Laser-Accelerated Ions and Their Applications



Ben Spiers

Lady Margaret Hall

University of Oxford

A thesis submitted for the degree of

Doctor of Philosophy

Michaelmas 2021

Dedicated to
two and a half thousand cups
of instant coffee

Acknowledgments

There are many people I would like to thank for their support during the course of my DPhil. Foremost, my undergraduate tutor and doctoral supervisor Peter Norreys has given me support, guidance, and many opportunities over the years, which have directly led to all the work I present in this thesis. Thank you for always encouraging me to have faith in my own abilities and work, and to always keep going until I can see for myself that the results are good.

The other members of Peter's research group have been a vital part of my time here as a postgraduate student. The company of Rami, Alex, Marko, Rob, and—when we can find him—Robin[1] has made the long hours and difficult times a lot easier than they would have been otherwise. Certain global events have unfortunately prevented me from getting to know Qingsong, Robin[2] and Iustin and well as I should do over a year after they all joined the group, but I have enjoyed what little time I have managed to spend with them all greatly. I wish you all the greatest success in the future; from what I've seen there's going to be a great deal of exciting science coming out of the Norreys group over the coming years.

Speaking of the conspicuous global events that have formed the background against which this thesis was written, I must thank my housemates Jake, Tom, Jess and Anna for making the years we've lived together so great, and making the successive lockdowns of the past couple of years a lot less hard than they could have been. I know many people found it very difficult suddenly working at home with those they live with, but in our case I actually quite enjoyed it. Regular virtual[3] pub trips with Rami, Rob and Aimee also made the darkest points of lockdown feel a lot brighter; thank you all for being such great friends.

A special mention is due to Oli, who despite having no scientific background volunteered to take on the onerous task of proofreadign all 150-odd pages of this thesis. He's responsible for a lot of the good grammar, sensible sentence constructions, and the majority of the Oxford commas[4] in this document. Tholi.

My parents are in many ways directly and indirectly responsible for my career in science at all. In particular, Dad, thank you for nurturing my interest in Physics

[1] Wang
[2] Timmis
[3] and then physical, then virtual again, then physical again
[4] except that one

in the first place from such a young age. Thank you both for always unflinchingly supporting me in all my endeavours, and for all the love and support you've given me over all the years of my life.

Most importantly, I'd like to thank Emily. For the seven and a half years we've been together, for the four and a half we spent living in different cities, for being there for me whenever I've needed it and for always trying to make me believe in myself more than really feels natural to me. Thank you for supporting me through the past few weeks of frantic writing, and the preceeding months of slightly less-frantic writing. Love you. U+2764[5]

Institutional & Funding

This work has been carried out within the framework of the EUROfusion Consortium and has received funding from the Euratom research and training programme 2014-2018 and 2019-2020 under grant agreement Nº 633053. The views and opinions expressed herein do not necessarily reflect those of the European Commission. I also acknowledge the support of STFC grant ST/P002048/1 and EPSRC grants EP/R029148/1 and EP/L000237/1. I would like to thank UKRI-EPSRC and AWE plc for financially supporting my DPhil research under iCASE studentship 1950656. Orion experimental data are © British Crown Owned Copyright 2021/AWE

[5]Unfortunately LaTeXwon't print the red heart emoji, so the Unicode codepoint will have to suffice

Abstract

Since the first studies of target-normal sheath acceleration around the year 2000, production of energetic ion beams from intense laser-matter interactions has been investigated in great depth. New acceleration mechanisms such as radiation pressure acceleration, collisionless shock acceleration, the hybrid breakout afterburner mechanism, and special regimes such as light-sail radiation pressure acceleration have been discovered and demonstrated. Applications of ion beams produced by these mechanisms range from common diagnostic techniques for laser plasma interactions to industrial radiography and even—albeit needing significant improvements to predictability and reproducibility before this can be realised—their proposed use in medical treatments like hadron therapy, which if successful could significantly reduce the cost to hospitals of hosting such therapies. I present work on three topics around laser-accelerated ion beams and their applications:

Firstly, I present an experiment which sought to demonstrate acceleration of ions from a laser-driven collisionless shock within a plasma produced from low-density plastic foam. This acceleration mechanism produces ions with a narrow energy spectrum, compared to the more well-established target-normal sheath acceleration. Despite limitations to the accessible experimental conditions, I present evidence that ions were accelerated with the narrow energy spectrum characteristic of this mechanism.

Secondly, I introduce the reader to an important application of fast ion beams: the proton radiography diagnostic. This diagnostic is used to measure the electromagnetic fields present in a plasma through their influence on a beam of protons as it traverses the plasma. I explain the theory of proton radiography and describe a laser-plasma channelling experiment on which proton radiography was used to probe the magnetic fields associated with a laser-produced plasma channel. I demonstrate the analysis of some proton radiographs using an improved solver I have implemented in a community analysis package.

Finally, I propose improving the information accessible to the proton radiography diagnostic by making observations of the plasma from several different directions, a procedure I term 'proton tomgraphy'. I then derive methods of improving the quality of data recovered from proton tomography. Compared to standard tomographic algorithms I achieve an orders-of-magnitude improvement in reconstrution quality of synthetic data when very few lines of sight are available.

Contents

1 Introduction **1**

1.1 Introduction to Plasma and Diagnostics 2
1.2 Thesis Outline 6
1.3 Author Contributions 8
1.4 Publications 9

2 Basic Laser and Plasma Theory **11**

2.1 Electromagnetism 12
2.1.1 Electromagnetic Waves 14
2.2 Units 16
2.3 Plasma Dynamics 18
2.3.1 The Vlasov-Maxwell System 18
2.3.2 Fluid Models 20
2.4 Waves in Plasma 27
2.5 Laser-Plasma Processes 32
2.5.1 Forming a Plasma 32
2.5.2 Ionisation 33
2.5.3 Energy Absorption 35
2.5.4 The Ponderomotive Effect 40

3 Numerical Modelling of Plasmas **43**

3.1 Large-Scale Simulations 43
3.1.1 Vlasov Simulations 44
3.1.2 Particle-in-Cell Simulations 46
3.1.3 Fluid Model Simulations 54
3.2 Semi-Analytic and Smaller-Scale Modelling 55

4 Ion Acceleration **57**

4.1 Target Normal Sheath Acceleration 58
4.1.1 Modelling TNSA 58
4.1.2 TNSA-produced beams 61
4.2 Additional Acceleration Mechanisms 62

4.2.1 Radiation Pressure Acceleration . . . 62
4.2.2 Collisionless Shock Acceleration . . . 65
4.3 Experiment: Collisionless Shock Acceleration in Foam Targets . . . 81
4.3.1 Experimental Conditions . . . 81
4.4 Simulation Post-Mortem . . . 85
4.5 Summary . . . 88

5 Proton Radiography 91
5.1 How It Works . . . 92
5.1.1 Proton Sources . . . 93
5.1.2 Detection . . . 94
5.1.3 Formation of a Radiograph . . . 95
5.1.4 The Confounding Effect of Electric Fields . . . 100
5.2 The Inverse Problem . . . 101
5.3 A New Solver for `PROBLEM` . . . 103
5.4 Experiment: Whole-Beam Self-Focusing at Orion . . . 105
5.4.1 Long-Pulse Interaction . . . 108
5.4.2 Short-Pulse Interaction . . . 112
5.5 Summary . . . 117

6 Proton Tomography 119
6.1 Tomography Basics . . . 120
6.1.1 Analytic Methods . . . 123
6.1.2 Algebraic Reconstruction Techniques . . . 126
6.2 Applying Tomographic Theory to Proton Radiography . . . 128
6.3 Improving Performance . . . 129
6.3.1 Hankel Reconstruction and Fourier Interpolation . . . 131
6.3.2 Aspect Ratio Compensation . . . 134
6.4 Results . . . 137
6.5 Summary . . . 138

7 Summary and Future Work 145
7.1 Summary of Work . . . 145
7.2 Future Work . . . 147

Appendices

A Producing Perturbed Density Maps **153**
A.1 The Problem . . . 154
A.2 Stochastic Methods . . . 155
A.3 Deterministic Methods . . . 157
A.3.1 Rectangular cells . . . 157
A.3.2 Triangular Cells . . . 159
A.3.3 Results . . . 161
A.4 Conclusions . . . 164

B Modelling Laser Pulse Propagation **165**
B.1 Problem Statement . . . 165
B.2 Collimated Propagation . . . 166
B.3 Focusing . . . 167
B.4 Focusing Beam Propagation . . . 173
B.4.1 Discrete Implementation . . . 178
B.5 Results . . . 179

C Drifting Boltzmann Relation **183**

References **185**

1 Introduction

Contents

1.1 Introduction to Plasma and Diagnostics 2
1.2 Thesis Outline 6
1.3 Author Contributions 8
1.4 Publications 9

In the 21st Century it's probably reasonable to describe the fact that solid, liquid and gas are not the only states of matter in the universe as common popular science knowledge. While the *weirdness* of the non-classical states—Bose-Einstein condensates, superconductors and superfluids to name just a few—can easily capture the imagination, the plasma state (often thought of simply as electrically conductive, ionized gas) exhibits a huge and fascinating diversity of phenomena.

Plasma is the most abundant form of matter in the Universe[1], comprising stars as well as the interstellar medium [1]. Terrestrially, plasmas are less common than they are in the cosmos, although they are still recognisable in lightning, plasma televisions, neon signs, the aurora borealis, and the everyday occurrence of high-power laser experiments.

From the latter half of the 20th Century, study of the plasma state has taken

[1] Here I am restricting myself to forms of matter understood by modern physics

on extra significance. Both major themes of research into controlled nuclear fusion take place in plasma (they differ by whether the plasma is confined inertially [2] or magnetically [3]), by necessity of the extreme conditions necessary to ignite a fusion reaction. Understanding the behaviour of plasma in these conditions is of vital importance in the search for fusion energy, and indeed the difficulty to date in achieving ignition has almost always been the result of incomplete understanding of the processes taking place within the plasma [4, 5]. Improvements in theoretical models and numerical simulation capability have worked in tandem to bring the world tantalisingly close to ignition, with construction of the magnetic-confinement tokamak experiment ITER underway in France [6] and recent results from the inertial-confinement National Ignition Facility in Livermore, California very nearly achieving ignition [7, 8].

While the work I present in this thesis does not represent direct progress towards successful nuclear fusion, it takes place against the background of continued efforts in laser-driven fusion research exemplified by the National Ignition Facility. My work has centred around the acceleration of ions from laser-driven plasma systems and applications of the resulting ion beams. These beams of energetic ions (as investigated in Chapter 4) have many potential uses, ranging from a source of heating in the so-called 'fast ignition' approach to inertially-confined fusion [9, 10], to possible uses in medical ion therapy treatments [11], given sufficient refinement. The application I have focused on in my research, however, is the use of ion beams (specifically protons) to probe the electromagnetic fields present in other laser-plasma interactions [12–14]. Present methods for proton beam probing of magnetic fields in a plasma (proton radiography) are demonstrated in Chapter 5, and the application of tomographic methods to extract more data from proton radiography is explored in Chapter 6.

1.1 Introduction to Plasma and Diagnostics

In the above introduction I have made frequent reference to plasma, but have so far only provided a very brief (and wholly incomplete) description of what a plasma *is*.

While plasma is popularly dismissed as 'ionized gas', this description belies the full complexity of the state quite significantly [1]. As pointed out by Vlasov [15], "for a system of charged particles the kinetic equation method which considers only binary interactions—interactions through collisions—is an approximation which is strictly speaking inadequate, so that in the theory of such systems an essential role must be played by the interaction forces, particularly at large distances and, hence, a system of charged particles is, in essence, not a gas but a distinctive system coupled by long-range forces". While plasma is, indeed, composed of at least two species of unbound, charged particles[2], the presence of long-range interactions between these charge carriers (mediated by the electric and magnetic fields) endow plasma with vastly different properties and dynamics to neutral gases.

A first important property of plasma is that of Debye shielding. It might be natural to assume that, in a form of matter composed of unbound charges, large electric field gradients would form naturally and over large distances. However, the opposite is true. Plasma actually evolves in such a way as to maintain neutrality, when measured over length scales larger than the 'Debye length':

$$\lambda_D = \sqrt{\frac{\epsilon_0 k_B T_e}{n_e e^2}}. \tag{1.1}$$

Here ϵ_0 is the permittivity of free space, k_B the Boltzmann constant, and e is the fundamental unit of charge. n_e is the electron density of the plasma. Probably its most important property, a plasma's electron density determines many of its other properties, from its natural frequency of oscillation to its response to different frequencies of light. Another important property of a plasma is its temperature[3] T_e, which enters into the Debye length and also into the thermal speed of electrons in the plasma. Electrostatic potentials in the plasma are 'screened' with an exponential decay radius equal to this Debye length [16]. Conceptually, because of the high electrical conductivity of plasma, electrons flow to rapidly neutralise any charge imbalance that occurs.

[2]These species are typically positively-charged ions and negatively-charged electrons, though more exotic plasmas exist, which may, for example, contain positrons in place of ions

[3]Usually quoted in energy units (such as eV, equivalent to 11 605 K), incorporating the Boltzmann constant k_B into the 'temperature' $k_B T_e$

There are also a multitude of oscillatory modes specific to plasma. For instance, electrons flowing to neutralise charge imbalance as described above will overshoot, setting up an oscillation with frequency given by the electron plasma frequency:

$$\omega_p = \sqrt{\frac{n_e e^2}{m_e \epsilon_0}} \tag{1.2}$$

These two phenomena—Debye shielding and Langmuir oscillations—exemplify the capacity of plasma to support collective behaviour, but are only the tip of the iceberg in terms of the diversity of plasma dynamics [16]. Unlike electric fields, magnetic fields do not experience damping in plasma, and the influence of strong magnetic fields on plasma dynamics can be dramatic. An important model for the description of magnetised plasma is magnetohydrodynamics (MHD) [17]. The MHD model generates several wave modes which are intimately linked to the magnetic field permeating a magnetised plasma. This model and the additional wave modes it supports are derived and described in some detail in Chapter 2, along with some other important models of plasma dynamics.

I will now briefly describe some approaches to measuring plasma properties experimentally. I will describe diagnostics used to measure electron density, temperature, and magnetic fields.

Electron Density

Commonly, electron density (or equivalently, plasma frequency) is measured indirectly through its influence on the plasma's index of refraction [18, 19]. Light of frequency ω_0 propagating in plasma with electron plasma frequency ω_p experiences a refractive index given by:

$$\eta = \sqrt{1 - \left(\frac{\omega_p}{\omega_0}\right)^2}. \tag{1.3}$$

If a wide laser beam propagates through a non-uniform plasma, different parts of the beam acquire different phase shifts due to fluctuations in refractive index. In low-density plasma ($\omega_p \ll \omega_0$), this phase shift can be measured by placing the plasma in one arm of an interferometer. At higher densities ($\omega_p \approx \omega_0$) the

effects of refraction become significant [18], so it is important to use imaging optics to compensate for the effects of refraction.

The (path-integrated) refractive index of a plasma (and therefore its plasma frequency and electron density) can be determined from the interferograms that are produced.

At high densities other methods such as shadowgraphy [20, 21] or angular filter refractometry [22, 23] can also be used to determine the plasma's density profile.

Electron Temperature

The electron temperature of laser-produced plasma can often reach the keV range, at which point the black-body radiation emitted by the electron population contains a detectable intensity of x-rays. The x-ray photon energies emitted by a plasma can therefore be used to measure the temperature of sufficiently hot plasma [24, 25]. Bragg reflection from a crystal is used to separate the x-rays by wavelength, and the x-ray spectrum is then used to determine the electron temperature of the sample. A more advanced variation of this diagnostic (the 'von Hamos' geometry) uses a curved crystal to image x-rays on the detector plate, producing a degree of spatial resolution in the temperature measurement [25].

Magnetic Fields

Finally, magnetic fields are often measured using the Faraday effect [26]. Sometimes referred to as "Faraday rotation", left- and right-circular polarisation components of light experience different refractive indices when travelling parallel to a magnetic field; the difference is proportional to both the magnitude of the longitudinal magnetic field component and the plasma's electron density. A phase difference accumulates between the two circular components of plane-polarised light, causing the direction of polarisation to rotate. With a known input direction of polarisation and for a known electron density, the measured output polarisation can be correlated directly to the magnetic field strength.

Magnetic fields can also be measured by their effect on a beam of charged particles which is fired through the plasma [27, 28]. The sensitivity of the diagnostic

is dependent on the ingoing particles' charge-to-momentum ratio; laser-accelerated protons are used most commonly as the probing species. Production of such proton beams is discussed in Chapter 4, and specifics of this diagnostic (termed proton radiography) are covered in detail in Chapter 5. Electric fields can also deflect the protons' trajectories, so if they are present in the plasma (and of comparable magnitude to the magnetic fields) the two effects must be decoupled. This problem has been explored in some depth in the literature, and is discussed in §5.1.4

In high-density plasma, both of these techniques run into difficulties. Plasma of high electron density will become opaque to the light used for a Faraday rotation diagnostic, and at very high density proton radiography probes will be scattered too much to be useful.

1.2 Thesis Outline

In this thesis, I work through a progression of topics which ultimately lead to the proposed application of tomographic methods (those involving probes directed along many lines of sight) to proton radiography and some techniques to improve the performance of this proposed 'proton tomography'. Following this introductory chapter, I will introduce the reader to some basic models of plasma behaviour—first the kinetic approach and then some fluid models—in Chapter 2, including the basic wave modes predicted by each of these models. I also introduce and summarise some important laser-plasma processes.

In Chapter 3 I explain how these theoretical models are translated to practical numerical techniques used in simulations of plasma dynamics. This explanation covers the two major schools of kinetic simulations—the Vlasov and the particle-in-cell methods—and fluid techniques including ideal and extended magnetohydrodynamics (MHD). Both fluid and particle-in-cell simulations are important components of the work I present later in the thesis. Following on from the discussion of large-scale, general plasma simulation methods, I also discuss some smaller-scale numerical techniques I have developed for use in support of my own work or to assist the work of others.

Chapter 4 begins with a high-level overview of some important mechanisms by which ions can be accelerated by laser-plasma interactions. This discussion covers the classical "target-normal sheath acceleration" (TNSA) mechanism, as well as the more novel "radiation pressure acceleration" (RPA) and "collisionless shock acceleration" (CSA). I then present a significant advancement to the theory of collisionless shock acceleration, which permits much more efficient and accurate mapping of the applicable parameter space compared to the previous state-of-the-art developed by Savin [29] and can be used to calculate exactly parameters such as the maximum Mach number at which CSA can exist, which were only estimated in that previous work. I also present an experiment which aimed to produce ion beams by collisionless shock acceleration, and interpret the results with the help of a simulation campaign.

In Chapter 5 I focus on the main application of ion beams in this thesis: the proton radiography diagnostic, which produces images encoding information about the magnetic field structure present in a plasma. I first describe the experimental implementation of this diagnostic and its various uses, and explain the process by which quantitative information about a plasma can be extracted from the resulting images. I then discuss the use of proton radiography in an experiment conducted by a collaboration including Prof. Norreys' research group at the Atomic Weapons Establishment's Orion laser facility, and present my own analysis of some proton radiographs which were produced there. Though analysis of the radiographs which were hoped to show evidence of channelling is inconclusive, I benchmark the analysis method against simulation-generated magnetic fields using radiographs of long pulse-only shots. The research presented in this chapter was published in Philosophical Transactions of the Royal Society A in December 2020 [12].

Chapter 6, my final results chapter, is an investigation into the applicability of tomographic techniques to proton radiography. I first introduce some background about tomography more generally, before explaining the ways in which the proposed 'proton tomography' differs from more conventional applications, the additional challenges posed by these differences, and some pre-processing methods which

can significantly alleviate these difficulties. This work was published in Physical Review E in October 2021 [30]

I conclude the thesis with a summary of the work presented and suggestions for the direction of future work in Chapter 7.

1.3 Author Contributions

The work presented in this thesis ranges from reproduction of existing background knowledge (in Chapters 1 and 2, as well as the initial sections of Chapters 3, 4, 5 and 6) to new research that I have conducted entirely independently, such as the latter half of Chapter 6.

As my DPhil supervisor, Prof. Peter Norreys has been largely responsible for the overall direction of my work throughout my doctoral studies. For example, while I independently carried out the experimental analysis presented in Chapter 5, I would not have had access to the experimental data if not for his influence and would not have been motivated to analyse it so thoroughly if not for his encouragement and insistence that there was more to be learned from the existing data.

The contributions of other members of the research group should also not be understated. Several of the small simulation codes presented at the end of Chapter 3 were motivated by questions posed to me by fellow DPhil student Marko Mayr. I hope they have proven as useful to him as they are enlightening to me.

The experiment described in Chapter 4 was originally proposed by former group member Dr. Alex Savin, and carrying it out would not have been possible without the support (and six-week time investment) of Drs. Ramy Aboushelbaya and Luke Ceurvorst, both alumni of the Norreys group. Current students Marko Mayr, Robert Paddock and Robin Wang also made valuable contributions to the experiment, as did visiting Professor Marta Fajardo, and, of course, the facility staff: Vulcan link scientists Dr. Chris Baird and Margaret Notley and target fabrication's Chris Spindloe and Sam Irving, without whose expertise the experiment would not have been possible.

While I carried out the analysis on data from the Orion experiment in Chapter 5 independently, the experiment itself was actually carried out before I started my DPhil. I must therefore thank all of the former group members who contributed to the success of the experimental campaign—Drs. Luke Ceurvorst, Naren Ratan, Alex Savin and James Sadler—as well as Dr. Matthew Hill for leading the campaign from the Orion side, and the rest of the Orion staff who took part. In particular I would like to thank Dr. Colin Brown for being an excellent contact at Orion and helping to facilitate my access to the shot sheets and diagnostic data necessary to carry out the analysis properly. Chapter 5 contains data collected at the Orion laser facility; Orion experimental data are © British Crown Owned Copyright 2021/AWE.

Unlike the previous two chapters, Chapter 6 is entirely theoretical in nature and therefore comes without a laundry list of experimental contributors. The work presented in the latter sections of this chapter is entirely my own.

Finally, the work presented in the three appendices to this thesis is entirely my own.

1.4 Publications

1. B. T. Spiers et al. “Methods for Extremely Sparse-Angle Proton Tomography”. In: *Physical Review E* 104.4 (Oct. 7, 2021), p. 045201. DOI: 10.1103/PhysRevE.104.045201

2. B. T. Spiers et al. “Whole-Beam Self-Focusing in Fusion-Relevant Plasma”. In: *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 379.2189 (Dec. 7, 2020), p. 20200159. DOI: 10.1098/rsta.2020.0159

3. R. W. Paddock et al. “One-Dimensional Hydrodynamic Simulations of Low Convergence Ratio Direct-Drive Inertial Confinement Fusion Implosions”. In: *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 379.2189 (Jan. 25, 2021), p. 20200224. DOI: 10.1098/rsta.2020.0224

4. R. Aboushelbaya et al. “Measuring the Orbital Angular Momentum of High-Power Laser Pulses”. In: *Physics of Plasmas* 27.5 (May 1, 2020), p. 053107. DOI: 10.1063/5.0005140

5. M. W. Mayr et al. “Nonlinear Wakefields and Electron Injection in Cluster Plasma”. In: *Physical Review Accelerators and Beams* 23.9 (Sept. 2, 2020), p. 093501. DOI: 10.1103/PhysRevAccelBeams.23.093501

6. M. W. Mayr et al. “Wakefields in a Cluster Plasma”. In: *Physical Review Accelerators and Beams* 22.11 (Nov. 15, 2019), p. 113501. DOI: 10.1103/PhysRevAccelBeams.22.113501

7. R. Aboushelbaya et al. “Orbital Angular Momentum Coupling in Elastic Photon-Photon Scattering”. In: *Physical Review Letters* 123.11 (Sept. 13, 2019), p. 113604. DOI: 10.1103/PhysRevLett.123.113604

8. James D. Sadler et al. “Kinetic Simulations of Fusion Ignition with Hot-Spot Ablator Mix”. In: *Physical Review E* 100.3 (Sept. 9, 2019), p. 033206. DOI: 10.1103/PhysRevE.100.033206

9. Alex F. Savin et al. “Energy Absorption in the Laser-QED Regime”. In: *Scientific Reports* 9.1 (June 20, 2019), p. 8956. DOI: 10.1038/s41598-019-45536-x

10. James D. Sadler et al. “Advantages to a Diverging Raman Amplifier”. In: *Communications Physics* 1.1 (May 17, 2018), pp. 1–7. DOI: 10.1038/s42005-018-0021-8

11. A. F. Savin et al. “Attosecond-Scale Absorption at Extreme Intensities”. In: *Physics of Plasmas* 24.11 (Nov. 1, 2017), p. 113103. DOI: 10.1063/1.4989798

2

Basic Laser and Plasma Theory

Contents

2.1 Electromagnetism 12
2.1.1 Electromagnetic Waves 14
2.2 Units 16
2.3 Plasma Dynamics 18
2.3.1 The Vlasov-Maxwell System 18
2.3.2 Fluid Models 20
2.4 Waves in Plasma 27
2.5 Laser-Plasma Processes 32
2.5.1 Forming a Plasma 32
2.5.2 Ionisation 33
2.5.3 Energy Absorption 35
2.5.4 The Ponderomotive Effect 40

In this chapter I introduce some of the important theoretical background used in the remainder of the thesis. In particular, models for laser beam propagation and plasma dynamics described in this chapter are used to motivate the numerical simulation techniques described in Chapter 3, and parts of these models are used in the original work I present in later chapters as well. First, I will describe the behaviour of beams of electromagnetic radiation (i.e. laser pulses) and some useful models for free-space modes: beam profiles which propagate through space without changing their form. I then explain various models of plasma behaviour in

widespread use, starting from the Vlasov equation and applying approximations to derive simplified models which are applicable in different situations. In §2.4 I derive dispersion relations for linear wave modes supported by some of these models, and finally, in §2.5, I discuss some additional phenomena important to the formation of plasma—ionisation and heating through laser energy absorption—as well as the ponderomotive force which is important in high-intensity laser-plasma interactions.

Some key texts for the physics described in this section include Chen [1], Nicholson [40], Mulser and Bauer [41], and Macchi [42]. In particular, derivations of magnetohydrodynamics from a kinetic model (the Vlasov equation) are harder to find good references for than derivations starting from the Navier-Stokes equations of fluid dynamics; my derivation in §2.3.2 is guided largely by that of Nicholson [40].

2.1 Electromagnetism

Maxwell's equations are a system of partial differential equations which provide the basic description of classical electromagnetism. In their 'microscopic' form, they are:

$$\boldsymbol{\nabla} \cdot \boldsymbol{E} = \frac{\rho_c}{\epsilon_0} \tag{2.1a}$$

$$\boldsymbol{\nabla} \cdot \boldsymbol{B} = 0 \tag{2.1b}$$

$$\boldsymbol{\nabla} \times \boldsymbol{E} = -\frac{\partial \boldsymbol{B}}{\partial t} \tag{2.1c}$$

$$\boldsymbol{\nabla} \times \boldsymbol{B} = \mu_0 \boldsymbol{J} + \frac{1}{c^2} \frac{\partial \boldsymbol{E}}{\partial t} \tag{2.1d}$$

Here, the conventional symbols $\boldsymbol{E}$ for electric field, $\boldsymbol{B}$ for magnetic flux density, ϵ_0 and μ_0 for the permittivity and permeability of free space respectively and $c = (\epsilon_0 \mu_0)^{-\frac{1}{2}}$ the speed of light in vacuum are used. The charge density ρ_c carries a subscript c to differentiate it from mass density ρ, which will be employed later in this chapter, and the current density is represented by $\boldsymbol{J}$.

The equations can also be expressed in terms of the scalar (or electric) potential

ϕ and the vector (or magnetic) potential $\boldsymbol{A}$, which are defined by:

$$\boldsymbol{B} = \boldsymbol{\nabla}\times\boldsymbol{A} \tag{2.2a}$$

$$= \boldsymbol{\nabla}\times\boldsymbol{A} \tag{2.2b}$$

$$\boldsymbol{E} = -\boldsymbol{\nabla}\phi - \frac{\partial \boldsymbol{A}}{\partial t} \tag{2.2c}$$

These definitions are invariant under a certain class of transformations known as 'gauge transformations',

$$\boldsymbol{A} \rightarrow \boldsymbol{A} + \boldsymbol{\nabla}\psi \tag{2.3a}$$

$$\phi \rightarrow \phi - \frac{\partial \psi}{\partial t} \tag{2.3b}$$

Commonly employed gauges are the Coulomb gauge, $\boldsymbol{\nabla}\cdot\boldsymbol{A} = 0$ and the Lorenz gauge,

$$c^2\boldsymbol{\nabla}\cdot\boldsymbol{A} + \frac{\partial \phi}{\partial t} = 0 \tag{2.4}$$

Combining pairs of Maxwell's equations, the electromagnetic wave equations can be derived:

$$\nabla^2\boldsymbol{E} - \frac{1}{c^2}\frac{\partial^2 \boldsymbol{E}}{\partial t^2} = \mu_0\frac{\partial \boldsymbol{J}}{\partial t} + \frac{\boldsymbol{\nabla}\rho_c}{\epsilon_0} \tag{2.5a}$$

$$\nabla^2\boldsymbol{B} - \frac{1}{c^2}\frac{\partial^2 \boldsymbol{B}}{\partial t^2} = -\mu_0\boldsymbol{\nabla}\times\boldsymbol{J} \tag{2.5b}$$

In vacuum, the right-hand sides of both wave equations vanish, yielding homogeneous wave equations for both the electric and the magnetic field. Similar wave equations for the potentials[1] are found by substituting Eqs. (2.2) into the Maxwell curl equations:

$$\nabla^2\phi - \frac{1}{c^2}\frac{\partial^2 \phi}{\partial t^2} = -\frac{\rho_c}{\epsilon_0} \tag{2.6a}$$

$$\nabla^2\boldsymbol{A} - \frac{1}{c^2}\frac{\partial^2 \boldsymbol{A}}{\partial t^2} = -\mu_0\boldsymbol{J} \tag{2.6b}$$

These equations are simpler in form than the non-vacuum equations for the fields, and separate the source contributions due to charges and currents. Just like the field wave equations, these reduce to homogeneous wave equations in vacuum. Therefore, to find free space modes I focus on the homogeneous wave equation.

[1] Only in the Lorenz gauge; in other gauges the expressions are more complicated

2.1.1 Electromagnetic Waves

As I am looking for wave modes, I assume an oscillatory component to solutions. Therefore I express the electric field as an envelope-modulated plane wave:

$$\boldsymbol{E}(\boldsymbol{x},t) = \tilde{\boldsymbol{E}}(\boldsymbol{x})e^{i(\boldsymbol{k}\cdot\boldsymbol{x}-\omega t)}, \tag{2.7}$$

with wave-vector (angular spatial frequency) $\boldsymbol{k}$ and angular (temporal) frequency ω

Substituting into the wave equation,

$$c^2\nabla^2\tilde{\boldsymbol{E}} - \frac{\partial^2\tilde{\boldsymbol{E}}}{\partial t^2} + 2i\left(c^2\boldsymbol{k}\cdot\boldsymbol{\nabla}\tilde{\boldsymbol{E}} + \omega\frac{\partial\tilde{\boldsymbol{E}}}{\partial t}\right) = (c^2k^2-\omega^2)\tilde{\boldsymbol{E}} \tag{2.8}$$

Considering the vacuum dispersion relation of light, $\omega = ck$, eliminates the right-hand side of this equation. Furthermore, as I am interested in steady-state solutions for the beam envelope I ignore the time derivatives, and under the assumption that the longitudinal evolution of the envelope is very slow compared to the laser wavelength, i.e.

$$\left|\frac{\partial^2\tilde{\boldsymbol{E}}}{\partial z^2}\right| \ll \left|k\frac{\partial\tilde{\boldsymbol{E}}}{\partial z}\right|, \tag{2.9}$$

the longitudinal component of the Laplacian in Eq. (2.8) can be neglected. Here, without loss of generality I have taken $\boldsymbol{k}$ to be directed along $\hat{z}$. Denoting the remaining spatial second derivatives as $\nabla_\perp^2 = \frac{\partial^2}{\partial x^2} + \frac{\partial^2}{\partial y^2}$, the paraxial Helmholtz equation results:

$$\nabla_\perp^2\tilde{\boldsymbol{E}} + 2i\boldsymbol{k}\cdot\boldsymbol{\nabla}\tilde{\boldsymbol{E}} = 0 \tag{2.10}$$

A special solution to this equation is the trivial 'plane wave', characterised by a constant transverse envelope $\tilde{\boldsymbol{E}} \perp \boldsymbol{k}$. A more interesting solution is the Gaussian beam:

$$\tilde{\boldsymbol{E}} = E_0\frac{q(0)}{q(z)}\exp\left(i\frac{kr^2}{2q(z)}\right)\hat{x} \tag{2.11}$$

where $q(z) = z - iz_R$ is the complex beam parameter, $z_R = kw_0^2/2$ is the Rayleigh length and the beam waist w_0 is the $1/e^2$ intensity radius of the beam at its tightest focus. I have chosen $\hat{x}$ as the direction of polarisation, but any direction transverse

to $\boldsymbol{k}$ is equally valid. This compact expression for the Gaussian beam is a little opaque, but defining the local beam width $w(z) = w_0\sqrt{1+(z/z_R)^2}$ and wavefront curvature $R(z) = z(1+(z_R/z)^2)$ and expanding the complex beam parameter yields the more comprehensible standard expression:

$$\tilde{\boldsymbol{E}} = E_0 \frac{w_0}{w(z)} \exp\left(-\frac{r^2}{w(z)^2}\right) \exp\left(i\left(k\frac{r^2}{2R(z)} - \arctan\frac{z}{z_r}\right)\right). \tag{2.12}$$

The transverse amplitude profile of the beam envelope is Gaussian in form (hence the name), and there is a parabolic transverse variability in the phase, along with a 'Guoy phase' term, $\psi(z) = \arctan z/z_R$ which effectively causes a (gradual) sign inversion as the beam passes through focus. Longitudinally, the width of the beam increases according to $w(z)$ and its intensity is diminished accordingly

More interesting still are the so-called higher-order Gaussian modes. These families of modes form a basis on the space of transverse beam profiles, meaning that arbitrary beam profiles can be expressed as sums of these modes. Each mode propagates without changing its transverse profile, similarly to the Gaussian beam. The difference is that higher-order modes accrue greater Guoy phase as they propagate.

The Hermite-Gaussian modes result from separation of Eq. (2.10) in Cartesian transverse coordinates x, y:

$$\begin{aligned}\tilde{\boldsymbol{E}}_{l,m} = E_0 \frac{w_0}{w(z)} H_l\left(\frac{\sqrt{2}x}{w(z)}\right) H_m\left(\frac{\sqrt{2}y}{w(z)}\right) \exp\left(-\frac{r^2}{w(z)^2}\right) \\ \exp\left(i\left(k\frac{r^2}{2R(z)} - (1+l+m)\psi(z)\right)\right),\end{aligned} \tag{2.13}$$

while the Laguerre-Gaussian modes result from a similar procedure in polar coordinates r, ϕ:

$$\begin{aligned}\tilde{\boldsymbol{E}}_p^{\ell} = E_0 \frac{w_0}{w(z)} \left(\frac{\sqrt{2}r}{w(z)}\right)^{|\ell|} L_p^{|\ell|}\left(\frac{2r^2}{w(z)^2}\right) \exp\left(-\frac{r^2}{w(z)^2}\right) \\ \exp\left(i\left(k\frac{r^2}{2R(z)} - (1+\ell+2p)\psi(z) + \ell\phi\right)\right).\end{aligned} \tag{2.14}$$

Laguerre-Gaussian modes with $\ell \neq 0$ are particularly interesting as they carry *orbital angular momentum* as well as the intrinsic spin angular momentum associated with the polarisation of beams [43]. This orbital angular momentum

is real and can be measured [32], and may have applications in investigations of fundamental physics [35].

Use of these modes and their properties is discussed in more detail in Appendix B, in the context of numerically evaluating the evolution of a laser beam with non-ideal transverse profile.

2.2 Units

Before I move on to plasma dynamics, now is a good time to define a special system of units, commonly referred to as 'PIC units' (After the particle-in-cell simulation model—see §3.1.2) but which are useful throughout plasma physics. These units normalise all relevant physical quantities based on a reference frequency ω_r and reference number density n_r. These quantities are related to each other using the fundamental unit of charge e and electron mass m_e, as well as the vacuum dielectric constant ϵ_0:

$$\omega_r^2 = \frac{n_r e^2}{m_e \epsilon_0}. \tag{2.15}$$

Where the relevant physics are solely a result of plasma dynamics it is conventional to use the background electron density n_e as a reference density, and in this case the corresponding frequency is the *electron plasma frequency* ω_p. On the other hand, for phenomena driven by an external oscillating driver (such as a laser pulse) of frequency ω_0, this frequency is used as the reference. Then the reference density is *critical density* for that frequency of excitation, a name which will be justified in §2.4.

Having defined this pair of interdependent reference units I can now derive

a set of reference units for other quantities:

$$\mathcal{E}_r = m_e c^2 \qquad \text{Energy} \tag{2.16a}$$

$$v_r = c \qquad \text{Speed} \tag{2.16b}$$

$$t_r = \omega_r^{-1} \qquad \text{Time} \tag{2.16c}$$

$$x_r = v_r t_r = \frac{c}{\omega_r} \qquad \text{Length} \tag{2.16d}$$

$$\rho_r = e n_r \qquad \text{Charge Density} \tag{2.16e}$$

$$J_r = \rho_r v_r = e n_r c \qquad \text{Current Density} \tag{2.16f}$$

$$B_r = \frac{m_e \omega_r}{e} \qquad \text{Magnetic Field} \tag{2.16g}$$

$$E_r = \frac{m_e \omega_r c}{e} \qquad \text{Electric Field} \tag{2.16h}$$

$$P_r = n_r \mathcal{E}_r = n_r m_e c^2 \qquad \text{Pressure} \tag{2.16i}$$

$$I_r = n_r \mathcal{E}_r v_r = n_r m_e c^3 \qquad \text{Intensity} \tag{2.16j}$$

$$\mathcal{P}_r = I_r x_r^2 = \frac{(m_e c^2)^2 c \epsilon_0}{e^2} \qquad \text{Power} \tag{2.16k}$$

Note that the units for energy, power and speed are constants: approximately $511\,\mathrm{keV}$, $693\,\mathrm{MW}$ and $3 \times 10^8\,\mathrm{m\,s^{-1}}$ respectively.

There is of course some subtlety in defining these reference quantities: the reference length is *not* the inverse cube root of the reference number density; as such the reference quantity for 'number' is not unity but rather

$$N_r = n_r c^3 \omega_r^{-3} = \frac{\omega_r^{-1} m_e c^3 \epsilon_0}{e^2}. \tag{2.17}$$

This quantity is not only not unity, it depends on the reference frequency. Because number is a dimensionless quantity, any of the above-defined reference values could be multiplied by arbitrary powers of N_r and remain dimensionally consistent (note for example that the definition of $\mathcal{P}_r$ is also expressible as $N_r \mathcal{E}_r \omega_r$). Though it is not immediately obvious, the definitions presented here turn out to be the most convenient.

2.3 Plasma Dynamics

To describe plasma dynamics I begin with Liouville's theorem, which states that phase-space density is conserved along the trajectories of a Hamiltonian system. What this means is that, if the state of a plasma is expressed using its distribution function $f(\boldsymbol{x}, \boldsymbol{p}, t)$ in 6-dimensional phase-space, the total derivative of this function with respect to time is zero:

$$\frac{Df}{Dt} = 0 \tag{2.18}$$

Expanding out this total derivative into convective form,

$$\frac{\partial f}{\partial t} + \frac{\partial \boldsymbol{x}}{\partial t} \cdot \boldsymbol{\nabla}_{\boldsymbol{x}} f + \frac{d\boldsymbol{p}}{dt} \cdot \boldsymbol{\nabla}_{\boldsymbol{p}} f = 0. \tag{2.19}$$

This is the collisionless Boltzmann equation[2], which describes the phase-space evolution of the state of a many-particle system.

2.3.1 The Vlasov-Maxwell System

Considering that a plasma is made up of charged particles, the effect of the Lorentz force will always be present, even in the absence of externally applied forces. Therefore the force term $\frac{d\boldsymbol{p}}{dt}$ of the Boltzmann equation is made up of the Lorentz force as a minimum. Including this form of the force term produces from the collisionless Boltzmann equation the Vlasov equation, and when this equation is coupled to the pair of Maxwell's equations which describe the temporal evolution of the electric and magnetic fields,the result is the Vlasov-Maxwell system:

$$\frac{\partial f_s}{\partial t} + \boldsymbol{v} \cdot \boldsymbol{\nabla}_{\boldsymbol{x}} f_s + \boldsymbol{F}_s \cdot \boldsymbol{\nabla}_{\boldsymbol{p}} f_s = 0 \tag{2.20a}$$

$$\boldsymbol{F}_s = q_s(\boldsymbol{E} + \boldsymbol{v} \times \boldsymbol{B}) \tag{2.20b}$$

$$\frac{\partial \boldsymbol{B}}{\partial t} + \boldsymbol{\nabla} \times \boldsymbol{E} = 0 \tag{2.20c}$$

$$\frac{\partial \boldsymbol{E}}{\partial t} = c^2(\boldsymbol{\nabla} \times \boldsymbol{B} - \mu_0 \boldsymbol{J}) \tag{2.20d}$$

This mathematical description of the dynamics of plasmas includes an equation describing the phase-space evolution of each species (indexed with the subscript

[2]Inclusion of collisions will be discussed briefly in the next section

s) present in the plasma. In the absence of inter-species collisions, the different species can only interact indirectly, via the electromagnetic fields. The only thing necessary to ensure this equation is relativistically correct is to use the relativistic relationship between momentum and velocity: $\boldsymbol{v} = \boldsymbol{p}c/\sqrt{m_s^2c^2 + |p|^2}$.

To include the effects of collisions, a term is introduced to the right-hand side of Eq. (2.20a) representing the Lagrangian-frame phase-space evolution due to collisions:

$$\frac{\partial f_s}{\partial t} + \boldsymbol{v} \cdot \boldsymbol{\nabla}_{\boldsymbol{x}} f_s + \boldsymbol{F}_s \cdot \boldsymbol{\nabla}_{\boldsymbol{p}} f_s = \left(\frac{\partial f_s}{\partial t}\right)_{\text{coll}} \tag{2.21}$$

Modelling collisions is in general quite difficult, but a particularly simple form of the collision operator $(\partial f_s/\partial t)_{\text{coll}}$ is the Bhatnagar-Gross-Krook (BGK) collision operator:

$$\left(\frac{\partial f_s}{\partial t}\right)_{\text{coll}} = \frac{f_s^M - f_s}{\tau_s}, \tag{2.22}$$

which forces the distribution function towards the best local Maxwellian approximation f_s^M at some rate characterised by the relaxation time τ_s, based on the observation that strongly collisional systems quickly approach local thermodynamic equilibrium and therefore a Maxwellian distribution function. This relaxation process typically occurs on picosecond time-scales, with [44]

$$\tau_s \approx 1\,\text{ps}\sqrt{\frac{m_s}{m_e}}\left(\frac{n_s}{10^{21}\,\text{cm}^{-3}}\right)^{-1}\left(\frac{k_B T_s}{\text{keV}}\right)^{\frac{3}{2}}\left(\frac{\ln\Lambda_s}{10}\right)^{-1} \tag{2.23}$$

for a species with density n_s, temperature T_s, mass m_s and Coulomb logarithm $\ln\Lambda_s$.

A more physically-based collision operator is the Fokker-Planck operator:

$$\left(\frac{\partial f_s}{\partial t}\right)_{\text{coll}} = \frac{1}{2}\frac{\partial}{\partial v_i v_j}\left(f_s \left\langle \delta v_i \delta v_j \right\rangle\right) - \frac{\partial}{\partial v_i}\left(f_s \left\langle \delta v_i \right\rangle\right) \tag{2.24}$$

which is often used in numerical schemes (see §3.1.1). The Fokker-Planck operator depends on two ensemble averages, calculated using the local distribution function:

$$\left\langle \delta v_i \right\rangle = \int \Delta v_i \frac{g_{\Delta t}(\Delta \boldsymbol{v}, \boldsymbol{v})}{\Delta t} d^3\Delta\boldsymbol{v} \tag{2.25}$$

$$\left\langle \delta v_i \delta v_j \right\rangle = \int \Delta v_i \Delta v_j \frac{g_{\Delta t}(\Delta \boldsymbol{v}, \boldsymbol{v})}{\Delta t} d^3\Delta\boldsymbol{v} \tag{2.26}$$

Here, the scattering probability $g_{\Delta t}(\Delta \boldsymbol{v}, \boldsymbol{v})$ represents the probability for a 'test' particle to be scattered from a velocity $\boldsymbol{v}$ to $\boldsymbol{v} + \Delta \boldsymbol{v}$ within a time Δt, and most properly should be calculated as a sum over each species present in the plasma. With this form of collision operator included the system of equations is usually referred to as the Vlasov-Fokker-Planck equation.

2.3.2 Fluid Models

A second important class of plasma model is the fluid models. Magnetohydrodynamics (MHD) is one such fluid model—and a particularly important one—but is by no means the only one: others include the two-fluid model [45] and the fluid treatment of electron dynamics I will present in §2.3.2.

The basic procedure involved in deriving a fluid model from the Vlasov-Maxwell or Vlasov-Fokker-Planck equation has a few steps.

First, fluid variables are defined as velocity moments of the distribution function. For example,

$$n_s = \int_{\boldsymbol{p}} f_s d^3\boldsymbol{p} \tag{2.27}$$

$$n_s \boldsymbol{u_s} = \int_{\boldsymbol{p}} \boldsymbol{v} f_s d^3\boldsymbol{p} \tag{2.28}$$

$$\underline{\underline{\mathrm{P}}}_s = m_s \int_{\boldsymbol{p}} (\boldsymbol{v} - \boldsymbol{u_s})(\boldsymbol{v} - \boldsymbol{u_s}) f_s d^3\boldsymbol{p} \tag{2.29}$$

Second, as the first term in the Vlasov equation is a partial time derivative, equivalent moments of the Vlasov equation produce evolution equations for these moments. In principle these evolution equations can be computed for arbitrarily high power moments of velocity. However, the evolution equation for the j^{th} velocity moment includes the $(j+1)^{\text{th}}$ moment due to the factor of $\boldsymbol{v}$ present in the Vlasov equation. Therefore any finite set of these moment equations is not closed, as the highest-order moment evolution equation always depends on an undefined higher-order moment.

Third, therefore, a constitutive relation is introduced to close the system by directly expressing the highest-order moment as a function of some lower-order

moment(s). Most commonly, the pressure is expressed as a function of density, and this relation is known as an equation of state. It is common to use isotropic equations of state, in which case the tensor pressure is some multiple of the identity tensor:

$$\underline{\underline{\mathrm{P}}}_s = P_s(n_s)\underline{\underline{\mathcal{I}}} \tag{2.30}$$

$$\boldsymbol{\nabla} \cdot \underline{\underline{\mathrm{P}}}_s = \boldsymbol{\nabla} P_s(n_s) \tag{2.31}$$

The resulting equations express—for each species—conservation of number[3] and momentum. I close the system by assuming an isotropic, adiabatic equation of state with adiabatic index γ:

$$\frac{\partial n_s}{\partial t} + \boldsymbol{\nabla} \cdot (n_s \boldsymbol{u_s}) = 0 \tag{2.32a}$$

$$\frac{\partial n_s \boldsymbol{u_s}}{\partial t} + \boldsymbol{\nabla} \cdot (n_s \boldsymbol{u_s} \boldsymbol{u_s}) = -\frac{\boldsymbol{\nabla} P_s}{m_s} + \frac{q_s n_s}{m_s}(\boldsymbol{E} + \boldsymbol{u}_s \times \boldsymbol{B}) \tag{2.32b}$$

$$P_s \propto n_s^{\gamma} \tag{2.32c}$$

The system now consists of a set of two moment equations and an equation of state for each species present, in addition to the Maxwell's equations (Eq. (2.20c) and Eq. (2.20d)) which were present in the Vlasov-Maxwell system. The plasma variables ρ_c and $\boldsymbol{J}$ from these equations are now easily expressible in terms of the moments above,

$$\rho_c = \sum_s q_s n_s \tag{2.33}$$

$$\boldsymbol{J} = \sum_s q_s n_s \boldsymbol{u_s}. \tag{2.34}$$

These definitions complete the general multi-fluid model of plasma dynamics. Because simultaneously considering the dynamics of every species present as a separate fluid can be difficult, other fluid models are derived from this starting point by making simplifying approximations. Which approximations are made determine which type of fluid model results.

[3] and equivalently mass and charge

Magnetohydrodynamics

Deriving the equations of magnetohydrodynamics from the starting point of kinetic theory is less common in the literature than derivations which start from the Navier-Stokes equations and add magnetic fields. Nicholson [40] is a good reference for the style of derivation undertaken in this section. Instead of evolving each fluid and the Maxwell's equations separately, magnetohydrodynamics considers a 'mass fluid' with properties:

$$\rho = \sum_s m_s n_s \qquad \text{Mass density} \tag{2.35a}$$

$$\rho \boldsymbol{U} = \sum_s m_s n_s \boldsymbol{u}_s \qquad \text{Mass-averaged velocity} \tag{2.35b}$$

$$\underline{\underline{\mathrm{P}}} = \sum_s \underline{\underline{\mathrm{P}}}_s^{\mathrm{CM}} \qquad \text{Total pressure} \tag{2.35c}$$

and a 'charge fluid' with properties:

$$\rho_c = \sum_s q_s n_s \qquad \text{Charge density} \tag{2.35d}$$

$$\boldsymbol{J} = \sum_s q_s n_s \boldsymbol{u}_s \qquad \text{Current density} \tag{2.35e}$$

When constructing equations of motion for these two sets of fluid variables, the pressure of each species is calculated in the centre-of-momentum frame of the entire plasma rather than for each species individually[4], by expressing the relevant integral using $\boldsymbol{v} = (\boldsymbol{v} - \boldsymbol{U}) + \boldsymbol{U}$ rather than $\boldsymbol{v} = (\boldsymbol{v} - \boldsymbol{u}_s) + \boldsymbol{u}_s$. The resulting pressure tensor (denoted $\underline{\underline{\mathrm{P}}}_s^{\mathrm{CM}}$) is the moment of $m_s(\boldsymbol{v} - \boldsymbol{U})(\boldsymbol{v} - \boldsymbol{U})$, and the corresponding species-wise momentum equations are:

$$\frac{\partial n_s \boldsymbol{u_s}}{\partial t} + \boldsymbol{\nabla} \cdot n_s(\boldsymbol{u_s U} + \boldsymbol{U u_s} - \boldsymbol{UU}) = -\frac{\boldsymbol{\nabla} \cdot \underline{\underline{\mathrm{P}}}_s^{\mathrm{CM}}}{m_s} + \frac{q_s n_s}{m_s}(\boldsymbol{E} + \boldsymbol{u}_s \times \boldsymbol{B}). \tag{2.36}$$

Summing Eq. (2.36) weighted by species mass or charge leads to equations expressing the evolution of the single-fluid velocity and the current density respectively:

$$\frac{\partial \rho \boldsymbol{U}}{\partial t} + \boldsymbol{\nabla} \cdot \rho \boldsymbol{UU} = \rho_c \boldsymbol{E} + \boldsymbol{J} \times \boldsymbol{B} - \boldsymbol{\nabla} \cdot \underline{\underline{\mathrm{P}}} \tag{2.37}$$

$$\frac{\partial \boldsymbol{J}}{\partial t} + \boldsymbol{\nabla} \cdot (\boldsymbol{JU} + \boldsymbol{UJ} - \rho_c \boldsymbol{UU}) = \sum_s \left(\frac{n_s q_s^2}{m_s}(\boldsymbol{E} + \boldsymbol{u}_s \times \boldsymbol{B}) - \frac{q_s}{m_s} \boldsymbol{\nabla} P_s^{\mathrm{CM}} \right) \tag{2.38}$$

[4]Performing calculations using the per-species drift velocity is also possible, but making this change simplifies the arithmetic dramatically

The right-hand side of Eq. (2.38) is still expressed in terms of individual fluid variables. In order to simplify these sums, the fluid is asssumed to consist of only two species: a light, negatively-charged electron species and a heavy, positively-charged ion species indexed with e and i respectively.

With only two particle species, the density and velocity of each species can be expressed in terms of the charge and mass fluid variables, and following some algebraic manipulations the current density evolution equation simplifies to:

$$\begin{aligned}\frac{\partial \boldsymbol{J}}{\partial t} + \boldsymbol{\nabla}\cdot(\boldsymbol{JU} + \boldsymbol{UJ} - \rho_c \boldsymbol{UU}) = &\left(\frac{q_e}{m_e} + \frac{q_i}{m_i}\right)(\rho_c \boldsymbol{E} + \boldsymbol{J}\times\boldsymbol{B}) \\ &- \frac{q_i q_e}{m_i m_e}\rho(\boldsymbol{E} + \boldsymbol{U}\times\boldsymbol{B}) \\ &- \boldsymbol{\nabla}\left(\frac{q_e}{m_e}P_e^{\mathrm{CM}} + \frac{q_i}{m_i}P_i^{\mathrm{CM}}\right) \end{aligned} \tag{2.39}$$

To simplify further, I will introduce some assumptions. The important assumptions of (ideal) magnetohydrodynamics are:

- Plasma is electrically neutral, i.e. $\rho_c \to 0$
- The effects of electron intertia are negligible, i.e. $m_e/m_i \to 0$
- Plasma evolution is slow, i.e. all velocities are much smaller than c
- Plasma features of interest have characteristic length and time scales which are long compared to the electron plasma frequency, i.e. $\frac{\partial(\cdot)}{\partial t} \ll \omega_p(\cdot)$ and $\boldsymbol{\nabla}(\cdot) \ll \frac{\omega_p}{c}(\cdot)$
- Plasma has negligible inter-species collisions, making it a near-perfect conductor

Applying the first two of these conditions to Eq. (2.39) allows a minor simplification:

$$\frac{\partial \boldsymbol{J}}{\partial t} + \boldsymbol{\nabla}\cdot(\boldsymbol{JU} + \boldsymbol{UJ}) = \frac{q_e}{m_e}\boldsymbol{J}\times\boldsymbol{B} + \frac{q_e^2 n_e}{m_e}(\boldsymbol{E} + \boldsymbol{U}\times\boldsymbol{B}) - \boldsymbol{\nabla}\frac{q_e}{m_e}P_e^{\mathrm{CM}} \tag{2.40}$$

This equation, while an improvement, is still unwieldy. The next two conditions will help to understand how readily each term can be ignored. To make this clearer,

I dimensionally normalise variables according to the plasma units Eqs. (2.16) for reference density equal to the background electron density n_0:

$$\tilde{n} = \frac{n}{n_0} \tag{2.41a}$$

$$\tilde{\boldsymbol{J}} = \frac{\boldsymbol{J}}{J_r} \tag{2.41b}$$

$$\boldsymbol{\beta} = \boldsymbol{U}/c \tag{2.41c}$$

$$\tilde{\boldsymbol{B}} = \frac{e\boldsymbol{B}}{m_e\omega_p} \tag{2.41d}$$

$$\tilde{\boldsymbol{E}} = \frac{e\boldsymbol{E}}{m_e\omega_p c} \tag{2.41e}$$

$$\tilde{P} = \frac{P_e^{\mathrm{CM}}}{P_r} \tag{2.41f}$$

$$\tilde{\boldsymbol{\nabla}} = \frac{c}{\omega_p}\boldsymbol{\nabla} \tag{2.41g}$$

$$\frac{\partial}{\partial\tilde{t}} = \frac{1}{\omega_p}\frac{\partial}{\partial t} \tag{2.41h}$$

In the following I will express the 'scale' of some dimensionless variable $\tilde{F}$ as $[F]$. From the third MHD condition, $[\boldsymbol{\beta}] \ll 1$ and from the fourth I write:

$$\frac{\partial\tilde{F}}{\partial\tilde{t}} \sim \epsilon[F] \tag{2.42}$$

$$\tilde{\boldsymbol{\nabla}}\tilde{F} \sim \epsilon[F] \tag{2.43}$$

for some $\epsilon \ll 1$ and where I use the $\sim$ symbol to represent order-of-magnitude equivalence.

Ampére's equation is normalised to:

$$\tilde{\boldsymbol{J}} = \tilde{\boldsymbol{\nabla}}\times\tilde{\boldsymbol{B}} - \frac{\partial\tilde{\boldsymbol{E}}}{\partial\tilde{t}} \tag{2.44}$$

$$[\boldsymbol{J}] \sim \epsilon[\boldsymbol{B}] + \epsilon[\boldsymbol{E}]. \tag{2.45}$$

The order of magnitude of $\boldsymbol{J}$ is smaller than that of (the larger of) $\boldsymbol{E}$ and $\boldsymbol{B}$. To relate the magnitude of the two fields, I now also rewrite Eq. (2.40) using these dimensionless variables:

$$\frac{\partial\tilde{\boldsymbol{J}}}{\partial t} + \boldsymbol{\nabla}\cdot\left(\tilde{\boldsymbol{J}}\boldsymbol{\beta} + \boldsymbol{\beta}\tilde{\boldsymbol{J}}\right) = \tilde{\boldsymbol{E}} + \boldsymbol{\beta}\times\tilde{\boldsymbol{B}} - \tilde{\boldsymbol{J}}\times\tilde{\boldsymbol{B}} + \boldsymbol{\nabla}\tilde{P}_e^{\mathrm{CM}} \tag{2.46}$$

$$\epsilon[\boldsymbol{J}] + \epsilon[\boldsymbol{J}][\boldsymbol{\beta}] \sim [\boldsymbol{E}] + [\boldsymbol{\beta}][\boldsymbol{B}] + [\boldsymbol{J}][\boldsymbol{B}] + \epsilon[P_e^{\mathrm{CM}}] \tag{2.47}$$

Combining with Eq. (2.45) and retaining only lowest-order terms in the small variables $\epsilon, [\boldsymbol{\beta}]$,

$$[\boldsymbol{E}] \sim (\epsilon[\boldsymbol{B}] + [\boldsymbol{\beta}])[\boldsymbol{B}] \ll [\boldsymbol{B}]. \tag{2.48}$$

The electric field is of a smaller order of magnitude than the magnetic field. This makes sense given the already-assumed quasineutrality. This revelation brings about a simple, closed-form expression for $\boldsymbol{J}$ by dropping the electric field dependence in Ampére's law, and the previous reasoning also indicates that the left-hand side of Eq. (2.40) is close to zero, compared to each term on the right:

$$\mu_0 \boldsymbol{J} = \boldsymbol{\nabla} \times \boldsymbol{B} \tag{2.49}$$

$$\boldsymbol{E} = \frac{\boldsymbol{\nabla} P_e^{\mathrm{CM}}}{q_e n_e} - \frac{\boldsymbol{J} \times \boldsymbol{B}}{q_e n_e} - \boldsymbol{U} \times \boldsymbol{B} \tag{2.50}$$

Eq. (2.50) (which has been restored to physical units) is known as a *generalised Ohm's law*, and in order to obtain *ideal* magnetohydrodynamics, it is simplified further by ignoring the pressure and Hall ("J cross B") terms. These two equations complete the set of single-fluid ideal MHD equations, which are reproduced here in their full glory (again returning to SI units):

$$\mu_0 \boldsymbol{J} = \boldsymbol{\nabla} \times \boldsymbol{B} \qquad \text{Definition of current} \tag{2.51a}$$

$$\boldsymbol{E} = -\boldsymbol{U} \times \boldsymbol{B} \qquad \text{Definition of electric field} \tag{2.51b}$$

$$\frac{\partial \rho}{\partial t} + \boldsymbol{\nabla} \cdot \rho \boldsymbol{U} = 0 \qquad \text{Conservation of mass} \tag{2.51c}$$

$$\frac{\partial \rho \boldsymbol{U}}{\partial t} + \boldsymbol{\nabla} \cdot \rho \boldsymbol{U}\boldsymbol{U} = \boldsymbol{J} \times \boldsymbol{B} - \boldsymbol{\nabla} P \qquad \text{Conservation of momentum} \tag{2.51d}$$

$$\frac{\partial \boldsymbol{B}}{\partial t} = -\boldsymbol{\nabla} \times \boldsymbol{E} \qquad \text{Induction Equation} \tag{2.51e}$$

$$P = P(\rho) \qquad \text{Equation of state} \tag{2.51f}$$

Note that only three of these equations are evolution equations; the others are direct relationships used to define the current density, electric field and pressure. It is possible to substitute these three relations into the evolution equations to make the system more compact, but I have chosen not to in the interest of readability.

