

University of Warwick institutional repository: <http://go.warwick.ac.uk/wrap>

A Thesis Submitted for the Degree of PhD at the University of Warwick

<http://go.warwick.ac.uk/wrap/77583>

This thesis is made available online and is protected by original copyright.

Please scroll down to view the document itself.

Please refer to the repository record for this item for information to help you to cite it. Our policy information is available from the repository home page.



**Short-Pulse Laser Interactions with High Density
Plasma**

by

Martin Ramsay

Thesis

Submitted to the University of Warwick

for the degree of

Doctor of Philosophy

Department of Physics

November 2015

THE UNIVERSITY OF
WARWICK

Contents

| | |
|--|------------|
| List of Tables | v |
| List of Figures | vi |
| Acknowledgments | ix |
| Declarations | x |
| Abstract | xi |
| Abbreviations | xii |
| Chapter 1 Introduction | 1 |
| 1.1 Overview | 2 |
| 1.2 Basic plasma physics | 3 |
| 1.2.1 Plasma equations of motion | 3 |
| 1.2.2 Ohm's law | 6 |
| 1.2.3 Debye length | 9 |
| 1.2.4 Plasma frequency | 10 |
| 1.2.5 Plasma dispersion relation | 11 |
| 1.2.6 Critical density | 13 |
| 1.3 High power lasers | 14 |
| 1.3.1 'Long-pulse' lasers | 14 |
| 1.3.2 'Short-pulse' lasers | 14 |
| 1.4 Short-pulse laser absorption | 15 |
| 1.4.1 Inverse bremsstrahlung | 16 |
| 1.4.2 Resonance absorption | 16 |
| 1.4.3 Vacuum heating | 17 |
| 1.4.4 Ponderomotive acceleration | 17 |
| 1.4.5 $\mathbf{J} \times \mathbf{B}$ force | 18 |

| | | |
|------------------|---|-----------|
| 1.4.6 | Other important processes | 18 |
| 1.5 | Energy transport | 19 |
| 1.5.1 | Return current generation | 19 |
| 1.5.2 | Ohmic heating | 19 |
| 1.5.3 | Filamentation | 20 |
| 1.5.4 | X-ray emission | 20 |
| 1.6 | Summary of short-pulse laser interactions | 21 |
| 1.7 | Short-pulse laser applications | 22 |
| 1.7.1 | Material heating | 22 |
| 1.7.2 | X-ray sources | 22 |
| 1.7.3 | Electron acceleration | 22 |
| 1.7.4 | Proton acceleration | 23 |
| 1.7.5 | Fast ignition | 23 |
| 1.8 | Overview of plasma modelling | 24 |
| 1.8.1 | Fluid models | 25 |
| 1.8.2 | Kinetic models | 26 |
| 1.8.3 | Hybrid models | 26 |
| 1.8.4 | Summary of plasma models | 27 |
| Chapter 2 | Overview of the EPOCH PIC code | 28 |
| 2.1 | Particle-in-cell codes | 28 |
| 2.1.1 | Explicit vs implicit schemes | 29 |
| 2.2 | Particle propagation | 30 |
| 2.2.1 | Particle shapes | 32 |
| 2.2.2 | Current density calculation | 33 |
| 2.3 | Field propagation | 34 |
| 2.4 | Boundary conditions | 37 |
| 2.5 | The laser package | 39 |
| 2.6 | The limits of PIC | 40 |
| 2.6.1 | Grid resolution | 41 |
| 2.6.2 | Particle resolution | 42 |
| 2.6.3 | Time-step constraints | 42 |
| 2.7 | PIC summary | 44 |
| 2.8 | Implementing short-range collisions in EPOCH | 44 |
| 2.8.1 | Introduction | 44 |
| 2.8.2 | Cumulative relativistic scattering of binary particle pairs | 45 |
| 2.8.3 | Limitations of the collision algorithm | 48 |

| | | |
|------------------|---|-----------|
| 2.8.4 | Test problems | 49 |
| Chapter 3 | An Ohmic field solver for high density PIC simulations | 53 |
| 3.1 | Introduction | 53 |
| 3.2 | The high density algorithm | 54 |
| 3.2.1 | Electric field calculation | 54 |
| 3.2.2 | Identifying particle species | 59 |
| 3.2.3 | Multiple electron/ion species | 61 |
| 3.2.4 | Magnetic field update | 61 |
| 3.2.5 | Particle position and velocity update | 61 |
| 3.3 | Implications for PIC modelling | 62 |
| Chapter 4 | Preliminary testing and comparisons | 63 |
| 4.1 | Introduction | 63 |
| 4.1.1 | THOR summary | 64 |
| 4.2 | Monoenergetic beam | 64 |
| 4.3 | Diffusion | 66 |
| 4.4 | Laser-plasma interaction | 68 |
| 4.4.1 | 1D | 68 |
| 4.4.2 | 2D | 69 |
| 4.5 | Run times and self-heating rates | 70 |
| 4.6 | Summary | 72 |
| Chapter 5 | The effect of high density transport phenomena on the TNSA process | 74 |
| 5.1 | Introduction | 74 |
| 5.2 | Simulation setup | 75 |
| 5.3 | EPOCH/EPOCH-C/EPOCH-H comparison | 77 |
| 5.3.1 | Proton energy spectra | 77 |
| 5.4 | Imprinting of fast electron transport effects | 79 |
| 5.4.1 | Control of proton beam filamentation via target morphology | 81 |
| 5.4.2 | Fast electron filamentation | 82 |
| 5.5 | Discussion | 86 |
| Chapter 6 | Return current heating of solid density targets | 88 |
| 6.1 | Modelling material heating | 88 |
| 6.2 | Simulation parameters | 89 |
| 6.2.1 | EPOCH-H | 89 |

| | | |
|-------------------|---|------------|
| 6.2.2 | THOR | 90 |
| 6.3 | Material heating comparison | 91 |
| 6.4 | Line emission comparison | 95 |
| 6.5 | Magnetic fields at layer interface | 98 |
| 6.6 | Summary and discussion | 99 |
| Chapter 7 | Summary | 102 |
| 7.1 | Background | 102 |
| 7.2 | The high density algorithm | 104 |
| 7.3 | Effect of high density transport phenomena on proton acceleration . . . | 106 |
| 7.4 | Heating of solid density matter | 108 |
| 7.5 | Future work | 108 |
| 7.6 | Implications for further short-pulse modelling | 109 |
| Appendix A | Derivation of Fokker-Planck equation | 111 |
| A.1 | The Landau equation | 113 |
| Appendix B | Derivation of radiating boundary conditions | 114 |
| Appendix C | Derivation of collision parameters | 117 |
| C.1 | Collision frequency | 117 |
| C.2 | Scattering angle | 120 |

List of Tables

| | | |
|-----|---|----|
| 3.1 | Tabulated values of the Braginskii electrical resistivity coefficient. . . . | 56 |
| 4.1 | Computational expense and electron heating rate associated with various versions of EPOCH for the case of a simple, uniform system. | 70 |
| 5.1 | Tabulated values of the average temperatures in the range $15\ \mu\text{m} < x < 25\ \mu\text{m}$ from figure 5.3. | 79 |
| 5.2 | Values, for various target widths (w), of the characteristic distance (λ_f) between filaments in the spatial distribution of protons recorded passing a $40\ \mu\text{m}$ -wide plane located $25\ \mu\text{m}$ behind the rear surface of each target. | 82 |

List of Figures

| | | |
|------|--|----|
| 1.1 | The plasma dispersion relations (left) and refractive indices (right) for electromagnetic waves in a cold plasma | 12 |
| 1.2 | The reflectivity of a plasma as a function of its electron density | 13 |
| 1.3 | Diagrammatic representation of the fast ignition concept | 24 |
| 1.4 | Visual representation of the disparate time-scales involved in short-pulse laser interactions | 25 |
| 2.1 | Diagrammatic representation of the dependencies for explicit and implicit methods | 30 |
| 2.2 | Diagrammatic representation of zeroth-order particle weighting | 33 |
| 2.3 | Diagrammatic representation of first-order particle weighting | 34 |
| 2.4 | An example of the 3D Yee mesh for the FDTD method | 35 |
| 2.5 | Diagram to illustrate how particle and field variables are used to update each other during a time-step | 37 |
| 2.6 | The dispersion relation on a discrete mesh | 43 |
| 2.7 | Diagram to illustrate the scattering angles associated with binary particle collisions | 47 |
| 2.8 | Results of a collision algorithm test problem aiming to reproduce Spitzer conductivity | 50 |
| 2.9 | Results of a collision algorithm test problem aiming to reproduce thermal conduction | 51 |
| 2.10 | Demonstration of the need for high particle resolution in ensuring accurate modelling of collisional phenomena | 52 |
| 4.1 | Results of the 1D monoenergetic beam test problem | 65 |
| 4.2 | Results of the 1D monoenergetic beam test problem | 66 |
| 4.3 | Lineouts of the background electron temperatures from the 2D diffusion test problem | 67 |

| | | |
|------|--|----|
| 4.4 | Comparison of the EPOCH-C and EPOCH-H background electron temperatures in the 1D LPI test problem | 68 |
| 4.5 | Lineout of the background electron temperatures from the 2D LPI test problem | 69 |
| 4.6 | Average electron temperature as a function of time, demonstrating the rate of numerical heating for the EPOCH-C and EPOCH-H simulations listed in table 4.1. | 71 |
| 5.1 | Lineouts of the initial electron and ion densities for the EPOCH-H and solid density EPOCH TNSA simulations. | 76 |
| 5.2 | Time-integrated proton energy spectra. | 77 |
| 5.3 | Lineouts of the background electron temperature and total electron densities from TNSA simulations after $t = 2.0$ ps. | 78 |
| 5.4 | Number of protons recorded by the probes as functions of position and direction of motion | 80 |
| 5.5 | Fast electron and proton number density demonstrating imprinting of fast electron filamentation onto the accelerated proton population . . . | 81 |
| 5.6 | Fourier transform of the spatial distribution of protons recorded by the particle probes | 82 |
| 5.7 | Magnetic field inside the target and electron density in the pre-plasma . | 83 |
| 5.8 | Spatial variation in the mean direction of motion of fast electrons recorded at the front and rear surfaces | 84 |
| 5.9 | Spatial variation in the mean direction of motion of fast electrons recorded at the front and rear surfaces | 85 |
| 5.10 | Spatial variation in the divergence of fast electrons recorded at the front and rear surfaces | 85 |
| 6.1 | Illustration of the positions of the EPOCH-H and THOR simulation domains relative to each other | 91 |
| 6.2 | Background electron temperature from EPOCH-H and THOR shortly after the peak laser interaction | 92 |
| 6.3 | Lineouts of the background electron temperature in the aluminium layer. | 93 |
| 6.4 | Time history of the average electrons temperature in the aluminium layer predicted by EPOCH-H and THOR | 94 |
| 6.5 | Time history of the average electrons temperature in the aluminium layer predicted by EPOCH-H and THOR | 94 |
| 6.6 | FLY-calculated emission spectrum using temperature and density conditions predicted by EPOCH-H | 96 |

| | | |
|------|--|-----|
| 6.7 | FLY-calculated emission spectrum using temperature and density conditions predicted by THOR | 96 |
| 6.8 | FLY-calculated emission spectrum using temperature and density conditions predicted by EPOCH-H | 97 |
| 6.9 | FLY-calculated emission spectrum using temperature and density conditions predicted by THOR | 97 |
| 6.10 | Lineouts of the z -component of the magnetic fields, parallel to the target normal, along a line $25\ \mu\text{m}$ from the target's central axis. | 98 |
| 7.1 | Visual representation of the disparate time-scales involved in short-pulse laser interactions | 103 |
| 7.2 | Diagram indicating the regions of a short-pulse laser interaction which can be effectively modelled using EPOCH, THOR and EPOCH-H | 106 |

Acknowledgments

Firstly, I would like to express my gratitude to Tony Arber and Nathan Sircombe for their assistance and guidance throughout the course of this project, and to AWE for providing the necessary funding.

My thanks also to Ben Williams for putting up with the alternating complaints of failure and celebrations of success that come with sitting next to someone performing code development over the last few years.

Finally, thank you to my friends and family for their constant, unwavering support and encouragement, and for smiling and nodding in all the right places whenever I talked about this work.

Declarations

This thesis is submitted to the University of Warwick in support of my application for the degree of Doctor of Philosophy. It has been composed by myself and has not been submitted in any previous application for any degree.

The work presented (including data generated and data analysis) was carried out by the author.

The collision algorithm tests discussed in section 2.8.4 have been published as a part of: T. D. Arber *et al.*, Contemporary particle-in-cell approach to laser-plasma modelling, *Plasma Phys. Control. Fusion*, 57:113001, 2015.

Abstract

The constraints on particle-in-cell (PIC) simulations of short-pulse laser interactions with solid density targets severely limit the spatial and temporal scales which can be modelled routinely. Although recent advances in high performance computing (HPC) capabilities have rendered collisionless simulations at a scale and density directly applicable to experiments tractable, detailed modelling of the fast electron transport resulting from the laser interaction is often only possible by sampling the fast electron populations and passing this information to a separate, dedicated transport code. However, this approach potentially neglects phenomena which take place or are seeded near the transition between the two codes. Consequently there is a need to develop techniques capable of efficiently modelling fast electron transport in high density plasma without being subject to the usual grid-scale and time-step constraints. The approach employed must also be compatible with retaining the standard PIC model in the laser interaction regions in order to model laser absorption and charged particle acceleration processes. Such an approach, proposed by Cohen, Kemp and Divol [*J. Comput. Phys.*, 229:4591, 2010], has been identified, adapted and implemented in EPOCH. The final algorithm, as implemented, is presented here. To demonstrate the ability of the adapted code to model high intensity laser-plasma interactions with peak densities at, and above, solid density, the results of simulations investigating filamentation of the fast electron population and heating of the bulk target, at high densities, are presented and compared with the results of recent experiments as well as other, similar codes.

Abbreviations

Abbreviations

| | |
|---------|---|
| ASE | Amplified Spontaneous Emission |
| AWE | Atomic Weapons Establishment |
| BGK | Bhatnagar-Gross-Krook |
| BOA | Break-out afterburner |
| CAEP | Chinese Academy of Engineering Physics |
| CFL | Courant-Friedrichs-Lewy |
| CPA | Chirped Pulse Amplification |
| EM | Electromagnetic |
| FDTD | Finite Difference Time Domain |
| FWHM | Full-Width-Half-Maximum |
| HEDP | High Energy Density Physics |
| HPC | High Performance Computing |
| IAP RAS | Institute of Applied Physics, Russian Academy of Sciences |
| ICF | Inertial Confinement Fusion |
| ILE | Institute for Laser Engineering, Osaka University |
| LANL | Los Alamos National Laboratory |
| LLE | Laboratory for Laser Energetics, University of Rochester |
| LLNL | Lawrence Livermore National Laboratory |
| LPI | Laser-Plasma Interaction |
| LTE | Local Thermodynamic Equilibrium |
| MHD | Magnetohydrodynamic(s) |

| | |
|-------------|---|
| NLTE | Non-Local Thermodynamic Equilibrium |
| OPCPA | Optical Parametric Chirped Pulse Amplification |
| PIC | Particle-In-Cell |
| QED | Quantum Electrodynamics |
| RAL | Rutherford Appleton Laboratory |
| RFNC-VNIIEF | Russian Federal Nuclear Centre — All-Russian Scientific Research Institute of Experimental Physics |
| RPA | Radiation pressure acceleration |
| SIOM | Shanghai Institute of Optics and Fine Mechanics |
| TNSA | Target Normal Sheath Acceleration |
| VFP | Vlasov-Fokker-Planck |
| WDM | Warm Dense Matter |

Chapter 1

Introduction

The strong electric and magnetic fields associated with a high intensity, short-pulse laser beam, (peak intensity, $I > 10^{18} \text{ Wcm}^{-2}$ and pulse duration, $\tau_L \lesssim 1 \text{ ps}$), when incident upon a solid density target, are capable of stripping the electrons from the atoms near the surface of the target, and accelerating them to relativistic average energies ($\langle K \rangle \gtrsim 1 \text{ MeV}$) [1, 2]. These ‘fast’ electrons (also commonly referred to as ‘hot’ or ‘superthermal’ electrons) stream away from the laser interaction region near the surface, travelling deeper into the target well ahead of any thermal conduction from the laser absorption region [3]. The electric field created by the resultant charge imbalance acts to draw a ‘return current’—the background electrons inside the target drift towards the laser interaction region to restore charge neutrality [4]. Since the background electron density is much greater than the fast electron density, their average velocity is much lower, and thus these electrons are much more collisional than their higher energy counterparts. Through collisions with each other, as well as the background ion population, the return current electrons ‘thermalise’ the energy they have gained from the electric field, resulting in target heating. It is in this manner that high intensity short-pulse lasers are capable of heating solid materials to high temperatures (many hundreds of eV) over time-scales which are much shorter than the hydrodynamic response time of the target, thereby permitting investigation of material properties at high temperatures and densities [5–10].

In addition to heating high density materials, short-pulse laser interactions also provide sources of high energy electrons, protons and x-rays, and thus possess a broad range of potential applications including fast ignition inertial confinement fusion [11, 12], cancer therapy [13, 14], proton [15–17] and x-ray radiography [18–25], material activation for nuclear physics [26], and the study of material properties in solar [27] and planetary [28] interiors.

1.1 Overview

The mechanisms by which a high power, short-pulse laser beam delivers its energy to a dense plasma, how this energy is distributed amongst the constituent particles and how it is transported deeper into the target are complex and interdependent. A detailed knowledge of such phenomena is essential for understanding high energy density physics (HEDP) experiments of the kind fielded on the Vulcan-PW, OMEGA-EP and Orion laser systems, and for supporting the development of a viable point design for fast ignition inertial confinement fusion (ICF).

The UK plasma physics community has of late gained a significant capability in modelling short-pulse laser experiments. A key element of this is the particle-in-cell (PIC) code EPOCH [29], developed under the auspices of the EPSRC's 'Collaborative Computational Projects' programme¹.

This modelling capability needs to be enhanced to include additional physics, such as a self-consistent model for the transport of fast electrons through the dense, cold material that makes up the bulk of the target, and optimised to make efficient use of available massively parallel high performance computing (HPC). To tackle the multi-scale problems of a combined long and short-pulse experiment (such as those planned for Orion) it is also necessary to integrate kinetic models with existing hydrodynamics codes [30, 31].

The project discussed in this thesis was undertaken with the aim of further developing the short-pulse modelling tools and techniques that will be required to support laser-plasma experimental campaigns on existing and planned facilities. This involved considerable algorithmic development of EPOCH to improve its ability to model high density plasmas and provide a multi-scale capability via 'in line' links to hydrocodes, as well as providing modelling support for short-pulse HEDP and fast ignition-relevant experiments.

In conclusion, this project has driven the development of EPOCH to meet the challenge of modelling future short-pulse laser-plasma experiments, so as to provide a better understanding of fundamental laser-matter interactions, and the subsequent energy transport and deposition in high density material.

In the remainder of this chapter the concepts and phenomena which will be referred to in subsequent chapters are introduced. Chapter 2 provides an overview of the PIC code EPOCH. An algorithm to permit more efficient modelling of high density regions in PIC simulations, based on that developed by Cohen et al. [32], has been implemented in EPOCH, and is presented in chapter 3. A selection of tests of the high

¹The EPOCH code was developed under UK Engineering and Physical Sciences Research Council (EPSRC) grants EP/G054940/1, EP/G055165/1 and EP/G056803/1.

density model, in one and two dimensions, are discussed in chapter 4. Standard and modified EPOCH simulations of fast electron transport through a solid density target, and the subsequent acceleration of protons from the rear surface, were performed and compared. Finally, the ability to simulate return current heating in solid targets using the high density model was demonstrated via comparisons with experimental results, as well as a bespoke fast electron transport code.

1.2 Basic plasma physics

1.2.1 Plasma equations of motion

The motion of an individual particle, i , of species j , with charge q_j and mass m_j , can be expressed as follows:

$$\begin{aligned}\frac{d\mathbf{r}_i(t)}{dt} &= \mathbf{V}_i(t), \\ \frac{d\mathbf{V}_i(t)}{dt} &= \frac{q_j}{m_j} [\mathbf{E}_i(\mathbf{r}_i(t), t) + \mathbf{V}_i(t) \times \mathbf{B}_i(\mathbf{r}_i(t), t)],\end{aligned}$$

where \mathbf{r}_i is the position, and \mathbf{V}_i the velocity, of particle i as a function of time. For a collection of N_j identical particles, N_j pairs of the above equations exist which define the trajectory of the particles through six-dimensional phase-space (\mathbf{x}, \mathbf{v}) , where \mathbf{x} and \mathbf{v} represent arbitrary spatial and velocity coordinates.

Each particle is represented by a delta function, δ , in phase-space (where the delta function is defined by $\int_{-\infty}^{\infty} f(x) \delta(x - a) dx = f(a)$). Thus the number density at an arbitrary point in phase-space and time is:

$$g_j(\mathbf{x}, \mathbf{v}, t) = \sum_{i=1}^{N_j} [\delta(\mathbf{x} - \mathbf{r}_i(t)) \delta(\mathbf{v} - \mathbf{V}_i(t))].$$

The charge and current densities, ρ and \mathbf{J} , respectively, are given by:

$$\begin{aligned}\rho(\mathbf{x}, t) &= \sum_j \left(q_j \int g_j(\mathbf{x}, \mathbf{v}, t) d\mathbf{v} \right), \\ \mathbf{J}(\mathbf{x}, t) &= \sum_j \left(q_j \int \mathbf{v} g_j(\mathbf{x}, \mathbf{v}, t) d\mathbf{v} \right),\end{aligned}$$

and the electric and magnetic fields are determined by solutions of Maxwell's equations:

$$\begin{aligned}\nabla \cdot \mathbf{E} &= \frac{\rho(\mathbf{x}, t)}{\epsilon_0}, \\ \nabla \cdot \mathbf{B} &= 0, \\ \nabla \times \mathbf{E} &= -\frac{\partial \mathbf{B}}{\partial t}, \\ \nabla \times \mathbf{B} &= \mu_0 \mathbf{J}(\mathbf{x}, t) + \frac{1}{c^2} \frac{\partial \mathbf{E}}{\partial t}.\end{aligned}$$

The above set of equations define the ' N -body problem' for charged particles. Their solution is an exact model for the dynamics of a single species (charge and mass combination) of charged particles, but multiple species can be easily included. Well-known analytic solutions exist for single particles. However due to the large number of particles in full-scale plasmas, obtaining an analytic solution for such systems is not tractable. For example, in the solid plastic foils typically used in short-pulse laser experiments (with dimensions $1 \text{ mm} \times 1 \text{ mm} \times 20 \text{ }\mu\text{m}$) N is of the order 10^{19} .

Taking the derivative of $g_j(\mathbf{x}, \mathbf{v}, t)$ with respect to time, and employing the chain rule, $\frac{d\delta(\mathbf{x}-\mathbf{r}_i)}{dt} = \frac{d\mathbf{r}_i}{dt} \cdot \nabla_{\mathbf{r}_i} \delta(\mathbf{x} - \mathbf{r}_i)$, yields:

$$\frac{\partial g_j}{\partial t} = \sum_{i=1}^{N_j} \left[\frac{d\mathbf{r}_i}{dt} \cdot \nabla_{\mathbf{r}_i} \delta(\mathbf{x} - \mathbf{r}_i) \delta(\mathbf{v} - \mathbf{V}_i) + \frac{d\mathbf{V}_i}{dt} \cdot \nabla_{\mathbf{V}_i} \delta(\mathbf{x} - \mathbf{r}_i) \delta(\mathbf{v} - \mathbf{V}_i) \right].$$

Using the identity $\frac{\partial f(a-b)}{\partial a} = -\frac{\partial f(a-b)}{\partial b}$ the spatial and velocity derivatives can be rewritten in terms of \mathbf{x} and \mathbf{v} . The original expressions for the time derivatives of \mathbf{r}_i and \mathbf{V}_i can also be substituted:

$$\begin{aligned}\frac{\partial g_j}{\partial t} &= - \sum_{i=1}^{N_j} \mathbf{V}_i \cdot \nabla \delta(\mathbf{x} - \mathbf{r}_i) \delta(\mathbf{v} - \mathbf{V}_i) \\ &\quad - \sum_{i=1}^{N_j} \frac{q_j}{m_j} [\mathbf{E}_i(\mathbf{r}_i, t) + \mathbf{V}_i \times \mathbf{B}_i(\mathbf{r}_i, t)] \cdot \nabla_{\mathbf{v}} \delta(\mathbf{x} - \mathbf{r}_i) \delta(\mathbf{v} - \mathbf{V}_i).\end{aligned}$$

The identity $f(a)\delta(a-b) = f(b)\delta(a-b)$ can be used to replace the instances of \mathbf{V}_i with \mathbf{v} . Finally, the definition of $g_j(\mathbf{x}, \mathbf{v}, t)$ can be substituted to obtain the Klimontovich equation:

$$\frac{\partial g_j}{\partial t} + \mathbf{v} \cdot \nabla g_j + \frac{q_j}{m_j} (\mathbf{E} + \mathbf{v} \times \mathbf{B}) \cdot \nabla_{\mathbf{v}} g_j = 0. \quad (1.1)$$

This equation follows directly from the N -body model without any assumptions or

approximations. The Klimontovich equation describes the motion of N_j particles in phase-space as a function, g_j , which consists of a sum of N_j delta functions. Solving equation 1.1 is generally not tractable, and so a more flexible description of particle phase-space must be sought.

If a range of values in phase-space is considered: $(\mathbf{x}, \mathbf{x} + d\mathbf{x})$, $(\mathbf{v}, \mathbf{v} + d\mathbf{v})$, then $g_j d\mathbf{x} d\mathbf{v}$ is the number of particles within this volume. A smooth, continuous distribution function can be introduced by taking the ensemble average of g_j :

$$f_j(\mathbf{x}, \mathbf{v}, t) \equiv \langle g_j(\mathbf{x}, \mathbf{v}, t) \rangle.$$

Now $f_j d\mathbf{x} d\mathbf{v}$ is the expected number of particles within the phase-space volume, rather than the absolute number. Similar smoothed functions can be used to replace the electric and magnetic fields.

The difference between the descriptions of the plasma provided by the exact and smooth functions can be attributed to discrete particle collisions. Thus:

$$g_j(\mathbf{x}, \mathbf{v}, t) = f_j(\mathbf{x}, \mathbf{v}, t) + \delta g_j(\mathbf{x}, \mathbf{v}, t),$$

where δg_j is the contribution from discrete particle interactions. Similarly:

$$\begin{aligned} \mathbf{E}(\mathbf{x}, t) &= \mathbf{E}_0(\mathbf{x}, t) + \delta \mathbf{E}(\mathbf{x}, t), \\ \mathbf{B}(\mathbf{x}, t) &= \mathbf{B}_0(\mathbf{x}, t) + \delta \mathbf{B}(\mathbf{x}, t). \end{aligned}$$

Substituting these expressions into equation 1.1, and noting that $f_j = \langle g_j \rangle$, $\mathbf{E}_0 = \langle \mathbf{E} \rangle$, $\mathbf{B}_0 = \langle \mathbf{B} \rangle$, and $\langle \delta g_j \rangle = \langle \delta \mathbf{E} \rangle = \langle \delta \mathbf{B} \rangle = 0$, gives the plasma kinetic equation:

$$\frac{\partial f_j}{\partial t} + \mathbf{v} \cdot \nabla f_j + \frac{q_j}{m_j} (\mathbf{E}_0 + \mathbf{v} \times \mathbf{B}_0) \cdot \nabla_{\mathbf{v}} f_j = -\frac{q_j}{m_j} \langle (\delta \mathbf{E} + \mathbf{v} \times \delta \mathbf{B}) \cdot \nabla_{\mathbf{v}} \delta g_j \rangle.$$

Collecting the terms on the right-hand-side into a single term results in:

$$\frac{\partial f_j}{\partial t} + \mathbf{v} \cdot \nabla f_j + \frac{q_j}{m_j} (\mathbf{E}_0 + \mathbf{v} \times \mathbf{B}_0) \cdot \nabla_{\mathbf{v}} f_j = \left(\frac{\partial f_j}{\partial t} \right)_{\text{coll}}. \quad (1.2)$$

The left-hand-side of equation 1.2 is the rate of change of the distribution function, f_j , due to the smooth, averaged fields, \mathbf{E}_0 and \mathbf{B}_0 , while the right-hand-side represents any additional contribution due to collisional effects.

The overall number, charge and current densities are defined as:

$$\begin{aligned} n_j(\mathbf{x}, t) &= \int f_j(\mathbf{x}, \mathbf{v}, t) d\mathbf{v}, \\ \rho(\mathbf{x}, t) &= \sum_j \left(q_j \int f_j(\mathbf{x}, \mathbf{v}, t) d\mathbf{v} \right), \\ \mathbf{J}(\mathbf{x}, t) &= \sum_j \left(q_j \int \mathbf{v} f_j(\mathbf{x}, \mathbf{v}, t) d\mathbf{v} \right). \end{aligned}$$

It should be noted that the plasma kinetic equation is still an exact description of the behaviour of a plasma. No assumptions or approximations have been made. In order to evolve the distribution function, a form for the collision operator needs to be substituted on the right-hand-side of equation 1.2. The simplest approach is the Krook collision operator [33], which assumes that the distribution function will tend towards some equilibrium distribution, f_m over a characteristic time, τ :

$$\left(\frac{\partial f}{\partial t} \right)_{\text{coll}} = \frac{f_m - f}{\tau}.$$

A more rigorous approach is to use the Fokker-Planck equation (derived in appendix A), resulting in the Vlasov-Fokker-Planck (VFP) equation:

$$\frac{\partial f_j}{\partial t} + \mathbf{v} \cdot \nabla f_j + \frac{q_j}{m_j} (\mathbf{E}_0 + \mathbf{v} \times \mathbf{B}_0) \cdot \nabla_{\mathbf{v}} f_j = -\nabla_{\mathbf{v}} \cdot \left[f_j \frac{\langle \Delta \mathbf{v} \rangle}{\Delta t} \right] + \frac{1}{2} \nabla_{\mathbf{v}}^2 : \left[f_j \frac{\langle \Delta \mathbf{v} \Delta \mathbf{v} \rangle}{\Delta t} \right]. \quad (1.3)$$

1.2.2 Ohm's law

The constituent particles which make up a short-pulse laser target can be divided into three 'species': the background electrons (b subscripts) and ions (i subscripts) which make up the bulk of the target, and the relativistic (fast) electrons (f subscripts) which have been accelerated by the laser interaction. The electric field associated with the fast electron current draws a 'return current' via the background electrons. An equation of motion for the background electrons can be obtained by taking the first velocity moment of equation 1.2:

$$\int m_e \mathbf{v} \left(\frac{\partial f_b}{\partial t} + \mathbf{v} \cdot \nabla f_b + \frac{q_e}{m_e} (\mathbf{E} + \mathbf{v} \times \mathbf{B}) \cdot \nabla_{\mathbf{v}} f_b \right) d^3v = \int m_e \mathbf{v} \left(\frac{\partial f_b}{\partial t} \right)_{\text{coll}} d^3v, \quad (1.4)$$

The first term of equation 1.4 can be evaluated as follows:

$$\int m_e \mathbf{v} \frac{\partial f_b}{\partial t} d^3v = m_e \frac{\partial}{\partial t} \int \mathbf{v} f_b d^3v = m_e \frac{\partial (n_b \mathbf{u}_b)}{\partial t} = m_e n_b \frac{\partial \mathbf{u}_b}{\partial t} + m_e \mathbf{u}_b \frac{\partial n_b}{\partial t}, \quad (1.5)$$

where $\mathbf{u}_b = \langle \mathbf{v} \rangle$ is the local background electron drift velocity.

In evaluating the second term, it is assumed that the electron velocity, \mathbf{v} , can be expressed as the sum of the drift velocity and a perturbation: $\mathbf{v} = \mathbf{u}_b + \mathbf{w}$. Since the mean electron velocity has already been defined as being equal to the drift velocity, it follows that the mean perturbation is zero: $\langle \mathbf{w} \rangle = 0$.

$$\begin{aligned} \int m_e \mathbf{v} (\mathbf{v} \cdot \nabla f_b) d^3v &= m_e \int \nabla \cdot (f_b \mathbf{v} \mathbf{v}) d^3v = m_e \nabla \cdot \int f_b \mathbf{v} \mathbf{v} d^3v, \\ &= m_e \nabla \cdot \int f_b (\mathbf{u}_b \mathbf{u}_b + \mathbf{u}_b \mathbf{w} + \mathbf{w} \mathbf{u}_b + \mathbf{w} \mathbf{w}) d^3w, \\ &= m_e \nabla \cdot (n_b \mathbf{u}_b \mathbf{u}_b) + \nabla \cdot \underline{\mathbf{P}}_b, \end{aligned}$$

where $\underline{\mathbf{P}}_b$ is the stress tensor and is defined as $\underline{\mathbf{P}}_b = m_e \int (\mathbf{v} - \mathbf{u}_b) (\mathbf{v} - \mathbf{u}_b) f_b d^3v$.

$$\therefore \int m_e \mathbf{v} (\mathbf{v} \cdot \nabla f_b) d^3v = m_e n_b (\mathbf{u}_b \cdot \nabla) \mathbf{u}_b - m_e \frac{\partial n_b}{\partial t} \mathbf{u}_b + \nabla \cdot \underline{\mathbf{P}}_b. \quad (1.6)$$

The third and fourth (electric and magnetic field) terms in equation 1.4 can be considered together:

$$\begin{aligned} q_e \int \mathbf{v} (\mathbf{E} + \mathbf{v} \times \mathbf{B}) \cdot \nabla_{\mathbf{v}} f_b d^3v &= q_e \int \nabla_{\mathbf{v}} \cdot (f_b \mathbf{v} [\mathbf{E} + \mathbf{v} \times \mathbf{B}]) \\ &\quad - f_b \mathbf{v} (\nabla_{\mathbf{v}} \cdot [\mathbf{E} + \mathbf{v} \times \mathbf{B}]) \\ &\quad - f_b (\mathbf{E} + \mathbf{v} \times \mathbf{B}) \cdot \nabla_{\mathbf{v}} \mathbf{v} d^3v. \end{aligned}$$

Using Gauss's theorem the first term on the right-hand-side can be rewritten as:

$$q_e \int \nabla_{\mathbf{v}} \cdot (f_b \mathbf{v} [\mathbf{E} + \mathbf{v} \times \mathbf{B}]) d^3v = q_e \int_S (f_b \mathbf{v} [\mathbf{E} + \mathbf{v} \times \mathbf{B}]) \cdot d\mathbf{S}_v,$$

where S is a surface in velocity space. If the surface is taken to infinity, the integral tends towards zero, since the electron distribution tends to zero with increasing velocity much faster than the area of the surface increases (e^{-v^2} cf. v^2). The second term on the right-hand-side also vanishes, since $\nabla_{\mathbf{v}} \cdot \mathbf{E} = 0$ and $\nabla_{\mathbf{v}} \cdot (\mathbf{v} \times \mathbf{B}) = 0$. The remaining, third term becomes:

$$-q_e \int f_b (\mathbf{E} + \mathbf{v} \times \mathbf{B}) d^3v = -q_e n_b \langle \mathbf{E} + \mathbf{u}_b \times \mathbf{B} \rangle.$$

$$\therefore \int m_e \mathbf{v} \frac{q_e}{m_e} (\mathbf{E} + \mathbf{v} \times \mathbf{B}) \cdot \nabla_{\mathbf{v}} f_b d^3v = -q_e n_b (\mathbf{E} + \mathbf{u}_b \times \mathbf{B}). \quad (1.7)$$

The collision term in equation 1.4 can be reduced to $\int m_e \mathbf{v} (\partial f_b / \partial t)_{\text{coll}} d^3v = m_e (\partial n_b \mathbf{u}_b / \partial t)_{\text{coll}} = m_e \mathbf{u}_b (\partial n_b / \partial t)_{\text{coll}} + m_e n_b (\partial \mathbf{u}_b / \partial t)_{\text{coll}}$. If ionisation and recombination are neglected, $(\partial n_b / \partial t)_{\text{coll}} = 0$, and this term reduces to:

$$\int m_e \mathbf{v} \left(\frac{\partial f_b}{\partial t} \right)_{\text{coll}} d^3v = m_e n_b \left(\frac{\partial \mathbf{u}_b}{\partial t} \right)_{\text{coll}}. \quad (1.8)$$

By combining equations 1.5–1.8, equation 1.4 becomes:

$$m_e n_b \frac{\partial \mathbf{u}_b}{\partial t} + m_e n_b (\mathbf{u}_b \cdot \nabla) \mathbf{u}_b + \nabla \cdot \mathbf{P}_b - q_e n_b (\mathbf{E} + \mathbf{u}_b \times \mathbf{B}) = m_e n_b \left(\frac{\partial \mathbf{u}_b}{\partial t} \right)_{\text{coll}}.$$

The collision term on the right-hand-side can be split into separate expressions for collisions with ions, $m_e n_b (\partial \mathbf{u}_b / \partial t)_{bi}$, and with a separate energetic electron species (hereafter referred to as fast electrons), $m_e n_b (\partial \mathbf{u}_b / \partial t)_{bf}$:

$$m_e n_b \left(\frac{\partial \mathbf{u}_b}{\partial t} \right)_{\text{coll}} = m_e n_b \nu_{bi} (\mathbf{u}_i - \mathbf{u}_b) + m_e n_b \left(\frac{\partial \mathbf{u}_b}{\partial t} \right)_{bf}.$$

Thus the equation of motion for the background electrons (those not considered to be fast electrons) is given by:

$$m_e \left(\frac{\partial}{\partial t} + \mathbf{u}_b \cdot \nabla \right) \mathbf{u}_b = q_e (\mathbf{E} + \mathbf{u}_b \times \mathbf{B}) - \frac{\nabla \cdot \mathbf{P}_b}{n_b} - m_e \nu_{bi} (\mathbf{u}_b - \mathbf{u}_i) + m_e \left(\frac{\partial \mathbf{u}_b}{\partial t} \right)_{bf}, \quad (1.9)$$

where the b , i and f subscripts denote background electron, ion and fast electron quantities, respectively.

