CCD by a microscope objective, the intensity of the beam having been reduced by two further reflections from wedges and transmission through neutral density filters. The recorded image was a measurement of the laser fluence at the exit plane of the waveguide. Each camera pixel represented approximately

¹Except in the case of the $L = 7$ mm capillary, for which there was a single gas slot feeding into the capillary half-way along its length.

1 μm in the object plane.

Another image of the capillary exit, from a different branch of the beam, was formed on the entrance slit of a Bentham grating spectrograph. The spectrum of the transmitted laser radiation was measured by a cooled CCD 16-bit Andor camera in the image plane of the spectrograph. The correspondence between wavelength and Andor pixel number was determined using a (mercury) spectral lamp. The procedure to determine the resolution and correct for the spectral response is detailed in sections 3.2.1 and 3.2.3.

The energies of the pulses entering and leaving the waveguide were measured by spectrally-flat photodiodes located behind a dielectric turning mirror prior to the paraboloid and behind a wedge after lens 2. A calibration was made between the energy input diode and a calorimeter placed at focus so that the input energy could be estimated. Regularly the waveguide would be removed and ‘calibration runs’ conducted in which the laser energy was varied by rotation of a half-wave plate through which a beam passed before passing through a polariser. Data was recorded for all diagnostics just as when the waveguide was in place. The input and output diode signals recorded during calibration runs were compared to one another so that when the waveguide was in place its transmission T could be determined from these diode signals.

In the original experimental set up the pick-off wedge could be wound out of position to allow electron beams to pass through to the electron diagnostics unimpeded. In this configuration no optical data could be recorded whilst electrons were being measured. For later runs a wedge with an on axis hole of diameter 4 mm through which the electrons could pass was used as the pick-off. This allowed for simultaneous measurement of electron beams and the properties of the laser radiation transmitted by the waveguide. The hole subtended an angle of 7 mrad relative to the capillary exit and was found to have no noticeable effect on the measurements of the transmitted laser radiation. The output diode was calibrated for both wedges.

3.1.4 Electron diagnostics

The electron diagnostics were isolated from the optical diagnostics by thin foils of aluminium that blocked all laser radiation but did not significantly impede the electron beam. As is now a standard technique, phosphor (Lanex) screens, imaged by a 12-bit CCD camera, were used to detect electron beams. The phosphor emitted visible light where electrons were incident, and this light was recorded by the camera to give 2-dimensional images of the electron beams as they passed through the detection plane. The use of this arrangement allows a spatially-resolved measurement of an electron beam without need for removal (for scanning) of the screen, which would have slowed the repetition rate of the experiment to such an extent that meaningful scans of parameters would have been impossible.

Phosphor screens were used for two different diagnostics, the electron pointing monitor and the electron spectrometer. Initially a phosphor screen was set up at 45° to the incoming electron beams, with the imaging camera orientated at 45° to the screen and 90° to the beam. This orientation meant that the profile and the positions of the electron beam were imaged without distortion due to the angle of the screen or camera. This was used to measure the electron beam divergence and pointing stability. However since the phosphor screens were found to disperse the transmitted

electron beams, the electron pointing monitor was removed once the electron spectrometer was installed, meaning that simultaneous measurements of beam pointing and energy were not possible.

The electron spectrometer has been used for several experiments other than that described here [4, 51, 67]. A DC current was flowed through two electromagnets of diameter 38 mm to produce a magnetic field of order 0.8 T (depending on the current). Electrons passing through the region of high magnetic field were deflected and the resulting displacement of the emitted signal from the phosphor screen at the back of the spectrometer served as a measure of the electron energy. The relationship between the electron energy and displacement, $E(x)$, was determined by a particle tracking code [68]. This relation was not linear meaning that the energy resolution, determined by the size of a perfectly monoenergetic beam on the lanex, was lower at high energy where the relative deflection of beams of differing energy was less. This experiment was conducted at the limits of the resolution of the spectrometer and the problem of deconvolving the spectrum effectively is addressed in section 3.2.4.

Figure 3.3 is an image used to calibrate camera pixels to beam displacement, shown with a shot of real data superimposed. In order to calibrate the electron spectrometer camera signal in terms of electron charge, electron-sensitive image plates were placed flat against the lanex screen so that electrons passed through the image plates before exciting the phosphor of the lanex screen. A more detailed description of how the image plate encodes information and the scanning process by which it was extracted for this experiment is given in [68, 69]. A spatially resolved signal would be recorded on the image plate and on the lanex screen for the same electron beam. A comparison of the two yielded an estimate of the electron charge per electron spectrometer count. The use of image plates is time consuming in that for each shot the plate needs removal, scanning and replacement, therefore image plates were used only to calibrate the lanex screen and thereafter the experiment proceeded without them.

3.1.5 Automated control and data collection system

All diagnostics were connected to a network of 5 computers controlled by the *Imagemaster* system originally written by D.J. Spence and adapted by A.J. Gonsalves. The code displayed and saved the data being collected and varied the delay t randomly within user specified ranges.

The control code was important for two reasons: firstly it allowed at least an order of magnitude more data to be collected than would have otherwise been possible, allowing more efficient use of the allocated laser time. Secondly, the ability of the control code to control the delay t automatically (whereas varying any other parameter required the target area to be opened, searched and locked) meant that t was scanned more finely than any other parameter². Typically all other parameters were set and then t varied. Such a subset of data would be saved as a ‘run’. The fact that data was collected in this way means that the data is also most naturally presented as the variation of measured parameters against t with all other conditions fixed (where multiple runs existed for the same set of conditions all will be included in such a plot).

²In order to mitigate the effects of drift in the properties of the laser t was varied randomly within the specified range. This way any trends observed as t varied could only be attributed to the particular properties of the plasma channel at the delays in question.




Figure 3.3: Electron spectrometer calibration image with electron beam data superimposed. The calibration image (greyscale) was recorded on the electron diagnostic camera with the room lights of the target area switched on. The back of the electron spectrometer and the edges of the mirror used to image it are clearly visible, as is the 16-bit camera itself. When experimental data was collected the room lights were switched off meaning that the only light detected would be that emitted from the lanex screen where electrons were incident. An example of a data image, with background subtracted, is overlaid (colour scale, 0 = transparent).

3.2 Analysis of diagnostic output

The use of automated control and data collection systems meant that over 10,000 shots were recorded during the three months of experiments on the Astra laser (of which about 7,000 were at high power) corresponding to over 250 GB of data. The number of shots was far in excess than those usually recorded in a laser wakefield acceleration experiment and as such represented an opportunity to gain valuable new insights into the key processes. Whilst it would have been easier to have selected a subset of shots for which the most interesting electron acceleration behaviour was observed and to have concentrated analysis upon those, much could also be learned about the processes of laser wakefield acceleration by analysis of the data for which electron acceleration activity was *not* apparent. The size of the body of data added two key requirements to the analysis. Firstly any processes applied to the output of a diagnostic needed to be sufficiently robust to be relied upon to work without individual verification, secondly the information from any given diagnostic needed to be condensed to one or two key numbers so that the statistical analyses of section 3.3 could be applied.

3.2.1 Processing of CCD images

The following diagnostics produced data in the form of CCD images: transmitted laser fluence monitor; laser spectrometer; lanex monitor. In all cases the cameras were either 12-bit or 16-bit, meaning that significant information could be recorded at very low signal levels, assuming that the images were processed correctly. In chapter 6, for example, an analysis of outlying spectral features of order 1/1000th of the peak signal allowed determination of the ionization state of the plasma channel. It was therefore important to define carefully the processes applied to CCD images so that key features of the data were not lost.

Much of the image processing was identical for different diagnostics. A laser spectrometer image is used as an example to illustrate how all CCD images were treated. Figure 3.4 shows an image from the laser spectrometer at different stages of image processing. These stages are discussed below.

Removing X-ray hits

On image 3.4(a) there are many isolated pixels for which the signal is much higher than those surrounding them. Experimentally these excited pixels are strongly correlated to the acceleration of electrons. They are commonly observed in laser wakefield experiments and result from the excitation of individual CCD pixels by short wavelength radiation generated by the acceleration or deceleration of electrons. Generally they are referred to as ‘X-ray hits’. When analysing data automatically X-ray hits can significantly distort the output.

A simple algorithm, effectively a median filter in one dimension, was written to remove X-ray hits which made use of the fact that they are usually restricted to individual pixels. The CCD image was analysed row by row. The signal of each pixel of that row, s_n was compared to the preceding and following pixel counts, s_{n-1} and s_{n+1} , and the new count value s'_n was determined

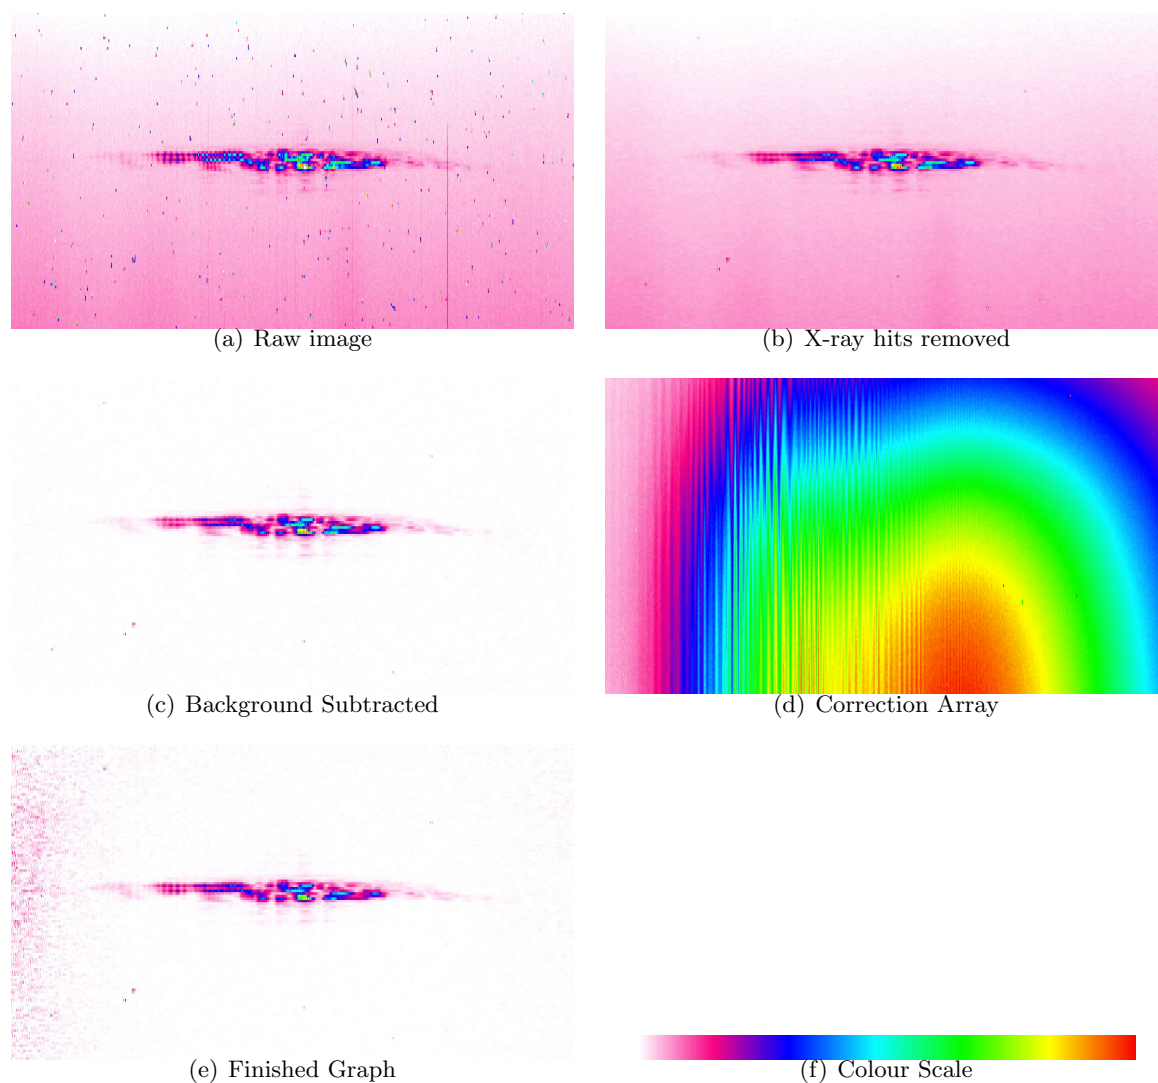


Figure 3.4: CCD image at different stages of processing.

The example shown is an image from the CCD camera on the grating spectrograph. To convert (a) to (b) an algorithm to remove X-ray hits was applied. To convert from (b) to (c) the background signal was subtracted. The processed image (e) was the result of dividing (c) by (d), which was the correction for the spectrometer sensitivity.

using the following rule:

$$s'_n = \begin{cases} \frac{1}{2}(C_{n-1} + s_{n+1}) & ; (s_n > s_{n-1}) \text{ AND } (s_n > s_{n+1}) = \text{TRUE} & \text{X-ray hit detected} \\ s_n & ; (s_n > s_{n-1}) \text{ AND } (s_n > s_{n+1}) = \text{FALSE} & \text{No X-ray hit} \end{cases} \quad (3.1)$$

The results of applying this rule can be seen when comparing figures 3.4(a) to 3.4(b). Application of (3.1) has removed the X-ray hits without distorting the image.

The only two situations that (3.1) would incorrectly identify pixels as being X-ray hits is when either there was a genuine peak in the signal, and when there were two adjacent X-ray hits. In the former case the consequence would be a slight lowering of the peak pixel, which is insignificant since real peaks in the data will span many pixels. If the latter case were to occur commonly enough to significantly affect the data then it would mean that the density of X-ray hits was such that any set of rules would make it impossible to isolate the signal effectively.

Subtracting background

The background signal was not necessarily uniform across the CCDs meaning that subtracting the background required careful consideration in order to avoid the loss of important data. In addition the background level of different CCDs changed slowly over time due to fluctuations in temperature, meaning that it was not satisfactory to simply subtract a ‘blank’ array taken at the start of the experiment. In order to emphasise the importance of correct background subtraction figure 3.5 outlines the process of background subtraction for a laser spectrometer image with important Raman data that is below the background level for some parts of the image.

In 3.5(a) the background across the CCD is not uniform — the background level is higher at the top of the CCD where the count is ~ 360 , whereas at the bottom the count is ~ 310 . This difference is likely due to non-uniform cooling of the CCD chip or data readout order. The background also varies longitudinally. There is a very low level feature marked in figure 3.5(a) that has a peak value ~ 360 . From 3.5(b) it is apparent that were the maximum background (ie 360) subtracted from the image the feature would be negative (and hence ‘lost’). On the other hand if 310 was subtracted much of the background would be at the same level and higher as the peak in question. This is not a problem when considering a single CCD such as this one, since changing the colour scale reveals the ‘hidden’ feature. However given the volume of data such individual attention was not possible, meaning that a robust means of subtracting background was required.

The solution to this problem is to isolate parts of the CCD where it is known that there is no laser signal (i.e. above and below the dotted lines of figure 3.5(a)) and to put a fit through that data column-by-column. The fit to a single column is shown in figure 3.5(b). The resulting array of column-by-column fits was taken to be the background and is shown in figure 3.5(c). Although a linear fit through each column satisfactorily reproduced the inhomogeneity of the background of the image, higher orders could have been used had it been necessary. Figure 3.5(c) was then taken to be the ‘local background’ and was subtracted from 3.5(a), producing 3.5(d), the background subtracted image, in which the low-level Raman feature has not been lost.

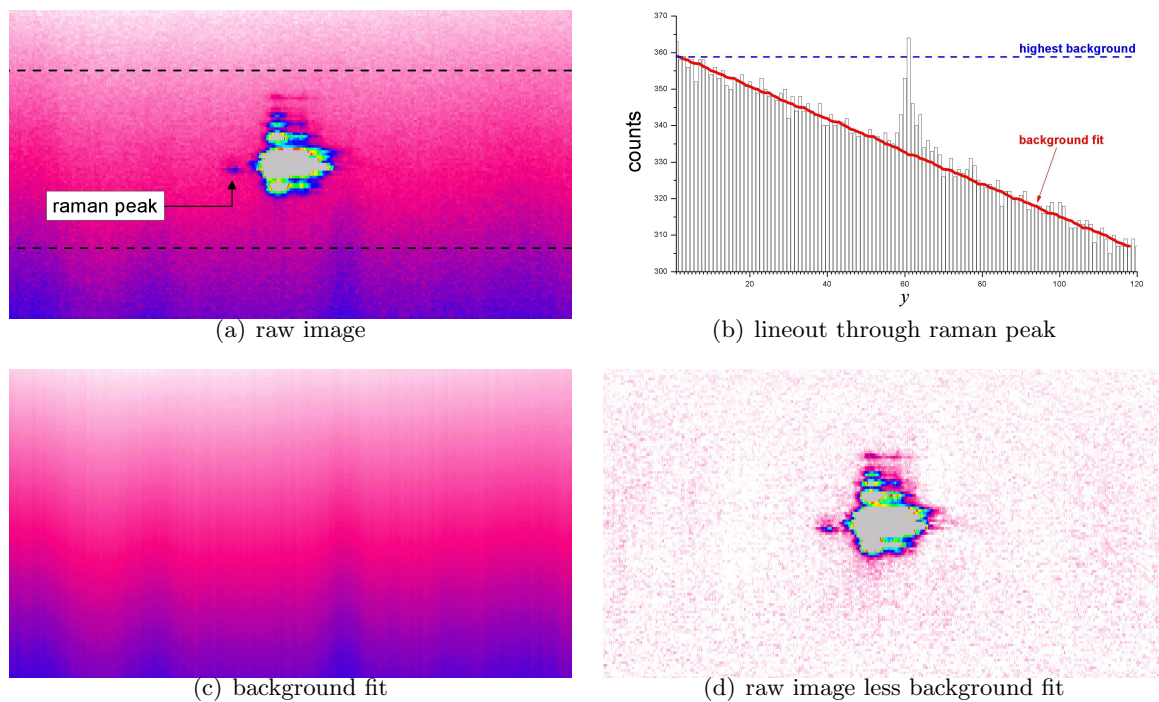


Figure 3.5: Background subtraction from a CCD array.

The example shown is an image from the CCD camera on the grating spectrograph. The colour scale has been set so that the majority of the image is saturated (grey) so that the variation in background across the CCD is visible. There is a low lying Raman Peak in (a) that is at the same level as the highest level of background. Therefore it would have been possible to have lost this data unless a ‘local background’ was subtracted. No signal was expected above and below the dotted lines. The data of an individual array column which passes through the Raman peak is plotted in figure (b). A linear fit through the top and bottom of the image for that column is marked. This is the ‘local’ background for that column. Image (c) is the local background for the entire array, and strongly resembles the background of (a). Image (d) is the (a) less (c), and the low level Raman feature is still visible.

Correcting for system response

The responsiveness of the spectrometer CCD was the result of several factors:

1. The wavelength dependence of the response of the CCD chip.
2. The sensitivity of the imaging system (including spectrometer) at differing wavelengths.
3. The variations in sensitivity across the CCD chip.

In addition there were possible effects due to chromatic aberration, which are discussed in appendix B. In order to compensate for these factors a white light source of known spectrum was placed at the laser focus position and its spectrum recorded by the spectrometer. A background image was recorded immediately afterwards. The background was subtracted from the white light spectrum. The counts in each column were divided by the spectral intensity of the white light source at the corresponding wavelength. Each column was then multiplied by the measured transmission of the neutral density filters³ that were used to attenuate the spectrometer signal. The resulting ‘correction’ image, by which the ‘background subtracted’ image would be divided, is shown in figure 3.4(d). Towards the left hand side of the image (which corresponds to the red end of the spectrum) the values of the array are close to zero: this reflects the fact that the Andor-CCD used is unable to detect radiation at wavelengths above around 950 nm, which corresponds to these columns of the CCD. There is also a ‘rippling’ effect towards the left hand side of the correction array. The origin of this rippling is unclear. It was removed through smoothing which reduced the spectrometer resolution to 5 nm.

3.2.2 Estimating fluence

For this experiment quantitative analysis of the fluence profile monitor was difficult for two reasons. Firstly in the case of the short pulse the spectral shifts that occurred within the waveguide were far greater than anticipated (and far greater than for gas jet experiments). Although achromatic lenses were used, the analysis presented in appendix B shows that the spectral shifts experienced by the short pulse took much of the radiation outside the range for which the achromat lenses were achromatic. The spectrum of the long pulse did not, in general broaden, meaning that optically imaging the exit mode was useful for those shots.

The second problem was a fault with the 12-bit camera used as a fluence monitor which was only discovered after the experiment. The camera was unable to detect a signal below a certain threshold which meant that there was a ‘missing depth’ of low level data that was lost. Assuming a linear response, the fluence at a given pixel was $r(o + s)$ where o is the undesired offset resulting from the camera fault, r is the response of the camera and s is the signal. The offset o could be estimated by comparing the input spot measured with different NDs before the camera and using the fact that the input spot should have had the same form in both cases. Such plots are shown in figure 3.6. Values of u_1 , u_2 and d_1 as defined in figure 3.6 were used to calculate o using

$$o = \frac{d_1 u_2}{u_1 - u_2 - d_1} \quad (3.2)$$

³Whose transmission was in fact wavelength dependent.




Figure 3.6: Comparing vacuum fluence profiles for different NDs.

The two plots are a lineout of counts against pixel averaged over 13 shots for two different runs for which no capillary was in place. In effect both are a measure of the input focus. The red run (plotted against the right hand y -axis) was conducted with stronger NDs in front of the camera than the blue run (plotted against the left hand y -axis). The two plots should have the same form but this is only possible if the red y -axis is offset from the blue y -axis. The offset o is calculated using the values of u_1 , u_2 and d_1 deduced from this figure and is found to be 30.2. The x -axis is situated at -30.2 (with respect to both y scales).

The deduced offset was $o = 30.2$. When estimating the laser spot FWHM that offset was taken into account, if it had not been the calculated spot sizes could have been wrong by up to a factor of two. Having established a value for the offset a calibration relating the number of counts recorded in a pixel to the fluence incident on that pixel was made. For pixels in which the pixel count was zero the fluence was known to be below a certain value.

The only CCD-process that was applied to the fluence monitor data was X-ray hit removal. Due to the non-zero offset o the background signal did not register. The response of the CCD was assumed to be even across the CCD, so no spatial correction was required.

3.2.3 Laser spectra

The processed image of figure 3.4(e) was summed column by column to produce a one dimensional lineout which was plotted against the corresponding wavelength for that pixel. In figure 3.4 the arrays are orientated such that the red end of the spectrum is on the left hand side of the array. The ‘correction array’ of figure 3.4(d) shows that longer wavelengths would not be detected. This was due to the response of the Andor CCD chip which was effectively zero for wavelengths longer than 950 nm. At those wavelengths the signal was effectively the noise (centered about zero) divided by a number very close to zero, which resulted in the large fluctuating signal visible at the left hand edge of figure 3.4(e). For this reason the signal at wavelengths longer than 950 nm were ignored having been deemed outside of the detection range of the diagnostic.

The detailed structure of individual spectra was of interest, and where this was the case the spectra for individual shots was considered. However, for the purposes of analysing a large number of such spectra and correlating their behaviour to the acceleration of electrons, a simple representation of the form of the spectrum was required. It was generally found that for the short pulse (where the interaction between laser and plasma was strong) the modulations within spectra varied from shot to shot. However the overall spread of the spectra were relatively stable when conditions were fixed. Two limits were defined to describe the range of wavelengths across which the spectra lay. The *red limit* (Λ^R) is the wavelength above which 15.9% of the energy resides, the *blue limit* (Λ^B) is the wavelength below which the shortest 15.9% lies. Accordingly 84.1% of the energy lies at wavelengths between the red and blue limits. These numbers are chosen because they correspond to the fraction of energy lying within one standard deviation of a normal distribution. Although the transmitted laser spectra were not at all like a normal distribution, defining such limits provides a useful way to describe the spectral spreading for a large number of shots concisely. If the total signal on the spectrometer was too low to calculate meaningful values of $\Lambda^{R,B}$, or it appeared that significant radiation lay outside the range of the spectrometer, the limits were not calculated.

The analysis of Raman satellites within the long pulse spectra is discussed in chapter 6 where that work is relevant.

3.2.4 Electron spectra

A typical electron spectrometer image, after processing, is shown in figure 3.3 superimposed onto the corresponding spatial calibration image. In addition to the processes of X-ray hit removal and

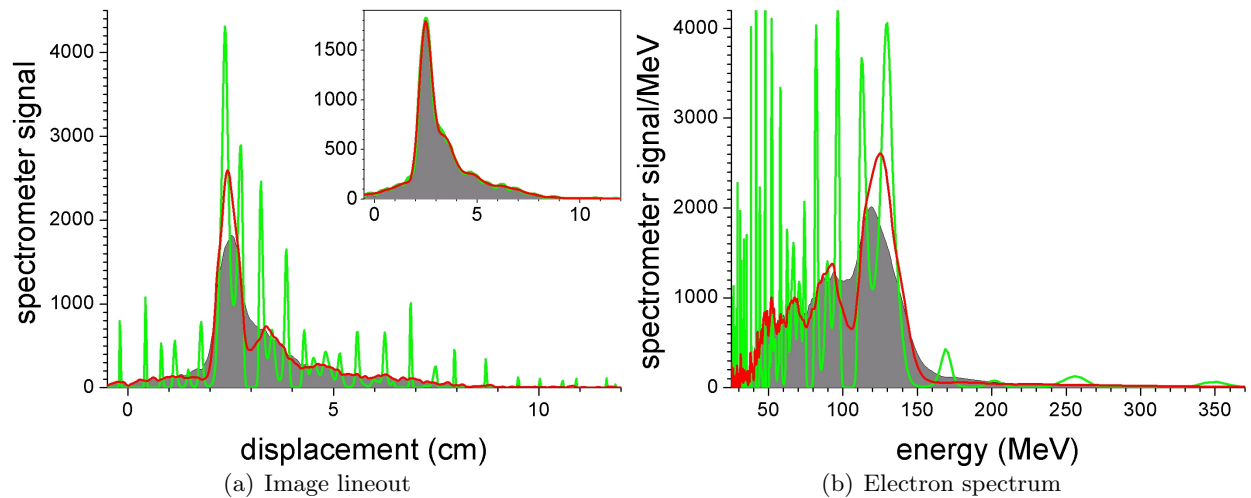


Figure 3.7: Methods of deconvolving the electron spectrum.

The image lineout (a) of the electron beam in figure 3.3 and the resultant electron spectrum (b) calculated using equation (3.3) are plotted (grey fill). The result of deconvolution of the image lineout using equation (3.5), the Richardson-Lucy algorithm [70, 71] is plotted (green line). The results of deconvolution of the spectrum using the adapted version of the Richardson-Lucy algorithm detailed in equations (3.7) to (3.11) is also plotted (red line). The inset graph of figure (a) plots the reconvolved lineouts $P(x) \otimes O(x)$ against the original lineout. Both reconvolved lines are in such good agreement with the original data that the green is obscured by the red.

background subtraction detailed in section 3.2.1, the electron spectrometer CCDs were rotated so that the vertical plane of the spectrometer was parallel to the edge of the array. The images were then cut down to an area corresponding to the lanex strip whose position could be determined by the spatial calibration image. The resulting array was summed column by column to give an image lineout array $\{x_k, s_k\}$ where x is the displacement and s is the electron signal. The electron energy spectrum was then found by the following conversion

$$\underbrace{\{x_k, s_k\}}_{\text{Image lineout array}} \Rightarrow \underbrace{\{E(x_k), s_k \frac{dx}{dE} \Big|_k\}}_{\text{Electron spectrum array}} \quad (3.3)$$

where k refers to each summed pixel column and $E(x_k)$ is calculated from the particle tracking code [68] describing the deviation of electrons by the magnetic field.

Spectrometer resolution limits

The lineout and spectrum of the electron beam shown in figure 3.3 are plotted in figure 3.7. What appeared to be a monoenergetic electron beam with a low energy tail in the original figure, is a rather broad energy spectrum with some poorly resolved peaks (with a FWHM max energy spread of 50%).

It is reasonable to expect that the true energy spectrum of the electron beam was narrower than that shown in figure 3.7(b). The beam in figure 3.3 is effectively circular on the lanex, suggesting

that the width of the beam in the dispersed direction is due mainly to the divergence of the incoming electron beam and not a spread in the electron energy. This effect is most significant at higher energies where changes of electron energy only translate into slight changes of deflection angle. If the instrument function of an electron beam was known it could be deconvolved from the observed lineout giving a truer idea of the extent to which the beams were monoenergetic.

The deconvolution problem

The following assumptions were made:

1. The instrument function, $P(x)$ is the profile of the electron signal in the undispersed dimension y .
2. The shape of a monoenergetic beam is not significantly distorted by deflection within the spectrometer.

The first assumption breaks down if the beam does not fill the aperture evenly and is not symmetric. In the majority of cases the beam filled the aperture evenly, and where it did not its profile was usually symmetric. The second assumption is equivalent to saying that a perfectly monoenergetic beam symmetrical in x and y should be symmetrical in x and y after dispersion by the bending magnet. The particle tracking code of [68] reveals that at high electron energies, where deconvolution is required, the assumption is reasonable.

The measured signal $I(x) = \int s(x, y)dy$ and the ‘underlying’ signal $O(x)$ are related by

$$I(x) = \int_{-\infty}^{\infty} P(x - x_1)O(x_1)dx_1 = P(x) \otimes O(x) \quad (3.4)$$

This problem has previously been addressed in other work [7, 8], where a Gaussians were fit to the electron signal in both the dispersed and undispersed dimension, allowing equation (3.4) to be solved analytically. However in that case the spectrometer dispersion was strong enough that individual features in the spectrum (such as the peaks and the low energy tail) did not overlap. It can be seen in figure 3.7 that a Gaussian fit would not have been appropriate for the data under consideration here. Whereas in the cited case deconvolution was used only to improve the estimate of the electron energy spread, in the case of this experiment the spectrometer dispersion was so low at the energies reached that deconvolution was required to even resolve individual features within the spectrum.

Theoretically this problem could be solved exactly using convolution theory. However, that approach does not work in practice because high frequency (noise) elements are amplified in Fourier space meaning that the resultant form of $O(x)$ is determined by the noise in $I(x)$.

Deconvolution for astronomy

The problem of deconvolving an image from a noisy sample using a known instrument function is common in Astronomy [72]. A typical approach to deal with the problem presented by noise is to use Poisson statistics to estimate the probability of $I(x)$ given a known $O(x)$ and $P(x)$, and then

maximise this probability with respect to $O(x)$. An iterative algorithm results of the form

$$O_{n+1}(x) = O_n(x) \left[\frac{I(x)}{(P(x) \otimes O_n(x))} \otimes P^*(x) \right] \quad (3.5)$$

where $P^*(x) = P(-x)$. This *maximum likelihood* solution is known as the *Richardson-Lucy* algorithm [70, 71] and is best known for its use in restoring images from the Hubble space telescope. The results of applying this algorithm are shown in figure 3.7. Note that this only satisfies equation (3.4) approximately. Since the instrument function is wider than the fluctuations of the noise it is impossible to satisfy equation (3.4) exactly. Instead convolving $O(x)$ and $P(x)$ should result in a ‘fit’ through the noise of $I(x)$. The condition satisfied by (3.5) is

$$I(x) \approx P(x) \otimes O(x) \quad (3.6)$$

The estimate of $O(x)$ that results from applying (3.5) clearly satisfies the condition of equation (3.6) but it is clear from figure 3.7 that the result is physically nonsensical — the low background signal of $I(x)$ has been deconvolved into a series of ‘monoenergetic’ peaks that extend to energies beyond infinity. Whilst this algorithm may well be suited to deconvolving astronomical data it is not directly applicable to dispersed electron beam signals.

Deconvolution for electron spectrometer images

Equation (3.5) does not converge. For no significant change in $P(x) \otimes O_n(x)$ considerable changes $O_n(x)$ can occur - the problem is ‘ill posed’ and there is no unique solution. Many solutions can be discounted as unphysical, but after that a range of diverse but physically satisfactory solutions remain. The choice of when to stop iterating would be little more than a matter of choosing a preferred solution, and this is clearly not sufficiently thorough for the result to be meaningful.

It could be argued that a unique solution to smooth fits through $I(x)$ and $P(x)$ can be found using convolution theory, but this would not change the demonstrated fact that many solutions can exist. The ‘unique’ solution would be strongly dependent on which particular fit was put through the data, and as such it would be wrong to treat this as a useful result. The same is true of windowing in Fourier space to overcome the problem of high frequency amplification - this is the same as putting a smooth fit through the data and the ‘unique’ solution would be strongly dependent on the form of the Fourier window. The same applies to any method that uses the ‘built in’ deconvolution algorithms in programs such as Matlab.

Despite this, attempting to deconvolve the electron dispersion is not pointless. If it is accepted that an array of physically valid solutions exist then at least deconvolution can be used to calculate the upper limit on the spread of the spectrum, that is the ‘widest possible’ solution that satisfies equation (3.4). The peak width measured from the unprocessed spectrum is an upper limit to the real energy spread. However $O(x) = I(x)$ does not satisfy equation (3.4) meaning that a narrower upper estimate of the energy spread can be found. The ‘widest possible’ solution should be the solution for $O(x)$ that is closest to $I(x)$.

An obvious way of finding this solution would be to set $O_0(x) = I(x)$ and then stop the iterative process as soon as equation (3.6) is satisfied to whatever level of precision is determined sufficient.

However, along x the condition of equation (3.6) is satisfied *locally* after differing numbers of iterations. For most x the condition is satisfied by simply saying $O(x) = I(x)$ and not iterating. At the peaks however, equation (3.6) is not satisfied by $O(x) = I(x)$. The iterations could be stopped once the condition has been satisfied at the peaks — but in that case for the x at which the condition was satisfied before iteration, a peaked structure such as that in figure 3.7 emerges.

If the iteration is started at $O_0(x) = I(x)$ and then slowed for x where $P(x) \otimes O_{n+1}(x) \approx P(x) \otimes O_n(x)$ (where iteration does not improve the solution) but allowed to continue unconstrained where the above condition is not met then a ‘widest possible’ solution, that is the least deviated from $I(x)$, emerges. An iterative algorithm that stably achieves this effect was devised and is detailed below.

$$O_{n+1}(x) = \eta(x)Q_n(x) + [1 - \eta(x)]O_n(x) \quad (3.7)$$

$$Q_n(x) = O_n(x) \left[\frac{I(x)}{(P(x) \otimes O_n(x))} \otimes P^*(x) \right] \quad (3.8)$$

$$\mu(x) = \tanh \left[\frac{\delta(x)^2}{\sigma_{\text{im}}^2} \right] \quad (3.9)$$

$$\delta(x) = (Q_n(x) \otimes P(x)) - (O_n(x) \otimes P(x)) \quad (3.10)$$

$$O_0(x) = I(x) \quad (3.11)$$

Equation (3.7) weights $O_{n+1}(x)$ between the Richardson-Lucy iteration, $Q_n(x)$ and $O_n(x)$ depending on the value of $\mu(x)$. Where $\eta(x)$ is approximately equal to 1, $O(x)$ evolves with continued iteration, for $\mu(x) \approx 0$ the local value of $O(x)$ does not evolve. The variance σ_{im}^2 is a measure of the background noise in the image, determined from a subset of the CCD image where it is known that there will be no electron signal. The parameter $\delta(x)$ measures how much the right hand side of equation (3.6) evolves with a step of the Richardson-Lucy algorithm. Where $\frac{\delta(x)^2}{\sigma_{\text{im}}^2} \ll 1$ the evolution is insignificant compared to the noise in the signal, meaning that equation (3.6) is satisfied. The resulting value of $\mu(x)$ is small. Where $\frac{\delta(x)^2}{\sigma_{\text{im}}^2}$ is not insignificant, meaning that equation (3.6) is not satisfied, $\mu(x) \sim 1$, allowing iteration to continue.

Iterations performed with equations (3.7–3.11) converge, and the solution found for the example being considered is plotted in 3.7. The FWHM half max of the deconvolved energy spectrum is 20 %. Because the algorithm of (3.7–3.11) found the most pessimistic solution to equation (3.6) it is possible to say that $\frac{\Delta E}{E} < 20\%$, whereas without deconvolution this upper limit was 50 %. Although other solutions to (3.6) with narrower energy spreads exist, they are effectively meaningless because there is no way of determining that they are the real solution. It is also important to note what can be seen qualitatively in the deconvolved spectrum: that there is a definite peak in the spectrum that is distinct from the background (unlike in the undeconvolved case). One would suspect from 3.3 that there was a monoenergetic peak in the spectrum of the beam, however this method puts that intuitive inference on a firmer footing.

For the data being analysed the method outlined above was useful to go from ‘very low’ resolution to ‘quite low’ resolution. It would be interesting to know if equations (3.7–3.11) can be applied to improve ‘quite high’ resolution to ‘very high’ resolution.

3.3 Statistical methods of analysis

As already stated there were over 10,000 ‘shots’ of data recorded during the Astra experiment. It would therefore be reasonable to expect all areas of ‘parameter space’ to be covered with a large number of data points making a statistical analysis straightforward. The variable parameters were: capillary length L ; capillary diameter D ; hydrogen pressure p ; discharge voltage V ; discharge timing t ; laser energy and laser pulse length. That is a total of seven dimensions to parameter space meaning that the data spread was thinner than might be expected, and did not cover all possible combinations of conditions — this is particularly true in cases where post-analysis of the data revealed certain parameters (such as D and t) to have a larger effect than expected. Moreover for the areas of parameter space that were covered by data the spread was not always even: for example if during a run it was found that an electron beam was generated for a particular delay more shots would be fired at similar values of t in order to produce a statistically useful number of electron beams. However it could have been the case that there was a wider range of delays which were just as conducive towards beam generation, but due to the manner in which it was collected the data would give an impression of exaggerated importance to the ‘magic’ delay. It is therefore important that a statistical approach was used that highlight genuine trends in the physical behaviour whilst not being affected by the bias arising from the manner in which data was collected.

3.3.1 Averaging for varying t

The effect of varying the delay t on the transmitted radiation and accelerated electron beam is of particular interest in chapters 4 and 6. For comparing the delay dependence of a particular measurement Y for two different sets of data, each composed of many discrete data points, it was useful to produce a continuous function $Y(t)$. To achieve this a Gaussian average was used with the following form

$$Y(t) = \frac{\sum_{i=1}^N Y_i e^{-\left[\frac{t-t_i}{\tau}\right]^2}}{\sum_{i=1}^N e^{-\left[\frac{t-t_i}{\tau}\right]^2}} \quad (3.12)$$

where $\{t_i, Y_i\}$ is a set of data for the conditions of interest and i is the laser shot number. The parameter τ defines the width of the averaging window, and was set to 5 ns throughout the analysis.

3.3.2 Averaging for varying n_e

The interaction between the laser and the plasma channel varied with density n_e . It was useful, therefore, to plot measurements, Y , against density, n_e . The obvious way to do this would be to plot all of the points of a set of data $\{n_{e,i}, Y_i\}$ or to apply equation 3.12 with n_e in place of t to form an average $Y(n_e)$. However as the values of different parameters were scanned much data was recorded at delays for which a plasma channel was not formed. The measured Y for shots at those delays were quite different from those recorded when the plasma channel was formed. As a result any plot of $\{n_{e,i}, Y_i\}$ would have a large scatter of values at each value of n_e , and the calculated average $Y(n_e)$ would have been meaningless.

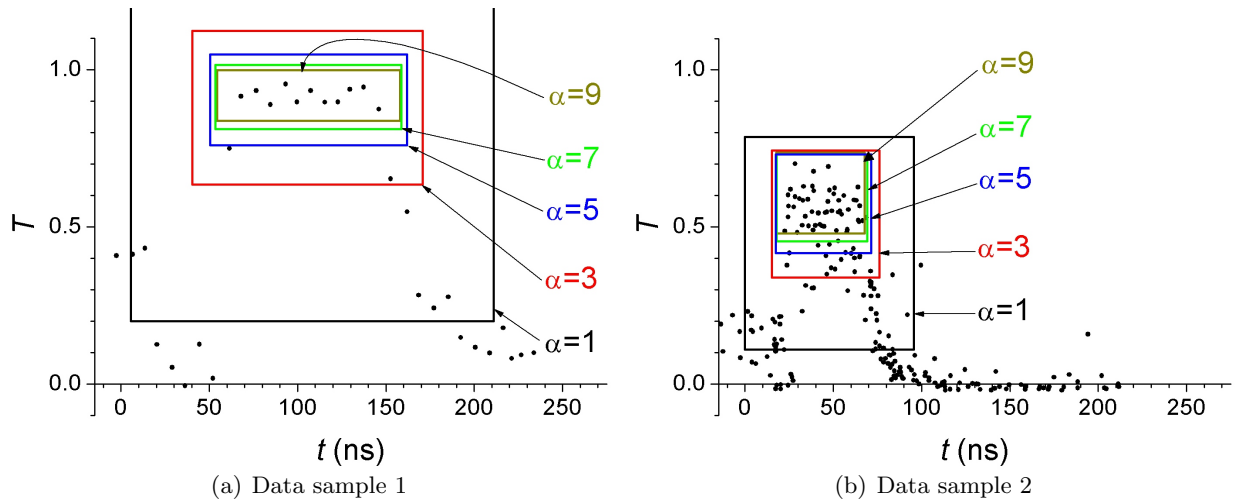


Figure 3.8: Applying transmission biased averaging to two samples of real data. Two datasets of $\{t_i, T_i\}$ are plotted (dots). Equations (3.13–3.15) were applied to the data for different values of α . The results are represented by the coloured squares: each is $4\sigma_t$ wide and $4\sigma_T$ high and centered at the calculated averages of t and T . For the data of figure (a) $N_{\text{eff}} = 20.3, 14.1, 12.5, 11.8, 11.3$ when $\alpha = 1, 3, 5, 7, 9$.

This is well illustrated by the data points in figure 3.8, which are examples of two data sets $\{t_i, T_i\}$ with fixed n_e . When looking at these graphs it is clear that the waveguide transmission is ~ 0.9 for the data in figure (a) and ~ 0.5 (with a certain amount of scatter) for figure (b). At delays outside of the transmission peak it is reasonable to assume that a plasma channel is not formed⁴. Were the datasets to form a point on a graph of T against n_e , whether as a group of points or as an average, the resulting plot would be misleading due to the large number of points at non-optimal delays.

Ideally data could be filtered before averaging to only include data within a range of delays for which the plasma channel was formed. This was not straightforward because the t at which T was highest varied with changes in conditions. This meant that the 600 runs recorded would need to be individually examined to determine a range of useful t for each — that would be very time consuming. Alternatively all shots for which T was below a certain threshold could have been excluded from the average. However for data sets for which the optimum T was relatively low, such as that in figure 3.8(b), all data would be excluded. Such a method would be unsatisfactory.

Transmission biased averaging

The desired ‘filtering’ of the data could be achieved using statistical methods. A useful mean, with no significant contribution from data at the ‘wrong’ delays, could be found by using an approach similar to equation (3.12) but weighted towards shots for which T was high. The following equations

⁴This assumption is confirmed by the interferometry measurements of chapter 6.

were defined to achieve such a result.

$$Y(n_e) = \frac{\sum_{i=1}^N Y_i T_i^\alpha e^{-\left[\frac{g-g_i}{\Delta g}\right]^2}}{\sum_{i=1}^N T_i^\alpha e^{-\left[\frac{g-g_i}{\Delta g}\right]^2}} \quad (3.13)$$

$$\sigma_Y(n_e) = \sqrt{\frac{N_{\text{eff}}(n_e)}{N_{\text{eff}}(n_e) - 1} \sqrt{Y^2(n_e) - [Y(n_e)]^2}} \quad (3.14)$$

$$N_{\text{eff}}(n_e) = \frac{1}{[T(n_e)]^\alpha} \sum_{i=1}^N T_i^\alpha e^{-\left[\frac{g-g_i}{\Delta g}\right]^2} \quad (3.15)$$

where $g = \log_{10}(n_e)$ and $\Delta g = 0.0212$ was the width of the averaging window⁵. The parameters $\sigma_Y(n_e)$ and $N_{\text{eff}}(n_e)$ are the standard deviation and effective number of shots within the average at n_e . The parameter α determines the degree of bias towards high T . A suitable value for α is determined below.

Equations (3.13–3.15) were applied to data-sets $\{n_{ei}, T_i, Y_i\}$, each of which included all data for a particular combination of L , D and pulse length. The Gaussian part of equation (3.13) provided a statistical window centered at a density n_e , and the factor T_i^α weights the average towards the highest transmission shots within that window. Equation (3.14) is simply an expression of the fact the uncertainty within a sample of data is the standard deviation multiplied by $\sqrt{N/(N-1)}$ where N is the number of shots. The effective number of shots N_{eff} calculated in equation (3.15) is simply the bottom of the right hand side of equation (3.13) normalised to the maximum statistical weight at that density, which is $T(n_e)^\alpha$.

Determining α

It is expected that α should be large since the aim was to prevent data shots for which T was relatively low from contributing to the average $Y(n_e)$. However should α have been too large the data would not have just been weighted towards the plateau of high T , but to those shots within the plateau that happened, through scatter in the measurement of T , to have a slightly higher transmission than others. The consequence of this would be to significantly reduce N_{eff} with many useful data points being excluded. Additionally the value of $T(n_e)$ would be an overestimate.

An appropriate value of α can be determined by using equations (3.13–3.15) to examples of real data, such as that plotted in figure 3.8. For both data-sets averaged values of T and t were calculated, as well as the uncertainties σ_T and σ_t , with different α . The boxes plotted on figure 3.8, each of which represents a different value of α , are $4\sigma_t$ wide, $4\sigma_T$ high and centered at the average calculated for t and T . For a normal distribution 98.2% of a sample lies within two standard deviations of the mean. That means that the range of delays included by the box should be approximately the width of the plateau of highest transmission of the data. The vertical position and size of the boxes should be the average value and range of the high transmission plateau respectively.

For $\alpha > 5$ the range of delays remains approximately fixed so it can be concluded that at that

⁵The use of the log of as an x -parameter rather than the density reflects the fact that fractional changes in density are of more interest than absolute changes. The chosen value of Δg corresponds to a Gaussian width $\Delta n_e \approx 0.05n_e$.

value of α and higher shots outside of the high transmission plateau do not contribute significantly to the mean. In figure 3.8(a) there are 11 shots within the plateau of high T . The calculated values of N_{eff} are only similar to this for $\alpha > 5$ suggesting that $\alpha < 5$ would not have been appropriate. In the case of 3.8(b) the mean value of T for $\alpha = 9$ appears to be an overestimate, and so it is concluded that that level of weighting is excessive. Therefore a value of $\alpha = 7$ was used⁶.

Filtering across conditions n_e

Where $N_{\text{eff}}(n_e)$ was low the average produced by equation (3.13) is effectively meaningless, and that would result in the large, infinite or imaginary values of $\sigma_Y(n_e)$. Rather than producing a continuous function, peaks in $N_{\text{eff}}(n_e)$ were detected above a threshold of $N_{\text{eff}} = 3$ and at these values of n_e discrete data points were output. This means that the process of transmission biased averaging was not just a method for filtering data and averaging across conditions — it was also an efficient means to identify the conditions for which a useful amount of data existed.

To illustrate the relationship between raw data and the averages produced by this method both are plotted in figure 3.9, with $Y = t$. The averages $t(n_e)$ and the ranges $\pm 2\sigma_t(n_e)$ are in agreement with the trends in the raw data. The process has effectively combined the data in the range $n_e = 3 - 4 \times 10^{18} \text{ cm}^{-3}$, which was not all at a single density, to a single point. On the other hand there is no value of $t(n_e)$ for $n_e \approx 9 \times 10^{18} \text{ cm}^{-3}$ even though data existed for that density. This is because after filtering $N_{\text{eff}}(n_e) < 3$, meaning that an average of the high transmission plateau is meaningless for that data. The process of transmission biased averaging has filtered 93 shots from 650 points of data in figure 3.9, and then combined those to create 5 data points at the densities at which it has determined (automatically) that there is sufficient data to create a meaningful average.

This method was also applied to data recorded at negative t , but in that case α was set to zero since there should be no delay dependence on the observed behaviour before the firing of the discharge, and so accordingly no need to weight towards a high plateau in transmission.