As an example, an equation of state depending only on density is included (this covers the common adiabatic and isothermal equations of state), but more complex equations of state are often necessary. Approaches to calculating non-ideal equations of state include the Chandrasekhar [46], Nadyozhin [47] and Timmes [48] equations of state.

Static-ion Hydrodynamics

Magnetohydrodynamics is, essentially, the limit of slow phenomena in plasma. However, fast processes are also important and interesting, and a separate model is needed if we wish to understand them in a fluid setting. Here, I will again assume that electron intertia is negligible compared to that of the ions, but this time I use this assumption to argue that the ions can be treated as essentially a stationary background when investigating behaviour on electron time-scales. This allows me to set both the equation of motion and the continuity equation to zero for the heavy ion species, and only consider electrons. The remaining equations are:

$$\frac{\partial n_e}{\partial t} + \boldsymbol{\nabla} \cdot n_e \boldsymbol{u}_e = 0 \tag{2.52a}$$

$$m_e \left(\frac{\partial n_e \boldsymbol{u}_e}{\partial t} + \boldsymbol{\nabla} \cdot n_e \boldsymbol{u}_e \boldsymbol{u}_e \right) = -\boldsymbol{\nabla} P_e + q_e n_e (\boldsymbol{E} + \boldsymbol{u}_e \times \boldsymbol{B}) \tag{2.52b}$$

$$\frac{\partial \boldsymbol{B}}{\partial t} = -\boldsymbol{\nabla} \times \boldsymbol{E} \tag{2.52c}$$

$$\frac{\partial \boldsymbol{E}}{\partial t} = c^2 \boldsymbol{\nabla} \times \boldsymbol{B} - \epsilon_0^{-1} \boldsymbol{J} \tag{2.52d}$$

Now that the ion density n_i is assumed to be constant it is simplifying to express the electron density in terms of this constant background density and the electric field divergence, using Gauss's law for the electric field:

$$n_e = \frac{\epsilon_0 \boldsymbol{\nabla} \cdot \boldsymbol{E} - q_i n_i}{q_e} \tag{2.53}$$

Substituting this expression for electron density into the electron continuity equation, and simplifying using $\boldsymbol{J} = q_e n_e \boldsymbol{u}_e$ and the above equation for electric field evolution yields the symbolic identity $\boldsymbol{\nabla} \cdot (\boldsymbol{\nabla} \times \boldsymbol{B}) = 0$, indicating that this equation is automatically fulfilled by the new definition of the electron density. This leaves

three equations describing the evolution of the three remaining independent variables $(\boldsymbol{u}_e, \boldsymbol{E}, \boldsymbol{B})$ and two equations which define the electron density and pressure:

$$m_e\left(\frac{\partial n_e(\boldsymbol{E})\boldsymbol{u}_e}{\partial t} + \boldsymbol{\nabla}\cdot n_e(\boldsymbol{E})\boldsymbol{u}_e\boldsymbol{u}_e\right) = -\boldsymbol{\nabla}P_e(\boldsymbol{E}) + q_e n_e(\boldsymbol{E})(\boldsymbol{E}+\boldsymbol{u}_e\times\boldsymbol{B}) \quad (2.54a)$$

$$\frac{\partial \boldsymbol{B}}{\partial t} = -\boldsymbol{\nabla}\times\boldsymbol{E} \quad (2.54b)$$

$$\frac{\partial \boldsymbol{E}}{\partial t} = c^2\boldsymbol{\nabla}\times\boldsymbol{B} - \frac{q_e}{\epsilon_0}n_e(\boldsymbol{E})\boldsymbol{u}_e \quad (2.54c)$$

$$n_e(\boldsymbol{E}) = \frac{\epsilon_0\boldsymbol{\nabla}\cdot\boldsymbol{E} - q_i n_i}{q_e} \quad (2.54d)$$

$$P_e(n_e) = P_0\left(\frac{n_e(\boldsymbol{E})}{n_0}\right)^{\gamma} \quad (2.54e)$$

assuming the adiabatic equation of state $Pn^{-\gamma} = \text{constant}$ with adiabatic index γ. This is the most distilled form of the static-ion limit of the two-fluid equations and is useful for studying waves in plama as will be done in the next section. Finally, it can be useful to recast these equations using the native plasma units defined in §2.2 (relative to the zero-order electron density $n_0 = Zn_i$ and the corresponding electron plasma frequency) to normalise all of the variables into dimensionless form (here explicitly including the species charges $q_i = Ze$ and $q_e = -e$), and expand the electron equation of motion into convection form rather than conservation form:

$$n_e(\boldsymbol{E})\left(\frac{\partial \boldsymbol{u}_e}{\partial t} + \boldsymbol{u}_e\cdot\boldsymbol{\nabla}\boldsymbol{u}_e\right) = -\boldsymbol{\nabla}P_e(\boldsymbol{E}) - (\boldsymbol{E}+\boldsymbol{u}_e\times\boldsymbol{B}) \quad (2.55a)$$

$$= -\gamma P_0 n_e^{\gamma-1}(\boldsymbol{E})\boldsymbol{\nabla}n_e(\boldsymbol{E}) - (\boldsymbol{E}+\boldsymbol{u}_e\times\boldsymbol{B}) \quad (2.55b)$$

$$\frac{\partial \boldsymbol{B}}{\partial t} = -\boldsymbol{\nabla}\times\boldsymbol{E} \quad (2.55c)$$

$$\frac{\partial \boldsymbol{E}}{\partial t} = \boldsymbol{\nabla}\times\boldsymbol{B} + n_e\boldsymbol{u}_e \quad (2.55d)$$

$$n_e(\boldsymbol{E}) = 1 + \boldsymbol{\nabla}\cdot\boldsymbol{E} \quad (2.55e)$$

where I have used the relation $\boldsymbol{\nabla}\log P = \gamma\boldsymbol{\nabla}\log n_e$ to evaluate the pressure gradient term.

2.4 Waves in Plasma

The above section describes the evolution of plasma in several different regimes: the general regime described by the Vlasov equation, which is valid even out of

local thermodynamic equilibrium (LTE); the multi-fluid model which describes a plasma with an assumed equation of state; MHD, used for large-scale, slow evolution dominated by ion dynamics and magnetic fields; and an electron hydrodynamic model which assumes static ions and describes electron dynamics in this limit. This section will derive some important waves and oscillations supported by plasma.

Fast Plasma Modes

I will begin with electron modes, for which the appropriate model to inspect is the static-ion model presented above. I first express the three variables in envelope-modulated plane wave form:

$$\boldsymbol{u}_e = \boldsymbol{\delta u} e^{i(\boldsymbol{k}\cdot\boldsymbol{x}-\omega t)} \tag{2.56}$$

$$\boldsymbol{E} = \boldsymbol{\delta E} e^{i(\boldsymbol{k}\cdot\boldsymbol{x}-\omega t)} \tag{2.57}$$

$$\boldsymbol{B} = \boldsymbol{\delta B} e^{i(\boldsymbol{k}\cdot\boldsymbol{x}-\omega t)} \tag{2.58}$$

With the dependent variables in this form, I now linearise the dimensionless static-ion model and evaluate all derivatives using complex envelope functions $(\boldsymbol{\delta u}, \boldsymbol{\delta E}, \boldsymbol{\delta B})$ which are constant in both space and time:

$$-i\omega\boldsymbol{\delta u} = i\boldsymbol{k}\gamma P_0 \left(i\boldsymbol{k}\cdot\boldsymbol{\delta E}\right) - \boldsymbol{\delta E} \tag{2.59a}$$

$$-i\omega\boldsymbol{\delta B} = -i\boldsymbol{k}\times\boldsymbol{\delta E} \tag{2.59b}$$

$$-i\omega\boldsymbol{\delta E} = i\boldsymbol{k}\times\boldsymbol{\delta B} + \boldsymbol{\delta u} \tag{2.59c}$$

In the linear regime the superposition principle applies, meaning the electric field can be decomposed into components which evolve independently. I therefore separate the treatment of these equations into solutions with $\boldsymbol{\delta E} \parallel \boldsymbol{k}$ and $\boldsymbol{\delta E} \perp \boldsymbol{k}$.

Considering first a *parallel* oscillation $\boldsymbol{\delta E} \parallel \boldsymbol{k}$, the above equations simplify to

$$-i\omega\delta u = -k^2\gamma P_0 \delta E - \delta E \tag{2.60}$$

$$-i\omega\delta B = 0 \tag{2.61}$$

$$-i\omega\delta E = \delta u \tag{2.62}$$

(2.61) states that such oscillations are electrostatic, and the other two equations combine to yield the dispersion relation for Langmuir waves:

$$\omega^2 = 1 + k^2\gamma P_0 \tag{2.63}$$

Converting this final result back into physical units,

$$\omega^2 = \omega_p^2 + k^2\gamma\frac{P_0}{m_e n_0} \tag{2.64}$$

$$= \omega_p^2 + k^2\gamma\frac{k_B T_e}{m_e} \tag{2.65}$$

$$= \omega_p^2 + k^2 v_{\text{th}}^2 \tag{2.66}$$

shows that a cold plasma oscillates at the electron plasma frequency, and dispersion starts to occur as the temperature of the plasma increases. Furthermore, the phase velocity v_ϕ and group velocity v_g of these waves are given by:

$$v_\phi = \frac{\omega}{k} = v_{\text{th}}\sqrt{\frac{\omega^2}{\omega^2 - \omega_p^2}} \tag{2.67}$$

$$v_g = \frac{\partial\omega}{\partial k} = v_{\text{th}}^2/v_\phi. \tag{2.68}$$

In the cold-electron limit these velocities, as well as k, cease to be meaningful – cold plasma only supports *plasma oscillations* rather than travelling *plasma waves*.

I now return to Eqs. (2.59) and consider a mode for which $\boldsymbol{E} \perp \boldsymbol{k}$

$$-i\omega\delta u = -\delta E \tag{2.69a}$$

$$-i\omega\delta B = -ik\delta E \tag{2.69b}$$

$$-i\omega\delta E = -ik\delta B + \delta u \tag{2.69c}$$

Where the first and third equations are found to be polarised along the direction parallel to $\boldsymbol{E}$ and the second is polarised along the direction mutually perpendicular to both $\boldsymbol{k}$ and $\boldsymbol{E}$. This leads to the conclusion that for this electromagnetic mode, the direction of propagation, electric field and magnetic field are all mutually orthogonal. I again solve for the dispersion relation, which comes out as $\omega^2 = 1 + k^2$. In physical units,

$$\omega^2 = \omega_p^2 + c^2k^2 \tag{2.70}$$

In vacuum, $\omega_p \to 0$ and the vacuum dispersion of light is recovered. As light propagates into increasingly dense plasma, the wave-vector k must decrease in order to satisfy the dispersion relation. Eventually the electron density increases sufficiently that $\omega_p(n_e) = \omega$ for a given frequency of light, and $k = 0$. At higher density, k becomes imaginary, changing the nature of the mode from a propagating wave to evanescent decay. For this reason the density at which this happens is called the *critical density* (n_c) for a given frequency of electromagnetic radiation.

The velocities of the plasma-electromagnetic wave are:

$$v_\phi = c\sqrt{\frac{\omega^2}{\omega^2 - \omega_p^2}} = \frac{c}{\sqrt{1 - \frac{n_e}{n_c}}} \tag{2.71}$$

$$v_g = \frac{c^2}{v_\phi} \qquad = c\sqrt{1 - \frac{n_e}{n_c}}, \tag{2.72}$$

where I have used the relation $\frac{n_e}{n_c} = \frac{\omega_p^2}{\omega^2}$ to simplify the end result. The expression $\eta = \sqrt{1 - n_e/n_c}$ is the refractive index of the plasma, as quoted in §1.1

MHD Modes

To study the slower modes of a plasma, I employ a similarly linearised version of Eqs. (2.51) which comprise the ideal MHD model. The equations become:

$$\mu_0 \boldsymbol{\delta J} = i\boldsymbol{k}\times\boldsymbol{\delta B} \tag{2.73a}$$

$$\boldsymbol{\delta E} = -\boldsymbol{\delta U}\times\boldsymbol{B}_0 \tag{2.73b}$$

$$-i\omega\delta\rho + i\rho_0\boldsymbol{k}\cdot\boldsymbol{\delta U} = 0 \tag{2.73c}$$

$$-i\omega\rho_0\boldsymbol{\delta U} = \boldsymbol{\delta J}\times\boldsymbol{B}_0 - i\boldsymbol{k}v_\mathrm{s}^2\delta\rho \tag{2.73d}$$

$$-i\omega\boldsymbol{\delta B} = -i\boldsymbol{k}\times\boldsymbol{\delta E}. \tag{2.73e}$$

Here I have included static, uniform unperturbed density and magnetic field, such that

$$\rho = \rho_0 + \delta\rho\ e^{i(kx-\omega t)} \tag{2.74a}$$

$$\boldsymbol{B} = \boldsymbol{B}_0 + \boldsymbol{\delta B}\ e^{i(kx-\omega t)} \tag{2.74b}$$

$$\boldsymbol{J} = \boldsymbol{\delta J}\ e^{i(kx-\omega t)} \tag{2.74c}$$

By substituting the expressions for $\delta\rho$, $\boldsymbol{\delta B}$ and $\boldsymbol{\delta J}$ into the transformed equation of motion, the general dispersion relation is

$$\omega^2\boldsymbol{\delta U} = k^2\frac{B_0^2}{\mu_0\rho_0}\hat{B}\times(\hat{k}\times(\hat{k}\times(\hat{B}\times\boldsymbol{\delta U}))) + k^2 v_s^2\hat{k}(\hat{k}\cdot\boldsymbol{\delta U}) \tag{2.75}$$

The unfortunate consequence of including a background magnetic field is a rather unsightly quadruple cross-product. Simplifying this using standard vector algebra identities and expressing the equation as a matrix multiplication gives

$$\left[\left(v_s^2+v_A^2\right)\hat{k}\otimes\hat{k} - v_A^2\left(\hat{k}\cdot\hat{B}\right)\left(\hat{k}\otimes\hat{B}+\hat{B}\otimes\hat{k}\right)\right]\boldsymbol{\delta U} = \left[\frac{\omega^2}{k^2} - v_A^2\left(\hat{k}\cdot\hat{B}\right)\right]\boldsymbol{\delta U} \tag{2.76}$$

where $v_A^2 = \frac{B_0^2}{\mu_0\rho_0}$ is the Alfvén speed. This is an eigenvalue equation: valid $\boldsymbol{\delta U}$ are eigenvectors of the matrix on the left-hand side, and have eigenvalues $\frac{\omega^2}{k^2} - v_A^2\left(\hat{k}\cdot\hat{B}\right)$. Each eigenvalue λ of this matrix therefore corresponds to a wave mode with dispersion relation $\omega^2 = (\lambda + v_A^2\left(\hat{k}\cdot\hat{B}\right))k^2$.

Important examples for which it is easy to compute the eigen-decomposition of the dispersion matrix are as follows:

For waves propagating along the magnetic field, $\hat{k} = \hat{B}$. Then our matrix is equation is reduced to:

$$\left[\left(v_s^2-v_A^2\right)\hat{k}\otimes\hat{k}\right]\boldsymbol{\delta U} = \left[\frac{\omega^2}{k^2} - v_A^2\right]\boldsymbol{\delta U} \tag{2.77}$$

The outer product of a vector with itself is always of unit matrix rank, implying that all but one of this matrix's eigenvalues are zero. It is easily verified that these zero eigenvalues correspond to transverse modes; those with $\boldsymbol{\delta U} \perp \hat{k}$. The other eigenvalue, corresponding to a longitudinal mode $\boldsymbol{\delta U} \parallel \hat{k}$ is $v_s^2 - v_a^2$. Rearranging into the usual form of a dispersion relation, the two modes are:

$$\omega^2 = k^2 v_A^2 \qquad \text{Alfvén wave} \tag{2.78}$$

$$\omega^2 = k^2 v_s^2 \qquad \text{Ion-acoustic (sound) wave} \tag{2.79}$$

The other easy case describes modes propagating orthogonal to the background field, $\hat{k} \perp \hat{B}$. Then the reduced form of the matrix is:

$$\left[\left(v_s^2+v_A^2\right)\hat{k}\otimes\hat{k}\right]\boldsymbol{\delta U} = \frac{\omega^2}{k^2}\boldsymbol{\delta U} \tag{2.80}$$

Again, the matrix multiplication on the left-hand side vanishes for transverse modes. Longitudinal modes produce eigenvalue $v_s^2 + v_A^2$. The longitudinal cross-field mode is described by dispersion relation:

$$\omega^2 = k^2(v_s^2 + v_A^2) \qquad \text{Magnetosonic wave} \tag{2.81}$$

Many intermediate cases are possible, corresponding to oblique angles between the propagation direction and the background magnetic field density. These are not important for the work presented in this thesis, so I will exclude deeper discussion of them. Information on the full variety of MHD modes can be found in most MHD textbooks, for example Goedbloed and Poedts [49].

2.5 Laser-Plasma Processes

The majority of this chapter so far has dealt with the existence of laser beams and plasma, each in isolation. The brief discussion of the electromagnetic dispersion relation in §2.4, and the cut-off behaviour it exhibits at critical density, only scratches the surface of the many forms of interaction possible between laser light and plasma.

In this section I will briefly discuss some of the other processes of importance in the interaction between laser pulses and pre-formed plasma, as well as the laser-matter interaction processes involved in the initial formation of a plasma.

2.5.1 Forming a Plasma

In order for any of the physics described above to occur, it is first required that we produce a plasma. As most forms of matter encountered on earth are not ionised, the first stage in this process is ionising some or all of the atoms in the target material. Merely being ionised is not enough to make a material a plasma, however - another important condition is that the density-Debye volume product is large:

$$N_D = n_e \frac{4\pi\lambda_D^3}{3} \gg 1 \tag{2.82}$$

This condition - N_D is easily interpreted as the number of particles in a 'Debye sphere', and this number must be large - is necessary for the emergence of the

'collective behaviour' that differentiates a plasma from a simple ionised gas. Note that due to the $\sqrt{T_e/n_e}$ scaling of the Debye length, increasing the density actually *reduces* the number of particles within this Debye sphere, so the way to increase N_D is by heating the plasma. As an example, material of critical density for micron-wavelength light, $n_e \approx 10^{21}\,\mathrm{cm}^{-3}$, would need a temperature much in excess of $k_B T_e = 7\,\mathrm{eV}$ to fulfil this condition. As an engineering formula:

$$k_B T_e \gg 7\,\mathrm{eV} \cdot \left(\frac{n_e}{10^{21}\,\mathrm{cm}^{-3}}\right)^{\frac{1}{3}} \tag{2.83}$$

In contrast to the relatively 'nice' linear interactions discussed in the previous sections, formation of a plasma is an inherently violent process involving very high intensities of laser radiation, and as such many of the constituent processes are highly nonlinear and accordingly harder to model analytically. As a result the below descriptions will be less rigorously mathematical and more phenomenological than those above.

2.5.2 Ionisation

There are many ways that an atom can be ionised[5]. A photon with energy in excess of the atom's ionisation energy could excite an electron into the continuum in one go, but the wavelengths of light utilised in high-intensity laser experiments are not usually sufficient to do so. Instead, the main *directly* laser-driven ionisation processes are usually multiphoton in nature. Due to the number of photons involved these can be conceptualised in terms of the classical laser electric field and its interaction with the atom's Coulomb potential, which is easier to reason about than considering quantum electronic energy levels and photon states. An illustration of this model is presented in Fig. 2.1. More detailed treatments of the ionisation dynamics in laser irradiated matter can be found in the literature, for example Robinson, Bell, and Kingham [50] or Bauer [51].

In this model, the Coulomb potential in which the electron is bound is distorted by the presence of the strong electric field of the laser. This 'tilt' of the potential

[5] and that an ion which is not completely stripped of electrons can be further ionised

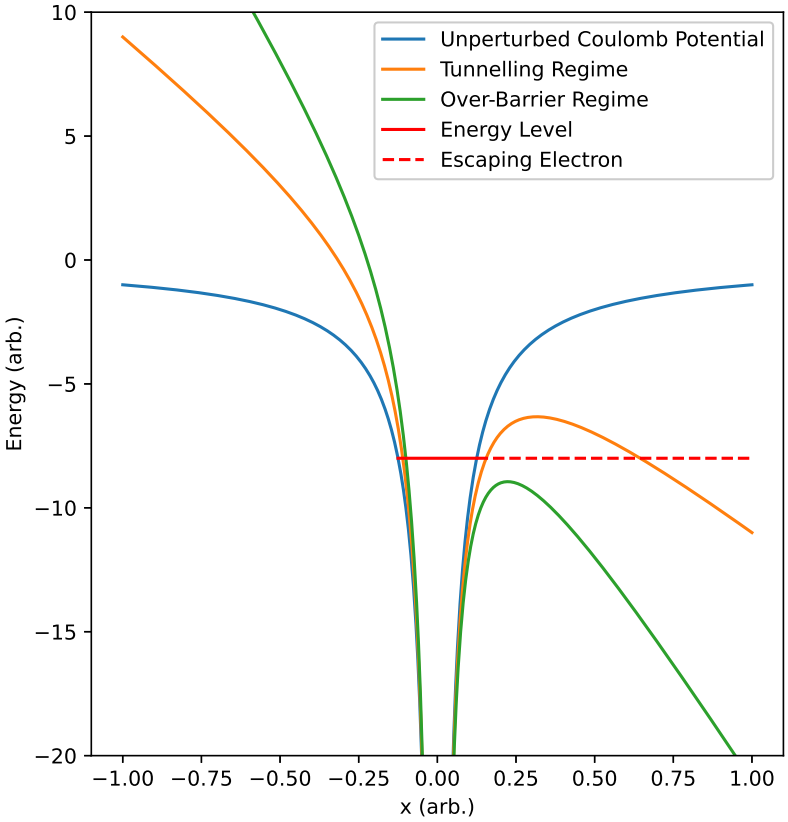


Figure 2.1: Schematic illustration of the energy landscape involved in field ionisation. Laser electric fields introduce an effective linear 'slant' to the electric potential. The electron energy level (red, solid) is fully bound by the unperturbed Coulomb potential of its atom's nucleus. Under the influence of a moderately strong quasistatic electric field, the net potential (orange) is lowered sufficiently for the electron to quantum-mechanically tunnel under the barrier (red, dashed). Very strong fields can lower the maximum of the net potential (green) to a great enough extent that the electron is classically allowed to pass over it, at which point an enhancement in ionisation rate is observed[51]

means that at some point 'downstream' of the atom, the net potential becomes lower than the energy of the electron's bound state energy, enabling it to quantum-mechanically tunnel through the barrier. The stronger the laser's electric field is, the more the Coulomb barrier will be suppressed and the higher the tunnelling probability becomes. This is because the distance to tunnel is shorter and the peak of the potential to tunnel through is lower.

In the limit of extremely intense laser irradiation, even the peak itself can

be lowered below the electron's bound state energy. In this regime—"over barrier" ionisation—escape of the electron is classically allowed, with no tunnelling necessary. As a result, ionisation rates increase greatly in the presence of over barrier ionisation [51].

Other processes may drive ionisation as well as the interaction with the electromagnetic field of the laser. Collisions of atoms (or ions) with already-free electrons can cause further ionisations to occur, and this process can cascade into an *ionisation avalanche*.

Typically, the interaction of an intense laser pulse with a solid target begins with multiphoton ionisations at the material's surface. Once some electrons are liberated and begin oscillating in the laser field, collisions between these electrons and the material's ions enhance the ionisation rate. Eventually the material will reach a steady-state degree of ionisation, with ionisations balanced by recombinations – at sufficiently high temperatures this steady state will be one of complete ionisation. The average ionisiation degree in local thermodynamic equilibrium is a function of density and temperature and can be included as a constitutive relation similar to the equation of state used to relate quantities like density, temperature and pressure in fluid models [52, 53].

2.5.3 Energy Absorption

Once the target material is ionised (and as it is being ionised), its temperature must increase to a degree sufficient to begin exhibiting plasma behaviour. The obvious source of energy in this system is of course the laser irradiation, and indeed it is by interacting with the incident laser fields that the plasma is heated.

It is well-known that an isolated electron oscillating in an electromagnetic plane wave returns to rest at the end of each laser cycle [42]. The electron oscilates perfectly in-phase with the plane wave, in the sense that its transverse momentum has the same time-dependence as the vector potential of the wave – energy given to

the electron by the field in one half-cycle is always taken away again in the next[6]:

$$\frac{\boldsymbol{p}_\perp(t)}{m_e c} = \frac{e\boldsymbol{A}(t)}{m_e c} = \boldsymbol{a}(t) \tag{2.84}$$

This motion is often labeled a 'quiver velocity', though I prefer to consider it a *quiver momentum* in the interest of easier relativistic generalisation. It is a concept which will reappear several times in the subsequent sections.

Because of the in-phase oscillations of electron momentum and vector potential described above, in the absence of other influences an electron absorbs no energy, on average, per laser cycle.

Laser energy absorption mechanisms are, essentially, processess by which this phase relationship between laser oscillations and electron motion are broken. Such mechanisms are numerous; in the interest of brevity I will provide brief phenomenological descriptions of some important mechanisms here; more detailed treatments may be found throughout the literature, for example in books by Colvin and Larsen [54] and by Macchi [42].

Inverse-Bremsstrahlung Absorption

One conceptually simple way of dephasing the electrons from the laser is by collisions with nearby ions, and energy absorption due to these collisions is termed inverse-Bremsstrahlung absorption[7].

As collisions effectively randomise the direction of the electron's momentum they break the phase relationship between the electron's motion and the driving field, and the electron escapes from the interaction still carrying some of the field's energy. Over the course of many collisions this randomisation of momenta manifests as an increase in the material's temperature.

Understanding this phenomenology, it is clear that the energy absorption rate due to inverse-Bremsstrahlung should be proportional to both the electron-ion

[6]The electron also gains longitudinal momentum, which also returns to zero every half-cycle. This is discussed some more in §2.5.3

[7]This naming is by analogy to Bremsstrahlung, the quantum-mechanical process by which an electron whose motion is disturbed by collision with an ion emits a photon. Conversely, inverse-Bremsstrahlung absorption can be thought of as the absorption of a photon by an electron, mediated by the nearby ion.

collision frequency τ_{ei}^{-1} and the average electron quiver energy, which I term the 'ponderomotive potential' in anticipation of §2.5.4:

$$\frac{dE}{dt} \sim \frac{\Phi_{\text{pond}}}{\tau_{ei}}, \tag{2.85}$$

where

$$\frac{\Phi_{\text{pond}}}{m_e c^2} = \left\langle \sqrt{1 + |\boldsymbol{a}(t)|^2} - 1 \right\rangle \tag{2.86}$$

$$\approx \left\langle 1 + \frac{|\boldsymbol{a}(t)|^2}{2} - 1 \right\rangle \tag{2.87}$$

$$\approx \frac{a_0^2}{4}, \tag{2.88}$$

for field amplitude of the form $\boldsymbol{a}(t) = a_0 \hat{x} e^{i\omega t}$ and where angle-brackets represent a cycle-average.

It is notable that as a collisional absorption mechanism, inverse-Bremsstrahlung inherits an unfavourable temperature scaling from the collision frequency. The hotter the electrons become, the slower the collision frequency and the less efficient this mechanism is. At high laser intensities therefore, other—collisionless—mechanisms must take over to bring a plasma to accordingly high temperatures.

Resonance Absorption

The inverse-Bremsstrahlung absorption mechanism described above is applicable where a laser pulse is travelling through under-dense plasma. It is worth considering what happens as laser energy approaches critical density, beyond which it cannot propagate (and therefore beyond which inverse Bremsstrahlung does not occur). Due to the relationship between electron density and refractive index:

$$\eta = \sqrt{1 - \frac{n_e}{n_c}}, \tag{2.89}$$

a laser pulse travelling at an angle to a density gradient will be turned away from regions of high density. It is well-known [42, 54] that a laser ray which approaches an electron density gradient at an angle θ from normal will reach a maximum density of $\cos^2\theta$ before being turned around. Light which is p-polarised[8] will, at the turning

[8] meaning that its electric field oscillates within the plane of incidence

point, have its direction of polarisation pointing parallel to the reflecting surface. This oscillating electric field is resonant with plasma oscillations (see §2.4) where the plasma frequency ω_p and the laser frequency ω_0 are the same, i.e. at the critical density surface. Macchi [42] notes that field oscillations can quantum-mechanically tunnel from the turning point $n_e = n_c \cos^2 \theta$ to the critical density surface $n_e = n_c$, and that this process is—unsurprisingly—more efficient in plasmas of shorter density scale-length, where these two densities are closer together. Thus this 'resonance absorption' allows the transfer of energy from laser light at its point of reflection into electron plasma oscillations at the critical density surface. A highly detailed treatment of resonance absorption processes can be found in Mulser and Bauer [41, ch. 4], along with discussion of the 'wave-breaking' processes attributed with transfer of energy from these oscillations into chaotic (i.e. thermal) energy.

The resonant excitation of plasma oscillations by the laser is a collisionless process, and as a result its efficiency is greater than that of inverse-Bremsstrahlung once the electron temperature is high [41, ch. 4].

Vacuum Heating

In the case of a physical system similar to the one described above, but taken to the limit of very high laser intensities, Macchi [42, §4.2.2] argues that a plasma can be modeled as an infinitely steep density step. Within this simplified model, they describe the dynamics of (fluid) electrons which are initially located within the plasma. At the point of reflection of the laser from the surface of the plasma, the electric field component normal to the surface acts to drag electrons out into the vacuum side of the interface. A half-laser cycle later, these electrons are accelerated back into the dense plasma. Once they have entered the dense plasma by more than the (very short) skin depth of evanescent waves in the plasma, they escape the influence of the laser's oscillating electromagnetic fields and continue to propagate, carrying energy gained from the laser into the target.

Laser energy absorption of this type was first proposed by Brunel, in the artfully-named "Not-so-Resonant, Resonant Absorption" [55]. It is often known as the 'Brunel effect' or 'vacuum heating' by scientists desiring more sober terminology.

$\boldsymbol{J}\times\boldsymbol{B}$ effects

A commonality of resonance absorption and Brunel heating is the mechanism by which motion of electrons parallel to the density gradient is induced: in both cases, the laser pulse must approach the plasma at oblique incidence, and must be polarised such that there is an electric field component parallel to the gradient. This electric field drives oscillations parallel to the gradient at the laser frequency, ω_0. At relativistic intensity however, there is another force which drives motion along the laser axis. The $\boldsymbol{v}\times\boldsymbol{B}$ force arising from the quiver motion of electrons in the laser field contributes a longitudinal momentum proportional to the square of the laser amplitude [42, §2.1.3]:

$$\frac{p_z(t)}{m_e c} = \frac{1}{2}\left|\frac{e\boldsymbol{A}(t)}{m_e c}\right|^2 = \frac{a(t)^2}{2} \sim a_0^2 \sin^2\omega_0 t \tag{2.90}$$

$$\frac{\boldsymbol{p}_\perp(t)}{m_e c} = \frac{e\boldsymbol{A}(t)}{m_e c} = \boldsymbol{a} \sim a_0 \sin\omega_0 t \tag{2.91}$$

This forward motion has a secular component and one which oscillates at twice the laser frequency (as $2\sin^2\omega_0 t = 1 - \cos(2\omega_0 t)$). Like the transverse oscillations, it also returns to rest at the end of each laser cycle (though the electron does end up displaced forwards in space).

In the same step-like plasma profile as Brunel absorption, then, electrons at the edge of the target can experience a fraction of a laser cycle of forwards acceleration by the $\boldsymbol{v}\times\boldsymbol{B}$ force before being pushed into the target and leaving the influence of the laser fields. Unlike resonance absorption and Brunel heating, the accelerating force is parallel to the laser axis and as a result this mechanism can be observed at normal incidence.

The $\omega = 2\omega_0$ frequency of the oscillating longitudinal force diminishes the resonant excitation of plasma oscillations in shallow density gradients. This can be seen by considering that the density is proportional to the square of the plasma

frequency, so that oscillations of frequency $2\omega_0$ are resonant at the quadruple-critical density surface, $n_e = 4n_c$. Resonances will therefore depend on evanescent fields reaching up a factor 4 in density and will accordingly be heavily suppressed. $\boldsymbol{J}\times\boldsymbol{B}$ heating is therefore more often considered in the context of 'not-so-resonant' absorption similar to that discussed in §2.5.3 .

It is also important to note that Eq. (2.90) for the longitudiunal momentum is only applicable for electrons experiencing linearly-polarised light. For circularly polarised light, there is no oscillating $\boldsymbol{J}\times\boldsymbol{B}$ force as the particle velocity and magnetic field remain parallel through an optical cycle.

2.5.4 The Ponderomotive Effect

The ponderomotive effect arises when charged particles are subjected to a non-homogeneous oscillating electromagnetic field, for instance a focused laser pulse. Over the course of a laser cycle, electric fields 'kick' charged particles, first in one direction (say, upwards) and then in the opposite direction (downwards). The key realisation that leads to the ponderomotive effect is that the 'upwards' and 'downwards' kicks are applied at different locations. If a particle is initially 'kicked' into a region of stronger fields, the 'kick' which returns it will be stronger, so it will move further during the 'downwards' half-cycle than it did in the 'upwards' half-cycle. This leads to a net cycle-averaged drift towards regions of lower field amplitude, and more careful mathematical analysis shows that this drift can be modelled as a force arising from the ponderomotive potential (defined here more generally than above, accounting for species of different mass and charge):

$$\boldsymbol{F}_{\text{pond}} = -\boldsymbol{\nabla}\Phi_{\text{pond}}, \tag{2.92}$$

$$\Phi_{\text{pond}} = \frac{q^2E^2}{4m\omega^2}, \tag{2.93}$$

$$= \frac{1}{4}mc^2 \cdot \left(\frac{Z}{M}\right)^2 a^2, \tag{2.94}$$

where as usual, a is the normalised amplitude of the oscillating field, $Z = q/e$ is the charge number of the particle species involved and $M = m/m_e$ is its (electron-normalised) mass. We see that the ponderomotive force is stronger for higher-charged

species, but also for lighter species – the acceleration induced is proportional to the square of the charge-mass ratio. Possibly most interestingly, its direction is the same for positively- and negatively-charged particle species. Ponderomotive effects therefore cause regions of high laser intensity to evolve towards evacuation of all charged species – an important factor in the formation of the laser-plasma channels studied in the latter sections of Chapter 5.

The above outline is non-relativistic in nature. A rigorous relativistic treatment was presented by Bauer, Mulser, and Steeb [56]. Their full derivation is rather dense and opaque, so would not be enlightening to reproduce here, but the end result is surprisingly elegant:

$$\boldsymbol{F}_{\text{pond}} = -mc^2 \boldsymbol{\nabla} \gamma, \tag{2.95}$$

where for a linearly ($\alpha = 2$) or circularly ($\alpha = 1$) polarised monochromatic wave in vacuum, the Lorentz factor γ is given by:

$$\gamma^2 = 1 + \frac{a^2}{\alpha}. \tag{2.96}$$

3

Numerical Modelling of Plasmas

Contents

3.1 Large-Scale Simulations 43
3.1.1 Vlasov Simulations 44
3.1.2 Particle-in-Cell Simulations 46
3.1.3 Fluid Model Simulations 54
3.2 Semi-Analytic and Smaller-Scale Modelling 55

In plasma physics it is common to employ numerical simulation methods to model plasma dynamics and laser-plasma interactions. These simulations evolve equations which describe the dynamics of a plasma system according to models such as those described in the previous chapter. The plasma regime being simulated is the primary factor determining which model is chosen, though it is also important to consider whether any additional physics needs to be described, and the computational resources available to run the simulation.

3.1 Large-Scale Simulations

Large-scale, 'full-physics' simulation of plasma systems is usually performed using one of three methods, each of which is in some way derived from the Vlasov-Maxwell

system of equations[1], the kinetic theoretic description of plasma dynamics detailed in §2.3.1. In the following discussion I will use the term 'Eulerian' to describe simulation variables which are defined on a fixed grid; 'Lagrangian' quantities, in contrast, are tracked in a frame which moves with the plasma.

3.1.1 Vlasov Simulations

Possibly the conceptually simplest method in use for numerical plasma simulations (though computationally the most demanding) is the fully-Eulerian kinetic phase-space model [57]. This model—commonly referred to as Vlasov simulation—evolves the full per-species distribution function on a grid which spans all of the momentum dimensions of the simulation as well as its spatial dimensions. Maxwell's equations, being independent of momentum in the phase-space treatment, are evolved on a grid which spans only the spatial dimensions of the system. The evolution is carried out using the Vlasov-Maxwell system, Eqs. (2.20). Vlasov methods of this kind are not used in this thesis, but I still discuss them here in order to demonstrate why other methods have been more appropriate for my work, and also because a discussion of numerical methods for plasma dynamics would be incomplete without mentioning them.

In the most complete case—three dimensions in each of space and momentum—the grid describing the distribution function is six-dimensional, which leads to very restrictive memory requirements. It therefore rapidly becomes infeasible to run Vlasov-model simulations in this way as the physical size of the system being studied increases. This model is therefore usually limited to relatively small-scale simulations where the compromises involved in the other models would not be acceptable.

It is quite common to reduce the effective dimensionality of the velocity space using a truncated series expansion of the distribution function. As a first example, `IMPACT` [58] and `VALIS` [59] both use a Taylor series in the Cartesian velocity

[1] Although the fluid approach can also be derived in the other direction, starting from the Navier-Stokes equations

components v_x, v_y, v_z:

$$f(\boldsymbol{x}, \boldsymbol{v}) = f_0(\boldsymbol{x}, |\boldsymbol{v}|) + \boldsymbol{f}_1(\boldsymbol{x}, |\boldsymbol{v}|) \cdot \hat{v} + \hat{v} \cdot \underline{\underline{f}}_2(\boldsymbol{x}, |\boldsymbol{v}|) \cdot \hat{v} + ... \tag{3.1}$$

where, as is conventional I have collected terms with the same 'sum of powers' $N = a+b+c$ in the expansion $v_x^a v_y^b v_z^c$ into symmetric coefficient tensors of order N in order to express the resulting expansion in a more compact, coordinate-free notation.

Alternatively, the distribution function may be discretised using spherical harmonic decomposition, such as in the code `OSHUN` [60]:

$$f(\boldsymbol{x}, \boldsymbol{v}) = \sum_l \sum_m f_l^m(\boldsymbol{x}, |\boldsymbol{v}|) Y_l^m(\hat{v}). \tag{3.2}$$

Once the Vlasov equation has been rewritten in terms of the desired series expansion (if one is indeed desired), it is common to solve it implicitly. Early codes [61] achieved this using an alternating direction implicit (ADI) method, made possible by ignoring the effect of magnetic fields, and resulting in significantly smaller computational cost than implicitly stepping the entire system at once. `IMPACT`[2] [58] accepted the cost required to solve the full matrix equation and was able to include magnetic field dynamics in two-dimensional simulations as a result. More recently, codes such as `VALIS` [59] emerged, which evolve a distribution function resolved on a full space-momentum grid (i.e. without use of a series expansion), albeit still only in two spatial and two momentum dimensions.

Overall, Vlasov solvers provide a useful capability, and have several advantages over the particle-in-cell (PIC) simulations I will describe in the next section. For example, they do not suffer from statistical noise issues to the same extent as PIC simulations, nor the grid-particle instabilities which limit energy conservation [62]. However, their memory consumption when evaluating physics in more than one spatial dimension can limit their practicality in many cases.

[2] With the advantage of 15 years of technological advancement

3.1.2 Particle-in-Cell Simulations

The particle-in-cell model is an alternative discretisation of the full Vlasov-Maxwell system. While the quantities involved in evolving Maxwell's equations (i.e. the fields themselves and the current densities generated by the plasma) are still considered in an Eulerian frame, each particle species' distribution function is discretised in a Lagrangian frame. Rather than modelling the full six-dimensional phase-space, the distribution function is represented as a sum over 'macroparticles'. These macroparticles act as clumps of particle density with a defined spatial distribution or 'shape function' $S(\boldsymbol{x} - \boldsymbol{x}_p)$ of finite extent (centred at the macroparticle position $\boldsymbol{x}_p$), well-defined momentum $\boldsymbol{p}_p$ and weight (proportional to the number of particles they represent) w_p:

$$f_s(\boldsymbol{x}, \boldsymbol{p}) = \sum_p w_p S(\boldsymbol{x} - \boldsymbol{x}_p)\delta(\boldsymbol{p} - \boldsymbol{p}_p). \tag{3.3}$$

The shape function is chosen to be some symmetric, compactly supported distribution centred at $\boldsymbol{x}_p$, such that

$$\iiint S(\boldsymbol{x} - \boldsymbol{x}_p)d^3\boldsymbol{x} = 1 \tag{3.4a}$$

$$\iiint \boldsymbol{x} S(\boldsymbol{x} - \boldsymbol{x}_p)d^3\boldsymbol{x} = \boldsymbol{x}_p. \tag{3.4b}$$

Commonly used shape functions are convolutions of multiple uniform distributions:

$$S(\boldsymbol{x}) = S^n(x)\ S^n(y)\ S^n(z) \tag{3.5a}$$

$$S^n(x) = S^{n-1}(x) \star S^0(x) \tag{3.5b}$$

$$S^0(x) = \begin{cases} \Delta x^{-1} & |x| \leq \Delta x/2 \\ 0 & |x| > \Delta x/2 \end{cases} \tag{3.5c}$$

Some one-dimensional shape functions $S^n(x)$ of this form are shown in Fig. 3.1

Taking the $\boldsymbol{x}$ and $\boldsymbol{p}$ moments of the Vlasov equation over a single macroparticle located at $(\boldsymbol{x}_p, \boldsymbol{p}_p)$:

$$\frac{\partial}{\partial t}(\boldsymbol{x} S(\boldsymbol{x} - \boldsymbol{x}_p)) + \boldsymbol{x}(\boldsymbol{v}_p \cdot \boldsymbol{\nabla})S(\boldsymbol{x} - \boldsymbol{x}_p) = 0 \tag{3.6}$$

$$\frac{\partial}{\partial t}(\boldsymbol{p}_p S(\boldsymbol{x} - \boldsymbol{x}_p)) + \boldsymbol{p}_p(\boldsymbol{v}_p \cdot \boldsymbol{\nabla})S(\boldsymbol{x} - \boldsymbol{x}_p) = S(\boldsymbol{x} - \boldsymbol{x}_p)\boldsymbol{F}_L(\boldsymbol{x}, \boldsymbol{p}_p) \tag{3.7}$$

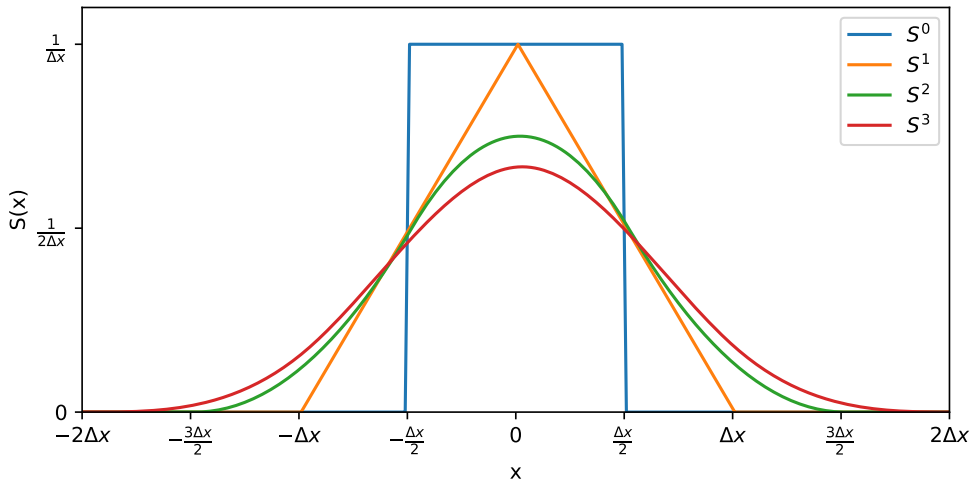


Figure 3.1: Shape functions S^n produced as n-fold convolutions of uniform distributions on $[-\frac{\Delta x}{2}, \frac{\Delta x}{2}]$. Due to properties of convolution, the n^{th} shape function S^n is compactly supported on the interval $[-(n+1)\frac{\Delta x}{2}, (n+1)\frac{\Delta x}{2}]$, reducing the cost of interpolation by limiting the number of contributing grid points to $n+1$.

Integrating also over $\boldsymbol{x}$, applying integration by parts to the second term of Eq. (3.6) and using the properties of shape functions given above, equations of motion for a macroparticle are obtained:

$$\frac{\partial \boldsymbol{x}_p}{\partial t} = \boldsymbol{v}_p \tag{3.8a}$$

$$\frac{\partial \boldsymbol{p}_p}{\partial t} = q\left(\bar{\boldsymbol{E}}(\boldsymbol{x}_p) + \boldsymbol{v}_p \times \bar{\boldsymbol{B}}(\boldsymbol{x}_p)\right), \tag{3.8b}$$

where the overbar indicates that the electric and magnetic fields have been spatially averaged by convolution with the shape function. These equations are the same as the single-particle equations of motion, supporting the interpretation of macroparticles as literal 'bigger particles'.

The finite spatial extent of macroparticle shape functions is used to couple their Lagrangian frame with the Eulerian grid used by the Maxwell solver: charge and current densities generated by the macroparticles are deposited (or 'projected') onto the grid points and electromagnetic field values which evolve on the grid are interpolated to the macroparticles' positions to determine the effect of the Lorentz force.

The four processes alluded to above take place in a loop, popularly known as "The PIC Loop" and subject to countless illustrations in talks, posters and

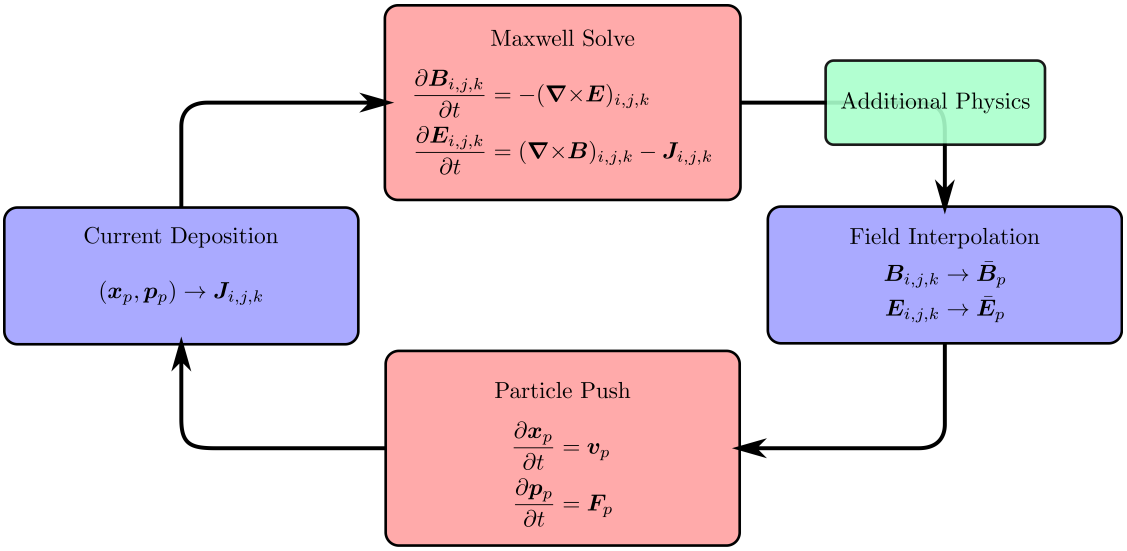


Figure 3.2: Depiction of the particle-in-cell evolution loop. Boxes in red represent physical models of the evolution of the two components of the PIC model: the particles and the fields. Those in blue represent communication between the cells of the grid (indexed i, j, k in three dimensions) and the particles (indexed p) necessary to evaluate the self-consistent dynamics of both sub-systems. In green, other parts of the simulation are represented. This includes optional additional physics processes such as collisions, ionisation, QED and radiation reaction as well as simulation processes such as diagnostic data outputs, handling a moving window and load balancing, if the user has enabled them.

theses. Fig. 3.2 is my own rendition of this loop.

Maxwell Evolution

Typically, Maxwell's equations are evolved using Yee's method, which defines staggered grids in space for the charge and current densities and electric and magnetic fields in such a way as to simplify the calculation of the curl operators in Maxwell's equations. Several more specialised solvers have been developed to offer improved performance in specific situations, for instance the method of Cowan et al. [63] was designed to eliminate numerical dispersion for beams propagating along coordinate axes, and that of Lehe et al. [64] was designed to minimise erroneous emissivity due to numerical Cherenkov radiation in laser wakefield acceleration simulations.

Field Interpolation

Once the fields have been advanced by a timestep, they must be evaluated at the particle positions. This is done according to the averaging present in Eq. (3.8b): Fields defined at grid points (i, j, k) are interpolated to the positions of each particle p using the particle's shape function:

$$\bar{\boldsymbol{E}}_p = \sum_{i,j,k} \boldsymbol{E}_{i,j,k} S(\boldsymbol{x}_p - \boldsymbol{x}_{i,j,k}) \tag{3.9}$$

An identical procedure is carried out for $\bar{\boldsymbol{B}}_p$.

Particle Push

Evolution of the macroparticle positions and momenta is conventionally undertaken using a leapfrog integrator, which considers positions (and therefore fields) offset by half-timesteps from momenta and velocities (in the following discussion of particle pushers I drop the overbar on fields and the subscript p on all quantities):

$$\frac{\boldsymbol{x}^{n+\frac{1}{2}} - \boldsymbol{x}^{n-\frac{1}{2}}}{\Delta t} = \boldsymbol{v}^n \tag{3.10a}$$

$$\frac{\boldsymbol{p}^{n+1} - \boldsymbol{p}^n}{\Delta t} = q\left(\boldsymbol{E}^{n+\frac{1}{2}} + \bar{\boldsymbol{v}}^{n+\frac{1}{2}} \times \boldsymbol{B}^{n+\frac{1}{2}}\right) \tag{3.10b}$$

Particle-in-cell codes typically include relativistic particle mechanics, so the velocity is expressed as a function of the momentum:

$$\boldsymbol{v} = \frac{\boldsymbol{p}}{\gamma m} = \frac{\boldsymbol{p}c}{\sqrt{m^2c^2 + p^2}} \tag{3.11}$$

Now, the intermediate velocity $\bar{\boldsymbol{v}}^{n+\frac{1}{2}}$ in Eq. (3.10b) must be defined. The canonical particle pusher is the Boris algorithm [65], which, despite its simplicity, shows impressive energy conservation properties [66]. It defines the intermediate velocity as:

$$\bar{\boldsymbol{v}}^{n+\frac{1}{2}} = \frac{\boldsymbol{p}^n + \boldsymbol{p}^{n+1}}{2\gamma^{n+\frac{1}{2}} m}, \tag{3.12}$$

where the intermediate Lorentz factor is calculated for the momentum after a half-electric field update, $|\boldsymbol{p}^{n+\frac{1}{2}}| = |\boldsymbol{p}^n + \frac{\Delta t}{2} q\boldsymbol{E}^{n+\frac{1}{2}}| = |\boldsymbol{p}^{n+1} - \frac{\Delta t}{2} q\boldsymbol{E}^{n+\frac{1}{2}}|$. While this definition appears to be implicit, it can be inverted to allow the algorithm to

proceed explicitly. This happens through a three-part update: A half-step electric field update, a full-step magnetic field update and finally the second half-step electric field update:

$$\boldsymbol{p}^{-} = \boldsymbol{p}^{n} + \frac{\Delta t}{2} q\boldsymbol{E}^{n+\frac{1}{2}} \tag{3.13a}$$

$$\frac{\boldsymbol{p}^{+} - \boldsymbol{p}^{-}}{\Delta t} = \frac{\boldsymbol{v}^{+} + \boldsymbol{v}^{-}}{2} \times q\boldsymbol{B}^{n+\frac{1}{2}} \tag{3.13b}$$

$$= \frac{(\boldsymbol{p}^{-} + \boldsymbol{p}^{-} \times \boldsymbol{\Omega}\frac{\Delta t}{2}) \times \boldsymbol{\Omega}}{1 + \left(\Omega\frac{\Delta t}{2}\right)^2} \tag{3.13c}$$

$$\boldsymbol{p}^{n+1} = \boldsymbol{p}^{+} + \frac{\Delta t}{2} q\boldsymbol{E}^{n+\frac{1}{2}} \tag{3.13d}$$

Here, for simplicity's sake I have introduced the vector cyclotron frequency $\boldsymbol{\Omega} = \frac{q\boldsymbol{B}}{\gamma m}$.

As with the Maxwell solvers, several more refined methods have been developed for particular purposes: the method developed by Vay [67] was designed to correctly handle systems with almost-opposing electric and magnetic components to the Lorentz force each other (which the Boris scheme does not), by using a different expression for the intermediate velocity in Eq. (3.10b):

$$\bar{\boldsymbol{v}}^{n+\frac{1}{2}} = \frac{1}{2}\left(\frac{\boldsymbol{p}^{n}}{\gamma^{n} m} + \frac{\boldsymbol{p}^{n+1}}{\gamma^{n+1} m}\right) = \frac{\boldsymbol{v}^{n} + \boldsymbol{v}^{n+1}}{2}, \tag{3.14}$$

i.e. the intermediate velocity is the arithmetic mean of the initial and final velocity.

Later, Higuera and Cary [68] developed a slightly different scheme, which also correctly computes this drift, but unlike the method of Vay also retains the Boris scheme's conservation of phase-space volume. This method uses a third form of the intermediate velocity:

$$\bar{\boldsymbol{v}}^{n+\frac{1}{2}} = \boldsymbol{v}\left(\frac{\boldsymbol{p}^{n} + \boldsymbol{p}^{n+1}}{2}\right) = \frac{\bar{\boldsymbol{p}}^{n+\frac{1}{2}}}{m\gamma\left(\bar{\boldsymbol{p}}^{n+\frac{1}{2}}\right)} \tag{3.15a}$$

$$\bar{\boldsymbol{p}}^{n+\frac{1}{2}} = \frac{\boldsymbol{p}^{n} + \boldsymbol{p}^{n+1}}{2}, \tag{3.15b}$$

i.e. the intermediate velocity is calculated using the arithmetic mean of the initial and final momentum. It is clear from the three definitions (3.12), (3.14) and (3.15) that the three integrators coincide in the non-relativistic limit.

Explicit algorithms for advancing the equations of motion using each of these modified schemes (analogous to Eqs. (3.13)) are given in their respective publications ([67] and [68]) and a comprehensive comparison of the three schemes (versus an implicit integrator and the analytic solution, for simple field configurations) by Ripperda et al. [69].

Current Deposition

Once the particles have moved, the currents associated with that motion must be deposited back onto the grid for use in the next step's Maxwell solve. While this can be done simply by using the particle shape function and velocity, this method is known not to conserve charge as it does not exactly satisfy the continuity equation [70]. Thankfully, a charge-conserving scheme for current deposition (attributed to Esirkepov [70]) does exist. The exact charge conservation of this scheme ensures that the electrostatic Poisson equation is conserved by the simulation loop, so that it remains satisfied (as long as it was satisfied by the initial conditions).

Extended PIC Physics

Many particle-in-cell simulation frameworks provide additional physics modules which extend the modelling capabilities beyond the limits of the collisionless Vlasov equation. Collisions are supported by many, often including the possibility of collisional ionisation of atoms [71]. Other ionisation mechanisms such as tunneling and barrier-suppression are supported by some particle-in-cell codes [71, 72]. Various quantum electrodynamics effects are supported by some codes, including 'macro-photons' which allow arbitrarily energetic photons to be represented within the simulation[3] and therefore electron-positron annihilation processes (which produce MeV-order gamma rays) as well as multiphoton Breit-Wheeler pair production by macrophotons. Other nonlinear QED effects such as the trident process are supported by codes which place a particular emphasis on QED modelling, such as `EPOCH` [72].

[3] In the absence of macro-photons the expressible frequencies of light are limited by the grid resolution.

Limitations

The elimination of the momentum grid can make the memory consumption of the particle-in-cell approach far more reasonable than the fully-Eulerian Vlasov method described in §3.1.1. However, this improvement comes with significant trade-offs in terms of statistical noise. Reducing noise in particle-in-cell simulations requires increasing the number of macroparticles used to discretise the distribution function, but this also leads to an increase in memory consumption. Finding a 'happy medium' between numerical noise and memory consumption is often possible however, and good results can often be obtained with far less severe memory requirements than Vlasov simulations.

Unfortunately, not all problems can be solved through superior particle statistics. As an example consider a basic feature of real physics: conservation laws. While Esirkepov current deposition [70] yields exact conservation of charge, it is not in general possible to simultaneously conserve both energy and momentum using an explicit particle pusher [62]. In practice momentum conservation is usually chosen, with the result that initially cold simulated plasma is subject to unphysical numerical heating [73].