The electron-ion drag term in the background electron equation of motion can be rewritten as follows:

$$\begin{aligned} m_e \nu_{bi} (\mathbf{u}_b - \mathbf{u}_i) &= \frac{m_e \nu_{bi}}{q_e n_b} (q_e n_b \mathbf{u}_b - q_e n_b \mathbf{u}_i) \\ &= q_e \eta (\mathbf{J}_b + \mathbf{J}_i), \end{aligned}$$

since $q_e n_b \mathbf{u}_b = \mathbf{J}_b$, $-q_e n_b \mathbf{u}_i = q_i n_i \mathbf{u}_i = \mathbf{J}_i$ and $\eta = m_e \nu_{bi} / q_e^2 n_b$, assuming a quasi-neutral plasma ($n_b \approx Z^* n_i$).

If the background electron response is assumed to be instantaneous on the time-scales of interest, then the inertia term (i.e. the left-hand-side of equation 1.9) can be

assumed to be zero:

$$\left(\frac{\partial}{\partial t} + \mathbf{u}_b \cdot \nabla \right) \mathbf{u}_b = \frac{d\mathbf{u}_b}{dt} = 0.$$

Re-arranging the remaining terms in equation 1.9 gives an expression for Ohm's law:

$$\mathbf{E} = \eta (\mathbf{J}_b + \mathbf{J}_i) - \mathbf{u}_b \times \mathbf{B} + \frac{\nabla \cdot \mathbf{P}_b}{q_e n_b} - \frac{m_e}{q_e} \left(\frac{\partial \mathbf{u}_b}{\partial t} \right)_{\text{bf}}.$$

1.2.3 Debye length

Consider a quasi-neutral plasma. If a charged particle is inserted its electric field, which usually scales as $1/r^2$ in vacuum, will act to attract the electrons in the plasma and repel the ions. Thus the plasma particles act to 'screen' the test particle's field. This results in an electric field which decays exponentially with distance, rather than the inverse square. The characteristic range over which the electric field drops by a factor of $1/e$ is referred to as the Debye length and can be derived as follows:

Let $n_0 = 0.5 (n_e + Z^* n_i)$, where Z^* is the ionisation state of the ions and quasi-neutrality has been assumed ($n_e = Z^* n_i$). Assume also that the electrons and ions are in thermal equilibrium so that $T_e = T_i = T$. For a positive test charge added to the plasma, the potential felt by an electron a distance r from the test charge is given by $\phi(r)$.

Using Boltzmann's law an expression for the number density of electrons at potential $\phi(r)$ can be found:

$$n_e(r) = n_0 \exp \left(-\frac{q_e \phi(r)}{k_B T} \right) \approx n_0 \left(1 - \frac{q_e \phi(r)}{k_B T} \right),$$

assuming $|q_e \phi| \ll k_B T$ (i.e. that the plasma is sufficiently hot that the thermal energy is much larger than the electric potential energy).

This indicates an excess of electrons: $n_e(r) = n_0 + n_e^+(r)$ where $n_e^+(r) = -n_0 q_e \phi(r) / k_B T$. Similarly, there is a deficit of ions, $n_i^-(r)$. Thus the excess charge density at a distance r from the test charge is:

$$\rho(r) = q_e (n_e^+(r) - Z^* n_i^-(r)) = 2n_0 \frac{q_e \phi(r)}{k_B T}.$$

This can then be inserted into Poisson's equation ($\nabla^2 \phi(r) = -\rho / \epsilon_0$) to give:

$$\nabla^2 \phi(r) = \frac{1}{r^2} \frac{\partial}{\partial r} \left(r^2 \frac{\partial \phi(r)}{\partial r} \right) = \frac{2n_0 q_e^2 \phi(r)}{\epsilon_0 k_B T}.$$

This has the solution $\phi(r) = (A/r) e^{-\sqrt{2}r/\lambda_D}$, where A is constant and λ_D is the Debye

length, defined as:

$$\lambda_D = \left(\frac{\epsilon_0 k_B T}{q_e^2 n_0} \right)^{1/2}.$$

For completeness, an expression for A can be found by requiring that the potential tend towards the value due to the test charge as r tends toward zero:

$$\phi(r) = \frac{q}{4\pi\epsilon_0 r} e^{-\sqrt{2}r/\lambda_D}.$$

1.2.4 Plasma frequency

Consider a charge-neutral cube of plasma, with length l and number density $n = n_e = Z^* n_i$. If all of the electrons are displaced by a distance x , an electric field is set up, attracting the electrons back towards the ions (and vice-versa). Since the mass of an ion is much greater than an electron, the electrons will respond to the electric field much more rapidly than the ions. Thus it can be assumed that the ions will be stationary over the time-scales of interest.

Using Gauss' theorem to obtain an expression for the electric field, it is possible to derive an equation of motion for the electrons:

$$\begin{aligned} E &= -\frac{n_e q_e x}{\epsilon_0}, \\ F &= m_e \ddot{x} = q_e E \\ \therefore \ddot{x} &= -\frac{n_e q_e^2 x}{\epsilon_0 m_e}. \end{aligned}$$

This is the equation of motion for simple harmonic oscillation ($\ddot{x} = -\omega_p^2 x$) with frequency:

$$\omega_p = \left(\frac{n_e q_e^2}{m_e \epsilon_0} \right)^{1/2}.$$

This oscillatory motion is known as an electron plasma wave. It should be noted that this is only valid if $\omega_p \tau_{ei} \gg 1$, where τ_{ei} is the average time between electron-ion collisions.

Note also that

$$\begin{aligned} \omega_p &= \left(\frac{n_e q_e^2}{m_e \epsilon_0} \right)^{1/2}, \\ &= \left(\frac{n_e q_e^2}{\epsilon_0 k_B T} \right)^{1/2} \left(\frac{k_B T_e}{m_e} \right)^{1/2}, \\ &= \frac{v_{th}}{\lambda_D}, \end{aligned}$$

where $v_{th} = \sqrt{k_B T_e / m_e}$ is the thermal velocity of the electrons.

1.2.5 Plasma dispersion relation

Due to the screening of electric fields by electrons in plasma, the dispersion relation for electromagnetic waves is modified. An expression for the plasma dispersion relation can be derived by taking the curl of Faraday's law:

$$\begin{aligned}\nabla \times (\nabla \times \mathbf{E}) &= -\frac{\partial}{\partial t} (\nabla \times \mathbf{B}), \\ \nabla (\nabla \cdot \mathbf{E}) - \nabla^2 \mathbf{E} &= -\frac{\partial}{\partial t} (\nabla \times \mathbf{B}).\end{aligned}$$

If a wave-like solution for the electric field is assumed (i.e. $\mathbf{E} = \mathbf{E}_0 e^{i(\omega t - \mathbf{k} \cdot \mathbf{x})}$), and Ampère's law is substituted into the right-hand-side, this becomes:

$$-\mathbf{k} (\mathbf{k} \cdot \mathbf{E}) + k^2 \mathbf{E} = -\mu_0 \frac{\partial}{\partial t} \left(\mathbf{J} + \epsilon_0 \frac{\partial \mathbf{E}}{\partial t} \right).$$

Substituting for the current density, \mathbf{J} , using Ohm's law ($\mathbf{J} = \sigma \mathbf{E}$), and using the wave-like solution to simplify the time derivative yields

$$-\mathbf{k} (\mathbf{k} \cdot \mathbf{E}) + k^2 \mathbf{E} = -\mu_0 \mathbf{E} (i\omega\sigma - \omega^2\epsilon_0). \quad (1.10)$$

To progress further, it is necessary to obtain an expression for the AC conductivity, σ . This is achieved by considering a separate uniform, oscillating electric field, $\mathbf{E} = \mathbf{E}_0(\omega) e^{i\omega t}$. The Lorentz force experienced by an electron in this field is:

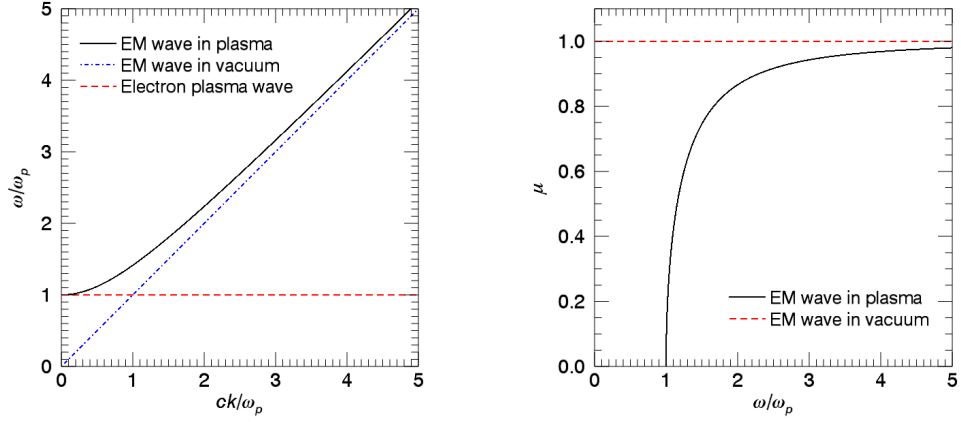
$$m_e \frac{\partial \mathbf{v}}{\partial t} = q_e \mathbf{E} = q_e \mathbf{E}_0(\omega) e^{i\omega t}.$$

If an oscillating velocity, resulting from this applied field is assumed ($\mathbf{v}(t) = \mathbf{v}(\omega) e^{i\omega t}$), this becomes:

$$\begin{aligned}i\omega m_e \mathbf{v}(\omega) &= q_e \mathbf{E}_0(\omega), \\ \frac{i\omega m_e \mathbf{J}(\omega)}{q_e n_e} &= q_e \mathbf{E}_0(\omega), \\ \therefore \mathbf{J}(\omega) &= -\frac{in_e q_e^2}{m_e \omega} \mathbf{E}_0(\omega).\end{aligned}$$

Thus, from Ohm's law it can be concluded that the AC conductivity is given by:

$$\sigma(\omega) = -\frac{in_e q_e^2}{m_e \omega}.$$



(a) The plasma dispersion relations for electromagnetic waves in plasma.

(b) The refractive index, μ , for electromagnetic waves propagating through plasma.

Figure 1.1: The plasma dispersion relations (a) and refractive indices (b) for electromagnetic (solid black) and electron plasma (dashed red) waves in a cold plasma. Note that electromagnetic waves below the plasma frequency are forbidden, as discussed in section 1.2.6.

Substituting this into equation 1.10 then yields:

$$-\mathbf{k}(\mathbf{k} \cdot \mathbf{E}) + k^2 \mathbf{E} = \mu_0 \epsilon_0 \mathbf{E} \left(\omega^2 - \frac{n_e q_e^2}{m_e \epsilon_0} \right).$$

Since the second term of the right-hand-side can be identified as the plasma frequency, this simplifies to:

$$(\omega^2 - \omega_p^2 - c^2 k^2) \mathbf{E} + c^2 \mathbf{k}(\mathbf{k} \cdot \mathbf{E}) = 0.$$

For longitudinal (e.g. Langmuir) waves, \mathbf{k} and \mathbf{E} are parallel. Thus $k^2 \mathbf{E} = \mathbf{k}(\mathbf{k} \cdot \mathbf{E})$, and so $\omega = \omega_p$. However, for transverse (e.g. electromagnetic) waves, \mathbf{k} and \mathbf{E} are perpendicular, and so $\mathbf{k} \cdot \mathbf{E} = 0$. The dispersion relation for electromagnetic waves in plasma is therefore given by:

$$\omega^2 = \omega_p^2 + c^2 k^2.$$

As discussed in the next section (1.2.6), and shown by figure 1.1, the dispersion relation of electromagnetic waves in plasma prohibits the propagation of waves with frequencies below the electron plasma frequency.




Figure 1.2: The reflectivity of a plasma as a function of its electron number density, normalised to the critical density (equation 1.11).

1.2.6 Critical density

Following the derivation of the dispersion relation for electrostatic and electromagnetic waves in plasma, it can be shown that electromagnetic waves whose frequencies are less than the plasma frequency cannot propagate through the plasma of the corresponding density due to screening of the wave's electric field by the electrons. It is possible to arrive at a condition for the electron density above which an electromagnetic wave of a given frequency cannot propagate.

Rearranging the plasma dispersion relation above gives:

$$c^2 = \frac{\omega^2}{k^2} \left(1 - \frac{\omega_p^2}{\omega^2} \right).$$

From this it can be seen that the refractive index of a plasma is given by $\mu(\omega) = (1 - \omega_p^2/\omega^2)^{1/2}$. The reflectivity for a given frequency is $R = |(1 - \mu) / (1 + \mu)|^2$, and is plotted as a function of electron number density in figure 1.2. Three distinct regimes can now be considered.

For waves with $\omega > \omega_p$, the reflectivity is less than one, and so the wave can propagate through the plasma. As ω tends to infinity, the refractive index tends toward unity, and the reflectivity to zero, as for vacuum. If $\omega < \omega_p$, however, μ is imaginary and $R = 1$. Thus electromagnetic waves cannot propagate through plasma with a plasma

frequency greater than the incident wave's frequency, and will instead be specularly reflected.

The electron number density corresponding to $\omega = \omega_p$ is referred to as the critical density, n_c :

$$\begin{aligned}\omega_p(n_e = n_c) &= \omega, \\ \therefore n_c &= \frac{m_e \epsilon_0 \omega^2}{q_e^2}.\end{aligned}\tag{1.11}$$

A wave interacting with plasma at the critical density will be reflected ($\mu = 0$, $R = 1$), and may drive resonant behaviour, an example of which is discussed in section 1.4.2.

1.3 High power lasers

1.3.1 'Long-pulse' lasers

High power lasers are generally divided into two broad categories based on the duration of the laser pulse produced. 'Long-pulse' lasers typically have pulse lengths on the order of a few nanoseconds. Many of the world's current and planned major high power laser facilities have multiple long-pulse laser beam lines: NIF (LLNL, 192), UFL-2M (RFNC-VNIIEF, 192 planned), LMJ (CEA, 176), OMEGA (LLE, 60), SG-III (CAEP, 48 planned), Gekko XII (ILE, 12), Orion (AWE, 10), Vulcan (RAL, 6).

Discussions of long-pulse laser interactions and their applications are beyond the scope of this work. Suffice it to say that when employed in conjunction with short-pulse lasers (discussed below), long-pulse lasers are generally used to compress a target to high densities, or deliberately ablate some of the surface material to create a region of low density plasma in front of the target which will potentially enhance the fraction of short-pulse laser energy absorbed.

1.3.2 'Short-pulse' lasers

The development of the chirped pulse amplification (CPA) process in the 1980s [34], and later, optical parametric chirped pulse amplification (OPCPA) [35], permitted the generation of much shorter (~ 1 ps), and thus higher power (~ 1 PW), laser pulses. These high power, short-pulse lasers can be focused to intensities of up to about 10^{21} Wcm^{-2} . When exposed to the electric fields associated with such intensities electrons are rapidly stripped from their parent atoms, and accelerated to relativistic energies. The high power of these lasers allows energy to be delivered over time-scales

much shorter than the normal hydrodynamic response time of the material. This rapid delivery of energy allows matter to be heated to extreme (hundreds of eV) temperatures, while remaining at a high (solid) density, resulting in conditions which would otherwise be very difficult to achieve under laboratory conditions.

Short-pulse laser pulses with energies between approximately 100 and 500 J typically have pulse durations between several hundred femtoseconds and several picoseconds. Examples of this type of laser can be found at several laser facilities: OMEGA-EP (LLE), Vulcan (RAL), Orion (AWE), Titan (LLNL), Trident (LANL), PETAL (CEA), GEKKO XII (ILE), SG-II-U (SIOM). Amplified spontaneous emission (ASE) and parametric fluorescence in the amplification chain are characteristic features of CPA and OPCPA lasers. These processes often result in the main pulses delivered by such lasers being preceded by a low intensity, nanosecond-scale duration pre-pulse [36], which will interact with the target in a similar way to a separate long-pulse laser. Several techniques have been developed to mitigate the pre-pulse on short-pulse lasers. One of which is to employ a non-linear frequency conversion crystal to convert the laser light from the first to the second harmonic [37]: since the conversion process has a minimum intensity threshold, the low intensity pre-pulse passes through the crystal without being converted, and is then spectrally filtered from the converted second harmonic light [38, 39].

Short-pulse lasers which deliver pulses carrying tens of joules of energy tend to have pulse durations on the order of tens of femtoseconds. In this way lasers such as Callisto (LLNL), PEARL (IAP RAS) and Astra Gemini (RAL) are able to achieve the same output power as higher energy short-pulse lasers. Due to the lower energy levels passing through the laser amplification chain, these lasers have a minimal pre-pulse, and are able to achieve a much higher shot rate (~ 10 shots per second cf. ~ 5 shots per day).

By employing a combination of short-pulse and long-pulse beam lines, it is possible to study the properties of matter compressed to intense pressures (by the long-pulse) and heated to extreme temperatures (by the short-pulse). In this way high power laser facilities facilitate investigations into a range of topics including, but not limited to: inertial fusion studies [11, 12, 40, 41], radiographic methods such as proton [15–17] and x-ray radiography [18–25], and material properties in astrophysical conditions [27, 28].

1.4 Short-pulse laser absorption

Many different absorption processes have been identified, each of which can play an important role depending on many variables such as the plasma density and temperature,

and the laser intensity, pulse length and angle of incidence. The abundance of absorption mechanisms, with no clear partitions between the regimes in which each occurs, necessitates a detailed numerical approach to modelling short-pulse laser-plasma interactions in order to capture as much of the pertinent physics as possible. Some of the main absorption processes are summarised below. For each process the conditions under which it becomes a major means of transferring energy out of the laser pulse are given.

1.4.1 Inverse bremsstrahlung

An electron, in plasma, oscillating in the electric field of a laser, may collide with a nearby ion, resulting in a transfer of energy. This process is often contrasted with bremsstrahlung, in which the electron-ion interaction results in the emission of radiation. Consequently this laser absorption process is commonly referred to as inverse bremsstrahlung. By balancing the rate of change of the electron distribution function with the collisional heating rate it can be shown that inverse bremsstrahlung acts to drive the electron distribution towards a supergaussian of the form $f_e \sim \exp(-v^5)$ [42].

Since the presence of an electromagnetic wave is required, inverse bremsstrahlung can only take place in underdense plasma (electron density is such that the electron plasma frequency is less than the laser frequency). The oscillatory velocity of an electron is proportional to the electric field amplitude, and since the electron-ion collision frequency scales inversely with the cube of velocity this process is usually only important for low intensity ($\lesssim 10^{16} \text{ Wcm}^{-2}$) laser-plasma interactions occurring over time-scales which are long compared to the electron-ion collision time. This results in inverse bremsstrahlung being the dominant absorption mechanism for many long-pulse (nanosecond-scale) interactions, such as inertial confinement fusion (ICF).

1.4.2 Resonance absorption

A p-polarised² electromagnetic wave obliquely incident on a plasma with a density gradient can propagate up to a density of $n_e = n_c \cos^2 \theta$ due to refraction, where θ is the angle of incidence of the laser measure relative to the target normal, and a linear density ramp has been assumed. Beyond this density the electric field decays evanescently, but is still non-zero at the critical surface. At this point the plasma frequency, by definition, matches the frequency of the incident laser, and thus a plasma wave is resonantly driven [43, 44]. The resonance acts to drive the plasma wave to

²P-polarisation is defined as being such that the electric field of the wave oscillates in the same plane as the target's density gradient. Consequently, the electric field of an s-polarised wave oscillates in a plane perpendicular to the density gradient.

very large amplitudes until it can no longer be supported and wave breaking occurs, resulting in the ejection of energetic (fast) electrons from the resonance region. The process of resonance absorption requires moderate (10^{14} – 10^{17} Wcm $^{-2}$) laser intensities in order for the evanescent field to be strong enough to drive the electrons at the critical surface, but avoid deformation of the plasma density profile due to the ponderomotive force (discussed later). This process also requires a long scale-length density ramp in order for the laser to be turned around via refraction.

1.4.3 Vacuum heating

If the scale-length of the plasma density ramp is very short compared to the laser wavelength, vacuum heating occurs [45]. Electrons in the overdense region (densities such that the electron plasma frequency is greater than the laser frequency) are directly exposed to the laser's electric field. Electrons near the surface of the plasma can be pulled, by the electric field of the laser, out of the plasma, beyond the Debye sheath. As the field reverses polarity the electrons are accelerated back into the plasma. By the time the field reverses again these electrons will have travelled beyond the skin depth ($l_s = c/\omega_p$) of the target, and thus are no longer influenced by the laser fields.

1.4.4 Ponderomotive acceleration

The finite size and duration of the laser pulse tends to result in a net force on the electrons, driving them away from regions of high laser intensity. An electron initially at rest in the middle of a Gaussian laser pulse will move under the influence of the laser's electric field. When the field reverses, however, the electron will have moved to a region of lower intensity (and thus field amplitude), and so experiences a weaker restoring force. Over time this leads to a net drift away from the peak electric field amplitude. An expression for the laser cycle-averaged ponderomotive force can be derived [46]:

$$\mathbf{F}_p = -\frac{q^2}{4m\omega^2} \nabla E_0^2,$$

where ω is the frequency, and E_0 the electric field amplitude, of the laser. This effect, parallel and perpendicular to the laser pulse, can lead to channelling (in underdense plasma) or hole boring (in overdense plasma). Note that the ponderomotive force is independent of the sign of the charge, and so will also act to expel ions from regions of high laser intensity. Due to their much higher inertia, however, ion acceleration takes place over longer time-scales than the electron interaction.

1.4.5 $J \times B$ force

At high laser intensities ($I \gtrsim 10^{18} \text{ Wcm}^{-2}$) the force due to the magnetic field of the laser acting on electrons oscillating in the electric field is non-negligible compared to the electric field force. The combination of oscillation in the electric field of the laser perpendicular to an oscillating magnetic field results in a force parallel to the laser axis which varies in time as $\cos 2\omega t$. For a lone electron this results in a lemniscate orbit rather than simple oscillation. If the laser is incident upon an overdense plasma, however, a similar effect to vacuum heating occurs, except with a longitudinal driving force, due to the magnetic field, at twice the laser frequency [47]. Therefore this acceleration is characterised by the production of energetic electron bunches, at the second laser harmonic, directed predominantly along the laser axis.

1.4.6 Other important processes

In addition to the absorption processes above, there are a number of other effects which manifest themselves during laser interactions. Two examples of such phenomena are given below.

Self-induced transparency

If the electrons in a plasma oscillate at relativistic velocities in the electric field of a laser, the local plasma frequency decreases. This relativistic behaviour allows a high intensity laser pulse to propagate beyond the classical critical density surface, and instead be reflected at relativistic critical density, γn_c [48].

Relativistic self-focusing

As discussed earlier, the finite size of a laser pulse results in electrons being driven away from the peak intensity regions by the ponderomotive force. This leads to a depletion of electrons in the centre of the pulse, and an excess in the wings. Since the dispersion relation for an electromagnetic wave in plasma is $\omega^2 = \omega_p^2 + c^2 k^2$, the phase velocity of the laser is lower in the low density channel carved by the ponderomotive force. The portions of the beam which span the density gradient surrounding the channel will undergo refraction, resulting in focussing of the beam to higher intensities in the channel. Thus the expulsion of electrons from the centre of a laser pulse can lead to focusing of the laser to higher intensities [49, 50].

1.5 Energy transport

1.5.1 Return current generation

The motion of the fast electrons produced by direct laser interaction generates an electrostatic field (via Ampère's law, $\partial \mathbf{E} / \partial t = c^2 \nabla \times \mathbf{B} - \mathbf{J} / \epsilon_0$). This field acts upon the background electrons present in the bulk of the target, causing them to drift toward the laser interaction region (via the Lorentz force, $\mathbf{F} = q(\mathbf{E} + \mathbf{v} \times \mathbf{B})$) in order to restore quasi-neutrality. This motion also helps to replenish the electron population available for acceleration by the laser.

The return current density due to an imposed fast electron current is related to the electric field via Ohm's law: $\mathbf{J}_b = \mathbf{E} / \eta$, where $\eta = m_e / q_e^2 n_b \tau_c$ is the plasma resistivity, and τ_c is the characteristic background electron-ion collision time. Substituting this into Ampère's law, and assuming magnetic field effects can be neglected (i.e. that the problem is one-dimensional), yields the following relation between the background and fast electron currents:

$$\mathbf{J}_f + \mathbf{J}_b = -\epsilon_0 \eta \frac{\partial \mathbf{J}_b}{\partial t},$$

where, for simplicity, a constant resistivity has been assumed.

A return current of the form $\mathbf{J}_b = \mathbf{J}_f (e^{-t/\tau} - 1)$ can be substituted into the above, and results in an expression for the response time of the plasma to changes in the fast electron current [51]: $\tau = \epsilon_0 \eta$.

For most conditions of interest this response time is much shorter than the plasma period, ω_p^{-1} , and so the return current is often assumed to respond instantaneously to changes in the fast electron current.

1.5.2 Ohmic heating

Since the background electron density is usually much greater than the density of the fast electrons, their drift velocity does not need to be very large in order for the counter-propagating currents to balance. Since their velocity is low, the background electrons are highly collisional compared to the fast electrons, and so act to heat the plasma. In this way laser energy can be coupled to the target beyond the laser interaction region, and any associated thermal conduction front, via the return current drawn by the fast electrons.

The rate of change of the background electron energy density, W_b , due to the return current, \mathbf{J}_b , drawn by the electric field, \mathbf{E} , generated by an imposed fast electron

current, is given by [52]:

$$\frac{\partial W_b}{\partial t} = \mathbf{E} \cdot \mathbf{J}_b.$$

By combining this with the material's specific heat ($C_V = (\partial W_b / \partial T)_V / n_b$), it is possible to calculate the Ohmic heating rate:

$$\frac{\partial T_b}{\partial t} = \frac{1}{n_b C_V} \frac{\partial W_b}{\partial t} = \frac{\mathbf{E} \cdot \mathbf{J}_b}{n_b C_V}. \quad (1.12)$$

1.5.3 Filamentation

The transport and deposition of energy inside the target is further complicated by the fast electrons' self-generated magnetic fields, which have been neglected thus far. In some cases these fields can act to collimate the fast electron beam [53]. More often, though, noise in the fast electron current density, background resistivity and/or small-scale magnetic fields inside the target lead to the fast electron population breaking up into filaments [54, 55]. The magnetic fields associated with these localised regions of higher current act to deflect fast electrons towards regions of higher current density, enhancing the filamentation. More recently, studies of the radiation and particles emitted from the rear of laser-irradiated targets have allowed for investigations into the effect of electrical conductivity [56, 57] and lattice structure [58] on fast electron filamentation.

1.5.4 X-ray emission

As fast electrons travel through cold, dense material they undergo elastic scattering off of the background ions. This results in the emission of bremsstrahlung ('braking radiation'), providing a broadband x-ray source. Line emission results from a fast electron colliding inelastically with a bound, inner shell electron, liberating or promoting it to a higher energy level. As an electron relaxes to fill the vacancy, a photon is emitted.

The bremsstrahlung emission is dependent upon the laser and plasma parameters (the former tend to determine the incident fast electron energy, while the latter determine the bremsstrahlung cross-section), and results in a continuous x-ray spectrum typically ranging from hundreds of eV to tens of MeV.

Line emission, however, is characterised by intense, narrowband peaks in intensity at energies up to approximately 100 keV. The precise energies depend upon the atomic structure of the target's constituent element(s). By examining the width of the lines, and their relative heights, it is often possible to infer the conditions, such as density and temperature, inside the emitting region of the target [59].

1.6 Summary of short-pulse laser interactions

Several examples of the principal means for short-pulse laser energy to be absorbed by a plasma have been discussed above. It is worth noting that all of these processes rely on the motion of the electrons in the laser's electromagnetic fields. Due to their much larger mass (and thus, inertia), the ions in the plasma respond very slowly compared to the laser pulse duration. Many of the ion dynamics occur as a result of fast electrons setting up large electrostatic fields via charge separation either at the front of the target, as they travel away from the laser interaction region, or at the rear surface, as they exit the target. Furthermore, many of the above processes are collisionless processes; the fast electrons they produce travel a significant distance away from the laser interaction region, and heating occurs mainly via Ohmic heating due to a return current. This is in contrast to long-pulse laser interactions which tend to be dominated by ion motion, with the electrons primarily acting as a means of transferring the laser energy to the plasma via collisional processes such as inverse bremsstrahlung, and heating occurring predominantly via thermal conduction [60].

It should be clear from the previous sections that the means by which short-pulse laser energy is delivered to a target, distributed amongst the constituent particles and transported into the target are complex and interdependent [61, 62]. For this reason analytic models are generally unable to encompass the full range of phenomena resulting from short-pulse laser interactions. A strong predictive computational modelling capability is essential in order to support, develop or challenge these models. Furthermore, such a capability enables the design of laser experiments to be optimised in order to make the best possible use of the capabilities of high power laser facilities such as OMEGA-EP, Orion and Vulcan-PW.

Computational models already exist which can accurately model the absorption of a short-pulse laser and production of energetic charged particles. Other methods also exist for modelling the transport of a pre-defined charged particle distribution and the related energy deposition in the bulk of a laser target. By coupling these approaches full-scale simulations of short-pulse laser experiments can be performed [30, 31]. Only recently, however, have algorithms been developed which make self-consistently modelling both the laser absorption and energy transport within a single simulation tractable [32]. These techniques (integrated modelling and novel high density PIC algorithms) are expected to become invaluable tools for modelling short-pulse laser experiments.

1.7 Short-pulse laser applications

1.7.1 Material heating

As discussed in sections 1.4 and 1.5, high power short-pulse lasers provide a means of delivering energy to material samples on time-scales shorter than the typical hydrodynamic response time. The associated rapid heating permits investigation into the properties of matter at high density (\sim solid) and high temperature (hundreds of eV) [5–9]. Using a separate long-pulse beam to compress the target prior to the short-pulse laser interaction further extends the upper limit on the range of densities at which high temperature material properties can be investigated.

1.7.2 X-ray sources

Some laser experiments make use of a bright x-ray source to backlight the experiment. This is achieved by fielding a separate target (typically a thin foil [21, 24] or wire [22, 23, 63]), which is hit by a short-pulse laser to produce x-rays via line emission (see section 1.5.4). The choice of backlighter material is an important consideration, as it should emit x-rays of sufficient energy that not all of the photons are stopped/scattered inside the target. Furthermore, the emission energy must not coincide with a peak in the material's absorption spectrum. For some experiments the x-ray energies produced by line emission are not sufficient, due to the desire to image very high density materials. In these cases high energy (\sim MeV) bremsstrahlung x-rays can be used for radiography [18–20, 25].

1.7.3 Electron acceleration

As a laser pulse propagates through underdense ($n_e < n_c$) plasma, it pondermotively drives electrons laterally out of the laser pulse, creating a wake of lateral plasma oscillations. As the laser pulse continues to propagate, these electrons are accelerated back toward the central axis of the laser pulse, and thus a longitudinal modulation in the electric field is created, trailing the laser pulse. This wakefield is maximised if the laser's pulse length is tuned so as to match half the plasma period ($\tau_L \approx \pi/\omega_p$). Electrons are then accelerated to very high energies (hundreds of MeV) over a few centimetres by 'surfing' the laser wakefield [64]. By exploiting this, and other similar phenomena, it is thought that it may be possible to use short-pulse lasers to accelerate electrons up to, and beyond, the energies currently achieved by conventional particle accelerators, and over much shorter distances (centimetres, rather than hundreds of metres) [65–67].

1.7.4 Proton acceleration

When the fast electrons produced by a short-pulse laser-plasma interaction reach the rear surface of the target, those with sufficient energy will leave the target. In doing so an electric field is set up perpendicular to the rear surface of the target. Protons initially present on the rear surface of the target (from water or oil surface contaminants) are accelerated to tens of MeV, over tens of microns, by this field [68–73]. This process is referred to as target normal sheath acceleration (TNSA).

Alternative methods for accelerating protons using high intensity short-pulse lasers have been proposed, such as the break-out afterburner (BOA) [74] and radiation pressure acceleration (RPA) [75, 76] approaches. These schemes aim to make use of the ponderomotive force of the laser pulse to accelerate thin ($< 1\ \mu\text{m}$) foil targets, either via the electrons (BOA) or by directly accelerating the ions (RPA).

Possible applications of laser-accelerated protons include cancer therapy [13, 14], radiography [15–17], material heating [28, 77–79], and material activation for nuclear physics [26].

1.7.5 Fast ignition

The pursuit of inertial confinement fusion (ICF) [40, 41] has led to the development of novel schemes to relax the stringent symmetry conditions imposed by the standard central hot spot ignition technique. The fast ignition approach [11] uses a high intensity, short-pulse laser to accelerate particles, usually electrons, up to energies such that they travel beyond the interaction region and deliver their energy to the core of a compressed fusion capsule, as shown in figure 1.3a. Ignition of the fuel capsule places constraints on the divergence, energy and duration of the fast electron beam.

Ablation of the capsule surface tends to result in the critical surface being located several hundred microns from the compressed core. Therefore a means is required by which the short-pulse laser can be allowed to interact as close to the anticipated hot spot as possible. One technique is to use a separate short-pulse beam, co-linear with the main pulse, and the ponderomotive hole-boring process [1] to carve a low density channel through the ablated coronal region, down which the main pulse can then propagate with minimal energy loss. An alternative concept, which has received a significant amount of attention, is to ensure a clear path for the ignition beam by having a high density cone (usually gold) inserted into the fuel capsule [12] as shown in figure 1.3b. The purpose of the cone being to prevent ablated material from the capsule from entering the path of the ignition pulse.




Figure 1.3: Diagrammatic representation of the fast ignition concept as originally envisaged (a) and with a gold cone insert (b). Reproduced with permission from Sircombe [80].

1.8 Overview of plasma modelling

Plasma physics modelling codes are widely employed in topics such as astrophysics, solar physics, magnetic and inertial confinement fusion, and laser-plasma interactions.

Detailed modelling of short-pulse (picosecond-scale) laser experiments requires capturing phenomena which occur on spatial and temporal scales which span several orders of magnitude (as illustrated by figure 1.4). The main pulse is typically preceded by an ASE (amplified spontaneous emission) pedestal, commonly referred to as a pre-pulse, of approximately nanosecond duration. Some experiments also make use of a separate long-pulse (nanosecond-scale) beam. These act to ablate the surface material creating a region of high temperature, low density plasma in front of the target (pre-plasma). Immediately behind the ablation front, the high thermal pressure results in a shock being driven into the target, compressing the material. Following the laser interaction the target undergoes hydrodynamic expansion and disassembly, while radiating heat (typically as x-rays), over several nanoseconds. In contrast, the fast electron transport and associated target heating typically occurs on time-scales comparable to the main short-pulse laser duration [81]: up to a few picoseconds, necessitating a separate, more detailed treatment than the other phases.

The mechanisms by which the short-pulse laser delivers its energy to the target, how this energy is distributed amongst the particles, and how it is transported deeper into the material are complex and interdependent. Often it is necessary to resolve spatial scales smaller than the laser wavelength and electron Debye length, and likewise




Figure 1.4: Visual representation of the disparate time-scales involved in short-pulse laser interactions.

temporal scales smaller than the laser period and the period of Langmuir waves in the plasma. Furthermore, it is essential that one accurately resolves the non-thermal nature of the electron distribution function. Key to short-pulse modelling is the kinetic nature of the plasma—but in treating this problem with such high fidelity it is not feasible to consider the much longer time-scale hydrodynamic phases.

Thus many disparate approaches to modelling laser-plasma physics phenomena are employed. These fall into two main categories: fluid and kinetic models.

1.8.1 Fluid models

Fluid models, in general, assume that the plasma at any point in space can be characterised by macroscopic quantities derived from moments of the particle distribution functions (namely, density, centre-of-mass velocity and temperature). Since each moment is a function of the next highest order moment, fluid models often invoke an equation of state, which relates the internal energy to the density and pressure, in order to obtain a closed set of equations. When attempting to model phenomena on large spatial and temporal scales this assumption is generally valid, and the magneto-hydrodynamic (MHD) equations [82] may be used to treat the plasma as a fluid moving under the

influence of electromagnetic fields. This makes hydrocodes powerful tools for simulating bulk behaviour during laser-plasma interactions. However, the fluid model is no longer valid when attempting to model phenomena which involve significant deviation from a Maxwellian particle distribution and/or motion which is dominated by non-linear wave-particle interactions, such as stimulated Raman scattering [83, 84], Landau damping [85], resonance absorption [43, 44], Brunel vacuum heating [45] and ponderomotive acceleration [46].

1.8.2 Kinetic models

The statistical treatment employed in kinetic models retains the individual particle properties, rather than integrating over the local distribution. The particles are considered to be evolving in time within a six-dimensional phase-space—three physical dimensions (x, y, z) and three momentum dimensions (p_x, p_y, p_z) . The evolution of the distribution function, f , of a particle species, in the presence of electric and magnetic fields, is governed by the plasma kinetic equation (see section 1.2.1):

$$\frac{\partial f}{\partial t} + \mathbf{v} \cdot \nabla f + \frac{q}{m} (\mathbf{E} + \mathbf{v} \times \mathbf{B}) \cdot \nabla_{\mathbf{v}} f = \left(\frac{\partial f}{\partial t} \right)_{\text{coll}}.$$

The additional physics embodied by kinetic codes carries with it an increase in the complexity of the codes. Thus the main limitation of kinetic codes is the high computational cost associated with modelling individual particle behaviour. This severely reduces the spatial and temporal scales that can be modelled efficiently compared with fluid methods.

1.8.3 Hybrid models

Occasionally both fluid and kinetic models are employed within the same code. These hybrid codes generally model an energetic particle species using a kinetic approach, and treat the background material (which is often at a lower temperature and higher number density) as a (multi-component) fluid. This approach has been used successfully to model energy transport in laser-plasma interactions [54, 86–88], but is unable to accurately model systems in which the kinetic and fluid populations cannot be treated as separate, for example the absorption of laser energy and production of fast electrons from a cold background material during short-pulse laser-plasma interactions.

