Imperial College
London

Diagnosis and Application of Laser Wakefield Accelerators

Jason Mark Cole

April 2016

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

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

Abstract

This thesis is concerned with experimental diagnostic methods of the laser-plasma interaction in laser wakefield accelerators, and how high energy photon beams from such accelerators may be exploited in a potential application.

Raman scattered light from the 20 TW Astra laser pulse was found to have shorter wavelength than expected given the interferometrically-measured plasma density, due to the relativistic motion of plasma electrons. Simulations indicated that the scattering helps to shape the pulse, removing laser energy not trapped in the wake. The signature of Raman scatter from the 200 TW Astra-Gemini pulse was instead the generation of extended plasma filaments, the direction of which indicated weaker relativistic effects. Altering the focal spot position or quality generated more filaments and reduced electron beam charge, linking Raman scattering, pulse focussing, and accelerator performance.

When a low intensity ($10^{10}\,\mathrm{Wcm}^{-2}$) probe pulse crossed a high intensity ($10^{19}\,\mathrm{Wcm}^{-2}$) drive pulse of equal frequency, the probe was observed to be amplified in intensity by up to 10^5 over a distance of tens of microns. No frequency shift of the probe was observed, and the effect required the polarisations of the pulses to be parallel. In simulations the pulses created a ponderomotive grating, trapping electrons which were able to scatter the drive into the probe in the superradiant regime of Raman amplification.

Laser wakefield-driven x-ray beams were used to perform a microtomographic scan of a human femoral trabecular bone sample. Under optimal conditions x-ray beams containing $1.3 \pm 0.5 \times 10^9$ photons were produced on 97% of laser shots, with critical energy $33 \pm 12\,\mathrm{keV}$ and source size of 2–3 µm. Image resolution was $36 \pm 7\,\mu\mathrm{m}$ and after 180 image acquisitions the 3D resolution of the tomogram was up to 50 µm. The average photon flux between 10 and 100 keV was $5.9 \pm 2.4 \times 10^5\,\mathrm{phs}^{-1}\mathrm{mrad}^{-2}$ which is comparable to microfocus sources capable of few-micron source sizes.

C Contents

Abstract 2

Role of the Author 10

Copyright 10

Declaration 10

Acknowledgements 11

1 Introduction 13

1.1 Particle accelerators 13
1.2 Plasma-based electron acceleration 15
1.3 X-ray light sources 19
1.4 Towards LWFA applications 21
1.5 Thesis outline 22

2 Theory 23

2.1 Plasma parameters 23
2.2 Gaussian laser beams 25
2.3 Single particle motion 28
2.4 Relativistic plasma optics 33
2.5 Wakefield generation 36
2.6 Limits on acceleration 47
2.7 Electron injection schemes 49
2.8 Betatron radiation 51
2.9 Parametric instabilities 54

3 Methods 69

3.1 Gas targets 69
3.2 Shadowgraphy 70
3.3 Interferometry 73
3.4 Electron spectrometer 93
3.5 Numerical simulation of laser-plasma interactions 100
3.6 Tomographic reconstruction 105

4 Stimulated Raman Scattering **115**
4.1 Spectral measurements of side scatter 115
4.2 Direct imaging of scatter 131
4.3 Conclusion 144

5 Interaction Between Crossing Laser Pulses **145**
5.1 Experimental results 145
5.2 Simulation results 152
5.3 Discussion 168
5.4 Conclusion 174

6 Bone Tomography **175**
6.1 Experimental motivation 175
6.2 Experimental setup 179
6.3 Results 181
6.4 Discussion 192
6.5 Conclusion 204

7 Conclusion **205**
7.1 Raman scattering as a diagnostic 205
7.2 Pulse amplification in a laser wakefield accelerator 206
7.3 Clinical applications of betatron radiation 207

A Appendix **208**

L List of figures

1 A brief summary of the subsequent chapters 12

1.1 Particle accelerators 14
1.2 The laser wakefield acceleration mechanism 16
1.3 Increase of laser wakefield electron energy with time 17
1.4 Laser wakefield accelerators compared to conventional accelerators 19

2.1 Gaussian beam 26
2.2 Aberrated focal spots 28
2.3 Drift velocities in plane wave 32
2.4 Self-focussing 34
2.5 Pulse compression 35
2.6 1D nonlinear wake 41
2.7 Nonlinear wake trajactories 41
2.8 Bubble 43
2.9 Injection of electrons into a bubble 45
2.10 Self-injection threshold 47
2.11 PIC simulation of pulse compression 48
2.12 Betatron spectrum 53
2.13 3 and 4 wave scattering processes 59
2.14 Plasma wave magnitude in SRS. 61
2.15 Solution of SRS dispersion relation 62
2.16 SRS growth rates 63
2.17 Regimes of instability 66
2.18 Effective growth rate as a function of damping rate 68
2.19 Spatiotemporal evolution under different damping rates 68

3.1 Shadowgraphy 71
3.2 Ray deflections in shadowgraphy 73
3.3 Numerical shadowgrams 73
3.4 Interferometry concept 74
3.5 Mach-Zehnder interferometer 75
3.6 LWFA interferogram 76
3.7 Interferogram Fourier transform. 78
3.8 Filtered interferograms 79
3.9 Hilbert transform 79

3.10 Suitable functions satisfying Bedrosian's theorem 80
3.11 Hilbert phase retrieval 80
3.12 Hilbert phase retrieval failure 81
3.13 Hilbert phase retrieval example 81
3.14 The real part of the Morlet wavelet 83
3.15 CWT scalogram 84
3.16 2D CWT phase retrieval 85
3.17 Comparison of phase retrieval methods 85
3.18 Wrapped phase 86
3.19 Phase map residue 87
3.20 Goldstein unwrap 88
3.21 Fourier unwrap 89
3.22 Abel transform 89
3.23 The first 10 basis functions $g_k(\rho)$ 92
3.24 Abel inversion by series expansion 92
3.25 Assumed density profile for gas jet 92
3.26 Fitted density profiles for gas jet 93
3.27 Electron spectrometer schematic 94
3.28 Image warping co-ordinates 96
3.29 Example image warping 97
3.30 General image transformation algorithm 98
3.31 Electron tracking through magnet 99
3.32 Experimental electron spectrometer 100
3.33 Experimental electron spectra 100
3.34 PIC algorithm 102
3.35 Numerical v_g for the 1D 2^{nd}-order EPOCH solver 103
3.36 Numerical v_g for the 1D 4^{th}-order EPOCH solver 104
3.37 Fourier slice theorem 106
3.38 Shepp-Logan phantom 107
3.39 900 projection reconstruction 108
3.40 45 projection reconstruction 108
3.41 Limit on reconstruction fidelity 109
3.42 Variation of error with projection number 109
3.43 Blurring by PSF of 10 pixels FWHM 110
3.44 Limit on reconstruction fidelity with finite PSF 110
3.45 Limit on reconstruction fidelity with spectral fluctuations 111
3.46 Streak artefact removal 112
3.47 Ring artefact removal 113
3.48 Beam hardening 113
3.49 Dishing artefacts 114

4.1 Plan view of the beamline layout for the 2012 Astra experiment. 116
4.2 In-vacuum Mach-Zehnder interferometer layout. 118

4.3 Low-density phase retrieval 118
4.4 A selection of transverse probing images displaying interesting features. 119
4.5 Side-scatter spectra 121
4.6 Side-scatter spectra as a function of density 121
4.7 Representative solution of dispersion relation 122
4.8 Measured frequency shifts 123
4.9 Measured bandwidths 124
4.10 Raman-scattered energy 126
4.11 Electron density plots from Astra simulation 127
4.12 Signature of Raman scattering in electron density 128
4.13 Laser pulse evolution in Astra simulation 129
4.14 Spectra of scattered light in Astra simulation 130
4.15 Plan view of the beamline layout for the 2012 Gemini experiment. 132
4.16 Observed scattered plasma channels 133
4.17 Side-scattered energy from the Gemini experiment 134
4.18 Maximum growth direction for RFS 135
4.19 Measured scatter angles 136
4.20 Small filaments 136
4.21 Filament occurence versus focal plane 137
4.22 Electron spectra versus focal plane 138
4.23 Focal spots with added coma 139
4.24 Filament presencse versus focal spot shape 139
4.25 Side and up-scattered light as a function of density 140
4.26 Laser pulse evolution from Gemini simulation 141
4.27 Scattered light level from Gemini simulation 142
4.28 Electron density from simulation of asymmetric focal spot 143

5.1 Shadowgrams from probe timing scan 147
5.2 Change in position of side scatter with density 148
5.3 Spatially-resolved spectrum of scattered light 149
5.4 Spatially-resolved spectra with and without plasma 149
5.5 Integrated spectra with and without plasma 150
5.6 Mean wavelength of scattered light 150
5.7 Scattered light as a function of crossing point 152
5.8 Electron density from CW simulation 153
5.9 Geometry for intersection of two beams forming a grating 154
5.10 Parallel and perpendicular momenta 155
5.11 Grating phase space 156
5.12 Electron motion 157
5.13 $(x_\perp, p_\perp)$ phase space on a logarithimic colour scale 158
5.14 Electron oscillation frequency as a function of oscillation amplitude 158
5.15 Change in Poynting flux with time 160
5.16 Electron density plots at increasing intensities 161

5.17 Amplification as a function of density and intensity 162
5.18 Temporal profile of amplified pulse 163
5.19 Amplification as a function of a_0 164
5.20 Spatially-resolved amplification of probe 165
5.21 Spectra of amplified pulses 166
5.22 Gain width as a function of intensity 166
5.23 Comparison of simulated to theoretical growth rates 167
5.24 Survey of Raman amplification experiments 173

6.1 Betatron imaging beamline layout 179
6.2 Gas cell 180
6.3 Femoral bone sample 181
6.4 Electron spectra 182
6.5 X-ray beam divergence 183
6.6 Correlation between experimental parameters 184
6.7 X-ray images from the tomographic scan 185
6.8 Locating wire tip 186
6.9 Motion of x-ray source 187
6.10 Rotation axis measurement 187
6.11 Calculating image offsets 188
6.12 Background removal 189
6.13 Sinogram filtering 190
6.14 Orthogonal slices of the 3D dataset 191
6.15 Renderings of the reconstruction 192
6.16 Brightness and photon flux of the source 192
6.17 SNR of absorption image 195
6.18 Synthetic absorption images 196
6.19 Comparison of projection to commercial microtomograph 196
6.20 Comparison of reconstruction to commercial microtomograph 197
6.21 Layout of back-illumination system 197
6.22 Spatial resolution as a function of imaging geometry 199
6.23 Spatial resolution for smaller pixels 199
6.24 Spatial resolution for a longer beamline 200
6.25 Beam-hardening correction 201
6.26 Beam-hardening correct example 201

A.1 Landau damping rates 209
A.2 Bedrosian's theorem 210
A.3 1D fluid expansion 211

L List of tables

4.1 Change in feature as a function of focal spot shape 140

5.1 Summary of parameters for the three experiments 146
5.2 A survey of plasma Raman amplification experiments 172

A.1 Shepp-Logan phantom parameters 211

Role of the Author

The experimental results presented in this thesis were gathered over the course of five experiments at the Central Laser Facility, Rutherford Appleton Laboratory, UK. The author designed, assembled and analysed the transverse diagnostics in all experiments. The author also played a leading role in the design and construction of the electron beam diagnostics in most experiments, and assisted with the experimental planning and construction of the main laser beamlines as part of the experimental team. During the tomography experiment the author assumed the role of deputy target area operator, a role which together with the target area operator assumes responsibility for the safety and day-to-day operation of the experiment. A significant contribution to the thesis were the results of some analysis of the betatron beam properties recorded during this experiment, performed by Jonathan Wood and reproduced with permission, described in chapter 6.

The author performed and analysed the particle-in-cell simulations described in the thesis, using the EPOCH code on the CX1 supercomputing cluster provided by the Imperial College HPC department.

Copyright

The copyright of this thesis rests with the author and is made available under a Creative Commons Attribution Non-Commercial No Derivatives licence. Researchers are free to copy, distribute or transmit the thesis on the condition that they attribute it, that they do not use it for commercial purposes and that they do not alter, transform or build upon it. For any reuse or redistribution, researchers must make clear to others the licence terms of this work.

Declaration

I herewith certify that all material in this dissertation which is not my own work has been properly acknowledged.

Jason Mark Cole

Acknowledgements

Before starting this thesis I had only the vaguest notion of what a laser or plasma was, never mind what you would want to do with one. After a fruitful summer spent in the basement during a 2010 undergraduate project, I discovered that I liked working with the Imperial plasma group very much, that a laser-wakefield accelerator was an interesting idea, and began an unhealthy relationship with Matlab, interferometry, and SCR muffins[1].
I've learned a tremendous amount over the past four years, and for that I have to thank my supervisors Stuart Mangles and Zulfikar Najmudin. Your patient explanation of physical concepts and experimental techniques has stood me in good stead in the lab and during the writing of this thesis. I'll feel lucky if I experience mentorship as useful and directed again in my career.
I've also felt fortunate to have worked alongside such a great bunch of people in the office and during experiments. Much of the work I've been involved in is a group effort, so to George, Martin, Christos, Drew, Steve, Jon, Kris, Nelson, Matt, Nick, Mike, Will, Tony, and many others over the years in different institutions and countries, thank you for the camaraderie and help despite the occasionally fraught nature of experiments!
To my family, thank you for your continuing enthusiasm and support for my efforts despite the polite incomprehension on what I actually *do* during all those months spent in rural Oxfordshire. You've always encouraged me to pursue my interests and trust that life will sort itself out along the way.
And thank you Hannah for seemingly endless cups of coffee, sandwiches, if anything too much cake, and for selflessly sleeping in on the weekends so I could analyse simulations undisturbed. Moving to London and doing a PhD has been challenging, but much less jarring than had you not been there with me.

[1] And footnotes

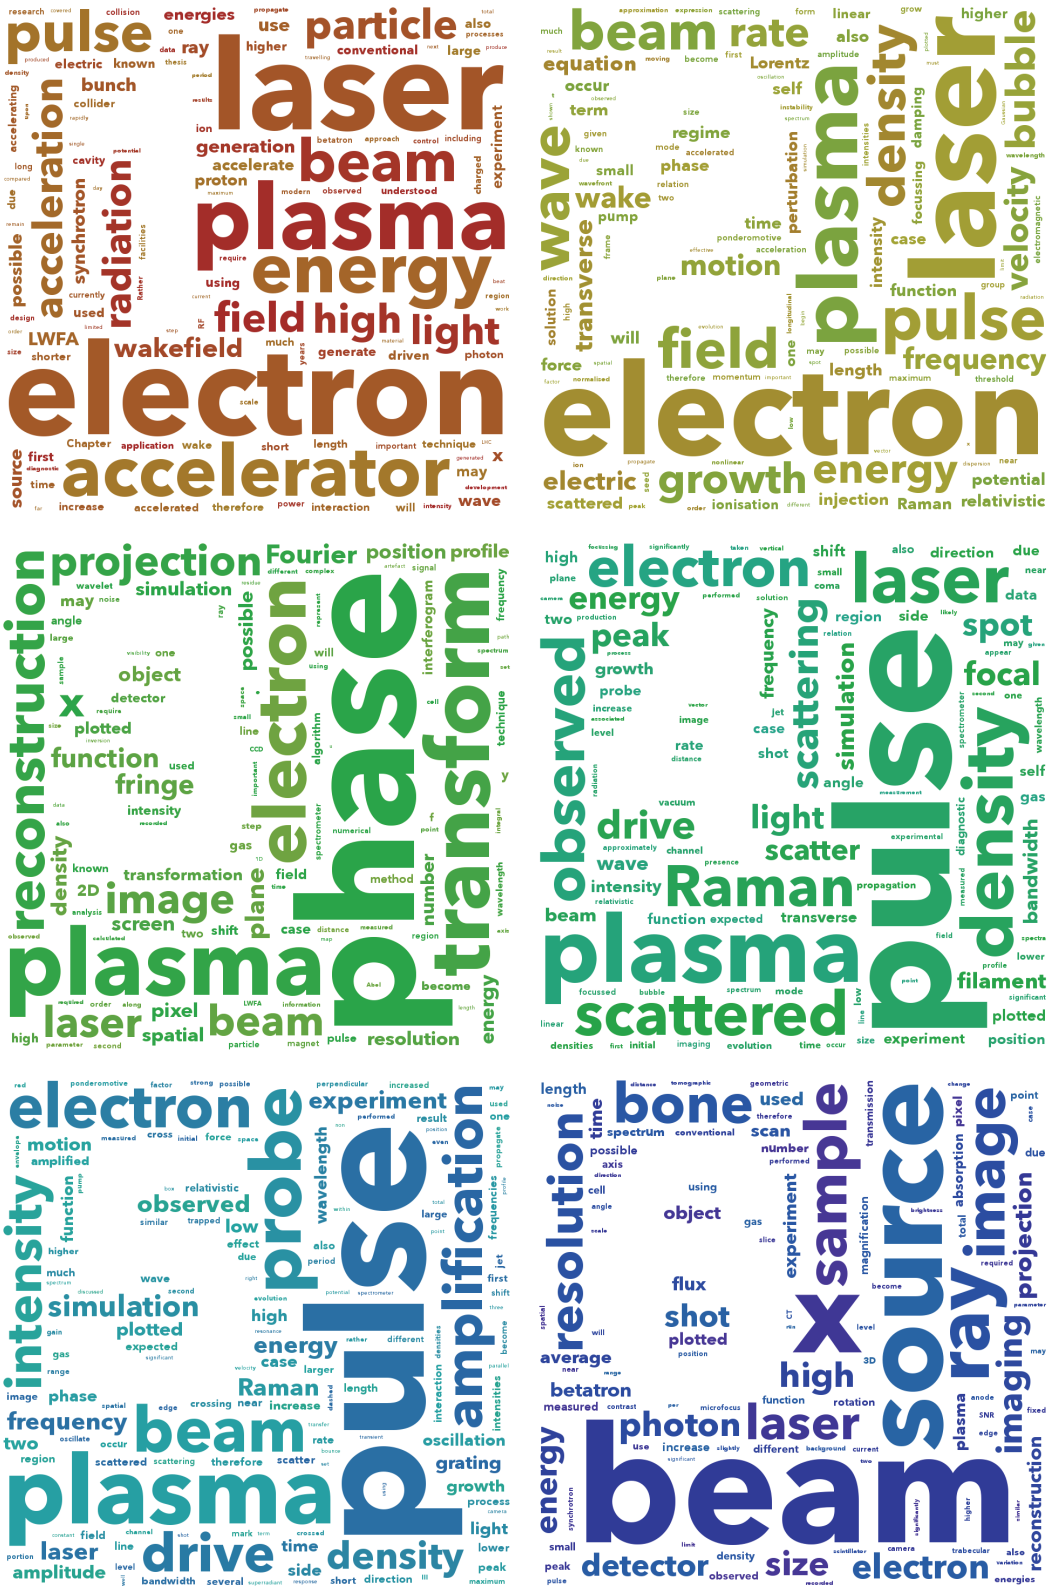


Figure 1: The first six chapters are represented here as 'word clouds', where the size of each word is proportional to the frequency of its occurrence.

1 Introduction

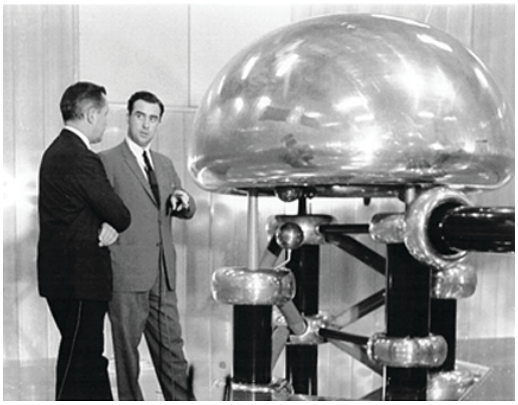
There is no better way to probe matter on the finest scales than to use high energy particle and photon beams. Rutherford calculated the size of the nucleus with a beam of helium nuclei[1], Crick, Watson, and Franklin unravelled the structure of DNA using x-ray diffraction[2], and more recently the collision of proton beams confirmed the existence of the Higgs boson[3], completing the Standard Model. These fundamental physics discoveries, heralded by a host of Nobel prizes, would not be possible without the decades of development of the particle accelerator. Particle accelerators have become more energetic over time, but also more prolific. They are fielded in hospitals to diagnose and treat patients[4], in semiconductor plants to build microchips[5], and at synchrotron light sources to generate x-ray beams[6].

At the energy frontier current accelerator technologies require progressively larger facilities to accelerate particles to higher energies, culminating in the 14 TeV Large Hadron Collider, which at 8.6 km in diameter and £3 billion in capital costs is one of the largest and most expensive machines ever constructed. The subject of this thesis is a complementary particle acceleration technique called laser-wakefield acceleration (LWFA), which promises to reduce the size of particle accelerators a thousand-fold or more. Given the ubiquity of accelerators in modern research the prospect of smaller and cheaper accelerator facilities is an enticing one, with the potential to have an impact on the broadest range of scientific disciplines.

1.1 Particle accelerators

Charged particles are accelerated when placed in an electric field, gaining an energy equal to the potential difference across the accelerating region. The first generation of particle accelerators were the Cockroft-Walton generators[7], consisting of a capacitor bank which produced a high-voltage electric discharge. While successful in accelerating helium nuclei and splitting the atom, this approach did not scale to accelerating voltages beyond 3 MV. The following years witnessed the rapid development of particle accelerator technology following technological advances in superconducting magnet design and radiofrequency

[1] [1] E. Rutherford. *Philosophical Magazine*. (1911).
[2] [2] J. D. Watson *et al. Nature*. (1953).
[3] [3] The ATLAS Collaboration. *Physics Letters B*. (2012); [4] The CMS Collaboration. *Physics Letters B*. (2012).
[4] [5] M. Silari. *Radiation Protection Dosimetry*. (2011).
[5] [6] J. F. Gibbons. *Proceedings of the IEEE*. (1968).
[6] [7] D. H. Bilderback *et al. Journal of Physics B*. (2005).
[7] [8] J. D. Cockroft *et al. Proceedings of the Royal Society*. (1932).



Cockroft-Walton generator



Stanford Linear Accelerator

Figure 1.1: While early particle accelerators were modest in size, the dimensions of modern conventional facilities are measured in kilometres.

generators[8]. Rather than using one very large electric field to accelerate particles in a single step, modern accelerators work incrementally by applying large voltages across small regions to gradually accelerate the particles. This is most simply achieved in a circular geometry, where the same fields are applied multiple times to the particle beam as it circulates and gradually gains energy. Such machines are known as synchrotrons and hold the record for the largest particle energies, including the 100 GeV Large Electron Positron collider (LEP) and the 14 TeV Large Hadron Collider (LHC).

Unfortunately, a particle of mass m travelling in a circular trajectory of radius R at energy E radiates synchrotron radiation of power $P \propto E^4/(m^4R^2)$ due to its continuous transverse acceleration. The m^{-4} dependence means that for a given energy electrons radiate at a power $1836^4 \approx 10^{13}$ that of protons, and synchrotron radiation is the limiting factor on electron acceleration. As the electron energy increases, the emitted synchrotron power increases rapidly as E^4 and eventually the acceleration ceases – the accelerator is unable to deliver energy to the electron more rapidly than it is being radiated. The solution is to increase the radius of the trajectory R, and to date the most energetic electron accelerator was LEP, built in the tunnel currently housing the LHC. If the same accelerator technology were extended to a ring circling the planet, the synchrotron radiation losses would limit the electron beam energy to 2.7 TeV. Accelerators capable of TeV energies therefore accelerate protons or ions, but this is unappealing from a particle physics standpoint because the internal structure of the proton is complex and currently only understood empirically. In addition the particle collisions occur between the quark or gluon constituents rather than the entire proton, and the centre-of-mass collision energy is reduced. The interpretation of proton collision results is then complicated and the collision inefficient, and researchers would prefer to collide simpler fundamental particles such as electrons. Another possibility is to accelerate muons, which are simple leptons like electrons but 200 times more massive. Significant challenges remain in the capturing and cooling of muons before they decay however,[9] so this is not an approach which will yield a collider in the short term.

[8] [9] A. R. Steere. *A Timeline of Particle Accelerators.* (2005).
[9] [10] M. M. Alsharo'a *et al. Physical Review ST-AB.* (2003).

Attaining TeV-scale electron beam energies therefore requires the use of a linear accelerator design where synchrotron radiation is not an issue. The planned International Linear Collider[10] (ILC) is based on such a scheme, with an initial design energy of 500 GeV and the possibility of a future increase to 1 TeV. It will use state-of-the-art accelerating structures based upon superconducting radiofrequency (RF) cavities, where a standing electromagnetic wave accelerates particles as they traverse the cavity. The electric field strength inside the cavity is limited by resistive dissipation of electromagnetic energy in the walls of the cavity, but more fundamentally by the field emission of electrons from the cavity material. If the electric field strength is too high, the electrons are pulled from the cavity surface and interact with the RF field, reflecting it and potentially damaging accelerator components. For the ILC this effect limits the average accelerating field of the cavity to 31.5 MVm^{-1}, leading to a total acceleration length of 11 km, total collider length of 31(50) km for the 500(1000) GeV design, and total proposed cost of £5.3 billion.
Despite extensive research, the average accelerating field of RF cavities has not increased for several decades[11], and is currently limited to 40-50 MeVm^{-1}. Alternative 'high-gradient' acceleration methods have been proposed, including direct-laser acceleration[12] and the inverse free electron laser[13], but the technique witnessing the most intensive research is perhaps the plasma accelerator.

1.2 Plasma-based electron acceleration

A plasma is a collection of positively and negatively charged particles which exhibits collective behaviour through electromagnetic interactions. A plasma may be generated by heating a gas to the point that electrons are ionised from the gas atoms, and so is commonly understood as the fourth state of matter. However plasmas can be composed of more exotic constituents, for example electron-positron or quark-gluon plasmas, and are therefore more general than this term would suggest.
The idea of using plasma to accelerate charged particles originated in a 1956 conference proceedings by Veksler[14] and expanded upon by Faĭnberg[15], where it was noted that because plasmas are already ionised they can support very large electric fields without further damage. If such electric fields can be made to move at relativistic velocities, then an electron bunch moving with the fields is continuously accelerated to high energies. Electric fields are generated in plasmas when the negatively and positively-charged components are separated, and the use of a laser pulse to generate a relativistically-travelling charge separation was proposed in an influential 1979 paper by Tajima and Dawson[16]. If a high-intensity laser is incident upon a plasma, the radiation pressure of the pulse sweeps the electrons aside, but the high-inertia ions remain approximately static. Soon afterwards, if the laser pulse is short the radiation pressure will have ceased and the electrons now feel an

[10] [11] T. Behnke *et al. ILC Technical Design Report Volume 1: Executive Summary.* (2013).
[11] [12] K. Saito *et al. The 10th Workshop on RF Superconductivity.* (2001).
[12] [13] E. Peralta *et al. Nature.* (2013).
[13] [14] R. B. Palmer. *Journal of Applied Physics.* (1972); [15] J. Duris *et al. Nature Communications.* (2014).
[14] [16] V. I. Veksler. *Proc. CERN Symp.* (1956).
[15] [17] Y. B. Fainberg. *Plasma Physics Reports.* (1967).
[16] [18] T. Tajima *et al. Physical Review Letters.* (1979).

electrostatic restoring force due to the remaining ions. The electrons are accelerated back but, like a simple harmonic oscillator, overshoot and continue in oscillatory motion at the characteristic plasma frequency ω_p. Associated with the electron motion is a density wave trailing the laser pulse and moving at the velocity of the pulse. A longitudinal electric field is generated by this wave due to Gauss' law, also moving with the pulse. This is known as the laser wakefield, analogously to a water wake driven by a boat. The acceleration scheme is therefore known as laser wakefield acceleration, illustrated schematically in Fig. 1.2. When this idea was proposed it was known that the laser pulse would need to be very short in order to generate a wakefield efficiently, far shorter than was possible at the time.

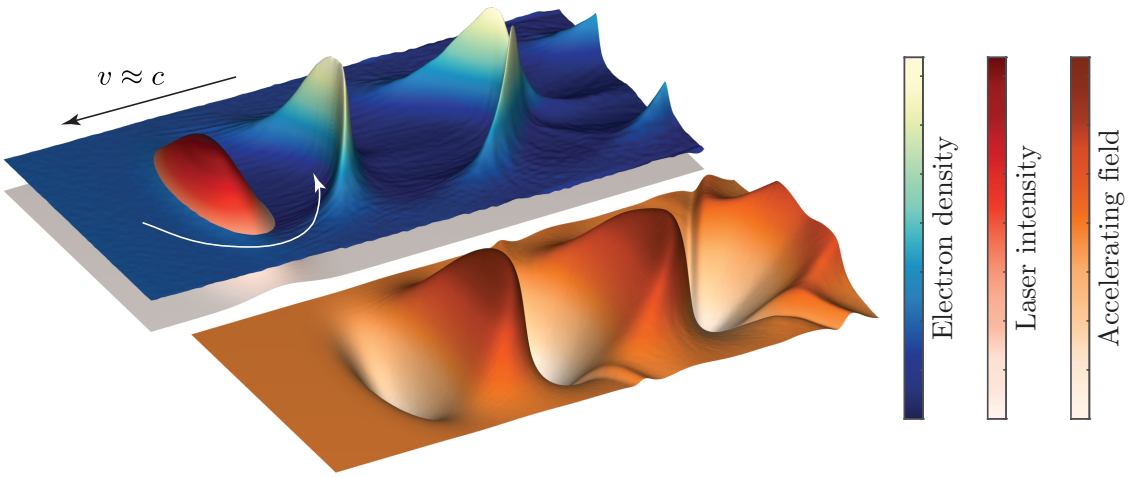


Figure 1.2: The laser wakefield acceleration mechanism, illustrated using data from a simulation. The short and highly intense laser pulse travelling at a velocity near the speed of light (red) forces plasma electrons aside (white arrow), before the electrons return due to the Coulomb attraction of the undisturbed plasma ions. The electron density wave (green) is associated with strong longitudinal accelerating fields (orange) which can trap and accelerate the electrons.

The first generation of laser-plasma accelerators used instead a scheme called plasma beat-wave acceleration[17] (PBWA) where two long laser pulses of frequencies ω_0 and ω_1 propagate colinearly such that their envelope is modulated through interference at the beat frequency $\Delta\omega = \omega_0 - \omega_1$. If $\Delta\omega \approx \omega_p$ then the beat wave can resonantly excite a large-amplitude plasma wave[18] and accelerate electrons[19]. At large wave amplitudes the frequency of the plasma wave changes, causing the resonance to 'de-tune' and the acceleration to be quenched. To overcome this issue the next generation of experiments operated in the self-modulated laser wakefield regime (SM-LWFA)[20], where a single long laser pulse develops modulation at ω_p through nonlinear optical processes in the plasma, and can therefore stay in resonance naturally. The first experiments on SM-LWFA

[17] [19] M. N. Rosenbluth. *Physical Review Letters.* (1972).
[18] [20] C. E. Clayton *et al. Physical Review Letters.* (1985).
[19] [21] C. E. Clayton *et al. Physical Review Letters.* (1993).
[20] [22] N. E. Andreev. *JETP Letters.* (1992); [23] J. Krall *et al. Physical Review E.* (1993).

demonstrated higher electron energy gains[21], extending up to the 100 MeV level[22]. During this period laboratory-scale high power laser systems had increased in intensity by orders of magnitude[23] due to the invention of the chirped pulse amplification technique[24] (CPA) in 1985. As the available laser pulses continued to shorten, by 2002 a single laser pulse could almost fit inside a single period of the plasma wave at the plasma density required for effective acceleration. By exploiting a nonlinear optical effect known as pulse compression, the pulse can be made shorter during its interaction with the plasma and generate a high amplitude wakefield. This is known as the forced laser wakefield regime (F-LWFA), differentiated experimentally from the SM-LWFA mechanism by the lack of modulations at ω_p in spectrum of the transmitted laser light[25].

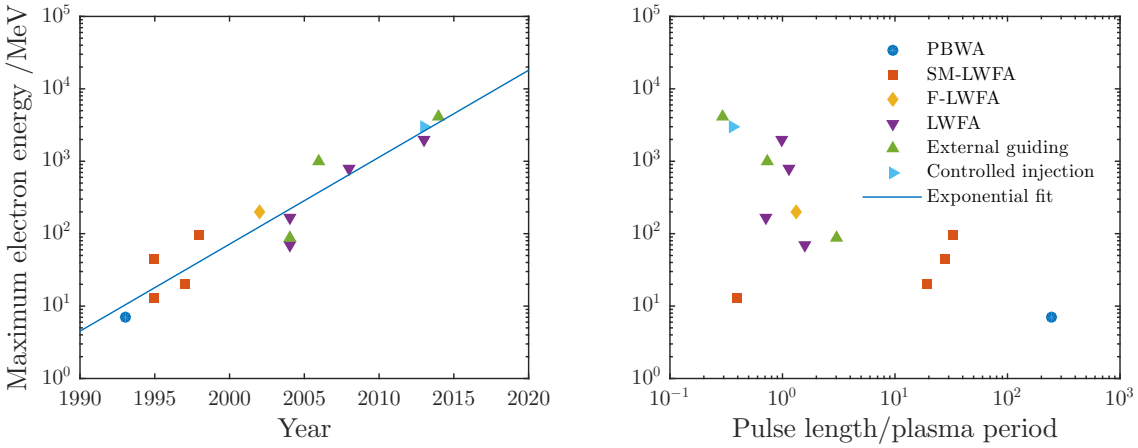


Figure 1.3: The maximum reported energies of laser-accelerated electrons in the literature are plotted as a function of publication date, and are observed to follow an exponential trend. Higher energies are correlated with pulse lengths short compared to a plasma period $2\pi\omega_p^{-1}$.

By the mid-2000s and the advent of multi-TW femtosecond lasers, available pulses were short enough and of sufficiently high intensity that the LWFA regime anticipated by Tajima and Dawson became a reality. A turning point in the field was marked by the simultaneous publication of three results in 2004 which demonstrated it was possible to generate beams of narrow energy spread[26]. This was in sharp contrast to previous results where the electron energy spectrum was broad and quasi-thermal, and an important step in proving that LWFA was capable of the production of high-quality electron beams. The measured average acceleration gradients were in the region of 100 GeVm^{-1} over acceleration lengths of millimetres, over a thousand times shorter than the acceleration length required for a state-of-the-art RF cavity. The GeV barrier was broken in 2006 where

[21] [24] K. Nakajima *et al. Physical Review Letters.* (1995); [25] A. Modena *et al. Nature.* (1995); [26] R. Wagner *et al. Physical Review Letters.* (1997).
[22] [27] D. Gordon *et al. Physical Review Letters.* (1998).
[23] [28] M. D. Perry *et al. Science.* (1994).
[24] [29] D. Strickland *et al. Optics Communications.* (1985); [30] P. Maine *et al. IEEE Journal of Quantum Electronics.* (1988).
[25] [31] V. Malka *et al. Science.* (2002).
[26] [32] S. P. D. Mangles *et al. Nature.* (2004); [33] C. Geddes *et al. Nature.* (2004); [34] J. Faure *et al. Nature.* (2004).

a density-tailored plasma waveguide was used to control the laser pulse propagation[27]. As laser facilities of higher power continue to come online the reported electron beam energies have increased, including the Astra-Gemini laser (800 MeV in 2009[28]), the Texas Petawatt (2 GeV in 2013[29]) and the Korean Petawatt (3 GeV in 2013[30]). The current record energy for LWFA electron beams is[31] 4.2 GeV, though as is suggested in Fig. 1.3, records do not last for very long.

Other characteristics of the electron bunch are important besides peak energy, and many applications for high energy electron beams require fine control of their parameters. Much experimental work in recent years has been dedicated to the control and understanding of the LWFA electron beam, including the injection[32], emittance[33], energy spread[34], charge[35], stability[36], duration[37] and tunability[38] of such beams. Once these issues are sufficiently well understood, the next step in building, for example, a particle collider, would be to chain together series of LWFA stages[39]. The coupling of two LWFA stages with a charge transfer efficiency of 3.5% has recently been demonstrated[40], which while an exciting development, demonstrates efficiencies far lower than will be required for a collider facility. The history of laser-plasma electron acceleration techniques is covered in a series of review articles by Esarey *et al.*[41].

1.2.1 Other plasma acceleration techniques

Rather than using a laser to generate a plasma wake, it is also possible to use a beam of high energy particles produced by conventional accelerators. A particle bunch with Lorentz factor $\gamma \gg 1$ travels much closer to the speed of light than a laser pulse in a plasma, and the 'witness' bunch may be accelerated to higher energies before catching up to the wake 'drive' bunch. The 42 GeV SLAC electron beam has been used to drive a 1 m-long plasma wake which accelerates both electrons[42] and positrons[43] at efficiencies exceeding[44] 30%. The electron beam of the FLASH accelerator will be used to drive a

[27] [35] W. P. Leemans *et al. Nature Physics.* (2006).
[28] [36] S. Kneip *et al. Physical Review Letters.* (2009).
[29] [37] X. Wang *et al. Nature Communications.* (2013).
[30] [38] H. T. Kim *et al. Physical Review Letters.* (2013).
[31] [39] W. P. Leemans *et al. Physical Review Letters.* (2014).
[32] [40] J. Faure *et al. Nature.* (2006); [41] S. P. D. Mangles *et al. Physical Review Letters.* (2006); [42] A. Buck *et al. Physical Review Letters.* (2013).
[33] [43] S. Kneip *et al. Physical Review ST-AB.* (2012); [44] R. Weingartner *et al. Physical Review ST-AB.* (2012); [45] M. Migliorati *et al. Physical Review ST-AB.* (2013).
[34] [46] B. B. Pollock *et al. Physical Review Letters.* (2011); [47] G. Golovin *et al. Physical Review ST-AB.* (2015).
[35] [48] C. McGuffey *et al. Physical Review Letters.* (2010).
[36] [49] S. P. D. Mangles *et al. Physics of Plasmas.* (2007); [50] J. Osterhoff *et al. Physical Review Letters.* (2008); [51] S. Banerjee *et al. Physical Review ST-AB.* (2013).
[37] [41] S. P. D. Mangles *et al. Physical Review Letters.* (2006); [52] O. Lundh *et al. Nature Physics.* (2011); [53] O. Lundh *et al. Physical Review Letters.* (2013).
[38] [54] A. J. Gonsalves *et al. Nature Physics.* (2011); [47] G. Golovin *et al. Physical Review ST-AB.* (2015).
[39] [55] W. P. Leemans *et al. Physics Today.* (2009).
[40] [56] S. Steinke *et al. Nature.* (2016).
[41] [57] E. Esarey *et al. IEEE Transactions on Plasma Science.* (1996); [58] E. Esarey *et al. Reviews of Modern Physics.* (2009).
[42] [59] I Blumenfeld *et al. Nature.* (2007).
[43] [60] S. Corde *et al. Nature.* (2015).
[44] [61] M. Litos *et al. Nature.* (2014).

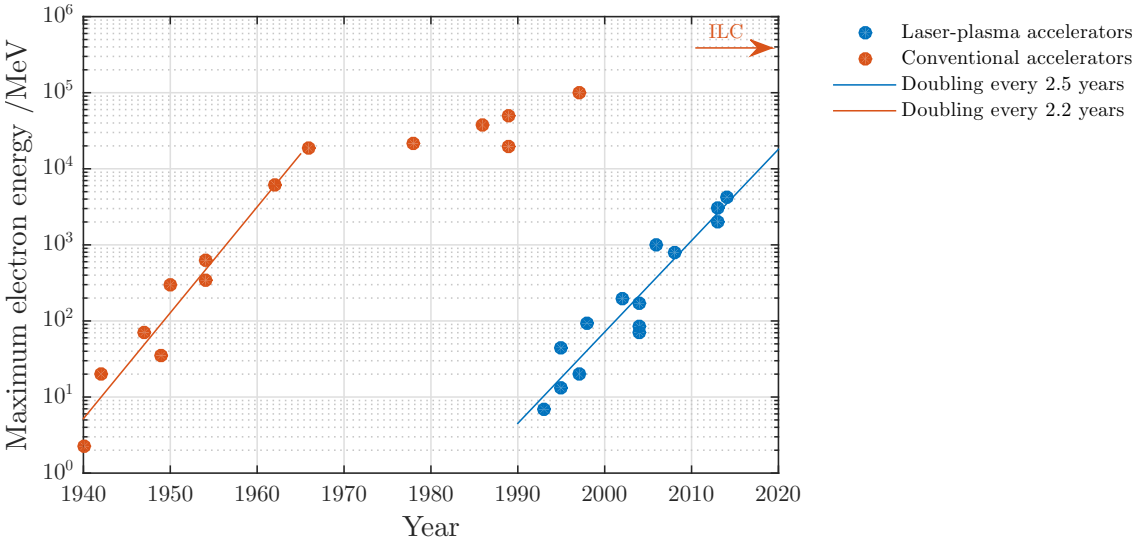


Figure 1.4: The increase with time of maximum electron energies attained with conventional and laser-plasma accelerators. In their early stages of development, both technologies develop rapidly, doubling in energy every 2 – 3 years.

wake in the FLASHForward experiment[45], and the proton beam of the Super Proton Synchrotron at CERN will uniquely be used to create a proton-driven wake during the AWAKE experiment[46]. The advantage of using a proton beam is that the drive bunch contains more energy for a given Lorentz factor, which would imply that the LHC proton beam could drive a wake which accelerates electron beams to TeV energies. Finally it is also possible to construct compact laser-plasma ion accelerators[47], which while not covered here, is a thriving research field in its own right.

1.3 X-ray light sources

Whenever a charged particle is accelerated it emits synchrotron radiation, first noticed in the visible light region at one of the early generation of electron accelerators where it was initially termed 'Schwinger radiation'[48]. As the electron energy increases so does the energy of the radiation, and modern synchrotron light sources exploit this effect to produce radiation far into the hard x-ray range. Far from a nuisance, here the synchrotron light is a useful source of photons and the primary goal of the accelerator. The first generation of experiments operated in a 'parasitic' mode where the radiation was a secondary concern compared to the primary purpose of accelerating electrons[49]. The next generation of light sources were synchrotrons specifically built for their radiation generation capabilities, and in current third-generation devices[50] such as the ESRF[51] and Diamond[52], the electron beam acceleration and storage processes are optimised to produce the highest photon

[45] [62] A. Aschikhin *et al. The FLASHForward Facility at DESY.* (2015).
[46] [63] R Assmann *et al. Plasma Physics and Controlled Fusion.* (2014).
[47] [64] H. Daido *et al. Reports on Progress in Physics.* (2012).
[48] [65] H. C. Pollock. *Am J. Phys.* (1983).
[49] [66] E. M. Rowe. *Particle Accelerators.* (1973).
[50] [67] D. H. Bilderback *et al. Journal of Physics B.* (2005).
[51] [68] S. Tazzari. *IEEE Transactions on Nuclear Science.* (1985).
[52] [69] URL: www.diamond.ac.uk. (2015).

brightnesses possible. The multi-GeV electron bunches traverse 'insertion devices' such as undulators[53] or wigglers[54] where they are forced to oscillate, emitting bright x-ray pulses along experimental beamlines. In pushing to higher x-ray brightnesses, fourth-generation light sources are built around the x-ray free electron laser concept (XFEL)[55], where a high-energy electron beam propagates linearly through a long undulator such that it may interact coherently with the radiated light field and lase. XFELs currently operating include the LCLS[56], FLASH[57], and SACLA[58], and new facilities such as the European XFEL[59], SCSS[60] and SwissFEL[61] are coming online in the next few years. The high brilliance, collimation and tunability of x-ray light sources have enabled a revolution in the diagnostic methods of x-ray diffraction, scattering, imaging and spectroscopy. These have in turn found use in the fields of medicine, materials science, condensed matter physics and biology, and demand for access to light sources continues to increase.

1.3.1 Plasma light sources

Building synchrotron light sources driven by conventional electron accelerators will remain a laborious and capital-intensive endeavour however, due to the long acceleration lengths required and necessary size of the facilities. By using a laser-wakefield accelerator to generate the electron bunch, it may be possible in the future to construct light sources which are much cheaper and smaller due to the shorter acceleration lengths.

If an electron bunch of Lorentz factor γ propagates through a region consisting of a plasma containing only positively-charged ions, the bunch naturally oscillates transversely at the frequency[62] $\omega_p/\sqrt{2\gamma}$. In the classical theory of particle accelerators this is known as betatron motion, where it is driven by gradients in the magnetic field strength of the bending magnets. In this context the synchrotron radiation emitted by an electron bunch is known as betatron radiation. The spatial oscillation period can be much shorter than that possible in conventional insertion devices, and so for the same electron beam energy, the emitted radiation has much shorter wavelength and higher energy. Betatron radiation was first observed with a conventional accelerator at SLAC[63], but it was not until high energy short laser pulses became available that it was observed in a LWFA experiment [64] where the plasma was able to both longitudinally accelerate and transversely wiggle the electron bunch. In the proceeding years the characteristics of this radiation were examined[65] and it was found to have femtosecond duration[66], small source size[67], extend to high photon

[53] [70] H. Motz *et al. Journal of Applied Physics.* (1953).
[54] [71] M. Berndt *et al. IEEE Transactions on Nuclear Science.* (1979).
[55] [72] B. W. J. McNeil *et al. Nature Photonics.* (2010).
[56] [73] P. Emma *et al. Nature Photonics.* (2010).
[57] [74] S Schreiber *et al. Proc. 1st Int. Particle Accelerator Conf.* (2010).
[58] [75] D. Pile. *Nature Photonics.* (2011).
[59] [76] A. Massimo. *Technical Design Report.* (2007).
[60] [77] T. Shintake. *Proceedings of IPAC'10.* (2010).
[61] [78] B. D. Patterson *et al. New Journal of Physics.* (2010).
[62] [79] D. H. Whittum. *Physics of Fluids B: Plasma Physics.* (1992).
[63] [80] S. Wang *et al. Physical Review Letters.* (2002).
[64] [81] A. Rousse *et al. Physical Review Letters.* (2004); [82] K. Ta Phuoc *et al. Physics of Plasmas.* (2005).
[65] [83] M. Schnell *et al. Journal of Plasma Physics.* (2015).
[66] [84] K. Ta Phuoc *et al. Physics of Plasmas.* (2007).
[67] [85] K. Ta Phuoc *et al. Physical Review Letters.* (2006); [86] F. Albert *et al. Physical Review E.* (2008).

energy[68], possess a synchrotron-like spectrum[69], and a peak brightness comparable to that of third-generation light sources[70]. It has also been shown that it is possible to use betatron radiation to diagnose properties of the LWFA-produced electron bunch[71].

Another approach to generating high-energy photons is to use the laser fields directly as the wiggler, rather than rely on the fields of the plasma. This method relies upon the interaction of the electron bunch with a counterpropagating laser pulse, either separate[72] or a reflection of the LWFA pulse[73], and can produce photons of MeV energy[74]. Other groups have incorporated aspects of conventional electron accelerators, using external undulators[75] or using a laser to wiggle an RF-accelerated electron bunch[76]. The history of laser-driven ultrafast x-ray sources is covered in a review by Corde *et al.*[77].

1.4 Towards LWFA applications

The development of laser-plasma accelerators has generated much excitement in recent decades, but in order to advance the field beyond the developmental stage it is important to investigate potential applications which leverage the unique capabilities of laser-driven electron and photon beams. This is understood by the research community[78], and the use of LWFA electron beams has been investigated in the context of radiological treatment[79], ultrafast electron diffraction[80] and the generation of dense antimatter beams[81]. Acting as a compact light source LWFA x-ray beams have been used as high-quality spatially coherent backlighters in the imaging of soft[82] and hard[83] materials, as probes in ultrafast spectroscopic measurements[84], and may have use in nuclear physics. For many applications it is important that the beams are stable and predictable, and so it is crucial that the electron beam injection and acceleration processes are well understood. For laser-driven accelerators, this means understanding how the laser pulse evolves as it propagates through the plasma, as it is the evolution of the laser pulse which dictates when and how the electrons are injected. It is therefore important that diagnostic techniques continue to develop, not only for the day-to-day operation of an accelerator but also to point towards ways in which the technique can be improved.

[68] [87] S. Kneip *et al. Physical Review Letters.* (2008).
[69] [88] S. Fourmaux *et al. New Journal of Physics.* (2011).
[70] [89] S. Kneip *et al. Nature Physics.* (2010).
[71] [43] S. Kneip *et al. Physical Review ST-AB.* (2012); [90] F. Albert *et al. Physical Review Letters.* (2013).
[72] [91] K. T. Phuoc *et al. Physical Review Letters.* (2003); [92] N. Powers *et al. Nature Photonics.* (2013); [93] K. Khrennikov *et al. Physical Review Letters.* (2015).
[73] [94] K. Ta Phuoc *et al. Nature Photonics.* (2012).
[74] [95] G. Sarri *et al. Physical Review Letters.* (2014).
[75] [96] H.-P. Schlenvoigt *et al. Nature Physics.* (2008); [97] M. Fuchs *et al. Nature Physics.* (2009).
[76] [98] K. Achterhold *et al. Scientific Reports.* (2013).
[77] [99] S. Corde *et al. Reviews of Modern Physics.* (2013).
[78] [100] F. Albert *et al. Plasma Physics and Controlled Fusion.* (2014).
[79] [101] M. Nicolai *et al. Applied Physics B: Lasers and Optics.* (2013).
[80] [102] Z. H. He *et al. Applied Physics Letters.* (2013).
[81] [103] G. Sarri *et al. Physical Review Letters.* (2013).
[82] [104] S. Fourmaux *et al. Optics Letters.* (2011); [105] S. Kneip *et al. Applied Physics Letters.* (2011); [106] Z. Najmudin *et al. Philosophical Transactions of the Royal Society.* (2014); [107] J. Wenz *et al. Nature Communications.* (2015).
[83] [108] J. M. Cole *et al. Scientific Reports.* (2015); [109] A. Döpp *et al. arXiv.* (2015).
[84] [110] M. Z. Mo *et al. Review of Scientific Instruments.* (2013).

1.5 Thesis outline

The contents of this thesis are somewhat eclectic, determined in part by the interests of the author but also by the relative success of the various experiments occurring over the previous years. The overarching theme is the diagnosis and application of laser-wakefield accelerators. Accurate diagnosis of processes occuring over micron spatial and femtosecond temporal scales is challenging, but important in understanding the physical processes at work. Once understood, it may one day be possible to move laser-plasma accelerators out of the laboratory and into the field where their impact may be most acutely felt. In this spirit, the first chapters are based around the theme of diagnostic techniques, either direct or inferred from the details of the laser-plasma interaction. The final chapters explore how the high-intensity laser-plasma interaction may be exploited for applications, either the generation of secondary intense laser pulses or stable and repeatable beams of x-ray light suitable for use in medical applications. The contents of the subsequent chapters are as follows.

Chapter 2: Theory

Theoretical results from the literature relevant to the experiments presented in this thesis are reviewed, including wakefield generation, electron acceleration, betatron radiation and laser-plasma parametric instabilities.

Chapter 3: Methods

Experimental and diagnostic methods used in the gathering and analysis of the data presented here are discussed. In particular, a thorough review of contemporary interferometric analysis techniques is included.

Chapter 4: Stimulated Raman Scattering

Results from two experiments are presented, focussing on how the scattering of laser light from a laser wakefield accelerator may be used to diagnose the laser pulse evolution.

Chapter 5: Interaction Between Crossing Laser Pulses

An unexpected interaction between laser pulses was observed during several experiments. In this chapter the data from the experiments are analysed, and interpreted with the help of a comprehensive set of simulations.

Chapter 6: Bone Tomography

The use of a laser-plasma x-ray source was explored in the context of the computed tomography of medical bone samples. The experimental data are analysed and compared to competing techniques, before a discussion on the relative merits and drawbacks of a laser-driven approach.

Chapter 7: Conclusion

Finally the conclusions of the thesis are summarised, and the prospects for future experimental investigations laid out.

2 Theory

In this chapter the necessary theory is reviewed in order to understand the formation of laser wakefields, the injection and acceleration of electron beams, and the evolution of the laser pulse through nonlinear optical processes. This basis is built upon to introduce practical considerations which must feed into the design of an accelerator, such as the limiting factors on electron acceleration, how the electron energy scales with laser and plasma parameters, and additional mechanisms for the injection of electrons. The final two sections cover the topics of betatron radiation and parametric instabilities in slightly more depth, as will required to appreciate the results of chapters 4, 5 and 6.

2.1 Plasma parameters

The plasmas supporting a laser-wakefield can be described as an ensemble of unbound electrons and heavy ions, interacting with one another electromagnetically. The ions are at least several thousand times more massive than the electrons, so it is the motion of electrons which dominates the response of a plasma to perturbations on a short timescale. An important example is the perturbation due to a transverse electromagnetic wave $\mathbf{E} = E_0 \cos(k_0 z - \omega_0 t)\hat{\mathbf{x}}$ entering a uniform plasma of electron density n_e. Assuming the electron velocity $\mathbf{v} \ll c$, the $\mathbf{v} \times \mathbf{B}$ force may be neglected and the motion of an electron is governed by the Lorentz force equation

$$m_e \ddot{x} = -eE_x \quad \Rightarrow \quad x = \frac{e}{m_e \omega_0^2} E_x \,. \tag{2.1}$$

The polarisation density P_x of the plasma, neglecting the motion of the ions, is given by

$$P_x = -en_e x = -\frac{n_e e^2}{m_e \omega_0^2} E_x \equiv \epsilon_0 \chi_e E_x \tag{2.2}$$

where χ_e is the (assumed isotropic) susceptibility of the plasma electrons. The relative permittivity $\epsilon = 1 + \chi_e$ and refractive index $\eta = \sqrt{\epsilon}$ are

$$\epsilon = 1 - \frac{n_e e^2}{\epsilon_0 m_e \omega_0^2} = 1 - \frac{\omega_p^2}{\omega_0^2} \quad \Rightarrow \quad \eta = \sqrt{1 - \frac{\omega_p^2}{\omega_0^2}} \tag{2.3}$$

where the plasma frequency ω_p has been defined as

$$\omega_p = \sqrt{\frac{n_e e^2}{m_e \epsilon_0}}. \tag{2.4}$$

This is the natural oscillation frequency of the electrons in a plasma – consider the

infinitesimal displacement of a sheet of electrons by δx. Over δx there is an electric field due to the imbalance between electron and ion densities. From Gauss' law, $E_x = \int -en_e/\epsilon_0 \, \mathrm{d}x \approx -en_e\delta x/\epsilon_0$. The Lorentz force on the electron sheet implies $m_e\ddot{\delta x} = -e^2 n_e \delta x/\epsilon_0$, which indicates simple harmonic motion at ω_p, known as a plasma electron or Langmuir wave.

The plasma wavelength is defined as $\lambda_p = 2\pi c/\omega_p$, which corresponds to the wavelength of a plasma wave driven by a perturbation moving at the speed of light through the plasma. The presence of the electron mass in the expression for ω_p justifies the assumption of a static ion background – the typical response time for an ion $\omega_{pi}^{-1} = (m_i\epsilon_0/n_i e^2)^{1/2}$ is $\sqrt{m_i/m_e}$ higher than that of an electron ($\omega_{pi}^{-1} \approx 800$ fs for a hydrogen plasma at $n_e = 10^{18}$ cm^{-3}).

The dispersion relation of an electromagnetic wave is given by $\omega_0 = ck_0/\eta$, and thus the phase and group velocities v_ϕ and v_g in a plasma are

$$v_\phi = \frac{\omega_0}{k_0} = \frac{c}{\sqrt{1-\frac{\omega_p^2}{\omega_0^2}}} \qquad v_g = \frac{\mathrm{d}\omega_0}{\mathrm{d}k_0} = c\sqrt{1-\frac{\omega_p^2}{\omega_0^2}} \qquad \gamma_g = \frac{\omega_0}{\omega_p} \tag{2.5}$$

where γ_g is the Lorentz factor corresponding to v_g. The electron density at which $\omega_p = \omega_0$ is known as the critical density n_c where electromagnetic waves cease to propagate, and is given by

$$n_c = \frac{m_e\epsilon_0\omega_0^2}{e^2}\,. \tag{2.6}$$

Plasmas with $n_e < n_c$ are 'underdense' and transmit laser light, which is the regime used in laser wakefield accelerators. Plasmas with $n_e > n_c$ are conversely 'overdense', often used in laser-driven ion acceleration. For a laser wavelength of 800 nm the critical density is 1.75×10^{21} cm^{-3}, approximately 70 times standard atmospheric density.

2.1.1 Plasma formation

An upper limit on the electric field required to ionise a material into a plasma can be estimated by the barrier-suppression ionisation mechanism. Given an electron sitting in an atomic Coulomb potential at depth $-V_e$, barrier suppression occurs when the magnitude of the electric field is sufficiently large to pull the Coulomb potential down to $-V_e$ in the vicinity of the electron. Consider as an example an atom of hydrogen in an electric field E. The potential V felt by the electron is

$$V(r) = -\frac{e}{4\pi\epsilon_0 r} - Er$$

with a peak in potential at radius $r_{\min} = (e/4\pi\epsilon_0 E)^{1/2}$. For barrier suppression ionisation to occur, $V(r_{\min}) \leq V_e$, or equivalently $E \geq \pi\epsilon_0 V_e^2/e^3$. For hydrogen $V_e = 13.6$ eV and complete ionisation will occur when $E \geq 3.2 \times 10^{10}$ Vm^{-1}, which corresponds to a laser intensity of 2.7×10^{14} Wcm^{-2}. The plasma will begin to form at fields lower than this, but above this intensity threshold hydrogen can be expected to be fully ionised. Typical peak intensities for laser-plasma acceleration experiments are above 10^{18} Wcm^{-2}, so the

plasma can usually be considered fully-formed by the time the main laser pulse arrives and ionisation effects ignored.

However, when using higher-Z materials or less intense lasers this assumption can fail to hold and ionisation can play an important part in the acceleration process, as discussed in section 2.7.1. There a more sophisticated model such as the Ammosov-Delone-Krainov scheme for tunnel-ionisation must be used to model the ionisation process[1]. A useful engineering formula for the ionisation rate R predicted by the ADK model is[2]

$$R[\mathrm{fs}^{-1}] \approx 1.52 \frac{4^n V_e[\mathrm{eV}]}{n\Gamma(2n)} \left(20.5 \frac{V_e[\mathrm{eV}]^{3/2}}{E[\mathrm{GV/m}]} \right)^{2n-1} \exp\left(-6.83 \frac{V_e[\mathrm{eV}]^{3/2}}{E[\mathrm{GV/m}]} \right) \tag{2.7}$$

where V_e is the ionisation potential of the electron, E is the laser electric field and $n = 3.69Z/\sqrt{V_e[\mathrm{eV}]}$ is the effective quantum number of the electron. Given a constant[3] ionisation rate R, the probability P that an electron is ionised during a period Δt is $P = 1 - e^{-R\Delta t}$, reaching 50% for $\Delta t = (\ln 2)/R$. A laser electric field of $E = 1$ GVm^{-1} corresponds to an intensity of $1.3{\times}10^{11}$ Wcm^{-2}, which is many orders of magnitude lower than the peak focussed intensities attainable at modern laser facilities. For the first electron of helium ($V_e = 23.6$ eV, $Z = 1$) the ionisation rate reaches 10 fs^{-1} at a laser intensity of $5{\times}10^{15}$ Wcm^{-2}. For the 6$^{\mathrm{th}}$ electron of nitrogen ($V_e = 552$ eV, $Z = 6$) reaching this ionisation rate requires a much higher laser intensity of $1.4{\times}10^{19}$ Wcm^{-2}.

2.2 Gaussian laser beams

The spatial profile of a laser beam produced by a solid-state laser cavity can be calculated from the resonant TEM modes of that cavity, and are known as gaussian modes[4]. These modes are important to understand as they retain their gaussian profile as they propagate, and are usually transformed by optical elements into related gaussian modes. A gaussian beam is a natural solution of the paraxial approximation to Maxwell's equations in cylindrical symmetry – consider

$$\nabla \times \nabla \times \mathbf{E} = \left\{ \begin{array}{c} \nabla(\nabla \cdot \mathbf{E}) - \nabla \cdot (\nabla \mathbf{E}) = \nabla^2 \mathbf{E} \\ -\mu \frac{\partial}{\partial t}(\nabla \times \mathbf{H}) = -\mu\epsilon \frac{\partial^2 \mathbf{E}}{\partial t^2} \end{array} \right\} \Rightarrow -\mu\epsilon \frac{\partial^2 \mathbf{E}}{\partial t^2} = \nabla^2 \mathbf{E}\,.$$

Assuming an electric field of the form $\mathbf{E} = \mathbf{E}_0(r,z)e^{i(kz-\omega t)}$ and making the approximation that $|k\mathbf{E}_0| \gg |\partial_z \mathbf{E}_0|$, this reduces to the paraxial form of the Helmholtz equation,

$$\left(\frac{\partial^2}{\partial r^2} + \frac{1}{r}\frac{\partial}{\partial r} - 2ik\frac{\partial}{\partial z} \right) \mathbf{E}_0 = 0\,. \tag{2.8}$$

If the electric field is assumed linearly-polarised and composed of a single transverse component E_x, the solution of Eq. 2.8 is

$$E_x(r,z,t) = E_x^{(0)} \frac{w_0}{w(z)} \exp\left(i(\omega t - kz) - \frac{r^2}{w^2(z)} - ik\frac{r^2}{2R(z)} + i\zeta(z) \right)\,. \tag{2.9}$$

[1] [111] M. Ammosov *et al. Sov Phys JETP.* (1986).

[2] [112] D. L. Bruhwiler *et al. Physics of Plasmas.* (2003).

[3] For non-constant $R(t)$, one can derive $P(t) = 1 - \exp\left(\int_{-\infty}^{t} R(t')\,\mathrm{d}t'\right)$

[4] [113] W. Chang. *Principles of Lasers and Optics.* (2005).

There are a number of important parametrisations in this expression which are related to physical characteristics of the beam. Fig. 2.1 plots the magnitude of the electric field of this solution, recognisable as an electromagnetic wave converging to a finitely-sized focus and expanding out again, the near and far-field profiles remaining gaussian.

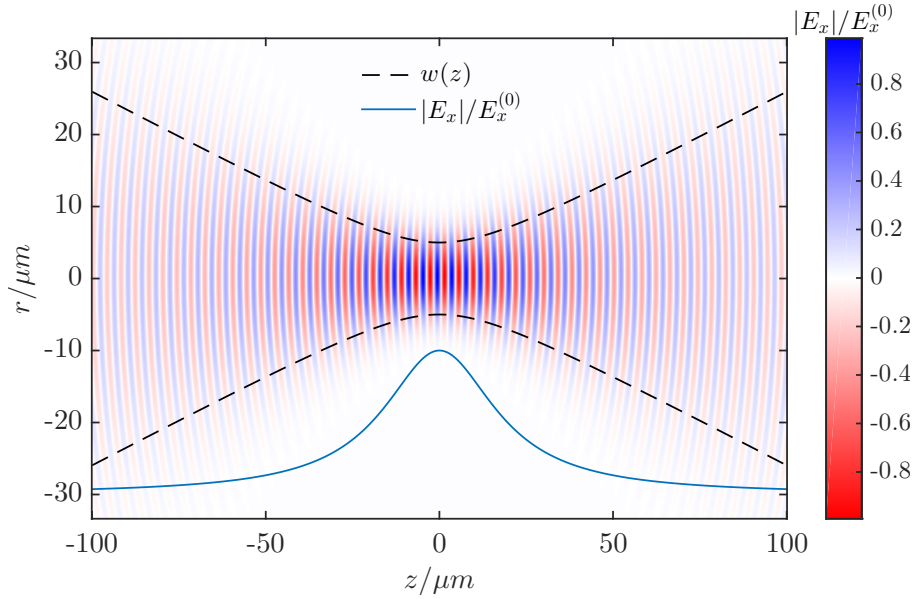


Figure 2.1: The gaussian beam solution for $E_x(z, r)$, with $\lambda = 800$ nm and $w_0 = 5$ μm.

Working through the expression, $E_x^{(0)}$ is the maximum amplitude of the electric field at focus and w_0 is the minimum $1/e^2$ intensity radius of the beam – the beam waist which defines the focal spot size. The beam width parameter $w(z)$ depends on the beam propagation direction z through $w(z) = w_0(1 + z^2/z_R^2)^{1/2}$, where the Rayleigh range is defined as $z_R = \pi w_0^2/\lambda$. The Rayleigh range of a beam is defined as the propagation distance required for the beam $1/e^2$ intensity radius to expand to $\sqrt{2}w_0$, and parametrises the distance over which the beam stays 'in focus'. In the context of a laser-plasma acceleration experiment the longer the Rayleigh range, the longer a laser pulse can remain intense and drive an accelerating plasma wave. However for a given laser pulse energy a smaller waist is required for higher peak intensity, reducing the Rayleigh range. Nonlinear focussing can extend the effective Rayleigh range – see section 2.4. The divergence angle of the beam $\partial_z w$ can be used to relate the f number of the focussing optic to the minimum beam waist as $w_0 = 2f\lambda/\pi$. Here the f number is defined as the focal length of the optic divided by the diameter of the laser beam on the optic, which increases with longer focal lengths and smaller beams.

The radius of curvature of the wavefronts is $R(z) = z(1 + z_R^2/z^2)$, and the paraxial approximation amounts to the condition that $R \gg \lambda$. The final phase function $\zeta(z) = \tan^{-1}(z/z_R)$ is known as the Gouy phase, which arises from the transverse spatial confinement of the beam[5] and causes an additional phase shift of π as the beam passes through focus. The gaussian beam solution is then specified entirely once w_0, λ and $E_x^{(0)}$ are known.

It is worth commenting on the validity of the above derivation when applied to tightly-

[5] [114] S. Feng *et al. Optics Letters.* (2001).

focussed beams. Calculating $\frac{\partial E_x}{\partial z}|_{r=0}$ (i.e. along the optical axis) and applying the paraxial condition at $z = 0$ yields

$$\left|\frac{\partial E_x}{\partial z}\right|^2_{r=0} = \left(\frac{1}{z^2+z_R^2} - \frac{2kz_R}{z^2+z_R^2} + k^2\right)|E_x|^2 \ll k^2|E_x|^2 \Rightarrow 2kz_R \gg 1, \quad w_0 \gg \frac{\lambda}{2\pi}.$$

Equivalent statements to this condition are that $z_R \gg \lambda/4\pi$ and that the minimum radius of curvature of the wavefronts $R_{\min} \gg \lambda/2\pi$. In ultrahigh intensity laser-plasma experiments, often the beam waist can reach 2.5 µm, $\approx 20 \times\lambda/2\pi$ which is perhaps approaching the validity of the gaussian approximation above and may motivate a higher-order description of the electric fields of the laser pulse. In particular a linearly-polarised pulse must contain a longitudinal component from the first of Maxwell's equations

$$\nabla \cdot \mathbf{E} = 0 \Rightarrow E_z = \int -\frac{\partial E_x}{\partial x}\,\mathrm{d}z$$

which is neglected in the solution above. Transverse and longitudinal components of gaussian beams have been calculated to several orders by Chen *et al.*[6], derived from an exact integral representation of the laser fields. The structure of a gaussian laser pulse near focus is also derived by Quesnel and Mora[7].

2.2.1 Wavefront aberrations

The expression for the focussing gaussian beam assumes a perfect gaussian intensity profile at the position of the focussing optic, along with a perfectly flat wavefront. In reality the beam profiles of high-intensity lasers are typically as uniform as possible so the beam is amplified evenly across its profile. Irregularities in the amplitude and wavefront arise from imperfections in the flatness and surface quality of the beam optics, and nonlinearities in the pulse amplification process. The spatial distribution of the phase deviations in the near-field can be decomposed into a sum over the orthonormal basis functions on the unit circle, known as Zernike polynomials. The electric field of the pulse in the near-field in polar coordinates is then represented as

$$E(r,\theta) = |E(r,\theta)| \exp\left(ik\sum_j a_j Z_j(r,\theta)\right) \tag{2.10}$$

where the coefficients a_j are often expressed in units of the laser wavelength. The first few Z_j correspond to specific classical aberrations, and are defined as

$$\begin{array}{ll} \text{Spherical} & \sqrt{5}(6\rho^4 - 6\rho^2 + 1) \\ \text{Astigmatism} & \sqrt{6}\rho^2\cos\theta \\ \text{Coma} & \sqrt{8}(3\rho^2 - 2)\rho\cos\theta \end{array} \tag{2.11}$$

where ρ is normalised such that $\rho = 1$ corresponds to the edge of the beam, and $\int_0^{2\pi}\int_0^1 Z_j^2 \rho\,\mathrm{d}\rho\mathrm{d}\theta = \pi$. Fig. 2.2 displays the calculated form of focal spots exhibiting

[6] [115] C. G. Chen *et al. JOSA A.* (2002).
[7] [116] B. Quesnel *et al. Physical Review E.* (1998).

these aberrations. Spherical tends to increase the prominence of the rings of the Airy pattern, astigmatism longitudinally separates the tangential and sagittal focal planes, and coma increases the asymmetry of the spot. Any aberration will reduce the peak intensity at focus, enlarge the size of the spot, and redistribute pulse energy outside of the spot, so should be mitigated as much as possible. One approach is the use of adaptive optics to selectively delay or advance spatial portions of the laser pulse. If the wavefront can be measured, an adaptive optic can work in a feedback loop to minimise wavefront distortions. Alternatively the feedback loop may respond to other parameters such as the pulse duration or accelerated electron beam properties, as has shown to be effective in the on-line optimisation of high-repetition-rate LWFA systems [8]. The effects of an aberrated wavefront on acceleration have been investigated experimentally[9] and numerically[10] finding that the electron beam charge and energy are lowered, especially when the laser pulse energy is not well confined to a central focal spot.

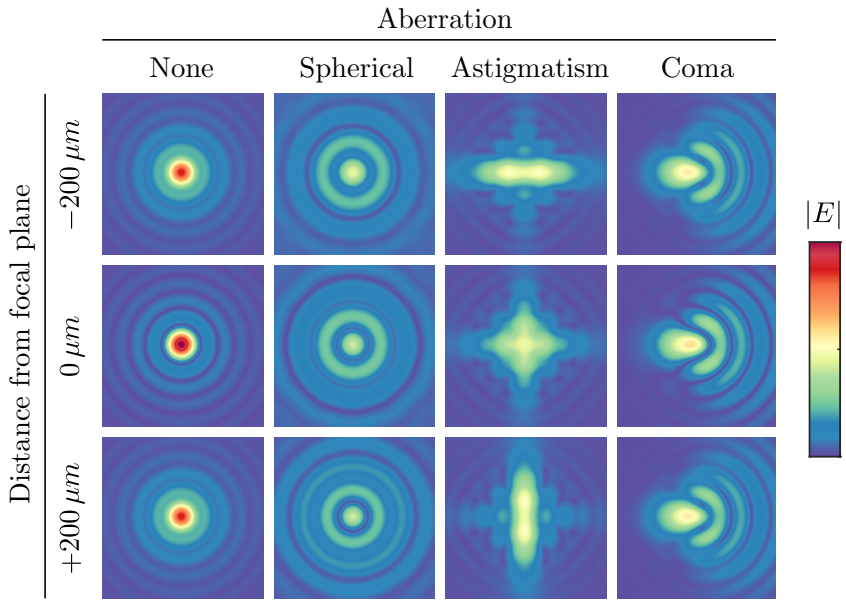


Figure 2.2: Focal spots corresponding to the $f/10$ focussing of a flat-top near field profile, with various wavefront phase errors added at the plane of the focussing optic. The colour scale represents the amplitude of the transverse laser electric field.

2.3 Single particle motion

Consider an infinite linearly-polarised plane wave described by the electric and magnetic fields

$$\mathbf{E} = E_0 \sin(kz - \omega t)\hat{\mathbf{x}} \tag{2.12}$$

$$\mathbf{B} = B_0 \sin(kz - \omega t)\hat{\mathbf{y}}\,. \tag{2.13}$$

[8] [117] Z. H. He *et al. Nature Communications.* (2015).

[9] [118] S. P. D. Mangles *et al. Applied Physics Letters.* (2009); [119] S. P. D. Mangles *et al. Physical Review ST-AB.* (2012); [120] B. Beaurepaire *et al. Physical Review X.* (2015).

[10] [121] J. Vieira *et al. Plasma Physics and Controlled Fusion.* (2012).

If the electron is assumed to be non-relativistic, the magnetic force can be neglected as

$$|\mathbf{v} \times \mathbf{B}| = \frac{v}{c} E_0 \ll E_0$$

and in this case the electron velocity is given by

$$v_x = -\frac{e}{m_e} \int E_0 \sin(kz - \omega t)\, \mathrm{d}t = -\frac{eE_0}{m_e \omega} \cos(kz - \omega t) \equiv -a_0 c \cos(kz - \omega t)$$

where the definition has been made

$$a_0 \equiv \frac{eE_0}{m_e c \omega}. \tag{2.14}$$

This dimensionless quantity is very important in determining the dynamics of plasmas under irradiation by intense laser pulses, as it directly determines the electron velocity in the laser fields. In the literature, a_0 is often interpreted as a normalised electron oscillation velocity $a_0 \equiv v_{\mathrm{osc}}/c$. If the electric field is expressed in terms of a vector potential $\mathbf{E} = -\partial_t \mathbf{A}$, then analogously $A_0 = E_0/\omega = m_e c a_0/e$ and hence a_0 also represents the normalised vector potential of the plane wave. The intensity of a plane wave is related to its electric field by $I = \epsilon_0 c E_0^2/2$, from which follows a convenient engineering formula for the normalised vector potential

$$a_0 \approx 0.855\, \lambda[\mu\mathrm{m}] \sqrt{I[10^{18}\ \mathrm{Wcm}^{-2}]}\,. \tag{2.15}$$

For a laser of wavelength 800 nm, a_0 reaches 1 for $I > 2{\times}10^{18}$ Wcm^{-2}, so for such intensities the relativistic motion of the electron in the laser fields must be taken into account. Here we follow the explicit method of Yang *et al.*[11]. The equations of motion for the electron are

$$\frac{\mathrm{d}\mathbf{x}}{\mathrm{d}t} = \frac{\mathbf{p}}{m_e \gamma} \tag{2.16}$$

$$\frac{\mathrm{d}\mathbf{p}}{\mathrm{d}t} = -e\left(\mathbf{E} + \frac{\mathbf{p}}{m_e \gamma} \times \mathbf{B}\right) \tag{2.17}$$

where $\mathbf{p}$ is the electron momentum and $\gamma = \sqrt{1 + \mathbf{p} \cdot \mathbf{p}/(m_e^2 c^2)}$ is the electron Lorentz factor. Expanding the momentum equation explicitly and taking $B_0 = E_0/c$

$$\frac{\mathrm{d}p_x}{\mathrm{d}t} = -eE_0 \sin(kz - \omega t)\left(1 - \frac{p_z}{\gamma m_e c}\right) \tag{2.18}$$

$$\frac{\mathrm{d}p_y}{\mathrm{d}t} = 0 \tag{2.19}$$

$$\frac{\mathrm{d}p_z}{\mathrm{d}t} = -eE_0 \sin(kz - \omega t)\frac{p_x}{\gamma m_e c}\,. \tag{2.20}$$

Without loss of generality take $p_y = 0$, and from the definition of γ

[11] [122] J.-H. Yang *et al. Plasma Physics and Controlled Fusion.* (2011).

$$\frac{\mathrm{d}p_x}{\mathrm{d}t} = \frac{m_e^2 c^2 \gamma\dot{\gamma} - p_z \dot{p}_z}{p_x}\,. \tag{2.21}$$

Substitution of this identity into the momentum equations yields the conservation law

$$\frac{\mathrm{d}}{\mathrm{d}t}\left(\gamma - \frac{p_z}{m_e c}\right) = 0\,. \tag{2.22}$$

If the electron begins its motion at rest, then the relation $p_z/(m_e c) = \gamma - 1$ holds throughout its subsequent motion. Inserting this into the definition of γ provides a second relation for the transverse momentum component

$$\frac{p_x}{m_e c} = \sqrt{2(\gamma - 1)}\,. \tag{2.23}$$

With these definitions, the momentum equations can be transformed into a single ordinary differential equation for γ as a function of proper time s. An infinitesimal interval of proper time is related to that of laboratory time by $\mathrm{d}s' = \mathrm{d}t'/\gamma(t')$. Assume the electron begins its motion at $s = 0$ at the laboratory co-ordinates (t_0, z_0), which implies the specification of an initial laser phase $\phi_0 = kz_0 - \omega t_0$. Laboratory time t at the electron can then be expressed as

$$\int_{t_0}^{t} \mathrm{d}t' = \int_0^s \gamma(s')\,\mathrm{d}s' \quad \Rightarrow \quad t = t_0 + \int_0^s \gamma(s')\,\mathrm{d}s' \tag{2.24}$$

and longitudinal position z of the electron as

$$z = z_0 + \int_{t_0}^{t} v_z(t')\,\mathrm{d}t' = z_0 + \int_{t_0}^{t} \frac{p_z(t')}{m_e c \gamma(t')}\,\mathrm{d}t' = z_0 + \int_0^s \frac{p_z(s')}{m_e}\,\mathrm{d}s'. \tag{2.25}$$

Combining these expressions shows that the phase of the wave as witnessed by the electron is a function of its proper time only

$$\begin{aligned} kz - \omega t &= kz_0 + \frac{k}{m_e}\int_0^s p_z(s')\,\mathrm{d}s' - \omega t_0 - \omega \int_0^s \gamma(s')\,\mathrm{d}s' && (2.26)\\ &= \phi_0 - \omega \int_0^s \gamma - \frac{p_z}{m_e c}\,\mathrm{d}s' = \phi_0 - \omega s && (2.27) \end{aligned}$$

where the conservation equation 2.22 is used in the second line. Substituting 2.22, 2.23 and 2.27 into 2.20 produces

$$\frac{1}{\omega a\sqrt{2(\gamma-1)}}\frac{\mathrm{d}(\gamma-1)}{\mathrm{d}s} = \pm\sin(\phi_0 - \omega s) \tag{2.28}$$

where a is the normalised vector potential of the plane wave. Taking $\gamma = 1$ at $s = 0$, an explicit expression for $\gamma(s)$ is

$$\gamma = 1 + \frac{a^2}{2}\left(\cos(\phi_0 - \omega s) - \cos\phi_0\right)^2. \tag{2.29}$$

Averaging Eq. 2.29 over a laser period, the average electron Lorentz factor is

$$\langle\gamma\rangle = 1 + \frac{a^2}{4}\left(1 + 2\cos^2\phi_0\right). \tag{2.30}$$

This is an important quantity to know as the effective mass of an electron inside a laser pulse will be increased by $\langle\gamma\rangle$, and therefore the effective plasma frequency decreased as $\langle\gamma\rangle^{-1/2}$. $\mathbf{p}(\gamma)$ is known, and the momentum components are therefore expressible in proper time as

$$\frac{p_x}{m_e c} = a\left(\cos(\phi_0 - \omega s) - \cos\phi_0\right) \tag{2.31}$$

$$\frac{p_z}{m_e c} = \frac{a^2}{2}\left(\cos(\phi_0 - \omega s) - \cos\phi_0\right)^2. \tag{2.32}$$

Both momenta are observed to possess constant components, and therefore electrons in a plane wave undergo an average drift. To find the drift velocity, note that $\frac{\mathrm{d}\mathbf{x}}{\mathrm{d}s} = \mathbf{p}(s)/m_e$, the components of which can be integrated to find

$$x = \frac{ca}{\omega}\left(\sin(\phi_0 - \omega s) - \sin\phi_0 + \omega s\cos\phi_0\right) \tag{2.33}$$

$$z = z_0 + \frac{a^2 c}{4} s\left(1 + 2\cos^2\phi_0\right) - \frac{a^2 c}{8\omega}\left(\sin(2\omega s + 2\phi_0) - \sin 2\phi_0\right) + \frac{a^2 c}{\omega}\cos\phi_0\left(\sin(\omega s + \phi_0) - \sin\phi_0\right). \tag{2.34}$$

From 2.27, $t(s)$ can be calculated from the relation

$$t = t_0 + s + \frac{z - z_0}{c}. \tag{2.35}$$

After a full laser period as witnessed by the electron (i.e. $\Delta s = 2\pi/\omega$), the electron co-ordinates have changed by

$$\Delta x = a\lambda\cos\phi_0 \tag{2.36}$$

$$\Delta z = \frac{a^2\lambda}{4}\left(1 + 2\cos^2\phi_0\right) \tag{2.37}$$

$$\Delta t = \frac{\lambda}{c}\left(1 + \frac{a^2}{4}\left(1 + 2\cos^2\phi_0\right)\right) \tag{2.38}$$

and so the drift velocities and drift Lorentz factor γ_d are

$$\beta_x = \frac{\Delta x}{c\Delta t} = \frac{4a\cos\phi_0}{4 + a^2\left(1 + 2\cos^2\phi_0\right)} \tag{2.39}$$

$$\beta_z = \frac{\Delta z}{c\Delta t} = \frac{a^2\left(1 + 2\cos^2\phi_0\right)}{4 + a^2\left(1 + 2\cos^2\phi_0\right)} \tag{2.40}$$

$$\gamma_d = \frac{4 + a^2\left(1 + 2\cos^2\phi_0\right)}{\sqrt{8(2 + a^2)}} \tag{2.41}$$

These quantities are plotted in Fig. 2.3 as a function of a and ϕ_0.

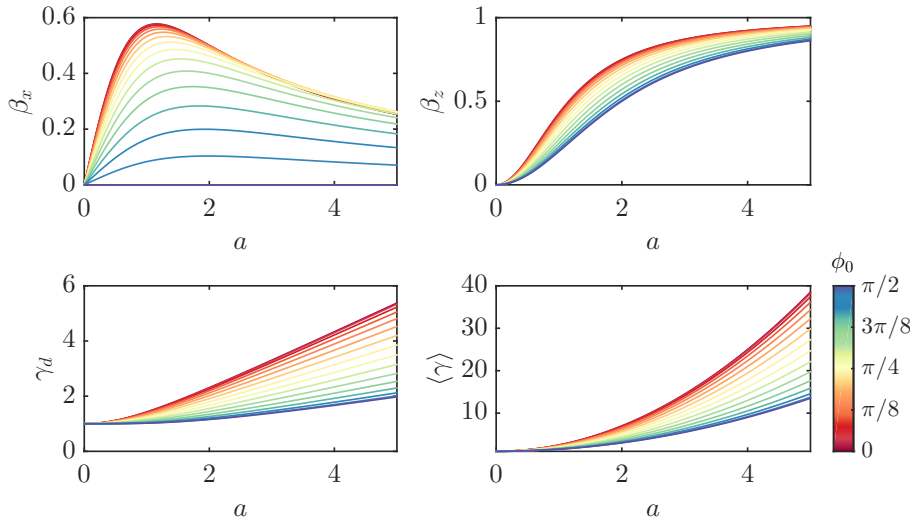


Figure 2.3: Transverse and longitudinal drift velocities of an electron of initial phase ϕ_0 in a plane wave of strength a.

2.3.1 Ponderomotive force

In addition to the drift motion of electrons in a plane wave, their transverse or 'quiver' motion due to the laser fields can be thought of as a source of potential energy, and therefore constitute a potential field. Like any other potential an associated force is directed down the potential gradient, which in this case is termed the ponderomotive force. In the non-relativistic limit, the electron motion consists of fast oscillatory motion (x_f), and slow ponderomotive drift motion (x_s) such that the Lorentz force equation reads

$$\ddot{x}_s + \ddot{x}_f = c\omega a(x_s + x_f)\cos(\omega t). \tag{2.42}$$

The oscillatory motion causes small excursions about the drift trajectory, so $|x_f| \ll |x_s|$ and $a(x_s + x_f) \approx a(x_s)$ is approximately constant for the fast motion. One can also take $\ddot{x}_f \gg \ddot{x}_s$ and integrate to find the fast motion explicitly, $x_f \approx (ca(x_s)/\omega)\cos(\omega t)$. The next order of approximation requires a Taylor expansion of a, $a(x_s+x_f) \approx a(x_s)+x_f a'(x_s)$. Substituting these expressions into 2.42

$$\ddot{x}_s = c\omega a(x_s)\cos(\omega t) - c^2 a(x_s)a'(x_s)\cos^2(\omega t) \tag{2.43}$$

and so averaging the ponderomotive force $F_p = m_e\ddot{x}_s$ over an optical cycle

$$\langle F_p\rangle = -\frac{m_e c^2}{2} a\frac{\mathrm{d}a}{\mathrm{d}x} = -\frac{m_e c^2}{4}\frac{\mathrm{d}a^2}{\mathrm{d}x} \tag{2.44}$$

where the identity $(a^2)' = 2aa'$ was used. The same expression is recovered if one takes the gradient of the quiver kinetic energy of the electron. The ponderomotive force acts to move electrons away from regions of high intensity ($a^2 \propto I$), so a laser pulse propagating through a plasma will push electrons and ions aside – this is an example of an electromagnetic interaction acting in the same direction on both positively and negatively charged species.

However the ponderomotive acceleration of a particle of mass m scales as m^{-2} (due to the mass dependence of the normalised vector potential) so is it the electrons which are predominantly swept aside by the laser pulse, while the ions can be regarded as a static background. The electromagnetic fields generated by the motion of the electrons will perturb the ions on a timescale longer than the laser pulse, but they can usually be considered static over several plasma periods $T_p = 2\pi\omega_p^{-1}$ (111 fs for $n_e = 10^{18}$ cm^{-3}). The relativistic expression for the ponderomotive force is[12]

$$\mathbf{F}_p = -\frac{m_e c^2}{\gamma}\nabla a^2 \tag{2.45}$$

where γ is the Lorentz factor of the electron due to its motion in the laser fields.

2.4 Relativistic plasma optics

At the high laser intensities encountered in LWFA, the usually transparent plasma medium drives a number of pulse modification processes, commonly grouped together as the relativistic nonlinear optical effects of plasmas. An overview is presented by Mori[13], where it is shown that a host of nonlinear processes nevertheless conserve the classical action (photon number) of the laser pulse, i.e.

$$\langle a^2\rangle\omega_0 w^2\tau = \text{const.} \tag{2.46}$$

where a is the normalised vector potential, ω_0 is the frequency, w is the pulse waist and τ is the pulse duration.

2.4.1 Relativistic self-focussing

As a function of the normalised vector potential a, $\langle\gamma\rangle$ varies transversely across the laser pulse and therefore so does the electron mass and relativistic plasma frequency $\omega_p/\langle\gamma\rangle^{1/2}$. The refractive index of the plasma η is then also a function of transverse position and the plasma can begin to act as a focussing device.

To see how this occurs, consider the flat wavefront of an electromagnetic wave as it propagates through its focal plane (e.g., the form of Eq. 2.9 at $z = 0$). The wavefront will move away at its local phase velocity, which will be a function of position – Fig. 2.4 shows how this leads to a focussing curvature of the wavefront. As the beam is near focus, $\delta\theta$ is small, so $\tan(\delta\theta) \approx \delta\theta = \partial_r v\delta t$. The appropriate velocity here is the phase velocity $v_\phi = c/\eta$ so $\delta\theta = -(c/\eta^2)\partial_r\eta\delta t$. There is now a component of the laser energy travelling perpendicularly to the propagation axis which can decrease the effective beam width. This transverse energy propagation occurs at the local group velocity $v_g = c\eta$, producing a plasma-induced change in beam width $\delta w_{\text{plasma}} = -v_g\sin(\delta\theta)\delta t \approx (c^2/\eta)\partial_r\eta\delta t^2$.

Natural diffraction of the spot counteracts this effect, and it is when these two processes are balanced that the critical conditions for self-focussing occur. Eq. 2.9 contains the beam width factor $w(z)$ which dictates the natural beam size as a function of propagation distance, and thus the rate of change of beam width due to diffraction. Expanding $w(z)$

[12] [123] P. Mora *et al. Physics Of Plasmas.* (1997); [116] B. Quesnel *et al. Physical Review E.* (1998).
[13] [124] W. B. Mori. *IEEE Journal of Quantum Electronics.* (1997).

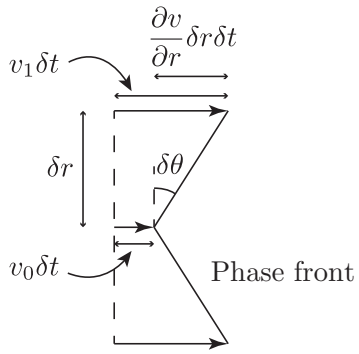


Figure 2.4: As a wavefront propagates to the right over a time δt, variations in the local phase velocity transverse to propagation may lead to a focussing effect on the wavefront

about focus

$$\delta w_{\text{diffract}} = w(\delta z) - w_0 \approx \delta z \left.\frac{\partial w}{\partial z}\right|_{z=0} + \frac{\delta z^2}{2}\left.\frac{\partial^2 w}{\partial z^2}\right|_{z=0} + \cdots = \frac{w_0}{z_R^2}\delta z^2 = \frac{4c^2}{\omega^2 w_0^3}\delta z^2 \,.$$

Taking $\delta z \approx c\delta t$, the threshold for self-focussing is found by allowing the plasma effects to balance those from diffraction

$$\delta w_{\text{plasma}} + \delta w_{\text{diffract}} = 0 \quad \Rightarrow \quad \frac{c^2}{\eta}\frac{\partial \eta}{\partial r} + \frac{4c^4}{\omega^2 w_0^3} = 0 \,.$$

To make progress, the assumption can be made that the self-focussing threshold occurs when the beam is weakly relativistic and $a_0 \ll 1$. The plasma may also be taken to be sufficiently underdense that $\omega_p \ll \omega$. The refractive index can be approximated as

$$\eta \approx 1 - \frac{\omega_p^2}{2\omega^2\langle\gamma\rangle} = 1 - \frac{\omega_p^2}{2\omega^2}\left(1 + \frac{a^2}{2}\right)^{-1} \approx 1 - \frac{\omega_p^2}{2\omega^2}\left(1 - \frac{a^2}{2}\right)$$

then the self-focussing threshold may be reformulated as a condition on the transverse laser intensity gradient

$$\frac{\omega_p^2}{4\omega^2}\frac{\partial a^2}{\partial r} + \frac{4c^2}{\omega^2 w_0^3} = 0 \,.$$

Taking $\partial_r a^2 \approx a_0^2/w_0$ shows that the threshold is determined by the combination $a_0^2 w_0^2$ which represents the power of the pulse. The end result is that there is a density-dependent laser power threshold P_c above which self-focussing occurs[14]

$$P_c = \frac{8\pi\epsilon_0 m_e^2 c^5}{e^2}\frac{\omega^2}{\omega_p^2} \approx 17.5\frac{n_c}{n_e}\,\text{GW}, \tag{2.47}$$

leading to the waist evolution equation[15]

[14] [125] P. Sprangle *et al. IEEE Transactions on Plasma Science.* (1987).
[15] [57] E. Esarey *et al. IEEE Transactions on Plasma Science.* (1996).

$$\frac{\mathrm{d}^2 w}{\mathrm{d}z^2} = \frac{w_0^2}{z_R^2 w^3}\left(1 - \frac{P}{P_c}\right). \tag{2.48}$$

It is perhaps surprising at first sight that the laser power, not the laser intensity, determines the onset of self-focussing. The reason being that were the beam focussed more loosely, resulting in a lower peak intensity, diffraction of the beam would be less severe. The transverse intensity gradient then need not be as steep to overcome the diffraction effects, allowing the lower intensity beam to self-focus as before.
Self-focussing has been found experimentally to allow the guiding of laser pulses over many Rayleigh ranges[16] with the laser spot oscillating in size as it propagates[17].

2.4.2 Pulse compression

Local variations in group velocity need not arise from purely relativistic effects, if the plasma electron density varies in space the shape of the laser pulse can be altered. Following Mori[18] a simple argument may be presented for the temporal compression of a laser pulse as it propagates up an electron density gradient. This case is important as the ponderomotive force of the laser expels electrons from the focal volume, so the rear of the pulse experiences a lower electron density than the front and the pulse witnesses a continual co-moving density gradient. Consider a laser pulse of length L as in Fig. 2.5, propagating in the x-direction with front and back group velocities v_{g2} and v_{g1} respectively. The change in length of the pulse after a time Δt will be $\Delta L = (v_{g2} - v_{g1})\Delta t \approx L\partial_x v_g \Delta t$.

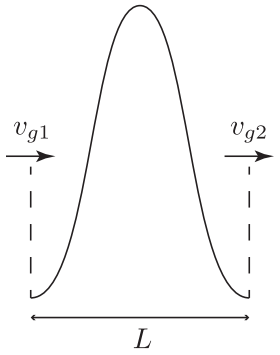


Figure 2.5: A laser pulse propagating through an inhomogeneous plasma at different local group velocities

Outside the main body of the pulse the ponderomotive motion of the electrons may be ignored and the variation in group velocity be taken to arise from the electron density gradient

$$\frac{\partial}{\partial x} v_g = \frac{\partial}{\partial x}(c\eta) \approx \frac{\partial}{\partial x} c\left(1 - \frac{n_e}{2n_c}\right) = -\frac{c}{2n_c}\frac{\partial n_e}{\partial x}. \tag{2.49}$$

If the electron density rises from 0 to n_0 over a distance l then $\partial_x n_e = n_0/l$ and

[16] [26] R. Wagner *et al. Physical Review Letters.* (1997); [126] A. G. R. Thomas *et al. Physical Review Letters.* (2007).
[17] [127] G.-Z. Sun *et al. Physics of Fluids.* (1987); [128] E. Esarey *et al. Physics of Fluids B: Plasma Physics.* (1993).
[18] [124] W. B. Mori. *IEEE Journal of Quantum Electronics.* (1997).

$$\frac{1}{L}\frac{\partial L}{\partial t} = -\frac{cn_0}{2n_c l} \quad , \quad L = L_0 e^{-cn_0 t/(2n_c l)} \, . \tag{2.50}$$

When self-focussing and pulse compression are tackled simultaneously Ren *et al.*[19] find the pulse length evolution equation

$$\frac{\partial^2 L}{\partial t^2} = \left(\frac{n_e}{n_c}\right)^2 \frac{4c^2}{k_0^2 L^3}\left(1 - \frac{L^2}{w^2}\frac{n_e}{n_c}\frac{P}{P_c}\right) \tag{2.51}$$

for $a_0 \gg 1$ and simulate pulse compression in a low-density plasma ($n_e = 0.09 n_c$). Importantly, pulse compression occurs more quickly as the pulse self-focusses and the two processes reinforce one another. It is also possible to calculate initial pulse compression rates from a photon kinetics model[20], where the spatial gradient in group velocity is driven by frequency shifts of the pulse.

The combination of self-focussing and pulse compression as a laser pulse enters a plasma are important processes for boosting the achievable a_0 beyond its vacuum value, and enabling the onset of wakefield acceleration[21]. As the pulse self-focusses it drives a deeper electron density depression. The pulse then self-compresses more rapidly, boosting a_0 further to drive an even deeper density depression and completing a feedback loop.

2.5 Wakefield generation

The ponderomotive force of an intense, focussed laser pulse is sufficient to expel a significant proportion of the plasma electrons from the focal volume, leaving the almost static ions behind. The large electric fields which result draw the electrons back to the optical axis, where they arrive with large transverse momentum, overshoot, and continue oscillating. A series of models are presented here for describing the electron motion, relevant at progressively higher laser intensities.