The physical extent of plasma which is practical to model is also limited with the PIC method, due to strict limits on minimum spatial and temporal resolution which result from the use of explicit methods to solve the particle motion and field evolution. For example, any light in the simulation will only propagate if its wavelength is well-resolved by the grid, and this also restricts the maximum temporal step size via the Courant–Friedrichs–Lewy (CFL) condition; with grid resolution Δx_j in N dimensions this condition is:

$$\frac{1}{c^2 \Delta t^2} \geq \sum_{j=1}^{N} \frac{1}{\Delta x_j^2} \tag{3.16}$$

This obviously implies that $c\Delta t < \Delta x_j$ - the temporal resolution must be finer than the spatial resolution in any of the simulation's dimensions. The CFL condition

may be more concisely expressed in terms of a 'grid spatial frequency' $K_j = \frac{2\pi}{\Delta x_j}$ and a 'timestep frequency' $\Omega = \frac{2\pi}{\Delta t}$, in terms of which the condition becomes $\Omega \geq c|\boldsymbol{K}|$

The aforementioned numerical heating arises from a ficticious interaction between macroparticles and the grid on which the electromagnetic field is defined. This numerical instability has the effect of an unphysical injection of energy into the system. The numerical grid-particle instability is particularly strong in regions where the Debye length of the plasma is not resolved adequately by the grid [73]. When this is the case the temperature rises exponentially in time until the Debye length becomes approximately equal to the grid spacing, and linearly after that point. Therefore it is necessary that the grid resolves the Debye length of the plasma:

$$|\boldsymbol{K}| > \left(k_D = \frac{\omega_p}{u_{\text{th}}} \right), \tag{3.17}$$

where the 'thermal speed' (which is not a physical velocity and so may exceed c) is defined by $u_{\text{th}}^2 = k_B T_e / m_e$. This condition limits the time-step via the electromagnetic CFL condition Eq. (3.16):

$$\Omega > \omega_p \frac{c}{u_{\text{th}}}, \tag{3.18}$$

For relativistically hot plasma ($u_{\text{th}} > c$), the timestep is more strictly limited by the necessity of resolving the plasma frequency:

$$\Omega > \omega_p, \tag{3.19}$$

In either case, simulations of supercritical plasma will be limited by the plasma frequency rather than the laser frequency. Plasma that is both cold and dense (and which is needed to remain cold) presents a particular challenge to particle-in-cell methods, and it has been argued [57] that the Vlasov method can actually provide better computational efficiency for plasmas of this type as they do not exhibit a grid-plasma instability and can therefore use larger grid spacing than an equivalent particle-in-cell simulation.

Some examples of general-purpose PIC simulation codes are `OSIRIS` [74], `EPOCH` [72] and `Smilei` [71].

3.1.3 Fluid Model Simulations

By taking moments of powers of particle velocity in the Vlasov equation—as in §2.3.2—it is possible to derive a set of governing equations for each species, in a 'fluid' description. Often, the plasma is assumed to maintain quasi-neutrality and the electron inertia is ignored, resulting in ideal magnetohydrodynamics (MHD). Within this model, magnetic fields are advected with the fluid flow in a way consistent with Maxwell's equations.

A particularly simple case is ideal MHD equipped with an initial condition of zero magnetic field. The equations of ideal MHD are not capable of sourcing mangetic fields spontaneously (note that the RHS of Eqs. (2.51) is linear in $\boldsymbol{B}$), so the magnetic field remains zero for the entire simulation. In this $\boldsymbol{B} \to 0$ limit, the MHD equations reduce to compressible Navier-Stokes equations. Therefore, unmagnetised plasma behaving according to the assumptions of MHD can in principle be modelled by standard gas dynamics simulation packages.

Extensions to the ideal MHD system are available as well as simplifications [75, 76]. Phenomena which can be added to MHD include non-zero resistivity (a consequence of inter-species collisions which acts as a diffusive term converting magnetic to thermal energy), the Hall current (which modifies the magnetic field dynamics and is commonly used in simulations of magnetic fusion plasma) and electron pressure gradients (which lead to the 'Biermann battery' spontaneous magnetic field source term). It is also possible to include multiple ion species in a hydrodynamic system, and different materials with equations of state which differ from the adiabatic ideal gas equation of state (EOS) used throughout Chapter 2, including tabulated equations of state based on experimental measurements. Codes may also model the evolution of a black-body radiation field, model separate ion, electron and radiation temperatures (each with a corresponding pressure term and equation of state) and deposit energy in the simulation via ray-traced incident laser irradiation in the geometric optics limit.

When the so-called radiation-hydrodynamic model is employed—tracking evolution of a radiation field 'fluid' as well as fluids representing matter—it is important to

understand the interaction between the two types of fluid. It is therefore important to model the opacity of each matter species present. Similarly to equations of state, these opacity functions are typically complicated and simple analytic models (such as collisional inverse-bremßtrahlung absorption [77]) are not always appropriate. For this reason opacities are often provided as a table of values based either on experimental results or numerical calculations [78].

3.2 Semi-Analytic and Smaller-Scale Modelling

As well as these large-scale simulations of the dynamics of entire plasma systems, it is often also useful to be able to model certain smaller sub-systems. These smaller-scale simulations can either provide useful initial conditions for large-scale simulations (such as the predicted beam profile of an aberrated laser focus) or be used to simulate real-world diagnostics applied to the simulated plasma to validate the simulation performance against real experiments.

An example of the former is presented in Appendix B: a laser beam propagation Python package, `pyrabola`, which I designed to be fast enough for use in the training loop of a reinforcement learning system (itself developed by my fellow DPhil student Marko Mayr) while still providing a realistic estimate of the beam profile produced by a misaligned parabolic mirror at an arbitrary location in space.

The latter case is a method I use in Chapter 5 to complement the direct analysis of diagnostic images in regimes where the direct analysis fails. For example a beam propagation code to model the optical interferometry diagnostic in §5.4.1 where probe light is lost due to diffraction and the proton radiography imaging code I present in Appendix A which produces noise-free results equivalent to the limiting case of Monte-Carlo approaches and works even in the caustic-forming regime which challenges other non-stochastic methods for rendering proton radiographs.

Another application of these smaller-scale methods is to evaluate reduced-physics models outside of the regime in which they are analytically solvable (e.g. a nonlinear plasma wakefield outside of the quasistatic approximation). I developed a method of this type which was used in Mayr et al. [34].

4
Ion Acceleration

Contents

4.1 Target Normal Sheath Acceleration 58
4.1.1 Modelling TNSA 58
4.1.2 TNSA-produced beams 61
4.2 Additional Acceleration Mechanisms 62
4.2.1 Radiation Pressure Acceleration 62
4.2.2 Collisionless Shock Acceleration 65
4.3 Experiment: Collisionless Shock Acceleration in Foam Targets 81
4.3.1 Experimental Conditions 81
4.4 Simulation Post-Mortem 85
4.5 Summary 88

Ion acceleration is a major branch of research in plasma physics. In addition to their potential to provide fast, isochoric heating of fusion plasmas for 'fast ignition' applications [9, 79], energetic beams of ions have been utilised to probe plasma dynamics (this will be the focus of Chapter 5 and Chapter 6), and sufficiently high-quality ion beams even have therapeutic applications, for example in radiotherapy for cancer treatment [11].

While initial interest in medical applications of laser-driven ion sources was driven by the promise of smaller, cheaper accelerators (as compared to conventional accelerators such as cyclotrons and synchrotrons [80]), more recent research has

suggested benefits to the ultra-high dose rates attainable using laser-driven ion sources. These sources seem to have reduced effect on healthy tissues around a tumor, while still stimulating the desired effect in cancerous tissues [81]. The cause of these benefits, known as the FLASH effect, is still under investigation.

4.1 Target Normal Sheath Acceleration

Target Normal Sheath Acceleration (TNSA)—first observed in experiments in the year 2000 by Clark et al. [82], Snavely et al. [83] and Maksimchuk et al. [84], and explained theoretically by Wilks et al. [85] a year later—is a very simple (experimentally, at least) mechanism for laser-driven ion acceleration. TNSA can be achieved with a wide range of laser and target parameters, an advantage compared to some more recently-developed methods (which will be discussed shortly), which are only applicable in a narrow range of plasma conditions and laser parameters.

TNSA takes place when a short (order 10 ps), relatively energetic (order 10 J) laser pulse impinges upon a thin (1 µm to 10 µm) metal foil. Laser pulse energy is absorbed by electrons in the metal target, heating them to MeV energies. These 'hot electrons' are accelerated through the target in the laser-forward direction, and eventually breach the rear surface of the thin target. This 'cloud' of un-neutralised negative space charge adjacent to the target's rear surface sources very strong quasistatic electric fields directed along the surface normal, which accelerate protons present in hydrocarbon impurities on the surface to high energy.

4.1.1 Modelling TNSA

The electrostatic fields generated in the rear-surface sheath are estimated by equating the 'temperature' of 'hot' electrons leaving the rear of the target with the potential difference needed to stop electrons of this energy:

$$eE_{\text{sheath}} \approx \frac{k_B T_{\text{hot}}}{L_{\text{sheath}}}. \tag{4.1}$$

An ion of charge Ze, 'yanked' from the target's rear surface by this field, can be expected to obtain energy of approximately:

$$\mathscr{E} \approx ZeE_{\text{sheath}}L_{\text{sheath}} = Zk_BT_{\text{hot}} \tag{4.2}$$

as it is accelerated down the field gradient [86]; employing the Wilks et al. [87] 'ponderomotive' scaling for hot electrons, $k_BT_{\text{hot}} \approx (\gamma - 1)m_ec^2$ with γ given by

$$\gamma^2 = 1 + \frac{a_0^2}{2}, \tag{4.3}$$

this expected energy scales $\propto I^{\frac{1}{2}}$ at high intensity.

More advanced modelling of the TNSA process can be undertaken by considering an electron population 'split' into 'hot' and 'cold' populations and the presence of 'heavy' and 'light' ion species. The treatment of these plasma components then separates a few common models:

Quasistatic Model

In the basic quasistatic model, the light ions are accelerated by the sheath field, but are assumed to be of sufficiently low density that they do not perturb the field. The field is sourced by charge separation between the much denser population of immobile heavy ions and isothermal, Boltzmann-distributed hot electrons, while the cold electrons are also modeled as stationary [86].

While this model is simple and can be solved (see for example Macchi, Borghesi, and Passoni [86] for an outline of the solution), the assumption that hot electrons are Boltzmann leads to the unfortunate prediction of infinite maximum accelerated ion energy.

Numerous enhancements and refinements to the quasistatic approach have (among other aims) attempted to rectify this issue, including introduction of a self-consistent hot electron cut-off energy (above which the electrons are assumed to escape the system entirely) [88, 89] or by limiting the longitudinal extent of the sheath [90, 91], or inclusion of the cold electrons' dynamics [92]. Further detailed discusssion can be found—for example—in the excellent review of Macchi, Borghesi, and Passoni [86].

Plasma Expansion Model

An example of a dynamic model for TNSA is that of plasma expansion into vacuum. In its simplest form, a one-dimensional quasineutral fluid with one electron and one ion species expands into vacuum. In this case, the self-similar solution of Gurevich, Pariiskaya, and Pitaevskii [93] applies, giving expressions for the ion density and velocity in regions $x + c_s t > 0$ [86]:

$$n_e = Z n_i = n_0 \exp\left(-\frac{x + c_s t}{L_{\text{sheath}}}\right), \tag{4.4}$$

$$v_i = c_s + \frac{x}{t} \tag{4.5}$$

for sound speed c_s related to the local density scale-length by $L_{\text{sheath}} = c_s t$. This solution is cut off at the point where $L_{\text{sheath}} = \lambda_D$; making explicit the dependence of the local Debye length on the local plasma density:

$$c_s t = \lambda_D(x,t) = \lambda_{D0}\sqrt{\frac{n_0}{n_e}} \tag{4.6}$$

$$= \lambda_{D0} \exp\left(\frac{x + c_s t}{2 c_s t}\right) \tag{4.7}$$

$$x_f(t) = c_s t\left[2\ln\left(\frac{c_s t}{\lambda_{D0}}\right) - 1\right]. \tag{4.8}$$

This is taken to be the position of the ion expansion front. The corresponding velocity is:

$$v_f(t) = c_s + \frac{x(t)}{t} = 2 c_s \ln \omega_{pi} t, \tag{4.9}$$

introducing the ion plasma frequency $\omega_{pi} = \frac{c_s}{\lambda_{D0}}$. A slight modification of the model's predictions to better fit simulation predictions was given by Mora [94], the result being

$$v_f(t) = 2 c_s \ln\left(\tau + \sqrt{1 + \tau^2}\right), \tag{4.10}$$

where $\tau = \omega_{pi} t / \sqrt{2\exp(1)}$, which for total acceleration time t_{acc} and corresponding τ_{acc} predicts maxmimum ion energy of

$$\mathcal{E}_{\max} = 2 k_B T_{\text{hot}}\left[\ln\left(\tau_{\text{acc}} + \sqrt{1 + \tau_{\text{acc}}^2}\right)\right]^2 \tag{4.11}$$

Now, employing the Wilks et al. [87] 'ponderomotive' scaling for hot electrons, $k_B T_{\text{hot}} \approx (\gamma - 1) m_e c^2$ with γ given by

$$\gamma^2 = 1 + \frac{a_0^2}{2}, \tag{4.12}$$

it is clear that—given the requirement of high intensity—this model predicts that the TNSA energy cut-off should scale with $a_0 \propto I^{\frac{1}{2}}$.

This plasma expansion model has been generalised to multiple ion species in a few studies, with a good example in the context of TNSA being the work of [95].

4.1.2 TNSA-produced beams

The quality of proton beams produced through TNSA was found to be significantly superior to the beams protons produced in earlier experiments conducted at lower laser intensity [86]. TNSA-produced proton beams were found to be highly laminar, directed orthogonally to the target's rear surface, and with sharp angular distribution[1] of energy-dependent opening half-angle of 10° to 20°[83]. Their high degree of laminarity corresponds to a low transverse emittance, quoted in the range from 0.1π mm mrad [97] down to 0.004 mm mrad [98] in more precise experiments.

The longitudinal emittance of TNSA-produced beams is also low, due to their very short duration and correlations between acceleration time and energy [99]. However, a focus on emittance in the longitudinal case could be somewhat misleading: the energy spread of TNSA beams is of the order of 100%, which places limits on their applicability even with superficially small longitudinal emittance values.

It is also interesting to consider the effect of target properties on the beams which can be produced by TNSA. As noted earlier, while use of conductive (i.e. metal) targets produces beams with sharp circular cross-section, beams produced from insulating targets (e.g plastics) display more structured cross-sections. This is attributed to hot electron transport from the laser-irradiated face of the target (where hot electrons are born) to the opposite face (where they drive formation of the sheath field and reflux back into the target): this transport is easy and takes

[1] at least in the case of conductive targets - when the target is insulating more structure is present in the beam [96].

place with minimal resistive losses in conductive material, but introduces significant perturbations to the hot electron population in insulators [96].

Once hot electrons have reached the far side of the target and been reflected by the sheath field, they propagate back towards the laser-irradiated side. In thin targets—and especially those which are conductive—a cycle of recirculation forms due to the minimal energy losses involved in a single round-trip. Experimental and simulation results [100, 101] show that this recircultion of electrons within the target strengthens the accelerating sheath field, increasing the cut-off energy observed in beams accelerated from thinner targets in proportion to the ratio of laser pulse duration to target thickness [101].

Finally, the optimal situation for TNSA is usually considered to be an initially cold foil of the target material. Presence of a pre-formed plasma was found by Mackinnon et al. [102] to lead to a reduction in maximum energy, and in fact to the beam being 'essentially destroyed', a fact that the experiment I present in §4.3 sought to take advantage of to eliminate background TNSA signal in ion acceleration.

4.2 Additional Acceleration Mechanisms

The particularly broad, 'thermal' spectrum produced by TNSA is undesirable in applications that call for protons of very narrow energy-spread. This limitation of TNSA has driven recent development of alternative acceleration mechanisms, which often aim to produce proton beams (or beams of heavier ions) with quasi-monoenergetic spectra.

4.2.1 Radiation Pressure Acceleration

Radiation Pressure Acceleration (RPA) is predicated on the concept of a laser pulse acting as a sort of 'piston', pushing the critical density surface of a plasma through the solid-density target via its radiation pressure. It is separated into two principal regimes: ions are accelerated by hole-boring RPA in thick targets, or in the light-sail

regime for thin targets. A key parameter [103] characterising the transition between these regimes is the dimensionless areal density-laser amplitude ratio:

$$\xi = \pi \frac{n_e}{n_c} \frac{d}{\lambda_L} a_0^{-1}. \tag{4.13}$$

For targets which are very thin ($\xi \ll 1$), electrons are expelled from the target and the ions disintegrate via Coulomb explosion. The optimal thickness for Light Sail RPA is found at $\xi = 1$, while thicker targets ($\xi \gg 1$) undergo hole-boring RPA. The latter two mechanisms are described below[2].

In both of the below regimes of radiation pressure acceleration, the laser pulse's polarisation is very important. Circularly-polarised beams produce only a steady ponderomotive force: the $\boldsymbol{J}\times\boldsymbol{B}$ force has no oscillating component in this case (see §2.5.3). This is important as it eliminates the $\boldsymbol{J}\times\boldsymbol{B}$ electron heating mechanism, limiting the electron temperature and preventing formation of a sheath field. Sheath acceleration would dominate RPA in the presence of a large sheath field, making RPA much easier to observe using circularly-polarised beams.

Hole Boring RPA

For a pulse with intensity I, the pressure exerted on the critical density surface as a result of reflection of light is $P = 2I/c^3$. This pressure pushes the electrons at the critical surface into the bulk of the plasma, and the charge separation between these electrons and the ions left behind sets up a capacitor-like electric field. The critical surface therefore propagates into the target as an electrostatic shock. Ions in the bulk are reflected by the shock, and the velocity of the shock front is found by equating the radiation pressure imparted by the laser pulse to the momentum flux of these ions,

$$\frac{2I}{c} = (m_i v_h)(2 n_i v_h) \tag{4.14}$$

$$\left(\frac{v_h}{c}\right)^2 = \frac{I}{m_i n_i c^3} = \Xi. \tag{4.15}$$

[2]Coulomb explosion is not best understood as a form of RPA and is not covered here, but good summaries can be found in, for example, Macchi [42]

[3]The mechanism by which this pressure acts is the ponderomotive force, §2.5.4 [86]

Here I have denoted the speed of the shock as v_h, the 'hole-boring velocity'. Ions upstream of the shock reflect specularly from it, being left with (in the non-relativistic limit, and assuming initially cold, static ions) twice the hole-boring velocity directed away from the shock front. The energy of these ions is therefore [104]

$$\mathcal{E} = \frac{1}{2} m_i (2v_h)^2 = \frac{2I}{n_i c} = 2 m_i c^2 \Xi \tag{4.16}$$

This derivation is non-relativistic; a relativistically-correct version was presented by Robinson et al. [104]. The final ion energy is corrected to

$$\mathcal{E} = m_i c^2 \frac{2\Xi}{1 + 2\sqrt{\Xi}}. \tag{4.17}$$

This scaling (and the prospect for monoenergetic ion beams) looks very promising, but unfortunately it has been predicted that the RPA mechanism will be dominated by TNSA for pulses with near-optical wavelength ($\sim 1\,\mu\text{m}$) unless the intensity is very high—in excess of $1 \times 10^{22}\,\text{W cm}^{-2}$ for linearly polarised light according to Esirkepov et al. [105], or $1 \times 10^{21}\,\text{W cm}^{-2}$ for circular polarisation according to Robinson et al. [106]. Hole-boring RPA has been demonstrated successfully using pulses of 10 µm light generated by a CO_2 laser, as reported by Palmer et al. [107] and [108]. Achieving hole-boring RPA with micron-class Nd:glass or Ti:sapphire lasers has proven more difficult owing to their higher critical density and the negative scaling of ion energy with density, but the doctoral thesis of Martin [109] reports on hole-boring acceleration from cryogenic hydrogen ribbon targets at the Vulcan laser facility (an Nd:glass laser operating at wavelengths near one micron).

The Light-Sail Regime

If a target thickness is of the order of the skin depth ($\xi \approx 1$), radiation pressure will accelerate the *entire* target. Within this regime the target is modeled as a relativistic flying mirror, and is accelerated for as long as the laser irradiation is maintained and the target remains intact. The energies expected from light-sail RPA are higher than those expected from hole-boring, as there is no unaccelerated background plasma to screen the accelerating field [86].

Modelling the target as an ideal mirror with one-dimensional dynamics (i.e. one that is planar and rigid) and areal density $\sigma = \rho\ell$ for mass density ρ and thickness ℓ, the relativistic equation of motion is easy to write down [86]:

$$\sigma \frac{d\gamma\beta c}{dt} = \frac{2I(t - x/c)}{c} \frac{1-\beta}{1+\beta}. \tag{4.18}$$

Key features of this equation are that the laser intensity should be evaluated at the retarded time $t_{\mathrm{ret}} = t - x/c$ [110] and that the intensity in the mirror's rest frame is reduced from the lab-frame intensity by the relativistic Doppler factor on the right-hand side.

Following the derivation of Simmons and McInnes [110], I define the dimensionless pulse fluence parameter $\mathcal{F}(t_{\mathrm{ret}})$ as a function of the retarded time:

$$\mathcal{F}(t_{\mathrm{ret}}) = \int_0^{t_{\mathrm{ret}}} \frac{2I(t')}{\sigma c^2} dt'. \tag{4.19}$$

The sail velocity and energy are then given by [110]

$$\beta(t_{\mathrm{ret}}) = \frac{[1 + \mathcal{F}(t_{\mathrm{ret}})] - 1}{[1 + \mathcal{F}(t_{\mathrm{ret}})] + 1} \tag{4.20}$$

$$E(t_{\mathrm{ret}}) = m_p c^2 \frac{\mathcal{F}^2(t_{\mathrm{ret}})}{2\left[\mathcal{F}(t_{\mathrm{ret}}) + 1\right]} \tag{4.21}$$

Rayleigh-Taylor-like instabilities have been found to limit light-sail RPA as the acceleration proceeds [111].

4.2.2 Collisionless Shock Acceleration

In collisionless shock acceleration, ions are accelerated by reflection from an electrostatic structure moving through the plasma at supersonic speed - such a structure is a collisionless electrostatic shock, from which the mechanism derives its name. Macchi, Nindrayog, and Pegoraro [112] consider such structures being driven directly by a laser pulse, and differentiate between acceleration by a solitary wave (or 'soliton')—which they compare to hole-boring RPA—and formation of a 'true' shock wave, which is only possible in the presence of non-zero initial ion temperature.

In contrast to this directly laser-driven regime of collisionless shock acceleration, Fiuza et al. [113] developed a model of a collisionless shock launched from the

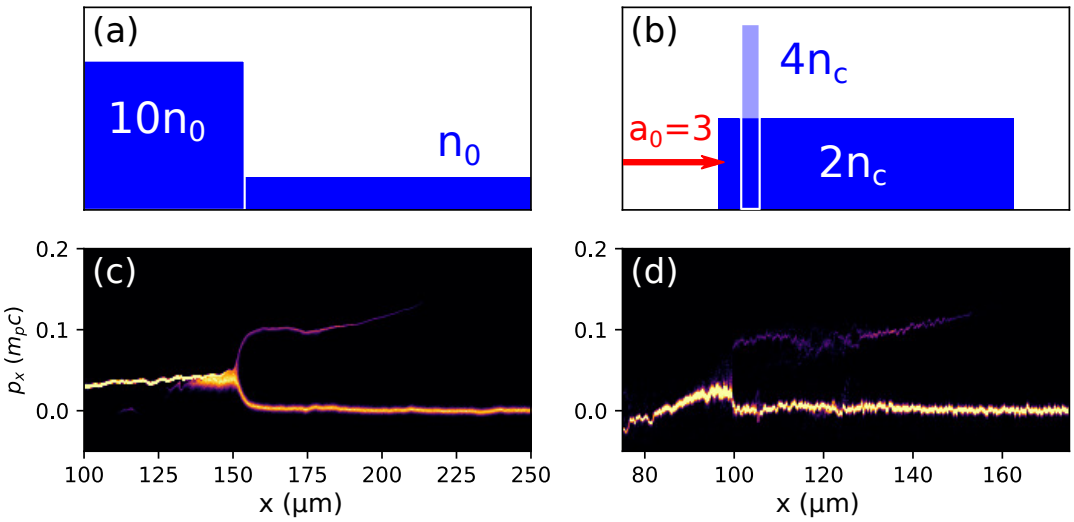


Figure 4.1: Results of two simulations that exhibit collisionless shock acceleration of protons. (a) Sketch of the "two-slab" simulation set-up. Two slabs with density ratio of $\Gamma = 10$ and electron temperature 1.5 MeV form a collisionless shock at their interface, as electrons at high pressure from the denser slab stream into the less dense one. Ions initially contained in the lower density slab are reflected from this shock structure, forming the characteristic proton phase-space structure shown in (b). (c) Sketch of "laser-driven" simulation. A single slab of twice critical density and initial temperature 5 keV is irradiated by a laser pulse of duration 2 ps and $a_0 = 3$, heating the electrons to MeV temperature and producing a region of density four times critical density. A similar (though noisier) phase-space structure results, and is shown in (d). The shock mach numbers and ion-electron temperature ratios for these two simulations are plotted in the parameter space compatible with CSA in figure Fig. 4.5

interface between two slabs of plasma with differing density and electron temperature, and cold ions. They find that for slabs of comparable electron temperature, a density ratio $\Gamma \geq 4$ is sufficient to drive a shock that can accelerate ions. Further, they argue that the required 'two-slab' density structure can be produced in laser-plasma experiments as the result of ponderomotive pressure on the critical density surface of a plasma profile with initially moderately supercritical peak density. Within this model, the predicted energy of the quasi-monoenergetic accelerated ions scales positively with the bulk electron temperature produced during the absorption of the laser pulse. Fig. 4.1(a-b) shows an example of the simple (in simulation at least; less so to implement experimentally) situation of two slabs of plasma with density ratio $\Gamma = 10$, electron temperature 1.5 MeV and ion temperature 5 keV. Fig. 4.1(c-d) demonstrates the laser-driven version of the model, in which a single slab of plasma

of 1.5 times critical density is irradiated by a 2 ps laser pulse with $a_0 = 3$. This heats the electrons to MeV temperatures while simultaneously, through radiation pressure applied to the front surface of the plasma, 'piling up' density around the front surface. Both models produce a characteristic structure in ion phase-space, indicative of partial ion reflection from the upstream plasma.

Collisionless Shock Acceleration (CSA) of this kind was first observed by Haberberger et al. [114] in the experiment that inspired the analysis of Fiuza et al. [113]. Further theoretical development was undertaken by Cairns et al. [115, 116] and in the DPhil thesis of Savin [29]. It requires homogeneous plasma with slightly supercritical electron density of approximately $1 \leq n_e/n_c \leq 10$; notably, for laser light of 1 µm wavelength, gas targets significantly undershoot this range and traditional solid targets significantly overshoot. For this reason much of the early research into CSA was carried out using CO_2 lasers, which operate at 10 µm and whose critical density is therefore a factor of 100 lower than that of micron-wavelength lasers. The first demonstration of CSA with Nd-glass lasers was by Pak et al. [117] at the Titan laser facility, using a preformed plasma ablated from a thin film to achieve the desired sub-solid density plasma.

Subsequently I participated in an experiment at the Vulcan laser facility using low-density plastic foams to produce correspondingly low-density plasma. The experiment was essentially designed to drive the collisionless shock acceleration mechanism shown in Fig. 4.1(c) is described in detail in §4.3. First though, I will describe the theoretical foundation of CSA and developments I have made in evaluating the phase-space of parameters for which the mechanism is applicable. This work was started by Cairns et al. [115, 116] and developed significantly by Savin [29], but here I provide significant further development.

The following derivation begins with a summary of the theoretical basis for a collisionless shock, as presented by Cairns et al. [115, 116] (though some of the derivation is conducted differently). In §4.2.2, a systematic method for determining sets of parameters which can self-consistently describe a collisionless shock is presented for the first time. This method is a significant deveopment to the work

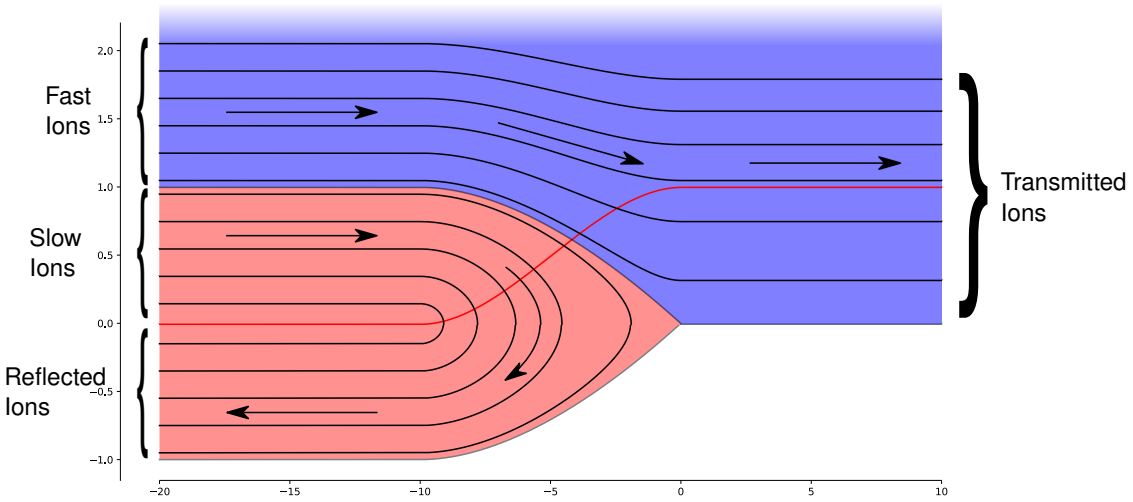


Figure 4.2: Ion trajectories in the shock frame. The red line represents a shock potential, which falls to zero at the left and rises to a maximum at the right. There is a bifurcation in ion phase-space which separates initially 'slow' ions (velocity smaller than $\sqrt{2m_i\phi_{\max}}$) from initially 'fast' ions. The former are reflected from the shock at some point dependent on their upstream velocity and the latter are slowed, but not reflected by the shock. On the phase diagram the slow and reflected ions are shaded in red and the fast and transmitted ions are shaded in blue.

presented by Savin [29]. The theory presented here is non-relativistic as this was deemed appropriate for the conditions of this experiment, but a relativistic treatment of some of the physics has been presented by Stockem et al. [118].

The Electrostatic Shock

An electrostatic shock is characterised by an electric potential profile $\phi(x)$ which rises monotonically[4], from zero far upstream of the shock to a maximum value $\phi_{\max}$ at the shock front. This structure moves through plasma at a given shock speed, V, which is taken to be in the negative x direction (i.e. a velocity of $-V$). The negative x-direction is then the upstream direction, and the positive x-direction is downstream. Ions with velocity v_0 in the rest frame of the shock have velocity $v_0 - V$ in the plasma rest frame, implying that the initial distribution of ion velocity in the shock frame, far upstream of the shock (where the electrostatic potential $\phi(-\infty) = 0$) is given by

$$f(v_0, \phi = 0) = n_0 \sqrt{\frac{m_i}{2\pi k_B T_i}} \exp\left(-\frac{m_i(v_0 - V)^2}{2k_B T_i}\right), \tag{4.22}$$

[4]'monotonically' is important here, as it makes it possible to use ϕ as a proxy for position x by ensuring that $\phi(x)$ is an invertible function

for an upstream ion population of mass m_i, temperature T_i and uniform density n_0. As expected this distribution drifts towards the shock front at the shock velocity V. The situation in the shock frame is shown in Fig. 4.2.

By conservation of energy, an ion's initial kinetic energy $\frac{1}{2}m_i v_0^2$ is diminished as it climbs the shock to a potential ϕ:

$$\frac{1}{2}m_i v^2 = \frac{1}{2}m_i v_0^2 - Ze\phi \tag{4.23}$$

In this expression I allow for multiply-charged ions (with charge state Z). This energy conservation equation is easily rearranged to yield the velocity an ion has left after climbing the potential barrier, and also the minimum initial velocity $v_{\min}$ required to reach a given potential ϕ

From simple energy-conservation arguments, an ion of charge state Z must have initial velocity greater than $v_{\min}(\phi) = \sqrt{2Ze\phi/m_i}$ to reach a potential ϕ. Having propagated up the shock to a potential $\phi < \phi_{\max}$, ions now have velocity $v = \sqrt{v_0^2 - v_{\min}^2}$. The steady-state form of the ion continuity equation leads to the condition

$$\frac{\partial}{\partial\phi}v(v_0,\phi)n_i(v_0,\phi) = 0. \tag{4.24}$$

Therefore the ion density at potential ϕ, including ions with initial velocity of at most $v_{\max}$ is given by:

$$n_i(\phi, v_{\max}) = n_0\sqrt{\frac{m_i}{2\pi k_B T_i}}\int_{v_{\min}}^{v_{\max}} \frac{v_0 dv_0}{\sqrt{v_0^2 - v_{\min}^2}}\exp\left(-\frac{m_i(v_0 - V)^2}{2k_B T_i}\right) \tag{4.25}$$

In order to improve the legibility of this integral (and rid myself of a free parameter), I now normalise quantities in this equation: number densities according to $\tilde{n}_i = n_i/n_0$, potentials according to $\varphi = \phi/\phi_{\max}$, and velocities according to $u^2 = m_i v^2/(2Ze\phi_{\max})$. The ion temperature is represented as the normalised inverse temperature[5] $\beta_i = Ze\phi_{\max}/(k_B T_i)$.

[5]The notation I use here is intended to evoke the use of the 'thermodynamic beta' in statistical physics to represent the Lagrange multiplier which ends up being the inverse temperature or 'coldness' factor $(k_B T)^{-1}$. It does *not* represent a relativistic normalised velocity as this analysis is wholly non-relativistic.

$$\tilde{n}_i(\varphi, u_{\max}) = b \int_{\sqrt{\varphi}}^{u_{\max}} \frac{u}{\sqrt{u^2 - \varphi}} e^{-\beta_i (u-U)^2} du \tag{4.26}$$

Here an additional constant $b = (\beta_i/\pi)^{1/2}$ has been introduced for the sake of brevity. This formulation differs from previous treatments which eliminate ion temperature and express integrals over the residual—rather than the initial—velocity [29, 115, 116].

Electrons are parametrised by the electron temperature, which is normalised similarly to the ion temperature[6]: $\beta_e = e\phi_{\max}/(k_B T_e)$, and are distributed according to a Boltzmann relation[7]:

$$\begin{aligned} \tilde{n}_e(\varphi, u_{\max}) &= Z n_i(0, u_{\max}) e^{\varphi \beta_e} \\ &= Z b e^{\varphi \beta_e} \int_0^{u_{\max}} e^{-\beta_i (u-U)^2} du \end{aligned} \tag{4.27}$$

At a given φ, two components contribute to the particle distributions: one component represents inflowing ions, for which $u_{\max} = \infty$; the other contains ions which have reflected from the shock and are travelling back upstream, whose $u_{\max} = 1$ such that all ions in this component have insufficient energy to pass through the shock. Therefore the total charge density at potential φ (normalised to Zen_0) is:

$$\rho(\varphi) = \left(\tilde{n}_i(\varphi, \infty) + \tilde{n}_i(\varphi, 1)\right) - Z^{-1} \left(\tilde{n}_e(\varphi, \infty) + \tilde{n}_e(\varphi, 1)\right). \tag{4.28}$$

The Sagdeev Potential

The Sagdeev potential Φ is a pseudo-potential introduced by Sagdeev [119] in his analysis of shock waves in plasma, which has also proved useful in the analysis of soliton solutions to the Korteweg–De Vries equation [120]. It has units of squared electric field and I normalise it using a base value $\Phi_0 = \phi_m^2/x_0^2$, where $x_0^2 = v_\phi^2/2\Omega_i^2$, $v_\phi^2 = 2Ze\phi_{\max}/m_i$ and $\Omega_i^2 = (n_0 Z^2 e^2)/(m_i \epsilon_0)$ is the standard ion plasma frequency far upstream.

[6]The temperature ratio emphasised by other treatments of this problem [29, 115, 116] is therefore given by $T = \beta_i/\beta_e$

[7]The use of a Boltzmann relation in a frame in which the plasma is drifting may cause you concern – it definitely caused me some. If you are concerned by this, you are correct. I will address these concerns once I have introduced the Sagdeev potential later in this section

The Sagdeev potential is defined by

$$\frac{d^2\varphi}{dx^2} = -\frac{d\Phi}{d\varphi}, \tag{4.29}$$

an equation identical in form to that of a particle moving under the influence of a potential Φ, whose 'time' variable is x and whose 'space' variable is φ. Accordingly this equation can also be expressed in an 'energy-conservation' form[8]:

$$\frac{1}{2}\left(\frac{d\varphi}{dx}\right)^2 + \Phi(\varphi) = 0 \tag{4.30}$$

Substituting an appropriately-normalised form of the electrostatic Poisson equation into Eq. (4.29) and integrating, the Sagdeev potential takes the form

$$\Phi(\varphi) = \int_0^{\varphi} \rho(\varphi')d\varphi'. \tag{4.31}$$

By comparison with Eq. 4.28, the Sagdeev potential is therefore a combination of four terms, each of which is a double integral. By considering that the 'pseudo-particle' moving according to Eq. 4.29 starts at $\varphi = 0$, and that by inspection of Eq. 4.31 the Sagdeev potential $\Phi(0) = 0$, we can deduce that the highest value of φ that the system evolving according to Eq. 4.29 will reach—by definition $\phi_{\max}/\phi_{\max} = 1$—also fulfils $\Phi(\varphi) = 0$. Therefore the condition $\Phi(1) = 0$ closes the system and ensures self-consistent dynamics.

The Sagdeev potential must have the form of a potential well (i.e. $\Phi(\varphi) < 0$ for $0 < \varphi < 1$) to ensure the pseudoparticle is able to travel between $\varphi = 0$ and $\varphi = 1$. A necessary condition is therefore:

$$\left.\frac{d\Phi}{d\varphi}\right|_{\varphi=1} > 0. \tag{4.32}$$

Now that I have introduced the Sagdeev potential, I briefly mention that the use of a Boltzmann relation to define the electron density in Eq. (4.27) is not strictly correct. The full derivation is lengthy and would break the flow of the present one too much to include here, but in Appendix C I explore the implications and applicability of a drifting Boltzmann relation. The upshot is that a Boltzmann density profile is

[8]The electric field can therefore be expressed as $E = \pm\sqrt{-2\Phi}$

unstable in plasma which is drifting relative to the electrostatic structure, with a first-order decay rate ω (normalised to the ion plasma frequency) given by

$$\omega = \sqrt{2}\beta_e U \frac{\partial \varphi}{\partial x} \tag{4.33}$$

$$= 2\beta_e U \sqrt{-\Phi}, \tag{4.34}$$

where I have simplified using Eq. (4.30).

What remains is to find sets of parameters (U, β_i, β_e) for which $\Phi(\varphi = 1) = 0$ and $\Phi'(\varphi = 1) > 0$.

The Valid Parameter Space

Firstly, observe that in Eq. 4.27, $\tilde{n}_e$ is trivially integrable with respect to φ:

$$N_e = \int_0^1 \tilde{n}_e(\varphi, u_{\max})/Z d\varphi \tag{4.35}$$

$$= \frac{e^{\beta_e} - 1}{\beta_e} \tilde{n}_i(0, u_{\max}) \tag{4.36}$$

$$= \frac{e^{\beta_e} - 1}{\beta_e} G_{U,\beta_i}(u_{\max}) \tag{4.37}$$

Integrating out the φ-dependence in Eq. 4.26 is somewhat more involved, but the integrals involved are not difficult. By splitting the domain of the velocity integral at $u = 1$, inverting the order of integration to integrate first over φ and then comparing the resulting integrand in the two regions of the velocity integral, it can be shown that the result is

$$N_i = \int_0^1 \tilde{n}_i(\varphi, u_{\max}) d\varphi \tag{4.38}$$

$$= b \int_0^{u_{\max}} 2u \left(u - \Re\sqrt{u^2 - 1} \right) e^{-\beta_i (u-U)^2} du$$

$$= F_{U,\beta_i}(u_{\max}) \tag{4.39}$$

Now, the Sagdeev self-consistency condition is given by

$$0 = (N_i(\infty) + N_i(1)) \tag{4.40}$$

$$- (N_e(\infty) + N_e(1)).$$

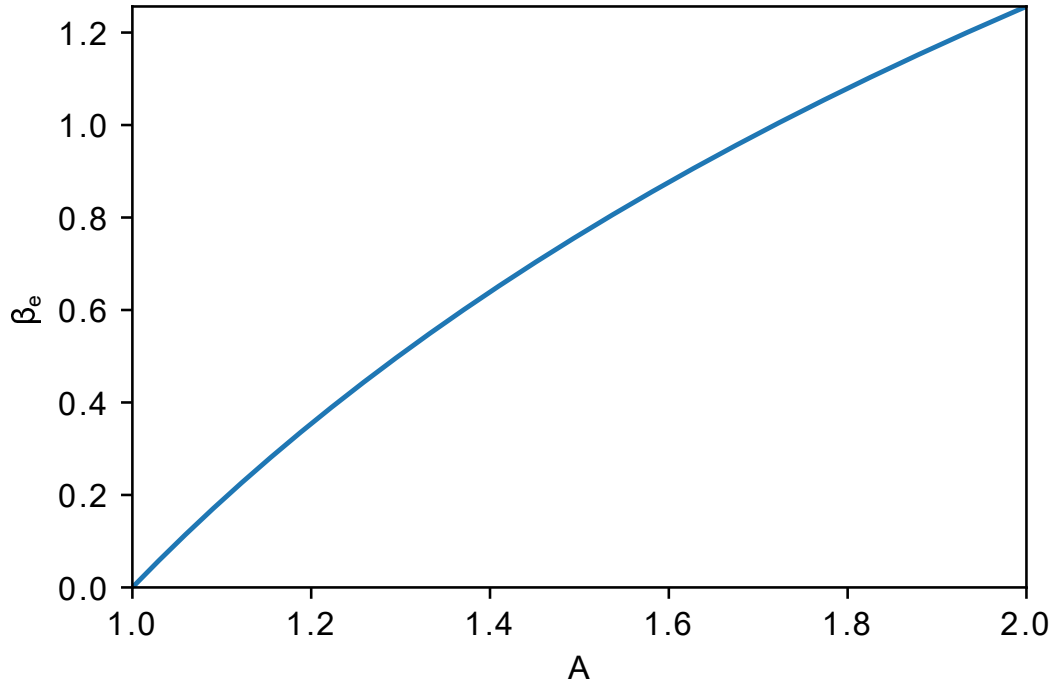


Figure 4.3: The relationship between β_e and A, the ratio of integrals defined in Eq. 4.41. It is clear that the constraints $\beta_e > 0$ and $A \leq 2$ correspond to $A \geq 1$ and $\beta_e \lessapprox 1.26$.

Rearranging to isolate β_e on one side of the equation,

$$\frac{e^{\beta_e} - 1}{\beta_e} = \frac{F_{U,\beta_i}(1) + F_{U,\beta_i}(\infty)}{G_{U,\beta_i}(1) + G_{U,\beta_i}(\infty)} = A(U, \beta_i) \tag{4.41}$$

This equation can be solved for β_e for given values of U and β_i to determine the electron temperature. Due to the simplified nature of the integrals involved, only $F_{U,\beta_i}(\infty)$ is not soluble analytically[9], and even that integral is sufficiently 'nice' to be evaluated numerically to relatively high accuracy[10].

Note that the condition $A > 1$ ensures that $\beta_e \geq 0$. The largest value A may be expected to take is roughly two—the maximum value of $2u(u - \Re\sqrt{u^2 - 1})$—corresponding to a maximum $\beta_e = 1.25643$. The function $\frac{e^{\beta_e}-1}{\beta_e}$, shown in Figure 4.3 has no inverse in terms of elementary functions, though a closed-form inverse exists and depends on the product logarithm special function. Given the small range

[9] $F_{U,\beta_i}(1)$ may be evaluated by noting that $\Re\sqrt{u^2 - 1} = 0$ where $0 \leq u \leq 1$.

[10] This integral (expressed as a function of U) takes the form of a convolution with a Gaussian kernel, a fact which could have potential for numerical evaluation, though this potential is not investigated here.

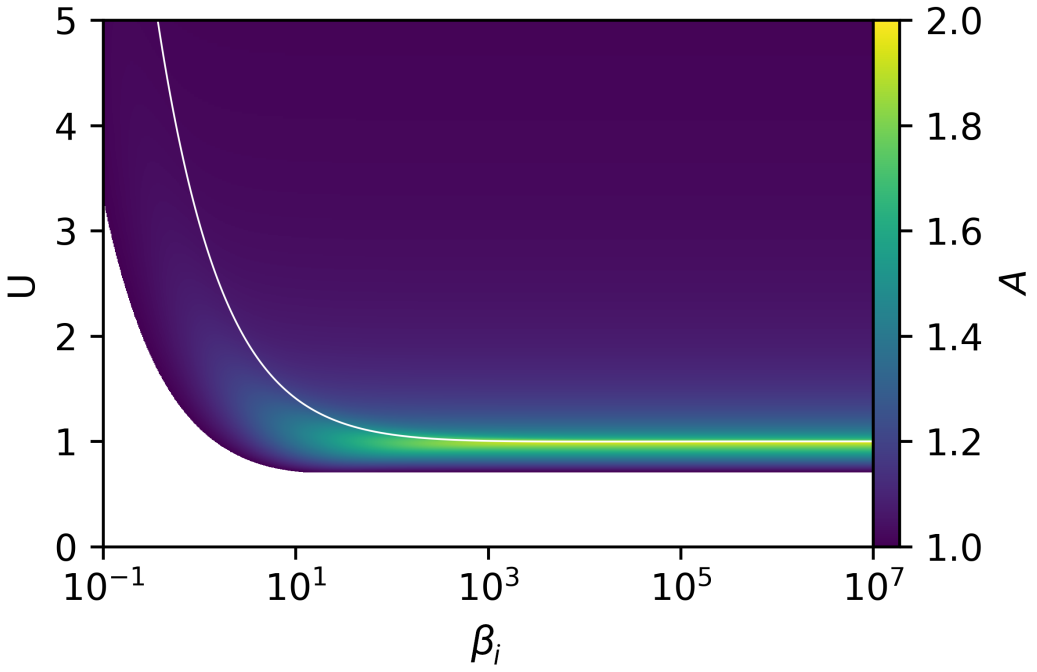


Figure 4.4: Values of the ratio of integrals A for a range of values of the independent variables U and β_i. Behaviour converges to a constant function $A(U)$ for large β_i in agreement with Eq. 4.44. The space is divided into two regions by an overlaid white line corresponding to the locus $\frac{d\Phi}{d\varphi}\Big|_1 = 0$. Above the line the condition of Eq. 4.32 is fulfilled, below it the gradient is negative and no shock forms. This condition restricts the valid phase-space to $U \geq 1$

of values taken by A and the smoothness of the function, calculating $\beta_e(A)$ by interpolation could also yield good accuracy.

The Cold-Ion Limit

I now examine the behaviour of the system for $\beta_i \gg 1$, corresponding to an ion temperature much smaller than the shock potential. For sufficiently large values of β_i, Equations 4.39 and 4.37 become

$$F(u_{\max}) = \begin{cases} 2U\left(U - \Re\sqrt{U^2-1}\right) & U < u_{\max} \\ 0 & U > u_{\max} \end{cases} \tag{4.42}$$

$$G(u_{\max}) = \begin{cases} 1 & U < u_{\max} \\ 0 & U > u_{\max} \end{cases} \tag{4.43}$$

This gives a value for A in closed form:

$$\frac{e^{\beta_e}-1}{\beta_e} = A = 2U\left(U - \Re\sqrt{U^2-1}\right) \tag{4.44}$$

For positive β_e, A must be larger than unity. This restricts the parameter space to $U^2 > \frac{1}{2}$. The minimum value of U^2 is further restricted by considering the condition of Eq. 4.32. Careful analysis shows that for $\beta_i \gg 1$, this condition reduces to $U \geq 1$. For finite β_i this condition has been numerically evaluated and the result is shown in Figure 4.4 as a line separating the region of space where this condition is fulfilled from that in which it is not. For all values of β_i, Eq. 4.32 is only fulfilled for $U \gtrapprox 1$. As $u = 1$ corresponds to the velocity required for an ion to pass the shock, this condition implies that the shock can never be sufficiently strong to reflect ions which propagate toward it in the plasma rest frame: these ions are always sufficiently energetic to pierce the shock front.

As seen in Figure 4.4, or by inspecting Eq. 4.44, the maximum value attained by A is two, at $U = 1$. For $A = 2$, $\beta_e \approx 1.25643$. From these constraints I derive the following constraints on the parameter space in which CSA is possible, in physical units

$$Ze\phi_{\max} < \frac{1}{2}m_i V^2 \tag{4.45}$$

$$e\phi_{\max} < 1.25643 k_B T_e \tag{4.46}$$

The maximum of $e\phi_{\max}/k_B T_e$ is attained at $Ze\phi_{\max} = \frac{1}{2}m_i V^2$.

Furthermore, I now introduce a dimensionless parameter which is popular for characterising the permitted phase-space of collisionless shock acceleration: the Mach number. In physical units this number is the ratio of the shock speed to the sound speed: $M^2 = m_i V^2/(Z k_B T_e)$. Within the normalisation in use here then, the Mach number is given by $M^2 = 2\beta_e U^2$. In the cold-ion limit under consideration, Eq. (4.44) is rearranged to yield an expression for U^2 in terms of

A, taken here as a function of the electron coldness β_e.

$$U^2 = \frac{1}{4}\frac{A^2}{A-1} \tag{4.47}$$

$$M^2 = 2\beta_e U^2 = \frac{1}{2}\frac{(\exp\beta_e - 1)^2}{\exp\beta_e - (1+\beta_e)} \tag{4.48}$$

When $U = 1$, both β_e and M attain their maximum values. Neither can be written in terms of elementary functions, but they can be expressed implicitly:

$$\beta_{\max} : e^{\beta_{\max}} - 1 = 2\beta_{\max} \qquad \beta_{\max} \approx 1.25643 \tag{4.49}$$

$$M_{\max} : e^{M_{\max}^2/2} - 1 = M_{\max}^2 \qquad M_{\max} \approx 1.5852 = \sqrt{2\beta_{\max}} \tag{4.50}$$

As U becomes large[11] on the other hand

$$A \to 1 + \frac{1}{4U^2} + \frac{1}{8U^4} \tag{4.51}$$

$$\beta_e \to 2(A-1) = \frac{1}{2U^2} + \frac{1}{4U^4} \tag{4.52}$$

$$M \to 1 + \frac{1}{2U^2} \tag{4.53}$$

β_e and M approach minimum values of zero and one respectively. Counter-intuitively, within the current normalisation, the highest Mach numbers correspond to the smallest shock speed U (relative to the shock potential). A different, perhaps less confusing, way of thinking of this is that for a constant physical shock speed, the Mach number increases with increasing shock potential.

The admissible parameter space of the system is shown in Figure 4.5. I display it in both electron temperature-normalised variables: $\sqrt{\beta_e/\beta_i}$ vs. M (because both of these variables are bounded) and for the purposes of comparison to previous work such as Savin [29] also in terms of the temperature ratio $T = \beta_i/\beta_e$ and the Mach number. By comparison to the work of Savin, it can be seen that the new method of calculation I have developed produces a more accurate treatment of the parameter space at high Mach number, where his model produced some odd artefacts. By inspecting these plots it is also clear that the bounds derived above for the parameters β_e and M in the cold-ion limit are also respected for finite ion temperature.

[11] in the sense that $(U-1) \gg 1/\sqrt{\beta_i}$ so that Eq. (4.44) may still be used

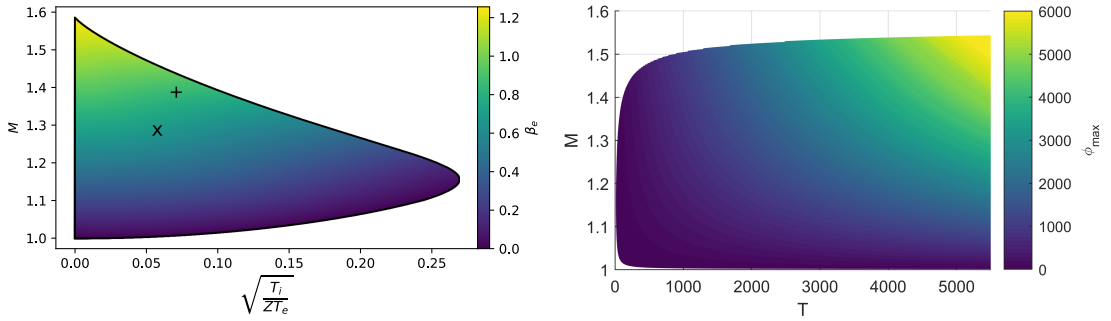


Figure 4.5: The valid parameter space of the collisionless shock system. Left: represented in electron temperature-normalised variables: The square root of the ion-electron pressure ratio $\sqrt{T_i/(ZT_e)}$ (where the root is taken solely to make the resulting plot easier to read), the Mach number $M = \sqrt{2\beta_e}U$ and β_e, which in this context serves as the electron temperature-normalised shock potential. Right: represented in the variables used in works such as Cairns et al. [116], which are normalised to the ion temperature. While the key parameters are bounded in relation to the electron temperature, they are not when normalised to the ion temperature. The simulation conditions presented in Fig. 4.1 are indicated on the left diagram by an x (two-slab simulation) and a + (laser-driven simulation).

Multiple Ion Species

The above theory assumes the existence of two species: electrons and ions. Many plasmas of practical interest are composed of more than a single ion species (for example, the carbon-hydrogen plasma produced in the experiment described at the end of this chapter consisted mostly of H^+ and C^{6+} ions, with some lower charge states), so it is desirable to know how to apply this theory to such plasmas.

First, note that the Sagdeev potential of a multispecies plasma—being linear in the charge density according to Eq. (4.31)—is simply the sum of the Sagdeev potentials associated with each ion species. As all electrons are interchangeable, the electron temperature of each species must be the same, and therefore the result Eq. (4.41) still applies. Now though, the numerator and denominator of the fraction are a sum over each ion species. Taking into account the scaling of the dimensionless parameters β_i and U with the charge state Z_i and mass number A_i of each ion i, and weighting by the charge density fraction of each ion, the modified form of Eq. (4.41) is:

$$\frac{e^{\beta_e}-1}{\beta_e} = \frac{\sum_i n_i Z_i \left(F_{U_i,\beta_i}(1) + F_{U_i,\beta_i}(\infty)\right)}{\sum_i n_i Z_i \left(G_{U_i,\beta_i}(1) + G_{U_i,\beta_i}(\infty)\right)} = A(U,\beta_i) \tag{4.54}$$

$$U_i^2 = \frac{A_i}{Z_i} U_H^2 \tag{4.55}$$

$$\beta_i = Z_i \beta_H \tag{4.56}$$

The normalised shock speed and ion coldness for each ion species are given in terms of the equivalent values, U_H and β_H respectively, normalised for a $Z = 1, A = 1$ ion species. This treatment assumes all ions have the same temperature, though it is trivial to modify it for differing temperature per ion species.

In my numerical experiments, the addition of a second ion species affects the precise limits of the parameter range for a collisionless shock to exist, but does not significantly affect the qualitative behaviour of the solutions. As an example, I show the parameter spaces for a pure-hydrogen plasma and for a plasma composed of 50% H^+ and 50% C^{6+} in Fig. 4.6. The parameters used here are the Mach number (for sound speed calculated using the average charge state and mass number across all ion species present) and the square root of the ratio β_e/β_i, which is equal to the ratio of ion pressure to electron pressure.

Ion Spectra

With the ability to find valid regions of $U - \beta_i - \beta_e$ phase-space, it is now desirable to find the energy spectrum of ions accelerated by a shock with these parameters. Dividing ions into three populations based on their initial velocity in the shock frame, it is possible to evaluate the velocities they are accelerated to, and therefore their energies in the lab frame. Ions with initial velocity $u < 0$, i.e. those travelling upstream, away from the shock, are not affected by the potential of the shock and so their velocity is unchanged. Ions with initial velocity $0 < u < 1$ reflect from the shock, and their final velocity is negated in the shock frame. Those with initial velocity $u > 1$ pass through the shock without being reflected – their velocity is reduced to $\sqrt{u^2 - 1}$ in the shock frame. Transforming these velocities into the lab frame, and negating the sign of velocities such that reflected ions have positive velocity:

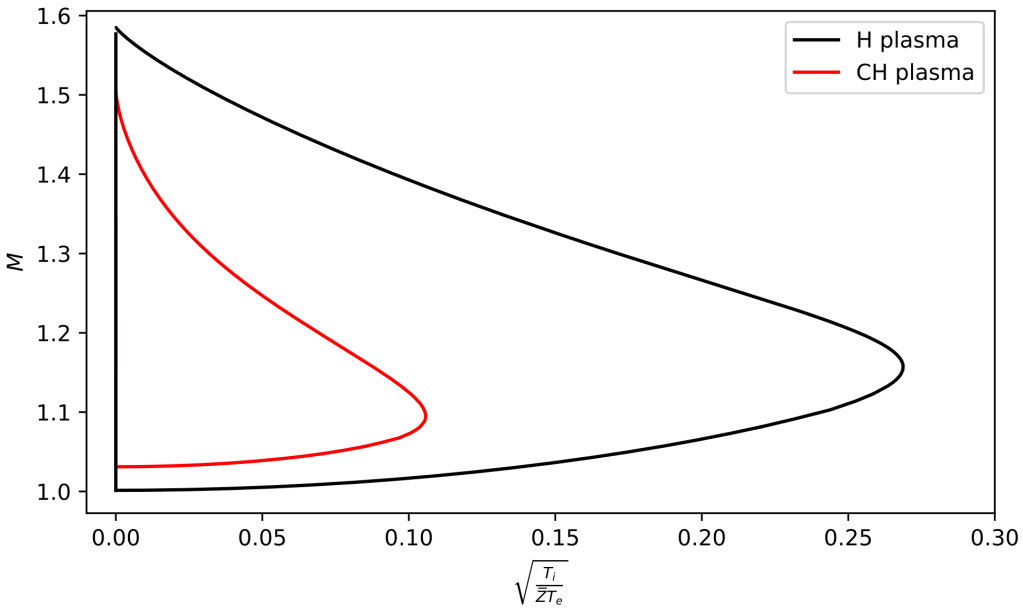


Figure 4.6: The phase space of the collisionless shock system for a hydrogen plasma (black) and for a fully-ionised 1:1 carbon-hydrogen plasma (red). The two-component plasma has a somewhat narrower range of parameters compatible with shock propagation, but the overall shape of its parameter space is comparable.

$$u_{\rm lab} = \begin{cases} u_{\rm E} - u_{\rm rf} & 0 > u_{\rm rf} + U \\ u_{\rm E} + 2U + u_{\rm rf} & 0 < u_{\rm rf} + U < 1 \\ u_{\rm E} + U - \sqrt{(u_{\rm rf} + U)^2 - 1} & 1 < u_{\rm rf} + U \end{cases} \tag{4.57}$$

for ions with initial velocity $u_{\rm rf}$ in the plasma rest frame, which is itself moving with an expansion velocity u_E relative to the lab frame. I note here that the potential downstream of the shock cannot be expected to stay around $\phi_{\rm max}$ infinitely far downstream. Some dissipative effect or another will serve to reduce it back to zero[12]. Therefore these initially 'fast' ions will at some point regain their original velocity $u_{\rm rf} > 1$ in the shock frame, so a more accurate set of velocities after the passage of the shock is

$$u_{\rm lab} = \begin{cases} u_{\rm E} - u_{\rm rf} & 0 > u_{\rm rf} + U \\ u_{\rm E} + 2U + u_{\rm rf} & 0 < u_{\rm rf} + U < 1 \\ u_{\rm E} - u_{\rm rf} & 1 < u_{\rm rf} + U \end{cases} \tag{4.58}$$

Transforming these into normalised energies, $E = u_{\rm lab}^2$:

[12]Otherwise a net electric field would form across the universe, which is not realistic.