3.3.3 Beam probability

As discussed in chapter 1, improving the stability of the electron beams produced by laser wakefield acceleration is an area of current focus. It was therefore of interest to calculate the probability P that a beam would be produced for a particular set of conditions. When using the hydrogen-filled capillary discharge waveguide there was a (sometimes sensitive) dependence of electron generation on t , which can be described by a delay dependent probability of the form $P(t)$. If $P(t)$ could be deduced from the data then the beam stability P_0 (the optimum value of $P(t)$) and the sensitivity S of beam generation to delay (which we define as the inverse of the standard deviation of $P(t)$) could both be determined.

Gaussian averaging

An obvious way of estimating $P(t)$ from the results would be to define a parameter B that is equal to 1 for those shots for which an electron beam was observed, and 0 for all others, and to then take

⁶Were this method to be applied to data from another experiment another value might be appropriate given that the scatter in T may be different

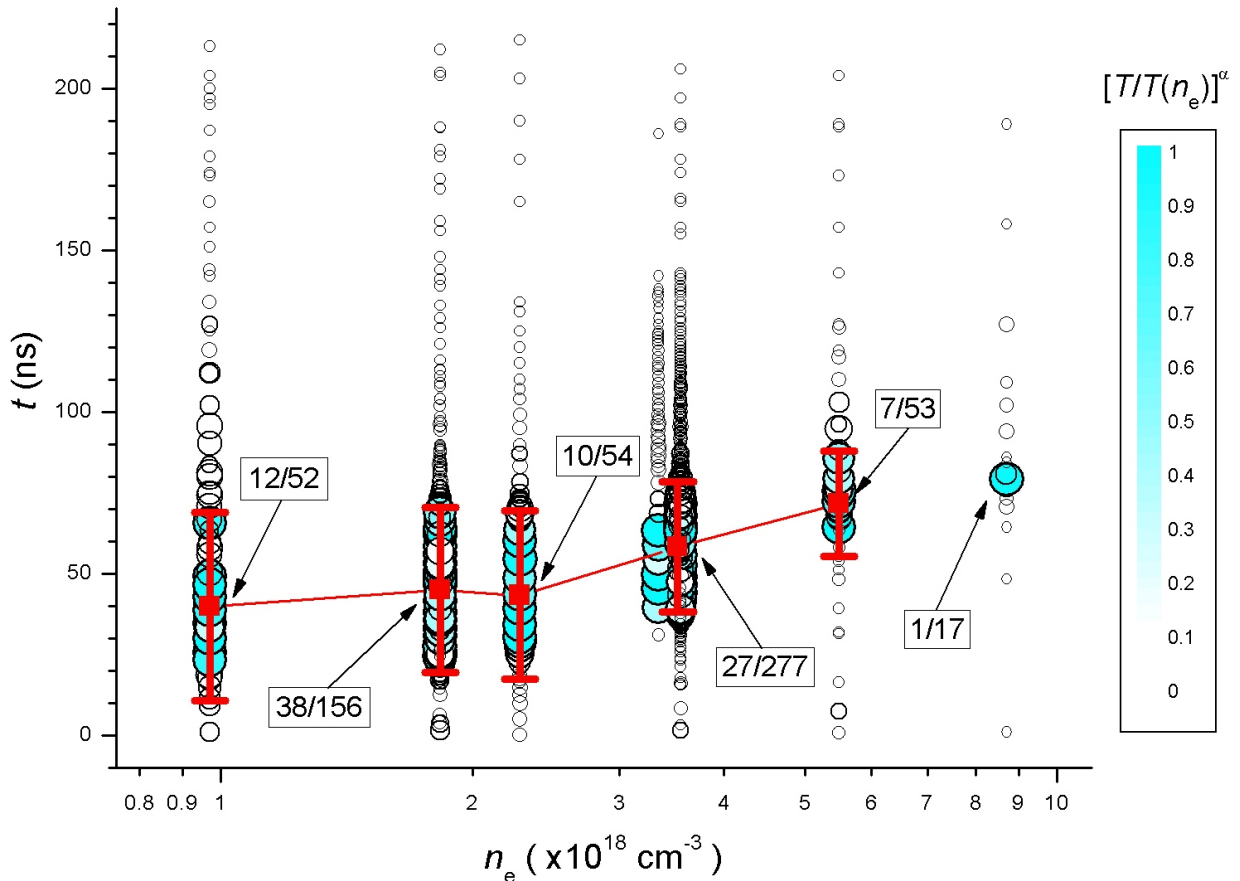


Figure 3.9: Applying transmission biased averaging across a range of densities.

The results of applying equations (3.13), (3.14) and (3.15) are plotted ($\alpha = 7$). Each circle represents a shot from the data-set $\{n_{ei}, T_i, t_i\}$. Note that t and n_e are independent variables so the position of the circles merely represents points at which data was collected. The radii of the circles represent the transmission of that shot and are proportional to $T_i/T(n_e) + 0.5$ (the constant is in place so that shots for which $T_i \approx 0$ are still visible). The fill-colour of the circles represents the statistical weight $[T_i/T(n_e)]^\alpha$. The red squares are the value of $t(n_e)$ at values of n_e at which $N(n_e)$ is peaked and greater than 3. The error bars are $\pm 2\sigma_t(n_e)$. The fractions in the boxes are $N_{\text{eff}}(n_e, \alpha = 7)/N_{\text{eff}}(n_e, \alpha = 0)$, that is the effective number of shots in the transmission biased average at the density indicated over the effective total number of shots at that density. This illustrates the extent to which the data is filtered through the process outlined above.

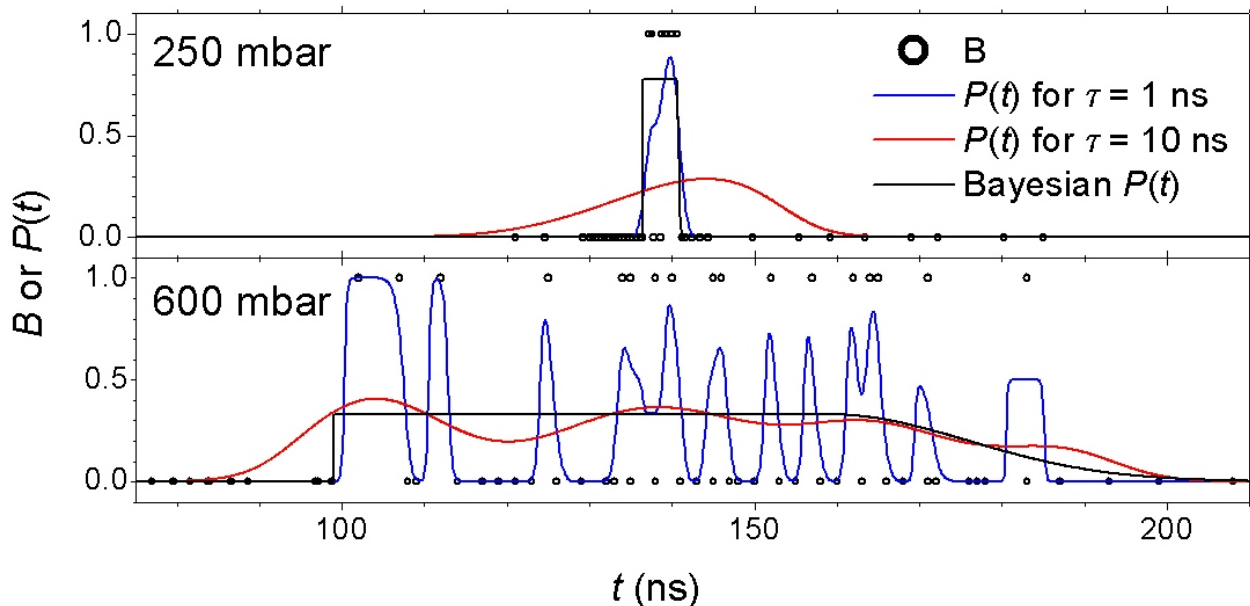


Figure 3.10: Comparison of methods for calculating $P(t)$.

For two sets of real data the binary parameter B (circles) and estimates of $P(t)$ (lines) are plotted against t . $P(t)$ was estimated using: Gaussian averaging with $\tau = 1$ ns (blue); Gaussian averaging with $\tau = 10$ ns (red); Bayes theorem (black).

a time dependent average of B using equation (3.12).

Figure 3.10 shows examples of $P(t)$ calculated in this way for (real) data at two different pressures, using $\tau = 1$ ns and $\tau = 10$ ns. The resulting values of P_0 and S are included in table 3.2. It is clear that the form of $P(t)$, and the resulting values of P_0 and S , depend strongly on τ . Ideally the value of τ chosen should be small enough to resolve any sharp changes in behaviour with respect to t , but large enough to include a reasonable number of shots in the effective average at each value of t . The estimate of $P(t)$ using $\tau = 1$ ns seems reasonable for the case of 250 mbar, whereas $\tau = 10$ ns is clearly too wide leading us to the conclusion that the generation of electron beams was less stable and less sensitive to timing than it was in reality. However for 600 mbar using $\tau = 1$ ns gives a result that is very strongly a function of the sampling density of data than of the data itself, and that the effective number of shots in the average is approximately 1 for many delays. In this case using $\tau = 10$ ns gives what seems like a reasonable estimate of $P(t)$.

Ideally in all cases the same density of data as in the 250 mbar case would have been available and then $\tau = 1$ ns could be employed consistently. However, with the limited rate of repetition rate of the system it would have been impossible to cover cases such as the 600 mbar data of figure 3.10 where the range of delays over which electrons were generated was considerably wider. Given that our values of P_0 and S are dependent on the value of τ chosen in our Gaussian average, it is not appropriate to use the method outlined above.

Bayesian approach

A method for estimating $P(t)$ was required that works equally well for all sets of data without being tailored to each individually. A way of achieving this is to find the most probable form of $P(t)$

		Dataset	
		250 mbar	600 mbar
Gaussian Average, $\tau = 1$ ns	P_0	0.89	1.0
	S (ns ⁻¹)	0.34	0.017
Gaussian Average, $\tau = 10$ ns	P_0	0.29	0.41
	S (ns ⁻¹)	0.052	0.015
Maximum $P(E P_0, t_a, t_b, \sigma_a, \sigma_b)$	P_0	0.78	0.33
	S (ns ⁻¹)	0.34	0.018

Table 3.2: Comparison of methods for calculating P_0 and S .

These numbers are calculated from the estimates of $P(t)$ plotted in figure 3.10. Numbers that are inconsistent with those calculated using the other two methods are highlighted in red.

given the available data. It is straightforward to estimate the probability of a set of data $\{t_i, B_i\}$ given a particular form of $P(t)$. This is written as

$$P[\{t_i, B_i\}|P(t)] = \prod_j P(t_j) \times \prod_k [1 - P(t_k)] \quad (3.16)$$

where t_j are the delays of those shots for which beams were observed, and t_k are the delays of the shots for which a beam was not observed. The expression required is $P[P(t)|\{t_i, B_i\}]$, the probability of a hypothesis for $P(t)$ given the data $\{t_i, B_i\}$. Maximising that would give the most probable form of $P(t)$.

Bayes theorem [73] provides a means of converting from $P[\{t_i, B_i\}|P(t)]$ to $P[P(t)|\{t_i, B_i\}]$. If a particular *experience* is E and a particular *hypothesis* used to describe that experience is H_0 then

$$P(H_0|E) = \frac{P[E|H_0]P[H_0]}{P[E]} \quad (3.17)$$

The term $P[H_0]$ is known as the *prior probability* and describes the probability that a particular hypothesis would be chosen given no prior information. $P[E]$ is the *marginal probability*, the probability of witnessing evidence E for all mutually exclusive hypotheses, meaning that $P[E] = \sum P[E|H_i]P[H_i]$. Since $P[E]$ is constant with respect to H_0 it is ignored.

Intuitively one might expect that the prior probability should be set equal to 1 for all H_0 . Philosophically it is impossible to argue that some hypotheses are more probable than others before even considering the data. Maximising equation (3.17) with $P[H_0] = P[P(t)] = 1$ is equivalent to maximising (3.16). The most probable form of $P(t)$ calculated from equation (3.16) is a function that is 1 at all t for which there were beams recorded, 0 for all t for which there were shots recorded with no beam, and undetermined for all t between. This is equivalent to the Gaussian solution for $\tau \rightarrow 0$. Common sense would suggest that this is not a realistic solution. That form of $P(t)$ is simply a recreation of $\{t_i, B_i\}$, without having condensed any information.

A non-uniform prior probability was necessary to deduce a believable estimate of $P(t)$. It may




Figure 3.11: The form used for making Bayesian estimates of $P(t)$.

seem that this prejudices the estimate of $P(t)$ (which is exactly what we hoped to avoid), but it is worth noting that the choice of any method for converting $\{t_i, B_i\}$ to an estimate $P(t)$ biases the results. The advantages of proceeding using Bayes theorem is that the prior expectations (prejudices) are formalised explicitly. So, the prior expectations for $P(t)$ are these: at very early delays $P(t) = 0$; at very late delays $P(t) = 0$; between those extreme values of t , $P(t)$ may rise and then drop away again; the rise and drop in $P(t)$ may be sharp or gradual, and there may be delays inbetween for which $P(t)$ remains steady.

The expected form of $P(t)$ is very general, and can be describe mathematically in the following way

$$P(t) = \begin{cases} P_0 \exp \left[-\frac{(t-t_a)^2}{2\sigma_a^2} \right] & ; t < t_a & \textcircled{1} \\ P_0 & ; t_a < t < t_b & \textcircled{2} \\ P_0 \exp \left[-\frac{(t-t_b)^2}{2\sigma_b^2} \right] & ; t > t_b & \textcircled{3} \end{cases} \quad (3.18)$$

An example of $P(t)$ with this form is plotted in figure 3.11. The function $P(t)$ is defined by 5 numbers; P_0 , t_a , t_b , σ_a and σ_b . Effectively $P[P(t)]$ is 0 for all $P(t)$ that cannot be described by these 5 numbers, and finite for all $P(t)$ that can. Combining that with equations (3.16) and (3.17) gives

$$P[P_0, t_a, t_b, \sigma_a, \sigma_b | \{t_i, B_i\}] = \frac{\prod_j P(t_j) \times \prod_k [1 - P(t_k)]}{\int \int \int \int \prod_j P(t_j) \times \prod_k [1 - P(t_k)] dP_0 dt_a dt_b d\sigma_a d\sigma_b} \quad (3.19)$$

An estimate of $P(t)$ can be found by maximising the top part of the right hand side of equation (3.19). A code was written to perform that operation. The results of using this method for the example data are plotted in figure 3.10, and the deduced values of P_0 and S are given in table 3.2.

For the 250 mbar data the Gaussian averaging method using $\tau = 1$ ns gave a believable curve $P(t)$, for 600 mbar this was achieved using $\tau = 10$ ns. The values of $P(t)$, P_0 and S calculated using those methods are both in approximate agreement with the values produced using the Bayesian method. It was concluded, therefore, that the Bayesian method is ideal for estimating $P(t)$, because the results produced were reasonable irrespective of the density of data in the sample $\{t_i, B_i\}$.

Chapter 4

Results from Astra experiments

The experimental setup and methods of data analysis for the Astra experiment are described in chapter 3. In this chapter the results of that work are presented, and significant trends are highlighted. Measurements of the transmitted laser pulse and measurements of the accelerated electron beams are first presented separately, before evidence demonstrating a correlation between them is given. The condition of two used capillaries is also documented for future reference. Due to the large volume of results the emphasis of this chapter is placed on presenting a full body of experimental evidence, with discussion and further investigation of the less obvious findings reserved for chapters 5 and 6.

4.1 Measurements of the transmitted laser pulse

4.1.1 Guiding

In figure 4.1 the waveguide transmission T of the short pulse¹ is plotted against delay t for the conditions [$L = 50$ mm, $D = 250$ μ m, $p = 160$ mbar, $V = 25$ kV]. Measurements of the transmitted fluence exit modes are included. The high peak transmission and compact exit modes for $t = 50 - 80$ ns show that beams of intensity $\sim 10^{18}$ W cm⁻² can be guided over at least $14Z_R$ using the hydrogen-filled discharge capillary waveguide. The fact that $T \approx 0.3$ for negative delays but $T \approx 1$ after the onset of the current demonstrates that the laser spot was guided by the plasma channel (formed by the current) and not by the walls of the capillary (present for all delays). The measured transmitted fluence was $\sim 1.6 \times 10^4$ J cm⁻², which is around half of the estimated input. The analysis of appendix B suggests that for the short pulse the difference between input and measured output fluence may have been due to chromatic aberration in the lens system imaging the waveguide exit. Estimates, which are included in appendix B, show that for the data of figure 4.1 the transmitted fluence profile was very similar to the input fluence profile.

Varying L and D

The effect of varying the length and diameter of the waveguide is seen in figure 4.2. Data for both the long pulse and short pulse are included. For all capillaries the peak long pulse transmission

¹Unfortunately no long pulse data was recorded for these conditions.

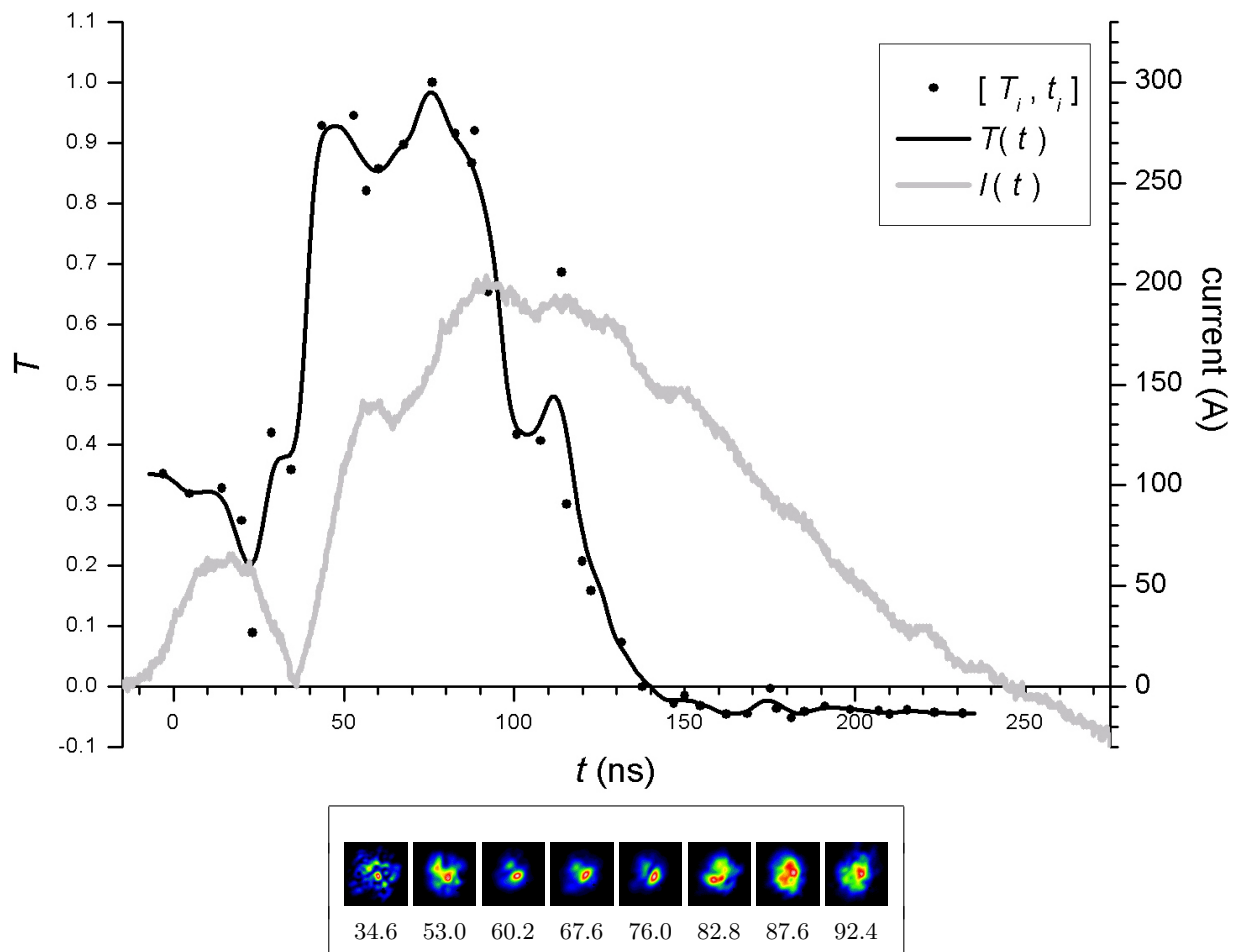
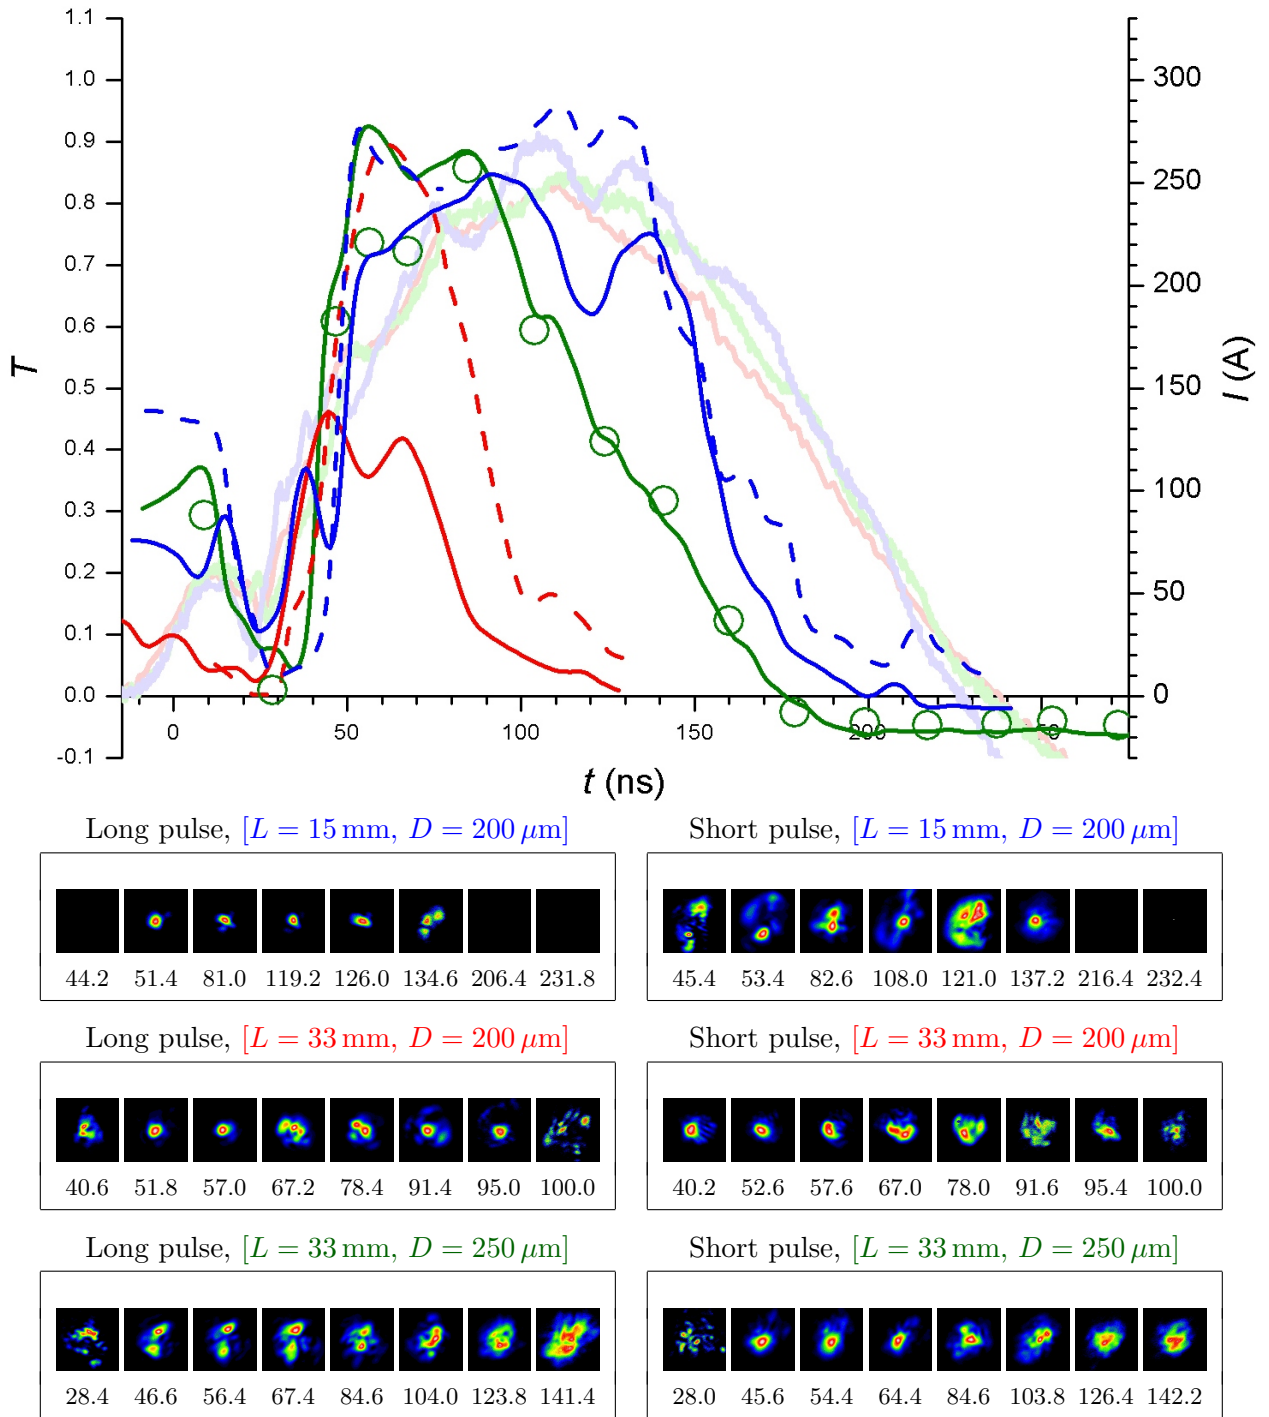


Figure 4.1: Short pulse guiding data for $[L = 50 \text{ mm}, D = 250 \mu\text{m}, p = 160 \text{ mbar}, V = 25 \text{ kV}]$. Transmission data $\{t_i, T_i\}$ (black dots), averaged curve $T(t)$ (black line), discharge current $I(t)$ (grey line) and selected exit plane fluence profiles are plotted. Each fluence profile plot represents a $200 \mu\text{m}$ square in the plane of the waveguide exit and the number below is t_i in ns.


 Figure 4.2: Guiding data for varying L and D .

Transmission T for the short pulse (solid lines) and long pulse (dotted lines). The waveguide conditions were: [$L = 15 \text{ mm}$, $D = 200 \mu\text{m}$, $p = 160 \text{ mbar}$, $V = 25 \text{ kV}$] (blue); [$L = 33 \text{ mm}$, $D = 200 \mu\text{m}$, $p = 160 \text{ mbar}$, $V = 30 \text{ kV}$] (red); [$L = 33 \text{ mm}$, $D = 250 \mu\text{m}$, $p = 160 \text{ mbar}$, $V = 30 \text{ kV}$] (green). In all cases the averaged form of T , $T(t)$ is plotted to maintain clarity, with the exception of the long pulse [$L = 33 \text{ mm}$, $D = 250 \mu\text{m}$] data where there were insufficient shots to form an average. Instead the data-set $\{t_i, T_i\}$ is used (green hoops). For each set of conditions the corresponding measurements of current (faint lines) and exit plane fluence are included. Each fluence profile plot represents a $200 \mu\text{m}$ square in the plane of the waveguide exit and the number below is t_i in ns.

was high ($T \approx 0.9$). The peak short pulse transmission, on the other hand, was affected by changes to L and D . For [$L = 33$ mm, $D = 250$ μm] the short pulse transmission was equal to that of the long pulse. That, as well as the high transmission observed in figure 4.1 for [$L = 50$ mm, $D = 250$ μm], demonstrates that the short pulse was not depleted during propagation through 250 μm wide capillaries, suggesting that for those waveguides interactions with the plasma were weak. For $D = 200$ μm the short pulse transmission was relatively low. The losses were greater after 33 mm than after 15 mm, suggesting that in these narrower capillaries the short pulse was losing energy during propagation. Because the corresponding long pulse transmission was high, the energy loss of the short pulse cannot be attributed to inefficient guiding, leading to the conclusion that in 200 μm capillaries the laser pulse was depleted through interaction with the plasma.

Figure 4.2 also contains a sample of output fluence profiles for each of the plotted curves. For both pulse lengths these exit modes were compact and symmetric when considering the $D = 200$ μm capillaries (at optimum delays). The equivalent long pulse exit modes for [$L = 33$ mm, $D = 250$ μm] appeared to be of lower quality, with a double peak appearing. It could be that the plasma channel that formed inside the wider capillary was not suitable for guiding a 34 μm spot. However the corresponding short pulse exit modes were compact and symmetric, suggesting that the plasma channel was not the problem. Fluence exit modes were more sensitive than transmission to waveguide alignment, so it may be that the poor quality of the long pulse exit modes for [$L = 33$ mm, $D = 250$ μm] were the result of a slight misalignment when that data was recorded.

Rise and fall of transmission

The dependence of T on t observed in figure 4.2 was the same for the short pulse as it was for the long pulse. For all capillaries the transmission increased sharply between $t = 40$ ns and $t = 50$ ns (when $p = 160$ mbar) corresponding, presumably, to the formation of a plasma channel suitable for guiding. For higher p that rise occurred at later delays, which is most likely due to the fact that more hydrogen atoms required ionisation during the channel formation process, requiring more energy from the current.

Theory [60, 62, 63] predicts that once a plasma channel has formed, a quasisteady state exists that is stable for the duration of the current pulse. It was unexpected, then, that T decreased at delays for which the current was still high. From figures 4.1 and 4.2 it would appear that the decrease occurs earlier for longer, narrower capillaries. To quantify this effect the delays t_r and t_f were defined, corresponding to the rise and fall of $T(t)$ above and below 60% of its peak value. The values of t_r and t_f at $p = 160$ mbar are plotted in figure 4.3 for different L and D (with other parameters fixed). Whilst t_r does not change significantly with L and D , t_f does. The fits to the data suggest that for a [$L = 50$ mm, $D = 200$ μm] capillary, t_f would be less than t_r . In such a case the transmission of the waveguide would never be high. In chapter 6 it is established that the unexpected fall in transmission at t_f is connected to an unexpected drop of the ionization level inside the waveguide. It is also found, fortunately, that the limits to the possible length and diameter of the capillary implied by figure 4.3 are not fundamental to the hydrogen-filled capillary discharge waveguide scheme, and guiding should be possible in longer narrower capillaries with a simple adaptation of the waveguide design.




Figure 4.3: Transmission rise and fall delays for varying L and D .

The delays t_r (green) and t_f (red) at $p = 160$ mbar are plotted for capillaries of varying length and diameter. The x -axis parameter is L_c/A , where L_c is the length of the capillary between gas slots, and A is the cross-sectional area of the capillary. This choice of x -axis is explained in chapter 6. Data is included for the short pulse (solid) and long pulse (hollow). The error bars in L_c/A was calculated by assuming a 10 % uncertainty in the capillary diameter. The error bars in t_f represents the delays for which $T(t)$ fell below 0.8 and 0.4 of the peak. Due to the sharp rise in $T(t)$ at the start of the current pulse, error bars were not necessary for t_r . The lines are linear fits to the data. Their formulae are $t_r = 52 \text{ ns} - 1.42 \times 10^{-5} \text{ ns m} \times L_c/A$ and $t_f = 189 \text{ ns} - 1.12 \times 10^{-4} \text{ ns m} \times L_c/A$.

4.1.2 Transmitted laser spectra

In chapter 1 the ways in which interaction with a plasma can distort the spectrum of an intense laser pulse were discussed. Analysis of the transmitted laser spectra can provide insight into the processes occurring as the laser pulse propagates through the waveguide. In table 1.1 and figure 1.3 several ways in which the laser spectrum might be modified by interacting with a plasma were described. Examples of all of these can be seen in figure 4.4.

Evidence of the Raman instability for the long pulse

Figures 4.4(a) and (b) show that at delays for optimum guiding ($t_r < t < t_f$), the spectrum of the transmitted long pulse did not broaden significantly, suggesting that wake excitation and photon acceleration were minimal. However peaks were present in these spectra either side of the central wavelength. Interpretation of the frequency downshifted peaks of figure 4.4(a) as Raman satellites leads to an estimated electron density $n_e = 5.85 \times 10^{18} \text{ cm}^{-3}$ whilst the value predicted from equation (2.13) is $5.36 \times 10^{18} \text{ cm}^{-3}$. For figure 4.4(b) the respective numbers are $n_e = 2.07 \times 10^{18} \text{ cm}^{-3}$ and $1.83 \times 10^{18} \text{ cm}^{-3}$. In both cases the estimated and predicted values are similar, confirming that the frequency downshifted peaks in the long pulse spectra should be interpreted as Raman satellites. The Raman satellite for $L = 33 \text{ mm}$ in figure 4.4(a) was ~ 80 times more intense than for $L = 15 \text{ mm}$. Equation (1.26) predicts that the extra length of the 33 mm capillary should have enhanced the satellite intensity by a factor ~ 50 . The agreement between the theoretical and measured ratio of the intensities is reasonable given the non-linear nature of the instability. The downshifted Raman satellite was used to deduce the waveguide axial density and the resulting data is presented in chapter 6.

The origin of the two upshifted peaks in figure 4.4(a) is unclear, but these were common in many of the transmitted long pulse spectra. These peaks appeared to be related because they were only ever observed together. The outer peak might have been the up-shifted Raman satellite because its frequency separation from the central peak was similar to that of the down-shifted Raman peak. However the inner peak was too close to the central laser frequency to be a Raman satellite. The intensity of the peaks was greater for the longer waveguide, by a factor that was approximately equal to the ratio of the capillaries lengths. That suggests that the upshifted peaks grew linearly and not exponentially as they would have had they been Raman satellites.

Photon acceleration and deceleration by the wakefield

Figure 4.4(d) shows the transmitted short pulse spectra for $t_r < t < t_f$ at $p = 250 \text{ mbar}$. For this pressure equation (2.13) predicts that $n_e \approx 5.4 \times 10^{18} \text{ cm}^{-3}$, meaning that for the short pulse, $c\tau/\lambda_p \approx 1$. That is the scenario depicted in figure 1.3(b), where the pulse is overlapping with the wakefield phases that both accelerated and decelerated photons. Accordingly the photon acceleration model predicts that the spectrum would have broadened if there was a wakefield being driven by the laser pulse, and that is observed for both capillary lengths in figure 4.4(d).

In figure 4.4(e) the equivalent data is plotted for $p = 80 \text{ mbar}$, corresponding to $n_e = 2 \times 10^{18} \text{ cm}^{-3}$. At this density $c\tau/\lambda_p \approx 0.55$, so as in 1.3(a) the laser pulse is almost entirely within the

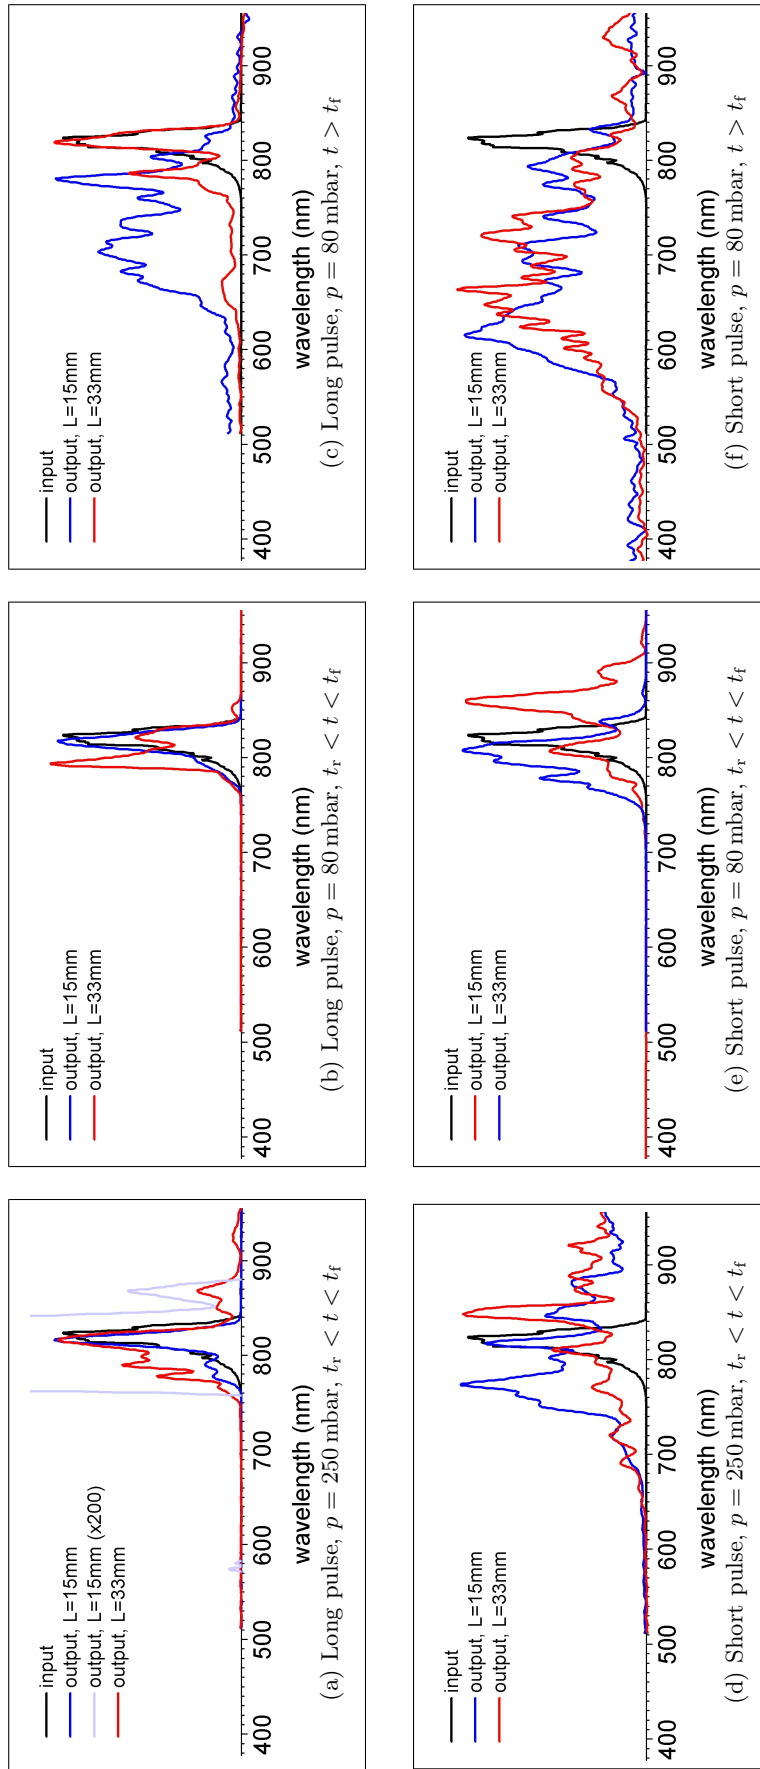


Figure 4.4: Spectra of transmitted laser pulses included in each plot are the input spectrum (black), $[L=15$ mm, $D = 200$ μ m] capillary transmitted spectrum (blue) and the $[L=33$ mm, $D = 200$ μ m] capillary transmitted spectrum (red). For $[p = 80$ mbar, $L = 33$ mm] there was a problem compressing the laser, so that the long pulse was over 450 fs in duration (rather than the usual long pulse duration of 150 fs).

phase of the wakefield that decelerates photons. The photon acceleration model predicts that if the laser pulse was driving a wakefield at this density the whole spectrum would have been red-shifted. Such a red-shift was observed for the 33 mm capillary, but not for 15 mm capillary, in fact there was a slight blue-shifted in that case. An interpretation of this evidence is that a certain degree of pulse evolution through self-focussing and pulse compression was necessary before the laser pulse reached a sufficiently high intensity to drive a significant wakefield, and that evolution length was longer than 15 mm but shorter than 33 mm.

Ionization blue-shifts

Figures 4.4(c) and (f) show the transmitted spectra for $t > t_f$. Both the long pulse and the short pulse spectra were spread towards higher frequencies at delays later than the plateau for optimum guiding. An overall shift of the spectrum towards shorter wavelengths is often the result of an ionization front co-propagating with the laser, as depicted in figure 1.3(c). The blue-shift is consistent with the interpretation of chapter 6 that for $t > t_f$ the ionization level on the waveguide axis is less than unity. However it is noticeable that the magnitude of the spreading is considerably greater than the ionization blue-shift observed in gas jet experiments for similar conditions [23]. It is very difficult to quantify the expected ionization blue-shift because of the growth of instabilities at the ionization front [24]. The magnitude of the blue-shift appears to be increase with the input laser intensity: it is strongest for the short pulse, and is weakest for the $L = 33$ mm long pulse plot in figure 4.4(c), which had a particularly low intensity². It may be that the shifting was greater inside the partially ionised plasma channel than a gas jet because the high intensity was maintained for a longer length.

Density limits to wakefield excitation

The limits Λ^R and Λ^B for transmitted spectra (calculated using the method set out in section 3.3.2) are plotted against plasma density in figure 4.5. A change of the red-limit, Λ^R from its vacuum value towards longer wavelengths corresponds to the spreading of some radiation to the red part of the spectrum. Since that implies that there are photons losing energy, it is assumed to be evidence of wakefield excitation. Therefore the minimum density at which the pulse excites a significant wakefield can be inferred from the density at which there is a noticeable shift in the value of Λ^R to longer wavelengths. For $[L = 33 \text{ mm}, D = 200 \mu\text{m}]$, Λ^R shifts towards the red when $n_e \gtrsim 2 \times 10^{18} \text{ cm}^{-3}$. For $[L = 15 \text{ mm}, D = 200 \mu\text{m}]$ the same occurs for $n_e \gtrsim 4.5 \times 10^{18} \text{ cm}^{-3}$. The dependence on length suggests that at lower densities the evolution phase of the laser propagation is similar to the capillary lengths.

For $[L = 15 \text{ mm}, D = 300 \mu\text{m}]$ there was no red-shift of Λ^R , even at the highest density for which data was taken. That would indicate the importance of the plasma channel, which was weaker in the wider capillary, to allow the excitation of wakefields at plasma densities $n_e \lesssim 15 \times 10^{18} \text{ cm}^{-3}$. These observations can be compared to the equivalent conditions without a plasma channel by considering the plots for Λ^R and Λ^B at negative delays. The red-limit Λ^R only shifted towards the

²When that latter spectrum was recorded the long pulse, due to an operational error, had a longer than usual pulse length that was in excess of 450 fs.

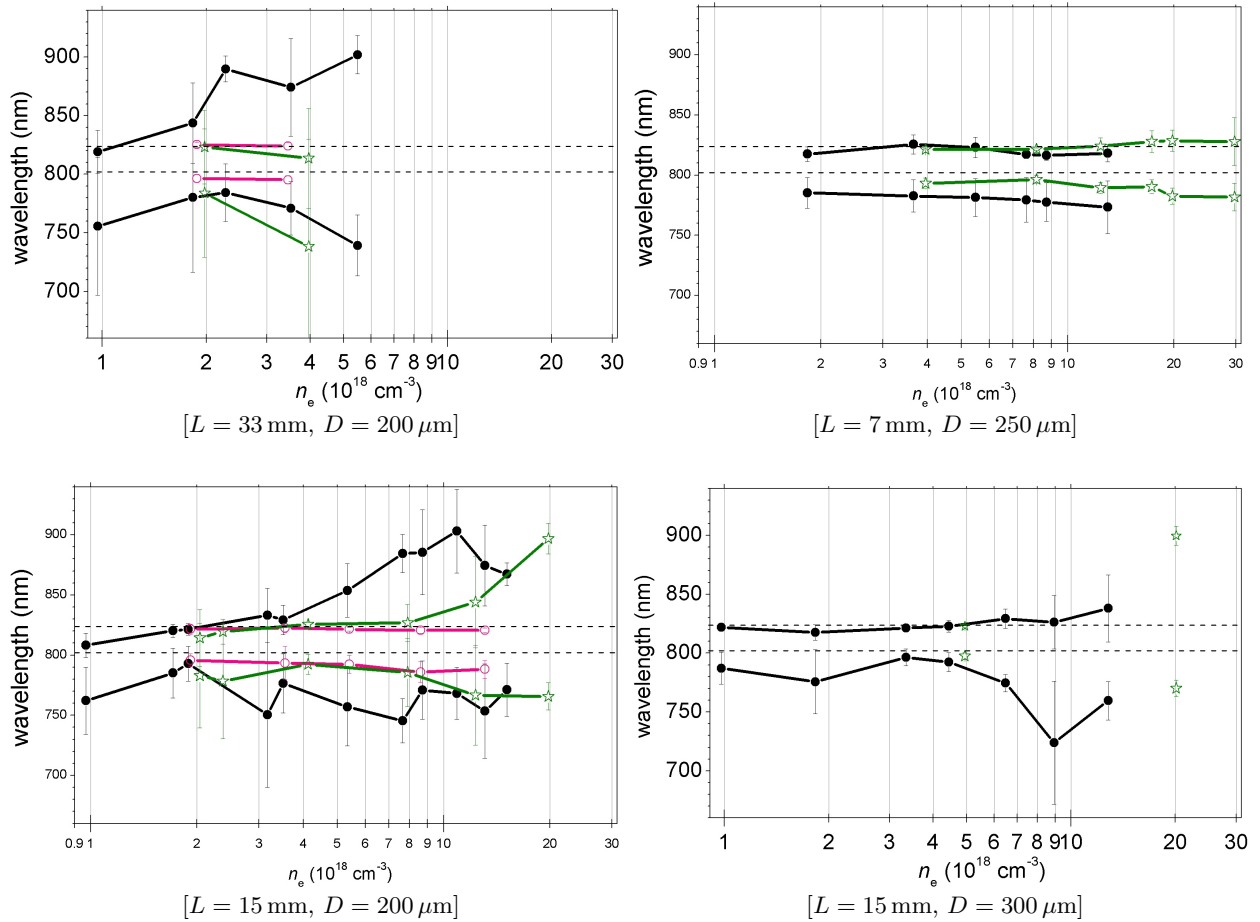


Figure 4.5: Spectral broadening versus plasma density.

Λ^R and Λ^B at peak transmission calculated using equation (3.13) for the long pulse (pink) the short pulse (black), or for the short pulse at negative t (green). For the first two of those plots the x -axis density is calculated from the recorded p using equation (2.13). The dotted lines are Λ^R and Λ^B for the spectrum of the input pulse. For the latter plot n_e was the mean density of the gas fill. Each graph shows data from a different combination of capillary parameters (L and D).

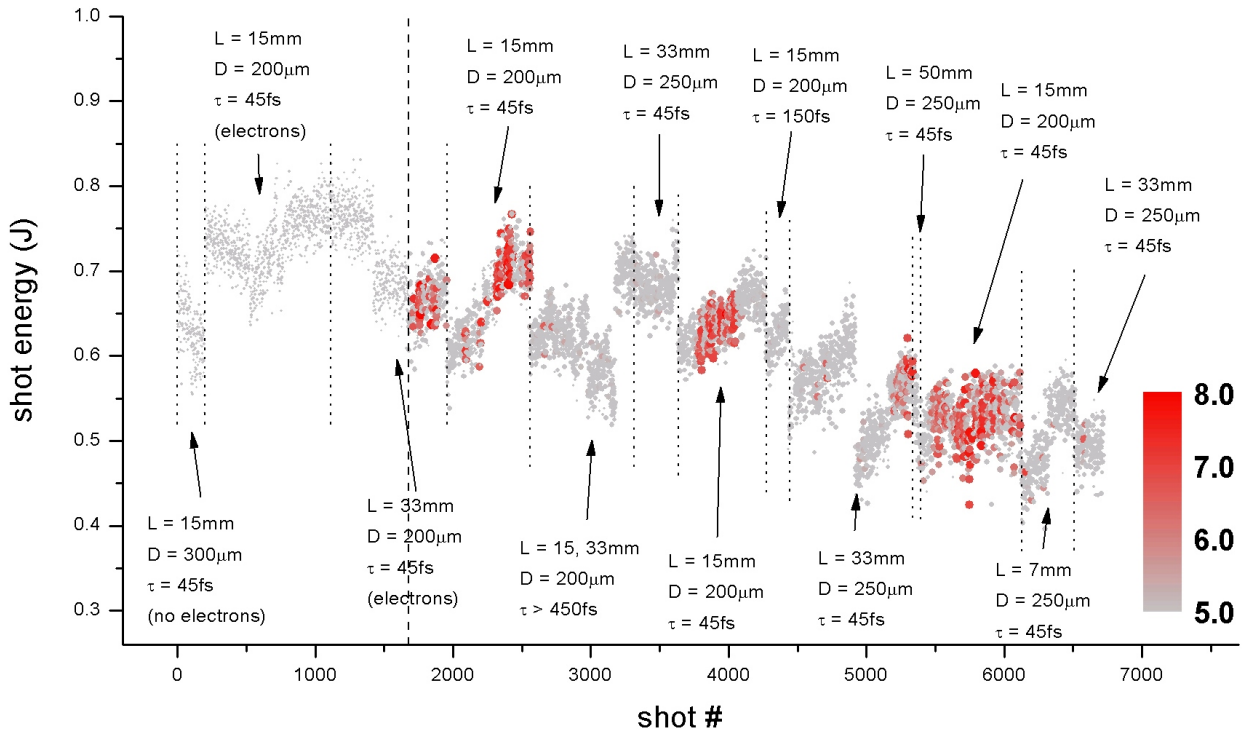


Figure 4.6: Overview of electron beam generation for the Astra experiments. Laser energy versus shot number for all full laser shots from July 25th 2006 to 18th August 2006. The size and colour of the dots represent the log of the electron spectrometer signal for those shots. Shots plotted to the left of the dashed line were taken before the installation of the electron spectrometer. Electrons were measured on the pointing monitor, which was not calibrated, so an electron signal is not plotted. It is indicated if electrons were detected on the plot labels. The general downwards trend of the laser input energy is due to a gradual degradation of the Astra compressor gratings, reducing the on target energy.

red for $n_e \gtrsim 15 \times 10^{18} \text{ cm}^{-3}$. Again, that implies that the excitation of wakefields was only possible at lower densities if there was a plasma channel present.