2.5.1 Linear 1D wake

Assuming a small-amplitude wakefield where the density perturbation $\delta n_e / n_e \ll 1$, it is possible to describe the 1D linear equation of motion of a cold electron wave using the fluid equations and Gauss' law[22]

$$m_e\left(\frac{\partial \mathbf{v}}{\partial t} + (\mathbf{v}\cdot\nabla)\,\mathbf{v}\right) = -e\left(\mathbf{E} + \mathbf{v}\times\mathbf{B}\right) \tag{2.52}$$

$$\frac{\partial n_e}{\partial t} + \nabla\cdot(n_e\mathbf{v}) = 0 \tag{2.53}$$

$$\nabla\cdot\mathbf{E} = -\frac{e(n_e - n_i)}{\epsilon_0}. \tag{2.54}$$

In the one dimensional approximation the equations of motion do not depend on the transverse co-ordinates, and are therefore translationally invariant. In this case the

[19] [129] C. Ren *et al. Physical Review E.* (2001).
[20] [130] J. Vieira *et al. New Journal of Physics.* (2010).
[21] [31] V. Malka *et al. Science.* (2002).
[22] [131] P. Gibbon. *Short Pulse Laser Interactions with Matter: An Introduction.* (2005).

transverse canonical momentum $\mathbf{p}_\perp - e\mathbf{A}$ is conserved, where $\mathbf{A}$ is the (purely transverse) vector potential of the laser fields. If as $t \to -\infty$ the plasma can be considered at rest and undisturbed, $\mathbf{p}_\perp = e\mathbf{A}\,\forall\, t$ and so in the nonrelativistic limit $\mathbf{v}_\perp = e\mathbf{A}/m_e$. Grouping the nonlinear terms in the momentum equation

$$-m_e\left(\mathbf{v}\cdot\nabla\right)\mathbf{v} - e\left(\mathbf{v}\times\mathbf{B}\right) \approx -\frac{e^2}{m_e}\left(\left(\mathbf{A}\cdot\nabla\right)\mathbf{A} + \left(\mathbf{A}\times\left(\nabla\times\mathbf{A}\right)\right)\right) = -\frac{e^2}{2m_e}\nabla|\mathbf{A}|^2 \quad (2.55)$$

which is the non-relativistic ponderomotive force. To derive the linear wake equation, the terms of Eqs. 2.52, 2.53 and 2.54 are expanded in small quantities about their equilibrium values. Explicitly, we set $n_e = n_0 + \delta n_e$ with $\delta n_e \ll n_0$, and assume that at equilibrium $\mathbf{v}$ and $\mathbf{E}$ are zero, so that in the linear equations they remain small. Linearity is enforced by setting to zero terms containing higher powers of small quantities. Taking second derivatives of Eqs. 2.52 and 2.53 and linearising,

$$m_e\frac{\partial^2 v}{\partial x\partial t} + e\frac{\partial E}{\partial x} = -\frac{1}{2}m_e c^2\frac{\partial^2 a^2}{\partial x^2} \quad (2.56)$$

$$m_e\frac{\partial^2 v}{\partial x\partial t} = -\frac{m_e}{n_0}\frac{\partial^2\delta n_e}{\partial t^2} \quad (2.57)$$

$$e\frac{\partial E}{\partial x} = -\frac{e^2\delta n_e}{\epsilon_0}\,. \quad (2.58)$$

Substituting for the derivatives of E and v in the first equation yields the equation of motion for a small electron density perturbation in a plasma

$$\left(\frac{\partial^2}{\partial t^2} + \omega_p^2\right)\delta n_e = \frac{n_0 c^2}{2}\frac{\partial^2 a^2}{\partial x^2}. \quad (2.59)$$

Eq. 2.59 indicates that the wake is a simple harmonic oscillator of natural frequency ω_p, and is driven by the gradient of the ponderomotive force. In the low-intensity regime the perturbations to density and longitudinal electric field then scale as a_0^2. Use of the ponderomotive force implicitly averages over the short-timescale oscillations of the laser fields, and so it is appropriate in this approximation to assume that the laser and wake vary slowly. In order to state this quasi-static assumption explicitly, it is possible to remove the dominant temporal changes caused by the motion of the laser by changing to co-moving co-ordinates. We make the transformation $\xi = x - ct$, $\tau = t$ such that the laser is approximately static in ξ. The derivatives transform as

$$\frac{\partial}{\partial t} = \frac{\partial\xi}{\partial t}\frac{\partial}{\partial\xi} + \frac{\partial\tau}{\partial t}\frac{\partial}{\partial\tau} = -c\frac{\partial}{\partial\xi} + \frac{\partial}{\partial\tau} \quad (2.60)$$

$$\frac{\partial}{\partial x} = \frac{\partial\xi}{\partial x}\frac{\partial}{\partial\xi} = \frac{\partial}{\partial\xi}. \quad (2.61)$$

The quasi-static approximation is that in the co-moving frame, quantities vary spatially much more rapidly than they do temporally, so $|c\partial_\xi| \gg |\partial_\tau|$ and $\partial_t \to -c\partial_\xi$ (note that

quantities still vary rapidly in laboratory time t through their travelling-wave $x - ct$ dependence). Under this transformation the wave equation for the density perturbation becomes

$$\left(\frac{\partial^2}{\partial\xi^2} + k_p^2\right)\delta n_e = \frac{n_e}{2}\frac{\partial^2 a_0^2}{\partial\xi^2} \tag{2.62}$$

where k_p is the plasma wavenumber. Gauss' law in this frame $\partial_\xi E = -e\delta n_e/\epsilon_0$ implies the electric field is also sinusoidal with wavelength λ_p in the comoving co-ordinate, and trails the laser pulse. Assuming a pulse intensity envelope of the form $\sin^2(\pi\xi/Ł)$, it can be shown[23] that the amplitude of the wake is maximised for a resonant pulse length of $L = \lambda_p$. The electron fluid can be thought to receive one 'kick' forward as the pulse first approaches, and a second kick backwards as the pulse passes. If these two impulses add in phase, the electron fluid motion is maximised. For a top-hat pulse envelope the resonant length is[24] $L = \lambda_p/2$, and for a gaussian envelope $e^{-\xi^2/L^2}$ it is $L = \lambda_p/(\pi\sqrt{2})$. To approximate the electron energy gain in the linear regime, Dawson shows that for a one-dimensional sinusoidal electron perturbation of wavelength λ_e, the maximum longitudinal electric field is[25]

$$E_{\max} = \frac{\omega_p^2 m_e}{e}\frac{\lambda_e}{2\pi}\,. \tag{2.63}$$

To generate electric fields beyond this limit, the electrons would require such a large oscillation amplitude that neighbouring electron sheets cross, the electron density tends towards infinity and the fluid approximation breaks down. For a laser pulse travelling at v_g the wavelength of the density perturbation is $\lambda_e = 2\pi v_g/\omega_p$, so the maximum electric field is

$$E_{\max} = \frac{m_e\omega_p v_g}{e} \equiv \beta_g E_0, \qquad E_0 = \frac{m_e c\omega_p}{e}, \tag{2.64}$$

where E_0 is the maximum attainable accelerating field in a cold non-relativistic plasma wave. E_0 scales with electron density as $n_e^{1/2}$, and for $n_e = 10^{18}\,\mathrm{cm}^{-3}$ is approximately $100\,\mathrm{GVm}^{-1}$. The field is moving at the phase velocity of the plasma wave v_p, which is the group velocity of the laser pulse. The electron energy is most simply calculated in the wave frame, so performing a Lorentz boost the magnitude of the longitudinal electric field stays constant, but the x co-ordinate transforms as

$$x' = \gamma_p(x - v_p t) \Rightarrow k_p x - \omega_p t = \frac{k_p x'}{\gamma_p} \tag{2.65}$$

and so the depth of the potential well of the wake in the boosted frame for $v_p \approx c$ becomes

$$E'(x') = E_0 \sin\left(\frac{k_p x'}{\gamma_p}\right) \Rightarrow \Delta\phi_{\max} = -\int E'\,\mathrm{d}x' = \frac{\gamma_p m_e c^2}{e}. \tag{2.66}$$

The maximum amount of energy the electron may gain in this frame is $W'_{\max} = e\Delta\phi_{\max} =$

[23] [132] P. Sprangle *et al. Applied Physics Letters.* (1988).
[24] [133] L. Gorbunov *et al. Sov Phys JETP.* (1987).
[25] [134] J. M. Dawson. *Physical Review.* (1959).

$\gamma_p m_e c^2$, corresponding to a momentum of $p' = \beta_p \gamma_p m_e c$. Boosting back to the lab frame,

$$\begin{pmatrix} W_{\max} \\ cp \end{pmatrix} = \begin{pmatrix} \gamma_p & \gamma_p\beta_p \\ \gamma_p\beta_p & \gamma_p \end{pmatrix} \begin{pmatrix} W'_{\max} \\ cp' \end{pmatrix} = \begin{pmatrix} \gamma_p^2(1+\beta_p^2)m_e c^2 \\ 2\gamma_p^2\beta_p m_e c^2 \end{pmatrix}.$$

The maximum energy of an accelerated electron in the laboratory frame is therefore

$$W_{\max} = 2\gamma_p^2 m_e c^2 = \frac{2\omega^2 m_e c^2}{\omega_p^2} = \frac{2\omega^2 m_e^2 c^2 \epsilon_0}{n_e e^2} \tag{2.67}$$

Somewhat counter-intuitively electrons are found to be accelerated to higher energies at lower plasma densities. Though the accelerating fields are weaker at lower densities, the higher velocity of the wake more than compensates. It should be noted that if the electron remains inside the wake after reaching $W_{\max}$ it will begin to lose energy. This is an important point to consider when aiming for electron beams of peak energy and small energy spread, and is discussed in section 2.6.2.

2.5.2 Nonlinear wake

It is possible to extend this model to incorporate the fully relativistic 1D wake, as is performed from first principles in the book of Gibbon[26]. An alternative hamiltonian method is presented by Esarey and Piloff[27], which considers the motion of a test electron moving in a 1D wake potential $\phi(\xi)$, only a function of the co-moving co-ordinate. The energy gain of an electron $m_e c^2 \delta\gamma$ in this potential can be related to its displacement by the electric field

$$m_e c^2 \delta\gamma = -eE_x \delta x \quad \Rightarrow \quad \frac{\mathrm{d}\gamma}{\mathrm{d}x} = \frac{e}{m_e c^2}\frac{\partial\phi}{\partial x} \equiv \frac{\partial\Phi}{\partial x} = \frac{\partial\Phi}{\partial\xi}. \tag{2.68}$$

where the potential Φ is normalised to $m_e c^2/e$. From the definition of the co-moving co-ordinate

$$\frac{\mathrm{d}\xi}{\mathrm{d}x} = \frac{\mathrm{d}}{\mathrm{d}x}(x - v_p t) = 1 - v_p\frac{\mathrm{d}t}{\mathrm{d}x} = 1 - \frac{\beta_p}{\beta}. \tag{2.69}$$

By keeping x as an independent variable (analogous to time in standard mechanics), γ may represent a generalised momentum and ξ a generalised position. It is then possible to construct a hamiltonian function $\mathcal{H}(\xi,\gamma)$ which is conserved throughout the electron motion, as long as it satisfies Hamilton's equations

$$\frac{\partial\mathcal{H}}{\partial\xi} = -\frac{\mathrm{d}\gamma}{\mathrm{d}x} \tag{2.70}$$

$$\frac{\partial\mathcal{H}}{\partial\gamma} = \frac{\mathrm{d}\xi}{\mathrm{d}x}. \tag{2.71}$$

The first is simply integrated to $\mathcal{H}(\xi,\gamma) = -\Phi(\xi) + f(\gamma)$ where $f(\gamma)$ is some arbitrary function. The second may also be integrated as

[26] [131] P. Gibbon. *Short Pulse Laser Interactions with Matter: An Introduction.* (2005).
[27] [135] E. Esarey *et al. Physics of Plasmas.* (1995).

$$\begin{aligned}\mathcal{H}(\xi,\gamma) &= \int 1-\beta_p\frac{\gamma}{\sqrt{\gamma^2-1}}\,\mathrm{d}\gamma \\ &= \gamma-\beta_p\sqrt{\gamma^2-1}+g(\xi) \\ &= \gamma(1-\beta_p\beta)-\Phi(\xi). \end{aligned} \tag{2.72}$$

This hamiltonian function is conserved along the electron trajectory, so if an electron begins its motion 'from infinity', $\mathcal{H}=1$ for the rest of its motion. This conservation law was also derived in the speed of light frame by Sprangle *et al*[28]. From the hamiltonian one can directly can calculate $\gamma(\Phi)$ and therefore $v(\Phi)$, from the quasistatic continuity equation in the moving frame the density is calculated as $n_e(\Phi)/n_0 = 1/(v_p - v(\Phi))$, and from the Poisson equation in the moving frame $\partial^2_\xi\Phi = k_p^2(n_e(\Phi)/n_0 - 1)$. The end result, including the contribution of the laser fields to the transverse component of γ, is an implicit differential equation for the plasma potential[29]

$$\frac{\partial^2\Phi}{\partial\xi^2} = \gamma_p^2k_p^2\left(\beta_p\left[1-\frac{1+a^2}{\gamma_p^2(1+\Phi)^2}\right]^{-1/2}-1\right). \tag{2.73}$$

While this expression can be solved analytically under the assumption of certain laser profiles[30], in general it must be solved numerically. Plotted in Fig. 2.6 are the density, electric field and fluid velocity of the wake as a function of wave-frame coordinate for progressively higher laser intensities. The ponderomotive force of the laser pulse initially accelerates the electrons forward, then they are drawn backwards and oscillate longitudinally. As the laser intensity increases the density wave becomes nonlinearly steepened and the longitudinal electric field assumes a sawtooth shape. The maximum electric field in the nonlinear regime becomes

$$E_{\max} = E_0\frac{a_0^2}{\sqrt{1+a_0^2/2}} \tag{2.74}$$

which scales with a_0 at high intensity. The fluid velocity becomes relativistic at large intensities, which leads to a lengthening of the plasma period due to the mass increase of the plasma electrons. In three dimensions the laser pulse driving the wake has a peaked transverse intensity profile which means the wake wavelength is longer on-axis. The wavefronts then assume a characteristic 'horseshoe' shape due to their relativistic motion[31].

Referring to Eq. 2.72, if an electron has some initial momentum before entering the wake potential, its path in the (ξ,γ) phase space is determined by the relation

$$\gamma_0-\beta_p\sqrt{\gamma_0^2-1} = \gamma-\beta_p\sqrt{\gamma^2-1}-\Phi(\xi) \tag{2.75}$$

[28] [136] P. Sprangle *et al. Physical Review Letters.* (1990).
[29] [137] V. I. Berezhiani *et al. Physics Letters A.* (1990); [58] E. Esarey *et al. Reviews of Modern Physics.* (2009).
[30] [138] S. V. Bulanov *et al. JETP Letters.* (1989).
[31] [139] S. V. Bulanov *et al. Physical Review Letters.* (1995).

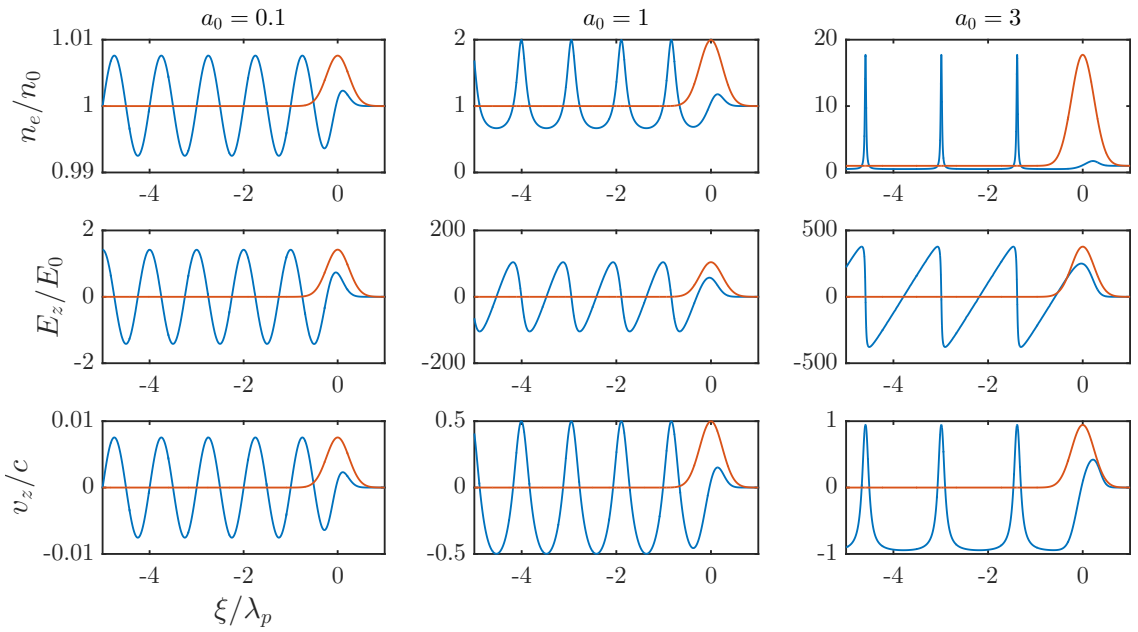


Figure 2.6: Numerically calculated density modulation, electric field and flow velocity (blue) as a function of normalised laser intensity (red).

where γ_0 is its initial Lorentz factor. Solving for Φ numerically, a set of these trajectories are plotted in Fig. 2.7.

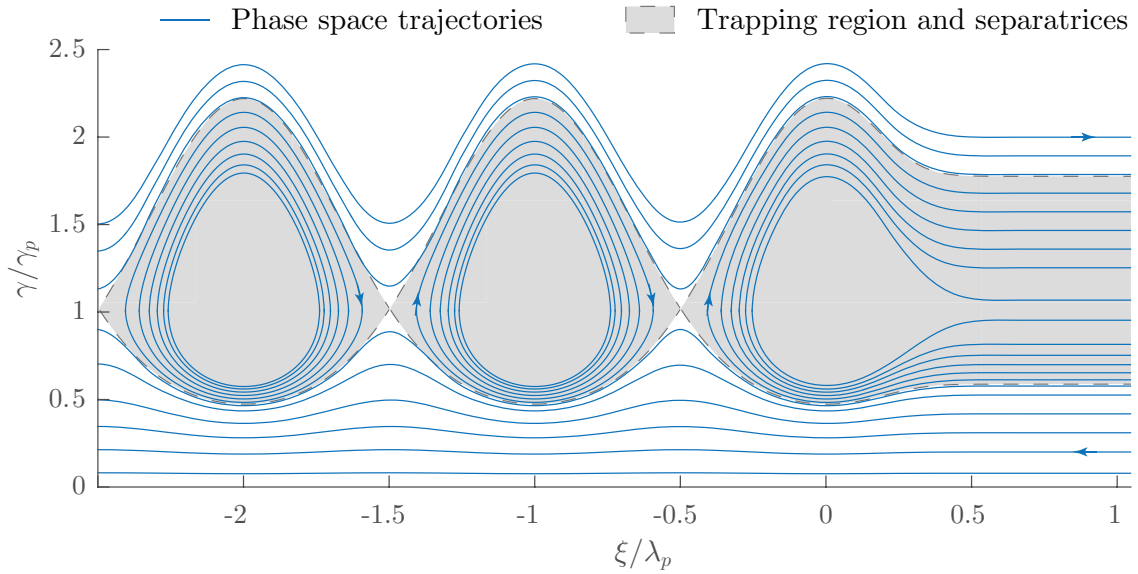


Figure 2.7: Electron trajectories in phase space for various Lorentz factors 'at infinity'. Here $a_0 = 0.1$, $n_e/n_c = 0.001$.

Electrons of low initial velocity relative to v_p are barely disturbed by the passing wake. Similarly, electrons travelling much more quickly than the wake are just accelerated forwards and backwards as the wake passes. However, those electrons with similar Lorentz factor to the wake are trapped by the wake potential and are observed to be initially decelerated then accelerated to significantly higher energies. In order for external electrons to gain energy from the plasma wake they must first be pre-accelerated, or injected,

into the wakefield. In the nonlinear one-dimensional approximation here, the maximum electron energy scales[32] as γ_p^3 rather than γ_p^2. However, as is discussed in the next section, multidimensional effects become important at high laser intensities, and this scaling is not achieved in practice.

The trajectories in this phase space form a hamiltonian flow which is symplectic, or area-preserving. This is a statement of Liouville's theorem that particle distribution functions are conserved as they evolve, or roughly $\Delta\xi\Delta\gamma \approx$ const. As the electrons rotate in phase space, those closer to the separatrices rotate more slowly and it is possible for the distribution function to rotate, exchanging spread in position for spread in energy and vice versa. This fact is important for understanding how laser wakefield accelerators can produce electron beams of narrow energy spread. Bulanov *et al.*[33] calculate the final energy spectrum of an electron bunch injected into a nonlinear wake, and find that the spectrum becomes narrowed during the acceleration process.

2.5.3 Wavebreaking

The maximum accelerating field which can be sustained by a plasma wake determines the energy gain of injected electrons. To calculate this in the nonlinear one-dimensional approximation, note from Eq. 2.73 that when $1 + \Phi \to 1/\gamma_p$ the right hand side diverges. It can also be shown that at this point the electron density diverges, and adjacent electron sheets begin to cross. This marks a breakdown in the fluid approximations thus far assumed, and indicates the onset of wavebreaking where the plasma wave phase-mixes and loses coherence. The longitudinal electric field can be calculated at this point by multiplying Eq. 2.73 by $\partial_\xi\Phi$ and integrating both sides

$$\left(\frac{\partial\Phi}{\partial\xi}\right)^2 = 2\gamma_p^2 k_p^2 \left(\frac{\beta_p}{\gamma_p}\sqrt{\gamma_p^2(1+\Phi)^2 - 1} - \Phi - \beta_p^2\right) . \tag{2.76}$$

Setting $1 + \Phi = 1/\gamma_p$ the electric field at this point is

$$E_{\max} = \frac{\partial\phi}{\partial\xi} = \sqrt{2(\gamma_p - 1)}E_0 \tag{2.77}$$

which is the relativistic cold wavebreaking limit as first derived by Akheizer and Polovin[34]. For $v_p \ll c$ this expression reduces to that derived by Dawson[35] $E_{\max} = \beta_p E_0$. For electrons of a warm waterbag distribution, the wavebreaking limit is derived by Katsouleas and Mori[36] and Bulanov *et al.*[37] and found to be lower in general, though a finite temperature can prevent singularities in the electron density.

In two or more dimensions wavebreaking can occur in distinct ways. In particular it is possible for waves to break transversely[38] due to the relativistic curvature of the wavefronts. After a number of wave periods the wavefront curvature increases until the

[32] [135] E. Esarey *et al. Physics of Plasmas.* (1995).
[33] [140] S. V. Bulanov *et al. Physics of Plasmas.* (2005).
[34] [141] A. I. Akhiezer *et al. Sov. Phys. JETP.* (1956).
[35] [134] J. M. Dawson. *Physical Review.* (1959).
[36] [142] W. B. Mori *et al. Physica Scripta.* (1990).
[37] [143] S. V. Bulanov *et al. Physics of Plasmas.* (2012).
[38] [144] S. V. Bulanov *et al. Physical Review Letters.* (1997).

electron trajectories begin to cross, causing a self-intersection of the wavefronts. At this point wave coherence is lost due to mixing of electron trajectories and the wave breaks. In 2D simulations the wake is observed to break transversely rather than longitudinally, implying the wavebreaking threshold is lower. This lowered threshold causes the maximum electron energy to drop below the formal γ_p^3 scaling derived under the assumption that the wave breaking is longitudinal.

2.5.4 Bubble regime

In 2002 Pukhov *et al.*[39] showed through 3D simulations that for a highly intense laser with $a_0 = 10$ the wake was able to trap background electrons and produce a monoenergetic electron beam with no external stimulus – this was termed the 'highly non-linear broken-wave regime'. The narrow energy spread of the accelerated beam was surprising as hitherto laser accelerated electron beams were observed to have a broad thermal energy spectrum. In this case a single electron-depleted and roughly spherical cavity is formed immediately behind and with approximately the same dimensions as the laser pulse. The electrons displaced from the cavity are seen to form a thin dense sheath around its edge prompting the analogy between the structure and a bubble. Fig. 2.8 shows a frame from a simulation demonstrating bubble formation soon after the laser pulse has entered the plasma. By tracking the trajectories in the wake frame of electrons impinging on the bubble, a proportion are seen to be channelled around the edge of the sheath to the back of the bubble and accelerated – this is the process of self-injection where static background plasma electrons are rapidly accelerated to the velocity of the bubble before further acceleration. The bubble structure is robust with respect to the evolution of the laser pulse and changes in plasma density, in contrast with the catastrophic connotations which might be associated with the term 'wavebreaking'.

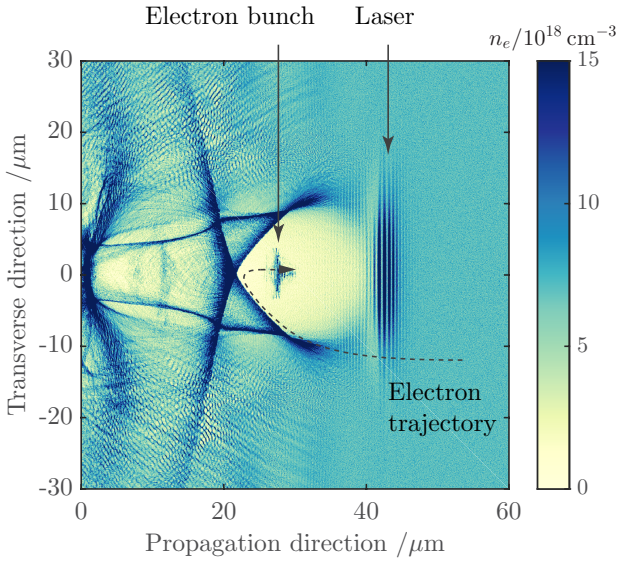


Figure 2.8: Electron density plot from a 2D EPOCH simulation showing the formation of the bubble and injection of electrons.

[39] [145] A. Pukhov *et al. Applied Physics B: Lasers and Optics.* (2002).

Under the assumption that the bubble is a perfect sphere moving relativistically in x at v_p, Kostyukov *et al.*[40] calculate the electromagnetic fields inside the bubble (neglecting terms of order γ_p^{-2})

$$\begin{array}{lll} E_x = \dfrac{n_0 e}{\epsilon_0}\dfrac{\xi}{2} & E_y = \dfrac{n_0 e}{\epsilon_0}\dfrac{y}{4} & E_z = \dfrac{n_0 e}{\epsilon_0}\dfrac{z}{4} \\ B_x = 0 & B_y = n_0 e \mu_0 \dfrac{z}{4} & B_z = -n_0 e \mu_0 \dfrac{y}{4} \\ F_x = \dfrac{m_e \omega_p^2}{4}\left(-2\xi + \beta_y y + \beta_z z\right) & F_y = -m_e \omega_p^2 \dfrac{y}{4}\left(1+\beta_x\right) & F_z = -m_e \omega_p^2 \dfrac{z}{4}\left(1+\beta_x\right) \end{array} \tag{2.78}$$

where $\xi = x - v_p t$, n_0 is the background ion density inside the bubble and β is the velocity of a test electron placed inside the bubble. The wake is electromagnetic in nature, with magnetic forces arising from the displacement current necessary to transport the electric fields forward with the bubble (these fields are different from the case of a relativistically moving ion sphere). The large curvature of the 'wavefront', i.e. electron sheath, causes the radial and longitudinal electromagnetic forces to stay in phase with one another, allowing simultaneous acceleration and focussing of the electron bunch over a region of length $r_b \approx \lambda_p/2$. In a quasi-linear wake this length is only[41] $\lambda_p/4$, increasing as the wake experiences greater relativistic curvature[42]. This is an important difference, and indicates that wakes in the bubble regime are able to trap a greater volume of electrons.

An estimate of the bubble radius r_b arises from the balancing of the outward ponderomotive force of the laser and the inward Lorentz force of the bubble on a plasma ($\beta_x \approx 0$) electron

$$\frac{m_e c^2}{\gamma}\nabla a_0^2 \approx \frac{m_e c^2 a_0}{r_b} \approx \frac{m_e \omega_p^2 r_b}{2} \quad \Rightarrow \quad r_b \approx \lambda_p \sqrt{a_0}/2\pi. \tag{2.79}$$

The bubble size scales slowly with laser intensity, and is near to a plasma wavelength in diameter. The size of the bubble is important in determining whether it can trap electrons, as for a fixed plasma density the magnitude of the Lorentz force at the edge is proportional to r_b. To see how trapping can occur, consider Fig. 2.9a) and the form of the Lorentz forces in Eqs. 2.78. This is a simplified picture of the electron motion, as the ponderomotive force due to the laser pulse is neglected.

As the electron first encounters the bubble, it is accelerated backwards by the longitudinal electric field. If this acceleration is too strong, $\beta_x \to -1$ and the transverse Lorentz forces vanish. The electron never gains significant transverse momentum and streams through the rear of the bubble. If this initial backwards acceleration is weak enough, i.e. the electron encounters the bubble at large enough radius (A), then the transverse Lorentz forces remain strong. In that case the electron rapidly accelerates transversely (B), boosting the forward Lorentz force, which in turn boosts the transverse Lorentz force. By the time the electron reaches the rear of the bubble (C) it has both large transverse and longitudinal

[40] [146] I. Kostyukov *et al. Physics of Plasmas.* (2004).
[41] [147] F. Dorchies *et al. Physics of Plasmas.* (1999).
[42] [148] S. Kalmykov *et al. Physics of Plasmas.* (2006).

momenta and can remain travelling with the bubble, gaining energy from the electric fields as it does so. A sample of numerically-integrated electron trajectories are plotted in Fig. 2.9b) in the co-moving frame, and it is clear that the majority of the transverse acceleration occurs in the small region near the back of the bubble.

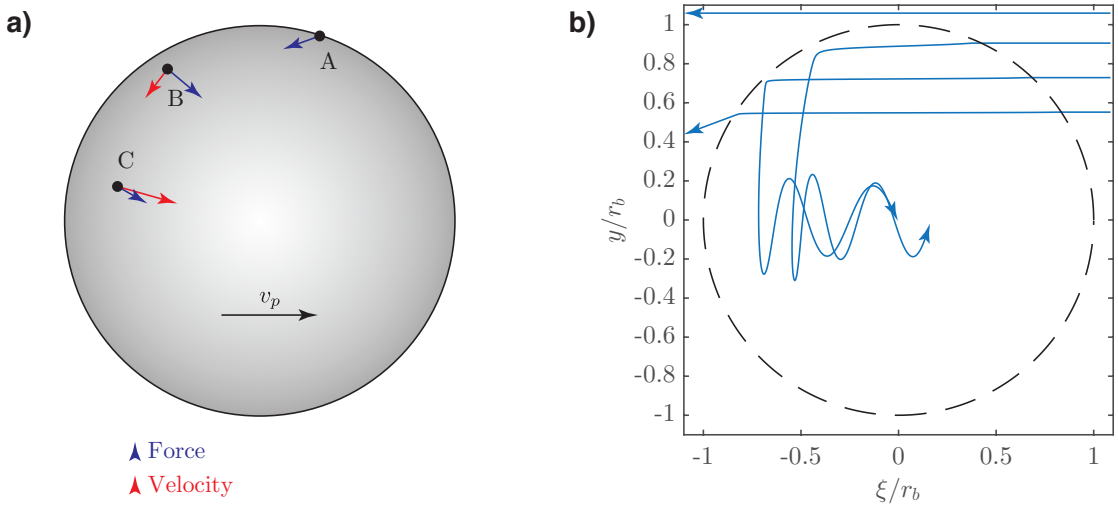


Figure 2.9: The trajectories of electrons arriving at a plasma bubble with different impact parameters.

Lu *et al.*[43] moved beyond the assumption of a static sphere and derived a self-consistent set of equations for the bubble shape and potentials (other models have since included explicit return currents around the bubble[44]). The bubble was found to be approximately spherical, being slightly flatter at the rear where the electron sheaths cross. It was found that optimal guiding of the laser pulse required the pulse transverse size r_m to be matched to the bubble size, with an empirical expression for r_m from 3D simulations found to be $r_m \approx \lambda_p\sqrt{a_0}/\pi$. This is larger than the simple estimate of the bubble size given above by approximately a factor of two. If the laser pulse is focussed too tightly then the ponderomotive force is increased and the local electron density in the focal volume drops. The relativistic guiding of the pulse is weaker and the natural diffraction stronger, so the pulse expands. If the pulse is larger than r_m then the opposite is true, and the pulse self-focusses until it reaches r_m. In both cases the radial envelopes of unmatched spots were observed to oscillate in size, with oscillations minimised when the spot is matched. The matching condition may also be reformulated as $a_0 \approx 2(P/P_c)^{1/3}$, where P is the laser power and P_c the critical power for self-focussing. If the condition for entering the bubble regime is taken to be $r_b > 2\lambda_p/\pi$ and $a_0 > 4$, then the required laser power in the matched case is $P/P_c > 8$. Accounting for the scaling of the bubble size with laser intensity, the maximum electron energy gain in the bubble regime is $W_{\max} = (2/3)m_e c^2(n_c/n_e)a_0$, which with a matched spot scales like $(P/P_c)^{1/3}$.

Experiments first showed in 2004 that by operating in the bubble regime it was possible to produce monoenergetic electron beams[45], a key prediction of the simulations and theory.

[43] [149] W. Lu *et al. Physics of Plasmas.* (2006); [150] W. Lu *et al. Physical Review ST-AB.* (2007).
[44] [151] S. A. Yi *et al. Physics of Plasmas.* (2013).
[45] [32] S. P. D. Mangles *et al. Nature.* (2004); [34] J. Faure *et al. Nature.* (2004); [33] C. Geddes *et al.*

Evidence for electrons being accelerated in the first wave period comes from experimental measurements of electron beam profiles,[46] or direct transverse imaging[47]. Simulations at much lower laser intensities also indicate that the theoretical scalings continue down to laser energies at the millijoule level[48].

Self-injection threshold

For a given laser, a lower limit on P/P_c is the same as a lower limit on plasma electron density, and so to operate in the bubble regime and accelerate self-injected monoenergetic electron beams the plasma density should be sufficiently high. An empirical condition $P/P_c > 3$ for self-injection was found to hold over a limited range of experimental parameters[49]. A number of theoretical models of self-injection in the bubble regime reformulate this threshold in terms of a minimum bubble size required for injection – if the bubble is too small the electrons don't reach sufficient energies in their transit time and aren't injected.

Kostyukov *et al.*[50] calculate a minimum bubble-radius criterion as $r_b > \sqrt{2}\gamma_b\lambda_p/2\pi$ where γ_b is the Lorentz factor associated with the bubble. Assuming a matched spot, this model predicts laser intensity thresholds which appear too large. By assuming that the injected electrons closely follow the shape of the bubble sheath, Thomas[51] calculates the bubble radius threshold

$$r_b > \frac{\lambda_p}{\pi}\sqrt{\ln\left(2\gamma_b^2\right) - 1}$$

which was reformulated by Mangles *et al.*[52] as an implicit inequality on n_e taking into account the pulse evolution prior to injection

$$\frac{\alpha P}{P_c} > \frac{1}{16}\left(\ln\left(\frac{2n_c}{3n_e}\right) - 1\right)^3 \tag{2.80}$$

where α is the fraction of the laser power within the FWHM spot size. The authors also compare the prediction of Eq. 2.80 and several experimentally-measured thresholds where agreement is found over a density range of 3 – 20 $\times 10^{18}$ cm^{-3}. The simpler condition $P/P_c > 3$ significantly overestimates the threshold above 15 $\times 10^{18}$ cm^{-3}.

In a model by Kalmykov[53] injection instead occurs during a period of bubble growth driven by evolution of the laser pulse, modelled by an adiabatically-varying hamiltonian. If the bubble then begins to contract, as is observed in simulations, injection is abruptly stopped and the bunch is monochromatised by the phase space rotation.

Fig. 2.10 displays the results of a series of simulations performed at increasing densities

Nature. (2004).

[46] [41] S. P. D. Mangles *et al. Physical Review Letters.* (2006).

[47] [152] A. Sävert *et al. Physical Review Letters.* (2015).

[48] [153] M. Geissler *et al. New Journal of Physics.* (2006).

[49] [154] D. H. Froula *et al. Physical Review Letters.* (2009).

[50] [155] I. Kostyukov *et al. Physical Review Letters.* (2009); [156] I. Kostyukov *et al. New Journal of Physics.* (2010).

[51] [157] A. G. R. Thomas. *Physics of Plasmas.* (2010).

[52] [158] S. P. D. Mangles *et al. Physical Review ST-AB.* (2011).

[53] [159] S. Kalmykov *et al. Physical Review Letters.* (2009).

with all other parameters being held constant. In the final frame the inset shows the appearance of a small injected bunch of electrons. For this combination of laser parameters, Eq. 2.80 predicts the self-injection threshold to occur at 4.2 $\times 10^{19}$ cm^{-3} which is consistent with the simulations. The $P/P_c > 3$ model overestimates the threshold as 7.8 $\times 10^{19}$ cm^{-3}.

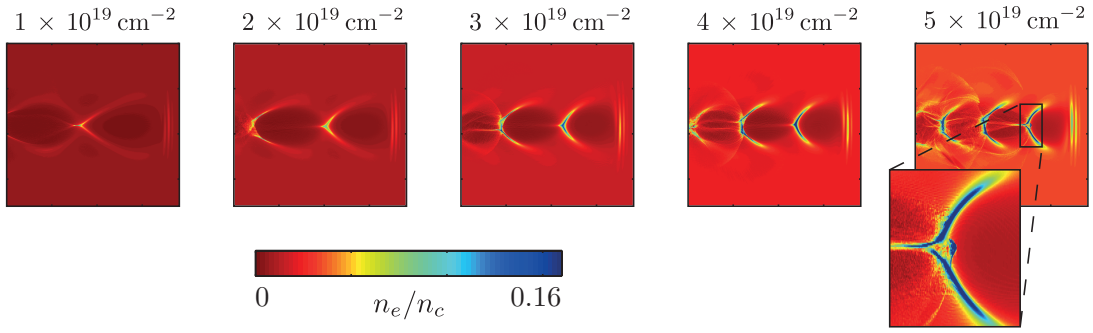


Figure 2.10: Simulation scan of electron density showing the onset of self-injection.

2.6 Limits on acceleration

2.6.1 Depletion Length

As the laser pulse propagates through the plasma it loses energy to the production of plasma waves and high-energy electrons. There will be a propagation distance over which the pulse will lose the majority of its energy – this is known as the depletion length L_{dep}. A simple order-of-magnitude estimate for L_{dep} can be derived by considering the energy density deposited into the wakefield. Solving the linear wake equation 2.62 for $a_0 \ll 1$ shows that the peak electric field of the plasma wave varies as[54] $E_p = a_0^2 E_0$ for a pulse length comparable to λ_p. The laser field E_L is related to E_0 by $E_L = a_0 \omega E_0/\omega_p$, and it is reasonable to assume that the laser spot and transverse dimensions of the wakefield are close to λ_p. Assuming depletion occurs when the laser has created a column of plasma containing the same electromagnetic energy as the pulse and neglecting any magnetic fields, the total energies of the pulse and plasma may be equated

$$L_{\mathrm{dep}} \langle E_p^2 \rangle = \lambda_p \langle E_L^2 \rangle \ , \ L_{\mathrm{dep}} = \lambda_p \frac{\langle E_L^2 \rangle}{\langle E_p^2 \rangle} = \lambda_p \frac{\omega^2}{\omega_p^2} \frac{1}{a_0^2} = \lambda_p \frac{n_c}{n_e} \frac{1}{a_0^2} \, . \tag{2.81}$$

As here $a_0 \ll 1$ it is appropriate to notice that $a_0 \to 0$ sends $L_{\mathrm{dep}} \to \infty$ and the transparency of underdense plasmas is recovered. As a_0 approaches and exceeds 1 nonlinear effects will arise and alter the functional form of this result. Shadwick *et al.*[55] find for a laser pulse with a gaussian envelope

$$L_{\mathrm{dep}} = \lambda_p \frac{n_c}{n_e} \begin{cases} 2.8/a_0^2 & a_0 \ll 1 \\ 1.4 & a_0 \gg 1 \end{cases} \tag{2.82}$$

where the depletion length is understood to be the $1/e$ decay length of the pulse energy. Similar expressions are found for a flat-top pulse, with the dimensionless numerical

[54] [133] L. Gorbunov *et al. Sov Phys JETP.* (1987).
[55] [160] B. A. Shadwick *et al. Physics of Plasmas.* (2009).

constants altered by factors of order unity. Decker *et al.*[56] calculate an etching velocity $v_{\text{etch}}/c = n_e/n_c$ where the leading edge of the pulse is eroded as it drives a wake. Taking the length of the pulse to be λ_p, the entire pulse is etched away after a distance $v_g\lambda_p/v_{\text{etch}} \approx \lambda_p n_c/n_e$, scaling in the same way as the above. This model also predicts that the velocity of the leading edge of the pulse is slower than would be expected from the linear plasma refractive index, and for underdense plasmas is actually

$$\frac{v_g}{c} = \sqrt{1 - \frac{n_e}{n_c}} - \frac{n_e}{n_c} \approx 1 - \frac{3}{2}\frac{n_e}{n_c}\,. \tag{2.83}$$

Fig. 2.11 shows the evolution of a laser pulse from a 2D EPOCH simulation in a frame moving at this modified group velocity. It is clear that the pulse initially travels at the faster linear group velocity, but by the time it is fully self-compressed it travels at the reduced depletion velocity.

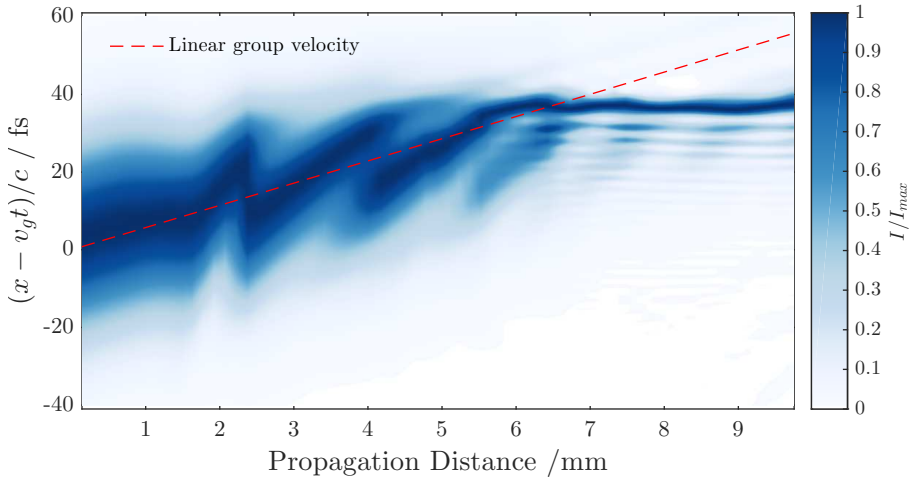


Figure 2.11: Each vertical slice is a lineout of laser intensity on-axis during a simulation frame. Each slice is normalised to its own peak intensity in order to illustrate the pulse evolution more clearly. Here Eq. 2.83 is used to calculate v_g.

2.6.2 Dephasing length

The low mass of electrons means that during acceleration they rapidly approach the speed of light. Acceleration will only continue so far before the electron bunch outruns the accelerating phase of the bubble and begins to be decelerated, as was discussed in section 2.5 in the Lorentz-boosted frame of the wake. The length over which this happens in known as the dephasing length, as the electrons 'dephase' from the wakefield. If a plasma is significantly longer than its dephasing length then the electron bunch begins to be decelerated chaotically and interact with the laser pulse, leading to a broader energy distribution. In the linear regime dephasing occurs once the electron bunch has moved $\lambda_p/2$ forward relative to the wake. If the electron bunch can be said to be travelling at approximately c in the lab frame, this distance is $\lambda_p c/2(c - v_p)$. Given that for $n_e/n_c \ll 1$, $\gamma_p \gg 1$ and so

[56] [161] C. D. Decker *et al. Physics of Plasmas.* (1996).

$$v_p = c\sqrt{1 - \frac{1}{\gamma_p^2}} \approx c\left(1 - \frac{1}{2\gamma_p^2}\right) \quad \Rightarrow \quad c - v_p \approx \frac{c}{2\gamma_p^2}\,. \tag{2.84}$$

The dephasing length L_d is then

$$L_d = \lambda_p \gamma_p^2 = \lambda_p \frac{n_c}{n_e} = \lambda \left(\frac{n_c}{n_e}\right)^{3/2} \tag{2.85}$$

where λ is the laser wavelength. In the matched bubble regime, the dephasing distance should be replaced by $r_b = \lambda_p\sqrt{a_0}/\pi$. The effective bubble velocity is also reduced by the etching speed, and so in this regime the dephasing length is

$$L_d = \lambda_p \frac{n_c}{n_e} \times \frac{2\sqrt{a_0}}{3\pi}. \tag{2.86}$$

If the pulse length is also r_b, then the dephasing length is shorter than the depletion length and in the 3D nonlinear regime the acceleration is dephasing limited.

2.7 Electron injection schemes

By Liouville's theorem the injection of electrons into a wake must be localised in space and time if the beam is to have small energy spread. Although self-injection allows this to occur naturally, relying on the self-focussing and self-compression of the pulse provides little control over the injection process, and so many experiments have been performed to examine alternative electron injection techniques.

If an external electron bunch taken from a conventional accelerator were to be injected, the bunch length must be much shorter than the bubble size if the whole of the bunch is to feel the same accelerating fields. Given the specifications of current generation laser systems, the plasma density range over which the bubble regime may be reached limits the plasma period to ≤ 100 fs. Producing high charge electron beams of this duration is challenging, though methods have been proposed which use the plasma wave to compress externally-injected bunches[57], but such injection is yet to be implemented.

2.7.1 Ionisation injection

A simple method to inject electrons using gas targets is ionisation injection, where the gas-jet target (typically H or He) contains a small fraction of higher-Z impurity. If the impurity contains electrons with a very high ionisation threshold those electrons are not ionised until the laser pulse arrives, with the H or He ionising long before and providing a background species of plasma electrons. Taking nitrogen as an example, the first 5 outer electrons have slowly increasing ionisation enthalpies of 1,400–9,500 kJmol^{-1}. For the next electron in a separate closed subshell the ionisation enthalpy increases to 53,000 kJmol^{-1}, requiring $a_0 > 1$ for ionisation.

Having originated from static ions, the newly-ionised electrons are deposited behind the laser pulse but are not moving with the bulk motion of the plasma. In the 1D nonlinear wake approximation, the condition that the ionised electron is trapped is[58]

[57] [162] A. G. Khachatryan *et al. Physical Review ST-AB.* (2004).
[58] [163] M. Chen *et al. Physics of Plasmas.* (2012).

$$a_0 \gtrsim 1.2\sqrt{1-\gamma_p^{-1}} \tag{2.87}$$

indicating that regardless of the ionising species relativistic laser intensities are required. Ionisation injection has been shown to improve the charge and divergence of laser-produced electron beams[59] and lower the laser intensity threshold at which electron beams are observed[60]. The intense electric fields produced by a focused electron beam have also been shown to be sufficient to demonstrate ionisation injection from He in a Li plasma, with injected electrons being accelerated up to 7.6 GeV[61].

The lower threshold of ionisation injection can increase the injection duration and therefore the energy spread of accelerated beams. Several methods are proposed to limit the region of injection, including confining the dopant to a small region[62] or triggering ionisation with a second laser pulse[63].

2.7.2 Colliding pulse injection

Rather than allowing the bubble to grow until background electrons are accelerated to γ_p, it has been proposed[64] that a second counterpropagating laser pulse could be used to produce a standing beatwave, the ponderomotive force of which rapidly accelerates background electrons to γ_p. Experiments have shown that the injection pulse is able to improve beam stability and energy spread, as well as allowing some degree of tunability by altering the crossing point of the two pulses[65].

2.7.3 Density modification

Sudden changes in plasma density have been observed analytically[66] and in simulation[67] to cause the localised injection of electrons. If the plasma density drops suddenly, the bubble elongates significantly in a short time and the effective velocity of the back of the bubble drops momentarily. The trapping threshold is relaxed briefly which allows a burst of injection. Experiments using multiple gas targets have shown that such beams have superior stability and tunability than self-injected beams and can allow the processes of injection and acceleration to be separated[68]. Alternatively, transient plasma density modifications are produced with secondary laser pulses[69] or solid obstructions are placed in the gas flow[70].

[59] [48] C. McGuffey *et al. Physical Review Letters.* (2010).
[60] [164] A. Pak *et al. Physical Review Letters.* (2010).
[61] [165] E Oz *et al. Physical Review Letters.* (2007).
[62] [46] B. B. Pollock *et al. Physical Review Letters.* (2011).
[63] [166] N. Bourgeois *et al. Physical Review Letters.* (2013); [167] L. L. Yu *et al. Physical Review Letters.* (2013).
[64] [168] E. Esarey *et al. Physical Review Letters.* (1997); [169] H. Kotaki *et al. Physics of Plasmas.* (2004); [170] G. Fubiani *et al. Physical Review E.* (2004).
[65] [40] J. Faure *et al. Nature.* (2006); [171] C. Rechatin *et al. Physical Review Letters.* (2009).
[66] [172] G. Fubiani *et al. Physical Review E.* (2006).
[67] [173] H. Suk *et al. Physical Review Letters.* (2001); [174] A. V. Brantov *et al. Physics of Plasmas.* (2008).
[68] [38] H. T. Kim *et al. Physical Review Letters.* (2013); [175] M. Hansson *et al. Physical Review ST-AB.* (2015).
[69] [176] T. Y. Chien *et al. Physical Review Letters.* (2005); [177] J. Faure *et al. Physics of Plasmas.* (2010).
[70] [178] K. Schmid *et al. Physical Review ST-AB.* (2010); [179] M. Burza *et al. Physical Review ST-AB.* (2013).

A density decrease over a long scale length has also been proposed as an injection mechanism[71], where here 'long' is taken to mean many plasma wavelengths. As the density falls the phase velocity of the wake decreases and eventually a portion of the wave breaks, injecting a small population of electrons. Experiments incorporating longer scale lengths have used inhomogenous gas jets for both the injection and acceleration of electrons[72], or separated the two stages[73]. Both mechanisms report much-improved stability and tunability of electron beam peak energies and energy spread.

2.8 Betatron radiation

In this section we move on from the physics of wake generation and electron acceleration and discuss the mechanism of betatron radiation. Laser-driven synchrotron sources have generated much interest during the previous decade, and are seen by many to represent a potential application for LWFA electron beams such as that presented in chapter 5.
The focussing fields of the bubble cavity, combined with the intrinsic high transverse momentum of self-injected electrons, means that LWFA electron beams oscillate transversely as they are accelerated. Betatron radiation was first observed at SLAC using wakes driven by much higher energy electron beams[74], and then in 2004 in a LWFA experiment[75].

2.8.1 Radiation parameters

The spectrum of radiation emitted by an electron as it moves through a path $\mathbf{r}(t)$ is given by[76]

$$\frac{\mathrm{d}^2 I}{\mathrm{d}\omega \mathrm{d}\Omega} = \frac{e^2}{16\pi^3 \epsilon_0 c} \left| \int_{-\infty}^{\infty} e^{i\omega(t-\mathbf{n}\cdot\mathbf{r}/c)} \frac{\mathbf{n} \times \left((\mathbf{n} - \beta) \times \dot{\beta} \right)}{(1 - \beta \cdot \mathbf{n})^2} \, \mathrm{d}t \right|^2 \tag{2.88}$$

where β is the electron velocity and $\mathbf{n}$ is the (assumed constant) normalised observation direction. The term in the denominator implies that radiation is maximised for $\beta \cdot \mathbf{n} \approx 1$, which occurs for relativistic electrons viewed from an angle within a narrow cone along the electron propagation axis. If the maximum angle between the electron trajectory and average propagation direction is ψ and the beam energy is γ, then the undulator parameter $K = \psi\gamma$ determines the character of the radiation. For $K \ll 1$ the spatial periodicity of the radiation is

$$\lambda = \frac{\lambda_u}{2\gamma^2} \left(1 + \frac{K^2}{2} + \gamma^2 \theta^2 \right) \tag{2.89}$$

where θ is the viewing angle relative to the propagation direction and λ_u is the spatial period of the electron trajectory. This is the undulator regime where the maximum electron angular excursion is smaller than the radiation divergence, and an observer on axis witnesses a beam of near-constant intensity of angular size $1/\gamma$. The radiation is

[71] [180] S. Bulanov *et al. Physical Review E.* (1998).
[72] [181] C. Geddes *et al. Physical Review Letters.* (2008).
[73] [54] A. J. Gonsalves *et al. Nature Physics.* (2011).
[74] [80] S. Wang *et al. Physical Review Letters.* (2002).
[75] [81] A. Rousse *et al. Physical Review Letters.* (2004).
[76] [182] J. D. Jackson. *Classical Electrodynamics.* (1998).

nearly monochromatic at the wavelength given by Eq. 2.89 with harmonics appearing for larger K. The regime $K \gg 1$ is known as the wiggler regime where an observer on axis witnesses temporally distinct bursts of radiation of duration $\tau \propto \lambda_u/(\gamma^2 K)$ with a broad synchrotron spectrum. The divergence angle in this case is determined by the maximum angular excursion $\Delta\theta = \psi = K/\gamma$.

To relate these quantities to parameters relevant to LWFA, consider the transverse Lorentz force on an electron of constant γ in the bubble regime

$$F_\perp = -\frac{m_e \omega_p^2 r_\perp}{2} \quad \Rightarrow \quad \ddot{r}_\perp + \frac{\omega_p^2}{2\gamma} r_\perp = 0\,. \tag{2.90}$$

In the laboratory frame the electron oscillates transversely at the betatron frequency $\omega_\beta = \omega_p/\sqrt{2\gamma}$, or equivalently the betatron wavelength $\lambda_\beta = \sqrt{2\gamma}\lambda_p$. For a plasma of $n_e = 10^{18}$ cm^{-3} supporting a GeV electron beam of $\gamma \approx 2000$, $\lambda_\beta \approx 2$ mm. This effective undulator wavelength is short compared to conventional magnetic undulators, indicating that high energy radiation in the hard x-ray spectrum may be produced.

From the assumption of sinusoidal motion, the maximum transverse displacement of the electron is related to the undulator parameter K by

$$r_\beta = \frac{K}{\gamma k_\beta} = \sqrt{\frac{2}{\gamma}}\frac{K}{k_p} \quad \Rightarrow \quad K = \sqrt{\frac{\gamma}{2}} r_\beta k_p = \gamma r_\beta k_\beta. \tag{2.91}$$

Assuming a matched bubble size, the amplitude of the transverse motion as a function of electron γ is derived by Thomas as[77]

$$k_p r_\beta = \left(\frac{2a_0^3}{\gamma\gamma_p}\right)^{1/4} \tag{2.92}$$

where using the same plasma parameters as above, $r_\beta \approx 1\,\mu$m and $K \approx 6$. Each individual electron radiates in the wiggler regime and produces a broad synchrotron-like spectrum, and the total emitted spectrum will be a sum of these spectra over the electron energy distribution. The fact that the oscillation radius is small means that the effective radiation source size is also small. The integrated spectrum radiated by a single electron can be shown to be in this regime[78]

$$\frac{\mathrm{d}^2 I}{\mathrm{d}E\mathrm{d}\Omega} \propto \frac{\gamma^2\xi^2}{1+\gamma^2\theta^2}\left(\frac{\gamma^2\theta^2}{1+\gamma^2\theta^2}\mathcal{K}_{1/3}^2(\xi) + \mathcal{K}_{2/3}^2(\xi)\right) \tag{2.93}$$

where $\xi = \frac{E}{2E_{\mathrm{crit}}}(1+\gamma^2\theta^2)^{3/2}$, $\mathcal{K}$ is a modified Bessel function of the second kind, and the critical energy E_{crit} of the spectrum is

$$E_{\mathrm{crit}} = \frac{3}{2} K \gamma^2 k_\beta \hbar c = \frac{3}{4}\gamma^2 r_\beta \omega_p^2 \hbar/c. \tag{2.94}$$

Note that the definition of E_{crit} varies by a factor of 2 in some publications, so care must be taken when comparing results. For betatron beams produced at high-power laser facilities, $E_{\mathrm{crit}} > 10\,$keV and the radiation spectrum is in the hard x-ray range. Eq. 2.93 is

[77] [157] A. G. R. Thomas. *Physics of Plasmas.* (2010).
[78] [183] E. Esarey *et al. Physical Review E.* (2002).

plotted in Fig. 2.12, along with the marginals over E and θ. The radiated spectrum is hardest on axis in this case, with the peak in the spectrum lying close to E_{crit}. Electron bunches with non-planar trajectories can instead produce x-ray beams with an intensity minimum on-axis[79]. The energy off-axis is observed to decay as $1/\theta$, though measurements made of real betatron beams indicate that the angular dependence of energy is much less strong, consistent with the presence of multiple electron populations oscillating in different planes[80].

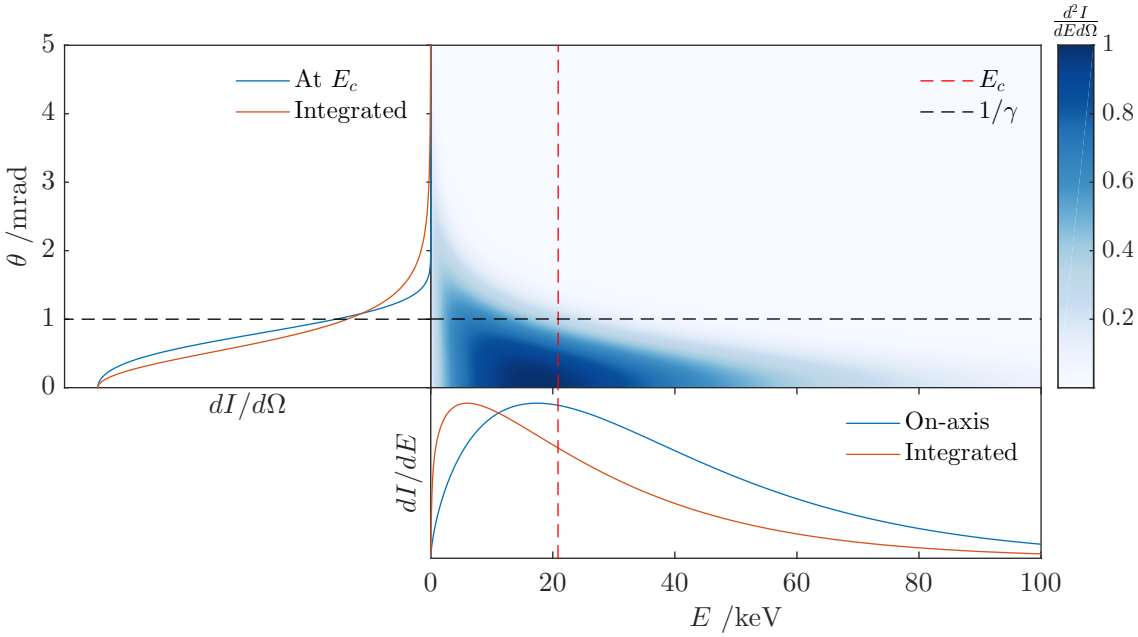


Figure 2.12: Betatron spectrum calculated according to Eq. 2.93 for $n_e = 2 \times 10^{18}\,\text{cm}^{-3}$, $\gamma = 1000$, $r_\beta = 2\,\mu\text{m}$.

The betatron oscillation radius, and therefore x-ray source size, has been measured directly through the examination of diffraction patterns from an opaque edge[81], and indirectly from the far-field divergence pattern of the x-ray beam[82]. The source sizes are measured to be in the micron range and so assuming the beam is produced by a bunch of incoherently emitting electrons, after propagating a distance L the x-ray beam has a transverse spatial coherence length

$$l_{\text{coh}} = \frac{\lambda L}{s} = \frac{hcL}{sE} \qquad l_{\text{coh}}[\text{mm}] = 1.2\frac{L[\text{m}]}{s[\mu\text{m}]}E[\text{keV}] \tag{2.95}$$

where s is the source size, λ is the radiation wavelength and E is the photon energy. Even for beams of very short wavelength, the transverse coherence length can quickly become macroscopically large for metre-scale propagation distances. Spatial coherence allows diffraction features to form and therefore betatron beams are suitable for in-line x-ray

[79] [85] K. Ta Phuoc *et al. Physical Review Letters.* (2006).
[80] [90] F. Albert *et al. Physical Review Letters.* (2013).
[81] [184] M. Schnell *et al. Physical Review Letters.* (2012); [43] S. Kneip *et al. Physical Review ST-AB.* (2012).
[82] [86] F. Albert *et al. Physical Review E.* (2008).

holography[83], also known as propagation-based phase-contrast imaging. Phase-contrast images of insects[84] have been recorded with betatron x-ray beams at critical energies extending to several tens of keV[85].

2.9 Parametric instabilities

When an intense laser pulse propagates through a plasma it can excite a wide variety of wave modes, which are able to interact with the laser pulse. If the interaction is unstable the mode can grow to macroscopic size, where it may produce a measurable effect. Examples of such parametric instabilities in plasmas are Brillouin and Raman scattering, where an ion wave or electron wave respectively grow to such an extent that it can scatter significant portions of the laser light. Raman scattering in particular was observed to be an important process during the experiments presented in chapters 4 and 5, and so the general theory is presented in this section. Crossing laser pulses may also couple into one another through the Raman process, discussed in section 2.9.3 in preparation for the results of chapter 5.

2.9.1 Coupling of oscillators

A parametric interaction between oscillators occurs when the motion of one is allowed to affect the motion of another. A parametric instability occurs when oscillators are parametrically coupled through a strong pump oscillator, the frequency and amplitude of which is assumed constant. The term 'parametric' arises from the notion that the pump can be thought of as a periodic modulation of the oscillator parameters. The coupled oscillators can gain energy from the pump, grow in amplitude, and couple more strongly, constituting an exponential feedback loop. Considering the general case, take two oscillators a and b interacting in the presence of pump E. The equations of motion are

$$\frac{\mathrm{d}^2 a}{\mathrm{d}t^2} + \omega_a^2 a = c_a b E \tag{2.96}$$

$$\frac{\mathrm{d}^2 b}{\mathrm{d}t^2} + \omega_b^2 b = c_b a E \tag{2.97}$$

where the coupling coefficients c_a, c_b determine the strength of the interaction, and the oscillators have natural frequencies ω_a, ω_b. Suppose E oscillates at a fixed frequency ω_0, and consider the amplitude of a at some frequency ω. The response of b at some second frequency ω' is found by substituting $b = b_0 \cos \omega' t$ in Eq. 2.97, with analogous substitutions for E and a.

$$\begin{aligned} \left(\omega_b^2 - \omega'^2\right) b_0 \cos \omega' t &= c_b E_0 a_0 \cos \omega_0 \cos \omega t \\ &= \frac{c_b E_0 a_0}{2} \left(\cos((\omega_0 - \omega)t) + \cos((\omega_0 + \omega)t)\right) . \end{aligned} \tag{2.98}$$

[83] [185] A. Krol. *Holography, Research and Technologies.* (2011).
[84] [104] S. Fourmaux *et al. Optics Letters.* (2011); [105] S. Kneip *et al. Applied Physics Letters.* (2011).
[85] [106] Z. Najmudin *et al. Philosophical Transactions of the Royal Society.* (2014).

The nonlinear interaction of the pump and a drives oscillations at $\omega_0 \pm \omega$. The response of b is largest if its resonant frequency lies near these drive frequencies, though with finite damping the resonance peak is broadened and b will be excited off-resonance too.
Assuming oscillator b has natural frequency ω_b near $\omega_0 \pm \omega$ and so has significant amplitudes $b(\omega_0 \pm \omega)$ there, through the same process it can be shown that a is driven at frequencies ω and $2\omega_0 \pm \omega$. If we assume that ω_a is near ω and $\omega_a \ll \omega_0$, the high-frequency $2\omega_0$ term is driving a far from resonance and can be neglected. Combining these observations we can relate the driven amplitudes of $a(\omega)$ and $b(\omega_0 \pm \omega)$ as follows

$$\begin{pmatrix} \omega_a^2 - \omega^2 & -c_a E_0 & -c_a E_0 \\ -c_b E_0 & \omega_b^2 - (\omega_0 - \omega)^2 & 0 \\ -c_b E_0 & 0 & \omega_b^2 - (\omega_0 + \omega)^2 \end{pmatrix} \begin{pmatrix} a(\omega) \\ b(\omega_0 - \omega) \\ b(\omega_0 + \omega) \end{pmatrix} = 0\,. \tag{2.99}$$

This system of equations has nontrivial solutions only if the determinant of the matrix vanishes, the condition for which defines the dispersion relation for oscillations of a at frequency ω. Let $D_a = \omega_a^2 - \omega^2$, $D_\pm = \omega_b^2 - (\omega_0 \pm \omega)^2$ represent the dispersion relations of the uncoupled oscillators such that $D \to 0$ corresponds to a resonance. The coupled dispersion relation takes the form

$$\frac{c_a c_b E_0^2}{D_a}\left(\frac{1}{D_+} + \frac{1}{D_-}\right) = 1\,. \tag{2.100}$$

If Eq. 2.100 is solved for ω and found to contain an imaginary component, then both a and b grow exponentially at the expense of E. Rewritten in the form $D_a D_+ D_- = c_a c_b E_0^2 (D_+ + D_-)$, it is clear that when the couplings $c_{a/b}$ or pump strength E_0 vanish, the oscillators become independent and the overall dispersion relation collapses to a product of the individual relations for each oscillator.

2.9.2 Raman scattering

Parametric instabilities occur in plasmas[86] when particle density perturbations scatter laser light. The scattered light can beat with the incident 'pump' light, reinforcing the density perturbation and completing a feedback loop. Stimulated Raman scattering[87] is caused by the rapid growth of parametrically unstable electron waves in a plasma, which grow exponentially and can scatter a significant fraction of incoming laser light. The natural electron oscillation frequency is ω_p, so the response to any perturbation will be to generate currents at this frequency. Any terms in the equations governing the laser-plasma interaction coupling the electron perturbation to the incident laser fields will act to generate currents at $\omega_0 \pm \omega_p$, where ω_0 is the laser frequency. These currents in turn generate electromagnetic radiation at $\omega_0 \pm \omega_p$. As the instability progresses, the pump is effectively shifted in frequency by ω_p and scattered into a different direction.
More explicitly, it is possible to perform an instability analysis of the Raman scattering process where a uniform background is assumed and small perturbations in the electron density and electromagnetic wave are allowed to grow. Here we follow the relativistic

[86] [186] J. F. Drake. *Physics of Fluids.* (1974).
[87] [187] J. F. Drake *et al. Physical Review Letters.* (1973); [188] M. N. Rosenbluth *et al. Physical Review Letters.* (1973); [189] D. W. Forslund *et al. Physics of Fluids.* (1975).

derivation by Sakharov *et al.*[88] which is expanded upon by Gibbon[89]. As a first approximation the incident laser and background plasma are assumed to be infinite and spatially homogeneous. Take the laser vector potential and plasma electron density to be

$$\mathbf{A} = \mathbf{A}_0 + \tilde{\mathbf{A}}, \qquad n_e = n_0 + \tilde{n} \tag{2.101}$$

where $\mathbf{A}_0$ and n_0 are the background laser vector potential and electron density respectively – both are taken to be constants. The small perturbations $\tilde{\mathbf{A}}$ and $\tilde{n}$ are allowed to vary temporally and spatially, and it is the exponential growth rate of these quantities which defines the growth of the Raman instability. The plasma ions are taken to be a fixed uniform background over the typical timescales of growth. From the translational symmetry of the assumed solution, transverse canonical momentum is conserved so $p_\perp = e\mathbf{A}$ and the electron Lorentz factor becomes

$$\gamma_e = \sqrt{1 + \frac{p_x^2}{m_e^2 c^2} + \frac{e^2 \mathbf{A}^2}{m_e^2 c^2}}\,. \tag{2.102}$$

From Gauss' law the longitudinal electric field is given by

$$\frac{\partial E_x}{\partial x} = -\frac{e\tilde{n}}{\epsilon_0} \tag{2.103}$$

and the continuity equation reads

$$\frac{\partial n_e}{\partial t} + \frac{\partial}{\partial x}\left(\frac{n_e p_x}{\gamma_e}\right) = 0\,. \tag{2.104}$$

From the time-dependent Maxwell equations and the definition of the vector potential

$$c^2 \nabla \times (\nabla \times \mathbf{A}) + \frac{\partial^2 \mathbf{A}}{\partial t^2} = \frac{\mathbf{J}}{\epsilon_0} - \nabla \frac{\partial \phi}{\partial t}. \tag{2.105}$$

It can be shown that the longitudinal component of the current density cancels the final term on the right hand side[90], so substituting $J_\perp = -en_e v_\perp$ and working in the Coulomb gauge $\nabla \cdot \mathbf{A} = 0$ the wave equation for the vector potential is

$$\left(\frac{\partial^2}{\partial t^2} - c^2\nabla^2 + \frac{\omega_p^2}{\gamma_e}\frac{n_e}{n_0}\right)\mathbf{A} = 0 \tag{2.106}$$

where the last term is due to the transverse plasma current. Finally, by differentiating Eq. 2.102 and substituting into the total derivative in the fluid momentum equation

$$\frac{\partial p_x}{\partial t} = -eE_x - m_e c^2 \frac{\partial \gamma_e}{\partial x} \tag{2.107}$$

The set 2.102, 2.103, 2.104, 2.106 and 2.107 is then linearised so terms including the small quantities $\tilde{\mathbf{A}}$ and $\tilde{n}$ are only included to first order. If the vector potential is normalised to $e/m_e c$ and $\tilde{n}$ to n_0 we have

[88] [190] A. S. Sakharov *et al. Physical Review E.* (1994).
[89] [131] P. Gibbon. *Short Pulse Laser Interactions with Matter: An Introduction.* (2005).
[90] [182] J. D. Jackson. *Classical Electrodynamics.* (1998).

$$\left(\frac{\partial^2}{\partial t^2} - c^2\nabla^2 + \Omega_p^2\right)\tilde{\mathbf{A}} = -\Omega_p^2\tilde{n}\mathbf{A}_0 + \Omega_p^2\left(\frac{\mathbf{A}_0}{\gamma_0^2}\right)\left(\mathbf{A}_0\cdot\tilde{\mathbf{A}}\right) \tag{2.108}$$

$$\left(\frac{\partial^2}{\partial t^2} + \Omega_p^2\right)\tilde{n} = \frac{c^2}{\gamma_0^2}\nabla^2\left(\mathbf{A}_0\cdot\tilde{\mathbf{A}}\right) \tag{2.109}$$

where, following the notation of Sahkarov, the relativistically-lowered plasma frequency is $\Omega_p = \omega_p/\sqrt{\gamma_0}$ and γ_0 is the relativistic correction due to the electron quiver motion. Eq. 2.108 contains two sources of nonlinearity – the direct coupling between plasma and electromagnetic waves through the current term, and the first-order relativistic modulation of the plasma frequency. From here on we begin to make assumptions about the structure of the solution. From Eq. 2.108 the electromagnetic perturbation $\tilde{\mathbf{A}}$ is driven by a term $\tilde{n}\mathbf{A}_0$ – this will naturally contain oscillatory parts at frequencies $\omega_0 \pm \omega_p$. Similarly Eq. 2.109 demonstrates that the electron density fluctuation is driven by the beat of the incident and scattered light, which will contain oscillatory terms at ω_p. The seed electron density fluctuation forms electromagnetic sidebands, either of which may beat with the incident light to reinforce the density fluctuation.

Bearing this in mind a trial set of solutions are adopted for the plasma density and scattered electromagnetic perturbations which evolve (for an underdense plasma) over very different timescales $\tau_0 = \omega_0^{-1} \ll \tau_p = \omega_p^{-1}$. The density fluctuation, pump field and scattered light are represented as

$$\tilde{n} = \frac{1}{2}\left(n_1 e^{i\psi} + n_1^\star e^{-i\psi}\right) \tag{2.110}$$

$$\mathbf{A}_0 = \frac{1}{2}\left(\mathbf{a}_0 e^{i\psi_0} + \mathbf{a}_0^\star e^{-i\psi_0}\right) \tag{2.111}$$

$$\tilde{\mathbf{A}} = \frac{1}{2}\left(\mathbf{a}_+ e^{i\psi_+} + \mathbf{a}_+^\star e^{-i\psi_+} + \mathbf{a}_- e^{i\psi_-} + \mathbf{a}_-^\star e^{-i\psi_-}\right). \tag{2.112}$$

where the phase terms are

$$\psi \equiv \mathbf{k}\cdot\mathbf{x} - \omega t \tag{2.113}$$

$$\psi_0 \equiv \mathbf{k_0}\cdot\mathbf{x} - \omega_0 t \tag{2.114}$$

$$\psi_\pm \equiv (\mathbf{k_0} \pm \mathbf{k})\cdot\mathbf{x} - (\omega_0 \pm \omega)t = \psi_0 \pm \psi \tag{2.115}$$

and the prefactors are assumed constant. The electron density perturbation evolves on the slow timescale set by ψ, and so will only be driven effectively by terms phase-matched at ψ also. The plasma wave parameters $(\omega, \mathbf{k})$ are those being determined in this analysis, and we assume the electromagnetic perturbation only consists of waves at $(\omega_0 \pm \omega, \mathbf{k}_0 \pm \mathbf{k})$. This construction means that higher-order Raman satellites are absent here, though it is possible to generalise this approach to arbitrarily many terms[91], which if continued *ad infinitum* would constitute an exact oscillatory solution of the coupled PDEs 2.108 and 2.109 by Bloch's theorem. The product $(\mathbf{A}_0\cdot\tilde{\mathbf{A}})$ driving the electron perturbation in

[91] [191] B. Quesnel *et al. Physical Review Letters.* (1997); [192] B. Quesnel *et al. Physics of Plasmas.* (1997).

Eq. 2.109 is expressed in terms of the trial solutions

$$\mathbf{A}_0 \cdot \tilde{\mathbf{A}} = \frac{1}{4}\Big(\mathbf{a}_0 \cdot \mathbf{a}_+ e^{i(\psi_0+\psi_+)} + \mathbf{a}_0^\star \cdot \mathbf{a}_+ e^{i(-\psi_0+\psi_+)} + \mathbf{a}_0 \cdot \mathbf{a}_+^\star e^{i(\psi_0-\psi_+)} + \mathbf{a}_0^\star \cdot \mathbf{a}_+^\star e^{i(-\psi_0-\psi_+)} \\ \mathbf{a}_0 \cdot \mathbf{a}_- e^{i(\psi_0+\psi_-)} + \mathbf{a}_0^\star \cdot \mathbf{a}_- e^{i(-\psi_0+\psi_-)} + \mathbf{a}_0 \cdot \mathbf{a}_-^\star e^{i(\psi_0-\psi_-)} + \mathbf{a}_0^\star \cdot \mathbf{a}_-^\star e^{i(-\psi_0-\psi_-)}\Big). \tag{2.116}$$

As the electron perturbation is forced to have a phase ψ, this must be matched in the ponderomotive driving term. Examining the second and seventh terms in Eq. 2.116

$$-\psi_0 + \psi_+ = -\psi_0 + (\psi_0 + \psi) = \psi$$
$$\psi_0 - \psi_- = \psi_0 - (\psi_0 - \psi) = \psi$$

while the other terms contain oscillations near ψ_0 or $2\psi_0$. Ignoring the fast terms, the surviving elements of Eq. 2.109 are

$$\left(\frac{\partial^2}{\partial t^2} + \Omega_p^2\right)\frac{n_1}{2}e^{i\psi} = \frac{c^2}{4\gamma_0^2}\nabla^2\left((\mathbf{a}_0^\star \cdot \mathbf{a}_+ + \mathbf{a}_0 \cdot \mathbf{a}_-^\star)e^{i\psi}\right). \tag{2.117}$$

Taking the derivatives $\partial_t \to -i\omega$, $\nabla \to i\mathbf{k}$ the amplitude of the density fluctuation is

$$n_1(\omega, \mathbf{k}) = -\frac{c^2\mathbf{k}^2}{2\gamma_0^2\left(\Omega_p^2 - \omega^2\right)}\left(\mathbf{a}_0^\star \cdot \mathbf{a}_+ + \mathbf{a}_0 \cdot \mathbf{a}_-^\star\right) \equiv -\frac{c^2\mathbf{k}^2}{2\gamma_0^2 D_e}\left(\mathbf{a}_0^\star \cdot \mathbf{a}_+ + \mathbf{a}_0 \cdot \mathbf{a}_-^\star\right) \tag{2.118}$$

where the electron dispersion relation $D_e = \Omega_p^2 - \omega^2$ has been defined. At resonance $D_e \to 0$ and the electron perturbation diverges, indicating that as expected the electron perturbation will be dominated by oscillations near the plasma frequency.

The same phase-matching procedure must also be carried out for Eq. 2.108, but there is an important difference between linear and circular polarisations. For circular polarisation γ_0 is constant in time, as is the scalar product $\mathbf{A} \cdot \tilde{\mathbf{A}}$, and there are only two phase-matched terms driving $\mathbf{a}_+$ or $\mathbf{a}_-$. Substituting trial solutions in Eq. 2.108 and separating the two scattered waves explicitly yields

$$\mathbf{a}_+ = \frac{\Omega_p^2\mathbf{a}_0}{4\gamma_0^2 D_+}\left(\frac{c^2\mathbf{k}^2}{D_e} + 1\right)\left(\mathbf{a}_0 \cdot \mathbf{a}_-^\star + \mathbf{a}_0^\star \cdot \mathbf{a}_+\right) \tag{2.119}$$

$$\mathbf{a}_- = \frac{\Omega_p^2\mathbf{a}_0}{4\gamma_0^2 D_-}\left(\frac{c^2\mathbf{k}^2}{D_e} + 1\right)\left(\mathbf{a}_0 \cdot \mathbf{a}_+^\star + \mathbf{a}_0^\star \cdot \mathbf{a}_-\right) \tag{2.120}$$

$$D_\pm(\omega, \mathbf{k}) = -(\omega_0 \pm \omega)^2 + c^2(\mathbf{k}_0 \pm \mathbf{k})^2 + \Omega_p^2 \tag{2.121}$$

where $D_\pm$ are the dispersion relations for the scattered electromagnetic waves which, similarly to D_e, go to zero at resonance. The first set of parentheses represent the two contributions to the scattering process – density modulations proportional to D_e^{-1} and relativistic mass modulations of the electrons. For linear polarisation the scalar products

contain fast terms oscillating at $2\psi_0$, as does γ_0. It is possible for many more terms to be phase-matched, complicating the expressions for $\mathbf{a}_{+/-}$. At high a_0 the additional harmonic components of the electron currents generate scattered waves at[92] $N\omega_0 \pm \Omega_p$ for integer N, but at low plasma densities there is little interaction between scattered harmonics. In the regime $\omega_p \ll \omega_0$ the $N = 1$ scattered waves considered here are equivalent to the circularly polarised case, but for $N > 1$ and higher densities the full problem must be solved. For $a_0 \lesssim 1$ the linear polarisation problem is tractable and an analytic dispersion relation follows[93].

To solve Eqs. 2.119 and 2.120 it is simplest to assume the polarisation of the scattered and incident waves are the same, valid for circular polarisation and linear polarisation when the scattering direction is out of the polarisation plane. In this case the scalar products become normal products and Eqs. 2.120 and 2.119 take the matrix form

$$\begin{pmatrix} \frac{\Omega_p^2 a_0^2}{4\gamma_0^2 D_+}\left(\frac{c^2\mathbf{k}^2}{D_e}+1\right)-1 & \frac{\Omega_p^2 a_0^2}{4\gamma_0^2 D_+}\left(\frac{c^2\mathbf{k}^2}{D_e}+1\right) \\ \frac{\Omega_p^2 a_0^2}{4\gamma_0^2 D_-}\left(\frac{c^2\mathbf{k}^2}{D_e}+1\right) & \frac{\Omega_p^2 a_0^2}{4\gamma_0^2 D_-}\left(\frac{c^2\mathbf{k}^2}{D_e}+1\right)-1 \end{pmatrix} \begin{pmatrix} a_+ \\ a_- \end{pmatrix} = \mathbf{0} \tag{2.122}$$

where $\mathbf{a} = \mathbf{a}^\star$ for real $\mathbf{a}$. Nontrivial solutions for the scattered wave amplitudes require the determinant of the matrix to vanish which yields the dispersion relation in standard form

$$D_{SRS}(\omega, \mathbf{k}) = \frac{\omega_p^2 a_0^2}{4\gamma_0^3}\left(\frac{c^2k^2}{D_e}+1\right)\left(\frac{1}{D_+}+\frac{1}{D_-}\right)-1=0\,. \tag{2.123}$$

It is interesting to note that this dispersion relation is very similar to that of the general coupled oscillator in Eq. 2.100 (aside from relativistic corrections), and emphasises the parametric nature of the instability. The coupling term is the plasma frequency, and the analogous pump strength is a_0. A particular plasma wave $(\omega, \mathbf{k})$ is resonant when it solves $D_{SRS}(\omega, \mathbf{k}) = 0$ and the SRS response of the plasma is dominated by these modes. As posed in Eq. 2.122 it is clear that the inclusion of extra modes requires the addition of rows and columns to the matrix, and so the algebraic analysis of higher-order terms in the dispersion relation rapidly becomes challenging. Solving this equation computationally is simplest when in the form of a polynomial, written in full in appendix A.1.

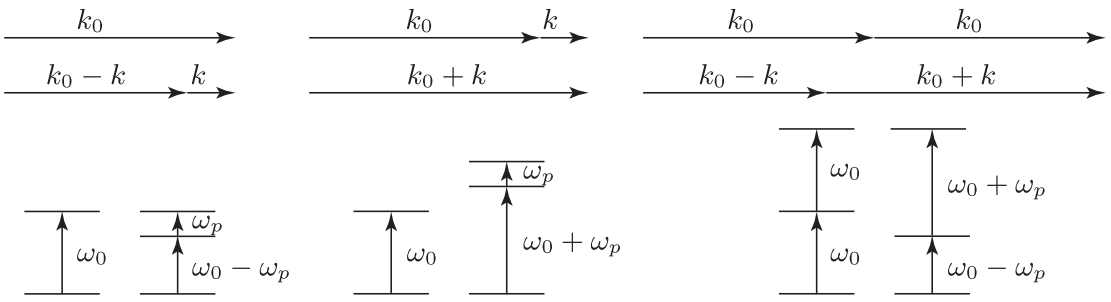


Figure 2.13: The momentum and energy relations schematically illustrated for 3 and 4-wave scattering processes

The matching conditions imposed on the scattered wave $\mathbf{a}_-$ can be interpreted as the

[92] [193] H. Barr *et al. Physics of Plasmas.* (2000).
[93] [194] C. J. McKinstrie *et al. Physics of Fluids B: Plasma Physics.* (1992).

conservation of energy and momentum, $\omega_- = \omega_0 - \omega_p$ and $\mathbf{k}_- = \mathbf{k}_0 - \mathbf{k}$. As illustrated in Fig. 2.13, conservation of energy and momentum for the up-shifted wave requires the interaction of four waves. The D_+ term is discarded for now as a non-resonant higher-order process, and we calculate the growth rate of the down-shifted wave $\mathbf{a}_-$ with wave vector $\mathbf{k}_s = \mathbf{k}_0 - \mathbf{k}$. At relativistic intensities a range of ω will solve the dispersion relation 2.123, but if $a_0 \ll 1$ then $\gamma_0 \to 1$ and $\omega \to \omega_p$. An instability will develop if for a particular $\mathbf{k}$, ω becomes complex – inserting into the phase term of the trial solution results in exponential growth. With the ansatz $\omega = \omega_p + i\gamma$ it is possible to solve to find the growth rate γ as a function of $\mathbf{k}$. The final assumption to make is that the scattered wave is exactly resonant and

$$c^2(\mathbf{k}_0 - \mathbf{k})^2 = (\omega_0 - \omega_p)^2 - \omega_p^2\,, \tag{2.124}$$

then the growth rate is given by[94]

$$\frac{\omega_p^2 c^2 k^2 a_0^2}{4} = \left(\omega_p^2 - (\omega_p + i\gamma)^2\right)\left((\omega_0 - \omega_p)^2 - (\omega_0 - (\omega_p + i\gamma))^2 + \omega_p^2\right) \tag{2.125}$$

$$\Rightarrow \gamma = \frac{a_0 ck}{4}\sqrt{\frac{\omega_p}{\omega_0 - \omega_p}}\,. \tag{2.126}$$

where in the last equality it was assumed that $\gamma \ll \omega_p$. Of course this limits the validity of Eq. 2.126 to situations involving very low growth rates. The resonance condition 2.124 determines the plasma wave wavevector $\mathbf{k}$ parametrised by a magnitude k and a direction relative to the laser pulse θ

$$k = k_0 \cos\theta \pm k_0\sqrt{\cos^2\theta - 2\left(\frac{\omega_p}{\omega_0} - \frac{\omega_p^2}{2\omega_0^2}\right)}. \tag{2.127}$$

For a given plasma wave angle θ there are two solutions for the magnitude k, illustrated in Fig. 2.14. Taking the positive sign in the expression for k requires the scattered wave to propagate at a much greater angle to the laser axis, and this solution is known as Raman back scatter (RBS). The opposite case is Raman forward scatter (RFS) where the scattered wave propagates much closer to the laser axis. For $\omega_p \ll \omega_0$ $k_{RBS} \approx 2k_0 \cos\theta$ and for small θ $k_{RFS} \approx k_p/\cos\theta$. The magnitude of the scattered wave vector k_s is the same in both cases.

From Eq. 2.126 for small γ, $\gamma \propto k$ and the RBS solution is expected to grow more rapidly. As θ increases the two solutions (intersections with the k_s circle) move closer together until they meet at a critical angle. This is the angle of maximum k_{RFS} and minimum k_{RBS}, or equivalently maximum/minimum growth rate for RFS/RBS respectively. The maximum growth rate solution for RFS would be expected to be the dominant RFS scattering mode, and has been observed in PIC simulations by Forslund *et al.*[95], and in experiment by Matsuoka *et al.*[96]. We have obtained similar results in our experiments as discussed in chapter 4. The maximum growth rate for RBS occurs when the scattered electromagnetic

[94] [195] W. L. Kruer. *The Physics of Laser Plasma Interactions.* (1988).
[95] [196] D. W. Forslund *et al. Physical Review Letters.* (1985).
[96] [197] T. Matsuoka *et al. Physical Review Letters.* (2010).

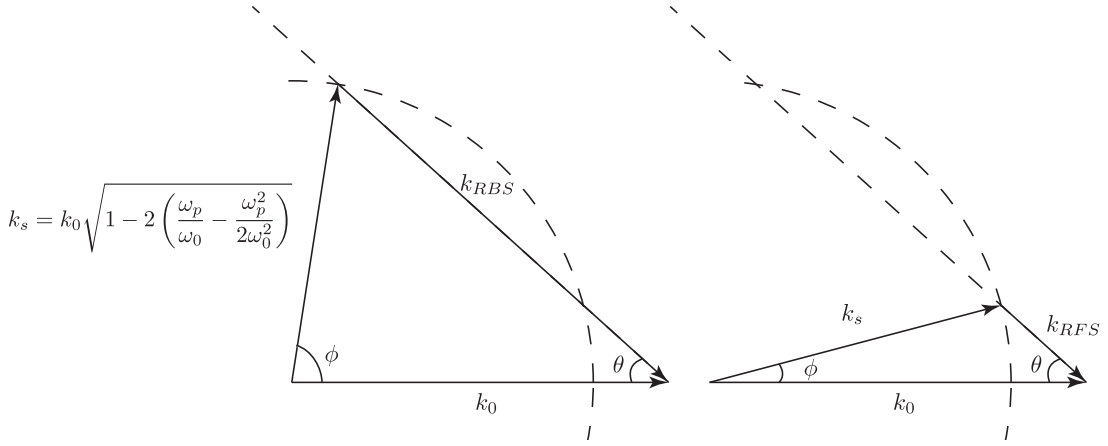


Figure 2.14: For a given plasma wave angle θ there are two solutions for the direction of the scattered light $\mathbf{k}_s$

wave propagates backwards relative to the laser pulse, i.e. $\theta \to 0°$ and $\phi \to 180°$. As a measure of the characteristic growth rate of the RBS mode (including side scatter), note that for pure back scatter k-matching requires $k \approx 2k_0 - \omega_p/c \approx 2k_0$ for underdense plasmas. From Eq. 2.126 $\gamma/\omega_p \approx a_0 c k_0/2\sqrt{\omega_0\omega_p} \propto a_0\sqrt{\omega_0/\omega_p}$ which sets the scale of the growth rate.

For higher intensities $a_0 \gg (\omega_p/\omega_0)^{1/2}$, so rather than specifying the form of ω it is instead possible to solve for it directly. Approximating in this way yields the growth rate[97]

$$\omega^3 = -\frac{\omega_0\omega_p^2 a_0^2}{2\gamma_0^2} \quad \Rightarrow \quad \frac{\gamma}{\omega_p} = \frac{\Im(\omega)}{\omega_p} = \frac{\sqrt{3}}{2}\left(\frac{\omega_0}{\omega_p}\right)^{1/3}\left(\frac{a_0^2}{2\gamma_0^3}\right)^{1/3} \tag{2.128}$$

which begins to fall for laser intensities greater than $a_0 = \sqrt{2}$. This is known as the 'strong coupling' regime of stimulated Raman scattering. The frequency of the unstable plasma mode is given by the real part of Eq. 2.128, and can be far from the plasma frequency. These modes are therefore not natural resonant modes of the plasma but instead are strongly modified by the laser fields. Scattering in the strong coupling regime is also known as stimulated Compton scattering (SCS), where the collective plasma response becomes irrelevant to the motion of the scattering electrons and can be neglected[98].

Solving Eq. 2.123 in full must be performed numerically – here we pick a particular plasma wave $\mathbf{k}$ parametrised by (k, θ) and find the ω roots. In general the root will be a complex number $\omega = \Re(\omega) + i\Im(\omega)$, the real part representing the oscillatory frequency of the plasma oscillation and the imaginary part the growth rate. Taking

$$(\mathbf{k}_0 \pm \mathbf{k})^2 = k_0^2 + k^2 \pm 2k_0 k\cos\theta \tag{2.129}$$

the growth rates are plotted in the (k, θ) plane for various different intensities in Fig. 2.15. At low intensities where Eq. 2.126 holds, growing modes (i.e. $\Im(\omega) \neq 0$) lie precisely along the resonance curve given by Eq. 2.127. Along this line $\Re(\omega) = \omega_p$ to a very good

[97] [190] A. S. Sakharov *et al. Physical Review E.* (1994).
[98] [186] J. F. Drake. *Physics of Fluids.* (1974).

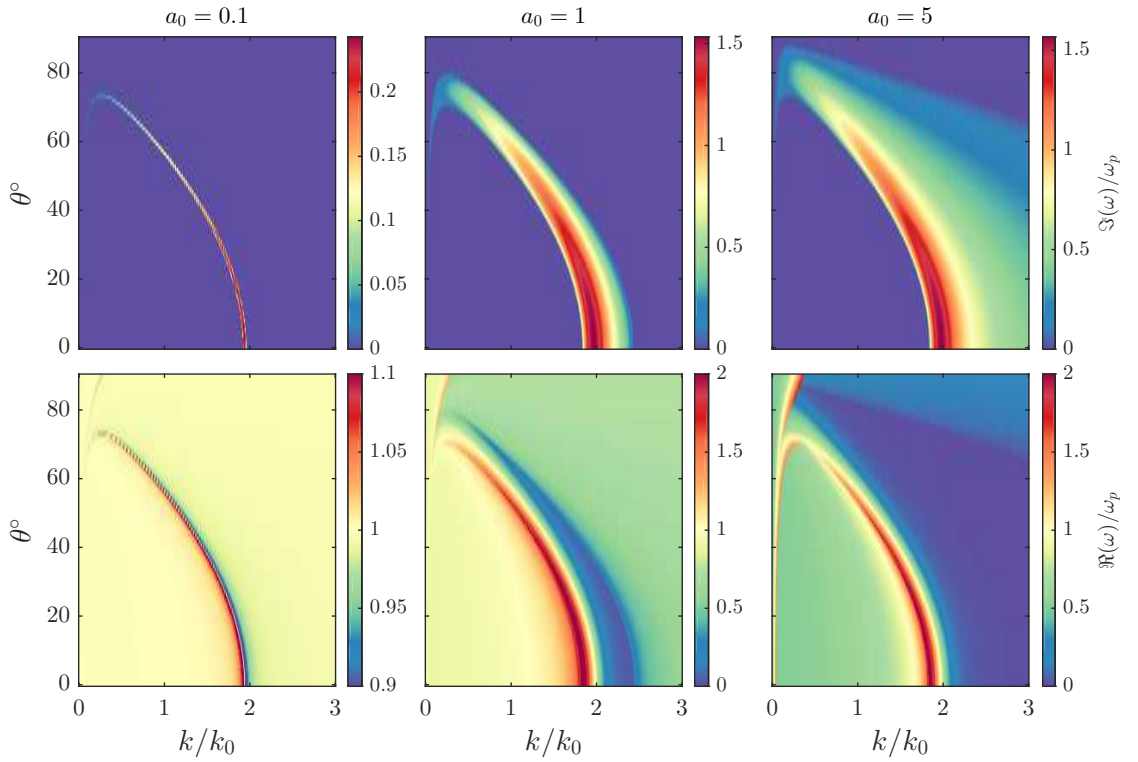


Figure 2.15: The growth rates (top) and oscillatory frequency (bottom) of plasma waves satisfying $D_{SRS} = 0$ are plotted as a function of k and θ for different laser intensities. Growth rates are largest for pure backscatter ($\theta = 0°$) and the unstable region of the (k, θ) plane grows with laser intensity.

approximation. At relativistic intensities $a_0 > 1$, growth occurs over a much broader range of k, extending to shorter plasma wavelengths $k/k_0 > 2$. In this region the plasma wave frequency drops towards zero, indicating that the scattered wave can have a frequency near the pump wave. At higher intensity still the peak growth rates begin to fall as $a_0^{-1/3}$, but the area of the (k, θ) plane where resonance can occur continues to broaden.

Assuming the scattered wave has frequency and k-vector $(\omega_0 - \Re(\omega), \mathbf{k}_0 - \mathbf{k})$, these growth rates can also be plotted as a function of the scattered wave parameters. In Fig. 2.16 the growth rate is plotted as a function of the frequency shift of the scattered wave for different scatter directions.

In the low intensity case, the frequency shift is close to ω_p as expected for SRS in the low-amplitude limit. Growth increases for greater scattering angles, peaking for pure backscatter. The frequency width of the growth region also increases with scattering angle. In the relativistic case growth rates are much higher, as are the frequency shifts. In the forward direction the frequency shift stays near ω_p, but at greater angles the frequency shift can be much different, extending down to zero.

It should be reiterated that the dispersion relation presented here is formally correct only for 1D SRS involving two scattered modes. It is shown by several authors[99] that transverse components of the plasma wave couple together higher-order scattered waves of the form

[99] [198] A. S. Sakharov *et al. Physics of Plasmas.* (1997); [191] B. Quesnel *et al. Physical Review Letters.* (1997); [192] B. Quesnel *et al. Physics of Plasmas.* (1997); [199] H. Barr *et al. Physics of Plasmas.* (2003).

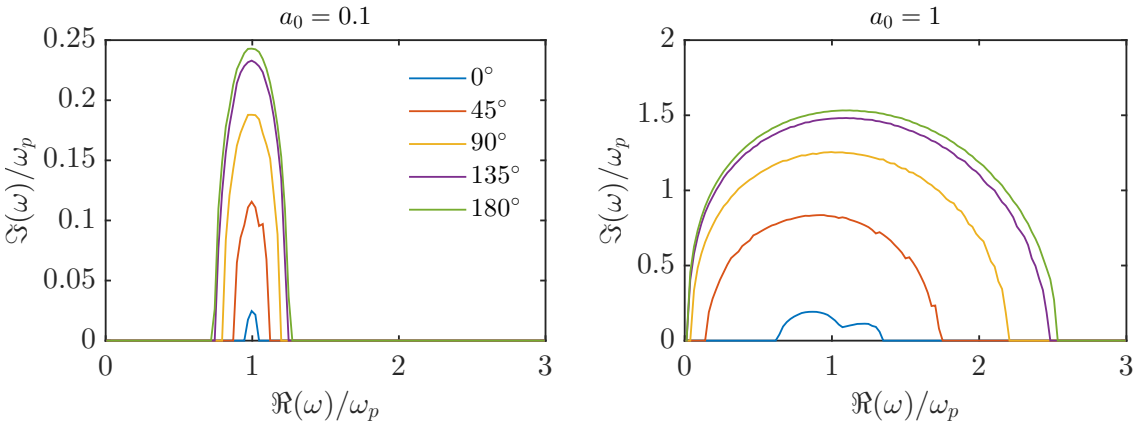


Figure 2.16: SRS growth rates are plotted as a function of scattered wave frequency shift. The angles correspond to the direction of the scatter relative to the pump.

$N\mathbf{k}_0 \pm \mathbf{k}$ where N is an integer. In high-density plasmas where ω_p is of similar magnitude to ω_0 this results in the merging of otherwise isolated growth regions in the (k, θ) plane, enhancing growth rates significantly. In the low density case, the effect is to introduce additional islands of growth along the curves $k \approx 2Nk_0 \cos\theta$, representing the growth of additional scattered harmonics of the pump wave. If $\omega_p \ll \omega_0$, these regions of growth are well separated and the growth rates of the $N = 1$ mode studied here are little affected.

2.9.3 Raman amplification

The Raman scattering mechanism has been proposed as a way to create laser pulses of unprecedented intensity[100], where the scheme is known as Raman amplification. The essentials of the relevant theory are covered here because the mechanism relies on the plasma-mediated interaction between two laser pulses. This interaction has been much studied in the context of Raman amplification, and the physical processes involved are relevant to the results and discussion of chapter 5.

The so-called backwards Raman amplifier (BRA) scheme uses a plasma as a laser amplifying medium, which crucially cannot be damaged as normal amplifying media can. A short, low energy seed pulse is overlapped with a counterpropagating long, high energy pump pulse which continually Raman back-scatters into the seed. Back-scattering is preferable as the growth rates are maximised and the interaction time between the two pulses is longest. Simulations[101] indicate that such a scheme is capable of producing significant amplification, but experimental implementation has thus far remained challenging. Growth rates are enhanced if the pump and seed lasers are of higher intensity, but the plasma response becomes nonlinear and difficult to control[102] and the amplification efficiency drops[103].

To derive the basic linear BRA equations in the initial regime where pump depletion is neglected, assume $a_0 \ll 1$ and that all evolution is one-dimensional. Now a is interpreted as the amplitude of a seed pulse down-shifted in frequency by ω_p in the presence of a

[100] [200] V. Malkin *et al. Physical Review Letters.* (1999).
[101] [201] R. M. G. M. Trines *et al. Nature Physics.* (2010); [202] R. M. G. M. Trines *et al. Physical Review Letters.* (2011).
[102] [203] V. Malkin *et al. European Physical Journal: Special Topics.* (2014).
[103] [204] N. A. Yampolsky *et al. Physics of Plasmas.* (2008).

constant-amplitude pump pulse a_0. Eqs. 2.108 and 2.109 become

$$\left(\frac{\partial^2}{\partial t^2} - c^2\frac{\partial^2}{\partial x^2} + \omega_p^2 + \nu\right) a = -\omega_p^2 n a \tag{2.130}$$

$$\left(\frac{\partial^2}{\partial t^2} + \omega_p^2 + \nu_p\right) n = c^2\frac{\partial^2}{\partial x^2}(a_0 a) \tag{2.131}$$

where the damping rates ν and ν_p have been inserted for the seed and plasma waves respectively. The quasi-static approximation is invoked so that a and n are represented by the (complex) envelopes

$$a \to a e^{i(k_0 x - \omega_0 t)} \tag{2.132}$$

$$n \to n e^{-i\omega_p t} \tag{2.133}$$

and these are assumed to vary slowly enough that second derivatives vanish, which also implies pulse lengths much longer than λ_0 or λ_p. The coupled equations become

$$\left(\frac{\partial}{\partial t} + c\frac{\partial}{\partial x} + \nu\right) a = i\gamma_R \left(\frac{\omega_p}{\omega_0}\right)^{3/2} n \tag{2.134}$$

$$\left(\frac{\partial}{\partial t} + \nu_p\right) n = -i\gamma_R a \left(\frac{\omega_0}{\omega_p}\right)^{3/2} \tag{2.135}$$

where γ_R is the linear growth rate for pure backscatter derived above. For notational convenience in the literature n is often normalised to $i(\omega_0/\omega_p)^{3/2}$ where it is then interpreted as a normalised longitudinal electric field amplitude. In the strong-damping limit where $\nu_p n \gg \partial_t n$ and $\nu a \gg c\partial_x a$, $n \approx -i\gamma a(\omega_0/\omega_p)^{3/2}/\nu_p$ and $a(t) \propto e^{\gamma^2 t/\nu_p}$. The seed is observed to grow exponentially, and in this initial stage it can be shown that it also lengthens temporally due to the low bandwidth of Raman growth in the linear regime. Eventually growth saturates through a variety of mechanisms, including pump-depletion, secondary ion-wave instabilities[104], electron wave breaking[105] and the presence of parasitic parametric instabilities[106]. As the pump depletes it can be periodically regenerated by parametric energy transfer from the (now energetic) seed, leading to a so-called π-pulse envelope for the seed[107] (also known as 'Burnham-Chiao ringing'[108]). The frequency of the plasma wave changes as it grows due to relativistic effects and particle trapping, leading to detuning of the 3-wave resonance and saturation. Chirping the pump pulse has been shown to be effective in mitigating these effects[109]. Lowering plasma density helps prevent unwanted forward Raman scattering of the seed, but too low and the electron plasma wave breaks too soon.

[104] [205] R. Kirkwood *et al. Physical Review Letters.* (1999).
[105] [206] N. A. Yampolsky *et al. Physics of Plasmas.* (2011).
[106] [207] D. Turnbull *et al. Physics of Plasmas.* (2012).
[107] [208] V. Malkin *et al. Physics of Plasmas.* (2000).
[108] [209] D. C. Burnham *et al. Physical Review.* (1969).
[109] [210] G. Vieux *et al. New Journal of Physics.* (2011).