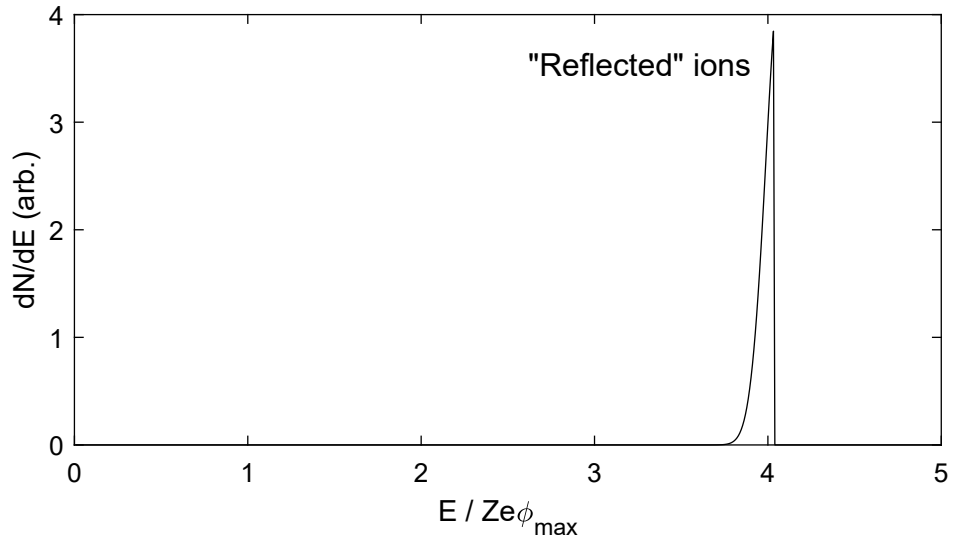


Figure 4.7: Ion spectrum predicted by the model of Section 4.2.2 for parameter values $U = 1.01,\ \beta_i = 1000$. Accelerated ions reflected from the shock form a narrow, sharp peak with energy approximately four times the shock potential. The energy axis is normalised according to the physical ion potential energy $Ze\phi_{\max} = \beta_i k_B T_i$.

$$f_{\rm E}(E) = \begin{cases} f_u(u_{\rm back}(E))\frac{d}{dE}u_{\rm back} & U_{\rm E}^2 < E \\ f_u(u_{\rm slow}(E))\frac{d}{dE}u_{\rm slow} & U_{\rm E}^2 < E < (1+U_{\rm E})^2 \\ f_u(u_{\rm fast}(E))\frac{d}{dE}u_{\rm fast} & 0 < E < (1-U_{\rm E})^2 \end{cases} \tag{4.59}$$

where the lab-frame shock speed $U_{\rm E} = U + u_{\rm E}$ has been introduced

Here values of $f_{\rm E}(E)$ in regions where two of the three conditions overlap are calculated by adding together values for the applicable regions, $f_u(u) \propto \exp(-\beta_i u^2)$ and the functions:

$$u_{\rm back}(E) = u_{\rm E} - \sqrt{E} \tag{4.60}$$

$$u_{\rm slow}(E) = \sqrt{E} - u_{\rm E} - 2U \tag{4.61}$$

$$u_{\rm fast}(E) = u_{\rm E} - \sqrt{E} \tag{4.62}$$

give the initial plasma-frame velocity in terms of the lab-frame energy. Figure 4.7 presents an example spectrum for parameters $U = 1.01$, $\beta_i = 1000$, and $u_{\rm E} = 0$.

4.3 Experiment: Collisionless Shock Acceleration in Foam Targets

I contributed to an experiment which was designed to produce collisionless shock acceleration in plasma formed from low-density plastic foam. It was hoped that this would be the first demonstration of collisionless shock acceleration using a micron-wavelength laser system, but this was pre-empted by the publication of results from the Titan laser system by Pak et al. [117] during experimental preparation. As will be discussed in the following, that experiment was able to produce more impressive ion spectra than the present one, and their more successful characterisation of the plasma conditions enabled more meaningful interpretation of their results than has been possible here. In the end, though the proton spectra produced in a few shots did exhibit a mono-energetic bump separated from a lower-energy population of protons with sheath-like energy distribution, we were unable to characterise the plasma well enough to claim this was due to collisionless shock acceleration. Indeed, simulation studies performed after-the-fact and described in §4.4 seem to indicate a different acceleration mechanism may have been responsible for these signals.

4.3.1 Experimental Conditions

As discussed in the previous sections—and demonstrated in the density profiles used in Pak et al. [117]—the density of plasma used for collisionless shock acceleration should be one to a few times critical density. For the wavelength of light used in this experiment, and assuming a CH material with charge-mass ratio $Z/A = 7/13$, critical density corresponds to a foam mass density of $3\,\mathrm{mg\,cm^{-3}}$. For comparison, achieving this density in a gas target would require tens of bars of pressure—prohibitively high—in pure hydrogen.

In this experiment, the required density of order $10\,\mathrm{mg\,cm^{-3}}$ was planned to be produced using a CH plastic foam target. We used cylindrical foam targets, 200 µm in diameter and 200 µm long, as shown in Fig. 4.8(a). This target size is significantly larger than the ideal thickness quoted by Fiuza et al. [113], $L_{\mathrm{target}} < \lambda\sqrt{m_i/m_e} \approx 40\,\mu\mathrm{m}$; this was imposed by limitations in the fabrication

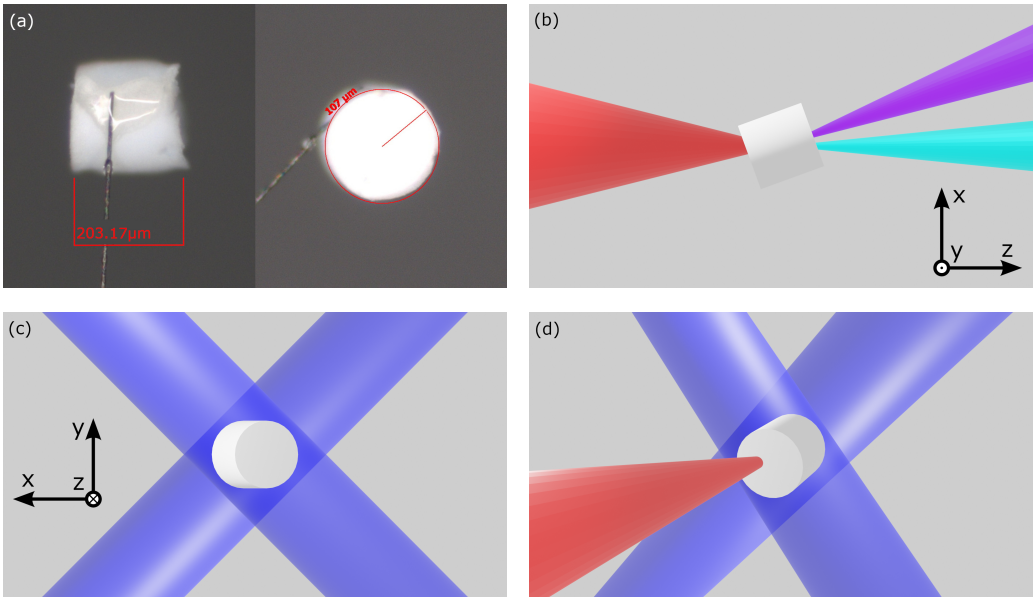


Figure 4.8: (a) Characterisation of the shape and size of a CH foam target. The foam is cylindrical with a length of 203 µm and a radius of 107 µm. (b) Top-down plan view sketch of the short pulse, target and ion beam geometry. The fundamental-frequency short-pulse (drawn in red) irradiates the target (white cylinder), causing emission of protons in the target-normal direction (purple, TNSA) and in the laser forward direction (cyan, CSA). (c) Front elevation sketch of the long-pulse heating geometry, from the point of view of the short pulse (not shown). The two long pulses (blue) irradiate the entirety of the foam target. (d) Perspective projection representation of the positioning of the three laser pulses relative to the target. All three beams are orthogonal, with the cylindrical target rotated 18° from the z-axis in the $x - z$ plane.

of free-standing foam structures, as was a higher than optimal target density of $50\,\mathrm{mg\,cm^{-3}}$ to $100\,\mathrm{mg\,cm^{-3}}$.

Foams are inherently highly inhomogeneous, as their low density is produced by regions of normal-densty material interspersed with voids. To produce homogeneous plasma, two of Vulcan's long pulse beams would be used to pre-form a plasma at the third-harmonic, 351 nm. The 1 ns heating beams carried 100 GW of power, focused to a spot of 300 µm diameter, corresponding to $1.4 \times 10^{14}\,\mathrm{W\,cm^{-2}}$. This was intended to heat the target to a temperature of approximately 1 keV, with the target expected to homogenise over the course of a few hundred picoseconds. The heating beams were delivered through the 45° windows of the target chamber, meaning the heating beams were orthogonal each other, as well as to the short pulse driving the shock generation. The heating geometry is rendered in Fig. 4.8. Due to

operational issues, the planned use of 3ω-converted long pulses was found not to be experimentally feasible within the time available. Second-harmonic heating was employed in its place, with associated higher target opacity to the heating beams.

After a delay of which varied between shots from 200 ps to 1.3 ns, the resulting supercritical plasma was struck by a fundamental frequency pulse of 2 picosecond duration, carrying 70 J of energy and focused to approximately $10^{19}\,\mathrm{W\,cm^{-2}}$, or $a_0 \approx 3$. This laser pulse was expected to "pile up" density at the critical density surface, sharpening that interface and producing the "two slab" density structure as in Fig. 4.1, which would then develop into the collisionless shock structure and accelerate ions in the laser-forward direction.

At the same time, the interaction of this laser pulse with the plasma needed to drive strong electron heating and achieve the large electron-ion temperature ratio necessary for shock formation according to the theory presented above. Using the relation derived in [113]:

$$\frac{k_B T_e}{m_e c^2} = \frac{\eta}{6}\frac{c t_{\mathrm{laser}}}{L_{\mathrm{target}}} a_0, \tag{4.63}$$

predicts, for $t_{\mathrm{laser}} = 2\,\mathrm{ps}$, $L_{\mathrm{target}} = 200\,\mu\mathrm{m}$ and $a_0 = 3$, that the electron temperature would reach hundreds of keV.

Ions travelling parallel to the drive beam axis were captured and recorded by a Thomson spectrometer, allowing protons and carbon ions to be distinguished. Ion spectra were also recorded in the target-normal direction (rotated 18 degrees from the drive beam axis), to determine whether there was a significant difference in the features of the spectrum due to the angular offset between the two axes. From prior work on TNSA [83] it was expected that some signal due to TNSA would be seen in the off-normal direction, albeit with fewer ions and a lower cut-off energy than in the target normal direction, but collisionless shock acceleration was only expected to produce a detectable ion signal parallel to the axis of the drive laser.

Ion spectra collected from the drive beam axis using a Thomson spectrometer on two different shots are shown in Fig. 4.9. Thermal (i.e. with $N \approx \exp(-E/kT)$ dependence) proton populations are clearly visible, with the spectra consistently

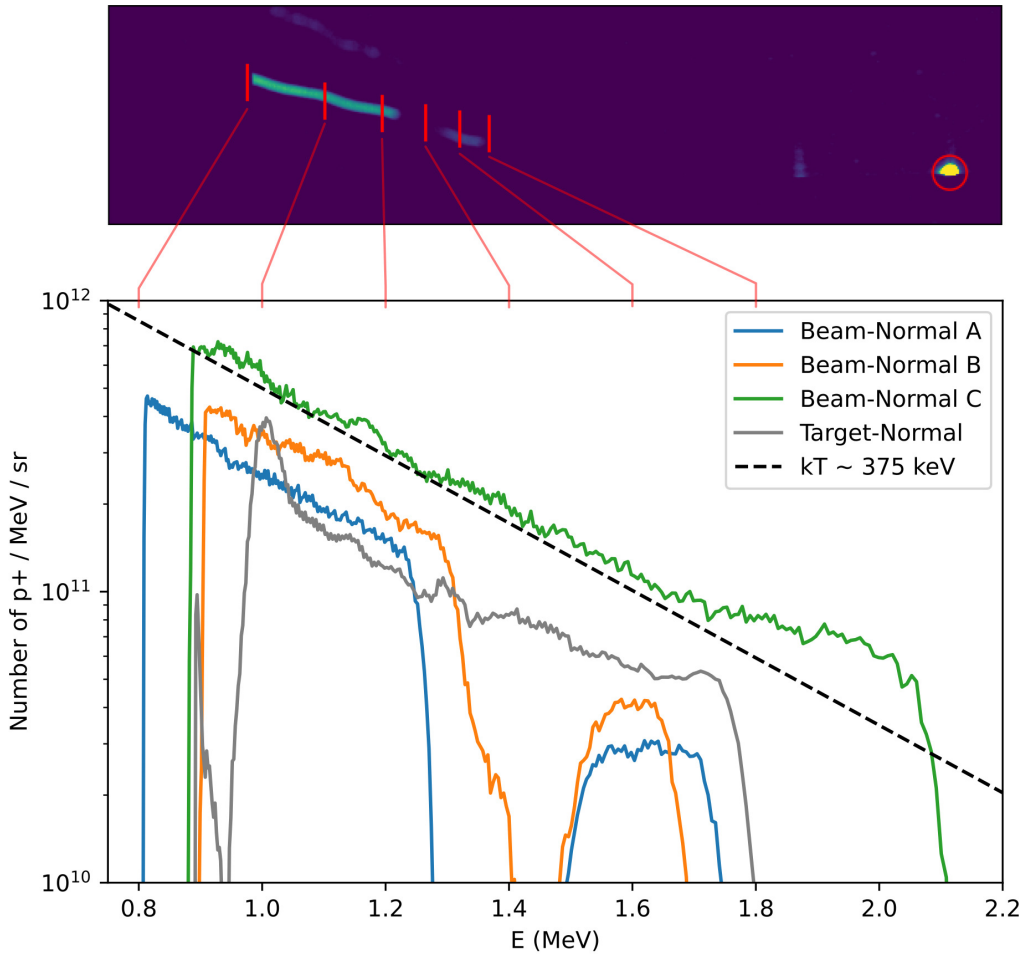


Figure 4.9: Top: The Thomson parabola spectrometer from which spectrum A below was extracted. The dispersion of the spectrometer is indicated by red lines at 0.2 MeV intervals from 0.8 MeV to 1.8 MeV, and the zero-deflection point (which receives neutral radiation such as x-rays and is the high-energy limit of the detector for charged particles) is marked with a red circle. Bottom: Energy spectra of ions, using the simple calibration of Martin et al. [121]. Beam-normal spectra A and B are from two shots which produced a proton peak with relatively narrow energy spread (1.5 MeV to 2.5 MeV) centred at approximately 1.6 MeV, beyond the cut-off of the broad, thermal spectrum component. Beam-normal spectrum C is from a shot on which this feature was not produced, and instead demonstrates a thermal spectrum with cut-off higher than the energy of the narrow peaks. Also shown is the target-normal spectrum collected on the same shot as spectrum A, once again demonstrating a thermal ion spectrum. For reference an exponential spectrum of temperature $k_B T = 375$ keV—similar to the thermal components of the spectra—is shown.

showing exponential decay temperatures around 375 keV. The energy dependence of these spectra is consistent with target-normal sheath acceleration. A narrower peak ranging from 1.5 MeV to 1.75 MeV is present in some shots along the drive beam axis, and on these shots the cut-off energy of the TNSA-like proton population is significantly lower than was found in shots without the narrower peak. No such narrow peak was ever observed in the target-normal direction, but TNSA-like spectra were detected with cut-off energy higher than the narrower peak, both in the target-normal direction and (on shots which did not exhibit the peak) along the drive beam axis.

Using formulas for TNSA cut-off energy scaling [94] and the acceleration time scaling proposed by Fuchs et al. [122]:

$$E_{\max} = 2k_B T_e \left(\ln\tau + \sqrt{1+\tau^2}\right)^2 \tag{4.64}$$

$$\tau = \omega_{pi} 1.3 t_{\text{laser}} / \sqrt{2\exp 1} \tag{4.65}$$

allows estimation of the electron temperature of the plasma (under the assumption that the TNSA-like proton spectra are well-described by a plasma expansion model):

$$k_B T_e = \frac{E_{\max}}{2\left(\ln\tau + \sqrt{1+\tau^2}\right)^2} \tag{4.66}$$

The predicted temperature varies from 100 keV to 160 keV for the cut-offs of the spectra shown in Fig. 4.9, using $\tau = 360$. This is much lower than the intended several MeV, and looking at the relation used by Fiuza et al. [113] (Eq. (4.63)). An x-ray crystal spectrometer was fielded to measure electron temperature, but unfortunately it did not produce usable data.

4.4 Simulation Post-Mortem

Simulations were carried out to confirm the initial conditions of the plasma prior to shock formation. First, the radiation-hydrodynamic code FLASH was used to model the long-pulse interaction in three dimensions: two beams, each carrying 130 J of energy at the second harmonic, irradiated a cylindrical target of diameter 200 µm, length 200 µm and density 100 mg cm^{-3} from orthogonal directions. The

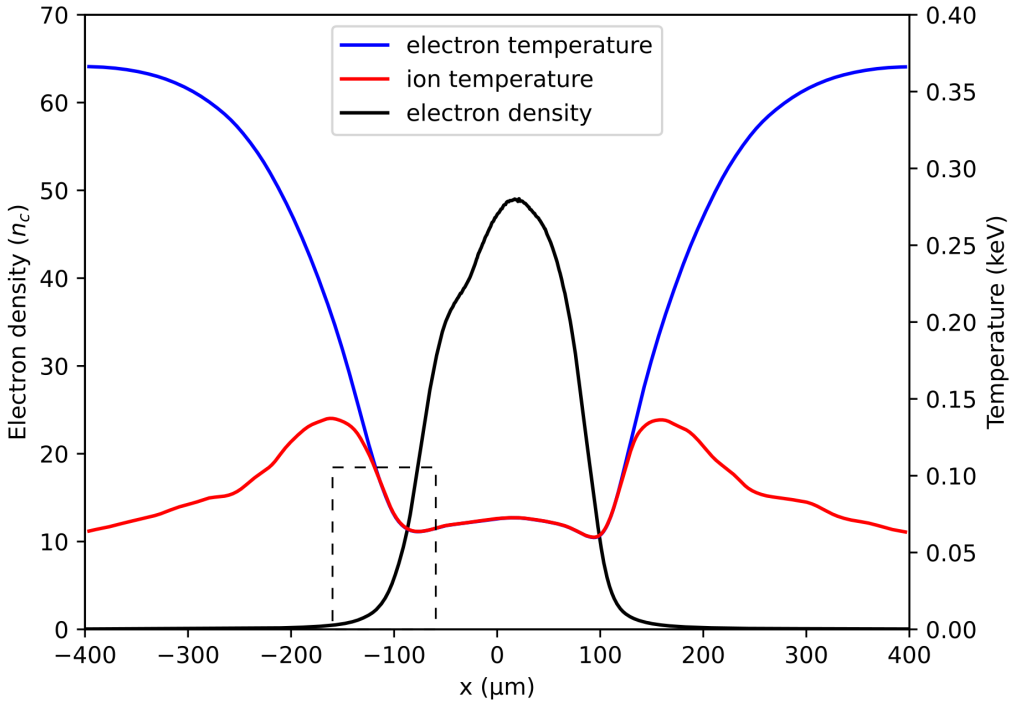


Figure 4.10: Initial plasma conditions, extracted along the laser axis at $t = 1200\,\mathrm{ps}$ from FLASH simulation. Contrary to the intended design of the experiment, the core of the target has a temperature below 100 eV and peak electron density of fifty times critical. Dashed box marks the region of interest of Fig. 4.11

.

target was rotated 18° from the x-axis. The target and beam geometry was set up as shown in Fig. 4.8. 1200 ps after the onset of heating, a snapshot of the target's ion and electron densities and temperatures, and the average charge state $\bar{Z}$ of the target, was extracted along the direction of short pulse propagation. A plot of these is shown in Fig. 4.10, showing that both electron and ion temperatures were much lower than the 1 keV that the experiment was designed to produce, and that the electron density of the target is around fifty times critical, far in excess of the few times critical desired to produce the shock structure. This is a consequence of the significantly larger target size and higher foam density used in the experiment, compared to the ideal parameters derived during the design of the experiment. The electron density of the target is in fact somewhat higher than would be expected from the target's ambient density, due to a small amount

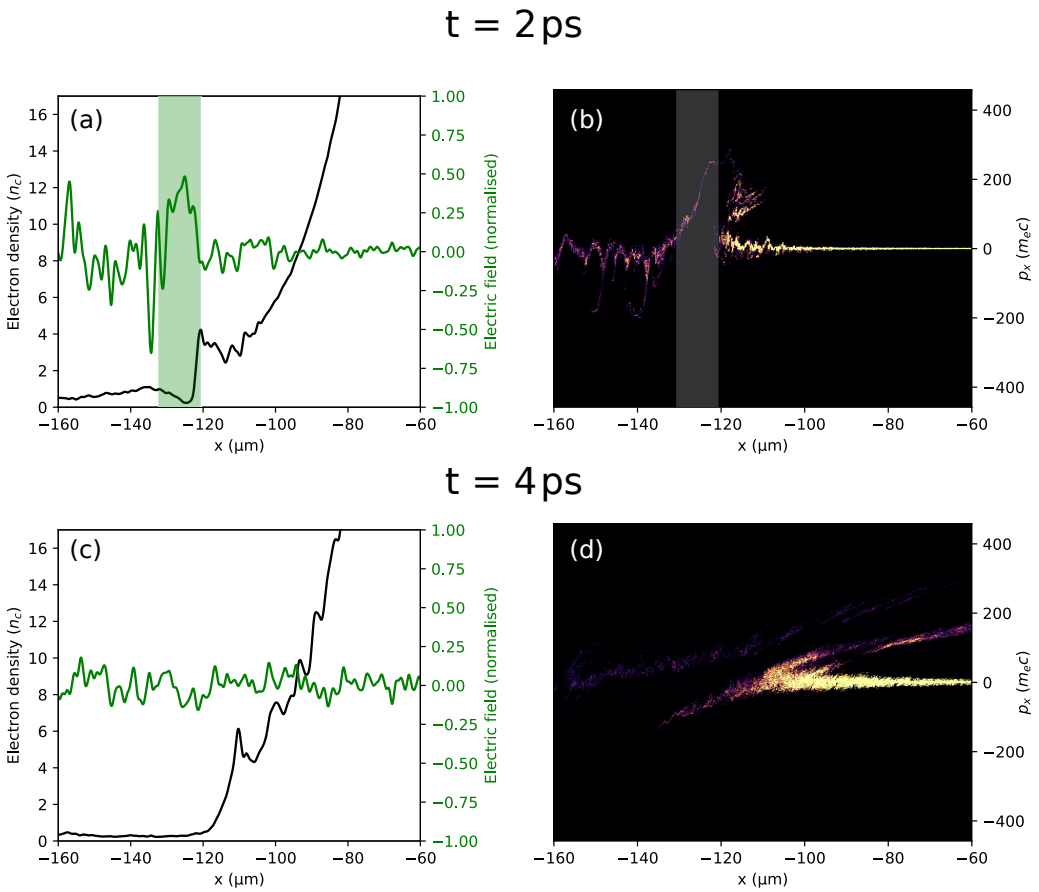


Figure 4.11: (a,c) Electron density and electric field profiles and (b,d) proton phase space at the point of laser reflection, 2 ps (a,b) and 4 ps (c,d) after simulation start. At the earlier timestep laser reflection is still underway, and the later timestep is after the end of laser reflection. The phase-space in (b) shows a similar structure to that seen in Fig. 4.1 at the shock front. However, the acceleration of ions is only maintained while the laser is incident on the reflection front, and the associated electric field highlighted in (a) has not persisted to the later timestep in (c). The CSA-like phase-space structure also collapses once laser irradiation ceases, as seen in (d). Note that unlike in Fig. 4.1, protons are principally accelerated from *behind* the density step, rather than by reflection from a moving electric field strucutre.

of target compression by the heating beams.

A Smilei particle-in-cell simulation was used to model the Vulcan short pulse (2 ps duration, 1 μm wavelength, $a_0 = 3$) propagating into CH plasma with these parameters. The initial state of the simulation is shown in Fig. 4.10(a). Within the relatively cold interior of the target, the majority ion charge states present were H^+ and C^{3+}, but the carbon was quickly further ionised to C^{4+} during the short pulse interaction.

The state of the plasma during and shortly after laser reflection from the critical density surface is shown in Fig. 4.11. Due to the very high electron density of the target (and the non-negligible scale length at the target front), the laser pulse was reflected from the density gradient some distance before the density plateau, in a region with an existing density gradient. Though some ions are accelerated from the region *behind* the critical density front, this process produces a proton phase-space that resembles that shown in [123] more than it does the 'forked' phase-space structure typical of self-sustaining shock propagation, which predominantly accelerates ions originating in the region upstream (ahead) of the shock. The accelerating electric field and phase-space structure are not maintained after shock reflection ceases (Fig. 4.11(c,d)), providing further evidence that in this case no self-sustaining shock was launched, and that this simulation shows proton acceleration by a laser-driven process rather than the desired mechanism of an electrostatic structure supported by the plasma itself. My proposed explanation of this difference is that the steep density gradient already present at the site of laser reflection suppressed the formation of a high-density plasma 'slab' as in Fig. 4.1, and that the best way to improve on this experiment would therefore be to use thinner targets at lower initial densities, so as to move the point of laser reflection closer to the density plateau and allow formation of the density structure necessary to drive a self-sustaining collisionless shock.

Finally, I observe that most of the protons accelerated in the simulation have forward momenta of $100 - 125 m_e c$. This translates to ion energies of 1.4 MeV to 2.2 MeV, somewhat higher than the energies of the peaks seen in experiment. It is possible that collisional stopping in the dense target material reduced these energies somewhat, though simple calculation based on the Bethe formula [124] suggests less than 100 keV would be lost to collisional stopping while moving through the target.

4.5 Summary

In this chapter, I have introduced the reader to the topic of laser-driven ion acceleration. I provided a brief overview of target-normal sheath acceleration and its

shortcomings, then introduced two alternative acceleration mechanisms that promise to produce ion beams with narrower energy spread than is possible using TNSA.

The first of these is radiation pressure acceleration, in which the laser pulse directly drives a shock through supercritical plasma. Ions reflected from the electric field of this shock are expected to be highly monochromatic. However, target-normal sheath acceleration is expected to dominate over radiation pressure acceleration unless the incident laser irradiation is extremely intense, limiting the utility of this acceleration mechanism.

Collisionless shock acceleration, on the other hand, involves a shock which continues to after the end of laser irradiation, supported by the plasma itself. The laser pulse produces conditions in the plasma which subsequently lead to shock formation. Because the short, intense pulse involved in collisionless shock acceleration propagates into pre-formed plasma, target-normal sheath acceleration can be suppressed more effectively than is possible when driving radiation pressure acceleration. The actual mechanism of ion acceleration in both of these methods is the same: reflection from the electric field of an electrostatic shock; it is the method of shock formation which differs.

The theory of collisionless shocks has seen significant development in recent years in the specific context of ion acceleration. First, 'numerical trial and error' by Cairns et al. [116] suggested a maximum Mach number at which collisionless shocks of this type can exist, and found this critical Mach number to be around 1.6. Later, Savin [29] presented significant advances in the understanding of the region of parameter space in which collisionless shock acceleration is viable, using an iterative search process to determine the ratio of shock potential to ion thermal energy, given a certain shock speed and electron-ion temperature ratio. This shed light on the range of applicable parameters, but the iterative search failed to converge for Mach numbers very close to critical and the resulting parameter space plots showed significant artefacts in the limit of high Mach number.

In this chapter I present a development to this theory: by expressing quantities normalised to the shock potential rather than the ion temperature, the integrals

involved become easier to solve, and by rearranging the problem, the electron temperature can be determined directly, as a function of the shock speed and ion temperature (both normalised to the shock potential). This eliminates the need for a slow iterative search, so is faster to evaluate, and also more accurate, stable, and reliable than the method used by Savin.

I then presented the results of an experiment conducted at Target Area West of the Vulcan laser facility, aiming to accelerate proton beams via collisionless shocks. Practical difficulties prevented this experiment from attaining the conditions predicted for Vulcan by Savin [29], and although some proton spectra resembling low-energy collisionless shock acceleration were seen on a few shots, simulations conducted after-the-fact suggest that these protons may in fact have been accelerated by a different process. Based on the results of these simulations I propose that the higher density and larger size of the foam targets relative to theoretical optimum likely prevented formation of the density structures necessary to launch a collisionless shock beyond the laser reflection front. Based on this I suggest that pushing to thinner, lower-density foam targets may improve the prospects of observing true collisionless shock acceleration in future experiments.

5
Proton Radiography

Contents

5.1 How It Works 92
5.1.1 Proton Sources 93
5.1.2 Detection 94
5.1.3 Formation of a Radiograph 95
5.1.4 The Confounding Effect of Electric Fields 100
5.2 The Inverse Problem 101
5.3 A New Solver for `PROBLEM` 103
5.4 Experiment: Whole-Beam Self-Focusing at Orion . . . 105
5.4.1 Long-Pulse Interaction 108
5.4.2 Short-Pulse Interaction 112
5.5 Summary 117

The proton radiography diagnostic uses deflection of protons in a beam to probe electromagnetic fields present in a plasma [125, 126]. It has been used to image magnetic fields in laboratory analogues of astrophysical collisionless shocks [127–129], the Weibel instability in interpenetrating plasma flows [130], stochastic magnetic fields amplified by the turbulent dynamo mechanism [131], and fields involved in laser-driven magnetic reconnection [132, 133]. In the higher-density regime, proton radiography has been used to probe plasma fields in imploding inertial fusion capsules in both direct-drive [27, 134, 135] and indirect-drive laser-hohlraum [28, 136, 137] configurations, and in studies of laser channelling physics relevant to

fast ignition [12]. Proton radiography has also been used to validate magnetic field measurements made using Faraday rotation at OMEGA [138]. It has also been used to probe electric fields, for example those associated with collisionless electrostatic shocks formed at the edge of an expanding plasma bubble [139–141]. In plasmas where both electric and magnetic fields are significant, special care is needed when interpreting radiographs as the effects of both can overlap; this is discussed below in §5.1.4.

Kugland et al. [142] presented a detailed theoretical treatment of proton radiography. Following this, algorithms for quantitative recovery of path-integrated magnetic field values were proposed by both Kasim et al. [21] and Bott et al. [143]. Later, Kasim et al. [144] addressed uncertainties in the initial source profile used as input to these algorithms. Chen et al. [145] investigated machine learning approaches to recover parametrised field structures from a proton radiograph, and proposed that imaging from several lines of sight could improve the fidelity of such three-dimensional reconstructions. While most proton radiography experiments to date use only one line of sight, Cecchetti et al. [146] and Li et al. [147] have used unusual experimental geometry to probe similar interactions both side-on and face-on in a single shot, and more recently Tubman et al. [133] used two proton probes separated by 45° to assist in separating electric and magnetic field contributions to proton deflection. The idea of multi-view proton radiography will be explored in detail in Chapter 6.

5.1 How It Works

Proton radiography measures electric and/or magnetic fields present in plasma indirectly, via the deflections imparted by those fields on a beam of protons as they pass through the plasma. Implementations of this diagnostic differ in several ways, such as the process used to produce the proton beam, the plasma regime being probed, and the detector used to record the beam after it has propagated through the plasma.

5.1.1 Proton Sources

The ion acceleration mechanism used to produce the proton probe beam varies, mostly based on the facility at which the experiment is being carried out. Facilities with multi-joule short-pulse capabilities—such as Vulcan's TW-scale Target Area West—use target-normal sheath acceleration as described in §4.1 to produce a proton beam from a thin metal foil, for example in Refs. [125–127, 132]. TNSA-generated proton beams have several properties which make them good candidates for proton radiography. Foremost is their good degree of laminarity - though actually emitted from a source area hundreds of microns in size, the beams behave as though produced by a much smaller virtual source [97]. This is essential for imaging performance, as any finite effective source size induces a corresponding blur in the produced image. The thermal energy spectrum of TNSA sources is also highly significant to their use in proton radiography. Taking into account the differing travel time of protons from their source to the region of plasma being probed, the broadband spectrum produces an element of time resolution. When used alongside a detector which can distinguish protons of different incident energy, this time resolution can allow experimenters to observe the temporal evolution of the system. Stacks of radiochromic film (see §5.1.2) are an appropriate energy-dependent detector for this kind of study, though their ability to resolve protons by energy is imperfect as it is a result of the energy-dependence of stopping power within the layers of the stack. Though protons lose much of their energy in the layer in which they stop, they also deliver non-negligible dose to the layers they pass through beforehand.

At ICF-scale facilities such as the National Ignition Facility, OMEGA and Laser Mégajoule—which are characterised by the provision of many kilojoule-scale long pulse beamlines—it is more common to use an exploding pusher-based platform for production of a proton beam [148]. A spherical capsule filled with deuterium and helium-3 is compressed by several of the available long pulse beams. Two

fusion reactions take place in the compressed fuel:

$$\mathrm{D} + {}^{3}\mathrm{He} \longrightarrow {}^{4}\mathrm{He} + \mathrm{p} + 18.35\,\mathrm{MeV} \tag{5.1}$$

$$\mathrm{D} + \mathrm{D} \longrightarrow \mathrm{T} + \mathrm{p} + 4.03\,\mathrm{MeV} \tag{5.2}$$

The protons produced by these reactions carry kinetic energies of 14.63 MeV and 3.02 MeV respectively. Note that this is not strictly speaking an ion *acceleration* mechanism: the protons are born from the thermonuclear reactions with the energy already present. Nor is the population of protons produced really a *beam*, so much as an isotropic expanding spherical shell covering a full 4π in solid angle. Notwithstanding these technicalities, these sources have been used to great effect in many of the aforementioned experiments at the NIF and OMEGA, for example those of Refs. [129–131, 147, 148].

5.1.2 Detection

Once the protons have passed through the plasma and propagated for some additional distance, they are detected using one of several charged particle detectors.

CR-39 is a polymer material that is damaged by the energy deposition associated with stopping charged particles. This damage is made visible through an etching process that produces 'pits' in the sheet of CR-39, which can then be counted individually if the particle flux was sufficiently low. As the depth of a proton's Bragg peak into a material increases with increasing incident energy, particles of differing energies can be distinguished in this process.

Radiochromic film (RCF) is a single-use, self-developing detection medium. RCF is commercially available, with one major supplier being GAFChromic™. The active detection medium of RCF is a suspension of monomers which are induced to polymerise when they absorb energy from ionising radiation, in this case energetic protons. The polymer produced has a different visible light absorption spectrum to the monomer, causing a colour change of the film related to the local energy deposition of the protons. The active medium is typically sandwiched between thin plastic substrate layers, which provide structural stability and protect the active

layer from physical damage; in some high-sensitivity films designed for low-dose applications, the front surface of the active layer is exposed to avoid radiation being stopped before reaching it.

RCF is often used to construct 'stacks', to maximise the information available for analysis when used in combination with a broadband proton source (i.e. TNSA). This entails the use of less sensitive radiochromic film at the front of the stack, where the relatively plentiful population of low-energy protons will be stopped, followed by films with higher sensitivity at the rear of the stack to detect the smaller population of higher-energy protons. Filters composed of thin metal foils are sometimes employed at various points in the stack, between layers of RCF, in order to slow protons and ensure that protons with high initial energies stop within the stack. Because these RCF stacks detect protons in an energy-dependent way, they are often used to obtain a time-resolved picture of an interaction, exploiting the variation in travel time from source to plasma for protons of different incident energy. An example RCF stack is described in Fig. 5.10.

5.1.3 Formation of a Radiograph

The first step towards finding an algorithm for deducing the plasma conditions which produced a given proton radiograph is to express mathematically the process by which such an image is formed. The formulation I present here is similar to that of Kugland et al. [142] and Bott et al. [143], although I express the geometric parameters of the imaging set-up differently, in order to separate the effect of image magnification from that of the plasma explicitly.

I define a coordinate system such that the centre of the plasma under observation is located at the origin, and the probe direction is the positive z-direction. The source-plasma distance is denoted z_s and the plasma-image plane distance is z_i; i.e. the proton beam is assumed to be generated by a point source at $[0, 0, -z_s]$ and the detector is the plane defined by $z = z_i$. This geometry is shown in schematic form in Fig. 5.1

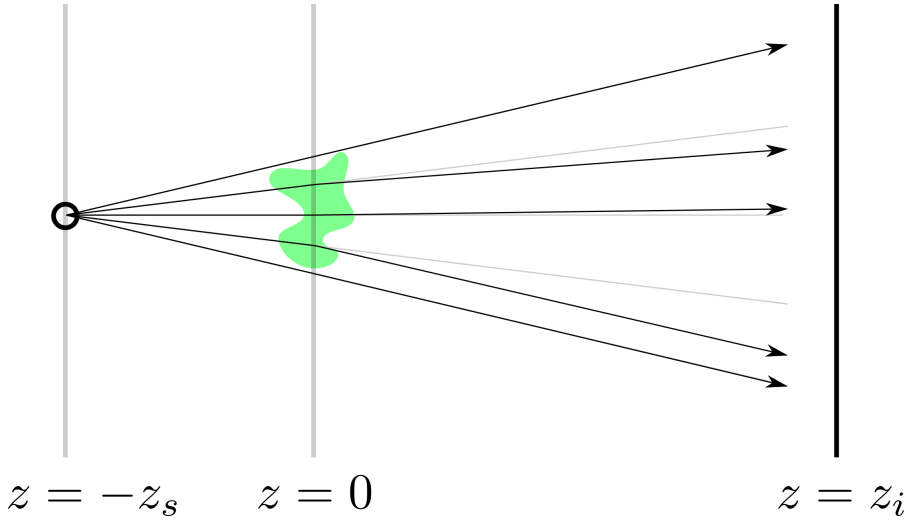


Figure 5.1: A schematic of proton radiography showing the distances involved. The point source is located at $z = -z_s$, the detector plane at $z = z_i$ and the plasma centre at the origin.

Proton Trajectories

In the zero-deflection limit, the protons travel ballistically and a proton travelling with momentum $[p_x, p_y, p_z]$ reaches the 'target plane' ($z = 0$) at position $[z_s p_x/p_z, z_s p_y/p_z, 0]$ and the image plane at position $[(z_s + z_i)p_x/p_z, (z_s + z_i)p_y/p_z, z_i]$. Using the subscript $\perp$ to denote the x and y components of a vector, this ballistic transport can be expressed concisely:

$$\boldsymbol{x}_\perp(z_i) = (1 + z_i/z_s)\ \boldsymbol{x}_\perp(0) = M\boldsymbol{x}_\perp(0) \tag{5.3}$$

$$\frac{\boldsymbol{p}_{\perp 0}}{p_{z0}} = \frac{\boldsymbol{x}_\perp(z_i)}{z_s + z_i} = \frac{\boldsymbol{x}_\perp(0)}{z_s}. \tag{5.4}$$

This equation both defines the geometric magnification $M = 1 + z_i/z_s$ of the system and motivates it as an important parameter.

Now consider the effect of changing a proton's (transverse) momentum by an amount $\Delta\boldsymbol{p}_\perp$ as it crosses the target plane. Due to the very small size of most plasmas compared to the distances involved in the diagnostics, this can be justified as a good approximation. A proton whose momentum is altered as it crosses the

$z = 0$ plane will cross the image plane at the transverse position:

$$\boldsymbol{x}_\perp(z_i) = z_s \frac{\boldsymbol{p}_{\perp 0}}{p_{z0}} + z_i \frac{\boldsymbol{p}_{\perp 0} + \Delta\boldsymbol{p}_\perp}{p_z} \tag{5.5}$$

$$\approx M\boldsymbol{x}_\perp(0) + z_i \frac{\Delta\boldsymbol{p}_\perp}{p_z} \tag{5.6}$$

$$= M\left[\boldsymbol{x}_\perp(0) + f_g \frac{\Delta\boldsymbol{p}_\perp}{p_z}\right], \tag{5.7}$$

under the approximation that $p_z \approx p_{z0}$.

This final expression introduces the 'geometric focal length', $f_g = z_i/M$, so named for its resemblance to a focal length when defined in terms of z_s and z_i:

$$\frac{1}{f_g} = \frac{1}{z_s} + \frac{1}{z_i} \tag{5.8}$$

In other words, a lens of focal length f_g would produce an image of the point source at the image plane if placed at the origin. This quantity parametrises the distance a deflected proton moves *relative to the magnified image*; two experiments with the same initial proton energy and the same plasma parameters will produce similar image features, even if the imaging geometry is changed significantly, as long as these changes maintain a constant geometric focal length and change only the magnification of the system.

I have so far neglected to explicitly specify the process by which the momentum deflection $\Delta\boldsymbol{p}_\perp$ is imparted upon the protons. In short, the effect of the Lorentz force within the plasma is to deflect the protons slightly from ballistic trajectories. This deflection is found by integrating the Lorentz force over the time that a proton spends inside the plasma:

$$\boldsymbol{\Delta p}(x, y) = \int_0^\infty \boldsymbol{F}_L dt \tag{5.9}$$

$$= q\int_{-z_0}^{z_0} (\boldsymbol{E}(x', y', z') + \boldsymbol{v}(z') \times \boldsymbol{B}(x', y', z')) \frac{dz'}{v_z(z')}, \tag{5.10}$$

where x' and y' are evaluated along the proton's trajectory inside the plasma, as functions of z'.

I now make some assumptions to simplify this expression. First, under the thin-plasma approximation the newly-introduced plasma thickness z_0 must be sufficiently

small that the (x, y) positions of protons within the plasma do not change by a significant factor while the protons are inside the plasma, i.e. $\boldsymbol{x}_\perp(\pm z_0) \approx \boldsymbol{x}_\perp(0)$. Under ballistic transport, this reduces to:

$$(z_s \pm z_0) \approx z_s, \tag{5.11}$$

and therefore $z_0 \ll z_s$.

It is also desirable for the initial proton trajectories to be paraxial, i.e. $p_x(0) \ll p_z(0)$. By considering the relationship between position and momentum at $z = 0$, it is easily seen that this condition implies that the transverse size r_0 of the plasma being observed must satisfy $r_0 \ll z_s$.

I also make assumptions that the electromagnetic fields producing the deflections are sufficiently weak. In particular, momentum deflections should be significantly smaller than the initial forward momentum of the protons:

$$\int_{-z_0}^{z_0} \boldsymbol{F}_L \frac{dz'}{v_z} \ll p_{z0} \tag{5.12}$$

$$2z_0 \langle \boldsymbol{F}_L \rangle \ll p_{z0} v_{z0} \tag{5.13}$$

$$2z_0 q \langle \boldsymbol{B} \rangle \ll p_{z0} \tag{5.14}$$

$$2z_0 q \langle \boldsymbol{E} \rangle \ll p_{z0} v_{z0} \approx 2\mathcal{E}_0, \tag{5.15}$$

where I have introduced angle-bracket notation to represent the mean value of a field across the plasma, and $\mathcal{E}_0$ is the initial kinetic energy of protons.

Radiograph Formation

Now considering a population of protons, it is clear that the total number of protons in any portion of the beam must be conserved, which implies that the areal density of a portion of the beam must scale inversely to its area as the beam propagates. Expressed mathematically, the density-area element product of a small region of the beam must be preserved:

$$\rho_0(\boldsymbol{x}_{0,\perp}) dA_0 = \rho_i(\boldsymbol{x}_{i,\perp}) dA_i \tag{5.16}$$

where ρ_0 and ρ_i are the plasma-plane and image-plane proton areal density respectively, and the point $\boldsymbol{x}_{0,\perp}$ has been transported to $\boldsymbol{x}_{i,\perp}$ according to the mapping Eq. (5.7). Converting this into a Jacobian equation,

$$\left|\frac{\partial(x_i, y_i)}{\partial(x_0, y_0)}\right| \rho_i(x_i, y_i) = \rho_0(x_0, y_0). \tag{5.17}$$

Here the transverse coordinate components $\boldsymbol{x}_\perp(0) = [x_0, y_0]$ and $\boldsymbol{x}_\perp(z_i) = [x_i, y_i]$ are written out explicitly. The meaning of this equation is that the local beam fluence scales inversely with the Jacobian determinant of the mapping from plasma-plane to image-plane coordinates. As both the Jacobian and the image-plane coordinates themselves are a function of the plasma-plane coordinates, this equation is a highly nonlinear PDE, which is known as the Monge-Ampère equation. Solvers for this equation often express the relation between the initial (plasma-plane) and final (image-plane) coordinates in terms of a 'deflection potential', $\Phi(x_0, y_0)$, such that $\boldsymbol{x}_\perp(z_i) = M(\boldsymbol{x}_\perp(0) + \boldsymbol{\nabla}\Phi)$ [21, 143, 149]. Thus, the image plane fluence mapping, Eq. (5.17), becomes:

$$\rho_i'(\boldsymbol{x}_\perp(0) + \boldsymbol{\nabla}\Phi) = \left|\frac{\partial x_i}{\partial x_0}\frac{\partial y_i}{\partial y_0} - \frac{\partial x_i}{\partial y_0}\frac{\partial y_i}{\partial x_0}\right|^{-1} \rho_0(\boldsymbol{x}_\perp(0)) \tag{5.18}$$

$$= \left|\left(1 + \frac{\partial^2\Phi}{\partial {x_0}^2}\right)\left(1 + \frac{\partial^2\Phi}{\partial {y_0}^2}\right) - \left(\frac{\partial^2\Phi}{\partial x_0\,\partial y_0}\right)^2\right|^{-1} \rho_0(\boldsymbol{x}_\perp(0)) \tag{5.19}$$

$$= \left|1 + \nabla^2\Phi + \frac{\partial^2\Phi}{\partial {x_0}^2}\frac{\partial^2\Phi}{\partial {y_0}^2} - \left(\frac{\partial^2\Phi}{\partial x_0\,\partial y_0}\right)^2\right|^{-1} \rho_0(\boldsymbol{x}_\perp(0)), \tag{5.20}$$

where in the interest of brevity I have defined the magnification-compensated image-plane fluence as $\rho_i'(\boldsymbol{x}) = M^2\rho_i(M\boldsymbol{x})$.

It is of theoretical interest that the expression within the absolute value operator on the right-hand side of Eq. (5.20) can be expressed more concisely in terms of the two eigenvalues $\lambda^\pm$ of the Hessian matrix of Φ:

$$\rho_i'(\boldsymbol{x}_\perp(0) + \boldsymbol{\nabla}\Phi) = \left|(1 + \lambda^+)(1 + \lambda^-)\right|^{-1} \rho_0(\boldsymbol{x}_\perp(0)). \tag{5.21}$$

While this expression does not appear to be of much practical utility for solving the Monge-Ampère equation, it will come in handy when classifying regimes in the next section.

5.1.4 The Confounding Effect of Electric Fields

In the above section, I derived the deflections of protons traversing a plasma as a sum of contributions due to electric fields and magnetic fields. In most applications it is desirable to measure electric and magnetic fields separately, rather than some combination of both. Reconstructions of proton radiography data, as discussed in the next section, recover the deflections of protons rather than the electric or the magnetic field, *per se*. Some model or technique must be used to then recover the desired electric and/or magnetic field data from the deflections.

In the simplest case, some systems may be assumed to have only electric or only magnetic deflections, rather than both. For instance, some fast laser-driven processes can generate magnetic field structures that persist for relatively long MHD timescales, whereas electric fields will decay much more quickly as the plasma restores quasineutrality [12, 150]. In these cases the proton deflections are dominated by the contribution of magnetic fields, as long as the probing is carried out at sufficiently late times.

Alternatively, some experiments are constructed such that the Lorentz force component due to one of the fields will be mostly parallel to the proton beam, and therefore not affect the observed deflections significantly. For example, Cecchetti et al. [146], Tubman et al. [133] and Li et al. [147] used a 'face-on' geometry to reduce the influence of electric fields and produce a radiograph mostly containing magnetic deflections. Cecchetti et al. [146] and Li et al. [147] both used additional 'side-on' of the interaction to measure electric fields without a significant magnetic contribution, while Tubman et al. [133] used a 45° probe which was sensitive to both electric and magnetic fields. Furthermore, Cecchetti et al. [146] was able to probe similar interactions from either the 'front' or the 'rear' side (relative to the face of the target being heated) by changing which side of the target foil was irradiated by the interaction laser beam. This provided particularly good evidence of mangetic field dominance, as the deflections reversed in direction when switching between front-side and rear-side probing geometry.

Du et al. [151] described a more sophisticated method decoupling electric and magnetic contributions to radiographs produced from a single shot and a single line of sight. This method exploits the differing velocity dependence of the Lorentz force due to electric and magnetic fields. Using a proton beam source which is not monoenergetic (both broadband TNSA and bi-chromatic fusion pellet sources fulfil this condition), it becomes possible to combine deflections of protons at different energies to isolate electrically- and magnetically-sourced deflections. However, due to the different probe times which arise for protons with different velocities, field evolution can complicate this kind of analysis.

5.2 The Inverse Problem

There are several approaches to solving the Monge-Ampère equation. For sufficiently small deflections, Kugland et al. [142] suggest solving this system by dropping from Eq. (5.20) all dependences on Φ other than its Laplacian, and thus solving the system as a Poisson equation. Graziani et al. [152] suggest that neglecting the $\boldsymbol{\nabla}\Phi$-dependence from the left-hand side of this equation causes the method to be very fragile, and derive an algorithm which includes a first-order expansion in $\boldsymbol{\nabla}\Phi$ of Eq. (5.20), thus ignoring only terms propotional to the square of the deflection strength, or equivalently the product of the two eigenvalues in Eq. (5.21). It is therefore clear that this 'linear' regime is applicable where $|\lambda^{\pm}| \ll 1$. This method has been implemented in the `python` package '*P*roton *Ra*diography *Line*ar reconstruction' or `PRaLine`[1] [153].

As the magnitude of deflections increases to be of order unity, the performance of this method will be increasingly degraded, since effects at second order in deflection strength are ignored. This is known as the 'nonlinear injective' regime. Significant advancements in the analysis of proton radiography images were made in publications by Kasim et al. [21] and Bott et al. [143]. Both recognised the form of the problem as a Monge-Ampère equation and used this fact to suggest analytical advancements.

[1]The package's authors note that the 'reconstruction' is silent

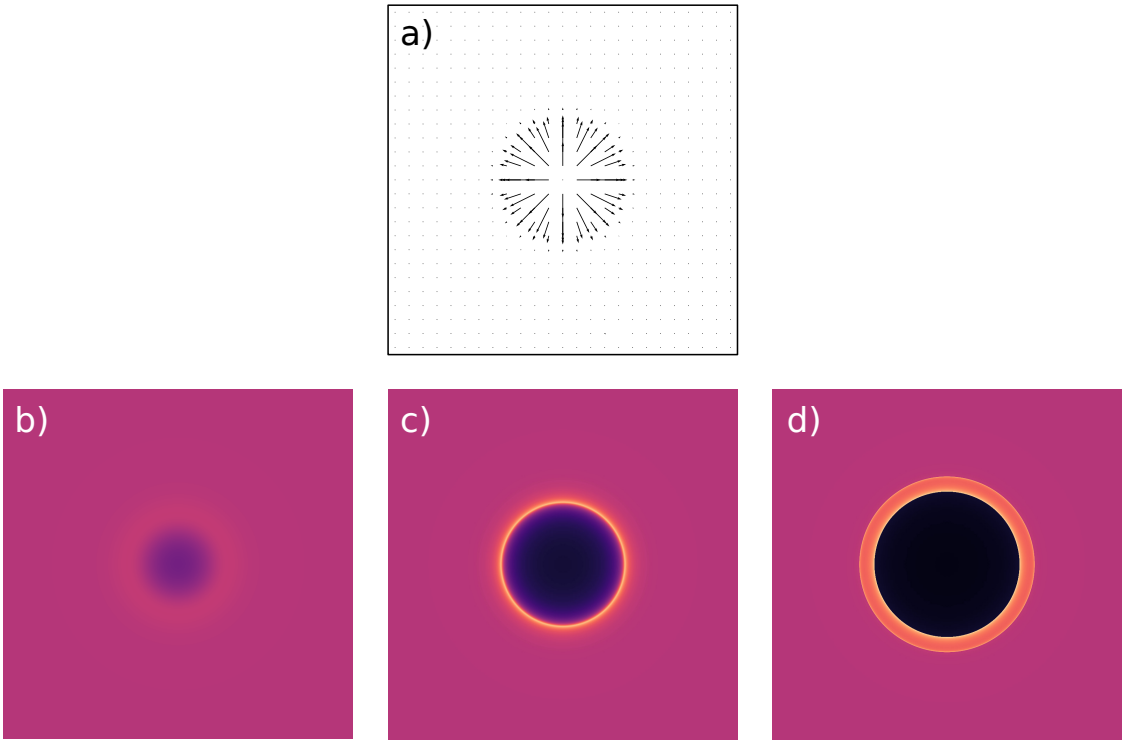


Figure 5.2: A deflection field and the fluence patterns it can cause: (a) a field of deflections represented by arrows, corresponding to a repelling Gaussian deflection potential. (b-d) Fluence patterns produced by simulated protons after they have been affected by this deflection field. (b) corresponds to weak fluence modulations in the linear regime, (c) to strong modulations in the nonlinear injective regime, and (d) to the caustic regime. Note the high-intensity structure of finite width in (d), which is a region with $(1 + \lambda^-) < 0$, where the 'proton sheet' has 'folded' over itself as it propagates. (b-d) can be thought of as the transverse fluence history of a single proton beam as it propagates through space after being deflected by a plasma.

Kasim et al. [21] used computational geometry techniques to solve the Monge-Kantorovich ideal transport problem, which is an equivalent reformulation of the Monge-Ampère equation. Bott et al. [143] suggested a finite-difference scheme based on the closely-related logarithmic parabolic Monge-Ampère equation using the algorithm of Sulman, Williams, and Russell [149], which was implemented by Laune, Feister, and Tzeferacos [154] in the `python` package '*PRO*ton-imaged *B*-field non*L*inear *E*xtraction *M*odule', or `PROBLEM`. This code computes the deflection potential Φ by evolving the logarithmic parabolic Monge-Ampère equation:

$$\frac{\partial \Phi}{\partial t} = \ln \frac{\det\left[1 + \boldsymbol{\nabla}\boldsymbol{\nabla}\Phi\right] \rho_i'(\boldsymbol{x}_0 + \boldsymbol{\nabla}\Phi)}{\rho_0(\boldsymbol{x}_0)} \tag{5.22}$$

While this method is significantly faster than that of Kasim et al. (due to

its lesser algorithmic complexity), it requires many algorithm steps to achieve a reasonable level of convergence and requires careful tuning of the algorithm parameters to avoid numerical instability. Worse, instabilities often only become apparent after many thousands of steps of the algorithm, forcing the user to change parameters with a restrictively slow feedback loop. To improve this situation I have implemented an improved solver, based on a linearisation of the Monge-Ampère equation, which displays much improved performance and robustness. This solver will be described in the next section.

By considering deflections with magnitude smaller than unity I have thus far managed to avoid discussion of the obvious divergence of Eq. (5.21) when either of $\lambda^{\pm} = -1$. In regions where either eigenvalue crosses the critical value of negative unity, injectivity of the mapping is lost. This leads to formation of *caustics*, defined as the locus of points where the local fluence is predicted to diverge according to Eq. (5.21). These caustic features can, for example, be formed by 'folding' of the population of protons such that protons in distinct positions before formation of the caustic are mapped to identical positions at the detector. In the presence of caustics, the theoretical correspondence between this problem and the Monge-Kantorovich problem (equivalently, the Monge-Ampère equation) is broken, as Eq. (5.21) must then be summed over each 'branch' of the caustic to compute the detector-plane proton fluence. An example of a caustic fluence distribution is shown in Fig. 5.2(d).

In the case of stochastic fields, further increasing the magnitude of deflections such that $|\lambda^{\pm}| \gg 1$ leads to what Bott et al. [143] termed the 'diffusive' regime. In this regime much of the structure of the image is blurred by the strength of deflections, and all that can be recovered is statistical information regarding their distribution.

5.3 A New Solver for `PROBLEM`

As discussed in the previous section, the original algorithm implemented in the `PROBLEM` code suffers from slow convergence and is not robust against numerical instability. In order to conduct the analysis presented in §5.4, I investigated other methods present in the literature and found that an algorithm based on the work

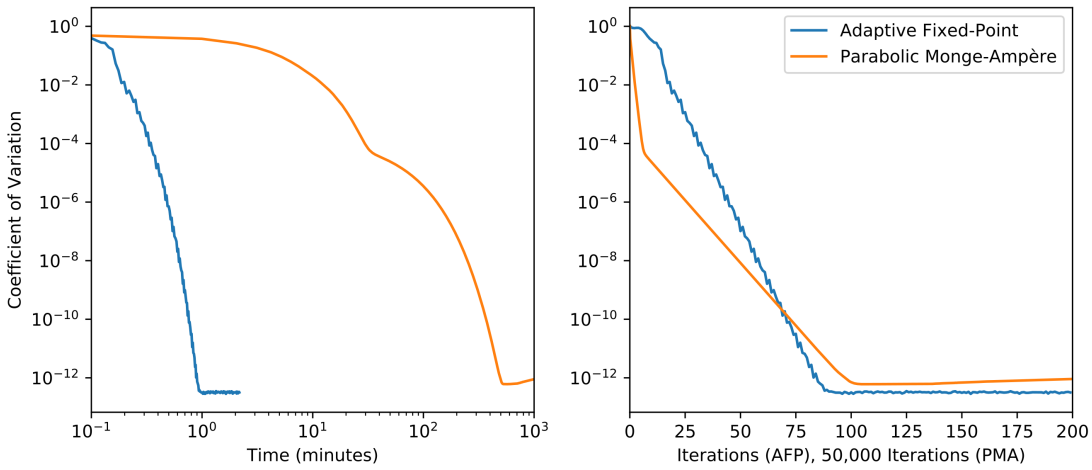


Figure 5.3: A comparison between the parabolic Monge-Ampère (PMA) approach previously used in the `PROBLEM` code and the adaptive fixed-point (AFP) approach I present here. Times are shown on a logarithmic scale, and, to make the graph of iterations legible, the number of iterations for AFP has been scaled by a factor of fifty thousand. PMA needs approximately five million iterations to reach a Coefficient of Variation of 10^{-12}, a level of accuracy reached by AFP in fewer than one hundred steps. The iteration time of AFP (100 iterations per minute) was two orders of magnitude longer than PMA (ten thousand iterations per minute), but due to the greater stability of the AFP method it can take much larger steps toward the solution, reducing the overall solution time from over fifteen hours to around one minute.

of Browne et al. [155] showed significantly improved convergence and stability characteristics. I wrote an implementation of this algorithm using the `python` programming language, including `numpy` to accelerate array-based operations and `scipy`'s sparse linear algebra interface to the `FORTRAN` library `umfpack`, as this method involves a single sparse matrix solve per iteration.

