Spatiotemporal Optical Injection in Laser-Plasma Accelerators



Emily Archer

Wadham College

University of Oxford

A thesis submitted for the degree of
Doctor of Philosophy
at the University of Oxford

Hilary 2024

Abstract

Laser-driven plasma wakefield accelerators (LWFAs) can accelerate charged particles to multi-gigaelectron volt energies over several centimetres due to the strong electric fields sustained in a plasma. Compared to conventional machines, which are often metres to kilometres-long, plasma-based accelerators have the potential to be the next generation of accelerators for particle physics research, in addition to providing convenient and cost-effective sources of high-energy electrons and radiation for applications in medicine, biology, chemistry and security. To be viable, there are challenges that remain to be addressed, particularly regarding the methods used to inject the witness electrons to be accelerated.

Controlled injection techniques aim to reliably generate electron bunches with desirable and repeatable bunch properties over long-term operation, by controlling how electrons from the plasma enter the wakefield. Optical injection techniques trigger injection using additional laser pulses.

In this thesis, we explore five different optical injection techniques. First, we compare three different effects arising from spatiotemporal couplings in ultrashort laser pulses. We show that when a radially-chirped laser pulse is focused by an achromatic lens, it forms a simultaneous space-time focus which exhibits a controllable group velocity, like a flying focus, and that unlike its linearly-chirped counterpart, it does not exhibit pulse-front tilt at the focus. We show that simultaneous space-time foci offer no higher degree of axial localisation of the intensity, contrary to as is often stated.

Accounting for novel polarisation states of the laser pulse, we present a new derivation of the wakefield quantities, to find that radially-polarised laser beams (RPLBs) may enhance electron trapping in LWFAs. We show that RPLBs have a higher degree of axial localisation compared to their linearly-polarised Gaussian counterparts.

Finally, we demonstrate injection and acceleration of $\sim 10-100\,\mathrm{pC}$, 600 MeV electron bunches using optically-generated shocks. This scheme may be promising for future light-source applications, due to the betatron oscillations that seem to be induced in a crossed-channel configuration.

Acknowledgements

Firstly, to Simon. Thank you for being such a wonderful supervisor and inspiring academic. Thank you for allowing me the autonomy as a researcher to pursue the topics that interested me, for being patient and enthused, and for always sending me your thoughts and derivations. You have been endlessly supportive and have given me a variety of opportunities that I would not have had with anyone else. The positive impact that you have had on my DPhil experience is completely undeniable. Roman, thank you for all of your brilliant insights.

To James and Linus, the ultimate post-doc tag team. Thank you for keeping me sane and for being so reassuring. To James Cowley, Alex, Aimee and Warren for showing me the ropes, particularly on RAL experiments. To everyone else in the Oxford LPAG group and our collaborators at Munich for making the inevitable ups and downs of experimental work far more bearable.

To the friends that are bound to me by collective experience. Joe and Cameron, there at the start, there at the end. I adore you both and I don't know what I will do without my martini-drinking, Vic-loving companions. Thanks for the manic and delirious evening escapades. Manuela, my serendipitous friend, you have been there through thick and thin (and library days that stretched forever). To Hank, thank you for being my only college friend, and a great one at that. Abigail, thank you for the sci-fi movie nights and long walks in the park. Thank you, Owen, for everything that you did for me; for 'Headstrong' when I was most in need of inspiration. Cornelia and Annabel, our conversations have brought me so much solace and insight, you are both incredible physicists.

To the friends that have supported me. Hannah and Adam, this is a formal apology for being so terrible at pub quizzes, thank you for all the fun and wholesome evenings. Eleanor, you bring so much joy and light to my life, thank you for always being there for me. Aimee, thank you for being somehow both a blissfully calm presence and also insanely inspiring. To my oldest and dearest trio of friends: Megan, Anna and Lauren. Thank you for always looking out for me over these 20+ years of friendship, I am very lucky to have you.

Finally, and most importantly, to my family. To Grandad, thank you for always sending kingfishers during long phone call walks along the canal. Nannie and Grandad, I love you lots. Mum, Dad and Hazel, you are my whole world and I love you so much. Thank you for always treating the things that matter to me as important.

Declaration

I, Emily Archer, do hereby certify that this project report has been written by me, is a record of work carried out by me, and is essentially different from work undertaken for any other purpose or assessment.

Date:

June 24, 2024

Signature of Candidate:



Contents

List of Figures **xi**

List of Tables **xiii**

1 Introduction **3**
1.1 Historical Overview and Motivation . . . 3
1.2 Laser-Driven Plasma Accelerators . . . 5
1.2.1 Basic Principles . . . 6
1.2.2 Operating Regimes . . . 7
1.3 Challenges . . . 9
1.3.1 Properties of the Electron Bunch . . . 9
1.3.2 Stability of the Electron Bunch Properties . . . 13
1.3.3 Accelerator Efficiency and Footprint . . . 13
1.4 Key Concepts . . . 14
1.4.1 Plasma Waveguides . . . 14
1.4.2 Electron Injection . . . 21
1.5 Thesis Outline . . . 24

2 Background Theory and Methods **27**
2.1 Propagation of Light in a Vacuum . . . 27
2.1.1 Wave and Geometric Optics . . . 27
2.1.2 Complex Representation of Electric Fields . . . 29
2.1.3 Scalar Wave Equation . . . 31
2.1.4 Paraxial Wave Equation . . . 34
2.1.5 Higher-Order Modes . . . 37
2.1.6 Modified Paraxial Wave Equation . . . 41
2.2 Spatiotemporal Couplings . . . 43
2.3 Collins Diffraction Integral . . . 44
2.3.1 Detailed Discussion of the Simulation Grids . . . 47
2.3.2 Comparison Between Collins and Analytic Methods . . . 51

3 The Flying Focus **53**
3.1 Introduction . . . 53
3.2 Theoretical Analysis of the FF . . . 55
3.2.1 Definition of Input bi-Gaussian Pulse . . . 55
3.2.2 Chromatic Lens . . . 56
3.2.3 Derivation of the Gaussian Flying Focus . . . 57
3.2.4 Derivation of the Focal Velocity . . . 58
3.3 Ray Matrix to Generate an FF . . . 61
3.4 Results . . . 62
3.4.1 Alternative Spectra: Super-Gaussian . . . 65

3.5 Conclusions and Future Work 66

4 The Linear Simultaneous Space-Time Focus 67
4.1 Introduction 67
4.2 Theoretical Analysis of the SSTF 70
4.2.1 Derivation of the Linear Simultaneous Space-Time Focus 72
4.3 Ray Matrix to Generate an SSTF 76
4.4 Results 78
4.4.1 Centroid, Ellipsoidal Spot Size, Bandwidth and Pulse Duration 79
4.4.2 Calculation of Peak Axial Intensity 81
4.4.3 Comparison for Equivalent Spatial Focusing 84
4.5 Conclusions and Future Work 88

5 The Radial Simultaneous Space-Time Focus 89
5.1 Introduction 89
5.2 Theoretical Analysis of the R-SSTF 90
5.2.1 Derivation of the Radial Simultaneous Space-Time Focus 91
5.2.2 Propagating Arbitrary Beams as a Decomposition of Laguerre-Gaussian Modes 94
5.3 Ray Matrix to Generate an R-SSTF 99
5.4 Results 101
5.4.1 Centroid, Ellipsoidal Spot Size, Bandwidth and Pulse Duration 102
5.4.2 Calculation of Peak Axial Intensity 104
5.5 Controllable Group Velocity 107
5.6 Conclusions and Future Work 108

6 The Radially-Polarised Focus 111
6.1 Introduction 111
6.2 Electromagnetic Fields of Different Polarisation States 114
6.2.1 Linearly-Polarised bi-Gaussian Beams 115
6.2.2 Radially-Polarised Laser Beams 117
6.3 Confinement of Intensity for Different Polarisation States 120
6.4 Group Velocity of Different Polarisation States 121
6.5 Ionisation and Trapping in Laser-Driven Wakefields 122
6.5.1 Trapping with Linearly-Polarised Laser Beams 122
6.5.2 Trapping with Radially-Polarised Laser Beams 126
6.6 Results 142
6.7 Conclusions and Future Work 144

7 Optical Injection with Crossed HOFI Channels 147
7.1 Introduction 147
7.1.1 Motivation 147
7.1.2 Objectives 148
7.2 Experimental Design 149
7.2.1 Astra-Gemini 149
7.3 Gas Jet Target 153
7.4 Diagnostics 155
7.4.1 Pulse Duration 155
7.4.2 Focal Spots 156
7.4.3 Guiding Diagnostic 163
7.4.4 Plasma Fluorescence 164
7.4.5 Transverse Interferometry 166

7.4.6 Electron Spectrometer 169
7.5 Results 178
7.5.1 Correlations Between Variables 180
7.5.2 Guiding 181
7.5.3 Electron Injection and Acceleration 183
7.5.4 Electron Betatron 186
7.6 Conclusions and Future Work 188

8 Conclusion **191**

References **195**

List of Figures

2.1 Comparison of numerically and analytically Fourier-transformed spectral and temporal profiles 50

3.1 Schematic of a flying focus 54
3.2 Schematic of a bi-Gaussian pulse 56
3.3 Plot of v_{ff} for different values of χ and β 61
3.4 Optical system and coordinate convention for generating a flying focus 62
3.5 Snapshots of a flying focus in the light-speed frame ξ 64
3.6 Propagation of flying foci with arbitrary spectra 65

4.1 Schematic of a linear simultaneous space-time focus 68
4.2 Input pulse for generating a linear simultaneous space-time focus 71
4.3 Optical system and coordinate convention for generating a linear simultaneous space-time focus 73
4.4 Evolution of a linear simultaneous space-time focus at different axial z positions . 77
4.5 Key parameters of the linear SSTF pulse as a function of z in the focal region . . . 79
4.6 Longitudinal variation of the axial peak intensity of the linear SSTF pulse calculated using two different methods 81
4.7 Schematic showing the issue with assuming an ellipsoidal focal volume 84
4.8 Spatiospectral profiles of the pre-chirp, SSTF and non-SSTF fair comparison pulses 85
4.9 Longitudinal variation of the peak intensity of the linear SSTF pulse calculated using the Collins method 87

5.1 Schematic of a radial simultaneous space-time focus 90
5.2 Input pulse for generating a radial simultaneous space-time focus 92
5.3 Optical system and 3D cylindrical coordinate convention for generating a radial simultaneous space-time focus 93
5.4 Transverse intensity profiles for the Laguerre-Gaussian modes with increasing radial index p 95
5.5 Amplitude $|c_p(\omega)|$ of the different modes with radial index p for the input R-SSTF shown in Figures 5.2a)–d), where $\alpha = -1$. 96
5.6 Evolution of a radial simultaneous space-time focus at different axial z positions . 100
5.7 Key parameters of the radial SSTF pulse as a function of z in the focal region . . . 102
5.8 Longitudinal variation of the axial peak intensity of the radial SSTF pulse 104
5.9 Snapshots of a radial simultaneous space-time focus in the light-speed frame ξ . . 106

6.1 Electric field components for a linearly-polarised Gaussian beam 116
6.2 Electric field components for a radially-polarised Gaussian beam 118
6.3 Ratios of the on-axis fields 119
6.4 Comparison of the longitudinal and transverse variation of intensity of a RPLB and a LPLB for an $f/2$ focusing geometry 119
6.5 Lineouts of the intensity at $z = 0$ and $x = 0$ for an $f/2$ focusing geometry 121

6.6 Calculated quasistatic wakefield quantities as a function of ξ/λ_p for a standard 2PII scheme . . . 123
6.7 Schematic of the R-2PII scheme with a radially-polarised trailing pulse . . . 126
6.8 Calculated quasistatic wakefield quantities as a function of ξ/λ_p for a R-2PII scheme 141
6.9 Calculated quasistatic wakefield quantities as a function of ξ/λ_p for a wakefield driven by a RPLB . . . 143

7.1 Simplified schematic of the experimental layout and the measured focal spot data 150
7.2 Schematic of the full experimental layout . . . 152
7.3 Gas jet target and the correlations between backing pressure and fluorescence intensity . . . 153
7.4 Gas jet density at 3 mm above the jet . . . 154
7.5 Pulse duration estimated from the FROG measurement . . . 155
7.6 Variation in the centroid position and spot size of the drive beam . . . 157
7.7 Variation in the centroid position and spot size of the channel-forming beam . . . 157
7.8 Variation in the centroid position and spot size of the injector beam . . . 158
7.9 Focal scan and beam quality analysis of the drive beam . . . 161
7.10 Schematic of the guiding diagnostic . . . 162
7.11 Scatter plot of the guided mode D4σ spot size against the counts on guided mode camera . . . 163
7.12 Plot of the output spot size w as a function of the matched spot size w_m . . . 164
7.13 Fluorescence measurement to determine the start position of the axicon plasma . 165
7.14 Schematic of the transverse interferometry diagnostic . . . 166
7.15 Exemplary analysis of the data from the transverse interferometry diagnostic . . . 167
7.16 Schematic of the electron spectrometer with electron trajectories and calibration curves . . . 170
7.17 Map of Pearson's correlation coefficients between experimental parameters and measured variables . . . 179
7.18 Key results from analysis of the data from the guiding diagnostic . . . 181
7.19 Pointing jitter of the guided beam for different channel combinations . . . 182
7.20 Exemplary electron spectrometer data for different channel conditions . . . 184
7.21 Key results from analysis of the electron spectrometer data . . . 185
7.22 Exemplary electron spectrometer data showing energy-divergence oscillations . . 187

List of Tables

3.1 Table of laser pulse parameters for FF simulations 62

4.1 Table of laser pulse parameters for SSTF simulations 78

5.1 Table of laser pulse parameters for R-SSTF simulations 101

6.1 Table of laser pulse parameters for RPLB simulations 115
6.2 Table of relevant timescales for a quasistatic approximation 136

7.1 Table of electron spectrometer calibrations for each screen 173
7.2 Comparison of Astra-Gemini laser pulse parameters with Glinec *et al.* 176

List of Symbols

Abbreviations

2PII	Two Pulse Ionisation Injection
ADK	Ammosov-Delone-Kraĭnov
AM	Annular Mirror
BSI	Barrier Suppression Ionisation
CCI	Crossed Channel Injection
CEP	Carrier-Envelope Phase
D-FFT	Double Fast Fourier Transform
DLA	Direct Laser Acceleration
DM	Deformable Mirror
EFV	Ellipsoidal Focal Volume
FFT	Fast Fourier Transform
FF	Flying Focus
FROG	Frequency-Resolved Optical Gating
FWHM	Full Width Half Maximum
HG	Hermite-Gaussian
HOFI	Hydrodynamic Optical Field Ionisation
IFFT	Inverse Fast Fourier Transform
L-2PII	Linear Two Pulse Ionisation Injection
LG	Laguerre-Gaussian
LHC	Large Hadron Collider
LPLB	Linearly-Polarised Laser Beam
LSE	Least-Squares Error
LWFA	Laser-Driven Plasma Wakefield Accelerator
OAM	Orbital Angular Momentum

OFI	Optical Field Ionisation
OTBI	Over-the-Barrier Ionisation
PBS	Polarising Beam Splitter
PBWA	Plasma Beat-Wave Accelerator
PFT	Pulse-Front Tilt
PIC	Particle-in-Cell
PI	Peak Intensity
PWFA	Beam-Driven Plasma Wakefield Accelerator
QWP	Quarter-Wave Plate
R-2PII	Radial Two Pulse Ionisation Injection
R-SSTF	Radial Simultaneous Space-Time Focusing
RF	Radio Frequency
RMSE	Root-Mean-Square Error
RPLB	Radially-Polarised Laser Beam
S-FFT	Single Fast Fourier Transform
SM-LWFA	Self-Modulated Laser-Driven Plasma Wakefield Accelerator
SSTF	Simultaneous Space-Time Focusing
STC	Spatiotemporal Coupling
TI	Tunneling Ionisation
TPEF	Two-Photon Excitation Fluorescence
X-FEL	X-Ray Free-Electron Laser

Physical Constants

$\hbar$	Reduced Planck's Constant $\hbar = h/2\pi$	$1.05457 \times 10^{-34}\,\mathrm{J\,s}$
μ_0	Vacuum Permeability	$1.25663 \times 10^{-6}\,\mathrm{N\,m^{-2}}$
ε_0	Vacuum Permittivity	$8.85418 \times 10^{-12}\,\mathrm{F\,m^{-1}}$
c	Speed of Light	$2.99792 \times 10^{8}\,\mathrm{m\,s^{-1}}$
e	Electron Charge	$1.60217 \times 10^{-19}\,\mathrm{C}$
h	Planck's Constant	$6.62607 \times 10^{-34}\,\mathrm{J\,s}$
m_e	Electron Mass	$9.10938 \times 10^{-31}\,\mathrm{kg}$
r_Bohr	Bohr Radius	$5.29177 \times 10^{-11}\,\mathrm{m}$
r_e	Classical Electron Radius	$2.81794 \times 10^{-15}\,\mathrm{m}$

Chemical Formulae

Al_2O_3	Sapphire

Ar	Argon	
$Ba(BO_2)_2$	Barium Borate (BBO)	
$Gd_2O_2S : Tb$	Gadolinium Oxysulphide Phosphor	
H_2	Hydrogen	
Nb	Niobium	
Ne	Neon	
N_2	Nitrogen	
$Ti : Al_2O_3$	Titanium:Sapphire (Ti:Sa)	
Ti	Titanium	
Xe	Xenon	

General Variables

ϵ	Perturbation Expansion Parameter	
η	Refractive Index	
γ	Lorentz Factor	
$\mathbf{A}$	Vector Potential	$\mathrm{T\,m}$
$\mathbf{a}$	Normalised Vector Potential	
$\mathbf{B}$	Magnetic Field	T
$\mathbf{E}$	Electric Field	$\mathrm{V\,m^{-1}}$
$\mathbf{j}$	Current Density	$\mathrm{A\,m^{-3}}$
$\mathbf{p}$	Generalised Momentum Vector	$\mathrm{kg\,m\,s^{-1}}$
$\mathbf{r}$	Generalised Coordinate Vector	m
$\mathbf{S}$	Poynting Vector	$\mathrm{W\,m^{-2}}$
$\mathbf{v}$	Generalised Velocity Vector	$\mathrm{m\,s^{-1}}$
$\mathscr{H}$	Hamiltonian	
$\mathscr{H}_s$	Separatrix Hamiltonian	
$\mathscr{L}$	Lagrangian	
$\mathfrak{H}$	Optical Transfer Function	$\mathrm{m\,s^{-1}}$
μ	Permeability	$\mathrm{N\,m^{-2}}$
$\Omega, \Pi, \Psi, \Theta, \Upsilon, \Xi$	Substitution Variables	
Φ	Scalar Potential	V
ϕ	Normalised Scalar Potential	
ρ, θ, z	Cylindrical Coordinate System	
ρ	Charge Density	$\mathrm{C\,m^{-3}}$

ε	Permittivity	$\mathrm{F\,m^{-1}}$
ξ	Moving-Frame Coordinate	
A, B, C, D	Optical Ray Matrix Elements	
k^i	Generalised Wave Four Vector	$\mathrm{m^{-1}}$
t	Time	s
u	Field Energy Density	
x, y, z	Cartesian Coordinate System	
x^i	Generalised Coordinate Four Vector	m
Laser Parameters		
α	Spatial Chirp	
α_{crit}	Spatial Chirp Normalisation Factor	$\mathrm{rad\,s^{-1}\,m^{-1}}$
β	Lens Dispersion	
β_{crit}	Lens Dispersion Normalisation Factor	
χ	Normalised Temporal Chirp	
Δx	Lateral Offset of Spectral Component	µm
$\Delta\omega$	Spectral Bandwidth	$\mathrm{rad\,s^{-1}}$
$\Delta\rho$	Radial Offset of Spectral Component	m
ℓ	Azimuthal Index	
Γ	Gaussian Spectral Parameter $\Gamma = a + ib$	
λ	Wavelength	nm
λ_0	Central Wavelength	nm
$\lvert\tilde{E}(\omega)\rvert$	Spectral Amplitude	$\mathrm{V\,m^{-1}}$
$\lvert E(t)\rvert$	Temporal Amplitude	$\mathrm{V\,m^{-1}}$
$\mathbf{k}$	Angular Wavevector	$\mathrm{m^{-1}}$
$\mathscr{E}$	Pulse Energy	J
ω	Angular Frequency	$\mathrm{rad\,s^{-1}}$
ω_0	Central Angular Frequency	$\mathrm{rad\,s^{-1}}$
ψ	Temporal Phase or Eikonal	rad
ψ_0	Temporal Phase Offset	rad
τ_{E}	Laser Evolution Time	fs
τ_{l}	Oscillation Period	fs
τ_{p0}	Unchirped Pulse Duration	fs
τ_{p}	Pulse Duration	fs

Θ_0	Gaussian Full Divergence Angle	rad
θ_0	Gaussian Half Divergence Angle	rad
$\tilde{E}(\omega)$	Spectral Field	$\mathrm{V\,m^{-1}}$
$\tilde{U}(\mathbf{r},\omega)$	Spectral Signal	$\mathrm{V\,m^{-1}}$
$\tilde{U}(\omega)$	Spectrum	$\mathrm{rad^{-1}\,s}$
φ	Spectral Phase	rad
φ_{G}	Guoy Phase	rad
$c_{p,\ell}$	Relative Mode Amplitude	
$E(t)$	Temporal Field	$\mathrm{V\,m^{-1}}$
F	Fluence	$\mathrm{J\,cm^{-2}}$
f	Lens Focal Length	m
f_0	Lens Focal Length for Central Wavelength	m
I	Intensity	$\mathrm{W\,cm^{-2}}$
k	Angular Wavenumber	$\mathrm{m^{-1}}$
p	Radial Index	
q	Gaussian Complex Parameter $q = z - z_0 - iz_{\mathrm{R}}$	m
q^+	Gaussian Inverted Complex Parameter $q = z - z_0 + iz_{\mathrm{R}}$	m
r_0	Annulus Radius for Central Wavelength	m
t_ω	Frequency-Dependent Arrival Time	s
$U(\mathbf{r},t)$	Temporal Signal	$\mathrm{V\,m^{-1}}$
$U(t)$	Temporal Profile	$\mathrm{s^{-1}}$
U_t	Temporal Normalisation Factor	$\mathrm{J\,s^{-1}}$
U_ω	Spectral Normalisation Factor	$\mathrm{J\,s^{-1}}$
v_{ff}	Flying Focus Focal Velocity	$\mathrm{m\,s^{-1}}$
$v_{\mathrm{g,L}}$	Group Velocity of Linearly-Polarised Beam	$\mathrm{m\,s^{-1}}$
$v_{\mathrm{g,R}}$	Group Velocity of Radially-Polarised Beam	$\mathrm{m\,s^{-1}}$
v_{g}	Group Velocity	$\mathrm{m\,s^{-1}}$
v_{ph}	Phase Velocity	$\mathrm{m\,s^{-1}}$
v_{sstf}	Simultaneous Space-Time Focal Velocity	$\mathrm{m\,s^{-1}}$
w	Gaussian Waist	m
w_0	Gaussian Waist at Focus	m
w_{m}	Matched Spot Size	m
z_0	Gaussian Waist Position	m

z_ω	Frequency-Dependent Focal Position	m
z_R	Gaussian Rayleigh Range	m

Mathematical Constants

e	Euler's Number	2.71828...
π	Pi	3.14159...
i	Imaginary Number $\sqrt{-1}$	

Mathematical Operators

$\boldsymbol{\nabla}$	Vector Gradient	m^{-1}
$\boldsymbol{\nabla}\cdot$	Divergence	m^{-1}
$\boldsymbol{\nabla}\times$	Curl	m^{-1}
$\hat{\{\}}$	Laser Envelope	
$\int\{\quad\}$	Integral	
$\langle\{\quad\}\rangle$	Time Average	
$\lvert\{\quad\}\rvert$	Absolute Value or Amplitude	
$\mathscr{F}\{\quad\}$	Fourier Transform	
$\mathscr{F}^{-1}\{\quad\}$	Inverse Fourier Transform	
∇	Scalar Gradient	m^{-1}
∇^2	Laplacian	m^{-2}
$\partial\{\quad\}/\partial$	Partial Differential	
$\rightarrow$	Tends To	
$\sqrt{\{\quad\}}$	Square Root	
$\arg\{\quad\}$	Argument $\tan\Im\{\quad\}/\Re\{\quad\}$	
$\mathrm{d}\{\quad\}/\mathrm{d}$	Differential	
$\max\{\quad\}$	Maximum	
$\tilde{\{\}}$	Fourier Conjugate Pair	
$\{\quad\}_\parallel$	Longitudinal Components	
$\{\quad\}_\perp$	Transverse Components	
$H\{\}$	Hilbert Transform Pair	

Plasma Parameters

β_p	Normalised Wake Phase Velocity $\beta_\mathrm{p} = v_\mathrm{p}/c$	
γ_k	Keldysh Parameter	
λ_p	Plasma Wavelength	µm
$\mathbf{F}_\mathrm{p}$	Ponderomotive Force	N

$\mathscr{E}_{\text{ion}}$	Ionisation Energy	$\mathrm{kJ\,mol^{-1}}$
$\mathscr{E}_{\text{k}}$	Kinetic Energy	J
$\mathscr{E}_{\text{p}}$	Ponderomotive Potential	J
ω_{p}	Langmuir or Plasma Frequency	$\mathrm{rad\,s^{-1}}$
E_{wb}	Wavebreaking Field Strength	$\mathrm{V\,m^{-1}}$
I_{app}	Ionisation Appearance Intensity	$\mathrm{W\,cm^{-2}}$
n_0	Equilibrium Plasma Electron Density	$\mathrm{cm^{-3}}$
n_{e}	Plasma Electron Density	$\mathrm{cm^{-3}}$
t'	Time in Retarded Frame	s
v_{p}	Wake Phase Velocity	$\mathrm{m\,s^{-1}}$

1 | Introduction

1.1 Historical Overview and Motivation

Over nearly a century of evolution, particle accelerators have become well-established as a linchpin tool for fundamental physics and modern life. Since the first generation of a potential difference by Wilcke in 1762[1] (popularised by Volta's electrophorus in 1775[2]), the acceleration of charged particles has served an essential role in numerous discoveries about the workings of our universe. Accelerators have developed in symbiosis with atomic and particle physics: from Rutherford's discovery of atomic structure in 1911[3] aided by Cockroft and Walton's generator,[4] to the discovery of the Higgs boson in 2012[5,6] at the Large Hadron Collider (LHC).[7] Alongside this, accelerators have also found many practical uses. With more than 40,000 particle accelerators around the world today,[8] uses range from medical imaging, radiotherapy and sterilisation of medical equipment, to the manufacture of semiconductors, synthesis of plastics and ceramics, radioisotope production and security.

Around the same time as Cockroft and Walton were developing their 800 kV generator,[4] Van de Graaff built his iconic electrostatic generator.[9] The first design in 1929 produced electrons with almost double the energy of Cockroft and Walton's. An alternative to direct voltage accelerators had already been proposed by Ising in 1924,[10] representing the first important shift in accelerator physics: from constant to time-varying electric fields. Ising's linear accelerator utilises the concept of resonant acceleration. Metal cavities are used to sustain a radio frequency (RF) electromagnetic field which oscillates at a constant frequency, the sign of the electric field flipping between accelerating and decelerating. A series of such cavities are spaced such that the particle continues to gain energy as it passes through the linear accelerator (linac), drifting between the cavities when the field is decelerating. The gaps between cavities must grow as the particle's speed increases, otherwise it would reach the next cavity too early and be decelerated. This

was never implemented by Ising, nonetheless, the paper was discovered by Widerøe in 1927 and he built a version of Ising's idea that could accelerate potassium and sodium ions from 25 keV to 50 keV. Rather than applying an alternating voltage across the cavities, Alvarez proposed applying the voltage across the whole machine and shielding the particles within 'drift tubes' during the decelerating portion of the RF cycle.[11] Protons were accelerated to a final energy of 31.5 MeV. Another approach was proposed by Lawrence,[12] due to the technical difficulties associated with building linacs at the time,[13] and was later realised by Livingston.[14] Moving from linear to circular geometries represents a second important shift in accelerator physics. A fixed-field cyclotron consists of two 'dees' with an RF driving electric field across the gap between them. Particles are steered on a 'spiral' trajectory (actually a series of concentric arcs) by a constant magnetic field applied perpendicular to the dees. As they gain energy, the radius of their trajectories increases, perfectly balancing the decrease in time taken to traverse the dee due to the increase in speed. This means that the particles always see an accelerating field as they cross the gap. The particles complete many revolutions, meaning that they can be accelerated to much higher energies than if they were to traverse the accelerating field only once. The problem with cyclotrons is that once the particles become relativistic, any natural energy spread in the beam will grow progressively as the spread in revolution frequencies leads to particles off-resonance with the accelerating field slipping further behind.[13] Phase stability, proposed by Veksler[15] and McMillan,[16] was a key concept for the realisation of synchrotrons. By picking a point in the RF cycle where the field is increasing with time (e.g. at a phase $0 < \phi < \pi/2$ for a sine curve), the particle beam is bunched together rather than spread out, since particles arriving later receive more energy from the RF field. In a synchrotron, the magnetic field ramps up in strength to keep the particles travelling along an orbit of constant radius as their speed increases. Similarly, the frequency of the RF accelerating fields must also ramp up with the increasing revolution frequency of the particles.

The extreme electric fields in radio frequency (RF) accelerating cavities can cause electrical breakdown as charges are pulled from the metallic walls, generating plasma and causing sparking within the cavity. Superconducting RF cavities can sustain higher electric field gradients without damage, however, there is still a fundamental limit to the electric field gradient that can be generated in RF accelerators, with 190 MV m^{-1} considered state of the art.[17] Even in this case, the accelerator has to be on the order of hundreds to thousands of meters long for the particle

beam to reach the desired GeV-level energies. In the case of a synchrotron, where the beam completes many revolutions, the field strength of the bending dipole magnets must increase to keep the particles on a trajectory of constant radius within the accelerator as their energy increases. To reach the highest energies of 7 TeV at the LHC, which has a bending radius of 2.805 km, the magnetic field strength must ramp up to a peak strength of 8.3 T, which necessitates cooling of the superconducting niobium-titanium (NbTi) magnets to 1.9 K.[18] Despite the continued development of superconducting RF cavities and magnets, it is clear that we are reaching an impasse posed by the fundamental technological limitations on the maximum energy that we can accelerate particles to in a vacuum. Not only does this motivate an exploration of novel technologies as a route to more compact accelerators, but it is also evidently vital if we want to build particle accelerators that can explore particle physics beyond the 13.6 TeV that is accessible today.

1.2 Laser-Driven Plasma Accelerators

In the most recent phase of their development, novel particle accelerators have shifted from operation in a vacuum to an exploration of controlled particle acceleration in the fourth state of matter: plasma. A plasma is an overall quasi-neutral ionised gas. The electrons and ions are no longer bound, such that they may independently exhibit collective effects under applied electromagnetic fields. Although it was in 1979 when Tajima and Dawson first proposed the laser-driven plasma accelerator as we know it today,[19] the idea of using a collection of charged particles or plasma as the accelerating medium can be traced back to Veksler[20] and Faĭnberg[21] in 1968 (the same Veksler of synchrotron phase stability).

Plasma-based accelerators, normally driven by either a laser pulse in a laser-driven plasma wakefield accelerator (LWFA) or by a particle beam in a beam-driven plasma wakefield accelerator (PWFA), can sustain accelerating gradients which are three orders of magnitude greater than most conventional accelerators, as high as $100\,\mathrm{GV\,m^{-1}}$. When a particle beam travels through a plasma, it creates a 'blow-out' plasma wakefield structure due to space-charge repulsion.[22] Although witness electrons which propagate in the beam-driven wakefield can gain significant energies in short distances, an energetic electron beam is still required to drive the wakefield.[23] Despite facing a unique set of challenges which will be discussed in Section 1.3, laser-driven plasma wakefield accelerators completely mitigate this necessity by using a laser to drive the wakefield instead. Electrons have been accelerated in laser-driven wakefields to energies of up

to 8 GeV on the scale of centimetres rather than kilometres.[24] The length of the accelerator required to reach the same energies in a plasma-based accelerator therefore offers the possibility for cheaper and more compact accelerators, which are ideal for universities and hospitals. Aside from radiotherapy and medical imaging, a key application is in driving x-ray free electron lasers (X-FELs). In the case of an LWFA, the production of ultra-short duration bunches which are synchronised to a femtosecond laser system has potential applications in ultra-fast science.

1.2.1 Basic Principles

In an LWFA, as a laser pulse propagates through a plasma, the ponderomotive force pushes electrons away from the region of the highest laser intensity. The ponderomotive force $\mathbf{F}_\mathrm{p}$ can be derived by taking the time-averaged motion of an electron in the electric field of a laser pulse $\mathbf{E}$ to give

$$\mathbf{F}_\mathrm{p} = -\frac{e^2\lambda_0^2}{16\pi^2 m_\mathrm{e} c^2}\nabla|\mathbf{E}|^2, \tag{1.1}$$

where λ_0 is the central wavelength of the laser pulse, m_e is the mass of an electron, e is the charge of an electron, and c is the speed of light.

Ions, which are more massive than electrons, are virtually unaffected by the ponderomotive force on the timescales we will consider in this thesis. As a result, the region behind the driving laser pulse, which is vacated of electrons, initially has a net positive charge due to the ions. This separation of charges causes a restoring force which pulls the electrons back towards the axis of propagation of the laser pulse; a collective oscillation which forms a longitudinal density wave behind the laser pulse. The frequency at which this oscillation occurs is characterised by the plasma frequency

$$\omega_\mathrm{p} = \sqrt{\frac{n_\mathrm{e} e^2}{m_\mathrm{e}\varepsilon_0}}, \tag{1.2}$$

where ε_0 is the permittivity of free-space and n_e is the number electron density. Somewhat easier to visualise is the action of a boat creating waves or a 'wake' as it travels through water. The peaks and troughs in the electron density generate large electric fields which can accelerate charged particles that have been trapped at the right position in the wakefield; much as a wakeboarder has to sit just ahead of the crest of a water wave to be pulled along with the boat. An estimate of the maximum accelerating field that can be sustained is the cold non-relativistic wavebreaking

field[25]

$$E_{\mathrm{wb}} = \frac{m_{\mathrm{e}} c \omega_{\mathrm{p}}}{e}, \tag{1.3}$$

$$\approx 96\sqrt{n_{\mathrm{e}}[10^{18}\,\mathrm{cm}^{-3}]}\,[\mathrm{GV\,m}^{-1}], \tag{1.4}$$

so for an electron density of $n_{\mathrm{e}} = 1 \times 10^{18}\,\mathrm{cm}^{-3}$, we see immediately that the maximum accelerating field is on the order of $100\,\mathrm{GV\,m}^{-1}$, which is three orders of magnitude higher than in a conventional accelerator.

1.2.2 Operating Regimes

Plasma Beat-Wave and Self-Modulated Laser-Driven Plasma Accelerators

In Tajima and Dawson's original paper,[19] they realised that to most effectively generate a wakefield, the length of the driving laser pulse should be half the length of the plasma wavelength. In other words

$$c\tau_{\mathrm{p}} = \frac{\lambda_{\mathrm{p}}}{2} = \frac{\pi c}{\lambda_{\mathrm{p}}} \qquad \rightarrow \qquad \tau_{\mathrm{p}} = \frac{\pi}{\omega_{\mathrm{p}}}, \tag{1.5}$$

where τ_{p} is the laser pulse duration, λ_{p} is the plasma wavelength and ω_{p} is the angular plasma frequency. For the nominal density of $n_{\mathrm{e}} = 1 \times 10^{18}\,\mathrm{cm}^{-3}$, the pulse duration required would be approximately 55 fs. This is considerably shorter than the shortest pulse duration of lasers that were available when Tajima and Dawson wrote their paper in 1979. The plasma beat-wave accelerator (PBWA) was proposed as an alternative scheme, using co-propagating laser pulses at two different frequencies ω_1 and ω_2, to generate a beat-wave with frequency $\Delta\omega = \omega_2 - \omega_1$ which resonantly excites a wakefield when $\Delta\omega = \omega_{\mathrm{p}}$. The generation of a PBWA was demonstrated in 1985[26] by a group at the University of California and was used to accelerate electrons by the same group in 1993.[27] However, the relativistic mass increase of the wakefield electrons causes a shift in the wakefield frequency (as can be seen from Eq. 1.2), leading saturation of the wakefield strength as it is driven off-resonance.[28] This can be mitigated using a self-modulated laser wakefield accelerator (SM-LWFA), which was first proposed by Andreev *et al.*[29] and Krall *et al.*[30] The laser pulse automatically drives a resonant wakefield because it is modulated by the wakefield driven by the leading edge of the laser pulse. This occurs via two mechanisms: i) focusing in regions of lower on-axis density and defocusing in regions of higher on-axis density, and ii) beating of the laser pulse with forward-Raman scattered laser light at $\Delta\omega = \omega_0 - \omega_{\mathrm{p}}$,

where ω_0 is the central frequency of the laser pulse.[31]

Linear and Non-Linear Regimes

Whether or not the wakefield electrons reach the energies required for relativistic effects to play a role provides an important demarcation between two operating regimes of plasma accelerators. In discussing the ponderomotive force (Eq. 1.1), we implicitly considered the average drift motion of electrons in the electromagnetic field of the laser pulse. But on timescales shorter than the pulse duration of the laser pulse envelope $\tau < \tau_p$, the electrons have a fast quiver motion in the instantaneous electric field of the laser pulse.[32] From the Lorentz equation, the maximum speed of an electron in the electric field of a linearly polarised plane wave is given by[33]

$$|\mathbf{v}| = \left| \frac{ie}{m_e \omega} \mathbf{E} \right| = \frac{eE_0}{m_e \omega}, \tag{1.6}$$

where E_0 is the peak electric field strength and ω is the angular frequency of the laser pulse. Using this we may define the normalised vector potential a_0 of the laser pulse as

$$a_0 = \frac{|\mathbf{v}|}{c} = \frac{eE_0}{m_e \omega c} = \frac{eA_0}{m_e c}. \tag{1.7}$$

We note that a more practical definition relates the normalised vector potential a_0 to the peak laser pulse intensity I_{peak} as

$$a_0 \approx 0.855 \, \lambda \, [\mu\text{m}] \sqrt{I_{\text{peak}} \, [10^{18} \text{W cm}^{-2}]}. \tag{1.8}$$

If the normalised vector potential of our laser pulse $a_0 \ll 1$, then (from Eq. 1.7) $|\mathbf{v}| \ll c$ and the system is in the non-relativistic linear regime. In this case, the axial accelerating field of the wakefield is sinusoidal and the accelerating electric field of the wake E is much less than the wavebreaking field E_{wb}. With the advent of chirped pulse amplification,[34] it became possible to generate ultrashort laser pulses at a petawatt level. For a Ti:Sa laser operating at a wavelength of $\lambda = 800\,\text{nm}$, the intensity required to reach $a_0 = 1$ is $I_{\text{peak}} \approx 2 \times 10^{18}\,\text{W cm}^{-2}$. For a $P = 6\,\text{TW}$ laser system, this peak intensity is reached for a $w_0 = 1\,\text{mm}$ spot size. Therefore, by using a single ultrashort laser pulse with a duration of half the plasma period, it became possible to drive wakefields not only in the linear regime that Tajima and Dawson originally envisioned, but also in the non-linear regimes where relativistic effects play a role.

The Bubble Regime

For non-linear wakefields with $a_0 \gg 1$, electrons are completely expelled by the laser pulse, creating a cavitated 'blow-out' region that is completely devoid of electrons.[35] In this regime, the axial accelerating field of the wakefield has a 'sawtooth' profile and a field strength E approximately equal to the wavebreaking field E_{wb}. Electrons can enter at the rear of the bubble near the axis, becoming trapped and accelerated in a process that is termed 'self-injection'. This will be discussed in greater detail in Section 1.4.2. In 2004, several groups experimentally produced quasi-monoenergetic electron bunches using self-injection in the bubble regime.[36–38]

1.3 Challenges

If LWFAs are to mature into a technology that can be operated similarly to the machines that are currently in use today, then there are challenges that remain to be addressed. These challenges will be considered in terms of three key areas:

1. Properties of the electron bunch,
2. Stability of these properties over long-term operation,
3. Efficiency and overall footprint of the accelerator.

In general, the requirements of any machine are specific to the applications that it is used for. For example, driving x-ray free-electron lasers (X-FELs) typically requires very high-quality bunches,[39] but the required energy may not be as high as that of a particle collider.[40,41] In this chapter, we outline the cutting-edge values and limitations faced by each property, and illustrate how simultaneous satisfaction of all desired requirements for a given application poses the ultimate challenge.

1.3.1 Properties of the Electron Bunch

Energy

At the time of writing, the maximum energy of an electron bunch produced by an LWFA is 10 GeV.[42] There are three key limitations to the maximum energy gain that may be achieved in an LWFA. These are diffraction, depletion and dephasing.

DIFFRACTION Any focusing laser pulse will diffract over the Rayleigh range. As already mentioned, to drive a wakefield in the quasi-linear regime, we require $a_0 \sim 1$. This means that the

input laser beam spot (normally on the order of a few centimetres) must be focused down to tens of micrometres. The spatial evolution of a Gaussian beam is given by

$$w(z) = w_0\sqrt{1+\left(\frac{z-z_0}{z_\mathrm{R}}\right)^2}, \tag{1.9}$$

where w_0 is the spot size of the Gaussian beam[†] at the focus position z_0. The Rayleigh range z_R for a Gaussian beam is given by

$$z_\mathrm{R} = \frac{\pi w_0^2}{\lambda}, \tag{1.10}$$

where λ is the wavelength of the light. We will derive Eq. 1.10 in Section 2.1.4. For an 800 nm beam focused to a spot size of 50 µm, the Rayleigh range is on the order of 10 mm. Clearly, this is not the many-centimetre interaction length that we require to accelerate electrons to the GeV energy level.

In the non-linear regime, the relativistic mass increase of the electrons changes the refractive index of the plasma. The refractive index η can be approximated as[43]

$$\eta \approx 1-\frac{1}{2}\left(\frac{\omega_\mathrm{p}}{\omega}\right)^2\left(1-\frac{a^2(r)}{2}+\frac{\delta n_\mathrm{e}(r)}{n_{e,0}}+\frac{\Delta n_\mathrm{ext}}{n_{e,0}}\right), \tag{1.11}$$

where $\delta n_\mathrm{e}(r)/n_{e,0}$ is the radial density perturbation from the ponderomotive force and $\Delta n_\mathrm{ext}/n_{e,0}$ is the normalised density perturbation of a preformed external plasma channel. For a laser pulse intensity profile which is peaked on axis, such as for a Gaussian beam, the refractive index will be lower on-axis, causing relativistic self-focusing. This occurs above the critical power of

$$P_\mathrm{c} = \frac{8\pi\varepsilon_0 m_\mathrm{e}^2 c^5 \omega^2}{e^2\omega_\mathrm{p}^2} \simeq 17.4\frac{\omega^2}{\omega_\mathrm{p}^2}[\mathrm{GW}]. \tag{1.12}$$

In this case, the spatial evolution of a Gaussian pulse is modified by the relativistic self-focusing and becomes

$$w(z) = w_0\sqrt{1+\left(1-\frac{P}{P_\mathrm{c}}\right)\left(\frac{z-z_0}{z_\mathrm{R}}\right)^2}, \tag{1.13}$$

such that a laser pulse with power P focused to a spot size of w_0 will diffract for $P < P_\mathrm{c}$, remain guided for $P = P_\mathrm{c}$, and be focused for $P > P_\mathrm{c}$.[43]

In cases where we want to drive a linear wakefield (for example, where we want to suppress

[†]Defined as the radial distance to a point where the intensity is $1/\mathrm{e}^2$ of the peak intensity.

self-injection to control the properties of the electron bunch, see Section 1.4.2), the power of the laser pulse will generally be less than the threshold power for self-focusing, where $P < P_c$. This means that in such regimes we need to provide an alternative guiding mechanism to keep the laser pulse intense enough to drive wakefields over long distances. This will be discussed in Section 1.4.1.

DEPLETION The laser pulse's energy is transferred to the wake as it propagates until it is completely depleted and acceleration ends. The length over which this occurs, or the pump depletion length L_{pd}, estimated by considering the energy density of the laser pulse $u_L = 1/2\varepsilon_0 E_0^2$ contained within the volume $V = \pi w_0^2 c\tau_p$ and the total energy density deposited in the wake $u_p = 1/2\varepsilon_0 E_z^2$ over the total volume $V = \pi w_0^2 L_{pd}$,[44] is

$$L_{pd} \approx \frac{\lambda_p^3}{\lambda_0^2 a_0^2}. \tag{1.14}$$

For a given laser pulse, with fixed values of λ_0 and a_0, this scales like $\sim n_e^{-3/2}$ by Eqs. 1.2 and 1.5.

DEPHASING Witness electrons accelerated in the wakefield quickly reach relativistic energies and speeds which are close to the speed of light c. As we will see in Section 2.1.5, the laser pulse has a subluminal group velocity, even in free-space. In a plasma, the group velocity is even lower and can be calculated using the refractive index $v_g = \eta c$, such that electrons very quickly outrun the accelerating portion of the wakefield. The distance over which an electron phase slips by $\lambda_p/4$ to enter the decelerating and defocusing portion of the wakefield is the dephasing length L_d, given by

$$(c - v_g)\frac{L_d}{c} = \frac{\lambda_p}{4} \qquad \rightarrow \qquad L_d = \frac{\lambda_p}{2}\frac{\omega^2}{\omega_p^2}. \tag{1.15}$$

Dephasing is generally the limiting factor for energy gain in an LWFA. It can be overcome to some extent by longitudinally tapering the plasma density profile.[45] In Chapter 3, we will study a technique for modifying the apparent group velocity of the laser pulse, as has been suggested for driving dephasingless laser wakefield accelerators.[46]

Energy Spread

Reliably producing electron bunches with percent-level energy spreads poses a significant challenge for LWFAs. Although the 'dream-beam' papers produced electron bunches with quasi-monoenergetic energy spreads at the percent-level, these bunches had fairly low energies, on the order of 100 MeV, and were trapped via self-injection, a process which is inherently uncontrollable and therefore not necessarily reproducible on every shot.[36–38] Preserving the energy spread of these bunches over longer acceleration lengths is not trivial, particularly in the non-linear regime.[47–50] Operation in linear regimes, combined with controlled injection techniques[51] to localise the injection region to a particular phase in the wake (as will be discussed in Section 1.4.2) and shaping of the injected charge to optimally 'load' or 'flatten' the wake,[52,53] can help to keep the energy spread low throughout the acceleration.

Bunch Length and Emittance

Accelerated electrons should sit within the portion of the wakefield where the fields are both accelerating and focusing, the length of this portion is $\lambda_p/4$. This corresponds to bunch durations of roughly $\tau < \lambda_p/4c \approx 10\,\mathrm{fs}$. Root-mean-square bunch durations as short as 2 fs have been measured experimentally.[54–57] This an order of magnitude shorter than the bunch lengths that can be achieved with conventional technology.[58] The emittance of LWFA electron bunches compares slightly less favourably. The normalised emittance is generally in the range of $0.1 - 2\pi\,\mathrm{mm\,mrad}$,[59–63] which is poor compared to the sub-mm mrad emittances of electron bunches from conventional accelerators. Some groups have demonstrated refocusing schemes which use plasma lenses[64,65] or more conventional transport lines with quadrupole magnets[66] to decrease the final bunch emittance. Others have shown that the initial emittance of the electron bunch can be reduced using controlled injection techniques.[67,68] An overview of these different controlled injection techniques will be given in Section 1.4.2.

Charge and Peak Current

The electron bunches produced by LWFAs generally have a total charge at the $10 - 100\,\mathrm{pC}$ level.[31] The high energy 10 GeV electron bunch[42] mentioned in Section 1.3.1 simultaneously had a high bunch charge of 340 pC. Another group demonstrated acceleration of a 640 pC bunch, corresponding to a peak current of 20 kA. Beam-loading requires careful tuning to prevent the injected charge from destroying the wakefield.[53,69,70] Due to the ultrashort electron bunch durations of LWFAs, bunches normally have peak currents at the 10s of kA-level. This is higher

than conventional accelerators operating at the kA-level, such that peak current requirements are generally less of a concern for LWFAs.

1.3.2 Stability of the Electron Bunch Properties

Typical shot-to-shot jitters are in the range of 1—5% for the energy, 5—50% for the charge, and $0.5-3$ mrad for the beam pointing.[31] These values are significantly higher than for conventional accelerators, which typically exhibit sub-percent level jitters. Finding a stable and controllable mechanism for injecting electrons into the wake is key for improving the stability of LWFAs. Operation in the linear regime could also be used to increase stability, since it reduces the risk of uncontrolled non-linear effects.[47–50]

1.3.3 Accelerator Efficiency and Footprint

The vast majority of LWFAs are based on Ti:Sa laser technology. Although such lasers are able to deliver high-power ultrashort laser pulses at repetition rates of $0.3-1$ Hz, many applications require repetition rates at a 1 kHz level, particularly for driving third-generation light sources, where the repetition rate of the source will have a huge impact on any imaging application which requires high temporal resolution. Operation of an LWFA at kHz repetition rate has been demonstrated, but the energy gain in this case was less than 10 MeV.[71]

Novel laser technologies, such as those based on optical parametric CPA (OPCPA),[72,73] fibre lasers[74] and diode-pumped solid-state or CW lasers,[75] have the potential to be very important in the development of future LWFAs,[31,76,77] but novel laser systems are currently limited to longer pulse durations at high kHz repetition rates.[78] This has motivated the study of LWFA schemes which can take a picosecond pulse and modulate it into a train of pulses using a plasma.[79,80] This train of pulses can then be used to resonantly drive a wakefield in the linear regime, which also reduces the energy requirements of the driving laser pulse by an order of magnitude or more.[81–84] One approach towards developing laser technologies which can provide sufficient energies at kHz repetition rates is to coherently combine the output of multiple fibre lasers.[85] Fibre lasers also tend to have wall-plug efficiencies in the range of $20-30$%, with the potential to reach 50% in the future,[76] whereas Ti:Sa lasers are generally very inefficient, often with < 0.1% wall-plug efficiencies.

When talking about a 'compact' particle accelerator, whose interaction or target length is on the order of several centimetres, the laser itself is often much larger, and in the case of Ti:Sa

laser systems, can often take up one or more rooms of lab space. Therefore, although the accelerator itself may be small, it is more important to consider the overall footprint of the machine. Future laser systems may yet be more compact, but even so, plasma-based accelerators are often as compact or more compact than conventional counterparts operating at similar energies. Furthermore, due to the steeper accelerating gradients, the length of any LWFA will scale at a fraction of the rate for a conventional accelerator.

1.4 Key Concepts

1.4.1 Plasma Waveguides

Plasma waveguides are crucial for extending the acceleration length of LWFAs, particularly those operating in the linear or quasi-linear regime. As discussed in Section 1.3.1, diffraction of the laser pulse ordinarily restricts the interaction length of the accelerator to the Rayleigh range, which is often on the order of a few millimetres or less. Therefore, to achieve multi-GeV electron beams, a waveguide is essential.

Normally, when we think about waveguides, we think of optical fibres. These are generally step-index fibres, where a higher refractive index core is surrounded by a cladding. At the boundary between the core and the cladding, total internal reflection occurs, such that the laser light remains trapped within the core of the fibre.[86] An alternative type of waveguide is the gradient-index waveguide, which has a refractive $\eta(\lambda, \rho)$ that is peaked on-axis, such that $\partial\eta/\partial\rho < 0$. This leads to a phase velocity which is the slowest on-axis and increases away from the axis, curving the wavefronts for a focusing effect. In an ideal waveguide, this perfectly cancels the effect of diffraction and keeps the pulse tightly focused over long distances.

Inverse Bremsstrahlung Heating

Using a plasma to generate an optical waveguide is ideal because it avoids the risk of damage to the waveguide material by the high-intensity laser pulse being guided.[87] Gradient-index plasma channels can be created from the hydrodynamic expansion of laser-produced channels.[88–90] The refractive index experienced by a laser pulse in a plasma is[91]

$$\eta(r) = \sqrt{1 - \frac{n_e(r)e^2}{\gamma m_e \varepsilon_0 \omega}}, \tag{1.16}$$

where γ describes the relativistic increase in electron mass owing to the quiver motion of the electron in the laser field, such that a plasma channel whose electron density $n_e(r)$ increases with radial distance from the optical axis will have a higher on-axis refractive index, creating an optical waveguide that keeps the laser pulse confined, as described in the previous section.

In early works, plasma channels were generated by ionisation and subsequent inverse bremsstrahlung heating caused by an axicon-focused laser pulse.[88] Inverse bremsstrahlung heats the plasma by allowing electrons to absorb energy from the laser pulse during collisions with the nucleus.[92] The on-axis hot plasma is surrounded by a neutral gas, causing a pressure gradient which drives a radial shock of the electrons outwards, with a shock-front where the ions (pulled by the negatively charged electrons) in the expanding plasma collide with the atoms in the surrounding gas. Subsequent collisional ionisation between the electrons and ions in the plasma can cause further heating. Relying on inverse bremsstrahlung heating requires long $\tau_p \approx 100\,\mathrm{ps}$, lower-intensity $10^{13} - 10^{14}\,\mathrm{W\,cm^{-2}}$ laser pulses and high on-axis densities of $n_e > 10^{18}\,\mathrm{cm^{-3}}$. These channels were used to demonstrate guiding over 24 Rayleigh ranges. However, there are a couple of drawbacks to using this technique.

Inverse bremsstrahlung necessitates the use of high-Z gas species if we want to perform both the ionisation and heating with the same low-intensity laser pulse. The outer electrons of high-Z gas species, such as argon (Ar), neon (Ne) or xenon (Xe) are less tightly bound and may therefore be ionised at lower intensities. However, this means that the channel may be destroyed or washed out by the wakefield driving laser pulse, which will have a much higher intensity and will therefore ionise the unionised inner shells of the gas atoms as it passes. Ionisation-induced phase modulation can also degrade the laser pulse focus.[93] This can be mitigated using the ignitor-heater method to ionise a low-Z gas, such as hydrogen (H_2), with a shorter $\tau_p = 75\,\mathrm{fs}$ pulse, before sending in a longer $\tau_p = 160\,\mathrm{ps}$ heater pulse, and this technique was used to demonstrate guiding of a laser pulse with an intensity of $I = 5 \times 10^{17}\,\mathrm{W\,cm^{-2}}$ over $\sim 5 - 10$ Rayleigh ranges. Another method instead used an electrical discharge to pre-ionise helium (He), demonstrating guiding of a laser pulse with an intensity of $I = 1.3 \times 10^{17}\,\mathrm{W\,cm^{-2}}$ laser pulse over 15 mm.[93] These methods still require two separate laser systems to produce the longer heater pulse and the shorter pulse to ionise the low-Z gas and drive the wakefield (and also to pre-ionise the plasma in the ignitor-heater method).

In addition to this, all of these methods rely primarily on inverse bremsstrahlung heating,

which is more effective at higher densities. Operation at high densities is less desirable for laser-wakefield applications, because the total energy gain scales inversely with the density.[44]

Clustered Gases

In an attempt to reduce the on-axis density of plasma waveguides, several groups employed the use of clustered gases, which have a low average density but a high local density. This was first demonstrated in 1998 by Ditmire *et al.* using a $\tau_{\mathrm{p}} \sim 2\,\mathrm{ps}$ to ionise the clustered gas by optical field ionisation.[94] It was later shown that a $\tau_p \sim 70\,\mathrm{fs}$ pulse can be used to perform both the optical field ionisation and heating, the peak on-axis density was less than $n_{\mathrm{e}} < 10^{18}\,\mathrm{cm}^{-3}$. Later work demonstrated again that a sub-picosecond $\tau_p \sim 400\,\mathrm{fs}$ pulse could be used for above threshold ionisation heating,[95] and the same group later showed that a circularly polarised laser pulse generated a faster blast wave, leading to more rapid channel formation.[96] It has also been shown that clustered gases could enhance ionisation injection in the bubble regime due to fluctuations in the shape and radius of the wakefield.[97]

Capillary Waveguides

In contrast to the gradient-index channels discussed so far, the dielectric walls of capillary waveguides act as step-index waveguides, such that the laser pulse can be guided in a vacuum or low-density gases.[98,99] This completely decouples the guiding mechanism from the laser pulse and plasma; we do not require any particular laser pulse or gas density for the ionisation and heating, such that the gas density can be tailored to optimise the wakefield acceleration. In earlier experiments, there was no gas within the capillary such that the guiding was mostly step-index.[98,99] Above the breakdown threshold, the guided laser pulse starts to ionise the capillary wall, creating a high-density plasma that further enhances the guiding efficiency. These experiments demonstrated guiding of $1-10\,\mathrm{TW}$ at efficiencies of 80%, but were multimode, meaning that higher-order Gaussian spatial modes were excited in the channel. Filling the capillary with gas allowed for the demonstration of monomode guiding,[24,100,101] giving a smoother transverse profile of the laser pulse. High-quality guiding of intense laser pulses in channels with on-axis plasma densities of $n_{\mathrm{e}} \sim 10^{17}\,\mathrm{cm}^{-3}$ has obvious applications for LWFAs.[102]

Capillary discharge waveguides rely primarily on gradient-index rather than step-index refraction, thus reducing the degree of damage to the walls of the capillary.[91,103–105] In this case, the capillary waveguide is filled with a gas which is ionised by a slow microsecond, high kiloampere current discharge. Hooker *et al.* filled sapphire (Al_2O_3, aluminium oxide, also termed

alumina) waveguides with hydrogen (H_2) to demonstrate this technique. An electron density gradient which increases away from the optical axis to the walls of the capillary forms due to the temperature gradient between the hot plasma on-axis and the cooler plasma at the walls, enabled by thermal conduction with the waveguide. Such capillary discharge waveguides were used to guide laser pulses with intensities of $I = 1.2 \times 10^{17}\,\mathrm{W\,cm^{-2}}$ over distances up to 50 mm, with coupling losses into the channel of $< 4\%$ and transmission losses of $7 \pm 1\,\mathrm{m^{-1}}$. These channels were also estimated to have a lifetime of $> 10^6$ shots.[106–108] A similar concave electron density profile can be produced using a Z-pinch discharge,[109,110] but this requires low jitter between the discharge and the arrival of the laser pulse to be guided because the channel only exists for a few nanoseconds.

The drawbacks of capillary discharge waveguides is the lack of optical access perpendicular to the channel, particularly if the inner wall of the channel becomes ablated, making it more difficult to probe and diagnose the properties of the channel. Despite this, Gonsalves *et al.* performed transverse interferometry on a capillary discharge waveguide[111] and found good agreement with earlier theoretical models,[112] which predict the matched spot size and axial electron density. For a parabolic channel, the radial electron density profile is

$$n_\mathrm{e}(r) = \Omega r^2, \tag{1.17}$$

with a constant coefficient Ω. The matched spot size w_m is defined as[113]

$$w_\mathrm{m} = \left(\pi r_\mathrm{e} \Omega\right)^{-1/4}, \tag{1.18}$$

where r_e is the classical electron radius. Any laser pulse which is focused into the waveguide with this 'matched' spot size w_m will remain confined and neither be focused nor defocused further. An unmatched input spot with spot size $w_0 \neq w_\mathrm{m}$ will oscillate according to[114,115]

$$w^2(z) = \frac{w_0^2}{2}\left[1 + \frac{w_\mathrm{m}^4}{w_0^4} + \left(1 - \frac{w_\mathrm{m}^4}{w_0^4}\right)\cos\left(2k_\beta z\right)\right], \tag{1.19}$$

where z is the propagation distance and $k_\beta = \lambda_0/(\pi w_\mathrm{m}^2)$ is the wavenumber of the laser pulse centroid oscillation about the plasma channel axis. Later studies used heater pulses to control the channel width and depth independently of the external capillary itself, allowing the channel to be

optimised for LWFA applications.[116] Operation at kilohertz repetition rates was demonstrated in 2016, although it should be noted that the laser pulse guided in this experiment was continuous-wave (CW) rather than a pulsed beam.[117] Laser-heated capillary discharge channels were used to demonstrate electron acceleration to 8 GeV,[24,118] with the theory about the axial uniformity of such channels following in 2021.[119] Finally, the length of these channels has been extended up to 40 cm long, with minimal fluctuations in both the on-axis electron density (less than $< 1\%$) and the focusing strength (less than $< 0.2\%$) across multiple shots.[120]

Hydrodynamic Optical Field Ionisation

An alternative to Section 1.4.1 is to use optical field ionisation (OFI) to ionise and heat the gas. OFI is the ionisation of atoms or molecules by a high-intensity optical field and can be modelled by solving the time-dependent Schrödinger equation. The Keldysh parameter γ_k is often used to distinguish between different regimes of ionisation and is defined in terms of the ponderomotive energy $\mathscr{E}_\mathrm{p}$ as

$$\gamma_\mathrm{k} = \sqrt{\frac{\mathscr{E}_\mathrm{ion}}{2\mathscr{E}_\mathrm{p}}}, \tag{1.20}$$

where the ionisation energy for hydrogen (H_2) is the Rydberg energy $\mathscr{E}_\mathrm{ion} = R_\mathrm{H} = m_\mathrm{e}e^4/8\varepsilon_0^2h^{2\dagger}$ and the ponderomotive energy $\mathscr{E}_\mathrm{p}$ is given by

$$\mathscr{E}_\mathrm{p} = \frac{e^2\lambda_0^2}{16\pi^2 m_\mathrm{e}c^2}|E|^2. \tag{1.21}$$

This is related to the ponderomotive force given in Eq. 1.1, noting the absence of ∇, since $\mathbf{F}_\mathrm{p} = -\nabla\mathscr{E}_\mathrm{p}$. This can also be related to the minimum appearance intensity I_app at which ionisation starts to occur by

$$\mathscr{E}_\mathrm{p} \approx 9.3 \times 10^{-14}\, I_\mathrm{app}\,[\mathrm{W\,cm^{-2}}]\,\lambda^2\,[\mu\mathrm{m}^2]. \tag{1.22}$$

Generally, $\gamma_\mathrm{k} \gg 1$ denotes regimes where multi-photon ionisation (MPI) dominates and $\gamma_\mathrm{k} \ll 1$ denotes regimes where tunnelling ionisation (TI) dominates. In the MPI regime, the time-dependent Schrödinger equation must be solved using perturbative methods, whereas TI can be modelled as quantum tunnelling of an electron in a modified Coulomb potential. In the region $\gamma_\mathrm{k} \sim 1$, the demarcation between different regimes is less distinct. For a highly-intense laser pulse, such that $\mathscr{E}_\mathrm{p} > \mathscr{E}_\mathrm{ion} \gg \hbar\omega$, perturbative models break down, so that the ionisation is better

[†]The ionisation energy $\mathscr{E}_\mathrm{ion}$ for other elements may be found using the NIST Atomic Spectra Database Ionization Energies Form.

described as TI.[121]

Solving the time-independent Schrödinger equation for hydrogen using the Wentzel-Kramers-Brillouin (WKB) approximation, with a Coulomb barrier suppressed by a DC (quasi-static) electric field of strength E, gives an estimated TI rate of

$$\Gamma_{\mathrm{TI}} = 8\frac{\mathscr{E}_{\mathrm{ion}}}{\hbar}\frac{2\mathscr{E}_{\mathrm{ion}}/a_0}{|eE|}\exp\left(-\frac{4}{3}\frac{\mathscr{E}_{\mathrm{ion}}/a_0}{|eE|}\right), \tag{1.23}$$

where $r_{\mathrm{Bohr}} = 4\pi\varepsilon_0\hbar^2/(m_{\mathrm{e}}e^2)$ is the Bohr radius. This is known as the Ammosov-Delone-Kraĭnov (ADK) rate.[122] We can note that the ionisation rate depends predominantly on the field strength E through the $\sim 1/|eE|$ term, and *not* the exponential term,[123] such that the threshold field strength E_{TI} is given by

$$E_{\mathrm{TI}} = \frac{2\mathscr{E}_{\mathrm{ion}}}{ea_0}. \tag{1.24}$$

In the limit of very strong applied electric fields, we can use a semi-classical approximation to find where the strength of the applied field suppresses the electrostatic field that binds the electron to the ion. Generally, for an element with atomic number Z, the potential $V(r)$ experienced by an electron orbiting an atom at a distance r within an intense laser field of strength E is given by

$$V(r) = -\frac{Ze}{4\pi\varepsilon_0|r|} - rE, \tag{1.25}$$

where Z is the atomic number[†]. Using the turning point $\partial V/\partial r = 0$ to find the position r_{max} where the potential is maximal, we obtain

$$r_{\mathrm{max}} = \sqrt{\frac{Ze}{4\pi\varepsilon_0 E}}, \tag{1.26}$$

and the potential at this point is

$$V_{\mathrm{max}} = V(r_{\mathrm{max}}) = -2\sqrt{\frac{Ze}{4\pi\varepsilon_0}E}. \tag{1.27}$$

Equating this to the known ionisation potential $\mathscr{E}_{\mathrm{ion}}/e$, we find that the threshold field strength for barrier-suppression is

$$E_{\mathrm{BSI}} = \frac{\pi\varepsilon_0\mathscr{E}_{\mathrm{ion}}^2}{Ze^3}, \tag{1.28}$$

[†]As has been used throughout in reference to 'high-Z' and 'low-Z' gases.

which can be related to the threshold intensity I_{app} as

$$I_{\text{app}} = \frac{c\varepsilon_0}{2}|E_{\text{BSI}}|^2 = (1+\epsilon)\frac{\pi^2\varepsilon_0^3 c}{2Z^2e^6}\mathscr{E}_{\text{ion}}^4, \tag{1.29}$$

where ϵ denotes the polarisation ellipticity of the field; $\epsilon = 0$ for linear polarisation and $\epsilon = 1$ for circular polarisation. This gives a threshold intensity for the ionisation of hydrogen with a circularly polarised laser pulse of $I_{\text{app,H}} = 2.8 \times 10^{14}\,\text{W}\,\text{cm}^{-2}$.

Unlike inverse bremsstrahlung heating, the electrons are heated directly from the laser field. This is not a collisional process and acts at the atomic level; it is therefore independent of the gas density. As already mentioned in Section 1.2.2, any electron in a laser field has both a fast quiver motion and a slower average drift. This slow average drift, determined by the ponderomotive potential, gives the resultant kinetic energy after interaction with the laser pulse. An electron will be freed as soon as $|E| \geq E_{\text{BSI}}$, and its kinetic energy $\mathscr{E}_{\text{k}}$ will be given by the ponderomotive energy at this point, such that

$$\mathscr{E}_{\text{k}} = \frac{e^2\lambda_0^2}{16\pi^2 m_{\text{e}}c^2}|E_{\text{BSI}}|^2, \tag{1.30}$$

which can again be written in terms of the ionisation energy $\mathscr{E}_{\text{ion}}$ as

$$\mathscr{E}_{\text{k}} = (1+\epsilon)\frac{\lambda_0^2\varepsilon_0^2}{16Z^2e^4m_{\text{e}}c^2}\mathscr{E}_{\text{ion}}^4. \tag{1.31}$$

For interaction with a given laser pulse, electrons released by BSI will exhibit distinct peaks depending on the gas species, through both the atomic number Z and the required ionisation energy of the element $\mathscr{E}_{\text{ion}}$. The ionised and heated plasma then expands hydrodynamically in the same manner as previously discussed, leading to them being termed 'hydrodynamic optical field ionisation' (HOFI) channels. HOFI channels were first formed experimentally using a spherical lens,[124] but axicon lenses are better suited to generating long channels, due to the line-like Bessel focus that they form.[125] Axicon-generated HOFI channels have been used to guide laser pulses with intensities up to $I = 6 \times 10^{17}\,\text{W}\,\text{cm}^{-2}$ at on-axis densities as low as $n_{\text{e}} = (1.0 \pm 0.3) \times 10^{17}\,\text{cm}^{-3}$ over lengths of 100 mm.[126,127] Furthermore, it has now been shown that a second laser pulse can ionise the neutral gas collar that surrounds the channel, 'conditioning' or deepening the channel to increase the attenuation length by two orders of magnitude compared to an unconditioned HOFI channel. Meter-long, low-loss channels can be generated from only

$1.2\,\mathrm{J\,m^{-1}}$ of laser energy.[128,129]

In summary, HOFI channels are low-loss, meter-scale guiding structures that are completely immune to damage. They can be formed in both low-Z and high-Z gases at any on-axis density. The properties of the channel can be tailored to the LWFA requirements on-shot by altering the gas density of the target and the relative timing between the channel-forming and wakefield-driving laser beams. They can be operated at kilohertz repetition rates[130] and are resistant to chromatic effects of the laser pulse.[131] It was shown in 2023 that they can be longitudinally structured for electron injection,[132] and in Chapter 7 we demonstrate that crossed HOFI channels can be used to increase the amount of charge injected into an LWFA.

1.4.2 Electron Injection

Electron injection has a huge impact both on the properties of the electron bunches produced, including their stability and repeatability over long-term operation.

Self-Injection

LWFAs often rely on self-injection, where in relativistic regimes the accelerating field is close to the wavebreaking field $E \sim E_{\mathrm{wb}}$. At high laser pulse intensities, it is possible for electrons to enter the back of the wakefield, this may be due to 'wave-breaking' or other interactions between the laser pulse and electrons in the plasma. If the amplitude of the plasma wave becomes so non-linear that it breaks, then this may inject electrons, provided that the structure of the wake is preserved.[133,134] If the injected electrons are in the right phase of the wake, they will be accelerated until they either overtake the driving laser pulse (dephasing) or until the laser pulse intensity decreases by diffraction or depletion such that it no longer drives a wake (see Section 1.3.1).

Self-injection often happens multiple times in the wake of one laser pulse and electrons that are injected at different times will witness different accelerating gradients, leading to bunches with very broad energy spreads[†]. It is inherently uncontrollable and therefore susceptible to variations across subsequent shots, due to the fact that it is a non-linear process. Even minute variations in other aspects, such as alignment, will give rise to a completely different electron bunches, leading to very large shot-to-shot energy variations with this method. There is often a large 'dark-current' of low energy background electrons which are released into the wrong phase

[†]Normally 100%, in other words, electrons of all energies up to the maximum are produced. Only a few electrons will experience the maximum possible accelerating electric field.

of the wake. Despite the experimental complexity posed, controlled injection techniques localise exactly when and where electrons are injected into the wake, leading to higher quality, more stable electron bunches in the long-term.

External Injection

Injection of electrons from within the plasma itself is often more desirable than external injection from an RF gun, as the latter produces electron bunches which are longer than the plasma wavelength λ_p. Timing an electron beam to arrive at the correct position in the wakefield relative to an independent laser system with femtosecond precision requires a high degree of control over the relative timing of both systems.[135] As a result, it is challenging to match the bunch structure of an RF beam into an LWFA, although some groups have shown successful acceleration of an externally matched electron beam.[136]

Density Control

Changing the density profile of the plasma longitudinally can be used to inject electrons using only a drive pulse. This can be loosely classified into two regimes. First, a gentle variation in density, or a 'density downramp', where the length of the density transition is greater than the plasma wavelength λ_p, can be used to slow down the wake phase velocity gradually. If this velocity drops below the fluid velocity of the electrons in the plasma, then the plasma wave effectively breaks, injecting electrons into the back of the wake.[43] This effect is compounded for subsequent buckets. It should be noted that this injection is prolonged, such that bunches produced this way have very large energy spreads. The reduction in wake phase velocity also means that electrons which are injected quickly dephase. Such density gradients can arise either naturally at the back of a gas jet[137] or can be created with a second transverse laser pulse.[138] In contrast, sharp density transitions on the scale of less than λ_p, also known as 'shock-injection', occur over such short distances that the wakefield appears to sharply jump at this transition. Electrons may suddenly find themselves in trapped orbits, such that they are accelerated in the wake. Compared to the gradual nature of a density downramp, electrons are injected for a much shorter period of time and therefore all witness the same accelerating gradient, resulting in low emittance electron beams with energy spreads of the order 10%.[139,140] Placing a razor in the nozzle of a gas jet has been shown to produce the sharp gradient required for localised shock injection of electrons.

Ionisation Injection

Ionisation injection relies on the release of electrons within the wakefield itself, by doping a hydrogen (H_2) plasma with a medium-Z dopant species, often nitrogen (N_2).[141] Only the very peak of the drive laser pulse has sufficient intensity to ionise the inner-shell electrons of the dopant species, meaning that these electrons are born close to the peak of the laser pulse and become trapped. Particularly in the case where a plasma channel is created by the pre-ionisation laser pulse, such as for a HOFI channel (see Section 1.4.1),[124,126,127] the driving laser pulse will be focused to a maximal intensity by the channel before diverging (particularly if the spot is not perfectly matched), creating a localised region where the maximum intensity of the laser pulse is high enough to ionise electron species from dopant atoms. These electrons are trapped if they are allowed to reach the phase velocity of the wake.[43] Ionisation injection may be combined with density ramps to further control trapping.[142,143]

Optical Injection

Optical injection has the potential to generate electron bunches with a high degree of control, leading to low energy spread and emittance bunches. In colliding pulse injection, a second laser pulse counter-propagates with the pulse that drives the wakefield. As they meet, they briefly create a beat wave with a ponderomotive force that can drive electrons to be trapped.[144] The drive pulse generally has a higher intensity than the counter-propagating pulse, but each has an intensity which is less than the threshold for self-injection, such that they only inject electrons when they interfere with each other. The parameters of the injected bunch can be controlled by adjusting the intensity of each pulse and the relative delay between them.[43] The difficulty here lies in the alignment and control of two beams rather than one.

In addition to a purely optical injection, certain types of density- and ionisation-injection can be classed also as optical injection. Optically-controlled density injection has already been mentioned; here a transverse laser pulse, which precedes the drive pulse, is used to carve out a region of lower density via HOFI.[145] In optically-controlled ionisation injection, a second laser pulse is used to ionise the dopant species. This has the benefit of decoupling the properties of the drive laser pulse from the injector, such that both can be optimised separately. An example of this is the two-pulse ionisation injection (2PII) scheme,[146] where rather than using the drive beam to ionise a dopant species in ionisation injection, a low-intensity driver is shadowed by a higher intensity injection pulse with a short Rayleigh range, such that it injects for a very short

distance before diffracting. This was shown experimentally to produce bunches with a small energy spread and emittance.[147] A modification of this scheme, which allows the trailing laser pulse to be radially-polarised, is discussed in Chapter 6.

1.5 Thesis Outline

This thesis will explore novel ideas based on all-optical electron injection. Chapters 3–6 are solely the author's own work, aided by discussions with Prof. Simon Hooker and Prof. Roman Walzcak. Chapter 4 was a study performed in collaboration with Prof. Martin Booth and Dr. Bangshan Sun. Chapter 7 presents the results of a collaborative experiment using the Astra-Gemini Laser at Rutherford Appleton Laboratories. In addition to taking a key role during the experiment, the author was solely responsible for coordinating all data analysis after the experiment.

CHAPTER 2 Introduces the relevant background theory and methods used throughout this thesis. This includes descriptions of the propagation of laser pulses in a vacuum under the paraxial and modified paraxial approximations and an introduction to spatiotemporal couplings in ultrashort laser pulses. The Collins diffraction integral is a Fourier optics method used to propagate non-monochromatic laser beams in free-space; an overview of the theory and a detailed description of the numerical simulation method is given.

CHAPTER 3 Reviews the concept of a flying focus: an ultrashort laser pulse which has a controllable apparent group velocity. A temporally chirped pulse is focused with a chromatic optic. This group velocity may be sub- or superluminal and may travel forwards or backwards with respect to the source. This defines the framework and simulation tools required for studying spatiotemporal couplings.

CHAPTER 4 Studies the concept of a linear simultaneous space-time focus: an ultrashort laser pulse which is temporally compressed during spatial focusing. A pulse which is spatially chirped along one axis is focused by an achromatic lens. We investigate whether a simultaneous space-time focus improves the axial localisation of the intensity, when compared to a conventional focus. We show that residual pulse-front tilt (PFT) prevents any improvement in the axial localisation.

CHAPTER 5 Studies the concept of a radial simultaneous space-time focus. A pulse which is radially-chirped in two dimensions is focused with an achromatic lens. We show that the radial simultaneous space-time focus has no residual pulse-front tilt, but that it still does not improve the axial localisation. However, we show that this pulse has a controllable group velocity, giving rise to a flying focus-like effect.

CHAPTER 6 Studies the effect of the laser pulse polarisation state on ionisation injection and trapping in quasi-linear wakes. We show that radially-polarised pulses have tighter confinement of the intensity in both the transverse and axial domains. An adaptation of the 2PII scheme, which includes alternative polarisation states, is presented.

CHAPTER 7 describes an experiment using the TA3 Gemini laser at the Rutherford Appleton Laboratory. We show that crossed HOFI channels can be used to inject and accelerate electron bunches with higher charge compared to a previous experiment using a single truncated HOFI channel. Using a simple model of the electron betatron motion, we use measured oscillations in the electron energy spectra to fit key parameters relating to the initial offset and energy of the electron bunch and the accelerating gradient of the wakefield.

2 | Background Theory and Methods

2.1 Propagation of Light in a Vacuum

2.1.1 Wave and Geometric Optics

The propagation of light was first understood geometrically based on Fermat's principle.[148] A much earlier version of Fermat's principle was proposed in the 11th century by Ibn al-Haytham.[149] From Fermat's principle, the basic ideas of reflection, refraction and diffraction, can be understood. The statement posed by Fermat in 1662 was that light always takes the path of least time.[148] Fermat's principle states that the path taken by light from a point P to P' is always such that the time taken is stationary in the calculus of variations sense. In other words, if the path is altered infinitesimally the resulting infinitesimal change in the time of propagation is zero.[150] Fermat's principle is analogous to the principle of least action in classical mechanics and Hamilton's equation of motion was originally formulated in 1827 as a characteristic function which defines the path length of an optical ray.[151] For a long time, light was thought of as streams of particles travelling along straight-line rays and this was sufficient to design optical systems of surprising elegance.[152]

Fermat's principle was somewhat heretical at the time because in a ray picture alone, it suggests that nature can test all possible routes and choose the fastest one. Therefore, it does not fully explain the nature of the propagation of light, although clearly, it does a good job of explaining its behaviour in many cases. When the wave-like nature of light was proposed by Huygens in 1678[153] and later demonstrated by Young in 1801,[154] Fermat's principle could be understood in all its glory; light takes every possible path but phase differences cause constructive interference only in the vicinity of the fastest possible path. It was not until 1911 that Sommerfeld and Runge[155] (aided by Debye[156]) made the connection between wave and geometric optics by

showing that the wave equation becomes the differential equation of the characteristic function in the limit of vanishingly small wavelengths.[150] We can derive the fundamental equation of geometric optics in the following way.