Superradiant amplification

If the pump and seed pulses are sufficiently strong that the ponderomotive force they produce overwhelms the Coulomb forces produced by the plasma, it is possible to enter the superradiant, or Compton regime of Raman amplification[110] (SRA). In this case the plasma electron motion isn't governed by a resonant natural mode of the plasma as in BRA, but by the ponderomotive potential of the two pulses. Electrons are trapped in the potential wells and perform oscillations at the bounce frequency $\omega_B = \sqrt{\omega_0 \omega_1 a_0 a_1}$ where $\omega_{0/1}$ and $a_{0/1}$ are the frequencies and amplitudes of the two laser pulses (see section 5.2.1). The condition for entering the superradiant regime is that the ponderomotive potential becomes deeper than the plasma wave potential, or $\omega_B > \omega_p$.
As the electrons oscillate at ω_B, they alternately transfer energy to and from the seed pulse. If the seed pulse length is comparable to ω_B^{-1} then the average energy transfer to the pulse is positive and as it propagates in x its intensity grows like[111]

$$a^2 \approx \frac{64}{9} \left(\frac{n_e}{n_c} \right)^2 \left(\frac{x}{\lambda_0} \right)^2 a_0^2. \tag{2.136}$$

The n_e^2 scaling is the origin of the term 'superradiant'[112], and this process has a direct analogue in FEL amplifiers[113] where in the electron bunch frame the Lorentz-boosted undulator fields and radiated fields take the places of pump and seed respectively.
The fact that SRA doesn't rely on the resonant response of the plasma means that it is less sensitive to frequency detuning and pulse length than BRA, and in simulations gain bandwidths can be larger than ω_p[114]. The downside is that SRA requires very large pump and seed amplitudes if the amplifier is to operate at densities where wavebreaking does not happen too soon.

2.9.4 Convective vs. absolute growth rates

As we are interested in understanding how energy is transferred from one beam to another we need to consider how these processes evolve in both space and time, and so it is important to understand the difference between an instability which grows at a fixed point in space (absolute instability) or one which grows as it propagates (convective instability). In Fig. 2.17 the spatiotemporal evolution of different instability classes are plotted, arising from some arbitrary perturbation to a plasma quantity. For the stable case, growth rates are zero or negative and perturbations don't grow. A perturbation may propagate or not but its value at any point surely decays. In the convective case the growth rates can be positive and instabilities develop. However it is not guaranteed that a perturbation grows at every point in space, but only over the range of characteristics $z = v_g t$, where $v_g(z,t)$ is the group velocity of the perturbation. A particular point in space will eventually witness local decay of the perturbation, but an observer moving at v_g will witness local exponential growth.

[110] [211] G. Shvets *et al. Physics of Plasmas.* (1997).
[111] [212] G. Shvets *et al. Physical Review Letters.* (1998).
[112] [213] R. H. Dicke. *Physical Review.* (1954).
[113] [214] R. Bonifacio *et al. Physical Review Letters.* (1994).
[114] [215] B. Ersfeld *et al. Physical Review Letters.* (2005).

For a general linearised perturbation, the Green's function determining the evolution of the perturbation away from $z = 0$, $t = 0$ is the solution of $\hat{\mathcal{L}}(\partial_z, \partial_t)G = \delta(z, t)$ where $\hat{\mathcal{L}}$ is the model-dependent linear operator governing the behaviour of G. Fourier transforming, $D(k, \omega)\tilde{G}(k, \omega) = e^{i(kz-\omega t)}$ where D is the dispersion relation of the instability, and so after an inverse transform the Green's function is

$$G(z, t) \propto \int\int \frac{e^{i(kz-\omega t)}}{D(k, \omega)} \, \mathrm{d}k \, \mathrm{d}\omega \, . \tag{2.137}$$

The modes dominating the evolution of G can be shown to satisfy[115] $D = 0$ and $\partial_k D = 0$, which for technical reasons are known as pinch points[116]. By Lorentz-transforming G and evaluating it at the pinch points it is possible to distinguish at which velocities the observer witnesses absolute growth. If growth continues in the negative direction, known as the Bers-Briggs criterion, then the instability is absolute and all spatial points become unstable.

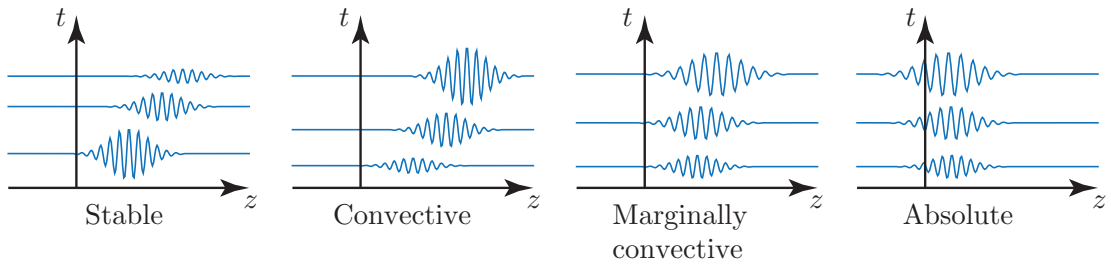


Figure 2.17: Space-time diagrams of perturbation propagation under different regimes of instability.

Turning to SRS, and assuming an absolute, or homogeneous, growth rate γ_0, a perturbation at a given position grows like $e^{\gamma_0 t}$. Considering the scattered wave propagating at $v_g \approx c$ in the direction z, the position of the perturbation is $z = ct$, and the temptation is to conclude convective growth like $e^{\gamma_0 z/c}$ of spatial rate γ_0/c. However the purely temporal growth rate is derived above under the assumption that all modes are spatially homogeneous, and therefore grow at the same rate over all space. One therefore has to be careful when applying this growth rate in a spatial sense, as the exponentiating perturbations it describes do not propagate.

In the strong-damping limit the steady-state convective spatial Raman amplification rate of a beam as it propagates in z can be derived to be[117]

$$a(z) \propto \exp\left(\frac{\gamma_0^2 z}{v_g \nu_p}\right), \tag{2.138}$$

where γ_0 is the temporal growth rate of the absolute instability, v_g is the group velocity of

[115] [216] R. J. Briggs. *Electron-Stream Interactions with Plasmas.* (1964); [217] A. Bers *et al. Physical Review Letters.* (1984); [218] A. Bers. *Basic Plasma Physics 1.* (1989).

[116] To satisfy causality one must carefully choose integration contours in Eq. 2.137, and when these contours meet, or 'pinch', the crossing point is known as a pinch point.

[117] [219] R. E. Giacone *et al. Physics of Plasmas.* (1995); [220] C. J. McKinstrie *et al. Physics of Plasmas.* (1996).

the beam and ν_p is the plasma wave damping rate. In this case the convective amplification rate is simply related to the temporal amplification rate by the group velocity of the amplified light – this was also observed in the simplified BRA equations.
Linking the weak and strong damping regimes requires a more general approach. Bobroff and Haus[118] derive a formalism for the coupling and spatio-temporal evolution of a pair of waves travelling in opposite directions, applicable to the backward Raman scattering of a weak seed in the linear regime where pump depletion is yet to occur. Malkin and Fisch apply their results to Raman amplification in the presence of damping[119] and find the spatiotemporal evolution of a seed beam in one dimension as

$$a(z,t) \propto \begin{cases} \exp\left(\frac{2\gamma_0}{c}\sqrt{z(ct-z)} - \frac{\nu}{c}(ct-z)\right) & 0 < z < ct \\ 0 & \text{otherwise} \end{cases} \tag{2.139}$$

where again γ_0 and ν represent the temporal growth and damping rates respectively. This approach was also followed by Kanaev *et al.*[120], though there the effect of damping was neglected. The effective amplification rate to consider is that of the peak of the amplified pulse, which by differentiating Eq. 2.139 has position

$$z_{\max}(t) = \frac{ct}{2}\left(1 + \frac{\nu}{\sqrt{\nu^2 + 4\gamma_0^2}}\right). \tag{2.140}$$

In the undamped case the amplified pulse has an effective group velocity $\dot{z}_{\max} = c/2$, which amounts to temporal broadening of the seed as expected in the linear regime of Raman amplification. As the damping rate increases with respect to the growth rate the group velocity tends to c and the pulse recompresses, as expected in the nonlinear regime of Raman amplification when the growth saturates. The effective exponential growth rate in general is given by the rate of growth of the peak of the amplified pulse, or

$$\frac{\partial}{\partial t}\log(a(z_{\max},t)) \equiv \gamma_{\mathit{eff}} = \frac{2\gamma_0^2}{\nu + \sqrt{\nu^2 + 4\gamma_0^2}}. \tag{2.141}$$

For weak damping $\nu/\gamma_0 \ll 1$, γ_{eff} tends to γ_0 and the homogeneous case is recovered – $a(z, t \to \infty) \to \infty\, \forall z$. For strong damping $\nu/\gamma_0 > 1$, γ_{eff} tends to γ_0^2/ν as expected. This approximate spatial growth rate is then valid in the strong-damping limit only, where it is inappropriate to take $\nu \to 0$. To illustrate these remarks, plotted in Fig. 2.19 are snapshots of the seed amplitude at different times as it propagates in z, under the influence of progressively higher damping rates. Here the full expression from Bobroff and Haus is used which permits propagation in both directions. With zero damping the seed grows everywhere, lengthening in duration and corresponding to an absolute instability. As the normalised damping rate approaches and exceeds a critical level the instability becomes convective, and takes the form of a single pulse with a well-defined group velocity. At significantly higher damping rates overall convective growth is reduced, but the group velocity of the seed is larger and the temporal spread is reduced.

[118] [221] D. L. Bobroff *et al. Journal of Applied Physics.* (1967).
[119] [222] V. Malkin *et al. Physical Review E.* (2009).
[120] [223] A. V. Kanaev *et al. Physics of Plasmas.* (1998).

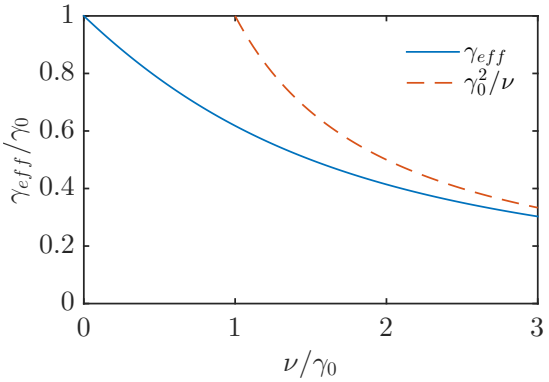


Figure 2.18: Effective growth rate as a function of damping rate

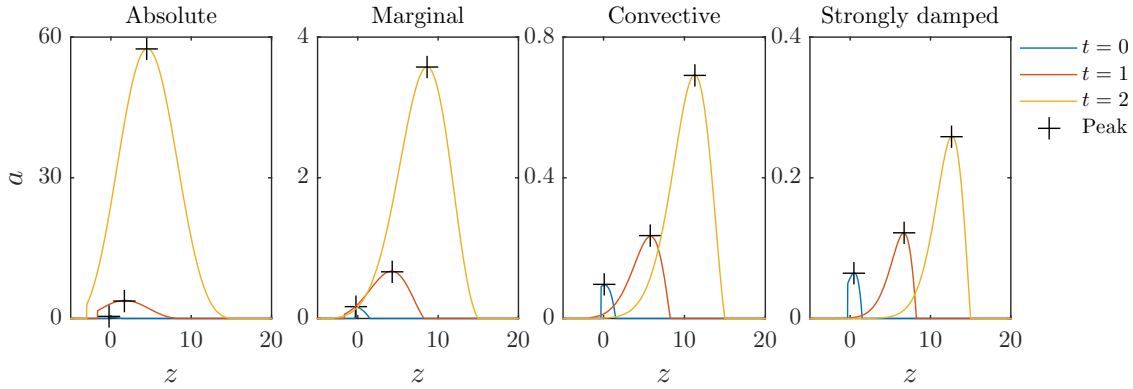


Figure 2.19: The one-dimensional spatial envelope of a perturbation at equal normalised temporal intervals for progressively higher damping rates. The peak of the envelope is marked with a cross, the group velocity is observed to increase with damping rate.

The transition from absolute to convective regimes in this formalism is restated by McKinstrie *et al.*[121] as a bound on the damping rate of the plasma wave

$$\nu_p > 2\gamma_0 \sqrt{\frac{v_e}{v_{\text{scat}}}} \tag{2.142}$$

where v_e is the group velocity of the electron plasma wave, and v_{scat} is the group velocity of the scattered electromagnetic wave. Since for plasmas produced by ultrashort laser pulses the electron temperature is low, typically $v_e \ll v_{\text{scat}} \approx c$ and this bound is not very restrictive in practice. In addition Landau damping rates are low for small temperatures and $k\lambda_D$ (appendix A.2), and most Raman instabilities in wakefield accelerators will therefore be convective in nature.

[121] [224] C. J. McKinstrie *et al. Physical Review E.* (1994).

3 Methods

In this chapter we review the experimental and numerical techniques involved in gathering and analysing the data presented in the following chapters. The generation of underdense plasma targets from neutral gas is discussed, followed by a description of the optical techniques used to diagnose the plasma shape and density profile. In particular a thorough review of numerical techniques for the analysis of interferograms is presented, all of which were incorporated into a software package used to perform the analysis of later chapters. The design and implementation of the electron spectrometer follows, along with an image processing technique to remove the effects of oblique imaging of the spectrometer plane. The chapter concludes with short sections on the PIC simulation technique used throughout this thesis, and the tomographic reconstruction technique applied in chapter 6.

3.1 Gas targets

For efficient particle acceleration, it is important to be able to produce a plasma of known composition and density profile. For wakefield acceleration it is important that the plasma is underdense and has a flat density profile. This allows the laser pulse to self-focus to a fixed matched spot size as discussed in section 2.4 and proceed with injection as soon as possible.

For >10 TW laser systems the required plasma densities for LWFA are below $10^{20}\,\mathrm{cm}^{-3}$, which is conveniently near atmospheric pressure. Perhaps the simplest target which can be conceived in this density range is then a gas jet, where a pressured gas supply is allowed to flow through a nozzle into vacuum. The atomic density witnessed by the laser is then determined by the nozzle backing pressure, the nozzle geometry, and the position of the laser focus with respect to the nozzle exit. Normal subsonic flow through a cylindrical section will tend to produce a parabolic density profile as a function of radius, and will expand into vacuum rapidly at a large expansion angle.

Supersonic flow on the other hand may be used to generate a fluid flow with a flat density profile[1], assuming the density is sampled sufficiently close to the exit of the nozzle that exit shocks (Prandtl-Meyer fans[2]) have not had time to propagate into the body of the flow. The Mach number of the flow is related to the divergence angle of the flow θ through $M = 1/\sin\theta$ and defines the sharpness of the density profile. Large Mach numbers help collimate the flow, but correspond to a rapid drop in the density of the gas as it traverses the length of the nozzle[3] (see appendix A.4) which may limit the density below a useful

[1] [225] S. Semushin *et al. Review of Scientific Instruments.* (2001).
[2] [226] F Sylla *et al. Review of Scientific Instruments.* (2013).
[3] [227] K. Schmid. *Supersonic Micro-Jets And Their Application to Few-Cycle Laser-Driven Electron*

level or necessitate large nozzle backing pressures.

The advantages of gas jets are their simplicity and robustness in a challenging environment involving hard vacuum, electrical discharges and large static magnetic fields. The fact that the gas freely expands into vacuum means that there are no fixed structures around the plasma blocking diagnostic access or the propagation of secondary laser pulses (e.g., as used in controlled injection techniques or electron-laser interaction experiments). Gas jets are limited however in the shape of density profile which can be produced. Different nozzle geometries produce profiles with sharper or softer edges but further customisation is limited, though multiple gas jet configurations have been used in several experiments to structure the density profile[4]. Alternative schemes make use of an obstruction placed in the gas flow to introduce sharp transitions in the density profile[5].

An important parameter for wakefield accelerators is the acceleration length, which is fixed for a given gas jet. More recently experimental groups have begun to use gas cells[6], where gas at much lower pressure is contained on a millisecond timescale in a short cavity. The effective length of such cells may be adjusted remotely which allows the exploration of a richer parameter space in experiments, and the gas density profile is expected to be more uniform than that produced by gas jets. It is possible to introduce longitudinal gradients in density or gas composition by altering gas inlet parameters, or separate the cell into compartments containing different gases[7] (useful for ionisation injection). Because gas is confined in a cell rather than expanding into vacuum, much lower gas pressures are required and the mechanical constraints on the material are lessened. It then becomes possible to cheaply 3D print gas cells in plastic[8].

3.2 Shadowgraphy

Plasma densities encountered in laser wakefield accelerators are low, and so the plasma is almost completely transparent to laser light (assuming, of course, the incident laser is of low enough intensity that it doesn't drive a wake or otherwise affect the plasma significantly). In order to image the structure of the plasma, the imaging technique must instead be sensitive to the phase imparted to the beam by the plasma. A simple phase-sensitive imaging technique is shadowgraphy, where a coherent laser source is used to backlight the plasma. Local gradients in the refractive index of the plasma deflect the laser wavefronts, causing parts of the beam to converge or diverge. The local beam intensity then increases or decreases respectively as the beam propagates farther from the plasma object.

To quantify this effect, it is possible to begin from Fermat's principle that the light paths $\mathcal{P}$ through a medium of inhomogeneous refractive index $\eta(x, y, z)$ minimise the total optical path length, leading to the variational statement[9]

Acceleration. (2009).

[4] [47] G. Golovin *et al. Physical Review ST-AB.* (2015); [38] H. T. Kim *et al. Physical Review Letters.* (2013); [228] W. Wang *et al. Applied Physics Letters.* (2013).

[5] [178] K. Schmid *et al. Physical Review ST-AB.* (2010); [42] A. Buck *et al. Physical Review Letters.* (2013); [179] M. Burza *et al. Physical Review ST-AB.* (2013).

[6] [50] J. Osterhoff *et al. Physical Review Letters.* (2008); [97] M. Fuchs *et al. Nature Physics.* (2009).

[7] [46] B. B. Pollock *et al. Physical Review Letters.* (2011).

[8] [229] M. Vargas *et al. Applied Physics Letters.* (2014).

[9] [230] P. K. Panigrahi *et al. Schlieren and Shadowgraph Methods in Heat and Mass Transfer.* (2012).

$$\frac{\delta}{\delta\mathcal{P}}\left(\int_{\mathcal{P}(s)}\eta(x,y,z)\,\mathrm{d}s\right)=0\,. \tag{3.1}$$

Assume that light rays are propagating predominantly in z, then $\mathrm{d}s=\sqrt{1+x'^2+y'^2}\,\mathrm{d}z$ where $x'=\mathrm{d}x/\mathrm{d}z$ etc. The Euler-Lagrange equations imply the coupled ODEs for the light path in the x and y co-ordinates

$$x'' = \frac{1}{\eta}\left(1+x'^2+y'^2\right)\left(\frac{\partial\eta}{\partial x}-x'\frac{\partial\eta}{\partial z}\right) \tag{3.2}$$

$$y'' = \frac{1}{\eta}\left(1+x'^2+y'^2\right)\left(\frac{\partial\eta}{\partial y}-y'\frac{\partial\eta}{\partial z}\right). \tag{3.3}$$

In general these equations must be solved numerically, but if the ray deflections are assumed small then x' and y' may be neglected and so

$$\frac{\mathrm{d}^2x}{\mathrm{d}z^2} = \frac{1}{\eta}\frac{\partial\eta}{\partial x}=\frac{\partial}{\partial x}\left(\log\eta\right) \tag{3.4}$$

$$\frac{\mathrm{d}^2y}{\mathrm{d}z^2} = \frac{1}{\eta}\frac{\partial\eta}{\partial y}=\frac{\partial}{\partial y}\left(\log\eta\right)\,. \tag{3.5}$$

In this limit it is possible to calculate the expected intensity modification of a beam due to ray deflections by a transparent object. Consider the setup sketched in Fig. 3.1, where parallel rays 1 and 2 of infinitesimal separation Δx are deflected by an object, then observed a distance L away.

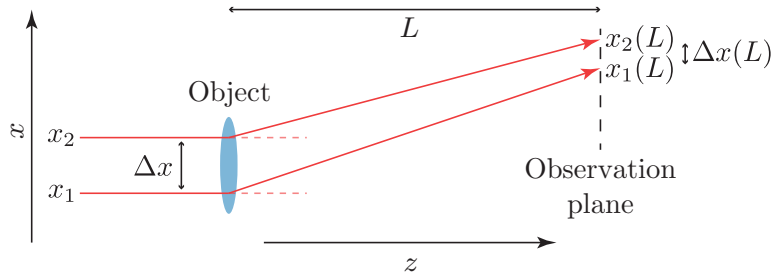


Figure 3.1: Rays initially separated by Δx are deflected by an object, and detected after propagation distance L with new separation $\Delta x(L)$.

Integrating Eq. 3.4 once over the object, the cumulative ray deflection is

$$x'=\int\frac{\partial\log\eta}{\partial x}\,\mathrm{d}z \tag{3.6}$$

so the position of ray 1 on the observation plane is

$$x_1(L)=x_1+L\int\frac{\partial\log\eta(x_1)}{\partial x}\,\mathrm{d}z\,. \tag{3.7}$$

Similarly, for ray 2 starting from position $x_2=x_1+\Delta x$

$$x_2(L)=x_2+L\int\frac{\partial\log\eta(x_2)}{\partial x}\,\mathrm{d}z=x_2+L\int\frac{\partial\log(\eta(x_1)+\Delta x\frac{\partial\eta}{\partial x})}{\partial x}\,\mathrm{d}z \tag{3.8}$$

where the refractive index is expanded to first order in x. The ray displacement at the observation plane is

$$\Delta x(L) = x_2(L) - x_1(L) = \Delta x \left(1 + L \int \frac{\partial^2 \log \eta(x_1)}{\partial x^2} \, \mathrm{d}z\right). \tag{3.9}$$

To find $\Delta y(L)$, it suffices to swap $x \leftrightarrow y$ in the above. The final step is to assume energy conservation such that the input energy $I_0 \Delta x \Delta y$ is equal to the received energy $I(L)\Delta x(L)\Delta y(L)$, where I_0 is the homogeneous input intensity and $I(L)$ is the observed intensity. Assuming the change in ray displacement is small, the relative change in local beam intensity is

$$\frac{I(x, y, L) - I_0}{I_0} = -L\nabla^2 \int \log \eta(x, y, z) \, \mathrm{d}z. \tag{3.10}$$

The shadowgram so observed is then approximately proportional to the Laplacian of the line-integrated refractive index of the object, with the intensity modulation strength increasing linearly with propagation distance. The presence of the Laplacian operator means that the shadowgraphy technique is particularly sensitive to rapid jumps in refractive index, such as those caused by shocks or boundaries between materials.

3.2.1 Experimental implementation

In laser-wakefield experiments the plasma is typically located inside a large vacuum chamber with the recording CCD far away. It is then impractical to simply allow the backlighting laser beam to propagate freely onto the CCD as the fine structure of the plasma channel diffracts rapidly and spatial information is quickly lost. Instead a lens is used to image the object plane of the plasma channel onto the CCD, which from Eq. 3.10 would imply $L = 0$ and no intensity modulation of the backlighting beam. This conclusion is only correct for an infinitesimally thin object in the geometric optics limit - in reality the plasma object is at least many laser wavelengths in size in all dimensions. Referring to Fig. 3.2 deflections occurring at the object plane do not create intensity modulations but deflections off of this plane will. In reality the laser rays are deflected continuously throughout the object and so in general intensity modulations persist despite the presence of an imaging optic.

In Fig. 3.3 a synthetic inhomogeneous plasma channel is imaged under various displacements of the object plane of the lens. At best focus (0 mm), intensity modulations persist but are minimised in magnitude. As the focal plane is translated to either side of the object the intensity modulations grow and reverse sign.

The state-of-the-art in the shadowgraphy of laser-wakefield plasmas requires the use of few-femtosecond laser pulses[10], where it becomes possible to resolve individual plasma wake periods[11], and even directly measure the dimensions of the bubble[12].

[10] [231] M. B. Schwab *et al. Applied Physics Letters.* (2013).
[11] [232] A. Buck *et al. Nature Physics.* (2011).
[12] [152] A. Sävert *et al. Physical Review Letters.* (2015).

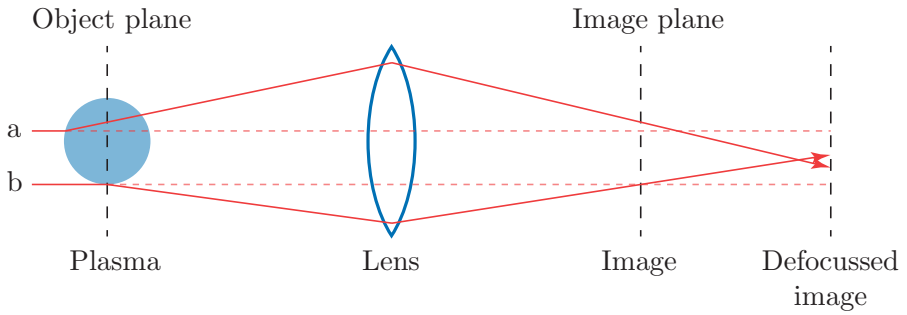


Figure 3.2: Laser rays a. and b. enter from the left, and are deflected by an extended plasma. Ray a. is deflected from a point far from the object plane of the lens, and so is displaced at the image plane. Ray b. is deflected from a point near the object plane of the lens, and so is not displaced at the image.

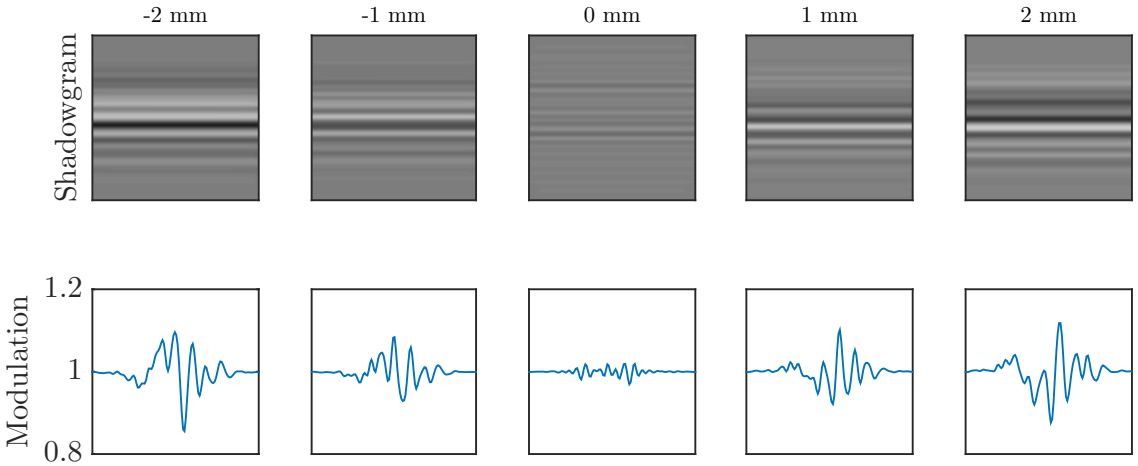


Figure 3.3: A series of synthetic shadowgrams of an inhomogeneous plasma channel, imaged at various distances out of focus. The modulation is stronger as the image is defocussed, but maintains finite amplitude at best focus due to the finite size of the object.

3.3 Interferometry

Interferometry is a general technique used wherever a signal of interest can be encoded in the phase of a stable carrier wave, which is then interfered with a second undisturbed wave. In this case the signal of interest is the electron density of a plasma, encoded into the phase of an optical laser wave propagating through the plasma. Unlike shadowgraphy, interferometry allows quantitative measurements directly. A plane wave of frequency ω_0 propagating in the z direction accumulates an integrated phase shift $\Delta\phi(x, y)$ with respect a reference wave as it traverses a plasma, where

$$\Delta\phi(x,y) = \frac{2\pi}{\lambda_0}\int\left(1-\frac{\omega_p(x,y,z)^2}{\omega_0^2}\right)^{1/2} - 1\,\mathrm{d}z \tag{3.11}$$

$$\approx \frac{\lambda_0 e^2}{4\pi m_e \epsilon_0 c^2}\int n_e(x,y,z)\,\mathrm{d}z. \tag{3.12}$$

In the second line the approximation was made that $\omega_p \ll \omega_0$ which is appropriate for

underdense LWFA plasmas. We now superpose a second identical beam which has suffered no phase shift but is travelling at an angle θ to the first, and measure the resulting intensity on a 2D plane nominally at $z = 0$ – see Fig. 3.4.

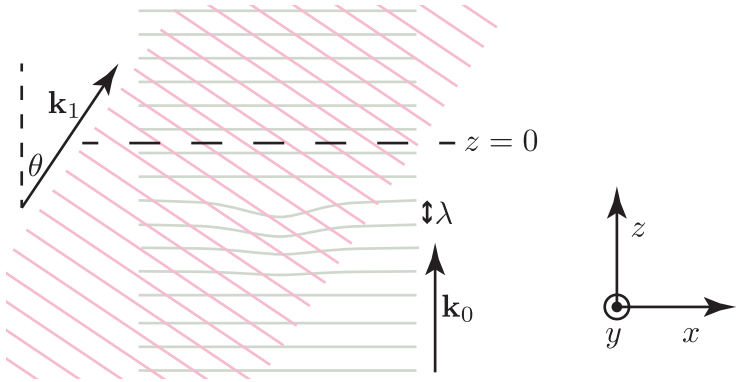


Figure 3.4: A beam experiencing some phase shift is overlapped on a detector with a reference beam travelling at relative angle θ.

It is also possible to measure phase information of a plane wave front without a second beam with the use of a wavefront sensor[13], though spatial resolution is currently low. For application to LWFA plasmas, assume each beam is pulsed with FWHM pulse duration τ_L and temporal intensity envelope $e^{-4\ln 2\, t^2/\tau_L^2}$, and that the path difference between the pulses is d. Defining $d_L = c\tau_L/\sqrt{2\ln 2}$, the instantaneous intensity is then

$$I(x,y,z,t) = \left| E_0 e^{i(\mathbf{k}_0\cdot\mathbf{x}+\Delta\phi(x,y))} e^{-\frac{(\mathbf{k}_0\cdot\mathbf{x}/k_0 - ct)^2}{d_L^2}} + E_0 e^{i(\mathbf{k}_1\cdot\mathbf{x})} e^{-\frac{(\mathbf{k}_1\cdot\mathbf{x}/k_0 - ct - d)^2}{d_L^2}} \right|^2 \tag{3.13}$$

Any CCD detector integrates signal over a time much longer than τ_L, so integrating over all time the measured intensity at the $z = 0$ plane is

$$I(x,y) \propto 1 + V(x)\cos(k_0 x \sin\theta + \Delta\phi(x,y)) \tag{3.14}$$

where the visibility function

$$V(x) = e^{-\frac{\ln 2 \sin^2\theta}{c^2\tau_L^2}\left(x - \frac{d}{\sin\theta}\right)^2}. \tag{3.15}$$

With no plasma present the detector will record simple fringes in the x-direction with fringe spacing $\lambda_f = \lambda_0/\sin\theta$. As the beams are combined at greater angles the fringe spacing shrinks in proportion. For modern CCD array-based cameras, pixel sizes are approximately 5 µm which limits $\theta < 4°$ for 800 nm light if the fringes are to be resolved. The presence of a plasma will cause a shift in the position of the fringes in the x-direction, or more generally in the direction $\mathbf{k}_0 - \mathbf{k}_1$ as projected on the (x, y) plane. Control of the fringe spacing and angle is achieved with the relative directions of the two beams.
The fringe visibility is defined in terms of the intensity maxima and minima of the interference pattern as $V = (I_{\max} - I_{\min})/(I_{\max} + I_{\min})$. $V = 1$ represents perfect

[13] [233] G. R. Plateau *et al. Review of Scientific Instruments.* (2010).

interference with the intensity dropping to zero at a fringe minimum. Here the fringe visibility varies across the image, but in principle can always be maximised to 1 at the position $x = d/\sin\theta$. If the interferometer is perfectly aligned and the path length difference is zero, peak visibility will occur in the centre of the detector ($x = 0$) and drop off towards the edges. If the interferometer is misaligned by $d = 20\,\mu\text{m}$ and the beams are at an angle $\theta = 1°$, the position for peak visibility will shift by over a millimetre and may necessitate realignment. The useful region of high-visibility fringes is determined by the pulse duration and beam angle, and is characterised by a FWHM distance on the detector of $c\tau_L/\sin\theta$, on the order of millimetres for 40 fs pulses. For few-cycle pulses and finely-spaced fringes the highly-visible region of the interferogram may shrink well below this and begin to cause problems.

Sensitivity

In measuring the plasma electron density it is important to record a phase shift with both high spatial and high phase accuracy. As the phase information is encoded in the spatial position of fringes these two issues are linked. For accurate phase retrieval it is shown in the following sections that the smallest spatial scale in $\Delta\phi(x, y)$ should be larger than a fringe spacing, so the spatial resolution is approximately bounded by λ_f. The phase resolution is determined by how accurately the location of a fringe can be established, and in principle the smallest increment of phase which can be measured is determined by the CCD pixel size Δ as $2\pi\Delta/\lambda_f$. If high visibility fringes are required across N pixels then $\lambda_f > \lambda_0 N\Delta/c\tau_L$. The best possible scenario for high visibility fringes with good phase and spatial resolution is then to use a CCD with pixels as small as possible. Spatial resolution is still limited by CCD size and pulse length, but phase resolution increases.

Practical realisation

In the laboratory the only way to ensure the required accuracy in synchronisation of the two beams is to take a single probe beam and split it in two with a 50-50 beamsplitter. In the classical Mach-Zehnder interferometry setup as sketched Fig. 3.5a), after splitting one beam passes through the plasma while the reference beam travels an otherwise equivalent path. Fine control of the relative beam timing is achieved with micrometre stages attached to the optics, typically the mirror in the path of the reference beam as indicated.

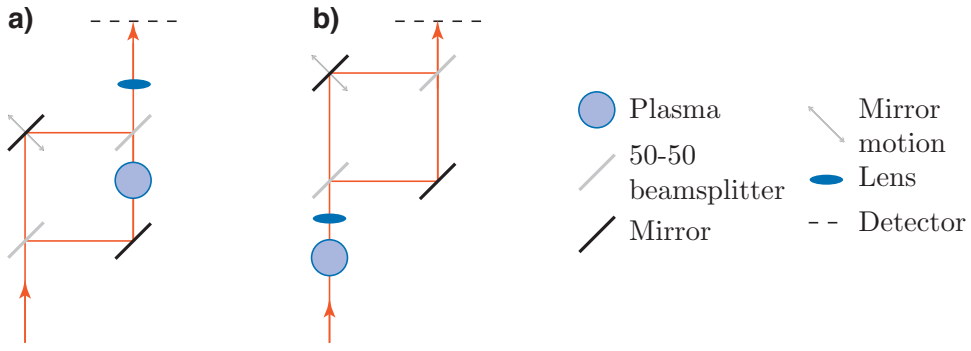


Figure 3.5: Two possible configurations for a Mach-Zehnder interferometer.

However this requires that all optics are situated inside a vacuum chamber near the plasma, which can be troublesome if realignment is required as the chamber will need to be opened. An alternative arrangement is possible if the plasma does not fill the field of view of the probe beam, often the case with LWFA plasmas which consist of a long, narrow channel. The beam may be split after it traverses the plasma as sketched in Fig. 3.5b), and offset in the interferometer such that a region of the beam experiencing no phase shift is overlaid on a region which has. Though slightly more difficult to align and not applicable to all plasmas, the advantage here is that all optics are outside of the vacuum chamber and easily accessible.

Once the beam has acquired a phase shift, as it continues to propagate towards the detector diffraction will destroy the phase information by spreading it into the rest of the beam profile, very rapidly for large phase gradients. The plane containing the plasma must be imaged onto the detector with a suitable lens system in a position where both beams pass – for best fringe visibility both beams should be of equal intensity, and hence equal size on the detector. The phase profile of the beam on the detector is then a magnified image of the phase profile imparted by the plasma. Note that the fringes are an artefact of the relative beam angles and therefore are not affected by the imaging arrangement. Adjustment of the fringe spacing may shift the position of peak fringe visibility, which can be moved back to the centre of the image by adjusting the relative pulse timing as discussed above. For short pulses and large detectors fringe visibility is not high for the entire image. If the probe pulse is sufficiently energetic it may be stretched by placing a bandpass filter centred at the laser wavelength after the plasma. The region of high visibility is then inversely proportional to the bandwidth of the pulse and can be widened easily. As the probe pulse was short when it encountered the plasma any temporal resolution in the interferogram is preserved. Fig. 3.6 displays an exemplary interferogram of a LWFA plasma employing the techniques discussed above. A 10 nm bandwidth bandpass filter was used to produce high-quality fringes with good visibility. Small pixel size and high spatial resolution allowed steep phase gradients to be observed.

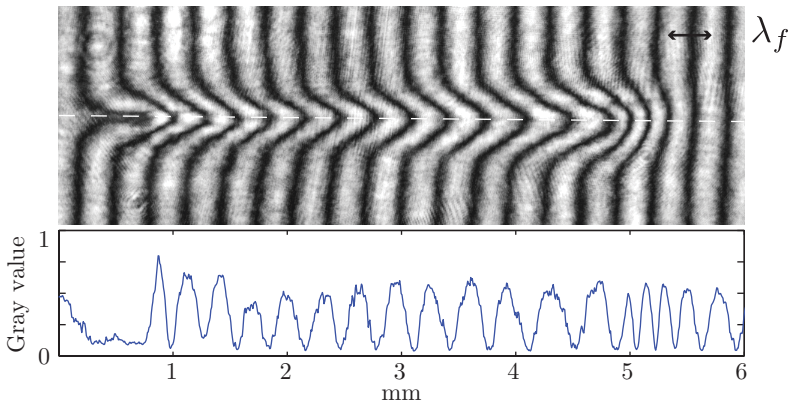


Figure 3.6: An interferogram of a LWFA plasma. The dotted white line indicates the position of the lineout plotted below.

Having recorded a satisfactory interferogram, the next steps are to

- Retrieve the phase profile,
- Unwrap the phase profile,
- Retrieve the plasma electron density.

Each step presents algorithmic challenges which are discussed in the following sections. A range of methods are covered, each with their own strengths and weaknesses. Depending on the interferogram under analysis it was often necessary to try different methods, which were incorporated into a GUI toolbox implemented in Matlab[14].

3.3.1 Phase retrieval

Fourier transform phase retrieval

Analysis of an interferogram is analysis of a sinusoidal signal, suggesting the use of Fourier-transform techniques. Indeed this was one of the first approaches in the computational analysis of interferograms[15]. In mathematical terms an interferogram $g(x, y)$ has the functional form

$$\begin{aligned} g(x,y) &= g_{bg}(x,y) + g_{\text{plasma}}(x,y)\cos(k_f x + \Delta\phi(x,y)) \\ &= g_{bg}(x,y) + \frac{1}{2}g_{\text{plasma}}(x,y)\left(e^{ik_f x}e^{i\Delta\phi(x,y)} + e^{-ik_f x}e^{-i\Delta\phi(x,y)}\right) \\ &\equiv g_{bg}(x,y) + \frac{1}{2}\left(e^{ik_f x}f(x,y) + e^{-ik_f x}f^*(x,y)\right) \end{aligned}$$

where fringes have been assumed in the x-direction and $k_f = 2\pi/\lambda_f$. $g_{bg}(x, y)$ is a noisy background to the fringes, the displacement of which encode the desired phase function $\Delta\phi(x, y)$. The term $g_{\text{plasma}}(x, y)$ represents changes to the intensity of the beam caused by the plasma rather than a phase shift. The shorthand $f(x, y) = g_{\text{plasma}}(x, y)e^{i\Delta\phi(x,y)}$ is used. After taking the two-dimensional Fourier transform to the (u, v) plane and denoting transformed functions as capitals, i.e. $\mathcal{F}(f(x, y)) = F(u, v)$, then

$$G(u,v) = G_{bg}(u,v) + \frac{1}{2}\left(F(u+k_f,v) + F^*(u-k_f,v)\right) . \tag{3.16}$$

The presence of the fringes creates copies of the Fourier transform of $f(x, y)$, translated by k_f in the Fourier plane as in Fig. 3.7.

It is important to separate the phase information from the noisy background, as a slight dip in intensity in the background can look like a fringe shift. The process is then to mask the Fourier plane by setting everything other than $F(u + k_f, v)$ to zero, generating $G_{\text{mask}}(u, v) = F(u + k_f, v)/2$. Upon performing the inverse transformation we recover $g_{\text{mask}}(x, y) = g_{\text{plasma}}(x, y)e^{ik_f x}e^{i\Delta\phi(x,y)}/2$ and the background noise is removed. However the phase influence of the fringes is still present. To remove this the same operations are performed on a reference interferogram where $\Delta\phi(x, y) = 0$, generating $g^{\text{ref}}_{\text{mask}}(x, y) = e^{ik_f x}/2$. The phase is then calculated as

[14] [234] URL: https://github.com/jasmcole/Interpret. (2015).
[15] [235] M. Takeda *et al. Journal of the Optical Society of America.* (1982).

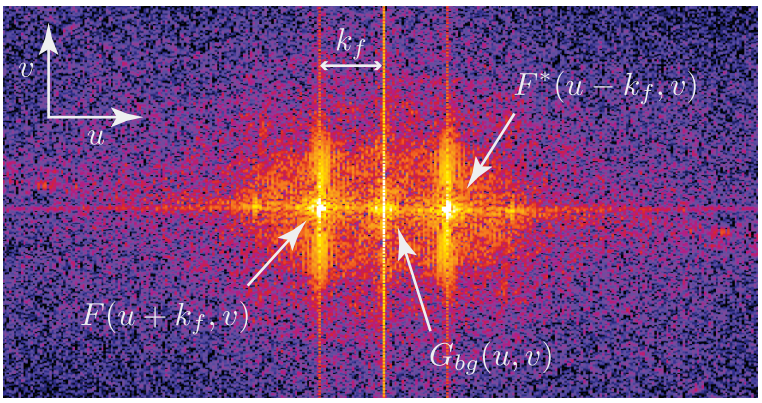


Figure 3.7: $\ln(|G(u,v)|)$ plotted in the (u,v) plane.

$$\Delta\phi(x,y) = \arg\left(\frac{g_{\text{mask}}(x,y)}{g^{\text{ref}}_{\text{mask}}(x,y)}\right) = \arg\left(\frac{g_{\text{plasma}}(x,y)e^{ik_f x}e^{i\Delta\phi(x,y)}}{e^{ik_f x}}\right) \tag{3.17}$$

and the intensity distribution as

$$g_{\text{plasma}}(x,y) = \text{abs}\left(\frac{g_{\text{mask}}(x,y)}{g^{\text{ref}}_{\text{mask}}(x,y)}\right) = \text{abs}\left(\frac{g_{\text{plasma}}(x,y)e^{ik_f x}e^{i\Delta\phi(x,y)}}{e^{ik_f x}}\right). \tag{3.18}$$

The effect of dividing by a reference set of fringes is to effectively shift F in the Fourier plane by k_f back to the origin. It is possible to perform this 'by hand' without a reference, but only when the fringes are perfectly uniform. In reality imaging imperfections cause local bending and stretching of the fringes, causing phase errors which are removed most simply with a reference image.

It should be clear from Fig. 3.7 that the size of the mask cannot be arbitrarily large. Ideally the mask should be as large as possible to capture the most spatial information, but too wide ($> 2k_f$) and it will start to encompass the noise background on one side, and a second harmonic of the fringes on the other. Fig. 3.8 shows what happens to the phase retrieval as the size of the mask changes. The spatial resolution in the retrieved phase is then limited to λ_f, so as discussed above the narrower the fringes the higher the spatial resolution.

Hilbert transform phase retrieval

The Hilbert Transform is a simple integral transform[16] commonly used in signal processing, and takes its simplest form in the frequency domain where it is a multiplication. If $\mathcal{F}$ denotes the Fourier transform from the x to the u domain, then the Hilbert transform of $f(x)$ is

$$\mathcal{F}(\mathcal{H}(f(x))) = \sigma(u)\mathcal{F}(f(x)) \tag{3.19}$$

where the function

[16] [236] S. G. Krantz. *Explorations in Harmonic Analysis.* (2009).

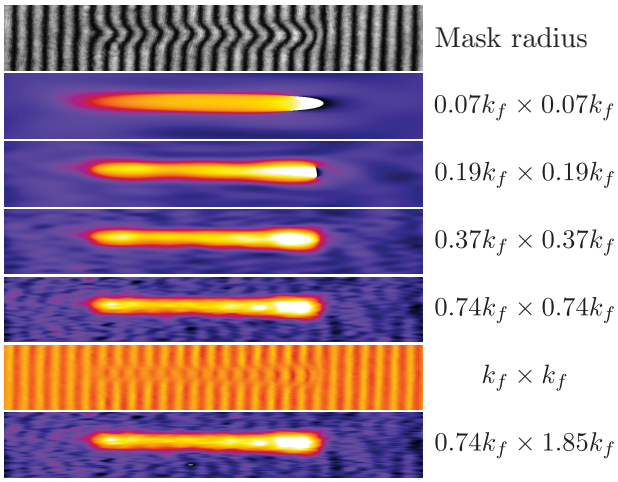


Figure 3.8: The original interferogram along with retrieved phases for different mask radii $u \times v$.

$$\sigma(u) = \begin{cases} i & u < 0 \\ -i & u > 0 \end{cases} \tag{3.20}$$

so the positive and negative frequency components of f are rotated in opposite directions – see Fig. 3.9. In the x domain this is equivalent to a convolution of f with $\mathcal{F}^{-1}(\sigma(u)) = (\pi x)^{-1}$, or the integral transform

$$\mathcal{H}(f(x)) = \frac{1}{\pi} PV \int_{-\infty}^{\infty} \frac{f(y)}{x - y}\, \mathrm{d}y\,. \tag{3.21}$$

where PV denotes the principle value of the integral.

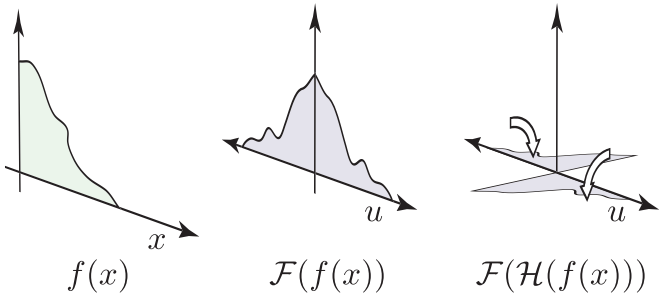


Figure 3.9: In the Fourier domain, upon a Hilbert transform the positive and negative frequency components of $f(x)$ are rotated by $\pi/2$ about the real axis in opposite directions.

A useful identity is $\mathcal{H}(e^{iux}) = -i\,\mathrm{sgn}(u)e^{iux}$ which implies $\mathcal{H}(\sin(ux)) = -\cos(ux)$ and $\mathcal{H}(\cos(ux)) = \sin(ux)$. Interpreted in the context of signal processing the Hilbert transform generates the quadrature component of a signal.

A second useful property of the Hilbert transform concerns the transform of products, known as Bedrosian's identity[17]. Define two generally complex functions $f(x)$ and $g(x)$ with Fourier transforms $F(u)$ and $G(u)$. Bedrosian's theorem states that if two functions

[17] [237] E. Bedrosian. *Proceedings of the IEEE.* (1962).

are supported on different intervals in the frequency domain, i.e. $F(u) = 0 \; \forall \; |u| > a$ and $G(u) = 0 \; \forall \; |u| < a$ (Fig. 3.10), then the Hilbert transform of the product is factored as

$$\mathcal{H}(f(x)g(x)) = f(x)\mathcal{H}(g(x)). \tag{3.22}$$

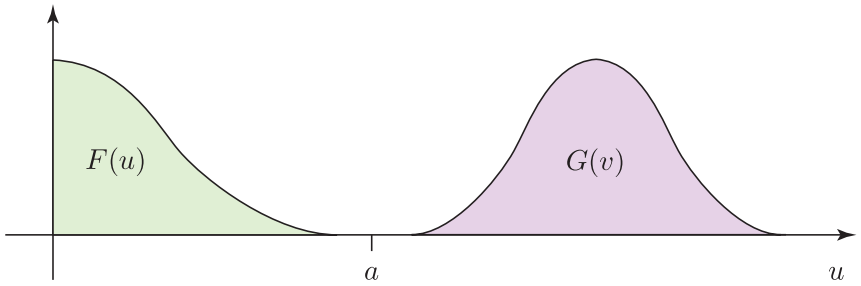


Figure 3.10: Bedrosian's theorem concerns functions separated completely in the Fourier domain, as indicated.

i.e. the 'low-pass' signal is factored out of the transform (see appendix A.3). This is very useful when analysing signals of the form

$$\begin{aligned} f(x) &= A(x)\cos(\omega x + \phi(x)) \\ &= \frac{1}{2}(A(x)e^{i\phi(x)}e^{i\omega x} + A(x)e^{-i\phi(x)}e^{-i\omega x}) \end{aligned}$$

where $A(x)$ and $\phi(x)$ are the amplitude and phase of the signal respectively. Here $f(x) = A(x)e^{\pm i\phi(x)}$ and $G(v) = \delta(v \pm \omega)$ so application of Bedrosian's identity requires that the spectra of A and ϕ have support $|u| < \omega$ – see Fig. 3.11.

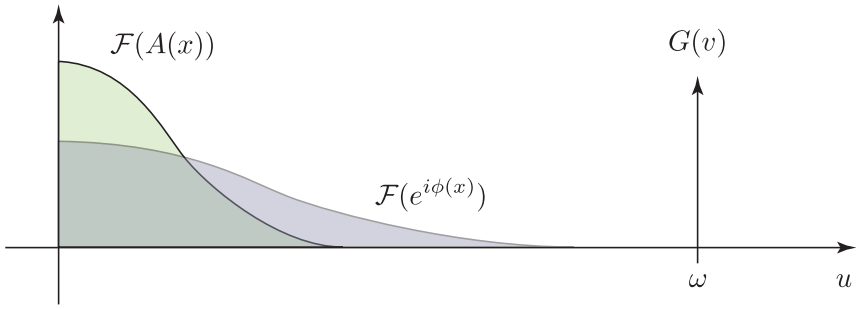


Figure 3.11: For accurate phase retrieval the spectra of $A(x)$ and $e^{i\Delta\phi(x)}$ should lie below ω in the Fourier domain.

To recover A and ϕ from a measured periodic signal, the first step is to sum the signal and its Hilbert transform

$$\begin{aligned} f(x) + i\mathcal{H}(f(x)) &= A(x)\cos(\omega x + \phi(x)) + iA(x)\mathcal{H}(\cos(\omega x + \phi(x))) \\ &= A(x)\cos(\omega x + \phi(x)) - iA(x)\sin(\omega x + \phi(x)) \\ &= A(x)e^{-i(\omega x + \phi(x))} \end{aligned}$$

where in the second line Bedrosian's identity was used to factor out $e^{i\phi(x)}$ terms from the Hilbert transform. The amplitude of the signal is then recovered as $\text{abs}(f(x) + i\mathcal{H}(f(x)))$, and the phase as $\arg(f(x) + i\mathcal{H}(f(x)))$. As an example consider the synthetic signal in Fig. 3.12 with a sinusoidal amplitude modulation and a Gaussian phase modulation. The width σ_ϕ of $\phi(x)$ is narrowed and its spectrum broadens. For large σ_ϕ the spectrum $\mathcal{F}(e^{i\phi(x)})$ lies below ω and both the amplitude and phase are retrieved correctly. As σ_ϕ decreases the spectrum begins to extend beyond ω and the retrievals drop in accuracy.

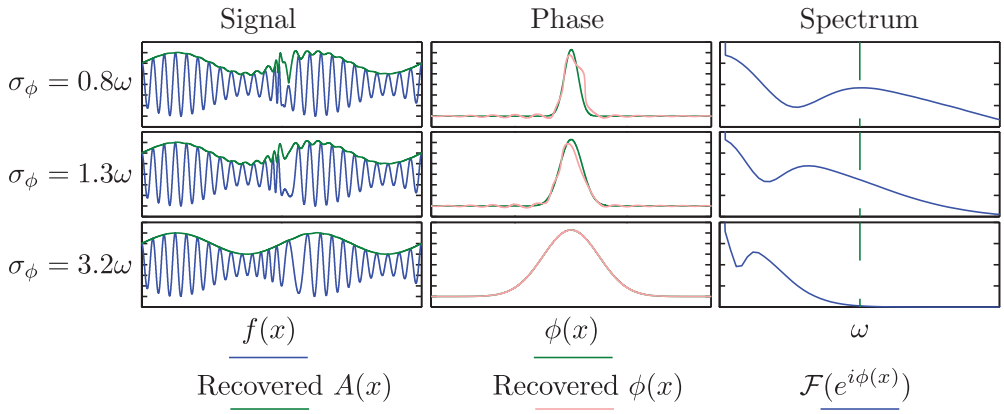


Figure 3.12: Attempts to recover $\phi(x) = e^{-x^2/2\sigma_\phi^2}$ fail when $\sigma_\phi < \omega$, i.e. the phase spatial scale length becomes shorter than a wavelength.

A particular 1D lineout of an interferogram may be interpreted as an oscillating signal suitable for analysis with a Hilbert transform[18]. Take for example the lineout indicated in Fig. 3.13a). The 'signal' is strictly positive but we require one centred at zero, so the moving average indicated in Fig. 3.13b) is subtracted. To help alleviate high-frequency noise the lineout is smoothed, at which point the amplitude and phase can be calculated. Repeating the process line-by-line generates a phase and amplitude map for the entire interferogram.

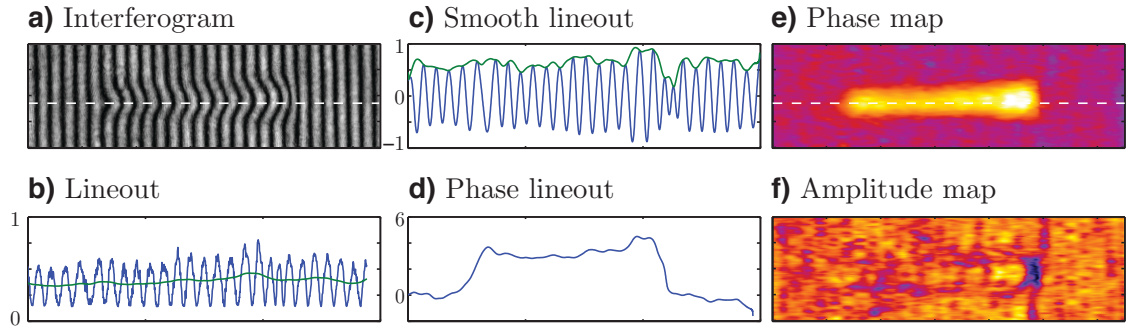


Figure 3.13: The interferogram in a) is analysed line-by-line. The lineout indicated by the dashed line is plotted and analysed in b), c), and d). e) and f) concatenate all lineouts into a 2D map.

Referring to Fig. 3.12, k_f takes the place of ω and it is clear the phase retrieval will fail for spatial gradients in $\phi < \lambda_f$, so again the fringe spacing is the limiting factor in spatial

[18] [238] K. Chauhan *et al. Measurement Science and Technology.* (2010).

resolution.

The phase resolution tends to be improved for the Hilbert method compared to the Fourier method, but as an intrinsically 1D process it can fail for interferograms with highly curved fringes. There is a third method which combines the phase sensitivity of the Hilbert transform with the 2D awareness of the Fourier transform.

Continuous wavelet transform

The continuous wavelet transform (CWT)[19] is a generalisation of the Fourier transform in that it acts to project a function $f(x)$ onto a set of labelled functions. In the Fourier case those functions are labelled by u as $\psi_u(x) = e^{iux}$, and the projection operator is the inner product over x. Labelling the projection of $f(x)$ on $\psi_u(x)$ as $F(u)$ then

$$F(u) = \int_{-\infty}^{\infty} f(x)\psi_u(x)\,\mathrm{d}x = \int_{-\infty}^{\infty} f(x)e^{iux}\,\mathrm{d}x \tag{3.23}$$

which is of course the normal Fourier transform. The fact that the $\psi_u(x)$ are mutually orthogonal introduces additional nice properties, but this need not be the case. In the CWT the labelling is extended to a two parameter set of functions $\psi_{ab}(x)$ called wavelets. These are derived from a 'mother' wavelet $\psi(x)$ by scaling and translation operations

$$\psi_{ab}(x) = \psi\left(\frac{x-b}{a}\right). \tag{3.24}$$

The CWT of $f(x)$ is then the two-dimensional function $F(a,b)$, often called the scalogram, generated by the projection onto the set of wavelets $\psi_{ab}(x)$

$$F(a,b) = \frac{1}{a}\int_{-\infty}^{\infty} f(x)\psi\left(\frac{x-b}{a}\right)\,\mathrm{d}x \tag{3.25}$$

where the normalisation has been introduced to ensure

$$\frac{1}{a}\int_{-\infty}^{\infty} \psi\left(\frac{x-b}{a}\right)\,\mathrm{d}x = \frac{1}{a}\int_{-\infty}^{\infty} \psi\left(y-\frac{b}{a}\right)a\,\mathrm{d}y = \int_{-\infty}^{\infty}\psi(y)\,\mathrm{d}y\,. \tag{3.26}$$

The choice of mother wavelet is dictated by the desired application and there are many possibilities. Choosing a wavelet localised in space allows the spatial dependence of $f(x)$ to be probed, down to a resolution dictated by the width of $\psi(x)$. The appropriate wavelet for space-frequency analysis is the Morlet wavelet, a Gaussian wave packet defined by a wavelength σ as

$$\psi_\sigma(x) \propto e^{-\frac{x^2}{2}}e^{i\sigma x}\,. \tag{3.27}$$

In the CWT this wavelet is translated across $f(x)$, and at every position scaled to different sizes. Varying σ alters the ratio of wavelength to wavepacket size, which allows probing at high spatial or high frequency resolution.

Interpreted in real space the labelling u of the Fourier transform corresponds to a frequency, and choosing the u which maximises the projection $F(u)$ is equivalent to finding the average frequency of $f(x)$. The Fourier basis functions are infinite in extent so all spatial

[19] [239] C. Heil *et al. SIAM Review.* (1989).

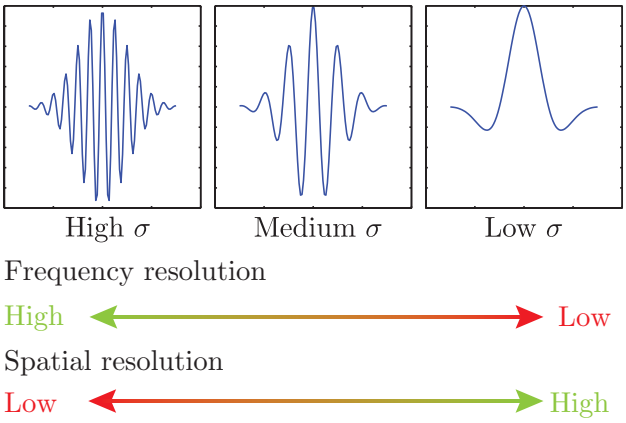


Figure 3.14: The real part of the Morlet wavelet $\Re(\psi_\sigma(x))$ for different σ

information is lost. Using the Morlet wavelet one should aim to maximise $F(a, b)$ over a at every position b, and this will recover the spatially-localised average frequency σ/a. Plotted in Fig. 3.15 are the scalograms of a function for different choices of σ. This is the approach followed by Tomassini *et al.*[20] for the analysis of laser-plasma interferograms (there the authors term the method 'ridge extraction'). At low σ the spatial distribtion of the frequency content of the function is well reproduced, but the retrieved wavelength isn't as accurate. As σ increases the finer spatial details are lost – see the small peaks at $x > 40$ – but the scalogram is maximised over a much narrower set of wavelengths, more accurately capturing the frequency content. For a given application it is important to choose the correct wavelet properties for accurate reconstruction.

To recover the signal in the form of $A(x)$ and $\phi(x)$ as before, the procedure is to take the value of maximum value of $F(a, b)$ over a at every b to form $F_{\max}(b) \equiv F(a_{\max}, b)$. Then $A(x) = \text{abs}(F_{\max}(b))$ and $\phi(x) = \arg(F_{\max}(b))$.

2D CWT

The 2D CWT is a straightforward extension of the 1D case, and is suited to the analysis of complex interferograms[21]. Rather than projecting $f(x)$ onto the wavelets $\psi_{ab}(x)$, the 2D function $f(x, y)$ is projected onto the set $\psi_{\mathbf{ab}}(x, y)$ where $\mathbf{a}$ and $\mathbf{b}$ are 2D vectors. In practice the scaling vector $\mathbf{a}$ is most conveniently treated as a linear scaling a along direction θ such that $\mathbf{a} = (a\cos\theta, a\sin\theta)$. The 2D CWT then takes the inner product of $f(x, y)$ and the mother wavelet over all translations, rotations and wavelengths. Mathematically, if R_θ represents the 2D rotation matrix, then

$$F(a, \theta, \mathbf{b}) = \frac{1}{a^2}\int_{-\infty}^{\infty} f(\mathbf{x})\psi_\sigma\left(\frac{\mathsf{R}_\theta(\mathbf{x} - \mathbf{b})}{a}\right) \mathrm{d}\mathbf{x}\,. \tag{3.28}$$

Formally this approach would require a minimisation over a four-dimensional parameter space, which could become time consuming. Fortunately the two integrals over $\mathbf{x}$ can be cast as a fast convolution as follows. Following Ma *et al.*[22], the Morlet wavelet is written

[20] [240] P Tomassini *et al. Applied Optics.* (2001).
[21] [241] J. Ma *et al. Applied Optics.* (2011).
[22] [242] H. Ma. *Optical Engineering.* (2006).

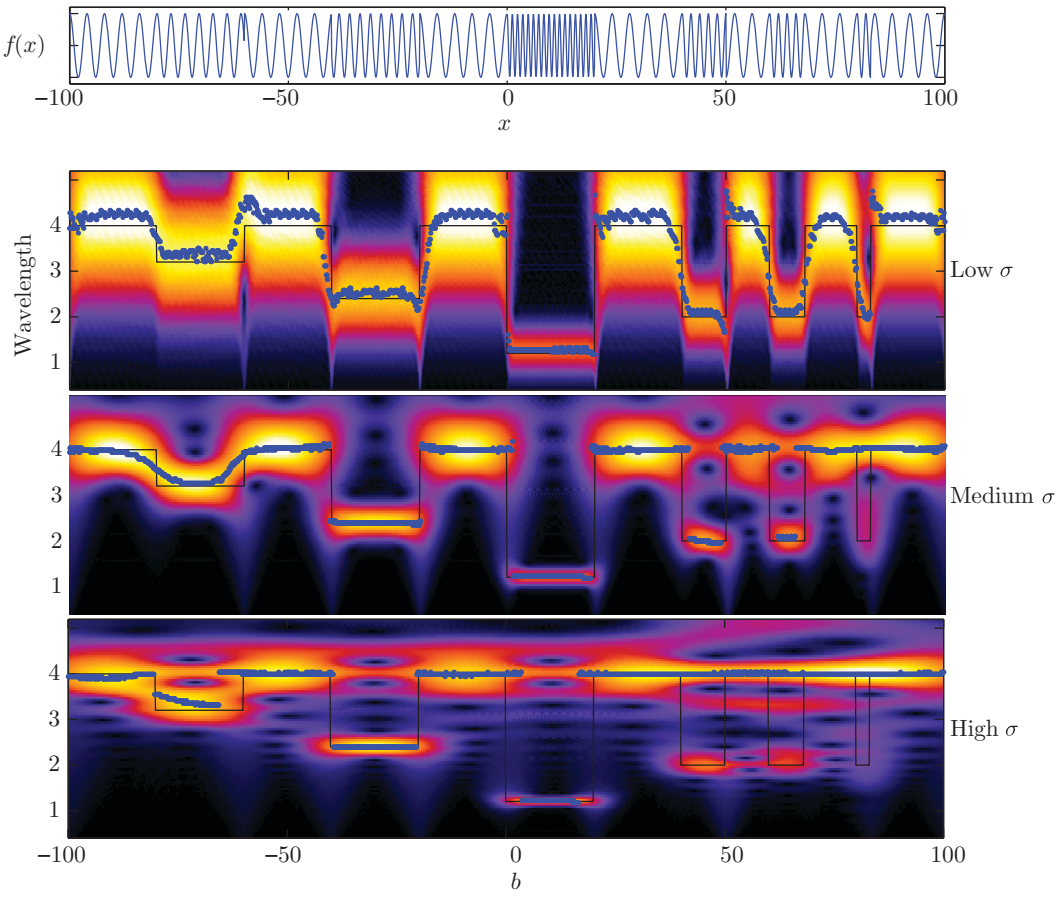


Figure 3.15: $F(a,b)$ - the scalogram - of the plotted function for various σ. For each position x there is a wavelength at which $F(a,b)$ is maximised, marked by a dot. The black line marks the actual wavelength.

in full as

$$\begin{aligned} \psi_\sigma\left(\frac{\mathsf{R}_\theta(\mathbf{x}-\mathbf{b})}{a}\right) &= e^{-\frac{(x-b_x)^2+(y-b_y)^2}{2a^2}} e^{\frac{i\sigma((x-b_x)\cos\theta+(y-b_y)\sin\theta)}{a}} \\ &\equiv \zeta_{a\theta}(\mathbf{b}-\mathbf{x}) \\ \text{where} \quad \zeta_{a\theta}(\mathbf{x}) &= e^{-\frac{x^2+y^2}{2a^2}} e^{\frac{i\sigma(x\cos\theta+y\sin\theta)}{a}}, \end{aligned}$$

then

$$F(a,\theta,\mathbf{b}) = \frac{1}{a^2}\int_{-\infty}^{\infty} f(\mathbf{x})\zeta_{a\theta}(\mathbf{b}-\mathbf{x})\,\mathrm{d}\mathbf{x} = f(\mathbf{b}) * \zeta_{a\theta}(\mathbf{b}) = \mathcal{F}^{-1}(\mathcal{F}(f)\mathcal{F}(\zeta))\,. \tag{3.29}$$

Finally to recover the amplitude and phase of $f(\mathbf{x})$ as before it remains to maximise over a and θ. If $f(\mathbf{x})$ is known to only contain a small range of wavelengths and angles, then computationally a and θ can be represented by $\approx$ 10-element vectors, and the whole process requires $\approx 10^2$ convolutions. Storing the values of $a_{\max}(\mathbf{b})$ and $\theta_{\max}(\mathbf{b})$ which maximise $F_{\max}(\mathbf{b}) \equiv F(a_{\max}(\mathbf{b}), \theta_{\max}(\mathbf{b}), \mathbf{b})$ also allows the calculation of local fringe wavelength and angle respectively.

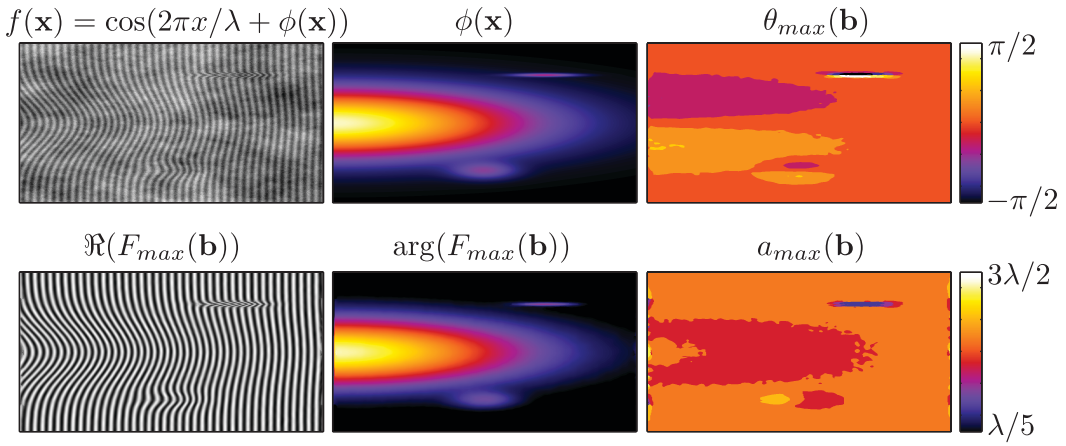


Figure 3.16: A synthetic noisy interferogram compared to a 2D CWT phase retrieval. Also plotted are the best-fit local angle and wavelength of the fringes.

Plotted in Fig. 3.16 is a synthetic noisy interferogram which represents $f(\mathbf{x})$. A phase profile has been imposed which mimics some features of interferograms of LFWA plasmas, namely a large smooth phase profile overlaid with small filamentary structure. Maximising the 2D CWT over a and θ as above produces complex $F_{\max}$, the real value of which represents the reconstructed interferogram. As plotted this is remarkably smooth and captures all of the detail of the synthetic phase profile. $\arg(F_{\max})$ is the recovered phase of the interferogram, and again matches the synthetic profile very well. The fact that the fitted angle and local wavelength of the fringes look reasonable means the correct set of wavelets was chosen to represent the interferogram, generating a good reconstruction.
A comparison of the different phase retrieval methods discussed here is plotted in Fig 3.17. The interferogram is experimental data, and was chosen to be particularly difficult to analyse. Not only is the maximum phase shift high at > 5 fringes, the plasma density profile is filamentary and asymmetric. There are also regions where the fringes are invisible behind bright self-emission. Nevertheless in this case the phase is recovered well by all methods, barring regions where the phase unwrapping has failed. The CWT retrieval is particularly high-resolution and artefact-free, though takes slightly longer than the other methods.

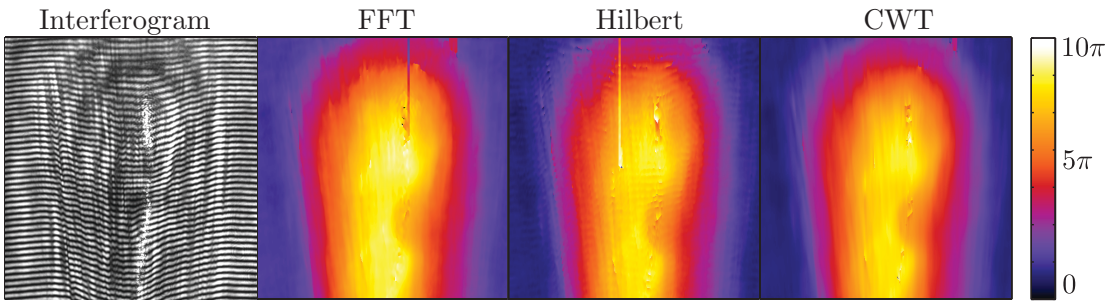


Figure 3.17: An experimental interferogram alongside the phase as retrieved by the methods described above.

3.3.2 Phase unwrapping

All phase retrieval methods rely on the $\arg(z)$ function and so can only return a phase measurement within its range, $(-\pi, \pi)$. Where an interferogram shifts by more than a fringe with respect to a reference, the retreived phase will be wrapped unphysically modulo 2π. Physical phase shifts can easily exceed this, as in Fig 3.17, so this is a common problem. The task of phase unwrapping is then to extract the physical phase shift $\Delta\phi$ from the measured wrapped phase shift $\Delta\phi_W$. In 1D the unwrapping algorithm is simple, suppose the phase is represented as a vector with elements ϕ_1 to ϕ_N.

1. Choose a position ϕ_i to start unwrapping where the phase is known not to suffer from wrapping errors, and a threshold phase jump $\Delta\phi_t$ above which the phase is considered wrapped.
2. For all $j > i$, if $|\phi_{j+1} - \phi_j| > \Delta\phi_t$ then $\phi_{j+1} \dots \phi_N \rightarrow \phi_{j+1} \dots \phi_N - (\phi_{j+1} - \phi_j)$
3. For all $j < i$, if $|\phi_{j-1} - \phi_j| > \Delta\phi_t$ then $\phi_1 \dots \phi_{j-1} \rightarrow \phi_1 \dots \phi_{j-1} - (\phi_{j-1} - \phi_j)$

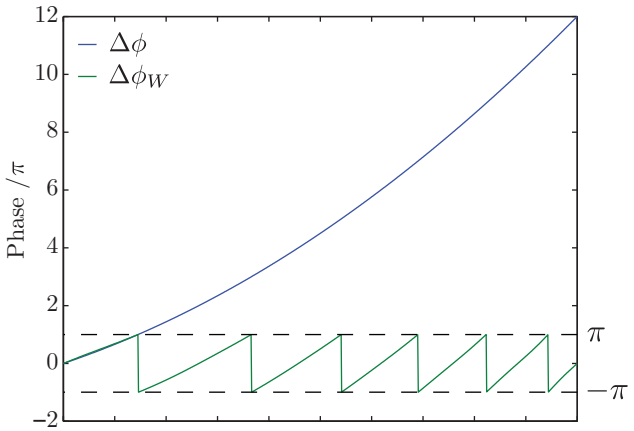


Figure 3.18: A synthetic 1D measured phase $\Delta\phi_W$ compared to the physical unwrapped phase $\Delta\phi$.

The algorithm then sweeps outwards from a known good starting position adding unphysical phase jumps back to the measured phases. In a 2D image the problem is more complex – starting from a known good position in which direction should the algorithm propagate? In some simple cases it is possible to apply the 1D algorithm line-by-line, but this is not always successful. With the presence of noise it is possible that phase ambiguities can arise, and the unwrapping becomes path-dependent. There is only one 'correct' solution, and several algorithms exist to find it.

Goldstein unwrapping

Goldstein *et al.*[23] developed what is now commonly known as the Goldstein unwrapping algorithm in the context of radar interferometry, where phase shifts are proportional to elevation. For a smooth phase map the property should hold that the phase residue

[23] [243] R. M. Goldstein *et al. Radio Science.* (1988).

integrated around any set of pixels should be zero – see Fig. 3.19. The residue at pixel (i, j) is the sum of the 4 phase differences

$$\begin{aligned}
\Delta_1 &= \phi_{i+1,j} - \phi_{i,j} \\
\Delta_2 &= \phi_{i+1,j+1} - \phi_{i+1,j} \\
\Delta_3 &= \phi_{i,j+1} - \phi_{i+1,j+1} \\
\Delta_4 &= \phi_{i,j} - \phi_{i,j+1} \\
\mathrm{Res}_{i,j} &= \Delta_1 + \Delta_2 + \Delta_3 + \Delta_4.
\end{aligned}$$

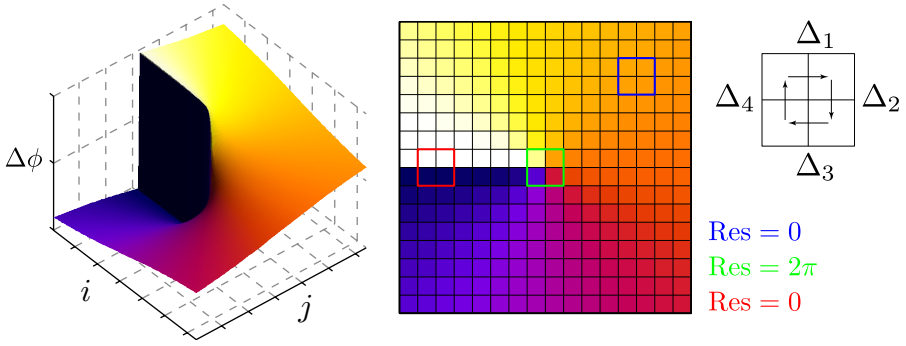


Figure 3.19: An example phase map with a single nonzero residue at the centre. A branch cut must be introduced connected to the residue - here it is directed to the left.

The terminology here is borrowed from complex analysis as the 'smooth phase' requirement is the same as the one for a complex function to be holomorphic – any loop integral on the complex plane is zero. As is well known from complex analysis, any loop integral around a residue (a simple pole of order 1) is non-zero. The existence of phase ambiguities in a phase map is then equivalent to the existence of loop integration ambiguities of a complex function with poles. The solution from complex analysis is to introduce branch cuts into the complex plane which define the principle branch of a function. As long as any integration path does not cross a branch cut the result is well defined and path independent. The Goldstein unwrapping algorithm then proceeds as follows:

1. Calculate residues clockwise for each pixel.
2. Classify nonzero residues as positive or negative depending on the sign of the residue.
3. Connect positive and negative residues with a branch cut. This object is now 'neutral' overall.
4. Starting from a position of known good phase, spiral outwards in a path which avoids all branch cuts.
5. If phase jumps are encountered remove as in the 1D case.

A synthetic wrapped phase map with a single negative and positive residue is plotted in Fig. 3.20. The Goldstein unwrapping algorithm inserts a single branch cut between the

two residues and returns a perfectly unwrapped output. The 1D algorithm produces severe artefacts in the region of the branch cut which persist if the unwrapping is performed from any direction. Also plotted in Fig. 3.20 is an example experimental wrapped phase map. The 1D unwrap is performed horizontally where it causes artefacts at a set of non-zero residues. In this region the Goldstein algorithm 'gives up' in that the branch cuts form a small loop. Inside this region it is impossible to unwrap the phase deterministically so it is sealed off from the rest of the phase map.

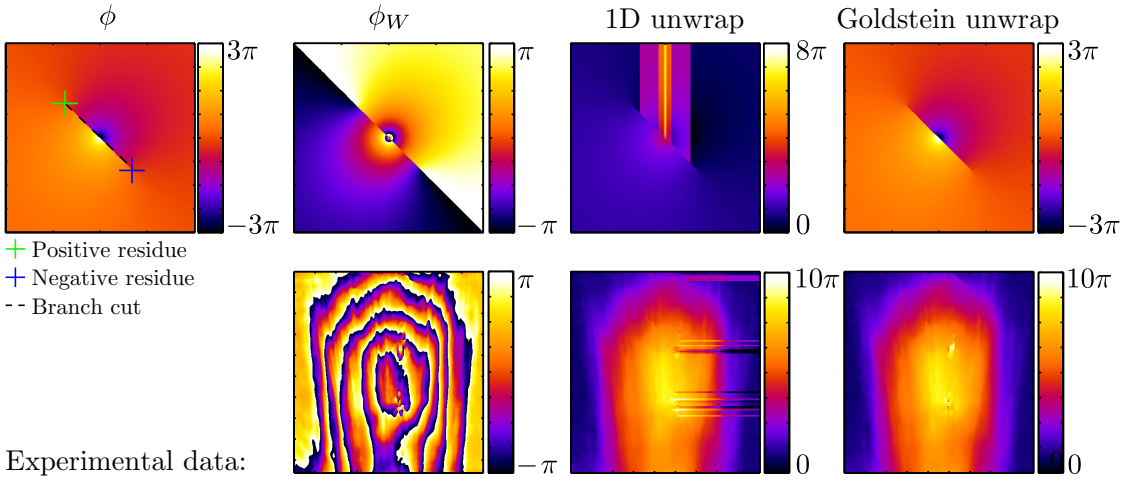


Figure 3.20: Here $\phi = 2\csc(e^{i\pi/4}z)$ for $z = x + iy$, where (x, y) are the coordinates of the image. Experimental data is taken from a LWFA plasma.

Fourier unwrapping

The Goldstein unwrap is effective, but can become very slow for large images. A faster method is proposed by Volkov *et al.*[24] based on Fourier transforms. The first key idea is to represent the wrapped phase shift as two components

$$\phi_W(x, y) = \phi(x, y) - 2\pi N(x, y), \quad N \in \mathbb{Z}\,. \tag{3.30}$$

The action of phase wrapping introduces a difference between the retrieved and actual phases $2\pi N(x, y)$, defined to ensure N is integer-valued. If $N(x, y)$ can be calculated, it can be added to the wrapped phase to remove the phase jumps. The next idea is to work with an auxiliary function that can be calculated from the measured phase shift

$$Z(x, y) = e^{i\phi_W} = e^{i\phi}e^{-2\pi iN} = e^{i\phi}. \tag{3.31}$$

Taking the 2D gradient of Eq. 3.30 and substituting for ϕ from Eq. 3.31 gives

$$\nabla N = \frac{1}{2\pi}(\nabla\phi - \nabla\phi_W) = \frac{1}{2\pi}\left(\Re\left(\frac{\nabla Z}{iZ}\right) - \nabla\phi_W\right) \tag{3.32}$$

which is in terms of measured quantities. The second step is to reconstruct N from ∇N using the Fourier identity

[24] [244] V. V. Volkov *et al. Optics Letters.* (2003).

$$N = \Re\left(\frac{1}{2\pi i}\mathcal{F}^{-1}\left(\frac{\mathcal{F}(\frac{\partial N}{\partial x})u + \mathcal{F}(\frac{\partial N}{\partial y})v}{u^2+v^2}\right)\right) \tag{3.33}$$

where as usual (u, v) are the conjugate variables to (x, y). The identity is exact up to a constant, so for effective reconstruction of N there should be a region of the phase map where it is known that no wrapping has occurred. N can then be offset to enforce $N = 0$ at that point. Also as defined N should only take integer values, so spatial smoothing from the $(u^2 + v^2)^{-1}$ term can be suppressed by rounding N to the nearest integer.

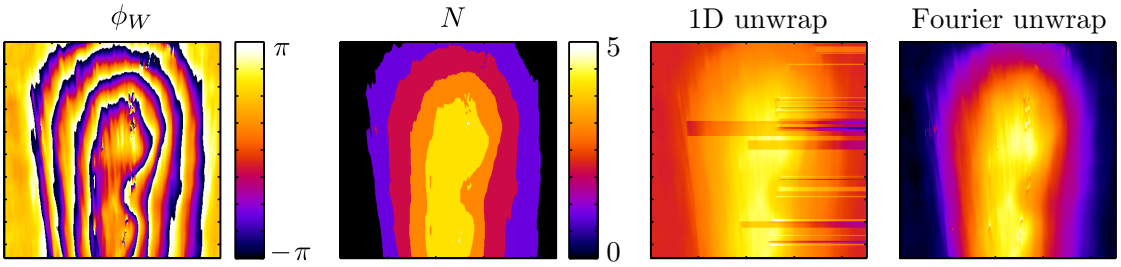


Figure 3.21: The Fourier unwrap is compared to a 1D unwrap for a wrapped experimentally recorded phase map. $N(x, y)$ is recovered with high fidelity, facilitating an effective unwrap.

3.3.3 Density retrieval

Referring back to Eq. 3.12 the recorded phase shift is proportional to the line-integrated electron density, assuming the plasma is very underdense and static over the transit timescale. The field of tomography covers the reconstruction of structures from projected data, and in general requires many projections to be recorded at different angles. This is not usually possible when probing LWFA plasmas, but can be applied to the characterisation of neutral gas jets[25]. Fortunately many laser produced plasmas exhibit axial symmetry, which in principle allows the full 3D structure to be recovered from a 2D projection.

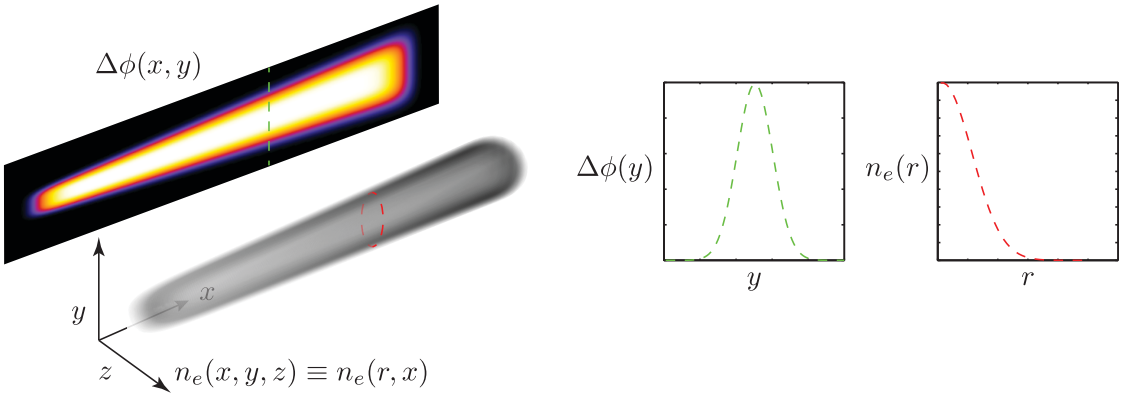


Figure 3.22: For each position x along the plasma, the phase shift $\Delta\phi(y)$ is proportional to the Abel transform of $n_e(r)$.

The projection of a radially symmetric function $f(r)$ is its Abel transform, denoted

[25] [245] G. Golovin *et al. Applied Optics.* (2015).

$$f_A(y) = 2\int_y^R \frac{r f(r)}{\sqrt{r^2 - y^2}}\,\mathrm{d}r\,. \tag{3.34}$$

The problem is then to recover $f(r)$ given knowledge of $f_A(y)$, and the required integral transform is the Abel inversion.

Abel inversion

The Abel inversion may be derived with Fourier techniques, following Bracewell [26]. First introduce new variables ρ and ξ such that

$$\rho = r^2 \quad \xi = y^2 \quad 2r\,\mathrm{d}r = \mathrm{d}\rho \quad f_A(y) \equiv \tilde{f}_A(\xi) \quad f(r) \equiv \tilde{f}(\rho)\,. \tag{3.35}$$

With these definitions the Abel transform is expressible as the convolution

$$\tilde{f}_A(\xi) = \int_0^\infty k(\xi - \rho)\tilde{f}(\rho)\,\mathrm{d}\rho\,, \quad \tilde{f}_A = k * \tilde{f} \tag{3.36}$$

for the kernel

$$k(x) = \begin{cases} \frac{1}{\sqrt{-x}} & x < 0 \\ 0 & x > 0 \end{cases} \tag{3.37}$$

Letting capital letters again denote the Fourier transform, we have

$$\widetilde{F_A} = K\widetilde{F} \Rightarrow \widetilde{F} = \frac{\widetilde{F_A}}{K} \tag{3.38}$$

so the inversion requires the Fourier transform of k

$$K(u) = \sqrt{\frac{\pi}{iu}}\,. \tag{3.39}$$

Substituting above, we get

$$\widetilde{F} = \sqrt{\frac{iu}{\pi}}\widetilde{F_A} = \frac{iu}{\pi}\widetilde{F_A}K\,. \tag{3.40}$$

The multiplication by u represents differentiation in real space, so performing the inverse Fourier transform and rewriting as a convolution integral we arrive at the Abel inversion formula

$$f(r) = \frac{1}{\pi}\int_r^\infty \frac{f_A'(y)}{\sqrt{y^2 - r^2}}\mathrm{d}y. \tag{3.41}$$

Written in terms of physical quantities, the electron density corresponding to a phase shift $\Delta\phi(x, y)$ is given by

$$n_e(x, r) = \frac{\lambda_0 e^2}{4\pi^2 m_e \epsilon_0 c^2}\int_r^\infty \frac{d\Delta\phi(x, y)}{dy}\frac{1}{\sqrt{y^2 - r^2}}\mathrm{d}y\,. \tag{3.42}$$

[26] [246] R. N. Bracewell. *The Fourier Transform and its Applications*. (1986).

While it is possible to numerically integrate this expression to any desired level of accuracy, it can be slow for large interferograms. The required differentiation of the phase information amplifies any noise present, and the denominator becomes small at the lower limit of the integrand. While both problems may be overcome with sufficiently careful treatment of the integral, other methods based on the Fourier-Hankel transform or algebraic reconstruction technique have been developed to cope with these issues. The existence of a radially symmetric electron density distribution also relies on symmetry of the phase distribution. No real phase data will be symmetric so an axis of symmetry must be added, either manually or with a fitting method. Some Abel inversion techniques can discriminate between symmetric and asymmetric parts of a function, keeping only the useful symmetric part.

Fourier methods

Kalal *et al.*[27] show that it is possible to speed up the calculation of the Abel inversion by expressing the solution in terms of a set of basis functions. To begin, assume now that that the projected density profile $f_A(y)$ is related to the radial density profile $f(r)$ by

$$f_A(y) = 2\int_y^R \frac{r f(r)}{\sqrt{r^2 - y^2}}\,\mathrm{d}r \tag{3.43}$$

where R is the radial extent of $f(r)$, and that the projection has the cosine expansion

$$f_A(y) = a_0 + \sum_{k=1}^{\infty} a_k \cos\left(\frac{k\pi y}{R}\right). \tag{3.44}$$

Expressing $f_A(y)$ in this way ensures it is symmetric and smooth at $x = 0$, even if the original data is noisy. The coefficients a_k may be rapidly calculated with a discrete cosine transform or similar Fourier methods. Substitution of this expression into the inverse Abel transform gives

$$f(r) = \frac{\pi}{2R}\sum_{k=1}^{\infty} k a_k g_k\left(\frac{r}{R}\right) \tag{3.45}$$

for the functions

$$g_k(\rho) = \frac{2}{\pi}\int_0^{\sqrt{1-\rho^2}} \frac{\sin\left(k\pi\sqrt{t^2+\rho^2}\right)}{\sqrt{t^2+\rho^2}}\,\mathrm{d}t. \tag{3.46}$$

The functions $g_k(\rho)$ are effectively a basis set for $f(r)$ corresponding to the inversion of $\cos(x)$. They share the nice property that $g_k'(0) = 0$, required for any function symmetric about 0. To avoid numerical errors when $\rho \to 0$ or $\rho \to 1$, it may be noted that $g_k(0) = \frac{2}{\pi}\,\mathrm{Si}(\pi k)$ and $g_k(1) = 0$. For any $f(y)$, $g_k(\rho)$ is only required over $0 < \rho < 1$ so these functions can be precomputed and stored in a lookup table. Operating in this fashion, the Abel inversion is sped up by a factor of 100 or more compared to direct numerical integration. Additionally, if high-frequency noise is a problem in the original

[27] [247] M. Kalal *et al. Applied Optics.* (1988).

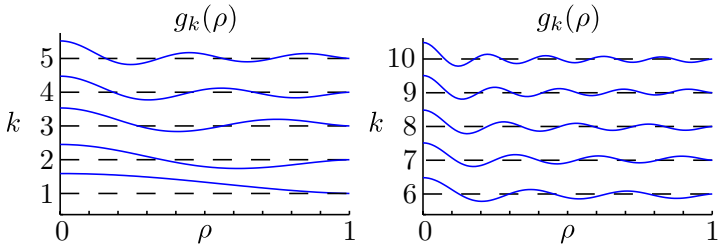


Figure 3.23: The first 10 basis functions $g_k(\rho)$.

projection then smoothing is achieved by limiting the number of a_k to include in the sum, as plotted in Fig. 3.24. Noise removal is then accompanied by a reduction in calculation time and avoids the requirement for additional preprocessing steps.

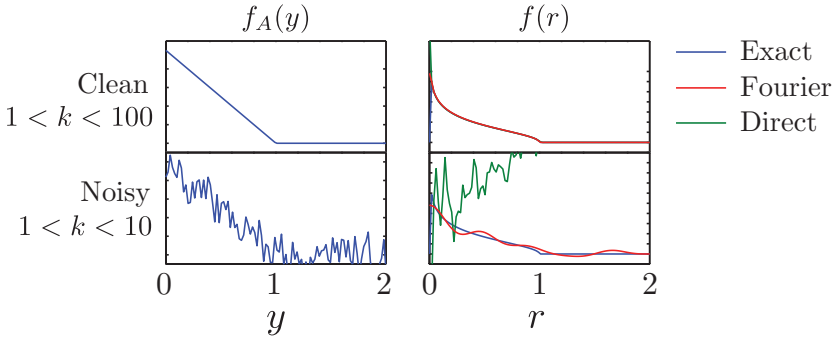
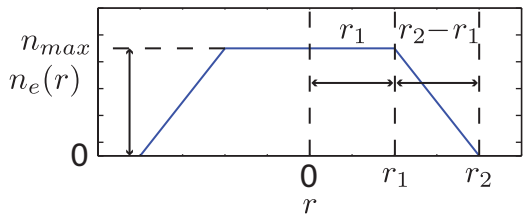


Figure 3.24: A comparison of Abel inversion by direct numerical integration and with the method described here.

Fitting methods

If some properties of $n_e(r)$ are known in advance, the simplest procedure may be to assume an analytic form for $n_e(r)$ and its Abel transform. A fit can then be performed to the phase profile, avoiding the need for a direct inversion. An example is the gas density distribution in a jet produced by a supersonic nozzle. Here the density profile is known to be approximately trapezoidal, as plotted in Fig 3.25. At the centre of the jet the density reaches a maximum $n_{\max}$, which remains approximately constant over a radius r_1. The density then falls approximately linearly to zero at a distance r_2.



$$n_e(r) = \begin{cases} n_{max} & r < r_1 \\ n_{max}\left(\frac{r_2 - r}{r_2 - r_1}\right) & r_1 < r < r_2 \end{cases}$$

Figure 3.25: The approximate shape of the density profile of a supersonic gas jet, parametrised with three variables.

The Abel transform of $n_e(r)$, and therefore the shape of the phase profile, is

$$\Delta\phi(y) = \frac{2\Delta\phi_{\max}}{r_1 + r_2} \begin{cases} f_1(r) & r < r_1 \\ \\ f_2(r) & r_1 < r < r_2 \end{cases} \tag{3.47}$$

where

$$\begin{aligned} f_1 = \quad & \sqrt{r_1^2 - y^2} + \frac{1}{r_2 - r_1}\left(r_2\sqrt{r_2^2 - y^2} - r_2\sqrt{r_1^2 - y^2}\right. \\ & \left. - \frac{1}{2}\left(r_2\sqrt{r_2^2 - y^2} - r_1\sqrt{r_1^2 - y^2} + y^2 \ln\left(\frac{r_2 + \sqrt{r_2^2 - y^2}}{r_1 + \sqrt{r_1^2 - y^2}}\right)\right)\right) \end{aligned} \tag{3.48}$$

$$f_2 = \frac{1}{r_2 - r_1}\left(r_2\sqrt{r_2^2 - y^2} - \frac{1}{2}\left(r_2\sqrt{r_2^2 - y^2} + y^2 \ln\left(\frac{r_2 + \sqrt{r_2^2 - y^2}}{y}\right)\right)\right) \tag{3.49}$$

and $\Delta\phi_{\max}$ is the maximum phase shift at the symmetry axis. For a neutral gas of refractive index $n = 1 + k$, $k \ll 1$, it can be shown that the peak molecular density is related to the phase profile by

$$n_{\max} = \frac{\Delta\phi_{\max} n_{\text{atm}} \lambda}{2\pi k(r_1 + r_2)} \tag{3.50}$$

where n_{atm} is atmospheric density at standard temperature and pressure, and λ is the wavelength of the probing laser. As plotted in Fig. 3.26 for a suitable density profile the fitting method works well and is less sensitive to noise.

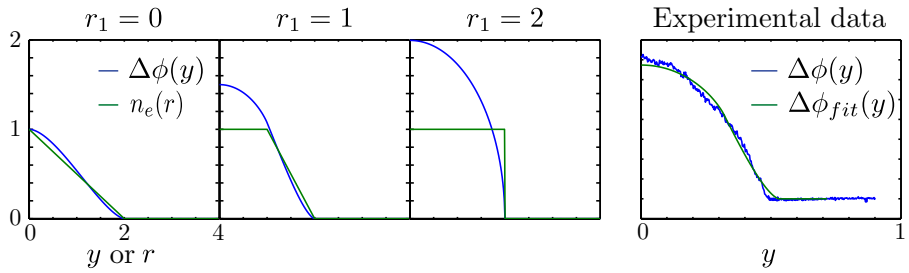


Figure 3.26: Synthetic density profiles with $n_{\max} = 1$, $r_2 = 2$, and phase profiles given by Eq. 3.47. Experimental data is taken from a 10 mm supersonic Argon gas nozzle.

3.4 Electron spectrometer

The electron spectrometer is a very important diagnostic for laser-wakefield accelerators, both in the commissioning of a new experiment and the characterisation of the performance of an accelerator under some given conditions. As the end product of the non-linear processes of laser pulse evolution and high-amplitude wake generation, the electron beam is sensitive to any problems in the accelerator. To detect electrons a scintillating Lanex screen is often used[28], consisting of a plastic substrate coated with a thin layer of the

[28] [248] J. Lindström. *Radioluminescence: A simple model for fluorescent layers.* (2011).

scintillator Gd_2O_2S:Tb. This scintillator emits photons at a range of wavelengths across the optical band, with a strong peak at 545 nm. The decay lifetime is 600 µs which allows the scintillator to be used at a high repetition rate. Importantly the energy deposited in the scintillator per electron remains constant[29] for electrons of energy larger than a few MeV. As long as these low energy electrons are blocked before they reach the Lanex screen, the emitted light at any point of the screen is then proportional to the total charge passing through it. If the light is additionally filtered around a narrow band near 545 nm, the wavelength-dependent quantum efficiency of the recording CCD may be neglected and the recorded charge is also proportional to the number of CCD counts.

3.4.1 Dipole spectrometer

In the experiments described in this thesis, all electron spectrometers were composed of one or more permanent dipole magnets and one or more Lanex screens. To assess the effectiveness of a real spectrometer, consider an idealised version as depicted in Fig 3.27. An electron travels along the z-axis with velocity and Lorentz factor v and γ respectively. It enters a magnet represented by a rectangular region of uniform magnetic field $\mathbf{B} = B\hat{\mathbf{x}}$ of length L_m and height H_m. The electron leaves the magnet at angle θ and passes through a detector plane, itself angled at α and a distance L_s from the end of the magnet. This intersection position occurs at a distance l from the optical axis as measured along the detector. The ability of the spectrometer to resolve electron energies accurately is determined by the way the dispersion position l depends on the electron energy E.

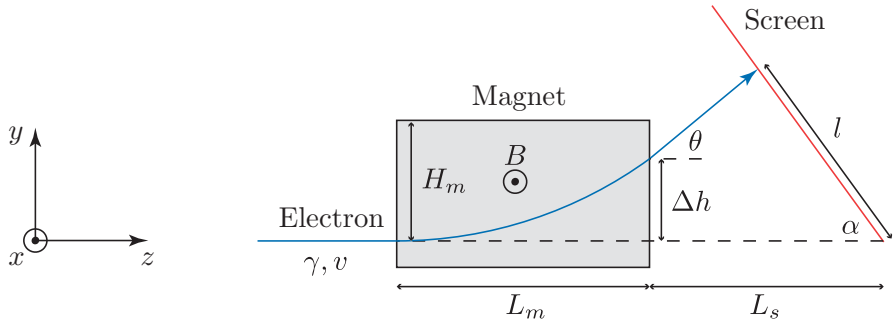


Figure 3.27: Permanent dipole magnet electron spectrometer, with relevant quantities indicated.