In short, the adaptive fixed-point algorithm works by expanding the Monge-Ampère equation around a current estimate $\Phi^{(k)}$ of the deflection potential, as a function of $d\Phi = \Phi^{(k+1)} - \Phi^{(k)}$. Browne et al. [155] found that linearising the Jacobian term was sufficient to produce good stability and accuracy. The linearisation is:

$$\boldsymbol{\nabla} \cdot \left(\underline{\underline{\mathrm{A}}} \cdot \boldsymbol{\nabla} d\Phi \right) = \frac{\rho_0(\boldsymbol{x}_0)}{\rho_i'(\boldsymbol{x}_0 + \boldsymbol{\nabla}\Phi)} - \det\left[1 + \boldsymbol{\nabla}\boldsymbol{\nabla}\Phi^{(k)} \right]. \tag{5.23}$$

This has the form of a diffusion equation with anisotropic, inhomogeneous diffusion tensor $\underline{\underline{\mathrm{A}}}$ whose components are:

$$\underline{\underline{\mathrm{A}}} = \begin{bmatrix} 1 + \Phi_{yy} & -\Phi_{xy} \\ -\Phi_{yx} & 1 + \Phi_{xx} \end{bmatrix}, \tag{5.24}$$

where I have used x and y subscripts to express partial derivatives.

While the new algorithm is significantly slower (by a factor 100 on tested hardware[2]) than the previous one to take a single step (due to the computational complexity of solving a large sparse matrix equation), it converges to a comparable extent in several orders of magnitude fewer iterations, and the overall effect is an improvement in runtime by a factor of order 500. A comparison between the two methods is shown in Fig. 5.3, using the coefficient of variation proposed by Browne et al. [155] as a measure of convergence. I also made improvements to the previously implemented parabolic Monge-Ampère algorithm, in order to make the comparison as fair as possible. I changed the discretisation scheme to ensure that the parabolic Monge-Ampère solver correctly obeys the Neumann boundary conditions imposed by the theory, and changed the discretisation of the second derivative to avoid checkerboarding artefacts in the solution.

Furthermore, I have integrated this new solver with the existing `PROBLEM` `python` package, allowing a user to employ the new, improved algorithm by simply providing a command-line option to the script or specifying a different parameter value when calling `PROBLEM`, if they are integrating it into a larger `python` program. In addition to the new algorithm, I have also developed a simple, optional plotting routine integrated into `PROBLEM`, which displays the current state of the solution in an intuitive way. This allows the user to assess how quickly the method is approaching convergence and easily spot numerical instabilities if they arise.

This new set of capabilities for `PROBLEM` was important for the successful analysis of proton radiographs I present in the next section.

5.4 Experiment: Whole-Beam Self-Focusing at Orion

The work presented in this section has previously been published in the peer-reviewed literature and can be found at Ref. [12]. I conducted new, quantitative analysis of proton radiographs produced at the Orion laser facility, during an experiment

[2]performance is most strongly bound by an Intel i7-8550U CPU

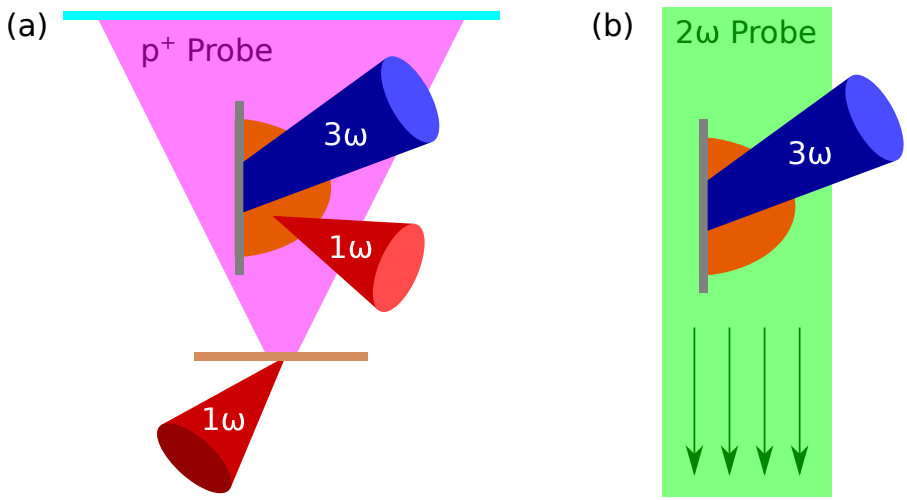


Figure 5.4: Sketch of the experimental setup. (a) A third-harmonic heating beam pre-forms a plasma from the surface of a 50 µm-thick Mylar foil. After a delay, a fundamental-frequency short pulse is focused into the plasma, to trigger whole-beam self-focusing. A second short pulse drives a TNSA to probe magnetic fields in the plasma by proton radiography. (b) A second-harmonic probe beam is used to produce interferograms of the formation and expansion of the pre-formed plasma. Due to operational constraints the short pulse and the probe beam were not used simultaneously; (a) and (b) are mutually-exclusive.

led by the Atomic Weapons Establishment's Dr. Matthew Hill, in collaboration with the Norreys group. The experiment investigated the whole-beam self-focusing mechanism of laser channelling, by which it is hoped that laser energy could be delivered efficiently to a compressed interial confinement fusion fuel assembly in future fast ignition experiments.

The whole-beam self-focusing mechanism[3] was proposed by Habara et al. [157], Kodama et al. [158], and Tanaka et al. [159] at the Institute for Laser Engineering, Osaka University. In essence it proposes combining two relativistic effects—relativistic self-focusing and relativistically-induced transparency—to improve the ability of ultra-intense laser pulses to penetrate into plasmas of long scale length. Such plasmas arise in fast ignition-relevant fusion regimes, and the proposed application of the mechanism was to bore a channel into a compressed fusion plasma with an intense laser pulse. This channel could then be used to guide an ignitor pulse into the compressed core, improving the efficiency of energy delivery.

[3]Referred to by the original authors as 'super-penetration', and sometimes termed the 'HKT' mechanism, after the authors. The name I use here corresponds to the terminology used in [156]

This mechanism has two components: first, relativistic self-focusing (RSF) is used to reduce the spot size of a laser pulse below its diffraction limit in vacuum. Matsuoka et al. [156] find that in order to suppress filamentation instabilities, this spot size (i.e. the beam diameter) must be smaller than the plasma wavelength when it reaches the quarter-critical point along the density gradient:

$$2w \leq \lambda_p \quad \text{at} \quad n_e = n_c/4 \tag{5.25}$$

Using the relationship between plasma frequency and density, this means the beam waist must be smaller than its wavelength when it reaches quarter-critical density, and it is this exceptionally tight focusing that RSF achieves. As the beam is focused by RSF, its intensity correspondingly increases. This causes an increase in relativistically-induced transparency of the plasma as the beam moves up the density gradient, improving its ability to penetrate into even higher-density plasma.

In this experiment, a 50 µm thick Mylar foil was heated by a 90 J laser pulse 200 ps in duration, at a wavelength of 351 nm (the third harmonic of Nd:glass). This pulse was focussed to a spot of 300 µm diameter, resulting in an intensity of $6 \times 10^{14}\,\mathrm{W\,cm^{-2}}$. This forms a plasma gradient that decreases from solid density at the original target surface to well below critical density near the expansion front, with a density scale-length of approximately 220 µm to emulate the properties of the plasma surrounding a compressed fusion plasma, in particular its long scale-length.

Orion's short-pulse capability was used to provide a short, intense channelling pulse. This pulse, operated at the fundamental frequency (wavelength of 1.053 µm), provided 100 J of laser energy in a typical pulse duration of 500 fs [160], though autocorrelator pulse duration measurements were unsuccessful during this experiment. The short pulse was focused to a spot of diameter 10 µm for an intensity of $2.5 \times 10^{20}\,\mathrm{W\,cm^{-2}}$. Some shots utilised a channelling pulse converted to the second harmonic (527 nm wavelength), but at the time the experiment was carried out, the available second-harmonic energy was limited to 70 J. This difference in the available energy prevents direct comparison of the effectiveness of the two

frequencies. No evidence of channelling was found using these second-harmonic pulses, so I consider only fundamental-frequency shots in the following analysis.

The primary diagnostic used to probe the short-pulse channel formation was proton radiography. During the channelling process, electrons in the channel are accelerated in the laser forward direction. The resulting current drives azimuthal magnetic fields around the interior of the channel, while a return current sheath along the outer wall of the channel prevents the fields extending beyond the channel itself. The magnetic field structure within the plasma therefore acts as an indicator of the shape and extent of the channel, and this field structure was probed by proton radiography.

5.4.1 Long-Pulse Interaction

Simulations

I used the radiation-hydrodynamics simulation code `FLASH` 4.6.2 to model the long-pulse interaction. For details about radiation-hydrodynamics simulations recall §3.1.3.

Because magnetic fields are the primary observable in this experimental campaign, it was essential to be able to model their evolution in the expanding plasma after the short-pulse interaction finished, as the proton beam probed the plasma several hundred picoseconds after the short pulse. For this reason I configured `FLASH` to enable magnetohydrodynamic (MHD) physics. I also wanted to model the generation of magnetic fields as fully as possible, to evaluate the relative effects of magnetic fields generated by the long-pulse heating and the short-pulse channelling interaction. To this end I enabled Biermann battery [161], the magnetic field source term supported by `FLASH` [162]. To improve numerical stability I used the relatively diffusive Local Lax-Friedrichs scheme [163] and three-temperature energy apportioning based on the method used by `CRASH` [164].

The target used in these simulations consists of a 50 µm foil of C-H plastic, chosen to match the thickness and material composition of the Mylar foil used at Orion. A laser pulse matching experimental parameters—90 J at third-harmonic

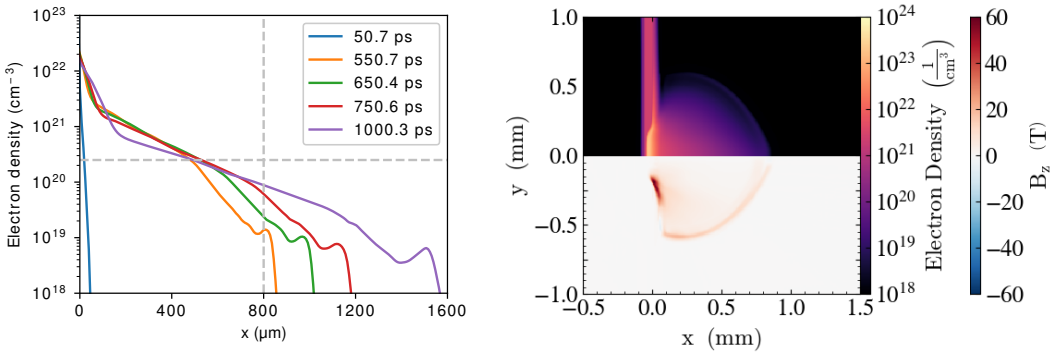


Figure 5.5: Left: Density profiles extracted from `FLASH` simulation of plasma formation by long-pulse heating, at times relevant to the short-pulse injection. Most shots fired the short pulse between 550 ps and 750 ps. Right: The density profile and Biermann battery-generated azimuthal magnetic field from `FLASH` at 550 ps. At its maximum, the magnetic field has attained a strength of 50 T.

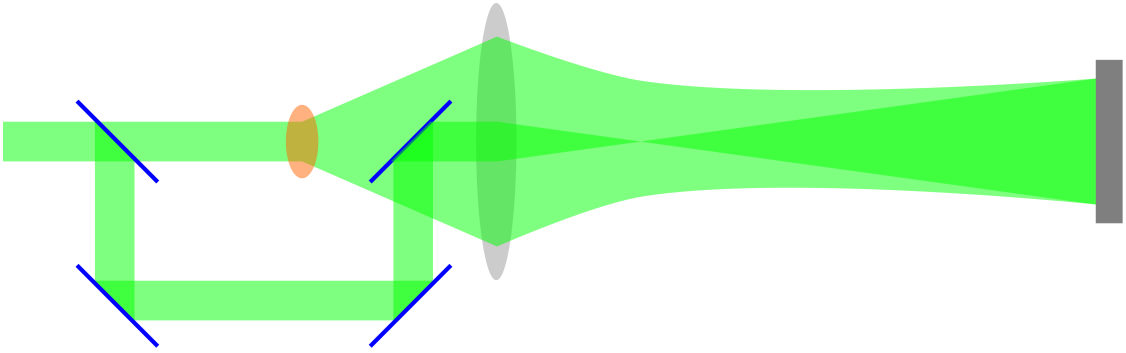


Figure 5.6: Simple schematic of an imaging interferometer. Probe light scattered by plasma refraction is gathered by a collector lens and imaged on the camera. The light is combined with a reference beam, producing fringes that shift according to the phase shift accrued by the probe beam in the plasma.

frequency, 200 ps duration, 300 µm-diameter spot—is incident on the target, ablating plasma from its front surface. The density profile of the plasma is shown in Fig. 5.5 at several times corresponding to times of short pulse interaction for different shots in the experiment.

Experimental Results

During the experiment the density profile was probed using both streaked and gated optical interferometry diagnostics. A simplified schematic of this diagnostic is shown in Fig. 5.6. Streaked interferometry gives an indication of the time evolution of the plasma density profile, over a period of 10 ns and a spatial extent of 2.25 mm

along the target normal, and gated interferograms correspond to a snapshot of the density profile at a specific time. Results of the streaked and $t_0 + 1\,\text{ns}$ gated interferometry diagnostics are shown in Fig. 5.7(a) and (c) respectively.

Common interferometric analysis techniques were rather limited in their utility here, as they rely on tracking the shifting of fringes across the image to recover the phase shift acquired by probe light as it propagates through the plasma. In the present interferograms however, there are large regions extending up to 650 µm from the target in which no fringes are formed. Because of the the large density gradients (and therefore refractive index gradients) present in the plasma, significant refraction occurs near the target's surface. In some places this causes the fringe spacing to become too dense for the camera to resolve, and closer to the target it seems to cause loss of light from this part of the probe beam altogether.

Unfortunately the regions of lost light are those of most interest when characterising the experimental conditions - observe that in Fig. 5.7(a), the entire region between the short pulse focus and the target surface exhibits fringe loss.

To allow some comparison to be made between the experimental interferometry and the density profiles predicted by simulation, I instead used the simulated density profiles to produce synthetic interferograms for comparison to the experimental interferometry diagnostics, shown in Fig. 5.7(b) and (d). These interferograms exibit qualitative similarities to those produced in the experiment, in particular in the interface between the regions with and without fringes, and those with apparent total probe beam loss.

The qualitative similarity between the experimental and synthetic radiographs provides only limited evidence in favour of the `FLASH` simulation-produced density profiles – especially as the key part of the profile nearest to the target is inaccessible by the interferometry diagnostic. In the absence of better characterisation of the density profile, the simulated density profiles are judged to be at least a reasonable approximation. If this experiment were to be repeated, use of a higher-frequency (e.g. fourth harmonic) probe beam, or a diagnostic designed specifically for high-density, long scale-length plasmas such as the angular filter refractometry technique

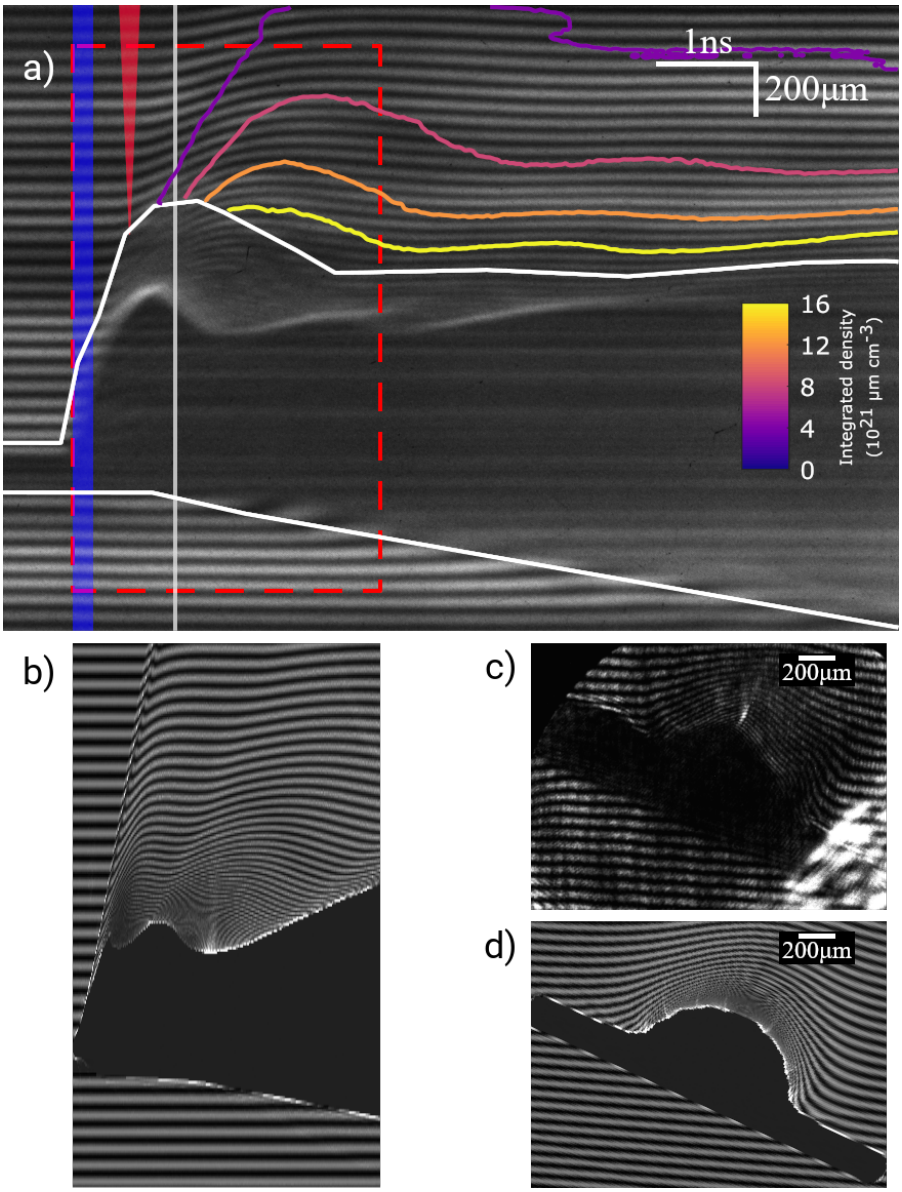


Figure 5.7: a) Experimental streaked optical interferometry image of the long-pulse interaction, using 10 ns sweep time. The first 1.5 ns recorded, before the interaction begins, are not shown. The region without defined fringes is outlined in white, and recovered density contours are shown outside of this region.These densities agree with the FLASH simulations to within a factor of 3-4. A blue bar indicates the time profile of the heating beam and a red triangle indicates the short pulse time and focal point. b) Synthetic streaked interferometry image produced from the plasma density profile extracted from `FLASH` simulations with 526.5 nm probe light. The dashed red box in a) indicates the region corresponding to b). c) Experimental gated interferogram and d) Synthetic gated interferogram taken at $t_0 + 1$ ns (indicated with white bar in a)). b) and d) simulate a misalignment of the collecting lens by 1 cm to optimize agreement with a) and c), respectively. The regions in which light is lost are 470 ± 10 µm in the gated images and 600 µm to 650 µm in the streaked images.

proposed by Haberberger et al. [22], may enable more meaningful information about the density profiles to be extracted.

5.4.2 Short-Pulse Interaction

Simulations

I used the particle-in-cell simulation code `Smilei` to model the short-pulse interaction. To reduce computational costs relative to simulations in full three-dimensional Cartesian geometry, I opted to use `Smilei`'s support for azimuthal decomposition of the simulation domain. This works by effectively evolving independent Fourier components of simulation quantities and can also be considered as a 'cylindrical harmonic' decomposition. This allows deviation from perfect cylindrical symmetry, while only increasing computational costs by a factor of order unity compared to two-dimensional simulations, whereas three-dimensional Cartesian geometry involves much greater increase in computational cost.

In the `Smilei` simulations, I initialised a plasma profile with density and temperature extracted from the above-mentioned `FLASH` simulations along the axis of propagation of the short pulse. I then injected a laser pulse matching the experimental parameters (100 J, 500 fs, focused to a 10 µm spot size) from the right (positive x) wall of the simulation box. The evolution of the laser pulse and the plasma was tracked until the pulse energy had been entirely absorbed by the plasma. The resulting density and magnetic field structures are shown in Fig. 5.8.

The confinement of laser energy close to the laser axis is of interest as it is one of the hallmarks of the whole-beam self-focusing model. Fig. 5.9 shows the integral of laser intensity over the duration of the simulation, as a demonstration of its reducing spot size. It shows that the majority of beam energy is confined within a plasma wavelength of the laser axis as the pulse penetrates into the plasma, and that very little energy extends beyond two plasma wavelengths from axis, supporting the interpretation of this as a channel produced by whole-beam self-focusing.

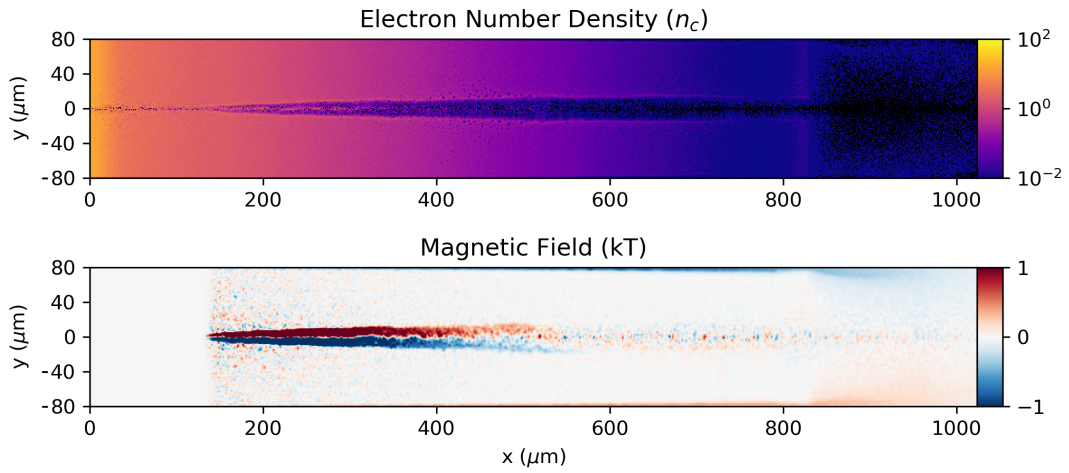


Figure 5.8: Electron density and magnetic field strength slices from a cylindrical particle-in-cell simulation, 4.4 ps after the injection of the short pulse at the right wall. At this time the pulse has been fully absorbed. A channel has formed and reaches approximately 130 µm from original target surface, which is located at $x = 0$. Moving away from the target, the channel has opened up to around 40 µm in diameter, but is closer to 10 µm at its point of deepest penetration.

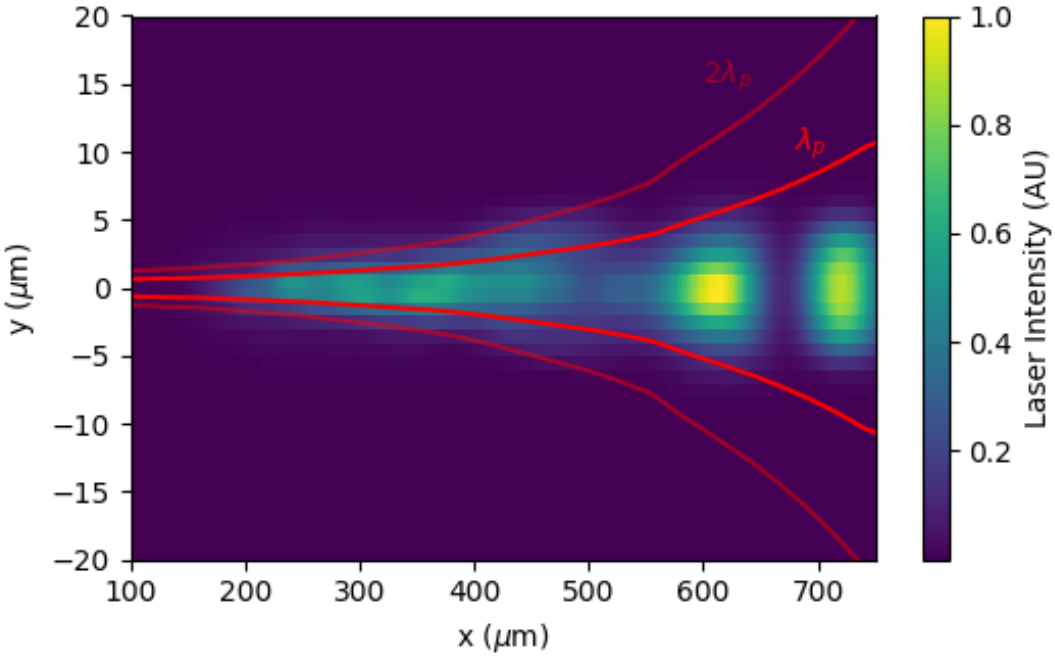


Figure 5.9: Laser intensity distribution, integrated over the simulation's duration. This gives an idea of the path taken by the laser pulse to determine whether it fulfils the HKT condition Eq. (5.25). As the beam propagates its energy is indeed mostly confined within a radius λ_p of the laser axis, and almost entirely within twice that distance from the axis. Note the unequal scaling of the x and y axes

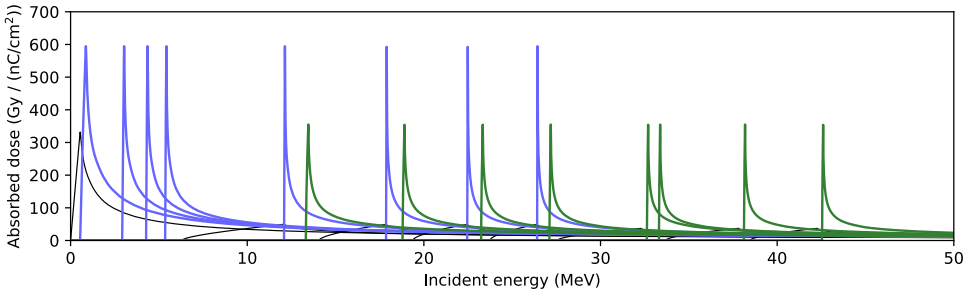


Figure 5.10: Dose delivered to each active layer of a typical radiochromic film stack (per unit charge fluence) as a function of incident proton energy. Taller, blue peaks are HD810 film and shorter green peaks are EBT3 film. Doses delivered to iron and aluminium filter layers are shown in thinner black lines. Dose delivered to inert substrate and adhesive layers of films are not shown.

Experimental Results

I now return to considering proton radiography, the subject of this chapter. The magnetic fields present in the plasma after injection of the short pulse were probed using a TNSA-generated proton beam, driven by the other of Orion's two short pulses. After passing through the plasma, the protons were detected using stacks of GAFChromic™ radiochromic film. These stacks were constructed as described in §5.1.2, using many layers of RCF and several metal filter layers. The stacks were constructed for an optimal sensitivity range of approximately 2 MeV to 40 MeV (see Fig. 5.10). The per-layer, energy-resolved sensitivity of a typical RCF stack in this experiment is shown in Fig. 5.10

An example of two layers of radiochromic film from a representative shot is shown in Fig. 5.11. These typify issues which are characteristic of proton radiography of short-pulse interactions. Many different processes produce competing electric and magnetic field structures in short-pulse interactions, with the result that proton radiographs tend to be rich in detail but difficult to interpret. Several filamentary structures are apparent, but it is not clear which of these might indicate the presence of a channel and which arise through other processes.

To try and interpret these radiographs more clearly, I used the algorithm described in §5.3 to recover the proton deflections that produced them. First, I calculated the dose distribution absorbed by each layer of radiochromic film using

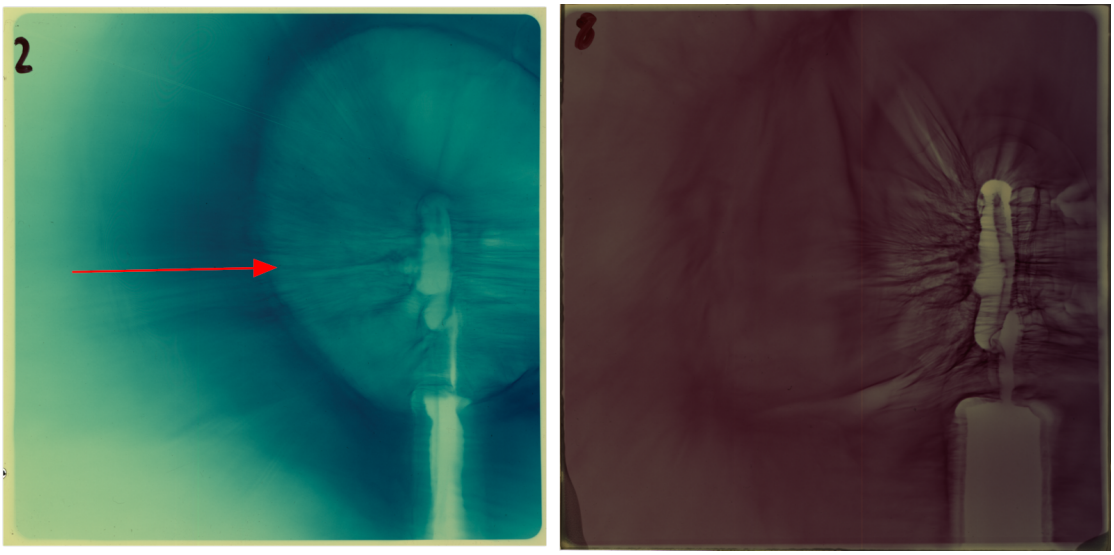


Figure 5.11: Example layers of radiochromic film from a representative shot. The layer labeled '2' (left) is HDV2 high-dose (low-sensitivity) film, and the one labeled '8' (right) is EBT3 low-dose (high-sensitivity) film. The numbers 2 and 8 relate to the layers' positions in the stack (refer to Fig. 5.10 for the stack configuration). Layer 2 was most sensitive to protons of 3 MeV energy (which arrived 630 ps after the channelling pulse was fired) and Layer 8 to 18 MeV protons (145 ps). Each layer of film is a 10 cm square, and the system magnification was 11, leading to a field of view at the plasma of 9 mm by 9 mm. The direction of short-pulse injection is indicated with a red arrow on layer 2.

known dose-optical density calibration data for the types of film present in the stack. This dose image was used as the input to the reconstruction algorithm, which computed the proton deflections necessary to produce the observed dose distribution. The algorithm requires a source profile to compare to the observed dose profile; as a basic appproximation I used a low-pass filtered version of the observed profile. The result of running the algorithm on several layers of the stack are shown in Fig. 5.12. The recovered line-integrated magnetic current profiles show some structure which may correspond to channel formation, but there are many electromagnetic features present in the plasma that could easily confound a claimed field structure. For instance, target charging is visible as a 'ballooning' of the proton-free region immediately around the target, and filamentary structures are present behind the target itself. In the presence of so much structure it is very difficult to attribute one piece of structure in particular to the presence of a channel.

A few shots produced proton radiographs of the pre-formed plasma without the

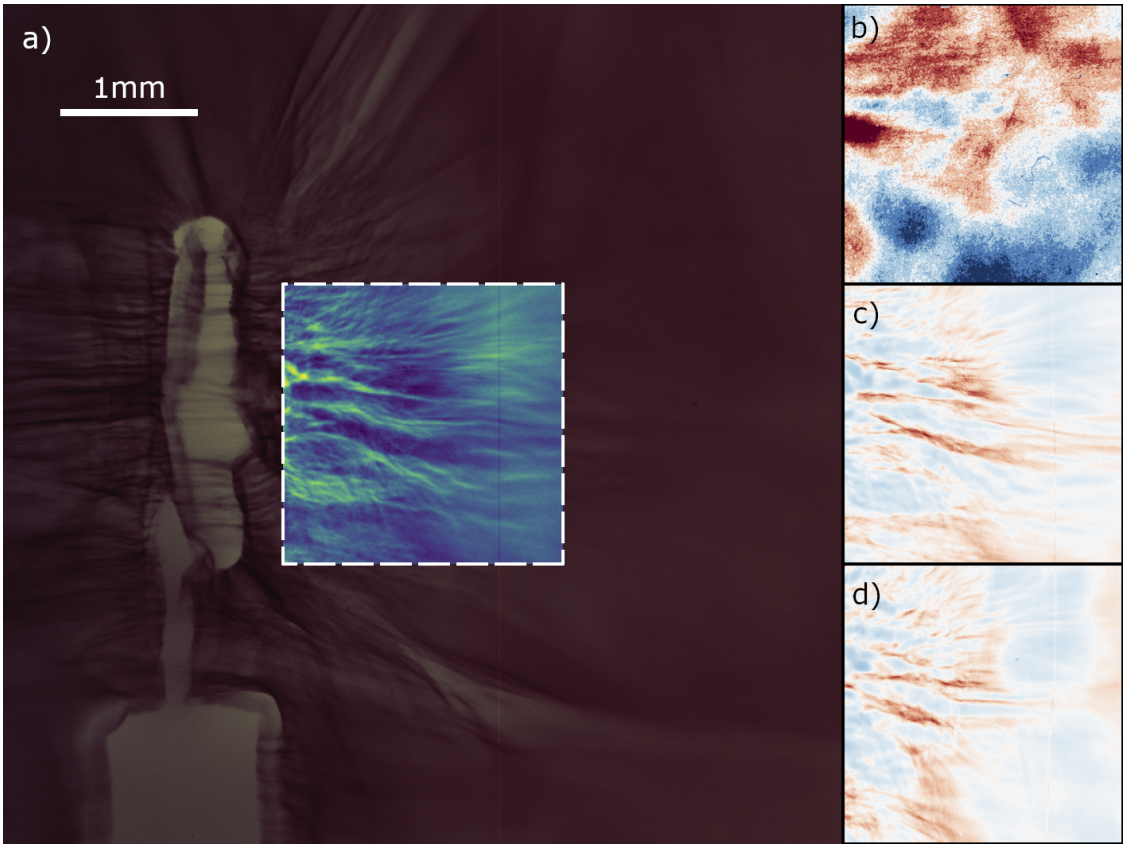


Figure 5.12: Results of quantitative analysis of a stack of radiochromic film. a) A section of an RCF layer. Dose levels (arb. units) are shown within the central region of interest, and the raw scanned image outside. Reconstructed longitudinal integrated MHD current densities are shown in b)-d), for three layers corresponding to energies of a) 1 MeV (1.25 ns), b) 13 MeV (200 ps) c) 23 MeV (105 ps), with times quoted the delay between the arrival of the channelling pulse and the proton probe time. Positive integrated current densities correspond to currents counter-propagating with the proton beam, and produce a defocusing effect.

channelling pulse, due to amplifier failure. The absence of the violent short-pulse interaction on these shots resulted in significantly less noisy radiographs, which are consequently easier to analyse. An example is shown in Fig. 5.13. This radiograph probed the plasma 670 ps after the onset of the long heating pulses (and so 470 ps after heating ended). There is little qualitative difference between the layers of this stack, indicating that this field structure remained steady for at least several hundred picoseconds. Combined with the fact that the lasers are operating at intensity of only $6 \times 10^{14}\,\mathrm{W\,cm^{-2}}$, this indicates that the plasma is in the MHD regime, and can therefore be assumed to contain negligible electric fields.

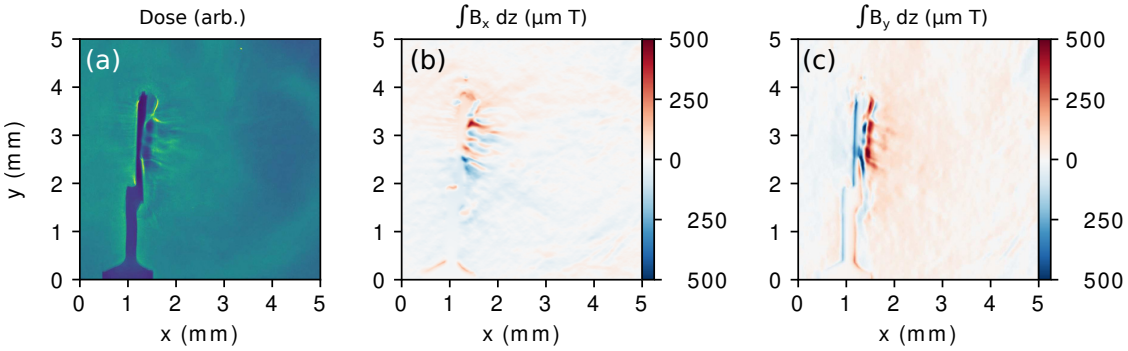


Figure 5.13: Quantitative analysis of a film layer from a long pulse only shot, showing structures associated with the plasma formation and expansion. This layer was most sensitive to 19 MeV protons, which probe the plasma 670 ps after the onset of long-pulse heating. Integrated magnetic field values of order 500 T µm are observed.

My analysis finds line-integrated magnetic field values with a maximum of order 500 T µm. The transverse extent of the magnetic field structure appears to be 1 mm, so a naïve estimate of the path-averaged magnetic field strength is approximately 0.5 T. This is small compared to the fields shown in Fig. 5.5, but this is to be expected: it is not the *peak* magnetic field strength, but rather its *average*, in a side-on geometry which should be expected to produce cancellation of magnetic fields (and is assumed to in several studies in the literature [133, 146, 147]). It would be concerning if the average value were found to be greater than the expected peak value, but an average significantly below the peak is to be expected.

A possibly more realistic estimate may be derived by examining the size of the magnetic field region in Fig. 5.5. It measures around 100 µm, leading to an average field strength estimate of 5 T. This still does not account for cancellation, however.

5.5 Summary

In this chapter I introduced the reader to the topic of proton radiography, an electromagnetic field diagnostic commonly used in laser-plasma experiments. After summarising the key theory of proton radiography, I described a new algorithm I have implemented in the community proton radiography analysis code `PROBLEM`, and demonstrated that it reaches convergence orders of magnitude faster than

the previous state-of-the-art while also being significantly more robust to large intensity modulations.

This code formed an important piece of my analysis of an experiment which took place at the Orion laser facility. In this experiment, a short, intense pulse of laser light was focused into a plasma profile that was pre-formed by long-pulse, third-harmonic laser irradiation. The short pulse was arranged to focus significantly prior to the original surface of the solid target, and before critical density. This was done to encourage the HKT mechanism of whole-beam self-focusing, a.k.a. the 'super-penetration' mode. In this mechanism a combination of relativistic effects serve to focus the beam to a smaller spot than its focusing geometry would produce in vacuum, while also protecting it from beam filamentation. A simulation study comprising a campaign of coupled fluid and particle-in-cell simulations indicated that whole-beam self focusing could take place in the plasma conditions produced by the long-pulse irradiation.

Unfortunately, proton radiographs produced on channelling shots showed a lot of confounding electromagnetic effects, making attribution of any particular feature to a laser-produced channel impossible and rendering this line of analysis inconclusive. A quantitative measurement of the magnetic fields present in the absence of the channelling pulse was made however, and found to be consistent with (albeit in a fairly loosely constrained sense) the magnetic field generated in a `FLASH` MHD simulation by the Biermann battery. When analysing these radiographs—especially of short-pulse interactions featuring lots of fine structure and large, rapidly-varying proton fluence modulations—the robustness improvements of my new adaptive fixed-point implementation for proton radiography reconstruction were vital. When analysing some of these images, the previous parabolic Monge-Ampere code would run for hours before eventually becoming numerically unstable. The adaptive fixed-point code on the other hand converged in under a minute to a solution with smaller residuals than the old algorithm had reached before numerical instability took over.

The analysis of this experimental data was previously published in Philosophical Transactions of the Royal Society in December 2020 [12].

6
Proton Tomography

Contents

6.1 Tomography Basics 120
6.1.1 Analytic Methods 123
6.1.2 Algebraic Reconstruction Techniques 126
6.2 Applying Tomographic Theory to Proton Radiography 128
6.3 Improving Performance 129
6.3.1 Hankel Reconstruction and Fourier Interpolation 131
6.3.2 Aspect Ratio Compensation 134
6.4 Results 137
6.5 Summary 138

In the previous chapter I presented work that I carried out in order to improve the speed, stability and accuracy of magnetic field reconstruction from proton radiography images. This method is still only capable of extracting line-integrated field values from radiographs, as it is these line-integrated values that determine the total deflection of protons. It would be a significant improvement to be able to recover the spatially-resolved values themselves, but this requires more information than is present in a single radiography image. Tomography is the science of reconstructing functions from sets of line integrals, so it is tomography I turn to here.

Some existing work on using multiple proton radiography lines of sight comes from Tubman et al. [133], Cecchetti et al. [146], and Li et al. [147], two experiments

which observed the same interaction (or two similar ones) using proton radiography from two different angles. In these experiments, however, the interaction produced significant electric fields as well as magnetic fields, and the different probe directions were chosen to separate the contributions of the two fields, rather than to achieve three-dimensional reconstruction of magnetic field structures. Chen et al. [145] alluded to a tomographic approach, suggesting that additional radiographs may lift the degeneracy of proton radiographs, in the course of developing machine learning-based techniques to reconstruct fields. Some of the work that I present here is similar to results published by Volegov et al. [165–168], who applied few-view reconstruction methods to experimental data gathered at the National Ignition Facility. My approach differs from theirs in that they use algebraic reconstruction techniques (ARTs, see §6.1.2) to implement their reconstructions, whereas I base my analysis in filtered back-projection, an analytic technique. I also go further in the interpretation of the resulting method as an interpolation, enabled by my analytic approach[1].

6.1 Tomography Basics

Tomography—often under the guise of computed tomography or CT—is widely used across medicine, science, and industry. The most widely recognised application to a general audience will almost certainly be x-ray CT. This employs an x-ray source which rotates around the subject taking x-ray attenuation radiographs at many different angles of probing. The radiographs are then used to reconstruct a three-dimensional model of the subject's x-ray attenuation, which can be used to identify different types of tissues in medicine or identify internal cracks and defects in industrial applications. Another example is magnetic resonance imaging, which is used to observe the distribution of nuclear magnetic moments within a patient. Positron emission tomography is performed similarly, but detects gamma rays produced by electron-positron annihilation events that indicate the presence of a radioactive 'tracer' chemical in tissues of the body.

[1]Here I am using 'analytic' in the mathematical sense, roughly meaning the study of functions, their properties and calculus, rather than the scientific meaning of inspection of data to make conclusions.

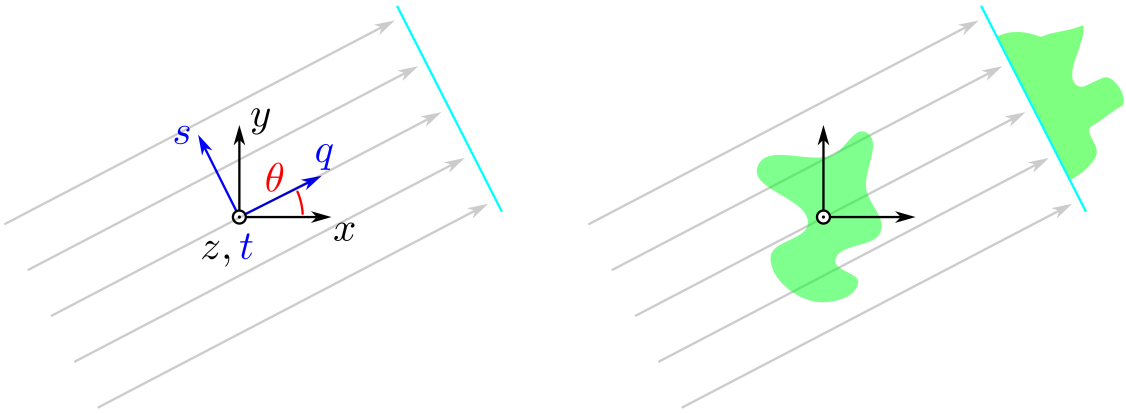


Figure 6.1: Geometry of tomographic projections. The fixed laboratory coordinates (x, y, z) and the rotated projection coordinates (q, s, t) are related by a rotation by the angle θ which defines a specific projection. The rotation is about the z-axis, such that $\hat{z}$ and $\hat{t}$ coincide. The detector is a screen (drawn in cyan) parallel to the s-axis. A single projection of a test function is shown in b), plotted as a filled-area plot at the detector screen.

In the field of laser-plasma interaction physics, there have been sporadic uses of tomographic methods. The Abel transform and its inversion see frequent use in diagnosis of laser-generated plasma, where the plasma electron density is assumed to carry cylindrical symmetry. In this limit, the inverse Abel transform is a simplified form of tomography that produces correct results from a single projection of the object[169–171].

In this section I will describe the basic theory of tomographic reconstruction and explain how the tomography of magnetic field structures might differ from the conventional applications listed above.

I will consider only *parallel-probe* tomography here. In the context of proton radiography, this means that the proton beam is assumed to be collimated until it interacts with the plasma. Target normal sheath acceleration and capsule implosion both produce divergent proton beams (in the latter case over a full 4π solid angle), but the distance between proton source and plasma is usually significantly larger than the transverse size of the plasma. The variation in incident proton velocity over the transverse extent of the plasma is therefore sufficiently small that treating the protons as collimated is a reasonable approximation for the purposes of tomographic reconstruction.

Before discussing tomography of magnetic fields in particular, I will consider the simpler case of tomographic probing of a scalar function $f(x, y, z)$ in order to provide a review of existing methods. Many of the above-mentioned computed tomography applications probe the x-ray attenuation of some object or part of a human body, and it may be helpful to think of f as the density of an object, measured, for example, by its x-ray attenuation. A tomographic projection of this function is parametrised by the probe angle θ. For a given θ we define a new Cartesian coordinate system (q, s, t) related to (x, y, z) by rotation about the z axis:

$$\begin{bmatrix} q \\ s \\ t \end{bmatrix} = \begin{bmatrix} x\cos\theta + y\sin\theta \\ y\cos\theta - x\sin\theta \\ z \end{bmatrix} \tag{6.1a}$$

$$\begin{bmatrix} x \\ y \\ z \end{bmatrix} = \begin{bmatrix} q\cos\theta - s\sin\theta \\ s\cos\theta + q\sin\theta \\ t \end{bmatrix}. \tag{6.1b}$$

Taking projections (i.e. computing line integrals) along the local q direction (as shown in Fig. 6.1) produces for each θ a function of s and t:

$$\begin{aligned} F_\theta(s,t) &= \int_{-\infty}^{\infty} f(q\cos\theta - s\sin\theta, s\cos\theta + q\sin\theta, t)dq \\ &= \mathcal{R}_\theta\left[f(x,y,z)\right](s,t) \\ &= \mathcal{R}_\theta\left[f(x,y,t)\right](s). \end{aligned} \tag{6.2}$$

This equation defines the Radon transform $\mathcal{R}_\theta$, the integral transform that is the theoretical basis of tomographic analysis. It is important to note that, due to our assumption of parallel probing, this is effectively a 'stack' of two-dimensional tomographs, one for each value in t. The values of $F_\theta(s,t)$ at fixed t are only influenced by the two-dimensional slice of the original function for which $z = t$: $f(x, y, t)$, for all s and θ. This allows two-dimensional inversion algorithms to be applied to the three-dimensional tomography problem.

The function $F_\theta(s)$ (whose dependence on t has been dropped following the previous argument) is often visualised as a 'sinogram': the parameter θ is promoted to a variable and the resulting two-variable function $F(s, \theta)$ is plotted as an image. An example of a sinogram is shown in Fig. 6.2. This sinogram was computed from the

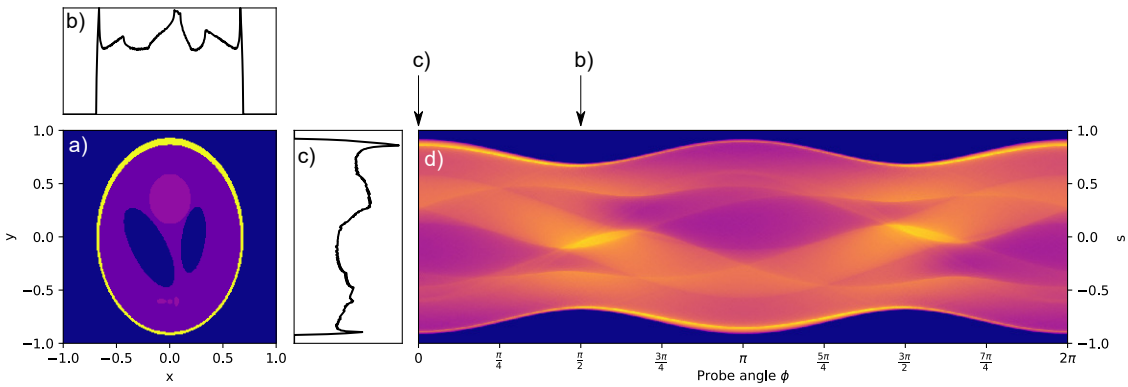


Figure 6.2: A demonstration of a sinogram. a) A pseudocolor plot of the function under observation: the Shepp-Logan phantom. b) and c) Projections of the function along the y and x directions respectively. d) Many such projections taken at a large number of finely-resolved angles are plotted as a function of both projection angle and the transverse spatial dimension. This form of image is known as a *sinogram*.

Shepp-Logan phantom, a function often used to test tomographic techniques [172]. A modified version of the Shepp-Logan phantom with both positive and negative values is employed in the following sections, as this better represents the nature of magnetic field structures.

I will now summarise two important classes of tomographic reconstruction algorithms. The first of these takes an analytic approach to reconstruction, using the Fourier projection-slice theorem as a starting point. The second is algebraic in nature. Higgins and Munson [173] consider direct Fourier reconstruction techniques and the filtered (or convolution) back-projection to be separate classes of algorithm, but here I have decided to group them together—along with the Hankel transform proposed in the same publication—because of their theoretical similarities.

6.1.1 Analytic Methods

The analytic methods for tomographic reconstruction all find their foundation in the Fourier projection-slice theorem. This theorem states that the one-dimensional Fourier transform $\tilde{F}_\theta(k_s)$ of the projection $F_\theta(s)$ is equal to the two-dimensional Fourier transform $\tilde{f}(k_x, k_y)$ of the original function $f(x, y)$, evaluated on a one-dimensional slice through the origin of frequency space normal to the probe direc-

tion[2]:

$$\tilde{F}_\theta(k) = \tilde{f}(-k\sin\theta, k\cos\theta) \tag{6.3}$$

Promoting the parameter θ to a variable, it is clear that $\tilde{F}(k,\theta)$ is in fact nothing more than a representation of $\tilde{f}$ in plane-polar coordinates, albeit with θ differing from the conventional polar angle variable by a quarter-cycle.

$$\tilde{F}(k,\theta-\pi/2) = \tilde{f}(k\cos\theta, k\sin\theta) \tag{6.4}$$

Therefore, in principle, all that is required for reconstruction of tomographic data is a two-dimensional inverse Fourier transform of the one dimensional Fourier transforms of projected data. This procedure is analytically exact (although, of course, only in the limit of perfect, complete data), but complications arise due to the discrete sampling of real data. The Fast Fourier Transform (`FFT`) provides an algorithm for Cartesian discrete Fourier transforms that is both fast and accurate, but how should the inverse Fourier transform of data sampled discretely on a *polar* grid be computed?

Direct Fourier Reconstruction

Possibly the simplest answer to this question is the 'direct Fourier' reconstruction technique. Direct Fourier methods simply use some two-dimensional interpolation between the known values (located on a polar grid) to evaluate the Fourier transform at Cartesian grid points, then perform the reconstruction using a two-dimensional inverse Fourier transform.

While this method is simple, the highly oscillatory nature of many functions in the Fourier domain can significantly degrade the accuracy of the interpolation routine used and therefore also the reconstruction in real space.

[2]It may be easier to visualise this as the Fourier transform of a sinogram along its s-axis.

Filtered Back-Projection

A canonical algorithm for the recovery of tomographic data sets is the 'filtered back-projection' (FBP) method [174]. This algorithm is derived by discretising the integral defining the inverse Fourier transform in polar coordinates, producing a result which is expressed as a sum over all projections.

First, this method filters each projection with a kernel proportional to $|k|$ in Fourier space, producing a new set of filtered projections $G_\theta(s)$:

$$\mathcal{F}_1 \left[G_\theta(s) \right] = |k| \mathcal{F}_1 \left[F_\theta(s) \right], \tag{6.5}$$

where $\mathcal{F}_1$ represents the one-dimesional Fourier transform along the s-axis.

These filtered projections are then 'smeared' along their probe directions, and the resulting 'back-projected' functions are summed (approximating an integral over angle) to recover an approximation of the original function:

$$f(x, y) \approx \frac{\Delta\theta}{2\pi} \sum_{n=0}^{N-1} G_{\theta_n}(s(x, y, \theta_n)) \tag{6.6}$$

where $s(x, y, \theta) = y \cos\theta - x \sin\theta$ as per Eq. (6.1).

FBP converges to an analytically correct result in the limit of many projections[3], but it suffers from severe 'streaking' artefacts where samples are few and sparse in the angular dimension. This behaviour is demonstrated in Figure 6.3. The quality of reconstructions is quantified using the strucural similarity index metric (`SSIM`), a popular measure of image similarity, between the input and reconstructed functions [175].

An honorable mention should be given to the Abel transform and its inversion. The Abel transform is a special case of the Radon transform in the case of an axially symmetric function. In this case, the resulting projection is independent of projecion angle, and the original function can be reconstructed from only a single projection (in fact, only half a projection is needed, due to the also-implied symmetry of projections around $s = 0$). Abel transform inversion performs much better than

[3] this is because it can be derived as a discretisation of the Fourier projection-slice theorem discussed above in Section 6.1.1

traditional tomography techniques designed for large numbers of projections when confronted with sparsely-sampled data (see the discussion of §6.3), but can only incorporate information from a single projection. Use of a slightly altered algorithm can account for asymmetry in the single projection [171], and Cormack [169, 170] derived relations applicable to higher-order Fourier modes as well, though these do not appear to have seen wide use.

6.1.2 Algebraic Reconstruction Techniques

The second major class of tomographic reconstruction algorithms is the Algebraic Reconstruction Techniques (ARTs) [176, 177]. These iterative methods dispense with the pretence of considering the observations and reconstructions as continuous functions and deal directly with the discretised versions. They differ in their specifics, but, in general, all ARTs seek to converge to an approximation of the target function by considering the forward projection process as a linear transformation, which takes the function being probed to the observed data. As integration is linear, it is possible to express (a discretised version of) a single forward projection (indexing projection angles with i) as a matrix equation:

$$\boldsymbol{P}^i = \underline{\underline{\mathrm{M}}}^i \boldsymbol{F}, \quad 0 \leq i < N_a \tag{6.7}$$

where the function vector $\boldsymbol{F}$ contains the target function discretised at N_f points, the projection vector $\boldsymbol{P}^i$ contains the observations from projection i discretised at N_p points, and $\underline{\underline{\mathrm{M}}}^i$ is a projection matrix of size $N_p \times N_f$ that contains weights. These weights depend on the geometry of the tomography problem in question and the projection angle θ_i and are chosen such that when used to form weighted averages of the discretised function values, they approximate the projection integral at angle θ_i.

In high-resolution tomography, both the number of observed points N_p and the number of desired reconstruction points N_f can be very large[4], meaning that Eq. (6.7) can be infeasible to solve directly. In addition, the matrices $\underline{\underline{\mathrm{M}}}^i$ will in many cases not be square ($N_f \neq N_p$), so the system of equations represented

[4]for three-dimensional reconstruction, even a relatively modest reconstruction resolution of 100^3 produces N_f of one million and N_p of order ten thousand

by the matrix equation will be under- or (less commonly) overdetermined and therefore likely impossible to solve exactly.

In order to find an approximate solution to Eq. (6.7) (i.e. minimise $\sum_i |\underline{\underline{\mathrm{M}}}^i \boldsymbol{F} - \boldsymbol{P}^i|^2$), an iterative method is adopted. Starting from an initial guess at the function (such as $\boldsymbol{F}^{(0)} = 0$), the original ART method consists of successive iterations to update the approximation by the Kaczmarz method (sometimes referred to as projection onto convex sets, or POCS) to incorporate information about each pixel:

$$\boldsymbol{F}^{(k)} - \boldsymbol{F}^{(k-1)} = \beta \, \frac{P_m^i - \boldsymbol{M}_m^i \cdot \boldsymbol{F}^{(k-1)}}{|\boldsymbol{M}_m^i|^2} \boldsymbol{M}_m^i. \tag{6.8}$$

This formula is first iterated over m for each projected pixel P_m^i and corresponding row of the weight matrix $\boldsymbol{M}_m^i$ of a given projection i, and then the iteration is repeated for each projection. This means that the iteration number k is given by $N_p i + m$, with one 'outer' iteration (after which all data have been considered) corresponding to $k = N_p N_a$. The parameter $\beta \leq 1$ controls iteration step size, with smaller values of β giving slower but more stable convergence. The iterative nature of this technique allows *a priori* knowledge to be incorporated into the solution by enforcing constraints: for example, positivity of a reconstructed mass density could be maintained by setting any negative elements of $\boldsymbol{F}^k$ to zero after every iteration. This principle can be adapted to enforce many constraints on the solution $\boldsymbol{F}$. For example, simultaneously minimising the total variation (i.e. the image gradient) of the solution can improve reconstruction quality, especially in cases where parts of the true function are piecewise constant [178].

Several refinements of the basic ART algorithm have emerged, such as the simultaneous algebraic reconstruction technique (SART), in which the correction given by the right-hand side of Eq. (6.8) is calculated for each pixel of a single projection using the current estimate $\boldsymbol{F}^{(k-1)}$, and the update is made using a weighted average of all the corrections simultaneously. In this scheme iterations are still performed over projections, but all pixels in a projection are considered simultaneously. Iteration number k is therefore equal to i in this scheme.

If performing a truly three-dimensional reconstruction is desired (for example, in cone-beam geometry, where the different z-planes lose their independence), it is conceptually simple to generalise the ARTs simply by changing the definitions of the weights contained in the matrices $\underline{\underline{\mathrm{M}}}^i$. This allows ARTs to handle projection geometry consisting of arbitrary sets of probe directions (unlike many other methods, which require probes that are confined to a single plane).

6.2 Applying Tomographic Theory to Proton Radiography

I now turn my attention to the application of this tomographic formalism to proton radiography. This problem is, on the surface, more difficult than most tomographic methods, because the quantity it observes is a vector field, rather than a scalar field like density. Furthermore, not all components of this vector field are accessible: for any given probe direction only the components of the field orthogonal to that direction have any effect on the deflections of protons. This prevents probing along some arbitrary set of vectors in three-dimensional space and then transforming the field components into a common coordinate system (e.g. the laboratory frame) in order to use conventional scalar approaches on each component separately.

However, the fact that proton radiographs encode components of the magnetic field *transverse* to the probe direction proves to be very useful. It is easy to choose a set of projections such that one of the transverse coordinates for each projection is always parallel to some vector fixed in the laboratory frame (I will call this fixed vector direction the 'common axis'). This condition is automatically fulfilled by ensuring that every probe direction is orthogonal to that common axis. For example, if 'up' is chosen as the common axis then all probe directions must be horizontal. Such probe directions fall along the unit circle in the plane whose normal is given by the common axis, and any suitable probe direction is related to the others simply by rotation around that common axis. It makes sense, therefore, to call this common axis the rotation axis defining the probe geometry.