Note that for the long pulse Λ^R and Λ^B were very similar to their input values for all densities. It is concluded that the long pulse was not sufficiently intense to drive a wakefield at any density in the range studied ($1 \times 10^{18} \text{ cm}^{-3} - 30 \times 10^{18} \text{ cm}^{-3}$).

Observant readers may also have noticed that for $[L = 7 \text{ mm}, D = 250 \mu\text{m}]$ the spectrum did not broaden at all. That waveguide had only a single gas slot feeding into the middle of the capillary, meaning that the pressure between the gas slot and the capillary end was not known. As a result, the actual density in that capillary was probably much lower than predicted from equation (2.13).

4.2 Electron beam measurements

The laser energy of every full power shot of the Astra experiments and the resultant electron spectrometer signal are plotted in 4.6. This graph gives a crude indication of the conditions for which electron acceleration occurred, although it should be emphasised that many other critical

parameters were also being changed. There was a scatter and overall downwards trend in the laser energy. The electron spectrometer signal did not appear to be affected by that, suggesting that other parameters were more critical than laser energy for accelerating electrons. Electrons were only generated for the short pulse in the $200\ \mu\text{m}$ wide capillary. That confirms what the evidence of section 4.1 also suggested: that the capillary diameter was critical to the LWFA process, and that although for $D > 200\ \mu\text{m}$ the laser could be guided, the interaction between the laser and the plasma was too weak for LWFA to occur. The reasons that a narrow plasma channel enabled LWFA at lower electron densities are explored in chapter 5.

4.2.1 Electron energy

Figure 4.7 shows examples of images of the lanex screen of monoenergetic electron spectra, and figure 4.8 plots the energy of all such electron spectra against electron density. Monoenergetic electron beams could be generated at electron densities ranging from $2 \times 10^{18}\ \text{cm}^{-3}$ to $25 \times 10^{18}\ \text{cm}^{-3}$.

Comparison to gas jet results

At negative t , for which no plasma channel was present, the maximum energy that could be achieved was $\sim 110\ \text{MeV}$ at an electron density of $20 \times 10^{18}\ \text{cm}^{-3}$. In work by another group, the maximum energy of electrons accelerated in a gas jet (of length 2 mm) using the same laser was $\sim 80\ \text{MeV}$ at an electron density of $20 \times 10^{18}\ \text{cm}^{-3}$ [4], suggesting that despite several technical differences in setup³ the shots recorded at negative delays are equivalent to those fired into a gas jet.

The highest energy of electrons accelerated in the plasma channel was approximately double the energy of electrons accelerated at negative delays or in gas jets. For $L = 15\ \text{mm}$ the maximum energy of 200 MeV is reached for $p = 400\ \text{mbar}$ ($n_e = 8.7 \times 10^{18}\ \text{cm}^{-3}$). The energy drops away either side of that pressure. The range of densities at which electrons were measured for $L = 33\ \text{mm}$ was not sufficiently broad for any trends to be identified. The maximum energy achieved with that capillary was also 200 MeV, but at a plasma density of only $2 \times 10^{18}\ \text{cm}^{-3}$.

Comparison to scaling laws

There were considerable shot-to-shot fluctuations in the measured electron energy. The likely source of that instability was the non-linear evolution of the laser pulse as it became sufficiently intense to induce electron trapping. The length over which that evolution occurs would have varied from shot-to-shot, so that the subsequent length over which electrons were accelerated (before either depletion of the laser pulse or the end of the capillary) also varied. The highest energy electron beams within the scatter would be those that were accelerated for exactly one dephasing length.

The scaling laws for energy gain in the linear and bubble regimes, derived in section 1.5, are plotted in figure 4.8. Because these expressions assume that electrons would be accelerated for exactly one dephasing length, their predictions should be in agreement with the highest energy electron beams. With that in mind, there is an excellent agreement between the data and equation (1.41) for $n_e > 8 \times 10^{18}\ \text{cm}^{-3}$: the highest energy beams at a particular density have an energy

³The gas jet experiments in question used helium rather than hydrogen gas and a spot size $W \sim 13\ \mu\text{m}$ rather than $34\ \mu\text{m}$.

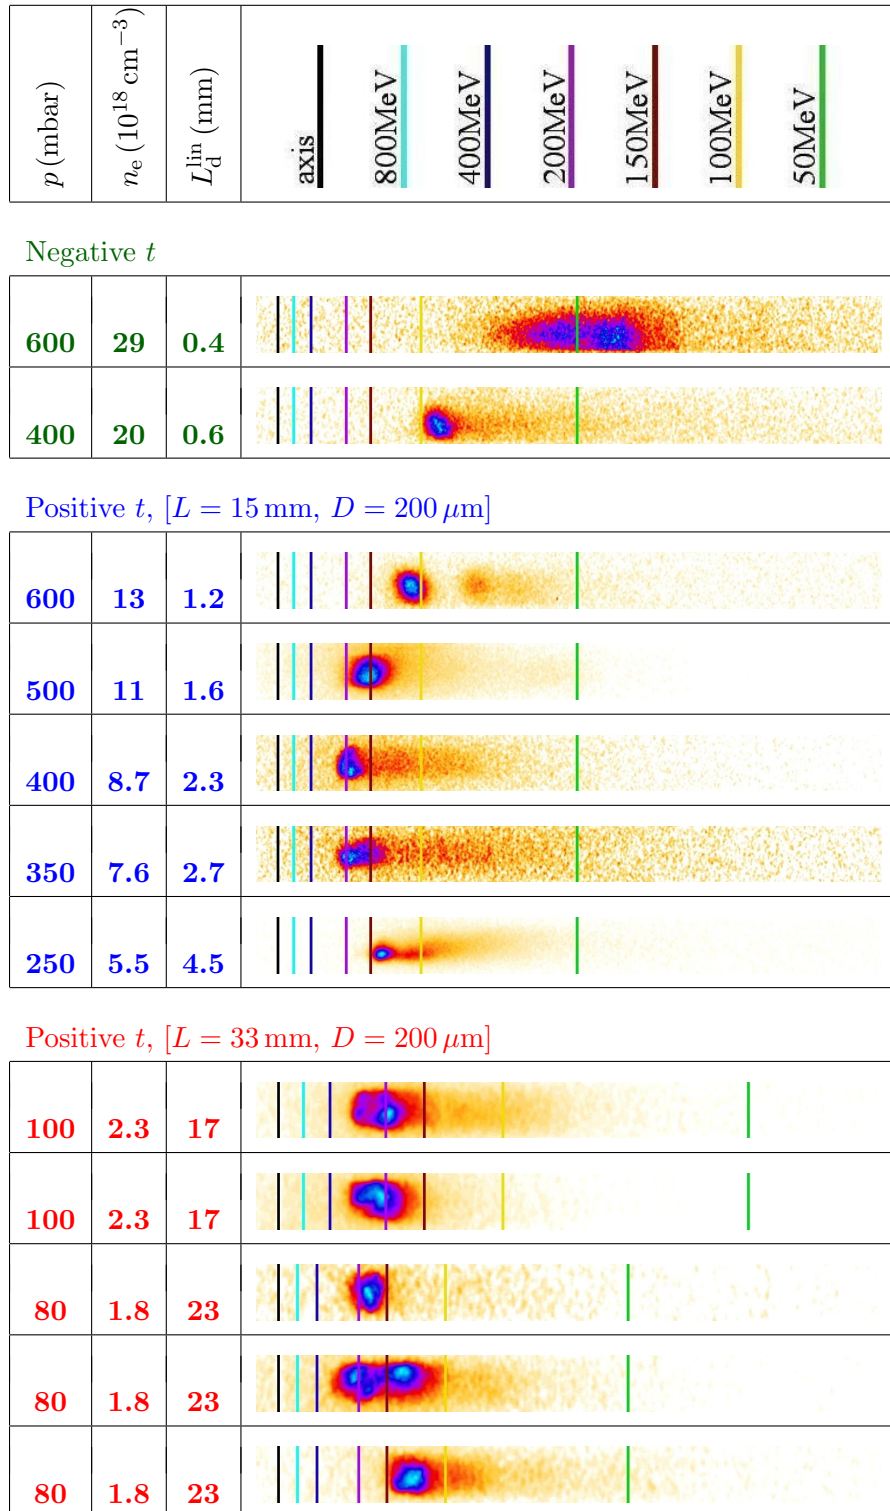


Figure 4.7: A selection of electron spectrometer lanex screen images. Each plot is an image from the lanex screen located at the back of the spectrometer. The coloured lines represent the relationship between displacement and electron energy, which varied as the spectrometer current varied. The data shown corresponds to the highest measured electron energies for the conditions stated.




Figure 4.8: Electron energy of monoenergetic beams versus density

The energy of the peak of monoenergetic beams is plotted against plasma density for negative delays (green dots), $[L = 15 \text{ mm}, D = 200 \mu\text{m}]$ waveguides (blue dots) and $[L = 33 \text{ mm}, D = 200 \mu\text{m}]$ waveguides (red dots). The predicted upper limit of electron beam energy are plotted for: the bubble regime, as predicted by equation (1.41) (black line); the linear regime, as predicted by (1.38) (grey line).

L (mm)	15	33
X -Pointing RMS (mrad)	6	6
Y -Pointing RMS (mrad)	3.5	6
X -Divergence FWHM (mrad)	5	3.5
Y -Divergence FWHM (mrad)	5	4.5
Minimum FWHM (mrad)	1.5	2.5

Table 4.1: Electron beam pointing and divergence.

Based on lanex screen pointing data of figure 4.9. X represents the dimension perpendicular to the laser polarisation, Y is parallel. The Minimum Beam FWHM is the divergence of the most collimated beams recorded.

very close to that predicted by the scaling laws. This suggests that at those densities the bubble regime description of LWFA applied. Decreasing to lower densities, there was a transition to linear regime scaling, and at $n_e = 2 \times 10^{18} \text{ cm}^{-3}$ the data is in agreement with equation (1.38), the linear prediction of maximum energy. According to theory [11, 32] the bubble regime scaling should have been applicable for $P/P_c \gtrsim 1$ which, for a 12 TW laser, corresponds to $n_e \gtrsim 2 \times 10^{18} \text{ cm}^{-3}$. The crossover between regimes occurred at a higher density than expected, and that is examined in chapter 5.

4.2.2 Electron beam profiles

Figure 4.9 shows selected profiles of electron beams emerging from the capillary. The pointing stability and divergence of the beams seemed to vary with delay t , but even at its most stable the electron beam has considerable shot-to-shot variations.

Beam divergence and pointing

Estimates of the pointing stability and divergence based on the data of the lanex screen pointing monitor are given in table 4.1. These estimates are based on a small sample of data, taken before acceleration conditions had been optimised and as such the conclusions which may be drawn are limited. The pointing stability was an improvement on that measured without a waveguide for the same laser [67], but not dramatically so. That work also suggested that electron beam pointing stability is a strong function of the laser parameters. Work with other lasers [17] has produced very low divergence electron beams with a high level of pointing stability without the presence of a plasma channel. That leads to the conclusion that the presence of a plasma channel does not significantly impact the beam divergence and pointing, and that these parameters are more sensitive to the quality and stability of the input laser pulse.

Beam quality in the linear and bubble regimes

A qualitative observation that can be made from the data in figure 4.9 is that the beams for [$L = 33 \text{ mm}$, $p = 100 \text{ mbar}$] were often fragmented with more than a single peak in the transverse

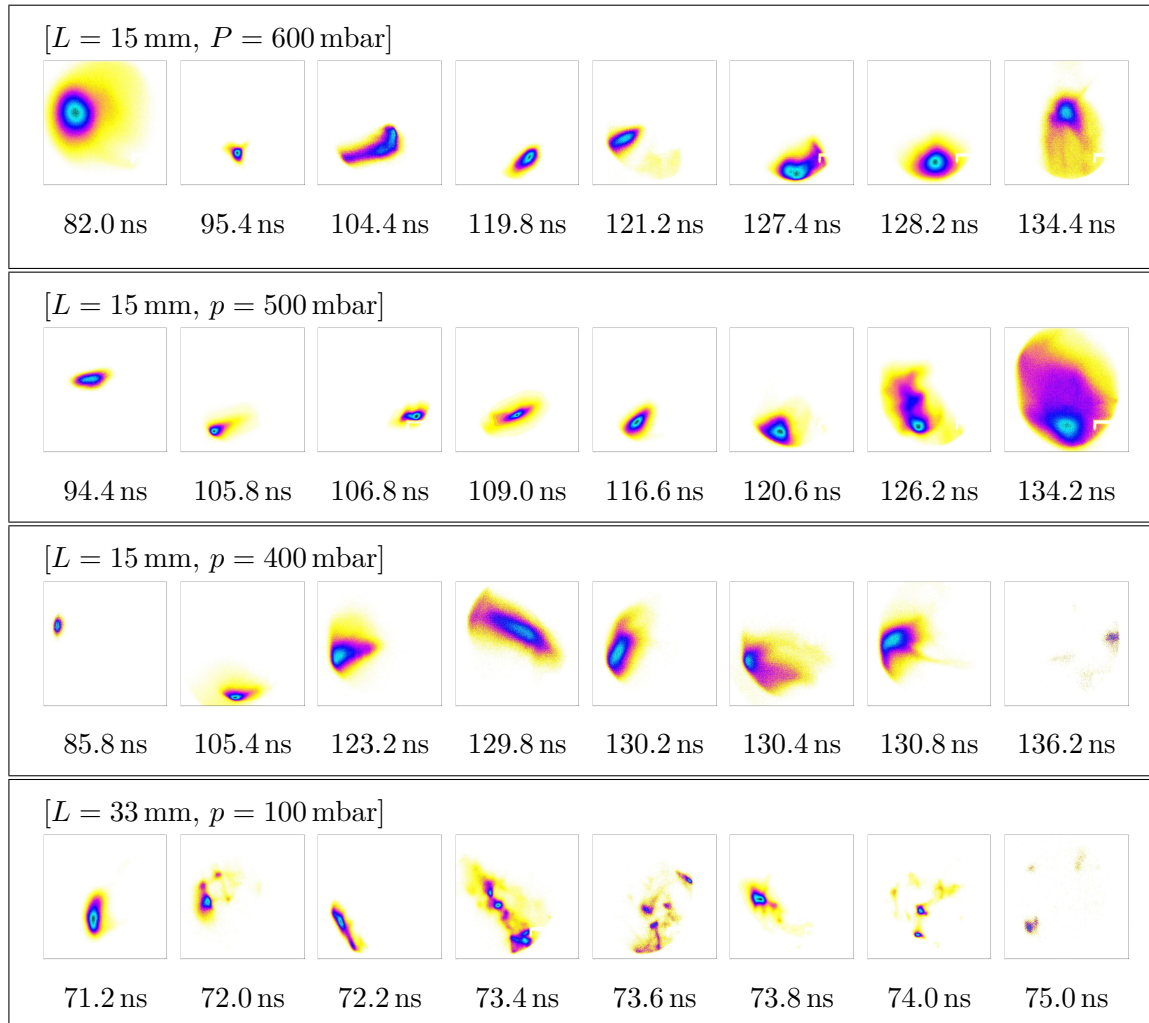


Figure 4.9: Electron beams recorded by the electron pointing monitor.

Each figure represents a 40 mrad square and is colour scaled to the peak electron signal for that shot. It should be noted that this pointing data is not necessarily representative of the optimum pointing stability possible using the experimental methods employed: the above measurements are the only direct measurements of the electron beam profile and were made at a point in the experimental run before the generation of electrons had been optimised.

profile. The evidence of section 4.2.1 suggests that the [$L = 33$ mm, $p = 100$ mbar] beams would have been generated in the linear regime, whereas the other beam profiles presented in figure 4.9, for which the density was higher, would have been generated in the bubble regime. It has been suggested [11, 31, 32, 35–37] that the beams generated in the bubble regime should have higher quality than those generated in the linear regime due to the ability of the plasma cavity to maintain its structure after beam-loading. Although the evidence of figure 4.9 is by no means conclusive, it supports that argument.

4.3 Correlating measurements of transmitted radiation to LWFA

4.3.1 Guiding and electron acceleration

Electrons were only generated in $D = 200$ μm capillaries⁴, which indicates that in these experiments the properties of the plasma channel were critical to the LWFA process. Comparing the dependence on delay of guiding and electron acceleration allows further insight as to how the presence of a plasma channel affected the LWFA process.

High pressure behaviour

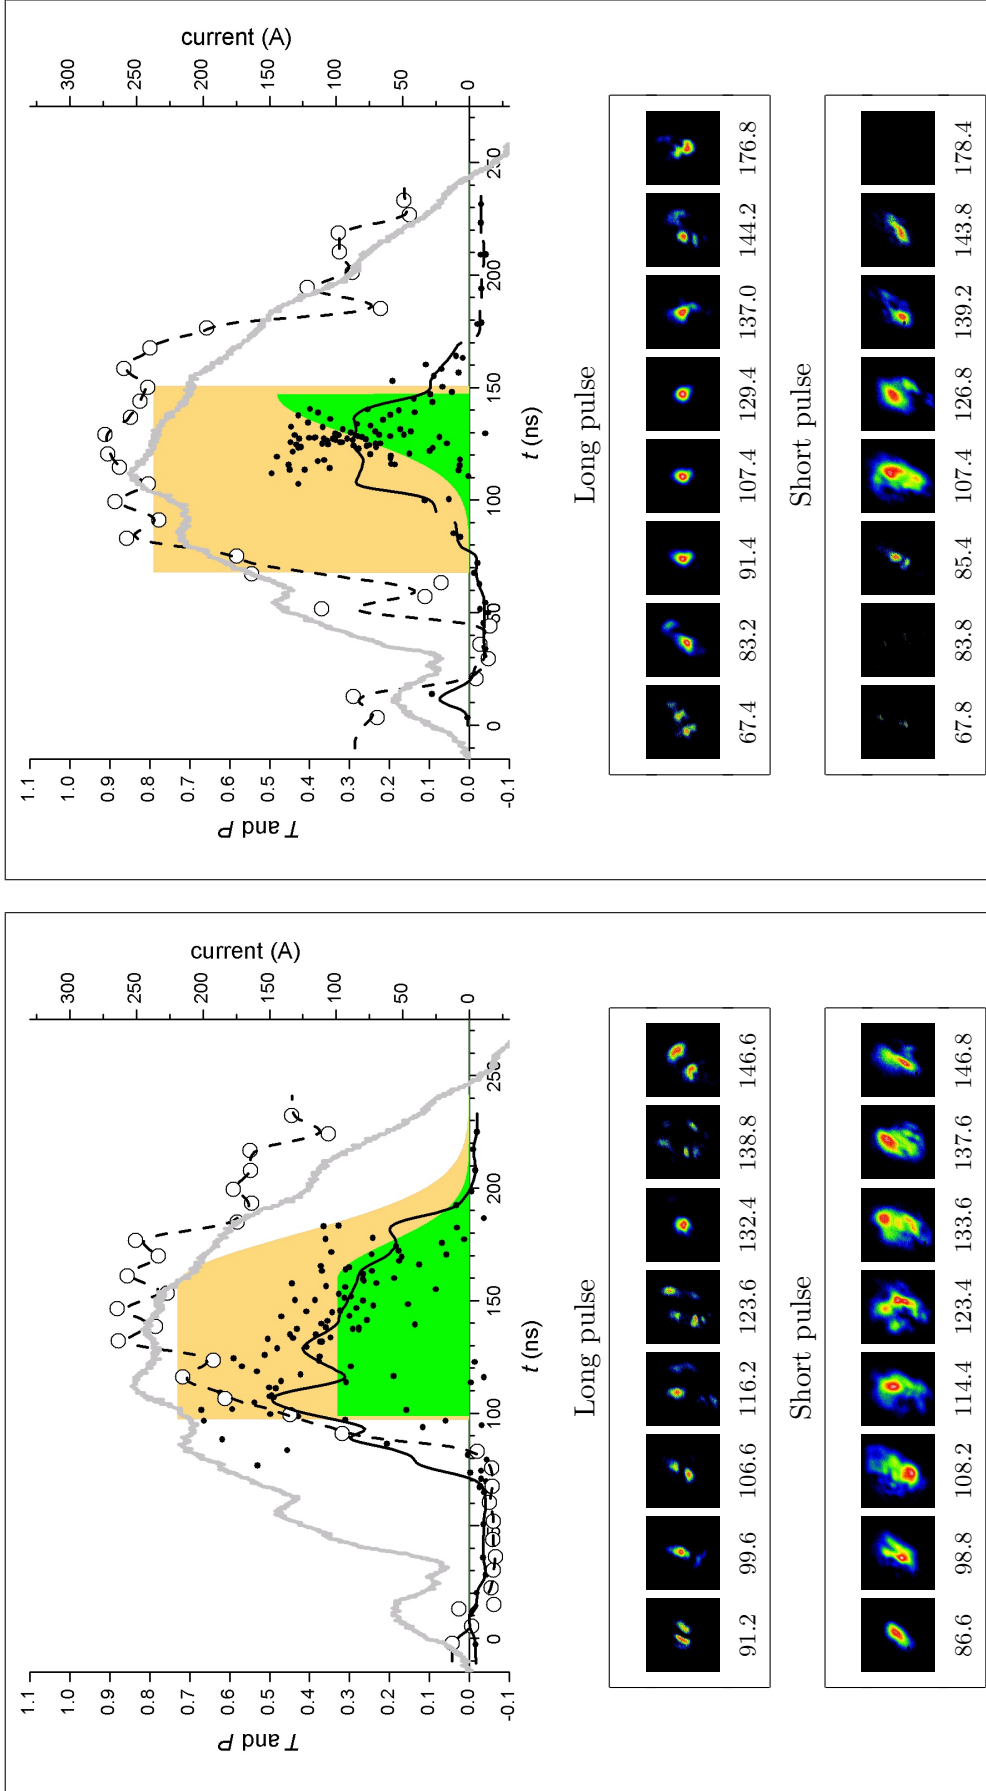
In figure 4.10 long and short pulse transmission data is plotted against delay for [$L = 15$ mm, $D = 200$ μm], alongside $P^e(t)$, the probability of accelerating electrons, and $P^b(t)$, the probability of accelerating monoenergetic electron beams. Fluence exit modes are included for a representative selection of delays. The two plots are for $p = 600$ mbar and $p = 400$ mbar, which were relatively high pressures. As with the data of section 4.1.1, the transmission of the long pulse was high, and the long pulse exit modes were symmetric and narrow, particularly for $p = 400$ mbar, demonstrating that a guiding plasma channel was present. The short pulse transmission was relatively low (~ 0.3) due, presumably, to pulse depletion, and the short pulse exit modes were distorted⁵. For $p = 400$ mbar the short pulse transmission was more sensitive to delay than the long pulse transmission. The range of delays for which $P^e(t)$ and $P^b(t)$ were non-zero matched the delays for which the waveguide transmission was high. This correlation between $P(t)$ and $T(t)$ is evidence that at these densities the presence of a plasma channel was critical to the acceleration of electrons. A physical explanation for this is given in chapter 5.

Low pressure behaviour

Figure 4.11 is equivalent to figure 4.10, but for [$L = 15$ mm, $D = 200$ μm , $p = 250$ mbar] and [$L = 33$ mm, $D = 200$ μm , $p = 80$ mbar]. These pressures were relatively low. The transmission data is similar to that in previous plots: high transmission of the long pulse with symmetric compact exit modes, lower transmission for the short pulse with distorted exit modes and for both pulse lengths an unexpected drop in the transmission during the current pulse. However $P^e(t)$ and $P^b(t)$

⁴Except for at negative delays, for which the capillary diameter had no effect.

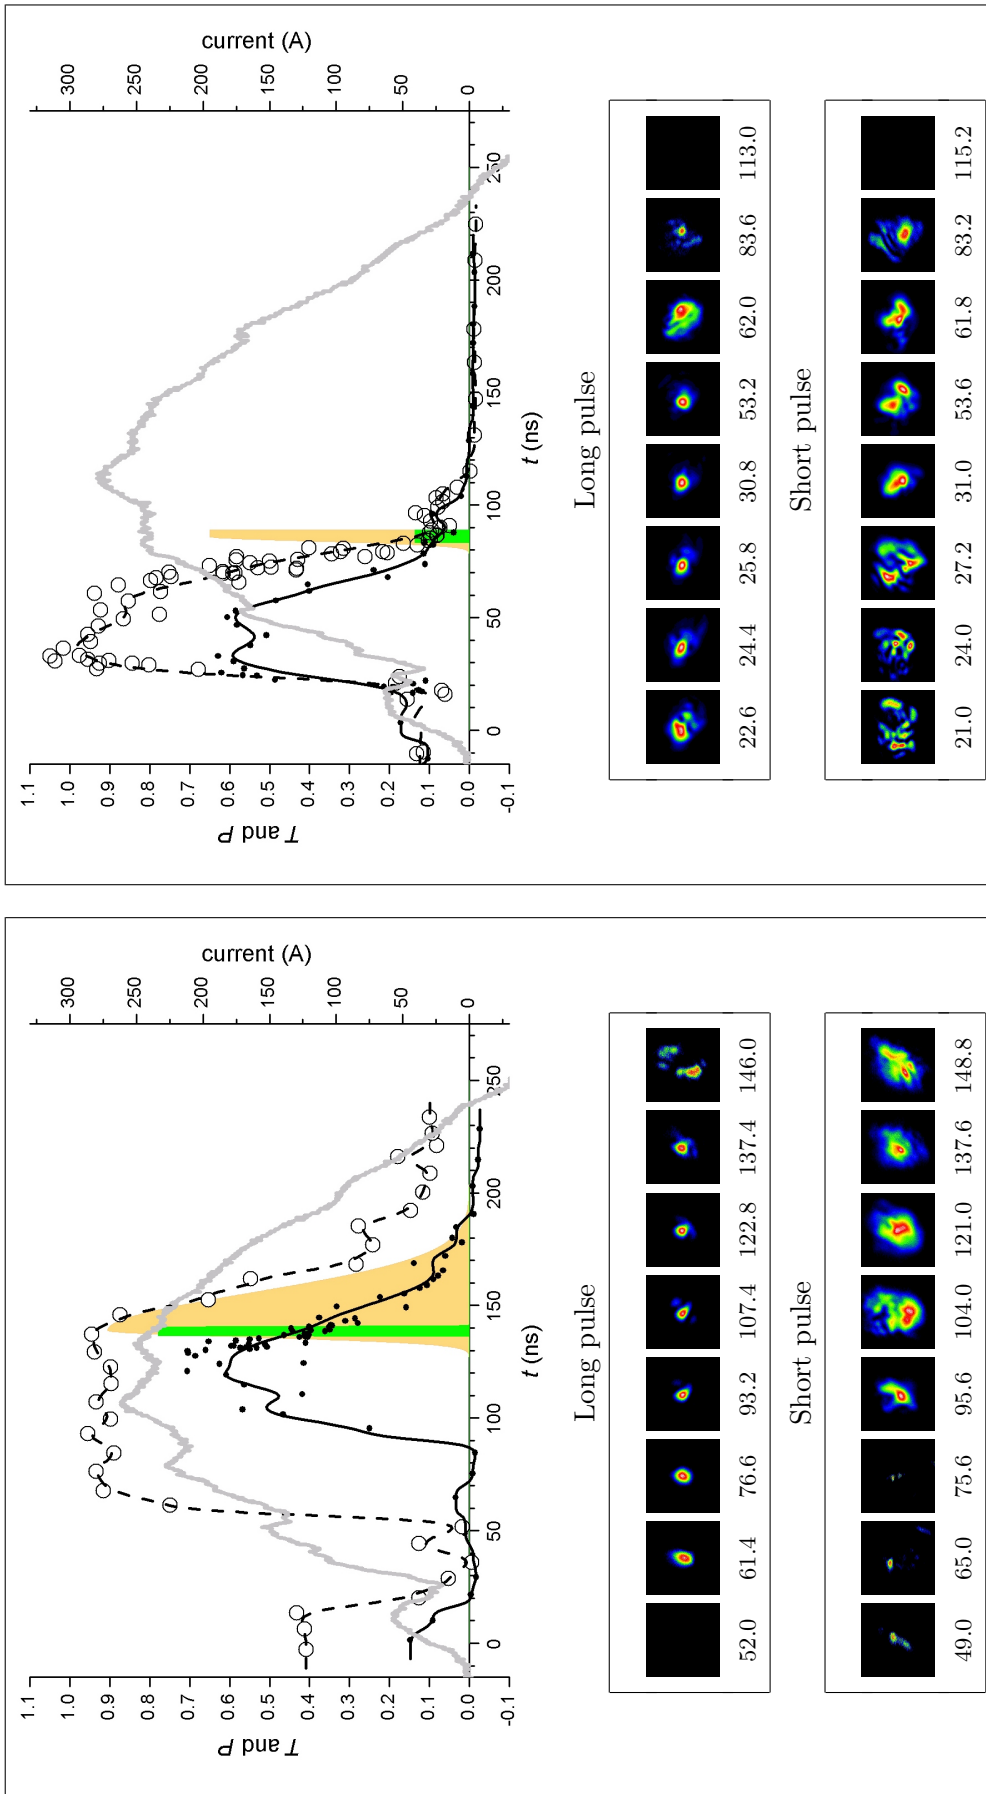
⁵As with other short pulse data, the distortion to the exit mode may have been the result of chromatic aberration in the lens system imaging the capillary exit. An analysis of the exit modes for 4.10(a) is presented in appendix B that suggests the real output mode had Gaussian width ~ 31 μm , which is slightly smaller than the input spot.



(b) [$L = 15$ mm, $D = 200$ μ m, $p = 400$ mbar]

(a) [$L = 15$ mm, $D = 200$ μ m, $p = 600$ mbar]

Figure 4.10: Transmission T and beam probability $P(t)$ versus delay t for $n_e \approx 6.5 \times 10^{18}$ cm^{-3} . Long/short pulse transmission data $\{t_i, T_i\}$ (hollow dots/solid line) and spot images for selected delays are plotted. Also plotted are $P^e(t)$ (yellow line and fill) and $P^b(t)$ (green line and fill). Respectively these are the probability of generating electrons, and the probability of generating a monoenergetic electric beam, deduced using the method described in section 3.3.3. The discharge current is plotted against the right hand axis (grey line). For figure (a) $n_e = 13 \times 10^{18}$ cm^{-3} and for figure (b) $n_e = 8.7 \times 10^{18}$ cm^{-3} .



(a) [$L = 15$ mm, $D = 200$ μ m, $p = 250$ mbar]

(b) [$L = 33$ mm, $D = 200$ μ m, $p = 80$ mbar]

Figure 4.11: Transmission T and beam probability $P(t)$ versus delay t for $n_e \lesssim 6.5 \times 10^{18} \text{ cm}^{-3}$. Long/short pulse transmission data $\{t_i, T_i\}$ (hollow dots/solid line) and spot images for selected delays are plotted. Also plotted are $P^e(t)$ (yellow line and fill) and $P^b(t)$ (green line and fill). Respectively these are the probability of generating electrons, and the probability of generating a monoenergetic electric beam, deduced using the method described in section 3.3.3. The discharge current is plotted against the right hand axis (grey line). For figure (a) $n_e = 5.5 \times 10^{18} \text{ cm}^{-3}$ and for figure (b) $n_e = 1.8 \times 10^{18} \text{ cm}^{-3}$.




Figure 4.12: Sensitivity of transmission and electron generation to delay against density. The sensitivity to delay of long pulse transmission (pink), short pulse transmission (black), P^b (green) and P^e (orange) are plotted against axial plasma density. For the transmissions S is defined as $(2dt)^{-1}$, where dt is calculated from equation (3.14). For the electron probabilities S is defined as the inverse of the standard deviation of the function $P(t)$. The values of n_e were calculated using equation (2.13).

were distinctly different at these low pressures. Electron beams were not generated for delays between the rise and fall of transmission, but were observed within 4 ns range of delays at $t \approx t_f$. The sensitivity of electron acceleration to delay suggests that there was a transient property of the waveguide plasma at $t \approx t_f$ which enabled electron acceleration. For $[L = 15 \text{ mm}, D = 200 \mu\text{m}, p = 250 \text{ mbar}]$ the most stable monoenergetic beams of the experiment were generated, with a peak probability $P_0 = 0.78$. In chapter 6 it is found that the unexpected fall in transmission was accompanied by a drop in the ionization level of the waveguide plasma, and reasons why that would enhance the LWFA process are discussed.

Sensitivity against density

The trends observed in figures 4.10 and 4.11 are summarised in figure 4.12, where the sensitivity of both electron acceleration and transmission to delay is plotted. For higher densities the sensitivity of electron generation to t was approximately equal to the sensitivity of transmission to timing, suggesting that the generation of electrons was dependent only on the formation of a plasma channel.

At lower densities the generation of electrons was much more sensitive to discharge timing than the transmission, suggesting that in that regime the formation of a plasma channel in the waveguide was not the only factor critical to the generation of electron beams. The cross over between these regimes occurred between $6 \times 10^{18} \text{ cm}^3$ and $7 \times 10^{18} \text{ cm}^3$.

4.3.2 Laser red-shifts and LWFA

As already discussed, the excitation of a wakefield by a laser pulse is accompanied by the broadening of part of the laser spectrum to longer wavelengths. The relationship between red-shifting and electron acceleration is well illustrated in figure 4.13. The electron spectrometer signal is plotted for every shot within the data-sets representing negative delays (no plasma channel), [$L = 15 \text{ mm}$, $D = 200 \mu\text{m}$] and [$L = 33 \text{ mm}$, $D = 200 \mu\text{m}$]. Many parameters, including t are varied within the datasets plotted. The red-limit of the spectrum, $\Lambda^{\text{R}}(n_e)$ calculated using the method described in section 3.3.2, is also plotted against the right hand axis. It is clear from the plot that the generation of electrons and red-shifting of the laser spectrum are correlated. For each data-set the minimum density at which the electron signal is non-zero is approximately the same density at which the value of transmitted $\Lambda^{\text{R}}(n_e)$ deviates from its original value towards longer wavelengths.

The data in figure 4.13 emphasises the role of the waveguide for accelerating electrons at low plasma densities. With no waveguide, that is at negative delays, the lowest density at which electrons were accelerated was $n_e \approx 20 \times 10^{18} \text{ cm}^{-3}$. For the 15 mm long capillary the minimum was $n_e \approx 5.5 \times 10^{18} \text{ cm}^{-3}$ and for 33 mm the minimum was $n_e \approx 1.8 \times 10^{18} \text{ cm}^{-3}$. The minimum density at which electrons could be accelerated depended on both the presence of a plasma channel, and the length of that plasma channel. The impact of the plasma channel on the LWFA process is explored in detail in chapter 5, where it is found that in addition to extending the laser-plasma interaction length the plasma channel allowed more efficient use to be made of the available laser energy.

Readers may have noticed that for the 33 mm capillary the electron signal returns to zero as the density is increased above the threshold, but the laser spectrum was still red-shifted. It is not obvious why, if electrons could be generated in the shorter capillary at a given density, they could not also be generated in the longer capillary. In fact under these conditions the peak transmission of the long pulse also dropped and it is believed that the waveguide was not operating as intended. The reason for this is attributed to the fact that t_f is earlier for long narrow capillaries, and t_r is later for high densities, due to the additional time necessary to ionize the gas-fill. The trends suggest that for high pressures in the long capillary $t_f < t_r$ — in other words, there were no values of t for which the device was operating as a waveguide for these conditions. Note that this effect is explored in detail in chapter 6 and it is found that the premature drop in transmission at t_r is not a fundamental limitation of the hydrogen-filled capillary discharge waveguide scheme.

4.4 Capillary durability

An advantage of the hydrogen-filled capillary discharge waveguide over other schemes is the long lifetime of the sapphire capillaries. Previous experimental work [45] found that the discharge

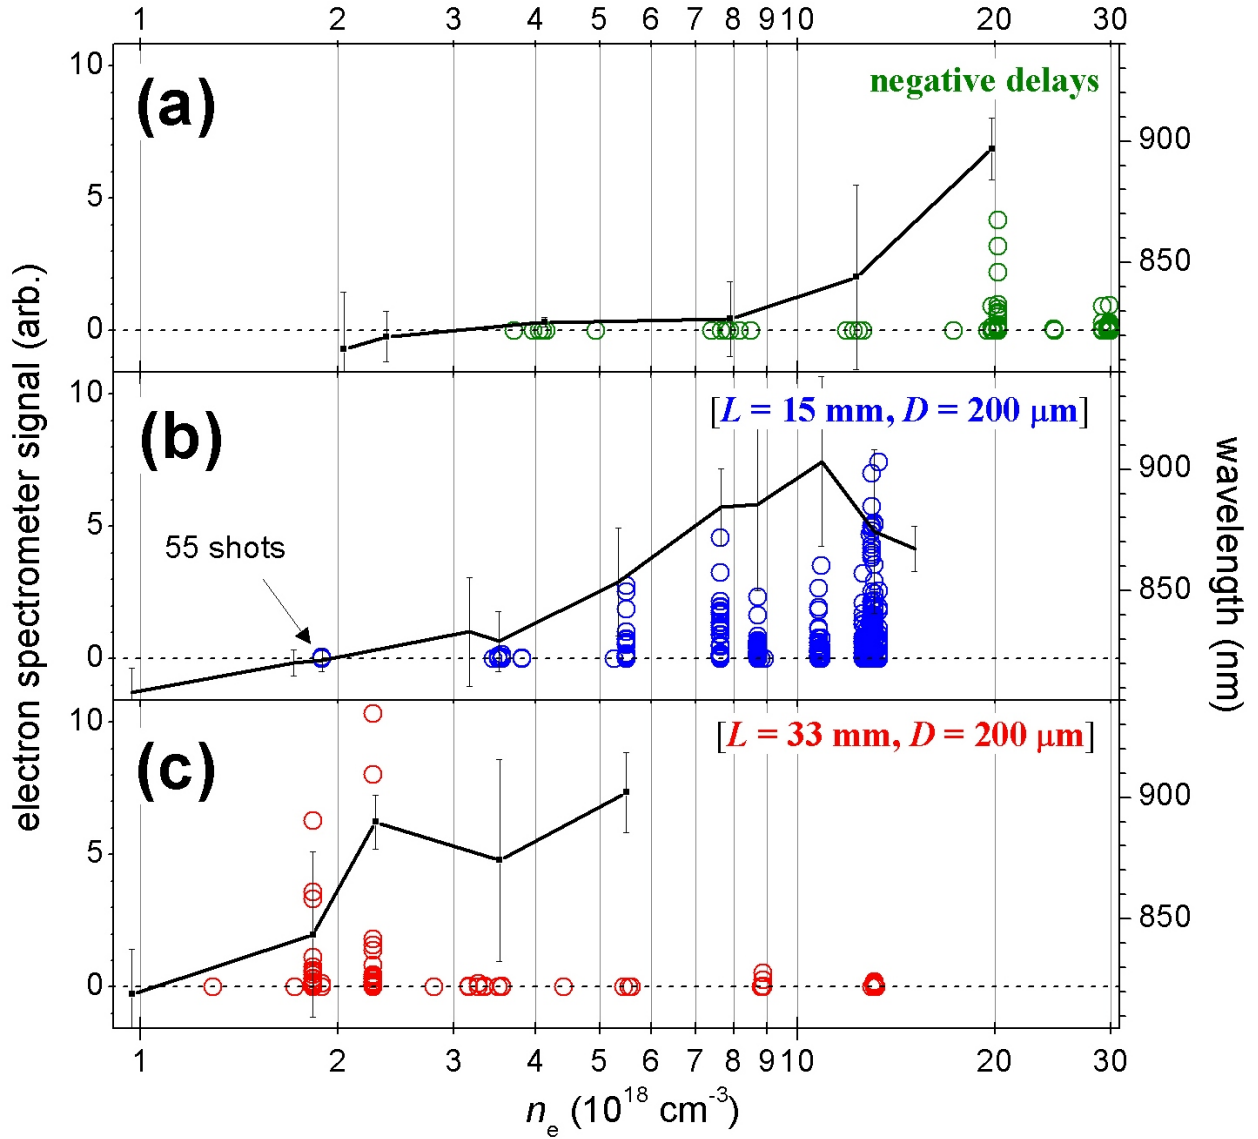


Figure 4.13: Electron signal and red-limit Λ^R versus density n_e .

The electron spectrometer signal is plotted against the left hand axis (rings). The mean red-limit $\Lambda^R(n_e)$, calculated using the method presented in section 3.3.2, is plotted against the right hand axis (line). Three different scenarios are represented: figure (a) is for negative delays, all capillary parameters; figure (b) is for positive delays, $[L = 15 \text{ mm}, D = 200 \mu\text{m}]$; figure (c) is for positive delays, $[L = 33 \text{ mm}, D = 200 \mu\text{m}]$. For each plot the electron spectrometer signal is plotted for every shot within the data-set under consideration. The dotted line represents $\Lambda^R = 822 \text{ nm}$ which is the value of $\Lambda^R = 822 \text{ nm}$ of the input spectrum. At densities for which no electrons were detected many shots with zero spectrometer signal can appear to be a single shot. As an example 55 shots that appear to be a single shot are indicated in figure (b). At all densities the number of shots was far greater than 1.

current widens a square capillary of side $200\ \mu\text{m}$ by approximately $20\ \mu\text{m}$ after 1.3 million shots, corresponding to approximately $0.2\ \text{\AA}$ per shot. Such damage is negligible. However the capillary can also be damaged by the laser pulse. In order to record the laser damage, the capillaries from the Astra experiments were scanned after use. The damage to these capillaries does not necessarily give a fair representation of the durability of the hydrogen-filled capillary discharge waveguide. It was estimated, using equation (3.15), that as different sets of conditions were explored 85 % of shots were fired at delays for which the plasma channel was not present. For those shots the unguided laser pulse would have diffracted into the capillary walls and ablated the sapphire. If the waveguide was only used at delays for which the laser pulse was channeled, as it would be in any LWFA device, the damage would no doubt have been less. Additionally variations in laser pointing would have increased damage to the front of the capillary. With that taken into consideration, the condition of two capillaries is documented for future reference.

The sapphire plates, into which the capillaries were machined, were scanned at a resolution of approximately $10\ \mu\text{m}$ per pixel. There were two types of scan: those that measured the transmission through the sapphire from a light source above the plates, and those that measured reflection from the sapphire using a light source below the plates. These scans were complementary, with the transmitted scan outlining the shape of the capillary, and the reflected scan revealing the contents of the capillary and gas slots. For each capillary both the plate with gas slots (side A) and that without gas slots (side B) were scanned. Unfortunately no scans were made of the capillaries before use to serve as a comparison.

Figure 4.14 shows the transmitted scans of both capillaries. The [$L = 15\ \text{mm}$, $D = 200\ \mu\text{m}$] capillary underwent 750 full power shots and the [$L = 33\ \text{mm}$, $D = 200\ \mu\text{m}$] underwent 1250 full power shots. At $\times 1$ magnification there is no visible damage to the capillaries, which are both in good condition. The only noteworthy observation is that for both capillaries the position of the gas slots are clearly visible on side B where the plates have been marked. It is likely that the marks were caused by plasma in the gas slots: interferometry experiments have shown that plasma is ejected from the capillary into the gas slots during the discharge [45]. The marks are superficial and would not have affected waveguide performance.

Capillary entrance damage

The full length of the [$L = 15\ \text{mm}$, $D = 200\ \mu\text{m}$] capillary at $\times 5$ magnification is shown in figure 4.15. It is clear from the transmitted scans that the entrance of the capillary was widened to a diameter of order $1\ \text{mm}$, with no similar effect at the capillary exit. That is most likely due to energy in the wings of the laser focus, which was not perfectly Gaussian, ablating the sapphire at the capillary entrance. Fluctuations and drifts in laser pointing may also have contributed to the damage. It is noteworthy that the majority of damage at the capillary entrance is to one side, suggesting that the laser was misaligned when the majority of that damage was inflicted. This ‘cone’ does not penetrate deeper than $0.5\ \text{mm}$. The bulk of the capillary appears to be undamaged. The capillary was working effectively as a waveguide when it was taken out of operation.

Figure 4.16 shows the entrance and exit of the [$L = 33\ \text{mm}$, $D = 200\ \mu\text{m}$] capillary at $\times 5$ magnification. As with the other capillary, there is a cone shape of damage at the capillary entrance.




Figure 4.14: $\times 1$ magnification scans for two used capillaries. The transmitted scans are shown. The capillary entrances are located at the left of the images. The $[L = 15 \text{ mm}, D = 200 \mu\text{m}]$ capillary is shown at a higher magnification in figure 4.15. The red boxes on the $[L = 33 \text{ mm}, D = 200 \mu\text{m}]$ capillary mark areas that are shown at higher magnification in figures 4.16 and 4.17.




Figure 4.15: $\times 5$ magnification scans of the used [$L = 15$ mm, $D = 200$ μm] capillary. The capillary entrance is located at the bottom of the images.




Figure 4.16: $\times 5$ magnification scans of the used [$L = 33$ mm, $D = 200 \mu\text{m}$] capillary. The areas scanned are marked by the red boxes in figure 4.16.




Figure 4.17: $\times 10$ magnification scans near the [$L = 33$ mm, $D = 200$ μm] capillary exit. The reflected scans are shown. The position of the scan relative to the capillary is marked by a red box on figure 4.14.

In this case the cone is more pronounced than that for the 15 mm capillary, having a depth of approximately 1 mm, a reflection of the fact that this capillary was used for almost twice as many high power laser shots.

Markings

An unusual feature that can also be seen in figure 4.16 occurs before the second gas slot. There are dark marks that forms a helical pattern inside the capillary, which ‘snakes’ from side to side. Similar marks can be seen in figure 4.15 in and around the gas slots of the [$L = 15$ mm, $D = 200$ μm] capillary. The helical pattern and snaking was also observed leading up to the second gas slot of another [$L = 15$ mm, $D = 200$ μm] capillary, so the effect was not isolated to the capillary in which it is seen here. It is difficult to interpret these observations because the conditions for which the snaking pattern formed are not known. Guiding did not appear to be affected, but the origin of the dark deposit and helical patterns are interesting and warrant further investigation.

Damage from the discharge

Figure 4.17 shows a $\times 10$ magnification scan of a narrow groove in the sapphire plates of the [$L = 33$ mm, $D = 200$ μm] capillary that ran from the second gas slot to the waveguide exit. It is likely that groove was formed by repeated ablation of the sapphire by the current to form an alternative route to the electrode at the end of the waveguide. This was the only capillary for which such a groove was observed, and it is unclear exactly when it was formed. It was noted, however, that although electrons were accelerated with that waveguide, at some point during the capillaries lifetime the LWFA process ceased to occur.