While inside the magnet the position of the electron is determined by the Lorentz force equation

$$\frac{\mathrm{d}\gamma m_e \mathbf{v}}{\mathrm{d}t} = q_e \mathbf{v} \times \mathbf{B}. \tag{3.51}$$

The magnetic field does no work on the electron so γ is constant. Expanding in Cartesian components then

[29] [249] Y. Glinec *et al. Review of Scientific Instruments.* (2006).

$$\begin{aligned}\frac{\mathrm{d}v_x}{\mathrm{d}t} &= 0\\ \frac{\mathrm{d}v_y}{\mathrm{d}t} &= -\frac{Bq_e}{\gamma m_e}v_z\\ \frac{\mathrm{d}v_z}{\mathrm{d}t} &= \frac{Bq_e}{\gamma m_e}v_y.\end{aligned}$$

If the electron enters the magnet with any velocity drift along the field direction, this component is unchanged by the field. The other two components follow the simple harmonic motion $\ddot{v}_i + \omega_B^2 v_i = 0$ for the cyclotron frequency $\omega_B = Bq_e/\gamma m_e$. Integrating a second time it can be shown that the path of the electron in the (y, z) plane is a circle of radius $R_B = v/\omega_B$. By simple geometry the exit angle of the electron is given by $\sin\theta = L_m/R_B$, and the change in height $\Delta h = R_B(1 - \cos\theta)$. This assumes the electron does not leave the top of the magnet, which is an implicit assumption that the electron energy is high. The intersection point on the detector is then

$$l = \frac{\sec\alpha}{\tan\alpha + \tan\theta}(L_s\tan\theta + \Delta h)\,. \tag{3.52}$$

To calculate the dispersive properties of the spectrometer it remains to find l as a function of electron energy E. Taking $E = (\gamma - 1)m_e c^2$ and assuming γ large, the product

$$\gamma v = c\sqrt{\frac{E^2}{m_e^2c^4} + \frac{2E}{m_e c^2}} \approx \frac{E}{m_e c}. \tag{3.53}$$

and $\sin\theta \approx L_m B q_e c/E$. For $L_m = 10$ cm and $B = 1$ T, $L_m B q_e c \approx 30$ MeV so for most LWFA electrons with energies much greater than this $\sin\theta \approx \theta$. Continuing with these approximate expressions, the dispersion of the spectrometer is found to be

$$\frac{dl}{dE} = -\frac{\sec\alpha\tan\alpha(L_s + L_m)L_m B q_e c}{(E\tan\alpha + L_m B q_e c)^2}. \tag{3.54}$$

For best energy discrimination this quantity should be as large as possible – increasing B, L_m or L_s is an obvious way to increase the dispersion. For high energies the dispersion scales as E^{-2} which implies the energy resolution worsens rapidly. However below $E = L_m B q_e c/\tan\alpha$ the dispersion is nearly constant, which means that tilting the detector plane towards the optical axis will improve energy resolution. Unfortunately the angles required for constant dispersion over 1 GeV are so small (2°) that the detector would need to be very large.

For the complete description of a rectangular magnet it is also possible to calculate l when the electron leaves through the top of the magnet. Making no high-energy approximations where now $\cos\theta = 1 - H_m/R_B$, the expression is

$$l = \frac{\sec\alpha}{\tan\alpha + \tan\theta}\left(L_s\tan\theta + \frac{1}{R_B/H_m - 1}\left(L_m\sqrt{\frac{2R_B}{H_m} - 1} - R_B\right)\right)\,. \tag{3.55}$$

and the dispersion scales more weakly as $E^{-1/2}$.

3.4.2 Image warping

When imaging a Lanex screen a 2D projection is recorded onto the camera CCD of an approximately 2D object embedded in 3D space. Due to experimental constraints it is not always possible to image the screen normally, and so the recorded image can appear distorted. It is difficult to extract the electron energy spectrum in this case, so here the procedure for 'undoing' this projection is detailed.
If the CCD co-ordinates are denoted (u, v), and the actual screen co-ordinates are (x, y), then the aim is to find the inverse of the 'imaging' transformation from (x, y) to (u, v), as depicted in Fig. 3.28.

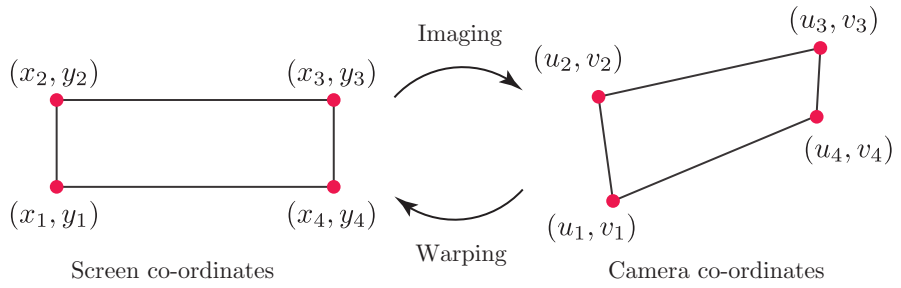


Figure 3.28: The co-ordinate systems used in screen-space and image-space.

The screen width and height are known in (x, y), and thus so are the positions of the screen corners (x_i, y_i). The challenge is then to construct an inverse transformation from (u, v) back to (x, y) given the measured positions of the corners in (u, v). This transformation may then be applied to the entire image to recover a 'flattened' version of the Lanex screen, which is much more amenable to analysis.

Projective transformation

The image of a 2D object embedded in 3D space is related to the real shape of the object by a projective transformation. A projective transform belongs to the class of affine transformations, which preserve straight lines and connectivity – a quadrilateral is always mapped to another quadrilateral. A linear transformation admits the general matrix form

$$\begin{pmatrix} x \\ y \end{pmatrix} = \begin{pmatrix} a & b \\ c & d \end{pmatrix} \begin{pmatrix} u \\ v \end{pmatrix} \tag{3.56}$$

where the 4-element matrix encodes a rotation (1 parameter), a scaling (1 parameter) and a 2D shear along a vector (2 parameters). This is not general enough for a projective transformation, and precludes the possibility of incorporating translations. Linear transformations are therefore affine, but affine transformations are not necessarily linear. The solution is to add a point 'at infinity', which represents the vanishing point of an image. The matrix formalism can be retained with the augmented matrix notation

$$\begin{pmatrix} x \\ y \\ 1 \end{pmatrix} = \lambda \begin{pmatrix} a & b & c \\ d & e & f \\ g & h & 1 \end{pmatrix} \begin{pmatrix} u \\ v \\ 1 \end{pmatrix}. \tag{3.57}$$

Represented in this way, compositons of affine transformations become matrix multiplications, as is the case for linear transformations. By convention the final element of the transformation matrix is defined as 1, with the degree of freedom shifted to a scaling parameter λ. With this definition, from the third line of Eq. 3.57

$$\lambda = \frac{1}{gu + hv + 1} \tag{3.58}$$

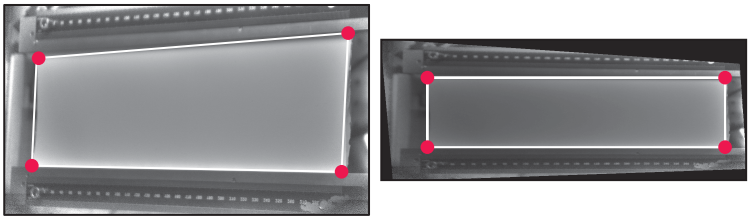
and so the transformation becomes

$$\begin{aligned} x(u,v) &= \lambda(au + bv + c) = \frac{au + bv + c}{gu + hv + 1} \\ y(u,v) &= \lambda(du + ev + f) = \frac{du + ev + f}{gu + hv + 1}\,. \end{aligned} \tag{3.59}$$

The task of finding the inverse transformation reduces to the calculation of the 8 parameters $\{a \ldots h\}$. This requires knowledge of at least 4 points in each image which are known to map to one another under the transformation. From the camera image $\{u_i\}$ and $\{v_i\}$ are known, and the knowledge of the real dimensions of the screen provide$\{x_i\}$ and $\{y_i\}$. The relations 3.59 between these n points can be cast into an $n \times 8$ matrix equation, the inversion of which recovers the transformation parameters

$$\begin{pmatrix} u_1 & v_1 & 1 & 0 & 0 & 0 & -x_1u_1 & -x_1v_1 \\ 0 & 0 & 0 & u_1 & v_1 & 1 & -y_1u_1 & -y_1v_1 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ u_n & v_n & 1 & 0 & 0 & 0 & -x_nu_n & -x_nv_n \\ 0 & 0 & 0 & u_n & v_n & 1 & -y_nu_n & -y_nv_n \end{pmatrix} \begin{pmatrix} a \\ b \\ \vdots \\ g \\ h \end{pmatrix} = \begin{pmatrix} x_1 \\ y_1 \\ \vdots \\ x_n \\ y_n \end{pmatrix}. \tag{3.60}$$

Note that it is possible to use more than 4 points, in which case a set of averaged transformation parameters will be recovered.



(u, v) (x, y)

Figure 3.29: An experimental reference image of an electron spectrometer screen, before and after image warping. The selected corners of the screen are depicted as red circles.

Transformation Jacobian

When performing image transformations such as these, it is important to conserve the physical quantities represented by the image. In this case the important quanitity is the brightness and location of image pixels representing electron charge. Upon transformation

the amount of charge represented by one camera pixel should remain the same, which can be non-trivial when the camera pixels end up contributing to more than one transformed pixel.

When performing a computational transform of a 2D array, most algorithms are created to work backwards – the centre of a desired array element is chosen in the (x, y) plane, and transformed back to the (u, v) plane. In general this position will not be at the centre of a particular element in the (u, v) plane, so some interpolation is required. The smoothest method is bicubic interpolation, when nearby pixels are weighted with a cubic spline and added to create a local average. This average value is then inserted at the appropriate point in the (x, y) array.

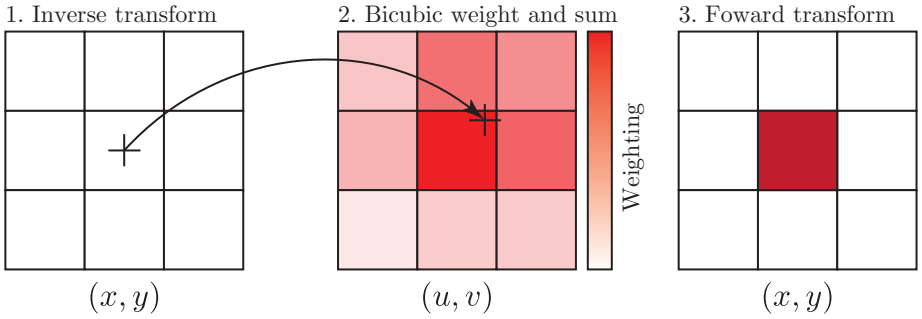


Figure 3.30: General image transformation algorithm.

If a bright region of the screen is significantly stretched in the transformation, that region will have the same brightness in the transformed image but spread over a larger number of pixels. The warping process would then appear to be creating extra charge on the screen, which is unphysical. The solution is to take into account how the size of each pixel changes, and adjust the brightness appropriately. If a bright pixel becomes two pixels in the transformed image, each of those pixels should inherit half of the value of the original. This is most simply accomplished by calculating the Jacobian of the transformation, defined as follows.

The total charge on the screen is given by an integral over the image I, and should be constant between co-ordinate systems

$$Q = \iint I(u, v)\,\mathrm{d}u\,\mathrm{d}v = \iint I(x, y)J(x, y)\,\mathrm{d}x\mathrm{d}y \tag{3.61}$$

where J is the Jacobian of the transformation which defines the ratio of area elements

$$J(x, y) = \frac{\partial(u, v)}{\partial(x, y)} = \begin{vmatrix} \frac{\partial u}{\partial x} & \frac{\partial u}{\partial y} \\ \frac{\partial v}{\partial x} & \frac{\partial v}{\partial y} \end{vmatrix}. \tag{3.62}$$

Calculation of J in the target (x, y) plane requires knowledge of the inverse transformation mapping (x, y) to (u, v), which is simply the inverse of the matrix in Eq. 3.57. Alternatively the Jacobian may be calculated in the (u, v) plane analytically and transformed to the (x, y) plane. In this case it has a simple explicit form in terms of the transformation

parameters

$$J(u,v) = \lambda^2(ae - bd) + \lambda^3((dh - eg)(au + bv + c) + (bg - ah)(du + ev + f)) . \quad (3.63)$$

3.4.3 Energy calibration

After correcting for perspective distortions in the image of the spectrometer screen, the mapping between CCD pixel and position on screen is simple. Before a spectrum can be calculated, the mapping between position on screen and energy must be calculated. While it is possible to obtain an estimate of this mapping from the calculations above, it is more accurate to measure the actual magnetic field of the spectrometer magnet. This step is important because fringe fields can extend far from the edge of the magnet, and affect the electron motion over relatively large distances. Once a field map is measured electron trajectories may be numerically integrated through the magnet and onto the measured position of the Lanex screen. Plotted in Fig. 3.31 is an example field map from the experiment discussed in chapter 6, with Lanex screen positions and numerical electron trajectories overlaid.

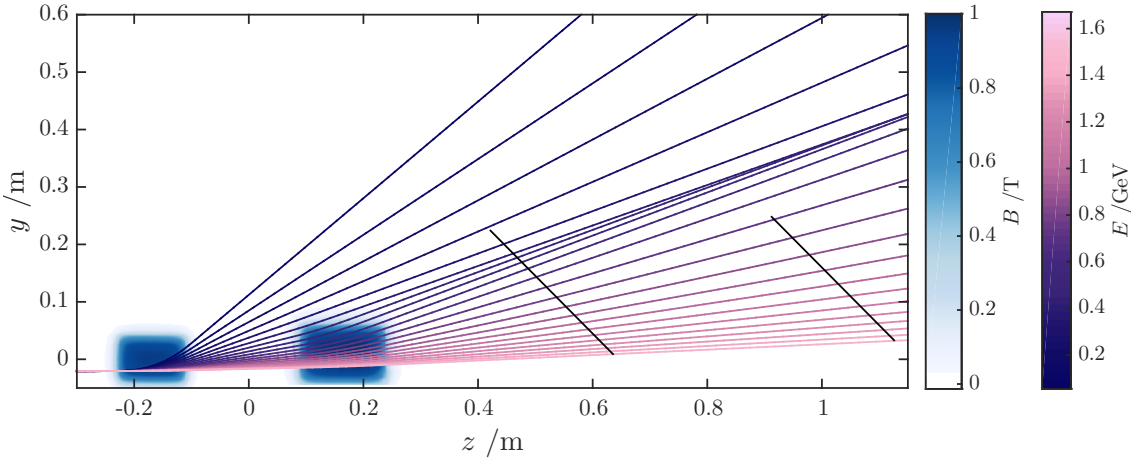


Figure 3.31: Numerically calculated electron trajectories for the measured magnetic field of a dipole spectrometer.

As a spectrometer designed, in this case, for high-energy electron beams, electrons of energy below 200 MeV miss the top of the spectrometer screens. A 30 cm-long Lanex screen positioned approximately 50 cm from the end of the magnets will have access to an energy range of 200-1500 MeV.

3.4.4 Experimental implementation

Pictured in Fig. 3.32 is part of the electron spectrometer from an experiment at the Gemini laser facility in 2012, discussed in Chapters 4 and 5. In this case the dipole magnet consisted of a pair of 150 mm×100 mm permanent 1 T magnets which could be moved in and out of the electron beam. The Lanex screen is placed above the laser axis and angled at approximately 45°, increasing the energy resolution as discussed above. Though not visible, a second Lanex screen is placed farther along the beam axis so that two spectra may be recorded separately. This allows the influence of pointing fluctuations of the

electron beam to be extracted from the energy spectrum.

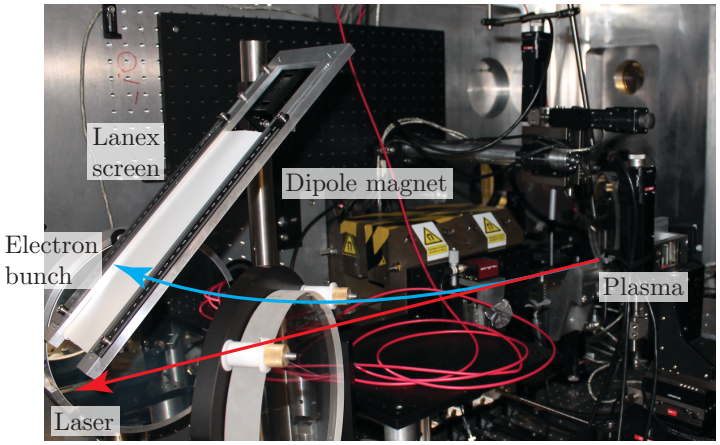


Figure 3.32: Photograph of a typical electron spectrometer configuration alongside other components of an experiment. Not pictured are the CCD cameras viewing the Lanex screen from above.

In Fig. 3.33 the steps in the analysis of an experimental electron beam are plotted. The raw images from the screens are foreshortened and rotated, though in this case the perspective distortion is minimal. After transforming the images the spectra may be calculated, and the two are overlaid on the right. Within the spectral window accessible to both screens the agreement is very good, indicating that this electron beam was travelling close to the laser axis. If the two spectra are displaced from one another, it is possible to calculate the electron beam pointing deviation from the optical axis and correct the spectra accordingly[30].

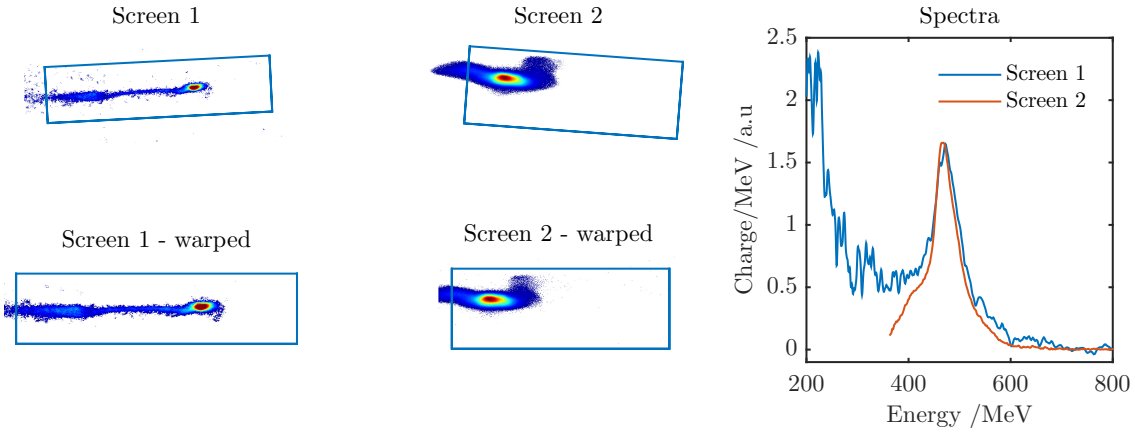


Figure 3.33: Example experimental images of electron spectrometer screens

3.5 Numerical simulation of laser-plasma interactions

The numerical simulation of laser-plasma interactions is a very useful tool in understanding the relevant physics and exploring a wider parameter space than may be possible experimentally. The very first theoretical exploration of the LWFA concept was supported by

[30] [250] A. A. Soloviev *et al. Review of Scientific Instruments.* (2011).

simulation[31], and simulations exploring LWFA at ultrahigh laser intensities predicted the possibility of monoenergetic beam production[32], an important experimental achievement in the field.

The laser-wakefield acceleration process is a complex one to simulate, consisting of an ultrashort laser pulse, complex plasma structures and the production of non-thermal relativistic electron bunches. Some simulation codes explicitly separate the components of the plasma which may be approximated as a fluid[33], or represent the laser fields as a slowly-varying envelope[34]. In order to represent the laser pulse evolution and particle injection/acceleration most accurately, the full electromagnetic fields of the laser and plasma must be directly evolved with Maxwell's equations. Maintaining maximum fidelity in the modelling of the plasma would require simulating the trajectory of every particle in response to these fields. To make this concept feasible the PIC (Particle-in-Cell) algorithm was introduced where the electromagnetic fields are defined on the vertices of a cartesian eulerian grid, and the trajectories of macroparticles (representing many real particles) through these grid cells are calculated. The macroparticles represent a discrete sampling of the species distribution function, and their trajectories the characteristics of the Vlasov equation. This sampling is efficient in that unoccupied regions of phase space are not included in the computation.

The simulation results presented in this thesis were produced with the EPOCH code[35], a fully relativistic particle-in-cell code written in FORTRAN with parallel support via MPI. The simulations took place on the CX1 supercomputing cluster based at Imperial College, either on a set of 6×12-core or 6×16-core Sandy Bridge Intel CPUs. The simulations typically required between 500 and 1100 CPU hours to run, depending on the spatial resolution, laser pulse parameters, and plasma length.

3.5.1 PIC algorithm

The general PIC algorithm is sketched in Fig. 3.34, and consists of four steps:

1. The position of each particle is updated according to the Lorentz force equation using a second-order particle push[36].
2. The charge and currents are deposited on the grid by summing over the particle positions. The particles may have different weights (i.e., represent different numbers of real particles), and different shapes – commonly top hat or smooth polynomials of various orders. Smoother particle weight functions decrease artificial heating at the expense of computation time.
3. The electromagnetic fields are advanced with Maxwell's equations using a finite-difference-time-domain method (FDTD) and stored on a Yee grid[37], where the

[31] [18] T. Tajima *et al. Physical Review Letters.* (1979).
[32] [145] A. Pukhov *et al. Applied Physics B: Lasers and Optics.* (2002).
[33] [123] P. Mora *et al. Physics Of Plasmas.* (1997).
[34] [251] C. Huang *et al. Journal of Computational Physics.* (2006).
[35] [252] T. D. Arber *et al. Plasma Physics and Controlled Fusion.* (2015).
[36] [253] J. P Boris. *Proceedings of the Conference on the Numerical Simulation of Plasmas.* (1971).
[37] [254] K. Yee. *IEEE Transactions on Antennas and Propagation.* (1966).

electric and magnetic fields are staggered by half a grid cell. This allows second-order accuracy in the central-difference approximation to the field derivatives.

4. The new fields are interpolated back to the positions of each particle. Depending on how this step is performed the PIC algorithm can exactly conserve energy or momentum, but not both simultaneously.

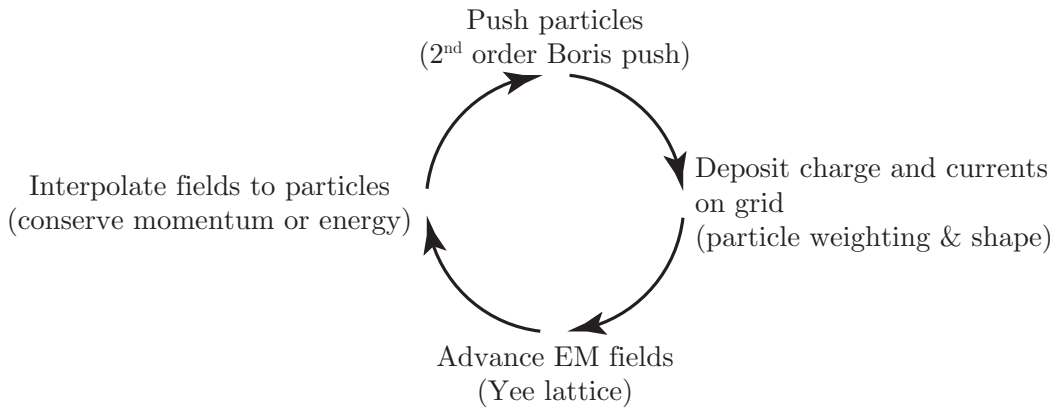


Figure 3.34: General PIC computation loop, performed per simulation step.

As every step in the PIC loop involves quantities local to each particle, it is a relatively simple step to parallelise the computation. In comparison, hybrid and implicit methods involve matrix inversions which are not as simple to scale. The simulation box is split into domains, each handled by a different CPU core. After each simulation step, which is handled by each core independently, only the information about the boundary of each domain need be communicated between cores. Some PIC codes also perform load balancing, where the domains are resized during the simulation such that each CPU core handles similar numbers of particles. Typically load balancing is performed when the peak load imbalance surpasses a user-defined level. In LWFA simulations $n_e \approx 0.001 \times n_c$, so $\lambda_p \approx 30 \times \lambda_0$ and the shortest length scale which must be resolved is the laser wavelength. However, the propagation distances involved are on the scale of millimetres and so there is a significant disparity in length scales involved. The solution in EPOCH, as in many PIC codes, is to use a 'sliding' window where the small simulation box travels at the group velocity of the laser pulse. A long propagation distance may then be simulated while discarding the unimportant information about the plasma which fails to propagate along with the wake. Physically this corresponds to a gallilean transformation and so requires minimal re-interpretation in the laboratory frame.

3.5.2 Resolution constraints

By substituting plane-wave solutions $\exp(i(kx - \omega t))$ into the 2D FDTD algorithm[38], the requirement that the electromagnetic fields do not grow exponentially in time leads to the Courant-Friedrichs-Lewy (CFL) condition on the simulation time-step Δt

[38] [255] A. Taflove. *Computational Electrodynamics: The Finite Difference Time Domain Method.* (1995).

$$c\Delta t < \frac{1}{\sqrt{\frac{1}{\Delta x^2} + \frac{1}{\Delta y^2}}}, \tag{3.64}$$

where Δx and Δy are the grid cell sizes in each direction. In the same manner, a numerical dispersion relation may be calculated for the plane wave as

$$\frac{1}{c^2\Delta t^2}\sin\left(\frac{\omega\Delta t}{2}\right)^2 = \frac{1}{\Delta x^2}\sin\left(\frac{k_x\Delta x}{2}\right)^2 + \frac{1}{\Delta y^2}\sin\left(\frac{k_y\Delta y}{2}\right)^2 \tag{3.65}$$

which leads directly to the numerical group velocity in the x-direction

$$\frac{\partial\omega}{\partial k_x} = v_g = c\frac{\cos(k_x\Delta x/2)}{\sqrt{1 - \frac{c^2\Delta t^2}{\Delta x^2}\sin^2(k_x\Delta x/2)}}. \tag{3.66}$$

In vacuum one would expect $v_g = c$, so numerical dispersion is minimised when $\Delta x \ll \lambda_x$. In practice it is often the case that $\Delta x \ll \Delta y$ and so the CFL condition reduces to $c\Delta t < \Delta x$, or that a wavefront should not be able to propagate more than one grid cell per time step.

Expression 3.66 is plotted in Fig. 3.35 and compared to a series of one dimensional EPOCH simulations, with spatial resolution expressed as a number of cells per laser wavelength. It is clear that finer spatial resolution improves numerical accuracy, at the expense of larger computational load. Accuracy is also improved as $\Delta t \to \Delta t_{\max} = \Delta x/c$ which is useful as larger time steps require fewer simulation steps overall.

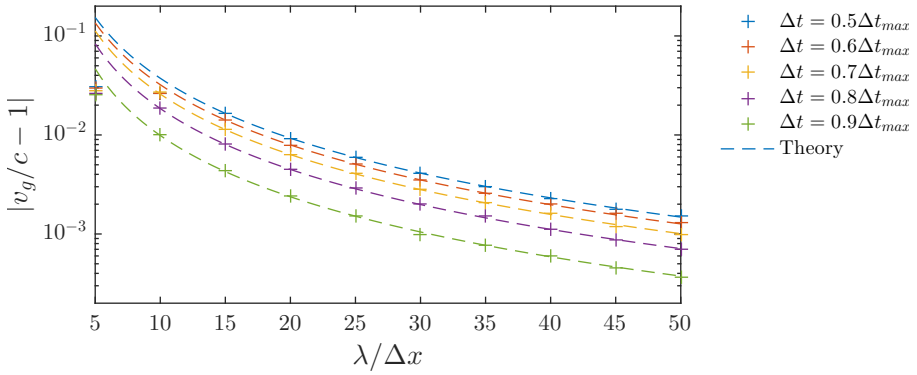


Figure 3.35: The magnitude of the group velocity error from a number of 1D EPOCH simulations compared to Eq. 3.66. $\Delta t_{\max}$ is the maximum timestep permitted by the CFL condition 3.64.

EPOCH has built-in capability for 4^{th} and 6^{th}-order field solvers too, which calculate spatial derivatives using a larger number of grid cells and are thus slightly more computationally intensive. For the 4^{th}-order field solver the CFL condition is altered to read

$$c\Delta t < \frac{6}{7}\frac{1}{\sqrt{\frac{1}{\Delta x^2} + \frac{1}{\Delta y^2}}} \tag{3.67}$$

and the numerical group velocity becomes

$$\frac{\partial\omega}{\partial k_x} = v_g = c\frac{\frac{9}{8}\cos(k_x\Delta x/2) - \frac{1}{8}\cos(3k_x\Delta x/2)}{\sqrt{1 - \frac{c^2\Delta t^2}{\Delta x^2}(\frac{9}{8}\sin(k_x\Delta x/2) - \frac{1}{24}\sin(3k_x\Delta x/2))^2}}\,. \tag{3.68}$$

The expressions for higher-order solvers become unwieldy so are omitted here. Plotted in Fig. 3.36 are the errors in group velocity from simulations with the fourth-order solver enabled. Errors are observed to be proportional to Eq. 3.68 but higher by 30%, a discrepancy with no obvious cause but presumably due to a subtlety in the field solver routines of the code. The numerical accuracy is higher than the second order solver, and again improves as grid cells become smaller. In this case however, the accuracy worsens as the timestep approaches the CFL limit.

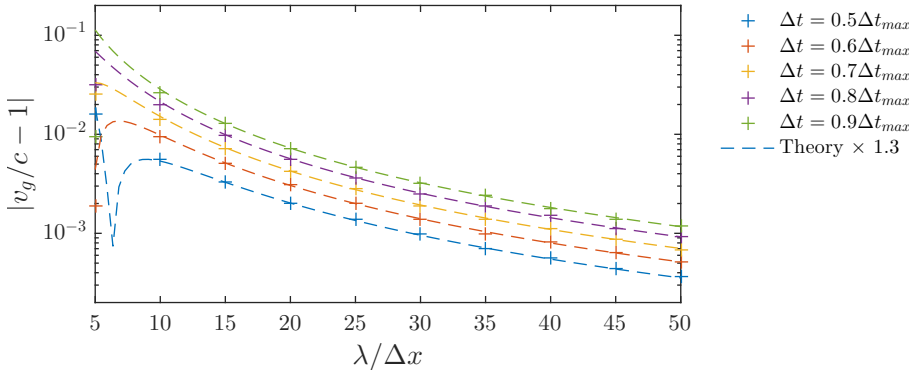


Figure 3.36: The magnitude of the group velocity error from a number of 1D EPOCH simulations compared to Eq. 3.68. Note that the error from EPOCH is approximately 30% higher than predicted.

While performing simulations for this thesis, it was observed that a relative error in group velocity $|v_g/c - 1| \gtrsim 10^{-3}$ over a propagation distance of several millimetres prevented correct self-compression of the laser pulse. In practice this equates to the stipulation that a 2D simulation using the (default) second-order field solvers should have a grid resolution of at least 30 cells per laser wavelength (a common rule-of-thumb), and that the transverse dimensions of the cell should be several times larger than the longitudinal resolution such that $\Delta t_{\max}$ is maximised. For simulations over shorter propagation distances this resolution threshold is likely less stringent.

3.5.3 PIC modifications

As particle trajectory information is available during a PIC simulation several groups have developed additions to the algorithm. Various codes now include the possibility of including ionisable species, particle collisions, synchrotron radiation emission[39], and QED processes[40]. LWFA-specific FDTD solvers have also been developed which help to improve numerical dispersion[41] and suppress artificial emittance growth by eliminating non-physical Čerenkov radiation[42]. In an effort to shrink the inherent disparity in spatial and temporal

[39] [256] A. G. R. Thomas. *Physical Review ST-AB.* (2010).
[40] [257] C. P. Ridgers *et al. Journal of Computational Physics.* (2014); [258] A. Gonoskov *et al. arXiv.* (2014).
[41] [259] B. M. Cowan *et al. Physical Review ST-AB.* (2013).
[42] [260] R. Lehe *et al. Physical Review ST-AB.* (2013).

scales, several codes support the use of a Lorentz-boosted simulation window[43] which can drastically reduce simulation time. It was not necessary to make use of these techniques for the simulations presented in this thesis, but for simulations involving longer plasmas, higher electron energies and higher laser intensities it is important to include next-order effects and efficient numerical methods.

3.6 Tomographic reconstruction

The results of chapter 6 are based upon the tomographic reconstruction of an extended three-dimensional object from a series of two-dimensional x-ray projections. The essentials of the reconstruction algorithm are presented here, as it is important to understand how the procedure works in order to mitigate the reconstruction artefacts caused by imperfect projection data.

Many tomographic techniques, including x-ray tomography, involve the use of projected data where the 3D signal of interest is integrated along a direction and recorded on a detector. For a given 2D slice of the object this process is equivalent to the Radon transform of the slice, so recovering the internal details of the object requires computation of the inverse Radon transform. Many techniques exist for this computation, one of the fastest, simplest and most robust is the filtered back projection (FBP), used extensively throughout the analysis of chapter 6. The fundamentals of FBP are presented below, adapted from the standard text by Kak and Slaney[44].

3.6.1 Fourier slice theorem

Define an (unknown) slice through a 3D object $f(x, y)$ where (x, y) are the 2D co-ordinates in the slice. The Radon transform of this slice is the projection at angle θ, or explicitly

$$P_\theta(t) = \int\int f(x,y)\delta(x\cos\theta + y\sin\theta - t)\,\mathrm{d}x\mathrm{d}y \tag{3.69}$$

where t is the radius. The set of all 1D projections of a slice f over projection angles θ is known as a sinogram, in reference to the fact that a feature in $f(x, y)$ will appear as a sinusoidal curve in the sinogram plane (t, θ). If a new rotated co-ordinate system (t, s) is used, related to (x, y) by

$$\begin{pmatrix} t \\ s \end{pmatrix} = \begin{pmatrix} \cos\theta & \sin\theta \\ -\sin\theta & \cos\theta \end{pmatrix} \begin{pmatrix} x \\ y \end{pmatrix}, \tag{3.70}$$

then the projection relation is more simply expressed as

$$P_\theta(t) = \int f(t,s)\,\mathrm{d}s\,. \tag{3.71}$$

Taking the Fourier transform of this projection, and substituting relation 3.71

$$S_\theta(w) = \int P_\theta(t)e^{iwt}\,\mathrm{d}t = \int\left[\int f(t,s)\,\mathrm{d}s\right]e^{iwt}\,\mathrm{d}t\,. \tag{3.72}$$

[43] [261] J.-L. Vay. *Physical Review Letters.* (2007); [262] S. F. Martins *et al. Nature Physics.* (2010).
[44] [263] A. C. Kak *et al. Principles of Computerized Tomographic Imaging.* (1988).

The transformation 3.70 is orthogonal and so has unit Jacobian, so the area elements between co-ordinate systems are equal and $\mathrm{d}s\mathrm{d}t = \mathrm{d}x\mathrm{d}y$. Rewriting in terms of (x, y) then,

$$S_\theta(w) = \int\int f(x,y)e^{i(w\cos\theta x + w\sin\theta y)}\,\mathrm{d}x\mathrm{d}y \equiv F(w\cos\theta, w\sin\theta) \tag{3.73}$$

i.e. the Fourier transform of the projection of f is equal to a slice of the 2D Fourier transform of f – this is the Fourier slice theorem.

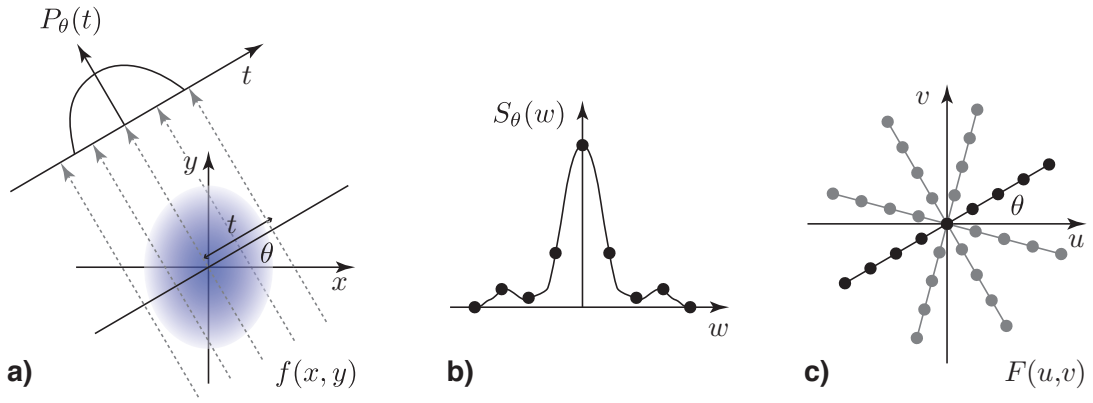


Figure 3.37: Schematic of the steps involved in the Fourier slice theorem **a)** The Radon transform of $f(x, y)$ at θ is $P_\theta(t)$ **b)** The Fourier transform of $P_\theta(t)$ is $S_\theta(w)$ **c)** $S_\theta(w)$ is related to the 2D Fourier transform of f as $F(w\cos\theta, w\sin\theta) = S_\theta(w)$.

From Fig. 3.37c) it is clear that as more projections of f are taken, the Fourier plane becomes more and more densely populated. If F is known, then so is f through the inverse 2D Fourier transform. A greater number of projection angles provide more knowledge of F, and therefore a better reconstruction of f.

3.6.2 Back projection

Starting from the explicit statement of the 2D inverse Fourier transform of F

$$f(x,y) = \int_{-\infty}^{\infty}\int_{-\infty}^{\infty} F(u,v)e^{-i(ux+vy)}\,\mathrm{d}u\mathrm{d}v\,, \tag{3.74}$$

we switch to the polar co-ordinate system $(u, v) = (w\cos\theta, w\sin\theta)$ such that

$$f(x,y) = \int_0^{2\pi}\int_{-\infty}^{\infty} F(w,\theta)e^{-iw(x\cos\theta + y\sin\theta)}w\,\mathrm{d}w\mathrm{d}\theta\,. \tag{3.75}$$

By splitting the integral into two parts from $[0, \pi]$ and $[\pi, 2\pi]$, and using the fact that $F(w, \theta + \pi) = F(-w, \theta)$, we have

$$f(x,y) = \int_0^{\pi}\left[\int_{-\infty}^{\infty} F(w,\theta)e^{-iwt}|w|\,\mathrm{d}w\right]\mathrm{d}\theta \tag{3.76}$$

where $t = x\cos\theta + y\sin\theta$ as before. Using the Fourier slice theorem,

$$f(x,y) = \int_0^{\pi}\left[\int_{-\infty}^{\infty} S_\theta(w)e^{-iwt}|w|\,\mathrm{d}w\right]\mathrm{d}\theta, \tag{3.77}$$

or defining

$$Q_\theta(t) = \int_{-\infty}^{\infty} S_\theta(w)|w|e^{-iwt}\,\mathrm{d}w \tag{3.78}$$

then

$$f(x,y) = \int_0^{\pi} Q_\theta(x\cos\theta + y\sin\theta)\,\mathrm{d}\theta\,. \tag{3.79}$$

Relations 3.78 and 3.79 are the key to the filtered back projection algorithm. The quantity $Q_\theta(t)$ is a equivalent to a filtered version of the projection data $P_\theta(t)$, where the filtering takes the form of a weighting in frequency space. Referring to Fig. 3.37c) one can see that this weighting is appropriate – sample points at large radius in the frequency plane are responsible for a larger region, and therefore must be weighted accordingly. The weighting increases linearly with distance, so it is correct that the points be weighted by $|w|$. Relation 3.79 also has a clear interpretation – the filtered projection data is 'smeared' across the (x, y) plane by the integral, or back-projected. It is also clear that complete information about f only requires projection information over $[0, \pi]$, and so projections need only be taken over 180° rather than a full rotation.
Many implementations of the FBP algorithm include the option to use additional filters in frequency space when constructing $Q_\theta(t)$. One of the most common is the Hamming window, which is the optimal choice for eliminating 'lobes' due to the finite sampling in frequency space imposed by the resolution of the detector. Most filters act to decrease the influence of high spatial frequencies which are commonly due to noise.

3.6.3 Reconstruction artefacts

Reconstruction of an object from a finite number of 2D projections will always introduce error into the reconstruction. In the following sequence of figures a variety of such artefacts are demonstrated, along with methods for their mitigation. The reconstruction object is the Shepp-Logan phantom[45], a standard analytically-defined tomography phantom composed of 10 ellipses (see appendix A.5 for details). The phantom was intended by its authors to represent a transverse slice through a human head.

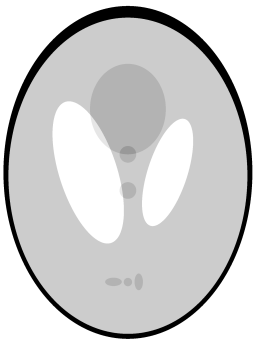


Figure 3.38: The Shepp-Logan phantom used to explore reconstruction artefacts in this section.

[45] [264] L. A. Shepp *et al. IEEE Transactions on Nuclear Science.* (1974).

In a two-dimensional digital image the discrete elements of the image are pixels (picture elements). In a three-dimensional tomographic reconstruction the analogous unit is the voxel (volume element). The voxel is cubic, and has the same linear dimensions as the detector pixel, e.g. a CCD with $10 \times 10\,\mu m$ pixels generates a reconstruction with $10 \times 10 \times 10\,\mu m$ voxels.

Finite number of projections

To explore the variation of reconstruction quality with projection number, the phantom is numerically projected along a finite number of directions N_{proj} and then reconstructed using the standard FBP technique. In Fig. 3.39 and subsequent figures, the upper and lower lineouts indicated by dashed lines are plotted in the second and third plots respectively. By subtracting the reconstruction from the initial (well-defined) phantom, an error level is plotted above. Using 900 projection angles in Fig. 3.39 produces a reconstruction which is observed to be very good, with small discrepancies around hard edges but capturing the detail from all features.

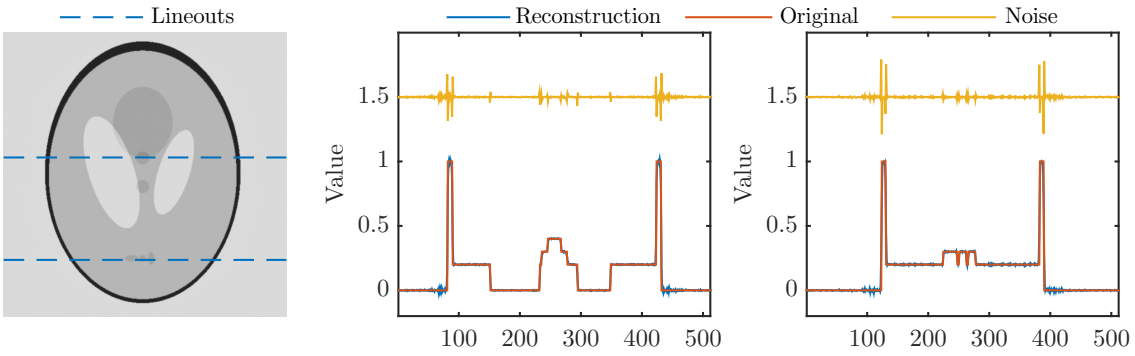


Figure 3.39: Reconstruction of the 512 × 512 voxel Shepp-Logan phantom with 900 projection angles.

Dropping N_{proj} by a factor of 20, Fig. 3.40 shows the reconstruction from 45 angles. It is clear that the reconstruction is affected by so-called streak artefacts which cause severe modulations in the gray values of the image towards the edges. The noise level increases substantially and the 3 small ellipses become difficult to differentiate.

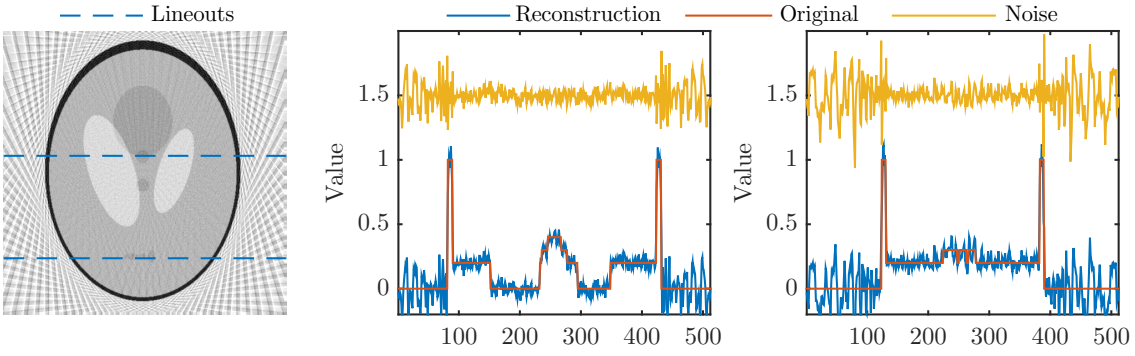


Figure 3.40: Reconstruction of the 512 × 512 voxel Shepp-Logan phantom with 45 projection angles.

It is obvious that fewer projections will degrade reconstruction quality, but there is also

an upper limit on the number of projections which it is reasonable to take. Consider a reconstruction volume which has transverse dimensions of N voxels, as sketched in Fig. 3.41. Two projections recorded an angle $\Delta\theta$ apart create features separated by at most $\approx N\Delta\theta/2$ voxels at the edge of the reconstruction volume. If this distance becomes less than 1, i.e. $\Delta\theta_{\text{min}} = 2/N$, then all voxels receive contributions from all projections and so no more information is gained by taking more projections. This can be converted to a maximum number of projections $N_{\text{max}} = \pi/\Delta\theta_{\text{min}} = N\pi/2$ which maximises the reconstruction quality.

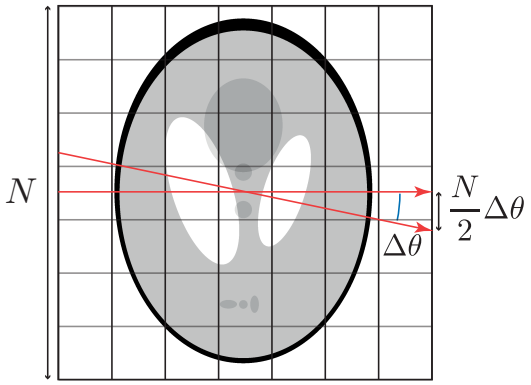


Figure 3.41: Projections offset by $\Delta\theta$ produce features separated by at most $N\Delta\theta/2$.

To check this threshold numerically, plotted in Fig. 3.42 are the standard deviations of smooth sections of the reconstruction as a function of N_{proj} and the number of transverse pixels in the phantom N_{pix}. It is observed that the noise level in the reconstruction drops quickly with the number of projections, then reaches a minimum beyond which no further improvement is seen. The fewer the number of pixels in the phantom, the sooner this threshold is reached. By overlaying the threshold derived above as dotted lines, it is clear that it is a good approximation. Before the threshold is reached the noise level is observed to decay like $N_{\text{proj}}^{-\alpha}$ for $\frac{1}{2} < \alpha < 1$.

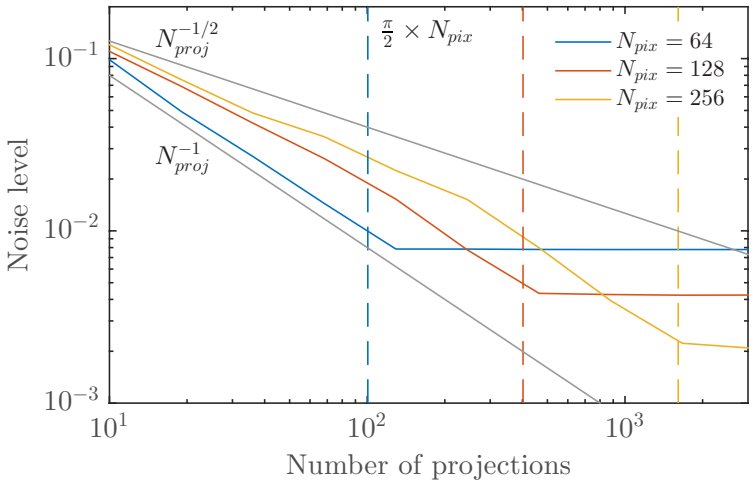


Figure 3.42: Standard deviation of phantom body as a function of N_{proj} for various numbers of voxels in the phantom.

Spatial resolution of detector

It is possible to simulate a finite detector spatial resolution by convolving the projection images with a point spread function (PSF) before reconstruction. An example is plotted in Fig. 3.43 and as might be expected the reconstruction becomes smoother and more blurred with finer features being lost.

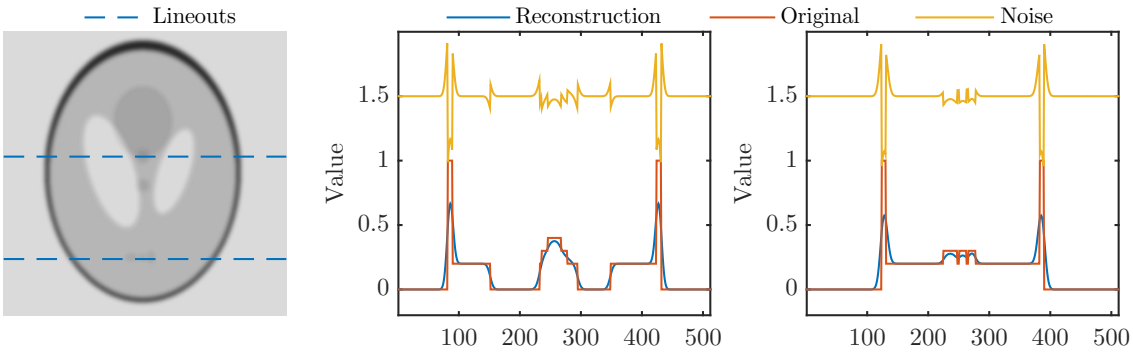


Figure 3.43: Blurring by PSF of 10 pixels FWHM.

An important consequence of this is that the effective number of pixels in the image is reduced, and therefore the maximum number of projections which should be taken is lower for a broader PSF. The change in reconstruction noise with N_{proj} is plotted in Fig. 3.44 for gaussian PSFs of successively larger FWHM. As the detector resolution worsens, the reconstruction noise is observed to reach a minimum at lower N_{proj}. In this case, for FWHM > 4 pixels the threshold is observed to scale like FWHM^{-2}.

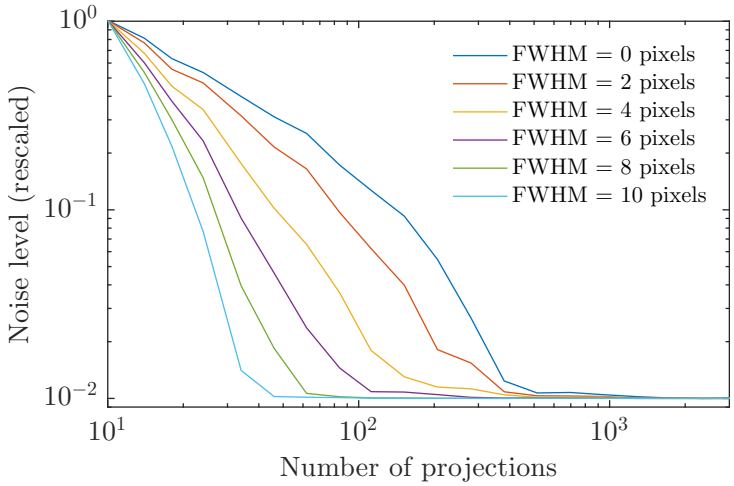


Figure 3.44: Standard deviation of phantom body as a function of N_{proj} for various detector resolutions.

Fluctuations in transmission

If the x-ray spectrum of the source varies from projection to projection, then so does the transmission of the object. Changes in x-ray intensity are less important, as it is the relative transmission of the object which is the input to the FBP algorithm. It is possible to simulate the impact of this effect by multiplying each projection by a random number,

and measuring the standard deviation of a region of the resulting reconstruction. A number of reconstructions were performed with different levels of shot-to-shot transmission fluctuation, plotted in Fig. 3.45. The noise in the reconstruction varies slowly with fluctuations in transmission, only increasing by 10% for a 50% transmission variation. Changes to transmission appear in the sinogram as vertical lines – see Fig. 3.46. An effective algorithm exists to remove such stripes[46], developed in the context of sinogram filtering. The image is first decomposed into a Daubechies wavelet representation[47], then the striped features are removed with a Fourier filter before reconstruction. As is clear in Fig. 3.46 the procedure works very well for removing stripe artefacts which are one pixel wide. In this case the improvement to the reconstruction is minimal though, as discussed above the transmission fluctuations have to become severe before the quality of the reconstruction is degraded. By looking at the difference between the filtered reconstructions the existence of streak artefacts appears to be reduced, potentially at the expense of uniformity.

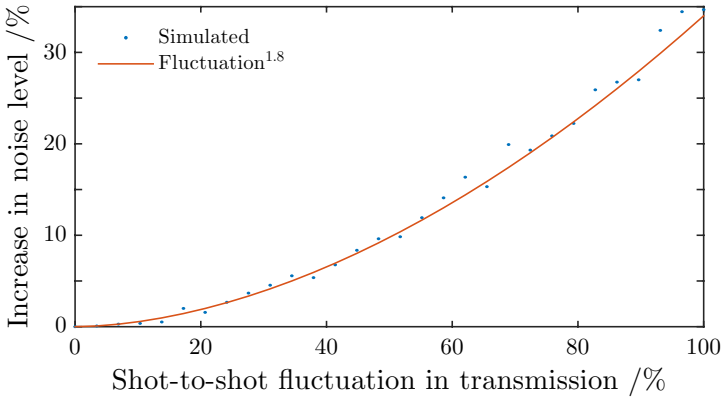


Figure 3.45: The noise level of the reconstruction plotted as a function of the level of shot-to-shot fluctuations in object transmission.

Ring artefacts

If while recording projection images a feature remains fixed on the detector, a horizontal streak will appear in the sinogram. After FBP reconstruction the streaks manifest as rings, visible in Fig. 3.47. The same algorithm as above may also be applied in this case. Here the number of pixels across the sinogram is greater than the number of projection angles so the streak artefacts in the sinogram are well separated from the image data in Fourier space. The filtering can therefore be much more aggressive on that region of Foureir space, resulting in a more efficient suppression of artefacts. See also Fig. 6.13 of chapter 6 for an example of an application of this filter to experimental data.

Beam hardening

Thus far the reconstructions have been produced simply by line-integrating the phantom to produce a projection, which is equivalent to using a monochromatic photon source. When imaging using an absorption-based technique, what is actually recorded on the detector for a photon energy E is

[46] [265] B. Münch *et al. EMPA Activities.* (2009).
[47] [266] I. Daubechies. *Ten Lectures on Wavelets.* (1992).

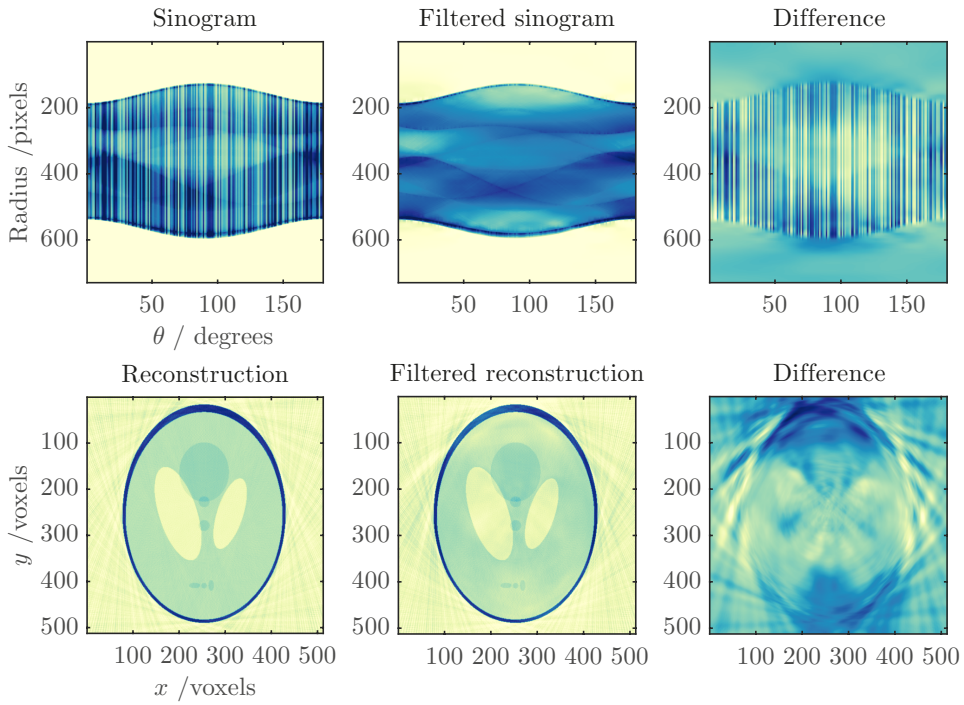


Figure 3.46: Streak artefact removal using a CWT-FFT filter.

$$I(x,y,E) = Q(E)I_0(E)\exp\left(-\int_{-\infty}^{\infty}\mu(x,y,z,E)\,\mathrm{d}z\right) \tag{3.80}$$

where $I(x,y,E)$ is the recorded intensity due to energy E, $I_0(E)$ is the unperturbed intensity at E (here assumed spatially constant), $Q(E)$ is the detector response at E and μ represents the absorption of the sample as a function of space and photon energy. Allowing multiple photon energies to be present in the source, as is the case in a broadband betatron spectrum, the total image is an integral over energy

$$I(x,y) = \int_0^{\infty} Q(E)I_0(E)\exp\left(-\int_{-\infty}^{\infty}\mu(x,y,z,E)\,\mathrm{d}z\right)\mathrm{d}E. \tag{3.81}$$

If the source is monochromatic at energy E_0, then $I_0(E) = I_0\delta(E-E_0)$ and the integral becomes simply

$$I(x,y) = I_0Q_0\exp\left(-\int_{-\infty}^{\infty}\mu_0(x,y,z)\,\mathrm{d}z\right) \tag{3.82}$$

where Q_0 is the detector response at E_0 and μ_0 is the absorption coefficient at E_0. What is actually required for the input to the FBP algorithm is the line integral of an unknown function, in this case μ_0. In the monochromatic case this can be found from the recorded intensity as

$$\int_{-\infty}^{\infty}\mu_0(x,y,z)\,\mathrm{d}z = -\ln\left(\frac{I(x,y)}{I_0Q_0}\right). \tag{3.83}$$

If the sample is assumed to be homogeneous and of thickness L, then the apparent thickness recovered from the FBP assuming a the sample is illuminated by monochromatically is

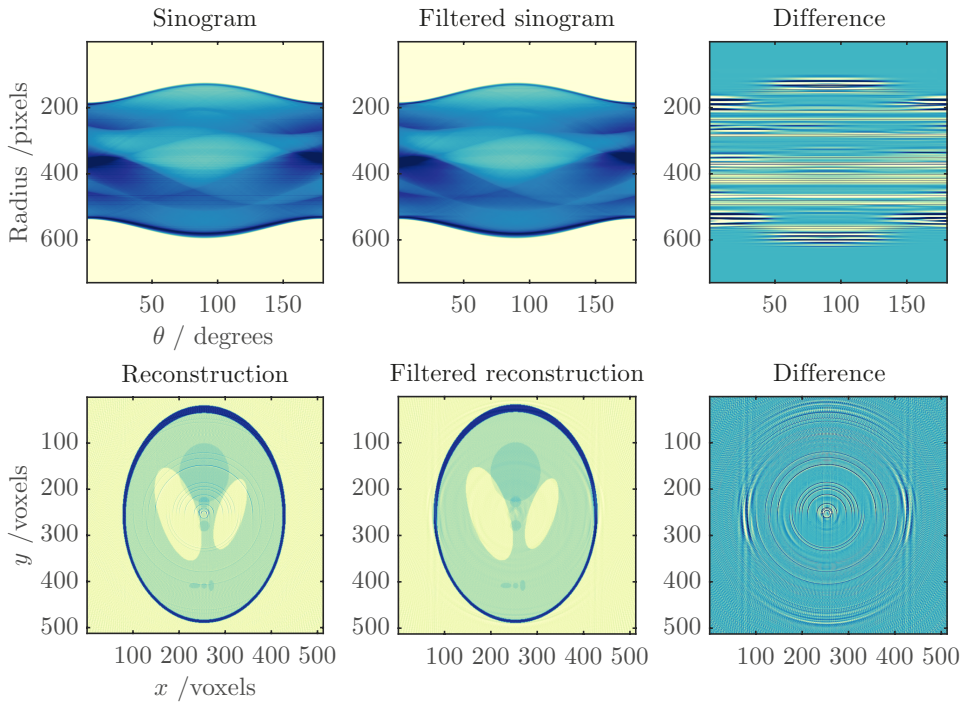


Figure 3.47: Ring artefact removal using a CWT-FFT filter.

$$L(x,y) = -\frac{1}{\mu_0} \ln\left(\frac{I(x,y)}{I_0 Q_0}\right) . \tag{3.84}$$

This expression is applied to calculated transmission factors for sources of gradually broader spectra in Fig. 3.48. The initial spectrum is monochromatic at 30 keV, then broadened until it reaches a synchrotron spectrum of $E_{\text{crit}} = 30\,\text{keV}$. The apparent thickness of the sample given by Eq. 3.84 as plotted on the right as the thickness of the sample is varied – in this case the absorption characteristics of the material are assumed to mimic human bone as appropriate for chapter 6.

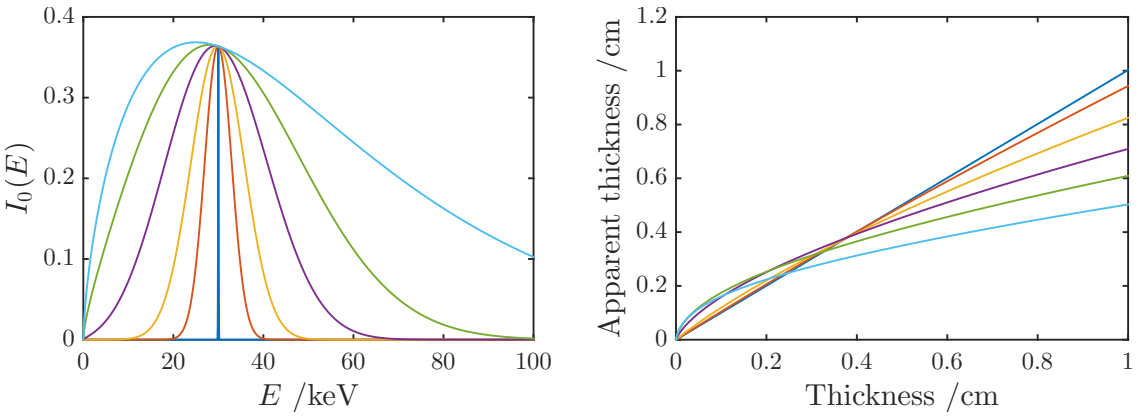


Figure 3.48: The spectra on the left are numerically transmitted through various thicknesses of material, and the apparent measured thicknesses plotted on the right. The colours of the curves match between the two plots, and the dark-blue line corresponding to a monochromatic spectrum marks the exact retrieval of the sample thickness.

In the monochromatic limit the measured thickness becomes the actual thickness, as the absorption characteristics of the sample are known to be constant. As the spectra become more broadband, Eq. 3.84 becomes a worse approximation and the apparent thickness no longer well approximates the actual thickness of the sample. In particular, short lengths appear longer and long lengths appear shorter. This is because the energy carried by the low energy photons is preferentially absorbed very quickly in the sample. The spectrum therefore becomes harder as it propagates through the sample, hence the expression 'beam hardening'. Likewise, the high energy photons are not absorbed efficiently and the transmission is higher than expected for longer samples. Assuming a mass absorption coefficient which scales like E^{-n} and a Gaussian beam spectrum centred at E_0 with width σ_E, the effective mass attenuation $\bar{\mu}$ coefficient can be shown, for $\sigma_E \ll E_0$, to become

$$\bar{\mu} = \mu_0 - \frac{n^2}{2}\left(\frac{\sigma_E}{E_0}\right)^2 \mu_0^2 L \tag{3.85}$$

where $\mu(E_0) = \mu_0$. For the energy range relevant to medical radiography, typically $n \approx 3$. The fact that the effective attenuation coefficient increases with decreasing L means thinner sections appear to attenuate more strongly.

These effects are demonstrated in the reconstruction of Fig. 3.49. The outer sections of the sample, encountered by photons traversing short sections of the sample, appear denser in the reconstruction. The inner parts of the sample, encountered by the more penetrating photons, appear less dense. This leads to a characteristic 'dishing' of the reconstruction, more clearly observed in the lineouts.

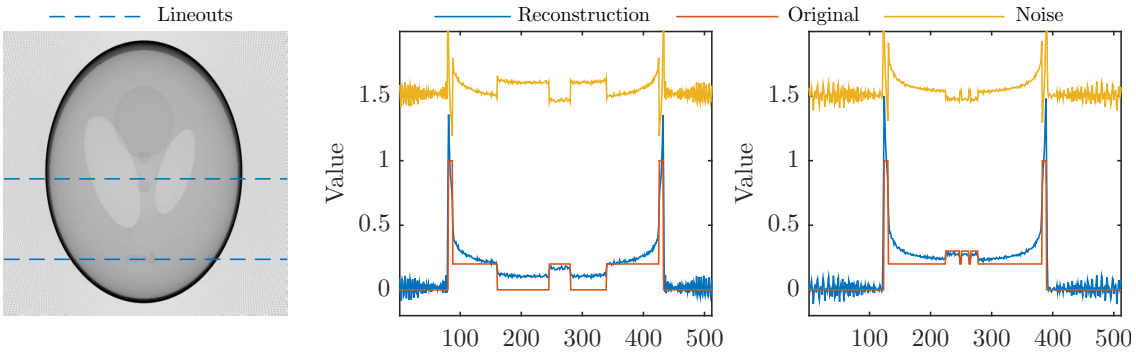


Figure 3.49: Beam hardening is observed to generate 'dishing' artefacts.

It is possible to mitigate the effects of beam hardening if the spectrum of the source is known and the object contains few materials[48], or by performing a calibration with objects of known uniform cross section[49]. Iterative correction techniques have also shown to be effective in removing beam-hardening effects due to the presence of highly absorbing materials[50].

[48] [267] J. P. Stonestrom *et al. IEEE transactions on bio-medical engineering.* (1981).
[49] [268] R. C. Roche *et al. Journal of Experimental Marine Biology and Ecology.* (2010).
[50] [269] G. Van Gompel *et al. Medical Physics.* (2011).

4 Stimulated Raman Scattering

In this chapter results are presented from two LWFA experiments performed during 2012 at the Central Laser Facility, Rutherford Appleton Laboratory, under the common theme of measurements of Raman scattering. The two experiments are discussed independently in separate halves of the chapter. The results of this chapter are framed in the context of the laser pulse evolution as it enters the plasma. This evolution is crucially important to the processes of injection and acceleration of electron bunches, and the consequent generation of betatron radiation. It is also a difficult process to observe directly, owing to its microscopic and ultrafast nature. Raman scatter, which by definition must originate from the volume containing the laser pulse, may then carry some information about the local properties of the plasma and the laser pulse from where it originated. It is therefore worthwhile to explore Raman scattering in laser wakefield accelerators, both as a useful diagnostic tool and as an interesting physical process in its own right.

Section 4.1 is concerned with spectral measurements of light which has been scattered from a 20 TW laser pulse by the Raman side-scattering instability. The light was found to have shorter wavelength than would be expected given the plasma densities present, measured independently with transverse interferometry. This observation is explained by assigning a moderate relativistic factor to the plasma electrons, implying the scattering occurs early in the plasma and near the peak of the laser pulse. From subsequent simulations, this scattering is observed to be associated with an initial period of self-compression of the laser pulse, and is therefore interpreted as aiding the initial pulse evolution.

Section 4.2 discusses the observation of plasma filaments in transverse shadowgraphic images. The angle made by the filaments with respect to the direction of the drive pulse is found to match that of the fastest-growing RFS mode, and the filaments are therefore attributed to RFS of the drive pulse. The probability of filament production is experimentally found to be correlated with poor quality focal spots and detuned focal plane positions, and to signify poor laser-plasma coupling efficiency. In simulations the generation of RSS is found to be minimised when the focal spot is symmetric and the focal plane is placed correctly.

4.1 Spectral measurements of side scatter

The first experiment was performed on the 20 TW Astra laser system. The primary purpose of the experiment was to commission a new short-duration probe beamline, and perform high resolution imaging of the wakefield structure. While that aim was not successfully achieved in this instance, direct imaging the wake has recently been performed

during a subsequent experiment using a similar setup at the JETI laser facility[1]. The results discussed here instead focus on spectral properties of Raman side-scattered light, and their relation to the independently measured plasma density.

4.1.1 Experimental setup

The experimental beamline layout is sketched in Fig. 4.1, along with the diagnostics relevant to the analysis presented here. During the experiment the Astra pulse delivered 610 ± 30 mJ to the target, had a FWHM duration of 41 ± 4 fs and was polarised horizontally (in the plane of the figure). From spectral measurements of the pulse the central wavelength was 808 nm with a FWHM bandwidth of 20 nm. The pulse was focussed with a 1 m length f/16 parabolic mirror to a FWHM spot size of 19×21 µm, reaching a peak intensity of $2.1 \pm 0.2 \times 10^{18}$ Wcm^{-2} corresponding to a normalised vector potential in vacuum of $a_0 = 1 \pm 0.1$. Approximately $35 \pm 4\%$ of the energy was contained within the FWHM contour of the focal spot. In contrast to previous experiments on the Astra laser, a folding mirror was placed between the parabolic mirror and the focal point which allowed the parabolic mirror to be mounted more stably, reducing fluctuations in the position of the focal spot.

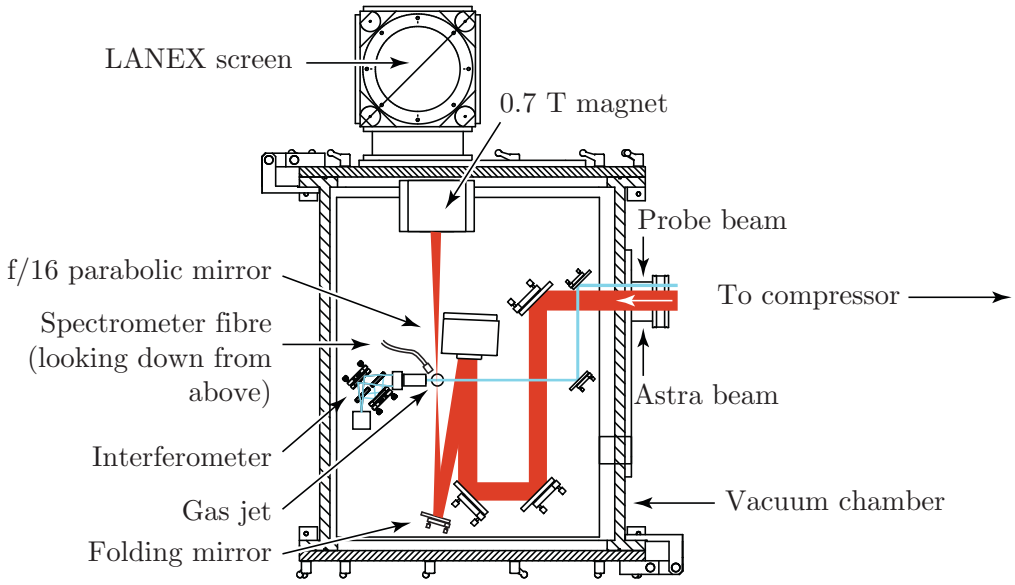


Figure 4.1: Plan view of the beamline layout for the 2012 Astra experiment.

The probe beam was formed from a small component of the drive beam, split off earlier in the laser amplification chain. The probe was compressed in air in its own compressor, then focussed onto the entrance of a hollow-core fibre filled with neon or argon where it was spectrally broadened through self-phase modulation [2]. The pulse was recompressed after a series of reflections from a pair of chirped mirrors before entering the target vacuum chamber. One of the aims of the experiment was to achieve a sub-10 fs duration for the probe pulse, which in practice fluctuated between 15 and 25 fs. The polarisation of the

[1] [231] M. B. Schwab *et al. Applied Physics Letters.* (2013); [152] A. Sävert *et al. Physical Review Letters.* (2015).
[2] [270] M. Nisoli *et al. Applied Physics Letters.* (1996).

probe was vertical (out of the plane of the figure), and a polariser was inserted before the transverse diagnostics in an effort to minimise the level of scattered pump light reaching the cameras. A high precision translation stage was inserted into the probe beamline to allow fine adjustments of the relative delay between the two pulses.

The plasma target was a 2 mm diameter conical helium gas jet, which when fully ionised was capable of producing electron densities of up to $10^{20}\,\mathrm{cm}^{-3}$. Relatively high plasma densities were required here when compared to the results in the latter half of the chapter due to the lower laser power, which increases the self-injection threshold[3]. The dephasing and depletion lengths are then shortened which necessitates the use of a short plasma. The typical neutral-gas density profile of the jet was trapezoidal, with 300 µm linear ramps at the edges leading to a constant-density plateau.

The electron spectrometer comprised of a 0.7 T 10 cm permanent magnet and a scintillating Lanex screen. High-energy electrons were dispersed across the screen and the scintillation light was imaged into a 14-bit camera, filtered to eliminate laser and ambient light. The accessible range of the spectrometer was 10 - 70 MeV, and typical peak electron beam energies were near 50 MeV. A second Lanex screen could be inserted before the magnet to characterise the beam size and pointing stability, which were approximately 10 mrad and 30 mrad RMS respectively.

4.1.2 Transverse imaging

In this experiment the interferometry took place entirely inside the vacuum chamber to minimise any imaging aberrations and group velocity dispersion of the probe pulse. In addition a 10 × magnification infinity-corrected microscope objective (Mitutoyo M Plan APO NIR 378-823-5) and tube lens (Mitutoyo MT-L4) were used to ensure high imaging quality across the field of view of the camera. The interferometer was designed in such a way that one of the arms was completely reflective – this was the arm used to form an image on the CCD. The second arm was used to offset a second image vertically to act as a reference for the interferometry. When assembled the diagnostic fit onto a small 175×225 mm breadboard, which was then mounted on a 3-axis translation stage. As sketched in Fig. 4.2 the configuration was compact and minimises the length of glass which must be traversed by the probe pulse. A motorised shutter was installed over the reference arm which allowed remote transition between interferometry and shadowgraphy.

Due to the high magnification the field of view was limited to 280 µm so the longest motion of the interferometer stage was along the pump laser propagation axis. Simultaneous adjustments of the probe timing and interferometer position allowed probing of the entire length of the gas jet. The depth of focus of the microscope objective was 4 µm, so before acquiring an image the interferometer was manually focused onto small structures on an alignment wire.

An example interferogram taken at low density is plotted in Fig. 4.3 alongside the retrieved phase shift. The short nature of the probe pulse caused the fringe visibility to vary significantly over the field of view of the camera. To alleviate this an 820 ± 5 nm bandpass filter was placed in front of the CCD. This lengthens the probe pulse enough to restore

[3] [158] S. P. D. Mangles *et al. Physical Review ST-AB.* (2011).

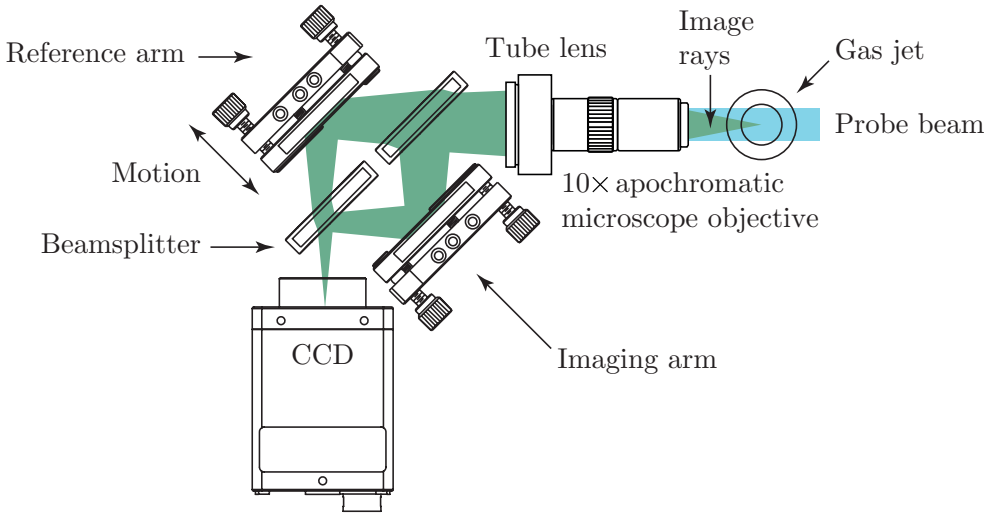


Figure 4.2: In-vacuum Mach-Zehnder interferometer layout.

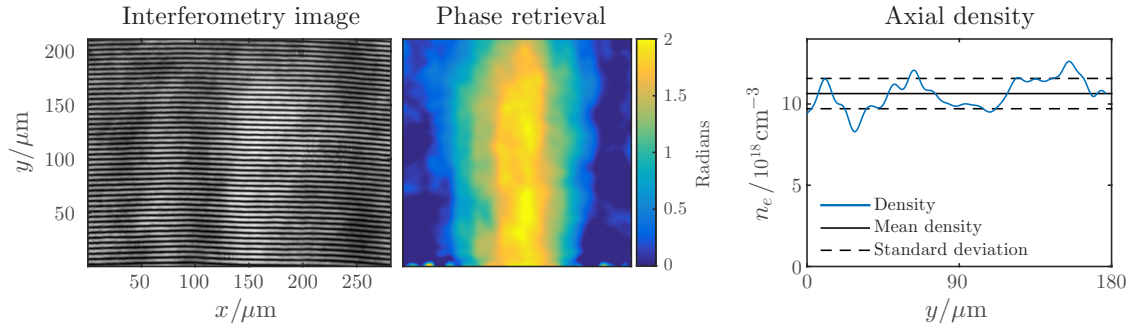


Figure 4.3: An interferometry image of a low density plasma channel, along with the phase and density retrievals. The drive pulse is propagating upwards in this image.

high-visibility fringes, but does not affect the imaging quality or temporal resolution[4]. A central denser plasma channel was often observed, surrounded by a less dense halo. When taking an average density measurement the central dense portion is used, averaged over a 20 µm diameter cylinder, as this is the region of the intense laser-plasma interaction. The axial density corresponding to the image of Fig. 4.3 is plotted alongside, with the mean and ± 1 standard deviation region indicated. If an average density is assigned to this image, the fractional error on the measurement is 9%.

The probe imaging was detailed enough to resolve fine filamentary structures in the channel, as well as the large-scale features. In Fig. 4.4 a selection of probe images which display a number of interesting features are displayed. In these images the drive pulse is propagating upwards, and approximately at the centre of the semicircular ionisation front where the plasma density stops increasing. The width of the plasma channel, approximately 150-170 um, is much larger than the size of the focal spot, and implies the existence of laser light sufficiently intense to ionise helium some distance away from the laser axis. For an Airy focal spot shape of peak intensity $2 \times 10^{18}\,\mathrm{Wcm}^{-2}$, the laser intensity drops below that required to have a 90% chance of ionising the first electron of helium at a radius

[4] [231] M. B. Schwab *et al. Applied Physics Letters.* (2013).

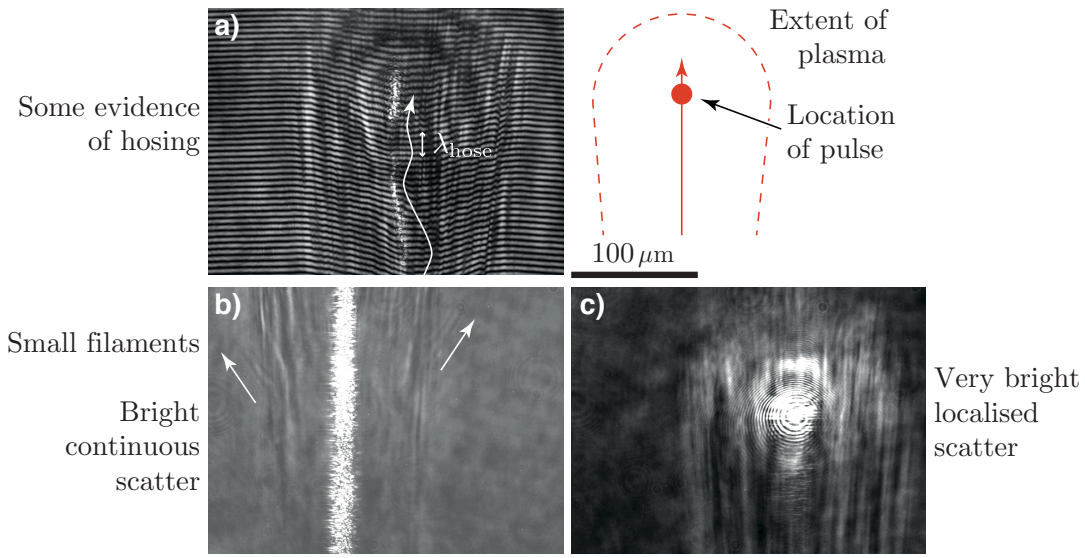


Figure 4.4: A selection of transverse probing images displaying interesting features.

of 88 µm, which is similar to the observed channel width. The channel is observed to expand in an f/16 cone, so as expected the outer regions of the focal spot are not guided. Additional bright features are observed which are due to light scattered from the drive pulse, which from the spectral measurements of section 4.1.3 would be expected to be caused by Raman scattering. Unlike the probe light, this is integrated over the entire exposure time of the camera and so constitutes a time history of the location of the drive pulse over the observed window. At high densities there is often some evidence of hosing, where the drive pulse oscillates transversely as it propagates. Kaluza *et al.* derive an expression for the hosing wavelength[5]

$$\lambda_{\text{hose}} \approx 6.7 \frac{\lambda_0^2}{a_0 r_{\text{FWHM}}} \left(\frac{n_c}{n_e}\right)^{3/2} \approx 20\,\mu\text{m}. \tag{4.1}$$

which assumes that the drive pulse oscillates transversely in its own ponderomotively-driven density depression. As overlaid on Fig. 4.4 a) the predicted wavelength is of the same order as that observed.

A more common occurrence was the continuous Raman side-scattering of light, visible in Fig. 4.4 b) as a bright vertical streak. The image here is slightly defocussed, so the fact that the scatter appears to be speckled suggests that it is formed from the superposition of many coherent wavefronts. Several models have been proposed for the 'bursting' nature of the Raman instability[6], but require many plasma periods for the wave to grow, saturate, and the plasma return to a quiescent state, and so it is unlikely that the instability can be periodic on the few-micron/femtosecond spatial/temporal scale observed here. The speckling is more likely a coherent imaging artefact due to the interference of continuously-emitted scattered wavefronts.

[5] [271] M. C. Kaluza *et al. Physical Review Letters.* (2010).

[6] [272] W. L. Kruer *et al. Physics of Fluids.* (1969); [273] K. Estabrook *et al. Physics of Fluids B: Plasma Physics.* (1989); [274] H. X. Vu *et al. Physical Review Letters.* (2001); [275] S. Brunner *et al. Physical Review Letters.* (2004).

Also visible in several shots are filaments escaping the main channel at an angle, and as is discussed in the following section 4.2 it is likely that these are formed by Raman scattering. In a minority of cases the transverse scatter took the form of a single bright spot as in Fig. 4.4c), localised near the position of the drive pulse and often associated with the presence of bright rings. As the ring spacing decreases with radius, this is likely not a diffraction artefact due to a misplaced object plane but instead the interference of a spherical wave (the scatter) with a plane wave (the probe). The fact that the fringe visibility is high means that the scattered light has significant intensity, comparable to the probe pulse, and that the light is scattered at the same time the probe is crossing the drive pulse. This effect was also witnessed in subsequent experiments, discussed in detail in chapter 5.

4.1.3 Spectral measurements

To measure the side-scattered light a spectrometer fibre was positioned directly above the gas jet at an angle close to 90°, with a collection lens attached to sample light from the entire gas jet. Light scattered in this direction should have the same polarisation as the drive pulse, and therefore the Raman side scatter growth rate should be largest. The fibre was connected to an Ocean Optics HR2000 spectrometer outside of the vacuum chamber which recorded the spectrum on every shot. Example scatter spectra from the spectrometer are plotted in Fig. 4.5, taken at low and high plasma densities. There is a small component of light from the unshifted laser spectrum which remains at very low plasma density, likely due to diffuse reflections from the gas nozzle and surrounding objects rather than Thomson scattering from the plasma. Also visible are several helium recombination lines – the spectrometer exposure time was significantly longer than the lifetime of the plasma. These lines act as a useful fiducial and increase in magnitude with plasma density. The dominant component of the spectra are peaks at longer wavelengths which are attributed to the Raman side scattering of the drive pulse, in most cases the brightest feature of the spectrum. In some cases the Raman peak displayed prominent spectral modulations. As plotted alongside, the spacing of these modulations $\Delta\omega$ increased with density, but the the spacing is much lower than ω_p as has been observed in previous experiments[7]. Some authors claim the onset of ion motion causes mixing of the Raman electron plasma waves with ion acoustic waves, leading to modulations in the spectrum of the scattered light at the ion plasma frequency[8]. Here the pulse is significantly shorter than ω_{pi} and $\Delta\omega$ too large for this to be an explanation. If converted to time intervals the modulations in the spectra correspond to $\Delta t = \pi/\Delta\omega = 100 - 300\,\text{fs}$. The pulse travels 30 - 90 µm in this time, which is similar to the spatial separation observed between scatter events in the transverse imaging (e.g, as in Fig. 4.4 a)). It may be possible then that the spectral modulations are due to the interference of side-scatter events separated in time by several hundred femtoseconds, and integrated temporally in the spectrometer.

A typical run of shots in this experiment was a plasma density scan, achieved in practice by filling the gas jet reservoir to a high pressure and taking shots while the pressure is

[7] [276] C. Rousseaux *et al. Physical Review Letters.* (1995).
[8] [277] A. Ting *et al. Optics letters.* (1996).

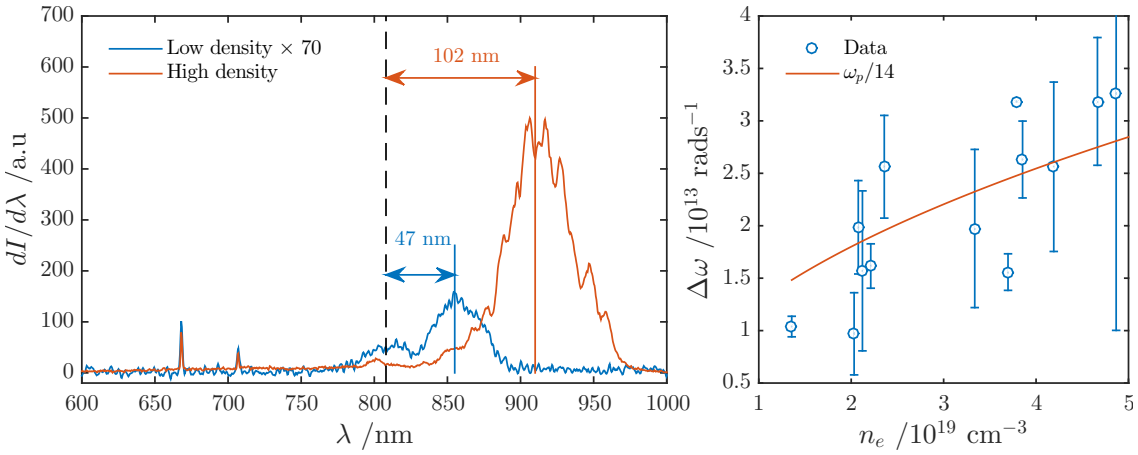


Figure 4.5: Measured side-scatter spectra and associated spacing of frequency modulations. Error bars correspond to ± 1 standard deviation in $\Delta\omega$ over repeat shots (if available) recorded at the same density

allowed to fall to zero. This method has the advantage that the fractional change in density between shots stays roughly constant throughout the scan. For a particular run of 104 shots analysed here, the gas jet pressure was varied from 23 bar down to 0.1 bar with spectra recorded on every shot. In Fig. 4.6 the spectrum from each shot is plotted as a column, centred at the interferometrically-measured plasma density corresponding to the pressure of the shot. The plot is coloured by the logarithm of the spectral intensity, normalised to the maximum signal observed during the run.

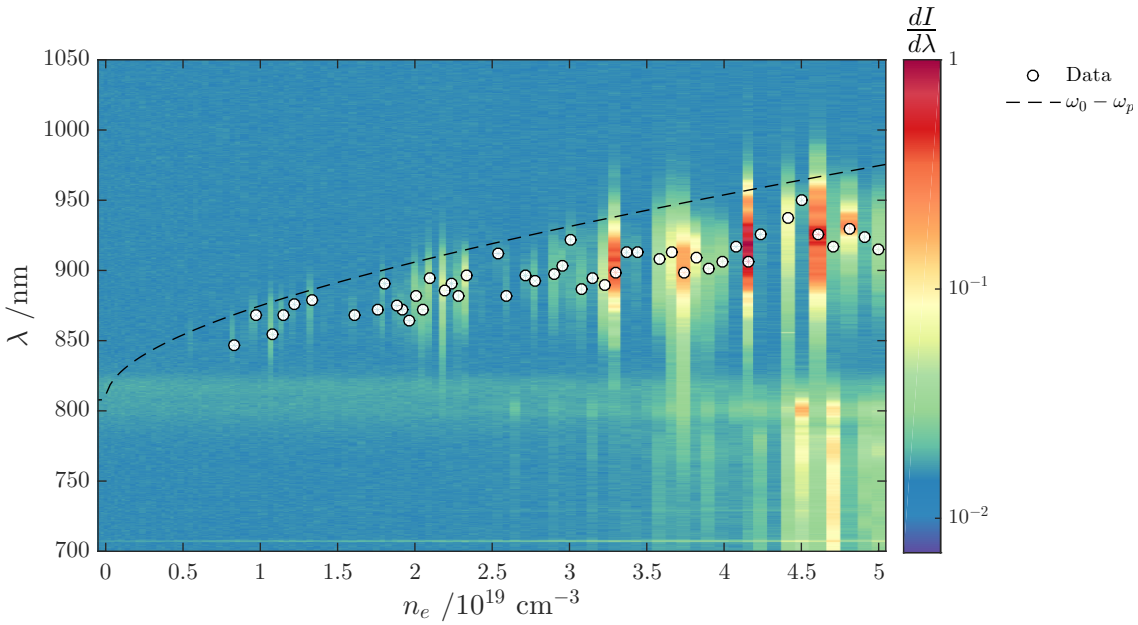


Figure 4.6: Recorded spectra of light scattered out of the polarisation plane of the drive pulse as a function of density.

Horizontal features are those which remain constant throughout the run, including the low-level reflection of fundamental laser light and the recombination lines. The down-shifted Raman peaks begin to appear above $n_e = 10^{19}\,\mathrm{cm}^{-3}$ and increase in wavelength

shift, intensity and bandwidth with density. For each shot, the wavelength of the peak in the Raman spectrum is indicated with a circle. Given the interferometrically-measured densities, the expected position of the Raman peaks according to linear theory $\omega_0 - \omega_p$ is plotted as a dashed line. The linear shift is significantly larger than the observed shifts, and would require a systematic 60% drop in the plasma density to match the data. Given the relative accuracy of the interferometric measurement, this seems unlikely.

4.1.4 Comparison to theory

The linear Raman theory is only valid for $a_0 \ll \sqrt{\omega_p/\omega_0}$, and the strong-coupling approximation only valid for $a_0 \gg \sqrt{\omega_p/\omega_0}$. In this dataset $a_0/\sqrt{\omega_p/\omega_0} \approx 2-10$ which means neither approximation is fully valid and the full dispersion relation of Eq. 2.123 should be solved directly. To accomplish this we pick plasma wave parameters (k, θ) in Eq. 2.123 which generate a scattered wave at 90° and solve for the complex frequency of the plasma wave $\omega = \omega_r + i\gamma$, where ω_r is the real part and represents the frequency shift $\Delta\omega$ of the scattered light, and γ is the growth rate of the mode. At relativistic intensities there are a range of solutions $\omega(k)$ – see Fig. 4.7 for a representative solution set. The wave frequency ω is represented by the fastest-growing mode as this will dominate the evolution of the instability. A bandwidth can be assigned to the solution from its FWMH width over ω_r.

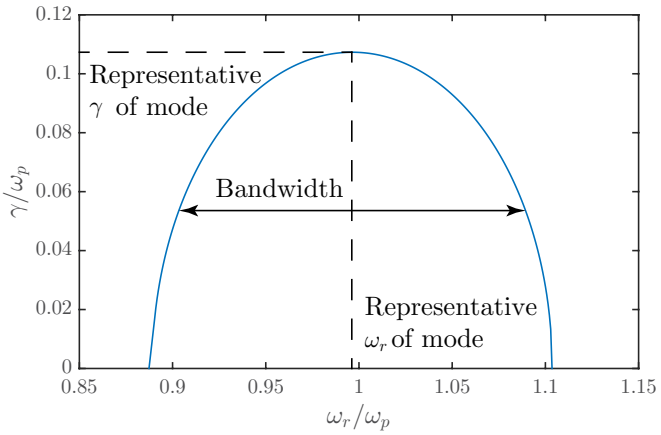


Figure 4.7: Solutions $\omega(k)$ of Eq. 2.123 plotted in the complex plane. The bandwidth is taken from the FWHM width of the set of solutions over frequency.

In general the solutions are as expected – growth rates increase with a_0 at low intensities as predicted by linear Raman theory[9], and decrease like $a_0^{2/3}/\gamma_0$ at high intensities as in the relativistic theory[10]. Here γ_0 is the Lorentz factor of the electrons due to their motion in the laser fields, where recalling Eq. 2.30 and averaging over laser phase, $\gamma_0 = 1 + a_0^2/2$. The plasma wave frequencies of the fastest-growing modes are observed to be well described by $\omega_r = \omega_p/\sqrt{\gamma_0}$ for $0.01 < a_0 < 10$ and $10^{18}\,\mathrm{cm}^{-3} < n_e < 10^{20}\,\mathrm{cm}^{-3}$. This linear scaling with ω_p is different from the $\omega_p^{2/3}$ scaling expected in the strong-coupling regime.

Another way in which density dependence of the growth rate may arise is through the self-focussing of the pulse. If the pulse self-focusses to a matched spot size then the peak

[9] [195] W. L. Kruer. *The Physics of Laser Plasma Interactions.* (1988).
[10] [190] A. S. Sakharov *et al. Physical Review E.* (1994).

intensity is not constant as the plasma density is increased. Assuming no depletion during self-focussing and no pulse compression (valid in the initial stages of the pulse evolution – see Fig. 4.13), the matched spot size w_m and vector potential a_m are related to the vacuum spot size and potential through $a_m w_m = a_0 w_0$. After incorporating the Lu scalings[11], when matched a_m scales with density as

$$a_m = \left(\frac{k_p^2 a_0^2 w_0^2}{4} \right)^{1/3} \propto \omega_p^{2/3}. \tag{4.2}$$

The experimentally-measured frequency shifts are plotted in Fig. 4.8 as a function of plasma density, along with the shifts expected for $\Delta\omega = \omega_p$, $\Delta\omega$ given by Eq. 2.123, and $\Delta\omega$ assuming peak $a_0 = a_m$.

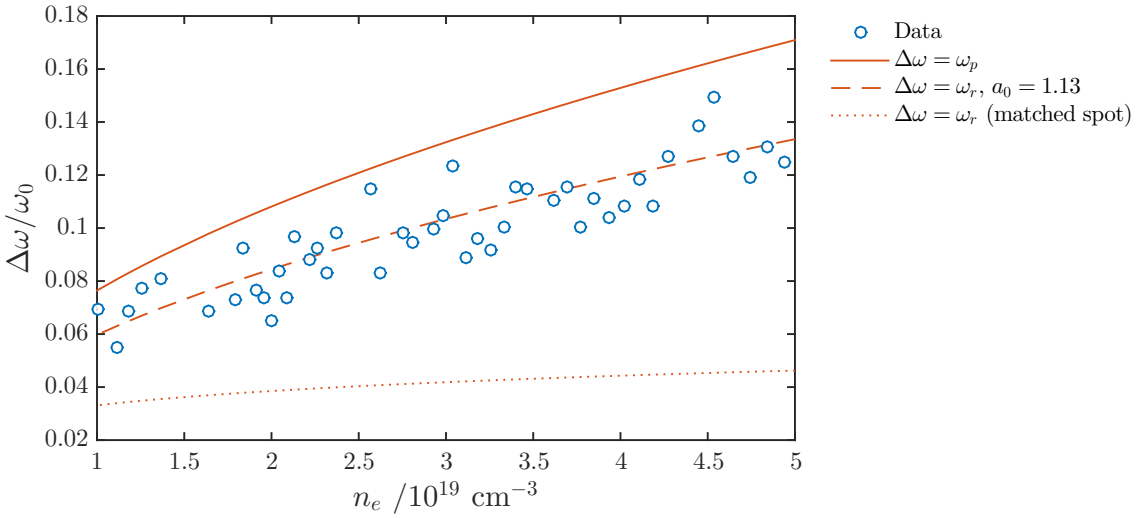


Figure 4.8: Measured frequency shifts as a function of plasma density.

The best fit to the data is $\gamma_0 = 1.64 \pm 0.10$, or $a_0 = 1.13 \pm 0.09$ (95% confidence intervals). Using the same value for a_0 in the strong-coupling expression produces frequency shifts which are too low, and which scale with density too weakly. If the matched values a_m are used instead, the theoretical frequency shifts are far too small and begin to fall with density. If some amount of depletion is assumed for the matched case the predicted shifts are higher but still far from those observed.

Bandwidth of scatter

Taking the best fit value for a_0 a theoretical bandwidth may be calculated. This is plotted in Fig. 4.9 as a function of frequency shift and density, along with the experimental data. The minimum observed bandwidths are approximately the same as that of the drive pulse. Here the correlation is much weaker, and the theoretical bandwidth is too high by approximately a factor of two. Taking $a_0 \approx 0.3$ improves the agreement between the data and theory, but is inconsistent with the previous fit to frequency shift. The maximum observed bandwidths for each density appear to scale with the theoretical bandwidths, but there are also shots with much lower bandwidths than expected. When plotted as a

[11] [150] W. Lu *et al. Physical Review ST-AB.* (2007).

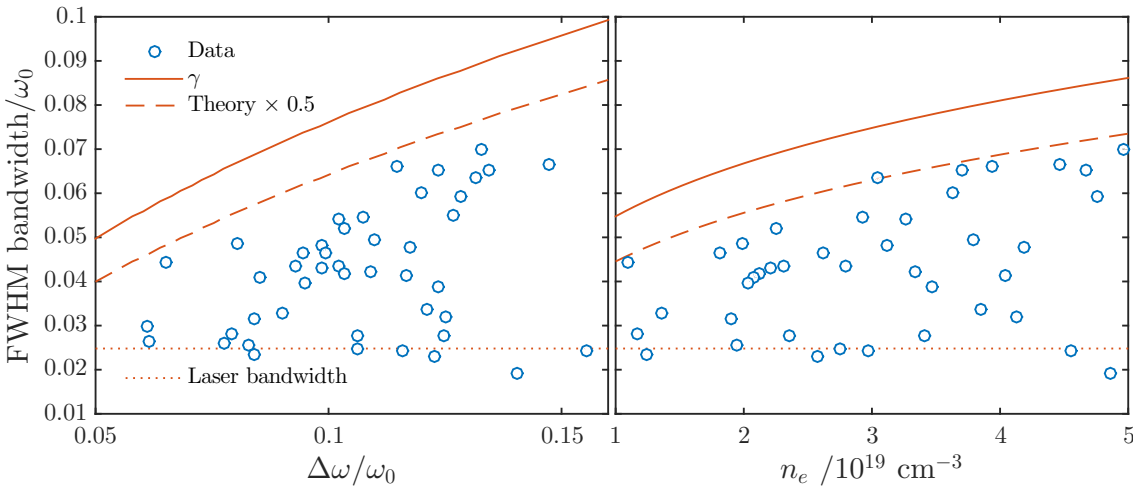


Figure 4.9: Experimental data are plotted as points, the dashed line represents the theoretical bandwidth scaled by 50%, and the solid line the theoretical growth rate. Both are calculated assuming $a_0 = 1.13$.

function of observed frequency shift rather than density the correlation is improved, and suggests that high-bandwidth scatter is associated with a high frequency shift.