Within a small region of space, we might consider that light can be represented by a plane wave.[157] For this to be true, the amplitude and direction of the wave must remain constant on the order of the wavelength. Then we can introduce the idea of a wave surface; a surface where the phase of the wave is the same at all points. The wave surfaces of a plane wave are planes perpendicular to the direction of propagation, whose tangents are the 'rays'.

As already mentioned, this ray picture of geometric optics is valid as $\lambda \to 0$. Consider E to be generally any quantity describing the field of the wave (e.g. E_x or B_z). A scalar monochromatic plane wave is given by

$$E = |E| e^{i(\mathbf{k}\cdot\mathbf{r} - \omega t + \psi_0)}, \tag{2.1}$$

which can also be written as

$$E = |E| e^{i\psi}, \tag{2.2}$$

where generally both $|E|$ and ψ are a functions of the spatial coordinates and time. The phase or eikonal ψ will not generally have the simple form given by Eq. 2.1 and is clearly a large quantity because it changes by 2π as we move through one wavelength. As already mentioned, we will be considering the limit where $\lambda \to 0$. Over small regions in space and time

$$\psi \approx \psi_0 + \mathbf{r} \cdot \frac{\partial \psi}{\partial \mathbf{r}} + t \frac{\partial \psi}{\partial t}, \tag{2.3}$$

to first-order approximation. Comparing this to Eq. 2.1, we see that

$$k_i = -\frac{\partial \psi}{\partial x_i}, \tag{2.4}$$

where $k^i = \left(\omega/c, k_x, k_y, k_z\right)$ is the wave four-vector and $x^i = \left(ct, x, y, z\right)$ and we know that the components of the four-vector are related by $k_i k^i = 0$, such that

$$\frac{\partial \psi}{\partial x_i} \frac{\partial \psi}{\partial x^i} = 0. \tag{2.5}$$

This is the *eikonal* equation and is fundamental to geometric optics.[157] The trajectories of Hamil-

ton's characteristic rays must satisfy Eq. 2.5. In Section 2.1.3, we will derive this from the wave equation.

We will now derive the wave equation by considering light as electromagnetic fields which must satisfy Maxwell's equations. Although the paraxial solution to the wave equation is commonly treated as a starting point for the propagation of beams and is usually valid in the case of well-behaved laser light, it is based on a set of assumptions and we must always be careful that we bear these in mind. Starting from the wave equation, we shall derive the paraxial equation for the spatial evolution of focusing light and the modified paraxial equation for the full spatiotemporal evolution of light as it propagates. Much of what is presented here is well-established and relevant texts are cited where possible.

2.1.2 Complex Representation of Electric Fields

The real electric field of an ultrashort laser pulse $E(t)$ is related to the electric field in the spectral domain $\tilde{E}(\omega)$ by a Fourier transform[158]

$$\tilde{E}(\omega) = \mathscr{F}[E(t)] = \sqrt{\frac{1}{2\pi}} \int_{-\infty}^{\infty} E(t) \mathrm{e}^{i\omega t} \, \mathrm{d}t. \tag{2.6}$$

Conversely, we can recover the real electric field from the inverse Fourier transform

$$E(t) = \mathscr{F}^{-1}[\tilde{E}(\omega)] = \sqrt{\frac{1}{2\pi}} \int_{-\infty}^{\infty} \tilde{E}(\omega) \mathrm{e}^{-i\omega t} \, \mathrm{d}\omega. \tag{2.7}$$

It is convenient to use a spectrum that is zero for negative frequencies and an electric field that is complex. In this case, they are both related to the spectrum $\tilde{E}(\omega)$ via

$$\tilde{E}^{+}(\omega) = \begin{cases} \tilde{E}(\omega), & \omega \geq 0 \\ 0, & \omega < 0 \end{cases} = \sqrt{\frac{1}{2\pi}} \int_{-\infty}^{\infty} E^{+}(t) \mathrm{e}^{i\omega t} \, \mathrm{d}t, \tag{2.8}$$

and

$$E^{+}(t) = \sqrt{\frac{1}{2\pi}} \int_{0}^{\infty} \tilde{E}(\omega) \mathrm{e}^{-i\omega t} \, \mathrm{d}\omega = \sqrt{\frac{1}{2\pi}} \int_{-\infty}^{\infty} \tilde{E}^{+}(\omega) \mathrm{e}^{-i\omega t} \, \mathrm{d}\omega. \tag{2.9}$$

The second parts of expressions 2.8 and 2.9 show that $\tilde{E}^{+}(\omega)$ and $E^{+}(t)$ may be written as Fourier transforms of each other. Conversely, the negative frequency components may be given in terms

of the whole spectrum $\tilde{E}(\omega)$ as

$$\tilde{E}^{-}(\omega) = \begin{cases} 0, & \omega > 0 \\ \tilde{E}(\omega), & \omega \leq 0 \end{cases} = \sqrt{\frac{1}{2\pi}} \int_{-\infty}^{\infty} E^{-}(t) \mathrm{e}^{i\omega t} \, \mathrm{d}t, \tag{2.10}$$

and

$$E^{-}(t) = \sqrt{\frac{1}{2\pi}} \int_{-\infty}^{0} \tilde{E}(\omega) \mathrm{e}^{-i\omega t} \, \mathrm{d}\omega = \sqrt{\frac{1}{2\pi}} \int_{-\infty}^{\infty} \tilde{E}^{-}(\omega) \mathrm{e}^{-i\omega t} \, \mathrm{d}\omega, \tag{2.11}$$

again noting that these parts share a Fourier relation to one another. The original spectral and temporal fields may then be recovered as

$$\tilde{E}(\omega) = \tilde{E}^{+}(\omega) + \tilde{E}^{-}(\omega), \tag{2.12}$$

and

$$E(t) = E^{+}(t) + E^{-}(t). \tag{2.13}$$

Throughout this thesis we will use U to denote any scalar part of the electric field signal in the temporal or spectral domains. For example, $\tilde{U}(\omega)$ could denote just the spectrum and $\tilde{U}(\mathbf{r})$ just the spatial envelope in the spectral domain along the direction of polarisation, but clearly neither of these are the full expression for the vectorial electric field $\mathbf{E}$.

Envelope Representation

The electric field may be written as an envelope function multiplied by a carrier wave

$$\tilde{E}(\omega) = |\tilde{E}(\omega)| \mathrm{e}^{i\varphi(\omega)}, \tag{2.14}$$

where $|\tilde{E}(\omega)|$ denotes the spectral amplitude and $\varphi(\omega)$ denotes the spectral phase. Note that because $E(t)$ is a real function, the relation $\tilde{E}(\omega) = \tilde{E}^{*}(-\omega)$ holds. Any pulse may be written in terms of an envelope function and carrier wave, however for pulses much shorter than few-cycle, which are defined by

$$\frac{\Delta\omega}{\omega_0} \gg 1, \tag{2.15}$$

where $\Delta\omega$ is the bandwidth and ω_0 is the central frequency of the laser pulse, care must be taken when using the envelope approximation to reconstruct the field. For a pulse with a central wavelength 800 nm, the period of one laser oscillation is $\tau_1 = 2.67$ fs. For an unchirped pulse

duration of $\tau_{p0} = 45\,\text{fs}$, the pulse is more than few-cycle since

$$n_{\text{cycles}} = \frac{\tau_p}{\tau_l} = \frac{45\,\text{fs}}{2.67\,\text{fs}} = 16.85. \tag{2.16}$$

Derivation of the Spectrum for a Linearly Chirped Pulse

Linear spectral chirp arises from the second-order temporal phase. The time signal $U^+(t)$ of a linearly-chirped pulse has the form

$$U^+(t) = \mathrm{e}^{-\Gamma t^2}\mathrm{e}^{-i\omega_0 t}, \tag{2.17}$$

where $\Gamma = a + ib$ is a complex parameter that characterises the temporal width and chirp of the pulse. As given by Eq. 2.8, the spectrum signal is the Fourier transform of this and can be calculated by completing the square as follows

$$\begin{aligned}
\tilde{U}^+(\omega) &= \sqrt{\frac{1}{2\pi}}\int_{-\infty}^{\infty} U^+(t)\mathrm{e}^{i\omega t}\,\mathrm{d}t, && (2.18)\\
&= \sqrt{\frac{1}{2\pi}}\int_{-\infty}^{\infty} \mathrm{e}^{-\Gamma t^2}\mathrm{e}^{i(\omega-\omega_0)t}\,\mathrm{d}t, && (2.19)\\
&= \sqrt{\frac{1}{2\pi}}\int_{-\infty}^{\infty} \mathrm{e}^{-\Gamma\left[t^2 - i\frac{(\omega-\omega_0)}{\Gamma}t\right]}\,\mathrm{d}t, && (2.20)\\
&= \sqrt{\frac{1}{2\pi}}\int_{-\infty}^{\infty} \mathrm{e}^{-\Gamma\left[t^2 - i\frac{(\omega-\omega_0)}{\Gamma}t + \left(i\frac{(\omega-\omega_0)}{2\Gamma}\right)^2 - \left(i\frac{(\omega-\omega_0)}{2\Gamma}\right)^2\right]}\,\mathrm{d}t, && (2.21)\\
&= \sqrt{\frac{1}{2\pi}}\mathrm{e}^{-\Gamma\left(i\frac{(\omega-\omega_0)}{2\Gamma}\right)^2}\int_{-\infty}^{\infty} \mathrm{e}^{-\Gamma\left[t - i\frac{(\omega-\omega_0)}{2\Gamma}\right]^2}\,\mathrm{d}t, && (2.22)\\
&= \sqrt{\frac{1}{2\pi}}\mathrm{e}^{-\Gamma\left(i\frac{(\omega-\omega_0)}{2\Gamma}\right)^2}\int_{-\infty}^{\infty} \mathrm{e}^{-\Gamma\Omega^2}\,\mathrm{d}\Omega, && (2.23)\\
&= \sqrt{\frac{1}{2\Gamma}}\mathrm{e}^{\frac{(\omega-\omega_0)^2}{4\Gamma}}, && (2.24)
\end{aligned}$$

using the substitution $\Omega = t - i(\omega - \omega_0)/2\Gamma$ to perform the integration. The real part of the complex parameter is related to the pulse duration τ_p by $a = 2\ln 2/\tau_p^2$.

2.1.3 Scalar Wave Equation

We start with Maxwell's equations in a homogeneous, isotropic, linear medium

$$\boldsymbol{\nabla}\cdot\mathbf{B} = 0, \tag{2.25}$$

$$\boldsymbol{\nabla} \times \mathbf{E} = -\frac{\partial \mathbf{B}}{\partial t}, \tag{2.26}$$

$$\boldsymbol{\nabla} \cdot \mathbf{E} = \frac{\rho}{\varepsilon}, \tag{2.27}$$

$$\boldsymbol{\nabla} \times \mathbf{B} = \varepsilon\mu\frac{\partial \mathbf{E}}{\partial t} + \mu\mathbf{j}, \tag{2.28}$$

where $\mathbf{E}$ is the electric field, $\mathbf{B}$ is the magnetic field, ρ is the charge density, $\mathbf{j}$ is the current density, ε is the permittivity and μ is the permeability. We take the curl of Faraday's Law

$$\boldsymbol{\nabla} \times \boldsymbol{\nabla} \times \mathbf{E} = -\boldsymbol{\nabla} \times \frac{\partial \mathbf{B}}{\partial t}, \tag{2.29}$$

$$\boldsymbol{\nabla}(\boldsymbol{\nabla} \cdot \mathbf{E}) - \nabla^2\mathbf{E} = -\frac{\partial}{\partial t}(\boldsymbol{\nabla} \times \mathbf{B}), \tag{2.30}$$

$$\boldsymbol{\nabla}\Big(\frac{\rho}{\varepsilon}\Big) - \nabla^2\mathbf{E} = -\frac{\partial}{\partial t}\Big(\varepsilon\mu\frac{\partial \mathbf{E}}{\partial t} + \mu\mathbf{j}\Big), \tag{2.31}$$

where we have used the vector identity $\boldsymbol{\nabla} \times \boldsymbol{\nabla} \times \mathbf{E} = \boldsymbol{\nabla}(\boldsymbol{\nabla} \cdot \mathbf{E}) - \nabla^2\mathbf{E}$. The velocity of the plane wave is given by $v = 1/\sqrt{\varepsilon\mu} = c/n$, where n is the refractive index of our medium. Clearly, if this wave instead propagates through free-space, we have no charges $\rho = 0$ or currents $\mathbf{j} = 0$ and $\varepsilon \to \varepsilon_0$ and $\mu \to \mu_0$, such that the wave travels at the speed of light $c = 1/\sqrt{\varepsilon_0\mu_0}$ and

$$\nabla^2\mathbf{E} = \frac{1}{c^2}\frac{\partial^2 \mathbf{E}}{\partial t^2}. \tag{2.32}$$

We can note that this is a particular case of the generalised wave equation[157]

$$\frac{\partial^2 f}{\partial x_i \partial x^i} = 0, \tag{2.33}$$

and that if we consider that the generalised field f takes the form of a plane wave such that $f = a\mathrm{e}^{i\psi}$, then

$$\frac{\partial^2 a}{\partial x_i \partial x^i}\mathrm{e}^{i\psi} + 2i\frac{\partial a}{\partial x_i}\frac{\partial \psi}{\partial x^i}\mathrm{e}^{i\psi} + ia\frac{\partial^2 \psi}{\partial x_i \partial x^i}\mathrm{e}^{i\psi} - f\frac{\partial \psi}{\partial x_i}\frac{\partial \psi}{\partial x^i} = 0, \tag{2.34}$$

but we know that ψ must be a large quantity in the limit $\lambda \to 0$, since the eikonal or phase ψ changes by 2π as we move through one wavelength. Therefore, in the limit of small wavelengths

$$\frac{\partial \psi}{\partial x_i}\frac{\partial \psi}{\partial x^i} = 0, \tag{2.35}$$

which defines a geometric ray of light, as we saw previously in Section 2.1.1.

We ordinarily consider separable solutions to the scalar wave equation (Eq. 2.32). This assumes that the behaviour of the wave in space is not coupled to its behaviour in time (later we shall consider what happens when this is not the case, see Section 2.1.6). In this case we assume that the electric field takes some unknown form $\mathbf{E}(\mathbf{r},t)=\mathbf{R}(\mathbf{r})T(t)$, such that

$$T(t)\boldsymbol{\nabla}^2\mathbf{R}(\mathbf{r})=\frac{1}{c^2}\mathbf{R}(\mathbf{r})\frac{\partial^2 T(t)}{\partial t^2}. \tag{2.36}$$

Rearranging this gives

$$c^2\frac{1}{\mathbf{R}(\mathbf{r})}\boldsymbol{\nabla}^2\mathbf{R}(\mathbf{r})-\frac{1}{T(t)}\frac{\partial^2 T(t)}{\partial t^2}=0, \tag{2.37}$$

and we can note that it will always be the case that these terms sum to zero, regardless of $\mathbf{r}$ and t. This is only possible if each term is constant, and we set the separation constant to be $-\omega^2$, such that the temporal separation is parameterised by ω and the spatial separation by $\omega/c=k$, as is the case for an electromagnetic wave. This gives

$$\frac{\partial^2}{\partial t^2}T(t)+\omega^2 T(t)=0, \tag{2.38}$$

$$\boldsymbol{\nabla}^2\mathbf{R}(\mathbf{r})+\frac{\omega^2}{c^2}\mathbf{R}(\mathbf{r})=0. \tag{2.39}$$

These equations can then be independently solved to give solutions which oscillate in space and time. One example is a spatially-dependant electric field amplitude multiplied by a time-harmonic phase

$$\mathbf{E}(\mathbf{r},t)=\mathbf{E}(\mathbf{r})\mathrm{e}^{-i\omega t}. \tag{2.40}$$

We can note that if we only want the field strength at any point, then

$$E(\mathbf{r},t)=E(\mathbf{r})\mathrm{e}^{-i\omega t}. \tag{2.41}$$

Substituting this into the wave equation, we find that the time-dependent phase is a common factor on both sides, such that we obtain

$$\nabla^2\big(E(\mathbf{r})\mathrm{e}^{-i\omega t}\big)=\frac{1}{c^2}\frac{\partial^2\big(E(\mathbf{r})\mathrm{e}^{-i\omega t}\big)}{\partial t^2}, \tag{2.42}$$

$$\mathrm{e}^{i\omega t}\nabla^2 E=-\frac{\omega^2}{c^2}\mathrm{e}^{i\omega t}E(\mathbf{r}), \tag{2.43}$$

$$\nabla^2 E = -\frac{\omega^2}{c^2} E(\mathbf{r}), \tag{2.44}$$

$$(\nabla^2 + k^2)E = 0, \tag{2.45}$$

where we have used the relation $k = \omega/c$. This is the Helmholtz equation, which clearly applies to harmonic solutions. The most general solution to the Helmholtz is the angular spectrum,[159] given by the 2D Fourier transform

$$E_{\text{out}}(x, y, z) = \sqrt{\frac{1}{2\pi}} \iint_{-\infty}^{\infty} \mathfrak{H}(k_x, k_y; z) \tilde{E}_{\text{in}}(k_x, k_y; 0) \mathrm{e}^{i(k_x x + k_y y)} \, \mathrm{d}k_x \, \mathrm{d}k_y, \tag{2.46}$$

where $\tilde{E}_{\text{in}}$ is the spatially-transformed input field and H is the transfer function up to an arbitrary distance z. The transfer function generally has the form

$$\mathfrak{H}(k_x, k_y; z) = \mathrm{e}^{i\sqrt{k^2 - k_x^2 - k_y^2} z}. \tag{2.47}$$

Section 2.1.4 covers the paraxial equation, which is a solution to the wave equation for a monochromatic beamlet. If we have a polychromatic wave with a broadband spectrum, then we would need to consider the Helmholtz equation for each frequency within the spectrum. The modified paraxial wave equation is equivalent to solving the paraxial equation for each frequency in the envelope and is studied in Section 2.1.6.

2.1.4 Paraxial Wave Equation

To derive the paraxial wave equation, we set the z-axis to lie in the direction of propagation of the wave. We neglect temporal behaviour for now. For a paraxial solution, the dominant term in the spatial phase is the z term, such that we can factor out the x and y dependent terms. This means assuming a solution of the form

$$\mathbf{E}(x, y, z) = |\mathbf{E}(x, y, z)| \, \mathrm{e}^{ikz}, \tag{2.48}$$

where $|\mathbf{E}|$ is the gradually varying spatial pulse envelope. Taking the Laplacian of this solution gives

$$\boldsymbol{\nabla}^2 \mathbf{E} = \boldsymbol{\nabla}_\perp^2 |\mathbf{E}| \mathrm{e}^{ikz} + \frac{\partial^2}{\partial z^2} \Big(|\mathbf{E}| \mathrm{e}^{ikz} \Big), \tag{2.49}$$

$$= \boldsymbol{\nabla}_\perp^2 |\mathbf{E}| e^{ikz} + \frac{\partial}{\partial z}\left(\frac{\partial |\mathbf{E}|}{\partial z} e^{ikz} + ik|\mathbf{E}| e^{ikz} \right), \tag{2.50}$$

$$= \boldsymbol{\nabla}_\perp^2 |\mathbf{E}| e^{ikz} + \frac{\partial^2 |\mathbf{E}|}{\partial z^2} e^{ikz} + 2ik\frac{\partial |\mathbf{E}|}{\partial z} e^{ikz} - k^2 |\mathbf{E}| e^{ikz}. \tag{2.51}$$

We can then make the paraxial approximation to neglect the second-order partial derivative in z, such that

$$\boldsymbol{\nabla}^2 \mathbf{E} = \left(\boldsymbol{\nabla}_\perp^2 |\mathbf{E}| + 2ik\frac{\partial |\mathbf{E}|}{\partial z} - k^2 |\mathbf{E}| \right) e^{ikz}. \tag{2.52}$$

Substituting this into the Helmholtz equation gives the paraxial equation

$$\left(\boldsymbol{\nabla}_\perp^2 |\mathbf{E}| + 2ik\frac{\partial |\mathbf{E}|}{\partial z} - k^2 |\mathbf{E}| + k^2 |\mathbf{E}| \right) e^{ikz} = 0, \tag{2.53}$$

$$\boldsymbol{\nabla}_\perp^2 |\mathbf{E}| + 2ik\frac{\partial |\mathbf{E}|}{\partial z} = 0. \tag{2.54}$$

Since this is a case of the Helmholtz equation under specific conditions, we can take the general solution of the Helmholtz equation given in Eq. 2.46. We can then apply our paraxial (propagation at very small angles along z axis) and Fresnel ($x^2 + y^2 \ll z^2$) approximations to the transfer function $\mathfrak{H}$, as these assumptions only govern the propagation of the ray, and therefore do not change the initial input field. Under the paraxial and Fresnel approximations, our transfer function Eq. 2.47 is

$$\mathfrak{H}(k_x, k_y; z) \approx e^{i\left(k - \frac{k_x^2}{2k} - \frac{k_y^2}{2k}\right)z} = e^{ikz} e^{\frac{-i(k_x^2 + k_y^2)z}{2k}}, \tag{2.55}$$

such that (substituting the above into Eq. 2.46) the field is

$$E_{\text{out}}(x, y, z) = \sqrt{\frac{1}{2\pi}} \iint_{-\infty}^{\infty} \mathfrak{H}(k_x, k_y; z) \tilde{E}_{\text{in}}(k_x, k_y; 0) e^{i(k_x x + k_y y)} dk_x dk_y, \tag{2.56}$$

$$= \sqrt{\frac{1}{2\pi}} \iint_{-\infty}^{\infty} e^{ikz} e^{\frac{-i(k_x^2 + k_y^2)z}{2k}} \tilde{E}_{\text{in}}(k_x, k_y; 0) e^{i(k_x x + k_y y)} dk_x dk_y, \tag{2.57}$$

$$= \sqrt{\frac{1}{2\pi}} e^{ikz} \iint_{-\infty}^{\infty} \tilde{E}_{\text{in}}(k_x, k_y; 0) e^{\frac{-i(k_x^2 + k_y^2)z}{2k}} e^{i(k_x x + k_y y)} dk_x dk_y. \tag{2.58}$$

Clearly, our envelope function is given by the k-space Fresnel integral

$$|\tilde{E}|_{\text{out}}(x, y, z) = \sqrt{\frac{1}{2\pi}} \iint_{-\infty}^{\infty} \tilde{E}_{\text{in}}(k_x, k_y; 0) e^{\frac{-i(k_x^2 + k_y^2)z}{2k}} e^{i(k_x x + k_y y)} dk_x dk_y. \tag{2.59}$$

We can derive an explicit expression in the case of a Gaussian beam. For a Gaussian input beam

$$E_{\text{in}}(x,y;0) = -\frac{|E_{\text{in}}|}{i\pi w_0^2/\lambda} e^{-\frac{x^2+y^2}{w_0^2}}, \tag{2.60}$$

where $|E_{\text{in}}|$ is a constant factor denoting the amplitude. Taking an inverse Fourier transform of the Gaussian spatial profile, by completing the square we can find the angular spectrum to be

$$\begin{aligned}
\tilde{E}_{\text{in}}(k_x,k_y;0) &= -\frac{|E_{\text{in}}|}{i\pi w_0^2/\lambda}\sqrt{\frac{1}{2\pi}}\iint_{-\infty}^{\infty} e^{-\frac{x^2+y^2}{w_0^2}} e^{-i(k_x x+k_y y)}\mathrm{d}x\,\mathrm{d}y, && (2.61)\\
&= -\frac{|E_{\text{in}}|}{i\pi w_0^2/\lambda}\sqrt{\frac{1}{2\pi}}\iint_{-\infty}^{\infty} e^{-\frac{1}{w_0^2}\left[(x^2+ik_x x w_0^2)+(y^2+ik_y y w_0^2)\right]}\mathrm{d}x\,\mathrm{d}y, && (2.62)\\
&= -\frac{|E_{\text{in}}|}{i\pi w_0^2/\lambda}\sqrt{\frac{1}{2\pi}} e^{-\frac{w_0^2}{4}(k_x^2+k_y^2)}\iint_{-\infty}^{\infty} e^{-\frac{\left(x+\frac{ik_x w_0^2}{2}\right)^2}{w_0^2}} e^{-\frac{\left(y+\frac{ik_y w_0^2}{2}\right)^2}{w_0^2}}\mathrm{d}x\,\mathrm{d}y, && (2.63)\\
&= -\frac{|E_{\text{in}}|}{i\pi w_0^2/\lambda}\sqrt{\frac{1}{2\pi}}(\sqrt{\pi}w_0)^2\exp\Big\{-\frac{w_0^2}{4}(k_x^2+k_y^2)\Big\}, && (2.64)\\
&= \frac{i\lambda|E_{\text{in}}|}{\sqrt{2\pi}}\exp\Big\{-\frac{w_0^2}{4}(k_x^2+k_y^2)\Big\}. && (2.65)
\end{aligned}$$

Using this in our equation for the output field (Eq. 2.58) we find that

$$\begin{aligned}
E_{\text{out}}(x,y,z) &= \sqrt{\frac{1}{2\pi}} e^{ikz}\iint_{-\infty}^{\infty}\frac{i\lambda|E_{\text{in}}|}{\sqrt{2\pi}} e^{-\frac{w_0^2}{4}(k_x^2+k_y^2)} e^{\frac{-i(k_x^2+k_y^2)z}{2k}} e^{i(k_x x+k_y y)}\mathrm{d}k_x\mathrm{d}k_y, && (2.66)\\
&= \sqrt{\frac{1}{2\pi}} e^{ikz}\iint_{-\infty}^{\infty}\frac{i\lambda|E_{\text{in}}|}{\sqrt{2\pi}} e^{-\left(\frac{w_0^2}{4}+\frac{iz}{2k}\right)(k_x^2+k_y^2)} e^{i(k_x x+k_y y)}\mathrm{d}k_x\mathrm{d}k_y, && (2.67)
\end{aligned}$$

where we have used the substitution that $\Omega = w_0^2/4 + iz/2k$ and we also note that the function we are Fourier transforming here has the same form as before, but with $w_0^2/4 \to w_0^2/4 + iz/2k$, such that conversely, we obtain

$$\begin{aligned}
E_{\text{out}}(x,y,z) &= \frac{i\lambda|E_{\text{in}}|}{2\pi} e^{ikz}\iint_{-\infty}^{\infty} e^{-\Omega\left[(k_x^2+ik_x x/\Omega)+(k_y^2+ik_y y/\Omega)\right]}\mathrm{d}k_x\mathrm{d}k_y, && (2.68)\\
&= \frac{i\lambda|E_{\text{in}}|}{2\pi} e^{ikz} e^{-\frac{1}{4\Omega}(x^2+y^2)}\iint_{-\infty}^{\infty} e^{-\Omega\left(k_x+\frac{ix}{2\Omega}\right)^2} e^{-\Omega\left(k_y+\frac{iy}{2\Omega}\right)^2}\mathrm{d}k_x\mathrm{d}k_y, && (2.69)\\
&= \frac{i\lambda|E_{\text{in}}|}{2\pi} e^{ikz} e^{-\frac{1}{4\Omega}(x^2+y^2)}\left(\sqrt{\frac{\pi}{\Omega}}\right)^2, && (2.70)\\
&= \frac{i\lambda|E_{\text{in}}|}{2\Omega} e^{ikz}\exp\Big(-\frac{x^2+y^2}{4\Omega}\Big). && (2.71)
\end{aligned}$$

We note that the substitution $\Omega = w_0^2/4 + iz/2k$ can be related to the complex q-parameter $q = z - i\pi w_0^2/\lambda$ as $\Omega = i\lambda q/4\pi = iq/2k$ and $k = 2\pi/\lambda$, such that

$$E_{\text{out}}(\rho, z) = 2\pi \frac{|E_{\text{in}}|}{q} e^{ikz} e^{\frac{ik\rho^2}{2q}}, \tag{2.72}$$

where we can write the Gaussian parameter as $q = z - iz_{\text{R}}$, with the Rayleigh range $z_{\text{R}} = \pi w_0^2/\lambda$, spot size w_0 and radial distance $\rho = \sqrt{x^2 + y^2}$. The Rayleigh range and spot size may be defined in terms of the intensity for a Gaussian beam. The Rayleigh range is the distance from the focal position where the peak intensity is highest, to the point at which the intensity is half of the peak intensity. The spot size is the radial distance from the center of the beam where the intensity is highest, to the point at which the intensity is $1/e^2$ of the peak intensity.

This is the form of Gaussian beam we shall assume in most cases as a starting point since the paraxial approximation is nearly always valid in situations of interest and the Gaussian solution is the easiest to manipulate mathematically.

2.1.5 Higher-Order Modes

In the previous section, we derived a solution to the free-space paraxial wave equation. This is only the lowest-order solution in an infinite family of higher-order solutions.[160] These higher-order solutions take the form of the Hermite-Gaussian functions in Cartesian coordinates, or the Laguerre-Gaussian functions in cylindrical coordinates.[161,162] In the context of LWFAs, a cylindrical geometry makes the most sense, due to the fact that the physical system is often approximately azimuthally symmetric. This is particularly pertinent if we consider plasma waveguides, as described in Section 1.4.1. Higher-order solutions in cylindrical coordinates generally take the form of Laguerre-Gaussian solutions[160]

$$E_p^{\ell}(\rho, \theta, z) = \sqrt{\frac{2p!}{\pi(|\ell| + p)!}} \frac{e^{ikz} e^{\frac{ik\rho^2}{2q(z)}}}{w(z)} \left[\frac{\sqrt{2}\rho}{w(z)}\right]^{|\ell|} e^{-i(2p+|\ell|+1)\varphi_{\text{G}}(z)} L_p^{|\ell|}\left[\frac{2\rho^2}{w(z)^2}\right] e^{i\ell\theta}, \tag{2.73}$$

where $L_p^{|\ell|}$ is the generalised Laguerre polynomial at a given radial index $p \geq 1$ and azimuthal mode index ℓ and

$$\rho = \sqrt{x^2 + y^2}, \qquad \theta = \tan^{-1}\left(\frac{y}{x}\right), \tag{2.74}$$

$$\varphi_{\text{G}}(z) = \tan^{-1}\left(\frac{z - z_0}{z_{\text{R}}}\right), \tag{2.75}$$

$$w(z) = w_0\sqrt{1+\left(\frac{z-z_0}{z_\mathrm{R}}\right)^2}, \tag{2.76}$$

$$q(z) = z - z_0 - iz_\mathrm{R}, \tag{2.77}$$

where ρ is the radial coordinate, θ is the polar angle coordinate, $\varphi_\mathrm{G}(z)$ is the Guoy phase, $w(z)$ is the beam waist and $q(z)$ is the complex q-parameter. We can also write this exclusively in terms of the complex q-parameter as follows.

Derivation of Waist-Independent Expression for a Laguerre-Gaussian

Generally, a standard Laguerre-Gaussian mode is given by[160]

$$E_p^\ell(\rho,\theta,z) = \sqrt{\frac{2p!}{\pi(|\ell|+p)!}}\frac{\mathrm{e}^{-\frac{ik\rho^2}{2q^+(z)}}}{w(z)}\left[\frac{\sqrt{2}\rho}{w(z)}\right]^{|\ell|}\mathrm{e}^{i(2p+|\ell|+1)\varphi_\mathrm{G}(z)}L_p^{|\ell|}\left[\frac{2\rho^2}{w(z)^2}\right]\mathrm{e}^{i\ell\theta}, \tag{2.78}$$

where $L_p^{|\ell|}$ is the generalised Laguerre polynomial at a given radial index $p \geq 1$ and azimuthal mode index ℓ. The inverted complex radius of curvature or q-parameter is given by $q^+(z) = z - z_0 + iz_\mathrm{R}$ (note the change in sign), where z_0 is the position of the beam waist and z_R is the Rayleigh range. The position-dependent spot size $w(z) = w_0\sqrt{1+\left(\frac{z-z_0}{z_\mathrm{R}}\right)^2}$ and the Guoy phase $\varphi_\mathrm{G}(z) = \tan^{-1}\left(\frac{z-z_0}{z_\mathrm{R}}\right)$ are the same as for a Gaussian beam. We will now derive an equation which eliminates these last two parameters to describe the beam entirely in terms of the q-parameter. We can first derive the position-dependent spot size in terms of the q-parameter as

$$w(z) = w_0\sqrt{1+\left(\frac{z-z_0}{z_\mathrm{R}}\right)^2}, \tag{2.79}$$

$$= \frac{w_0}{z_\mathrm{R}}\sqrt{z_\mathrm{R}^2+(z-z_0)^2}, \tag{2.80}$$

$$= \frac{w_0}{z_\mathrm{R}}|q(z)|, \tag{2.81}$$

$$= \sqrt{\frac{2}{kz_\mathrm{R}}}|q(z)|, \tag{2.82}$$

where we have noted that the beam waist spot size is $w_0 = \sqrt{2z_\mathrm{R}/k}$, since $z_\mathrm{R} = \pi w_0^2/\lambda$. The Guoy phase shift term may also be written as

$$\mathrm{e}^{i\varphi_\mathrm{G}(z)} = \frac{i|q(z)|}{q(z)}. \tag{2.83}$$

Substituting these into our original expression, we get

$$E_p^\ell(\rho,\theta,z) = \sqrt{\frac{p!}{\pi(|\ell|+p)!}}\frac{\sqrt{kz_\mathrm{R}}\mathrm{e}^{-\frac{ik\rho^2}{2q(z)}}}{|q(z)|}\left[\frac{\rho\sqrt{kz_\mathrm{R}}}{|q(z)|}\right]^{|\ell|}\left[\frac{i|q(z)|}{q(z)}\right]^{2p+|\ell|+1} L_p^{|\ell|}\left[\frac{kz_\mathrm{R}\rho^2}{|q(z)|^2}\right]\mathrm{e}^{i\ell\theta}, \quad (2.84)$$

$$= \frac{\sqrt{kz_\mathrm{R}}}{|q(z)|}\left[\frac{\rho\sqrt{kz_\mathrm{R}}}{|q(z)|}\right]^{|\ell|}\left[\frac{i|q(z)|}{q(z)}\right]^{2p+|\ell|+1}\sqrt{\frac{p!}{\pi(|\ell|+p)!}}L_p^{|\ell|}\left[\frac{kz_\mathrm{R}\rho^2}{|q(z)|^2}\right]\mathrm{e}^{-\frac{ik\rho^2}{2q(z)}}\mathrm{e}^{i\ell\theta}. \quad (2.85)$$

We now want to simplify everything before the normalisation term containing factorials. Separating this out

$$\frac{\sqrt{kz_\mathrm{R}}}{|q(z)|}\left[\frac{\rho\sqrt{kz_\mathrm{R}}}{|q(z)|}\right]^{|\ell|}\left[\frac{i|q(z)|}{q(z)}\right]^{2p+|\ell|+1} = \frac{1}{\rho}\left[\frac{\rho\sqrt{z_\mathrm{R}k}}{|q(z)|}\right]^{|\ell|+1}\left[\frac{i|q(z)|}{q(z)}\right]^{2p+|\ell|+1}, \quad (2.86)$$

$$= \frac{1}{\rho}\left[\frac{\rho\sqrt{z_\mathrm{R}k}}{|q(z)|}\frac{i|q(z)|}{q(z)}\right]^{|\ell|+1}\left[\frac{i|q(z)|}{q(z)}\right]^{2p}, \quad (2.87)$$

$$= \frac{1}{\rho}\left[\frac{i\rho\sqrt{z_\mathrm{R}k}}{q(z)}\right]^{|\ell|+1}\left[\left(\frac{i\sqrt{q^*(z)q(z)}}{q(z)}\right)^2\right]^p, \quad (2.88)$$

$$= \frac{1}{\rho}\left[\frac{i\rho\sqrt{z_\mathrm{R}k}}{q(z)}\right]^{|\ell|+1}\left[-\frac{q^*(z)q(z)}{q(z)^2}\right]^p, \quad (2.89)$$

$$= \frac{1}{\rho}\left[\frac{i\rho\sqrt{z_\mathrm{R}k}}{q(z)}\right]^{|\ell|+1}\left[-\frac{q^*(z)}{q(z)}\right]^p, \quad (2.90)$$

such that our final expression for a Laguerre-Gaussian beam only in terms of the q-parameter is

$$E_p^\ell(\rho,\theta,z) = \sqrt{\frac{p!}{\pi(|\ell|+p)!}}\frac{\mathrm{e}^{-\frac{ik\rho^2}{2q(z)}}}{\rho}\left[\frac{i\rho\sqrt{z_\mathrm{R}k}}{q(z)}\right]^{|\ell|+1}\left[-\frac{q^*(z)}{q(z)}\right]^p L_p^{|\ell|}\left[\frac{kz_\mathrm{R}\rho^2}{|q(z)|^2}\right]\mathrm{e}^{i\ell\theta}, \quad (2.91)$$

as is consistent with existing literature.[163]

Group Velocity of Structured Light

It is so often stated that the velocity of light is a constant c in free-space, despite the fact that this is only true of a plane wave, which happens to be a very poor model for our laser pulse. As seen already, our laser pulses always have a transverse extent and in cases where this is much less than the propagation distance, we can use the paraxial approximation to simplify our models for the propagation. However, the finite transverse extent of our beams has consequences that should not be overlooked. The introduction of transverse structure changes the group velocity of the beam, even in free-space.[164] Defining the group velocity of structured light is still an ongoing area of research,[165] therefore we will not go into detail, but instead we derive a simple estimate of the on-axis group velocity of a Laguerre-Gaussian mode from the spectral phase (where clearly

this boils down to the Gaussian case for $p = 0$, $\ell = 0$). An alternative derivation was presented by Bareza *et al.*[166] and leads to the same result. Starting from Eq. 2.73, we can note that the spectral phase of this beam can be written as

$$\varphi(\omega) = \frac{\omega}{c} z - (2p + |\ell| + 1)\varphi_{\mathrm{G}} + \frac{\omega}{c}\frac{\rho^2}{2}\frac{z}{z^2 + (w_0^2\omega/2c)^2}, \tag{2.92}$$

where we have noted that $\Im\{1/q\} = z^2/(z^2 + z_{\mathrm{R}}^2)$ and $z_{\mathrm{R}} = w_0^2\omega/2c$, where we have assumed that the spot size is frequency-independent, such that by this definition the Rayleigh range is frequency-dependent. An expansion of the spectral phase about the central angular frequency ω_0 yields

$$\varphi(\omega) = \varphi(\omega_0) + \left.\frac{\partial\varphi}{\partial\omega}\right|_{\omega_0}(\omega - \omega_0) + \cdots. \tag{2.93}$$

Considering just the on-axis group velocity, such that $\rho = 0$ and matching terms we see that

$$\left.\frac{\partial\varphi}{\partial\omega}\right|_{\omega_0} = \varphi^{(1)} = \frac{z}{v_{\mathrm{g}}} = \frac{z}{c} - (2p + |\ell| + 1)\left.\frac{\partial\varphi_{\mathrm{G}}}{\partial\omega}\right|_{\omega_0}. \tag{2.94}$$

Since

$$\tan(\varphi_{\mathrm{G}}) = \frac{z}{z_R} = \frac{2zc}{w_0^2\omega}, \tag{2.95}$$

we can differentiate implicitly to find that

$$\sec^2(\varphi_{\mathrm{G}})\left.\frac{\partial\varphi_{\mathrm{G}}}{\partial\omega}\right|_{\omega_0} = -\frac{2zc}{w_0^2\omega_0^2}, \tag{2.96}$$

such that

$$\left.\frac{\partial\varphi_{\mathrm{G}}}{\partial\omega}\right|_{\omega_0} = -\left[1 + \left(\frac{z}{z_{\mathrm{R}}}\right)^2\right]^{-1}\frac{2zc}{w_0^2\omega_0^2}. \tag{2.97}$$

Substituting this into Eq. 2.94, we find that

$$v_{\mathrm{g}} = c\left\{1 + \frac{2}{w_0^2k_0^2}(2p + |\ell| + 1)\left[1 + \left(\frac{z}{z_{\mathrm{R}}}\right)^2\right]^{-1}\right\}^{-1}. \tag{2.98}$$

In the focal region, where $z \ll z_{\mathrm{R}}$, our expression reduces to the expression in Bareza *et al.* Away from the focus where $z \gg z_{\mathrm{R}}$, we find that $v_{\mathrm{g}} \to c$, such that this effect is localised for focused beams.

2.1.6 Modified Paraxial Wave Equation

We will now derive an equation governing the electric field of a polychromatic laser pulse, whose wavelengths propagate according to the paraxial equation.[167,168] This is referred to as the modified paraxial equation (MPE) and it evolves the spatiotemporal profile of a polychromatic laser pulse in free-space. As described in Eq. 2.14, we can describe the pulse as a plane wave which is modulated by a more slowly varying envelope

$$\mathbf{E}(x,y,z,t) = |\mathbf{E}(x,y,z,t)|\mathrm{e}^{i(kz-\omega t)}. \tag{2.99}$$

Our field now has a spectral or temporal dependence, such that it is easier to write this field in terms of the moving-frame coordinate $\xi = z - ct$ and either the explicit space z dependence or the time t dependence. Here we write the time-dependence explicitly for

$$\mathbf{E}(r_\perp,\xi,t) = |\mathbf{E}(r_\perp,\xi,t)|\mathrm{e}^{ik\xi}, \tag{2.100}$$

where the field is split in terms of the longitudinal space-time coordinate ξ and the transverse coordinate $r_\perp$. In a vacuum, the central angular wavenumber k is related to the angular frequency by $k = \omega/c$. Starting from the free-space wave equation

$$\nabla^2\mathbf{E} = \frac{1}{c^2}\frac{\partial^2\mathbf{E}}{\partial t^2}, \tag{2.101}$$

we now want to derive this relation for a field in terms of the moving-frame coordinate $\xi(t) = z - ct$. Our envelope amplitude takes the form $|\mathbf{E}(r_\perp,\xi(t),t)|$, such that the time derivative will be given by

$$\frac{\mathrm{d}\mathbf{E}(\xi(t),t)}{\mathrm{d}t} = \frac{\partial\mathbf{E}}{\partial t} + \frac{\partial\mathbf{E}}{\partial\xi}\frac{\partial\xi}{\partial t} = \frac{\partial\mathbf{E}}{\partial t} - c\frac{\partial\mathbf{E}}{\partial\xi}. \tag{2.102}$$

Assuming that our field takes the form in Eq. 2.100, we can take the second-order time derivative explicitly as follows

$$\frac{\partial^2\mathbf{E}}{\partial t^2} = \frac{\partial^2}{\partial t^2}\Big(|\mathbf{E}|\mathrm{e}^{ik\xi}\Big), \tag{2.103}$$

$$= \Big(\frac{\partial}{\partial t} - c\frac{\partial}{\partial\xi}\Big)\Big(\frac{\partial|\mathbf{E}|}{\partial t}\mathrm{e}^{ik\xi} - c\frac{\partial|\mathbf{E}|}{\partial\xi}\mathrm{e}^{ik\xi} - ikc|\mathbf{E}|\mathrm{e}^{ik\xi}\Big), \tag{2.104}$$

$$= \frac{\partial^2|\mathbf{E}|}{\partial t^2}\mathrm{e}^{ik\xi} - c\frac{\partial^2|\mathbf{E}|}{\partial\xi\partial t}\mathrm{e}^{ik\xi} - ikc\frac{\partial|\mathbf{E}|}{\partial t}\mathrm{e}^{ik\xi} - c\frac{\partial^2|\mathbf{E}|}{\partial t\partial\xi}\mathrm{e}^{ik\xi} + c^2\frac{\partial^2|\mathbf{E}|}{\partial\xi^2}\mathrm{e}^{ik\xi}$$

$$\ldots + ikc^2\frac{\partial|\mathbf{E}|}{\partial\xi}e^{ik\xi} - ikc\frac{\partial|\mathbf{E}|}{\partial t}e^{ik\xi} + ikc^2\frac{\partial|\mathbf{E}|}{\partial\xi}e^{ik\xi} - k^2c^2|\mathbf{E}|e^{ik\xi}, \tag{2.105}$$

$$= \frac{\partial^2|\mathbf{E}|}{\partial t^2}e^{ik\xi} - 2c\frac{\partial^2|\mathbf{E}|}{\partial t\partial\xi}e^{ik\xi} - 2ikc\frac{\partial|\mathbf{E}|}{\partial t}e^{ik\xi} + c^2\frac{\partial^2|\mathbf{E}|}{\partial\xi^2}e^{ik\xi} + 2ikc^2\frac{\partial|\mathbf{E}|}{\partial\xi}e^{ik\xi} - k^2c^2|\mathbf{E}|e^{ik\xi}, \tag{2.106}$$

assuming that $\partial^2/\partial\xi\partial t \equiv \partial^2/\partial t\partial\xi$. We now need to find the Laplacian of our field (Eq. 2.100), again using the partial derivative with respect to ξ. With $\xi(z) = z - ct$ such that

$$\frac{\mathrm{d}\mathbf{E}(\xi(z))}{\mathrm{d}z} = \frac{\partial\mathbf{E}}{\partial z} = \frac{\partial\mathbf{E}}{\partial\xi}\frac{\partial\xi}{\partial z} = \frac{\partial\mathbf{E}}{\partial\xi}, \tag{2.107}$$

and therefore

$$\boldsymbol{\nabla}^2\mathbf{E} = \boldsymbol{\nabla}_\perp^2|\mathbf{E}|e^{ik\xi} + \frac{\partial^2}{\partial z^2}\Big(|\mathbf{E}|e^{ik\xi}\Big), \tag{2.108}$$

$$= \boldsymbol{\nabla}_\perp^2|\mathbf{E}|e^{ik\xi} + \frac{\partial}{\partial\xi}\left(\frac{\partial|\mathbf{E}|}{\partial\xi}e^{ik\xi} + ik|\mathbf{E}|e^{ik\xi}\right), \tag{2.109}$$

$$= \boldsymbol{\nabla}_\perp^2|\mathbf{E}|e^{ik\xi} + \frac{\partial^2|\mathbf{E}|}{\partial\xi^2}e^{ik\xi} + 2ik\frac{\partial|\mathbf{E}|}{\partial\xi}e^{ik\xi} - k^2|\mathbf{E}|e^{ik\xi}. \tag{2.110}$$

Using both Eqs. 2.106 and 2.110 in the free-space wave equation (Eq. 2.101), we get

$$\boldsymbol{\nabla}_\perp^2|\mathbf{E}|e^{ik\xi} + \frac{\partial^2|\mathbf{E}|}{\partial\xi^2}e^{ik\xi} + 2ik\frac{\partial|\mathbf{E}|}{\partial\xi}e^{ik\xi} - k^2|\mathbf{E}|e^{ik\xi} = \frac{1}{c^2}\frac{\partial^2|\mathbf{E}|}{\partial t^2}e^{ik\xi} - \frac{2}{c}\frac{\partial^2|\mathbf{E}|}{\partial t\partial\xi}e^{ik\xi} - \frac{2ik}{c}\frac{\partial|\mathbf{E}|}{\partial t}e^{ik\xi} + \frac{\partial^2|\mathbf{E}|}{\partial\xi^2}e^{ik\xi} + 2ik\frac{\partial|\mathbf{E}|}{\partial\xi}e^{ik\xi} - k^2|\mathbf{E}|e^{ik\xi}. \tag{2.111}$$

Cancelling all terms common to both sides, neglecting any remaining second order time derivatives and dividing by $e^{ik\xi}$ gives

$$\boldsymbol{\nabla}_\perp^2|\mathbf{E}| + \frac{2ik}{c}\frac{\partial|\mathbf{E}|}{\partial t} + \frac{2}{c}\frac{\partial^2|\mathbf{E}|}{\partial t\partial\xi} = 0. \tag{2.112}$$

We can write this as one operator such that

$$\left[\boldsymbol{\nabla}_\perp^2 + \frac{2}{c}\frac{\partial}{\partial t}\left(ik + \frac{\partial}{\partial\xi}\right)\right]|\mathbf{E}| = 0. \tag{2.113}$$

This is the modified paraxial equation, as in previous literature.[168] Another approach involves writing the longitudinal spatial dependence explicitly,[167] but is not investigated here. As before,

the most general solution to this type of equation is the Fresnel integral[167]

$$\mathbf{E}(\mathbf{r}_\perp, \xi, t) = \frac{k}{2\pi i z} \int \mathrm{e}^{\frac{ik}{2z}(\mathbf{r}-\mathbf{r}')^2} \mathbf{E}_\perp(\mathbf{r}'_\perp, \xi, t=0)\mathrm{d}\mathbf{r}', \tag{2.114}$$

which, unlike solutions to the paraxial wave equation, is dependent on the angular frequency ω through the angular wavenumber $k = \omega/c$. This allows us to study the propagation of non-monochromatic light, accounting for the full spatiotemporal evolution of the fields, which is particularly relevant when considering ultrashort pulses with appreciable bandwidth.

2.2 Spatiotemporal Couplings

Spatiotemporal couplings (STCs) are relations between the behaviour of a pulse in time and space. The nature of these couplings can be complicated because different effects are not independent due to the Fourier transform relations of pulses in the near and far fields. Without a thorough treatment of the underpinning mathematics, it is easy to be misled and to draw incorrect conclusions[†]. Although generally speaking STCs arise with any coupling between the spatial and temporal properties of the beam (for example, a stepped mirror might add a delay to different spatial regions of the beam), these are often introduced through the spectrum, for example, by chirping a pulse and passing it through a chromatic optic to link spatiospectral behaviour with temporal.[170] In the context of this thesis, we define a spatiotemporally coupled laser pulse as follows.

DEFINITION Spatiotemporally coupled pulse: $U(\mathbf{r}, t) \neq R(\mathbf{r})T(t)$.
Any pulse defined mathematically in space and time as $U(\mathbf{r}, t)$ has a behaviour in space $R(\mathbf{r})$ which is coupled to the behaviour in time $T(t)$ such that they can no longer be treated independently.

Introducing STCs is often unavoidable in the case of short pulse laser beams, and can have detrimental effects such as increased pulse duration and reduced intensity at focus.[171,172] These are often compensated for, but it has come to light that STCs could have effects which may be incredibly useful in certain situations. In this thesis, we shall primarily explore two types of STC pulse in the context of optical control in LWFAs: the flying focus (Chapter 3) and the simultaneous space-time focus (Chapters 4 and 5).

[†]For example, it has often been stated that pulse-front tilt and angular dispersion are equivalent, when in fact one can generate pulse-front tilt from beams with no angular dispersion.[169]

2.3 Collins Diffraction Integral

Fourier optics methods are widely used to study the diffraction and propagation of light. They are normally based on solutions to the paraxial wave equation (see Section 2.1.4), ordinarily for monochromatic light, but this is insufficient if we wish to study effects that are only present in polychromatic beams (such as effects arising from spatiotemporal couplings). It is normally assumed that for polychromatic light, the solution may be taken as a superposition of the solutions for each wavelength; this is based on solutions to the modified paraxial wave equation, as was shown in Section 2.1.6. Throughout this thesis, we shall focus on one particular diffraction integral: the Collins integral.[173] This diffraction is particularly intuitive and neat due to its intersection with conventional ray-tracing methods. There has been considerable work done to understand spatiotemporal couplings in terms of dispersive ray matrices,[169,174–177] but here we use the usual 2×2 ray matrices to describe the optical system, where generally the elements of these matrices may be frequency-dependent to allow for spatiotemporal couplings. The Collins method has been used to propagate exotic beams such as hypergeometric beams[178] and reciprocating flying-foci,[179] showing the versatility of this method. Additionally, it allows for the simulation of a more realistic laser pulse propagation, in cases where one can measure the amplitude and the phase of the beam at the input plane.

The Collins diffraction integral relates the elements of a conventional ABCD ray matrix to a diffraction integral. A conventional ray matrix has the general form

$$O = \begin{pmatrix} A & B \\ C & D \end{pmatrix}. \tag{2.115}$$

The Collins diffraction integral is defined in terms of these elements as[173]

$$\tilde{E}_{\text{out}}(\mathbf{r},\omega) = \frac{1}{i\lambda B} e^{ik(z-z_0)} \int_{-\infty}^{\infty}\int_{-\infty}^{\infty} \tilde{E}_{\text{in}}(\mathbf{r}_{\text{in}},\omega) e^{ikS} dx_{\text{in}} dy_{\text{in}}, \tag{2.116}$$

$$\text{with} \quad S = \left\{ \frac{1}{2B}\Big[A(x_{\text{in}}^2 + y_{\text{in}}^2) + D(x^2+y^2) - 2(x x_{\text{in}} + y y_{\text{in}})\Big] \right\}, \tag{2.117}$$

where the input signal $\tilde{E}_{\text{in}}(\mathbf{r}_{\text{in}},\omega)$ may have any spatial profile and spectrum. The angular wavenumber is given by $k = 2\pi/\lambda$. One could model dispersion by including the dispersion relation $k = \omega/v$, where v is dependant on the refractive index $v = c/n(\lambda)$, but for propagation

in free-space we assume here that $k = \omega/c$. This allows for the Collins integral to be written directly in terms of ω, as will be seen below.

S-FFT Method

The numerical computation of diffraction integrals, or digital holography, is a well-established field. Previous texts give extensive detail on the various methods available and their limitations.[180] Here we present a short summary of the methods used in this thesis.

The Collins equation (Eq. 2.117) may be rewritten in terms of a single Fourier transform (where 'single' denotes the number of Fourier transforms performed per computation, rather than the number of transform variables) as[180,181]

$$\tilde{E}_{\text{out}}(x,y) = \frac{e^{ik(z-z_0)}}{i\lambda B} e^{\frac{ik}{2B}D(x^2+y^2)} \mathscr{F}\left\{\tilde{E}_{\text{in}}(x_{\text{in}},y_{\text{in}}) \exp\left[\frac{ik}{2B}A(x_{\text{in}}^2+y_{\text{in}}^2)\right]\right\}_{f_x=\frac{x}{\lambda B}, f_y=\frac{y}{\lambda B}}. \tag{2.118}$$

This may be calculated numerically at different axial positions z using a single fast Fourier transform (S-FFT)

$$\begin{aligned}\tilde{E}_{\text{out}}\left(p\frac{\lambda B}{\Delta L_{\text{in}}}, q\frac{\lambda B}{\Delta L_{\text{in}}}\right) = \frac{e^{ik(z-z_0)}}{i\lambda B} e^{i\pi\frac{\lambda BD}{\Delta L_{\text{in}}^2}(p^2+q^2)} \\ \times \text{FFT}\left\{\tilde{E}_{\text{in}}\left(m\frac{\Delta L_{\text{in}}}{N}, n\frac{\Delta L_{\text{in}}}{N}\right) \exp\left[i\pi\frac{\Delta L_{\text{in}}^2}{\lambda}\frac{A}{B}(m^2+n^2)\right]\right\}, \\ \text{with} \quad p,q,m,n = -\frac{N}{2},\ldots,\frac{N}{2}-1,\end{aligned} \tag{2.119}$$

where $\tilde{E}_{\text{in}}(\mathbf{r},\omega)$ is any arbitrary input field and A, B, C and D are the matrix elements of the optical system. ΔL_{in} and ΔL are the widths of the input and output grids respectively and N is the number of grid points. For a detailed discussion of the simulation grids and the conditions used to decide the sampling grids, see Section 2.3.1. We may then reduce the number of spatial coordinates and assume propagation in free-space to rewrite this discrete integral in terms of x and ω only to give

$$\begin{aligned}\tilde{E}_{\text{out}}\left(n\frac{\Delta L}{N},\omega\right) = -\frac{i\omega}{2\pi cB} e^{\frac{i\omega}{2c}\frac{D}{B}\left(\frac{n\Delta L}{N}\right)^2} \text{IFFT}\left\{\tilde{E}_{\text{in}}\left(m\frac{\Delta L_{\text{in}}}{N},\omega\right) \exp\left(\frac{i\omega}{2c}\frac{A}{B}\frac{m^2\Delta L_{\text{in}}^2}{N^2}\right)\right\}, \\ \text{with} \quad m,n = -\frac{N}{2},\ldots,\frac{N}{2}-1.\end{aligned} \tag{2.120}$$

In the same way as for the analytic solution, the field may be calculated at different axial positions z, this time for any input field and optical system.

D-FFT Method

Equation 2.117 may also be rewritten in terms of a double Fourier transform as[180,181]

$$\tilde{E}_{\text{out}}(x,y) = \mathrm{e}^{ik(z-z_0)}\mathrm{e}^{-\frac{ikC}{2A}(x^2+y^2)}\mathscr{F}^{-1}\left\{\mathscr{F}\left\{\frac{1}{A}\tilde{E}_{\text{in}}\Big(\frac{x_{\text{in}}}{A},\frac{y_{\text{in}}}{A}\Big)\right\}\exp\Big[-i\pi\lambda BA(f_x^2+f_y^2)\Big]\right\}_{f_x=\frac{x}{\lambda B},f_y=\frac{y}{\lambda B}}. \tag{2.121}$$

This may be calculated numerically at different axial positions z using a double fast Fourier transform (D-FFT)

$$\tilde{E}_{\text{out}}\Big(p\frac{\Delta L}{N},q\frac{\Delta L}{N}\Big) = \mathrm{e}^{ik(z-z_0)}\mathrm{e}^{-\frac{ikC}{2A}(\frac{\Delta L}{N})^2\left(p^2+q^2\right)}\text{IFFT}\left\{\text{FFT}\left\{\frac{1}{A}\tilde{E}_{\text{in}}\Big(m\frac{\Delta L_{\text{in}}}{AN},n\frac{\Delta L_{\text{in}}}{AN}\Big)\right\}\mathrm{e}^{-i\pi\lambda BA\frac{m^2}{\Delta L_{\text{in}}^2}}\right\},$$
$$\text{with}\quad p,q,m,n=-\frac{N}{2},\ldots,\frac{N}{2}-1. \tag{2.122}$$

where FFT denotes our fast Fourier transform and IFFT denotes the inverse operation. In the case of the D-FFT, the scale of the simulation grid is preserved such that the output grid size ΔL is the same as the input grid size ΔL_{in}. The following expression is often termed the 'transfer function' as it is convolved with the input field to 'transfer' or propagate it to the output plane

$$\mathfrak{H}\Big(m\frac{\Delta L}{N}\Big) = \exp\Big\{-i\pi\lambda BA\frac{m^2}{\Delta L_0^2}\Big\}. \tag{2.123}$$

Notably, for an optical system consisting of a lens followed by propagation in free-space, at the focus $A=0$, since $A=1-z/f$ and $z=f$. This causes issues when the field is rescaled by $1/A$, as it can lead to numerical singularities about the focus. Therefore, since we are generally studying the fields in the vicinity of the focus throughout this thesis, the S-FFT method is used instead. As before, we can reduce the number of spatial coordinates and assume propagation in free-space to rewrite this discrete integral in terms of x and ω only, to give

$$\tilde{E}\Big(n\frac{\Delta L}{N},\omega\Big) = \mathrm{e}^{\frac{i\omega}{2c}\frac{C}{A}\frac{\Delta L}{N}\left(p^2+q^2\right)^2}\text{IFFT}\left\{\text{FFT}\left\{\frac{1}{A}\tilde{E}_0\Big(m\frac{\Delta L_0}{AN},\omega\Big)\right\}\mathrm{e}^{\frac{2\pi^2 c}{i\omega}BA\frac{m^2}{\Delta L_0^2}}\right\},$$
$$\text{with}\quad m,n=-\frac{N}{2},\ldots,\frac{N}{2}-1. \tag{2.124}$$

2.3.1 Detailed Discussion of the Simulation Grids

FFTs are well suited to many programming languages[†] and allow for quick computation of the output field at different axial positions z. All of the analytic models discussed in this thesis calculate the output field in the spectral domain, and do not Fourier transform analytically to the time domain because this requires making an assumption that the pulse duration is long. Such an assumption is not valid for ultrashort laser pulses, therefore we always Fourier transform the field to the time domain numerically.

When calculating the output fields using the Collins method, one must also convolve the input field with a propagating function at the input plane. This means taking a numerical Fourier transform of the input field, multiplied by the propagator term e^{ikS}. Correct discretization of the electric field is therefore crucial to ensure that the entire volume of the pulse is correctly resolved.

Discrete Fourier Transforms in Python

In Python's Scipy module[†], the discrete Fourier transform of the spectral electric field is given by

$$\tilde{E}(\omega)[n] = \sum_{m=0}^{N-1} \mathrm{e}^{-2\pi j \frac{m*n}{N}} E(t)[m], \tag{2.125}$$

and conversely, the inverse Fourier transform for the temporal electric field is given by

$$E(t)[m] = \frac{1}{N} \sum_{n=1}^{N-1} \mathrm{e}^{2\pi j \frac{m*n}{N}} \tilde{E}(\omega)[n], \tag{2.126}$$

where m and n index the spectral and temporal grids respectively and j is a complex number. This is a different convention from that stated in Section 2.1.2, therefore care must be taken to ensure that the correct fft is called in each instance, with the correct prefactors. A forwards fft should be multiplied by $\mathrm{d}t/\sqrt{2\pi}$, where $\mathrm{d}t$ is the size of the grid-spacing in time. The inverse fft should be multiplied by $\sqrt{2\pi}/\mathrm{d}t$.

Temporal and Spectral Grids

For an unchirped pulse with a pulse duration $\Delta\tau_{\mathrm{p0}}$, the laser pulse bandwidth $\Delta\omega$ is given by

$$\Delta\omega = \frac{4\log 2}{\Delta\tau_{\mathrm{p0}}}. \tag{2.127}$$

[†]This simulation was originally written by the author in both Python and MATLAB, with Python proving to be significantly faster. Therefore, all simulations described in this thesis were performed using Python.

[†]https://docs.scipy.org/doc/scipy/tutorial/fft.html

A chirped pulse with this bandwidth will then have a duration of

$$\Delta\tau_{\mathrm{p}} = \frac{4\log 2}{\Delta\omega}\sqrt{1+\chi^2}, \tag{2.128}$$

where χ parameters the chirp strength. To capture all spectral behaviour, we need a window that is approximately $W \geq N\Delta\omega_l = \frac{4N\log 2}{\Delta\tau_{\mathrm{p}}}$ and to capture the behaviour in time we need a window that is approximately $T \geq N\Delta\tau_{\mathrm{p}} = \frac{4N\log 2}{\Delta\omega}\sqrt{1+\chi^2}$, where N is some integer sufficient to capture the entire pulse in both frequency and time. Some experimentation is needed to determine the value of N which captures the entire pulse. Once this is decided, we have the required widths of our windows.

Thinking now about the resolution required, we can estimate the laser period from the central frequency of the laser pulse

$$\tau_{\mathrm{l}} = \frac{2\pi}{\omega_0}. \tag{2.129}$$

Again we can parameterise this with an integer such that $t_{\mathrm{res}} \leq \frac{\tau_{\mathrm{l}}}{M}$. Equivalently, we would like to sample at least M points across the spectrum of the laser pulse, such that $\omega_{\mathrm{res}} \leq \frac{\omega_0}{M}$. By Nyquist's theorem, the minimum sampling frequency must be at least twice the highest frequency of a given signal, such that $M > 2$. However, ideally, we would like to sample many points across one period of the signal (closer to about 10 points). There is no hard and fast rule, but care must be taken to ensure that the signal is sampled properly.

In summary, we have four conditions on our spectral and temporal grids in terms of the parameters of our pulse

1. $W \geq N\Delta\omega_l = \frac{4N\log 2}{\Delta\tau_{\mathrm{p}}}$,
2. $T \geq N\Delta\tau_{\mathrm{p}} = \frac{4N\log 2}{\Delta\omega_l}\sqrt{1+\chi^2}$,
3. $t_{\mathrm{res}} \leq \frac{\tau_{\mathrm{p}}}{M}$,
4. $\omega_{\mathrm{res}} \leq \frac{\omega_l}{M}$.

The most important thing to note is that our time and frequency grids are connected when we take the Fourier transform, such that

$$\mathrm{d}t = \frac{2\pi}{P\mathrm{d}\omega}. \tag{2.130}$$

From this, we can write four restrictions that come about from the relation in the spectral and temporal grids

1. $P\mathrm{d}t = \frac{2\pi}{\mathrm{d}\omega} > T$,
2. $P\mathrm{d}\omega = \frac{2\pi}{\mathrm{d}t} > W$,
3. $\mathrm{d}t = \frac{2\pi}{P\mathrm{d}\omega} < t_{\mathrm{res}}$,
4. $\mathrm{d}\omega = \frac{2\pi}{P\mathrm{d}t} > \omega_{\mathrm{res}}$,

which we may then combine for two inequalities; one for the spectral domain, one for the time domain

$$\max\left[\frac{W}{P}, \frac{2\pi}{Pt_{\mathrm{res}}}\right] < \mathrm{d}\omega < \min\left[\frac{2\pi}{T}, \omega_{\mathrm{res}}\right], \tag{2.131}$$

$$\max\left[\frac{T}{P}, \frac{2\pi}{P\omega_{\mathrm{res}}}\right] < \mathrm{d}t < \min\left[\frac{2\pi}{P}, t_{\mathrm{res}}\right]. \tag{2.132}$$

In theory, these are equivalent, but we will always ensure that these conditions are both met for the same minimum number of grid points. We can then use the values of our grid size and spacing requirements based on the pulse parameters (given in the first set of restrictions). In all cases, we found P, then found the lowest $P^{'} = 2^{Q} + 1 > P$ number of grid points. As expected, these conditions were always met for the same number of grid points.

For all simulations, we always checked that the Fourier transforms were calculated as expected. We can derive an exact expression for both the temporal profile of the input Gaussian pulse and its spectrum; see Section 2.1.2. Taking the numerical Fourier transform of each allows direct comparison with the converse analytic form, thus verifying the FFT. Figure 2.1 shows the result of these calculations starting from either the temporal profile or the spectrum and Fourier transforming in both directions to ensure that the converse is recovered. The pulse parameters used are consistent with those used throughout this thesis.

Spatial and Spatial Frequency Grids

Using the Collins method requires taking the discrete spatial Fourier transform of the input field. The spatial grid at the input plane can be discretised as

$$x_0 \in \left[-\frac{\Delta L_{x0}}{2}, \frac{\Delta L_{x0}}{2}\right] \quad \rightarrow \quad x_0 = \frac{m\Delta L_{x0}}{N}; \quad m = -\frac{N}{2}, \ldots, \frac{N}{2} - 1. \tag{2.133}$$

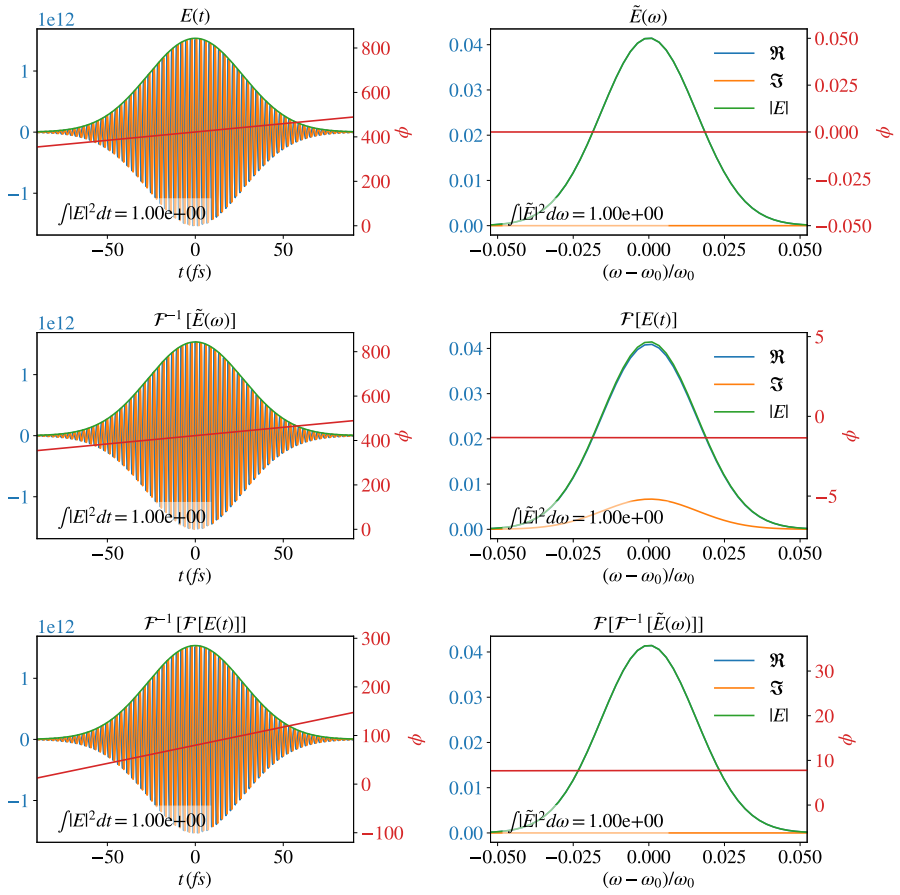


Figure 2.1: Comparison of starting with both analytic expressions for temporal profile $E(t)$ and spectrum $\tilde{E}(\omega)$ and Fourier transforming or inverse Fourier transforming to recover the other.

Considering an input grid of width ΔL_{x0} and an output grid of width ΔL_x, both sampled by N grid points, then the spatial frequency bandwidth of the input frequency spectrum is $N/\Delta L_{x0}$. This relates to the output grid spacing by

$$\Delta L_x = \frac{B\lambda}{\Delta L_{x0}/N}, \tag{2.134}$$

where B is the upper-right matrix element. We can therefore calculate the output spatial grid using the `fftfreq` function (native to Python) of the input grid, multiplied by λB. However, we would rather that the grid spacing would not change with propagation distance, so we must be careful to calculate the requirements based on the grid of values for different axial positions z, before taking the overall largest number of grid points at any position. This ensures that our requirements are met throughout the entire simulation.

The spot size of the central wavelength component λ_0 at the focus w_{0f} relates to the input

spot size w_0 by

$$w_{0f} = \frac{\lambda_0 f_0}{\pi w_0}, \tag{2.135}$$

where f_0 is the focal length of the lens. Therefore we can calculate the widths of our input and output grids as

$$\Delta L_{x0} > N w_0, \tag{2.136}$$

$$\Delta L_x > N w_{0f} = \frac{N \lambda_0 f_0}{\pi w_0}. \tag{2.137}$$

As before, N is an integer deemed to be large enough to view the behaviour of the pulse. This defines the width of our simulation grid. It can be shown[180] that the number of grid points required to correctly perform the discrete spatial Fourier transform is given by

$$N = \frac{\Delta L_{x0}\, \Delta L_x}{\lambda_0 B}. \tag{2.138}$$

This was used to calculate the number of grid points required to properly map the input and output grids to each other. In cases where the pulse is more spatially-structured than a lowest-order Gaussian beam, such as for a focusing, multimode, Laguerre-Gaussian beam (see Chapter 5), a more involved consideration of the spatial grids was be required. Trial and error was used to ensure that the fields were correctly spatially-resolved.

2.3.2 Comparison Between Collins and Analytic Methods

Where an analytic solution can be derived, the least-squares error (LSE) of the 3D spatiotemporal fields for the two methods is calculated at every axial location as

$$\mathrm{LSE}(z) = \sum_M \frac{\left|\tilde{E}_{\mathrm{analytic}} - \tilde{E}_{\mathrm{Collins}}\right|^2}{M}, \tag{2.139}$$

where M is the total number of grid points in the simulation (e.g. $M_x * M_y * M_t$). The least-squares error $\mathrm{LSE}(z)$ over the entire window is checked to make sure that it is $< 1\%$, to ensure that the pulse is correctly resolved and that there is good agreement between the Collins and analytic methods.

3 | The Flying Focus

3.1 Introduction

A flying focus (FF) is a laser focal spot which can appear to travel at any arbitrary velocity.[172,182,183] This may include superluminal velocities (without violating relativity[172]) and negative velocities, such that the laser pulse focus appears to travel backwards towards the source.[167,172,179,182,184,185] Laser- and plasma-based applications of flying foci[186] include driving ionisation waves with arbitrary velocities,[167,187] dephasingless laser wakefield acceleration,[46,188,189] studies of radiation reaction[190] and the generation of light from Raman amplification,[191–193] photon acceleration,[194] and Thomson scattering.[195] Flying foci (FFs) can be produced using standard chromatic optics or plasma mirrors.[196,197] It is possible to generate flying foci using colourless optics, for example via cross-phase modulation[198] or by using stepped mirrors.[199] A paper published in 2024 demonstrated this by using an axiparabola to control the apparent focal velocity, without introducing temporal chirps which lengthen the pulse duration.[200] However, often a coupling between space and time is introduced via the wavelength. For example, a standard approach for generating FFs is to pass a temporally-chirped pulse[†] through a chromatic optic.[172,182]

Consider Figure 3.1, which shows a positively temporally-chirped pulse (in which longer wavelengths come first in time) that is focused by a chromatic lens which focuses longer wavelengths more strongly. One could imagine that longer wavelengths focus both earlier and closer to the lens. This creates an FF which propagates forwards (towards the right in Figure 3.1) at an arbitrary velocity which may be controlled by adjusting the input chirp and chromaticity of the lens.

[†]This is different to the *spatially* chirped pulses which shall be discussed in Chapters 4 and 5, we will therefore refer to it as temporal chirp to make the distinction clear; in most textbooks and papers this is what is meant when they say 'chirp'

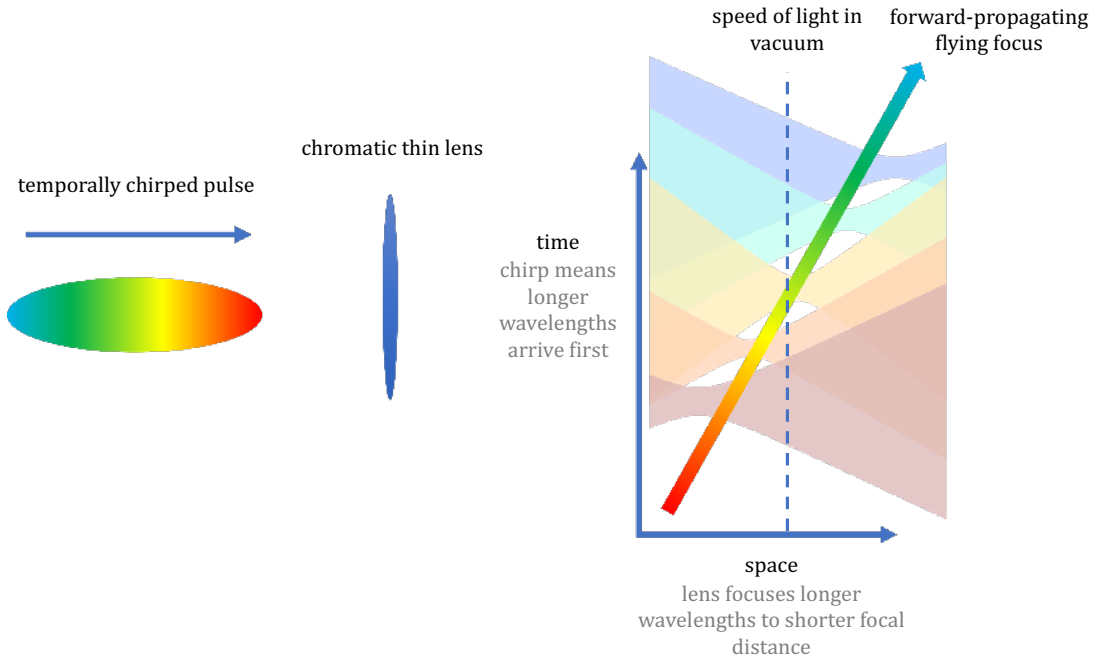


Figure 3.1: Schematic of a flying focus in the light-speed frame. A positively temporally-chirped pulse passes through a chromatic lens which focuses longer wavelengths more strongly.

It has been shown that flying foci (FFs) can overcome dephasing[†],[46] a critical limitation to energy gain in laser wakefield accelerators (LWFAs). FFs may also have applications for electron injection by temporarily decreasing the wake velocity[188,189] or by allowing for an optical injector pulse which is distinct from the wakefield driver and can move at a controlled speed relative to the wakefield. One can imagine that with free control over how the optical injector moves, we could 'paint' electrons into the phase space of the accelerator.

In Section 3.2.3, we shall first present a derivation of the analytic FF based on a bi-Gaussian beam. This derivation requires that simplifying assumptions are made and is only presented for a bi-Gaussian beam. Although perfectly Gaussian beams do not exist in nature, they do allow for exact mathematical solutions which may not be possible for other spectral or spatial profiles. In Section 3.2.4, an expression for the velocity of the focal spot is derived. Section 3.3 then defines the optical ray matrix required to propagating flying foci with arbitrary spectra, based on Fourier optics methods; specifically the Collins diffraction integral (details on the Collins diffraction integral were given in Section 2.3). Both methods are used to propagate FF beams and the results are presented in Section 3.4. In the final section, we propagate an FF with a

[†]One of the three key limitations for a LPA: defocusing, depletion and dephasing. Dephasing occurs when electrons (travelling at almost the speed of light) outrun the plasma wave, which has a group velocity less than the speed of light in plasma. See Section 1.3.1 for more detail.

non-Gaussian spectra.