4.5 Conclusion

The results from the Astra experiment establish that the hydrogen-filled capillary discharge waveguide can channel laser pulses with peak input intensities of order $10^{18} \text{ W cm}^{-2}$, and significantly improved LWFA performance. The waveguide enabled electrons to reach around double the energy of electrons accelerated by the same laser in a gas jet. A comparison of the electron energy scaling laws to the data suggests that the principal reason that higher electron energies could be reached in the waveguide was that it allowed LWFA at lower densities. With a waveguide electrons could be accelerated at densities as low as $n_e = 2 \times 10^{18} \text{ cm}^{-3}$; without a waveguide (at negative delays) electrons were not accelerated below $n_e = 20 \times 10^{18} \text{ cm}^{-3}$. The plasma channel also improved the stability of the accelerated electrons, with mono-energetic beams being generated on 80 % of shots for certain conditions.

Although these results were encouraging, there were significant trends in the data that suggested the contribution of the waveguide was not fully understood, and went beyond extending the laser-plasma interaction length. That is demonstrated by the sensitivity of electron acceleration to capillary diameter. According to scaling laws the matched spot of a $D = 200 \mu\text{m}$ waveguide is only 10–15 % smaller than that of a $D = 250 \mu\text{m}$ waveguide. The laser pulses were guided for both of these diameters. For $D = 250 \mu\text{m}$ the short pulse transmission was high and no electrons were accelerated. For $D = 200 \mu\text{m}$ the short pulse was depleted, presumably by driving a wakefield, and accelerated electrons were observed. Given the small difference in matched spot size, such a clear distinction in behaviour was unexpected. Of course one may argue that the axial laser intensity would have been higher for the narrower capillary, but that difference would have been small compared to the fluctuations and drifts in laser energy plotted in figure 4.6 which did not appear to affect electron acceleration. This emphasises the importance of the waveguide, and that its role in electron acceleration is complex. An overview of the evidence presented in this chapter suggests that the waveguide enabled LWFA at lower densities through two distinct effects.

Firstly, the presence of a plasma channel enabled the excitation of wakes at lower densities. That is the conclusion drawn from figures 4.5 and 4.13. Red-shifts of the short pulse laser spectra, which result from the excitation of a wakefield, were not observed without a plasma channel below $n_e \approx 15 \times 10^{18} \text{ cm}^{-3}$. With a plasma channel red-shifts were observed at densities as low as $2 \times 10^{18} \text{ cm}^{-3}$. The reasons that the plasma channel enhanced wake excitation are explored in chapter 5.

The second effect occurred at the densities below $6.5 \times 10^{18} \text{ cm}^{-3}$. In figure 4.11 it was observed that electron acceleration was more sensitive to delay than guiding, implying that at these densities the presence of a plasma channel alone was not sufficient for electron acceleration. For a narrow range of delays coinciding with an unexpected drop in waveguide transmission, electrons were trapped and accelerated. In chapter 6 it is established that at this critical delay the ionization level on the waveguide axis decreased. Reasons for such a drop in ionization, and how it may have been connected to electron trapping are also explored in that chapter.

Chapter 5

The role of the plasma channel in LWFA

In chapter 4 it was observed that electrons could be accelerated at lower plasma densities in a plasma channel than in a ‘flat profile’, that is a plasma of even density. This was established by comparing the electron signal at positive t , for which a plasma channel was formed, to the electron signal at negative delays, for which there was no plasma channel formed. For the Astra laser a density threshold of electron acceleration without a plasma channel was deduced as $n_e \sim 20 \times 10^{18} \text{ cm}^{-3}$. The equivalent threshold with a plasma channel was $n_e \sim 2 \times 10^{18} \text{ cm}^{-3}$. Electrons were only accelerated below the flat-profile density threshold in the narrowest capillaries, emphasising that the plasma channel was critical to the process. Because electron energy gain increases as density decreases, acceleration of electrons to higher energy was possible in the plasma channel.

The experimental results of 4 also showed that for $n_e \lesssim 20 \times 10^{18} \text{ cm}^{-3}$ no part of the transmitted laser spectrum was shifted towards longer wavelengths in the absence of a plasma channel. Such red-shifting is conventionally interpreted as evidence that a wake of significant amplitude is being driven, suggesting that the key to understanding why electrons could be generated at lower density in the plasma channel lies in establishing why it enabled a wake to be driven at lower density. The aim of this chapter is to establish the role of a plasma channel such as that of the hydrogen-filled capillary discharge waveguide in allowing wakes to be driven at densities below $20 \times 10^{18} \text{ cm}^{-3}$. The effect of extending the laser-plasma interaction is considered first, and then the contribution of a plasma channel in the self-focussing regime is examined.

At the very lowest densities ($n_e \lesssim 6.5 \times 10^{18} \text{ cm}^{-3}$), the generation of electron beams was found to also depend very sensitively on the conditions of the plasma channel. That is a separate effect that relates to electron trapping rather than wake excitation, and is the subject of chapter 6.

5.1 Extending interaction length

The initial motivation for using a waveguide in LWFA experiments is to extend the laser-plasma interaction length beyond the rayleigh range, Z_R , of the laser pulse. The highest energy electrons that can be produced in a single stage at a particular density are those accelerated over exactly a dephasing length, L_d . Therefore, if its effect is considered as solely an extension of the interaction

length, the waveguide should only be beneficial at plasma densities for which $Z_R < L_d \lesssim L$, where L is the length of the waveguide. At higher densities for which $L_d < Z_R$, the presence of a plasma channel should, according to these arguments, not have an effect.

The dephasing length for linear wakefield theory, and 3D non-linear ‘bubble-regime’ theory [11], is plotted against density in figure 5.1. The experimentally observed threshold densities for laser wakefield acceleration (above which electrons were generated) were deduced from figure 4.13. These densities are marked in figure 5.1 for the 33 mm waveguide, the 15 mm waveguide and the flat plasma profile. Comparing these thresholds to the plot of L_d shows that the above analysis is partially true. Electrons could be generated at lower densities, for which L_d is longer, in the longer waveguide. At the threshold density for the 33 mm capillary $L \sim L_d$. The same is very approximately true for the 15 mm waveguide.

However the above analysis does not hold at all when considering the threshold density for the flat profile. At that density $L_d = 0.5 - 1$ mm, far less than Z_R , which was 4.5 mm. It might be argued that the laser required several mm to evolve at these densities and the extra length required for that evolution was provided by the waveguide. However those arguments can be dismissed by considering the experimental findings of [75]. In that work experiments were conducted using the Astra laser with a 2 mm gas jet to determine the point of injection of electrons through spatially resolved measurement of the side scattered radiation spectrum. The laser conditions used were identical to those of the Astra experiments of this thesis except that the laser spot of the gas-jet experiment had a slightly smaller radius ($W \approx 25 \mu\text{m}$, as opposed to $W \approx 34 \mu\text{m}$). At densities around the flat profile threshold (which was found to also be $\sim 20 \times 10^{18} \text{ cm}^{-3}$ in that experiment) electrons were injected into the wakefield a distance $\sim 200 \mu\text{m}$ after the start of the gas jet, which is negligible.

It can be concluded that at the threshold density for electron acceleration in the flat plasma the limiting factor was not the laser-plasma interaction length. Had it been, it would have been possible to have generated electrons with the flat profile at densities for which $L_d \approx Z_R$. That corresponds to $n_e \approx 5 \times 10^{18} \text{ cm}^{-3}$, which is four times lower than the experimentally observed flat-profile threshold for electron acceleration.

5.2 Self-focussing regime

The experimentally observed benefit of the plasma channel for densities below $20 \times 10^{18} \text{ cm}^{-3}$ could not be explained in terms of the longer acceleration length afforded by the plasma channel. The benefit of the waveguide becomes even less clear when self-focussing is considered. Laser wakefield acceleration in gas jets relies strongly on self-focussing of the laser pulse to form an intense filament that drives a non-linear wakefield [4, 5]. The ratio P/P_c is plotted in figure 5.1, where $P = 12 \text{ TW}$ is the laser power and P_c is the critical power for self-focussing [18]. It would be expected that at densities for which $P/P_c > 1$ the effect of a plasma channel should become less significant. This is not merely because the channel would no longer be required to guide the laser pulse, but because theory predicts that at these densities the self-focussed laser pulse has radius $\sim \lambda_p/2$ [11] which is far narrower than the matched spot radius of the channel. In that case the effect of the guiding channel would be negligible. However at the threshold density for the flat profile $P/P_c \sim 6$ meaning




Figure 5.1: Scaling L_d and P/P_c against density

The dephasing length predicted by both linear theory [74] (black line) and 3D non-linear theory [11] (violet line) are plotted against the left hand scale as a function of plasma density. The following lengths of interest are marked by horizontal dotted cursors: Z_R (black); 15 mm (blue); 33 mm (red). The ratio P/P_c (orange line) is plotted against the right hand axis, where P is the peak laser power of the Astra laser (12 TW) and P_c is the critical power for self-focussing [18]. The shaded regions are deduced from figure 4.13 and represent the minimum plasma density at which a wakefield was excited and electrons generated for: a flat plasma profile (green); a 15 mm waveguide (blue); a 33 mm waveguide (red).

that electrons could only be generated far above the self-focussing power in the absence of a plasma channel.

These questions then arise. Why could electrons not be generated at densities for which $L_d \ll Z_R$ and $P/P_c \gg 1$? Why did the presence of a plasma channel then allow electron acceleration at these densities.

Simulations

The most obvious route to answering these questions would be to run PIC simulations comparing flat plasma profiles and plasma channels. Such simulations were performed by C. Kampberidis using a 2-dimensional geometry but they did not show a large qualitative difference between the plasma channel and the flat profile at the densities of interest [69]. Strong wakes were excited in both cases, and the laser depleted long before the end of the waveguide.

It may be that 2D calculations underestimate the importance of diffraction and focussing. That is because the changes of intensity associated with changes in the radial profile of the laser scale with r^2 in three dimensions but r in two dimensions, where r is the radial co-ordinate. Since diffraction widens the radial profile of the laser pulse, and the plasma channel counteracts diffraction and narrows the radial profile it may be that the need for and impact of a plasma channel is lessened in a 2D calculation, and that might account for the fact that the simulations did not reflect the experimental findings.

The results of 3D PIC simulations in a plasma channel have been published, and they show that a plasma channel allows laser wakefield acceleration with a lower plasma density (and higher electron energy) due to guiding effects [11], and enhanced trapping of electrons at the edges of the plasma channel [76]. However in those cases the matched spot size of the channel was approximately equal to the self-focussed spot size, which is $\sim \lambda_p/2$. This is qualitatively very different to the case under consideration here — for the hydrogen filled capillary discharge waveguide the matched spot is more like $3\lambda_p/2$. To perform 3D PIC-simulations for Astra conditions would have required access to the relevant simulation codes and a computer cluster able to hand the heavy calculation. Neither were available.

Adapting self-focussing theory

Instead we proceed by attempting to simplify the problem to gain insight into the effect of the plasma channel. Self-focussing in a plasma is due to two different effects: the relativistic self-focussing due to mass changes of quivering electrons at high intensity, and the refraction of the laser pulse by its own wake. The latter effect can be both focussing and defocussing depending on the phase of the plasma wake. Theoretical treatment of self-focussing in the form of the quasistatic approximation [19, 20, 77] combines these effects into a single term that is a function of both the radial and temporal co-ordinates of the laser pulse. Self-focussing can be seen, therefore, as a convolution of radial and temporal effects. In the following sections of the chapter both are considered individually in order to gain insight into the effect that a plasma channel has on the propagation of relativistically intense laser pulses in plasmas. Temporal effects are considered in section 5.3, and radial effects are considered in section 5.4.

5.3 Temporal effects

Before proceeding it is important to clarify that the expression ‘temporal self-focussing effects’ is used in the following sections to mean the variation in the strength of radial self-focussing along the laser pulse. It does not refer to temporal pulse compression.

The principal result of temporal self-focussing effects is that the front of a laser pulse will not self-guide [19, 20]. At the leading edge of a laser pulse the change in plasma density due to the ponderomotive force is greater where the laser intensity is peaked, leading to a defocussing density profile which exactly cancels the relativistic self-focussing effect, meaning that the front of the pulse diffracts as if it was of low intensity. At a later phase in the plasma wave the plasma density profile is inverted to a focussing profile meaning that the part of the laser pulse at that phase will experience a focussing effect that is greater than that of relativistic self-focussing alone. If the laser is much shorter than a plasma wavelength then there will be no self-focussing or self guiding. In the case of most wakefield acceleration experiments, including those described in this thesis, the laser pulse length was of similar order to the plasma wavelength meaning that the front of the pulse would not have self guided but would have contributed towards guiding at the back of the laser pulse.

It has been suggested [11] that the benefit of using a plasma channel might be that it prevents the diffraction of the front of the laser pulse, which would then be depleted more slowly. For certain conditions and channel widths this will certainly allow LWFA acceleration at lower densities; and the applicability of this interpretation to the results from the Astra laser will be discussed in section 5.3.5. Additionally, where only a fraction of the laser pulses temporal profile is self-focussed the amplitude of the resulting wakefield would be less than expected due to the lower driving intensity. For consideration of both of these effects it is useful to derive expressions that describe self-focussing and external channel guiding simultaneously.

5.3.1 Quasistatic equations

The starting point for considering the temporal self-focussing effects are the 1D quasistatic equations derived by Sprangle et al [19, 20], which are given in this section. According to the quasistatic approximations the evolution of the laser pulse is assumed to occur on a much slower timescale than oscillations of the plasma wave. A coordinate frame (ξ, t) is adopted where $\xi = z - ct$. The laser pulse and plasma wave are described completely by the normalised transverse vector and scalar potentials $\underline{a} = |e|\underline{A}_\perp/m_e c^2$ and $\phi = |e|\Phi/m_e c^2$. By adapting the wave equation and Poisson’s equation the following formulae result

$$\eta = \left[1 - \frac{\omega_p^2}{\omega^2} \frac{1}{1 + \phi} \right]^{1/2} \quad (5.1)$$

where η is the refractive index experienced by the wave, ω_p is the plasma frequency of the unperturbed plasma and ϕ is given by

$$\frac{\partial^2 \phi}{\partial \xi^2} = \frac{k_p^2}{2} \left[\frac{1 + a^2}{(1 + \phi)^2} - 1 \right] \quad (5.2)$$

where $a^2 = \underline{a} \cdot \underline{a}$ (which means that a^2 is oscillating). The quantities ϕ and \underline{a} describe completely the wakefield properties through the following equations

$$dn/n_e = \frac{1}{2} [(1 + a^2)/(1 + \phi)^2 - 1] \quad (5.3)$$

$$\gamma = [1 + a^2 + (1 + \phi)^2] \quad (5.4)$$

$$\beta_z = [1 + a^2 - (1 + \phi)^2] / [1 + a^2 + (1 + \phi)^2] \quad (5.5)$$

These equations are one-dimensional, however they are often used to consider focussing effects [11, 19, 20, 78], which by their very nature include transverse dimensions. Equation (5.2) can be adapted to the 3-d case by considering the evolution of ϕ at different radii individually. This approach holds so long as the laser spot size W is much greater than the plasma wavelength. However, it shall be shown that even when $W \sim \lambda_p$ the expression derived for the self-focussed matched spot is in good agreement with 3D-PIC simulations. A three dimensional version of these equations has also been derived [77], but are considerably more complex. The arguments used below could be applied to those equations in future if it is deemed necessary.

5.3.2 Deriving a self-focussed matched spot size

To the knowledge of the author, and as stated in [11], there is not currently an analytic expression for the matched spot of a self-focussed laser pulse, whether in a plasma channel or in a flat profile. Such an expression is necessary to consider the importance of temporal self-focussing effects. We will proceed to derive an expression. The electromagnetic wave equation can be written as

$$\left[\nabla^2 - \frac{1}{c^2} \frac{\partial^2}{\partial t^2} \right] \underline{a} = \frac{\omega^2}{c^2} \underbrace{(1 - \eta^2)}_h \underline{a} \quad (5.6)$$

where \underline{a} is the normalised electric field of the laser radiation. This can be simplified by making the envelope approximation [18], with the substitution $\underline{a} = \text{Re} \{ \hat{\underline{a}}(x, y, z, ct) \exp[i(kz - \omega t)] \}$ and approximations $|\partial \hat{\underline{a}} / \partial z| \ll |k \hat{\underline{a}}|$ and $|\partial \hat{\underline{a}} / \partial t| \ll |\omega \hat{\underline{a}}|$ to give

$$[\nabla_{\perp}^2 + 2ik\partial/\partial z] \hat{\underline{a}} = k^2 h \hat{\underline{a}} \quad (5.7)$$

Note that $k = \omega/c$ was chosen, meaning that k is the wavenumber in vacuum. Introducing cylindrical coordinates and assuming symmetry about the axis leads to

$$\frac{\partial \hat{\underline{a}}}{\partial z} = -\frac{i}{2k} \left(k^2 h \hat{\underline{a}} - \frac{1}{r} \frac{\partial \hat{\underline{a}}}{\partial r} - \frac{\partial^2 \hat{\underline{a}}}{\partial r^2} \right) \quad (5.8)$$

This equation describes the evolution of the pulse envelope. If the refractive index profile is matched so that the beam is being perfectly guided then $\partial \hat{\underline{a}} / \partial z$ should be constant for all r . Assuming a Gaussian beam of radius W_m , and that the radial form of $h(r)$ is approximately parabolic for all r of interest this leads to the condition for guiding that

$$h'' = \frac{\partial^2 h}{\partial r^2} \Big|_{r=0} = \frac{8}{k^2 W_m^4} \quad (5.9)$$

This is a generally applicable expression for the guiding of a Gaussian beam. In the case in question we use equation (5.1) to say that

$$h = \frac{n(r)}{n_c} \frac{1}{1 + \phi} \quad (5.10)$$

where $n(r) = n_e + \frac{1}{2}n''r^2$. Cylindrical symmetry can be used to say that on axis

$$h'' = \frac{1}{n_c} \left(\frac{n''(1 + \phi) - n_e\phi''}{(1 + \phi)^2} \right) \quad (5.11)$$

When there is no channel $n'' = 0$ meaning that equation (5.9) and (5.11) can be used to give the self-focussed matched spot W_s . Where $\phi = 0$ and $\phi'' = 0$ there is no plasma response which means equation (5.9) and (5.11) gives the channel matched spot, W_c . This leads to the expressions

$$W_s = \left[\frac{-8(1 + \phi)^2}{k_p^2\phi''} \right]^{1/4}; W_c = \left[\frac{8n_e}{k_p^2n''} \right]^{1/4} \quad (5.12)$$

The expression for the channel matched spot is the same as that given in Chapter 2. Where n'' , ϕ and ϕ'' are all non-zero an expression for a self-focussed matched spot in a plasma channel results. This can be expressed in terms of W_s and W_c as

$$\frac{1}{W_m^4} = \frac{1}{1 + \phi} \frac{1}{W_c^4} + \frac{1}{W_s^4} \quad (5.13)$$

The expressions in (5.12) and (5.13) are deceptively simple looking, because W_s is a function of ϕ , which is itself a function of ξ (this is the temporal effect). However (5.13) does give an idea of how the channel guiding and self-focussing effects combine. The powers of 4 emphasise the fact that where the self-focussed matched spot is smaller than the channel matched spot one would expect the channel to have very little effect. All that remains then is to derive an expression for ϕ'' , which is easily found from equation (5.2)

$$\frac{\partial^2\phi''}{\partial\xi^2} = \frac{\partial^2}{\partial r^2} \frac{\partial^2\phi}{\partial\xi^2} \Big|_{r=0} = \frac{n''}{n_e} \frac{k_p^2}{2} \left[\frac{1 + a^2}{(1 + \phi)^2} - 1 \right] - \frac{k_p^2}{2} \left[\frac{4a^2/W(\xi)^2}{(1 + \phi)^2} + \frac{2\phi''(1 + a^2)}{(1 + \phi)^3} \right] \quad (5.14)$$

where $W(\xi)$ is the Gaussian radius of the laser pulse, which varies along the temporal profile. The first term on the right hand side results from adding a term $n(r)/n_e$ to the front of the right hand side of (5.2) to account for the change in plasma wavelength at different radii.

Comparison to 3D-PIC results

Before proceeding it is worth testing the validity of these expressions against a published result. A general, time independent expression for a self-focussed matched spot can be derived by considering the long pulse case for which ϕ is approximately constant with respect to ξ , as is $|\hat{a}|$. The left hand sides of equations (5.2) and (5.14) are set to zero and a^2 set to $\langle a^2 \rangle = a_0^2/2$. Setting $W = W_s$ the

following expression for the matched spot results

$$W_s^{\text{steady}} = \frac{\sqrt{8}(1 + a_0^2/2)^{3/4}}{k_p a_0} \approx \begin{cases} \frac{2.8}{a_0 k_p}; & a_0^2 \ll 1 \\ \frac{1.7\sqrt{a_0}}{k_p}; & a_0^2 \gg 1 \end{cases} \quad (5.15)$$

In references [11, 32] 3D simulations were used to determine a self-focussed matched spot size empirically. For a fixed $a_0 \gg 1$ the Gaussian pulse radius was varied until oscillations in the spot radius during propagation were minimized [79] (since there would have been a temporal dependence along the pulse it would have been impossible to have found a radius for which there was no oscillation). The deduced expression for the self-focussed matched spot is $W_s \approx \frac{2\sqrt{a_0}}{k_p}$. This means that the result of equation (5.15) is in very good agreement with 3D-PIC simulations suggesting that estimates of the laser profile made with the above expressions are reasonable, even where $W \sim \lambda_p$.

Aside from checking the above approach against empirical results the above equation can be reformulated usefully in terms of the laser power, which rather than a_0 , is the fixed parameter in a real experiment. For a Gaussian beam $a_0 W = \frac{2\sqrt{8}}{k_p} \sqrt{\frac{P}{P_c}}$, which means that the ideal self-focussed spot size is

$$W_s^{\text{steady}} = \frac{4}{k_p \sqrt{4^{2/3}(P/P_c)^{-1/3} - (P/P_c)^{-1}}} \quad (5.16)$$

5.3.3 The quasi-steady state

The expressions derived above simplify the radial aspect of focussing to two numbers, $\underline{a}(\xi)$ and $W(\xi)$. Since the on axis intensity is inversely proportional to W^2 the product $\underline{a}(\xi)W(\xi)$ should not change as $W(\xi)$ evolves during the propagation of the laser pulse¹. The laser pulse and plasma can reach a quasi-steady state for which diffraction is balanced by the combination of channel guiding and self-focussing for all ξ . For an initial temporal profile $\underline{a}_{t=0}(\xi)W_{t=0}(\xi)$, channel matched spot W_c and axial density n_e that quasi-steady state can be calculated by setting $W(\xi) = W_m(\xi)$, $\underline{a}(\xi) = \underline{a}_{t=0}(\xi)W_{t=0}(\xi)/W(\xi)$ and solving equations (5.2), (5.12), (5.13) and (5.14).

If it is assumed that as the laser is self-focussed from its initial radius it evolves towards the self-consistent quasi-steady state, then equation (5.3) can be used to give a simple estimate of $dn_e(\xi)$, the axial density perturbation after evolution. A value of $\mathcal{E}_z(\xi)$, the accelerating field experienced by the electrons, could also be deduced from ϕ . Examples of the deduced steady state, dn_e and E_z are plotted in figure 5.2. The form of the initial pulse was $\underline{a}(\xi) = a_0 \sin(k\xi) \sin(\xi\pi/2c\tau)\hat{\underline{i}}$ between $\xi = 0$ and $\xi = 2c\tau$ (where τ is the FWHM of the temporal intensity profile and $\hat{\underline{i}}$ is a unit vector representing the polarization of the laser). The matched spot size of the channel was set to $W_c = 34 \mu\text{m}$. At the two highest densities ($n_e = 8$ and $16 \times 10^{18} \text{ cm}^{-3}$) only the very front of the laser is not self-focussed. The very back of the pulse at $16 \times 10^{18} \text{ cm}^{-3}$ is defocussed by the plasma wake. At the lower densities ($n_e = 2$ and $5 \times 10^{18} \text{ cm}^{-3}$) the length at the front of the laser pulse that is not self-focussed is more significant. This provides an explanation as to why the energy of

¹Assuming that the temporal effects such as pulse compression and etching of the pulse front occur on a longer time-scale than self-focussing.

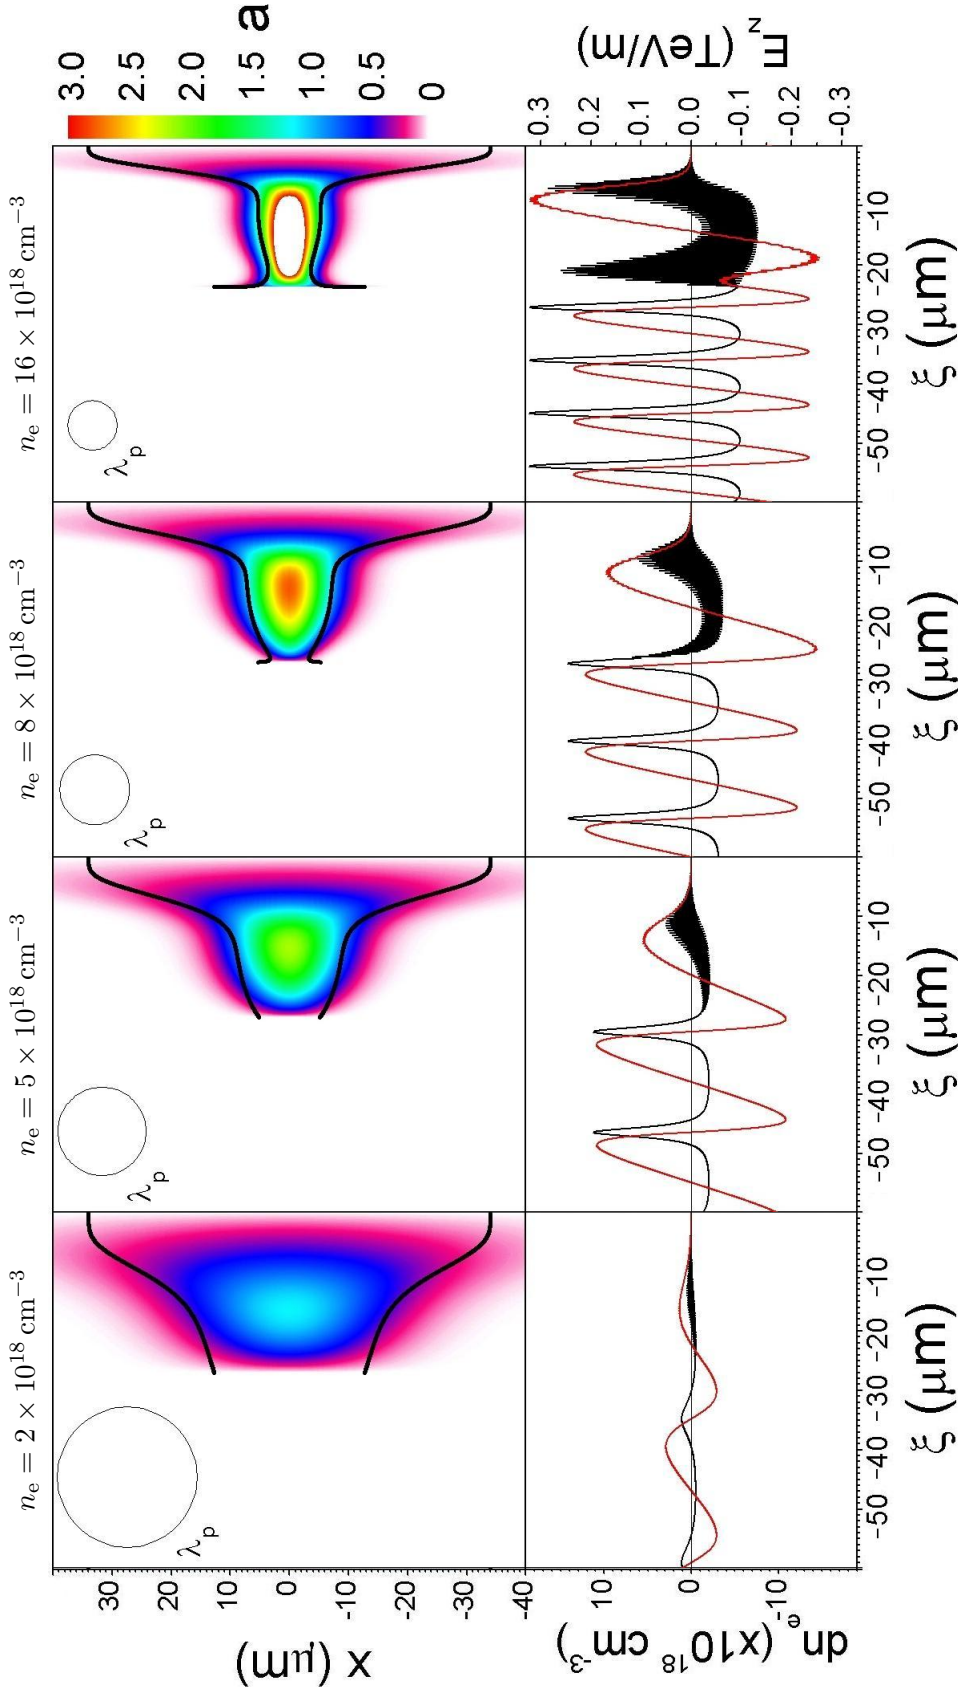


Figure 5.2: Laser pulses and wakefields in the quasi-steady state.

The form of the Astra laser pulse and the plasma wake in the quasi-steady state outlined in section 5.3.3, for four selected densities. In the top panel the normalised laser pulse amplitude is plotted across the pulse and along ξ , using the colour scale to the right of the figure. Circles whose diameter represents one plasma wavelength are plotted, as is the Gaussian spot radius $\pm W(\xi)$ (black lines). The axial wakefield, calculated from the top panel profile and the quasistatic equations of section 5.3.1, are plotted against ξ in the bottom panel. The perturbation of the plasma density $dn_e(\xi)$ is plotted against the left hand axis (black line), and the axial electric field $\mathcal{E}_z(\xi)$ is plotted against the right hand axis (red line). Note that the accelerating phase of the wakefield is that for which $\mathcal{E}_z(\xi)$ is negative. For all plots a plasma channel of $W_c = 34 \mu\text{m}$ was used.

electrons observed at densities below $8 \times 10^{18} \text{ cm}^{-3}$ began to drop away from the predictions of bubble regime scaling towards linear scaling. Also included in the figures are plots of dn and \mathcal{E}_z on axis. In keeping with all theoretical predictions the amplitude of the wakefield decreases as density decreases. Even at the lowest plasma density the waves are not sinusoidal meaning they do not fit with linear theory.

5.3.4 Electron energy scaling

It is stated in reference [11] that despite the physical differences, the form and field strength of the wakes calculated using 1D equations are approximately the same as those calculated in 3D, if the axial intensity is correct. The calculations of the quasi-steady state give an estimate of the intensity on axis after self-focussing. Presumably many λ_p behind the main pulse the plasma oscillations die out in the 3D case, but the calculated value of $\mathcal{E}_z(\xi)$ for the first plasma wave after the laser can be treated quantitatively.

By assuming that the laser pulse reaches the quasi-steady state, it is possible to estimate the expected electron energy. Following the lead of [11], the phase velocity of the wake is the group velocity of the laser pulse, which is approximately $c(1 - \frac{1}{2} \frac{\omega_p^2}{\omega^2})$, minus v_{etch} , the rate at which the front of the pulse is etched which is estimated in [27] to be approximately $c \frac{\omega_p^2}{\omega^2}$. Hence using a wakefield phase velocity of $v_\phi = c(1 - \frac{3}{2} \frac{\omega_p^2}{\omega^2})$ an estimate of the maximum possible energy gain of electrons co-propagating with the wake at a velocity very close to c is given by

$$\Delta E = -m_e c^2 \frac{c}{c - v_\phi} (\phi_b - \phi_e) = -m_e c^2 \frac{2n_c}{3n_e} (\phi_b - \phi_e) \quad (5.17)$$

where ϕ_b and ϕ_e are the values of ϕ at the beginning and end of the accelerating phase in the first plasma wave behind the laser pulse.

Computing the self-consistent form of the laser pulse takes a small fraction of a second. Therefore it is possible to predict the maximum energy that can be achieved by electrons across a range of densities. The results are plotted in figure 5.3(a). The calculation of maximum energy gain agrees well with the maximum energy of the data². Unlike the bubble regime scaling and the linear scaling this prediction is in approximate agreement with the data across a broad range of densities, meaning that it successfully predicts the transition between the two regimes. The only exception to this is at $n_e \approx 5 \times 10^{18} \text{ cm}^{-3}$, where the measured energy was lower than the predictions by a factor of approximately two. It could be the case that at this density, which was the lowest for which electrons were detected in the 15 mm capillary, the laser exited the waveguide before electrons could reach their maximum possible energy. In chapter 6 it is found that at that density the ionization state of the channel had a role to play in injection of electrons into the wakefield, and it could also have been that in this case the assumption that electrons are injected at approximately the back of the first plasma wave is wrong.

It is also interesting to note figure 5.3(b), which plots the predicted maximum accelerating field against density. For $n_e > 8 \times 10^{18} \text{ cm}^{-3}$, which are the densities corresponding to bubble regime

²An unusual feature of 5.3(b) is that at very high density is a discontinuity in the curve for predicted energy. That is an artifact of the automated method of identifying the back of the first plasma wave.

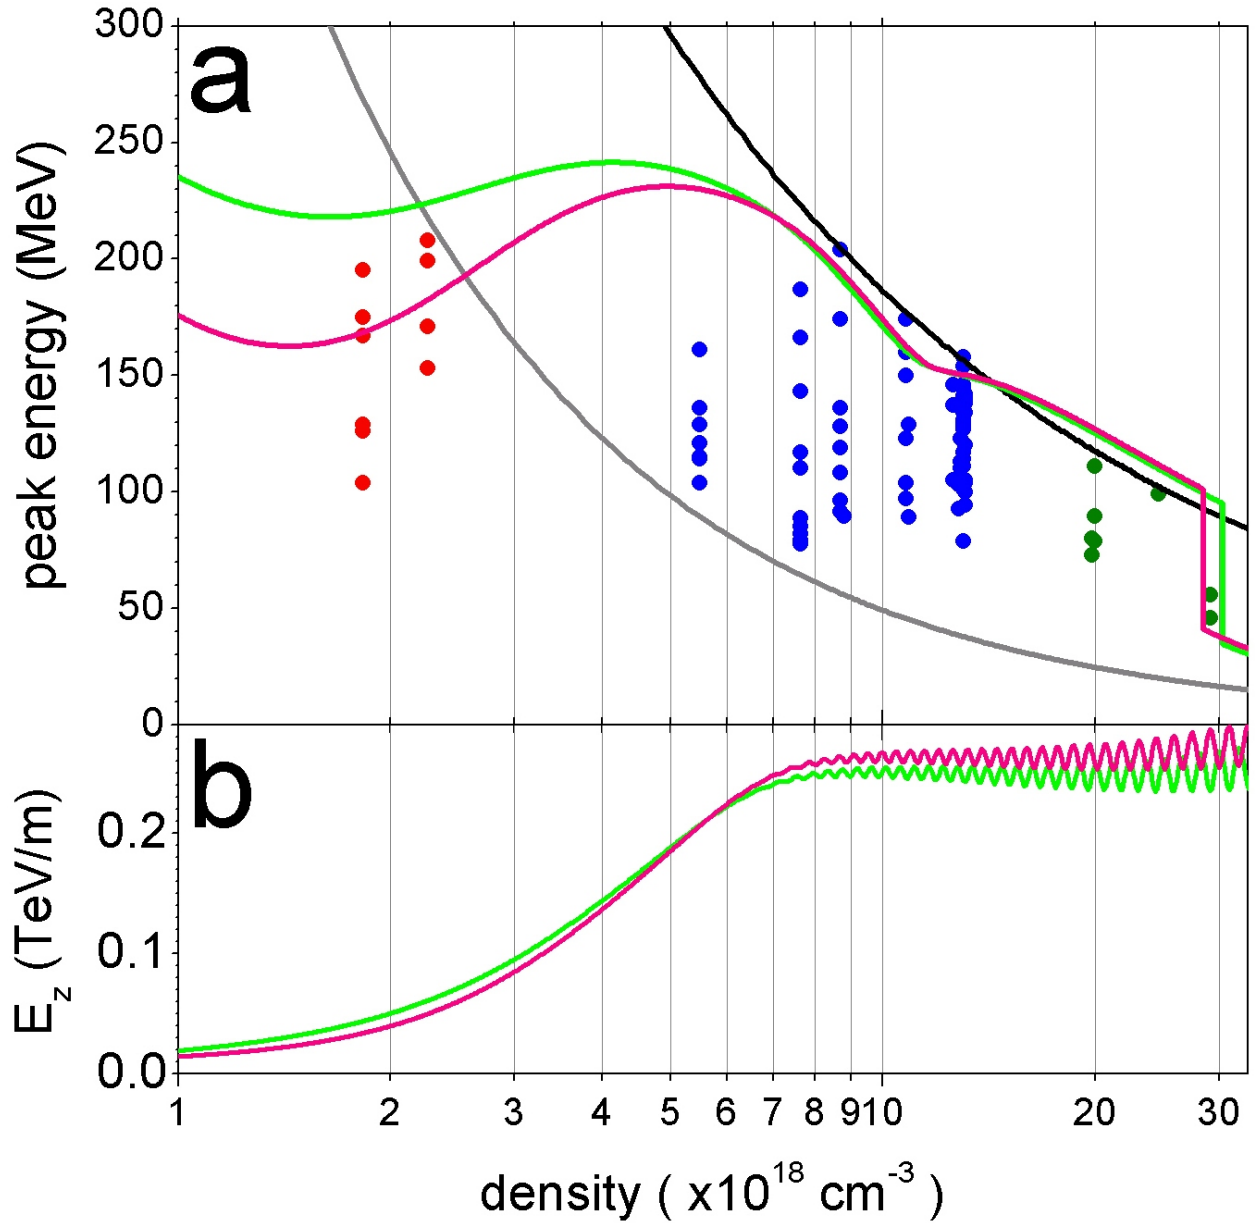


Figure 5.3: Scaling for the quasi-steady state

In the top panel the electron energy data of figure 4.8 is plotted against density with the predicted maximum energy from scaling laws in the linear regime [74] (grey line) and the 3D-non linear regime [11] (black line). The maximum electron energy predicted from the quasi-steady state in a plasma channel (outlined in sections 5.3.3 and 5.3.4) is plotted for the waveguide (green line). The maximum electron energy predicted by the quasi-steady state for a flat plasma profile (outlined in section 5.3.5) is also plotted (pink line). In the bottom panel the peak electric field of the wake in the quasi-steady state is plotted against density for the plasma channel (green line) and the flat profile (pink line).

scaling, the calculated peak electric field is approximately constant³. As the density is decreased and the energy prediction drops away from bubble regime scaling the value of E_z also begins to drop off. A comparison of figure 5.3(b) to figure 4.12 also reveals that it was at these densities that the generation of electron beams was very sensitive to discharge timing t . This is considered in more detail in Chapter 6.

It can be concluded then that by considering how the spot size of the laser pulse varies along its length a reasonable prediction of the expected electron energy can be made using the quasistatic equations. In particular, the observed transition between bubble regime scaling at high density and linear regime scaling at low density was predicted quantitatively.

5.3.5 Comparison to flat profile

Predicting a steady state for a gas jet using the above arguments is not possible because at the front of the pulse $W_m = \infty$. This means that at all successive ξ , W_m is equal to infinity. In effect the self consistent state with no external guiding channel is a completely diffracted pulse. However over distances less than a rayleigh range the laser pulse can reach a quasi-steady state for which the back of the laser pulse is self-focussed and the front slowly diffracts. An estimate of this state can be made in exactly the same way as the estimates made for the channel if it is assumed that the length-scale over which the laser reaches the self consistent state is shorter than the rayleigh range. So, for the flat profile the quasi-steady state is defined in the following way: at the front of the pulse $W_s(\xi) > W_0$, so we set $W(\xi) = W_0$, where W_0 is the input spot size. Further back the pulse can self-guide, so that for $W_s(\xi) < W_0$ we set $W(\xi) = W_s(\xi)$. Other than that difference the calculation was exactly the same as those for the plasma channel. The predictions of maximum electron energy and peak electric field using this method are also plotted in figure 5.3. The predicted energy barely deviates from that in a waveguide. That means that with the approximations made the theoretical strength of the wakes produced in a uniform plasma are not very different from that produced in the plasma channel. This then does not explain why the threshold density for LWFA in the flat profile was so high.

We return to the argument that an external guiding structure can allow electron acceleration at lower densities for a given power by preventing diffraction of the front of the pulse. Arguments are put forwards in references [11, 19, 20] all of which calculate a velocity, in ξ co-ordinates, at which the front of the laser pulse recedes through diffraction. The expression for the diffraction velocity, v_d is

$$v_d = \frac{\Delta\xi}{Z_r} c \quad (5.18)$$

where $\Delta\xi$ is the length from the front of the pulse to the ξ for which $W_s < W$. References [19, 20] estimate $\Delta\xi$ to be $\sim \lambda_p$. Reference [11] estimates $\Delta\xi$ to be $(k_p^2 W a_0)^{-1}$ where a step temporal profile is assumed. Using the definitions for W_s and ϕ'' from equations (5.15) and (5.14) we can argue that for small ϕ and ϕ'' that $\phi'' \approx \frac{-k_p^2 a^2}{W^2} \xi^2$ and setting $W(\Delta\xi) = W_0$ and using the time

³The oscillations in the value of E_z at these densities result from the overlapping of the laser pulse with the back of the wakefield causing slight oscillations in the calculate field.

averaged value for a^2 we find

$$\Delta\xi = \frac{4}{k_p^2 a_0 W} \quad (5.19)$$

This approach is adapted from that used in [11]. The factor of 4 difference arises from a difference in interpretation of the term a^2 and the definition of W_s . However for the purposes of explaining the observed trends in the data these differences are largely academic: using the above expression for $\Delta\xi$, a density of $n_e = 20 \times 10^{18} \text{ cm}^{-3}$ and applying equation (5.18) to the front of the pulse where $W_0 = 34 \mu\text{m}$ it is found that $v_d/v_e = 0.007$, where v_e is the velocity at which the front of the pulse is etched through pulse depletion. Since v_d is negligible compared to v_e , it can be concluded that at $n_e = 20 \times 10^{18} \text{ cm}^{-3}$ diffraction of the front of the pulse was not the limiting factor in the acceleration process.

We conclude that although the approach used in this section, which simplified the radial aspects of self-focussing and focussed on temporal aspects, does explain observed relationship between electron energy and plasma density, it does not explain why electrons could not be accelerated at lower densities without a plasma channel. It would appear that the expressions calculated provide a reasonable description of the laser pulse in a plasma channel, but not in a flat plasma profile.

5.4 Radial effects

A key simplification to the approach of section 5.3 was that the radial profile of the laser intensity could always be described as a Gaussian, and that $h = 1 - \eta^2$ was always parabolic (unless flat). An ODE method outlined in [18] shows that a Gaussian laser pulse experiencing a parabolic h will remain Gaussian even though the width of the Gaussian profile may change. A consequence of this is that in the work in section 5.3 wherever the laser was self focussed to an intense filament, all of the energy in the original laser profile would be focussed down into that central filament.

In reality not all of the laser energy is focussed into the central filament. Equation (5.15) applies on axis, but at wider radii h deviates from a parabolic form. This means that different parts of the transverse profile of the pulse are focussed inwards at different rates, and for large r , where the intensity is low and the refractive index is virtually flat, the energy will diffract and be lost to the interaction. This then raises two questions. What fraction of energy will form part of the central filament? Does less energy escape if the process occurs within a guiding channel?

It could be the case that a significant amount of energy that would not otherwise have formed part of the self-focussed central filament is focussed inwards by a guiding channel to be ‘caught’ and included in the interaction. If this were the case then using a plasma channel would effectively be like using a more powerful laser. This would then provide an answer to the question of why electrons could be generated at such lower densities when a plasma channel was employed.

5.4.1 Self-focussed propagation code

In section 5.3 the radial aspects of the self-focussing process were simplified in order to isolate trends caused by temporal effects. In this section temporal effects are simplified in order to concentrate on effects relating to the evolution of the radial profile of the laser pulse. Equation (5.6) was solved

numerically with the following expression for h [18–20]

$$h = 1 - \eta^2 = \frac{1}{\gamma} \frac{\omega_p^2}{\omega^2} = \underbrace{\frac{1}{(1 + a \cdot a)^{1/2}}}_{\text{Instantaneous}} \frac{\omega_p^2}{\omega^2} = \underbrace{\frac{1}{(1 + \langle \hat{a} \rangle^2 / 2)^{1/2}}}_{\text{Time-averaged}} \frac{\omega_p^2}{\omega^2} \quad (5.20)$$

As in section 5.3 the envelope approximation [18] was applied, however in this case $k(r, z)$ was defined as

$$k(r, z) = \eta(r, z) \frac{\omega}{c} \quad (5.21)$$

which was found to minimise numerical instabilities. For non-zero z this led a complicated expression, however for $z = 0$ the envelope approximation gives

$$\frac{\partial \hat{a}}{\partial z} = \frac{i}{2k} \nabla_{\perp}^2 \hat{a} \quad (5.22)$$

The effects of self-focussing and channel guiding are all included in the expression for k . For an initial pulse profile the calculation proceeded by using equation (5.22) to calculate the evolution of the laser profile in steps of dz . Rather than use a more complicated expression the complex envelope \hat{a} was multiplied by $e^{-ik(r,z)\Delta z}$ at each z step of the calculation in order to ‘reset’ z to zero. The reduction in numerical stabilities due to this definition was essential in allowing this calculation to be extended to 3-dimensions without need for symmetry.

The validity of the code was tested by simulating known cases, such as a diffracting pulse and a mismatched Gaussian pulse scalloping in a parabolic channel as described in [18]. The total energy of the radiation was measured at each step so as to detect false results caused by numerical instabilities. If the total energy deviated by more than 1% from its initial value the calculation was stopped.

The method outlined represents a CW pulse propagating in a plasma, so the self-focussing effects are exaggerated. They do however provide a reasonable description of the back of the pulse, particularly at higher densities where $c\tau \gtrsim \lambda_p$. Therefore any qualitative difference observed between the gas jet and the waveguide in these simulations will in fact be an underestimate of the real experimental difference since inclusion of temporal effects in the simulation would reduce the effect of self-focussing.

5.4.2 Propagation of an initially Gaussian laser spot

Calculations were performed to recreate the conditions of the Astra experiment by introducing Gaussian laser intensity profiles to flat or parabolic plasma density profiles. In all cases the Gaussian width of the focal spot was $34 \mu\text{m}$ and the parabolic curvature of the channel was based on the matched spot scaling laws of equation (2.14) for a $200 \mu\text{m}$ wide capillary. A laser power of 12 TW was used. Initially the calculation ran until the laser had reached a maximum intensity on axis and had started to diffract, to give an idea of how much energy was trapped as the energy focussed down to a central filament. It was possible to run this code repeatedly for a range of densities so as uncover any threshold density for which the behaviour of the laser pulse in the plasma channel deviated from that of the laser pulse in the flat profile.




Figure 5.4: Fraction of laser energy trapped within a plasma wavelength radius. The results of running the calculation outlined in section 5.4.1 for a range of densities are plotted. At each density the laser intensity profile at the value of z for which the intensity was peaked was isolated, and the fraction of the total laser energy trapped within half a plasma wavelength radius was calculated. The resulting fraction of energy is plotted against density for calculations using a plasma channel (blue line) and a flat plasma profile (green line). The threshold density for electron acceleration in the flat profile is also marked (green area).

The results of these simulations are plotted in figure 5.4. For every density the fraction of the laser energy trapped within half a plasma wavelength radius is plotted (it is assumed that if a filament were to form it would have radius $\sim \lambda_p/2$). Although for densities $n_e \sim 2 \times 10^{18} \text{ cm}^{-3}$ ($P/P_c \approx 1$), there is significant difference between the flat profile calculation and the plasma channel calculation, at $n_e \sim 20 \times 10^{18} \text{ cm}^{-3}$, the threshold density for electron acceleration in the flat profile, there is little difference in the calculation for the plasma channel and that for the gas jet.

For selected densities the calculations were continued along what would have been the full length of the waveguide, although of course it is understood that at higher densities the pulse would have depleted long before reaching the end of the plasma. The results are plotted in figure 5.5. There is a qualitative difference between the plots for the plasma channel and those for the flat profile: the low intensity light towards the edges of the laser profile diffracts away for the gas jet, but is channeled by the waveguide. This does not appear to be significant however: for all densities except $2 \times 10^{18} \text{ cm}^{-3}$ the intensity of the light on axis is not significantly different when the two cases are compared, and changes in density affect the behaviour far more than the presence of a plasma channel. Self-focussing effects are dominant, and the presence of a plasma channel appears to make little difference to the behaviour.