The expected bandwidth of the scattered light can be calculated in different ways. One method is to examine the frequency spectrum of a growing electron wave mode, i.e. one oscillating in time at ω_r and growing at γ. The temporal form is then

$$a(t) \propto \cos(\omega_r t) e^{\gamma t}\,, \tag{4.3}$$

which has the Fourier transform

$$a(\omega) \propto \frac{1}{(\gamma^2 + \omega_r^2 - \omega^2)^2 + 4\gamma^2\omega^2}\,. \tag{4.4}$$

For small growth rates $\gamma < \omega_r/2$ the bandwidth in frequency space is approximately twice the growth rate. For large growth rates $\gamma > \omega_r/2$ the mode is no longer a well-separated peak in frequency space and the bandwidth increases more slowly, but is still approximately linear with γ.

In the strong-coupling regime $\gamma = \sqrt{3}\omega_r$ and so the bandwidth should be proportional to frequency shift[12]. From the general nonlinear dispersion relation[13] in the linear regime of a three-wave parametric instability[14], a frequency detuning by δ from exact resonance decreases the growth rate from $\gamma_{\max}$ to

$$\gamma(\delta) = \gamma_{\max}\sqrt{1 - \frac{\delta^2}{4\gamma_{\max}^2}} \tag{4.5}$$

and the FWHM bandwidth is $2\sqrt{3}\gamma_{\max}$, again proportional to growth rate. Solving the

[12] [278] C. B. Darrow *et al. Physical Review Letters.* (1992).
[13] [279] K. Nishikawa. *Journal of the Physical Society of Japan.* (1968); [280] K. Nishikawa. *Journal of the Physical Society of Japan.* (1968).
[14] [281] P. M. Bellan. *Fundamentals of Plasma Physics.* (2006).

SRS dispersion relation exactly shows that γ/ω_r drops slowly as a function of density, changing by $\approx 20\%$ over the density range considered here. We therefore expect some concordance between frequency shift, growth rate and bandwidth in all regimes. The fact that no simple relation holds for all of the data is perhaps an indication of the complex nature of the instability between weakly and strongly-coupled regimes, and its inherent nonlinearity.

In addition, it is curious that some of the scattered spectra had such low bandwidths, the smallest of which were near the laser bandwidth. One explanation may be gain narrowing, where the frequency-dependence of the growth rate causes narrowing of the spectral peaks. If the frequency dependence of the growth rate is assumed to have the form of Eq. 4.5, then after N e-foldings an initially flat noise spectrum becomes $e^{N\gamma(\omega)/\gamma_{\max}}$, with FWHM bandwidth $\gamma_{\max}\sqrt{8\log 2/N}$. If gain narrowing were the mechanism behind the relatively small bandwidths observed here, one might expect that the brightest scatter events are associated with narrower bandwidths. This is not observed in the data, but other effects may come into play for strongly exponentiated spectra. For example the fact that the scattered light is propagating down a plasma density gradient as it grows may lower the growth rates below those given above for the absolute instability. Rosenbluth[15] shows that density inhomogeneities limit overall growth and lower the effective growth rate. If the scattered light originates from a region of high plasma density near the pump pulse, its frequency is set by this density. It then rapidly propagates away and growth is quenched, so its bandwidth can be lower than that expected from its frequency shift and linear theory.

4.1.5 Total scattered flux

From a separate experimental dataset where a density scan was performed over a larger density range, the integrated scattered light at wavelengths longer than the laser wavelength is plotted as a function of density in Fig. 4.10.

On logarithmic axes the majority of the data are observed to lie on a straight line, and therefore the relation between them expressible as a power law. In this case the best fit is to $\approx n_e^{2.8}$ or $\omega_p^{5.6}$. At $n_e = 4 \times 10^{19}\,\mathrm{cm}^{-3}$ and a growth rate $\approx \omega_p$ there are tens of e-folding lengths in 20 µm (the size of the focal spot), and so it is reasonable to expect the instability to saturate. The total scattered power at saturation is theoretically expected to scale as[16] $\omega_p^{8/3}$, which is a significantly weaker scaling than observed here.

One could also take a saturation threshold at fixed $\delta n_e/n_e$, where it is assumed the plasma wave can grow no further. If the scattered light has normalised vector potential a_s, then $\delta n_e/n_e = a_0 a_s n_c/n_e = \text{const.}$ implies $a_s^2 \propto \omega_p^4$. This is a faster scaling with density, but still slower than that observed. It could be that the time over which the laser pulse is scattering is longer at higher densities, causing the scaling of scattered *energy* with density to be quicker than the scattered *power*.

A final point to take into consideration is the existence of significant levels of blueshifted scattered light at high densities. At high densities the spectrum of this light is much

[15] [282] M. N. Rosenbluth. *Physical Review Letters*. (1972).

[16] [283] E. Esarey *et al. Physical Review A*. (1992); [211] G. Shvets *et al. Physics of Plasmas*. (1997).

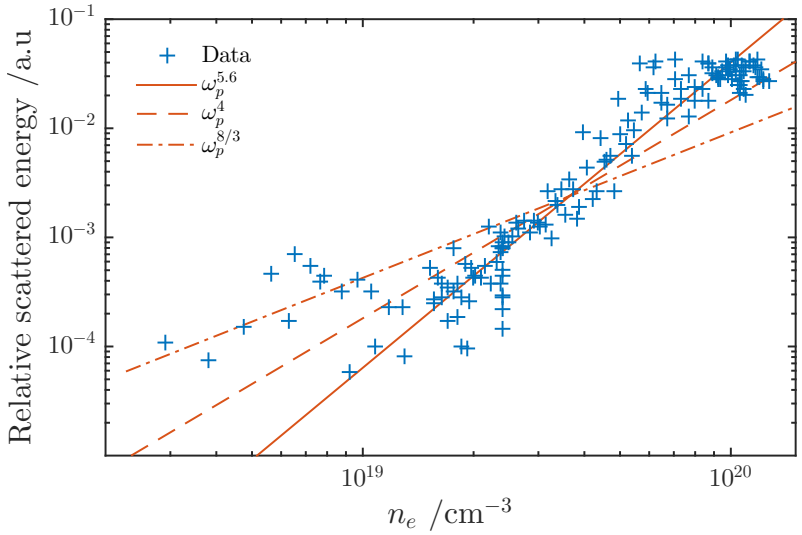


Figure 4.10: The integrated energy of the Raman scattered light as a function of plasma density. The solid line is an empirical power law scaling $\omega_p^{5.6}$, and the dashed lines those of theory.

broader than that of the redshifted light, consisting of a smooth continuum down to progressively shorter wavelengths rather than a well-defined peak. Such broad spectra have been previously observed by Darrow *et al.*[17] and attributed to the strongly-coupled instability. The authors also note that in the case that the scattered modes are far from the normal modes of the plasma, a coupling must occur between the region where the laser intensity is high (non-normal modes) and where it is low (normal modes). The time-varying boundary conditions on the scattered light due to the motion of this region acts to increase the blueshift further. If a broad, smooth spectrum is indicative of an increase in the number of excited modes with density, then the total amount of scattered light might be expected to increase more quickly with density than the single-mode scalings above.

4.1.6 EPOCH simulation

To examine the Raman scattering process further, a 2D EPOCH simulation was performed using the same parameters as the experiment. A gaussian laser pulse of spatial and temporal intensity FWHM 20 µm and 40 fs respectively was focussed into a plasma of peak density 2×10^{19} cm^{-3}. The plasma density profile was trapezoidal with density ramps 300 µm long and a total length of 2 mm, and the vacuum focus of the laser was placed at the peak of the first density ramp. The plasma electrons were initiated at zero temperature with 2 particles per cell. The spatial resolution was 30 cells per laser wavelength in the longitudinal direction and 15 cells per wavelength in the transverse direction. The laser pulse was polarised out of the plane of the simulation box – this was to ensure the observable Raman scatter would propagate out of the laser polarisation plane as in the experiment. The simulation window was 110×140 µm in the longitudinal and transverse directions respectively, and moved along with the laser pulse at the linear group velocity.

[17] [278] C. B. Darrow *et al. Physical Review Letters.* (1992).

The simulation boundaries were open to both particles and fields, the particles were represented as a first-order (triangular) shape function with variable weighting, and other physical effects such as ionisation and Coulomb collisions were ignored.

Simulation frames

Electron density maps are plotted in Fig. 4.11 from four points in the simulation - the x-axes mark the overall progress through the 2000 µm-long simulation domain. Overlaid on the density map is a contour plot marking the shape of the laser pulse which propagates in the positive x-direction. The oscillations in the electric field of the pulse were removed with a Hilbert transform filter and isocontours generated for $a_0 = 0.5$, 1 and 2 (not accounting for frequency shifts).

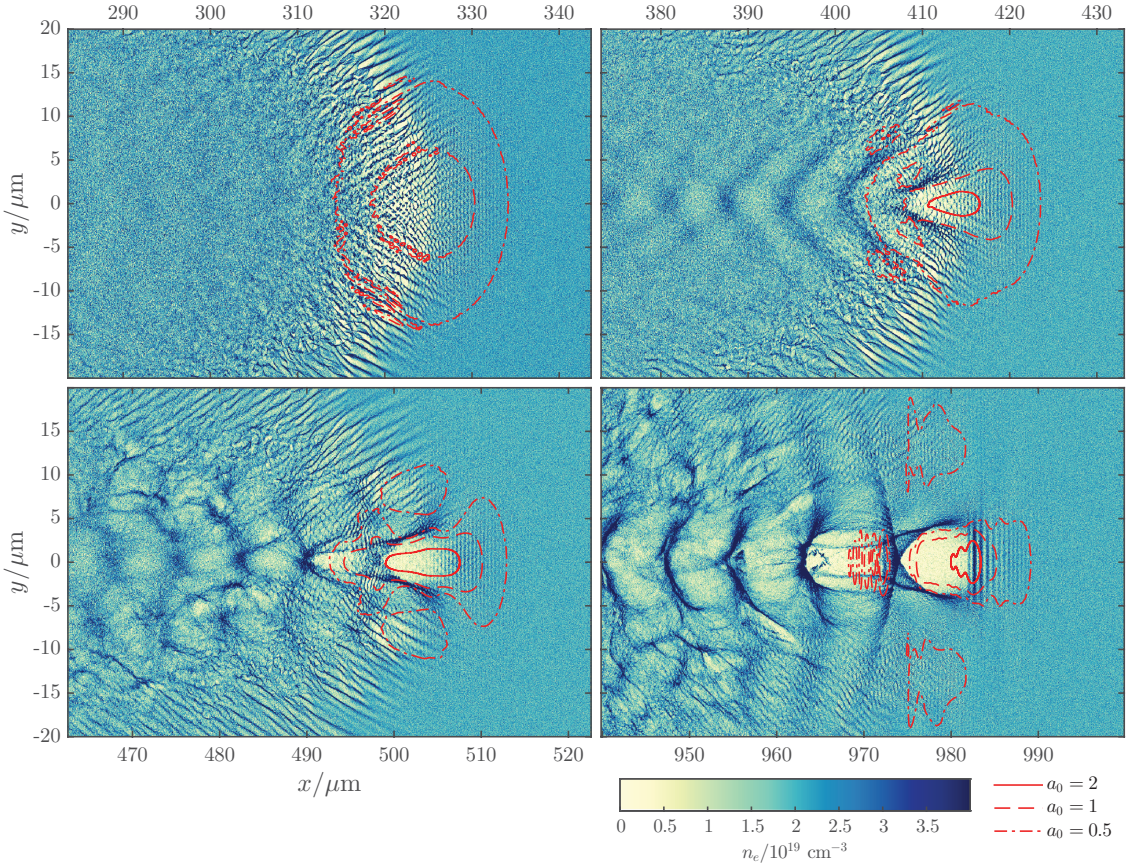


Figure 4.11: Electron density plots from four points in a 2D EPOCH simulation. Isocontours marking the normalised vector potential of the laser are overlaid in red. Raman-scattered plasma waves are observed to be associated with the self-focussing of the pulse.

Shortly after the pulse enters the plasma the electron density distribution shows obvious signs of Raman-generated plasma waves, generated at an oblique angle with respect to the drive pulse direction. As the simulation window is moving in a gallilean sense the apparent direction of the plasma waves in the simulation box is not the same as their direction in the laboratory frame. In the early stages of the pulse evolution the plasma remains fairly uniform, as the ponderomotive force of the pulse is insufficient to expel the electrons from the focal volume. The pulse soon begins to self-focus from the rear,

with electron waves continuing to be generated from near the peak of the pulse. By the third frame the peak intensity of the pulse is large inside the bubble, but has developed lobes of lower intensity outside the electron sheath. Plasma waves continue to originate from this region until the fourth frame when the pulse is mostly contained within the now fully-cavitated bubble, and the emission of scatter ceases.

Plasma waves

To prove that these electron waves are indeed a signature of Raman scattering, the 2D Fourier transform of the electron density may be examined. From Eq. 2.127 in the low density limit, for k-matching the electron waves must satisfy $k_x \approx 2k_0 \cos^2\theta$ and $k_y \approx 2k_0 \sin\theta\cos\theta$. It can be verified that in this case $(k_x - k_0)^2 + k_y^2 = k_0^2$, which is the equation for a circle in the Fourier plane of radius k_0 centred at $(k_0, 0)$. Fig. 4.12 displays the Fourier transform of electron density from two temporal snapshots of the simulation with this circle overlaid. The k-spectrum of the electron density clearly has components which obey this relation, indicating that Raman-scattering is indeed responsible for their generation

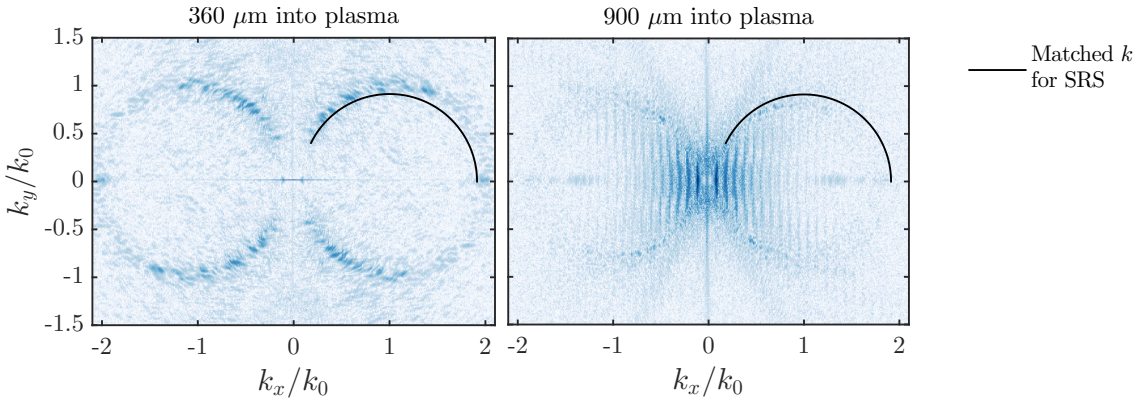


Figure 4.12: The 2D Fourier transforms of electron density are plotted for the cases when the laser pulse has just entered and is midway through the plasma. The semicircle marks the k-matching constraint of Eq. 2.127.

The other four quadrants of the Fourier plane follow because n_e is real and the simulation is symmetric about $y = 0$, and thus two characteristic circles appear. The circle is not complete, but has a gap for small k near the origin. This is because Eq. 2.127 does not have real solutions for $k < \sqrt{2\omega_0\omega_p}/c$, and this cutoff is reproduced well in the simulation output. While prominent early in the simulation, at a propagation distance of 900 µm the circles have almost vanished which indicates the disappearance of k-matched electron waves. In their place are many harmonics of k_p in the x-direction, generated by the large-amplitude plasma wave. This also corresponds with the cessation of the obvious electron plasma waves emanating from the front of the bubble, again supporting the conclusion that these waves are due to Raman scattering.

Pulse evolution

The properties of the drive pulse are simple to extract as it is solely responsible for any out-of-plane E_z fields in the simulation. Plotted in Fig. 4.13 are several quantities

important in the evolution of the pulse – the peak electric field of the pulse E_z, the rate of change of laser energy (integrated E_z^2) in the simulation domain, the integrated ratio of transverse to longitudinal electromagnetic flux, and the FWHM pulse duration. The energy loss curve is shifted by 70 µm backwards with respect to the other curves to account for the time it takes scattered energy to propagate transversely out of the simulation box. The fact that the initial peaks in the transverse flux fraction and energy loss rate match is a *post hoc* justification that the initial energy loss is due to transverse propagation of electromagnetic radiation out of the simulation domain.

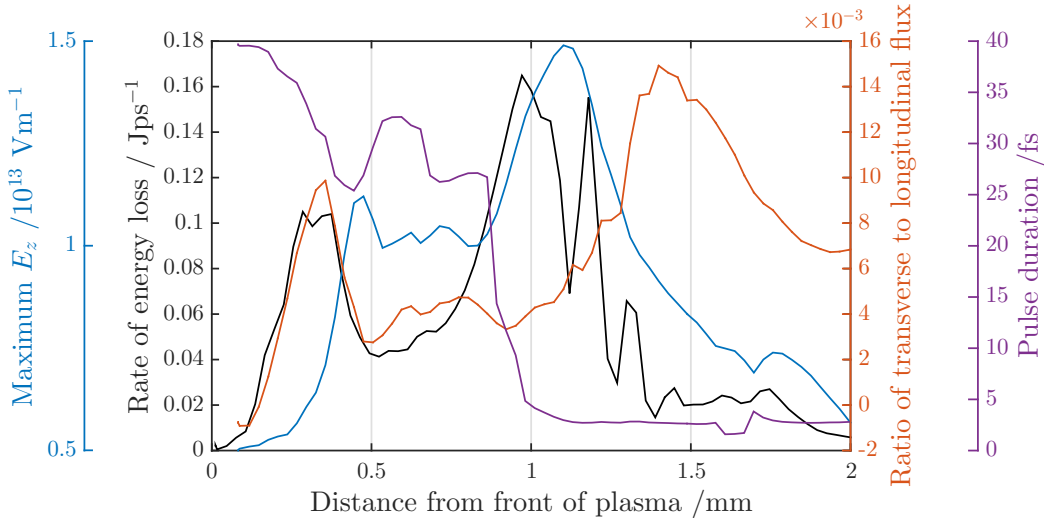


Figure 4.13: Laser parameters are plotted as a function of propagation distance: peak vector potential, energy loss rate, fractional transverse flux and FWHM pulse duration.

As soon as the pulse enters the plasma it begins to lose energy, even when the density is fairly low. This energy loss is well correlated with a spike in the level of transverse flux leaving the pulse, and is in the region where Raman-generated electron waves are observed. This energy loss is therefore attributed to the Raman side-scattering of pulse energy, which occurs predominantly in the first 500 µm of the plasma. The scattering originates preferentially from the rear half of the pulse, which contributes towards the initial burst of self-compression down to ≈ 25 fs. The conclusion of side scattering is accompanied by an initial period of self-focussing where the peak electric field of the pulse is boosted to several times its vacuum value. Soon afterwards the pulse begins to self-compress dramatically which boosts the intensity further at the expense of a rapid drop in pulse energy - at this point the pulse is driving the wake in the bubble regime. Through depletion the power of the pulse drops until it can no longer relativistically self-guide, whereupon diffraction sets in. The signature of diffraction is an increase in the transverse flux fraction without an associated increase in the rate of pulse energy loss. The pulse remains short as it exits the plasma but continues to expand, and its peak electric field falls monotonically.

Side scatter

The scattered radiation at 90° is of much lower intensity than the drive pulse, and is most easily distinguished in Fourier space. Here the drive pulse takes the form of two

peaks along the $k_y = 0$ axis, with the scattered radiation lying approximately on a circle centred at the origin. To create Fig. 4.14, for each simulation output the E_z field is 2D-Fourier transformed and regions outside $85° < \phi = \tan^{-1}(k_y/k_x) < 95°$ set to zero. The remaining elements are then binned over $k = \sqrt{k_x^2 + k_y^2}$. To convert k back to a wavelength as detected on the spectrometer, it is assumed that the radiation in the simulation box satisfies the dispersion relation in a plasma such that its frequency ω_s is related to k by $\omega_s^2 = \omega_p^2 + c^2k^2$. If it is also assumed that ω_s remains constant as the radiation propagates out of the plasma (neglecting the blueshift theorised by Darrow *et al.*[18]), then the detected wavelength is related to the scattered frequency by the vacuum dispersion relation $\lambda_s = 2\pi c/\omega_s$. A relativistic correction to the plasma frequency in the dispersion relation is not required as the radiation sampled from the simulation lies far outside the drive pulse (though the frequency shift itself is subject to relativistic effects).

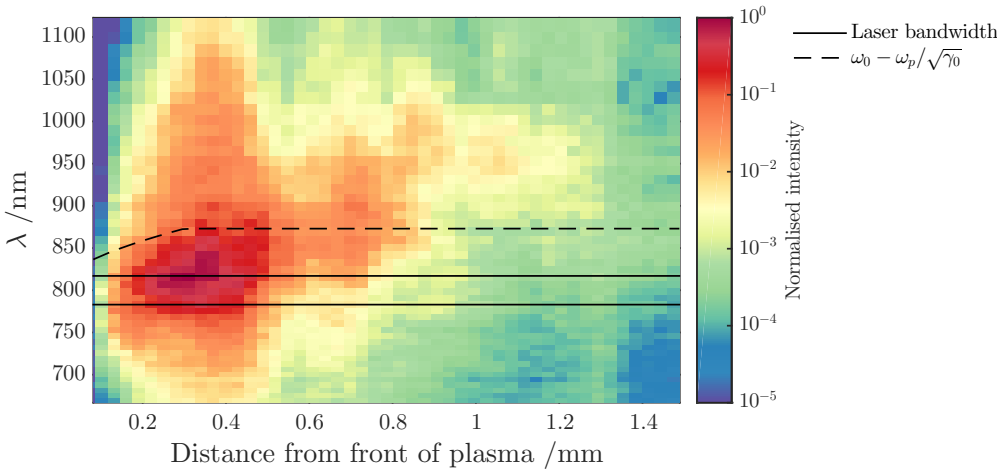


Figure 4.14: Spectra of z-polarised radiation propagating at an angle of $90 \pm 5°$ in the simulation box as a function of drive pulse propagation distance. The bandwidth of the drive pulse and expected wavelength of scattered light at this density are overlaid.

The initial scattering event apparent in Fig. 4.13 is observed here to be composed of a broad spectrum, extending down to 700 nm and up to 1100 nm. This lies underneath a narrow peak, which is only slightly shifted from the laser wavelength. The region between the initial scattering and the self-compression of the drive pulse is associated with the lower level scattering of narrower bandwidth radiation. The initial broad and secondary narrow bandwidth events are approximately centred around the expected shifted wavelength, but there is a significant spread in wavelength. During self-compression the scattering rapidly diminishes in magnitude, fading into noise beyond a propagation distance of ≈ 1.2 mm.

4.1.7 Discussion

The initial purpose of the Raman scattering spectrometer in this experiment was to act as a simple and accurate density diagnostic, with the plasma density encoded in the (assumed linear) frequency shifts of the scattered light. The fact that the measured shifts departed from those expected limits the utility of Raman measurements as a density diagnostic in a

[18] [278] C. B. Darrow *et al. Physical Review Letters.* (1992).

high intensity laser plasma experiment. On the other hand, the imprinting of a systematic frequency shift on the scattered light is interesting, and suggests that some information may be recovered from the spectra about the region of the plasma near the drive pulse. The lack of a strong correlation in the data between frequency shift and bandwidth is an indicator that the growth processes taking place are more complex than linear theory would suggest.

Indeed, from 2D PIC simulation the Raman scattering process is observed to evolve in a complex fashion, tightly coupled with the laser pulse evolution and wake generation. When the pulse is larger growth may occur convectively over a longer length scale, but scattering is observed to etch the rear of the pulse. As the pulse self-focusses and self-compresses scattering is limited to the outer regions of the pulse, the strong density inhomogeneities of the bubble inhibiting propagation of the instability. This is consistent with the experimental observation that the frequency shifts are associated with an effective $a_0 \approx 1.1$. The scatter does not originate from the peak of the self-focussed pulse, rather the outer low-intensity wings around the bubble. In this way the scattering is assisting in the shaping of the pulse, removing light where it is not guided by the plasma density gradients.

The light scattered from the simulation appears to be composed of two distinct events - the initial burst accompanying the self-focussing of the pulse, and a second intermediate burst of lower bandwidth. The fact that different bandwidths are observed during the simulation is interesting, as it may explain the surprisingly large variation in bandwidths in the experimental data. If the initial burst is seeded by the transverse k-vectors introduced during self-focussing, then it can be considered a forced process for which growth can occur far from resonance and the bandwidth is increased. At the conclusion of self-focussing the scatter reverts to a 'natural' unforced state, and the frequency shift and bandwidth return to that expected. Though a simulation parameter scan was not undertaken, it is likely that the relative importance of these two scattering phases can fluctuate significantly depending on the plasma density profile and laser properties. It seems possible then that from shot-to-shot in the experimental data, the observed bandwidth and frequency shift could also vary widely depending on the details of the laser pulse evolution.

4.2 Direct imaging of scatter

The second experiment was performed in 2012 over several weeks on the 200 TW Astra-Gemini laser, with the aim of investigating the LWFA trapping and injection processes. An extensive suite of diagnostics were fielded, including a transverse probing line making use of the second Gemini beam. Fig. 4.15 contains a schematic of the beamline layout for the experiment.

Both beams had a pulse duration of 44 ± 3 fs and were vertically polarised (out of the plane in Fig. 4.15). The drive pulse was focussed with a 3 m focal length f/20 parabolic mirror to a spot of FWHM size 19×21 µm and $33 \pm 2\%$ of the pulse energy was contained within the FWHM contour. The pulse energy was 13.1 ± 1.4 J generating a peak intensity of $4.1 \pm 0.5 \times 10^{19}$ Wcm^{-2} and normalised vector potential of $a_0 = 4.4 \pm 0.3$. An adaptive optic was placed in the drive beam and the wavefront optimised in a closed loop process,

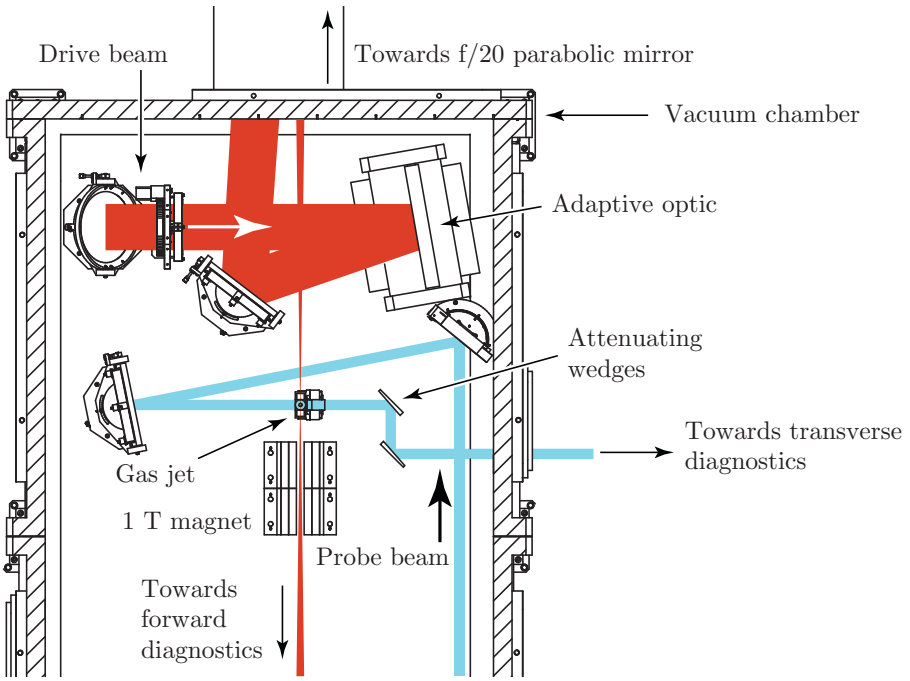


Figure 4.15: Plan view of the beamline layout for the 2012 Gemini experiment.

where the effects of the adaptive optic on the wavefront were measured and the wavefront distortions minimised. The probe beam was derived from the second Gemini beam, apertured to a 5 cm-diameter circular profile and with a lower energy than the drive beam. The plasma target here was a helium supersonic gas jet of diameter 10 mm or 15 mm, and the energy of the electron beams produced measured with a magnetic spectrometer. The gas jet was imaged from above with a high resolution optical camera ('top view'), and from the side with the probe beam. The probe was attenuated using a pair of wedged reflectors, each reflecting a small fraction of the pulse energy. The gas jet was then imaged with an f/11 achromatic lens onto shadowgraphy and interferometry cameras outside of the vacuum chamber. The results discussed in this section are taken from the top view and shadowgraphic diagnostics.

4.2.1 Shadowgrams of plasma channels

Plotted in Fig. 4.16 are a series of shadowgrams taken from the initial stages of the experiment, during one of a number of scans over various experimental parameters in order to optimise the electron beam charge and energy. At this point the interferometer was inoperative so shadowgraphy was the primary transverse diagnostic. In these images the drive beam propagates from left to right and the probe beam is propagating out of the page. Here the plasma target is a 10 mm diameter supersonic helium gas jet, and the drive pulse is focussed near the leading edge of the gas. As the pulse propagates to the right the low-intensity unguided wings of the pulse continue to expand in an $f/20$ cone, while the high-intensity portion of the pulse remains guided. The presence of plasma is signified in the shadowgrams by modulations of the probe intensity. Probe light is deflected outwards by the plasma, causing a darkening of the image inside the channel.
Smaller deflection angles cause local changes in image intensity, typically due to smaller

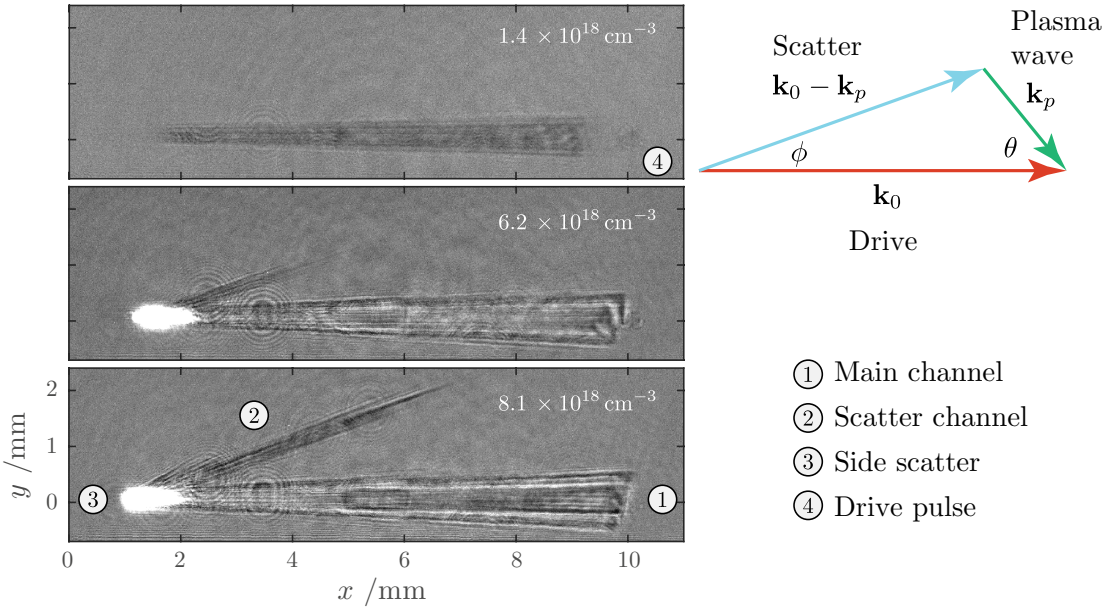


Figure 4.16: Plasma channels as observed on the shadowgraphy diagnostic. The shadowgrams here are divided by a reference shadowgram with no plasma present.

scale filamentary structure. Imperfections in the intensity profile of the probe beam were removed by dividing by a reference image where no plasma was present. However parts of the probe beam near to the channel suffer an average deflection and so the imperfections are translated, causing the background correction to fail. The rings near the channel do not have any physical significance, and are only uncorrected imperfections due to dust on the imaging optics. In this initial stage of the experiment a faulty gas nozzle was used, and the structure visible between $x = 5$ and 6 mm is due to a shock in the gas flow. Finally, it is interesting to note that the drive pulse is often visible as a modulation in the probe beam intensity, even when the two beams cross in the very tenuous plasma near the edge of the jet. This observation foreshadows the results of chapter 5 where the two-pulse interaction becomes important.

At the beginning of the experiment, it was frequently observed that significant secondary channels were generated. These secondary channels originated from the beginning of the plasma, where the drive pulse first encounters the gas. This point is also associated with the production of side-scattered radiation, which appears as a (time-integrated) bright spot on the shadowgram. From our own previous experiments and elsewhere in the literature[19] this spot is known to be due to Raman side scatter, which given the size of the collection optics, must be propagating at an angle of $90 \pm 2.5°$ to the drive pulse. The total scattered energy was observed to increase with plasma density, as plotted in Fig. 4.17.

The frequency of occurrence and size of the secondary channels also increases with density, as is clear from Fig. 4.16. The natural hypothesis to make is that the secondary channel is also due to the Raman scattering of the drive pulse, given that it originates from a position known to host the generation of Raman side scatter. This phenomenon has been

[19] [284] A. G. R. Thomas *et al. Physical Review Letters.* (2007).

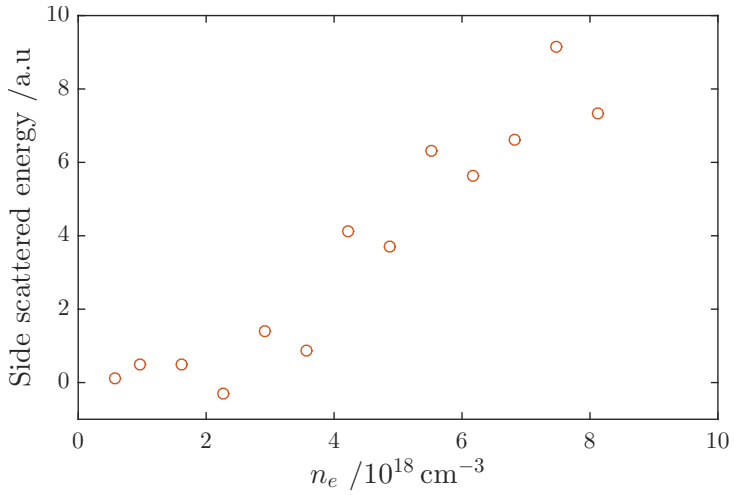


Figure 4.17: Integrated side-scattered energy as a function of plasma density from the 10 mm nozzle.

observed previously by Matsuoka *et al.*[20], where the same conclusion was reached.

4.2.2 Angle of scatter

The secondary channels appear at an angle to the drive pulse propagation direction which can be directly measured from the shadowgrams. If caused by Raman scattering, the angle of the scattered light should be a function of density. To calculate the expected angle, recall Eq. 2.127 which relates the magnitude of the plasma wave vector k to its direction θ assuming that the Raman-scattered wave satisfies the dispersion relation for an electromagnetic wave in a plasma

$$k = k_0 \cos\theta \pm k_0 \sqrt{\cos^2\theta - 2\left(\frac{\omega_p}{\omega_0} - \frac{\omega_p^2}{2\omega_0^2}\right)}. \tag{4.6}$$

Letting $\mathbf{k}_s = \mathbf{k}_0 - \mathbf{k}$ represent the k-vector of the scattered light,

$$k_{s\parallel} = k_0 - k\cos\theta \tag{4.7}$$

$$k_{s\perp} = k\sin\theta \tag{4.8}$$

where $k_{s\parallel}$ and $k_{s\perp}$ are the components of $\mathbf{k}_s$ parallel and perpendicular to $\mathbf{k}_0$ respectively. The scattering angle ϕ between $\mathbf{k}_0$ and $\mathbf{k}_s$ is then

$$\phi = \arctan\left(\frac{|k_{s\perp}|}{|k_{s\parallel}|}\right) = \arctan\left(\frac{k\sin\theta}{k_0 - k\cos\theta}\right). \tag{4.9}$$

The direction of the scattered light is fixed once the direction of the plasma wave is known. In general the observed scattered mode will be that with the largest growth rate, as it will exponentially dominate other excited modes. As Raman growth rates are approximately proportional to k, the observed scatter angle ϕ should correspond to the value of θ which

[20] [197] T. Matsuoka *et al. Physical Review Letters.* (2010).

maximises Eq. 4.6. Some care must be taken here as there are two solutions for the plasma wavevector magnitude for a given angle θ – see Fig. 4.18.

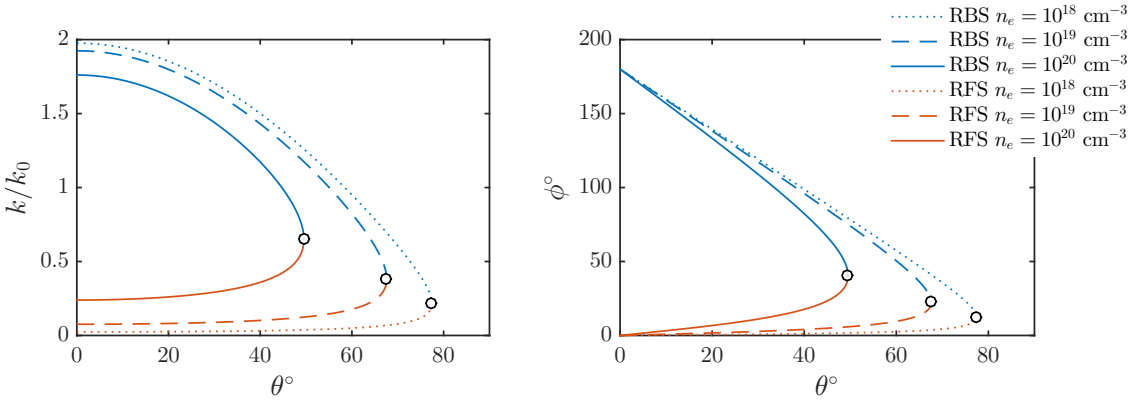


Figure 4.18: $k(\theta)$ from Eq. 4.6 and $\phi(\theta)$ from Eq. 4.9. Blue/red lines correspond to the $+/-$ signs in Eq. 4.6 respectively, which correspond to RBS/RFS. Circles mark the maximum growth mode for RFS.

The backwards-scattering solution corresponding to greater ϕ is notated as RBS (Raman back scatter) and the forward-scattering solution corresponding to smaller ϕ as RFS (Raman forward scatter). The RBS solutions are observed to peak in k at $\theta = 0°$ and $\phi = 180°$, i.e. a forward-going plasma wave of short wavelength generating pure back scatter. The RFS solutions are observed to peak at different angles depending on the plasma density, which occurs when the radical of Eq. 4.6 becomes zero. At this point assuming $\omega_p \ll \omega_0$, $\cos^2\theta \approx 2\omega_p/\omega_0$ and $k = k_0 \cos\theta$. Following some manipulation, the scatter angle is

$$\sin\phi = \left(\frac{2\omega_p}{\omega_0} + \frac{\omega_p^2}{\omega_0^2} - \frac{2\omega_p^3}{\omega_0^3} + \frac{\omega_p^4}{\omega_0^4}\right)^{1/2} \approx \left(\frac{4n_e}{n_c}\right)^{1/4} \tag{4.10}$$

where higher order terms are neglected due to the low plasma density. To account for relativistic effects the change to make here is to use the relativistic dispersion relation for an electromagnetic wave in a plasma, which amounts to replacing ω_p with $\omega_p/\sqrt{\gamma_0}$. Expression 4.10 is plotted in Fig. 4.19 assuming Lorentz factors corresponding to $a_0 = 0$, 0.5 and 1, along with experimental data from a range of plasma densities.

The data were taken from shadowgrams recorded on two different days of the experiment, the 10 mm nozzle data on the same day as the shadowgrams above, and the 15 mm nozzle data on the next day. At the backing pressures used in these experimental runs the density ranges explored with the two nozzles do not quite overlap. The correspondence between data and theory is good, with Eq. 4.10 predicting reasonably similar scattering angles to those observed over a wide range of plasma densities. Relativistic corrections to the expected scattering angles are small, but it appears that the electrons taking part in the scattering plasma wave feel little influence of the drive pulse fields. Inserting $a_0 > 1$ into Eq. 4.10 generates predicted angles which are too low, whereas $0 < a_0 \lesssim 0.5$ is more plausible depending on the plasma density. This fact suggests that the scattering occurs away from the peak of the pulse, in the wings outside of the bubble as was also suggested

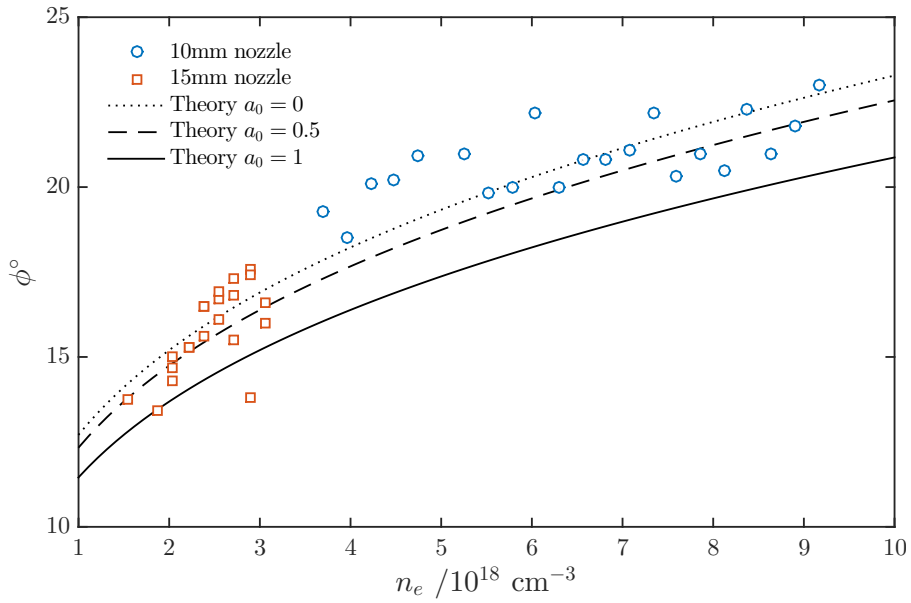


Figure 4.19: Measured scatter angles from transverse imaging of the two nozzles, overlaid with the theoretical predictions of Eq. 4.10.

in the previous section.

Filaments may be produced at the same ϕ but not in the vertical plane, which would cause the filament to appear to have smaller ϕ in the shadowgram. It would be expected then that a range of ϕ would be measured for each density, with a clear maximum. This was not observed, which may be due to the difficulty of observing filaments travelling towards the transverse diagnostics, or due to an asymmetry in the drive pulse which caused filaments in the vertical plane to be preferential.

4.2.3 Scatter as a function of focal position

Later in the experiment the very large secondary channels stopped appearing, likely due to further optimisation of the focal spot shape and gas target. However much smaller filaments were often observed – see Fig. 4.20 – and it was assumed that these filaments were caused by a lower level manifestation of the same mechanism as before.

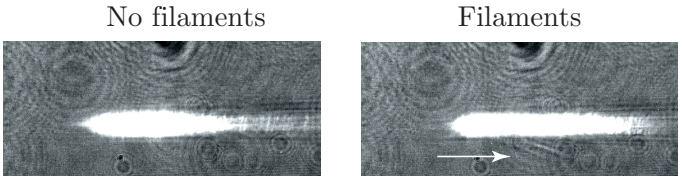


Figure 4.20: Small filaments were often observed leaving the side-scatter region.

The angles of the smaller filaments were similar to the larger ones, 10° to 30° with respect to the drive beam axis, and continued to originate from the side scatter region. It was observed that the presence of filaments was correlated with the position of the focal spot relative to the leading edge of the gas jet. A scan of this position was performed later in the experiment in an effort to optimise the electron beam quality, and the presence of filaments was confirmed or not for each shot using the shadowgraphy diagnostic. For

each focal position a number of shots were taken, and the number of shots displaying obvious filaments noted. The interesting quantity here is the probability of observing filaments, which must be estimated from the small dataset for each position. If N shots were recorded with n of them displaying filaments, then the inferred filament production probability f is distributed as[21]

$$P(f) \propto f^n (1-f)^{N-n} . \tag{4.11}$$

This peak of this distribution is the maximum likelihood estimate for f, trivially $f_{\mathrm{MLE}} = n/N$. The width of the distribution quantifies the uncertainty on f_{MLE} – even if five shots were taken and no filaments observed, there is a finite probability that the 'real' chance of filament production was higher and the shots were 'unlucky'. It should also be noted that filaments produced in the transverse direction are invisible to this diagnostic, so the probabilities reported here are likely underestimates.

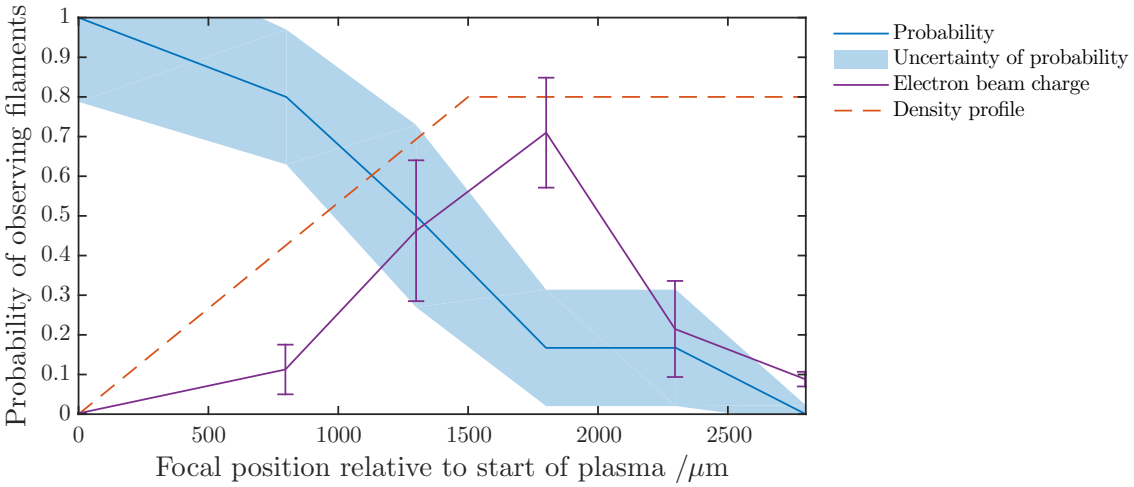


Figure 4.21: The probability of observing filaments as a function of focal spot position is plotted as a shaded area. Also overlaid are the relative electron beam charge and approximate density profile.

Plotted in Fig. 4.21 are the peaks and ± 1 standard deviation widths of Eq. 4.11 for the experimental data as a function of focal position relative to the beginning of the gas jet. Also plotted are the average and ± 1 standard deviation electron beam charges as observed on the electron spectrometer, and the approximate gas density profile. When operating the accelerator under optimal conditions, the laser focal position is situated some distance inside the plasma, often between 1 and 2 mm into the gas jet. It is obvious in Figs. 4.21 and 4.22 that in these regions the electron beam charge is significantly higher and beams are produced more consistently. If the pulse is not correctly matched to the density profile which it is propagating through due to an incorrect focal plane position, there must be a rapid relocation and redirection of pulse energy before matching can occur. As suggested in the previous section this likely acts to seed any of the high growth rate transverse Raman modes, and lead to the formation of scattered filaments.

[21] This is a form of the beta distribution, the Bayesian conjugate prior for this binomial process.

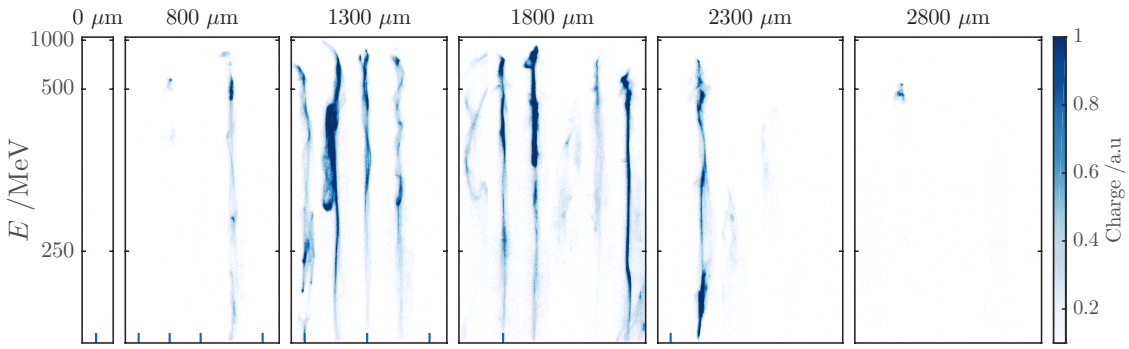


Figure 4.22: Raw spectrometer images showing electron beams observed on consecutive shots as the focal position is varied with respect to the leading edge of the gas jet. Small dashes at the bottom indicate the presence of filaments for each shot.

If the focal position is translated near to the front edge of the jet such that the pulse is re-expanding as it enters the bulk of the plasma, then filaments are observed on nearly every shot. As the focal position is translated into the plasma the occurrence of filaments falls monotonically and the electron beam charge begins to rise. Too far and the acceleration of electron beams begins to be suppressed, but the filaments continue to be absent. The side scatter region does not move longitudinally as a function of focal spot position, suggesting that the location of the Raman side scatter is linked to the plasma density profile.

4.2.4 Scatter as a function of focal spot symmetry

One of the goals of this experiment was to investigate the effects of focal spot asymmetries on the electron acceleration and x-ray production processes, previously shown to be important[22]. Here we focus instead on how the production of scattered filaments changes as the focal spot is altered. The aim was to produce a significant asymmetry of the focal spot which remained as the beam travelled through focus, thereby stimulating the preferential transverse injection of electrons along one axis. Recalling Fig. 2.2, of the low-order aberrations coma has this property – spherical aberration is symmetric and astigmatism rotates through 90° about focus. The effect of adding half a wavelength of coma to the focal spot is displayed in Fig. 4.23. Here the aberration was introduced with the adaptive optic, the effect of which on the wavefront was previously characterised in terms of the Zernike components of the beam phase profile. It is clear that the addition of this level of aberration has a significant impact on the peak intensity of the pulse, which falls by a factor of two. Significant lobes are produced preferentially on one side of the focal spot, the location of which depends on the direction of the added coma.

A series of scans over plasma density were performed using the 15 mm nozzle, reaching densities of $1.6 - 3.6 \times 10^{18}\,\mathrm{cm}^{-3}$. The focal spot was either left unchanged, or corrupted with the addition of 0° and 90° coma. With five shots recorded at each density and the presence of filaments confirmed using the shadowgraphy diagnostic, a probability of filament formation was determined as described above.

[22] [118] S. P. D. Mangles *et al. Applied Physics Letters.* (2009).

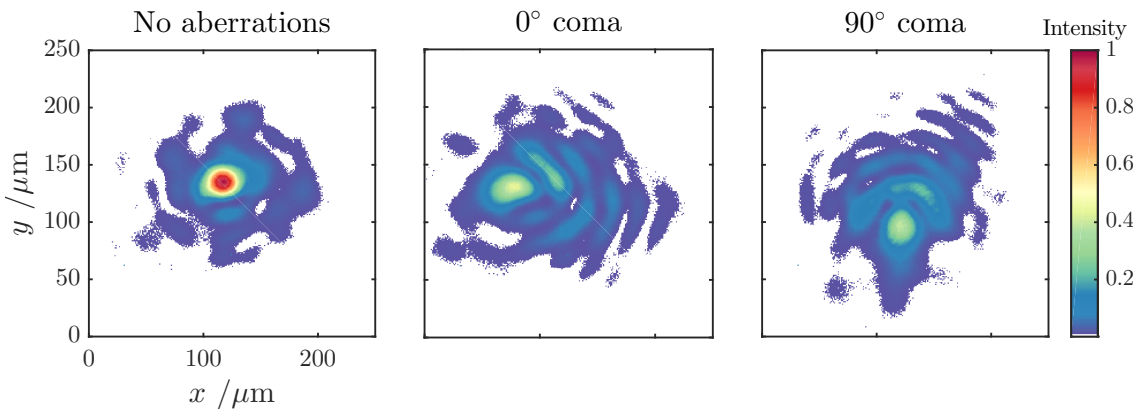


Figure 4.23: Focal spots intentionally aberrated by adding $\lambda/2$ of coma, where the amplitude of the aberration is that defined in Eq. 2.10.

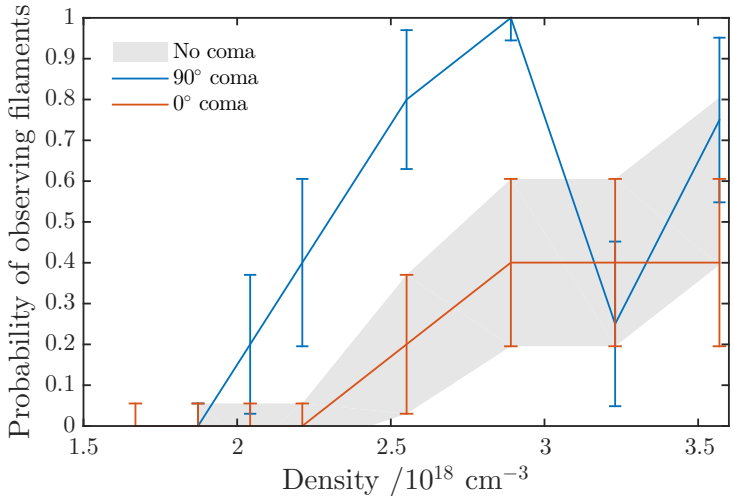


Figure 4.24: The inferred probability of filament production as a function of plasma density and focal spot quality.

The data are plotted in Fig. 4.24, with the results of the unaberrated spot represented by the shaded area. The addition of 0° coma is not observed to alter the probability of filament production, lying on almost the same curve as the unaberrated spot. Adding 90° coma on the other hand significantly increases the chances of producing filaments across most of the densities measured. The shadowgraphy diagnostic is only sensitive to those filaments produced in the vertical direction, and the fact that 90° coma causes laser energy to be displaced above the spot suggests that this asymmetry is a seed for the scattering process. If the transverse scattering and filament production is dominated by the unguided laser energy, then an increase in the unguided portion of the spot should be expected to enhance this process. It is possible that the presence of 0° coma caused filament production in the orthogonal plane, which is invisible on the transverse diagnostics.

In addition the scattered energy at 90° was measured in two orthogonal directions – transversely towards the shadowgraphy camera and upwards towards the top-view camera. Plotted in Fig. 4.25 are the integrated contributions from the bright side-scatter spot on the shadowgraphy and top-view cameras, normalised to the incident laser energy. As was

previously observed for the case of the 10 mm nozzle, in general the amount of scattered energy increases with increasing plasma density.

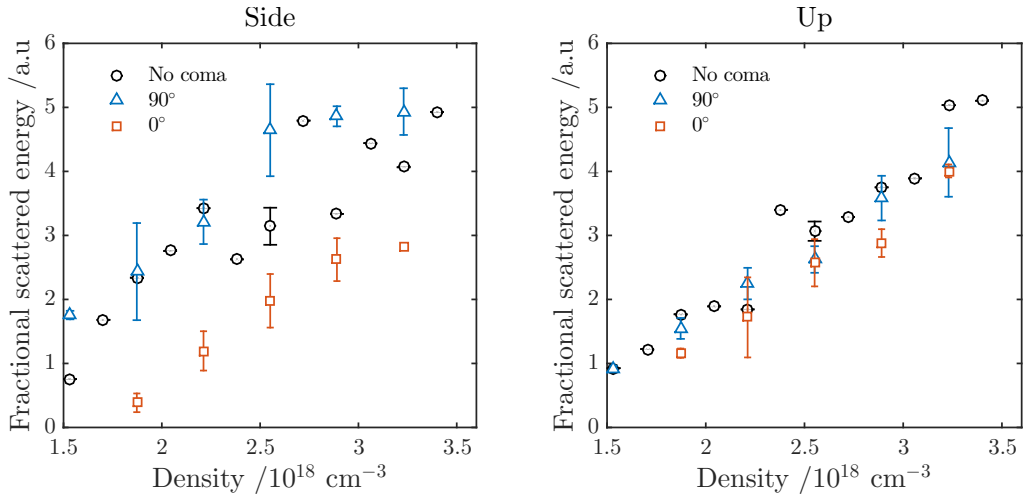


Figure 4.25: Integrated scattered energy in the sidewards and upwards directions as a function of plasma density.

Here there are some surprising differences. Alterations of the focal spot appear to have minimal influence on the amount of energy scattered vertically. For the 0° coma spot there is a slight reduction in scattered energy in two out of five cases compared to the 90° coma spot, but the difference is small. In the transverse direction however, the addition of 0° coma significantly reduces the level of scatter. The experimental observations of this section are summarised in table 4.1.

Feature	Horizontal coma	Vertical coma
Vertical filaments	-	↑
Vertical 90° scatter	-	-
Horizontal filaments	?	?
Horizontal 90° scatter	↓	-

Table 4.1: Change in feature as a function of focal spot shape.

4.2.5 EPOCH simulations

Several simulations were performed to investigate the evolution of the drive pulse under various conditions. In this section the simulated laser pulse energy was 15 J and the peak electron density $4 \times 10^{18}\,\mathrm{cm}^{-3}$ with 1500 µm-long density ramps. The other simulation parameters are the same as in section 4.1.6. Due to the computational cost involved the laser energy and plasma density are slightly higher than the experimental values in order to be sure of the occurrence of Raman scatter. In addition it is well known that in 2D slab geometry the laser intensity must be artificially increased to account for the slower scaling of a_0 with self-focussed spot size[23].

[23] [285] F. S. Tsung *et al. Physical Review Letters.* (2004); [286] C. Kamperidis *et al. Plasma Physics and Controlled Fusion.* (2014).

Focal position scan

Keeping the laser energy fixed, the vacuum focus position was placed at $x = 0\,\mu\text{m}$, $x = 500\,\mu\text{m}$ and $x = 1500\,\mu\text{m}$, where x is the laser propagation direction and $x = 0\,\mu\text{m}$ marks the start of the plasma. The laser is injected into the simulation box in a region of vacuum, and a change in focal position is achieved by adjusting the size of the pulse and the curvature of the wavefronts as appropriate for a gaussian beam.

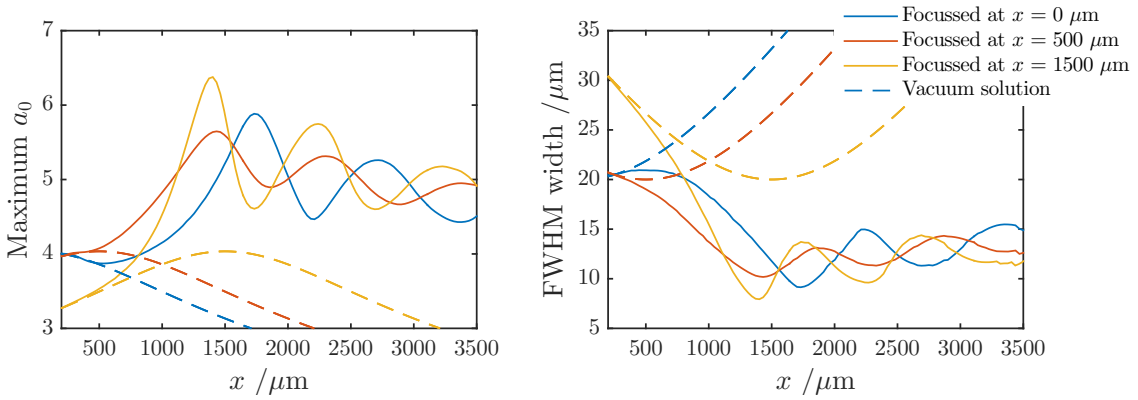


Figure 4.26: The peak vector potentials and transverse widths are plotted as a function of propagation distance for the three cases.

As the pulse enters the plasma it very quickly begins to self-focus. The transverse widths and peak vector potentials of the pulses for the three cases are plotted in Fig. 4.26 as a function of propagation distance through the plasma. The dashed lines mark the gaussian-beam evolution of the pulse in vacuum as appropriate for 2D geometry. In all cases over the first 1.5 mm of propagation the growth in peak a_0 is driven by self-focussing, as the pulse energies and durations change little in this region. The pulse focussed at $x = 0\,\mu\text{m}$ initially begins to diffract before the rising plasma density causes it to be recaptured and focussed further. After over-focussing the pulses experience oscillations in transverse size which slowly damp in amplitude. The oscillation amplitude is observed to be smallest in the intermediate $x = 500\,\mu\text{m}$ case where the pulse is focussed near the middle of the density ramp, implying a smoother transition between vacuum and plasma. This conclusion is reinforced when examining the light scattered from the drive pulse, plotted in Fig. 4.27. Here the $x = 500\,\mu\text{m}$ simulation is observed to produce minimal scattering throughout the self-focussing process. The other simulations generate significantly higher levels of scattered light at different positions in the plasma. Where the pulse is focussed earlier the level of scatter peaks near $x = 1400\,\mu\text{m}$, and as plotted beneath prominent electron density waves are observed at this location. When the pulse is focussed later the scatter peaks later near $x = 1800\,\mu\text{m}$, and is correlated with a maximum in transverse pulse size. Again waves are visible in the electron density map. In both cases the scatter is produced over an extended region across several millimetres of plasma.

Focal spot asymmetry

To mimic the aberrated focal spot produced in the experiment, a simulation was performed where the laser pulse was given an asymmetric profile. As the simulation was two

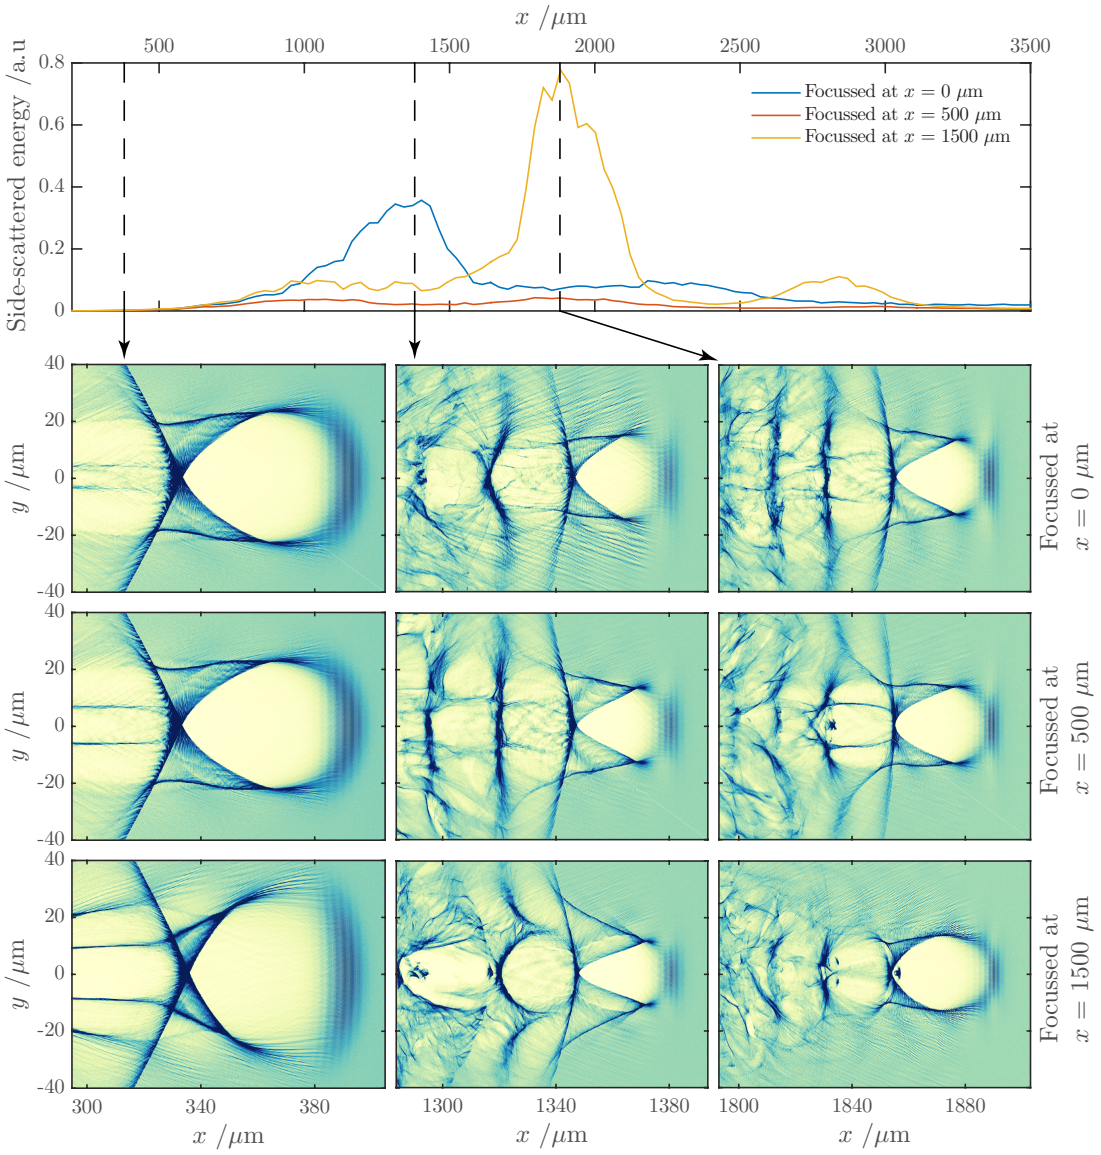


Figure 4.27: The integrated level of scattered light between 750 and 850 nm into angles $90 \pm 10°$ is plotted as a function of propagation distance for the three cases where the drive pulse is focussed (in vacuum) at $x = 0\,\mu$m, $500\,\mu$m and $1500\,\mu$m. Electron density plots for the three cases are displayed below for positions indicated by the vertical black lines.

dimensional it was only possible to add asymmetry in one dimension, which in this case was achieved by forming the pulse as a sum of two gaussian pulses. The first pulse had the same parameters as previous simulations, the second was offset from the $y = 0$ axis by $10\,\mu$m, and had half the transverse size and 1/15 the energy of the first. The total laser energy in the simulation box was kept at the same level as previous simulations. As plotted in the first frame of Fig. 4.28 the pulse energy is biased towards the positive y direction. The electron density is represented by colour with the overlaid laser intensity isocontours marking the shape and location of the pulse. In the first frame the longitudinal ponderomotive force is larger for $y > 0$ and the density cavity formed is noticeably asymmetric. As the pulse progresses through the plasma the wake continues to show

signs of asymmetry, with additional electron sheaths in the second frame and disrupted sheaths in the third. The pulse evolves towards a symmetric state and the wake eventually assumes a symmetric form. In the intervening period however the pulse is observed to 'leak' outside the bubble as in the fourth frame. When this occurs scattering is strongly enhanced, demonstrated by the stronger density modulations for $y > 0$ compared to $y < 0$. The strong scattering depletes the laser energy there, and the scatter eventually ceases.

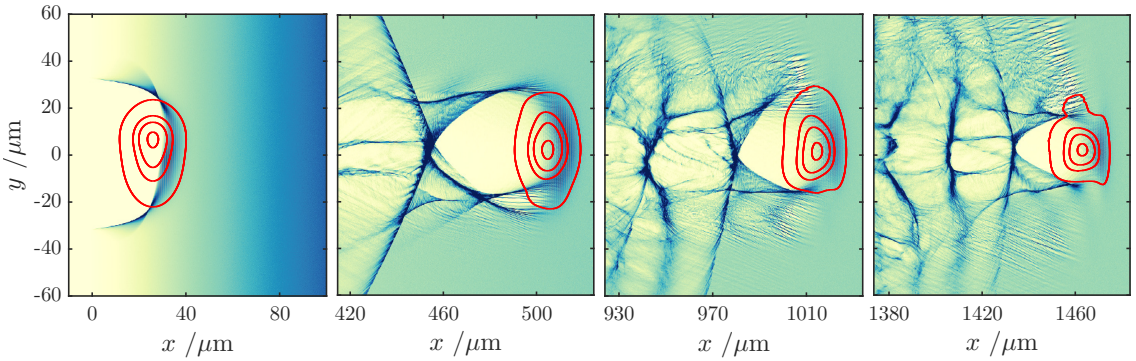


Figure 4.28: Electron density maps taken from four points in the simulation of the asymmetric focal spot. Overlaid in red are laser intensity isocontours at 2%, 25%, 50% and 90% of the peak intensity in a given frame.

4.2.6 Discussion

The plasma filaments under examination here were determined to be consistent with production via Raman scattering of the drive pulse. However unlike Matsuoka *et al.*[24] and the results of section 4.1, a significant relativistic shift is not required in order to fit the data. The common difference between those two experiments and this one is that the plasma densities here are much lower, below $10^{19}\,\mathrm{cm}^{-3}$. As is evident from the PIC simulations the pulse immediately cavitates a large region of plasma behind it, forming sharp electron sheaths in the transverse direction. This is to be contrasted with the pulse evolution in the low intensity case in Fig. 4.11 and Fig. 4 of the Matsuoka paper where the plasma remains relatively uniform near the pulse for significant propagation distances. In this case Raman growth can occur near the peak of the pulse for a limited period of time where the plasma electrons experience strong relativistic mass shifts. In the high intensity bubble regime growth only occurs outside of the bubble where the laser intensity is low. At lower densities then the pulse is less accessible to convective Raman growth and the effective intensity felt by the growing modes is lower.

The level of scatter was observed to increase with density as in section 4.1, but behaved differently when stimulated with focal spot asymmetries depending on whether it was observed at 90° or not. One systematic asymmetry to take into account here is the vertical polarisation direction of the drive pulse. Scatter into the horizontal direction should be preferred as the polarisations of the drive pulse and scattered light are parallel. It is therefore surprising that adding horizontal asymmetry to the focal spot did not stimulate 90° side scattering in the horizontal direction. It is possible that scattering into the

[24] [197] T. Matsuoka *et al. Physical Review Letters.* (2010).

'resonant' direction $\arcsin(\sqrt{2\omega_p/\omega_0})$ was increased at the expense of 90° scatter, though this was unobservable here.

Whether due to an incorrect laser focal position or an asymmetric focal spot, the presence of Raman-scattered filaments was correlated with an imperfection in the wakefield accelerator. Both cases were associated with a decrease in accelerator performance as measured by the charge and peak energy of accelerated electron beams. From PIC simulations it appears that scatter is generated when the drive pulse is forced to evolve rapidly through nonlinear optical processes. If this evolution is smoother then minimal pulse energy is scattered away and the fraction contained within the high-intensity region driving the wake is maximised. The presence or not of significant levels of Raman-generated filaments is therefore a useful diagnostic in the operation of a wakefield accelerator, and provides a simple binary evaluation of the coupling efficiency between the laser pulse energy and the high amplitude plasma wave. One could also imagine that on a high-repetition-rate system the presence of filaments could form part of the feedback into an automatic accelerator optimisation routine, whether a genetic algorithm or other general gradient-descent technique.

4.3 Conclusion

Stimulated Raman scattering in a laser wakefield accelerator is a complex process, accompanied by rapid changes in plasma density profile and laser pulse intensity. Spectroscopic and direct imaging measurements have indicated that under the experimental conditions used to drive a LWFA on Astra-Gemini the bulk of the scattered radiation is produced some distance away from the peak of the drive laser pulse. From PIC simulation this effect is attributed to the presence of the bubble, the strong density gradients of which limit instability growth. Nevertheless, in the high intensity strongly-coupled regime of the instability convective growth occurs rapidly over micron length scales and broad bandwidths, and significant levels of scattered light may leave the pulse. In some cases the scattering is likely to be useful, shaping the pulse and allowing a resonant wake to be driven. In other cases it is a mark of wasted laser energy and indicative of poor laser-plasma coupling. As a density diagnostic its interpretation is more complex than direct interferometry due to the relativistic motion of the plasma electrons. As a wakefield operation diagnostic it is shown to be useful, sensitive to focal spot aberrations and laser focal position offsets of fractions of a Rayleigh range.

It has long been known that laser pulse in a wakefield accelerators will evolve through the processes of self-focusing and pulse compression[25], and recent results have suggested that this evolution phase plays a crucial role in injection[26]. The measurements and simulations presented in this chapter show that the processes of pulse evolution are intimately linked with Raman side scattering. It is also shown that Raman side scattering is a mechanism for removing pulse energy outside of the 'matched conditions', allowing the plasma to shape the pulse and matched conditions to be reached.

[25] [31] V. Malka *et al. Science.* (2002); [32] S. P. D. Mangles *et al. Nature.* (2004).

[26] [158] S. P. D. Mangles *et al. Physical Review ST-AB.* (2011); [152] A. Sävert *et al. Physical Review Letters.* (2015).

5 Interaction Between Crossing Laser Pulses

As discussed in chapters 3 and 4 a common diagnostic of laser wakefield interactions is a transverse optical probe. Such probes are typically much less intense than the wakefield drive pulse, and the parametric interaction between the drive and probe is not usually considered important. Chapter 4 shows that Raman processes, in particular Raman side scatter, occur while the pump pulse evolves via self-focussing and compression, and it is therefore interesting to consider how the simultaneous presence of the drive and the probe modify the parametric processes occurring.

In this chapter experimental and simulation results are discussed on the theme of the parametric interaction between these crossing laser pulses. It was first noticed in a 2011 Astra-Gemini experiment that when the pulses cross near the beginning of the plasma channel, the level of side-scattered light is increased. This phenomenon was investigated during portions of subsequent experiments on Astra-Gemini, the results of which are presented in section 5.1. The interaction of a low intensity ($10^{10}\,\mathrm{Wcm}^{-2}$) and high intensity ($10^{19}\,\mathrm{Wcm}^{-2}$) pulse of equal wavelength is observed to be surprisingly dramatic, with the low-intensity pulse increasing in intensity by a factor of up to 10^5 over tens of microns. No frequency shift of this pulse was observed, and the interaction disappeared when the polarisation of the pulses are made orthogonal to one another.

Section 5.2 discusses a number of PIC simulations performed in an effort to understand the physical basis of the interaction. The plasma electrons are observed, at sufficiently high laser intensities, to become trapped in the wells of a ponderomotive potential generated by the crossing pulses, a mark of entry to the so-called superradiant regime. Energy is transferred to the probe pulse when, by the classical theory of Raman amplification, no net energy transfer would be expected on average due to the violation of the 3-wave matching conditions. In the final section 5.3 these results are compared to various reported Raman amplification experiments in the literature. The amplification rate and penetration into the superradiant regime is found to be the highest observed in the laboratory to date.

5.1 Experimental results

Data from three separate experiments are presented here, all of which took place on the Astra-Gemini laser of the Central Laser Facility. The experiments were all LWFA experiments which, despite having different aims, used the same design for the wakefield acceleration beamline – the Gemini laser pulse was focussed with a 3 m focal length $f/20$

parabolic mirror into a helium gas jet. The experimental layout for the drive pulse was therefore the same as that shown in Fig. 4.15.

Experiment	I_{drive} /Wcm^{-2}	w_{drive} /µm	I_{probe} /Wcm^{-2}	Polarisations	L_{jet} /mm	n_e /cm^{-3}
I: Late 2012	4.4×10^{19}	19×21	6×10^{10}	Parallel	15	3×10^{18}
II: Early 2013	4.1×10^{19}	19×21	6×10^{10}	Orthogonal	15	3×10^{18}
III: Late 2013	3.3×10^{19}	23×28	3×10^{12}	Parallel	5	7×10^{18}

Table 5.1: Summary of parameters for the three experiments – drive intensity, drive FWHM focal spot size, probe intensity, relative polarisation between drive and probe, gas jet length and plasma density.

The important difference here between the experiments was the way in which the probe beam was generated and delivered to the gas jet, summarised in table 5.1. In the first experiment the second Gemini beam was used as a probe, amplified to a lesser extent and therefore at a relatively low intensity. The geometry of the probe beamline caused the polarisations of the two pulses to be parallel and vertical when they met at the gas jet, out of the plane of Fig. 4.15. The second experiment was very similar to the first, except the probe beamline design caused the two pulses to have approximately orthogonal polarisations. Finally, in the third experiment the probe pulse was generated as a pick-off from the drive pulse and so had a much higher intensity, but the polarisations of the two pulses were again parallel. In all cases the duration of the drive and probe pulses were the same as reported in section 4.2 at 44 ± 3 fs

Experiments I and II used as a plasma target a 15 mm diameter gas jet operating at densities near $n_e = 3 \times 10^{18}\,\text{cm}^{-3}$. Experiment III used a shorter 5 mm diameter gas jet at densities near $n_e = 7 \times 10^{18}\,\text{cm}^{-3}$. For both jets the radial density profile was measured to be approximately trapezoidal, comprised of an approximately constant-density plateau with linear gradients to vacuum at the edges. The results here are taken again from the transverse imaging diagnostics. In all experiments the probe beam was imaged onto shadowgraphy and interferometry cameras, and in experiment III also onto a Czerny-Turner imaging spectrometer.

5.1.1 Shadowgraphy

A typical shadowgram from experiment I is displayed in Fig. 5.1a). Here the 15 mm gas jet is viewed from the side and the dark region is the upper end of the nozzle. The drive beam propagates from left to right, generating Raman side scatter when it first enters the plasma as discussed in the previous chapter. The sketches on the right-hand side of Fig. 5.1 illustrate the pulse geometry at the snapshot in time recorded when the probe pulse crosses the plasma. As viewed from above, the probe pulse is effectively a thin sheet while the drive pulse is focussed and much smaller. Due to the 90° transverse probing geometry, the pulses cross for a short time at some location in the plasma. This crossing point can be translated inside the jet by adjusting the delay between the drive and probe.

Throughout this chapter the delay between the probe and the pump is a key experimental variable. We set zero delay to correspond to the two beams crossing when the drive pulse is at the leading edge of the gas jet. Increasing the delay between the drive and probe moves the crossing point towards the rear of the gas jet.

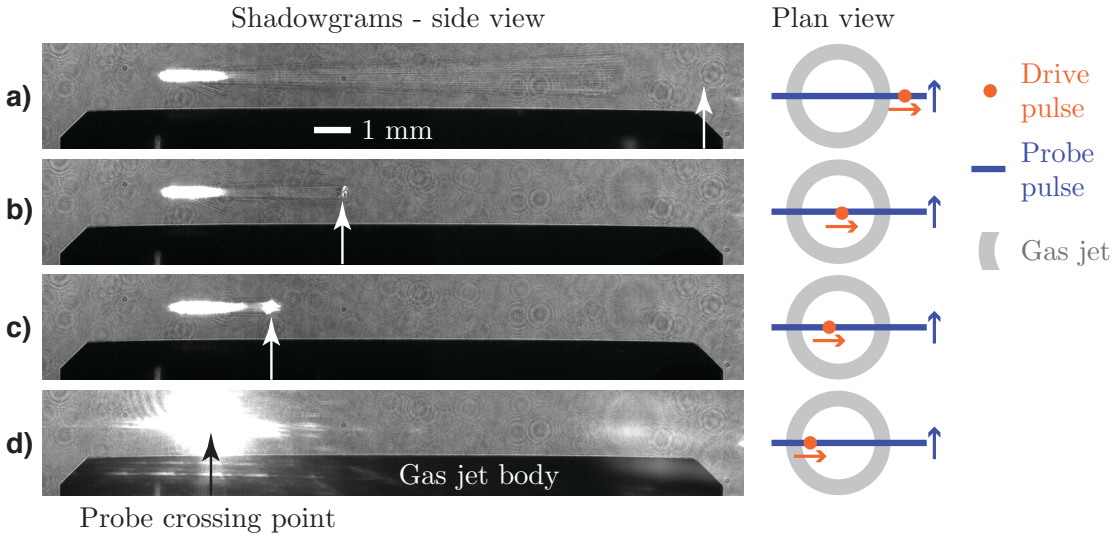


Figure 5.1: a) - d) Images recorded on the shadowgraphy camera over four separate shots as the delay between the drive and probe is progressively reduced. This moves the crossing point earlier in the plasma channel, as illustrated in the cartoons on the right.

Typically during an experiment the probe pulse is delayed with respect to the drive pulse such that the two pulses cross soon after the drive pulse exits the plasma. This is late enough that the entire plasma channel is visible, but early enough that hydrodynamic expansion has not yet occurred and therefore the density measurement of the plasma channel is an accurate one. During the initial stages of experiment I the probe pulse delay was too low, causing the probe to arrive early and the pulses to cross earlier in the channel. At this point it was noticed that the drive pulse was visible as a bright region on the shadowgram, often associated with detailed structure. As the probe delay was reduced this region was observed to increase in brightness, and when the two pulses crossed near the start of the plasma channel as in Fig. 5.1d) the light level increased hugely, saturating the camera. The same procedure was attempted during experiment II where the laser polarisations were approximately crossed. Some increase in side scatter was observed, but to a far lesser extent than that visible in Fig. 5.1d).

During experiment III the probe beam was much more intense, and its brightness on the diagnostic cameras more comparable to that of the side-scatter. It was therefore possible to observe more detail in the side scatter region within the dynamic range of the cameras. A plasma density scan was performed while the two pulses crossed in a region known to produce a large increase in side-scatter. Sections of the shadowgrams from this scan are plotted in Fig. 5.2. The dashed line marks the spatial location at which the two pulses cross for each image, and the solid line marks the position at which side-scatter is observed to be generated. Like the previous shadowgrams, the drive pulse is travelling across the

image from left to right.

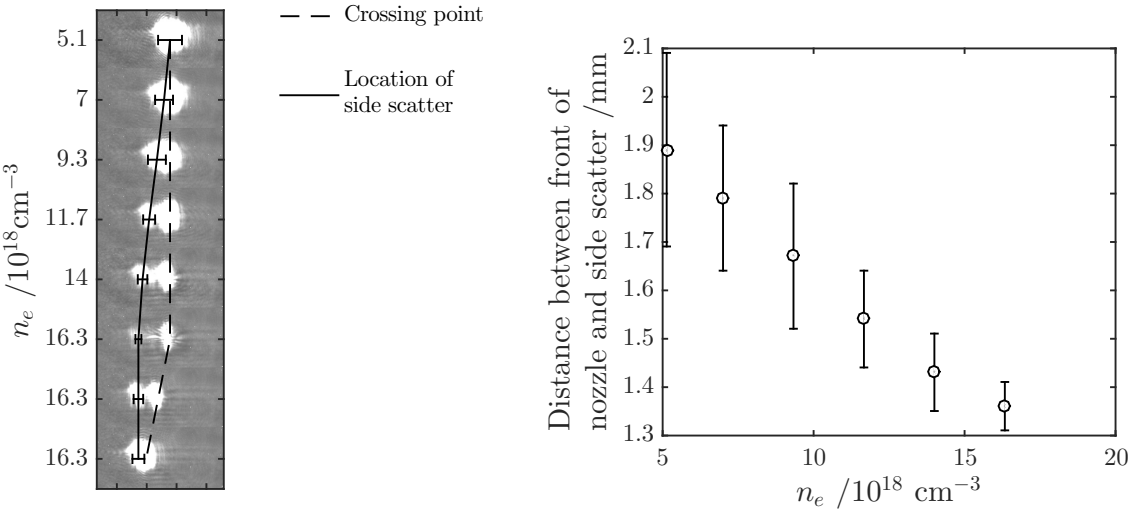


Figure 5.2: Sections of shadowgrams recorded during experiment III. As the plasma density was increased the location of the side scatter was observed to move earlier in the gas jet.

The probe pulse was initially delayed such that it maximised the scatter for $n_e = 5.1 \times 10^{18}\,\mathrm{cm}^{-3}$. As the plasma density was increased the side scatter region was observed to move to the left, closer to the leading edge of the plasma. At a density of $n_e = 16.3 \times 10^{18}\,\mathrm{cm}^{-3}$ the location of the side scatter was shifted far enough that two bright regions are observed, one for the side-scatter and one for the pulse crossing point. Reducing the probe delay again such that the pulses crossed in the side-scatter region increased the total amount of scattered light. The distance of the side-scatter region from the edge of the gas nozzle is plotted alongside as a function of density. If, following the previous chapter, the side scatter is assumed to arise from the unguided portion of the laser pulse, these data support the hypothesis that the pulse evolution is more rapid at higher densities.

5.1.2 Imaging spectrometry

During experiment III an imaging spectrometer was added to the probe diagnostic suite. An additional periscope rotated the image by 90° , and a demagnifying lens used to image the plasma channel onto the vertical entrance slit of the spectrometer. A spectrally-resolved image of the side-scatter is plotted in Fig. 5.3, where in this case there was no probe beam present and the drive pulse is propagating upwards. For the same gas jet backing pressure, on a second shot the probe beam was enabled and the plasma density profile measured interferometrically. The expected location of the Raman-shifted light as a function of position is plotted as a solid line. The dashed line marks the expected frequency shift assuming an electron Lorentz factor corresponding to $a_0 = 3.9$. The spatially-integrated spectrum is plotted beneath.

This is then the normal Raman side-scattered light from the beginning of the plasma channel as presented in the previous chapter. Unlike the results of the low-power Astra experiment in section 4.1, here the spectra are very broad and extend below the central wavelength of the drive pulse. As discussed in chapter 4 this is because the bandwidth

of the scattered light increases at high drive intensity. The peak of the spectrum occurs at shorter wavelengths than would be expected given the ambient plasma density, which suggests that a significant relativistic shift is occurring. This is different to the previous Gemini results of section 4.2 where it was found that minimal relativistic shifts were required in order to fit the data. The majority of the scattered light is observed to be emitted in the region of the gas jet where the density is approximately constant, rather than in the density ramp.

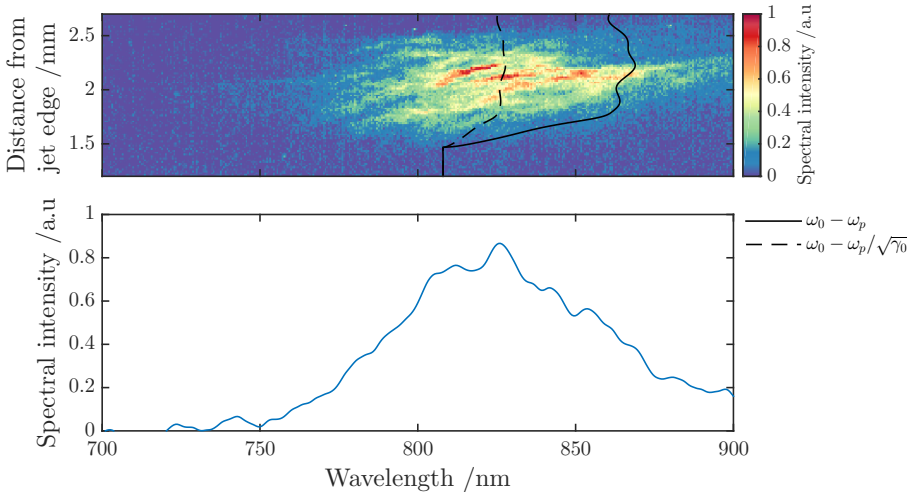


Figure 5.3: An example spectrum of side-scattered light as recorded on the imaging spectrometer during experiment III for $n_e = 6 \times 10^{18}$ cm^{-3}. The solid and dashed lines mark the expected non-relativistic and relativistic frequency shifts respectively.

The spectrum of this 'natural' Raman scattered light should be compared to that produced when the probe pulse is timed to cross the the drive pulse near the start of the gas jet. With the probe beam enabled now, in Fig. 5.4 spectrometer images are displayed for shots with and without the presence of a plasma.

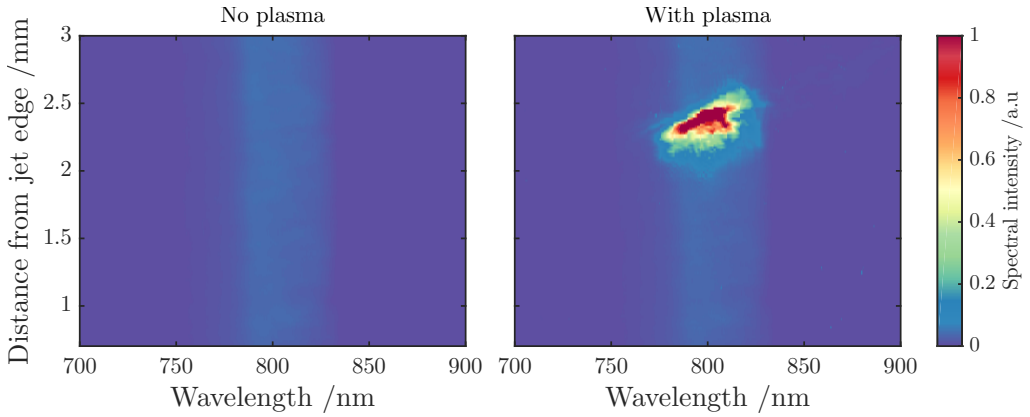


Figure 5.4: Comparison of spectrometer images recorded with and without the presence of plasma. The vertical streak is due to the probe beam, and the plasma density was $n_e = 1.6 \times 10^{19}$ cm^{-3}.

The additional light is observed to have similar spectral properties to the probe beam,

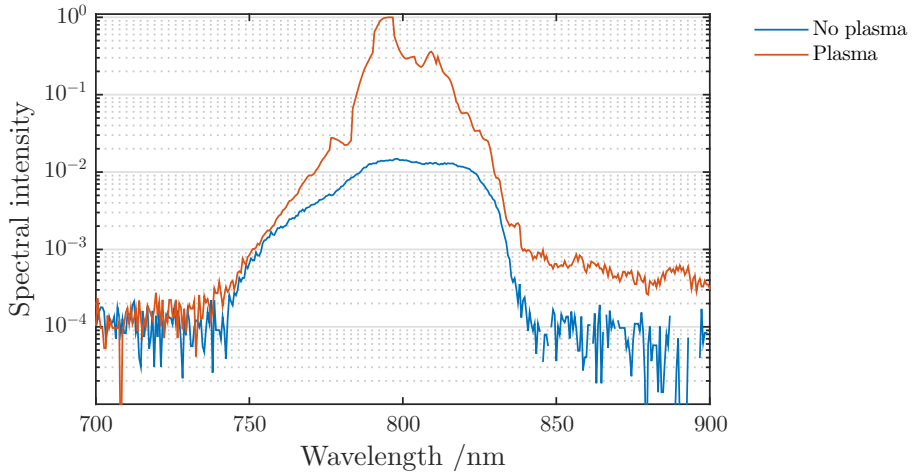


Figure 5.5: Comparison of spectral lineouts with and without the plasma present. The presence of plasma causes an increase in light across the probe bandwidth, and the addition of low levels of Raman side scatter.

in that the central wavelength and bandwidth are similar. Lineouts taken from the spectrometer images at 2.3 mm are plotted in Fig. 5.5. The scattered light is again observed to be much brighter than the probe beam, by a factor of a hundred or more on the spectrometer CCD. Red-shifted Raman scatter is observed to be present, albeit at a much lower intensity than the probe beam or the bright scattered light. The average central and FWHM bandwidths of the scattered light are plotted in Fig. 5.6 as a function of plasma density. The expected nonrelativistic and relativistic wavelengths of Raman-scattered light are again plotted as solid and dashed lines. The constant FWHM bandwidth of the probe beam is indicated by the shaded region.

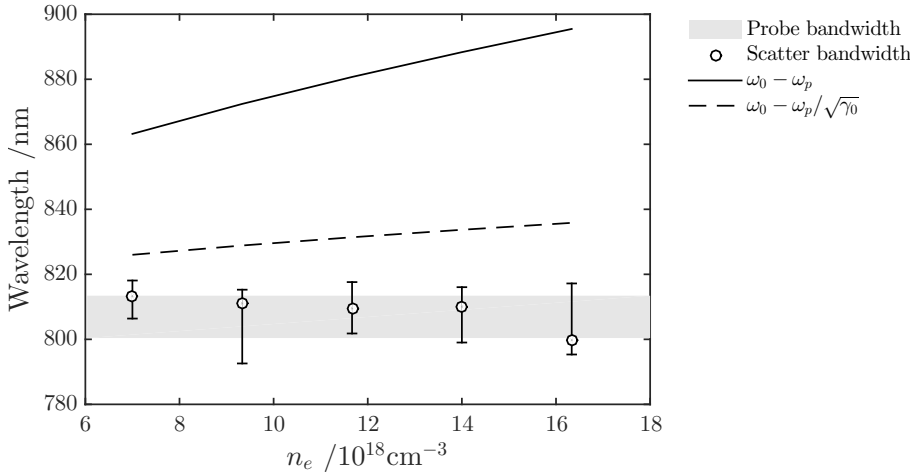


Figure 5.6: The mean wavelength of the amplified light as a function of plasma density. The light is observed to remain near the probe wavelength and does not increase with plasma density as would be expected for Raman scatter, even in the strongly relativistic case.

If the additional scattered light were due to Raman scattering, perhaps stimulated in some way by the probe pulse, it would be expected that its frequency would be shifted. This is not observed to be the case, even at high enough densities that a shift should be clear for

strong relativistic modifications. The scattered light is then in the same spatial location as the probe pulse and appears to be composed of the same frequencies, so may reasonably be considered to be a component of the probe pulse. It is therefore appropriate to term this an amplification process rather than a scattering process.

5.1.3 Amplification as a function of position

During all three experiments the amplification of the probe beam was examined as a function of pulse crossing position. In each case the shadowgraphy camera was used to measure the total amount of energy emitted from the plasma channel. This is because the CCD chip was large and there were no apertures in the imaging line, minimising the possibility that light was obstructed before it hit the CCD. When the image became very bright, secondary images appeared on the CCD due to reflections from the glass cover protecting the chip. If the CCD began to saturate at any point, these secondary images were used instead to measure the total amount of energy, boosting the effective dynamic range of the camera. It was also necessary to subtract the probe light from each image, which was accomplished on a case-by-case basis by fitting a smooth beam profile behind the bright region(s) before subtracting the entire beam from the image.

The results of experiments I and II using the 15-mm nozzle are plotted in the top panel of Fig. 5.7, and the results of experiment III in the bottom panel. In each case the energy measurements were calibrated using the known near-field fluence of the probe beam, and the error bars composed of shot-to-shot fluctuations and the combined uncertainties of the energy calibration and background subtraction processes. At the plasma densities used for each scan, the integrated energy contained in the Raman side-scatter was measured and plotted here as a grey shaded region representing the RMS shot-to-shot variation. A schematic of the longitudinal plasma density profile is plotted as a solid red line, and all distances are measured with respect to the outer solid edge of the gas jet nozzle.

For experiment I, as was already clear from the shadowgraphs of Fig. 5.1, when the two pulses cross late in the gas jet there is no clear amplification of the probe. As the pulses cross progressively earlier the amplification increases rapidly, and the total scattered energy may increase by a factor of 25 or more. If it is assumed that the region of the probe pulse experiencing amplification has a Gaussian profile of 30 µm FWHM and that the probe pulse duration is unchanged, the peak intensity of the probe beam is increased by a factor of up to 10^5. This level of peak amplification is witnessed over a ≈ 1 mm-wide region near the start of the plasma. Where the pulses start crossing in the density ramp the amplification is diminished. An identical procedure was carried out in experiment II, which occurred shortly after experiment I and made use of a very similar diagnostic setup. With the beam polarisations approximately orthogonal the amplification process was almost absent. There was some slight enhancement in the light level visible between 3 and 4 mm, but the effect was much less obvious. This is an indication that the amplification relies on the interference between the two pulses.

During experiment III the plasma density was higher and the gas jet shorter. This also has the effect of shortening the length of the density ramp at the leading edge of the plasma, and as is clear from the lower panel of Fig. 5.7, shortening the region over which

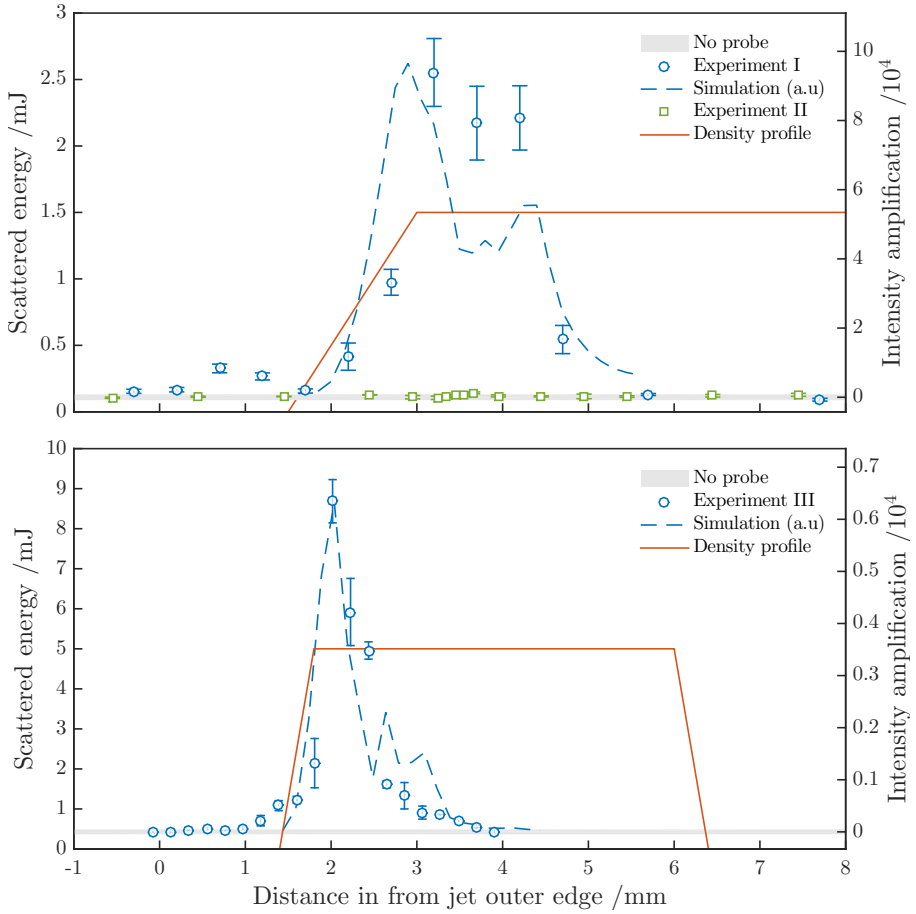


Figure 5.7: The integrated level of scattered light is plotted as a function of the pulse crossing position for the three experiments. When the polarisation of the beams are crossed the scattering is hugely diminished. When the plasma density is higher and the density ramp length shorter amplification occurs over a smaller interval. Assuming the amplified light originates from a Gaussian profile of 30 µm FWHM, the inferred intensity amplification is indicated on the right axes.

amplification is observed. The total energy of the scattered light was larger in this case, but the probe beam was also 50 times more intense (though still vastly less intense than the drive pulse). Though the fractional amplification was lower, after amplification the peak probe intensity was higher than that of experiment I.

The dashed lines are the results of 2D PIC simulations, discussed further in the following section. Though the simulation results have been rescaled vertically here, the simulations are in good agreement when it comes to the change in the level of scatter as a function of pulse crossing position, suggesting that the relevant physics is captured.