1.8.4 Summary of plasma models

The choice of code to use when simulating plasma physics is dependent on the complexity of the system and the spatial and temporal scales of interest. Fluid models are useful in modelling the relatively slow, large-scale bulk motion of the plasma. In contrast, when modelling non-Maxwellian, non-linear phenomena a kinetic model is essential.

Chapter 2

Overview of the EPOCH PIC code

2.1 Particle-in-cell codes

The most rigorous approach to kinetic modelling is that employed by Vlasov codes, such as IMPACT [89], KALOS [90], VALIS [91], FIDO [92, 93] and OSHUN [94]. These simulate the plasma by solving the plasma kinetic equation directly, using a variety of approaches. For example, in Eulerian codes such as VALIS and FIDO, the particle distribution functions are evolved on an N -dimensional phase-space grid, where N is the number of spatial and momentum dimensions being simulated. Thus a 2D3P Eulerian Vlasov code requires a five-dimensional grid (two space and three momentum dimensions— x, y, p_x, p_y, p_z ; or $x, y, |\mathbf{p}|, \theta, \phi$ in the case of FIDO). The requirements of high dimensionality and high resolution in each dimension mean that such codes can be very computationally intensive. IMPACT also employs an Eulerian phase-space grid, but employs the diffusion approximation [89] to reduce the number of momentum dimensions required. Although this approach is well suited to simulating fast electron transport through dense plasma, it is not conducive to modelling oscillatory phenomena such as laser-plasma interactions. Codes such as KALOS and OSHUN, however, represent the distribution function in momentum space using an series expansion in spherical harmonics. This approach is also well suited to modelling fast electron transport. The use of spherical harmonics also reduces the computational expense associated with the distribution function representation in momentum space, but can suffer from difficulties with ensuring that mass is conserved, and that the distribution function is positive everywhere [90].

For most ‘everyday’ purposes a more general ‘workhorse’ code is required which can perform a broad range of kinetic simulations to an adequate level of precision and be adapted quickly for new problems, but without the intense computing requirements

of a Vlasov code. The most common solution is a particle-in-cell (PIC) code. PIC codes, unlike Vlasov codes, do not simulate a continuous distribution function. Instead they employ Monte-Carlo methods to randomly sample the plasma particles and apply a weight to the sampled particles in order to approximate the original distribution function [82, 95]. These pseudo-particles, in the collisionless limit, are free to move through each other and can be considered as representing clouds of charge and mass rather than point particles. It should be noted that the pseudo-particles have the same charge-to-mass ratio (and thus same behaviour in electromagnetic fields) as their constituent particles, but do not change their shape or undergo rotation under normal motion, nor do they have any internal degrees of freedom. This pseudo-particle method exploits Liouville theorem—the distribution function is constant along any trajectory in phase-space—by assuming that real particles which are initially close together in phase-space will remain close together throughout the duration of the simulation.

The use of non-uniform particle weighting allows PIC codes to simulate systems in which the density can vary by several orders of magnitude without requiring much larger numbers of particles in more dense regions of the simulation domain.

The motion of the particles is treated using a Lagrangian approach, which allows them to move freely in position and momentum space. Meanwhile the moments of the distribution function (e.g. density) and electromagnetic field variables are calculated on a fixed, Eulerian grid.

The key features of a PIC code are discussed below, using EPOCH [29] as an example.

2.1.1 Explicit vs implicit schemes

In explicit schemes the new state at a given point depends on the current state at that point and its neighbours: $A_k^{n+1} = f(A_k^n, A_{k-1}^n, A_{k+1}^n)$ (see also figure 2.1). A simple example of an explicit scheme is the forward-time centered-space (FTCS) scheme [96]:

$$\frac{f_i^{n+1} - f_i^n}{\Delta t} + u \frac{f_{i+1}^n - f_{i-1}^n}{2\Delta x} = 0$$

Such schemes have a finite numerical information propagation speed: $c_n = \Delta x / \Delta t$. In order for an explicit scheme to be stable it usually requires that the numerical information propagation speed be greater than the maximum speed at which information will propagate through the system (for an electromagnetic field solver, the speed of light in vacuum). This condition is commonly referred to as the CFL limit [97]. For a one-dimensional system this is $\Delta x / \Delta t > v_{\max}$.

In contrast, the new state at a given point using an implicit scheme is




Figure 2.1: Diagrammatic representation of the dependencies for explicit (left) and implicit (right) methods. When using an explicit method, the updated state is calculated from the previous state at the same and neighbouring points. The implicit method, however, uses the previous state at the same point and the updated states of its neighbours.

dependent upon the current state at that point and the neighbouring points' new state: $A_k^{n+1} = f(A_k^n, A_{k-1}^{n+1}, A_{k+1}^{n+1})$, resulting in an infinite numerical information propagation speed. The implicit equivalent of the FTCS scheme is the backward-time centered-space (BTCS) scheme [96]:

$$\frac{f_i^{n+1} - f_i^n}{\Delta t} + u \frac{f_{i+1}^{n+1} - f_{i-1}^{n+1}}{2\Delta x} = 0$$

Such schemes tend to be unconditionally stable, permitting arbitrarily large time-steps. This helps to offset the higher computational cost per time-step associated with an implicit method compared with an explicit approach. However, since implicit schemes tend to require a large amount of inter-processor communication, they do not scale well to large numbers of processors on massively parallel computational platforms compared with explicit schemes.

In order to accurately capture the evolution of the distribution function, it is generally desirable to resolve high frequency phenomena, such as Langmuir waves. Thus, for the purposes of modelling kinetic laser-plasma interaction phenomena, explicit schemes tend to be preferable to implicit methods.

2.2 Particle propagation

The pseudo-particles in a PIC code are advanced along the characteristics of Vlasov's equation (equation 1.2, with $(\partial f_j / \partial t)_{\text{coll}} = 0$) in response to the Lorentz force, often

using a leapfrog scheme:

$$\mathbf{x}^{n+1/2} = \mathbf{x}^{n-1/2} + \frac{\Delta t}{m\gamma^n} \mathbf{p}^n, \quad (2.1)$$

$$\mathbf{p}^{n+1} = \mathbf{p}^n + q\Delta t \left(\mathbf{E}^{n+1/2} + \frac{1}{m\gamma^{n+1/2}} \mathbf{p}^{n+1/2} \times \mathbf{B}^{n+1/2} \right). \quad (2.2)$$

In EPOCH equation 2.1 is split into two half-updates across the time-step in order to ensure that the particle positions and momenta are both known at the end of each step ($n + 1$):

$$\begin{aligned} \mathbf{x}^{n+1/2} &= \mathbf{x}^n + \frac{\Delta t}{2m\gamma^n} \mathbf{p}^n, \\ \mathbf{x}^{n+1} &= \mathbf{x}^{n+1/2} + \frac{\Delta t}{2m\gamma^{n+1}} \mathbf{p}^{n+1}. \end{aligned}$$

EPOCH evaluates the momentum update using the ‘Boris’ algorithm [95] to decouple the electric and magnetic field effects:

$$\mathbf{p}^- = \mathbf{p}^n + \frac{q\Delta t}{2} \mathbf{E}^{n+1/2}, \quad (2.3)$$

$$\mathbf{p}^+ = \mathbf{p}^- + \frac{q\Delta t}{\gamma^{n+1/2}m} \left(\frac{\mathbf{p}^+ + \mathbf{p}^-}{2} \times \mathbf{B}^{n+1/2} \right), \quad (2.4)$$

$$\mathbf{p}^{n+1} = \mathbf{p}^+ + \frac{q\Delta t}{2} \mathbf{E}^{n+1/2}, \quad (2.5)$$

where it has been assumed that $\mathbf{p}^{n+1/2} = (\mathbf{p}^+ + \mathbf{p}^-)/2$ in equation 2.4.

Although equation 2.4 is a second order implicit update, it can be solved explicitly by noting that the magnetic force is conservative. This can be shown by taking the dot product of equation 2.4 with $\mathbf{p}^+ + \mathbf{p}^-$:

$$\begin{aligned} \mathbf{p}^+ \cdot (\mathbf{p}^+ + \mathbf{p}^-) &= \mathbf{p}^- \cdot (\mathbf{p}^+ + \mathbf{p}^-) \\ &\quad + \frac{q\Delta t}{2\gamma^{n+1/2}m} \left([\mathbf{p}^+ + \mathbf{p}^-] \times \mathbf{B}^{n+1/2} \right) \cdot (\mathbf{p}^+ + \mathbf{p}^-), \\ (\mathbf{p}^+)^2 + \mathbf{p}^+ \cdot \mathbf{p}^- &= \mathbf{p}^- \cdot \mathbf{p}^+ + (\mathbf{p}^-)^2, \\ \therefore |\mathbf{p}^+| &= |\mathbf{p}^-|. \end{aligned}$$

Since, the update due to the magnetic field is conservative, equation 2.4 can be replaced with a rotation of the particle’s momentum vector [95]:

$$\mathbf{p}^+ = \hat{A} \mathbf{p}^-, \quad (2.6)$$

where

$$\hat{A} = \begin{pmatrix} 1 + \tau_x^2 - \tau_y^2 - \tau_z^2 & 2\tau_x\tau_y + 2\tau_z & 2\tau_x\tau_z - 2\tau_y \\ 2\tau_x\tau_y - 2\tau_z & 1 - \tau_x^2 + \tau_y^2 - \tau_z^2 & 2\tau_y\tau_z + 2\tau_x \\ 2\tau_x\tau_z + 2\tau_y & 2\tau_y\tau_z - 2\tau_x & 1 - \tau_x^2 - \tau_y^2 + \tau_z^2 \end{pmatrix} \left(\frac{1}{1 + \boldsymbol{\tau} \cdot \boldsymbol{\tau}} \right),$$

and

$$\boldsymbol{\tau} = \frac{q\Delta t}{2} \frac{c\mathbf{B}^{n+1/2}}{\sqrt{m^2c^2 + \mathbf{p}^- \cdot \mathbf{p}^-}}.$$

2.2.1 Particle shapes

Since the electric and magnetic fields required to evaluate equations 2.3–2.5 are calculated on a discrete grid, it is necessary to interpolate the fields to the particle positions, and vice-versa to obtain properties such as density, temperature and current density. Using zeroth-order weighting, the fields at the nearest grid point to the particle are assigned to the particle position. This, however, implies that the particles are in fact uniform regions of charge of size $\Delta x \times \Delta y \times \Delta z$, as shown by figure 2.2. The step-wise changes in density as particles enter and leave cells can introduce noise in the current density, and thus electric field (via Ampère’s law—see below). This can be mitigated by the use of first-order shape functions, which employ linear interpolation from the grid points surrounding the cell in which the particle is currently located. This is equivalent to convolving zeroth-order weighting with itself—i.e. assuming the particles are regions of charge of size $\Delta x \times \Delta y \times \Delta z$, and performing zero-th order weighting on each portion of the particle. In 1D, the electric fields interpolated from the surrounding grid points (X_j, X_{j+1}), at the particle’s position (x_i), using first-order weighting are given by:

$$\mathbf{E}(x_i) = \left(\frac{X_{j+1} - x_i}{\Delta x} \right) \mathbf{E}(X_j) + \left(\frac{x_i - X_j}{\Delta x} \right) \mathbf{E}(X_{j+1}).$$

This results in a ‘triangular’ particle shape of size $2\Delta x \times 2\Delta y \times 2\Delta z$ (see figure 2.3). Higher order weighting schemes (such as cubic splines, which imply particles of size $4\Delta x \times 4\Delta y \times 4\Delta z$) are often used to further reduce the noise levels in PIC simulations. As with first-order weighting, the shape functions at higher orders can be obtained by repeated convolutions of zeroth-order weighting with itself.




Figure 2.2: Diagrammatic representation of zeroth-order particle weighting. The fields from the nearest grid point are applied directly to the particle, and vice-versa when calculating properties such as density on the grid. This implies a rectangular particle shape of width Δx in 1D.

2.2.2 Current density calculation

In order to ensure that Gauss's law ($\nabla \cdot \mathbf{E} = \rho/\epsilon_0$) is satisfied, the current density using in the electric field update (see section 2.3) must satisfy the charge continuity equation:

$$\frac{\partial \rho}{\partial t} + \nabla \cdot \mathbf{J} = 0.$$

EPOCH uses the 'density decomposition' scheme developed by Esirkepov [98], which is a generalisation of the Villasenor and Buneman [99] current deposition scheme, to deposit the particle currents on the mesh. After updating the particle positions and momenta, an additional half-step in position is performed: $\mathbf{x}^{n+3/2} = \mathbf{x}^{n+1} + \Delta t \mathbf{p}^{n+1}/2m\gamma^{n+1}$. The fraction of each cell which passes through each cell face as a result of this update, combined with the particle weights and choice of particle shape function (discussed above) allows for the components of the current density, \mathbf{J}^{n+1} , to be accumulated onto the grid.




Figure 2.3: Diagrammatic representation of first-order particle weighting. The fields from the surrounding grid points are linearly interpolated to the particle, and vice-versa when calculating properties such as density on the grid. This is equivalent to performing zeroth-order weighting to a uniform region of charge and mass. This implies a triangular particle shape of width $2\Delta x$ in 1D.

2.3 Field propagation

Evaluating the electric and magnetic fields on a discrete grid, and over discrete time-steps, makes the finite difference approach a natural choice. A commonly used technique for computationally modelling electromagnetic phenomena is the finite-difference-time-domain (FDTD) method [100]. However other schemes are equally valid.

Field values are updated using the time-dependent Maxwell's equations:

$$\frac{\partial \mathbf{E}}{\partial t} = c^2 (\nabla \times \mathbf{B}) - \frac{\mathbf{J}}{\epsilon_0}, \quad (2.7)$$

$$\frac{\partial \mathbf{B}}{\partial t} = -\nabla \times \mathbf{E}. \quad (2.8)$$

The FDTD method is performed in such a way that provided $\nabla \cdot \mathbf{B} = 0$ is true at the start of the simulation, it will remain valid throughout. Furthermore, Gauss's law does not need to be solved since the charge conserving scheme for calculating the current




Figure 2.4: An example of the 3D Yee mesh for the finite-difference-time-domain method. The arrows indicate the positions at which the fields are defined. Grid-based particle properties, such as density and temperature, are defined at the centre of the cell. The components of the current densities are defined at the same positions as the corresponding electric field components.

densities ensures that $\nabla \cdot \mathbf{E} = \rho/\epsilon_0$ also holds throughout the simulation (see section 2.2.2).

If the spatial derivatives are discretized on a standard grid, there is a risk of a ‘chequerboard’ instability developing, in which two competing solutions to Maxwell’s equations develop on adjacent grid cells. This instability can arise if the gradients at each point are not a function of the value in that cell (e.g. $\partial F/\partial x|_j = (F_{j+1} - F_{j-1})/2\Delta x$). This can be avoided through careful choice of the differencing scheme. Alternatively, it can be noted that the only gradient required to update the electric fields is in the magnetic fields, and vice-versa. Thus, the calculation of the field gradients can be simplified, and the chequerboard instability avoided by defining the electric and magnetic fields at points which are staggered relative to the grid (i.e. a Yee mesh [101]). An example of such a staggered mesh in 3D is presented in figure 2.4.

Equations 2.3–2.5 require knowledge of the electric and magnetic fields at half time-step intervals (this is displayed diagrammatically in figure 2.5). The half time-step

field values are obtained using the discretised Maxwell's equations as follows:

$$\mathbf{E}_{jkl}^{n+1/2} = \frac{c^2 \Delta t}{2} (\nabla^- \times \mathbf{B}_{jkl}^n - \mu_0 \mathbf{J}_{jkl}^n) + \mathbf{E}_{jkl}^n, \quad (2.9)$$

$$\mathbf{B}_{jkl}^{n+1/2} = -\frac{\Delta t}{2} (\nabla^+ \times \mathbf{E}_{jkl}^{n+1/2}) + \mathbf{B}_{jkl}^n. \quad (2.10)$$

Following the particle update, as discussed in section 2.2, the fields are advanced to the full time-step (i.e. from $n + 1/2$ to $n + 1$) using the current densities calculated at step $n + 1$:

$$\mathbf{B}_{jkl}^{n+1} = -\frac{\Delta t}{2} (\nabla^+ \times \mathbf{E}_{jkl}^{n+1/2}) + \mathbf{B}_{jkl}^{n+1/2}, \quad (2.11)$$

$$\mathbf{E}_{jkl}^{n+1} = \frac{c^2 \Delta t}{2} (\nabla^- \times \mathbf{B}_{jkl}^{n+1} - \mu_0 \mathbf{J}_{jkl}^{n+1}) + \mathbf{E}_{jkl}^{n+1/2}. \quad (2.12)$$

In equations 2.9–2.12, ∇^- and ∇^+ are defined as:

$$\begin{aligned} \nabla^- F_{jkl}^n &= \left(\frac{F_{jkl}^n - F_{j-1kl}^n}{\Delta x}, \frac{F_{jkl}^n - F_{jk-1l}^n}{\Delta y}, \frac{F_{jkl}^n - F_{jkl-1}^n}{\Delta z} \right), \\ \nabla^+ F_{jkl}^n &= \left(\frac{F_{j+1kl}^n - F_{jkl}^n}{\Delta x}, \frac{F_{jk+1l}^n - F_{jkl}^n}{\Delta y}, \frac{F_{jkl+1}^n - F_{jkl}^n}{\Delta z} \right), \end{aligned}$$

and the subscripts denote the grid point indices and the superscripts the time-step index. Thus in 3D:

$$\nabla^- \times \mathbf{B}_{jkl}^n = \begin{pmatrix} \frac{(B_z)_{jkl}^n - (B_z)_{jk-1l}^n}{\Delta y} - \frac{(B_y)_{jkl}^n - (B_y)_{jkl-1}^n}{\Delta z} \\ \frac{(B_x)_{jkl}^n - (B_x)_{jkl-1}^n}{\Delta z} - \frac{(B_z)_{jkl}^n - (B_z)_{j-1kl}^n}{\Delta x} \\ \frac{(B_y)_{jkl}^n - (B_y)_{j-1kl}^n}{\Delta x} - \frac{(B_x)_{jkl}^n - (B_x)_{jk-1l}^n}{\Delta y} \end{pmatrix},$$

and

$$\nabla^+ \times \mathbf{E}_{jkl}^n = \begin{pmatrix} \frac{(E_z)_{jk+1l}^n - (E_z)_{jkl}^n}{\Delta y} - \frac{(E_y)_{jkl+1}^n - (E_y)_{jkl}^n}{\Delta z} \\ \frac{(E_x)_{jkl+1}^n - (E_x)_{jkl}^n}{\Delta z} - \frac{(E_z)_{j+1kl}^n - (E_z)_{jkl}^n}{\Delta x} \\ \frac{(E_y)_{j+1kl}^n - (E_y)_{jkl}^n}{\Delta x} - \frac{(E_x)_{jk+1l}^n - (E_x)_{jkl}^n}{\Delta y} \end{pmatrix}.$$

The FDTD method is time-centred and all values are known at the beginning and end of each time-step. This provides the method with an advantage over more basic leapfrog schemes [82, 95] in that the time-step duration can be changed dynamically if necessary.

The full time-stepping algorithm is detailed in figure 2.5. The ordering (shown by the numbering) is important in ensuring that the scheme is both stable and accurate in space and time to at least second order.




Figure 2.5: Diagram to illustrate how the particle position (x), particle momentum (p), electric fields (E) and magnetic fields (B) are used to propagate each other over the course of a single time-step.

2.4 Boundary conditions

In order for the computational domain to reside in memory, it must be constrained to a finite region. The imposition of artificial boundaries necessitates careful consideration of the particles and fields at the boundaries of the domain.

The simplest, and most frequently used, boundary conditions for the particles are reflecting and periodic boundaries. The reflecting boundary condition acts to specularly reflect any particles which attempt to pass through it. Thus for a particle whose x position is less than x_{\min} :

$$\begin{aligned} x &\rightarrow 2x_{\min} - x, \\ p_x &\rightarrow -p_x. \end{aligned}$$

For a particle with $x > x_{\max}$:

$$\begin{aligned} x &\rightarrow 2x_{\max} - x, \\ p_x &\rightarrow -p_x. \end{aligned}$$

Similar conditions can be applied for the y and z directions.

For periodic boundaries, the particle boundary conditions are:

$$\text{If } x < x_{\min} : x \rightarrow x + x_{\max} - x_{\min} + \Delta x,$$

$$\text{If } x > x_{\max} : x \rightarrow x - x_{\max} + x_{\min} - \Delta x.$$

In addition to the standard reflecting and periodic boundary conditions, EPOCH also includes an option for 'thermal' boundaries. In this case particles which attempt to pass through the boundary are reinjected into the simulation, with their momentum resampled based on an assumed Maxwellian distribution with a temperature equal to the initial value at the boundary. Thus for $x < x_{\min}$:

$$p_x \rightarrow |p'_x(T)|,$$

$$p_y \rightarrow p'_y(T),$$

$$p_z \rightarrow p'_z(T),$$

and $x > x_{\max}$:

$$p_x \rightarrow -|p'_x(T)|,$$

$$p_y \rightarrow p'_y(T),$$

$$p_z \rightarrow p'_z(T).$$

A region of 'ghost cells' is generally added to the boundaries of the domain to allow the field gradients to be calculated for cells close to the boundary. The value of the fields in these cells, however, is determined through the choice of field boundary condition. Often the fields at the boundary are fixed to zero:

$$F_{0jk} = 0,$$

$$F_{-1jk} = -F_{1jk},$$

$$F_{-2jk} = -F_{2jk},$$

$$F_{n_x+1jk} = 0,$$

$$F_{n_x+2jk} = -F_{n_xjk},$$

$$F_{n_x+3jk} = -F_{n_x-1jk},$$

or varied in time so as to ensure that the field gradient across the boundary is zero:

$$\begin{aligned}
F_{0jk} &= F_{1jk}, \\
F_{-1jk} &= F_{2jk}, \\
F_{-2jk} &= F_{3jk}, \\
F_{n_x+1jk} &= F_{n_xjk}, \\
F_{n_x+2jk} &= F_{n_x-1jk}, \\
F_{n_x+3jk} &= F_{n_x-2jk}.
\end{aligned}$$

Periodic field boundary conditions can also be applied as follows:

$$\begin{aligned}
F_{0jk} &= F_{n_xjk}, \\
F_{-1jk} &= F_{n_x-1jk}, \\
F_{-2jk} &= F_{n_x-2jk}, \\
F_{n_x+1jk} &= F_{1jk}, \\
F_{n_x+2jk} &= F_{2jk}, \\
F_{n_x+3jk} &= F_{3jk}.
\end{aligned}$$

Corrections to both sets of field boundary conditions must be applied to take account of the staggering of the electric and magnetic fields.

While many PIC simulations employ reflecting and/or periodic boundaries, the laser package in EPOCH also allows for the possibility of open boundaries which allow particles and fields to leave the system. Particles passing through an open boundary are removed from memory, and thus lost from the simulation. The field boundary conditions associated with open boundaries are discussed below.

2.5 The laser package

The laser package in EPOCH adds two boundary conditions: `simple_laser` and `simple_outflow`. These conditions can be thought of as radiating and absorbing boundaries respectively. The conditions for a radiating boundary at $x = x_{\min}$ are [102]

(see appendix B):

$$\begin{aligned}
(B_z)_{0kl}^n &= \left[4(F_{yz})_{1kl}^n - 2(E_y)_{1kl}^{n-1/2} - L_z ((B_x)_{1kl}^n - (B_x)_{1kl-1}^n) \right. \\
&\quad \left. + \frac{\Delta t}{\epsilon_0} (J_y)_{1kl}^n + (L_x - c) (B_z)_{1kl}^n \right] \left(\frac{1}{L_x + c} \right), \\
(B_y)_{0kl}^n &= \left[4(F_{zy})_{1kl}^n - 2(E_z)_{1kl}^{n-1/2} - L_y ((B_x)_{1kl}^n - (B_x)_{1kl-1}^n) \right. \\
&\quad \left. + \frac{\Delta t}{\epsilon_0} (J_z)_{1kl}^n + (L_x - c) (B_y)_{1kl}^n \right] \left(\frac{1}{L_x + c} \right),
\end{aligned}$$

where F_{yz} and F_{zy} are the contributions to the electric field from incoming laser pulses polarised in the y and z directions respectively, $L_x = c\Delta t/\Delta x$ and $L_z = c\Delta t/\Delta z$. The values of F_{yz} and F_{zy} are defined by the spatial and temporal profiles of the amplitude of the laser's electric field, as well as the polarisation.

The equivalent radiating boundary conditions at the x_{\max} boundary are:

$$\begin{aligned}
(B_z)_{N_x kl}^n &= \left[-4(F_{yz})_{N_x kl}^n + 2(E_y)_{N_x kl}^{n-1/2} + L_z ((B_x)_{N_x kl}^n - (B_x)_{N_x kl-1}^n) \right. \\
&\quad \left. - \frac{\Delta t}{\epsilon_0} (J_y)_{N_x kl}^n + (L_x - c) (B_z)_{N_x-1kl}^n \right] \left(\frac{1}{L_x + c} \right), \\
(B_y)_{N_x kl}^n &= \left[-4(F_{zy})_{N_x kl}^n + 2(E_z)_{N_x kl}^{n-1/2} + L_y ((B_x)_{N_x kl}^n - (B_x)_{N_x kl-1}^n) \right. \\
&\quad \left. - \frac{\Delta t}{\epsilon_0} (J_z)_{N_x kl}^n + (L_x - c) (B_y)_{N_x-1kl}^n \right] \left(\frac{1}{L_x + c} \right),
\end{aligned}$$

where N_x is the number of grid cells along the x -axis.

Absorbing boundary conditions can be simulated by setting $F_{yz} = F_{zy} = 0$ at the desired boundary.

Similar radiating and absorbing boundary conditions can be derived for the y and z -boundaries by following the derivation given in appendix B.

2.6 The limits of PIC

The strength of PIC codes lies in the simplicity of the core algorithm. This makes them stable, flexible and applicable to a wide range of problems. However, there are several limitations and potential pitfalls that users must bear in mind.

Even though they are more intuitive than Vlasov codes, PIC codes still require a sizable amount of computational effort (100s to 1000s of CPU hours) in comparison with codes employing a fluid or hybrid approach. This can place limits on the size of

the system and the period of time that can be modelled. For example, 2D simulations of a short-pulse laser incident on a high density target are typically limited to tens of microns surrounding the interaction region and a few picoseconds. Reducing the grid resolution or number of pseudo-particles can relieve some of the computational strain, but carry drawbacks, which are discussed below. Thus running a PIC code, as with any code, requires one to strike a balance between fidelity and practicality.

2.6.1 Grid resolution

Finely resolved simulation grids in PIC codes are very important. By evaluating the field values on a discrete grid PIC codes place a lower bound on the spatial scale over which electromagnetic phenomena can be resolved. The smallest length-scale that it is generally desirable to resolve is the electron Debye length, λ_D . However, as λ_D is often several orders of magnitude smaller than the size of the system to be simulated, resolving the Debye length is not always appropriate. Failure to resolve the Debye length tends to lead to self-heating of the plasma; an undesirable, but often unavoidable, effect at high densities [103].

An order-of-magnitude estimate of the self-heating rate for the PIC code EPOCH is given by Arber et al. [29]:

$$\frac{\partial T_{\text{eV}}}{\partial t_{\text{ps}}} = \alpha_H \frac{n_{29}^{3/2} \Delta x_{\text{nm}}^2}{N_{\text{ppc}}}, \quad (2.13)$$

where T_{eV} is the temperature in eV, t is the time in picoseconds, n_{29} is the electron number density in units of 10^{29} m^{-3} , Δx_{nm} is the grid spacing in nanometres, and N_{ppc} is the number of pseudo-particles per cell. The value of α_H depends upon the particle shape function (see section 2.2.1), and whether current smoothing is applied. Typical values for α_H are given in table 1 of Arber et al. [29].

Numerical self-heating can be mitigated by employing implicit methods, rather than the explicit methods discussed so far.

In order to correctly simulate the behaviour of an incident laser pulse it is vital that the code is able to resolve the laser wavelength. Any electromagnetic waves with a wavelength of the order, or less than, a few grid cells will not propagate correctly and will rapidly dissipate.

It should be noted that the use of a discrete grid carries with it many risks, and if the user is not careful can lead to the emergence of unphysical phenomena. Any periodic behaviour or features observed at or close to the grid-scale during a simulation should be treated with caution.

2.6.2 Particle resolution

In simulations with large density ranges, spanning many orders of magnitude, uniform particle weighting will require very high particles numbers. To overcome this variable particle weights are often employed: Each grid cell is populated with a fixed number of pseudo-particles, each individually weighted so as to produce the desired density in the cell. This approach has the limitation, however, that the particle distribution function is often poorly resolved in the highest density cells, threatening the code's ability to treat physical processes in such cells correctly. In simulations with steep density gradients, this can result in the generation of large levels of noise in the field variables due to the rapid increase in density when small numbers of 'heavy' particles drift into lower density cells.

The desire to have each pseudo-particle represent as few real particles as possible, to reduce the computational expense, means that not only is it difficult to accurately model large systems ($> 10^4 \mu\text{m}^2$ in 2D), but also those containing regions at or near solid density, using standard PIC techniques. PIC codes can generally cope with electron densities up to about 10^{29} m^{-3} ($\sim 1/3$ of solid density for plastic). Simulations of higher density materials are often performed with a reduced peak density, and the assumption that the results will be a weak function of density (i.e. that the system is dominated by collisionless effects).

It has recently been noted that since the stopping of relativistic particles travelling through a cold background plasma scales as q^2/m , pseudo-particles with a large particle weight will stop more rapidly than lower weight particles with the same charge-to-mass ratio [104]. This enhanced stopping is associated with the wakefields generated by the energetic particles, and can be mitigated by reducing the particle weights (more pseudo-particles per grid cell) and/or employing a pseudo-particle splitting scheme for high energy particles [104, 105].

2.6.3 Time-step constraints

The Courant-Friedrichs-Lewy (or CFL) condition [97] described in section 2.1.1 sets the condition for which the algorithmic solutions to partial differential equations on a discrete grid will be stable when employing an explicit field solver.

For the case of a PIC code with a full electromagnetic field solver, the time-step is therefore constrained by the CFL limit. However, if the time-step is too short relative to the grid scale, the phase velocity at short wavelengths is less than the speed of light, as shown by figure 2.6. This can also potentially lead to a situation where charged particles are able to travel faster than the code-calculated phase velocity (since particles




Figure 2.6: The dispersion relations for an electromagnetic wave propagating on a discrete, 1D mesh, with varying time-step durations [95].

are not restricted by the grid). This in turn can lead to the growth of short wavelength modes via numerical Cerenkov radiation [106].

Consequently, for a wave crossing a discrete grid, the length of a time-step must be less than the time taken for the wave to travel between adjacent grid points. Therefore, increasing the resolution of the grid also imposes a more stringent upper limit on the length of each time-step.

In addition, it is often necessary to resolve the plasma (and laser) period to accurately resolve kinetic plasma phenomena. Thus, if $\omega_p^{-1} < \Delta x/c$, the time-step will be limited by the plasma period, rather than the CFL limit. If the mesh resolution is expressed as a multiple of the minimum Debye length in the simulation ($\Delta x = N\lambda_D$), the condition under which the time-step is determined by the plasma period can be written as:

$$k_B T_e > \frac{m_e c^2}{N^2}.$$

This indicates that the plasma-frequency-based limit is only required if N is very large

(i.e. $\Delta x \gg \lambda_D$). However, under these conditions, PIC codes are susceptible to numerical self-heating as discussed previously in section 2.6.1.

2.7 PIC summary

Provided the user is aware of the potential pitfalls above, PIC codes are a valuable tool for modelling kinetic plasma behaviour. They are generally robust, flexible, reliable and can perform detailed simulations of many fundamental plasma physics phenomena. PIC codes can be, and have been, applied to a wide range of plasma physics problems, making them the ideal tool for use in fields as diverse as solar physics [107, 108], astrophysics [109], magnetic confinement fusion [110] and laser-plasma interaction [111], where they can model wave-particle interactions [112], charged particle acceleration [66], parametric instabilities [113, 114], and various other non-linear phenomena which are absent from fluid models.

It should be noted, however, that a standard PIC codes will only solve the Vlasov equation (plasma kinetic equation with the collisional term assumed to be zero):

$$\frac{\partial f_j}{\partial t} + \frac{\mathbf{u}}{\gamma} \cdot \nabla f_j + \frac{q_j}{m_j} (\mathbf{E}_0 + \mathbf{v} \times \mathbf{B}_0) \cdot \nabla_{\mathbf{v}} f_j = 0.$$

In order to obtain a more accurate description of the plasma it is necessary to implement an algorithm which attempts to model the contribution from discrete particle collisions.

2.8 Implementing short-range collisions in EPOCH

2.8.1 Introduction

The particle-in-cell (PIC) method is a robust and reliable approach to the kinetic modelling of plasmas which are dominated by collective effects. However, these codes generally neglect particle interactions over very short (less than grid-scale) ranges. At high temperatures ($\gtrsim 1$ keV) and relatively low densities ($\lesssim 10^{27} \text{ m}^{-3}$) collisional effects in plasma are generally considered minimal (i.e., the mean time between collisions is comparable to the time-scales of interest¹), and the collisionless approximation used in PIC codes is valid. However, at lower temperatures and/or higher densities the effect of sub-grid-scale particle-particle interactions on the evolution of the system can become non-negligible.

¹For example, the electron-ion collision time for a 1 keV hydrogen plasma at a density of $n_e = 10^{27} \text{ m}^{-3}$ is of the order 1 ps.

As a means of accounting for discrete particle collisions, many PIC codes make use of some form of collision algorithm [115–122] which stochastically scatters particles in momentum space. The majority of algorithms in use are derived from Takizuka and Abe’s binary collision approach [115] which performs Rutherford scattering on pairs of particles which reside in the same grid cell to solve the Landau equation (equation A.5 in appendix A.1), and is equivalent to the Fokker-Planck approach [123]. Alternative methods using, for example, an additional force term in the Lorentz force calculation (a ‘collision field’) [116] have also been developed.

2.8.2 Cumulative relativistic scattering of binary particle pairs

A collision algorithm, based on the approach of Sentoku and Kemp [121], has been implemented in the 1D version of EPOCH, and subsequently ported to the 2D and 3D versions of the code [29]. Here it is assumed, without loss of generality, that a particle i of species α is being scattered off of a particle j of species β (with $\alpha = \beta$ a possibility). Unless otherwise specified, SI units should be assumed.

Collisions in EPOCH are assumed to be elastic. Therefore, in order to simplify momentum conservation, collisions are calculated in the centre-of-momentum reference frame of the two particles. Lorentz transformations are employed in order to evaluate the particles’ momenta in the centre-of-momentum frame.

It should be noted that although the collision algorithm is not fully relativistic, it can still be considered valid for weakly-relativistic particles, where collisional effects may be important. An extension to approach which is valid in the ultra-relativistic limit has been developed by Peano et al. [124], but is not included in EPOCH since highly relativistic particles can generally be considered to be collisionless.

Collision frequency and Coulomb logarithm

An expression based on the approach of Manheimer et al. [125] can be used to calculate the collision frequency in EPOCH. The ‘Manheimer collision frequency’ matches the low and high energy limits from the NRL Plasma Formulary [126], and varies smoothly in between:

$$\nu = \frac{\nu_s}{1 + \frac{\nu_s}{\nu_f}}, \quad (2.14)$$

where:

$$\begin{aligned} \nu_s &= 0.23 (\mu/T_\beta)^{3/2} n_\beta \ln \Lambda Z_j^2, \\ \nu_f &= 3.9 \times 10^{-6} K_r^{-3/2} n_\beta \ln \Lambda Z_j^2, \end{aligned}$$

are the slow and fast limits of the collision frequency given by Huba [126]. T_β is the local temperature (in eV), μ the mass (in units of proton mass) and n_β the number density (in cm^{-3}) of the particles of species β . K_r (also in eV) is the kinetic energy of particle i in the rest frame of particle j .

Although the Manheimer approach allows for a more accurate treatment of particle collisions at low velocities, it is only valid in its current form for electrons as species α (since the expressions for ν_s and ν_f are the limits of the electron-ion collision frequency). For most applications, however, the standard v^{-3} expression (see section C.1) is sufficient:

$$\nu = \frac{q_i^2 q_j^2 n_\beta \ln \Lambda}{4\pi \epsilon_0^2 \mu^2 v_r^3}. \quad (2.15)$$

Note that here $\mu = (m_i m_j) / (m_i + m_j)$ and v_r is the relative velocity between particles i and j .

The Coulomb logarithm ($\ln \Lambda$) is calculated using the standard definition of the natural logarithm of the ratio between the upper and lower limits for the distance of closest approach between the two colliding particles [60, 126]. The lower limit is determined by the maximum of the Landau length, $q_\alpha q_\beta / 12\pi \epsilon_0 k_B T_\alpha$, and the de Broglie wavelength, h/p . Thus, for slow particles, the Coulomb logarithm can be expressed as:

$$\ln \Lambda_s = \ln \left(\frac{12\pi (\epsilon_0 k_B T_\alpha)^{3/2}}{q_\alpha^2 q_\beta \sqrt{n_\alpha}} \right), \quad (2.16)$$

where T_α and n_α are the local temperature (here in K) and local number density (in m^{-3}) of species α . For fast-moving particles, however, the Coulomb logarithm takes the form:

$$\ln \Lambda_f = \ln \left(\frac{m_\alpha c}{q_\alpha h} \sqrt{\frac{\epsilon_0 k_B T_\alpha (\gamma_\alpha^2 - 1)}{n_\alpha}} \right), \quad (2.17)$$

where $\gamma_\alpha = 1 + k_B T_\alpha / m_\alpha c^2$.