5.4.3 Adapting to non-ideal conditions

The approach of section 5.3, in which the temporal aspect of self-focussing was emphasised, could not explain why the density threshold for LWFA in a flat plasma was so high relative to that for the plasma channel. Nor did the initial consideration of the radial aspects of self-focussing in 5.4.2. We are then left to ask, what other effects could there be?

A clue is provided by considering in detail the experimentally measured exit modes of the long pulse at the densities in question, and comparing these to the input spot. Four examples of each are plotted in figure 5.6, along with the peak camera signal, the full width half-maximum of each spot and the total integrated signal within that FWHM. Note that for the long pulse, where the intensity is lower, the effect of self-focussing is negligible, meaning that the observed pattern at the capillary exit is determined largely by the effect of the plasma channel. It is interesting then that the peak signal measured on the camera is almost twice as high after 15 mm of guiding than at the input of the capillary.

An obvious explanation for this would be that the matched spot size of the channel was smaller than the input pulse, and that the beam is focussed down to a smaller size at the capillary exit, resulting in a higher intensity. This is not the case however. Aside from the fact that the matched spot size of the channel is predicted to be approximately the same as the size of the input spot at these densities, the measured full width at half maximum is also approximately the same for the output spots as for the input spots. It would appear that the beam has become more intense without becoming narrower. In fact the integrated signal within the FWHM was almost a factor of two greater where the waveguide was employed.

Clearly the energy can not have been ‘created’ within the waveguide. As discussed in section 3.2.2 where a stronger ND was employed the camera was unable to detect light at the lowest intensities, and the input and output spots of figure 5.6 were recorded with that filter in place.

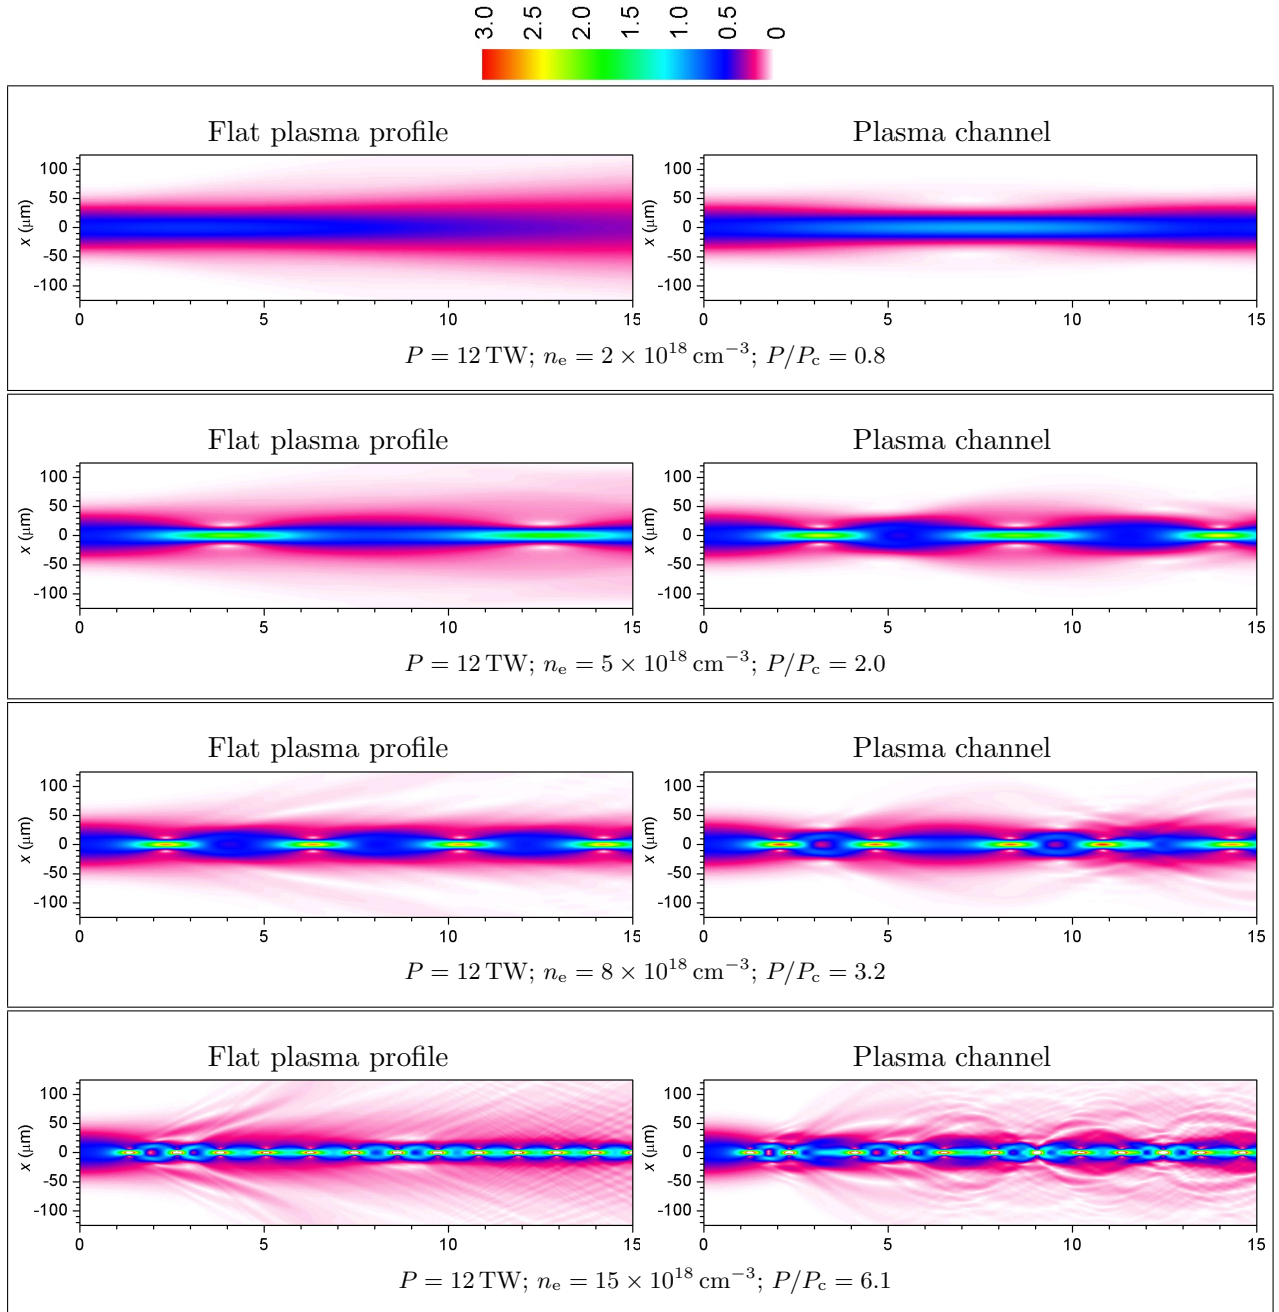


Figure 5.5: Propagation of an intense Gaussian laser in flat plasmas and channels. The propagation of a continuous 12 TW, laser pulse with a Gaussian focal spot profile was calculated in cylindrical geometry for a range of plasma densities using the method outlined in section 5.4.1. The profile of the normalised laser amplitude (colour scale) is plotted against z for the case of a flat plasma profile (left hand plots) and the parabolic plasma channel (right hand plots).

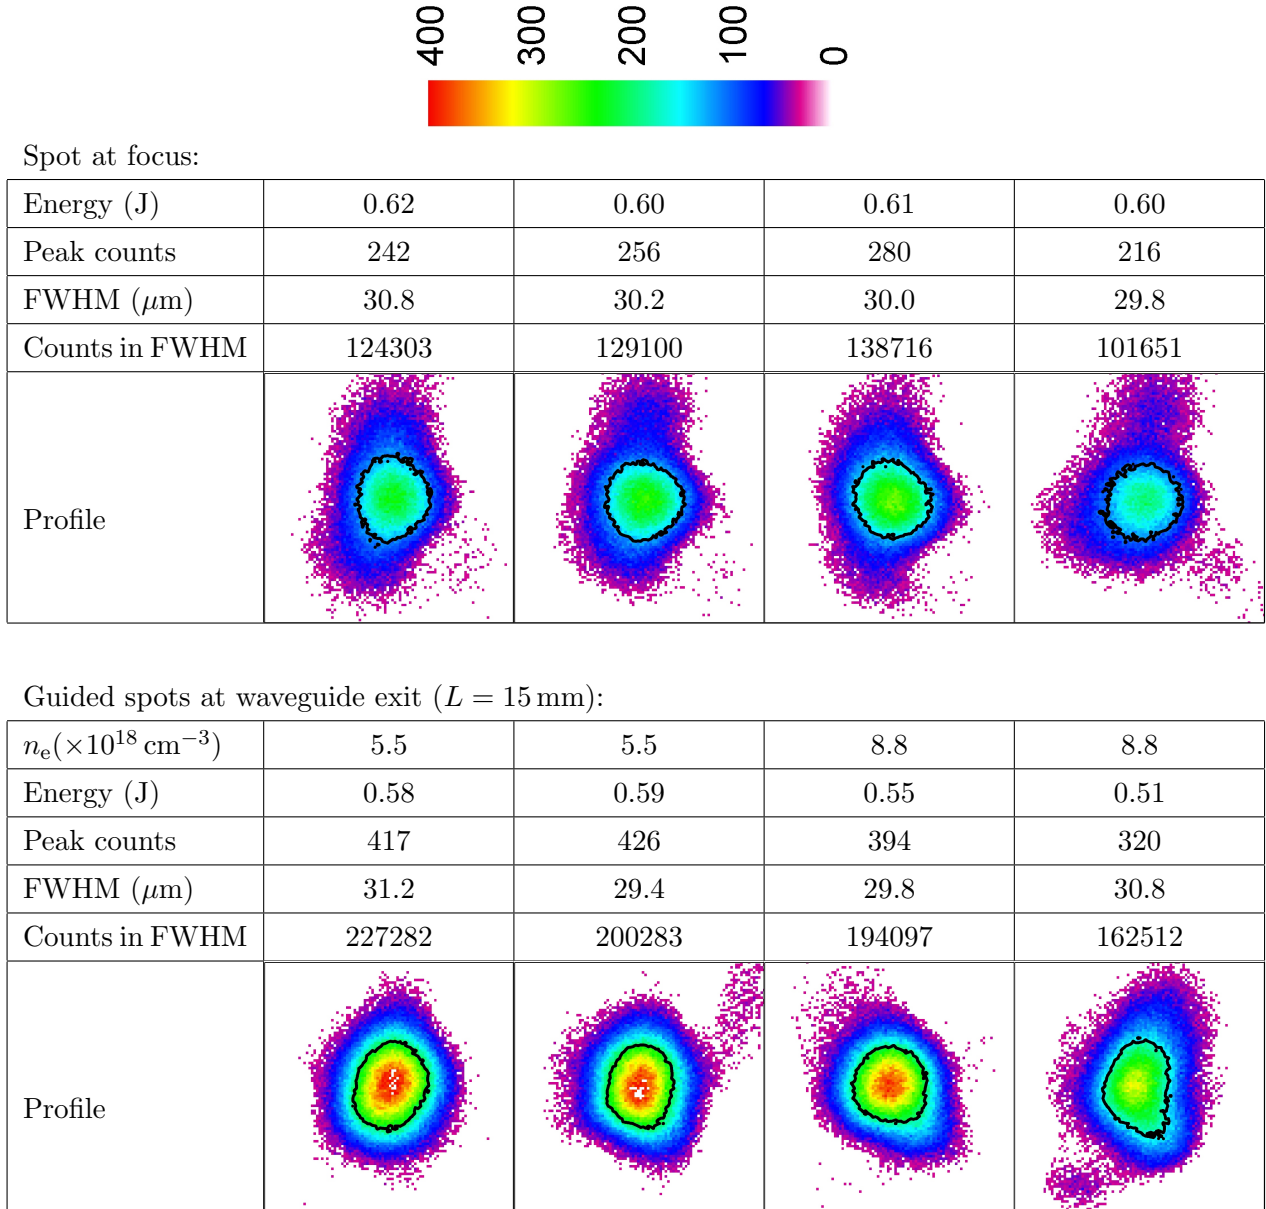


Figure 5.6: Measured input and output transverse fluence profiles of the long pulse. Each plot represents a square of side $200 \mu\text{m}$ in the plane of either the input spot focus (top table) or in the plane of the exit of a 15 mm waveguide (bottom table). The measured signal (colour scale) on the fluence monitor is plotted for each case. For each spot the contour representing half of the peak signal is marked (black line). The input spots were recorded in a ‘dummy run’ immediately after the output spots without change to the experimental setup (except removal of the waveguide from the beam-line). The energy included for each shot is that measured from the output energy diode. This value is included to show that changes in the signal recorded on the camera were not due to fluctuations in the laser energy. The data plotted includes a corrective offset as outlined in section 3.2.2 meaning that signals between 0 and 30 were not detected.

However the focal scan of figure 3.2, which was conducted with a 12-bit camera placed at focus, shows that there is a significant amount of energy lying at low intensities outside the main central peak. Figure 5.6 suggests that for the long pulse the plasma channel ‘collects’ the outlying light and brings it into the main peak.

Realistic laser and channel profiles

For the short pulse, self-focussing alone would not necessarily be enough to ‘collect’ the outlying energy, because of its low intensity and large area spread. Were that the case, the outlying energy would not be included in any central self-focussed filament, and would not contribute to driving a wake. It is reasonable to expect that if the waveguide ‘collected’ outlying energy for the long pulse, it would also be able to do so for the short pulse. That energy then could form part of the central filament, and would contribute towards driving a wake. If real, this effect would explain the effect of the plasma channel on experimental results. In order to test this, the code that solved equations (5.6) and (5.7) was adapted to include both x and y dimensions of the fluence profile. This meant that the measured focal spot could be used for the simulation.

In addition to changing the laser profile used in simulation to better reflect the experimental reality, the density profile of the plasma channel was adapted. Theory [60, 62, 63, 80] and experiment [64] show that the channel is not parabolic, but is flat on axis and steeper towards the wall of the capillary. The matched spot for a non-parabolic channel is determined by satisfying the condition from [81] that the matched spot size is the minimum radius r for which the plasma density exceeds the axial density by $(\pi r_e r^2)^{-1}$, where r_e is the classical radius of the electron.

The lowest order description of a non-parabolic channel is given by

$$n(r) = n_e + \frac{1}{2}n''r^2 + \frac{1}{24}n''''r^4 \quad (5.23)$$

Because there are three terms in that expression, three pieces of information about the channel are required in order to determine the channel profile. The form of the parabolic channel was determined from n_e and W_c predicted by equations (2.13) and (2.14). A third piece of information can be deduced by assuming that the laser pulse ionizes any unionized hydrogen in the channel, in which case the mean number of electrons inside the channel should be equal to twice n_{H_2} , the initial density of hydrogen molecules. From n_{H_2} , n_e and W_c the terms in equation (5.23) are given by

$$n_e = n_e \quad (5.24)$$

$$\frac{1}{2}n'' = \frac{1}{\pi r_e W_c^4} - \frac{W_c^2 n''''}{24} \quad (5.25)$$

$$\frac{1}{24}n'''' = \frac{\frac{3}{4}D^2 - 6\pi r_e W_c^4 (2n_{\text{H}_2} - n_e)}{\pi D^2 W_c^4 r_e (3W_c^4) - \frac{1}{2}D^2} \quad (5.26)$$

Note that for a given set of waveguide parameters, this represents the lowest order estimate for the plasma channel profile. The parabolic estimate does not satisfy the condition that the total number of electrons should remain unchanged before and after the formation of the channel. Therefore the

estimate that includes a quartic term must be a closer estimate of reality.

Due to the increased computation of an extra dimension it was not possible to repeat simulations across a range of densities like in section 5.4.2. Simulations were run for $n_e = 8 \times 10^{18} \text{ cm}^{-3}$. At this density $P/P_c \approx 4$ and $L_d \ll Z_r$ meaning that electrons *should* have been accelerated for the flat profile, but its a long way below the threshold density for which electrons actually *were* accelerated in the flat profile.

The results of simulations for the first 2.4 mm of propagation are plotted in figure 5.7. For the flat profile a central filament forms but the energy at wider radii is not drawn in as this is happening. Where a waveguide was employed the energy away from the centre of the capillary is forced inwards and becomes part of the main interaction. For the gas jet the peak a_0 was 2.5 and occurred after 2.3 mm. In the plasma channel the corresponding numbers were 3.4 and 2.0 mm. Note that a_0 is proportional to the square root of intensity meaning that there was almost a factor of two difference in peak intensity. Most tellingly is the fact that at the peak 48% of the laser energy lay within a FWHM max for the plasma channel, compared to 31% for the flat profile. It should be taken into account that for a perfect Gaussian 50% energy lies within the FWHM, meaning that the plasma channel forces almost 100% of the laser energy to form part of the central filament, whilst for a flat plasma profile the number is more like 60%.

More qualitative difference can be seen between the two cases when comparing the propagation of the two pulses beyond the initial focus. This is plotted in figure 5.8. Self-focussing effects cause the laser pulse in the gas jet to scallop between high and low intensity without ever sustaining high intensity. Two principal observations may be made in the case of the waveguide. Firstly energy that escapes the filament after the central focus is refocussed back into the middle focus by the waveguide. Secondly a sustained central filament forms that, unlike that in the gas jet case, does not scallop but maintains a constantly high intensity.

By using more realistic profiles for the laser focus and plasma channel it has been shown that at the density of interest there is a significant difference between the behaviour of the laser propagating in a flat plasma and that propagating in a plasma channel. It is also interesting to note that an ideal (i.e. flat) phase was used for the initial profile in the simulations. In fact the focal scan of the laser focus showed that the beam was astigmatic meaning that the phase was certainly not flat. This may very well have further increased the discrepancy between the plasma channel and the flat profile since at the focal plane not all light would have been travelling in the overall direction of propagation. This would further reduce light trapped within a central filament by self-focussing alone.

5.5 Conclusion

The aim of this chapter was to establish a reason why, for the experiments performed with the Astra laser, the presence of the plasma channel enabled the acceleration of electrons at low density.

The effect of extending the laser-plasma interaction length with the plasma channel was considered. At the very lowest densities, approaching the linear regime, it was found that the length of the plasma channel was the limiting factor on electron acceleration. However that could not account for the fact that the plasma channel was also necessary for the acceleration of electrons




Figure 5.7: Simulated initial laser pulse evolution using the real laser profile and quartic channel.

The propagation of a continuous 12 TW, laser pulse with the measured Astra focal spot profile was calculated in 3-d geometry for a plasma density of $n_e = 8 \times 10^{18} \text{ m}^{-3}$ using the method outlined in section 5.4.1. The profile of the normalised laser amplitude (colour scale) is plotted for each 0.3 mm of the first 2.4 mm of the plasma. Each plot is a $210 \mu\text{m}$ square, although the calculation window was wider than this. Results for the flat plasma profile are plotted in the top row, results for the plasma channel (with a 4th order profile determined by the method outlined in section 5.4.3) are plotted in the bottom row.

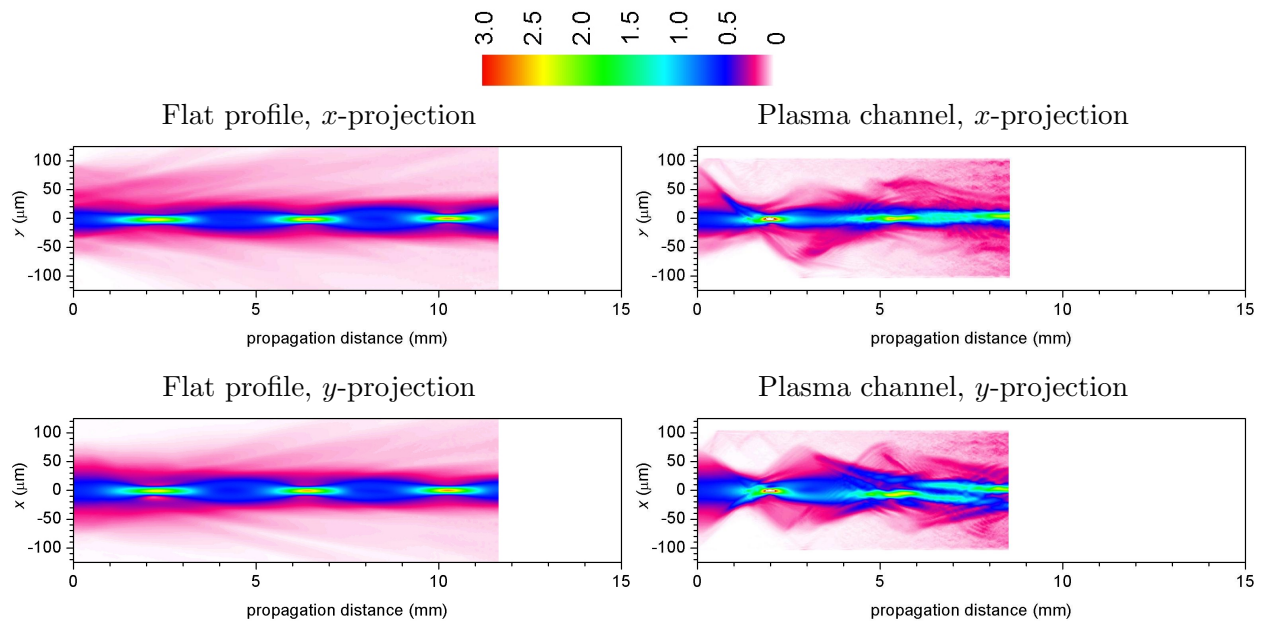


Figure 5.8: Propagation of an intense laser with the Astra profile in flat plasmas and channels. The propagation of a continuous 12 TW laser pulse with the real laser profile in a plasma was calculated in 3D geometry for a plasma density of $n_e = 8 \times 10^{18} \text{ cm}^{-3}$ using the method outlined in section 5.4.1. In order to represent the propagation in 2 dimensions side projections were created. For the x -projection the peak value of the x -amplitude profile was plotted for each (y, z) . The y projection was the same but with x and y exchanged. Both projections are plotted for the case of a flat plasma profile (left hand plots) and the plasma channel (right hand plots). Note that the development of instabilities meant that the simulations were not continued beyond ~ 10 mm however they are plotted on the same scales as the plots in figure 5.5 for ease of comparison.

at densities immediately below $n_e \approx 20 \times 10^{18} \text{ cm}^{-3}$ for which the required interaction length was much shorter than the rayleigh range of the laser, and the power of the laser was far in excess of the critical power for self-focussing.

There was not an opportunity to run fully three dimensional simulations to analyse the role of the plasma channel in the self-focussing regime. In order to gain insight into the key effects, self-focussing was considered in two different ways.

Firstly the temporal aspect of self-focussing, in other words the fact that the degree of self-focussing varies along the temporal profile of the laser pulse, was investigated. Analytic expressions for the matched spot size resulting from both self-focussing and the presence of a channel were derived. These were found to be in agreement with empirical observations of 3D-PIC simulations [11, 32], and were used to calculate the form of the laser pulse in a self-consistent quasi-steady state. From this the accelerating field on axis could be estimated, and the maximum possible energy gain predicted. The predictions of electron energy agreed well with the data. According to this model, the cross-over between the bubble and linear regimes was due the fact that at intermediate densities a significant fraction of the leading edge of the pulse was not self-focussed causing a reduction to the axial laser intensity and wakefield amplitude. However, despite being able to account for the cross-over between bubble regime and linear behaviour, considering temporal effects, like considering the acceleration length, could not explain why electrons could not be generated in a flat plasma profile at densities immediately below $n_e \approx 20 \times 10^{18} \text{ cm}^{-3}$.

In order to further investigate the combined roles of the plasma channel and self-focussing the evolution of the radial laser profile was considered. A code was written that simulated the propagation of a continuous-wave laser in a plasma. Initial simulations using ‘ideal’ initial conditions (a Gaussian pulse and a parabolic plasma channel) showed that the plasma channel made little difference to the propagation of the laser in the self-focussing regime. However the measured laser profile was not ideal and had significant amounts of energy lying at low intensity over a large area outside of the main peak. Further, examination of the transmitted long pulse fluence profiles suggested that the plasma channel ‘collected’ that energy and brought it to the laser axis. Therefore the code was adapted to simulate ‘non-ideal’ measured input conditions: the measured intensity profile was used and the plasma channel profile was adapted to better represent the real waveguide. For both profiles an intense filament formed due to self-focussing. However for the flat plasma profile the outlying energy did not form part of that filament, but diffracted away. In the plasma channel that energy was forced inwards and did form part of the self-focussed filament. Only 60% of the total laser energy was included in the self-focussed filament that formed in the flat plasma profile, compared to almost 100% in the plasma channel. The importance of the channel at that plasma density was not that it extended the laser-plasma interaction length, but that it allowed more efficient use of all of the available laser energy.

Chapter 6

Electron trapping at low plasma densities

The experimental data from the Astra experiments presented in chapter 4 shows that the generation of electron beams in the hydrogen-filled discharge capillary waveguide depended on discharge-laser delay t . This is not surprising: changes to t meant changes to the state of the plasma channel encountered by the laser. At plasma densities above $6.5 \times 10^{18} \text{ cm}^{-3}$ it could be seen that the delays for which electron beams were likely to be generated were also those for which the waveguide transmission was highest suggesting that the delays for which electron beams were generated were those for which a quasi-stable guiding channel was formed. The importance of the channel in allowing efficient use of the full laser energy was explored in chapter 5, and the apparent correlation between the presence of a plasma channel and the acceleration of electrons is not unexpected.

However at densities below $6.5 \times 10^{18} \text{ cm}^{-3}$ the data could not be interpreted so simply. The range of t for which beams were accelerated was much narrower than the range for which the waveguide transmission was high. In terms of guiding, the behaviour was similar to that observed at higher densities, but the generation of beams became extremely sensitive to delay. The narrow range of delays for which beams were generated also appeared to coincide with t_f , the delay at which the waveguide transmission was observed to fall.

The sensitivity of electron acceleration to t , and the apparently related fall in transmission at t_f , were unexpected. Interestingly, it was also found that the most stable beams of the experiments were generated under such conditions. Figure 6.1 plots all electron-spectrometer images for [$p = 250 \text{ mbar}$, $L = 15 \text{ mm}$, $D = 200 \mu\text{m}$] in order of t . At that pressure $n_e \approx 5.5 \times 10^{18} \text{ cm}^{-3}$, given by equation (2.13). Of 13 shots in the range $130 \text{ ns} < t < 135 \text{ ns}$, none yielded electron beams; in contrast for $137 \text{ ns} < t < 141 \text{ ns}$ 7 of 9 shots gave quasi-monoenergetic beams. In that latter interval the mean and standard deviation of the beam energies were 128 and 15 MeV.

In this chapter the conditions inside the waveguide are studied in detail in order to establish how the evolution of the plasma channel during the discharge related to guiding and electron acceleration. Particular attention is paid to the data for [$p = 250 \text{ mbar}$, $L = 15 \text{ mm}$, $D = 200 \mu\text{m}$] to determine what property of the plasma channel at $t \approx t_f$ enabled such consistent generation of electron beams. Details of experimental and analytical methods are given in 6.1. The resulting data is compared to the scaling laws of [45, 64] in section 6.2.1, and in section 6.2.2 it is used to




Figure 6.1: Electron spectra for [$L = 15$ mm, $D = 200$ μ m, $p = 250$ mbar]. Figure (a) shows every electron spectrometer image recorded for [$L = 15$ mm, $D = 200$ μ m, $p = 250$ mbar] arranged in order of delay. The colour scale of each image is normalized individually. Figure (b) plots the energy of the peak electron signal against t . Where no electrons were detected the energy is plotted as 0 MeV.

show that the observed generation of stable electrons coincides with an unexpected drop in the ionization level within the waveguide. Reasons why the waveguide plasma should have ceased to be fully ionized during the firing of the discharge are given in section 6.3, and several explanations for the apparent dependence of beam generation on the ionization state of the plasma channel are discussed in 6.4. The principal results of this chapter have been published in Physical Review Letters [82].

6.1 Diagnosing plasma conditions within the waveguide

6.1.1 Interferometry measurements

The plasma profile inside the waveguide can be measured by transverse interferometry. The interferometry experiments were conducted in the Clarendon laboratory in Oxford, re-creating the conditions used on the Astra experiments.

Experimental setup

The experiments were performed with the same apparatus as those described in [45, 64]. Since the waveguide was being probed transversely it was necessary to use a capillary of square cross section — the probe beam would have been refracted by a circular capillary making interferometry impossible. Simulations show that plasma profile inside a square capillary of side D are very similar to those of a cylindrical capillary of diameter D [80]. Two spacer plates were sandwiched between two sapphire blocks, with a square channel left between to form a capillary, as shown in figure 6.2(a).

The layout of the experiment is given in figure 6.2(b). A Nd:YAG laser operating at 1064 nm was used to produce 5 ns laser pulses which were then passed through a doubling crystal to create a probe beam with radiation at both 532 nm and 1064 nm. The square capillary was placed in one arm of a Mach-Zender interferometer through which the probe beam was passed. A beam splitter and band pass filters after the interferometer were used to re-separate the two wavelengths so that interference measurements could be made for both 1064 nm and 532 nm radiation simultaneously. The delay between the firing of the discharge and the arrival of the probe beam was controlled in exactly the same way as the delay t on the Astra experiment. This meant that images could be recorded for different t with 5 ns resolution.

Measuring the evolution of the plasma channel

A reference image was recorded for both colours at negative delays. The phase deduced from these images (by a fourier analysis method outlined in [45]) was subtracted from the phase deduced from images recorded at subsequent delays. The resulting difference in phase, $\Delta\phi(y, z)$, is the change in phase accumulated by the probe beam across the capillary due to the firing of the discharge — this includes phase changes due to the refractive index of the plasma and changes of refractive index in the sapphire walls. The co-ordinate z points along the length of the capillary, and for these experiments $\Delta\phi(y, z)$ was averaged along axis to give $\Delta\phi(y)$.

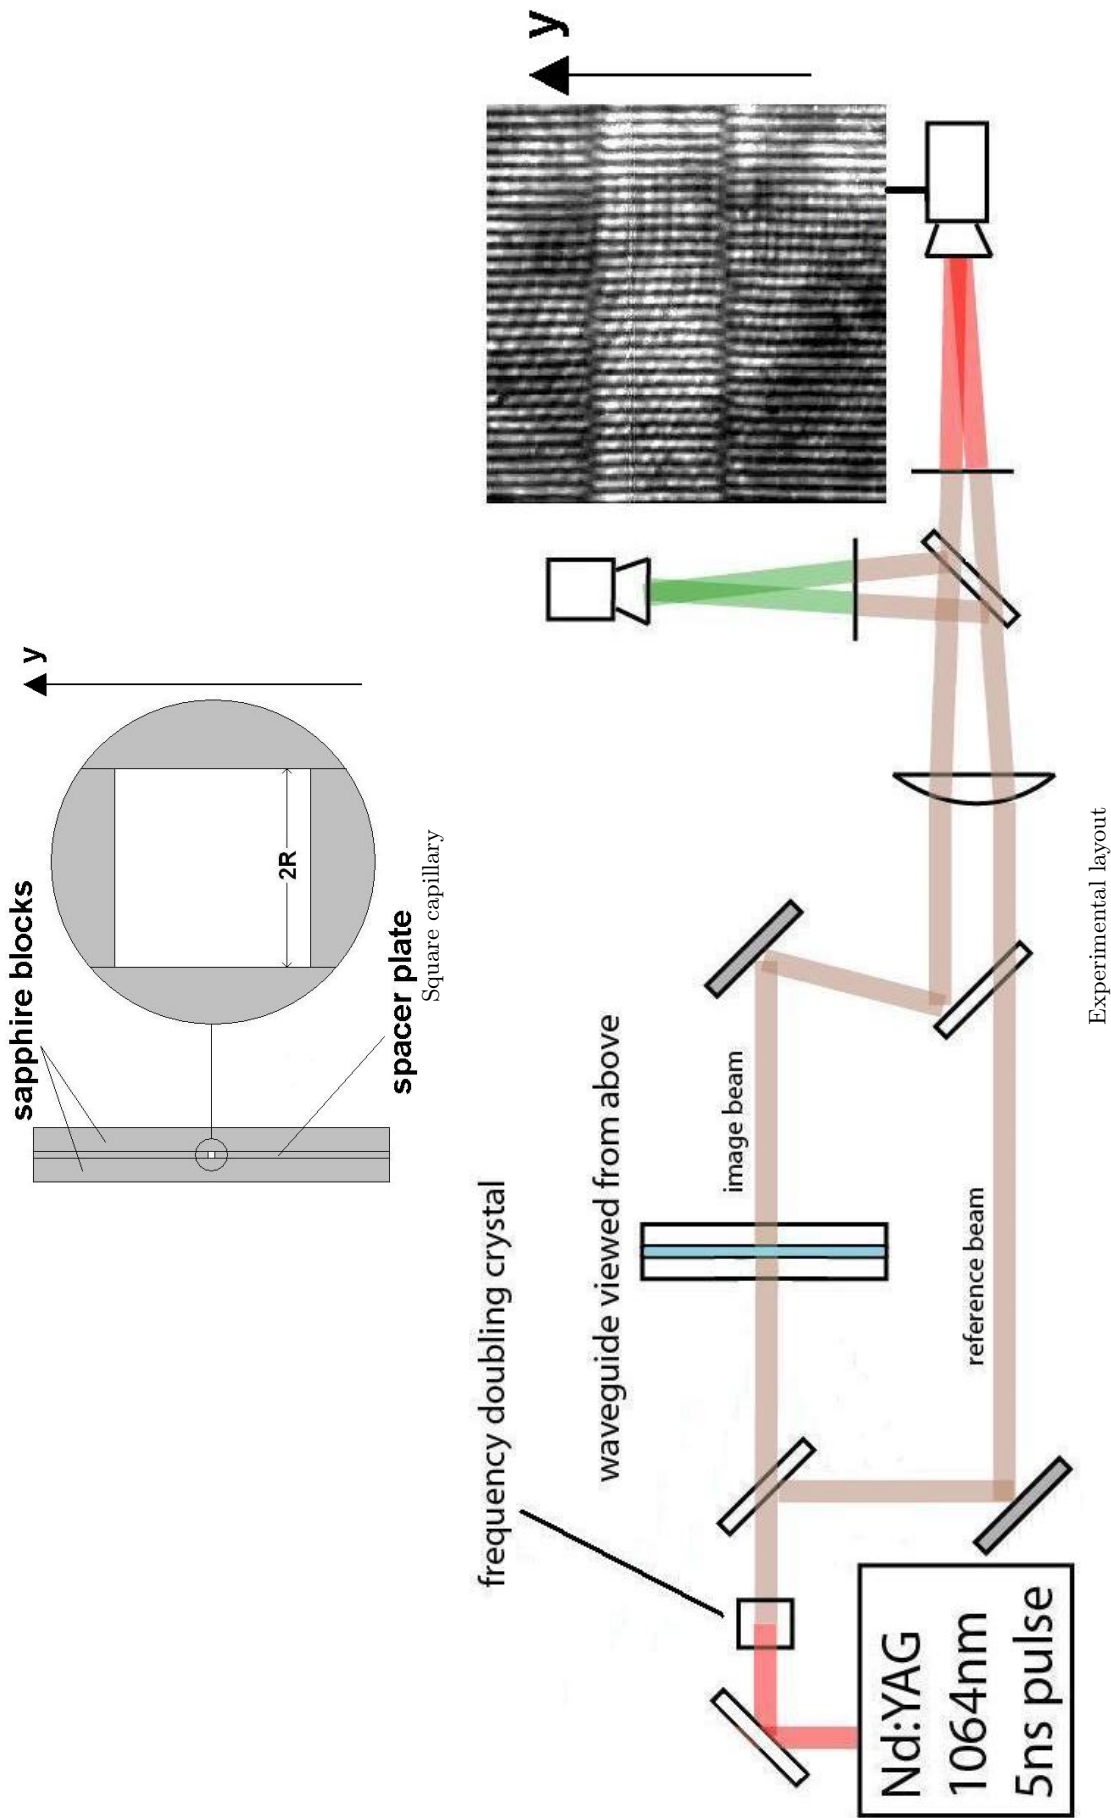


Figure 6.2: Schematic diagram of transverse interferometry experiments. Figure (a) shows the cross-section of a capillary created from sapphire blocks and sapphire spacer plates. Figure (b) shows the experimental layout. Red lines represent 1064 nm radiation, green lines represent 532 nm radiation and brown lines represent a mixture of 532 nm and 1064 nm radiation. Also shown is an example of a recorded image, with the capillary and interference fringes clearly visible.

The motivation for using two wavelengths in the probe beam was to isolate the phase shifts due to the plasma evolution from those caused by temperature changes in the sapphire blocks. In references [45, 64] the following relationship was deduced

$$\Delta\phi_e(y) = \frac{4}{3}\Delta\phi_{1064}(y) - \frac{2}{3}\Delta\phi_{532}(y) \quad (6.1)$$

where $\Delta\phi_e(y)$ is the phase that would have been accumulated by the 1064 nm beam had only the refractive index of the plasma contributed. This can be related to the electron density profile $n(x, y)$ by considering the relationship between the plasma density and its refractive index, giving

$$\int_{-R}^{+R} n(x, y)dx = G(y) = \frac{\Delta\phi_e(y)}{r_e\lambda} \quad (6.2)$$

Simulations of the plasma channel suggest that the density profile within the waveguide can be described approximately by a function of the form $n(x, y) = f(x)f(y)$ [45, 64, 80]. Assuming such a form for $n(x, y)$ it is straightforward to show that

$$n(x, y) = \frac{G(x)G(y)}{\int_{-R}^{+R} G(y)dy} \quad (6.3)$$

In order to represent the evolution of the plasma channel graphically a continuous (in t) density profile was calculated by applying equation (3.12) to the profiles $n(x, 0, t_i)$ across $y = 0$, giving $n(x, 0, t)$. The interferometric measurement of axial density, $n_e^1(t_i)$ is defined as the minimum of each of the measured profiles.

6.1.2 Raman analysis of long pulse spectra

Section 1.3.4 described how the raman instability led to the growth of peaks in the spectrum of a laser pulse as it propagates through a plasma. The frequency of these ‘raman satellites’ is $\omega + n\omega_p$, where n is an integer that can be positive or negative. In chapter 4 it was shown that the $n = -1$ raman satellites were present at low intensity in the transmitted spectrum of the long pulse. Occasionally the $n = -2$ peak was also observed. Because ω_p is a function of n_e , the frequency of the raman satellite can be used to deduce the plasma density within the waveguide.

Examples of this are shown in figure 6.3. For each plot there is a clear peak towards the infra-red end of the spectrum. Also plotted are the expected wavelengths of the $n = -1$ raman satellites based on the results of the interferometry measurements of section 6.1.1. For delays $t < 120$ ns the position of the raman satellites are in agreement with the values predicted from interferometry. At later delays this was not the case, and this is discussed in detail in section 6.2.2.

For each spectrum a value of the density within the capillary n_e^R was deduced from the difference in frequency between the peak of the input spectrum and that of the raman satellite. The error in n_e^R was calculated from width of the peaks at three-quarters of their peak value. The long pulse fluence profiles at the capillary exit are included with each spectrum of figure 6.3. From these it can be seen that within the range of delays for which transmission was high (including the delay at which electrons were generated), the pulse was guided. This means that laser radiation would

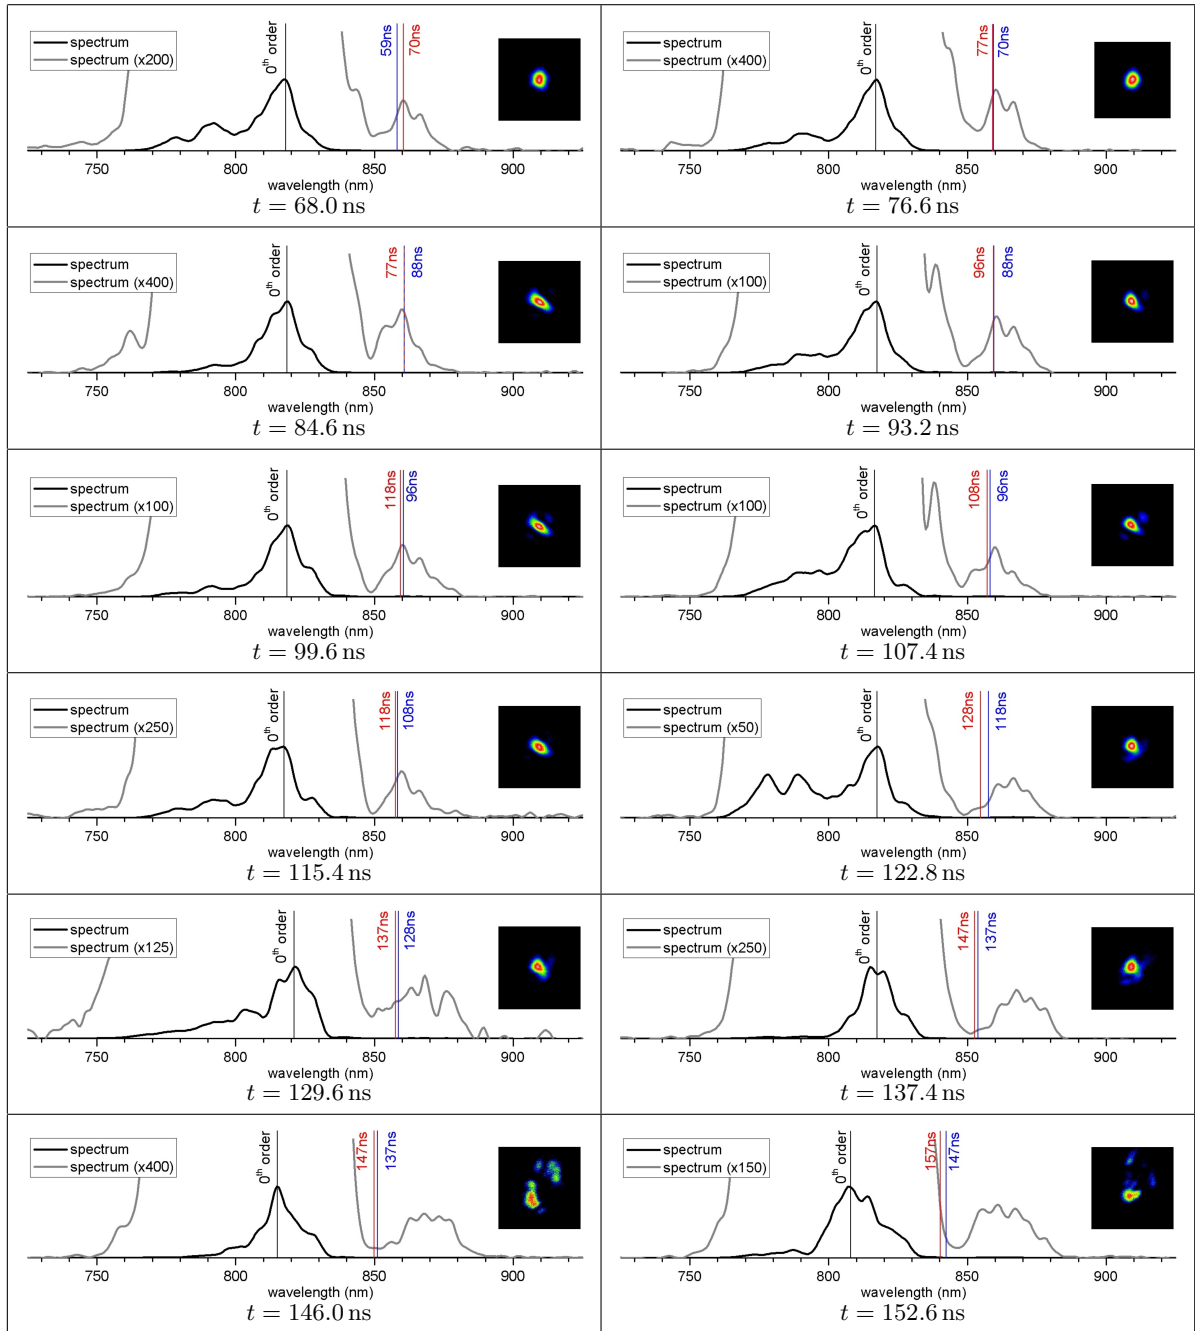


Figure 6.3: Long pulse spectra and fluence profiles for $[L = 15 \text{ mm}, D = 200 \mu\text{m}, p = 250 \text{ mbar}]$. All long pulse spectra recorded for $[L = 15 \text{ mm}, D = 200 \mu\text{m}, p = 250 \text{ mbar}]$ for delays between 65 and 155 ns are plotted (black line). Each spectrum is also plotted multiplied by a factor between 50 and 400 (grey) to emphasise the raman satellite. The cursors represent: the central wavelength (black); the expected wavelengths of the Raman peak calculated from n_e^I measured at the delays closest to that of the spectrum (red and blue). The exit fluence profiles are included with each spectrum to confirm that the laser radiation was confined to the axis as it passed through the waveguide.

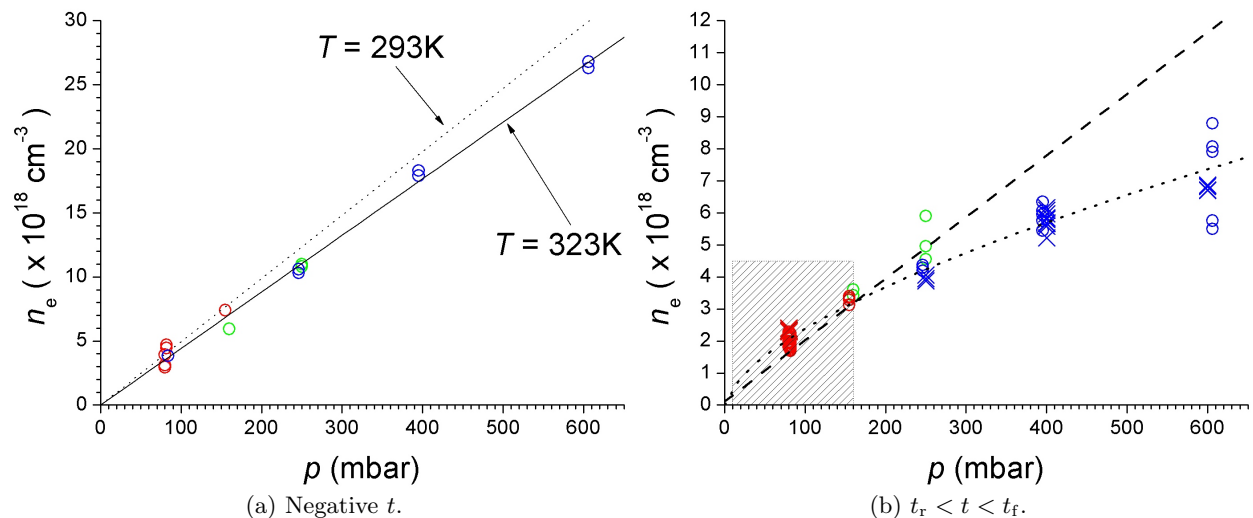


Figure 6.4: Comparison of n_e^I and n_e^R to predictions.

Results are from three capillaries: [$L = 33$ mm, $D = 250$ μm] (green); [$L = 33$ mm, $D = 200$ μm] (red); [$L = 15$ mm, $D = 200$ μm] (blue). **Figure (a)**: n_e^R measured for negative delays. The dotted line represents the value calculated from the ideal gas law assuming $T = 293$ K. The solid line is a fit to the data from which the ideal gas law was used to deduce a temperature of $T = 323$ K. **Figure (b)**: n_e^I (crosses) and n_e^R (dots) are plotted for delays during the high transmission plateaux. The dashed line is the axial density predicted equation (2.13) [64]. The shaded area represents the range within which that law was deduced. The dotted line is a fit of the form $n_e = Ap^\alpha$.

have been confined to the axis as it propagated through the waveguide. For these delays n_e^R is an approximate measurement of the axial density of the plasma channel.

The leading edge of the long pulse is assumed to ionize any gas in the waveguide, meaning that for $t < 0$, n_e^R is a measurement of the initial density of hydrogen atoms. The initial hydrogen density could also be estimated from the measured value of p and the ideal gas law, with $T \approx 293$ K (20°C) assumed. These two independent estimates of the initial hydrogen density can be used to test the reliability of n_e^R as a measurement. This is plotted in figure 6.4(a), where n_e^R for $t < 0$ is plotted against p . As expected the relationship between n_e^R and p is linear. The values of n_e^R are slightly less than predicted by the ideal gas law for $T = 293$ K, and a fit yields a temperature of $T = 323$ K (50°C). This small discrepancy is most likely due to the fact that the capillary was being fired regularly, meaning that it did not cool completely between shots. The results confirm that n_e^R is a valid measurement of the electron density in the waveguide.