In the case of proton radiography, the cross product present in the Lorentz force means that the relevant proton deflection to observe is actually that *orthogonal* to the common axis, as these deflections are generated by the magnetic field component parallel to the common axis.

6.3 Improving Performance

Now that I have developed a mathematical framework under which tomographic methods can be applied to observations of magnetic fields extracted from proton radiography images, I will discuss some practical difficulties that I expect to affect the quality of these reconstructions. These difficulties would also apply to more conventional applications of tomography, but I will explain why they are harder to mitigate in this context. I will then present pre-processing methods that can help to improve the quality of reconstructions, given these difficulties.

The first issue htat must be overcome to guarantee good reconstruction quality is that the number of probe directions it is possible to deploy simultaneously is limited by the number of laser pulses available to drive proton beams. On short-pulse systems the number of pulses available with sufficient energy to produce TNSA-driven beams is commonly one, two or (in the particular case of the NIF-ARC system) four. This number is, of course, reduced further if the interaction being investigated is itself short pulse-driven.

Conventional CT systems often use only a single beam source (usually of x-rays) and detector, but rotate these around the target (or vice versa), taking a very large number of observations with probe directions orthogonal to the axis of rotation. This takes time on the order of minutes to complete, which is obviously not compatible with inspecting interactions of sub-nanosecond duration. Instead, it is necessary to find ways to maximise reconstruction quality and legibility with very small numbers of projections: orders of magnitude fewer than conventional CT considers to be a sparse sampling rate. As an example, Wu et al. [179] quote 75 as a 'very sparse number of projection angles' when compared to the Nyquist sampling rate of 2800 (for the system they consider).

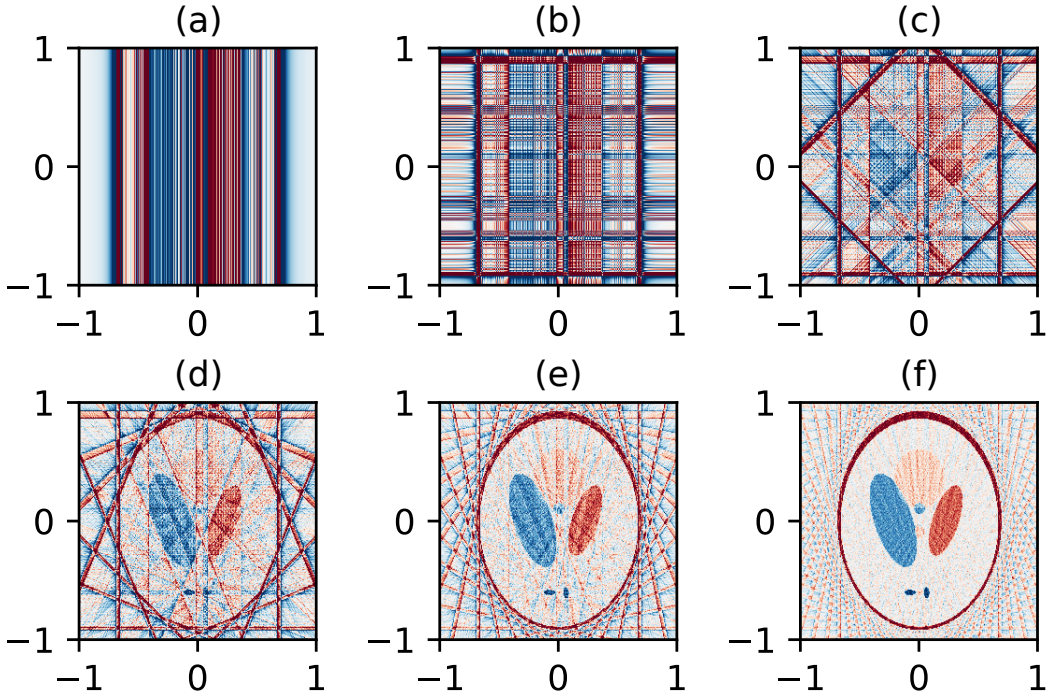


Figure 6.3: Example of the typical failure mode of filtered back-projection, in the case of few projections. These images were produced using a) 1, b) 2, c) 4, d) 8, e) 16 and f) 32 projections. The structural similarity index metrix (`SSIM`) was used to evaluate the quality of these reconstructions, and ranges from 5.64% in a) to 35.0% in e). Linear streaking artefacts are dominant in a) to c), with little to no resemblance to the original function. Even for the larger numbers of projections, checkerboarding is still observed in d) to f)

The characteristic behaviour of the FBP algorithm when provided with a decreasing number of projections is displayed in Fig. 6.3. Although the function being observed is quite simple, FBP produces results with streaking structures directed along the probe directions. Even for the case of a perfect circle (shown in Fig. 6.9) that is fully specified by even a single projection[5] (as all projections are identical) the filtered back-projection algorithm struggles, producing 'streak' artefacts along the directions of probing. In the limit of many projections these streaks cancel each other, leaving only the correctly-reconstructed circle.

Studying the derivation of filtered back-projection, it is easy to see why this is the case. The integral equation derived from the Fourier projection-slice theorem includes three integrals: over s, the projection transverse direction; and over k and ϕ

[5]Notably, an Abel transform inversion would reconstruct the circle perfectly.

the polar coordinates in the two-dimensional spatial frequency domain. By changing the order of integration the s and k integrals are trivially resolved analytically: the k-integral resolves to a Dirac delta depending on s, which is then used to evaluate the s-integral. The only integral that remains to be converted into a summation is over observation angle ϕ, and this conversion from an integral to a sum over angle produces all of FBP's streaking issues at low sampling rate. Note that the dependence on this angle arises in two places inside the integrand/summand: the angular dependence of the projections themselves (only known at the sampled angles) and the angular dependence of the relation between polar and Cartesian coordinates (known at any angle). The angular dependence of Cartesian coordinates is, however only sampled at the probe angles, leading to the streaks seen along the corresponding directions.

6.3.1 Hankel Reconstruction and Fourier Interpolation

The first improvement I propose aims to solve the problem of discretely-sampled angles. I will start this section by presenting a derivation of a reconstruction algorithm which employs integer-order Hankel transforms, a family integral transform measuring a function's overlap with Bessel functions of the corresponding order. I then discuss the implementation of the required Hankel transforms and suggest two different approaches to calculating it numerically. The first algorithm was proposed by Higgins and Munson [180] for the purpose of tomographic reconstruction and is more applicable to very small numbers of observations; the second is an algorithm I suggested in Ref. [14], which provides an alternative interpretation of this method as a Fourier transform-preserving interpolation of the projections to arbitrary angles.

The key observation is that of Eq. (6.4): performing a one-dimensional Fourier transform of a sinogram along its s-axis returns a polar-coordinate representation of the two-dimensional Fourier transform of the reconstructed function. Writing out the inverse Fourier transform in polar coordinates, with $k_x = -k\sin\phi, k_y =$

$k\cos\phi$ following Eq. (6.3) gives:

$$f(r,\theta) \approx \mathcal{F}_2^{-1}\left[\tilde{f}\right] = \frac{1}{2\pi}\iint e^{i\boldsymbol{k}\cdot\boldsymbol{x}}\tilde{f}(\boldsymbol{k})d^2\boldsymbol{k} \tag{6.9}$$

$$= \frac{1}{2\pi}\int_0^{2\pi}\int_0^{\infty} e^{ikr\sin(\theta-\phi)}\tilde{f}(k,\phi)kdkd\phi \tag{6.10}$$

The angular integral may be simplified if the angular dependence of $\tilde{f}$ is known. To achieve this, expand $\tilde{f}$ into a Fourier series in ϕ.

$$\tilde{f}(k,\phi) = \sum_{m=-M}^{M}\tilde{f}_m(k)e^{im\phi}, \tag{6.11}$$

where the bandlimit M is equal to the number of probe directions present in the data.

Then the inverse Fourier transform is:

$$f(r,\theta) \approx -\frac{1}{2\pi}\sum_{m=-M}^{M} e^{im\theta}\int_0^{\infty}\tilde{f}_m(k)k\int_0^{2\pi} e^{i(kr\sin\psi-m\psi)}d\psi dk \tag{6.12}$$

$$= \sum_{m=-M}^{M} f_m(r)e^{im\theta} \tag{6.13}$$

where I have employed the change of variable $\psi = \theta - \phi$ and noted that the 2π-periodicity in ψ of the integrand makes altering the limits of the angular integral unnecessary. A factor of $\exp(im\theta)$ arises from the change of variables, indicating that a Fourier mode in ϕ is transformed to a Fourier mode in θ of the same order. I have therefore also expressed the reconstructed function as a Fourier series in θ. The angular integral can be related to an integral definition of the Bessel functions of the first kind,

$$J_m(x) = \frac{1}{2\pi}\int_{-\pi}^{\pi} e^{i(x\sin\tau-m\tau)}d\tau. \tag{6.14}$$

The inverse Fourier transform is then expressible in terms of the Hankel transform:

$$f_m(r) = -\int_0^{\infty}\tilde{f}_m(k)J_m(kr)kdk \tag{6.15}$$

$$= -\mathcal{H}_m\left[\tilde{f}_m(k)\right] \tag{6.16}$$

This result—that angular Fourier modes of the observations transform to matching angular modes of the reconstructed function, and that the two modes are

related by a Hankel transform—is the main result of Higgins and Munson [173]. In this work the authors suggest the use fo a specialised algorithm that they had presented previously [180] to compute the integer-order Hankel transforms involved in this approach. As calculating a different order of transform is necessary for each of the $2M + 1$ angular modes present in the system, this procedure is most efficient when the number of observations is small.

An alternative approach, which I suggest in a paper published in Physical Review E in 2021 [30], looks back at the filtered back-projection algorithm discussed earlier. Both FBP and this Hankel approach are derived from the same expression: the inverse Fourier transform in polar coordinates. Therefore, in the continuous limit the two approaches should converge to the same result. This implies that the Hankel transform of order m can be approximated by filtered back-projection of a two-dimensional function with angular dependence given by $e^{im\phi}$. As the angular dependence is known, the function can be sampled with arbitrary angular precision to improve the performance of FBP. With this approach the reconstruction is

$$f(x, y) \approx \sum_{m=-M}^{M} \texttt{FBP}\left[F_m(s)e^{im\phi}\right]. \tag{6.17}$$

A further simplification can be achieved by noting that the FBP operator is linear and therefore commutes with the sum in the previous equation. Moving the summation inside FBP, I arrive at a new interpretation of this approach:

$$f(x, y) \approx \texttt{FBP}\left[\hat{F}(s, \phi)\right] \tag{6.18}$$

$$\hat{F}(s, \phi) = \sum_{m=-M}^{M} F_m(s)e^{im\phi} \tag{6.19}$$

Here, $\hat{F}$ is a function which by definition has the same Fourier series expansion as F, but is extended to continuous ϕ. The properties of Fourier series guarantee that this enhanced function agrees with the original observations at the discrete sample points (making it an interpolation of the original function), that it is explicitly 2π-periodic, and that it is smooth[6]. It is also the interpolation which minimises the average gradient.

[6] C^∞-smooth, in fact

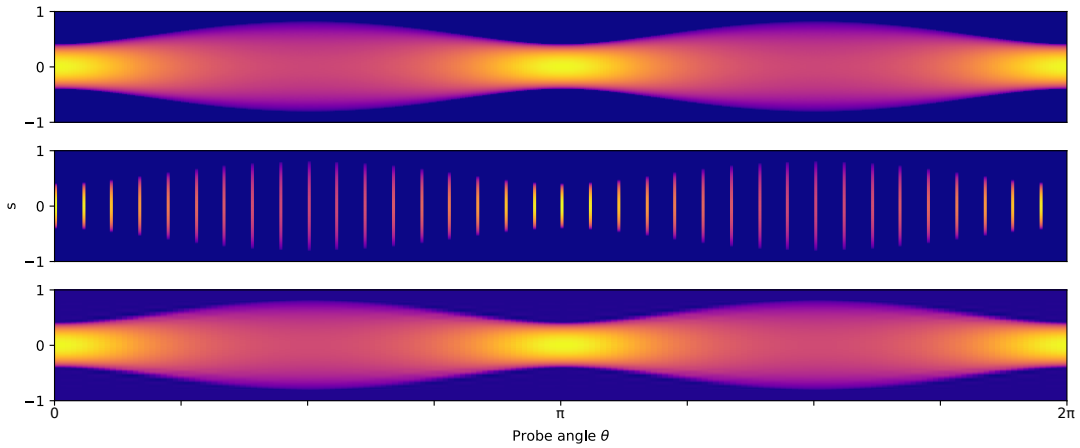


Figure 6.4: The effect of Fourier series interpolation on a sinogram of an elliptical function. a) a well-sampled sinogram, sampled at one degree increments. b) a poorly-sampled sinogram, with ten degree increments. c) A sinogram produced by interpolating b) using the Fourier series method.

Any number of samples can be drawn from $\hat{F}$ in order to obtain the desired performance from the FBP algorithm, and samples can be efficiently computed using a fast Fourier transform. First the `FFT` of the projected data is calculated, then the result is zero-padded to the desired number of samples and the inverse `FFT` is used to recover samples of $\hat{F}$. An example of this interpolation scheme is shown in Fig. 6.4.

6.3.2 Aspect Ratio Compensation

From the results shown in Fig. 6.5 it is clear that objects which are longer than they are wide (or vice versa) present a challenge to the Fourier interpolation scheme I developed in the previous section. Furthermore, this problem worsens as the aspect ratio of the object being observed is made more uneven. Conceptually it is easy to understand why this happens: the lowest-order assumption of the method is that objects are basically circular, and additional observations allow refinement away from circularity. The further from circular an object is—that is, the further from unity its aspect ratio—the more refinement is needed.

To try to mitigate this problem I now develop a sampling and pre-processing protocol which realises a non-uniform scaling of space as seen by the reconstruction

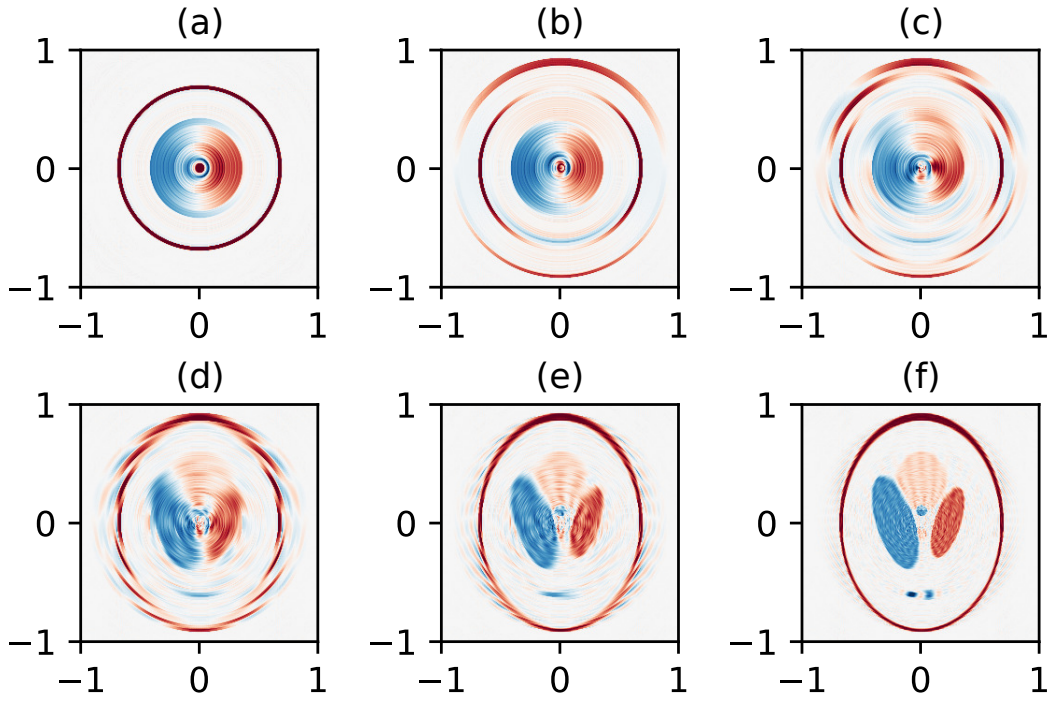


Figure 6.5: Demonstration of the improvement attainable by using Fourier interpolation techniques. Images were produced identically to Fig. 6.3, except that the set of projections was enhanced by interpolation to an angular frequency of 1024 view angles prior to reconstruction. Noise and streaking are significantly reduced, though performance is still quite poor for very small numbers of projections. Linear streaking has been replaced by circular smearing, reflecting this method's polar-coordinates formulation as opposed to the Cartesian formulation of standard FBP. The range of `SSIM` reconstruction quality is now from 74.6% ($N = 1$) in (a) to 85.9% ($N = 32$) in (f).

algorithm. In the following derivations I will use capital letters (X, Y, R, Θ) to refer to computational space and lower-case letters (x, y, r, θ) for physical space as before.

Suppose I wish to observe an elliptical object using tomography. The boundary of the object in physical space is given by:

$$x^2 + \left(\frac{y}{A}\right)^2 = 1 \tag{6.20}$$

Following the above discussion, I seek a coordinate transformation leading to the computational-space equation:

$$X^2 + Y^2 = 1 \tag{6.21}$$

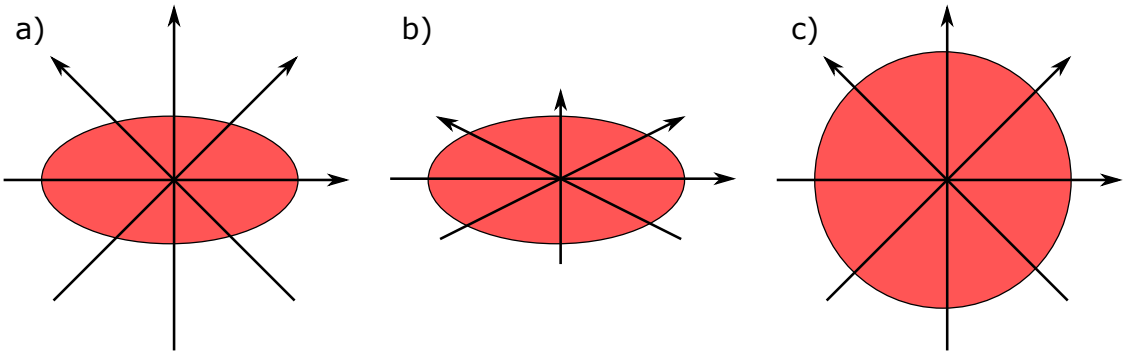


Figure 6.6: A schematic of the geometry of aspect ratio compensation. In a), an ellipse is sampled conventionally. In b), probe directions are chosen to match the non-uniform sampling expression Eq. (6.26). c) The situation in b) is transformed into computational space: the uniformity of sampling is recovered and the ellipse is transformed into a circle.

Finding such a transformation is trivial:

$$X = x \tag{6.22}$$

$$Y = \frac{y}{A} \tag{6.23}$$

Now, 'uniform' sampling is defined in computational coordinates, so

$$\Theta_m = \frac{m\pi}{M} \qquad (0 \leq m < M) \tag{6.24}$$

The angles are related to the Cartesian coordinates through the tangent function:

$$\tan\Theta = \frac{Y}{X} \tag{6.25}$$

$$\tan\theta = \frac{y}{x} = \frac{AY}{X} = A\tan\Theta \tag{6.26}$$

This implies that the physical probe directions should be arranged at angles with non-uniform spacing $\theta_m = \arctan(A\tan(m\pi/M))$. A schematic of this 'stretching' effect is shown in Fig. 6.6.

Now, projections need to be converted from physical to computational units. This will require a scaling of the s-axis to compensate for this transformation. The

physical transverse width of the ellipse, observed from some angle θ, is:

$$w(\theta) = \sqrt{A^2 \cos^2\theta + \sin^2\theta} \tag{6.27}$$

$$= A\cos\theta\sqrt{1 + \tan^2\theta/A^2} \tag{6.28}$$

$$= A\cos\theta\sqrt{1 + \tan^2\Theta} \tag{6.29}$$

$$= A\frac{\cos\theta}{\cos\Theta} \tag{6.30}$$

This width can also be expressed in terms of Θ alone:

$$w(\Theta) = \frac{A}{\sqrt{A^2\sin^2\Theta + \cos^2\Theta}} \tag{6.31}$$

Therefore the m^{th} projection is sampled with angle, transverse scaling and overall magnitude scaling (necessary to preserve normalisation) given by:

$$F(S, \Theta_m) \rightarrow F(s, \theta_m) = F\left(w(\Theta_m)S, \arctan(A\tan\Theta_m)\right) w(\Theta_m) \tag{6.32}$$

The effectiveness of this method at reconstructing elliptical functions is demonstrated in Fig. 6.7. Reconstruction is perfect where the correct aspect ratio is used, and intermediate aspect ratios (between no compensation and the correct ratio) also give some improvement. This indicates that this technique may also improve reconstruction quality for elongated functions that have less well-defined aspect ratios, or whose exact aspect ratio is not known *a priori* as is likely to be the case in laser-driven plasma structures.

6.4 Results

I will now demonstrate the result of applying tomographic reconstruction to a more interesting function than the relatively simple one considered so far. I performed a particle-in-cell simulation (similar to the one discussed in §5.4.2) in which an intense ($500\,\text{J}, 500\,\text{fs}, a_0 = 21$) laser pulse propagated into plasma with a shallow density gradient of $500\,\mu\text{m}$. The simulation was stopped once all of the laser energy had been absorbed by the plasma, and then the residual magnetic field associated with the laser channel was extracted. I used this magnetic field structure (shown in Fig. 6.8) to test the effectiveness of the tomographic methods presented above.

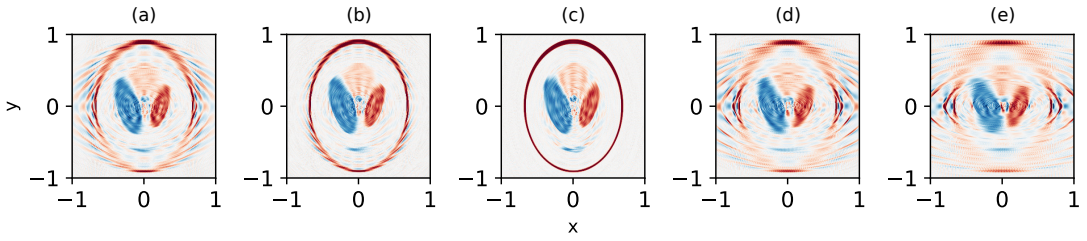


Figure 6.7: The effect of aspect ratio compensation on reconstruction quality. All reconstructions were carried out using 15 projections, and expected aspect ratios which are integer powers of the correct 3:4 ratio: a) 4:3 (SSIM: 47.5%), b) 1:1 (SSIM: 69.0%), c) 3:4 (SSIM: 87.2%), d) 9:16 (SSIM: 39.5%), and e) 27:64 (SSIM: 35.3%). b) therefore shows no compensation; c) shows the correct level of compensation and resultantly the highest quality of reconstruction; a) shows the correct ratio but in the wrong direction, worsening the quality of the reconstruction; d) and e) 'squeeze' the function due to excessive compensation. It is clear that compensation for aspect ratio can improve reconstruction quality, if applied appropriately.

In addition to this more realistic field, I also show full results for reconstructing a circle (Fig. 6.9), an ellipse (Fig. 6.10) and the modified Shepp-Logan phantom (Fig. 6.11) for completeness. These figures show a direct comparison between the naïve FBP algorithm, each of my proposed enhancments individually, and both in combination, at a range of probe direction counts. The first case, the circle, shows how well Fourier interpolation can improve results for highly symmetric functions; the second, the ellipse, shows that aspect ratio compensation can transform less symmetric functions into a form more amenable to the Fourier-enhanced reconstruction technique; and the third shows how both methods improve the result even for non-ideal functions.

6.5 Summary

In this chapter I proposed a new approach to proton radiography. By employing tomographic methods, fully three-dimensional information can be recovered from proton radiographs. However, the expected shortage of available lines-of-sight severely limits the performance of common algorithms such as filtered back-projection (FBP) and algebraic reconstruction techniques (ARTs).

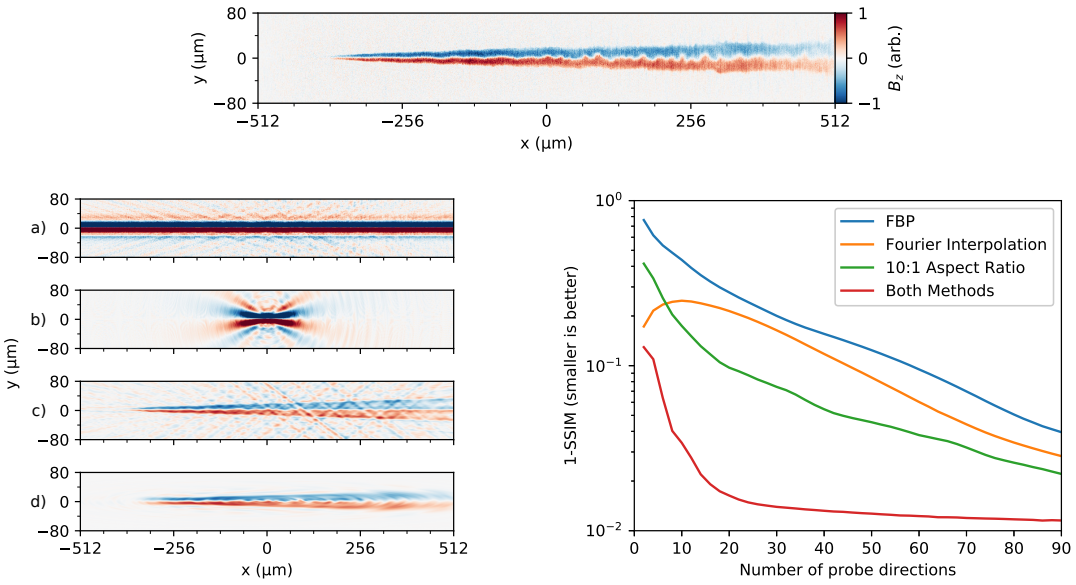


Figure 6.8: Above: Z-component of magnetic field extracted from a particle-in-cell simulation of a laser-plasma channelling interaction, of kilotesla order of magnitude. The field is shown in two dimensions, averaged over the z-axis of the simulation to improve legibility. All of the methods presented here are however applicable to three-dimensional data sets as well as two-dimensional ones. Below left: Reconstructions produced using a sampling rate $N = 11$ views. a) presents the results of applying naïve filtered back-projection, b) includes aspect ratio-compensation, c) includes Fourier-series interpolation and d) includes both. Aspect ratio compensation uses a ratio of 10:1. Below right: quantitative image reconstruction quality [175], plotted as $1 - \texttt{SSIM}$ such that smaller values are better, with 0 representing perfect reconstruction. While each technique struggles with this field when used individually, the composition of both techniques produces results of good quality even for very sparse sampling, equivalent to FBP using numbers of views an order of magnitude larger: 90% SSIM is achieved for five view angles when both enhancements are applied; without them this quality level requires fifty-eight. All reconstructions are plotted using the same colour scaling as the original field.

To improve this situation, I have suggested a Fourier-domain method for interpolation of observations which can be applied as a pre-processing step before a reconstruction is carried out. I present rigorous theoretical justification for the particular interpolation I propose, which effectively turns FBP into an algorithm for computation of integer-order Hankel transforms. I have shown using some simple example functions that the interpolation transforms the linear, streaking artefacts seen in sparsely-sampled FBP reconstructions into circular streaks, which are in many cases far less noticeable and problematic. Because of the interpretation of this as a pre-processing step, it can also be applied to algorithms other than

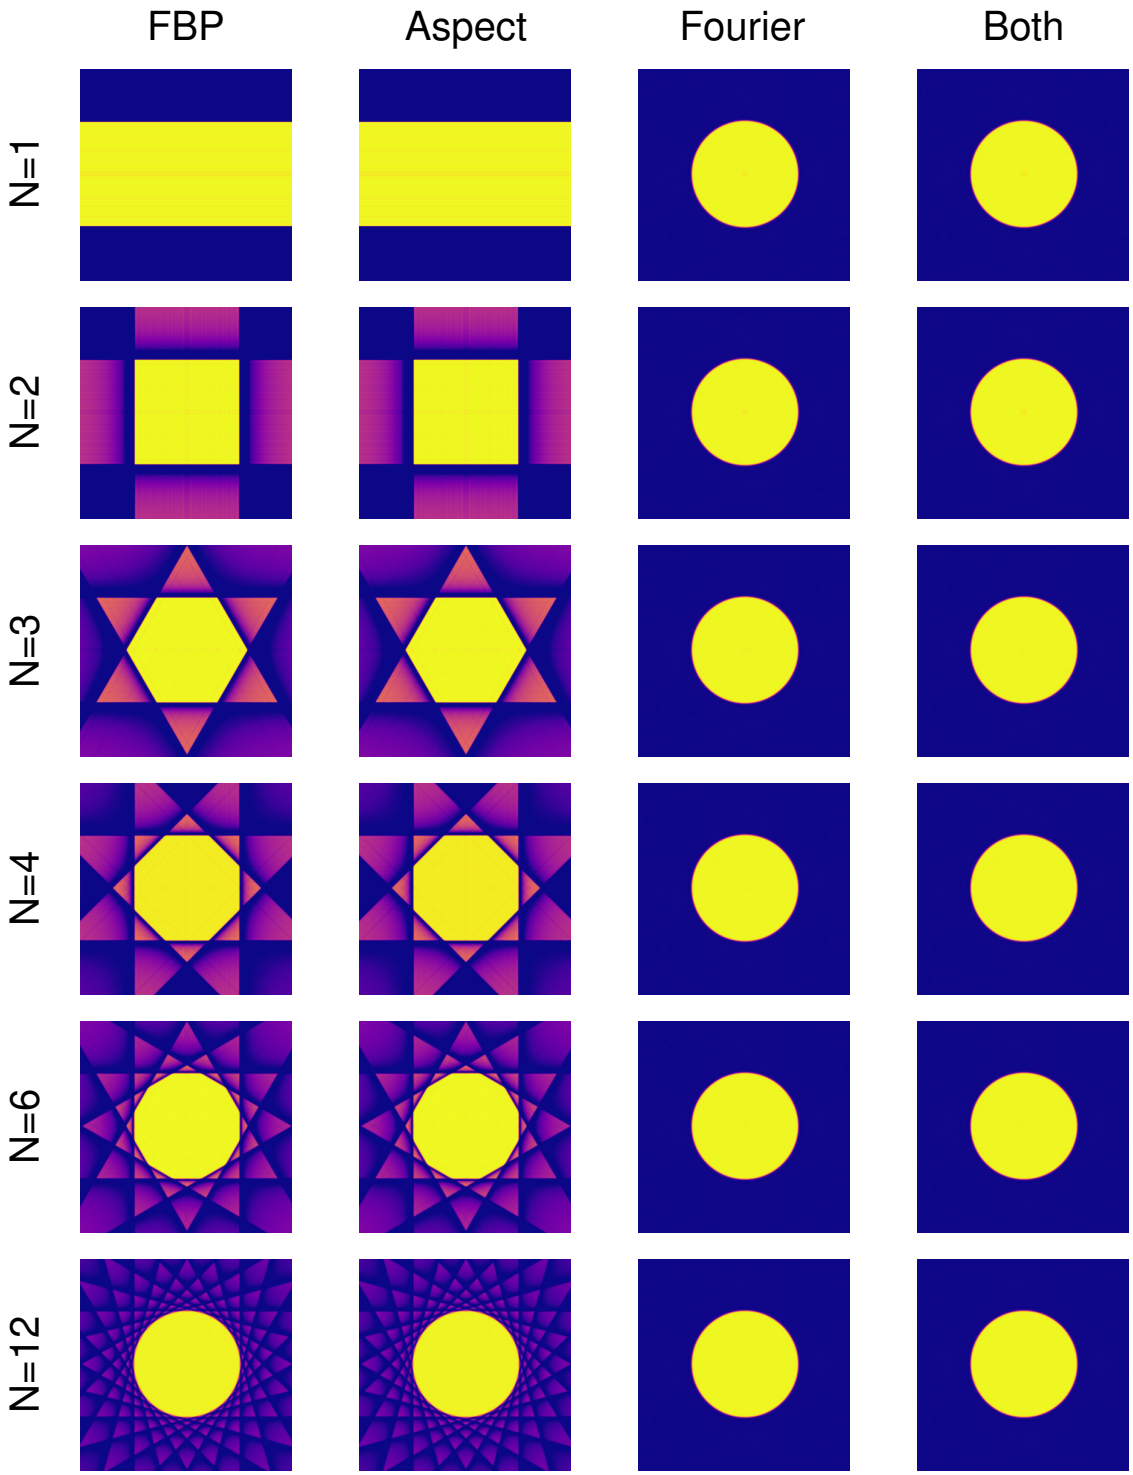


Figure 6.9: Circle reconstructions. Because the aspect ratio of a circle is unity, aspect ratio compensation has no effect. Due to the symmetry of a circle the Fourier series interpolation produces perfect reconstruction for any number of projections.

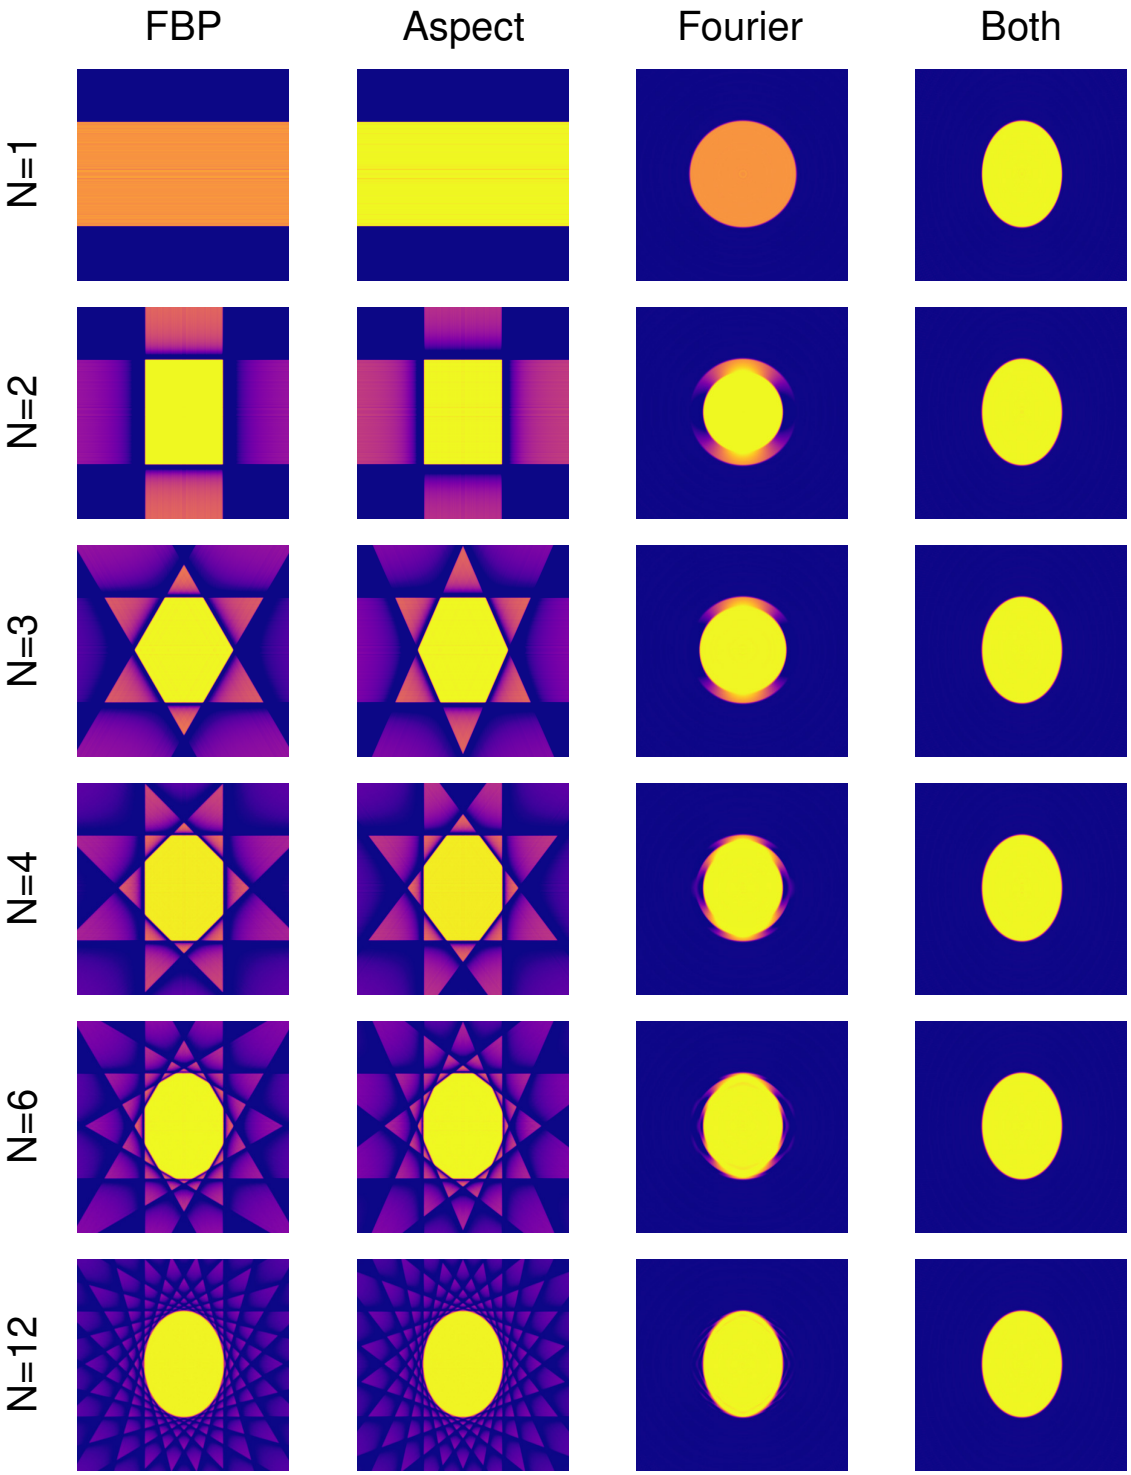


Figure 6.10: Ellipse reconstructions. Angular fringing is seen in the Fourier-interpolated result for small numbers of probe directions, and is eliminated by use of the correct aspect ratio compensation.

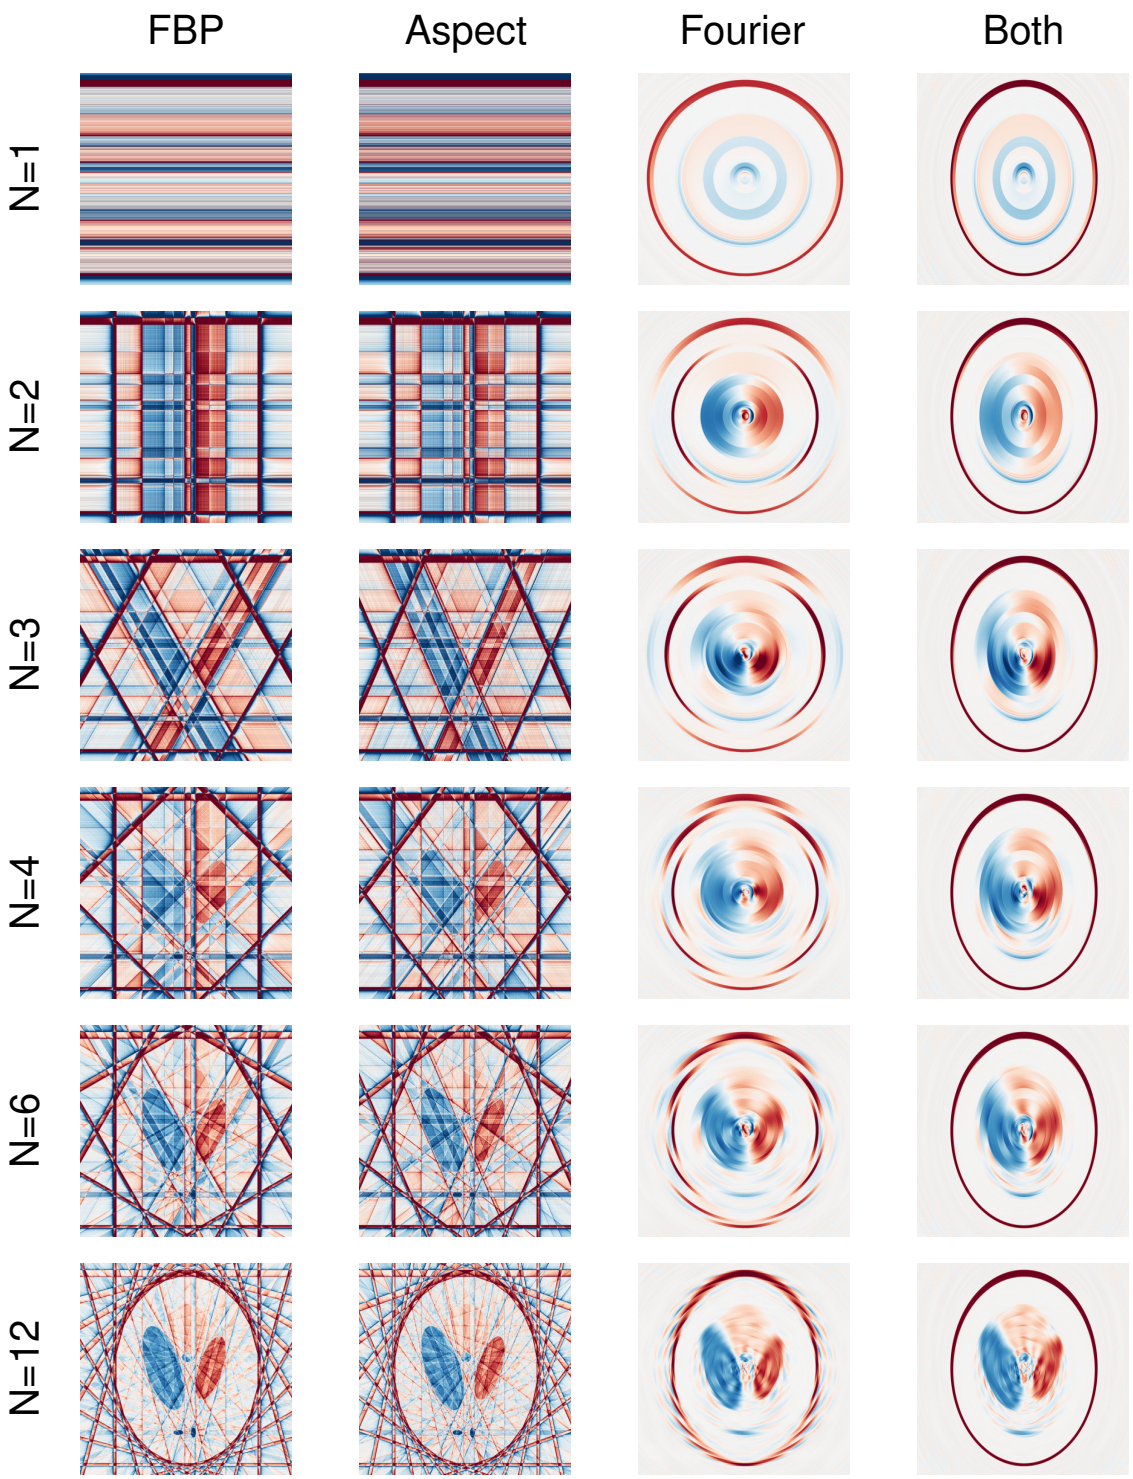


Figure 6.11: Modified Shepp-Logan reconstructions. This less symmetric function nonetheless benefits significantly from both Fourier interpolation. Without the interpolation the benefit of aspect ratio compensation is not clear, but the improvement can be seen in interpolated reconstructions.

FBP, such as the aforementioned ARTs, and in that context could enhance a regularised reconstruction.

Further, I used the Fourier series interpolation method to explore the tomography of elliptical object, or more generally objects with non-uniform aspect ratio. I have shown how such objects violate the key assumption of circularity inherent to the Fourier method and decrease the efficiency of the Fourier series representation as a result. To remedy this problem I propose a new sampling scheme, which uses controlled non-uniformity in the probe directions to make the object under inspection appear to the reconstruction algorithm to be more circular than it actually is. The result is then returned to its physical aspect ratio by a simple non-uniform scaling.

The work presented in this chapter was published in Physical Review E in October 2021 [30]

7 Summary and Future Work

Contents

7.1 Summary of Work 145
7.2 Future Work 147

7.1 Summary of Work

In this thesis I have presented several pieces of work - some theoretical, some computational and some experimental. This work covers a diverse range of topics related to laser-plasma interactions and the way they are modelled.

In Chapter 5, I presented a new solver I have implemented in an existing community proton radiography analysis code. This new solver offers improved accuracy and stability compared to the existing approach, as well as speed-up by a factor of the order of several thousand (in terms of iterations) or about a hundred (in terms of wall-clock time). The improved stability of the new solver was essential for the experimental analysis I present in the same chapter—the previous solver encountered numerical instability issues when processing some of the radiographs—and represents a significantly improved capability for the community.

The first experimental results I present are in Chapter 4. In §4.3 I report an experiment which I carried out at the Vulcan laser facility along with other members of Prof. Norreys' research group and external collaborators. This experiment was the first attempt to drive collisionless shock acceleration—a novel mechanism for ion acceleration—in plasma formed from a solid foam. It was limited by the sizes and densities that could be achieved in free-standing foam targets, and simulation modelling conducted after-the-fact suggested that the observed bunches of protons with narrow energy spread may not have been the result of collisionless shock acceleration. As well as the experiment, I presented significant advances to the theoretical basis of collisionless shock acceleration, which was first attempted by Cairns et al [115, 116] and previously expanded on in the DPhil thesis of Savin [29].

More experimental analysis is shown in the latter sections of Chapter 5. An experiment conducted by the Norreys group and Atomic Weapons Establishment collaborators at AWE's Orion laser facility produced proton radiography images as a major diagnostic. Using the aforementioned new algorithm for proton radiography I conducted the first quantitative analysis of these radiographs, looking for evidence that the laser pulse had channelled deep into the pre-formed plasma. While simulation studies suggested this may have been the case, the proton radiographs of channelling shots contained enough confounding structure to render the analysis inconclusive. Instead, I reconstructed path-integrated magnetic field densities for long-pulse only shots and found the results to be consistent with the magnetic field strength predicted by magnetohydrodynamics simulations. This work was published in Philosophical Transactions of the Royal Society A in 2020 [12].

Next, in Chapter 6 I present the results of my work investigating the application of tomographic techniques to data obtained by proton radiography. I motivate, derive mathematically and provide enlightening interpretation of a new approach to pre-processing tomographic data based on the Hankel transform. On the way I also derive a new, generalised form of the Fourier-Hankel-Abel cycle of integral transforms (of which the classical FHA cycle is the zero-order special case). Based on shortcomings of the reconstructions obtained using this new approach, I then

propose that a 'squeezing' transformation between physical space and computational space may improve its performance, elaborate on how such a transformation can be implemented and demonstrate how significantly results can be improved.

Finally, in Appendix A and Appendix B I have given an overview of software I have written to model specific parts of physical systems. In Appendix A I present a specialised rendering system designed to produce high-quality images for a simulated proton radiography diagnostic. These images are the noise-free limiting case of popular stochastic approaches to the same problem. In Appendix B I present an optical propagation code which uses a novel coordinate transformation of the paraxial Helmholtz equation to efficiently model the evolution of focusing light, coupled to a simple model I developed to describe the aberrations induced by a misaligned parabolic mirror. These are integrated as components in the python package `pyrabola`, which provides a user-friendly interface for modelling the propagation of a laser beam through simple optical systems consisting of several planar beam transport mirrors, a focusing mirror and a camera.

7.2 Future Work

The work presented in this DPhil thesis also leaves room for future development. Most obviously, the collisionless shock acceleration experiment presented in Chapter 4 was limited by technical difficulties during the run, as well as limitations in fabrication capabilities for foam targets. A refinement of that experiment, using a facility with well-tested capacity to provide third-harmonic long pulses and a more complete set of diagnostics (to measure, for example, the plasma electron temperature), is expected to provide greatly improved—and better-characterised—results, especially as more mature foam production capabilities of target fabrication groups enable them to provide lower-density foams. With these improvements in place, theoretical modelling (as well as experimental results from Pak et al. [117]) predicts that significantly higher ion energies would be produced by collisionless shock acceleration. This acceleration mechanism has also seen recent interest as a driver for ion-based fast ignition approaches, such as in a particle-in-cell simulation

study published by Boella et al. [181], and looking further into the future this would be an interesting avenue for experimental research.

On the topic of fast approaches to inertial confinement fusion, the HKT mechanism of whole-beam self-focusing I have presented in Chapter 5 shows great promise in both fast ignition experiments and the 'auxiliary heating' approach, which was proposed by Norreys et al. [182] based on theory by Ratan et al. [183], then expanded upon in ongoing simulation studies and upcoming experimental work by Paddock et al. [31]. The auxiliary heating scheme looks somewhat like a hybrid between conventional hot-spot approaches to fusion and fast ignition: the compression phase proceeds similarly to that of a conventional implosion (though possibly using slightly unusual capsule designs, such as wetted-foam constructions [31]), resulting in an isobaric central hot-spot at stagnation. At this point, a fast heating is applied to the capsule, for example using the kinetic crossing-beam electron instability discovered by Ratan et al. [183] to cascade energy from electron beams into the hot-spot ions, via plasma electrons. The whole-beam self-focusing mechanism shows promise for channel formation (if used to evacuate plasma ahead of a subsequent, more intense pulse) or itself allowing a single-beam solution to penetrate deep into fusion plasma and deposit energy much nearer to the core than would otherwise be possible.

The experiment whose results I showed in Chapter 5 was performed in planar geometry, as a proof of concept. Obviously there is much work to be done if the results are to be generalised to irradiation of the spherical targets found in laser fusion experiments. It is obvious that light propagating into a spherical density gradient is far more unstable to refraction than the equivalent planar situation; a clear and interesting direction for further work is to investigate the extent to which whole-beam self-focusing can act to counteract this instability.

The experimental implementation of the methods proposed in Chapter 6 is likely to be difficult. Either a tedious process of realigning proton sources and detectors after every shot must be undertaken, or to avoid the detrimental effect of jitter multiple proton radiorgraphy diagnostics must be fielded concurrently. This

poses a particular challenge to facilities which employ short pulse-generated proton beams, as most laser facilities have no more than two short pulses. In order to address this difficulty I have been investigating novel target geometries which could produce multiple proton beams from a single short, intense laser pulse. While this work was not mature enough for inclusion in this thesis, preliminary work is published on the arχiv [13], and I introduced the concept in the peer-reviewed article based on Chapter 6 [30].

In conclusion, I have presented three major topics of research related to laser-accelerated proton beams and their application to probing magnetic fields in laser-plasma experiments.

Appendices

A

Producing Perturbed Density Maps

Contents

A.1 The Problem 154
A.2 Stochastic Methods 155
A.3 Deterministic Methods 157
A.3.1 Rectangular cells 157
A.3.2 Triangular Cells 159
A.3.3 Results 161
A.4 Conclusions 164

Some of the work involved in this thesis has been focused on obtaining perturbation maps which produce a certain density distribution at the image plane, but the 'forward' problem corresponding to this difficult 'inverse' problem is also of interest. It is useful to be able to produce density maps from a given perturbation field affecting a given initial density distribution not only to validate the numerical techniques used to solve the inverse problem (detailed in §5.2) but also to produce fully synthetic radiographs based on magnetic fields produced by simulations. The latter of these use cases can be helpful in informing experimental design, for example by enabling experimenters to predict ahead of time which combinations of plasma conditions and proton beam parameters will produce well-behaved radiographs which do not exhibit the formation of caustics.

A.1 The Problem

The problem at hand is as described in §5.2. A two-dimensional population of particles with known initial density $n(\boldsymbol{x})$ is given some velocity perturbation in the transverse direction $\Delta\boldsymbol{v}(\boldsymbol{x})$. After some time passes, particles have moved to new positions $\boldsymbol{x} + \Delta\boldsymbol{v}\Delta t$. In regions where the velocity perturbation causes particles to converge, their density increases accordingly, and in regions where they diverge their density is reduced. Therefore, particles initially located in the neighborhood of a point $\boldsymbol{x}_0 = (x_0, y_0)$ contribute to the final density at a location

$$\boldsymbol{x}' = \boldsymbol{x}_0 + \Delta\boldsymbol{v}(\boldsymbol{x}_0)\Delta t. \tag{A.1}$$

The magnitude of their contribution is

$$n(\boldsymbol{x}', \Delta t) = n(\boldsymbol{x}_0, 0)\left|\frac{\partial(x', y')}{\partial(x_0, y_0)}\right|^{-1}, \tag{A.2}$$

or, defining $n_0(\boldsymbol{x}) = n(\boldsymbol{x}, 0)$, $n'(\boldsymbol{x}) = n(\boldsymbol{x}, \Delta t)$,

$$n'(\boldsymbol{x}') = n_0(\boldsymbol{x}_0)\left|\frac{\partial(x', y')}{\partial(x_0, y_0)}\right|^{-1}, \tag{A.3}$$

with the velocity Jacobian given by

$$\frac{\partial(x', y')}{\partial(x_0, y_0)} = \left(1 + \Delta t\frac{\partial\Delta v_x}{\partial x_0}\right)\left(1 + \Delta t\frac{\partial\Delta v_y}{\partial y_0}\right) - \Delta t^2\frac{\partial\Delta v_y}{\partial x_0}\frac{\partial\Delta v_x}{\partial y_0}. \tag{A.4}$$

In the general case, any given location (x, y) may contain contributions from multiple initial locations (x_0, y_0). In this case there will also be regions where the Jacobian vanishes, leading to locally divergent density. However, the Jacobian in Eq. (A.3) implies that the integral over each distribution is the same:

$$\int n_0(\boldsymbol{x}_0)dx_0dy_0 = \int n'(\boldsymbol{x}')\left|\frac{\partial(x', y')}{\partial(x_0, y_0)}\right|dx_0dy_0 \tag{A.5}$$

$$= \int n'(\boldsymbol{x}')dx'dy' \tag{A.6}$$

As the input velocity perturbation field and the output density image will be discretised, it makes sense to average the end result over each pixel in the final

image to suppress divergences associated with caustics. For a pixel p centred at $\boldsymbol{x}_p$ and square in shape with dimensions $\Delta x, \Delta y$, this average takes the following form:

$$\bar{n}_p = \frac{1}{\Delta x \Delta y} \iint_p n'(\boldsymbol{x}') dx' dy' \tag{A.7}$$

$$= \frac{1}{\Delta x \Delta y} \int_{x_p - \Delta x/2}^{x_p + \Delta x/2} \int_{y_p - \Delta y/2}^{y_p + \Delta y/2} n'(\boldsymbol{x}') dx' dy' \tag{A.8}$$

This can also be expressed in terms of the pre-transformation co-ordinates $\boldsymbol{x}_0$, making use of the Jacobian in Eq. (A.5) to enable the change of variables:

$$\bar{n}_p = \frac{1}{\Delta x \Delta y} \iint_p n_0(\boldsymbol{x}_0) dx_0 dy_0 \tag{A.9}$$

$$= \frac{1}{\Delta x \Delta y} \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \Pi^2 \left(\frac{\boldsymbol{x}'(\boldsymbol{x}_0) - \boldsymbol{x}_p}{\Delta \boldsymbol{x}} \right) n_0(\boldsymbol{x}_0) dx_0 dy_0 \tag{A.10}$$

Here, I have introduced the rectangle function $\Pi(u)$, which is one for argument $|u| < \frac{1}{2}$ and zero outside of this range. The notation $\Pi^2(\boldsymbol{x})$ represents the product of the function applied to each component of the vector argument, i.e.

$$\Pi^2(\boldsymbol{x}) = \Pi(x)\Pi(y), \tag{A.11}$$

which is one on a unit square (centred on the origin) and zero outside of it. The expression $\Pi^2 \left(\frac{\boldsymbol{x}' - \boldsymbol{x}_p}{\Delta \boldsymbol{x}} \right)$ is therefore one within the pixel centred on $\boldsymbol{x}_p$ and zero elsewhere, and so encodes the domain of the integral which may be difficult to express directly.

A.2 Stochastic Methods

A simple approach, which has seen use in existing literature (for example, by Graziani et al. [152]), is to randomly generate initial particle positions within the original density distribution and then apply binning to the corresponding final positions to produce a histogram (weighted by the input density at the sampled initial location). This method is easy to implement and gives a good result when the sampling rate is high, but suffers from noise if not enough samples can be drawn to ensure many samples contribute to each pixel of the output image. This is especially problematic in images with large density fluctuations which can result in

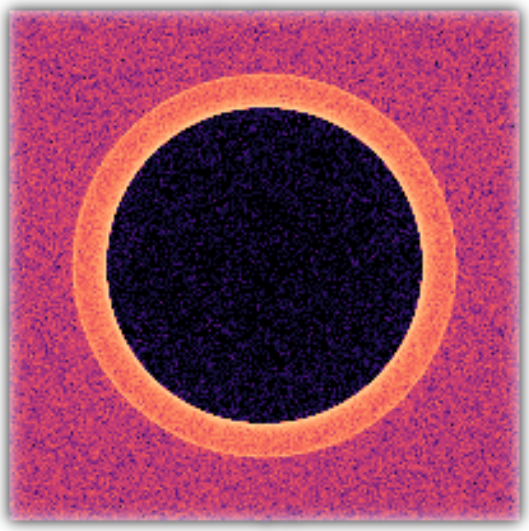


Figure A.1: Stochastic proton radiography image. This image was computed using a spatial resolution of 200×200 pixels, and 10^8 pseudo-particles (mean of 25 per pixel). Poisson noise is notable, especially in the centre of the image where density is low.

very low-density regions in the output image. Because these regions of low density receive only a small proportion of the simulated protons, they exhibit larger amounts of Poisson noise. This method can be understood as replacing n_0 with an initial density distribution n_1 which is a sum over K Dirac delta density distributions (or 'pseudo-particles'), located at random positions (i.e. a set of random samples of n_0):

$$n_1(\boldsymbol{x}_0) = \frac{A_{\text{tot}}}{K} \sum_{k=0}^{K-1} n_0(\boldsymbol{x}_k) \delta^2(\boldsymbol{x}_0 - \boldsymbol{x}_k), \tag{A.12}$$

where the total area A_{tot} over which samples are produced is used to normalise the distribution and ensure the results scale correctly.

When this initial density profile is used to evaluate the integral Eq. (A.10), the sum commutes outside of the integral and the integral can be trivially evaluated for each pseudo-particle:

$$\bar{n}_p = \frac{A_{\text{tot}}}{K \Delta x \Delta y} \sum_{k=0}^{K} n_0(\boldsymbol{x}_k) \Pi^2 \left(\frac{\boldsymbol{x}'(\boldsymbol{x}_k) - \boldsymbol{x}_p}{\Delta \boldsymbol{x}} \right) \tag{A.13}$$

$$= \frac{A_{\text{tot}}}{K A_p} \sum_{k=0}^{K} \begin{cases} n_0(\boldsymbol{x}_k) & \text{pseudo-particle } k \text{ lands in pixel } p \\ 0 & \text{otherwise} \end{cases} \tag{A.14}$$

where $A_p = \Delta x \Delta y$ is the cell area. This is clearly a binning operation, which produces a histogram. Note that if the 'final position' function $\boldsymbol{x}'(\boldsymbol{x}_0)$ and/or

the initial density $n_0(\boldsymbol{x}_0)$ are only known at specific sampled locations, some two-dimensional interpolation scheme will be necessary to evaluate this expression at the random pseudo-particle locations.