Scattering angles

The scattering angles for particle i in the centre-of-momentum frame are defined as shown in figure 2.7: θ being the angle within the scattering plane relative to particle i 's initial momentum vector, and ϕ as the angle in the plane perpendicular to the particle's initial momentum vector. Since scattering is a stochastic process, θ is assumed to have a Gaussian distribution by the central limit theorem. Consequently, for an individual scattering event, $\theta \in [-\pi, \pi)$ is randomly sampled from a Gaussian distribution of width $\langle \theta^2 \rangle \approx \langle \tan^2 \theta \rangle = 2\nu \Delta t$ (see section C.2) using a Box-Muller transform [127].




Figure 2.7: Diagram to illustrate the scattering angles, θ and ϕ , that particle i (red) is scattered through when colliding with particle j (blue) in the centre-of-momentum frame.

Meanwhile, $\phi \in (0, \pi]$ is randomly selected from a uniform distribution.

The post-collision momentum of particle i in the centre-of-momentum frame can subsequently be calculated:

$$\mathbf{p}'_i = |\mathbf{p}_i| (\hat{\mathbf{e}}_1 \cos \theta + \hat{\mathbf{e}}_2 \sin \theta \cos \phi + \hat{\mathbf{e}}_3 \sin \theta \sin \phi),$$

and since momentum must be conserved, $\mathbf{p}'_j = -\mathbf{p}'_i$. The vectors $(\hat{\mathbf{e}}_1, \hat{\mathbf{e}}_2, \hat{\mathbf{e}}_3)$ correspond to unit vectors defining an orthonormal basis set, with $\hat{\mathbf{e}}_1$ parallel to particle i 's initial momentum vector.

Collision pairing

To simplify particle pairing, the main particle list on each processor is split into an array of secondary lists, one per species per grid cell. For intra-species collisions (when $\alpha = \beta$) each odd numbered particle in the list is scattered off of the next particle in the list. For an odd number of particles, the last three particles in the list collide with each other (two collisions per particle), but using a collision frequency at half of its normally calculated value. Essentially the last three particles each undergo two half-collisions. Thus all particles within the cell have undergone one full collision. When performing inter-species collisions, all particles from the more numerous species undergo one and only one collision. While running through the more numerous list, the code loops over the second species list to form the successive collision pairs. The result is that

some (or all) of the particles from the second species may undergo multiple collisions. This is corrected for by the application of a common collisional time interval, which also corrects for the effect of non-uniform particle weights. The collision frequency is multiplied through by a correction factor F [118]:

$$F = \begin{cases} \frac{1}{f_{\alpha\beta}} \sum_i^{N_\alpha} w_i & \text{for } N_\alpha > N_\beta \\ \frac{1}{f_{\alpha\beta}} \sum_j^{N_\beta} w_j & \text{for } N_\alpha < N_\beta \end{cases},$$

$$f_{\alpha\beta} = (1 + \delta_{\alpha\beta}) \sum_k \min(w_i, w_j),$$

where the sum over k is over all collision pairs for species α and β . Thus while equations 2.14 and 2.15 each define a single frequency for each pair of particles in a binary collision event these are modified, and hence potentially different for the two particles, due to variable weights and different particle numbers per cell.

2.8.3 Limitations of the collision algorithm

The collision algorithm employed by EPOCH can be considered to be valid in the case of a ‘well-behaved’ plasma. However, if the upper limit on the impact parameter, b_{\max} , is approximately equal to the lower limit, b_{\min} , the expression for the Coulomb logarithm in equation 2.16 or 2.17 is no longer valid. This is generally the case for the low temperature and high density conditions which mark the strongly coupled regime. Consequently, most collisional PIC simulations tend to assume that the target is rapidly heated to a temperature beyond that of the strongly coupled regime (i.e. of the order 100 eV for solid density materials).

Furthermore, since targets heated by a short-pulse laser interaction are often only heated to a depth of a few tens of microns [8, 9], parts of the target which lie beyond this depth remain strongly coupled, and thus cannot be modelled accurately using EPOCH.

The collision algorithm also assumes a homogeneous plasma. EPOCH cannot currently model targets which possess some form of lattice structure within the target, as this would require incorporating a directional dependence in the collision frequency calculation. Often the lattice structure of the target can be assumed to have broken down by the time the target has been heated sufficiently that it has transited the strongly coupled regime discussed above.

2.8.4 Test problems

Spitzer conductivity

The current drawn by a constant, uniform electric field applied to an infinite plasma will tend to increase until the electric field force on the particles is equal to their rate of loss of momentum parallel to the field due to collisions. The steady state towards which the system tends can be expressed by a simplified Ohm's law: $J = \sigma E$, where $\sigma = q^2 n_e \tau / \alpha_0 m_e$ is the plasma conductivity [60], and $1/\alpha_0$ is the appropriate electrical conductivity transport coefficient [128, 129] (see section 3.2.1 for more detail), with $\tau = \nu^{-1}$ being given by equation 2.15. The variation of the current density with time can be obtained via a time-dependent solution to the Drude electron transport model (see section 1.5):

$$\mathbf{J}(t) = \sigma \mathbf{E} \left(1 - e^{-t/\tau}\right) + \mathbf{J}_0 e^{-t/\tau} \quad (2.18)$$

Such a system provides a potential test case for the ability of a collision algorithm to reproduce the 'Spitzer conductivity', which is itself a prerequisite for modelling the electron transport and Ohmic return current heating which dominate the transport processes in high density plasma. The system used for this test problem consisted of a periodic system with a constant external electric field.

In order to ensure that the applied field was constant, and that stopping was due to collisions only, the field updates within EPOCH were disabled for the purposes of this test. Ion position and momentum updates were also disabled so that the ions served as a constant stopping medium. A fixed Coulomb log of $\ln \Lambda = 10$ was used throughout.

Results for the case of $n_e = 10^{29} \text{ m}^{-3}$ and 100,000 particles per cell are shown in figure 2.8, and are in reasonable agreement with equation 2.18. The right-hand plot of figure 2.8 demonstrates that the current densities at time $t = 100\omega_p^{-1}$ are in good agreement with the predicted values from equation 2.18, with a root-mean-square fractional difference of 3.1%. Taking the gradient of a linear fit to the simulation data returns a numerical conductivity of $\sigma_n = 1.83 \times 10^6 \Omega^{-1} \text{ m}^{-1}$. This is comparable to the value returned by the analytic expression for the plasma conductivity: $\sigma_a = 1.71 \times 10^6 \Omega^{-1} \text{ m}^{-1}$ ($(\sigma_n - \sigma_a) / \sigma_a = 6.9\%$). The error in the conductivity arises due to the plasma experiencing greater Ohmic heating by the induced current in simulations with higher applied fields (since $\partial T / \partial t \propto \mathbf{E} \cdot \mathbf{J}$ —see section 1.5.2), and thus the late-time current densities not being a strict linear function of the applied field. This is evidenced by the fact that the current densities for applied fields of $E \geq 4 \times 10^9 \text{ V m}^{-1}$ are continuing to increase at time $t = 100\omega_p^{-1}$ in the left-hand plot of figure 2.8.




Figure 2.8: Left: Numerical (solid lines) and analytic (dashed lines) current densities as a function of time resulting from applied electric fields of 5×10^8 (blue), 10^9 (green), 2×10^9 (orange), 4×10^9 (red) and 6×10^9 V/m (black). Right: Late-time current densities as a function of applied electric field (crosses). The error bars indicate the range of values over the time period $90 < \omega t \leq 100$. Squares denote the expected values from equation 2.18. The solid line represents a linear fit to the simulation data.

Thermal conduction

In the previous section the use of periodic boundary conditions and averaging over the entire system tends to increase the particle statistics, making it easier to recover the correct behaviour. A more rigorous test can be constructed based on work by Bell et al. [130] in which a finite plasma of uniform density is modelled with a temperature gradient across part of the simulation box. Since the heat flux, Q , is a function of the temperature gradient, by defining the temperature profile using a \tanh function (ranging from 100 to 400 eV) the predicted temperature gradient, and thus heat flux, varies with position. In order for a PIC code to recover the correct collisional heat flux limiting [130], its collision algorithm must be accurate on the individual grid cell level (rather than relying on averaging over a collection of cells). Consequently, this is a much more demanding test of a collision algorithm's accuracy.

For this test the plasma was modelled as having a density of $n_e = 10^{28} \text{ m}^{-3}$, with stationary ions ($Z = 4$, $A = 8$). The simulation mesh was chosen to have a resolution equal to the Debye length ($\Delta x = \lambda_D$) and populated with 50000 pseudo-particles per cell (50% electrons). A fixed Coulomb log of $\ln \Lambda = 10$ was used throughout.

Figure 2.9 shows the resulting heat flux as a function of the temperature gradient




Figure 2.9: Left: Electron heat flux (normalised to the free-streaming flux, Q_f) as a function of the temperature gradient scale-length (normalised to the electron mean free path, λ) at a time $t = 50\omega_p^{-1}$. The straight line indicates the Spitzer-Harm values. Right: Electron heat flux (solid black) and temperature (dashed red) as functions of position (normalised to electron Debye length, λ_D). Both images are taken from an EPOCH simulation with 50000 particles per cell.

scale-length, L , after $t = 50\omega_p^{-1}$, where ω_p is the usual electron plasma frequency. The solid line indicates the Spitzer-Harm values.

Note that the heat flux for some temperature gradients appears to be double-valued (particularly in the range $10 \lesssim L/\lambda \lesssim 100$, where λ is the electron mean free path). This is due to the temperature profile being such that each value of the gradient occurs twice in the simulation domain. The lower band corresponds to values near the higher temperature region, and exhibit flux limiting due to collisions [130]. The upper band of values are from the colder region and are artificially inflated by high energy particles which have moved through from the hot region. This leads to the calculated heat flux values exceeding the Spitzer-Harm values.

Additional simulations were also performed with fewer particles per grid cell. The results of a simulation with 500 particles per cell are shown in figure 2.10. Comparing these results with those in figure 2.9 indicate that high particle statistics are required for PIC codes to accurately model electron transport in high density plasma.

At high particle resolution, the thermal conduction results are qualitatively similar to those obtained using a Vlasov code with a BGK collision operator [91, 131] or using a full Vlasov-Fokker-Planck (VFP) code [130]. However, the large numbers of particles required to accurately model collision-dominated phenomena mean that such problems can be very computationally intensive for a standard, collisional PIC code. For this




Figure 2.10: Left: Electron heat flux (normalised to the free-streaming flux) as a function of the temperature gradient scale-length (normalised to the electron mean free path) at a time $t = 50\omega_p^{-1}$. The straight line indicates the Spitzer-Harm values. Right: Electron heat flux (solid black) and temperature (dashed red) as functions of position (normalised to electron Debye length). Both images are taken from an EPOCH simulation with 500 particles per cell.

reason alternative approaches to modelling transport phenomena at high density, akin to hybrid models, are required in order to self-consistently model the absorption and subsequent transport of short-pulse laser energy in PIC codes [32].

Chapter 3

An Ohmic field solver for high density PIC simulations

3.1 Introduction

For typical laser interaction simulations with peak densities at, or approaching, solid density, the constraints on using standard PIC techniques in solving Maxwell's equations and the particle equations of motion (e.g. the CFL limit [97]) can severely limit the spatial and temporal scales which can be easily modelled. Continued increases in HPC (high performance computing) capabilities have helped to reduce the limitations on collisionless PIC simulations, but detailed modelling of collision-dominated regimes remains too computationally expensive to be performed on a regular basis.

To overcome this the direct laser interaction is often modelled using PIC methods at densities ranging from vacuum up to some multiple of the relativistic critical density (γn_c), the distribution of energetic electrons accelerated by the laser interaction (hereafter referred to as fast electrons) is sampled at a plane beyond the laser interaction region, and the transport of these electrons into the remainder of the target is modelled using a separate code with a hybrid kinetic-fluid model [30, 31]. However, this approach introduces its own potential sources of error. Often there is a 'transport gap' of several microns between the laser interaction region and the fast electron probe plane in which collisional effects may begin to have an important effect, and beam-plasma instabilities (e.g. Weibel instability [132]) may be seeded.

Since each code is designed to model a different range of densities and temperatures, the minimal overlap in densities for which PIC and hybrid codes are valid may result in phenomena such as magnetic field formation, beam filamentation and seeding of instabilities, which might occur at the intermediate densities, not currently

being captured accurately. Furthermore, the lack of feedback from the hybrid code to the PIC code prevents transport phenomena, via the return current, affecting the laser absorption.

Changes made by Cohen et al. [32] to the PSC [133], and also incorporated into OSIRIS [134, 135], allow for a more integrated approach by using a composite code which employs a conventional, fully electromagnetic PIC model in the laser interaction region (electron densities $\lesssim 100n_c$) and a hybrid-like model in the high density region. A similar algorithm has been developed for the PIC code EPOCH. It is expected that the algorithm will allow the ‘transport gap’ to be bridged by increasing the density regime overlap between PIC and Monte-Carlo transport codes.

Furthermore it is hoped that this algorithm will relax the need for high particle numbers in modelling collisional phenomena as indicated in section 2.8.4, since the model implicitly assumes a high degree of collisionality between the background electron and ion populations.

To easily distinguish between EPOCH simulations using only the standard PIC algorithms and those using the new field solver at high densities, the code used for the latter will herein be referred to as EPOCH-H.

3.2 The high density algorithm

3.2.1 Electric field calculation

At high electron density ($n_e \gtrsim 100n_c$) the plasma is sufficiently collisional that high plasma frequency waves are critically damped, and the relatively low energy (compared with the fast electrons) background electrons are well approximated by a drifting Maxwellian distribution which acts to cancel the fast electron currents, and so maintain quasi-neutrality (see section 1.5).

By taking the first moment of the Vlasov-Fokker-Planck equation, and neglecting electron inertia, it is possible to arrive at an expression for Ohm’s law (see section 1.2.2). This provides a means of self-consistently calculating the electric field based on various plasma properties (densities, temperatures, currents, etc.):

$$\mathbf{E} = \underline{\eta} \cdot (\mathbf{J}_b + \mathbf{J}_i) - \mathbf{u}_b \times \mathbf{B} - \frac{\nabla \cdot \mathbf{P}_b}{en_b} + \frac{1}{q_e} \left(\frac{\partial \langle \mathbf{p}_b \rangle}{\partial t} \right)_{\text{bf}}, \quad (3.1)$$

where:

$\underline{\eta}$ is the plasma resistivity tensor,

\mathbf{J}_b is the background electron current density,

\mathbf{J}_i is the ion current density,
 $\underline{\mathbf{P}}_b$ is the background electron stress tensor,
 n_b is the background electron number density, and
 \mathbf{u}_b is the background electron fluid velocity.

The resistive terms

The first term of equation 3.1, $\underline{\boldsymbol{\eta}} \cdot \mathbf{J}_b$, tends to be the dominant term in most high density plasma physics applications, with the remaining terms in Ohm's law providing corrections to the resistive electric field. The sum of the background electron and ion currents is evaluated via Ampère's law as detailed below.

Although strictly a tensor quantity, EPOCH-H currently uses a scalar plasma resistivity [60] for simplicity:

$$\eta = \frac{\alpha_0 m_e \nu_{bi}}{n_b q_e^2}, \quad (3.2)$$

where

$$\nu_{bi} = \frac{q_e^2 q_i^2 n_i \ln \Lambda}{3\epsilon_0^2 \sqrt{\mu} (2\pi k_B T_b)^{3/2}}, \quad (3.3)$$

is the background electron-ion collision frequency and $\mu = (m_e m_i) / (m_e + m_i)$.

With $\alpha_0 = 1$, the expression for η is that which results from assuming that the background electron energy distribution is a shifted Maxwellian. However, when the electrons are accelerated by an electric field (such as that produced by a local fast electron current), the resultant electron distribution diverges from a true shifted Maxwellian. Due to the v^{-3} dependence on the collision frequency, lower energy electrons experience a greater 'friction' force, and thus the return current is more dependent upon the electrons in the tail of the distribution. To account for this a correction factor, the electrical resistivity transport coefficient, α_0 , is often required. Values of α_0 have been calculated by Braginskii [128] as a function of the ions' charge state for the case where the electron cyclotron frequency, Ω_e , is much less than the electron-ion collision frequency. Further corrections to α_0 have been evaluated by Epperlein and Haines [129]. The latter can be approximated by the following function for the case of $\Omega_e \ll \nu_{bi}$:

$$\alpha_0 \approx 0.2945 + 0.2116 Z^{*-0.7}, \quad (3.4)$$

where the -0.7 exponent is the result of taking the average of the values of the expression $\log_Z^*(\alpha_{EH}(Z^*) - \alpha_{EH}(\infty))$, weighted by $1/Z^*$, with α_{EH} being the resistivity coefficients calculated by Epperlein and Haines [129]. The values of α_0 given by equation 3.4 are tabulated alongside the Braginskii [128] and Epperlein and Haines

Table 3.1: Tabulated values of the Braginskii electrical resistivity coefficient, α_0 , from Braginskii [128], the values based on corrections by Epperlein and Haines [129], and the corresponding approximate values using equation 3.4.

| α_0 | $Z^* = 1$ | $Z^* = 2$ | $Z^* = 3$ | $Z^* = 4$ | $Z^* \rightarrow \infty$ |
|----------------------------|-----------|-----------|-----------|-----------|--------------------------|
| Braginskii [128] | 0.5129 | 0.4408 | 0.3965 | 0.3752 | 0.2949 |
| Epperlein and Haines [129] | 0.5061 | 0.4295 | 0.3950 | 0.3750 | 0.2945 |
| Equation 3.4 | 0.5061 | 0.4248 | 0.3926 | 0.3747 | 0.2945 |

[129] values in table 3.1.

Recent work by Sherlock et al. [136] has demonstrated that fast electron bunches accelerated into dense targets by a short-pulse laser interaction induce plasma waves in the background electron population. The collisional damping of these waves provides an additional heating source which can exceed the heating due to the resistive return current. EPOCH-H, like many traditional hybrid codes which employ Ohm's law, assume that the inertia term in the background electron equation of motion can be considered to be zero:

$$\left(\frac{\partial}{\partial t} + \mathbf{u}_b \cdot \nabla \right) \mathbf{u}_b = \frac{d\mathbf{u}_b}{dt} = 0.$$

As a result, the form of Ohm's law used by these codes does not permit the production of electron plasma waves in the background material, and thus the additional source of heating associated with these waves is absent from many of the simulations using these codes. Sherlock et al. have developed a model which is capable of reproducing the plasma wave heating, and the feasibility of implementing their improved model in EPOCH-H will likely be investigated once further details of the algorithm are available.

Calculating the return current

Ampère's law relates the net current to the curl of the magnetic field. If the current contribution due to the fast electrons, \mathbf{J}_f , is calculated from the local particles (as described in section 2.2.2), Ampère's law provides a computationally inexpensive method for evaluating the sum of the background electron and ion currents without also accumulating these particles onto the grid (a time-consuming process). At high densities quasi-neutrality can be assumed, and thus the displacement current can be neglected. Consequently:

$$\mathbf{J}_b + \mathbf{J}_i = \frac{\nabla \times \mathbf{B}}{\mu_0} - \mathbf{J}_f. \quad (3.5)$$

Assuming quasi-neutrality implies that the force on the background electrons

from the electric field is balanced by a collisional drag force. In EPOCH-H this is provided by the binary collision algorithm. In order to ensure that the background electron current density is limited to a value which is consistent with the electric field associated with the resistive term in equation 3.1, the action of the collision algorithm needs to imply a resistivity which is equal to that calculated by the field solver (i.e. equation 3.2).

The Hall term

Although the Hall term, $\mathbf{u}_b \times \mathbf{B}$, is included in EPOCH-H's calculation of Ohm's law, this term only appears to offer a minor correction to the resistive term. However, its effect may become important during laser interactions in which the inward travelling fast electron current is interacting with its self-generated magnetic field over long time-scales.

The pressure term

As with the resistivity, EPOCH-H uses an isotropic, scalar pressure, rather than evaluating the full stress tensor. Since the pressure term, $\nabla \cdot \mathbf{P}_b / en_b$, tends to contribute little to the electric field, this approximation was deemed valid for most situations in which EPOCH-H is likely to be employed. At the time of writing, the pressure is calculated using an ideal gas equation of state ($P_b = n_b k_B T_b$), but EPOCH-H can be easily amended to use a user-defined equation of state, or look up values from an external tabulated library.

During testing it was noted that the pressure term in Ohm's law could potentially lead to the generation of magnetic fields from PIC noise. The source of these fields is the Biermann battery effect [137]:

$$\begin{aligned} \frac{\partial \mathbf{B}_P}{\partial t} &= -\nabla \times \mathbf{E}_P = \nabla \times \left(\frac{\nabla P_b}{q_e n_b} \right), \\ \Rightarrow \frac{\partial \mathbf{B}_P}{\partial t} &= \frac{\nabla P_b \times \nabla n_b}{q_e n_b^2}, \\ \Rightarrow \frac{\partial \mathbf{B}_P}{\partial t} &= \frac{k_B}{q_e n_b} \nabla T_b \times \nabla n_b. \end{aligned} \tag{3.6}$$

The accumulation of particle properties onto a discrete grid to derive variables such as density and temperature tends to introduce noise. From equation 3.6 it can be seen that this noise will result in the generation of magnetic fields, which will in turn introduce gradients in the electron density and temperature via the Hall term in Ohm's law, thereby enhancing the magnetic fields. This can be mitigated by employing higher order particle shape functions. However, the pressure gradient term in Ohm's law is a function of both the background electron density, and the gradient of the product of density and temperature. As a result it is much more sensitive to particle noise than the

other terms.

For most applications, the pressure term (neglecting any noise) will be small compared to the resistive term in Ohm's law, and thus can be neglected.

The collision term

The final term, $e^{-1}(\partial\langle\mathbf{p}_b\rangle/\partial t)_{\text{bf}}$, is associated with the momentum exchange between background and fast electrons via collisions. If required, it is possible to adapt the existing collision algorithm to calculate this contribution to the electric field. For the majority of anticipated applications, however, the fast electron collision frequency is much less than the background electron-ion collision frequency, and thus this term is negligible compared to the resistive terms. Consequently the collision term in equation 3.1 is not currently included in EPOCH-H.

Updating the electric field

In order to ensure that the magnetic field update remains accurate to second order, it is necessary to evaluate the electric field at the mid-point of each time-step. However, Ohm's law can only provide the full time-step values for the electric fields. To obtain the half time-step values a first order linear extrapolation method is used. Specifically, to advance the solution from step n to $n + 1$, it is assumed that the change in the electric fields between $n - 1/2$ and n is equal to the change between n and $n + 1/2$:

$$\mathbf{E}^{n+1/2} = 2\mathbf{E}^n - \mathbf{E}^{n-1/2},$$

where

$$\mathbf{E}^n = \eta^n (\mathbf{J}_b^n + \mathbf{J}_i^n) - \mathbf{u}_b^n \times \mathbf{B}^n - \frac{\nabla P_b^n}{q_e n_b^n}, \quad (3.7)$$

is the electric field calculated at time n via Ohm's law (equation 3.1).

The interface region

In order to reduce the effect of any potential issues in having a step-change from the PIC to the high density model, a gradual, density-based interface was added. The definitions of 'low' and 'high' density were amended to cases where $n_b < n_l$ and $n_b > n_u$ respectively, where n_l and n_u are set by the user in the input deck. At the intermediate densities ($n_l < n_b < n_u$) EPOCH-H interpolates between the two possible solutions based on the background electron density. This ensures that the electric field solution

transitions smoothly between the two field solvers. Thus:

$$\mathbf{E}^n = \mathbf{E}^{n-1/2} + (1 - \delta) \Delta_M^n + \delta \Delta_O^n, \quad (3.8)$$

$$\mathbf{E}^{n+1/2} = \mathbf{E}^n + (1 - \delta) \Delta_M^n + \delta \Delta_O^n, \quad (3.9)$$

where

$$\delta = \begin{cases} 0 & \text{if } n_b < n_l \\ 0.5 \left(1 - \cos \left(\pi \frac{n_b - n_l}{n_u - n_l} \right) \right) & \text{if } n_l \leq n_b \leq n_u \\ 1 & \text{if } n_b > n_u \end{cases},$$

$$\Delta_M^n = \frac{\Delta t}{2} \left(c^2 \nabla \times \mathbf{B}^n - \frac{\mathbf{J}^n}{\epsilon_0} \right),$$

$$\Delta_O^n = \left(\eta^n (\mathbf{J}_b^n + \mathbf{J}_i^n) - \mathbf{u}_b^n \times \mathbf{B}^n - \frac{\nabla P_b^n}{q_e n_b^n} \right) - \mathbf{E}^{n-1/2}.$$

3.2.2 Identifying particle species

In order to differentiate between fast and background electrons two particle species are required for the electrons. For most simulations, the fast electron particle list is initially empty, with all of the electrons considered to be background electrons. Each time-step the kinetic energy of each electron (fast as well as background) is compared with the local temperature of the background population. Any background electron whose kinetic energy, K , is greater than some multiple, α , of the local thermal energy is promoted to the fast electron population. As the simulation progresses, and electrons are accelerated beyond this threshold, the simulation code transfers them from the background electron species list to the fast electron species list. A similar criterion, governed by the free parameter β , is also used to determine whether a fast electron has lost sufficient energy (via field stopping or collisions, for example), to warrant returning it to the background population. Thus for fast electrons:

$$K_f > \frac{3}{2} \alpha k_B T_b,$$

and similarly for background electrons:

$$K_b < \frac{3}{2} \beta k_B T_b.$$

Note that α and β need not necessarily be equal. For most cases it may be desirable to run with $\alpha > \beta$ so that some overlap in phase-space between the two electron populations is possible. This also ensures that fast electrons are well within

the background electron energy distribution before ‘demoting’ them, thus reducing the possibility of the demotion of large numbers of fast electrons resulting in a significant departure from the assumed Maxwellian background electron energy distribution.

Criteria for promotion/demotion of electrons based on the background electron density are also employed in EPOCH-H so that it is possible to run simulations such that electrons are only promoted to the fast population in the low density pre-plasma, and only returned to the background population in the high density regions. The former condition (only promote if $n_e < \xi n_c$, where ξ is a free parameter) can be used to ensure that electrons are not promoted to fast electrons in regions that the laser is unable to reach (i.e. $n_e \gg n_c$).

Although the ‘particle migration’ described above is necessary for EPOCH-H’s Ohmic field solver, it is also a useful means of differentiating between background and fast electrons in standard PIC simulations [31]. Consequently, the particle migration algorithm has been written so as to be independent of the field solver.

Whether a particle species should be treated by the Ohmic field solver as a background electron, fast electron or ion species is determined by the `species_type` variable. This can be assigned by the user using making use of the species identity functionality added by the QED module [138]. More specifically, the user can set the `identify` deck option in the species blocks of the input deck to either `electron`, `hot_electron` or `ion`. If the `identify` option is not used in the deck, EPOCH-H attempts to infer the intended role for the species based on the particles’ charge and migration options:

- If $q > 0$, assume the particles are ions.
- If $q < 0$ and the species can be demoted, assume the particles are fast electrons.
- If $q < 0$ and the species cannot be demoted, assume the particles are background electrons.

Particle demotion at thermal boundaries

Any particles which exit through a thermalising boundary in EPOCH have their momenta re-sampled from the initial temperature conditions, and are re-injected. For fast electrons this can result in large numbers of particles collecting near the boundaries. This is inconsistent with the concept of these particles as being energetic and nearly collisionless. Furthermore, since the fields in EPOCH-H are indirectly calculated from the fast electron current, such behaviour can potentially result in noisy electric fields near the boundaries which will rapidly heat the background electron population.

To prevent these phenomena the thermalising boundary conditions were modified so that any particles attempting to leave the simulation domain were automatically

demoted to the lowest level species in the particle migration chain. This ensured that any fast electrons leaving the domain were replaced with background electrons. This modification is consistent with the behaviour that the thermalising boundary conditions aim to emulate; namely the supply of particles to form a return current from beyond the region being simulated. Furthermore, the potential creation of void regions at the boundaries, due to the electric fields drawing background electrons away from the boundaries, is mitigated by the modified boundary conditions.

3.2.3 Multiple electron/ion species

In some simulations, such as the proton acceleration and buried layer heating simulations discussed in chapters 5 and 6 respectively, multiple ion species may be present. Similarly, the user may choose to employ multiple background and/or fast electron species. To account for this, the number and current densities used to calculate the resistive and pressure terms in equation 3.7 are the sum over all species of the relevant particle type (background electron, fast electron or ion). The temperature used in these terms is a density-weighted average over all species of the relevant type. The ion charge and mass used to calculate the electron-ion collision frequency when evaluating the resistive term (equation 3.3) is a density-weighted average, evaluated on a per-cell basis.

3.2.4 Magnetic field update

The magnetic field update in EPOCH-H is unchanged from the standard PIC method using Faraday's law, with the magnetic field at step $n + 1$ being advanced from step n using the time-centred electric field at step $n + 1/2$, as discussed in section 2.3:

$$\begin{aligned} \mathbf{B}^{n+1/2} &= \mathbf{B}^n - \frac{\Delta t}{2} \nabla \times \mathbf{E}^{n+1/2}, \\ \mathbf{B}^{n+1} &= \mathbf{B}^{n+1/2} - \frac{\Delta t}{2} \nabla \times \mathbf{E}^{n+1/2}. \end{aligned}$$

This is the standard finite-difference-time-domain (FDTD) method for solving Maxwell's equations [100, 101]. This method is second order accurate in space and time, and is mathematically equivalent to the commonly used leapfrog method [95].

3.2.5 Particle position and velocity update

Unlike many other 'hybrid' methods, there is no continuum fluid in EPOCH-H. Therefore all species retain the standard PIC treatment (see section 2.2) in updating their positions

and momenta via the Lorentz force:

$$\begin{aligned} \mathbf{x}^{n+1/2} &= \mathbf{x}^n + \frac{\Delta t}{2m\gamma^n} \mathbf{p}^n, \\ \mathbf{p}^{n+1} &= \mathbf{p}^n + q\Delta t \left(\mathbf{E}^{n+1/2} + \frac{1}{2m\gamma^{n+1/2}} (\mathbf{p}^{n+1} + \mathbf{p}^n) \times \mathbf{B}^{n+1/2} \right), \\ \mathbf{x}^{n+1} &= \mathbf{x}^{n+1/2} + \frac{\Delta t}{2m\gamma^{n+1}} \mathbf{p}^{n+1}. \end{aligned}$$

The momentum update above is performed by decoupling the electric and magnetic field effects using the 'Boris' algorithm [95]. The update is split into two half-impulses from the electric field applied before and after the magnetic field effect. Since the $\mathbf{p} \times \mathbf{B}$ force is conservative, the magnetic field effect can be calculated by applying a rotation to the particle's momentum vector (see equation 2.6). It should be noted that the conservative nature of the Boris rotation means that the value of γ above is constant during this step of the particle update.

3.3 Implications for PIC modelling

In standard PIC simulations, failure to resolve the Debye length results in the plasma self-heating until it reaches a temperature at which the Debye length is comparable to the grid resolution. By neglecting electron inertia the algorithm employed in EPOCH-H assumes quasi-neutrality, and as such will not self-heat if the Debye length is not resolved. Furthermore the algorithm does not permit the propagation of electromagnetic or electron plasma waves. As a result there is no need to resolve the electron plasma frequency in the high density region. The method developed by Cohen et al. [32], therefore, relaxes the time-step and grid resolution conditions such that they are set by the peak electron density in the standard PIC region ($\sim 100n_c$) rather than the maximum density in the simulation domain (likely to be $\gg 100n_c$).

The high particle numbers required in standard PIC codes to accurately model collisional effects (as shown in section 2.8.4) are not required in EPOCH-H. Many of the collisional effects which would normally need to be explicitly modelled using a standard collisional PIC code (such as the background electron population forming a drifting Maxwellian population) are assumed by the model.

Chapter 4

Preliminary testing and comparisons

4.1 Introduction

In order to test the accuracy of the Ohmic field solver in EPOCH-H, a series of test simulations were performed using the one and two dimensional versions of the code. The tests were performed by comparing the results of EPOCH-H simulations with those obtained using collisionless EPOCH, collisional EPOCH (hereafter referred to as EPOCH-C to avoid confusion), analytic models and/or the fast electron transport code, THOR (see section 4.1.1 below).

Unless otherwise stated, first order ('triangular') particle shapes were used in the EPOCH-H simulations. Furthermore, in all cases a fixed Coulomb logarithm of $\ln \Lambda = 5$ was used. In order to ensure consistency between EPOCH-H and THOR, the latter was run using an ideal gas equation of state and Spitzer resistivity (also with a fixed Coulomb log of 5). Since the fast electron particle statistics are generally not sufficiently high in EPOCH-H to accurately model fast electron collisions using binary scattering, they were assumed to be collisionless.

In order to allow for easy comparison with THOR, the 2D tests were performed with a particle probe recording the fast electrons passing a plane in the simulation. The recorded population was post-processed onto a four-dimensional grid (position, energy, direction of motion and time) to create a fast electron source for THOR. This ensured that the fast electron distribution being injected in the THOR simulations was as similar as possible to the distribution travelling across the low/high density interface in EPOCH-H.

4.1.1 THOR summary

THOR [31] models the effects of fast electron transport using a hybrid approximation based on that of Davies et al. [87, 139]. Monte-Carlo methods are used to sample energetic ($> 10 \text{ keV}$) electrons from a defined distribution function. These particles are injected through a plane in the simulation domain, with their positions and momenta being updated over a time-step (set by the CFL limit) in a similar manner to a PIC code (section 2.2).

Electric fields are calculated on an Eulerian mesh using a simplified Ohm's law, assuming that the background return current exactly cancels the fast electron current, and magnetic fields induced via Faraday's law (equation 2.8).

The effect of collisions on the fast electrons are modelled by splitting the collision operator into drag (Δp) and scatter ($\Delta \theta$) terms (see equations C.1 and C.2 in appendix C):

$$\begin{aligned}\langle \Delta p \rangle &= -\frac{Z q_e^4 n \gamma m \ln \Lambda_l}{4\pi \epsilon_0^2 p^2} \Delta t, \\ \langle \Delta \theta^2 \rangle &= \frac{Z^2 q_e^4 n \gamma m \ln \Lambda_s}{2\pi \epsilon_0^2 p^3} \Delta t,\end{aligned}$$

where the Coulomb logarithm-like terms $\ln \Lambda_l$ and $\ln \Lambda_s$ are as defined by Davies et al. [139]. Note that the expression for $\Delta \theta^2$ is used above, since $\langle \Delta \theta \rangle = 0$ (see appendix C.2).

Ohmic heating of the background plasma, which is modelled as a static, single temperature fluid, is calculated and added to the energy lost by fast electrons via the drag term, with the subsequent increase in the background material temperature being calculated in accordance with the estimated specific heat.

4.2 Monoenergetic beam

The most simple test of the high density algorithm is to simulate a uniform, periodic plasma with a fraction of the electron population travelling with some initial momentum in one direction. Since these electron will constitute a 'fast current', the new field solver should generate an electric field which draws a return current. Provided the collision algorithm is sufficiently accurate, the return current electrons will be inhibited by collisions and the current will be limited in accordance with Ohm's law, as discussed in section 3.2.1. Note that this test case is very similar to the Spitzer resistivity tests described in section 2.8.4.

The background electron current, as shown by figure 4.1, increased over the first

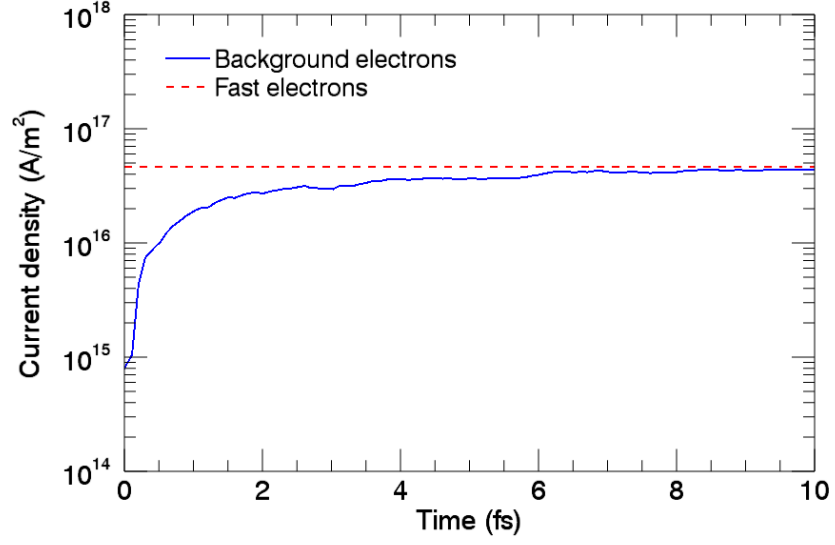


Figure 4.1: The absolute values of fast electron current (red) and background electron current (blue) as functions of time.

few femtoseconds of the simulation until the effect of collisions led to it being limited to a value which approximately balanced the fast electron current, as seen in the collision algorithm tests in section 2.8.4.

The x -components of the electric field and electron current densities late in the simulation ($t = 100$ fs) are shown in figure 4.2. The values of the current densities in figure 4.2 have been multiplied by the local values of resistivity, and demonstrate that the return current drawn is in accordance with the resistive term of Ohm's law ($\mathbf{E} = \eta \mathbf{J}_b$).

The expected Ohmic heating rate (equation 1.12), when the currents had reached equilibrium but before heating had begun to slow ($6 \text{ fs} \leq t \leq 8 \text{ fs}$), was calculated to be $2\langle J_x E_x \rangle / 3n_e q_e = 11.27 \text{ eV/fs}$. The measured heating over the same period was found to imply a heating rate of $\Delta T_e / \Delta t = 10.77 \text{ eV/fs} = 0.96 (2\langle J_x E_x \rangle / 3n_e q_e)$, which is in good agreement with the expected value.