3.2 Theoretical Analysis of the FF

3.2.1 Definition of Input bi-Gaussian Pulse

The envelope $\tilde{U}(\mathbf{r},\omega)$ of a bi-Gaussian pulse at the input plane $z=0$ may be generally given as (see Sections 2.1.2 and 2.1.4)

$$\tilde{U}_{\text{in}}(\rho,z=0,\omega)=U_{\omega}\tilde{U}(\omega)\exp\left\{-\frac{\rho^2}{w_0^2}\right\}, \tag{3.1}$$

where $U_{\omega}=|\mathbf{U}_{\text{in}}|/\omega_0$ is a normalisation constant related to the peak electric field strength and central frequency of the laser pulse. We shall assume that our pulse has a Gaussian spectrum with some temporal chirp. Linear spectral chirp arises from second-order temporal phase

$$U(t)=\mathrm{e}^{-\Gamma t^2}\mathrm{e}^{-i\omega_0 t}, \tag{3.2}$$

where $\Gamma=a+ib$ is a complex parameter that characterises the temporal width and temporal chirp of the pulse. Therefore the spectrum is (see Section 2.1.2)

$$\tilde{U}(\omega)=\sqrt{\frac{1}{2\Gamma}}\mathrm{e}^{\frac{(\omega-\omega_0)^2}{4\Gamma}}, \qquad \Gamma=a+ib, \qquad \chi=b/a, \tag{3.3}$$

where we have used the parameter χ to describe the strength of the temporal chirp. It is possible to show that the full width half maximum (FWHM) pulse duration (of the pulse intensity) is given by

$$\tau_{\text{p}}=\sqrt{\frac{2\ln 2}{\Re\{\Gamma\}}}=\sqrt{\frac{2\ln 2}{a}}, \tag{3.4}$$

and the FWHM pulse bandwidth (of the pulse intensity) is given by

$$\Delta\omega=\sqrt{\frac{8\ln 2}{\Re\{1/\Gamma\}}}=\sqrt{8\ln 2}\sqrt{a\left[1+\left(\frac{b}{a}\right)^2\right]}, \tag{3.5}$$

such that the time-bandwidth product is given by

$$\tau_{\text{p}}\Delta\omega=4\ln 2\sqrt{1+\left(\frac{b}{a}\right)^2}=4\ln 2\sqrt{1+\chi^2}, \tag{3.6}$$

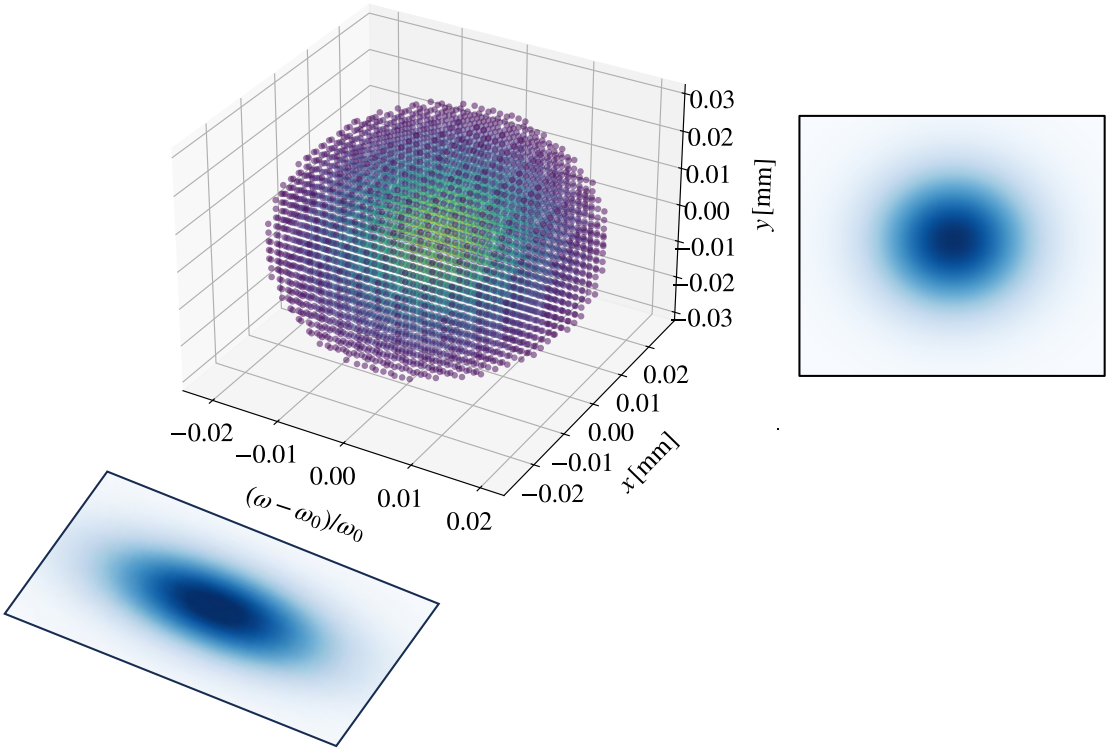


Figure 3.2: Schematic of a bi-Gaussian pulse in the spectral domain, lower panel shows the projection onto the $x-\omega$ plane and right panel shows the projection onto the $x-y$ plane.

by writing this as

$$\tau_{\mathrm{p}} = \frac{4\ln 2}{\Delta\omega}\sqrt{1+\chi^2} = \tau_{\mathrm{p0}}\sqrt{1+\chi^2}, \tag{3.7}$$

it is clear to see that increasing the temporal chirp of the pulse always increases the pulse duration from the fully-compressed limit where $\tau_{\mathrm{p0}} = 4\ln 2/\Delta\omega$. This completes our definition of a Gaussian input pulse, which is shown in Figure 3.2.

KEY PARAMETER χ: describes the temporal chirp of the incoming pulse.

3.2.2 Chromatic Lens

Next we pass this pulse through a lens which has a frequency-dependent focal length, assuming a linear dependence on frequency of the form

$$f(\omega) = f_0 + \beta(\omega-\omega_0), \tag{3.8}$$

where f_0 is the position of the nominal focus for a frequency ω_0. Clearly β parameterises the strength of the chromaticity; the higher the magnitude of β, the greater the distance between the focal position f of a non-central frequency component and the nominal focal position f_0. In this case, a positive chromaticity $\beta > 0$ causes longer wavelengths to be focused more strongly, as shown in Figure 3.1.

KEY PARAMETER β: describes the chromatic aberration of the lens.

3.2.3 Derivation of the Gaussian Flying Focus

Assuming that Eq. 3.1 describes the input spatiospectral field $\tilde{U}_{\text{in}}$ at the plane of the chromatic lens, we can multiply the field given by Eq. 3.1 by the phase term introduced at the lens such that

$$\tilde{U}_{\text{post-lens}}(\rho) = U_\omega \tilde{U}(\omega) \exp\left\{\frac{-\rho^2}{w_0^2}\right\} \exp\left(-\frac{ik\rho^2}{2f(\omega)}\right). \tag{3.9}$$

We then need to propagate the field a distance z to the focal region using a Fresnel integral. The Fresnel diffraction integral is given in Cartesian coordinates by[201]

$$\tilde{U}_{\text{out}}(x,y) = \frac{e^{ikz}}{\sqrt{i\lambda z}} \iint_{-\infty}^{\infty} \tilde{U}_{\text{post-lens}}(x_{\text{in}}, y_{\text{in}}) e^{\frac{ik}{2z}\left[(x-x_{\text{in}})^2+(y-y_{\text{in}})^2\right]} dx_{\text{in}} dy_{\text{in}}. \tag{3.10}$$

We could equivalently write the Fresnel integral in cylindrical coordinates, but here we use the 1D Fresnel integral along either Cartesian axis to describe the behaviour of the flying focus. We can show this since

$$\begin{aligned}\tilde{U}_{\text{out}}(x,y) = U_\omega \tilde{U}(\omega) \frac{e^{ikz}}{\sqrt{i\lambda z}} &\int_{-\infty}^{\infty} \exp\left\{-\left[\frac{1}{w_0^2} - \frac{ik}{2f(\omega)}\right] x_{\text{in}}^2\right\} e^{\frac{ik}{2z}(x-x_{\text{in}})^2} dx_{\text{in}} \\ &\int_{-\infty}^{\infty} \exp\left\{-\left[\frac{1}{w_0^2} - \frac{ik}{2f(\omega)}\right] y_{\text{in}}^2\right\} e^{\frac{ik}{2z}(y-y_{\text{in}})^2} dy_{\text{in}}, \qquad (3.11)\\ &= \tilde{U}_{\text{out}}(x)\, \tilde{U}_{\text{out}}(y), \qquad (3.12)\end{aligned}$$

such that the propagation in the x and y-directions is independent. We shall drop the y-dependence, knowing that it will be exactly the same. This gives

$$\tilde{U}_{\text{out}}(x) = U_\omega \tilde{U}(\omega) \frac{e^{ikz}}{\sqrt{i\lambda z}} \int_{-\infty}^{\infty} \exp\left\{-\left[\frac{1}{w_0^2} - \frac{ik}{2f(\omega)}\right] x_{\text{in}}^2\right\} e^{\frac{ik}{2z}(x-x_{\text{in}})^2} dx_{\text{in}}, \tag{3.13}$$

$$= U_\omega \tilde{U}(\omega) \frac{e^{ikz}}{\sqrt{i\lambda z}} \int_{-\infty}^{\infty} \exp\left\{ -\left[\frac{1}{w_0^2} - \frac{ik}{2f(\omega)} + \frac{ik}{2z} \right] x_{\text{in}}^2 - \frac{ikx}{z} x_{\text{in}} + \frac{ikx^2}{2z} \right\} \mathrm{d}x_{\text{in}}, \quad (3.14)$$

where we have collected like-powers of x. Eq. 3.14 fully describes the propagation of the flying focus. We can then use the following standard integral for the Gaussian[202]

$$\int_{-\infty}^{\infty} e^{-ax^2+bx+c} \, \mathrm{d}x = \sqrt{\frac{\pi}{a}} e^{\frac{b^2}{4a}+c}, \quad (3.15)$$

to obtain

$$\tilde{U}_{\text{focal region}}(x,z,\omega) = U_\omega \tilde{U}(\omega) \frac{e^{ikz}}{\sqrt{i\lambda z}} \sqrt{\frac{\pi}{\frac{1}{w_0^2} - \frac{ik}{2f(\omega)} + \frac{ik}{2z}}}$$
$$\exp\left\{ \frac{1}{4} \left(\frac{ikx}{z} \right)^2 \left[\frac{1}{w_0^2} - \frac{ik}{2f(\omega)} + \frac{ik}{2z} \right]^{-1} + \left(\frac{ikx^2}{2z} \right) \right\}, \quad (3.16)$$

$$= U_\omega \tilde{U}(\omega) \frac{e^{ikz}}{\sqrt{\Omega}} \exp\left\{ \frac{ikx^2}{2z} \left(1 - \frac{1}{\Omega} \right) \right\}, \quad (3.17)$$

where

$$\Omega = \frac{z}{f(\omega)} - 1 + \frac{2iz}{kw_0^2}. \quad (3.18)$$

Generally, the spectrum $\tilde{U}(\omega)$ and frequency-dependent focal length $f(\omega)$ may take any form, parameterised by χ and β respectively. In the next section, we shall derive the focal velocity, as dependent on both these parameters.

3.2.4 Derivation of the Focal Velocity

The focal velocity is given by

$$v_{\text{ff}} = \frac{\mathrm{d}z_\omega}{\mathrm{d}t_\omega} \simeq \frac{\partial z_\omega}{\partial \omega} \Big/ \frac{\partial t_\omega}{\partial \omega}, \quad (3.19)$$

where we can note that t_ω is the frequency-dependent arrival time and z_ω is the frequency-dependent focal position. We know that

$$z_\omega = f(\omega) = f_0 + \beta(\omega - \omega_0), \quad (3.20)$$

$$t_\omega = t_{\text{delay}} + \frac{z_{\text{path}}}{c} = t_{\text{delay}} + \frac{\beta(\omega - \omega_0)}{c}, \quad (3.21)$$

where we have recognised that the axial distance z_{path} from the nominal focal position f_0 to the focal position of a component with frequency ω is given by $\beta(\omega-\omega_0)$. We can also note that t_{delay} is the relative delay of each frequency component due to the temporal chirp of the input pulse. This time delay is related to the spectral phase φ by $t_{\text{delay}} = \partial\varphi/\partial\omega$. The imaginary part of the spectrum given by Eq. 3.3 is related to the spectral phase φ of the pulse. We can write the spectrum as

$$\tilde{U}(\omega) = \sqrt{\frac{1}{2\Gamma}} e^{\Re\left\{\frac{(\omega-\omega_0)^2}{4\Gamma}\right\}} e^{i\Im\left\{\frac{(\omega-\omega_0)^2}{4\Gamma}\right\}}, \tag{3.22}$$

$$= \sqrt{\frac{1}{2\Gamma}} e^{\Re\left\{\frac{(\omega-\omega_0)^2}{4\Gamma}\right\}} e^{-i\varphi}, \tag{3.23}$$

such that the spectral phase φ from the spectrum is

$$\varphi = -\frac{1}{\Im\{4\Gamma\}}(\omega-\omega_0)^2. \tag{3.24}$$

The total spectral phase from Eq 3.17 is therefore

$$\varphi(x,\omega) = -\frac{1}{\Im\{4\Gamma\}}(\omega-\omega_0)^2 + \frac{kx^2}{2z}\left\{1 - \frac{1}{\Omega^*\Omega}\left[\frac{z}{f(\omega)}+1\right]\right\}, \tag{3.25}$$

where Ω^* is the complex conjugate of Ω. Considering the on-axis time delay, where $x = 0$, we obtain

$$t_{\text{delay}} = \left.\frac{\partial\varphi}{\partial\omega}\right|_{x=0} = -\frac{1}{\Im\{2\Gamma\}}(\omega-\omega_0). \tag{3.26}$$

Using our definitions of the pulse duration τ_{p} and bandwidth $\Delta\omega$ of the pulse intensity from Eqs. 3.4 and 3.5 respectively, we find that

$$\frac{\tau_{\text{p}}}{\Delta\omega} = \sqrt{\frac{2\ln 2}{8\ln 2}}\frac{1}{a\sqrt{1+\chi^2}} = \frac{1}{2a\sqrt{1+\chi^2}}. \tag{3.27}$$

Writing this in terms of the fully-compressed pulse duration $t_{p0} = t_p/\sqrt{1+\chi^2}$ from Eq. 3.7, we obtain

$$\frac{\tau_{\text{p0}}}{\Delta\omega} = \frac{1}{2a(1+\chi^2)}, \tag{3.28}$$

which allows us to write

$$-\Im\left\{\frac{1}{2\Gamma}\right\} = -\Im\left\{\frac{a-ib}{2(a^2+b^2)}\right\} = \frac{b}{2a^2(1+\chi^2)} = \frac{\tau_{\mathrm{p0}}}{\Delta\omega}\chi, \tag{3.29}$$

such that

$$t_{\mathrm{delay}} = \frac{\tau_{\mathrm{p0}}}{\Delta\omega}\chi(\omega-\omega_0). \tag{3.30}$$

Returning to Eqs. 3.20 and 3.21, we find that

$$z_\omega = f_0 + \beta(\omega-\omega_0), \tag{3.31}$$

$$t_\omega = \left(\frac{\tau_{\mathrm{p0}}}{\Delta\omega}\chi + \frac{\beta}{c}\right)(\omega-\omega_0). \tag{3.32}$$

Taking the partial derivatives of both expressions with respect to ω gives

$$\frac{\partial z_\omega}{\partial\omega} = \beta, \tag{3.33}$$

$$\frac{\partial t_\omega}{\partial\omega} = \left(\frac{\tau_{\mathrm{p0}}}{\Delta\omega}\chi + \frac{\beta}{c}\right). \tag{3.34}$$

Therefore, the focal velocity v_{ff} is

$$v_{\mathrm{ff}} \simeq \frac{\partial z_\omega}{\partial\omega}\Big/\frac{\partial t_\omega}{\partial\omega} = \frac{\beta}{\frac{\tau_{\mathrm{p0}}}{\Delta\omega}\chi + \frac{\beta}{c}}, \tag{3.35}$$

which may be rewritten as a fraction of the speed of light c as

$$\frac{v_{\mathrm{ff}}}{c} \simeq \left(1 + \frac{c\tau_{\mathrm{p0}}}{\Delta\omega}\frac{\chi}{\beta}\right)^{-1}. \tag{3.36}$$

Figure 3.3a) shows the expected focal velocity v_{ff} calculated for a range of different temporal chirp χ and chromaticity β strengths according to Eq. 3.36. It is useful to rescale β relative to β_{crit}

$$\beta_{\mathrm{crit}} = \left|\frac{c\tau_{\mathrm{p0}}^2}{4\ln 2}\chi\right|. \tag{3.37}$$

which corresponds to the value where using Eqs. 3.36 and 3.6 with $-\beta_{\mathrm{crit}}$ we get

$$\frac{v_{\mathrm{ff}}}{c} \simeq \left(1 - \frac{\tau_{\mathrm{p0}}}{\Delta\omega}\frac{4\ln 2}{\tau_{\mathrm{p0}}^2}\right)^{-1} = \left(1 - \frac{4\ln 2}{\tau_{\mathrm{p0}}\Delta\omega}\right)^{-1} = \left(1 - \frac{1}{\sqrt{1+\chi^2}}\right)^{-1} \to \infty. \tag{3.38}$$

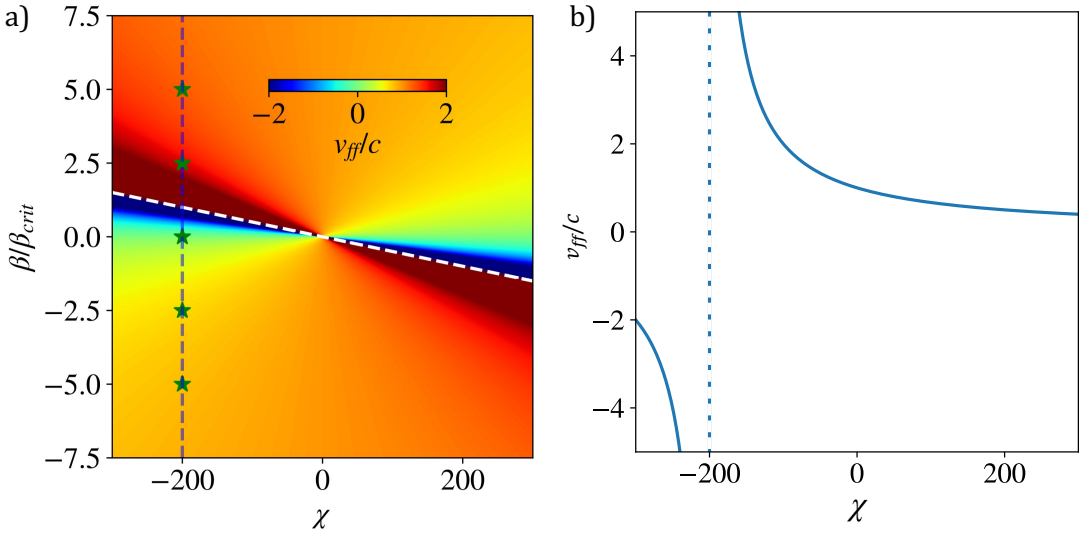


Figure 3.3: Plot of v_{ff} for different values of χ and β. a) Shows a map of the flying focus velocity v_{ff} as a function of temporal chirp χ and lens chromaticity β. Dashed white line denotes the asymptote to infinite focal velocity. Blue dashed line indicates the fixed value $\chi = -200$ where samples for Figure 3.5 were taken, these samples are indicated with green stars. b) Shows the variation of v_{ff} as a function of the temporal chirp parameter χ, according to Eq. 3.36, for the fixed value of $\beta = \beta_{\mathrm{crit}}$.

This is where the velocity becomes infinite as the temporal chirp tends to zero. As the temporal chirp becomes infinite our focal spot travels at c. A lineout at a fixed chromaticity of $\beta = \beta_{\mathrm{crit}}$ is shown in Figure 3.3b). It is evident from these plots that any arbitrary focal velocity (whether in the forward or backward direction or at sub- or super-luminal velocities), is theoretically possible.

3.3 Ray Matrix to Generate an FF

In the previous section, we found an analytic method to propagate temporally-chirped pulses. This method allows for a direct calculation of the SSTF field at any z position and is therefore quick to implement. We can see directly the form that the fields take and the way that they evolve. An alternative approach is to use diffraction integrals, such as the Collins integral, to propagate the fields. With the Collins method, we only need to specify the input field and the ray matrices that describe the optical system that this field propagates through. The benefit of this method is that one is not restricted to assuming any particular form of the input field, and we can easily simulate more complex systems with multiple optics. As shown in Figure 3.4, we send the temporally-chirped input pulse through a chromatic thin lens with a frequency-dependent focal length $f(\omega)$, before propagating it to a focus in free-space. This is described by the following

matrices

$$O = \begin{pmatrix} A & B \\ C & D \end{pmatrix} = \begin{pmatrix} 1 & z - z_0 \\ 0 & 1 \end{pmatrix} \begin{pmatrix} 1 & 0 \\ -1/f(\omega) & 1 \end{pmatrix} = \begin{pmatrix} 1 - (z - z_0)/f(\omega) & z - z_0 \\ -1/f(\omega) & 1 \end{pmatrix}. \tag{3.39}$$

These matrix elements depend on z and ω. This means that each matrix element must be calculated for all relevant values of z and ω, and programatically this is done by calculating each matrix element as a 2D grid. Full details of the Collins method were given in Section 2.3.

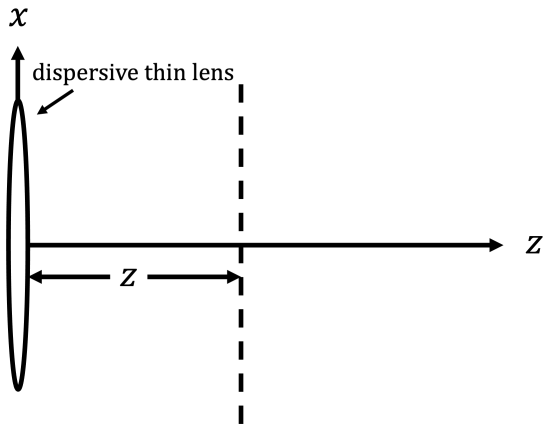


Figure 3.4: Diagram of the 2D Cartesian coordinate system. A thin lens with a frequency-dependent focal length $f(\omega)$ is positioned at the origin. The output pulse is then propagated after the lens by a distance z which is in a narrow range about the focus f_0.

3.4 Results

We assume the following laser pulse parameters:

Pulse energy, $\mathscr{E}$	1 J
Central wavelength, λ_0	800 nm
Uncompressed pulse duration, τ_{p0}	10 fs
Input spot size, w_0	25 mm
Nominal focal length, f_0	1 m

Table 3.1: Table of laser pulse parameters for FF simulations

In this section, we study the evolution of the FF beam, in particular, the apparent focal velocity. Eq. 3.17 gives an expression for the spatiospectral field after the chromatic lens. Numerical calculation of the spatiospectral fields is performed using the Collins method described in the previous section. To avoid unnecessary duplication, results are plotted for the analytic method only (see Section 2.3.2). Both methods give $\tilde{U}_{\text{out}}(x, y, \omega)$ at a particular axial position z. The

spatiospectral field $\tilde{U}_{\text{out}}(x, y, \omega)$ is then numerically Fourier transformed to the time-domain to give $U_{\text{out}}(x, y, t)$ at each axial position z; repeating this for different longitudinal z-positions in a grid that spans a number of Rayleigh ranges around the focus gives the full spatiotemporal profile $U_{\text{out}}(x, y, z, t)$. The standard definition of the Rayleigh range after a lens $z_{R,\text{out}} = f_0^2 z_{\text{R}}/(z_{\text{R}}^2 + f_0^2)$ is used to normalise the longitudinal axis, where $z_{\text{R}} = \pi w_0^2/\lambda_0$ is the Rayleigh range of the central frequency component before the lens.

In a frame that is moving at the speed of light, the total propagation distance $|z - z_0|$ from the input plane at z_0 corresponds to a time of occurrence $t_0 = |z - z_0|/c$. This is the time at which $t = 0$ in the moving-frame coordinate coincides with the point z. We have calculated the temporal profile of the pulse passing through this point z, and the field in the light-speed moving-frame coordinate $\xi = z - ct$ can therefore be calculated from temporal profile of the pulse $U_{\text{out}}(x, y, t)$ at that z position, by simply rescaling the time grid by a factor of c, such that $\xi_{\text{grid}} = ct_{\text{grid}}$, since $t_{\text{grid}} = 0$ is the time of arrival of a light-speed pulse. In this way we can convert our evolution of the pulse at different axial z-positions to snapshots at different times, allowing us to visualise the evolution of the pulse in a fixed moving-frame window ξ.

Figure 3.5 shows how the apparent velocity of the focal spot changes when we vary the chromaticity of the lens. The temporal chirp is fixed at a value of $\chi = -200$, although we could equivalently vary the chirp χ for a fixed value of the lens chromaticity β. Here the velocity of the flying focus ranges from $v_{\text{ff}} = 0.83c$ to $v_{\text{ff}} = 1.67c$. The velocity of the flying focus is calculated from the simulated focal spot data as follows. Before transformation to the moving-frame coordinate, for each axial distance $|z - z_0|$ at which the fields are calculated, the relative time t_0 at which the peak intensity occurs is found by taking a lineout of the intensity at $x = 0$. The flying focus velocity is estimated to be

$$\frac{v_{\text{ff}}}{c} = \frac{|z - z_0|}{\big(|z - z_0|/c - t_0\big)/c}. \tag{3.40}$$

The flying focus velocity calculated from Eqs. 3.36 and 3.40 agree with an error of $< 1\%$. The calculated velocities of all focal spots are given in Figure 3.5. The offset time t_0 is converted to a position in the moving-frame coordinate and is plotted as dashed lines in Figure 3.5. The solid lines denote the expected position of the focal spot without spatiotemporal couplings.

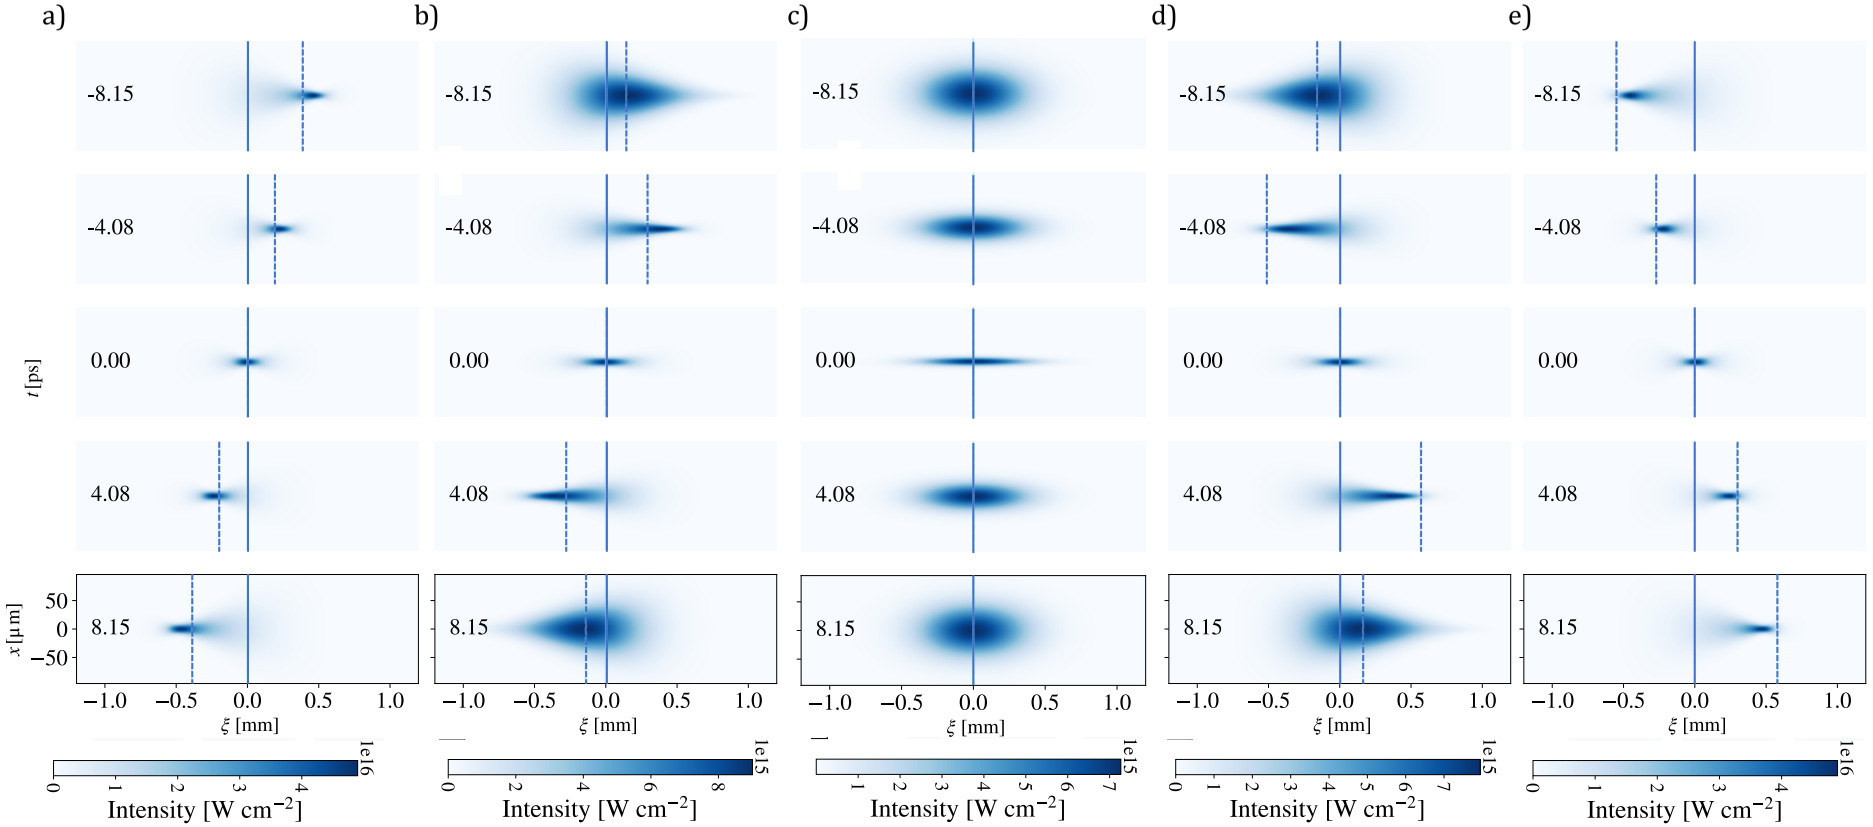


Figure 3.5: Figure showing snapshots of the flying focus in the light-speed frame around the focus for a fixed input chirp of $\chi = -200$ and varied lens chromaticity β. a) $\beta = -5\beta_{\text{crit}}$ with $v_{\text{ff}} = 0.83c$, b) $\beta = -2.5\beta_{\text{crit}}$ with $v_{\text{ff}} = 0.71c$, c) the non-flying focus formed with an achromatic lens where $\beta = 0$ with $v_{\text{ff}} = c$, d) $\beta = 2.5\beta_{\text{crit}}$ with $v_{\text{ff}} = 1.67c$ and e) $\beta = 5\beta_{\text{crit}}$ with $v_{\text{ff}} = 1.25c$. Dashed lines show the position of the flying focus peak intensity, solid lines show the position of a non-flying focus where $\beta = 0$.

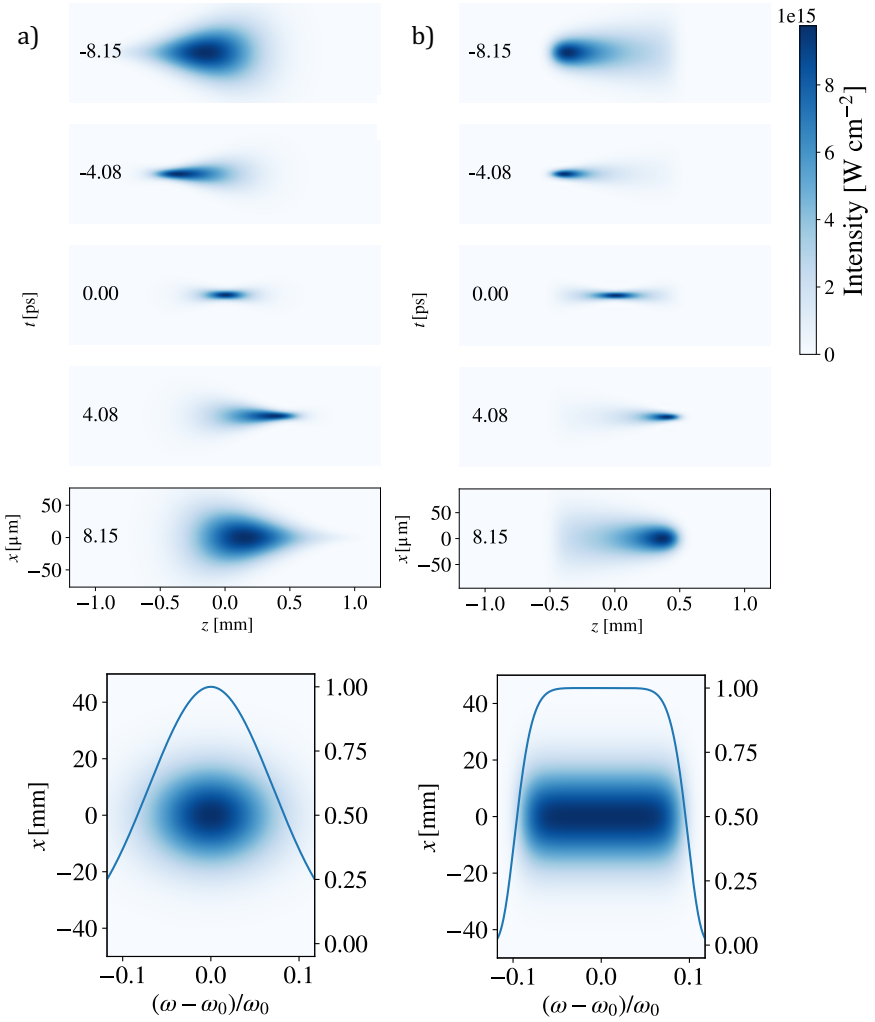


Figure 3.6: Propagation of flying foci with arbitrary spectra, for $\chi = -200$ and $\beta = 2.5\beta_{\text{crit}}$. The focal spot has a calculated focal velocity of $v_{\text{ff}} = 1.67c$ a) Shows the flying focus generated using a Gaussian input spectrum and b) shows the flying focus generated using a super-Gaussian input spectrum. The spectra are shown in the lowermost panes.

3.4.1 Alternative Spectra: Super-Gaussian

We can see from Eq. 3.17 that our derivation did not depend on the form of the spectrum $\tilde{U}(\omega)$. Therefore, it is possible to find the output fields of the flying focus for a non-Gaussian input spectrum. Figure 3.6 compares the analytic solution for the flying focus parameters of $\chi = -200$ and $\beta = 2.5\beta_{\text{crit}}$ for two different spectra. Figure 3.6a) shows the result for a Gaussian spectrum[†] and Figure 3.6b) shows the propagation of an alternative beam with a super-Gaussian spectrum to the 8^{th} power in ω. As always the Collins method shows good agreement. Although the focal velocity does not change, it is apparent from Figure 3.6 that the spatial properties of the focal

[†]Also plotted in Figure 3.5d)

spot change when we change the spectrum. The key difference is that the super-Gaussian has a focal spot which is extended longitudinally. Although these differences are not marked, it is an illustration that the analytic solution derived is valid for arbitrary spectra and could even be used to calculate flying foci using experimentally measured spectra. In addition, one could imagine that shaping of the laser input spectrum could be used to tailor the longitudinal intensity profile of the laser at the focus, with applications for beam-loading in an ionisation injection scheme.

3.5 Conclusions and Future Work

In this chapter, the fields of a bi-Gaussian flying focus laser pulse were derived using the Fresnel diffraction integral. From the spectral phase in the analytic expression for the spatiospectral field of the flying focus, we found an expression for the apparent velocity of the focal spot, which depended on the strength of the input chirp of the pulse χ and the chromaticity β of the thin lens used to focus the beam. This expression used to map the parameter space, and we found that by varying the strength of the input chirp of the pulse χ and the lens chromaticity β we could achieve focal spots with arbitrary velocities, including focal spots which travel backwards with respect to the source and superluminal focal spots which appear to travel faster than the speed of light. The spatiospectral fields were calculated both analytically and using the Collins diffraction integral. We also calculated the focal velocities directly from the spatiotemporal fields in the light-speed frame coordinate, and we found an agreement with our analytic expression for the focal velocities of $< 1\%$. Finally, we showed that the analytic solution is general to arbitrary spectra and that a super-Gaussian spectrum gives rise to a flying focus with a modified longitudinal profile.

In the future, we plan to investigate the maximum possible propagation distance of a bi-Gaussian flying focus analytically. This will be used to investigate whether flying foci could allow for longitudinal density tailoring of hydrodynamic optical-field-ionised (HOFI) channels, which should be tens of centimetres long. An alternative study of ionisation injection with a flying focus will also be attempted using particle-in-cell (PIC) codes, to investigate whether it is possible to paint electrons into the phase-space, using a flying focus to control beam-loading. It is anticipated that flying focus LWFA drivers will become popular for dephasingless laser-wakefield acceleration; the techniques developed here could also be used to design experiments such that the Oxford LPAG group can implement such techniques.

4 | The Linear Simultaneous Space-Time Focus

4.1 Introduction

Simultaneous space-time focusing (SSTF) is a phenomenon that has been used for well over a decade in the fields of microscopy and laser machining.[203–207] It localises time-dependent processes to a particular axial depth by 'focusing' the pulse in time to achieve full temporal recompression (shortest pulse duration) only at the focus. Applications include multiphoton imaging[203–205,207,208] and laser machining.[206,207]

SSTF is supposed to lead to tighter localisation of the region of high intensity near the focus of the laser pulse. The localisation of the axial intensity can be defined by identifying the axial distance from the point of highest intensity to where the intensity has dropped to half of the peak intensity, where a shorter distance indicates a higher degree of axial localisation. Intuitively, this would make sense; as the pulse focuses spatially, the increase in local bandwidth causes the pulse duration to decrease towards the geometric focus. We would therefore expect that the peak intensity increases as the pulse is compressed into a shorter time window. However, introducing a spatial chirp also increases the transverse size of the input pulse by spreading out the different frequency components across the input plane, which reduces the focal spot size at focus (therefore also the Rayleigh range z_R). It is therefore important to understand to what extent the improved localisation that is seen arises simply from the effect of tighter focusing and therefore whether SSTF itself offers any advantage. There is considerable confusion in the literature, with conflicting statements from different authors. In the original literature on the topic[203–205] it is never stated that SSTF offers a greater degree of axial localisation, only that it can localise a time-dependant process due to the reduction in pulse duration. Indeed, in

one paper it is explicitly stated that "the axial characteristics of SSTF [in two-photon excitation fluorescence (TPEF)] are determined by the input spatial profile".[204] This seems to be in contrast with other papers which all make similar statements about the "axial localisation of intensity with SSTF",[209,210] apparently without justification or citation to corroborate such statements. In addition, a later paper by Durfee *et al.* appears to show a greater degree of axial localisation of an SSTF pulse in one of their figures.[211]

In this chapter, we aim to establish whether SSTF does or does not improve the axial localisation of the high intensity region of the focus, provided that one compares the SSTF and non-SSTF cases in a fair way. We show that one has to be mindful of spatiotemporal couplings when characterising or describing ultrashort laser pulses, to avoid reaching conclusions which are at best paradoxical and at worst false. By simulating the intensity profile of the pulse using multiple techniques and defining a fair comparison case which undergoes equivalent spatial focusing and is defined mathematically in Section 4.4.3, we show that SSTF does *not* improve intensity localisation and furthermore, that simplistic estimates of the intensity profile of the pulse can lead to an overestimate of the degree of intensity localisation. To simulate the propagation of SSTF and non-SSTF pulses, we use both the analytic solution for the full evolution, alongside a complementary numerical simulation, to study the degree of axial localisation of the laser field. We show that the degree of axial localisation is the same for SSTF and non-SSTF pulses with equivalent spatial focusing, and demonstrate that heuristic methods used to analyse focal spots fail when spatiotemporal couplings are present.

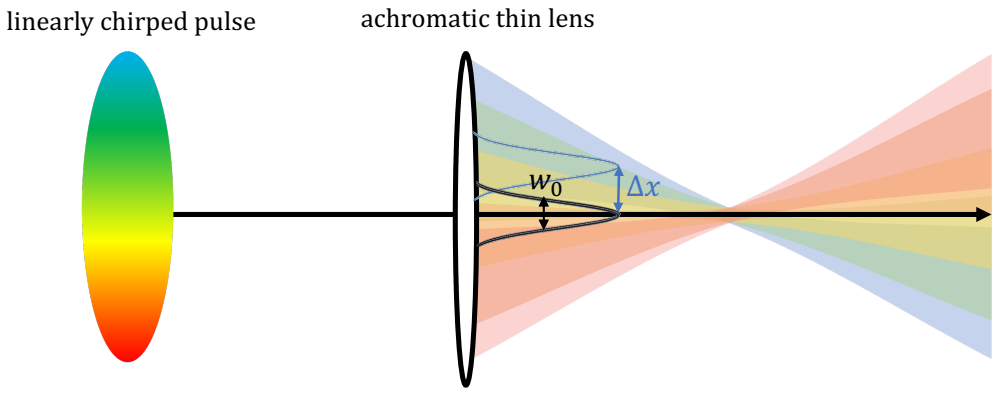


Figure 4.1: Schematic of a linear simultaneous space-time focus. A pulse which is chirped transversely along one axis passes through an achromatic lens. The nominal Gaussian beamlet is shown for the central frequency ω_0 in black and for an arbitrary beamlet with a different frequency ω which is transversely offset by a distance Δx in blue.

Generating simultaneous space-time foci (SSFT) usually involves chirping the input pulse transversely, before passing it through an achromatic lens or reflecting off a focusing mirror, as shown in Figure 4.1. This case will be referred to as an SSTF or 'linear' SSTF, to distinguish it from the case where the input pulse is chirped radially, as will be studied in Chapter 5. At the input plane, the local bandwidth at any point in space is low, such that the pulse duration is long by the time-bandwidth product. As the pulse comes to a focus, the different spectral components begin to overlap. This decreases the pulse duration. It is important to note that SSTF does not reduce the pulse duration at the temporal focus compared to the fully-compressed pulse duration that would result from the global spectral width, instead, it lengthens the pulse duration everywhere else.

Supposing that a process of interest can only occur where the pulse duration is sufficiently short, then this process will be localised to a particular axial depth for an SSTF. There are a few processes which are dependent on the pulse duration. An example of such a process is two photon excitation fluorescence (TPEF), where the molecules' two-photon absorption cross section is inversely proportional to the excitation pulse duration.[212] This means that fluorescence microscopy which utilises SSTF will have an enhanced signal-to-noise ratio, since molecules at the imaging focus experience a shorter pulse duration and are more likely to fluoresce, thereby suppressing the out-of-focus background fluorescence. An addendum to this is that one can scan the sample axially by adjusting the second-order spectral phase of the incident pulse,[204] which is particularly useful in endoscopic and biological applications where moving the sample (particularly where samples are in solution or not affixed to a substrate) is not desirable.

If the region of highest intensity is also more localised using SSTF, then this has implications for purely intensity-dependent processes too. An example of a purely intensity-dependent process is the ionisation of gases for ionisation injection in laser-wakefield accelerators (LWFA). We shall study whether the intensity is more localised for an SSTF pulse by comparing the on-axis peak intensity as a function of z (distance along the optical axis) for an SSTF pulse and a fair comparison case. To do this we shall simulate the propagation of these pulses using two methods, the first is by plotting the analytic solution and the second is by performing a numerical simulation using the Collins diffraction integral.[173]

4.2 Theoretical Analysis of the SSTF

SSTF can be modelled by considering the input pulse as a summation of monochromatic Gaussian beamlets of $1/e^2$ input spot size w_0 (defined in Section 2.1.4), which are offset from the optical axis by a transverse distance $\Delta x(\omega) = \alpha(\omega - \omega_0)/\alpha_{\text{norm}}$. This distance is therefore proportional to the frequency of each beamlet ω relative to the on-axis central frequency ω_0. We assume that the superposition of these monochromatic beamlets has a Gaussian spectrum, peaked at the central frequency ω_0. The chirp scaling factor α is used to control the direction and strength of the spatial chirp and is scaled by the ratio $\alpha_{\text{norm}} = \Delta\omega/w_0$, where $\Delta\omega$ is the FWHM input bandwidth.

KEY PARAMETER α: describes the spatial chirp of the incoming pulse.

Mathematically, the input field at the axial location $z = 0$ may then be written as

$$\tilde{U}_{\text{in}}(x, y, z = 0, \omega) = U_\omega \tilde{U}(\omega) \exp\left\{ - \frac{[x - \Delta x(\omega)]^2 + y^2}{w_0^2} \right\}, \tag{4.1}$$

where

$$\Delta x(\omega) = \frac{\alpha}{\alpha_{\text{norm}}}(\omega - \omega_0), \tag{4.2}$$

and $U_\omega = |\mathbf{U}_{\text{in}}|/\omega_0$ is a normalisation constant related to the peak electric field strength and central frequency of the laser pulse. The Gaussian input spectrum is $\tilde{U}(\omega) = 1/\sqrt{2\Gamma} \exp[(\omega - \omega_0)^2/4\Gamma]$, where $\Gamma = a + ib$ is a complex parameter that characterises the temporal width and chirp of the pulse, with $a = 2\ln 2/\tau_{\text{p}}^2$ and $b = \chi a$ (see Section 2.1.2 for more details). The pulse duration is τ_{p} and χ is a normalised parameter that defines the input chirp. For simplicity, we assume the nominal case that the pulse is not chirped, such that $b = 0$. We could assume otherwise without affecting any subsequent derivation, but then the focal plane of the pulse no longer coincides with the geometric focus,[204] making our fair comparison more complicated.

The input fields are shown in the spatiospectral and spatiotemporal domains in Figure 4.2, where in the spectral domain shown in Figure 4.2a), the correlation between frequency and transverse position can clearly be seen. Figure 4.2b) shows that in the temporal domain, the linearly-chirped pulse has a longer FWHM pulse duration, 83.18 fs, compared to the unchirped case, as shown in Figure 4.2b′), where the fully-compressed FWHM pulse duration is 45 fs. Com-

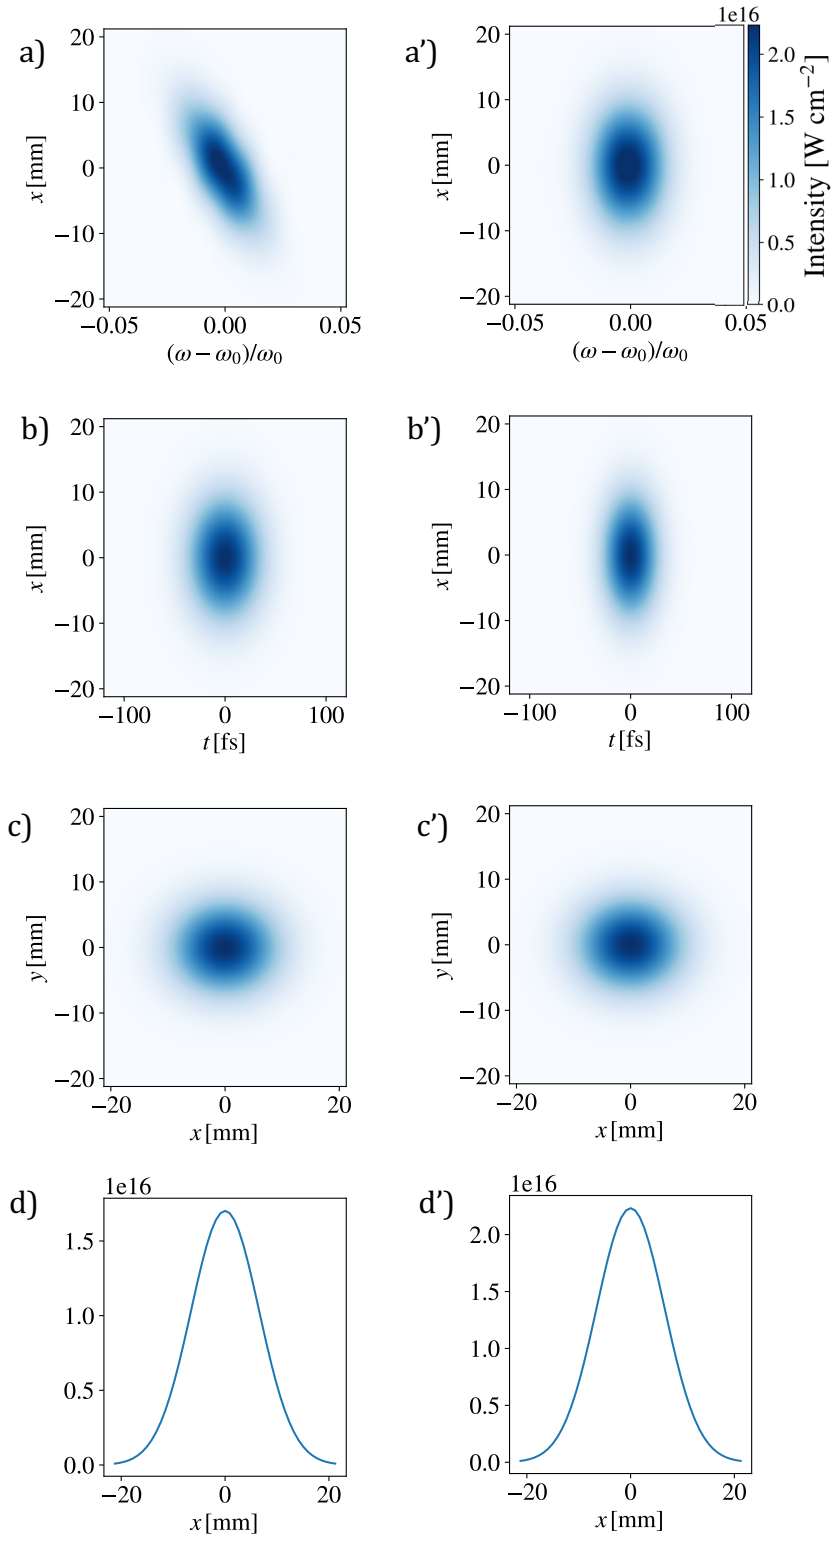


Figure 4.2: Comparison of the intensity distributions of the input pulses for linearly-chirped pulses with $\alpha = -1$ and $\alpha = 0$ (i.e. an unchirped pulse). Figures a)–d) show the chirped case and Figures a′)–d′) show the unchirped fair-comparison (see Section 4.4.3). a) shows the spatiospectral domain where the correlation between spatial and frequency offset can be seen, compared to a′) where there is no correlation, b) shows the spatiotemporal domain, where it can be seen that the FWHM pulse duration is longer for all transverse positions in the chirped case at 83.18 fs compared to the unchirped case b′) where the fully-compressed FWHM pulse duration is 45 fs. Comparing c) and d) with c′) and d′), we see that the spatial profile is exactly the same, so our spatial focusing will be equivalent in both cases.

paring Figures 4.2c) and c′), we see that the spatial profile of the chirped pulse is the same as with the unchirped fair comparison case (see Section 4.4.3). Figures 4.2d) and d′) show a lineout of the intensity at $y = 0$. The fair comparison will be discussed in Section 4.4.3.

4.2.1 Derivation of the Linear Simultaneous Space-Time Focus

We will now derive the output field of a linear SSTF pulse, following previous approaches.[203,211] Like the derivation for a Flying Focus (see Chapter 3), this is based on a Fresnel integral. There are three steps to the derivation: i) propagating the spatially-chirped pulse a certain distance f_0, ii) passing it through a lens and iii) propagating the pulse again some distance z beyond the lens. We will follow all three steps to avoid a loss of generality and to allow for comparison to previous texts, but it should be noted that in the case where step i) is skipped and the input pulse is defined at the plane of the lens, the derivation is simpler (see Figure 4.3).

The Fresnel diffraction integral is given by[201]

$$\tilde{U}_{\text{out}}(x,y) = \frac{e^{ikz}}{\sqrt{i\lambda z}} \iint_{-\infty}^{\infty} \tilde{U}_{\text{in}}(x_{\text{in}}, y_{\text{in}}) \exp\left\{\frac{ik}{2z}\Big[(x-x_{\text{in}})^2 + (y-y_{\text{in}})^2\Big]\right\} \mathrm{d}x_{\text{in}}\mathrm{d}y_{\text{in}}. \tag{4.3}$$

Starting from the input spatiospectral field $\tilde{U}_{\text{in}}$ of the spatially chirped pulse given by Eq. 4.1, the integral can be written

$$\begin{aligned}\tilde{U}_{\text{out}}(x,y) = U_\omega \tilde{U}(\omega) \frac{e^{ikz}}{\sqrt{i\lambda z}} \int_{-\infty}^{\infty} \exp\left\{-\frac{[x_{\text{in}} - \Delta x_{\text{in}}(\omega)]^2}{w_0^2}\right\} \exp\left\{\frac{ik}{2z}(x-x_{\text{in}})^2\right\} \mathrm{d}x_{\text{in}} \\ \int_{-\infty}^{\infty} \exp\left\{-\frac{y_{\text{in}}^2}{w_0^2}\right\} \exp\left\{\frac{ik}{2z}(y-y_{\text{in}})^2\right\} \mathrm{d}y_{\text{in}}, \quad (4.4)\end{aligned}$$

$$= \tilde{U}_{\text{out}}(x)\, \tilde{U}_{\text{out}}(y), \tag{4.5}$$

such that the propagation in the x and y-directions is independent. Since the variation with z in the y direction is well-known from the theory of Gaussian pulses, we shall drop y dependence and consider only the x-dependent integral given by

$$\tilde{U}_{\text{out}}(x) = U_\omega \tilde{U}(\omega) \frac{e^{ikz}}{\sqrt{i\lambda z}} \int_{-\infty}^{\infty} \exp\left\{-\frac{[x_{\text{in}} - \Delta x_{\text{in}}(\omega)]^2}{w_0^2}\right\} e^{\frac{ik}{2z}(x-x_{\text{in}})^2} \mathrm{d}x_{\text{in}}. \tag{4.6}$$

The assumptions used to calculate the SSTF propagation in previous literature are that i) propagation can be modelled by the paraxial approximation and ii) $k \approx k_0 = \omega_0/c$, but this is only valid for pulse durations of approximately $\sim 100\,\text{fs}$ or longer.[205] Even in a vacuum, $k = \omega/c \neq \omega_0/c$.

Therefore, unlike previous texts, we will not make this approximation and will assume that the angular wavenumber is related to the angular frequency as $k = \omega/c$. This only becomes problematic if we want to find an explicit closed-form solution for the pulse in the time domain, because it is not possible to take the Fourier transform analytically without setting $k = k_0$. However, to find the pulse in the time domain, we can still take a Fourier transform of the full spectral profile, and we will do this numerically.

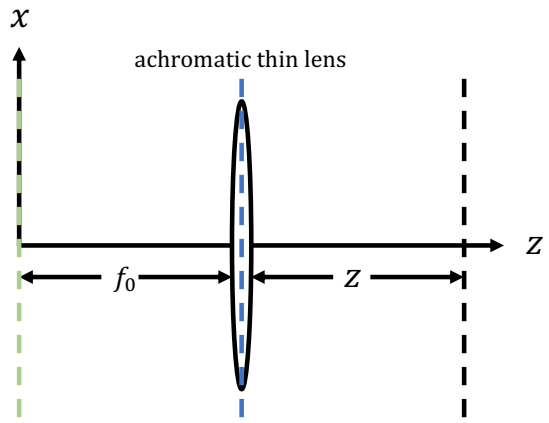


Figure 4.3: Diagram of the 2D Cartesian coordinate system. A thin lens is positioned a distance f_0 away from the origin. The input pulse travels a distance f_0 from the input plane (green dashed line) before propagating through the lens. The output pulse is then propagated after the lens by a distance z which is within a narrow range about the focus f_0. For simplicity, the derivation is done in 1D; there are no spatiotemporal couplings in the y dimension so the pulse propagates independently as a Gaussian beam. A simplification can be made where the pulse is defined at the input plane of the lens (dashed blue line).

As described by Eq. 4.1, the linearly-chirped input pulse is

$$\tilde{U}_{\text{in}}(x) = U_\omega \tilde{U}(\omega) \exp\left\{\frac{-[x - \Delta x(\omega)]^2}{w_0^2}\right\}. \tag{4.7}$$

We will use Eq. 4.6 to propagate this field the focal distance to the lens such that

$$\tilde{U}_{\text{pre-lens}}(x) = U_\omega \tilde{U}(\omega) \frac{e^{ikf_0}}{\sqrt{i\lambda f_0}} \int_{-\infty}^{\infty} \exp\left\{\frac{-[x_{\text{in}} - \Delta x(\omega)]^2}{w_0^2}\right\} e^{\frac{ik}{2f_0}(x - x_{\text{in}})^2} \, dx_{\text{in}}, \tag{4.8}$$

$$= U_\omega \tilde{U}(\omega) \frac{e^{ikf_0}}{\sqrt{i\lambda f_0}} \int_{-\infty}^{\infty} \exp\left\{-\left[\frac{1}{w_0^2} - \frac{ik}{2f_0}\right] x_{\text{in}}^2 + \left[\frac{2\Delta x(\omega)}{w_0^2} - \frac{ikx}{f_0}\right] x_{\text{in}} + \left[\frac{ikx^2}{2f_0} - \frac{\Delta x(\omega)^2}{w_0^2}\right]\right\} dx_{\text{in}}, \tag{4.9}$$

where we have collected like powers of x_{in}. We can then use the following standard integral for the Gaussian[202]

$$\int_{-\infty}^{\infty} e^{-ax^2 + bx + c} \, dx = \sqrt{\frac{\pi}{a}} e^{\frac{b^2}{4a} + c}, \tag{4.10}$$

to obtain

$$\tilde{U}_{\text{pre-lens}}(x) = U_\omega \tilde{U}(\omega) \frac{e^{ikf_0}}{\sqrt{i\lambda f_0}} \sqrt{\frac{\pi}{\frac{1}{w_0^2} - \frac{ik}{2f_0}}} \exp\left\{ \frac{1}{4}\left[\frac{2\Delta x(\omega)}{w_0^2} - \frac{ikx}{f_0} \right]^2 \left(\frac{1}{w_0^2} - \frac{ik}{2f_0} \right)^{-1} + \frac{ikx^2}{2f_0} - \frac{\Delta x(\omega)^2}{w_0^2} \right\}, \tag{4.11}$$

$$= U_\omega \tilde{U}(\omega) \frac{e^{ikf_0}}{\sqrt{1 + \frac{2if_0}{kw_0^2}}} \exp\left\{ \frac{1}{4}\left[\frac{2\Delta x(\omega)}{w_0^2} - \frac{ikx}{f_0} \right]^2 \left(\frac{1}{w_0^2} - \frac{ik}{2f_0} \right)^{-1} + \frac{ikx^2}{2f_0} - \frac{\Delta x(\omega)^2}{w_0^2} \right\}. \tag{4.12}$$

The argument of the exponential can be written as

$$\frac{\left[\frac{2\Delta x(\omega)}{w_0^2} - \frac{ikx}{f_0} \right]^2}{4\left(\frac{1}{w_0^2} - \frac{ik}{2f_0} \right)} + \frac{ikx^2}{2f_0} - \frac{\Delta x(\omega)^2}{w_0^2}, \tag{4.13}$$

$$= \frac{\frac{\Delta x(\omega)^2}{w_0^4} - \frac{ikx\Delta x(\omega)}{f_0 w_0^2} - \frac{k^2x^2}{4f_0^2} + \frac{ikx^2}{2f_0 w_0^2} - \frac{\Delta x(\omega)^2}{w_0^4} + \frac{k^2x^2}{4f_0^2} + \frac{ik\Delta x(\omega)^2}{2f_0 w_0^2}}{\frac{1}{w_0^2} - \frac{ik}{2f_0}}, \tag{4.14}$$

$$= \frac{\frac{ik}{2f_0 w_0^2}\left[x^2 - 2x\Delta x(\omega) + \Delta x(\omega)^2 \right]}{\frac{1}{w_0^2} - \frac{ik}{2f_0}}, \tag{4.15}$$

$$= \frac{-k[x - \Delta x(\omega)]^2}{2if_0 + kw_0^2}, \tag{4.16}$$

such that

$$\tilde{U}_{\text{pre-lens}}(x) = U_\omega \tilde{U}(\omega) \frac{e^{ikf_0}}{\sqrt{1 + \frac{2if_0}{kw_0^2}}} \exp\left\{ \frac{-k[x - \Delta x(\omega)]^2}{2if_0 + kw_0^2} \right\}. \tag{4.17}$$

At the lens, the field is multiplied by the phase term introduced by the lens such that

$$\tilde{U}_{\text{post-lens}}(x) = U_\omega \tilde{U}(\omega) \frac{e^{ikf_0}}{\sqrt{1 + \frac{2if_0}{kw_0^2}}} \exp\left\{ \frac{-k[x - \Delta x(\omega)]^2}{2if_0 + kw_0^2} \right\} \exp\left(- \frac{ikx^2}{2f_0} \right). \tag{4.18}$$

Propagating the field a distance z to the focal region using a second Fresnel integral gives

$$\tilde{U}_{\text{focal region}}(x) = U_\omega \tilde{U}(\omega) \frac{e^{ikf_0}}{\sqrt{1 + \frac{2if_0}{kw_0^2}}} \frac{e^{ikz}}{\sqrt{i\lambda z}} \int_{-\infty}^{\infty} \exp\left\{ \frac{-k[x_{\text{in}} - \Delta x(\omega)]^2}{2if_0 + kw_0^2} \right\} \exp\left(- \frac{ikx_{\text{in}}^2}{2f_0} \right) e^{\frac{ik}{2z}(x - x_{\text{in}})^2} \, dx_{\text{in}}. \tag{4.19}$$

Using the substitution $w_1^2 = (2if_0 + kw_0^2)/k$, this becomes very similar to our previous integral, so we can approach it in exactly the same way, obtaining

$$= U_\omega \tilde{U}(\omega) \frac{e^{ikf_0}}{\sqrt{1+\frac{2if_0}{kw_0^2}}} \frac{e^{ikz}}{\sqrt{i\lambda z}} \int_{-\infty}^{\infty} \exp\left\{ -\left(\frac{1}{w_1^2} - \frac{ik}{2z} + \frac{ik}{2f_0}\right)x_{\text{in}}^2 + \left[\frac{2\Delta x(\omega)}{w_1^2} - \frac{ikx}{z}\right]x_{\text{in}} + \left[\frac{ikx^2}{2z} - \frac{\Delta x(\omega)^2}{w_1^2}\right]\right\} \mathrm{d}x_{\text{in}}, \quad (4.20)$$

$$= U_\omega \tilde{U}(\omega) \frac{e^{ikf_0}}{\sqrt{1+\frac{2if_0}{kw_0^2}}} \frac{e^{ikz}}{\sqrt{i\lambda z}} \sqrt{\frac{\pi}{\frac{1}{w_1^2} - \frac{ik}{2z} + \frac{ik}{2f_0}}} \exp\left\{ \frac{1}{4}\left[\frac{2\Delta x(\omega)}{w_1^2} - \frac{ikx}{f_0}\right]^2 \left(\frac{1}{w_1^2} - \frac{ik}{2z} + \frac{ik}{2f_0}\right)^{-1} + \frac{ikx^2}{2f_0} - \frac{\Delta x(\omega)^2}{w_1^2}\right\}, \quad (4.21)$$

$$= U_\omega \tilde{U}(\omega) \frac{e^{ikf_0}}{\sqrt{1+\frac{2if_0}{kw_0^2}}} \frac{e^{ikz}}{\sqrt{(1-\frac{z}{f_0}) + \frac{2iz}{kw_1^2}}} \exp\left\{ \frac{1}{4}\left[\frac{2\Delta x(\omega)}{w_1^2} - \frac{ikx}{f_0}\right]^2 \left(\frac{1}{w_1^2} - \frac{ik}{2z} + \frac{ik}{2f_0}\right)^{-1} + \frac{ikx^2}{2f_0} - \frac{\Delta x(\omega)^2}{w_1^2}\right\}. \quad (4.22)$$

The argument of the exponential can be written as follows

$$= \frac{\left[\frac{2\Delta x(\omega)}{w_1^2} - \frac{ikx}{f_0}\right]^2}{4\left(\frac{1}{w_1^2} - \frac{ik}{2z} + \frac{ik}{2f_0}\right)} + \frac{ikx^2}{2f_0} - \frac{\Delta x(\omega)^2}{w_1^2}, \quad (4.23)$$

$$= \frac{ik}{2f_0}\left\{ \frac{ik}{2f_0}\left[\frac{1}{w_1^2} - \frac{ik}{2z}\left(1-\frac{z}{f_0}\right)\right]^{-1}\left[x + \frac{2if_0}{kw_1^2}\Delta x(\omega)\right]^2 + x^2 + \frac{2if_0}{kw_1^2}\Delta x(\omega)^2\right\}, \quad (4.24)$$

$$= \frac{ik}{2f_0}\left\{ \frac{-z}{\Omega f_0}\left[x + \Xi\Delta x(\omega)\right]^2 + x^2 + \Xi\Delta x(\omega)^2\right\}, \quad (4.25)$$

where

$$\Omega = \left(1 - \frac{z}{f_0}\right) + \frac{2iz}{kw_1^2}, \quad (4.26)$$

$$\Xi = \frac{2if_0}{kw_1^2} = i\frac{f_0}{z}\Im\{\Omega\} = \frac{f_0}{z}\frac{\Omega - \Omega^*}{2}, \quad (4.27)$$

and we can note that we can write the exponential argument solely in terms of Ω

$$= \frac{ik}{2f_0}\left\{ \left(1 - \frac{z}{\Omega f_0}\right)x^2 - \left(\frac{\Omega - \Omega^*}{\Omega}\right)x\Delta x(\omega) + \frac{f_0}{z}\left(\frac{\Omega - \Omega^*}{2}\right)\left(1 - \frac{\Omega - \Omega^*}{2\Omega}\right)\Delta x(\omega)^2\right\}. \quad (4.28)$$

The field after the lens is therefore

$$\tilde{U}_{\text{focal region}}(x,z,\omega) = U_\omega \tilde{U}(\omega) \frac{e^{ik(z+f_0)}}{\sqrt{\Omega(1+\frac{2if_0}{kw_0^2})}} \exp\left(\frac{ik}{2f_0}\left\{\frac{-z}{\Omega f_0}\Big[x+\Xi\Delta x(\omega)\Big]^2 + x^2 + \Xi\Delta x(\omega)^2\right\}\right), \tag{4.29}$$

with Ω and Ξ defined above, and $w_1^2 = (2if_0 + kw_0^2)/k$. We now have two options. As already mentioned, if we wish to derive an analytic expression for the time-dependent field, we require the approximation that $k \approx k_0$. Although this allows us to derive an analytic expression, it is only valid for pulses longer than $\sim 100\,\text{fs}$. Since we are interested in pulses shorter than this, we leave this step to other texts and instead retain the full spectral dependence as given by Eq. 4.29. The spatiospectral field given by Eq. 4.29 is then numerically Fourier transformed into the time domain for arbitrary $k(\omega)$.

4.3 Ray Matrix to Generate an SSTF

In the previous section, we found an analytic method to propagate linearly-chirped pulses. This method allows for a direct calculation of the SSTF field at any z position. We can see directly the form that the fields take and the way that they evolve. An alternative approach is to use diffraction integrals, such as the Collins integral, to propagate the fields. With the Collins method, we only need to specify the input field and the ray matrices that describe the optical system that this field propagates through. The benefit of this method is that one is not restricted to assuming any particular form of the input field, and we can also simulate more complex systems with multiple optics. As shown in Figure 4.3, we propagate the spatially chirped input pulse a distance f_0 from the the input plane (shown by the dashed green line) before passing it through a non-dispersive thin lens. We then propagate it to a focus in free-space. This is described by the following matrices

$$O = \begin{pmatrix} A & B \\ C & D \end{pmatrix} = \begin{pmatrix} 1 & z-z_0 \\ 0 & 1 \end{pmatrix} \begin{pmatrix} 1 & 0 \\ -1/f_0 & 1 \end{pmatrix} \begin{pmatrix} 1 & f_0 \\ 0 & 1 \end{pmatrix} = \begin{pmatrix} 1-(z-z_0)/f_0 & f_0 \\ -1/f_0 & 0 \end{pmatrix}. \tag{4.30}$$

These matrix elements depend on z and ω. This means that each matrix element must be calculated for all relevant values of z and ω, and programatically this is done by calculating each matrix element as a 2D grid. Full details of the Collins method were given in Section 2.3.

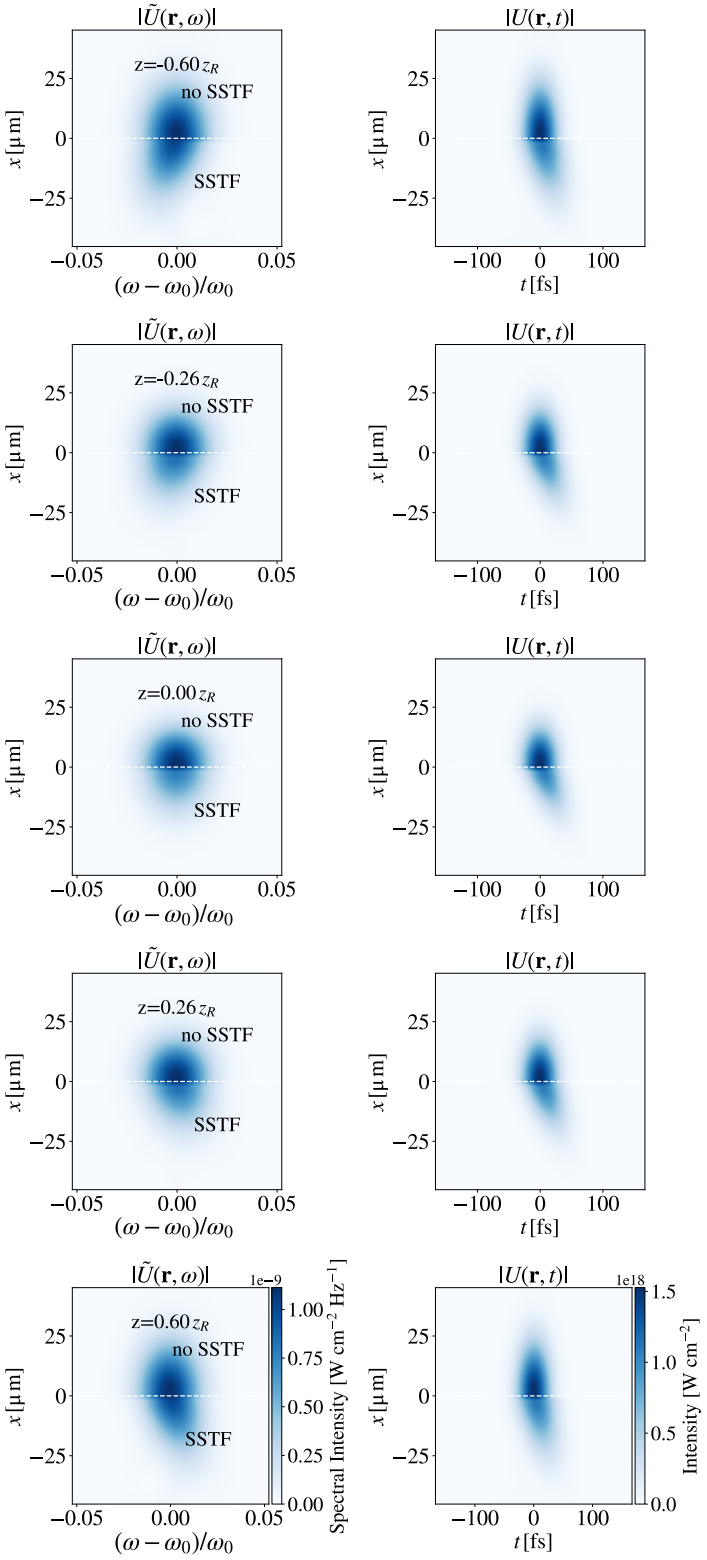


Figure 4.4: Figure showing the evolution of the linear SSTF around the focus. Spatiospectral intensity (left column) is Fourier transformed to the time domain to give the spatiotemporal intensity (right column). The upper half of each plot shows the intensity of the field without SSTF ($\alpha = 0$) and the lower half shows the intensity of the field with SSTF ($\alpha = -1$). Pulse-front tilt (seen as a slant of the spatiotemporal field) is only present for the pulse with SSTF and is strongest at the focus.

4.4 Results

We assume the following laser pulse parameters:

Pulse energy, $\mathscr{E}$	1 J
Central wavelength, λ_0	800 nm
Uncompressed pulse duration, τ_{p0}	45 fs
Input spot size, w_0	10 mm
Nominal focal length, f_0	1 m

Table 4.1: Table of laser pulse parameters for SSTF simulations

In this section, we study the evolution of some key features and parameters of the SSTF pulse, before studying the degree of axial localisation. Eq. 4.29 gives an analytic expression for the spatiospectral field after the lens. Numerical calculation of the spatiospectral fields is performed using both Eq. 4.29 and the Collins method described in the previous section. To avoid unnecessary duplication, results are plotted for the Collins method only (see Section 2.3.2). Both methods give $\tilde{U}_{out}(x,y,\omega)$ at a particular axial position z. The spatiospectral field $\tilde{U}_{out}(x,y,\omega)$ is then numerically Fourier transformed to the time-domain to give $U(x,y,t)$ at each axial position z; repeating this for different longitudinal axial positions z in a grid that spans a number of Rayleigh ranges around the focus gives the full spatiotemporal profile $U_{out}(x,y,z,t)$. The expected Rayleigh range $z'_R = f_0^2 z_R/(z_R^2+f_0^2)$ of a Gaussian pulse of Rayleigh range z_R focused by a lens of focal length f_0 is used to normalise the longitudinal axis for both the SSTF and non-SSTF fields, where $z_R = \pi w_0^2/\lambda_0$ is the Rayleigh range of the beamlet before the lens. In performing the calculations, a hybrid grid is used which has a finer mesh for the region within half a Rayleigh range of the focus.

Figure 4.4 shows the spatiospectral (left column) and spatiotemporal (right column) fields at different axial positions z relative to the focus f_0. The upper half of each plot shows the evolution of a focus which has no SSTF ($\alpha = 0$), whereas the lower half of each plot shows a focus with SSTF ($\alpha = -1$). In both cases, the pulse comes to a focus at $z = 0$ before diverging again. For the SSTF case, all spectral components overlap on the optical axis only at the focus. With increasing spectral overlap, the on-axis pulse duration shortens; maximum temporal compression occurs at the focus as expected (this is more obvious in Figure 4.5b) and will be discussed in Section 4.4.1). Pulse-front tilt (PFT),[174] as previously observed for an SSTF pulse,[213] is evident as a correlation

in the $x-t$ domain that is strongest at the focus. It has previously been shown that the PFT cannot be compensated for by applying a counter-acting PFT to the input pulse.[214]

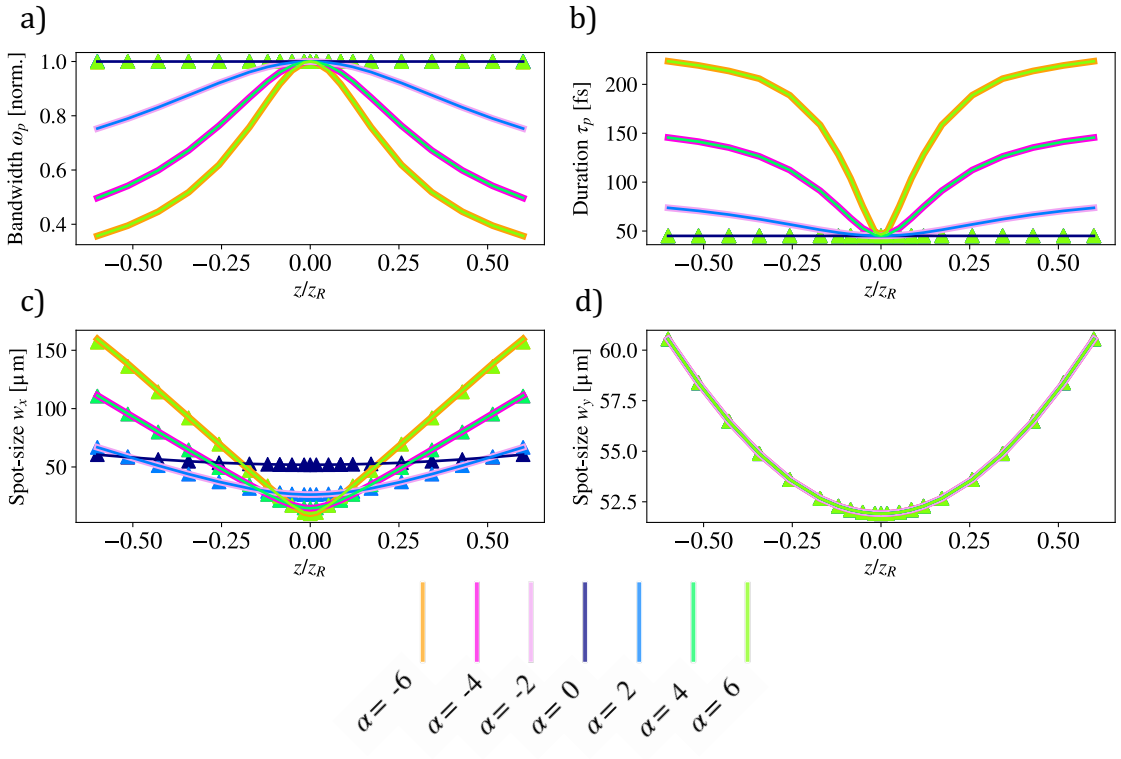


Figure 4.5: Key parameters of the linear SSTF pulse calculated using the Collins method as a function of z in the focal region showing: a) the FWHM bandwidth $\Delta\omega$, b) the FWHM pulse duration τ_p, c) the $1/e^2$ spot size $w_x(z)$ and d) the $1/e^2$ spot size $w_y(z)$. Solid lines denote the SSTF pulse ($\alpha \neq 0$) and triangles are the corresponding non-SSTF fair comparison ($\alpha = 0$) defined in Section 4.4.3.

4.4.1 Centroid, Ellipsoidal Spot Size, Bandwidth and Pulse Duration

In order to understand the evolution of the peak intensity of the SSTF pulse in the focal region, we first determine the longitudinal dependence of the key parameters of the pulse — i.e. the spot size, bandwidth, and pulse duration — as follows.