6.2 Results of electron density measurements

6.2.1 Comparison to existing scaling laws

Throughout this thesis the scaling laws of (2.13) and (2.14), deduced from interferometry measurements [45, 64], are used to convert from p and D to n_e and W_c . The range of pressures used for those measurements was 10 mbar to 150 mbar. The applicability of these scaling laws at higher pressure is unknown. The pressure range for the Astra experiments was 40 mbar to 600 mbar, and

the investigations of section 6.1 produced new data in the form of n_e^I and n_e^R allowing for a test of the scaling laws for $p > 150$ mbar.

In figure 6.4(b) n_e^R and n_e^I for $t_f < t < t_r$, the range of delays corresponding to high transmission, are plotted against p . It is seen that n_e^R and n_e^I agree within the expected error. The new data suggests that equation (2.13) is valid at pressures below $p \sim 200$ mbar, but at higher pressures the axial density is less than predicted.

The fit to the data in figure 6.4(b) is:

$$n_e [\text{m}^{-3}] = 1.079 \times 10^9 \times (n_{H_2}^i [\text{m}^{-3}])^{0.6303} \quad (6.4)$$

Given the weight of data below 200 mbar in reference [64], equations (2.13) and (2.14) are still the most valid scaling laws for W_c and n_e within that range; however above $p = 200$ mbar equation (6.4) should be applied. Despite these differences, continuing improvements in the available laser power for laser wakefield acceleration suggest that future experiments will be conducted with pressures below 200 mbar, making the need to replace equations (2.13) and (2.14) small (at least for laser wakefield acceleration experiments). Measurements of how the plasma channel scales at lower pressure may be a subject for future work.

6.2.2 On axis ionization levels

Figures 6.5–6.8 combine measurements presented in chapter 4 with the additional data resulting from the work of 6.1 for four different sets of conditions. For all of these graphs comparing panel (a) to panel (c) shows that, as expected, the delays for which the transmission of the waveguide was high were those for which a plasma channel was formed. The rise and fall in transmission coincide with the formation and decay of the plasma channel. The data in panel (b) also illustrates that the drop in transmission and decay of the plasma channel corresponds to the onset of blue-shifting of the transmitted laser spectra. The blue-shift was stronger for the short pulse than the long pulse, and appears to have been more dramatic at low pressure.

From panel (d) it can be seen that for many delays the two methods of measuring n_e were not in agreement. This discrepancy is also seen in figure 6.3 for late delays. Since the leading edge of the long pulse laser would have ionized any neutral or partially ionized species that were encountered, n_e^R is a measurement of the density of electrons ionized by the discharge plus the density of electrons ionized by the laser. The intensity of the probe beam in interferometry was sufficiently low that it would not have ionized any species within the waveguide, meaning that n_e^I corresponds to the density of electrons ionized by the discharge only. This means that comparison of n_e^R and n_e^I can be used to determine the level of ionization of the waveguide on axis before the arrival of a laser pulse.

Figures 6.5–6.7 all follow the same pattern. Where the transmission was high, $n_e^R \approx n_e^I$ suggesting that at those delays the plasma channel was fully ionized on axis. At earlier delays $n_e^R > n_e^I$, so the plasma was not fully ionized. That is unsurprising given that some time was required before all of the initial gas fill was ionized. As the transmission falls away and the spectrum starts to blue shift for $t > t_f$, the values of n_e^R and n_e^I are seen to diverge, with a decrease in n_e^I and an increase in n_e^R . That suggests that the ionization level on axis began to drop at those delays. There is no data

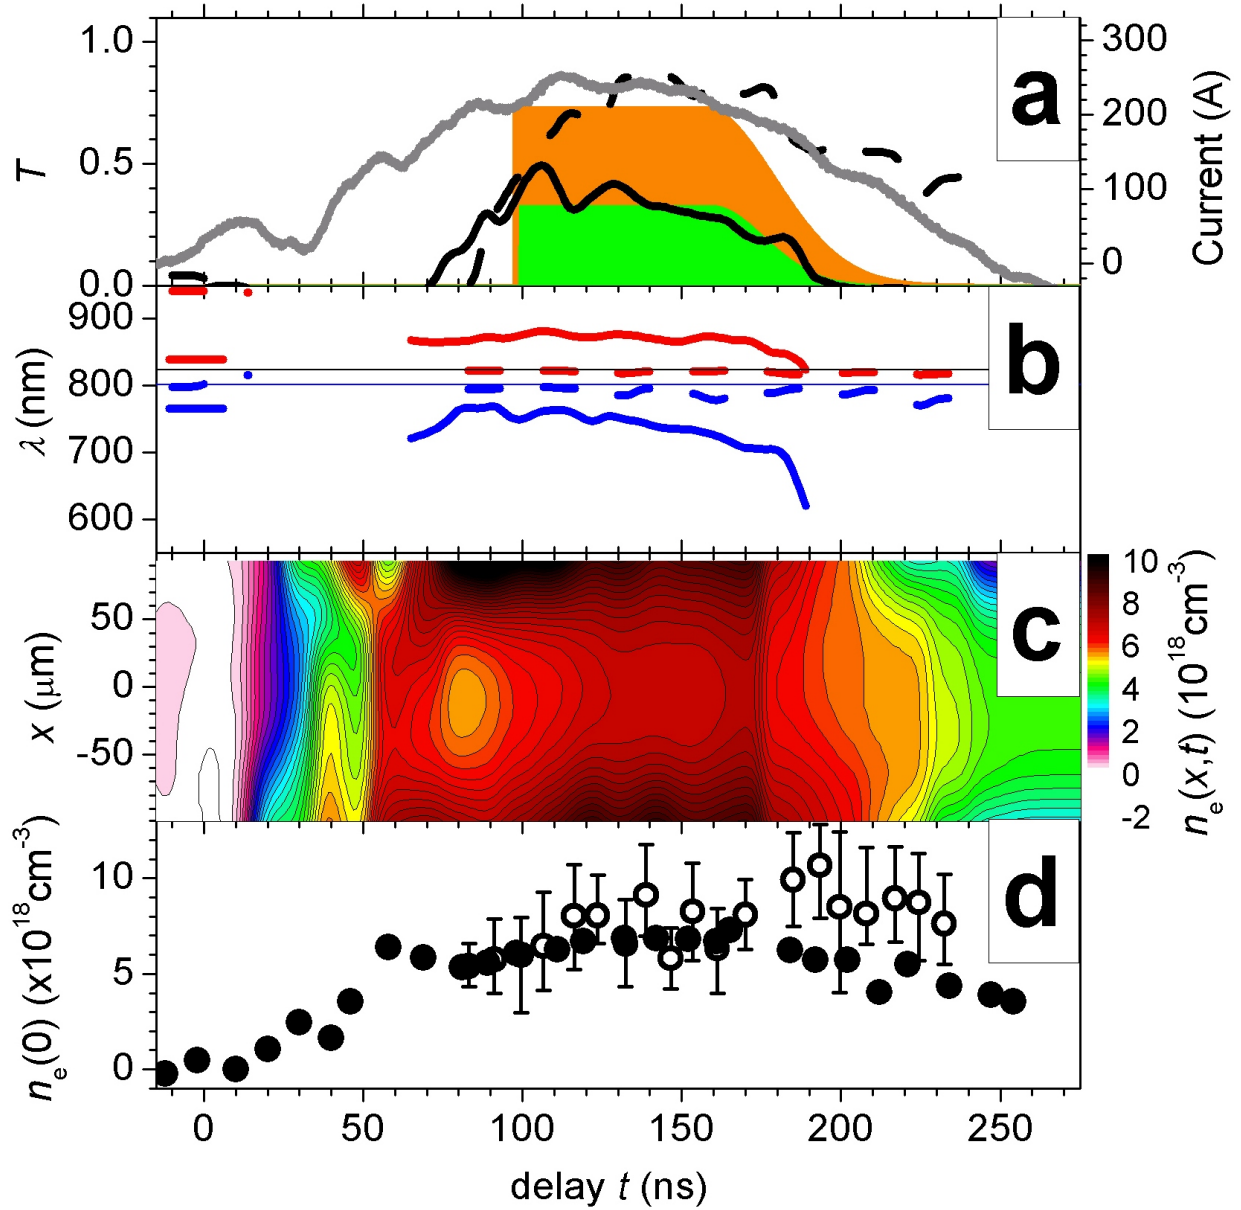


Figure 6.5: Measured data as a function of t for [$L = 15$ mm, $D = 200 \mu\text{m}$, $p = 600$ mbar]. Panel (a): against the left hand axis are long/short pulse transmission measurement $T_{av}(t)$ (dotted/solid black lines) and probabilities $P_{beam}(t)$ (green fill) and $P_{electron}(t)$ (orange fill); against the right hand axis the current (grey line). Panel (b): the red and blue limits of the transmitted laser spectrum $\Lambda^R(t)$ (red) and $\Lambda^B(t)$ (blue) are plotted for the long/short pulse (dotted/solid). Panel (c): the density profile $n(x, 0, t_i)$ measured by interferometry. Panel (d): density measurements n_e^I (solid) and n_e^R (hoops).




Figure 6.6: Measured data as a function of t for [$L = 15$ mm, $D = 200$ μm , $p = 400$ mbar]. This figure is equivalent to figure 6.7.

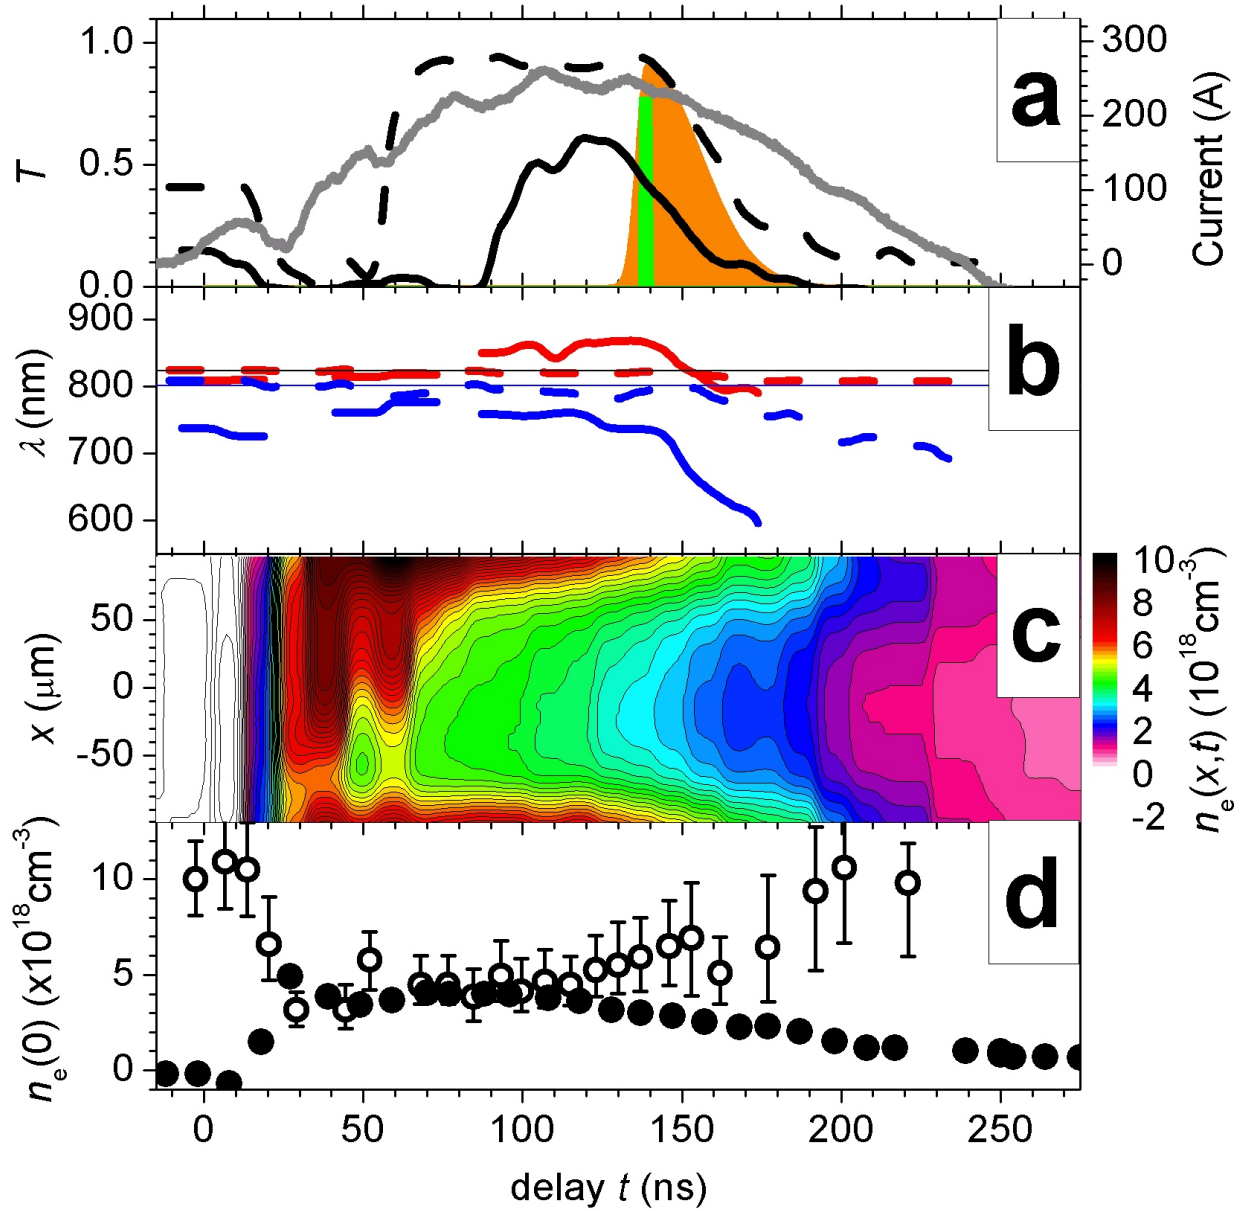


Figure 6.7: Measured data as a function of t for [$L = 15$ mm, $D = 200 \mu\text{m}$, $p = 250$ mbar]. This figure is equivalent to figure 6.7.

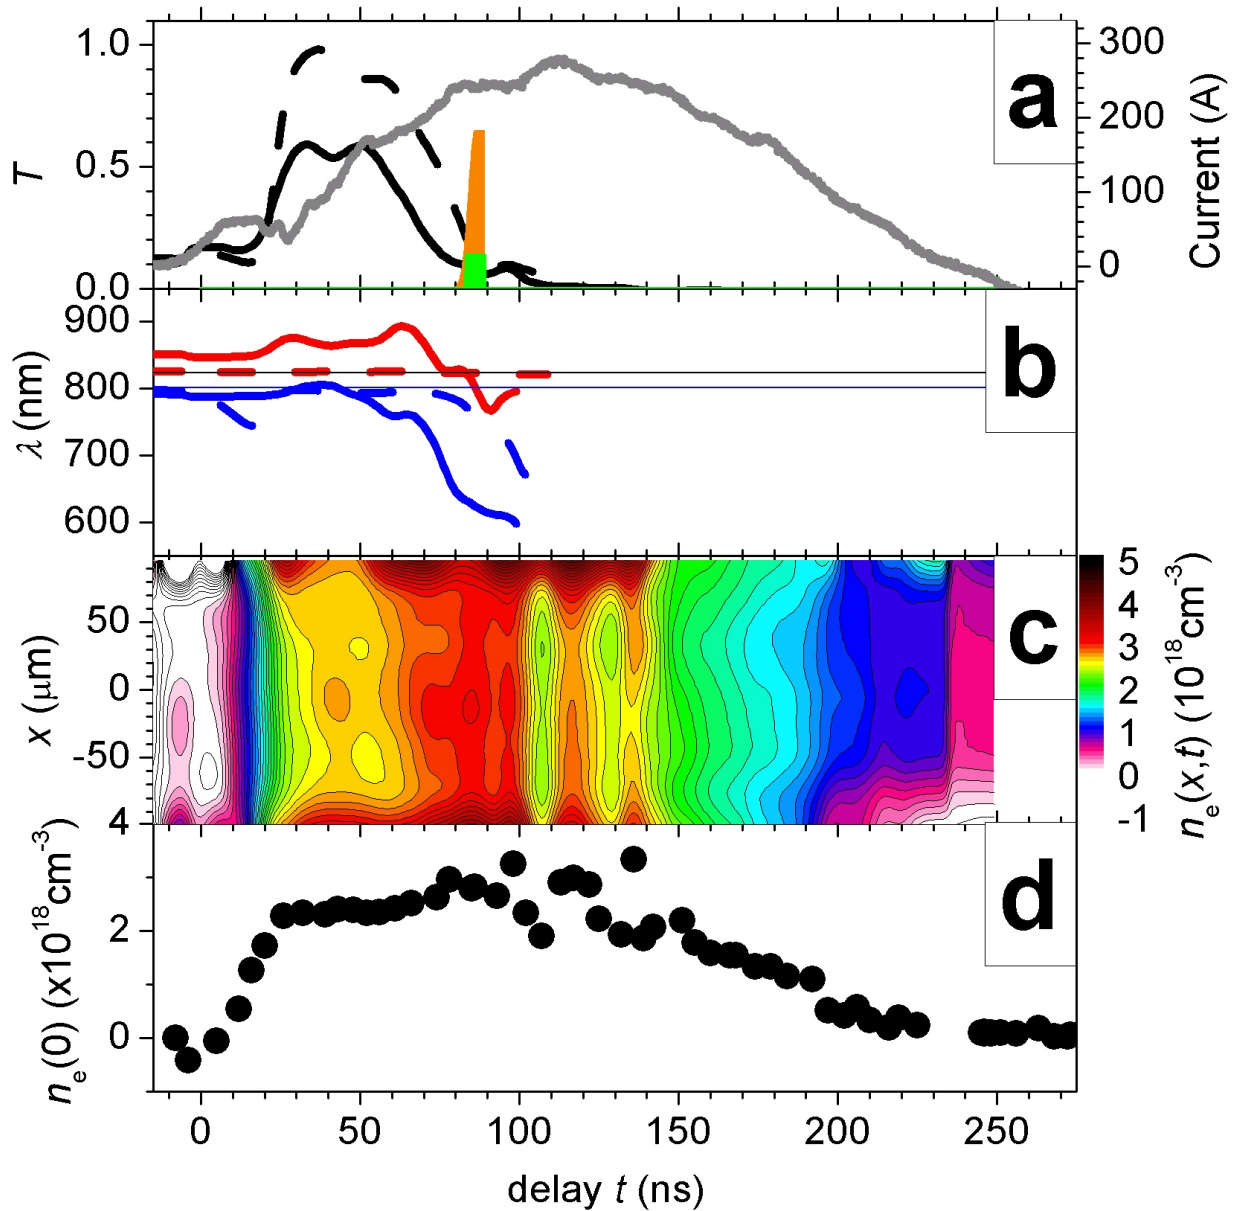


Figure 6.8: Measured data as a function of t for [$L = 33$ mm, $D = 200 \mu\text{m}$, $p = 80$ mbar]. This figure is equivalent to figure 6.7, except that there was no data for n_e^R due to the fact that the raman satellites could not be resolved from the central peak of the transmitted long pulse spectrum.

for n_e^R in figure 6.8 because the density was too low for the raman satellite to be resolved from the central peak. However, n_e^I does decrease at $t \approx t_f$, suggesting that the ionization level also dropped under those conditions.

In figures 6.5 and 6.6, which represent behaviour for $n_e > 6.5 \times 10^{18} \text{ cm}^{-3}$, electrons were accelerated across the entire range of delays for which the transmission was high. In figures 6.7 and 6.8, representing behaviour for $n_e < 6.5 \times 10^{18} \text{ cm}^{-3}$, electrons were only generated at the delay for which the degree of ionization in the channel began to decrease. It was established in chapters 4 and 5 that a plasma channel was required for efficient wake excitation when $n_e \lesssim 20 \times 10^{18} \text{ cm}^{-3}$. It would appear that for $n_e \lesssim 6 \times 10^{18} \text{ cm}^{-3}$ there was an additional requirement for electron acceleration: that the plasma channel was only partially ionized on axis. Since the drop in the degree of ionization coincides with the decay of the plasma channel both these conditions are satisfied for a very narrow range of delays only. That gives a likely explanation for the observed sensitivity of beam generation to t .

6.3 Un-ionized Species

It is unexpected that the ionization levels in the waveguide should have dropped when the discharge current was near its peak. In this section possible explanations for the effect are discussed. There are two likely sources of unionized species: recombined hydrogen, and ablated aluminium oxide from the sapphire capillary walls. In reference [82] neither possibility was ruled out. Whilst there has been some debate regarding the origin of un-ionized species on axis, it is firmly the opinion of the author that the effects observed are due to recombination of the axial plasma to form hydrogen. The case for the aluminium oxide ‘wall hypothesis’ shall be considered in section 6.3.1, before arguments for the ‘recombination hypothesis’ are put forwards in sections 6.3.2 and 6.3.3.

6.3.1 Wall hypothesis

Unsurprisingly, theory suggests that the plasma temperature inside the waveguide is higher when the discharge current is higher [60]. The data in figures 6.5–6.8 shows that the delay at which the ionization level dropped was near the peak of the current. That would appear to rule out the possibility that the plasma recombined to form hydrogen at those delays. This is a very strong argument that leaves one possibility for the source of the unionized axial species, and that is material ablated from the wall.

The chemical formula for sapphire is Al_2O_3 . The thermal velocity of an Al ion in a 5 eV plasma is 6 km s^{-1} . That corresponds to 0.8 mm in 140 ns, meaning that if ions were ablated from the capillary wall it would be reasonable to expect that some would have traveled to the waveguide axis during the discharge. The same also applies to the lighter (and hence faster) O ion.

The presence of aluminium and oxygen ions in the channel would have caused the drop-off in transmission: on-axis where the laser was most intense they would have produced more electrons than at lower intensities further from the axis. That would have negated the refractive index profile of the plasma channel, and, when concentrations of these impurities were high enough, inverted the channel to de-focus the laser.

6.3.2 Recombination hypothesis

Arguments against the wall hypothesis

Despite the above arguments for the wall hypothesis, there are a number of observations that independently suggest that the observed drop in the ionization level within the waveguide cannot be explained in terms of the ablation of material from the capillary wall.

When considered carefully the evidence of figures 6.5–6.8 is not consistent with the wall hypothesis. Firstly, the drop in ionization level is characterised by both an increase in n_e^R and a decrease n_e^I . An influx of wall material to the axis would have led to an increase of n_e^R , but it should not have led to a reduction in the number of electrons measured by interferometry. Some aluminium and oxygen would be partially ionized by the discharge so if anything n_e^I would have increased at these delays if the wall hypothesis was correct.

In defending the wall hypothesis it has been suggested that n_e^I dropped due to the ejection of plasma from the ends of the waveguide. However there is no physical reason that the drop in n_e^I should have occurred at the same delay as the increase in n_e^R . If the above argument was true the fact that the two effects were simultaneous could only be attributed to coincidence. Moreover, the velocity at which plasma is ejected into the gas-slots has been measured as 12.2 km s^{-1} [45], which is approximately 1 mm in 100 ns. This should be similar to the velocity at which plasma is ejected from the waveguide ends suggesting that n_e^I , measured in the middle of the capillary, should not have been affected by end effects for the timescales in question. Moreover, the drop in n_e^I occurs earliest in figure 6.8, which represents the longest capillary. If end effects were responsible for the drop in n_e^I then that would have been the capillary for which the effect occurred last.

An alternative explanation for the observed drop in n_e^I that has been discussed in defence of the wall hypothesis is that ions ablated from the wall could have altered the refractive index of the plasma. That effect has not been estimated quantitatively, but the argument can be dismissed by observing that in figures 6.7 and 6.8 the value of n_e^I at late delays goes to zero. The same also occurred for the data of figures 6.5 and 6.6, but at later delays than are plotted on the graphs. It would be an extremely unlikely coincidence that for all of these pressures the volume of material ablated from the capillary wall was enough to exactly cancel the refractive index due to plasma electrons.

A final argument against the wall hypothesis relates to the volume of Al_2O_3 that would have to have been ablated from the wall by the discharge in order to produce the observed increase in n_e^R . The barrier ionization thresholds for certain elements of interest are listed in appendix C. From these we can infer that the number of electrons contributed by each ablated Al_2O_3 molecule after ionization by the long pulse, would have been ~ 30 . The molecular mass and density of sapphire are 102 g mol^{-1} and 3.98 g cm^{-3} respectively. At late delays $n_e^R \approx 10^{19} \text{ cm}^{-3}$ and $n_e^I \approx 0$, meaning that, according to the wall hypothesis, there must have been a density $\sim 3 \times 10^{17} \text{ cm}^{-3}$ of Al_2O_3 molecules in the capillary¹. To fill the capillary with that density would have required the ablation of 1 nm of wall material for every shot². The ablation of the wall by the discharge was measured

¹Or of the constituent parts of that number of molecules

²This figure was calculated assuming that none of the wall material would be ejected during the discharge. Were that effect included, even more ablation would be necessary to replace the ejected hydrogen plasma.




Figure 6.9: Schematic diagrams of ‘clean’ and ‘unclean’ discharges.

Grey represents the sapphire into which the capillary and gas slots are machined, orange represents the electrodes, and blue represents the observed path of the discharge current.

in [45] by firing the discharge, with a peak current of 300 A, 1.3 million times in a $200\ \mu\text{m}$ square capillary. The depth of the ablation was $20\ \mu\text{m}$, which corresponds to a widening of $0.015\ \text{nm}$ per shot. The wall hypothesis is therefore inconsistent with that experimental measurement by two orders of magnitude.

Arguments for the recombination hypothesis

Almost all of the data in figures 6.5–6.8 supports the supposition that current ceases flowing through the waveguide at $t \approx t_f$. The increase in n_e^R and drop in T are consistent with the flattening of the guiding plasma channel. The return of n_e^R to its original value at late delays seen in figure 6.7 can be seen as conservation of particles³. The only evidence that is not consistent with this supposition is, as pointed out in section 6.3.1, the measurement of the current itself. A possible explanation for this is given in section 6.3.3.

6.3.3 Gas-slot discharge model

Discharge switch for $t_f \rightarrow 0$

The delay $t = t_f$ appears to mark the transition between two distinct phases of behaviour within the waveguide. For the earlier phase ($t < t_f$) the behaviour is exactly as expected — the current flows and a plasma channel forms. For the later phase ($t > t_f$) the behaviour is not as expected, because the plasma appears to recombine near the peak of the measured current. Isolating that later phase is therefore the key to understanding the observed effects.

It was observed in chapter 4 that t_f was less for longer, narrower capillaries. Figure 4.3 summarises this trend for $p = 160\ \text{mbar}$: t_f is plotted against L_c/A , where L_c is the length of capillary between the gas slots (normally $L_c = L - 4\ \text{mm}$) and A is the cross-sectional area of the capillary. The reason for this choice of x -axis shall become clear later in the discussion. As L_c/A increases, t_f decreases. The points all conform very closely to a linear fit.

For $t_f \rightarrow 0$ the latter phase is isolated because for all positive delays $t > t_f$. Should the linear trend of figure 4.3 continue it would be found that for $L_c/A = 1.5 \times 10^6\ \text{m}^{-1}$, $t_f \approx 0\ \text{ns}$. This

³It having been established that the bulk of the capillary would be unaffected by ejection of the plasma from the capillary ends

is true for a [$L = 50$ mm, $D = 200$ μm] capillary. Although no such capillaries were used during the Astra experiments, in more recent experiments at the upgraded Astra-Gemini laser facility an [$L = 50$ mm, $D = 200$ μm] capillary was employed. It was found that for that capillary most shots the discharge would not fire ‘cleanly’. Generally the clean discharge current is visible as a bright spark inside the capillary, as depicted in figure 6.9(a). For [$L = 50$ mm, $D = 200$ μm] there was no spark observed in the capillary. Instead the light of the discharge, although fainter, could be seen in the gas slots, as depicted in figure 6.9(b). In this case the discharge current took an alternative route between the two electrodes: through the capillary before the first gas slot, through the gas slot, through the gas reservoir to the second gas slot, through the second gas slot to the capillary and through the final section of the capillary to the electrode. It was observed that when this capillary was replaced with an [$L = 50$ mm, $D = 300$ μm] capillary a bright discharge was observed to flow through the capillary once more.

Taken together the evidence of figures 4.3 and 6.9 suggests that although current flowed through the capillary for the earlier ($t < t_f$) phase, at $t = t_f$ it switched to the gas-slot path. The fact that the discharge could take both routes during a single shot would not have been visible by eye because the discharge light in the gas slots was much fainter than that in the capillary.

Gas-slot break down

It is not obvious that the current could have changed path once it has begun to flow through the capillary. The resistance of the capillary path, $R_c(t)$ would have changed from $R_c(0)$, the resistance of the capillary filled with hydrogen, to R_c^* , the resistance of the ionized capillary. Whilst the exact value of R_c^* may have varied slightly as the plasma channel evolved, it would have been that $R_c^* \ll R_c(0)$. Since the current flowed through the capillary initially it can be assumed that $R_c(0) < R_g(0)$, where $R_g(0)$ is the initial resistance of the gas slot path. It follows that $R_c^* \ll R_g(0)$. Therefore one would expect that once current began to flow through the capillary, there would have been no cause for a breakdown of the plasma in the alternative path.

However, comparison with an earlier design of the hydrogen-filled capillary discharge waveguide shows that such a breakdown was to be expected. The waveguide used for the Astra experiments was designed by A. J. Gonsalves with electrodes situated at either end of the capillary: one at high voltage, the other earthed. An earlier version, designed by A. Butler, is shown in figure 6.10. That used three electrodes [65]: two earthed electrodes at either end of the capillary, and a high voltage electrode in the center. Whilst the configuration of electrodes in the two waveguide designs was different, electrically they are equivalent: in the Gonsalves waveguide there are two identifiable paths between the high voltage electrode and Earth (the capillary path and the gas-slot path); in the Butler waveguide there are also two paths between the charged electrode and Earth (the left-hand capillary path and the right-hand capillary path).

Using the same notation for the resistances of the left and right hand paths in the Butler waveguide it can be said that $R_L^* \ll R_R(0)$. If the left hand side were to break down first⁴, a naive argument would predict that no current should flow through the right hand side because its resistance was so much higher. Therefore the right hand side would never break down. However,

⁴All the same arguments can be applied irrespective of which half breaks down first.




Figure 6.10: Diagram of the Butler waveguide. After [1].

experiments showed that on a timescale of 100 ns after the breakdown of one path within the Butler waveguide, breakdown would occur in the other half [65]. It is possible that plasma ejected from the side that broke down first seeded the break down of the other side, and it has been shown that in the Gonsalves waveguide plasma is ejected into the gas slots during the discharge. Whatever the mechanism responsible, the fact that both paths eventually broke down in the Butler waveguide shows that both paths of the Gonsalves waveguide could also have broken down.

One can say that $R_c^* \approx \rho_p L_c/A$, where ρ_p is the resistivity of the plasma. As the gas slot broke down its resistance began to drop from $R_g(0)$ to R_g^* . Since the current eventually switched to the gas slot path we can say that $R_g^* < R_c^*$. It follows that t_f was the delay where $R_g(t_f) = R_c^* \approx \rho_p L_c/A$. This means that t_f was a function of L_c/A , which is the reason for its use as the x -axis in figure 4.3. For two distinctly different capillaries, [$L = 33\text{mm}$, $D = 200\ \mu\text{m}$] and [$L = 50\text{mm}$, $D = 250\ \mu\text{m}$] the values L_c/A were similar. The fact for both of those capillaries the value of t_f were also similar strongly supports the argument that t_f was a function of the resistance between the capillary gas slots.

The above arguments by no means prove that the current switched path during the discharge, but there is enough evidence to suggest that it might be the case. Most importantly, it is the only explanation (so far put forward) for the observed drop in ionization of figure 6.7 that is consistent with all of the available evidence. Further investigation is required to fully establish whether or not the above description is correct, and that would allow amendments to the waveguide design to better control the path taken by the current.

6.4 Effect of un-ionized species on laser wakefield acceleration

It has been established that the ionization level on the waveguide axis drops during the firing of the discharge, and a probable explanation has been given for that effect. At the time of writing there are several explanations as to why the drop in the ionization level appeared to facilitate the trapping of electrons in the waveguide. These are outlined below, and future work that might help eliminate certain possibilities is discussed. It has also been suggested that the sensitivity to delay is not connected to the ionization state of the plasma channel but is simply related to changes in density within the plasma channel. Indeed in figure 6.7 the value of n_e^R is higher at the critical delay than during the plateau of high transmission, and it is well known that electron trapping occurs more easily at high plasma densities. However the same sensitivity to delay can be seen in figure 6.8, for which $p = 80\text{ mbar}$. The mean density is less than a third in that case, ruling out the possibility that the sensitivity to delay was due to variations in the axial density within the waveguide.

6.4.1 Electrons born within the wakefield

Experiments using particle driven wakefield accelerators have shown that the trapping threshold is lower for a partially ionized plasma [83]. This was attributed to the fact that electrons ionized within the wakefield are more easily trapped. If plasma electrons moving in the wake reach a forward velocity approaching the phase velocity of the wakefield, they are trapped and accelerated.

However, before being reached by the accelerating phase of the wakefield, electrons from the plasma background gain negative momentum from the initial decelerating phase. For these electrons to be trapped the accelerating gradient must be strong enough to first cancel this negative momentum and then accelerate the electrons to the wakefield phase velocity. Electrons ionized within the wakefield, whose initial momentum is effectively zero, experience less (or none) of the decelerating phase. This means that the accelerating gradients do not need to be as strong to trap these electrons.

This principle is shown to apply at the density under consideration. Figure 6.11 plots the path of axial electrons born at different phases (ξ) of a wakefield. The wakefield is that calculated in section 5.3 for $n_e = 5 \times 10^{18} \text{ cm}^{-3}$. Relativistic effects were included. If it is assumed that the plasma is stationary before the arrival of the laser pulse then the path of the electron born at $\xi = 0$ represents the path of the background plasma electrons as the wakefield passes. The path of that electron is affected by the accelerating field, but not enough to prevent it from exiting the first accelerating phase. Hence it is not trapped. However electrons born at later ξ , which experience only some or none of the decelerating field, are trapped.

In figure 6.11 it is seen that the phase at which an electron is born determines the phase at which it is trapped. That in turn affects its maximum possible energy gain. Electrons that are most easily trapped (that is the ones that did not reach the back of the accelerating phase) had a lower maximum energy, of about 125 MeV. It is interesting to compare this to figure 5.3, where is seen that 125 MeV was the energy reached by electrons in the Astra experiments for $n_e \approx 5 \times 10^{18} \text{ cm}^{-3}$.

It is simple to argue that this principle would apply for particle driven accelerators. However it is more complicated for the LWFA because one would expect the front of the laser to ionize any un-ionized species, so that none survived to the phase of the wakefield that was favourable to electron trapping. Referring again to the ionization thresholds of appendix C it is clear why this model of ionization induced electron trapping can be more easily reconciled to the wall hypothesis of section 6.3.1: there are many states of aluminium and oxygen that would be ionized above and around the peak intensity of the pulse. As demonstrated in [69] electrons ionized at the barrier threshold for high charge states of aluminium and oxygen would be trapped. Hydrogen, on the other hand, has a barrier ionization threshold four or five orders of magnitude lower than the laser intensity, suggesting that neutral atoms would be ionized by the pulse front far away from the critical phase.

This does not change the fact that the experimental evidence is not consistent with the presence of aluminium and oxygen atoms on axis. We note that the required density of species donating electrons is low: if injection occurs in a volume at least equal to the plasma wavelength cubed, a 100 pC electron bunch corresponds to a donor density of less than 3% of the axial plasma density. This percentage is considerably lower if the injection process occurs over many λ_p . Therefore even if the majority of hydrogen is ionized by the pulse front a statistical approach may show that a sufficient fraction of hydrogen species may survive to the critical phase of the wakefield.

This possibility is more likely when it is considered that the ionization rates of hydrogen estimated using a BSI-model are, in probability, an overestimate. The theoretical work of [84, 85] predicts that in fields much higher than the barrier ionization threshold the rate at which atoms are ionized is lowered considerably by quantum interference effects that conventional ionization rates

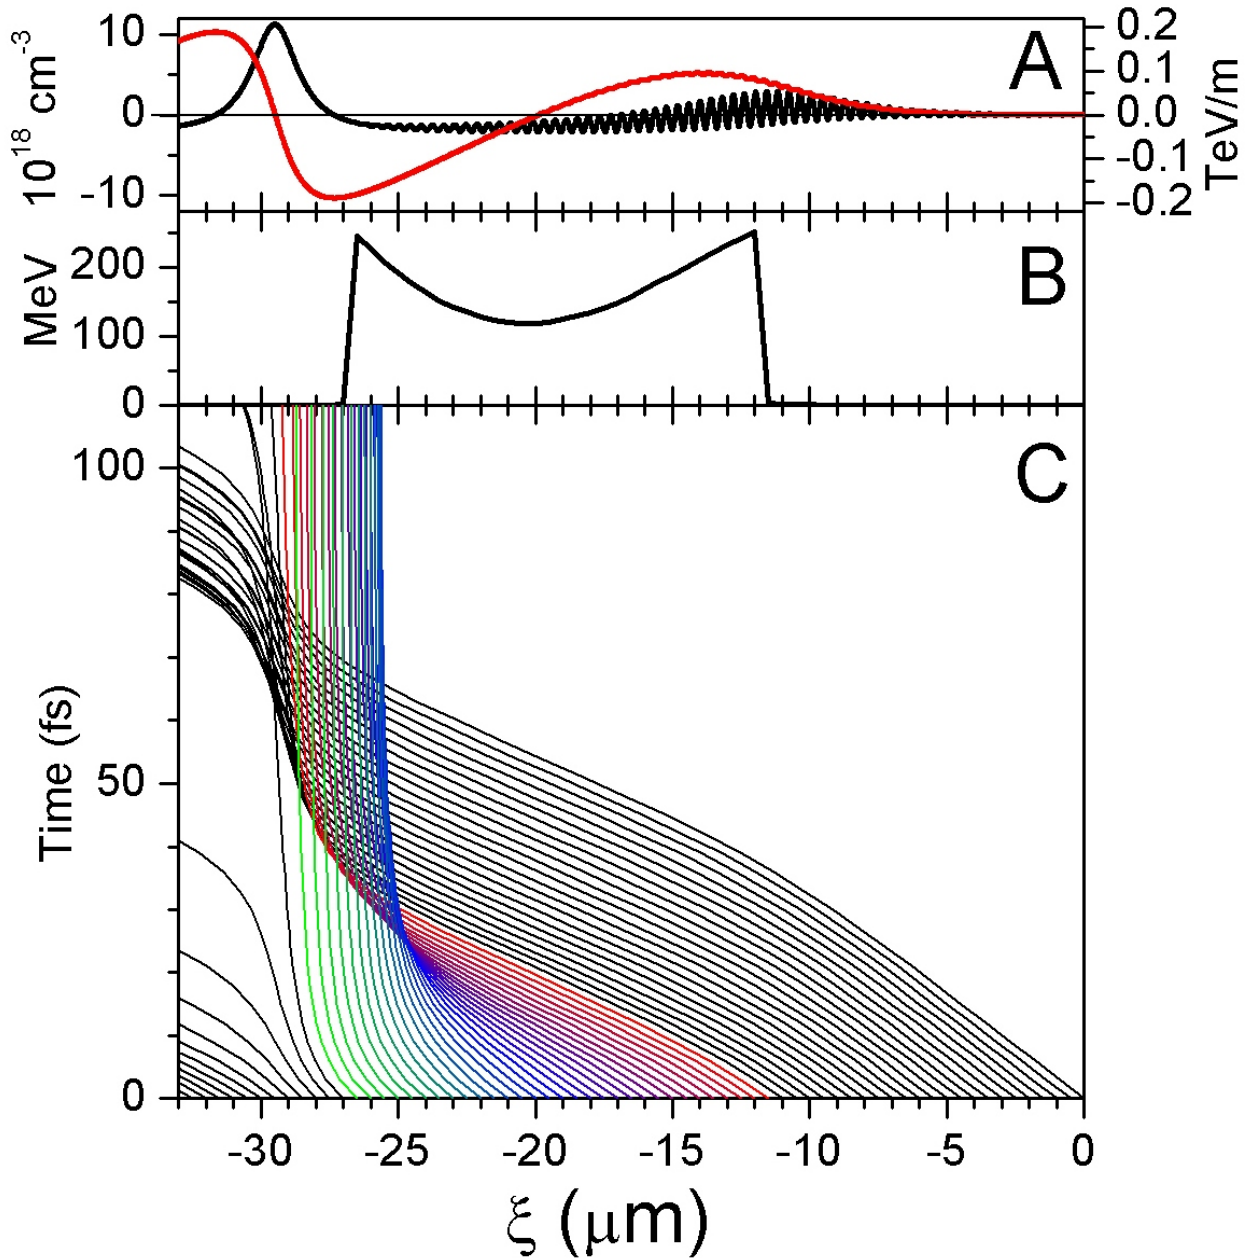


Figure 6.11: Trajectories of axial electrons ionized for varying ξ .

Panel (a): the density perturbation (black line) and electric field (red line) are plotted against ξ . This is the wakefield calculated in figure 5.2 for $n_e = 5 \times 10^{18} \text{ cm}^{-3}$. Panel (b): the maximum energy achieved by electrons against the ξ at which they were born (not trapped). Panel (c): the paths of electrons starting with zero momentum from different points in the wakefield is plotted in terms of the co-ordinates (ξ, t) . The paths of electrons that are eventually trapped are marked in colour to allow comparison between their initial phase and the phase at which they are trapped.

do not describe. Unfortunately the exact conditions of the Astra experiment were not simulated in that work, but [84] predicts that 20% of Hydrogen atoms would survive to the back end of a 10 fs, 800 nm pulse of peak intensity $10^{16} \text{ W cm}^{-3}$. Whilst the intensity is less than would have been present on axis during the Astra experiment, this certainly demonstrates that significant numbers of hydrogen atoms can survive in the fringes of the pulse.

Another factor affecting the ionization rate is that the BSI-model is only applicable to monatomic species. For a diatomic molecule the ionization rate is a function of the orientation of the molecular axis relative to the laser field [86], and where they are perpendicular the rate is suppressed considerably [87]. The picture is complicated further by the presence of H_2^+ and H_3^+ ions in a partially ionized hydrogen plasma [88]. Whilst it might be reasonable to expect the dominant species in a partially ionized plasma to be electrons, protons and hydrogen molecules, the dominant cation is in fact H_3^+ [89]. Therefore as the plasma recombined there could be a brief phase dominated by molecular hydrogen ions, the ionization rates of which would be quite different from those of atomic hydrogen.

More work is required to establish whether or not this model can explain the observed correlation between acceleration of electrons and recombination of the plasma. The most important aspect of that investigation is to establish realistic ionization rates for H , H_2 , H_2^+ and H_3^+ and to determine the relative populations of these species as the plasma recombines. It may be that the observed effect is due to only one of these species that exists briefly as an intermediate stage during the recombination of the plasma: that would certainly explain why electron acceleration was so extremely sensitive to delay under these conditions.

Irrespective of whether this model is the true explanation for the observations of figure 6.7, it suggests that the trapping threshold of LWFA in the waveguide can be lowered considerably if the hydrogen gas is doped with another gas to donate atoms at the critical phase of the wakefield. For practical reasons a noble gas would be ideal. The barrier thresholds of appendix C suggest that argon would be more suitable than neon because more states have a threshold near the typical peak laser intensity.

6.4.2 Increased wakefield amplitude due to an ionization front

Another mechanism that might have enhanced electron acceleration in a partially ionized plasma is the increased ponderomotive force associated with an ionization front. In reference [90] the following formulae were derived for the ponderomotive forces F due to the laser profile and an ionization front

$$n_e F_{\text{las}} \propto -h \frac{\partial \langle a^2 \rangle}{\partial z} \quad (6.5)$$

$$n_e F_{\text{ion}} \propto -2 \langle a^2 \rangle \frac{\partial h}{\partial z} \quad (6.6)$$

These formulae are expressed in terms of the notation used in Chapter 5. The second formula averages the term $\sin^2 \phi$ describing the phase of the laser at which electrons are born to $1/2$.

Since the additional ponderomotive force depends on the intensity of the laser at the ionization front this effect will not be significant for the initial laser profile — the ionization front would

overlap with the front of the pulse where the intensity is low. However if the front of the laser pulse was steepened sufficiently by the wake, the ionization front could overlap with an intense part of the pulse. In the most extreme case the normalised intensity at the front of the pulse is described by a step function from 0 to a_0^2 at $z = 0$. With no ionization front the force at $z = 0$ is proportional to $h_0 a_0^2/4$. If there is an ionization front the density goes from 0 to n_e and h from 0 to h_0 at $z = 0$. Their values at $z = 0$ are $\frac{1}{2}n_e$ and $\frac{1}{2}h$ respectively. This would give a total force F proportional to

$$F \propto \underbrace{h_0 a_0^2/4}_{\text{Force from laser front}} + \underbrace{2h_0 a_0^2/4}_{\text{Force from ionization front}} \quad (6.7)$$

In this extreme case the force at the leading edge of the pulse is increased by a factor of 3. Whilst this is significant, it is difficult to argue that the pulse can become as steep as that described above. In recent 2-d simulations by R. Trines it was seen that shortly after a laser pulse had entered a plasma channel the ionization level made a small difference to the propagation of the laser but as the laser pulse evolved it became insignificant [91], which is the opposite of what would be expected if the pulse front were to steepen. In light of that simulation, it is likely that the ponderomotive force due to the ionization front is not responsible for the enhanced generation of electrons in the plasma channel.

6.4.3 Thermal effects

Electrons ionized by the laser or its pre-pulse have a different temperature to electrons ionized by the discharge. Pre-ionized electrons oscillate in the laser field in phase with one another. The momentum distribution of these electrons oscillates from side to side but its width (which defines the temperature) does not change. Electrons ionized by the laser will not oscillate in phase with the rest of the plasma. Their phase is determined by the phase of the laser oscillation at which they were born, and that could be assumed to be random. Therefore the width of the momentum distribution is affected by ionized electrons, and once the electrons thermalise that can be seen as an increase to the electron temperature.

It is known that electrons are more easily trapped in a laser wakefield accelerator if the initial plasma is hot, because electrons with an initial forwards momentum are more easily trapped [16]. In [69] 2-d simulations were used to show that for the conditions in question trapping occurred for $T_e \gtrsim 1$ keV. When the thermal velocity was restricted to the transverse dimension no trapping took place, confirming that the effect was due to the initial longitudinal velocity of some electrons. It was also seen that the propagation of the laser was affected by the plasma temperature.

Initially the additional ‘thermal’ energy of ionized electrons oscillating in the laser field would be in the transverse dimension. This means that for ionized ‘hot’ electrons to contribute to the trapping process, they must undergo collisions to scatter them in the forwards direction before the wakefield has passed. The frequency at which an electron with kinetic energy E_k will collide with

ions and electrons is given by [92–94]

$$\nu_{ei} = \frac{e^4 n_i \Lambda_{ei}}{\sqrt{72} \pi^{\frac{3}{2}} \epsilon_0^2 m_e E_k^{\frac{3}{2}}} \quad (6.8)$$

$$\nu_{ee} = \frac{e^4 n_e \Lambda_{ee}}{\sqrt{72} \pi^{\frac{3}{2}} \epsilon_0^2 m_e E_k^{\frac{3}{2}}} \quad (6.9)$$

where Λ is the coulomb logarithm given by

$$\Lambda_{ei} = \frac{144 \pi^2 \epsilon_0^3 T_e^3}{n_e e^6 (1 + T_e/T_i)} \quad (6.10)$$

$$\Lambda_{ee} = \frac{36 \pi^2 \epsilon_0^3 T_e^3}{n_e e^6} \quad (6.11)$$

The frequency ν is defined as the rate at which the electron is scattered through 90° . For $E_k = 1$ keV, $1/\nu_{ei}$ and $1/\nu_{ee}$ are both approximately 200 ps. Since this timescale is so long compared to τ and $1/\omega_p$ (which are on a fs timescale) it is unlikely that transverse momentum from the main laser pulse could be scattered into the forwards direction before the passing of the plasma wake.