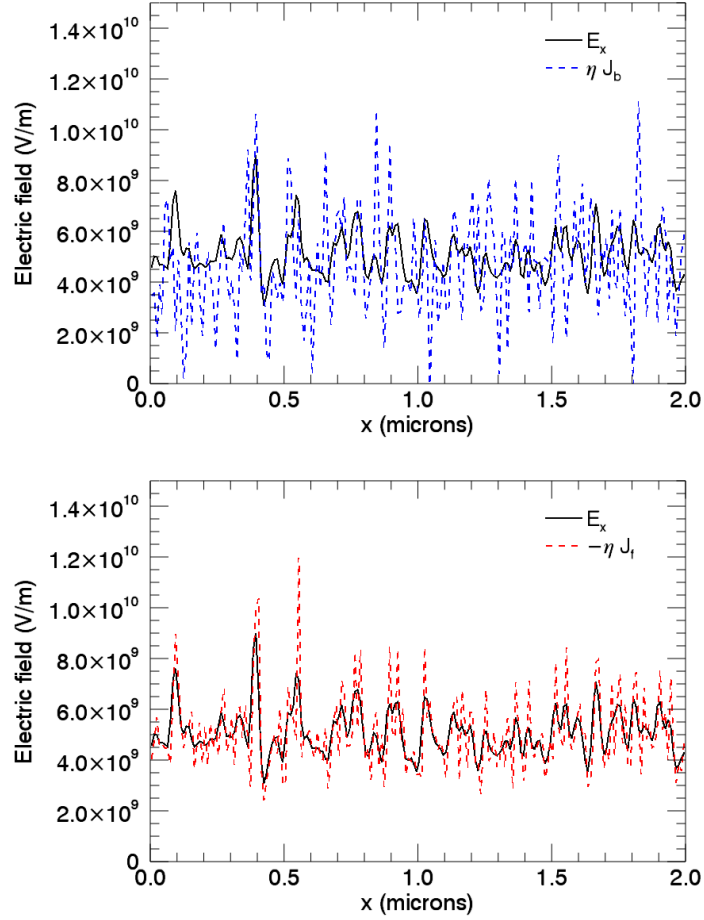


Figure 4.2: Results of the 1D monoenergetic beam test problem. Note that the return current drawn is in good agreement with Ohm's law ($\mathbf{E} = \eta \mathbf{J}$), and balances the fast electron current ($\mathbf{J}_f \approx -\mathbf{J}_b$). The x -component of the electric field (black), fast electron current (red, $\eta \mathbf{J}_f$) and background electron current (blue, $\eta \mathbf{J}_b$) as functions of position after 100 fs.

4.3 Diffusion

The monoenergetic beam test above provides a means of ensuring that the Ohmic field solver performs as expected in isolation. However, the main motivation in adding the high density algorithm in EPOCH-H was to reduce the computational overhead associated with modelling solid density laser-plasma interactions. Since Ohm's law does not permit electromagnetic waves to propagate, the standard Maxwell field solver must still be used in regions where the electron density is sufficiently low that the laser fields are not strongly damped (i.e. where $\omega_p \lesssim \omega$). Consequently, great care must be taken in




Figure 4.3: Lineouts of background electron temperature along $y = 0$ from 2D diffusion simulations using EPOCH-C (dot-dashed blue), EPOCH-H (solid red) and THOR using an EPOCH-H source (dashed green) at $t = 100$ fs.

handling the interface between the two field solver regions.

To test the interface between the Maxwell and Ohmic field solvers simulations were performed of a uniform deuterium plasma with initial temperature $T = 100$ eV and electron density $n_e = 10^{28} \text{ m}^{-3}$, split into two regions: $x < 0$ using the standard PIC Maxwell field solver, and $x \geq 0$ employing the high density model. A population of fast electrons, with an initial temperature of 100 keV, was added in the PIC region at 10% the background electron density, and allowed to diffuse into the $x > 0$ region.

Simulations were performed of a uniform deuterium plasma with fast electrons in the region $x < 0$, $-1.5 \mu\text{m} < y < 1.5 \mu\text{m}$ at 10% the background density, and the interface between the Maxwell and Ohmic field solvers in EPOCH-H was set as being the plane $x = 0$.

Additionally, a particle probe plane was placed at the interface ($x = 0$) to record the fast electrons which drifted into the Ohmic field solver region, so as to provide a source term for a THOR simulation of this region.

Although the fast electron behaviour was found to be very similar in all cases, there was a marked difference in the background electron properties between EPOCH-C and EPOCH-H, as shown in figure 4.3.

The EPOCH-C simulation exhibited significant heating of the background electron population due, in part, to numerical self-heating. In contrast, the EPOCH-H simulation displayed good agreement with THOR, thus providing further confidence in the employment of EPOCH-H for modelling energy transport in high density plasma.

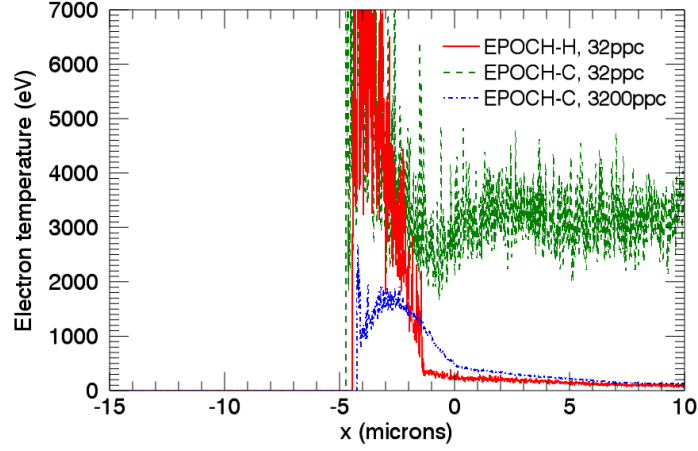


Figure 4.4: Background electron temperatures at $t = 150$ fs from 1D EPOCH-C (dashed green) and EPOCH-H (solid red) simulations. The self-heating of the background electrons in the EPOCH-C simulation demonstrates the advantage of EPOCH-H's high density algorithm ($x > 0$). A 1D EPOCH-C simulation with the particle numbers increased by a factor of 100 is also shown (dot-dashed, blue) to demonstrate the effect of particle numbers on numerical heating.

4.4 Laser-plasma interaction

4.4.1 1D

Since the main motivation for developing EPOCH-H was modelling high intensity laser interactions with high density plasma in 2D, an additional test problem to test the code's ability to model such a system was developed. A further simulation was performed of a laser pulse incident upon an exponential density ramp ($n_e \sim e^{x/L}, \forall x < 0$, where $L = 1.5 \mu\text{m}$) preceding a high density plasma ($n_e = 360n_c, \forall x \geq 0$). The interface region between the low and high density regions was defined as being at $x = 0$.

The advantage of using EPOCH-H for high density PIC simulation is evident when numerical self-heating is considered. Figure 4.4 shows the result of simulations using EPOCH-C and EPOCH-H. To ensure that any differences were due solely to the Ohmic field solver, the EPOCH-C simulation was performed with particle migration settings identical to the EPOCH-H simulation.

The EPOCH-C simulation with 32 particles per cell shown above exhibits a significant amount of self-heating, whereas very little heating is observed in the EPOCH-H results prior to the arrival of the fast electron population. While the EPOCH-H simulation required an additional 12% run time to reach completion, an EPOCH-C simulation with the particle numbers increased by a factor of 100 to mitigate numerical self-heating

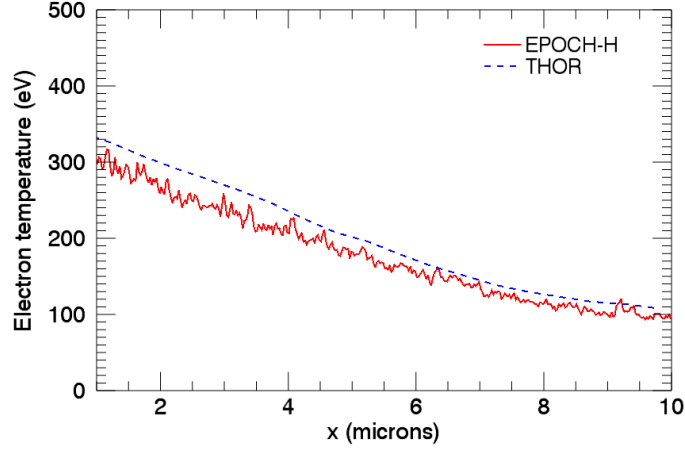


Figure 4.5: Lineout of background electron temperature from 2D laser-plasma interaction simulations at $t = 200$ fs taken along $y = 0$.

required more than 100 times the run time of the original EPOCH-C simulation. Thus it is clear that the increase in computational workload associated with solving Ohm's law is negligible compared to that required by standard methods of reducing self-heating, such as increasing particle numbers and/or grid resolution. Note, however, that this increase in run time does not account for any savings made possible by the relaxation of the time-step constraints in EPOCH-H simulations, since both the EPOCH-C and EPOCH-H simulations were run with the same time-step.

4.4.2 2D

For this test the system used in section 4.4.1 was extended $25 \mu\text{m}$ in the y direction, the incident laser was set up so as to focus onto a spot of diameter $5 \mu\text{m}$ at $x = 0$ (the top of the density ramp), and the peak electron density was reduced to $120n_c$ (to reduce run time). The transition between the two regions in EPOCH-H was set so as to take place between initial background electron densities of $90n_c$ and $100n_c$ (see equations 3.8 and 3.9 in section 3.2.1). For the purposes of this test the interface was not permitted to vary with time.

Simulations were performed using EPOCH-H, again with a probe plane at $x = 0$ to provide a source for THOR simulations via LOKI. In the high density plateau EPOCH-H again demonstrates minimal self-heating compared with EPOCH-C. The results shown in figure 4.5 demonstrate qualitative agreement between EPOCH-H's and THOR's heated regions.

As with the 1D test, the EPOCH-H simulation was found to have a longer run time

Table 4.1: Computational expense and electron heating rate associated with various versions of EPOCH for the case of a simple, uniform system.

| Particles per cell | Cell size | Comments | Run time (CPU hours) | Heating rate (eV/ps) |
|-----------------------|---------------|----------------------------|-------------------------|-------------------------|
| 64 | $10\lambda_D$ | Collisions disabled | 2.66 | 75.2 |
| 64 | $10\lambda_D$ | Zeroth order particles | 9.16 | 896.4 |
| 64 | $10\lambda_D$ | Standard EPOCH-C | 10.17 | 86.8 |
| 64 | $10\lambda_D$ | Third order particles | 12.40 | 7.5 |
| 64 | $10\lambda_D$ | Smoothed currents | 10.14 | 31.6 |
| 6400 | $10\lambda_D$ | High particle resolution | 1372.08 | 1.5 |
| 64 | $1\lambda_D$ | High mesh resolution | 163.53 | 1.9 |
| 64 | $10\lambda_D$ | Particle migration enabled | 11.44 | 86.8 |
| 64 | $10\lambda_D$ | Ohmic field solver | 12.44 | 0.7 |

(an increase of 9%) than an equivalent simulation performed using standard EPOCH-C.

4.5 Run times and self-heating rates

A uniform, periodic, 32×32 -cell region of fully-ionised deuterium ($Z = 1$, $A = 2$), with initial temperature 100 eV and density 10^{29} m^{-3} , was simulated using EPOCH, EPOCH-C and EPOCH-H to investigate the rate of self-heating without energy being added by an external source, as well as estimate the computational expense associated with the various code features required to accurately model transport phenomena (particle migration, collisions, Ohm's law). Additional simulations of solid density plastic using EPOCH-C were performed with higher order particle shape functions, current spatial smoothing, increased particle numbers and a higher mesh resolution to investigate the relative efficiency of each in reducing self-heating to a similar level to that observed using EPOCH-H. The run time required for each simulation is given in table 4.1.

Although the number of grid cells was kept constant in all cases, the reduced grid spacing used in the high mesh resolution simulation resulted in a reduction in the time-step duration. Hence the factor of ten increase in run time. It should also be noted that in order to simulate the same physical extent as the other simulations, the high mesh resolution case would have needed 100 times as many cells, resulting in a similar increase in the run time.

The gradients of the average electron temperatures (figure 4.6) indicate the heating rate for each case. The gradients obtained from linear fits to the curve in figure 4.6 are listed in table 4.1. The rates observed in the EPOCH and EPOCH-C simulations

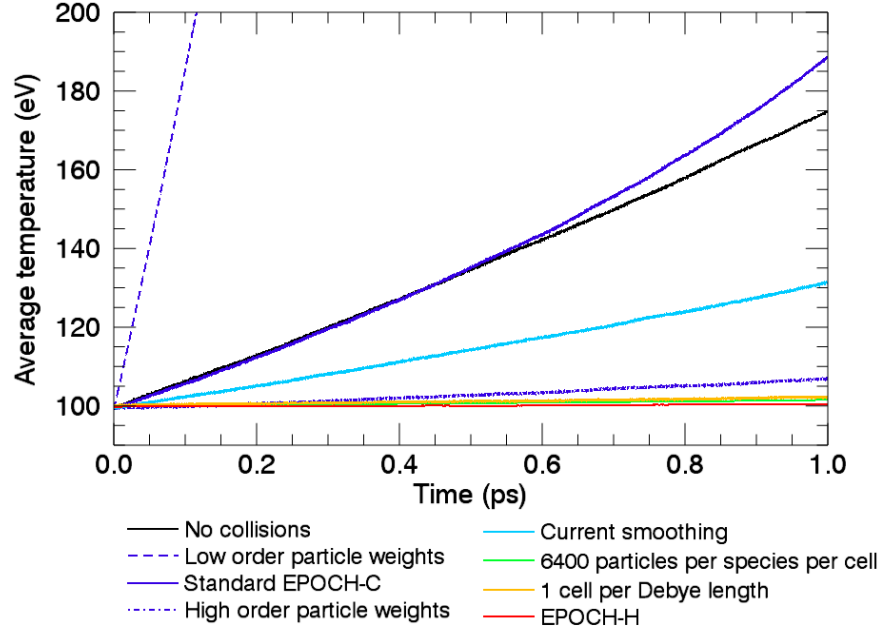


Figure 4.6: Average electron temperature as a function of time, demonstrating the rate of numerical heating for the EPOCH, EPOCH-C and EPOCH-H simulations listed in table 4.1. Note that the temperature for the low order particle shape case increased linearly to approximately 1 keV after 1 ps.

were found to all be within a factor of 4 of the expected values based on the order-of-magnitude estimates for numerical self-heating given by equation 2.13 [29]. As noted in the earlier tests, no significant self-heating was evident in the EPOCH-H simulation, in sharp contrast to the EPOCH and EPOCH-C results which exhibited much higher heating rates (approximately 90 eV/ps). The high particle and high mesh resolution simulations also displayed self-heating, but with much slower rates of 1.5 eV/ps and 1.9 eV/ps, respectively.

Overall, the results of the tests described above demonstrate that the fraction of simulation run time associated with the high density model in EPOCH-H (8%) is small compared to that required by the particle collision routines (60%). Furthermore, the increase in run time required by EPOCH-H is much smaller than that required to mitigate self-heating via increasing the particle and/or mesh resolution.

4.6 Summary

A series of one and two-dimensional tests problems were investigated to evaluate the ability of EPOCH-H to model conditions and phenomena relevant to solid density laser-plasma interactions.

Simulations of a monoenergetic beam of energetic electrons travelling through a uniform density plasma confirmed that the electric field calculated using Ohm's law will act to accelerate the background electrons to provide return current of sufficient magnitude to balance the imposed fast electron current. Furthermore, collisions between the background electrons and ions acted to limit the return current, thus confirming that the value of the resistivity calculated when evaluating Ohm's law is consistent with the value implied by the effects of the collision algorithm.

To test EPOCH-H's ability to simultaneously employ differing approaches to updating the electric fields in separate regions, two-dimensional simulations of fast electrons drifting through a uniform, dense background plasma were performed. A particle probe plane placed at the interface between the two regions in EPOCH-H recorded all of the fast electrons which passed into the Ohmic field solver region. These electrons were used to form a source function for a simulation of the Ohmic field region using a dedicated fast electron transport code (THOR). The lack of numerical self-heating in the Ohmic region, resulted in good agreement between EPOCH-H and THOR compared with the portion of the domain using the standard Maxwell field solver. This further increased confidence in EPOCH-H's ability to accurately model high density regions.

Since the main purpose of developing the high density model in EPOCH-H was to assist with the simulation of short-pulse laser interactions with solid density targets, one and two-dimensional simulations of laser interactions were also performed. These simulations consisted of an intense laser pulse incident upon a short density ramp preceding a region of high density plasma. In a similar manner to the diffusion tests, fast electrons were recorded by a probe plane placed at the top of the density ramp in the 2D simulation, and used to provide a source term for a comparative simulation of the dense region in THOR. As with the previous tests, negligible self-heating and good agreement with THOR was observed in the EPOCH-H high density regions.

Overall, the tests discussed in the preceding sections confirm that the Ohmic field solver employed by EPOCH-H allows the range of densities which can be modelled in a laser-plasma interaction simulation to be extended so as to provide a greater overlap with the densities typically modelled using dedicated transport codes such as THOR. The results presented in section 4.5 also clearly demonstrate that EPOCH-H exhibits significantly reduced levels of numerical self-heating in high density plasmas, with a

much lower computational cost than the standard approaches of mitigating self-heating, such as increasing the mesh and/or particle resolution.

In order to demonstrate the enhanced modelling capability for short-pulse laser experiments provided by EPOCH-H, two situations which have previously proven challenging for PIC codes were investigated: proton acceleration from the rear surface of a target, and return current heating of a solid density foil. Each of these cases is discussed individually in chapters 5 and 6, respectively.

Chapter 5

The effect of high density transport phenomena on the TNSA process

5.1 Introduction

The target normal sheath acceleration (TNSA) process [70] is an effective means of producing directional proton beams with typical energies of tens of MeV, using high intensity short-pulse lasers. Fast electrons, produced by the laser interaction, travel through the target and create a large electrostatic field at the rear surface [69, 71, 73]. This field then accelerates the carbon and hydrogen ions, which constitute the hydrocarbon contaminant layer present on most laser targets, up to multi-MeV energies. Protons can also be accelerated from the front surface of the target [68, 72], but upon encountering the fields at the rear surface, these protons are accelerated up to similar energies so that they are indistinguishable from protons originating from the rear surface [140]. Protons accelerated by the TNSA mechanism have been used for a variety of purposes such as radiographic imaging [16, 17], probing of electric and magnetic fields [15], and isochoric heating of a secondary target for equation of state measurements [77–79].

The results of TNSA simulations using EPOCH-H and standard EPOCH are presented and compared below to demonstrate the improved high density modelling capability provided by EPOCH-H. To more easily distinguish between the different versions of EPOCH discussed in this chapter, standard EPOCH running with the collision algorithm enabled will be referred to as EPOCH-C.

5.2 Simulation setup

A 100 J Gaussian pulse with a 0.5 ps full-width-half-maximum (FWHM) temporal profile was injected from the left boundary of the simulation domain with a wavelength of $1.056 \mu\text{m}$. The incident pulse was modelled as having a Gaussian transverse intensity profile, with 90% of the laser energy contained within a $30 \mu\text{m}$ diameter spot.

The target was modelled as solid density aluminium (2.7 g/cc), with an assumed constant, uniform ionisation state of $Z^* = 10$ to represent the expected average ionisation of the target based on calculations by the atomic physics code CASSANDRA [141]. In order to provide a smooth density profile, the following definitions were used for the ion, proton and electron number densities, respectively (see also figure 5.1):

$$n_i = 0.5n_0 \left[\tanh\left(\frac{x}{L_1}\right) - \tanh\left(\frac{x-w}{L_2}\right) \right], \quad (5.1)$$

$$n_p = 0.5n_0 \left[\tanh\left(\frac{x}{L_1}\right) - \tanh\left(\frac{x}{L_2}\right) \right] + 0.5n_0 \left[\tanh\left(\frac{x-w+L_2}{L_2}\right) - \tanh\left(\frac{x-w-L_2}{L_2}\right) \right], \quad (5.2)$$

$$n_e = Z^*n_i + n_p, \quad (5.3)$$

where:

n_i is the aluminium ion density,

n_p is the proton number density,

n_e is the electron number density,

$n_0 = 6 \times 10^{28} \text{ m}^{-3}$ is the ion number density of solid aluminium,

$L_1 = 4 \mu\text{m}$ is the front surface density scale-length,

$L_2 = 0.5 \mu\text{m}$ is the rear surface density scale-length,

$w = 30 \mu\text{m}$ is the target width, and

$Z^* = 10$ is the assumed aluminium ionisation state.

The simulation domain measured $125 \times 80 \mu\text{m}^2$, with a mesh resolution of $\Delta x = \Delta y = 24.4 \text{ nm}$ (~ 40 cells per micron), providing several tens of microns of vacuum region behind the target to allow the protons to be accelerated by the rear sheath field. The simulation was performed with 32 electrons, 8 ions and 16 protons per cell in which the density of the relevant species was non-zero, with non-uniform particle weighting being employed. The interface between the low density, Maxwell, and high density, Ohmic, field solvers was assigned to the density range $80n_c \leq n_e \leq 90n_c$, where $n_c = 10^{27} \text{ m}^{-3}$ is the critical electron density. The updates to the electric field for cells




Figure 5.1: Lineouts of the initial electron and ion densities for the EPOCH-H and solid density EPOCH TNSA simulations. The capped simulations (not shown) applied an upper limit to the electron density, and corresponding upper limit to the ion densities.

whose density lies within this range were interpolated between the two solutions in the manner described by equations 3.8 and 3.9 in chapter 3.

A particle probe placed at $x = 95 \mu\text{m}$ was used to record the energy, direction of motion and transverse (y) position of all protons accelerated from the target over the course of the simulation.

In order to investigate the effect of the high density model, and associated modelling of electron transport through a solid density target, on the acceleration of protons from the rear surface of the target, four further simulations were performed using standard EPOCH and EPOCH-C. The first was identical to the EPOCH-H simulation, but did not make use of the updated field solver. Note that in this case there was no change in the time-step duration, since the high mesh resolution resulted in the time-step being constrained by the CFL limit [97].

To assess the efficacy of using EPOCH-H's high density model to mitigate numerical self-heating compared to varying the choice of particle shape function, an additional solid density EPOCH-C simulation was performed using higher order particle shapes.

The remaining two simulations were performed with the electron density limited to a maximum value of $90n_c$ (the upper limit on the interface region in the EPOCH-H simulation), to allow for a comparison of more typical PIC simulations with EPOCH-H. These simulations were performed with and without collisions enabled, and are hereafter referred to as 'capped' simulations.




Figure 5.2: Time-integrated energy spectra of protons passing a plane $65\ \mu\text{m}$ behind the rear target surface, taken from simulations using EPOCH-H (thick, black), solid density EPOCH-C (dot-dash, blue), solid density EPOCH-C with high order particle shape functions (dot-dash, green), capped density EPOCH-C (solid, red) and collisionless, capped density EPOCH (dashed, orange). Most of the results shown have been integrated over 5 ps. However, the solid density EPOCH-C simulation, with first order particle shapes, (dot-dash, blue) was stopped after 3 ps due to the excessive self-heating.

5.3 EPOCH/EPOCH-C/EPOCH-H comparison

5.3.1 Proton energy spectra

Particle probe planes placed $65\ \mu\text{m}$ behind the targets' rear surfaces were used to record the proton populations accelerated from the target in each simulation. The time-integrated energy spectra recorded by the probe data are shown in figure 5.2. The EPOCH-H simulation was found to produce a similar proton energy spectrum to the capped EPOCH simulations, suggesting that transport phenomena have little effect upon the energy distribution of protons accelerated via the TNSA process.

Significant levels of numerical self-heating resulted in a rapid increase in temperature early in the solid density EPOCH-C simulations (before the peak laser interaction). The large numbers of artificially heated background electrons drifting beyond the initial rear surface resulted in large electric fields, which accelerated the protons to much higher energies than the EPOCH-H and capped EPOCH simulations. Although self-heating was observed in the capped EPOCH simulations (see the background electron temperature lineouts in figure 5.3), this was to a much lesser extent, and thus did not induce significant target expansion. The use of high order particle shape functions (labelled as “EPOCH-C + splines” in figures 5.2 and 5.3) was also found to reduce the

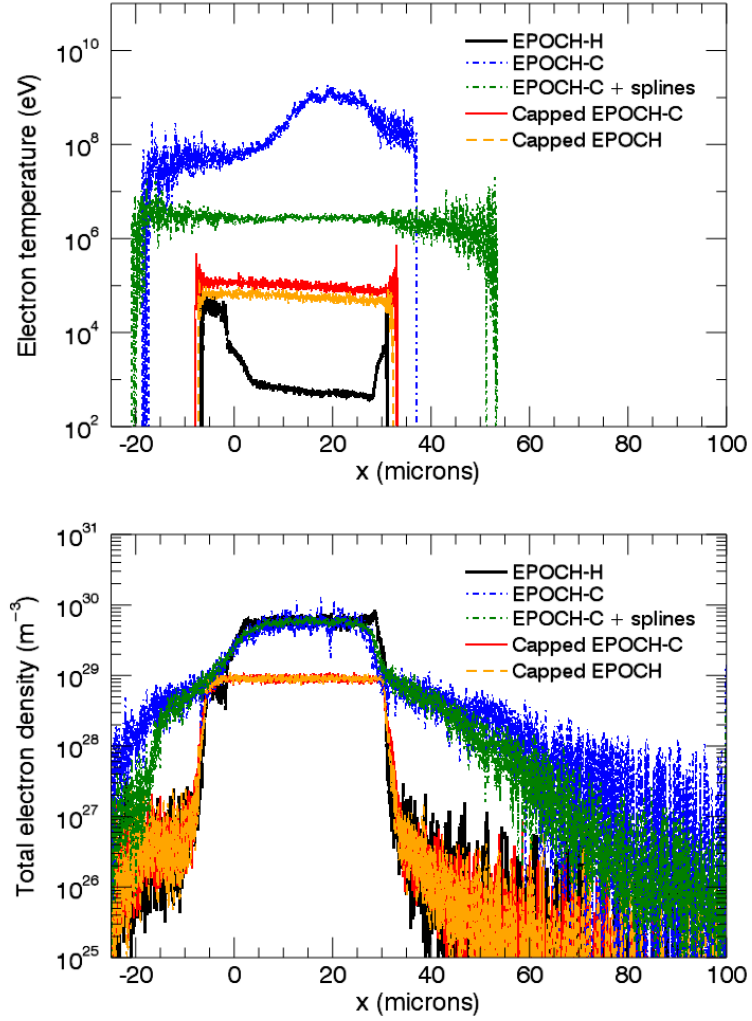


Figure 5.3: Lineouts of the background electron temperatures (top) and total electron densities (bottom) after 2.0 ps, taken from simulations using EPOCH-H (thick, black), solid density EPOCH-C (dot-dash, blue), solid density EPOCH-C with high order particle shape functions (dot-dash, green), capped density EPOCH-C (solid, red) and collisionless, capped density EPOCH (dashed, orange).

level of self-heating, but to a lesser extent than either reducing the peak electron density or making use of EPOCH-H's high density model.

The excess heating in the EPOCH and EPOCH-C simulations can be clearly seen from the lineouts of background electron temperature shown in figure 5.3, and the average temperatures in table 5.1. The higher electron temperature observed in the capped EPOCH-C simulation compared to the collisionless equivalent can be attributed

| Code Version | Peak electron density | Particle shape function | Temperature (keV) | Debye length |
|-----------------|--------------------------|----------------------------|----------------------|-----------------|
| EPOCH | $90n_c$ | First order | 53.6 | $0.24\Delta x$ |
| EPOCH-C | $90n_c$ | First order | 88.7 | $0.30\Delta x$ |
| EPOCH-C | $599n_c$ | First order | 1,036,764.3 | $12.67\Delta x$ |
| EPOCH-C | $599n_c$ | Third order | 2,789.9 | $0.66\Delta x$ |
| EPOCH-H | $599n_c$ | First order | 0.5 | $0.01\Delta x$ |

Table 5.1: Tabulated values of the average temperatures in the range $15\mu\text{m} < x < 25\mu\text{m}$ from figure 5.3.

to the electron-ion scattering acting to thermalise the kinetic energy associated with the return current drift velocity. Without collisions enabled, the return current energy spectrum is more beam-like, and thus at a lower temperature (once the net drift velocity has been subtracted).

The top image of figure 5.3 suggests that the background electrons drift further into the vacuum behind the target when using high order particle shapes. When the density profile resulting from the sum of the background and fast electron populations is plotted (the bottom image in figure 5.3) similar density profiles result for the two solid density EPOCH-C simulations. This is due to the less extreme heating rate resulting in fewer background electrons being promoted to the fast electron population, and thus the rear density ramp appearing to extend further behind the target when only the background electrons are shown.

5.4 Imprinting of fast electron transport effects

The number of protons passing through the probe planes, as a function of position (y) and direction of motion ($\tan\theta = p_y/p_x$), was plotted for each of the simulations (figure 5.4). Collisionless, capped EPOCH returned a narrow range of angles at each position along the probe. Enabling particle collisions resulted in scattering of the protons, and so a more divergent population. The EPOCH-H simulation recorded additional structure on top of the same general behaviour as observed in the EPOCH-C simulation.

The features in the EPOCH-H proton distribution can be attributed to filamentation of the fast population inside the target. Since the field responsible for accelerating protons from the rear surface is generated by fast electrons attempting to leave the target, filamentation of the fast electron population leads to localised enhancements of the sheath fields at the rear surface. These in turn imprint filamentary structures on the accelerated proton distribution [55, 57, 58], as shown in figure 5.5.

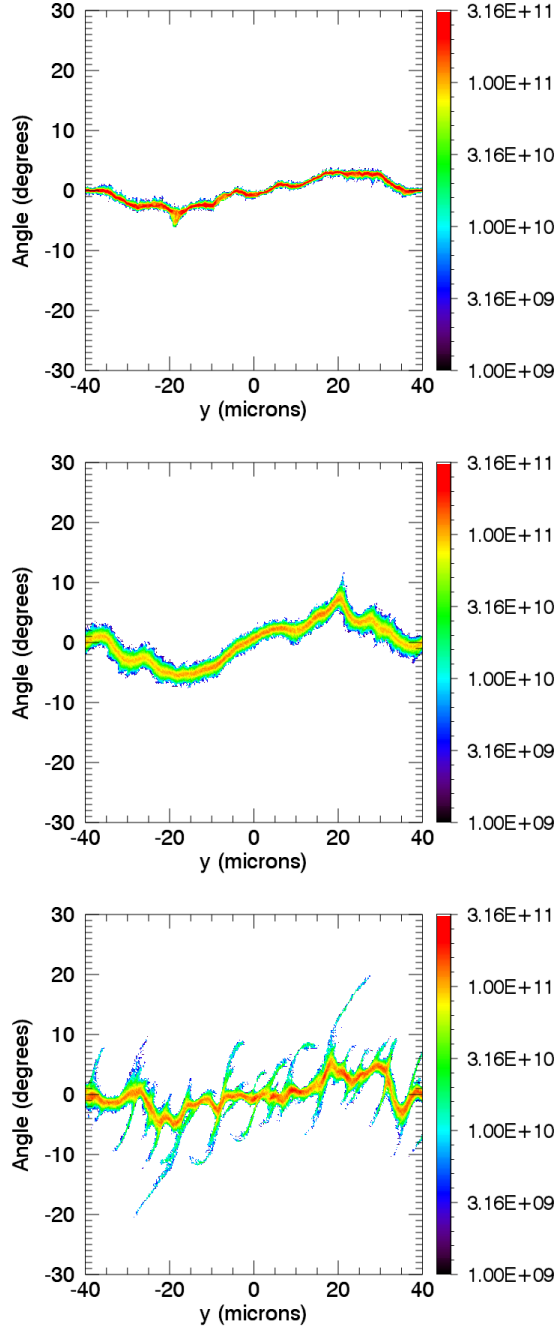


Figure 5.4: Plots of the number of protons (in arbitrary units) recorded by the probes as functions of position along the probe (y) and direction of motion ($\theta = \tan^{-1}(p_y/p_x)$). Results are shown for collisionless, capped EPOCH (top), capped EPOCH-C (middle) and EPOCH-H (bottom) at $t = 2.5$ ps.




Figure 5.5: Fast electron number density inside the target ($0 \leq x \leq 30 \mu\text{m}$), and proton number density near the rear surface ($x > 30 \mu\text{m}$), shortly after the peak laser interaction ($t = 1.5 \text{ ps}$, to demonstrate the imprinting of fast electron filamentation on the proton population accelerated from the rear surface.

5.4.1 Control of proton beam filamentation via target morphology

To investigate the effect of target thickness on the smoothness of the proton population accelerated from the rear surface, additional simulations of 10, 20 and $30 \mu\text{m}$ thick targets were performed. The extent of the simulation domain was reduced in these simulations, since any imprinting of fast electron filamentation on the accelerated protons manifests close to the rear surface. Consequently, the probe plane was placed $25 \mu\text{m}$ behind the rear surface of each target.

A characteristic inter-filament spacing, λ_f , was calculated using the maximum amplitude wavelength resulting from a Fourier transform of the spatial distribution of the protons recorded by the probe plane over 3 ps. To minimise boundary effects, only the results from the middle $40 \mu\text{m}$ of the probe were used. The values of λ_f for various target widths are given in table 5.2 and figure 5.6.

Although higher modes are present in the Fourier transform, the distribution of protons from thinner targets ($\sim 10 \mu\text{m}$) is dominated by longer wavelength modes, and




Figure 5.6: Fourier transforms of the time-integrated transverse spatial distribution of protons recorded passing probe planes $25\text{ }\mu\text{m}$ behind the rear surface of targets with thicknesses of 10 (solid, red), 20 (dashed, green) and $30\text{ }\mu\text{m}$ (dot-dashed, blue).

| Target thickness, $w\text{ (}\mu\text{m)}$ | Filament spacing, $\lambda_f\text{ (}\mu\text{m)}$ |
|---|---|
| 10 | 20.0 |
| 20 | 10.0 |
| 30 | 5.7 |

Table 5.2: Values, for various target widths (w), of the characteristic distance (λ_f) between filaments in the spatial distribution of protons recorded passing a $40\text{ }\mu\text{m}$ -wide plane located $25\text{ }\mu\text{m}$ behind the rear surface of each target.

thus exhibits less filamentation than the distributions recorded from thicker targets. This behaviour is to be expected based on the fast electron density profile shown in figure 5.5, which indicates that the fast electrons must travel some distance into the target (and become more dispersed) before any clear filaments can be observed. The magnetic field strength associated with the filaments was also observed to increase with time and depth into the target. This is consistent with a filamentation instability, in which magnetic fields associated with current filaments act to enhance the filamentation, which reinforce the magnetic fields.

5.4.2 Fast electron filamentation

By examining the magnetic fields inside the target in figure 5.7, the spacing between fast electron filaments was found to be of the order $1\text{ }\mu\text{m}$. This is consistent with the




Figure 5.7: The magnetic field strength inside the target ($x > 0$) and electron density in the pre-plasma ($x < 0$) at $t = 1.0$ ps.

fastest growing mode of the resistive filamentation instability [54]:

$$k_F \sim \frac{\omega_{pf}}{c} \sqrt{\frac{v_f}{v_{th}}}, \quad (5.4)$$

where ω_{pf} is the fast electron plasma frequency, v_f is the fast electron velocity and v_{th} is the background electron thermal velocity. For the conditions found in the EPOCH-H simulations, equation 5.4 predicts an average wavelength of $1.1 \mu\text{m}$. This is in contrast to the Weibel instability [132], which would result in a much shorter wavelength due to the fastest growing mode being related to the background electron skin depth rather than that associated with the fast electrons: $k_W \sim \omega_{pb}/c \rightarrow \lambda_W \sim 43 \text{ nm}$.

The absence of filaments due to the Weibel instability can be attributed to the mesh resolution being too low to resolve the fastest growing mode ($\lambda_W \sim 2\Delta y$). Furthermore, in collisional plasmas resistive filamentation tends to dominate over the Weibel instability due to high wavenumber modes being suppressed by collisional effects [142], and thus even if a higher mesh resolution were used it is unlikely that filaments on the scale of λ_W would be observed in the EPOCH-H and EPOCH-C simulations.

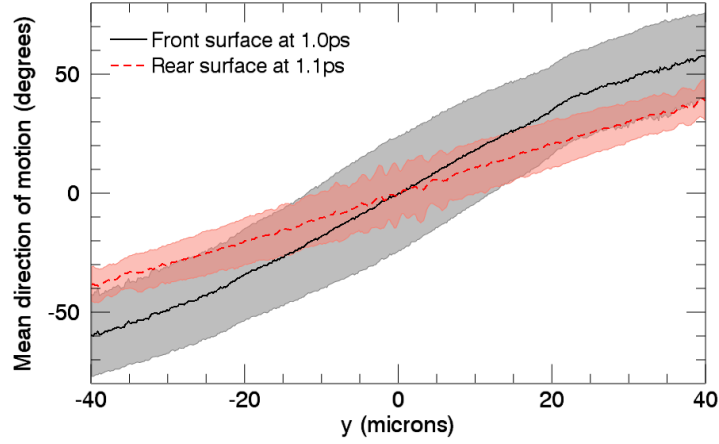


Figure 5.8: Spatial variation in the mean direction of motion of fast electrons recorded by particle probes placed at the top of the front surface density ramp (solid, black) and rear surface (dashed, red) of the target, with the latter being delayed by 0.1 ps to account for the transit time through the target of the fast electrons. The width of the shaded regions indicates the local divergence of the fast electrons.

Figure 5.8 shows the average direction, $\theta = \tan^{-1}(p_y/p_x)$, that the fast electrons are travelling as a function of their lateral (y) position through particle probes placed at the front and rear surfaces of the target ($x = 0$ and $30 \mu\text{m}$, respectively). Any electrons which have been reflected from the y boundaries have been ignored. The extent of the shaded region represents the standard deviation in angles at each position, and can be considered to be a measure of the fast electron divergence as a function of position. In order to account for the time required for the fast electrons recorded at the front surface to reach the rear surface probe, the results recorded by the latter are taken from the output 100 fs later in the simulation.

The reduction in fast electron divergence between the front and rear surfaces seen in figure 5.8 can be shown a purely geometric effect. When the fast electron population recorded by the front surface is projected onto the rear surface, assuming ballistic trajectories, the resultant distribution is in good agreement with that recorded by the rear surface probe, as shown by figures 5.9 and 5.10.