For each longitudinal position z, the time at which the peak axial intensity occurs is found. At this time, a cross-section of the pulse in the xy-plane is taken. Along one transverse axis, the D4σ beam diameter is given in terms of the second moment of the intensity distribution $I = \varepsilon_0 c|\mathbf{E}|^2$ (we will describe in Section 4.4.2 how the intensity is calculated from the fields) as[215,216]

$$\mathrm{D4}\sigma_x = 4\sqrt{\frac{\int_{-\infty}^{\infty} I(x,y)(x-\bar{x})^2\,\mathrm{d}x\,\mathrm{d}y}{\int_{-\infty}^{\infty} I(x,y)\,\mathrm{d}x\,\mathrm{d}y}}, \tag{4.31}$$

where the x-coordinate centroid is given by

$$\bar{x} = \frac{\int_{-\infty}^{\infty} I(x,y)x\,\mathrm{d}x\,\mathrm{d}y}{\int_{-\infty}^{\infty} I(x,y)\,\mathrm{d}x\,\mathrm{d}y}. \tag{4.32}$$

The $1/\mathrm{e}^2$ Gaussian spot size w_x can be calculated from the D4σ_x beam diameter as

$$w_x = \frac{\mathrm{D4}\sigma_x}{2}. \tag{4.33}$$

The FWHM bandwidth and pulse duration are calculated at each axial distance z by sitting on the optical axis and recording the intensity profile that passes in the spectral and temporal domains, respectively. The width of the pulse between the two values that are half of the peak intensity is calculated to give the FWHM spectral bandwidth and pulse duration. Calculation of the FWHM bandwidth and pulse duration could be improved by instead fitting a Gaussian to the data at each axial location.

The evolution of key parameters is plotted in Figures 4.5a)–d). We can see the increase in spot size in the chirped direction in Figure 4.5c). Figures 4.5a) and b) confirm that, for the SSTF case, the bandwidth increases (and the pulse duration decreases) as the focus is approached. Away from the focus, the pulse duration can increase from the fully-compressed value of 45 fs by as much as a factor of 4.5 for the highest chirp considered ($\alpha = \pm 6$). Therefore, time-dependent processes will still be localised using the SSTF beam.

We can now look at what happens to the degree of axial localisation of intensity when we increase the magnitude of the spatial chirp. Figure 4.6 shows the absolute peak intensity $I_{\mathrm{peak}} = \max\{\varepsilon_0 c|\mathbf{E}|^2\}$ (denoted by solid lines, further detail on how this is calculated will be given in Section 4.4.2). We see that as the magnitude of the spatial chirp increases, the focus becomes more axially localised and the intensity falls away more steeply from $z = 0$. The direction of the spatial chirp does not matter, which makes sense because when we reverse the direction of the chirp, the transverse distance to the optical axis does not change for each frequency component. One might conclude from Figure 4.9 that a linear SSTF *does* lead to a more localised region of high intensity.

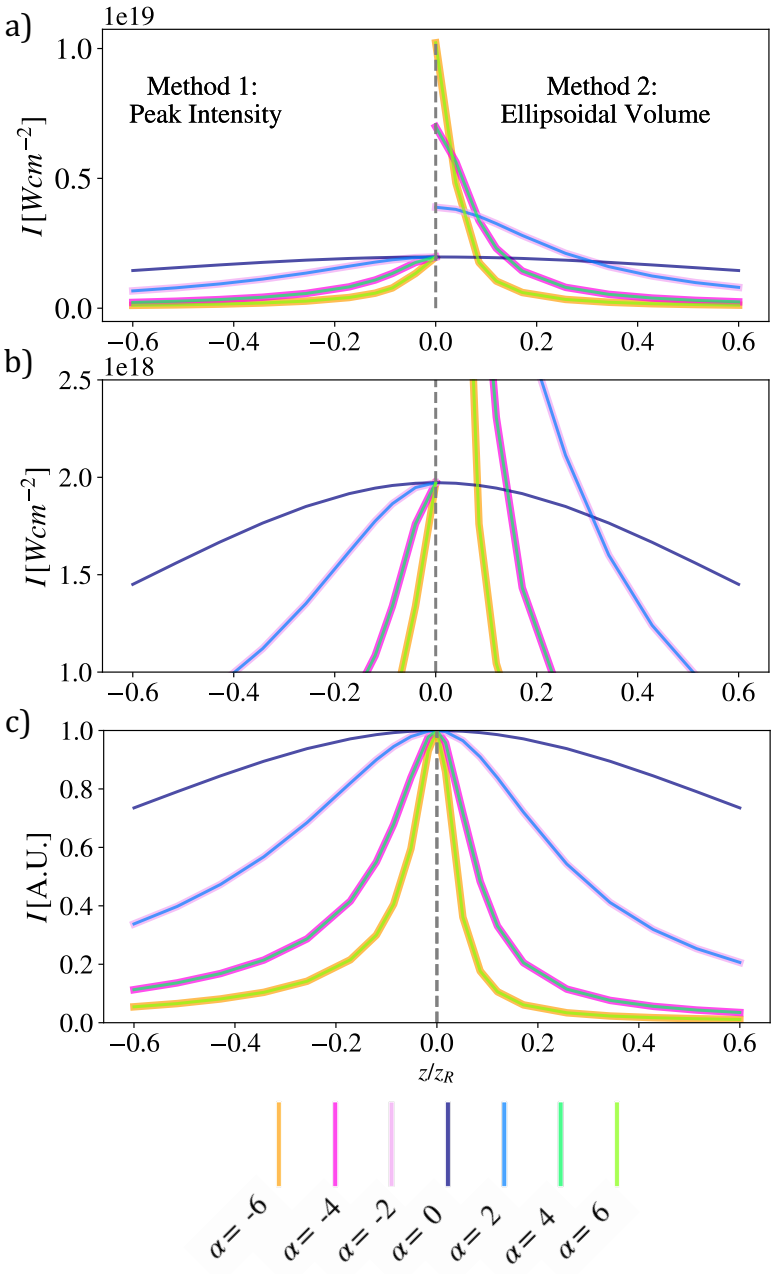


Figure 4.6: Longitudinal variation of the axial peak intensity of the linear SSTF pulse calculated using two different methods. Method 1: Extracting the peak intensity (PI) of the fields at all axial positions. Method 2: assuming an ellipsoidal focal volume (EFV). a) Shows the raw axial intensity, b) shows the same data as in a) but on a different scale to highlight that the peak axial intensity of the SSTF pulse is always the same and c) shows the normalised axial intensity. Note that for all values of α the axial intensity profiles are symmetric about $z = 0$, with only half plotted here to allow comparison. The positive and negative chirps perfectly overlap due to the fact that it does not matter which direction you chirp the pulse.

4.4.2 Calculation of Peak Axial Intensity

Method 1: Peak Intensity The first method to characterise the degree of axial localisation is to calculate the peak intensity at each axial position directly from the amplitude of the electric fields of the pulse. The intensity (also sometimes referred to as the irradiance or power density)

of a laser pulse in free-space is[217]

$$I \equiv c\langle u \rangle = \frac{c}{2}\Big(\varepsilon_0 |\mathbf{E}|^2 + \frac{1}{\mu_0}|\mathbf{B}|^2\Big). \tag{4.34}$$

We know that $E = cB$, such that

$$I = \varepsilon_0 c |\mathbf{E}|^2. \tag{4.35}$$

Therefore, given our calculation of the spatiotemporal electric fields $E(x, y, t) \approx U_{\text{out}}$, we can calculate the intensity of the field. For a fixed plane z, the total energy of the pulse is given by the energy density integrated for all space and time

$$\mathscr{E} = \int I(x, y, t)\, \mathrm{d}x\, \mathrm{d}y\, \mathrm{d}t. \tag{4.36}$$

Throughout the simulation, we assume an input pulse energy of $\mathscr{E} = 1\,\mathrm{J}$. At each axial step of the simulation, we check energy conservation. The total energy of the pulse is the same in the spectral and temporal domains as a consequence of Parseval's relation, and we find that at every axial step

$$\mathscr{E} = c\varepsilon_0 \int |\tilde{E}|^2 \mathrm{d}x\, \mathrm{d}y\, d\omega = c\varepsilon_0 \int |E|^2 \mathrm{d}x\, \mathrm{d}y\, \mathrm{d}t = 1\,\mathrm{J}. \tag{4.37}$$

Knowing that the total energy of the pulse is the same for the SSTF and non-SSTF comparisons, we can then calculate the absolute peak intensity $I_{\text{peak}} = \max\{I\}$ (given in W cm^{-2}), which is found by taking the maximum of the calculated intensity I. This method will henceforth be referred to as the peak intensity method, or PI for short. This quantity is most relevant in the context of ionisation injection because the ionisation of gases depends on the intensity of the laser field. As already discussed, this is plotted in Figure 4.6.

METHOD 2: ELLIPSOIDAL FOCAL VOLUME Durfee *et al.*[211] obtain an expression for the axial intensity profile by multiplying together several factors that contribute to the axial peak intensity. Their expression for the peak intensity "when the contributions to the axial dependence of the intensity are multiplied" is

$$I(\xi) = I_0 \sqrt{\frac{1}{(1+\xi^2)(1+\beta_{\text{BA}}^4 \xi^2)}}, \tag{4.38}$$

where I_0 is the peak intensity at the focus $z = 0$, $\xi = z/z_R$ is the normalised longitudinal position, and $\beta_{BA} = \sqrt{1+\beta^2}$ is their beam aspect ratio, where their β is our α_{norm}.

No derivation of Eq. 4.38 is given, but they appear to be using a common heuristic approach for calculating the peak intensity, where the total energy is divided by the z-dependent cross-sectional area and pulse duration.

Implicit in such a derivation is that we assume our focal spot has an untilted ellipsoidal focal volume and input energy of 1J. This can be used to calculate the axial peak intensity I_{peak}. The peak fluence F_{peak} of an elliptical Gaussian pulse is

$$F_{peak}(z) = \frac{2\mathscr{E}}{\pi w_x(z) w_y(z)}. \tag{4.39}$$

For a pulse with a Gaussian spectrum, the peak fluence F_{peak} is related to the peak intensity I_{peak} by

$$I_{peak} = 2\sqrt{\frac{\ln 2}{\pi}}\frac{F_{peak}}{\tau_p} \tag{4.40}$$

assuming a Gaussian temporal intensity profile with FWHM pulse duration τ_p, such that the peak intensity in terms of the input energy $\mathscr{E}$, pulse duration τ_p and pulse waists w_x and w_y is given by

$$I_{peak} = 4\sqrt{\frac{\ln 2}{\pi^3}}\frac{\mathscr{E}}{w_x(z) w_y(z) \tau_p}. \tag{4.41}$$

We shall call this method for calculating the peak axial intensity the the ellipsoidal focal volume (EFV) method. The EFV method requires knowledge of the key on-axis parameters that were described in Section 4.4.1, namely the pulse duration τ_p and spot sizes w_x and w_y. Extracting these parameters also allowed us to better understand the nature of SSTF, as we saw how the increase in on-axis bandwidth towards the focus gave rise to temporal compression. Here we shall see if these key on-axis parameters can also be used to deduce the intensity of the pulse through Eq. 4.41.

Figure 4.7 highlights the issues with using the EFV method for pulses with spatiotemporal couplings, in this case PFT. It also explains how a decrease in local pulse duration does not necessarily lead to an increase in peak intensity. In Figure 4.7a) there is only modest PFT, so we can see that the EFV approximation, shown by the dashed blue lines, works reasonably well. However, as seen in Figure 4.7b), near the focus the PFT is substantial, so the EFV method

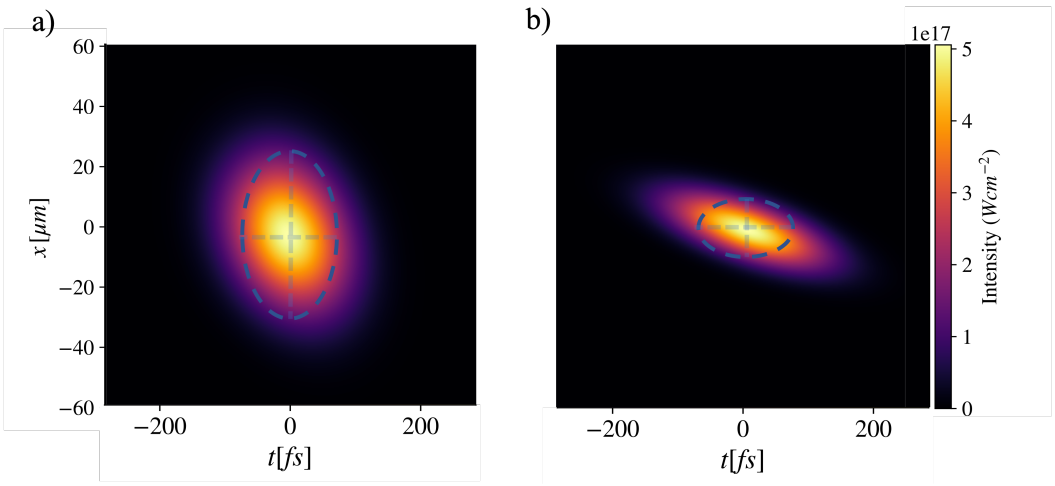


Figure 4.7: Schematic showing the issue with assuming an ellipsoidal focal volume (EFV). a) Pulse in the temporal domain away from the focus at $z = 2z_R$ and b) pulse in the temporal domain at focus $z = 0$.

fails. The spatiotemporal distribution of the pulse intensity is not well described by an *untilted* ellipsoid, especially close to the focus and leads to underestimates of the focal volume, and therefore overestimates of the peak intensity at the focus. In turn, this causes the EFV method to overestimate of the degree of axial localisation of the SSTF pulse. This can be seen in Figure 4.6, where in Figure 4.6a), we see that the peak axial intensity of the SSTF is vastly overestimated, and in Figure 4.6c) where the axial intensity appears to decrease at a faster rate away from the focus compared to the intensity calculated using the PI method. This may explain why Figure 4 in Durfee *et al.*[211] shows a higher degree of axial localisation when the pulse duration is taken into account. It remains unclear why Durfee *et al.* see a higher degree of axial localisation of the SSTF calculated using the full Fresnel integral, as contradicts the results of the present study and other previous literature. In any case, commonly-used, intuitive methods for calculating peak axial intensity fail for pulses with spatiotemporal couplings.

4.4.3 Comparison for Equivalent Spatial Focusing

By spatially chirping the input pulse at the plane of the lens, we increase the spot size of the input pulse along the direction of the chirp. This is shown in Figure 4.8. In the absence of any other effects, this will tend to decrease the spot size and Rayleigh range of the focused pulse, as can be seen in Figure 4.5c). Therefore, it is not immediately evident whether the greater degree of axial localisation seen in Figure 4.9 for $\alpha \neq 0$ is from the larger size of the input pulse (hence decreasing the effective f-number), or from the SSTF effect, or both. An explicit comparison of the degree of axial localisation of the peak intensity for an SSTF pulse and a *spatially-equivalent*

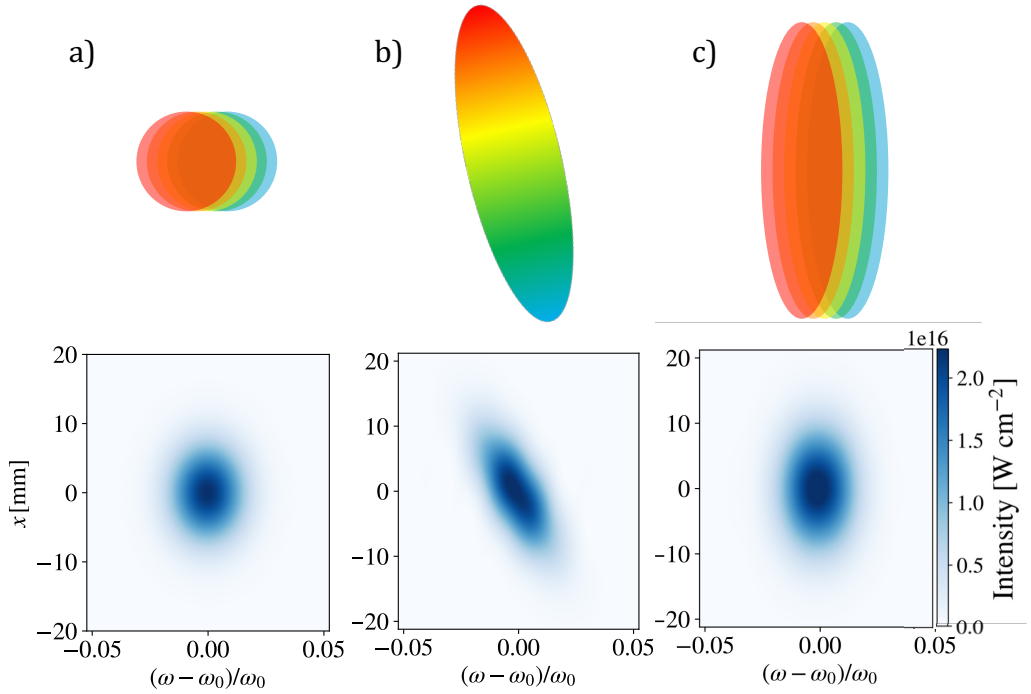


Figure 4.8: Spatiospectral profiles of the pre-chirp, SSTF and non-SSTF fair comparison pulses. a) Polychromatic beam prior to spatial chirp; b) SSTF pulse with frequency components displaced in x; c) Fair comparison case with same spatial amplitude as the SSTF but all frequency components are overlapped spatially.

pulse with no transverse chirp has not been presented in earlier work. In this section we provide this comparison. First, we define a 'fair comparison' case for the unchirped pulse by integrating the amplitude of the electric field (without phase) over the spectral domain to obtain a projection of the pulse at the input plane. This transverse profile is then multiplied by a chirp-free Gaussian spectrum, to give the input amplitude of the fair comparison case as

$$\tilde{U}_{\text{SSTF-free}}(x, y, z = z_{\text{lens}}, \omega) = \frac{\tilde{U}(\omega)}{M} \int_{-\infty}^{\infty} |\tilde{U}_{\text{SSTF}}(x, y, z, \omega)| \, \mathrm{d}\omega, \tag{4.42}$$

where $\tilde{U}_{\text{SSTF}}$ is the amplitude of the SSTF pulse at the input plane of the lens and M is a normalisation coefficient that preserves the energy of the pulse. The fair comparison case is shown in Figure 4.8c). Experimentally, this is a comparison between chirping the beam transversely and using a telescope which increases the size of the beam in one direction to have the same spatial extent as the SSTF beam, but without chirp.

With this definition of the fair comparison case, we find that *the degree of axial localisation is exactly the same for the SSTF and non-SSTF case*. This is shown on the left-hand side of Figure 4.9c) as a perfect agreement between the longitudinal variation of the peak axial intensity of the SSTF

(solid lines) and non-SSTF pulses (dashed lines and triangles). In other words, the greater degree of axial localisation of the high-intensity region apparently observed for a SSTF pulse arises solely from the increased size of the input pulse. This is seemingly paradoxical. If the pulse duration increases away from the focus, why doesn't the intensity decrease at a faster rate than for the non-SSTF pulse? In the previous section we saw that this is because our common understanding of intensity, based on the spot size and duration of the pulse, fails when spatiotemporal couplings are present. In other words, the degree of pulse-front tilt, which increases towards the focus, seems to prevent the intensity from decreasing at a faster rate.

The normalised intensity roll-off looks the same for the SSTF and non-SSTF comparison, suggesting that SSTF offers no improvement in axial intensity localisation. However, if we look at the absolute intensity, we see that the peak intensity of the SSTF field is always equal to the peak intensity of the unchirped case. It is a well-known fact that PFT decreases the peak intensity and, surprisingly, this is often not acknowledged in the context of SSTF, particularly for papers which study intensity-dependent processes.[210]

The fact that SSTF does not localise the high intensity region any more than a fair comparison non-SSTF case agrees with the original literature on the topic, which states that the degree of axial localisation depends only on the input spatial profile.[204,205] Durst *et al.* have derived an expression for the two-photon excitation fluorescence (TPEF) signal produced by an SSTF pulse[205]

$$I_{\mathrm{TPEF}}(z) = \left[1 + \frac{f_0 - (z - z_0)}{z_{R,\mathrm{mod}}}\right]^{-1/2}, \tag{4.43}$$

where $I_{\mathrm{TPEF}} = \int\int |E_t|^4 \,\mathrm{d}x\,\mathrm{d}y\,\mathrm{d}t$ is the two-photon excitation signal, related to the electric field strength at every axial position. Eq. 4.43 might suggest that the degree of axial localisation is independent of the spatial chirp α. However, on closer inspection, Durst *et al.* define a new Rayleigh range for the SSTF pulse after the lens[205]

$$z_{R,\mathrm{mod}} = \frac{2f_0^2/k_0}{w_0^2 + \alpha^2(\omega - \omega_0)^2} = \frac{2f_0^2/k_0}{w_{x,\mathrm{eff}}^2}, \tag{4.44}$$

where $w_{x,\mathrm{eff}}^2 = w_0^2 + \alpha^2(\omega - \omega_0)^2$ is the effective spot size of the input pulse along the spatially-chirped axis. From this, we see that SSTF only leads to a greater degree of axial localisation through the change in input spot size that arises when we chirp the pulse. There is no chirp in the y dimension, so in that plane $z_{R,\mathrm{mod}}$ will be the same as in the unchirped case. These

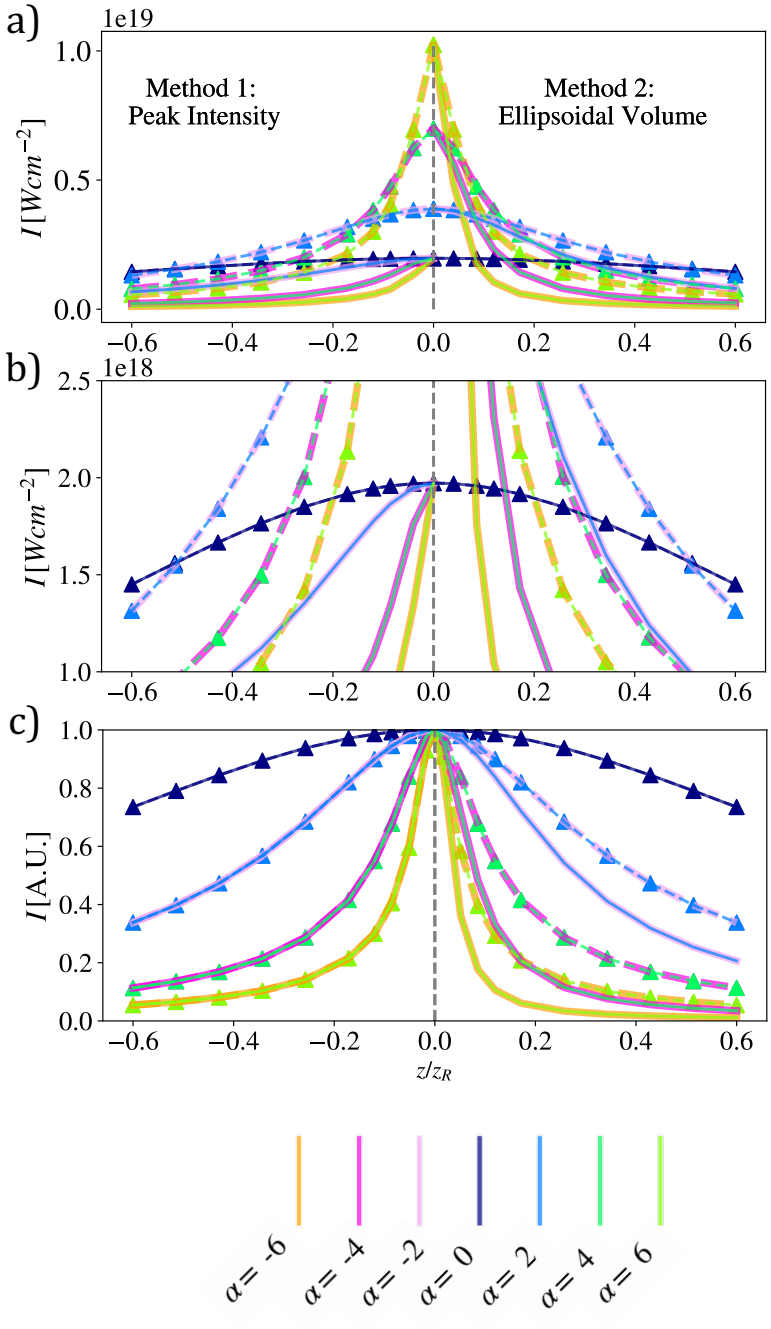


Figure 4.9: Longitudinal variation of the peak intensity of the linear SSTF pulse calculated using the Collins method as a function of z in the focal region showing: a) Absolute peak intensity I_0 calculated using two different methods. Method 1: Extracting the peak intensity (PI) of the fields at all axial positions. Method 2: assuming an ellipsoidal focal volume (EFV). Note that all intensity profiles are symmetric about $z = 0$, with only half plotted here to allow comparison. b) Same data as a) but plotted on a different scale to show that the peak intensity of the SSTF pulse is always limited to this value. c) Comparison of the longitudinal variation of the peak axial intensity, normalized to the value at focus. Solid lines denote the SSTF pulse ($\alpha \neq 0$) and dashed lines / triangles are the corresponding non-SSTF fair comparison ($\alpha = 0$) defined in Section 4.4.3.

definitions agree with the results shown in Figure 4.5 and confirm our findings that SSTF does not result in stronger localisation of the high-intensity region, relative to the fair-comparison case.

One final point to note is that for the fair-comparison case (dashed lines and triangles), the calculated axial intensity is exactly the same for both the peak intensity and ellipsoidal volume methods; this is shown by the symmetry about $z = 0$ of the dashed lines and triangles plotted in Figures 4.9a) and b). This is because the fair-comparison pulse does not have PFT, such that in this case the untilted ellipsoid approximation works well.

4.5 Conclusions and Future Work

To conclude, we investigated the degree of axial localisation of the peak intensity of a linearly-chirped SSTF pulse and found that the improved localisation of the intensity could be entirely explained by the increase in the transverse size of the input pulse caused by the spatial chirp. This was established from calculations of the spatiotemporal fields of the SSTF, using both an analytic expression for the SSTF in the focal region and numerical methods based on the Collins diffraction integral. We found that simplistic methods used to characterise focal spots fail when spatiotemporal couplings are present and that employing these can lead to erroneous conclusions about the effectiveness of SSTF in localising the focal intensity. We presented a method that can be used to benchmark the degree of axial localisation of such pulses by creating a fair-comparison pulse that undergoes the same spatial focusing. We also found that the peak intensity of the linear SSTF beam was always equal to the peak intensity of the unchirped polychromatic beam, where $\alpha = 0$.

In future work, we will use principle component analysis to fit tilted ellipsoids to focal spots with first-order spatiotemporal couplings, to investigate whether tilted ellipsoids can be used as a more accurate model of focal spot volumes when first-order spatiotemporal couplings are present. We also aim to develop phase-space arguments in terms of the global bandwidth to understand why the peak intensity of the SSTF beam appears to be fundamentally limited to the peak intensity of the pre-chirped spot.

5 | The Radial Simultaneous Space-Time Focus

5.1 Introduction

In this chapter, we study the properties of the radial simultaneous space-time focus (R-SSTF): an alternative to the more widely used 'linear SSTF', as was discussed in the previous chapter. Much less common in existing literature is the case where an SSTF is generated using an input pulse which is chirped radially, instead of along one axis.[206] One method for generating radially chirped laser pulses is to use concentric-circle gratings. A 'positive' radial chirp indicates that the wavelength decreases with distance from the optical axis, as shown in Figure 5.1. As the pulse comes to a focus, much like for a linear SSTF (see Chapter 4), the different frequency components begin to overlap, leading to an increase in local bandwidth. This decreases the pulse duration. We saw in the previous chapter that counter-intuitively this does *not* lead to localisation of the region of highest intensity about the focus. We showed that heuristic models used to estimate the intensity of laser foci fail when spatiotemporal couplings are present, and that residual spatiotemporal couplings limit the peak intensity at the focus. Here we study the properties of an SSTF generated using a spatially chirped input pulse which is axially symmetric, and we assess whether any advantages are offered in comparison to a conventional linear SSTF.

We first attempt to derive the output fields of an R-SSTF, in a similar manner used in the previous chapter. We find that we are not able to integrate the result for a closed-form analytic solution. Instead, we propagate the R-SSTF by describing the beam as a decomposition of Laguerre-Gaussian modes. Each mode propagates independently within the modified paraxial approximation and all modes are summed at the output plane to give the resulting R-SSTF beam. By calculating the intensity profile of the pulse using both the mode-decomposition and

the Collins technique, we show that the degree of axial localisation of an R-SSTF beam is worse than for the fair comparison annular beam (defined using the method from the previous chapter, see Section 4.4.3). However, an R-SSTF may have several advantages for the laser machining of solids, such as the absence of pulse-front tilt (PFT) at the focus and a controllable group velocity.

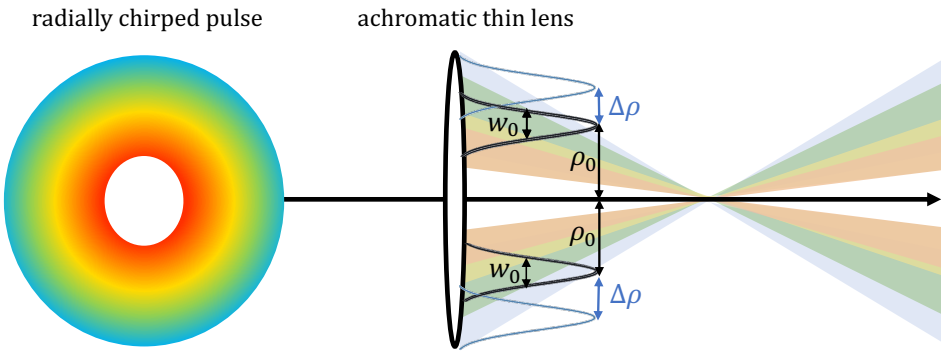


Figure 5.1: Schematic of a radial simultaneous space-time focus. A pulse which is positively chirped transversely along the radius passes through an achromatic lens. The nominal Gaussian annulus with a radius of ρ_0 is shown for the central frequency ω_0 in black and for an arbitrary beamlet with a different frequency ω, which has a radius $\rho_0 + \Delta\rho$ in blue. Each annulus has a Gaussian width of w_0.

5.2 Theoretical Analysis of the R-SSTF

The radially-chirped input pulse can be described by a summation of monochromatic Gaussian annuli of $1/e^2$ widths w_0, which are offset radially about the nominal annulus size ρ_0 at the central frequency ω_0. The radius of each annulus $\Delta\rho(\omega) = \rho_0 + \alpha(\omega - \omega_0)/\alpha_{\text{norm}}$ is proportional to the relative frequency $\omega - \omega_0$ of that component. Assuming that the different frequency components can be represented by annuli that have different radii but the same widths presents a simple case which ensures that frequency components are spatially separated at the input plane.

KEY PARAMETER α: describes the spatial chirp of the incoming pulse.

The chirp scaling factor α is used to control the direction and strength of the spatial chirp and is scaled by the ratio $\alpha_{\text{norm}} = \Delta\omega/w_0$, where $\Delta\omega$ is the FWHM input bandwidth. Mathematically, the input field may then be written as

$$\tilde{U}_{\text{in}}(\rho, z = 0, \omega) = U_\omega \tilde{U}(\omega) \exp\left\{\frac{-[\rho - \Delta\rho(\omega)]^2}{w_0^2}\right\}, \tag{5.1}$$

where

$$\Delta\rho(\omega) = \rho_0 + \frac{\alpha}{\alpha_{\text{norm}}}(\omega - \omega_0), \tag{5.2}$$

and $U_\omega = |\mathbf{U}_{\text{in}}|/\omega_0$ is a normalisation constant related to the peak electric field strength and central frequency of the laser pulse. The Gaussian input spectrum is $\tilde{U}(\omega) = 1/\sqrt{2\Gamma}\exp[(\omega-\omega_0)^2/4\Gamma]$, where $\Gamma = a + ib$ is a complex parameter that characterises the temporal width and chirp of the pulse, with $a = 2\ln 2/\tau_{\text{p}}^2$ and $b = \chi a$ (see Section 2.1.2 for more details). The pulse duration is τ_{p} and χ is a normalised parameter that defines the input chirp. For simplicity, we assume the nominal case that the pulse is not chirped, such that $b = 0$. We could assume otherwise without affecting any subsequent derivation, but then the focal plane of the pulse may not coincide with the geometric focus,[204] making our fair comparison more complicated.

The input fields in the spatiospectral and spatiotemporal domains are shown in Figure 5.2, where the radial chirp can be seen as a correlation between frequency and radial position in Figure 5.2a). In the temporal domain, the radially-chirped pulse, shown in Figure 5.2b), has a longer FWHM pulse duration of 83.18 fs, compared to the unchirped case, shown in Figure 5.2b′), where the fully-compressed FWHM pulse duration is 45 fs. Comparing Figures 5.2c) and c′), we see that the spatial profile of the chirped pulse is the same as with the unchirped fair comparison case. Note that since the fair comparison which was defined in the previous chapter (see Section 4.4.3) is only concerned with measuring the spatial projection of the beam at the input plane and multiplying this with an unchirped equivalent spectrum without spatiotemporal couplings, it is also valid for the R-SSTF too. This method could therefore provide a potentially powerful tool for identifying the contribution of spatiotemporal effects in focusing beams. Figures 5.2d) and d′) show a lineout of the intensity at $y = 0$. The fair comparison was discussed in Section 4.4.3 and the same method is used here for the R-SSTF.

5.2.1 Derivation of the Radial Simultaneous Space-Time Focus

An explicit derivation of a radial SSTF does not currently exist in the literature, so we shall see if one exists here. The Fresnel diffraction integral in Cartesian coordinates is given by[201]

$$\tilde{U}_{\text{out}}(x,y) = \frac{\mathrm{e}^{ikz}}{\sqrt{i\lambda z}} \iint_{-\infty}^{\infty} \tilde{U}_{\text{in}}(x_{\text{in}}, y_{\text{in}}) \exp\left\{\frac{ik}{2z}\Big[(x-x_{\text{in}})^2 + (y-y_{\text{in}})^2\Big]\right\} \mathrm{d}x_{\text{in}}\mathrm{d}y_{\text{in}}. \tag{5.3}$$

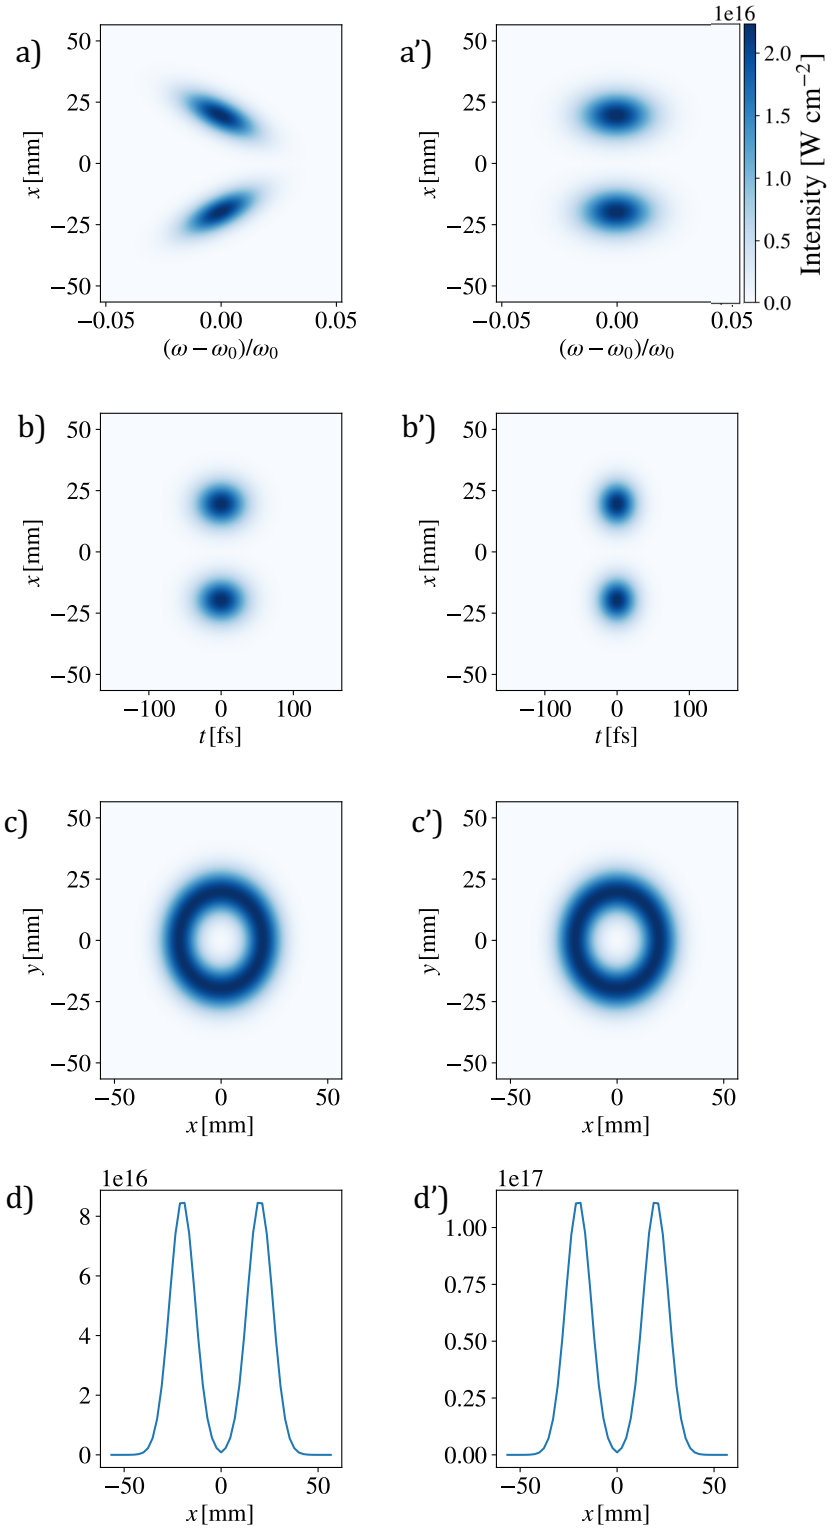


Figure 5.2: Comparison of the intensity distributions of the input pulses for radially-chirped pulses with $\alpha = -1$ and $\alpha = 0$ (i.e. an unchirped pulse). Figures a)–d) show the chirped case and Figures a′)–d′) show the unchirped fair-comparison (see Section 4.4.3). a) shows the spatiospectral domain where the correlation between spatial and frequency offset can be seen, compared to a′) where there is no correlation, b) shows the spatiotemporal domain, where it can be seen that the FWHM pulse duration is longer for all transverse positions in the chirped case at 83.18 fs compared to the unchirped case b′) where the fully-compressed FWHM pulse duration is 45 fs. Comparing c) and d) with c′) and d′), we see that the spatial profile is exactly the same, so our spatial focusing will be equivalent in both cases.

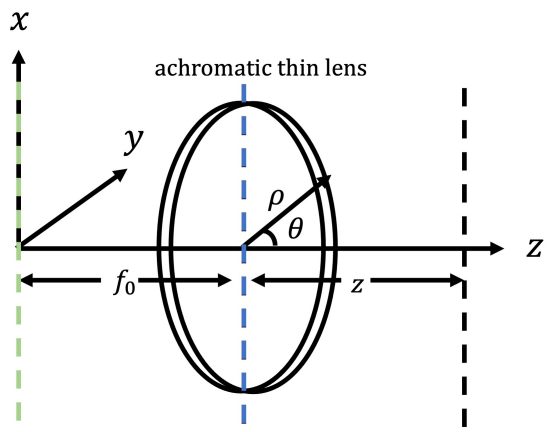


Figure 5.3: Diagram of the 3D cylindrical coordinate system. A thin lens is positioned a distance f_0 away from the origin. The input pulse travels a distance f_0 from the input plane (green dashed line) before propagating through the lens. The output pulse is then propagated after the lens by a distance z which is in a narrow range about the focus f_0. A simplification can be made where the pulse is defined at the input plane of the lens (dashed blue line).

In contrast to the linear SSTF, here frequency components are chirped radially in two-dimensions, as described by Eq. 5.1. Therefore, we rewrite the Fresnel integral in cylindrical coordinates. Under a transformation of variables, we obtain

$$\tilde{U}_{\text{out}}(\rho,\theta) = \frac{e^{ikz}}{\sqrt{i\lambda z}} \int_0^{2\pi}\int_0^{\infty} \tilde{U}_{\text{in}}(\rho_{\text{in}},\theta_{\text{in}}) \exp\left\{\frac{ik}{2z}\left[\left(\rho\cos\theta-\rho_{\text{in}}\cos\theta_{\text{in}}\right)^2+\left(\rho\sin\theta-\rho_{\text{in}}\sin\theta_{\text{in}}\right)^2\right]\right\}\rho_{\text{in}}\mathrm{d}\rho_{\text{in}}\,\mathrm{d}\theta_{\text{in}}, \quad (5.4)$$

$$= \frac{e^{ikz}}{\sqrt{i\lambda z}} \int_0^{2\pi}\int_0^{\infty} \tilde{U}_{\text{in}}(\rho_{\text{in}},\theta_{\text{in}}) \exp\left\{\frac{ik}{2z}\left[\rho^2+\rho_{\text{in}}^2-2\rho\rho_{\text{in}}\left(\cos\theta\cos\theta_{\text{in}}+\sin\theta\sin\theta_{\text{in}}\right)\right]\right\}\rho_{\text{in}}\mathrm{d}\rho_{\text{in}}\,\mathrm{d}\theta_{\text{in}}, \quad (5.5)$$

$$= \frac{e^{ikz}}{\sqrt{i\lambda z}} \int_0^{2\pi}\int_0^{\infty} \tilde{U}_{\text{in}}(\rho_{\text{in}},\theta_{\text{in}}) \exp\left\{\frac{ik}{2z}\left[\rho^2+\rho_{\text{in}}^2-2\rho\rho_{\text{in}}\cos\left(\theta-\theta_{\text{in}}\right)\right]\right\}\rho_{\text{in}}\mathrm{d}\rho_{\text{in}}\,\mathrm{d}\theta_{\text{in}}, \quad (5.6)$$

which may be written as

$$\tilde{U}_{\text{out}}(\rho,\theta) = \frac{e^{ikz}e^{\frac{ik\rho^2}{2z}}}{\sqrt{i\lambda z}} \int_0^{\infty} \tilde{U}_{\text{in}}(\rho_{\text{in}},\theta_{\text{in}}) \exp\left\{\frac{ik\rho_{\text{in}}^2}{2z}\right\}\rho_{\text{in}}\mathrm{d}\rho_{\text{in}} \int_0^{2\pi} \exp\left\{\frac{ik\rho\rho_{\text{in}}}{z}\cos\left(\theta-\theta_{\text{in}}\right)\right\}\mathrm{d}\theta_{\text{in}}, \quad (5.7)$$

we can use the known result that[218]

$$\int_0^{2\pi} \exp\left\{\frac{ik\rho\rho_{\text{in}}}{z}\cos(\theta-\theta_{\text{in}})\right\}\mathrm{d}\theta_{\text{in}} = 2\pi J_0\left(\frac{k\rho\rho_{\text{in}}}{z}\right), \tag{5.8}$$

where J_0 is the zeroth order Bessel function of the first kind. Therefore, the output field after a propagation distance z is

$$\tilde{U}_{\text{out}}(\rho,\theta) = \frac{2\pi \mathrm{e}^{ikz}\mathrm{e}^{\frac{ik\rho^2}{2z}}}{\sqrt{i\lambda z}}\int_0^{\infty}\tilde{U}_{\text{in}}(\rho_{\text{in}},\theta_{\text{in}})\exp\left\{\frac{ik\rho_{\text{in}}^2}{2z}\right\}J_0\left(\frac{k\rho\rho_{\text{in}}}{z}\right)\rho_{\text{in}}\mathrm{d}\rho_{\text{in}}, \tag{5.9}$$

which is proportional to a Hankel transform. Writing our radially-chirped input field, we obtain

$$\begin{aligned}\tilde{U}_{\text{out}}(\rho,\theta) = U_\omega\tilde{U}(\omega)\frac{2\pi \mathrm{e}^{ikz}\mathrm{e}^{\frac{ik\rho^2}{2z}}}{\sqrt{i\lambda z}}\int_0^{\infty}\exp\left\{\frac{-[\rho_{\text{in}}-\Delta\rho(\omega)]^2}{w_0^2}\right\}& \\ \exp\left\{\frac{ik\rho_{\text{in}}^2}{2z}\right\}J_0\left(\frac{k\rho\rho_{\text{in}}}{z}\right)\rho_{\text{in}}\mathrm{d}\rho_{\text{in}},& \quad (5.10)\\ = U_\omega\tilde{U}(\omega)\frac{2\pi \mathrm{e}^{ikz}\mathrm{e}^{\frac{ik\rho^2}{2z}}}{\sqrt{i\lambda z}}\int_0^{\infty}\exp\left\{-\left(\frac{1}{w_0^2}-\frac{ik}{2z}\right)\rho_{\text{in}}^2+\frac{2\Delta\rho}{w_0^2}\rho_{\text{in}}-\frac{\Delta\rho^2}{w_0^2}\right\}& \\ \rho_{\text{in}}J_0\left(\frac{k\rho\rho_{\text{in}}}{z}\right)\mathrm{d}\rho_{\text{in}},& \quad (5.11)\end{aligned}$$

which we are not able to integrate for a closed-form analytic solution. It should be noted that even for the simplified geometry of a lens positioned at the origin (dashed blue line in Figure 5.3), there still does not appear to be a closed-form analytic solution.

5.2.2 Propagating Arbitrary Beams as a Decomposition of Laguerre-Gaussian Modes

Since it seems like an analytic solution does not exist, we use an alternative approach to propagate radial SSTF beams analytically. One method, which can also be used to propagate orbital angular momentum (OAM) beams, is to decompose the input pulse as a sum of Laguerre-Gaussian modes. Each of these Laguerre-Gaussian modes may then be propagated within the paraxial approximation. The complex spatiospectral field of a Laguerre-Gaussian mode is (see Section 2.1.5)

$$U_p^\ell(\rho,\theta,z) = \sqrt{\frac{2p!}{\pi(|\ell|+p)!}}\frac{\mathrm{e}^{ikz}\mathrm{e}^{\frac{ik\rho^2}{2q(z)}}}{w(z)}\left[\frac{\sqrt{2}\rho}{w(z)}\right]^{|\ell|}\mathrm{e}^{-i(2p+|\ell|+1)\varphi_{\text{G}}(z)}L_p^{|\ell|}\left[\frac{2\rho^2}{w(z)^2}\right]\mathrm{e}^{i\ell\theta}, \tag{5.12}$$

where $L_p^{|\ell|}$ is the generalised Laguerre polynomial at a given radial index $p \geq 1$ and azimuthal mode index ℓ and

$$\rho = \sqrt{x^2 + y^2}, \qquad \theta = \tan^{-1}\left(\frac{y}{x}\right), \tag{5.13}$$

$$\varphi_{\mathrm{G}}(z) = \tan^{-1}\left(\frac{z - z_0}{z_{\mathrm{R}}}\right), \tag{5.14}$$

$$w(z) = w_0\sqrt{1 + \left(\frac{z - z_0}{z_{\mathrm{R}}}\right)}, \tag{5.15}$$

$$q(z) = z - z_0 - iz_{\mathrm{R}}. \tag{5.16}$$

The Rayleigh range z_{R} is a parameter that can be tuned to optimise the mode decomposition.

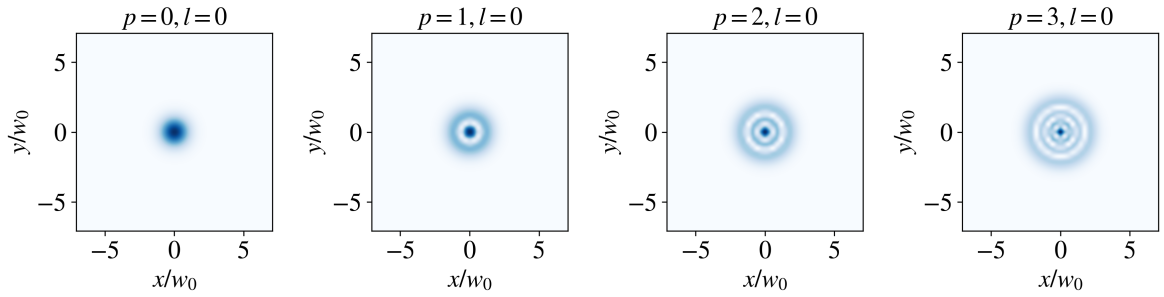


Figure 5.4: Transverse intensity profiles for the Laguerre-Gaussian modes with increasing radial index p

Defining the Rayleigh range as $z_{\mathrm{R}} = k_0 w_0^2/2$, dictated by the width of the Gaussian annulus w_0 and the central frequency $k_0 = \omega_0/c$, is found to work well (with only 15 modes required in this case), so this has not been investigated further here. In our case, we want a flat phase front with no OAM, so we will consider a sum over increasing radial mode number p, where the azimuthal mode number ℓ is always zero. The overlap integral of each mode U_p^0 with the desired incident field $\tilde{U}_{\mathrm{in}}$ (given by Eq. 5.1 in the main text) gives the complex mode coefficient as

$$c_p(\omega) = 2\pi \int_{-\infty}^{\infty} \tilde{U}_{\mathrm{in}}(\rho, z = 0, \omega) * U_p^0(\rho, z = 0)\,\mathrm{d}\rho. \tag{5.17}$$

Note that the θ dependence is dropped because $\tilde{U}_{\mathrm{in}}$ and U_p^0 are both azimuthally symmetric. The reconstructed field is then

$$\tilde{U}_{\mathrm{in},N_p}(\rho, z, \omega) = \sum_{0}^{N_p} c_p(\omega) * U_p^0(\rho, z). \tag{5.18}$$

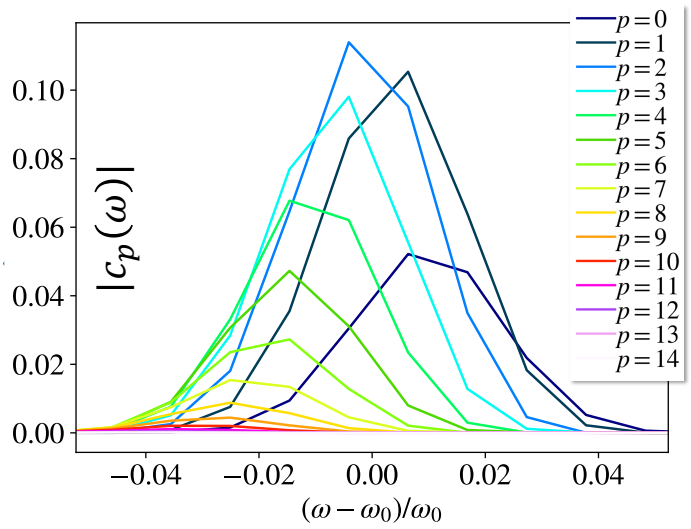


Figure 5.5: Amplitude $|c_p(\omega)|$ of the different modes with radial index p for the input R-SSTF shown in Figures 5.2a)–d), where $\alpha = -1$.

To calculate the optimum number of modes to include, the least squares error (LSE) between the amplitude of the incident pulse and the reconstructed field is found. The LSE for any given number of included modes N_p is given by

$$\mathrm{LSE}(N_p) = \sum_{m}^{M} \left| \tilde{U}_{\mathrm{in}}(m) - \tilde{U}_{\mathrm{in},N_p}(m) \right|^2, \tag{5.19}$$

where each grid point is denoted by m and M is the total number of grid points in the simulation (e.g. $M_x * M_y * M_\omega$, note that the simulation is performed on a Cartesian grid). This error is calculated up to a certain maximum number of modes. Usually, the error $\mathrm{LSE}(N_p)$ is a convex function of the number of included modes N_p, such that a minimum can be found within this range. The value $N_{p,\mathrm{min}}$ where this minimum error occurs is then the number of modes to be included in the simulation, here $N_{p,\mathrm{min}} = 15$. Figure 5.5 shows the amplitude of these 15 modes.

After finding the relative contribution of each mode, these modes need to be independently propagated. We will now derive the output field of each mode after the lens. For simplicity, we will assume that the focusing geometry is as shown in Figure 5.3, with the input plane defined by the blue dashed line, and assume that the beam waist of the input pulse coincides with the lens (a collimated input) at $z_0 = z_1 = 0$. At the lens, the field is modified by a factor

$$\times \mathrm{e}^{-\frac{ik\rho^2}{2f_0}}, \tag{5.20}$$

such that the field of the mode at the output plane of the lens is given by

$$U_p^{\ell}(\rho,\theta,z)=\sqrt{\frac{2p!}{\pi(|\ell|+p)!}}\frac{e^{\frac{ik\rho^2}{2}\left(\frac{1}{q_{\text{in}}(0)}-\frac{1}{f_0}\right)}}{w(0)}\left[\frac{\sqrt{2}\rho}{w(0)}\right]^{|\ell|}e^{-i(2p+|\ell|+1)\varphi_{\text{G,in}}(0)}L_p^{|\ell|}\left[\frac{2\rho^2}{w(0)^2}\right]e^{i\ell\theta}. \quad (5.21)$$

The modification to the complex 'q'-parameter at the lens is given by

$$q_{\text{out}}=\frac{(Aq_{\text{in}}+B)}{(Cq_{\text{in}}+D)}, \quad (5.22)$$

where the ray matrix for the lens is

$$\begin{pmatrix}A & B\\ C & D\end{pmatrix}=\begin{pmatrix}1 & 0\\ -1/f_0 & 1\end{pmatrix}. \quad (5.23)$$

Calculating q_{out} directly from the ABCD law, we get

$$q_{\text{out}}=\frac{q_{\text{in}}}{1-q_{\text{in}}/f_0}, \quad (5.24)$$

such that

$$\frac{1}{q_{\text{out}}}=\frac{1}{q_{\text{in}}}-\frac{1}{f_0}. \quad (5.25)$$

We can also note that

$$q_{\text{out}}=\left(\frac{1}{q_{\text{in}}}-\frac{1}{f_0}\right)^{-1}=z-z_{0,\text{out}}-iz_{R,\text{out}}. \quad (5.26)$$

It is clear to see that the real and imaginary parts calculated from the thin lens equation relate to the new Rayleigh range and spot size of the beam. We may rearrange the expression above, using $q_{\text{in}}=-iz_{\text{R}}$, in order to split q_{out} into real and imaginary terms as follows

$$q_{\text{out}}=\frac{1}{\frac{-i}{z_{\text{R}}}-\frac{1}{f_0}}, \quad (5.27)$$

$$=-\frac{f_0 z_{\text{R}}}{z_{\text{R}}^2+f_0^2}\Big(z_{\text{R}}+if_0\Big). \quad (5.28)$$

The imaginary part gives the modified Rayleigh range

$$z_{R,\text{out}}=\frac{f_0^2 z_{\text{R}}}{z_{\text{R}}^2+f_0^2}\approx\frac{f_0^2}{z_{\text{R}}}, \quad (5.29)$$

and the real part gives the beam waist position

$$z_{0,\text{out}} = \frac{f_0 z_\text{R}^2}{z_\text{R}^2 + f_0^2} \approx f_0. \tag{5.30}$$

The complex 'q'-parameter may also written in terms of the radius of curvature $R(z)$ and the beam waist $w(z)$ as

$$\frac{1}{q_\text{in}(z)} = \frac{1}{R(z)} + i\frac{\lambda}{\pi w^2(z)}, \tag{5.31}$$

such that from Eq. 5.25

$$\frac{1}{q_\text{out}(0)} = \frac{1}{R(0)} + i\frac{\lambda}{\pi w_0^2} - \frac{1}{f_0}, \tag{5.32}$$

$$\frac{1}{q_\text{out}(0)} = \left(\frac{1}{R(0)} - \frac{1}{f_0}\right) + i\frac{\lambda}{\pi w^2(0)}. \tag{5.33}$$

In Eqs. 5.31 and 5.33, the imaginary parts are identical such that the pulse waists must be the same on either side of the lens $w_\text{in}(0) = w_\text{out}(0)$. Using this change of variables in our original expression for the mode after the lens (Eq. 5.21) we obtain

$$U_p^\ell(\rho,\theta,z) = \sqrt{\frac{2p!}{\pi(|\ell|+p)!}}\frac{\mathrm{e}^{\frac{ik\rho^2}{2q_\text{out}(0)}}}{w_\text{out}(0)}\left[\frac{\sqrt{2}\rho}{w_\text{out}(0)}\right]^{|\ell|}\mathrm{e}^{-i(2p+|\ell|+1)\varphi_{\text{G,in}}(0)}L_p^{|\ell|}\left[\frac{2\rho^2}{w_\text{out}^2(0)}\right]\mathrm{e}^{i\ell\theta}. \tag{5.34}$$

This leaves us only with the Guoy phase $\varphi_{\text{G,in}}(0)$ to deal with. We want to find an expression for the Laguerre-Gaussian pulse after the lens depending only on the modified 'q'-parameter. Eq. 5.34 can equivalently be written

$$U_p^\ell(\rho,\theta,z) = \sqrt{\frac{2p!}{\pi(|\ell|+p)!}}\frac{\mathrm{e}^{\frac{ik\rho^2}{2q_\text{out}(0)}}}{w_\text{out}(0)}\left[\frac{\sqrt{2}\rho}{w_\text{out}(0)}\right]^{|\ell|}\mathrm{e}^{-i(2p+|\ell|+1)\varphi_{\text{G,out}}(0)}L_p^{|\ell|}\left[\frac{2\rho^2}{w_\text{out}^2(0)}\right]\mathrm{e}^{i\ell\theta} \times \mathrm{e}^{i(2p+|\ell|+1)\Delta\varphi_\text{G}}, \tag{5.35}$$

such that the field takes the same form before and after the lens, with a modified 'q'-parameter and an *additional mode-dependent phase shift* $\Delta\varphi_\text{G} = \varphi_{G,\text{out}}(0) - \varphi_{G,\text{in}}(0)$. This appears in the term on the second line of Eq. 5.35. This phase shift is just a constant phase (i.e. it is not spatially or frequency dependent) and ordinarily could be ignored as a global phase, however, the phase shift varies for different mode numbers and will manifest as a relative phase between different modes, so we need to account for this. Although we are still looking for a physical interpretation

of this phase shift, it is required in the simulations. Neglecting this phase shift leads to fields which are unphysical. We can derive this phase shift using our definition of the 'q'-parameter before and after the lens

$$\Delta\varphi_{\mathrm{G}} = \varphi_{G,\mathrm{out}}(0) - \varphi_{G,\mathrm{in}}(0), \tag{5.36}$$

$$= \tan^{-1}\left(-\frac{z_{0,\mathrm{out}}}{z_{R,\mathrm{out}}}\right) - \tan^{-1}\left(0\right), \tag{5.37}$$

$$= \tan^{-1}\left(\frac{z_{\mathrm{R}}}{f_0}\right). \tag{5.38}$$

Therefore, the field of each mode at an arbitrary axial distance z after the lens is given by

$$U_p^{\ell}(\rho,\theta,z) = \sqrt{\frac{2p!}{\pi(|\ell|+p)!}}\frac{\mathrm{e}^{\frac{ik\rho^2}{2q_{\mathrm{out}}(z)}}}{w_{\mathrm{out}}(z)}\left[\frac{\sqrt{2}\rho}{w_{\mathrm{out}}(0)}\right]^{|\ell|}\mathrm{e}^{-i(2p+|\ell|+1)\varphi_{\mathrm{G,out}}(z)}L_p^{|\ell|}\left[\frac{2\rho^2}{w_{\mathrm{out}}(z)^2}\right]\mathrm{e}^{i\ell\theta} \times \mathrm{e}^{i(2p+|\ell|+1)\Delta\varphi_{\mathrm{G}}}. \tag{5.39}$$

5.3 Ray Matrix to Generate an R-SSTF

In the previous section, we found an analytic method to propagate radial simultaneous space-time foci. This method allows for a calculation of the R-SSTF field at any z position, by propagating the modes and taking the superposition of these propagated modes at the output plane. An alternative approach is to use diffraction integrals, such as the Collins integral, to propagate the fields. With the Collins method, we only need to specify the input field and the ray matrices that describe the optical system that this field propagates through. The benefit of this method is that one is not restricted to assuming any particular form of the input field, and we can also simulate more complex systems with multiple optics. As shown in Figure 5.3, we propagate the spatially chirped input pulse a distance f_0 from the the input plane (shown by the dashed green line) before passing it through a non-dispersive thin lens. We then propagate it to a focus in free-space. This is described by the following matrices

$$O = \begin{pmatrix} A & B \\ C & D \end{pmatrix} = \begin{pmatrix} 1 & z-z_0 \\ 0 & 1 \end{pmatrix}\begin{pmatrix} 1 & 0 \\ -1/f_0 & 1 \end{pmatrix} = \begin{pmatrix} 1-(z-z_0)/f_0 & z-z_0 \\ -1/f_0 & 1 \end{pmatrix}. \tag{5.40}$$

Full details of the Collins method are given in Section 2.3.

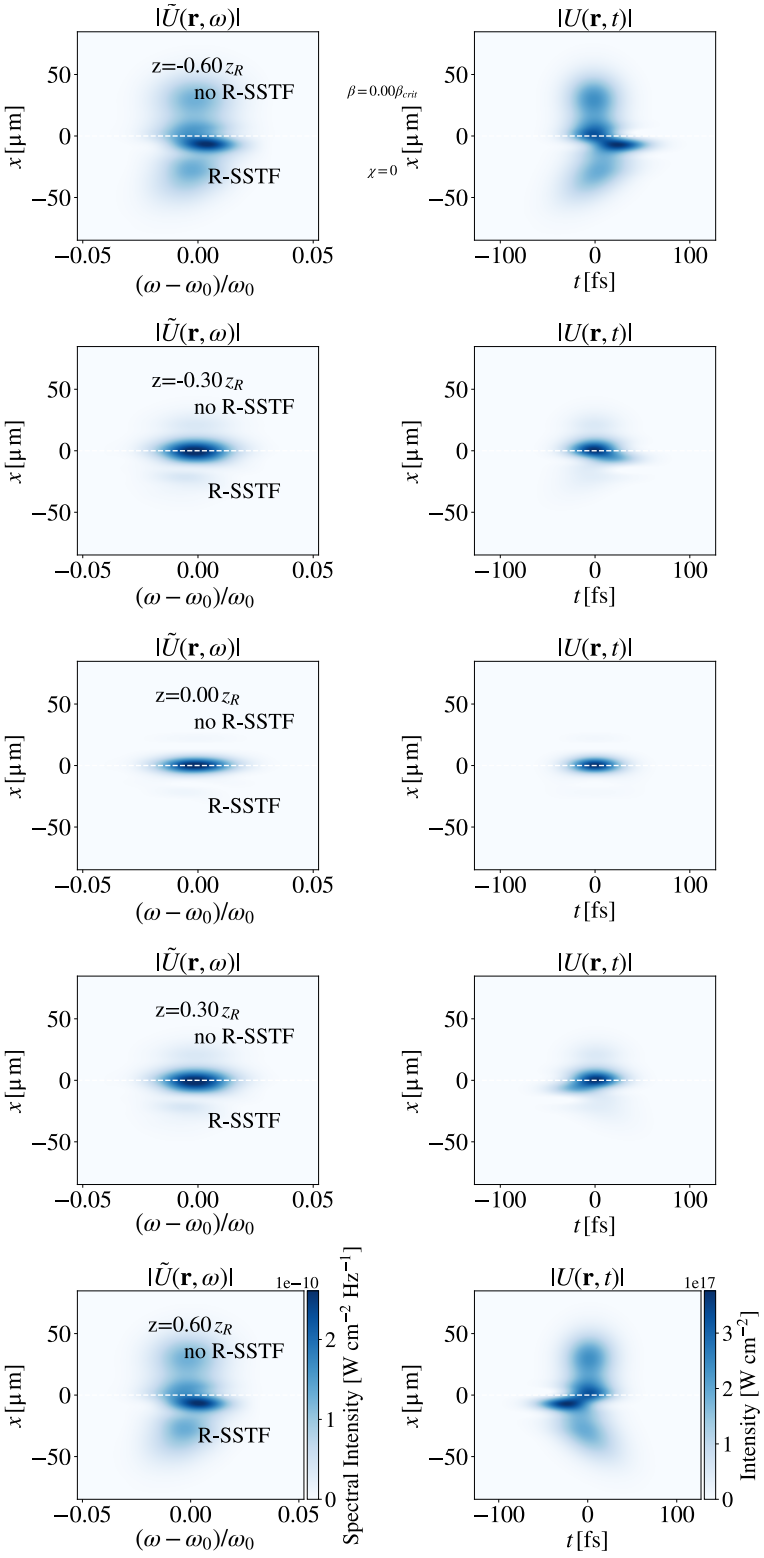


Figure 5.6: Figure showing the evolution of the radial SSTF around the focus, for an input beam with annular radius $r_0 = 20\,\text{mm}$. Spatiospectral intensity (left column) is Fourier transformed to the time domain to give the spatiotemporal intensity (right column). The upper half of each plot shows the intensity of the field without R-SSTF ($\alpha = 0$) and the lower half shows the intensity of the field with R-SSTF ($\alpha = -1$). Pulse-front tilt (seen as a slant of the spatiotemporal field) is only present off-axis for the pulse with R-SSTF, but unlike the linear SSTF (Chapter 4) this is negligible at the focus.

5.4 Results

We assume the following laser pulse parameters:

Pulse energy, $\mathscr{E}$	1 J
Central wavelength, λ_0	800 nm
Uncompressed pulse duration, τ_{p0}	45 fs
Input annulus size, r_0	60 mm[†]
Input spot size, w_0	10 mm
Nominal focal length, f_0	1 m

[†]20 mm for Figure 5.6 only.

Table 5.1: Table of laser pulse parameters for R-SSTF simulations

In this section, we study the evolution of some of the key features and parameters of the R-SSTF pulse, before studying the degree of axial localisation. Eq. 5.39 gives an analytic expression for the spatiospectral field of each mode after the lens. Multiplying each mode U_p^0 by the corresponding mode coefficient $c_p(\omega)$ given by Eq. 5.17 and summing over all modes gives the full spatiospectral field of the R-SSTF. Numerical calculation of the spatiospectral fields is also performed using the Collins method described in the previous section. To avoid unnecessary duplication, results are plotted for the Collins method only (see Section 2.3.2). Both methods give $\tilde{U}_{\mathrm{out}}(x,y,\omega)$ at a particular axial position z. The spatiospectral field $\tilde{U}_{\mathrm{out}}(x,y,\omega)$ is then numerically Fourier transformed to the time-domain to give $U(x,y,t)$ at each axial position z; repeating this for different longitudinal positions in a grid that spans a number of Rayleigh ranges around the focus gives the full spatiotemporal profile $U_{\mathrm{out}}(x,y,z,t)$. The expected Rayleigh range $z'_{\mathrm{R}} = f_0^2 z_{\mathrm{R}}/(z_{\mathrm{R}}^2 + f_0^2)$ of a lowest-order Gaussian pulse of Rayleigh range z_R focused by a lens of focal length f_0 is used to normalise the longitudinal axis for both the SSTF and non-SSTF fields, where $z_{\mathrm{R}} = \pi w_0^2/\lambda_0$ is the Rayleigh range of the Laguerre-Gaussian beamlet before the lens. In performing the calculations a hybrid grid is used which has a finer mesh for the region within half a Rayleigh range of the focus.

Figure 5.6 shows the spatiospectral (left column) and spatiotemporal (right column) fields at different axial positions z relative to the focus f_0. The upper half of each plot shows the evolution of a focus which has no R-SSTF ($\alpha = 0$), whereas the lower half of each plot shows a focus with R-SSTF ($\alpha = -1$). In both cases, the pulse comes to a focus at $z = 0$ before diverging

again. For the R-SSTF case, the spectral components overlap on the optical axis only at the focus. With increasing spectral overlap, the on-axis pulse duration shortens; maximum temporal compression occurs at the focus as expected (this is more obvious in Figure 5.7b) and will be discussed in Section 5.4.1). Pulse-front tilt (PFT)[174] of the R-SSTF can be seen off-axis away from the focus, but appears to be negligible at the focus.

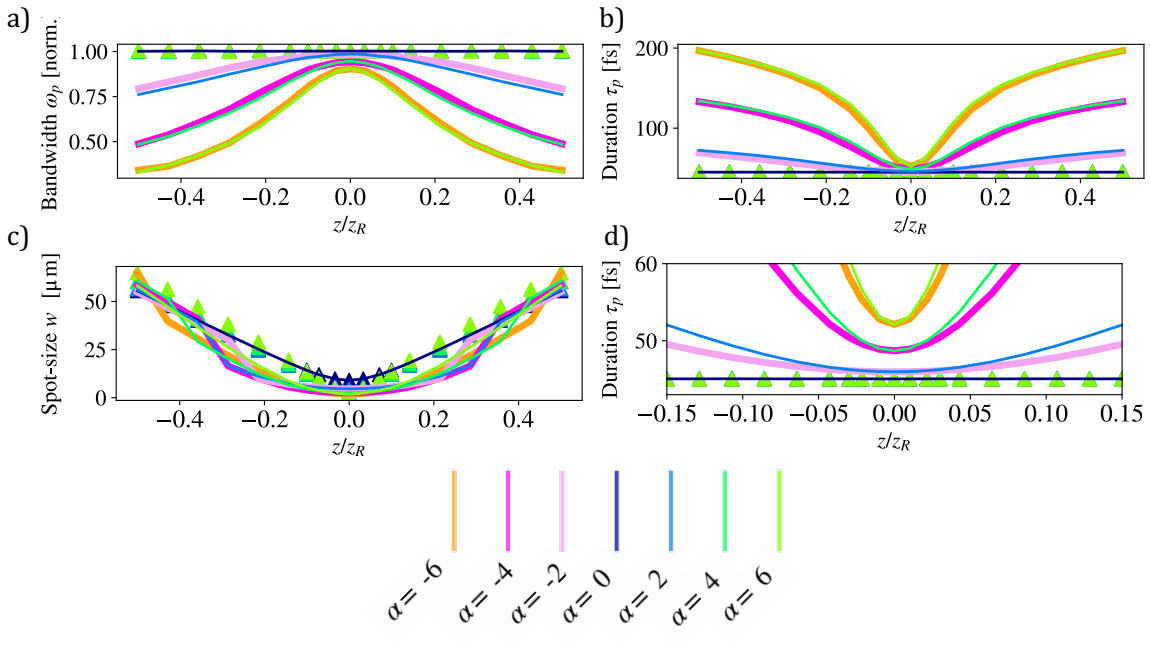


Figure 5.7: Key parameters of the radial SSTF pulse calculated using the Collins method as a function of z in the focal region showing: a) the FWHM bandwidth $\Delta\omega$, b) the FWHM pulse duration τ_p, c) the $1/e^2$ spot size $w(z)$ and d) same data as b) but plotted on a different scale to show that the pulse is not fully-recompressed at the focus. Solid lines denote the SSTF pulse ($\alpha \neq 0$) and triangles are the corresponding non-SSTF fair comparison ($\alpha = 0$) defined in Section 4.4.3.

5.4.1 Centroid, Ellipsoidal Spot Size, Bandwidth and Pulse Duration

In order to understand the evolution of the peak intensity of the R-SSTF pulse in the focal region, we first determine the longitudinal dependence of the key parameters of the pulse — i.e. the spot size, bandwidth, and pulse duration — as follows.

For each longitudinal position z, the time at which the peak axial intensity occurs is found. At this time, a cross-section of the pulse in the xy-plane is taken. For the axially symmetric R-SSTF, the D4σ beam diameter is given in terms of the second moment of the intensity distribution I as[215,216]

$$\mathrm{D4}\sigma = 4\sqrt{\frac{\int_{-\infty}^{\infty} I(x,y)(x-\bar{x})^2\,\mathrm{d}x\,\mathrm{d}y}{\int_{-\infty}^{\infty} I(x,y)\,\mathrm{d}x\,\mathrm{d}y}} = 4\sqrt{\frac{\int_{-\infty}^{\infty} I(x,y)(y-\bar{y})^2\,\mathrm{d}x\,\mathrm{d}y}{\int_{-\infty}^{\infty} I(x,y)\,\mathrm{d}x\,\mathrm{d}y}}, \tag{5.41}$$

where we have written this in Cartesian coordinates since the simulation is performed using a Cartesian grid. The x- and y-coordinate centroids are given by

$$\bar{x} = \frac{\int_{-\infty}^{\infty} I(x,y)x\,\mathrm{d}x\,\mathrm{d}y}{\int_{-\infty}^{\infty} I(x,y)\,\mathrm{d}x\,\mathrm{d}y}, \qquad \bar{y} = \frac{\int_{-\infty}^{\infty} I(x,y)y\,\mathrm{d}x\,\mathrm{d}y}{\int_{-\infty}^{\infty} I(x,y)\,\mathrm{d}x\,\mathrm{d}y}. \tag{5.42}$$

The D4σ beam diameter can be converted to the $1/\mathrm{e}^2$ Gaussian spot size w as follows

$$w = \frac{\mathrm{D4}\sigma_x}{2} = \frac{\mathrm{D4}\sigma_y}{2}. \tag{5.43}$$

The FWHM bandwidth and pulse duration are calculated at each axial distance z by sitting on the optical axis and recording the intensity profile that passes, whether in the spectral or temporal domain. The width of the pulse between the two values that are half of the peak intensity is calculated to give either the FWHM spectral bandwidth or pulse duration (depending on whether we are looking at the pulse pre- or post-Fourier transform from the spectral domain to the time domain). Calculation of the FWHM bandwidth and pulse duration could be improved by fitting a Gaussian to the data at each axial location.

The evolution of key parameters is plotted in Figures 5.7a)–d). In contrast to the linear SSTF, the evolution of the spot size $w(z)$ is very similar for all different values of the chirp α, as shown in Figure 5.7c). This is due to the fact that the radius r_0 (which is fixed to be the same value for all input spatial chirps), rather than the width w_0, has a greater impact on the overall spatial extent of the annular input beam, such that the spatial focusing of the R-SSTF beams only changes slightly as the input chirp is varied. The variation in spot size w as a function of axial position z is not perfectly smooth, and we believe that this is due to the fact that the R-SSTF is composed of higher-order Laguerre-Gaussian modes, such that overall the R-SSTF pulse has a higher M^2 beam quality factor. The Gaussian spot size, calculated from the D4σ beam diameter is therefore a much poorer measure of the spatial extent of the beam.[216] As we can see in Figure 5.6, the transverse spatial structure of the beam changes significantly as the R-SSTF propagates and this seems to be reflected in the non-smooth variation of the spot size in Figure 5.7c).

Figures 5.7a) and b) confirm that, for the SSTF case, the bandwidth increases (and the pulse duration decreases) as the focus is approached. Away from the focus, the pulse duration can increase from the maximally-compressed value of 53 fs by as much as a factor of 3.6 for the highest chirp considered $\alpha = \pm 6$. This time, however, the pulse is no longer fully-recompressed to 45 fs

at the focus, as shown in Figure 5.7d). This would indicate that there are residual spatiotemporal couplings at the focus and we will see in Section 5.5 that these give rise to a focal spot with an arbitrary group velocity.

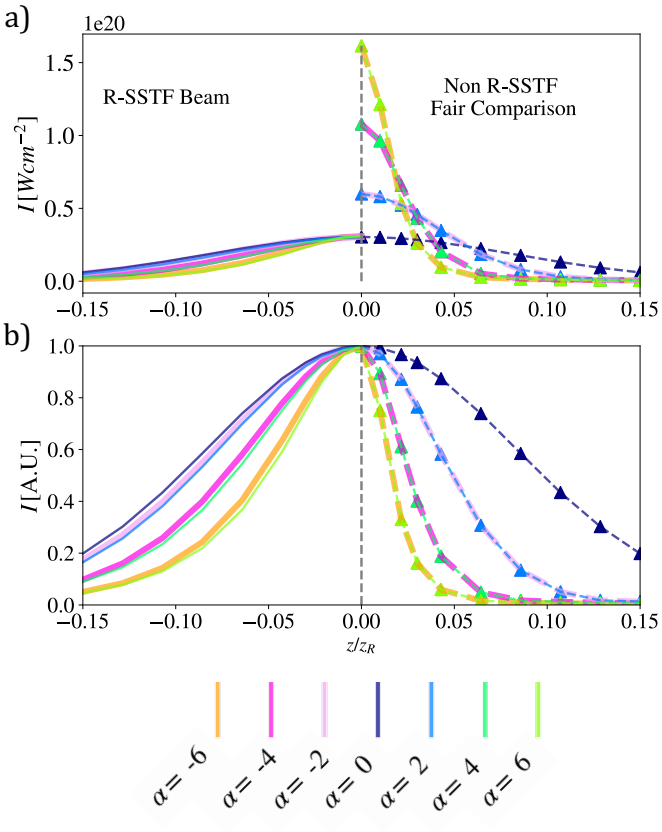


Figure 5.8: Longitudinal variation of the axial peak intensity of the radial SSTF pulse as a function of z in the focal region. a) Absolute peak intensity I_0 calculated by extracting the peak intensity (PI) (see Section 4.4.2) of the fields at all axial positions. b) Comparison of the longitudinal variation of the peak axial intensity, normalised to the value at focus. Solid lines on the left denote the SSTF pulse ($\alpha \neq 0$) and dashed lines / triangles on the right are the corresponding non-SSTF fair comparison ($\alpha = 0$) defined in Section 4.4.3.

5.4.2 Calculation of Peak Axial Intensity

To calculate the peak intensity I_{peak} of the R-SSTF beam, the same Method 1 is used as in Section 4.4.2. The peak intensity is calculated from the spatiotemporal electric fields $E(x, y, z)$ as

$$I_{\text{peak}} = \max\left\{\varepsilon_0 c|E|^2\right\}. \tag{5.44}$$

We do not attempt any estimation of the peak intensity using the spot size and pulse duration, as we found in Section 4.4.2 that such methods do not work for beams with spatiotemporal couplings. In addition to this, our beam is no longer a lowest-order Gaussian. Therefore, simplified models assuming an untilted, ellipsoidal, Gaussian beam profile cannot be used to calculate the

intensity of the beam, and we present only the calculation of the intensity directly from the field.