It is possible that electrons quivering in a pre-pulse would have time to scatter in the forwards direction. However to give electrons transverse kinetic energies of order keV requires a normalised field $a \sim 0.06$, which would mean that the pre-pulse would have to have an intensity that was 1% of the peak intensity. The pre-pulse of the Astra laser was not measured for these experiments, but 1% would be unusually high.

Although a more detailed analysis would be welcome the simple calculations above would suggest that arguments based on thermal effects are unlikely to explain the correlation between recombination of the plasma channel and the generation of electron beams.

6.4.4 Current-switch effects

Having established a correlation between electron acceleration and a drop in ionization level (for the given conditions), we have considered possible ways that the ionization level could affect electron acceleration. However, it is also possible that the observed enhancement of electron acceleration is not caused by the drop in ionization level, rather both are the result of another, third effect. The favoured explanation for the drop in ionization level is that the current switches path during the discharge. A possible link between the switch in the current path and the observed enhancement in electron acceleration is considered below.

The calculations of figure 6.11 were repeated but with the initial forward velocity of the electrons varied rather than their initial position, which was set to $\xi = 0$. It was found that electrons with a forwards kinetic energy ~ 10 keV were trapped⁵. Since the applied voltage is 25 kV the potential exists to accelerate electrons to the energy required to be trapped, if they do not undergo collisions. This possibility is discussed in the following sections.

⁵This is not inconsistent with the 1 keV threshold found in [69], since that threshold is a temperature and, only a low fraction of plasma electrons are needed to contribute to the electron bunch.

Electron drift

With a current flowing the electrons have an overall drift velocity. If it is assumed that, due to the greater mass of the protons, electrons are the principle current carriers in the medium then an estimate of the mean velocity of electrons whilst current is flowing through the capillary is

$$v_{\text{dr}} = \frac{4I}{en_e\pi D^2} \approx 4 \times 10^{-5} c \quad (6.12)$$

where $D = 200 \mu\text{m}$, $n_e = 5 \times 10^{18} \text{cm}^{-3}$ and $I = 300 \text{A}$ were used. The velocity corresponds to a kinetic energy of $\sim 1 \text{meV}$, meaning that this drift is negligible compared to the thermal energy of the electrons. By considering the drift velocity only one would conclude that the discharge current should have no effect on the trapping process.

Superthermal electrons

It is wrong to assume that the electron velocity distribution is simply a thermal distribution offset by v_{dr} . The electron velocity distribution of a plasma with a current flowing is anisotropic, subject to instability and depends on the history (starting at breakdown) of the plasma [92]. If a relative minority of electrons avoided collision and were accelerated to the keV level then this would contribute towards the trapping process.

It is beyond the scope of this work to estimate the velocity distribution within the plasma during the flow of the current. However, it is interesting to consider the dependence on velocity of the electron collision frequency. In the non-relativistic case $E_k = \frac{1}{2}m_e v^2$, meaning that from (6.8) and (6.9) it can be seen that $\nu_e \propto v^{-3}$. It is well known that in a fully ionized plasma the fastest electrons undergo collisions less frequently [93]. Since the time for an effective collision is the time required for an electrons momentum to deviate by an angle of 90° , the change in momentum in the initial propagation direction can be written as $\Delta p = m_e v$. The average collisional force on an electron is then $F_c = \nu_e \Delta p \propto 1/v^2$. When accelerated by an applied electric field, the fastest electrons experience less drag than slower electrons.

Above a certain velocity threshold in the positive (current) direction, the force due to the electric field outweighs the drag force, F_c . For $n_e = 5 \times 10^{18} \text{cm}^{-3}$, $\mathcal{E} = 25 \text{kV}/15 \text{mm}$, equations (6.8) and (6.9) were used to find that this threshold was $\sim 250 \text{eV}$. In a 10eV plasma there would be very few electrons above this energy. However, if electrons at lower energy were to survive without collision for longer than expected, they could exceed 250eV . The probability P that an electron with an initial velocity \underline{v}_0 survives without collision to a time⁶ t is

$$P = e^{-\int_0^t \nu(|\underline{v}|) dt'} \quad (6.13)$$

$$\underline{v} = \underline{v}_0 - \frac{e\mathcal{E}t}{m_e} \quad (6.14)$$

where $\mathcal{E} = 25 \text{kV}/15 \text{mm}$ is the electric field due to the applied voltage. The result of applying these equations to an electron with an initial forwards kinetic energy of 10eV is plotted in figure

⁶Note that in this instance the symbol t is not the delay between the firing of the discharge and the arrival of the laser pulse.




Figure 6.12: Probability a plasma electron does not collide whilst being accelerated. The results of equations (6.13) and (6.14) for an electron with initial $E_k = 10$ eV, travelling in the forwards direction. For this calculation $n_e = 5 \times 10^{18} \text{ cm}^{-3}$ and $\mathcal{E} = 25 \text{ kV}/15 \text{ mm}$.

6.12. As predicted, above a certain energy (~ 250 eV) the electrons continue to gain energy from the applied electric field with relatively little risk of collision.

Enhanced electric field

For $\mathcal{E} = 25 \text{ kV}/15 \text{ mm}$ approximately 0.06 % of forward scattered electrons with an initial energy of 10 eV could survive to reach keV energies. This is negligible, especially when considering that the fraction would be considerably less when integrating over all angles and energies of a 10 eV thermal distribution.

However, a switch in the path of the current would have affected the value of \mathcal{E} in the sections of the capillary before and after the gas slots. Since the model assumes $R_g^* \ll R_c^*$ it follows that the voltage drop across the gas slot path is negligible. It then follows that the field in the beginning and end sections of the capillary was $\mathcal{E} \approx 25 \text{ kV}/8 \text{ mm}$ (since for the data of figure 6.7 the gas slots were located 4 mm from the capillary ends). Repeating the calculations of equations (6.13) and (6.14) with this value it is found that 2 % of forward scattered electrons with an initial energy of 10 eV become superthermal — an increase by a factor of 400. The effect of switching the current path is even more striking for the data of figure 6.8 — for that 33 mm capillary the gas slots were inset by 2 mm and a voltage of 30 kV was used. For $\mathcal{E} \approx 30 \text{ kV}/33 \text{ mm}$ (normal operation) only 0.0002 % of the forward 10 eV electrons become superthermal, whereas for $\mathcal{E} \approx 30 \text{ kV}/4 \text{ mm}$ the figure is 18 %, an increase by a factor of 90000.

A switch in the current path dramatically increases the number of superthermal electrons in the beginning and end sections of the capillary. If the current-switch model is correct and the electric

field was applied in the correct direction (this was not checked), the electron trapping would have been enhanced as the laser pulse passed through the initial section of the capillary.

Further investigation into current effects

The rough calculations above are sufficient to show that the current may have a larger role to play in electron trapping than previously expected, and the arguments based on drift velocity are over simplistic. It would be useful to perform a more rigorous theoretical analysis to better establish the feasibility of the above picture.

Fortunately a very straightforward experimental test exists to establish the validity of this hypothesis, and that is to observe the effect of swapping the direction of current flow on laser wakefield acceleration in the waveguide at densities near the electron generation threshold. Were it to be found that switching the current can enhance (or suppress) the electron trapping process the waveguide design could easily be adapted to include an injection sub-section in which the electric field gradient was high, increasing the population of superthermal electrons⁷. This may also help control the breakdown of the discharge.

6.5 Conclusion

For densities above $6.5 \times 10^{18} \text{ cm}^{-3}$ the plasma wake driven by the Astra laser was sufficiently strong to trap electrons from the background plasma. For densities below $6.5 \times 10^{18} \text{ cm}^{-3}$ the plasma wake driven by the Astra laser did not, in general, have a large enough amplitude to trap electrons. However for these densities electron bunches *could* be trapped and accelerated very stably within a narrow range of delays t . Interferometry measurements of the plasma channel and analysis of raman satellites in the long pulse spectra showed that at those delays the ionization level of the channel had decreased. From this it was concluded that the generation of electrons was connected to the ionization state of the plasma channel.

The source of ionized species was considered carefully. The experimental data does not appear to be consistent with the hypothesis that the drop in ionization level was caused by material ablated from the capillary walls. Instead, the data indicates that the hydrogen plasma was recombining prematurely, despite the fact that at those delays the discharge current was near its peak. There is some evidence to suggest that during the discharge pulse, the current changed its path from the capillary to the gas slots and reservoir. Effectively this was equivalent to abruptly switching off the current in the capillary. Although this is consistent with all available data, more experiments are required to confirm this interpretation directly.

The connection between the ionization state of the channel and the trapping of electrons is not yet fully understood. Several possible explanations were considered. Of these the most likely is that some electrons were ionized in the wakefield and were, as a result more easily trapped. Another strong possibility is that if the current switched path during the discharge, there would have been a

⁷Indeed even if the above arguments are not in fact the correct explanation for the observed effects of this thesis, the calculations still suggest that the possibility of using high fields to create superthermal electron populations could be a way of controlling electron trapping, and merits further investigation.

significant population of superthermal electrons in the beginning (and end) sections of the capillary with sufficient forwards momentum to be trapped in the wakefield.

Regarding the first of these a statistical analysis is required that includes the true ionization rates for hydrogen at high intensity — conventional ADK rates are not applicable, for a variety of reasons. Irrespective of whether this is the true mechanism assisting electron trapping at the lowest densities of the Astra experiments, the evidence suggests that electron trapping could be enhanced in future experiments if the hydrogen fill is doped with a suitable species, such as argon.

Testing the ‘superthermal’ hypothesis is very simple since the model assumes that during the Astra experiments the direction of the current flow (which was not recorded) was in the same direction as the propagation of the laser. If this is true then in future experiments it should be found that around the threshold for electron trapping switching the direction of the current should have an effect on the generation of electron beams. Were it to be found that the discharge current can contribute to electron trapping the waveguide design could be adapted to include a high applied field ‘trapping section’ to maximise the effect.

Chapter 7

Conclusions

7.1 Summary

Laser wakefield acceleration and the role of the hydrogen-filled capillary discharge waveguide, or similar waveguides, in enabling the acceleration of electrons to high energy has been studied. A combination of experimental and theoretical work was presented.

The experiments performed on the Astra laser confirmed that electrons are accelerated to higher energy when a waveguide is used: monoenergetic electron beams were accelerated to 200 MeV inside 15 mm and 33 mm long waveguides using a 0.6 J laser. That is approximately double the energy that electrons could be accelerated to without a waveguide. Scaling laws predict that the maximum possible energy gain from wakefield acceleration increases as density decreases. Analysis of the data showed that a wake was driven at lower plasma densities when the waveguide was employed, and that allowed acceleration of electrons to higher energy.

To understand the role of the hydrogen-filled capillary discharge waveguide in laser wake-field acceleration it is necessary to understand why a wake could be driven at lower densities in the waveguide. The obvious explanation would be that the extended laser-plasma interaction length afforded by the waveguide allowed acceleration over the full laser-electron dephasing length, which is longer for lower density. At the very lowest densities in the ‘linear regime’ this was indeed shown to be the case. However, for the majority of the density range in question this explanation did not hold: the dephasing length was much shorter than the Rayleigh range of the laser focal spot, and the power was far in excess of the self-focussing threshold. Under these conditions the benefit of the plasma channel, whose matched spot was considerably wider than the predicted size of the laser pulse after self-focussing, was not obvious. However, a close examination of the transmitted fluence profiles for the long pulse suggested that as well as guiding the laser pulse the plasma channel collected outlying laser energy that did not form part of the initial laser focal spot. For a real laser focus the fraction of energy lying outside the central focal spot is significant. Simulations using the measured fluence profile of the Astra focus showed that the waveguide channelled outlying energy into the intense central filament that formed after self-focussing. Simulations with a flat plasma profile showed that a self-focussed filament also formed, but did not include the outlying energy. As a result using the hydrogen-filled capillary discharge waveguide was equivalent to using a more powerful laser, and that allowed acceleration of electrons at lower plasma densities.

A separate effect, observed for the lowest plasma densities, was that the acceleration of electrons was very sensitive to the conditions of the plasma channel. The work of chapter 6 found that this was due to an unexpected dependence of electron trapping on the ionization level of the plasma. The data suggested that at densities below the electron trapping threshold for a fully ionized plasma channel, electrons could be trapped by a wakefield driven in a partially ionized plasma channel. There are several explanations currently being considered for why the presence of un-ionized species was connected to the electron trapping process, and further investigation may reveal that the effect can be controlled and used to improve the stability of laser wakefield devices. That may include mixing a buffer gas such as argon into the hydrogen gas-fill, or customising the discharge current.

In order to quantify the role of the hydrogen-filled capillary discharge waveguide, a set of coupled equations was derived from quasistatic theory [19]. These describe a self-focussed steady state of a laser propagating through a plasma, and can be used to give an estimate of the shape of the laser pulse and wakefield after self-focussing. An analytic solution to the equations gave an expression for the self-focussed spot size that agreed with the empirical expression deduced from simulations by Lu et al [11, 35]. The equations were solved numerically to predict the energy gain in the Astra experiments, and the results were in good agreement with the data, describing the transition from the non-linear scaling laws of Lu et al to the linear regime. These equations are by no means an alternative to simulation, but their simplicity and the relative speed with which they can be solved means that they are useful as a predictive tool.

In the remaining sections of this chapter future directions for research are suggested, based on what has been learned from the work of this thesis.

7.2 Short term future work

A definitive explanation is still required for the drop in the waveguide plasma ionization level measured in chapter 6. Currently there are two reasonable hypotheses for the effect: the presence of aluminium and oxygen ions ablated from the capillary wall, and the recombination of hydrogen due to a change in the current path inside the waveguide housing. Although there is considerable circumstantial evidence to support the latter, it remains to be confirmed directly. Experiments to confirm that hypothesis were suggested in chapter 6. Following on from that, the connection between the drop of the plasma ionization level and electron trapping remains a subject of discussion. Several explanations for the effect have been proposed, and further work should clarify the mechanisms responsible. Several ways of investigating that are also suggested in chapter 6. The understanding of these mechanisms may allow control of the electron injection process in future experiments.

7.3 Multi-GeV electron accelerators

The generation of multi-GeV electron beams that could replace those used by synchrotron facilities is the current goal for laser wakefield acceleration. The highest energy that has yet been reached through laser wakefield acceleration is 1 GeV. The scaling laws of Lu et al [11], which were found to agree excellently with the data of this thesis, are used to assess the laser and waveguide specifications

required to reach those energies. It is shown that with an appropriate waveguide, existing facilities could be used to accelerate electron energies to energies as high as 40 GeV in a single stage. Existing waveguide technology is reviewed to assess which schemes can meet the requirements of future experiments and devices, and a way of adapting the hydrogen-filled capillary discharge waveguide to meet those needs is suggested.

Only single stage acceleration in the non-linear regime are considered in the following sections. Staging a series of linear regime laser wake-field accelerators would be an alternative way of accelerating electrons to high energy that is not discussed below.

7.3.1 Scaling to high energy

The simulations and theoretical work of reference [11] found that for a given laser power P the optimum density for accelerating electrons was that for which $P/P_c = 1$. That represents the low density limit of the bubble regime. In order to operate at that density a plasma channel of matched spot size $W_c \approx \lambda_p/2$ was required to prevent diffraction of the front of the laser pulse. Without such an ‘ideal’ waveguide the optimum density for accelerating electrons was that for which $P/P_c = \frac{1}{8} \left(\frac{n_c}{n_e} \right)^{3/5}$. That represents the minimum density for which depletion of the laser pulse front through diffraction was less than the depletion of the laser pulse by its wake. For a given laser power, those expressions can be used to calculate the optimum density for electron acceleration, and by then substituting the results into equation (1.41), the optimum energy gain.

The results of those calculations are plotted in the top two panels of figure 7.1(a) and (b). It is evident that as the laser power increases, so does the benefit of using a waveguide: a 20 TW laser can accelerate electrons to 800 MeV with the ideal waveguide, but only 325 MeV without; a 1 PW laser can drive acceleration to 40 GeV in the waveguide, but only 6 GeV without. As laser technology advances and higher power lasers become available, the benefit of using a waveguide increases.

For a particular laser power and density there are limits on the usable pulse durations for bubble regime acceleration. The laser pulse duration should be less than λ_p/c to allow the formation of a plasma bubble without the need for pulse evolution. Matching the non-linear expressions for the dephasing and depletion lengths, given in equations (1.29) and (1.31), shows that the pulse must also be greater than $0.3\lambda_p/c$ if it is to not deplete in less than a dephasing length. These limits on the useful pulse length are plotted in figure 7.1(c). Also plotted are areas representing the capabilities of a selection of existing laser facilities. These areas were calculated by using the quoted laser specifications (minimum pulse length, maximum shot energy), and then assuming that if desired the pulse length could be lengthened, or the shot energy reduced.

Where the areas representing the laser capabilities overlap with the regions representing the ideal pulse length, optimal electron acceleration is possible. Figure 7.1 can be used to predict the optimum conditions for electron acceleration by a given laser. The conditions and energy gain for a selection of lasers are tabulated in table 7.1. For a given set of conditions the required accelerator length is given by the non-linear dephasing length, the expression for which is given in equation (1.29). The resulting values are included in table 7.1 and plotted in figure 7.1(d).

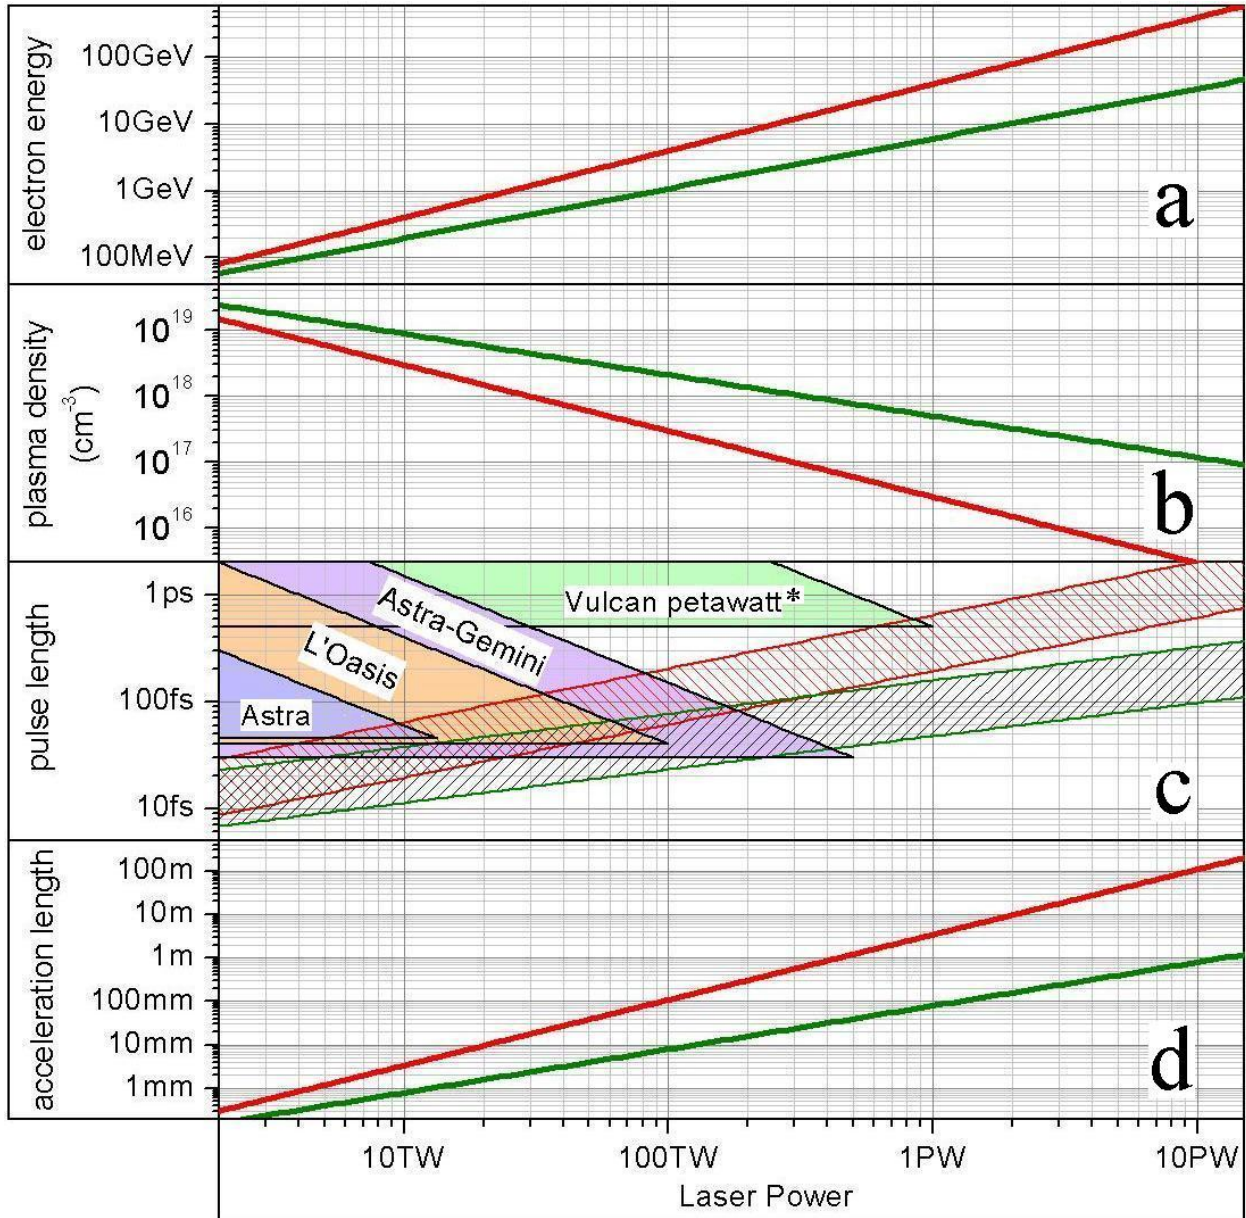


Figure 7.1: Optimum acceleration parameters for future experiments.

The parameters were calculated by adapting the scaling of [11] for an ideal waveguide (red), and for no waveguide (green) using a $\lambda = 800$ nm laser. Figure (a): maximum electron energy. Figure (b): optimum operating density. Figure (c): the upper (λ_p/c) and lower ($0.3\lambda_p/c$) limits on the required laser pulse lengths. Also plotted are areas representing the capabilities of existing laser systems to meet those requirements. Figure (d): required accelerator length, which is the non-linear dephasing length. (*The Vulcan Petawatt laser has wavelength 1055 nm, however this does not significantly affect the position of its plot.)

Laser	Chan.	E_L (J)	τ (fs)	P (TW)	W (μm)	n_e (cm^{-3})	L (mm)	E_e (MeV)
Astra	No	0.6	45	12	6	7.8×10^{18}	1	227
	Yes	0.6	45	12	11	2.4×10^{18}	5	497
L'Oasis (2006)	No	1.6	40	40	9	3.7×10^{18}	3	550
	Yes	1.6	40	40	19	7.4×10^{17}	28	1613
L'Oasis (2008-)	No	4	40	100	12	2.1×10^{18}	8	1088
	Yes	4	53	75	27	3.9×10^{17}	72	3060
Astra-Gemini	No	16	40	400	18	8.8×10^{17}	32	3092
	Yes	16	80	200	43	1.5×10^{17}	306	8040
Vulcan PW	Yes	500	500	1000	127	1.7×10^{16}	4500	40100

Table 7.1: Optimum electron accelerations for a selection of existing laser systems. All the lasers listed have a central wavelength of 800 nm, except for the Vulcan Petawatt laser, for which the wavelength is 1055 nm. The values in the column labeled W are equal to $\lambda_p/2$, which is the approximate size of the laser spot after self-focussing at optimum density.

Reaching high energy with the discharge waveguide

The first two lasers considered in table 7.1 have been used recently in conjunction with the hydrogen-filled capillary discharge waveguide. These are the Astra laser, which produced the results presented in this thesis, and the L'Oasis system which was used to accelerate electrons to 1 GeV (and at the time had a peak power of 40 TW) [7, 8]. It is interesting to note that the optimum conditions found through experimental use of the hydrogen-filled capillary discharge waveguide agree better with the ‘no channel’ scaling. That is particularly true for the Astra results: there was a peak in the accelerated electron energy at a density of $8 \times 10^{18} \text{ cm}^{-3}$, at which the electron energy was 200 MeV. For L'Oasis the optimum density was found to be $4.3 \times 10^{18} \text{ cm}^{-3}$, which agrees reasonably well with the predictions of gas jet scaling, although the maximum energy reached exceeded expectations.

The spot size of the hydrogen-filled capillary discharge waveguide is wider than that of an ideal waveguide. For the Astra experiments its role, as already discussed, was to ensure that all of the laser energy is included in the intense filament that forms after self-focussing. Once the energy was self-focussed into that filament the relatively wide plasma channel did not, in probability, play a significant role. That would explain why the ‘no channel’ scaling better applies to the results obtained with the discharge waveguide. Without any waveguide at all those same scaling laws apply, but the laser energy should be reduced significantly to take account of the non-ideal focal spot.

With that in mind, a suitable waveguide for use with the Astra-Gemini laser would have length of at least 32 mm and the capability to operate at a density $\sim 9 \times 10^{17}$. The matched spot size of the channel would not need to be as small as the value W given in table 7.1, which represents the size of the laser after self-focussing. Judging from the results of this thesis, a spot size two or three times that would be suitable. That would require a capillary of inner diameter $D \approx 250 \mu\text{m}$. Fortunately the current design of the hydrogen-filled capillary discharge waveguide would

not require modification to meet these specifications. If the laser can be operated at the optimum of its specifications then energy gains of 3 GeV are a realistic short-term goal. That in itself would match the energy of electrons in the large scale Diamond Light Source synchrotron facility, which is shown in figure 1.1(b).

Reaching high energy with an ideal waveguide

A glance at table 7.1 highlights the fact that the development of an ‘ideal’ waveguide, for which $W_c = \lambda_p/2$, would be extremely worthwhile. Where scaling laws predict that the Astra-Gemini laser could be used to accelerate electrons to 3 GeV in the hydrogen-filled capillary discharge waveguide (or possibly a gas jet), the ideal waveguide would enable the acceleration of electrons to 8 GeV over 30 cm. More intriguing is that an ideal waveguide would allow the Vulcan Petawatt laser to be used in the bubble regime. That laser has been used to accelerate electrons to energies as high as ~ 300 MeV in both the self-modulated regime and the direct laser acceleration regime [68, 95, 96]. With an ideal waveguide the optimum density for acceleration would be $n_e = 1.7 \times 10^{16}$, for which $\lambda_p/c = 850$ fs. At 500 fs the laser pulse is shorter, allowing access to the bubble regime without need for laser evolution. For those conditions the energy gain would be 40 GeV, which matches the 3.2 km Stanford Linear Accelerator, shown in figure 1.1(a).

The waveguide that could support such acceleration would need to be at least 4.5 m long and able to guide a $122 \mu\text{m}$ at the required plasma density. Clearly the development of such a device is a considerable technical challenge, in addition to the problems of focussing the laser to the correct spot size (which would require roughly 70 m) and the measurement of electron energies in the 40 GeV range, as well as the safety issues relating to radiation shielding. Whilst it may be believed that suitable lasers do not yet exist to accelerate electrons to the high energy frontier, the scaling laws suggest that existing lasers are sufficient, so long as a suitable waveguide device can be developed. Of course, the low repetition rate of a Petawatt would limit its use as an operational facility. However the demonstration that LWFA can be used to reach the high energy frontier in a single stage would be sufficient motivation to try to develop systems with higher repetition rates.

7.3.2 Developing an ideal waveguide

The ideal waveguide would be capable of guiding spots with a Gaussian radius $W \approx \lambda_p/2$. In order to assess the potential of some existing schemes to meet these requirements their matched spots are plotted in figure 7.2 for a range of plasma densities in terms of λ_p .

Existing waveguide schemes that may be suitable

Across all densities the step-index waveguide [40, 41] can guide a spot of $\lambda_p/2$ for the required distance. If it is possible to fabricate hollow glass capillaries that meet the requirements set out in table 7.1 the relative simplicity of this scheme would make it ideal for initial experimental work.

A single point representing the published results for the newly developed fast discharge gas cell [57] is plotted in figure 7.2. This is certainly a promising scheme, and current experimental results show that it might meet the requirements of an ideal waveguide for plasma densities $\sim 2 \times 10^{17} \text{ cm}^{-3}$,




Figure 7.2: Matched spots of different waveguide schemes against plasma density. The matched spots are plotted in terms of the plasma wavelength at the density in question. The waveguide schemes plotted are: the step-index waveguide [40, 41] (blue line); the fast discharge gas cell (red star); the hydrogen-filled capillary discharge waveguide (black line); the hydrogen-filled capillary discharge waveguide with an applied axial magnetic field (green dots).

which would be ideal for generating 8 GeV electron beams with the Astra-Gemini laser. The published results were for a prototype that was only 17 mm long, whereas the required length is 30 cm. It is likely that the fabrication of a longer device is relatively straightforward since the dividing layers of dielectric can be drilled individually and then assembled. It is not known if there are issues relating to the breakdown of the plasma discharge when the length of the electrodes is extended. Even if that were the case, it is likely that the problem could be overcome. It remains to be seen if this waveguide scheme scales to lower densities.

Adapting the hydrogen-filled capillary discharge waveguide

The relation $W_c \approx 3\lambda_p/2$ for the hydrogen-filled capillary discharge waveguide comes from setting $D = 5W_c$ and scaling using equations (2.9) and (2.14). It is not possible to further reduce the diameter before the interaction between the capillary wall and the laser dominates. These predictions are valid for densities down to $n_e \approx 2.5 \times 10^{17} \text{ cm}^{-3}$ which is the lowest density at which a plasma channel has been measured [64]. For lower densities it is not known if the matched spot size of the channel continues to scale as $3\lambda_p/2$. Presuming that it does, the hydrogen-filled capillary discharge waveguide does not guide a sufficiently small spot to meet the requirements for an ideal waveguide. However it could be that adapting the scheme will lead to a waveguide that can guide spots as small as $\lambda_p/2$.

Ivanov et al [97] suggested that the application of magnetic field along the axis of the hydrogen-filled capillary discharge waveguide would reduce thermal transport by restricting electron motion in the radial direction, and so increase the thermal gradients that give rise to the plasma channel. In that approach the calculations of Bobrova et al [60] for the steady state plasma channel were adapted to include magnetic fields. It was found that for $n_e \sim 10^{18} \text{ cm}^{-3}$ an applied axial magnetic field of order 10 T could significantly reduce the matched spot size of the waveguide. The hydrogen-filled capillary discharge waveguide was not adapted by the group to include an axial magnetic field because for the densities of interest at that time, an extremely large magnet would be required to have a significant impact.

The effect of the magnetic field on the thermal conductivity is described by the Hall parameter x , given by

$$x = \omega_c \tau_{ei} \quad (7.1)$$

where ω_c is the cyclotron frequency of the electrons and τ_{ei} is the electron-ion collision time. For $x \ll 7$ the effect of the magnetic field on electron transport are small, but for $x \gtrsim 7$ they affect the flow of heat perpendicular to the magnetic field significantly. In a 10 eV plasma (typical for the axial plasma)

$$x \approx 4 \times 10^{17} \text{ T}^{-1} \text{ cm}^{-3} \frac{B}{n_e} \quad (7.2)$$

from this it can be seen that at densities $\sim 10^{18} \text{ cm}^{-3}$ a 10 T field is necessary to affect electron transport. However at densities in the 10^{16} cm^{-3} – 10^{17} cm^{-3} the required field is much lower.

In order to quantify the effect, the numerical calculations of the steady state plasma profile of [97] were repeated at a lower density. The results for $D = 300 \mu\text{m}$, $I = 350 \text{ A}$ and an initial density of hydrogen molecules $n_{\text{H}_2} = 2.3 \times 10^{17} \text{ cm}^{-3}$ are plotted in figure 7.2. As the magnetic

field is increased the axial density and matched spot size decrease. With a 2 T magnetic field the matched spot was less than $\lambda_p/2$. The calculations suggest that with the addition of a relatively modest axial magnetic field the hydrogen-filled capillary discharge waveguide scheme could meet the requirements for an ideal waveguide.

Non-LTE simulations of the evolution of the waveguide plasma during the discharge current [62, 63] have shown a better agreement with experimental results than the calculations of the plasma steady state [64]. Simulations have not been performed for axial densities less than 10^{18} cm^{-3} . Before attempting to redesign the waveguide to include an axial magnetic field it would be interesting to perform non-LTE simulations in the 10^{16} cm^{-3} – 10^{17} cm^{-3} range, and to adapt the code to include an axial magnetic field to allow a better quantitative estimate of the effect.

The development of a 30 cm device to operate at densities $\sim 2 \times 10^{18} \text{ cm}^{-3}$ for use with the Astra-Gemini laser is a reasonable goal for the near future. The resulting electron energy would be 8 GeV. If successful, finding ways to extend the scheme to 4.5 m for use with the Vulcan Petawatt laser could be considered.

Appendix A

Imaging the waveguide exit

In principle the method used for imaging the exit plane of the waveguide is straightforward. A three lens system is employed comprising 2 achromatic lenses, which collimate and then refocus light originating from the image plane, and a microscope objective, which magnifies the image of the exit plane onto the CCD array of a 12-bit camera. However there are two main methods for implementing this system that have not, in the past, always arrived at the same result. It is of practical use to consider whether or not these methods are equivalent, and where errors might arise. This may prevent uncertainty and save time in future experiments.

Initially the lenses are positioned with no waveguide in place, and the plane in which the beam is focussed is imaged, as shown in figure A.1(a). When the waveguide is installed its entrance is positioned in the plane of the laser focus. The exit plane of the waveguide can then be imaged simply by translating the first achromat lens L away from the focus, as showing in figure A.1(b).

To set this up the achromat lenses are placed as accurately as possible¹ and the camera (to which the microscope objective is mounted) is translated carefully until the plane of the laser focus is being imaged on the CCD. On figure A.1(a) this corresponds to adjusting z with all other lengths fixed. There are two possible methods for determining that the laser focus has been imaged and it is these that have produced differing results in the past.

For the ‘target method’, a target (such as a thin wire) is placed in the plane of the laser focus and back-lit by light from the laser in low-power mode. The camera is then translated until the image of the target comes into sharp focus. This method is theoretically sound but it is not always convenient. The imaging system is designed for high power and there are many wedge reflections to attenuate the beam. In low power mode the signal from the laser focus is low. To back light the target it is necessary to temporarily change the focal position of the laser, dropping the signal on the camera even further, so that it is sometimes too low to be useful.

The alternative ‘quick method’ is to adjust z until the image of the spot is as bright and small as possible. This method is relatively convenient since it does not require the addition of a target or any optics, and the brightness of the signal is not an issue. Intuitively this should give the same result as the target method — changing z changes the plane being imaged without significantly changing the magnification of the system, so that when the brightness of the image is maximised

¹The system is relatively insensitive to the exact placement of the achromat lenses and so this should not be a time consuming operation.




Figure A.1: Lens system for imaging the waveguide exit.

In each subfigure the object plane is marked with an ‘o’ and the image plane (the camera CCD chip) is marked with an ‘i’.

it is logical to expect that the laser focus is being imaged.

Unfortunately it is found experimentally that these methods for setting z do not always give the same result. A possible cause for the discrepancy was thought to be differences between geometric optics (which best describe the ‘target method’) and Gaussian optics (which best describe the ‘quick method’). The correct value of z is that described by geometric optics, and so it is useful to estimate how much the position found by z using the ‘quick method’ differs from this.

For a Gaussian beam of wavelength λ , radius of wavefront curvature R_g and $1/e^2$ beam radius W , a complex beam parameter q is defined as

$$\frac{1}{q} = \frac{1}{R_g} - \frac{i\lambda}{\pi W^2} \quad (\text{A.1})$$

Ray transfer matrices can then be used to relate q_f at the laser focus to q_{CCD} at the CCD chip through the equations

$$\frac{1}{q_{\text{CCD}}} = \frac{C + D/q_f}{A + B/q_f} \quad (\text{A.2})$$

where

$$\begin{pmatrix} A & B \\ C & D \end{pmatrix} = \underbrace{\begin{pmatrix} 1 & d_2 \\ 0 & 1 \end{pmatrix}}_{\text{objective to CCD}} \overbrace{\begin{pmatrix} 1 & 0 \\ -1/f_2 & 1 \end{pmatrix}}^{\text{microscope objective}} \underbrace{\begin{pmatrix} 1 & z' \\ 0 & 1 \end{pmatrix}}_{\text{lens pair to objective}} \overbrace{\begin{pmatrix} 1 & 0 \\ -1/f_1 & 1 \end{pmatrix}}^{\text{lens pair}} \underbrace{\begin{pmatrix} 1 & d_1 \\ 0 & 1 \end{pmatrix}}_{\text{focus to lens pair}} \quad (\text{A.3})$$

is the ray transfer matrix for the system, itself the matrix product of the individual ray transfer matrix for each element. Where there is a lens pair in the real system there is a single lens in equation (B.7). That is because the transfer matrix for a pair of lenses of focal length f_a and f_b separated by a space S is equivalent to the transfer matrix for a space $S_a = \frac{Sf_b}{f_a+f_b-S}$, followed by a lens of focal length $f_1 = \frac{f_a f_b}{f_a+f_b-S}$, followed by another space $S_b = \frac{Sf_a}{f_a+f_b-S}$. It is therefore more simple to treat the lens pair as a single lens and the problem as two lenses rather than three lenses. This is shown in figure A.1(c), with $d_1 = f_a + S_a$ and $z' = S_b + z$. When scanning z , S_b remains fixed, meaning that $dz' = dz$.

The value of z' for which the spot size on the CCD chip will be minimised is that for which the imaginary part of $1/q_{\text{CCD}}$ is maximised. Because q_f represents the focus it is completely imaginary meaning that from equation (A.1)

$$\text{Im} \left\{ \frac{1}{q_{\text{CCD}}} \right\} = \frac{\frac{1}{q_f} (\overbrace{AD - BC})}{A^2 - \frac{B^2}{q_f^2}} = \frac{\frac{1}{q_f}}{A^2 - \frac{B^2}{q_f^2}} \quad (\text{A.4})$$

Since the top half of this expression is independent of z' , the derivative of $\text{Im} \left\{ \frac{1}{q_{\text{CCD}}} \right\}$ will be equal to zero when the derivative of the bottom part is equal to zero, that is when

$$\frac{d}{dz'} \left[A^2 - \frac{B^2}{q_f^2} \right] = 0 \quad (\text{A.5})$$

therefore

$$A \frac{dA}{dz'} = \frac{1}{q_f^2} B \frac{dB}{dz'} \quad (\text{A.6})$$

The values of A and B from equation (B.7) are

$$A = 1 - \frac{d_2}{f_2} - \frac{d_2}{f_1} - \frac{z'}{f_1} \left(1 - \frac{d_2}{f_2} \right) \quad (\text{A.7})$$

$$B = d_1 \left(1 - \frac{d_2}{f_2} \right) + d_2 \left(1 - \frac{d_1}{f_1} \right) + z' \left(1 - \frac{d_1}{f_1} \right) \left(1 - \frac{d_2}{f_2} \right) \quad (\text{A.8})$$

An expression for z' may be derived by combining these with equation (A.6) and using the fact that $1/q_f^2 = -Z_R^2$:

$$z' = \frac{d_1 f_1 (d_2 - f_2) + d_2 f_2 (d_1 - f_1) + Z_R^2 \left[\frac{-f_1 f_2 + d_2 f_1 + d_2 f_2}{d_1 f_1} \right]}{(d_1 - f_1) (d_2 - f_2) + Z_R^2 \left[\frac{d_2 - f_2}{d_1 - f_1} \right]} \quad (\text{A.9})$$

For $Z_R = 0$ the geometric expression for z' results. The two methods of imaging the plane of the focus are not equivalent and theoretically the quick method is incorrect. The important question however, is not whether the two methods are equivalent theoretically, to which the answer is no, but whether or not they are equivalent in practical terms. The difference between z' found using

Method	Systematic Error	Random error in imaging focus	Random error in imaging exit mode
Target method	0	$\sqrt{Z_R^2 + d_f^2}$	d_f
Quick method	$\frac{S-2f_a}{f_a^2} Z_R^2$	Z_R	$\sqrt{2}Z_R$

Table A.1: Magnitude of errors for the target method and quick method.

the target method and z' found using the quick method is found by subtracting the expression for z' from equation (A.9) with $Z_R = 0$ from the expression for z' with $Z_R \neq 0$ to give

$$\Delta z' = \frac{-f_1^2 Z_R^2}{(d_1 - f_1)^3 + Z_R^2 (d_1 - f_1)} \quad (\text{A.10})$$

For the Astra, and most similar experiments, $f_a = f_b \gg Z_R$ meaning that the systematic error in setting the camera position using the quick method is

$$\Delta z \approx \frac{S - 2f_a}{f_a^2} Z_R^2 \quad (\text{A.11})$$

For the Astra experiment

$$\begin{aligned} f_a &= 1200 \text{ mm} \\ S &\approx 2000 \text{ mm} \\ Z_R &= 4.5 \text{ mm} \\ \therefore \Delta z &\approx 0.0056 \text{ mm} \end{aligned}$$

This is small compared to both the depth of field of the imaging system, $d_f \sim 0.5 \text{ mm}$, and the random errors associated with setting up the system. The systematic difference between the quick method and target method is not significant.

Past discrepancies between the target method and quick method most likely arose from the fact that random errors enter into the system at different stages of the setup. In the target method the random error in setting the position of z is d_f whilst it is Z_R for the quick method. However there is an additional error in the target method, which arises from determining the plane of the laser focus in which the target should be placed. That error is of order Z_R . Having said that the plane that is determined to be the laser focus is also the plane in which the capillary entrance is placed, so that the error in imaging the capillary exit plane when the target method has been used is just d_f . The random error Z_R in choosing the location of the capillary entrance then contributes to the error in imaging the capillary exit for the system set up using the quick method. The errors associated with the two methods are summarised in table A.1. The random errors are more significant than any systematic difference between the two methods. In reality the position of a laser focus can be determined to much greater precision than Z_R so long as care is taken. In recent experiments where the waveguide was mounted inside the vacuum chamber on motorised stages it was possible to use other methods that were both fast and accurate, but where this is not possible the quick

method should suffice.

Appendix B

Chromatic aberration

Results from the Astra experiment are presented in chapter 4 which show that the measured fluence profile at the capillary exit was distorted for the short pulse, but not the long pulse. An interpretation of this could be that the non-linear interaction between the intense short pulse and the plasma affected the transverse profile of the pulse. The results also showed that the short pulse laser spectrum was broadened considerably, whereas the long pulse was not. Another interpretation, therefore, is that the distortion of the short pulse fluence profile was a result of chromatic aberration in the lens system imaging the waveguide exit. In this appendix an attempt is made to quantify the effect of chromatic aberration on the measured fluence profile, and deduce the real spot size of the transmitted radiation.

B.1 Quantifying chromatic aberration

A schematic diagram of the imaging system is shown in figure A.1. It was found, from analysis of the lens system, that the contribution of the objective lens to chromatic aberration was relatively small. In the following analysis that lens is ignored. The other two lenses, which were used to collimate and then refocus radiation emerging from the capillary exit were achromatic doublets designed specifically for aberration free imaging at wavelengths centered on 800 nm. The range of wavelengths over which the lenses were aberration free was limited, however, meaning that their focal length f was a function of λ .

Data for the achromat lenses was not available from the manufacturer, but by considering the way an achromatic doublet lens works, $f(\lambda)$ can be estimated. Figure B.1 shows the lens design. The refractive index of the glass from which the two different parts are made are $n_1(\lambda)$ and $n_2(\lambda)$. The radii of curvature of the 3 marked surfaces are R_1 , R_2 and R_3 . The focal length at the central wavelength λ_c , is f_0 . Because the wavelength dependence of $n_1(\lambda)$ and $n_2(\lambda)$ differ, the values of R can be chosen so that $f(\lambda_b) = f(\lambda_r)$ where λ_b and λ_r lie either side of λ_c . The thickness of the lenses was negligible compared to $f(\lambda)$ and can be ignored, so that in general

$$\frac{1}{f(\lambda)} = \frac{\alpha_1}{R_1} + \frac{\alpha_2}{R_2} + \frac{\alpha_3}{R_3} \quad (\text{B.1})$$

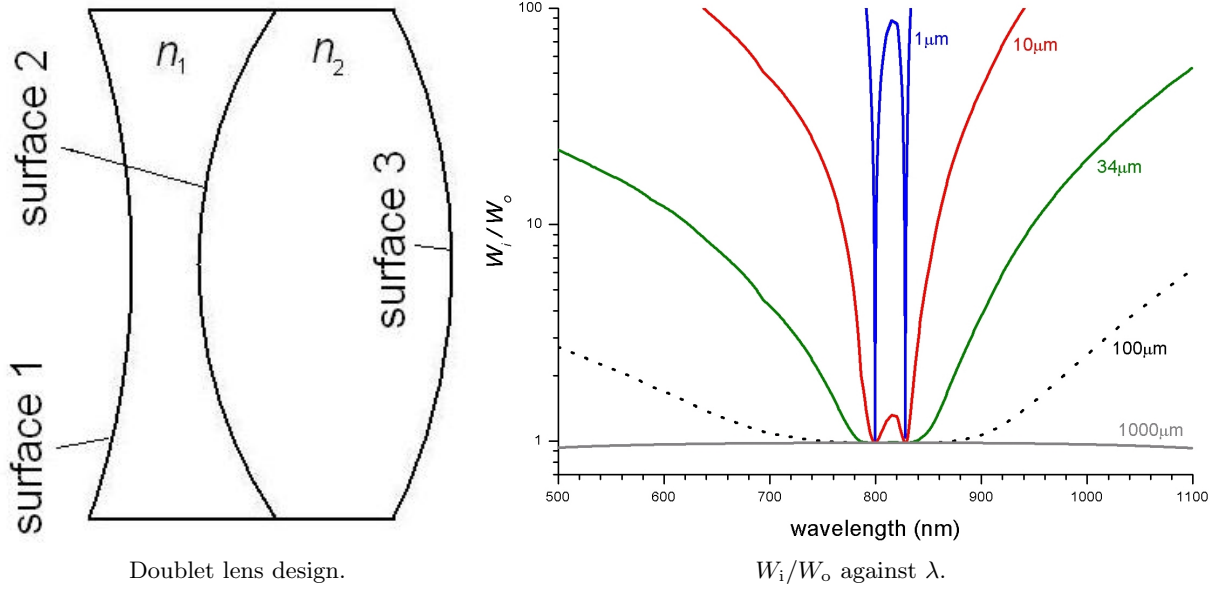


Figure B.1: The design of a doublet lens.

R_1 , R_2 and R_3 must satisfy the equations

$$\frac{1}{f_0} = \frac{\alpha_1^c}{R_1} + \frac{\alpha_2^c}{R_2} + \frac{\alpha_3^c}{R_3} \quad (\text{B.2})$$

$$\frac{1}{f_1} = \frac{\alpha_1^r}{R_1} + \frac{\alpha_2^r}{R_2} + \frac{\alpha_3^r}{R_3} \quad (\text{B.3})$$

$$= \frac{\alpha_1^b}{R_1} + \frac{\alpha_2^b}{R_2} + \frac{\alpha_3^b}{R_3} \quad (\text{B.4})$$

where $\alpha_1 = n_1 - 1$, $\alpha_2 = n_2 - n_1$ and $\alpha_3 = 1 - n_2$. These were rearranged to give

$$\frac{1}{R_2} = \frac{\frac{1}{R_2} \Gamma_2^b + \frac{1}{f_1} [\alpha_3^b - \alpha_3^r]}{\Gamma_1^b} \quad (\text{B.5})$$

$$\frac{1}{R_3} = \frac{\frac{1}{R_3} \Gamma_3^r + \frac{1}{f_1} [\alpha_2^r - \alpha_2^b]}{\Gamma_1^b} \quad (\text{B.6})$$

where $\Gamma_i^b = \alpha_j^r \alpha_k^b - \alpha_j^b \alpha_k^r$ and so on.

For $n_1(\lambda)$ and $n_2(\lambda)$ refractive index data for a crown glass (B270) and a flint glass (SF11) were obtained. These are glasses of they types typically used in doublet lenses. The lenses used were designed so that $\lambda^c = 800$ nm and $f_0 = 1200$ mm. Having chosen¹ values for R_1 , λ^b and λ^r , equations (B.5) and (B.6) could be used to calculate values for R_2 and R_3 . Then an estimate of $f(\lambda)$ was made using equation (B.1).