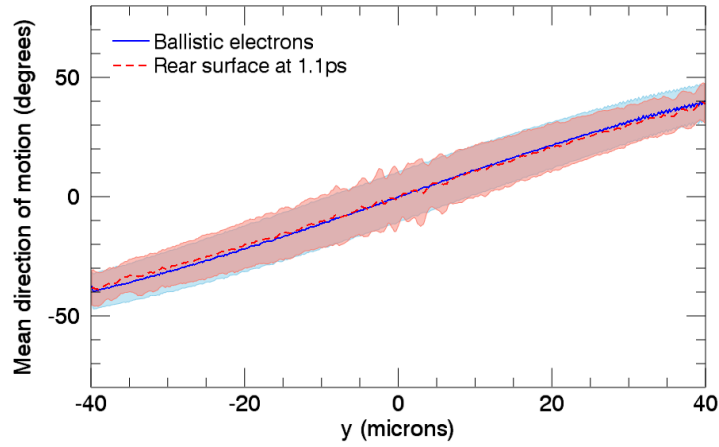


Figure 5.9: Spatial variation in the mean direction of motion of fast electrons recorded by particle probes placed at the top of the front surface density ramp (solid, blue) and rear surface (dashed, red) of the target, with the front surface population having been projected onto the rear surface, assuming ballistic trajectories.

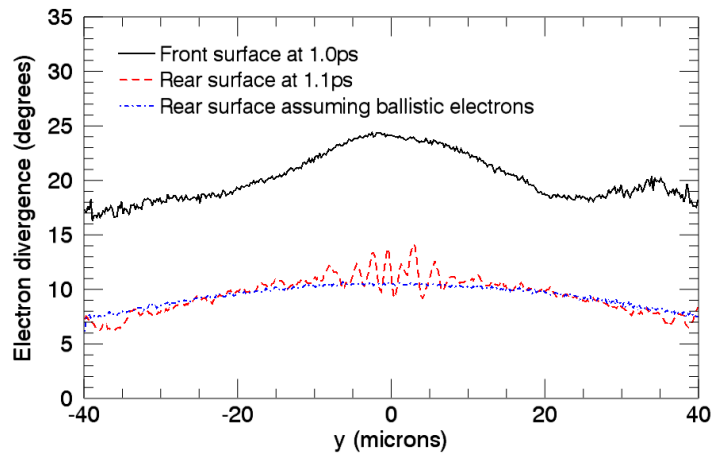


Figure 5.10: Spatial variation in the divergence half-angle of fast electrons recorded by particle probes placed at the top of the front surface density ramp (solid, black) and rear surface (dashed, red) of the target. Also shown are the values of the divergence expected for the fast electrons at the rear surface assuming ballistic trajectories between the front and rear surface probe planes (dot-dash, blue).

5.5 Discussion

The simulations discussed in section 5.3.1 clearly demonstrate the dramatically reduced numerical self-heating present in EPOCH-H simulations, compared to EPOCH-C simulations with the same peak density. The use of high order particle shape functions was found to be less effective than either reducing the peak electron density in the simulation or using EPOCH-H's high density model.

The capped EPOCH, capped EPOCH-C and EPOCH-H fast electron probe data recorded at the front surface of the targets displayed very similar overall behaviour. However, the increased prominence of features in the proton distribution (see section 5.4), and more pronounced filamentation of the fast electron population recorded at the rear surface (as evidenced by the results of section 5.4), suggest that there may be perturbative effects acting on the electron transport in EPOCH-H's high density regions. Observing such effects in standard EPOCH-C simulations would likely require very high mesh and/or particle resolutions in order to reduce the noise in the density, temperature and electric and magnetic fields. However, doing so will impose a significantly larger increase in the computational cost of the simulations than that incurred by using EPOCH-H's high density model. Thus the use of EPOCH-H permits a more efficient means of investigating transport phenomena such as fast electron filamentation.

Care should be taken when comparing the above results to filamentary behaviour observed in experiments. The simulations discussed above were performed in 2D. Thus the particles were constrained to the simulation plane. The inability to spread in the unmodelled z -direction may have resulted in filamentary phenomena manifesting sooner than would otherwise be expected. Furthermore, in 3D the filaments are capable of more complex behaviour, such as twisting around each other.

Recent investigations into fast electron filamentation [57, 58] indicate that phenomena such as ionisation front instabilities and lattice structure can have an important effect on filamentation. Enabling the ionisation module in EPOCH [143] may allow ionisation effects to be included in future simulations. The effect of lattice structure is not currently modelled by PIC codes, but it may be possible to incorporate the anisotropic resistivity by introducing a directional dependence to the collision frequency calculated by the code's collision algorithm. However, care would need to be taken to ensure that the collision frequency returns to being independent of the particles' direction of motion once the target has been heated to a sufficiently high temperature that any lattice structure has disintegrated. Consequently, the simulation results above can be considered as analogous to the vitreous carbon results presented by McKenna et al. [58].

Many of these effects are dominant primarily in the warm dense matter (WDM) regime ($T \sim 1\text{--}100\text{ eV}$). However, the simulations discussed above were performed with an initial temperature of 100 eV . Furthermore, EPOCH-C and EPOCH-H currently assume a ‘Spitzer’ resistivity, with a minimum temperature of $T = 100\text{ eV}$ used in the calculation of the collision frequency and resistivity to reduce the error in the resistivity at low temperatures. Thus, in order to incorporate WDM phenomena more rigorous, but more computationally expensive, low temperature transport models will also be required (e.g. using *ab initio* quantum molecular dynamics simulations to calculate the low temperature resistivity [144]).

Chapter 6

Return current heating of solid density targets

6.1 Modelling material heating

To validate EPOCH-H, it is necessary to compare simulations with experimental results, in addition to standard PIC and/or hybrid techniques. Recent experiments [9, 10] on the Orion laser facility have involved heating a thin sample foil to hundreds of eV in order to make x-ray opacity measurements (see section 1.7.1). To delay hydrodynamic expansion of the sample, and thus permit opacity measurements to be performed at high densities and as close to local thermodynamic equilibrium (LTE) as possible, the foil is usually tamped with several microns of material with a low x-ray opacity, such as plastic or diamond [7–10].

Recent simulations [31], aimed at developing an integrated approach to modelling such experiments, used a combination of a PIC code (EPOCH), fast electron transport code (THOR) and radiation-hydrodynamics code (CORVUS) to model the short-pulse laser-plasma interaction, fast electron energy deposition and hydrodynamic expansion of the target, respectively.

The EPOCH simulations performed by Sircombe et al. [31] required high mesh resolutions ($\Delta x = 13 \text{ nm}$), high particle numbers (84 particles per cell) and linking to a hybrid transport code, such as THOR, in an attempt to reduce the impact of numerical self-heating. This non-physical heating results from a failure to resolve the Debye length at high densities (see section 2.6). Since this requirement does not apply to the high density regions in EPOCH-H simulations, they are much less susceptible to self-heating. Consequently, this adapted version of EPOCH can potentially be used to model the heating of solid density material directly from the short-pulse laser interaction, rather

than having to transfer the fast electron population to a separate, dedicated transport code. Such an approach ensures that the influences of transport phenomena on the laser absorption region are not neglected.

Furthermore, the much higher spatial resolution employed in PIC codes compared to hybrid codes makes it comparatively straightforward to include small-scale features, such as a sub-micron layer of high density material.

6.2 Simulation parameters

It has already been demonstrated that the Ohmic field solver in EPOCH-H drastically reduces the level of numerical self-heating in high density PIC simulations (see chapter 5). This potentially permits investigation of material heating via the return current without requiring any of the more computationally intensive, standard numerical heating mitigation methods (e.g. increasing the mesh and/or particle resolution).

To investigate the ability of EPOCH-H to model return current heating, a detailed 2D simulation was performed with the intention of replicating the heating of aluminium layers embedded in plastic foil targets observed in recent experiments [9, 10] on the Orion laser facility. It was anticipated that this would serve as a means of validating the modified code against experimental results. Consequently, the simulation parameters were chosen so as to reproduce as closely as possible the typical experimental conditions immediately prior to the main short-pulse interaction. Additionally, a THOR simulation of the material heating by the fast electron distribution recorded near the front of the EPOCH-H target was used to provide a point of comparison with previously employed integrated modelling methodologies [31].

6.2.1 EPOCH-H

A 100 J supergaussian pulse with a 0.5 ps full-width-half-maximum (FWHM) was injected from the left boundary of the simulation with a wavelength of 528 nm. The incident pulse was modelled as having a Gaussian transverse intensity profile, with 90% of the laser energy contained within a 50 μm diameter spot.

The target was modelled as fully-ionised plastic at 1.0 g/cc and with an initial temperature of 100 eV. For simplicity, the plastic was modelled using composite ($A = 6.5$, $Z = 3.5$) ions to represent the mix of hydrogen and carbon (hereafter referred to as plastic ions). A 150 nm thick layer, 10 μm into the plastic, was replaced with aluminium at 2.7 g/cc, with an assumed constant and uniform ionisation state of $Z^* = 12$. This value was chosen based on the expected ionisation state for solid density aluminium at 600 eV (the approximate temperature observed experimentally) calculated using the

atomic physics code CASSANDRA [141]. Preceding the target, an exponential density profile of the form $ae^{x/L_1} + (1-a)e^{x/L_2}$ was added, with $a = n_c / (3.5n_{CH})|_{x=0}$. Density scale-length of $L_1 = 3\mu\text{m}$ from vacuum to critical density, and $L_2 = 0.5\mu\text{m}$ from critical density to solid, were assumed to be representative of the front surface conditions. Finally, the target was rotated 10° relative to the incident laser pulse.

The simulation was performed using approximately 6×10^8 pseudo-particles (64 electrons, 16 aluminium ions and 16 plastic ions per cell), with high order particle shapes and non-uniform weights, and a mesh of 2200×4000 cells covering $55 \times 100\mu\text{m}^2$ (approximately half the mesh resolution used previously [31]). The interface between the low and high density regions was assigned to the density range $8 \times 10^{28} \leq n_e \leq 9 \times 10^{28} \text{ m}^{-3}$.

To distinguish between fast and background electrons, any electrons whose velocity exceeded five times the local background electron thermal velocity were moved to the fast electron species list (i.e. $\alpha = 5$, where α is as defined in section 3.2.2). Demotion of fast electrons back to the background population was disabled in order to prevent large amounts of energy being added to the tails of the energy distributions, potentially resulting in deviation from a Maxwellian distribution, and thus unreliable temperature calculations.

6.2.2 THOR

To provide a point of comparison of the heating predicted by EPOCH-H with that from THOR, the fast electrons passing a particle probe $3\mu\text{m}$ behind the top of the pre-plasma density ramp were recorded. The link code LOKI [31] was used to convert the EPOCH-H probe data into a source term for the fast electron transport code THOR [31], using a $600 \times 100 \times 100$ mesh covering position, energy and angle ranges of -75 to $75\mu\text{m}$, 0.1 to 40MeV and -100 to 100° , respectively.

The THOR simulation was performed with a spatial extent of $30 \times 150\mu\text{m}^2$ to encompass the high density plateau in the EPOCH-H simulation. A mesh resolution of $\Delta x = 0.2\mu\text{m}$ and $\Delta y = 0.5\mu\text{m}$ was employed. The background plasma was modelled as solid density plastic (1g/cc) at a constant, uniform ionisation state of 3.5 and initial temperature of 100eV . A $0.2\mu\text{m}$ thick solid aluminium (2.7g/cc) layer with constant ionisation $Z^* = 12$ was embedded $7\mu\text{m}$ into the system (equivalent to $10\mu\text{m}$ into the target). 3200 (200 per core) fast electrons were injected per time-step into the system with their position, energy and direction of motion sampled from the LOKI output described above. A Spitzer resistivity and external tabulated equation of state library were used in the THOR simulation.

Figure 6.1 illustrates the positions of the EPOCH-H and THOR domains in relation




Figure 6.1: Illustration of the positions of the EPOCH-H and THOR domains relative to each other.

to each other. In the comparisons with EPOCH-H discussed below, only the region which overlaps with the simulated domain in the EPOCH-H simulation was considered.

6.3 Material heating comparison

Figure 6.2 displays the background electron temperatures predicted by EPOCH-H and THOR. These plots correspond to the approximate time at which the peak of the laser reaches the front of the target, $t = 0.9$ ps. Similar temperature profiles for both the plastic tamper and aluminium layer are predicted in the two codes.

Radial temperature (figure 6.3) and density variations in the simulated aluminium layer posed a challenge to comparing the simulation results with the experimental value. In an attempt to obtain a single temperature value for comparison with experimental

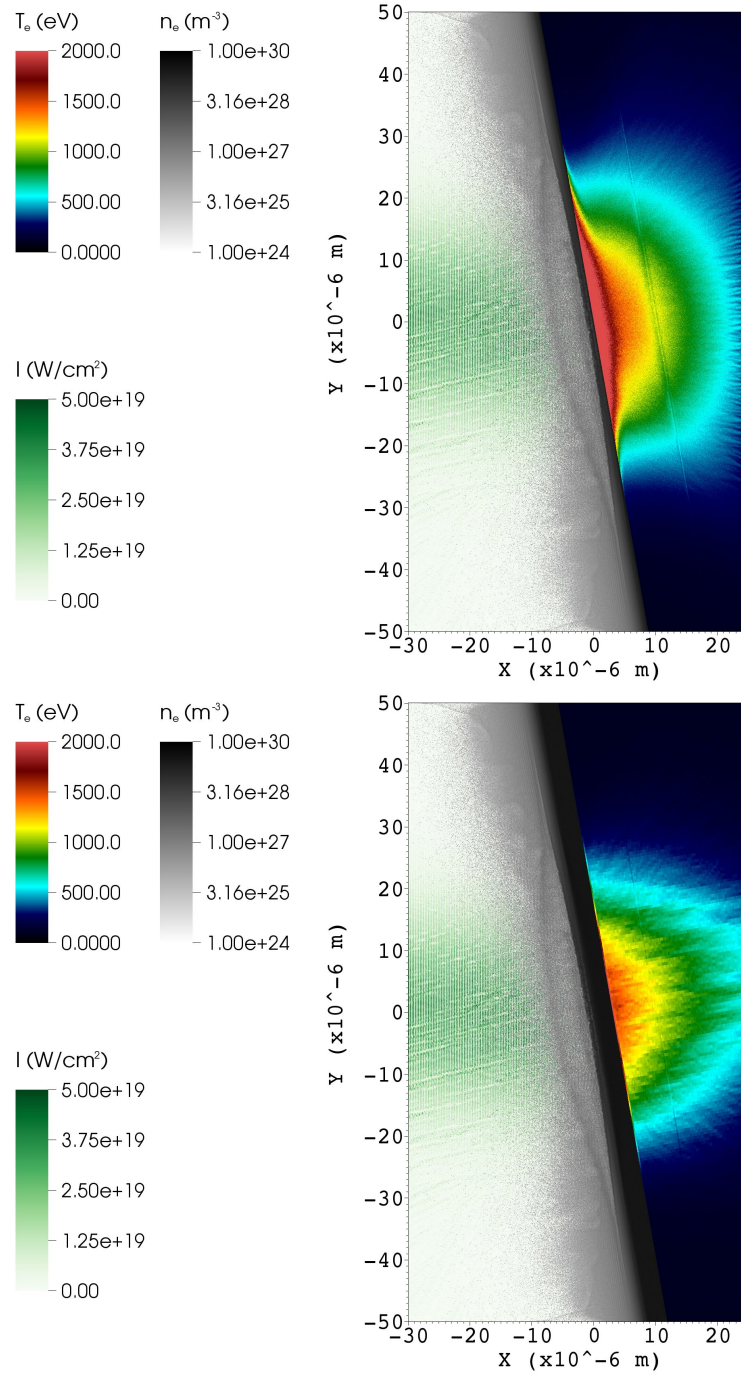


Figure 6.2: Total electron density (n_e) and Poynting flux of the laser (I) from EPOCH-H, and the background electron temperature (T_e) from EPOCH-H (top) and THOR (bottom) shortly after the peak of the laser reaches the target surface ($t = 0.9$ ps).




Figure 6.3: Transverse profiles of the background electron temperature along the middle of the aluminium layer. Results are shown for both the EPOCH-H (solid black) and THOR (dashed red) simulations.

results, the background electron temperature within the layer was spatially averaged (see figure 6.4). Doing so yielded a peak density-weighted average electron temperature of 666 eV from EPOCH-H, and 709 eV from THOR at the same point in time ($t = 1.5$ ps). These are in much closer agreement with the reported experimental values of 600–700 eV [9, 10] than the temperatures predicted in the same region in a similar, collisionless, standard PIC simulation (dot-dash green curve in figure 6.4): $\langle T_e \rangle = 8898$ eV.

Passing the fast electrons recorded in the collisionless EPOCH simulation to THOR resulted in an average electron temperature of 664 eV at $t = 1.5$ ps, in good agreement with the EPOCH-H and related THOR simulations. This suggests that, in this case, transport effects in the EPOCH-H simulation's high density regions did not have a significant impact on the short-pulse laser interaction and fast electron production in the pre-plasma.

The average temperatures, as a function of time, in EPOCH-H and THOR were found to be in good agreement, as shown in figure 6.4. These results suggest that the energy added to the background electrons in EPOCH-H via the electric field set up by the fast electrons, and subsequently 'thermalised' by the code's collision algorithm, was equivalent to the Ohmic heating calculation employed in THOR.

Increased heating, close to the interface region, was observed in the EPOCH-H simulation. This was determined to be due to energy exchange from energetic ions, accelerated by the laser and/or electric field set up by the fast electrons, which entered the high density region late in the simulation. This resulted in additional late-time




Figure 6.4: Evolution, with time, of the spatially-averaged electron temperature in the aluminium layer predicted by EPOCH-H (black, solid) and THOR (red, dashed). Also shown are the temperatures from a similar, collisionless EPOCH simulation (green, dot-dash) linked to THOR (blue, dot-dot-dot-dash).




Figure 6.5: Comparison of the time-averaged background electron temperature evolution from figure 6.4 with that predicted by EPOCH-H (green, dot-dash) and THOR (blue, dot-dot-dot-dash) with energetic plastic ions in EPOCH-H being promoted to a separate, collisionless fast ion population.

($t \gtrsim 2$ ps) heating of the layer. Subsequent simulations, which permitted plastic ions to be promoted to a separate ‘fast ion’ population, in a similar manner to the electron

promotion, did not exhibit any significant late-time heating, and returned average temperatures which were in closer agreement with the THOR-calculated values throughout the simulation, as shown in figure 6.5.

6.4 Line emission comparison

Despite the good agreement noted in the previous section, using the average simulation temperature as a point of comparison may be unreliable, since temperature and density cannot be directly measured during a short-pulse laser experiment. Instead, the experimental temperature and density conditions are usually inferred from the x-ray emission [59], such as the ratio of the intensity of the He- β and Ly- β emission lines. For this reason an alternate means of comparison was adopted which aimed to emulate the experimental approach by calculating the emission spectrum expected from the array of conditions predicted by the EPOCH-H simulation. Comparisons were performed with various combinations of single temperatures and densities to determine the combination which best reproduced the relative heights and widths of the aluminium He- β (1.86 keV) and Ly- β (2.05 keV) emission lines.

The EPOCH-H cells containing aluminium ions were assigned to 50 eV temperature and 0.1 g/cc density bins. The atomic kinetics code FLY [145] was used to calculate the x-ray spectrum from each bin in the range 1.8–2.2 keV, to encompass the emission lines of interest. These were subsequently combined, weighted by the mass in each bin, to produce a total expected emission spectrum. The spectra were time-integrated over a 1.0 ps period centred about the time of peak average temperature ($t = 1.5$ ps).

The simulated spectrum was compared with that calculated using a single temperature and density. The expected emission from the EPOCH-H-predicted conditions showed good agreement with that of aluminium at 740 eV and 2.2 g/cc, as shown in figure 6.6. These values are broadly consistent with the reported experimental values [8–10]. However, uncertainties due to the difficulty in measuring the experimental conditions at the front surface of the target (i.e. the pre-plasma density and temperature profile) immediately prior to the short-pulse interaction, combined with the current inability to model radiative cooling processes in EPOCH-H, present a significant challenge to any further refinement.

Repeating the above procedure with the results of the THOR simulation returned characteristic temperature and density conditions of 770 eV and 2.7 g/cc, respectively (see figure 6.7). It should be noted that since THOR was not coupled to a hydrodynamics code, such as CORVUS [31], no hydrodynamic motion or radiative processes were modelled, and as a result the aluminium layer remained at solid density and high




Figure 6.6: FLY-calculated emission spectrum using temperature and density conditions predicted by EPOCH-H (black, solid). The spectrum for aluminium at 740 eV and 2.2 g/cc (red, dashed) is also shown for comparison. Note that the spectra are normalised to the maximum value within the energy range shown.




Figure 6.7: FLY-calculated emission spectrum using temperature and density conditions predicted by THOR (black, solid) using the fast electron distribution recorded in the EPOCH-H simulation. The spectrum for aluminium at 770 eV and 2.7 g/cc (red, dashed) is also shown for comparison. Note that the spectra are normalised to the maximum value within the energy range shown.

temperature throughout the simulation.

Additional comparisons of the predicted emission spectra, including the He- α and Ly- α lines, are provided in figures 6.8 and 6.9. These figures indicate that the peaks and




Figure 6.8: FLY-calculated emission spectrum using temperature and density conditions predicted by EPOCH-H (black, solid). The spectrum for aluminium at 740 eV and 2.2 g/cc (red, dashed) is also shown for comparison. Note that the spectra are normalised to the value of the Ly- β peak.




Figure 6.9: FLY-calculated emission spectrum using temperature and density conditions predicted by THOR (black, solid) using the fast electron distribution recorded in the EPOCH-H simulation. The spectrum for aluminium at 770 eV and 2.7 g/cc (red, dashed) is also shown for comparison. Note that the spectra are normalised to the value of the Ly- β peak.

widths of the He- α and Ly- α lines, predicted from the EPOCH-H and THOR simulations, also correspond to the same temperature and density conditions given above.




Figure 6.10: Lineout of the z -component of the magnetic fields, parallel to the target normal, along a line $25\text{ }\mu\text{m}$ from the target's central axis. Results are shown for both the EPOCH-H (solid black) and THOR (dashed red) simulations.

6.5 Magnetic fields at layer interface

By considering the resistive term in Ohm's law, it can be shown that magnetic fields can be produced due to gradients in resistivity [53, 146]:

$$\begin{aligned}\frac{\partial \mathbf{B}}{\partial t} &= -\nabla \times \mathbf{E} \approx -\nabla \times (\eta \mathbf{J}_b), \\ &= -\eta \nabla \times \mathbf{J}_b - \nabla \eta \times \mathbf{J}_b.\end{aligned}$$

The resistive magnetic fields resulting from the $\nabla \eta$ term above can potentially deflect fast electrons, inhibiting their transport deeper into the material [147]. Peak magnetic fields of $\sim 1000\text{ T}$ were observed in the EPOCH-H and THOR simulations (see figure 6.10). Similar fields have been observed in simulations of targets with micron-scale layers [148–150], and have resulted in deflection and trapping of fast electrons in the layer, resulting in localised, enhanced heating.

In the case considered here, however, the aluminium layer was much thinner than those previously simulated, resulting in the Larmor radius for electrons at the average predicted energy being much greater than the width of the layer and associated interface fields:

$$r_L = \frac{\gamma m_e v}{q_e B} = 2.86\text{ }\mu\text{m}.$$

Following the approach of Bell et al. [147], it is possible to calculate a magnetisation factor, $M = \int B_z dx / r_M B_c$, in the region of the material interface,

where B_c is the characteristic magnetic field within the region of interest. For the EPOCH-H and THOR simulations discussed above $M \approx 0.04$ and 0.07 , respectively, which is well below the threshold of $M = 1$ for significant fast electron transport inhibition.

The peak magnetic fields in the EPOCH-H and THOR simulations were found to be in good agreement. However, as can be seen in figure 6.10, the grid spacing in the THOR simulation was too coarse to accurately resolve the spatial extent of the resistive fields. Thus the fields in THOR extend over a longer range and result in a higher calculated magnetisation factor than the EPOCH-H simulation.

Thus it appears that the combination of a thin sample layer and relatively low resistivity gradients (compared with, for example, a copper layer of similar thickness embedded in plastic) results in resistive magnetic fields at the material interfaces which are unlikely to have a significant effect of the fast electron transport.

6.6 Summary and discussion

Two-dimensional simulations of short-pulse laser-driven heating of a thin aluminium layer buried in a plastic tamper have been performed in an attempt to validate EPOCH-H against experimental results, as well as the fast electron transport code THOR. The heating observed in the two codes was found to be in good agreement, with similar heating rates observed in the EPOCH-H and THOR simulations. This indicates that, in the absence of numerical self-heating, the scattering and thermalisation of the background electrons which constitute the return current, as performed by EPOCH's collision algorithm, is consistent with THOR's Ohmic heating calculation.

Although the average temperature in the layer in both codes displayed a closer agreement with the experimental results than similar, standard PIC simulations, the presence of radial temperature and density gradients meant that a more rigorous form of comparison was required than that of comparing the average simulated temperature with an inferred experimental value. The x-ray emission spectrum predicted by the atomic kinetics code FLY [145] was found to show good agreement with the spectrum from an aluminium sample at a uniform temperature and density. The temperatures which characterised these spectra were also found to be in reasonable agreement with the average values predicted by EPOCH-H and THOR. Further investigation will be required to determine the form, if any, of the correlation between the mass-weighted average and spectrally-derived temperatures. Since these results suggest that it is possible to characterise the x-ray emission from a non-uniformly heated layer with a single temperature and density, the presence of temperature gradients within the layer may only be discernible experimentally through the use of spatially-resolved diagnostics.

Both the average and spectrally-inferred temperatures were of the order 100 eV higher than those inferred from experimental results. However, these values were much more in keeping with the experimental values than previous PIC modelling attempts. The absence of radiative cooling models in both EPOCH-H and THOR resulted in the peak temperatures being sustained for much longer than would otherwise be expected. Sub-cycling THOR within a radiation-hydrodynamics code, such as CORVUS [31], will allow for the inclusion of radiative effects, and thus provide a more rigorous means of modelling the buried layer's temperature and density history.

Energetic ions accelerated from the pre-plasma in the high density region caused additional heating of the aluminium layer at later times. This heating can be reduced by permitting ions to be promoted to a collisionless energetic population in a similar manner to the electrons. This allowed for closer comparison between the EPOCH-H and THOR-calculated temperatures at later ($t \gtrsim 2$ ps) times, since THOR does not currently include a kinetic ion population.

The need to accumulate particle data onto the grid to evaluate densities and temperatures for the high density field solver led to a much higher computational cost associated with running EPOCH-H with high order particle shapes compared to standard EPOCH at a reduced density. With high order particle shapes the EPOCH-H simulation used approximately 270% of the CPU hours required for a similar standard PIC simulation. This is a significant increase compared to previous tests, with 'triangular' particle shapes, which exhibited a 20% increase. However, it is still a less severe increase than that associated with standard measure for reducing numerical self-heating in solid density PIC simulations.

EPOCH-H provides an efficient means of performing one-off, self-consistent, laser absorption and fast electron transport simulations with peak densities approaching, and in excess of, solid. However, if multiple, similar simulations are required, with the variations between each being located primarily in the transport region (e.g. varying the position/width/composition of a buried layer), it is less computationally intensive to run EPOCH or EPOCH-H once and re-use the fast electron source in multiple THOR simulations. Furthermore, THOR can be sub-cycled within a hydrodynamics code, such as CORVUS, to provide access to radiation transport models, material libraries and NLTE (non-local thermodynamic equilibrium) models, as well as the longer time-scale evolution.

The results presented above have demonstrated the potential for EPOCH-H's use as a tool for modelling the initial stages of a short-pulse laser experiment. Specifically, the interaction of the main laser pulse, production of energetic (fast) electrons, generation of a return current and the associated heating away from the interaction region, without being subject to numerical heating. Some of the wider limitations of

PIC codes are still present, however, and prevent EPOCH-H from being sufficient for modelling the full range of phenomena associated with short-pulse laser experiments. In order to accurately model ablation of the target and shock generation by long-pulse beams, radiative processes, and the comparatively long time-scale ($\sim \text{ns}$) heating and hydrodynamic expansion, additional simulations codes are required. Efforts to link suitable codes together in order to provide a full end-to-end short-pulse experiment modelling capability have met with much initial success [31]. Using models such as that developed by Cohen et al. [32], and implemented in EPOCH-H, it is possible to extend the range of conditions which can be modelled by PIC codes up to, and beyond, solid density. This in turn strengthens the linking from PIC to transport codes.

Chapter 7

Summary

7.1 Background

In order to support, develop and challenge analytic models of short-pulse (picosecond-scale) laser interactions, and optimise experiments so as to make the best possible use of high power laser facilities, an accurate, detailed computational modelling capability is crucial. The various phenomena which take place during short-pulse laser-solid interactions occur across a wide range of spatial and temporal scales, ranging from femtosecond- and sub-micron-scale kinetic effects to nanosecond- and millimetre-scale hydrodynamic motion (see figure 7.1).

A pre-pulse, produced via amplified spontaneous emission (ASE), often precedes the main pulse. This and/or separate long-pulse (nanosecond-scale) beams can ablate the surface of the target to produce a high temperature, low density region (pre-plasma). The high thermal pressure generated by the ablation often results in a shock being driven into the target. After the main pulse interaction, the target undergoes hydrodynamic expansion and disassembly over several nanoseconds, while radiating heat (predominantly as x-rays). Most of these effects can be investigated using a fluid model, which assumes that at each point in space the plasma can be treated as a fluid, characterised by the moments of the distribution function (i.e. density, centre-of-mass velocity and temperature) [82]. When considering the main short-pulse laser absorption and energy transport, however, the system rapidly deviates from the assumed Maxwellian velocity distribution, and the particle motion can become dominated by non-linear wave-particle interactions such as Raman scattering, Landau damping, resonance absorption, Brunel vacuum heating and ponderomotive acceleration. Under these circumstances the fluid approximation employed by hydrocodes is no longer valid, and a more rigorous, kinetic treatment of the plasma is required.

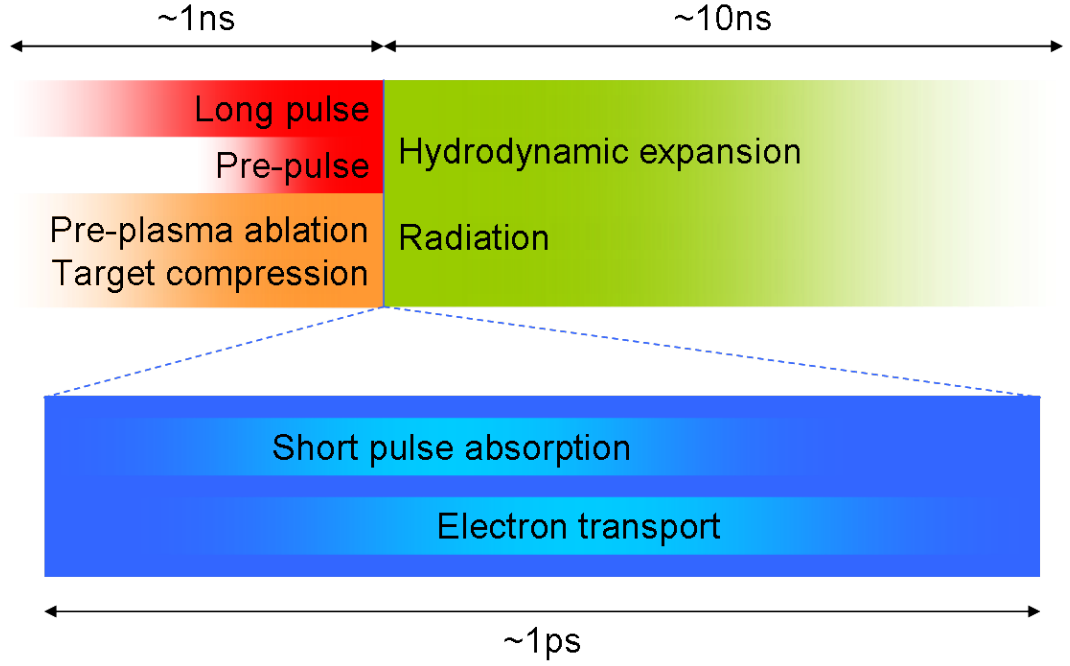


Figure 7.1: Visual representation of the disparate time-scales involved in short-pulse laser interactions.

The fast electron transport and target heating caused by the associated return current typically take place on a similar time-scale to that of the short-pulse laser duration [81]. The myriad processes which may take place during a typical high intensity ($I \gtrsim 10^{18} \text{ Wcm}^{-2}$), picosecond-scale laser interaction with a solid density target present a complex web of interconnected pathways by which the laser energy may be absorbed, distributed amongst the target's constituent particles and subsequently transported deep into the material. Modelling these processes generally requires resolving spatial scales of the order the laser wavelength or Debye length, and temporal scales shorter than the period of the laser oscillation and/or Langmuir waves in the plasma. Since the system often deviates from thermodynamic equilibrium, it is also necessary to resolve the particle distribution functions. The high degree of fidelity required to accurately model the short-pulse laser absorption and fast electron transport results in kinetic models being much more computationally intensive than fluid models. Consequently it is usually not viable to simulate the longer time-scale hydrodynamic phases using a kinetic approach.

Rather than averaging over the local particle distributions, kinetic models employ a statistical treatment which retains more detail about the individual particle properties. Eulerian Vlasov codes, such as IMPACT [89], VALIS [91] and FIDO [92, 93], evolve the

particle distribution functions on an N -dimensional phase-space grid, where $2 \leq N \leq 6$ is the number of spatial and momentum dimensions being simulated (x, y, z, p_x, p_y, p_z), in accordance with the plasma kinetic equation:

$$\frac{\partial f}{\partial t} + \mathbf{v} \cdot \nabla f + \frac{q}{m} (\mathbf{E} + \mathbf{v} \times \mathbf{B}) \cdot \nabla_{\mathbf{v}} f = \left(\frac{\partial f}{\partial t} \right)_{\text{coll}}.$$

Meanwhile, other codes employ more novel approaches, such as using spherical harmonics to represent the distribution function in momentum space (e.g. KALOS [90] and OSHUN [94]).

Particle-in-cell (PIC) codes, such as EPOCH [29], OSIRIS [134], the PSC [133], PICLS [121] and VLPL [151], use Monte-Carlo methods to randomly sample the distribution function. The resultant pseudo-particles are updated based on the Lorentz force, with the electric and magnetic fields being calculated on a discrete, Eulerian mesh:

$$\begin{aligned} \frac{\partial \mathbf{p}}{\partial t} &= q (\mathbf{E} + \mathbf{v} \times \mathbf{B}), \\ \frac{\partial \mathbf{E}}{\partial t} &= c^2 (\nabla \times \mathbf{B}) - \frac{\mathbf{J}}{\epsilon_0}, \\ \frac{\partial \mathbf{B}}{\partial t} &= -\nabla \times \mathbf{E}. \end{aligned}$$

Particle interactions below the grid-scale can be accounted for by the use of a collision algorithm [115, 118, 121] which scatters particles in momentum space, often based on the Spitzer collision frequency [60] of particle i in the rest frame of particle j :

$$\langle \tan^2 \theta \rangle = 2\nu \Delta t = \frac{2q_i^2 q_j^2 n_\beta \ln \Lambda \Delta t}{4\pi \epsilon_0^2 \mu^2 v_r^3},$$

where all of the above terms are as defined in section 2.8.

7.2 The high density algorithm

Standard collisional PIC methods require large numbers of particles in order to accurately model collisional transport phenomena, as demonstrated by the results presented in section 2.8.4. This requirement currently renders detailed PIC modelling of collisional phenomena unfeasible for all but the simplest 1D problems on current HPC platforms. In order to accurately model high intensity short-pulse laser interactions with solid targets, however, 2D (and 3D) PIC simulations are often required which may include collision-dominated regimes. Furthermore, failure to resolve the Debye length in PIC codes causes the plasma to numerically self-heat, with the temperature rising towards a value where

the Debye length is comparable to the grid-scale [95], as discussed in section 2.6. In order to be able to model the intermediate regime between the direct laser interaction (usually modelled using PIC codes, such as EPOCH) and fast electron transport (modelled by hybrid codes, such as THOR) regions an alternative field solver has been developed for EPOCH, based on work by Cohen et al. [32], to produce an extended version of the code: EPOCH-H.

The high density model introduced in chapter 3, which is employed by EPOCH-H, overcomes the need to resolve the Debye length and electron plasma frequency by updating the electric fields via a simplified form of Ohm's law instead of Ampère's law:

$$\mathbf{E} = \eta (\mathbf{J}_b + \mathbf{J}_i) - \mathbf{u}_b \times \mathbf{B} - \frac{\nabla P_b}{q_e n_b}.$$

By retaining a kinetic model for all of the particle species throughout the simulation domain, rather than following the standard hybrid approach of a fluid model for the background material, EPOCH-H is easily able to employ this model in high density grid cells concurrently with the standard Maxwell field solver in the lower density regions of the simulation.

Since many of the effects which must be explicitly modelled in a PIC code for accurate collisional modelling are also implicitly assumed in this algorithm (e.g. that the return current electrons have a drifting Maxwellian energy distribution), it is able to easily model conditions in which collisions are dominant without the high particle numbers required by standard PIC codes; requiring only that the resistivity manifested via the collision algorithm be consistent with the expression evaluated by the field solver. The results presented in sections 2.8.4 and 4.2 demonstrate that the collision algorithm employed in EPOCH-H satisfies this condition.