The left-hand side of Figure 5.8 shows the intensity of the R-SSTF beam calculated directly from the electric field strength. As the magnitude of the spatial chirp increases, the focus appears to be more axially localised, as the intensity falls away more steeply from $z = 0$. On the right-hand side of Figure 5.8, we plot the non-SSTF fair comparison. This was discussed in Section 4.4.3, and is calculated by integrating the spatial profile of the input radially-chirped beam then multiplying by a chirp-free Gaussian spectrum, to give a non-STC equivalent. This accounts for any change in the spatial profile of the input beam as the chirp strength α is varied. Looking at Figure 5.8a), we see that the peak intensity of the R-SSTF beam is equal to the peak intensity of the pre-chirped case ($\alpha = 0$), which we also found to be true for the linear SSTF beam. Comparing the normalised axial intensity of the R-SSTF and the fair comparison in Figure 5.8b), we see that in this case, the axial localisation of the R-SSTF is worse compared to the fair comparison.

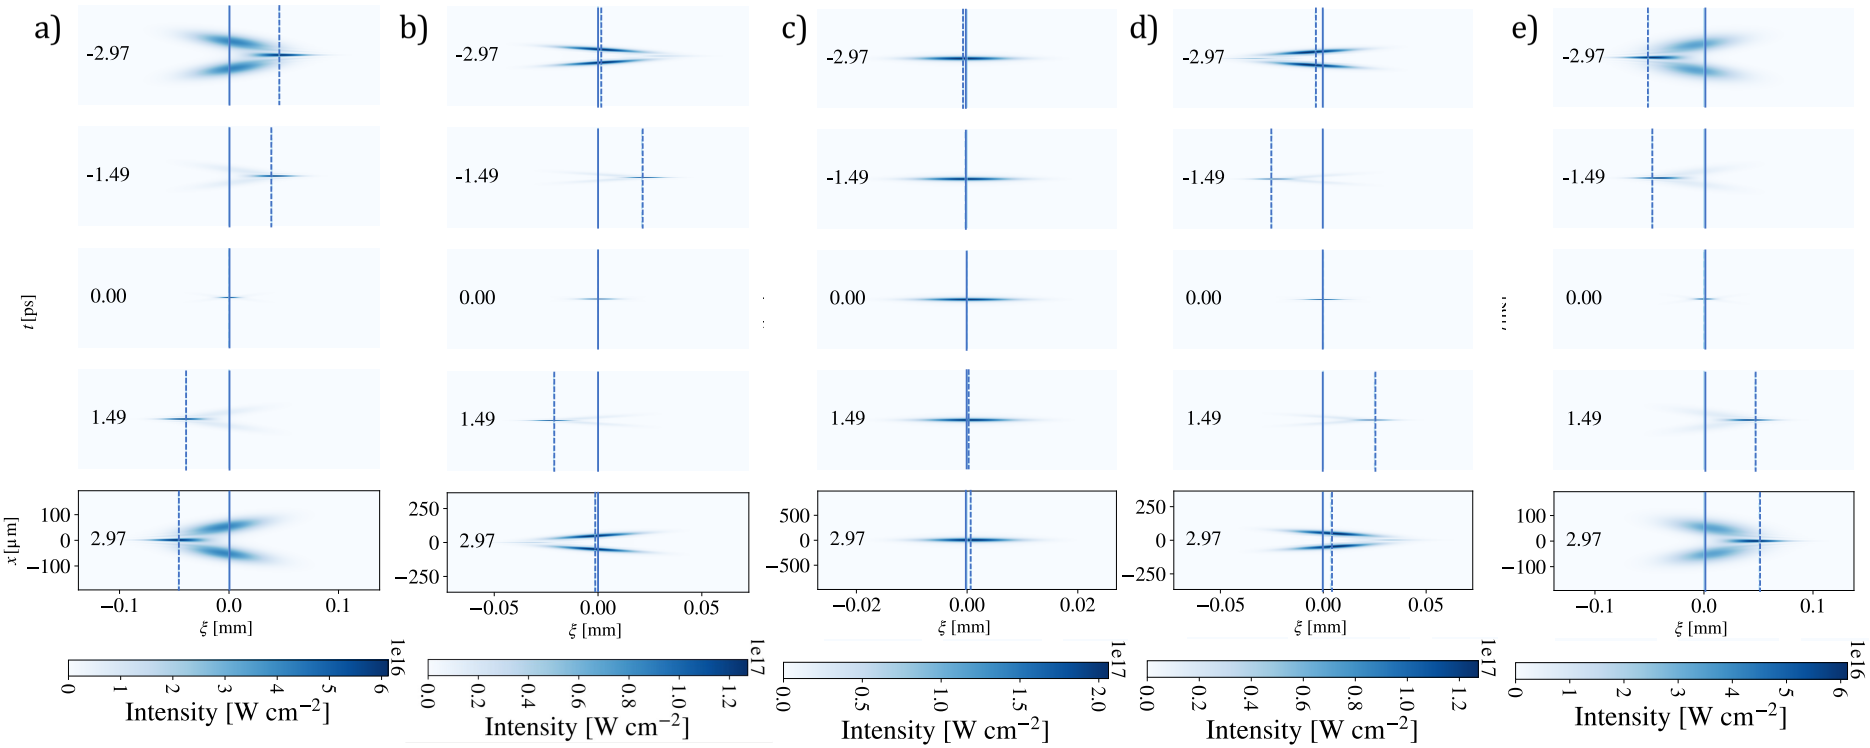


Figure 5.9: Figure showing snapshots of the radial simultaneous space-time focus around the focus for different input spatial chirps α. a) $\alpha = -5\alpha_{\text{crit}}$ with $v_{\text{sstf}} = 0.90c$, b) $\alpha = -2.5\alpha_{\text{crit}}$ with $v_{\text{sstf}} = 0.96c$, c) the non-SSTF beam where $\alpha = 0$ with $v_{\text{sstf}} = c$, d) $\alpha = 2.5\alpha_{\text{crit}}$ with $v_{\text{sstf}} = 1.06c$ and e) $\alpha = 5\alpha_{\text{crit}}$ with $v_{\text{sstf}} = 1.11c$. Dashed lines show the position of the R-SSTF peak intensity used to calculate estimated group velocities v_{sstf} of the focal spots, solid lines show the expected position of a non-flying focus where $\alpha = 0$.

5.5 Controllable Group Velocity

Although the degree of axial localisation of the intensity of the R-SSTF beam is worse than for the fair comparison case, R-SSTF beams may still have applications for processes which depend on the duration of the laser pulse, such as in multiphoton imaging[203,205,207] and laser machining.[206,207] We found that although the pulse does not fully-recompress to a pulse duration of 45 fs, the pulse duration can increase from the maximally-compressed value of 53 fs by as much as a factor of 3.6 for the highest chirp considered ($\alpha = \pm 6$). Therefore, the R-SSTF can still be used to localise time-dependent processes, much like the linear SSTF. The residual pulse-front tilt of the more commonly-used linear SSTF has been shown to cause non-reciprocal or 'quill' effects in the laser machining of solids, such that the material is modified asymmetrically depending on the scanning direction relative to the direction of the pulse-front tilt.[219–225] If we wish to modify our sample uniformly as we scan back and forth across it, then we require a pulse without pulse-front tilt at the focus. We saw in Figure 5.6 that the R-SSTF pulse has negligible pulse-front tilt at the focus. The R-SSTF could therefore have an advantage over the linear SSTF for such applications.

An additional effect that is observed for the R-SSTF beam is a sub- or superluminal apparent group velocity, which depends on the strength of a input spatial chirp α. As was described in Chapter 3, the spatiotemporal field at different longitudinal positions may be converted to snapshots in time. In a frame that is moving at the speed of light, the total propagation distance $|z - z_0|$ from the input plane at z_0 corresponds to a time of occurrence $t_0 = |z - z_0|/c$. This is the time at which the $t = 0$ of the calculated pulse coincides with the point z. We have calculated the temporal profile of the pulse passing through this point z, and the field in the light-speed moving-frame coordinate $\xi = z - ct$ can therefore be calculated from temporal profile of the pulse $U_{\text{out}}(x, y, t)$ at that z position, by simply rescaling the time grid by a factor of c, such that $\xi_{\text{grid}} = ct_{\text{grid}}$, since $t_{\text{grid}} = 0$ is the time of arrival of a light-speed pulse. In this way we can convert our evolution of the pulse at different axial z positions to snapshots at different times, allowing us to visualise the evolution of the pulse in a fixed moving-frame window ξ.

Figure 5.9 shows how the apparent velocity of the focal spot changes when we vary the input spatial chirp. This effect is similar to what we saw in Figure 3.5 of Chapter 3. Figure 5.9a) would indicate that negative chirps can generate focal spots with subluminal apparent group

velocities, whereas the positive chirp shown in Figure 5.9e) gives rise to a superluminal focal spot. Here the velocity of the R-SSTF ranges from $v_{\text{sstf}} = 0.90c$ to $v_{\text{sstf}} = 1.11c$. The velocity of the simultaneous space-time focus is calculated from the simulated focal spot data as follows. Before transformation to the moving-frame coordinate, for each axial distance $|z - z_0|$ at which the fields are calculated, the relative time t_0 at which the peak intensity occurs is found by taking a lineout of the intensity at $x = 0$. The R-SSTF flying focus velocity is estimated to be

$$\frac{v_{\text{ff}}}{c} = \frac{|z - z_0|}{\big(|z - z_0|/c - t_0\big)/c}. \tag{5.45}$$

The calculated velocities of all focal spots are given in Figure 5.9. The offset time t_0 is converted to a position in the moving-frame coordinate and is plotted as dashed lines in Figure 5.9. The solid lines denote the expected position of the focal spot without spatiotemporal couplings.

5.6 Conclusions and Future Work

To summarise, in this chapter we investigated whether a simultaneous space-time focus formed by focusing a radially-chirped input beam with an achromatic lens offered any improvement in the degree of axial localisation of the peak intensity. We found that the degree of axial localisation of the R-SSTF is worse when compared to the equivalent fair-comparison pulse, as defined in the previous chapter, and that although the R-SSTF exhibits temporal recompression, the pulse was not fully-recompressed at the focus. It did not seem possible to obtain a closed-form analytic expression for the R-SSTF in the focal region, therefore an alternative method which decomposed the radially-chirped input beam as a superposition of Laguerre-Gaussian modes was used to calculate the output field analytically. Each mode was propagated independently and the superposition of these modes gave the fields in the vicinity of the focus. The spatiospectral fields were calculated both using a superposition of the modes and the Collins diffraction integral. We found that although the degree of axial localisation of the R-SSTF is worse than for the fair comparison, R-SSTFs have properties that may make them useful for application in laser-machining. We found that the R-SSTF beam does not have appreciable pulse-front tilt at the focus, making it possible, in theory, to achieve uniform material modification, regardless of scanning direction, and also that the R-SSTF has a controllable group-velocity which depends on the strength of the input spatial chirp.

In the future, we will investigate the dependence of the apparent focal velocity of the R-SSTF

on the input spatial chirp and other laser pulse properties, to understand the range of different focal velocities that are possible with an R-SSTF. We will also see whether it is possible to derive an analytic expression for the focal velocity in terms of the spectral phase of the Laguerre-Gaussian modes that comprise the R-SSTF.

6 | The Radially-Polarised Focus

6.1 Introduction

In the introduction to this thesis, we explored the concept of accelerating charged particles in a plasma wakefield driven by a high-intensity laser pulse. Direct laser acceleration (DLA) is the acceleration of charged particles, in this case electrons, in the electric fields of the laser pulse itself. The laser intensities required for direct laser acceleration are generally greater than $10^{18}\,\mathrm{W\,cm^{-2}}$. We will review a few different mechanisms by which DLA can occur. Considering the most basic theoretical models of DLA, the first studies considered the interaction between a laser pulse and a very small number of electrons surrounded by a vacuum.[226,227] Such studies often made the assumption that these electrons are cold and fixed at a particular location defined with respect to the laser pulse. Using the relativistic Lorentz equation, we can derive the single particle motion of the electron as it quivers in the oscillating electric field of the laser pulse and any subsequent drifts on longer timescales by Alfvén averaging.[32,228]

The Lawson-Woodward theorem states that no net energy gain is possible using the fields of a laser beam.[229–231] The Lawson-Woodward theorem rests on the assumptions that[44]

1. The interaction region is infinite,
2. The laser fields are in a vacuum without walls or boundaries,
3. The electron is highly-relativistic,
4. There are no static electromagnetic fields,
5. Non-linear effects (such as ponderomotive or radiation reaction forces) are neglected.

Charged particles can be accelerated in laser fields if one or more of the above assumptions are broken. We can note that the Lawson-Woodward theorem is already broken by the ponderomo-

tive force of the laser pulse. Clearly, single particle motion of an electron in a vacuum is already an idealised picture of any real interaction, so we will instead consider a target gas species which is ionised by the high-intensity laser pulse that is performing the DLA[†].[232]

In the ultra-low density regimes where we can consider the motion of individual electrons in free-space, the electron density is insufficient to support a wakefield. In this case, the longitudinal component of the electric field of the laser pulse is crucial for direct laser acceleration. Tightly-focused linearly-polarised Gaussian beams have a non-zero component of the longitudinal electric field (as will be discussed in Section 6.2.1), however the phase velocity of the field near the optical axis is greater than c, such that the accelerated electrons rapidly phase slip relative to the field and are then decelerated.[44] Gaussian beams also have an intensity profile which is peaked on-axis, driving electrons away from the axis via the ponderomotive force. In order to counteract this effect, early in the development of DLA concepts it was suggested that the higher-order modes of a laser beam, such as linearly-polarised Hermite-Gaussian (HG)[226] and Laguerre-Gaussian (LG) modes,[233] could be used to accelerate electrons instead.

Radially-polarised laser beams (RPLBs) can accelerate electrons to substantial energies, even in a vacuum, due to their large longitudinal electric field component at focus.[227,234,235] RPLBs have a doughnut-like transverse intensity profile and can be written as a superposition of the $HG^{0,1}$ or $LG^{0,1}$ modes.[236,237] They can be generated using laser resonators,[238–240] axicons,[241,242] interferometers,[236,243,244] waveguides[245–247] as well as specially designed optics[248,249] and devices.[250,251] On-axis, the transverse components of the field are zero and the longitudinal electric field is maximal. The exact form that the fields of a RPLB take will be discussed in Section 6.2.2. Numerous studies have investigated the effect of different parameters and the scaling laws for direct laser acceleration of electrons using RPLBs.[234,235,252,253] In a simple approximation, for an electron initially at rest, the maximum possible net energy gain in a RPLB scales like the square-root of the incident laser power P[252]

$$\Delta W_{\mathrm{max}}[\mathrm{GeV}] \simeq \Big(P[\mathrm{PW}]\Big)^{1/2}, \tag{6.1}$$

which makes sense since the laser electric field strength scales like $E \sim \sqrt{I}$, where I is the in-

[†]An intensity of $10^{18}\,\mathrm{W\,cm^{-2}}$ is four orders of magnitude larger than the ionisation threshold for hydrogen (see Section 1.4.1), such that the laser pulse will ionise most target gas species. Therefore, ionisation and acceleration can be done using a single laser pulse.

tensity. For the 6 TW Ti:Sa laser system considered in Section 1.2.2, the maximum possible net energy gain is on the order of 80 MeV. It is also important to consider temporal and spectral properties of the laser pulse, such as CEP phase, chirps and spatiotemporal couplings, which can have a significant effect on the final energy of accelerated electron bunches, particularly for ultrashort RPLBs.[254,255]

The interaction between a laser pulse and a gas increases the overall energy gain compared to that obtained in a vacuum.[226] Several regimes can be identified that depend on the laser pulse intensity and plasma density. In 1995, a radially-polarised laser beam generated using an axicon was used to accelerate electrons using the inverse Cherenkov effect, where the target gas was used to reduce the phase velocity of the laser pulse in order to achieve phase matching with the electron beam.[256] Esarey *et al.* published a theoretical study the following year which modelled the ionisation of a gas and propagation in the subsequent plasma by a RPLB.[257] The acceleration of electrons from a low-density plasma using RPLBs has been studied in great detail[258–260] and numerous theoretical studies indicate that it might be possible to produce sub-attosecond electron bunches using RPLBs.[261–263] Although the acceleration of electrons from a 40 keV electron gun by a RPLB was demonstrated in 2016,[264] the net energy gain from this experiment is not completely clear. Another experiment generated 23 keV electrons by ionising oxygen in an ambient gas using a RPLB.[234,235]

As the electron density increases, the laser pulse may also drive a wakefield. Even if the laser pulse duration is long, self-modulated laser wakefield acceleration (SM-LWFA) can occur (this was discussed in Section 1.2.2). There have been numerous studies of the interplay between LWFA and DLA, whether by a single linearly-polarised laser pulse (LPLB)[265,266] or a pair of linearly-polarised laser pulses,[267] some of which indicate that DLA can increase the maximum energy gain relative to LWFA-only schemes. Others have studied the effect of different properties of the laser pulse in DLA:LWFA hybrid schemes, such as the relative polarisation and wavelength,[268] or the frequency-chirp.[269] In 2020, a scheme which used a RPLB to drive an inverted wake was proposed, with the potential application of accelerating positrons.[270] It is also known that DLA in a trailing, co-propagating, linearly-polarised pulse can instigate betatron motion of the accelerated electrons, due to their transverse quiver motion,[271] potentially with applications for light-source generation. All of these schemes consider LWFA in the bubble regime (as discussed in Section 1.2.2).

At high, but still sub-critical plasma densities, sufficiently intense ($I > 10^{18}\,\mathrm{W\,cm^{-2}}$) picosecond laser pulses can drive electrons away from the optical axis to form a plasma ion channel.[272–276] It has been shown that in such regimes, DLA can occur by a completely different mechanism: namely that the electrostatic fields of the ion channel support strong electron betatron oscillations,[275] which can couple resonantly to the oscillations in the field of the LPLB,[277] leading to acceleration beyond the vacuum limit.[278–282] However, in this chapter, we shall instead consider regimes where wakefields, not ion channels, are formed.

We suggest that the radial symmetry and more localised intensity of the RPLB may offer an advantage for generating ultrashort, low-emittance electron bunches. Additionally, we derive a new model for laser-driven wakefield acceleration (LWFA) which includes the longitudinal field component of a laser pulse. We do not believe that such a derivation has been attempted before. An alternative scheme, related to the two-pulse ionisation injection (2PII) scheme[146] but utilising novel polarisation states, is proposed. It was previously shown that a trailing laser pulse can ionize electrons from a dopant species, and that these can be trapped if the driving and trailing pulses are separated so that the wakefields from the two pulses add coherently. Injection of electrons into a quasilinear wakefield using 2PII was shown to produce quasi-monoenergetic electron bunches on the order of 100 MeV with a relative energy spreads of $\sim 10\%$.[147] Here we perform a preliminary analytic study, rederiving the wakefield quantities and Hamiltonian from first principles and allowing for alternative polarisation states of the laser pulse, to investigate whether changing the polarisation state of the trailing 'injector' pulse from linear to radial polarisation could offer any advantage for trapping of electrons in quasilinear wakes. We find that the RPLB seems to enhance the wakefield and allows electrons to be trapped when they otherwise wouldn't be in a standard 2PII configuration. The model only considers the motion of on-axis electrons in a quasistatic approximation. The limitations of such a model are discussed and future work is proposed.

6.2 Electromagnetic Fields of Different Polarisation States

Throughout this section, numerical results are shown for a $w_{\mathrm{in}} = 10\,\mathrm{cm}$ spot focused by a $f_0 = 20\,\mathrm{cm}$ lens (same f-number as for the previous 2PII study[146]) and the optical axis points in the z-direction. The calculated spot size at the focus is $w_0 \simeq (\lambda_0 f_0)/\pi w_{\mathrm{in}} = 0.51\,\mu\mathrm{m}$, and it is worth noting that this is a 'sub-wavelength' focal spot, as has been previously studied[283,284]

and demonstrated experimentally.[285] We will later assume a peak normalised vector potential of $a_0 = 1$ for our RPLB (see Section 6.6), which corresponds to a peak intensity of approximately $I_{\text{peak}} \simeq 2 \times 10^{18}\,\text{W cm}^{-2}$ by Eq. 1.8. The peak electric field strength E_0 is therefore

$$E_0 = \sqrt{\frac{I_{\text{peak}}}{\varepsilon_0 c}} \sim 2.7 \times 10^{12}\,\text{V m}^{-1}, \tag{6.2}$$

and the input pulse energy $\mathscr{E}$ required to reach this intensity for our chosen focusing geometry is

$$\mathscr{E} = \frac{w_0^2 \tau_{\text{p0}}}{4} \sqrt{\frac{\pi^3}{\ln 2}} I_{\text{peak}} = 0.4\,\text{mJ}, \tag{6.3}$$

assuming a bi-Gaussian laser pulse which is free from spatiotemporal couplings for this first approximation. We assumed a pulse duration of $\tau_{\text{p0}} = 45\,\text{fs}$, a laser pulse central wavelength of 800 nm and set the peak electric field strength E_0 to be the same for the LPLB and the RPLB.

To summarise, we assume the following laser pulse parameters for all laser pulse pulses:

Pulse energy, $\mathscr{E}$	0.4 mJ
Central wavelength, λ_0	800 nm
Uncompressed pulse duration, τ_{p0}	45 fs
Input spot size, w_{in}	10 cm
Nominal focal length, f_0	20 cm
Estimated parameters at the focus:	
Spot size, w_0	0.5 µm
Peak Intensity, I_{peak}	$2 \times 10^{18}\,\text{W cm}^{-2}$
Peak Electric Field Strength, E_0	$2.7 \times 10^{12}\,\text{V m}^{-1}$
Peak Normalised Vector Potential, a_0	1

Table 6.1: Table of laser pulse parameters for RPLB simulations

6.2.1 Linearly-Polarised bi-Gaussian Beams

A tightly focused Gaussian laser pulse has a longitudinal field component,[286] in addition to the component that points in the direction of polarisation in the transverse plane. The fields of a linearly-polarised bi-Gaussian laser pulse at focus are[287]

$$E_x(\rho, z, t') = E_0 \left[\frac{w_0}{w(z)}\right] \text{e}^{-\frac{\rho^2}{w(z)^2}} \cos(\psi) f(t'), \tag{6.4}$$

$$E_z(\rho, z, t') = -\frac{i\theta_0}{1 + i\frac{z}{z_{\text{R}}}} \left(\frac{\rho}{w_0}\right) E_x, \tag{6.5}$$

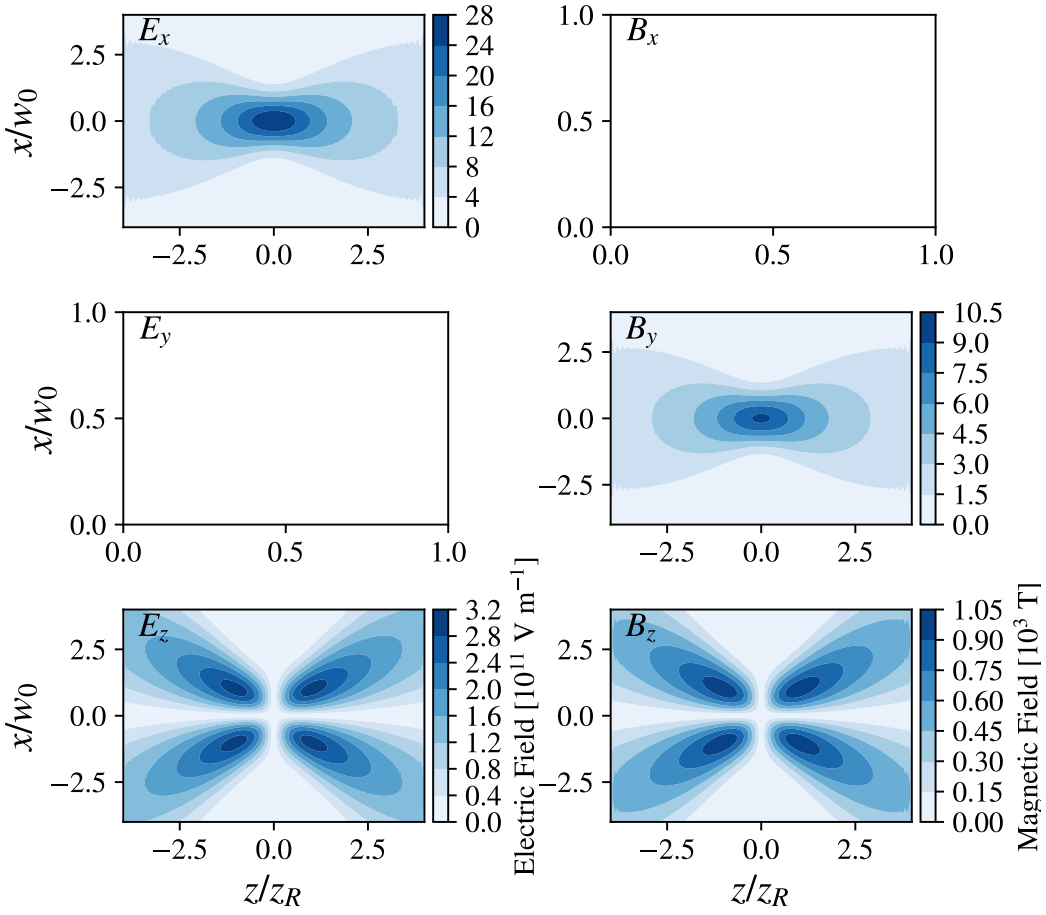


Figure 6.1: Spatial components of the electromagnetic fields of a linearly-polarised Gaussian beam at focus, for the laser pulse parameters given in Table 6.1. Lineouts of the electric field at $x = 0$ and $z = 0$ are shown in Figure 6.3.

$$B_y(\rho, z, t') = E_x/c, \tag{6.6}$$

$$B_z(\rho, z, t') = E_z/c, \tag{6.7}$$

where

$$\psi = \omega_0 t - k_0 z + \psi_{\text{CEP}}, \tag{6.8}$$

$$= \omega_0 t - k_0 z + 2\tan^{-1}\left(\frac{z_\text{R}}{z}\right) - \frac{k_0 \rho^2}{2z(1 + z_\text{R}^2/z^2)} - \psi_0, \tag{6.9}$$

is the phase, E_0 is the peak electric field strength on-axis, $w(z) = w_0\sqrt{1 + (z/z_\text{R})^2}$ is the beam radius in terms of the $1/\text{e}^2$ Gaussian spot size w_0 at the focal position and $\theta_0 = 2/(k_0 w_0)$ is the far-field divergence half angle. The temporal pulse envelope is given by $f(t') = \exp\{-2\ln(2)t'^2/\tau_{\text{p0}}^2\}$,

with a pulse duration τ_{p0}, where $t' = t - z/v_{g,L}$ and $v_{g,L}$ is the group velocity of a Gaussian beam in free-space. All other components of the field are zero.

The fields of the linearly-polarised bi-Gaussian beam are plotted in Figure 6.1. The longitudinal field has a maximum amplitude at $w_0/\sqrt{2}$ given by[286]

$$|E_z|_{\text{peak}} \simeq 0.606\theta_0 |E_x|_{\text{peak}}, \tag{6.10}$$

where the peak $|E_x|_{\text{peak}}$ is the peak transverse electric field at focus. For the fields shown in Figure 6.1, the estimated ratio between the peak longitudinal electric field and peak transverse electric field, calculated from Eq. 6.10, is $|E_z|_{\text{peak}}|/|E_x|_{\text{peak}} = 0.152$ and the calculated ratio from taking the maxima of the fields plotted in Figure 6.1 is $|E_z|_{\text{peak}}|/|E_x|_{\text{peak}} = 0.108$, giving a difference in the field ratio of 0.0435 between these two methods.

6.2.2 Radially-Polarised Laser Beams

The fields of a RPLB are given as[252]

$$E_r(\rho, z, t') = E_0 \frac{\sqrt{2}\rho}{w_0} \exp(1/2) \left[\frac{w_0}{w(z)}\right]^2 \mathrm{e}^{-\frac{\rho^2}{w(z)^2}} \cos(\psi) g(t'), \tag{6.11}$$

$$E_z(\rho, z, t') = E_0 \frac{2\sqrt{2}}{k_0 w_0} \exp(1/2) \left[\frac{w_0}{w(z)}\right]^2 \mathrm{e}^{-\frac{\rho^2}{w(z)^2}} \left\{\left[1 - \frac{\rho^2}{w(z)^2}\right] \sin(\psi) - \frac{k_0 \rho^2}{2z\left(1 - \frac{z_R^2}{z^2}\right)} \cos(\psi)\right\} g(t'), \tag{6.12}$$

$$B_\theta(\rho, z, t') = E_r/c, \tag{6.13}$$

where all parameters were defined in Section 6.2.1. The only difference is that now the temporal pulse envelope is given by $g(t') = \exp\{-2\ln(2)t'^2/\tau_{p0}^2\}$, with $t' = t - z/v_{g,R} - \Delta t$ where $v_{g,R}$ is the group velocity of a RPLB in free-space (which will be derived in Section 6.4) and Δt is an arbitrary time delay between the linearly-polarised drive beam and trailing radially-polarised injector beam in a 2PII scheme (as will be discussed in Section 6.5). All other components of the field are zero.

By plotting these fields, we can see from Figure 6.2 that radially-polarised laser beams (RPLBs) have a strong axial electric field component E_z at focus. The longitudinal electric field has a maximum amplitude in terms of the peak electric field strength of the transverse component E_r at

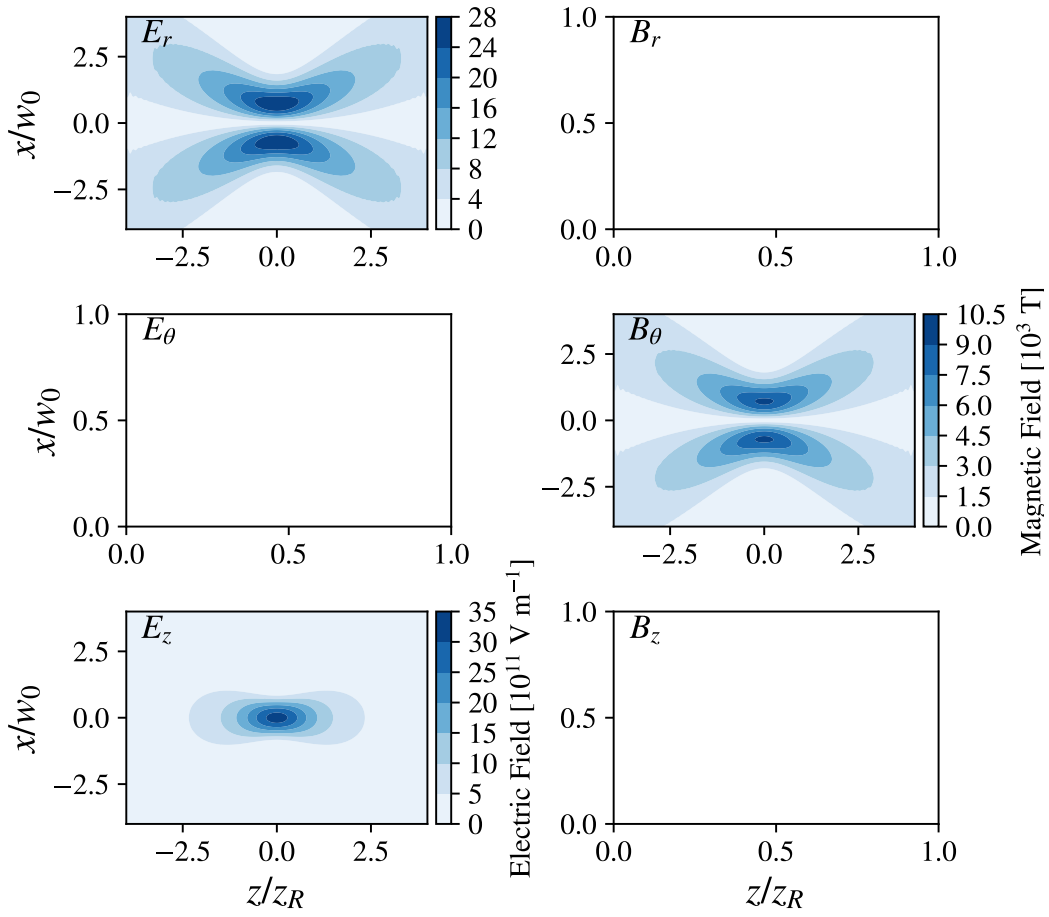


Figure 6.2: Spatial components of the electromagnetic fields of a radially-polarised laser beam at focus, for the laser pulse parameters given in Table 6.1. Lineouts of the electric field at $x = 0$ and $z = 0$ are shown in Figure 6.3.

$w_0/\sqrt{2}$ given by[235, 252, 288†]

$$|E_z|_{\text{peak}} = \frac{2\sqrt{2}e^{1/2}}{k_0 w_0}|E_r|_{\text{peak}}. \tag{6.14}$$

For the fields shown in Figure 6.2, the estimated ratio between the peak longitudinal electric field and peak transverse electric field, calculated from Eq. 6.14, is $|E_z|_{\text{peak}}|/|E_r|_{\text{peak}} = 1.166$ and the calculated ratio from taking the maxima of the fields plotted in Figure 6.2 is $|E_z|_{\text{peak}}|/|E_r|_{\text{peak}} = 1.151$, giving a difference of 0.0152 in the ratio of the fields between these two methods.

Figure 6.3 compares the on-axis electric fields of the LPLB and the RPLB. We see that the LPLB only has a transverse electric field component and the RPLB only has a longitudinal electric field component. The other components are negligible. This will later inform our decision to

†The third reference here is missing the factor of 2 in their Eq. 6.14. It seems like this is a mistake.

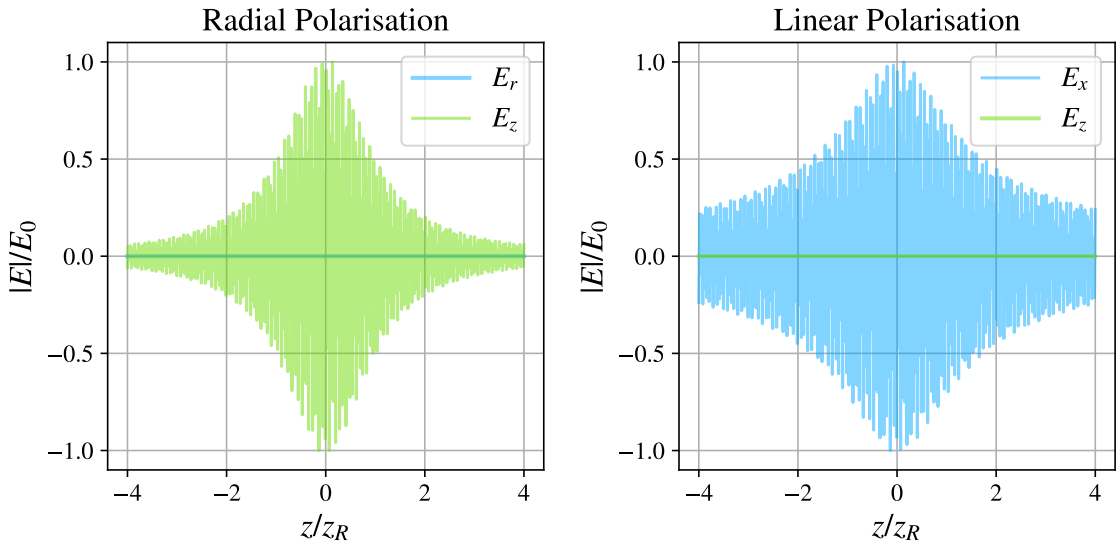


Figure 6.3: Relative amplitude of longitudinal and transverse electric field components along the optical axis. A radially-polarised beam has a negligible on-axis transverse electric field and a significant longitudinal electric field, whereas the converse is true for a linearly-polarised beam.

include only E_x for the LPLB and only E_z for the RPLB when we derive the 1D on-axis wakefield quantities in Section 6.5.2.

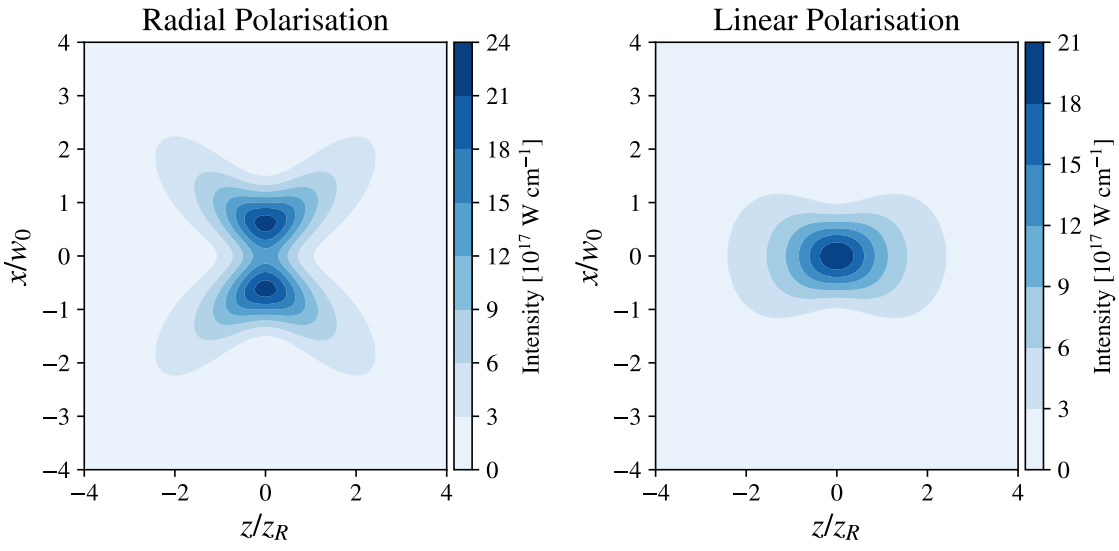


Figure 6.4: Comparison of the longitudinal and transverse variation of intensity of a RPLB and a LPLB for an $f/2$ focusing geometry. All assumed laser pulse parameters are given in Table 6.1.

6.3 Confinement of Intensity for Different Polarisation States

In particular focusing regimes, the spot size of a focused radially or azimuthally polarised beam has been shown to be less than the diffraction limit both theoretically[283,289–291] and experimentally,[284,285,292] potentially leading to smaller electron bunches if such a technique is applied to ionisation injection. Higher-order beams also have higher divergence than their pure Gaussian counterparts,[293] leading to shorter Rayleigh ranges. We will now study whether the intensity of the RPLB is more localised than the bi-Gaussian LPLB counterpart. If the region of highest intensity is more localised for the RPLB, then a RPLB could be used to inject localised electron bunches which have smaller emittances and energy spreads compared to those generated from ionisation injection using a standard LPLB or circularly-polarised laser beam (CPLB).

In Chapter 4, Section 4.4.2 the intensity of an electromagnetic field was defined as

$$I = \frac{c}{2}\Big(\varepsilon_0|\mathbf{E}|^2 + \frac{1}{\mu_0}|\mathbf{B}|^2\Big). \tag{6.15}$$

This gives the intensity of the LPLB as

$$I_{\mathrm{LPLB}} = \frac{c}{2}\Bigg[\varepsilon_0\Big(|E_x|^2 + |E_y|^2 + |E_z|^2\Big) + \frac{1}{\mu_0}\Big(|B_x|^2 + |B_y|^2 + |B_z|^2\Big)\Bigg], \tag{6.16}$$

and similarly for the RPLB

$$I_{\mathrm{RPLB}} = \frac{c}{2}\Bigg[\varepsilon_0\Big(|E_r|^2 + |E_\theta|^2 + |E_z|^2\Big) + \frac{1}{\mu_0}\Big(|B_r|^2 + |B_\theta|^2 + |B_z|^2\Big)\Bigg], \tag{6.17}$$

where we have noted that the components given by Eqs. 6.4–6.7 and 6.11–6.13 are the real electric fields. The Hilbert transformed components $E_i = H(E_i)$ can be used to calculate the complex electric field from the real electric field (see Section 2.1.2), to give an intensity profile that looks smooth.

Figure 6.4 shows the intensity of the RPLB and the LPLB. Lineouts of the intensity along the optical axis and transversely at the focus are shown in Figure 6.5. The grey lines in Figure 6.5b) show the relative contributions of the summed annular $|E_r|^2$ and $|B_\theta|^2$ components (solid grey line) and the peaked on-axis $|E_z|^2$ component (dashed grey line). For an $f/2$ focusing geometry,

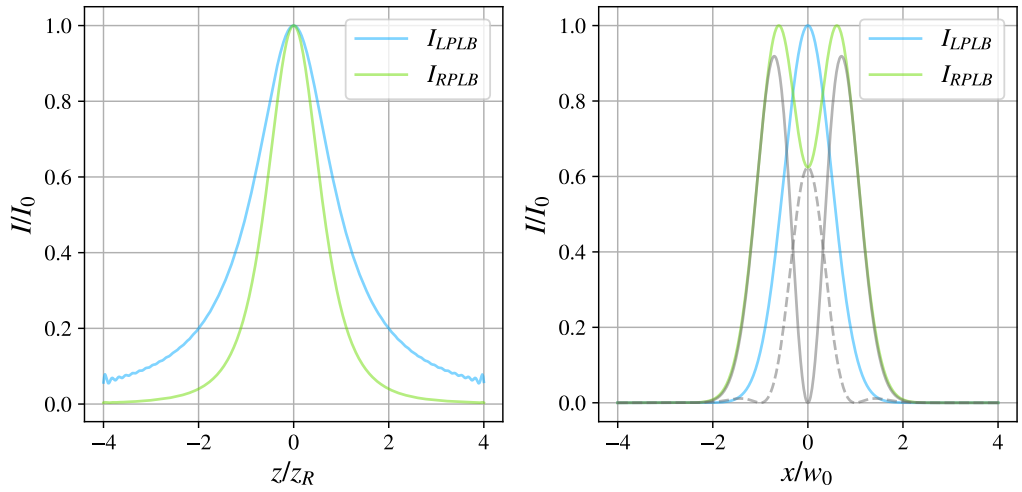


Figure 6.5: Lineouts of the intensity taken a) along the z axis at $x = 0$ (longitudinally) and b) along the x axis at $z = 0$ (transversely) for an $f/2$ focusing geometry. The Rayleigh range of the RPLB, defined as the axial distance from the focus to the point where the intensity is half the peak intensity, is $0.642z_R$ and the spot size, defined as the radial distance from the axis to the point where the intensity is $1/e^2$ the peak intensity, is $1.502w_0$, where z_R and w_0 are the Rayleigh range and spot size of the bi-Gaussian LPLB, defined in the usual way. In b) the solid grey line denotes the summed $|E_r|^2$ and $|B_\theta|^2$ components and the dashed grey line denotes the $|E_z|^2$ component.

the $|E_z|^2$ component has a lower peak amplitude than the annular components, leading to a beam that is overall annular and has a wider transverse extent than the LPLB.

The spot size of the RPLB was calculated by finding the radial distance to the point at which the intensity is $1/e^2$ of the peak intensity. This gave a spot size of $1.502w_0$ for the RPLB, where $w_0 = 0.51\,\mu m$ is the spot size of the focused LPLB, defined in the usual way for a Gaussian beam.

The Rayleigh range of the RPLB was calculated by finding the axial distance to the point where the intensity has fallen to half of the peak intensity. However, the Rayleigh range is considerably shorter for the RPLB, at $0.642z_R$, where z_R is the Rayleigh range of the focused LPLB, defined in the usual way for a Gaussian beam. This is ~ 3.5 times shorter than the Rayleigh range predicted when we account for the larger spot size at focus $z_R \propto (1.502w_0)^2 = 2.256w_0$.

6.4 Group Velocity of Different Polarisation States

In Section 2.1.5, we showed that the group velocity of a Laguerre-Gaussian mode with radial mode number p and azimuthal mode number $|\ell|$ is given by

$$v_g = \frac{c}{1 + \left(\frac{\theta_0}{4}\right)^2 (2p + |\ell| + 1)}, \tag{6.18}$$

where $\theta_0 = 4/w_0 k_0$ is the beam divergence half angle. For our linearly-polarised, lowest-order Gaussian beam, where $p = 0, |\ell| = 0$, the group velocity is therefore

$$v_{\mathrm{g,L}} = \frac{c}{1 + \left(\frac{\theta_0}{4}\right)^2}, \tag{6.19}$$

whereas for the radially-polarised beam (assuming construction from a sum of lowest order Laguerre-Gaussians with opposite vorticities and $p = 0, |\ell| = 1$)[237] the group velocity is

$$v_{\mathrm{g,R}} = \frac{c}{1 + 2\left(\frac{\theta_0}{4}\right)^2}. \tag{6.20}$$

Assuming the parameters given in Table 6.1, we can calculate the beam divergence half angle $\theta_0 = 4/w_0 k_0$. The spot size at the focus is estimated to be $w_0 = 0.51\,\mu\mathrm{m}$ (smaller than the wavelength, therefore ideally a full non-paraxial calculation is required), which gives an approximate beam divergence half angle of $\theta_0 = 57.3°$. The group velocity in free-space is therefore $v_{\mathrm{g,L}} = 0.941c$ for the LPLB and $v_{\mathrm{g,R}} = 0.889c$ for the RPLB. We will see that a derivation of the wake quantities requires an assumption about the temporal form that the fields take in the moving-frame coordinate, so we will account for this difference in group velocity in our equations for the fields (Eqs. 6.4–6.7 and 6.11–6.13) through the temporal profiles $f(t')$ of the LPLB and $g(t')$ of the RPLB, where these are defined in Section 6.2. Later, we shall see that this assumption is not our primary concern, since dephasing due to the difference in group velocity occurs on longer timescales compared to diffraction. Nonetheless, we shall bear it in mind when deriving the wake equation in Section 6.5.2.

6.5 Ionisation and Trapping in Laser-Driven Wakefields

6.5.1 Trapping with Linearly-Polarised Laser Beams

Starting from a 1D quasistatic approximation in the non-linear regime, the normalised scalar wake potential $\phi = e\Phi/m_e c^2$ can be derived for a relativistic cold plasma from the Lorentz equation, continuity equation and Poisson's equation.[30,43,44,295–298] In this case, the differential equation which describes the wakefield is

$$\frac{\partial^2 \phi}{\partial \xi^2} = k_{\mathrm{p}}^2 \gamma_{\mathrm{p}}^2 \left\{ \beta_{\mathrm{p}} \left[1 - \frac{1 + a_x^2}{\gamma_{\mathrm{p}}^2 (1+\phi)^2} \right]^{-1/2} - 1 \right\}, \tag{6.21}$$

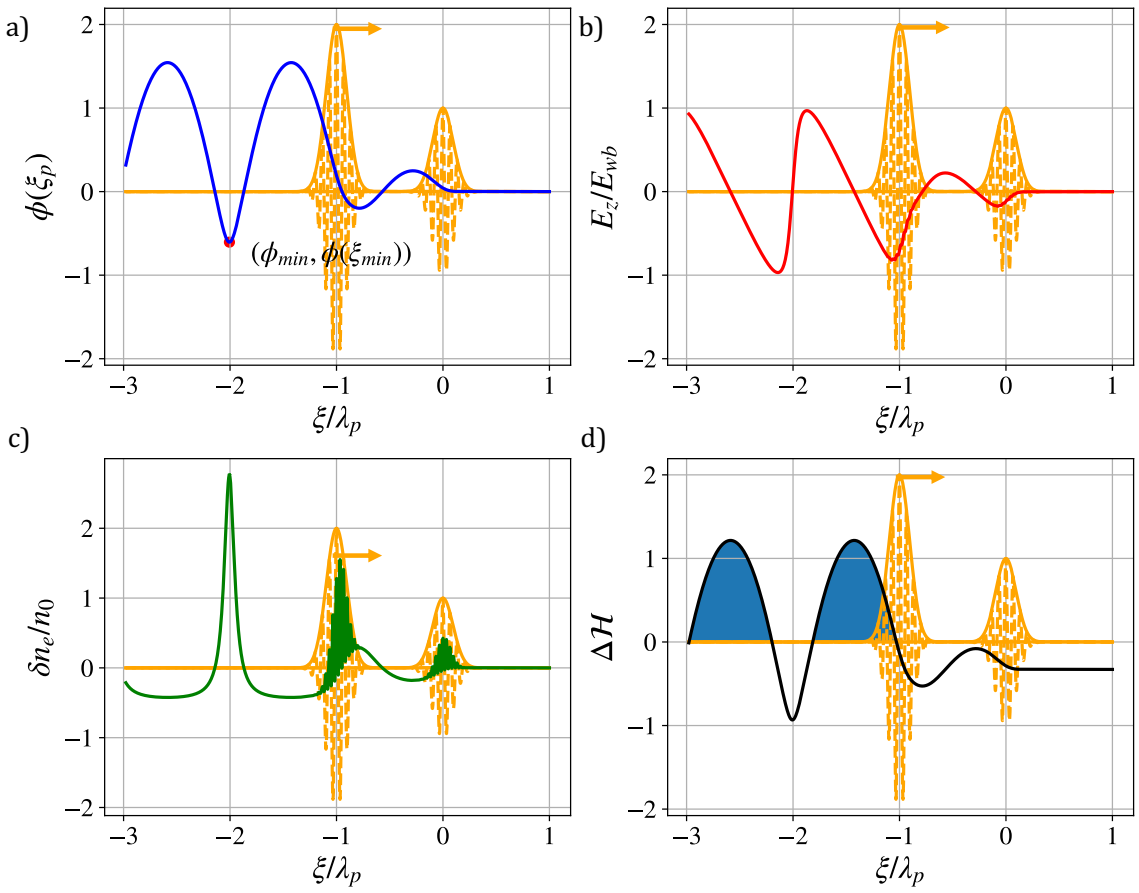


Figure 6.6: Calculated quasistatic wakefield quantities as a function of the moving-frame coordinate ξ/λ_p behind the leading pulse, assuming a 1D relativistic cold fluid in the non-linear regime.[294] Plotted for the case of a linearly-polarised drive pulse with $a_{0,1} = 1.0$ followed by a linearly-polarised injector pulse with $a_{0,2} = 2.0$, travelling to the right spaced by $\xi/\lambda_\mathrm{p} = 1$. The length of all pulses was $L_{\mathrm{p}0} = 1/k_p$. a) Normalised wake potential $\phi(\xi)$ with the minimum wake ϕ_{min} and trapping phase ξ_{min} indicated by a red dot, b) on-axis longitudinal component of the electric field $E_z(\xi)/E_{\mathrm{wb}}$, c) normalised electron density $\delta n_\mathrm{e}(\xi)/n_0 = n_\mathrm{e}/n_0 - 1$ and d) trapping parameter $\Delta\mathscr{H}(\xi) = \mathscr{H}_\mathrm{s} - \mathscr{H}_\mathrm{i}$, blue shading denotes regions where trapping may occur. Also shown are the normalised vector potential a_x (orange dashed line), and pump strength $|a_x|$ (orange solid line).

where $\xi = z - v_\mathrm{g} t$ is the normalised relativistic moving-frame coordinate, $k_\mathrm{p} = 2\pi/\omega_\mathrm{p}$ is the plasma angular wavenumber, $\gamma_\mathrm{p} = (1-\beta_\mathrm{p}^2)^{-1/2} = \omega_0/\omega_\mathrm{p}$ is the longitudinal momentum, $\gamma_x = 1 + u_x^2 = 1 + a_x^2$ is the transverse momentum, $\beta_\mathrm{p} = v_\mathrm{g}/c = [1-(\omega_\mathrm{p}/\omega_0)^2]^{1/2}$ is the normalised phase velocity of the plasma wave and v_g is the group velocity of the driving laser pulse (all derived in the previous section). All other quantities, including Eq. 6.21, will be derived for arbitrary polarisation states in Section 6.5.2. This second-order, non-linear, differential equation for the wake can be solved numerically to give the scalar wake potential ϕ, assuming an arbitrary laser beam profile a_x[299] which may include multiple linearly-polarised laser pulses.[146] We assume

that both laser pulses are Gaussian, such that the transverse vector potential a_x takes the form

$$a_x(\xi) = \frac{1}{2}a_{0,1}\exp\left\{-2\ln(2)\frac{\xi^2}{L_{\mathrm{p0}}^2}\right\}e^{ik_0\xi} + c.c. \\ + \frac{1}{2}a_{0,2}\exp\left\{-2\ln(2)\frac{\left[\xi-\Delta L\right]^2}{L_{\mathrm{p0}}^2}\right\}e^{ik_0\xi} + c.c., \tag{6.22}$$

where $c.c.$ denotes the complex conjugate, L_{p0} is the FWHM pulse length and ΔL is the pulse spacing. Solving Eq. 6.21 numerically using Python's Scipy `odeint` function[†], we obtain the normalised wake potential ϕ, which is plotted in Figure 6.6a). The longitudinal wakefield E_z/E_{wb}, normalised to the wave-breaking field $E_{\mathrm{wb}} = m_e c\omega_p/e$, and the electron density $\delta n_e/n_0$, normalised to the unperturbed on-axis electron density n_0, can be calculated from the wake potential as[43]

$$\frac{E_z}{E_{\mathrm{wb}}} = -\frac{1}{k_p}\frac{\partial\phi}{\partial\xi}, \tag{6.23}$$

$$\frac{\delta n_e}{n_0} = \frac{1}{k_p^2}\frac{\partial^2\phi}{\partial\xi^2} - 1, \tag{6.24}$$

also using Python's `odeint` function. These quantities are plotted in Figures 6.6b) and c), respectively. We note that "the electrostatic potential ϕ is predominantly slowly varying within the laser pulse even though $\delta n_e/n_0$ has rapidly varying components", as was previously observed by Ting *et al.*[300] This is a well-established theory, and our implementation of this model was cross-checked against work by other authors.[43,146,300–302]

The Hamiltonian determines the motion of electrons in the wakefield and is given by[294]

$$\mathscr{H}(u_z,\xi) = \sqrt{1+a_x(\xi)^2+u_z(\xi)^2} + \phi(\xi) - \beta_p u_z(\xi). \tag{6.25}$$

This can be solved for the trajectories $u_z(\xi)$, given by[43]

$$u_z(\xi) = \beta_p\gamma_p^2[\mathscr{H}+\phi(\xi)] \pm \gamma_p\sqrt{\gamma_p^2[\mathscr{H}+\phi(\xi)]^2-\gamma_x^2}, \tag{6.26}$$

where the conserved quantity $\mathscr{H}(u_z,\xi) = \mathscr{H}_0 = \text{const.}$ denotes distinct orbits. These orbits may be open or closed, with the separatrix demarcating the boundary between the two. Looking at

[†]https://docs.scipy.org/doc/scipy/reference/generated/scipy.integrate.odeint.html

the trajectories, the condition for a closed orbit is

$$\mathscr{H} < \frac{\gamma_x(\xi)}{\gamma_\mathrm{p}} - \phi(\xi). \tag{6.27}$$

Conversely, for a background electron on an open orbit, the condition on the Hamiltonian must satisfy

$$\mathscr{H} > \frac{\gamma_x(\xi)}{\gamma_\mathrm{p}} - \phi(\xi). \tag{6.28}$$

The separatrix is the boundary between the open and closed orbits, given by

$$\mathscr{H}_s = \frac{\gamma_x(\xi_\mathrm{min})}{\gamma_\mathrm{p}} - \phi(\xi_\mathrm{min}), \tag{6.29}$$

where we have noted that this condition will be met for a particular trapping phase ξ_min (plotted as a red dot in Figure 6.6). When an electron is born at rest it has no longitudinal or transverse momentum ($u_x, u_z = 0$). Therefore, the initial Hamiltonian of the ionised electron is given by

$$\mathscr{H}_i = 1 - \phi(\xi_i), \tag{6.30}$$

where ξ_i is the phase of the wakefield at the position where the electron under consideration is ionised. From Eq. 6.27 we can see that for the ionised electron to be trapped, we require $\mathscr{H}_i < \mathscr{H}_s$, so we can denote all regions where an ionised electron will be trapped using the trapping parameter $\Delta\mathscr{H} = \mathscr{H}_s - \mathscr{H}_i > 0$. This is plotted in Figure 6.6d), with regions where trapping may occur shaded in blue. An important thing to note is that regions of trapping occur after the peak of the laser pulse, where everything has already been ionised.[146] This is a fundamental problem for ionisation injection in the linear regime, and unless the region where it is possible to trap an electron can be overlapped with a region where the laser pulse intensity is high enough to ionise an unionised gas, no electron trapping can occur.

Chen *et al.* showed that trapping of electrons in a quasilinear wakefield driven by a single Gaussian pulse is impossible with a laser pulse with $a_0 \leq 1.7$,[299] but this restriction can be circumvented by introducing a second laser pulse, as shown in Figure 6.6.[146,147] The trailing pulse i) allows ionisation of species with ionisation potentials too high to be ionised by the leading drive pulse and ii) modifies the wake in the focal region, so that trapping can occur. The trailing pulse may also be chosen have a short Rayleigh range z_R, such that electrons are only ionised

and injected over a short axial region. In the following section, we consider what happens to the wakefield if this trailing pulse is radially-polarised.

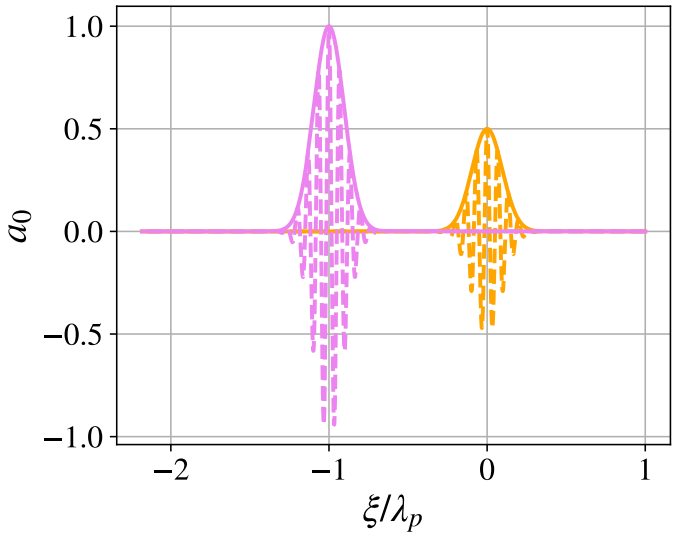


Figure 6.7: Schematic of the R-2PII scheme with a radially-polarised trailing pulse. Plotted for the case of a linearly-polarised drive pulse with $a_{0,1} = 0.5$ followed by a radially-polarised injector pulse with $a_{0,2} = 1$, travelling to the right spaced by $\xi/\lambda_p = 1$. Also shown are the normalised vector potential a_x (orange dashed line), and pump strength $|a_x|$ of the linearly-polarised drive beam and the normalised vector potential a_z (pink dashed line) and pump strength $|a_z|$ (pink solid line) of the radially-polarised injector beam.

6.5.2 Trapping with Radially-Polarised Laser Beams

In this particular example, we will consider a two-pulse scheme[146] where a linearly-polarised drive beam (orange) is trailed by a radially-polarised injector beam (pink). We shall call this the radial two pulse ionisation injection (R-2PII) scheme; as shown in Figure 6.7. After deriving the wake equation and Hamiltonian from first principles, allowing for a longitudinal component of the laser field, we will then solve the new wake equation numerically to give the wakefield quantities, as already described in Section 6.5.1. This allows us to investigate the effect of changing the polarisation state of the laser pulse on the wakefield. We can write Maxwell's equations, given by

$$\nabla \cdot \mathbf{E} = \frac{\rho}{\varepsilon_0}, \qquad \nabla \times \mathbf{E} = -\frac{\partial \mathbf{B}}{\partial t}, \tag{6.31}$$

$$\nabla \cdot \mathbf{B} = 0, \qquad \nabla \times \mathbf{B} = \mu_0 \varepsilon_0 \frac{\partial \mathbf{E}}{\partial t} + \mu_0 \mathbf{j}, \tag{6.32}$$

in terms of the scalar potential Φ and vector potential $\mathbf{A}$ as

$$\mathbf{E} = -\nabla\Phi - \frac{\partial \mathbf{A}}{\partial t}, \qquad \mathbf{B} = \boldsymbol{\nabla} \times \mathbf{A}. \tag{6.33}$$

The Lorentz equation

$$\frac{\mathrm{d}\mathbf{p}}{\mathrm{d}t} = -e\Big(\mathbf{E} + \mathbf{v} \times \mathbf{B}\Big), \tag{6.34}$$

is written in Cartesian coordinates as

$$\frac{\mathrm{d}\mathbf{p}}{\mathrm{d}t} = -e\left[\begin{pmatrix} E_x \\ E_y \\ E_z \end{pmatrix} + \begin{pmatrix} v_y B_z - v_z B_y \\ v_z B_x - v_x B_z \\ v_x B_y - v_y B_x \end{pmatrix}\right]. \tag{6.35}$$

We can include only the field components that were defined by Eqs. 6.4–6.7 and 6.11–6.13. We showed that on-axis, only the E_x and B_y components of the LPLB and the E_z component of the RPLB are appreciable. This can be seen in Figures 6.1 and 6.2, respectively. In this case, the Lorentz equation becomes

$$\frac{\mathrm{d}\mathbf{p}}{\mathrm{d}t} = -e\left[\begin{pmatrix} E_x \\ 0 \\ E_z \end{pmatrix} + \begin{pmatrix} -v_z B_y \\ 0 \\ v_x B_y \end{pmatrix}\right]. \tag{6.36}$$

From this, the coupled equations for the momenta are

$$\frac{\mathrm{d}p_x}{\mathrm{d}t} = -e\Big(E_x - v_z B_y\Big), \tag{6.37}$$

$$\frac{\mathrm{d}p_z}{\mathrm{d}t} = -e\Big(E_z + v_x B_y\Big). \tag{6.38}$$

Using Maxwell's equations (Eq. 6.33), we can rewrite the electric and magnetic fields in terms of the vector and scalar potentials to give

$$E_x = -\frac{\partial A_x}{\partial t}, \tag{6.39}$$

$$E_z = -\frac{\partial \Phi}{\partial z} - \frac{\partial A_z}{\partial t}, \tag{6.40}$$

$$B_y = \frac{\partial A_x}{\partial z}, \tag{6.41}$$

where $\partial\Phi/\partial x = 0$ and $\partial A_z/\partial x = 0$, since we consider only the on-axis fields. Substituting these into the Lorentz equation, we obtain

$$\frac{\mathrm{d}p_x}{\mathrm{d}t} = e\left(\frac{\partial A_x}{\partial t} + v_z\frac{\partial A_x}{\partial z}\right), \tag{6.42}$$

$$\frac{\mathrm{d}p_z}{\mathrm{d}t} = e\left(\frac{\partial \Phi}{\partial z} + \frac{\partial A_z}{\partial t} - v_x\frac{\partial A_x}{\partial z}\right), \tag{6.43}$$

where the first component integrates directly for

$$p_x = \gamma m_e v_x = eA_x. \tag{6.44}$$

Following a similar approach to Sprangle *et al.*,[300,301,303] we will now derive the non-linear differential equation that describes the behaviour of the scalar potential of the wake. Firstly, we normalise the vector $a_0 = eA_0/mc$ and scalar potentials $\phi = e\Phi/mc^2$ (using the same conventions in SI units as previous authors[43]). Using the normalised quantities, the transverse component of the Lorentz equation (Eq. 6.44) is

$$u_x = p_x/m_e c = eA_x/mc = a_x, \tag{6.45}$$

and the longitudinal component of the Lorentz equation (Eq. 6.43) is

$$\frac{\mathrm{d}u_z}{\mathrm{d}t} = c\frac{\partial \phi}{\partial z} + \frac{\partial a_z}{\partial t} - \frac{c}{2\gamma}\frac{\partial a_x^2}{\partial z}. \tag{6.46}$$

The Moving-Frame Coordinate

We now want to make a change of variables to a reference frame which is co-moving with the linearly-polarised drive beam. Our moving-frame coordinates (ξ, τ) and the derivatives are related to the lab-frame coordinates (z, t) according to

$$\tau = t, \qquad \frac{\partial}{\partial t} = \frac{\partial}{\partial \tau} - v_{\mathrm{g,L}}\frac{\partial}{\partial \xi}, \tag{6.47}$$

$$\xi = z - v_{\mathrm{g,L}}t, \qquad \frac{\partial}{\partial z} = \frac{\partial}{\partial \xi}. \tag{6.48}$$

Taking a change of variables in Eq. 6.46, we obtain

$$\frac{\mathrm{d}u_z}{\mathrm{d}t} = \frac{\partial u_z}{\partial t} + v_z \frac{\partial u_z}{\partial z}, \tag{6.49}$$

$$= \frac{\partial u_z}{\partial \tau} - v_{\mathrm{g,L}} \frac{\partial u_z}{\partial \xi} + v_z \frac{\partial u_z}{\partial \xi}. \tag{6.50}$$

The right-hand side of Eq. 6.46 can be written in terms of the new coordinates as

$$c\frac{\partial \phi}{\partial \xi} + \frac{\partial a_z}{\partial \tau} - v_{\mathrm{g,L}} \frac{\partial a_z}{\partial \xi} - \frac{c}{2\gamma}\frac{\partial a_x^2}{\partial \xi}. \tag{6.51}$$

We can split the relativistic factor γ into longitudinal and transverse components using $\beta_x = v_x/c = u_x/\gamma = a_x/\gamma$ by rearranging

$$\gamma = (1-\beta_x^2-\beta_z^2)^{-1/2} = \left(1 - \frac{a_x^2}{\gamma^2} - \beta_z^2\right)^{-1/2}, \tag{6.52}$$

for

$$\gamma^2\left(1 - \frac{a_x^2}{\gamma^2} - \beta_z^2\right) = 1 \qquad \rightarrow \qquad \gamma = \gamma_x\gamma_z = \frac{\sqrt{1+a_x^2}}{\sqrt{1-\beta_z^2}}, \tag{6.53}$$

such that

$$a_x^2 = \gamma^2(1-\beta_z^2) - 1, \tag{6.54}$$

and

$$\frac{\partial a_x^2}{\partial \xi} = \frac{\partial}{\partial \xi}\left[\gamma^2\left(1 - \frac{u_z^2}{\gamma^2}\right) - 1\right], \tag{6.55}$$

$$= \frac{\partial}{\partial \xi}\left[\gamma^2 - u_z^2 - 1\right], \tag{6.56}$$

$$= 2\gamma\frac{\partial \gamma}{\partial \xi} - 2u_z\frac{\partial u_z}{\partial \xi}. \tag{6.57}$$

We can put this all together to rewrite Eq. 6.46 as

$$\frac{1}{c}\frac{\partial}{\partial \tau}\left(u_z - a_z\right) = \frac{\partial}{\partial \xi}\left[\phi - \gamma + \beta_{\mathrm{p}}\left(u_z - a_z\right)\right], \tag{6.58}$$

where $\beta_{\mathrm{p}} = v_{\mathrm{g,L}}/c$ is the normalised phase velocity of the wakefield driven by the leading LPLB. It should be noted that Eq. 6.58 may look different to other texts due to different conventions

and normalisation[†].

In the moving-frame coordinates of the LPLB, the 1D mass continuity fluid equation is

$$\frac{1}{c}\frac{\partial n_\mathrm{e}}{\partial \tau} = \frac{\partial}{\partial \xi}\Big[n_\mathrm{e}\big(\beta_\mathrm{p} - \beta_z\big)\Big], \tag{6.59}$$

where the electron density is normalised to the unperturbed electron density n_0 throughout. Similarly, Poisson's equation is

$$\frac{\partial^2 \phi}{\partial \xi^2} = k_\mathrm{p}^2(n_\mathrm{e} - 1). \tag{6.60}$$

This gives us a system of three partial differential equations which describe the evolution of the plasma.

The Quasistatic Approximation

Until this point in the derivation, we have discussed the equations that govern the behaviour of the plasma (namely, the cold relativistic fluid equations consisting of the Lorentz equation, the mass continuity equation and Poisson's equation), but we have not considered how our laser pulses will evolve as they propagate through the plasma. This is described by the wave equation, which in the moving-frame coordinates is given by[301]

$$\left(\frac{1}{\gamma_\mathrm{p}^2}\frac{\partial^2}{\partial \xi^2} + \frac{2\beta_\mathrm{p}}{c}\frac{\partial^2}{\partial \xi \partial \tau} - \frac{1}{c^2}\frac{\partial^2}{\partial \tau^2}\right)\mathbf{a} \simeq \frac{k_\mathrm{p}^2}{\gamma}\frac{n_\mathrm{e}}{n_0}\mathbf{a}. \tag{6.61}$$

We could solve Eqs. 6.58–6.61 iteratively, to study the evolution of both the laser pulses and the plasma. However, an alternative approach is to make a quasistatic approximation, which assumes that the laser pulses do not evolve in the time that it takes the plasma to pass through the laser pulse. This allows us to assume a particular fixed solution to Eq. 6.61, where the laser pulses are static and non-evolving in the moving-frame coordinate. In this case, the plasma quantities do not change with respect to τ and we can neglect the partial derivatives $\partial u_z/\partial \tau = 0$ and $\partial n_\mathrm{e}/\partial \tau = 0$. Eqs. 6.58–6.60 then reduce to a single non-linear second-order differential equation, which can be solved to give the quasistatic wakefield quantities in the moving-frame coordinate, as described in Section 6.5.1.

We will now consider all relevant timescales of the system, to allow us to understand whether a quasistatic approximation is valid in this case. There are three aspects that we need to consider:

[†]In particular, certain texts, including Gibbon,[295] use CGS units.

1. The time taken for the two laser pulses to CEP phase slip τ_{CEP}, due to the difference between the phase and group velocities,

2. The time taken for the two laser pulses to dephase with each other τ_{de}, due to the difference in group velocities between the LPLB and RPLB,

3. The time taken for the laser pulses to diffract τ_{d}.

If these timescales are long compared to the duration of the laser pulse envelope τ_{p0} then we can make a quasistatic approximation. Throughout, we shall assume an unperturbed plasma density of $n_0 = 1 \times 10^{17}\,\mathrm{cm}^{-3}$, a laser pulse central wavelength of $\lambda_0 = 800\,\mathrm{nm}$ and a pulse duration of $\tau_{\mathrm{p0}} = 45\,\mathrm{fs}$. From Eq. 6.61, the characteristic timescale over which the pulse envelope is modified by the plasma τ_{e} is given by[300]

$$\tau_{\mathrm{e}} \sim 2\gamma \left| \frac{n_0}{n_{\mathrm{e}}} \right| \frac{\omega_0}{\omega_{\mathrm{p}}^2}. \tag{6.62}$$

We can note that the evolution time of the laser pulse envelopes is three orders of magnitude larger than the pulse duration $\tau_{\mathrm{e}} \sim 1000\tau_{\mathrm{p0}}$, for the assumed parameters above. So certainly we can neglect the modification of the envelope due to the wakefield.

1: Carrier-Envelope Phase Slip A possible solution to Eq. 6.61 is that the fields of the laser pulse take the well-behaved form of a slowly-varying envelope multiplied by a carrier wave (see Section 2.1.2).[43,44,295,304] Any pulse can be described as an envelope and a carrier wave, as was assumed in Sections 6.2.1 and 6.2.2. By assuming that the pulse envelopes are Gaussian and transforming the fields in the lab-frame to the moving-frame coordinate, the vector potentials a_x and a_z of the pulses are

$$a_x(\xi, \xi^0) = \frac{1}{2} a_{0,x} \exp\left\{ -2\ln(2) \frac{\xi^2}{L_{\mathrm{p0}}^2} \right\} \mathrm{e}^{ik_0\xi^0} + c.c., \tag{6.63}$$

$$a_z(\xi, \xi^0, \tau) = \frac{1}{2} a_{0,z} \exp\left\{ -2\ln(2) \frac{\left[\xi - \Delta v_{\mathrm{g}}\tau - \Delta L\right]^2}{L_{\mathrm{p0}}^2} \right\} \mathrm{e}^{ik_0\xi^0} + c.c., \tag{6.64}$$

where $\Delta v_{\mathrm{g}} = v_{\mathrm{g,L}} - v_{\mathrm{g,R}}$ is the difference in group velocities between the LPLB and RPLB. We will assume that both pulses have the same wavelength and therefore the same carrier wave vector k_0. The way that we have written the pulse envelopes can therefore account for the difference

in group velocity of the two laser pulses Δv_g. It should be noted that

$$\xi^0 = z - \frac{\omega_0}{k_0} t, \tag{6.65}$$

$$= z - v_{ph}\tau, \tag{6.66}$$

where v_{ph} is the phase velocity of light in a plasma. The refractive index of a plasma is given by[43]

$$\eta = \sqrt{1 - \frac{\omega_p^2}{\omega_0^2}} < 1. \tag{6.67}$$

A plane wave travelling through the plasma will have a phase velocity of

$$v_{ph} = \frac{\omega_0}{k_0} = \frac{c}{\eta}, \tag{6.68}$$

and a group velocity of

$$v_g \equiv \frac{\partial \omega_0}{\partial k_0} = \eta c. \tag{6.69}$$

The time taken for the carrier-envelope phase (CEP) to slip so that a point of constant phase moves through the entire pulse is

$$\tau_{CEP} = \frac{c\tau_{p0}}{(v_{ph} - v_{g,L})}. \tag{6.70}$$

Most texts assume that $v_g \approx v_p$, which is often not explicitly stated but is instead implied in assuming a laser pulse of the form $\hat{a}e^{ik_0\xi}$, where ξ is the usual moving-frame coordinate, moving at the group velocity of the pulse.[295] We can relate these two frames of reference for the LPLB according to

$$\xi = z - v_{g,L}t, \tag{6.71}$$

$$= z - v_{ph}t + (v_{ph} - v_{g,L})t, \tag{6.72}$$

$$= \xi^0 + \Delta v_{p,L}\tau, \tag{6.73}$$

where $\Delta v_{p,L} = v_{ph} - v_{g,L} > 0$, since in a plasma the refractive index $\eta < 1$. For the assumed parameters in Table 6.1, $\Delta v_{p,L} = 0.059c$, 6% of the speed of light. If we assume that our pulses are monochromatic plane waves, then in a plasma density of $n_0 = 1 \times 10^{17}\,\mathrm{cm}^{-3}$, the time taken for the pulse to phase slip is $\tau_{CEP} \sim 1.74 \times 10^4 \tau_{p0}$. Therefore, the assumption that $v_p \approx v_g$ is cer-

tainly true if our pulses are plane waves, but in this chapter so far, we have considered spatially-structured light, which generally has a subluminal group velocity, even in vacuum. Assuming that both pulses have the same central wavelength λ_0 and pulse duration τ_{p0}, then the time that it takes for a point of constant phase to move through the entire pulse in the moving-frame ξ is $\tau_{CEP} \sim 16.9\tau_{p0}$ (where we have assumed a group velocity of $v_{g,L} = 0.941c$, see Section 6.4). This means that a point of constant phase will move approximately a seventeenth of the way through the envelope in the time that it takes a plasma electron to move through the pulse.

2: TWO PULSE DEPHASING As mentioned in Section 6.4, the trailing RPLB has a group velocity $v_{g,R} = 0.945v_{g,L}$. The time taken for the RPLB to dephase with the LPLB is given by

$$\tau_{de} = \frac{\lambda_p}{4(v_{g,L} - v_{g,R})}, \tag{6.74}$$

which for the assumed parameters and the group velocities derived in Section 6.4 is $\tau_{de} \sim 37.6\tau_{p0}$. We can actually take the difference in group velocities into account in Eq. 6.58. In the moving-frame coordinate of the RPLB $\xi' = z - v_{g,R}t$. We can relate these two frames of reference according to

$$\xi = z - v_{g,L}t, \tag{6.75}$$

$$= z - v_{g,R}t - (v_{g,L} - v_{g,R})t, \tag{6.76}$$

$$= \xi' - \Delta v_g \tau, \tag{6.77}$$

where we have noted that the coordinate τ is the same in both frames. From Section 6.4, we know that $\Delta v_g = v_{g,L} - v_{g,R} > 0$, since the LPLB always travels faster than the RPLB. For the assumed parameters in Table 6.1, $\Delta v_g = 0.052c$, 5% of the speed of light. In the moving-frame of the LPLB, infinitesimal variation of a_z (Eq. 6.64) is given by

$$da_z = \left(\frac{\partial a_z}{\partial \tau}\right)_{\xi',\xi^0} d\tau + \left(\frac{\partial a_z}{\partial \xi'}\right)_{\tau,\xi^0} d\xi' + \left(\frac{\partial a_z}{\partial \xi^0}\right)_{\tau,\xi'} d\xi^0, \tag{6.78}$$

such that

$$\left(\frac{\partial a_z}{\partial \tau}\right)_{\xi} = \left(\frac{\partial a_z}{\partial \tau}\right)_{\xi',\xi^0}\left(\frac{\partial \tau}{\partial \tau}\right)_{\xi} + \left(\frac{\partial a_z}{\partial \xi'}\right)_{\tau,\xi^0}\left(\frac{\partial \xi'}{\partial \tau}\right)_{\xi} + \left(\frac{\partial a_z}{\partial \xi^0}\right)_{\tau,\xi'}\left(\frac{\partial \xi^0}{\partial \tau}\right)_{\xi}, \tag{6.79}$$

$$= \left(\frac{\partial a_z}{\partial \xi}\right)_{\tau,\xi^0}\left(\frac{\partial \xi}{\partial \xi'}\right)_{\tau,\xi^0}\left(\frac{\partial \xi'}{\partial \tau}\right)_{\xi} + \left(\frac{\partial a_z}{\partial \xi}\right)_{\tau,\xi'}\left(\frac{\partial \xi}{\partial \xi^0}\right)_{\tau,\xi'}\left(\frac{\partial \xi^0}{\partial \tau}\right)_{\xi}. \tag{6.80}$$

We can note that

$$\xi' = \xi + \Delta v_g \tau, \tag{6.81}$$

$$\xi^0 = \xi - \Delta v_{p,L} \tau. \tag{6.82}$$

Assuming that v_{ph}, $v_{g,L}$ and $v_{g,R}$ are all constant, then

$$\left(\frac{\partial a_z}{\partial \tau}\right)_{\xi} = \Delta v_g \left(\frac{\partial a_z}{\partial \xi}\right)_{\tau} - \Delta v_{p,L}\left(\frac{\partial a_z}{\partial \xi}\right)_{\tau}, \tag{6.83}$$

$$= (\Delta v_g - \Delta v_{p,L})\left(\frac{\partial a_z}{\partial \xi}\right)_{\tau}. \tag{6.84}$$

Taking this into account, our equation for the wakefield potential (Eq. 6.58) becomes

$$\frac{1}{c}\frac{\partial u_z}{\partial \tau} = \frac{\partial}{\partial \xi}\Big[\phi - a_z \Delta\beta - \gamma + \beta_p u_z\Big], \tag{6.85}$$

where

$$\Delta\beta = \frac{1}{c}(\Delta v_{p,L} - \Delta v_g + v_{g,L}), \tag{6.86}$$

$$= \frac{1}{c}\Big[(v_{ph} - v_{g,L}) - (v_{g,L} - v_{g,R}) + v_{g,L}\Big], \tag{6.87}$$

$$= \beta_{ph} + \beta_{g,R} - \beta_{g,L}, \tag{6.88}$$

$$\approx \beta_p, \tag{6.89}$$

which can also be noted from the fact that $\Delta v_{p,L} \approx \Delta v_g \approx 0.05c$. Nonetheless, we shall use $\Delta\beta$ to fully account for the difference between the phase and group velocities in what remains of our derivation. Eq. 6.85 agrees with Esarey *et al.*,[301] where they state that $\psi = \phi - a_z$ is the wake potential, but here we have also accounted for a difference in the phase velocity and group velocity of the two pulses.