Having found $f(\lambda)$ it is possible to consider the propagation of Gaussian beams of differing wavelengths through the imaging system by use of ray transfer matrices. Using the same notation as in appendix A, q_0 , the complex beam parameter in the object plane is converted to q_i , the

¹Although R_1 , λ^b and λ^r are not known, it was found that the function $f(\lambda)$ did not strongly depend on the values chosen.

complex beam parameter in the image plane of the lens system by equation (A.2) where

$$\begin{pmatrix} A & B \\ C & D \end{pmatrix} = \underbrace{\begin{pmatrix} 1 & f_0 \\ 0 & 1 \end{pmatrix}}_{\text{achromat to image}} \overbrace{\begin{pmatrix} 1 & 0 \\ -1/f(\lambda) & 1 \end{pmatrix}}^{\text{achromat}} \underbrace{\begin{pmatrix} 1 & S \\ 0 & 1 \end{pmatrix}}_{\text{between achromats}} \overbrace{\begin{pmatrix} 1 & 0 \\ -1/f(\lambda) & 1 \end{pmatrix}}^{\text{achromat}} \underbrace{\begin{pmatrix} 1 & f_0 \\ 0 & 1 \end{pmatrix}}_{\text{object to achromat}} \quad (\text{B.7})$$

For simplicity it was assumed that the beam emerging from the capillary is approximately collimated so that

$$\frac{1}{q_o} \approx -\frac{i\lambda}{\pi W_o} \quad (\text{B.8})$$

where W_o was the beam radius in the object plane. The beam radius in the image plane W_i , which was a function W_o and λ , could be extracted from the imaginary part of $1/q_i$. In figure B.1 the ratio W_i/W_o is plotted against λ for a selection of W_o . Between 785 nm and 835 nm aberration was not an issue when imaging the 34 μm spot. This was approximately the bandwidth of the Astra laser, meaning that the input spot and the transmitted spot of the long pulse (which was generally not broadened) were imaged correctly. At other wavelengths W_i/W_o rises far above 1, meaning that measurement of the transmitted fluence profiles for the spectrally broadened short pulse would have been affected very significantly by chromatic aberration.

B.2 Deducing the real transmitted fluence profile

Having quantified the effect of chromatic aberration, it is possible correct the measured fluence profile to give an estimate of the real profile at the capillary exit. In order to make the correction it was assumed that the fluence profile at the capillary exit was Gaussian with a width W_o for all wavelengths. The Gaussian width $W_i(W_o, \lambda)$ of the fluence in the image plane for every wavelength element was calculated using equations (A.9) and (B.7). The measured laser spectrum $g(\lambda)$, which also needed to be corrected to account for chromatic aberration, was then used to integrate the Gaussian profile of each wavelength element to give an estimate of the image profile. This was repeated for different values of W_o until the RMS difference between the deduced image profile and the measured image profile was minimised. The calculation and its results are outlined in the remainder of the section.

B.2.1 Correcting measurements of the laser spectrum

The slit of the spectrometer was $\kappa = 12 \mu\text{m}$ wide and positioned in the image plane of the achromat system. The fraction of energy entering the spectrometer at a particular wavelength depended on the value of W_i at that wavelength. The energy from a wavelength element that passes through

the spectrometer slit is

$$\begin{aligned} B(W_o, \lambda) &= \int_{-\infty}^{+\infty} \int_{-\kappa/2}^{+\kappa/2} A \exp \left[-\frac{2(x^2 + y^2)}{W_i^2} \right] dx dy \\ &= AW_i \sqrt{\frac{\pi}{2}} \int_{-\kappa/2}^{+\kappa/2} \exp \left[-\frac{2x^2}{W_i^2} \right] dx \end{aligned} \quad (\text{B.9})$$

where $A(W_o, \lambda)$ is the peak fluence of that wavelength element in the image plane. The total energy of a wavelength element before the spectrometer is $\frac{\pi}{2}AW_i^2$. That means that the fraction of a wavelength element entering the spectrometer is

$$\begin{aligned} C(W_o, \lambda) &= \frac{2B}{\pi AW_i^2} \\ &= \frac{1}{W_i} \sqrt{\frac{2}{\pi}} \int_{-\kappa/2}^{+\kappa/2} \exp \left[-\frac{2x^2}{W_i^2} \right] dx \end{aligned} \quad (\text{B.10})$$

$$= \frac{1}{\sqrt{\pi}} \int_{-\kappa/\sqrt{2}W_i(W_o, \lambda)}^{+\kappa/\sqrt{2}W_i(W_o, \lambda)} \exp[-m^2] dm \quad (\text{B.11})$$

The corrected spectrum $g'(W_o, \lambda)$ is therefore given by

$$g'(W_o, \lambda) = \frac{g(\lambda)}{C(W_o, \lambda)} \quad (\text{B.12})$$

B.2.2 Estimating the image plane fluence for an assumed W_o

For each shot of the experiment a measurement was made of the transmitted spectrum $g(\lambda)$ and energy E_T of the laser pulse. Assuming a value for W_o , the transmitted energy between λ and $\lambda + d\lambda$ is $D(W_o, \lambda)d\lambda$ where

$$D(W_o, \lambda) = E_T \frac{g'(W_o, \lambda)}{\int_0^\infty g'(W_o, \lambda)d\lambda} \quad (\text{B.13})$$

The peak image plane fluence of radiation between λ and $\lambda + d\lambda$ is $Ad\lambda$ and

$$A = \frac{2D}{\pi W_i^2} \quad (\text{B.14})$$

which means that the fluence profile in the image plane is

$$F_i(W_o, r) = \int_0^\infty A \exp \left[-\frac{2r^2}{W_i^2} \right] d\lambda \quad (\text{B.15})$$

Note that both W_i and A are functions of W_o and λ .

B.2.3 Estimating W_o from the measured fluence profile

Equations (B.12–B.15) estimate F_i , the fluence profile in the image plane for an assumed value of W_o . A program was written to repeat these calculations for many W_o to find the value that minimised the RMS difference between the measurement and estimate of F_i . This provides an

estimate for W_o . The resulting estimate of the fluence profile at the capillary exit is then given by

$$F_o(W_o, r) = \frac{2E_T}{\pi W_o^2} \exp \left[-\frac{2r^2}{W_o^2} \right] \quad (\text{B.16})$$

Examples of the results of using this method are plotted in figure B.2. In all cases the measured and predicted F_i are in reasonable agreement. For the input spot and long pulse, F_o is the same as the measured F_i , because the majority of the pulse energy was at wavelengths unaffected by chromatic aberration. However for the short pulse F_o is considerably narrower than F_i , and has a Gaussian width similar to that of the guided long pulse and input spot. That would suggest that the distortion to the fluence profiles measured for the short pulse was solely the result of chromatic aberration, and not the result of distortion of the transverse laser profile by non-linear interactions. In the case of the $L = 15$ mm short pulse data, which was collected at a relatively high electron density, the peak fluence is considerably lower than at the capillary entrance. That is due to depletion of the laser pulse energy as through wakefield excitation, not an inability to guide the laser profile, which is narrow at the capillary exit. For the $L = 50$ mm data the transmission was high and F_o is very similar to the fluence profile of the input spot. This suggests that even at intensities of around $10^{18} \text{ W cm}^{-3}$ controlled guiding is possible over many Z_R using the hydrogen-filled capillary discharge waveguide.

Although it was possible to try to recover an estimate of the real exit fluence profile using the above methods, it would be preferable if the measurement was aberration free. For that reason future high intensity guiding experiments, where spectral broadening is likely, would benefit from the use of focussing mirrors (which are not subject to chromatic aberration) in place of achromat lenses, which are of limited use.

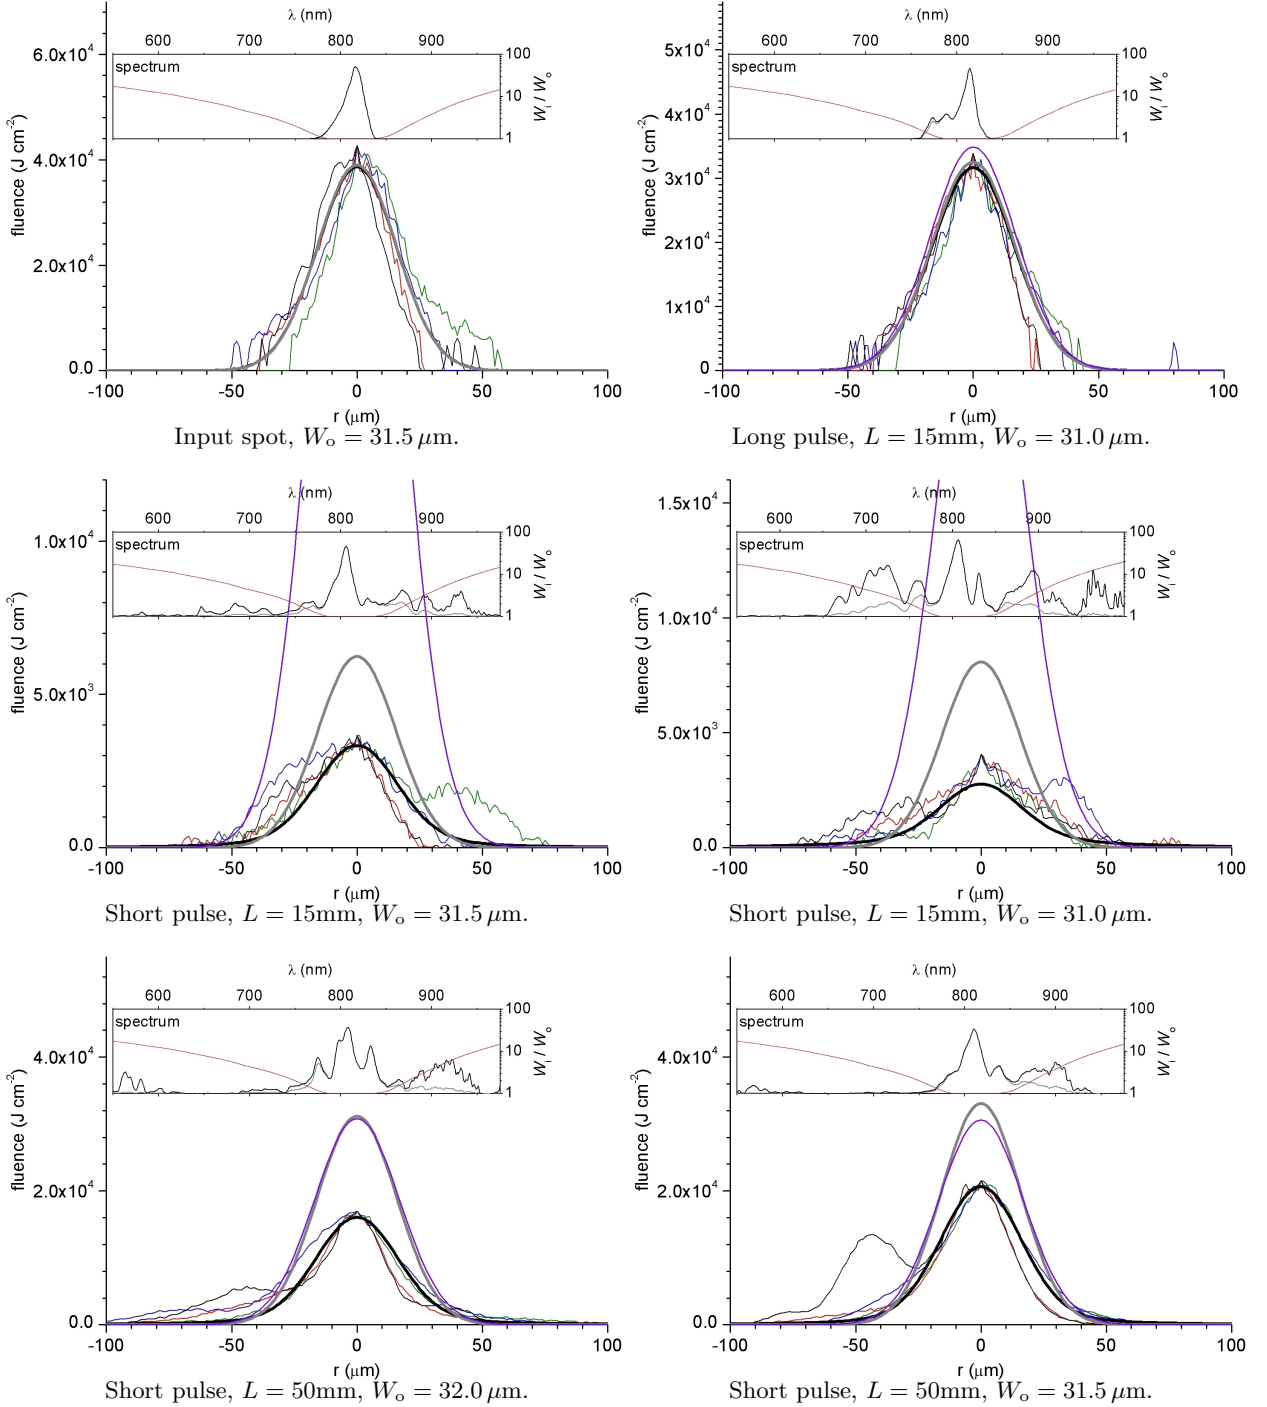


Figure B.2: Estimates of the real exit fluence profile and corrected laser spectrum. The following fluences are plotted against radius r : measured F_1 (coloured thin lines); estimated F_1 (black line); estimated F_0 (grey line); estimated F at the capillary entrance, using the measurement of the input energy and assuming $W = 34 \mu\text{m}$. The inset figure plots: $g(\lambda)$ (grey line); $g'(\lambda)$; W_1/W_0 (red line, against the right hand axis). In each figure caption the deduced value of W_0 is given. The full waveguide conditions were: for $L = 15 \text{ mm}$, [$L = 15 \text{ mm}$, $D = 200 \mu\text{m}$, $p = 600 \text{ mbar}$]; for $L = 50 \text{ mm}$, [$L = 50 \text{ mm}$, $D = 250 \mu\text{m}$, $p = 150 \text{ mbar}$].

Appendix C

Ionization thresholds for elements of interest

Values for the energy required to ionize different atomic species is widely available [98]. The following formula, from reference [85], was used to estimate from this data the laser intensity required for above barrier ionization for several different elements of interest.

$$I_{\text{ABI}} = \frac{c\pi^2\epsilon_0^3}{2e^6} \frac{E_{\text{I}}^4}{Z^2} \quad (\text{C.1})$$

where E_{I} is the energy required to ionize a single ion of the species in question, and Z is the overall charge of the resulting ion species, in atomic units. In practical units this is

$$I_{\text{ABI}} [\text{W cm}^{-2}] = 4.0467 \times 10^9 \times \frac{(E_{\text{I}} [\text{eV}])^4}{Z^2} \quad (\text{C.2})$$

The resulting values are shown in table C.1 and figure C.1.

Hydrogen ($Z = 1$)

Charge	0
Energy	14
Threshold	14.1

Helium ($Z = 2$)

Charge	0	1
Energy	25	54
Threshold	15.2	15.9

Oxygen ($Z = 8$)

Charge	0	1	2	3	4	5	6	7
Energy	14	35	55	77	114	138	739	871
Threshold	14.1	15.2	15.6	16.0	16.4	16.6	19.4	19.6

Neon ($Z = 10$)

Charge	0	1	2	3	4	5	6	7	8	9
Energy	22	41	63	97	126	158	207	239	1196	1362
Threshold	14.9	15.5	15.9	16.4	16.6	16.8	17.2	17.3	20.0	20.1

Aluminium ($Z = 13$)

Charge	0	1	2	3	4	5	6	7	8	9	10	11	12
Energy	6	19	28	120	154	190	242	285	330	399	442	2086	2304
Threshold	12.7	14.1	14.5	16.7	17.0	17.2	17.5	17.6	17.8	18.0	18.1	20.7	20.8

Argon ($Z = 18$)

Charge	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17
Energy	16	28	41	60	75	91	124	143	422	479	539	618	686	756	855	918	4121	4426
Threshold	14.4	14.8	15.1	15.5	15.7	15.9	16.3	16.4	18.2	18.3	18.5	18.6	18.7	18.8	19.0	19.1	21.6	21.7

Table C.1: Above Barrier Ionization Thresholds for selected elements

The values of the row named charge are the charge of a species prior to ionization. The values of the row named energy are the energies in eV required to ionize that species from its ground state, taken from [98]. The barrier ionization threshold calculated using this data and equation (C.1) is expressed in the bottom row, labelled threshold, as the base-10 log of the threshold intensity in units of W cm^{-2} . Note that the intensity of the long and short pulses expressed in this way are 17.4 and 17.9 respectively.

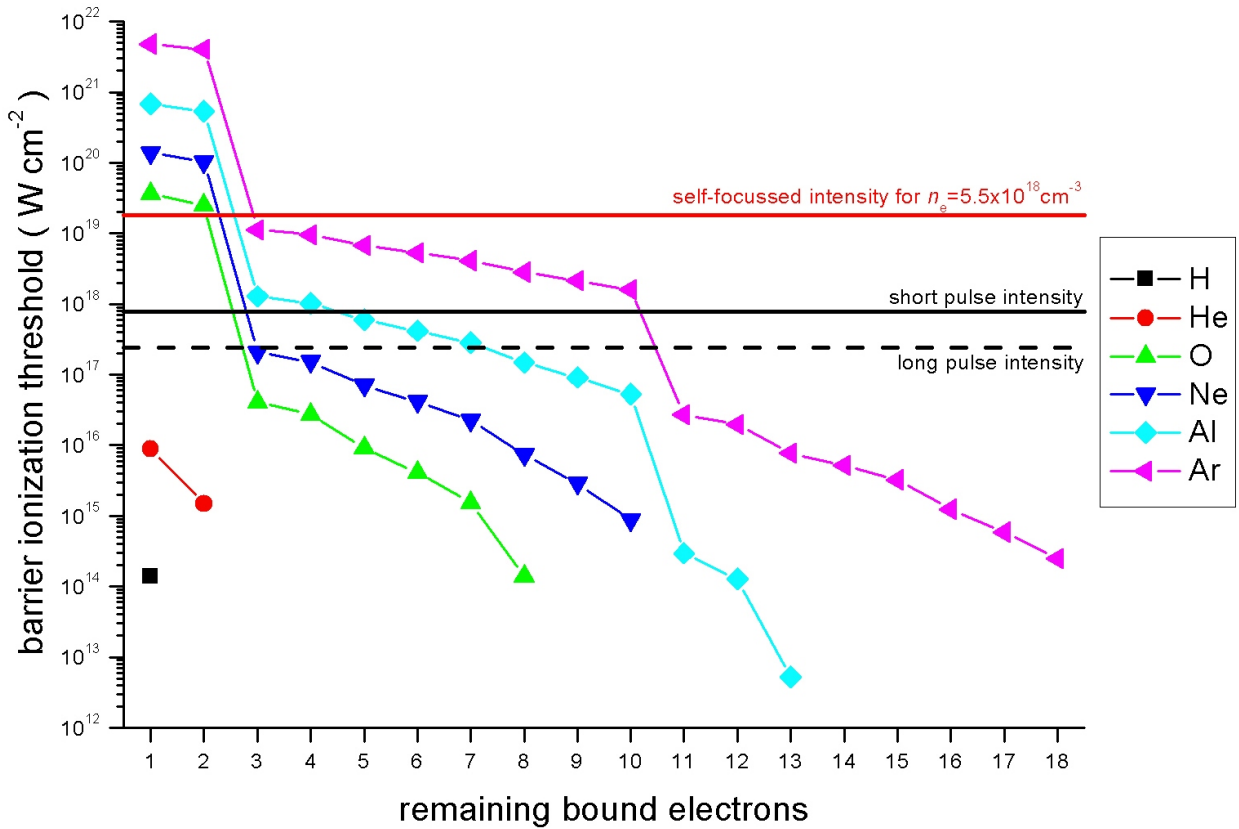


Figure C.1: Above Barrier Ionization Thresholds for selected elements

The above barrier ionization intensities of table C.1 are plotted against the number of bound electrons of the ionization species (before ionization). The following intensities of interest are also marked: Astra experiment long pulse intensity (dotted line); Astra experiment short pulse intensity (solid black line); Astra experiment short pulse intensity after self-focussing for $n_e = 5.5 \times 10^{18} \text{ cm}^{-3}$, calculated using bubble regime theory of reference [11]. The unorthodox choice of x -axis is simply to allow data for all the selected elements to be easily visible without excessive overlapping.

Bibliography

- [1] A. Butler, D. J. Spence, and S. M. Hooker. Guiding of high-intensity laser pulses with a hydrogen-filled capillary discharge waveguide. *Physical Review Letters*, 89(18):–, 2002.
- [2] T. Tajima and J. M. Dawson. Laser electron-accelerator. *Physical Review Letters*, 43(4): 267–270, 1979.
- [3] D. Strickland and G. Mourou. Compression of amplified chirped optical pulses. *Optics Communications*, 55:447–449, October 1985.
- [4] S. P. D. Mangles, C. D. Murphy, Z. Najmudin, A. G. R. Thomas, J. L. Collier, A. E. Dangor, E. J. Divall, P. S. Foster, J. G. Gallacher, C. J. Hooker, D. A. Jaroszynski, A. J. Langley, W. B. Mori, P. A. Norreys, F. S. Tsung, R. Viskup, B. R. Walton, and K. Krushelnick. Monoenergetic beams of relativistic electrons from intense laser-plasma interactions. *Nature*, 431(7008):535–538, 2004.
- [5] J. Faure, Y. Glinec, A. Pukhov, S. Kiselev, S. Gordienko, E. Lefebvre, J. P. Rousseau, F. Burgy, and V. Malka. A laser-plasma accelerator producing monoenergetic electron beams. *Nature*, 431(7008):541–544, 2004.
- [6] C. G. R. Geddes, C. Toth, J. van Tilborg, E. Esarey, C. B. Schroeder, D. Bruhwiler, C. Nieter, J. Cary, and W. P. Leemans. High-quality electron beams from a laser wakefield accelerator using plasma-channel guiding. *Nature*, 431(7008):538–541, 2004.
- [7] W. P. Leemans, B. Nagler, A. J. Gonsalves, C. Toth, K. Nakamura, C. G. R. Geddes, E. Esarey, C. B. Schroeder, and S. M. Hooker. Gev electron beams from a centimetre-scale accelerator. *Nature Physics*, 2(10):696–699, 2006.
- [8] K. Nakamura, B. Nagler, C. Toth, C. G. R. Geddes, C. B. Schroeder, E. Esarey, W. P. Leemans, A. J. Gonsalves, and S. M. Hooker. Gev electron beams from a centimeter-scale channel guided laser wakefield accelerator. *Physics of Plasmas*, 14(5):–, 2007.
- [9] J. Faure, C. Rechatin, A. Norlin, A. Lifschitz, Y. Glinec, and V. Malka. Controlled injection and acceleration of electrons in plasma wakefields by colliding laser pulses. *Nature*, 444(7120): 737–739, 2006.
- [10] H. P. Schlenvoigt, K. Haupt, A. Debus, F. Budde, O. Jackel, S. Pfotenhauer, H. Schwoerer, E. Rohwer, J. G. Gallacher, E. Brunetti, R. P. Shanks, S. M. Wiggins, and D. A. Jaroszynski. A compact synchrotron radiation source driven by a laser-plasma wakefield accelerator. *Nat Phys*, 4(2):130–133, February 2008. doi: <http://dx.doi.org/10.1038/nphys811>. URL <http://dx.doi.org/10.1038/nphys811>.
- [11] W. Lu, M. Tzoufras, C. Joshi, F. S. Tsung, W. B. Mori, J. Vieira, R. A. Fonseca, and L. O. Silva. Generating multi-gev electron bunches using single stage laser wakefield acceleration in a 3d nonlinear regime. *Physical Review Special Topics-Accelerators and Beams*, 10(6):–, 2007.
- [12] Pisin Chen, J. M. Dawson, Robert W. Huff, and T. Katsouleas. Acceleration of electrons by the interaction of a bunched electron beam with a plasma. *Phys. Rev. Lett.*, 54(7):693–696, Feb 1985. doi: 10.1103/PhysRevLett.54.693.

- [13] I. Blumenfeld, C. E. Clayton, F. J. Decker, M. J. Hogan, C. K. Huang, R. Ischebeck, R. Iverson, C. Joshi, T. Katsouleas, N. Kirby, W. Lu, K. A. Marsh, W. B. Mori, P. Muggli, E. Oz, R. H. Siemann, D. Walz, and M. M. Zhou. Energy doubling of 42 gev electrons in a metre-scale plasma wakefield accelerator. *Nature*, 445(7129):741–744, 2007.
- [14] P. Sprangle, E. Esarey, A. Ting, and G. Joyce. Laser wakefield acceleration and relativistic optical guiding. *Applied Physics Letters*, 53(22):2146–2148, 1988. doi: 10.1063/1.100300. URL <http://link.aip.org/link/?APL/53/2146/1>.
- [15] C. Joshi, T. Tajima, J. M. Dawson, H. A. Baldis, and N. A. Ebrahim. Forward raman instability and electron acceleration. *Phys. Rev. Lett.*, 47(18):1285–1288, Nov 1981. doi: 10.1103/PhysRevLett.47.1285.
- [16] W. L. Kruer. *Physics of Laser Plasma Interactions*. Perseus Publishing, 2003.
- [17] J. Osterhoff, A. Popp, Zs. Major, B. Marx, T. P. Rowlands-Rees, M. Fuchs, M. Geissler, R. Hörlein, B. Hidding, S. Becker, E. A. Peralta, U. Schramm, F. Grüner, D. Habs, F. Krausz, S. M. Hooker, and S. Karsch. Generation of stable, low-divergence electron beams by laser-wakefield acceleration in a steady-state-flow gas cell. *Physical Review Letters*, 101(8):085002, 2008. doi: 10.1103/PhysRevLett.101.085002. URL <http://link.aps.org/abstract/PRL/v101/e085002>.
- [18] E. Esarey, P. Sprangle, J. Krall, and A. Ting. Self-focusing and guiding of short laser pulses in ionizing gases and plasmas. *Ieee Journal of Quantum Electronics*, 33(11):1879–1914, 1997.
- [19] P. Sprangle, E. Esarey, and A. Ting. Nonlinear-theory of intense laser-plasma interactions. *Physical Review Letters*, 64(17):2011–2014, 1990.
- [20] P. Sprangle, E. Esarey, and A. Ting. Nonlinear-interaction of intense laser-pulses in plasmas. *Physical Review A*, 41(8):4463–4467, 1990.
- [21] S. C. Wilks, J. M. Dawson, W. B. Mori, T. Katsouleas, and M. E. Jones. Photon accelerator. *Phys. Rev. Lett.*, 62(22):2600–2603, May 1989. doi: 10.1103/PhysRevLett.62.2600.
- [22] E. Esarey, A. Ting, and P. Sprangle. Frequency-shifts induced in laser-pulses by plasma-waves. *Physical Review A*, 42(6):3526–3531, 1990.
- [23] C. D. Murphy, R. Trines, J. Vieira, A. J. W. Reitsma, R. Bingham, J. L. Collier, E. J. Divall, P. S. Foster, C. J. Hooker, A. J. Langley, P. A. Norreys, R. A. Fonseca, F. Fiuza, L. O. Silva, J. T. Mendonca, W. B. Mori, J. G. Gallacher, R. Viskup, D. A. Jaroszynski, S. P. D. Mangles, A. G. R. Thomas, K. Krushelnick, and Z. Najmudin. Evidence of photon acceleration by laser wake fields (vol 13, pg 033108, 2006). *Physics of Plasmas*, 13(7):–, 2006.
- [24] Thomas M. Antonsen and Zhigang Bian. Ionization induced scattering of short intense laser pulses. *Phys. Rev. Lett.*, 82(18):3617–3620, May 1999. doi: 10.1103/PhysRevLett.82.3617.
- [25] H. Jang, M. S. Hur, J. M. Lee, M. H. Cho, W. Namkung, and H. Suk. A method to measure the electron temperature and density of a laser-produced plasma by raman scattering. *Applied Physics Letters*, 93(7):071506, 2008. doi: 10.1063/1.2973395. URL <http://link.aip.org/link/?APL/93/071506/1>.
- [26] C. D. Decker and W. B. Mori. Group velocity of large amplitude electromagnetic waves in a plasma. *Phys. Rev. Lett.*, 72(4):490–493, Jan 1994. doi: 10.1103/PhysRevLett.72.490.
- [27] C. D. Decker, W. B. Mori, K. C. Tzeng, and T. Katsouleas. Evolution of ultra-intense, short-pulse lasers in underdense plasmas. *Physics of Plasmas*, 3(5):2047–2056, 1996.
- [28] W.B. Mori. The physics of the nonlinear optics of plasmas at relativistic intensities for short-pulse lasers. *Quantum Electronics, IEEE Journal of*, 33(11):1942–1953, Nov 1997. ISSN 0018-9197. doi: 10.1109/3.641309.

- [29] D. F. Gordon, B. Hafizi, R. F. Hubbard, J. R. Peñano, P. Sprangle, and A. Ting. Asymmetric self-phase modulation and compression of short laser pulses in plasma channels. *Phys. Rev. Lett.*, 90(21):215001, May 2003. doi: 10.1103/PhysRevLett.90.215001.
- [30] S. V. Bulanov, I. N. Inovenkov, V. I. Kirsanov, N. M. Naumova, and A. S. Sakharov. Nonlinear depletion of ultrashort and relativistically strong laser pulses in an underdense plasma. *Physics of Fluids B: Plasma Physics*, 4(7):1935–1942, 1992. doi: 10.1063/1.860046. URL <http://link.aip.org/link/?PFB/4/1935/1>.
- [31] A. Pukhov and J. Meyer-ter-Vehn. Laser wake field acceleration: the highly non-linear broken-wave regime. *Applied Physics B: Lasers and Optics*, 74:355–361, 2002. doi: 10.1007/s003400200795.
- [32] W. Lu, C. Huang, M. Zhou, M. Tzoufras, F. S. Tsung, W. B. Mori, and T. Katsouleas. A nonlinear theory for multidimensional relativistic plasma wave wakefields. *Physics of Plasmas*, 13(5):–, 2006.
- [33] T. Katsouleas, S. Wilks, P. Chen, J. M. Dawson, and J. J. Su. BEAM LOADING IN PLASMA ACCELERATORS. *Part. Accel.*, 22:81–99, 1987.
- [34] John M. Dawson. Nonlinear electron oscillations in a cold plasma. *Phys. Rev.*, 113(2):383–387, Jan 1959. doi: 10.1103/PhysRev.113.383.
- [35] W. Lu. *Nonlinear Plasma Wakefield Theory and Optimum Scaling for Laser Wakefield Acceleration in the Blowout Regime*. PhD thesis, UCLA, 2006.
- [36] M. Tzoufras, W. Lu, F. S. Tsung, C. Huang, W. B. Mori, T. Katsouleas, J. Vieira, R. A. Fonseca, and L. O. Silva. Beam loading in the nonlinear regime of plasma-based acceleration. *Physical Review Letters*, 101(14):145002, 2008. doi: 10.1103/PhysRevLett.101.145002. URL <http://link.aps.org/abstract/PRL/v101/e145002>.
- [37] A Pukhov, S Gordienko, S Kiselev, and I Kostyukov. The bubble regime of laser–plasma acceleration: monoenergetic electrons and the scalability. *Plasma Physics and Controlled Fusion*, 46(12B):B179–B186, 2004. URL <http://stacks.iop.org/0741-3335/46/B179>.
- [38] G. Fubiani, E. Esarey, C. B. Schroeder, and W. P. Leemans. Beat wave injection of electrons into plasma waves using two interfering laser pulses. *Phys. Rev. E*, 70(1):016402, Jul 2004. doi: 10.1103/PhysRevE.70.016402.
- [39] H. Kotaki, S. Masuda, M. Kando, J. K. Koga, and K. Nakajima. Head-on injection of a high quality electron beam by the interaction of two laser pulses. *Physics of Plasmas*, 11(6):3296–3302, 2004. doi: 10.1063/1.1751171. URL <http://link.aip.org/link/?PHP/11/3296/1>.
- [40] B. Cros, C. Courtois, G. Matthieussent, A. Di Bernardo, D. Batani, N. Andreev, and S. Kuznetsov. Eigenmodes for capillary tubes with dielectric walls and ultraintense laser pulse guiding. *Phys. Rev. E*, 65(2):026405, Jan 2002. doi: 10.1103/PhysRevE.65.026405.
- [41] B. Cros, C. Courtois, G. Malka, G. Matthieussent, J.R. Marques, F. Dorchies, F. Amiranoff, S. Rebibo, G. Hamoniaux, N. Blanchot, and J.L. Miquel. Extending plasma accelerators: guiding with capillary tubes. *Plasma Science, IEEE Transactions on*, 28(4):1071–1077, Aug 2000. ISSN 0093-3813. doi: 10.1109/27.893291.
- [42] R. Abrams. Coupling losses in hollow waveguide laser resonators. *Quantum Electronics, IEEE Journal of*, 8(11):838–843, Nov 1972. ISSN 0018-9197.
- [43] S. Jackel, R. Burris, J. Grun, A. Ting, C. Manka, K. Evans, and J. Kosakowski. Channeling of terawatt laser pulses by use of hollow waveguides. *Optics Letters*, 20:1086–1088, May 1995.
- [44] A. G. R. Thomas, Z. Najmudin, S. P. D. Mangles, C. D. Murphy, A. E. Dangor, C. Kamperidis, K. L. Lancaster, W. B. Mori, P. A. Norreys, W. Rozmus, and K. Krushelnick. Effect of laser-focusing conditions on propagation and monoenergetic electron production in laser-wakefield accelerators. *Physical Review Letters*, 98(9):–, 2007.

- [45] A. J. Gonsalves. *Investigation of a Hydrogen-filled Capillary Discharge Waveguide for Laser-driven Plasma Accelerators*. D-phil, Oxford, 2006.
- [46] C. G. Durfee and H. M. Milchberg. Light pipe for high intensity laser pulses. *Phys. Rev. Lett.*, 71(15):2409–2412, Oct 1993. doi: 10.1103/PhysRevLett.71.2409.
- [47] T. R. Clark and H. M. Milchberg. Time- and space-resolved density evolution of the plasma waveguide. *Phys. Rev. Lett.*, 78(12):2373–2376, Mar 1997. doi: 10.1103/PhysRevLett.78.2373.
- [48] C. G. R. Geddes, Cs. Tóth, J. van Tilborg, E. Esarey, C. B. Schroeder, D. Bruhwiler, C. Nieter, J. Cary, and W. P. Leemans. Production of high-quality electron bunches by dephasing and beam loading in channeled and unchanneled laser plasma accelerators. *Physics of Plasmas*, 12(5):056709, 2005. doi: 10.1063/1.1882352. URL <http://link.aip.org/link/?PHP/12/056709/1>.
- [49] C. G. R. Geddes, Cs. Toth, J. van Tilborg, E. Esarey, C. B. Schroeder, J. Cary, and W. P. Leemans. Guiding of relativistic laser pulses by preformed plasma channels. *Physical Review Letters*, 95(14):145002, 2005. doi: 10.1103/PhysRevLett.95.145002. URL <http://link.aps.org/abstract/PRL/v95/e145002>.
- [50] P. Volfbeyn, E. Esarey, and W. P. Leemans. Guiding of laser pulses in plasma channels created by the ignitor-heater technique. volume 6, pages 2269–2277. AIP, 1999. doi: 10.1063/1.873503. URL <http://link.aip.org/link/?PHP/6/2269/1>.
- [51] A. G. R. Thomas, C. D. Murphy, S. P. D. Mangles, A. E. Dangor, P. Foster, J. G. Gallacher, D. A. Jaroszynski, C. Kamperidis, K. L. Lancaster, P. A. Norreys, R. Viskup, K. Krushelnick, and Z. Najmudin. Monoenergetic electronic beam production using dual collinear laser pulses. *Physical Review Letters*, 100(25):–, 2008.
- [52] Y. Ehrlich, C. Cohen, A. Zigler, J. Krall, P. Sprangle, and E. Esarey. Guiding of high intensity laser pulses in straight and curved plasma channel experiments. *Phys. Rev. Lett.*, 77(20):4186–4189, Nov 1996. doi: 10.1103/PhysRevLett.77.4186.
- [53] T. Palchan, D. Kaganovich, P. Sasorov, P. Sprangle, C. Ting, and A. Zigler. Electron density in low density capillary plasma channel. *Applied Physics Letters*, 90(6):061501, 2007. doi: 10.1063/1.2472525. URL <http://link.aip.org/link/?APL/90/061501/1>.
- [54] D.J. Spence and S.M. Hooker. Simulations of the propagation of high-intensity laser pulses in discharge-ablated capillary waveguides. *Lasers and Electro-Optics, 2000. (CLEO 2000). Conference on*, pages 293–294, 2000. doi: 10.1109/CLEO.2000.907031.
- [55] T. Kameshima, W. Hong, K. Sugiyama, X. Wen, Y. Wu, C. Tang, Q. Zhu, Y. Gu, B. Zhang, H. Peng, S.-i. Kurokawa, L. Chen, T. Tajima, T. Kumita, and K. Nakajima. 0.56 GeV Laser Electron Acceleration in Ablative-Capillary-Discharge Plasma Channel. *Applied Physics Express*, 1(6):066001–+, June 2008.
- [56] N. C. Lopes, G. Figueira, L. O. Silva, J. M. Dias, R. Fonseca, L. Cardoso, C. Russo, C. Carias, G. Mendes, J. Vieira, and J. T. Mendonça. Plasma channels produced by a laser-triggered high-voltage discharge. *Phys. Rev. E*, 68(3):035402, Sep 2003. doi: 10.1103/PhysRevE.68.035402.
- [57] R.A. Bendoyro, R.I. Onofrei, J. Sampaio, R. Macedo, G. Figueira, and N.C. Lopes. Plasma channels for electron accelerators using discharges in structured gas cells. *Plasma Science, IEEE Transactions on*, 36(4):1728–1733, Aug. 2008. ISSN 0093-3813. doi: 10.1109/TPS.2008.927145.
- [58] D. J. Spence, A. Butler, and S. M. Hooker. First demonstration of guiding of high-intensity laser pulses in a hydrogen-filled capillary discharge waveguide. *Journal of Physics B-Atomic Molecular and Optical Physics*, 34(21):4103–4112, 2001.
- [59] D. J. Spence and S. M. Hooker. Investigation of a hydrogen plasma waveguide. *Physical Review E*, 6302(2):–, 2001.

- [60] N. A. Bobrova, A. A. Esaulov, J. I. Sakai, P. V. Sasorov, D. J. Spence, A. Butler, S. M. Hooker, and S. V. Bulanov. Simulations of a hydrogen-filled capillary discharge waveguide. *Physical Review E*, 65(1):–, 2002.
- [61] S.I. Braginskii. Transport processes in plasma. In M.A. Leontovich, editor, *Review of Plasma Physics*, volume 1 of *Review of Plasma Physics*, pages 201–311. Consultants Bureau, New York, U.S.A., 1965.
- [62] B. H. P. Broks, K. Garloff, and J. J. A. M. van der Mullen. Nonlocal-thermal-equilibrium model of a pulsed capillary discharge waveguide. *Physical Review E*, 71(1):–, 2005.
- [63] B. H. P. Broks, W. van Dijk, and J. J. A. M. van der Mullen. Parameter study of a pulsed capillary discharge waveguide. *Journal of Physics D-Applied Physics*, 39(11):2377–2383, 2006.
- [64] A. J. Gonsalves, T. P. Rowlands-Rees, B. H. P. Broks, J. J. A. M. van der Mullen, and S. M. Hooker. Transverse interferometry of a hydrogen-filled capillary discharge waveguide. *Physical Review Letters*, 98(2):–, 2007.
- [65] A. Butler. ??? D-phil, Oxford, 2004.
- [66] A. Butler, A. J. Gonsalves, C. M. McKenna, D. J. Spence, S. M. Hooker, S. Sebban, T. Mocek, I. Bettiabi, and B. Cros. Demonstration of a collisionally excited optical-field-ionization xuv laser driven in a plasma waveguide. *Physical Review Letters*, 91(20):–, 2003.
- [67] S. P. D. Mangles, A. G. R. Thomas, M. C. Kaluza, O. Lundh, F. Lindau, A. Persson, Z. Najmudin, C. G. Wahlstrom, C. D. Murphy, C. Kamperidis, K. L. Lancaster, E. Divall, and K. Krushelnick. Effect of laser contrast ratio on electron beam stability in laser wakefield acceleration experiments. *Plasma Physics and Controlled Fusion*, 48(12B):B83–B90, 2006.
- [68] S. P. D. Mangles. *Measurements of Relativistic Electrons from Intense Laser-Plasma Interactions*. PhD thesis, University of London, 2005.
- [69] C. Kamperidis. *Investigation of Electron Laser Wakefield Acceleration in Novel Plasma Structures*. PhD thesis, University of London, 2008.
- [70] Richards.Wh. Bayesian-based iterative method of image restoration. *Journal of the Optical Society of America*, 62(1):55, 1972.
- [71] L. B. Lucy. Iterative technique for rectification of observed distributions. *Astronomical Journal*, 79(6):745–754, 1974.
- [72] J. L. Starck, E. Pantin, and F. Murtagh. Deconvolution in astronomy: A review. *Publications of the Astronomical Society of the Pacific*, 114(800):1051–1069, 2002.
- [73] D.J.C. MacKay. *Information Theory, Inference, and Learning Algorithms*. Cambridge University Press, 2003.
- [74] W. Leemans, E. Esarey, C. Geddes, C. Schroeder, and C. Toth. Laser guiding for gev laser-plasma accelerators. *Philosophical Transactions of the Royal Society a-Mathematical Physical and Engineering Sciences*, 364(1840):585–600, 2006.
- [75] A. G. R. Thomas, S. P. D. Mangles, Z. Najmudin, M. C. Kaluza, C. D. Murphy, and K. Krushelnick. Measurements of wave-breaking radiation from a laser-wakefield accelerator. *Physical Review Letters*, 98(5):–, 2007.
- [76] F. S. Tsung, R. Narang, W. B. Mori, C. Joshi, R. A. Fonseca, and L. O. Silva. Near-gev-energy laser-wakefield acceleration of self-injected electrons in a centimeter-scale plasma channel. *Physical Review Letters*, 93(18):–, 2004.
- [77] P. Mora and T. M. Antonsen. Kinetic modeling of intense, short laser pulses propagating in tenuous plasmas. *Physics of Plasmas*, 4(1):217–229, 1997.


- [78] A. Ting, E. Esarey, and P. Sprangle. Nonlinear wake-field generation and relativistic focusing of intense laser-pulses in plasmas. *Physics of Fluids B-Plasma Physics*, 2(6):1390–1394, 1990.
- [79] M. Tzoufras.
- [80] B. H. P. Broks, W. Van Dijk, J. J. A. W. van der Mullen, A. J. Gonsalves, T. P. Rowlands-Rees, and S. M. Hooker. Modeling of a square pulsed capillary discharge waveguide for interferometry measurements. *Physics of Plasmas*, 14(2):–, 2007.
- [81] C. G. Durfee III, J. Lynch, and H. M. Milchberg. Mode properties of a plasma waveguide for intense laser pulses. *Optics Letters*, 19(23):1937–1939, 1994. URL <http://www.opticsinfobase.org/abstract.cfm?URI=OL-19-23-1937>.
- [82] T. P. Rowlands-Rees, C. Kamperidis, S. Kneip, A. J. Gonsalves, S. P. D. Mangles, J. G. Gallacher, E. Brunetti, T. Ibbotson, C. D. Murphy, P. S. Foster, M. J. V. Streeter, F. Budde, P. A. Norreys, D. A. Jaroszynski, K. Krushelnick, Z. Najmudin, and S. M. Hooker. Laser-driven acceleration of electrons in a partially ionized plasma channel. *Physical Review Letters*, 100(10):–, 2008.
- [83] E. Oz, S. Deng, T. Katsouleas, P. Muggli, C. D. Barnes, I. Blumenfeld, F. J. Decker, P. Emma, M. J. Hogan, R. Ischebeck, R. H. Iverson, N. Kirby, P. Krejcik, C. O’Connell, R. H. Siemann, D. Walz, D. Auerbach, C. E. Clayton, C. Huang, D. K. Johnson, C. Joshi, W. Lu, K. A. Marsh, W. B. Mori, and M. Zhou. Ionization-induced electron trapping in ultrarelativistic plasma wakes. *Physical Review Letters*, 98(8):–, 2007.
- [84] B. B. Wang, X. F. Li, P. M. Fu, J. Chen, and J. Liu. Coulomb potential recapture effect in above-barrier ionization in laser pulses. *Chinese Physics Letters*, 23(10):2729–2732, 2006.
- [85] J. Gorlinger, L. Plagne, and H. J. Kull. Above-barrier ionization and quantum interference in strong laser fields. *Applied Physics B-Lasers and Optics*, 71(3):331–336, 2000.
- [86] K. Mishima, K. Nagaya, M. Hayashi, and S. H. Lin. Theoretical studies of high-power laser ionization of molecules in the tunneling region. *Phys. Rev. A*, 70(6):063414, Dec 2004. doi:10.1103/PhysRevA.70.063414.
- [87] A Talebpour, S Larochele, and S L Chin. Suppressed tunnelling ionization of the molecule in an intense ti:sapphire laser pulse. *Journal of Physics B: Atomic, Molecular and Optical Physics*, 31(2):L49–L58, 1998. URL <http://stacks.iop.org/0953-4075/31/L49>.
- [88] J. J. Thomson. Bakerian Lecture: Rays of Positive Electricity. *Royal Society of London Proceedings Series A*, 89:1–20, August 1913.
- [89] B. J. McCall. *Spectroscopy of trihydrogen(+) in laboratory and astrophysical plasmas*. PhD thesis, AA(Department of Chemistry and Department of Astronomy and Astrophysics, University of Chicago, 5735 S. Ellis Ave., Chicago, IL 60637; jEMAIL;bjmccall@h3plus.uchicago.edu; jEMAIL;), 2001.
- [90] W. B. Mori and T. Katsouleas. Ponderomotive force of a uniform electromagnetic-wave in a time-varying dielectric medium. *Physical Review Letters*, 69(24):3495–3498, 1992.
- [91] R. Trines.
- [92] P. M. Bellan. *Fundamentals of Plasma Physics*. Cambridge University Press, 1st edition, 2006.
- [93] P. H. Rutherford R. J. Goldston. *Introduction to Plasma Physics*. IOP, 1995.
- [94] A. Bhattacharjee D. A. Gurnett. *Introduction to Plasma Physics*. Cambridge University Press, 2005.

- [95] S. Kneip, S. R. Nagel, C. Bellei, N. Bourgeois, A. E. Dangor, A. Gopal, R. Heathcote, S. P. D. Mangles, J. R. Marquès, A. Maksimchuk, P. M. Nilson, K. Ta Phuoc, S. Reed, M. Tzoufras, F. S. Tsung, L. Willingale, W. B. Mori, A. Rousse, K. Krushelnick, and Z. Najmudin. Observation of synchrotron radiation from electrons accelerated in a petawatt-laser-generated plasma cavity. *Physical Review Letters*, 100(10):105006, 2008. doi: 10.1103/PhysRevLett.100.105006. URL <http://link.aps.org/abstract/PRL/v100/e105006>.
- [96] D. Neely, K. Krushelnick, S. Mangles, R. Clarke, J. Collier, A.E. Dangor, C.N. Danson, R.D. Edwards, S. Fritzler, A. Gopal, P. Hatton, R. Heathcote, C. Hernandez-Gomez, K.W.D. Ledingham, P. McKenna, Z. Najmudin, B. Walton, and M. Wei. Particle acceleration using the vulcan petawatt laser. *Lasers and Electro-Optics, 2003. CLEO '03. Conference on*, pages 1 pp.–, June 2003.
- [97] V V Ivanov, K N Koshelev, E S Toma, and F Bijkerk. Influence of an axial magnetic field on the density profile of capillary plasma channels. *Journal of Physics D: Applied Physics*, 36(7): 832–836, 2003. URL <http://stacks.iop.org/0022-3727/36/832>.
- [98] *CRC Handbook of Chemistry and Physics*. CRC, 84 edition, 2003.

DIFFRACTION AFFECTING YOUR RESEARCH?

UNABLE TO MAINTAIN A WAKEFIELD FOR MORE THAN 5mm?

EMBARRASSED TO TALK ABOUT YOUR RESULTS WITH PEERS?



WAVEGUIDE TM

SCIENTIFICALLY PROVEN. NO NEED FOR BIG EXPENSIVE LASERS. 80% SAFE.

CALL 01865 272 301