Tests in one and two dimensions (chapter 4) have shown that EPOCH-H simulations are simultaneously consistent with standard EPOCH in the low density laser interaction regime ($n_e \lesssim 100n_c$) and THOR in the high density region ($n_e \gtrsim 100n_c$). The results of the comparisons with THOR indicate that the Ohmic field solver has extended the density regime which can be modelled effectively so as to bridge the 'transport gap' between the fast electron probe plane and the injection plane of a Monte-Carlo transport code, as indicated in figure 7.2. Although the two-region model employed by EPOCH-H requires an additional $\sim 10\%$ run time compared to equivalent simulations performed using standard EPOCH, these tests demonstrated that employing EPOCH-H is significantly more efficient at reducing numerical self-heating than increasing the EPOCH simulation resolution. However, the retention of the standard PIC model at lower densities still permits some self-heating to occur. Thus despite the advantages granted by EPOCH-H's




Figure 7.2: Diagram indicating the regions of a short-pulse laser interaction which can be effectively modelled using a PIC code, EPOCH (blue), and a hybrid code such as, THOR (red). The Ohmic field solver implemented in EPOCH-H (green) allows the code to bridge the gap between simulations using collisionless PIC and hybrid codes. Underlying image courtesy of N. J. Sircombe, AWE.

two-region model, the usual precautions required for standard PIC simulations must still be taken in the low density, PIC region.

7.3 Effect of high density transport phenomena on proton acceleration

When the fast electrons produced by a high intensity, short-pulse laser-plasma interaction reach, and exit, the rear surface of the target they set up a large electrostatic field ($\sim 10^{12} \text{ Vm}^{-1}$) [69, 71, 73]. This field is capable of accelerating protons, and heavier ions, from the rear surface to produce a low divergence, multi-MeV beam directed normal to the rear surface. Consequently, this process is referred to as target normal sheath acceleration (TNSA) [70]. The protons accelerated from the target in this way can be used for applications including radiographic imaging [16, 17], probing of electric and magnetic fields [15] and heating of a secondary target [77–79].

As a demonstration of EPOCH-H's ability to model transport effects in high density

laser-plasma interactions, a series of simulations were performed comparing collisionless and collisional EPOCH simulations of the TNSA process with EPOCH-H. The results, discussed in section 5.3.1, provided a clear demonstration of the mitigation of numerical self-heating in EPOCH-H simulations.

The proton distributions observed in the capped EPOCH, capped EPOCH-C and EPOCH-H simulations exhibited very similar overall behaviour. However, the inclusion of particle collisions in EPOCH-C resulted in a more divergent proton source compared to EPOCH. The proton distribution in EPOCH-H exhibited not only the higher divergence seen in the EPOCH-C simulation, but also additional filamentary structure. The filaments in the EPOCH-H proton population were attributed to filamentation of the fast electrons as they travelled through the high density target: the filamented fast electron distribution reaching the rear surface of the target resulted in localised enhancements of the rear sheath field, which in turn imprinted a filamentary structure on the protons accelerated by the field.

The absence of fast electron and proton filaments in the EPOCH-C simulation may have been the result of the higher noise levels in the fields and densities. With sufficiently high mesh and/or particle resolutions it should be possible to observe similar filamentation to that found in the EPOCH-H simulations. However, the higher computational expense associated with such simulations means that EPOCH-H may be a more efficient means of investigating high density transport phenomena.

In order to effectively compare fast electron filamentation simulations with experimental observations, three-dimensional simulations are required. The inability of the fast electron population to spread in the unmodelled z direction may have forced the population to undergo filamentation sooner than would normally be expected. In addition, in 2D Cartesian simulations it is not possible to account for more complex behaviour, such as the filaments becoming intertwined.

Recent investigations into the filamentation of fast electron populations produced by short-pulse laser interactions [57, 58] have indicated that effects such as ionisation front instabilities and lattice structure, which are most significant during in the warm dense matter (WDM) regime ($T \sim 1\text{--}100\text{ eV}$), may seed filamentation. Further studies [152, 153] have also demonstrated hollowing of the fast electron beam due to low temperature resistivity effects. The simulations discussed in chapter 5 assumed that the WDM regime was transited quickly, and had little impact on the evolution of the system, by starting with an initial temperature of 100 eV . The collision algorithm and resistivity model currently employed by EPOCH-C and EPOCH-H assume a ‘Spitzer’ resistivity model, which is only valid above temperatures of approximately 100 eV . In order to accurately model the evolution of the system in the WDM regime more rigorous, low temperature

transport models will be required.

7.4 Heating of solid density matter

Following recent experiments on the Orion laser facility [9, 10], simulations were performed in an attempt to replicate the observed heating of an aluminium buried layer. 2D simulations of the heating of an aluminium layer embedded in plastic were performed using EPOCH-H, and the hot electron transport code THOR. As in the previous simulations, the high density model employed by the former was found to dramatically reduce numerical self-heating. Consequently, the average temperature in the layer displayed a much closer agreement with the recent experimental results than a similar, collisionless, standard EPOCH calculation.

A more rigorous comparison was performed by using the atomic kinetics code FLY [145] to calculate the expected x-ray emission spectra for the densities and temperatures predicted by EPOCH-H and THOR. These spectra were found to be in good agreement with the spectrum from an aluminium sample at a similar temperature and density to that inferred from the experimentally recorded spectra.

Although the temperatures predicted by EPOCH-H and THOR were much closer to the reported experimental values than a standard PIC simulation, the temperatures calculated through either spatial averaging or spectral fitting were still of the order 100 eV higher than the typical experimental results. The lack of radiative and hydrodynamic models within both codes will have resulted in higher temperatures over longer time-scales than would otherwise be expected. Until such models can be included, it is still necessary to employ an integrated modelling approach to simulating short-pulse laser experiments, such as those described by Strozzi et al. [30] and Sircombe et al. [31]. The ability to model fast electron transport through the first few microns of a solid target using EPOCH-H, however, provides more confidence in the use of such an approach, by overcoming some of the limitations related to linking EPOCH and THOR, such as the lack of feedback from the transport region to the laser interaction region.

7.5 Future work

A natural extension of the work discussed above would be to implement the high density model in the 3D version of EPOCH. The high computational expense associated with three-dimensional PIC codes make it very difficult to perform high density simulations at a sufficiently high resolution that self-heating effects are not a concern. Extending EPOCH-H to 3D may provide a means by which high density simulations can be performed

at lower mesh resolutions. Furthermore, a 3D version of EPOCH-H will permit a more rigorous investigation of the filamentary phenomena discussed in chapter 5 which did not appear to manifest in the equivalent standard PIC simulations.

A recent publication by Sherlock et al. [136] discussed the possibility that the use of the Spitzer resistivity model [60] in simulations of fast electron transport in high density plasmas may not be entirely valid, and that wakefields in the background electron population, generated by the fast electron bunches accelerated by the laser at twice the laser frequency (see section 1.4.5), may contribute heavily to the material heating. The authors presented a potential algorithm which may account for this source of additional heating. Implementation of a similar plasma wave heating model in EPOCH-H will likely be investigated once further details of the model are available, and pending the outcome of the currently in progress development and testing by Sherlock et al.

At high densities, and particularly in targets with moderate-to-high Z materials, radiative cooling can become an important effect over time-scales greater than a few picoseconds. The implementation of models to include energy losses via thermal bremsstrahlung, fast electron bremsstrahlung and line emission will likely be considered in the future. In addition to providing a more accurate estimate of material heating, such models will allow for the production of synthetic diagnostic images to facilitate comparisons with experimental results.

7.6 Implications for further short-pulse modelling

By modelling both the laser interaction and fast electron transport within the same code the effects of the presence of a high temperature laser interaction region near the high density transport region, and the feedback to the interaction region of phenomena within the transport region, which have generally been neglected by the methods of coupling PIC and hybrid codes used to date, can be captured with EPOCH-H.

Updating the electric fields via Ohm's law, rather than Ampère's law, does not permit electromagnetic or electron plasma waves to propagate. This approximation is valid at high densities, since such waves would normally be heavily damped by Debye screening and collisional effects, respectively. Since these phenomena are transient in the high density plasma, it is generally not necessary to resolve the minimum electron plasma frequency within the high density regions of EPOCH-H simulations. Therefore, when evaluating the plasma frequency-based time-step limit, EPOCH-H uses the maximum value within the PIC region, namely $\omega_{\max} = (n_u e^2 / m_e \epsilon_0)^{1/2}$ (where n_u is as defined in section 3.2.1). In simulations where the plasma frequency is the limiting factor, EPOCH-H's relaxation of this time-step constraint potentially permits larger time-

steps, and thus a reduced run time, compared to similar simulations which employ the traditional PIC method throughout the simulation domain. It should be noted, however, that the potential reduction in simulation run time will not apply to systems where the time-step is constrained by the CFL limit [97] (i.e. where $\Delta x/c < \omega_p^{-1}$).

Based on the results presented in chapters 4–6, EPOCH-H appears to be capable of providing a short-pulse laser modelling capability which:

- is less prone to non-physical numerical effects, such as self-heating, than standard PIC codes;
- is capable of modelling high density plasma kinetics with more relaxed time-step constraints compared to standard PIC codes;
- can efficiently model collision-dominated regimes using currently available computing resources;
- provides a means of ensuring that no significant transport phenomena are neglected when using collisionless PIC simulation data as a fast electron source in Monte-Carlo transport codes;
- can model the effect of fast electron transport phenomena and target heating on the laser interaction physics; and
- is capable of simulating both the production of energetic electrons via laser interaction and their subsequent transport through high density plasma within a single simulation code.

Despite the advantages to using EPOCH-H, it should be noted that the code remains too computationally expensive to be run sub-cycled within a radiation-hydrodynamics code. Thus linking via a dedicated transport code (e.g. THOR) is still required, but, by using EPOCH-H, this can be performed with a greater degree of confidence that no pertinent phenomena have been neglected. Furthermore, EPOCH-H lacks the appropriate material models to handle low temperature effects, and thus relies on the assumption that the target is either initially in the classical plasma state (due to the laser pre-pulse or a separate long-pulse interaction), or transitions the warm dense matter regime sufficiently quickly that it can be ignored. Despite the recent addition of a QED module [138], containing photon tracking, to EPOCH (and thus, by extension, EPOCH-H), the emission of radiation via line emission and bremsstrahlung, which become a significant energy loss mechanism in high Z materials, is not currently modelled.

Overall, it is expected that EPOCH-H's Ohmic field solver will provide an efficient and effective means of simulating short time-scale, high density plasma kinetics, especially compared to the approach employed thus far of increasing the grid and/or particle resolution in a standard collisional PIC simulation.

Appendix A

Derivation of Fokker-Planck equation

Let $\psi(\mathbf{v}, \Delta\mathbf{v})$ be the probability that a particle with initial velocity, \mathbf{v} , experiences a change in velocity, $\Delta\mathbf{v}$, in a time Δt due to short-range scattering. The distribution function, f at a time $t + \Delta t$ can therefore be written as the product of the distribution function at time t with ψ , and then integrated over all possible values of $\Delta\mathbf{v}$:

$$\begin{aligned} f(\mathbf{x}, \mathbf{v}, t) + \Delta t \left(\frac{\partial f(\mathbf{x}, \mathbf{v}, t)}{\partial t} \right)_{\text{coll}} &= f(\mathbf{x}, \mathbf{v}, t + \Delta t) \\ &= \int f(\mathbf{x}, \mathbf{v} - \Delta\mathbf{v}, t) \psi(\mathbf{v} - \Delta\mathbf{v}, \Delta\mathbf{v}) d^3\Delta\mathbf{v}, \end{aligned} \quad (\text{A.1})$$

where all other forces which might evolve the distribution function have been ignored.

The product of f and ψ can be expanded about \mathbf{v} using Taylor's theorem, assuming $\Delta\mathbf{v}$ is small (this is equivalent to the small angle scattering assumption employed later in appendix C):

$$\begin{aligned} f(\mathbf{x}, \mathbf{v}, t + \Delta t) &= \int f(\mathbf{x}, \mathbf{v}, t) \psi(\mathbf{v}, \Delta\mathbf{v}) \\ &\quad - \Delta\mathbf{v} \cdot \nabla_{\mathbf{v}} [f(\mathbf{x}, \mathbf{v}, t) \psi(\mathbf{v}, \Delta\mathbf{v})] \\ &\quad + \frac{1}{2} \Delta\mathbf{v} \Delta\mathbf{v} : \nabla_{\mathbf{v}}^2 [f(\mathbf{x}, \mathbf{v}, t) \psi(\mathbf{v}, \Delta\mathbf{v})] d^3\Delta\mathbf{v}, \end{aligned} \quad (\text{A.2})$$

to second order.

Since $f(\mathbf{x}, \mathbf{v}, t)$ is not a function of $\Delta\mathbf{v}$ and, by definition, $\int \psi(\mathbf{v}, \Delta\mathbf{v}) d^3\Delta\mathbf{v} \equiv$

1, the first term on the right-hand-side of equation A.2 reduces to:

$$\int f(\mathbf{x}, \mathbf{v}, t) \psi(\mathbf{v}, \Delta \mathbf{v}) d^3 \Delta \mathbf{v} = f(\mathbf{x}, \mathbf{v}, t)$$

Re-arranging the second term yields:

$$\begin{aligned} \int -\Delta \mathbf{v} \cdot \nabla_{\mathbf{v}} (f \psi) d^3 \Delta \mathbf{v} &= -\nabla_{\mathbf{v}} \cdot f \int \psi \Delta \mathbf{v} d^3 \Delta \mathbf{v} \\ &= -\nabla_{\mathbf{v}} \cdot (f \langle \Delta \mathbf{v} \rangle), \end{aligned}$$

where $\langle \Delta \mathbf{v} \rangle = \int \psi \Delta \mathbf{v} d^3 \Delta \mathbf{v}$.

Similarly, the third term becomes:

$$\begin{aligned} \int \frac{1}{2} \Delta \mathbf{v} \Delta \mathbf{v} : \nabla_{\mathbf{v}}^2 (f \psi) d^3 \Delta \mathbf{v} &= \frac{1}{2} \nabla_{\mathbf{v}}^2 : f \int \psi \Delta \mathbf{v} \Delta \mathbf{v} d^3 \Delta \mathbf{v} \\ &= \frac{1}{2} \nabla_{\mathbf{v}}^2 : (f \langle \Delta \mathbf{v} \Delta \mathbf{v} \rangle), \end{aligned}$$

where $\langle \Delta \mathbf{v} \Delta \mathbf{v} \rangle = \int \psi \Delta \mathbf{v} \Delta \mathbf{v} d^3 \Delta \mathbf{v}$.

Therefore, equation A.2 can be rewritten as:

$$\begin{aligned} f(\mathbf{x}, \mathbf{v}, t + \Delta t) &= f(\mathbf{x}, \mathbf{v}, t) \\ &\quad - \nabla_{\mathbf{v}} \cdot [f(\mathbf{x}, \mathbf{v}, t) \langle \Delta \mathbf{v} \rangle] \\ &\quad + \frac{1}{2} \nabla_{\mathbf{v}}^2 : [f(\mathbf{x}, \mathbf{v}, t) \langle \Delta \mathbf{v} \Delta \mathbf{v} \rangle]. \end{aligned} \quad (\text{A.3})$$

Substituting equation A.3 into equation A.1 yields an expression for the rate of change of the distribution function due to collisions:

$$\left(\frac{\partial f}{\partial t} \right)_{\text{coll}} = -\nabla_{\mathbf{v}} \cdot \left[f \frac{\langle \Delta \mathbf{v} \rangle}{\Delta t} \right] + \frac{1}{2} \nabla_{\mathbf{v}}^2 : \left[f \frac{\langle \Delta \mathbf{v} \Delta \mathbf{v} \rangle}{\Delta t} \right]. \quad (\text{A.4})$$

The $\langle \Delta \mathbf{v} \rangle / \Delta t$ (coefficient of dynamical friction) term in equation A.4 is a drag term, which slows/accelerates fast/slow-moving particles towards the mean velocity of the distribution. Meanwhile, the $\langle \Delta \mathbf{v} \Delta \mathbf{v} \rangle / \Delta t$ (coefficient of dynamical diffusion) term acts as a scattering term, broadening the distribution in velocity-space. Together, these terms act to drive f towards an equilibrium distribution (often Maxwellian).

A.1 The Landau equation

The Fokker-Planck coefficients can be recast in terms of the Rosenbluth potentials [154]:

$$\begin{aligned}\frac{\langle \Delta \mathbf{v} \rangle}{\Delta t} &= \Gamma \left(1 + \frac{m_i}{m_j} \right) \int f'(\mathbf{v}') \frac{\mathbf{u}}{|\mathbf{u}|^3} d^3 v', \\ \frac{\langle \Delta \mathbf{v} \Delta \mathbf{v} \rangle}{\Delta t} &= \Gamma \int f'(\mathbf{v}') \left(\frac{\mathbf{I}}{|\mathbf{u}|} - \frac{\mathbf{u}\mathbf{u}}{|\mathbf{u}|^3} \right) d^3 v'.\end{aligned}$$

where $\Gamma = q_i^2 q_j^2 \ln \Lambda / 4\pi \epsilon_0^2 m_i^2$ and $\mathbf{u} = \mathbf{v} - \mathbf{v}'$. Here \mathbf{v} is the velocity of particle i which is scattering off of particle j , moving with velocity \mathbf{v}' .

Using the following identity:

$$\nabla_{\mathbf{v}} \cdot \left(\frac{\mathbf{I}}{|\mathbf{u}|} - \frac{\mathbf{u}\mathbf{u}}{|\mathbf{u}|^3} \right) = -\nabla_{\mathbf{v}'} \cdot \left(\frac{\mathbf{I}}{|\mathbf{u}|} - \frac{\mathbf{u}\mathbf{u}}{|\mathbf{u}|^3} \right) = -\frac{2\mathbf{u}}{|\mathbf{u}|^3}$$

and integrating the expression for \mathbf{H} by parts, the Fokker-Planck equation becomes the Landau equation:

$$\left(\frac{\partial f(\mathbf{v})}{\partial t} \right)_{\text{coll}} = \nabla_{\mathbf{v}} \cdot \frac{\Gamma m_i}{2} \int \left(\frac{\mathbf{I}}{|\mathbf{u}|} - \frac{\mathbf{u}\mathbf{u}}{|\mathbf{u}|^3} \right) \left(\frac{f(\mathbf{v})}{m_j} \nabla_{\mathbf{v}'} f'(\mathbf{v}') - \frac{f'(\mathbf{v}')}{m_i} \nabla_{\mathbf{v}} f(\mathbf{v}) \right) d^3 v'. \quad (\text{A.5})$$

Appendix B

Derivation of radiating boundary conditions

For simplicity, we normalise Maxwell's equations using the following:

$$\begin{aligned}\tilde{\mathbf{x}} &= \mathbf{x}\omega/c, \\ \tilde{t} &= \omega t, \\ \tilde{\mathbf{E}} &= \mathbf{E}/E_0, \\ \tilde{\mathbf{B}} &= \mathbf{B}c/E_0, \\ \tilde{\mathbf{J}} &= \mathbf{J}/\epsilon_0\omega E_0,\end{aligned}$$

where ω is the laser frequency and E_0 is the peak laser amplitude, and the tildes mark the normalised values.

Thus equations 2.7 and 2.8, written in component form, become:

$$\begin{aligned}\partial_{\tilde{t}}\tilde{E}_x &= \partial_{\tilde{y}}\tilde{B}_z - \partial_{\tilde{z}}\tilde{B}_y - \tilde{J}_x, \\ \partial_{\tilde{t}}\tilde{E}_y &= \partial_{\tilde{z}}\tilde{B}_x - \partial_{\tilde{x}}\tilde{B}_z - \tilde{J}_y, \\ \partial_{\tilde{t}}\tilde{E}_z &= \partial_{\tilde{x}}\tilde{B}_y - \partial_{\tilde{y}}\tilde{B}_x - \tilde{J}_z, \\ \partial_{\tilde{t}}\tilde{B}_x &= \partial_{\tilde{z}}\tilde{E}_y - \partial_{\tilde{y}}\tilde{E}_z, \\ \partial_{\tilde{t}}\tilde{B}_y &= \partial_{\tilde{x}}\tilde{E}_z - \partial_{\tilde{z}}\tilde{E}_x, \\ \partial_{\tilde{t}}\tilde{B}_z &= \partial_{\tilde{y}}\tilde{E}_x - \partial_{\tilde{x}}\tilde{E}_y.\end{aligned}$$

These can be combined to give:

$$\partial_{\tilde{t}} \left(\tilde{E}_y \pm \tilde{B}_z \right) \pm \partial_{\tilde{x}} \left(\tilde{E}_y \pm \tilde{B}_z \right) = \pm \partial_{\tilde{y}} \tilde{E}_x + \partial_{\tilde{z}} \tilde{B}_x - \tilde{J}_y, \quad (\text{B.1})$$

$$\partial_{\tilde{t}} \left(\tilde{E}_z \mp \tilde{B}_y \right) \pm \partial_{\tilde{x}} \left(\tilde{E}_z \mp \tilde{B}_y \right) = \pm \partial_{\tilde{z}} \tilde{E}_x - \partial_{\tilde{y}} \tilde{B}_x - \tilde{J}_z, \quad (\text{B.2})$$

$$\partial_{\tilde{t}} \left(\tilde{E}_x \mp \tilde{B}_z \right) \pm \partial_{\tilde{y}} \left(\tilde{E}_x \mp \tilde{B}_z \right) = \pm \partial_{\tilde{x}} \tilde{E}_y - \partial_{\tilde{z}} \tilde{B}_y - \tilde{J}_x, \quad (\text{B.3})$$

$$\partial_{\tilde{t}} \left(\tilde{E}_z \pm \tilde{B}_x \right) \pm \partial_{\tilde{y}} \left(\tilde{E}_z \pm \tilde{B}_x \right) = \pm \partial_{\tilde{z}} \tilde{E}_y + \partial_{\tilde{x}} \tilde{B}_y - \tilde{J}_z, \quad (\text{B.4})$$

$$\partial_{\tilde{t}} \left(\tilde{E}_x \pm \tilde{B}_y \right) \pm \partial_{\tilde{z}} \left(\tilde{E}_x \pm \tilde{B}_y \right) = \pm \partial_{\tilde{x}} \tilde{E}_z + \partial_{\tilde{y}} \tilde{B}_z - \tilde{J}_x, \quad (\text{B.5})$$

$$\partial_{\tilde{t}} \left(\tilde{E}_y \mp \tilde{B}_x \right) \pm \partial_{\tilde{z}} \left(\tilde{E}_y \mp \tilde{B}_x \right) = \pm \partial_{\tilde{y}} \tilde{E}_z - \partial_{\tilde{x}} \tilde{B}_z - \tilde{J}_y. \quad (\text{B.6})$$

If the gradients of the x -components of the fields in equation B.1 are assumed to be zero, and the case of $\mathbf{J} = 0$ (i.e. vacuum) is considered, then this equation becomes: $\partial_{\tilde{t}} \left(\tilde{E}_y \pm \tilde{B}_z \right) = \mp \partial_{\tilde{x}} \left(\tilde{E}_y \pm \tilde{B}_z \right)$, which, after taking the gradient in x and substituting, results in the wave equation:

$$\partial_{\tilde{t}}^2 \left(\tilde{E}_y \pm \tilde{B}_z \right) = \mp \partial_{\tilde{x}}^2 \left(\tilde{E}_y \pm \tilde{B}_z \right). \quad (\text{B.7})$$

Equation B.7 can be considered as describing a forward/backward-travelling wave, polarised in the y -direction, propagating parallel to the x -axis (i.e. a laser pulse propagating through the x_{\min}/x_{\max} boundary).

Following a similar approach with equation B.2 results in the wave polarised in the z -direction. Repeating this approach for equations B.3–B.6 leads to wave equations for forward and backward propagating waves in the y and z -directions.

Let $2\tilde{F}_{yz} = \tilde{E}_y \pm \tilde{B}_z$ be the normalised electric field of a wave travelling parallel to the x -axis, and polarised along the y -axis. On a discrete mesh, with F_{yz} being defined at the same positions as the electric field components, this becomes:

$$2(\tilde{F}_{yz})_{jkl}^n = (\tilde{E}_y)_{jkl}^n \pm \left((\tilde{B}_z)_{jkl}^n + (\tilde{B}_z)_{j-1kl}^n \right) / 2,$$

where the magnetic fields have been interpolated to the electric field positions.

Using the FDTD forms of Faraday's law (i.e. equation 2.12), and substituting

the expressions for the normalised fields, this can be rewritten as:

$$\begin{aligned}
4(F_{yz})_{jkl}^n &= c^2 \Delta t \left(\nabla_z^- (B_x)_{jkl}^n - \nabla_x^- (B_z)_{jkl}^n - \mu_0 (J_y)_{jkl}^n \right) + 2(E_y)_{jkl}^{n-1/2} \\
&\quad \pm c \left((B_z)_{jkl}^n + (B_z)_{j-1kl}^n \right), \\
&= \frac{c^2 \Delta t}{\Delta z} \left((B_x)_{jkl}^n - (B_x)_{jkl-1}^n \right) - \frac{c^2 \Delta t}{\Delta x} \left((B_z)_{jkl}^n - (B_z)_{j-1kl}^n \right) \\
&\quad - c^2 \Delta t \mu_0 (J_y)_{jkl}^n + 2(E_y)_{jkl}^{n-1/2} \pm c (B_z)_{jkl}^n \pm c (B_z)_{j-1kl}^n, \\
\therefore 4(F_{yz})_{jkl}^n &= L_z \left((B_x)_{jkl}^n - (B_x)_{jkl-1}^n \right) - (L_x \mp c) (B_z)_{jkl}^n + (L_x \pm c) (B_z)_{j-1kl}^n \\
&\quad - \Delta t (J_y)_{jkl}^n / \epsilon_0 + 2(E_y)_{jkl}^{n-1/2}.
\end{aligned}$$

This can then be re-arranged to obtain a boundary condition on B_z at the x_{\min} ($j = 1$) boundary:

$$\begin{aligned}
(B_z)_{0kl}^n &= \left[4(F_{yz})_{1kl}^n - 2(E_y)_{1kl}^{n-1/2} - L_z \left((B_x)_{1kl}^n - (B_x)_{1kl-1}^n \right) \right. \\
&\quad \left. + \frac{\Delta t}{\epsilon_0} (J_y)_{1kl}^n + (L_x - c) (B_z)_{1kl}^n \right] \left(\frac{1}{L_x + c} \right),
\end{aligned}$$

where the $2F_{yz} = E_y + cB_z$ solution has been used, since $2F_{yz} = E_y - cB_z$ implies a wave heading away from the simulation domain. Thus, at the $x = x_{\max}$ boundary ($j = N_x$):

$$\begin{aligned}
(B_z)_{N_x kl}^n &= \left[-4(F_{yz})_{N_x kl}^n + 2(E_y)_{N_x kl}^{n-1/2} + L_z \left((B_x)_{N_x kl}^n - (B_x)_{N_x kl-1}^n \right) \right. \\
&\quad \left. - \frac{\Delta t}{\epsilon_0} (J_y)_{N_x kl}^n + (L_x - c) (B_z)_{N_x-1kl}^n \right] \left(\frac{1}{L_x + c} \right).
\end{aligned}$$

Repeating this approach with equation B.2 will yield x -boundary conditions for waves polarised along the z -axis. Similarly, equations B.3–B.6 can be used to derive conditions for the y and z -boundaries.

Appendix C

Derivation of collision parameters

C.1 Collision frequency

Consider an electron with velocity $\mathbf{v} = (|\mathbf{v}| \ll c, 0, 0)$ approaching an ion, offset by some distance, b (the impact parameter). Assume that the deflection of the electron's velocity due to the influence of the ion's electric field is small compared to the electron's initial velocity ($|\Delta \mathbf{v}| / |\mathbf{v}| \ll 1$). Also assume an infinitely massive ion (the Lorentz approximation).

The electron's position as a function of time is given by $\mathbf{r} = (|\mathbf{v}|t, b, 0)$, where $t = 0$ is the time of closest approach, at which the electron-ion separation is equal to b . The force on the electron due to the ion's electric field is given by the standard Coulomb force:

$$\mathbf{F} = q_e \mathbf{E} = -q_e \nabla \phi = -\frac{Z^* q_e^2 \mathbf{r}}{4\pi\epsilon_0 |\mathbf{r}|^3}.$$

Therefore, as the electron passes the ion, it experiences an impulse, $m_e \Delta \mathbf{v} = \int_{t \rightarrow -\infty}^{t \rightarrow \infty} \mathbf{F} dt$.

The component of the change in the electron's velocity perpendicular to its initial velocity (i.e. in the y - z plane) can be calculated as follows:

$$\begin{aligned} m_e \Delta \mathbf{v} &= - \int_{-\infty}^{\infty} \frac{Z^* q_e^2 \mathbf{r}}{4\pi\epsilon_0 |\mathbf{r}|^3} dt, \\ \Delta \mathbf{v}_{\perp} &= - \frac{Z^* q_e^2 b \hat{\mathbf{y}}}{4\pi\epsilon_0 m_e} \int_{-\infty}^{\infty} \frac{dt}{(b^2 + v^2 t^2)^{3/2}}, \\ \Delta \mathbf{v}_{\perp} &= - \frac{2Z^* q_e^2}{4\pi\epsilon_0 m_e b |\mathbf{v}|} \hat{\mathbf{y}}. \end{aligned}$$

Since the Lorentz approximation assumes an infinitely massive ion, the ion does not recoil from the interaction. Therefore, conservation of energy can be used to show

that $\mathbf{v} \cdot \Delta \mathbf{v} = -\frac{1}{2} \Delta \mathbf{v} \cdot \Delta \mathbf{v}$. Since $\mathbf{v} \cdot \Delta \mathbf{v} \equiv |\mathbf{v}| \Delta v_{\parallel}$, an expression for the parallel (x) component of $\Delta \mathbf{v}$ can be obtained:

$$\begin{aligned} \Delta v_{\parallel} &\approx -\frac{1}{2|\mathbf{v}|} \Delta \mathbf{v}_{\perp} \cdot \Delta \mathbf{v}_{\perp}, \\ &= -\frac{1}{2|\mathbf{v}|} \frac{4Z^{*2}q_e^4}{(4\pi\epsilon_0)^2 m_e^2 b^2 |\mathbf{v}|^2}, \\ &= -\frac{2Z^{*2}q_e^4}{(4\pi\epsilon_0)^2 m_e^2 b^2 |\mathbf{v}|^3}. \end{aligned}$$

Now consider an electron moving through a collection of ions with number density n_i . The number of ions passed per unit time in cylindrical geometry is $n_i |\mathbf{v}| \int_0^{2\pi} d\phi \int_0^{\infty} b db$. The ensemble average Coulomb collision force experienced by the electron parallel to its direction of motion is then given by:

$$\begin{aligned} \langle F_{\parallel} \rangle &\equiv m_e \frac{\langle \Delta v_{\parallel} \rangle}{\Delta t}, \\ &= n_i |\mathbf{v}| \int_0^{2\pi} d\phi \int_0^{\infty} b db m_e \Delta v_{\parallel}, \\ &= n_i |\mathbf{v}| m_e 2\pi \left(-\frac{2Z^{*2}q_e^4}{(4\pi\epsilon_0)^2 m_e^2 |\mathbf{v}|^3} \right) \int_0^{\infty} \frac{db}{b}, \\ &= -\frac{Z^{*2}q_e^4 n_i}{4\pi\epsilon_0^2 m_e |\mathbf{v}|^2} \int_0^{\infty} \frac{db}{b}, \end{aligned}$$

where $\Delta t \sim b/v$ is the typical interaction time for an individual collision. It should be noted that since the scattering in the y - z plane is assumed to be isotropic, $\langle \mathbf{F}_{\perp} \rangle = 0$.

In order to evaluate the integral with respect to the impact parameter, b , it is necessary to restrict the range of possible values. Since the plasma acts to screen electric fields over length-scales greater than the Debye length, λ_D (see section 1.2.3), ions beyond this distance will have a negligible effect upon the electron's motion. Consequently, this distance can be taken as the upper limit of the integral range: $b_{\max} = \lambda_D = (\epsilon_0 k_B T_e / q_e^2 n_e)^{1/2}$.

If the electron passes sufficiently close to an ion, the Coulomb potential energy, $Z^* q_e^2 / 4\pi\epsilon_0 |\mathbf{r}|$, is comparable to the electron's kinetic energy, $\Delta v_{\parallel} \sim |\Delta \mathbf{v}_{\perp}|$, resulting in a scattering angle of $\sim 90^\circ$. This invalidates the small angle scattering approximation invoked above. Therefore the minimum impact parameter can be taken to be the value

of $|\mathbf{r}|$ for which $\Delta v_{\parallel} = |\Delta \mathbf{v}_{\perp}|$:

$$\begin{aligned} -\frac{2Z^*q_e^2}{4\pi\epsilon_0 m_e b_{\min} |\mathbf{v}|} &= -\frac{2Z^{*2}e^4}{(4\pi\epsilon_0)^2 m_e^2 \lambda_L^2 |\mathbf{v}|^3}, \\ \lambda_L &= \frac{Z^*q_e^2}{4\pi\epsilon_0 m_e |\mathbf{v}|^2}, \\ &\approx \frac{Z^*q_e^2}{12\pi\epsilon_0 k_b T_e}, \end{aligned}$$

where a Maxwellian electron distribution and equipartition theorem ($mv^2/2 = 3k_b T/2$) have been assumed.

For relativistic particles, however, the minimum impact parameter is given by the de Broglie wavelength:

$$\lambda_{dB} = \frac{h}{\sqrt{\gamma^2 - 1}mc},$$

where $\gamma = 1 + k_B T_e / m_e c^2$. Consequently, the value of b_{\min} is taken as the maximum of λ_L and λ_{dB} :

$$b_{\min} = \text{MAX} \left(\frac{Z^*q_e^2}{12\pi\epsilon_0 k_b T_e}, \frac{h}{\sqrt{(1 + k_B T_e / m_e c^2)^2 - 1}m_e c} \right)$$

As a result, $\int_{b_{\min}}^{b_{\max}} b^{-1} db = \ln b_{\max} / b_{\min} = \ln \Lambda$, where

$$\ln \Lambda = \ln \left(\frac{12\pi (\epsilon_0 k_b T_e)^{3/2}}{Z^* q_e^3 \sqrt{n_e}} \right),$$

in the slow particle limit and

$$\ln \Lambda = \ln \left(\frac{m_e c}{q_e h} \sqrt{\frac{\epsilon_0 k_b T_e (\gamma^2 - 1)}{n_e}} \right),$$

for relativistic particles, where $\gamma = 1 + k_B T_e / m_e c^2$.

Thus:

$$\begin{aligned}
m_e \frac{\partial v_{\parallel}}{\partial t} &= \langle F_{\parallel} \rangle = m_e \frac{\Delta v_{\parallel}}{\Delta t} = -\frac{Z^{*2} q_e^4 n_i \ln \Lambda}{4\pi \epsilon_0^2 m_e |\mathbf{v}|^2}, \\
&= -\frac{Z^{*2} q_e^4 n_i \ln \Lambda}{4\pi \epsilon_0^2 m_e^2 |\mathbf{v}|^3} m_e v_{\parallel} = -\nu m_e v_{\parallel}. \\
\therefore \nu &= \frac{Z^{*2} q_e^4 n_i \ln \Lambda}{4\pi \epsilon_0^2 m_e^2 |\mathbf{v}|^3}.
\end{aligned} \tag{C.1}$$

C.2 Scattering angle

Since scattering is a random process, central limit theorem indicates that the scattering angles relative to the initial velocity vector, $\Delta\theta = \arcsin |\Delta\mathbf{v}_{\perp}| / |\mathbf{v} + \Delta\mathbf{v}|$, should form a Gaussian distribution:

$$f(\Delta\theta) \propto \exp\left(-\frac{\Delta\theta^2}{2s^2}\right),$$

where $s^2 = \langle \Delta\theta^2 \rangle - \langle \Delta\theta \rangle^2$ is the variance. Since $\langle \Delta\theta \rangle = 0$, $s^2 = \langle \Delta\theta^2 \rangle$.

It can be shown that $\Delta v_{\parallel} / |\Delta\mathbf{v}_{\perp}| = \tan \theta / 2$. Using small angle approximation, $\Delta\theta \approx 2\Delta v_{\parallel} / |\Delta\mathbf{v}_{\perp}|$, and thus

$$\begin{aligned}
s^2 &= \langle \Delta\theta^2 \rangle \approx \left\langle \frac{4\Delta v_{\parallel}^2}{|\Delta\mathbf{v}_{\perp}|^2} \right\rangle = \Delta t n_i |\mathbf{v}| \int_0^{2\pi} d\phi \int_{b_{\min}}^{b_{\max}} \frac{4\Delta v_{\parallel}^2}{|\Delta\mathbf{v}_{\perp}|^2} b db, \\
&= \frac{2Z^{*2} q_e^4 n_i \Delta t}{4\pi \epsilon_0^2 m_e^2 |\mathbf{v}|^3} \int_{b_{\min}}^{b_{\max}} \frac{db}{b}, \\
&= 2 \frac{Z^{*2} q_e^4 n_i \ln \Lambda}{4\pi \epsilon_0^2 m_e^2 |\mathbf{v}|^3} \Delta t. \\
\therefore s &\approx \sqrt{2\nu \Delta t}.
\end{aligned} \tag{C.2}$$

This is equivalent to the Fokker-Planck approach [121, 123, 155], since the method employed here of multiple, small angle scattering events is equivalent to considering random deflections due to the simultaneous effects of the presence of many nearby particles [156].