3: DIFFRACTION The timescale for the diffraction of a focused laser pulse can be estimated to be[300,301,303]

$$\tau_{\mathrm{d}} \sim \frac{z_{\mathrm{R}}}{c} = \frac{\pi w_0^2}{\lambda_0 c}, \tag{6.90}$$

for the LPLB. Assuming the laser pulse parameters given in Table 6.1, the diffraction time of the LPLB is on the order of $\tau_{\mathrm{d}} \sim 0.073\tau_{\mathrm{p0}}$. As we saw in Section 6.3, the Rayleigh range of the RPLB was 0.642 times the Rayleigh range of the LPLB. The estimated diffraction time of the RPLB is therefore $\tau_{\mathrm{d}} \sim 0.047\tau_{\mathrm{p0}}$.

To summarise, the relevant timescales for the evolution of the laser pulse in an $f/2$ focusing geometry are given by:

Pulse duration, τ_{p0}	45 fs
Envelope evolution time, τ_e	$1000\tau_{p0}$
Dephasing time, τ_{de}	$37.6\tau_{p0}$
Carrier-envelope phase slip time, τ_{CEP}	$16.9\tau_{p0}$
Diffraction time, τ_d	$0.05\tau_{p0}$

Table 6.2: Table of relevant timescales for a quasistatic approximation

Clearly, without a mechanism to guide the laser pulse, we are not able to make a quasistatic approximation in an $f/2$ focusing geometry, since the laser pulse will diffract on a much shorter timescale than it takes the plasma to traverse it. In this case, we need other methods to fully simulate the interaction of a plasma and a tightly-focused laser beam. To study the effect of the RPLB in a quasistatic approximation, one option is to relax the focusing geometry. In the case where the spot size at the focus is greater than the plasma wavelength, $w_0 \gg \lambda_p$, then both laser pulses can be approximated as plane waves with group velocities $v_g \approx v_p$. In this case, diffraction, carrier-envelope phase slip and dephasing are all negligible. Since the point of this study is to ascertain in the first instance whether a longitudinal component of the laser field has any impact on the wakefield, we shall fix the time to be $\tau = 0$, assuming that the laser pulses do not evolve to study the quasistatic wakefield quantities in this case. This allows us to assess whether this scheme is worth investigating further.

Quasistatic Approximation: Plasma Quantities As already mentioned, in the quasistatic approximation, the quantities u_z and n_e vary slowly, such that $\partial u_z/\partial\tau = 0$ and $\partial n_e/\partial\tau = 0$ in the fluid equations Eqs. 6.85 and 6.59, respectively. Setting $\partial u_z/\partial\tau = 0$ in Eq. 6.85 we obtain

$$\frac{\partial}{\partial\xi}\Big[\phi - a_z\Delta\beta - \gamma + \beta_p u_z\Big] = 0. \tag{6.91}$$

In the usual manner, we integrate to obtain the conservation relations

$$\phi - a_z\Delta\beta - \gamma + \beta_p u_z + 1 = 0, \tag{6.92}$$

$$n_e = \frac{\beta_p}{\beta_p - \beta_z}, \tag{6.93}$$

where we assume that away from the region of the interaction where $\xi \to \infty$ the boundary conditions are given by

$$\phi = 0, \qquad u_z = 0, \qquad a_z = 0, \qquad \gamma = 1, \qquad n = 1. \tag{6.94}$$

This has reduced our set of three partial differential equations to one non-linear ordinary differential equation. We will now express all fluid quantities in terms of the laser pulse potentials a_x and a_z. Using $u_z = \gamma\beta_z$ and Eq. 6.53 in Eq. 6.92 we can eliminate γ

$$(\phi + 1) - a_z\Delta\beta - \frac{\sqrt{1+a_x^2}}{\sqrt{1-\beta_z^2}}(1 - \beta_p\beta_z) = 0, \tag{6.95}$$

$$\to \qquad \frac{(1-\beta_p\beta_z)^2}{(1-\beta_z^2)} = \frac{\left[\phi + 1 - a_z\Delta\beta\right]^2}{(1+a_x^2)}. \tag{6.96}$$

This gives an equation that can be solved for β_z in terms of β_p, a_x, a_z and ϕ[†]. The solution β_z is

$$\beta_z = \frac{\beta_p(1+a_x^2) - \sqrt{\left(\phi + 1 - a_z\Delta\beta\right)^2\left\{\beta_p^2 + a_z^2\Delta\beta^2 + a_x^2(-1+\beta_p^2) + 2\phi + \phi^2 - 2a_z\Delta\beta(1+\phi)\right\}}}{\beta_p^2(1+a_x^2) + (\phi + 1 - a_z\Delta\beta)^2}, \tag{6.97}$$

$$= \frac{\beta_p(1+a_x^2) - \sqrt{\left(\phi + 1 - a_z\Delta\beta\right)^2\left\{\left(\phi + 1 - a_z\Delta\beta\right)^2 + \beta_p^2(1+a_x^2) - (1+a_x^2)\right\}}}{\beta_p^2(1+a_x^2) + (\phi + 1 - a_z\Delta\beta)^2}, \tag{6.98}$$

$$= \frac{\beta_p\gamma_x^2 - \Omega^2\sqrt{1 - \gamma_x^2/(\gamma_p^2\Omega^2)}}{\beta_p^2\gamma_x^2 + \Omega^2}, \tag{6.99}$$

$$= \frac{\beta_p\gamma_x^2 - \Omega^2\Psi}{\beta_p^2\gamma_x^2 + \Omega^2}, \tag{6.100}$$

where we have noted that $\gamma_p = (1-\beta_p^2)^{-1/2}$ and

$$\Omega = \phi - a_z\Delta\beta + 1, \tag{6.101}$$

$$\Psi = \sqrt{1 - \gamma_x^2/(\gamma_p^2\Omega^2)}. \tag{6.102}$$

[†]This was solved using Mathematica to avoid tedious manipulation and was cross-checked by inputting both Eqs. 6.95 and 6.96. The result was then factorised and tidied by hand.

Substituting the mass continuity equation (Eq. 6.93) into Poisson's equation (Eq. 6.60) gives

$$\frac{\partial^2\phi}{\partial\xi^2} = k_\mathrm{p}^2(n_\mathrm{e}-1), \tag{6.103}$$

$$= k_\mathrm{p}^2\left(\frac{\beta_\mathrm{p}}{\beta_\mathrm{p}-\beta_z}-1\right), \tag{6.104}$$

and substituting the explicit form of β_z given by Eq. 6.100, we obtain

$$\frac{\partial^2\phi}{\partial\xi^2} = k_\mathrm{p}^2\left(\frac{\beta_\mathrm{p}}{\beta_\mathrm{p}-\frac{\beta_\mathrm{p}\gamma_x^2-\Omega^2\Psi}{\beta_\mathrm{p}^2\gamma_x^2+\Omega^2}}-1\right), \tag{6.105}$$

$$= k_\mathrm{p}^2\left(\frac{\beta_\mathrm{p}(\beta_\mathrm{p}^2\gamma_x^2+\Omega^2)-\beta_\mathrm{p}(\beta_\mathrm{p}^2\gamma_x^2+\Omega^2)+\beta_\mathrm{p}\gamma_x^2-\Omega^2\Psi}{\beta_\mathrm{p}(\beta_\mathrm{p}^2\gamma_x^2+\Omega^2)-\beta_\mathrm{p}\gamma_x^2+\Omega^2\Psi}\right), \tag{6.106}$$

$$= k_\mathrm{p}\left(\frac{\beta_\mathrm{p}\gamma_x^2-\Omega^2\Psi}{\Omega^2\Psi+\beta_\mathrm{p}\Omega^2\Psi^2}\right). \tag{6.107}$$

This gives the final form of the differential equation which describes the wakefield in the case where we also have a longitudinal component of the laser field a_z. In the limit that $a_z \to 0$, then $\Omega \to (\phi+1)$, and we recover the original form of Poisson's equation as derived in previous texts,[297] given by Eq. 6.21 in Section 6.5.1. Eq. 6.107 may equivalently be written as

$$\frac{\partial^2\phi}{\partial\xi^2} = k_\mathrm{p}^2\gamma_\mathrm{p}^2\left\{\beta_\mathrm{p}\left[1-\frac{1+a_x^2}{\gamma_\mathrm{p}^2\Omega^2}\right]^{-1/2}-1\right\}. \tag{6.108}$$

The Hamiltonian

The final step is to derive the Hamiltonian in the case of arbitrary laser pulse polarisation. The action of a relativistic electron in electromagnetic fields is Lorentz invariant and the sum of the action of a free particle and the interaction between a charge and the electromagnetic fields. It is given by[157]

$$S = \int_t^{t+\delta t}\mathscr{L}\,\mathrm{d}t = \int_t^{t+\delta t}\left(-m_e c^2\sqrt{1-\frac{v^2}{c^2}}-q\Phi+q\mathbf{A}\cdot\mathbf{v}\right)\mathrm{d}t, \tag{6.109}$$

such that the Lagrangian for an electron is

$$\mathscr{L} = -m_e c^2\sqrt{1-\frac{v^2}{c^2}}+e\Phi-e\mathbf{A}\cdot\mathbf{v}. \tag{6.110}$$

The canonical momentum is given by

$$\mathbf{P} = \frac{\partial \mathscr{L}}{\partial \mathbf{v}} = \frac{m_e \mathbf{v}}{\sqrt{1-\frac{v^2}{c^2}}} - e\mathbf{A} = \mathbf{p} - e\mathbf{A}. \tag{6.111}$$

If the fields are constant, as is the case in a quasistatic approximation, then the Lagrangian does not depend on time and the energy is conserved and coincides with the Hamiltonian.[157] Generally, the Hamiltonian is the Legendre transformation of the Lagrangian and is given in terms of the canonical momentum by

$$\mathscr{H} = \mathbf{v}\cdot\mathbf{P} - \mathscr{L}, \tag{6.112}$$

$$= c\sqrt{m_e c^2 + (\mathbf{P}+e\mathbf{A})^2} - e\Phi, \tag{6.113}$$

$$= \sqrt{1+(\mathbf{U}+\mathbf{a})^2} - \phi, \tag{6.114}$$

where we have used the normalised units given in Section 6.5.2, dividing through by mc^2. This agrees with previous work by Esirkepov *et al.*,[302] but unlike their work it does not assume transverse polarisation of the laser pulse. We now need to make a canonical transformation from the coordinates (z, u_z) to the moving-frame coordinates (ξ, u_z). This is the same transformation as was performed in Section 6.5.2. Generally, a transformation $(q,p) \to (p',q')$ is canonical, where q are the generalised coordinates and p are the generalised canonical momenta, if there exists a Hamiltonian $\tilde{\mathscr{H}}$ such that

$$\dot{q}' = [q', \tilde{\mathscr{H}}], \qquad \dot{p}' = [p', \tilde{\mathscr{H}}], \tag{6.115}$$

gives the same trajectory as

$$\dot{q} = [q, \mathscr{H}], \qquad \dot{p} = [p, \mathscr{H}], \tag{6.116}$$

where only in this instance (Eqs. 6.115 and 6.116) do the square brackets represent Poisson brackets. Choosing a generating function F which satisfies

$$\frac{\partial F}{\partial q} = p, \qquad \frac{\partial F}{\partial p'} = q', \qquad \tilde{\mathscr{H}} = \mathscr{H} + \frac{1}{c}\frac{\partial F}{\partial t} \tag{6.117}$$

ensures that Eqs. 6.115 and 6.116 are satisfied.[157] Using the coordinates given by Eqs. 6.47–6.48, we can try a generating function F of the form[43]

$$F = u_z(z - v_{\mathrm{g,L}}t), \tag{6.118}$$

and we find that this satisfies the requirements given by Eq. 6.117 for a canonical transformation. Under this transformation, the Hamiltonian is

$$\tilde{\mathscr{H}}(\xi, u_z) = \sqrt{1+\mathbf{u}^2} - \phi - \beta_g u_z, \tag{6.119}$$

$$= \sqrt{1+u_x^2+u_z^2} - \phi - \beta_g u_z, \tag{6.120}$$

$$= \sqrt{1+a_x(\xi)^2+u_z(\xi)^2} - \phi(\xi) - \beta_g u_z(\xi) \tag{6.121}$$

where we have again noted that $\beta_\mathrm{p} = v_{\mathrm{g,L}}/c$ is the normalised phase velocity of the wakefield driven by the LPLB and that $u_x = a_x$ from Eq. 6.45. *This is exactly the same as the Hamiltonian given by Eq. 6.25.* Therefore, the definition of the Hamiltonian does not change depending on whether we have a linearly- or radially-polarised laser pulse, and the trapping parameter remains as given in Section 6.5.1

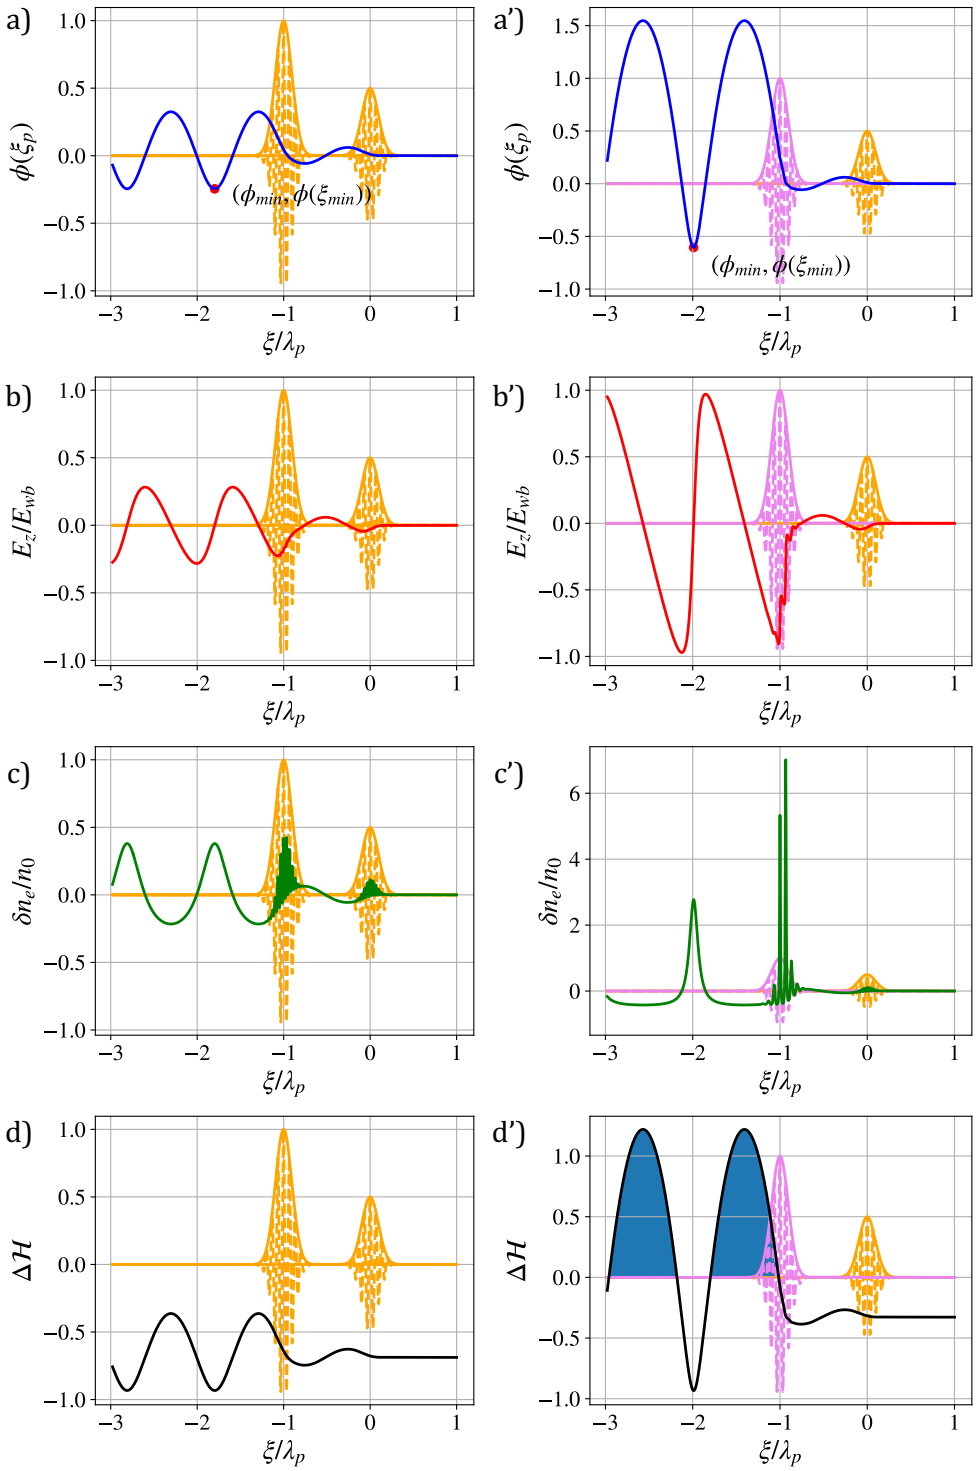


Figure 6.8: Calculated quasistatic wakefield quantities as a function of the moving-frame coordinate ξ/λ_p behind the leading pulse, assuming a 1D relativistic cold fluid in the non-linear regime.[294] Plotted for the case of a linearly-polarised drive pulse with $a_{0,1} = 0.5$ followed by a radially-polarised injector pulse with $a_{0,2} = 1$, travelling to the right spaced by $\xi/\lambda_\mathrm{p} = 1$. The length of all pulses was $L_{\mathrm{p}0} = 1/k_p$. a)–d) Show results for a conventional 2PII scheme with two linearly-polarised beams (orange). a) Normalised wake potential $\phi(\xi)$ with the minimum wake ϕ_min and trapping phase ξ_min indicated by a red dot, b) on-axis longitudinal component of the electric field $E_z(\xi)/E_\mathrm{wb}$, c) normalised electron density $\delta n_\mathrm{e}(\xi)/n_0 = n_\mathrm{e}/n_0 - 1$ and d) trapping parameter $\Delta\mathscr{H}(\xi) = \mathscr{H}_\mathrm{s} - \mathscr{H}_\mathrm{i}$, blue shading denotes regions where trapping may occur. a′)–d′) Show the same results for the novel 2PII scheme with a linearly-polarised drive pulse (orange), followed by a radially-polarised pulse (pink).

6.6 Results

In exactly the same manner as in Section 6.5.1, we can solve Eq. 6.108 numerically using Python's Scipy `odeint` function[†], assuming that the normalised vector potentials a_x and a_z are given by Eqs. 6.63 and 6.64 (in this case at $\tau = 0$). This accounts for a different polarisation state of the trailing laser pulse. The wake potential ϕ was then used to calculate the normalised longitudinal wakefield E_z/E_{wb}, normalised electron density $\delta n_{\mathrm{e}}/n_0$ and trapping parameter $\Delta\mathscr{H}$ (all given in Section 6.5.1). As before, it was assumed that $L_{\mathrm{p0}} = 1/k_{\mathrm{p}}$ for both laser pulses and that the two pulses were spaced by $k_p\xi = 2\pi$. The first LPLB had a peak normalised vector potential $a_{0,1} = 0.5$ and the second RPLB had a peak normalised vector potential $a_{0,2} = 1.0$.

Figure 6.8 shows a comparison of the wakefield quantities for a conventional 2PII scheme with two linearly-polarised pulses (L-2PII) and the novel scheme with a linearly- and radially-polarised pulse (R-2PII). Figures 6.8a)–d) show key plasma quantities for the L-2PII scheme. As in Section 6.5.1, we see slow variation of the wake potential ϕ, which is positive within both laser pulses, and fast variation of the electron density $\delta n_{\mathrm{e}}/n_0$. However, in this case the two linearly-polarised pulses drive a wakefield which does not trap electrons. This is indicated by the lack of a shaded region for the trapping parameter $\Delta\mathscr{H}$ plotted in Figure 6.8d). The difference between Figures 6.8a)–d) and Figures 6.6a)–d) is due to the fact that the peak normalised vector potential of the two pulses are now lower at $a_{0,1} = 0.5$ and $a_{0,2} = 1.0$, compared to $a_{0,1} = 1.0$ and $a_{0,2} = 2.0$ as we had previously.

Figures 6.8a′)–d′) show the same key plasma quantities for the R-2PII scheme. All laser pulse parameters remain the same as for Figures 6.8a)–d). The only thing that changes is the polarisation state of the trailing laser pulse. We can see that in Figure 6.8b′) the wakefield is highly non-linear, where the longitudinal wakefield E_z/E_{wb} has a characteristic 'sawtooth' profile. Figure 6.8a′) shows that the peak wake potential ϕ is 4.8 times higher than in the L-2PII scheme and Figure 6.8b′) shows that the longitudinal accelerating fields strength E_z/E_{wb} is 3.4 times higher than in the L-2PII scheme. In Figure 6.8c′), we still see the fast variation of the electron density $\delta n_{\mathrm{e}}/n_0$, but we note that this variation has a significantly higher amplitude for the RLPB, and occurs once rather than twice per optical cycle as it does for the LPLB. Looking at Figure 6.8d′), we see that the enhancement of the wakefield allows for electrons to be trapped when they oth-

[†]https://docs.scipy.org/doc/scipy/reference/generated/scipy.integrate.odeint.html

erwise wouldn't be for the comparative L-2PII case. Therefore, it seems from this first study that it may be potentially advantageous to use a RPLB to promote electron trapping in LWFA.

For the case of a wakefield driven by a single radially-polarised laser pulse, we find that the peak normalised vector potential a_0 required to trap electrons is only $a_0 \sim 0.75$ for a wakefield driven by a RPLB, compared to $a_0 \sim 1.7$ for a wakefield driven by a LPLB (the latter is in agreement with work by Chen *et al.*[299]). The wakefield quantities for the RPLB wake are shown in Figure 6.9. In the case where there is a longitudinal component a_z of the laser field in the wake-

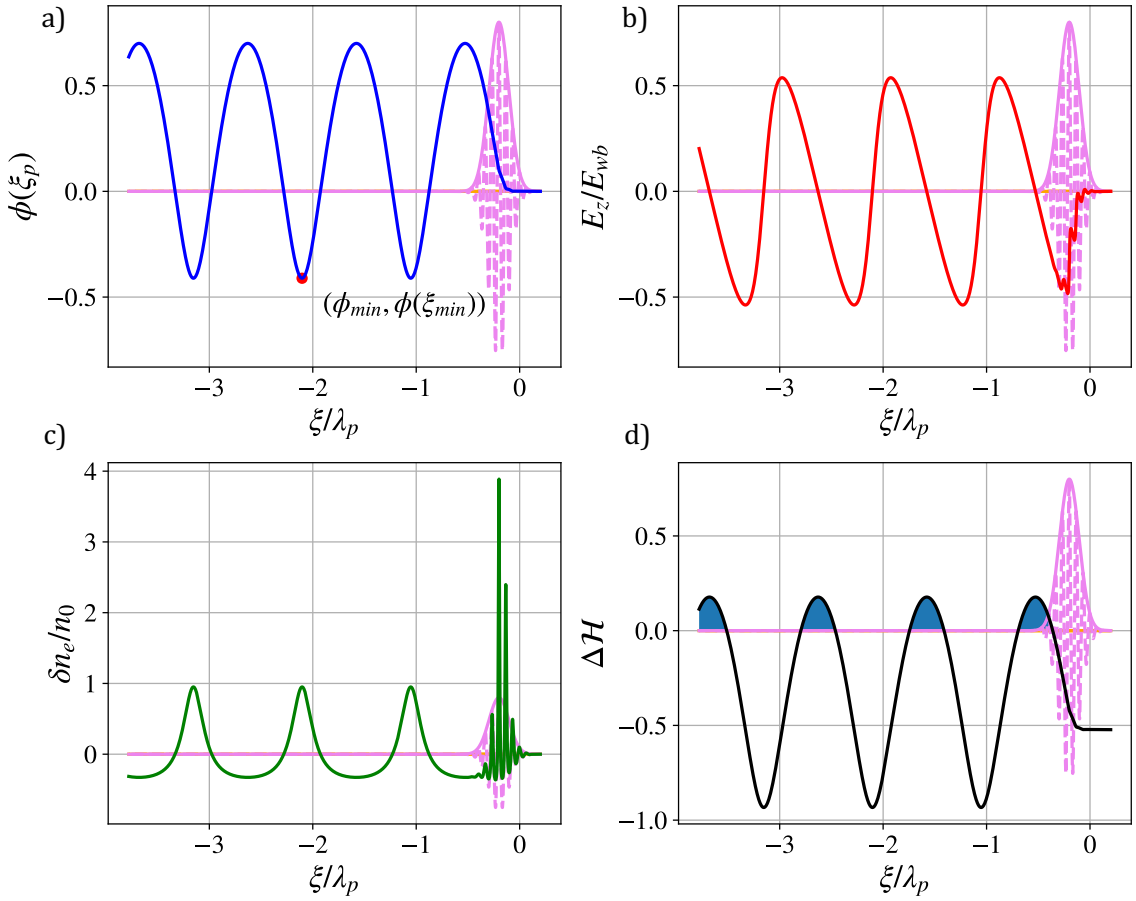


Figure 6.9: Calculated quasistatic wakefield quantities as a function of the moving-frame coordinate ξ/λ_p behind the radially-polarised drive pulse, assuming a 1D relativistic cold fluid in the non-linear regime.[294] Plotted for the case of a radially-polarised drive pulse with $a_{0,1} = 0.75$, travelling to the right. The length of the pulse was $L_{p0} = 1/k_p$. a) Normalised wake potential $\phi(\xi)$ with the minimum wake ϕ_{min} and trapping phase ξ_{min} indicated by a red dot, b) on-axis longitudinal component of the electric field $E_z(\xi)/E_{wb}$, c) normalised electron density $\delta n_e(\xi)/n_0 = n_e/n_0 - 1$ and d) trapping parameter $\Delta\mathcal{H}(\xi) = \mathcal{H}_s - \mathcal{H}_i$, blue shading denotes regions where trapping may occur. Also shown are the normalised vector potential a_z (pink dashed line), and pump strength $|a_z|$ (pink solid line).

field equation, the wake potential ϕ is modified on the right hand side of Eq. 6.108 according to $\phi \rightarrow \phi - a_z\Delta\beta$. It seems like this term introduces a stronger non-linear coupling between the wake potential ϕ and the normalised vector potential a_z of the RPLB, compared to the coupling

between the wake potential ϕ and the normalised vector potential a_x of the LPLB which is introduced in Eq. 6.21 through division of the $(1+\phi)^2$ term by $1+a_x^2$. Further studies are certainly required to understand the reason for this and the conditions under which it may occur.

6.7 Conclusions and Future Work

We presented a preliminary theoretical investigation into whether alternative laser pulse polarisation states could be used to enhance the injection of electrons in LWFAs. We found that RPLBs have a larger spot size in paraxial focusing, but that they have a higher degree of axial confinement which could be used to produce electron bunches with lower energy spreads in ionisation injection. In future work, we intend to also apply the fair comparison case defined in Section 4.4.3 to the RPLB. We found that even in paraxial focusing regimes, the longitudinal field strength of the RPLB can be an appreciable fraction of the transverse field strength. We investigated whether this longitudinal electric field component could be used to enhance electron trapping in an LWFA, by deriving the wakefield quantities and Hamiltonian from first-principles, allowing for alternative polarisation states of the laser pulse. It should be noted that although Bingham states that his derivation is valid for arbitrary polarisation states of the laser pulse,[304] the laser pulse polarisation is still assumed to be in the transverse plane, and therefore does not account for longitudinal components of the laser field. This is the first derivation that allows the laser pulse polarisation to be truly arbitrary, where it is trivial to modify a_x in our derivations here to also allow for $\mathbf{a}_\perp = a_x\hat{x} + a_y\hat{y}$ instead of $\mathbf{a}_\perp = a_x\hat{x}$ (since this has already been done in previous texts, it is omitted here[304]).

By making a quasistatic approximation, with all our assumptions clearly stated, we derived a modified version of Poisson's equation, which we solved numerically to calculate key quantities related to the wakefield. We found that changing the polarisation state of the laser pulse from linearly-polarised to radially-polarised appears to enhance the wakefield and therefore the trapping potential, even when all other parameters of the laser pulse remain the same. We hypothesised that this was due to the way in which the longitudinal component of the laser field a_z appears in the wake equation, as $\phi - a_z\Delta\beta$, rather than multiplying or dividing the wake potential ϕ, as is the case for the transverse component a_x. We hope that there is a mathematical explanation for this and plan to compare this to other known non-linear differential equations in future work.

This study aims to present a framework for how we might understand the impact that the polarisation state of the laser pulse has on the wakefield, but it is very clear that there are many more aspects that we could investigate in the future to try and understand this further. We plan to examine the form of Eq. 6.108 more closely, to try and explain some of the observations made, including the enhancement in the wakefield and the fast variation in the electron density that occurs once per optical cycle for the RPLB. We will study the effect of laser pulse frequency and pulse length in the single-pulse RPLB-driven wakefield scheme. For the R-2PII scheme, we could vary the relative amplitude of both pulses and their spacing, to identify whether the condition for resonance is the same as for the L-2PII scheme.

It is also important to recognise the limitations of such a model. To make a quasistatic approximation, we had to assume that our laser pulses do not evolve in the time that it takes the plasma to transit them. This meant neglecting the effects of CEP phase slip and dephasing of the two pulses, but most crucially we had to neglect diffraction, which happens on short timescales for small focal spots. A full iterative approach to solving Eqs. 6.58–6.61 could be used to model the evolution of the laser fields, rather than assuming a particular time-dependence (this approach is given in Esarey *et al.*[301]). Even so, any 1D model cannot fully account for either diffraction or for the transverse fields of the RPLB away from the axis which, as we saw in Section 6.3, may be even stronger than the longitudinal electric field component in some paraxial focusing regimes. On the other hand, we may benefit from an effect known as relativistic ponderomotive trapping, whereby the annular intensity profile of the RPLB actually drives electrons towards to optical axis.[305] To understand the full spatiotemporal evolution of the laser fields and their interaction with a plasma, we plan to use particle-in-cell (PIC) simulations, to both assess the validity of the model presented here and to firmly establish whether the RPLB can indeed be used to promote trapping in LWFAs. On the basis of the preliminary results obtained here, it would also be interesting to use PIC simulations to study the ionisation of high-Z gas species using RPLBs, to identify whether the tighter confinement of the laser pulse intensity can be used to tune the initial phase-space of ionisation injected electrons.

7 | Optical Injection with Crossed HOFI Channels

This chapter is dedicated in memory of Oscar Jakobsson

7.1 Introduction

7.1.1 Motivation

Hydrodynamic optical-field ionized (HOFI) plasma channels allow for operation of an LWFA in low-density regimes by ensuring that the laser pulse stays focused over distances that are much longer than the Rayleigh range. As already discussed in Chapter 1, operation in low-density regimes may be preferential for a few reasons; from the prevention of dark-current and uncontrolled electron injection that degrades the quality of electron bunches produced, to compatibility with more efficient and higher repetition-rate laser technology. Free-standing channels within a gas target (whether of cell[306] or jet[307] design) are created by the ionisation, heating and subsequent expansion of a column of gas by an optical field,[123,124] for example, the long Bessel-like focus generated by an axicon.[125–127] Such channels are immune to damage,[128] have very low losses with attenuation lengths of the order of 10 m[129] and have been shown to operate at kHz repetition rates.[130] In previous work, it was shown that these channels can also provide a mechanism that injects electrons directly into the wakefield, simply by translating the onset of the axicon focus downstream with respect to the start of the gas target.[128,132] A hemispherical shock wave at the start of the axicon focus forms a natural density transition from the neutral gas into the channel. This density downramp was used to inject electrons which were accelerated to 1.2 GeV with percent-level energy spreads, but was only demonstrated for low total electron bunch charges of a few pC.

In this chapter, we extend this result by showing that in a crossed-channel injection (CCI) scheme, the walls of an additional transverse channel can increase the charge of the electron bunches injected and give additional, independent control over the injection conditions. This transverse channel was formed by focusing an additional laser pulse with independent timing control into the target using an astigmatic focus to generate a plasma sheet.[308] The astigmatic focus was formed upstream of an independent truncated channel used to guide the drive laser pulse, and provided independent control of the injection region, as shown in Figure 7.1. An astigmatic focus which is has a much greater extent along the vertical axis generates a vertical plasma sheet which relaxes the alignment constraints between the drive beam and the transverse astigmatic beam used to generate the injection ramps.[309] We show that it was possible to couple the drive laser pulse into a 27 mm-long HOFI channel through the walls of a transverse channel without significant losses, to achieve accelerating gradients of the laser-driven wakefield which are comparable to our previous experiment.[132] Furthermore, we observe evidence of strong betatron oscillations of the injected electron bunch. Using simplistic models of the betatron motion, we estimate the initial offset of the electron bunch and suggest that this scheme could be used to deliberately instigate betatron oscillations of the electron bunch as an alternative to using external undulators for generating broadband radiation. Undulators are still likely to be required for generating narrowband radiation. We discuss the limitations of such a simplistic model and propose future studies.

The experimental design and setup was led by Dr. James Chappell and Dr. Linus Feder, who were the target area operators alongside Dr. James Cowley for the duration of the experiment. The experiment was carried out by the Oxford LPA group, aided by Prof. Stefan Karsch and Dr. Jinpu Lin from Ludwig-Maximilians-Universität München. Our link person while at Central Laser Facility (CLF) was Dr. Nicolas Bourgeois and we received extensive support from the CLF team throughout the experiment. The original design of the experiment was informed by particle-in-cell simulations carried out by Oscar Jakobsson and Johannes van de Wetering. The analysis of all experimental results was carried out by the author and specific contributions from other members of the group are acknowledged on a case-by-case basis for each diagnostic.

7.1.2 Objectives

The objectives of this chapter are to:

1. Demonstrate injection and acceleration of higher charge electron bunches using a transverse injector pulse;

2. Demonstrate that propagation through the transverse injector pulse does not significantly impact guiding of the laser pulse;

3. Assess the validity of an analytic model for the electron betatron motion in predicting energy-divergence oscillations in the electron spectrometer data.

7.2 Experimental Design

Figure 7.2 shows the experimental layout employed at TA3 of the Rutherford Appleton Laboratory's Central Laser Facility using their Astra-Gemini Ti:Sa laser.

7.2.1 Astra-Gemini

Astra-Gemini is a laser at Central Laser Facility; a part of the Rutherford Appleton Laboratories, located near Oxford. Astra-Gemini is a Ti:Sa laser system which provides two linearly-polarised beams to Target Area 3 (TA3), each up to 15 J in energy with full width half maximum duration of around 40 fs, with each beam capable of providing up to ~ 300 TW on-target. Both beams have a central wavelength of $\lambda_0 = 800$ nm and the pulse duration was measured to be $\tau_p = (45 \pm 14\,\text{fs})$ during our experiment (the details of this measurement are given in Section 7.4.1). These two beams will be termed the 'South' beam and the 'North' beam. The whole South beam was used as the drive beam. The North beam had two cut-outs, each with a 50 mm diameter. One cut-out was used for the channel-forming beam and the other was further split to form the injector and probe beams.

CHANNEL-FORMING BEAM A simplified schematic showing the layout of all beams is shown in Figure 7.1. The channel-forming beam (red) had an estimated energy on target of (117.9 ± 30.5)mJ. After entering the vacuum chamber, a delay stage was used to control the arrival time of the channel-forming beam relative to the drive. The channel-forming beam arrived first, with a relative arrival time $\Delta t_{axi} = (2.85 \pm 0.25)$ns prior to the drive beam. The beam was then focused by a transmissive axicon lens with an approach angle of $\alpha = 2.5°$. The focusing axicon light was reflected into the target from the surface of an annular mirror with an aperture radius of $r_A = 9$ mm. The drive beam was aligned to be collinear with the central axis defined by the axicon

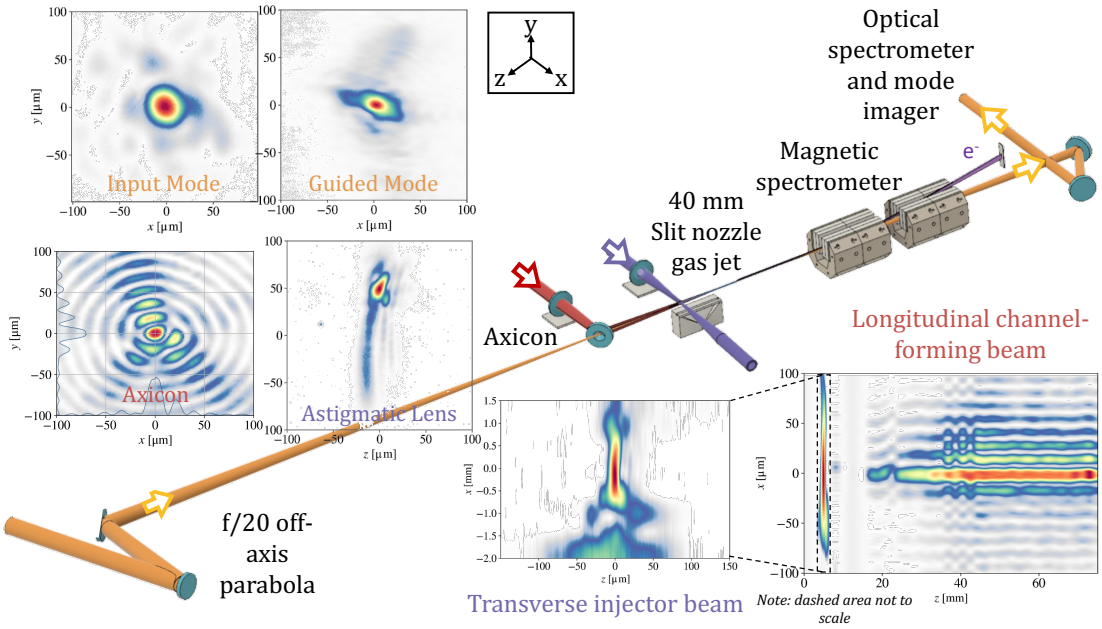


Figure 7.1: Simplified schematic of the experimental layout and the measured focal spot data. An $f/20$ parabola focuses the drive beam (orange) into a slit nozzle gas jet. The longitudinal channel was formed by reflecting the focusing light from an axicon (red) from the annulus of a holey mirror. An astigmatic focus (purple) was used to generate a perpendicular vertical plasma channel. Top inset: measured focal spots for all three beams, with the drive laser pulse shown both at the focus and at the exit of the longitudinal channel. Bottom inset: top-down view of the crossed channel-forming beams with the astigmatic injector beam crossing $z_{\mathrm{inj}} = 5\,\mathrm{mm}$ and the axicon guiding beam starting $z_{\mathrm{axi}} = 13.03\,\mathrm{mm}$ from the start of the jet.

line focus. The target was also aligned to this axis by optimisation of the plasma fluorescence signal.

INJECTOR AND PROBE BEAMS The injector (purple) and probe lines (blue) were split from the same line (purple) using a quarter-wave plate (QWP) as shown in Figure 7.2a). This generates an elliptically polarised beam with a varying relative amplitude along each axis that depends on the polar angle of the QWP. The injector beam (purple) was formed by reflection from a polarising beam-splitter (PBS) and was retro-reflected from a delay stage which was used to control the arrival time relative to the channel-forming and drive beams. The injector beam came 1 ns later than the channel-forming beam, with a relative arrival time of $\Delta t_{\mathrm{inj}} = (1.85 \pm 0.15)\mathrm{ns}$ prior to the drive beam. The delay stage was followed by a deformable mirror (DM), which was used for additional control over the spatial profile of the focus. The beam was focused into the target by an achromatic lens with a focal length of $f_{\mathrm{inj}} = 300\,\mathrm{mm}$ set at a slight angle to generate an astigmatic focus. The astigmatic focus has a much greater extent vertically (y-axis in Figure 7.1) compared to horizontally (z-axis in Figure 7.1). This generates a plasma sheet which is large in the x- and y-directions, but narrow along the z-axis, generating shocks which are close together but easy to

align the drive laser pulse to transversely. The focusing beam was reflected into the target by an 800 nm mirror, with the beam crossing the jet $z_{\text{inj}} = 5$ mm from the upstream end. The astigmatic transverse injector beam had an energy on target that could be varied from (20.7 ± 26.2)mJ to (386.3 ± 58.2)mJ using the QWP and PBS combination. The rest of the energy from this North beam cut-out went to the probe. The probe line was formed from transmission through the PBS and was frequency-doubled from 800 nm to 400 nm using a barium borate (BBO) crystal and passed through the back of the 800 nm mirror that was used to couple the injector beam into the target. Frequency-doubling of the probe allowed us to filter out all 800 nm light from the transverse injector after the target. This was done by reflecting the probe light off a 400 nm dielectric mirror and allowing the injector beam to pass through and be dumped. A delay stage on the probe line allowed us to probe a range of evolution times of the transverse injector and longitudinal axicon channels.

DRIVE BEAM As shown in Figure 7.2b), after entering the target chamber, the drive beam was reflected from mirrors at the roof of the chamber, above the beamlines shown in Figure 7.2a). Reflecting the beam off a deformable mirror (DM) allowed us to minimise aberrations on the focal spot and also to control the longitudinal position of the focus by adjusting the second-order spatial phase. The beam then reflected off an $f/20$ parabola at the end of an external arm of the chamber and was focused into the target through a $r_{\text{A}} = 9$ mm aperture in an annular mirror (AM). The $f/20$ off-axis parabola focused the beam to a focal spot with a $1/e^2$ intensity radius spot size $w_0 \approx 25\,\mu$m, a Rayleigh range of $z_{\text{R}} \approx 2.5$ mm and an M^2 beam quality factor of 1.65. The drive laser pulse had an average energy on target of (6.9 ± 0.3)J.

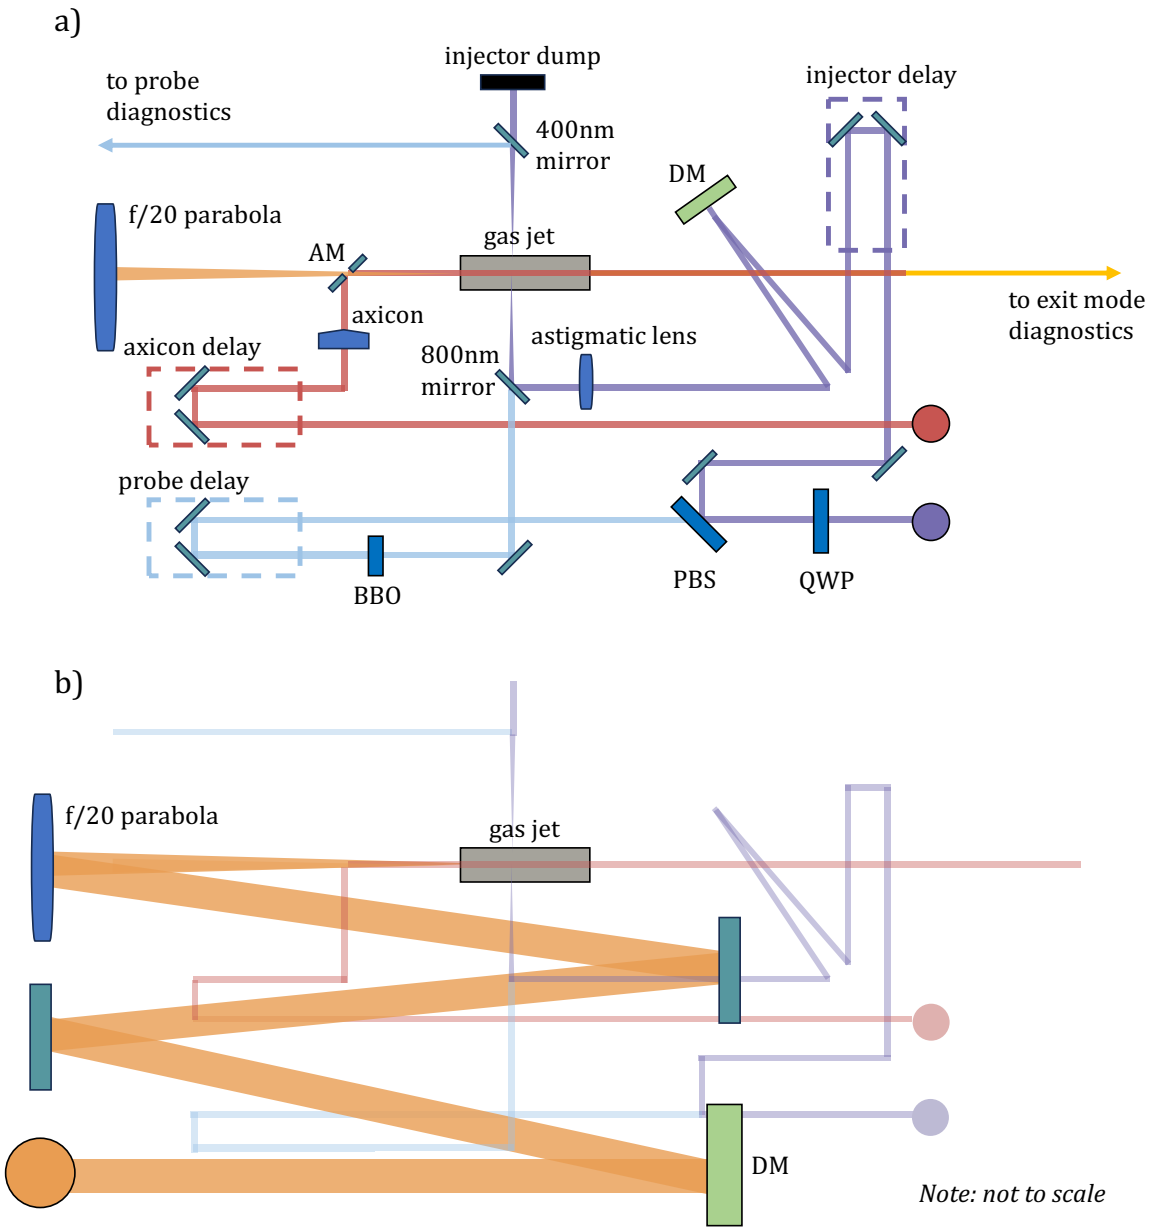


Figure 7.2: Schematic of the full experimental layout. a) The two collimated North beam cut-outs enter from TA3 at the circles. The axicon-focused beam (red) was reflected into the target from the surface of an annular mirror (AM). The injector and probe lines (purple and blue) were split from the same cut-out using a quarter-wave plate (QWP) and polarising beam-splitter (PBS) to control the relative energy in each line. All three beams had separate delay stages to control all relative timings. b) The collimated South beam (orange) entered from TA3 at the circle and was focused into the target by an $f/20$ parabola.

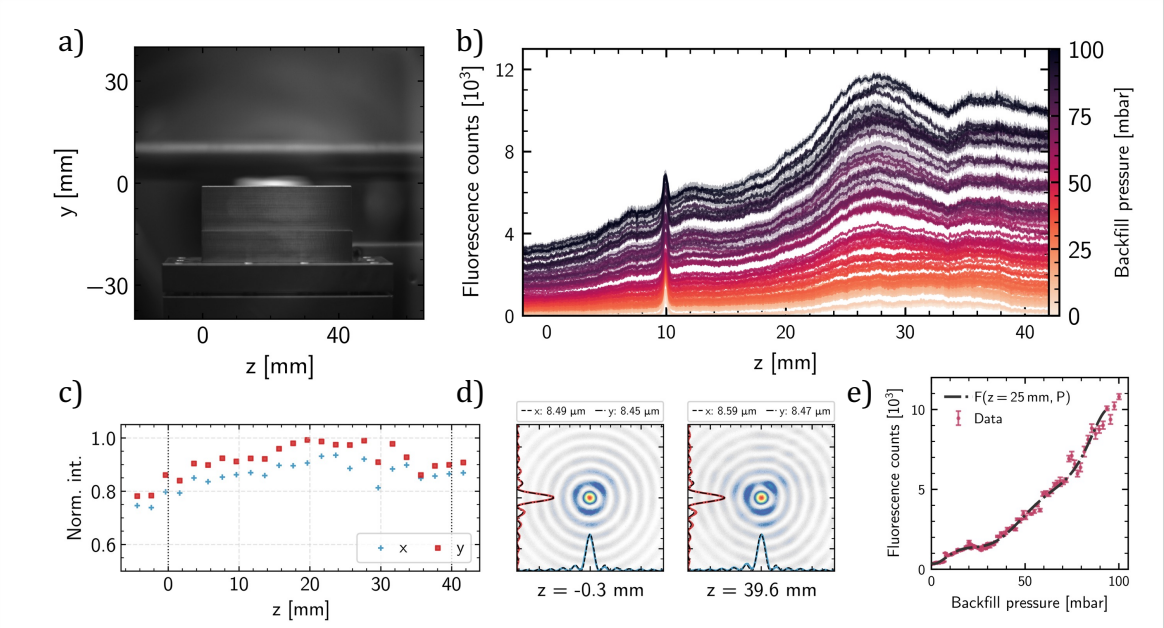


Figure 7.3: Gas jet target and correlations between backing pressure and fluorescence intensity. Figure reproduced with permission from Dr. James Chappell and Dr. Linus Feder. a) Photograph of the gas jet; b) calibration curves of the fluorescence intensity as a function of backfill pressure; c) measured axicon longitudinal intensity profile; d) axicon foci measured at the front and back of the jet; (e) example calibration curve (for $z = 25\,\text{mm}$) that was used to relate the measured fluorescence intensity to the gas pressure.

7.3 Gas Jet Target

The target was a 40 mm-long slit nozzle gas jet, with a hydrogen (H_2) gas species doped with 1% argon (Ar) dopant. The backing pressure could be varied from 0 bar to 80 bar, with a gate valve opening time of 20 ms. Figure 7.3a) shows a photo of the gas jet target. Prior to the Astra-Gemini experiment, the plasma fluorescence intensity was recorded for different backfill pressures using the Oxford LPA Group's TW laser system. A camera positioned directly above the gas jet recorded the 2D $x-z$ intensity profile of the fluorescence of the longitudinal HOFI channel generated by the axicon-focused channel-forming beam. These calibrations were performed by Dr. James Chappell and Dr. Linus Feder. Figure 7.3b) shows the counts recorded on the fluorescence camera at different backing pressures. The counts were integrated transversely in x across the ~ 20 pixels that correspond to the region of interest of the HOFI plasma itself, to give a 1D lineout as a function of the longitudinal position z. These lineouts can be used to fit the correlation between the backfill pressure and the recorded fluorescence counts for each longitudinal position, as shown in Figure 7.3e).

Figure 7.3c) shows normalised measurements of the axicon peak intensity, taken by recording

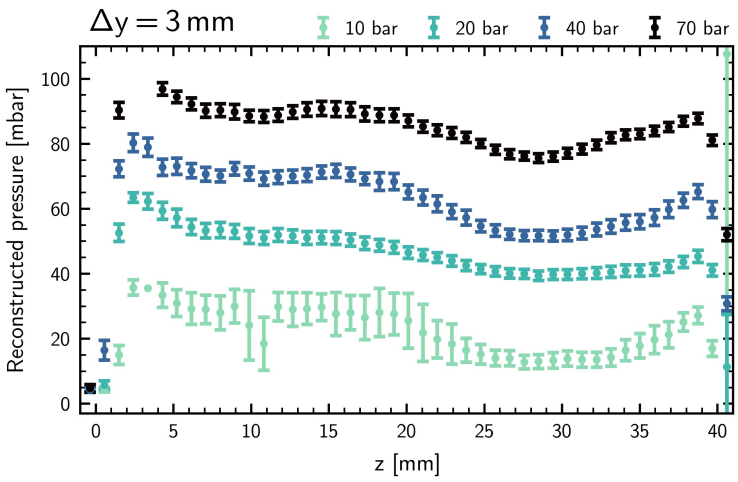


Figure 7.4: Gas jet density at 3 mm above the jet. Gas pressure is shown for different backing pressures for all longitudinal z positions along the jet. Figure reproduced with permission from Dr. James Chappell and Dr. Linus Feder.

the axicon focal spot for different z-positions along the optical axis. The recorded peak intensities were normalised to the maximum peak intensity recorded over the entire scan. The axicon foci recorded at the front and back of the jet can be seen in Figure 7.3d).

The correlation plotted in Figure7.3e) was used to convert the plasma fluorescence intensity into a gas pressure measurement, as shown in Figure 7.4. Measurements taken in the Oxford LPAG lab for the same conditions as the Astra-Gemini experiment allow us to characterise the gas jet and estimate the gas density during the experiment for different backing pressures, since the axicon focus was also positioned at the same height $\Delta y = 3\,\text{mm}$ above the top surface of the gas jet during the Astra-Gemini experiment. The gas density was actually more uniform with respect to z during the experiment than for the data shown here; the dip on the right hand side of Figure 7.4 was due to a small obstruction in the jet that was removed prior to the experiment. In previous experiments, transverse interferometry was used to estimate the on-axis density in the HOFI channel.[125,128] This method will be discussed in Section 7.4.5. However, we found that it was not possible to retrieve the on-axis plasma density from the interferograms taken in this experiment; in future work magnetohydrodynamic (MHD) simulations will instead be used to estimate the plasma density. Knowing the density of the ambient gas prior to ionisation will be crucial for these simulations.

The use of a gas jet, rather than a gas cell, allowed for full optical access without the need for entrance and exit windows or a pinhole for the drive beam. Another advantage to using a gas jet instead of a gas cell is that there is no front pinhole to aperture the axicon-focused channel-

forming beam. Aperturing the axicon beam reduces the length of the axicon focus and therefore the channel, placing a minimum requirement on the size of the front pinhole for a certain channel length.[125] The length of density up-ramp at the start of the target is proportional to the size of the pinhole,[128] such that the longer the channel that we require, the larger the density up-ramps are likely to be. Gas jet are known to have very short up-ramps and this prevents unwanted effects that occur where the length of the up-ramp is longer than the Rayleigh range of the drive beam.[128] We can see from Figure 7.4 that the up-ramp was on the order of $\sim 2\,\mathrm{mm}$, which was shorter than the Rayleigh range of the drive laser pulse $z_R \approx 2.5\,\mathrm{mm}$, meaning that we should have had no detrimental effects on the laser pulse, since the density up-ramp at the start of the target is short.

7.4 Diagnostics

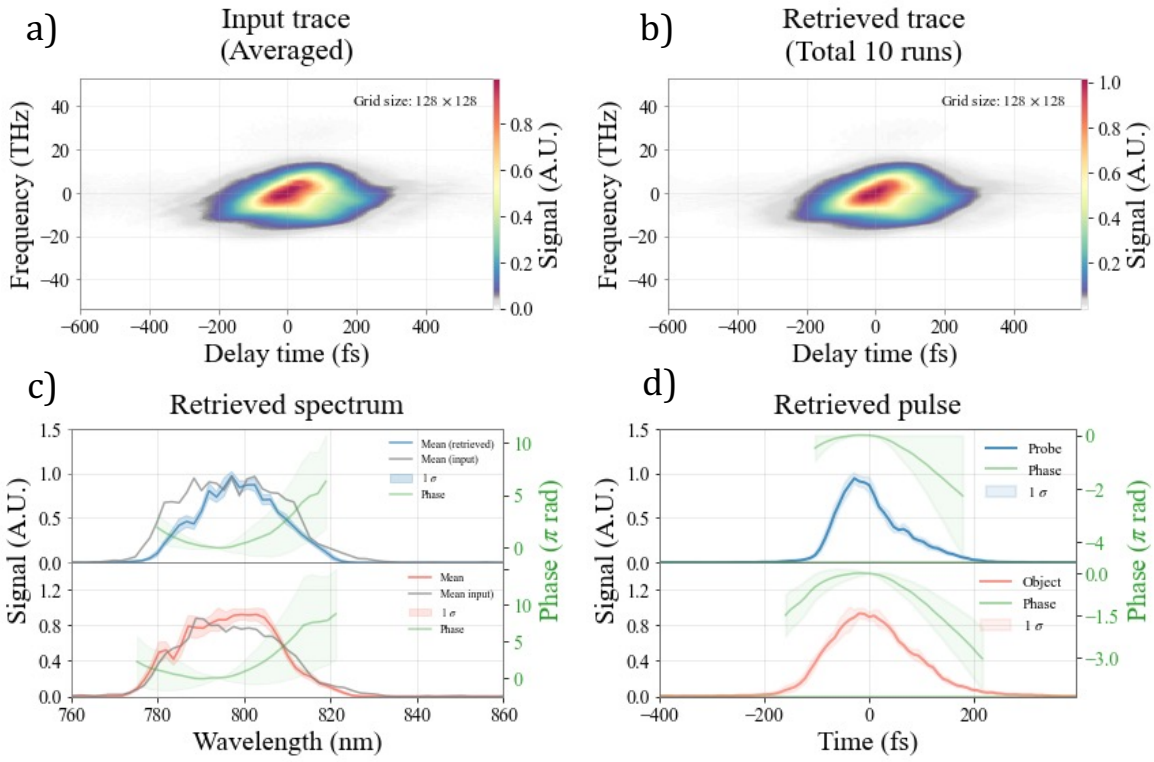


Figure 7.5: Pulse duration estimated from the frequency-resolved optical gating (FROG) measurement. a) Experimental FROG trace, taken by averaging 10 shots. b) Retrieved FROG trace after mPIE algorithm. c) Retrieved spectrum and spectral phase (green) of reference (blue) and object (red) pulses. d) Retrieved temporal profile and temporal phase (green) of refererence (blue) and object (red) pulses.

7.4.1 Pulse Duration

The pulse duration was the same for both the North and South beams and was estimated using the frequency-resolved optical gating (FROG) technique. All measurements and analysis

were carried out by Warren Wang using a custom-built setup.[310] Figure 7.5a) shows the input trace averaged over 10 shots. Figure 7.5b) shows the trace retrieved using the mPIE retrieval algorithm. Figure 7.5c) shows the retrieved spectrum and spectral phase. The two traces represent the probe (blue) and object (red) traces respectively. The spectrum was peaked at around 800 nm as expected, with a FWHM bandwidth of around 30 nm. Figure 7.5d) shows the retrieved temporal profile and phase of the pulse. The estimated pulse duration directly from the phase retrieval was (132.1 ± 29.8 fs). There was a certain amount of BK7 glass in the FROG system that disperses the pulse, increasing the pulse duration. Accounting for this, the estimated pulse duration on target was $\tau_{\mathrm{p}} = (45 \pm 14\,\mathrm{fs})$.

7.4.2 Focal Spots

Focal spot cameras were positioned to image the foci of the channel-forming, injector and drive beams. It is worth noting that the foci could only be imaged in the low-power mode due to the risk of damaging the cameras. Therefore, all focal spot analyses cannot account for shifts in the beam when switching from low-power to full-power shots, although in theory any changes should be minimal.

The focal spot cameras used were all 12-bit CMOS cameras from Allied Vision (Manta G-033B). Spatial calibrations were taken by imaging USAF targets and then calculating the resolution per pixel. The resolutions were 1.43 µm/pixel for the injector focal camera and 2.4 µm/pixel for the focal camera which was used to image both the drive and axicon channel-forming beams (this camera was also used to overlap these beams in the transverse plane). Exemplary focal spot profiles for all beams can be seen in Figure 7.1. Images taken at different z-positions along the optical axis of each focus were stitched together for the top-down view of both channel-forming beams, this is shown as the inset at the lower-right of Figure 7.1.

Peak Intensity

Drive and Injector Beams For the drive and injector beams, the peak intensity I_{peak} can be estimated from the focal spot data. The intensity I at any point in the transverse plane can be estimated by scaling the counts I_{counts} at a particular pixel against the total integrated counts, summed over all pixels. The intensity $I(x_i, y_i)$ is therefore

$$I(x_i, y_i) = \frac{\mathscr{E}}{\tau_{\mathrm{p}} \mathrm{d}x^2} \frac{I_{\mathrm{counts}}(x_i, y_i)}{\sum_{x_i, y_i} I_{\mathrm{counts}}(x_i, y_i)}, \tag{7.1}$$

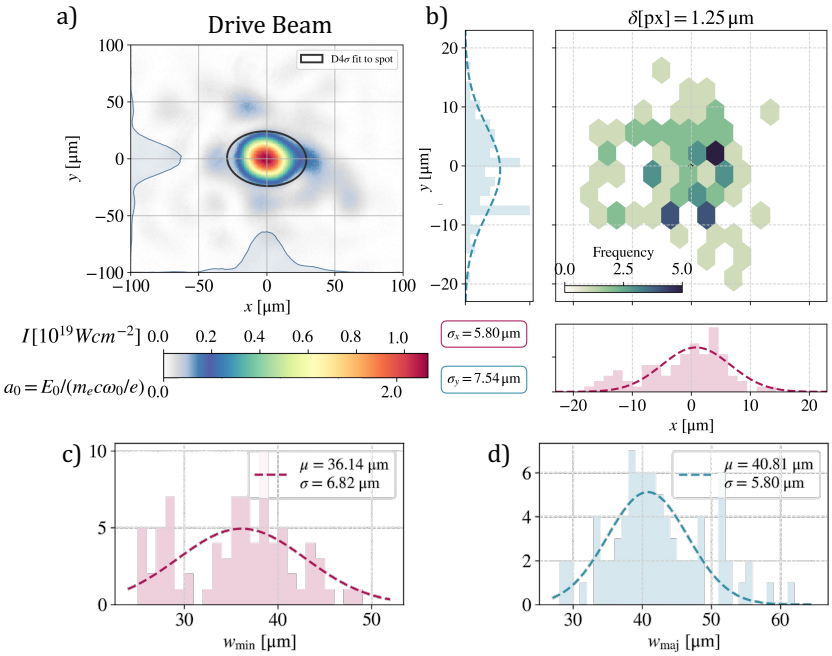


Figure 7.6: Variation in the centroid position and spot size of the drive beam. a) Shows the focal spot of the drive beam with a D4σ fit to the spot. The intensity I and normalised vector potential a_0 are shown on the colorbar; b) shows a histogram of the centroid position of the $f/20$ drive beam, darker shades indicate a higher frequency of shots at this position and summed histograms are shown below and to the left, with the same scales as the 2D plot; c) shows a histogram of the D4σ minor axis length and d) shows the corresponding histogram for the major axis.

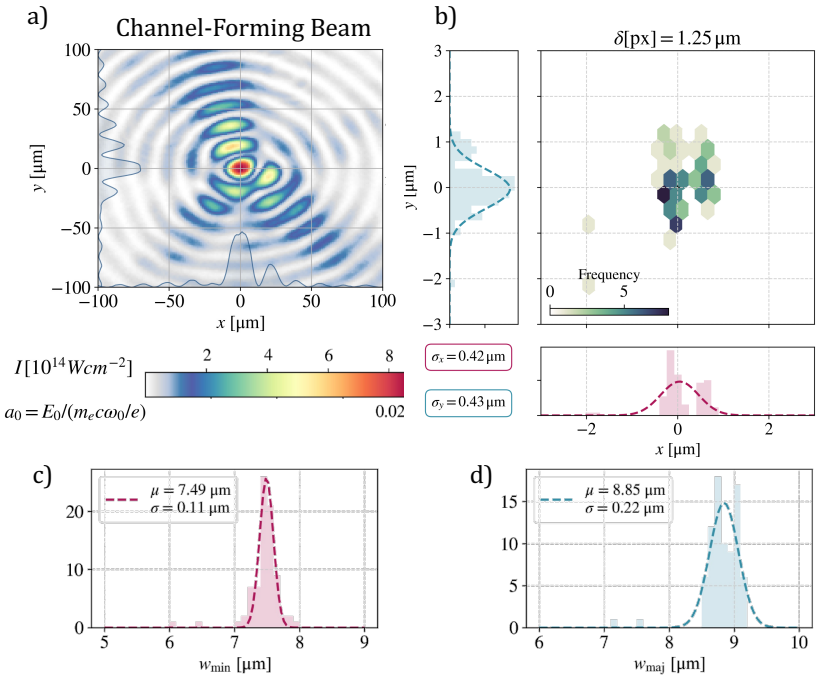


Figure 7.7: Variation in the centroid position and spot size of the channel-forming beam. a) Shows the focal spot of the channel-forming beam. The intensity I and normalised vector potential a_0 are shown on the colorbar; b) shows a histogram of the centroid position of the axicon channel-forming beam, darker shades indicate a higher frequency of shots at this position and summed histograms are shown below and to the left, with the same scales as the 2D plot; c) shows a histogram of the D4σ minor axis length and d) shows the corresponding histogram for the major axis.

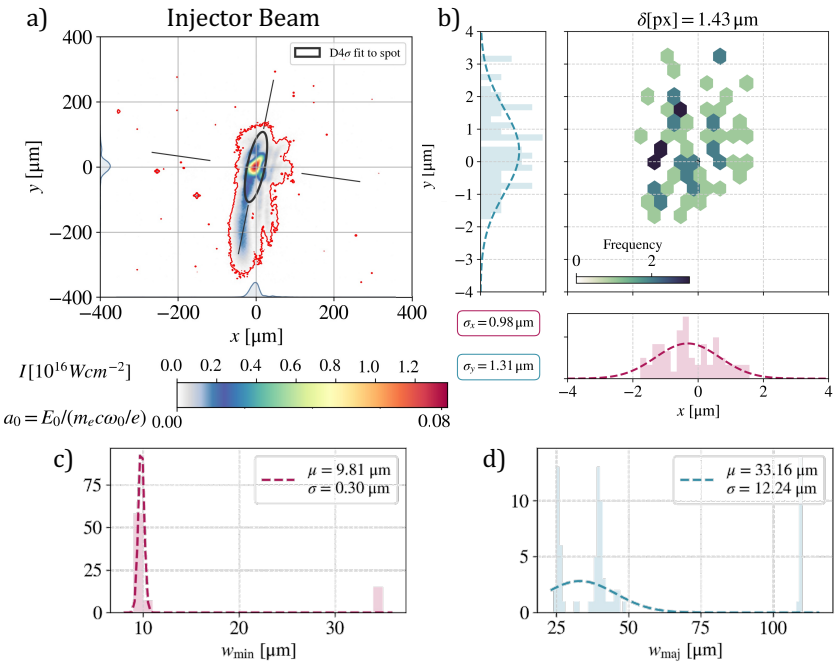


Figure 7.8: Variation in the centroid position and spot size of the injector beam. a) Shows the focal spot of the injector beam with a D4σ fit to the spot. The intensity I and normalised vector potential a_0 are shown on the colorbar. The red contour indicates the $I \sim 10^{14}\,\mathrm{W\,cm^{-2}}$ ionisation threshold for hydrogen; b) shows a histogram of the centroid position of the injector beam, darker shades indicate a higher frequency of shots at this position and summed histograms are shown below and to the left, with the same scales as the 2D plot; c) shows a histogram of the D4σ minor axis length and d) shows the corresponding histogram for the major axis.

where dx is the width of one pixel (it is assumed that all pixels were square such that dx^2 is the area of one pixel) and $\mathscr{E}$ was the total laser pulse energy on target, calculated from the pre-compressor energy which was recorded on every shot and a throughput calibration from a measure of the incident laser pulse power both pre-compressor and at the target. Assuming that the sensor of the camera was large enough to capture the whole beam, we can assume that the total counts was proportional to the energy of the beam on target, such that dividing the energy of the beam by the total number of counts summed over all pixels gives an estimate of the energy per count. This can be converted to the flux per count if this value is then divided by the area per pixel. Dividing again by the pulse duration gives an estimate of the intensity per count. This is a conversion factor which gives the intensity if we multiply the background-subtracted image by this factor. This intensity can then be converted to the normalised vector potential using Eq. 1.7. In the case of a laser spot which is better described by a Bessel function, there is an additional step to the analysis as follows.

CHANNEL-FORMING BEAM The intensity of a Bessel-like axicon beam can be estimated by counting the total number of rings that can be seen in the focal spot data, assuming that the energy contained in each ring is the same.[125] The total number of rings can be estimated according to

$$N = \frac{2\alpha r}{\lambda_0}, \tag{7.2}$$

where $\lambda_0 = 800\,\text{nm}$, $\alpha = 2.5^\circ$ was the angle of approach and $r = 25\,\text{mm}$ was the radius of the input beam. Dividing the energy on target by the total number of rings N gives the estimated energy per ring. By counting the number of full and partial rings incident on the sensor M, we can approximate the total energy on the sensor by multiplying the number of visible rings M by the energy per ring. We then performed the same calibration as in the previous section; dividing this total energy by the total number of counts, the pixel area and the pulse duration.

The intensity I and normalised vector potential a_0 of the drive beam are shown in Figure 7.6a). The drive beam was found to have a peak intensity $I_0 \sim 10^{19}\,\text{W cm}^{-2}$, which corresponds to $a_0 \sim 2.3$. The intensity I and normalised vector potential a_0 of the transverse injector beam are shown in Figure 7.8a), where the peak intensity was $I_0 \sim 10^{16}\,\text{W cm}^{-2}$ corresponding to $a_0 \sim 0.08$. Similarly, the longitudinal channel-forming axicon beam is shown in Figure 7.7, where the peak intensity was $I_0 \sim 10^{15}\,\text{W cm}^{-2}$ corresponding to $a_0 \sim 0.02$. The intensity required to ionise hydrogen (H_2) and the four outer electrons of argon (Ar) is just over $I_0 \sim 10^{14}\,\text{W cm}^{-2}$, such that the intensity of both channel-forming beams was sufficiently high to perform the optical field ionisation for the formation of HOFI channels.[125,128]

Spot Size and Jitter

Codes that were written for previous experiments by Dr. James Chappell were used to analyse the focal spot sizes and jitter. The ISO standard D4σ method was used to calculate the focal spot size.[215] Along one transverse axis, the D4σ beam diameter is given in terms of the second moment of the background-subtracted intensity distribution I as

$$\text{D4}\sigma_x = 4\sqrt{\frac{\int_{-\infty}^{\infty} I(x,y)(x-\bar{x})^2\,\text{d}x\,\text{d}y}{\int_{-\infty}^{\infty} I(x,y)\,\text{d}x\,\text{d}y}}, \tag{7.3}$$

where the x-coordinate centroid is given by

$$\bar{x} = \frac{\int_{-\infty}^{\infty} I(x,y)x\,\mathrm{d}x\,\mathrm{d}y}{\int_{-\infty}^{\infty} I(x,y)\,\mathrm{d}x\,\mathrm{d}y}. \tag{7.4}$$

Figure 7.6b) shows the spatial jitter of the drive beam focal spot over 80 subsequent shots. The centroid position $(\bar{x},\bar{y})$ was calculated for each shot and a Gaussian distribution was fitted to a histogram of the centroid positions over all shots. The standard deviation of the spatial jitter was 5.80 µm in the x-direction and 7.54 µm in the y-direction. The D4σ spot size calculated in the x- and y- directions was used to calculate the elliptical D4σ major and minor axes

$$\mathrm{D4}\sigma_{\mathrm{min}} = \min\left\{\mathrm{D4}\sigma_{+}, \mathrm{D4}\sigma_{-}\right\}, \tag{7.5}$$

$$\mathrm{D4}\sigma_{\mathrm{maj}} = \max\left\{\mathrm{D4}\sigma_{+}, \mathrm{D4}\sigma_{-}\right\}, \tag{7.6}$$

where

$$\mathrm{D4}\sigma_{+} = \left\{2\Big[\mathrm{D4}\sigma_x + \mathrm{D4}\sigma_y + \sqrt{(\mathrm{D4}\sigma_x - \mathrm{D4}\sigma_y)^2 + 4\mathrm{D4}\sigma_{xy}^2}\Big]\right\}, \tag{7.7}$$

$$\mathrm{D4}\sigma_{-} = \left\{2\Big[\mathrm{D4}\sigma_x + \mathrm{D4}\sigma_y + \sqrt{(\mathrm{D4}\sigma_x - \mathrm{D4}\sigma_y)^2 + 4\mathrm{D4}\sigma_{xy}^2}\Big]\right\}, \tag{7.8}$$

with the cross-term D4σ_{xy} given by

$$\mathrm{D4}\sigma_{xy} = 4\sqrt{\frac{\int_{-\infty}^{\infty} I(x,y)(x-\bar{x})(y-\bar{y})\,\mathrm{d}x\,\mathrm{d}y}{\int_{-\infty}^{\infty} I(x,y)\,\mathrm{d}x\,\mathrm{d}y}}. \tag{7.9}$$

For the drive beam, the mean D4σ_{min} minor axis beam diameter was 36.14 µm with a standard deviation of 6.82 µm, as shown in Figure 7.6c), and the D4σ_{maj} major axis beam diameter was 40.81 µm with a standard deviation of 5.80 µm as shown in Figure 7.6d).

The same calculations were repeated for the channel-forming and injector beams, again over 80 subsequent shots each. In this case, the spatial jitter was found to be much smaller. For the channel-forming beam, the spatial jitter was 0.42 µm in the x-direction and 0.43 µm in the y-direction, as shown in Figure 7.7b), suggesting that the position of the channel was very stable from shot-to-shot. The mean D4σ_{min} minor axis beam diameter was 7.49 µm with a standard deviation of 0.11 µm as shown in Figure 7.7c), and the D4σ_{maj} major axis beam diameter was 8.85 µm with a standard deviation of 0.22 µm, as shown in Figure 7.7d). For the transverse

injector beam, the spatial jitter was 0.98 µm in the x-direction and 1.31 µm in the y-direction. The beam was highly astigmatic to generate a vertical plasma sheet and to therefore improve the shot-to-shot probability of spatial overlap with the focus of the drive beam. The astigmatism can be seen in the strong asymmetry between the major D4σ_{maj} and minor D4σ_{min} axis lengths. The mean D4σ_{min} minor axis beam diameter was 9.81 µm with a standard deviation of 0.30 µm and the D4σ_{maj} major axis beam diameter was over three times larger at 33.16 µm with a standard deviation of 12.24 µm. The standard deviation in the beam diameters indicate that the D4σ spot size calculation was not particularly effective in this case, so we have also plotted a contour of the ionisation threshold intensity at $\sim 10^{14}\,\mathrm{W\,cm^{-2}}$ (shown by the red line in Figure 7.8), which indicates that the injector beam had a width of $\sim$ 100 µm and a height of $\sim$ 450 µm. The spot size of the drive beam was around 20× smaller than the vertical height of the injector beam. As already mentioned, this made it easier to ensure that the drive beam overlapped with the plasma generated by the transverse injector in the y-direction, increasing the probability of electron injection and thereby making this technique more reliable.

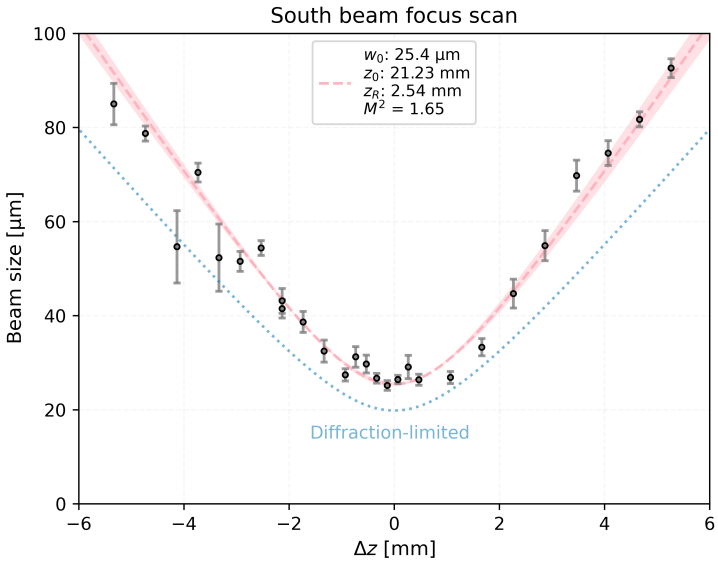


Figure 7.9: Focal scan and beam quality analysis of the drive beam. Data points are shown in grey, a fit to the data is shown in red and comparison with the theoretical diffraction-limited spot is shown by the dashed blue line.

Beam Quality of Drive Beam

A focal scan of the $f/20$ drive beam was performed in the low-power mode. The Gaussian spot size was estimated from the calculated $\mathrm{D4}\sigma_{\mathrm{min}}$ and $\mathrm{D4}\sigma_{\mathrm{maj}}$ axis lengths as

$$w_0 = \frac{\mathrm{D4}\sigma_{\mathrm{min}} + \mathrm{D4}\sigma_{\mathrm{maj}}}{4}. \tag{7.10}$$

This calculation was repeated at each longitudinal z-position and this data is plotted in Figure 7.9. At the focus, the drive beam had an $1/\mathrm{e}^2$ intensity radius spot size of $w_0 \approx 25\,\mu\mathrm{m}$. The Rayleigh range was found to be $z_{\mathrm{R}} \approx 2.5\,\mathrm{mm}$ and the M^2 beam quality factor was 1.65, suggesting a beam which was much poorer quality than a diffraction-limited spot.[216] The M^2 beam quality factor was determined by a fit of the $w(z)$ spot size data found from the D4σ spot sizes by Eq. 7.10 to

$$w^2(z) = w_0^2 + M^4\left(\frac{\lambda_0}{\pi w_0}\right)^2 (z - z_0), \tag{7.11}$$

where w_0 was the D4σ spot size at the focus position $z = z_0$, given by Eq. 7.10 where $w_0 = w(z_0)$, and $\lambda_0 = 800\,\mathrm{nm}$ was the central wavelength of the laser pulse. It should be noted that other methods may be used to calculate the spot size,[125] but since a thorough review of these methods has already been given, we chose the method that is known to be most robust.