A refinement to the stochastic approach is to weight the generation of samples such that cells with a large Jacobian (Eq. (A.4)) receive more samples, and adjust the binned weight contributed by samples by the inverse of their Jacobian. This measure ensures uniform sampling of space in the final image, distributing Poisson noise more uniformly. However, in the presence of caustics this method is problematic as it could generate samples for which the Jacobian vanishes, causing an infinite contribution to the output image.

An example of the performance of stochastic methods such as these is shown in Fig. A.1.

A.3 Deterministic Methods

Another approach is to compute the integral in equation Eq. (A.10) directly. If the the 'final position' function $\boldsymbol{x}'(\boldsymbol{x}_0)$ is sampled at the vertices of a grid of cells, then this task can be undertaken for each cell individually. The contribution of density originating in a cell c to the value of a single pixel p is then

$$\bar{n}_p(c) = \frac{n_c}{A_p} \iint_c \Pi^2 \left(\frac{\boldsymbol{x}'(\boldsymbol{x}_0) - \boldsymbol{x}_p}{\Delta \boldsymbol{x}} \right) dx_0 dy_0, \tag{A.15}$$

where I have assumed that within a single cell, $n_0(\boldsymbol{x}_0) = n_c$.

Unlike the stochastic method discussed above, implementation of this method is not transparent to the method used to interpolate the final position from the defined cell vertices to locations within a cell. This is because evaluating Eq. (A.15) efficiently requires knowing the inverse function to the interpolation used, so that the limits of integration can be chosen appropriately.

A.3.1 Rectangular cells

As a first example, consider dividing the (x_0, y_0) plane into rectangular cells of size $(\Delta x_c, \Delta y_c)$ whose product is the cell area A_c. Parametrising each cell with

the coordinates $(u, v) \in [0, 1]^2$ and (for notational brevity) defining $\lambda = \frac{x'-x_p}{\Delta x}$, $\mu = \frac{y'-y_p}{\Delta y}$, the above integral becomes

$$\bar{n}_p(c) = n_c \frac{A_c}{A_p} \int_0^1 \int_0^1 \Pi(\lambda(u,v))\Pi(\mu(u,v)) dv du \tag{A.16}$$

If the region of (u, v)-space covered by the pixel is simply-connected, this can be simplified using Green's theorem into a contour integral over the edges of the pixel. For example, in the case of bilinear interpolation,

$$\begin{aligned} \lambda(u,v) = (1-v)((1-u)\lambda(0,0) + u\lambda(1,0)) \\ +v((1-u)\lambda(0,1) + u\lambda(1,1)) \end{aligned} \tag{A.17}$$

$$\begin{aligned} \mu(u,v) = (1-v)((1-u)\mu(0,0) + u\mu(1,0)) \\ +v((1-u)\mu(0,1) + u\mu(1,1)) \end{aligned} \tag{A.18}$$

For fixed u, v can be found in terms of either λ or μ:

$$v(u;\lambda) = \frac{\lambda - \lambda(u,0)}{\lambda(u,1) - \lambda(u,0)} \tag{A.19}$$

$$v(u;\mu) = \frac{\mu - \mu(u,0)}{\mu(u,1) - \mu(u,0)} \tag{A.20}$$

Green's theorem can then be used to write Eq. (A.16) as the integral of $v du$ over the boundary of the pixel:

$$\begin{aligned} \frac{\bar{n}_p(c)}{n_c}\frac{A_p}{A_c} = &\int_{u^{--}}^{u^{+-}} v\left(u;\mu=-\frac{1}{2}\right) du \\ &+ \int_{u^{+-}}^{u^{++}} v\left(u;\lambda=-\frac{1}{2}\right) du \\ &+ \int_{u^{++}}^{u^{-+}} v\left(u;\mu=\frac{1}{2}\right) du \\ &+ \int_{u^{-+}}^{u^{--}} v\left(u;\lambda=\frac{1}{2}\right) du \end{aligned} \tag{A.21}$$

where $u^{\pm\pm}, v^{\pm\pm}$ are defined such that $\lambda(u^{\pm *}, v^{\pm *}) = \pm\frac{1}{2}$ and $\mu(u^{*\pm}, v^{*\pm}) = \pm\frac{1}{2}$. Implicit here is the assumption that each of the $u^{\pm\pm}$ is single-valued, which is not necessarily true in the presence of caustics. For pixels that cross caustics one or more of the pixel corners could correspond to two values of u (corresponding to

the roots of the quadratic obtained by inverting the bilinear interpolation map); in this case the contour integral becomes more complicated. Even more complication can arise with use of higher-order (e.g. bicubic) interpolation. Pixels which cross the edges of the cell being considered will also need special handling, to limit the (u, v) values included in the integral to the appropriate $[0, 1]$ interval.

A.3.2 Triangular Cells

Handling all of the edge-cases noted above unfortunately limits the efficiency of this method, so I now propose a method based on a different discretisation of the (x_0, y_0) plane. Using triangular instead of quadrilateral cells facilitates use of linear (rather than *bi*linear, or higher-order) interpolation. This has a number of advantages, such as the fact that the inverse mapping is also linear and so cannot be multi-valued. Because the mapping between (u, v) coordinates of a given cell and the (λ, μ) coordinates of a pixel is linear (more strictly, affine), the image of a (rectangular) pixel under the inverse mapping is always a parallelogram.

These facts simplify the evaluation of the equivalent of Eq. (A.21) greatly: where the triangle is parametrised using barycentric coordinates $(u, v, 1-u-v)$, it is well known that integrals over the triangle (of area A_c) can be expressed as

$$\int_{\text{triangle}} f(\boldsymbol{x}) d^2\boldsymbol{x} = 2A_c \int_0^1 \int_0^{1-v} f(u, v) du dv. \tag{A.22}$$

The coordinates $u^{\pm\pm}$ of the pixel corners are straightforward to express in terms of the coordinates(λ_i, μ_i) of each of the triangle's three vertices (where i runs from 0 to 2):

$$u^{\pm\pm} = \frac{\Delta\lambda^{\pm}\Delta\mu_1 - \Delta\mu^{\pm}\Delta\lambda_1}{\Delta\lambda_0\Delta\mu_1 - \Delta\mu_0\Delta\lambda_1}, \tag{A.23}$$

where the first (second) $\pm$ on the left side relates to the first (second) $\pm$ on the right, and where:

$$\begin{pmatrix}\Delta\lambda_0 \\ \Delta\mu_0\end{pmatrix} = \begin{pmatrix}\lambda_0 - \lambda_2 \\ \mu_0 - \mu_2\end{pmatrix} \tag{A.24}$$

$$\begin{pmatrix}\Delta\lambda_1 \\ \Delta\mu_1\end{pmatrix} = \begin{pmatrix}\lambda_1 - \lambda_2 \\ \mu_1 - \mu_2\end{pmatrix} \tag{A.25}$$

$$\begin{pmatrix}\Delta\lambda^{\pm} \\ \Delta\mu^{\pm}\end{pmatrix} = \begin{pmatrix}\pm\frac{1}{2} - \lambda_2 \\ \pm\frac{1}{2} - \mu_2\end{pmatrix} \tag{A.26}$$

Finally, the v-values in the integral are given by:

$$v(u;\lambda) = \frac{(\lambda - \lambda_2) - u\Delta\lambda_0}{\Delta\lambda_1} \tag{A.27}$$

$$v(u;\mu) = \frac{(\mu - \mu_2) - u\Delta\mu_0}{\Delta\mu_1} \tag{A.28}$$

With these definitions in place it is straightforward (though tedious) to show that the sum of the four integrals in Eq. (A.21) evaluates to give:

$$\frac{\bar{n}_p(c)}{n_c}\frac{A_p}{2A_c} = \frac{1}{\Delta\lambda_0\Delta\mu_1 - \Delta\mu_0\Delta\lambda_1}, \tag{A.29}$$

where the factor 2 on the left-hand side originates from Eq. (A.22). The expression

$$\Delta\lambda_0\Delta\mu_1 - \Delta\mu_0\Delta\lambda_1 \tag{A.30}$$

is in fact twice the post-transform area of the triangle in (λ, μ) coordinates (as it's the wedge product of two of the triangle's sides in these coordinates), so is equal to $2\frac{A'_c}{A_p}$. Therefore, the pixel density is:

$$\bar{n}_p(c) = n_c\frac{2A_c}{A_p}\frac{A_p}{2A'_c} = \frac{A_c}{A'_c}n_c. \tag{A.31}$$

This result should be obvious (or at least expected) with the application of some geometric insight: the density deposited by a single triangular cell in any pixel it completely covers is equal to the original density in that triangular cell, scaled by the inverse to the factor by which its area changes under the coordinate transformation. If a cell is stretched its density decreases and if it is compressed its density increases, to keep its total 'mass' constant. This constant density property associated with triangular cells makes the implementation of this method much simpler, as the only special case that must be handled is the possibility of pixels not being completely covered. Even in those cases, the fraction of density received by the cell is simple to calculate.

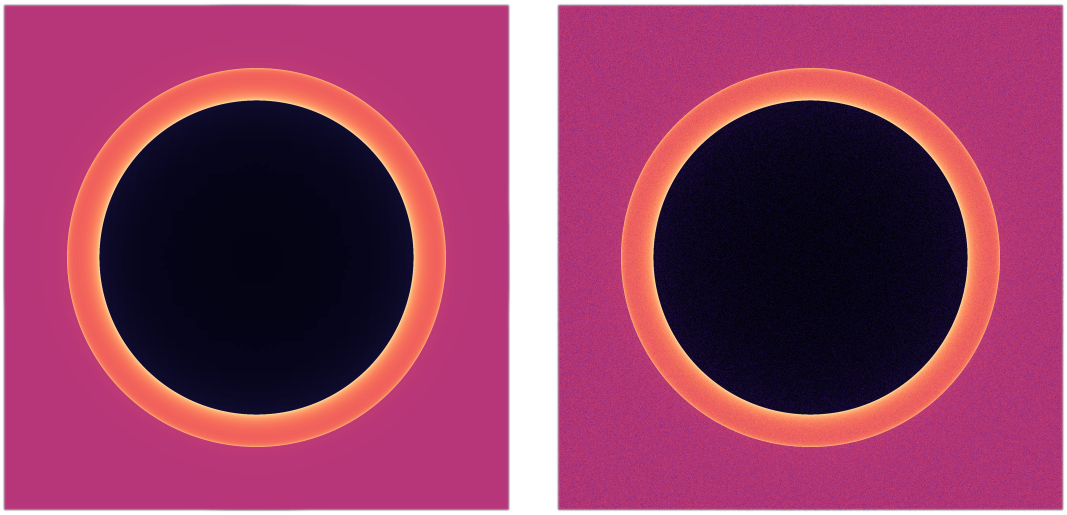


Figure A.2: Result of the direct integration method applied to the same perturbation shown in Fig. A.1, which is reproduced here at the same 2000×2000 resolution as the directly integrated image for comparison. On identical hardware, the directly integrated image was produced in 800 ms, and the stochastic image took 11 s. This speedup is possible because the direct integration method can use a cell resolution (in this case 200×200 cells, a factor of 10 fewer than the pixels in each dimension) which is smaller than the pixel resolution, whereas the stochastic method needs many (in this case 10) pseudo-particle samples per pixel.

A.3.3 Results

I now present the results of two implementations of this direct integration method for generating images. The former calculates image fluence under bilinear interpolation, applied to a rectangular grid of cells. The latter symmetrically divides each rectangular cell into four triangles by addition of a fifth point at the cell centre, and applies linear interpolation to each triangular sub-cell. Due to the complexities inherent in the bilinear interpolation case, that implementation makes use of a generic numerical integration facility provided by `MATLAB`; the implementation using linear interpolation is much better-optimised[1] and uses just-in-time compilation available in the `numba` library for python. Results produced using the linear implementation (matching the deflection map used in Fig. A.1) are shown in Fig. A.2; comparisons between linear interpolation, bilinear interpolation, and

[1] based on exact antialiasing algorithms used by vector graphics programs such as `libart`, `font-rs`, `skia` and `Pathfinder`

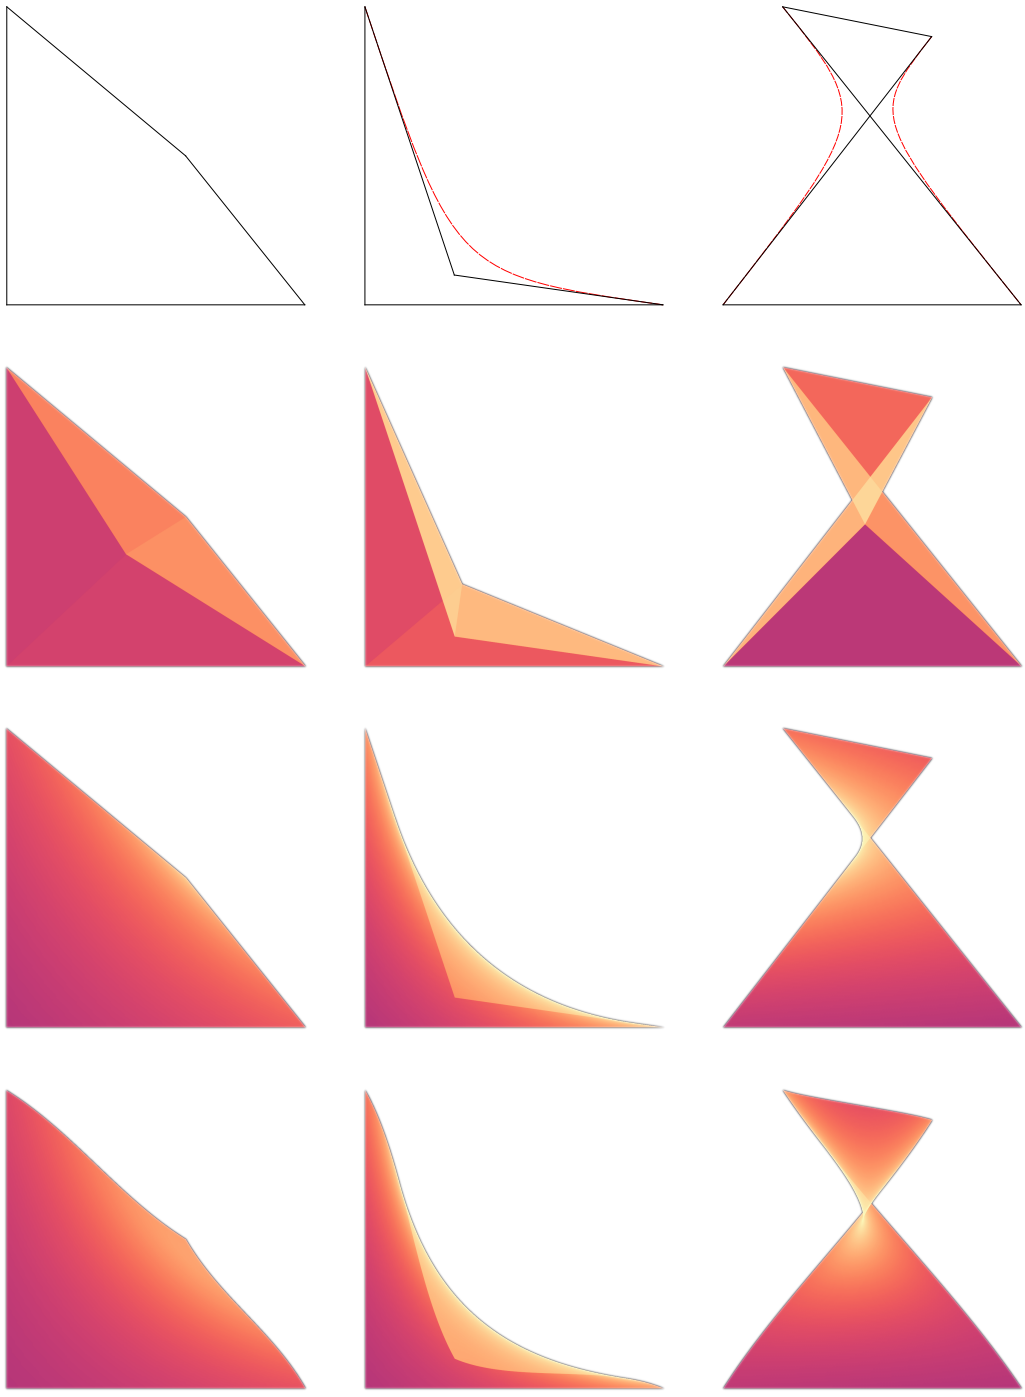


Figure A.3: The result of different interpolations on a single cell, with different perturbations. The top row shows a line diagram of the outline of the deformed shape, with possible locations for caustics drawn in red. The subsequent rows display the result of rendering this deformed quadrilateral using linear, bilinear and bicubic interpolation respectively. In the left column, the deformation results in a convex quadrilateral, so remains in the injective regime. In the middle, the resultant quadrilateral is no longer convex, but is still simple. One caustic may form. The right column shows deformation resulting in a self-intersecting 'bowtie' quadrilateral, which has two potential caustics.

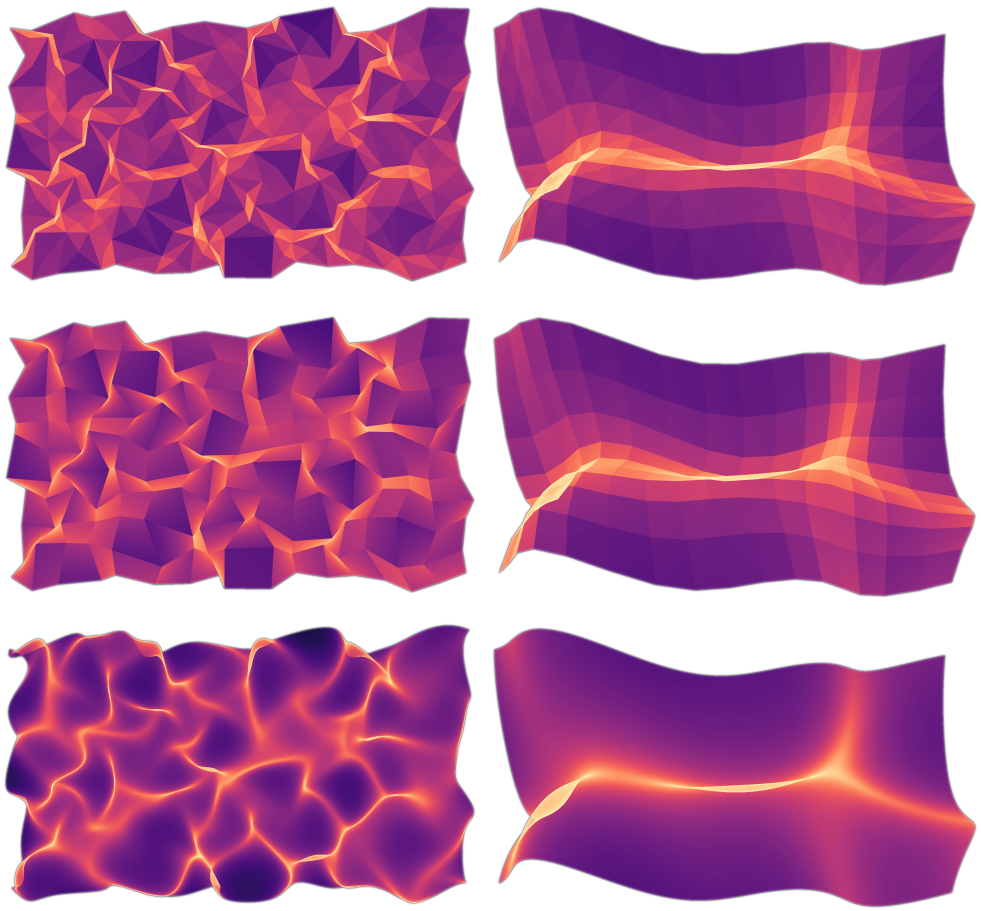


Figure A.4: Illustration of the results of the exact scheme. The left column shows rough input (no correlations between adjacent grid points, cell vertex displacements generated as random noise), and the right shows smoother input produced by convolution of random noise with a Gaussian filter. For rough input, significant differences can be seen between only using the four-triangle-per-cell linear interpolation, bilinear interpolation and bicubic interpolation (top, middle and bottom, respectively). As bilinear interpolation is not differentiable at cell boundaries, density jumps can be observed at these boundaries. Because bicubic interpolation is designed to ensure the result is differentiable everywhere, the resulting density map is smooth, even for very rough input. When the input is smoother (right column) each method gives quite similar results, but bicubic interpolation is still the only option which produces a smooth density map.

bicubic interpolation (achieved by providing supersampled inputs to the linear implementation) are shown for rough and smoother input in Fig. A.4. Further, a detailed rendering of a single cell under linear and bilinear interpolation is shown in Fig. A.3. All of the images presented here were produced using the aforementioned python code, which produces images from $\mathcal{O}(10^6)$ cells, at multi-megapixel resolution, in just a few seconds[2] on a laptop CPU.

A.4 Conclusions

The problem of producing synthetic proton radiography images has been explored, existing solutions have been presented, and a new method has been demonstrated. This new method is able to produce high-quality images, even in regions where the sampling is much sparser than the pixel grid (in which current methods suffer from large amounts of Poisson noise). Because this new method relaxes sampling requirements relative to the pixel resolution of the final image it can produce images in a much shorter time than stochastic approaches, especially at high resolutions. The new method is also entirely free of Poisson noise and is sufficiently fast that inputs can be upsampled using (e.g.) bicubic interpolation to produce smoother results, which may be necessary where inputs are noisy (see Fig. A.4).

[2]the un-interpolated images use far fewer cells: 16×9 in Fig. A.4 and one in Fig. A.3, and are produced in milliseconds.

B

Modelling Laser Pulse Propagation

Contents

B.1 Problem Statement 165
B.2 Collimated Propagation 166
B.3 Focusing 167
B.4 Focusing Beam Propagation 173
B.4.1 Discrete Implementation 178
B.5 Results 179

In this appendix I will describe in detail a piece of code written to assist in the development of a reinforcment learning-based automatic alignment system. This code was written to solve problems present in available optical propagation codes, namely that the existing codes use an excessive number of fast Fourier transforms in their propagation of beams which should only experience weak diffraction, and where this diffraction is expected to make negligible contribution to the results of the propagation.

B.1 Problem Statement

The system under consideration has the following components:

- A laser source, which produces a monochromatic beam of laser light with

a given wavelength, transverse intensity profile, position and direction of pointing

- An ordered set of planar mirrors, defined by the spatial position of their centres, their normal vectors and their radii. The interaction of the beam with each mirror is considered in turn, resulting in a change of the beam's axis of propagation and possibly clipping of its transverse profile.
- A focusing optic, namely a parabolic mirror. This is defined by the same parameters as the planar mirrors, in addition to a focal length and a measure of how 'off-axis' the mirror is.
- A camera screen, placed at a given location and with a given normal vector. The camera detects the intensity of the beam on a grid of virtual pixels.

The beam profile $A(u, v)$ is expressed in the transverse components of a beam-aligned coordinate system,

$$\hat{U} = \frac{\hat{y}\times\hat{n}}{|\hat{y}\times\hat{n}|}; \quad \hat{V} = \hat{n}\times\hat{U}; \quad \hat{n} \tag{B.1}$$

B.2 Collimated Propagation

Until the beam is reflected from the focusing mirror, it is assumed to propagate without diffraction. This approximation is made as the size of the beam is assumed to be much larger than its wavelength until it is brought to focus. While this assumption may be invalidated in the case of extreme clipping it is expected that in the case of such severe misalignment the beam is unlikely to be successfully focussed anyway.

With this assumption made the interaction of beams with planar mirrors has two parts: First, the position of the beam is updated to the intersection of its current axial vector with the mirror surface, and its normal vector is reflected in the plane of the mirror by negating its component parallel to the mirror normal. Second, the transverse profile of the beam may be clipped if it overlaps the edge of the mirror. The mirror area is projected onto the transverse plane of the beam and used as a mask by which the electric field profile is multiplied.

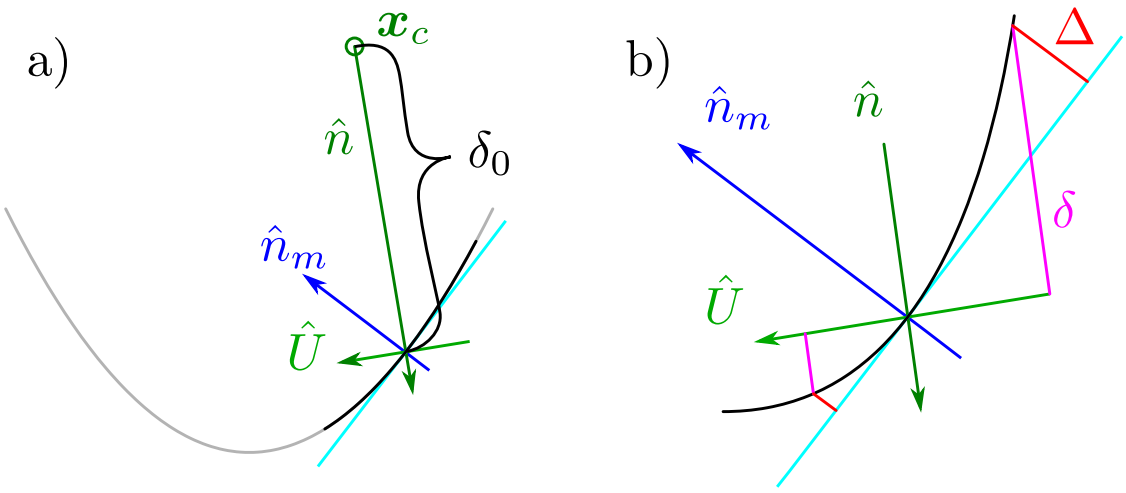


Figure B.1: Illustration of the geometry described in §B.3. (a) the projection of the beam centre onto the mirror surface, by advancing its position along its normal direction by a distance δ_0 calculated according to Eq. (B.7). The a large section of the full mathematical parabola is shown in light grey, with the actual surface of the mirror in black. The plane tangent to the parabola at the hit centre is shown in light blue. Other lines are labeled in the diagram. (b) a close-up of the geometry which defines the path difference Δ. A point in the beam tangent space (spanned by $\hat{U}, \hat{V}$) is projected along the beam normal direction to hit the surface of the mirror, and the perpendicular distance from this point on the mirror to the planar approximation is Δ. The curvature of the parabola is exaggerated in (b) to make the diagram clearer.

In addition to these processes checks are carried out to make sure that the beam hits the front face of the mirror rather than the rear and that it hits the mirror at all. If these checks fail the algorithm exits early noting that the beam has been lost.

B.3 Focusing

Once the beam has propagated through the system of planar mirrors it arrives at the parabolic focusing mirror. Once it has passed the aforementioned checks, a subroutine calculates the aberration phase of reflection from the parabolic mirror. This represents the phase difference imparted on the beam by the mirror, minus the isotropic, parabolic component which is handled separately.

For the purposes of this work, I defined the phase difference like so: For a beam whose centre ray impinges on the focusing mirror at a point $\boldsymbol{x}_h$ (the hit centre), consider a planar mirror tangent to the focusing mirror at that point. This mirror intersects the focusing mirror at the hit centre and has a normal $\hat{n}_m$ equal to the the

focusing mirror's normal at that point. If a ray in the beam is incident on a point $\boldsymbol{x}$, it picks up an optical path difference equal to the distance from that point to the planar tangent mirror. This distance is measured along the mirror's normal (and is therefore the shortest distance from the point to the plane), so can be expressed as the projection of the vector between the hit centre and $\boldsymbol{x}$ onto the mirror normal.

$$\boldsymbol{\Delta} = (\boldsymbol{x} - \boldsymbol{x}_h) \cdot \hat{n}_m \hat{n}_m \tag{B.2}$$

The phase difference is then given by the inner product of the vector displacement with the wavevector of the beam:

$$\phi = 2\boldsymbol{k} \cdot \boldsymbol{\Delta} = 2k((\boldsymbol{x} - \boldsymbol{x}_h) \cdot \hat{n}_m)(\hat{n}_m \cdot \hat{n}) \tag{B.3}$$

for a beam with normal $\hat{n}$ and wavevector $\boldsymbol{k} = k\hat{n}$ Finding $\boldsymbol{x}_h$ from a known beam centre location $\boldsymbol{x}_c$ and beam normal $\hat{n}$ makes use of the functional form of a parabola aligned—without loss of generality—along the z-axis, with its focus located at the origin:

$$z(x, y) = \frac{x^2 + y^2}{4f} - f \tag{B.4}$$

From this point onwards I will normalise all spatial dimensions according to a factor $2f$, leading to a general parabola equation

$$z(x, y) = \frac{x^2 + y^2 - 1}{2} \tag{B.5}$$

Now, finding the intersection of the central ray of a beam with the parabola simply involves equating the expression for the ray with that of the parabola

$$\boldsymbol{x}_c + \delta_0 \hat{n} = \begin{pmatrix} x_h \\ y_h \\ \frac{x_h^2 + y_h^2 - 1}{2} \end{pmatrix} \tag{B.6}$$

Substituting the x- and y-components of this vector equation into the z-component and some rearrangement leads to an equation defining δ_0:

$$\delta_0^2(n_x^2 + n_y^2) + 2\delta_0(n_x x_c + n_y y_c - n_z) + x_c^2 + y_c^2 - 1 - 2z_c = 0 \tag{B.7}$$

The mirror normal at the point of intersection is given by differentiating Eq. (B.5):

$$\hat{n}_m = \begin{pmatrix} -x_h \\ -y_h \\ 1 \end{pmatrix} \cos\Phi = \boldsymbol{a}\cos\Phi, \tag{B.8}$$

where for convenience I have defined the angle Φ between the z-axis and the mirror normal, i.e. $\cos\Phi = \hat{n}_m \cdot \hat{z} = (1+x_c^2+y_c^2)^{-\frac{1}{2}}$ and the 'un-normalised' mirror normal $\boldsymbol{a}$.

Defining also the beam-mirror angle $\cos\Phi_{bm} = \hat{n}\cdot\hat{n}_m$, the phase can be expressed as

$$\phi = 2k\cos\Phi_{bm}\cos\Phi((\boldsymbol{x}-\boldsymbol{x}_h)\cdot\boldsymbol{a}) \tag{B.9}$$

The far-right term here is the interesting one. Expanding out the dot product:

$$((\boldsymbol{x}-\boldsymbol{x}_h)\cdot\boldsymbol{a}) = x_h^2 - xx_h + y_h^2 - yy_h + \frac{x^2+y^2-1}{2} - \frac{x_h^2+y_h^2-1}{2} \tag{B.10}$$

$$= \frac{1}{2}\left(\left(x_h^2 - 2x_h x + x^2\right) + \left(y_h - 2y_h y + y^2\right)\right) \tag{B.11}$$

$$= \frac{1}{2}\left((x-x_h)^2 + (y-y_h)^2\right) \tag{B.12}$$

$$\phi = k\cos\Phi_{bm}\cos\Phi\left((x-x_h)^2 + (y-y_h)^2\right) \tag{B.13}$$

Now, the procedure for evaluating the phase-shift applied to a beam on reflection from the mirror is as follows:

Starting from the initial beam centre position $\boldsymbol{x}_c$, the beam profile is advanced a distance δ_0 (according to Eq. (B.7)) such that the beam centre is transported to the surface of the mirror at the point $\boldsymbol{x}_h$. Now, points in the beam-transverse plane are expressed using the beam-transverse coordinates (u,v) and have three-dimensional positions $\boldsymbol{x}(u,v) = \boldsymbol{x}_h + u\hat{U} + v\hat{V}$. These three-dimensional positions are also projected onto the surface of the parabolic mirror by calculating $\delta(u,v)$ for each original position. As the beam centre has already been traced to the mirror surface, $\delta(0,0) = 0$.

Then, the path difference in beam-transverse coordinates is

$$\Delta(u,v) = \frac{1}{2}\cos\Phi\left((uU_x + vV_x + \delta(u,v)n_x)^2 + (uU_y + vV_y + \delta(u,v)n_y)^2\right) \tag{B.14}$$

While it is possible to evaluate this expression directly, it's not particularly enlightening and doesn't shed much light on the way the phase difference depends on the misalignment components of the beam normal n_x, n_y. It's also going to be necessary to calculate (and exclude) the isotropic focusing component (proportional to $u^2 + v^2$) as this will be handled geometrically by the focusing propagation code I present in the next section.

To elucidate the effect of small misalignments on the phase I will perform a Taylor expansion of $\Delta(u, v)$ in n_x, n_y. To first order in these variables,

$$\hat{U} \approx \begin{pmatrix} -1 \\ 0 \\ -n_x \end{pmatrix} \quad \hat{V} \approx \begin{pmatrix} 0 \\ 1 \\ n_y \end{pmatrix} \quad \hat{n} \approx \begin{pmatrix} n_x \\ n_y \\ -1 \end{pmatrix}, \tag{B.15}$$

where the zero-order normal vector is the negative z unit vector[1]. This ensures that the beam moves toward the mirror, not away from it.

The distance $\Delta(u, v)$ is similarly Taylor-expanded:

$$\Delta(u, v) = \frac{1}{2} \cos\Phi \left((-u + \delta(u, v) n_x)^2 + (v + \delta(u, v) n_y)^2 \right) \tag{B.16}$$

$$= \frac{1}{2} \cos\Phi \left(u^2 + v^2 + 2\delta(u, v)(v n_y - u n_x) \right) \tag{B.17}$$

This expression uses the zero-order value of $\delta(u, v)$:

$$\delta = u x_0 - v y_0 - \frac{u^2 + v^2}{2} \tag{B.18}$$

and yields:

$$\frac{1}{2} \cos\Phi \left[A\rho^2 + (B(u^2 - v^2) + 2Cuv) + (u n_x - v n_y)\rho^2 \right]. \tag{B.19}$$

Here $\rho^2 = u^2 + v^2$, the magnitude of the isotropic focusing component of the phase is given by

$$A = 1 - (x_0 n_x + y_0 n_y), \tag{B.20}$$

and the anisotropic component causing astigmatism is described by two constants,

$$B = y_0 n_y - x_0 n_x \tag{B.21}$$

$$C = x_0 n_y + y_0 n_x \tag{B.22}$$

[1] These vectors are of course no longer exact unit vectors, but their magnitude is second-order in the small components n_x, n_y

A third-order component causing comatic aberration is also present:

$$(un_x - vn_y)(u^2 + v^2) \tag{B.23}$$

Recall that these are only computed to first order in the misalignment, represented by the transverse normal components n_x, n_y. Because of the importance of the quadratic (focusing) phase to the position of the focal point of a beam after reflection from the mirror, I seek to improve the accuracy of that component. In order to achieve this I expand Eq. (B.14) as a series in (u, v), while allowing (n_x, n_y) to have arbitrary magnitude[2].

To begin, I denote $x' = uU_x + vV_x + \delta n_x$ and $y' = uU_y + vV_y + \delta n_y$. Then,

$$\Delta = \frac{1}{2}\cos\Phi\left(x'^2 + y'^2\right) \tag{B.24}$$

$$\boldsymbol{\nabla}\Delta = \cos\Phi\left(x'\boldsymbol{\nabla}x' + y'\boldsymbol{\nabla}y'\right), \tag{B.25}$$

where the gradient is taken with respect to u and v.

Evaluating these expressions at $u = 0, v = 0$, I find that δ is also zero (as the hit centre is by definition on the mirror surface) and therefore $x', y' = 0$, $\Delta = 0$ and $\boldsymbol{\nabla}\Delta = \mathbf{0}$. This is expected, as the construction of the problem ensures that there is no phase component linear in the beam-transverse coordinates (u, v)[3]. Taking the second derivatives, evaluated again at the point $(u, v) = 0$ (and so dropping terms proportional to x', y' and δ)

$$\frac{\partial^2\Delta}{\partial u^2} = \cos\Phi\left[\left(\frac{\partial x'}{\partial u}\right)^2 + \left(\frac{\partial y'}{\partial u}\right)^2\right] \tag{B.26}$$

$$\frac{\partial^2\Delta}{\partial v^2} = \cos\Phi\left[\left(\frac{\partial x'}{\partial v}\right)^2 + \left(\frac{\partial y'}{\partial v}\right)^2\right] \tag{B.27}$$

$$\frac{\partial^2\Delta}{\partial u\,\partial v} = \cos\Phi\left[\frac{\partial x'}{\partial u}\frac{\partial x'}{\partial v} + \frac{\partial y'}{\partial u}\frac{\partial y'}{\partial v}\right] \tag{B.28}$$

[2]Though of course, remaining constrained by $n_x^2 + n_y^2 \leq 1$

[3]The aberration the opthalmic profession refers to as 'prismatic', and scientists refer to as 'laser going in the wrong direction'

Some algebraic manipulation leads to expressions for the relevant derivatives:

$$\frac{\partial x'}{\partial u} = \frac{-\hat{V}\cdot\begin{pmatrix}0\\1\\y_h\end{pmatrix}}{\hat{n}\cdot\boldsymbol{a}} \qquad \frac{\partial y'}{\partial u} = \frac{\hat{V}\cdot\begin{pmatrix}1\\0\\x_h\end{pmatrix}}{\hat{n}\cdot\boldsymbol{a}} \tag{B.29}$$

$$\frac{\partial x'}{\partial v} = \frac{\hat{U}\cdot\begin{pmatrix}0\\1\\y_h\end{pmatrix}}{\hat{n}\cdot\boldsymbol{a}} \qquad \frac{\partial y'}{\partial v} = \frac{-\hat{U}\cdot\begin{pmatrix}1\\0\\x_h\end{pmatrix}}{\hat{n}\cdot\boldsymbol{a}} \tag{B.30}$$

This leads to the following expressions for the second derivatives:

$$C_{uu} = \frac{\partial^2\Delta}{\partial u^2} = \cos\Phi\frac{\hat{V}\cdot\underline{\underline{\mathrm{M}}}\cdot\hat{V}}{(\hat{n}\cdot\boldsymbol{a})^2} \tag{B.31}$$

$$C_{vv} = \frac{\partial^2\Delta}{\partial v^2} = \cos\Phi\frac{\hat{U}\cdot\underline{\underline{\mathrm{M}}}\cdot\hat{U}}{(\hat{n}\cdot\boldsymbol{a})^2} \tag{B.32}$$

$$C_{uv} = \frac{\partial^2\Delta}{\partial u\,\partial v} = -\cos\Phi\frac{\hat{U}\cdot\underline{\underline{\mathrm{M}}}\cdot\hat{V}}{(\hat{n}\cdot\boldsymbol{a})^2}, \tag{B.33}$$

where the matrix $\underline{\underline{\mathrm{M}}}$ has components

$$\underline{\underline{\mathrm{M}}} = \begin{pmatrix}1 & 0 & x_h\\ 0 & 1 & y_h\\ x_h & y_h & r_h^2\end{pmatrix} \tag{B.34}$$

and $r_h^2 = x_h^2 + y_h^2$.

The path difference evaluated at second order in the beam coordinates is therefore:

$$\Delta_2(u,v) = \frac{1}{2}\left[u^2\frac{\partial^2\Delta}{\partial u^2} + 2uv\frac{\partial^2\Delta}{\partial u\,\partial v} + v^2\frac{\partial^2\Delta}{\partial v^2}\right] \tag{B.35}$$

$$= \frac{1}{2}\cos\Phi\frac{(v\hat{U}-u\hat{V})\cdot\underline{\underline{\mathrm{M}}}\cdot(v\hat{U}-u\hat{V})}{(\hat{n}\cdot\boldsymbol{a})^2} \tag{B.36}$$

Alternatively, separating the focusing and astigmatic elements of the phase:

$$\Delta_2(u,v) = \frac{1}{2}\cos\Phi\left[\frac{C_{uu}+C_{vv}}{2}(u^2+v^2) + \frac{C_{uu}-C_{vv}}{2}(u^2-v^2) + 2C_{uv}uv\right] \tag{B.37}$$

By defining the mean and half-difference constants $\bar{C} = (C_{uu}+C_{vv})/2$ and $C_\Delta = (C_{uu}-C_{vv})/2$ and diagonalising the 'astigmatism matrix'

$$\begin{pmatrix}C_\Delta & C_{uv}\\ C_{uv} & -C_\Delta\end{pmatrix} = \begin{pmatrix}\cos\theta & -\sin\theta\\ \sin\theta & \cos\theta\end{pmatrix}\begin{pmatrix}A & 0\\ 0 & -A\end{pmatrix}\begin{pmatrix}\cos\theta & \sin\theta\\ -\sin\theta & \cos\theta\end{pmatrix} \tag{B.38}$$

$$\tan 2\theta = \frac{C_{uv}}{C_\Delta} \tag{B.39}$$

$$A^2 = C_\Delta^2 + C_{uv}^2 \tag{B.40}$$

it is clear that this can be expressed most cleanly as

$$\Delta_2(u,v) = \frac{1}{2}\cos\Phi\left[\bar{C}\rho^2 + A(\xi^2 - \zeta^2)\right], \tag{B.41}$$

in the rotated coordinate system given by

$$\begin{pmatrix}\xi\\ \zeta\end{pmatrix} = \begin{pmatrix}\cos\theta & \sin\theta\\ -\sin\theta & \cos\theta\end{pmatrix}\begin{pmatrix}u\\ v\end{pmatrix} \tag{B.42}$$

$$\rho^2 = u^2 + v^2 = \xi^2 + \zeta^2 \tag{B.43}$$

The final expression for the approximate phase difference is now:

$$\phi \approx k\cos\Phi_{bm}\cos\Phi\left[\bar{C}\rho^2 + A(\xi^2 - \zeta^2) + (un_x - vn_y)\rho^2\right] \tag{B.44}$$

Here the first term in the square brackets is responsible for focusing, the second produces astigmatism and the third causes coma.

Comparing to the expression for focusing from a perfectly aligned, on-axis parabolic mirror, $\phi = k\rho^2$ it is easy to see that the effective focal length of the misaligned aligned off-axis parabola is

$$f' = \sec\Phi\sec\Phi_{bm}\bar{C}^{-1}f_0. \tag{B.45}$$

If the alignment is corrected, $\Phi_{bm} \to \Phi$ and $\bar{C} \to 1$, so the effective focal length of an off-axis parabola is the parent focal length multiplied by $\sec^2\Phi$, which is indeed the distance from a point on the parabola to the focus.

B.4 Focusing Beam Propagation

It is now necessary to propagate the focusing beam from the mirror to an observation point. In principle this could be done in a number of ways, such as employing a Fresnel propagation approach. However, many methods would require fully resolving the very quickly varying phase which leads to focusing across the entire, large beam profile. This phase has the form

$$\phi_{\text{foc}} = \frac{kr^2}{2f} \tag{B.46}$$

$$|\boldsymbol{\nabla}\phi_{\text{foc}}| = 2\pi\frac{r}{\lambda f} \tag{B.47}$$

The grid used to evolve the beam would need to have grid spacing Δx satisfying $|\boldsymbol{\nabla}\phi_{\text{foc}}|\Delta x < 2\pi$ to resolve this phase variation without aliasing. As the phase gradient increases with radius this condition is strictest at the outer edge of the beam, $r = R$. Using the f-number of the beam, $F/\# = f/(2R)$ the condition can be rewritten

$$\Delta x < 2\lambda F/\# \tag{B.48}$$

This implies that the beam profile should be resolved with a resolution of the order of the wavelength, as f-numbers of order unity are common in high-energy laser experiments. The number of grid points per dimension would therefore need to be larger than R/λ, and as these experiments also employ large (order 10 cm) expanded beams to avoid damage to optical elements this poses a significant computational workload. I therefore seek another approach.

I take inspiration from the 'laser envelope model' approach, which reduces the computational demands of modelling a laser pulse by factoring out the fast oscillating factor $\exp i(kx - \omega t)$ and deriving new evolution equations for the remaining, relatively slowly-varying envelope of the pulse. In this case however I will also factor out the parabolic phase component that causes focusing, which will require re-defining the spatial variables transverse and longitudinal to the pulse.

I begin from the paraxial Helmholtz equation[4]

$$\nabla_\perp^2 E + 2ik_0 \frac{\partial E}{\partial z} = 0 \tag{B.49}$$

And the mathematical form of the Gaussian beam, which solves this equation[5]:

$$E_{0,0}(x, y, z) = E_0 \frac{w_0}{w(z)} \exp\left(-\frac{r^2}{w(z)^2}\right) \exp\left(i\left(k_0 \frac{r^2}{2R(z)} - \arctan\left(\frac{z}{z_R}\right)\right)\right). \tag{B.50}$$

This standard expression for the Gaussian beam employs the wavefront curvature $R(z) = z(1 + (\frac{z_R}{z})^2)$. Observing that the reciprocal curvature may also be expressed

[4]The Helmholtz equation is the steady-state form of the laser envelope model in vacuum
[5]The meaning of the subscript $0, 0$ will become clear shortly

as $R^{-1}(z) = z_R^{-1}\frac{w_0^2}{w(z)^2}\frac{z}{z_R}$ and employing the relation defining the Rayleigh range $z_R = k_0 w_0^2/2$ the Gaussian beam is instead expressed as:

$$E_{0,0}(x,y,z) = E_0 \frac{w_0}{w(z)} \exp\left(-\frac{r^2}{w(z)^2}\right) \exp\left(i\left(\frac{r^2}{w(z)^2}\frac{z}{z_R} - \arctan\left(\frac{z}{z_R}\right)\right)\right). \tag{B.51}$$

In this expression, observe that wherever the transverse coordinate $r = \sqrt{x^2+y^2}$ appears it is divided by the beam width $w(z)$, and the longitudinal coordinate z is always divided by the Rayleigh length z_R, motivating re-parametrisation of the beam in terms of normalised transverse coordinates $\bar{x} = x/w(z),\ \bar{y} = y/w(z)$ and $\bar{r} = r/w(z)$ and the Guoy phase $\psi = \arctan(z/z_R)$:

$$E_{0,0}(x,y,z) = E_0 \cos\psi \exp\left(-\bar{r}^2\right) \exp\left(i\left(\bar{r}^2\tan\psi - \psi\right)\right). \tag{B.52}$$

I now re-express the derivatives involved in the paraxial Helmholtz equation in terms of the $\bar{x}-$, $\bar{y}-$ and $\psi-$derivatives:

$$\frac{\partial}{\partial x} = w(z)^{-1}\frac{\partial}{\partial \bar{x}} = w_0^{-1}\cos\psi\frac{\partial}{\partial\bar{x}} \tag{B.53a}$$

$$\frac{\partial^2}{\partial x^2} = w(z)^{-2}\frac{\partial^2}{\partial(\bar{x})^2} = w_0^{-2}\cos^2\psi\frac{\partial^2}{\partial(\bar{x})^2} \tag{B.53b}$$

$$\frac{\partial}{\partial y} = w(z)^{-1}\frac{\partial}{\partial \bar{y}} = w_0^{-1}\cos\psi\frac{\partial}{\partial\bar{y}} \tag{B.53c}$$

$$\frac{\partial^2}{\partial y^2} = w(z)^{-2}\frac{\partial^2}{\partial(\bar{y})^2} = w_0^{-2}\cos^2\psi\frac{\partial^2}{\partial(\bar{y})^2} \tag{B.53d}$$

$$\frac{\partial}{\partial z} = z_R^{-1}\cos^2\psi\left(\frac{\partial}{\partial\psi} - \tan\psi\left(\bar{x}\frac{\partial}{\partial\bar{x}} + \bar{y}\frac{\partial}{\partial\bar{y}}\right)\right) \tag{B.53e}$$

and therefore the paraxial Helmholtz equation itself:

$$4i\frac{\partial E}{\partial\psi} = -\bar{\nabla}_\perp^2 E + 4i\tan\psi\ \bar{\boldsymbol{r}}\cdot\boldsymbol{\nabla}E. \tag{B.54}$$

Now I inspect the more general *Hermite-Gaussian* beam modes, which can be expressed similarly to the Gaussian beam:

$$E_{l,m}(x,y,z) \propto E_{0,0}(x,y,z) H_l\left(\frac{\sqrt{2}x}{w(z)}\right) H_m\left(\frac{\sqrt{2}y}{w(z)}\right) \tag{B.55}$$

$$\exp\left(-i(l+m)\arctan\left(\frac{z}{z_R}\right)\right) \tag{B.56}$$

$$\propto E_0 \cos\psi \; h_l\left(\sqrt{2}\bar{x}\right) h_m\left(\sqrt{2}\bar{y}\right) \tag{B.57}$$

$$\exp\left(i\left(\bar{r}^2\tan\psi - (1+l+m)\psi\right)\right), \tag{B.58}$$

introducing the Hermite polynomials $H_n(x)$ and the so-called Hermite functions $h_n(x) = H_n(x)\exp(-x^2/2)$[6]. The Gaussian beam is the (0,0) case of Hermite-Gaussian beams, and the set of all (l,m) Hermite-Gaussian beams forms an Orthogonal basis of the space of beams of any transverse profile.

Observing that every Hermite-Gaussian beam has phase dependence including a common $\bar{r}^2\tan\psi$ term, and also $\cos\psi$ amplitude dependence, I seek to factor out these common factors. I now represent a general transverse electric field profile as

$$E(\bar{x},\bar{y},\psi) = A(\bar{x},\bar{y},\psi)\cos\psi\exp\left(i\bar{r}^2\tan\psi\right). \tag{B.59}$$

Evaluating the normalised paraxial Helmholtz equation B.54 on an electric field of this form and eliminating terms yields:

$$4i\frac{\partial A}{\partial\psi} = -\bar{\nabla}_\perp^2 A + 4\bar{r}^2 A. \tag{B.60}$$

This equation has the same form as the laser envelope model in the presence of a parabolic susceptibility function. Parabolic channels of this kind are well-known to counteract diffraction of laser pulses, and it therefore makes sense that a mathematical transformation which compensates for focusing would reproduce a similar governing equation. This correspondence suggests that algorithms for beam propagation in media—for example, the 'Fresnel slicing' method—may be applicable to solving this problem also. Approaches I have taken to solving the propagaton problem are described in more detail in the next section.

[6]The hermite functions are most commonly denoted ψ_n in the literature, but I use h_n to avoid confusion with the Guoy phase ψ.

Eq. (B.60) also is also identical in form to the time-dependent Schrödinger equation for a quantum harmonic oscillator with angular frequency ω_0 and mass $m = 2\hbar/\omega_0$, when the Guoy phase is interpreted as $\psi = \omega_0 t$. This makes explicit the link between the Helmholtz equation of classical electromagnetism and the quantum harmonic oscillator which is implied by the fact that the Hermite functions which form the transverse dependence of the Hermite-Gaussian modes are also (when appropriately scaled) eigenstates of the quantum harmonic oscillator.

Some well-known results applying to evolution under the Schrödinger equation therefore directly translate to evolution of the transverse field profile in these "focus-relative" coordinates. Most notably the evolution is unitary (that is, norm-preserving) - this relates to conservation of beam energy in the same way as unitarity of quantum systems relates to conservation of probability.

It is also interpretable as a (classical) Hamiltonian system, when the real and imaginary parts of A act as the canonical coordinates analogous to the position and momentum of classical Hamiltonian mechanics respectively. The "Hamiltonian" of the system interpreted this way[7] is separable and given by:

$$\mathcal{H} = \int A_r \bar{\nabla}_\perp^2 A_r + A_i \bar{\nabla}_\perp^2 A_i - 4\bar{r}^2 |A|^2 dxdy \tag{B.61}$$

The canonical equations relating to this Hamitonian are the real and imaginary parts of Eq. (B.60).

The symplectic structure of this Hamiltonian system may be used to apply symplectic integration schemes to Eq. (B.60), improving the stability of solutions. However, I have found an 'analytic' solution technique founded in decomposition of the complex beam profile to be superior to numerical solution of Eq. (B.60) for evaluating the evolution of this system. I will now briefly describe this analytic method. Recall that the Hermite-Gaussian beam modes solve exactly the paraxial Helmholtz equation (and are the eigenfunctions of the quantum harmonic oscillator). It is therefore obvious that appropriately-scaled products of the Hermite functions,

[7]This is the Hamiltonian in the mathematical sense, and is unrelated to usual physical interpretations of Hamiltonians

$A_{lm} = h_l(\sqrt{2}\bar{x})h_m(\sqrt{2}\bar{y})$ which describe the transverse profiles of these modes can be used to diagonalise the right-hand side of Eq. (B.60):

$$-\bar{\nabla}_{\perp}^2 A_{lm} + 4\bar{r}^2 A_{lm} = 4(1+l+m)A_{lm} \tag{B.62}$$

In this Hermite basis Equation B.60 therefore reduces to

$$i\frac{\partial A_{lm}}{\partial \psi} = (1+n+m)A_{lm}, \tag{B.63}$$

which is solved by

$$A_{lm}(\psi) = e^{-i(1+l+m)\psi}A_{lm}(\psi = 0). \tag{B.64}$$

B.4.1 Discrete Implementation

To make use of Equation B.64 it is necessary to decompose A into a sum of products of Hermite functions of different orders. Fortunately these functions form an orthonormal basis on $\mathbb{L}_2$, making it possible to decompose arbitrary functions using the matrix inner product. Firstly, Hermite functions are enumerated up to a desired bandlimit N, using the form of the first two functions and then a standard recurrence relation to calculate as many higher functions as are desired:

$$h_0(\sqrt{2}\bar{x}) = e^{-\bar{x}^2} \tag{B.65}$$

$$h_1(\sqrt{2}\bar{x}) = 2xe^{-\bar{x}^2} \tag{B.66}$$

$$h_{n+1}(\sqrt{2}\bar{x}) = (2xh_n - \sqrt{n}h_{n-1})/\sqrt{n+1} \tag{B.67}$$

The functions are then normalised, muiltiplying by a factor $\sqrt{\Delta\bar{x}}(\pi/2)^{-1/4}$ to ensure that each discretised hermite function $h_n(\sqrt{2}\bar{x}_j)$ satisfies the normalisation $\sum_j h_n^2(\sqrt{2}\bar{x}_j) = 1$. Denoting the matrices containing discretised hermite functions as $\mathbb{H}_{nj} = h_n(\sqrt{2}\bar{x}_j)$ and $\mathbb{H}_{mk} = h_m(\sqrt{2}\bar{y}_k)$ and the discretised normalised laser field as $\mathbb{A}_{jk} = A(\bar{x}_j, \bar{y}_k)$, expansion coefficients $c_n m$ are efficiently calculated and manipulated, and the evolved field pattern reconstructed as:

$$c_{nm} = \mathbb{H}_{nj} \cdot \mathbb{A}_{jk} \cdot \mathbb{H}^T_{km} \tag{B.68}$$

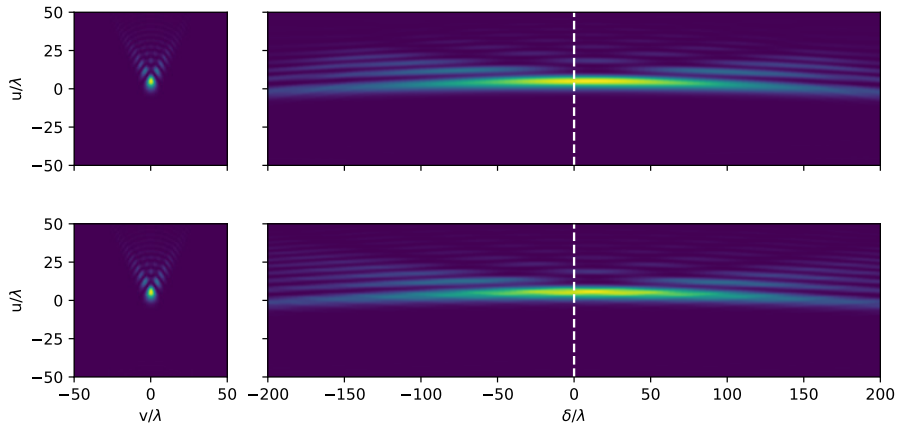
$$c'_{nm} = e^{-i(1+n+m)\Delta\psi} c_{nm} \tag{B.69}$$

$$\mathbb{A}'_{jk} = \mathbb{H}^T_{jn} \cdot c_{nm} \cdot \mathbb{H}_{mk} \tag{B.70}$$

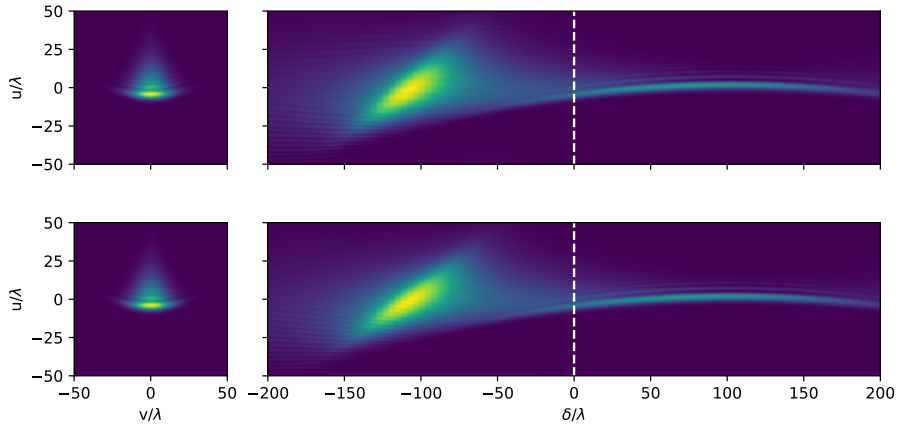
This technique has the advantage of being effectively random-access: the field at any z-position can be evaluated in the same amount of time, as opposed to the sequential calculation involved in evolving Equation B.60 directly, which must use small enough step sizes to maintain accuracy.

B.5 Results

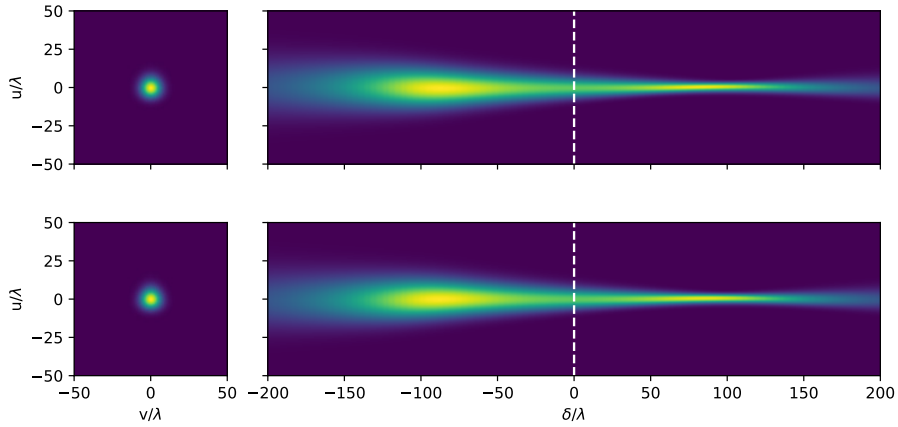
I now present a comparison between the results of the paraxial phase model of focusing by a parabolic mirror presented in this chapter and the prediction of the Rayleigh-Sommerfeld diffraction integral, which is taken as the 'ground truth' result (for monochromatic, scalar wave propagation) as it exactly solves the Helmholtz equation which my method produces an approximation to. The diffraction integral is evaluated at positions in absolute space, so to align the result with that of my paraxial method I use the focal point and beam normal predicted by the paraxial method. Figures B.2, B.3 and B.4 compares the focal spot shapes predicted by the two methods for a variety of off-axis angles and a variety of beam misalignments, showing excellent agreement throughout. The evolution away from focus is also shown, along the direction of beam propagation. Both the location of the focus and the direction of propagation are predicted correctly by the approximate method.



a) Misalignment of 5 mrad, on-axis; coma-dominated.