5.2 Simulation results

In an effort to qualitatively understand the response of a plasma to crossing laser pulses a series of 2D simulations were performed. The first set presented here explore the physical processes occurring in the plasma, and as such are heavily idealised. The second set are closer approximations to the experiment and examine how the amplification process changes as a function of plasma density, drive pulse intensity and probe pulse wavelength.

All were performed with the EPOCH code at a resolution of 40 cells per wavelength in both directions. The results were observed to be insensitive to the number of particles per cell between 2 and 16. The simulation boundaries were open to both particles and fields, the particles were represented as a first-order (triangular) shape function with variable weighting, and other physical effects such as ionisation and Coulomb collisions were ignored.

5.2.1 CW simulations

As a first approximation, two constant-wave (CW) lasers were injected into a square simulation box from adjacent edges, incident on a square $30\times 30\,\mu$m slab of plasma. The plasma density was $2\times10^{18}\,\mathrm{cm}^{-3}$ and the laser intensities were $I_0 = 10^{16}\,\mathrm{Wcm}^{-2}$ and $I_1 = 10^{13}\,\mathrm{Wcm}^{-2}$, polarised out of the plane of the simulation. The simulation box was larger than the plasma at $100\times100\,\mu$m, and the Poynting flux of the low-intensity beam was measured at the exit of the simulation box after it had traversed the plasma.

A map of the electron density is plotted in Fig. 5.8, taken shortly after the two lasers fully cross the plasma. Here the high-intensity beam is propagating to the right, and the low-intensity beam to the top. Immediately apparent is the presence of short-wavelength electron waves, which are spatially static and oscillate in amplitude at the plasma frequency. The initial phase of the oscillation is set by the time at which the two laser beams first cross, and so the amplitude of the waves appears to propagate across the simulation box to the right and top at the group velocity of the beams.

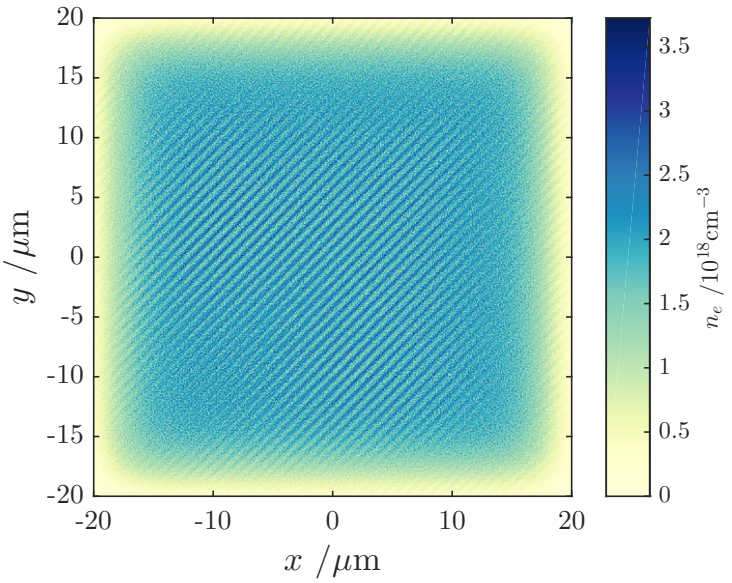


Figure 5.8: Electron density plot taken 100 fs after the beams first fully propagate through the plasma. Beam crossing is marked by the immediate appearance of short wavelength plasma density modulations.

To see how these waves are generated by the ponderomotive force of the crossing laser beams, let the normalised vector potentials of the two beams be represented as $\mathbf{a}_0 = \hat{\mathbf{z}}a_0\sin(\mathbf{k}_0\cdot\mathbf{r}-\omega_0 t)$ and $\mathbf{a}_1 = \hat{\mathbf{z}}a_1\sin(\mathbf{k}_1\cdot\mathbf{r}-\omega_0 t)$, where the beams have the same frequency and propagate in the $\mathbf{k}_{0/1}$ directions. The ponderomotive force $\mathbf{F}_p$ is proportional to the

gradient of the time average of the square of the total vector potential,

$$\begin{aligned} \mathbf{F}_p &\propto \nabla\langle(\mathbf{a}_0+\mathbf{a}_1)\cdot(\mathbf{a}_0+\mathbf{a}_1)\rangle = 2\nabla\langle\mathbf{a}_0\cdot\mathbf{a}_1\rangle \\ &= a_0a_1\nabla\langle\cos\left((\mathbf{k}_0-\mathbf{k}_1)\cdot\mathbf{r}\right)-\cos\left((\mathbf{k}_0+\mathbf{k}_1)\cdot\mathbf{r}-2\omega_0t\right)\rangle \\ &= -a_0a_1(\mathbf{k}_0-\mathbf{k}_1)\sin\left((\mathbf{k}_0-\mathbf{k}_1)\cdot\mathbf{r}\right) \end{aligned} \tag{5.1}$$

where the temporal averaging removes the oscillating terms. Referring to Fig. 5.8 the ponderomotive force is directed along $y=-x$ line, and has spatial wavelength $\lambda_0/\sqrt{2}$. The effect is to periodically bunch and un-bunch electrons, causing the observed density structure. As the frequencies of the two beams are equal, the ponderomotive force field is static in space and a single electron can feel its effects for a long time. Even when one beam is very much weaker than the other, the fact that the spatial scale length of the static fields is only a laser wavelength means the gradient of these fields, and therefore the ponderomotive force, can be large. If the laser polarisations are set perpendicular, as in experiment II, the time-averaged ponderomotive force is zero and a grating is not formed.

Electron motion

Given the existence of the ponderomotive grating, it is possible to calculate the non-relativistic motion of a plasma electron. Consider the rotated co-ordinate system sketched in Fig. 5.9.

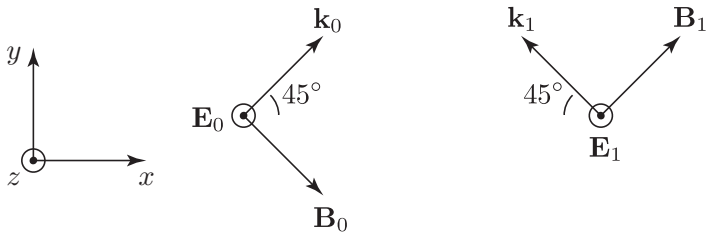


Figure 5.9: Geometry for intersection of two beams forming a grating.

The explicit form of the laser fields is then

$$\begin{aligned} \mathbf{k}_0 &= \tfrac{k_0}{\sqrt{2}}(\hat{\mathbf{x}}+\hat{\mathbf{y}}) & \mathbf{k}_1 &= \tfrac{k_0}{\sqrt{2}}(-\hat{\mathbf{x}}+\hat{\mathbf{y}}) \\ \mathbf{B}_0 &= \tfrac{B_0}{\sqrt{2}}(\hat{\mathbf{x}}-\hat{\mathbf{y}})\sin\psi_0 & \mathbf{B}_1 &= \tfrac{B_1}{\sqrt{2}}(\hat{\mathbf{x}}+\hat{\mathbf{y}})\sin\psi_1 \\ \mathbf{E}_0 &= \hat{\mathbf{z}}E_0\sin\psi_0 & \mathbf{E}_1 &= \hat{\mathbf{z}}E_1\sin\psi_1 \\ \psi_0 &= \mathbf{k}_0\cdot\mathbf{r}-\omega_0t & \psi_1 &= \mathbf{k}_1\cdot\mathbf{r}-\omega_0t \end{aligned} \tag{5.2}$$

As the velocity of the electron is nonrelativistic, $|\mathbf{v}|\ll c$ and in the z direction the electric fields dominate the Lorentz force so $v_z = c(a_0\cos\psi_0 + a_1\cos\psi_1)$.
Using this expression in the x-component of the $\mathbf{v}\times\mathbf{B}$ force in the x direction, the equation of motion in this direction is

$$\frac{\mathrm{d}v_x}{\mathrm{d}t} = \frac{\omega_0 c}{2\sqrt{2}}\left(-a_0^2\sin 2\psi_0 + 2a_0a_1(\sin(\psi_0+\psi_1)-\sin(\psi_0-\psi_1)) + a_1^2\sin 2\psi_1\right). \tag{5.3}$$

The terms oscillating at high frequency can be neglected over the electron motion timescale, so the slow electron motion is governed by the $\psi_0 - \psi_1$ term which yields

$$\frac{\mathrm{d}^2 x}{\mathrm{d}t^2} + \frac{\omega_0 c}{\sqrt{2}} a_0 a_1 \sin(\sqrt{2} k_0 x) = 0. \tag{5.4}$$

The electron oscillates about its starting position analogously to a large-amplitude pendulum. If the electron is close to a node in the standing wave, then $\sqrt{2}k_0 x \ll 1$ and the approximate low-amplitude equation of motion is

$$\frac{\mathrm{d}^2 x}{\mathrm{d}t^2} + \omega_B^2 x = 0 \qquad \omega_B \equiv \omega_0 \sqrt{a_0 a_1}\,. \tag{5.5}$$

This is simple harmonic motion at the 'bounce' frequency ω_B. In the case that $a_0 \gg a_1$ and is relativistically intense, it is possible to derive the relativistic bounce frequency[1] $\omega_B = \omega_0 \sqrt{a_0 a_1}/\gamma_0$, where γ_0 is the Lorentz factor of the electron motion in the fields of a_0. If the two beams have different frequencies ω_0 and ω_1, the bounce frequency is modified to $\omega_B = \sqrt{\omega_0 \omega_1}\sqrt{a_0 a_1}/\gamma_0$. The condition that the laser fields dominate the motion of the plasma electrons is $\omega_B > \omega_p$, and the grating is established over a timescale of ω_B^{-1} which can be short relative to a plasma period if the lasers are sufficiently intense. Plasma gratings are also formed at the surfaces of solid-density targets, and their growth can be calculated in a similar fashion[2].

While Eq. 5.4 does not have solutions expressible in simple closed form, it is possible to calculate the electron momentum as a first integral. Here the components along the grating wavefronts ($p_\parallel$) or across them ($p_\perp$) are explicitly separated. Fig. 5.10 illustrates the relevant geometry.

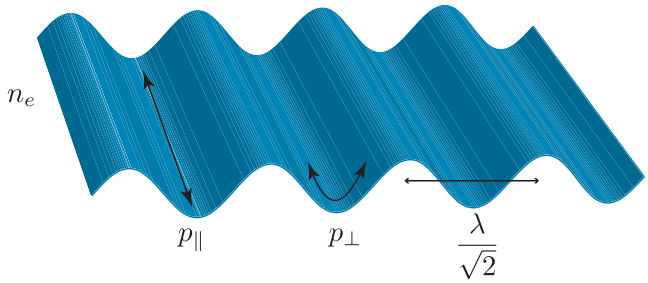


Figure 5.10: Parallel and perpendicular momenta directions relative to the plasma grating.

Using the relation

$$\frac{\mathrm{d}v_\perp}{\mathrm{d}x_\perp} = \frac{\dot{v}_\perp}{\dot{x}_\perp} = -\frac{\omega_0 c a_0 a_1}{\sqrt{2} v_\perp} \sin(\sqrt{2} k_0 x_\perp), \tag{5.6}$$

the perpendicular velocity component can be expressed as

$$v_\perp = \sqrt{c^2 a_0 a_1 \cos(\sqrt{2} k_0 x_\perp) + \text{const.}} \tag{5.7}$$

[1] [211] G. Shvets *et al. Physics of Plasmas.* (1997).
[2] [287] L. Plaja *et al. Physical Review E.* (1997).

The value of the constant depends on the initial position of the electron, assuming that the grating is formed in a time which is small compared to the plasma frequency such that the electron does not move significantly while the grating is established. If the electron performs closed oscillations, and the oscillation amplitude in the grating is x_0, then

$$v_\perp = c\sqrt{a_0 a_1}\sqrt{\cos(\sqrt{2}k_0 x_\perp) - \cos(\sqrt{2}k_0 x_0)}. \tag{5.8}$$

As Eq. 5.8 describes the large-amplitude motion of a pendulum, the effective frequency of large amplitude oscillations is well known as

$$\omega_{eff} = \omega_B \frac{\pi}{2} \frac{1}{\mathcal{K}(\sin(\sqrt{2}\pi x_0/\lambda))} \tag{5.9}$$

where $\mathcal{K}$ is the elliptic integral of the first kind and $\lambda = 2\pi/k_0$ is the laser wavelength. For $x_0 \ll \lambda$, i.e. oscillations which sit near the bottom of the grating, $\omega_{eff} \to \omega_B$ as derived above. As the oscillation amplitude increases towards $\lambda/(2\sqrt{2})$, the motion becomes increasingly anharmonic and the oscillation frequency decreases.

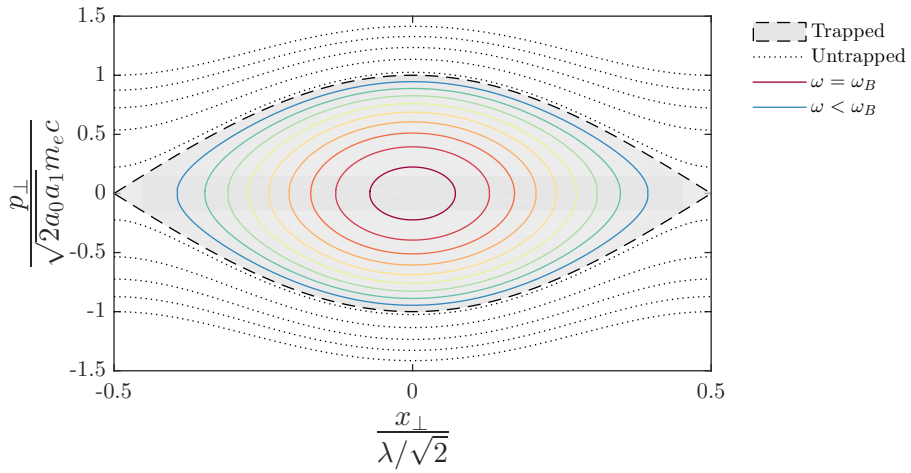


Figure 5.11: Perpendicular phase space for electrons oscillating in a ponderomotive grating.

Depending on the initial momenta of electrons encountering the ponderomotive grating, they may or may not become trapped in the potential. To find the trapping threshold, note that for $v_\perp$ to have real solutions it must be that $\cos(\sqrt{2}k_0 x_\perp) > \cos(\sqrt{2}k_0 x_0)$ for all $x_\perp < x_0$. The last point at which this remains possible is when $\cos(\sqrt{2}k_0 x_0)$ reaches a minimum, i.e. when $\sqrt{2}k_0 x_0 = \pi$ and $x_0 = \lambda/(2\sqrt{2})$. Specifying this value of the oscillation amplitude therefore determines the orbit of the separatrix, from which the maximum (nonrelativistic) velocity of a trapped electron is $v_{\max} = c\sqrt{2a_0 a_1}$, equivalent to a maximum trapped momentum of $p_\perp = \sqrt{2a_0 a_1} m_e c$. In Fig. 5.11 the phase space orbits of several electrons are plotted inside one spatial period of the grating. Low-energy electrons occupy the closed trapped orbits, and oscillate near ω_B. Higher energy electrons orbit nearer the separatrix, performing oscillations at frequencies lower than ω_B. Those with $p_\perp > \sqrt{2a_0 a_1} m_e c$ are not trapped in the grating and freely cross it in the perpendicular direction. The electron motion in the orthogonal direction is unaffected by the laser fields on a slow timescale, and therefore any oscillations occur at the plasma frequency.

A pair of 2D simulations were performed where again a slab of uniform plasma is irradiated by two CW laser beams of equal intensities. One simulation was low intensity where $\omega_B = 0.4\omega_p$, the second was high intensity where $\omega_B = 4\omega_p$. A plasma grating is rapidly formed in both cases, and test electrons are injected into this grating and their motion tracked. The test electrons respond to electromagnetic fields, but are ignored during the charge and current deposition processes of the PIC algorithm and have no effect on the rest of the simulation. The paths of several electrons from the high-amplitude simulation are plotted in Fig. 5.12

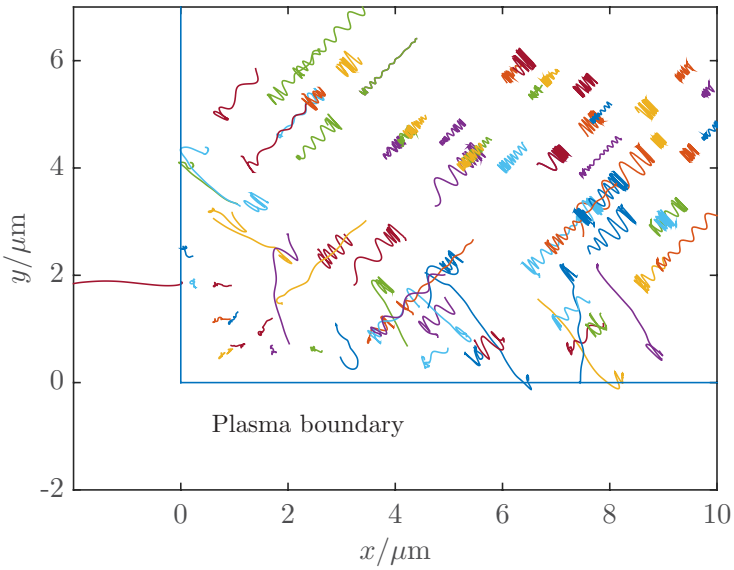
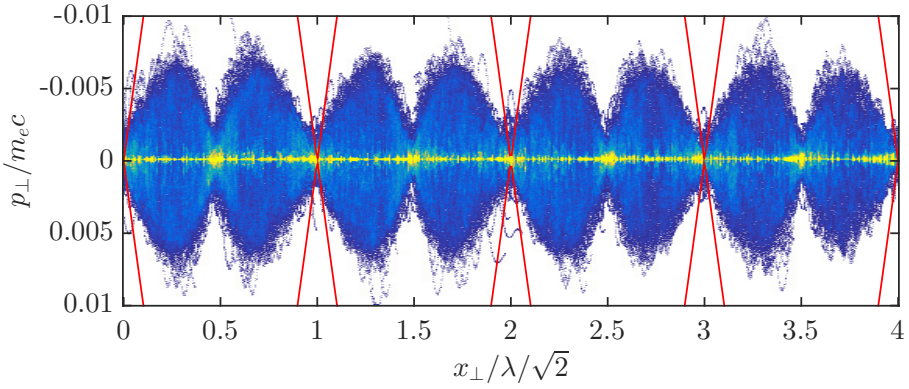
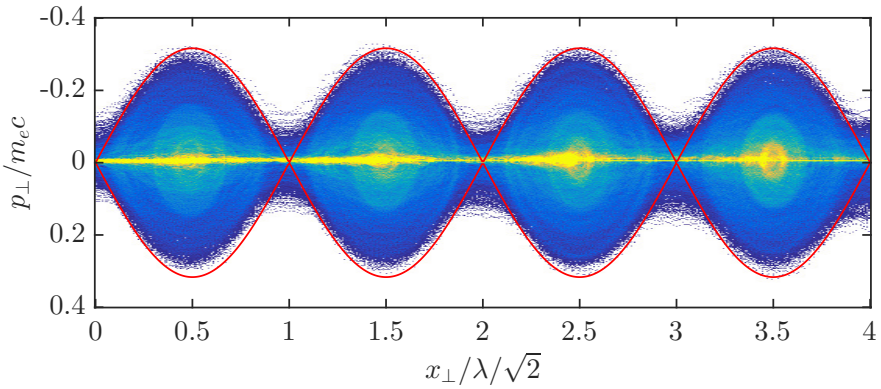


Figure 5.12: Spatial motion of test electrons over 200 fs during the high-amplitude simulation.

It is clear that the electrons are largely confined to the wells in the grating, even when the laser amplitudes are large. Some electrons gain enough energy to become untrapped and cross the well boundaries, but these are the exception in general. Most electrons perform oscillations perpendicular to the grating direction, while drifting parallel to the grating.
The $(x_\perp, p_\perp)$ phase spaces from the two simulations are plotted in Figs. 5.13a) and 5.13b) on a logarithmic colour scale, averaged along the $x_\parallel$ direction. The red line in each case marks the separatrix predicted above for $a_0 = a_1$, $p_\perp = m_e c a_0(\cos(\sqrt{2}k_0 x_\perp))^{1/2}$.
In the low amplitude case the electron momenta are much lower than the maximum allowable trapped momenta, the separatrix extending far outside the plotted axes. There is some reduction of the electron momenta near the edges of the grating wells, but also near the centres. This second effect is unexpected, and may be due to the fact that the plasma frequency is near to an integer multiple of the bounce frequency, leading to a resonance in the electron motion. In the high amplitude case the perpendicular electron phase space is in good agreement with predictions. Electrons are accelerated by the ponderomotive potential to the maximum allowed momentum, and then largely trapped in the grating. There is evidence of some motion between wells, also visible in Fig. 5.12, but the majority of electrons execute bounded oscillations.
By Fourier-transforming the perpendicular electron motion in time, it is possible to



(a) Low-amplitude simulation



(b) High-amplitude simulation

Figure 5.13: $(x_\perp, p_\perp)$ phase space on a logarithimic colour scale. The red lines mark the separatricies as predicted by Eq. 5.8.

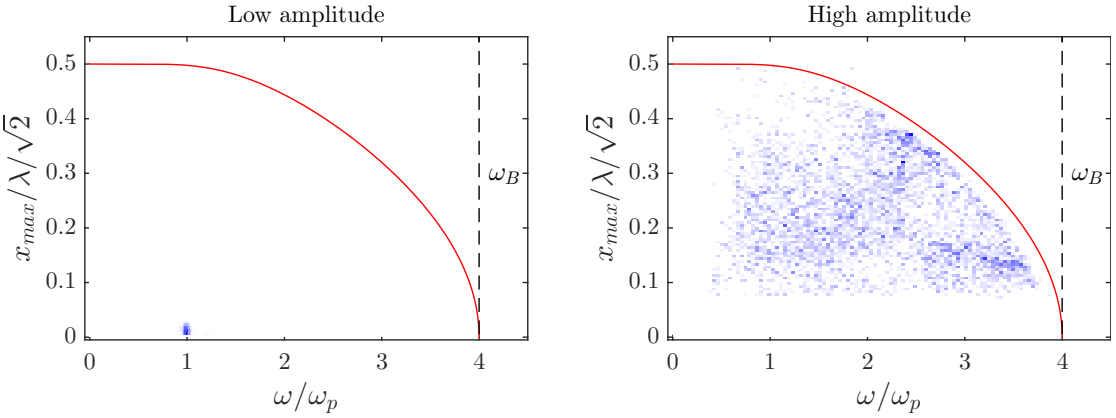


Figure 5.14: 2D histogram of electron oscillation amplitude and oscillation frequency from the low-amplitude simulation. The red line is the prediction of Eq. 5.9.

measure the dominant frequency ω of the oscillatory motion. Significant oscillatory motion is signalled by the presence of a peak in the Fourier spectrum of the motion of a single electron. For a subset of test electrons, the frequency and amplitude of the main oscillatory component of the motion is calculated and binned into a 2D histogram. The duration of the simulation was too short to reliably resolve frequencies below $\approx 0.5\omega_p$. In Fig. 5.14 the test electrons from both simulations are plotted. The red line is the prediction of Eq. 5.9, and as expected it bears little relation to the electron motion in the low-amplitude

case, which is instead dominated by small amplitude oscillations at the plasma frequency. The high amplitude case is markedly different. Electrons oscillate at many frequencies, and lower frequencies correspond to larger oscillation amplitudes in general. The overall envelope matches the prediction of Eq. 5.9 fairly well, suggesting that the electron motion can be approximated by a pendulum model even when mildly relativistic.

Laser modulation

Return now to the simulation of Fig. 5.8 where the two beams have widely differing intensities. The change in Poynting flux of the low-intensity beam as a function of time normalised to $T_p = 2\pi/\omega_p$ is plotted in Fig. 5.15. It can be shown (see appendix A.7) that the vector potential of the low-intensity beam is modified as it propagates in y as

$$\frac{\partial a_1}{\partial y} = -i\frac{\omega_p^2 a_0}{2c\omega_0}\frac{\delta n}{n}e^{i\omega_p t} = -i\frac{1}{c}\left(\frac{\omega_p}{\omega_0}\right)^{3/2}\frac{\delta n}{n}\gamma_R e^{i\omega_p t} \tag{5.10}$$

where $\delta n/n$ is the amplitude of the plasma grating, γ_R is the linear growth rate for Raman backscatter and t marks the time since the head of the beam entered the plasma. The oscillating exponential term represents the phase of the plasma grating – as an oscillator driven by a constant force, the grating oscillates at its natural frequency at a fixed amplitude. As the simulation progresses, new portions of the beam continually encounter the grating. A given (temporal) portion of the beam witnesses the grating in the same phase as it progresses across the plasma, as the grating was initiated at different times in y during the initial propagation of the beam. Some portions of the beam witness the grating in an amplifying phase, where the higher-intensity beam is scattered into the direction of the low-intensity beam. Other portions witness the grating in a depleting phase, where the scattering direction is reversed. Upon exiting the plasma the intensity remains constant until the exit of the simulation box, and this is why in Fig. 5.15 the amplitude of the beam oscillates as a function of time at the plasma frequency. In steady-state no amplification would be expected for beams of equal frequency, and averaging the Poynting flux over many plasma periods confirms this.

If instead the low-intensity beam is down-shifted in frequency by ω_p, the phase term in Eq. 5.10 vanishes and the resulting growth rate is exactly that of the Raman amplifier (see Eq. 2.134 of section 2.9.3). The low-intensity beam should be expected to grow continuously as the plasma grating does not oscillate but is instead driven to increasingly larger amplitudes. In the initial period of the simulation this is indeed the case, and the low-intensity beam increases in intensity by a factor of 4. After several plasma periods, the relatively high intensities of the two beams (in the context of Raman amplification) cause the electron waves to break and the amplification efficiency rapidly drops.

If the plasma electrons are given a finite temperature the above processes are inhibited. Examination of the electron density indicates that the grating is of lower amplitude, likely due to the homogenising effect of the additional thermal pressure. Those simulations were also observed to rapidly become numerically unstable.

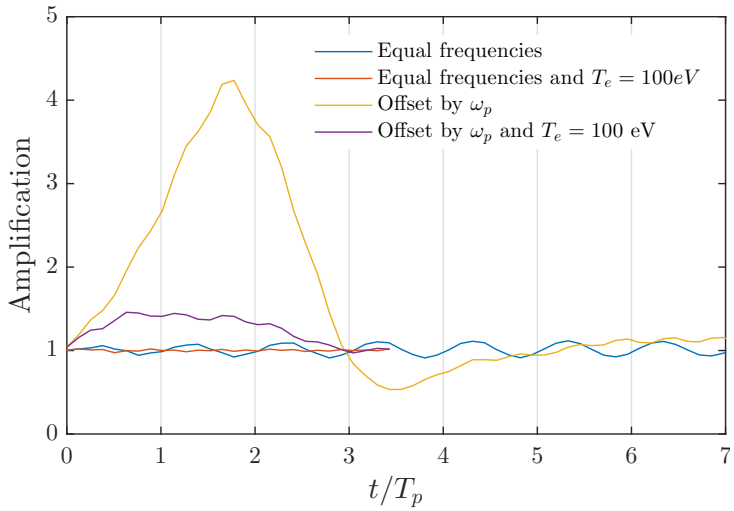


Figure 5.15: The relative increase in Poynting flux at the exit of the simulation box is plotted as a function of time normalised to the plasma period for different conditions. Lines are terminated early when the simulation became numerically unstable.

5.2.2 Parameter scan

The amplification of a lower-intensity beam was observed to be a transient process in the CW simulations, occurring over a timescale of ω_p^{-1}. This is similar to the pulse length of the lasers in a LWFA experiment, so it is important to include finite pulse length effects as here the transient behaviour will dominate the plasma response. In this series of simulations the drive beam is a pulse, of FWHM pulse length and width 40 fs and 20 µm respectively. The probe beam continues to be represented as a continuous plane wave, and is allowed to propagate fully across the simulation box before the drive beam is injected. This remains a good approximation because the drive pulse continually encounters fresh plasma and probe beam, and as both beams travel near the speed of light it is impossible for events early in the simulation to affect those occurring later. The additional benefit is that the amplification process may be examined as a function of drive pulse propagation in a single simulation. A series of simulations using pulsed probe beams was compared to those using continuous beams, and the measured amplifications agreed to within 5%.

An extensive parameter scan was performed, for 13 plasma densities between 1 and $4 \times 10^{18}\text{cm}^{-3}$ and 20 drive pulse intensities between $1 \times 10^{16}\,\text{Wcm}^{-2}$ and $3 \times 10^{19}\,\text{Wcm}^{-2}$. In all cases the probe beam intensity was $1 \times 10^{12}\,\text{Wcm}^{-2}$. The plasma density was assumed uniform, with a short 50 µm density ramp at the leading edge. Snapshots of each simulation were recorded at the same time, after the drive pulse had propagated ≈ 100 µm through the plasma. Three snapshots are plotted in Fig. 5.16, for drive pulse intensities of $2 \times 10^{17}\,\text{Wcm}^{-2}$, $2 \times 10^{18}\,\text{Wcm}^{-2}$ and $2 \times 10^{19}\,\text{Wcm}^{-2}$. In each case the plasma electron density is represented on a greyscale colour map, with contours of laser intensity overlaid. The red lines mark isocontours of the drive pulse intensity envelope, extracted from the out-of-plane E_z field map with a Hilbert transform filter. The probe fields are many orders of magnitude weaker, so are extracted from the B_x magnetic fields. Magnetic fields due to transverse propagation of drive pulse energy are filtered with a mask in the Fourier plane,

and the probe envelope again extracted with a Hilbert transform filter. The probe beam propagates upwards and the drive pulse propagates to the right. The 45° streak in the probe envelope is due to the 90° motion of the beams with respect to one another.

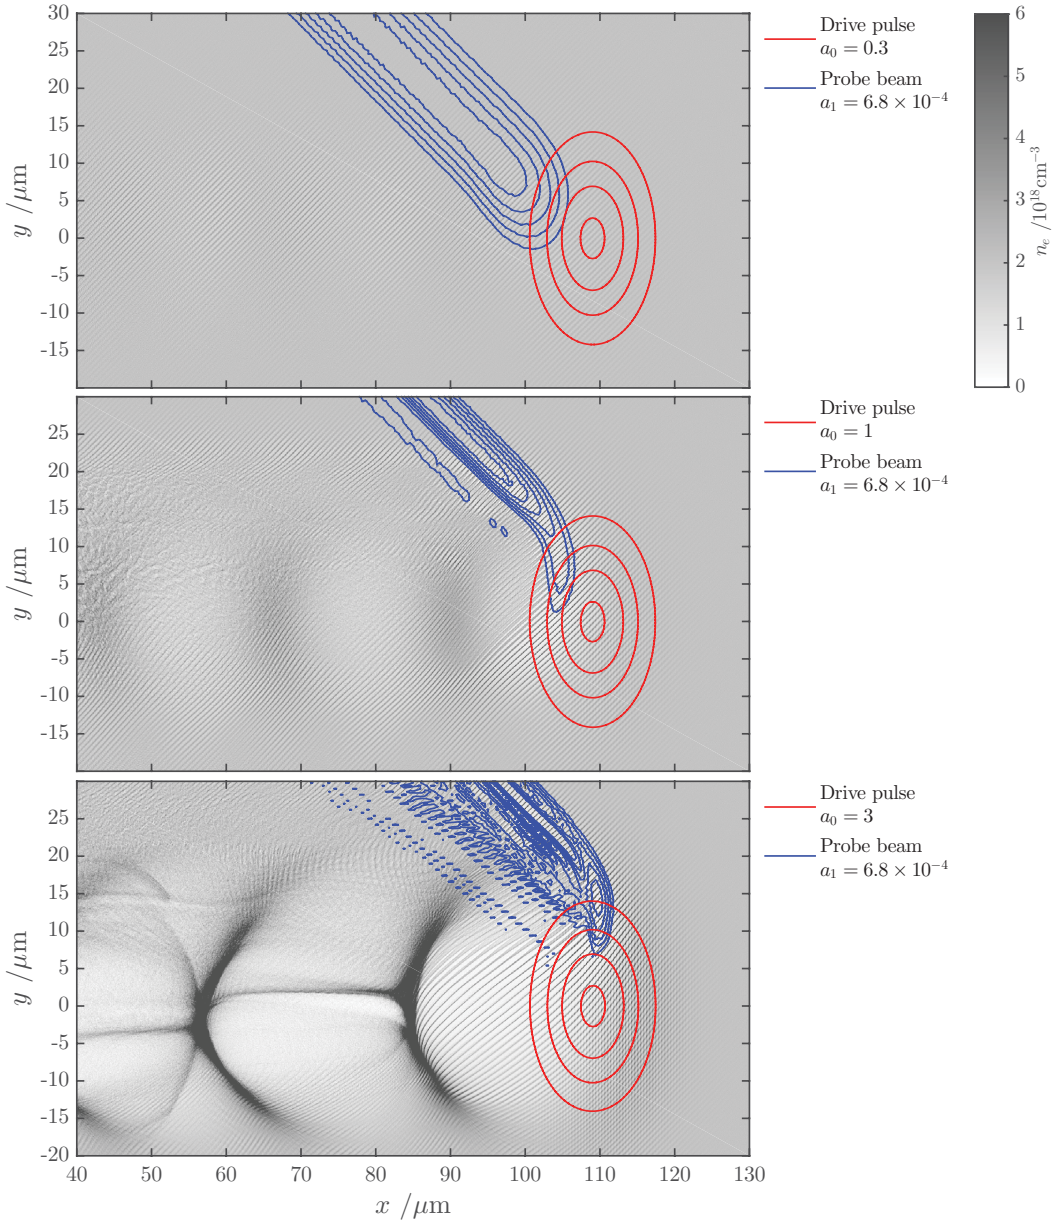


Figure 5.16: Electron density maps in greyscale from three simulations ordered by increasing drive pulse intensity. In each case the red contours mark equal spacings in drive pulse intensity, and the blue contours mark equal spacings in the range $[a_1^{\max}/2, a_1^{\max}]$ where $a_1^{\max}$ is the maximum value of the vector potential of the probe beam in the simulation box.

The probe amplification process is observed to change significantly as the drive pulse intensity is increased. For $a_0 = 0.3$ the amplification is broadly as expected. The plasma remains mostly homogeneous, with the exception of a low-amplitude density grating. The parts of the probe beam which pass through the peak of the drive pulse witness maximum amplification, and overall the amplification of the probe is simply a function of which

portion of the drive pulse it experiences. The probe reaches its maximum intensity near the outer edges of the drive beam, after which it propagates to the edge of the simulation box.

At a higher intensity corresponding to $a_0 = 1$ the amplification is more complex. Portions of the probe passing through the peak of the pulse are initially amplified and then subsequently depleted. The most amplified portion of the probe pulse transits the drive pulse some distance away from its peak, and continues to be amplified for some distance in the low intensity 'wings' surrounding the core of the pulse.

As a_0 continues to increase the region of peak amplification moves away from the centre of the drive pulse. The amplified pulse becomes more filamented, but also grows to significantly larger amplitudes. Unlike simulations at lower intensity, the grating structure is observed to be destroyed by the electron sheath around the bubble, likely due to electron sheet crossing and subsequent breakdown of fluid behaviour. The wake is also observed to become significantly asymmetric. This is surprising as the probe intensity in the third panel of Fig. 5.16 is 20 million times smaller than the peak drive pulse intensity, and yet clearly has a macroscopic effect on the wakefield structure.

The results of this set of simulations are summarised in Fig. 5.17, which shows the relative amplification of the probe energy as a function of plasma density and drive pulse intensity. The amplification is plotted on the z-axis and depicted in the colourmap with a logarithmic scale.

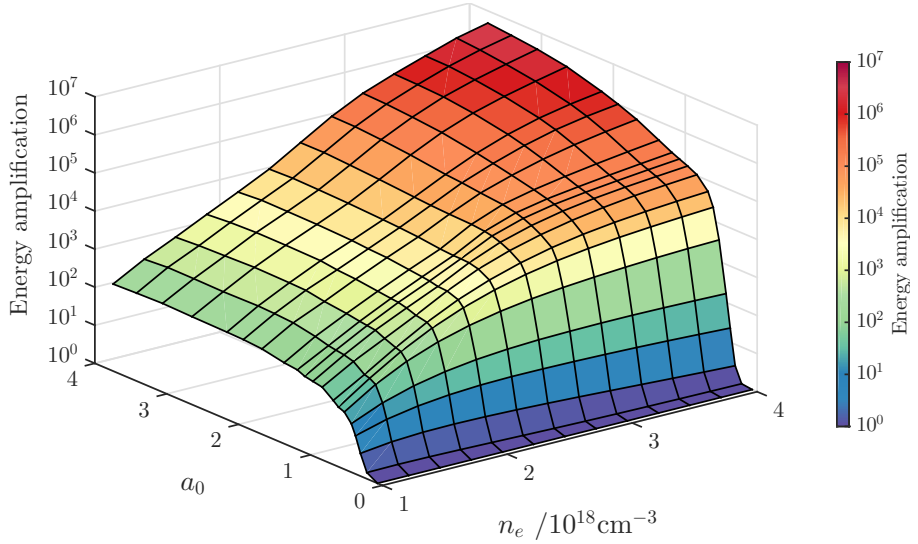


Figure 5.17: Energy amplification of probe as a function of n_e and drive pulse a_0.

In general the amplification is observed to increase with plasma density, with the exception of the very lowest intensities. To explain this effect, it is instructive to extract the temporal envelope of the probe beam as a spatial slice along the simulation y-coordinate. Several such envelopes are plotted for different plasma densities at two drive pulse intensities in Fig. 5.18. Larger times correspond to earlier interactions with the drive pulse.

In the low intensity case, for low plasma densities and long plasma periods the probe is amplified over the duration of its interaction with the drive pulse. As the plasma density

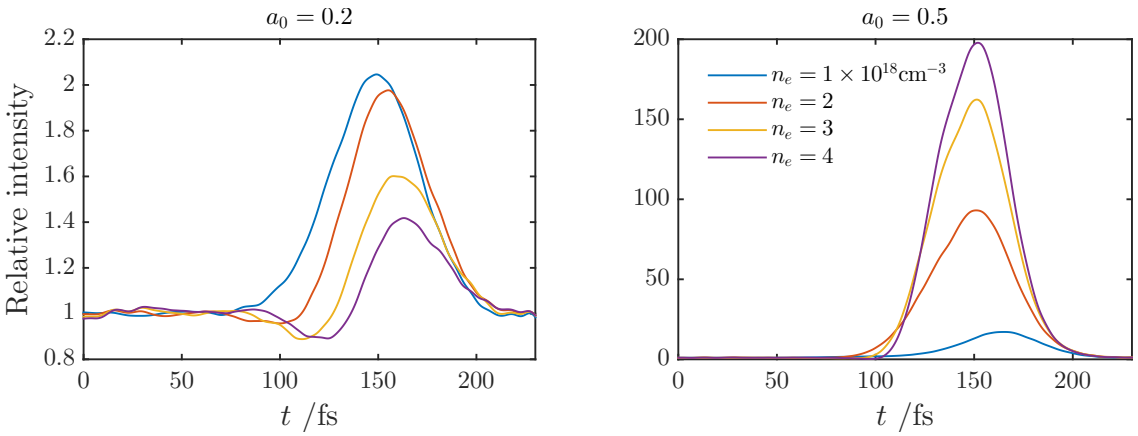


Figure 5.18: The temporal profile of the amplified probe beam as a function of density and drive pulse intensity. In the low intensity limit the plasma oscillations deplete the rear of the probe.

is increased, later in the interaction the rear of the probe is depleted. This is because at higher plasma densities, when the plasma oscillations occur more quickly, there is time over the transit of the probe pulse for the phase of the grating to advance such that probe energy is scattered into the pump. Energy is extracted from the rear of the probe, and the average intensity would tend to zero as the density is increased. For small amplification factors, integrating Eq. 5.10 over a Gaussian probe pulse envelope of duration τ shows that the probe energy amplification ratio scales as

$$\frac{E_{\mathrm{out}}}{E_{\mathrm{in}}} \approx a_0^2 \omega_0 \tau e^{-\tau^2 \omega_p^2/4} \tag{5.11}$$

which quickly tends towards 1 when the pulse length exceeds several plasma periods. For peak amplification one should choose $\tau \approx \omega_p^{-1}$, which is similar to the condition for driving a high-amplitude wake. In the high intensity case the probe grows so quickly over the first plasma period that any subsequent depletion by later oscillation periods is negligible. For all but the lowest intensities simulated here, this second effect is dominant and the probe amplitude always grows with plasma density.

Plotting the amplification results as separate line plots in Fig. 5.19 reveals a second curiosity – at low intensities the energy amplifications and scalings with a_0 are almost independent of the plasma density. To explain this, from appendix A.6 the expected amplitude of the grating in the low-intensity limit is

$$\frac{\delta n}{n} = \frac{\omega_0^2}{\omega_p^2} a_0 a_1 \tag{5.12}$$

and therefore when inserting into Eq. 5.10 the dependence on the plasma density drops out. Integrating Eq. 5.10 for a Gaussian envelope of spatial size σ, then

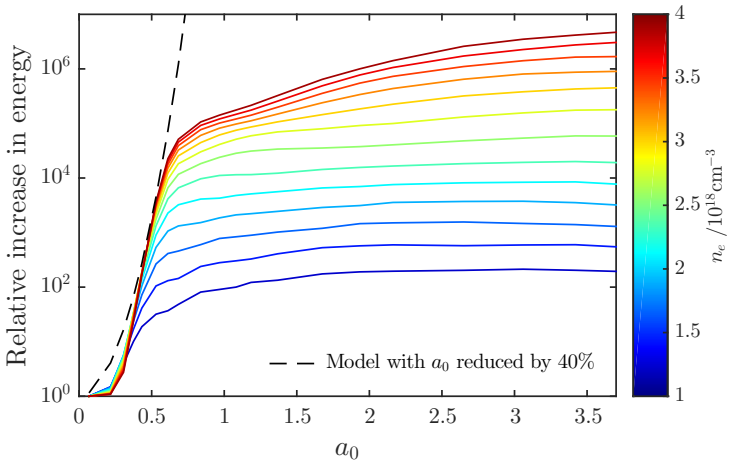


Figure 5.19: Probe beam energy amplification is plotted as a function of drive pulse amplitude for different densities. At low a_0, the amplification is observed to be almost independent of plasma density, and is comparable to the prediction of Eq. 5.14 which incorporates this observation.

$$\frac{\partial \log a_1}{\partial y} = \frac{k_0}{2} a_0^2 \quad \Rightarrow \quad a_1(y) = a_{1,in} \exp\left(\frac{k_0}{2}\int_{\infty}^{y} a_0^2 e^{-y'^2/\sigma^2}\,\mathrm{d}y'\right)$$
$$a_1(y) = a_{1,in} \exp\left(\frac{k_0 a_0^2 \sigma\sqrt{\pi}}{4}\left(1+\mathrm{erf}\left(\frac{y}{\sigma}\right)\right)\right) \tag{5.13}$$

where $a_{1,in}$ is the input probe amplitude. After fully traversing the drive pulse, $y \to \infty$ and the expected energy amplification ratio is

$$\frac{E_{\mathrm{out}}}{E_{\mathrm{in}}} = e^{2\pi^{3/2} a_0^2 \sigma/\lambda_0} \approx e^{86 a_0^2}, \tag{5.14}$$

which scales strongly with drive pulse amplitude but does not depend on plasma density. The appropriate value for σ is $(\sigma_x^{-2} + \sigma_y^{-2})^{-1/2} \approx 6.2\,\mu\mathrm{m}$, where $\sigma_{x,y}$ are the standard deviations of the drive pulse vector potential envelope in the longitudinal and transverse directions. As the beams are propagating through one another, the pulse size experienced by the probe beam is smaller than the actual pulse size. Expression 5.14 is plotted in Fig. 5.19 as a dashed line. To achieve reasonable agreement with the simulations the values of a_0 used in Eq. 5.14 were lowered by 40%. This would indicate that the probe does not pass through the peak of the drive pulse, which is observed in the simulations.

It is also possible to measure the growth of the probe directly as a function of propagation distance, plotted in Fig. 5.20. The analytic expression 5.13 is overlaid where again it was necessary to reduce the value of a_0 to achieve a good fit, but the spatial profile of the growth is well reproduced.

Finally, an intensity threshold is observed above which growth occurs much more slowly with increasing a_0. The first effect to consider at high intensity is a relativistic modification to the plasma electrons. The ω_p^2 term of Eq. 5.10 arises from the transverse currents induced by the drive pulse. At high intensity the effective electron mass is reduced by

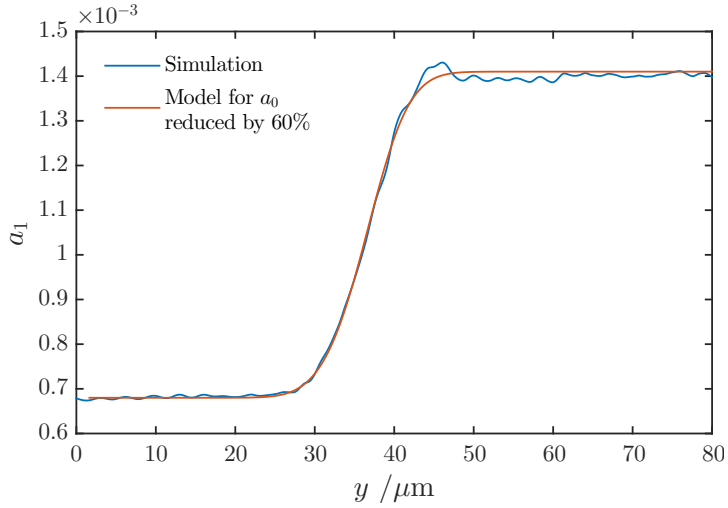


Figure 5.20: The spatial amplification of the probe beam amplitude as it crosses the drive pulse of amplitude $a_0 = 0.2$ is plotted from a simulation. Overlaid is the analytic expression 5.13.

$\gamma_0 \approx a_0^2/2$, and so Eq. 5.10 becomes

$$\frac{\partial a_1}{\partial y} = \frac{\omega_p^2 a_0}{2\omega_0 c\gamma_0}\frac{\delta n}{n} \approx \frac{\omega_p^2}{a_0\omega_0 c}\frac{\delta n}{n}. \tag{5.15}$$

The second effect is a saturation of the grating amplitude $\delta n/n$, which cannot increase indefinitely, implying that $\partial_y a_1 \to 0$ at high intensity and the probe growth saturates at some point as it transits the high-intensity drive pulse. If $\delta n/n$ remains approximately constant at high intensities throughout the interaction, then the amplified probe energy should increase with density. This is a reversal of the scaling of Eq. 5.11 expected at low intensities, and is observed in the simulations.

5.2.3 Probe wavelength scan

In the experiment and simulations discussed thus far, both lasers were at the same frequency. It is well established that a density grating is formed, which begs the question whether the 'amplification' observed is simply the Bragg scattering of the drive pulse from the grating. To address this question a final set of simulations were performed where the wavelength of the probe beam was varied. If the amplification were a direct reflection of the drive pulse, then the amplified light should be within the bandwidth of the drive pulse. The spectra of the amplified probe beams are plotted in Fig. 5.21 as a function of probe wavelength λ_{in}. The plasma density was 10^{19}cm^{-3} to ensure a relatively large spectral separation between the drive pulse wavelength (800 ± 20 nm) and the expected wavelength of Raman scattered light (865 nm). The peak intensity of the drive pulse was varied over several orders of magnitude, from $a_0 = 0.2$ to $a_0 = 5$.

In the low-intensity case where $a_0 = 0.2$, the probe spectrum appears to be unaffected during the amplification process. When the probe frequency fulfils the Raman matching condition $\omega_{\text{probe}} = \omega_{\text{drive}} - \omega_p$, i.e. $\lambda_{\text{in}} = 865$ nm, the maximum amplification is observed. This is then the normal resonant Raman amplification of the probe beam, albeit in a

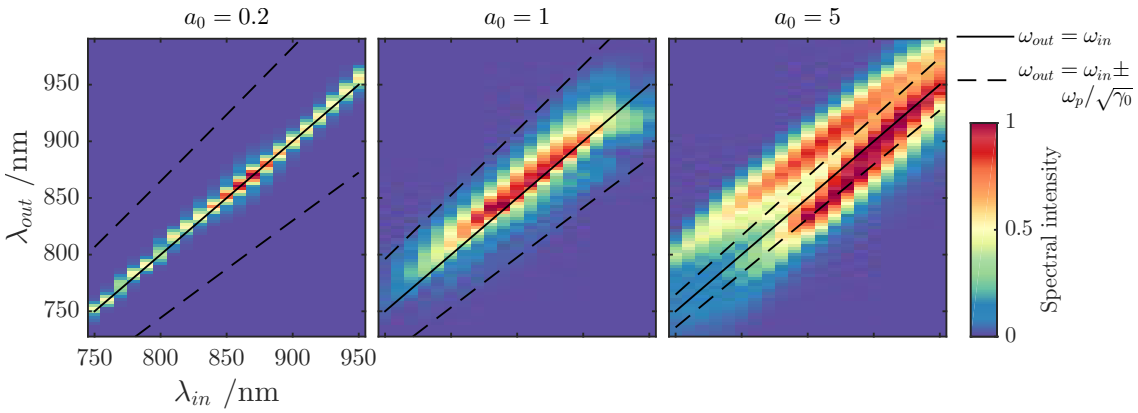


Figure 5.21: Simulated spectra of amplified probe beams as a function of probe beam input wavelength λ_{in}. The output wavelengths λ_{out} are observed to be related to the probe wavelength rather than the drive wavelength.

90° geometry. At relativistic intensities $a_0 = 1$ the amplification of the probe occurs over a broader range of wavelengths, but the amplified beam is observed to have similar wavelength to the incoming beam. There is a slight shift towards longer wavelengths, but the shift is very much lower than the plasma frequency, indicated by dashed lines. At an even higher intensity corresponding to $a_0 = 5$ the amplified wavelengths continue to increase with λ_{in}, but the spectra are quite different. When the probe has a shorter wavelength than the drive it tends to be shifted to longer wavelengths. As the probe wavelength increases beyond the Raman matching condition the amplified wavelengths lie nearer λ_{in}. There is not observed to be any significant stimulation of Raman scatter by the probe at any intensity, which agrees with the experimental observation that the RSS was unaffected.

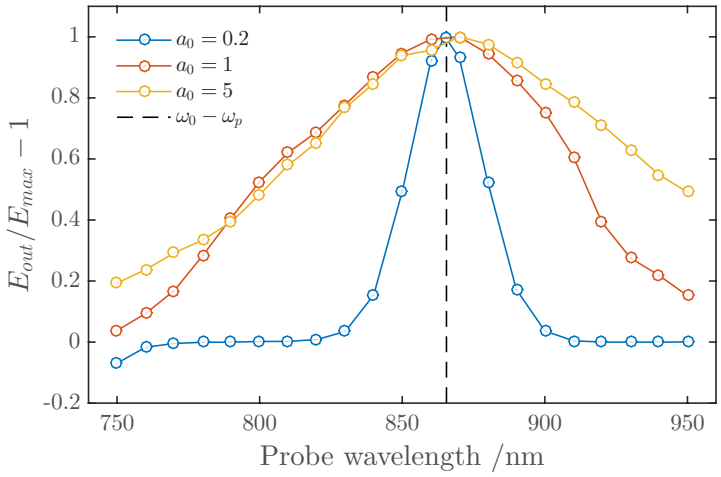


Figure 5.22: The amplified energy of the probe beam is plotted as a function of probe wavelength for different drive amplitudes, in each case normalised to the peak probe energy for a given drive amplitude. The bandwidth over which growth occurs is observed to increase with drive pulse intensity.

The energy amplification of the probe beam is plotted in Fig. 5.22 as a function of wavelength for the three cases. For each drive pulse intensity the curves were normalised

to their peak amplification factor. Here the interesting quantity is the spectral 'gain region' over which amplification is observed rather than the absolute amplification level. The dashed line marks the resonant frequency $\omega_0 - \omega_p$, and for all intensities amplification is maximised near there. At low intensity the amplification falls off rapidly as the probe wavelength is moved away from resonance, and near $\omega_0 + \omega_p$ the probe begins to transfer energy to the drive pulse. This is a consequence of the Manley-Rowe relations for three-wave interactions which state that energy is transferred to the lower-frequency wave. It is interesting that this relation continues to hold despite the many orders of magnitude difference in the intensities of the two lasers. As the pump intensity increases the bandwidth over which growth can occur is rapidly widened, including at and below the wavelength of the drive pulse.

A number of mechanisms exist by which the gain region may be broadened beyond that of perfect frequency matching. For the $a_0 = 0.2$ simulations, growth rates γ are plotted as a function of frequency mismatch from exact resonance $\delta\omega = \omega_{\text{probe}} - (\omega_0 - \omega_p)$. Here the growth rates are inferred from the total gain $G = \gamma T = \log(E_{\text{out}}/E_{\text{in}})/2$ where T is the growth time, assumed the same across simulations. Also plotted are several models for the dependence of the growth rate on the mismatch. In each case the growth rates are normalised to their peak values at resonance γ_{max}. The general nonlinear dispersion relation for parametric instabilities predicts the growth rate dependence[3]

$$\gamma(\delta\omega) = \gamma_{\text{max}}\sqrt{1 - \frac{\delta\omega^2}{4\gamma_{\text{max}}^2}} \tag{5.16}$$

which has hard cutoffs at $|\delta\omega| = 2\gamma_{\text{max}}$. In this case the broadening is due to the intrinsic bandwidth of a growing mode. For $\gamma_{\text{max}} = 0.16\omega_p$ and small $\delta\omega$ the shape of the gain curve is in reasonable agreement with the simulation data, but fails at larger $\delta\omega$.

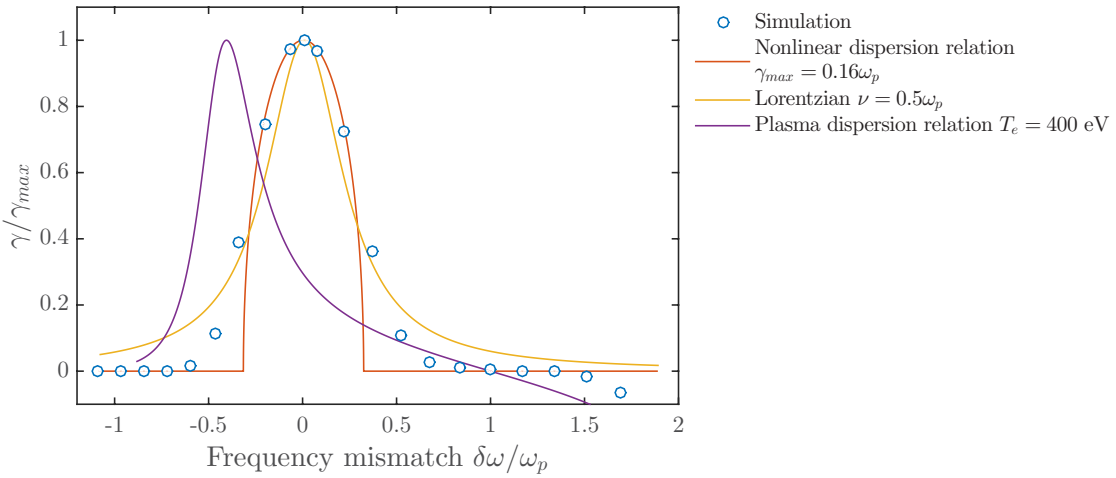


Figure 5.23: Comparison of inferred growth rates from the low-intensity simulation to several models.

In the analogous process of energy transfer between laser beams via Brillouin scattering[4],

[3] [279] K. Nishikawa. *Journal of the Physical Society of Japan.* (1968).
[4] [288] W. L. Kruer *et al. Physics of Plasmas.* (1996); [220] C. J. McKinstrie *et al. Physics of Plasmas.*

in the strong-damping limit the growth rate dependence is a Lorentzian

$$\gamma(\delta\omega) = \frac{\gamma_{max}}{1 + (2\delta\omega/\nu)^2} \tag{5.17}$$

where ν is the damping rate of the plasma wave. As in any oscillator, finite levels of damping act to broaden the resonance region. Generating a growth curve of similar width requires large plasma wave damping rates however, which would require an electron temperature $T_e > 500$ eV if Landau damping were to be responsible (see appendix A.2). This is a very high temperature for LWFA plasmas which typically have temperatures in the few eV range after field ionisation[5].

Finally, broadening can occur via thermal effects. Growth rates may be calculated numerically for a warm maxwellian electron distribution under the influence of crossing laser beams[6] — see appendix A.6. In this case sufficient broadening again requires high electron temperatures, so high that the plasma frequency is increased significantly beyond that of a cold plasma and the peak of the growth curve is shifted. It is possible that a relativistic reduction could cancel this effect, but such a coincidence seems unlikely.

If a gain curve of bandwidth B_g is assumed to take the same form as Eq. 5.16 and peak at a frequency ω_g, then after a growth time t the peak of the (assumed gaussian) probe spectrum is shifted by

$$\Delta\omega_{\text{peak}} = \frac{\omega_g - \omega_0}{1 + 2B_g^2/(\gamma_{\max} t B_p^2)} \tag{5.18}$$

where B_p is the bandwidth of the probe. For a moderate number of e-folding times $\gamma_{\max} t$, the requirement that the shift is small is $B_g \gg B_p$, i.e. the gain bandwidth is much larger than the probe bandwidth. This is clearly the case in the high-intensity regime, and helps to explain why a frequency shift is not observed.

5.3 Discussion

The exchange of energy between crossing laser pulses is not a new phenomenon in general, and has been discussed extensively in the literature in the context of Raman amplification, and also cross-beam energy transfer (CBET). This is a technique used in inertial confinement fusion (ICF) experiments for redirecting laser energy from one beam to another via Brillouin scattering[7]. ICF makes use of many beams incident on a low-density plasma, so multi-beam parametric interactions have to be taken into account[8]. As the ion acoustic frequency is typically small compared to the electron plasma frequency, it is much more likely that beams of similar frequency can be phase matched to a Brillouin scattering process rather than a Raman scattering process, and significant levels of energy transfer are possible[9]. Much of the theoretical work in multiple-beam interactions has therefore been driven by a need to understand the energy transfer rates in an ICF context.

(1996).

5 [58] E. Esarey *et al. Reviews of Modern Physics.* (2009).

6 [289] P. Michel *et al. Physics of Plasmas.* (2013); [290] P. Michel *et al. Physical Review Letters.* (2014).

7 [291] J. D. Moody *et al. Nature Physics.* (2012).

8 [292] J. F. Myatt *et al. Physics of Plasmas.* (2014); [293] L. Yin *et al. Physics of Plasmas.* (2014).

9 [294] R. K. Kirkwood *et al. Physical Review Letters.* (1996).

While ion motion, and therefore Brillouin scattering, is not a factor in the short-pulse experiments discussed here, the physical concepts are also applicable to Raman scattering and so CBET results from the literature are included in the following discussion.

5.3.1 Why is frequency matching not important?

One of the puzzling observations from the Gemini experiments and simulations was that the two-pulse interaction was so strong, even when classically the energy transfer should drop sharply when away from resonance. The problem of the non-resonant interaction between beams of equal frequencies has been discussed in the context of CBET, in most detail by McKinstrie *et al.*[10]. The resonant interaction, where the two beams have their frequencies offset by the ion acoustic frequency, maximises the steady-state energy transfer as expected. Before reaching steady-state the plasma develops a transient response where it is observed that energy is alternately transferred between the two beams at the ion acoustic frequency[11]. In a separate simulation paper Kruer notes that[12] "*the two large waves beat to set up a density modulation, which, in fact, develops a component at the acoustic frequency via the transient response*". Though neglected in the linear theory, the transient response of the plasma is important in determining the energy transfer between beams in the initial interaction period. This phenomenon was also observed in the simulations performed here, where because the laser crossing time was similar to the plasma period, the entire interaction effectively occurs during the transient plasma response.

As available laser pulse lengths decreased it became more important to understand the transient response, and the initial spatio-temporal behaviour of the Raman and Brillouin instabilities is discussed in a series of theoretical papers by Hinkel *et al.*[13] and Mounaix *et al.*[14]. In the frame of the drive laser, fresh plasma is continuously streaming into the pulse. For a sufficiently short pulse, the long-term evolution of the plasma wave and side or back-scattered light is unimportant as it rapidly convects out of the pulse. The early-time growth rate of the instability is converted to a constant convective growth from the leading edge of the drive pulse in the moving frame. Importantly the spatial growth rate in this frame is found to be weakly affected by significant frequency detuning, and the resulting gain curves are found to have a broad frequency width. This is another statement of the fact that exact frequency matching is unimportant for the evolution of the instability during pulses short comparable to the growth time. Simulations of Raman backscatter in the superradiant regime also indicate the existence of a brief superradiant 'spike' in the scattered energy, the intensity of which is many times larger than that of the seed amplified in the steady-state[15].

The second effect to consider is the motion of the plasma electrons which, irrespective of

[10] [220] C. J. McKinstrie *et al. Physics of Plasmas.* (1996).
[11] [295] V. V. Eliseev *et al. Physics of Plasmas.* (1996).
[12] [288] W. L. Kruer *et al. Physics of Plasmas.* (1996).
[13] [296] D. E. Hinkel *et al. Physics of Plasmas.* (1994); [297] D. E. Hinkel *et al. Physics of Plasmas.* (1995).
[14] [298] P. Mounaix *et al. Physics of Fluids B: Plasma Physics.* (1993); [299] P. Mounaix *et al. Physics of Plasmas.* (1994).
[15] [211] G. Shvets *et al. Physics of Plasmas.* (1997).

the pulse length, becomes highly nonlinear at high laser intensities. When the crossing laser pulses have equal frequencies, the electrons are observed to become trapped in the ponderomotive potential and perform anharmonic oscillations. This trapping process has been previously explored in the literature where it is known as the 'trapped particle instability'[16]. The subsequent electron motion has previously been studied theoretically[17] where it was noted that the mismatch between the plasma and electron oscillation frequencies caused a slow nonlinear evolution of the plasma wave. From the simulations of section 5.2.1 the electron motion was found to consist of a broad range of frequencies, distributed fairly evenly up to the bounce frequency. The width of the Raman growth curve in frequency space is due to the distribution of electron oscillation frequencies. If electrons only oscillate at frequencies near ω_p, then superposing motion at ω_0 can only generate new frequencies near some combination of the two. If the electron population oscillates at a range of frequencies extending towards zero, which is only possible if they are trapped in a strong potential, then the possible range of radiation frequencies is increased.

The radiation emitted by electrons trapped in gratings has been calculated by several authors[18], where there is found to be substantial interplay between the electron motion at the laser and bounce frequencies, generating a broad spectrum of radiation when the motion is relativistic. The same process occurring in the relativistically-boosted frame of a high-energy electron bunch is described by Balcou *et al.* as a Raman free-electron x-ray laser[19]. In this case the bounce motion couples to the laser motion through the collective oscillations of the bunch, leading to a Raman-type instability which generates radiation in the forward direction (this differs from the usual definition of the 'Raman' regime in XFELs in that the electron oscillations are transverse to the beam direction rather than parallel to it).

Finally, other nonlinearities which couple the motion of oscillators can come into play at high intensity. One example is known as parametric autoresonance[20] where if at any point a coupled oscillator system is near a resonance, the oscillators can 'lock' and stay in phase despite any subsequent detuning. This effect has been observed in simulations to enhance the Raman scattering process[21], where there is a characteristic phase locking of the three waves and increase in Langmuir wave amplitude despite the presence of a strong density gradient.

The combination of short pulse lengths and high field amplitudes used in the present experiments then causes the plasma response to be non-resonant and the oscillatory electron motion to be non-harmonic. The conventional definitions of Raman growth rates and gain bandwidths are derived under the assumption that the electron wave grows resonantly and the electron motion is oscillatory at ω_p, neither of which can be

[16] [272] W. L. Kruer *et al. Physics of Fluids.* (1969); [300] K. Mima *et al. Journal of the Physical Society of Japan.* (1971).
[17] [301] B. I. Cohen *et al. Physics of Fluids.* (1977).
[18] [302] S. Sepke *et al. Physical Review E.* (2005); [303] S. B. Dabagov *et al. Physical Review ST-AB.* (2015).
[19] [304] P. Balcou. *European Physical Journal D.* (2010); [305] I. A. Andriyash *et al. Physical Review Letters.* (2012); [306] I. A. Andriyash *et al. Physical Review ST-AB.* (2015).
[20] [307] E Khain *et al. Physical Review E.* (2001).
[21] [308] T. Chapman *et al. Physical Review Letters.* (2012).

assumed here. The phase-mixed electron population supports the radiation of a range of frequencies, and the plasma grating is always encountered by the probe pulse in its initial amplifying phase. The probe pulse is then effectively a high-amplitude seed within the broad gain bandwidth of the transient Raman response, and is able to be amplified without a significant change in frequency. The observed amplification is likely occurring in the superradiant regime, which is a purely transient process.

5.3.2 Raman amplification experiments

To put the measured amplification into context, the results of recent Raman amplification experiments available in the literature are summarised in table 5.2. The group working at Princeton were the first to begin experiments[22] on Raman amplification in 2002, though at the time the accessible lasers suffered a large frequency difference which necessitated the use of a high-density plasma to achieve phase matching. The consequent high growth rates of deleterious plasma instabilities limited amplification, even in very short plasmas. In the following years the group employed lasers with much more similar frequencies[23], facilitating the use of lower density plasmas and eventually reaching amplification factors of up to[24] 10,000. These levels of amplification required reaching the non-linear stage of Raman amplification, where the seed pulse becomes appreciably compressed and pump depletion becomes significant. Subsequent experiments at Lawrence Livermore National Laboratory explored the amplification of longer seed pulses using longer pump pulses[25], though with limited success compared to earlier experiments. The use of a waveguide to increase coupling efficiency between the pulses has also been explored[26] with some degree of success.

In contrast, there has been limited experimental investigation into the superradiant Raman amplification regime, which requires pump pulses of higher intensity. One experiment at MPQ looked explicitly for signs of superradiant amplification in relatively low-density plasmas[27], thus keeping ω_B/ω_p large. The signature of superradiance was the formation of several amplified pulses of duration π/ω_B separated by $2\pi/\omega_B$, and the fact the seed was spectrally broadened beyond that expected from the linear Raman regime. An experiment was also performed at the IAP in Russia [28], ostensibly investigating standard backwards Raman amplification. However the available pump and seed beams were at the same wavelength, and thus the plasma density was chosen to be very low ($1\times\ 10^{16}$ cm^{-3}) to ensure phase-matching within the available bandwidth of the pump. This compromise actually caused ω_B/ω_p to be larger than any other reported experiment, and thus far the Balakin experiment is the one which has probed the farthest into the superradiant regime. The principle results of these experiments, namely the reported amplification, are plotted in Fig. 5.24 as a function of the ratio ω_B/ω_p. Blue circles mark the ratio as calculated for

[22] [309] Y. Ping *et al. Physical Review E.* (2002).
[23] [310] Y. Ping *et al. Physical Review E.* (2003); [311] Y. Ping *et al. Physical Review Letters.* (2004); [312] W. Cheng *et al. Physical Review Letters.* (2005).
[24] [313] J. Ren *et al. Nature Physics.* (2007).
[25] [314] R. K. Kirkwood *et al. Physics of Plasmas.* (2007); [315] Y. Ping *et al. Physics of Plasmas.* (2009).
[26] [316] C.-H. Pai *et al. Physical Review Letters.* (2008).
[27] [317] M. Dreher *et al. Physical Review Letters.* (2004).
[28] [318] A. A. Balakin *et al. Journal of Experimental and Theoretical Physics Letters.* (2004).

Experiment	λ_0/nm	$a_0/10^{-3}$	λ_1/nm	$a_1/10^{-3}$	L_{pl}/mm	n_e/cm^{-3}	$I_{\mathrm{out}}/I_{\mathrm{in}}$
Ping, 2002	532	1.3	745	0.9	1-3	1×10^{20}	5-10
Ping, 2003	800	7.1	850-920	2.6	1	1.2×10^{19}	2
Balakin, 2004	800	5	800	0.7	15	2×10^{16}	100
Ping, 2004	800	8.7	840-940	1.4	2	1×10^{19}	10-100
Dreher, 2004	793	53	815	8.2	1	3.5×10^{18}	17
Cheng, 2005	800	7.1	873	1	2	1.1×10^{19}	1,000
Ren, 2007	800	10.7	878	0.9	2	1.3×10^{19}	10,000
Kirkwood, 2007	1054	31	1200	2.4	2	1×10^{19}	25
Pai, 2008	810	30	862	2	9	6×10^{18}	910
Ping, 2009	1054	10	1200	0.6	3	1×10^{19}	11
Gemini, 2012	808	4,500	808	0.2	0.05	3×10^{18}	100,000
Gemini, 2013	808	3,800	808	1.2	0.05	7×10^{18}	7,000

Table 5.2: A survey of plasma Raman amplification experiments where $\lambda_{0/1}$ and $a_{0/1}$ are the central wavelengths and amplitudes of the drive/probe beams, L_{pl} is the length of the plasma, n_e is the average plasma density and $I_{\mathrm{out}}/I_{\mathrm{in}}$ is the reported intensity amplification of the probe pulse.

the input parameters reported for the laser pulses, and red triangles for the amplified pulses. In all cases pump depletion is assumed negligible, which is perhaps a poor assumption for the 2007 Ren experiment but otherwise reasonable. In comparison, the experiments reported on here are notable for several reasons. Compared to previous experiments the intensity of the drive pulse is several orders of magnitude larger, and the pulse length much shorter. Coupled with the relatively low plasma density employed, the ratio ω_B/ω_p was larger for these experiments than any other, and the superradiant approximation that the plasma motion is dominated by the laser fields should be a very good one. The fact that the pulses crossed at 90° rather than 180° limited the interaction length to be much shorter than in Raman amplification experiments, and the spatial intensity amplification rate is up to a thousand-fold more rapid than previously published results.

While this sounds impressive, one should question the utility of such a result. The fact that there is such a huge disparity between the intensities of the two pulses here means that the probe can become hugely amplified at little expense to the drive. The efficiency of the process, defined as the fraction of the drive energy transferred to the probe, was maximised in experiment III and approached 0.1%. The highest reported efficiency for Raman amplification is[29] 6.4%, and was achieved with a much lower intensity drive pulse. On the other hand, there are some advantages to this 90° geometry. One of the limiting factors[30] on maximum achievable intensity in BRA is the relativistic detuning of the 3-wave resonance. As the probe pulse propagates through the plasma there is a cumulative nonlinear phase shift between the probe and plasma waves due to the slight relativistic motion of the plasma electrons. At large enough phase shift the direction of energy transfer is reversed and the probe begins to be depleted. This reversal happens over much longer

[29] [313] J. Ren *et al. Nature Physics.* (2007).
[30] [203] V. Malkin *et al. European Physical Journal: Special Topics.* (2014).

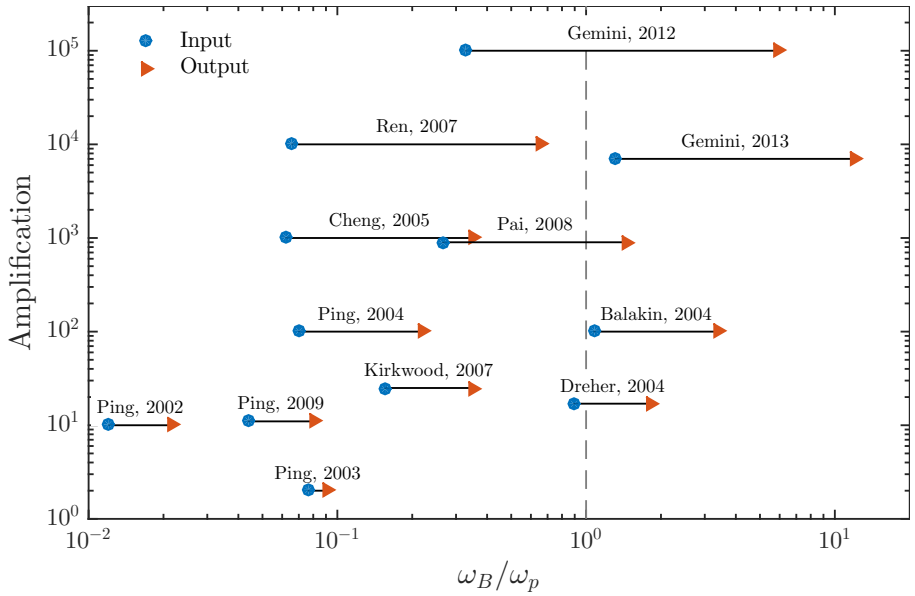


Figure 5.24: A survey of Raman amplification experiments, where circles and triangles represent pre– and post–amplification respectively. The dashed line marks the conventional transition to the superradiant regime.

timescales than ω_p^{-1}, and eventually the amplification of the probe can be completely negated. The cumulative phase shift can be reduced by lowering the plasma density, but a limit is eventually reached as the plasma wave breaks more easily at lower density. To prevent premature wavebreaking, the electron density should satisfy

$$\frac{n_e}{n_c} > \left(\frac{I_{\text{drive}}}{2.6 \times 10^{17}\,\text{Wcm}^2} \right)^{2/3}. \tag{5.19}$$

At lower density the Raman growth rate is also lower, so to achieve the same total gain the length of the plasma must be increased. This comes at the expense of decreased drive energy due to premature Raman backscattering of the drive pulse, and a compromise must be found. The ideal situation is then to have large growth in a short region of plasma where wavebreaking, detuning and pump depletion effects are not important. By using a highly intense and short pulse the experiments here have circumvented these issues, which helps to explain why the observed scattering was so strong.

From simulations the amplification was observed to be limited here by the presence of the bubble, as the strong density gradients appeared to inhibit the propagation and growth of the probe. If the same drive pulse laser energy is focussed into a larger focal spot, the peak intensity is smaller and the bubble takes some time to develop. The distance over which growth can occur is larger, and simulations show that the total amplification is greater despite the lower peak intensity of the drive pulse.

5.3.3 Use as a diagnostic

The measured probe amplification was a strong function of the pulse crossing position, and was maximised when the pulses crossed in the region where Raman side scatter was produced. From simulation the probe is observed to grow over the regions outside of the

bubble, and is therefore sensitive to the unguided pulse energy. From the previous chapter, the unguided portion of the pulse is also responsible for the Raman side scatter, and so it should be expected that the two occur in the same region of the plasma. As the drive pulse propagates it self-focusses and scatters unguided energy until it is enclosed within the bubble. At this point Raman side-scattering and probe amplification cease, as is observed experimentally. The timing scans of Fig. 5.7 are therefore sensitive to the unguided portion of the drive pulse and so the probe amplification level can act as a diagnostic of the pulse evolution. While not a single-shot diagnostic, a set of timing scans performed over a range of plasma densities would provide useful information about the pulse evolution rate as a function of density. This is an important quantity to understand as it determines the electron injection point, and therefore the effective acceleration length of the accelerator. From the shadowgrams of Fig. 5.2 this evolution length is observed to increase with decreasing density, but it would be interesting to measure this quantity more precisely. It is probable that the length of the density ramp is also important in determining the pulse evolution length, which could be measured by comparing timing scans between different gas jets at the same density.

Finally if the transverse imaging resolution and camera dynamic range were high enough, it should be possible to locate the position of the drive pulse in the shadowgrams to within several tens of microns using this technique. By changing the probe delay by known amounts it would then be possible to directly measure the group velocity of the drive pulse as it propagates. The electron injection threshold and energy gain are linked to the group velocity, and knowing how it evolves through the plasma would provide useful information on the wakefield dynamics.

5.4 Conclusion

The interaction between two laser pulses in a laser wakefield accelerator is observed to be a surprisingly dramatic process, even though the fractional perturbation to the laser intensity was as low as 1 part in a billion. The weak probe pulse is observed to increase in intensity by a factor of up to 10^5 within its own bandwidth, though it only extracts up to 0.1% of the energy of the drive pulse in the process. From simulation the amplification is attributed to the non-resonant Raman scattering of the drive pulse from a ponderomotively-driven electron density grating. The drive pulse is of sufficiently high intensity that the amplification process is expected to be well within the superradiant regime, for which phase-matching of the two beams is unimportant. The 90° geometry of the interaction differs from the design of conventional Raman amplification schemes, and may mitigate the effects of detuning.

By changing the point at which the two pulses cross, the observed amplification is a function of the amount of unguided drive pulse energy, and thus probes the evolution of the drive pulse. This mechanism may then function as a useful diagnostic tool for laser wakefield accelerators, in addition to being by far the most rapid reported example of Raman amplification.

6 Bone Tomography

In this chapter we discuss the use of a laser wakefield accelerator in a potential clinical application, namely the 3D x-ray imaging of a human bone sample. The data presented herein were recorded during an experiment at the end of 2013 using the Astra-Gemini laser at the Central Laser Facility, and published in 2015[1].

This work is motivated in the context of the diagnosis of osteoporosis in section 6.1, which is presently driving a need for high brightness, high resolution, compact x-ray light sources. The details of the experimental design are described in section 6.2, before a discussion of the results in section 6.3. When optimised for stability and brightness of x-ray beams, the LWFA was found to generate beams containing $1.3 \pm 0.5 \times 10^9$ photons at critical energy $33 \pm 12\,\mathrm{keV}$ on 97% of laser shots. This led to the successful tomographic scan of the bone sample at a 3D resolution of up to $50\,\mu\mathrm{m}$.

In section 6.4 the quality of the tomogram is discussed and compared to that recorded with a conventional x-ray source. With future development and optimisation of the LWFA source in mind, the chapter is closed with a discussion on the potential future steps towards clinical applications of this technique.

6.1 Experimental motivation

Since the first experiments involving betatron x-ray sources[2] it was observed that LWFAs were naturally capable of generating hard, collimated and bright x-ray beams. Soon afterwards the betatron x-ray source size was additionally measured to be extremely small[3], and therefore the intrinsic resolution of point-projection betatron radiography was known to be high. An important corollary is that the spatial coherence length of the betatron beam becomes large after a short propagation distance, leading to the first experimental phase-contrast images of insects using betatron radiation[4].

While these demonstrations were an important advance in the field, the chosen x-ray imaging targets were of less interest to the wider scientific community. The relatively low energies of the betatron beams at the time limited application to thin and transparent (low-Z) targets. Moving to thicker and more opaque targets is interesting as it increases the range of objects which can be imaged. Modern x-ray radiography is applied in a

[1] [108] J. M. Cole *et al. Scientific Reports.* (2015); [319] J. M. Cole *et al. Plasma Physics and Controlled Fusion.* (2016).
[2] [81] A. Rousse *et al. Physical Review Letters.* (2004).
[3] [184] M. Schnell *et al. Physical Review Letters.* (2012); [89] S. Kneip *et al. Nature Physics.* (2010).
[4] [105] S. Kneip *et al. Applied Physics Letters.* (2011); [104] S. Fourmaux *et al. Optics Letters.* (2011).

vast number of fields, the imaging of batteries[5], composite materials[6], fossils[7], coral[8] and mineral samples[9]. For the purposes of demonstrating the capabilities of betatron beams as a new light source, here the imaging of medically-relevant tissue samples was chosen as a goal. These samples require x-ray energies of several tens of keV, reasonable flux and a small source size. Assuming the success of 2D radiography, the second goal was to move to 3D imaging and full computerised tomography (CT)[10], where 2D projection images are acquired over a range of projection angles and the sample digitally reconstructed. The stability and consistency of betatron beams compared to conventional x-ray sources has always been a question, so experimental proof that tomography is possible goes some way towards demonstrating the capabilities of the source, as it is required to perform well during the acquisition of hundreds of images. While 3D imaging is possible for both absorption-contrast and phase-contrast radiography, we restrict ourselves here to absorption contrast. The team at MPQ Garching have recently demonstrated that phase-contrast tomography with betatron beams is also possible[11], albeit of small and transparent objects.

6.1.1 Choice of imaging object

There is a widespread need for micron scale x-ray imaging of human bone. As the population ages, degenerative diseases such as osteoporosis are becoming increasingly common and costing health services billions of pounds per year[12]. This disease is characterised by bone tissue deterioration caused by an imbalance in bone remodelling rates, leading to lowered bone density and increased fracture risk. Trabecular, or cancellous bone is a porous form of bone responsible for the transmission of forces from the joints to load-bearing surfaces. It is composed of a connected network of trabeculae, thin 100-500 µm diameter rods of calcified material, surrounded by soft marrow tissue. Its high surface area means it is host to many of the bone osteoclasts and osteblasts (cells responsible for bone remodelling), and therefore the site of a large fraction of bone metabolic activity. As such trabecular bone is particularly sensitive to any metabolic disruption[13], as occurs in old age or when using drugs to combat osteoporosis.

Traditional diagnosis of osteoporosis is based on the deviation from the mean of an individual's average bone density, which is an accurate predictor of fracture risk over the population. However for a particular patient up to 90% of the variation in bone strength is unexplained by average bone density [14], and it is becoming increasingly clear that the trabecular micro-architecture and distribution of mineralisation in the bone play an important role in determining individual fracture risk[15]. In addition, certain classes of

[5] [320] I. Manke *et al. Applied Physics Letters.* (2007).
[6] [321] P. J. Schilling *et al. Composites Science and Technology.* (2005).
[7] [322] P. C. J. Donoghue *et al. Nature.* (2006).
[8] [268] R. C. Roche *et al. Journal of Experimental Marine Biology and Ecology.* (2010).
[9] [323] F Mees *et al. The Geological Society of London.* (2003).
[10] [324] G. N. Hounsfield. *The British Journal of Radiology.* (1973).
[11] [107] J. Wenz *et al. Nature Communications.* (2015).
[12] [325] S. W. Blume *et al. Osteoporosis International.* (2011).
[13] [326] B. D. Snyder *et al. Calcified Tissue International.* (1993).
[14] [327] M. J. Ciarelli *et al. Journal of Orthopaedic Research.* (1991); [328] J. Hazrati Marangalou *et al. Osteoporosis International.* (2014).
[15] [329] M. Kleerekoper *et al. Calcified Tissue International.* (1985); [330] L. D. Carbonare *et al. Journal of Endocrinological Investigation.* (2004).

drugs traditionally used in treating osteoporosis may affect patients in different ways, and it is important to understand how their bone properties are affected[16].

If the internal microstructure of bone is known then additional metrics are available to use when comparing samples, such as trabecular volume, trabecular thickness, average porosity and even topological details such as the Euler characteristic of the trabecular connectivity network. If spatial resolution is sufficiently high it become possible to map the distribution of micro-cracks[17], and if the x-ray spectrum is well-controlled also the bone mineralisation[18].

An open question in current clinical research is how the interaction of the bone mineralisation and micro-crack repair mechanisms affects the load-bearing capacity of certain bone structures. Drugs currently used to treat osteoporosis increase the average mineral content of bone, but may also increase the number of micro-cracks in the bone, the prevalence of which are thought to correlate with fracture risk. There is however some debate on whether or not micro-cracks facilitate dissipation of stress within the bone, and actually prevent larger fractures, and which metrics should be used when defining bone quality[19]. These questions can only be answered with longitudinal studies of the 3D properties of bone at the micron scale.

Detailed structural and compositional information of trabecular bone is therefore required on a per-patient basis, which cannot be accomplished with a purely 2D examination[20] such as that delivered by traditional radiography or histology. As a highly anisotropic and opaque object composed of micron-scale structural elements, trabecular bone presents a significant imaging challenge. Micro-computed tomography[21] (µCT), commonly defined as CT with a spatial resolution below 100µm, has rapidly become the dominant method for imaging the internal microstructure of bone[22] and has proven to provide precise quantitative evaluations of its 3D morphology[23]. Alternative non-invasive imaging techniques such as ultrasound or MRI lack the combination of high spatial resolution and penetrating power which is required.

The µCT imaging of human trabecular bone then presents an ideal experimental challenge for a proposed light source. The resolution, photon energy and photon flux need to be high, and there is pressing clinical need for access to suitable laboratory-scale sources.

6.1.2 Suitable x-ray sources

The dominant x-ray source in use today for µCT is the microfocus x-ray tube, composed of a 30-200 keV electron gun focussed onto a high-Z anode material. The anode radiates x-ray photons at a characteristic K_α energy for lower-Z materials, along with a broad bremsstrahlung component which increases in flux with[24] Z. Scanners used for bone

[16] [331] B. Borah *et al. Bone.* (2005).
[17] [332] P. J. Thurner *et al. Bone.* (2006); [333] A. Larrue *et al. PloS one.* (2011).
[18] [334] S. Nuzzo *et al. Medical Physics.* (2002).
[19] [335] M. L. Bouxsein. *Osteoporosis International.* (2003).
[20] [336] A Odgaard *et al. Bone.* (1993).
[21] [337] J. C. Elliott *et al. Journal of Microscopy.* (1982); [338] B. P. Flannery *et al. Science.* (1987).
[22] [339] E. L. Ritman. *Annual Review of Biomedical Engineering.* (2011).
[23] [340] C. Chappard *et al. Osteoarthritis and Cartilage.* (2006).
[24] [341] M Green *et al. Proceedings of the Physical Society.* (1961).

imaging typically use tungsten anodes where the bremsstrahlung dominates the x-ray spectrum. By focussing the electron beam more strongly, the x-ray source size is reduced and spatial resolution is increased. If the beam is too strongly focussed the anode will begin to melt[25], which places an upper limit on the electron beam density and therefore a lower limit on the x-ray source size. Some scanners employ rotating anodes [26] to lower average heat loads at the expense of larger source sizes. Others use anodes formed of liquid metal jets[27], which places strong restrictions on anode materials. This brightness limitation places a lower bound on total scan times – typically one to several hours for the type of bone sample described here.

Where there is a need for higher-quality imaging, researchers are increasingly performing µCT scans at synchrotron light sources[28], where the high average source brightness is useful in several ways. More projections may be acquired in a given time, increasing scan resolution, or the same scan may be performed more quickly, increasing sample throughput. It is also possible to monochromate the x-ray beam while maintaining a high photon flux, which readily allows quantitative tomographic measurements. Bone tomography has been performed at the TOMCAT beamline at a rate of 600 exposures per second with a flux of 10^5 photons per detector pixel per exposure[29], which is clearly far in excess of what is possible on any other laboratory-based source.

Unfortunately access for routine µCT imaging is limited due to the relative scarcity and over-subscription of suitable light sources. As these synchrotron sources are driven by conventional accelerators, it is difficult to see how the size, and therefore cost, of such facilities will be reduced in the near future. Additionally, for applications such as bone imaging, the high brightness of a synchrotron source may be significantly more than can be used effectively, and samples can be damaged by the x-ray beam if care is not taken[30]. This suggests the need for an x-ray source which can deliver higher flux than conventional microfocus tubes, without necessarily reaching the extreme average brightnesses of large synchrotrons.

A recent hybrid approach which also shows promise is based on the inverse Compton scattering of laser photons from a 25–45 MeV electron beam[31], with an average flux of[32] 2.4×10^8 phs^{-1}mrad^{-2}. The source operates near 20 keV and produces an x-ray beam of similar relative bandwidth to the electron beam. Another laser-based approach where a few millijoule pulse is focussed onto a mercury jet at kHz repetition rates[33] achieves a photon flux of 5×10^5 phs^{-1}mrad^{-2}, which is similar to that of a low-power microfocus tube.

[25] [342] D. E. Grider *et al. Journal of Physics D: Applied Physics.* (1986).
[26] [343] A. Taylor. *Journal of Scientific Instruments.* (1949).
[27] [344] O. Hemberg *et al. Applied Physics Letters.* (2003).
[28] [345] F. Peyrin *et al. Technology and Health Care.* (1998); [346] M. Salome *et al. Medical Physics.* (1999).
[29] [347] R. Mokso *et al. AIP Conference Proceedings.* (2010).
[30] [348] W.-H. Kim *et al. Korean Circulation Journal.* (2008).
[31] [98] K. Achterhold *et al. Scientific Reports.* (2013).
[32] [349] E. Eggl *et al. Proceedings of the National Academy of Sciences.* (2015).
[33] [350] C. M. Laperle *et al. Applied Physics Letters.* (2007).

6.2 Experimental setup

The experiment reported in this chapter was performed over three weeks at the end of 2013, and was dedicated to the imaging of biological and medical objects with betatron x-ray radiation. The experimental configuration altered slightly between different imaging targets, so here we describe the layout specifically for the bone imaging section of the experiment, sketched in Fig. 6.1.

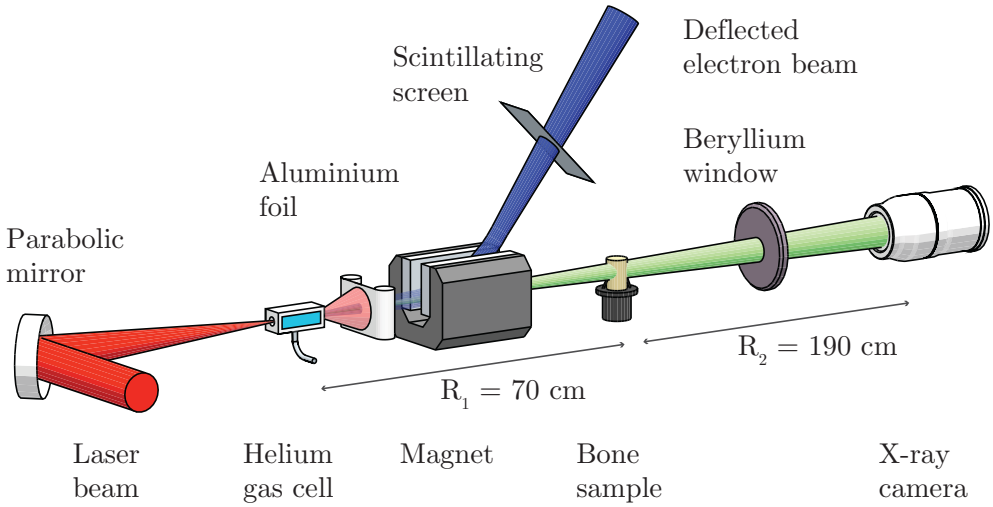
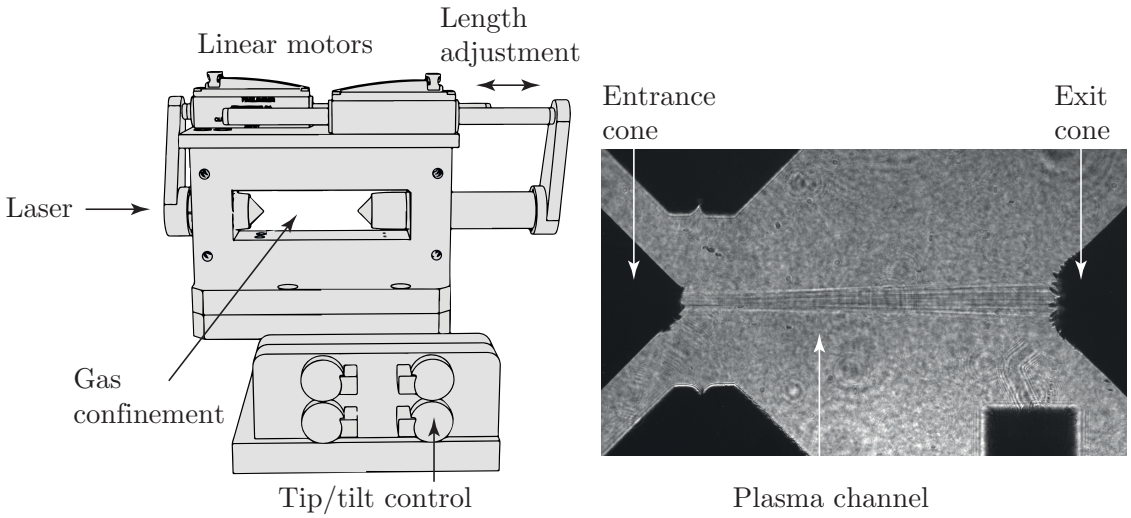


Figure 6.1: Beamline layout (not to scale)

The experiment took place at the Central Laser Facility and used the Astra-Gemini laser. The laser pulse was linearly polarised with a central wavelength of 800 nm and a temporal FWHM of 40 ± 3 fs. The pulse energy delivered to the plasma was 11.4 ± 0.4 J and was focussed with a 3 m focal length parabolic mirror, corresponding to an f-number of 20. The focal spot was typically slightly elliptical with a FWHM size of 23×28 µm. The resulting peak intensity at vacuum focus was $(3.2 \pm 0.1) \times 10^{19}$ Wcm^{-2}, or a normalised vector potential of $a_0 = 3.9 \pm 0.05$. The FWHM contour of the focal spot contained $27 \pm 3\%$ of the laser energy.

The pulse was focussed into the entrance of a gas cell, the design of which is illustrated in Fig. 6.2a). Several tens of milliseconds before the laser shot, fast gas valves were opened and gas allowed to flow into the cell at a pressure of hundreds of millibar. The gas pressure in the cell defines the plasma density. Pressure transducers close to the cell allowed the timing of the gas entry with respect to the laser pulse to be measured, ensuring the gas was delivered at the correct time. The laser enters and exits the cell through two small holes in the cone tips. The holes are small (≈ 250 µm diameter) to ensure the gas stays confined to the cell, though during the course of the experiment the holes can sustain laser damage and become enlarged, necessitating replacement. The length of the plasma is then varied by adjusting the distance between the two cones. Typically the latter cone was adjusted, to ensure the laser focal position with respect to the start of the plasma remained constant.

The plasma electron density was recorded on every shot with transverse Mach-Zehnder



a) Gas cell design

b) Shadowgram of gas cell

Figure 6.2

interferometry, where the second Gemini pulse was used as a probe beam. A highly intense probe beam was not required, so for this experiment the probe beam energy was approximately 250 mJ. Shadowgrams were simultaneously recorded on a separate imaging channel, an example of which is in Fig. 6.2b). Shortly after the exit of the gas cell was a 25 µm thick aluminium foil. This was used to completely block the laser pulse while still transmitting the majority of the x-ray energy. The transmission of 25 µm thick aluminium foil reaches 50% at a photon energy of 6 keV and 90% at 11 keV. The foil was translated after every shot to ensure that it stayed intact and that no laser light reached the samples. Following the foil was a 30 cm long permanent magnet, with a peak magnetic field of 1 T. The magnet performs two useful functions here – the deflected electrons are detected with a scintillating Lanex screen which allows the electron energy spectrum to be recovered as described in section 3.4. It also deflects the electron beam away from the sample, avoiding the production of bremsstrahlung which would otherwise degrade the x-ray image.

After the magnet the betatron x-ray beam freely propagates in vacuum to the bone sample. The sample is a 7 mm diameter cylindrical section of human femoral bone, interesting as an imaging target due to the high occurrence of femoral fractures in osteoporitic patients. A composite macro photograph of the sample is shown in Fig. 6.3, courtesy of Susan Parker. The bone was fixed to a rotation stage 70 cm from the exit of the gas cell. Attached to the sample rotation stage were two wires which rotated with the sample – see Fig. 6.7. By tracing the projected position of the ends of the wires, it is possible to accurately determine the axis of rotation. A third fiducial was fixed independently of the stage allowing correction for the small shot-to-shot fluctuations in the position of the x-ray source. The transverse and longitudinal positions of the sample were chosen to achieve high magnification while keeping the entire sample in view over all possible rotation angles.

A point projection image was formed on an x-ray camera a further 1.2 m away, a geometric magnification of 2.7. The camera was a Princeton Instruments PIXIS-XF, composed of a



Figure 6.3: Femoral bone sample, composite macro photograph courtesy of Susan Parker.

CsI scintillator fibre-coupled to a cooled 16-bit 2048×2048 array of $13.6 \times 13.6\,\mu m$ pixels. The detector does not operate in vacuum, so the end of the vacuum chamber on the beam axis was closed with a 180 µm thick beryllium window. The transmission of this window reaches 50% at 3 keV and 90% at 5.5 keV. To characterise the x-ray spectrum a 64-element filter array composed of various elemental metals with different K-edges was placed directly in front of the x-ray camera. Comparing the measured transmission through each filter with that expected from a synchrotron spectrum, the best-fit critical energy was obtained[34]. The beam profile was measured by optically imaging a CsI scintillator placed along the beam axis. The source size was measured by imaging a 1 mm-thick Si crystal at high magnification, and deconvolving the assumed gaussian intensity profile from the recorded line spread function. The analyses to recover the x-ray beam properties were performed by Jonathan Wood, the results reproduced here with permission.

6.3 Results

6.3.1 Electron and x-ray beams

One of the challenges associated with tomographic imaging is the need to produce a reliable x-ray beam. If a significant number of photons are not produced on every shot the total scan time would rapidly increase. Prior to the imaging effort a number of scans were performed over plasma density and gas cell length in an attempt to find a regime where energetic electron beams were reliably produced. For longer gas cell lengths the x-ray beam tended to become brighter, but the apparent source size increases due to the geometric smearing of the divergent x-ray beam being produced at different lengths within the gas cell[35]. At shorter lengths the electron acceleration was limited and peak electron

[34] [89] S. Kneip *et al. Nature Physics.* (2010).
[35] [351] S. Corde *et al. Plasma Physics and Controlled Fusion.* (2012).

beam energies reduced. At low plasma density the electron beam energies tended to be higher, but the charge was lower and the beams more inconsistent. Similarly the number of x-ray photons decreased with beam charge. At high densities the beam charge was higher but more noise was observed on the x-ray detector, possibly due to the more divergent electron beam producing bremsstrahlung radiation from the exit cone of the gas cell.
An optimum was found for a plasma density near $(2.9 \pm 0.2) \times 10^{18}\,\mathrm{cm}^{-3}$ and gas cell length 1.2 cm, and it was under these conditions that the tomographic scan was performed. The raw electron spectra for each shot in the imaging run are plotted in Fig. 6.4.

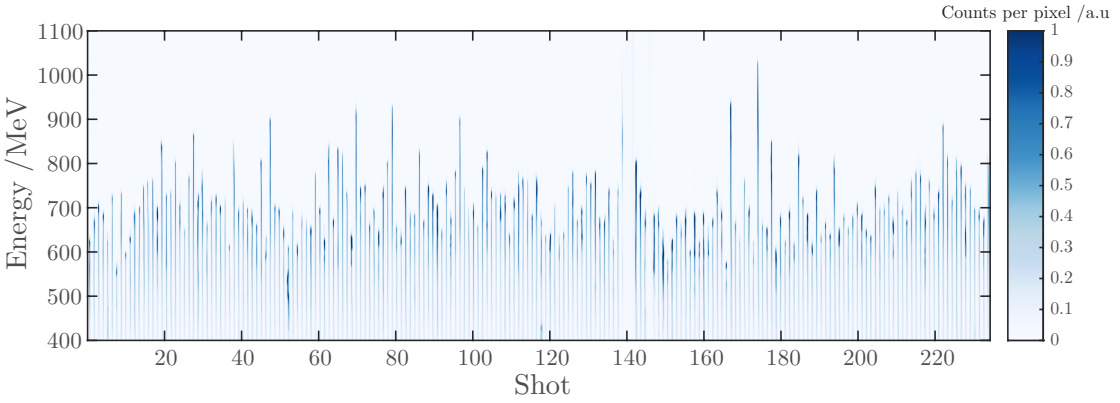


Figure 6.4: Raw spectrometer images scaled in the vertical direction such that the energy spacing is linear. All 234 images from the tomographic run are plotted.