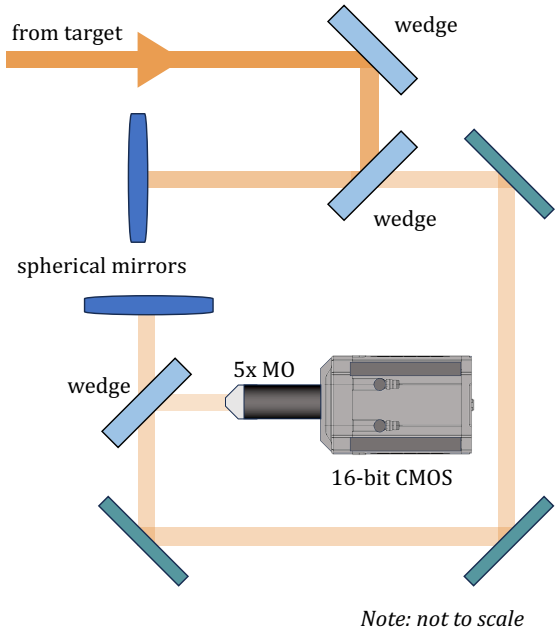


Figure 7.10: Schematic of the guiding diagnostic. The drive beam (after propagation through the target) was attenuated using the reflections of BK7 wedges. A pair of spherical mirrors were used to image the beam onto a 16-bit CMOS camera with a 5× microscope objective.

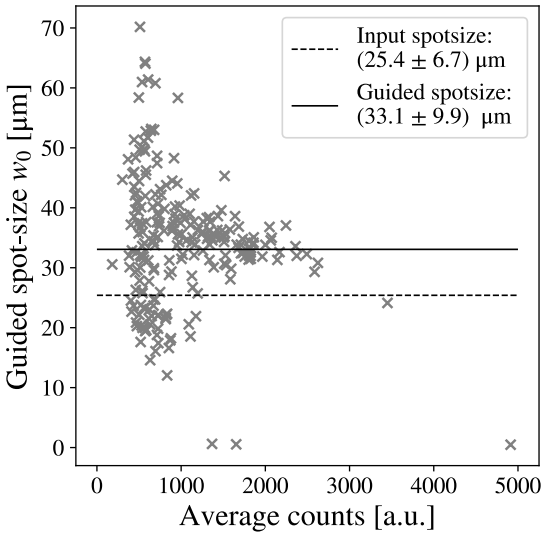


Figure 7.11: Scatter plot of the guided mode D4σ spot size against the counts on guided mode camera. Dashed black line shows the calculated spot size of the input drive beam and the solid black line shows the count-weighted average spot size of the guided beam. The average guided spot size $w_{0,\text{guided}}$ weighted by the number of counts was $w_{0,\text{guided}} = (33.8 \pm 9.9)\,\mu\text{m}$ (solid line) and the calculated input spot size of the drive beam $w_0 = (25.4 \pm 6.7)\,\mu\text{m}$ (dashed line).

7.4.3 Guiding Diagnostic

After interaction with the target, the drive beam was attenuated by reflection off two 10″ BK7 wedges, each with an angle of a few degrees. A pair of 4″ spherical mirrors were used to image the beam onto a 16-bit sCMOS camera from Andor (Neo 5.5). An additional filter wheel with variable ND was used for final attenuation of the beam and a 5× microscope objective was used to enlarge the focus. The final resolution of the camera looking at the guided mode was 0.495 µm/pixel. The first spherical mirror was on a stage to optimise the imaging of the focus of the drive beam at the exit of the target. Additional spectrometers would have allowed for a diagnosis of the wakefield,[84] however in this experiment we only had a camera to image the guided drive beam at the end of the target. The exit mode camera had a narrow spectral range which did not capture all guided light, and since the spectrum of the drive beam changes from shot-to-shot, we were not able to perform an absolute calibration of the throughput from the target to the camera. This meant that we were not able to calculate the absolute guided energy. Instead, the D4σ spot size of the guided spot and the uncalibrated energy from the number of counts on the camera within this region were used to correlate the guiding quality with the electron injection. The aim of this experiment was not to study the guiding in quality in great detail, since this had already been performed in previous experiments.[125,128]

Figure 7.11 shows a scatter plot of the guided spot size against the average counts. For particular values of the guided spot size, the average number of counts on the guiding diagnostic camera increases. The count-weighted average of the guided spot size was used to estimate the matched spot size of the channel. The average guided spot size $w_{0,\text{guided}}$, weighted by the number of counts, was $w_{0,\text{guided}} = (33.8 \pm 9.9)\,\mu\text{m}$, which was 8.4 µm larger than the calculated input spot size of the drive beam $w_0 = (25.4 \pm 6.7)\,\mu\text{m}$. From Eq. 1.19, we can estimate the matched

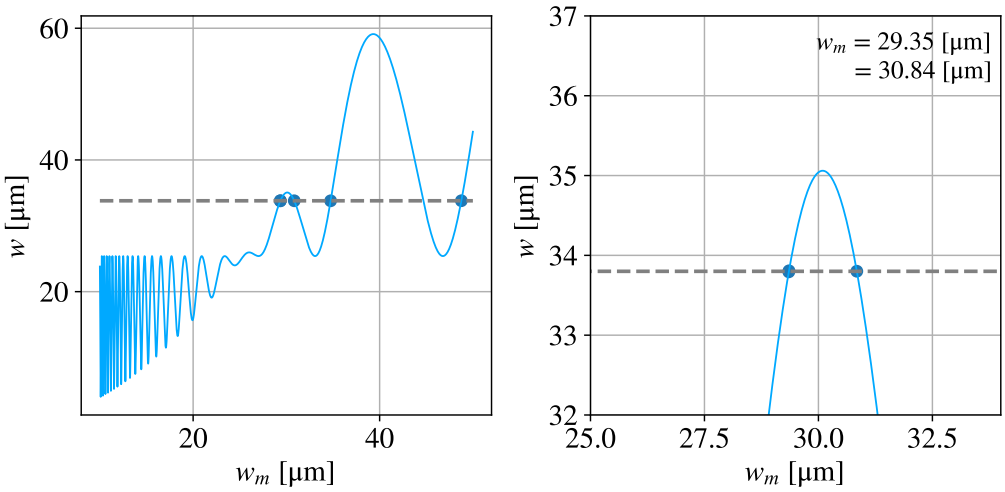


Figure 7.12: Plot of the output spot size w as a function of the matched spot size w_m for a channel length of $L = 27\,\text{mm}$ and an input spot size of $w_0 = 25.4\,\mu\text{m}$. a) Shows the predicted output spot size w as a function of the matched spot size w_m, where the dashed grey line shows the measured output spot size $w = w_{0,\text{guided}} = 33.8\,\mu\text{m}$ and the blue dots show the intersection between the measured value and the model. b) Shows the same data as a) but plotted on a different scale to show the region where the measured value first intersects with the spot sizes predicted by Eq. 1.19. From these two intersection points, we estimate that the minimum matched spot size was $\sim 30\,\mu\text{m}$.

spot size of the channel w_m. Figure 7.12 shows the predicted output spot size w as a function of the matched spot size w_m given by Eq. 1.19, for a channel length of $L = 27\,\text{mm}$ (calculated in Section 7.4.4). By looking at where the measured value of the output guided spot size $w_{0,\text{guided}}$ is equal to the spot size w predicted by Eq. 1.19, we estimate that the matched spot size of the channel was on the order of $\sim 30\,\mu\text{m}$.

7.4.4 Plasma Fluorescence

A camera positioned on the roof of the chamber, looking directly down onto the target, was used to record the fluorescence of the plasma. This was primarily used to optimise the position the gas jet, and to determine the position of the start of the longitudinal channel with respect to the start of the gas cell. The fluorescence was recorded for a scan of different axicon z stage positions. As the axicon was translated downstream (increasing z position), the start of the channel also

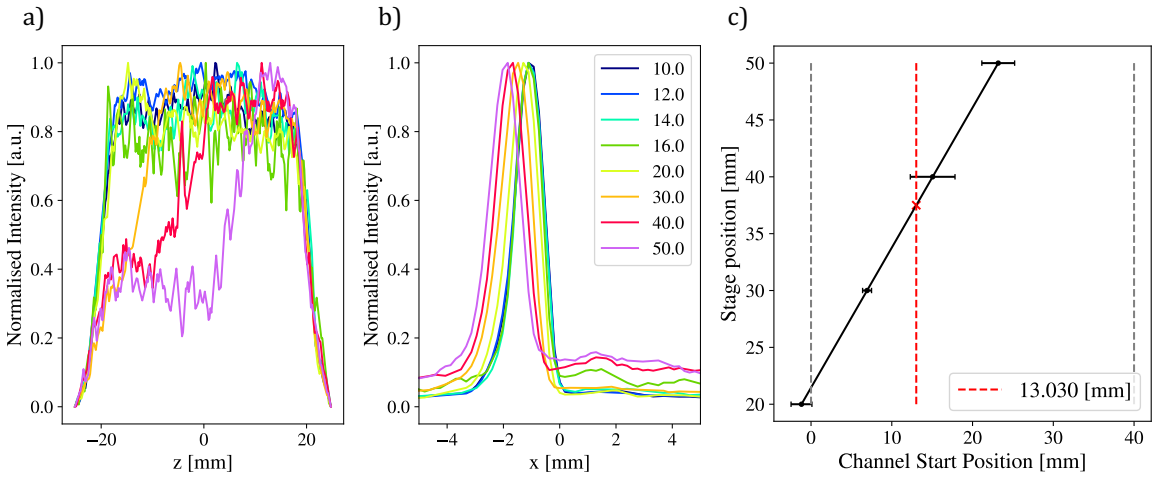


Figure 7.13: Fluorescence measurement to determine the start position of the axicon plasma. a) Shows integrated fluorescence intensity against longitudinal z position for different axicon z stage positions, with the same legend as b), which shows the integrated fluorescence intensity against transverse x position for different axicon z stage positions; c) shows the calibrated start position of the channel for different axicon z stage positions. Dashed grey lines show the start and end of the gas jet, the red dashed line shows the calculated start position of the channel $z_{\text{axi}} = 13.03\,\text{mm}$ for the axicon $z = 37.5\,\text{mm}$ stage position used during the final data-taking runs.

moved downstream with respect to the jet (increasing z_{start} position). Figure 7.13a) shows an integrated lineout of the fluorescence intensity for different axicon z positions. The start position of the channel was estimated from the integrated lineout of the plasma fluorescence intensity by finding the points at which the intensity was half of the maximum intensity. It was possible to vary the length of the longitudinal guiding channel from having a channel that spanned the entire jet for the stage position of $z = 20\,\text{mm}$, to a channel that was less than half of the length of the jet for the maximum stage position of $z = 50\,\text{mm}$. A linear fit was applied to a scatter plot of the z position of the stage against the start position of the channel z_{start}. The correlation between the axicon stage position and the start position of the channel is shown in Figure 7.13c). The dashed grey lines show the start and end of the gas jet. The dashed red line shows the channel start position $z_{\text{axi}} = 13.03\,\text{mm}$ for the axicon stage position $z = 37.5\,\text{mm}$ that was used during the last three crucial data-taking runs. This gives a channel length of $L = 27\,\text{mm}$. It should be noted that the channel also shifted transversely across the jet by about 1 mm for the total range of the axicon z stage. The axicon remained at $z = 37.5\,\text{mm}$ for the majority of the data taking and the drive beam was aligned to the axicon beam at regular intervals, although this shift might have been the cause of a slight offset between the focal position of the injector beam and the axicon guiding channel.

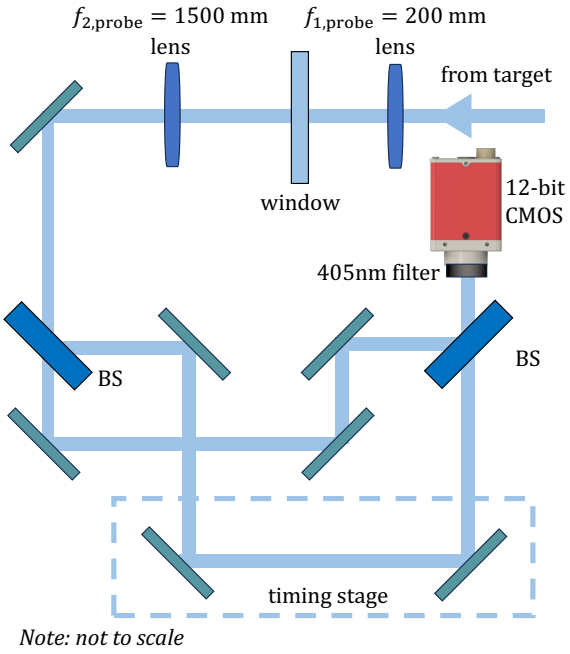


Figure 7.14: Schematic of the transverse interferometry diagnostic. The probe beam (after propagation through the target and reflection off a 400 nm dielectric mirror) was recollimated using a $f_{1,\text{probe}} = 200$ mm lens, before exiting the vacuum chamber through a BK7 window. The probe beam then passed through another lens with a focal length of $f_{2,\text{probe}} = 1500$ mm, which was the distance to the sensor of a 12-bit CMOS camera with a 405 nm bandpass filter. The two arms were split at the first beamsplitter (BS) and recombined at the second forming a Mach-Zehnder interferometer. One arm contained a timing stage to optimise interference between the two arms.

7.4.5 Transverse Interferometry

A folded wavefront Mach-Zehnder interferometer was used to probe the crossed HOFI channels. A schematic of the interferometer is shown in Figure 7.14. The probe beam propagated transversely across the target (co-linearly with the injector pulse). After interaction with the target, the 400 nm probe beam reflected off a 400 nm dielectric mirror, which was used to dump the transmitted 800 nm light from the injector beam. A pair of lenses was used to image the plane of the plasma onto the plane of the detector, with a magnification $\sim 7.5\times$. A beamsplitter was used to split the beam into two identical copies, which were offset transversely from one another, before recombining at a second beamsplitter. A 12-bit CMOS camera from Allied Vision (Mako G-234B) was used to image the interference between the two copies of the beam. The transverse offset between the two beams meant that the portion of the beam which acquired a phase shift from propagation through the plasma was interfered with a portion of the other beam that had not propagated through the plasma and therefore did not have this phase shift. Each arm therefore acted as a reference for the other arm. This led to two copies of the phase shift

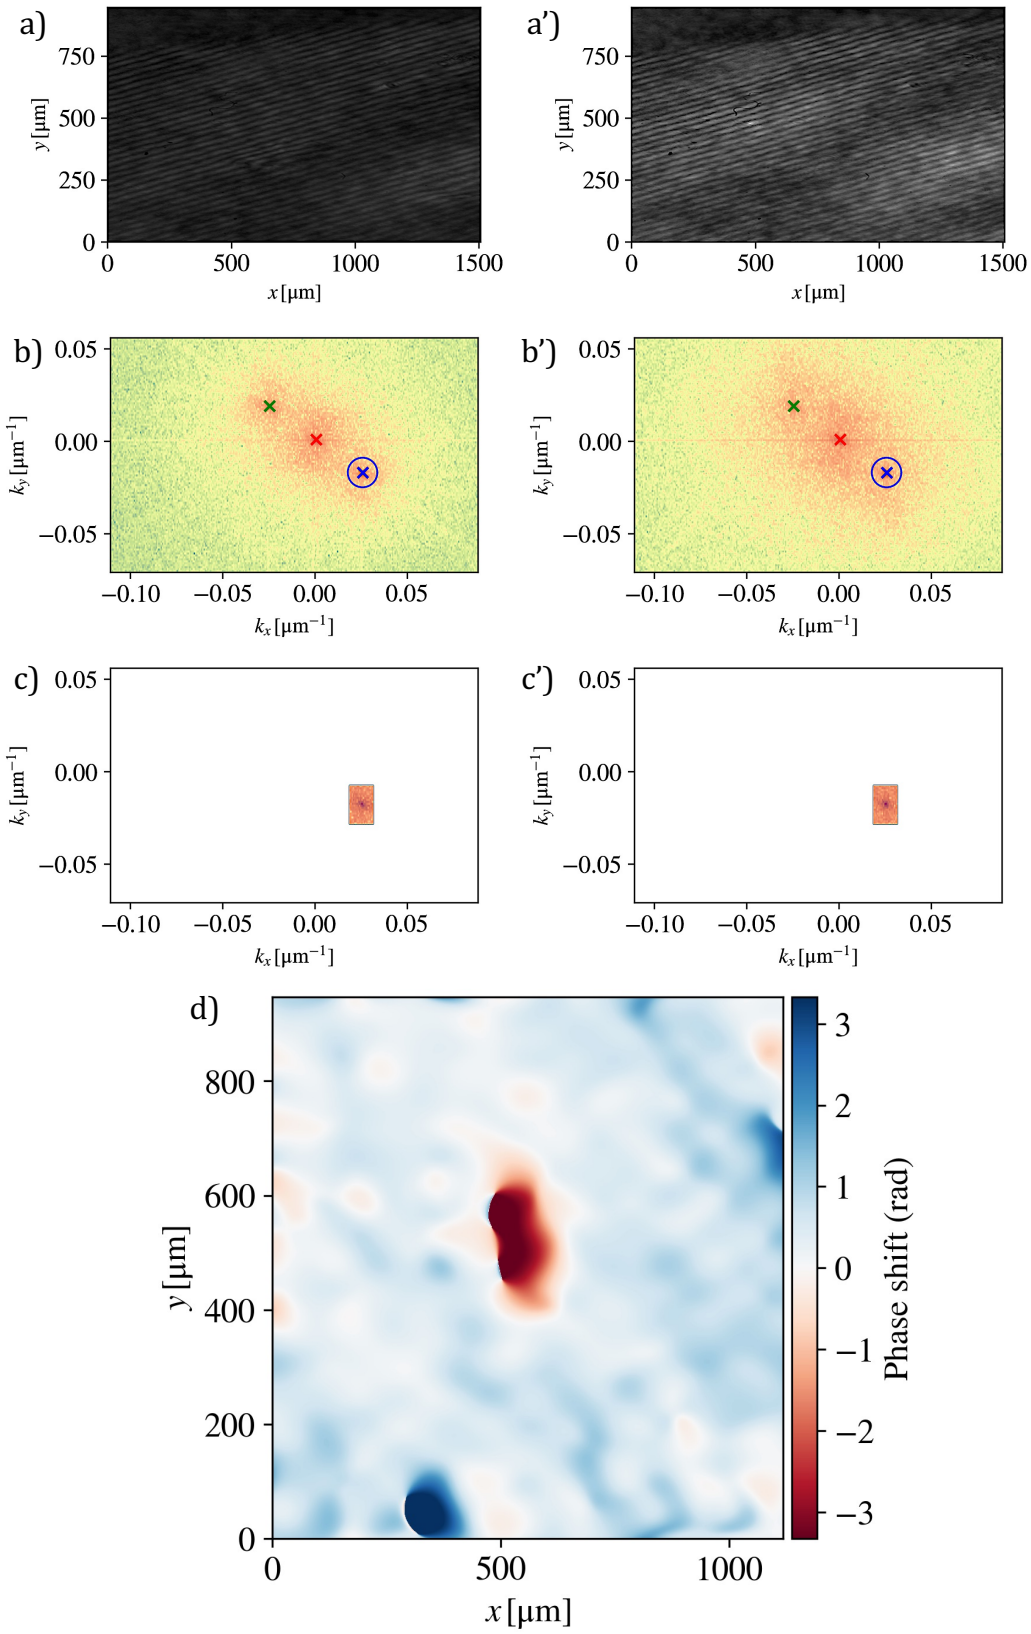


Figure 7.15: Exemplary analysis of the data from the transverse interferometry diagnostic. For the background signal a) shows the raw interferogram; b) the Fourier transformed data; c) a 2D window was applied to the data to isolate a first-order sideband. The corresponding plots for the data are shown in a′)–c′). d) Shows the final background-subtracted phase map.

from the plasma. The final resolution of the probe data was 0.7772 μm/pixel.

Codes that were written for previous experiments by Dr. Alex Picksley and Darren Chan were modified by the author to analyse the parallel probe data. A detailed description of the methods used can be found in previous work,[125,128] therefore we only present a brief summary here. Figure 7.15 shows the analysis of the transverse interferometry data. Figure 7.15a) and Figure 7.15a′) show the raw interferograms for the background and signal data. The background data was the nearest shot taken without gas. The interferogram was transformed to the frequency domain as shown in Figure 7.15b) and Figure 7.15b′). The zeroth-order central fringe containing the background information is plotted in red. The two sidebands containing information related to the plasma phase shift are plotted in green (negative sideband) and blue (positive sideband). The positive sideband (also circled) was isolated using a 2D rectangular window to set all Fourier-transformed data to zero outside this window. The results of applying this filter to both the background and signal data are shown in Figures 7.15c) and 7.15c′), respectively. The windowed data was then inverse Fourier transformed back into the spatial domain to give a 2D map of complex values. Taking the argument of the image and unwrapping this angle gave the 2D phase map, shown in Figure 7.15d).

The probe data was mostly used for alignment between shots. Clearly shown in red in Figure 7.15d) is the HOFI plasma from the transverse injector beam. As already mentioned, there were two copies of the phase shift; the second copy can be seen in blue at the very bottom of the phase map. The phase shift can be estimated from

$$\Delta\phi = \frac{2\pi}{\lambda_0}[\eta - 1]L, \tag{7.12}$$

where L is the distance travelled through a medium with a refractive index η. For a plasma

$$\eta = \left[1 - \left(\frac{\omega_p}{\omega_0}\right)^2\right]^{1/2}, \tag{7.13}$$

$$\simeq 1 - \frac{1}{2}\left(\frac{\omega_p}{\omega_0}\right)^2 + \cdots, \tag{7.14}$$

$$\simeq 1 - \frac{1}{2}\frac{n_\mathrm{e}}{n_\mathrm{c}} + \cdots, \tag{7.15}$$

such that

$$|\Delta\phi| \simeq \frac{\pi}{\lambda_0}\frac{n_\mathrm{e}}{n_\mathrm{c}}L, \tag{7.16}$$

where n_e is the plasma density and n_c is the critical density. The critical density n_c for a probe beam with $\lambda_0 = 400\,\text{nm}$ is

$$n_c = \frac{m_e \varepsilon_0}{e^2}\left(\frac{2\pi c}{\lambda_0}\right)^2, \tag{7.17}$$

$$= 6.97 \times 10^{21}\,\text{cm}^{-3}. \tag{7.18}$$

The probe co-propagated parallel to the transverse injector beam, perpendicular to the longitudinal guiding channel. From previous work, we can estimate that the longitudinal axicon-generated HOFI channel will have a total diameter $L_{\text{guide}} = 100\,\mu\text{m}$.[128] This gives a phase shift of

$$|\Delta\phi| \simeq \frac{\pi}{\lambda_0}\frac{n_e}{n_c}L, \tag{7.19}$$

$$\simeq 0.06\,\text{rad}, \tag{7.20}$$

for an estimated on-axis channel density of $\sim 5 \times 10^{17}\,\text{cm}^{-3}$. This very approximate estimate suggests that the longitudinal HOFI channel therefore cannot be seen in the phase map shown in Figure 7.15d), due to the fact that the estimated phase shift from the longitudinal guiding channel was an order of magnitude smaller than the phase shift due to noise, which was $\simeq 0.3\,\text{rad}$.

7.4.6 Electron Spectrometer

The energy spectrum, divergence and charge of accelerated electrons are key measures of the performance of the LWFA. An electron spectrometer can be used to deduce these bunch parameters. The electron spectrometer was designed and built by Warren Wang, with the help of Dr. James Chappell.[310] Analysis codes used to perform the energy, charge and divergence calculations were written by Warren Wang, based on previous work by Dr. Gavin Cheung.[311] A schematic of the electron spectrometer is shown in Figure 7.16a). Two permanent dipole magnets deflected electrons onto a set of four Kodak Lanex Regular screens, placed to minimise the energy uncertainty resulting from vertical pointing variations in the energy range $80-350\,\text{MeV}$. Lanex screens are primarily composed of gadolinium oxysulphide phosphor ($Gd_2O_2S:Tb$), which fluoresces at 545 nm when subject to ionising radiation. This light was captured by four different 16-bit sCMOS cameras from Andor (Neo 5.5) positioned outside the chamber, each looking at a different screen via BK7 optical windows. Central wavelength filters at 546 nm were used to block out all stray light not emitted by the Lanex; these filters had a 10 nm full width-half

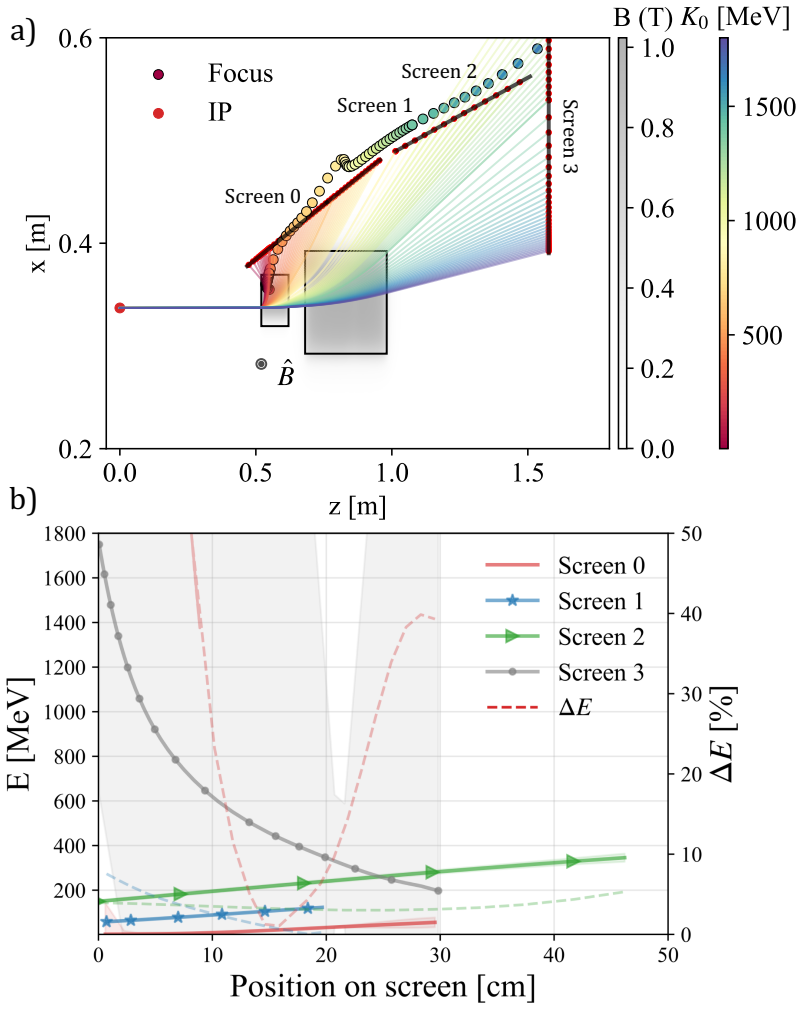


Figure 7.16: Schematic of the electron spectrometer with simulated electron trajectories and calibration curves. a) After the interaction point (IP) electrons propagated through two permanent dipole magnets (black boxes). The magnetic fields of the magnets (grey) deflected the electrons transversely according to their momentum. Electron trajectories are plotted and colour-coded according to energy. b) The simulated electron trajectories were used to convert the position on four different Lanex screens to energy of the incident electrons.

maximum and a transmission of approximately 98%.

Pre-Processing

Pre-processing of the electron spectrometer data was performed by Warren Wang and the author. The raw images of the Lanex screens were first cropped to only show the region of the Lanex itself. A distortion correction was applied by using a chequerboard reference image to calculate the projection matrix. The undistorted image was then background-subtracted. This was done by selecting a region of the Lanex screen where there was never an obvious fluorescence signal from the electron beam. A histogram of the counts was fitted with a Gaussian to determine the mean number of counts. A reference image taken during each run was scaled to this number of counts and was subtracted from the data.

Energy Calibration

Using the measured magnetic field maps of the permanent dipole magnets, the electron trajectories were determined by the Lorentz equation.

$$\mathbf{F} = \frac{d(\gamma m_e \boldsymbol{v})}{dt} = -e(\mathbf{E} + \boldsymbol{v} \times \mathbf{B}). \tag{7.21}$$

Using the scaling $\xi = \boldsymbol{p}/m_e c$, such that $\gamma = \sqrt{1+|\xi|^2}$, the equation of motion can instead be written as a pair of first order differential equations[311]

$$\frac{d\boldsymbol{x}}{dt} = \frac{c\xi}{\sqrt{1+|\xi|^2}}, \tag{7.22}$$

$$\frac{d\xi}{dt} = -\frac{e}{m_e}\left(\frac{\mathbf{E} + \boldsymbol{v} \times \mathbf{B}}{\sqrt{1+|\xi|^2}}\right), \tag{7.23}$$

where in this case $\mathbf{E} = 0$. A fourth-order Runge Kutta method was used to solve the system of ordinary differential equations at each timestep. The electron trajectory was calculated until it intersected with one of the Lanex screens. A set of calibration curves were generated from the list of intersection positions on the Lanex screens, corresponding to electron trajectories which have predetermined initial kinetic energies upon exiting the target. These calibration curves are shown in Figure 7.16b). These calibration curves were then used to convert the longitudinal z position on each Lanex screen to the corresponding initial electron energy.

Divergence Calibration

The transverse x position of an electron on the Lanex screen can be converted to a divergence angle, given that the total path length $\ell(z)$ of the the trajectory is known. The total path length $\ell(z)$ depends on the longitudinal z position on the Lanex screen and therefore on the initial energy E_0 according to the trajectories calculated from Eqs. 7.22 and 7.23, such that the conversion factor will vary for different final electrons energies. The divergence angle θ can therefore be estimated for any position on the Lanex by

$$\theta(x,z) = \tan^{-1}\left(\frac{x}{\ell(z)}\right), \tag{7.24}$$

where $\ell(z)$ is the calculated trajectory length for a final screen position z. It should be noted that as z increases, the trajectories become longer, such that for the same width of screen the divergence becomes smaller. This means that after the divergence calibration the screen narrows towards the high-energy end. Since there was always a blank region at the wider edges, the image was re-cropped to a uniform divergence width.

Charge Calibration

The scaling between the camera counts and the charge incident on the Lanex screen was estimated according to[311–318]

$$\frac{dQ}{dE}(E_0) \approx \left(\frac{\mathscr{C}(E_0)}{\delta s}\right) \times \left(\frac{ds}{dE}\right) \times \left(\frac{\mathscr{Q}\, T\, \Omega \cos[\theta_{\text{out}}]}{\mathscr{S}}\right)^{-1} \times \left(\frac{\kappa_{\text{L}}}{\cos[\theta_{\text{in}}]} \exp\left[-\frac{t_{\text{d}}}{\tau_{\text{L}}}\right]\right)^{-1}, \tag{7.25}$$

where $\mathscr{C}(E_0)$ is the number of counts for a region on the screen between $s \to \delta s$, where s is the longitudinal position on the screen, assuming that this region corresponds to the central energy E_0. The energy dispersion ds/dE can be calculated from the calibration shown in Figure 7.16b). The central term in this equation accounts for the quantum efficiency $\mathscr{Q} = 0.57$ and sensitivity $\mathscr{S} = 0.49$ of the 16-bit sCMOS cameras from Andor (Neo 5.5) at 546 nm. The total transmission of the intermediate windows and filters T and the solid collection angle Ω were used to estimate the fraction of photons emitted by the screen that reach the detector. Assuming a Lambertian emission from the Lanex, the angle θ_{out} accounts for the decrease in intensity for viewing angles which were not directly perpendicular to the screen. The final term accounts for the emission of the Lanex. The total number of photons emitted by the screen per incident unit charge per unit

solid angle $\kappa_L = 6.95 \times 10^9$ photons/sr/pC was taken from an absolute calibration of the Kodak Lanex Regular screen.[316] The angle θ_{in} was the angle of the incident electron at a particular energy E_0, calculated from the intersection of the trajectory with the Lanex screen. The time of detection $t_d = 0$ and the fluorescence decay lifetime $\tau_L = 660\,\mu s$ of the Lanex, such that the exponential term was equal to unity in our case.

The final calibrations for each screen are:

Screen	Energy Range ΔE	Divergence Range $\Delta\theta$	Charge Conversion $dQ/d\mathscr{C}$
Screen 0	$2.63 - 57.0\,\text{MeV}$	$\pm 0.074\,\text{rad}$	$5.10 \times 10^{-7}\,\text{pC/count}$
Screen 1	$57.0 - 110.2\,\text{MeV}$	$\pm 0.061\,\text{rad}$	$8.85 \times 10^{-7}\,\text{pC/count}$
Screen 2	$149.9 - 358.9\,\text{MeV}$	$\pm 0.038\,\text{rad}$	$6.57 \times 10^{-7}\,\text{pC/count}$
Screen 3	$191.6 - 599.4\,\text{MeV}$	$\pm 0.025\,\text{rad}$	$5.36 \times 10^{-7}\,\text{pC/count}$

Table 7.1: Table of electron spectrometer calibrations for each screen

Betatron Model

After these calibrations were applied, the electron spectrometer data gave us a 2D map of the charge density $Q\,[\text{pC mrad}^{-1}\,\text{MeV}^{-1}]$ plotted for different electron energies and divergences. As we will see in Section 7.5, for the condition where we have both the longitudinal guiding channel and the transverse injector pulse, strong energy-divergence oscillations were observed in the electron spectrometer data. These oscillations have been previously observed by others and are generally attributed to strong betatron oscillations of the accelerated electrons.[319,320] Using simple models based on the electron motion, presented by Glinec *et al.*,[319] we fit the observed electron oscillations to this model in order to estimate key parameters related to the betatron oscillations of the accelerated electron bunch. In Section 7.4.6, we will explain all assumptions of such a model in order to understand its limitations.

The equation of motion of the electrons is given by[319]

$$m\frac{\mathrm{d}\gamma\mathbf{v}}{\mathrm{d}t} = e\mathbf{E}_z - \frac{m\omega_p^2\mathbf{r}_\perp}{2}, \tag{7.26}$$

splitting this vectorially, in the longitudinal direction we obtain

$$m\frac{\mathrm{d}\gamma v_z}{\mathrm{d}t} = e\mathbf{E}_z, \tag{7.27}$$

which may be integrated for the energy gain

$$\gamma\beta = \frac{eE_z}{mc}t + \gamma_0\beta_0. \tag{7.28}$$

In the transverse direction, we obtain

$$m\frac{\mathrm{d}\gamma v_\perp}{\mathrm{d}t} = -\frac{m\omega_\mathrm{p}^2 r}{2}, \tag{7.29}$$

which describes the betatron motion. These radial oscillations obey

$$(ut+v)\frac{\mathrm{d}^2 r}{\mathrm{d}t^2} + u\frac{\mathrm{d}r}{\mathrm{d}t} + wr = 0, \tag{7.30}$$

where $u = eE_z/mc$, $v = \gamma_0\beta_0$ and $w = \omega_\mathrm{p}^2/2$. An analytical solution to this differential equation exists if $\sqrt{2\gamma_0\beta_0}E_0/E_z \gg 1$, in which case the output angle $\theta \approx v_r/v_z \approx c^{-1}\mathrm{d}r/\mathrm{d}t$ simplifies to[319]

$$\theta(\gamma) \simeq \frac{\theta_0}{\pi}\frac{(\gamma_0\beta_0)^{1/4}}{(\gamma\beta)^{3/4}}\sin\left\{\frac{E_0}{E_z}\left(\sqrt{2\gamma\beta} - \sqrt{\gamma_0\beta_0}\right)\right\} + \Theta, \tag{7.31}$$

where $E_0 = mc\omega_\mathrm{p}/e$ and $\theta_0 = \pi\omega_\mathrm{p}r_0/2c$. This allows us to fit three independent parameters that describe the betatron motion. These parameters are:

1. γ_0, which describes the initial energy of the electron and can be used to estimate the electron density n_e,
2. θ_0, which describes the amplitude of the betatron oscillations and was used to deduce the initial offset of the electron r_0,
3. E_0/E_z, which can be used to calculate the average accelerating field E_z.

We also fit a fourth parameter Θ which was a constant offset for all energies. This allows us to account for the fact that the electron spectrometer data was not necessarily centred transversely, although this is omitted in Glinec *et al.* We will now describe how the model was fitted to the experimental data.

Fitting the Betatron Model Firstly, the energy axis E[MeV] was normalised for

$$\gamma \simeq \frac{E[\mathrm{MeV}]}{0.5110}, \tag{7.32}$$

which can also be used to deduce

$$\beta = \sqrt{1 - \frac{1}{\gamma^2}}. \tag{7.33}$$

For each energy slice in the 2D electron spectrometer data, the center weighted average divergence angle was found

$$\theta_{\mathrm{COM}}(\gamma) = \frac{\int_{-\infty}^{\infty} Q(\gamma,\theta)\,\theta\,\mathrm{d}\theta}{\int_{-\infty}^{\infty} Q(\gamma,\theta)\,\mathrm{d}\theta}. \tag{7.34}$$

Eq. 7.31 was then fitted to the data using the Python Scipy `curve_fit` function to find the parameters γ_0, θ_0 and E_0/E_z. The fitted parameters were then used to check that the condition $\sqrt{2\gamma_0\beta_0}E_0/E_z \gg 1$ was satisfied. These three parameters were then used to calculate key wakefield parameters as follows.

Accounting for erosion of the laser pulse as it drives the wake[321] and assuming that γ_0 is determined by the trapping condition for a background electron, such that the initial relativistic gamma factor of an individual electron is equal to the relativistic gamma factor of the plasma wave, we find from the wake phase velocity

$$v_{\mathrm{p}} \simeq c\left(1 - \frac{3\omega_{\mathrm{p}}^2}{2\omega_0^2}\right) = c\left(1 - \frac{3\gamma_0^2}{2}\right), \tag{7.35}$$

such that the fitted parameter $\gamma_0 = \omega_0/\sqrt{3}\omega_{\mathrm{p}}$ can be used, in theory, to deduce the electron density. Since the central angular frequency of the laser pulse $\omega_0 = 2\pi c/\lambda_0$ and the plasma frequency is

$$\omega_{\mathrm{p}} = \sqrt{\frac{n_{\mathrm{e}}e^2}{m_e\varepsilon_0}}, \tag{7.36}$$

we can note that

$$n_{\mathrm{e}} = \frac{\omega_{\mathrm{p}}^2 m_e \varepsilon_0}{e^2} = \frac{\omega_0^2 m_e \varepsilon_0}{3\gamma_0^2 e^2}, \tag{7.37}$$

and we can see that electron density n_{e} can be estimated from the initial electron energy γ_0. The fit of θ_0 can be used directly to calculate the injection radius r_0 as

$$r_0 = \frac{2c\theta_0}{\pi\omega_{\mathrm{p}}}, \tag{7.38}$$

and noting that $E_0 = mc\omega_{\mathrm{p}}/e$, we can calculate the accelerating field E_z from the fitted parameter

E_0/E_z as

$$E_z = \frac{1}{E_0/E_z}\left(\frac{mc\omega_p}{e}\right). \tag{7.39}$$

LIMITATIONS OF THE BETATRON MODEL This approach is based on solving the equation of motion for a single electron in a constant accelerating field and radial restoring force. It is therefore a simplified model, and the assumptions made are as follows.

The first assumption is that the accelerator operates in the bubble regime. Eq. 7.26 describes the fields within the bubble, where the longitudinal field is constant and the radial field is proportional to the offset. In linear or quasi-linear regimes, the longitudinal accelerating fields are sinusoidal with respect to the propagation axis z, such that this assumption no longer holds. The parameters of the laser pulse used in our experiment comparable to Glinec *et al.*,[319] justifying the validity of the first approximation.

Parameter	Glinec *et al.*	Astra-Gemini
λ_0	820 nm	800 nm
τ_p	30 fs	45 fs
$\mathscr{E}$	1.3 J	6.9 J
$f/$	$f/18$	$f/20$
w_0	18 µm	25 µm
I	$2.55 \times 10^{18}\,\mathrm{W\,cm^{-2}}$	$7.81 \times 10^{18}\,\mathrm{W\,cm^{-2}}$
a_0	1.09	1.91
n_e	$6 \times 10^{18}\,\mathrm{cm^{-3}}$	$\sim 1 \times 10^{17}\,\mathrm{cm^{-3}}$
P/P_{crit}	1.343	0.475

Table 7.2: Comparison of Astra-Gemini laser pulse parameters with Glinec *et al.* The value of the intensity I for Glinec *et al.* given in this table is our calculation from their laser pulse energy, spot size and pulse duration assuming a bi-Gaussian laser pulse, in their paper the intensity is stated to be $I \sim 3 \times 10^{18}\,\mathrm{W\,cm^{-2}}$.[319] The value of the normalised vector potential a_0 for Glinec *et al.* given in this table was calculated according to Eq. 1.8, this is different to their stated value of 1.3. The laser pulse power is given by $P = 2\sqrt{\ln 2/\pi}(\mathscr{E}/\tau_p)$ and the critical power P_{crit} was calculated using Eq. 1.12. Unlike Glinec *et al.* we were not in a regime where self-focusing was likely to occur.

The second approximation of the model by Glinec *et al.* is that we were operating in the paraxial regime.[319] The focusing geometry employed during the Astra-Gemini experiment was certainly in the paraxial regime (with a larger focal spot size and longer focal length compared to Glinec *et al.*) such that we can assume that this assumption was also valid.

The third approximation is that the longitudinal acceleration is constant throughout the accelerator. This neglects any dephasing, depletion or other propagation effects of the laser pulse. In Glinec *et al.* the parameters calculated from the model were compared with the results of

particle-in-cell (PIC) simulations studying injection and acceleration for the same parameters. Reasonable agreement was found between the results of the PIC simulations and the analytic model. In our Astra-Gemini experiment, our laser pulse propagated in low-density plasma channels. Although we were unable to measure the on-axis electron density n_e during the experiment, it was likely to be lower than the $n_e = 6 \times 10^{18}\,\mathrm{cm}^{-3}$ of Glinec *et al.* An assumed density of $n_e = 1 \times 10^{17}\,\mathrm{cm}^{-3}$ was used as an initial guess in the curve fitting of the oscillations on the electron spectrometer. Therefore, we can be fairly sure that this approximation also holds.

Since both the initial electron offset r_0 and longitudinal accelerating field E_z depend on the plasma angular frequency ω_p in the model for the betatron motion of the electron, the fitting of the initial electron energy γ_0 was crucial; this is used to deduce the plasma angular frequency ω_p in the absence of an accurate measurement of the on-axis plasma density n_e. If there is large error in the fitting of the initial electron energy γ_0 then clearly this will give an inaccurate value of the plasma density n_e, leading to incorrect calculations of the offset radius and accelerating field. In Glinec *et al.* it is stated that the fit is not sensitive to the initial electron energy. Therefore, they use their known value for the on-axis electron plasma density, which they assume to be $n_e = 6 \times 10^{18}\,\mathrm{cm}^{-3}$. This presents a problem since we were unable to measure the on-axis electron density in the channel during the experiment. For now, we fit the initial electron energy γ_0 as a parameter in the model, assuming an initial guess of $n_e = 1 \times 10^{17}\,\mathrm{cm}^{-3}$, similar to the conditions used in our previous studies of truncated-channel injection.[132] A simulations campaign is planned as future work; magnetohydrodynamic (MHD) simulations of the expansion of the longitudinal HOFI channel should allow us to deduce more accurately the on-axis electron density in the channel, improving our estimates of the injection offset r_0 and average accelerating field E_z.

Two further limitations of this model are that it does not account for spot size or transverse oscillations of the laser pulse. Any oscillation of the laser pulse will couple to the betatron motion of the electrons. Again, this will be investigated in a future PIC simulations campaign, to fully understand how the evolution of the drive laser pulse affects the motion of injected electrons and whether charge was lost throughout the acceleration.

7.5 Results

The data for the last 409 shots (corresponding to the three final data-taking runs) of the experiment were analysed and key parameters were extracted. These parameters were analysed in the context of certain independent variables that were varied during the course of the data-taking, namely:

1. The backing pressure of the gas in the target,
2. The focal position of the drive beam,
3. The transverse position of the injector beam focus,
4. The longitudinal position of the axicon beam focus,
5. The energy in the drive beam and
6. The energy in the channel-forming beams.

Dependent variables relating to the propagation of the drive laser pulse, measured using the guiding diagnostic, and the injection and acceleration of electrons, measured using the electron spectrometer, were analysed as described in Sections 7.4.3 and 7.4.6 respectively. The key dependent parameters related to the guiding were:

1. The D4σ_{maj} and D4σ_{min} elliptical spot size axes (described in Section 7.4.2) fitted to the guided mode and,
2. The total counts on the camera within the D4σ spot, as an estimate of the guided energy,

and the key dependent parameters related to the injection and acceleration of electrons were:

1. The total charge measured by the electron spectrometer,
2. The mean energy of the electrons incident on the electron spectrometer,
3. The peak energy of the electrons incident on the electron spectrometer,
4. The parameters estimated from a fitting of the betatron oscillations in the electron spectra, as described in the previous section.

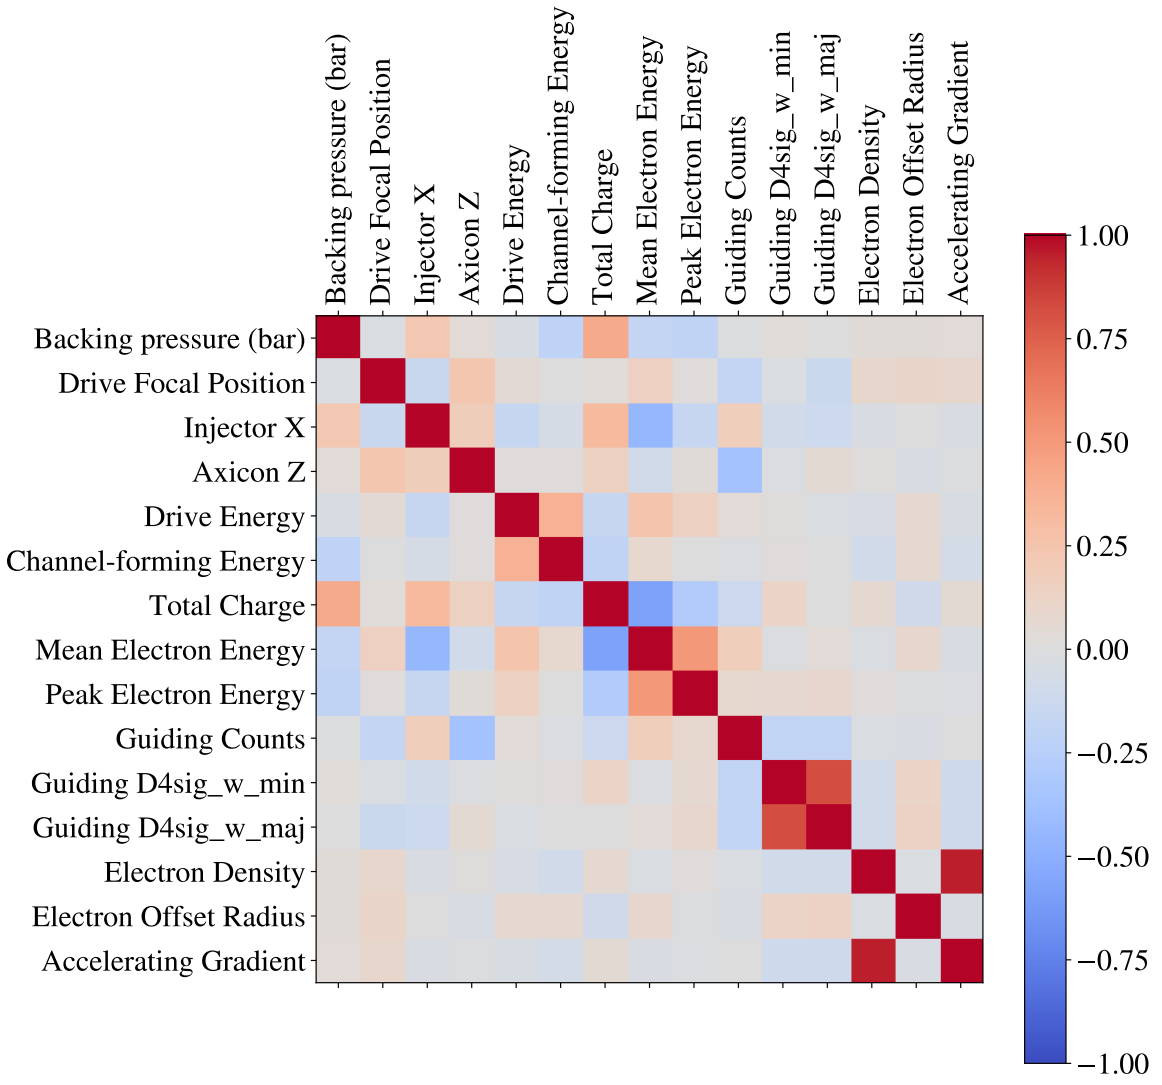


Figure 7.17: Map of Pearson's correlation coefficients between experimental parameters and measured variables, for the last 409 data-taking shots of the experiment. The Pearson correlation coefficient was calculated for each pair of variables to measure the strength of a linear relationship. A value of -1 indicates a perfect negative correlation, a value of 0 indicates no correlation and a value of 1 indicates a perfect positive correlation.

7.5.1 Correlations Between Variables

Figure 7.17 shows the correlation between all independent and dependant variables. The Pearson correlation coefficients were calculated using the Python Pandas' function `corr`[†] for each pair of variables. The Pearson coefficient r is a measure of the strength of a linear relationship between two variables and for a pair of variables (x_i, y_i) is given by[322]

$$r = \frac{\sum (x_i - \bar{x})(y_i - \bar{y})}{\sqrt{\sum (x_i - \bar{x})^2 \sum (y_i - \bar{y})^2}}, \tag{7.40}$$

where $\bar{x}$ is the mean of the set (x_i) and $\bar{y}$ is the mean of the set (y_i). A value of -1 indicates a perfect negative correlation, a value of 0 indicates no correlation and a value of 1 indicates a perfect positive correlation. This can be roughly divided into three categories: i) $|r| < 0.3$ signifies a weak correlation, ii) $0.3 \geq |r| \leq 0.5$ a moderate correlation and iii) $|r| > 0.5$ indicates a strong correlation between the variables. There are a few key points to be noted from Figure 7.17:

- There was a moderate positive correlation between the backing pressure and the total charge injected, suggesting that a higher backing pressure and therefore a higher initial gas density led to higher availability of electrons to be injected into the wakefield.

- There was a strong, inverse correlation between the total charge injected and the mean energy of the electrons, indicating that it was more difficult to accelerate high-charge electron bunches, due to beam-loading effects.

- There was a strong positive correlation between the transverse x position of the injector focus and the total charge injected, suggesting that the amount of charge injected was sensitive to the position of the transverse injector beam (and therefore the transverse injector pulse formed by this beam).

- There was a weak inverse correlation between the longitudinal start position of the axicon focus z_{axi} and the mean electron energy. As z_{axi} increases, the guiding channel was pushed downstream and was therefore shorter, meaning that the distance over which a wakefield was driven was shorter and the average energy gain was lower. This was also reflected as an even weaker negative correlation in the peak electron energy.

[†]https://pandas.pydata.org/docs/reference/api/pandas.DataFrame.corr.html

- There was a moderate inverse correlation between the longitudinal start position of the axicon focus z_{axi} and the guided counts within a D4σ spot size. This suggests that it was more difficult to couple the drive laser pulse into the channel as it was pushed downstream.

- There was a very strong positive correlation between the D4σ_{maj} and D4σ_{min} spot size axes, suggesting that the guided spot was not being squashed in either direction, but instead grew and shrunk circularly as expected.

- In the fitting of the betatron model, there is an almost perfect positive correlation between the estimated on-axis density n_{e} and average accelerating gradient of the wakefield E_z/E_0. This makes sense because we know from Eq. 7.39 that $E_z \propto \omega_{\text{p}}$.

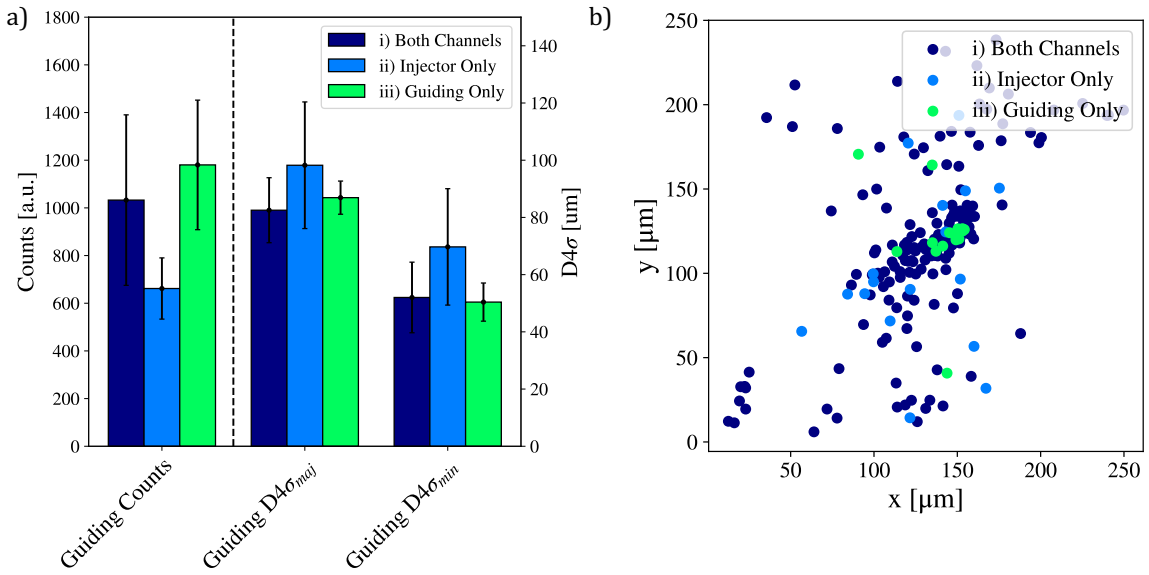


Figure 7.18: Key results from analysis of the data from the guiding diagnostic. The 364 shots were divided according to three conditions: i) a transverse injector HOFI channel which crossed before the start of a longitudinal guiding HOFI channel, ii) only a transverse injector HOFI channel, iii) only a longitudinal guiding HOFI channel. a) Shows the mean number of counts within a D4σ spot size, the mean length of the D4σ_{maj} major spot size axis and the mean length of the D4σ_{min} minor spot size axis and b) shows the centroid position of the guided mode. Error bars in a) show the standard error on the mean.

7.5.2 Guiding

Here we compare results of the guiding analysis for three different cases: i) a transverse injector HOFI channel which crosses before the start of a longitudinal guiding HOFI channel, ii) only a transverse injector HOFI channel and iii) only a longitudinal guiding HOFI channel. First the data was filtered to ignore any shots with a total charge on the electron spectrometer of less than 1 pC. This reduced the 409 shots from the last three data-taking runs to 364 shots.

Figure 7.18a), shows the mean D4σ major and minor axes fitted to the guiding diagnostic data (to the right of the dotted line) and the mean number of counts within this region (to the left of the dotted line). The guiding diagnostic was described in Section 7.4.3 and the method used to fit the D4σ spot size was given in Section 7.4.2. Figure 7.18 shows that that the presence of the transverse injector HOFI channel used to inject electrons had no significant impact on the quality of the guiding in the longitudinal channel. Firstly, from Figure 7.18a), we can see that the size of the guided spot was virtually unchanged, regardless of the presence of the transverse injector channel when we always had a longitudinal guiding channel. The energy within the guided spot when the transverse injector channel was present was 87.5% of the energy in the guided spot with the longitudinal guiding channel alone, suggesting that approximately 12.5% of the drive beam energy was lost during propagation through a transverse injector channel. A proper throughput calibration and the use of a laser with better energy and pointing stability would have allowed for a better estimate of the energy coupled through the transverse channel.

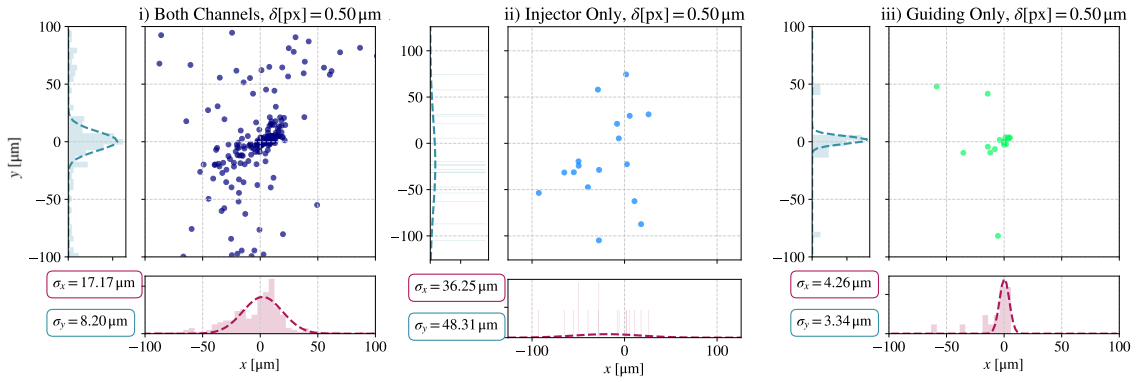


Figure 7.19: Pointing jitter of the guided beam for different channel combinations. The data was divided according to three conditions: i) a transverse injector HOFI channel which crossed before the start of a longitudinal guiding HOFI channel, ii) only a transverse injector HOFI channel, iii) only a longitudinal guiding HOFI channel.

Figure 7.18b) shows the centroid position of the guided beam, with and without the transverse injector and longitudinal guiding channels. A statistical analysis of this data is shown in Figure 7.19. It can be seen that when only the transverse channel was present, i.e. there was no longitudinal guiding channel, the position of the guided drive beam was not localised to the center of the camera, with a centroid variation of $\sim 40\,\mu$m in both directions. When a guiding channel was present, the spatial jitter of the guided drive beam was greatly reduced and was $\sim 2\,\mu$m in the x-direction and $\sim 4\,\mu$m in the y-direction. The spatial jitter was twice as large in the y-direction, suggesting that the transverse position of the longitudinal channel or the point-

ing of the drive beam was varying more in the vertical direction. When both channels were present, the pointing jitter was still greatly reduced compared to the case without a longitudinal guiding channel, with a variation $\sim 17\,\mu$m in the x-direction and $\sim 7\,\mu$m in the y-direction.

7.5.3 Electron Injection and Acceleration

The core aim of this experiment was to study whether crossed HOFI channels can enhance electron injection in LWFAs. As described in Section 7.4.6, the electron spectrometer data was processed and used to extract key parameters pertaining to the accelerated electron bunches. Figure 7.20 compares the processed electron spectrometer data for the three different conditions, namely for i) crossed-channel injection, where a high-charge electron bunch was accelerated to 500 MeV and exhibited characteristic oscillations (see also Figure 7.22), ii) high-charge, but low-energy 200 MeV electron bunches injected with the transverse channel alone and iii) low-charge and low-energy electron bunches injected with the longitudinal channel alone, due to ionisation injection prior to or during propagation of the drive laser pulse into the guiding channel.

A summary of the key parameters extracted from the electron spectrometer data are shown in Figure 7.21. The data to the left of the dashed line in Figure 7.21 show the amount of charge injected. We see that the mean charge injected with a transverse injector channel was 64.5 pC, compared to 20.2 pC for the guiding channel alone. In the crossed channel configuration, the mean charge injected was reduced to 40.1 pC. This was still an order of magnitude higher than our previous experiment.[132] The peak charge injected was highest for the crossed channel configuration at 251.5 pC, with a similar peak charge for the longitudinal guiding channel alone at 247.9 pC. The maximum charge that could be seen with a transverse injector channel alone was 122.6 pC. This was probably due to the fact that without a longitudinal guiding structure, the energy gain of electrons in the wakefield was limited. The electron spectrometer data was only processed for energies greater than 100 MeV, so it may be possible that there was more low-energy charge injected, but this is not of particular relevance.

Considering now the data to the right of the dashed line in Figure 7.21, we see that the average electron energy was slightly lower for the crossed channel configuration. This is shown for both the mean electron energy at 189.9 MeV, taken by averaging the electron signal, and the modal electron energy 263.3 MeV, corresponding to the energy of the peak charge density. The maximum peak energy that could be reached was highest for the crossed channel case at 599.4 MeV, compared to 549.1 MeV for the guiding channel only and 494.2 MeV for the injector

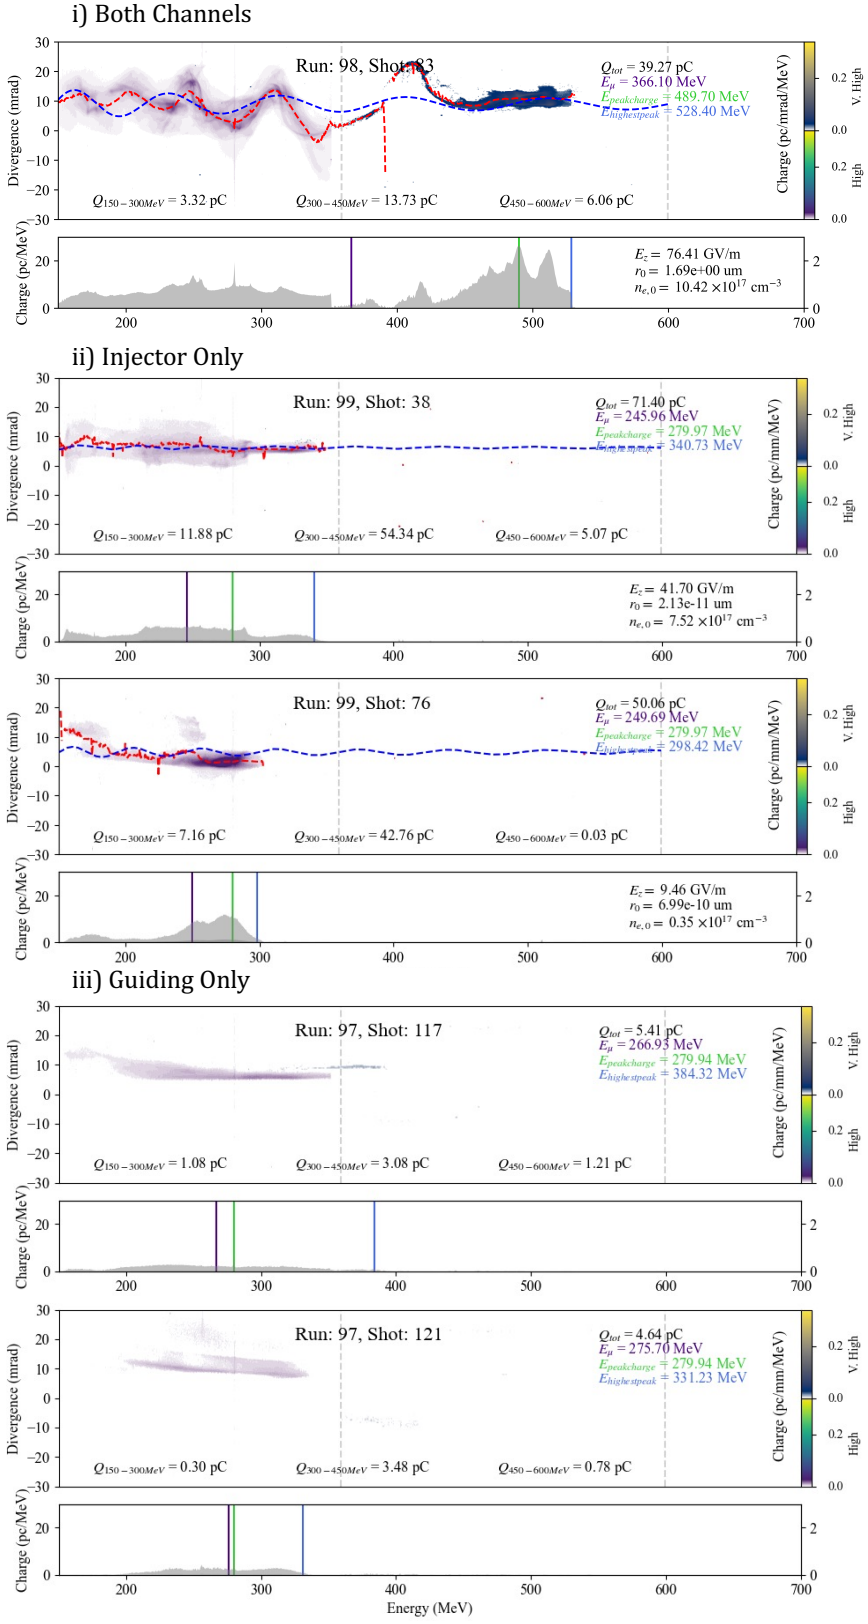


Figure 7.20: Exemplary electron spectrometer data for different channel conditions, where i) shows crossed-channel injection, where a high-charge electron bunch was accelerated to 500 MeV and exhibited characteristic oscillations (see also Figure 7.22), ii) shows high-charge but low-energy 200 MeV electron bunches injected with the transverse channel alone and iii) shows low-charge and low-energy electron bunches injected with the longitudinal channel alone.

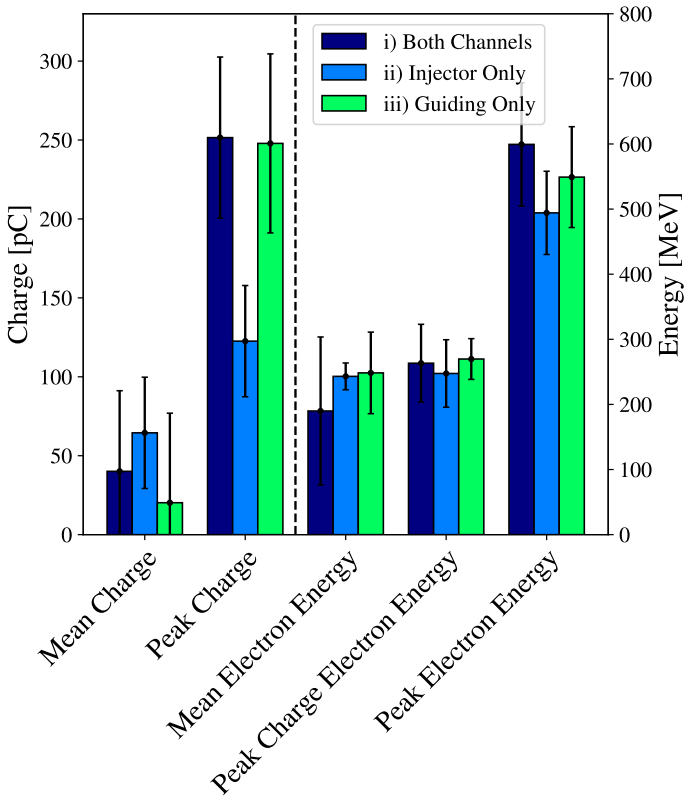


Figure 7.21: Key results from analysis of the electron spectrometer data. The 364 shots were divided according to three conditions: i) a transverse injector HOFI channel which crossed before the start of a longitudinal guiding HOFI channel, ii) only a transverse injector HOFI channel, iii) only a longitudinal guiding HOFI channel. Plotted to the left hand side of the dashed line are the mean and peak charge of the electron bunches produced for these different conditions. To the right hand side of the dashed line are the mean, modal (the energy of the peak charge density), and peak (taken to be the maximum energy where the charge density is greater than 10% of the peak charge density) energies of the electron bunches produced. Error bars show the standard error on the mean.

channel only. This was found by finding the highest energy where the charge density was 10% of the peak charge density (the charge density is plotted as grey integrated lineouts below the 2D electron spectrometer data in Figure 7.20). These three key metrics of electron energy: the mean electron energy, modal electron energy and peak electron energy, are also plotted for the exemplary electron spectrometer data shown in Figure 7.22 as green, purple and blue vertical lines, respectively.

7.5.4 Electron Betatron

In Section 7.4.6, a model for the betatron oscillations of electrons in the wakefield was presented and we attempted to use this model to fit key parameters such as the average longitudinal accelerating field E_z, initial offset radius r_0 and estimated on-axis electron density $n_{e,0}$ to the observed energy-divergence oscillations in the electron spectrometer data. Five exemplary shots which exhibit these oscillations are shown in Figure 7.22. Oscillations were only present with the combination of the two channels and hence the enhanced betatron motion was likely induced by transverse misalignment between the injector and guiding sections. We can see that across these five shots which exhibited clear energy-divergence oscillations, using the betatron model the average total charge was estimated to be $Q_{\text{tot}} = 37.35 \pm 11.21\,\text{pC}$, the average longitudinal accelerating field $E_z = 46.24 \pm 37.14\,\text{GV m}^{-1}$, the average initial electron offset $|r_0| = 1.84 \pm 0.47\,\mu\text{m}$ and the average on-axis electron density $n_{e,0} = 5.68 \pm 6.01 \times 10^{17}\,\text{cm}^{-3}$. The large uncertainties in average fitted values across these five shots (for the same experimental conditions) highlight the issue with such a simplistic model for the electron motion. We can see from Figure 7.22 that the quality of the fits of the simple betatron model to the electron spectrometer data are poor. In particular, the shots shown in Figure 7.22 were all for a fixed backing pressure of 30 bar, such that, as we predicted in Section 7.4.6, the fitting of the on-axis electron density $n_{e,0}$ is not particularly meaningful, with a larger variation in the predicted values than the approximate on-axis density itself. This impacted the fitting of the longitudinal accelerating field E_z, which also had a high standard deviation of almost 100%. As already mentioned, MHD simulations are planned for future work, which should allow us to more accurately deduce the on-axis density and average accelerating gradient.

In addition to the initial offset radius r_0 of injected electrons with respect to the drive laser pulse, there may have also been a misalignment between the injection location and the axis of the longitudinal guiding channel. As already mentioned, this offset could have caused trans-

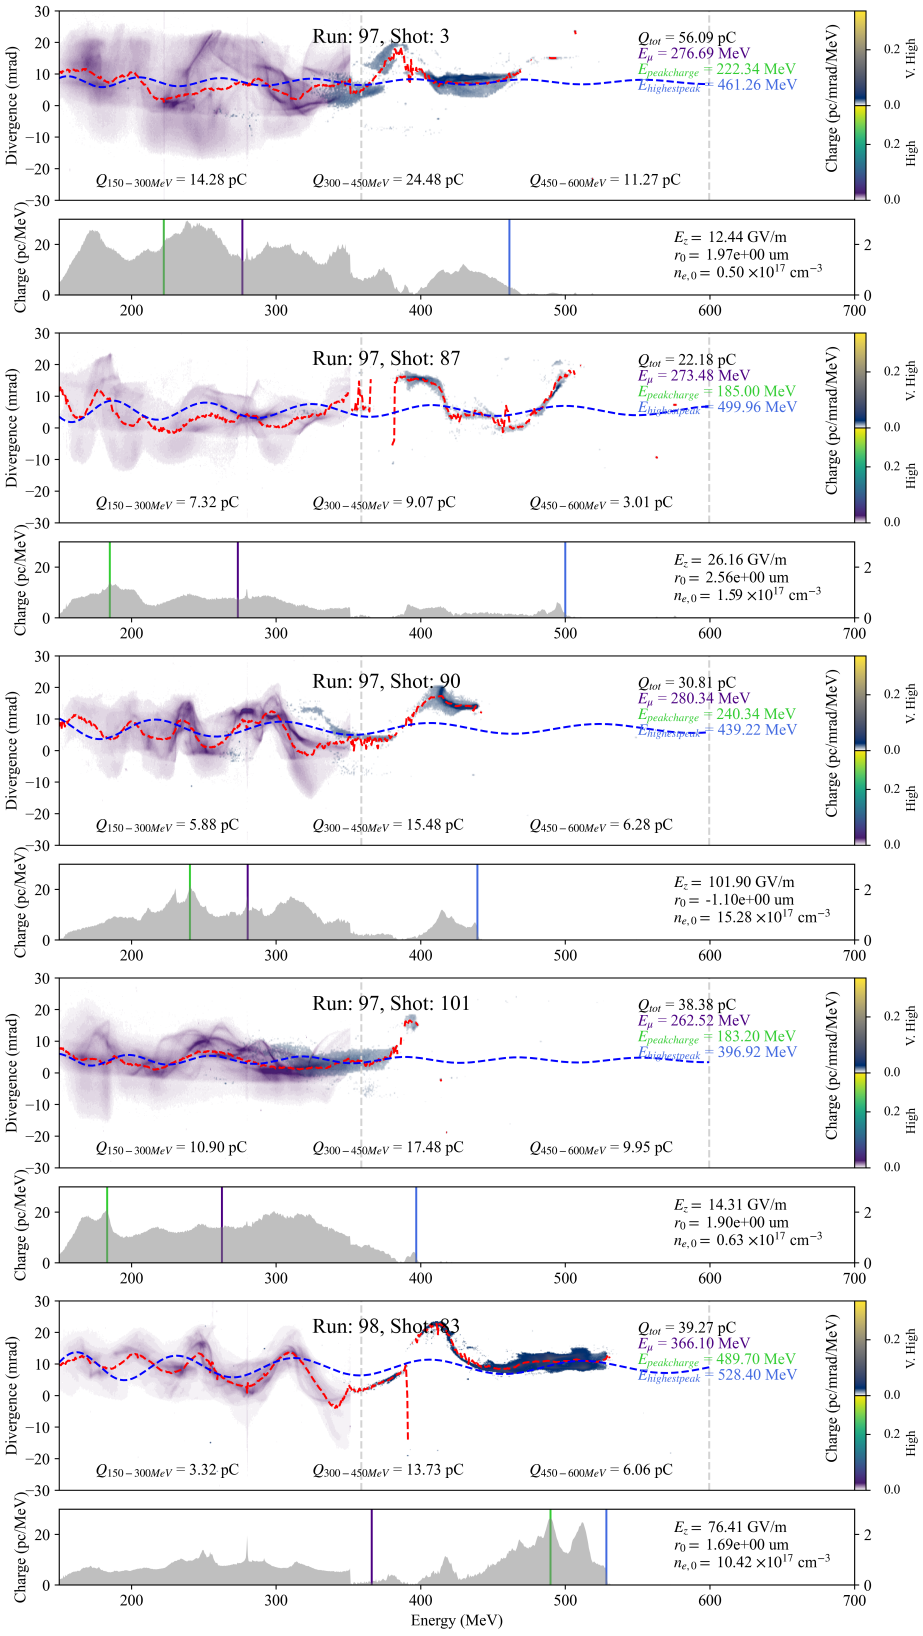


Figure 7.22: Exemplary electron spectrometer data showing energy-divergence oscillations. The center weighted signal $\theta_{\mathrm{COM}}(\gamma)$ is shown as a dashed red line, and the fitted theoretical energy-divergence oscillation $\theta(\gamma)$ is shown as a dashed blue line. The key electron parameters such as the total charge Q_{tot}, mean E_μ, modal $E_{\mathrm{peak\ charge}}$ and peak $E_{\mathrm{highest\ peak}}$ energies are shown in the upper panel and the fitted betatron model parameters, such as the average longitudinal accelerating field E_z, initial offset radius r_0 and estimated on-axis electron density $n_{e,0}$ are shown in the lower panel.

verse oscillations in the centroid position of the drive laser pulse; these oscillations in the laser beam spot size can couple to the electron bunch and further enhance the betatron oscillations of the injected electron bunch compared to the betatron model. The radiation generated through betatron oscillations has potential applications for future light-sources, and if these oscillations can be deliberately instigated and controlled using the scheme proposed here then this is highly relevant for radiation generation in a single stage. However, the origin of these oscillations needs to be further investigated using MHD and PIC simulations. Alternatively, if we wish to perform staging of LWFAs, then this dual-channel scheme allows us to estimate the alignment tolerances. A campaign of PIC simulations which thoroughly investigate these tolerances is planned as future work.

7.6 Conclusions and Future Work

A controlled injection scheme to inject and accelerate 600 MeV electron bunches with $\sim 10-100$ pC charges was demonstrated. A transverse channel generated by HOFI using an astigmatic focus was used to create two longitudinal density shocks prior to the onset of a second guiding channel which was collinear to the target. Using a slit nozzle gas jet that was 40 mm long, we were able to couple both the laser pulse and the injected electrons into a longitudinal HOFI channel, formed by an axicon focus, to accelerate the electrons to up to 600 MeV. This suggests that our accelerating gradient was on the order of $\sim 20\,\mathrm{GV\,m^{-1}}$.

The experimental layout and parameters were described and an overview of the key diagnostics and associated analysis were given. The laser pulse data were analysed to understand the key parameters associated with the three key laser beams incident upon the target; the two channel-forming beams, one focused by an axicon and one focused by a lens to generate an astigmatic focus, forming the longitudinal guiding and transverse injector channels respectively, and the drive beam used to drive a wakefield. This included understanding the laser pulse energy, intensity, pulse duration, relative timings, spot sizes and jitters. We found that the guided spot size was slightly larger than the input spot size of the drive beam, indicating only a slight mismatch between the drive beam and the channel. We found that guiding of the drive beam was not obstructed by the presence of the transverse channel used for electron injection, with only $\sim 12.5\%$ of the drive beam energy lost from propagation through the transverse injector channel and the guided spot size completely unaffected. We demonstrated higher mean and peak

charge of the electron bunches injected with an injector pulse and found that the energy gain of these bunches was comparable or higher than using a longitudinal guiding channel alone. This would again suggest that we were able to couple the laser pulse and injected electrons into the channel successfully. We found that there was a positive correlation between the transverse x position of the injector focus and the total charge injected, suggesting that the amount of charge injected was sensitive to the transverse position of the injector channel, therefore confirming that electrons were injected at the density ramps of the transverse channel. We observed strong energy-divergence oscillations in the electron spectrometer data, only when both channels are present, as is indicative of betatron oscillations of the electron bunch during acceleration. Using a basic model for the betatron motion of an electron in the wakefield, we fitted key parameters of the wakefield such as the average longitudinal accelerating field E_z, initial offset radius r_0 and estimated on-axis electron density $n_{e,0}$, but found that such a model does not provide a useful fit without prior knowledge of the on-axis electron density. Furthermore, that such a model cannot account for the evolution of the drive laser pulse within a guiding channel and therefore cannot account for transverse oscillations of the drive laser pulse due to a transverse offset when entering into the longitudinal guiding channel. Such oscillations of the laser pulse couple to the betatron motion of the electrons and are therefore likely to affect the observed energy-divergence oscillations significantly. However, initial observation of these oscillations could be used to understand the injection mechanism and suggests that crossed-channel schemes for electron injection and acceleration in LWFAs could not only have applications for generating high-charge, high-energy electron bunches, but also could be used to generate radiation in a single stage, as is relevant for applications in future light-sources.