Bibliography

- [1] S. C. Wilks, W. L. Kruer, M. Tabak, and A. B. Langdon. Absorption of ultra-intense laser pulses. *Phys. Rev. Lett.*, 69:1383, 1992.
- [2] F. N. Beg, A. R. Bell, A. E. Dangor, C. N. Danson, A. P. Fews, M. E. Glinsky, B. A. Hammel, P. Lee, P. A. Norreys, and M. Tatarakis. A study of picosecond laser-solid interactions up to 10^{19} Wcm^{-2} . *Phys. Plasmas*, 4:447, 1997.
- [3] R. F. Benjamin, G. H. McCall, and A. Wayne Ehler. Measurement of return current in a laser-produced plasma. *Phys. Rev. Lett.*, 42:890, 1979.
- [4] A. R. Bell, J. R. Davies, S. Guérin, and H. Ruhl. Fast-electron transport in high-intensity short-pulse laser-solid experiments. *Plasma Phys. Control. Fusion*, 39:653, 1997.
- [5] B. K. F. Young, H. A. Baldis, R. P. Drake, E. M. Campbell, and K. G. Estabrook. Measurement of x-ray emission and thermal transport in near-solid-density plasmas heated by 130 fs laser pulses. *Phys. Rev. E*, 58:4929, 1998.
- [6] K. Eidmann, U. Andiel, F. Pisani, P. Hakel, R. C. Mancini, G. C. Junkel-Vives, J. Abdallah, and K. Witte. K-shell spectra from hot dense aluminum layers buried in carbon and heated by ultrashort laser pulses. *J. Quant. Spectrosc. Radiat. Transfer*, 81:133, 2003.
- [7] D. J. Hoarty, S. F. James, C. R. D. Brown, B. M. Williams, H. K. Chung, J. W. O. Harris, L. Upcraft, B. J. B. Crowley, C. C. Smith, and R. W. Lee. Measurements of emission spectra from hot, dense germanium plasma in short pulse laser experiments. *HEDP*, 6:105, 2010.
- [8] C. R. D. Brown, D. J. Hoarty, S. F. James, D. Swatton, S. J. Hughes, J. W. Morton, T. M. Guymer, M. P. Hill, D. A. Chapman, J. E. Andrew, A. J. Comley, R. Shepherd, J. Dunn, H. Chen, M. Schneider, G. Brown, P. Beiersdorfer, and

- J. Emig. Measurements of electron transport in foils irradiated with a picosecond time scale laser pulse. *Phys. Rev. Lett.*, 106:185003, 2011.
- [9] D. J. Hoarty, P. Allan, S. F. James, C. R. D. Brown, L. M. R. Hobbs, M. P. Hill, J. W. O. Harris, J. Morton, M. G. Brookes, R. Shepherd, J. Dunn, H. Chen, E. Von Marley, P. Beiersdorfer, H. K. Chung, R. W. Lee, G. Brown, and J. Emig. The first data from the Orion laser; measurements of the spectrum of hot, dense aluminium. *HEDP*, 9:661, 2013.
- [10] D. J. Hoarty, P. Allan, S. F. James, C. R. D. Brown, L. M. R. Hobbs, M. P. Hill, J. W. O. Harris, J. Morton, M. G. Brookes, R. Shepherd, J. Dunn, H. Chen, E. Von Marley, P. Beiersdorfer, H. K. Chung, R. W. Lee, G. Brown, and J. Emig. Observations of the effect of ionization-potential depression in hot dense plasma. *Phys. Rev. Lett.*, 110:265003, 2013.
- [11] M. Tabak, J. Hammer, M. E. Glinsky, W. L. Kruer, S. C. Wilks, J. Woodworth, E. M. Campbell, and M. D. Perry. Ignition and high gain with ultrapowerful lasers. *Phys. Plasmas*, 1:1626, 1994.
- [12] R. Kodama, P. A. Norreys, K. Mima, A. E. Dangor, R. G. Evans, H. Fujita, Y. Kitagawa, K. Krushelnick, T. Miyakoshi, N. Miyanaga, T. Norimatsu, S. J. Rose, T. Shozaki, K. Shigemori, A. Sunahara, M. Tampo, K. A. Tanaka, V. Toyama, T. Yamanaka, and M. Zepf. Fast heating of ultrahigh-density plasma as a step towards laser fusion ignition. *Nature*, 412:798, 2001.
- [13] U. Linz and J. Alonso. What will it take for laser driven proton accelerators to be applied to tumor therapy? *Phys. Rev. ST—Accel. Beams*, 10:094801, 2007.
- [14] S. A. Gaillard, T. Kluge, K. A. Flippo, M. Bussmann, B. Gall, T. Lockard, M. Geissel, D. T. Offermann, M. Schollmeier, Y. Sentoku, and T. E. Cowan. Increased laser-accelerated proton energies via direct laser-light-pressure acceleration of electrons in microcone targets. *Phys. Plasmas*, 18:056710, 2011.
- [15] M. Borghesi, A. Schiavi, D. H. Campbell, M. G. Haines, O. Willi, A. J. MacKinnon, L. A. Gizzi, M. Galimberti, R. J. Clarke, and H. Ruhl. Proton imaging: a diagnostic for inertial confinement fusion/fast ignitor studies. *Plasma Phys. Control. Fusion*, 43:A267, 2001.
- [16] J. A. Cobble, R. P. Johnson, T. E. Cowan, N. Renard-Le Galloudec, and M. Allen. High resolution laser-driven proton radiography. *J. Appl. Phys.*, 92:1775, 2002.

- [17] G. Sarri, C. A. Cecchetti, L. Romagnani, C. M. Brown, D. J. Hoarty, S. James, J. Morton, M. E. Dieckmann, R. Jung, O. Willi, S. V. Bulanov, F. Pegoraro, and M. Borghesi. The application of laser-driven proton beams to the radiography of intense laser-hohlraum interactions. *New J. Phys.*, 12:045006, 2010.
- [18] M. D. Perry, J. A. Sefcik, T. Cowan, S. Hatchett, A. Hunt, M. Moran, D. Pennington, R. Snavely, and S. C. Wilks. Hard x-ray production from high intensity laser solid interactions. *Rev. Sci. Instrum.*, 70:265, 1999.
- [19] R. D. Edwards, M. A. Sinclair, T. J. Goldsack, K. Krushelnick, F. N. Beg, E. L. Clark, A. E. Dangor, Z. Najmudin, M. Tatarakis, B. Walton, M. Zepf, K. W. D. Ledingham, I. Spencer, P. A. Norreys, R. J. Clarke, R. Kodama, Y. Toyama, and M. Tampo. Characterization of a gamma-ray source based on a laser-plasma accelerator with applications to radiography. *Appl. Phys. Lett.*, 80:2129, 2002.
- [20] Y. Glinec, J. Faure, L. Le Dain, S. Darbon, T. Hosokai, J. J. Santos, E. Lefebvre, J. P. Rousseau, F. Burgy, B. Mercier, and V. Malka. High-resolution γ -ray radiography produced by a laser-plasma driven electron source. *Phys. Rev. Lett.*, 94:025003, 2005.
- [21] H.-S. Park, D. M. Chambers, H.-K. Chung, R. J. Clarke, R. Eagleton, E. Giraldez, T. Goldsack, R. Heathcote, N. Izumi, M. H. Key, J. A. King, J. A. Koch, O. L. Landen, A. Nikroo, P. K. Patel, D. F. Price, B. A. Remington, H. F. Robey, R. A. Snavely, D. A. Steinman, R. B. Stephens, C. Stoeckl, M. Storm, M. Tabak, W. Theobald, R. P. J. Town, J. E. Wickersham, and B. B. Zhang. High-energy $K\alpha$ radiography using high-intensity, short-pulse lasers. *Phys. Plasmas*, 13:056309, 2006.
- [22] H.-S. Park, B. R. Maddox, E. Giraldez, S. P. Hatchett, L. T. Hudson, N. Izumi, M. H. Key, S. Le Pape, A. J. MacKinnon, A. G. MacPhee, P. K. Patel, T. W. Phillips, B. A. Remington, J. F. Seely, R. Tommasini, R. Town, J. Workman, and E. Brambrink. High-resolution 17–75keV backlighters for high energy density experiments. *Phys. Plasmas*, 15:072705, 2008.
- [23] E. Brambrink, H. G. Wei, B. Barbrel, P. Audebert, A. Benuzzi-Mounaix, T. Boehly, T. Endo, C. Gregory, T. Kimura, R. Kodama, N. Ozaki, H.-S. Park, M. Rabec le Gloahec, and M. Koenig. X-ray source studies for radiography of dense matter. *Phys. Plasmas*, 16:033101, 2009.
- [24] B. R. Maddox, H. S. Park, B. A. Remington, C. Chen, S. Chen, S. T. Prisbrey, A. Comley, C. A. Back, C. Szabo, J. F. Seely, U. Feldman, L. T. Hudson, S. Seltzer,

- M. J. Haugh, and Z. Ali. Absolute measurements of x-ray backlighter sources at energies above 10 keV. *Phys. Plasmas*, 18:056709, 2011.
- [25] A. Compant La Fontaine, C. Courtois, and E. Lefebvre. Production of multi-MeV bremsstrahlung x-ray sources by petawatt laser pulses on various targets. *Phys. Plasmas*, 19:023104, 2012.
- [26] J. M. Yang. Nuclear reactions in copper induced by protons from a petawatt laser-foil interaction. *Appl. Phys. Lett.*, 84:675, 2004.
- [27] S. J. Rose. Set the controls for the heart of the Sun. *Contemp. Phys.*, 45:109, 2004.
- [28] A. Pelka, G. Gregori, D. O. Gericke, J. Vorberger, S. H. Glenzer, M. M. Günther, K. Harres, R. Heathcote, A. L. Kritcher, N. L. Kugland, B. Li, M. Makita, J. Mithen, D. Neely, C. Niemann, A. Otten, D. Riley, G. Schaumann, M. Schollmeier, An. Tauschwitz, and M. Roth. Ultrafast melting of carbon induced by intense proton beams. *Phys. Rev. Lett.*, 105:265701, 2010.
- [29] T. D. Arber, K. Bennett, C. S. Brady, A. Lawrence-Douglas, M. G. Ramsay, N. J. Sircombe, P. Gillies, R. G. Evans, H. Schmitz, A. R. Bell, and C. P. Ridgers. Contemporary particle-in-cell approach to laser-plasma modelling. *Plasma Phys. Control. Fusion*, 57:113001, 2015.
- [30] D. J. Strozzi, M. Tabak, D. J. Larson, L. Divol, A. J. Kemp, C. Bellei, M. M. Marinak, and M. H. Key. Fast-ignition transport studies: Realistic electron source, integrated particle-in-cell and hydrodynamic modeling, imposed magnetic fields. *Phys. Plasmas*, 19:072711, 2012.
- [31] N. J. Sircombe, S. J. Hughes, and M. G. Ramsay. Integrated calculations of short-pulse laser interactions with matter. *New J. Phys.*, 15:025025, 2013.
- [32] B. I. Cohen, A. J. Kemp, and L. Divol. Simulation of laser-plasma interactions and fast electron transport in inhomogeneous plasma. *J. Comput. Phys.*, 229:4591, 2010.
- [33] P. L. Bhatnagar, E.P. Gross, and M. Krook. A model for collision processes in gases. I. Small amplitude processes in charged and neutral one-component systems. *Phys. Rev.*, 94:511, 1954.
- [34] D. Strickland and G. Mourou. Compression of amplified chirped optical pulses. *Op. Commun.*, 56:219, 1985.

- [35] A. Dubietis, G. Jonušauskas, and A. Piskarskas. Powerful femtosecond pulse generation by chirped and stretched pulse parametric amplification in BBO crystal. *Opt. Commun.*, 88:433, 1992.
- [36] D. I. Hillier, S. J. Duffield, S. P. Elsemere, M. T. Girling, N. W. Hopps, D. Hussey, M. J. Norman, S. J. Parker, P. A. Treadwell, and D. N. Winter. Nanosecond contrast of the Orion short pulse beamlines. *AWE Plasma Physics Department Annual Report*, page 15, 2011. Copies of the report can be obtained from the Corporate Communications Office at the Atomic Weapons Establishment, Aldermaston, Reading, Berkshire RG7 4PR, UK.
- [37] D. Hillier, C. Danson, S. Duffield, D. Egan, S. Elsmere, M. Girling, E. Harvey, N. Hopps, M. Norman, S. Parker, P. Treadwell, D. Winter, and T. Bett. Ultrahigh contrast from a frequency-doubled chirped-pulse-amplification beamline. *Appl. Opt.*, 52:4258, 2013.
- [38] P. A. Franken, A. E. Hill, C. W. Peters, and G. Weinreich. Generation of optical harmonics. *Phys. Rev. Lett.*, 7:118, 1961.
- [39] J. A. Armstrong, N. Bloembergen, J. Ducuing, and P. S. Pershan. Interactions between light waves in a nonlinear dielectric. *Phys. Rev.*, 127:1918, 1962.
- [40] J. Nuckolls, L. Wood, A. Thiessen, and G. Zimmerman. Laser compression of matter to super-high densities: Thermonuclear (CTR) applications. *Nature*, 239:139, 1972.
- [41] J. Lindl. Development of the indirect-drive approach to inertial confinement fusion and the target physics basis for ignition and gain. *Phys. Plasmas*, 2:3933, 1995.
- [42] A. B. Langdon. Nonlinear inverse bremsstrahlung and heated-electron distributions. *Phys. Rev. Lett.*, 44:575, 1980.
- [43] D. W. Forslund, J. M. Kindel, K. Lee, E. L. Lindman, and R. L. Morse. Theory and simulation of resonant absorption in a hot plasma. *Phys. Rev. A*, 11:679, 1975.
- [44] K. G. Estabrook, E. J. Valeo, and W. L. Kruer. Two-dimensional relativistic simulations of resonance absorption. *Phys. Fluids*, 18:1151, 1975.
- [45] F. Brunel. Not-so-resonant, resonant absorption. *Phys. Rev. Lett.*, 59:52, 1987.
- [46] E. A. Startsev and C. J. McKinstrie. Multiple scale derivation of the relativistic ponderomotive force. *Phys. Rev. E*, 55:7527, 1997.

- [47] W. L. Kruer and K. Estabrook. $J \times B$ heating by very intense laser light. *Phys. Fluids*, 28:430, 1985.
- [48] P. Kaw and J. Dawson. Relativistic nonlinear propagation of laser beams in cold overdense plasmas. *Phys. Fluids*, 13:472, 1970.
- [49] C. E. Max, J. Arons, and A. B. Langdon. Self-modulation and self-focusing of electromagnetic waves in plasmas. *Phys. Rev. Lett.*, 33:209, 1974.
- [50] K. G. Estabrook, W. L. Kruer, and D. S. Bailey. Two-dimensional ray-trace calculation of thermal whole beam self-focusing. *Phys. Fluids*, 28:19, 1985.
- [51] R. B. Miller. *An Introduction to the Physics of Intense Charged Particle Beams*. Plenum Press, 1982.
- [52] J. R. Davies. Electric and magnetic field generation and target heating by laser-generated fast electrons. *Phys. Rev. E*, 68:056404, 2003.
- [53] A. R. Bell and R. J. Kingham. Resistive collimation of electron beams in laser-produced plasmas. *Phys. Rev. Lett.*, 91:035003, 2003.
- [54] L. Gremillet, G. Bonnaud, and F. Amiranoff. Filamented transport of laser-generated relativistic electrons penetrating a solid target. *Phys. Plasmas*, 9:941, 2002.
- [55] J. Fuchs, T. E. Cowan, P. Audebert, H. Ruhl, L. Gremillet, A. Kemp, M. Allen, A. Blazevic, J.-C. Gauthier, M. Geissel, M. Hegelich, S. Karsch, P. Parks, M. Roth, Y. Sentoku, R. Stephens, and E. M. Campbell. Spatial uniformity of laser-accelerated ultrahigh-current MeV electron propagation in metals and insulators. *Phys. Rev. Lett.*, 91:255002, 2003.
- [56] M. Manclossi, J. J. Santos, D. Batani, J. Faure, A. Debayle, V. T. Tikhonchuk, and V. Malka. Study of ultraintense laser-produced fast-electron propagation and filamentation in insulator and metal foil targets by optical emission diagnostics. *Phys. Rev. Lett.*, 96:125002, 2006.
- [57] M. N. Quinn, D. C. Carroll, X. H. Yuan, M. Borghesi, R. J. Clarke, R. G. Evans, J. Fuchs, P. Gallegos, L. Lancia, K. Quinn, A. P. L. Robinson, L. Romagnani, G. Sarri, C. Spindloe, P. A. Wilson, D. Neely, and P. McKenna. On the investigation of fast electron beam filamentation in laser-irradiated solid targets using multi-MeV proton emission. *Plasma Phys. Control. Fusion*, 53:124012, 2011.

- [58] P. McKenna, A. P. L. Robinson, D. Neely, M. P. Desjarlais, D. C. Carroll, M. N. Quinn, X. H. Yuan, C. M. Brenner, M. Burza, M. Coury, P. Gallegos, R. J. Gray, K. L. Lancaster, Y. T. Li, X. X. Lin, O. Tresca, and C.-G. Wahlström. Effect of lattice structure on energetic electron transport in solids irradiated by ultraintense laser pulses. *Phys. Rev. Lett.*, 106:185004, 2011.
- [59] R. H. Huddleston and S. L. Leonard. *Plasma Diagnostic Techniques*. Academic Press, 1965.
- [60] L. Spitzer. *Physics of Ionized Gases*. John Wiley & Sons, 1962.
- [61] J. R. Davies. Laser absorption by overdense plasmas in the relativistic regime. *Plasma Phys. Control. Fusion*, 51:014006, 2009.
- [62] M. C. Levy, S. C. Wilks, M. Tabak, S. B. Libby, and M. G. Baring. Petawatt laser absorption bounded. *Nat. Comms.*, 5:4149, 2014.
- [63] K. Vaughan, A. S. Moore, V. Smalyuk, K. Wallace, D. Gate, G. Glendinning, S. W. McAlpin, H.-S. Park, C. Sorce, and R. M. Stevenson. High resolution 22 – 52 keV backlighter sources with applications for X-ray radiography. *AWE Plasma Physics Department Annual Report*, page 83, 2011. Copies of the report can be obtained from the Corporate Communications Office at the Atomic Weapons Establishment, Aldermaston, Reading, Berkshire RG7 4PR, UK.
- [64] T. Tajima and J. M. Dawson. Laser electron accelerator. *Phys. Rev. Lett.*, 43: 267, 1979.
- [65] S. P. D. Mangles, C. D. Murphy, Z. Najmudin, A. G. R. Thomas, J. L. Collier, A. E. Dangor, E. J. Divali, P. S. Foster, J. G. Gallacher, C. J. Hooker, D. A. Jaroszynski, A. J. Langley, W. B. Mori, P. A. Norreys, F. S. Tsung, R. Viskup, B. R. Walton, and K. Krushelnick. Monoenergetic beams of relativistic electrons from intense laser-plasma interactions. *Nature*, 431:535, 2004.
- [66] J. Faure, Y. Glinec, A. Pukhov, S. Kiselev, S. Gordienko, E. Lefebvre, J.-P. Rousseau, F. Burgy, and V. Malka. A laser-plasma accelerator producing monoenergetic electron beams. *Nature*, 431:541, 2004.
- [67] R. Bingham, J. T. Mendonça, and P. K. Shukla. Plasma based charged-particle accelerators. *Plasmas Phys. Control. Fusion*, 46:R1, 2004.
- [68] E. L. Clark, K. Krushelnick, J. R. Davies, M. Zepf, M. Tatarakis, F. N. Beg, A. Machacek, P. A. Norreys, M. I. K. Santala, I. Watts, and A. E. Dangor.

Measurements of energetic proton transport through magnetized plasma from intense laser interactions with solids. *Phys. Rev. Lett.*, 84:670, 2000.

- [69] R. A. Snavely, M. H. Key, S. P. Hatchett, T. E. Cowan, M. Roth, T. W. Phillips, M. A. Stoyer, E. A. Henry, T. C. Sangster, M. S. Singh, S. C. Wilks, A. MacKinnon, A. Offenberger, D. M. Pennington, K. Yasuike, A. B. Langdon, B. F. Lasinski, J. Johnson, M. D. Perry, and E. M. Campbell. Intense high-energy proton beams from petawatt-laser irradiation of solids. *Phys. Rev. Lett.*, 85:2945, 2000.
- [70] S. C. Wilks, A. B. Langdon, T. E. Cowan, M. Roth, M. Singh, S. Hatchett, M. H. Key, D. Pennington, A. MacKinnon, and R. A. Snavely. Energetic proton generation in ultra-intense laser-solid interactions. *Phys. Plasmas*, 8:542, 2001.
- [71] A. Pukhov. Three-dimensional simulations of ion acceleration from a foil irradiated by a short-pulse laser. *Phys. Rev. Lett.*, 86:3562, 2001.
- [72] M. Zepf, E. L. Clark, F. N. Beg, R. J. Clarke, A. E. Dangor, A. Gopal, K. Krushelnick, P. A. Norreys, M. Tatarakis, U. Wagner, and M. S. Wei. Proton acceleration from high-intensity laser interactions with thin foil targets. *Phys. Rev. Lett.*, 90:064801, 2003.
- [73] P. McKenna, K. W. D. Ledingham, J. M. Yang, L. Robson, T. McCanny, S. Shimizu, R. J. Clarke, D. Neely, K. Spohr, R. Chapman, R. P. Singhal, K. Krushelnick, M. S. Wei, and P. A. Norreys. Characterization of proton and heavier ion acceleration in ultrahigh-intensity laser interactions with heated target foils. *Phys. Rev. E*, 70:036405, 2004.
- [74] L. Yin, B. J. Albright, B. M. Hegelich, and J. C. Fernández. GeV laser ion acceleration from ultrathin targets: The laser break-out afterburner. *Laser Part. Beams*, 24:291, 2006.
- [75] A. P. L. Robinson, M. Zepf, S. Kar, R. G. Evans, and C. Bellei. Radiation pressure acceleration of thin foils with circularly polarized laser pulses. *New J. Phys.*, 10:013021, 2008.
- [76] A. P. L. Robinson, R. M. G. M. Trines, N. P. Dover, and Z. Najmudin. Hole-boring radiation pressure acceleration as a basis for producing high-energy proton bunches. *Plasma Phys. Control. Fusion*, 54:115001, 2012.
- [77] Anna Tauschwitz, E. Brambrink, J. A. Maruhn, M. Roth, M. Schollmeier, T. Schlegel, and Andreas Tauschwitz. Laser-produced proton beams as a tool for equation-of-state studies of warm dense matter. *HEDP*, 2:16, 2006.

- [78] E. Lefebvre, M. Carrié, and R. Nuter. Ion acceleration with high-intensity lasers and application to isochoric heating. *C. R. Physique*, 10:197, 2009.
- [79] D. J. Hoarty, T. Guymier, S. F. James, E. Gumbrell, C. R. D. Brown, M. Hill, J. Morton, and H. Doyle. Equation of state studies of warm dense matter samples heated by laser produced proton beams. *HEDP*, 8:50, 2012.
- [80] N. J. Sircombe. *Vlasov code development with inertial confinement fusion applications*. PhD thesis, University of Warwick, 2006.
- [81] S. Inoue, S. Tokita, K. Otani, M. Hashida, M. Hata, H. Sakagami, T. Taguchi, and S. Sakabe. Autocorrelation measurement of fast electron pulses emitted through the interaction of femtosecond laser pulses with a solid target. *Phys. Rev. Lett.*, 109:185001, 2012.
- [82] D. Potter. *Computational Physics*. John Wiley & Sons, 1973.
- [83] J. F. Drake, P. K. Kaw, Y. C. Lee, G. Schmid, C. S. Liu, and M. N. Rosenbluth. Parametric instabilities of electromagnetic waves in plasmas. *Phys. Fluids*, 17:778, 1974.
- [84] D. W. Forslund, J. M. Kindel, and E. L. Lindman. Theory of stimulated scattering processes in laser-irradiated plasmas. *Phys. Fluids*, 18:1002, 1975.
- [85] J. Dawson. On Landau damping. *Phys. Fluids*, 4:869, 1961.
- [86] M. E. Glinsky. Regimes of suprathermal electron transport. *Phys. Plasmas*, 2:2796, 1995.
- [87] J. R. Davies. How wrong is collisional Monte Carlo modeling of fast electron transport in high-intensity laser-solid interactions? *Phys. Rev. E*, 65:026407, 2002.
- [88] J. J. Honrubia and J. Meyer-ter-Vehn. Fast ignition of fusion targets by laser-driven electrons. *Plasma Phys. Control. Fusion*, 51:014008, 2009.
- [89] 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. *J. Comput. Phys.*, 194:1, 2004.
- [90] A. R. Bell, A. P. L. Robinson, M. Sherlock, R. J. Kingham, and W. Rozmus. Fast electron transport in laser-produced plasmas and the KALOS code for solution of the Vlasov-Fokker-Planck equation. *Plasma Phys. Control. Fusion*, 48:R37, 2006.

- [91] N. J. Sircombe and T. D. Arber. VALIS: A split-conservative scheme for the relativistic 2D Vlasov-Maxwell system. *J. Comput. Phys.*, 228:4773, 2009.
- [92] M. Sherlock. Universal scaling of the electron distribution function in one-dimensional simulations of relativistic laser-plasma interactions. *Phys. Plasmas*, 16:103101, 2009.
- [93] R. J. Kingham, M. Sherlock, C. P. Ridgers, and R. G. Evans. Vlasov-Fokker-Planck simulations of fast-electron transport with hydrodynamic plasma response. *J. Phys.: Conf. Ser.*, 244:022042, 2010.
- [94] M. Tzoufras, A. R. Bell, P. A. Norreys, and F. S. Tsung. A Vlasov-Fokker-Planck code for high energy density physics. *J. Comput. Phys.*, 230:6475, 2011.
- [95] C. K. Birdsall and A. B. Langdon. *Plasma Physics via Computer Simulation*. Taylor & Francis Group, 2005.
- [96] J. D. Hoffman. *Numerical Methods for Engineers and Scientists*. CRC Press, 2001.
- [97] R. Courant, K. Friedrichs, and H. Lewy. Über die partiellen differenzengleichungen der mathematischen physik (On the partial difference equations of mathematical physics). *Mathematische Annalen*, 100:32, 1928. English translation published in *IBM Journal*, pg 215, 1967.
- [98] T. Zh. Esirkepov. Exact charge conservation scheme for Particle-in-Cell simulation with an arbitrary form-factor. *Comput. Phys. Comm.*, 135:144, 2001.
- [99] J. Villasenor and O. Buneman. Rigorous charge conservation for local electromagnetic field solvers. *Comput. Phys. Comm.*, 69:306, 1992.
- [100] K. S. Kunz and J. Luebbers. *The Finite Difference Time Domain Method for Electromagnetics*. CRC Press, 1993.
- [101] K. Yee. Numerical solution of initial boundary value problems involving Maxwell's equations in isotropic media. *IEEE Trans. Ant. Prop.*, 14:302, 1966.
- [102] H. Ruhl. *Classical particle simulations with the PSC code*. An introduction to the PSC.
- [103] A. B. Langdon. Effects of the spatial grid in simulation plasmas. *J. Comput. Phys.*, 6:247, 1970.

- [104] J. May, J. Tonge, I. Ellis, W. B. Mori, F. Fiuza, R. A. Fonseca, L. O. Silva, and C. Ren. Enhanced stopping of macro-particles in particle-in-cell simulations. *Phys. Plasmas*, 21:052703, 2014.
- [105] A. J. Kemp, F. Fiuza, A. Debayle, T. Johzaki, W. B. Mori, P. K. Patel, Y. Sentoku, and L. O. Silva. Laser-plasma interactions for fast ignition. *Nucl. Fusion*, 54:054002, 2014.
- [106] B. B. Godfrey. Numerical Cherenkov instabilities in electromagnetic particle codes. *J. Comput. Phys.*, 15:504, 1974.
- [107] P. C. Birch and S. C. Chapman. Two dimensional particle-in-cell simulations of the lunar wake. *Phys. Plasmas*, 9:1785, 2002.
- [108] D. Tsiklauri and T. Haruki. Magnetic reconnection during collisionless, stressed, X-point collapse using particle-in-cell simulation. *Phys. Plasmas*, 14:112905, 2007.
- [109] M. E. Dieckmann, G. Sarri, D. Doria, H. Ahmed, and M. Borghesi. Evolution of slow electrostatic shock into a plasma shock mediated by electrostatic turbulence. *New J. Phys.*, 16:073001, 2014.
- [110] W. N. Lai, S. C. Chapman, and R. O. Dendy. Self-consistent nonlinear kinetic simulations of the anomalous Doppler instability of suprathermal electrons in plasmas. *Phys. Plasmas*, 20:102122, 2013.
- [111] A. Pukhov and J. Meyer-ter-Vehn. Relativistic laser-plasma interaction by multi-dimensional particle-in-cell simulations. *Phys. Plasmas*, 5:1880, 1998.
- [112] S. Bulanov, N. Naumova, F. Pegoraro, and J. Sakai. Particle injection into the wave acceleration phase due to nonlinear wake wave breaking. *Phys. Rev. E*, 58:R5257, 1998.
- [113] C. D. Decker, W. B. Mori, and T. Katsouleas. Particle-in-cell simulations of Raman forward scattering from short-pulse high-intensity lasers. *Phys. Rev. E*, 50:R3338, 1994.
- [114] O. Klimo, S. Weber, V. T. Tikhonchuk, and J. Limpouch. Particle-in-cell simulations of laser-plasma interaction for the shock ignition scenario. *Plasma Phys. Control. Fusion*, 52:055013, 2010.
- [115] T. Takizuka and H. Abe. A binary collision model for plasma simulation with a particle code. *J. Comput. Phys.*, 25:205, 1977.

- [116] M. E. Jones, D. S. Lemons, R. J. Mason, V. A. Thomas, and D. Winske. A grid-based Coulomb collision model for PIC codes. *J. Comput. Phys.*, 123:169, 1996.
- [117] W. Manheimer, M. Lampe, and G. Joyce. Langevin representation of Coulomb collisions in PIC simulations. *J. Comput. Phys.*, 138:563, 1997.
- [118] K. Nanbu and S. Yonemura. Weighted particles in Coulomb collision simulations based on the theory of a cumulative scattering angle. *J. Comput. Phys.*, 145:639, 1998.
- [119] Y. Sentoku, K. Mima, Y. Kishimoto, and M. Honda. Effects of relativistic binary collisions on PIC simulation of laser plasmas. *J. Phys. Soc. Jpn.*, 67:4084, 1998.
- [120] D. J. Larson. A Coulomb collision model for PIC plasma simulation. *J. Comput. Phys.*, 188:123, 2003.
- [121] Y. Sentoku and A. J. Kemp. Numerical methods for particle simulations at extreme densities and temperatures: Weighted particles, relativistic collisions and reduced currents. *J. Comput. Phys.*, 227:6846, 2008.
- [122] D. S. Lemons, D. Winske, W. Daughton, and B. Albright. Small-angle Coulomb collision model for particle-in-cell simulations. *J. Comput. Phys.*, 228:1397, 2009.
- [123] C. Wang, T. Lin, R. Caflisch, B. I. Cohen, and A. M. Dimits. Particle simulation of Coulomb collisions: Comparing the methods of Takizuka & Abe and Nanbu. *J. Comput. Phys.*, 227:4308, 2008.
- [124] F. Peano, M. Marti, L.O. Silva, and G. Coppa. Statistical kinetic treatment of relativistic binary collisions. *Phys. Rev. E*, 79:025701, 2009.
- [125] W. Manheimer, D. Colombant, and V. Goncharov. The development of a Krook model for nonlocal transport in laser produced plasmas. I. Basic theory. *Phys. Plasmas*, 15:083103, 2008.
- [126] J. D. Huba. *NRL Plasma Formulary*. Naval Research Laboratory, Washington D.C., 2006.
- [127] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery. *Numerical Recipes: The Art of Scientific Computing*. Cambridge University Press, 2007.
- [128] S. I. Braginskii. Transport processes in a plasma. In M. A. Leontovich, editor, *Reviews of Plasma Physics, Vol. 1*. Consultants Bureau, New York, 1965.

- [129] E. M. Epperlein and M. G. Haines. Plasma transport coefficients in a magnetic field by direct numerical solution of the Fokker-Planck equation. *Phys. Fluids*, 29: 1029, 1986.
- [130] A. R. Bell, R. G. Evans, and D. J. Nicholas. Electron energy transport in steep temperature gradients in laser-produced plasmas. *Phys. Rev. Lett.*, 46:243, 1981.
- [131] T. D. Arber and N. J. Sircombe. Simple collision operators for direct Vlasov simulations of laser plasma interaction and transport. *J. Phys.: Conf. Ser.*, 244: 022017, 2010.
- [132] E. S. Weibel. Spontaneously growing transverse waves in a plasma due to an anisotropic velocity distribution. *Phys. Rev. Lett.*, 2:83, 1959.
- [133] M. Bonitz and D. Semkat, editors. *Introduction to Computational Methods in Many Body Physics*. Rinton Press, 2004.
- [134] R. A. Fonseca, L. O. Silva, F. S. Tsung, V. K. Decyk, W. Lu, C. Ren, W. B. Mori, S. Deng, S. Lee, T. Katsouleas, and J. C. Adam. Osiris: A three-dimensional, fully relativistic particle in cell code for modeling plasma based accelerators. *Lect. Notes Comput. Sci.*, 2331:342, 2002.
- [135] F. Fiuza, M. Marti, R. A. Fonseca, L. O. Silva, J. Tonge, J. May, and W. B. More. Efficient modeling of laser-plasma interactions in high energy density scenarios. *Plasma Phys. Control. Fusion*, 53:074004, 2011.
- [136] M. Sherlock, E. G. Hill, R. G. Evans, and S. J. Rose. In-depth plasma-wave heating of dense plasma irradiated by short laser pulses. *Phys. Rev. Lett.*, 113: 255001, 2014.
- [137] R. M. Kulsrud. The origin of galactic magnetic fields. In R. Wielebinski and R. Beck, editors, *Cosmic magnetic fields*, volume 664 of *Lecture notes in physics*, pages 69–88. Springer-Verlag, Berlin Heidelberg, 2005. Section 4.
- [138] C. P. Ridgers, C. S. Brady, R. Duclous, J. G. Kirk, K. Bennett, T. D. Arber, A. P. L. Robinson, and A. R. Bell. Dense electron-positron plasmas and ultraintense γ rays from laser-irradiated solids. *Phys. Rev. Lett.*, 108:165006, 2012.
- [139] J. R. Davies, A. R. Bell, M. G. Haines, and S. M. Guérin. Short-pulse intensity laser-generated fast electron transport into thick solid targets. *Phys. Rev. E*, 56: 7193, 1997.

- [140] M. G. Ramsay and A. J. Simons. Proton acceleration at the front surface of laser-irradiated targets and the resultant material activation. *AWE Plasma Physics Department Annual Report*, page 74, 2011. Copies of the report can be obtained from the Corporate Communications Office at the Atomic Weapons Establishment, Aldermaston, Reading, Berkshire RG7 4PR, UK.
- [141] B. J. B. Crowley and J. W. Harris. Modelling of plasmas in an average-atom local density approximation: the CASSANDRA code. *J. Quant. Spectrosc. Radiat. Transfer*, 71:257, 2001.
- [142] A. Bret, L. Gremillet, and M. E. Dieckmann. Multidimensional electron beam-plasma instabilities in the relativistic regime. *Phys. Plasmas*, 17:120501, 2010.
- [143] A. Lawrence-Douglas. *Ionisation effects for laser-plasma interactions by particle-in-cell code*. PhD thesis, University of Warwick, 2013.
- [144] D. A. MacLellan, D. C. Carroll, R. J. Gray, N. Booth, M. Burza, M. P. Desjarlais, F. Du, B. Gonzalez-Izquierdo, D. Neely, H. W. Powell, A. P. L. Robinson, D. R. Rusby, G. G. Scott, X. H. Yuan, C.-G. Wahlström, and P. McKenna. Annular fast electron transport in silicon arising from low-temperature resistivity. *Phys. Rev. Lett.*, 111:095001, 2013.
- [145] R. W. Lee and J. T. Larsen. A time-dependent model for plasma spectroscopy of K-shell emitters. *J. Quant. Spectrosc. Radiat. Transfer*, 56:535, 1996.
- [146] A. P. L. Robinson and M. Sherlock. Magnetic collimation of fast electrons produced by ultraintense laser irradiation by structuring the target composition. *Phys. Plasmas*, 14:083105, 2007.
- [147] A. R. Bell, J. R. Davies, and S. M. Guerin. Magnetic field in short-pulse high-intensity laser-solid experiments. *Phys. Rev. E*, 58:2471, 1998.
- [148] S. Chawla, M. S. Wei, R. Mishra, K. U. Akli, C. D. Chen, H. S. McLean, A. Morace, P. K. Patel, H. Sawada, Y. Sentoku, R. B. Stephens, and F. N. Beg. Effect of target material on fast-electron transport and resistive collimation. *Phys. Rev. Lett.*, 110:025001, 2013.
- [149] H. B. Zhuo, X. H. Yang, C. T. Zhou, Y. Y. Ma, X. H. Li, and M. Y. Yu. Effect of resistivity gradient on laser-driven electron transport and ion acceleration. *Phys. Plasmas*, 20:093103, 2013.

- [150] X. H. Yang, H. Xu, H. B. Zhuo, Y. Y. Ma, F. Q. Shao, Y. Yin, and M. Borghesi. Effects of buried high-Z layers on fast electron propagation. *Eur. Phys. J. D.*, 68: 30, 2014.
- [151] A. Pukhov. Three-dimensional electromagnetic relativistic particle-in-cell code VLPL (Virtual Laser Plasma Lab). *J. Plasma Phys.*, 61:425, 1999.
- [152] D. A. MacLellan, D. C. Carroll, R. J. Gray, A. P. L. Robinson, M. P. Desjarlais, D. Neely, and P. McKenna. Influence of laser-drive parameters on annular fast electron transport in silicon. *Plasma Phys. Control. Fusion*, 56:084002, 2014.
- [153] D. A. MacLellan, D. C. Carroll, R. J. Gray, N. Booth, M. Burza, M. P. Desjarlais, F. Du, D. Neely, H. W. Powell, A. P. L. Robinson, G. G. Scott, X. H. Yuan, C.-G. Wahlström, and P. McKenna. Tunable mega-ampere electron current propagation in solids by dynamic control of lattice melt. *Phys. Rev. Lett.*, 113:185001, 2014.
- [154] M. N. Rosenbluth, W. M. MacDonald, and D. L. Judd. Fokker-Planck equation for an inverse-square force. *Phys. Rev.*, 107:1, 1957.
- [155] R. Shanny, J. M. Dawson, and J. M. Greene. One-dimensional model of a Lorentz plasma. *Phys. Fluids*, 10:1281, 1967.
- [156] P. C. Clemmow and J. P. Dougherty. *Electrodynamics of particles and plasmas*. Addison-Wesley Publishing Company, 1969.