In this regime the electron beams were reasonably stable, reaching energies above 400 MeV on 97% of the shots. The spectra were typically composed of a sharp peak at the maximum energy of the beam and a broad lower energy background. The mean energy of the peak was 711 MeV ± 9.5%, and in between 600 and 800 MeV on 87% of the shots. X-ray beams were detected on 227 shots, which amounts to 97% of the total laser shots. Without this level of consistency it would rapidly become impossible to perform a tomographic scan in a reasonable time period.
The x-ray beam energy was measured by comparing the relative transmission of different elemental filters and fitting a one-parameter synchrotron spectrum in a least-squares sense (performed by Jonathan Wood). Under the same parameters as the tomographic imaging run, E_{crit} was measured to be 33 ± 12 keV. The x-ray beam pointing fluctuations were $\pm 2.1 \times 1.5$ mrad (horizontal × vertical), and the total number of photons above 1 keV was measured as $1.3 \pm 0.5 \times 10^{9}$. On a different experiment under similar conditions the FWHM x-ray source size was estimated to have an upper bound in the range of 2 - 3 µm, consistent with the electron oscillation amplitude required to produce the observed spectrum as predicted by Eq. 2.94. The x-ray photon number and energy both varied by ±35%, significantly higher variations than the electron beam energy, which in turn varied more than the input laser energy. The total charge of the electron beam at energies above 400 MeV fluctuated by ±51%, which from Fig. 6.6 is slightly correlated with laser pulse energy. The nonlinearities present in the self-injection and betatron emission processes likely act to amplify any slight differences in initial conditions, emphasising the requirement that laser and plasma properties should be carefully controlled for imaging applications.

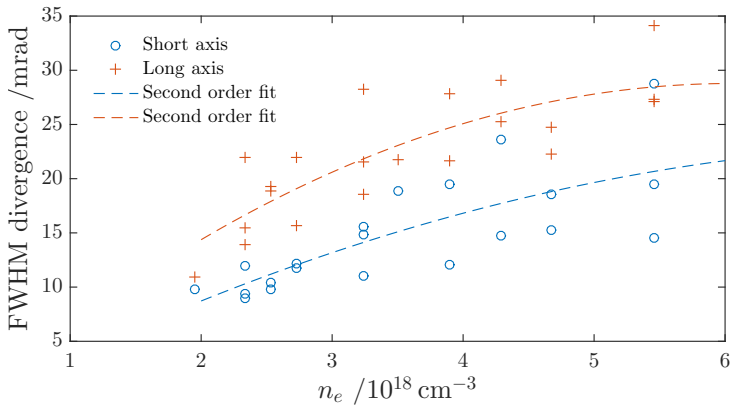


Figure 6.5: X-ray FWHM beam divergence measured for a 10 mm length cell as a function of plasma density.

The beam divergence was recorded on a separate day by imaging a scintillating screen placed on the beam axis. By fitting an ellipse to the FWHM contour of the beam shape, the long and short axes of the beam are plotted in Fig. 6.5. No significant x-ray beams were observed for densities below $2 \times 10^{18}\,\mathrm{cm}^{-3}$. From section 2.8 the divergence is expected to scale as $\Delta\theta = K/\gamma = r_\beta k_\beta$. If the electron energy is assumed to scale as n_e^{-1} then $\Delta\theta \propto r_\beta n_e$. The scaling observed here is slightly slower than linear with n_e, suggesting r_β falls with density. Thomas predicts the scaling[36] $r_{\beta\,\mathrm{min}} \propto n_e^{-1/8}$, the minimum betatron oscillation radius at peak electron beam energy. The x-ray beams tended to have significant ellipticity of 1.6 ± 0.4 with the longer axis oriented at $7 \pm 20°$ to the vertical (laser polarisation) direction. These properties were not observed to change significantly with density.
A number of metrics were recorded for each laser shot during the tomographic run, namely

1. the on-axis plasma density n_e from interferometry - first 103 shots,
2. the laser energy (before compression) E_L - all shots,
3. the laser FWHM duration from FROG measurement τ_L - 55 shots,
4. the energy of peak in electron spectrum E_{max} - all shots,
5. the electron beam charge above 400 MeV - all shots,
6. the number of counts on x-ray detector within the camera field of view - all shots.

These quantities are plotted against one another in Fig. 6.6. The laser and plasma parameters were nominally held fixed, so variation in E_L, τ_L and n_e comes from a combination of natural variation and measurement error. The raw data are plotted as grey points for every shot on which it was available. The data are binned into variable-width histograms such that each bin contains the same number of points. The mean of the bins are plotted as lines, and the shaded area represents ± 1 standard deviation for each bin. The laser duration and plasma density vary little over the course of the run, and no significant correlations are observed with respect to these parameters. The plasma density was observed to decrease steadily with time, probably because the entrance and exit

[36] [157] A. G. R. Thomas. *Physics of Plasmas*. (2010).

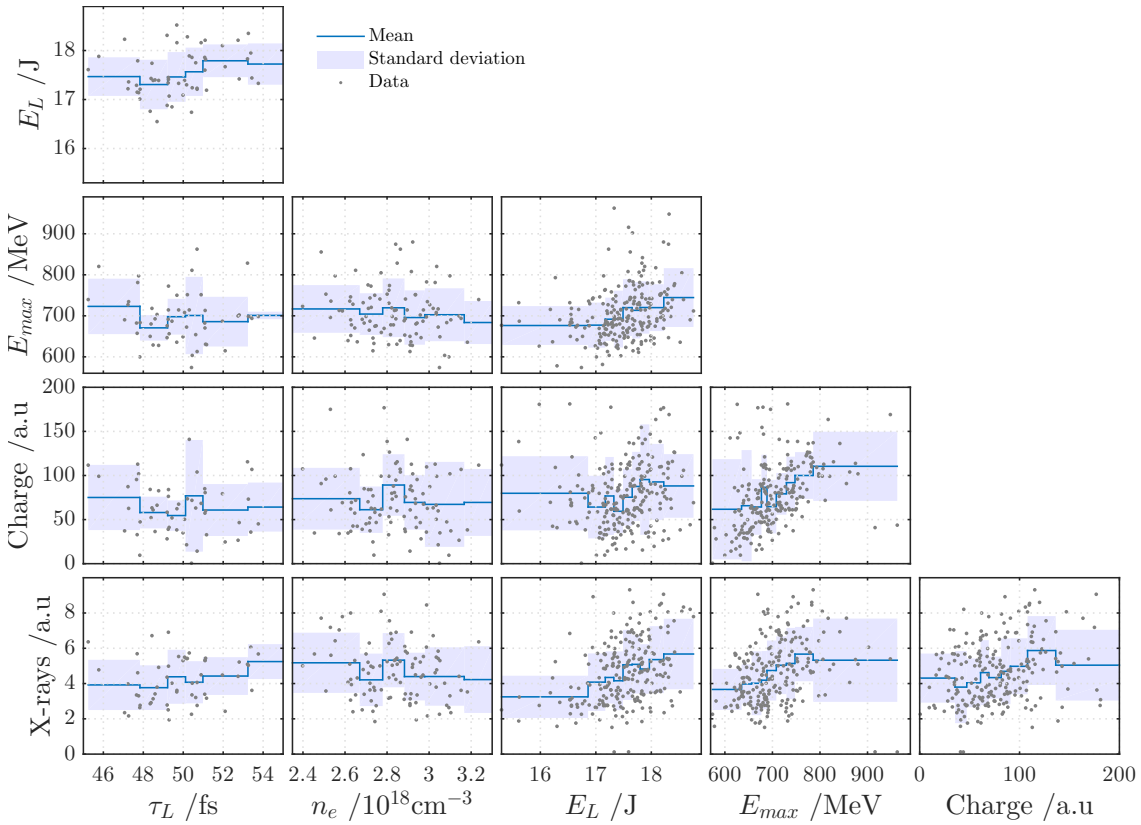


Figure 6.6: Correlations between laser, plasma, electron beam and x-ray beam parameters during the tomographic imaging run. The variations in the laser parameters are intrinsic, the laser energy was nominally held fixed. The plasma density dropped monotonically by 16% over the course of the scan.

apertures of the gas cell were slowly eroded over the course of the scan, allowing gas to leak more rapidly. Laser energy on the other hand varies significantly and appears to have a relatively strong impact on the peak electron beam energy, electron beam charge, and counts on the x-ray detector. The increased a_0 at larger pulse energy will tend to increase the matched size of the bubble and the depth of its potential well, accelerating injected electrons more rapidly. The electrons will likely be injected sooner which will increase the maximum beam energy, assuming the plasma length is shorter than the dephasing length. The electrons will also enter the bubble with larger transverse momentum which increases the total emitted x-ray energy, predicted to scale quite strongly[37] with a_0. The x-ray signal increases with electron beam energy to a point, but does not appear to increase significantly for electron energies above 800 MeV. Several of the highest energy electron beams were not observed to be associated with a significant betatron beam, which is surprising because the emitted synchrotron power is a strong function of γ. It is possible that for these beams the transverse motion of the electron bunch was less significant because it was injected with a smaller transverse momentum[38].

[37] [157] A. G. R. Thomas. *Physics of Plasmas.* (2010).
[38] [352] S. Corde *et al. Nature Communications.* (2013).

6.3.2 X-ray images

The mass absorption coefficient of the sample $\mu(x, y, z)$ determines the x-ray transmission $e^{-\int \mu \, \mathrm{dl}}$. Spatial variations in line-integrated μ therefore generate contrast in the radiographic absorption images. Four samples of the raw x-ray camera data are displayed in Fig. 6.7 with a number of features present

1. The horizontal 250 µm-thick tungsten wire, fixed independently of the sample
2. The upright 250 µm-thick tungsten wires rotating with the sample
3. The opaque edge of the steel beam tube
4. The hexagonal pattern due to the fibre-coupling of the x-ray scintillator to the camera
5. The non-uniform x-ray beam illumination

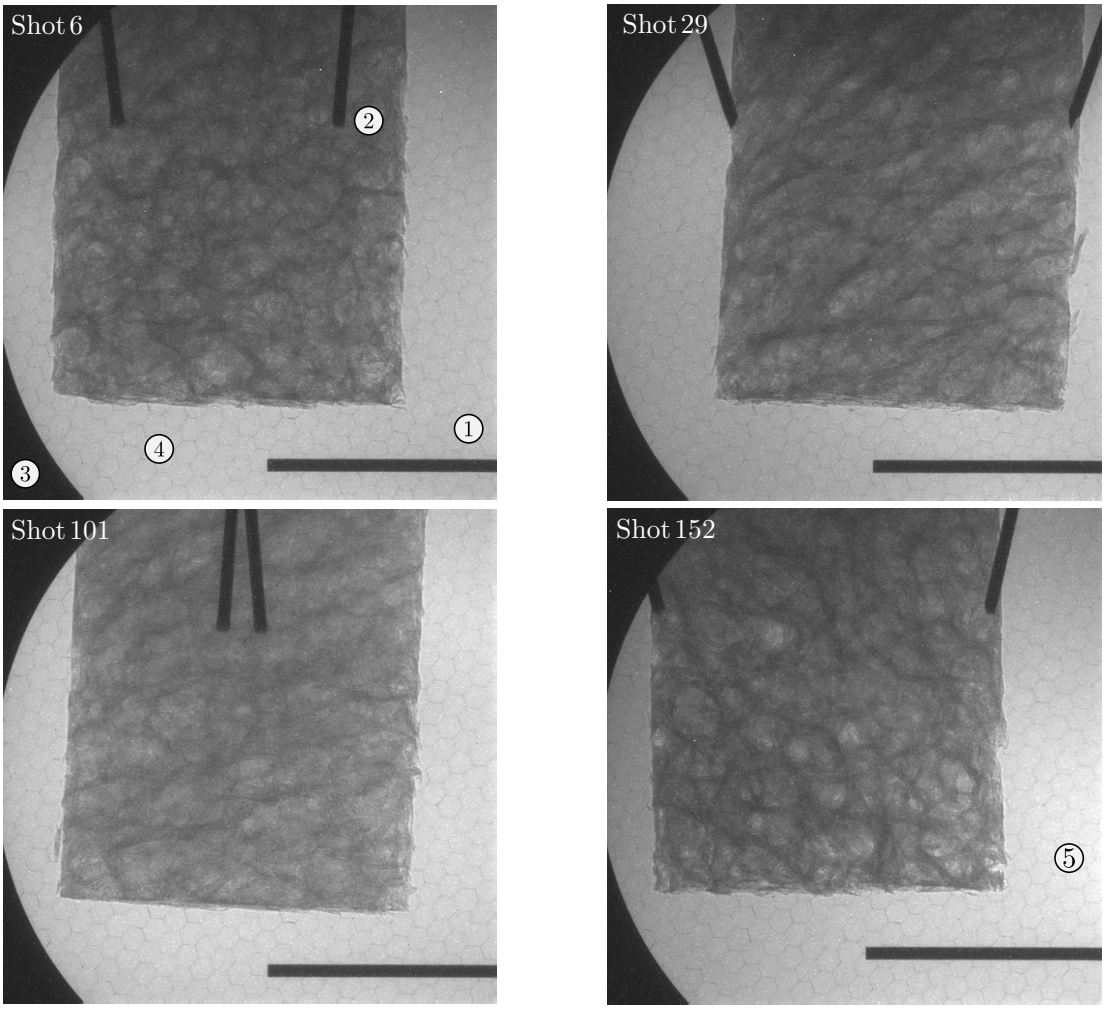


Figure 6.7: Raw x-ray camera images from the tomography run.

The use or elimination of these features during the reconstruction process is discussed in the following sections. The FWHM of the line spread function of the betatron image is a measure of the spatial resolution, and can be measured from the hard edge of the fixed fiducial wire. For the imaging run this was measured to be approximately 8 pixels, or

$36 \pm 7\,\mu m$ at the object plane. This is close to the reported spatial resolution of the x-ray camera scintillator, implying that it is the limiting factor on resolution here.

Image translation

A change in the pointing direction of the x-ray beam does not change the position of the image, rather the position of the back-illumination. However translations of the x-ray source δ cause image translations $\Delta = (R_2/R_1)\delta$. These translations must be removed before reconstruction, but the translation distances are useful in that they provide quite sensitive information about the position of the source. In Fig. 6.8 the analysis procedure is illustrated – the location of the tip of the tungsten pin is deduced from smoothed image lineouts, and indicated with a white cross. After performing this procedure for the 234 shots of the tomography run, the motion of the source may be visualised as in Fig. 6.9. There is a slight complication to make note of here, in that mid-way through the tomographic run a period of beam realignment took place. This is visible in Fig. 6.9 as a gap of just over an hour between shots 119 and 120.

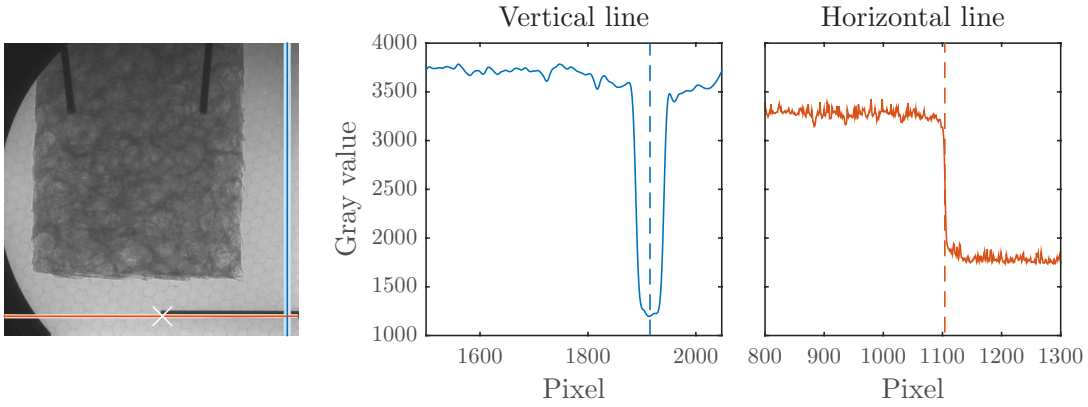


Figure 6.8: Image lineouts used in the automatic location of the wire tip.

In both directions the source is observed to undergo a systematic drift over time by $22\,\mu m$ in x and $40\,\mu m$ in y over 4.2 hours, where x and y are the horizontal and vertical directions respectively. A drift is also observed in the pointing direction of the laser pulse as it entered the experimental area, so it is likely that the focal point of the laser, and therefore the location of the plasma channel, was drifting over time and causing the slow translation of the source. After the gap in shooting this drift is observed to slow slightly in the x-direction, so in order to subtract the drift motion from the random fluctuations a quadratic fit is performed to the whole dataset, with short linear fits on either side of the gap. Having subtracted the drift, the resulting fluctuations about the average motion are binned and plotted as histograms below. In x the source fluctuates by $\pm 2.1\,\mu m$ about the drift motion, and by $\pm 4.9\,\mu m$ in y. This is the laser polarisation direction, so the fact that the source motion is greater in this direction is perhaps an indication that significant levels of radiation are emitted while the electron bunch is under the influence of the laser fields, or that the electron injection event inherits some asymmetry from the laser polarisation.

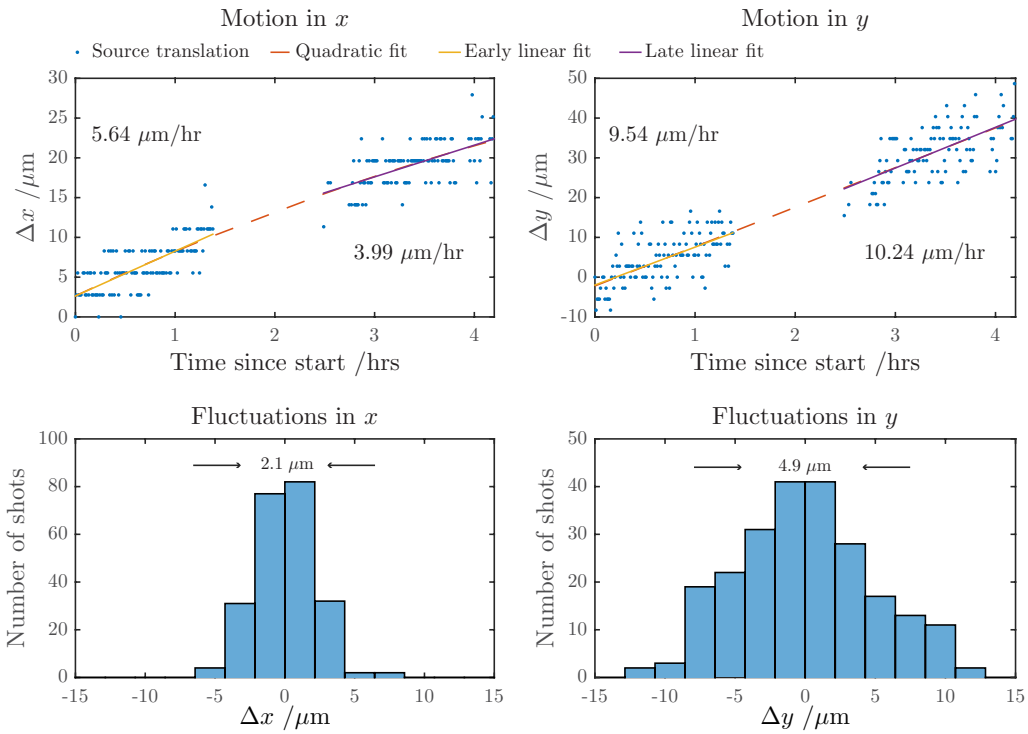


Figure 6.9: Drift motion and random fluctuations of source position.

Rotation axis location

As the rotation stage used in the experiment is not mechanically fixed relative to the camera, it is important to account for the possibility that the sample rotation axis as viewed on the camera is not centred or aligned vertically. To test this, pairs of images acquired 180° apart were compared, i.e. projections at (0° , 180°), (12° , 192°) etc. If the rotation axis were perfectly centred and aligned vertically, then mirroring the second image of each pair and overlaying with the first would produce a perfect correspondence.

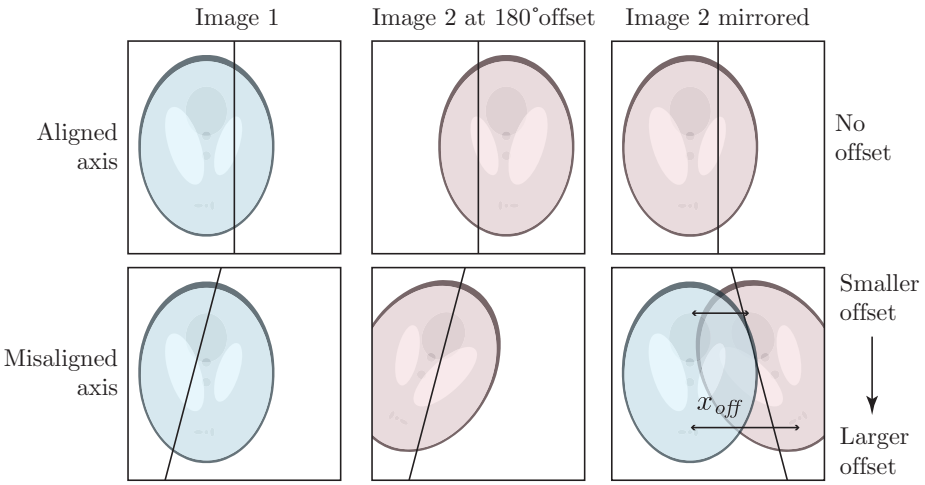


Figure 6.10: Rotating the sample and mirroring the resulting image produces measurable discrepancies if the rotation axis is misaligned.

If not, as illustrated schematically in Fig. 6.10, upon rotation and mirroring the two images

will be offset horizontally. If the rotation axis has some tilt, then the apparent offset also changes as a function of the vertical image co-ordinate. Explicitly, if a feature appears initially at (x_0, y_0) in an image of size $(N \times N)$ pixels, and the rotation axis passes through the point $(x_a, N/2)$ at angle θ to the vertical, then the apparent offset in x after the above procedure is

$$x_{\text{off}} = N - 2x_0 + 2\cos^2\theta(x_0 - x_a) - \sin 2\theta(y_0 - N/2) \tag{6.1}$$

which allows the angle of the rotation axis to be recovered for small θ as

$$\theta \approx -\frac{1}{2}\frac{\mathrm{d}x_{\text{off}}}{\mathrm{d}y_0}. \tag{6.2}$$

For the bone sample, the change in x_{off} was approximately +6 pixels for a change in y_0 of +1,100 pixels, implying $\theta \approx -0.16°$. As a secondary check, the observed offset was calculated after the images were rotated a small amount. As plotted in Fig. 6.11 the offsets recorded at different parts of the image agree if the images are rotated by $0.14 \pm 0.4°$, which is consistent with the estimate of Eq. 6.1. A slight complication here is that the features being compared between image pairs are the edges of the bone sample. As the sample is slightly elliptical in cross-section the edge appears to move as the bone is rotated, necessitating an averaging over bone rotation angles to produce the solid lines in Fig. 6.11. Although this is a small correction, and implies the initial alignment of the rotation stage was very good, uncorrected it would nevertheless lead to a rotation axis misalignment of up to 10 pixels at the edges of the reconstruction and produce obvious artefacts.

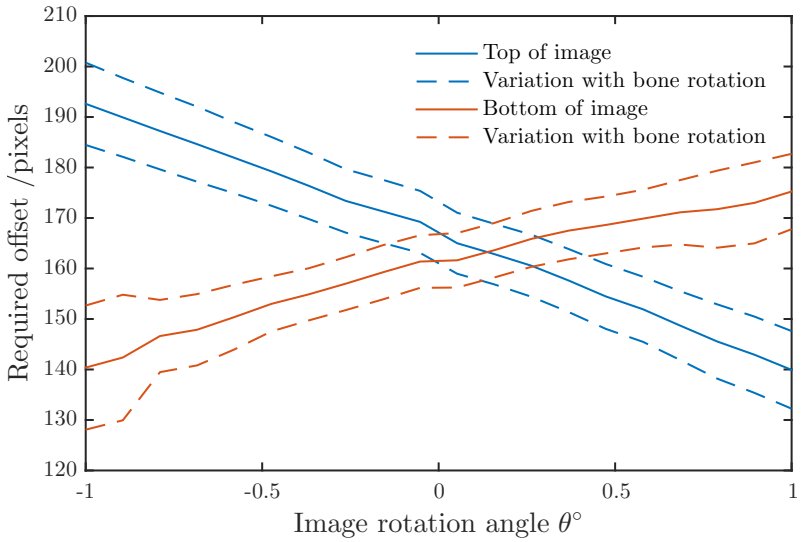


Figure 6.11: The offset x_{off} is constant across the image if the images are rotated by $0.14 \pm 0.4°$.

Background removal

After translation the spatial variation of the beam profile must be removed individually from each image. This is quite different to a conventional µCT scanner where the beam overfills the detector and is therefore much more uniform and consistent between shots.

Nevertheless, from the simulations of shot-to-shot transmission stability in section 3.6.3 it is expected that the variation in beam profile will not have a large effect on reconstruction quality. For each x-ray image, an edge-finding algorithm approximates the projection of the bone sample as a quadrilateral and locates its corners. After dilating the size of the quadrilateral, this shape is now used as a mask so that the parts of the image containing the bone are not used to generate a background – see the first panel in Fig. 6.12. Row-by-row, a smooth 1D spline is fit across the image and then smoothed itself with a 2D spline to produce the beam profile $I_0(x, y)$. The quantity of interest in FBP is the transmission $I(x, y)/I_0(x, y)$, displayed in the third panel.

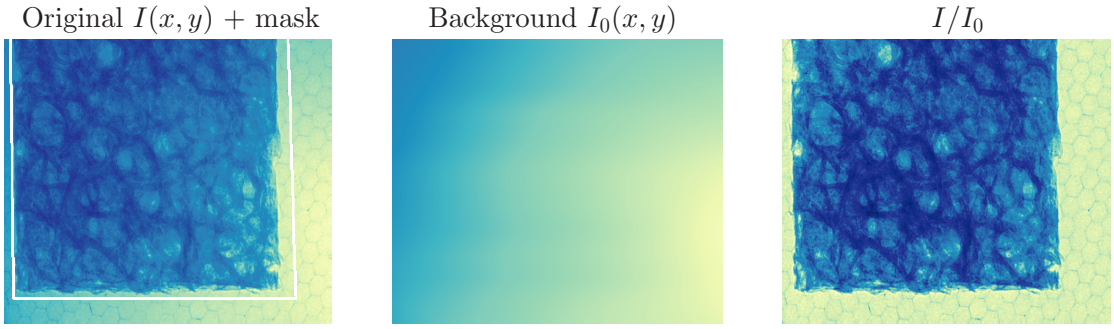


Figure 6.12: The bone sample is masked in the raw image and a smooth background interpolated behind it.

6.3.3 Reconstruction

After the corrections discussed above it is now possible to construct the sinogram for each slice of the sample y as $S_y(r, \theta) = -\ln(I_\theta(r, y)/I_{0,\theta}(r, y))$, where I_θ and $I_{0,\theta}$ denote the projection and background images at angle θ. The raw sinograms display obvious vertical and horizontal features which affect the quality of the reconstruction – see Fig. 6.13.

The vertical features are due to the shot-to-shot fluctuations in sample transmission, and the horizontal features due to features fixed on the detector, in this case dominated by the hexagonal pattern due to the fibre-coupler. By filtering the sinograms as discussed in section 3.6.3 the rings in the reconstruction of each slice are removed reasonably well.

The reconstruction algorithm is the filtered-back-projection approach to the inverse Radon transform[39], here using the `iradon` function of Matlab. Each reconstructed slice of the sample is stacked together into a 3D array, the components of which are known as voxels (volume elements). Orthogonal slices through this dataset are displayed in Fig. 6.14, produced with the ImageJ application.

All three slices demonstrate the ability to penetrate entirely through the sample and observe voids enclosed by bone tissue. The coronal and sagittal slices are observed to possess different structures – the former dominated by prominent north-east to south-west trabeculae, and the latter by finer and more evenly dispersed trabecular sections. The sample is therefore anisotropic with respect to the average trabecular direction, which presumably represents the load-bearing requirements of its particular location in the

[39] [263] A. C. Kak *et al. Principles of Computerized Tomographic Imaging.* (1988).

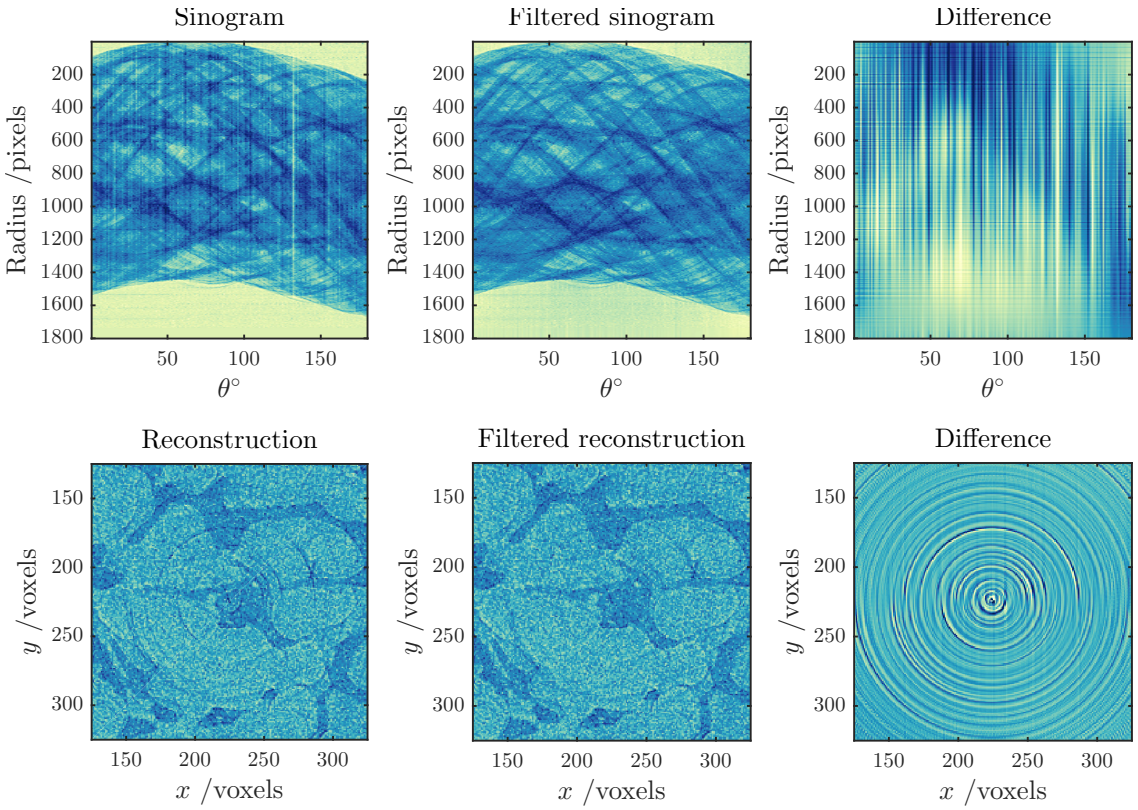


Figure 6.13: Sinogram filtering improves the quality of the bone sample reconstruction. Reconstructions zoomed for clarity.

femoral head. This distinction is reflected in some of the projection images, for example shot 29 of Fig. 6.7.

The bone tissue is observed, at the resolution level of this scan, to have broadly uniform absorption. Typical variation of mineralised bone density is approximately[40] 10%, which is below the noise level of this reconstruction. Quantitatively observing mineral density variations due to physiological processes will require the noise level to be reduced significantly. Beam hardening effects are minimal, with slight enhancement near some outer edges of the sample but otherwise absent. This is likely due to the effective filtering of the x-ray spectrum due to the aluminium foil shields and scintillator response – see section 6.4.5.

At this level of geometric magnification the voxel size is $4.8 \times 4.8 \times 4.8\,\mu m^3$ limited by the pixel size of the detector. The 3D spatial resolution is worse than this, at best reaching approximately 50 µm in some sections of the reconstruction. The Poisson photon noise and relatively small number of projections are the major contributing factors here, which are simply enough improved in the future by the recording of more projection images. Other effects which impact resolution are the parallax errors due to the x-ray beam divergence, diffuse scattering by the sample and blurring due to diffraction. The parallax errors are minimal here, causing transverse averaging across $\approx 70\,\mu m$ regions near the edge of the sample, but may become an issue in the future. So called cone-beam reconstruction methods exist in this case[41], which are well understood due to their commonplace use in

[40] [353] B. Hesse *et al. Journal of Bone and Mineral Research.* (2015).
[41] [354] L. A. Feldkamp *et al. Journal of the Optical Society of America A.* (1984).

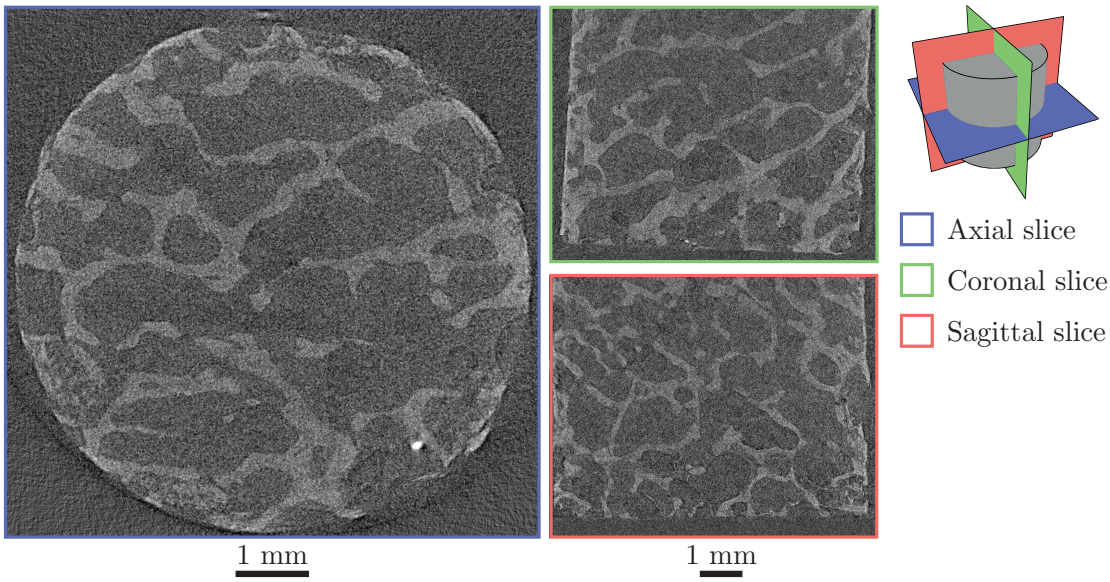


Figure 6.14: Orthogonal slices of the 3D dataset. The slice names are anatomical in origin, where the sagittal plane denotes the symmetry axis of the sample, the axial plane is perpendicular to the tomographic rotation axis, and the coronal plane orthogonal to both.

conventional µCT and can be implemented if necessary. The diffraction of the beam is more of an issue in that it will become more detrimental at higher geometric magnifications when the detector resolution is unimportant. It may be possible in that case to numerically back-propagate the recorded image[42] to recover the sharp absorption image.

Visualisation

It is difficult to convey the full 3D structure of the sample using 2D slices only, so 3D visualisation methods were explored. The first approach was to use isosurface rendering, where a mesh is fitted to the surface of the sample. Defining a threshold grey value as the bone surface fails here due to the relatively noisy background of the reconstruction, so the slices are first thresholded[43] to form an image of bit depth 1 where black is background and white is bone. The dividing surface between 0 and 1 in the resulting 3D array is calculated using a marching cubes algorithm, converted to an `.stl` mesh and rendered using the Blender raytracer[44].

An alternative method is to use volume rendering, where a mesh is not generated and the voxel data are rendered directly. The volume renderer used here is Drishti[45], an open source data visualisation tool designed for 3D datasets. A custom exporter was written to convert from Matlab `.mat` format to Dristi `.pvl.nc` format. Each voxel contributes to the rendered opacity of an image in proportion to a combination of the voxel value and the magnitude of the local intensity gradient. By altering the transfer function parameters the object can be rendered transparent with highlighted edges, or opaque as in conventional

[42] [355] J. Wu *et al. Optik.* (2013).
[43] [356] J. Sauvola *et al. Pattern Recognition.* (2000).
[44] [357] URL: `http://www.blender.org`.
[45] [358] URL: `https://github.com/AjayLimaye/drishti`.

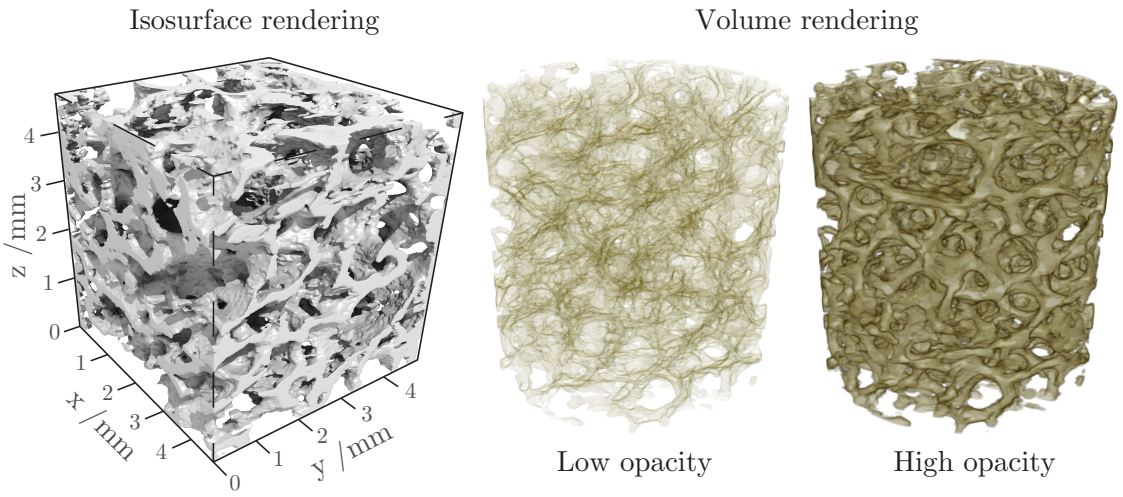


Figure 6.15: Renderings of the reconstruction using three different methods.

rendering. The resulting renders corresponding to these three approaches are displayed in Fig. 6.15.

6.4 Discussion

6.4.1 Flux and brightness

Since the source here is pulsed the exposure time for each projection is effectively zero, and the limiting factor on the total scan time is the repetition rate of the laser. With the Gemini laser, operating presently at a shot every 40 seconds, the full µCT scan shown here could be performed in 2 hours of laser time. Despite this relatively low repetition rate, the high collimation and photon number of the x-ray beam means that this is comparable to the time required for similar investigations with existing microfocus x-ray sources. This is demonstrated in Fig. 6.16 which shows the average photon flux and brightness of our source as compared to microfocus x-ray sources.

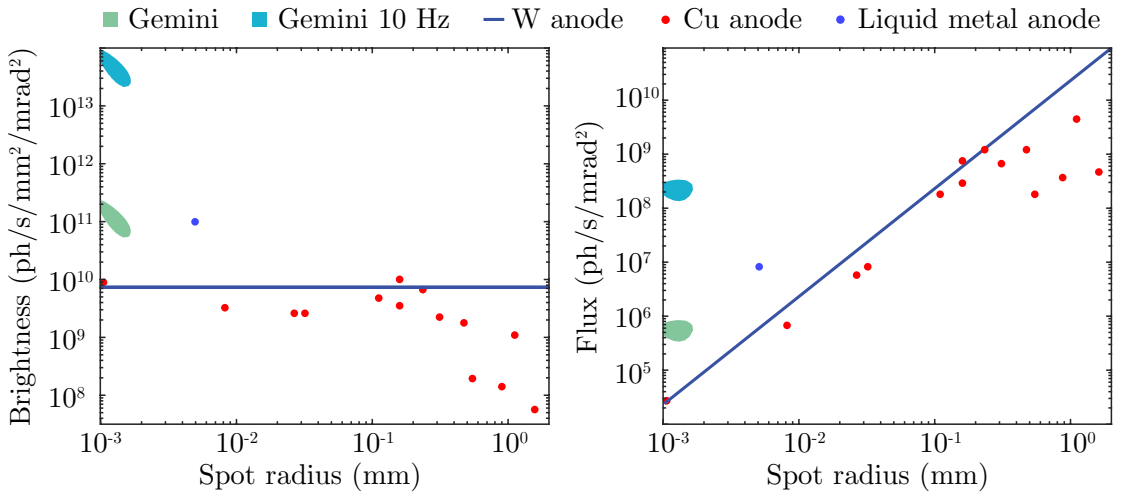


Figure 6.16: Brightness and photon flux as a function of source size for a variety of x-ray sources.

Given the measured fluctuation in photon number and source size, the 1σ brightness

contour is plotted in the brightness-spot size plane, in units of photons per second per square millimetre of source area per square milliradian of beam solid angle. Experimental data from microfocus sources are plotted as points, adapted from Verman *et al.*[46] and the PhD thesis of Tuohimaa[47]. The solid line represents the upper limit of a tungsten anode, using measured anode melting points[48] and simulated spectra[49] using the SpekCalc tool[50]. Since the plasma is almost instantly replenishable, there is no equivalent anode-dependent limitation on source size and x-ray power in the betatron case. The high photon energies required for bone tomography require that microfocus sources use high-Z anode materials, which result in a broad bremsstrahlung-dominated x-ray spectrum. In comparison to conventional sources there is then no disadvantage of the broad betatron spectrum. We integrate over photon energies of 10-100 keV (comprising 64% of the photons above 1 keV at $E_{\mathrm{crit}} = 33$ keV) when comparing the x-ray sources. For the laser-driven source we measure an average brightness of $11^{+12}_{-6} \times 10^{10}$ s^{-1}mm^{-2}mrad^{-2} and an average flux of $5.9 \pm 2.4 \times 10^{5}$ s^{-1}mrad^{-2}. These figures correspond to the average of the peak values within the FWHM area of the x-ray beam and we estimate that the total dose received by the sample over the scanning period was 400 mGy, comparable to conventional imaging modalities[51].

As a comparison to recent conventional sources, Palacio-Mancheno *et al.*[52] provide some details on the scanning of a bone sample with a Bruker SkyScan 1172 microtomograph. Their detector is a 4000×4000 array of 11×11 µm pixels coupled to a 10 W electron gun operating at 100 kV and 100 µA. The x-ray beam is filtered by a 500 µm-thick aluminium filter to mitigate beam hardening effects, and has a broad spectrum peaking near 25 keV. The field of view is 16×16 mm with a similar effective pixel size of 4 µm, and a projection image is recorded in 1.7 s. The rotation step is 0.3° and 5 images are acquired for each projection angle such that a two hour scan contains 3000 images. At a higher geometric magnification of 1 µm per pixel the scan time increases to 5 hours. Though the average flux of the betatron source is high on axis compared to conventional sources, the beam size is much smaller than that of microfocus tubes. Transmission anodes used for high-resolution conventional imaging typically have beam divergences of over 170°, so integrating this flux over a larger solid angle increases the total number of photons received by the sample. Collimated beams are preferable in other situations of course, including those which require significant propagation distances or precision application of radiation.

If the Astra-Gemini laser were to operate at a repetition rate of 10 Hz, pumped by a diode-pumped system as is planned for the future[53], the average betatron flux and brightness would be increased by a factor of 400. As indicated in Fig. 6.16, this source would possess a unique combination of small source size and high average flux. 200-300 TW

[46] [359] B Verman *et al. The Rigaku Journal.* (2002).
[47] [360] T. Tuohimaa PhD Thesis. *Liquid-Jet-Target Microfocus X-Ray Sources: Electron Guns, Optics and Phase-Contrast Imaging.* (2008).
[48] [342] D. E. Grider *et al. Journal of Physics D: Applied Physics.* (1986).
[49] [361] G. Poludniowski *et al. Physics in Medicine and Biology.* (2009).
[50] [362] URL: http://spekcalc.weebly.com/.
[51] [363] W. A. Kalender. *Physics in Medicine and Biology.* (2014).
[52] [364] P. E. Palacio-Mancheno *et al. Journal of Bone and Mineral Research.* (2014).
[53] [365] S. Banerjee *et al. Optics Letters.* (2012).

systems operating at 10 Hz are now commercially available, and with such a system the scan described here could be completed in less than a minute if coupled to a suitably fast detector. Additionally, the resolution of the 3D reconstruction would benefit from the recording of more projection angles in a given time. Furthermore, scientists have already considered the kHz operation of LWFA[54], and consequently of this type of synchrotron source. Beyond considerations of average photon flux the very high *peak* flux of LWFA sources provides the high temporal resolution required of *in vivo* imaging, as the exposure time of each projection is only tens of femtoseconds. Temporal resolution may have benefits in imaging of bone structure, where it may become possible to study microcrack formation in real-time. Of course femotsecond temporal resolution is vastly more than is necessary for the imaging of biological specimens, and recent experiments at synchrotrons have shown that with a sufficiently high-flux source (10^{14} - 10^{15} phs^{-1}mrad^{-2}) *in vivo* time-resolved tomography is possible at kHz image acquisition rates[55].

6.4.2 Signal to noise ratio

One of the main advances made in this experiment when compared to previous work is the successful scaling of the photon number and E_{crit} to a medically relevant level by using a higher-power laser system. The importance of this fact is underlined by the form of the signal-to-noise ratio (SNR) of the tomographically-reconstructed bone density[56]

$$SNR \propto \sqrt{\frac{N\mu^2}{Le^{\mu L}}}, \tag{6.3}$$

where N is the number of photons incident on the sample, μ is the mass absorption coefficient of the sample, here assumed spatially constant, and L is the sample thickness. Maximising the SNR minimises the required dose for a given measurement precision, and therefore is an important quantity to optimise when considering *in vivo* imaging. For a fixed photon flux and sample size (increasing photon flux always increases SNR), maximising SNR by differentiating Eq. 6.3 with respect to μ shows that the optimal absorption coefficient is $\mu_{\text{opt}} = 2/L$. If μ is near μ_{opt} the integrated transmission of the sample changes rapidly as a function of L, and therefore the absorption image is sensitive to changes in sample thickness. Otherwise the image is either too opaque or too transparent and the dynamic range of the detector is under-utilised. Over the medically relevant energy range, the mass absorption coefficient of a sample typically varies as $\mu(E) \propto E^{-3}$, so finding a μ which maximises the SNR is equivalent to using photons of an optimal energy. The transmission curve of the bone sample is plotted in Fig. 6.17, alongside the SNR scaling predicted by Eq. 6.3.

Photons with energies below $\approx 20\,\text{keV}$ are almost completely absorbed, and so do not contribute towards the generation of absorption contrast. Similarly photons above $\approx 80\,\text{keV}$ are almost completely transmitted and contribute a near-constant background. In the intermediate energy range there is a strong peak in SNR near 30 keV, and it is in this energy range that image contrast is maximised and dose levels are minimised. Samples of

[54] [366] S. M. Hooker *et al. Journal of Physics B.* (2014).
[55] [367] S. M. Walker *et al. PLoS Biology.* (2014).
[56] [368] M. J. Flynn *et al. Nuclear Instruments and Methods in Physics Research.* (1994).

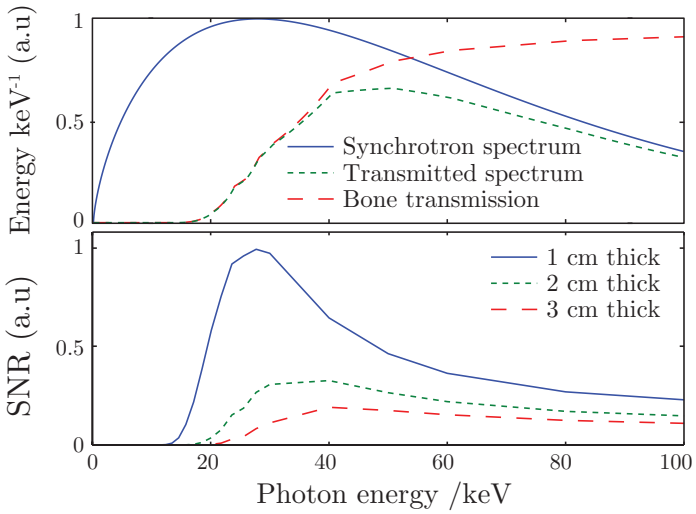


Figure 6.17: The SNR of the absorption contrast image peaks when the absorption of the sample changes rapidly with photon energy.

different thicknesses and compositions have different optimal imaging energies, though the optimum is less pronounced for thicker samples.

The importance of being able to optimise E_{crit} is demonstrated in Fig. 6.18a). Using the reconstructed structure and x-ray absorption characteristics of human bone, absorption contrast images were simulated for incident synchrotron spectra of varying E_{crit} assuming the detection efficiency of the detector used in the experiment. For $E_{\text{crit}} = 10\,\text{keV}$, the sample would be almost completely opaque. For the most energetic beams generated in the present experiment, where $E_{\text{crit}} = 55\,\text{keV}$, the sample becomes nearly 50% transparent at its thickest point and again image contrast is reduced. The images are optimised for intermediary E_{crit}, in the range measured for the experimentally determined optimum contrast: $E_{\text{crit}} = 33 \pm 12\,\text{keV}$. Plotted beneath in Fig. 6.18b) are the histograms of the grey levels in each image. The dotted line marks the mid-point of the measurable range which is where the majority of the image should lie for maximum sensitivity.

6.4.3 Comparison to conventional source

The same bone sample was scanned on a conventional microfocus tube based scanner after the betatron experiment. The scanner was a Bruker SkyScan 1172, often used for the high resolution imaging of small clinical samples, and the scan took place at the Harwell campus of the Medical Research Council. The detector is a Hamamatsu C9300 camera with $9 \times 9\,\mu\text{m}$ pixels, and at a geometric magnification of 3.8 and 2×2 binning of the CCD the effective pixel size at the sample is 4.75 µm – very similar to the current experiment. Projections were spaced by 0.2° over a full 360° and the exposure time per projection was 450 ms. Each projection corresponded to a sum of four images, over which the sample was translated slightly, for a total of 7200 acquisitions over a scanning period of 1 hour. The electron gun power was 10 W, operated at 100 keV and 100 µA with filtering by 500 µm of aluminium.

Comparable projection images are displayed alongside one another in Fig. 6.19. The

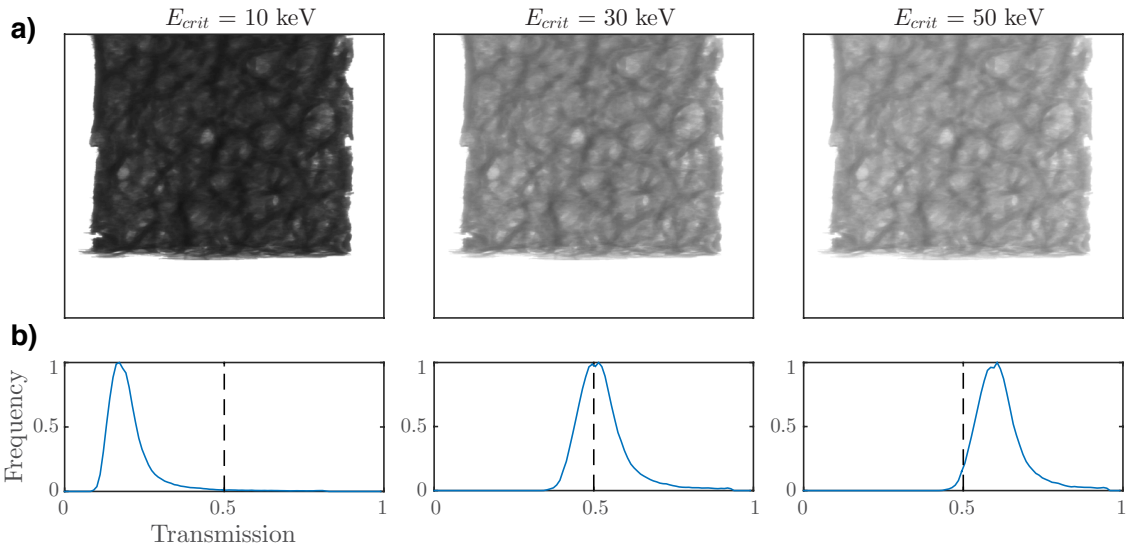


Figure 6.18: a) Synthetic absorption images assuming incident synchrotron spectra of varying E_{crit}. **b)** Histograms of image grey values demonstrating that E_{crit} near 30 keV maximises available dynamic range on the detector.

betatron image appears to be of similar quality, albeit with the addition of an imperfect background subtraction and some evidence of diffraction-enhancement around hard edges. The noise level of the betatron image appears slightly higher, though this may be an artefact of the various correction procedures applied to it.

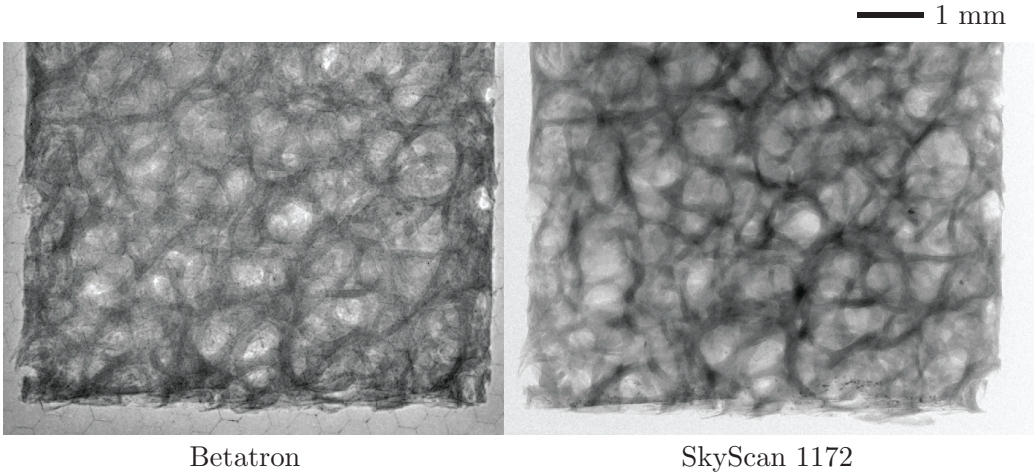


Figure 6.19: Comparison of projections from commercial microtomograph and the current experiment.

A reconstruction of the SkyScan data was performed using the proprietary GPU-accelerated software NRecon developed by Bruker[57]. Similar sections from the two reconstructions are plotted in Fig. 6.20, alongside the lineouts at the indicated positions.
In this case the comparison is less flattering – the betatron reconstruction suffers from significantly higher levels of noise with an SNR of 2 compared to an SNR of 11 from the commercial microtomograph. Given that the commercial scan required the acquisition of

[57] [369] URL: http://bruker-microct.com/products/downloads.htm.

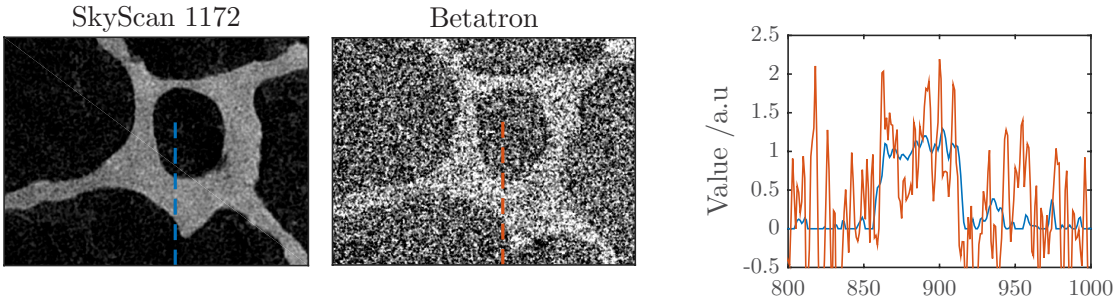


Figure 6.20: Comparison of reconstructions from commercial microtomograph and the current experiment.

3600 images over 180° compared to the 180 images of the betatron scan, one would expect from Eq. 6.3 that the SNR should be $\sqrt{3600/180} = 4.5$ times higher. This is in rough agreement with the measured values from the two reconstructions, and is encouraging in that implies it should be possible to increase the fidelity of the betatron scan by simply acquiring more projection images. The scan time would be significantly increased in that case, and again points to the need for high repetition rate high power lasers if this approach is to become competitive with current technology.

6.4.4 Resolution limit

In this experiment the source size was smaller than a voxel and so does not contribute to the measured resolution, which is instead dominated by the Poisson photon noise and scintillator PSF. Increasing the resolution to the CCD limit would be possible by acquiring more projection images. The ultimate resolution limit depends on the imaging geometry and detector – consider the layout sketched in Fig. 6.21 where a source of size s is backlighting an object R_1 away, forming an image on a detector at a further distance R_2.

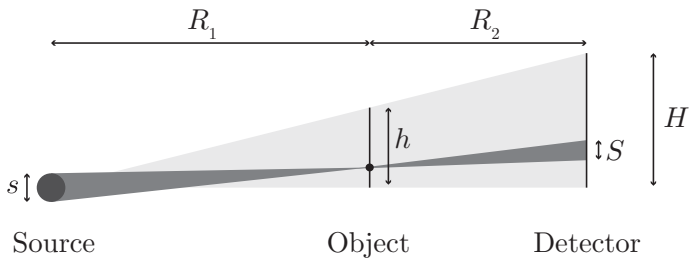


Figure 6.21: Layout of a generic back-illumination system. The imaging distances R_1 and R_2 in lensless x-ray imaging determine the geometric magnification and source size-limited resolution.

By considering the similar triangles between the source and the object and the source and the detector, the geometric magnification $M = H/h = (R_1 + R_2)/R_1$. Considering the darker similar triangles either side of the object, a point in the object is projected over a transverse distance $S = sR_2/R_1$ on the detector – the source-size limited resolution. Taking into account the geometric magnification, this limit at the object plane becomes

$$\text{Source limit} = \frac{S}{M} = \frac{sR_2}{R_1 + R_2} = s\frac{R_2}{R_1}\frac{1}{M}. \tag{6.4}$$

For small R_2 or very high magnification, the source size plays little role in the resolution of the projection image. The second limit is the resolution of the detector – if the detector has pixel size p, the effective minimum resolution determined by the detector is

$$\text{Detector limit} = \frac{p}{M} = \frac{pR_1}{R_1 + R_2}. \tag{6.5}$$

To determine the maximum resolution given s and p, fix the total beamline length $L = R_1 + R_2$ and assume the overall resolution r is governed by some generalised ℓ^n norm

$$r = \left(\left(\frac{S}{M}\right)^n + \left(\frac{p}{M}\right)^n\right)^{1/n}. \tag{6.6}$$

With L fixed, minimising r with respect to R_1 shows that there is an optimal distance between the source and object which generates the minimum resolution r

$$\frac{1}{r} = \left(\left(\frac{1}{s}\right)^{\frac{n}{n-1}} + \left(\frac{1}{p}\right)^{\frac{n}{n-1}}\right)^{\frac{n-1}{n}}. \tag{6.7}$$

and at this point

$$\frac{R_1}{R_2} = \left(\frac{s}{p}\right)^{\frac{n}{n-1}}. \tag{6.8}$$

The form of Eqs. 6.7 and 6.8 shows that if s and p are significantly different then the smaller of the two effectively determines the resolution limit. For example if the source is large then $s \gg p$ and $r \to p$ as long as $R_1 \gg R_2$. This is the contact geometry often used on synchrotron beamlines, where the source is large (hundreds of microns) and the object is placed directly against the scintillator. High resolution in that case comes from the minimisation of the effective detector pixel size, achieved in practice through optical magnification of a thin high resolution scintillator. Detector efficiency drops with scintillator thickness, so it is important that the source also has high brightness if high resolution is to be achieved. There is then a lower bound on R_2 and the total beamline length $L \approx R_2(s/p)^{n/(n-1)}$ scales with the source size.

The opposite situation applies to the present experiment, where $s \ll p$ and $r \to s$ – the 'microscope' configuration where high resolution requires high geometric magnification. In this case $R_1 \ll R_2$ and magnification is large but the requirements on the detector are less strict. This is generally the configuration encountered when using a commercial microtomograph, and limits the field of view of the detector. The problem for betatron imaging is that the high energy electron beam needs to be removed from the laser axis before encountering the sample. For GeV electron beams and macroscopic ($\approx$ cm) deflections, the required length for a 1 T permanent magnet is several tens of centimetres (recall Fig. 3.31). This places a strict lower bound on R_1 and then the total beamline length scales as $L \approx R_1(p/s)^{n/(n-1)}$, i.e. for smaller source sizes longer beamlines are required

to maximise the available resolution. As the photon flux drops like L^{-2}, at some point the Poisson photon noise will dominate the spatial resolution rendering any further gains moot.

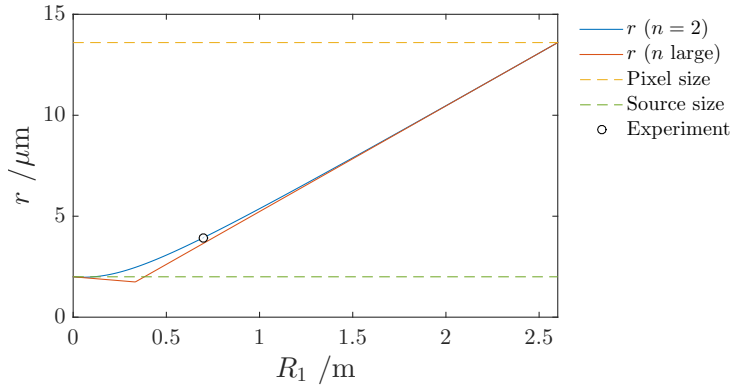


Figure 6.22: Spatial resolution for the current experiment as a function of object position. $s = 2\,\mu\text{m}$, $p = 13.6\,\mu\text{m}$.

To illustrate these remarks we plot Eq. 6.7 as a function of object position for various fixed-length beamlines. In Fig. 6.22 the parameters corresponding to the current experiment are used, with the actual object position indicated by the white marker. In this 'microscope' mode the resolution almost linearly worsens as the object as moved farther from the source. In this case the region $R_1 \lesssim 0.5\,\text{m}$ was inaccessible due to the magnet, so the imaging in the present experiment is not too far from the highest resolution which could have been achieved.

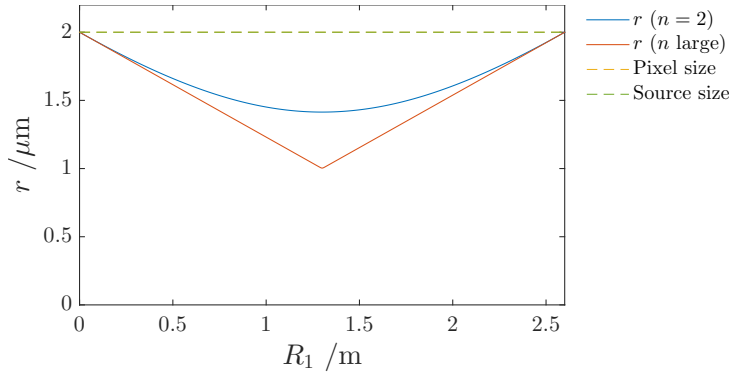


Figure 6.23: Spatial resolution for the current experiment as a function of object position if the detector pixels were smaller. $s = 2\,\mu\text{m}$, $p = 2\,\mu\text{m}$.

If the current experiment had employed a detector with smaller pixels close to the source size, the resolution profile would look instead like Fig. 6.23. When the source and detector contribute equally to the resolution, the geometry is symmetric and the optimal configuration is to place the sample halfway to the detector at a magnification of 2.

Finally, if the current experiment had been performed with a beamline length of 10 metres the resulting resolution profile is plotted in Fig. 6.24. With the sample the same distance from the source as before the spatial resolution is more than twice as high, but the magnification is much larger. Depending on exactly how the two resolution constraints

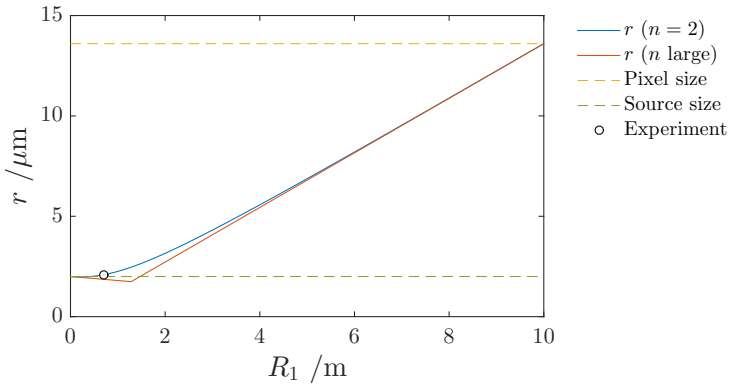


Figure 6.24: Spatial resolution for the current experiment as a function of object position for a longer beamline. $s = 2\,\mu\text{m}$, $p = 2\,\mu\text{m}$.

contribute to the overall resolution, it might even be preferable to have the object slightly farther away from the source. In this case the photon flux would be reduced by a factor of ≈ 15 compared to the current experiment.

6.4.5 Beam hardening correction

If the sample is assumed to be composed of a single material, it should be possible to correct for the broadband nature of the x-ray spectrum. Recalling Eq. 3.84, if the x-ray spectrum is monochromatic and the sample is composed of a material of attenuation coefficient μ_0, then the projected thickness of the sample at a pixel (x, y) in the absorption image I is

$$L(x, y) = -\frac{1}{\mu_0} \ln \frac{I(x, y)}{I_0} \tag{6.9}$$

where I/I_0 is the relative transmission of the sample at (x, y). If the spectrum is actually polychromatic, then the logarithm of the absorption image fails to recover the actual projected thickness of the sample and beam hardening artefacts appear in the reconstruction. If the spectrum is known however, the mapping between L and I can be calculated and the beam hardening effects eliminated.

In Fig. 6.25 this mapping is plotted assuming the spectral absorption characteristics of human bone, the detection efficiency of the detector used in the experiment and an incident synchrotron spectrum of $E_{\text{crit}} = 33\,\text{keV}$. Also plotted is the length-transmission curve for a monochromatic spectrum incident on the same sample and detector, which as expected appears linear on a logarithmic x-axis. The broadband transmission curve is amenable to a simple polynomial fit which can be applied to a sinogram in a computationally efficient manner.

To assess the effects of this transmission curve correction, here a slice of the reconstructed sample is used as a test object. Synthetic projection images are calculated, assuming the slice is composed of uniform human bone, and then reconstructed a second time with or without the beam hardening correction. These reconstructions are plotted in Fig. 6.26 above histograms of the reconstructed values.

The naive reconstruction, where the logarithmic mapping is used, displays prominent

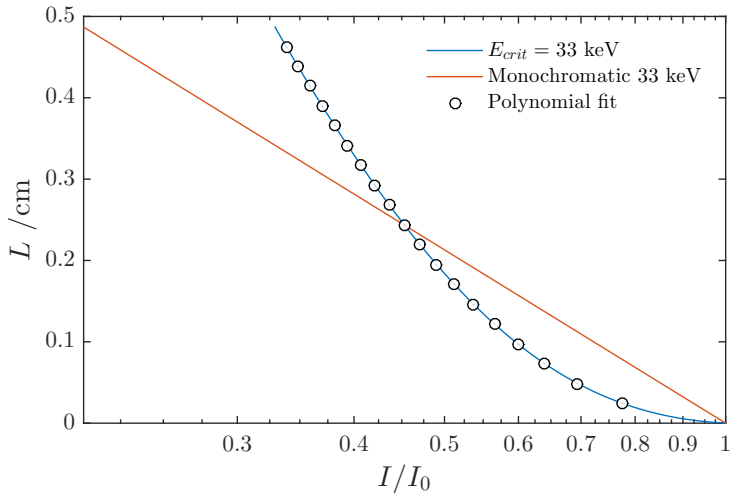


Figure 6.25: Sample length-transmission curves are plotted for monochromatic and broadband spectra. By fitting a polynomial to the broadband curve, it is possible to reduce the impact of beam hardening by transforming the projection images in an inverse fashion.

beam hardening artefacts around the edges of the sample. Also present are characteristic streaks through the centre of the reconstruction volume. The object and background peaks in the histogram are observed to be broadened which limits the accuracy of any quantitative measurement of bone density. With the correction applied to the sinogram the reconstruction is much more uniform and the background streaks have been eliminated. The broadening of the peaks in the histogram in this case is predominantly due to the finite number of projection angles used.

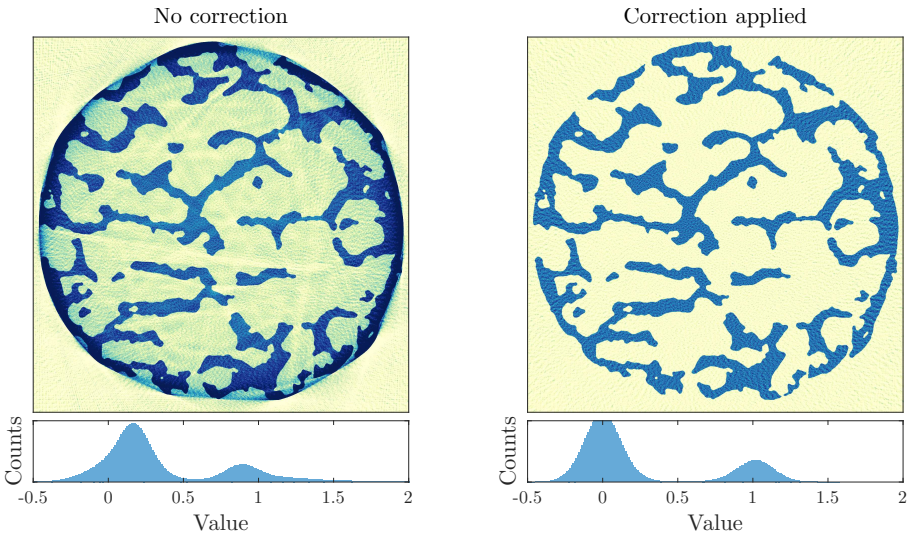


Figure 6.26: Applying a fixed beam hardening correction flattens the profile of the reconstruction, which is also observed in the histogram plotted below each image. A narrower peak represents a more accurate measurement of average μ, or equivalently average bone density.

6.4.6 Towards clinical applications

The results presented in this chapter represent the first application of a laser-wakefield driven x-ray source towards the tomographic imaging of a hard medical sample, exciting in its own right. However, before this technique can be used in a diagnostic or investigatory capacity there are a number of challenges to be overcome. The emerging consensus on the factors constituting bone quality is that it is important to diagnose how the distribution of mineralisation in the bone affects the likelihood of generating microcracks and full fractures. The spatial resolution reported here is not high enough to observe microcracks directly, and the joint spatial and compositional resolution is not high enough to map the spatial distribution of low-amplitude density variations throughout the bone, typically 1–10%[58]. This final section briefly discusses some of the potential steps towards clinical tomographical application of a betatron source.

Quantitatively, synchrotron µCT has demonstrated the possibility of observing microcracks with high fidelity at a voxel size of[59] 1.4 µm. Mapping the size and distribution of microcracks does not rely on precision density measurements, and so the spectral content of the x-ray beam is not very important compared to the resolution of the projection images. The laser-wakefield source enjoys a naturally small source size, so the route to high resolution is

- the acquisition of more projection images – this approach scales until the resolution limit discussed in section 3.6.3, and benefits from the use of a high-repetition-rate laser
- greater geometric magnification – the use of smaller samples placed closer to the source, or a detector placed farther away, is a simple way to improve the geometric magnification. If the beam propagation distance becomes too large then the image will begin to diffract, which may necessitate a further processing step to back-propagate the image before reconstruction. If the image over-fills the detector then more sophisticated 'interior tomography' reconstruction algorithms need to be employed[60].
- higher-resolution detectors – absorptive x-ray scintillators typically have spatial resolution similar to their thickness, limited in practice to $\approx 5\,\mu m$, due to the scattering of visible photons as they traverse the substrate. If a transparent scintillator such as YAG is used, resolution can be increased by employing an optical imaging setup with short depth-of-field to image the front surface of the scintillator. X-ray detection efficiency is decreased in this case, so the average x-ray flux should be increased to compensate.

Better compositional resolution comes from better knowledge and control of the x-ray spectrum. If the spectrum is known and the energy dependence of the sample transmission

[58] [370] P. Roschger *et al. Current Osteoporosis Reports.* (2014); [353] B. Hesse *et al. Journal of Bone and Mineral Research.* (2015).
[59] [333] A. Larrue *et al. PloS one.* (2011).
[60] [371] L. Ritschl *et al. IEEE Nuclear Science Symposium Conference Record.* (2010).

is known, from the procedure outlined in the previous section it is always possible to recover $\int \mu \, \mathrm{d}l$ from a projection image. Furthermore, if the sample is known to contain a single material of particular constant attenuation coefficient per unit density (μ/ρ), then one also knows the projected density as

$$\int \mu \, \mathrm{d}l = \left(\frac{\mu}{\rho}\right) \int \rho \, \mathrm{d}l, \tag{6.10}$$

and so after tomographic reconstruction the material density is retrieved. In the laser-wakefield case the spectrum can be measured on shot by, for example, the placement of small filters around the imaging region (though the filters should not be placed too far from the beam centre as the photon spectrum may vary across the beam[61]). The spectrum can be narrowed by absorptive filtering at the low energy end, and employing the decreased detector sensitivity at the high energy end.

If the sample contains multiple materials of differing (μ/ρ) then it is not possible in general to convert the absorption of the sample into a mass density. In this case the solution is to use a multi-spectral method where multiple absorption images are recorded for different incident x-ray spectra. If a sample is illuminated by different spectra i and contains materials j with density distributions ρ_j, the transmission of a particular spectrum T_i is

$$-\ln T_i = \int \mu \, \mathrm{d}l = \sum_j \left(\frac{\mu}{\rho}\right)_{ij} \int \rho_j \, \mathrm{d}l \tag{6.11}$$

and so recovery of the projected densities of the different components involves the inversion of a matrix of the (known) attenuation coefficients $(\mu/\rho)_{ij}$. Clearly the distinguishing of n components requires the use of at least n different x-ray spectra. For $n = 2$ this technique is known as dual-energy x-ray absorptiometry (DEXA) where the bone is assumed to consist of a soft and hard component, widely used in the diagnosis of osteoporosis. Here the x-ray spectra are altered by differentially filtering the x-ray beam or applying different accelerating voltages in the x-ray tube. For the laser-wakefield source the spectrum may also be changed through filtering, or if sufficient control is established by changing the laser power or plasma density.

Finally, if this technique were to become commonplace in hospitals it would need to be performed *in vivo*. The high collimation of the source is then a benefit, as the radiation dose is localised to a smaller volume. The x-ray energy needs to be higher to penetrate centimetres of soft tissue ($> 100\,\mathrm{keV}$), and high-resolution detection at this energy is challenging. In addition there are significant technological steps to be taken before a high power ($> 100\,\mathrm{TW}$) high repetition-rate ($> 10\,\mathrm{Hz}$) laser system could be installed and run in a 'turn key' fashion by non-experts at a hospital. Nevertheless, the fact that this plasma-based approach sidesteps the material limitations facing current laboratory-based sources is a key advantage, rendering it amenable to rapid progression in brightness and stability.

[61] [90] F. Albert *et al. Physical Review Letters.* (2013).

6.5 Conclusion

In summary, we have demonstrated that it is possible to use a LWFA to drive an x-ray source at the high photon energies required for the imaging of medically-relevant bone samples. The source size remains small at these high energies, and is sufficiently stable that we were able to perform micro-computed tomography in a reasonable time. We measured an average photon flux which is already comparable to commercial microfocus sources, but note that the use of a high-repetition-rate laser will enable a dramatic increase in flux without compromising source size. Coupled to the proof-of-principle demonstrated here, the prospect of bringing laser-based µCT sources to the laboratory seems bright.

7 Conclusion

In laser wakefield acceleration the evolution of the laser pulse as it enters the plasma is of utmost importance, determining how soon the electrons are injected, the charge of the injected bunch, and the eventual energy gain. Optimisation of the laser-plasma coupling allows the beam energy and charge to be maximised, which causes the subsequent betatron radiation to become brighter and more energetic. Diagnosis of this evolution has formed a central part of this thesis, through direct optical diagnostics and inferred from the details of the unstable laser-plasma interaction.

A study of the interplay between the transverse probe and wakefield driver delivered insights into the ultrahigh intensity regime of Raman amplification, and by finding a stable regime of electron acceleration and x-ray generation the first tomographic scan of a clinical tissue sample was performed with a betatron source. These achievements demonstrate the rapid evolution of laser wakefield accelerators as a source of high-energy electron and photon beams, and the wealth of interesting physics yet to be explored surrounding the high-intensity laser-plasma interaction.

7.1 Raman scattering as a diagnostic

The spectral properties of Raman-scattered light were investigated in order to evaluate its use as a simple non-invasive diagnostic of plasma density. The spectra were often modulated, of widely varying bandwidth, and systematically shifted in frequency compared to that expected. PIC simulations demonstrated the complex interplay between the Raman scattering and self-focussing processes, which while complicating the interpretation of Raman spectra as a density diagnostic, indicate its use in providing information about the laser-plasma interaction. The frequency shift indicates that the scatter originates from the centre of the laser pulse and can freely propagate out of the wake. The laser pulse is shaped as it scatters, and this pulse evolution in turn alters the character of the scattered light, changing its bandwidth and central frequency. Crucially late in the interaction the detected light is observed to originate from the region outside of the plasma wave, and ceases when the pulse becomes matched. That the scattered light increases in brightness so rapidly with density suggests the pulse must undergo a more extensive period of self-focussing and pulse compression at high densities, and therefore lower densities are preferable for coupling laser energy into the wake.

At higher laser intensities and longer plasma lengths the signature of Raman scatter was the formation of extended filaments leaving the plasma channel at a small angle. The fact that little relativistic correction was required to fit the data indicates the filaments

are generated by the outer low intensity regions of the pulse not captured by the wake. The connection to the laser coupling efficiency was made by intentionally detuning the accelerator – moving the focal plane position and aberrating the focal spot. The presence of filaments was correlated with a decrease in accelerator performance, reducing the electron beam charge and stability and correlated with increased scattering of the laser pulse energy. It is hypothesised that a perfectly-matched laser focal spot and pulse length will minimise these effects, scattering little energy and producing no filaments such that the maximum fraction of energy is trapped in the bubble. If this is the case the oscillations in the bubble size are minimised and injection can occur as soon as possible. Spectral and optical measurements of Raman side-scatter are therefore proposed as a useful on-line diagnostic of accelerator performance, which importantly for a high-repetition-rate system would represent data which is quick to acquire and simple to analyse.

7.2 Pulse amplification in a laser wakefield accelerator

The fact that crossing femtosecond laser pulses interact so strongly in a laser wakefield accelerator was surprising, especially considering the billion-fold disparity in intensity between the pulses. Initially suspected to be a stimulation of the Raman scattering process, the lack of a frequency shift in the observed light sparked a more thorough investigation into the physical process underpinning the effect. Rotating the probe beam polarisation destroyed the amplification, which must therefore on interference between the two pulses. By changing the time delay between the pulses it was also observed that the amplification was maximised near the beginning of the plasma. Building upon the conclusions of the previous investigation into Raman scattering, this was consistent with the requirement for laser light to exist outside of the bubble. Further simulations confirmed this to be the case, showing that the strong density gradients of the bubble sheath prevented effective propagation of the amplified light. One use of the mechanism is then as a diagnostic, probing the level of light untrapped in the wake and the position of the drive pulse.
Interpreted in a Raman amplification context the spectral gain region was found to become very broad at high drive intensities, caused by the transient nature of the interaction and the anharmonic motion of the plasma electrons. This explained why the observed light was little shifted from its original frequency and why the observed growth rate was so large. The amplification factor of the probe was inferred to be extremely large and the growth length extremely short, which even interpreted conservatively represents the most rapid example of Raman amplification observed to date.
Many questions remain to be answered here however, and further experiments would help to clarify the physical picture and determine whether or not this mechanism may become useful in practice in the future. Specifically,

- How does the amplification efficiency change as a function of the pulse intensities and polarisations?
- Does the amplification increase with larger or aberrated focal spots?
- Can the gain bandwidth be probed by changing the wavelength and/or bandwidth of the probe?
- Can the interaction introduce asymmetry into the injected electron beams?

As a largely non-invasive measurement, it is probable that some of these questions can be answered in the future in a parasitic fashion during other experiments. Exploring these issues more deeply will require dedicated experimental time and further theoretical effort.

7.3 Clinical applications of betatron radiation

Finally the application of laser wakefield-driven betatron beams to a contemporary clinical imaging problem was successfully attempted, resulting in the high resolution tomographic reconstruction of a human femoral bone sample. The high-fidelity 3D imaging of human bone is a pressing issue as the medical community works to understand the causes and symptoms of osteoporosis, and there is need for high-flux x-ray sources in the 30 keV range with micron source sizes. Synchrotron light sources fill this niche for state-of-the-art imaging studies, but a laboratory scale solution will be required in the future if microtomographic imaging is to become routine in the treatment of osteoporosis. The unique properties of betatron sources are their small source size, short pulse duration, high energy, and simple scaling with laser repetition rate. By positioning proposed high repetition rate betatron sources between synchrotrons and microfocus tubes their advantages are maximised, representing a real chance of becoming a disruptive technology. Finding a stable region of operation of the source was crucial to the successful outcome of the experiment. Under optimal conditions x-ray beams were produced on 97% of laser shots, with critical energy 33 ± 12 keV, pointing jitter 2.1×1.5 mrad, containing $1.3 \pm 0.5 \times 10^9$ photons above 1 keV from a source size of 2–3µm. Projection images were recorded at a resolution of 36 ± 7 µm and after 180 image acquisitions it was possible to reconstruct a tomogram at a 3D resolution of up to 50 µm. At a repetition rate of 0.025 Hz the scan was completed in two hours of laser time, which in the future could be compressed into one hour of Astra-Gemini laser shots at the maximum rate of 0.05 Hz. The average photon flux between 10 and 100 keV was $5.9 \pm 2.4 \times 10^5\,\mathrm{s}^{-1}\mathrm{mrad}^{-2}$ which is comparable to microfocus sources capable of few-micron source sizes.

While an impressive proof-of-principle, it is important to consider how future experiments will improve upon the technique. Increasing the source repetition rate, and therefore average flux, is important in several respects. A greater number of projections can be recorded in a given time, which will deliver substantial improvements in the SNR and spatial resolution of the reconstruction. Multiple images may also be integrated for each projection, allowing the use of thinner scintillators which decrease the influence of the detector PSF on the image resolution and maximise the benefits of the small source size. If the betatron spectrum is measured on every shot the effects of beam hardening may be mitigated, facilitating quantitative measurements of bone density with greater precision. Finally betatron x-ray beams enjoy significant levels of spatial coherence, and phase-contrast tomography of tissue samples is possible. The benefits of phase-contrast imaging occur when the photon energy is high and absorption, and therefore dose, is low. The flux must also be high to compensate for the drop in detector sensitivity at high energies, and therefore this technique will also benefit from the use of a high repetition rate driving laser.

A Appendix

A.1 1D SRS dispersion relation

With the following quantities defined,

$$\Omega = \frac{\omega}{\omega_p}, \quad \Omega_0 = \frac{\omega_0}{\omega_p}, \quad K_\pm = \frac{c}{\omega_p}\sqrt{(\mathbf{k}_0 \pm \mathbf{k})^2}, \quad K_e = \frac{c\mathbf{k}}{\omega_p}, \quad \gamma = 1 + \frac{a_0^2}{2}, \tag{A.1}$$

Eq. 2.123 of the main text takes the explicit form

$$\begin{aligned}
&\Omega^6 \\
&+ \Omega^4\left[\tfrac{a_0^2}{2\gamma^3} - \tfrac{3}{\gamma} - 2\Omega_0^2 - (K_+^2 + K_-^2)\right] \\
&+ \Omega^3\left[2\Omega_0(K_+^2 - K_-^2)\right] \\
&+ \Omega^2\left[\tfrac{2}{\gamma^2} + \tfrac{2\Omega_0^2}{\gamma} + \tfrac{1}{\gamma}(K_+^2 + K_-^2) + \left(\tfrac{1}{\gamma} - \Omega_0^2\right)^2 + \left(\tfrac{1}{\gamma} - \Omega_0^2\right)(K_+^2 + K_-^2)\right. \\
&\qquad \left. + K_+^2K_-^2 - \tfrac{a_0^2}{4\gamma^3}\left(\tfrac{2}{\gamma} - 2\Omega_0^2 + K_+^2 + K_-^2\right) - \tfrac{a_0^2}{2\gamma^3}\left(K_e^2 + \tfrac{1}{\gamma}\right)\right] \\
&- \Omega\left[\tfrac{2\Omega_0}{\gamma}(K_+^2 - K_-^2)\right] \\
&+ \tfrac{a_0^2}{4\gamma^3}\left(K_e^2 + \tfrac{1}{\gamma}\right)\left(\tfrac{2}{\gamma} - 2\Omega_0^2 + K_+^2 + K_-^2\right) - \tfrac{1}{\gamma}\left(\tfrac{1}{\gamma} - \Omega_0^2\right)^2 \\
&\qquad - \tfrac{1}{\gamma}\left(\tfrac{1}{\gamma} - \Omega_0^2\right)(K_+^2 + K_-^2) - \tfrac{K_+^2K_-^2}{\gamma} \\
&= 0.
\end{aligned} \tag{A.2}$$

While cumbersome and opaque, this polynomial form is most suitable for computational root finding.

A.2 Landau damping rates

Electron Landau damping rates can be derived from the plasma dispersion relation

$$\epsilon(\omega, k) = 1 - \frac{1}{2k^2\lambda_D^2} Z'\left(\frac{\omega}{\sqrt{2}\omega_p\lambda_D k}\right) = 0 \tag{A.3}$$

where a maxwellian plasma is assumed and Z' is defined as in appendix A.6. For a given k, a mode satisfying this dispersion relation corresponds to a complex root ω, where the damping rate is $\nu_L = \Im(\omega)$ and the resonant frequency is $\Re(\omega)$. Exact and approximate expressions for these quantities are derived by McKinstrie *et al.*[1] as

[1] [372] C. J. McKinstrie *et al. Physics of Plasmas.* (1999).

$$\begin{aligned}\Re(\omega) &= \omega_p\Big(1+\tfrac{3}{2}(k\lambda_D)^2+\tfrac{15}{8}(k\lambda_D)^4+\tfrac{147}{16}(k\lambda_D)^6\Big)\\ \nu_L &= \omega_p\sqrt{\pi/8}\Big((k\lambda_D)^{-3}-6k\lambda_D\Big)\exp\left(-\tfrac{1}{2}(k\lambda_D)^{-2}-\tfrac{3}{2}-3(k\lambda_D)^2-12(k\lambda_D)^4\right),\end{aligned} \tag{A.4}$$

which are reasonably accurate for $k\lambda_D \lesssim 0.3$. As plotted in Fig. A.1, it is clear that the Landau damping rate is very small for $k\lambda_D < 0.5$.

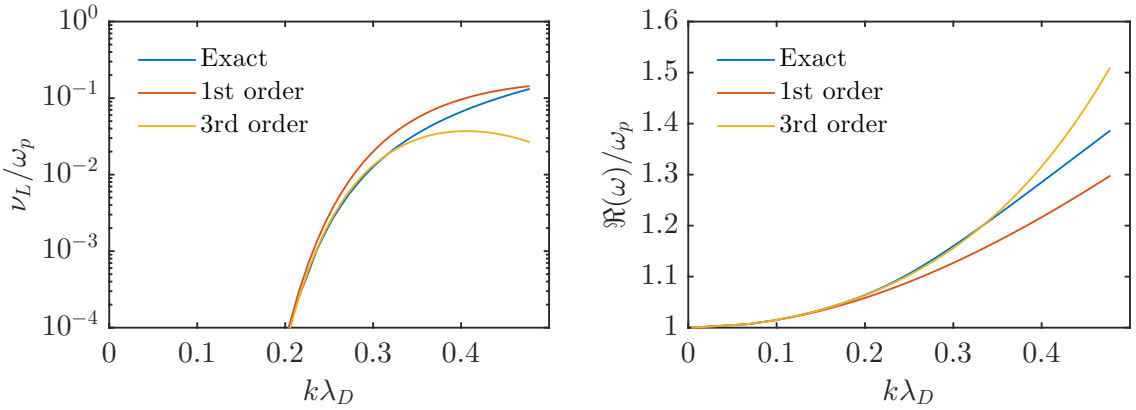


Figure A.1: Landau damping rate and resonant electron oscillation frequency as a function of $k\lambda_D$.

A.3 Bedrosian's theorem

Start with the product of f and g expressed as a product of inverse Fourier transforms

$$f(x)g(x) = \frac{1}{2\pi}\int_0^\infty \mathrm{d}u \int_0^\infty \mathrm{d}v\, F(u)G(v)e^{-i(u+v)x} \tag{A.5}$$

and take the Hilbert transform with respect to x. Remembering that $\mathcal{H}(e^{iux}) = -i\,\mathrm{sgn}(u)e^{iux}$ this becomes

$$\mathcal{H}(f(x)g(x)) = \frac{1}{(2\pi)^2}\int_0^\infty \mathrm{u} \int_0^\infty \mathrm{d}v\, F(u)G(v)i\,\mathrm{sgn}(u+v)e^{-i(u+v)x}. \tag{A.6}$$

The fact that the functions have different support now allows factorisation – see Fig. A.2. The product $F(u)G(v)$ is only nonzero in the boxed regions in the (u,v) plane. In these regions, above and below the $u+v=0$ line, we can make the identification that $\mathrm{sgn}(u+v) = \mathrm{sgn}(v)$, and factorise the above as

$$\begin{aligned}\mathcal{H}(f(x)g(x)) &= \frac{1}{(2\pi)^2}\int_0^\infty F(u)e^{-iux}\,\mathrm{d}u \int_0^\infty G(v)i\,\mathrm{sgn}(v)e^{-ivx}\,\mathrm{d}v\\ &= f(x)\mathcal{H}(g(x)),\end{aligned}$$

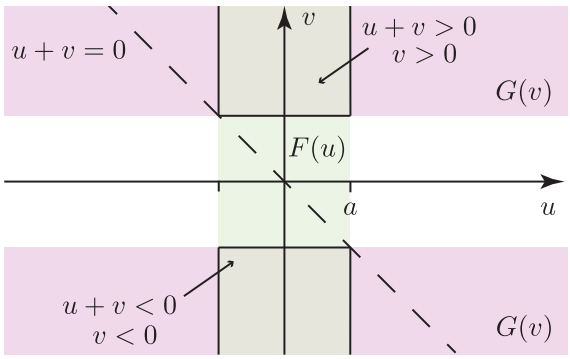


Figure A.2: Not to be confused with the 2D Fourier plane, in the 'product plane' of the two 1D Fourier domains of F and G their product is nonzero only in the boxed regions.

A.4 One dimensional fluid flow

Assuming flow of a fluid through a cavity, it is possible to calculate the fluid properties at any point as a function of the cavity cross-sectional area A. Assume the fluid starts at rest at infinity, with properties labelled with zero subscript. Conserving enthalphy per mol along the flow implies the relation

$$C_p T_0 = C_p T + \frac{1}{2} N_A m M^2 c_s^2 \tag{A.7}$$

where the particle mass is m, Avogadro's constant is N_A and the Mach number is M. The specific heat at constant pressure can be rewritten in terms of the ratio of specific heats for the gas $C_p = \gamma R/(\gamma - 1)$. Assuming a perfect barotropic gas such that $P = \rho RT/N_A$, $P \propto \rho^\gamma$ and $c_s^2 = \gamma P/m\rho$ (where ρ is the particle density), then

$$\frac{T_0}{T} = 1 + \frac{\gamma - 1}{2} M^2 \Rightarrow \frac{P_0}{P} = \left(1 + \frac{\gamma - 1}{2} M^2\right)^{\frac{\gamma}{\gamma - 1}}. \tag{A.8}$$

Conservation of particle number along the flow implies the condition $\rho u A = \text{const} \equiv \dot{N}$. Substituting for ρ and u in terms of M and quantities at infinity, one can show

$$\frac{\dot{N}}{A} = \rho u = P_0 \sqrt{\frac{\gamma N_A}{R m T_0}} M \left(1 + \frac{\gamma - 1}{2} M^2\right)^{-\frac{\gamma+1}{2(\gamma-1)}} \tag{A.9}$$

From the Bernoulli relation $u^2/2 + \gamma P/(\gamma - 1)m\rho = \text{const.}$ and particle number conservation, one can derive the differential relations

$$u\mathrm{d}u + \frac{\gamma}{\gamma - 1} \frac{1}{m} \left(\frac{1}{\rho} \mathrm{d}P - \frac{P}{\rho^2} \mathrm{d}\rho\right) \tag{A.10}$$

$$\frac{\mathrm{d}\rho}{\rho} = -\frac{\mathrm{d}u}{u} - \frac{\mathrm{d}A}{A} \tag{A.11}$$

which are combined to form

$$\left(u - \frac{c_s^2}{u}\right) \mathrm{d}u = \frac{\mathrm{d}A}{A}. \tag{A.12}$$

At the position of minimal cross-sectional area $dA = 0$, and at this point therefore $u = c_s \Rightarrow M = 1$. Call this minimal cross section A^* and substitute the sonic condition of Eq. A.12 into Eq. A.9, and arrive at the relation

$$\frac{A}{A^*} = \frac{1}{M}\left(\frac{2 + (\gamma - 1)M^2}{\gamma + 1}\right)^{\frac{\gamma+1}{2(\gamma-1)}} . \tag{A.13}$$

M can be calculated as a function of A/A^*, and from Eq. A.8 all other fluid properties follow.

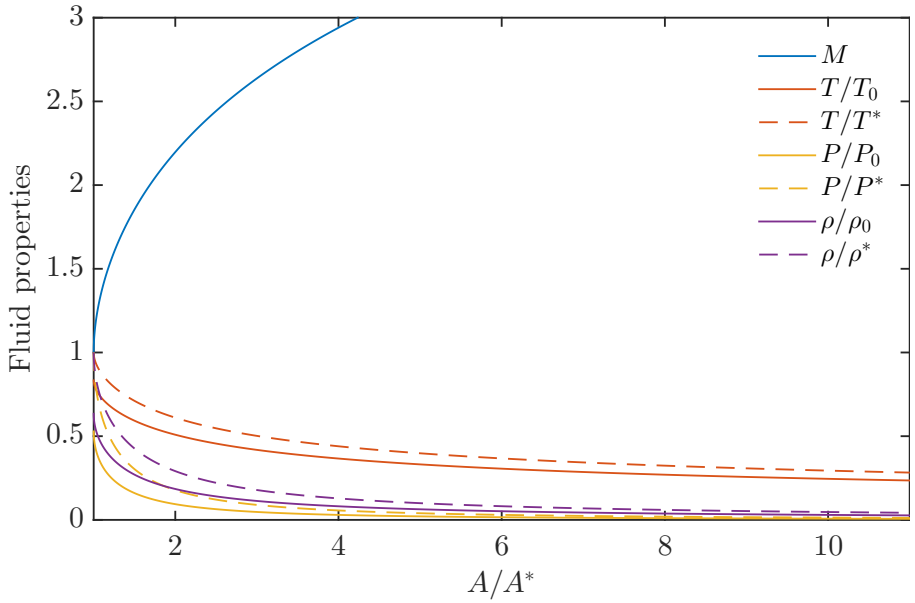


Figure A.3: Properties of a fluid undergoing 1D expansion in a nozzle.

A.5 Shepp-Logan phantom

The Shepp-Logan phantom is composed of 10 superimposed uniform ellipses, summarised in Table A.1. A is the value inside the ellipse, a, b are the horizontal and vertical semi-major axes respectively, x_0, y_0 are the co-ordinates of the centre, and ϕ is the rotation angle between the horizontal axis and the x-axis. The co-ordinate range of the phantom is $[-1, 1]$ in x and y.

A	a	b	x_0	y_0	ϕ
1	0.69	0.92	0	0	0
-0.98	0.6624	0.874	0	-0.0184	0
-0.02	0.11	0.31	0.22	0	-18
-0.02	0.16	0.41	-0.22	0	18
0.01	0.21	0.25	0	0.35	0
0.01	0.046	0.046	0	0.1	0
0.01	0.046	0.046	0	-0.1	0
0.01	0.046	0.023	-0.08	-0.605	0
0.01	0.023	0.023	0	-0.606	0

Table A.1: Shepp-Logan phantom parameters

A.6 Thermal electron response

Following Michel *et al.*[2], it is possible to derive the general non-relativistic response of plasma electrons to a slow perturbation caused by the beating of two intense laser pulses. Consider two beams of frequencies and wavevectors $(\omega_0, \mathbf{k}_0)$ and $(\omega_1, \mathbf{k}_1)$. As they beat they generate a slowly-moving ponderomotive potential at frequency $\omega = \omega_0 - \omega_1$ and wavevector $\mathbf{k} = \mathbf{k}_0 - \mathbf{k}_1$ (and also a fast beat, which we don't consider due to the inertia of the electrons). Assume that the electrons only respond to the slow beat, and that their distribution function is modified so that it has the form $f_e(\mathbf{r}, \mathbf{v}, t) = f_0(\mathbf{v}) + \delta f(\mathbf{r}, t)$. Further assume that δf only varies in space and time as $e^{i(\omega t - \mathbf{k}\cdot\mathbf{r})}$. Begin with the Vlasov equation for the electrons

$$\frac{\partial f_e}{\partial t} + \mathbf{v}\cdot\nabla f_e + \frac{\mathbf{F}_e}{m_e}\cdot\frac{\partial f_e}{\partial \mathbf{v}} = 0 \tag{A.14}$$

where the force acting on them is $\mathbf{F}_e$. We may substitute the above approximations and find

$$i(\omega - \mathbf{k}\cdot\mathbf{v})\delta f = -\frac{\mathbf{F}_e}{m_e}\cdot\frac{\partial f_0}{\partial \mathbf{v}}. \tag{A.15}$$

The forces on the electrons are the ponderomotive force due to the beat of the lasers, and the Coulomb forces arising from charge separation, so decompose $\mathbf{F}_e = -e\nabla(\phi_p + \phi)$ where ϕ_p is the ponderomotive potential due to the lasers and ϕ is the electrostatic potential due to the electron perturbation. These have the same wavelike spatial dependence as δf, so $\mathbf{F}_e = ie\mathbf{k}(\phi_p + \phi)$. Substituting and integrating over velocity, the perturbation to the electron density is

$$\delta n = \int \delta f \mathrm{d}^3 v = -\frac{e}{m_e}(\phi_p + \phi)\int \mathbf{k}\cdot\frac{\partial f_0}{\partial \mathbf{v}}\frac{1}{\omega - \mathbf{k}\cdot\mathbf{v}}\mathrm{d}^3 v. \tag{A.16}$$

The integral on the right hand side is proportional to the electron susceptibility χ_e, which we may substitute in

$$\delta n = -\frac{e}{m_e}(\phi_p + \phi)\chi_e \frac{m_e \epsilon_0 k^2}{e^2}. \tag{A.17}$$

To eliminate ϕ, note from Poisson's equation we have

$$\nabla^2\phi = -\frac{e\delta n}{\epsilon_0} \quad \Rightarrow \quad \phi = \frac{e\delta n}{\epsilon_0 k^2} \tag{A.18}$$

and so

$$\delta n = -\frac{\chi_e}{1+\chi_e}\frac{\epsilon_0 k^2}{e}\phi_p. \tag{A.19}$$

For a maxwellian plasma the electron susceptibility evaluates to

$$\chi_e = -\frac{1}{2k^2\lambda_D^2} Z'\left(\frac{\omega}{\sqrt{2}\omega_p k \lambda_D}\right) \tag{A.20}$$

[2] [373] P. Michel *et al. Physics of Plasmas.* (2013); [290] P. Michel *et al. Physical Review Letters.* (2014).

where λ_D is the Debye length, and the derivative of the plasma dispersion function is defined as

$$Z'(x) = -2\left(1 + x\sqrt{\pi}e^{-x^2}(i - \mathrm{erfi}(x))\right). \tag{A.21}$$

As an approximation for LWFA plasmas, take $T_e \to 0$ then in the large-argument limit of $|Z'(x)| \to x^{-2}$, $\chi_e \to -\omega_p^2/\omega^2$. For two lasers with $a_0 \gg a_1$ crossing at 90° the ponderomotive potential is

$$\phi_p = \frac{m_e \omega_0 c a_0 a_1}{2ek_0}, \tag{A.22}$$

so substituting into the expression for δn in the cold plasma limit

$$\delta n = \frac{\omega_p^2}{\omega - \omega_p^2}\frac{\epsilon_0 k^2}{e} \times \frac{m_e \omega_0 c a_0 a_1}{2ek_0} \tag{A.23}$$

$$\Rightarrow \quad \frac{\delta n}{n} = \frac{\omega_0^2}{\omega^2 - \omega_p^2} a_0 a_1 \tag{A.24}$$

where $k^2 = 2k_0^2$ as appropriate for this geometry. For beams of equal frequency where $\omega = 0$, $|\delta n/n| \approx a_0 a_1 n_c/n_e$. For underdense plasmas with $n_e \approx 10^{-3} n_c$ it is then possible have $\delta n/n = 1$ even for $a_1 < 10^{-3} a_0$, or for pulses of a million times difference in intensity.