Planned future work includes full magnetohydrodynamic (MHD) simulations to understand the evolution and final structure of the crossed-channels up to the time at which the drive beam enters the target. The results of these simulations will then be used to estimate the on-axis electron density and to define initial plasma density profiles for particle-in-cell (PIC) simulations, which will be used to study the electron injection and acceleration. The relative energy spread of the injected electron bunches is still reasonably high and results from simulations should aid us in understanding both the origin of the electron betatron oscillations and whether the large energy spread is due to multiple injection events. Such simulations could then inform tolerances for the staging of LWFAs and the design of experiments aiming to generate radiation by deliberately

instigating oscillations like those seen in the experiment presented here.

8 | Conclusion

The prospect of accelerating high-quality electron bunches to GeV energies in centimetre-long distances is clearly a promising one, with widespread applications. However, we established that the development of controlled injection techniques remains central to improving the shot-to-shot reproducibility of laser-wakefield accelerators and the quality of the electron bunches that they produce. In this thesis, we focused primarily on optical electron injection techniques due to the high degree of control and tunability that they potentially offer. Chapters 3, 4 and 5 of this thesis studied the unique properties of different ultrashort laser pulses with spatiotemporal couplings and discussed their viability for applications in laser-driven plasma wakefield acceleration, especially for the ionisation injection of electrons. Chapter 6 studied how the wakefield is altered when we allow the laser pulse to take novel polarisation states; specifically we rederived the wakefield quantities following a two-pulse ionisation injection scheme with a radially-polarised injector beam. In Chapter 7 we presented an experiment where we used optically-generated density shocks to demonstrate injection and acceleration of 600 MeV, $\sim 10-100$ pC electron bunches, with spectral brightnesses on the order of $\sim 10\,\mathrm{pC\,MeV^{-1}}$.

In Chapter 2, an overview of the background theory and the methods used to model the propagation of ultrashort laser pulses in the vacuum were given. Starting from Maxwell's equations, we derived the paraxial and modified paraxial equations which govern the propagation of focusing beams in free-space under simplifying assumptions. The Collins method, a diffraction integral written in terms of the ray matrices of an optical system, was presented as a method for propagating polychromatic beams in free-space, including those with spatiotemporal couplings. We discussed the details related to performing these simulations numerically and the simulation grid requirements to properly capture and resolve ultrashort, broadband laser pulses.

Chapter 3 applied these techniques to the first of three phenomena observed for ultrashort laser pulses with spatiotemporal couplings. Flying foci are laser pulse focal spots which appear

to propagate at arbitrary velocities, including those which are sub- or superluminal and those which appear to propagate backwards relative to the source. Flying foci have applications for laser-driven plasma wakefield acceleration since they may be used to drive 'dephasingless' wakefields. One method of generating flying foci is to pass a temporally-chirped laser pulse through a dispersive optic. We derived the fields of the flying focus in the vicinity of the focus and found an analytic expression for the apparent velocity of the focal spot v_{ff} in terms of the chirp χ and the chromaticity of the lens β. We found that focal spots with arbitrary velocities could be generated and there was good agreement between the analytic expression for the focal velocity and the focal velocities calculated from the spatiotemporal fields in the focal region simulated using both an analytically derived expression for the fields and the Collins diffraction integral. We also showed that the analytic expression can be used to calculate flying foci fields for input pulses with arbitrary spectra. Chapter 3 lays a theoretical foundation for being able to implement flying focus techniques experimentally in the Oxford Laser-Plasma Accelerator Group's laboratory, including the dephasingless laser wakefield accelerator scheme which has been presented by other groups. Future investigations may also include studying whether it is possible to generate hydrodynamic optical field ionised channels using flying foci and whether a flying focus could be used to control beam loading in ionisation injection schemes. Both of these studies would need to be performed using particle-in-cell (PIC) codes.

Chapter 4 studied another established technique: the simultaneous space-time focus. We found that, contrary to statements often made in the literature, the degree of axial localisation of the intensity of the simultaneous space-time focus could be entirely explained by the increase in transverse size of the input pulse caused by the spatial chirp. By defining a fair-comparison pulse that undergoes comparative spatial focusing, we showed that i) the fair comparison case shows that no advantage is provided by SSTF and ii) simplistic methods commonly used to characterise focal spots fail when spatiotemporal couplings are present. The lack of localisation of the peak intensity of the simultaneous space-time focus and the residual pulse-front tilt at the focus makes such pulses generally unsuitable for ionisation injection in laser-driven plasma wakefield accelerators. Although we found that the focal volume of a laser pulse cannot be approximated by an ellipsoid when spatiotemporal couplings are present, we plan to investigate whether tilted ellipsoids fitted using principle component analysis can be used to model the focal volume and estimate the degree of different first-order couplings present in the beam. Principal manifolds

could be investigated as a method for modelling beams with arbitrary higher-order spatiotemporal couplings.

An alternative to the linear SSTF scheme was presented in Chapter 5. Instead of chirping the input pulse along one axis, the pulse was radially-chirped in two-dimensions. We found that although the axial localisation of the radial simultaneous space-time focus is worse than the fair-comparison and that the pulse does not quite fully-recompress at the focus, a radial simultaneous space-time focus still undergoes strong temporal compression, and furthermore it does not exhibit pulse-front tilt at the focus. The radial simultaneous space-time focus may therefore have applications for laser machining in cases where we require 'reciprocal' or uniform laser writing, regardless of the scanning direction of the laser pulse. In addition to this, we found that radial simultaneous space-time foci have an arbitrary focal velocity which can be varied using the chirp of the input beam and we plan to investigate the range of focal velocities that can be achieved. An investigation into whether it is possible to derive an analytic expression for the focal velocity of a radial simultaneous space-time focus is planned as future work.

In Chapter 6, we studied whether exotic polarisation states of the laser pulse could be used to enhance ionisation injection and electron trapping in laser-driven plasma wakefield accelerators. We found that although the spot size is larger, the high-intensity region of radially-polarised beams is more axially confined in paraxial focusing regimes. We investigated whether the longitudinal field component of the RPLB could be used to promote electron trapping in laser-driven wakefields. We presented a derivation of the wakefield quantities based on simplified 1D models for the wakefield in a quasistatic approximation, allowing for a longitudinal electric field component of the laser pulse. This preliminary study seemed to suggest that changing the trailing laser pulse to be radially-polarised in a two-pulse ionisation injection scheme could enhance the wakefield and promote electron trapping, and we also showed that it is hypothetically possible to promote electron trapping in laser-driven wakefields driven by a single radially-polarised laser beam. PIC simulations of the ionisation dynamics using radially-polarised laser beams are already underway and we also plan to replicate the analytic study presented here using both quasistatic and non-quasistatic PIC codes, to assess the validity of both the model and the R-2PII technique.

Finally, in Chapter 7 we presented the results of an experimental campaign using the Astra-Gemini laser system at Rutherford Appleton Laboratories. We demonstrated the acceleration of

high-charge $\sim 10-100$ pC electron bunches to 600 MeV in a 40 mm-long slit-nozzle gas jet. Longitudinal density ramps from a transverse HOFI channel were used to inject electron bunches with a higher mean and peak charge compared to those injected by ionisation injection for a truncated guiding channel alone. Electron beams were generally more bunch-like compared to those generated without the transverse injector channel. We were able to couple both the drive laser pulse and electron bunch into a longitudinal guiding channel. We observed strong oscillations of the electron beam pointing as a function of the bunch energy, suggesting that the electron bunches underwent strong betatron oscillations when the injector beam was present. This may have applications for single-stage radiation generation. PIC simulations are also required to understand the experimental results presented in the final chapter, and we plan to complement these with magnetohydrodynamic simulations (MHD) to understand the evolution of the crossed-channels. It is hoped that this will explain the origin of the electron betatron oscillations that we saw experimentally, allowing us to develop this scheme for light-source applications or to suppress it, if required. We will also attempt to understand the large energy spread of the injected bunches.

Although it is clear that work needs to be done to achieve the control and repeatability of conventional accelerators, we hope to have demonstrated that novel and imaginative techniques which utilise the properties of ultrashort laser pulses offer a wealth of potential and variety for generating electron bunches in plasma-based accelerators. We hope that in the future some of these techniques may be demonstrated experimentally, in aid of the eventual goal to miniaturise particle accelerators.

References

[1] J. C. Wilcke. Ytterligare Rön Och Försök Om Contraira Electriciteterne Vid Laddningen Och Därtil Hörande Delar | Additional Findings and Experiments on the Opposing Electric Charges During Charging, and Parts Related Thereto. *Kungliga Svenska Vetenskaps Academiens Handlingar | Proceedings of the Royal Swedish Science Academy*, 23(206229):245266, 1762. 3

[2] G. Pancaldi. *Volta: Science and Culture in the Age of Enlightenment*. Princeton University Press, 2005. 3

[3] E. Rutherford. The Scattering of α and β Particles by Matter and the Structure of the Atom. *Philosophical Magazine*, 21(669-688), 1911. 3

[4] J. D. Cockroft and E. T. S. Walton. Disintegration of Lithium by Swift Protons. *Nature*, 129(649), 1932. 3

[5] P. W. Higgs. Broken Symmetries and the Masses of Gauge Bosons. *Phys. Rev. Lett.*, 13:508–509, 1964. 3

[6] G. Aad, T. Abajyan, B. Abbott, et al. Observation of a New Particle in the Search for the Standard Model Higgs Boson with the ATLAS Detector at the LHC. *Physics Letters B*, 716(1):1–29, 2012. 3

[7] L. Evans and P. Bryant. LHC Machine. *JINST*, 3:S08001, 2008. 3

[8] A. P. Chernyaev and S. M. Varzar. Particle Accelerators in Modern World. *JINST*, 77:1203–1215, 2014. 3

[9] R. J. Van de Graaff, K. T. Compton, and L. C. Van Atta. The Electrostatic Production of High Voltage for Nuclear Investigations. *Phys. Rev.*, 43:149–157, 1933. 3

[10] G. Ising. Prinzip einer Methode zur Herstellung von Kanalstrahlen hoher Voltzahl | Principle of a Method for Producing High-Voltage Channel Beams. *Ark. Mat. Astron. Fys.*, 18(30):1–4, 1924. 3

[11] L. W. Alvarez, H. Bradner, J. V. Franck, et al. Berkeley Proton Linear Accelerator. *Review of Scientific Instruments*, 26(2):111–133, 1955. 4

[12] E. O. Lawrence. On the Production of High Speed Protons. *Science*, 72(1867), 1930. 4

[13] P. J. Bryant. A Brief History and Review of Accelerators. 4

[14] E. O. Lawrence and M. S. Livingston. The Production of High Speed Light Ions Without the Use of High Voltages. *Phys. Rev.*, 40:19–35, 1932. 4

[15] V. I. Veksler. *A New Method of Acceleration of Relativistic Particles*. Institute for Physical Problems, Academy of Sciences of the USSR, Moscow, 1945. 4

[16] E. M. McMillan. The Synchrotron–A Proposed High Energy Particle Accelerator. *Phys. Rev.*, 68:143–144, 1945. 4

[17] W. Wuensch, C. Achard, S. Dobert, et al. Demonstration of High-Gradient Acceleration. In *Proceedings of the 2003 Particle Accelerator Conference*, volume 1, pages 495–497, 2003. 4

[18] *Operational Challenges of the LHC*. CEA, 2007. 5

[19] T. Tajima and J. M. Dawson. Laser Electron Accelerator. *Phys. Rev. Lett.*, 43:267–270, 1979. 5, 7

[20] *Coherent Principle of Acceleration of Charged Particles*, 1956. 5

[21] Y. B. Faĭnberg. Acceleration of Charged Particles in a Plasma. *Soviet Physics Uspekhi*, 10(6):750, 1968. 5

[22] J. B. Rosenzweig, B. Breizman, T. Katsouleas, and J. J. Su. Acceleration and Focusing of Electrons in Two-Dimensional Nonlinear Plasma Wake Fields. *Phys. Rev. A*, 44:R6189–R6192, Nov 1991. 5

[23] C. Joshi. Plasma-Based Accelerators: Then and Now. *Plasma Physics and Controlled Fusion*, 61(10):104001, aug 2019. 5

[24] A. J. Gonsalves, K. Nakamura, J. Daniels, et al. Petawatt Laser Guiding and Electron Beam Acceleration to 8 GeV in a Laser-Heated Capillary Discharge Waveguide. *Phys. Rev. Lett.*, 122:084801, 2019. 6, 16, 18

[25] J. M. Dawson. Nonlinear Electron Oscillations in a Cold Plasma. *Phys. Rev.*, 113:383–387, 1959. 7

[26] C. E. Clayton, C. Joshi, C. Darrow, et al. Relativistic Plasma-Wave Excitation by Collinear Optical Mixing. *Phys. Rev. Lett.*, 54:2343–2346, 1985. 7

[27] C. E. Clayton, K. A. Marsh, A. Dyson, et al. Ultrahigh-Gradient Acceleration of Injected Electrons by Laser-Excited Relativistic Electron Plasma Waves. *Phys. Rev. Lett.*, 70:37–40, 1993. 7

[28] M. N. Rosenbluth and C. S. Liu. Excitation of Plasma Waves by Two Laser Beams. *Phys. Rev. Lett.*, 29:701–705, 1972. 7

[29] N. E. Andreev, L. M. Gorbunov, V. I. Kirsanov, et al. Resonant Excitation of Wakefields by a Laser Pulse in a Plasma. *Soviet Journal of Experimental and Theoretical Physics Letters*, 55:571, 1992. 7

[30] J. Krall, A. Ting, E. Esarey, et al. Enhanced Acceleration in a Self-Modulated-Laser Wake-Field Accelerator. *Phys. Rev. E*, 48:2157–2161, 1993. 7, 122

[31] S. Hooker. Developments in Laser-Driven Plasma Accelerators. *Nature Photon*, 7:775782, 2013. 8, 12, 13

[32] F. F. Chen. *Introduction to Plasma Physics and Controlled Fusion*. Springer Cham, 2016. 8, 111

[33] A. Macchi. *A SuperIntense Laser-Plasma Interaction Theory Primer*. SpringerBriefs in Physics. Springer Dordrecht, 1 edition, 2013. 8

[34] D. Strickland and G. Mourou. Compression of Amplified Chirped Optical Pulses. *Optics Communications*, 56(3):219–221, 1985. 8

[35] A. Pukhov and J. Meyer-ter Vehn. Laser Wake Field Acceleration: The Highly Non-Linear Broken-Wave Regime. *Applied Physics B*, 74(4):355–361, 2002. 9

[36] C. G. R. Geddes, C. Toth, J. van Tilborg, et al. High-Quality Electron Beams from a Laser Wakefield Accelerator using Plasma-Channel Guiding. *Nature*, 431(7008):538–541, 2004. 9, 12

[37] S. P. D. Mangles, C. D. Murphy, Z. Najmudin, et al. Monoenergetic Beams of Relativistic Electrons from Intense Laser-Plasma Interactions. *Nature*, 431(7008):535–538, 2004. 9, 12

[38] J. Faure, Y. Glinec, A. Pukhov, et al. A Laser-Plasma Accelerator Producing Monoenergetic Electron Beams. *Nature*, 431(7008):541–544, 2004. 9, 12

[39] F. Grüner, S. Becker, U. Schramm, et al. Design Considerations for Table-Top, Laser-Based VUV and X-Ray Free Electron Lasers. *Applied Physics B*, 86(3):431–435, 2007. 9

[40] K. Nakajima, A. Deng, X. Zhang, et al. Operating Plasma Density Issues on Large-Scale Laser-Plasma Accelerators Toward High-Energy Frontier. *Phys. Rev. ST Accel. Beams*, 14:091301, 2011. 9

[41] C. B. Schroeder, E. Esarey, and W. P. Leemans. Beamstrahlung Considerations in Laser-Plasma-Accelerator-Based Linear Colliders. *Phys. Rev. ST Accel. Beams*, 15:051301, 2012. 9

[42] C. Aniculaesei, T. Ha, S. Yoffe, et al. The Acceleration of a High-Charge Electron Bunch to 10 GeV in a 10-cm Nanoparticle-Assisted Wakefield Accelerator. *Matter and Radiation at Extremes*, 9(1):014001, 2023. 9, 12

[43] J. Wenz and Karsch S. Physics of Laser-Wakefield Accelerators (LWFA), 2020. 10, 22, 23, 122, 124, 128, 131, 132, 140

[44] E. Esarey, C. B. Schroeder, and W. P. Leemans. Physics of Laser-Driven Plasma-Based Electron Accelerators. *Rev. Mod. Phys.*, 81:1229–1285, 2009. 11, 16, 111, 112, 122, 131

[45] P. Sprangle, B. Hafizi, J. R. Peñano, et al. Wakefield Generation and GeV Acceleration in Tapered Plasma Channels. *Phys. Rev. E*, 63:056405, 2001. 11

[46] J. P. Palastro, J. L. Shaw, P. Franke, et al. Dephasingless Laser Wakefield Acceleration. *Phys. Rev. Lett.*, 124:134802, 2020. 11, 53, 54

[47] C. Huang, W. Lu, M. Zhou, et al. Hosing Instability in the Blow-Out Regime for Plasma-Wakefield Acceleration. *Phys. Rev. Lett.*, 99:255001, 2007. 12, 13

[48] T. J. Mehrling, R. A. Fonseca, A. Martinez de la Ossa, et al. Mitigation of the Hose Instability in Plasma-Wakefield Accelerators. *Phys. Rev. Lett.*, 118:174801, 2017. 12, 13

[49] T. J. Mehrling, R. A. Fonseca, A. Martinez de la Ossa, et al. Mechanisms for the Mitigation of the Hose Instability in Plasma-Wakefield Accelerators. *Phys. Rev. Accel. Beams*, 22:031302, 2019. 12, 13

[50] R. Ariniello, C. E. Doss, V. Lee, et al. Chromatic Transverse Dynamics in a Nonlinear Plasma Accelerator. *Phys. Rev. Res.*, 4:043120, 2022. 12, 13

[51] C. S. Hue, Y. Wan, E. Y. Levine, et al. Control of Electron Beam Current, Charge, and Energy Spread using Density DownRamp Injection in Laser Wakefield Accelerators. *Matter and Radiation at Extremes*, 8(2):024401, 2023. 12

[52] G. Li, Q. Ain, S. Li, et al. Control of Electron Beam Energy-Spread by Beam Loading Effects in a Laser-Plasma Accelerator. *Plasma Physics and Controlled Fusion*, 62(5):055004, 2020. 12

[53] M. Kirchen, S. Jalas, P. Messner, et al. Optimal Beam Loading in a Laser-Plasma Accelerator. *Phys. Rev. Lett.*, 126:174801, 2021. 12

[54] J. van Tilborg, C. B. Schroeder, C. V. Filip, et al. Temporal Characterization of Femtosecond Laser-Plasma-Accelerated Electron Bunches Using Terahertz Radiation. *Phys. Rev. Lett.*, 96:014801, 2006. 12

[55] T. Ohkubo, A. Maekawa, R. Tsujii, et al. Temporal Characteristics of Monoenergetic Electron Beams Generated by the Laser Wakefield Acceleration. *Phys. Rev. ST Accel. Beams*, 10:031301, 2007. 12

[56] A. D. Debus, M. Bussmann, U. Schramm, et al. Electron Bunch Length Measurements from Laser-Accelerated Electrons Using Single-Shot THz Time-Domain Interferometry. *Phys. Rev. Lett.*, 104:084802, 2010. 12

[57] O. Lundh, J. Lim, C. Rechatin, et al. Few Femtosecond, Few Kiloampere Electron Bunch Produced by a LaserPlasma Accelerator. *Nature Physics*, 7(3):219–222, 2011. 12

[58] A. Dabrowski, M. Velasco, H.H. Braun, et al. Non-Destructive Single Shot Bunch Length Measurements for the CLIC Test Facility 3. In *2007 IEEE Particle Accelerator Conference (PAC)*, pages 4069–4071, 2007. 12

[59] S. Fritzler, E. Lefebvre, V. Malka, et al. Emittance Measurements of a Laser-Wakefield-Accelerated Electron Beam. *Phys. Rev. Lett.*, 92:165006, 2004. 12

[60] C. M. S. Sears, A. Buck, K. Schmid, et al. Emittance and Divergence of Laser Wakefield Accelerated Electrons. *Phys. Rev. ST Accel. Beams*, 13:092803, 2010. 12

[61] E. Brunetti, R. P. Shanks, G. G. Manahan, et al. Low Emittance, High Brilliance Relativistic Electron Beams from a Laser-Plasma Accelerator. *Phys. Rev. Lett.*, 105:215007, 2010. 12

[62] G. R. Plateau, C. G. R. Geddes, D. B. Thorn, et al. Low-Emittance Electron Bunches from a Laser-Plasma Accelerator Measured using Single-Shot X-Ray Spectroscopy. *Phys. Rev. Lett.*, 109:064802, 2012. 12

[63] R. Weingartner, S. Raith, A. Popp, et al. Ultralow Emittance Electron Beams from a Laser-Wakefield Accelerator. *Phys. Rev. ST Accel. Beams*, 15:111302, 2012. 12

[64] I. Dornmair, K. Floettmann, and A. R. Maier. Emittance Conservation by Tailored Focusing Profiles in a Plasma Accelerator. *Phys. Rev. ST Accel. Beams*, 18:041302, 2015. 12

[65] Y. Zhao, W. An, X. Xu, et al. Emittance Preservation Through Density Ramp Matching Sections in a Plasma Wakefield Accelerator. *Phys. Rev. Accel. Beams*, 23:011302, 2020. 12

[66] T. André, I. A. Andriyash, A. Loulergue, et al. Control of Laser Plasma Accelerated Electrons for Light Sources. *Nature Communications*, 9(1):1334, 2018. 12

[67] C. B. Schroeder, J.-L. Vay, E. Esarey, et al. Thermal Emittance from Ionization-Induced Trapping in Plasma Accelerators. *Phys. Rev. ST Accel. Beams*, 17:101301, 2014. 12

[68] S. K. Barber, J. van Tilborg, C. B. Schroeder, et al. Measured Emittance Dependence on the Injection Method in Laser Plasma Accelerators. *Phys. Rev. Lett.*, 119:104801, 2017. 12

[69] J. Götzfried, A. Döpp, M. F. Gilljohann, et al. Physics of High-Charge Electron Beams in Laser-Plasma Wakefields. *Phys. Rev. X*, 10:041015, 2020. 12

[70] J. P. Couperus, R. Pausch, A. Köhler, et al. Demonstration of a Beam Loaded Nanocoulomb-Class Laser Wakefield Accelerator. *Nature Communications*, 8(1):487, 2017. 12

[71] L. Rovige, J. Huijts, I. Andriyash, et al. Demonstration of Stable Long-Term Operation of a Kilohertz Laser-Plasma Accelerator. *Phys. Rev. Accel. Beams*, 23:093401, 2020. 13

[72] I. N. Ross, P. Matousek, M. Towrie, et al. The Prospects for Ultrashort Pulse Duration and Ultrahigh Intensity using Optical Parametric Chirped Pulse Amplifiers. *Optics Communications*, 144(1):125–133, 1997. 13

[73] A. Dubietis, R. Butkus, and A. P. Piskarskas. Trends in Chirped Pulse Optical Parametric Amplification. *IEEE Journal of Selected Topics in Quantum Electronics*, 12(2):163–172, 2006. 13

[74] T. Eidam, J. Rothhardt, F. Stutzki, et al. Fiber Chirped-Pulse Amplification System Emitting 3.8 GW Peak Power. *Opt. Express*, 19(1):255–260, 2011. 13

[75] G. D. Goodno, H. Komine, S. J. McNaught, et al. Coherent Combination of High-Power, Zigzag Slab Lasers. *Opt. Lett.*, 31(9):1247–1249, 2006. 13

[76] W. International Committee for Future Accelerators, edited by Leemans, W. Chou, and M. Uesaka. Particles and Fields Commission of IUPAP Beam Dynamics Newsletter No. 56. In *Proceedings of the International Committee for Future Accelerators*, 2011. 13

[77] J. W. Dawson, J. K. Crane, M. J. Messerly, et al. High Average Power Lasers for Future Particle Accelerators. *AIP Conference Proceedings*, 1507(1):147–153, 2012. 13

[78] D. J. Richardson, J. Nilsson, and W. A. Clarkson. High Power Fiber Lasers: Current Status and Future Perspectives. *J. Opt. Soc. Am. B*, 27(11):B63–b92, 2010. 13

[79] O. Jakobsson, S. M. Hooker, and R. Walczak. Gev-Scale Accelerators Driven by Plasma-Modulated Pulses from Kilohertz Lasers. *Phys. Rev. Lett.*, 127:184801, 2021. 13

[80] J. J. van de Wetering, S. M. Hooker, and R. Walczak. Stability of the Modulator in a Plasma-Modulated Plasma Accelerator. *Phys. Rev. E*, 108:015204, 2023. 13

[81] S. M. Hooker, R. Bartolini, S. P. D. Mangles, et al. Multi-Pulse Laser Wakefield Acceleration: a New Route to Efficient, High-Repetition-Rate Plasma Accelerators and High Flux Radiation Sources. *Journal of Physics B: Atomic, Molecular and Optical Physics*, 47(23):234003, 2014. 13

[82] R. J. Shalloo, L. Corner, C. Arran, et al. Generation of Laser Pulse Trains for Tests of Multi-Pulse Laser Wakefield Acceleration. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 829:383–385, 2016. 2nd European Advanced Accelerator Concepts Workshop - EAAC 2015. 13

[83] J. Jonnerby. *Multi-Pulse Laser Wakefield Acceleration*. PhD Thesis, University of Oxford, 2021. 13

[84] A. J. Ross, J. Chappell, J. J. van de Wetering, et al. Resonant Excitation of Plasma Waves in a Plasma Channel (Accepted). *Phys. Rev. Rea.*, 2024. 13, 163

[85] G. Mourou, B. Brocklesby, T. Tajima, et al. The Future is Fibre Accelerators. *Nature Photonics*, 7(4):258–261, 2013. 13

[86] N. S. Kapany and J. J. Burke. *Optical Waveguides*. Quantum Electronics–Principles and Applications. Academic Press, New York, 1972. 14

[87] L. C. Steinhauer and H. G. Ahlstrom. Propagation of Coherent Radiation in a Cylindrical Plasma Column. *The Physics of Fluids*, 14(6):1109–1114, 1971. 14

[88] C. G. Durfee and H. M. Milchberg. Light Pipe for High Intensity Laser Pulses. *Phys. Rev. Lett.*, 71:2409–2412, 1993. 14, 15

[89] C. G. Durfee, J. Lynch, and H. M. Milchberg. Development of a Plasma Waveguide for High-Intensity Laser Pulses. *Phys. Rev. E*, 51:2368–2389, 1995. 14

[90] H. M. Milchberg, T. R. Clark, C. G. Durfee, et al. Development and Applications of a Plasma Waveguide for Intense Laser Pulses. *Physics of Plasmas*, 3(5):2149–2155, 1996. 14

[91] S. M. Hooker, D. J. Spence, and R. A. Smith. Guiding of High-Intensity Picosecond Laser Pulses in a Discharge-Ablated Capillary Waveguide. *J. Opt. Soc. Am. B*, 17(1):90–98, 2000. 14, 16

[92] J. F. Seely. *Quantum Theory of Inverse Bremsstrahlung Absorption and Pair Production*, pages 835–847. Springer US, Boston, MA, 1974. 15

[93] E. W. Gaul, S. P. Le Blanc, A. R. Rundquist, et al. Production and Characterization of a Fully Ionized He Plasma Channel. *Applied Physics Letters*, 77(25):4112–4114, 2000. 15

[94] T. Ditmire, R. A. Smith, and M. H. R. Hutchinson. Plasma Waveguide Formation in Predissociated Clustering Gases. *Opt. Lett.*, 23(5):322–324, 1998. 16

[95] N. Lemos, T. Grismayer, L. Cardoso, et al. Plasma Expansion into a Waveguide Created by a Linearly Polarized Femtosecond Laser Pulse. *Physics of Plasmas*, 20(6):063102, 2013. 16

[96] N. Lemos, T. Grismayer, L. Cardoso, et al. Effects of Laser Polarization in the Expansion of Plasma Waveguides. *Physics of Plasmas*, 20(10):103109, 2013. 16

[97] M. W. Mayr, B. Spiers, R. Aboushelbaya, R. W. Paddock, J. D. Sadler, C. Sillett, R. H. W. Wang, K. Krushelnick, and P. A. Norreys. Nonlinear Wakefields and Electron Injection in Cluster Plasma. *Phys. Rev. Accel. Beams*, 23:093501, Sep 2020. 16

[98] S. Jackel, R. Burris, J. Grun, et al. Channeling of Terawatt Laser Pulses by Use of Hollow Waveguides. *Opt. Lett.*, 20(10):1086–1088, 1995. 16

[99] M. Borghesi, A. J. Mackinnon, R. Gaillard, et al. Guiding of a 10-TW Picosecond Laser Pulse Through Hollow Capillary Tubes. *Phys. Rev. E*, 57:R4899–r4902, 1998. 16

[100] F. Dorchies, J. R. Marquès, B. Cros, et al. Monomode Guiding of $10^{16}\,\mathrm{W\,cm^{-2}}$ Laser Pulses Over 100 Rayleigh Lengths in Hollow Capillary Dielectric Tubes. *Phys. Rev. Lett.*, 82:4655–4658, 1999. 16

[101] C. Courtois, A. Couairon, B. Cros, et al. Propagation of Intense Ultrashort Laser Pulses in a Plasma Filled Capillary Tube: Simulations and Experiments. *Physics of Plasmas*, 8(7):3445–3456, 2001. 16

[102] B. Cros, C. Courtois, G. Malka, et al. Extending Plasma Accelerators: Guiding with Capillary Tubes. *IEEE Transactions on Plasma Science*, 28(4):1071–1077, 2000. 16

[103] A. Zigler, Y. Ehrlich, C. Cohen, et al. Optical Guiding of High-Intensity Laser Pulses in a Long Plasma Channel Formed by a Slow Capillary Discharge. *J. Opt. Soc. Am. B*, 13(1):68–71, 1996. 16

[104] Y. Ehrlich, C. Cohen, A. Zigler, et al. Guiding of High Intensity Laser Pulses in Straight and Curved Plasma Channel Experiments. *Phys. Rev. Lett.*, 77:4186–4189, 1996. 16

[105] Y. Ehrlich, C. Cohen, D. Kaganovich, et al. Guiding and Damping of High-Intensity Laser Pulses in Long Plasma Channels. *J. Opt. Soc. Am. B*, 15(9):2416–2423, 1998. 16

[106] D. J. Spence and S. M. Hooker. Investigation of a Hydrogen Plasma Waveguide. *Phys. Rev. E*, 63:015401, 2000. 17

[107] 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, 2001. 17

[108] D. J. Spence, A. Butler, and S. M. Hooker. Gas-Filled Capillary Discharge Waveguides. *J. Opt. Soc. Am. B*, 20(1):138–151, 2003. 17

[109] T. Hosokai, M. Kando, H. Dewa, et al. Optical Guidance of Terrawatt Laser Pulses by the Implosion Phase of a Fast Z-Pinch Discharge in a Gas-Filled Capillary. *Optics letters*, 25(1):10–2, 2000. 17

[110] Tomonao H., Masaki K., Hideki D., et al. Fast Z-Pinch Optical Guiding for Laser Wakefield Acceleration. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 455(1):155–160, 2000. Proceedings of the Int. Symp. on New Visions in Laser-Beam. 17

[111] A. J. Gonsalves, T. P. Rowlands-Rees, B. H. P. Broks, et al. Transverse Interferometry of a Hydrogen-Filled Capillary Discharge Waveguide. *Phys. Rev. Lett.*, 98:025002, 2007. 17

[112] B. H. P. Broks, K. Garloff, and J. J. A. M. van der Mullen. Nonlocal-Thermal-Equilibrium Model of a Pulsed Capillary Discharge Waveguide. *Phys. Rev. E*, 71:016401, 2005. 17

[113] P. S. Antsiferov, M. R. Akdim, and H. T. van Dam. Direct Measurement of the Matched Spot Size in a Slow Capillary Discharge Optical Waveguide. *Review of Scientific Instruments*, 78(12):123107, 2007. 17

[114] A. J. Gonsalves, K. Nakamura, C. Lin, et al. Plasma Channel Diagnostic Based on Laser Centroid Oscillations. *AIP Conference Proceedings*, 1299(1):150–155, 2010. 17

[115] A. J. Gonsalves, K. Nakamura, C. Lin, et al. Accurate Alignment of Plasma Channels Based on Laser Centroid Oscillations. In *Proceedings of 2011 Particle Accelerator Conference*, volume 1, pages 328–330, 2011. 17

[116] N. A. Bobrova, P. V. Sasorov, C. Benedetti, et al. Laser-Heater Assisted Plasma Channel Formation in Capillary Discharge Waveguides. *Physics of Plasmas*, 20(2):020703, 2013. 18

[117] A. J. Gonsalves, F. Liu, N. A. Bobrova, et al. Demonstration of a High Repetition Rate Capillary Discharge Waveguide. *Journal of Applied Physics*, 119(3):033302, 2016. 18

[118] A. J. Gonsalves, K. Nakamura, C. Benedetti, et al. Laser-Heated Capillary Discharge Plasma Waveguides for Electron Acceleration to 8 GeV. *Physics of Plasmas*, 27(5):053102, 2020. 18

[119] G. A. Bagdasarov, N. A. Bobrova, O. G. Olkhovskaya, et al. Creation of an Axially Uniform Plasma Channel in a Laser-Assisted Capillary Discharge. *Physics of Plasmas*, 28(5):053104, 2021. 18

[120] M. Turner, A. J. Gonsalves, S. S. Bulanov, et al. Radial Density Profile and Stability of Capillary Discharge Plasma Waveguides of Lengths up to 40 cm. *High Power Laser Science and Engineering*, 9:e17, 2021. 18

[121] R. R. Freeman and P. H. Bucksbaum. Investigations of Above-Threshold Ionization using Subpicosecond Laser Pulses. *Journal of Physics B: Atomic, Molecular and Optical Physics*, 24(2):325, 1991. 19

[122] M. V. Ammosov, N. B. Delone, and V. P. Krainov. Tunnel Ionization of Complex Atoms and of Atomic Ions in an Alternating Electromagnetic Field. *Soviet Journal of Experimental and Theoretical Physics*, 64(6):1191, 1986. 19

[123] M. V. Ammosov, Nikolai B. Delone, and Vladimir P. Krainov. Tunnel Ionization Of Complex Atoms And Atomic Ions In Electromagnetic Field. In *Other Conferences*, 1986. 19, 147

[124] R. J. Shalloo, C. Arran, L. Corner, et al. Hydrodynamic Optical-Field-Ionized Plasma Channels. *Phys. Rev. E*, 97:053203, 2018. 20, 23, 147

[125] R. Shalloo. *Hydrodynamic Optical-Field-Ionized Plasma Waveguides for Laser Plasma Accelerators*. PhD Thesis, University of Oxford, 2018. 20, 147, 154, 155, 159, 162, 163, 168

[126] R. J. Shalloo, C. Arran, A. Picksley, et al. Low-Density Hydrodynamic Optical-Field-Ionized Plasma Channels Generated with an Axicon Lens. *Phys. Rev. Accel. Beams*, 22:041302, 2019. 20, 23, 147

[127] A. Picksley, A. Alejo, J. Cowley, et al. Guiding of High-Intensity Laser Pulses in 100-mm-long Hydrodynamic Optical-Field-Ionized Plasma Channels. *Phys. Rev. Accel. Beams*, 23:081303, 2020. 20, 23, 147

[128] A. Picksley. *Low Density Plasma Waveguides for Multi-GeV Laser Wakefield Accelerators*. PhD Thesis, University of Oxford, 2021. 21, 147, 154, 155, 159, 163, 168, 169

[129] A. Picksley, A. Alejo, R. J. Shalloo, et al. Meter-Scale Conditioned Hydrodynamic Optical-Field-Ionized Plasma Channels. *Phys. Rev. E*, 102:053201, 2020. 21, 147

[130] A. Alejo, J. Cowley, A. Picksley, et al. Demonstration of Kilohertz Operation of Hydrodynamic Optical-Field-Ionized Plasma Channels. *Phys. Rev. Accel. Beams*, 25:011301, 2022. 21, 147

[131] A. J. Ross, A. Alejo, A. von Boetticher, et al. Numerical Modelling of Chromatic Effects on Axicon-Focused Beams used to Generate HOFI Plasma Channels. *Journal of Physics: Conference Series*, 1596(1):012049, 2020. 21

[132] A. Picksley, J. Chappell, E. Archer, et al. All-Optical GeV Electron Bunch Generation in a Laser-Plasma Accelerator via Truncated-Channel Injection. *Phys. Rev. Lett.*, 131:245001, 2023. 21, 147, 148, 177, 183

[133] S. Mangles, C. Murphy, and Z. Najmudin. Monoenergetic Beams of Relativistic Electrons from Intense Laser-Plasma Interactions. *Nature*, 431:535538, 2004. 21

[134] S. Bulanov, N. Naumova, F. Pegoraro, et al. Particle Injection into the Wave Acceleration Phase due to Nonlinear Wake Wave Breaking. *Phys. Rev. E*, 58:R5257–r5260, 1998. 21

[135] E. Svystun, R. W. Assmann, U. Dorda, et al. Beam Quality Preservation Studies in a Laser-Plasma Accelerator with External Injection for EuPRAXIA. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 909:90–94, 2018. 3rd European Advanced Accelerator Concepts workshop (EAAC2017). 22

[136] Y. Wu, J. Hua, and Z. Zhou. High-Throughput Injection-Acceleration of Electron Bunches from a Linear Accelerator to a Laser Wakefield Accelerator. *Nature Physics*, 17:801806, 2021. 22

[137] A. Gonsalves, K. Nakamura, and C. Lin. Tunable Laser Plasma Accelerator Based on Longitudinal Density Tailoring. *Nature Phys*, 7:862866, 2011. 22

[138] J. Faure, C. Rechatin, O. Lundh, et al. Injection and Acceleration of Quasimonoenergetic Relativistic Electron Beams using Density Gradients at the Edges of a Plasma Channel. *Physics of Plasmas*, 17(8):083107, 2010. 22

[139] A. Buck, J. Wenz, J. Xu, et al. Shock-Front Injector for High-Quality Laser-Plasma Acceleration. *Phys. Rev. Lett.*, 110:185006, 2013. 22

[140] K. Schmid, A. Buck, C. M. S. Sears, et al. Density-Transition Based Electron Injector for Laser Driven Wakefield Accelerators. *Phys. Rev. ST Accel. Beams*, 13:091301, 2010. 22

[141] A. Pak, K. A. Marsh, S. F. Martins, et al. Injection and Trapping of Tunnel-Ionized Electrons into Laser-Produced Wakes. *Phys. Rev. Lett.*, 104:025003, 2010. 23

[142] C. Thaury, E. Guillaume, A. Lifschitz, et al. Shock Assisted Ionization Injection in Laser-Plasma Accelerators. *Sci Rep*, 5:16310, 2015. 23

[143] M. Hansson, T. L. Audet, H. Ekerfelt, et al. Localization of Ionization-Induced Trapping in a Laser Wakefield Accelerator using a Density Down-Ramp. *Plasma Physics and Controlled Fusion*, 58(5):055009, 2016. 23

[144] J. Faure, C. Rechatin, A. Norlin, et al. Controlled Injection and Acceleration of Electrons in Plasma Wakefields by Colliding Laser Pulses. *Nature*, 444:737739, 2006. 23

[145] A. Deng, O.S. Karger, and T. Heinemann. Generation and Acceleration of Electron Bunches from a Plasma Photocathode. *Nature Physics*, 15:11561160, 2019. 23

[146] N. Bourgeois, J. Cowley, and S. M. Hooker. Two-Pulse Ionization Injection into Quasilinear Laser Wakefields. *Phys. Rev. Lett.*, 111:155004, 2013. 23, 114, 123, 124, 125, 126

[147] M. W. von der Leyen, J. Holloway, Y. Ma, et al. Observation of Monoenergetic Electrons from Two-Pulse Ionization Injection in Quasilinear Laser Wakefields. *Phys. Rev. Lett.*, 130:105002, 2023. 24, 114, 125

[148] P. de Fermat. *Œuvres de Fermat, Vol. 2: Correspondance*. Gauthier-Villars, 1891. [English translation from the French of 'Œuvres de Fermat' (1662)]. 27

[149] I. al Haytham. *Book of Optics*. The Warburg Institute, 1021. [English translation from the Arabic of كتاب المناظر (1983)]. 27

[150] R. J. Pegis. I The Modern Development of Hamiltonian Optics. In E. Wolf, editor, *Progress in Optics*, volume 1 of *Progress in Optics*, pages 1–29. Elsevier, 1961. 27, 28

[151] W. R. Hamilton. Theory of Systems of Rays. *Transactions of the Royal Irish Academy*, 15:69–174, 1828. 27

[152] J. Kepler. *Optics: Paralipomena to Witelo and the Optical Part of Astronomy*. Green Lion Press, 2000. [English translation from the Latin of 'Ad Vitellionem Paralipomena, Quibus Astronomiae Pars Optica Traditur' (1604)]. 27

[153] C. Huygens. *Treatise on Light: In which are Explained the Causes of that which Occurs in Reflexion, & in Refraction. And Particularly in the Strange Refraction of Iceland Crystal*. MacMillan and Company, limited, 1912. [English translation from the French of 'Traité de la Lumière' (1690)]. 27

[154] T. Young. The Bakerian Lecture: On the Theory of Light and Colours. [Abstract]. *Abstracts of the Papers Printed in the Philosophical Transactions of the Royal Society of London*, 1:63–67, 1800. 27

[155] A. Sommerfeld and J. Runge. The Application of Vector Calculus to the Foundations of Geometrical Optics. *Annalen der Physik*, 340(7):277–298, 1911. [English translation from the German of 'Anwendung der Vektorrechnung auf die Grundlagen der geometrischen Optik' (1911)]. 27

[156] M. Born, E. Wolf, and A. B. Bhatia. *Principles of Optics*. Cambridge University Press, Cambridge, 2019. 27

[157] L. D. Landau and E. M. Lifshitz. *The Classical Theory of Fields*. Addison-Wesley Physics Series. Butterworth-Heinemann, Cambridge, Massachusetts, second edition, 1951. 28, 32, 138, 139, 140

[158] J.-C. Diels, W. Rudolph, P. F. Liao, et al. *Ultrashort Laser Pulse Phenomena: Fundamentals, Techniques, and Applications on a Femtosecond Time Scale*. Elsevier Science & Technology, San Diego, 2006. 29

[159] G. C. Sherman and H. J. Bremermann. Generalization of the Angular Spectrum of Plane Waves and the Diffraction Transform∗. *J. Opt. Soc. Am.*, 59(2):146–156, 1969. 34

[160] A. E. Siegman. *Lasers*. University Science Books, Oxford, 2006. 37, 38

[161] G. Goubau and F. Schwering. On the Guided Propagation of Electromagnetic Wave Beams. *IRE Transactions on Antennas and Propagation*, 9(3):248–256, 1961. 37

[162] H. Kogelnik and T. Li. Laser Beams and Resonators. *Appl. Opt.*, 5(10):1550–1567, Oct 1966. 37

[163] G. Vallone. Role of Beam Waist in Laguerre-Gauss Expansion of Vortex Beams. *Optics Letters*, 42(6):1097–1100, 2017. 39

[164] D. Giovannini, J. Romero, V. Potoek, et al. Spatially Structured Photons that Travel in Free Space Slower than the Speed of Light. *Science*, 347(6224):857–860, 2015. 39

[165] P. Saari. Reexamination of Group Velocities of Structured Light Pulses. *Phys. Rev. A*, 97:063824, Jun 2018. 39

[166] N. D. Bareza and N. Hermosa. Subluminal Group Velocity and Dispersion of Laguerre-Gauss Beams in Free Space. *Scientific Reports*, 6(1):26842, 2016. 40

[167] J. P. Palastro, D. Turnbull, S.-W. Bahk, et al. Ionization Waves of Arbitrary Velocity Driven by a Flying Focus. *Phys. Rev. A*, 97:033835, 2018. 41, 42, 43, 53

[168] P. Mora and T. M. Antonsen, Jr. Kinetic Modeling of Intense, Short Laser Pulses Propagating in Tenuous Plasmas. *Physics of Plasmas*, 4(1):217–229, 1997. 41, 42

[169] S. Akturk, X. Gu, E. Zeek, et al. Pulse-Front Tilt Caused by Spatial and Temporal Chirp. *Opt. Express*, 12(19):4399–4410, 2004. 43, 44

[170] S. W. Jolly, O. Gobert, and F. Quéré. Spatio-Temporal Characterization of Ultrashort Laser Beams: A Tutorial. *Journal of Optics*, 22(10):103501, sep 2020. 43

[171] A. Popp, J. Vieira, J. Osterhoff, et al. All-Optical Steering of Laser-Wakefield-Accelerated Electron Beams. *Phys. Rev. Lett.*, 105:215001, 2010. 43

[172] A. Sainte-Marie, O. Gobert, and F. Quéré. Controlling the Velocity of Ultrashort Light Pulses in Vacuum Through Spatio-Temporal Couplings. *Optica*, 4(10):1298–1304, 2017. 43, 53

[173] S. A. Collins. Lens-System Diffraction Integral Written in Terms of Matrix Optics. *J. Opt. Soc. Am.*, 60(9):1168–1177, 1970. 44, 69

[174] S. Akturk, X. Gu, P. Gabolde, et al. The General Theory of First-Order Spatio-Temporal Distortions of Gaussian Pulses and Beams. *Opt. Express*, 13(21):8642–8661, 2005. 44, 78, 102

[175] O. E. Martinez. Matrix Formalism for Pulse Compressors. *IEEE Journal of Quantum Electronics*, 24(12):2530–2536, 1988. 44

[176] O. E. Martinez. Matrix Formalism for Dispersive Laser Cavities. *IEEE Journal of Quantum Electronics*, 25(3):296–300, 1989. 44

[177] A. G. Kostenbauder. Ray-Pulse Matrices: a Rational Treatment for Dispersive Optical Systems. *IEEE Journal of Quantum Electronics*, 26(6):1148–1157, 1990. 44

[178] G. Zhang and K. Shi. A Numerical Method of Improved Bandwidth Adaptability for Simulating Diffraction Integral of a Spatial Light Modulator. *Optics Communications*, 486:126781, 2021. 44

[179] Z. Li, Y. Gu, and J. Kawanaka. Reciprocating Propagation of Laser Pulse Intensity in Free Space. *Communications Physics*, 4(1):87, 2021. 44, 53

[180] J.-C. Li and P. Picart. *Calculating Diffraction by Fast Fourier Transform*, chapter 3, pages 77–114. John Wiley & Sons, Ltd, 2012. 45, 46, 51

[181] J. Li and C. Li. Algorithm Study of Collins Formula and Inverse Collins Formula. *Appl. Opt.*, 47(4):97–102, 2008. 45, 46

[182] D. H. Froula, D. Turnbull, A. S. Davies, et al. Spatiotemporal Control of Laser Intensity. *Nature Photonics*, 12(5):262–265, 2018. 53

[183] D. Ramsey, A. Di Piazza, M. Formanek, et al. Exact Solutions for the Electromagnetic Fields of a Flying Focus. *Phys. Rev. A*, 107:013513, 2023. 53

[184] Tanner T. Simpson, Dillon Ramsey, Philip Franke, et al. Nonlinear Spatiotemporal Control of Laser Intensity. *Opt. Express*, 28(26):38516–38526, 2020. 53

[185] M. V. Ambat, J. L. Shaw, J. J. Pigeon, et al. Programmable-Trajectory Ultrafast Flying Focus Pulses. *Opt. Express*, 31(19):31354–31368, 2023. 53

[186] D. H. Froula, J. P. Palastro, D. Turnbull, et al. Flying Focus: Spatial and Temporal Control of Intensity for Laser-Based Applications. *Physics of Plasmas*, 26(3):032109, 2019. 53

[187] D. Turnbull, P. Franke, J. Katz, et al. Ionization Waves of Arbitrary Velocity. *Phys. Rev. Lett.*, 120:225001, 2018. 53

[188] F. Li, T. N. Dalichaouch, J. R. Pierce, et al. Ultrabright Electron Bunch Injection in a Plasma Wakefield Driven by a Superluminal Flying Focus Electron Beam. *Phys. Rev. Lett.*, 128:174803, 2022. 53, 54

[189] P.-F. Geng, M. Chen, X.-Y. An, et al. Plasma Density Transition-Based Electron Injection in Laser Wake Field Acceleration Driven by a Flying Focus Laser. *Chinese Physics B*, 32(4):044101, 2023. 53, 54

[190] M. Formanek, D. Ramsey, J. P. Palastro, et al. Radiation Reaction Enhancement in Flying Focus Pulses. *Phys. Rev. A*, 105:L020203, 2022. 53

[191] D. Turnbull, S. Bucht, A. Davies, et al. Raman Amplification with a Flying Focus. *Phys. Rev. Lett.*, 120:024801, 2018. 53

[192] D. Turnbull, S.-W. Bahk, I. A. Begishev, et al. Flying Focus and its Application to Plasma-Based Laser Amplifiers. *Plasma Physics and Controlled Fusion*, 61(1):014022, 2018. 53

[193] Z. Wu, Y. Zuo, Z. Zhang, et al. Self-Compression of Stimulated Raman Backscattering by a Flying Focus. *Phys. Rev. E*, 106:035209, 2022. 53

[194] A. J. Howard, D. Turnbull, A. S. Davies, et al. Photon Acceleration in a Flying Focus. *Phys. Rev. Lett.*, 123:124801, 2019. 53

[195] H. Ye, Y. Gu, Q. Fan, et al. Enhanced Thomson Scattering X-Ray Sources with Flying Focus Laser Pulse. *AIP Advances*, 13(3):035330, 2023. 53

[196] T. M. Jeong, S. V. Bulanov, P. Valenta, et al. Relativistic Flying Laser Focus by a Laser-Produced Parabolic Plasma Mirror. *Phys. Rev. A*, 104:053533, 2021. 53

[197] D. Li, K. G. Miller, J. R. Pierce, et al. Spatiotemporal Control of High-Intensity Laser Pulses with a Plasma Lens, 2023. 53

[198] T. T. Simpson, D. Ramsey, P. Franke, et al. Spatiotemporal Control of Laser Intensity Through Cross-Phase Modulation. *Opt. Express*, 30(6):9878–9891, 2022. 53

[199] Q. Fan, S. Wen, S. Wang, et al. Colorless Movement of Focal Spot with Arbitrary Velocity Based on Mirrors with Special Shape. *Physica Scripta*, 96(6):065508, 2021. 53

[200] A. Liberman, R. Lahaye, S. Smartsev, et al. Use of Spatiotemporal Couplings and an Axiparabola to Control the Velocity of Peak Intensity. *Opt. Lett.*, 49(4):814–817, 2024. 53

[201] J. W. Goodman. *Introduction to Fourier Optics*. Roberts & Co. Publishers, 2005, 3rd edition, 2005. 57, 72, 91

[202] C. F. Gauss and E. J. Schering. *Theoria Motus Corporum Coelestium In Sectionibus Conicis Solem Ambientium*, volume 7 of *Cambridge Library Collection - Mathematics*, page 12. Cambridge University Press, 2011. 58, 73

[203] G. Zhu, J. v Howe, M. Durst, et al. Simultaneous Spatial and Temporal Focusing of Femtosecond Pulses. *Opt. Express*, 13(6):2153–2159, 2005. 67, 72, 107

[204] M. E. Durst, G. Zhu, and C. Xu. Simultaneous Spatial and Temporal Focusing for Axial Scanning. *Opt. Express*, 14(25):12243–12254, 2006. 67, 68, 69, 70, 86, 91

[205] M. E. Durst, G. Zhu, and C. Xu. Simultaneous Spatial and Temporal Focusing in Nonlinear Microscopy. *Optics Communications*, 281:1796–1805, 2008. 67, 72, 86, 107

[206] T. E. Lanier and Jeremy R. G. Nonlinear Space-Time Focusing and Filamentation of Annular Femtosecond Pulses in Dielectrics. *Journal of the Optical Society of America B*, 33(2):292, 2016. 67, 89, 107

[207] B. Sun, P. S Salter, C. Roider, et al. Four-Dimensional Light Shaping: Manipulating Ultrafast Spatiotemporal Foci in Space and Time. *Light: Science & Applications*, 7(1):17117–17117, 2018. 67, 107

[208] Q. Song, A. Nakamura, K. Hirosawa, K. Isobe, K. Midorikawa, and F. Kannari. Two-Dimensional Spatiotemporal Focusing of Femtosecond Pulses and its Applications in Microscopy. *Review of Scientific Instruments*, 86(8):083701, 08 2015. 67

[209] G. Li, F. Lu, H. Liu, et al. Analysis of Spectral Chirps and Intensities Evolution of a Spatiotemporally Focused Femtosecond Laser Beam. *Results in Physics*, 7:2898–2906, 2017. 68

[210] R. Kammel, R. Ackermann, J. Thomas, et al. Enhancing Precision in fs-Laser Material Processing by Simultaneous Spatial and Temporal Focusing. *Light: Science & Applications*, 3, 2014. 68, 86

[211] C. G. Durfee, M. Greco, E. Block, et al. Intuitive Analysis of Space-Time Focusing with Double-ABCD Calculation. *Optics Express*, 20:14244–14259, 2012. 68, 72, 82, 84

[212] M. Sibai, H. Mehidine, F. Poulon, et al. The Impact of Compressed Femtosecond Laser Pulse Durations on Neuronal Tissue Used for Two-Photon Excitation Through an Endoscope. *Scientific Reports*, 8(1):11124, 2018. 69

[213] M. A. Coughlan, M. Plewicki, and R. J. Levis. Parametric Spatio-Temporal Control of Focusing Laser Pulses. *Opt. Express*, 17(18):15808–15820, 2009. 78

[214] S. Zhang, D. Asoubar, R. Kammel, et al. Analysis of Pulse Front Tilt in Simultaneous Spatial and Temporal Focusing. *J. Opt. Soc. Am. A*, 31(11):2437–2446, 2014. 79

[215] International Organization for Standardization. Lasers and Laser-Related Equipment Test Methods for Laser Beam Widths, Divergence Angles and Beam Propagation Ratios Part 3: Intrinsic and Geometrical Laser Beam Classification, Propagation and Details of Test Methods. Standard, International Organization for Standardization, Geneva, CH, 2002. 79, 102, 159

[216] A. E. Siegman. How to (Maybe) Measure Laser Beam Quality. In *OSA Trends in Optics and Photonics*. Optica Publishing Group, 1998. 79, 102, 103, 162

[217] D. J. Griffiths. *Introduction to Electrodynamics*. Pearson, 2013. 82

[218] A. Papoulis. *Systems and Transforms with Applications in Optics*. McGraw-Hill series in systems science. McGraw-Hill, New York, 1968. 94

[219] P. G. Kazansky, W. Yang, E. Bricchi, et al. "Quill" Writing with Ultrashort Light Pulses in Transparent Materials. *Applied Physics Letters*, 90(15):151120, 2007. 107

[220] W. Yang, P. G. Kazansky, Y. Shimotsuma, et al. Ultrashort-Pulse Laser Calligraphy. *Applied Physics Letters*, 93(17):171109, 2008. 107

[221] D. N. Vitek, E. Block, Y. Bellouard, et al. Spatio-Temporally Focused Femtosecond Laser Pulses for Nonreciprocal Writing in Optically Transparent Materials. *Opt. Express*, 18(24):24673–24678, 2010. 107

[222] P. G. Kazansky, Y. Shimotsuma, M. Sakakura, et al. Photosensitivity Control of An Isotropic Medium Through Polarization of Light Pulses with Tilted Intensity Front. *Opt. Express*, 19(21):20657–20664, 2011. 107

[223] Y. Dai, G. Wu, X. Lin, et al. Femtosecond Laser Induced Rotated 3D Self-Organized Nanograting in Fused Silica. *Opt. Express*, 20(16):18072–18078, 2012. 107

[224] P. S. Salter and M. J. Booth. Dynamic Control of Directional Asymmetry Observed in Ultrafast Laser Direct Writing. *Applied Physics Letters*, 101(14):141109, 2012. 107

[225] H. Song, Y. Dai, J. Song, et al. Femtosecond Laser-Induced Structural Difference in Fused Silica with a Non-Reciprocal Writing Process. *Applied Physics A*, 123(4):255, 2017. 107

[226] E. J. Bochove, G. T. Moore, and M. O. Scully. Acceleration of Particles by an Asymmetric Hermite-Gaussian Laser Beam. *Phys. Rev. A*, 46:6640–6653, 1992. 111, 112, 113

[227] E. Esarey, P. Sprangle, and J. Krall. Laser Acceleration of Electrons in Vacuum. *Phys. Rev. E*, 52:5443–5453, 1995. 111, 112

[228] T. J. M. Boyd and J. J. Sanderson. *The Physics of Plasmas*. Cambridge University Press, 2003. 111

[229] P. M. Woodward and J. D. Lawson. The Theoretical Precision with which an Arbitrary Radiation-Pattern may be Obtained from a Source of Finite Size. *Journal of the Institution of*

Electrical Engineers - Part III: Radio and Communication Engineering, 95:363–370(7), 1948. 111

[230] J. D. Lawson. Lasers and Accelerators. *IEEE Transactions on Nuclear Science*, 26(3):4217–4219, 1979. 111

[231] R. B. Palmer. *A Laser Driven Grating Linac*, pages 375–384. Birkhäuser Basel, Basel, 1980. 111

[232] L. V. Keldysh. Ionization in the Field of a Strong Electromagnetic Wave. *J. Exp. Theor. Phys.*, 20(5):1307–1314, 1965. 112

[233] D. Batani, E. Molnár, and D. Stutman. Direct Laser-Driven Electron Acceleration and Energy Gain in Helical Beams. *IEEE Transactions on Plasma Science*, 2021:6645668, 2021. 112

[234] S. Payeur, S. Fourmaux, B. E. Schmidt, J. P. MacLean, C. Tchervenkov, F. Légaré, M. Piché, and J. C. Kieffer. Generation of a Beam of Fast Electrons by Tightly Focusing a Radially Polarized Ultrashort Laser Pulse. *Applied Physics Letters*, 101(4):041105, 07 2012. 112, 113

[235] C. Varin, S. Payeur, V. Marceau, et al. Direct Electron Acceleration with Radially Polarized Laser Beams. *Applied Sciences*, 3(1):70–93, 2013. 112, 113, 118

[236] C. Maurer, A. Jesacher, S. Fürhapter, et al. Tailoring of Arbitrary Optical Vector Beams. *New Journal of Physics*, 9(3):78, 2007. 112

[237] S. W. Jolly and M. A. Porras. Clarification for the Fields of Different Radially Polarized Laguerre-Gaussian Light Beams. *Opt. Lett.*, 47(15):3632–3635, 2022. 112, 122

[238] Y. Mushiake, K. Matsumura, and N. Nakajima. Generation of Radially Polarized Optical Beam Mode by Laser Oscillation. *Proceedings of the IEEE*, 60(9):1107–1109, 1972. 112

[239] G. He, J. Guo, B. Wang, et al. Generation of Radially Polarized Beams Based on Thermal Analysis of a Working Cavity. *Opt. Express*, 19(19):18302–18309, 2011. 112

[240] D. Didychenko, S. Esser, F. Beirow, et al. Generation of a Radially Polarized Beam in a Polycrystalline Ceramic Yb:Lu_2O_3 Thin-Disk Laser. *Applied Physics B*, 129(9):146, 2023. 112

[241] F. P. Schäfer. On some Properties of Axicons. *Applied Physics B*, 39(1):1–8, 1986. 112

[242] J. Xia, Z. Yang, H. Chen, et al. Tangentially and Radially Polarized Nd:YAG Hollow Lasers with Two Pairs of Axicons. *Infrared Physics & Technology*, 107:103301, 2020. 112

[243] S. C. Tidwell, D. H. Ford, and W. D. Kimura. Generating Radially Polarized Beams Interferometrically. *Appl. Opt.*, 29(15):2234–2239, 1990. 112

[244] S. C. Tidwell, G. H. Kim, and W. D. Kimura. Efficient Radially Polarized Laser Beam Generation with a Double Interferometer. *Appl. Opt.*, 32(27):5222–5229, 1993. 112

[245] S. Carbajo, E. Granados, D. Schimpf, et al. Efficient Generation of Ultra-Intense Few-Cycle Radially Polarized Laser Pulses. *Opt. Lett.*, 39(8):2487–2490, 2014. 112

[246] A. Shulga. Generation of Radially Polarized Beams by Intracavity Waveguide Technique. In *2019 IEEE 8th International Conference on Advanced OptoElectronics and Lasers (CAOL)*, pages 170–173, 2019. 112

[247] R. Zhao, L. Qiao, D. Wang, et al. Spectral Broadening of Radially Polarized Pulses Through a Gas-Filled Hollow-Core Fiber. *Optics & Laser Technology*, 112:222–228, 2019. 112

[248] A. V. Nesterov, V. G. Niziev, and V. P. Yakunin. Generation of High-Power Radially Polarized Beam. *Journal of Physics D: Applied Physics*, 32(22):2871, 1999. 112

[249] Y. Kozawa and S. Sato. Generation of a Radially Polarized Laser Beam by Use of a Conical Brewster Prism. *Opt. Lett.*, 30(22):3063–3065, 2005. 112

[250] M. Stalder and M. Schadt. Linearly Polarized Light with Axial Symmetry Generated by Liquid-Crystal Polarization Converters. *Opt. Lett.*, 21(23):1948–1950, 1996. 112

[251] M. A. A. Neil, F. Massoumian, R. Juškaitis, et al. Method for the Generation of Arbitrary Complex Vector Wave Fronts. *Opt. Lett.*, 27(21):1929–1931, 2002. 112

[252] P.-L. Fortin, M. Piché, and C. Varin. Direct-Field Electron Acceleration with Ultrafast Radially Polarized Laser Beams: Scaling Laws and Optimization. *Journal of Physics B: Atomic, Molecular and Optical Physics*, 43(2):025401, 2009. 112, 117, 118

[253] L. J. Wong and F. X. Kärtner. Direct Acceleration of an Electron in Infinite Vacuum by a Pulsed Radially-Polarized Laser Beam. *Opt. Express*, 18(24):25035–25051, 2010. 112

[254] S. W. Jolly. Influence of Longitudinal Chromatism on Vacuum Acceleration by Intense Radially Polarized Laser Beams. *Opt. Lett.*, 44(7):1833–1836, 2019. 113

[255] S. W. Jolly. On the Importance of Frequency-Dependent Beam Parameters for Vacuum Acceleration with Few-Cycle Radially Polarized Laser Beams. *Opt. Lett.*, 45(14):3865–3868, 2020. 113

[256] W. D. Kimura, G. H. Kim, R. D. Romea, et al. Laser Acceleration of Relativistic Electrons Using the Inverse Cherenkov Effect. *Phys. Rev. Lett.*, 74:546–549, 1995. 113

[257] P. Sprangle, E. Esarey, and J. Krall. Self-Guiding and Stability of Intense Optical Beams in Gases Undergoing Ionization. *Phys. Rev. E*, 54:4211–4232, 1996. 113

[258] P. Sprangle, E. Esarey, and J Krall. Laser Driven Electron Acceleration in Vacuum, Gases, and Plasmas. *Physics of Plasmas*, 3(5):2183–2190, 1996. 113

[259] B. Rau, T. Tajima, and H. Hojo. Coherent Electron Acceleration by Subcycle Laser Pulses. *Phys. Rev. Lett.*, 78:3310–3313, 1997. 113

[260] K.-J. Kim, K. T. McDonald, G. V. Stupakov, et al. Comment on 'Coherent Acceleration by Subcycle Laser Pulses'. *Phys. Rev. Lett.*, 84:3210–3210, 2000. 113

[261] C. Varin and M. Piché. Relativistic Attosecond Electron Pulses from a Free-Space Laser-Acceleration Scheme. *Phys. Rev. E*, 74:045602, 2006. 113

[262] A. Karmakar and A. Pukhov. Collimated Attosecond GeV Electron Bunches from Ionization of High-Z Material by Radially Polarized Ultra-Relativistic Laser Pulses. *Laser and Particle Beams*, 25(3):371377, 2007. 113

[263] A. Sell and F. X. Kärtner. Attosecond Electron Bunches Accelerated and Compressed by Radially Polarized Laser Pulses and Soft-X-Ray Pulses from Optical Undulators. *Journal of Physics B: Atomic, Molecular and Optical Physics*, 47(1):015601, 2013. 113

[264] S. Carbajo, E. A. Nanni, L. J. Wong, et al. Direct Longitudinal Laser Acceleration of Electrons in Free Space. *Phys. Rev. Accel. Beams*, 19:021303, 2016. 113

[265] X. Zhang, V. N. Khudik, and G. Shvets. Synergistic Laser-Wakefield and Direct-Laser Acceleration in the Plasma-Bubble Regime. *Phys. Rev. Lett.*, 114:184801, 2015. 113

[266] T. Wang, V. Khudik, A. Arefiev, et al. Direct Laser Acceleration of Electrons in the Plasma Bubble by Tightly Focused Laser Pulses. *Physics of Plasmas*, 26(8):083101, 2019. 113

[267] X. Zhang, V. N Khudik, A. Pukhov, et al. Laser Wakefield and Direct Acceleration with Ionization Injection. *Plasma Physics and Controlled Fusion*, 58(3):034011, 2016. 113

[268] X. Zhang, T. Wang, V. N. Khudik, et al. Effects of Laser Polarization and Wavelength on Hybrid Laser Wakefield and Direct Acceleration. *Plasma Physics and Controlled Fusion*, 60(10):105002, 2018. 113

[269] H. S. Ghotra. Laser Wakefield and Direct Laser Acceleration of Electron by Chirped Laser Pulses. *Optik*, 260:169080, 2022. 113

[270] Z. Y. Xu, C. F. Xiao, H. Y. Lu, et al. New Injection and Acceleration Scheme of Positrons in the Laser-Plasma Bubble Regime. *Phys. Rev. Accel. Beams*, 23:091301, 2020. 113

[271] J. L. Shaw, F. S. Tsung, N. Vafaei-Najafabadi, et al. Role of Direct Laser Acceleration in Energy Gained by Electrons in a Laser Wakefield Accelerator with Ionization Injection. *Plasma Physics and Controlled Fusion*, 56(8):084006, 2014. 113

[272] C. Gahn, G. D. Tsakiris, A. Pukhov, et al. Multi-MeV Electron Beam Generation by Direct Laser Acceleration in High-Density Plasma Channels. *Phys. Rev. Lett.*, 83:4772–4775, 1999. 114

[273] A. V. Arefiev, V. N. Khudik, A. P. L. Robinson, et al. Beyond the Ponderomotive Limit: Direct Laser Acceleration of Relativistic Electrons in Sub-Critical Plasmas. *Physics of Plasmas*, 23(5):056704, 2016. 114

[274] Z. Gong, F. Mackenroth, T. Wang, et al. Direct Laser Acceleration of Electrons Assisted by Strong Laser-Driven Azimuthal Plasma Magnetic Fields. *Phys. Rev. E*, 102:013206, 2020. 114

[275] A. E. Hussein, A. V. Arefiev, T. Batson, et al. Towards the Optimisation of Direct Laser Acceleration. *New Journal of Physics*, 23(2):023031, 2021. 114

[276] E. Starodubtseva, I. Tsymbalov, D. Gorlova, et al. Phase Space Consideration of Low Energy Electron Injection for Direct Laser Acceleration, 2022. 114

[277] R. Babjak, L. Willingale, A. Arefiev, et al. Direct Laser Acceleration in Underdense Plasmas with Multi-PW Lasers: A Path to High-Charge, GeV-Class Electron Bunches. *Phys. Rev. Lett.*, 132:125001, 2024. 114

[278] A. Pukhov, Z. M. Sheng, and J. Meyer-Ter-Vehn. Particle Acceleration in Relativistic Laser Channels. *Phys. Plasmas*, 6:2847–2854, 1999. 114

[279] G. Shvets and J. S. Wurtele. Instabilities of Short-Pulse Laser Propagation Through Plasma Channels. *Phys. Rev. Lett.*, 73:3540–3543, 1994. 114

[280] A. P. L. Robinson, A. V. Arefiev, and D. Neely. Generating 'Superponderomotive' Electrons due to a Non-Wake-Field Interaction Between a Laser Pulse and a Longitudinal Electric Field. *Phys. Rev. Lett.*, 111:065002, 2013. 114

[281] A. V. Arefiev, A. P. L. Robinson, and V. N. Khudik. Novel Aspects of Direct Laser Acceleration of Relativistic Electrons. *Journal of Plasma Physics*, 81(4):475810404, 2015. 114

[282] T. Wang, V. Khudik, A. Arefiev, et al. Direct Laser Acceleration of Electrons in the Plasma Bubble by Tightly Focused Laser Pulses. *Phys. Plasmas*, 26(8):083101, 2019. 114

[283] J. Lin, R. Chen, H. Yu, et al. Analysis of Sub-Wavelength Focusing Generated by Radially Polarized Doughnut Gaussian Beam. *Optics & Laser Technology*, 64:242–246, 2014. 114, 120

[284] K. Huang, H. Ye, H. Liu, et al. Subwavelength Focusing of Azimuthally Polarized Beams with Vortical Phase in Dielectrics by Using an Ultra-Thin Lens, 2014. 114, 120

[285] K. Kitamura, K. Sakai, and S. Noda. Sub-Wavelength Focal Spot with Long Depth of Focus Generated by Radially Polarized, Narrow-Width Annular Beam. *Opt. Express*, 18(5):4518–4525, 2010. 115, 120

[286] W. H. Carter. Longitudinal Electric Field in a Gaussian Beam near Focus. *Appl. Opt.*, 12(8):1732–1732, 1973. 115, 117

[287] K. T. McDonald. Gaussian Laser Beams with Radial Polarization, 2012. 115

[288] N. Zaïm, D. Guénot, L. Chopineau, et al. Interaction of UltraIntense Radially-Polarized Laser Pulses with Plasma Mirrors. *Phys. Rev. X*, 10:041064, 2020. 118

[289] S. Quabis, R. Dorn, M. Eberler, et al. Focusing Light to a Tighter Spot. *Optics Communications*, 179(1):1–7, 2000. 120

[290] R. Dorn, S. Quabis, and G. Leuchs. Sharper Focus for a Radially Polarized Light Beam. *Phys. Rev. Lett.*, 91:233901, 2003. 120

[291] Yuichi K. and Shunichi S. Sharper Focal Spot Formed by Higher-Order Radially Polarized Laser Beams. *J. Opt. Soc. Am. A*, 24(6):1793–1798, 2007. 120

[292] G. Endale, D. Mohan, and S. Yadav. Focusing of Radially Polarized Light Using Gaussian Laser Beam Near its Focal Length. *IOP Conference Series: Materials Science and Engineering*, 1225(1):012027, 2022. 120

[293] W. B. Bridges. Divergence of High Order Gaussian Modes. *Appl. Opt.*, 14(10):2346–2347, 1975. 120

[294] E. Esarey and M. Pilloff. Trapping and Acceleration in Nonlinear Plasma Waves. *Physics of Plasmas*, 2(5):1432–1436, 1995. 123, 124, 141, 143

[295] P. Gibbon. *Short Pulse Laser Interactions with Matter*. Imperial College Press, 2005. 122, 130, 131, 132

[296] V. I. Berezhiani and I. G. Murusidze. Interaction of Highly Relativistic Short Laser Pulses with Plasmas and Nonlinear Wake-Field Generation. *Physica Scripta*, 45(2):87, 1992. 122

[297] D. Teychenné, G. Bonnaud, and J.-L. Bobin. Wave-Breaking Limit to the Wake-Field Effect in an Underdense Plasma. *Phys. Rev. E*, 48:R3248–r3251, 1993. 122, 138

[298] E. Esarey, R. F. Hubbard, W. P. Leemans, et al. Electron Injection into Plasma Wakefields by Colliding Laser Pulses. *Phys. Rev. Lett.*, 79:2682–2685, 1997. 122

[299] M. Chen, E. Esarey, C. B. Schroeder, et al. Theory of Ionization-Induced Trapping in Laser-Plasma Accelerators. *Physics of Plasmas*, 19(3):033101, 2012. 123, 125, 143

[300] A. Ting, E. Esarey, and P. Sprangle. Nonlinear WakeField Generation and Relativistic Focusing of Intense Laser Pulses in Plasmas. *Physics of Fluids B: Plasma Physics*, 2(6):1390–1394, 1990. 124, 128, 131, 135

[301] E. Esarey, P. Sprangle, J. Krall, et al. Optically Guided Laser WakeField Acceleration. *Physics of Fluids B: Plasma Physics*, 5(7):2690–2697, 1993. 124, 128, 130, 134, 135, 145

[302] T. Esirkepov, S. V. Bulanov, M. Yamagiwa, et al. Electron, Positron, and Photon Wakefield Acceleration: Trapping, Wake Overtaking, and Ponderomotive Acceleration. *Phys. Rev. Lett.*, 96:014803, 2006. 124, 139

[303] P. Sprangle, E. Esarey, and A. Ting. Nonlinear Interaction of Intense Laser Pulses in Plasmas. *Phys. Rev. A*, 41:4463–4469, 1990. 128, 135

[304] Paul McKenna, David Neely, Robert Bingham, and Dino Jaroszynski, editors. *Laser-Plasma Interactions and Applications*, chapter 7. Scottish Graduate Series. Springer, 2013. 131, 144

[305] J. L. Chaloupka, T. J. Kessler, and D. D. Meyerhofer. *A Single-Beam, Ponderomotive-Optical Trap for Energetic Free Electrons*, chapter 28, pages 179–184. Springer US, Boston, MA, 1998. 145

[306] C. Aniculaesei, H. T. Kim, B. J. Yoo, et al. Novel Gas Target for Laser Wakefield Accelerators. *Rev Sci Instrum*, 89(2):025110, 2018. 147

[307] B. Miao, L. Feder, J. E. Shrock, et al. Optical Guiding in Meter-Scale Plasma Waveguides. *Phys. Rev. Lett.*, 125:074801, 2020. 147

[308] K. v. Grafenstein, F. M. Foerster, F. Haberstroh, et al. Laser-Accelerated Electron Beams at 1 GeV Using Optically-Induced Shock Injection. *Scientific Reports*, 13(1):11680, 2023. 148

[309] K. v. Grafenstein, F. M. Foerster, F. Haberstroh, D. Campbell, F. Irshad, F. C. Salgado, G. Schilling, E. Travac, N. WeiSSe, M. Zepf, A. Döpp, and S. Karsch. Laser-Accelerated Electron Beams at 1 GeV using Optically-Induced Shock Injection. *Scientific Reports*, 13(1):11680, July 2023. 148

[310] W. Wang. *Pulse Characterisation Techniques for Multi-Pulse Laser Plasma Wakefield Accelerators (in Preparation)*. PhD Thesis, University of Oxford, 2023. 156, 169

[311] G. Cheung. *X-Ray Powder Diffraction as a Non-Destructive Tool for Characterising Laser Wakefield Betatron Radiation*. PhD Thesis, University of Oxford, 2018. 169, 171, 172

[312] G. E. Giakoumakis and D. M. Miliotis. Light Angular Distribution of Fluorescent Screens Excited by X-Rays. *Physics in Medicine and Biology*, 30(1), 1985. 172

[313] J.-P. Bissonnette, I. A. Cunningham, and P. Munro. Optimal Phosphor Thickness for Portal Imaging. *Medical Physics*, 24(6), 1997. 172

[314] Y. Glinec and et al. Absolute Calibration for a Broad Range Single Shot Electron Spectrometer. *Review of Scientific Instruments*, 77(10), 2006. 172

[315] R. Nowotny and A. Taubeck. A Method for the Production of Composite Scintillators for Dosimetry in Diagnostic Radiology. *Physics in Medicine and Biology*, 54(6), 2009. 172

[316] A. Buck and et al. Absolute Charge Calibration of Scintillating Screens for Relativistic Electron Detection. *Review of Scientific Instruments*, 81(3), 2010. 172, 173

[317] K. Svensson. *Laser-Produced Plasmas for Particle Acceleration*. PhD Thesis, Lund University, 2012. 172

[318] P. A. Walker. *Laser Wakefield Acceleration of Electrons to GeV Energies and Temporal Laser Pulse Compression Characterization in a Capillary Discharge Waveguide*. PhD Thesis, University of Oxford, 2013. 172

[319] Y. Glinec, J. Faure, A. Lifschitz, et al. Direct Observation of Betatron Oscillations in a Laser-Plasma Electron Accelerator. *Europhysics Letters*, 81(6):64001, 2008. 173, 174, 176

[320] S. V. Bulanov, M. Yamagiwa, T. Zh. Esirkepov, et al. Spectral and Dynamical Features of the Electron Bunch Accelerated by a Short-Pulse High Intensity Laser in an Underdense Plasma. *Physics of Plasmas*, 12(7):073103, 2005. 173

[321] C. D. Decker, W. B. Mori, K.C. Tzeng, et al. The Evolution of UltraIntense, ShortPulse Lasers in Underdense Plasmas. *Physics of Plasmas*, 3(5):2047–2056, 1996. 175

[322] D. Freedman, R. Pisani, and R. Purves. *Statistics (International Student Edition)*. Norton & Company, New York, fourth edition, 2007. 180