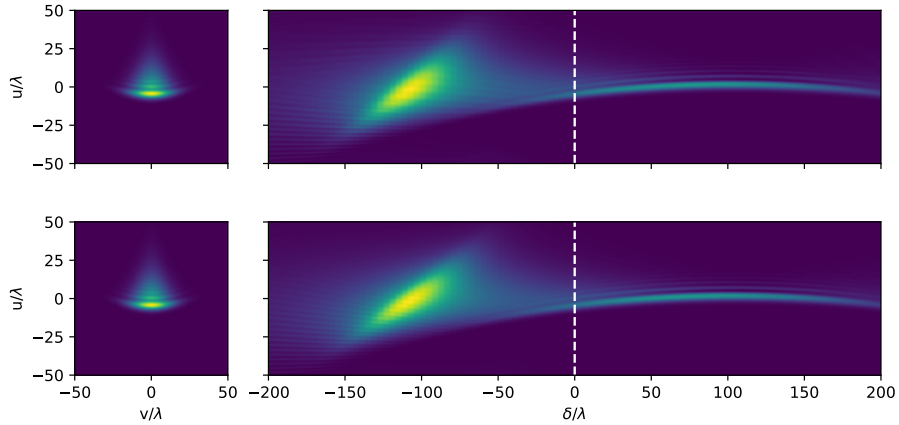


b) Misalignment of 1 mrad, 22.5° off-axis; coma and astigmatism are both apparent.

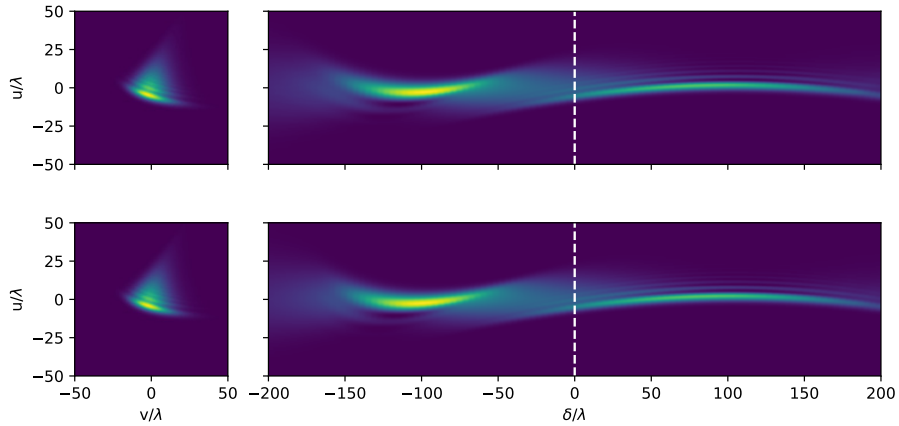


c) Misalignment of 0.1 mrad, 90° off-axis; astigmatism-dominated.

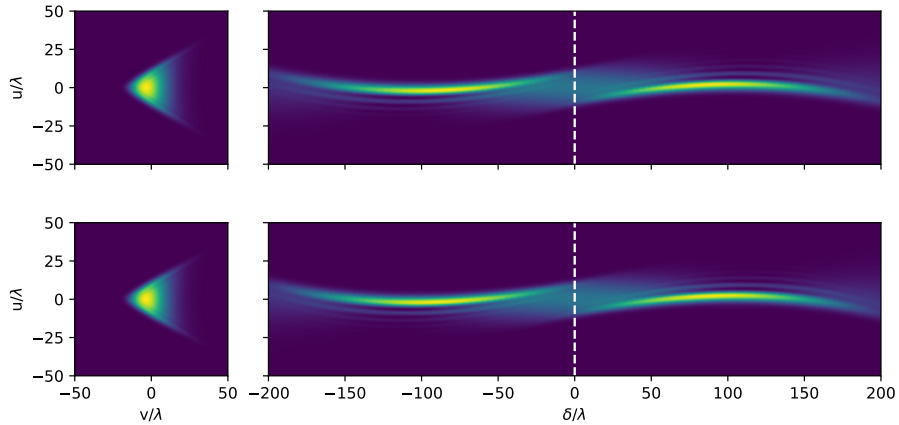
Figure B.2: Varying off-axis angle. In each sub-figure, the upper panes show the Rayleight-Sommerfeld integral result and those on the bottom show the prediction of the method presented here. As predicted by §B.3, astigmatism becomes increasingly significant relative to coma as the mirror is used further from its axis. All misalignments are horizontal (in off-axis plane). All input beams have $F_{\#} = 3$



a) Horizontal misalignment of 1 mrad.

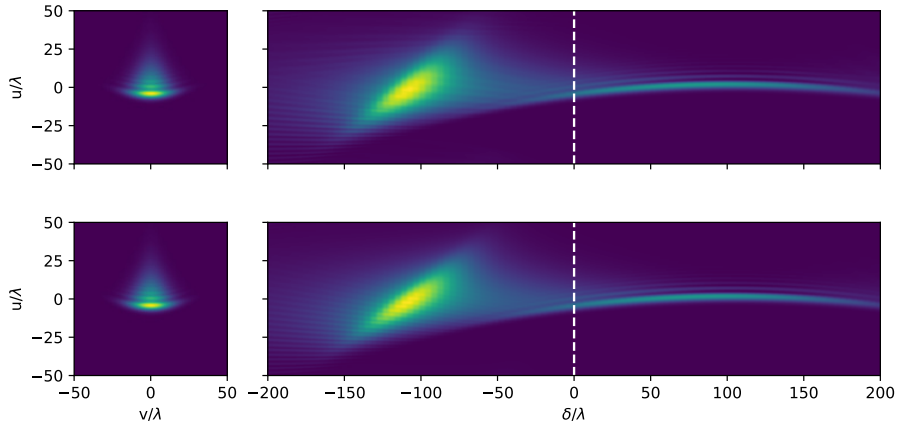


b) Diagonal misalignment of 1 mrad.

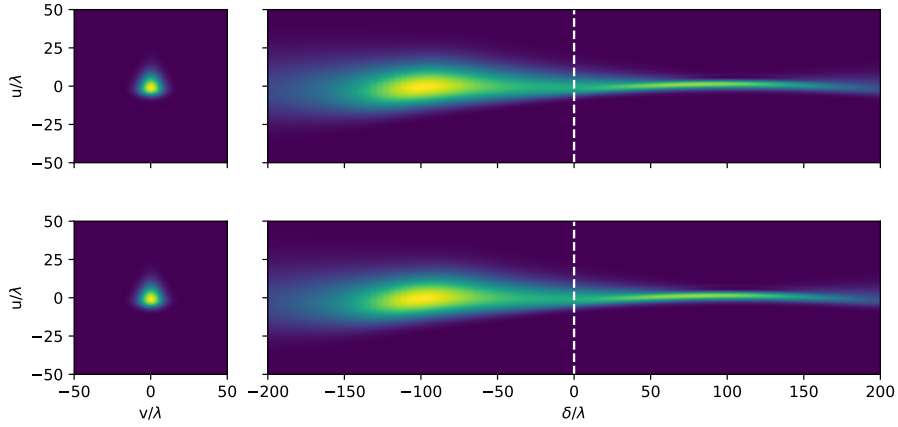


c) Vertical misalignment of 1 mrad.

Figure B.3: Varying misalignment direction. In each sub-figure, the upper panes show the Rayleight-Sommerfeld integral result and those on the bottom show the prediction of the method presented here.



a) $F_{\#} = 3$ beam.



b) $F_{\#} = 5$ beam.

Figure B.4: Varying $F_{\#}$. In each sub-figure, the upper panes show the Rayleight-Sommerfeld integral result and those on the bottom show the prediction of the method presented here. The smaller input beam b) shows less aberration in the presence of the same 1 mrad misalignment.

C

Drifting Boltzmann Relation

In Chapter 4 a Boltzmann relation is used as a substitute for describing the full electron dynamics. At first glance this seems like a reasonable approximation, however it is unclear that the Boltzmann approximation is applicable in a plasma which is drifting relative to the electrostatic structure providing the electric potential which appears in the Boltzmann relation.

As the Boltzmann relation is a fluid relation rather than a kinetic one it doesn't explicitly refer to the bulk drift velocity of the electrons or their phase-space distribution function, only to their number density and their pressure (or equivalently temperature).

Taking the electron fluid equation of motion (Eq. (2.52c)) and setting the electron mass to zero, the remaining terms (in the absence of magnetic field) give

$$\boldsymbol{\nabla} P_e = q_e n_e \boldsymbol{E} \tag{C.1}$$

Taking an isothermal ($\gamma = 1$) equation of state and writing the electric field in terms of the electric potential ϕ,

$$k_B T_e \boldsymbol{\nabla} n_e = e n_e \boldsymbol{\nabla} \phi \tag{C.2}$$

Integrating both sides yields the usual form of the Boltzmann relation:

$$n_e = n_0 \exp\left(\frac{e\phi}{k_B T_e}\right) \tag{C.3}$$

I will now investigate the effect of a drift velocity V on this form of the electron density. The one-dimensional, drifting, isothermal electron distribution function following this density relationship is

$$f_e(x,v) = \frac{n_0}{v_{\rm th}\sqrt{\pi}} \exp\left(-\frac{\frac{1}{2}m_e(v-V)^2 - e\phi}{k_B T_e}\right), \tag{C.4}$$

where $v_{\rm th} = \sqrt{2k_B T_e/m_e}$ is the thermal speed.

I now insert this distribution function into the Vlasov equation (Eq. (2.20a)) and evaluate the x- and v-derivatives.

$$\frac{\partial f_e}{\partial t} + v\frac{\partial f_e}{\partial x} + \frac{e}{m_e}\frac{\partial \phi}{\partial x}\frac{\partial f_e}{\partial v} = 0 \tag{C.5}$$

$$\frac{\partial f_e}{\partial t} + v\frac{e}{k_B T_e}\frac{\partial \phi}{\partial x} f_e - \frac{e}{m_e}\frac{\partial \phi}{\partial x}\frac{m_e}{k_B T_e}(v-V) f_e = 0 \tag{C.6}$$

$$\frac{\partial f_e}{\partial t} = -V\frac{e}{k_B T_e}\frac{\partial \phi}{\partial x} f_e \tag{C.7}$$

$$= V\frac{eE}{k_B T_e} f_e \tag{C.8}$$

This result shows that the Boltzmann relation holds for non-drifting plasma: when $V = 0$, the time-derivative of the distribution function vanishes and the distribution function Eq. (C.4) is therefore a stationary solution of the Vlasov equation. When $V \neq 0$ however, the distribution will experience exponential growth (in regions where V is parallel to E) or decay (in regions where they are antiparallel). Local thermodynamic equilibrium is still maintained to first order, as the growth/decay rate is not a function of phase-space velocity (only of the fluid drift velocity). To quantify the failure of the Boltzmann relation I now claim that for plasma drifting into an electrostatic structure with velocity V, the Boltzmann relation is applicable for times $t \ll \tau_B$, where:

$$\tau_B = \frac{k_B T_e}{e|E||V|} \tag{C.9}$$

For the conditions of Haberberger et al. [114] (1 MeV electrons, maximum electric field of order $1 \times 10^{13}\,\mathrm{V\,m^{-1}}$, expansion velocity $0.1c$) this expression gives a stability time of a few fs.

References

[1] Francis F. Chen. *Introduction to Plasma Physics and Controlled Fusion*. Cham: Springer International Publishing, 2016. DOI: 10.1007/978-3-319-22309-4.

[2] S Nakai and H Takabe. "Principles of Inertial Confinement Fusion - Physics of Implosion and the Concept of Inertial Fusion Energy". In: *Reports on Progress in Physics* 59.9 (Sept. 1, 1996), pp. 1071–1131. DOI: 10.1088/0034-4885/59/9/002.

[3] George Neilson. *Magnetic Fusion Energy*. 1st edition. Woodhead Publishing, 2016. URL: https://learning.oreilly.com/library/view/-/9780081003268/?ar.

[4] David H Crandall. *Review of the National Ignition Campaign*. July 16, 2012. URL: https://fire.pppl.gov/NIF_NIC_report_rev5_koonin_2012.pdf.

[5] John Wesson. *The Science of JET*. Culham Centre for Fusion Energy, 1999, p. 184. URL: https://www.euro-fusion.org/fileadmin/user_upload/Archive/wp-content/uploads/2012/01/the-science-of-jet-2000.pdf.

[6] Bernard Bigot. "ITER Construction and Manufacturing Progress toward First Plasma". In: *Fusion Engineering and Design* 146 (Sept. 2019), pp. 124–129. DOI: 10.1016/j.fusengdes.2018.11.052.

[7] Jeff Tollefson. "US Achieves Laser-Fusion Record: What It Means for Nuclear-Weapons Research". In: *Nature* 597.7875 (Aug. 27, 2021), pp. 163–164. DOI: 10.1038/d41586-021-02338-4.

[8] Editorial. "Fusion News Ignites Optimism". In: *Nature Photonics* 15.10 (Oct. 2021), pp. 713–713. DOI: 10.1038/s41566-021-00890-z.

[9] M. Roth et al. "Fast Ignition by Intense Laser-Accelerated Proton Beams". In: *Physical Review Letters* 86.3 (Jan. 15, 2001), pp. 436–439. DOI: 10.1103/PhysRevLett.86.436.

[10] M. Tabak et al. "Review of Progress in Fast Ignition". In: *Physics of Plasmas* 12.5 (May 2005), p. 057305. DOI: 10.1063/1.1871246.

[11] Ashutosh Sharma and Christos Kamperidis. "High Energy Proton Micro-Bunches from a Laser Plasma Accelerator". In: *Scientific Reports* 9.1 (Sept. 25, 2019), p. 13840. DOI: 10.1038/s41598-019-50348-0.

[12] B. T. Spiers et al. "Whole-Beam Self-Focusing in Fusion-Relevant Plasma". In: *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 379.2189 (Dec. 7, 2020), p. 20200159. DOI: 10.1098/rsta.2020.0159.

[13] Ben T. Spiers et al. "Enabling the Realisation of Proton Tomography". Mar. 20, 2021. arXiv: 2103.11268v3 [physics.plasm-ph].

[14] Ben T. Spiers et al. "Methods for Extremely Sparse-Angle Proton Tomography". Mar. 20, 2021. arXiv: 2103.11266v2 [physics.plasm-ph].

[15] A. A. Vlasov. “THE VIBRATIONAL PROPERTIES OF AN ELECTRON GAS”. In: *Soviet Physics Uspekhi* 10.6 (June 30, 1968), p. 721. DOI: 10.1070/PU1968v010n06ABEH003709.

[16] David Pines and David Bohm. “A Collective Description of Electron Interactions: II. Collective vs Individual Particle Aspects of the Interactions”. In: *Physical Review* 85.2 (Jan. 15, 1952), pp. 338–353. DOI: 10.1103/PhysRev.85.338.

[17] J. P. Hans Goedbloed and Stefaan Poedts. *Principles of Magnetohydrodynamics: With Applications to Laboratory and Astrophysical Plasmas*. Cambridge: Cambridge University Press, 2004. DOI: 10.1017/CBO9780511616945.

[18] Y. Matsumoto et al. “Interferometric Electron Density Measurements of High-density Plasmas”. In: *Journal of Applied Physics* 52.5 (May 1981), pp. 3259–3263. DOI: 10.1063/1.329143.

[19] E. A. McLean and S. A. Ramsden. “Optical Interferometric and Spectroscopic Measurements of Electron Density in a Plasma”. In: *Physical Review* 140 (Nov. 15, 1965), A1122–A1129. DOI: 10.1103/PhysRev.140.A1122.

[20] E Siminos et al. “Modeling Ultrafast Shadowgraphy in Laser-Plasma Interaction Experiments”. In: *Plasma Physics and Controlled Fusion* 58.6 (June 1, 2016), p. 065004. DOI: 10.1088/0741-3335/58/6/065004.

[21] Muhammad Firmansyah Kasim et al. “Quantitative Shadowgraphy and Proton Radiography for Large Intensity Modulations”. In: *Physical Review E* 95.2 (Feb. 16, 2017), p. 023306. DOI: 10.1103/PhysRevE.95.023306.

[22] D. Haberberger et al. “Measurements of Electron Density Profiles Using an Angular Filter Refractometer”. In: *Physics of Plasmas* 21.5 (May 2014), p. 056304. DOI: 10.1063/1.4873899.

[23] P. Angland et al. “Angular Filter Refractometry Analysis Using Simulated Annealing”. In: *Review of Scientific Instruments* 88.10 (Oct. 2017), p. 103510. DOI: 10.1063/1.4991511.

[24] D.M. Chambers et al. “Detailed Hydrodynamic and X-ray Spectroscopic Analysis of a Laser-Produced Rapidly-Expanding Aluminium Plasma”. In: *Journal of Quantitative Spectroscopy and Radiative Transfer* 71.2-6 (Oct. 2001), pp. 237–247. DOI: 10.1016/S0022-4073(01)00071-1.

[25] M. M. Notley et al. “Development of Time Resolved X-Ray Spectroscopy in High Intensity Laser-Plasma Interactions”. In: *Review of Scientific Instruments* 77.10 (Oct. 1, 2006). DOI: 10.1063/1.2222340.

[26] W. S. Porter and E. M. Bock. “Faraday Effect in a Plasma”. In: *American Journal of Physics* 33.12 (Dec. 1965), pp. 1070–1073. DOI: 10.1119/1.1971154.

[27] A. J. Mackinnon et al. “Proton Radiography of a Laser-Driven Implosion”. In: *Phys. Rev. Lett.* 97.4 (July 2006), p. 045001. DOI: 10.1103/PhysRevLett.97.045001.

[28] G. Sarri et al. “The Application of Laser-Driven Proton Beams to the Radiography of Intense Laser–Hohlraum Interactions”. In: *New Journal of Physics* 12.4 (Apr. 2010), p. 045006. DOI: 10.1088/1367-2630/12/4/045006.

[29] Alex Savin. "Modelling Laser-Plasma Interactions for the next Generation of High-Power Laser Experiments". University of Oxford, 2019. URL: https://ora.ox.ac.uk/objects/uuid:88d7de4a-9878-442a-a737-1f37a437c585.

[30] B. T. Spiers et al. "Methods for Extremely Sparse-Angle Proton Tomography". In: *Physical Review E* 104.4 (Oct. 7, 2021), p. 045201. DOI: 10.1103/PhysRevE.104.045201.

[31] R. W. Paddock et al. "One-Dimensional Hydrodynamic Simulations of Low Convergence Ratio Direct-Drive Inertial Confinement Fusion Implosions". In: *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 379.2189 (Jan. 25, 2021), p. 20200224. DOI: 10.1098/rsta.2020.0224.

[32] R. Aboushelbaya et al. "Measuring the Orbital Angular Momentum of High-Power Laser Pulses". In: *Physics of Plasmas* 27.5 (May 1, 2020), p. 053107. DOI: 10.1063/5.0005140.

[33] M. W. Mayr et al. "Nonlinear Wakefields and Electron Injection in Cluster Plasma". In: *Physical Review Accelerators and Beams* 23.9 (Sept. 2, 2020), p. 093501. DOI: 10.1103/PhysRevAccelBeams.23.093501.

[34] M. W. Mayr et al. "Wakefields in a Cluster Plasma". In: *Physical Review Accelerators and Beams* 22.11 (Nov. 15, 2019), p. 113501. DOI: 10.1103/PhysRevAccelBeams.22.113501.

[35] R. Aboushelbaya et al. "Orbital Angular Momentum Coupling in Elastic Photon-Photon Scattering". In: *Physical Review Letters* 123.11 (Sept. 13, 2019), p. 113604. DOI: 10.1103/PhysRevLett.123.113604.

[36] James D. Sadler et al. "Kinetic Simulations of Fusion Ignition with Hot-Spot Ablator Mix". In: *Physical Review E* 100.3 (Sept. 9, 2019), p. 033206. DOI: 10.1103/PhysRevE.100.033206.

[37] Alex F. Savin et al. "Energy Absorption in the Laser-QED Regime". In: *Scientific Reports* 9.1 (June 20, 2019), p. 8956. DOI: 10.1038/s41598-019-45536-x.

[38] James D. Sadler et al. "Advantages to a Diverging Raman Amplifier". In: *Communications Physics* 1.1 (May 17, 2018), pp. 1–7. DOI: 10.1038/s42005-018-0021-8.

[39] A. F. Savin et al. "Attosecond-Scale Absorption at Extreme Intensities". In: *Physics of Plasmas* 24.11 (Nov. 1, 2017), p. 113103. DOI: 10.1063/1.4989798.

[40] Dwight R Nicholson. *Introduction to Plasma Theory*. Wiley Series in Plasma Physics. Wiley, 1983.

[41] Peter Mulser and Dieter Bauer. *High Power Laser-Matter Interaction*. Vol. 238. Springer Tracts in Modern Physics. Berlin, Heidelberg: Springer Berlin Heidelberg, 2010. DOI: 10.1007/978-3-540-46065-7.

[42] Andrea Macchi. *A Superintense Laser-Plasma Interaction Theory Primer*. SpringerBriefs in Physics. Dordrecht: Springer Netherlands, 2013. DOI: 10.1007/978-94-007-6125-4.

[43] L. Allen et al. "Orbital Angular Momentum of Light and the Transformation of Laguerre-Gaussian Laser Modes". In: *Physical Review A* 45.11 (June 1, 1992), pp. 8185–8189. DOI: 10.1103/PhysRevA.45.8185.

[44] Nicolas Crouseilles, Pierre Degond, and Mohammed Lemou. "A Hybrid Kinetic–Fluid Model for Solving the Vlasov–BGK Equation". In: *Journal of Computational Physics* 203.2 (Mar. 1, 2005), pp. 572–601. DOI: 10.1016/j.jcp.2004.09.006.

[45] D. A. Frank-Kamenetskii. "The Two-Fluid Model". In: *Plasma: The Fourth State of Matter*. Ed. by D. A. Frank-Kamenetskii. Boston, MA: Springer US, 1972, pp. 59–64. DOI: 10.1007/978-1-4684-1896-5_17.

[46] Subrahmanyan Chandrasekhar. *An Introduction to the Study of Stellar Structure*. Courier Corporation, Jan. 1, 1957. 516 pp. Google Books: VwPLAgAAQBAJ.

[47] S. I. Blinnikov, N. V. Dunina-Barkovskaya, and D. K. Nadyozhin. "Equation of State of a Fermi Gas: Approximations for Various Degrees of Relativism and Degeneracy". In: *The Astrophysical Journal Supplement Series* 106 (Sept. 1, 1996), p. 171. DOI: 10.1086/192334.

[48] F. X. Timmes and F. Douglas Swesty. "The Accuracy, Consistency, and Speed of an Electron-Positron Equation of State Based on Table Interpolation of the Helmholtz Free Energy". In: *The Astrophysical Journal Supplement Series* 126.2 (Feb. 1, 2000), p. 501. DOI: 10.1086/313304.

[49] "Waves and Characteristics". In: *Principles of Magnetohydrodynamics: With Applications to Laboratory and Astrophysical Plasmas*. Ed. by J. P. Hans Goedbloed and Stefaan Poedts. Cambridge: Cambridge University Press, 2004, pp. 186–229. DOI: 10.1017/CBO9780511616945.006.

[50] A P L Robinson, A R Bell, and R J Kingham. "Fast Electron Transport and Ionization in a Target Irradiated by a High Power Laser". In: *Plasma Physics and Controlled Fusion* 48.8 (Aug. 1, 2006), pp. 1063–1076. DOI: 10.1088/0741-3335/48/8/002.

[51] D. Bauer. "Plasma Formation through Field Ionization in Intense Laser–Matter Interaction". In: *Laser and Particle Beams* 21.4 (Oct. 2003), pp. 489–495. DOI: 10.1017/S0263034603214026.

[52] Kun Wang et al. "The Equation of State and Ionization Equilibrium of Dense Aluminum Plasma with Conductivity Verification". In: *Physics of Plasmas* 22.6 (June 2015), p. 062709. DOI: 10.1063/1.4922909.

[53] Petros Tzeferacos. *FLASH User's Guide, Version 4.6.2*. FLASH User's Guide, Version 4.6.2. Dec. 23, 2021. URL: https://flash.rochester.edu/site/flashcode/user_support/flash4_ug_4p62.pdf.

[54] Jeff Colvin and Jon Larsen. *Extreme Physics: Properties and Behavior of Matter at Extreme Conditions*. Cambridge Core. Nov. 2013. DOI: 10.1017/CBO9781139095150.

[55] F. Brunel. "Not-so-Resonant, Resonant Absorption". In: *Physical Review Letters* 59.1 (July 6, 1987), pp. 52–55. DOI: 10.1103/PhysRevLett.59.52.

[56] D. Bauer, P. Mulser, and W. -H. Steeb. "Relativistic Ponderomotive Force, Uphill Acceleration, and Transition to Chaos". In: *Physical Review Letters* 75.25 (Dec. 18, 1995), pp. 4622–4625. DOI: 10.1103/PhysRevLett.75.4622.

[57] A.G.R. Thomas et al. “A Review of Vlasov–Fokker–Planck Numerical Modeling of Inertial Confinement Fusion Plasma”. In: *Journal of Computational Physics* 231.3 (Feb. 2012), pp. 1051–1079. DOI: 10.1016/j.jcp.2011.09.028.

[58] R.J. Kingham and A.R. Bell. “An Implicit Vlasov–Fokker–Planck Code to Model Non-Local Electron Transport in 2-D with Magnetic Fields”. In: *Journal of Computational Physics* 194.1 (Feb. 2004), pp. 1–34. DOI: 10.1016/j.jcp.2003.08.017.

[59] N.J. Sircombe and T.D. Arber. “VALIS: A Split-Conservative Scheme for the Relativistic 2D Vlasov–Maxwell System”. In: *Journal of Computational Physics* 228.13 (July 2009), pp. 4773–4788. DOI: 10.1016/j.jcp.2009.03.029.

[60] Adam Tableman et al. “Kinetic Simulations - Oshun (Vlasov-Fokker-Planck) and PIC (Osiris) - Physics and Open Source Software In The UCLA PICKSE Initiative”. In: 2016 (Oct. 1, 2016), GP10.012. URL: https://ui.adsabs.harvard.edu/abs/2016APS..DPPG10012T.

[61] E.M. Epperlein, G.J. Rickard, and A.R. Bell. “A Code for the Solution of the Vlasov-Fokker-Planck Equation in 1-D or 2-D”. In: *Computer Physics Communications* 52.1 (Dec. 1988), pp. 7–13. DOI: 10.1016/0010-4655(88)90165-8.

[62] J. U. Brackbill. “On Energy and Momentum Conservation in Particle-in-Cell Plasma Simulation”. In: *Journal of Computational Physics* 317 (July 15, 2016), pp. 405–427. DOI: 10.1016/j.jcp.2016.04.050.

[63] Benjamin M. Cowan et al. “Generalized Algorithm for Control of Numerical Dispersion in Explicit Time-Domain Electromagnetic Simulations”. In: *Physical Review Special Topics - Accelerators and Beams* 16.4 (Apr. 4, 2013), p. 041303. DOI: 10.1103/PhysRevSTAB.16.041303.

[64] R. Lehe et al. “Numerical Growth of Emittance in Simulations of Laser-Wakefield Acceleration”. In: *Physical Review Special Topics - Accelerators and Beams* 16.2 (Feb. 28, 2013), p. 021301. DOI: 10.1103/PhysRevSTAB.16.021301.

[65] Jay P Boris and Ramy A Shanny. *Proceedings [of the] 4th Conference on Numerical Simulation of Plasmas*. Washington: Naval Research Laboratory, 1972.

[66] Hong Qin et al. “Why Is Boris Algorithm so Good?” In: *Physics of Plasmas* 20.8 (Aug. 2013), p. 084503. DOI: 10.1063/1.4818428.

[67] J.-L. Vay. “Simulation of Beams or Plasmas Crossing at Relativistic Velocity”. In: *Physics of Plasmas* 15.5 (May 2008), p. 056701. DOI: 10.1063/1.2837054.

[68] A. V. Higuera and J. R. Cary. “Structure-Preserving Second-Order Integration of Relativistic Charged Particle Trajectories in Electromagnetic Fields”. In: *Physics of Plasmas* 24.5 (May 2017), p. 052104. DOI: 10.1063/1.4979989.

[69] B. Ripperda et al. “A Comprehensive Comparison of Relativistic Particle Integrators”. In: *The Astrophysical Journal Supplement Series* 235.1 (Mar. 15, 2018), p. 21. DOI: 10.3847/1538-4365/aab114.

[70] T. Zh. Esirkepov. “Exact Charge Conservation Scheme for Particle-in-Cell Simulation with an Arbitrary Form-Factor”. In: *Computer Physics Communications* 135.2 (Apr. 1, 2001), pp. 144–153. DOI: 10.1016/S0010-4655(00)00228-9.

[71] J Derouillat et al. “Smilei: A Collaborative, Open-Source, Multi-Purpose Particle-in-Cell Code for Plasma Simulation”. In: *Computer Physics Communications* (2018), p. 23.

[72] T D Arber et al. “Contemporary Particle-in-Cell Approach to Laser-Plasma Modelling”. In: *Plasma Physics and Controlled Fusion* 57.11 (Nov. 1, 2015), p. 113001. DOI: 10.1088/0741-3335/57/11/113001.

[73] A. Pukhov. “Particle-In-Cell Codes for Plasma-based Particle Acceleration”. In: *CERN Yellow Reports* 1 (Feb. 16, 2016), pp. 181–181. DOI: 10.5170/CERN-2016-001.181.

[74] R. A. Fonseca et al. “OSIRIS: A Three-Dimensional, Fully Relativistic Particle in Cell Code for Modeling Plasma Based Accelerators”. In: *Computational Science — ICCS 2002*. Ed. by Peter M. A. Sloot et al. Lecture Notes in Computer Science. Berlin, Heidelberg: Springer, 2002, pp. 342–351. DOI: 10.1007/3-540-47789-6_36.

[75] B. Fryxell et al. “FLASH: An Adaptive Mesh Hydrodynamics Code for Modeling Astrophysical Thermonuclear Flashes”. In: *The Astrophysical Journal Supplement Series* 131.1 (Nov. 1, 2000), p. 273. DOI: 10.1086/317361.

[76] Anshu Dubey et al. “Extensible Component-Based Architecture for FLASH, a Massively Parallel, Multiphysics Simulation Code”. In: *Parallel Computing* 35.10 (Oct. 1, 2009), pp. 512–522. DOI: 10.1016/j.parco.2009.08.001.

[77] R. Cauble and W. Rozmus. “The Inverse Bremsstrahlung Absorption Coefficient in Collisional Plasmas”. In: *Physics of Fluids* 28.11 (1985), p. 3387. DOI: 10.1063/1.865338.

[78] Carlos A. Iglesias and Forrest J. Rogers. “Updated Opal Opacities”. In: *The Astrophysical Journal* 464 (June 1, 1996), p. 943. DOI: 10.1086/177381.

[79] J. C. Fernández et al. “Fast Ignition with Laser-Driven Proton and Ion Beams”. In: 54.5 (Apr. 2014), p. 054006. DOI: 10.1088/0029-5515/54/5/054006.

[80] International Atomic Energy Agency. *Regulatory Control of the Safety of Ion Radiotherapy Facilities: A Guide for Best Practice.* 2020.

[81] Pankaj Chaudhary et al. “Radiobiology Experiments With Ultra-high Dose Rate Laser-Driven Protons: Methodology and State-of-the-Art”. In: *Frontiers in Physics* 9 (2021). URL: https://www.frontiersin.org/article/10.3389/fphy.2021.624963.

[82] E. L. Clark et al. “Measurements of Energetic Proton Transport through Magnetized Plasma from Intense Laser Interactions with Solids”. In: *Physical Review Letters* 84.4 (Jan. 24, 2000), pp. 670–673. DOI: 10.1103/PhysRevLett.84.670.

[83] R. A. Snavely et al. “Intense High-Energy Proton Beams from Petawatt-Laser Irradiation of Solids”. In: *Physical Review Letters* 85.14 (Oct. 2, 2000), pp. 2945–2948. DOI: 10.1103/PhysRevLett.85.2945.

[84] A. Maksimchuk et al. “Forward Ion Acceleration in Thin Films Driven by a High-Intensity Laser”. In: *Physical Review Letters* 84.18 (May 1, 2000), pp. 4108–4111. DOI: 10.1103/PhysRevLett.84.4108.

[85] S C Wilks et al. “Energetic Proton Generation in Ultra-Intense Laser–Solid Interactions”. In: *Phys. Plasmas* 8.2 (2001), p. 9.

[86] Andrea Macchi, Marco Borghesi, and Matteo Passoni. “Ion Acceleration by Superintense Laser-Plasma Interaction”. In: *Reviews of Modern Physics* 85.2 (May 10, 2013), pp. 751–793. DOI: 10.1103/RevModPhys.85.751.

[87] S. C. Wilks et al. “Absorption of Ultra-Intense Laser Pulses”. In: *Physical Review Letters* 69.9 (Aug. 31, 1992), pp. 1383–1386. DOI: 10.1103/PhysRevLett.69.1383.

[88] Maurizio Lontano and Matteo Passoni. “Electrostatic Field Distribution at the Sharp Interface between High Density Matter and Vacuum”. In: *Physics of Plasmas* 13.4 (Apr. 2006), p. 042102. DOI: 10.1063/1.2184067.

[89] M. Passoni and M. Lontano. “Theory of Light-Ion Acceleration Driven by a Strong Charge Separation”. In: *Physical Review Letters* 101.11 (Sept. 10, 2008), p. 115001. DOI: 10.1103/PhysRevLett.101.115001.

[90] M. Passoni and M. Lontano. “One-Dimensional Model of the Electrostatic Ion Acceleration in the Ultraintense Laser–Solid Interaction”. In: *Laser and Particle Beams* 22.2 (June 2004), pp. 163–169. DOI: 10.1017/S026303460422211X.

[91] M. Nishiuchi et al. “The Laser Proton Acceleration in the Strong Charge Separation Regime”. In: *Physics Letters A* 357.4 (Sept. 18, 2006), pp. 339–344. DOI: 10.1016/j.physleta.2006.04.053.

[92] M. Passoni et al. “Charge Separation Effects in Solid Targets and Ion Acceleration with a Two-Temperature Electron Distribution”. In: *Physical Review E* 69.2 (Feb. 27, 2004), p. 026411. DOI: 10.1103/PhysRevE.69.026411.

[93] A V Gurevich, L V Pariiskaya, and L P Pitaevskii. “SELF-SIMILAR MOTION OF RAREFIED PLASMA”. In: (), p. 6.

[94] P. Mora. “Plasma Expansion into a Vacuum”. In: *Physical Review Letters* 90.18 (May 7, 2003), p. 185002. DOI: 10.1103/PhysRevLett.90.185002.

[95] V. T. Tikhonchuk et al. “Ion Acceleration in Short-Laser-Pulse Interaction with Solid Foils”. In: *Plasma Physics and Controlled Fusion* 47 (Nov. 2005), B869–B877. DOI: 10.1088/0741-3335/47/12B/S69.

[96] M. Manclossi et al. “Study of Ultraintense Laser-Produced Fast-Electron Propagation and Filamentation in Insulator and Metal Foil Targets by Optical Emission Diagnostics”. In: *Physical Review Letters* 96.12 (Mar. 27, 2006), p. 125002. DOI: 10.1103/PhysRevLett.96.125002.

[97] M. Borghesi et al. “Multi-MeV Proton Source Investigations in Ultraintense Laser-Foil Interactions”. In: *Physical Review Letters* 92.5 (Feb. 5, 2004), p. 055003. DOI: 10.1103/PhysRevLett.92.055003.

[98] T. E. Cowan et al. “Ultralow Emittance, Multi-MeV Proton Beams from a Laser Virtual-Cathode Plasma Accelerator”. In: *Physical Review Letters* 92.20 (May 19, 2004), p. 204801. DOI: 10.1103/PhysRevLett.92.204801.

[99] Hans-Peter Schlenvoigt et al. *Laser-Based Particle Acceleration*. IntechOpen, Feb. 1, 2010. DOI: 10.5772/7965.

[100] A. J. Mackinnon et al. “Enhancement of Proton Acceleration by Hot-Electron Recirculation in Thin Foils Irradiated by Ultraintense Laser Pulses”. In: *Physical Review Letters* 88.21 (May 14, 2002), p. 215006. DOI: 10.1103/PhysRevLett.88.215006.

[101] Y. Sentoku et al. “High Energy Proton Acceleration in Interaction of Short Laser Pulse with Dense Plasma Target”. In: *Physics of Plasmas* 10.5 (May 2003), pp. 2009–2015. DOI: 10.1063/1.1556298.

[102] A. J. Mackinnon et al. “Effect of Plasma Scale Length on Multi-MeV Proton Production by Intense Laser Pulses”. In: *Physical Review Letters* 86.9 (Feb. 26, 2001), pp. 1769–1772. DOI: 10.1103/PhysRevLett.86.1769.

[103] Sang Yun Shin, Sang Yoon Park, and Sang June Hahn. “Investigation of Light-sail and Hole-boring Radiation Pressure Accelerations upon the Interaction of Ultra-intense Laser Pulses with Thin Targets”. In: *Journal of the Korean Physical Society* 75.12 (Dec. 1, 2019), pp. 968–977. DOI: 10.3938/jkps.75.968.

[104] A P L Robinson et al. “Relativistically Correct Hole-Boring and Ion Acceleration by Circularly Polarized Laser Pulses”. In: *Plasma Physics and Controlled Fusion* 51.2 (Feb. 1, 2009), p. 024004. DOI: 10.1088/0741-3335/51/2/024004.

[105] T. Esirkepov et al. “Highly Efficient Relativistic-Ion Generation in the Laser-Piston Regime”. In: *Physical Review Letters* 92.17 (Apr. 27, 2004), p. 175003. DOI: 10.1103/PhysRevLett.92.175003.

[106] A P L Robinson et al. “Radiation Pressure Acceleration of Thin Foils with Circularly Polarized Laser Pulses”. In: *New Journal of Physics* 10.1 (Jan. 21, 2008), p. 013021. DOI: 10.1088/1367-2630/10/1/013021.

[107] Charlotte A. J. Palmer et al. “Monoenergetic Proton Beams Accelerated by a Radiation Pressure Driven Shock”. In: *Physical Review Letters* 106.1 (Jan. 4, 2011), p. 014801. DOI: 10.1103/PhysRevLett.106.014801.

[108] Z. Najmudin et al. “Observation of Impurity Free Monoenergetic Proton Beams from the Interaction of a CO2 Laser with a Gaseous Target”. In: *Physics of Plasmas* 18.5 (May 2011), p. 056705. DOI: 10.1063/1.3562926.

[109] Philip Martin. “Schemes of Ion Acceleration Employing High Energy, Petawatt Laser Pulses”. Queen’s University, Belfast, 2020. URL: https://pure.qub.ac.uk/en/studentTheses/schemes-of-ion-acceleration-employing-high-energy-petawatt-laser--2.

[110] J. F. L. Simmons and C. R. McInnes. “Was Marx Right? Or How Efficient Are Laser Driven Interstellar Spacecraft?” In: *American Journal of Physics* 61.3 (Mar. 1993), pp. 205–207. DOI: 10.1119/1.17291.

[111] C. A. J. Palmer et al. “Rayleigh-Taylor Instability of an Ultrathin Foil Accelerated by the Radiation Pressure of an Intense Laser”. In: *Physical Review Letters* 108.22 (May 29, 2012), p. 225002. DOI: 10.1103/PhysRevLett.108.225002.

[112] Andrea Macchi, Amritpal Singh Nindrayog, and Francesco Pegoraro. “Solitary versus Shock Wave Acceleration in Laser-Plasma Interactions”. In: *Physical Review E* 85.4 (Apr. 2, 2012), p. 046402. DOI: 10.1103/PhysRevE.85.046402.

[113] F. Fiuza et al. “Laser-Driven Shock Acceleration of Monoenergetic Ion Beams”. In: *Physical Review Letters* 109.21 (Nov. 20, 2012), p. 215001. DOI: 10.1103/PhysRevLett.109.215001.

[114] Dan Haberberger et al. “Collisionless Shocks in Laser-Produced Plasma Generate Monoenergetic High-Energy Proton Beams”. In: *Nature Physics* 8.1 (Jan. 2012), pp. 95–99. DOI: 10.1038/nphys2130.

[115] R. A. Cairns et al. “Laminar Shocks in High Power Laser Plasma Interactions”. In: *Physics of Plasmas* 21.2 (Feb. 2014), p. 022112. DOI: 10.1063/1.4864328.

[116] R A Cairns et al. “Weak Collisionless Shocks in Laser-Plasmas”. In: *Plasma Physics and Controlled Fusion* 57.4 (Apr. 1, 2015), p. 044008. DOI: 10.1088/0741-3335/57/4/044008.

[117] A. Pak et al. “Collisionless Shock Acceleration of Narrow Energy Spread Ion Beams from Mixed Species Plasmas Using 1 μ m Lasers”. In: *Physical Review Accelerators and Beams* 21.10 (Oct. 25, 2018), p. 103401. DOI: 10.1103/PhysRevAccelBeams.21.103401.

[118] A. Stockem et al. “Relativistic Generalization of Formation and Ion-Reflection Conditions in Electrostatic Shocks”. In: *Physical Review E* 87.4 (Apr. 30, 2013), p. 043116. DOI: 10.1103/PhysRevE.87.043116.

[119] R. Z. Sagdeev. “COOPERATIVE PHENOMENA AND SHOCK WAVES IN COLLISIONLESS PLASMAS”. In: *Reviews of Plasma Physics*. Ed. by M. A. Leontovich. CONSULTANTS BUREAU · NEW YORK, Dec. 1965, p. 37.

[120] M. A. Hellberg et al. “Acoustic Solitary Waves and Sagdeev Potential Triple Roots”. In: NEW FRONTIERS IN ADVANCED PLASMA PHYSICS. Trieste, (Italy), 2010, pp. 50–60. DOI: 10.1063/1.3533193.

[121] P. Martin et al. “Absolute Calibration of Fujifilm BAS-TR Image Plate Response to Laser Driven Protons up to 40 MeV”. In: *Review of Scientific Instruments* 93.5 (May 2022), p. 053303. DOI: 10.1063/5.0089402.

[122] J. Fuchs et al. “Laser-Driven Proton Scaling Laws and New Paths towards Energy Increase”. In: *Nature Physics* 2.1 (Jan. 2006), pp. 48–54. DOI: 10.1038/nphys199.

[123] Andrea Macchi et al. “Laser Acceleration of Ion Bunches at the Front Surface of Overdense Plasmas”. In: *Physical Review Letters* 94.16 (Apr. 29, 2005), p. 165003. DOI: 10.1103/PhysRevLett.94.165003.

[124] S. N. Chen et al. “Experimental Evidence for the Enhanced and Reduced Stopping Regimes for Protons Propagating through Hot Plasmas”. In: *Scientific Reports* 8.1 (Oct. 1, 2018), p. 14586. DOI: 10.1038/s41598-018-32726-2.

[125] M. Borghesi et al. “Proton Imaging: A Diagnostic for Inertial Confinement Fusion/Fast Ignitor Studies”. In: *Plasma Physics and Controlled Fusion* 43 (Nov. 2001), A267–A276. DOI: 10.1088/0741-3335/43/12A/320.

[126] A. J. Mackinnon et al. “Proton Radiography as an Electromagnetic Field and Density Perturbation Diagnostic (Invited)”. In: *Review of Scientific Instruments* 75.10 (Oct. 1, 2004), pp. 3531–3536. DOI: 10.1063/1.1788893.

[127] N. L. Kugland et al. "Visualizing Electromagnetic Fields in Laser-Produced Counter-Streaming Plasma Experiments for Collisionless Shock Laboratory Astrophysics". In: *Physics of Plasmas* 20.5 (May 1, 2013), p. 056313. DOI: 10.1063/1.4804548.

[128] T. Morita et al. "Proton Imaging of an Electrostatic Field Structure Formed in Laser-Produced Counter-Streaming Plasmas". In: *Journal of Physics: Conference Series* 688 (Mar. 2016), p. 012071. DOI: 10.1088/1742-6596/688/1/012071.

[129] Hye-Sook Park et al. "Laboratory Astrophysical Collisionless Shock Experiments on Omega and NIF". In: *Journal of Physics: Conference Series* 688 (Mar. 2016), p. 012084. DOI: 10.1088/1742-6596/688/1/012084.

[130] C. M. Huntington et al. "Magnetic Field Production via the Weibel Instability in Interpenetrating Plasma Flows". In: *Physics of Plasmas* 24.4 (Apr. 1, 2017), p. 041410. DOI: 10.1063/1.4982044.

[131] P. Tzeferacos et al. "Laboratory Evidence of Dynamo Amplification of Magnetic Fields in a Turbulent Plasma". In: *Nature Communications* 9.1 (Feb. 9, 2018), p. 591. DOI: 10.1038/s41467-018-02953-2.

[132] C. A. J. Palmer et al. "Field Reconstruction from Proton Radiography of Intense Laser Driven Magnetic Reconnection". In: *Physics of Plasmas* 26.8 (Aug. 1, 2019), p. 083109. DOI: 10.1063/1.5092733.

[133] E. R. Tubman et al. "Observations of Pressure Anisotropy Effects within Semi-Collisional Magnetized Plasma Bubbles". In: *Nature Communications* 12.1 (Jan. 12, 2021), p. 334. DOI: 10.1038/s41467-020-20387-7.

[134] J. R. Rygg et al. "Proton Radiography of Inertial Fusion Implosions". In: *Science* 319.5867 (Feb. 29, 2008), pp. 1223–1225. DOI: 10.1126/science.1152640.

[135] C. K. Li et al. "Proton Radiography of Dynamic Electric and Magnetic Fields in Laser-Produced High-Energy-Density Plasmas". In: *Physics of Plasmas* 16.5 (May 1, 2009), p. 056304. DOI: 10.1063/1.3096781.

[136] C. K. Li et al. "Observations of Electromagnetic Fields and Plasma Flow in Hohlraums with Proton Radiography". In: *Physical Review Letters* 102.20 (May 18, 2009), p. 205001. DOI: 10.1103/PhysRevLett.102.205001.

[137] C. K. Li et al. "Charged-Particle Probing of X-ray–Driven Inertial-Fusion Implosions". In: *Science* 327.5970 (Mar. 5, 2010), pp. 1231–1235. DOI: 10.1126/science.1185747.

[138] A. Rigby et al. "Implementation of a Faraday Rotation Diagnostic at the OMEGA Laser Facility". In: *High Power Laser Science and Engineering* 6 (2018/ed). DOI: 10.1017/hpl.2018.42.

[139] M. Borghesi et al. "Electric Field Detection in Laser-Plasma Interaction Experiments via the Proton Imaging Technique". In: *Physics of Plasmas* 9.5 (May 2002), pp. 2214–2220. DOI: 10.1063/1.1459457.

[140] L. Romagnani et al. "Observation of Collisionless Shocks in Laser-Plasma Experiments". In: *Physical Review Letters* 101.2 (July 10, 2008), p. 025004. DOI: 10.1103/PhysRevLett.101.025004.

[141] H. Ahmed et al. "Time-Resolved Characterization of the Formation of a Collisionless Shock". In: *Physical Review Letters* 110.20 (May 14, 2013), p. 205001. DOI: 10.1103/PhysRevLett.110.205001.

[142] N. L. Kugland et al. "Invited Article: Relation between Electric and Magnetic Field Structures and Their Proton-Beam Images". In: *Review of Scientific Instruments* 83.10 (2012), p. 101301. DOI: 10.1063/1.4750234.

[143] A. F. A. Bott et al. "Proton Imaging of Stochastic Magnetic Fields". In: *Journal of Plasma Physics* 83.6 (2017), p. 905830614. DOI: 10.1017/S0022377817000939.

[144] M. F. Kasim et al. "Retrieving Fields from Proton Radiography without Source Profiles". In: *Physical Review E* 100.3 (Sept. 24, 2019), p. 033208. DOI: 10.1103/PhysRevE.100.033208.

[145] Nicholas F. Y. Chen et al. "Machine Learning Applied to Proton Radiography of High-Energy-Density Plasmas". In: *Physical Review E* 95.4 (Apr. 17, 2017), p. 043305. DOI: 10.1103/PhysRevE.95.043305.

[146] C. A. Cecchetti et al. "Magnetic Field Measurements in Laser-Produced Plasmas via Proton Deflectometry". In: *Physics of Plasmas* 16.4 (Apr. 2009), p. 043102. DOI: 10.1063/1.3097899.

[147] C. K. Li et al. "Measuring E and B Fields in Laser-Produced Plasmas with Monoenergetic Proton Radiography". In: *Physical Review Letters* 97.13 (Sept. 28, 2006), p. 135003. DOI: 10.1103/PhysRevLett.97.135003.

[148] J. R. Rygg et al. "Note: A Monoenergetic Proton Backlighter for the National Ignition Facility". In: *Review of Scientific Instruments* 86.11 (Nov. 2015), p. 116104. DOI: 10.1063/1.4935581.

[149] Mohamed M. Sulman, J. F. Williams, and Robert D. Russell. "An Efficient Approach for the Numerical Solution of the Monge–Ampère Equation". In: *Applied Numerical Mathematics* 61.3 (Mar. 1, 2011), pp. 298–307. DOI: 10.1016/j.apnum.2010.10.006.

[150] A. G. Smyth et al. "Magnetic Field Generation during Intense Laser Channelling in Underdense Plasma". In: *Physics of Plasmas* 23.6 (June 1, 2016), p. 063121. DOI: 10.1063/1.4953547.

[151] Bao Du et al. "Separating the Contributions of Electric and Magnetic Fields in Deflecting the Probes in Proton Radiography with Multiple Proton Energies". In: *Matter and Radiation at Extremes* 6.3 (Apr. 8, 2021), p. 035903. DOI: 10.1063/5.0033834.

[152] Carlo Graziani et al. "Inferring Morphology and Strength of Magnetic Fields from Proton Radiographs". In: *Review of Scientific Instruments* 88.12 (2017), p. 123507. DOI: 10.1063/1.5013029.

[153] Alemayehu Bogale, Scott Feister, and Petros Tzeferacos. *PRaLine Code (Proton Radiography Linear Reconstruction, the "Reconstruction" Is Silent)*. Flash Center for Computational Science, Nov. 13, 2017. URL: https://github.com/flash-center/PRaLine.

[154] Jordan Laune, Scott Feister, and Petros Tzeferacos. *PROBLEM Solver (PROton-imaged B-field nonLinear Extraction Module)*. Flash Center for Computational Science, Oct. 29, 2017. URL: https://github.com/flash-center/PROBLEM.

[155] P. A. Browne et al. "Nonlinear Solution Techniques for Solving a Monge-Ampère Equation for Redistribution of a Mesh". Sept. 30, 2016. arXiv: 1609.09646 [math].

[156] T. Matsuoka et al. "Focus Optimization of Relativistic Self-Focusing for Anomalous Laser Penetration into Overdense Plasmas (Super-Penetration)". In: *Plasma Physics and Controlled Fusion* 50.10 (Sept. 2008), p. 105011. DOI: 10.1088/0741-3335/50/10/105011.

[157] H. Habara et al. "Efficient Propagation of Ultra-Intense Laser Beam in Dense Plasma". In: *Plasma Physics and Controlled Fusion* 57.6 (Apr. 2015), p. 064005. DOI: 10.1088/0741-3335/57/6/064005.

[158] R. Kodama et al. "Fast Ignitor Research at the Institute of Laser Engineering, Osaka University". In: *Physics of Plasmas* 8.5 (2001), pp. 2268–2274. DOI: 10.1063/1.1352598.

[159] K. A. Tanaka et al. "Studies of Ultra-Intense Laser Plasma Interactions for Fast Ignition". In: *Physics of Plasmas* 7.5 (2000), pp. 2014–2022. DOI: 10.1063/1.874023.

[160] Stefan Parker et al. "400 TW Operation of Orion at Ultra-High Contrast". In: *High Power Laser Science and Engineering* 6 (2018/ed). DOI: 10.1017/hpl.2018.44.

[161] L. Biermann. "Über Den Ursprung Der Magnetfelder Auf Sternen Und Im Interstellaren Raum (Miteinem Anhang von A. Schlüter)". In: *Zeitschrift Naturforschung Teil A* 5 (Jan. 1, 1950), p. 65. URL: https://ui.adsabs.harvard.edu/abs/1950ZNatA...5...65B.

[162] Carlo Graziani et al. "THE BIERMANN CATASTROPHE IN NUMERICAL MAGNETOHYDRODYNAMICS". In: *The Astrophysical Journal* 802.1 (Mar. 19, 2015), p. 43. DOI: 10.1088/0004-637X/802/1/43.

[163] "The Numerical Flux Function". In: *Numerical Solutions of Partial Differential Equations*. Ed. by Silvia Bertoluzza et al. Advanced Courses in Mathematics - CRM Barcelona. Basel: Birkhäuser, 2009, pp. 89–95. DOI: 10.1007/978-3-7643-8940-6_7.

[164] B. van der Holst et al. "CRASH: A BLOCK-ADAPTIVE-MESH CODE FOR RADIATIVE SHOCK HYDRODYNAMICS—IMPLEMENTATION AND VERIFICATION". In: *The Astrophysical Journal Supplement Series* 194.2 (May 2011), p. 23. DOI: 10.1088/0067-0049/194/2/23.

[165] P. L. Volegov et al. "On Three-Dimensional Reconstruction of a Neutron/X-ray Source from Very Few Two-Dimensional Projections". In: *Journal of Applied Physics* 118.20 (Nov. 28, 2015), p. 205903. DOI: 10.1063/1.4936319.

[166] P. L. Volegov et al. "Three-Dimensional Reconstruction of Neutron, Gamma-Ray, and X-ray Sources Using Spherical Harmonic Decomposition". In: *Journal of Applied Physics* 122.17 (Nov. 7, 2017), p. 175901. DOI: 10.1063/1.4986652.

[167] P. L. Volegov et al. "Density Determination of the Thermonuclear Fuel Region in Inertial Confinement Fusion Implosions". In: *Journal of Applied Physics* 127.8 (Feb. 24, 2020), p. 083301. DOI: 10.1063/1.5123751.

[168] P. L. Volegov et al. "Three-Dimensional Reconstruction of Neutron, Gamma-Ray, and X-ray Sources Using a Cylindrical-Harmonics Expansion". In: *Review of Scientific Instruments* 92.3 (Mar. 1, 2021), p. 033508. DOI: 10.1063/5.0042860.

[169] A. M. Cormack. "Representation of a Function by Its Line Integrals, with Some Radiological Applications". In: *Journal of Applied Physics* 34.9 (Sept. 1, 1963), pp. 2722–2727. DOI: 10.1063/1.1729798.

[170] A. M. Cormack. "Representation of a Function by Its Line Integrals, with Some Radiological Applications. II". In: *Journal of Applied Physics* 35.10 (Oct. 1, 1964), pp. 2908–2913. DOI: 10.1063/1.1713127.

[171] P. Tomassini and A. Giulietti. "A Generalization of Abel Inversion to Non-Axisymmetric Density Distribution". In: *Optics Communications* 199.1 (Nov. 15, 2001), pp. 143–148. DOI: 10.1016/S0030-4018(01)01575-9.

[172] L. A. Shepp and B. F. Logan. "The Fourier Reconstruction of a Head Section". In: *IEEE Transactions on Nuclear Science* 21.3 (June 1974), pp. 21–43. DOI: 10.1109/TNS.1974.6499235.

[173] W.E. Higgins and D.C. Munson. "A Hankel Transform Approach to Tomographic Image Reconstruction". In: *IEEE Transactions on Medical Imaging* 7.1 (Mar. 1988), pp. 59–72. DOI: 10.1109/42.3929.

[174] D.E. Dudgeon, D.D. E, and R.M. Mersereau. *Multidimensional Digital Signal Processing*. Prentice-Hall Signal Processing Series. Prentice-Hall, 1984. URL: https://books.google.co.uk/books?id=-4woAQAAMAAJ.

[175] Zhou Wang et al. "Image Quality Assessment: From Error Visibility to Structural Similarity". In: *IEEE Transactions on Image Processing* 13.4 (Apr. 2004), pp. 600–612. DOI: 10.1109/TIP.2003.819861.

[176] Richard Gordon, Robert Bender, and Gabor T. Herman. "Algebraic Reconstruction Techniques (ART) for Three-Dimensional Electron Microscopy and X-ray Photography". In: *Journal of Theoretical Biology* 29.3 (Dec. 1, 1970), pp. 471–481. DOI: 10.1016/0022-5193(70)90109-8.

[177] Avinash Kak and Malcolm Slaney. "7. Algebraic Reconstruction Algorithms". In: *Principles of Computerized Tomographic Imaging*. Classics in Applied Mathematics. Society for Industrial and Applied Mathematics, Jan. 1, 2001, pp. 275–296. DOI: 10.1137/1.9780898719277.ch7.

[178] Ludwig Ritschl et al. "Improved Total Variation-Based CT Image Reconstruction Applied to Clinical Data". In: *Physics in Medicine & Biology* 56.6 (Feb. 16, 2011), p. 1545. DOI: 10.1088/0031-9155/56/6/003.

[179] Ziling Wu et al. "Robust X-ray Sparse-view Phase Tomography via Hierarchical Synthesis Convolutional Neural Networks". Jan. 30, 2019. arXiv: 1901.10644 [cs.CV].

[180] W. Higgins and D. Munson. "An Algorithm for Computing General Integer-Order Hankel Transforms". In: *IEEE Transactions on Acoustics, Speech, and Signal Processing* 35.1 (Jan. 1987), pp. 86–97. DOI: 10.1109/TASSP.1987.1165032.

[181] E. Boella et al. “Collisionless Shock Acceleration in the Corona of an Inertial Confinement Fusion Pellet with Possible Application to Ion Fast Ignition”. In: *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 379.2189 (Jan. 25, 2021). DOI: 10.1098/rsta.2020.0039.

[182] P. A. Norreys et al. “Preparations for a European R&D Roadmap for an Inertial Fusion Demo Reactor”. In: *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 379.2189 (Jan. 25, 2021), p. 20200005. DOI: 10.1098/rsta.2020.0005.

[183] N. Ratan et al. “Dense Plasma Heating by Crossing Relativistic Electron Beams”. In: *Physical Review E* 95.1 (Jan. 23, 2017), p. 013211. DOI: 10.1103/PhysRevE.95.013211.