A.7 Probe amplification model

Begin from the equations 2.108 and 2.109 in the nonrelativistic limit.

$$\left(\frac{\partial^2}{\partial t^2} - c^2\nabla^2 + \omega_p^2\right)\tilde{\mathbf{A}} = -\omega_p^2 \tilde{n}\mathbf{A}_0 \tag{A.25}$$

$$\left(\frac{\partial^2}{\partial t^2} + \omega_p^2\right)\tilde{n} = c^2\nabla^2\left(\mathbf{A}_0 \cdot \tilde{\mathbf{A}}\right) \tag{A.26}$$

Represent quantities by their envelopes, where $\tilde{\mathbf{A}}$ represents the probe beam, $\mathbf{A}_0$ represents the drive pulse and $\tilde{n}$ the normalised plasma density,

$$\tilde{\mathbf{A}} = a_1 e^{i(k_1 y - \omega_1 t)} + c.c \tag{A.27}$$

$$\mathbf{A}_0 = a_0 e^{i(k_0 x - \omega_0 t)} + c.c \tag{A.28}$$

$$\tilde{n} = n_0 e^{i(k_x x - k_y y - \omega_2 t)} + c.c, \tag{A.29}$$

and substitute into Eq. A.26

$$\left((\omega_p^2 - \omega_2^2) - 2i\omega_2\frac{\partial}{\partial t} + \frac{\partial^2}{\partial t^2}\right) n_0 = c^2\left(2ik_0 a_1\frac{\partial a_0}{\partial x} - 2ik_1 a_1\frac{\partial a_0}{\partial y} - 2ik_1 a_0\frac{\partial a_1}{\partial y} - a_0 a_1(k_0^2 + k_1^2)\right) \times e^{i(\Delta k_x x - \Delta k_y y - \Delta\omega t)} \tag{A.30}$$

where $\Delta k_x = k_0 - k_x$, $\Delta\omega = \omega_0 - \omega_1 - \omega_2$. Let $\omega_2 = \omega_p$, i.e. the grating oscillates at the plasma frequency, $k_x = k_y = k_0$, i.e. the grating period is set by the ponderomotive force and the two beams have equal wavelength. Then in the drive pulse frame $\xi = x - v_g t$

$$\left(-2i\omega_p\frac{\partial}{\partial t} + 2iv\omega_p\frac{\partial}{\partial\xi} + \frac{\partial^2}{\partial t^2} - 2v\frac{\partial}{\partial t}\frac{\partial}{\partial\xi} + v^2\frac{\partial^2}{\partial\xi^2}\right) n_0 = \tag{A.31}$$

$$2c^2\left(ik_0a_1\frac{\partial a_0}{\partial\xi} - ik_0a_1\frac{\partial a_0}{\partial y} - ik_0a_0\frac{\partial a_1}{\partial y} - k_0^2a_0a_1\right)e^{-i\omega_p t} \tag{A.32}$$

Approximate that late in time the grating reaches steady-state and $\partial_t \to 0$. If the pulses are also much longer than a laser wavelength then the ponderomotive term is dominated by the short-wavelength beat and

$$\left(\frac{\partial^2}{\partial\xi^2} + 2ik_p\frac{\partial}{\partial\xi}\right) n_0 = -2k_0^2a_0a_1. \tag{A.33}$$

This equation describes the convective growth of the grating from the head of the drive pulse, which in the moving frame is static. Roughly, the characteristic rate of growth for each term is

$$\frac{1}{L} \approx \frac{k_0^2a_0a_1}{k_p} = \frac{\omega_B^2}{c\omega_p} \to L \approx 5\lambda_0 \tag{A.34}$$

$$\frac{1}{L} \approx k_0\sqrt{2a_0a_1} = \frac{\sqrt{2}\omega_B}{c} \to L \approx 3\lambda_0 \tag{A.35}$$

The grating should then be expected to grow rapidly from the front of the pulse over a length scale of a few laser wavelengths. It can therefore be considered established by the time it is met by the probe pulse.

We now expand Eq. A.25 as an envelope, inserting the envelope for $\tilde{n}$

$$\left(-2i\omega_0\frac{\partial}{\partial t} + \frac{\partial^2}{\partial t^2} + 2ic^2k_0\frac{\partial}{\partial y} - c^2\frac{\partial^2}{\partial y^2}\right) a_1 = -\omega_p^2 n_0 a_0 e^{i\omega_p t_0} \tag{A.36}$$

where t_0 represents the phase of the grating at the time it is met by the probe pulse. The probe sees this phase as it propagates across the grating, so t_0 is a constant. Now switch to co-ordinates moving with probe $\eta = y - v_g t$ and let $a_1 = T(t)H(\eta)$, i.e. a constant shape growing in time. The LHS becomes

$$-2i\omega_0\dot{T}H + 2i\omega_0 v_g TH' + \ddot{T}H - 2v_g\dot{T}H' + v_g^2TH'' + 2ic^2k_0TH' - c^2TN'' \tag{A.37}$$

where dots denote time derivatives and primes denote spatial derivatives. Assume that $k_0 \gg \partial_\eta H$, i.e. the pulse is much longer than a wavelength, and that the pulse grows significantly over a transit time such that $cTH' \ll \dot{T}H$. Take the real part to find

$$\frac{\partial a_1}{\partial t} \equiv \dot{T}H = \frac{\omega_p^2 a_0}{2\omega_0}\frac{\delta n}{n}\sin(\omega_p t_0) = \left(\frac{\omega_p}{\omega_0}\right)^{3/2}\frac{\delta n}{n}\gamma_R\sin(\omega_p t_0) \tag{A.38}$$

where γ_R is the linear growth rate for RBS and $\omega_p t_0$ is the phase difference between the grating and the probe. Equivalent expressions are derived by Michel *et al.*[3] and McKinstrie *et al.*[4]. Also note the similarity to Eq. 2.134 of section 2.9.3, which is the same if transferred to the moving frame and $\omega_p t_0 = \pi/2$.

[3] [290] P. Michel *et al. Physical Review Letters.* (2014).
[4] [220] C. J. McKinstrie *et al. Physics of Plasmas.* (1996).

Bibliography

[1] E. Rutherford. *Philosophical Magazine*. **6,** pp. 669–688. (1911). (See p. 13)
"The scattering of α and β particles by matter and the structure of the atom"
DOI: 10.1080/14786435.2011.617037.

[2] J. D. Watson *et al. Nature*. **171,** 4356, pp. 737–738. (1953). (See p. 13)
"Molecular structure of nucleic acids"
DOI: 10.1097/BLO.0b013e3181468780.

[3] The ATLAS Collaboration. *Physics Letters B*. **716,** 1, pp. 1–29. (2012). (See p. 13)
"Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC"
DOI: 10.1016/j.physletb.2012.08.020.

[4] The CMS Collaboration. *Physics Letters B*. **716,** 1, pp. 30–61. (2012). (See p. 13)
"Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC"
DOI: 10.1016/j.physletb.2012.08.021.

[5] M. Silari. *Radiation Protection Dosimetry*. **146,** 4, pp. 440–450. (2011). (See p. 13)
"Applications of particle accelerators in medicine"
DOI: 10.1093/rpd/ncr243.

[6] J. F. Gibbons. *Proceedings of the IEEE*. **56,** 3, pp. 295–319. (1968). (See p. 13)
"Ion implantation in semiconductors-Part I Range distribution theory and experiments"
DOI: 10.1109/PROC.1968.6273.

[7] D. H. Bilderback *et al. Journal of Physics B*. **38,** 9, S773–S797. (2005). (See p. 13)
"Review of third and next generation synchrotron light sources"
DOI: 10.1088/0953-4075/38/9/022.

[8] J. D. Cockroft *et al. Proceedings of the Royal Society*. **129,** p. 477. (1932). (See p. 13)
"Experiments with High Velocity Positive Ions. -(I) Further Developments in the method of obtaining High Velocity Positive Ions."
URL: http://rspa.royalsocietypublishing.org/content/royprsa/136/830/619.full.pdf.

[9] A. R. Steere, pp. 1–82. (2005). (See p. 14)
A Timeline of Particle Accelerators.
URL: http://bt.pa.msu.edu/pub/papers/steeremsc/steeremsc.pdf.

[10] M. M. Alsharo'a *et al. Physical Review ST-AB*. **6,** 8, p. 081001. (2003). (See p. 14)
"Recent progress in neutrino factory and muon collider research within the Muon Collaboration"
DOI: 10.1103/PhysRevSTAB.6.081001.

[11] T. Behnke *et al.*, pp. 1–60. (2013). (See p. 15)
ILC Technical Design Report Volume 1: Executive Summary.
URL: http://arxiv.org/pdf/1306.6327v1.pdf.

[12] K. Saito *et al. The 10th Workshop on RF Superconductivity*, pp. 583–587. (2001). (See p. 15)
"Critical Field Limitation of the Niobium Superconducting RF Cavity"
URL: https://accelconf.web.cern.ch/accelconf/srf01/papers/ph003.pdf.

[13] E. Peralta *et al. Nature*. **503,** V, pp. 91–94. (2013). (See p. 15)
"Demonstration of electron acceleration in a laser-driven dielectric microstructure."
DOI: 10.1038/nature12664.

[14] R. B. Palmer. *Journal of Applied Physics*. **43,** 7, pp. 3014–3023. (1972). (See p. 15)
"Interaction of relativistic particles and free electromagnetic waves in the presence of a static helical magnet"
DOI: 10.1063/1.1661650.

[15] J. Duris *et al.* *Nature Communications.* **5,** p. 4928. (2014). (See p. 15)
"High-quality electron beams from a helical inverse free-electron laser accelerator"
DOI: 10.1038/ncomms5928.

[16] V. I. Veksler. *Proc. CERN Symp.* Pp. 80–83. (1956). (See p. 15)
"Coherent principle of acceleration of charged particles"
DOI: 10.5170/CERN-1956-025.80.

[17] Y. B. Fainberg. *Plasma Physics Reports.* **10,** 6, p. 750. (1967). (See p. 15)
"Acceleration of charged particles in plasma"
URL: http://iopscience.iop.org/article/10.1070/PU1968v010n06ABEH003715/pdf.

[18] T. Tajima *et al.* *Physical Review Letters.* **43,** 4, pp. 267–270. (1979). (See . 15, 101)
"Laser electron accelerator"
URL: http://link.aps.org/doi/10.1103/PhysRevLett.43.267.

[19] M. N. Rosenbluth. *Physical Review Letters.* **29,** 11. (1972). (See p. 16)
"Excitation of plasma waves by two laser beams"
URL: http://link.aps.org/doi/10.1103/PhysRevLett.29.701.

[20] C. E. Clayton *et al.* *Physical Review Letters.* **54,** 21, pp. 2343–2346. (1985). (See p. 16)
"Relativistic Plasma-Wave Excitation by Collinear Optical Mixing"
DOI: 10.1103/PhysRevLett.54.2343.

[21] C. E. Clayton *et al.* *Physical Review Letters.* **70,** 1, pp. 37–40. (1993). (See p. 16)
"Ultrahigh-Gradient Acceleration of Injected Electrons By Laser-Excited Relativistic Electron-Plasma Waves"
DOI: 10.1103/PhysRevLett.70.37.

[22] N. E. Andreev. *JETP Letters.* **55,** 10, p. 571. (1992). (See p. 16)
"Resonant excitation of wakefields by a laser pulse in a plasma"
URL: http://www.jetpletters.ac.ru/ps/1277/article{_}19308.pdf.

[23] J. Krall *et al.* *Physical Review E.* **48,** 3, pp. 2157–2161. (1993). (See p. 16)
"Enhanced acceleration in a self-modulated laser wakefield accelerator".

[24] K. Nakajima *et al.* *Physical Review Letters.* **75,** 5, pp. 984–984. (1995). (See p. 17)
"Observation of Ultrahigh Gradient Electron Acceleration by a Self-Modulated Intense Short Laser Pulse"
DOI: 10.1103/PhysRevLett.75.984.2.

[25] A. Modena *et al.* *Nature.* **377,** 6550, pp. 606–608. (1995). (See p. 17)
"Electron acceleration from the breaking of relativistic plasma waves"
DOI: 10.1038/377606a0.

[26] R. Wagner *et al.* *Physical Review Letters.* **78,** 16, pp. 3125–3128. (1997). (See . 17, 35)
"Electron Acceleration by a Laser Wakefield in a Relativistically Self-Guided Channel"
DOI: 10.1103/PhysRevLett.78.3125.

[27] D. Gordon *et al.* *Physical Review Letters.* **80,** 10, pp. 2133–2136. (1998). (See p. 17)
"Observation of Electron Energies Beyond the Linear Dephasing Limit from a Laser-Excited Relativistic Plasma Wave"
DOI: 10.1103/PhysRevLett.80.2133.

[28] M. D. Perry *et al.* *Science.* **264,** 5161, pp. 917–24. (1994). (See p. 17)
"Terawatt to petawatt subpicosecond lasers."
DOI: 10.1126/science.264.5161.917.

[29] D. Strickland *et al.* *Optics Communications.* **56,** 3, pp. 219–221. (1985). (See p. 17)
"Compression of amplified chirped optical pulses"
URL: http://www.sciencedirect.com/science/article/pii/0030401885901518.

[30] P. Maine *et al.* *IEEE Journal of Quantum Electronics.* **24,** 2, pp. 398–403. (1988). (See p. 17)
"Generation of ultrahigh peak power pulses by chirped pulse amplification"
DOI: 10.1109/3.137.

[31] V. Malka *et al.* *Science.* **298,** 5598, pp. 1596–1600. (2002). (See . 17, 36, 144)
"Electron acceleration by a wake field forced by an intense ultrashort laser pulse."
URL: http://www.ncbi.nlm.nih.gov/pubmed/12446903.

[32] S. P. D. Mangles *et al.* *Nature.* **431,** 7008, pp. 535–538. (2004). (See . 17, 45, 144)
"Monoenergetic beams of relativistic electrons from intense laser-plasma interactions."
URL: http://www.nature.com/nature/journal/v431/n7008/abs/nature02939.html.

[33] C. Geddes *et al. Nature.* **431,** 7008, pp. 538–541. (2004). (See . 17, 45)
"High-quality electron beams from a laser wakefield accelerator using plasma-channel guiding."
URL: http://www.nature.com/nature/journal/v431/n7008/abs/nature02900.html.

[34] J. Faure *et al. Nature.* **431,** 7008, pp. 541–544. (2004). (See . 17, 45)
"A laser-plasma accelerator producing monoenergetic electron beams."
URL: http://www.nature.com/nature/journal/v431/n7008/abs/nature02963.html.

[35] W. P. Leemans *et al. Nature Physics.* **2,** 10, pp. 696–699. (2006). (See p. 18)
"GeV electron beams from a centimetre-scale accelerator"
DOI: 10.1038/nphys418.

[36] S. Kneip *et al. Physical Review Letters.* **103,** 3, p. 035002. (2009). (See p. 18)
"Near-GeV Acceleration of Electrons by a Nonlinear Plasma Wave Driven by a Self-Guided Laser Pulse"
DOI: 10.1103/PhysRevLett.103.035002.

[37] X. Wang *et al. Nature Communications.* **4,** May. (2013). (See p. 18)
"Quasi-monoenergetic laser-plasma acceleration of electrons to 2 GeV"
DOI: 10.1038/ncomms2988.

[38] H. T. Kim *et al. Physical Review Letters.* **111,** 16, pp. 1–5. (2013). (See . 18, 50, 70)
"Enhancement of electron energy to the multi-gev regime by a dual-stage laser-wakefield accelerator pumped by petawatt laser pulses"
DOI: 10.1103/PhysRevLett.111.165002.

[39] W. P. Leemans *et al. Physical Review Letters.* **113,** p. 245002. (2014). (See p. 18)
"Multi-GeV Electron Beams from Capillary-Discharge-Guided Subpetawatt Laser Pulses in the Self-Trapping Regime"
DOI: 10.1103/physrevlett.113.245002.

[40] J. Faure *et al. Nature.* **444,** 7120, pp. 737–739. (2006). (See . 18, 50)
"Controlled injection and acceleration of electrons in plasma wakefields by colliding laser pulses."
URL: http://www.ncbi.nlm.nih.gov/pubmed/17151663.

[41] S. P. D. Mangles *et al. Physical Review Letters.* **96,** 21, p. 215001. (2006). (See . 18, 46)
"Laser-wakefield acceleration of monoenergetic electron beams in the first plasma-wave period."
URL: http://link.aps.org/doi/10.1103/PhysRevLett.96.215001.

[42] A. Buck *et al. Physical Review Letters.* **110,** 18, pp. 1–5. (2013). (See . 18, 70)
"Shock-front injector for high-quality laser-plasma acceleration"
DOI: 10.1103/PhysRevLett.110.185006.

[43] S. Kneip *et al. Physical Review ST-AB.* **15,** 2, pp. 2–6. (2012). (See . 18, 21, 53)
"Characterization of transverse beam emittance of electrons from a laser-plasma wakefield accelerator in the bubble regime using betatron x-ray radiation"
DOI: 10.1103/PhysRevSTAB.15.021302.

[44] R. Weingartner *et al. Physical Review ST-AB.* **15,** 11, p. 111302. (2012). (See p. 18)
"Ultralow emittance electron beams from a laser-wakefield accelerator"
DOI: 10.1103/PhysRevSTAB.15.111302.

[45] M. Migliorati *et al. Physical Review ST-AB.* **16,** 1, p. 011302. (2013). (See p. 18)
"Intrinsic normalized emittance growth in laser-driven electron accelerators"
DOI: 10.1103/PhysRevSTAB.16.011302.

[46] B. B. Pollock *et al. Physical Review Letters.* **107,** 4, pp. 1–4. (2011). (See . 18, 50, 70)
"Demonstration of a narrow energy spread, 0.5GeV electron beam from a two-stage laser wakefield accelerator"
DOI: 10.1103/PhysRevLett.107.045001.

[47] G. Golovin *et al. Physical Review ST-AB.* **18,** 1, p. 011301. (2015). (See . 18, 70)
"Tunable monoenergetic electron beams from independently controllable laser-wakefield acceleration and injection"
DOI: 10.1103/PhysRevSTAB.18.011301.

[48] C. McGuffey *et al. Physical Review Letters.* **104,** 2, pp. 1–4. (2010). (See . 18, 50)
"Ionization Induced Trapping in a Laser Wakefield Accelerator"
DOI: 10.1103/PhysRevLett.104.025004.

[49] S. P. D. Mangles *et al. Physics of Plasmas.* **14,** 5, p. 056702. (2007). (See p. 18)
"On the stability of laser wakefield electron accelerators in the monoenergetic regime"
DOI: 10.1063/1.2436481.

[50] J. Osterhoff *et al. Physical Review Letters.* **101,** 8, p. 085002. (2008). (See . 18, 70)
"Generation of Stable, Low-Divergence Electron Beams by Laser-Wakefield Acceleration in a Steady-State-Flow Gas Cell"
DOI: 10.1103/PhysRevLett.101.085002.

[51] S. Banerjee *et al. Physical Review ST-AB.* **16,** 3, pp. 1–12. (2013). (See p. 18)
"Stable, tunable, quasimonoenergetic electron beams produced in a laser wakefield near the threshold for self-injection"
DOI: 10.1103/PhysRevSTAB.16.031302.

[52] O. Lundh *et al. Nature Physics.* **7,** 3, pp. 219–222. (2011). (See p. 18)
"Few femtosecond, few kiloampere electron bunch produced by a laser plasma accelerator"
DOI: 10.1038/nphys1872.

[53] O. Lundh *et al. Physical Review Letters.* **110,** 6, pp. 1–5. (2013). (See p. 18)
"Experimental measurements of electron-bunch trains in a laser-plasma accelerator"
DOI: 10.1103/PhysRevLett.110.065005.

[54] A. J. Gonsalves *et al. Nature Physics.* **7,** 11, pp. 862–866. (2011). (See . 18, 51)
"Tunable laser plasma accelerator based on longitudinal density tailoring"
DOI: 10.1038/nphys2071.

[55] W. P. Leemans *et al. Physics Today.* **62,** 3, pp. 44–49. (2009). (See p. 18)
"Laser-driven plasma-wave electron accelerators"
DOI: 10.1063/1.3099645.

[56] S. Steinke *et al. Nature.* (2016). (See p. 18)
"Multistage coupling of independent laser-plasma accelerators"
DOI: 10.1038/nature16525.

[57] E. Esarey *et al. IEEE Transactions on Plasma Science.* **24,** 2, pp. 252–288. (1996). (See . 18, 34)
"Overview of plasma-based accelerator concepts"
DOI: 10.1109/27.509991.

[58] E. Esarey *et al. Reviews of Modern Physics.* **81,** 3, pp. 1229–1285. (2009). (See . 18, 40, 168)
"Physics of laser-driven plasma-based electron accelerators"
DOI: 10.1103/RevModPhys.81.1229.

[59] I Blumenfeld *et al. Nature.* **445,** 7129, pp. 741–4. (2007). (See p. 18)
"Energy doubling of 42 GeV electrons in a metre-scale plasma wakefield accelerator."
DOI: 10.1038/nature05538.

[60] S. Corde *et al. Nature.* **524,** 7566, pp. 442–445. (2015). (See p. 18)
"Multi-gigaelectronvolt acceleration of positrons in a self-loaded plasma wakefield"
DOI: 10.1038/nature14890.

[61] M. Litos *et al. Nature.* **515,** pp. 92–95. (2014). (See p. 18)
"High-efficiency acceleration of an electron beam in a plasma wakefield accelerator"
DOI: 10.1038/nature13882.

[62] A. Aschikhin *et al.*, pp. 1–19. (2015). (See p. 19)
The FLASHForward Facility at DESY.

[63] R Assmann *et al. Plasma Physics and Controlled Fusion.* **56,** 8, p. 084013. (2014). (See p. 19)
"Proton-driven plasma wakefield acceleration: a path to the future of high-energy particle physics"
DOI: 10.1088/0741-3335/56/8/084013.

[64] H. Daido *et al. Reports on Progress in Physics.* **75,** 5, p. 056401. (2012). (See p. 19)
"Review of laser-driven ion sources and their applications"
DOI: 10.1088/0034-4885/75/5/056401.

[65] H. C. Pollock. *Am J. Phys.* **51,** 3, p. 278. (1983). (See p. 19)
"The discovery of synchrotron radiation"
DOI: doi:10.1119/1.13289.

[66] E. M. Rowe. *Particle Accelerators.* **4,** p. 211. (1973). (See p. 19)
"TANTALUS I : A Dedicated Storage Ring Synchrotron Radiation Source".

[67] D. H. Bilderback *et al. Journal of Physics B.* **38,** 9, S773–S797. (2005). (See p. 19)
"Review of third and next generation synchrotron light sources"
DOI: 10.1088/0953-4075/38/9/022.

[68] S. Tazzari. *IEEE Transactions on Nuclear Science*. **32,** 5, pp. 3400–3402. (1985). (See p. 19)
"The European Synchrotron Radiation Facility".

[69] *Diamond*. 2015. URL: www.diamond.ac.uk (see p. 19).

[70] H. Motz *et al. Journal of Applied Physics*. **24,** 7, pp. 826–33. (1953). (See p. 20)
"Experiments on Radiation by Fast Electron Beams"
DOI: 10.1063/1.1721389.

[71] M. Berndt *et al. IEEE Transactions on Nuclear Science*. **26,** 3, pp. 3812–3815. (1979). (See p. 20)
"Initial Operation of SSRL Wiggler in Spear"
DOI: 10.1017/CBO9781107415324.004.

[72] B. W. J. McNeil *et al. Nature Photonics*. **4,** 12, pp. 814–821. (2010). (See p. 20)
"X-ray free-electron lasers"
DOI: 10.1038/nphoton.2010.239.

[73] P. Emma *et al. Nature Photonics*. **4,** 9, pp. 641–647. (2010). (See p. 20)
"First lasing and operation of an ångstrom-wavelength free-electron laser"
DOI: 10.1038/nphoton.2010.176.

[74] S Schreiber *et al. Proc. 1st Int. Particle Accelerator Conf.* Pp. 2149–2151. (2010). (See p. 20)
"FEL User Facility FLASH"
URL: http://accelconf.web.cern.ch/Accelconf/IPAC10/papers/tupe004.pdf.

[75] D. Pile. *Nature Photonics*. **5,** 8, pp. 456–457. (2011). (See p. 20)
"X-rays: First light from SACLA"
DOI: 10.1038/nphoton.2011.178.

[76] A. Massimo. *Technical Design Report*. (2007). (See p. 20)
"The European X-ray Free-Electron laser".

[77] T. Shintake. *Proceedings of IPAC'10*, pp. 1285–1289. (2010). (See p. 20)
"Status Report on Japanese XFEL Construction Project at Spring-8".

[78] B. D. Patterson *et al. New Journal of Physics*. **12,** 3, p. 035012. (2010). (See p. 20)
"Coherent science at the SwissFEL x-ray laser"
DOI: 10.1088/1367-2630/12/3/035012.

[79] D. H. Whittum. *Physics of Fluids B: Plasma Physics*. **4,** 3, p. 730. (1992). (See p. 20)
"Electromagnetic instability of the ion-focused regime"
DOI: 10.1063/1.860271.

[80] S. Wang *et al. Physical Review Letters*. **88,** 13, p. 135004. (2002). (See . 20, 51)
"X-ray emission from betatron motion in a plasma wiggler."
DOI: 10.1103/PhysRevLett.88.135004.

[81] A. Rousse *et al. Physical Review Letters*. **93,** 13, p. 135005. (2004). (See . 20, 51, 175)
"Production of a keV x-ray beam from synchrotron radiation in relativistic laser-plasma interaction."
URL: http://link.aps.org/doi/10.1103/PhysRevLett.93.135005.

[82] K. Ta Phuoc *et al. Physics of Plasmas*. **12,** 2, p. 023101. (2005). (See p. 20)
"Laser based synchrotron radiation"
DOI: 10.1063/1.1842755.

[83] M. Schnell *et al. Journal of Plasma Physics*. **81,** 04, p. 475810401. (2015). (See p. 20)
"Characterization and application of hard x-ray betatron radiation generated by relativistic electrons from a laser-wakefield accelerator"
DOI: 10.1017/S0022377815000379.

[84] K. Ta Phuoc *et al. Physics of Plasmas*. **14,** 8, p. 080701. (2007). (See p. 20)
"Demonstration of the ultrafast nature of laser produced betatron radiation"
DOI: 10.1063/1.2754624.

[85] K. Ta Phuoc *et al. Physical Review Letters*. **97,** 22, p. 225002. (2006). (See . 20, 53)
"Imaging Electron Trajectories in a Laser-Wakefield Cavity Using Betatron X-Ray Radiation"
DOI: 10.1103/PhysRevLett.97.225002.

[86] F. Albert *et al. Physical Review E*. **77,** 5, pp. 1–6. (2008). (See . 20, 53)
"Betatron oscillations of electrons accelerated in laser wakefields characterized by spectral x-ray analysis"
DOI: 10.1103/PhysRevE.77.056402.

[87] S. Kneip *et al. Physical Review Letters.* **100,** 10, p. 105006. (2008). (See p. 21)
"Observation of synchrotron radiation from electrons accelerated in a petawatt-laser-generated plasma cavity."
URL: http://link.aps.org/doi/10.1103/PhysRevLett.100.105006.

[88] S. Fourmaux *et al. New Journal of Physics.* **13,** (2011). (See p. 21)
"Demonstration of the synchrotron-type spectrum of laser-produced Betatron radiation"
DOI: 10.1088/1367-2630/13/3/033017.

[89] S. Kneip *et al. Nature Physics.* **6,** 12, pp. 980–983. (2010). (See . 21, 175, 181)
"Bright spatially coherent synchrotron X-rays from a table-top source"
DOI: 10.1038/nphys1789.

[90] F. Albert *et al. Physical Review Letters.* **111,** 23, p. 235004. (2013). (See . 21, 53, 203)
"Angular Dependence of Betatron X-Ray Spectra from a Laser-Wakefield Accelerator"
DOI: 10.1103/PhysRevLett.111.235004.

[91] K. T. Phuoc *et al. Physical Review Letters.* **91,** 19, p. 195001. (2003). (See p. 21)
"X-ray radiation from nonlinear Thomson scattering of an intense femtosecond laser on relativistic electrons in a helium plasma"
DOI: 10.1103/PhysRevLett.91.195001.

[92] N. Powers *et al. Nature Photonics.* November, pp. 1–4. (2013). (See p. 21)
"Quasi-monoenergetic and tunable X-rays from a laser-driven Compton light source"
DOI: 10.1038/nphoton.2013.314.

[93] K. Khrennikov *et al. Physical Review Letters.* **114,** 19, pp. 1–5. (2015). (See p. 21)
"Tunable All-Optical Quasimonochromatic Thomson X-Ray Source in the Nonlinear Regime"
DOI: 10.1103/PhysRevLett.114.195003.

[94] K. Ta Phuoc *et al. Nature Photonics.* **6,** 5, pp. 308–311. (2012). (See p. 21)
"All-optical Compton gamma-ray source"
DOI: 10.1038/nphoton.2012.82.

[95] G. Sarri *et al. Physical Review Letters.* **113,** 22, pp. 1–5. (2014). (See p. 21)
"Ultrahigh brilliance multi-MeV γ-ray beams from nonlinear relativistic thomson scattering"
DOI: 10.1103/PhysRevLett.113.224801.

[96] H.-P. Schlenvoigt *et al. Nature Physics.* **4,** 2, pp. 130–133. (2008). (See p. 21)
"A compact synchrotron radiation source driven by a laser-plasma wakefield accelerator"
DOI: 10.1038/nphys811.

[97] M. Fuchs *et al. Nature Physics.* **5,** 11, pp. 826–829. (2009). (See . 21, 70)
"Laser-driven soft-X-ray undulator source"
DOI: 10.1038/nphys1404.

[98] K. Achterhold *et al. Scientific Reports.* **3,** p. 1313. (2013). (See . 21, 178)
"Monochromatic computed tomography with a compact laser-driven X-ray source."
DOI: 10.1038/srep01313.

[99] S. Corde *et al. Reviews of Modern Physics.* **85,** 1. (2013). (See p. 21)
"Femtosecond x rays from laser-plasma accelerators"
DOI: 10.1103/RevModPhys.85.1.

[100] F. Albert *et al. Plasma Physics and Controlled Fusion.* **56,** 8, p. 084015. (2014). (See p. 21)
"Laser wakefield accelerator based light sources: potential applications and requirements"
DOI: 10.1088/0741-3335/56/8/084015.

[101] M. Nicolai *et al. Applied Physics B: Lasers and Optics*, pp. 1–9. (2013). (See p. 21)
"Realizing a laser-driven electron source applicable for radiobiological tumor irradiation"
DOI: 10.1007/s00340-013-5747-0.

[102] Z. H. He *et al. Applied Physics Letters.* **102,** 6, pp. 1–5. (2013). (See p. 21)
"Electron diffraction using ultrafast electron bunches from a laser-wakefield accelerator at kHz repetition rate"
DOI: 10.1063/1.4792057.

[103] G. Sarri *et al. Physical Review Letters.* **110,** 25, p. 255002. (2013). (See p. 21)
"Table-Top Laser-Based Source of Femtosecond, Collimated, Ultrarelativistic Positron Beams"
DOI: 10.1103/PhysRevLett.110.255002.

[104] S. Fourmaux *et al. Optics Letters.* **36,** 13, pp. 2426–8. (2011). (See . 21, 54, 175)
"Single shot phase contrast imaging using laser-produced Betatron x-ray beams."
URL: http://www.ncbi.nlm.nih.gov/pubmed/21725433.

[105] S. Kneip *et al. Applied Physics Letters.* **99,** 9, p. 093701. (2011). (See . 21, 54, 175)
"X-ray phase contrast imaging of biological specimens with femtosecond pulses of betatron radiation from a compact laser plasma wakefield accelerator"
DOI: 10.1063/1.3627216.

[106] Z. Najmudin *et al. Philosophical Transactions of the Royal Society.* **372,** p. 20130032. (2014). (See . 21, 54)
"Compact laser accelerators for X-ray phase-contrast imaging"
URL: http://www.ncbi.nlm.nih.gov/pmc/articles/PMC3900035/.

[107] J. Wenz *et al. Nature Communications.* **6,** May, p. 7568. (2015). (See . 21, 176)
"Quantitative X-ray phase-contrast microtomography from a compact laser-driven betatron source"
DOI: 10.1038/ncomms8568.

[108] J. M. Cole *et al. Scientific Reports.* **5,** p. 13244. (2015). (See . 21, 175)
"Laser-wakefield accelerators as hard x-ray sources for 3D medical imaging of human bone"
DOI: 10.1038/srep13244.

[109] A. Döpp *et al. arXiv*, pp. 2–5. (2015). (See p. 21)
"An all-optical Compton source for single-exposure X-ray imaging"
URL: http://arxiv.org/abs/1510.08316.

[110] M. Z. Mo *et al. Review of Scientific Instruments.* **84,** 12, p. 123106. (2013). (See p. 21)
"Laser wakefield generated X-ray probe for femtosecond time-resolved measurements of ionization states of warm dense aluminum"
DOI: 10.1063/1.4842237.

[111] M. Ammosov *et al. Sov Phys JETP.* **64,** 6, pp. 1191 –1194. (1986). (See p. 25)
"Tunnel ionization of complex atoms and of atomic ions in an alternating electromagnetic field"
URL: http://www.jetp.ac.ru/cgi-bin/dn/e{_}064{_}06{_}1191.pdf.

[112] D. L. Bruhwiler *et al. Physics of Plasmas.* **10,** 5 II, pp. 2022–2030. (2003). (See p. 25)
"Particle-in-cell simulations of tunneling ionization effects in plasma-based accelerators"
DOI: 10.1063/1.1566027.

[113] W. Chang, p. 245. (2005). (See p. 25)
Principles of Lasers and Optics.

[114] S. Feng *et al. Optics Letters.* **26,** 8, pp. 485–7. (2001). (See p. 26)
"Physical origin of the Gouy phase shift"
URL: http://www.ncbi.nlm.nih.gov/pubmed/18040360.

[115] C. G. Chen *et al. JOSA A.* (2002). (See p. 27)
"Analyses of vector Gaussian beam propagation and the validity of paraxial and spherical approximations"
URL: https://www.osapublishing.org/josaa/abstract.cfm?uri=josaa-19-2-404.

[116] B. Quesnel *et al. Physical Review E.* **58,** 3, pp. 3719–3732. (1998). (See . 27, 33)
"Theory and simulation of the interaction of ultraintense laser pulses with electrons in vacuum"
DOI: 10.1103/PhysRevE.58.3719.

[117] Z. H. He *et al. Nature Communications.* **6,** May, p. 7156. (2015). (See p. 28)
"Coherent control of plasma dynamics"
DOI: 10.1038/ncomms8156.

[118] S. P. D. Mangles *et al. Applied Physics Letters.* **95,** 18, p. 181106. (2009). (See . 28, 138)
"Controlling the spectrum of x-rays generated in a laser-plasma accelerator by tailoring the laser wavefront"
DOI: 10.1063/1.3258022.

[119] S. P. D. Mangles *et al. Physical Review ST-AB.* **15,** 1, p. 011302. (2012). (See p. 28)
"Self-injection threshold in self-guided laser wakefield accelerators"
DOI: 10.1103/PhysRevSTAB.15.011302.

[120] B. Beaurepaire *et al. Physical Review X.* **5,** 3, p. 031012. (2015). (See p. 28)
"Effect of the Laser Wave Front in a Laser-Plasma Accelerator"
DOI: 10.1103/PhysRevX.5.031012.

[121] J. Vieira *et al. Plasma Physics and Controlled Fusion.* **54,** 5, p. 055010. (2012). (See p. 28)
"Influence of realistic parameters on state-of-the-art laser wakefield accelerator experiments"
DOI: 10.1088/0741-3335/54/5/055010.

[122] J.-H. Yang *et al. Plasma Physics and Controlled Fusion.* **53,** 12, p. 125006. (2011). (See p. 29)
"Explicit general solutions to relativistic electron dynamics in plane-wave electromagnetic fields and simulations of ponderomotive acceleration"
DOI: 10.1088/0741-3335/53/12/125006.

[123] P. Mora *et al. Physics Of Plasmas.* **4,** 1, pp. 217–229. (1997). (See . 33, 101)
"Kinetic modeling of intense, short laser pulses propagating in tenuous plasmas"
DOI: 10.1063/1.872134.

[124] W. B. Mori. *IEEE Journal of Quantum Electronics.* **33,** 11, pp. 1942–1953. (1997). (See . 33, 35)
"The physics of the nonlinear optics of plasmas at relativistic intensities for short-pulse lasers"
DOI: 10.1109/3.641309.

[125] P. Sprangle *et al. IEEE Transactions on Plasma Science.* **15,** 2, pp. 145–153. (1987). (See p. 34)
"Relativistic Self-Focusing of Short-Pulse Radiation Beams in Plasmas"
DOI: 10.1109/TPS.1987.4316677.

[126] A. G. R. Thomas *et al. Physical Review Letters.* **98,** 9, p. 095004. (2007). (See p. 35)
"Effect of laser-focusing conditions on propagation and monoenergetic electron production in laser-wakefield accelerators."
URL: http://link.aps.org/doi/10.1103/PhysRevLett.98.095004.

[127] G.-Z. Sun *et al. Physics of Fluids.* **30,** 2, p. 526. (1987). (See p. 35)
"Self-focusing of short intense pulses in plasmas"
DOI: 10.1063/1.866349.

[128] E. Esarey *et al. Physics of Fluids B: Plasma Physics.* **5,** 7, p. 2690. (1993). (See p. 35)
"Optically guided laser wake-field acceleration"
DOI: 10.1063/1.860707.

[129] C. Ren *et al. Physical Review E.* **63,** 2, pp. 1–8. (2001). (See p. 36)
"Compressing and focusing a short laser pulse by a thin plasma lens"
DOI: 10.1103/PhysRevE.63.026411.

[130] J. Vieira *et al. New Journal of Physics.* **12,** 4, p. 045025. (2010). (See p. 36)
"Onset of self-steepening of intense laser pulses in plasmas"
DOI: 10.1088/1367-2630/12/4/045025.

[131] P. Gibbon, p. 328. (2005). (See . 36, 39, 56)
Short Pulse Laser Interactions with Matter: An Introduction.

[132] P. Sprangle *et al. Applied Physics Letters.* **53,** 1988, pp. 2146–2148. (1988). (See p. 38)
"Laser wakefield acceleration and relativistic optical guiding"
DOI: 10.1063/1.100300.

[133] L. Gorbunov *et al. Sov Phys JETP.* **66,** 2, pp. 290–294. (1987). (See . 38, 47)
"Excitation of plasma waves by an electromagnetic wave packet"
URL: http://www.jetp.ac.ru/cgi-bin/dn/e{_}066{_}02{_}0290.pdf.

[134] J. M. Dawson. *Physical Review.* **113,** 2, pp. 383–387. (1959). (See . 38, 42)
"Nonlinear electron oscillations in a cold plasma"
DOI: 10.1103/PhysRev.113.383.

[135] E. Esarey *et al. Physics of Plasmas.* **2,** 5, p. 1432. (1995). (See . 39, 42)
"Trapping and acceleration in nonlinear plasma waves"
DOI: 10.1063/1.871358.

[136] P. Sprangle *et al. Physical Review Letters.* **64,** 17, pp. 2011–2014. (1990). (See p. 40)
"Nonlinear Theory of Intense Laser-Plasma Interactions"
DOI: 10.1103/PhysRevLett.64.2011.

[137] V. I. Berezhiani *et al. Physics Letters A.* **148,** 6-7, pp. 338–340. (1990). (See p. 40)
"Relativistic wake-field generation by an intense laser pulse in a plasma"
DOI: 10.1016/0375-9601(90)90813-4.

[138] S. V. Bulanov *et al. JETP Letters.* **50,** 4, pp. 176–178. (1989). (See p. 40)
"Excitation of ultrarelativistic plasma waves by pulse of electromagnetic radiation"
URL: http://jetpletters.ac.ru/ps/1127/article{_}17078.pdf.

[139] S. V. Bulanov *et al. Physical Review Letters.* **74,** 5, pp. 710–713. (1995). (See p. 40)
"Two-Dimensional Regimes of Self-Focusing, Wake Field Generation, and Induced Focusing of a Short Intense Laser Pulse in an Underdense Plasma"
DOI: 10.1103/PhysRevLett.74.710.

[140] S. V. Bulanov *et al. Physics of Plasmas.* **12,** 7, pp. 1–11. (2005). (See p. 42)
"Spectral and dynamical features of the electron bunch accelerated by a short-pulse high intensity laser in an underdense plasma"
DOI: 10.1063/1.1948347.

[141] A. I. Akhiezer *et al. Sov. Phys. JETP.* **3,** p. 696. (1956). (See p. 42)
"Theory of Wave Motion of an Electron Plasma"
URL: http://www.jetp.ac.ru/cgi-bin/dn/e{_}003{_}05{_}0696.pdf.

[142] W. B. Mori *et al. Physica Scripta.* **T30,** pp. 127–133. (1990). (See p. 42)
"Wavebreaking of longitudinal plasma oscillations"
DOI: 10.1088/0031-8949/1990/T30/018.

[143] S. V. Bulanov *et al. Physics of Plasmas.* **19,** 11. (2012). (See p. 42)
"On the breaking of a plasma wave in a thermal plasma. I. the structure of the density singularity"
DOI: 10.1063/1.4764052.

[144] S. V. Bulanov *et al. Physical Review Letters.* **78,** 22, pp. 4205–4208. (1997). (See p. 42)
"Transverse-Wake Wave Breaking"
DOI: 10.1103/PhysRevLett.78.4205.

[145] A. Pukhov *et al. Applied Physics B: Lasers and Optics.* **74,** 4-5, pp. 355–361. (2002). (See . 43, 101)
"Laser wake field acceleration: the highly non-linear broken-wave regime"
DOI: 10.1007/s003400200795.

[146] I. Kostyukov *et al. Physics of Plasmas.* **11,** 11, p. 5256. (2004). (See p. 44)
"Phenomenological theory of laser-plasma interaction in bubble regime"
DOI: 10.1063/1.1799371.

[147] F. Dorchies *et al. Physics of Plasmas.* **6,** 7, p. 2903. (1999). (See p. 44)
"Acceleration of injected electrons in a laser wakefield experiment"
DOI: 10.1063/1.873248.

[148] S. Kalmykov *et al. Physics of Plasmas.* **13,** 11, p. 113102. (2006). (See p. 44)
"Injection, trapping, and acceleration of electrons in a three-dimensional nonlinear laser wakefield"
DOI: 10.1063/1.2363172.

[149] W. Lu *et al. Physics of Plasmas.* **13,** 5, p. 056709. (2006). (See p. 45)
"A nonlinear theory for multidimensional relativistic plasma wave wakefields"
DOI: 10.1063/1.2203364.

[150] W. Lu *et al. Physical Review ST-AB.* **10,** 6, pp. 1–12. (2007). (See . 45, 123)
"Generating multi-GeV electron bunches using single stage laser wakefield acceleration in a 3D nonlinear regime"
DOI: 10.1103/PhysRevSTAB.10.061301.

[151] S. A. Yi *et al. Physics of Plasmas.* **20,** 1, p. 013108. (2013). (See p. 45)
"Analytic model of electromagnetic fields around a plasma bubble in the blow-out regime"
DOI: 10.1063/1.4775774.

[152] A. Sävert *et al. Physical Review Letters.* **115,** 5, p. 055002. (2015). (See . 46, 72, 116, 144)
"Direct Observation of the Injection Dynamics of a Laser Wakefield Accelerator Using Few-Femtosecond Shadowgraphy"
DOI: 10.1103/PhysRevLett.115.055002.

[153] M. Geissler *et al. New Journal of Physics.* **8,** 9, pp. 186–186. (2006). (See p. 46)
"Bubble acceleration of electrons with few-cycle laser pulses"
DOI: 10.1088/1367-2630/8/9/186.

[154] D. H. Froula *et al. Physical Review Letters.* **103,** 21, pp. 1–4. (2009). (See p. 46)
"Measurements of the Critical Power for Self-Injection of Electrons in a Laser Wakefield Accelerator"
DOI: 10.1103/PhysRevLett.103.215006.

[155] I. Kostyukov *et al. Physical Review Letters.* **103,** 17, p. 175003. (2009). (See p. 46)
"Electron self-injection in multidimensional relativistic-plasma wake fields."
DOI: 10.1103/PhysRevLett.103.175003.

[156] I. Kostyukov *et al. New Journal of Physics.* **12,** (2010). (See p. 46)
"A multidimensional theory for electron trapping by a plasma wake generated in the bubble regime"
DOI: 10.1088/1367-2630/12/4/045009.

[157] A. G. R. Thomas. *Physics of Plasmas.* **17,** 5, p. 056708. (2010). (See . 46, 52, 183, 184)
"Scalings for radiation from plasma bubbles"
DOI: 10.1063/1.3368678.

[158] S. P. D. Mangles *et al. Physical Review ST-AB*, pp. 2–7. (2011). (See . 46, 117, 144)
"The self-injection threshold in laser wakefield accelerators"
DOI: 10.1103/PhysRevSTAB.15.011302.

[159] S. Kalmykov *et al. Physical Review Letters.* **103,** 13, p. 135004. (2009). (See p. 46)
"Electron Self-Injection and Trapping into an Evolving Plasma Bubble"
DOI: 10.1103/PhysRevLett.103.135004.

[160] B. A. Shadwick *et al. Physics of Plasmas.* **16,** 5, p. 056704. (2009). (See p. 47)
"Nonlinear laser energy depletion in laser-plasma accelerators"
DOI: 10.1063/1.3124185.

[161] C. D. Decker *et al. Physics of Plasmas.* **3,** 5, p. 2047. (1996). (See p. 48)
"The evolution of ultra-intense, short-pulse lasers in underdense plasmas"
DOI: 10.1063/1.872001.

[162] A. G. Khachatryan *et al. Physical Review ST-AB.* **7,** 12, pp. 1–10. (2004). (See p. 49)
"Extremely short relativistic-electron-bunch generation in the laser wakefield via novel bunch injection scheme"
DOI: 10.1103/PhysRevSTAB.7.121301.

[163] M. Chen *et al. Physics of Plasmas.* **19,** 3. (2012). (See p. 49)
"Theory of ionization-induced trapping in laser-plasma accelerators"
DOI: 10.1063/1.3689922.

[164] A. Pak *et al. Physical Review Letters.* **104,** 2, pp. 1–4. (2010). (See p. 50)
"Injection and Trapping of Tunnel-Ionized Electrons into Laser-Produced Wakes"
DOI: 10.1103/PhysRevLett.104.025003.

[165] E Oz *et al. Physical Review Letters.* **98,** 8, p. 84801. (2007). (See p. 50)
"Ionization-induced electron trapping in ultrarelativistic plasma wakes"
DOI: 10.1103/PhysRevLett.98.084801.

[166] N. Bourgeois *et al. Physical Review Letters.* **111,** 15, pp. 1–5. (2013). (See p. 50)
"Two-pulse ionization injection into quasilinear laser wakefields"
DOI: 10.1103/PhysRevLett.111.155004.

[167] L. L. Yu *et al. Physical Review Letters.* **112,** 12, pp. 1–5. (2013). (See p. 50)
"Two-color laser-ionization injection"
DOI: 10.1103/PhysRevLett.112.125001.

[168] E. Esarey *et al. Physical Review Letters.* **79,** 14, pp. 2682–2685. (1997). (See p. 50)
"Electron Injection into Plasma Wakefields by Colliding Laser Pulses"
DOI: 10.1103/PhysRevLett.79.2682.

[169] H. Kotaki *et al. Physics of Plasmas.* **11,** 6, pp. 3296–3302. (2004). (See p. 50)
"Head-on injection of a high quality electron beam by the interaction of two laser pulses"
DOI: 10.1063/1.1751171.

[170] G. Fubiani *et al. Physical Review E.* **70,** 1 2, pp. 1–12. (2004). (See p. 50)
"Beat wave injection of electrons into plasma waves using two interfering laser pulses"
DOI: 10.1103/PhysRevE.70.016402.

[171] C. Rechatin *et al. Physical Review Letters.* **102,** 16, pp. 1–4. (2009). (See p. 50)
"Controlling the phase-space volume of injected electrons in a laser-plasma accelerator"
DOI: 10.1103/PhysRevLett.102.164801.

[172] G. Fubiani *et al. Physical Review E.* **73,** 2, pp. 1–8. (2006). (See p. 50)
"Improvement of electron beam quality in optical injection schemes using negative plasma density gradients"
DOI: 10.1103/PhysRevE.73.026402.

[173] H. Suk *et al. Physical Review Letters.* **86,** 6, pp. 1011–1014. (2001). (See p. 50)
"Plasma electron trapping and acceleration in a plasma wake field using a density transition"
DOI: 10.1103/PhysRevLett.86.1011.

[174] A. V. Brantov *et al. Physics of Plasmas.* **15,** 7. (2008). (See p. 50)
"Controlled electron injection into the wake wave using plasma density inhomogeneity"
DOI: 10.1063/1.2956989.

[175] M. Hansson *et al. Physical Review ST-AB.* **18,** 7, p. 071303. (2015). (See p. 50)
"Down-ramp injection and independently controlled acceleration of electrons in a tailored laser wakefield accelerator"
DOI: 10.1103/PhysRevSTAB.18.071303.

[176] T. Y. Chien *et al. Physical Review Letters.* **94,** 11, pp. 1–4. (2005). (See p. 50)
"Spatially localized self-injection of electrons in a self-modulated laser-wakefield accelerator by using a laser-induced transient density ramp"
DOI: 10.1103/PhysRevLett.94.115003.

[177] J. Faure *et al. Physics of Plasmas.* **17,** 8. (2010). (See p. 50)
"Injection and acceleration of quasimonoenergetic relativistic electron beams using density gradients at the edges of a plasma channel"
DOI: 10.1063/1.3469581.

[178] K. Schmid *et al. Physical Review ST-AB.* **13,** 9, pp. 1–5. (2010). (See . 50, 70)
"Density-transition based electron injector for laser driven wakefield accelerators"
DOI: 10.1103/PhysRevSTAB.13.091301.

[179] M. Burza *et al. Physical Review ST-AB.* **16,** 1, p. 011301. (2013). (See . 50, 70)
"Laser wakefield acceleration using wire produced double density ramps"
DOI: 10.1103/PhysRevSTAB.16.011301.

[180] S. Bulanov *et al. Physical Review E.* **58,** 5, R5257–R5260. (1998). (See p. 51)
"Particle injection into the wave acceleration phase due to nonlinear wake wave breaking"
DOI: 10.1103/PhysRevE.58.R5257.

[181] C. Geddes *et al. Physical Review Letters.* **100,** 21, pp. 1–4. (2008). (See p. 51)
"Plasma-density-gradient injection of low absolute-momentum-spread electron bunches"
DOI: 10.1103/PhysRevLett.100.215004.

[182] J. D. Jackson, p. 841. (1998). (See . 51, 56)
Classical Electrodynamics.

[183] E. Esarey *et al. Physical Review E.* **65,** 5, p. 056505. (2002). (See p. 52)
"Synchrotron radiation from electron beams in plasma-focusing channels"
DOI: 10.1103/PhysRevE.65.056505.

[184] M. Schnell *et al. Physical Review Letters.* **108,** 7, p. 075001. (2012). (See . 53, 175)
"Deducing the Electron-Beam Diameter in a Laser-Plasma Accelerator Using X-Ray Betatron Radiation"
DOI: 10.1103/PhysRevLett.108.075001.

[185] A. Krol. (2011). (See p. 54)
In-Line Hard X-ray Holography for Biomedical Imaging.
URL: http://www.intechopen.com/books/holography-research-and-technologies.

[186] J. F. Drake. *Physics of Fluids.* **17,** 4, p. 778. (1974). (See . 55, 61)
"Parametric instabilities of electromagnetic waves in plasmas"
DOI: 10.1063/1.1694789.

[187] J. F. Drake *et al. Physical Review Letters.* **3,** 19, pp. 3–6. (1973). (See p. 55)
"Temporally Growing Raman Backscattering Instabilities in an Inhomogenous Plasma"
DOI: 10.1103/PhysRevLett.31.1197.

[188] M. N. Rosenbluth *et al. Physical Review Letters.* **31,** 19. (1973). (See p. 55)
"Temporal evolution of a three-wave parametric instability"
DOI: 10.1103/PhysRevLett.31.1190.

[189] D. W. Forslund *et al. Physics of Fluids.* **18,** 8, p. 1002. (1975). (See p. 55)
"Theory of stimulated scattering processes in laser-irradiated plasmas"
DOI: 10.1063/1.861248.

[190] A. S. Sakharov *et al. Physical Review E.* **49,** 4. (1994). (See . 56, 61, 122)
"Theory of Raman Scattering for a Short Ultrastrong Laser Pulse in a Rarefied Plasma"
DOI: 10.1103/PhysRevE.49.3274.

[191] B. Quesnel *et al. Physical Review Letters.* **78,** 11, pp. 2132–2135. (1997). (See . 57, 62)
"Electron Parametric Instabilities of Ultraintense Short Laser Pulses Propagating in Plasmas"
DOI: 10.1103/PhysRevLett.78.2132.

[192] B. Quesnel *et al. Physics of Plasmas.* **4,** p. 3358. (1997). (See . 57, 62)
"Electron parametric instabilities of ultraintense laser pulses propagating in plasmas of arbitrary density"
DOI: 10.1063/1.872494.

[193] H. Barr *et al. Physics of Plasmas.* **7,** 6, p. 2604. (2000). (See p. 59)
"Electron parametric instabilities of relativistically intense laser light in under and overdense plasma"
DOI: 10.1063/1.874102.

[194] C. J. McKinstrie *et al. Physics of Fluids B: Plasma Physics.* **4,** 8, p. 2626. (1992). (See p. 59)
"Stimulated Raman forward scattering and the relativistic modulational instability of light waves in rarefied plasma"
DOI: 10.1063/1.860178.

[195] W. L. Kruer. (1988). (See . 60, 122)
The Physics of Laser Plasma Interactions.

[196] D. W. Forslund *et al. Physical Review Letters.* **1,** 6. (1985). (See p. 60)
"Two-Dimensional Simulations of Single-Frequency and Beat-Wave Laser-Plasma Heating"
URL: http://link.aps.org/doi/10.1103/PhysRevLett.54.558.

[197] T. Matsuoka *et al. Physical Review Letters.* **105,** 3, p. 034801. (2010). (See . 60, 134, 143)
"Stimulated Raman Side Scattering in Laser Wakefield Acceleration"
DOI: 10.1103/PhysRevLett.105.034801.

[198] A. S. Sakharov *et al. Physics of Plasmas.* **4,** 9, p. 3382. (1997). (See p. 62)
"Stimulated scattering of relativistically strong radiation from an underdense plasma at high-frequency harmonics"
DOI: 10.1063/1.872479.

[199] H. Barr *et al. Physics of Plasmas.* **10,** 4, p. 1135. (2003). (See p. 62)
"Stimulated scattering and harmonic generation from relativistically intense laser light in plasma"
DOI: 10.1063/1.1559732.

[200] V. Malkin *et al. Physical Review Letters.* **82,** 22, pp. 4448–4451. (1999). (See p. 63)
"Fast Compression of Laser Beams to Highly Overcritical Powers"
DOI: 10.1103/PhysRevLett.82.4448.

[201] R. M. G. M. Trines *et al. Nature Physics.* **7,** 1, pp. 87–92. (2010). (See p. 63)
"Simulations of efficient Raman amplification into the multipetawatt regime"
DOI: 10.1038/nphys1793.

[202] R. M. G. M. Trines *et al. Physical Review Letters.* **107,** 10, p. 105002. (2011). (See p. 63)
"Production of Picosecond, Kilojoule, and Petawatt Laser Pulses via Raman Amplification of Nanosecond Pulses"
DOI: 10.1103/PhysRevLett.107.105002.

[203] V. Malkin *et al. European Physical Journal: Special Topics.* **223,** pp. 1157–1167. (2014). (See . 63, 172)
"Key plasma parameters for resonant backward Raman amplification in plasma"
DOI: 10.1140/epjst/e2014-02168-0.

[204] N. A. Yampolsky *et al. Physics of Plasmas.* **15,** 11, pp. 1–10. (2008). (See p. 63)
"Demonstration of detuning and wavebreaking effects on Raman amplification efficiency in plasma"
DOI: 10.1063/1.3023153.

[205] R. Kirkwood *et al. Physical Review Letters.* **83,** 15, pp. 2965–2968. (1999). (See p. 64)
"Observation of the Nonlinear Saturation of Langmuir Waves Driven by Ponderomotive Force in a Large Scale Plasma"
DOI: 10.1103/PhysRevLett.83.2965.

[206] N. A. Yampolsky *et al. Physics of Plasmas.* **18,** 5. (2011). (See p. 64)
"Limiting effects on laser compression by resonant backward Raman scattering in modern experiments"
DOI: 10.1063/1.3587120.

[207] D. Turnbull *et al. Physics of Plasmas.* **19,** 8. (2012). (See p. 64)
"Simultaneous stimulated Raman, Brillouin, and electron-acoustic scattering reveals a potential saturation mechanism in Raman plasma amplifiers"
DOI: 10.1063/1.4748290.

[208] V. Malkin *et al. Physics of Plasmas.* **7,** May 2015, p. 2232. (2000). (See p. 64)
"Ultra-powerful compact amplifiers for short laser pulses"
DOI: 10.1063/1.874051.

[209] D. C. Burnham *et al. Physical Review.* **188,** 2, pp. 667–675. (1969). (See p. 64)
"Coherent resonance fluorescence excited by short light pulses"
DOI: 10.1103/PhysRev.188.667.

[210] G. Vieux *et al. New Journal of Physics.* **13,** (2011). (See p. 64)
"Chirped pulse Raman amplification in plasma"
DOI: 10.1088/1367-2630/13/6/063042.

[211] G. Shvets *et al. Physics of Plasmas.* **4,** 5, p. 1872. (1997). (See . 65, 125, 155, 169)
"Analysis and simulation of Raman backscatter in underdense plasmas"
DOI: 10.1063/1.872330.

[212] G. Shvets *et al. Physical Review Letters.* **81,** 22, pp. 4879–4882. (1998). (See p. 65)
"Superradiant Amplification of an Ultrashort Laser Pulse in a Plasma by a Counterpropagating Pump"
DOI: 10.1103/PhysRevLett.81.4879.

[213] R. H. Dicke. *Physical Review.* **93,** 1, pp. 99–110. (1954). (See p. 65)
"Coherence in spontaneous radiation processes"
DOI: 10.1103/PhysRev.93.99.

[214] R. Bonifacio *et al. Physical Review Letters.* **73,** 1, pp. 70–73. (1994). (See p. 65)
"Spectrum, temporal structure, and fluctuations in a high-gain free-electron laser starting from noise"
DOI: 10.1103/PhysRevLett.73.70.

[215] B. Ersfeld *et al. Physical Review Letters.* **95,** 16, p. 165002. (2005). (See p. 65)
"Superradiant Linear Raman Amplification in Plasma Using a Chirped Pump Pulse"
DOI: 10.1103/PhysRevLett.95.165002.

[216] R. J. Briggs. (1964). (See p. 66)
Electron-Stream Interactions with Plasmas.

[217] A. Bers *et al. Physical Review Letters.* **53,** 15, pp. 1457–1460. (1984). (See p. 66)
"Relativistic Analysis of Absolute and Convective Instability Evolutions in Three Dimensions"
DOI: 10.1103/PhysRevLett.53.1457.

[218] A. Bers. (1989). (See p. 66)
Basic Plasma Physics 1.

[219] R. E. Giacone *et al. Physics of Plasmas.* **2,** 12, p. 4596. (1995). (See p. 66)
"Angular dependence of stimulated Brillouin scattering in homogeneous plasma"
DOI: 10.1063/1.871464.

[220] C. J. McKinstrie *et al. Physics of Plasmas.* **3,** 7, p. 2686. (1996). (See . 66, 167, 169, 215)
"Two-dimensional analysis of the power transfer between crossed laser beams"
DOI: 10.1063/1.871721.

[221] D. L. Bobroff *et al. Journal of Applied Physics.* **38,** 1, pp. 390–403. (1967). (See p. 67)
"Impulse response of active coupled wave systems"
DOI: 10.1063/1.1708986.

[222] V. Malkin *et al. Physical Review E.* **80,** 4, p. 046409. (2009). (See p. 67)
"Quasitransient regimes of backward Raman amplification of intense x-ray pulses"
DOI: 10.1103/PhysRevE.80.046409.

[223] A. V. Kanaev *et al. Physics of Plasmas.* **5,** 12, p. 4511. (1998). (See p. 67)
"Exact Green's function for a class of parametric instabilities"
DOI: 10.1063/1.873139.

[224] C. J. McKinstrie *et al. Physical Review E.* **50,** 3, pp. 2182–2185. (1994). (See p. 68)
"Two-dimensional stimulated Brillouin scattering"
DOI: 10.1103/PhysRevE.50.2182.

[225] S. Semushin *et al. Review of Scientific Instruments.* **72,** 7, p. 2961. (2001). (See p. 69)
"High density gas jet nozzle design for laser target production"
DOI: 10.1063/1.1380393.

[226] F Sylla *et al. Review of Scientific Instruments.* **033507,** 2012. (2013). (See p. 69)
"Development and characterization of very dense submillimetric gas jets for laser-plasma interaction"
DOI: 10.1063/1.3697859.

[227] K. Schmid, PhD Thesis. P. 202. (2009). (See p. 69)
Supersonic Micro-Jets And Their Application to Few-Cycle Laser-Driven Electron Acceleration.

[228] W. Wang *et al. Applied Physics Letters.* **103,** 24. (2013). (See p. 70)
"Control of seeding phase for a cascaded laser wakefield accelerator with gradient injection"
DOI: 10.1063/1.4842236.

[229] M. Vargas *et al. Applied Physics Letters.* **104,** 17, p. 174103. (2014). (See p. 70)
"Improvements to laser wakefield accelerated electron beam stability, divergence, and energy spread using three-dimensional printed two-stage gas cell targets"
DOI: 10.1063/1.4874981.

[230] P. K. Panigrahi *et al.* (2012). (See p. 70)
Schlieren and Shadowgraph Methods in Heat and Mass Transfer.
DOI: 10.1007/978-1-4614-4535-7.

[231] M. B. Schwab *et al. Applied Physics Letters.* **103,** 19, p. 191118. (2013). (See . 72, 116, 118)
"Few-cycle optical probe-pulse for investigation of relativistic laser-plasma interactions"
DOI: 10.1063/1.4829489.

[232] A. Buck *et al. Nature Physics.* **7,** 7, pp. 543–548. (2011). (See p. 72)
"Real-time observation of laser-driven electron acceleration"
DOI: 10.1038/nphys1942.

[233] G. R. Plateau *et al. Review of Scientific Instruments.* **81,** 3, p. 033108. (2010). (See p. 74)
"Wavefront-sensor-based electron density measurements for laser-plasma accelerators."
DOI: 10.1063/1.3360889.

[234] *Interpret.* 2015. URL: https://github.com/jasmcole/Interpret (see p. 77).

[235] M. Takeda *et al. Journal of the Optical Society of America.* **72,** 1, p. 156. (1982). (See p. 77)
"Fourier-transform method of fringe-pattern analysis for computer-based topography and interferometry"
DOI: 10.1364/JOSA.72.000156.

[236] S. G. Krantz, p. 362. (2009). (See p. 78)
Explorations in Harmonic Analysis.
DOI: 10.1007/978-0-8176-4669-1.

[237] E. Bedrosian. *Proceedings of the IEEE.* **51,** 5, pp. 868–869. (1962). (See p. 79)
"A product Theorem for Hilbert Transforms"
DOI: 10.1109/PROC.1963.2308.

[238] K. Chauhan *et al. Measurement Science and Technology.* **21,** 10, p. 105405. (2010). (See p. 81)
"Empirical mode decomposition and Hilbert transforms for analysis of oil-film interferograms"
DOI: 10.1088/0957-0233/21/10/105405.

[239] C. Heil *et al. SIAM Review.* **31,** 4, pp. 628–666. (1989). (See p. 82)
"Continuous and discrete wavelet transforms".

[240] P Tomassini *et al. Applied Optics.* **40,** 35, pp. 6561–8. (2001). (See p. 83)
"Analyzing laser plasma interferograms with a continuous wavelet transform ridge extraction technique: the method."
URL: http://www.ncbi.nlm.nih.gov/pubmed/18364963.

[241] J. Ma *et al. Applied Optics.* **50,** 16, pp. 2425–30. (2011). (See p. 83)
"Two-dimensional continuous wavelet transform for phase determination of complex interferograms."
URL: http://www.ncbi.nlm.nih.gov/pubmed/21629323.

[242] H. Ma. *Optical Engineering.* **45,** 4, p. 045601. (2006). (See p. 83)
"Advanced continuous wavelet transform algorithm for digital interferogram analysis and processing"
DOI: 10.1117/1.2188399.

[243] R. M. Goldstein *et al. Radio Science.* **23,** 4, pp. 713–720. (1988). (See p. 86)
"Satellite radar interferometry: Two-dimensional phase unwrapping"
DOI: 10.1029/RS023i004p00713.

[244] V. V. Volkov *et al. Optics Letters.* **28,** 22, pp. 2156–8. (2003). (See p. 88)
"Deterministic phase unwrapping in the presence of noise."
URL: http://www.ncbi.nlm.nih.gov/pubmed/14649926.

[245] G. Golovin *et al. Applied Optics.* **54,** 11, pp. 3491–3497. (2015). (See p. 89)
"Tomographic imaging of non-symmetric multi-component tailored supersonic flows from structured gas nozzles"
DOI: 10.1364/AO.54.003491.

[246] R. N. Bracewell. (1986). (See p. 90)
The Fourier Transform and its Applications.
DOI: 10.1049/ep.1965.0268.

[247] M. Kalal *et al. Applied Optics.* **27,** 10, pp. 1956–9. (1988). (See p. 91)
"Abel inversion using fast Fourier transforms."
URL: http://www.ncbi.nlm.nih.gov/pubmed/20531689.

[248] J. Lindström, PhD Thesis. P. 57. (2011). (See p. 93)
Radioluminescence: A simple model for fluorescent layers.

[249] Y. Glinec *et al. Review of Scientific Instruments.* **77,** 10. (2006). (See p. 94)
"Absolute calibration for a broad range single shot electron spectrometer"
DOI: 10.1063/1.2360988.

[250] A. A. Soloviev *et al. Review of Scientific Instruments.* **82,** 4. (2011). (See p. 100)
"Two-screen single-shot electron spectrometer for laser wakefield accelerated electron beams"
DOI: 10.1063/1.3585862.

[251] C. Huang *et al. Journal of Computational Physics.* **217,** 2, pp. 658–679. (2006). (See p. 101)
"QUICKPIC: A highly efficient particle-in-cell code for modeling wakefield acceleration in plasmas"
DOI: 10.1016/j.jcp.2006.01.039.

[252] T. D. Arber *et al. Plasma Physics and Controlled Fusion.* **57,** p. 113001. (2015). (See p. 101)
"Contemporary particle-in-cell approach to laser-plasma modelling"
DOI: 10.1088/0741-3335/57/11/113001.

[253] J. P Boris. *Proceedings of the Conference on the Numerical Simulation of Plasmas*, p. 757. (1971). (See p. 101)
"Proceedings of the Conference on the Numerical Simulation of Plasmas"
DOI: 10.1016/0038-1101(61)90037-5.

[254] K. Yee. *IEEE Transactions on Antennas and Propagation.* **14,** 3, pp. 302–307. (1966). (See p. 101)
"Numerical solution of initial boundary value problems involving Maxwell's equations in isotropic media"
DOI: 10.1109/TAP.1966.1138693.

[255] A. Taflove. (1995). (See p. 102)
Computational Electrodynamics: The Finite Difference Time Domain Method.

[256] A. G. R. Thomas. *Physical Review ST-AB.* **13,** p. 020702. (2010). (See p. 104)
"Algorithm for calculating spectral intensity due to charged particles in arbitrary motion"
DOI: 10.1103/PhysRevSTAB.13.020702.

[257] C. P. Ridgers *et al. Journal of Computational Physics.* **260,** pp. 273–285. (2014). (See p. 104)
"Modelling gamma-ray photon emission and pair production in high-intensity lasermatter interactions"
DOI: 10.1016/j.jcp.2013.12.007.

[258] A. Gonoskov *et al. arXiv*, pp. 1–21. (2014). (See p. 104)
"Extending PIC schemes for the study of physics in ultra-strong laser fields".

[259] B. M. Cowan *et al. Physical Review ST-AB.* **16,** 4, p. 041303. (2013). (See p. 104)
"Generalized algorithm for control of numerical dispersion in explicit time-domain electromagnetic simulations"
DOI: 10.1103/PhysRevSTAB.16.041303.

[260] R. Lehe *et al. Physical Review ST-AB.* **16,** 2, p. 021301. (2013). (See p. 104)
"Numerical growth of emittance in simulations of laser-wakefield acceleration"
DOI: 10.1103/PhysRevSTAB.16.021301.

[261] J.-L. Vay. *Physical Review Letters.* **98,** 13, pp. 2–5. (2007). (See p. 105)
"Noninvariance of Space- and Time-Scale Ranges under a Lorentz Transformation and the Implications for the Study of Relativistic Interactions"
DOI: 10.1103/PhysRevLett.98.130405.

[262] S. F. Martins *et al. Nature Physics.* **6,** 4, pp. 311–316. (2010). (See p. 105)
"Exploring laser-wakefield-accelerator regimes for near-term lasers using particle-in-cell simulation in Lorentz-boosted frames"
DOI: 10.1038/nphys1538.

[263] A. C. Kak *et al.*, p. 69. (1988). (See . 105, 189)
Principles of Computerized Tomographic Imaging.
URL: http://www.slaney.org/pct/pct-toc.html.

[264] L. A. Shepp *et al. IEEE Transactions on Nuclear Science.* **21,** 3, p. 23. (1974). (See p. 107)
"The Fourier reconstruction of a head section"
DOI: 10.1109/TNS.1974.6499235.

[265] B. Münch *et al. EMPA Activities*. **17,** 2009-2010 EMPA ACTIVITIES, pp. 34–35. (2009). (See p. 111)
"Stripe and ring artifact removal with combined wavelet-Fourier filtering"
DOI: 10.1364/OE.17.008567.

[266] I. Daubechies. (1992). (See p. 111)
Ten Lectures on Wavelets.

[267] J. P. Stonestrom *et al. IEEE transactions on bio-medical engineering*. **28,** 2, pp. 128–141. (1981). (See p. 114)
"A framework for spectral artifact corrections in x-ray CT."
DOI: 10.1109/TBME.1981.324786.

[268] R. C. Roche *et al. Journal of Experimental Marine Biology and Ecology*. **396,** 1, pp. 1–9. (2010). (See . 114, 176)
"Quantification of porosity in Acropora pulchra (Brook 1891) using X-ray micro-computed tomography techniques"
DOI: 10.1016/j.jembe.2010.10.006.

[269] G. Van Gompel *et al. Medical Physics*. **38,** S1, S36. (2011). (See p. 114)
"Iterative correction of beam hardening artifacts in CT"
DOI: 10.1118/1.3577758.

[270] M. Nisoli *et al. Applied Physics Letters*. **68,** 20, p. 2793. (1996). (See p. 116)
"Generation of high energy 10 fs pulses by a new pulse compression technique"
DOI: 10.1063/1.116609.

[271] M. C. Kaluza *et al. Physical Review Letters*. **105,** 9, p. 095003. (2010). (See p. 119)
"Observation of a Long-Wavelength Hosing Modulation of a High-Intensity Laser Pulse in Underdense Plasma"
DOI: 10.1103/PhysRevLett.105.095003.

[272] W. L. Kruer *et al. Physics of Fluids*. **23,** 15, pp. 838–841. (1969). (See . 119, 170)
"Trapped-Particle Instability"
DOI: 10.1103/PhysRevLett.23.838.

[273] K. Estabrook *et al. Physics of Fluids B: Plasma Physics*. **1,** 6, p. 1282. (1989). (See p. 119)
"Nonlinear features of stimulated Brillouin and Raman scattering"
DOI: 10.1063/1.858952.

[274] H. X. Vu *et al. Physical Review Letters*. **86,** 19, pp. 4306–4309. (2001). (See p. 119)
"Transient Enhancement and Detuning of Laser-Driven Parametric Instabilities by Particle Trapping"
DOI: 10.1103/PhysRevLett.86.4306.

[275] S. Brunner *et al. Physical Review Letters*. **93,** 14, p. 145003. (2004). (See p. 119)
"Trapped-Particle Instability Leading to Bursting in Stimulated Raman Scattering Simulations"
DOI: 10.1103/PhysRevLett.93.145003.

[276] C. Rousseaux *et al. Physical Review Letters*. **74,** 23, pp. 4655–4658. (1995). (See p. 120)
"Experimental Validation of the Linear Theory of Stimulated Raman Scattering Driven by a 500-fs Laser Pulse in a Preformed Underdense Plasma"
DOI: 10.1103/PhysRevLett.76.4649.

[277] A. Ting *et al. Optics letters*. **21,** 15, pp. 1096–1098. (1996). (See p. 120)
"Backscattered supercontinuum emission from high-intensity laser-plasma interactions."
DOI: 10.1364/OL.21.001096.

[278] C. B. Darrow *et al. Physical Review Letters*. **69,** 3. (1992). (See . 124, 126, 130)
"Strongly coupled stimulated Raman backscatter from subpicosecond laser-plasma interactions"
URL: http://link.aps.org/doi/10.1103/PhysRevLett.69.442.

[279] K. Nishikawa. *Journal of the Physical Society of Japan*. **24,** pp. 916–922. (1968). (See . 124, 167)
"Parametric Excitation of Coupled Waves I. General Formulation"
DOI: 10.1143/JPSJ.24.916.

[280] K. Nishikawa. *Journal of the Physical Society of Japan*. **24,** 5, pp. 1152–1158. (1968). (See p. 124)
"Parametric Excitation of Coupled Waves. II. Parametric Plasmon-Photon Interaction"
DOI: 10.1143/JPSJ.24.1152.

[281] P. M. Bellan. (2006). (See p. 124)
Fundamentals of Plasma Physics.

[282] M. N. Rosenbluth. *Physical Review Letters*. **29,** 9, pp. 565–567. (1972). (See p. 125)
"Parametric Instabilities in Inhomogeneous Media"
DOI: 10.1103/PhysRevLett.29.565.

[283] E. Esarey *et al.* *Physical Review A*. **45,** 8, p. 5872. (1992). (See p. 125)
"Generation of stimulated backscattered harmonic radiation from intense-laser interactions with beams and plasmas Eric"
DOI: 10.1103/PhysRevA.45.5872.

[284] A. G. R. Thomas *et al.* *Physical Review Letters*. **98,** 5, p. 054802. (2007). (See p. 133)
"Measurements of wave-breaking radiation from a laser-wakefield accelerator."
URL: http://link.aps.org/doi/10.1103/PhysRevLett.98.054802.

[285] F. S. Tsung *et al.* *Physical Review Letters*. **93,** 18, p. 185002. (2004). (See p. 140)
"Near-GeV-Energy Laser-Wakefield Acceleration of Self-Injected Electrons in a Centimeter-Scale Plasma Channel"
DOI: 10.1103/PhysRevLett.93.185002.

[286] C. Kamperidis *et al.* *Plasma Physics and Controlled Fusion*. **56,** 8, p. 084007. (2014). (See p. 140)
"Low energy spread electron beams from ionization injection in a weakly relativistic laser wakefield accelerator"
DOI: 10.1088/0741-3335/56/8/084007.

[287] L. Plaja *et al.* *Physical Review E*. **56,** 6, pp. 7142–7146. (1997). (See p. 155)
"Analytical description of a plasma diffraction grating induced by two crossed laser beams"
DOI: 10.1103/PhysRevE.56.7142.

[288] W. L. Kruer *et al.* *Physics of Plasmas*. **3,** 1, p. 382. (1996). (See . 167, 169)
"Energy transfer between crossing laser beams"
DOI: 10.1063/1.871863.

[289] P. Michel *et al.* *Physics of Plasmas*. **20,** 5. (2013). (See p. 168)
"Saturation of multi-laser beams laser-plasma instabilities from stochastic ion heating"
DOI: 10.1063/1.4802828.

[290] P. Michel *et al.* *Physical Review Letters*. **113,** 20, p. 205001. (2014). (See . 168, 212, 215)
"Dynamic Control of the Polarization of Intense Laser Beams via Optical Wave Mixing in Plasmas"
DOI: 10.1103/PhysRevLett.113.205001.

[291] J. D. Moody *et al.* *Nature Physics*. **8,** 4, pp. 344–349. (2012). (See p. 168)
"Multistep redirection by cross-beam power transfer of ultrahigh-power lasers in a plasma"
DOI: 10.1038/nphys2239.

[292] J. F. Myatt *et al.* *Physics of Plasmas*. **21,** 5, p. 055501. (2014). (See p. 168)
"Multiple-beam laserplasma interactions in inertial confinement fusion"
DOI: 10.1063/1.4878623.

[293] L. Yin *et al.* *Physics of Plasmas*. **21,** 9, p. 092707. (2014). (See p. 168)
"Stimulated scattering in laser driven fusion and high energy density physics experiments"
DOI: 10.1063/1.4895504.

[294] R. K. Kirkwood *et al.* *Physical Review Letters*. **76,** 12, pp. 2065–2068. (1996). (See p. 168)
"Observation of energy transfer between frequency-mismatched laser beams in a large-scale plasma."
URL: http://www.ncbi.nlm.nih.gov/pubmed/10060597.

[295] V. V. Eliseev *et al.* *Physics of Plasmas*. **3,** 6, p. 2215. (1996). (See p. 169)
"Interaction of crossed laser beams with plasmas"
DOI: 10.1063/1.871703.

[296] D. E. Hinkel *et al.* *Physics of Plasmas*. **1,** 9, p. 2987. (1994). (See p. 169)
"Temporal evolution of stimulated Brillouin backscatter"
DOI: 10.1063/1.870539.

[297] D. E. Hinkel *et al.* *Physics of Plasmas*. **2,** 9, p. 3447. (1995). (See p. 169)
"Stimulated Brillouin backscatter of a short-pulse laser"
DOI: 10.1063/1.870539.

[298] P. Mounaix *et al.* *Physics of Fluids B: Plasma Physics*. **5,** 9, p. 3304. (1993). (See p. 169)
"Space and time behavior of parametric instabilities for a finite pump wave duration in a bounded plasma"
DOI: 10.1063/1.860627.

[299] P. Mounaix *et al.* *Physics of Plasmas*. **1,** 8, p. 2579. (1994). (See p. 169)
"Space and time behavior of backscattering instabilities in the modified decay regime"
DOI: 10.1063/1.870585.

[300] K. Mima *et al.* *Journal of the Physical Society of Japan*. **30,** 6, pp. 1722–1733. (1971). (See p. 170)
"On the Theory of Trapped Particle Instability. I. General Formulation and Analysis for Weak Coupling Limit"
DOI: 10.1143/JPSJ.30.1722.

[301] B. I. Cohen *et al. Physics of Fluids.* **20,** 7, p. 1113. (1977). (See p. 170)
"Nonlinear plasma waves excited near resonance"
DOI: 10.1063/1.861671.

[302] S. Sepke *et al. Physical Review E.* **72,** 2, pp. 1–8. (2005). (See p. 170)
"Thomson scattering and ponderomotive intermodulation within standing laser beat waves in plasma"
DOI: 10.1103/PhysRevE.72.026501.

[303] S. B. Dabagov *et al. Physical Review ST-AB.* **18,** 6, p. 064002. (2015). (See p. 170)
"Channeling of electrons in a crossed laser field"
DOI: 10.1103/PhysRevSTAB.18.064002.

[304] P. Balcou. *European Physical Journal D.* **59,** 3, pp. 525–537. (2010). (See p. 170)
"Proposal for a Raman X-ray free electron laser"
DOI: 10.1140/epjd/e2010-00185-5.

[305] I. A. Andriyash *et al. Physical Review Letters.* **109,** 24, pp. 1–5. (2012). (See p. 170)
"X-ray amplification from a Raman free-electron laser"
DOI: 10.1103/PhysRevLett.109.244802.

[306] I. A. Andriyash *et al. Physical Review ST-AB.* **18,** 5, p. 050704. (2015). (See p. 170)
"Amplified short-wavelength light scattered by relativistic electrons in the laser-induced optical lattice"
DOI: 10.1103/PhysRevSTAB.18.050704.

[307] E Khain *et al. Physical Review E.* **64,** 3, p. 036619. (2001). (See p. 170)
"Parametric autoresonance."
DOI: 10.1103/PhysRevE.64.036619.

[308] T. Chapman *et al. Physical Review Letters.* **108,** 14, pp. 1–5. (2012). (See p. 170)
"Driven spatially autoresonant stimulated raman scattering in the kinetic regime"
DOI: 10.1103/PhysRevLett.108.145003.

[309] Y. Ping *et al. Physical Review E.* **66,** 4, pp. 1–6. (2002). (See p. 171)
"Raman amplification of ultrashort laser pulses in microcapillary plasmas"
DOI: 10.1103/PhysRevE.66.046401.

[310] Y. Ping *et al. Physical Review E.* **67,** 1, p. 016401. (2003). (See p. 171)
"Raman backscattering and amplification in a gas jet plasma."
DOI: 10.1103/PhysRevE.67.016401.

[311] Y. Ping *et al. Physical Review Letters.* **92,** 17, p. 175007. (2004). (See p. 171)
"Amplification of Ultrashort Laser Pulses by a Resonant Raman Scheme in a Gas-Jet Plasma"
DOI: 10.1103/PhysRevLett.92.175007.

[312] W. Cheng *et al. Physical Review Letters.* **94,** 4, p. 045003. (2005). (See p. 171)
"Reaching the Nonlinear Regime of Raman Amplification of Ultrashort Laser Pulses"
DOI: 10.1103/PhysRevLett.94.045003.

[313] J. Ren *et al. Nature Physics.* **3,** 10, pp. 732–736. (2007). (See . 171, 172)
"A new method for generating ultraintense and ultrashort laser pulses"
DOI: 10.1038/nphys717.

[314] R. K. Kirkwood *et al. Physics of Plasmas.* **14,** 11. (2007). (See p. 171)
"Amplification of an ultrashort pulse laser by stimulated Raman scattering of a 1 ns pulse in a low density plasma"
DOI: 10.1063/1.2804083.

[315] Y. Ping *et al. Physics of Plasmas.* **16,** 12. (2009). (See p. 171)
"Development of a nanosecond-laser-pumped Raman amplifier for short laser pulses in plasma"
DOI: 10.1063/1.3276739.

[316] C.-H. Pai *et al. Physical Review Letters.* **101,** 6, p. 065005. (2008). (See p. 171)
"Backward Raman Amplification in a Plasma Waveguide"
DOI: 10.1103/PhysRevLett.101.065005.

[317] M. Dreher *et al. Physical Review Letters.* **93,** 9, pp. 1–4. (2004). (See p. 171)
"Observation of superradiant amplification of ultrashort laser pulses in a plasma"
DOI: 10.1103/PhysRevLett.93.095001.

[318] A. A. Balakin *et al. Journal of Experimental and Theoretical Physics Letters.* **80,** 1, pp. 12–16. (2004). (See p. 171)
"Laser pulse amplification upon Raman backscattering in plasma produced in dielectric capillaries"
DOI: 10.1134/1.1800205.

[319] J. M. Cole *et al.* *Plasma Physics and Controlled Fusion.* **58,** 1, p. 014008. (2016). (See p. 175)
"Tomography of human trabecular bone with a laser-wakefield driven x-ray source"
DOI: 10.1088/0741-3335/58/1/014008.

[320] I. Manke *et al.* *Applied Physics Letters.* **90,** 21, p. 214102. (2007). (See p. 176)
"In situ investigation of the discharge of alkaline ZnMnO2 batteries with synchrotron x-ray and neutron tomographies"
DOI: 10.1063/1.2742283.

[321] P. J. Schilling *et al.* *Composites Science and Technology.* **65,** 14, pp. 2071–2078. (2005). (See p. 176)
"X-ray computed microtomography of internal damage in fiber reinforced polymer matrix composites"
DOI: 10.1016/j.compscitech.2005.05.014.

[322] P. C. J. Donoghue *et al.* *Nature.* **442,** 7103, pp. 680–3. (2006). (See p. 176)
"Synchrotron X-ray tomographic microscopy of fossil embryos."
DOI: 10.1038/nature04890.

[323] F Mees *et al.* *The Geological Society of London.* **215,** pp. 1 –6. (2003). (See p. 176)
"Applications of X-ray computed tomography in the geosciences"
DOI: 10.1144/GSL.SP.2003.215.01.01.

[324] G. N. Hounsfield. *The British Journal of Radiology.* **46,** 552, pp. 1016–1022. (1973). (See p. 176)
"Computerized transverse axial scanning (tomography): Part I. Description of system."
DOI: 10.1259/0007-1285-46-552-1016.

[325] S. W. Blume *et al.* *Osteoporosis International.* **22,** 6, pp. 1835–1844. (2011). (See p. 176)
"Medical costs of osteoporosis in the elderly Medicare population"
DOI: 10.1007/s00198-010-1419-7.

[326] B. D. Snyder *et al.* *Calcified Tissue International.* **53,** pp. 14–22. (1993). (See p. 176)
"Role of Trabecular Morphology in the Etiology of Age-Related Vertebral Fractures"
URL: http://link.springer.com/article/10.1007/BF01673396.

[327] M. J. Ciarelli *et al.* *Journal of Orthopaedic Research.* **9,** 5, pp. 674–682. (1991). (See p. 176)
"Evaluation of orthogonal mechanical properties and density of human trabecular bone from the major metaphyseal regions with materials testing and computed tomography"
DOI: 10.1002/jor.1100090507.

[328] J. Hazrati Marangalou *et al.* *Osteoporosis International.* **25,** 4, pp. 1285–1296. (2014). (See p. 176)
"Locally measured microstructural parameters are better associated with vertebral strength than whole bone density"
DOI: 10.1007/s00198-013-2591-3.

[329] M. Kleerekoper *et al.* *Calcified Tissue International.* **37,** 6, pp. 594–7. (1985). (See p. 176)
"The role of three-dimensional trabecular microstructure in the pathogenesis of vertebral compression fractures."
URL: http://www.ncbi.nlm.nih.gov/pubmed/3937580.

[330] L. D. Carbonare *et al.* *Journal of Endocrinological Investigation.* **27,** 1, pp. 99–105. (2004). (See p. 176)
"Bone microarchitecture as an important determinant of bone strength"
DOI: 10.1007/BF03350919.

[331] B. Borah *et al.* *Bone.* **37,** 1, pp. 1–9. (2005). (See p. 177)
"The effect of risedronate on bone mineralization as measured by micro-computed tomography with synchrotron radiation: correlation to histomorphometric indices of turnover."
DOI: 10.1016/j.bone.2005.03.017.

[332] P. J. Thurner *et al.* *Bone.* **39,** 2, pp. 289–99. (2006). (See p. 177)
"Time-lapsed investigation of three-dimensional failure and damage accumulation in trabecular bone using synchrotron light."
DOI: 10.1016/j.bone.2006.01.147.

[333] A. Larrue *et al.* *PloS one.* **6,** 7, e21297. (2011). (See . 177, 202)
"Synchrotron radiation micro-CT at the micrometer scale for the analysis of the three-dimensional morphology of microcracks in human trabecular bone."
DOI: 10.1371/journal.pone.0021297.

[334] S. Nuzzo *et al.* *Medical Physics.* **29,** 11, p. 2672. (2002). (See p. 177)
"Quantification of the degree of mineralization of bone in three dimensions using synchrotron radiation microtomography"
DOI: 10.1118/1.1513161.

[335] M. L. Bouxsein. *Osteoporosis International*. **14,** 5, S118–S127. (2003). (See p. 177)
"Bone quality: where do we go from here?"
DOI: 10.1007/s00198-003-1489-x.

[336] A Odgaard *et al*. *Bone*. **14,** 2, pp. 173–82. (1993). (See p. 177)
"Quantification of connectivity in cancellous bone, with special emphasis on 3-D reconstructions."
URL: http://www.ncbi.nlm.nih.gov/pubmed/8334036.

[337] J. C. Elliott *et al*. *Journal of Microscopy*. **126,** 2, pp. 211–213. (1982). (See p. 177)
"X-ray microtomography"
DOI: 10.1111/j.1365-2818.1982.tb00376.x.

[338] B. P. Flannery *et al*. *Science*. **237,** 4821, pp. 1439–1444. (1987). (See p. 177)
"Three-Dimensional X-ray Microtomography."
DOI: 10.1126/science.237.4821.1439.

[339] E. L. Ritman. *Annual Review of Biomedical Engineering*. **13,** pp. 531–52. (2011). (See p. 177)
"Current status of developments and applications of micro-CT."
DOI: 10.1146/annurev-bioeng-071910-124717.

[340] C. Chappard *et al*. *Osteoarthritis and Cartilage*. **14,** 3, pp. 215–223. (2006). (See p. 177)
"Subchondral bone micro-architectural alterations in osteoarthritis: a synchrotron micro-computed tomography study."
DOI: 10.1016/j.joca.2005.09.008.

[341] M Green *et al*. *Proceedings of the Physical Society*. **1206,** (1961). (See p. 177)
"The Efficiency of Production of Characteristic X-radiation in Thick Targets of a Pure Element"
URL: http://iopscience.iop.org/0370-1328/78/6/315.

[342] D. E. Grider *et al*. *Journal of Physics D: Applied Physics*. **2281,** 19. (1986). (See . 178, 193)
"Electron beam melting in microfocus x-ray tubes"
URL: http://iopscience.iop.org/0022-3727/19/12/008.

[343] A. Taylor. *Journal of Scientific Instruments*. **26,** 7, pp. 225–229. (1949). (See p. 178)
"A 5 kW. Crystallographic X-ray Tube with a Rotating Anode"
DOI: 10.1088/0950-7671/26/7/301.

[344] O. Hemberg *et al*. *Applied Physics Letters*. **83,** 7, pp. 1483–1485. (2003). (See p. 178)
"Liquid-metal-jet anode electron-impact x-ray source"
DOI: 10.1063/1.1602157.

[345] F. Peyrin *et al*. *Technology and Health Care*. **6,** 5-6, pp. 391–401. (1998). (See p. 178)
"Micro-CT examinations of trabecular bone samples at different resolutions: 14, 7 and 2 micron level."
URL: http://europepmc.org/abstract/MED/10100941.

[346] M. Salome *et al*. *Medical Physics*. **26,** 10, p. 2194. (1999). (See p. 178)
"A synchrotron radiation microtomography system for the analysis of trabecular bone samples"
DOI: 10.1118/1.598736.

[347] R. Mokso *et al*. *AIP Conference Proceedings*. **1365,** pp. 38–41. (2010). (See p. 178)
"Following dynamic processes by X-ray tomographic microscopy with sub-second temporal resolution"
DOI: 10.1063/1.3625299.

[348] W.-H. Kim *et al*. *Korean Circulation Journal*. **38,** 9, p. 462. (2008). (See p. 178)
"Synchrotron Microangiography of the Rat Heart Using the Langendorff Model"
DOI: 10.4070/kcj.2008.38.9.462.

[349] E. Eggl *et al*. *Proceedings of the National Academy of Sciences*. **112,** 18, p. 201500938. (2015). (See p. 178)
"X-ray phase-contrast tomography with a compact laser-driven synchrotron source"
DOI: 10.1073/pnas.1500938112.

[350] C. M. Laperle *et al*. *Applied Physics Letters*. **91,** 17, pp. 1–4. (2007). (See p. 178)
"Propagation based differential phase contrast imaging and tomography of murine tissue with a laser plasma x-ray source"
DOI: 10.1063/1.2802728.

[351] S. Corde *et al*. *Plasma Physics and Controlled Fusion*. **54,** 12, p. 124023. (2012). (See p. 181)
"Betatron emission as a diagnostic for injection and acceleration mechanisms in laser plasma accelerators"
DOI: 10.1088/0741-3335/54/12/124023.

[352] S. Corde *et al. Nature Communications.* **4,** p. 1501. (2013). (See p. 184)
"Observation of longitudinal and transverse self-injections in laser-plasma accelerators."
DOI: 10.1038/ncomms2528.

[353] B. Hesse *et al. Journal of Bone and Mineral Research.* **30,** 2, pp. 346–356. (2015). (See . 190, 202)
"Canalicular Network Morphology Is the Major Determinant of the Spatial Distribution of Mass Density in Human Bone Tissue: Evidence by Means of Synchrotron Radiation Phase-Contrast nano-CT"
DOI: 10.1002/jbmr.2324.

[354] L. A. Feldkamp *et al. Journal of the Optical Society of America A.* **1,** 6, p. 612. (1984). (See p. 190)
"Practical cone-beam algorithm"
DOI: 10.1364/JOSAA.1.000612.

[355] J. Wu *et al. Optik.* **124,** 9, pp. 864–866. (2013). (See p. 191)
"A phase retrieval algorithm for X-ray phase contrast imaging"
DOI: 10.1016/j.ijleo.2012.02.030.

[356] J. Sauvola *et al. Pattern Recognition.* **33,** pp. 225–236. (2000). (See p. 191)
"Adaptive document image binarization"
URL: http://www.sciencedirect.com/science/article/pii/S0031320399000552.

[357] *Blender.org.* URL: http://www.blender.org (see p. 191).

[358] *Drishti.* URL: https://github.com/AjayLimaye/drishti (see p. 191).

[359] B Verman *et al. The Rigaku Journal.* **19,** 1, pp. 4–13. (2002). (See p. 193)
"Microfocusing source and multilayer optics based x-ray diffraction systems"
URL: http://www.rigaku.com/downloads/journal/Vol19.1.2002/verman.pdf.

[360] T. Tuohimaa, PhD Thesis. (2008). (See p. 193)
Liquid-Jet-Target Microfocus X-Ray Sources: Electron Guns, Optics and Phase-Contrast Imaging.

[361] G. Poludniowski *et al. Physics in Medicine and Biology.* **54,** 19, N433–8. (2009). (See p. 193)
"SpekCalc: a program to calculate photon spectra from tungsten anode x-ray tubes."
DOI: 10.1088/0031-9155/54/19/N01.

[362] *SpekCalc.* URL: http://spekcalc.weebly.com/ (see p. 193).

[363] W. A. Kalender. *Physics in Medicine and Biology.* **59,** 3, R129–R150. (2014). (See p. 193)
"Dose in x-ray computed tomography"
DOI: 10.1088/0031-9155/59/3/R129.

[364] P. E. Palacio-Mancheno *et al. Journal of Bone and Mineral Research.* **29,** 1, pp. 142–150. (2014). (See p. 193)
"3D Assessment of Cortical Bone Porosity and Tissue Mineral Density Using High-Resolution CT: Effects of Resolution and Threshold Method"
DOI: 10.1002/jbmr.2012.

[365] S. Banerjee *et al. Optics Letters.* **37,** 12, p. 2175. (2012). (See p. 193)
"High-efficiency 10 J diode pumped cryogenic gas cooled Yb:YAG multislab amplifier"
DOI: 10.1364/OL.37.002175.

[366] S. M. Hooker *et al. Journal of Physics B.* **47,** 23, p. 234003. (2014). (See p. 194)
"Multi-pulse laser wakefield acceleration: a new route to efficient, high-repetition-rate plasma accelerators and high flux radiation sources"
DOI: 10.1088/0953-4075/47/23/234003.

[367] S. M. Walker *et al. PLoS Biology.* **12,** 3, e1001823. (2014). (See p. 194)
"In Vivo Time-Resolved Microtomography Reveals the Mechanics of the Blowfly Flight Motor"
DOI: 10.1371/journal.pbio.1001823.

[368] M. J. Flynn *et al. Nuclear Instruments and Methods in Physics Research.* **353,** 1-3, pp. 312–315. (1994). (See p. 194)
"Microfocus X-ray sources for 3D microtomography"
DOI: 10.1016/0168-9002(94)91664-0.

[369] *NRecon.* URL: http://bruker-microct.com/products/downloads.htm (see p. 196).

[370] P. Roschger *et al. Current Osteoporosis Reports.* **12,** 3, pp. 338–350. (2014). (See p. 202)
"Changes in the degree of mineralization with osteoporosis and its treatment"
DOI: 10.1007/s11914-014-0218-z.

[371] L. Ritschl *et al.* *IEEE Nuclear Science Symposium Conference Record.* **1545,** pp. 3231–3240. (2010). (See p. 202)
"Improved sparsity-constrained image reconstruction applied to clinical CT data"
DOI: 10.1109/NSSMIC.2010.5874402.

[372] C. J. McKinstrie *et al.* *Physics of Plasmas.* **6,** 2, pp. 463–466. (1999). (See p. 208)
"Accurate formulas for the Landau damping rates of electrostatic waves"
DOI: 10.1063/1.873212.

[373] P. Michel *et al.* *Physics of Plasmas.* **20,** 5. (2013). (See p. 212)
"Saturation of multi-laser beams laser-plasma instabilities from stochastic ion heating"
DOI: 10.1063/1.4802828.